



Inventory Control for Supply Chains with Service Level Constraints: A Synergy between Large Deviations and Perturbation Analysis

IOANNIS CH. PASCHALIDIS * yannisp@bu.edu
*Center for Information and Systems Engineering (CISE) and Department of Manufacturing Engineering,
Boston University, Boston, MA 02215, USA*

YONG LIU yong.liu.2002@alum.bu.edu
SmartOps Corporation, One North Shore Center, Suite 400, Pittsburgh, PA 15212, USA

CHRISTOS G. CASSANDRAS ** cgc@bu.edu
*Center for Information and Systems Engineering (CISE) and Department of Manufacturing Engineering,
Boston University, Boston, MA 02215, USA*

CHRISTOS PANAYIOTOU *** christosp@ucy.ac.cy
Department of Electrical and Computer Engineering, University of Cyprus, Cyprus

Abstract. We consider a model of a supply chain consisting of n production facilities in tandem and producing a single product class. External demand is met from the finished goods inventory maintained in front of the most downstream facility (stage 1); unsatisfied demand is backlogged. We adopt a base-stock production policy at each stage of the supply chain, according to which the facility at stage i produces if inventory falls below a certain level w_i and idles otherwise. We seek to optimize the hedging vector $\mathbf{w} = (w_1, \dots, w_n)$ to minimize expected inventory costs at all stages subject to maintaining the stockout probability at stage 1 below a prescribed level (service level constraint). We make rather general modeling assumptions on demand and production processes that include autocorrelated stochastic processes. We solve this stochastic optimization problem by combining analytical (large deviations) and sample path-based (perturbation analysis) techniques. We demonstrate that there is a natural synergy between these two approaches.

Keywords: inventory control, supply chain management, service levels, large deviations, perturbation analysis

1. Introduction

In an era of time-based competition, the management and control of *supply chains* has emerged as a critical component of manufacturing and distribution enterprises (National

* Corresponding author. url: <http://ionia.bu.edu/>
** url: <http://vita.bu.edu/cgc>
*** This research was conducted while the author was with the Center for Information and Systems Engineering (CISE) and Department of Manufacturing Engineering, Boston University, Boston, MA 02215, USA.

Research Council, 1998; The Economist, 1998). Customers have become more demanding and require customized products delivered in a consistently timely manner. As competition intensifies, product shortages and stockouts significantly affect companies' reputations. As a result, and in addition to inventory cost reduction, *Quality of Service* (QoS) considerations have come to the foreground.

In this paper we consider a supply chain consisting of a tandem of production or distribution facilities. Demand is met from a *finished goods inventory* (FGI) that is accumulated in front of the most downstream facility. The supply chain strives to maintain this inventory nonempty to avoid stockouts, which lead either to backordered demand or simply lost sales. The fundamental trade-off is between *producing*, which accumulates inventory and incurs inventory costs, and *idling*, which leads to stockouts and unsatisfied demand. The objective is to devise a production policy at each stage of the supply chain which optimizes some measure of the integrated system's performance. More precisely, we will formulate a *stochastic optimization problem* aiming to minimize expected inventory costs subject to constraints that ensure that probabilities of stockout events or delays stay bounded below given desirable levels. We believe that such *service-level* constraints provide a natural representation of customer satisfaction; indeed their use is pervasive in practice. This is in contrast to most of the work in the literature which so far has considered policies minimizing expected linear inventory and backorder costs. Such an objective can be handled in our framework and permits trading-off the level of service with the resulting expected inventory cost. In practice, though, it is hard to obtain data which will help quantify customer satisfaction via linear backorder costs. Moreover, service-level constraints are often fixed by customers and not left to the discretion of the supply chain manager.

There is a large literature on production inventory systems (see (Kapuściński and Tayur, 1999) for a survey). The single-stage, single-class, version of the problem is significantly simpler. It has been shown in a variety of settings (all without set-up costs) that a so called *base-stock* policy minimizes expected linear inventory and backorder costs (e.g., (Federgruen and Zipkin, 1986)). Such a policy produces when inventory falls below a prescribed level and idles, otherwise; this threshold level is usually referred to as *hedging point* or *safety stock*. In multiclass single-stage systems the optimal policy is not known in general. Only special cases or asymptotic results have appeared in the literature (see (Bertsimas and Paschalidis, 2001) for a detailed literature review). Combining fluid and large deviations techniques, the single-stage multiclass problem was analyzed in (Bertsimas and Paschalidis, 2001) and two families of production policies have been proposed (one for linear and the other for quadratic inventory cost assumptions).

In a multi-stage single-class system *without capacity limits*, and for an expected linear inventory and backorder cost objective, Clark and Scarf (1960) in their seminal paper have shown the optimality of a production policy where each facility follows a base-stock policy based on the total inventory available locally and in the downstream facilities (*echelon* inventory). In the more general case where capacity limits exist and demand and service processes are autocorrelated, such a policy is not necessarily optimal. Under a similar policy Glasserman and Tayur (1995) proposed a perturbation

analysis approach to compute the hedging points in a capacitated single-class multi-stage system, and Glasserman (1997) has developed large deviations asymptotics to approximate stockout probabilities under renewal demand and constant production capacities. For a similar capacitated single-class multi-stage supply chain, and under autocorrelated demand and production processes, Paschalidis and Liu (2003) explores the full power of large deviations asymptotics to compute hedging points under two base-stock policies: one that operates based on local inventory information and an echelon base-stock policy. The large deviations analysis in the present paper follows the local inventory case of (Paschalidis and Liu, 2003).

In this paper we will combine Large Deviations (LD) (employed in (Bertsimas and Paschalidis, 2001; Paschalidis and Liu, 2003)) and Perturbation Analysis (PA) techniques (Cassandras and Lafortune, 1999). We will demonstrate that there is a natural synergy between these rather distinct approaches in tackling the complex stochastic optimization problem we are considering.

LD analysis is an *off-line* approach, based on asymptotics, intended to evaluate performance measures of interest involving “rare events”. A case in point arises with stockout probabilities which, by design, are relatively small. One can obtain asymptotically tight approximations of such probabilities for a variety of models, including some that capture dependencies in the demand and production processes. The approximation becomes tighter as these probabilities become smaller. LD analysis is computationally fast; it can be used to answer “what-if” questions and understand what aspects of demand and production affect stockouts. Further, it provides a characterization of the most likely path leading to stockouts, which leads to a more intuitive understanding of how stockouts occur. On the other hand, it does require detailed knowledge of the underlying model (e.g., the characteristics of stochastic processes representing demand and production).

PA is an *on-line* approach intended to estimate gradients of performance metrics with respect to control parameters by observing an actual system sample path (or through simulation). In contrast to LD analysis, it requires data, the collection of which may be time-consuming; in particular, estimating small probabilities generally requires long sample paths. On the other hand, PA does not rely on detailed model knowledge.

The origins of our approach that combines LD and PA are in (Paschalidis et al., 2001), where we considered a single-stage single-class system operating under a base-stock policy with a scalar hedging point equal to w . We addressed the problem of determining the smallest value of w which maintains the stockout probability below some given level η . The multi-stage supply chain problem we are considering in this paper is much harder. We now seek to determine a vector of hedging points $\mathbf{w} = (w_1, \dots, w_n)$, where n is the number of stages, so as to minimize a total expected inventory cost metric while still maintaining the stockout probability (at the most downstream stage 1) below some given level η_1 . This is a multidimensional constrained stochastic integer optimization problem. In this case, the synergy between LD and PA manifests itself as follows. LD analysis provides an approximation of the stockout probability based on which a feasible point \mathbf{w} is obtained. This becomes the initial point of an itera-

tive optimization scheme driven by cost sensitivity estimates obtained through PA. In (Paschalidis et al., 2001) we relied on a Finite Perturbation Analysis (FPA) approach, which may be extended to multiple stages, as in (Panayiotou and Cassandras, 1999). We are still, though, left with the challenge of a constrained stochastic integer optimization problem. Such problems are notoriously hard, despite some recent advances (e.g., (Gokbayrak and Cassandras, 2001; Shi and Olafsson, 2000)). Moreover, the fact that there are now multiple stages increases the computational burden of updating FPA-based estimates with every event in the system. Thus, we choose to convert the queueing model of the supply chain into a Stochastic Fluid Model (SFM) in which the parameters w_1, \dots, w_n are real-valued, enabling us to make use of much more efficient Infinitesimal Perturbation Analysis (IPA) methods to estimate a cost *gradient*, and employ a gradient-based optimization algorithm based on fairly standard stochastic approximation schemes (e.g., (Kushner and Clark, 1978)). Such SFMs have been recently used with success in communication network applications (Liu and Gong, 1999; Cassandras et al., 2002). We emphasize that an SFM is used to determine only the *form* of cost gradient estimators, while their *actual values* are still based on the original queueing model. Note that while a SFM may be too “crude” for performance analysis purposes, it is often able to accurately capture sensitivity information for control and optimization purposes (see (Cassandras et al., 2002)), which is precisely the context in which the proposed SFM is used in this paper.

Our analysis is general enough to handle models of demand and production that capture inherent temporal dependencies in these processes. This is key in modeling realistic demand scenarios, e.g., demand levels over a certain time period might have common underlying causes (weather, state of the economy, etc.). In addition, manufacturing facilities are *stochastic* and *failure-prone*. In particular, machine break downs affect capacity over extended periods which implies temporal dependencies in the production process.

To summarize, our main contributions are: (i) addressing an important inventory control problem in supply chains under realistic modeling assumptions that include the handling of service level constraints and the use of autocorrelated models for demand and production processes, and (ii) on the methodological side, the successful combination of LD and PA techniques.

The remainder of this paper is organized as follows: section 2 introduces the model and formulates the optimization problem. Section 3 contains the LD analysis, while the PA analysis is in section 4. The optimization algorithm we develop is presented in section 5 and numerical results are reported in section 6. Our concluding remarks are in section 7.

