

IPA for Loss Volume and Buffer Workload in Tandem SFM Networks *

Yorai Wardi
George F. Riley

School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0250
{ywardi,riley}@ece.gatech.edu
(404)894-8326, Fax: (404)385-9959

July 10, 2002

Abstract

This paper considers congestion-related performance metrics in tandem networks of Stochastic Fluid Models (SFM), and derives their IPA gradient estimators with respect to buffer sizes. Specifically, the performance metrics in question are the total loss volume and the cumulative buffer workload (buffer contents), and the control parameter consists of buffer limits at both the node where the performance is measured and at an upstream node. The IPA estimators are unbiased and nonparametric, and hence can be computed on-line from field measurements as well as off-line from simulation experiments. The IPA derivatives are applied to packet-based networks, where simulation results support the theoretical developments. Possible applications to congestion management in telecommunications networks are discussed.

Key words. IPA, stochastic flow networks, telecommunications networks, congestion management.

1 Introduction

The deployment of Infinitesimal Perturbation Analysis (IPA) to queuing networks has been fraught with technical difficulties, largely due to the fact that IPA gradients are generally biased. To get around this problem, Stochastic Flow Models (SFM) recently have been considered as an alternative paradigm to queuing networks for modeling and simulation of telecommunications networks [1, 7, 2, 8]. SFM networks offer two distinct advantages over their queuing-network counterparts: they can be faster to simulate, and they give unbiased IPA gra-

dient estimators for a large number of network configurations, queuing disciplines, and performance functions [6]. On-going research in IPA for SFM has been focused on congestion-related performance measures as functions of various operational parameters and traffic parameters, like inflow rates, transmission bandwidth, and buffer limits [1, 7, 2, 8]. The performance metrics in question were either the total loss volume, or the cumulative buffer contents (workload), over a given time interval. The derived IPA gradient estimators have had an additional appealing property: they were nonparametric, namely, computable by formulas that are independent of the probability law underlying the network. This suggests that they can be deployed not only in off-line simulation, but also in on-line provisioning, management and control.

Most of the results obtained so far pertain to single-node systems with either a single-flow class [1, 7] or several multi-flow classes [2]. Tandem networks have been considered in [8], where the loss volume at a given node is considered as a function of the buffer limit at an upstream node. We will summarize these results and supplement them with the IPA derivative of the buffer workload at a node as function of the buffer limit at an upstream node. Since the derived IPA derivative estimators are nonparametric, they can be used not only in SFM networks, but also in packet-based networks. We demonstrate this point by presenting simulation results showing the efficacy of the IPA derivatives in performance prediction. We mention that, based on queuing-network models, the IPA derivatives are biased and hence cannot be adequately used in performance prediction. What we do is to derive the IPA derivatives in the setting of fluid-flow networks, where they are unbiased and nonparametric, and then apply them to queuing-based models. The results exhibit good performance prediction.

Section 2 presents the basic SFM paradigm, and Section 3 surveys the existing results derived for single-node

*This work is supported in part by NSF under contract number ANI-9977544 and DARPA under contract numbers N66002-00-1-8934 and F30602-00-2-0556.




Figure 1: The Basic SFM

systems. Section 4 presents the results for tandem networks, and Section 5 contains simulation results. Finally, Section 6 concludes the paper and discusses some possible applications to network management and control.

2 The Basic SFM

The basic single-server, single-flow SFM is shown in Figure 1. It consists of a fluid server preceded by a storage tank. Its inflow rate and service rate at time t are denoted by $\alpha(t)$ and $\beta(t)$, respectively, where the time t is assumed confined to a given interval $[0, T]$. The processes $\{\alpha(t)\}$ and $\{\beta(t)\}$ are stochastic processes defined over a suitable common probability space, (Ω, \mathcal{F}, P) . Let c denote the buffer size, and suppose that $c > 0$. The inflow-rate process $\{\alpha(t)\}$ and service-rate process $\{\beta(t)\}$ define, together with the buffer size, much of the behavior of the SFM, and hence are labeled the *defining processes* (see [6]).

