Proceedings of the 1994 Winter Simulation Conference
ed. J. D. Tew, S. Manivannan, D. A. Sadowski, and A. F. Seila

SIMULATION OPTIMIZATION VIA SIMULTANEOUS PERTURBATION STOCHASTIC APPROXIMATION

Stacy D. Hill

Johns Hopkins University Applied Physics Laboratory
Laurel, Maryland 20723

Michael C. Fu

University of Maryland at College Park
College Park, Maryland 20742

ABSTRACT

Stochastic approximation is a simulation optimization technique that has received much attention recently. Traditional finite difference-based stochastic approximation schemes require a large number of simulations when the number of parameters of interest is large. We apply simultaneous perturbation stochastic approximation (SPSA), which requires only two simulations per gradient estimate, regardless of the number of parameters of interest. We report simulation experiments conducted on a single-server queue, comparing the algorithm with finite differences and with perturbation analysis (PA). We then consider a transportation problem and formulate it as a stochastic optimization problem to which we propose to apply SPSA.

1 INTRODUCTION

Consider the problem of optimizing some performance measure of a discrete event system (see, e.g., Cassandras 1993 or Fu 1994). Under suitable conditions, optimization requires finding the zero of the performance measure gradient. Techniques such as perturbation analysis or likelihood ratio provide an efficient means of computing the performance measure gradient from a single sample path of the system. Such techniques, however, require detailed knowledge of the system dynamics and model. For example, one must usually assume known the form of the input distributions. For some important discrete event systems that arise in practice — for example, complex transportation networks — such information is often unknown, in which case direct estimation of the gradient is not possible.

This paper considers a stochastic approximation technique for optimizing DEDS under minimal assumptions on the system of interest. To be more specific, let $\theta \in \Theta \subset \mathcal{R}^p$ denote a vector of controllable (adjustable) parameters and ω the stochastic effects. Let $L(\theta, \omega)$ denote the sample path performance of interest and $J(\theta) = E[L(\theta, \omega)]$ expected system performance. The problem is to find $\arg\min\{J(\theta) : \theta \in \Theta\}$. The stochastic approximation algorithm for solving $\nabla J = 0$ is given by the following iterative scheme:

$$\theta_{(n+1)} = \theta_{(n)} - a_n \widehat{g}_n, \qquad (1)$$

where $\widehat{g}_n$ represents an estimate of the gradient ∇J at θ_n, and $\{a_n\}$ is a positive sequence of numbers converging to 0.

Direct gradient estimates are the most efficient. However, in many cases this may not be feasible, in which case gradient estimates based on (noisy) measurements of the performance measure itself are the only recourse. The purpose of our work is to see if simultaneous perturbation stochastic approximation (SPSA) can optimize discrete event systems at significant reductions in computations over the standard approach based on finite differences (FD).

SPSA uses the simultaneous perturbation (SP) method to estimate the gradient (Spall 1992, Spall and Cristion 1994) and stochastic approximation (SA) to find the zero of the gradient. The SP method does not require detailed knowledge of system dynamics and input distributions. Rather, it only requires measurements (which may contain measurement noise) of the performance measure. In fact, in each stochastic approximation update step SPSA requires only two measurements of the performance measure to calculate a gradient estimate, regardless of p (the dimension of the vector of parameters). This contrasts sharply with the method of finite differences for estimating gradients, where the number of required function evaluations is at least p. Thus, SP requires substantially less data — in our application, meaning significantly fewer simulations — than finite differences for estimating gradients in high dimensions. In real systems, where data acquisition can often be expensive and time-consuming, this reduction in computation translates into a cost savings.

The SP technique has been applied to nonlinear control problems using neural networks (Spall and Cristion 1994). Here, we illustrate its application to discrete-event systems, by considering a single-server queueing problem and a transportation problem.

