2011 American Control Conference
on O'Farrell Street, San Francisco, CA, USA
June 29 - July 01, 2011

Spectral and Graph-Theoretic Bounds on Steady-State-Probability Estimation Performance for an Ergodic Markov Chain

Mengran Xue, Sandip Roy

Abstract— **We pursue a spectral and graph-theoretic performance analysis of a classical estimator for Markov-chain steady-state probabilities. Specifically, we connect a performance measure for the estimate to the structure of the underlying graph defined on the Markov chain's state transitions. To do so, 1) we present a series of upper bounds on the performance measure in terms of the subdominant eigenvalue of the state transition matrix, which is closely connected with the graph structure; 2) as an illustration of the graph-theoretic analysis, we then relate the subdominant eigenvalue to the connectivity of the graph, including for the strong-connectivity case and the weak-link case. We also apply the results to characterize estimation in Markov chains with rewards.**

I. INTRODUCTION

Markov chains have proved useful for modeling stochastics in a very wide range of application areas. Unfortunately, many Markov chain models in real applications are "hidden" from outside (i.e., the full transition matrix is unknown, or some part of the matrix is unknown), or the updating rules are very complex (i.e., the number of states is very large), both of which make it impossible or difficult to directly analyze the Markov chain model. Therefore, a broad class of *inference* problems for Markov chains have been intensively studied [1]–[4]. Within this class of inference problems, the problem of estimating the steady-state probabilities of a Markov chain is often of particular interest, because the asymptotic behavior of a Markov chain must be characterized in a wide range of applications. This article is concerned with the steady-state probability estimation problem for ergodic Markov chains.

A rich literature has been specifically dedicated to the steady-state estimation problem [5]–[10]. For instance, many estimators for the steady state of a Markov chain have been proposed and thoroughly analyzed [11], [12]. Of particular interest to us, several efforts have sought to characterize the performance of common steady-state-probability estimators, including specifically estimators based on sample state-occupancy frequencies. Specifically, these estimators' performances can be related to long-run time averages or *central-limit theorems* for Markov chains, and in this way the performance can be connected with Markov chain metrics. For instance, in the recent work [13], Doyle summarizes some results for the *Kemeny constant* (or the seek time from any initial state to the steady state) of an ergodic Markov chain, and presents a connection between the Kemeny constant and the *central limit theorem* (CLT) for Markov chains. From an entirely different perspective, performance analysis of steady-state probability estimators has been motivated by the importance of Markov chains in simulations (e.g., MCMC). In this context, steady-state probability estimator performance has been related to some extent to the spectrum of the Markov chain's transition matrix [14]–[16]. We note that the convergence rate of a Markov chain and the mixing time to the steady state have also been frequently studied [20], [21] in the context of MCMC simulations.

Here, we revisit the Markov chain steady-state-probability estimation problem, but take a new graph-theoretic perspective. Specifically, we relate the underlying graphical topology of the Markov chain (i.e., the graph associated with its state transitions) to the state-occupancy-frequency-based steady state estimator and its performance (as measured by its error covariance). As a further step, we also apply the results to Markov chains with rewards, in particular giving spectral and graphical characterizations of expected reward estimator (or, equivalently, of sample-average reward values). Although the estimators themselves are already well known, spectral and graphical characterizations of performance are valuable in that they provide broad intuition into estimator structure/design without requiring specific knowledge of model parameters. Such a connection is sensible and meaningful for several reasons, including because graphical structure of state transitions is often more available than actual transition probabilities in many applications (especially for a Markov chain requiring estimation or one with a large number of states); because graph-based analyses can provide more accurate bounds than other methods in predicting certain behaviors of a Markov chain; and because graph-theoretic analysis and in turn proper graphical design can help us achieve desired asymptotic inference properties.

Of note, the steady state estimation problem studied here is in fact connected to a broad class of recent work on network estimation and estimator performance analysis (see e.g. [22], [23]). Also, we stress that the results here have specific application in several domains, including in characterizing Markov Chain Monte Carlo methods, queueing-network analysis, jump-Markov modeling, and network control (see e.g. [18]). We omit the details of these connections/applications in the interest of space, and ask the reader to see [31].

The rest of the article is organized as follows. In Section II, we formulate the steady-state-probability estimation problem for Markov chains in generality, and introduce a

The authors are with Department of Electrical Engineering, Washington State University, Pullman, WA. This project was partially supported by NSF grants ECS-0725589, ECS-0901137, CNS-1035369, and CNS-1058124. Please send correspondence to mxue@eecs.wsu.edu .

classical unbiased estimator and its performance analysis. In Section III, we give a full eigenstructure-based analysis for the performance measures (Section III-A); we then present several graphical results on the estimator performance based on the spectral results (Section III-B). We also enhance the development for the reward-inference application (Section IV).

II. Problem Formulation

In this section, we first review the Markov chain steady-state-probability estimation problem (Section II-A). We then invoke a classical unbiased estimator for this problem, and present algebraic expressions for two performance measures of the estimator (Section II-B).

