Proceedings of the 1991 Winter Simulation Conference
Barry L. Nelson, W. David Kelton, Gordon M. Clark (eds.)

APPLICATION OF PERTURBATION ANALYSIS TO (s,S) INVENTORY SYSTEMS

Sridhar Bashyam and Michael C. Fu

College of Business and Management
University of Maryland
College Park, Maryland 20742

ABSTRACT

A considerable amount of interest has been generated in recent years on the use of the Robbins-Monro procedure to optimize stochastic systems via simulation. An important condition for this procedure to converge to the optimal value of the parameter is that the gradient estimator be unbiased. The technique known in the literature as Pertubation Analysis (PA) is especially promising in this respect. While this technique has been predominantly applied to queueing models, Fu (1990) for the first time used PA to derive sample path derivatives of (s,S) inventory systems. His results, however, require the distribution of the aggregated demand during each review period. In this paper, we extend his results to (s,S) systems where the demand has an associated renewal arrival process. We derive a methodology to determine sample path derivatives of the average cost per period with respect to S and s. Our methodology uses the distribution of individual demands rather than the aggregated demand. Preliminary results from simulation experiments indicate that our estimators are unbiased and have very low variance.

1 INTRODUCTION

Consider the standard infinite horizon, single product, periodic review inventory model with full backlogging and independent demands. Under these conditions, Scarf (1960) and Iglehart (1963) showed that an optimal policy can be found within the class of (s, S) policies. However, determining the actual optimal values of (s, S) can be computationally quite complex, and hence, is rarely attempted in real operating systems (Nahmias,1989). The usual means of finding the optimal values is either through dynamic programming or through stationary analysis. The stationary analysis approach can only be applied to restricted cases and usually involves numerical methods which do not allow for easy sensitivity analysis or optimization. The dynamic programming method is a recursive means of finding the optimal values, but is not well-suited for sensitivity analysis. Algorithms to determine approximately optimal policies can be found extensively in literature, but the majority apply only to systems with restrictive demand distributions. A comprehensive comparison of several approximate (s,S) policies can be found in Porteus (1985).

Given the analytical complexity of (s,S) inventory systems, one obvious way to analyze them is via simulation. While evaluating alternate systems through simulation is fairly routine, optimization through simulation is a challenging problem. Suri and Leung (1991) review some of the common methods to be found in the simulation literature on stochastic optimization. The methods they mention are all basically alternative approaches to span the search space of the system parameters of interest, and the list includes Complete Enumeration, Random Search, Coordinate Search, Pattern Search, and Response Surface Methods.

One other approach that has evoked considerable interest within the simulation community in recent years is stochastic approximation, which is based on gradient search techniques. The pioneering work in this area was done by Robbins and Monro (1951). Central to the stochastic approximation approach are the sample path estimates of the derivatives of the objective function with respect to the parameters of interest. Blum (1954) showed that, under certain conditions, the process converges to the optimal value of the parameter with probability 1. One important condition for convergence is that the gradient estimate be unbiased. The difficulty in obtaining such an unbiased estimator, in order to achieve convergence in reasonable time, perhaps explains why this approach - until recently - was rarely adopted. Rubinstein (1986) used the finite difference estimator, which gives a biased estimate of the gradient. The technique of perturbation analysis (PA) originated by Ho and Cao (1983) is one promising technique for obtaining unbiased derivative estimators from a single simulation of the system. Suri (1989) provides an excellent overview of this technique via the G/G/1 queue, and a comprehensive monograph on PA with an upto date bibliography can be found in Ho and Cao (1991). Suri and Leung (1991) used infinitesimal (IPA) estimates within a Robbins-Monro procedure to investigate empirically the optimization of a M/M/1 queue in a single run. Fu (1989) gave a theoretical proof of convergence for a similar IPA-based algorithm.

Application of perturbation analysis to (s,S) inventory systems was addressed, for the first time, by Fu (1990). He considered a periodic review system with full backlogging and general holding and shortage costs. In addition, the demands were assumed to (i) occur once each period, (ii) be independent and identically distributed (iid), and (iii) have a general but continuous distribution function. For this system, he derived the PA derivative estimators for the inventory level and average cost per period with respect to s and q (=S−s), and provided strong consistency proofs by comparison with analytical results.

The primary objective of this paper is to extend his work to the more general setting where the demand is also characterized by a renewal arrival process. Since the demand process will always have an associated arrival process, the distribution of aggregated demand per period has to be viewed as a convolution of a random number of occurrences per period. For the majority of individual distributions, specifying the distribution of such a convolution will be computationally intractable. In this paper, we address this problem by extending the methodology developed by Fu (1990), such that the individual demand distribution and not the aggregated demand distribution is used to derive the PA gradient estimates of the objective function with respect to s and S.

The rest of the paper is organized as follows. Section 1 basically presents the theoretical foundation laid down by Fu (1990). The concept of perturbation analysis as applied to (s,S) systems is explained via the single demand per period model and expressions for the sample path estimates of the gradients are derived. In Section 2, we focus our attention on the general setting where the demand arrival process is specified. We use the basic principles developed in Section 1 to develop a methodology that could be effectively used to derive sample path gradient estimates for this more general setting. In developing a methodology for the above system, a major difficulty is the absence of any closed-form solutions, even for special cases. Consistency of the estimates can be verified through the use of "sample path proofs" (see, e.g., Glasserman and Gong,1990), but we do not attempt such proofs in this paper. Instead, in Section 3 we provide some computational results which compare our PA estimates with those using the finite difference method. Finally, in Section 4, we discuss related problems that we are currently investigating, as well as further potentially promising avenues of research.

