Perturbation Analysis for Production Control and Optimization of Manufacturing Systems

Haining Yu [a], Christos G. Cassandras [a],

[a] *Dept. of Manufacturing Engineering, Boston University, Brookline, MA 02446*

Abstract

We use Stochastic Fluid Models (SFM) to capture the operation of threshold-based production control policies in manufacturing systems without resorting to detailed discrete event models. By applying Infinitesimal Perturbation Analysis (IPA) to a SFM of a workcenter, we derive gradient estimators of throughput and buffer overflow metrics with respect to production control parameters. It is shown that these gradient estimators are unbiased and independent of distributional information of supply and service processes involved. In addition, based on the fact that they can be evaluated using data from the observed actual (discrete event) system, we use them as approximate gradient estimators in simple iterative schemes for adjusting thresholds (hedging points) on line seeking to optimize an objective function that trades off throughput and buffer overflow costs.

Key words: Manufacturing system, Perturbation Analysis, Performance Optimization

1 Introduction

Production control problems in manufacturing systems have been widely studied, starting with the pioneering work in [1]. In a typical manufacturing setting, a machine may either occasionally break down, or, in a multi-product environment, it may be temporarily inaccessible to a certain part buffer because it is serving another one. In the latter case, from the point of view of a buffer, the machine appears to be failure prone (in queueing theory, this is also referred to as a server that "takes vacations"). As a result, the buffer content experiences fluctuations and occasionally overflows causing blocking phenomena that disrupt the smooth operation of the system and incur significant costs. To compensate, one can control the flow of parts into a buffer so as to maintain a satisfactory throughput while minimizing buffer overflow. A common approach is to formulate appropriate stochastic control problems so that control-theoretic techniques can be applied, as in [2],[3],[4],[5]. For certain problem formulations, under specific modeling assumptions, production control policies based on *thresholds* or *hedging points* have been identified as being optimal [4],[6],[7]; for a general overview and a recent survey see [8],[9]. Although in general such policies do not guarantee optimality, their implementation simplicity also makes them widely appealing in practice. These facts motivate us to further study their application to manufacturing systems. Unfortunately, the determination of optimal values for these hedging points is a difficult problem; see [2],[4],[7],[5]. In this paper,we address this particular problem, aiming at methodologies which can be applied on-line and without any knowledge of the stochastic characteristics of machine behavior or supply processes.

The manufacturing system we consider consists of a source supplying parts to a server with a buffer where the parts can be stored while awaiting to be processed. The source feeds the server with a variable rate which is controllable, but may be constrained. The server may be in one of two states: a "functional" or "on" state and an "unavailable" or "off" state; the latter represents the fact that the server may have failed or that it is simply unavailable to the source because it is busy processing different part buffers. The amount of time spent at each state is generally random and assumed not to depend on the buffer content. The goal of a production control policy is to maximize the throughput of the system, while minimizing the cost of buffer overflow, measured with respect to a given level B. Clearly, there is a trade-off between these two objectives.

⋆ Supported in part by the National Science Foundation under grants ACI-98-73339 and EEC-00-88073, by AFOSR under grant F49620-01-0056, by ARO under grant DAAD19-01-0610, and by the Air Force Research Laboratory under contract F30602-99-C-0057

Email addresses: fernyu@bu.edu (Haining Yu), cgc@bu.edu (Christos G. Cassandras).

A natural modeling framework for manufacturing systems is provided through queueing systems or, more generally, Discrete Event Systems (DES). A higher level of modeling abstraction is provided through fluid models which have been successfully adopted in the literature to analyze manufacturing systems (e.g., [2],[3],[4],[10],[11]). In this paper, we shall adopt a *Stochastic Fluid Model* (SFM) for the manufacturing system under study. We should stress that, in contrast to the fluid models mentioned above, a SFM captures *stochastic* fluctuations in the supply and service processes by treating all fluid rates as random processes. For the purpose of *performance analysis* the accuracy of SFMs (compared to a DES model) depends on factors such as the structure of the underlying system, and the nature of the performance metrics of interest. For the purpose of *control and optimization*, on the other hand, as long as a SFM captures the salient features of the underlying "real" system it is possible to obtain accurate solutions to problems even if we cannot estimate the corresponding performance with accuracy. This point of view was recently taken in [12] for the purpose of admission and flow control in communication networks. Using such a SFM, we shall invoke Infinitesimal Perturbation Analysis (IPA) methods [13],[14] adapted to fluid models (see also [15],[16],[12]). IPA is a well-developed approach that has been shown to yield unbiased gradient estimators for a class of DES which, unfortunately, does not include interesting phenomena such as blocking or overflow due to finite buffer capacities. Moreover, IPA can obviously not be applied to discrete parameters such as buffer capacities. For all these reasons, the use of IPA in this context has been limited. However, the use of SFMs opens up a broad new range of possibilities, as it allows us to derive unbiased estimators under very mild technical conditions. The simplicity and sometimes nonparametric nature of the estimators (see [12]) also renders them suitable for on-line control use.

The first contribution of the paper is to use IPA in order to derive unbiased gradient estimators of performance metrics related to throughput and buffer overflow with respect to threshold parameters (i.e., hedging points) which can be used for production control in the SFM under study. Compared to earlier work in [12],[15] that involves IPA of SFMs, note that establishing unbiasedness here is significantly more challenging due to the presence of the feedback effect of the threshold parameters, which complicates the SFM dynamics. In addition, the performance metrics considered are different than those in the SFMs studied in [12],[15]. As we shall see, the estimators obtained are simple to implement and are independent of the stochastic characteristics of the state switching process at the server or the supply process. Note that the analysis leading to the optimality of hedging point production control policies assumes that server state holding times are exponentially distributed; see [2],[3],[4],[6],[7]. Our analysis imposes no such limitations and also allows for stochastically varying server processing rates. In addition, the controllable supply flow rate may be constrained within arbitrarily time-varying lower and upper bounds. Finally, our analysis is also easily extendable to the case where the server has a finite number of different operating states. A significant difference between our approach here and the work in [4] is that we develop an *on-line* optimization scheme, while in [4] optimal hedging points are obtained *off line* through a set of nonlinear equations.

A second contribution of this paper is to make use of the IPA gradient estimators derived in order to tackle production control as an optimization problem. In particular, we seek to determine the hedging points that maximize a given performance metric of the SFM. This can be accomplished using a standard gradient-based stochastic optimization scheme, where we estimate the gradient of the performance function with respect to the hedging points. In addition, however, we propose an approximation method we apply to the actual (discrete-event) system from which the SFM is derived. This is based on the observation that the IPA gradient estimates obtained from the SFM can be evaluated from data directly observable on a sample path of the "real" system. Thus, we may use the SFM only to obtain the *form* of the gradient estimator; the associated *value* at any operating point is obtained from real system data. The result is of course only an approximate gradient estimator for the performance of this system, based on which we present several simulation-based results.

The paper is organized as follows. First, in Section 2, we present the system model and the basic flow control problem for the SFM. In Section 3, we derive IPA estimators for the throughput gradient with respect to the hedging points in the SFM and show that they are unbiased. In Section 4 we repeat this process for an overflow metric. In Section 5,we show how the SFM-based gradient estimators can be used on line for optimization purposes, including a proposed approximation method where they are applied to the actual system (not the SFM). Finally, in Section 6 we outline a number of open problems and future research directions motivated by this work.

2 System Model and Problem Formulation

We consider a Stochastic Fluid Model (SFM) for a manufacturing system as shown in Fig. 1. The server switches between two states, 0 (OFF) or 1 (ON). When in state $s(t) = 1$, the maximal processing rate of the server is generally time-varying and denoted by $\mu(t) > 0$; when $s(t) = 0$, the server is unavailable to the buffer either because it has failed or because it is busy providing service to different queues (not shown in Fig.

1). These switches occur randomly or they may be prescribed through some scheduling policy that assigns the server to the buffer for specific time intervals, as long as this assignment is not dependent on the state of the buffer. The buffer content at time t is denoted by $x(t)$. The buffer is fed by a source with controllable rate $u(t)$ chosen so that it always satisfies $0 < \lambda_{\min}(t) \leq u(t) \leq \lambda_{\max}(t)$, $0 \leq t \leq T$, where $\lambda_{\min}(t)$ and $\lambda_{\max}(t)$ are the *minimal* and *maximal feed rates* respectively, also generally time-varying.

