Perturbation Analysis for Online Control and Optimization of Stochastic Fluid Models

Christos G. Cassandras, *Fellow, IEEE*, Yorai Wardi, *Member, IEEE*, Benjamin Melamed, *Fellow, IEEE*, Gang Sun, and Christos G. Panayiotou, *Member, IEEE*

***Abstract*—This paper uses stochastic fluid models (SFMs) for control and optimization (rather than performance analysis) of communication networks, focusing on problems of buffer control. We derive gradient estimators for packet loss and workload related performance metrics with respect to threshold parameters. These estimators are shown to be unbiased and directly observable from a sample path without any knowledge of underlying stochastic characteristics, including traffic and processing rates (i.e., they are nonparametric). This renders them computable in online environments and easily implementable for network management and control. We further demonstrate their use in buffer control problems where our SFM-based estimators are evaluated based on data from an actual system.**

***Index Terms*—Communication network, perturbation analysis, stochastic fluid network.**

I. Introduction

A NATURAL modeling framework for packet-based communication networks is provided through queuing systems. However, the huge traffic volume that networks are supporting today makes such models highly impractical. It may be impossible, for example, to simulate at the packet level a network slated to transport packets at gigabit-per-second rates. If, on the other hand, we are to resort to analytical techniques from classical queuing theory, we find that traditional traffic models, largely based on Poisson processes, need to be replaced by more sophisticated stochastic processes that capture the bursty nature of realistic traffic; in addition, we need to explicitly model buffer overflow phenomena which typically defy tractable analytical derivations.

An alternative modeling paradigm, based on stochastic fluid models (SFMs), has been recently considered for the purpose of analysis and simulation. Introduced in [1] and later proposed in [2] for the analysis of multiplexed data streams and network performance [3], SFMs have been shown to be especially useful for simulating various kinds of high-speed networks [4]–[10]. The fluid-flow worldview can provide either approximations to queuing-based models or primary models in their own right. In any event, its justification rests on a molecular view of packets in moderate-to-heavy loads over high-speed transmission links, where the effect of an individual packet or cell on the entire traffic process is virtually infinitesimal, not unlike the effect of a water molecule on the water flow in a river.

The efficacy of a SFM rests on its ability to aggregate multiple events. For example, a discrete event simulation run of an ATM link operating at 622 Mb/s may have to process over a million events per second. On the other hand, if traffic arrives from the source at rates that are piecewise-constant functions of time, then a simulation run would process only one event per rate change. Thus, 30 rate changes per second (as in certain video encoders) may require the processing of only 30 events per second. In effect, the SFM paradigm allows the aggregation of multiple events, associated with the movement of individual packets/cells over a time period of a constant flow rate, into a single event associated with a rate change. It foregoes the identity and dynamics of individual packets and focuses instead on the aggregate flow rate.

For the purpose of performance analysis with quality of service (QoS) requirements, the accuracy of SFMs depends on traffic conditions, the structure of the underlying system, and the nature of the performance metrics of interest. By foregoing the identity of individual packets, the SFM paradigm is more suitable for network-related measures, such as buffer levels and packet loss volumes, rather than packet-related measures such as sojourn times (although it is still possible to define fluid-based sojourn times [11]). A QoS metric that depends on the identity of certain packets, for example, cannot be obviously captured by a fluid model. Moreover, some metrics may depend on higher-order statistics of the distributions of the underlying random variables involved, which a fluid model may not be able to accurately capture.

In this paper, our goal is to explore the use of SFMs for the purpose of *control and optimization* rather than *performance analysis*. In this case, it is reasonable to expect that the solution of an optimization problem can be identified through a model which captures only those features of the underlying "real" system that are needed to lead to the right solution, even though the corresponding optimal performance may not be accurately estimated. Even if the exact solution cannot be obtained by such "lower-resolution" models, one can still obtain near-optimal

Manuscript received June 13, 2001; revised October 26, 2001. The work of C. G. Cassandras, G. Sun, and C. G. Panayiotou was supported in part by the National Science Foundation under Grants EEC-95-27422 and ACI-98-73339, in part by the Air Force Office of Scientific Research under Contract F49620-98-1-0387, in part by the Air Force Research Laboratory under Contract F30602-99-C-0057, and in part by EPRI/ARO under Contract WO8333-03. The work of Y. Wardi and B. Melamed was supported in part by the National Science Foundation under Grant DMI-0085659 and in part by DARPA under Contract F30602-00-2-0556.

C. G. Cassandras, G. Sun, and C. G. Panayiotou are with the Department of Manufacturing Engineering, Boston University, Brookline, MA 02446 USA (e-mail: cgc@bu.edu; gsun@bu.edu; panayiot@bu.edu).

Y. Wardi is with the School of Electrical Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: wardi@ee.gatech.edu).

B. Melamed is with the Department of Management Science and Information Systems, Rutgers University-Newark and New Brunswick, Piscataway, NJ 08854-8054 USA (e-mail: melamed@rbs.rutgers.edu).

Publisher Item Identifier 10.1109/TAC.2002.800739.

points that exhibit robustness with respect to certain aspects of the model they are based on. Such observations have been made in several contexts (e.g., [12]), including recent results related to SFMs reported in [13] and [21] where a connection between the SFM and queuing-system-based solution is established for various optimization problems in queuing systems.

With this in mind, we consider here optimization problems for single-node SFMs involving loss volume and workload levels; both are network-related performance metrics associated with buffer control or call-admission control. In a typical buffer control problem, for instance, the optimization problem involves the determination of a threshold (measured in packets or bytes) that minimizes a weighted sum of loss volume and buffer content. As the motivating example presented in Section II illustrates, a solution of this problem based on a SFM gives a close approximation to the solution of the associated queuing model. Since solving such problems usually relies on gradient information, estimating the gradient of a given cost function with respect to the aforementioned threshold parameters in a SFM becomes an essential task. Perturbation analysis (PA) methods [14], [15] are therefore suitable, if appropriately adapted to a SFM viewed as a discrete-event system. Liu and Gong [16], for example, have used PA to analyze an infinite-capacity SFM, with incoming traffic rates as the parameters of interest. In this paper, we show that infinitesimal perturbation analysis (IPA) yields remarkably simple sensitivity estimators for packet loss and workload metrics with respect to threshold or buffer size parameters. These estimators also turn out to be *nonparametric* in the sense that they are computable from data directly observable along a sample path, requiring no knowledge of the underlying probability law, including distributions of the random processes involved, or even parameters such as traffic or processing rates. In addition, the estimators obtained are unbiased under very weak structural assumptions on the defining traffic processes. Therefore, the IPA gradient estimators that we derive can be readily used for online control purposes to perform periodic network management functions in order to guarantee negotiated QoS parameters and to improve performance. For instance, a network can monitor its relative loss rate and mean buffer contents for a period of time, and then adjust admission parameters, provision transmission capacities, or reassign threshold levels in order to improve performance. Such management functions have not been standardized, and typically are performed in ad-hoc ways by monitoring performance levels. Aside from solving explicit optimization problems, IPA gradient estimators simplify the implementation of sensitivity analysis.

The contributions of this paper are as follows. First, we consider a single-node SFM and derive IPA gradient estimators for performance metrics related to loss and workload levels with respect to threshold parameters (equivalently, buffer sizes). One can derive such estimators by either 1) considering the finite difference of a performance metric as a function of the finite difference of a parameter and then use explicit limit arguments to obtain an unbiased estimate of the performance metric derivative, or 2) deriving the sample derivative for the performance metric involved, and then proving that it indeed yields an unbiased estimate. The former approach provides clear insights into the dynamic process of generation and propagation of perturbations, which is very helpful in understanding how to extend the approach to multiple fluid classes and multiple nodes. In addition, it requires no technical conditions on the traffic processes or the sample functions involved. However, this approach is tedious, even for a simple single-node model. The latter approach is simpler and more elegant, at the expense of some mild technical conditions needed to justify the evaluation of the sample derivative. It requires, however, some results from the first approach in order to prove unbiasedness of the derived estimators. Thus, in this paper, we start with the former approach, and then show that the estimators derived are equivalent to the latter, which we subsequently adopt. Based on these estimators, we also present simple algorithms for implementing them on line, taking advantage of their nonparametric nature.

The second contribution of this paper is to make use of the IPA gradient estimators derived to tackle buffer control as an optimization problem. In particular, we seek to determine the threshold value that minimizes a given performance metric. Packet-by-packet buffer control can be applied after the session admission decision is made in order to dynamically adjust network resources so as to minimize some cost based on the promised QoS. We use a standard gradient-based stochastic optimization scheme, where we estimate the gradient of the performance function with respect to the threshold parameter on the SFM; however, due to the simplicity of this gradient estimator, we evaluate it *based on data observed on a sample path of the actual (discrete-event) system*. Thus, we use the SFM only to obtain a gradient estimator; the associated value at any operating point is obtained from real system data.

The paper is organized as follows. First, in Section II, we motivate our approach with a buffer control problem in the SFM setting and show the application of IPA to it. In Section III, we describe the detailed SFM setting and define the performance metrics and parameters of interest. In Section IV, we derive IPA estimators for the sensitivities of the expected loss rate and workload with respect to threshold parameters (equivalently, buffer sizes) and show their unbiasedness. This is first demonstrated by a direct approach based on finite differences. The IPA approach is then generalized by evaluating sample derivatives (at the expense of introducing some mild technical conditions); these are shown to provide unbiased performance derivative estimators which are of nonparametric nature. Algorithms for implementing the derivative estimators obtained are also provided. In Section V, we show how the SFM-based derivative estimates can be used on line using data from the *actual* system (not the SFM) in order to solve buffer control problems. Finally, in Section VI we outline a number of open problems and future research directions motivated by this work.

II. A Motivating Example: Threshold-Based Buffer Control

This section presents a motivating example of buffer control in the setting of both a queuing model and a SFM and then compares the two. Consider a network node where buffer control at the packet level takes place using a simple threshold-based policy: when a packet arrives and the queue length is below a

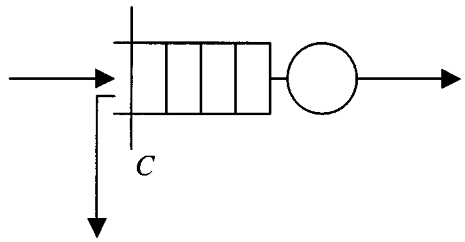

Fig. 1. Buffer control in a single node.

given amount C, it is accepted; otherwise, it is rejected. Let $\overline{L}(C)$ denote the expected loss rate, i.e., the expected rate of packet overflow at steady state, and let $\overline{Q}(C)$ denote the mean queue length when the threshold is C. We then define the cost function

$$J(C) = \overline{Q}(C) + R \cdot \overline{L}(C) \tag{1}$$

where R is a penalty associated with rejecting a packet. Thus, $J(C)$ captures the tradeoff between providing satisfactory service (low delay) and rejecting too many packets. Since, arguably, the notion of steady state is hard to justify in many networks, and since control decisions need to be made periodically or in response to apparent adverse network conditions, a more realistic performance measure is one where $\overline{L}(C)$ and $\overline{Q}(C)$ are replaced by $\overline{L}_T(C)$ and $\overline{Q}_T(C)$, the expected loss rate and mean queue length, respectively, over the time-interval $[0, T]$. We then consider

$$J_T(C) = \overline{Q}_T(C) + R \cdot \overline{L}_T(C) \tag{2}$$

to be the cost function of interest. Care must be taken in defining the previous expectations over a finite time-horizon, since they generally depend on the initial conditions; for the time being, we shall assume that the queue is empty at time $t = 0$, and revisit this point later. Fig. 1 depicts the queuing system under consideration.

