

Robust measurement-based buffer overflow probability estimators for QoS provisioning and traffic anomaly prediction applicationm

Paschalidis, Ioannis Ch.

Spyridon Vassilaras, Ioannis Ch Paschalidis. "Robust measurement-based buffer overflow probability estimators for QoS provisioning and traffic anomaly prediction applicationm."

<https://hdl.handle.net/2144/18013>

Boston University

Robust measurement-based buffer overflow probability estimators for QoS provisioning and traffic anomaly prediction applications

Spyridon Vassilaras and Ioannis Ch. Paschalidis

Abstract

Suitable estimators for a class of Large Deviation approximations of rare event probabilities based on sample realizations of random processes have been proposed in our earlier work [PV01]. These estimators are expressed as non-linear multi-dimensional optimization problems of a special structure. In this paper, we develop an algorithm to solve these optimization problems very efficiently based on their characteristic structure. After discussing the nature of the objective function and constraint set and their peculiarities, we provide a formal proof that the developed algorithm is guaranteed to always converge. The existence of efficient and provably convergent algorithms for solving these problems is a prerequisite for using the proposed estimators in real time problems such as call admission control, adaptive modulation and coding with QoS constraints, and traffic anomaly detection in high data rate communication networks.

Keywords: Large Deviations, Markov-modulated processes, Estimation, Non-linear optimization, QoS, Traffic anomaly detection.

I. INTRODUCTION

Large Deviations (LD) theory (see [DZ98]) is an important analytical tool that has been applied to provide approximate solutions to a variety of queuing and stochastic decision problems for which analytical solutions are not practical (see for example [BPT98b], [BPT98a], [Pas99] and references therein). Most LD results provide an asymptotic approximation for the probability of a rare event that is formulated as a non-linear optimization problem. Solving this optimization problem is preferable to calculating the exact solution, which is often intractable or requires computing a definite integral in a high dimensional space. This computation typically requires numerical integration since the integrand is not given in closed form.

In prior work [PV01], we have proposed a class of LD estimators of small overflow probabilities at a queue fed by Markovian arrival processes when the statistics of the arrival process are estimated from a finite realization of the process. The proposed estimators are expressed as non-linear multi-dimensional optimization problems of a special structure. In [PV01], a heuristic, custom-made algorithm to solve this type of problems was presented. Unfortunately, this algorithm, albeit working very well in most cases, did not exhibit guaranteed convergence. In this paper, we improve

Spyridon Vassilaras. Athens Information Technology, 0.8km Markopoulo Ave., P.O. Box 68, Peania, 19002, Greece, e-mail: svas@ait.edu.gr

Ioannis Ch. Paschalidis, Dept. of Electrical & Computer Eng., Division of Systems Eng., and Center for Information & Systems Eng., Boston University, 8 St. Mary's St., Boston, MA 02215, e-mail: yannisp@bu.edu, Web: <http://ionia.bu.edu/>. Research partially supported by the NSF under grants EFRI-0735974, CNS-1239021, and IIS-1237022, by the ARO under grants W911NF-11-1-0227 and W911NF-12-1-0390, and by the ONR under grant N00014-10-1-0952.

on the previous algorithm to develop an algorithm that maintains the efficiency of the previous one while guaranteeing convergence. Proving that this algorithm always converges to the global optimum is the main result of this paper.

To better understand the context of this optimization problem, we start by a more detailed description of its perceived applications. In modern high-speed fixed communication networks congestion manifests itself as buffer overflows; the *Quality of Service (QoS)* faced by various connections can be quantified by the buffer overflow probability. To provide QoS guarantees the so called *effective bandwidth* admission control mechanism has been proposed [Hui88], [GH91], [Kel91], [GAN91], [Kel96], [Pas99]. Briefly, effective bandwidth is a number between the peak and average rate of a connection such that when connections are allocated their effective bandwidth in an appropriately dimensioned buffer, the buffer overflow probability stays below a given small level (say on the order of 10^{-6}). Real-time applications can tolerate such small frequencies of congestion phenomena. In wireless networks, which experience higher Bit Error Rate (BER) and time varying link quality, the targeted buffer overflow probability is much higher while the effective bandwidth of the source cannot be offered to a connection on a constant basis (since the available link capacity is a function of the link quality). In order to characterize the time-varying behavior of a wireless link the dual concept of *effective capacity* has been proposed by [WN03], [Wu03]. The effective capacity of a wireless link is the maximum constant source rate that can be offered to this link such that the buffer overflow probability stays below a given level. As the data rate and BER over a wireless link depend on the modulation and coding scheme applied, adaptive modulation and coding can achieve maximum effective capacity under a given BER constraint [LZG04], [LZG05], [LZG06], [TZ06], [TZ07b], [TZ07a]. More recently, a combined buffer overflow probability and BER minimization scheme has been proposed in [Vas10].

In make-to-stock manufacturing systems and supply chains, the objective is to control the inventory in order to avoid stockouts (see [BP01], [Gla97], [PL03]). In such systems demand is met from a finished goods inventory, and it is backlogged if inventory is not available. The stockout probability quantifies the QoS encountered by customers. It can be shown that this probability is equal to a buffer overflow probability in a corresponding make-to-order system [BP01], [PL03]. Thus, the problem of estimating the stockout probability can be transformed into one of estimating a buffer overflow probability.

Similar LD techniques have been applied in traffic anomaly detection where real time monitoring and analysis of the aggregate traffic in various points of a network can identify and classify unusual traffic fluctuations that can be caused by network intrusion, malicious attacks, network malfunction or irregular use of network resources [PS09], [PC10]. LD techniques for traffic anomaly detection can provide automatic alerts to previously unknown suspicious events, unlike signature based techniques that rely on known malicious content in data packets.

In each one of the above applications, we are interested in estimating buffer overflow or excessive queuing delay or stockout or traffic anomaly indicator probabilities that are very small. Moreover, arrival and service processes are typically autocorrelated (to model bursty traffic in communication networks and time varying capacity of wireless links, and to accommodate realistic demand scenarios and model failure-prone production facilities in make-to-stock manufacturing systems). As a result, obtaining exact analytic expressions is intractable; it is, therefore, natural to focus on approximate solutions such as the ones developed using LD techniques.

Most of the initial LD work assumed the detailed knowledge of models for the arrival and service processes. In practice however, such traffic models are not known a priori and have to be estimated from real observations. Consequently, one approach for estimating buffer overflow probabilities is to assume a certain traffic model, estimate the parameters of the model from real observations, and then calculate the overflow probability using large deviations results. Consider for example the case of a *deterministic Markov-modulated* source (i.e., a source which has a constant rate at every state of the underlying Markov chain) and let $g(\Xi)$ be the overflow probability when this source is fed in a certain buffer, where Ξ denotes the transition probability matrix of the underlying Markov-chain. (We assume that the rate at every state, and the characteristics of the buffer are given, thus, we do not explicitly denote the dependence of the overflow probability on these quantities). Let $\hat{\Xi}$ be an unbiased estimator of the transition probability matrix Ξ , that is, $\mathbf{E}[\hat{\Xi}] = \Xi$. Suppose that we use $g(\hat{\Xi})$ as an estimator for the overflow probability. An important observation is that due to the nonlinearity of $g(\cdot)$, $\mathbf{E}[g(\hat{\Xi})]$ is not necessarily equal to $g(\mathbf{E}[\hat{\Xi}]) = g(\Xi)$. That is, a *certainty equivalence* approach can lead to an erroneous estimate.

Subsequent work in the field has recognized this practical issue and developed LD estimators for measurement-based, model-free admission control [GT99a], [GT99b], [Duf99]. Our prior work [PV01] proposed new estimators of overflow probabilities in queues fed by Markov-modulated arrival processes. Intuitively, the proposed estimators suggested that we should quote a quantity that is larger than $g(\hat{\Xi})$ to guard against estimation errors. Hence, these estimators are “safe” even when based on relatively few observations, meaning that they do not lead to substantial underestimation of the overflow probability that can compromise QoS. Still, they are *consistent* in the sense that they converge to $g(\Xi)$ with probability 1 (w.p.1) as the number of observations tends to infinity.

The development of efficient non-linear optimization algorithms to compute the proposed estimators is of paramount importance. Given that these estimators are meant to be used for real time applications in fixed and wireless telecommunication networks, they should be computed in very short time. This means that the non-linear optimization algorithm must guarantee convergence to the minimum in a small number of steps. The high dimensionality of the decision variable space and the nature of the objective function (which is computationally expensive to calculate) makes standard optimization algorithms not efficient enough. The algorithm proposed in this paper takes advantage of the special structure of the objective function and the constraint set to achieve fast and guaranteed convergence.

This paper is organized as follows. We start in Section II with a summary of the key results and notation from [PV01] including the optimization problems that we want to solve. In Section III we analyze the structure of the objective functions and constraint sets of these optimization problems and discuss their key properties such as differentiability and convexity. In Section IV we develop two algorithms for solving the optimization problems and show that the second one (which is an improved version of the first one) exhibits guaranteed convergence. Finally, conclusions are in Section V.

On a notational remark, in this paper, matrices are denoted by bold uppercase and vectors by bold lowercase characters. Moreover, all vectors will be assumed to be column vectors, unless otherwise specified.

II. ESTIMATES OF THE OVERFLOW PROBABILITY BASED ON AN ARRIVAL PROCESS REALIZATION

In this Section we summarize the key results and notation from [PV01] in order to establish our notation and understand the significance of the optimization problems at hand. Consider a single class G/G/1 queue with a Markov-modulated arrival process. We will be using a discrete-time model, where time is divided into time slots. We let A_k , $k \in \mathbb{Z}$, denote the aggregate number of “customers” that enter the queue at time k . The queue has an infinite buffer and is serviced according to a general service process which can clear up to B_k customers during the time interval $[k, k + 1]$. We assume that the stochastic processes $\{A_k; k \in \mathbb{Z}\}$ and $\{B_k; k \in \mathbb{Z}\}$ are stationary, possibly autocorrelated and mutually independent processes. For such a discrete time stochastic process $\{X_k; k \in \mathbb{Z}\}$ let us define the partial sum process $S_n = \sum_{k=1}^n X_k$. We assume that the partial sum processes of both the arrival and service processes $\{A_k; k \in \mathbb{Z}\}$ and $\{B_k; k \in \mathbb{Z}\}$ satisfy some technical conditions (see Assumption A in [PV01]) among which is the property that the limit

$$\Lambda_X(\theta) \triangleq \lim_{n \rightarrow \infty} \Lambda_n^X(\theta) = \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{E}[e^{\theta S_n}],$$

exists for all θ , where $\pm\infty$ are allowed both as elements of the sequence $\Lambda_n^X(\theta)$ and as limit points. Note that the set of processes satisfying these conditions is large enough to include renewal, Markov-modulated, and stationary processes with mild mixing conditions. Such processes can model “burstiness” and are commonly used in modeling the input traffic to communication networks and the link capacity fluctuations due to wireless channel fading. They have also been used in modeling demand and the production process in manufacturing systems [BP01], [PL03].

Let us define:

$$\Lambda_X^*(a) \triangleq \sup_{\theta} (\theta a - \Lambda_X(\theta)), \quad (1)$$

which is the Legendre transform of $\Lambda_X(\cdot)$. The function $\Lambda_X^*(\cdot)$ is convex and lower semicontinuous (see [DZ98]). In Large Deviations parlance, $\Lambda_X(\cdot)$ and $\Lambda_X^*(\cdot)$ are called the limiting log-moment generating function and the large deviations rate function, respectively, of the process $\{X_k; k \in \mathbb{Z}\}$.

We denote by L_k the queue length at time k (without counting arrivals at time k). We assume that the server uses a work-conserving policy (i.e., the server never stays idle when there is work in the system) and that

$$\mathbf{E}[A_1] < \mathbf{E}[B_1], \quad (2)$$

which by stationarity carries over to all k . We also assume that the queue length process $\{L_k; k \in \mathbb{Z}\}$ is stationary. To simplify the analysis we consider a discrete-time “fluid” model, meaning that we will be treating A_k , L_k , and B_k as nonnegative real numbers (the amount of fluid entering the buffer, in queue, and the service capacity, respectively).