2 GRADIENT ESTIMATORS FOR AVERAGE COST PER PERIOD

We consider the single demand per period model, and develop the expressions for the gradient estimators using perturbation analysis. The objective function is the average cost per period, which we shall denote as $\bar{L}$(s,S). Thus, we are interested in estimators of the partial gradients $\delta\bar{L}/\delta S$ and $\delta\bar{L}/\delta s$. We shall first derive the expressions for estimating $\delta\bar{L}/\delta S$. We shall then show that this solves the problem of estimating $\delta\bar{L}/\delta s$ as well.

2.1 Infinitesimal Perturbation Analysis : The IPA Component of $\delta\bar{L}/\delta S$

The intuitive idea behind the derivation of sample path estimators using perturbation analysis is a thought experiment of introducing a perturbation into the sample path and tracing its effect. Figure 1 shows the sample paths for the inventory system operating at S (called the Nominal Path) and $S + \Delta S$ (called the Perturbed Path). The vertical lines in the figure represent review periods. At the beginning of each period, depending upon the observed inventory level, a decision is taken on whether to order or not. It is assumed that order lead time is zero. After this decision is taken, the aggregated demand for the period is subtracted out. As can be seen, the figure depicts the situation where an infinitesimal change in S produces an infinitesimal change in the objective function—in this case, the average cost per period. The PA literature refers to this case as one in which there is no change in the sequence of events in the perturbed path. In this case IPA can be applied (see Ho and Cao 1983), the derivation of which turns out to be quite straightforward.

Derivation of $(\delta\bar{L}/\delta S)_{IPA}$

Let N = length (in terms of number of periods) of the simulation run,

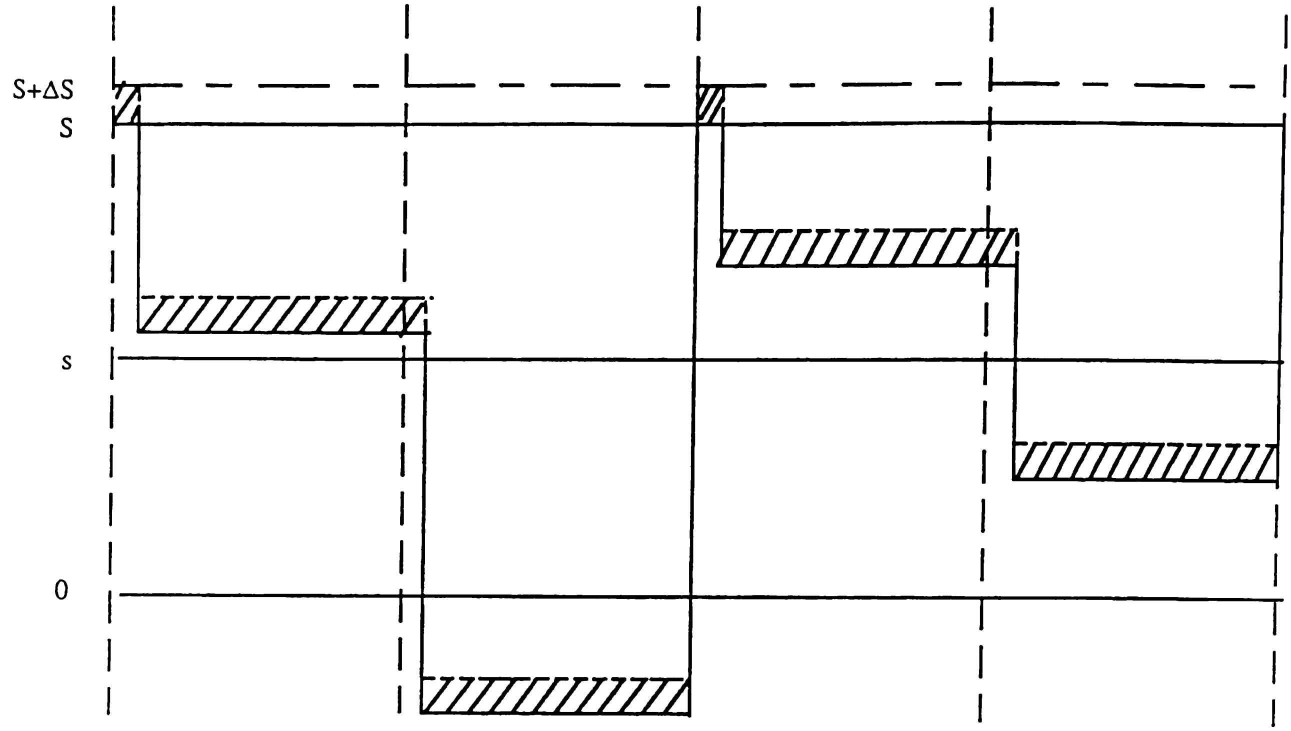

Figure 1: Infinitesimal Change between Nominal and Perturbed Paths

N^+ = number of periods in which the inventory level is positive,
N^- = number of periods in which the inventory level is negative,
h = holding cost per period,
and p = penalty cost per period.

From the sample path it is clear that :

$$\Delta \bar{L} = \frac{h[(\Delta S) * N^+] - p[(\Delta S) * N^-]}{N}$$

Therefore

$$(\delta \bar{L}/\delta S)_{IPA} = \lim_{\Delta S \to 0} \frac{\Delta \bar{L}}{\Delta S} = \frac{h[N^+] - p[N^-]}{N} \quad (1)$$

The above expression is actually derived under the additional assumption that during any period, the nominal and perturbed paths are both either positive or negative. This assumption is justified since the probability of the inventory levels of the two paths having opposite signs in any given period is negligible.