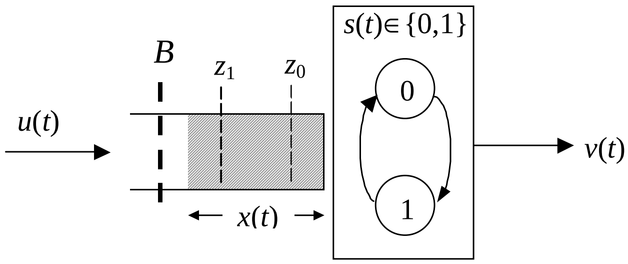

Fig. 1. Stochastic Fluid Model

We define a *threshold function* dependent on the server state $s(t)$ as follows:

$$z(s(t)) = \begin{cases} z_1 & \text{if } s(t)=1 \\ z_0 & \text{if } s(t)=0 \end{cases}$$

where z_0 and z_1 are threshold parameters to be determined. These parameters affect the input rate according to a *threshold-based production control policy*:

$$u(z_0,z_1;t) = \begin{cases} \lambda_{\max}(t) & \text{if } x(t) < z(t) \\ \lambda_{\min}(t) & \text{if } x(t) > z(t) \\ \max\{v(t), \lambda_{\min}(t)\} & \text{if } x(t) = z(t) \end{cases} \tag{1}$$

where the case $x(z_0, z_1; t) = z(s(t))$ is included to prevent "chattering" of the production rate between $\lambda_{min}(\tau)$ and $\lambda_{max}(\tau)$, whenever $x(\tau) = z_1$. Such chattering behavior is due to the nature of the SFM and does not occur in the actual discrete event system where buffer levels are maintained for finite periods of time (for details, see [17]). In (1), the notational dependence on the parameters z_0, z_1 indicates that we will analyze system performance metrics as functions of z_0, z_1. However, for simplicity, we will drop these arguments in the sequel and write $u(t)$ and $x(t)$ unless we need to stress the dependence on z_0, z_1. Similarly, we will write $z(t)$ instead of $z(s(t))$. We will assume that the real-valued parameters z_0, z_1 are confined to a closed and bounded (compact) interval $Z \subset \mathbb{R}^2$; to avoid unnecessary technical complications, we assume that $z_0, z_1 > 0$ for all $z_0, z_1 \in Z$. We also assume that $z_0 \leq z_1$; this condition is not necessary, but it is made so as to conform with the model used in [4].

Given (1), we can see that the output rate $v(t)$ is given by

$$v(t) = \begin{cases} \min\{\lambda_{\max}(t), \mu(t)\} & \text{if } s(t)=1 \text{ and } x(t)=0 \\ 0 & \text{if } s(t)=0 \\ \mu(t) & \text{otherwise} \end{cases} \tag{2}$$

and the buffer content $x(t)$ is determined by the following one-sided differential equation:

$$\frac{dx(t)}{dt^+} = u(t) - v(t) \tag{3}$$

with the initial condition $x(0) = x_0$ for some given x_0; for simplicity, we set $x_0 = 0$ throughout the paper.

To summarize, the underlying (uncontrollable) *rate processes* in the SFM are $\lambda_{\max}(t)$, $\lambda_{\min}(t)$, and $\mu(t)$. These are independent of the buffer content $x(t)$, the server state $s(t)$, and the threshold function $z(t) = z(s(t))$, which is a function of $s(t)$. We are interested in studying sample paths of this SFM over a time interval $[0,T]$ for a given fixed $0 < T < \infty$. We assume that the random processes $\{s(t)\}$, $\{\lambda_{\max}(t)\}$, $\{\lambda_{\min}(t)\}$, and $\{\mu(t)\}$ are independent of z_0, z_1 and they are right-continuous piecewise continuously differentiable w.p.1. Note that a typical sample path can be decomposed into two kinds of alternating intervals: *empty periods* and *buffering periods.* Empty Periods (EP) are intervals during which the buffer is empty, while Buffering Periods (BP) are intervals during which the buffer is nonempty. Observe that during an EP the system is not necessarily idle since the server may be active and processing at a rate that does not exceed $\mu(t)$, i.e., $\lambda_{\max}(t) \leq \mu(t)$, as seen in (2).

Let $\mathcal{L}(z_0, z_1) : Z \to \mathbb{R}$ be a random function defined over the underlying probability space $(\Omega, \mathcal{F}, P)$. When necessary, we shall also use the symbol ω to denote an observed sample path of the SFM, following standard conventions of the PA literature (e.g., [13],[14]) regarding the construction of sample functions. In what follows, we will consider two performance metrics, the *Throughput* $Q_T(z_0, z_1)$ and the *Overflow Rate* $R_T(z_0, z_1)$, both defined for a given interval $[0,T]$, defined as follows:

$$Q_T(z_0,z_1) = \frac{1}{T}\int_0^T u(t)dt \tag{4}$$

$$R_T(z_0,z_1) = \frac{1}{T}\int_0^T \mathbf{1}[x(t) \geq B]dt \tag{5}$$

in which $\mathbf{1}[\cdot]$ is the usual indicator function so that $\mathbf{1}[x(t) \geq B] = 1$ if $x(t) \geq B$ and 0 otherwise. Therefore, $Q_T(z_0, z_1)$ measures the average feed rate of the supply source and $R_T(z_0, z_1)$ is the fraction of time over $[0,T]$ during which the buffer level exceeds some given value B, which we assume to be such that $B > z_1 \geq z_0$. We can now formulate an optimization problem seeking the determination of $\mathbf{z}^* = [z_0^* \ z_1^*]$ that maximizes an objective function of the form:$J_T(z_0, z_1) = E[Q_T(z_0, z_1)] - \gamma E[R_T(z_0, z_1)]$ where γ is an overflow penalty per unit time. Since it is not feasible to obtain analytical solutions for this type of problem (except for very simple models), we rely on standard stochastic approximation algorithms [18], an approach that has been recently applied with success in the control and optimization of certain SFMs [12]. Details will be given in Section 5. The key to this approach is the efficient estimation of $\nabla J_T(\mathbf{z}_n) = [\partial J_T/\partial z_0 \ \partial J_T/\partial z_1]$, for which Perturbation Analysis (PA) methods [13],[14] are suitable, if appropriately adapted to a SFM viewed as a discrete-event system. In particular, we shall pursue the estimation of $\partial J_T/\partial z_0$ and $\partial J_T/\partial z_1$ through Infinitesimal Perturbation Analysis(IPA) techniques, similar to [12],[15].

3 IPA for Throughput

3.1 Notation and Definitions

Viewing a SFM as a discrete-event system, an *event* in a sample path of the above SFM may be either *exogenous* or *endogenous*. An *exogenous event* occurs at any time instant corresponding to (i) a jump in $s(t)$, which affects the threshold function $z(s(t))$ and the output rate $v(t)$, or (ii) a change in the partial order of the values of $\mu(t)$, $\lambda_{\max}(t)$ and $\lambda_{\min}(t)$, which affects the sign of $u(t)-v(t)$ in (3). For example, an exogenous event occurs at time β if $\mu(t) > \lambda_{\max}(t) > \lambda_{\min}(t)$ when $t < \beta$, and $\lambda_{\max}(t) > \mu(t) > \lambda_{\min}(t)$ when $t \geq \beta$. Clearly, the occurrence times of all exogenous events are independent of z_0 and z_1.

An *endogenous event* occurs when the value of the controllable input rate changes from one to any other value among $\mu(t)$, $\lambda_{\max}(t)$ and $\lambda_{\min}(t)$ in (1). In particular, there are four kinds of endogenous events that we are interested in and define them as follows:

(1) $z^\uparrow$ event: This occurs when $x(t)$ is increasing (which implies that $\lambda_{\max}(t) \geq \lambda_{\min}(t) > \mu(t)$) and crosses the $z(s(t))$ level from below.
(2) $z^\downarrow$ event: This occurs when $x(t)$ is decreasing (which implies $\mu(t) > \lambda_{\max}(t) \geq \lambda_{\min}(t)$) and crosses the z_1 level from above. Note that a $z^\downarrow$ event cannot occur when $s(t) = 0$, since in this case we have $v(t) = 0$ and, by (3) and (1), the buffer content may only be nondecreasing.
(3) z^+ event: This occurs when $x(t)$ hits $z(s(t))$ from below, i.e., $x(t^-) < z(s(t^-))$, and $x(t) = z(s(t))$) for some time interval of strictly positive length (which, as seen in (1), implies that $\lambda_{\max}(t) > \mu(t) > \lambda_{\min}(t)$).
(4) z^- event: This occurs when $x(t)$ hits $z(t)$ from above, i.e., $x(t^-) > z(s(t))$, and $x(t) = z(s(t))$ for some time interval of strictly positive length (which, again, implies that $\lambda_{\max}(t) > \mu(t) > \lambda_{\min}(t)$).