A Large Deviations Principle (LDP) for the queue length process has been established in [GW94], [BPT98b] and is given in the next proposition. In preparation for the result, consider a convex function $g(u)$ with the property $g(0) = 0$. We define the *largest root* of $g(u)$ to be the solution of the optimization problem $\sup_{u: g(u) < 0} u$. If $g(\cdot)$ has negative derivative at $u = 0$, there are two cases: either $g(\cdot)$ has a single positive root or it stays below the horizontal axis for all $u > 0$. In the latter case, we will say that $g(\cdot)$ has a root at $u = \infty$.

Proposition II.1 *The steady-state queue length process L_k satisfies*

$$\lim_{U \rightarrow \infty} \frac{1}{U} \log \mathbf{P}[L_k \geq U] = -\theta^*, \quad (3)$$

where $\theta^* > 0$ is the largest root of the equation

$$\Lambda_A(\theta) + \Lambda_B(-\theta) = 0. \quad (4)$$

More intuitively, for large enough U we have

$$\mathbf{P}[L_k \geq U] \sim e^{-U\theta^*}.$$

This expression can be used to estimate the overflow probability in a queue with a finite buffer of size U . Kelly [Kel96] establishes that the latter probability has the same asymptotic decay rate (same exponent) with $\mathbf{P}[L_k \geq U]$.

A. Markov-modulated arrivals

Calculating the limiting log-moment generating function can be made easier than computing infinite sums when dealing with certain special stochastic processes. The most common such process is the Markov-modulated process, which is defined as follows: Consider an irreducible Markov chain with M states $1, 2, \dots, M$ and transition probability matrix $\Xi = \{p(i, j)\}_{i, j=1}^M$. We will be using the notation $\mathbf{p}^T = (\mathbf{p}_1^T, \dots, \mathbf{p}_M^T)$, where \mathbf{p}_i^T is the i -th row of Ξ and \mathbf{A}^T denotes the transpose of \mathbf{A} . The Markov chain makes one transition per time slot; let Y_k be the state at time k . The number of arrivals at time k is a random function of the state, i.e., it is drawn according to a p.d.f. $f_{Y_k}(\cdot)$ associated with that state. Let us denote by $\eta_i(\theta)$ the moment generating function of $f_i(\cdot)$. Note that for a deterministic amount of arrivals r_i per time slot at state i , the moment generating function is $\eta_i(\theta) = e^{\theta r_i}$.

In [DZ98, Sec. 3.1.1] it is established that the limiting log-moment generating function of the arrival process $\{A_k; k \in \mathbb{Z}\}$, is given by

$$\Lambda_A(\theta, \mathbf{p}) = \log \rho(\mathbf{\Pi}_{\theta, \mathbf{p}}^A), \quad (5)$$

where $\rho(\mathbf{\Pi}_{\theta, \mathbf{p}}^A)$ denotes the Perron-Frobenius eigenvalue of the $M \times M$ matrix:

$$\mathbf{\Pi}_{\theta, \mathbf{p}}^A = \{p(i, j)\eta_j(\theta)\}_{i, j=1}^M. \quad (6)$$

(In this Markov-modulated case we are using notation that explicitly denotes the dependence of $\Lambda_A(\theta, \mathbf{p})$ and $\mathbf{\Pi}_{\theta, \mathbf{p}}^A$ on the transition probabilities \mathbf{p} .) Notice that because the quantities $\eta_j(\theta)$ are always positive the irreducibility of Ξ implies that $\mathbf{\Pi}_{\theta, \mathbf{p}}^A$ is irreducible.

B. Estimating the Overflow Probability

Consider now the more realistic case where the transition probabilities $p(i, j)$ are not known in advance, but need to be estimated by observing the arrival process. In particular, we will be assuming that we have perfect knowledge of the service process $\{B_k; k \in \mathbb{Z}\}$, and that we can

observe the states of the Markov chain associated with the arrival process. That is, we do know M and the probability density functions $f_i(\cdot)$ (the supports of whom do not overlap), but the transition probability matrix Ξ is unknown. Suppose that we observe a sequence $\mathbf{Y} = Y_1, Y_2, \dots, Y_n$ of states that the unknown Markov chain visits with the initial state being $Y_0 = \sigma$. Consider the empirical measures:

$$q_n^{\mathbf{Y}}(\mathbf{y}) = \frac{1}{n} \sum_{k=1}^n \mathbf{1}_{\mathbf{y}}(Y_{k-1}Y_k),$$

where $\mathbf{y} \in \mathcal{A}^2 \triangleq \{1, \dots, M\} \times \{1, \dots, M\}$. Note that when $\mathbf{y} = (i, j) \in \mathcal{A}^2$ the empirical measure $q_n^{\mathbf{Y}}(\mathbf{y})$ denotes the fraction of times that the Markov chain makes transitions from i to j in the sequence \mathbf{Y} . Let now $\mathcal{A}_{\mathbf{p}}^2 \triangleq \{(i, j) \in \mathcal{A}^2 \mid p(i, j) > 0\}$ denote the set of pairs of states that can appear in the sequence Y_1, Y_2, \dots, Y_n and denote by $M_1(\mathcal{A}_{\mathbf{p}}^2)$ the standard $|\mathcal{A}_{\mathbf{p}}^2|$ -dimensional probability simplex, where $|\mathcal{A}_{\mathbf{p}}^2|$ denotes the cardinality of $\mathcal{A}_{\mathbf{p}}^2$. Note that the vector of $q_n^{\mathbf{Y}}(\mathbf{y})$'s denoted by $\mathbf{q}_n^{\mathbf{Y}} = (q_n^{\mathbf{Y}}(\mathbf{y}); \mathbf{y} \in \mathcal{A}_{\mathbf{p}}^2)$ is an element of $M_1(\mathcal{A}_{\mathbf{p}}^2)$. For any $\mathbf{q} \in M_1(\mathcal{A}_{\mathbf{p}}^2)$, let

$$q_1(i) \triangleq \sum_{j=1}^M q(i, j) \quad \text{and} \quad q_2(i) = \sum_{j=1}^M q(j, i) \quad (7)$$

be its marginals. Whenever $q_1(i) > 0$, let $q_f(j \mid i) \triangleq q(i, j)/q_1(i)$. We will be using the notation $\mathbf{q}_f = (q_f(1 \mid 1), \dots, q_f(M \mid 1), q_f(1 \mid 2), \dots, q_f(M \mid 2), \dots, q_f(1 \mid M), \dots, q_f(M \mid M))$. (To avoid overburdening the notation, we will suppress the dependence of the estimators on the sequence \mathbf{Y} . We will also often omit the subscript n and simply write \mathbf{q} , \mathbf{q}_f and $q_1(i)$.)

Note that \mathbf{p} (the vector of the actual but unknown transition probabilities) is an element of $(M_1(\mathcal{A}))^M$ (i.e., the M -times cartesian product of $M_1(\mathcal{A})$). As the transition probabilities \mathbf{p} are not known, we assume a prior pdf $\phi_{\mathbf{p}^-} \in (M_1(M_1(\mathcal{A})))^M$, which assigns probability mass only to \mathbf{p} 's corresponding to irreducible Markov chains. Let $\underline{\mathbf{p}}^-$ denote the support of $\phi_{\mathbf{p}^-}$.

Let us define:

$$I_3(\mathbf{p}) = \begin{cases} \sum_{i=1}^M q_1(i) H(q_f(\cdot \mid i) \mid p(i, \cdot)), & \text{if } \mathbf{p} \in \underline{\mathbf{p}}^-, \\ \infty, & \text{otherwise.} \end{cases} \quad (8)$$

where $H(q_f(\cdot \mid i) \mid p(i, \cdot))$ is the relative entropy defined as

$$H(q_f(\cdot \mid i) \mid p(i, \cdot)) = \sum_{j=1}^M q_f(j \mid i) \log \frac{q_f(j \mid i)}{p(i, j)}.$$

Now given the empirical measure \mathbf{q}_n , a maximum likelihood estimator of the transition probabilities is given by

$$\hat{p}_n(i, j) = q_{nf}(j \mid i) = \frac{q_n(i, j)}{q_{n1}(i)}, \quad i, j = 1, \dots, M, \quad (9)$$

where the extra n in the subscripts explicitly denotes the dependence on the length n of the sample sequence. Let $\hat{\mathbf{p}}_n$ denote the vector of these estimates. We can now construct a matrix $\mathbf{\Pi}_{\theta, \hat{\mathbf{p}}_n}^A$ with

elements

$$\pi_{\theta, \hat{\mathbf{p}}_n}^A(i, j) = \hat{p}_n(i, j) \eta_j(\theta) \quad i, j = 1, \dots, M,$$

and obtain an estimate, $\Lambda_A(\theta, \hat{\mathbf{p}}_n)$, of the limiting log-moment generating function for the arrival process by computing the Perron-Frobenius eigenvalue of $\mathbf{\Pi}_{\theta, \hat{\mathbf{p}}_n}^A$. The following 3 estimators of the overflow probability $\mathbf{P}[L_i \geq U]$ have been described in [PV01] (the reader is referred to [PV01] for more detailed explanation of these estimators).

The traditional ‘‘certainty equivalent’’ estimator:

$$\mathcal{P}_n^I \triangleq e^{-U\theta^*(\hat{\mathbf{p}}_n)}, \quad (10)$$

where $\theta^*(\hat{\mathbf{p}}_n)$ is the largest root of the equation $\Lambda_A(\theta, \hat{\mathbf{p}}_n) + \Lambda_B(-\theta) = 0$.

And two improved estimators:

$$\mathcal{P}_n^{II} \triangleq \exp \left\{ -n \inf_{\mathbf{p} \in (M_1(\mathcal{A}))^M} [s\theta^*(\mathbf{p}) + I_3(\mathbf{p})] \right\}. \quad (11)$$

and

$$\mathcal{P}_n^{IV}(\mu) \triangleq \mathcal{P}_n^{II} + \mu \sqrt{\mathcal{P}_n^{III}(2) - (\mathcal{P}_n^{II})^2}, \quad (12)$$

where μ is some scalar and

$$\mathcal{P}_n^{III}(\ell) \triangleq \exp \left\{ -n \inf_{\mathbf{p} \in (M_1(\mathcal{A}))^M} [\ell s\theta^*(\mathbf{p}) + I_3(\mathbf{p})] \right\} \quad \ell = 1, 2, \dots \quad (13)$$

$\mathcal{P}_n^{IV}(\mu)$ can be interpreted as expectation plus μ times standard deviation. The idea is that by selecting μ large enough $\mathcal{P}_n^{IV}(\mu)$ can be adequately safe, i.e., the likelihood that $\mathcal{P}_n^{IV}(\mu)$ underestimates the true probability of loss can be made adequately small.

III. COMPUTING THE ESTIMATORS

Computing the estimators in (11) and (13) requires solving nonlinear optimization problems. In this section we examine the structure of these problems and devise efficient algorithms for their solution.

The optimization problems in (11) and (13) have the following form

$$\begin{aligned} & \text{minimize} && \ell s\theta^*(\mathbf{p}) + I_3(\mathbf{p}) \\ & \text{s.t.} && p(i, j) \geq 0, \quad i, j = 1, \dots, M, \\ & && \sum_{j=1}^M p(i, j) = 1, \quad i = 1, \dots, M, \end{aligned} \quad (14)$$

where ℓ is some positive scalar and $I_3(\mathbf{p})$ is given by (cf. Eq. (8))

$$I_3(\mathbf{p}) = \begin{cases} \sum_{i=1}^M q_1(i) \sum_{j=1}^M q_f(j | i) \log \frac{q_f(j|i)}{p(i,j)}, & \text{if } \mathbf{p} \in \underline{\mathbf{p}}^-, \\ \infty, & \text{otherwise.} \end{cases} \quad (15)$$

We are interested in solving relatively large instances of (14); typically M can be in the range of 5-100 which brings the number of decision variables in the range of 25-10000. Since one major application for these estimators is Call Admission Control in telecommunication networks that provide QoS guarantees, the optimization problems at hand should be solved in real time. As a result, computational efficiency is critical and special purpose algorithms that exploit the special structure of these optimization problems are of interest.