In the remainder of the article, we limit ourselves to discrete-time, finite-state ergodic Markov chains (i.e., ones that are irreducible and aperiodic), which are very common models in applications requiring steady-state probability estimation (see e.g. [24] for an effort in the infinite-state-space case). We note that state occupancy probabilities for ergodic finite-state Markov chains necessarily approach fixed constants asymptotically, i.e. the probabilities have a steady-state.

A. Markov Chain Steady State Estimation: Overview

We consider a discrete-time ergodic Markov chain with m states, labeled $1, \cdots, m$. Let the matrix $D \in \mathbf{R}^{m\times m}$ be the state transition matrix of the Markov chain. We refer to the i,jth entry of D, d_{ij}, as the transition probability from state i to state j. We note that D is a *row-stochastic matrix* here. Moreover, we define a graph G for this Markov chain, as is classical in the study of Markov chains [25]. We consider a weighted and directed graph G consisting of m nodes, each of which represents one state of the Markov chain. Between any ordered pair (i,j) of the nodes, a directed edge is drawn from node i to node j if and only if $d_{ij} > 0$, and the edge is given weight d_{ij}. The Markov chain's state is assumed to evolve along a discrete time axis $k = 0, 1, \ldots$, according to the specified transition probabilities. For notational convenience, here we use a $0-1$ indicator vector $\mathbf{s}[k]$ (referred to as the *state vector*) to represent the state at time k. That is, if the state at time k is i, the state vector $\mathbf{s}[k]$ has ith entry equal to 1 and is zero elsewhere. We also define a **probability vector** $\mathbf{p}[k]$ which specifies the state occupancy probabilities of the Markov chain at each time k; we note that $\mathbf{p}[k] = E[\mathbf{s}[k]]$. Finally, we use the notation $\pi = [\pi_1, \cdots, \pi_m]^T$ to represent the *steady state probability vector* of the ergodic Markov chain, i.e. $\pi = \lim_{k\to\infty} \mathbf{p}[k]$.

We will study the problem of estimating the steady-state probabilities of the Markov chain, from observations of its state. Specifically, let us assume that the Markov chain's transition matrix is unknown, or that the steady-state probabilities or statistics defined thereof are too cumbersome to compute analytically for the Markov chain. Instead, the steady-state probabilities must be obtained from observations of the state. Let us also assume that we are able to observe the state of the Markov chain at each time over a time-interval of N steps, i.e. we observe $\mathbf{s}[k]$ for $k = 1, \cdots, N$. We are concerned with estimation of the steady-state probability vector π of the Markov chain from the sequence of observations $\mathbf{s}[1], \cdots, \mathbf{s}[N]$, and evaluation of the estimator's performance in terms of the graphical structure of G. That is, we seek to construct an estimate $\mathbf{p}$ through a functional mapping from the observation sequence, and further to characterize the error in this estimate in terms of the graph structure. We refer this estimation problem as the **Markov chain steady state (MCSS) estimation** problem.

The MCSS estimation problem has been thoroughly studied, and numerous estimators have been developed for the problem. However, our main focus is on the graphical characterizations of the estimator and its performance. Therefore, we focus on the classical unbiased estimator for the MCSS estimation problem, and pursue extensive spectral and graph-theoretic characterization of this estimator. We introduce the estimator, and develop its basic performance analysis, in the following subsection.

B. Estimator and Performance Measures

Let us first recall a classical estimator for the MCSS estimation problem, that we will use in our development: $\mathbf{p} = \frac{1}{N}\sum_{k=1}^{N} \mathbf{s}[k]$. That is, the estimator simply averages all the available states to form the estimate $\mathbf{p}$, or in other words estimates each state's asymptotic probability as the frequency that the Markov chain is in the state during the observation period. More details about this sample-frequency-based estimator can be found in [12]. Of particular note, the sample-frequency-based estimator is asymptotically optimal, in the sense that its performance approaches the Cramer-Rao bound in the limit of large N.

We will characterize the performance of this estimator, under the assumption that the Markov chain is initially in steady state, i.e. the state occupancy probabilities at the initial time, $\mathbf{p}[0]$, are equal to the steady-state probabilities π. We focus on this special case 1) because unknown Markov processes in the environment typically satisfy this assumption, and 2) to explicitly delineate the role of the data stochasticity in estimator performance as compared to the role of the settling to steady-state (or *mixing*). Even in cases where a Markov chain is not in steady-state when observations are commenced (such as in the MCMC applications), this assumption is often not greatly limiting because: 1) the mixing time of the Markov chain is well characterized (including in terms of the graph structure) and is usually small compared to the estimation time [14], [15]; and 2) in fact algorithms for perfect mixing within a small number of steps can be designed in a range of computational applications. It is worth noting that a range of MCMC applications in which rare events are captured require use of slowly-mixing chains; while in some of these cases settling is still relatively fast compared to steady-state probability estimation, the conditions for fast mixing and fast estimation are distinct (as our results show) and indeed mixing rather than steady-state probability computation may be the dominant temporal constraint. We caution that the mixing-time analysis must be considered with care in these cases (e.g. by drawing on [14], [15]); we hope to return to this issue in future work, but do not consider it further here in the interest of focusing on steady-state probability estimation.

