Exponential Bounds and Stopping Rules for MCMC and General Markov Chains

I. Kontoyiannis
Dept of Informatics
Athens Univ of Econ & Business
Patission 76, Athens 10434, Greece
Email: yiannis@aueb.gr

L.A. Lastras-Montaño
IBM TJ Watson Research Center
1101 Kitchawan Rd
Yorktown Heights, NY, 10598
Email: lastrasl@us.ibm.com

S.P. Meyn
Dept of Electrical & Computer Eng
& Coordinated Sciences Laboratory
University of Ill at Urbana-Champaign
Urbana, IL 61801, USA
Email: meyn@uiuc.edu

Abstract— We develop explicit, general bounds for the probability that the empirical sample averages of a function of a Markov chain on a general alphabet will exceed the steady-state mean of that function by a given amount. Our bounds combine simple information-theoretic ideas together with techniques from optimization and some fairly elementary tools from analysis. In one direction, motivated by central problems in simulation, we develop bounds for the general class of "geometrically ergodic" Markov chains. These bounds take a form that is particularly suited to simulation problems, and they naturally lead to a new class of sampling criteria. These are illustrated by several examples. In another direction, we obtain a new bound for the important special class of Doeblin chains; this bound is *optimal*, in the sense that in the special case of independent and identically distributed random variables it essentially reduces to the classical Hoeffding bound.

I. INTRODUCTION

In many important scientific applications, we wish to compute the expected value $\pi(F) = \mathsf{E}_\pi[F] = \int F d\pi$ of a given function F, with respect to a probability distribution π which is known explicitly, but which cannot be computed numerically. In several areas, e.g., Bayesian statistics, image processing, and statistical mechanics, this is the rule rather than the exception; see [3], [25], [35] and the references therein.

Perhaps the most commonly used approach, is to estimate $\pi(F)$ through the sample averages,

$$\Gamma_n(F) := \frac{1}{n} \sum_{i=0}^{n-1} F(\Phi_i), \qquad n \geq 1, \qquad (1)$$

where the sequence $\mathbf{\Phi} = \{\Phi_i : i \geq 0\}$ is a Markov chain which is known to have stationary distribution π. This is the celebrated Markov Chain Monte Carlo (MCMC) paradigm. Under appropriate conditions, it is typically easy to check that the Markov chain $\mathbf{\Phi}$ is positive recurrent, and then the law of large numbers implies that the estimates in (1) will converge to the

required value $\pi(F)$, with probability one, as $n \to \infty$. Furthermore, the central limit theorem tells us that the estimates converge at a rate of $O(1/\sqrt{n})$, in that,

$$\sqrt{n}\big[\Gamma_n(F) - \pi(F)\big] \xrightarrow{\mathcal{D}} N(0, \sigma^2),$$

where $\sigma^2 = \lim_n n \mathrm{Var}(\Gamma_n(F))$ is the asymptotic variance of F.

The asymptotic results mentioned above are of limited use if we want to have any kind of "guarantee" that, after a certain number of steps in the simulation, the estimate $\Gamma_n(F)$ will indeed be close to $\pi(F)$. There are several well-studied approaches in the very extensive literature on this subject. Partly motivated by this discussion, here we consider the problem of providing simple, computable bounds for probabilities of "large deviations" type,

$$\Pr\big\{\Gamma_n(F) \geq r\big\}, \qquad (2)$$

for general classes of Markov chains $\mathbf{\Phi}$ and functions F.

Sanov's celebrated theorem for i.i.d. random variables states that, for appropriate sets $\mathcal{S}$ of probability measures on X, for large n we have,

$$-\log \mathsf{P}\{\Gamma_n \in \mathcal{S}\} \approx n \inf\{H(\mu\|\pi) : \mu \in \mathcal{S}\}, \quad (3)$$

where the relative entropy between two probability distributions μ, π on $(\mathsf{X}, \mathcal{B})$ is defined as usual by

$$H(\mu\|\pi) = \begin{cases} \int \log \frac{d\mu}{d\pi} \, d\mu & \text{if } d\mu/d\pi \text{ exists} \\ \infty & \text{otherwise.} \end{cases} \quad (4)$$

Donsker and Varadhan's classical large deviations principle [11][32] provides the natural extension of Sanov's theorem to the case of Markov chains. It states that, under appropriate conditions on the Markov chain and on F, as $n \to \infty$ we have,

$$-\log \mathsf{P}_x\big\{\Gamma_n(F) \geq r\big\} \approx n \inf_\gamma H(\gamma\|\gamma^1 \odot P), \quad (5)$$

where $\gamma^1 \odot P$ denotes the bivariate measure $(\gamma^1 \odot P)(x, dy) = \gamma^1(dx)P(x, dy)$, and the infimum is over all bivariate distributions γ with marginals γ^1 and γ^2 that satisfy $\gamma^1 = \gamma^2$ and $\gamma^1(F) \geq r$.

In this work, simple information-theoretic ideas are combined together with techniques from optimization and some fairly elementary tools from analysis, to obtain explicit *non-asymptotic* bounds for the probabilities (2). Information-theoretic methods have been very influential in the development of asymptotic results as well as non-asymptotic inequalities for probabilities as in (2). Combinatorial techniques based on the method of types have provided some of the simplest proofs as well as some of the strongest results when $\mathbf{\Phi}$ is an independent and identically distributed (i.i.d.) sequence [7], [8], [6]. But in the case of Markov chains they have been much less successful, their applicability limited essentially only to the more elementary case of finite state-space Markov chains; see [9] and the references therein. This difficulty reflects, to some extent, the intrinsic limitations of the information-theoretic methods, but it is also due to the much greater complexity of the field of large deviations for general Markov chains. For example, the elegant theorem of Sanov, which holds in complete generality in the i.i.d. case, requires extremely strong assumptions in order to be translated to Markov chains; see Donsker and Varadhan's classic results [11][32], as well as the numerous counter-examples indicating that such strong assumptions are indeed necessary, e.g., [2], [4], [13].

In Section II and Section III we prove a series of *non-asymptotic* deviation bounds, which apply to a variety to different problems. Then in Section III, motivated by the form these bounds naturally take in their most general form, we propose a new, general-purpose sampling criterion, which can easily be applied to a great variety of different MCMC problems. In particular, this sampling criterion comes with a precise, non-asymptotic performance guarantee.

The exponential deviation bounds obtained in Section II are established in four steps. We obtain an information-theoretic bound for the probability of interest in terms of relative entropy; this is done in Theorem 2.2, which is inspired by an argument used by Csiszár in the proof of [7, Theorem 1]. We next apply Pinsker's inequality to simplify the bound obtained in Step 1. In Step 3 a further relaxation of this bound is obtained, which is expressed as a linear program over the space of probability measures. The final step is to construct bounds on this linear program. This is achieved by constructing a dual linear program. Similar tools are used to obtain worst-case exponential bounds

for constrained problems in the i.i.d. case in [31].