2.2 Smoothed Perturbation Analysis : The SPA Component of $\delta \bar{L}/\delta S$

If there were no possibility of a change in the order of events in the perturbed path, then the above expression for $(\delta \bar{L}/\delta S)_{IPA}$ would, in fact, be an unbiased estimator for $(\delta \bar{L}/\delta S)$. However, consider the situation shown in Figure 2. The figure shows the two paths conforming to the IPA assumptions until the epoch t_0. At that point, we observe an ordering change in the perturbed path. More specifically, it can be observed that in the nominal path an order is placed which raises the inventory level to S. In the perturbed path, the inventory level is greater than s and, hence, an order is not placed. From this point on, we can observe a finite change in the objective function. At epoch t_1, the two paths converge, and from then on proceed according to IPA assumptions. To derive the sample path derivative estimator for this case, we use the technique known as Smoothed Perturbation Analysis (SPA) formulated by Gong and Ho (1987). We derive the SPA component of $\delta \bar{L}/\delta S$ by the following expected conditional contribution, under the limit $\Delta S \longrightarrow 0$:

$$\text{E[ordering change effect]} = \sum_{i=1}^{N} E[\Delta \bar{L} \mid \omega \in \Omega_i(\Delta S)] * P\{\omega \in \Omega_i(\Delta S)\} \quad (2)$$

where $\Omega_i(\Delta S)$ = {sample paths where ordering change occurs in period i due to ΔS}. The first term

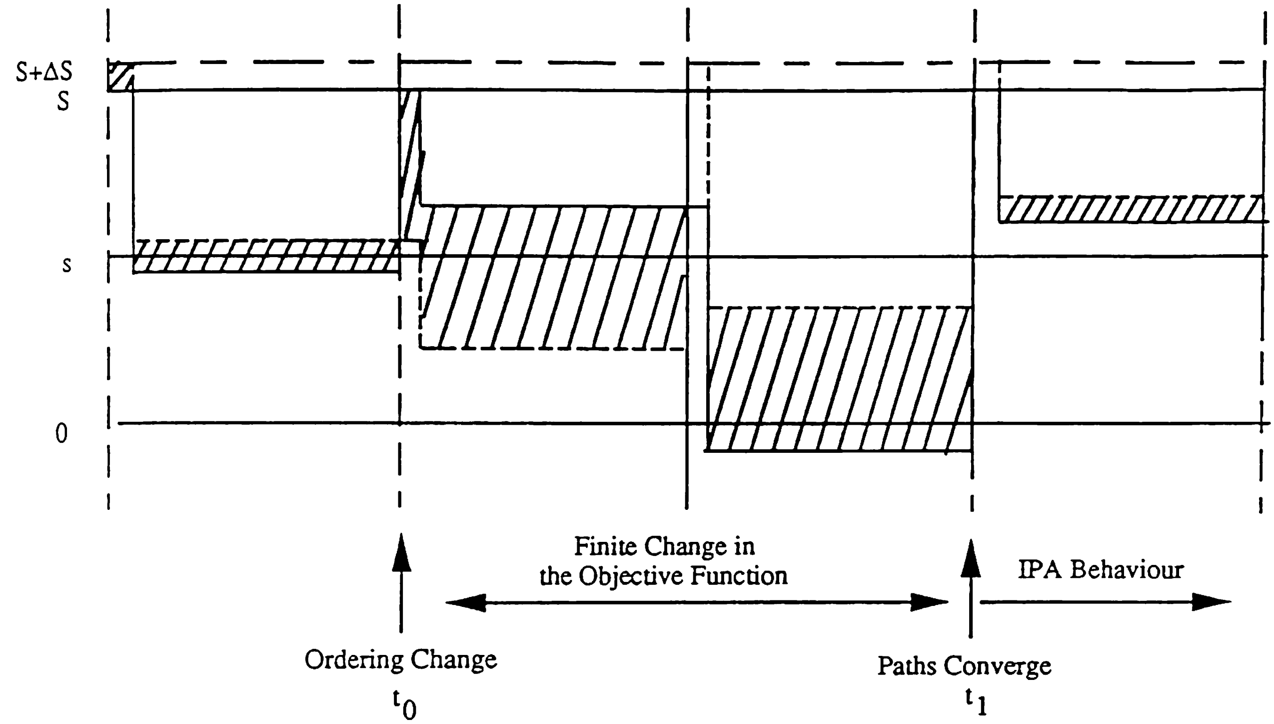

Figure 2: Finite Change between Nominal and Perturbed Paths

in the above expression is the expected finite change in cost caused by an ordering change, and can in general be obtained through a single "off-line" simulation run. Figure 3 shows how this is done and is pretty much self-explanatory. We start the run by initializing the inventory levels of the perturbed and nominal paths to be infinitesimally above and below s, respectively. We then take the reordering decision immediately, which thus sets up an ordering change between the two paths. The sample paths are then simulated until they converge at epoch t_1. We denote the period t_0 to t_1 as a cycle. At t_1 we re-initialize and restart the sample paths. As can be seen, the cycles here form a renewal process. The simulation is run for an adequately large number of cycles, from which the "off-line" estimator denoted by $E[\Delta L \mid \Delta S \to 0]$, $\Delta \bar{L} = \Delta L/N$, is computed as

$$E[\Delta L \mid \Delta S \longrightarrow 0] = \frac{(L_p - L_n)}{C}$$

where L_p = total cost of the perturbed path,
L_n = total cost of the nominal path,
C = total number of cycles.