On the structure of the objective function, a first observation is that $I_3(\mathbf{p})$ is a convex function if $\underline{\mathbf{p}}^-$ is a convex set. This can be easily established directly from the definition of convexity. Furthermore, $I_3(\mathbf{p})$ is strictly convex in $\underline{\mathbf{p}}^-$ (where it is finite). In Section II we have assumed that the prior assigns probability mass only to \mathbf{p} 's corresponding to irreducible Markov chains. Let \mathcal{I} be the set of \mathbf{p} 's corresponding to irreducible Markov chains. The set \mathcal{I} is a convex set (for a proof see [PV01]).

Based on the convexity of \mathcal{I} and assuming that the support of the prior $\underline{\mathbf{p}}^-$ is a convex subset of \mathcal{I} , $I_3(\mathbf{p})$ is a convex function. Henceforth, and in the absence of more information on the true transition probabilities of the Markov chain we wish to estimate, we will be making the following assumption.

Assumption A

The support of the prior, $\underline{\mathbf{p}}^-$, is the set \mathcal{I} of transition probability vectors \mathbf{p} corresponding to irreducible Markov chains.

Assumption A implies that n can always be taken large enough so that at least one exit transition for each state has been observed and therefore \mathbf{q}_f is the transition probability vector of an irreducible Markov chain. Hence we assume that if the observed sequence \mathbf{Y} results to $\mathbf{q}_f \notin \mathcal{I}$ then we prolong the observation period by a few samples to achieve $\mathbf{q}_f \in \mathcal{I}$.

Consider now the following optimization problem:

$$\begin{aligned} & \text{minimize} && \ell s \theta^*(\mathbf{p}) + \hat{I}_3(\mathbf{p}) \\ & \text{s.t.} && p(i, j) \geq 0, \quad i, j = 1, \dots, M, \\ & && \sum_{j=1}^M p(i, j) = 1, \quad i = 1, \dots, M, \end{aligned} \tag{16}$$

where

$$\hat{I}_3(\mathbf{p}) = \sum_{i=1}^M q_1(i) \sum_{j=1}^M q_f(j | i) \log \frac{q_f(j | i)}{p(i, j)}. \tag{17}$$

The following property of the optimal solution of the above optimization problem is established in [PV01]:

Lemma III.1 *Let \mathbf{p}^* be an optimal solution of the optimization problem in (16). Then \mathbf{p}^* is the transition probability vector of an irreducible Markov chain.*

The result of this lemma suggests that under Assumption A the optimization problem in (16) is equivalent to the problem in (14). Hence, we will focus on solving (16).

A very important issue for any nonlinear programming problem is the form of the objective function. A convex objective function, especially under polyhedral constraints (as in (16)), can

lead to more efficient algorithms. Let

$$\bar{\theta}(\mathbf{p}) \triangleq \ell s \theta^*(\mathbf{p}) + \hat{I}_3(\mathbf{p})$$

denote the objective function in (16). Notice that it is the weighted sum of a strictly convex function, $\hat{I}_3(\mathbf{p})$, and $\theta^*(\mathbf{p})$, which is not necessarily convex. To see this, we consider an example where $\ell = 1$ and the arrival process is a two-state Markov-modulated process with deterministic number of arrivals per state $(r_1, r_2) = (0.042, 0.077)$. This arrival process is fed into a buffer which is served at fixed service rate $c = 0.058$. The parameter s was set equal to 0.002. In Figure 1 (a) we plot $\theta^*(\mathbf{p})$ and in (b) $\bar{\theta}(\mathbf{p})$ versus the two (arbitrarily chosen as) independent decision variables $p(1, 1)$ and $p(2, 2)$. It can be seen that although $\theta^*(\mathbf{p})$ is not convex (it is convex only along some directions), $\bar{\theta}(\mathbf{p})$ is convex. Consequently, $\bar{\theta}(\mathbf{p})$ is not convex in general. Nevertheless, since $\hat{I}_3(\mathbf{p})$ is strictly convex, it can be seen that $\bar{\theta}(\mathbf{p})$ will be convex for small enough values of s . Recalling that $s = U/n$ and assuming that the buffer size U is given, we will be dealing with a convex objective function if we can afford a large number n of measurements.

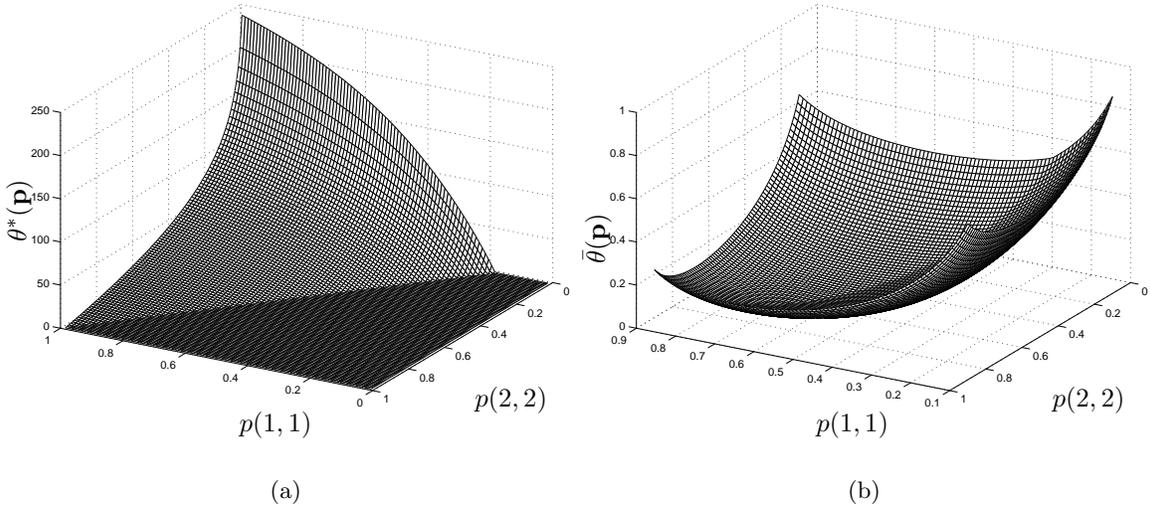


Fig. 1

PLOTS OF $\theta^*(\mathbf{p})$ (A) AND $\bar{\theta}(\mathbf{p})$ (B) VERSUS $p(1, 1)$ AND $p(2, 2)$ WHEN THE ARRIVAL PROCESS IS A MARKOV CHAIN WITH 2 STATES.

Another important observation, which will affect the convergence of the optimization algorithms presented in the next section, is that $\theta^*(\mathbf{p})$ might be equal to 0 (which translates into $\mathcal{P}_n^I = 1$) for some values of \mathbf{p} . This happens when the mean arrival rate is larger or equal to the mean service rate $\mathbf{E}[B_1]$ (which is equal to c in the most common case of a deterministic service process). Therefore, if we denote by $p_1(i)$ the steady state probability for state i corresponding to the transition probability vector \mathbf{p} and by $\mathbf{E}[A_1 | i]$ the conditional mean number of arrivals given that we are at state i , the condition for $\theta^*(\mathbf{p}) = 0$ can be written as:

$$\sum_{i=1}^M p_1(i) \mathbf{E}[A_1 | i] \geq \mathbf{E}[B_1] \quad (18)$$

Let us now define the following subsets of \mathcal{I} :

- $\mathcal{I}_1 \triangleq \{\mathbf{p} \in \mathcal{I} \mid \sum_{i=1}^M p_1(i) \mathbf{E}[A_1 | i] < \mathbf{E}[B_1]\}$
- $\mathcal{I}_2 \triangleq \{\mathbf{p} \in \mathcal{I} \mid \sum_{i=1}^M p_1(i) \mathbf{E}[A_1 | i] > \mathbf{E}[B_1]\}$
- $\mathcal{I}_3 \triangleq \{\mathbf{p} \in \mathcal{I} \mid \sum_{i=1}^M p_1(i) \mathbf{E}[A_1 | i] = \mathbf{E}[B_1]\}$
- $\mathcal{I}_4 \triangleq \{\mathbf{p} \in \mathcal{I} \mid \sum_{i=1}^M p_1(i) \mathbf{E}[A_1 | i] \leq \mathbf{E}[B_1]\} = \mathcal{I}_1 \cup \mathcal{I}_3$
- $\mathcal{I}_5 \triangleq \{\mathbf{p} \in \mathcal{I} \mid \sum_{i=1}^M p_1(i) \mathbf{E}[A_1 | i] \geq \mathbf{E}[B_1]\} = \mathcal{I}_2 \cup \mathcal{I}_3$

Obviously, \mathcal{I}_1 , \mathcal{I}_2 and \mathcal{I}_3 form a partition of \mathcal{I} into disjoint sets. Furthermore, the dimension of \mathcal{I}_3 is always equal to the dimension of \mathcal{I} minus one. \mathcal{I}_3 forms a boundary between \mathcal{I}_1 and \mathcal{I}_2 . Note that if $\mathbf{q}_f \in \mathcal{I}_5$, then $\mathcal{P}_n^I = 1$, which implies that $\mathcal{P}_n^{II} = 1$ (since $\mathcal{P}_n^{II} \geq \mathcal{P}_n^I$), which is a case of hardly any interest. In the sequel, we will assume that $\mathbf{q}_f \in \mathcal{I}_1$. The next lemma reveals an important property of \mathcal{I}_1 with respect to the optimization problem in (16):

Lemma III.2 *If $\mathbf{q}_f \in \mathcal{I}_1$ then $\mathbf{p}^* \in \mathcal{I}_4$, where \mathbf{p}^* is the minimizing \mathbf{p} in (16).*

Proof : Assume $\mathbf{p}^* \notin \mathcal{I}_4$, which is equivalent to $\mathbf{p}^* \in \mathcal{I}_2$. Now, denote by l_1 the line segment that connects the two points \mathbf{q}_f and \mathbf{p}^* . For every point $\mathbf{p} \in l_1$ different than \mathbf{p}^* , $\hat{I}_3(\mathbf{p}) < \hat{I}_3(\mathbf{p}^*)$ due to the strict convexity of $\hat{I}_3(\cdot)$ whose global minimum is at \mathbf{q}_f . Also, for every point $\mathbf{p} \in l_1 \cap \mathcal{I}_5$, $\theta^*(\mathbf{p}) = 0$. Furthermore, there is at least one point $\mathbf{p} \in l_1 \cap \mathcal{I}_5 \setminus \{\mathbf{p}^*\}$, since $\mathbf{p}^* \in \mathcal{I}_2$. For this point we have that $\bar{\theta}(\mathbf{p}) = ks\theta^*(\mathbf{p}) + \hat{I}_3(\mathbf{p}) = \hat{I}_3(\mathbf{p}) < \hat{I}_3(\mathbf{p}^*) = \bar{\theta}(\mathbf{p}^*)$. But this contradicts our initial assumption that \mathbf{p}^* is the minimizer in (16). Thus, $\mathbf{p}^* \in \mathcal{I}_4$. ■

Figure 1 seems to suggest that \mathcal{I}_3 is a straight line, which makes \mathcal{I}_1 , \mathcal{I}_2 , \mathcal{I}_4 and \mathcal{I}_5 convex sets. Indeed, for the case where $M = 2$, $p_1(1)$ and $p_1(2)$ can be expressed as a function of $p(1, 1)$ and $p(2, 2)$ as follows:

$$p_1(1) = \frac{1 - p(2, 2)}{2 - p(1, 1) - p(2, 2)}$$

and

$$p_1(2) = \frac{1 - p(1, 1)}{2 - p(1, 1) - p(2, 2)}$$

Therefore the condition for $\mathbf{p} \in \mathcal{I}_3$ becomes:

$$\sum_{i=1}^2 \frac{1 - p(i, i)}{2 - p(1, 1) - p(2, 2)} \mathbf{E}[A_1 | 3 - i] = \mathbf{E}[B_1]$$

which after some routine algebraic manipulations can be written as:

$$(\mathbf{E}[B_1] - \mathbf{E}[A_1 | 2])p(1, 1) + (\mathbf{E}[B_1] - \mathbf{E}[A_1 | 1])p(2, 2) = 2\mathbf{E}[B_1] - \mathbf{E}[A_1 | 1] - \mathbf{E}[A_1 | 2]$$