2. The model and problem formulation

Figure 1 depicts the supply chain model which consists of n production facilities in tandem. We will be referring to these facilities as *stages* of the supply chain and say that production consists of n stages. External demand is met from the finished goods




Figure 1. The model of the supply chain.

inventory (FGI) maintained in front of the most downstream facility (stage 1), and is backordered if inventory is not available. Every production facility is fed by its upstream facility; in particular, to produce one unit facility i , $i = 1, \dots, n-1$, requires one unit of the product of facility $i+1$. We assume that facility n is fed with an infinite supply of raw material, that is, no material requirement constraints are in effect there. In front of every facility i , $i = 2, \dots, n$, there is an inventory buffer which holds the final product of that facility and from which facility $i-1$ draws material for its production. We assume a discrete-time model where time is divided into time slots of equal duration. For all $i = 1, \dots, n$ and t we let B_t^i denote the amount that the facility at stage i can produce during time slot t (production capacity). We also let D_t^1 denote the amount of external orders arriving at stage 1 during time slot t . Finally, we let X_t^i , $i = 1, \dots, n$, denote the inventory in front of stage i at the beginning of time slot t . In intermediate stages $i = 2, \dots, n$ the inventory X_t^i is constrained to be nonnegative. In contrast, we allow the inventory at stage 1, X_t^1 , to take negative values to denote backordering; when X_t^1 is negative $-X_t^1$ is equal to the amount of backordered demand.

We assume that the demand process $\{D_t^1; t \in \mathbb{Z}\}$ and the production processes $\{B_t^i; t \in \mathbb{Z}\}$, $i = 1, \dots, n$, are arbitrary, stationary, and mutually independent stochastic processes that satisfy certain mild technical conditions (a large deviations principle, see (Bertsimas and Paschalidis, 2001; Bertsimas, Paschalidis, and Tsitsiklis, 1998) for details). These assumptions are satisfied by a fairly large class of stochastic processes, which includes renewal processes, Markov-modulated processes (where D_t^1 , for example, is a function of an underlying Markov process), and general stationary processes with mild mixing conditions. For stability purposes, we further assume that

$$\mathbf{E}[D_t^1] < \min_{i=1, \dots, n} \mathbf{E}[B_t^i], \quad (1)$$

which, by stationarity, carries over to all time slots t . Stability can be shown under the base-stock policy we will consider in this paper by using techniques from (Baccelli and Liu, 1992).

We propose a *base-stock* policy that maintains a safety stock equal to w_i for the (local) inventory of every stage i , $i = 1, \dots, n$. The inventory process becomes a function of $\mathbf{w} = (w_1, \dots, w_n)$; we will write $X_t^i(\mathbf{w})$ for $i = 1, \dots, n$ and all t . According to this policy, stage i produces until the local inventory $X_t^i(\mathbf{w})$ reaches the hedging point w_i and idles if $X_t^i(\mathbf{w}) \geq w_i$. The amount produced by stage i constitutes demand for the upstream stage $i+1$, for $i = 1, \dots, n-1$; we will denote it by $D_t^{i+1}(\mathbf{w})$. Note that the demand for stage i , $D_t^i(\mathbf{w})$, $i = 2, \dots, n$, is constrained by the production capacity B_t^{i-1}

and the available inventory $X_t^i(\mathbf{w})$. Demand at stage 1 is simply the external demand, so $D_t^1(\mathbf{w}) \equiv D_t^1$ for all t . According to this policy the inventory evolves as follows:

$$\begin{aligned} X_{t+1}^i(\mathbf{w}) &= \min\{X_t^i(\mathbf{w}) - D_t^i(\mathbf{w}) + B_t^i, X_t^i(\mathbf{w}) - D_t^i(\mathbf{w}) + X_t^{i+1}(\mathbf{w}), w_i\}, \\ i &= 1, \dots, n-1, \\ X_{t+1}^n(\mathbf{w}) &= \min\{X_t^n(\mathbf{w}) - D_t^n(\mathbf{w}) + B_t^n, w_n\}. \end{aligned}$$

The demand for stage i (or, equivalently, production of stage $i-1$) is given by

$$D_t^i(\mathbf{w}) = X_{t+1}^{i-1}(\mathbf{w}) - X_t^{i-1}(\mathbf{w}) + D_t^{i-1}(\mathbf{w}), \quad i = 2, \dots, n.$$

We quantify customer dissatisfaction by the steady-state probability, $\mathbf{P}[X_t^1(\mathbf{w}) \leq 0]$, of not being able to meet incoming demand immediately (*stockout probability*). We are interested in selecting a hedging point vector \mathbf{w} to minimize total expected inventory costs subject to the service level constraint of maintaining the stockout probability below some given η_1 . Namely, we will be solving the following optimization problem:

$$\begin{aligned} \min_{\mathbf{w} \in \mathbb{R}^+} C(\mathbf{w}) &= h_1 \mathbf{E}[(X_t^1(\mathbf{w}))^+] + \sum_{i=2}^n h_i \mathbf{E}[X_t^i(\mathbf{w})], \\ \text{subject to } \mathbf{P}[X_t^1(\mathbf{w}) \leq 0] &< \eta_1, \end{aligned} \tag{2}$$

where \mathbf{w} is the decision vector, h_i is a given holding cost of stage i , $i = 1, \dots, n$, $(X_t^1(\mathbf{w}))^+$ denotes $\max\{X_t^1(\mathbf{w}), 0\}$, and all expectations and probabilities are with respect to the steady-state probability distribution. Note that a steady-state distribution exists under a base-stock policy and the condition (1).

3. Large deviations analysis

In this section we will present our LD analysis. We start in section 3.1 with some background material on large deviations. In section 3.2 we address the special case of a single-stage problem. In this case, the problem in (2) is simplified into finding the minimum $w \equiv w_1$ that satisfies the service level constraint. Finally, in section 3.3 we develop a *decomposition approach* to address the general multi-stage supply chain problem. Our approach provides approximations for the stockout probabilities at all stages of the supply chain and allows us to find a hedging point vector \mathbf{w} that satisfies the service level constraint. This feasible point will be used in conjunction with the PA analysis to solve the problem in (2).

3.1. Preliminaries

Before we proceed with our agenda and in the form of background on large deviations we first review some basic results, which will also help in establishing some of our notation. Consider a sequence of i.i.d. random variables X_i , $i \geq 1$, with mean $\mathbf{E}[X_1] = \bar{X}$. The strong law of large numbers asserts that $(1/n) \sum_{i=1}^n X_i$ converges to \bar{X} , as $n \rightarrow \infty$,

with probability 1. Thus, for large n the event $\sum_{i=1}^n X_i > na$, where $a > \bar{X}$, (or $\sum_{i=1}^n X_i < na$, for $a < \bar{X}$) is a rare event. Specifically, its probability behaves as $e^{-nr(a)}$, as $n \rightarrow \infty$, where the function $r(\cdot)$ determines the rate at which the probability of this event is diminishing. Cramér's theorem (Dembo and Zeitouni, 1993) determines $r(\cdot)$, and is considered the first Large Deviations statement. In particular,

$$r(a) = \sup_{\theta} (\theta a - \log \mathbf{E}[e^{\theta X_1}]).$$

Consider next a sequence $\{S_1, S_2, \dots\}$ of random variables, with values in \mathbb{R} and define

$$\Lambda_n(\theta) \triangleq \frac{1}{n} \log \mathbf{E}[e^{\theta S_n}].$$

For the applications we have in mind, S_n is a partial sum process. Namely, $S_n = \sum_{i=1}^n X_i$, where X_i , $i \geq 1$, are identically distributed, possibly dependent, random variables. Let

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

(We assume that the limit exists for all θ , where $\pm\infty$ are allowed both as elements of the sequence $\Lambda_n(\theta)$ and as limit points.) We will refer to $\Lambda(\cdot)$ as the *limiting log-moment generating function*. Let us also define

$$\Lambda^*(a) \triangleq \sup_{\theta} (\theta a - \Lambda(\theta)),$$

which will be referred to as the *large deviation rate function*. Under a technical assumption on the form of $\Lambda(\theta)$ (see (Dembo and Zeitouni, 1993)) it has been shown (Gärtner–Ellis theorem) that for large enough n and for small $\epsilon > 0$,

$$\mathbf{P}[S_n \in (na - n\epsilon, na + n\epsilon)] \sim e^{-n(\Lambda^*(a) + O(\epsilon))}.$$

This can be viewed as an extension of Cramér's theorem to autocorrelated stochastic processes. We say that $\{S_n\}$ satisfies a *Large Deviations Principle* (LDP) with *rate function* $\Lambda^*(\cdot)$. The notation “ \sim ” should be interpreted as “asymptotically behaves”; more rigorously, the logarithm of the probability divided by n converges to $-\Lambda^*(a)$, as $n \rightarrow \infty$.¹ For the case that S_n is a partial sum process, i.e., $S_n = \sum_{i=1}^n X_i$, it has been shown (see (Dembo and Zeitouni, 1993)) that an LDP is satisfied when X_i 's are functions of an underlying Markov chain (Markov modulated process) or come from a general stationary process with mild mixing conditions.

In the sequel, we are also estimating the tail probabilities of the form $\mathbf{P}[S_n \leq na]$ or $\mathbf{P}[S_n \geq na]$. We therefore define large deviation rate functions associated with such probabilities. Consider the case where $S_n = \sum_{i=1}^n X_i$, the random variables X_i being identically distributed, and let $m = \mathbf{E}[X_1]$. It is easily shown (see (Dembo and Zeitouni, 1993)) that $\Lambda^*(m) = 0$. Let us define

$$\Lambda^{*+}(a) \triangleq \begin{cases} \Lambda^*(a) & \text{if } a > m, \\ 0 & \text{if } a \leq m \end{cases} \quad \text{and} \quad \Lambda^{*-}(a) \triangleq \begin{cases} \Lambda^*(a) & \text{if } a < m, \\ 0 & \text{if } a \geq m. \end{cases}$$

Notice that $\Lambda^{*+}(a)$ is a non-decreasing and $\Lambda^{*-}(a)$ is a non-increasing function of a , respectively. The convex duals of these functions are (see (Rockafellar, 1970) on convex duality)

$$\Lambda^+(\theta) \triangleq \begin{cases} \Lambda(\theta) & \text{if } \theta \geq 0, \\ +\infty & \text{if } \theta < 0 \end{cases} \quad \text{and} \quad \Lambda^-(\theta) \triangleq \begin{cases} \Lambda(\theta) & \text{if } \theta \leq 0, \\ +\infty & \text{if } \theta > 0, \end{cases}$$

respectively. In particular, $\Lambda^{*-}(a) = \sup_{\theta}(\theta a - \Lambda^-(\theta))$ and $\Lambda^{*+}(a) = \sup_{\theta}(\theta a - \Lambda^+(\theta))$.