The packet arrival process is modeled as an ON–OFF source so that packets arrive at a peak rate α during an ON period, followed by an OFF period during which no packets arrive. The packet processing rate is β. For the example used here and illustrated in Fig. 1, the number of arrivals in each ON period is geometrically distributed with parameter $p = 0.05$ and arrival rate $\alpha = 2$; the OFF period is exponentially distributed with parameter $\mu = 0.1$; and the service rate is $\beta = 1.01$. Thus, the traffic intensity of the system is $\alpha(1/\alpha p)/\beta((1/\alpha p) + (1/\mu)) = 0.99$, where $1/\alpha p$ is the average length of an ON period and $1/\mu$ is the average length of an OFF period. The cost function $J_T(C)$ in this problem is piecewise constant, hence gradient-based algorithms cannot be used. However, by exhaustively simulating this queuing system and averaging over 25 sample paths of length $T = 100\,000$ time units and estimating $J_T(C)$ over different discrete values of C, we obtained the curve labeled "DES" in Fig. 2, using a rejection penalty $R = 50$. One can see that the optimal threshold value in this example is $C^* = 15$.

Next, we adopt a simple SFM for the same system, treating packets as "fluid." During an ON period, the fluid volume in the buffer, $x(t)$, increases at rate $\alpha - \beta$ (we assume $\alpha > \beta$, otherwise there would be no buffer accumulation), while during an OFF period it decreases at a rate β. The cost function in this model is

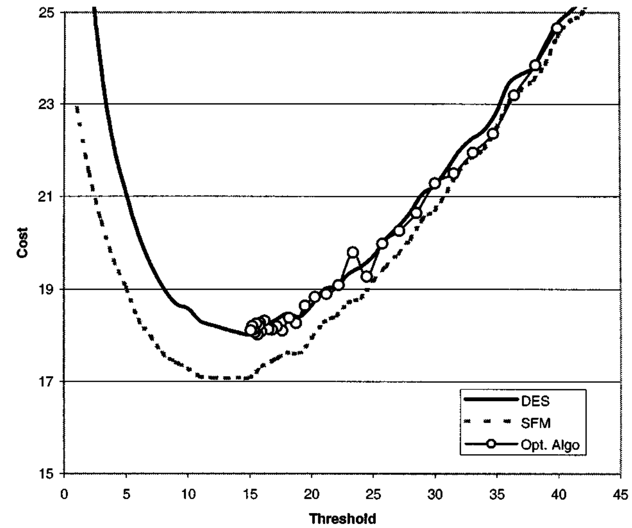

Fig. 2. Cost versus threshold comparison for DES and SFM.

$$J_T^{SFM}(\theta) = \overline{Q}_T^{SFM}(\theta) + R \cdot \overline{L}_T^{SFM}(\theta) \tag{3}$$

where $\theta \in \mathbb{R}^+$ is the threshold used to reject incoming fluid when the buffer fluid volume reaches level θ. The corresponding expected loss rate and mean buffer fluid volume over the time-interval $[0, T]$ are denoted by $\overline{L}_T^{SFM}(\theta)$ and $\overline{Q}_T^{SFM}(\theta)$, respectively. Simulating this model under the same ON–OFF conditions as before over many values of θ results in the curve labeled "SFM" in Fig. 2. The important observation is that the two optima are close, whereas the difference in the actual cost estimates can be substantial (especially for a lightly loaded system). In fact, $\theta^* = 13$ and $J_T^{SFM}(13) = 17.073$, as compared to $J_T(13) = 18.127$ and the optimal $J_T(C^*) = J_T(15) = 18.012$.

Based on this observation, we are motivated to study means for efficiently identifying solutions to problems formulated in a SFM setting. It is still difficult to obtain analytical solutions, however, since expressions for $\overline{Q}_T^{SFM}(\theta)$ and $\overline{L}_T^{SFM}(\theta)$ are unavailable, unless the arrival and service processes in the actual system are very simple. Therefore, one needs to resort to iterative methods such as stochastic approximation algorithms (e.g., [17]), which are driven by estimates of the gradient of a cost function with respect to the parameters of interest.

In the case of the aforementioned simple buffer control problem, we are interested in estimating $dJ_T/d\theta$ based on directly observed (simulated) data. We can then seek to obtain θ^* such that it minimizes $J_T(\theta)$ through an iterative scheme of the form

$$\theta_{n+1} = \theta_n - \nu_n H_n(\theta_n, \omega_n^{SFM}), \qquad n = 0, 1, \ldots \tag{4}$$

where $\{\nu_n\}$ is a step size sequence and $H_n(\theta_n, \omega_n^{SFM})$ is an estimate of $dJ_T/d\theta$ evaluated at $\theta = \theta_n$ and based on information obtained from a sample path of the SFM denoted by ω_n^{SFM}. However, as we will see, the simple form of $H_n(\theta_n, \omega_n^{SFM})$ to

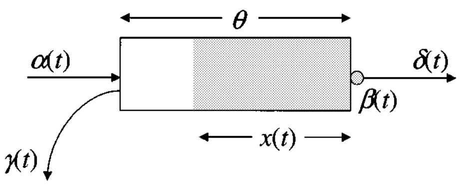

Fig. 3. The basic SFM.

be derived also enables us to apply the same scheme to the original discrete-event system

$$C_{n+1} = C_n - \nu_n H_n(C_n, \omega_n^{DES}), \qquad n = 0, 1, \ldots \tag{5}$$

where C_n is the threshold used for the nth iteration and ω_n^{DES} is a sample path of the discrete event system. In other words, analyzing the SFM provides us with the *structure* of a gradient estimator whose actual value can be obtained based on data from the actual system. In Fig. 2, the curve labeled "Opt.Algo." corresponds to this process and illustrates how one can indeed recover the optimal threshold $C^* = 15$.

III. SFM Setting

The SFM setting is based on the fluid-flow worldview, where "liquid molecules" flow in a continuous fashion. The basic SFM, used in [11] and shown in Fig. 3, consists of a single-server (spigot) preceded by a buffer (fluid storage tank), and it is characterized by five stochastic processes, all defined on a common probability space $(\Omega, \mathcal{F}, \mathcal{P})$ as follows:

- $\{\alpha(t)\}$: the input flow (inflow) rate to the SFM;
- $\{\beta(t)\}$: the service rate, i.e., the maximal fluid discharge rate from the server;
- $\{\delta(t)\}$: the output flow (outflow) rate from the SFM, i.e., the actual fluid discharge rate from the server;
- $\{x(t)\}$: the buffer occupancy or buffer content, i.e., the volume of fluid in the buffer;
- $\{\gamma(t)\}$: the overflow (spillover) rate due to excessive incoming fluid at a full buffer.

The previous processes evolve over a time interval $[0, T]$ for a given fixed $T > 0$. The inflow process $\{\alpha(t)\}$ and the service-rate process $\{\beta(t)\}$ are assumed to be right-continuous piecewise constant, with $0 \le \alpha_{\min} \le \alpha(t) \le \alpha_{\max} < \infty$ and $0 \le \beta_{\min} \le \beta(t) \le \beta_{\max} < \infty$. Let θ denote the size of the buffer, which is the variable parameter we will concentrate on for the purpose of IPA. The processes $\{\alpha(t)\}$ and $\{\beta(t)\}$, along with the buffer size θ, define the behavior of the SFM. In particular, they determine the buffer content, $x(\theta; t)$, the overflow rate $\gamma(\theta; t)$, and the output flow $\delta(\theta; t)$. The notational dependence on θ indicates that we will analyze performance metrics as functions of the given θ. We will assume that the real-valued parameter θ is confined to a closed and bounded (compact) interval Θ; to avoid unnecessary technical complications, we assume that $\theta > 0$ for all $\theta \in \Theta$.

The buffer content $x(\theta; t)$ is determined by the following one-sided differential equation:

$$\frac{dx(\theta;t)}{dt^+} = \begin{cases} 0, & \text{if } x(\theta;t)=0 \text{ and } \alpha(t)-\beta(t)\le 0 \\ 0, & \text{if } x(\theta;t)=\theta \text{ and } \alpha(t)-\beta(t)\ge 0 \\ \alpha(t)-\beta(t), & \text{otherwise} \end{cases} \tag{6}$$

with the initial condition $x(\theta; 0) = x_0$ for some given x_0; for simplicity, we set $x_0 = 0$ throughout the paper. The outflow rate $\delta(\theta; t)$ is given by

$$\delta(\theta;t) = \begin{cases} \beta(t), & \text{if } x(\theta;t) > 0 \\ \alpha(t), & \text{if } x(\theta;t) = 0 \end{cases} \tag{7}$$

where we point out that if we allow $\theta = 0$, then $\delta(\theta; t) = \min\{\alpha(t), \beta(t)\}$. The overflow rate $\gamma(\theta; t)$ is given by

$$\gamma(\theta;t) = \begin{cases} \max\{\alpha(t) - \beta(t), 0\}, & \text{if } x(\theta;t) = \theta \\ 0, & \text{if } x(\theta;t) < \theta. \end{cases} \tag{8}$$

This SFM can be viewed as a dynamic system whose input consists of the two *defining* processes $\{\alpha(t)\}$ and $\{\beta(t)\}$ along with the buffer size θ, its state is comprised of the buffer content process, and its output includes the outflow and overflow processes. The state and output processes are referred to as *derived* processes, since they are determined by the defining processes. Since the input sample functions (realizations) of $\{\alpha(t)\}$ and $\{\beta(t)\}$ are piecewise constant and right-continuous, the state trajectory $x(\theta; t)$ is piecewise linear and continuous in t, and the output function $\gamma(\theta; t)$ is piecewise constant. Moreover, the state trajectory can be decomposed into two kinds of intervals: *empty periods* and *busy periods*. Empty periods (EPs) are maximal intervals during which the buffer is empty, while busy periods (BPs) are supremal intervals during which the buffer is nonempty. Observe that during an EP the system is not necessarily idle since the server may be active; see (7). Note also that since $x(\theta; t)$ is continuous in t, EPs are always closed intervals, whereas BPs are open intervals unless containing one of the end points 0 or T. The outflow process $\{\delta(t)\}$ becomes important in modeling networks of SFMs and it will not concern us any further here, since our interest in this paper lies in single-node systems.

Let $\mathcal{L}(\theta)$: $\Theta \to \mathbb{R}$ be a random function defined over the underlying probability space $(\Omega, \mathcal{F}, P)$. Strictly speaking, we write $\mathcal{L}(\theta, \omega)$ to indicate that this sample function depends on the sample point $\omega \in \Omega$, but will suppress ω unless it is necessary to stress this fact. In what follows, we will consider two performance metrics, the *Loss Volume* $L_T(\theta)$ and the *Cumulative Workload* (or just *Work*) $Q_T(\theta)$, both defined on the interval $[0, T]$ via the following equations:

$$L_T(\theta) = \int_0^T \gamma(\theta; t)\, dt \tag{9}$$

$$Q_T(\theta) = \int_0^T x(\theta; t)\, dt \tag{10}$$

where, as already mentioned, we assume that $x(\theta; 0) = 0$. Observe that $(1/T)E[L_T(\theta)]$ is the *expected loss rate* over the interval $[0, T]$, a common performance metric of interest (from which related metrics such as *loss probability* can also be derived). Similarly, $(1/T)E[Q_T(\theta)]$ is the *expected buffer content* over $[0, T]$. We may then formulate optimization problems such as the determination of θ^* that minimizes a cost function of the form

$$J(\theta) = \frac{1}{T} E[Q_T(\theta)] + \frac{R}{T} E[L_T(\theta)] \equiv \frac{1}{T} J_Q(\theta) + \frac{R}{T} J_L(\theta)$$

where R represents a rejection cost due to overflow. In order to accomplish this task, we rely on estimates of $dJ_L(\theta)/d\theta$ and $dJ_Q(\theta)/d\theta$ provided by the sample derivatives $dL_T(\theta)/d\theta$ and $dQ_T(\theta)/d\theta$ for use in stochastic gradient-based schemes. Accordingly, the objective of the next section is the estimation of the derivatives of $J_L(\theta)$ and $J_Q(\theta)$, which we will pursue through infinitesimal perturbation analysis (IPA) techniques [14], [15]. Henceforth, we shall use the "prime" notation to denote derivatives with respect to θ, and will proceed to estimate the derivatives $J'_L(\theta)$ and $J'_Q(\theta)$. The corresponding sample derivatives are denoted by $L'_T(\theta)$ and $Q'_T(\theta)$, respectively.