Under the assumption that the initial state occupancy probabilities are the steady-state ones, it can be easily shown that $E[\mathbf{p}] = \pi$ for any Markov chain and any N, thus indicating that the estimator is unbiased. Here, we will consider two sensible (and classical) performance measures for the estimate $\mathbf{p}$, namely its **error covariance matrix** COV (defined as $E[(\mathbf{p}-\pi)(\mathbf{p}-\pi)^T]$) and the *expected total squared error* (defined as $E[(\mathbf{p}-\pi)^T(\mathbf{p}-\pi)]$). We note that the expected total squared error is equal to the trace of the covariance matrix, and so we use the notation $tr(COV)$ for it. Before we derive expressions for the two performance measures COV and $tr(COV)$, for convenience's sake, let us define the following diagonal matrix here: $\Delta = diag\{\pi_i\}$, $i = 1,\cdots,m$.

Let us develop an expression for the error covariance matrix COV first, as follows: $COV = E[(\mathbf{p}-\pi)(\mathbf{p}-\pi)^T] = E[\mathbf{p}\mathbf{p}^T] - \pi\pi^T = E\left[\left(\frac{1}{N}\sum_{k=1}^{N}\mathbf{s}[k]\right)\left(\frac{1}{N}\sum_{k=1}^{N}\mathbf{s}^T[k]\right)\right] - \pi\pi^T = \frac{1}{N^2}E\left[\left(\sum_{k=1}^{N}\mathbf{s}[k]\right)\left(\sum_{k=1}^{N}\mathbf{s}^T[k]\right)\right] - \pi\pi^T$. We also have $E[\mathbf{s}[k]\mathbf{s}^T[j]] = \begin{cases} \Delta, & \text{for } k=j; \\ (D^T)^{k-j}\Delta, & \text{for } k>j; \\ \Delta D^{j-k}, & \text{for } k<j. \end{cases}$ With a little algebraic effort, we can rewrite COV in the following form:

$$\begin{aligned} & COV \\ = \; & \frac{1}{N^2}E\Big[\sum_{r=1}^{N-1}\sum_{k=1}^{N-r}\mathbf{s}[k]\mathbf{s}^T[k+r] + \sum_{r=1}^{N-1}\sum_{k=1}^{N-r}\mathbf{s}[k+r]\mathbf{s}^T[k] \\ & + \sum_{k=1}^{N}\mathbf{s}[k]\mathbf{s}^T[k]\Big] - \pi\pi^T \\ = \; & \frac{1}{N^2}\left[\sum_{r=1}^{N-1}(N-r)\Delta D^r + \sum_{r=1}^{N-1}(N-r)(D^T)^r\Delta + N\Delta\right] - \pi\pi^T \end{aligned} \quad (1)$$

Based on the expression for the error covariance matrix, the total squared error $tr(COV)$ can be automatically written as: $tr(COV) = \frac{1}{N^2}\left(2\sum_{r=1}^{N-1}(N-r)\sum_{i=1}^{m}e_iD^re_i\pi_i + N\right) - \pi^T\pi$. We have thus presented algebraic expressions for a sample-frequency estimator of a Markov chain's steady-state probabilities, and for two performance measures of the estimator.

III. Graphical Results on MCSS Estimation

In this section, we first give an eigenstructural analysis of the performance of the MCSS estimator. Then, based on this spectral performance analysis, we develop several graph-theoretic bounds on the performance by exploiting relationships between the non-unity eigenvalues and their associated eigenvectors and the graph topology. We note that eigenstructural and consequent graphical characterizations have been obtained comprehensively for Markov chain mixing times [14], [15], but to the best of our knowledge a comprehensive study of MCSS estimation in this direction has not been attempted. We point the reader to [13], [20], [21] for relevant preliminary work (which comes from the MCMC and MC CLT literature).

A. Spectral Performance Analysis

Before presenting the spectral analysis, let us clarify several notations. Let us again consider an ergodic Markov chain as described in the MCSS estimation problem, and invoke the MCSS estimator introduced above. Let $\lambda_1 = 1, \lambda_2, \cdots, \lambda_m$ be the m eigenvalues of the transition matrix D, ordered for convenience in increasing distance from the point $1 + j0$ in the complex plane (i.e., so that $|1-\lambda_1| < |1-\lambda_2| \le \cdots \le |1-\lambda_m|$). In developing graph-theoretic bounds, we will actually find it convenient to perform the spectral analysis in the case that the eigenvalues of D are simple (in Jordan blocks of size 1), and to subsequently argue that the bounds encompass the general case through a perturbation analysis. In the simple-eigenvalue case, each λ_i has an associated right eigenvector and a left eigenvector, which we call $\mathbf{v}_i$ and $\mathbf{w}_i^T$, for $i = 1,\cdots,m$. Since D is the transition matrix of an ergodic Markov chain, the unity eigenvalue $\lambda_1 = 1$ is unique and strictly dominant, with $\mathbf{v}_1 = \vec{1}$ and $\mathbf{w}_1 = \pi > 0$, where $\vec{1}$ is a column vector with all unity entries. Let us also define matrix $V = \begin{bmatrix}\mathbf{v}_1 & \cdots & \mathbf{v}_m\end{bmatrix}$, matrix $W = \begin{bmatrix}\mathbf{w}_1 & \cdots & \mathbf{w}_m\end{bmatrix}^T$, and matrix $\Lambda = diag\{\lambda_i\}$ (and where $W = V^{-1}$ also). Then, we immediately have $D = V\Lambda W$.