The above equation is clearly a line equation on the plane and that is in agreement with the straight line border between \mathcal{I}_1 and \mathcal{I}_2 shown in Figure 1. However, for $M > 2$, it can be easily shown by example that both sets \mathcal{I}_1 and \mathcal{I}_2 are not necessarily convex. For a case with $M=3$, where \mathcal{I}_1 is not convex, take the following two elements of \mathcal{I} given in matrix form

$$\mathbf{\Xi}_1 = \begin{bmatrix} 0.2 & 0.6 & 0.2 \\ 0.6 & 0.2 & 0.2 \\ 0.2 & 0.4 & 0.4 \end{bmatrix} \quad \text{and} \quad \mathbf{\Xi}_2 = \begin{bmatrix} 0.7 & 0.1 & 0.2 \\ 0.2 & 0.1 & 0.7 \\ 0.3 & 0.6 & 0.1 \end{bmatrix},$$

and assume

$$[\mathbf{E}[A_1 | 1], [\mathbf{E}[A_1 | 2], [\mathbf{E}[A_1 | 3]] = [0.2, 1.0, 2.5].$$

Figure 2(a) shows the mean arrival rate of a Markov-modulated process driven by an underlying Markov chain with transition probability matrix along the line segment connecting $\mathbf{\Xi}_1$ to $\mathbf{\Xi}_2$. More specifically, it plots a graph of

$$\mathbf{E}[A_1] = \sum_{i=1}^3 p_1(i) \mathbf{E}[A_1 | i],$$

where $p_1(i)$ is the steady state probability of state i of a Markov chain with transition probability matrix $\alpha\mathbf{\Xi}_1 + (1 - \alpha)\mathbf{\Xi}_2$, as a function of the parameter α . It can be seen that for an appropriate choice of $\mathbf{E}[B_1]$ (e.g., the one depicted by the horizontal dashed line) $\mathbf{\Xi}_1$ and $\mathbf{\Xi}_2$ belong to \mathcal{I}_1 but there are points on the line segment connecting $\mathbf{\Xi}_1$ and $\mathbf{\Xi}_2$ that don't. Hence, in this case \mathcal{I}_1 is not convex.

Similarly, for a case where \mathcal{I}_2 is not convex, consider the transition probability matrices

$$\mathbf{\Xi}_1 = \begin{bmatrix} 0.2 & 0.6 & 0.2 \\ 0.2 & 0.2 & 0.6 \\ 0.3 & 0.4 & 0.3 \end{bmatrix} \quad \text{and} \quad \mathbf{\Xi}_2 = \begin{bmatrix} 0.2 & 0.1 & 0.7 \\ 0.7 & 0.1 & 0.2 \\ 0.4 & 0.2 & 0.4 \end{bmatrix},$$

and keep the same conditional mean arrival rates as in the previous example. A plot of the mean arrival rate of the MMP along the line segment connecting $\mathbf{\Xi}_1$ and $\mathbf{\Xi}_2$ is shown in Figure 2(b). Clearly, for an appropriate choice of $\mathbf{E}[B_1]$ (e.g., the one depicted by the horizontal dashed line) $\mathbf{\Xi}_1$ and $\mathbf{\Xi}_2$ belong to \mathcal{I}_2 but there are points on the line segment connecting $\mathbf{\Xi}_1$ and $\mathbf{\Xi}_2$ that don't. Hence, in this case \mathcal{I}_2 is not convex.

In the next section we will see that $\bar{\theta}(\mathbf{p})$, the objective function in problem (16), is not differentiable at \mathcal{I}_3 . If \mathcal{I}_1 was always convex, we could take advantage of Lemma (III.2) and the continuity of $\bar{\theta}(\mathbf{p})$ to constraint the optimization into \mathcal{I}_1 . Unfortunately this is not the case.

IV. ALGORITHMS FOR SOLVING THE OPTIMIZATION PROBLEM

To solve the problem in (16) we have developed a heuristic algorithm which performs very well in practice, typically giving first decimal digit approximations to the optimal value by the third

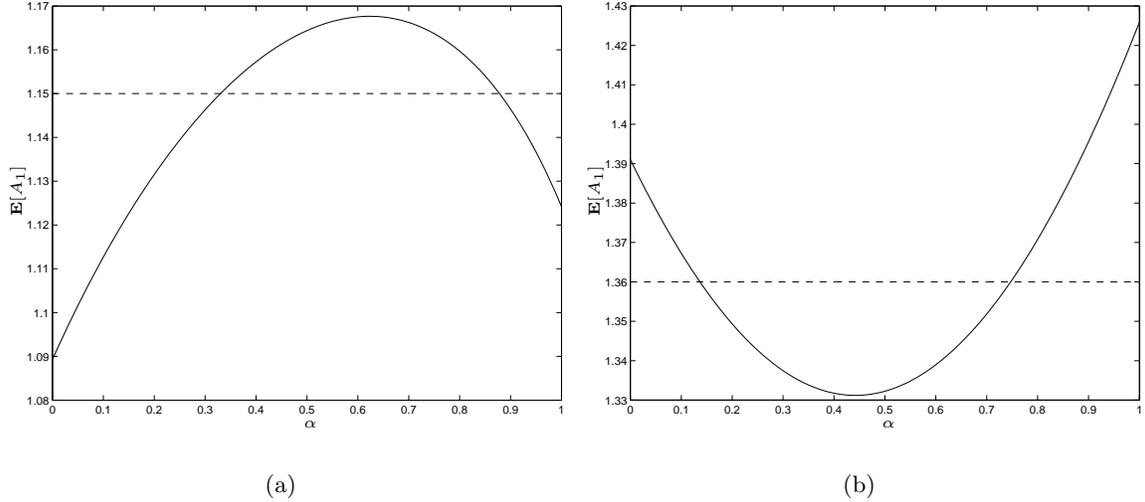


Fig. 2

PLOTS OF $\mathbf{E}[A_1]$ ALONG THE LINE SEGMENTS CONNECTING TWO POINTS IN \mathcal{I} FOR EXAMPLES WHERE (A) \mathcal{I}_1 IS NOT CONVEX AND (B) \mathcal{I}_2 IS NOT CONVEX.

iteration and third decimal digit by the fourth.

To describe the heuristic, note that the constraint set \mathcal{I} is the Cartesian product of simplices, and that $q_f(j | i) > 0$ implies $p(i, j) > 0$ for all $i, j = 1, \dots, M$. This feasible set implies the following optimality conditions ([Ber95, pg. 178-179])

$$\frac{\partial \bar{\theta}(\mathbf{p})}{\partial p(i, j)} = \frac{\partial \bar{\theta}(\mathbf{p})}{\partial p(i, j')} \quad \forall i, j, j' = 1, \dots, M, \text{ satisfying } q_f(j | i), q_f(j' | i) > 0. \quad (19)$$

The heuristic algorithm iterates as follows:

Algorithm A:

1. Initialize with $\mathbf{p}^{(0)} = \mathbf{q}_f$ and $m=0$.
2. Form an approximation of $\bar{\theta}(\mathbf{p})$

$$\bar{\theta}_{\text{app}}^{(m)}(\mathbf{p}) \triangleq \ell_s \theta^*(\mathbf{p}^{(m)}) + \ell_s \nabla \theta^*(\mathbf{p}^{(m)})' (\mathbf{p} - \mathbf{p}^{(m)}) + \hat{I}_3(\mathbf{p}). \quad (20)$$

Minimize the expression in Eq. (20) subject to the constraints of problem (16) to obtain an optimal solution $\mathbf{p}_{\mathbf{p}^{(m)}}^*$.

3. Set $\mathbf{p}^{(m+1)} := \mathbf{p}_{\mathbf{p}^{(m)}}^*$ and then make $m=m+1$.
4. If $\bar{\theta}(\mathbf{p}^{(m-1)}) - \bar{\theta}(\mathbf{p}^{(m)}) < \epsilon$, where ϵ is the desired accuracy, terminate with optimal solution $\mathbf{p}^{(m)}$. Otherwise, return to Step 2.

The intuition behind this algorithm is that as we get closer to the real minimum, \mathbf{p}^* , of problem

(16) the approximations $\bar{\theta}_{\text{app}}^{(m)}(\mathbf{p})$ of $\bar{\theta}(\mathbf{p})$ improve and yield $\mathbf{p}^{(m+1)}$ that are closer to \mathbf{p}^* . Because at each step we solve the approximate problem exactly, the algorithm typically needs much fewer iterations than a standard gradient-based algorithm.

The minimization of the expression in Eq. (20) can be performed by solving the system of optimality conditions:

$$\frac{\partial \bar{\theta}_{\text{app}}^{(m)}(\mathbf{p})}{\partial p(i, j)} = \frac{\partial \bar{\theta}_{\text{app}}^{(m)}(\mathbf{p})}{\partial p(i, j')} \quad \forall i, j, j' = 1, \dots, M \text{ satisfying } q_f(j | i), q_f(j' | i) > 0 \quad (21)$$

which, by introducing the variables w_i , can be written equivalently as:

$$\frac{\partial \bar{\theta}_{\text{app}}^{(m)}(\mathbf{p})}{\partial p(i, j)} = w_i \quad \forall i, j = 1, \dots, M \text{ satisfying } q_f(j | i) > 0 \quad (22)$$

subject to the conditions:

$$\sum_{j=1}^M p(i, j) = 1 \quad \forall i = 1, \dots, M \quad (23)$$

and

$$p(i, j) > 0 \quad \forall i, j = 1, \dots, M \text{ satisfying } q_f(j | i) > 0 \quad (24)$$

At the same time, $\forall i, j = 1, \dots, M$ satisfying $q_f(j | i) = 0$, $p(i, j)$ must be equal to zero as well. Let us define the sets:

$$\mathcal{J}(i) \triangleq \{j = 1, \dots, M \mid q_f(j | i) > 0\}.$$

The partial derivatives in Eq. (22) are:

$$\frac{\partial \bar{\theta}_{\text{app}}^{(m)}(\mathbf{p})}{\partial p(i, j)} = \ell_s \frac{\partial \theta^*(\mathbf{p}^{(m)})}{\partial p^{(m)}(i, j)} + \frac{\partial \widehat{I}_3(\mathbf{p})}{\partial p(i, j)}$$

where

$$\frac{\partial \widehat{I}_3(\mathbf{p})}{\partial p(i, j)} = -\frac{q_1(i)q_f(j|i)}{p(i, j)} \quad (25)$$

Therefore Eq. (22) becomes:

$$\ell_s \frac{\partial \theta^*(\mathbf{p}^{(m)})}{\partial p^{(m)}(i, j)} - \frac{q_1(i)q_f(j|i)}{p(i, j)} = w_i \quad \forall i = 1, \dots, M \text{ and } j \in \mathcal{J}(i) \quad (26)$$

which is equivalent to

$$p(i, j) = \frac{q_1(i)q_f(j|i)}{\ell_s \frac{\partial \theta^*(\mathbf{p}^{(m)})}{\partial p^{(m)}(i, j)} - w_i} \quad \forall i = 1, \dots, M \text{ and } j \in \mathcal{J}(i) \quad (27)$$

In order for the $p(i, j)$ to be positive we need:

$$w_i < \ell s \frac{\partial \theta^*(\mathbf{p}^{(m)})}{\partial p^{(m)}(i, j)} \quad \forall i = 1, \dots, M \text{ and } j \in \mathcal{J}(i)$$

or equivalently

$$w_i < \min_{j \in \mathcal{J}(i)} \ell s \frac{\partial \theta^*(\mathbf{p}^{(m)})}{\partial p^{(m)}(i, j)} \triangleq w_{max} \quad \forall i$$

We can then use Eq. (23) to get:

$$\sum_{j \in \mathcal{J}(i)} \frac{q_1(i) q_f(j|i)}{\ell s \frac{\partial \theta^*(\mathbf{p}^{(m)})}{\partial p^{(m)}(i, j)} - w_i} = 1, \quad \forall i$$

The latest equation is a scalar equation in w_i which can be solved numerically in $(-\infty, w_{max})$. A unique solution is guaranteed since the function $\bar{\theta}_{app}^{(m)}(\mathbf{p})$ is strictly convex and therefore has a unique minimum. We need to solve M such equations one for each w_i . Then by substituting the w_i 's back in Eq. (27) we get all $p(i, j)$ that minimize the expression in Eq. (20) to form the optimal solution $\mathbf{p}_{\mathbf{p}^{(m)}}^*$.