IV. IPA With Respect to Buffer Size or Threshold

As already mentioned, we will concentrate on the buffer size θ in the SFM previously described or, equivalently, a threshold parameter used for buffer control. We assume that the processes $\{\alpha(t)\}$ and $\{\beta(t)\}$ are independent of θ and of the buffer content. Thus, we consider network settings operating with protocols such as ATM and UDP, but not TCP. Our objective is to estimate the derivatives $J'_L(\theta)$ and $J'_Q(\theta)$ through the sample derivatives $L'_T(\theta)$ and $Q'_T(\theta)$ which are commonly referred to as infinitesimal perturbation analysis (IPA) estimators; comprehensive discussions of IPA and its applications can be found in [14] and [15]. The IPA derivative-estimation technique computes $L'_T(\theta)$ and $Q'_T(\theta)$ along an observed sample path ω. An IPA-based estimate $\mathcal{L}'(\theta)$ of a performance metric derivative $dE[\mathcal{L}(\theta)]/d\theta$ is *unbiased* if $dE[\mathcal{L}(\theta)]/d\theta = E[\mathcal{L}'(\theta)]$. Unbiasedness is the principal condition for making the application of IPA practical, since it enables the use of the sample (IPA) derivative in control and optimization methods that employ stochastic gradient-based techniques.

We consider sample paths of the SFM over $[0, T]$. For a fixed $\theta \in \Theta$, the interval $[0, T]$ is divided into alternating EPs and BPs. Suppose there are K busy periods denoted by $\mathcal{B}_k$, $k = 1, \ldots, K$, in increasing order. Then, by (9) and (10), the sample performance functions assume the following form:

$$L_T(\theta) = \sum_{k=1}^{K} \int_{\mathcal{B}_k} \gamma(\theta; t)\, dt \tag{11}$$

$$Q_T(\theta) = \sum_{k=1}^{K} \int_{\mathcal{B}_k} x(\theta; t)\, dt. \tag{12}$$

As mentioned earlier, the processes $\{\alpha(t)\}$ and $\{\beta(t)\}$ are assumed piecewise constant. This implies that, w.p. 1, there exist a random integer $N(T) > 0$ and an increasing sequence of time points $0 = t_0 < t_1 < \cdots < t_{N(T)} < t_{N(T)+1} = T$, generally dependent upon the sample path ω, such that t_i is a jump (discontinuity) point of $\alpha(t) - \beta(t)$; clearly, $\alpha(t) - \beta(t)$ is continuous at all points other than $t_0, \ldots, t_{N(T)}$. We will assume that $N(T)$ has a finite expectation, i.e., $E[N(T)] < \infty$.

Viewed as a discrete-event system, an *event* in a sample path of the SFM may be either *exogenous* or *endogenous*. An exogenous event is a jump in either $\{\alpha(t)\}$ or $\{\beta(t)\}$. An endogenous event is defined to occur when the buffer becomes full or empty. We note that the times at which the buffer *ceases to be* full or empty are locally independent of θ, because they correspond to a change of sign in the difference function $\alpha(t) - \beta(t)$ (by a random function $f(\theta)$ being "locally independent" of θ we mean that for a given θ there exists $\Delta\theta > 0$ such that for every $\overline{\theta} \in (\theta - \Delta\theta, \theta + \Delta\theta)$, w.p. 1 $f(\overline{\theta}) = f(\theta)$, where $\Delta\theta$ may depend on both θ and on the sample path). Thus, given a BP $\mathcal{B}_k$, its starting point is one where the buffer ceases to be empty and is therefore locally independent of θ, while its end point generally depends on θ. Denoting these points by ξ_k and $\eta_k(\theta)$ we express $\mathcal{B}_k$ as

$$\mathcal{B}_k = (\xi_k, \eta_k(\theta)), \qquad k = 1, \ldots, K$$

for some random integer K. The BPs can be classified according to whether some overflow occurs during them or not. Thus, we define the random set

$$\Phi(\theta) := \{k \in \{1, \ldots, K\}: x(t) = \theta,\ \alpha(t) - \beta(t) > 0 \text{ for some } t \in (\xi_k, \eta_k(\theta))\}.$$

For every $k \in \Phi(\theta)$, there is a (random) number $M_k \geq 1$ of *overflow periods* in $\mathcal{B}_k$, i.e., intervals during which the buffer is full and $\alpha(t) - \beta(t) > 0$. Let us denote these overflow periods by $\mathcal{F}_{k,m}$, $m = 1, \ldots, M_k$, in increasing order and express them as $\mathcal{F}_{k,m} = [u_{k,m}(\theta), v_{k,m}]$, $k = 1, \ldots, K$. Observe that the starting time $u_{k,m}(\theta)$ generally depends on θ, whereas the ending time $v_{k,m}$ is locally independent of θ, since it corresponds to a change of sign in the difference function $\alpha(t) - \beta(t)$, which has been assumed independent of θ. Finally let

$$B(\theta) = |\Phi(\theta)| \tag{13}$$

where $|\cdot|$ denotes the cardinality of a set, i.e., $B(\theta)$ is the number of BPs in $[0, T]$ during which some overflow is observed. To summarize, the following points hold true.

- There are K busy periods in $[0, T]$, with $\mathcal{B}_k = (\xi_k, \eta_k(\theta))$, $k = 1, \ldots, K$.
- $k \in \Phi(\theta)$ iff some overflow occurs during $\mathcal{B}_k$; we set $B(\theta) = |\Phi(\theta)|$.
- For each $k \in \Phi(\theta)$, there are M_k overflow periods in $\mathcal{B}_k$, i.e., $\mathcal{F}_{k,m} = [u_{k,m}(\theta), v_{k,m}]$, $m = 1, \ldots, M_k$.

A typical sample path is shown in Fig. 4, where $K = 3$, $\Phi = \{1, 3\}$, $M_1 = 2$, $M_2 = 0$, and $M_3 = 1$.

As mentioned in Section I, we present two ways of deriving IPA estimators: 1) by evaluating the finite differences $\Delta L_T(\theta)$ and $\Delta Q_T(\theta)$ as functions of $\Delta\theta$, obtaining left and/or right sample derivatives (depending on whether $\Delta\theta < 0$ or $\Delta\theta > 0$), taking limits as $\Delta\theta \to 0$, and finally exploring if they yield unbiased estimates of $J'_L(\theta)$ and $J'_Q(\theta)$; or 2) by explicitly evaluating $L'_T(\theta)$ and $Q'_T(\theta)$, which requires some additional technical assumptions. We will first proceed with the former approach and consider only the loss volume metric $L_T(\theta)$; the analysis for $Q_T(\theta)$ is similar, though a bit more involved (see also [18]). In pursuing this approach, we will also derive some results that will be used to establish the unbiasedness of the estimators $L'_T(\theta)$ and $Q'_T(\theta)$ obtained through the latter approach.

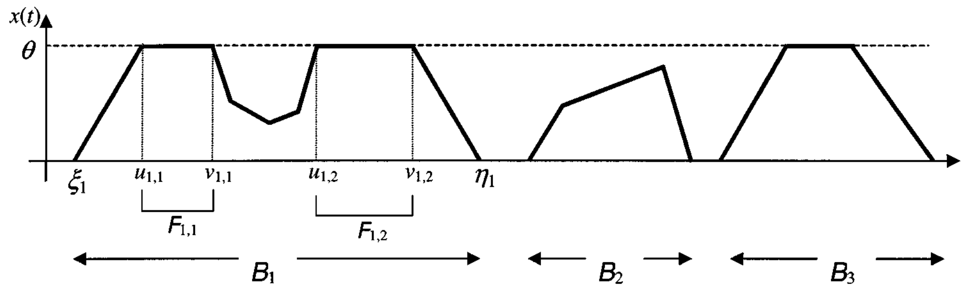

Fig. 4. A typical sample path of a SFM.

A. IPA Using Finite Difference Analysis

The stochastic component of the SFM manifests itself in the duration of the intervals defined by exogenous event occurrences corresponding to jumps in either $\alpha(t)$ or $\beta(t)$. Let $\{A_i\}$, $i = 1, 2, \ldots$, be the point process defined by these exogenous event times. For convenience, let α_i and β_i denote the (constant) inflow rate and service rate, respectively, over the interval $[A_i, A_{i+1})$. Note that we do not impose any restrictions on the probability law of the intervals defined by these events.

The main result of this section is to show that the sample derivative $L'_T(\theta)$, i.e., the sensitivity of the loss volume with respect to θ, is given by $-B(\theta)$, and that this is an unbiased estimator of $J'_L(\theta)$. Recall that $B(\theta)$ is simply the count of busy periods in which at least one overflow period is observed. Moreover, this remarkably simple estimator is independent of any assumptions on the traffic process or service process, as well as of the rates involved and even θ, i.e., it is *nonparametric*.

The starting point in IPA is to consider a *nominal* sample path under some buffer size (equivalently, admission threshold) θ and a *perturbed* sample path resulting from perturbing θ by $\Delta\theta$, while keeping the realizations of the processes $\{\alpha(t)\}$ and $\{\beta(t)\}$ unchanged, hence, leaving $\{A_i\}$, $i = 1, 2, \ldots$, unchanged. For simplicity, we limit ourselves to the case where $\Delta\theta > 0$, leading to an estimate of the right sample derivative of $L_T(\theta)$; the case where $\Delta\theta < 0$ is similar, leading to an estimate of the left sample derivative of $L_T(\theta)$. We then define

$$\Delta x_i(\theta, \Delta\theta) = x_i(\theta + \Delta\theta) - x_i(\theta)$$

where $x_i(\theta)$ denotes the nominal sample buffer content at time A_i and $x_i(\theta + \Delta\theta)$ denotes the perturbed sample buffer content at the same time. Similarly, we define perturbations for some additional sample path quantities as follows. First, setting $A_0 = 0$, let

$$L_i(\theta) = \int_{A_{i-1}}^{A_i} \gamma(\theta; t)\, dt, \qquad i = 1, 2, \ldots \tag{14}$$

be the total loss volume observed over an interevent interval $[A_{i-1}, A_i)$, and define

$$\Delta L_i(\theta, \Delta\theta) = L_i(\theta + \Delta\theta) - L_i(\theta). \tag{15}$$

In addition, let

$$y_{i+1}(\theta) = x_i(\theta) + (\alpha_i - \beta_i)[A_{i+1} - A_i] \tag{16}$$

and note that $(\alpha_i - \beta_i)[A_{i+1} - A_i]$ is simply the amount of change in the buffer content from time A_i to time A_{i+1}. Therefore, $y_{i+1}(\theta)$ is the queue content obtained at time A_{i+1} if the queue were allowed to become negative or to exceed θ. We may then define

$$\Delta y_i(\theta, \Delta\theta) = y_i(\theta + \Delta\theta) - y_i(\theta).$$

Finally, we define a perturbation in the ending time of a BP as

$$\Delta\eta_k(\theta, \Delta\theta) = \eta_k(\theta + \Delta\theta) - \eta_k(\theta), \qquad k = 1, 2, \ldots.$$

For notational simplicity, we shall henceforth suppress the arguments of all quantities Δx_i, Δy_i, ΔL_i, $\Delta\eta_k$.

Consider a typical BP $\mathcal{B}_k$, and all possible events that can take place in it, so as to determine how associated perturbations are either generated (due to $\Delta\theta$) or propagated from the previous event. The kth busy period is initiated by an exogenous event at time $\xi_k = A_i$, for some i, such that $\alpha_i - \beta_i > 0$, and let us assume that $\Delta x_i = 0$. Regarding the next exogenous event at time A_{i+1} there are two possible cases to consider:

Case I: $y_{i+1}(\theta) \le \theta$. In this case, $y_{i+1}(\theta)$ is given by (16) and we have [see also Fig. 5(a)]

$$x_{i+1}(\theta) = y_{i+1}(\theta)$$
$$L_i(\theta) = 0.$$

Clearly, $\Delta x_{i+1} = \Delta y_{i+1} = \Delta L_{i+1} = 0$.