Now, we are ready to use the eigendecomposition of matrix D to analyze the estimator performance in the case that eigenvalues are simple. Specifically, in the expression for COV (Equation 1), the two summation terms can be rewritten as follows: $\sum_{r=1}^{N-1}(N-r)\Delta D^r = \Delta\sum_{r=1}^{N-1}(N-r)V\Lambda^rW = \Delta V\left(\sum_{r=1}^{N-1}(N-r)\Lambda^r\right)W = \Delta VBW$, and similarly $\sum_{r=1}^{N-1}(N-r)(D^T)^r\Delta = W^TBV^T\Delta$, where B is a diagonal matrix $diag\{\beta_i\}$, with $\beta_1 = \frac{N(N-1)}{2}$ and $\beta_i = \frac{\lambda_i^{N+1} - N\lambda_i^2 + (N-1)\lambda_i}{(1-\lambda_i)^2}$, for $i = 2,\cdots,m$.

Let $\gamma_i = \frac{\beta_i}{N^2} + \frac{1}{2N}$, for all i, and $\Gamma = diag\{\gamma_i\}$. We then can rewrite the matrix COV as follows:

$$COV = \Delta V\Gamma W + W^T\Gamma V^T\Delta - \pi\pi^T. \quad (2)$$

Moreover, noticing that $\beta_1 = \frac{N(N-1)}{2}$ and $\gamma_1 = \frac{\beta_1}{N^2} + \frac{1}{2N} = \frac{1}{2}$ for all N, we can obtain an expression for COV that more explicitly reflects its eigenstructure. In particular, COV now becomes the following:

$$\begin{aligned} COV &= \Delta V\Gamma W + W^T\Gamma V^T\Delta - \pi\pi^T \\ &= \sum_{i=1}^{m}\gamma_i(\Delta\mathbf{v}_i\mathbf{w}_i^T + \mathbf{w}_i\mathbf{v}_i^T\Delta) - \pi\pi^T \\ &= \sum_{i=2}^{m}\gamma_i(\Delta\mathbf{v}_i\mathbf{w}_i^T + \mathbf{w}_i\mathbf{v}_i^T\Delta) + \left(\frac{1}{2}\Delta\vec{1}\pi^T + \frac{1}{2}\pi\vec{1}^T\Delta - \pi\pi^T\right) \\ &= \sum_{i=2}^{m}\gamma_i(\Delta\mathbf{v}_i\mathbf{w}_i^T + \mathbf{w}_i\mathbf{v}_i^T\Delta), \end{aligned} \quad (3)$$

where $\gamma_i = \frac{\lambda_i^{N+1} - N\lambda_i^2 + (N-1)\lambda_i}{N^2(1-\lambda_i)^2} + \frac{1}{2N}$, $i = 2,\cdots,m$.

Using the expression for COV, the second performance measure $tr(COV)$ can be expressed as follows (in the case that the eigenvalues of D are simple):

$$tr(COV) = \sum_{i=2}^{m}2\gamma_i tr(\Delta\mathbf{v}_i\mathbf{w}_i^T) = 2\sum_{i=2}^{m}\gamma_i\left(\sum_{j=1}^{m}v_{ij}w_{ij}\pi_j\right), \quad (4)$$

where v_{ij} and w_{ij} are the jth entry in right eigenvector $\mathbf{v}_i$ and left eigenvector $\mathbf{w}_i$.

As a preliminary step in characterizing the error covariance, let us study the asymptotics of the covariance (i.e., the value of COV as $N \to \infty$). Since $\Gamma \to \begin{bmatrix} \frac{1}{2} & \\ & \mathbf{0} \end{bmatrix}$ as $N \to \infty$, we have $COV = \frac{1}{2}\Delta \mathbf{v}_1 \mathbf{w}_1^T + \frac{1}{2}\mathbf{w}_1 \mathbf{v}_1^T \Delta - \pi\pi^T$. Noting that $\mathbf{v}_1 = \vec{1}$ and $\mathbf{w}_1 = \pi$, we have $COV \to \frac{1}{2}\Delta\vec{1}\pi^T + \frac{1}{2}\pi\vec{1}^T\Delta - \pi\pi^T = \mathbf{0}$. Thus, as expected, the estimate becomes increasingly accurate as more data is used; in fact, it is easy to check that asymptotically the error covariance decreases with $\frac{1}{N}$, see also e.g. [20].