We shall derive expressions for both the right-hand derivative ($\Delta S > 0$), as well as the left-hand derivative ($\Delta S < 0$) of $\bar{L}$. This becomes useful since, depending on the system parameters, one may yield a lower variance than the other. This aspect is discussed further in section 2. Denoting these derivatives as $(\delta\bar{L}/\delta S)^{+}_{SPA}$ and $(\delta\bar{L}/\delta S)^{-}_{SPA}$ respectively, we have

$$(\delta\bar{L}/\delta S)^{+}_{SPA} = \frac{1}{N} E[\Delta L \mid \Delta S \to 0] * \sum_{i=1}^{N} \lim_{\Delta S \to 0+} \frac{P\{\omega \in \Omega_i(\Delta S^{+})\}}{\Delta S} \tag{3}$$

$$(\delta\bar{L}/\delta S)^{-}_{SPA} = \frac{1}{N} E[\Delta L \mid \Delta S \to 0] * \sum_{i=1}^{N} \lim_{\Delta S \to 0-} \frac{P\{\omega \in \Omega_i(\Delta S^{-}\}}{\Delta S} \tag{4}$$

To derive an expression for the probability term, we have to identify the situation where there exists the *potential* for an ordering change. Once this is identified, we compute the probability that an ordering change would *actually* occur, *given* such a situation. The situation for a potential ordering change for the right-hand derivative is shown in Figure 4, while that for the left-hand derivative is shown Figure 5. Figures 4 and 5 define some sample path quantities that are required for the derivations. For the right-hand derivative, Figure 4 shows that a potential change exists at reorder point n if the inventory level X_n is below s and the inventory level X_{n-1} is above s. For

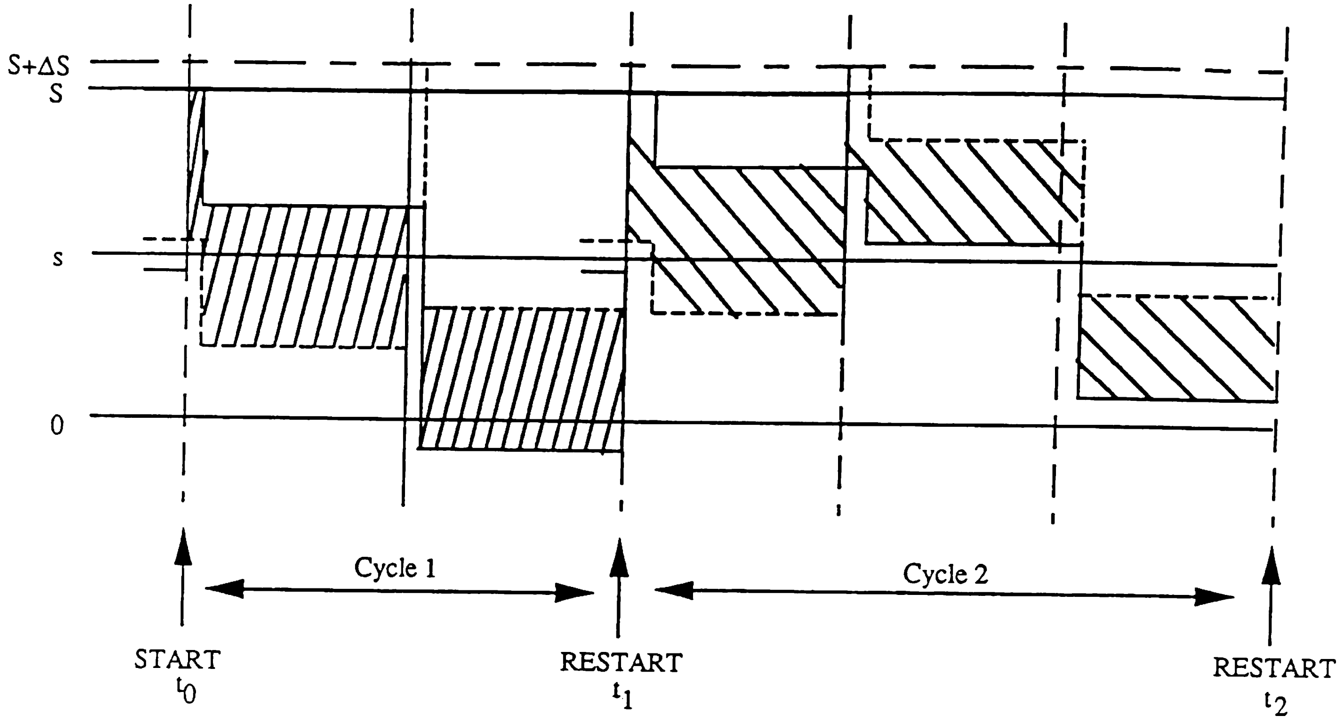

Figure 3: Offline Simulation to determine E[ordering change effect]

the left-hand derivative, the requirement is that both X_n and X_{n-1} be above s. In both cases, the change would actually have occurred if $\alpha \leq \Delta S$. We now derive the required expressions for the probability term as follows :

Let the demand size be denoted by the random variable D whose density and distribution functions are g(.) and G(.) , respectively.

For the case of the right-hand derivative

$$\begin{aligned} \text{P}\{\text{Ordering Change}\} &= \text{P}[\,\alpha \leq \Delta S \mid D > Z] \\ &= \text{P}[\,D - Z \leq \Delta S \mid D > Z] \\ &= \text{P}[\,D \leq Z + \Delta S \mid D > Z] \\ &= \frac{G(Z+\Delta S) - G(Z)}{1 - G(Z)} \end{aligned}$$

Therefore,

$$\begin{aligned} &\lim_{\Delta S \to 0+} \frac{P[\alpha \leq \Delta S \mid D > Z]}{\Delta S} \\ &= \lim_{\Delta S \to 0+} \left[\frac{G(Z+\Delta S) - G(Z)}{\Delta S}\right]\left[\frac{1}{1-G(Z)}\right] \\ &= \left[\frac{g(Z)}{1-G(Z)}\right] \end{aligned}$$