Associated with these defining processes are three *derived processes*, namely the workload (buffer contents) $\{x(t)\}$, overflow rate $\{\gamma(t)\}$, and outflow rate $\{\delta(t)\}$ (see Figure 1, where the dependence of the various processes on t is implicit). The workload $x(t)$ is the amount of fluid in the buffer at time t , and it depends on the defining processes via the following one-sided differential equation (where $\nu(t)$ is defined by $\nu(t) = \alpha(t) - \beta(t)$):

$$\frac{dx(t)}{dt^+} = \begin{cases} 0, & \text{if } x(t) = 0 \text{ and } \nu(t) \leq 0, \\ 0, & \text{if } x(t) = c \text{ and } \nu(t) \geq 0, \\ \nu(t), & \text{otherwise,} \end{cases} \quad (2.1)$$

whose boundary condition is set to $x(0) = 0$ to simplify the presentation. The overflow rate $\gamma(t)$ is the spillover rate due to full buffer, and it is defined by

$$\gamma(t) = \begin{cases} \nu(t), & \text{if } x(t) = c \text{ and } \nu(t) \geq 0, \\ 0, & \text{otherwise.} \end{cases} \quad (2.2)$$

Finally, the outflow rate $\delta(t)$ is the fluid discharge rate from the server, defined by

$$\delta(t) = \begin{cases} \beta(t), & \text{if } x(t) > 0, \\ \alpha(t), & \text{if } x(t) = 0. \end{cases} \quad (2.3)$$

We will be concerned with two performance measures: the loss volume, and the buffer workload. The loss vol-

ume, denoted by L_T , is defined by

$$L_T := \int_0^T \gamma(t) dt, \quad (2.4)$$

and the buffer workload, denoted by Q_T , is defined by

$$Q_T = \int_0^T x(t) dt. \quad (2.5)$$

We mention that a number of congestion-related performance measures of interest are related to L_T and Q_T , like the average loss rate $T^{-1}E[L_T]$ ($E[\cdot]$ denoting expectation), the loss probability defined by $E[(\int_0^T \alpha(t) dt)^{-1} L_T]$, and the expected delay (sojourn time) which can be obtained from $E[Q_T]$ via a fluid variant of Little's Law [3, 4].

We note that the processes $\{\alpha(t)\}$ and $\{\beta(t)\}$ can have fairly general form and, in particular, need not be statistically independent. Their realizations, namely the functions $\alpha(t)$ and $\beta(t)$, may be discontinuous, and are only required to be piecewise continuously differentiable.

3 IPA Derivatives: Single-Stage

Consider the loss volume and buffer workload in the basic SFM described in Section 2 as functions of the buffer size (also called the buffer limit). In order to conform to the notation commonly used in the literature on IPA, we denote the buffer size by θ , $\theta = c$. We observe that the defining processes $\{\alpha(t)\}$ and $\{\beta(t)\}$ are independent of θ , but the derived processes, $\{x(t)\}$, $\{\gamma(t)\}$, and $\{\delta(t)\}$, certainly depend on θ , and hence are denoted by $\{x(\theta; t)\}$, $\{\gamma(\theta; t)\}$, and $\{\delta(\theta; t)\}$, respectively. This dependence is apparent by Eqs. (2.1)–(2.3) which now merit the notational modification of replacing $x(t)$, $\gamma(t)$, and $\delta(t)$ by $x(\theta; t)$, $\gamma(\theta; t)$, and $\delta(\theta; t)$. Subsequently, the performance measures of loss volume and buffer workload also depend on θ . Respectively denoted by $L(\theta)$ and $Q(\theta)$, they assume the following forms,

$$L(\theta) = \int_0^T \gamma(\theta; t) dt, \quad (3.1)$$

and

$$Q(\theta) = \int_0^T x(\theta; t) dt, \quad (3.2)$$

where an explicit notational dependence on T is omitted for the sake of simplicity.