Most of our results are present here without proofs, or with brief proof outlines. The complete arguments will be given in a longer version of this paper.

II. GENERAL BOUNDS AND DOEBLIN CHAINS

In the following two sections we present several non-asymptotic versions of exponential upper bounds for the probability that the sample averages of a function of a Markov chain deviate from their (asymptotic) mean by a certain amount.

Throughout the paper, we consider a discrete-time Markov chain $\mathbf{\Phi} = \{\Phi_i : i \geq 0\}$ with values in the Polish space $(\mathsf{X}, \mathcal{B})$, where $\mathcal{B}$ is the natural Borel σ-algebra associated with X. The distribution of $\mathbf{\Phi}$ is specified by its initial state $\Phi_0 = x \in \mathsf{X}$ and its transition kernel P, where $P(x, B) := \Pr\{\Phi_{i+1} \in B \mid \Phi_i = x\}$. We write P_x for the distribution of the chain conditional on the initial state $\Phi_0 = x$, and E_x for the corresponding expectation. For an arbitrary (measurable) function $F : \mathsf{X} \to \mathbb{R}$ we write $\Gamma_n(F)$ for the sample averages (1).

We begin in Section II-A by deriving an elementary bound which applies to *all* Markov chains, not even requiring positive recurrence; this is given in Theorem 2.1, and in a more general form in Theorem 2.2. These bounds show that for deviation probabilities of the form,

$$\mathsf{P}_x\{\Gamma_n(F) \geq r\} = \Pr\left\{\frac{1}{n}\sum_{i=0}^{n-1} F(\Phi_i) \geq r \,\Big|\, \Phi_0 = x\right\},$$

we can always find a natural exponential upper bound, where the exponent is expressed in terms of relative entropy. The proofs are based on simple information-theoretic ideas.

Although Theorems 2.1 and 2.2 hold in complete generality, the exponent in these bounds is defined in an implicit way and it cannot be easily evaluated in specific cases. In order to obtain concrete, computable bounds, in Section II-B we turn to the special class of Doeblin (or "uniformly ergodic") chains, and in Theorem 2.4 we give an explicit bound for the deviation probabilities of the sample averages of any bounded function of a Doeblin chain.

Finally, in Section III we return to simulation and consider the problem of empirically estimating the mean $\pi(F) := \int F d\pi$ of a given function F using the sample averages of a Markov chain with stationary distribution π. For such simulation problems we introduce a new *sampling criterion*, which gives a precise guideline regarding which parts of the simulation give

more accurate estimates with high probability. The theoretical justification for this sampling criterion is given in Theorem 3.1, where we present a general, explicit bound for the probability of large deviations of the sample averages from their mean. This is established for the more general class of "geometrically ergodic" chains, and it also allows us to consider the sample averages of unbounded functions.

A. A General Information-Theoretic Bound

Recall the definition (1) for the sample averages $\Gamma_n(F)$ of a function F of a Markov chain $\mathbf{\Phi} = \{\Phi_i : i \geq 0\}$. The sup norm of any such function $F : \mathsf{X} \to \mathbb{R}$ is denoted by $\|F\|_\infty = \sup_{x \in \mathsf{X}} |F(x)|$, and F_+ denotes the function $\max\{F, 0\}$. For any two probability measures μ, ν on $(\mathsf{X}, \mathcal{B})$, the L^1 distance $\|\mu - \nu\|$ between them is twice the total variation distance,

$$\|\mu - \nu\| := 2 \sup_{A \in \mathcal{B}} |\mu(A) - \nu(A)|,$$

and the relative entropy $H(\mu\|\nu)$ between μ and ν is defined as in (4).

Our first result is a non-asymptotic version of the upper bound in the Donsker-Varadhan large deviations principle (5).

Theorem 2.1: Let $\mathbf{\Phi} = \{\Phi_i : i \geq 0\}$ be an *arbitrary* Markov chain with values on X. For any function $F \colon \mathsf{X} \to \mathbb{R}$ which is bounded above, any $r \geq \pi(F)$ and any initial condition $x \in \mathsf{X}$, we have,

$$-\log \mathsf{P}_x\big\{\Gamma_n(F) \geq r\big\} \geq (n-1)H(\gamma\|\gamma^1 \odot P), \quad (6)$$

where γ is a bivariate probability measure on $(\mathsf{X} \times \mathsf{X}, \mathcal{B} \times \mathcal{B})$ with respective marginals γ^1 and γ^2, and where the marginals γ^1, γ^2 satisfy,

$$\|\gamma^1 - \gamma^2\| \leq \frac{2}{n-1} \quad \text{and} \quad \gamma^1(F) \geq r - \frac{\|F_+\|_\infty - r}{n-1}. \quad (7)$$

A more general version of this result is given in Theorem 2.2 below, along with a proof.

Remark. As mentioned in the Introduction, the classical extension of Sanov's theorem to Markov chains is Donsker and Varadhan's [11], [32] large deviations principle. It states that, under appropriate conditions on the Markov chain and on F, as $n \to \infty$ we have,

$$-\log \mathsf{P}_x\big\{\Gamma_n(F) \geq r\big\} \approx n \inf_\gamma H(\gamma\|\gamma^1 \odot P), \quad (8)$$

where the infimum is over all bivariate measures γ with marginals γ^1 and γ^2 that satisfy

$$\gamma^1 = \gamma^2 \quad \text{and} \quad \gamma^1(F) \geq r. \quad (9)$$

Comparing the Donsker-Varadhan result in equations (8)-(9) with the bound of Theorem 2.1 in equations (6)-(7), we see that the theorem provides a *non-asymptotic* version of the upper bound in (8), and thus offers an elementary "explanation" for the Donsker-Varadhan rate function. Of course the intimate relationship between large deviations and the entropy functional has been noted by many authors in the literature, and several other alternative points of view have been offered. For example, as argued in [12], the minimization of the entropy in (8) can be seen as an optimal control problem, where the entropy naturally appears as the cost functional. For more on this relationship, see, e.g., the monographs [14][12], the review [33], and the references therein.[1]

For some of our subsequent results, we will need a somewhat more general version of the above theorem. Instead of a single function F, we now consider a finite collection of functions $\{F_1, F_2, \ldots, F_m\}$, and we bound the probability of the following event,

$$\mathcal{E} = \{\Gamma_n(F_j) \geq r_j \text{ for all } j, \text{ and } \Phi_{n-1} \in B\}, \quad (10)$$

where the sample average $\Gamma_n(F_j)$ of each function F_j is taken to exceed some constant r_j, and we also "pin" the position of the Markov chain during the last time step Φ_{n-1} to be in some fixed set B.