On a notational remark, in the sequel we will be denoting by $\Lambda_X(\cdot)$ and $\Lambda_X^*(\cdot)$ the limiting log-moment generating function and the large deviations rate function, respectively, of the process X .

3.2. Large deviations analysis for single-stage systems

For the single-stage system, we discard the superscript i that indicates the stage for all quantities. Under a base-stock policy with a hedging point w , the inventory level evolves as $X_{t+1}(w) = \min\{X_t(w) - D_t + B_t, w\}$, where $\{D_t\}$ and $\{B_t\}$ denote the demand and production processes. Define the *shortfall* $L_t(w)$, at time slot t , as the gap between the hedging point w and the current inventory $X_t(w)$:

$$L_t(w) \triangleq w - X_t(w). \quad (3)$$

In terms of $L_t(w)$, the dynamics of the system can be written as

$$L_{t+1}(w) = \max\{L_t(w) + D_t - B_t, 0\}.$$

It can be seen that $L_t(w)$ does not in fact depend on w , thus, for the remaining of this subsection we will drop the reference to w in its notation. We can interpret L_t as the queue length of a discrete-time G/G/1 queue with D_t arrivals and B_t server's capacity during time slot t . We will refer to this queue as the make-to-order system corresponding to the make-to-stock system we are studying.

Stockout probability

Using the equivalence in (3), we can express the stockout probability in the make-to-stock system as the overflow probability in the make-to-order system, i.e.,

$$\mathbf{P}[X_t(w) \leq 0] = \mathbf{P}[L_t \geq w]. \quad (4)$$

Assuming that the demand and production processes satisfy an LDP, Bertsimas and Paschalidis (2001) establish the following result.

Theorem 3.1. The steady-state queue length process L_t satisfies

$$\lim_{w \rightarrow \infty} \frac{1}{w} \log \mathbf{P}[L_t \geq w] = -\theta^*,$$

where $\theta^* > 0$ is the largest root of the equation $\Lambda_D(\theta) + \Lambda_B(-\theta) = 0$.

Intuitively, for large enough w we have

$$\mathbf{P}[X_t(w) \leq 0] = \mathbf{P}[L_t \geq w] \sim e^{-w\theta^*}, \quad (5)$$

thus, based on this approximation, the minimum w that guarantees a stockout probability below η is

$$w^* = -\frac{\log \eta}{\theta^*}. \quad (6)$$

Notice that $\Lambda_D(\theta) + \Lambda_B(-\theta)$ is zero at the origin and has negative derivative at the same point (due to (1)). In the extreme case that $\Lambda_D(\theta) + \Lambda_B(-\theta) < 0$ for all $\theta > 0$ we will say that $\theta^* = \infty$. In this case, no stockouts occur and a safety stock of zero should be maintained (*Just-in-Time* (JIT) policy).

We can improve the accuracy of the asymptotic in (5) by introducing a constant in front of the exponential. More specifically, we will be estimating the stockout probability using the expression

$$\mathbf{P}[X_t(w) \leq 0] \approx \alpha e^{-w\theta^*}. \quad (7)$$

Thus, based on this approximation, (6) is modified according to

$$w^* = -\frac{\log(\eta/\alpha)}{\theta^*}. \quad (8)$$

An estimate of α can be obtained by using an idea from (Abate, Choudhury, and Whitt, 1995) and assuming that (7) is based on the *exact* distribution of the inventory process. Using (4), the right-hand side of (7) also approximates the tail of L_t . Hence, integrating the right-hand side of (7) from zero to infinity yields $\mathbf{E}[L_t]$, which implies

$$\alpha = \theta^* \mathbf{E}[L_t]. \quad (9)$$

Thus, to determine w^* we need the exponent θ^* and the expectation of the queue length L_t . The latter one is independent of the choice of w and can be obtained either by analytical approximations or by simulation (see (Bertsimas and Paschalidis, 2001)). In particular, $\mathbf{E}[L_t]$ can be obtained as a byproduct of the PA work (discussed in section 4). A key point here is that large deviations analysis is used to determine the exponent, and simulation or actual data may be used only to estimate $\mathbf{E}[L_t]$. This is beneficial since it is much easier to obtain a reliable estimate for $\mathbf{E}[L_t]$ than $\mathbf{P}[X_t(w) \leq 0]$, which might be small and require a large sample size.

Inventory cost

Next we consider the inventory cost. Using (3), the objective function in (2) becomes

$$\begin{aligned} C(w) &= h\mathbf{E}[(w - L_t)^+] \\ &= h\mathbf{E}[\max(w - L_t, 0)] \\ &= h(w - \mathbf{E}[L_t] + \mathbf{E}[\max(L_t - w, 0)]). \end{aligned}$$

Using the asymptotic in (7) we have

$$\begin{aligned}\mathbf{E}[\max(L_t - w, 0)] &= \int_0^\infty \mathbf{P}[\max(L_t - w, 0) > x] dx \\ &= \int_0^\infty \mathbf{P}[L_t - w > x] dx \approx \alpha \frac{e^{-w\theta^*}}{\theta^*}.\end{aligned}$$

Using (9) we obtain the following approximation for the expected inventory cost

$$C(w) \approx h(w - \mathbf{E}[L_t] + \mathbf{E}[L_t]e^{-w\theta^*}). \quad (10)$$

Summarizing the results for single-stage systems, LD techniques yield (asymptotically exact) approximations for the stockout probability in (7), and expected inventory cost in (10).

3.3. Multiple stages

As in the single stage case, we define the inventory shortfall for stage i as follows:

$$L_t^i(\mathbf{w}) \triangleq w_i - X_t^i(\mathbf{w}), \quad i = 1, \dots, n.$$

The dynamics of the supply chain can be written as

$$\begin{aligned}L_{t+1}^i(\mathbf{w}) &= \max\{L_t^i(\mathbf{w}) + D_t^i(\mathbf{w}) - B_t^i, L_t^i(\mathbf{w}) + D_t^i(\mathbf{w}) + L_t^{i+1}(\mathbf{w}) - w_{i+1}, 0\}, \\ i &= 1, \dots, n-1,\end{aligned} \quad (11)$$

$$L_{t+1}^n(\mathbf{w}) = \max\{L_t^n(\mathbf{w}) + D_t^n(\mathbf{w}) - B_t^n, 0\}. \quad (12)$$

The demand for stage i can now be expressed as

$$D_t^i(\mathbf{w}) = L_t^{i-1}(\mathbf{w}) - L_{t+1}^{i-1}(\mathbf{w}) + D_t^{i-1}(\mathbf{w}), \quad i = 2, \dots, n. \quad (13)$$

The major difficulty for analyzing this model and characterizing the stockout probabilities is that the production is constrained not only by its own capacity, but also by the upstream inventory. To bypass this difficulty we will *decouple* the various stages by ignoring the upstream inventory constraint on the downstream production. We can intuitively argue that this decomposition is in fact accurate when the inventory level of the upstream stage is high enough; then the influence of the upstream inventory constraint will be insignificant when compared to the capacity constraint. More precisely, the proposed decomposition amounts to assuming that the system operates according to a policy which satisfies

$$X_t^{i+1}(\mathbf{w}) \geq B_t^i, \quad i = 1, \dots, n-1, \quad (14)$$

almost surely for all time slots t . As a result, the dynamics of the supply chain can be simplified as follows:

$$\hat{X}_{t+1}^i(\mathbf{w}) = \min\{\hat{X}_t^i(\mathbf{w}) - \hat{D}_t^i + B_t^i, w_i\}, \quad i = 1, \dots, n,$$




Figure 2. The equivalent G/G/1 queue of stage i , $i = 1, \dots, n$, in a decoupled multi-stage supply chain.

where we use hats ($\hat{\cdot}$) to denote quantities of interest in the decomposed system. Then, (cf. (11), (12))

$$\hat{L}_{t+1}^i = \max\{\hat{L}_t^i + \hat{D}_t^i - B_t^i, 0\}, \quad i = 1, \dots, n. \quad (15)$$

That is, each stage behaves exactly as a single-stage system. As we will see \hat{L}_{t+1}^i and \hat{D}_t^i do not depend on \mathbf{w} , thus we have dropped the reference to \mathbf{w} from their notation.

Next note that the dynamics in (15) are exactly the dynamics of n decoupled make-to-order G/G/1 queues. In particular, as in the single-stage problem discussed above, \hat{L}_t^i can be interpreted as the queue length in a discrete-time G/G/1 queue with arrival process $\{\hat{D}_t^i\}$ and service process $\{B_t^i\}$ (see figure 2). Hence, theorem 3.1 holds. To apply it, however, we need the large deviations rate functions of the processes $\{\hat{D}_t^i\}$. For $i = 1$, $\{\hat{D}_t^1\}$ is the external demand process, whose large deviations rate function is assumed known. For the remaining stages $i = 2, \dots, n$, recall that \hat{D}_t^i is the demand for stage i generated by stage $i - 1$. In the equivalent make-to-order version of the system \hat{D}_t^i can be interpreted as the number of departures from the stage $i - 1$ queue during time slot t . To see that, consider the queue corresponding to stage $i - 1$ which has queue length equal to \hat{L}_t^{i-1} at time slot t . Equation (13) simply states that the queue length at slot t (\hat{L}_t^{i-1}) plus the number of arrivals at slot t (\hat{D}_t^{i-1}) is equal to the queue length at slot $t + 1$ (\hat{L}_{t+1}^{i-1}) plus the number of departures during slot t (\hat{D}_t^i).

The following theorem characterizes the large deviations behavior of the departure process $\{\hat{D}_t^i\}$, for all $i = 2, \dots, n$ (Paschalidis and Liu, 2003). This theorem is a corollary of a result in (Bertsimas, Paschalidis, and Tsitsiklis, 1998) which characterizes the departure process of a G/GI/1 queue using a continuous-time model.