The IPA derivatives, described below, will be denoted by $L'(\theta)$ and $Q'(\theta)$, respectively, where “prime” denotes derivative with respect to θ . To describe these derivatives, we observe that a typical trajectory of the buffer workload, $x(\theta; t)$, switches between three states: empty, partial, and full. In the empty state the buffer is empty and the server discharges fluid “molecules” as soon as they arrive without any buffer buildup; in the full state the buffer is full




Figure 2: Buffer State Trajectory

and spillover occurs as a result; and in the partial state the buffer is neither full nor empty. An *empty period* is a maximal subinterval of $[0, T]$ in which the buffer is empty; a *full period* is a maximal subinterval of $[0, T]$ in which the buffer is full; and a *partial period* is a supremal interval in which the buffer is neither full nor empty. Since the function $x(\theta; t)$ is generally continuous in t for a fixed θ , it follows that empty periods and full periods are closed intervals, whereas partial periods are open intervals in the relative topology induced by $[0, T]$. A *buffering period*, also called *busy period* in [1], or *nonempty period* in [7], is defined as a supremal subinterval of $[0, T]$ during which the buffer is not empty. Observe that buffering periods are unions of partial periods and full periods. A buffering period is said to be *lossy* if some loss occurs during any time in it, that is, it is the union of some partial periods and at least one full period incurring loss.

Now fix $\theta \in \Theta$, and consider the state trajectory $x(\theta; t)$. Suppose there are $K = K(\theta)$ lossy buffering periods in the interval $[0, T]$. Then the following propositions were proved in [1, 7] under mild assumptions.

Proposition 3.1. $L'(\theta) = -K$, and this IPA derivative is unbiased. \square

Regarding the IPA derivative $Q'(\theta)$, let us denote by $v_{k,1}$ the first time the buffer becomes full in the k^{th} lossy buffering period, and let $\eta_k(\theta)$ be the last time-point in that buffering period. The following proposition was proved in [1, 7].

Proposition 3.2. The IPA derivative $Q'(\theta)$ has the following form,

$$Q'(\theta) = \sum_{k=1}^K [\eta_k(\theta) - v_{k,1}(\theta)],$$

and it is unbiased. \square

We remark that these two IPA derivatives are clearly nonparametric: $L'(\theta)$ requires the counting of lossy buffering periods in the interval $[0, T]$, and $Q'(\theta)$ requires identifying those time-points in which the buffer becomes full or empty. To illustrate, consider Figure 2, where $L'(\theta) = -1$ and $Q'(\theta) = \eta_1 - v_1$.

4 IPA Derivatives: Tandem

Consider the network shown in Figure 3, consisting of two SFMs in tandem. Denoting them by SFM_1 and SFM_2 ,




Figure 3: Tandem Network




Figure 4: Buffer State Trajectory in SFM2

we note that the inflow process to SFM_2 consists of the outflow process from SFM_1 . The various traffic processes and other variables will be labeled with the subscript 1 or 2 according to the relevant SFM. Thus, we see in Figure 3 that $\alpha_2(\theta; t) = \delta_1(\theta; t)$. The network is assumed to operate over a fixed time horizon $[0, T]$, and both SFMs are assumed to be empty at time $t = 0$.

Let us consider the loss volume and buffer workload at SFM_2 as functions of the buffer size at SFM_1 . Thus, denoting the above buffer limit by θ , we define the functions $L(\theta)$ and $Q(\theta)$, respectively, by

$$L(\theta) = \int_0^T \gamma_2(\theta; t) dt, \quad (4.1)$$

and

$$Q(\theta) = \int_0^T x_2(\theta; t) dt. \quad (4.2)$$

The dependence of these functions on θ comes through the inflow rate at SFM_2 , which is identical to the outflow rate at SFM_1 . Indeed, by Eq. (2.3), $\alpha_2(\theta; t)$ has the following form,

$$\alpha_2(\theta; t) = \delta_1(\theta; t) = \begin{cases} \beta_1(t), & \text{if } x_1(\theta; t) > 0 \\ \alpha_1(t), & \text{if } x_1(\theta; t) = 0. \end{cases} \quad (4.3)$$

Thus, θ affects L and Q through the switchover points (times) of $\alpha_2(\theta; t)$ back and forth between $\beta_1(t)$ and $\alpha_1(t)$. A switch from $\alpha_1(t)$ to $\beta_1(t)$ corresponds to the event that the buffer at SFM_1 ceases to be empty. Such events generally depend on the defining processes $\alpha_1(t)$ and $\beta_1(t)$ but not on θ , and hence the switching time is locally independent of θ . On the other hand, a switch in the opposite direction, namely from $\beta_1(t)$ to $\alpha_1(t)$, corresponds to the end of a buffering period at SFM_1 . Such a time-point generally depends on θ as long as the buffering period ending at SFM_1 experienced some loss. Thus, an infinitesimal perturbation in θ is propagated from SFM_1 to SFM_2 only through the end points of lossy buffering periods at SFM_1 (this will become apparent from the formulas for the IPA derivatives, presented later). Consequently, we label such time-points as *active switchover points*.