For the case of the left-hand derivative, we proceed likewise to get

$$\lim_{\Delta S \to 0-} \frac{P[\alpha \leq \Delta S \mid D < Z]}{\Delta S} = \left[\frac{g(Z)}{G(Z)}\right]$$

Thus, if we denote $\Re^+$ and $\Re^-$ to be the set of all reorder points representing a potential ordering change for the right-hand and left-hand derivatives respectively, the final expressions for the SPA contribution are :

$$\begin{aligned} (\delta\bar{L}/\delta S)^+_{SPA} &= \frac{1}{N} E[\Delta L \mid \Delta S \longrightarrow 0] \\ &\quad * \sum_{j \in \Re^+} \left[\frac{g(Z_j)}{1-G(Z_j)}\right] \end{aligned} \tag{5}$$

$$\begin{aligned} (\delta\bar{L}/\delta S)^-_{SPA} &= \frac{1}{N} E[\Delta L \mid \Delta S \longrightarrow 0] \\ &\quad * \sum_{j \in \Re^-} \left[\frac{g(Z_j)}{G(Z_j)}\right] \end{aligned} \tag{6}$$

The total gradient $(\delta\bar{L}/\delta S)$ is finally estimated as the sum of the IPA and SPA components, where the SPA component can be represented by either the right-hand or the left-hand derivatives.

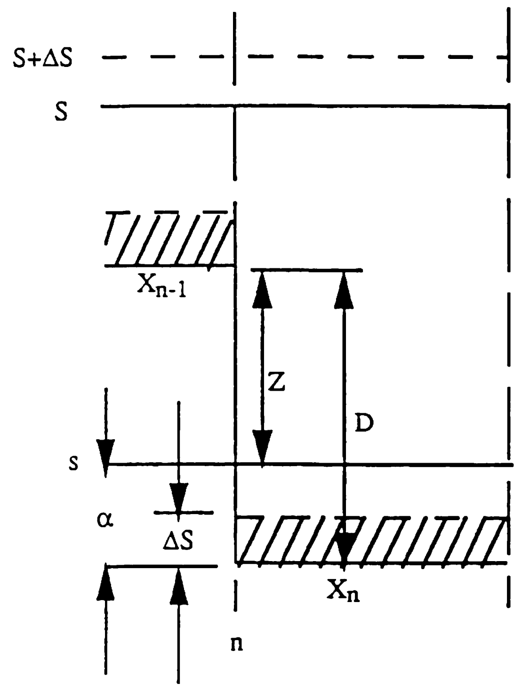

Figure 4: Potential Ordering Change situation for the Right-Hand Derivative

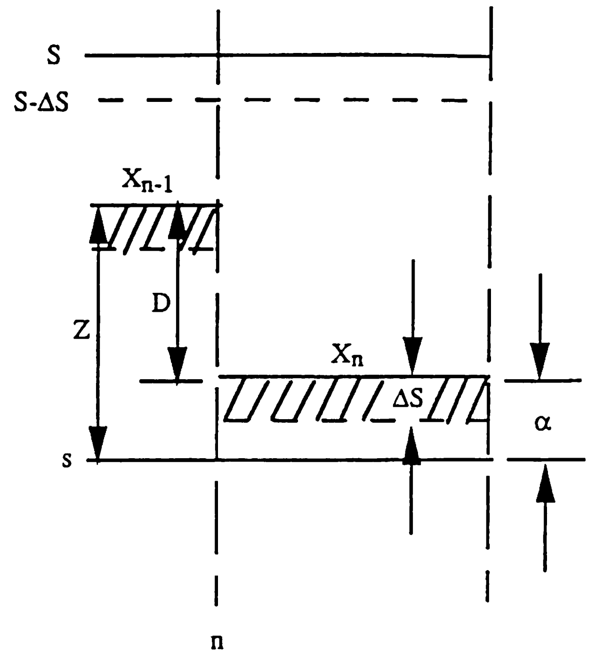

Figure 5 : Potential Ordering Change situation for the Left-Hand Derivative

2.3 PA Estimators of $\delta \bar{L}/\delta s$

For the estimators of $(\delta \bar{L}/\delta s)$, we do the IPA and SPA analysis in a similar fashion. Let us first consider the case when there is no ordering change between the nominal and perturbed paths. For this case, the nominal and perturbed paths are, in fact, identical. A change from s to $s+\Delta s$ does not change the sample path at all, as long as the sequence of events remain unaltered. Hence, it follows that

$$(\delta \bar{L}/\delta s)_{IPA} = 0 \tag{7}$$

For the SPA component, potential ordering change situations for the right-hand and left-hand derivatives are shown in Figures 6 and 7 respectively. From the above Figures, and following the derivations resulting in equations (5) and (6), we get

$$\begin{aligned} (\delta \bar{L}/\delta s)^{+}_{SPA} &= -\frac{1}{N} E[\Delta L \mid \Delta s \longrightarrow 0] \\ &\quad * \sum_{j \in \Re^{+}} \left[\frac{g(Z_j)}{G(Z_j)}\right] \end{aligned} \tag{8}$$

$$\begin{aligned} (\delta \bar{L}/\delta s)^{-}_{SPA} &= -\frac{1}{N} E[\Delta L \mid \Delta s \longrightarrow 0] \\ &\quad * \sum_{j \in \Re^{-}} \left[\frac{g(Z_j)}{1-G(Z_j)}\right] \end{aligned} \tag{9}$$

In the above equations, the sets $\Re^{+}$ and $\Re^{-}$ have a similar interpretation as those in equations (5) and (6).

3 Methodology for a General Arrival Process

In this section we consider the more realistic situation where the demands are also characterized by an associated renewal arrival process. The approach used here is largely similar to that discussed in the previous section.