For notational economy, we shall henceforth use $\mathbf{z} = [z_0, z_1]$. We shall also make **Assumption 1**: *For every* $\mathbf{z} = [z_0 \ z_1] \in \mathbf{Z}$, *w.p. 1, no two events occur at the same time.* As already mentioned, a sample path of the SFM may be decomposed into Empty Periods (EP) and Buffering Periods (BP). Let us assume there are K BPs in the sample path, where K is a random number which, because of **Assumption 1**, is locally independent of $\mathbf{z}$ (since no two events may occur simultaneously, there exists a neighborhood of $\mathbf{z}$ within which, w.p.1, the number of BPs in $[0, T]$ is constant). Given the kth BP, denoted by $\mathcal{B}_k$, $k = 1, \ldots, K$, its starting point, denoted by α_k, occurs when the buffer ceases to be empty; this corresponds to an exogenous event causing a change in the sign of $u(t) - v(t)$ in (3) from nonpositive to strictly positive; therefore α_k is locally independent of $\mathbf{z}$ under **Assumption 1**. The end point of $\mathcal{B}_k$, however, generally depends on $\mathbf{z}$. Denoting this end point by $\eta_k(\mathbf{z})$, we express $\mathcal{B}_k$ as $\mathcal{B}_k = [\alpha_k, \eta_k(\mathbf{z}))$, $k = 1, \ldots, K$. Figure 2 shows a typical sample path, including two BPs and one EP. There are three endogenous events in this example: two $z^\uparrow$ events when $x(t)$ crosses z_0 from below with $s(t) = 0$ at time τ_1, τ_3, and a $z^\downarrow$ event when $x(t)$ crosses z_1 from above with $s(t) = 1$ at time τ_2. Recall that when the buffer content $x(t)$ crosses the threshold $z(t)$ (either z_0 or z_1 depending on $s(t)$), the feed rate from the supply source switches from $\lambda_{\max}(t)$ to $\lambda_{\min}(t)$ (or vice versa), causing a jump in $\dot{x}(t)$. Note also that $x(t)$ is not necessarily piecewise linear, reflecting the time-varying nature of $\lambda_{\max}(t), \lambda_{\min}(t)$ and $\mu(t)$ and in contrast to other, simpler, fluid models (e.g., [2],[4],[11]). An overflow period $[\delta_1, \xi_1)$ is also included in the sample path of Fig. 2.

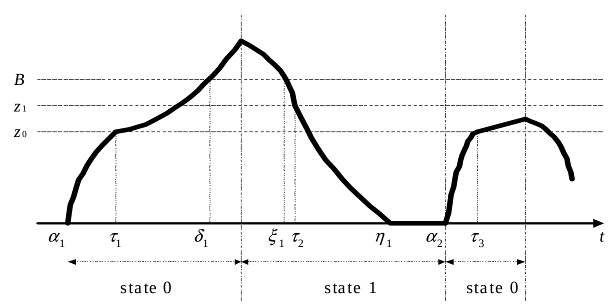

Fig. 2. A Typical Sample Path

The BPs can be classified according to whether they include any *endogenous* event (as defined above) or not. Thus, we define the random set $\Phi(\mathbf{z}) := \{k \in \{1, \ldots, K\} : x(t) = z(s(t)), u(t) - v(t) > 0$ for some $t \in (\alpha_k, \eta_k(\mathbf{z}))\}$ since the first possible endogenous event in a BP is one occurring when $z(s(t))$ is reached from below, i.e., a $z^\uparrow$ or z^+ event. For each $k \in \Phi(\mathbf{z})$, let $L_k \geq 1$ be the (random) number of endogenous events contained in $\mathcal{B}_k$, and let $\{\tau_{k,j}(\mathbf{z})\}$ denote their corresponding occurrence times, $j = 1, \ldots, L_k$, which are generally dependent on $\mathbf{z}$.

Next, consider $[\alpha_k, \alpha_{k+1})$, i.e., an interval consisting of $\mathcal{B}_k$ and the ensuing EP $[\eta_k(\mathbf{z}), \alpha_{k+1})$. For each such $k = 1, \ldots, K$, there is a (random) number $N_k \geq 1$ of *exogenous* events in $[\alpha_k, \alpha_{k+1})$, and let $\{\beta_{k,i}\}$, $i = 1, 2, \ldots, N_k$ denote the sequence of all corresponding event times. We observe that two endogenous events cannot occur consecutively without an exogenous event between them. Therefore, $L_k \leq N_k$, for all $k = 1, \ldots, K$. For simplicity, we also define $\alpha_k = \beta_{k,1}$ and $\alpha_{k+1} = \beta_{k,N_k+1}$. Finally, for any $k \in \Phi(\mathbf{z})$, we define the index set for $z^{\uparrow}$ events: $\Phi_1^k(\mathbf{z}) := \{j \in \{1, ..., L_k\} : \text{a } z^{\uparrow} \text{ event occurs at time } \tau_{k,j}(\mathbf{z}) \in (\alpha_k, \eta_k(\mathbf{z}))\}$. Similarly we define $\Phi_2^k(\mathbf{z})$, $\Phi_3^k(\mathbf{z})$ and $\Phi_4^k(\mathbf{z})$ for $z^{\downarrow}$, z^{+} and z^{-} events respectively. It is important to keep in mind that only endogenous event times $\tau_{k,j}(\mathbf{z})$ generally depend on $\mathbf{z}$. In the sequel, however, we shall drop this explicit dependence and write $\tau_{k,j}$.

Returning to (4), and using the definitions of $\{\alpha_k\}$ and $\{\beta_{k,i}\}$, we may write

$$Q_T(\mathbf{z}) = \frac{1}{T}\int_0^T u(t)dt = \frac{1}{T}\sum_{k=1}^{K}\int_{\alpha_k}^{\alpha_{k+1}} u(t)dt$$

and by further decomposing each $[\alpha_k, \alpha_{k+1}) \equiv [\beta_{k,1}, \beta_{k,N_k+1})$ into N_k intervals $[\beta_{k,1}, \beta_{k,2}), \cdots, [\beta_{k,N_k}, \beta_{k,N_k+1})$, we get

$$Q_T(\mathbf{z}) = \frac{1}{T}\sum_{k=1}^{K}\sum_{i=1}^{N_k}\int_{\beta_{k,i}}^{\beta_{k,i+1}} u(t)dt \qquad (6)$$

where we note again that all $\beta_{k,i}$ are locally independent of $\mathbf{z}$. Note that there is at most one endogenous event in any interval $[\beta_{k,i}, \beta_{k,i+1})$.

Define $q_{k,i}(\mathbf{z}) = \int_{\beta_{k,i}}^{\beta_{k,i+1}} u(t)dt$. We get

$$Q_k(\mathbf{z}) = \int_{\alpha_k}^{\alpha_{k+1}} u(t)dt = \sum_{i=1}^{N_k}\int_{\beta_i}^{\beta_{i+1}} u(t)dt = \sum_{i=1}^{N_k} q_{k,i}(\mathbf{z})$$

Assuming that the derivatives $\partial q_{k,i}(\mathbf{z})/\partial z_r$, $r = 0, 1$, exist (we return to this issue later in this section), and recalling that N_k is the number of exogenous events in $[\alpha_k, \alpha_{k+1})$ and is therefore locally independent of $\mathbf{z}$, it follows that $\frac{\partial Q_k(\mathbf{z})}{\partial z_r} = \sum_{k=1}^{N_k}\frac{\partial q_{k,i}(\mathbf{z})}{\partial z_r}$, $r = 0, 1$, and from (6):

$$\frac{\partial Q_T(\mathbf{z})}{\partial z_r} = \frac{1}{T}\sum_{k=1}^{K}\frac{\partial Q_k(\mathbf{z})}{\partial z_r}, \quad r = 0, 1 \qquad (7)$$

Therefore, the evaluation of the sample derivatives $\partial Q_T(\mathbf{z})/\partial z_r$ reduces to evaluating $\partial q_{k,i}(\mathbf{z})/\partial z_r$.

3.2 Sample Derivatives

In this section, we show that the evaluation of the sample derivative $\partial Q_T(\mathbf{z})/\partial z_r$ requires the derivatives $\partial\tau_{k,j}/\partial z_r$ of the event times $\tau_{k,j}$ within some $\mathcal{B}_k$, for all $k = 1, \ldots, K$. In what follows, we shall concentrate on a typical $\mathcal{B}_k = [\alpha_k, \eta_k(\mathbf{z}))$ and drop the index k from $Q_k(\mathbf{z}), q_{k,i}(\mathbf{z}), \Phi_i^k(\mathbf{z})$ for $i = 1, \ldots, 4$, as well as for all $\tau_{k,j}$, $\beta_{k,i}$ in order to simplify notation. We obtain for $r = 0, 1$, the following lemma (proofs of all lemmas in the paper may be found in [17]).