The core question of interest to us is how the performance of the estimator—i.e., the rate of convergence of the error to nil—depends on the graphical structure of the Markov chain. Equation 3 and Equation 4 provide a route for connecting the estimator performance to the underlying graph structure defined from the transition matrix D, because the non-unity eigenvalues/eigenvectors of D can be connected to the graph structure using algebraic graph theory constructs. In the following part of this section, the graphical characterizations are developed by first developing simpler bounds in terms of particular eigenvalues, and then exploiting the connections between eigenvalues and graph structure.

We note that the above analysis can easily be extended to the case that some non-unity eigenvalues of D are not simple (i.e., are in Jordan blocks of size 2 or larger). We kindly ask the reader to see [31] for the details, which are essentially similar to the simple-eigenvalue case but are needed for proofs of subsequent theorems.

B. Main Results

Let us here focus on developing spectral and graphical bounds for the total squared error measure, $tr(COV)$, since 1) the trace measure aggregates the error covariance performance information into a single, useful scalar; and 2) it permits a more comfortable analysis, which can naturally be generalized to the matrix case. In the interest of space, we omit all the proofs. Please see [31] for the proofs.

First, let us develop general lower and upper bounds for $tr(COV)$, in terms of the eigenvalues of the state transition matrix. We then modify these general bounds in order to relate them to the second largest eigenvalue of matrix D. Here is the first theorem, which presents the general bounds.

Theorem 1: We consider the MCSS estimation problem described above. The trace of the error covariance matrix can be bounded as $2\pi_{min}\sum_{i=2}^{m}\gamma_i \le tr(COV) \le 2\pi_{max}\sum_{i=2}^{m}\gamma_i$, where $\pi_{min} = \min_j \{\pi_j\}$, and $\pi_{max} = \max_j \{\pi_j\}$, for $j = 1, \cdots, m$.

Next, building on Theorem 1, we focus on developing upper bounds on $tr(COV)$ that are phrased in terms of the "second-largest" eigenvalue λ_2 of the transition matrix D (i.e., the non-unity eigenvalue that is closest to $1 + j0$ in the complex plane); many relationships between this second-largest-magnitude eigenvalue and the network's topological structural are known, and so these bounds provide a stepping-stone toward graph-theoretic analysis. We begin with a basic upper bound on $tr(COV)$ in terms of the second largest eigenvalue, that is applicable to all ergodic Markov chains.

Theorem 2: We consider the MCSS estimation problem described above. Then the trace of the error covariance matrix can be bounded as $tr(COV) < \frac{2(m-1)}{N|1-\lambda_2|} + \frac{4(m-1)}{N^2|1-\lambda_2|^2}$.

Let us make several observations about the obtained bound. First, for large N, we note that squared error is inversely proportional to N, with the proportionality constant depending on the distance of λ_2 from 1. Meanwhile, the term scaling with $\frac{1}{N^2}$ is negligible for large N, but may contribute significantly to the error for small N, if λ_2 is near 1. It is interesting to note that the bound depends on the distance of λ_2 from 1, not the distance of the eigenvalues from the unit circle.

Also of interest, for many classes of Markov chains, the entries in the steady-state probability vector π_i can be upper bounded away from 1. The reader is referred to Minc's classical work [26] as well as the more recent work [27] for examples. For these cases, we note that a tighter bound on the performance measure in terms of the immediately follows from Theorem 1: $tr(COV) < (\max_i \pi_i)(\frac{2(m-1)}{N|1-\lambda_2|} + \frac{4(m-1)}{N^2|1-\lambda_2|^2})$.

Next, let us present a series of results for several special classes of Markov chains. We first consider the case where the eigenvalues of the matrix D are real. This class is of particular interest because it includes reversible Markov chains. For this case, we obtain a tighter bound:

Theorem 3: Let us assume that the Markov chain described in the MCSS estimation problem has a state transition matrix whose eigenvalues are all real. Then the trace of the error covariance matrix can be bounded as $tr(COV) < \frac{2(m-1)}{N(1-\lambda_2)} - \frac{m-1}{N} + \frac{4(m-1)}{N^2(1-\lambda_2)^2}$.

For the real eigenvalue case (and hence the reversible case), the bound in Theorem 3 may yield a significant improvement over the bound in Theorem 2 when N is small. Next, by building on the bound in Theorem 3, we obtain an even tighter bound in the case that the Markov chain's transition matrix is symmetric. We note that, if a Markov chain's transition matrix is known to be symmetric (i.e., the eigenvalues are simple), estimation of the steady state probabilities from data is not of interest, since the steady-state probabilities are already known to be identically $\frac{1}{m}$. However, the following bound is interesting in that is shows how accurately the steady-state probabilities would be obtained if the chain happened to be symmetric, even though it was not known to be so.

Theorem 4: We consider the Markov chain described in the MCSS problem. Here, we assume that the state transition matrix D is symmetric. Then, the trace of the error covariance matrix can be bounded as $tr(COV) < \frac{2(m-1)}{mN(1-\lambda_2)} - \frac{m-1}{mN} + \frac{4(m-1)}{mN^2(1-\lambda_2)^2}$.