Large deviation limit theory for events with pinning of the form (10) has been studied extensively; see, e.g., [10][23][29][30] and the references therein. In the present context, the value of the additional pinning constraint $\{\Phi_{n-1} \in B\}$ is that it allows us to treat unbounded functions F_j. Theorem 2.2 forms the basis for analysis of the simulation algorithm introduced in Section III.

Theorem 2.2: Let $\mathbf{\Phi} = \{\Phi_i : i \geq 0\}$ be an *arbitrary* Markov chain with values on X, let $B \in \mathcal{B}$ be fixed measurable set, and let $F_1, F_2, \ldots, F_m$ be an arbitrary (finite) collection of functions from X to $\mathbb{R}$. For any constants $r_1, r_2, \ldots, r_m$, and any initial condition $x \in \mathsf{X}$, we have that

$$-\log \mathsf{P}_x\Big\{\Gamma_n(F_j) \geq r_j \text{ for all } j \text{ , and } \Phi_{n-1} \in B\Big\}$$

is bounded below by

$$(n-1)H(\gamma\|\gamma^1 \odot P),$$

where γ is a bivariate probability measure on $(\mathsf{X} \times \mathsf{X}, \mathcal{B} \times \mathcal{B})$ whose marginals γ^1 and γ^2 satisfy,

$$\|\gamma^1 - \gamma^2\| \leq \frac{2}{n-1}$$

---

[1]In fact, Varadhan in [33] remarks that "It is not hard to see that any 'large deviation' has to be related to 'entropy'."

and

$$\gamma^1(F_j) \geq r_j - \frac{\|(F_j)_+\|_{\infty,B} - r_j}{n-1}, \quad \text{for all } j,$$

where $\|F\|_{\infty,B}$ is the sup norm of the restriction of the function F to the set B, i.e., $\|F\|_{\infty,B} := \sup_{x \in B} |F(x)|$.

The first step in the proof of Theorem 2.2 is an application of the following observation due to Csiszár [7]. It states that the probability of *any* event can be expressed as the exponential of a relative entropy. The proof is immediate from the definitions.

Lemma 2.3 (Csiszár's Lemma): Let p be an arbitrary probability distribution on any probability space, and $\mathcal{E}$ any event with $p(\mathcal{E}) > 0$. Let $p|_{\mathcal{E}}$ denote the conditional distribution $p|_{\mathcal{E}}(\cdot) = p(\cdot \cap \mathcal{E})/p(\mathcal{E})$. Then:

$$-\log p(\mathcal{E}) = H(p|_{\mathcal{E}}\|p).$$

PROOF OF THEOREM 2.2. Fix an initial state $x \in \mathsf{X}$. Let p denote the measure on X^n induced by the distribution of $(\Phi_0, \Phi_1, \ldots, \Phi_{n-1})$ conditioned on $\{\Phi_0 = x\}$, and let μ_n denote the conditional measure $\mu_n = p|_{\mathcal{E}}$ with $\mathcal{E}$ for the event of interest expressed in (10). From Csiszár's Lemma,

$$-\log \mathsf{P}_x\{\mathcal{E}\} \geq H(\mu_n\|p). \tag{11}$$

For any $k \geq 1$, we write $x_0^k = (x_0, x_1, \ldots, x_k)$ for a $(k+1)$-dimensional vector from X^{k+1}. Writing μ_n and p as products of conditional distributions, $\mu_n(dx_0^{n-1})$ is

$$\mu_1(dx_0)\mu_2(dx_1|x_0) \cdots \mu_n(dx_{n-1}|x_0^{n-2}),$$

and $p(dx_0^{n-1})$ is

$$\delta_x(dx_0)P(x_0, dx_1) \cdots P(x_{n-2}, dx_{n-1}).$$

Then the relative entropy in (11) can be expanded,

$$\sum_{i=1}^{n-1} \int_{x_0^{i-1}} H\Big(\mu_{i+1}(\cdot\,|x_0^{i-1})\|P(x_{i-1},\cdot)\Big)\, \mu_i(dx_0^{i-1}).$$

Denote by μ^i the one-dimensional marginal of μ_n corresponding to x_{i-1}, and $\mu^{i,i+1}$ the two-dimensional marginal corresponding to (x_{i-1}, x_i). The above expansion for $H(\mu_n\|p)$ can be expanded as follows,

$$\sum_{i=1}^{n-1} \int_{x_{i-1}} \int_{x_0^{i-2}} H\Big(\mu_{i+1}(\cdot\,|x_0^{i-1})\|P(x_{i-1},\cdot)\Big)$$
$$\mu_i(dx_0^{i-2}|x_{i-1})\mu^i(dx_{i-1}).$$

Now, using the joint convexity of the relative entropy in its two arguments, we can take the innermost integral

inside the relative entropy, to that $H(\mu_n\|p)$ is bounded below by

$$\sum_{i=1}^{n-1} \int_{x_{i-1}} H\Big(\mu_{i+1}(\cdot\,|x_{i-1})\|P(x_{i-1},\cdot)\Big)\mu^i(dx_{i-1})$$

which equals

$$\sum_{i=1}^{n-1} H\big(\mu^{i,i+1}\|\mu^i \odot P\big).$$

Using the joint convexity of H again,

$$\begin{aligned} H(\mu_n\|p) &\geq (n-1)\sum_{i=1}^{n-1}\frac{1}{n-1}H\big(\mu^{i,i+1}\|\mu^i \odot P\big) \\ &\geq (n-1)H(\gamma\|\gamma^1 \odot P), \end{aligned}$$

where the bivariate measure γ and its first marginal γ^1 are,

$$\gamma = \frac{1}{n-1}\sum_{i=1}^{n-1}\mu^{i,i+1} \quad \text{and} \quad \gamma^1 = \frac{1}{n-1}\sum_{i=1}^{n-1}\mu^i.$$