Theorem 3.2 (Departure process). The departure process of the G/G/1 queue of stage $i - 1$ satisfies

$$\lim_{T \rightarrow \infty} \frac{1}{T} \log \mathbf{P} \left[\sum_{t=1}^T \hat{D}_t^i \geq Ta \right] = -\Lambda_{\hat{D}^i}^{*+}(a), \quad i = 2, \dots, n,$$

where

$$\begin{aligned} \Lambda_{\hat{D}^i}^{*+}(a) &= \Lambda_{B^{i-1}}^{*+}(a) + \Lambda_{\Gamma^{i-1}}^{*+}(a) \quad \text{and} \\ \Lambda_{\Gamma^{i-1}}^{*+}(a) &= \sup_{\{\theta | \Lambda_{\hat{D}^{i-1}}^+(\theta) + \Lambda_{B^{i-1}}(-\theta) < 0\}} [\theta a - \Lambda_{\hat{D}^{i-1}}^+(\theta)]. \end{aligned}$$

Proof. We prove the result by establishing a correspondence with a continuous-time G/G/1 queue and invoking the result in (Bertsimas, Paschalidis, and Tsitsiklis, 1998). Consider first the queue corresponding to stage $i - 1$ with queue length \hat{L}_t^{i-1} at time slot t . Recall that the Lindley equation for this queue length is

$$\hat{L}_{t+1}^{i-1} = \max\{\hat{L}_t^{i-1} + \hat{D}_t^{i-1} - B_t^{i-1}, 0\}. \quad (16)$$

Adding equation (13) for stage i over all time slots $1, \dots, T$, we obtain

$$\sum_{t=1}^T \hat{D}_t^i = \hat{L}_1^{i-1} - \hat{L}_{T+1}^{i-1} + \sum_{t=1}^T \hat{D}_t^{i-1}. \quad (17)$$

Consider next a continuous-time G/G/1 queue and let us denote by A_t the t th interarrival (interval between the arrivals of the $(t - 1)$ st and t th customer), and by S_t the service time of the t th customer. The waiting time, W_t , of the t th customer satisfies the following Lindley equation:

$$W_t = \max\{W_{t-1} + S_{t-1} - A_t, 0\}. \quad (18)$$

The interdeparture time, Z_t , of the t th customer (time interval between the departure times of customers $t - 1$ and t) can be expressed as

$$Z_t = W_t - W_{t-1} + A_t + S_t - S_{t-1}.$$

Summing up the above over all customers $1, \dots, T$ we obtain

$$\sum_{t=1}^T Z_t = W_T - W_0 + S_T - S_0 + \sum_{t=1}^T A_t. \quad (19)$$

Compare now equations (16) and (17) with (18) and (19), respectively. By making the substitutions

$$W_t := \hat{L}_{t+1}^{i-1}, \quad A_t := -\hat{D}_t^{i-1}, \quad S_{t-1} := -B_t^{i-1}, \quad Z_t := -\hat{D}_t^i,$$

for all t , we observe that $\sum_{t=1}^T \hat{D}_t^i$ has the same large deviations behavior as $-\sum_{t=1}^T Z_t$.² Hence, invoking the result for the departure process of a continuous-time G/G/1 queue from (Bertsimas, Paschalidis, and Tsitsiklis, 1998) we obtain the large deviations characterization of the process $\{\hat{D}_t^i, t \in \mathbb{Z}\}$ as it appears in the statement of the theorem. \square

We now have all the ingredients to analyze \hat{L}_t^i for every stage i in isolation. The result is summarized in the following theorem.

Theorem 3.3. For every stage $i = 1, \dots, n$ of the decoupled system, the steady-state queue length \hat{L}^i satisfies

$$\lim_{w_i \rightarrow \infty} \frac{1}{w_i} \log \mathbf{P}[\hat{L}^i \geq w_i] = -\theta_{\hat{L},i}^*,$$

where $\theta_{\hat{L},i}^*$ is the largest root of the equation $\Lambda_{\hat{D}^i}^+(\theta) + \Lambda_{B^i}(-\theta) = 0$, and $\Lambda_{\hat{D}^i}^+(\theta)$, for $i = 2, \dots, n$, is given by

$$\Lambda_{\hat{D}^i}^+(\theta) = \sup_a (\theta a - \Lambda_{\hat{D}^i}^{*+}(a)),$$

where $\Lambda_{\hat{D}^i}^{*+}(a)$ is as specified in theorem 3.2.

Assume now that the stockout probability for stage 1 needs to be below some η_1 . We can then select the requirement for the stockout probability of stage i , η_i , to be the same as, or an order of magnitude less than, the corresponding requirement, η_{i-1} , for its downstream stage $i - 1$. By doing so we force the system to operate in a regime where the decomposition approach is more accurate, that is, (14) would be rarely violated. As in section 3.2, we can obtain the hedging points (cf. (6))

$$w_i^* = -\frac{\log \eta_i}{\theta_{\hat{L},i}^*}, \quad i = 1, \dots, n.$$

We can again refine our approximations as follows

$$\mathbf{P}[X^i(\mathbf{w}) \leq 0] = \mathbf{P}[L^i(\mathbf{w}) \geq w_i] \approx \alpha_i e^{-w_i \theta_{\hat{L},i}^*}, \quad i = 1, \dots, n,$$

where $\alpha_i = \theta_{\hat{L},i}^* \mathbf{E}[\hat{L}^i]$. Note that in the decoupled system $\mathbf{E}[\hat{L}^i]$ is independent of w_i , and can be obtained either by approximations of the expected queue length in a G/G/1 queue (as in (Bertsimas and Paschalidis, 2001)) or from simulation or actual system data. Estimation of $\mathbf{E}[\hat{L}^i]$ can also be a by-product of PA. Hence, the hedging point satisfies

$$w_i^* = -\frac{\log(\eta_i/\alpha_i)}{\theta_{\hat{L},i}^*}, \quad i = 1, \dots, n. \quad (20)$$

4. Perturbation analysis

In this section, we turn our attention to the use of PA in seeking to solve the optimization problem (2). Closed-form expressions for the cost function $C(\mathbf{w})$ and the stockout probability $\mathbf{P}[X_t^1(\mathbf{w}) \leq 0]$ are generally not available, short of some approximations such as those obtained through LD analysis. FPA techniques may be used to estimate these quantities under multiple values of \mathbf{w} while observing a *single* sample path of the supply chain under one particular value of \mathbf{w} . This approach was used in (Paschalidis et al., 2001) for the single-stage case, where solving (2) is reduced to a relatively simple problem of determining the minimum w_1 that satisfies the service level constraint. The FPA techniques used in (Paschalidis et al., 2001) may be extended to multiple stages. However, aside from the fact that they become computationally more intensive, they do not directly lead to a solution of (2), which requires the use of some sort

of iterative scheme making use of the FPA-based estimates of $\mathbf{E}[(X_t^1(\mathbf{w}))^+]$, $\mathbf{E}[X_t^i(\mathbf{w})]$, and $\mathbf{P}[X_t^1(\mathbf{w}) \leq 0]$ under multiple different values of \mathbf{w} . General-purpose algorithms for such stochastic integer optimization problems are not available, despite recent advances for some classes of such problems (e.g., as in (Gokbayrak and Cassandras, 2001; Shi and Olafsson, 2000)).

The approach we take here is to replace the queueing model of the supply chain by a Stochastic Fluid Model (SFM) in which the parameters w_1, \dots, w_n are real-valued. Then, we can tackle (2) through stochastic approximation algorithms (e.g., (Kushner and Clark, 1978)) that rely on estimates of the gradient of the cost function $C(\mathbf{w})$. Such estimates may be obtained through Infinitesimal Perturbation Analysis (IPA) methods. As we shall see, even though the IPA estimators are derived based on the SFM of the supply chain, they can be implemented using data from a single sample path of the *actual* system. Arguably, a SFM may not always provide sufficient accuracy for the purpose of performance analysis; yet it can accurately capture *sensitivity* information for control and optimization purposes so as to provide the same or a very close approximation of the optimal solution in (2); for more details on this approach and some successful applications see (Cassandras et al., 2002).

4.1. The Stochastic Fluid Model (SFM)

Figure 3 shows the Stochastic Fluid Model (SFM) for the supply chain with n stages shown in figure 1. Each stage M_1, \dots, M_n can process fluid at a rate $\mu_i(t)$, $i = 1, \dots, n$. The n th stage M_n is assumed to have an infinite supply while the finished goods inventory buffer F is drained by the demand at a rate $\rho(t)$. The random functions $\mu_i(t)$, $i = 1, \dots, n$, and $\rho(t)$ are assumed piecewise constant with probability 1, though extensions to piecewise differentiable functions are possible (see (Wardi et al., 2001)).

Let $\mathbf{x}(t) = [x_1(t), \dots, x_n(t)]$ denote the state of the system, where $x_1(t) \leq w_1$, and $0 \leq x_i(t) \leq w_i$, $i = 2, \dots, n$, correspond to the amount of fluid present in each of the buffers; w_i , $i = 1, \dots, n$, correspond to the hedging points of our policy (equivalently, they may be thought of as buffer capacities). Note that $x_1(t)$ can take negative values which represents the amount of backlogged demand as in the discrete-time model of




Figure 3. Stochastic fluid model.




Figure 4. Typical sample path for a two-stage system.

section 2. The dynamics of the system are given by

$$\frac{dx_1(t)}{dt} = \begin{cases} \gamma_1(t) - \rho(t) & \text{if } x_1(t) < w_1, \\ 0 & \text{if } x_1(t) = w_1 \end{cases}$$

and

$$\frac{dx_i(t)}{dt} = \begin{cases} \gamma_i(t) - \gamma_{i-1}(t) & \text{if } 0 < x_i(t) < w_i, \\ 0 & \text{if } x_i(t) = 0 \text{ or } x_i(t) = w_i, \end{cases}$$

for $i = 2, \dots, n$, where $\gamma_i(t)$ and $\gamma_{i-1}(t)$ are the inflow and outflow rates of buffer i capturing the coupling between the individual dynamics of the n buffers, and we have:

$$\gamma_i(t) = \begin{cases} \gamma_{i+1}(t) & \text{if } x_{i+1}(t) = 0, \\ \mu_{i+1}(t) & \text{if } x_{i+1}(t) > 0, \end{cases} \quad i = 1, \dots, n-1, \quad \text{and} \quad \gamma_n(t) = \mu_n(t).$$

Figure 4 shows a typical sample path for a system with $n = 2$. We observe that the trajectories of $x_1(t)$ and $x_2(t)$ can be divided into non-empty (busy) and empty periods. Figure 4 also illustrates the coupling between the $x_1(t)$ and $x_2(t)$ sample paths due to the occurrence of either *exogenous* or *endogenous* events. Exogenous events are those not associated with the structure of the system, i.e., the random changes in the rates $\mu_i(t)$ and $\rho(t)$, which we assume are independent of \mathbf{w} . On the other hand, endogenous events are generated due to the dynamics of the system.