2 SIMULTANEOUS PERTURBATIONS

Let e_i denote the unit vector in the ith direction, and t_n the simulation "duration" (e.g., number of customer completions in a queueing simulation) of the nth iteration. Let $\{\Delta_1, ..., \Delta_p\}$ be a set of i.i.d. perturbations satisfying the conditions given in Spall (1992), and define the vector $\Delta = [\Delta_1 ... \Delta_p]$. We took the Δ_i's to be symmetric Bernoulli in all of our simulation experiments. Let $(\widehat{g}_n)_i$ denote the ith component of $\widehat{g}_n$. Then, the simultaneous perturbation (SP) estimator is given by

$$(\widehat{g}_n)_i = \frac{J^+_{t_n}(\theta_{(n)} + c_n\Delta, \omega_n^+) - J^-_{t_n}(\theta_{(n)} - c_n\Delta, \omega_n^-)}{2c_n\Delta_i}. \tag{2}$$

Compare these estimators with symmetric difference (SD) estimators. Symmetric differences require a different pair of estimates in the numerator for each parameter, thus requiring $2p$ simulations, whereas in SPSA, the *same* pair is used in the numerator for all parameters, and instead the denominator changes; thus, only two discrete-event simulations are required at each iteration.

In the stochastic approximation implementation given by (1), both SP and SD estimators also require a positive sequence $\{c_n\}$ converging to 0 at an appropriate rate. In our experiments, we took $a_n = a/n$ (a to be selected), $c_n = c/n^{0.25}$ (c to be selected). For the single-server queue example, we considered four different gradient estimates: SP, FD, SD, and IPA (infinitesimal perturbation analysis).

3 SINGLE-SERVER QUEUE EXAMPLE

Consider a single-server queue with Poisson arrivals and service times from a uniform distribution (an M/U/1 queue). The goal is to minimize a customer's mean steady-state system time T, under penalty costs on the service time. Specifically, we wish to determine the values of the two parameters $\theta = (\theta_1, \theta_2)$ in the uniform service time distribution $U(\theta_1 - \theta_2, \theta_1 + \theta_2)$ to minimize the objective function

$$J(\theta) = E[T] - c_1\theta_1 - c_2\theta_2, \quad \theta \in \Theta, \tag{3}$$

where c_1 and c_2 are costs on reducing the service time mean and "variability," respectively, $\Theta = \{(\theta_1, \theta_2) : 0 < \theta_2 \leq \theta_1 < 1/\lambda\}$, and λ is the arrival rate. This example was considered in Fu and Ho (1988).

We compare the SP estimator with finite difference estimators (both one-sided and symmetric) and with perturbation analysis (PA) estimators, which require only a *single* simulation per estimate. We present simulation results for implementation in the SA algorithm given by (1) to minimize the objective function (3), with the estimate of the gradient given by

$$\widehat{g} = \left(\frac{\partial E[T]}{\partial \theta}\right)_{est} - c_1 e_1 - c_2 e_2. \tag{4}$$

We considered six cases. Table 1 gives the values of c_1 and c_2, the resulting optimal values and the corresponding values of the objective function and the two partial second derivatives, from which the value of a is determined, as described in the next paragraph. As noted in Fu and Ho (1988), for some values of c_1 and c_2, the theoretical optimal solution could lie arbitrarily close to the boundary of the constraint set. A minimum at which $dJ/d\theta = 0$ exists if $c_1 > 6c_2^2 + 3c_2 + 1$. Our cases were selected such that this condition held, for which the minimum occurs at

$$\theta^* = \left(1 - \frac{1}{\sqrt{C}}, \frac{3c_2}{\sqrt{C}}\right)\frac{1}{\lambda}, \quad C = 2c_1 - 3c_2^2 - 1. \tag{5}$$

Further implementation values are as follows: $\lambda = 1$; $c = 0.001$; starting point: $\theta_1 = 0.5, \theta_2 = 0.3$; observation length per iteration: 100 customers; 40 independent replications; number of iterations per replication: 1000 (total budget of 100,000 customers/replication); value of a: geometric mean of second derivatives (approximated to one significant figure). In general, of course, the parameter a could not be calculated "optimally" in advance, since the objective function is unknown.