3.1 The IPA Estimator

The expression for the IPA estimator remains the same. The only change here is that the terms N^{+} and N^{-} in equation (1) are replaced by the total time that the inventory level is positive (T^{+}) and negative (T^{-}) respectively. Defining each period to have a length equal to one time unit, we have

$$(\delta \bar{L}/\delta S)_{IPA} = \frac{h[T^{+}] - p[T^{-}]}{N} \tag{10}$$

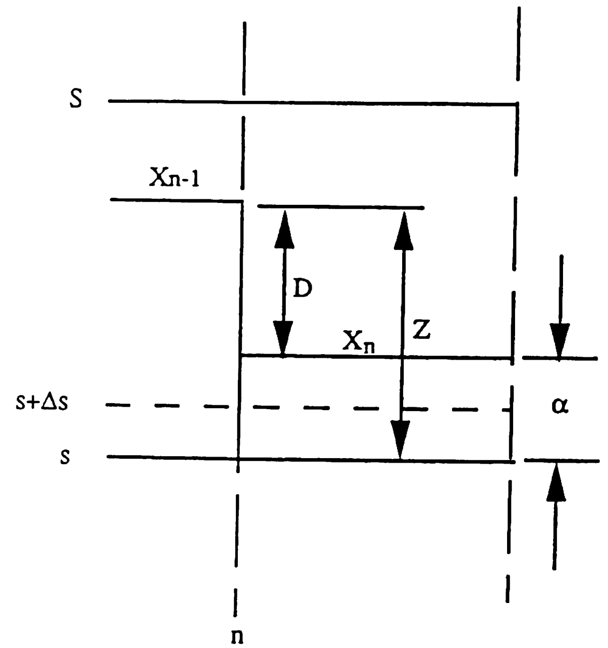

Figure 6: Potential Ordering Change situation for the Right-Hand Derivative

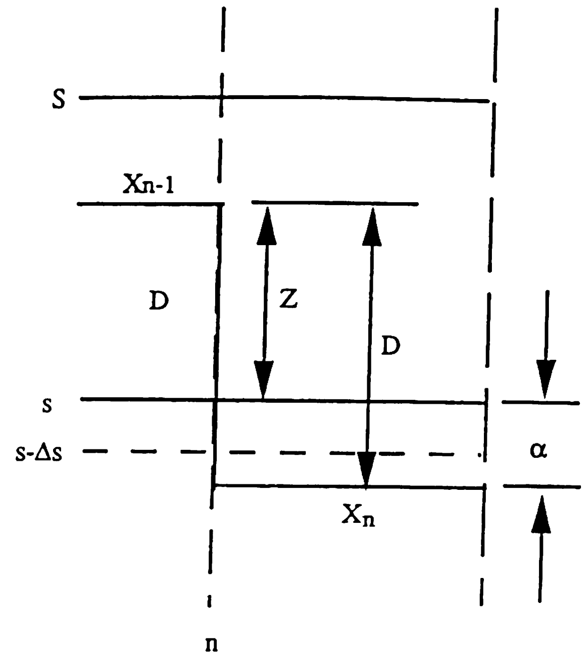

Figure 7 : Potential Ordering Change situation for the Left-Hand Derivative

3.2 The SPA Estimator

As before, we derive estimators for both the right-hand and left-hand derivatives. The expressions are of the same general form as equations (3) and (4). In fact, the term $E[\Delta L \mid \Delta S \longrightarrow 0]$ is estimated in exactly the same manner via an offline run. Difficulties here can arise if the inter-arrival process has a general distribution, the reason for which will be clarified in the next section.

To estimate the probability term, we adopt a different approach here. The difference lies in the way we identify a situation that has the potential to cause a potential change. This is explained via Figure 8 shown below.

After each demand occurrence, we inspect the inventory levels just prior to the occurrence, and just after the occurrence. As before, in the case of the right-hand derivative, there is a potential for an ordering change if these levels are above and below s, respectively. Such a situation is represented by epoch t_3. Similarly, for the left-hand derivative, potential for ordering change exists if both levels are above s, and is represented by epochs t_1 and t_2. However, the actual occurrence of the ordering change is dependent on both the demand size as well as the occurrence of the next demand. With reference to Figure 8, an ordering change would actually occur if $\alpha \leq \Delta S$, and $A > Y$, where the random variable A represents the time between successive demands. Thus if F(.) denotes the distribution function of the time between successive demand arrivals, the expressions for the SPA contributions are now computed as

$$(\delta\bar{L}/\delta S)^{+}_{SPA} = \frac{1}{N}E[\Delta L \mid \Delta S \rightarrow 0] * \sum_{j\in\Re+}\left\{\left[\frac{g(Z_j)}{1-G(Z_j)}\right] * [1-F(Y_j)]\right\} \quad (11)$$

$$(\delta\bar{L}/\delta S)^{-}_{SPA} = \frac{1}{N}E[\Delta L \mid \Delta S \longrightarrow 0] * \sum_{j\in\Re-}\left\{\left[\frac{g(Z_j)}{G(Z_j)}\right] * [1-F(Y_j)]\right\} \quad (12)$$

In equation (11), $\Re^+$ represents the set of demand arrival epochs that are potential ordering change situations in the case of the right-hand derivative. Similarly, $\Re^-$ in equation (12) represents the corresponding set for the left- hand derivative. We conclude this section with the following comments :

(i) Having the opportunity to compute both the right-hand as well as the left-hand derivatives is important from the viewpoint of getting estimates that have as low a variance as possible.