Consider now a typical state trajectory in SFM_2 (namely, the graph of $x_2(t)$ for $t \in [0, T]$; see Figure 4 as a visual aid). Suppose it has N lossy buffering periods, denoted by B_n , $n = 1, \dots, N$, in increasing order (in Figure 4, $N = 1$). Let B_n start at the point ξ_n , and let s_n be the last time-point in B_n when the buffer is full. The IPA derivative $L'(\theta)$ consists of the sum of all the contributions of the various active switchover points. Under mild assumptions, this contribution can be summarized as follows.

- If an active switchover point is in one of the intervals $[\xi_n, s_n)$, then its contribution is 1.
- If an active switchover point is one of the points s_n , then its contribution is

$$\frac{\beta_1(s_n) - \beta_2(s_n)}{\beta_1(s_n) - \alpha_1(s_n)}.$$

- The contribution of all other active switchover points is 0.

To formalize, let M denote the number of active switchover points in the set

$$\cup_{n=1}^N [\xi_n, s_n),$$

and define the index set Φ by

$$\Phi = \{n = 1, \dots, N : s_n \text{ is an active switchover point}\}.$$

The following result was derived in [8] under mild assumptions.

Proposition 4.1. $L'(\theta)$ has the following form,

$$L'(\theta) = M + \sum_{n \in \Phi} \frac{\beta_1(s_n) - \beta_2(s_n)}{\beta_1(s_n) - \alpha_1(s_n)}, \quad (4.4)$$

and this IPA derivative is unbiased. \square

It is apparent that the term M in the right-hand side of Eq. (4.4) is nonparametric: all that its computation requires is to determine where the active switchover points fall in the state trajectory at SFM_2 . For example, consider the state trajectory shown in Figure 4, and suppose that the active switchover points are u_1, u_2 and u_3 . Then the contributions of u_1 and u_2 are 1 each while the contribution of u_3 is 0, and therefore, $L'(\theta) = 2$. On the other hand, the second term in the right-hand side of Eq. (4.4) is nonparametric only to the extent that instantaneous flow rates can be computed on-line. Each term in the sum was shown in [8] to be in the interval $[0, 1]$, and therefore it may be possible to approximate it by 0.5. In fact, in our simulation experiments we neglected this term outright without apparent loss of precision.

Finally, consider the workload $Q(\theta)$. As for $L(\theta)$, the active switchover points carry forward the perturbations from SFM_1 to SFM_2 . Let us denote by P_m , $m =$

$1, \dots, M$, the partial periods at SFM_2 , in increasing order, and suppose that each partial period P_m starts at a point σ_m and ends at a point τ_m . That is, $P_m = (\sigma_m, \tau_m)$. The contribution of each active switchover point to $Q'(\theta)$ can be summarized as follows.

- If an active switchover point u is in P_m , then its contribution is $\tau_m - u$.
- If σ_m is an active switchover point, then its contribution is

$$\frac{\beta_2(\sigma_m) - \alpha_1(\sigma_m)}{\beta_1(\sigma_m) - \alpha_1(\sigma_m)} \times [\tau_m - \sigma_m].$$

- The contribution of all other active switchover points is 0.

Formally, let $u_{k,m}$, $k = 1, \dots, K_m$ denote the active switchover points in the partial period $P_m = (\sigma_m, \tau_m)$, and let Ψ denote the following index set,

$$\Psi = \{m = 1, \dots, M : \sigma_m \text{ is an active switchover point}\}.$$

Under mild assumption, the following is in force (see [9] for a proof).

Proposition 4.2. $Q'(\theta)$ has the following form,

$$Q'(\theta) = \sum_{m=1}^M \sum_{k=1}^{M_m} [\tau_m - u_{k,m}] + \sum_{m \in \Psi} \frac{\beta_2(\sigma_m) - \alpha_1(\sigma_m)}{\beta_1(\sigma_m) - \alpha_1(\sigma_m)} \times [\tau_m - \sigma_m], \quad (4.5)$$

and this IPA derivative is unbiased. \square

As an example, in figure 4 we have that $Q'(\theta) = (\tau_1 - u_1) + (\tau_3 - u_3)$.