Although Algorithm A performs well in practice, it does not guarantee convergence. This is the case because the approximations $\bar{\theta}_{app}^{(m)}(\mathbf{p})$ are good in a small region around the expansion point, but may be well off away from that point. This might lead to unpleasant situations where $\bar{\theta}(\mathbf{p}^{(m)}) > \bar{\theta}(\mathbf{p}^{(m-1)})$. For the same reason, a step of the algorithm might produce a $\mathbf{p}^{(m)}$ that belongs to \mathcal{I}_2 , independently of whether it is improving the value of the objective function. In this case the next step will yield $\mathbf{p}^{(m+1)} = \mathbf{q}_f$, since $\bar{\theta}_{app}^{(m+1)}(\mathbf{p}) = \hat{I}_3(\mathbf{p})$. Clearly, this will bring the algorithm into an infinite loop. Lastly, we might be really unlucky and land on a point $\mathbf{p}^{(m)} \in \mathcal{I}_3$. As we will see later on, $\nabla \theta^*(\mathbf{p}^{(m)})$ and consequently $\bar{\theta}_{app}^{(m)}(\mathbf{p})$ are not well defined at any $\mathbf{p}^{(m)} \in \mathcal{I}_3$ and hence the algorithm won't be able to continue.

Before we discuss an algorithm that does guarantee convergence, we need to calculate the gradient of $\bar{\theta}(\mathbf{p})$. The partial derivatives of $\hat{I}_3(\mathbf{p})$ are given in Eq. (25) and are continuous everywhere in \mathcal{I} . The calculation of the gradient of $\theta^*(\mathbf{p})$ in \mathcal{I} is a bit more involved:

Obviously, for $\mathbf{p} \in \mathcal{I}_2$, $\nabla \theta^*(\mathbf{p}) = 0$. In the following, we will calculate the gradient $\nabla \theta^*(\mathbf{p})$, for $\mathbf{p} \in \mathcal{I}_1$, i.e., in the region where $\theta^*(\mathbf{p}) > 0$:

Let us denote by λ_i the i -th eigenvalue of $\mathbf{\Pi}_{\theta, \mathbf{p}}^A$, and by $\mathbf{u}^{(i)}$ and $\mathbf{v}^{(i)}$ the (normalized) left and right eigenvectors corresponding to this eigenvalue respectively (thus the spectral radius $\rho(\mathbf{\Pi}_{\theta, \mathbf{p}}^A) = \lambda_1$). Note that all the above eigenvalues and eigenvectors are functions of \mathbf{p} , θ and the marginal probability densities of arrival rates per state in the arrival process A. We will often omit these dependences for brevity. Assuming that the entries of $\mathbf{\Pi}_{\theta, \mathbf{p}}^A$ are functions of some parameter δ and since λ_1 has always multiplicity 1 (Perron-Frobenius theorem), we have (see [MH88] and [Wil65, Sec. 2.5-2.8])

$$\frac{\partial \lambda_1}{\partial \delta} = \frac{\mathbf{u}^{(1)T} \left(\frac{\partial}{\partial \delta} \mathbf{\Pi}_{\theta, \mathbf{p}}^A \right) \mathbf{v}^{(1)}}{\mathbf{u}^{(1)T} \mathbf{v}^{(1)}}. \quad (28)$$

Using the definition of $\mathbf{\Pi}_{\theta, \mathbf{p}}^A$, i.e., $\mathbf{\Pi}_{\theta, \mathbf{p}}^A = \{\pi_{\theta, \mathbf{p}}^A(i, j)\}_{i, j=1, \dots, M}^M \triangleq \{p(i, j)\eta_j(\theta)\}_{i, j=1, \dots, M}^M$, the above equation for $\delta = p(i, j)$ becomes:

$$\frac{\partial \rho}{\partial p(i, j)} = \frac{u_i^{(1)} v_j^{(1)} \eta_j(\theta)}{\mathbf{u}^{(1)T} \mathbf{v}^{(1)}}. \quad (29)$$

and for $\delta = \theta$:

$$\frac{\partial \rho}{\partial \theta} = \sum_{i, j} \left[\frac{u_i^{(1)} v_j^{(1)}}{\mathbf{u}^{(1)T} \mathbf{v}^{(1)}} p(i, j) \frac{\partial \eta_j(\theta)}{\partial \theta} \right] \quad (30)$$

Now recall that $\Lambda_A(\theta, \mathbf{p}) = \log \rho(\mathbf{\Pi}_{\theta, \mathbf{p}}^A)$. Therefore:

$$\frac{\partial \Lambda_A(\theta, \mathbf{p})}{\partial p(i, j)} = \frac{\partial \log \rho(\mathbf{\Pi}_{\theta, \mathbf{p}}^A)}{\partial p(i, j)} = \frac{1}{\rho(\mathbf{\Pi}_{\theta, \mathbf{p}}^A)} \frac{\partial \rho(\mathbf{\Pi}_{\theta, \mathbf{p}}^A)}{\partial p(i, j)} = \frac{\eta_j(\theta)}{\rho(\mathbf{\Pi}_{\theta, \mathbf{p}}^A)} \frac{u_i^{(1)} v_j^{(1)}}{\mathbf{u}^{(1)T} \mathbf{v}^{(1)}} \quad (31)$$

Similarly,

$$\begin{aligned} \frac{\partial \Lambda_A(\theta, \mathbf{p})}{\partial \theta} &= \frac{1}{\rho(\mathbf{\Pi}_{\theta, \mathbf{p}}^A)} \frac{\partial \rho(\mathbf{\Pi}_{\theta, \mathbf{p}}^A)}{\partial \theta} \\ &= \frac{1}{\rho(\mathbf{\Pi}_{\theta, \mathbf{p}}^A)} \sum_{i, j} \left[\frac{u_i^{(1)} v_j^{(1)}}{\mathbf{u}^{(1)T} \mathbf{v}^{(1)}} p(i, j) \frac{\partial \eta_j(\theta)}{\partial \theta} \right]. \end{aligned} \quad (32)$$

Recall also that $\theta^*(\mathbf{p})$ is the largest root of the equation $\Lambda_A(\theta, \mathbf{p}) + \Lambda_B(-\theta) = 0$. Based on the discussion in Section II, $\theta^*(\mathbf{p})$ can be alternatively written as

$$\begin{aligned} \theta^*(\mathbf{p}) &= \sup_{\{\theta | \Lambda_A(\theta, \mathbf{p}) + \Lambda_B(-\theta) < 0\}} \theta \\ &= \sup_{\{\theta | \Lambda_A(\theta, \mathbf{p}) + \Lambda_B(-\theta) \leq 0\}} \theta \end{aligned} \quad (33)$$

$$= \inf_{a \geq 0} \sup_{\theta} [\theta - a \Lambda_A(\theta, \mathbf{p}) - a \Lambda_B(-\theta)], \quad (34)$$

where the second equality is due to the convexity of the limiting log-moment generating functions and the fact that $\Lambda_A(\theta, \mathbf{p}) + \Lambda_B(-\theta)$ is negative for sufficiently small $\theta > 0$. The third equality above is due to strong duality which holds since we are dealing with a convex programming problem (see [Ber95, Chapter 5]). Using the envelope theorem [Var92], we obtain

$$\nabla \theta^*(\mathbf{p}) = -a^* \nabla_{\mathbf{p}} \Lambda_A(\theta^*, \mathbf{p}), \quad (35)$$

where a^*, θ^* are optimal solutions of the optimization problems in (34), and the elements of $\nabla_{\mathbf{p}} \Lambda_A(\theta^*, \mathbf{p})$ are given by (31). Note that a^* is a Lagrange multiplier for the original optimization problem in (33) and therefore satisfies the first order optimality condition

$$1 - a^* \frac{\partial (\Lambda_A(\theta^*, \mathbf{p}) + \Lambda_B(-\theta^*))}{\partial \theta^*} = 0,$$

which implies

$$a^* = \left[\frac{\partial(\Lambda_A(\theta^*, \mathbf{p}) + \Lambda_B(-\theta^*))}{\partial \theta^*} \right]^{-1}. \quad (36)$$

Another way of deriving this result is by using the fact that $\Lambda_A(\theta^*(\mathbf{p}), \mathbf{p}) + \Lambda_B(-\theta^*(\mathbf{p}))$ is equal to zero for all \mathbf{p} satisfying $0 < \theta^*(\mathbf{p}) < +\infty$. Thus,

$$\nabla_{\mathbf{p}} \Lambda_A(\theta^*(\mathbf{p}), \mathbf{p}) + \nabla_{\mathbf{p}} \Lambda_B(-\theta^*(\mathbf{p})) = 0,$$

and by using the chain rule we obtain

$$\nabla \theta^*(\mathbf{p}) = - \frac{\nabla_{\mathbf{p}} \Lambda_A(\theta^*, \mathbf{p})}{\frac{\partial}{\partial \theta^*} (\Lambda_A(\theta^*, \mathbf{p}) + \Lambda_B(-\theta^*))}, \quad (37)$$

which is in agreement with (35)-(36).

It is crucial to understand the behavior of $\nabla \bar{\theta}(\mathbf{p})$ and the directional derivatives of $\bar{\theta}(\mathbf{p})$ along feasible directions as \mathbf{p} approaches the boundary \mathcal{I}_3 from within \mathcal{I}_1 . We will denote this by $\mathbf{p} \xrightarrow{\mathcal{I}_1} \mathbf{r}$, where \mathbf{r} is some arbitrary point belonging to \mathcal{I}_3 . Let us start by noting that due to the continuity of $\theta^*(\mathbf{p})$,

$$\lim_{\mathbf{p} \xrightarrow{\mathcal{I}_1} \mathbf{r}} \theta^*(\mathbf{p}) = \theta^*(\mathbf{r}) = 0$$

On the same token,

$$\lim_{\mathbf{p} \xrightarrow{\mathcal{I}_1} \mathbf{r}} \eta_j(\theta^*(\mathbf{p})) = \eta_j(0) = 1$$

which implies that:

$$\lim_{\mathbf{p} \xrightarrow{\mathcal{I}_1} \mathbf{r}} \mathbf{\Pi}_{\theta^*, \mathbf{p}}^A = \lim_{\mathbf{p} \xrightarrow{\mathcal{I}_1} \mathbf{r}} \{p(i, j) \eta_j(\theta^*(\mathbf{p}))\} = \{p(i, j)\} = \mathbf{\Xi}$$

It is known that the spectral radius of $\mathbf{\Xi}$, $\rho(\mathbf{\Xi}) = 1$ and that its non-normalized right eigenvector has all its elements equal to each other (the normalized one is $(1, 1, \dots, 1)$) while its normalized left eigenvector comprises of the steady state probabilities for the Markov chain represented by $\mathbf{\Xi}$.