Lemma 3.1

$$\frac{\partial Q(\mathbf{z})}{\partial z_r} = \sum_{j\in\Phi_1} A_j\frac{\partial\tau_j}{\partial z_r} - \sum_{j\in\Phi_2} A_j\frac{\partial\tau_j}{\partial z_r} + \sum_{j\in\Phi_3} B_j\frac{\partial\tau_j}{\partial z_r} + \sum_{j\in\Phi_4} C_j\frac{\partial\tau_j}{\partial z_r} \qquad (8)$$

in which $A_j \equiv \lambda_{\max}(\tau_j) - \lambda_{\min}(\tau_j)$, $B_j \equiv \lambda_{\max}(\tau_j) - v(\tau_j)$, $C_j \equiv \lambda_{\min}(\tau_j) - v(\tau_j)$.

The derivatives $\partial\tau_j/\partial z_r$ exist as long as $\tau_j(\mathbf{z})$ is not a jump point of the function $u(t) - v(t)$, which can be guaranteed if $\tau_j(\mathbf{z})$ is not a jump point of $s(t)$, $\lambda_{\max}(t)$, $\lambda_{\min}(t)$, or $\mu(t)$. This, in turn, is ensured by **Assumption 1**. We will also need two mild technical conditions, i.e., **Assumption 2**: *$0 < \lambda_{\min}(t) \leq \lambda_{\max}(t) \leq c_1 < \infty$ and $\mu(t) \leq c_1 < \infty$, w.p. 1, for some positive constant c_1 and for all $t \in [0, T]$*, and **Assumption 3**: *There exists a positive constant c_2 such that, w.p. 1, $|\lambda_{\max}(t) - \mu(t)| \geq c_2$ and $|\lambda_{\min}(t) - \mu(t)| \geq c_2$ for all $t \in [0, T]$.* Returning to (8), note that the sets $\Phi_1, \cdots, \Phi_4$ are also locally independent of $\mathbf{z}$ due to **Assumption 1**. Therefore, the existence of $\partial Q(\mathbf{z})/\partial z_r$ is ensured, and so is that of the sample throughput derivative $\partial Q_T(\mathbf{z})/\partial z_r$ in (7).

3.3 Endogenous Event Time Derivatives $\partial\tau_j/\partial z_0$ and $\partial\tau_j/\partial z_1$

As seen in Lemma 3.1, the form of the sample derivatives $\partial Q(\mathbf{z})/\partial z_r$ requires the event time derivatives $\partial\tau_j/\partial z_r, r = 0, 1$, the evaluation of which is the central task of this section. Let e_j, e_{j+1} denote the endogenous events that occur at τ_j and τ_{j+1} respectively, with $e_j, e_{j+1} \in \{z^{\uparrow}, z^{\downarrow}, z^{+}, z^{-}\}$. The main result, stated as Lemma 3.2, shows that $\partial\tau_j/\partial z_r$ can be evaluated through a simple linear recursion over $j = 1, \ldots, L$ depending only on B_j, C_j defined in Lemma 3.1.

Lemma 3.2 *Let τ_j, $j = 1, \ldots, L$, be the endogenous event times in a BP and let $e_j \in \{z^{\uparrow}, z^{\downarrow}, z^{+}, z^{-}\}$ be the* *corresponding event. Then, for* $j = 1, \ldots, L-1$,

$$\frac{\partial \tau_{j+1}}{\partial z_r} = \frac{F_j}{G_{j+1}} \frac{\partial \tau_j}{\partial z_r} + \frac{1}{G_{j+1}} \varphi_{j+1}(z_r), \tag{9}$$
$$\frac{\partial \tau_1}{\partial z_r} = \frac{1}{B_1} \frac{\partial z(\tau_1)}{\partial z_r}$$

where $\varphi_{j+1}(z_r) \equiv \frac{\partial z(\tau_{j+1})}{\partial z_r} - \frac{\partial z(\tau_j)}{\partial z_r} \in \{-1, 0, 1\}$, $r = 0, 1$, *and*

$$F_j = \begin{cases} B_j & \text{if } e_j = z^{\downarrow} \\ C_j & \text{if } e_j = z^{\uparrow} \\ 0 & \text{otherwise} \end{cases},$$
$$G_{j+1} = \begin{cases} B_{j+1} & \text{if } e_{j+1} = z^{\uparrow}, z^{+} \\ C_{j+1} & \text{otherwise} \end{cases} \tag{10}$$

Thus, within a BP, $\partial \tau_{j+1}/\partial z_r$, $r = 0, 1$, is readily evaluated through (9) upon observing the endogenous event e_{j+1} and using $\partial \tau_j/\partial z_r$ along with the rate information $\lambda_{\max}(t)$, $\lambda_{\min}(t)$, $\mu(t)$ at $t = \tau_j$ and τ_{j+1}, which are needed to calculate F_j, G_{j+1} depending on the event types e_j, e_{j+1}. In addition, the information $s(\tau_j)$ and $s(\tau_{j+1})$ allows us to evaluate $z(s(\tau_j))$ and $z(s(\tau_{j+1}))$, from which $\varphi_{j+1}(z_r) \in \{-1, 0, 1\}$ is immediately obtained.

We can now summarize the IPA estimator $\partial Q_T(\mathbf{z})/\partial z_r$ of the performance derivative $\partial E[Q_T(\mathbf{z})]/\partial z_r$, $r = 0, 1$, as follows: (i) Within each observed BP $\mathcal{B}_k$ in $[0, T]$, evaluate $\partial \tau_{k,j}/\partial z_r$ for all $j = 1, \ldots, L_k$ through (9), (ii) At the end of $\mathcal{B}_k$, evaluate $\partial Q_k(\mathbf{z})/\partial z_r$ through (8), and (iii) At time T, the IPA estimator is given by (7). Note that, except for the lower and upper bounds of the supply rate and the service rate *at endogenous event time instants only*, no other information regarding the service or supply processes is involved. In the case where $\lambda_{\max}, \lambda_{\min}, \mu$ are time-invariant and known over $[0, T]$, the IPA estimator becomes extremely simple to implement since A_j, B_j, and C_j are reduced to known constants.

3.4 Unbiasedness of IPA Estimators

In this section the unbiasedness of the IPA estimators $\partial Q_T(\mathbf{z})/\partial z_r$ is established. An IPA-based estimate $\partial \mathcal{L}(\theta)/\partial \theta$ of a performance metric derivative $dE[\mathcal{L}(\theta)]/d\theta$ is *unbiased* if $dE[\mathcal{L}(\theta)]/d\theta = E[\partial \mathcal{L}(\theta)/\partial \theta]$. Unbiasedness is the principal condition for making the application of IPA useful in practice, since it enables the use of the sample (IPA) derivative in control and optimization methods that employ stochastic gradient-based techniques. Note that we are only interested in gradient estimation over a finite interval $[0, T]$, so that we do not concern ourselves here with the issue of estimator consistency (as, for instance, in [11]). In general, the unbiasedness of an IPA derivative of some sample function $\mathcal{L}(\theta)$ with respect to θ has been shown to be ensured by the following two conditions (see [19], Lemma A2, p.70): (i) For every $\theta \in \Theta$, the sample derivative of $\mathcal{L}(\theta)$ exists w.p.1, (ii) 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 $\partial Q_T(\mathbf{z})/\partial z_r$ as estimators of $\partial E[Q_T(\mathbf{z})]/\partial z_r$, $r = 0, 1$, reduces to verifying the Lipschitz continuity of $Q_T(\mathbf{z})$ with appropriate Lipschitz constants. We accomplish this in two steps. First we prove the that any buffer content perturbation $\Delta x(t)$ resulting from a parameter perturbation Δz_r, $r = 0, 1$, is bounded so that $0 \le \Delta x(t) \le \Delta z_r$ for all $t \in [0, T]$. Next, unbiasedness is established in Theorem 3.1.

We begin by defining the *buffer content perturbation* at time t due to a perturbation Δz_r in z_r, $r = 0$ or 1. For simplicity, let us limit ourselves to a perturbation $\Delta z_0 > 0$; the cases where $\Delta z_0 < 0$ or z_1 is perturbed instead of z_0 are similarly analyzed. Thus, set $\Delta x(t) = x(\mathbf{z} + \Delta \mathbf{z}_0; t) - x(\mathbf{z}; t)$ where $\Delta \mathbf{z}_0 = [\Delta z_0 \; 0]$. To keep notation simple, we also denote the nominal buffer content $x(\mathbf{z}; t)$ by $x(t)$ and the perturbed one by $x'(t) = x(\mathbf{z} + \Delta \mathbf{z}_0; t)$

Lemma 3.3 *Assuming a perturbation* $\Delta z_0 > 0$, *the buffer content perturbation* $\Delta x(t)$ *satisfies:*

$$0 \le \Delta x(t) \le \Delta z_0 \quad \textit{for all}\; t \in [0, T]$$

The proof of this lemma (see [17]) involves no specific information about the server state $s(t)$. Therefore, the same bound can be proved when the perturbed sample path is due to Δz_1, instead of Δz_0, in the same way. The same is true for a model in which more than two server states are defined.