Case II: $y_{i+1}(\theta) > \theta$. In this case, the queue content in the perturbed path can increase beyond θ up to the perturbed value $\theta + \Delta\theta$. Then, as also seen in Fig. 5(b)

$$\Delta x_{i+1} = \Delta\theta \tag{17}$$

$$\Delta L_{i+1} = -\Delta\theta \tag{18}$$

provided that $\Delta\theta$ is such that $0 < \Delta\theta \le y_{i+1}(\theta) - \theta$. To consider the case where $\Delta\theta > y_{i+1}(\theta) - \theta$, let the length of the overflow period in the nominal path be F_i and note that

$$F_i = \frac{y_{i+1} - \theta}{\alpha_i - \beta_i}.$$

Thus, if $\Delta\theta > y_{i+1}(\theta) - \theta = (\alpha_i - \beta_i)F_i$, then it is easy to see that the shaded area in Fig. 5(b) reduces to a triangle with area $(1/2)(\alpha_i - \beta_i)F_i^2$. We then get

$$\Delta x_{i+1} = (\alpha_i - \beta_i)F_i \tag{19}$$

$$\Delta L_{i+1} = -(\alpha_i - \beta_i)F_i. \tag{20}$$

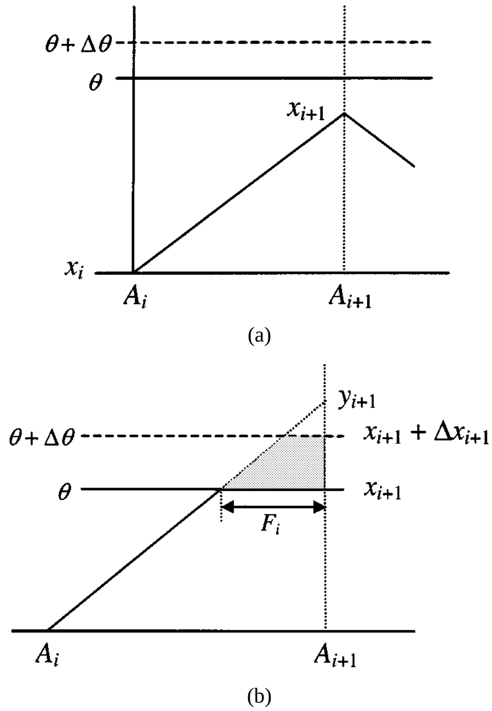

Fig. 5. (a) Case I: No perturbation generation ($y_{i+1}(\theta) \le \theta$). (b) Case II: Perturbation generation for $0 < \Delta\theta \le y_{i+1}(\theta) - \theta$.

Using the standard notation $[x]^+ = \max(x, 0)$, we can combine (17) and (18) with (19) and (20) to write

$$\Delta x_{i+1} = \Delta\theta - [\Delta\theta - (\alpha_i - \beta_i)F_i]^+ \tag{21}$$

$$\Delta L_{i+1} = -\Delta\theta + [\Delta\theta - (\alpha_i - \beta_i)F_i]^+. \tag{22}$$

Equations (21) and (22) capture the perturbation *generation* process due to $\Delta\theta$. The next step is to study how perturbations can be *propagated*, assuming the general situation $\Delta x_i \ge 0$. Doing so leads to the following result, which describes the complete queue content perturbation dynamics and establishes bounds for Δx_i.

Lemma 1: For all $i = 1, 2, \ldots$

$$0 \le \Delta x_i \le \Delta\theta \tag{23}$$

and

$$\Delta x_{i+1} = \begin{cases} [\Delta x_i - (\beta_i - \alpha_i)I_i]^+, & \text{if } \alpha_i - \beta_i < 0 \\ \Delta\theta - [\Delta\theta - \Delta x_i - (\alpha_i - \beta_i)F_i]^+, & \text{if } \alpha_i - \beta_i \ge 0 \end{cases} \tag{24}$$

where I_i is the length of an EP ending at A_{i+1} with $I_i = 0$ if no such period exists, and F_i is the length of an overflow period ending at A_{i+1} with $F_i = 0$ if no such period exists.

Proof: See the Appendix.

An immediate consequence of Lemma 1 is that a queue content perturbation may propagate across busy periods depending on the length of the EP separating these busy periods. This is because $\Delta x_{i+1} = [\Delta x_i - (\beta_i - \alpha_i)I_i]^+ \ge 0$ when an event occurs at time A_{i+1} that ends an EP of length I_i. Moreover, recalling that the endpoints of busy periods are denoted by $\eta_k(\theta)$, $k = 1, 2, \ldots$, the perturbation in $\eta_k(\theta)$ can be easily obtained by noticing in Fig. 8(a) [Case 1.2) in the proof of Lemma 1] that

$$\Delta\eta_k(\theta) = \frac{\Delta\theta}{\beta_i - \alpha_i} \tag{25}$$

provided that $\Delta\theta \le (\beta_i - \alpha_i)I_i$, where α_i and β_i are the inflow rate and service rate at the time the BP ends. To account for the fact that the kth BP may contain an overflow interval of length F_i with $\Delta\theta > (\alpha_i - \beta_i)F_i + \Delta x_i$, $\Delta\theta$ in (25) can be replaced by $\Delta\theta - [\Delta\theta - \Delta x_i - (\alpha_i - \beta_i)F_i]^+ \le \Delta\theta$ in view of (24). If, on the other hand, $\Delta\theta > (\beta_i - \alpha_i)I_i$, then the kth and $(k+1)$th busy periods are merged, which implies that $\Delta\eta_k(\theta)$ includes the entire length of the $(k+1)$th busy period.

Next, we identify bounds for ΔL_i (a generalization of the bounds for Δx_i and ΔL_i can also be found in [11]).

Lemma 2: For all $i = 1, 2, \ldots$

$$-\Delta\theta \le \Delta L_i \le 0. \tag{26}$$

Proof: See the Appendix.

Recall that if at least one overflow period is observed in the kth BP, then $k \in \Phi(\theta)$. Making use of the standard indicator function $\mathbf{1}[k \in \Phi(\theta)] = 1$ if $k \in \Phi(\theta)$ and zero otherwise, we have the following result, which allows us to characterize the cumulative loss perturbation at the end of a BP, which we will denote by $\Lambda_k(\Delta\theta)$, $k = 1, \ldots, K$.

Lemma 3: Consider a BP $\mathcal{B}_k = (\xi_k, \eta_k(\theta))$ with $\xi_k = A_j$, $\Delta x_j = 0$, and $A_m < \eta_k(\theta) \le A_{m+1}$. Assuming $\Delta\theta - \Delta x_i - (\alpha_i - \beta_i)F_i \le 0$ for all $i = j, \ldots, m$, the cumulative loss perturbation at the end of this busy period is

$$\Lambda_k(\Delta\theta) = -\Delta\theta \mathbf{1}[k \in \Phi(\theta)], \qquad k = 1, \ldots, K. \tag{27}$$

Proof: See the Appendix.

In simple terms, the loss perturbation depends *only* on the presence of an overflow within the observed busy period and not their number. It is noteworthy that this perturbation does not explicitly depend on any values that $\alpha(t)$ or $\beta(t)$ may take or the nature of the stochastic processes involved. Considering Lemma 3, note that it allows us to analyze all busy periods separately and accumulate loss perturbations at the end of the sample path over all busy periods observed; this, however, is contingent on the fact that $\Delta x_i = 0$ when a BP starts with an exogenous event at A_i. On the other hand, we saw that a consequence of Lemma 1 is $\Delta x_{i+1} = [\Delta x_i - (\beta_i - \alpha_i)I_i]^+$ following an EP of length I_i, i.e., the buffer content perturbation may not be zero when a BP starts, depending on the length of the EP separating it from the preceding BP.

We can now derive an unbiased derivative estimate for our performance metric by establishing the following result.

Theorem 4: The (right) derivative of the expected loss, $E[L_T(u)]$, is given by

$$\frac{dE[L_T(\theta)]}{d\theta} = -E\left[\sum_{k=1}^{K} \mathbf{1}[k \in \Phi(\theta)]\right] = -E[B(\theta)] \tag{28}$$

where K is the (random) number of busy periods contained in $[0, T]$, including a possibly incomplete last busy period.

Proof: We have

$$\frac{dE[L_T(\theta)]}{d\theta} = \lim_{\Delta\theta\to 0}\frac{1}{\Delta\theta}E[\Delta L_T(\theta)] = \lim_{\Delta\theta\to 0}\frac{1}{\Delta\theta}E\left[\sum_{k=1}^{K}\Lambda_k(\Delta\theta)\right]$$

where $\Lambda_k(\Delta\theta) = -\Delta\theta\mathbf{1}[k \in \Phi(\theta)]$ from Lemma 3, provided $\Delta\theta - \Delta x_i - (\alpha_i - \beta_i)F_i \le 0$ for all $A_i \in [0, T]$. It follows that

$$\frac{dE[L_T(\theta)]}{d\theta} = -E\left[\sum_{k=1}^{K}\mathbf{1}[k \in \Phi(\theta)]\right] = -E[B(\theta)]$$

where we have used the definition in (13).

If $\Delta\theta - \Delta x_i - (\alpha_i - \beta_i)F_i > 0$ for some $A_i \in [0, T]$, then the only additional effect comes from $\Delta L_{i+1} = -(\alpha_i - \beta_i)F_i < 0$ in (53). Then, consider

$$E[-(\alpha_i - \beta_i)F_i | \Delta\theta - \Delta x_i > (\alpha_i - \beta_i)F_i] = \int_0^{\Delta\theta - \Delta x_i} -xf(x)\,dx$$

where $f(\cdot)$ is the conditional pdf of $(\alpha_i - \beta_i)F_i$ given $\Delta\theta - \Delta x_i > (\alpha_i - \beta_i)F_i$, and let $f(\cdot) \le c < \infty$. Recalling that $0 \le \Delta x_i \le \Delta\theta$ from Lemma 1, we get

$$\int_0^{\Delta\theta - \Delta x_i} -xf(x)\,dx \ge \int_0^{\Delta\theta} -xf(x)\,dx \ge \int_0^{\Delta\theta} -\Delta\theta f(x)\,dx \ge \Delta\theta\int_0^{\Delta\theta} -c\,dx = -c(\Delta\theta)^2$$

and it follows that

$$E[-(\alpha_i - \beta_i)F_i] \ge -c(\Delta\theta)^2. \tag{29}$$

The cumulative loss perturbation due to events such that $\Delta L_{i+1} = -(\alpha_i - \beta_i)F_i$ is bounded from below by

$$\sum_{i=1}^{N(T)} -(\alpha_i - \beta_i)F_i$$

where F_i is the length of an overflow interval after the ith exogenous event, with $F_i = 0$ if no such overflow interval is present, and $N(T)$ is the total number of exogenous events in $[0, T]$. This cumulative loss perturbation is also bounded from above by zero, since $\Delta L_i \le 0$ from Lemma 2. Using (29), we get, given some $N(T)$

$$\lim_{\Delta\theta\to 0}\frac{1}{\Delta\theta}E\left[\sum_{i=1}^{N(T)} -(\alpha_i - \beta_i)F_i\right] \ge \lim_{\Delta\theta\to 0}\frac{1}{\Delta\theta}\sum_{i=1}^{N(T)} -c(\Delta\theta)^2 = \lim_{\Delta\theta\to 0}[-c\Delta\theta N(T)]$$

and

$$\lim_{\Delta\theta\to 0}[-c\Delta\theta]E[N(T)] = 0$$

where, by assumption, $E[N(T)] < \infty$. This completes the proof. ∎

An immediate implication of this theorem is that $-B(\theta)$ is an unbiased estimator of $dE[L_T(\theta)]/d\theta$

$$\left[\frac{dE[L_T(\theta)]}{d\theta}\right]_{est} = -B(\theta). \tag{30}$$

This estimator is extremely simple to implement: (30) is merely a counter of all busy periods observed in $[0, T]$ in which at least one overflow takes place. Again, no knowledge of the traffic or processing rates is required, nor does (30) depend on the nature of the random processes involved.

Using the aforementioned finite-difference approach, it is also possible to derive an unbiased estimator for $dE[Q_T(\theta)]/d\theta$ (see [18]), but it is considerably more tedious; we will see how to derive the same estimator in the next section by simpler means. Finally, note that (28) was derived using $\Delta\theta > 0$; thus, the analysis has to be repeated for $\Delta\theta < 0$ in order to evaluate the left sample derivative, and, although this does not present any conceptual difficulties, it adds to the tediousness of the finite difference analysis we have pursued thus far.