5 Simulation Experiments

In this section, we describe the simulation experiments we performed to validate that the IPA methodology can be applied to packet based networks and make meaningful predictions regarding the sensitivity of network performance to small adjustments in network configuration parameters. The simulation experiments were performed using the popular and publicly available *ns2* network simulator [5]. The *ns2* simulator is a packet based, discrete event simulation environment that models a wide variety of network protocols and traffic characteristics.

Using the *ns2* simulator, we created the network model shown in Figure 5, consisting of the following elements.

1. n data sources, each of which is modeled as a bursty on/off data generator with average rate λ_s bits per second. The on and off periods for the generators are sampled from a Pareto distribution with an α of 1.5. All results reported here are for $n = 20$, but we obtained similar results for other values of n . The value used for λ_s is discussed below.




Figure 5: Simple Two-Stage Topology

2. A single-queue, single-server system S1 at stage 1. The service rate at this queue is λ_1 bits per second, and the buffer size is Θ_1 , specified in units of packets. For all experiments, λ_1 is 1 Megabit per second, and the packet size is 512 bytes, a typical IP packet size. The buffer size is varied from 10 packets to 50 packets as described below.
3. A single-queue, single-server system S2 at stage 2. The service rate at this queue is $\lambda_2 = 2\lambda_1$. The buffer size Θ_2 is fixed at 20 packets for all experiments.
4. A data generator modeling competing, background traffic (cross traffic) into the stage 2 queue. The background traffic rate λ_b , representing the data rate in bits per second, is periodically sampled from the uniform distribution on the interval $[0.5\lambda_1 \dots 1.5\lambda_1]$. Since the overall service rate at stage 2 is $2\lambda_1$, and the background traffic consumes an average of λ_1 , the remaining bandwidth available for servicing traffic from stage 1 is λ_1 bits per second.
5. A multiple-queue, multiple-server system modeling the traffic sinks. The service rate for each is large relative to the amount of traffic offered, and thus little queuing or delay is seen at this point.
6. The average data rate for each data source λ_s is defined to be $\rho_1(\lambda_1/n)$. In other words, the average aggregate traffic intensity offered at S1 is ρ_1 . For our experiments, we used values for ρ_1 of 0.85, 0.95, and 1.00.
7. The simulator was run for a total of 100 simulation seconds, and the results analyzed as described below.

The purpose of this simulation experiment was to demonstrate that, by using simple counting processes described previously, we can make accurate predictions of the affect of changing the buffer size Θ_1 at stage 1 on the packet loss rate at stage 2. The realization of the counting process is defined by the following algorithm.

1. Initialize an accumulator to 0.
2. Initialize a temporary variable to 0.

3. Monitor the queue at stage 1 for an *active switchover point*. At the time of each active switchover point found in stage 1:
 - (a) If the queue at stage 2 is full, add 1 to the accumulator.
 - (b) If the queue at stage 2 is in a buffering period, add 1 to the temporary variable.
4. Monitor the queue at stage 2 for transitions from a buffering period to a full period. On each such transition, clear the temporary variable to 0.
5. Monitor the queue at stage 2 for transitions from a buffering period to a lossy period. On each such transition, add the temporary variable to the accumulator and clear the temporary variable to 0.

At the completion of the simulation, the negative value of the accumulator is the IPA sensitivity estimator of the loss volume at stage 2 as function of the buffer size at stage 1. In other words, increasing the buffer size Θ_1 by one packet should increase the number of lost packets at stage 2 roughly by the number of packets specified in the accumulator.