(ii) We can now clarify one major difficulty that

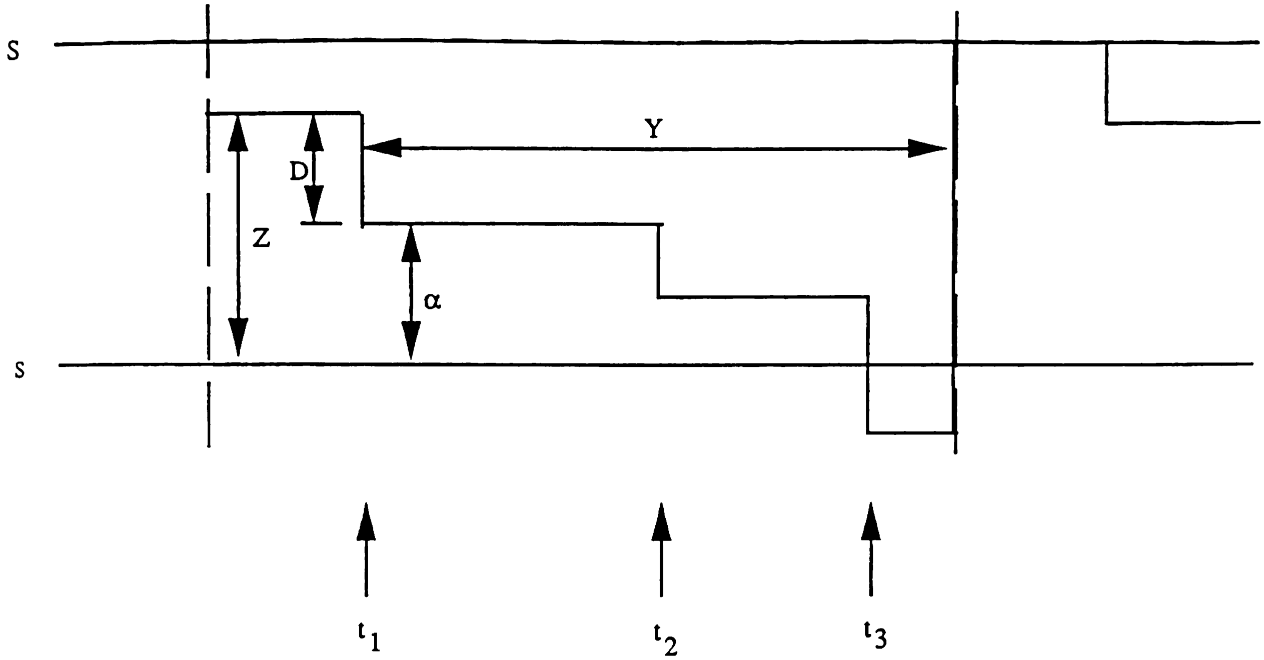

Figure 8: Conditions for an ordering change in a general (s,S) model having a renewal arrival process

arises if the inter-arrival times have a general distribution. While doing the off-line simulation, it is not strictly correct to start the run by generating an inter-arrival time from the given distribution. This can be seen by referring to Figure 8. The figure shows a potential ordering change for the left-hand derivative at epoch t_1. When we perform the thought experiment of determining the expected change in cost given that ordering change does occur at t_1, it becomes clear that the arrival time of the first demand in the off-line simulation would have to be generated from the distribution of the residual life, i.e. $P[A \leq Y + a \mid A > Y]$. Thus, the expected effect per ordering change is no longer constant, and hence, cannot be estimated via a single run. We are, at this time, unable to comment on how serious the inaccuracy would be if we were to ignore the residual nature of the first arrival. For the computational results reported in Section 3, a Poisson arrival process is assumed for which, due to the memoryless property, it is valid to generate a fresh exponential arrival time for the first demand.

4 COMPUTATIONAL RESULTS

(1) For the single demand per period discussed in Section 1, analytical expressions for the IPA and SPA components are available for the special case where the demand size has an exponential distribution with mean 1. It can be shown that the following analytical expressions hold :

$$(\delta\bar{L}/\delta S)_{IPA} = \frac{h\left(1+S-s-e^{-S}\right)-pe^{-S}}{1+S-s}$$

$$(\delta\bar{L}/\delta S)_{SPA}$$

$$= -\left[\frac{k+h\left\{(S-s)\left(1+\frac{S-s}{2}-e^{-S}\right)\right\}-pe^{-S}(S-s)}{(1+S-s)^2}\right]$$

For the right-hand derivative, the term $g(Z)/[1-G(Z)]$ is simply the reciprocal of the mean demand, which, in this case happens to be 1. Thus, we can expect the term $g(Z)/G(Z)$ for the left-hand derivative to be also equal to one in this case. Furthermore, it can be shown that for $S - s = 1$,
the total number of "samples" for the right-hand derivative
= the total number of "samples" for the left-hand derivative
$= N/2$.
Using the above expressions, we tested the validity of the IPA and SPA expressions for the

following (s,S) system :

* $S = 3$
* $s = 2$
* Demand $\sim$ exp(1) as required.
* holding cost per period = 10
* penalty cost per period = 50
* set-up cost = 10

We performed 10 replications and the results are shown in Table 1, in the form of 95% Confidence Intervals. As can be seen, the PA estimators exhibit a very low variance, and are strongly consistent.

(2) To test the PA estimators for the general setting discussed in Section 2, we considered the following (s,S) system :

* $S = 10$
* $s = 5$
* holding cost per unit time = 10
* penalty cost per unit time = 50
* setup cost = 10
* inter-arrival times $\sim$ exp(mean = 0.3125)

For the demand size, we considered three different distributions :

(i) Exponential with mean = 2,

(ii) Uniform [0 , 10]

(iii) Weibull (4, 3)

For the Weibull distribution, the value of 3 was chosen for the shape parameter to make the distribution resemble a normal distribution. The scale parameter was given a value of 4 so that the generated demands would be reasonable compared to the values of S and s.
For each case, we ran 10 replications and compared the PA derivatives with those derived via the finite difference method (FDM) using common random numbers. The results are presented in Table 2, again in the form of 95% Confidence Intervals. From the table, it can be seen that there appears to be good agreement between the two approaches. One can also observe very tight confidence intervals for the PA estimates as compared to the noisy estimates produced by the finite difference method.