B. IPA Using Sample Derivatives

In this section, we derive explicitly the sample derivatives $L'_T(\theta)$ and $Q'_T(\theta)$ of the loss volume and work, defined in (11) and (12), respectively. We then show that they provide unbiased estimators of the expected loss volume sensitivity $dE[L_T(\theta)]/d\theta$ and the expected work sensitivity $dE[Q_T(\theta)]/d\theta$.

Since we are concerned with the sample derivatives $L'_T(\theta)$ and $Q'_T(\theta)$, we have to identify conditions under which they exist. Observe that any endogenous event time (a time point when the buffer becomes full or empty) is generally a function of θ; see also (6). Denoting this point by $t(\theta)$, the derivative $t'(\theta)$ exists as long as $t(\theta)$ is not a jump point of the difference process $\{\alpha(t) - \beta(t)\}$. Recall that the times at which the buffer ceases to be full or empty are locally independent of θ, because they correspond to a change-of-sign of the difference sample function $\alpha(t) - \beta(t)$, which does not depend on θ. Excluding the possibility of the simultaneous occurrence of two events, the only situation preventing the existence of the sample derivatives $L'_T(\theta)$ and $Q'_T(\theta)$ involves an interval during which $x(t) = \theta$ and $\alpha(t) - \beta(t) = 0$, as seen in (8); in this case, the one-sided derivatives of $L_T(\theta)$and $Q_T(\theta)$ exist and can be obtained with the approach of the previous section. In order to keep the analysis simple, we will focus only on the differentiable case. Therefore, the analysis that follows rests on the following technical conditions.

Assumption 1:

a) W.p. 1, $\alpha(t) - \beta(t) \ne 0$.
b) For every $\theta \in \Theta$, w.p. 1, no two *events* may occur at the same time.

Remark: We stress the fact that the aforementioned conditions for ensuring the existence of the sample derivatives $L'_T(\theta)$ and $Q'_T(\theta)$ are very mild. Part b) is satisfied whenever the cumulative distribution functions (cdfs) (or conditional cdfs) characterizing the intervals between exogenous event occurrences are continuous. For example, in the simple case where $\beta(t) = \beta$ and $\alpha(t)$ can only take two values, 0 and $\alpha > \beta$, suppose that the inflow process switches from α to 0 after $\theta/(\alpha-\beta)$ time units w.p. 1. The buffer then becomes full exactly when an exogenous event occurs, and the loss volume sample function experiences a discontinuity w.p. 1. Such situations can only arise for a small-finite subset of Θ (for which one can still calculate either the left or right derivatives) and they are of limited practical consequence.

We next derive the IPA derivatives of $L_T(\theta)$ and $Q_T(\theta)$. Recall that $B(\theta)=|\Phi(\theta)|$, i.e., the number of BPs containing at least one overflow period.

Theorem 5: For every $\theta\in\Theta$

$$L_T'(\theta)=-B(\theta). \tag{31}$$

Proof: Recalling that $\mathcal{B}_k=(\xi_k,\eta_k(\theta))$, we have, from (11)

$$L_T(\theta)=\sum_{k\in\Phi(\theta)}\int_{\xi_k}^{\eta_k(\theta)}\gamma(\theta;t)\,dt \tag{32}$$

which after differentiation yields

$$L_T'(\theta)=\sum_{k\in\Phi(\theta)}\frac{d}{d\theta}\int_{\xi_k}^{\eta_k(\theta)}\gamma(\theta;t)\,dt. \tag{33}$$

Note that the derivative in (33) is taken along a sample path. The set $\Phi(\theta)$, though depending on θ, can be viewed as a constant for the purpose of taking the derivative. The reason is that, by virtue of Assumption 1b), it is locally independent of θ, similarly to the endogenous event times discussed in the first part of Section IV (i.e., for every fixed θ, w.p. 1 there exists $\Delta\theta>0$, such that, for every $\overline{\theta}\in[\theta-\Delta\theta,\,\theta+\Delta\theta]$, $\Phi(\overline{\theta})=\Phi(\theta)$; although this $\Delta\theta$ generally depends on the given sample path, our derivative is taken along a specific sample path, hence, (33) is justified).

Next, we focus on a particular $\mathcal{B}_k$with $k\in\Phi(\theta)$ and we shall suppress the index k to simplify the notation. Accordingly, the BP in question is denoted by $\mathcal{B}=(\xi,\eta(\theta))$, and there are $M\geq 1$ overflow periods in $\mathcal{B}$, denoted by $\mathcal{F}_m=[u_m(\theta),\,v_m]$, $m=1,\,\ldots,\,M$. A typical scenario is depicted in Fig. 4, where in the first BP we have $M=2$. The loss volume over $\mathcal{B}$ is given by the function

$$\lambda(\theta)=\int_{\xi}^{\eta(\theta)}\gamma(\theta;t)\,dt. \tag{34}$$

We next prove that

$$\lambda'(\theta)=-1 \tag{35}$$

from which (31) immediately follows in view of (32)–(34). From the definition of $\gamma(\theta;t)$ in (8), we can rewrite (34) as

$$\lambda(\theta)=\sum_{m=1}^{M}\int_{u_m(\theta)}^{v_m}[\alpha(t)-\beta(t)]\,dt. \tag{36}$$

Since the points $u_m(\theta)$, $m=1,\,\ldots,\,M$, and the jump points of $\alpha(t)-\beta(t)$ constitute *events*, and since w.p. 1 no two events can occur at the same time by Assumption 1b), the function $\alpha(t)-\beta(t)$ must be continuous w.p. 1 at the points $u_m(\theta)$, $m=1,\,\ldots,\,M$. Consequently, by taking derivatives with respect to θ in (36) we obtain

$$\lambda'(\theta)=-\sum_{m=1}^{M}[\alpha(u_m(\theta))-\beta(u_m(\theta))]u_m'(\theta). \tag{37}$$

Next, consider the individual terms in the previous sum (see also Fig. 4 for an illustration).

1) If $m=1$, then the buffer is neither full nor empty in the interval $(\xi,\,u_1(\theta))$. Since the buffer content evolves from $x(\xi)=0$ to $x(u_1(\theta))=\theta$, (6) implies

$$\int_{\xi}^{u_1(\theta)}[\alpha(t)-\beta(t)]\,dt=\theta$$

and, upon taking derivatives with respect to θ

$$[\alpha(u_1(\theta))-\beta(u_1(\theta))]u_1'(\theta)=1. \tag{38}$$

2) If $m>1$, then the buffer is neither full nor empty in the interval $(v_{m-1},\,u_m(\theta))$. Since $x(v_{m-1})=x(u_m(\theta))=\theta$ we obtain, by (6)

$$\int_{v_{m-1}}^{u_m(\theta)}[\alpha(t)-\beta(t)]\,dt=0$$

and upon differentiating with respect to θ

$$[\alpha(u_m(\theta))-\beta(u_m(\theta))]u_m'(\theta)=0. \tag{39}$$

Finally, (37)–(39) imply (35), which immediately implies (31) and the proof is complete. ■

Note that Theorem 5 is consistent with Theorem 4. However, Theorem 4 includes a direct proof of the unbiasedness of the estimator $-B(\theta)$, whereas the present approach requires a separate proof that the sample derivative $L_T'(\theta)=-B(\theta)$ is in fact unbiased. The unbiasedness of this IPA derivative will be proven later, after we establish the IPA derivative of the work $Q_T(\theta)$ defined in (12).

Theorem 6: For every $\theta\in\Theta$,

$$Q_T'(\theta)=\sum_{k\in\Phi(\theta)}[\eta_k(\theta)-u_{k,1}(\theta)]. \tag{40}$$

Proof: We focus on a particular BP $\mathcal{B}_k$ with $k\in\Phi(\theta)$, and again suppress the notational dependency on k for the sake of simplicity. Accordingly, consider a BP $\mathcal{B}_k=(\xi,\eta(\theta))$, and denote its overflow periods by $\mathcal{F}_m=[u_m(\theta),\,v_m]$, $m=1,\,\ldots,\,M$, for some $M\geq 1$ (e.g., $M=2$ in the first BP of Fig. 4). Define the function

$$q(\theta)=\int_{\xi}^{\eta(\theta)}x(\theta;t)\,dt. \tag{41}$$

It suffices to prove that

$$q'(\theta)=\eta(\theta)-u_1(\theta) \tag{42}$$

since this would immediately imply (40). Since $x(\theta;t)$ is continuous in t, taking the derivative with respect to θ in (41) and letting $x'(\theta;t)$ denote the partial derivative with respect to θ yields

$$\begin{aligned}q'(\theta)&=\int_{\xi}^{\eta(\theta)}x'(\theta;t)\,dt+x(\theta;\eta(\theta))\eta'(\theta)\\&=\int_{\xi}^{\eta(\theta)}x'(\theta;t)\,dt\end{aligned} \tag{43}$$

since the BP ends at $\eta(\theta)$, hence, $x(\theta;\eta(\theta)) = 0$. To evaluate this partial derivative (which exists at all t except $t = u_m$ and $t = v_m$) we consider all possible cases regarding the location of t in the BP $\mathcal{B}_k = (\xi, \eta(\theta))$ (see Fig. 4).

1) $t \in (\xi, u_1(\theta))$. In this case, the buffer is neither empty nor full in this interval. It follows, using (6), that

$$x(\theta;t) = \int_{\xi}^{t} [\alpha(\tau) - \beta(\tau)]\, d\tau.$$

Since the right-hand side is independent of θ, we have $x'(\theta;t) = 0$.

2) $t \in (u_m(\theta), v_m)$, $m = 1 \ldots, M$. Since $(u_m(\theta), v_m)$ is an overflow period, $x(\theta, t) = \theta$ in these intervals, hence, $x'(\theta;t) = 1$.

3) $t \in (v_m, u_{m+1}(\theta))$, $m = 1, \ldots, M-1$. Here, the buffer is neither empty nor full in the interval (v_m, t), while $x(\theta;v_m) = \theta$. It follows, using (6), that

$$x(\theta;t) = \theta + \int_{v_m}^{t} [\alpha(\tau) - \beta(\tau)]\, d\tau$$

and upon differentiating with respect to θ, we obtain $x'(\theta;t) = 1$.

4) $t \in (v_M, \eta(\theta))$. This case is identical to the previous one, yielding $x'(\theta;t) = 1$.

In summary, $x'(\theta;t) = 0$ for all $t \in (\xi, u_1(\theta))$ (Case 1), and $x'(\theta;t) = 1$ for all $t \in (u_1(\theta), \eta(\theta))$ (Cases 2–4). Therefore, it follows from (43) that (42) holds, implying (40) and completing the proof. ■

In simple terms, the contribution of a BP, $\mathcal{B}_k$, to the sample derivative $Q'_T(\theta)$ in (40) is the length of the interval defined by the first point at which the buffer becomes full and the end of the BP. Once again, as in (31), observe that the IPA derivative $Q'_T(\theta)$ is nonparametric, since it requires only the recording of times at which the buffer becomes full [i.e., $u_{k,1}(\theta)$] and empty [i.e., $\eta_k(\theta)$] for any $\mathcal{B}_k$ with $k \in \Phi(\theta)$. We also remark that the same IPA derivative can be obtained through the finite difference analysis of the previous section (see [18]), but with considerably more effort.

1) IPA Unbiasedness: We next prove the unbiasedness of the IPA derivatives $L'_T(\theta)$ and $Q'_T(\theta)$ previously obtained. Although we have already shown in (28) that $-B(\theta)$ is an unbiased estimate of $dE[L_T(\theta)]/d\theta$, we supply an alternative and greatly simplified proof based on the direct derivation of the IPA estimator in this section and on some of the results of the finite-difference analysis in Section IV-A. By a similar technique, we also supply a proof of the unbiasedness of the IPA estimator $Q'_T(\theta)$ in (40). These proofs, jointly with the sample-derivative technique for obtaining the estimators, suggest the possibility of extensive generalizations to the functional forms of $\alpha(t)$ and $\beta(t)$ (beyond piecewise constant), to be explored in a forthcoming paper (also, see [19] and [10]).