Since analytical results are available, performance of the algorithms was simply measured by the performance measure $J(\theta_n)$, lower being better. The results are summarized in Table 2. The headings SDSA, FDSA, and PASA refer to stochastic approximation algorithms based on symmetric differences, one-sided (forward) finite differences, and perturbation analysis, respectively; J^* refers to the (true) minimum. The table gives the mean of the estimated minimum $\pm$ standard error, based on 40 independent replications, for the algorithm after 500 customers simulated and after 1000 customers simulated. In terms of simulation budget, the $n = 1000$ case of SPSA corresponds approximately to the $n = 500$ case of SDSA, since each iteration of SDSA requires twice ($p = 2$) as many simulations as each iteration of SPSA. In this limited set of cases, SPSA performs comparably to

Table 1: Optimization Cases for $M/U/1$ Queue ($\lambda = 1$)

Case	c_1	c_2	optimal θ_1^*	optimal θ_2^*	J^*	$\partial^2 J^*/\partial\theta_1^2$	$\partial^2 J^*/\partial\theta_2^2$	a
1	1.28125	0.00125	0.2	0.003	-0.03125	1.953	0.4167	1.0
2	1.28969	0.075	0.2	0.180	-0.03969	1.974	0.4167	1.0
3	2.5	0.002	0.5	0.003	-0.5000	8.000	0.6667	0.4
4	2.6536	0.32	0.5	0.480	-0.6536	8.614	0.6667	0.4
5	13.0	0.005	0.8	0.003	-8.000	125	1.6667	0.1
6	15.535	1.3	0.8	0.780	-10.535	150	1.6667	0.1

Table 2: SA Results: J as a function of the number of iterations

Case	SPSA		SDSA		FDSA		PASA	
J^*	500	1000	500	1000	500	1000	500	1000
1	-0.029890	-0.029391	-0.030901	-0.031060	-0.030900	-0.031058	-0.030900	-0.031058
-0.03125	± 0.000340	± 0.000353	± 0.000240	± 0.000139	± 0.000240	± 0.000140	± 0.000238	± 0.000139
2	-0.039420	-0.039648	-0.039018	-0.039316	-0.039112	-0.039414	-0.039194	-0.039451
-0.03969	± 0.000119	± 0.000031	± 0.001873	± 0.001171	± 0.001379	± 0.000734	± 0.001306	± 0.000707
3	-0.490221	-0.490450	-0.498311	-0.498733	-0.498367	-0.498773	-0.498296	-0.498720
-0.5000	± 0.005581	± 0.004888	± 0.001458	± 0.001091	± 0.001251	± 0.000929	± 0.001482	± 0.001111
4	-0.652158	-0.652730	-0.652633	-0.652782	-0.652394	-0.652603	-0.652387	-0.652607
-0.6536	± 0.001316	± 0.000787	± 0.001122	± 0.000781	± 0.001831	± 0.001367	± 0.001876	± 0.001368
5	-7.840566	-7.823904	-7.911220	-7.903313	-7.911106	-7.903152	-7.910641	-7.902878
-8.000	± 0.100562	± 0.107669	± 0.045980	± 0.040336	± 0.046555	± 0.040226	± 0.046049	± 0.040332
6	-10.345714	-10.328994	-10.380064	-10.359260	-10.379407	-10.358351	-10.378743	-10.358410
-10.535	± 0.089023	± 0.084015	± 0.083043	± 0.077733	± 0.083706	± 0.078061	± 0.083636	± 0.077930

the other techniques. Compared to the finite differences, it does so with half as many computations.

4 TRANSPORTATION APPLICATION

We consider the application of SPSA to a transportation problem, in particular to transfer optimization in a transit network, where the goal is to reduce waiting time during transfers. We formulate this problem as a DEDS optimization problem to which we propose to apply SPSA.

The model we consider is a transit network with bus lines traveling in four directions on a grid: east, west, north, and south. Transfers occur, for instance, from a west-bound line to a north-bound line. Multiple transfers are possible. Undesirable delays occur for passengers due to waiting for a transfer. As summarized in Bookbinder and Désilets (1992), there are two basic approaches to this problem: timed transfer and transfer optimization. The former focuses on coordinating the transfer points, and is more applicable for networks where transfers constitute a relatively smaller proportion of overall traffic, e.g., intercity trains and planes. This approach would not be appropriate for a large transit network, such as is found in a downtown bus network, where transfers are decentralized. In this case, transfer optimization is usually employed, whereby the decisions to be made have to do with the departure times of the first bus on a line.