In particular, we define the following endogenous events. (I) β_i^q : the trajectory of $x_q(t)$, $q \in \{1, \dots, n\}$, *ceases to be empty* for the i th time. This event occurs at time instants b_i^q . (II) ϵ_i^q : the trajectory of $x_q(t)$, $q \in \{1, \dots, n\}$, *becomes empty* for the i th time. This event occurs at time instants e_i^q . (III) ϕ_{ik}^q : the trajectory of $x_q(t)$, $q \in \{1, \dots, n\}$, *becomes full* for the k th time, $k = 1, \dots, K_i^q$, in the i th busy period, where $K_i^q \geq 0$ is the number of times that $x_q(t)$ has reached its limit w_q . This event

occurs at time instants $f_{ik}^q = f_{n_{ik}}^q$. Here, n_{ik} is a single index that counts the total number of times that ϕ_{ik}^q has occurred and is given by $n_{ik} = k + \sum_{j=1}^{i-1} K_j^q$. (IV) ω_{ik}^q : the trajectory of $x_q(t)$, $q \in \{1, \dots, n\}$, *ceases being full* for the k th time in the i th busy period. Examples of these events are shown in figure 4. For $x_q(t)$, $B_j^q = (b_j^q, e_j^q)$ denotes the j th busy period for all $t \in B_j^q$, $q \in \{2, \dots, n\}$. Given an interval $[0, T]$, N_T^q and \tilde{N}_T^q , $q \in \{1, \dots, n\}$, denote the random number of busy and empty periods, respectively (including a possibly incomplete last one). Finally, N_T denotes the total number of endogenous events that occur in the interval $[0, T]$, assuming $\mathbf{E}[N_T] < \infty$.

4.2. Infinitesimal perturbation analysis

In this section we consider the *expected workload* $\mathbf{E}[Q_T(\mathbf{w})]$ and *stockout probability* $P_T(\mathbf{w})$ over an interval $[0, T]$ as functions of our control parameter vector \mathbf{w} . The expected workload cost is defined as

$$C_T(\mathbf{w}) = \mathbf{E}[Q_T(\mathbf{w})] = \sum_{i=1}^n h_i \mathbf{E}[Q_T^i(\mathbf{w})], \quad (21)$$

where

$$Q_T^1(\mathbf{w}) = \frac{1}{T} \int_0^T [x_1(\mathbf{w}; t)]^+ dt \quad \text{and} \quad Q_T^i(\mathbf{w}) = \frac{1}{T} \int_0^T x_i(\mathbf{w}; t) dt, \quad i = 2, \dots, n. \quad (22)$$

The stockout probability is given by $P_T(\mathbf{w}) = \mathbf{E}[S_T(\mathbf{w})]$, where

$$S_T(\mathbf{w}) = \frac{1}{T} \int_0^T \mathbf{1}[x_1(\mathbf{w}; t) < 0] dt. \quad (23)$$

Our objective is to evaluate the derivatives of the sample functions $Q_T(\mathbf{w})$ and $S_T(\mathbf{w})$ with respect to w_i , $i = 1, \dots, n$. These are referred to as the IPA estimators of the derivatives of $\mathbf{E}[Q_T(\mathbf{w})]$ and $\mathbf{E}[S_T(\mathbf{w})]$. A key issue is that of establishing the unbiasedness of these estimators, i.e., verifying that $\partial \mathbf{E}[Q_T(\mathbf{w})] / \partial w_i = \mathbf{E}[\partial Q_T(\mathbf{w}) / \partial w_i]$ and $\partial \mathbf{E}[S_T(\mathbf{w})] / \partial w_i = \mathbf{E}[\partial S_T(\mathbf{w}) / \partial w_i]$. Unbiasedness enables the use of the sample (IPA) derivatives in stochastic gradient-based techniques, which is needed in solving (2).

IPA proceeds by assuming that w_j is increased by $\Delta w_j > 0$ and determining all state changes

$$\Delta x_{ij}(t) = x_i(\mathbf{w} + \mathbf{e}_j \Delta w_j; t) - x_i(\mathbf{w}; t), \quad 0 \leq t \leq T, \quad i, j = 1, \dots, n,$$

where \mathbf{e}_j is an n -dimensional vector with all of its elements equal to 0 except the j th one which is equal to 1. Subsequently, we determine the corresponding changes in $Q_T(\mathbf{w})$ and $S_T(\mathbf{w})$, $\Delta Q_T(\mathbf{w})$ and $\Delta S_T(\mathbf{w})$, respectively. By taking limits as $\Delta w_i \rightarrow 0$ under appropriate conditions, we can then obtain the corresponding *right* sample derivatives. Note that in a similar fashion we can perform the analysis with $\Delta w_j < 0$, $j \in \{1, \dots, n\}$ and obtain the *left* sample derivative. Under some mild technical conditions (see (Cassandras et al., 2002; Wardi et al., 2001)) the two derivatives are equal, i.e., the sample

derivative exists, however, for the purposes of this paper we will not worry about those assumptions and simply use the *right* derivative. In the sequel all derivatives are actually *right* derivatives. In what follows, we shall limit ourselves to the cases of one and two stages.

Single-stage system

Clearly, perturbations in $x_1(t)$ due to Δw_1 , and consequently perturbations in $Q_T(\mathbf{w})$ and $S_T(\mathbf{w})$, respectively, are *only* generated when a ϕ^1 event occurs (see also figure 5). Therefore, the first perturbation is generated at the occurrence of the *first* ϕ^1 event at time

$$\tau_1 = f_1^1 = \inf\{t: x_1(t) = w_1, t \geq 0\}. \quad (24)$$

The perturbation in $x_1(t)$ due to Δw_1 is given by

$$\Delta x_{1,1}(t) = \begin{cases} 0 & \text{if } t < \tau_1, \\ (t - \tau_1)(\gamma_1(t) - \rho(t)) & \text{if } \tau_1 \leq t \leq \tau_1 + \Delta\tau_1, \\ \Delta w_1 & \text{if } t = \tau_1 + \Delta\tau_1, \end{cases} \quad (25)$$

where $\Delta\tau_1 = \tau_1(w_1 + \Delta w_1) - \tau_1(w_1)$ (shown in figure 5), and $\gamma_1(t) = \mu_1(t)$ is the inflow rate to buffer F . We point out that the second term in (25) is of the order $O(\Delta w_1^2)$ thus, for simplicity, in the sequel such terms will always be ignored, and we write

$$\Delta x_{1,1}(\tau_1^+) = \Delta w_1. \quad (26)$$

Since, $x_1(t)$ can also take negative values, the perturbation $\Delta x_{1,1}(\tau_1)$ persists for all $t > \tau_1$. Using (22) and (26) we immediately obtain

$$\Delta Q_T(w_1) = \frac{1}{T} \int_0^T \Delta x_{1,1}(t) dt = (T - \tau_1)\pi(\tau_1, T)\Delta w_1, \quad (27)$$

where we define

$$\pi(t_1, t_2) = \frac{1}{t_2 - t_1} \int_{t_1}^{t_2} \mathbf{1}[x_1(t) \geq 0] dt \quad (28)$$




Figure 5. Perturbed sample path for single-stage system.

to be the fraction of time in an interval (t_1, t_2) such that there is no stockout. This is easy to determine from direct observation of the supply chain operation by simply recording the times when $x_1(t) = 0$.

Next, we obtain $\Delta S_T(w_1)$. From (23) note that

$$S_T(w_1) = \frac{1}{T} \int_0^T \mathbf{1}[x_1(\mathbf{w}; t) < 0] dt = \frac{1}{T} \sum_{k=1}^{\bar{N}_T^1} (b_{k+1}^1 - e_k^1),$$

where, if the last empty period in $[0, T]$ is not complete, then we set $b_{\bar{N}_T^1+1}^1 = T$. As a result

$$\Delta S_T(w_1) = \frac{1}{T} \sum_{k=1}^{\bar{N}_T^1} (\Delta b_{k+1}^1 - \Delta e_k^1), \quad (29)$$

where $\Delta b_k^1 = b_k^1(w_1 + \Delta w_1) - b_k^1(w_1)$ and $\Delta e_k^1 = e_k^1(w_1 + \Delta w_1) - e_k^1(w_1)$, $k = 1, \dots, \bar{N}_T^1$. By inspection in figure 5 we can easily see that

$$\Delta b_k^1 = -\frac{\Delta x_{1,1}(b_k^1)}{\mu_1(b_k^1) - \rho(b_k^1)} \leq 0 \quad \text{and} \quad \Delta e_k^1 = -\frac{\Delta x_{1,1}(e_k^1)}{\mu_1(e_k^1) - \rho(e_k^1)} \geq 0.$$

Next, for notational convenience, let us define the sequence $\{v_k, k = 1, \dots, 2\bar{N}_T^1\}$ which consists of the elements of $\{b_k\}$ and $\{e_k\}$ such that $v_1 < \dots < v_{2\bar{N}_T^1}$. Then, using (26), we get

$$\Delta v_k = \begin{cases} 0 & \text{if } v_k < \tau_1, \\ -\frac{\Delta w_1}{\mu_1(v_k) - \rho(v_k)} & \text{if } v_k \geq \tau_1. \end{cases} \quad (30)$$

Substituting, (30) in (29) and noticing that $\Delta v_k \leq 0$ if v_k corresponds to a b_k event time while $\Delta v_k \geq 0$ if v_k corresponds to an e_k event time, we get

$$\Delta S_T(w_1) = -\frac{\Delta w_1}{T} \sum_{i=k^*}^{2\bar{N}_T^1} \frac{1}{|\mu_1(v_i) - \rho(v_i)|}, \quad (31)$$

where $k^* = \min\{k: v_k > \tau_1\}$. Dividing (27) and (31) by Δw_1 and taking the limit $\Delta w_1 \rightarrow 0$ we get

$$\frac{dQ_T(w_1)}{dw_1} = \frac{T - \tau_1}{T} \pi(\tau_1, T) \quad \text{and} \quad \frac{dS_T(w_1)}{dw_1} = -\frac{1}{T} \sum_{i=k^*}^{2\bar{N}_T^1} \frac{1}{|\mu_1(v_i) - \rho(v_i)|}. \quad (32)$$

Both estimators are simple to implement from direct observation of a sample path of the SFM. More importantly, though, the same information may also be obtained from a sample path of the *actual* supply chain. In particular, the workload derivative estimator simply requires (i) detecting τ_1 , the time when $x_1(t)$ reaches the hedging point w_1 for the first time, and (ii) measuring the fraction of time after τ_1 for which $x_1(t) \geq 0$. The stockout probability derivative estimator simply requires counting the number of

times that a stockout period either starts or ends weighted by the absolute value of the net instantaneous inflow at that time. Moreover, the estimators are independent of any assumptions on demand or production processes. Finally, it can be formally shown that IPA estimates in (32) are unbiased; see (Panayiotou, Cassandras, and Zhang, 2002) for details.

