

MONTE CARLO OPTIMIZATION OF STOCHASTIC SYSTEMS: TWO NEW APPROACHES

P. W. Glynn and J. L. Sanders
Department of Industrial Engineering
University of Wisconsin-Madison
Madison, Wisconsin

ABSTRACT

The design of modern manufacturing systems presents a number of challenges. In particular, the stochastic nature of machine failures in combination with the large number of decision variables makes optimization of such systems difficult. In this paper, we present two new approaches to optimization of the complex stochastic systems that arise in a manufacturing context; both are Monte Carlo simulation-oriented, and are therefore broadly applicable. The first technique involves using a likelihood ratio gradient estimate to drive a Robbins-Monro algorithm, and is relevant to problems in which the decision variables are continuous. The second idea employs homotopy methods to follow an "optimal path" in decision variable space, and can be used for both discrete and continuous optimization.

1. INTRODUCTION

The problem of deriving an optimal system design for a modern automatic assembly or manufacturing system is a formidable task. A machine may involve up to 100 separate assembly or fabrication stations. For each station, there are design variables which affect the corresponding cycle time and buffer area. In addition, there are questions of station sequencing, the number of pallets to be placed on the line, whether or not parallel processing is to be used and if so of what type. Furthermore, when inspection and repair are required, one must decide whether or not to use a multiple loop system or to include the functions in a single loop design. Finally, there are a variety of gating and control rules that must be selected for implementation by the programmable controllers that operate various sectors of the machine. In total, there may be as many as 200 to 300 decision variables that will affect the machine's ultimate productivity. In addition, the complexity of the problem is greatly increased by the stochastic nature of the station failures in such systems.

Optimizing the design of these complex machines is therefore a major focus of current research in the

area. In this paper, we will present two new approaches for optimizing the type of complex stochastic systems which arise in manufacturing. Both methods are Monte Carlo simulation-oriented, and are therefore applicable to an enormous variety of different problems. In Section 2, we describe, in a precise sense, the stochastic optimization problem to be considered here. Sections 3 through 5 discuss our new techniques, while Section 6 summarizes our results.

2. FORMULATION OF THE STOCHASTIC OPTIMIZATION PROBLEM

An essential requirement for the development of stochastic optimization algorithms is a sound probabilistic framework for their analysis. The following problem will serve to motivate our discussion.

(2.1) EXAMPLE. Consider an $M/M/1/\infty$ queue (see [3], p. 273) in which customers arrive at rate λ . The goal of the engineer is to determine an optimal service rate μ .

We formalize the stochastic optimization problem as follows. Let $\theta \in \mathbb{R}^d$ represent the decision variable, and let $K \subseteq \mathbb{R}^d$ correspond to feasible choices of the decision variable θ . As is standard in the analysis of stochastic systems, Ω shall represent the set of possible sample outcomes ω of the stochastic system and F corresponds to a suitable σ -field of subsets (events) contained in Ω .

(2.1) EXAMPLE (continued). Here $d = 1$, $\theta = \mu$, and $K = (0, \infty)$. The sample space Ω consists of the set of right-continuous functions $\omega = \omega(\cdot)$ with left limits, taking values in $Z^+ = \{0, 1, \dots\}$; the function $\omega(\cdot)$ should be viewed as the queue-length process. Specifically, $\Omega = D[0, \infty)$ and F is the associated Borel σ -field (see [4], p. 116-127).

The concept of optimality requires an objective function of some kind. Given the outcome ω and decision parameter θ , let $Z(\theta) = Z(\theta, \omega)$ be the "cost" incurred by the system.

(2.1) EXAMPLE (continued). Suppose that the system pays out one dollar per unit of time that each customer is in the system, and that the server is paid $b\mu$ dollars per unit of time. The long-run cost of operating the queue is then given by

$$Z(\theta, \omega) = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \omega(s) ds + b\theta.$$

The distribution on Ω is clearly affected by the choice of θ . Specifically, the probability measure P_θ on (Ω, \mathcal{F}) is allowed to depend on θ .