This final bound combined with (11) gives the required bound, and it only remains to verify that γ satisfies the stated properties. Indeed, since the second marginal of γ is $\gamma^2 = \frac{1}{n-1}\sum_{i=1}^{n-1}\mu^{i+1}$, their difference is $\gamma^1 - \gamma^2 = \frac{\mu^1 - \mu^n}{n-1}$, and since the L^1-norm is bounded by 2 it follows that $\|\gamma^1 - \gamma^2\| \leq 2/(n-1)$. Finally, by the definition of γ^1 and the event $\mathcal{E}$, for any j we have that $\gamma^1(F_j)$ is equal to

$$\frac{1}{n-1}\sum_{i=1}^{n-1}\mu^i(F_j) = \mathsf{E}_x\Big[\frac{1}{n-1}\sum_{i=1}^{n-1}F_j(\Phi_{i-1})\,\big|\,\mathcal{E}\Big]$$

which becomes,

$$\frac{n}{n-1}\mathsf{E}_x\Big[\frac{1}{n}\sum_{i=0}^{n-1}F_j(\Phi_i)\,\big|\,\mathcal{E}\Big] - \mathsf{E}_x\Big[\frac{F_j(\Phi_{n-1})}{n-1}\,\big|\,\mathcal{E}\Big],$$

and this is bounded below by $r_j - \frac{1}{n-1}\big[\|(F_j)_+\|_{\infty,B} - r_j\big]$, as required. $\qquad\square$

B. Doeblin Chains

As before, we start with a Markov processes $\Phi = \{\Phi_i : i \geq 0\}$ taking values in a general state space $(\mathsf{X}, \mathcal{B})$, with initial state $\Phi_0 = x \in \mathsf{X}$ and transition kernel P. Throughout the rest of the paper we assume that Φ is *ψ-irreducible* and *aperiodic*. This means that there is a σ-finite measure ψ on $(\mathsf{X}, \mathcal{B})$ such that, for any $A \in \mathcal{B}$ satisfying $\psi(A) > 0$ and any initial condition x,

$$P^n(x, A) > 0, \qquad \text{for all } n \text{ sufficiently large,}$$

where $P^n(x, \cdot)$ is the probability measure that describes the conditional distribution of Φ_n, given that $\Phi_0 = x$.

Recall that any kernel $P(x, dy)$ acts linearly on functions $F : \mathsf{X} \to \mathbb{R}$ on the right and signed measures μ on $(\mathsf{X}, \mathcal{B})$ on the left, via,

$$PF(x) = \int F(y) P(x, dy)$$

and

$$\mu P(dy) = \int P(x, dy) \mu(dx),$$

respectively, whenever the above integrals are well-defined.

The Markov chains we consider are also positive recurrent. This means that there is a unique invariant probability measure π satisfying $\pi P = \pi$. The first such class of Markov chains we consider are *Doeblin* chains. This is probably the "nicest" class of general-alphabet Markov chains. For the aperiodic, ψ-irreducible chains that we consider, Doeblin's condition is equivalent to requiring that Φ is "uniformly ergodic," namely, that it has a unique stationary distribution π such that, as $n \to \infty$,

$$d_n := \sup_x \| P^n(x, \cdot) - \pi \| \to 0 \quad \text{geometrically fast,} \tag{12}$$

where, as before, $\| \mu - \nu \|$ denotes the L^1-distance between two probability measures μ, ν. Alternatively, an aperiodic and ψ-irreducible chain Φ is a Doeblin chain if and only if either one of the following two conditions is satisfied:

(i) There exists some $n \geq 1$ for which $d_n < 2$.
(ii) *Doeblin's minorization condition* holds: There exists an integer $n_0 \geq 1$, a constant $a_0 > 0$, and a probability measure ϱ on $(\mathsf{X}, \mathcal{B})$, such that,

$$P^{n_0}(x, A) \geq a_0 \varrho(A), \ \forall \, x \in \mathsf{X}, A \in \mathcal{B}. \tag{13}$$

The main advantage of the minorization condition (13) is that it can usually be easily verified in applications, and with explicit values for the constants n_0 and a_0. See [28] for a detailed discussion of Doeblin chains and their characterizations.

Our next result, Theorem 2.4 below, gives a non-asymptotic, exponential bound, on the probability that the sample averages $\Gamma_n(F)$ of a bounded function of a Doeblin chain deviate from their steady-state mean $\pi(F)$ by a certain amount. This bound improves upon a recent result of Glynn and Ormoneit [17] by a factor of 2 in the exponent. The proof technique of [17] is completely different, relying on martingale methods

and a generalization of Hoeffding's original argument [22].[2]

Theorem 2.4: Suppose the Markov chain $\Phi = \{ \Phi_i : i \geq 0 \}$ is ψ-irreducible, aperiodic, and satisfies the Doeblin condition (13), so that it has a unique stationary distribution π. Then, for any bounded function $F : \mathsf{X} \to \mathbb{R}$ and any $\epsilon > 0$ we have that

$$\log \mathsf{P}_x \big\{ \Gamma_n(F) \geq \pi(F) + \epsilon \big\}$$

is bounded above by

$$\leq -\frac{n-1}{2} \Big[\frac{a_0}{n_0 \| F \|_{\mathrm{sp}}} \epsilon - \frac{4}{n-1} \Big]_+^2,$$

where $\| F \|_{\mathrm{sp}}$ denotes the normalized span norm of F, defined by $\| F \|_{\mathrm{sp}} := \inf_{u \in \mathbb{R}} \| [F - u] \|_\infty$.

Remarks.

1) When Φ is a sequence of i.i.d. random variables with common distribution π, then the Doeblin condition holds with $n_0 = a_0 = 1$ and $\varrho = \pi$, and the bound of Theorem 2.4 reduces to

$$-\frac{n-1}{2} \Big[\frac{\epsilon}{\| F \|_{\mathrm{sp}}} - \frac{4}{n-1} \Big]_+^2.$$

This is essentially identical to the exponent $-\frac{n}{2} [\epsilon / \| F \|_\infty]^2$ of the classical Hoeffding bound [22], which is known to be tight in the i.i.d. case.

2) Although Doeblin chains form a very restricted sub-class of all ergodic chains, it is perhaps the most natural class to consider in terms of large deviations properties. To see that, recall that Bryc and Dembo [4] have provided a counter-example of a stationary Doeblin chain for which the regular large deviations principle fails to hold with any rate function. Moreover, if it were possible to obtain meaningful exponential bounds as in Theorem 2.4, with exponents that were independent of the initial condition, this would mean that the ergodic theorem would hold for all bounded functions uniformly in the initial condition, a fact which is known to *imply* that the chain is Doeblin [28].

3) Once the Doeblin condition (13) has been verified, all the constants appearing in the result of Theorem 2.4 are explicitly known. Moreover, if the span norm cannot be easily evaluated, we always have the trivial bound $\| F \|_{\mathrm{sp}} \leq \| F \|_\infty$.

[2]Note also that the result of [17] upon which we improve is actually stated slightly incorrectly there; there should have been an extra factor of 2 in the estimate of the norm of the function g, which translates to an extra factor of $1/2$ in the exponent they finally obtain. A factor of 2 is also missing in the ergodic theorem for uniformly ergodic chains, [28, Theorem 16.2.4]. This ergodic theorem is a key ingredient of the proof in [17].