Remark: The above result holds, more broadly, whenever the transition matrix is doubly stochastic.

We also provide a tighter bound in the case that the Markov chain has only two states. We note that the two-state case is particularly common in many applications (e.g., in communications applications).

Theorem 5: Consider the MCSS estimation problem, for a Markov chain with two states ($m = 2$). The performance measure $tr(COV)$ can be bounded as $tr(COV) < \frac{1}{2N(1-\lambda_2)} - \frac{1}{4N} + \frac{1}{N^2(1-\lambda_2)^2}$.

Remark: The above analyses/bounds on the performance measures also automatically yield bounds on the probability that the estimate deviates by a large amount from the true value, through the Chebyshev inequality. In the asymptotic case (N large), the CLT results for Markov chains further show that the estimate error is normally distributed, and so much tighter bounds on the probability of large deviation can be obtained, see e.g. [13] for some other results of this form.

In sum, we have shown that, for any fixed time N, $tr(COV)$ becomes small (since the performance bound becomes small) when the other eigenvalues of D are far away from the unity eigenvalue. These results provide a foundation for connecting the performance (or at least performance bound) with the graph-structural features of the Markov chain. A very wide range of results already exist, and many more can be imagined, that relate the Markov chain's graph structure with the subdominant eigenvalue λ_2, and hence with the estimator's performance. Here, let us only pursue a couple such results, with the primary aim of showing that highly connected Markov chains enjoy fast estimation. These graphical characterizations are all fundamentally based on showing the closeness or distance of λ_2 from unity in the complex plane. We build the first result using the following lemma (please see [28] for the detailed proof):

Lemma 1: For a stochastic matrix $D = \{d_{ij}\}$, the non-unity eigenvalues can be bounded as $|\lambda| \leq \frac{1}{2}\sum_{k=1}^{m} \max_{i,j} |d_{ik} - d_{jk}|$.

Lemma 1 indicates an upper bound on $|\lambda_i|$, for $i = 2, \cdots, m$, in terms of column-entry differences in D. We see that, when the graph is strongly connected, especially when the edge weights are evenly distributed, the bound turns out to be small. Hence, the second largest eigenvalue, as well as all the other non-unity eigenvalues, are far from unity, which also makes the upper bound on $tr(COV)$ in Theorem 2 small. In other words, strong connectivity of the graph defined on D makes the estimation easier, since transitions to all states are more frequent and temporally uncorrelated, and hence the data contains more information for estimation.

Unfortunately, the result based on Lemma 1 is fairly restrictive, in that it requires that all pairwise differences in a column of D are small. The following alternate characterization shows that λ_2 is not too close to 1 as long as products of probabilities between nodes are not too small.

Theorem 6: Consider an ergodic Markov chain described as in the MCSS estimation problem. We assume that, in the underlying directed graph G defined by the state transitions, the product of edge weights along a shortest path between any pair of vertices is lower bounded by some positive constant q. Then, the non-unity eigenvalues of D can not be close to 1 if q is large.

Our graph-theoretic results thus far have shown that strong connectivity in a Markov chain's graph implies relatively good estimator performance. It is reasonable to conjecture that weakly connected Markov chains may display poor estimator performance. However, the performance analysis of weakly-connected chains turns out somewhat more intricate. In particular, one can show that the non-unity eigenvalues of D are near 1 if and only if the Markov chain is weakly connected, and so indeed the bounds developed above become large for weakly-connected chains [29]. However, weakly connected chains can be constructed such that the estimator performance is good (even though the subdominant eigenvalue is near 1); we omit the details.

So far, we have used the spectral upper bounds to relate the total squared error performance measure to the Markov chain's graph in limiting cases, i.e. for strong or weak connectivities. While these strong- and weak- connectivity cases are often of particular interest in engineering applications, it is worth noting that the bounds on λ_2, as well as the exact expression for $tr(COV)$, can be related to a variety of other graph features using results from algebraic graph theory. We omit the details to save space: please see [31] for these details.

IV. Application: Average Reward Estimation

In this part, we briefly discuss an enhancement of the estimator and performance-measure computations for the MCSS estimation problem, to the case of Markov chains with rewards. Specifically, noting the interest in estimating expected rewards from data in numerous applications [17] (and also noting the similarity of this problem with the statistics-computation problem for Markov Chain Monte Carlo, or MCMC), we find it useful to develop and evaluate an estimator for expected rewards.

Again, we consider the same ergodic Markov chain pursued in the MCSS estimation problem. In addition, we assume that the Markov chain has a reward vector $\mathbf{r} \in \mathbf{R}^m$, where the ith entry in $\mathbf{r}$, r_i, represents an accumulated reward when the Markov chain enters state i. Obviously, the average reward per unit time can be represented as $r_{ss} = \mathbf{r}^T\pi$. For the reward case, suppose that we are interested in finding an estimator for the average reward per unit time (again for an unknown Markov chain), from measures of the actual rewards $r[k] = \mathbf{r}^T\mathbf{s}[k]$ at times $k = 1, \cdots, N$.