(2.1) EXAMPLE (continued). The probability P_θ is that associated with a continuous-time Markov chain on Z^+ having initial distribution $\mu(\theta)$ and generator $Q(\theta)$, where

$$Q_{ij}(\theta) = \begin{cases} \lambda & , & j = i + 1, & i > 0 \\ -\lambda - \theta & , & j = i, & i > 0 \\ -\theta & , & j = i - 1, & i > 1 \\ 0 & , & \text{else.} \end{cases}$$

The goal of the engineer is to minimize, over the decision variable θ , the expected cost of running the system. Mathematically, the stochastic optimization problem is:

$$\min_{\theta \in K} r(\theta) \quad (2.2)$$

where $r(\theta) = E_\theta Z(\theta) \equiv \int_\Omega Z(\theta, \omega) P_\theta(d\omega)$.

(2.1) EXAMPLE (continued). For the $M/M/1/\infty$ queue, it is easy to calculate that $r(\theta) = \lambda(\theta - \lambda)^{-1} + b\theta$. The minimizer is then given by $\theta^* = \lambda + (\lambda/b)^{1/2}$.

We shall say that (2.2) is a continuous parameter stochastic optimization problem if K is an open set, and discrete parameter if K is a discrete set.

3. CONTINUOUS PARAMETER MONTE CARLO OPTIMIZATION

For the next two sections, we specialize (for the sake of simplicity) to the case where the decision variable is a scalar (i.e. $d = 1$). In the continuous parameter setting, the optimization problem (2.2) can be analyzed in terms of finding the root θ^* to the equation $r'(\theta) = 0$ (provided, of course, that r is differentiable with a minimizer on K). One iterative method for accomplishing this (when r is convex) then takes the form

$$\theta_{k+1} = \theta_k - a_k r'(\theta_k) \quad (3.1)$$

where a_k is a suitably chosen sequence of constants tending to zero. A difficulty in implementing (3.1) for the complex stochastic systems under consideration here is that $r'(\theta)$ is not derivable in closed form. In fact, $r(\theta)$ is, in general, not known and must itself be estimated via Monte Carlo techniques. Such algorithms are well-known and easily constructed; on the other hand, Monte Carlo procedures for estimating the derivative $r'(\theta)$ are not widely available.

As a consequence, it seems (at first) reasonable to limit oneself to optimization algorithms in which only estimators for $r(\theta)$ are required (as opposed to the derivative $r'(\theta)$). The Kiefer-Wolfowitz (KW) algorithm [6] circumvents the need to estimate $r'(\theta)$

by instead estimating a central difference approximation to $r'(\theta)$. Specifically, let

$$\theta_{k+1} = \theta_k - a_k \left[\frac{1}{2c_k} (Y_k(\theta_k + c_k) - Y_k(\theta_k - c_k)) \right]$$

where $Y_k(\theta_k + c_k)$ and $Y_k(\theta_k - c_k)$ are conditionally independent given $\theta_1, \dots, \theta_k$ and $E(Y_k(\theta_k \pm c_k) | \theta_1, \dots, \theta_k) = r(\theta_k \pm c_k)$ (a_k and c_k are deterministic constants tending to zero). The convergence rate of θ_n to θ^* has been extensively studied (see, for example, [9]) and the resulting analysis has shown that convergence (except in rare cases) is slower than $O(n^{-1/2})$. Since the typical convergence rate for Monte Carlo algorithms is $O(n^{-1/2})$, this is somewhat unsatisfactory.

Suppose, however, that Monte Carlo estimates for $r'(\theta)$ are available. The Robbins-Monro (RM) algorithm [7] is the direct stochastic analog of (3.1), namely

$$\theta_{k+1} = \theta_k - a_k Y_k(\theta_k) \quad (3.2)$$

where $E(Y_k(\theta_k) | \theta_1, \dots, \theta_k) = r'(\theta_k)$. It turns out that (3.2) enjoys a convergence rate of $O(n^{-1/2})$ (see p. 381 of [9]); furthermore, the improved convergence rate of RM is manifested empirically. This discussion suggests that one should focus attention on optimization algorithms that iterate on the basis of Monte Carlo sampling of the derivative (as in RM), as opposed to the function value itself (as in KW).