In general, the unbiasedness of an IPA derivative $\mathcal{L}'(\theta)$ has been shown to be ensured by the following two conditions (see [20, Lemma A2, p. 70]).

Condition 1: For every $\theta \in \Theta$, the sample derivative $\mathcal{L}'(\theta)$ exists w.p. 1.

Condition 2: W.p. 1, the random function $\mathcal{L}(\theta)$ is Lipschitz continuous throughout Θ, and the (generally random) Lipschitz constant has a finite first moment.

Consequently, establishing the unbiasedness of $L'_T(\theta)$ and $Q'_T(\theta)$ as estimators of $dE[L_T(\theta)]/d\theta$ and $dE[Q_T(\theta)]/d\theta$, respectively, reduces to verifying the Lipschitz continuity of $L_T(\theta)$ and $Q_T(\theta)$ with appropriate Lipschitz constants. Recall that $N(T)$ is the random number of all exogenous events in $[0, T]$ and that we have assumed $E[N(T)] < \infty$.

Theorem 7: Under Assumption 1

1) if $E[N(T)] < \infty$, then the IPA derivative $L'_T(\theta)$ is an unbiased estimator of $dE[L_T(\theta)]/d\theta$;
2) the IPA derivative $Q'_T(\theta)$ is an unbiased estimator of $dE[Q_T(\theta)]/d\theta$.

Proof: Under Assumption 1, Condition 1 holds for $L_T(\theta)$ and $Q_T(\theta)$. Therefore, it only remains to establish Condition 2.

First, consider $L_T(\theta)$. Recalling (14) and (15), we can write

$$\Delta L_T(\theta) = \sum_{i=1}^{N(T)} \Delta L_i(\theta)$$

by partitioning $[0, T]$ into intervals $[A_{i-1}, A_i)$ defined by successive exogenous events. Then, by Lemma 2, $-\Delta\theta \le \Delta L_i \le 0$, so that

$$|\Delta L_T(\theta)| \le N(T)|\Delta\theta|$$

i.e., $L_T(\theta)$ is Lipschitz continuous with constant $N(T)$. Since $E[N(T)] < \infty$, this establishes unbiasedness.

Consider next the sample function $Q_T(\theta)$, defined by (12) and fix θ and $\Delta\theta > 0$. By Lemma 1, $0 \le \Delta x_i \le \Delta\theta$, hence the difference $\Delta x(\theta, \Delta\theta;t) := x(\theta + \Delta\theta;t) - x(\theta;t)$ satisfies the inequalities

$$0 \le \Delta x(\theta, \Delta\theta;t) \le \Delta\theta.$$

Consequently, in view of (12)

$$|\Delta Q_T(\theta)| = \left|\int_0^T \Delta x(\theta, \Delta\theta;t)\, dt\right| \le T|\Delta\theta|$$

that is, $Q_T(\theta)$ is Lipschitz continuous with constant T. This completes the proof. ■

Remark: For the more commonly used performance metrics $(1/T)E[L_T(\theta)]$ (the expected loss rate over $[0, T]$) and $(1/T)E[Q_T(\theta)]$ (the expected buffer content over $[0, T]$), the Lipschitz constants in Theorem 7 become $N(T)/T$ and 1, respectively. As $T \to \infty$, the former quantity typically converges to the exogenous event rate.

V. Optimal Buffer Control Using SFM-Based IPA Estimators

As suggested in Section II and illustrated in Fig. 2, the solution to an optimization problem defined for an actual network node (i.e., a node that operates as a queuing system) may be accurately approximated by the solution to the same problem based on a SFM of the node. However, this may not be always the case. On the other hand, the simple form of the IPA estimators of the expected loss rate and expected buffer content obtained through (31) and (40) allows us to use data from the *actual* (real-world) system in order to estimate sensitivities that, in turn, may be used to solve an optimization problem of interest. In other words, the *form* of the IPA estimators is obtained by analyzing the system as a SFM, but the associated *values* are based on real data. In particular, an algorithm for implementing the estimators (31) and (40) is given as follows.

IPA Estimation Algorithm

- Initialize a counter $\mathcal{C} := 0$ and a cumulative timer $\mathcal{T} := 0$.
- Initialize $\tau := 0$.
- If an overflow event is observed at time t and $\tau = 0$:
  –Set $\tau := t$.
- If a busy period ends at time t and $\tau > 0$:
  –Set $\mathcal{C} := \mathcal{C} - 1$ and $\mathcal{T} := \mathcal{T} + (t - \tau)$
  –Reset $\tau := 0$.
- If $t = T$, and $\tau > 0$:
  –Set $\mathcal{C} := \mathcal{C} - 1$ and $\mathcal{T} := \mathcal{T} + (t - \tau)$.

The final values of $\mathcal{C}$ and $\mathcal{T}$ provide the IPA derivatives $L'_T(\theta)$ and $Q'_T(\theta)$ respectively. We remark that the "overflow" and "end of BP" events are readily observable during actual network operation. In addition, we point out once again that these estimates are independent of all underlying stochastic features, including traffic and processing rates. Finally, the algorithm is easily modified to apply to any interval $[T_1, T_2]$.

Let us now return to the buffer control problem presented in Section II, where the objective was to determine a threshold C that minimizes a cost function of the form

$$J_T(C) = \overline{Q}_T(C) + R \cdot \overline{L}_T(C)$$

trading off the expected loss rate with a rejection penalty R for the expected queue length. If a SFM is used instead, then the cost function of interest becomes

$$J(\theta) = \frac{1}{T} E[Q_T(\theta)] + \frac{R}{T} E[L_T(\theta)]$$

and the optimal threshold parameter, θ^*, may be determined through a standard stochastic approximation algorithm based on (4). The gradient estimator $H_n(\theta, \omega_n^{SFM})$ is the IPA estimator of $dJ/d\theta$ based on (31) and (40)

$$H_n(\theta, \omega_n^{SFM}) = \frac{1}{T} \sum_{k \in \Phi(\theta)} [\eta_k(\theta) - u_{k,1}(\theta)] - \frac{R}{T} B(\theta) \quad (44)$$

evaluated over a simulated sample path ω_n^{SFM} of length T, following which a control update is performed through (4) based on the value of $H_n(\theta, \omega_n^{SFM})$.

The interesting observation here is that the same estimator may be used in (5) as follows. If a packet arrives and is rejected, the time this occurs is recorded as τ in the aforementioned algorithm. At the end of the current busy period, the counter $\mathcal{C}$ and timer $\mathcal{T}$ are updated. Thus, the exact same expression as in the right-hand side of (44) can be used to update the threshold

$$C_{n+1} = C_n - \nu_n H_n(C_n, \omega_n^{DES}), \qquad n = 0, 1, \ldots. \quad (45)$$

Note that, after a control update, the state must be reset to zero, in accordance with our convention that all performance metrics are defined over an interval $[0, T]$ with an initially empty buffer. In the case of offline control, this simply amounts to simulating the system after resetting its state to 0. In the more interesting case of online control, we proceed as follows. Suppose that the nth iteration ends at time τ_n and the state is $x(C_n; \tau_n)$ [in general, $x(C_n; \tau_n) > 0$]. At this point, the threshold is updated and its new value is C_{n+1}. Let $\tau_n^0 > \tau_n$ be the next time that the buffer is empty, i.e., $x(C_{n+1}; \tau_n^0) = 0$. At this point, the $(n+1)$th iteration starts and the next gradient estimate is obtained over the interval $[\tau_n^0, \tau_n^0 + T]$, so that $\tau_{n+1} = \tau_n^0 + T$ and the process repeats. The implication is that over the interval $[\tau_n, \tau_n^0]$ no estimation is carried out while the controller waits for the system to be reset to its proper initial state; therefore, sample path information available over $[\tau_n, \tau_n^0]$ is effectively wasted as far as gradient estimation is concerned.

Fig. 6 depicts examples of the application of this scheme to a single-node SFM under six different parameter settings (scenarios), summarized in Table I. As in Fig. 2, "DES" denotes curves obtained by estimating $J_T(C)$ over different (discrete) values of C, "SFM" denotes curves obtained by estimating $J(\theta)$ over different values of θ, and "Opt.Algo." represents the optimization process (45), where we maintain real-valued thresholds throughout. The first three scenarios correspond to a high traffic intensity ρ compared to the remaining three. For each example, C^* is the optimal threshold obtained through exhaustive simulation. In all simulations, an ON–OFF traffic source is used with the number of arrivals in each ON period geometrically distributed with parameter p and arrival rate α; the OFF period is exponentially distributed with parameter μ; and the service rate is fixed at β. Thus, the traffic intensity of the system ρ is $\alpha(1/\alpha p)/\beta((1/\alpha p) + (1/\mu))$, where $1/\alpha p$ is the average length of an ON period and $1/\mu$ is the average length of an OFF period. The rejection cost is $R = 50$. For simplicity, ν_n in (45) is taken to be a constant $\nu_n = 5$. Finally, in all cases $T = 100\,000$. As seen in Fig. 6, the threshold value obtained through (45) using the SFM-based gradient estimator in (44) either recovers C^* or is close to it with a cost value extremely close to $J_T(C^*)$; since in some cases the cost function is nearly constant in the neighborhood of the optimum, it is difficult to determine the actual optimal threshold, but it is also practically unimportant since the cost is essentially the same. We have also implemented (45) with $H_n(C_n, \omega_n^{DES})$ estimated over shorter interval lengths $T = 10\,000$ and $T = 5000$, with virtually identical results. Looking at Fig. 6, it is worth observing that determining θ^* as an approximation to C^* through offline analysis of the SFM would also yield good approximations, further supporting the premise of this paper that SFMs provide an attractive modeling framework for control and optimization (not just performance analysis) of complex networks.

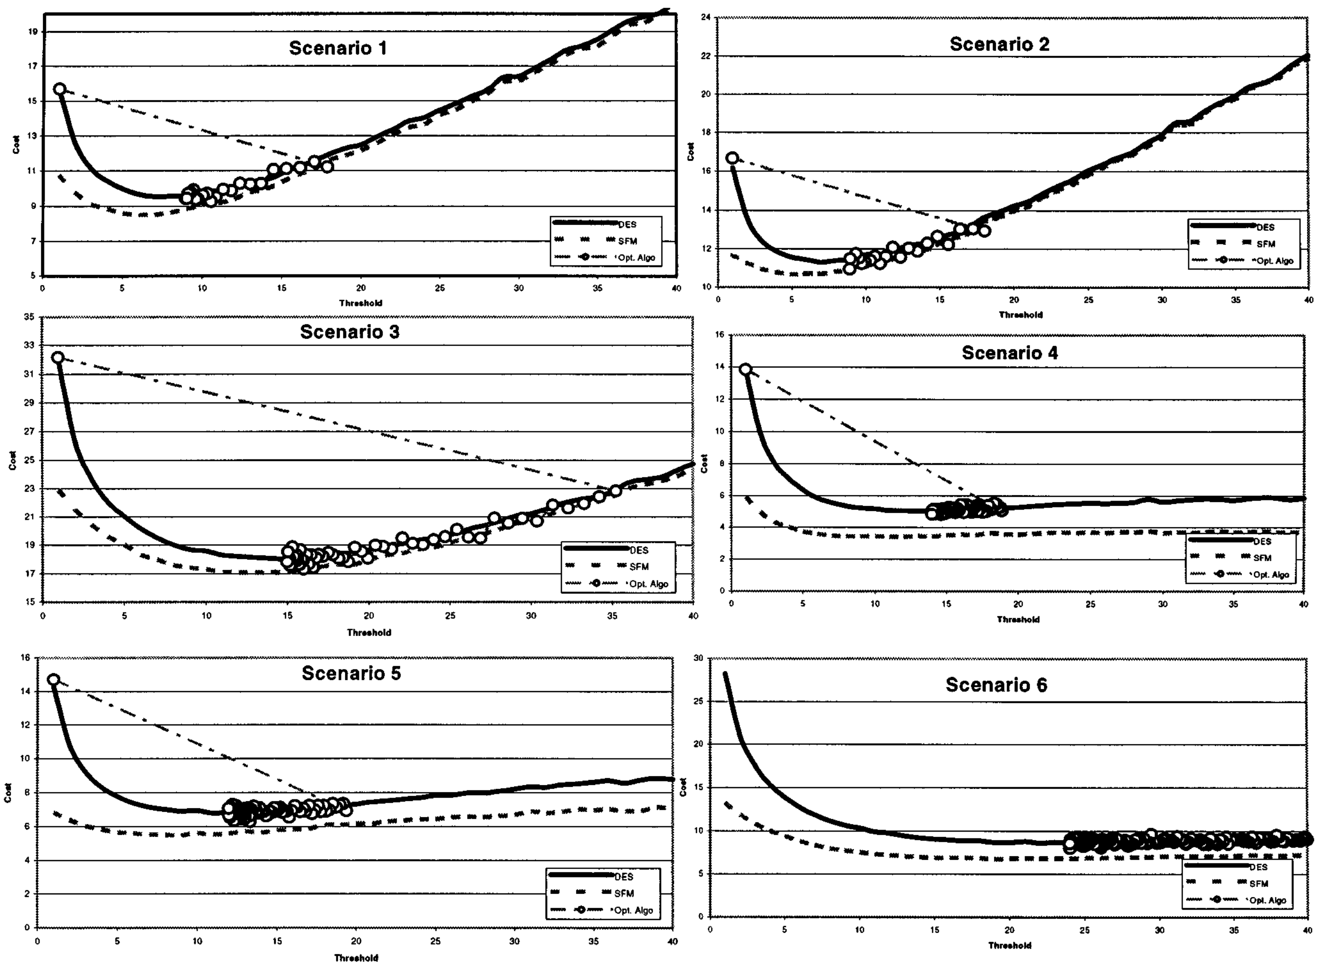

Fig. 6. Optimal threshold determination in an actual system using SFM-based gradient estimators—Scenarios 1–6.