Theorem 3.1 *Let* $N(T)$ *be the number of all exogenous events in* $[0, T]$ *and assume* $E[N(T)] < \infty$. *Then,* $\partial Q_T(\mathbf{z})/\partial z_0$ *in (8) is an unbiased IPA estimator of* $\partial E[Q_T(\mathbf{z})]/\partial z_0$.

Proof. See Appendix.

4 IPA for Overflow Rate

We now turn our attention to the second performance metric of interest defined in (5). We introduce two additional endogenous events:

(1) $B^{\uparrow}$ event: This occurs when $x(t)$ is increasing (which implies that $\lambda_{\max}(t) \ge \lambda_{\min}(t) > \mu(t)$) and crosses the B level from below. Recall that $B > z_1 \ge z_0$.

(2) $B^{\downarrow}$ event: This occurs when $x(t)$ is decreasing (which implies $\mu(t) > \lambda_{\max}(t) \geq \lambda_{\min}(t)$) and crosses the B level from above.

Note that $B^{\uparrow}$ and $B^{\downarrow}$ events cannot occur in sequence without an exogenous event between them to cause a sign change in $u(t) - v(t)$ in (3). We define an *overflow interval* as an interval that starts with a $B^{\uparrow}$ event and ends with the next $B^{\downarrow}$ event. We can then define the random set $\Phi(\mathbf{B}) := \{k \in \{1, \ldots, K\} \; : \; x(t) = B \text{ for some } t \in (\alpha_k, \eta_k(\mathbf{z}))\}$ which includes all BPs with at least one overflow interval. Let $n = 1, \ldots, O_k$ index these overflow intervals for any $k \in \Phi(\mathbf{B})$ and let $I_{k,n}$ denote the nth overflow interval in the kth BP. The starting and ending time of $I_{k,n}$ are denoted by $\delta_{k,n}$ and $\xi_{k,n}$ respectively. From (5) we can write:

$$R_T(\mathbf{z}) = \frac{1}{T} \sum_{k \in \Phi(B)} \sum_{n=1}^{O_k} (\xi_{k,n} - \delta_{k,n}) \qquad (11)$$

Therefore, for $r = 0, 1$,

$$\frac{\partial R_T}{\partial z_r} = \frac{1}{T} \sum_{k \in \Phi(B)} \sum_{n=1}^{O_k} \left(\frac{\partial \xi_{k,n}}{\partial z_r} - \frac{\partial \delta_{k,n}}{\partial z_r}\right) \qquad (12)$$

where we note that $\Phi(\mathbf{B})$ and O_k are locally independent of $\mathbf{z}$ by **Assumption 1**. Proceeding as in Section 3.3, we now seek to determine the derivatives of the event times $\delta_{k,n}$ and $\xi_{k,n}$ with respect to z_r, $r = 0, 1$. For simplicity, let us concentrate on a specific BP and drop the index k.

Proceeding as in Section 3.3., it is straightforward to determine the derivatives $\frac{\partial \delta_{k,n}}{\partial z_r}$ and $\frac{\partial \xi_{k,n}}{\partial z_r}$, $r = 0, 1$ (see [17]). Next, we establish that this estimator is unbiased. As in the case of the throughput, we limit ourselves to $\partial R_T(\mathbf{z})/\partial z_0$ since the proof is independent of $s(t)$ and therefore applies to $\partial R_T(\mathbf{z})/\partial z_1$ as well.

Theorem 4.1 *Let $N(T)$ be the number of all exogenous events in $[0, T]$ and assume $E[N(T)] < \infty$. Then, $\partial R_T(\mathbf{z})/\partial z_0$ in (12) is an unbiased IPA estimator of $\partial E[R_T(\mathbf{z})]/\partial z_0$.*

Proof. See Appendix.

5 Optimal Control Using IPA Estimators

Let us return to the optimization problem presented in Section 2. Our first objective is to determine a pair of thresholds (z_0, z_1) so as to maximize a function

$$J_T^{SFM}(z_0, z_1) = E\left[Q_T^{SFM}(z_0, z_1)\right] - \gamma E\left[R_T^{SFM}(z_0, z_1)\right]$$

reflecting the trade-off between the throughput and overflow rate over $[0, T]$ in a SFM. The superscript SFM is added here to stress the fact that the expected performance above is defined on a SFM, rather than the underlying DES that we consider later in this section. We implement a standard Stochastic Approximation (SA) algorithm

$$\mathbf{z}_{n+1} = \mathbf{z}_n + \nu_n \mathbf{H}_n(\mathbf{z}_n; \omega_n^{SFM}), \;\; n = 0, 1, \ldots \qquad (13)$$

where $\{\nu_n\}$ is a step-size sequence, and the gradient estimator $\mathbf{H}_n(\mathbf{z}_n; \omega_n^{SFM})$ is the IPA estimator of $\nabla J_T^{SFM}(\mathbf{z}_n)$ evaluated over a simulated sample path of the SFM, denoted by ω_n^{SFM}, of length T. The output of this algorithm is denoted by $\mathbf{z}^*_{SFM}$ and is expected to converge to the optimal solution of the above optimization problem under certain standard technical conditions; details on SA algorithms, including conditions required for convergence to an optimum (generally local, unless the form of the cost functions ensures the existence of a single optimum) may be found, for example, in [18].

As mentioned earlier, our work is partly motivated by [4], where a fluid model with fixed $\lambda_{\max}$ and $\lambda_{\min}$ and a fixed-rate server spending an exponentially distributed amount of time in state 0(OFF) or 1(ON) has been studied. For this model, the optimal hedging point pair can be determined through a set of nonlinear equations given in [4]. We denote the solution to these equations by $\mathbf{z}^*_{theo}$. Our second objective in this section is to compare $\mathbf{z}^*_{SFM}$ to $\mathbf{z}^*_{theo}$ for a case where we reduce our SFM to the simpler fluid model in [4].

Our third objective is to determine a pair of thresholds (z_0, z_1) so as to maximize a function

$$J_T^{DES}(z_0, z_1) = E\left[Q_T^{DES}(z_0, z_1)\right] - \gamma E\left[R_T^{DES}(z_0, z_1)\right]$$

where the superscript DES indicates that the expected throughput and overflow rate above are now defined on a DES. A natural way to proceed in this case is to resort again to a SA algorithm of the form

$$\mathbf{z}_{n+1} = \mathbf{z}_n + \nu_n \mathbf{H}_n(\mathbf{z}_n; \omega_n^{DES}), \;\; n = 0, 1, \ldots \qquad (14)$$

where ω_n^{DES} represents a sample path of the DES of length T. Since we have no means of deriving an unbiased estimator for the gradient of $J_T^{DES}(z_0, z_1)$, we shall make use of an approximation, $\mathbf{H}_n(\mathbf{z}_n; \omega_n^{DES})$, motivated by the following observation. Notice that the form of the SFM-based IPA estimators we have derived enables their values to be obtained from data of an *actual* (discrete-event) system: The expressions in (9), (8), and (7) for the Throughput IPA estimator simply require (i) detecting when the buffer level crosses $z(s(t))$ given the observed server state $s(t)$, and (ii) the values of flow rates at these instants so as to evaluate A_j, B_j, C_j in Lemma 3.1; similarly, for the Overflow Rate IPA estimator in (12). In other words, the *form* of the IPA estimators is obtained by analyzing the system as a SFM, but the associated *values* can be obtained from real data from the underlying DES. Obviously, the resulting gradient estimator, denoted by $\mathbf{H}_n(\mathbf{z}_n;\omega_n^{DES})$ in (14), is merely an approximation. The output of (14) is denoted by $\mathbf{z}_{DES}$ and represents a sub-optimal solution of the above optimization problem, based on the heuristic described above. Obtaining the actual solution is generally very hard. For our purposes, we have used exhaustive simulation under all possible values of integer-valued pairs (z_0, z_1) over a given range to estimate this solution; we denote the result by z^*_{DES}. Our objective then is to compare the output $\mathbf{z}_{DES}$ of (14) to z^*_{DES}, as well as to $\mathbf{z}^*_{SFM}$.

We point out that in obtaining z^*_{DES} the response surface was estimated by initializing each simulation under a different (z_0, z_1) pair with the same random seed (for variance reduction purposes). However, in implementing the SA algorithms (13) and (14), our goal was to emulate an on-line controller which does not have the luxury of a common random number approach; therefore, no such action was taken. Upon completion of one iteration of the algorithm, the next iteration was carried out from the final state of the last one.