To accomplish this goal, one therefore requires Monte Carlo estimates of $r'(\theta)$. Recall that $r(\theta) = E_\theta Z(\theta)$. If the expectation did not depend on θ (i.e. if $r(\theta) = EW(\theta)$), then (assuming the interchange of derivative and expectation was valid) it would follow that $r'(\theta) = EW'(\theta)$. Thus, by sampling $W'(\theta)$, one would have a Monte Carlo estimate for $r'(\theta)$.

To eliminate dependence of the expectation on θ , two basic approaches are available. Both are most easily illustrated when $\Omega = \mathbb{R}^1$, in which case $P_\theta(\cdot)$ is determined by a distribution function $F(\theta, \cdot)$. The first idea involves observing that

$$\begin{aligned} r(\theta) &= \int_{-\infty}^{\infty} Z(\theta, x) F(\theta, dx) \\ &= \int_0^1 Z(\theta, F^{-1}(\theta, x)) dx \\ &= EZ(\theta, F^{-1}\theta, U) \end{aligned}$$

where $F^{-1}(\theta, \cdot)$ is the inverse distribution function defined by $F^{-1}(\theta, x) = \sup\{y: F(\theta, y) \leq x\}$, and U is a uniform $(0, 1)$ r.v. Thus, if $Z(\cdot, x)$ and $F^{-1}(\cdot, x)$ are differentiable, it follows that $r(\theta) = EW(\theta)$ where $W'(\theta)$ exists. This idea of representing the response in terms of a single uniform r.v. is basically the method of common random numbers (see [8]). It turns out that this approach applies also to discrete-event systems; the estimator $W'(\theta)$ developed in the discrete-event context is the perturbation analysis estimator due to Ho and Suri (see, for example, [10]).

For the second approach, assume that $F(\theta, \cdot)$ has a density $f(\theta, \cdot)$. Let $g(\cdot)$ be the density of any r.v. X for which the density is positive everywhere

(e.g. $X \sim N(0,1)$). Note that

$$\begin{aligned} r(\theta) &= \int_{-\infty}^{\infty} Z(\theta, x) f(\theta, x) dx \\ &= \int_{-\infty}^{\infty} Z(\theta, x) \frac{f(\theta, x)}{g(x)} g(x) dx \\ &= E(Z(\theta, X) f(\theta, X) / g(X)) \end{aligned} \quad (3.3)$$

If $Z(\cdot, x)$ and $f(\cdot, x)$ are differentiable, then clearly $r(\theta)$ can be represented as $E W(\theta)$ where $W'(\theta)$ exists. The idea described above is an application of importance sampling [8] to the derivative estimation problem. For those familiar with statistics, one may interpret the function $f(\theta, X)/g(X)$ as a likelihood ratio; as a consequence, we refer to Monte Carlo gradient estimators based on (3.3) as likelihood ratio gradient estimators.

4. LIKELIHOOD RATIO GRADIENT ESTIMATORS

We now wish to describe the likelihood ratio gradient estimation algorithm, as applied in the discrete-time countable state Markov chain context. Assume the state space S of the chain is a subset of Z^+ . Here, we let the sample space $\Omega = \{\omega = (\omega_0, \omega_1, \dots) : \omega_i \in S\}$. Analogous to the situation in Example 2.1, a probability measure P_θ is induced on Ω by specifying an initial distribution $\mu(\theta) = (\mu(\theta, i) : i \in S)$ and a transition matrix $P(\theta) = (P(\theta, i, j) : i, j \in S)$. The probability P_θ is then defined by