TABLE I
PARAMETER SETTINGS FOR SIX EXAMPLES

Scenario	ρ	α	p	μ	β	C^*
1	0.99	1	0.1	0.1	0.505	7
2	0.99	1	0.05	0.05	0.505	7
3	0.99	2	0.05	0.1	1.01	15
4	0.71	1	0.1	0.1	0.7	13
5	0.71	1	0.05	0.05	0.7	11
6	0.71	2	0.05	0.1	1.4	22

VI. CONCLUSION AND FUTURE WORK

SFMs can adequately describe the dynamics of high-speed communication networks, where they may be used to approximate discrete event models or constitute primary models in their own right. When control and optimization are of primary importance (rather than performance analysis), a SFM may be used as a means for accurately determining an optimal parameter setting, even though the corresponding performance evaluated through the SFM may not be particularly accurate. With this premise in mind, we have considered single-node SFMs from the standpoint of IPA derivative estimation. In particular, we have developed IPA estimators for the loss volume and work as functions of the buffer size, and shown them to be unbiased and nonparametric. The simplicity of the estimators and their nonparametric property suggest their application to online network management. Indeed, for a class of buffer control problems, we have shown how to use an optimization scheme (and illustrated it through numerical examples) for a discrete-event model (viewed as a real, queuing-based single-node system) using the IPA gradient obtained from its SFM counterpart. Interestingly, there is no IPA derivative for the discrete event model, since its associated control parameter is discrete.

For the loss volume performance function, the IPA derivative has been developed by two separate techniques: finite difference analysis, and a sample derivative analysis. The former method is more elaborate, but sheds light on the structure of the derivative estimator. The second method is more direct and elegant, but its unbiasedness proof requires some results obtained by the analysis of the former method. The sample-derivative method was also applied to the IPA estimator of the buffer workload performance function.

The sample derivative analysis holds the promise of considerable extensions to multiple SFMs as models of actual networks and to multiple flow classes that can be used for differentiating traffic classes with different quality-of-service (QoS) requirements. Ongoing research has already led to very encouraging results, reported in [18], involving IPA estimators and associated optimization for flow control purposes in multinode models. Finally, for the purpose of session-by-session admission control, preliminary work suggests that one can use sensitivity information with respect to inflow rates (which can be obtained through an approach similar to the one presented in this paper) and contribute to the development of effective algorithms, yet to be explored.

APPENDIX

Proof of Lemma 1: Looking at any segment of the sample path over an interval $[A_i, A_{i+1})$, there are two possibilities: either $\alpha_i - \beta_i < 0$ or $\alpha_i - \beta_i \geq 0$. First, suppose that $\alpha_i - \beta_i < 0$ and consider the event which occurs at time A_{i+1}. There are three cases to analyze.

Case 1.1) $y_{i+1}(\theta) \geq 0$. In this case, as seen in Fig. 7, we have

$$\Delta y_{i+1} = \Delta x_{i+1} = \Delta x_i. \quad (46)$$

Case 1.2) $y_{i+1}(\theta) < 0$ and $y_{i+1}(\theta) + \Delta y_{i+1} \leq 0$. In this case, as seen in Fig. 8(a), the kth BP ends and it is followed by an EP of length I_i, which in turn ends at time A_{i+1}. Clearly

$$\Delta x_{i+1} = 0. \quad (47)$$

Case 1.3) $y_{i+1}(\theta) < 0$ and $y_{i+1}(\theta) + \Delta y_{i+1} > 0$. This represents a situation where an EP of length I_i is eliminated in the perturbed path, i.e., $I_i < \Delta x_i/(\beta_i - \alpha_i)$. As seen in Fig. 8(b), the buffer content perturbation becomes

$$\Delta x_{i+1} = \Delta x_i - (\beta_i - \alpha_i) I_i. \quad (48)$$

Next, let us assume that $\alpha_i - \beta_i \geq 0$. We then have three cases as follows.

Case 2.1) $y_{i+1}(\theta) \leq \theta$ and $y_{i+1}(\theta) + \Delta y_{i+1} \leq \theta$. It is easy to see [Fig. 9(a)] that this is identical to Case 1.1) yielding (46).

Case 2.2) $y_{i+1}(\theta) \leq \theta$ and $y_{i+1}(\theta) + \Delta y_{i+1} > \theta$. The perturbed buffer content cannot exceed $\theta + \Delta\theta$, since $\Delta y_{i+1} = \Delta x_i \leq \Delta\theta$ from (21); therefore, $y_{i+1} + \Delta y_i \leq \theta + \Delta\theta$ and the situation is identical to that of Fig. 9(a), again yielding (46).

Case 2.3) $y_{i+1}(\theta) > \theta$. As seen in Fig. 9(b)

$$\Delta x_{i+1} = \Delta\theta$$

as in Case II where perturbation generation was considered. Once again, however, it is possible that $\Delta\theta > (\alpha_i - \beta_i)F_i + \Delta x_i$, so that we write, similarly to Case II

$$\Delta x_{i+1} = \Delta\theta - [\Delta\theta - \Delta x_i - (\alpha_i - \beta_i)F_i]^+. \quad (49)$$

We may now establish (23) by combining (46)–(49) and by observing that i) in (48), $I_i < \Delta x_i/(\beta_i - \alpha_i)$ with $\beta_i - \alpha_i > 0$, therefore $0 < \Delta x_{i+1} < \Delta x_i$, and ii) in (49), $0 \leq \Delta x_{i+1} \leq \Delta\theta$.

Next, by combining (46)–(48) we obtain the first part of (24), observing that $I_i = 0$ in Case 1.1). To obtain the second part, we combine (46) and (49), observing that when $F_i = 0$ in (49), we get $\Delta x_{i+1} = \Delta x_i$, since $\Delta\theta - \Delta x_i \geq 0$ from (23), which reduces to (46) corresponding to Cases 2.1)–2.2). ■

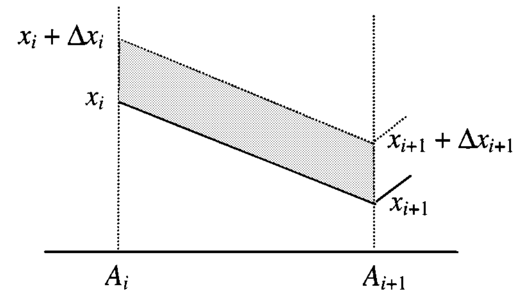

Fig. 7. Case 1.1): $y_{i+1}(\theta) \geq 0$.

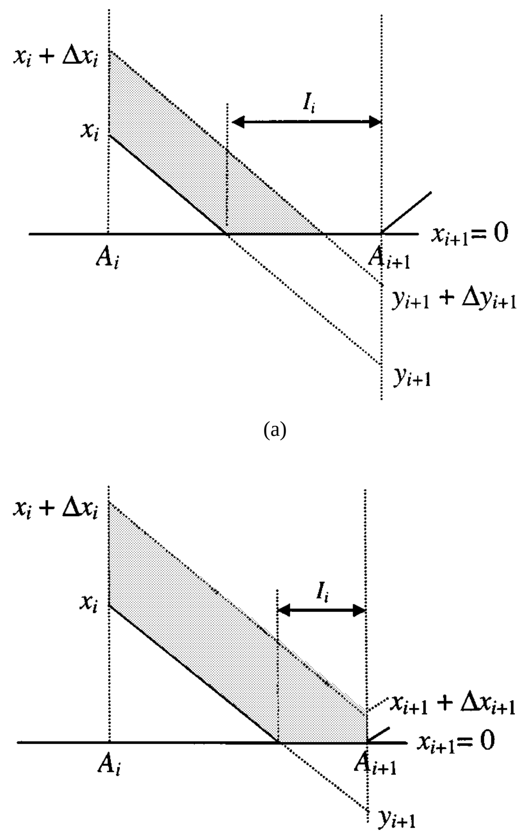

Fig. 8. (a) Case 1.2): $y_{i+1}(\theta) < 0$ and $y_{i+1}(\theta) + \Delta y_{i+1} \leq 0$. (b) Case 1.3): $y_{i+1}(\theta) < 0$ and $y_{i+1}(\theta) + \Delta y_{i+1} > 0$.

Proof of Lemma 2: Proceeding as in the proof of Lemma 1, we first consider the case $\alpha_i - \beta_i < 0$ and get the following.

Case 1.1) $y_{i+1}(\theta) \geq 0$. In this case, as seen in Fig. 7, we have

$$\Delta L_{i+1} = 0. \quad (50)$$

Case 1.2) $y_{i+1}(\theta) < 0$ and $y_{i+1}(\theta) + \Delta y_{i+1} \leq 0$. Clearly, as seen in Fig. 8(a)

$$\Delta L_{i+1} = 0. \quad (51)$$

Case 1.3) $y_{i+1}(\theta) < 0$ and $y_{i+1}(\theta) + \Delta y_{i+1} > 0$. This represents a situation where an EP of length I_i is eliminated in the perturbed path, i.e., $I_i < \Delta x_i/(\beta_i - \alpha_i)$. As seen in Fig. 8(b), no loss is involved in either path

$$\Delta L_{i+1} = 0. \quad (52)$$

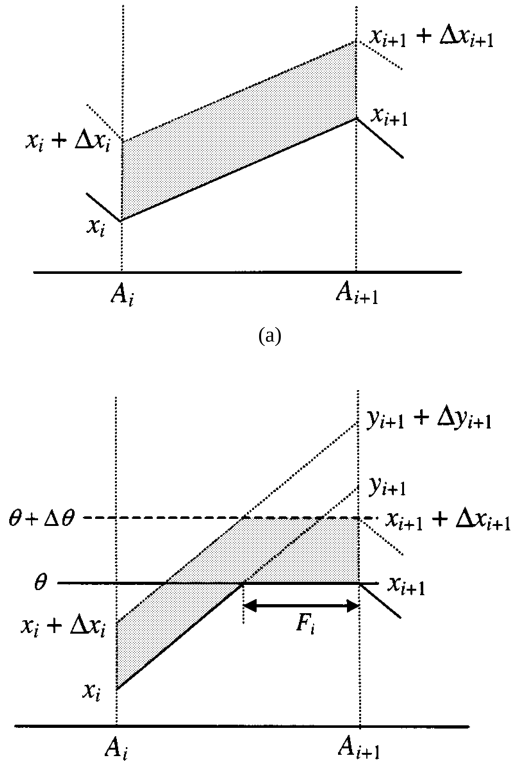

Fig. 9. (a) Cases 2.1)–2.2): $y_{i+1}(\theta) \le \theta$. (b) Case 2.3): $y_{i+1}(\theta) > \theta$.