The simulation experiments described were executed 41 times, varying the queue limit Θ_1 from 10 to 50 inclusive. For each run k , the value of the accumulator A_k was noted, along with the total number of packets lost (L_k) at the stage 2 queue. The predicted value for L_{k+1} for each run is $L_{k+1} = L_k + A_k$. The results from one set of experiments with $\rho_1 = 0.95$, are shown in Figure 6. The graph shows the observed value of the number of lost packets for each run, and the predicted values for the numbers of lost packets. As can be seen, the observed values very closely match the predictions, with the two graph lines essentially on top of each other. The largest error is 7 packets out of 1734 packets. For all Θ_1 queue limits (buffer sizes) greater than 24, the prediction error is 1 packet or less. Figure 7 shows the same data presented as a percentage error relative to the total number of packets dropped. As can be seen, the percentage error is never more than 0.5%, and typically less than 0.05%.

6 Conclusions and Future Research

We have demonstrated the efficacy of IPA derivative estimators for performance predictions in tandem queuing networks. The performance functions of interest were the loss volume and workload over a given time horizon. In the setting of packet based models, the IPA estimators for these performance functions are biased. Consequently, we derived the IPA derivatives in the setting of stochastic fluid models, where they were shown to be unbiased and nonparametric, and then applied them to the discrete setting of packet based networks. In some practical situations, the tandem configuration can be used to model paths in general-topology networks, where, owing to the nonparametric nature of the IPA derivatives, it would be




Figure 6: Predicted vs. Observed Packet Loss

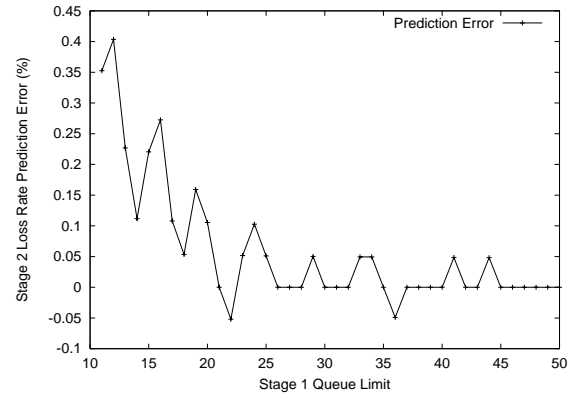


Figure 7: Loss Prediction Error, Percent

possible to assess the impact of cross traffic by direct observations. In these situations it may be possible to use the IPA derivative estimators in on-line congestion management and control.

References

- [1] C.G. Cassandras, Y. Wardi, B. Melamed, G. Sun and C.G. Panayiotou, "Perturbation Analysis for On-Line Control and Optimization of Stochastic Fluid Models", *IEEE Transactions on Automatic Control*, to appear.
- [2] C.G. Cassandras, G. Sun, C.G. Panayiotou, and Y. Wardi, "Perturbation Analysis for On-Line Control and Optimization of Multiclass Stochastic Fluid Models", submitted, *IEEE Transactions on Automatic Control*, January 2002.
- [3] T. Konstantopoulos and G. Last, "On the Dynamics and Performance of Stochastic Fluid Systems", *Journal of Applied Probability*, Vol. 37, pp. 652-667, 2000.
- [4] T. Konstantopoulos, M. Zazanis, and G. de Veciana, "Conversion Laws and Reflection Mappings with an Application to Multiclass Mean Value Analysis of Stochastic Fluid Queues", *Stochastic Processes and their Applications*, Vol. 65, pp. 139-146, 1996.
- [5] S. McCanne and S. Floyd, "The LBNL Network Simulator", Software on-line: <http://www.isi.edu/nsnam>, Lawrence Berkeley Laboratory, 1997.
- [6] Y. Wardi and B. Melamed, "Variational Bounds and Sensitivity Analysis of Continuous Flow Models", *J. of Discrete Event Dynamic Systems*, Vol. 11, No. 3, pp. 249-282, 2001.
- [7] Y. Wardi, B. Melamed, C.G. Cassandras, and C.G. Panayiotou, "On-Line IPA Gradient Estimators in Stochastic Continuous Fluid Models", *J. of Optimization Theory and Applications*, to appear.
- [8] Y. Wardi and B. Melamed, "Estimating Nonparametric IPA Derivatives of Loss Functions in Tandem Fluid Models", in *Proc. IEEE Conf. on Decision and Control*, Orlando, Florida, December 4-7, 2001.
- [9] Y. Wardi and G. Riley, "IPA for Congestion Management in Tandem Fluid-Flow Networks", Technical Memorandum, School of Electrical and Computer Engineering, Georgia Institute of Technology, January 2002.