PROOF OUTLINE. The proof of Theorem 2.4 consists of four steps.

The first step is an application of the general bound of Theorem 2.1, implying that

$$-\log \mathsf{P}_x\big\{\Gamma_n(F) \geq \pi(F) + \epsilon\big\}$$

is bounded below by $(n-1)H(\gamma\|\gamma^1 \odot P)$, for some bivariate probability measure γ satisfying the conditions in (7) with $r = \pi(F) + \epsilon$.

The second step is to obtain a bound on the above relative entropy using Pinsker's inequality:

Lemma 2.5 (Pinsker's inequality): The bound,

$$H(\mu\|\nu) \geq \tfrac{1}{2}\|\mu - \nu\|^2,$$

holds for any two probability measures μ, ν on $(\mathsf{X}, \mathcal{B})$. Consequently, the result of the previous step yields,

$$-\log \mathsf{P}_x\big\{\Gamma_n(F) \geq \pi(F) + \epsilon\big\} \geq \frac{n-1}{2}\|\gamma - \gamma^1 \odot P\|^2. \tag{14}$$

Step 3 amounts to the following construction of a *linear program* whose value bounds the right hand side of (14). Write $\gamma = \gamma^1 \odot Q$ for some transition kernel Q. An application of Jensen's inequality gives that,

$$
\begin{aligned}
\|\gamma^1 \odot Q &- \gamma^1 \odot P\| \\
&= \int \gamma^1(dx)\|P(x,\,\cdot\,) - Q(x,\,\cdot\,)\| \\
&\geq \|\gamma^1[P - Q]\|. \tag{15}
\end{aligned}
$$

Moreover $\gamma^1[P-Q] = \gamma^1 P - \gamma^2$, which we write as,

$$\gamma^1[P - Q] = \gamma^1 P - \gamma^2 = [\gamma^1 - \gamma^1 P] - [\gamma^2 - \gamma^1].$$

We have $\|\gamma^1 - \gamma^2\| \leq 2/(n-1)$ by Theorem 2.1, which together with (14) and (15) gives,

$$
\begin{aligned}
-\log \mathsf{P}_x&\big\{\Gamma_n(F) \geq \pi(F) + \epsilon\big\} \\
&\geq \frac{n-1}{2}\big[\|\gamma^1 - \gamma^1 P\| - 2/(n-1)\big]_+^2 \tag{16}
\end{aligned}
$$

A bound on the right hand side of (16) follows from any lower bound on the solution to the following linear program over the space of probability measures:

$$\min \quad \|\mu - \mu P\| \qquad \text{s.t.} \quad \mu(F) - \pi(F) \geq \epsilon^n, \tag{17}$$

where $\epsilon^n := \epsilon - (\|F_+\|_\infty - r)/(n-1)$, with $\epsilon = r - \pi(F)$. Letting L^* denote the value of this linear program, we obtain from (16),

$$
\begin{aligned}
-\log \mathsf{P}_x&\big\{\Gamma_n(F) \geq \pi(F) + \epsilon\big\} \\
&\geq \frac{n-1}{2}\big[L^* - 2/(n-1)\big]_+^2. \tag{18}
\end{aligned}
$$

The fourth and final step is to obtain a bound on L^*. The details involve linear programming ideas and duality. $\qquad\square$

III. SIMULATION BOUNDS AND STOPPING RULES

In this section we present a new, general-purpose stopping rule, which can be easily applied to a variety of Markov Chain Monte Carlo problems. In Section III-A we treat a simple example illustrating the stopping rule in a simple setting, Section III-B contains the description of the general stopping rule, and in Section III-C we present a non-asymptotic bound that justifies this stopping rule and also offers precise performance guarantees for its use.

A. A Stopping Rule for Simulating a Simple Queue

We begin by considering the example of the simple discrete-time queue, described by the recursion,

$$\Phi_{k+1} = [\Phi_k - S_{k+1}]_+ + A_{k+1}, \qquad k \geq 0, \ \Phi_0 \in \mathbb{R}_+, \tag{19}$$

where the processes $\boldsymbol{A} = \{A_i : i \geq 0\}$ and $\boldsymbol{S} = \{S_i : i \geq 0\}$ are mutually independent, each consisting of i.i.d. random variables with values in $\mathbb{R}_+$, and with marginals that possess finite means denoted α, μ, respectively. If the 'load condition' $\rho := \alpha/\mu < 1$ is satisfied, then it is known that $\boldsymbol{\Phi} = \{\Phi_i : i \geq 0\}$ is a positive recurrent Markov chain, and we denote its stationary distribution by π. If, in addition, the marginal of $\boldsymbol{A}$ has a finite second moment, then the steady state mean of Φ_i is finite, i.e., $\pi(F) < \infty$ with $F(x) \equiv x$.

Simulation of the steady-state mean appears to be straightforward: Simply compute $\Gamma_n(F)$ for large n. However, it is well known that this can be computationally intensive since the variance of $\Gamma_n(F)$ is very large, actually of order $(1-\rho)^{-4}/n$ for $\rho \sim 1$; see [1], [34], [18], [27]. Moreover, it is *impossible* to obtain an exponential bound on the error probability: It is shown in [27] that for any $\varepsilon > 0$ we have,

$$\lim_{n\to\infty} n^{-1}\log \mathsf{P}_x\big\{\Gamma_n(F) \geq \pi(F) + \epsilon\big\} = 0, \tag{20}$$

provided the $\{A_i\}$ are not identically zero.

One approach to improve this estimator is through the application of the control variate method [24], [15]. Suppose that there is a function $\mathcal{D} : \mathsf{X} \to \mathbb{R}$ for which it is known a priori that $\pi(\mathcal{D}) = 0$. For any such $\mathcal{D}$, the following will also be a consistent estimator,

$$\Gamma_n(F - \mathcal{D}) = \frac{1}{n}\sum_{i=0}^{n-1}\big(F(\Phi_i) - \mathcal{D}(\Phi_i)\big), \qquad n \geq 1. \tag{21}$$

The function $\mathcal{D}$ is called a *shadow function*, since it is meant to 'eclipse' the function F to be simulated.