$$P_\theta\{X_0 = i_0, \dots, X_n = i_n\} = \mu(\theta, i_0) \cdot \prod_{k=0}^{n-1} P(\theta, i_k, i_{k+1}) \quad (4.1)$$

where $X_i(\omega)$ is the coordinate r.v. given by $X_i(\omega) = \omega_i$. We assume throughout our discussion that $\mu(\theta)$ and $P(\theta)$ are (element-by-element) continuously differentiable in θ in a neighborhood of θ_0 , and that $\Lambda(\theta) = \{(i, j) : P(i, j, \theta) > 0\}$ is independent of θ in a neighborhood of θ_0 .

The "cost" r.v. $Z(\theta, \omega)$ will initially be assumed to be of the form $Z(\theta, \omega) = g(\theta, X_0(\omega), \dots, X_T(\omega), T(\omega))$ where T is a stopping time for X (this basically means that the event $\{T = n\}$ is determined by the history of X up to time n , for all $n > 0$; see [2], p. 95) and $g(\cdot, i_0, \dots, i_n, n)$ is continuously differentiable.

(4.2) EXAMPLE. If $Z(\theta, \omega) = bX_n(\omega)$, the cost is proportional to the state of the chain at time n .

(4.3) EXAMPLE. If $Z(\theta, \omega) = bT(A, \omega)$ where $T(A) = \inf\{n > 0 : X_n \in A\}$, the cost is proportional to the first hitting time of A . This is useful when one wants to determine a θ under which the chain will enter A as quickly as possible. In particular, if A corresponds to states of a manufacturing system in which the system is operating efficiently, this may be of interest from a real-time control standpoint.

Our goal is to obtain Monte Carlo estimators for $r'(\theta_0)$, where $r(\theta) = E_\theta Z(\theta)$. Now, it is easily seen that

$$E_\theta Z(\theta) = E_\theta Z(\theta) L(\theta) \quad (4.4)$$

where $L(\theta) = \mu(\theta, X_0) \cdot \prod_{k=0}^{T-1} L(\theta, X_k, X_{k+1})$ and $L(\theta, i, j) = P(\theta, i, j) / P(\theta_0, i, j)$. (Decompose both sides of (4.4) over the possible sample outcomes of (X_1, \dots, X_T) and use (4.1).) Thus, $r(\theta) = E_\theta W(\theta)$ where

$$W'(\theta_0) = Z'(\theta_0) + Z(\theta_0) L'(\theta_0)$$

and $L'(\theta_0) = \sum_{k=0}^{T-1} L'(\theta_0, X_k, X_{k+1})$ (observe that $L(\theta_0, i, j) = 1$). We now illustrate the method with an example.

(4.5) EXAMPLE. Suppose that we wish to minimize the expected time to empty an $M/M/1/\infty$ queue given that the queue initially contains m customers. The queue is charged $b\theta$ dollars per unit time to run the server at rate θ . This problem can be posed in terms of the associated embedded discrete-time Markov chain $X = \{X_n : n \geq 0\}$. In particular, let $\mu(\theta, i) = \delta_{im}$,

$$P(\theta, i, j) = \begin{cases} \lambda/(\lambda + \theta), & j = i + 1, i > 1 \\ \theta/(\lambda + \theta), & j = i - 1, i > 1 \\ 1, & i = 0, j = 1 \\ 0, & \text{else,} \end{cases}$$

$$Z(\theta) = T(1 + b\theta)/(\lambda + \theta)$$

where $T = \inf\{n > 0 : X_n = 0\}$. To generate $W_1^i(\theta_0)$:

- 1) Simulate the chain X up to time T , under initial distribution $\mu(\theta_0)$ and transition matrix $P(\theta_0)$.
- 2) Calculate

$$W_1^i(\theta_0) = T \frac{(b\lambda - 1)}{(\lambda + \theta_0)^2} + \frac{T(1 + b\theta_0)}{(\lambda + \theta_0)} \cdot \sum_{k=0}^{T-1} L'(\theta_0, X_k, X_{k+1})$$

where $L'(\theta_0, i, i + 1) = -(\lambda + \theta_0)^{-1}$ and $L'(\theta_0, i, i - 1) = \lambda/\theta_0(\lambda + \theta_0)$.