The two examples that follow are referred to as Scenario 1 and 2 respectively. For both scenarios, the supply process for the DES simulated is Poisson with rate $\lambda_{\max}(t)$ or $\lambda_{\min}(t)$. The server remains in the ON state for an exponentially distributed amount of time with rate q_u and in the OFF state for an exponentially distributed amount of time with rate q_d. In the ON state, the service rate is μ. In the OFF state, the server does not work. Finally, in both cases $T = 5,000,000$ time units and the step sequence $\{\nu_n\}$ in (14) is selected so that $\nu_n = \frac{1}{50 \times n^{0.99}}$, $n = 1, 2, \ldots$

Scenario 1: $B = 100$, $\mu = 20$, $\lambda_{\max} = 12$, $\lambda_{\min} = 6$, $q_u = 0.1$, $q_d = 0.14$, $\gamma = 9$.

In this case, $\lambda_{\max}(t)$ and $\lambda_{\min}(t)$ are fixed over time and the service rate μ is constant. This enables us to calculate the optimal hedging point pair for the corresponding fluid model through the set of nonlinear equations provided in [4]. As already mentioned, the solution to this equation set is denoted by $\mathbf{z}^*_{theo}$ and we found $\mathbf{z}^*_{theo} = (19.91, 77.51)$. In addition, we implemented (13) by simulating the SFM and obtained $\mathbf{z}^*_{SFM} = (19.81, 78.37)$.

Through exhaustive discrete-event simulation of this system, the response surface of the objective function $J_T^{DES}(\mathbf{z})$ was obtained as shown in Fig. 3 and the associated optimal point was found to be $\mathbf{z}^*_{DES} = (21, 80)$. We then implemented (14) by simulating the DES and obtained $\mathbf{z}_{DES} = (20.34, 78.20)$.

Scenario 2: $B = 100$, $\mu = 20$, $\lambda_{\max,0} = 22$, $\lambda_{\max,1} = 4$, $\lambda_{\min,0} = 20.002$, $\lambda_{\min,1} = 2$, $q_u = 0.1$, $q_d = 0.15$,

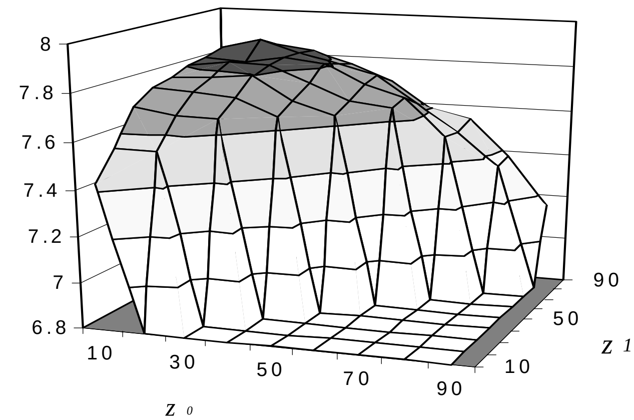

Fig. 3. Objective Function $J_T(z_0, z_1)$ (Scenario 1)

$q_{l,0} = 0.025$, $q_{l,1} = 0.005$, $\gamma = 12$.

In this scenario, $\lambda_{\max}$ ($\lambda_{\min}$) switches between $\lambda_{\max,0}$ and $\lambda_{\max,1}$ ($\lambda_{\min,0}$ and $\lambda_{\min,1}$). The time interval over which $\lambda_{\max} = \lambda_{\max,0}$ ($\lambda_{\min} = \lambda_{\min,0}$) is exponentially distributed with rate $q_{l,0}$, and the time interval over which $\lambda_{\max} = \lambda_{\max,1}$ ($\lambda_{\min} = \lambda_{\min,1}$) is exponentially distributed with rate $q_{l,1}$. By allowing $\lambda_{\max}(t)$ and $\lambda_{\min}(t)$ to be random processes, we illustrate the use of our approach to systems with complex rate processes beyond the fixed ones found in [4]. Since $\lambda_{\max}(t)$ and $\lambda_{\min}(t)$ are no longer fixed, the method for determining an optimal hedging point pair $\mathbf{z}^*_{theo}$ for the corresponding fluid model provided in [4] cannot be used. All other notation here is the same as that of Scenario 1. The service time is still exponentially distributed with rate μ.

Using the SFM for this scenario, we determined $\mathbf{z}^*_{SFM} = (63.74, 71.95)$ through (13). Using the DES, we found $\mathbf{z}_{DES} = (63.49, 74.79)$ through (14) and by exhaustive simulation we found $\mathbf{z}^*_{DES} = (68, 80)$. Similar to scenario 1, the response surface of the objective function $J_T^{DES}(\mathbf{z})$ is shown in Fig. 4. In Fig. 5 we show the convergence behavior of $J_T^{SFM}(\mathbf{z}_n)$ for the optimization of the SFM using (13) in the curve labeled 'SFM', and of $J_T^{DES}(\mathbf{z}_n)$ for the optimization of the DES using (14) in the curve labeled 'DES'. In the same figure, we also show the value of $J_T^{DES}(\mathbf{z}^*_{DES})$, labeled 'OPT'. The corresponding convergence behavior of the two hedging points is shown in Fig. 6 using the same notation. Note that because the response surfaces are relatively insensitive to small changes in z_0, z_1 in the vicinity of the optimal point, there is a set of hedging points all yielding performances that are hard to distinguish in the presence of noise in Figs. 3 and 4. This is consistent with the observation that the objective function is quite robust with respect to the two hedging point parameters (e.g., in Fig. 4, the performance range is limited to about 20% of the attainable optimal value).

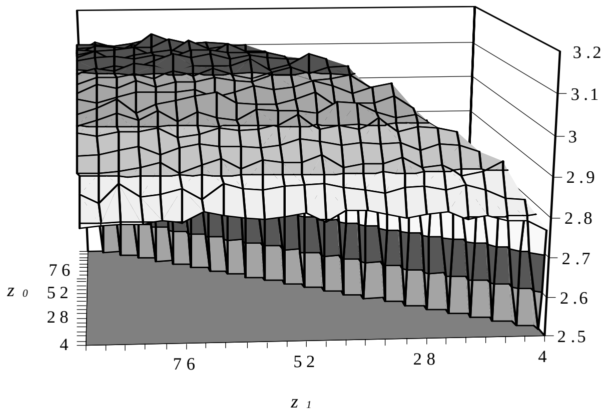

Fig. 4. Objective Function $J_T(z_0, z_1)$ (Scenario 2)

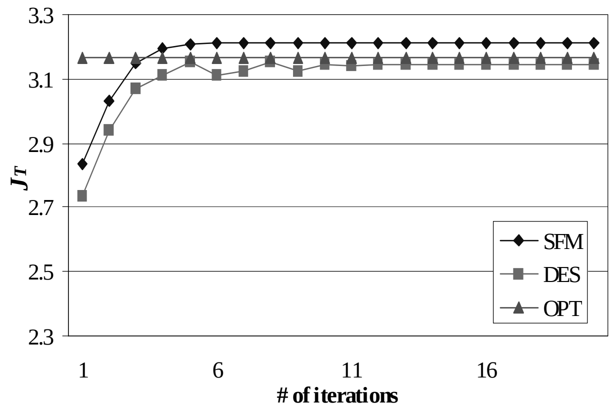

Fig. 5. Objective Function Convergence (Scenario 2)

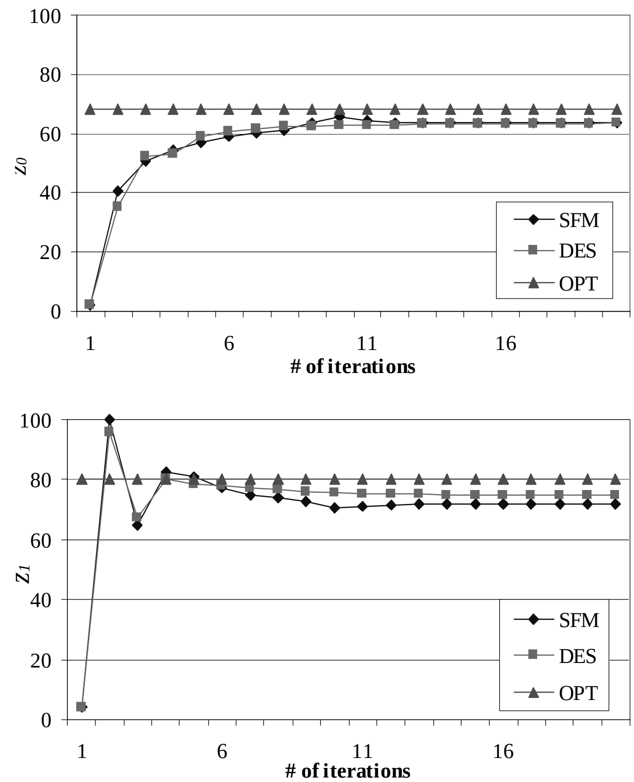

Fig. 6. Hedging Point Convergence (Scenario 2)

6 Conclusions and Future Work

We have adopted Stochastic Fluid Models (SFM) for control and optimization of manufacturing systems in order to capture the main features of production control policies without requiring a detailed discrete-event model for our analysis. Our approach is based on the observation that a SFM can be used to accurately determine optimal settings for control and optimization purposes, even when it fails to provide adequately accurate performance estimates. We also stress that using a detailed queueing model for the purpose of adjusting integer-valued hedging points leads to *discrete* stochastic optimization and the use of elaborate *finite* PA methods, a task that becomes highly and, as our work suggests, unnecessarily complex in practice.