For queueing models, there is a specialized version of the control-variate method corresponding to a particular choice of the shadow function $\mathcal{D}$, called the
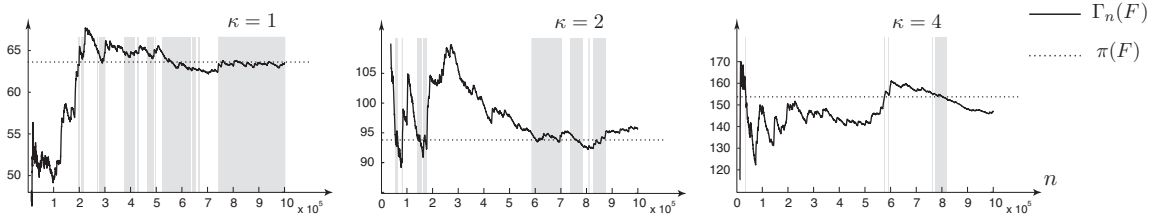
Fig. 1: Application of the shadow function sampling criterion for the simple queue (19) with $\mathcal{D} = J - PJ$, $J(x) \equiv \frac{1}{2}x^2/(\mu - \alpha)$, $F(x) \equiv x$, and for $\kappa = 1, 2, 4$. The grey regions indicate time periods during which $|\Gamma_n(\mathcal{D})| \leq 1$ and $\Phi_{n-1} \in B := [0, 10^4]$.

smoothed estimator. The starting point is the use of a natural "value function" $J(x)$, obtained from a deterministic fluid model for the network to be simulated. In our example, suppose that each A_k takes the values $(1 + \kappa)\alpha$ and zero, with probabilities $(1 + \kappa)^{-1}$ and $1 - (1 + \kappa)^{-1}$, respectively, and that each S_k takes the values 2μ and zero, each with probability $1/2$. The parameters α and μ describe the average arrival rate and service rate, respectively, and κ describes the variability of the arrival process. In this case, the value function $J(x)$ takes the form,

$$J(x) = \frac{x^2}{2(\mu - \alpha)}, \qquad x \in \mathbb{R}_+. \tag{22}$$

The smoothed estimator corresponds to following choice for the shadow function $\mathcal{D}$,

$$\begin{aligned}
\mathcal{D}(x) &= \mathsf{E}[J(\Phi_k) - J(\Phi_{k+1}) \mid \Phi_k = x] \\
&= J(x) - PJ(x), \qquad x \in \mathbb{R}_+. \tag{23}
\end{aligned}$$

The smoothed estimator was introduced for the single queue in [20], [18], where it was shown that the variance is brought down to $(1 - \rho)^{-2}/n$ for $\rho \sim 1$. Generalizations, analysis, and various numerical experiments can be found in [21], [19], [5], [26], [27].

Here we will introduce a completely new application of shadow functions. As we shall see, the result of Theorem 2.2 motivates the formulation of *sampling rules* for simulation. In the context of the present example, we have the following sampling criterion; its more general version is in the next section.

SHADOW FUNCTION SAMPLING CRITERION
Fix a set $B \in \mathbb{R}_+$ on which $\|F\|_{\infty, B}$ is finite, and also fix a constant $u > 0$. Then sample from the estimator $\Gamma_n(F)$ only at times $n \geq 1$ such that,

$$|\Gamma_n(\mathcal{D})| \leq u \quad and \quad \Phi_{n-1} \in B,$$

where the shadow function $\mathcal{D}$ is defined by $\mathcal{D}(x) = J(x) - PJ(x)$, $x \in \mathsf{X}$, with J given by (22).

The probability of interest in evaluating the estimate at a sampling time is,

$$\mathsf{P}_x\{\Gamma_n(F) \geq \pi(F) + \epsilon, \ |\Gamma_n(\mathcal{D})| \leq u, \ \Phi_{n-1} \in B\}. \tag{24}$$

Given a precise bound on this probability, we thus obtain a precise guarantee on the probability that the sample averages of interest are within ϵ of $\pi(F)$ at any of the sampling times defined by the shadow function sampling rule. The simulation is stopped when that guarantee is satisfactory.

Simulation experiments were performed for the simple discrete-time queue (19), as described above. Shown in Figure 1 is a plot of the sample path average $\Gamma_n(F)$ with $F(x) \equiv x$. Observe that whenever $\Gamma_n(F)$ is sampled at time-points in the grey region, its value is close to the steady state mean $\pi(F)$.

Note that the sampling rule can also be used for simultaneous simulation of several functions of Φ. Shown in Figure 2 are results from an experiment using the two functions $F(x) \equiv x$ and $G(x) \equiv x \sin(0.1x)$. The function G is not globally Lipschitz continuous, a property that that is used in the construction and evaluation of the smoothed estimator. Hence it is not obvious how one would construct a shadow function for estimating the mean $\pi(G)$.
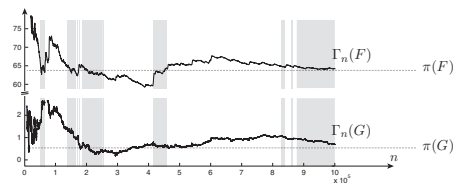

Fig. 2: Application of the shadow function sampling criterion with multiple functions. The simple queue (19) with $\kappa = 1$ was simulated to obtain estimates of the steady state mean of the two functions $F(x) \equiv x$, and $G(x) \equiv x \sin(0.1x)$. The grey regions indicate time periods during which $|\Gamma_n(\mathcal{D})| \leq 1$ and $\Phi_{n-1} \in B := [0, 10^4]$.

B. The General Sampling Criterion

As mentioned already, our starting point for the development of a general sampling rule for the estimator

$\Gamma_n(F)$ will be to obtain a bound on the deviation probability (24) with an appropriately chosen shadow function $\mathcal{D}$. This bound, presented in the following section, will be derived for the class of "geometrically ergodic" Markov chains, a family much broader than the Doeblin chains considered earlier. A ψ-irreducible, aperiodic chain $\mathbf{\Phi} = \{\Phi_i : i \geq 0\}$ is called *geometrically ergodic*, if there exists a function $V: \mathsf{X} \to [1, \infty)$, positive constants δ, b, and a *small* set $S \in \mathcal{B}$ such that

$$PV(x) - V(x) \leq -\delta V(x) + b\mathbb{I}_S(x), \quad x \in \mathsf{X}. \quad (25)$$

Recall that the set S is called small if there exists an integer $n_0 \geq 1$, an $a_0 > 0$ and a probability distribution ϱ on X such that

$$P^{n_0}(x, A) \geq a_0\varrho(A), \quad \text{for all } x \in S, \ A \in \mathcal{B}. \quad (26)$$

The function V satisfying (25) is called a *Lyapunov function* for the chain, and the bound (25) is the *Foster-Lyapunov drift condition (V4)* of [28]. Note that the Doeblin property of the previous section is equivalent to the requirement that (V4) is satisfied with V bounded, and also to the condition that the state space X is itself a small set; see [28] for details.

It is easy to verify that in most practical applications, the underlying Markov chain does indeed satisfy condition (V4). Moreover, in many applications where detailed quantitative results are sought, it is useful to find a Lyapunov function V satisfying (V4), even if the chain is Doeblin or finite-valued.