The sampling technique described above could now be used at each step of the RM algorithm given by (3.2), in order to minimize $r(\theta)$.

In many application settings, the goal is to minimize $r(\theta)$, where $Z(\theta)$ is a steady-state limit (see Example 2.1). In general, the steady-state depends on the infinite history of the process and therefore $Z(\theta)$ cannot be chosen to depend on X only up to the stopping time T (in other words, $Z(\theta)$ cannot be represented as $g(\theta, X_0, \dots, X_T, T)$). However, in the current context, we are in luck since the regenerative structure of Markov chains allows one to reduce the analysis from the infinite horizon to that of a single regenerative cycle. Specifically, suppose

$$Z(\theta) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} f(\theta, X_k)$$

Regenerative process theory [3] shows that if X is an irreducible positive recurrent Markov chain,

$$Z(\theta) = u(\theta)/\ell(\theta) \text{ a.s.}$$

where $u(\theta) = E_{\theta} \left(\sum_{k=0}^{T-1} f(\theta, X_k) \right)$ and $l(\theta) = E_{\theta} T$.

($P_{\theta}(\cdot)$ is the probability on Ω under which $\mu(\theta, i) = \delta_{0i}$ and the transition matrix $P(\theta)$ is used; $T = \inf\{n > 1: X_n = 0\}$ is then a regenerative stopping time for X). Hence, $r(\theta) = u(\theta)/l(\theta)$ and $r'(\theta) = [u'(\theta)l(\theta) - l'(\theta)u(\theta)] \cdot l(\theta)^{-2}$. The quantities $l(\theta)$ and $u(\theta)$ can be estimated using standard Monte Carlo methods. The estimation of $u'(\theta)$ and $l'(\theta)$ is easily accomplished using the analysis discussed above, since $u(\theta)$ and $l(\theta)$ are the expectations of r.v.'s which depend on X only up to a stopping time T . Furthermore, the estimator obtained in this way is suitable for evaluation at each step of a RM algorithm for minimizing the steady-state cost $r(\theta)$ over θ .

Further work (in particular, extensions to more general processes) may be found in [5].

5. A HOMOTOPY METHOD FOR STOCHASTIC OPTIMIZATION

In this section, we shall briefly describe a new Monte Carlo algorithm which holds considerable promise as a technique for optimizing stochastic systems in both the discrete and continuous context. The idea focuses on the fact that in virtually all stochastic optimization problems, the system contains both controllable and uncontrollable parameters.

(2.1) EXAMPLE (continued). In the $M/M/1/\infty$ queue, the system depends on both λ (uncontrollable parameter), and μ (the controllable decision parameter).

Returning to our formulation of Section 2, this suggests that the probability measure P on Ω depends both on the uncontrollable decision variable θ and a vector of uncontrollable parameters, say $\alpha \in R^m$, so that $P = P(\theta, \alpha)$. Our goal is to find θ^* to minimize $r(\theta, \alpha_0)$ over $\theta \in K$ (α_0 is a given prescribed setting for the uncontrollable factors), where

$$r(\theta, \alpha) = \int_{\Omega} Z(\theta, \omega) P(\theta, \alpha, d\omega).$$

Our basic idea is that there may be a region Γ in the α -space R^m in which the optimization problem is particularly easy to solve (either analytically, numerically, or via Monte Carlo). Specifically, assume that one can find a "good" solution $\hat{\theta}$ at the point $\alpha_1 \in \Gamma$. For $\lambda \in [0, 1]$, let $\alpha(\lambda) = (1 - \lambda)\alpha_1 + \lambda\alpha_0$. One expects that the minimizer $\theta^*(\lambda)$ for the problem $r(\theta, \alpha(\lambda))$ will deform continuously from $\theta^*(0)$ (i.e. the known point $\hat{\theta}$) into the desired minimizer $\theta^*(1)$, as λ is increased from 0 to 1; the parameter λ is called the homotopy parameter. The minimizer $\theta^* = \theta^*(1)$ is, of course, the minimizer of the original problem.