A classical unbiased estimator for the average reward assuming that the chain has reached steady-state is $\bar{r} = \frac{1}{N}\sum_{k=1}^{N} r[k] = \frac{1}{N}\sum_{k=1}^{N} \mathbf{r}^T\mathbf{s}[k] = \mathbf{r}^T\left(\frac{1}{N}\sum_{k=1}^{N}\mathbf{s}[k]\right) = \mathbf{r}^T\mathbf{p}$. The error variance of this estimate $\bar{r}$, which we call C_r, can straightforwardly be obtained and related to the error covariance COV for MCSS estimator:

$$\begin{aligned} C_r &= E[(\bar{r} - r_{ss})(\bar{r} - r_{ss})^T] \\ &= E[\mathbf{r}^T(\mathbf{p} - \pi)(\mathbf{p} - \pi)^T\mathbf{r}] \\ &= \mathbf{r}^T E[(\mathbf{p} - \pi)(\mathbf{p} - \pi)^T]\mathbf{r} \\ &= \mathbf{r}^T COV \mathbf{r}. \end{aligned} \tag{5}$$

Thus, we see that the performance computation for the reward case (Equation 5) is very similar to the one for the MCSS estimation problem. The series of algebraic and graphical characterizations of COV developed in Section III can also be applied for the reward case. However, an interesting difference here is that the reward vector $\mathbf{r}$ plays a role of scaling/averaging the entries in the error covariance matrix COV. If $\mathbf{r}$ can be chosen flexibly, its choice also has an influence on the estimation performance. For example, if $\mathbf{r}$ is chosen as vector $[1,0,\cdots,0]^T$ (i.e., reward is only considered when state 1 is hit), C_r equals the top-left entry in COV. Meanwhile, if $\mathbf{r}$ is (approximately) $\vec{1}$ (i.e., the reward is identical or nearly identical in each state), $\bar{\mathbf{r}}$ is (approximately) constant and hence C_r is (close to) zero.

Let us present one result that illustrates and quantifies the role of the reward vector $\mathbf{r}$ in determining the estimator's error variance. For this first result, we limit ourselves to the case that the Markov chain's graph is undirected (i.e., the state transition matrix D is symmetric). Then, in the following theorem, we identify the maximum possible value of the error variance of the estimate $\bar{r}$ (C_r), with respect to all possible normalized reward vectors. Here is the theorem:

Theorem 7: We consider an ergodic Markov chain just as in the MCSS estimation problem, and assume that the state transition matrix D is symmetric. Then, the maximum possible error variance of the unbiased estimator $\bar{r}$ for the average reward among all normalized reward vectors is $\max_{\mathbf{r}\, s.t.\, \|\mathbf{r}\|=1} C_r = \frac{2\gamma_2}{m}$, where $\gamma_2 = \frac{\lambda_2^{N+1} - N\lambda_2^2 + (N-1)\lambda_2}{N^2(1-\lambda_2)^2} + \frac{1}{2N}$. A reward vector that achieves the maximum is $\mathbf{r} = \mathbf{v}_2$, where $\mathbf{v}_2$ is any normalized eigenvector of D associated with the eigenvalue λ_2.

The above result bounds the variability of the expected reward estimate at each time-step N in terms of only the eigenvalue λ_2 of D (in the symmetric case), and shows that the reward vector that maximizes the estimate's variability is proportional to the eigenvector of D associated with λ_2. In fact, the eigenvalue λ_2 of D is intimately related to the graph's connectivity, and hence we can obtain graph-theoretic characterizations of reward estimation performance directly. Please see [31] for details.

References

[1] J. L. Gauvain, "Maximum a posteriori estimation for multivariate Gaussian mixture observations of Markov chains," *IEEE Transactions on Speech and Audio Processing 2*, pp. 291-298, 1994.

[2] S. T. Garren and R. L. Smith, "Estimating the second largest eigenvalue of a Markov transition matrix," *Bernoulli*, vol. 6, no. 2, pp. 215-242, April 2000.

[3] L. E. Baum and T. Petrie, "Statistical inference for probabilistic functions of finite state Markov chains," *The Annals of Mathematical Statistics*, Vol. 37, No. 6, pp. 1554-1563, Dec., 1966.

[4] O. Capp*é*, E. Moulines, T. Ryd*é*n, *Inference in Hidden Markov Models*, Springer, New York, 2005.

[5] A. Nasroallah, "Monte Carlo simulation of Markov chain steady-state distribution," *Extracta Mathematicae*, vol. 19 no. 2, pp. 279-288, 2004.

[6] H. Damerdji, "Strong consistency of the variance estimator in steady-state simulation output analysis," *Mathematics of Operations Research*, vol. 19, no. 2, pp. 494-512, May 1994.

[7] A. Doucet, A. Logothetis, V. Krishnamurthy, "Stochastic sampling algorithms for state estimation of jump Markov linear systems," *IEEE Transactions on Automatic Control*, Vol.45, No.2, 188-202, 2000.