Two-stage system

Following the main ideas used in the single-stage system, we can derive the required derivatives with respect to w_1 and w_2 in a two-stage system, though the derivation becomes more tedious due to the coupling between the two stages. Since the technical details are beyond the scope of this paper, we limit ourselves to presenting the IPA estimators and refer interested readers to (Panayiotou, Cassandras, and Zhang, 2002) for complete derivations.

First, for notational convenience, define the sequence $\{\tau_i\}$, $i = 1, \dots, N_T^*$, such that

$$\tau_1 < \tau_2 < \dots < \tau_{N_T^*} \equiv \tau^*, \quad (33)$$

where

- τ_1 is the time that the very first ϕ^1 event occurred (see (24));
- τ_{2n} , $n = 1, \dots, (N_T^* - 1)/2$, is the time of the *first* ϵ^2 event following a ϕ^1 event (i.e., index of τ is even);
- τ_{2n-1} , $n = 2, \dots, (N_T^* + 1)/2$, is the time of the *first* ϕ^1 event following an ϵ^2 event (i.e., index of τ is odd);
- $N_T^* = \min\{2n - 1: x_2(t) = w_2 \text{ for some } \tau_{2n-1} \leq t < \tau_{2n}, n = 1, 2, \dots\}$. In case the event $x_2(t) = w_2$ for some $\tau_{2n-1} \leq t < \tau_{2n}$, $n = 1, 2, \dots$, does not occur in $[0, T]$, we set $N_T^* = \max\{2n - 1: \tau_{2n-1} \leq T, n = 1, 2, \dots\}$ and also set $\tau^* = T$.

The first and third items above together define a sequence of event times $\{\hat{\phi}_i^1: i = 1, \dots, (N_T^* + 1)/2\}$ such that stage 1 reaches its hedging point w_1 for the first time after a busy period ends at stage 2. The second item defines a sequence of events, $\{\hat{\epsilon}_i^2: i = 1, \dots, (N_T^* - 1)/2\}$, such that a busy period at stage 2 ends after some $\hat{\phi}_i^1$. Both are subsequences of $\{\phi_i^1\}$ and $\{\epsilon_i^2\}$, respectively. We stop adding elements in $\{\hat{\phi}_i^1\}$ and $\{\hat{\epsilon}_i^2\}$ after the occurrence of a $\hat{\phi}^1$ at time τ^* where either (i) $x_2(\tau^*) = w_2$, or (ii) $x_2(\tau^*) < w_2$ and a ϕ^2 event occurs *before* an ϵ^2 . Thus, N_T^* is the random number of events included in the two subsequences in the interval $[0, T]$. Finally, note that since the sequence terminates after a $\hat{\phi}^1$ event, then N_T^* must be an odd number. With these definitions in mind, we now present the following IPA derivatives:

$$\frac{\partial Q_T^1(\mathbf{w})}{\partial w_1} = \frac{T - \tau^*}{T} + \frac{1}{T} \sum_{n=1}^{(N_T^*-1)/2} (\tau_{2n} - \tau_{2n-1}) \pi(\tau_{2n-1}, \tau_{2n}), \quad (34)$$

$$\frac{\partial Q_T^2(\mathbf{w})}{\partial w_1} = -\frac{f^{2*} - \tau^*}{T} \cdot \mathbf{1}[x_2(\tau^*) < w_2] - \frac{1}{T} \sum_{n=1}^{(N_T^*-1)/2} (\tau_{2n} - \tau_{2n-1}), \quad (35)$$

$$\frac{\partial Q_T^1(\mathbf{w})}{\partial w_2} = \frac{1}{T} \sum_{j=1}^{N_T^2} (f_{m_j}^1 - e_j^2) \pi(e_j^2, f_{m_j}^1) \mathbf{1}[x_2(t) = w_2 \text{ for some } t \in B_j^2] \times \mathbf{1}[x_1(e_j^2) < w_1], \quad (36)$$

$$\frac{\partial Q_T^2(\mathbf{w})}{\partial w_2} = \frac{1}{T} \sum_{j=1}^{N_T^2} (e_j^2 - f_{j1}^2) \mathbf{1}[x_2(t) = w_2 \text{ for some } t \in B_j^2], \quad (37)$$

where τ^* and $\tau_i, i = 1, \dots, N_T^*$, are defined in (33), $f^{2*} = \inf\{t: x_2(t) = w_2, t \geq \tau^*\}$, $f_{m_j}^1 = \inf\{t: x_1(t) = w_1, t \geq e_j^2, j = 1, \dots, N_T^2\}$, $f_{j1}^2 = \inf\{t: x_2(t) = w_2, b_j^2 \leq t \leq e_j^2, j = 1, \dots, N_T^2\}$, and, $\pi(\cdot, \cdot)$ is defined in (28). Furthermore,

$$\frac{\partial S_T(\mathbf{w})}{\partial w_1} = - \sum_{i=1}^{2\bar{N}_T^1} \frac{\Delta_1(v_i)}{|\gamma_1(v_i) - \rho(v_i)|} \quad \text{and} \quad \frac{\partial S_T(\mathbf{w})}{\partial w_2} = - \sum_{i=1}^{2\bar{N}_T^1} \frac{\Delta_2(v_i)}{|\gamma_1(v_i) - \rho(v_i)|}, \quad (38)$$

where

$$\Delta_1(v_i) = \begin{cases} 0 & \text{if } v_i < \tau_1 \text{ or } \tau_{2n} \leq v_i < \tau_{2n+1}, n = 1, \dots, \frac{N_T^* - 1}{2}, \\ 1 & \text{if } v_i \geq \tau^* \text{ or } \tau_{2n-1} \leq v_i < \tau_{2n}, n = 1, \dots, \frac{N_T^* - 1}{2}, \end{cases} \quad (39)$$

and $\Delta_2(v_i)$ is the number of ϵ^2 events since the latest ϕ^1 event that occurred before v_i , i.e., $\Delta_2(\cdot)$ represents a counter that is increased by one when an ϵ^2 event occurs and is reset to zero every time a ϕ^1 event occurs.

Though the IPA estimates may appear complicated, in practice they are really simple to implement. All of them are either counters or accumulators that are updated at the occurrence of various events that are easy to detect on the actual supply chain, e.g., when a busy period at some stage starts or ends. Finally, it can also be shown that the IPA estimates in (34)–(38) are unbiased (see (Panayiotou, Cassandras, and Zhang, 2002)).

Remark. Given the discrete-time model presented in section 2, we obtain a fluid equivalent where the system state in the interval from t_n to t_{n+1} , $n = 1, 2, \dots$, is linearly interpolated:

$$x_i(\mathbf{w}; t) = X_{t_n}^i(\mathbf{w}) + (X_{t_{n+1}}^i(\mathbf{w}) - X_{t_n}^i(\mathbf{w}))(t - t_n) \quad \text{for all } t \in [t_n, t_{n+1}).$$

This SFM approximation of the discrete-time model enables us to immediately use the IPA estimates presented above, which, in turn, may be used to solve our optimization problem (2) as described next.

5. The optimization algorithm

Returning to (2), we shall treat \mathbf{w} as real-valued, which will enable us to make use of the gradient of $C(\mathbf{w})$ with respect to \mathbf{w} . To apply a stochastic optimization algorithm driven by gradient estimates, set

$$C_T(\mathbf{w}) = \mathbf{E} \left[\sum_{i=1}^n h_i Q_T^i(\mathbf{w}) \right]$$

with $Q_T^i(\mathbf{w})$ as in (22). Clearly, the steady state cost $C(\mathbf{w})$ in (2) is approximated by the finite horizon cost $C_T(\mathbf{w})$ above and we rely on standard ergodicity conditions to guarantee that the latter will asymptotically recover the former as $T \rightarrow \infty$.

Next, we formulate an unconstrained version of the problem by invoking standard multiplier techniques and define

$$Q_T^U(\mathbf{w}) = \sum_{i=1}^n h_i Q_T^i(\mathbf{w}) + \lambda(S_T(\mathbf{w}) - \eta_1) + \frac{p}{2} \|S_T(\mathbf{w}) - \eta_1\|^2, \quad (40)$$

where $S_T(\mathbf{w})$ was defined in (23) and we use λ as a Lagrange multiplier and p as a penalty factor. We then use the following gradient-based stochastic optimization type of algorithm, for $k = 0, 1, \dots$:

$$\begin{aligned} \mathbf{w}_{k+1} &= \mathbf{w}_k + v_k \nabla Q_T^U(\mathbf{w}_k), \\ \lambda_{k+1} &= \lambda_k + p_k (S_T(\mathbf{w}_k) - \eta_1), \\ p_{k+1} &= p_k + \delta p, \end{aligned} \quad (41)$$

where $\nabla Q_T^U(\mathbf{w}_k)$ is the IPA gradient estimate of $\mathbf{E}[Q_T^U(\mathbf{w})]$ and is given by

$$\nabla Q_T^U(\mathbf{w}_k) = \nabla Q_T(\mathbf{w}_k) + \lambda_k \nabla S_T(\mathbf{w}_k) + p_k (S_T(\mathbf{w}_k) - \eta_1) \nabla S_T(\mathbf{w}_k). \quad (42)$$

Moreover, $\{v_k\}$ is the step size sequence and δp is the increasing rate of the penalty which is taken to be constant for simplicity. Furthermore, $Q_T(\mathbf{w})$ and $S_T(\mathbf{w})$ are the sample functions corresponding to the workload and stockout probability defined in (21) and (23), respectively. The algorithm we proposed is based on the so called *method of multipliers* (see (Bertsekas, 1999)). The natural question that arises is whether this algorithm converges to a stationary point (i.e., a point where optimality conditions are satisfied) with probability 1. This can be shown in two steps. First show that for appropriate stepsize, $\{v_k\}$, and penalty, $\{p_k\}$, sequences, and for bounded $\{\lambda_k\}$ the iteration in (41) converges to a stationary point of the augmented Lagrangean in (40) with probability 1. Subsequently, one can follow the steps of (Bertsekas, 1999, theorem 4.2.2) to show that the standard penalty method converges to a stationary point of the original constrained problem with probability 1. This latter argument can be generalized as in (Bertsekas, 1999) to establish convergence of the method of multipliers. To establish convergence of the unconstrained problem (i.e., minimizing the augmented Lagrangean in (40)) one needs to show that certain technical conditions relating to the step size sequence and

properties of the IPA estimators are satisfied (see (Gokbayrak and Cassandras, 2001) for details), however, this is beyond the scope of this paper and thus such proof is omitted. Instead, we present several numerical results in the next section that we believe provide ample empirical evidence of the practical value of this approach.