(5.1) EXAMPLE. Consider an asynchronous automatic assembly system, as indicated in Figure 1. Pallets containing partially completed assemblies circulate in a clockwise fashion on the machine. Completed assemblies are offloaded at station n ; at the next station (#1), a new assembly base is loaded onto the fixture on the circulating pallets and the assembly process begins anew. We wish to optimize the expected number of completed assemblies per hour, as a function of the number of buffer units b_i in front of station i and the number of pallets p on the machine (so $\theta = (b_1, \dots, b_r, p)$). Although the machine cycle and

transport times are assumed deterministic here, a stochastic effect due to jamming makes this problem hard to solve.




Fig. 1 An asynchronous automatic assembly system

The idea is to observe that this problem is greatly simplified when the jam rate parameter α is zero. For example, if the machine cycle times are all identically equal to c , the machine is then synchronous. Thus, the optimal throughput $1/c$ is obtained by setting $b_i = 0$ and $p = n$. The homotopy algorithm works by increasing the jam rate slowly from zero to the "true" jam rate $\alpha_0 > 0$; the discreteness of the problem is advantageous here since for relatively small increments in λ , the optimizer can move only to neighboring values in the space of decision variables θ .

This type of approach has met with considerable success in the deterministic mathematical programming setting (see [1], for example). The primary difficulty in applying the method involves following the path $\theta^*(\lambda)$ as λ ranges from 0 to 1. Appropriate path-following algorithms in the Monte Carlo optimization context described here are currently under investigation.

6. CONCLUSION

Two new approaches to Monte Carlo optimization of stochastic systems have been presented here. The first is based on utilization of likelihood ratio gradient estimators in RM procedures, whereas the second uses homotopy methods to follow an "optimal path" in decision variable space. Both methods are under further investigation, at both a theoretical and empirical level.

7. ACKNOWLEDGEMENTS

The research of the first author was supported by NSF Grant ECS-840-4809 and sponsored by the United States Army under Contract No. DAAG29-80-C-0041.

REFERENCES

- [1] Allgower, E. L. and Georg, K., "Predictor-Corrector and Simplicial Methods for Approximating Fixed Points and Zero Points of Nonlinear Mappings," Mathematical Programming - The State of the Art - Bonn 1982, Springer-Verlag, 1983.
- [2] Breiman, L., "Probability," Addison-Wesley, Reading, Mass., 1968.
- [3] Cinlar, E., "Introduction to Stochastic Processes," Prentice-Hall, Englewood Cliffs, NJ, 1975.
- [4] Ethier, S. N. and Kurtz, T. G., "Markov Processes: Characterization and Convergence," Wiley, New York, 1986.
- [5] Glynn, P. W., "Likelihood Ratio Gradient Estimators for Discrete-Event Systems," forthcoming technical report, 1986.
- [6] Kiefer, J. and Wolfowitz, J., "Stochastic Estimation of the Maximum of a Regression Function," Ann. Math. Statist., Vol. 23, 1952, pp. 462-466.
- [7] Robbins, H. and Monro, S., "A Stochastic Approximation Method," Ann. Math. Statist., Vol. 22, 1951, pp. 400-107.
- [8] Rubinstein, R. Y., "Simulation and the Monte Carlo Methods," Wiley, New York, 1981.
- [9] Sacks, J., "Asymptotic Distribution of Stochastic Approximation Procedures," Ann. Math. Statist., Vol. 29, 1958, pp. 373-405.
- [10] Suri, R., "Implementation of Sensitivity Calculations on a Monte Carlo Experiment," J. Optim. Theory Applic., Vol. 40, 1983, p. 4.