Thus:

$$\lim_{\mathbf{p} \xrightarrow{\mathcal{I}_1} \mathbf{r}} \frac{u_i^{(1)} v_j^{(1)}}{\mathbf{u}^{(1)T} \mathbf{v}^{(1)}} = \frac{u_i^{(1)} v_j^{(1)}}{\sum_i u_i^{(1)} v_i^{(1)}} = \frac{u_i^{(1)} v_j^{(1)}}{v_j^{(1)} \sum_i u_i^{(1)}} = \frac{u_i^{(1)}}{\sum_i u_i^{(1)}} = p_1(i)$$

Furthermore, from the basic property of moment generating functions,

$$\lim_{\mathbf{p} \xrightarrow{\mathcal{I}_1} \mathbf{r}} \left. \frac{\partial \eta_j(\theta)}{\partial \theta} \right|_{\theta=\theta^*(\mathbf{p})} = \left. \frac{\partial \eta_j(\theta)}{\partial \theta} \right|_{\theta=0} = \mathbf{E}[A_1 | j]$$

Putting all these together we have that for all $\mathbf{r} \in \mathcal{I}_3$:

$$\begin{aligned}
\lim_{\substack{\mathcal{I}_1 \\ \mathbf{p} \rightarrow \mathbf{r}}} \frac{\partial \Lambda_A(\theta^*, \mathbf{p})}{\partial \theta^*} &= \lim_{\substack{\mathcal{I}_1 \\ \mathbf{p} \rightarrow \mathbf{r}}} \frac{1}{\rho(\mathbf{\Pi}_{\theta^*, \mathbf{p}}^A)} \sum_{i,j} \frac{u_i^{(1)} v_j^{(1)}}{\mathbf{u}^{(1)T} \mathbf{v}^{(1)}} p(i, j) \frac{\partial \eta_j(\theta^*)}{\partial \theta^*} = \sum_{i,j} p_1(i) p(i, j) \mathbf{E}[A_1 | j] \\
&= \sum_j \mathbf{E}[A_1 | j] \sum_i p_1(i) p(i, j) = \sum_j p_1(j) \mathbf{E}[A_1 | j] = \mathbf{E}[B_1]
\end{aligned}$$

Moreover,

$$\frac{\partial \Lambda_B(-\theta^*)}{\partial \theta^*} = -\mathbf{E}[B_1]$$

Therefore the denominator in Eq. (37) goes to zero for $\mathbf{p} \xrightarrow{\mathcal{I}_1} \mathbf{r}$. At the same time the components of the numerator become:

$$\lim_{\substack{\mathcal{I}_1 \\ \mathbf{p} \rightarrow \mathbf{r}}} \frac{\partial \Lambda_A(\theta^*, \mathbf{p})}{\partial p(i, j)} = \lim_{\substack{\mathcal{I}_1 \\ \mathbf{p} \rightarrow \mathbf{r}}} \frac{\eta_j(\theta^*)}{\rho(\mathbf{\Pi}_{\theta^*, \mathbf{p}}^A)} \frac{u_i^{(1)} v_j^{(1)}}{\mathbf{u}^{(1)T} \mathbf{v}^{(1)}} = p_1(i) \neq 0$$

We conclude that all partial derivatives tend to (plus or minus) infinity as $\mathbf{p} \xrightarrow{\mathcal{I}_1} \mathbf{r}$. However, this is not necessarily the case for directional derivatives along feasible directions. Let us denote by $\mathcal{D}(\mathbf{p})$ the set of all feasible directions \mathbf{d} at point \mathbf{p} . Note that all feasible directions are of the form:

$$\mathbf{d} = (d_{11}, d_{12}, \dots, d_{1M}, \dots, d_{M1}, d_{M2}, \dots, d_{MM}) \quad \text{with} \quad \sum_j d_{ij} = 0, \quad \forall i.$$

In order to calculate the directional derivative we consider a normalized feasible direction vector which satisfies the additional constraint:

$$\sqrt{\sum_{i,j} d_{ij}^2} = 1,$$

so that moving along the feasible direction is expressed as $\mathbf{p} + \beta \mathbf{d}$

We denote by \mathbf{D} the $M \times M$ matrix corresponding to \mathbf{d} and by \mathbf{H} the diagonal matrix with elements $\eta_i(\theta)$. We also denote by $\mathbf{H}_\theta \triangleq \frac{\partial \mathbf{H}}{\partial \theta}$, by $\mathbf{H}_{\theta\theta} \triangleq \frac{\partial^2 \mathbf{H}}{\partial \theta^2}$, by $\mathbf{\Pi}_\theta \triangleq \frac{\partial \mathbf{\Pi}}{\partial \theta}$ and by $\mathbf{\Pi}_\beta \triangleq \frac{\partial \mathbf{\Pi}}{\partial \beta}$. Note that:

$$\mathbf{\Pi} = \mathbf{P} \cdot \mathbf{H} \Rightarrow \begin{cases} \mathbf{\Pi}_\theta = \mathbf{P} \cdot \mathbf{H}_\theta \\ \mathbf{\Pi}_\beta = \mathbf{D} \cdot \mathbf{H} \end{cases}$$

Based on the above notation we can also write:

$$\frac{\partial \Lambda_A}{\partial \theta} = \frac{1}{\rho} \frac{\mathbf{u}^{(1)T} \mathbf{P} \mathbf{H}_\theta \mathbf{v}^{(1)}}{\mathbf{u}^{(1)T} \mathbf{v}^{(1)}}. \quad (38)$$

Now, the directional derivative of $\rho(\mathbf{\Pi}_{\theta, \mathbf{p}}^A)$ is derived from (28) as:

$$\frac{\partial \lambda_1}{\partial \beta} = \frac{\mathbf{u}^{(1)T} \left(\frac{\partial}{\partial \beta} \mathbf{\Pi}_{\theta, \mathbf{p}}^A \right) \mathbf{v}^{(1)}}{\mathbf{u}^{(1)T} \mathbf{v}^{(1)}} = \frac{\mathbf{u}^{(1)T} \mathbf{D} \mathbf{H} \mathbf{v}^{(1)}}{\mathbf{u}^{(1)T} \mathbf{v}^{(1)}}. \quad (39)$$

Therefore:

$$\frac{\partial \Lambda_A}{\partial \beta} = \frac{1}{\rho} \frac{\mathbf{u}^{(1)T} \mathbf{D} \mathbf{H} \mathbf{v}^{(1)}}{\mathbf{u}^{(1)T} \mathbf{v}^{(1)}}. \quad (40)$$

By the same arguments used above to obtain (37) we obtain:

$$\frac{\partial \theta^*}{\partial \beta} = - \frac{\frac{\partial \Lambda_A(\theta^*, \mathbf{p})}{\partial \beta}}{\frac{\partial}{\partial \theta^*} (\Lambda_A(\theta^*, \mathbf{p}) + \Lambda_B(-\theta^*))} = - \frac{\frac{1}{\rho} \frac{\mathbf{u}^{(1)T} \mathbf{D} \mathbf{H} \mathbf{v}^{(1)}}{\mathbf{u}^{(1)T} \mathbf{v}^{(1)}}}{\frac{\partial}{\partial \theta^*} (\Lambda_A(\theta^*, \mathbf{p}) + \Lambda_B(-\theta^*))}, \quad (41)$$

The denominator of this expression goes to zero for $\mathbf{p} \xrightarrow{\mathcal{I}_1} \mathbf{r}$, as discussed above. But the numerator goes to zero as well, since:

$$\lim_{\mathbf{p} \xrightarrow{\mathcal{I}_1} \mathbf{r}} \frac{1}{\rho} \frac{\mathbf{u}^{(1)T} \mathbf{D} \mathbf{H} \mathbf{v}^{(1)}}{\mathbf{u}^{(1)T} \mathbf{v}^{(1)}} = \lim_{\mathbf{p} \xrightarrow{\mathcal{I}_1} \mathbf{r}} \frac{1}{\rho} \sum_{i,j} d_{ij} \eta_j(\theta^*) \frac{u_i^{(1)} v_j^{(1)}}{\mathbf{u}^{(1)T} \mathbf{v}^{(1)}} = \sum_{i,j} d_{ij} p_1(i) = \sum_i p_1(i) \sum_j d_{ij} = 0$$

It can be proven using De L' Hôpital's theorem that the limit of all directional derivatives towards directions pointing into \mathcal{I}_1 as $\mathbf{p} \xrightarrow{\mathcal{I}_1} \mathbf{r}$ is zero. However, the proof is quite lengthy and is omitted due to space limitations.

Thus, the necessary and sufficient (due to convexity) condition for a point $\mathbf{p}^* \in \mathcal{I}$ to be the global minimum of $\bar{\theta}(\mathbf{p})$ can be written as:

$$\lim_{a \downarrow 0} \frac{\bar{\theta}(\mathbf{p}^* + a \mathbf{d}) - \bar{\theta}(\mathbf{p}^*)}{a} \geq 0, \quad \forall \mathbf{d} \in \mathcal{D}(\mathbf{p})$$

Lemma IV.1 *The direction sequence $\{\mathbf{d}^{(m)}\} = \{\mathbf{p}_{\mathbf{p}^{(m)}}^* - \mathbf{p}^{(m)}\}$ is gradient related to $\{\mathbf{p}^{(m)}\}$ in $\mathcal{I}_1 \cup \mathcal{I}_2$, that is for any subsequence $\{\mathbf{p}^{(m)}\}_{m \in \mathcal{M}}$ such that $\mathbf{p}^{(m)} \notin \mathcal{I}_3$, that converges to a non-stationary point $\mathbf{p}^{(0)} \in \mathcal{I}$ of $\bar{\theta}(\mathbf{p})$, the corresponding subsequence $\{\mathbf{d}^{(m)}\}_{m \in \mathcal{M}}$ is bounded and satisfies*

$$\limsup_{m \rightarrow \infty, m \in \mathcal{M}} \nabla \bar{\theta}(\mathbf{p}^{(m)})' \mathbf{d}^{(m)} < 0$$

Proof : By definition, $\mathbf{p}_{\mathbf{p}^{(m)}}^* \in \mathcal{I}$ which is bounded. Therefore, $\{\mathbf{d}^{(m)}\}$ is bounded, too.

Note that for all $\mathbf{p}^{(m)}$, $\nabla \bar{\theta}(\mathbf{p}^{(m)}) = \nabla \bar{\theta}_{\mathbf{p}^{(m)}}^{\text{app}}(\mathbf{p}^{(m)})$.

By definition:

$$\nabla \bar{\theta}_{\mathbf{p}^{(m)}}^{\text{app}}(\mathbf{p}^{(m)})' (\mathbf{p}_{\mathbf{p}^{(m)}}^* - \mathbf{p}^{(m)}) = \lim_{a \downarrow 0} \frac{\bar{\theta}_{\mathbf{p}^{(m)}}^{\text{app}}(\mathbf{p}^{(m)} + a(\mathbf{p}_{\mathbf{p}^{(m)}}^* - \mathbf{p}^{(m)})) - \bar{\theta}_{\mathbf{p}^{(m)}}^{\text{app}}(\mathbf{p}^{(m)})}{a} \quad (42)$$

Now note that $\bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\cdot)$ is a strictly convex function whose global minimum is $\mathbf{p}_{\mathbf{p}^{(m)}}^*$. Thus,

$$\begin{aligned}
& \frac{\bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)} + a(\mathbf{p}_{\mathbf{p}^{(m)}}^* - \mathbf{p}^{(m)})) - \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)})}{a} = \\
& \frac{\bar{\theta}_{\mathbf{p}^{(m)}}^{app}((1-a)\mathbf{p}^{(m)} + a\mathbf{p}_{\mathbf{p}^{(m)}}^*) - \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)})}{a} < \\
& \frac{(1-a)\bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)}) + a\bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}_{\mathbf{p}^{(m)}}^*) - \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)})}{a} = \\
& \frac{-a\bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)}) + a\bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}_{\mathbf{p}^{(m)}}^*)}{a} = \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}_{\mathbf{p}^{(m)}}^*) - \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)})
\end{aligned} \tag{43}$$

Equations (42) and (43) imply that:

$$\nabla \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)})'(\mathbf{p}_{\mathbf{p}^{(m)}}^* - \mathbf{p}^{(m)}) \leq \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}_{\mathbf{p}^{(m)}}^*) - \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)})$$

Therefore:

$$\limsup_{m \rightarrow \infty, m \in \mathcal{M}} \nabla \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)})'(\mathbf{p}_{\mathbf{p}^{(m)}}^* - \mathbf{p}^{(m)}) \leq \limsup_{m \rightarrow \infty, m \in \mathcal{M}} \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}_{\mathbf{p}^{(m)}}^*) - \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)})$$

We now distinguish between two cases: $\mathbf{p}^{(0)} \in \mathcal{I}_3$ and $\mathbf{p}^{(0)} \notin \mathcal{I}_3$.