Next, let us assume that $\alpha_i - \beta_i \ge 0$ and we have the following.

Case 2.1) $y_{i+1}(\theta) \le \theta$ and $y_{i+1}(\theta)+\Delta y_{i+1} \le \theta$. It is easy to see [Fig. 9(a)] that this is identical to Case 1.1) yielding (50).

Case 2.2) $y_{i+1}(\theta) \le \theta$ and $y_{i+1}(\theta)+\Delta y_{i+1} > \theta$. As argued in the proof of Lemma 1, the situation is identical to that of Fig. 9(a), again yielding (50).

Case 2.3) $y_{i+1}(\theta) > \theta$. If $\Delta\theta > (\alpha_i - \beta_i)F_i + \Delta x_i$, then $\Delta L_{i+1} = 0 - (y_{i+1} - \theta) = -(\alpha_i - \beta_i)F_i$. Otherwise, $L_{i+1}(\theta + \Delta\theta) = y_{i+1} + \Delta x_i - \theta - \Delta\theta$, and we get $\Delta L_{i+1} = \Delta x_i - \Delta\theta$. Thus

$$\Delta L_{i+1} = (\Delta x_i - \Delta\theta) + [\Delta\theta - \Delta x_i - (\alpha_i - \beta_i)F_i]^+. \quad (53)$$

We may now combine (50)–(53). Observe that in (53) $\Delta L_{i+1} \ge -\Delta\theta$, since we have already established that $\Delta x_i \ge 0$ in Lemma 1. Moreover, $\Delta L_{i+1} = -(\alpha_i - \beta_i)F_i < 0$ if $\Delta\theta - \Delta x_i - (\alpha_i - \beta_i)F_i > 0$, and $\Delta L_{i+1} = \Delta x_i - \Delta\theta$ if $\Delta\theta - \Delta x_i - (\alpha_i - \beta_i)F_i \le 0$, where $\Delta x_i - \Delta\theta \le 0$ from Lemma 1. This yields (26). ■

Proof of Lemma 3: Proceeding as in the proof of Lemma 2, we get $\Delta L_{i+1} = 0$ in (50)–(52), i.e., in all cases except Case 2.3) where (53) applies

$$\Delta L_{i+1} = (\Delta x_i - \Delta\theta) + [\Delta\theta - \Delta x_i - (\alpha_i - \beta_i)F_i]^+.$$

Suppose the first overflow interval in the BP ends at A_r. Under the assumption $\Delta\theta - \Delta x_i - (\alpha_i - \beta_i)F_i \le 0$ for all $i = j, \ldots, m$, it follows from (17) and (18) that $\Delta x_r = \Delta\theta$ and $\Delta L_r = -\Delta\theta$. Moreover, from Lemma 1, (24) gives $\Delta x_i = \Delta\theta$ for all $i = r+1, \ldots, m$. Therefore, (53) gives $\Delta L_{i+1} = 0$ after every subsequent overflow interval, and we get $\Lambda_k(\Delta\theta) = \Delta L_r = -\Delta\theta$. ■

References

[1] D. Anick, D. Mitra, and M. Sondhi, "Stochastic theory of a data-handling system with multiple sources," *Bell System Tech. J.*, vol. 61, pp. 1871–1894, 1982.

[2] H. Kobayashi and Q. Ren, "A mathematical theory for transient analysis of communications networks," *IEICE Trans. Commun.*, vol. E75-B, pp. 1266–1276, 1992.

[3] R. Cruz, "A calculus for network delay, Part I: Network elements in isolation," *IEEE Trans. Inform. Theory*, vol. 37, pp. 114–131, Jan. 1991.

[4] Y. Wardi and B. Melamed, "IPA gradient estimation for the loss volume in continuous flow models," in *Proc. Hong Kong Int. Workshop on New Directions of Control and Manufacturing*, Hong Kong, Nov. 1994, pp. 30–33.

[5] G. Kesidis, A. Singh, D. Cheung, and W. Kwok, "Feasibility of fluid-driven simulation for ATM network," in *Proc. IEEE Globecom*, vol. 3, 1996, pp. 2013–2017.

[6] K. Kumaran and D. Mitra, "Performance and fluid simulations of a novel shared buffer management system," in *Proc. IEEE INFOCOM*, Mar. 1998.

[7] N. Miyoshi, "Sensitivity estimation of the cell-delay in the leaky bucket traffic filter with stationary gradual input," in *Proc. Int. Workshop on Discrete Event Systems, WoDES'98*, Cagliari, Italy, Aug. 1998, pp. 190–195.

[8] B. Liu, Y. Guo, J. Kurose, D. Towsley, and W. Gong, "Fluid simulation of large scale networks: Issues and tradeoffs," in *Proc. Int. Conf. on Parallel and Distributed Processing Techniques and Applications*, Las Vegas, NV, June 1999.

[9] A. Yan and W. Gong, "Fluid simulation for high-speed networks with flow-based routing," *IEEE Trans. Inform. Theory*, vol. 45, pp. 1588–1599, July 1999.

[10] Y. Wardi and B. Melamed, "Loss volume in continuous flow models: Fast simulation and sensitivity analysis via IPA," in *Proc. 8th IEEE Mediterranean Conf. on Control and Automation (MED 2000)*, Patras, Greece, July 17–19, 2000.

[11] ——, "Variational bounds and sensitivity analysis of traffic processes in continuous flow models," *Discrete Event Dyna. Syst.: Theory Appl.*, vol. 11, pp. 249–282, 2001.

[12] B. Mohanty and C. Cassandras, "The effect of model uncertainty on some optimal routing problems," *J. Optim. Theory Appl.*, vol. 77, pp. 257–290, 1993.

[13] S. Meyn, "Sequencing and routing in multiclass networks. Part I: Feedback regulation," in *Proc. IEEE Int. Symp. on Information Theory*, 2000, pp. 4440–4445.

[14] Y. C. Ho and X. Cao, *Perturbation Analysis of Discrete Event Dynamic Systems*. Norwell, MA: Kluwer, 1991.

[15] C. G. Cassandras and S. Lafortune, *Introduction to Discrete Event Systems*. Norwell, MA: Kluwer, 1999.

[16] Y. Liu and W. Gong, "Perturbation analysis for stochastic fluid queueing systems," in *Proc. 38th IEEE Conf. Decision Control*, 1999, pp. 4440–4445.

[17] H. Kushner and D. Clark, *Stochastic Approximation for Constrained and Unconstrained Systems*. Berlin, Germany: Springer-Verlag, 1978.

[18] C. Cassandras, G. Sun, and C. Panayiotou, "Stochastic fluid models for control and optimization of systems with quality of service requirements," in *Proc. 40th IEEE Conf. Decision Control*, Orlando, FL, 2001, pp. 1917–1922.

[19] Y. Wardi and B. Melamed, "Continuous flow models: Modeling, simulation and continuity properties," in *Proc. 38th IEEE Conf. Decision Control*, Phoenix, AZ, Dec. 7–10, 1999, pp. 34–39.

[20] R. Y. Rubinstein and A. Shapiro, *Discrete Event Systems: Sensitivity Analysis and Stochastic Optimization by the Score Function Method*. New York: Wiley, 1993.

[21] S. Meyn, "Sequencing and routing in multiclass networks. Part I: Feedback regulation," *SIAM J. Control Optim.*, 2002, to be published.

Christos G. Cassandras (M'82–SM'91–F'96) received the B.S. degree from Yale University, New Haven, CT, the M.S.E.E. degree from Stanford University, Stanford, CA, and the S.M. and Ph.D. degrees from Harvard University, Cambridge, MA, in 1977, 1978, 1979, and 1982, respectively.

From 1982 to 1984, he was with ITP Boston, Inc., Boston, MA, where he worked on the design of automated manufacturing systems. From 1984 to 1996, he was a Faculty Member at the Department of Electrical and Computer Engineering, University of Massachusetts, Amherst. Currently, he is Professor of Manufacturing Engineering and Professor of Electrical and Computer Engineering at Boston University, Boston, MA. He specializes in the areas of discrete-event and hybrid systems, stochastic optimization, and computer simulation, with applications to computer networks, manufacturing systems, transportation systems, and command-control systems. He has published over 190 papers in these areas, and two textbooks.

Dr. Cassandras received the 1999 Harold Chestnut Prize by the IFAC for one of his textbooks. He currently serves as Editor-in-Chief of the IEEE TRANSACTIONS ON AUTOMATIC CONTROL, and has served on several editorial boards and as Guest Editor for various journals. He is a Member of the IEEE Control Systems Society Board of Governors and an IEEE Distinguished Lecturer. He was awarded a 1991 Lilly Fellowship, and is also a Member of Phi Beta Kappa and Tau Beta Pi.

Yorai Wardi (S'80–M'81) received the Ph.D. degree in electrical engineering and computer sciences from the University of California at Berkeley in 1982.

From 1982 to 1984, he was a Member of Technical Staff at Bell Telephone Laboratories and Bell Communications Research, Holmdel, NJ. Since 1984, he has been with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, where he currently is an Associate Professor. He spent the 1987–1988 academic year at the Department of Industrial Engineering and Management, Ben Gurion University of the Negev, Israel. His interests are in discrete-event dynamic systems, perturbation analysis, and network control.

Dr. Wardi was a Member of the Conference Editorial Board of the IEEE Control Systems Society from 1999 to 2001, and is currently an Associate Editor of the IEEE TRANSACTIONS ON AUTOMATIC CONTROL.

Benjamin Melamed (M'85–SM'91–F'94) received the B.Sc. degree in mathematics and statistics from Tel Aviv University, Tel Aviv, Israel, and the M.S. and Ph.D. degrees in computer science from the University of Michigan, Ann Arbor, in 1972, 1973, and 1976, respectively.

From 1977 to 1981, he taught in the Department of Industrial Engineering and Management Science, Northwestern University, Evanston, IL. He joined the Performance Analysis Department at Bell Laboratories, Holmdel, NJ, in 1981, and later became an AT&T Fellow. He moved to NEC, Princeton, NJ, in 1989 and served there as a Deputy Director, Head of the Performance Analysis Department, and NEC Fellow. He consulted at Bellcore, Morristown, NJ, in 1995, and joined Rutgers—The State University of New Jersey in 1996, where he is currently a Professor in the School of Business, Newark, NJ, and the Department of Management Science and Information Systems, Piscataway, NJ (Rutgers Business School-Newark and New Brunswick). He is also a member of the Rutgers Center for Operations Research (RUTCOR). His research interests include system modeling and analysis (especially telecommunications systems and traffic modeling), simulation, stochastic processes, and visual modeling environments, and his research has been supported by DARPA and the National Science Foundation. He has authored or coauthored over 80 papers, and the book *Modern Modeling and Simulation* (with R. Rubinstein) (New York: Wiley, 1998.

Dr. Melamed became an AT&T Fellow in 1988, and an IFIP WG7.3 Member in 1997. He was elected to Beta Gamma Sigma in 1998.

Gang Sun received the B.E. degree in electrical engineering from Xian Jiaotong University, Xian, China, and the M.E. degree in system and control from the Institute of Automation, Chinese Academy of Science, Beijing, China, in 1994 and 1999, respectively. He is currently working toward the Ph.D. degree in the Department of Manufacturing Engineering, Boston University, Boston, MA.

From 1994 to 1996, he worked as a System Engineer at Xian Instruments (Group), Xian, China. His current research interests lie in the analysis, optimization, and simulation of discrete-event systems, with applications to communication networks and manufacturing systems.

Christos G. Panayiotou (S'94–M'99) received the B.S. and Ph.D. degrees in electrical and computer engineering from the University of Massachusetts at Amherst, in 1994 and 1999, respectively, and the M.B.A. degree from the Isenberg School of Management, University of Massachusetts at Lowell, in 1999.

He is currently a Research Associate at the Manufacturing Engineering Department, Boston University, Boston, MA. His research interests include optimization and control of discrete-event systems with emphasis on computer communication networks, manufacturing systems, and transportation networks.

Dr. Panayiotou is an Associate Editor for the Conference Editorial Board of the IEEE Control Systems Society.