In this paper, we limit ourselves to a threshold-based flow control policy in which the objective is to adjust the threshold parameters (hedging points) so as to optimize an objective function combining throughput and overflow rate metrics. For a single workcenter model, we derive IPA gradient estimators based on the SFM, show them to be unbiased, and subsequently use them for optimization purposes. Exploiting the simple structure of these estimators, we have also proposed an approximation method aimed at optimizing a similar objective function for the actual single server system. Although our current work is limited to a single workcenter model, we believe that the nature of the IPA estimators enables extending the analysis to multiple workcenters in series, modeled as SFMs. Indeed, recent work [20] has achieved such results in tandem networks with similar settings, but in the absence of feedback. In addition, we believe that the same modeling framework may be used to study scheduling policies in which the server is shared by multiple competing buffers, by developing IPA estimators with respect to parameters that determine the amount of time that any given buffer sees the server at the ON state.

Appendix

Proof of Theorem 3.1. Recall (4) and let $u'(t) = u(\mathbf{z}+\Delta\mathbf{z}_0;t)$ and $Q'_T = Q_T(\mathbf{z}+\Delta\mathbf{z}_0)$ denote the input flow and throughput respectively when $\mathbf{z}$ is perturbed by $\Delta\mathbf{z}_0 = [\Delta z_0 \ 0]$. Thus,

$$\Delta Q_T = \frac{1}{T}\int_0^T \Delta u(t)dt \tag{15}$$

where we set $\Delta u(t) = u'(t) - u(t)$. Using (3), and setting $v'(t) = v(\mathbf{z} + \Delta\mathbf{z}; \mathbf{t})$, we obtain:

$$x(T) = x(0) + \int_0^T [u(t) - v(t)]\, dt,$$
$$x'(T) = x'(0) + \int_0^T [u'(t) - v'(t)]\, dt$$

and, since $x(0) = x'(0) = 0$, we get

$$\int_0^T \Delta u(t) dt = \int_0^T \Delta v(t) dt + \Delta x(T) \qquad (16)$$

where $\Delta v(t) = v'(t) - v(t)$. Recalling (2), where we see that $v(t)$ can take one of the three values 0, $\mu(t)$, or $\lambda_{\max}(t)$, we get:

$$\Delta v(t) = \begin{cases} \lambda_{\max}(t) - \mu(t) & \text{if } v(t) = \mu(t), v'(t) = \lambda_{\max}(t) \\ \mu(t) - \lambda_{\max}(t) & \text{if } v(t) = \lambda_{\max}(t), v'(t) = \mu(t) \\ 0 & \text{otherwise} \end{cases}$$

Observe that: (i) $\Delta v(t) = \mu(t) - \lambda_{\max}(t)$ implies that $s(t) = 1$, $x(t) = 0$, $\lambda_{\max}(t) < \mu(t)$, and $x'(t) > 0$. (ii) $\Delta v(t) = \lambda_{\max}(t) - \mu(t)$ implies that $s(t) = 1$, $x'(t) = 0$, $\lambda_{\max}(t) < \mu(t)$, and $x(t) > 0$. It follows that $\Delta x(t) < 0$. However, by Lemma 3.3 we have $\Delta x(t) \geq 0$, therefore this case is infeasible.

With these observations in mind, let us decompose $[0, T]$ into intervals according to the value of $\Delta v(t)$. Assume there are V intervals (V is a random variable) in which $\Delta v(t) = \mu(t) - \lambda_{\max}(t)$, and let each such interval be $[a_i, b_i)$, $i = 1, \ldots, V$. We can then write:

$$\int_0^T \Delta v(t) dt = \sum_{i=1}^{V} \int_{a_i}^{b_i} [\mu(t) - \lambda_{\max}(t)]\, dt \qquad (17)$$

where, for every such interval, for all $t \in [a_i, b_i)$, $s(t) = 1$, $x(t) = 0$, $x'(t) > 0$ and

$$\mu(t) > \lambda_{\max}(t) > \lambda_{\min}(t) \quad \text{for all } t \in [a_i, b_i) \quad (18)$$

Moreover, for every such interval, the perturbed sample path is such that $x'(t) > 0$, $v'(t) = \mu(t) > \lambda_{\max}(t)$, and, from (1), we have

$$u'(t) \leq \lambda_{\max}(t) \qquad (19)$$

Thus, for the perturbed sample path we obtain:

$$x'(b_i) = x'(a_i) + \int_{a_i}^{b_i} [u'(t) - v'(t)]\, dt$$
$$= x'(a_i) + \int_{a_i}^{b_i} [u'(t) - \mu(t)]\, dt$$

which gives $\int_{a_i}^{b_i} [u'(t) - \mu(t)]\, dt = x'(b_i) - x'(a_i)$.

Since $x(t) = 0$ in any such interval, we have $\Delta x(a_i) = x'(a_i)$ and $\Delta x(b_i) = x'(b_i)$ where, by Lemma 3.3, $\Delta x(a_i) \leq \Delta z_0$ and $\Delta x(b_i) \leq \Delta z_0$. It follows that

$$\left| \int_{a_i}^{b_i} [u'(t) - \mu(t)]\, dt \right| = |x'(b_i) - x'(a_i)| \leq 2\Delta z_0 \qquad (20)$$

Regarding the left-hand-side of (20), observe that $\mu(t) - u'(t) \geq \mu(t) - \lambda_{\max}(t) > 0$, where we have used (18) and (19). Therefore,

$$\left| \int_{a_i}^{b_i} [u'(t) - \mu(t)]\, dt \right| = \int_{a_i}^{b_i} [\mu(t) - u'(t)]\, dt$$
$$\geq \int_{a_i}^{b_i} [\mu(t) - \lambda_{\max}(t)]\, dt \qquad (21)$$

Looking at the right-hand-side of (21), we have

$$\int_{a_i}^{b_i} [\mu(t) - \lambda_{\max}(t)]\, dt \geq c_2(b_i - a_i) \qquad (22)$$

where c_2 is a positive constant by **Assumption 3**. Combining the three inequalities (20), (21), and (22), we get $c_2(b_i - a_i) \leq 2\Delta z_0$ or

$$b_i - a_i \leq \frac{2\Delta z_0}{c_2} \qquad (23)$$

Returning to (17) and recalling that $\lambda_{\max}(t) \leq c_1$, $\mu(t) \leq c_1$ for some $c_1 < \infty$ from **Assumption 2**, we get

$$\left| \int_0^T \Delta v(t) dt \right| \leq \sum_{i=1}^{V} \left| \int_{a_i}^{b_i} 2c_1 dt \right|$$
$$\leq 2c_1 V \left(\frac{2\Delta z_0}{c_2} \right) = \frac{4c_1 \Delta z_0}{c_2} V \qquad (24)$$

where the second inequality is due to (23). Finally, returning to (16), we get $\left| \int_0^T \Delta u(t) dt \right| \leq \left| \int_0^T \Delta v(t) dt \right| + |\Delta x(T)| \leq \left(\frac{4c_1}{c_2} V + 1 \right) \Delta z_0$ where we have used the bound in (24), as well as Lemma 3.3. Note that all V intervals start and end with exogenous or endogenous events, and the number of endogenous events is bounded by the number of exogenous events. Thus, we have $V \leq N(T)$, and it follows that $\left| \int_0^T \Delta u(t) dt \right| \leq \left(\frac{4c_1}{c_2} N(T) + 1 \right) \Delta z_0$. Therefore, from (15), we have $|\Delta Q_T| \leq \frac{1}{T} (\frac{4c_1}{c_2} N(T) + 1) |\Delta z_0|$, i.e.,

$Q_T(\mathbf{z})$ is Lipschitz continuous in z_0 with constant $\frac{1}{T}(\frac{4c_1}{c_2}N(T)+1)$. Since, by assumption, $E[N(T)] < \infty$, this establishes unbiasedness. ■