If $\mathbf{p}^{(0)} \notin \mathcal{I}_3$ then by continuity arguments we can conclude that:

$$\limsup_{m \rightarrow \infty, m \in \mathcal{M}} \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}_{\mathbf{p}^{(m)}}^*) - \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)}) = \bar{\theta}_{\mathbf{p}^{(0)}}^{app}(\mathbf{p}_{\mathbf{p}^{(0)}}^*) - \bar{\theta}_{\mathbf{p}^{(0)}}^{app}(\mathbf{p}^{(0)})$$

But $\bar{\theta}_{\mathbf{p}^{(0)}}^{app}(\cdot)$ is also a strictly convex function whose global minimum is $\mathbf{p}_{\mathbf{p}^{(0)}}^*$ and therefore unless $\mathbf{p}^{(0)} = \mathbf{p}_{\mathbf{p}^{(0)}}^*$ (in which case $\mathbf{p}^{(0)}$ is the global minimum of both $\bar{\theta}_{\mathbf{p}^{(0)}}^{app}(\cdot)$ and $\bar{\theta}(\cdot)$), $\bar{\theta}_{\mathbf{p}^{(0)}}^{app}(\mathbf{p}_{\mathbf{p}^{(0)}}^*) - \bar{\theta}_{\mathbf{p}^{(0)}}^{app}(\mathbf{p}^{(0)}) < 0$.

If $\mathbf{p}^{(0)} \in \mathcal{I}_3$, we have three sub-cases:

1. $\exists m_1 \in \mathcal{M}$ s.t. $\forall m \in \mathcal{M}$ with $m > m_1$, $\mathbf{p}^{(m)} \in \mathcal{I}_2$.

In this case $\forall m \in \mathcal{M}$ with $m > m_1$ we have $\mathbf{p}_{\mathbf{p}^{(m)}}^* = \mathbf{q}_f$ which yields $\bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}_{\mathbf{p}^{(m)}}^*) = 0$ and $\bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)}) = \hat{I}_3(\mathbf{p}^{(m)})$. Hence:

$$\limsup_{m \rightarrow \infty, m \in \mathcal{M}} \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}_{\mathbf{p}^{(m)}}^*) - \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)}) = -\hat{I}_3(\mathbf{p}^{(0)})$$

But this is always less than zero, since $\mathbf{q}_f \in \mathcal{I}_1$ and $\mathbf{p}^{(0)} \in \mathcal{I}_3$.

2. $\exists m_2 \in \mathcal{M}$ s.t. $\forall m \in \mathcal{M}$ with $m > m_2$, $\mathbf{p}^{(m)} \in \mathcal{I}_1$.

Consider a function $\bar{\varphi}(\mathbf{p})$ which is equal to $\bar{\theta}(\mathbf{p})$, $\forall \mathbf{p} \in \mathcal{I}_4 = \mathcal{I}_1 \cup \mathcal{I}_3$ and continuously differentiable at any $\mathbf{r} \in \mathcal{I}_3$. Define $\bar{\varphi}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p})$ in the same way that $\bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p})$ is defined. Then:

$$\limsup_{m \rightarrow \infty, m \in \mathcal{M}} \bar{\varphi}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}_{\mathbf{p}^{(m)}}^*) - \bar{\varphi}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)}) = \bar{\varphi}_{\mathbf{p}^{(0)}}^{app}(\mathbf{p}_{\mathbf{p}^{(0)}}^*) - \bar{\varphi}_{\mathbf{p}^{(0)}}^{app}(\mathbf{p}^{(0)})$$

But $\bar{\varphi}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p})$ is also strictly convex (by definition) and therefore $\bar{\varphi}_{\mathbf{p}^{(0)}}^{app}(\mathbf{p}_{\mathbf{p}^{(0)}}^*) - \bar{\varphi}_{\mathbf{p}^{(0)}}^{app}(\mathbf{p}^{(0)}) < 0$ unless $\mathbf{p}^{(0)}$ is the global minimum of $\bar{\varphi}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p})$. In that case $\mathbf{p}^{(0)}$ would also be the global minimum of $\bar{\varphi}(\mathbf{p})$ which in turn would make it the global minimum of $\bar{\theta}(\mathbf{p})$ in \mathcal{I}_4 . But as seen in the proof of Lemma III.2 the global minimum of $\bar{\theta}(\mathbf{p})$ is always in \mathcal{I}_4 which means that $\mathbf{p}^{(0)}$ should be the global minimum of $\bar{\theta}(\mathbf{p})$ in \mathcal{I} . This clearly contradicts our assumption that $\mathbf{p}^{(0)}$ is not a stationary point of $\bar{\theta}(\mathbf{p})$. Hence, $\bar{\varphi}_{\mathbf{p}^{(0)}}^{app}(\mathbf{p}_{\mathbf{p}^{(0)}}^*) - \bar{\varphi}_{\mathbf{p}^{(0)}}^{app}(\mathbf{p}^{(0)})$ is always less than 0.

3. \mathcal{M} can be partitioned into two unbounded sets \mathcal{M}_1 and \mathcal{M}_2 with $\mathbf{p}^{(m)} \in \mathcal{I}_1, \forall m \in \mathcal{M}_1$ and $\mathbf{p}^{(m)} \in \mathcal{I}_2, \forall m \in \mathcal{M}_2$. Then:

$$\begin{aligned} & \limsup_{m \rightarrow \infty, m \in \mathcal{M}} \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}_{\mathbf{p}^{(m)}}^*) - \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)}) = \\ & = \max[\limsup_{m \rightarrow \infty, m \in \mathcal{M}_1} \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}_{\mathbf{p}^{(m)}}^*) - \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)}), \\ & \quad \limsup_{m \rightarrow \infty, m \in \mathcal{M}_2} \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}_{\mathbf{p}^{(m)}}^*) - \bar{\theta}_{\mathbf{p}^{(m)}}^{app}(\mathbf{p}^{(m)})] \end{aligned}$$

But each one of these two limits superior belongs to the special cases analyzed above. Therefore, each one is less than zero and the maximum is always less than zero.

This completes the proof. ■

Lemma IV.1 can be used to prove that the following algorithm always converges:

Algorithm B:

1-3. Steps 1-3 are exactly the same as in Algorithm A.

4. Set $\mathbf{p}^{(m+1)} = \mathbf{p}^{(m)} + a^{(m)}(\mathbf{p}_{\mathbf{p}^{(m)}}^* - \mathbf{p}^{(m)})$ where $a^{(m)}$ is the stepsize. Use a slightly modified Armijo rule [Ber95] to select the stepsize. To implement this rule, fixed scalars $\beta \in (0, 1)$ and $\sigma \in (0, 1)$ are chosen and we set $a^{(m)} = \beta^{l'}$, where l' is the first nonnegative integer l for which

$$\mathbf{p}^{(m)} + \beta^l(\mathbf{p}_{\mathbf{p}^{(m)}}^* - \mathbf{p}^{(m)}) \notin \mathcal{I}_3$$

and

$$\bar{\theta}(\mathbf{p}^{(m)}) - \bar{\theta}(\mathbf{p}^{(m)} + \beta^l(\mathbf{p}_{\mathbf{p}^{(m)}}^* - \mathbf{p}^{(m)})) \geq -\sigma \beta^l \nabla \bar{\theta}(\mathbf{p}^{(m)})'(\mathbf{p}_{\mathbf{p}^{(m)}}^* - \mathbf{p}^{(m)})$$

Such a stepsize is guaranteed to exist since the set of stepsizes that satisfy the above inequality will always contain an interval of the form $[0, \delta]$ with $\delta > 0$ and $\mathbf{p}^{(m)} \notin \mathcal{I}_3 \Rightarrow \exists \varepsilon > 0$ s.t. $\forall 0 \leq a < \varepsilon, \mathbf{p}^{(m)} + a(\mathbf{p}_{\mathbf{p}^{(m)}}^* - \mathbf{p}^{(m)}) \notin \mathcal{I}_3$. Return to step 2.

Note that for $l = 0$, $\mathbf{p}^{(m+1)} = \mathbf{p}_{\mathbf{p}^{(m)}}^*$ and therefore if the above conditions are satisfied at first trial, Algorithm B behaves exactly as Algorithm A. Thus, Algorithm B guarantees convergence without sacrificing much in performance.

Theorem IV.2 *Let $\{\mathbf{p}^{(m)}\}$ be a sequence generated by Algorithm B. Then $\{\mathbf{p}^{(m)}\}$ converges to \mathbf{p}^* , the global minimum of $\bar{\theta}(\mathbf{p})$.*

Proof : (based on the proof of Proposition 1.2.1 in [Ber95]):

Assume that \mathbf{p}_0 is a limit point of $\{\mathbf{p}^{(m)}\}$ that is not the global minimum, i.e.,

$$\exists \mathbf{d} \in \mathcal{D}(\mathbf{p}) \text{ s.t. } \lim_{a \downarrow 0} \frac{\bar{\theta}(\mathbf{p}_0 + a\mathbf{d}) - \bar{\theta}(\mathbf{p}_0)}{a} < 0 \quad (44)$$

Note that since $\{\bar{\theta}(\mathbf{p}^{(m)})\}$ is monotonically non-increasing, it either converges to a finite value or diverges to $-\infty$. Since $\bar{\theta}(\mathbf{p})$ is continuous, $\bar{\theta}(\mathbf{p}_0)$ is a limit point of $\{\bar{\theta}(\mathbf{p}^{(m)})\}$ and it follows that the entire sequence $\{\bar{\theta}(\mathbf{p}^{(m)})\}$ converges to $\bar{\theta}(\mathbf{p}_0)$. Hence, $\bar{\theta}(\mathbf{p}^{(m)}) - \bar{\theta}(\mathbf{p}^{(m+1)}) \rightarrow 0$.

By the definition of the stepsize selection rule in Algorithm B, we have that all $\mathbf{p}^{(m)} \notin \mathcal{I}_3$ (although \mathbf{p}_0 can belong to \mathcal{I}_3) and:

$$\bar{\theta}(\mathbf{p}^{(m)}) - \bar{\theta}(\mathbf{p}^{(m+1)}) \geq -\sigma a^{(m)} \nabla \bar{\theta}(\mathbf{p}^{(m)})' \mathbf{d}^{(m)}$$

Hence, $a^{(m)} \nabla \bar{\theta}(\mathbf{p}^{(m)})' \mathbf{d}^{(m)} \rightarrow 0$. Let $\{\mathbf{p}^{(m)}\}_{\mathcal{M}}$ be a subsequence converging to \mathbf{p}_0 . From Lemma IV.1, we have

$$\limsup_{m \rightarrow \infty, m \in \mathcal{M}} \nabla \bar{\theta}(\mathbf{p}^{(m)})' \mathbf{d}^{(m)} < 0$$

and therefore $\{a^{(m)}\}_{\mathcal{M}} \rightarrow 0$.