Finally, a key issue in such iterative schemes is the speed of convergence, which greatly depends on the choice of an initial point \mathbf{w}^0 in (41). The LD analysis of section 3 provides this capability with the added benefit that this point lies in the feasible region of the problem with respect to the service level constraint.

6. Numerical results

Next, we provide some illustrative numerical results for: (a) a single-stage, and (b) a two-stage system.

In case (a), satisfying the service level constraint yields a hedging point. The numerical results we report demonstrate the consistent accuracy of all the approaches we used: the large deviations approach, the finite perturbation analysis (FPA) approach, and the SFM-based approach of section 4.

In case (b), the large deviations analysis only provides a feasible hedging point vector, i.e., one that satisfies the service level constraint. Starting from this initial feasible point, we rely on the gradient-based optimization algorithm of section 5 (where the gradients are supplied by the SFM-based method of section 4) to solve the problem in (2). We compare the results with a brute-force heuristic that uses only simulation.

6.1. A single-stage example

Our first example is a single-stage system. The demand and production processes are discrete-time Markov modulated processes (see figure 6). Both $\{D_t\}$ and $\{B_t\}$ are modulated by a two-state Markov chain. By \mathbf{r} we denote the (constant) vector of demand or production amounts at each state of the corresponding Markov chain. That is, D_t can be either 5 or 10, and B_t can be either 0 (machine down) or 14 (machine up). The load of the system is nearly 0.8.

Using the results of section 3.2 (theorem 3.1) we obtain $\theta^* = 0.120$. Thus, the hedging point suggested by the LD analysis (cf. equation (8)) is

$$w^* = -\frac{\log(\eta/(0.120 \cdot 6.402))}{0.120},$$




Figure 6. The models of the demand and production processes.

where we have used a value of $E[L_t] = 6.402$ obtained as a by-product of the PA algorithm.

Table 1 compares the LD results for w^* with the corresponding ones obtained by an FPA approach developed in (Paschalidis et al., 2001) and by the simple version of the algorithm presented in the previous section when $n = 1$, using the IPA derivative estimates (32) of section 4. All three approaches are very accurate and the results are very close for most service levels η . It is surprising that the LD results are accurate even for rather large η , since the asymptotics are taken as $w \rightarrow \infty$, or equivalently, as the stockout probability tends to zero. This is due to the refined approximation of equation (7) that led to the hedging point of equation (8). Only for the case $\eta = 0.3$ the LD hedging point slightly violates the service level constraint. In this case, the result obtained through the SFM-based method is most accurate; in general, it is somewhat surprising that such a “simplified” model of the true system dynamics gives such good results. When η is very small, however, FPA (and to a lesser degree the SFM-based IPA) requires longer sample paths to yield good (i.e., sufficiently low-variance) estimates, while the LD approximation remains fast and reliable.

6.2. A two-stage example

Next, we consider a two-stage system. Figure 7 depicts the model of demand and production processes. We use the result of theorem 3.3 to compute

$$\theta_{L,1}^* = 0.343, \qquad \theta_{L,2}^* = 0.213.$$

Table 1
Comparing the analytical LD expression for the stockout probability with PA (FPA and SFM-based IPA) and simulated values. The 1st column lists the required stockout probabilities, ranging from 0.3 to 10^{-5} . The 2nd column provides the hedging points obtained by LD. The 3rd and 4th columns are the hedging points obtained by the FPA algorithm and the optimization algorithm driven by SFM-based IPA derivative estimates. The next two columns provide the hedging points used in the simulation and the stockout probability obtained.

η	LD results	FPA results	SFM results	Simulation results	
	w^*	w^*	w^*	w	$P[X_t(w) \leq 0]$
0.3	7.84	7.60	9.00	8	0.326
0.2	11.22	11.00	11.00	11	0.194
0.1	16.99	17.05	17.27	17	0.0996
0.05	22.77	22.70	22.25	23	$4.852 \cdot 10^{-2}$
0.01	36.18	36.10	36.40	36	$1.040 \cdot 10^{-2}$
$5 \cdot 10^{-3}$	41.96	41.90	42.80	42	$5.061 \cdot 10^{-3}$
10^{-3}	55.37	54.85	55.38	55	$1.072 \cdot 10^{-3}$
$5 \cdot 10^{-4}$	61.14	60.60	61.71	62	$5.220 \cdot 10^{-4}$
10^{-4}	74.56	73.35	75.12	75	$0.964 \cdot 10^{-4}$
$5 \cdot 10^{-5}$	80.33	80.85	81.65	80	$5.283 \cdot 10^{-5}$
10^{-5}	93.74	95.50	93.70	94	$0.983 \cdot 10^{-5}$




Figure 7. The models of demand and production processes in a two-stage system.

To compute analytically the hedging points we use the expression in (20). To compute the prefactors α_i we simulated the system to obtain the expected shortfalls $\mathbf{E}[\hat{L}^i]$, which are independent from the hedging points (since we decomposed the system). These expectations can also be obtained as a by-product of the PA work: while observing a sample path in order to evaluate the IPA estimators used for optimization purposes, information from the same sample path can be readily used to estimate the expectations involved.

Table 2 reports our numerical results. The first two columns list the desired service level for the stage 1 stockout probability and the inventory holdings cost rates for stages 1 and 2, respectively. The 4th column reports the hedging point \mathbf{w}^0 obtained from the LD analysis (cf. (20)). To obtain \mathbf{w}^0 we used in (20) an η_2 service level requirement for stage 2 which we report in the 3rd column. It can be seen that η_2 was selected to be the same or one order of magnitude less than η_1 . As we explained in section 3.3, in this regime our decomposition approach becomes more accurate. The 5th and 6th columns report simulated values for the stockout probability and the expected inventory cost, respectively, under the policy with hedging point equal to \mathbf{w}^0 . Using \mathbf{w}^0 as an initial point we solve problem (2) by employing the algorithm of section 5. The optimal solution \mathbf{w}^* is reported in the 7th column. Columns 8 and 9 list simulated values of the stockout probability and the expected inventory cost, respectively, at the optimal solution. We have also devised a heuristic “descent” approach for approximating the optimal solution of (2) by brute-force simulation. Specifically, for each \mathbf{w} encountered in the course of the heuristic algorithm we simulate the system for several hedging points close to \mathbf{w} and move to the one that satisfies the service level constraint and improves $C(\mathbf{w})$. Again, we initialize the heuristic with \mathbf{w}^0 . In columns 10, 11, and 12 of table 2 we report the hedging point vector obtained by the heuristic (denoted by \mathbf{w}^s) and the corresponding simulated values of the stockout probability and the expected inventory cost, respectively.

Several comments on these results are in order. First, it can be seen that \mathbf{w}^0 is “approximately” feasible, that is, the stockout probability at \mathbf{w}^0 is close to η_1 (in all cases they agree in the order of magnitude and the first significant digit). Moreover, in cases

Table 2
Numerical results for the two-stage supply chain.

η_1	(h_1, h_2)	LD results			SFM results			Simulation results			
		η_2	(w_1^0, w_2^0)	$\mathbf{P}[X^1(\mathbf{w}^0) < 0]$	$C(\mathbf{w}^0)$	(w_1^*, w_2^*)	$\mathbf{P}[X^1(\mathbf{w}^*) < 0]$	$C(\mathbf{w}^*)$	(w_1^s, w_2^s)	$\mathbf{P}[X^1(\mathbf{w}^s) < 0]$	$C(\mathbf{w}^s)$
0.1	(10, 1)	0.01	(7, 46)	0.109	87.42	(8.75, 29)	0.0996	84.37	(8, 33)	0.0985	82.36
0.1	(1, 1)	0.01	(7, 46)	0.109	42.87	(17.4, 12.6)	0.103	20.17	(17, 13)	0.100	20.21
0.1	(1, 10)	0.01	(7, 46)	0.109	384.10	(37.75, 6)	0.0981	63.63	(31, 7)	0.0941	65.17
0.01	(10, 1)	0.01	(16, 47)	$1.36 \cdot 10^{-2}$	174.23	(17.5, 47)	$1.02 \cdot 10^{-2}$	189.07	(17, 52)	$0.994 \cdot 10^{-2}$	189.42
0.01	(5, 1)	0.01	(16, 47)	$1.36 \cdot 10^{-2}$	106.62	(18.4, 42.6)	$1.08 \cdot 10^{-2}$	113.89	(18, 45)	$0.995 \cdot 10^{-2}$	114.40
0.01	(1, 1)	0.01	(16, 47)	$1.36 \cdot 10^{-2}$	52.50	(38.2, 16.6)	$0.968 \cdot 10^{-2}$	44.25	(38, 16)	$1.019 \cdot 10^{-2}$	43.45
0.01	(1, 5)	0.01	(16, 47)	$1.36 \cdot 10^{-2}$	208.39	(52.2, 8.4)	$1.12 \cdot 10^{-2}$	69.72	(56, 8)	$0.919 \cdot 10^{-2}$	72.05
10^{-3}	(5, 1)	10^{-3}	(25, 72)	$1.11 \cdot 10^{-3}$	176.70	(26, 67)	$0.959 \cdot 10^{-3}$	176.66	(26, 69)	$0.918 \cdot 10^{-3}$	178.70
10^{-3}	(1, 1)	10^{-3}	(25, 72)	$1.11 \cdot 10^{-3}$	86.37	(58, 21)	$0.890 \cdot 10^{-3}$	68.41	(58, 20)	$0.986 \cdot 10^{-3}$	67.40
10^{-3}	(1, 5)	10^{-3}	(25, 72)	$1.11 \cdot 10^{-3}$	341.57	(67, 15)	$0.886 \cdot 10^{-3}$	109.77	(67, 14)	$1.010 \cdot 10^{-3}$	105.79
10^{-4}	(5, 1)	10^{-4}	(34, 96)	$1.10 \cdot 10^{-4}$	245.69	(35, 82)	$1.00 \cdot 10^{-4}$	236.69	(35, 83)	$0.979 \cdot 10^{-4}$	237.64
10^{-4}	(1, 1)	10^{-4}	(34, 96)	$1.10 \cdot 10^{-4}$	119.36	(47, 57)	$1.06 \cdot 10^{-4}$	93.36	(47, 58)	$0.986 \cdot 10^{-4}$	94.37
10^{-4}	(1, 5)	10^{-4}	(34, 96)	$1.10 \cdot 10^{-4}$	470.51	(88, 16)	$1.06 \cdot 10^{-4}$	134.82	(89, 16)	$0.965 \cdot 10^{-4}$	135.82