[8] A. Doucet and C. Andrieu, "Iterative algorithm for state estimation for jump Markov linear systems," *IEEE Transaction on Signal Processing*, vol. 49, no. 6, June 2001.

[9] L. A. Liporace, "Maximum likelihood estimation for multivariate observations of Markov sources," *IEEE Transactions on Information Theory*, Vol. IT-28, pt. 1, pp. 729-734. Sept. 1982.

[10] M. Rudemo, "State estimation for partially observed Markov chains," *Mathematical Analysis and Applications*, vol. 44, issue 3, pp. 581-611, December 1973.

[11] T. W. Anderson and L. A. Goodman, "Statistical inference about Markov chains," *The Annuals of Mathematical Statistics*, vol. 28, no. 1, pp. 89-110, March 1957.

[12] A. D. Sokal, "Monte Carlo methods in statistical mechanics: foundations and new algorithms," *Cours de Troisieme Cycle de la Physique*, Suisse Romande, 1989.

[13] P. Doyle, "The Kemeny constant of a Markov chain," Arxiv preprint arXiv:0909.2636, September 2009.

[14] M. Jerrum and A. Sinclair, "Conductance and the rapid mixing property for Markov chains: the approximation of permanent resolved," *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, pp. 235-244, Chicago, IL, 2-4 May 1988.

[15] S. Boyd, P. Diaconis and L. Xiao, "Fastest mixing Markov chain on a graph," *SIAM Review*, vol. 46, no. 4, pp. 667-689, Dec., 2004.

[16] C. D. Meyer, "Sensitivity of the stationary distribution of a Markov chain," *SIAM Joournal on Matrix Analysis and Applications*, vol. 15, issue 3, pp. 715-728, July 1994.

[17] B. D. Theelen, J. P. M. Voeten, and Y. Pribadi, "Accuracy analysis of long-run average performance metrics," *Proceeding of the 2nd workshop on Embedded Systems*, pp. 261-269, Utrecht, STW, 2001.

[18] W. Landowski, M. Xue, S. Roy, "Comparison of Schedule-Based and Flow-Based Paradigms for Air Traffic Management," submitted to the *AIAA Aviation Technology, Integrations and Operations Conference*, Virginia Beach, VA, 20-22 Sep 2011.

[19] S. Mannor, D. Simester, P. Sun, J. N. Tsitsiklis, "Bias and variance approximations in value function estimates," *Management Science*, vol. 53, no. 2, pp. 308-322, 2007.

[20] R. M. Neal, "Probabilistic Inference Using Markov chain Monte Carlo Methods," technical report, University of Toronto, 1993.

[21] J. Besag, "Markov chain Monte Carlo for statistical inference," technical report, University of Washington, Centre for Statistical and the Social Sciences, working paper no. 9, 2000.

[22] M. Xue, E. Yeung, A. Rai, S. Roy, Y. Wan, S. Warnick, "Initial-condition estimation in network synchronization processes: algebraic and graphical characterizations of the estimator," submitted to the *AIAA Guidance, Navigation and Control Conference*, Portland, OR, 8-11 Aug 2011. .

[23] Y. Wan, S. Roy, "On inference of network time constants from impulse response data: graph-theoretic Cramer-Rao bounds," in *Proceedings of the 2009 IEEE Conference on Decision and Control*, Shanghai, China, pp. 4111-4116, Dec. 2009.

[24] G. L. Jones, "On the Markov chain central limit theorem," *Probability Surveys*, vol. 1, pp. 299-320, 2004.

[25] R. G. Gallager, *Discrete Stochastic Processes*, Kluwer Academic Publishers: Boston, 1996.

[26] H. Minc, "On the maximal eigenvector of a positive matrix," *SIAM Journal on Numerical Analysis*, vol. 7, no. 3, pp. 366-379, Oct. 2007.

[27] S. Cioaba and D. A. Gregory, "Principal eigenvectors of irregular graphs," *Electronic Journal of Linear Algebra*, vol. 16, pp. 366-379, Oct. 2007.

[28] C. Zenger, "A comparison of some bounds for the nontrivial eigenvalues of stochastic matrices," *Numerische Mathematik*, vol. 19, no. 3, June 1972.

[29] D. J. Hartfiel and Carl D. Meyer, "On the structure of stochastic matrices with a subdominant eigenvalue near 1," *Linear Algebra and its Applications*, vol. 272, issues 1-3, pp. 193-203, March 1998.

[30] M. Fiedler, "A Property of Eigenvectors of Non-negative Symmetric Matrices and its Application to Graph Theory", *Czech. Math. J.*, vol. 85, pp. 619-633, 1975.

[31] M. Xue and S. Roy, "Spectral and Graph-Theoretic Bounds on Steady-State-Probability Estimation Performance for an Ergodic Markov Chain," submitted to *Journal of the Franklin Institute*, available at `http://www.eecs.wsu.edu/~mxue/j2.pdf`.