The main difference between the earlier queueing example and the general stopping rule we seek, is that, in general, we cannot easily find a value function for an associated fluid model. Instead, we define a shadow function $\mathcal{D}$ by analogy to the earlier definition in (23), but with J replaced by the Lyapunov function V:

$$\mathcal{D}(x) = V(x) - PV(x), \quad x \in \mathsf{X}. \quad (27)$$

Observe that the mean of $\mathcal{D}$ under π is always equal to zero, since $\pi(V) < \infty$ under (V4).

With this terminology in place, we can now state:

SHADOW FUNCTION SAMPLING CRITERION

(I) ASSUMPTIONS

Let $\mathbf{\Phi}$ be a geometrically ergodic chain, so that (V4) *holds with respect to the Lyapunov function V. Suppose that the function F is dominated by $\sqrt{V}$ in the sense that $\sup_x[F^2(x)/V(x)] < \infty$, and assume that the set $B \in \mathcal{B}$ is such that $\pi(B) > 0$ and F is bounded on B.*

(II) SAMPLING

Fix a constant $u > 0$, and sample from the estimator $\Gamma_n(F)$ only at times $n \geq 1$ such that,

$$|\Gamma_n(\mathcal{D})| \leq u \quad and \quad \Phi_{n-1} \in B,$$

where the shadow function $\mathcal{D}$ is defined by $\mathcal{D}(x) = V(x) - PV(x)$, $x \in \mathsf{X}$.

C. Finite-n Performance of the Sampling Criterion

In Theorem 3.1 below, we provide a non-asymptotic bound, which furnishes our sampling rule with a precise performance guarantee, and at the same time explains the motivation for its form.

For a geometrically ergodic chain the Lyapunov function V is used to construct the Banach space of measurable functions $F: \mathsf{X} \to \mathbb{R}$ for which the norm

$$\|F\|_V := \sup_{x \in \mathsf{X}} \frac{|F(x)|}{V(x)}, \quad (28)$$

is finite [16], [28].

Writing $1 \otimes \pi$ for the rank-one kernel defined by,

$$[1 \otimes \pi](x, A) = \pi(A), \quad x \in \mathsf{X}, \ A \in \mathcal{B},$$

the *fundamental kernel Z* is defined to be the inverse $Z = [I - P + 1 \otimes \pi]^{-1}$. That is, for any function F in the domain of Z, the function $\widehat{F} := ZF$ satisfies *Poisson's equation*,

$$P\widehat{F} = \widehat{F} - F + \pi(F). \quad (29)$$

The domain of Z is equal to the range space of $[I - P + 1 \otimes \pi]$, of which appropriate subsets can be identified under reasonable conditions. For any function $W: \mathsf{X} \to [1, \infty)$ we define the induced operator norm of a linear operator Q on L_∞^W, as,

$$\|Q\|_W = \sup\{\|QF\|_W : \|F\|_W \leq 1\}. \quad (30)$$

Under geometric ergodicity, the fundamental kernel is a bounded linear operator on L_∞^V, and for any $F \in L_\infty^V$ we have the representation,

$$ZF(x) = \sum_{k=0}^{\infty} [P^k F(x) - \pi(F)]. \quad (31)$$

We assume moreover that the function F^2 is dominated by the function V in the sense that $\|F^2\|_V < \infty$. These assumptions are identical to those used in Theorem 17.5.4 of [28] to show that the central limit theorem (CLT) holds for the sample averages $\Gamma_n(F)$. While these conditions are non-minimal for the CLT, they are easily verified in many cases.

We now present a non-asymptotic bound for the probabilities appearing in (24):

Theorem 3.1: Suppose that the chain $\mathbf{\Phi}$ is geometrically ergodic, so that (25) holds for some function V, and the function F satisfies $\|F^2\|_V < \infty$. Then with the shadow function $\mathcal{D}$ defined as in (27), we have for any $\epsilon > 0$, $u > 0$, that the logarithm of the probability,

$$\mathsf{P}_x\{\Gamma_n(F) - \pi(F) \geq \epsilon, \ |\Gamma_n(\mathcal{D})| \leq u, \ \Phi_{n-1} \in B\}$$

is bounded above by

$$-\frac{n-1}{2}\left[\frac{1}{8}\frac{\delta}{\xi}\frac{1}{\|\widetilde{F^2}\|_V}\left(\frac{[\epsilon^n]_+^2}{u^n+b}\right)-\frac{2}{n-1}\right]_+^2,$$

for all $n \geq 4b/\delta$, where $\epsilon^n := \epsilon - \|F\|_{\infty,B}/(n-1)$ and $u^n := u + \|\mathcal{D}\|_{\infty,B}/(n-1)$; b and δ are defined in *(V4)*; $\widetilde{F} = F - \pi(F)$; and $\xi = \|Z[I - 1 \otimes \pi]\|_{\sqrt{V}}$.

It is somewhat remarkable that it is possible to establish a large deviations upper bound at this level of generality, since even the standard asymptotic large deviations principle may well fail even for a bounded function of a stationary Doeblin chain. Moreover, as shown in the limit (20) for the simple queue, it may be impossible to have a meaningful (positive, that is) exponent for the deviations of an unbounded function of a geometrically ergodic chain. Theorem 3.1 is especially surprising since $\{|\Gamma_n(\mathcal{D})| \leq u\}$ is *not* a rare event when $u > 0$.

PROOF OUTLINE. The proof of Theorem 3.1 is based on essentially the same four steps used in the proof of Theorem 2.4.

Step 1 is an application of Theorem 2.2 to establish that the negative logarithm of the probability of interest is bounded below by

$$(n-1)H(\gamma\|\gamma^1 \odot P), \qquad (32)$$

where γ is a bivariate probability measure on $(\mathsf{X} \times \mathsf{X}, \mathcal{B} \times \mathcal{B})$, whose marginals γ^1 and γ^2 satisfy,

$$\|\gamma^1 - \gamma^2\| \leq \frac{2}{n-1}, \\ \gamma^1(F) - \pi(F) \geq \epsilon^n, \quad |\gamma^1(\mathcal{D})| \leq u^n. \qquad (33)$$

Consider the remaining three steps: We again apply Pinsker's inequality to obtain a lower bound analogous to (14), and Step 3 is the construction of a linear program to obtain a lower bound on (14). Following the same steps presented in Section II-B we find that a bound analogous to (16) holds, and applying (32) we conclude that the solution of the following linear program provides a bound on (16):

$$\begin{aligned} \min \quad & \|\mu - \mu P\| \\ \text{s.t.} \quad & \mu(F) - \pi(F) \geq \epsilon^n, \quad |\mu(\mathcal{D})| \leq u^n. \end{aligned} \qquad (34)$$

On letting $L^* = L^*(\epsilon^n, u^n)$ denote the value of this linear program, we again arrive at the bound (18), for any γ satisfying the constraints (33).