where $h_1 \gg h_2$ we can see that $C(\mathbf{w}^0)$ is quite close to the optimal cost obtained either by the SFM-guided optimization algorithm or by our “descent” heuristic. This is to be expected because in those cases it is profitable to keep a lot of inventory at stage 2 and not starve stage 1, i.e., the system is essentially decomposed. Of course, when $h_2 \gg h_1$ our initial point \mathbf{w}^0 can be far away from optimal, since our decomposition assumption is no longer valid. Still, being able to quickly compute \mathbf{w}^0 is useful since, otherwise, we have no way of determining a meaningful range for \mathbf{w} and a random initial point can be much further away from the optimal.

Finally, we can observe that the optimal solution \mathbf{w}^* obtained by the SFM-guided algorithm of section 5 is in close agreement with \mathbf{w}^s obtained by the descent heuristic. This provides a basic “sanity check” of the SFM-based optimization algorithm we developed. There is, of course, a very substantial efficiency gain by using the SFM-based algorithm vs. the heuristic. At each step of the heuristic we simulate the system for a number of \mathbf{w} 's in the vicinity of the current one to determine where it is profitable to move. Moreover, our steps are always short since we move to one of the \mathbf{w} 's in the immediate vicinity of the current one. In contrast, the SFM-guided algorithm requires no additional simulation from the one we use at each point to evaluate the values of the objective and the constraint. In addition, knowledge of the gradients leads to much faster convergence since it allows us to take potentially very long steps towards the optimum.

7. Conclusions

We have addressed an important and practical inventory control problem in a single-product, serial supply chain. We made rather realistic modeling assumptions. In particular, (i) we imposed a service level constraint on the stockout probability, which, increasingly, becomes a requirement in practice as Quality of Service (QoS) considerations gain significance; and (ii) we accounted for temporal dependencies in the demand and production processes which is critical in accurately modeling demand and failure-prone production facilities.

On the methodological front, we have advanced our agenda to successfully combine large deviations (LD) and perturbation analysis (PA) techniques. We have explored synergies between LD and PA, arising in several domains. Specifically,

- PA provides accurate estimates for relatively large stockout probabilities (e.g., 5%) within reasonable time, and LD approximations become very reliable for relatively small ones (e.g., 10^{-4}). Hence, their combination enables us to accurately estimate a large range of such probabilities.
- LD analysis was used to quickly obtain an initial feasible point for our proposed optimization algorithm which relied on gradients obtained by SFM-based IPA estimates. We provided strong numerical evidence that the proposed algorithm yields a hedging point that minimizes expected linear inventory costs subject to the service level constraint. In fact, the LD analysis can be based on less than perfect models of demand and production processes. We can rely on our IPA-guided algorithm to optimize,

based on the actual realization of these processes. This allows us to avoid modeling errors in representing demand and production stochastic processes.

- PA was also seen to be of direct use in the LD analysis. We use the expected inventory position to fine-tune our asymptotics. This expectation was obtained as a by-product of the IPA estimation.

It should be noted that the IPA approach, which in the paper is limited to a supply chain with up to two stages ($n = 1$ or $n = 2$), can be extended to the more general multistage case ($n > 2$). Of course, the resulting optimization problem will be computationally rather challenging for large values of n .

We conclude by noting that our work can potentially be extended to a multiclass setting. We can consider a serial supply chain that produces multiple products and implements a so called *generalized processor sharing* (GPS) policy at each stage. In particular, at each stage we allocate a constant fraction of the capacity to each class. This can be viewed as “fair” since the performance of a class cannot be compromised at times that other classes are congested, as might happen for example with a priority policy. The GPS dynamics can be approximated by decomposing the system across classes. Hence, the multiclass supply chain is decomposed in single-class chains, one for each class, and the results we have developed here are immediately applicable.

Acknowledgments

Research partially supported by the NSF under a CAREER award ANI-9983221 and grants ACI-9873339, EEC-0088073, DMI-0300359, and DMI-0330171, by the ARO under the ODDR&E MURI2001 Program Grant DAAD19-01-1-0465 to the Center for Networked Communicating Control Systems, by the AFOSR under grant F49620-01-0056, by the AFRL under contract F30603-99-C-0057, and by the EPRI/DOD under contract WO8333-03.

Notes

1. For a more rigorous statement of the Gärtner–Ellis theorem see (Dembo and Zeitouni, 1993).
2. Note that for large values of T , S_T and S_0 in (19) are constants and do not affect the large deviations rate function.

References

- Abate, J., G. Choudhury, and W. Whitt. (1995). “Exponential Approximations for Tail Probabilities in Queues, I: Waiting Times.” *Operations Research* 43(5), 885–901.
- Baccelli, F. and Z. Liu. (1992). “On a Class of Stochastic Recursive Sequences Arising in Queueing Theory.” *Annals of Probability* 20, 350–374.
- Bertsekas, D. (1999). *Nonlinear Programming*, 2nd edn. Belmont, MA: Athena Scientific.
- Bertsimas, D. and I.C. Paschalidis. (2001). “Probabilistic Service Level Guarantees in Make-to-Stock Manufacturing Systems.” *Operations Research* 49(1), 119–133.

- Bertsimas, D., I.C. Paschalidis, and J.N. Tsitsiklis. (1998). "On the Large Deviations Behaviour of Acyclic Networks of G/G/1 Queues." *The Annals of Applied Probability* 8(4), 1027–1069.
- Cassandras, C. and S. Lafortune. (1999). *Introduction to Discrete Event Systems*. Dordrecht: Kluwer.
- Cassandras, C., Y. Wardi, B. Melamed, G. Sun, and C. Panayiotou. (2002). "Perturbation Analysis for On-Line Control and Optimization of Stochastic Fluid Models." *IEEE Transactions on Automatic Control* 47(8), 1234–1248.
- Clark, A. and H. Scarf. (1960). "Optimal Policies for a Multi-Echelon Inventory Problem." *Management Science* 6, 475–490.
- Dembo, A. and O. Zeitouni. (1993). *Large Deviations Techniques and Applications*. Jones & Bartlett.
- Federgruen, A. and P. Zipkin. (1986). "An Inventory Model with Limited Production Capacity and Uncertain Demands I. The Average Cost Criterion." *Mathematics of Operations Research* 11(2), 193–207.
- Glasserman, P. (1997). "Bounds and Asymptotics for Planning Critical Safety Stocks." *Operations Research* 45(2), 244–257.
- Glasserman, P. and S. Tayur. (1995). "Sensitivity Analysis for Base-Stock Levels in Multiechelon Production–Inventory Systems." *Management Science* 45(2), 263–281.
- Gokbayrak, K. and C.G. Cassandras. (2001). "An On-Line 'Surrogate Problem' Methodology for Stochastic Discrete Resource Allocation Problems." *Journal of Optimization Theory and Applications* 108(2), 349–376.
- Kapuściński, R. and S. Tayur. (1999). "Optimal Policies and Simulation Based Optimization for Capacitated Production Inventory Systems." In S. Tayur, R. Ganeshan, and M. Magazine (eds.), *Quantitative Models for Supply Chain Management*, chapter 2, pp. 7–40. Dordrecht: Kluwer.
- Kushner, H. and D. Clark. (1978). *Stochastic Approximation for Constrained and Unconstrained Systems*. Berlin: Springer.
- Liu, Y. and W. Gong. (1999). "Perturbation Analysis for Stochastic Fluid Queueing Systems." In *Proceedings of 38th IEEE Conference on Decision and Control*, pp. 4440–4445.
- National Research Council. (1998). "Visionary Manufacturing Challenges for 2000." Washington, DC: National Academy Press.
- Panayiotou, C.G. and C.G. Cassandras. (1999). "Optimization of Kanban-Based Manufacturing Systems." *Automatica* 35, 1521–1533.
- Panayiotou, C.G., C.G. Cassandras, and P. Zhang. (2002). "On-Line Inventory Cost Minimization for Make-to-Stock Manufacturing Systems." In *Proceedings of the American Control Conference*, Vol. 6, pp. 4469–4474.
- Paschalidis, I. and Y. Liu. (2003). "Large Deviations-Based Asymptotics for Inventory Control in Supply Chains." *Operations Research* 51(3), 437–460.
- Paschalidis, I., Y. Liu, C. Cassandras, and P. Zhang. (2001). "Threshold-Based Control for Make-to-Stock Models: A Synergy between Large Deviations and Perturbation Analysis." In *Proceedings of the 40th IEEE Conference on Decision and Control*, Orlando, FL, pp. 4523–4528.
- Rockafellar, R. (1970). *Convex Analysis*. Princeton University Press.
- Shi, L. and S. Olafsson. (2000). "Nested Partitions Method for Global Optimization." *Operations Research* 48, 390–407.
- The Economist (June 20th, 1998). "A Survey of Manufacturing: Meet the Global Factory."
- Wardi, Y., B. Melamed, C. Cassandras, and C. Panayiotou. (2001). "IPA Gradient Estimators in Single-Node Stochastic Fluid Models." *Journal of Optimization Theory and Applications* 115(2), 369–406.