5 CONCLUSIONS

In this paper, we have made a preliminary attempt at applying perturbation analysis to general (s,S) inventory models. The results presented in this paper indicate that a vast potential exists for such an application. Our major motivation lies in the fact that PA estimators of the required derivatives seem to exhibit remarkably low variance. If, in addition, consistency proofs of these estimators can be established, then use of PA within a Robbins-Monro type procedure for stochastic optimization of (s,S) models becomes extremely attractive. It is generally accepted that the finite differences method can yield unreliable estimates of sample path derivatives, which would result in slow convergence of the optimization procedure. However, a lot of additional research remains to be done. Some directions presently under way include :

- Deriving a sample path proof for the consistency of the PA estimators in Section 2.
- Application of PA for other objective function criteria, such as service levels.
- Application of PA to continuous review models.
- Application of PA to random lead time models.

REFERENCES

Blum, J.R. 1954. Approximation Methods which Converge with Probability One.*Annals of Mathematical Statistics* **25**: 382-386.

Fu, M.C. 1990. Convergence of a Stochastic Approximation Algorithm for the GI/G/1 Queue Using Infinitesimal Perturbation Analysis. *Journal of Optimization Theory and Applications* **65**: 149- 160.

Fu, M.C. 1990. Sample Path Derivatives for (s,S) Inventory Systems. Submitted to Operations Research.

Glasserman, P. and W.B. Gong. 1990. Smoothed Perturbation Analysis for a Class of Discrete-Event Systems. *IEEE Transactions on Automatic Control* **35**: 1218-1230.

Gong, W.B. and Y.C. Ho. 1987. Smoothed Perturbation Analysis of Discrete- Event Dynamic Systems. *IEEE Transactions on Automatic Control* **32**: 858-867.

Ho, Y.C. and X.R. Cao. 1983. Optimization and Perturbation Analysis of Queueing Networks. *Journal of Optimization Theory and Applications* **40**: 559-582.

Ho, Y.C. and X.R. Cao. 1991. *Pertubation Analysis and Discrete Event Dynamic Systems.* Kluwer Academic.

Table 1
Computational Results for the Analytical Case
Derivatives of Average Cost/Period with respect to S

	IPA	SPA	Total Gradient	E[g(Z)/G(Z)]	Prob. of Order Change: LH Derivative
Sample Path Estimates	5.933±0.037	-4.234±0.042	1.700±0.056	1.000±0.015	0.501±0.001
True Value	5.940	-4.220	1.720	1.000	0.500

Table 2
Computational Results for the general (s,S) Inventory Model
Gradient of Average Cost/Period with respect to S

	IPA	SPA	Total Gradient	FDM(1) ΔS=+0.02	FDM(2) ΔS=-0.02	FDM(3) ΔS=±0.01
Exponential	3.956±0.026	-0.863±0.041	3.093±0.049	3.325±0.636	3.097±0.532	3.202±0.524
Uniform	-9.993±0.038	1.844±0.023	-8.109±0.033	-8.292±0.351	-8.476±0.509	-8.468±0.354
Weibull	-3.766±0.029	1.589±0.025	-2.177±0.025	-2.235±0.466	-2.269±0.394	-2.319±0.504

Iglehart, D. 1963. Optimality of (s,S) Policies in the Infinite Horizon Dynamic Inventory Problem. *Management Science* 9: 259-267.

Nahmias, S. 1989. *Production and Operations Analysis.* Boston:Irwin Press.

Porteus, E.L. 1985. Numerical Comparisons of Inventory Policies for Periodic Review Systems. *Operations Research* **33**: 134-152.

Robbins, H. and S. Monro. 1951. A Stochastic Approximation Method. *Annals of Mathematical Statistics* **22**: 400-407.

Rubinstein, Y.R. and J. Kreimer. 1986. Inventory Models Under Uncertainty: an Adaptive Approach. *Mathematics and Computers in Simulation* **28**: 169-188.

Scarf, H. 1960. The Optimality of (S,s) Policies in the Dynamic Inventory Problem. In: *Mathematical Methods in the Social Sciences*, eds. K. Arrow, S. Karlin and P. Suppes. Stanford: Stanford University Press.

Suri, R. 1989. Perturbation Analysis: The State of the Art and Research Issues Explained Via the G/G/1 Queue. *Proceedings of the IEEE* **77**: 114-137.

Suri, R. and Y.T. Leung. 1991. Single Run Optimization of Discrete Event Simulations - An Empirical Study using the M/M/1 Queue. *Proceedings of IIE* **21**: 35-49.

AUTHOR BIOGRAPHIES

SRIDHAR BASHYAM is a doctoral student in management science in the College of Business and Management at the University of Maryland, College Park. He received a B.Tech degree in Mechanical Engineering from the Indian Institute of Technology, Madras (India) in 1979 and a MBA from the Indian Institute of Management, Calcutta (India) in 1983. His doctoral thesis deals with stochastic optimization of inventory systems.

MICHAEL C. FU is an assistant professor in management science in the College of Business and Management at the University of Maryland. He received an S.B. and S.M. in electrical engineering and an S.B. in mathematics from MIT in 1985 and a Ph.D. in applied mathematics from Harvard University in 1989. His research interests include stochastic derivative estimation and stochastic optimization, particularly with applications towards manufacturing systems and inventory control.