Hence, we arrive at Step 4: to complete the proof of Theorem 3.1 it is enough to establish the following:

Proposition 3.2: Suppose that *(V4)* holds and that $\|F\|_{\sqrt{V}} < \infty$. Then the following lower bound on the

value of the linear program in *(34)* holds for all $\epsilon > 0$, $u > 0$:

$$L^*(\epsilon, u) \geq \frac{1}{8}\frac{\delta}{\xi}\frac{1}{\|\widetilde{F^2}\|_V}\left(\frac{\epsilon^2}{u+b}\right).$$

$\square$

REFERENCES

[1] S. Asmussen. Queueing simulation in heavy traffic. *Math. Operations Res.*, 17:84–111, 1992.

[2] J.R. Baxter, N.C. Jain, and S.R.S. Varadhan. Some familiar examples for which the large deviation principle does not hold. *Comm. Pure Appl. Math.*, 44(8-9):911–923, 1991.

[3] P. Brémaud. *Markov chains*, volume 31 of *Texts in Applied Mathematics*. Springer-Verlag, New York, 1999. Gibbs fields, Monte Carlo simulation, and queues.

[4] W. Bryc and A. Dembo. Large deviations and strong mixing. *Ann. Inst. H. Poincaré Probab. Statist.*, 32(4):549–569, 1996.

[5] M. Chen, I.-K. Cho, and S.P. Meyn. Reliability by design in a distributed power transmission network. *Automatica*, 42:1267–1281, August 2006. (invited).

[6] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. J. Wiley, New York, 1991.

[7] I. Csiszár. Sanov property, generalized I-projection and a conditional limit theorem. *Ann. Probab.*, 12(3):768–793, 1984.

[8] I. Csiszár. The method of types. *IEEE Trans. Inform. Theory*, 44(6):2505–2523, 1998. Information theory: 1948–1998.

[9] I. Csiszár, T.M. Cover, and B.S. Choi. Conditional limit theorems under Markov conditioning. *IEEE Trans. Inform. Theory*, 33(6):788–801, 1987.

[10] A. de Acosta and P. Ney. Large deviation lower bounds for arbitrary additive functionals of a Markov chain. *Ann. Probab.*, 26(4):1660–1682, 1998.

[11] M.D. Donsker and S.R.S. Varadhan. Asymptotic evaluation of certain Markov process expectations for large time. I. II. *Comm. Pure Appl. Math.*, 28:1–47; ibid. **28** (1975), 279–301, 1975.

[12] P. Dupuis and R.S. Ellis. *A weak convergence approach to the theory of large deviations*. John Wiley & Sons Inc., New York, 1997.

[13] P. Dupuis and O. Zeitouni. A nonstandard form of the rate function for the occupation measure of a Markov chain. *Stochastic Process. Appl.*, 61:249–261, 1996.

[14] R.S. Ellis. *Entropy, large deviations, and statistical mechanics*. Springer-Verlag, New York, 1985.

[15] P. Glynn and R. Szechtman. Some new perspectives on the method of control variates. In K.T. Fang, F.J. Hickernell, and H. Niederreiter, editors, *Monte Carlo and Quasi-Monte Carlo Methods 2000: Proceedings of a Conference held at Hong Kong Baptist University, Hong Kong SAR, China*, pages 27–49, Berlin, 2002. Springer-Verlag.

[16] P. W. Glynn and S. P. Meyn. A Liapounov bound for solutions of the Poisson equation. *Ann. Probab.*, 24(2):916–931, 1996.

[17] P.W. Glynn and D. Ormoneit. Hoeffding's inequality for uniformly ergodic Markov chains. *Statist. Probab. Lett.*, 56(2):143–146, 2002.

[18] S. G. Henderson and P. W. Glynn. Approximating martingales for variance reduction in Markov process simulation. *Math. Operations Res.*, 27(2):253–271, 2002.

[19] S. G. Henderson, S. P. Meyn, and V. B. Tadić. Performance evaluation and policy selection in multiclass networks. *Discrete Event Dynamic Systems: Theory and Applications*, 13(1-2):149–189, 2003. Special issue on learning, optimization and decision making (invited).

[20] S.G. Henderson. *Variance Reduction Via an Approximating Markov Process*. PhD thesis, Stanford University, Stanford, California, USA, 1997.

[21] S.G. Henderson and S.P. Meyn. Variance reduction for simulation in multiclass queueing networks. To appear, *IIE Trans. on Operations Engineering*, 2004.

[22] W. Hoeffding. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.*, 58:13–30, 1963.

[23] I. Iscoe, P. Ney, and E. Nummelin. Large deviations of uniformly recurrent Markov additive processes. *Adv. Appl. Math.*, 6:373–412, 1985.

[24] A. M. Law and W. D. Kelton. *Simulation Modeling and Analysis*. McGraw-Hill, New York, 3rd edition, 2000.

[25] J.S. Liu. *Monte Carlo strategies in scientific computing*. Springer Series in Statistics. Springer-Verlag, New York, 2001.

[26] S. P. Meyn. Large deviation asymptotics and control variates for simulating large functions. *Ann. Appl. Probab.*, 16:310-339 2006.

[27] S. P. Meyn. Workload models for stochastic networks: Value functions and performance evaluation. *IEEE Trans. Automat. Control*, 50(8):1106– 1122, August 2005.

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

[29] P. Ney and E. Nummelin. Markov additive processes. I. Eigenvalue properties and limit theorems. *Ann. Probab.*, 15(2):561–592, 1987.

[30] P. Ney and E. Nummelin. Markov additive processes II. Large deviations. *Ann. Probab.*, 15(2):593–609, 1987.

[31] C. Pandit. *Robust Statistical Modeling Based On Moment Classes With Applications to Admission Control, Large Deviations and Hypothesis Testing*. PhD thesis, Dept. of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, 2004.

[32] S.R.S. Varadhan. *Large Deviations and Applications*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, Pa., 1984.

[33] S.R.S. Varadhan. Large deviations and entropy. In *Entropy*, Princeton Ser. Appl. Math., pages 199–214. Princeton Univ. Press, Princeton, NJ, 2003.

[34] W. Whitt. Planning queueing simulations. *Management Science*, 35:1341–1366, 1994.

[35] G. Winkler. *Image Analysis, Random Fields and Dynamic Monte Carlo Methods: A Mathematical Introduction*. Springer-Verlag, Berlin, 1995.