Proof of Theorem 4.1. Recall (5) and let $R'_T = R_T(\mathbf{z} + \Delta\mathbf{z}_0)$ denote the overflow rate when $\mathbf{z}$ is perturbed by $\Delta\mathbf{z}_0 = [\Delta z_0\ 0]$. Thus,

$$R_T(\mathbf{z}) = \frac{1}{T}\int_0^T \mathbf{1}[x(t) \geq B]dt$$

$$R'_T = \frac{1}{T}\int_0^T \mathbf{1}[x'(t) \geq B]dt$$

and we get

$$\Delta R_T = \frac{1}{T}\int_0^T \psi(t)dt \tag{25}$$

where we set $\psi(t) = \mathbf{1}[x'(t) \geq B] - \mathbf{1}[x(t) \geq B]$. Observe that the case of $\psi(t) = -1$ implies that $\Delta x(t) < 0$, which contradicts the result $\Delta x(t) \geq 0$ of Lemma 3.3. This case is, therefore, infeasible. With this in mind, let us decompose $[0, T]$ into intervals according to the value of $\psi(t)$. Assume there are P intervals (P is a random variable) in which $\psi(t) = 1$ and no exogenous event occurs within the interval, and let each such interval be $[a_i, b_i)$, $i = 1, \ldots, P$. We can then write:

$$\int_0^T \psi(t)dt = \sum_{i=1}^P \int_{a_i}^{b_i} \psi(t)dt = \sum_{i=1}^P (b_i - a_i) \tag{26}$$

where, for every such interval,

$$x(t) < B \leq x'(t) \quad \text{for all } t \in [a_i, b_i) \tag{27}$$

Moreover, in each interval we have $u(t) = \lambda_{\min}(t)$, since $x'(t) \geq B > z_1 \geq z_0$. Recalling that each interval is defined so that no exogenous event is included, there are two possible cases to consider regarding the sign of $\lambda_{\min}(t) - v(t)$:

Case 1: $\lambda_{\min}(t) > v(t)$. We have

$$x(b_i) = x(a_i) + \int_{a_i}^{b_i} [\lambda_{\min}(t) - v(t)]\, dt \tag{28}$$

and it follows that

$$x(b_i) > x(a_i) \tag{29}$$

Using **Assumption 3** and the fact that $v(t) \leq \mu(t)$ for all $t \in [0, T]$, we obtain

$$\left| \int_{a_i}^{b_i} [\lambda_{\min}(t) - v(t)]\, dt \right| = \int_{a_i}^{b_i} [\lambda_{\min}(t) - v(t)]\, dt$$
$$\geq \int_{a_i}^{b_i} [\lambda_{\min}(t) - \mu(t)]\, dt \geq c_2(b_i - a_i) \tag{30}$$

Using (28), (29), (27) we also obtain:

$$\left| \int_{a_i}^{b_i} [\lambda_{\min}(t) - v(t)]\, dt \right| = x(b_i) - x(a_i)$$
$$\leq B - x(a_i) \leq x'(a_i) - x(a_i) \leq \Delta z_0 \tag{31}$$

where the last inequality follows from Lemma 3.3. Combining (30) and (31) gives

$$b_i - a_i \leq \frac{\Delta z_0}{c_2} \tag{32}$$

Case 2: $\lambda_{\min}(t) < v(t)$. This implies that $v(t) = \mu(t)$, i.e., $s(t) = 1$, and $v'(t) = \mu(t)$. Then, for the perturbed path, $x'(b_i) = x'(a_i) + \int_{a_i}^{b_i} [\lambda_{\min}(t) - \mu(t)]\, dt$,

and it follows that

$$x'(b_i) < x'(a_i) \tag{33}$$

Using **Assumption 3**, we get

$$\left| \int_{a_i}^{b_i} [\lambda_{\min}(t) - \mu(t)]\, dt \right| = \int_{a_i}^{b_i} [\mu(t) - \lambda_{\min}(t)]\, dt \tag{34}$$
$$\geq c_2(b_i - a_i)$$

In addition, using (33) and (27), we get

$$\left| \int_{a_i}^{b_i} [\lambda_{\min}(t) - \mu(t)]\, dt \right| = |x'(b_i) - x'(a_i)| \tag{35}$$
$$= x'(a_i) - x'(b_i) \leq x'(a_i) - B \leq x'(a_i) - x(a_i) \leq \Delta z_0$$

where the last inequality follows from Lemma 3.3. Therefore, combining (34) and (35) we obtain (32) once again.

Combining both cases above, from (26) we get $\int_0^T \psi(t)dt = \sum_{i=1}^P (b_i - a_i) \leq \frac{P}{c_2}\Delta z_0$. Note that all P intervals start and end with exogenous or endogenous events, and the number of endogenous events is bounded by the number of exogenous events. Thus, we have $P \leq N(T)$, and it follows that $\left|\int_0^T \psi(t)dt\right| \leq \frac{N(T)}{c_2}\Delta z_0$. Therefore, from (25), we have $|\Delta Q_T| \leq \frac{1}{T}\frac{N(T)}{c_2}|\Delta z_0|$, i.e., $R_T(\mathbf{z})$ is Lipschitz continuous in z_0 with constant $\frac{1}{T}\frac{N(T)}{c_2}$. Since, by assumption, $E[N(T)] < \infty$, this establishes unbiasedness. ■

References

[1] J.G.Kimemia and S.G.Gershwin, "An algorithm for the computer control of production in flexible manufacturing

systems," *IEEE Trans. Automatic Control*, vol. AC-15, pp. 353–362, Dec. 1983.

[2] R. Akella and P.R.Kumar, "Optimal control of production rate in a failure prone manufacturing system," *IEEE Trans. Automatic Control*, vol. AC-31, pp. 116–126, Feb. 1986.

[3] T.Bielecki and P.R.Kumar, "Optimality of zero-inventory policies for unreliable manufacturing systems," *Operations Research*, vol. 36, pp. 532–541, July-August 1988.

[4] J. Perkins and R. Srikant, "The role of queue length information in congestion control and resource pricing," *Proc. of the 38th Conference on Decision and Control*, Dec. 1999.

[5] Y. Feng and H. Yan, "Optimal production control in a discrete manufacturing system with unreliable machines and random demands," *IEEE Trans. Automatic Control*, vol. 35, pp. 2280–2296, December 2000.

[6] J.Q.Hu, P.Vakili, and G.X.Yu, "Optimality of hedging point policies in the production control of failure prone manufacturing systems," *IEEE Trans. Automatic Control*, vol. 39, pp. 1875–1994, Sept. 1994.

[7] J.R.Perkins and R.Srikant, "Hedging policies for failure-prone manufacturing systems: Optimality of JIT and bounds on buffer levels," *IEEE Trans. Automatic Control*, vol. 43, pp. 953–957, July 1998.

[8] S. P. Sethi and Q. Zhang, *Hierarchical Decision Making in Stochastic Manufacturing Systems.* Cambridge, MA: Birkhauser Boston, 1994.

[9] S. P. Sethi, H. Yan, H. Zhang, and Q. Zhang, "Optimal and hierarchical controls in dynamic stochastic manufacturing systems: A review," *Manufacturing and Service Operations Management*, vol. 4, no. 2, pp. 133–170, 2002.

[10] R. Suri and B. Fu, "On using continuous flow lines for performance estimation of discrete production lines," in *Proc. 1991 Winter Simulation Conference*, (Piscataway, NJ), pp. 968–977, 1991.

[11] H. Yan, X. Y. Zhou, and G. Yin, "Approximating an optimal production policy in a continuous flow line: Recurrence and asymptotic properties," *Operations Research*, vol. 47, pp. 535–549, 1999.

[12] C. G. Cassandras, Y. Wardi, B. Melamed, G. Sun, and C. G. Panayiotou, "Perturbation analysis for online control and optimization of stochastic fluid models," *IEEE Trans. Automatic Control*, vol. 47, August 2002.

[13] Y. Ho and X. Cao, *Perturbation Analysis of Discrete Event Dynamic Systems.* Boston, Massachusetts: Kluwer Academic Publishers, 1991.

[14] C. G. Cassandras and S. Lafortune, *Introduction to Discrete Event Systems.* Kluwer Academic Publishers, 1999.

[15] Y. Wardi, B. Melamed, C. G. Cassandras, and C. G. Panayiotou, "IPA gradient estimators in single-node stochastic fluid models," *Journal of Optimization Theory and Applications*, vol. 115, no. 2, pp. 369–406, 2002.

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

[17] H. Yu and C. G. Cassandras, "Perturbation analysis for production control and optimization of manufacturing systems," Tech. Rep., Boston Univ. Dept. of Manuf. Eng., 2002.

[18] H. Kushner and G. Yin, *Stochastic Approximation Algorithms and Applications.* New York, NY: Springer-Verlag, 1997.

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

[20] G. Sun, C. G. Cassandras, Y. Wardi, and C. G. Panayiotou, "Perturbation analysis of stochastic flow networks," in *Proceedings of 42th IEEE Conf. On Decision and Control*, pp. 4831–4836, Dec. 2003.