Now, by the definition of the stepsize selection rule and the fact that $\{a^{(m)}\}_{\mathcal{M}}$ goes to zero, there must be some \bar{m} such that the initial stepsize s will be reduced for at least two times for all $m \in \mathcal{M}$, $m \geq \bar{m}$. But due to the shape and dimension of \mathcal{I}_3 , its intersection with the line segment connecting $\mathbf{p}^{(m)}$ and $\mathbf{p}_{\mathbf{p}^{(m)}}^*$ can be at most a single point. Therefore, at most one of the stepsize reductions occurs because $\mathbf{p}^{(m)} + \beta^l (\mathbf{p}_{\mathbf{p}^{(m)}}^* - \mathbf{p}^{(m)}) \in \mathcal{I}_3$, and at least one of the following two inequalities holds:

$$\bar{\theta}(\mathbf{p}^{(m)}) - \bar{\theta}(\mathbf{p}^{(m)} + (a^{(m)}/\beta)\mathbf{d}^{(m)}) < -\sigma(a^{(m)}/\beta) \nabla \bar{\theta}(\mathbf{p}^{(m)})' \mathbf{d}^{(m)} \quad (45)$$

or

$$\bar{\theta}(\mathbf{p}^{(m)}) - \bar{\theta}(\mathbf{p}^{(m)} + (a^{(m)}/\beta^2)\mathbf{d}^{(m)}) < -\sigma(a^{(m)}/\beta^2) \nabla \bar{\theta}(\mathbf{p}^{(m)})' \mathbf{d}^{(m)} \quad (46)$$

Let us denote by \mathcal{M}_1 the set of m 's for which Eq. (45) holds and by \mathcal{M}_2 , the set of m 's for which Eq. (45) does not hold and therefore Eq. (46) does hold. Define:

$$\mathbf{x}^{(m)} = \frac{\mathbf{d}^{(m)}}{\|\mathbf{d}^{(m)}\|}$$

and

$$\bar{a}^{(m)} = \begin{cases} \frac{a^{(m)} \|\mathbf{d}^{(m)}\|}{\beta} & \text{if } m \in \mathcal{M}_1 \\ \frac{a^{(m)} \|\mathbf{d}^{(m)}\|}{\beta^2} & \text{if } m \in \mathcal{M}_2 \end{cases}$$

Since $\{\mathbf{d}^{(m)}\}$ is gradient related, $\{\|\mathbf{d}^{(m)}\|\}_{\mathcal{M}}$ is bounded and therefore $\{\bar{a}^{(m)}\}_{\mathcal{M}} \rightarrow 0$. Since $\|\mathbf{x}^{(m)}\| = 1$ for all $m \in \mathcal{M}$, there exists a subsequence $\{\mathbf{x}^{(m)}\}_{\bar{\mathcal{M}}}$ of $\{\mathbf{x}^{(m)}\}_{\mathcal{M}}$ such that $\{\mathbf{x}^{(m)}\}_{\bar{\mathcal{M}}} \rightarrow \bar{\mathbf{x}}$ where $\bar{\mathbf{x}}$ is some vector with $\|\bar{\mathbf{x}}\| = 1$ (see Prop. A.5(c) in Appendix A of [Ber95]).

Given the above definitions, Eq. (45) and (46) can be combined into:

$$\frac{\bar{\theta}(\mathbf{p}^{(m)}) - \bar{\theta}(\mathbf{p}^{(m)} + \bar{a}^{(m)}\mathbf{x}^{(m)})}{\bar{a}^{(m)}} < -\sigma \nabla \bar{\theta}(\mathbf{p}^{(m)})' \mathbf{x}^{(m)}, \quad \forall m \in \bar{\mathcal{M}}, m \geq \bar{m}$$

Since $\bar{\theta}(\mathbf{p})$ is a convex and continuous function there must be some $\tilde{a}^{(m)} \in [0, \bar{a}^{(m)}]$ with $\mathbf{p}^{(m)} + \tilde{a}^{(m)}\mathbf{x}^{(m)} \notin \mathcal{I}_3$ such that:

$$\begin{aligned} -\nabla \bar{\theta}(\mathbf{p}^{(m)} + \tilde{a}^{(m)}\mathbf{x}^{(m)})' \mathbf{x}^{(m)} &\leq \frac{\bar{\theta}(\mathbf{p}^{(m)}) - \bar{\theta}(\mathbf{p}^{(m)} + \tilde{a}^{(m)}\mathbf{x}^{(m)})}{\tilde{a}^{(m)}} \\ &< -\sigma \nabla \bar{\theta}(\mathbf{p}^{(m)})' \mathbf{x}^{(m)}, \quad \forall m \in \bar{\mathcal{M}}, m \geq \bar{m} \end{aligned}$$

Taking limits in the above equation we obtain

$$-\lim_{a \downarrow 0} \frac{\bar{\theta}(\mathbf{p}_0 + a\bar{\mathbf{x}}) - \bar{\theta}(\mathbf{p}_0)}{a} \leq -\sigma \lim_{a \downarrow 0} \frac{\bar{\theta}(\mathbf{p}_0 + a\bar{\mathbf{x}}) - \bar{\theta}(\mathbf{p}_0)}{a}$$

or

$$0 \leq (1 - \sigma) \lim_{a \downarrow 0} \frac{\bar{\theta}(\mathbf{p}_0 + a\bar{\mathbf{x}}) - \bar{\theta}(\mathbf{p}_0)}{a}$$

And since $\sigma < 1$, it follows that

$$\lim_{a \downarrow 0} \frac{\bar{\theta}(\mathbf{p}_0 + a\bar{\mathbf{x}}) - \bar{\theta}(\mathbf{p}_0)}{a} \geq 0$$

But this contradicts Eq. (44) and therefore \mathbf{p}_0 can only be the global minimum of $\bar{\theta}(\mathbf{p})$. ■

V. CONCLUSIONS

Building on our earlier work on suitable estimators for a wide set of Large Deviation approximations for the probabilities of rare events, we have analyzed the non-linear multi-dimensional optimization problems that need to be solved in order to calculate these estimators. The special structure of the objective function of the optimization problems at hand, suggests that a custom-made algorithm can be much more efficient than any generic non-linear optimization algorithm. However, proving that such an algorithm always converges is not straightforward. In order to do so, we had to modify the simpler algorithm presented in our earlier work so that its selected step size follows a modified Armijo rule. This modified algorithm is proven to exhibit guaranteed convergence.

The LD estimators considered in this paper have found many applications in real-time problems such as Call Admission Control through buffer overflow probability prediction, adaptive modulation and coding with QoS constraints for transmission of data over wireless links, make-to-stock manufacturing systems and supply chains, and traffic anomaly detection in high data rate communication networks. In all the above applications, the existence of efficient and provably convergent algorithms for solving the estimation problem is a prerequisite for using the proposed LD estimators for real-time decision and control. The algorithm developed in this paper meets these requirements.

REFERENCES

- [Ber95] D.P. Bertsekas, *Nonlinear programming*, Athena Scientific, Belmont, MA, 1995.
- [BP01] D. Bertsimas and I. Ch. Paschalidis, *Probabilistic service level guarantees in make-to-stock manufacturing systems*, *Operations Research* **49** (2001), no. 1, 119–133.
- [BPT98a] D. Bertsimas, I. Ch. Paschalidis, and J. N. Tsitsiklis, *Asymptotic buffer overflow probabilities in multiclass multiplexers: An optimal control approach*, *IEEE Transactions on Automatic Control* **43** (1998), no. 3, 315–335.
- [BPT98b] D. Bertsimas, I. Ch. Paschalidis, and J. N. Tsitsiklis, *On the large deviations behaviour of acyclic networks of G/G/1 queues*, *The Annals of Applied Probability* **8** (1998), no. 4, 1027–1069.
- [Duf99] N.G. Duffield, *A large deviation analysis of errors in measurement based admission control to buffered and bufferless resources*, *Proceedings of the IEEE INFOCOM*, 1999.
- [DZ98] A. Dembo and O. Zeitouni, *Large deviations techniques and applications*, 2nd ed., Springer-Verlag, NY, 1998.
- [GAN91] R. Guérin, H. Ahmadi, and Naghshineh, *Equivalent capacity and its applications to bandwidth allocation in high-speed networks*, *IEEE Journal on Selected Areas in Communications* **9** (1991), 968–981.
- [GH91] R.J. Gibbens and P.J. Hunt, *Effective bandwidths for the multi-type UAS channel*, *Queueing Systems* **9** (1991), 17–28.
- [Gla97] P. Glasserman, *Bounds and asymptotics for planning critical safety stocks*, *Operations Research* **45** (1997), no. 2, 244–257.
- [GT99a] M. Grossglauser and D. Tse, *A framework for robust measurement-based admission control*, *IEEE/ACM Transactions on Networking* **7** (1999), no. 3, 293–309.
- [GT99b] ———, *A time-scale decomposition approach to measurement-based admission control*, *Proceedings of the IEEE INFOCOM*, 1999.
- [GW94] P.W. Glynn and W. Whitt, *Logarithmic asymptotics for steady-state tail probabilities in a single-server queue*, *Journal of Applied Probability* **31A** (1994), 131–156.
- [Hui88] J. Y. Hui, *Resource allocation for broadband networks*, *IEEE Journal on Selected Areas in Communications* **6** (1988), no. 9, 1598–1608.
- [Kel91] F. P. Kelly, *Effective bandwidths at multi-class queues*, *Queueing Systems* **9** (1991), 5–16.
- [Kel96] ———, *Notes on effective bandwidths*, *Stochastic Networks: Theory and Applications* (S. Zachary, I.B. Ziedins, and F.P. Kelly, eds.), vol. 9, Oxford University Press, 1996, pp. 141–168.
- [LZG04] Q. Liu, S. Zhou, and G. B. Giannakis, *Cross-layer combining of adaptive modulation and coding with truncated arq over wireless links*, *IEEE Transactions on Wireless Communications* **3** (2004), no. 5, 1746–1755.
- [LZG05] ———, *Queueing with adaptive modulation and coding over wireless links: cross-layer analysis and design*, *IEEE Transactions on Wireless Communications* **4** (2005), no. 3, 1142–1153.
- [LZG06] ———, *Cross-layer modeling of adaptive wireless link for qos support in heterogeneous wired-wireless networks*, *Wireless Networks*, Kluwer Academic Publishers **12** (2006), no. 4, 427–437.
- [MH88] D.V. Murthy and R.T. Haftka, *Derivatives of eigenvalues and eigenvectors of a general complex matrix*, *Intl. Journal for Numerical Methods in Engineering* **26** (1988), no. 2, 293–311.
- [Pas99] I. Ch. Paschalidis, *Class-specific quality of service guarantees in multimedia communication networks*, *Automatica*, (Special Issue on Control Methods for Communication Networks), Anantharam and Walrand Eds., **35** (1999), no. 12, 1951–1968.
- [PC10] I. Ch. Paschalidis and Y. Chen, *Statistical anomaly detection with sensor networks*, *ACM Transactions on Sensor Networks* **7** (2010), no. 3, 17:1–17:23.
- [PL03] I.Ch. Paschalidis and Y. Liu, *Large deviations-based asymptotics for inventory control in supply chains*, *Operations Research* **51** (2003), no. 3, 437–460.
- [PS09] I. Ch. Paschalidis and G. Smaragdakis, *Spatio-temporal network anomaly detection by assessing deviations of empirical measures*, *IEEE/ACM Transactions on Networking* **17** (2009), no. 3, 685–697.
- [PV01] I. Ch. Paschalidis and S. Vassilaras, *On the estimation of buffer overflow probabilities from measurements*, *IEEE Transactions on Information Theory* **47** (2001), no. 1, 178–191.
- [TZ06] J. Tang and X. Zhang, *Cross-layer-model based adaptive resource allocation for statistical qos guarantees in mobile wireless networks*, *QShine06 Third International Conference on Quality of Service in Heterogeneous Wired/Wireless Networks* (Waterloo, Canada), August 2006.

- [TZ07a] ———, *Cross-layer modeling for quality of service guarantees over wireless links*, IEEE Transactions on Wireless Communications **6** (2007), no. 12.
- [TZ07b] ———, *Quality-of-service driven power and rate adaptation over wireless links*, IEEE Transactions on Wireless Communications **6** (2007), no. 8.
- [Var92] H.R. Varian, *Microeconomic analysis*, 3rd ed., W. W. Norton and Company, 1992.
- [Vas10] S. Vassilaras, *A cross-layer optimized adaptive modulation and coding scheme for transmission of streaming media over wireless links*, Springer / ACM Wireless Networks **16** (2010), no. 4, 903–914.
- [Wil65] J.H. Wilkinson, *The algebraic eigenvalue problem*, Oxford University Press, 1965.
- [WN03] D. Wu and R. Negi, *Effective capacity: a wireless link model for support of quality of service*, IEEE Transactions on Wireless Communications **2** (2003), no. 4, 630–643.
- [Wu03] D. Wu, *Providing quality-of-service guarantees in wireless networks*, Ph.D. thesis, Department of Electrical and Computer Engineering, Carnegie Mellon University, 2003.