In transfer optimization, the following are usually assumed to be given: the network, i.e., no re-routing is allowed; the headways, defined as the times between adjacent buses on the same line (assumed to be constant and equal); the transfer points; the passenger traffic and transfers. Traffic on a route can be given either as a point-to-point total or equivalently as a Markovian routing matrix at each stop. Stochastic elements of the network — incorporated indirecly in Bookbinder and Désilets (1992) — include the arrival process of passengers at each stop, both timing and number; and the travel times of buses.

Let N be the number of transit lines, M be the number of transfer points, H_i be the headway for transit line i, $i = 1, \ldots, N$, Θ_i be the set of allowable offset times for transit line i, $i = 1, \ldots, N$, $\theta = \{\theta_i\}_{i=1}^N$ be the timetable for the transit network, $\Theta = \{\Theta_i\}_{i=1}^N$ be the allowable timetables for the transit network. Note that a *transfer point* in the network model is quite different from a *stop* in the physical real world. In particular, if a given "single" stop occurs at an intersection of two bi-directional routes and allows all possible transfers, then this would generate eight separate transfer points in the network model.

We wish to minimize the total expected waiting time for transfers in the network. This problem is usually formulated as a mathematical program, which requires the assumption that the sets $\Theta_i, i = 1, ..., N$ be **discrete and finite**, yielding the integer program:

$$\min_{\theta \in \Theta} \sum_{i=1}^{N} \sum_{j=1}^{N} C_{ij\theta_i\theta_j},$$

where n_k is the transfer flow at transfer connection k, $k = 1, ..., M$, $C_{ijrs} = \sum_{k \in A_{ij}} n_k W_k(r, s)$, $r \in \Theta_i, s \in \Theta_j$, $A_{ij} = \{k$: connection k goes from line i to line $j\}$, $W_k(r, s)$ is the mean waiting time at connection k, for offset times $r \in \Theta_i$ and $s \in \Theta_j$, for lines i and j, respectively. The key elements are the waiting times, which must somehow be estimated. This problem is equivalent to the 0-1 quadratic optimization problem:

$$\min \sum_{i=1}^{N} \sum_{j=1}^{N} \sum_{r \in \Theta_i} \sum_{s \in \Theta_j} C_{ijrs} x_{ir} x_{js},$$

subject to $\sum_{r \in \Theta_i} x_{ir} = 1, \; i = 1, ..., N, \quad x_{ir} \in \{0, 1\}.$

The 0-1 variables x_{is} take the value 1 if and only if offset time s is chosen for link i, and the equality constraints insure that exactly one of the allowable offset times is chosen for each line. This formulation is equivalent to the well-known quadratic assignment problem (QAP) in facilities layout planning, and hence is NP-complete.

A more realistic model should probably include the following features: the feasible set of offset times is continuous; headways need not be constant nor deterministic, e.g., they could be closer during rush hours; travel times need not be constant nor deterministic, i.e., they are likely to be random and higher during rush hours; the passenger arrival process need not be deterministic. However, incorporating such factors into a model leads to analytical intractability in determining the mean waiting times, in which case the best approach is a stochastic discrete-event simulation model.

Taking Θ_i to correspond to intervals $[0, K_i]$, where K_i is the maximum allowable offset time on transit line i, and assuming that the optimum (at least local) is found by $\nabla_\theta g(\theta) = 0$, we propose to apply SPSA to

$$g(\theta) = \sum_{i=1}^{N} \sum_{j=1}^{N} C_{ij\theta_i\theta_j}.$$

REFERENCES

Bookbinder, J.H. and Désilets. A., Transfer optimization in a transit network, *Transportation Science*, Vol. 26, No. 2, May 1992, 106-118.

Cassandras, C.G. 1993. *Discrete Event Systems: Modeling and Performance Analysis*, Irwin.

Fu, M.C.. 1994. Optimization via simulation: a review, to appear in *Annals of Operations Research*.

Fu, M.C. and Y. C. Ho. 1988. Using perturbation analysis for gradient estimation, averaging, and updating in a stochastic approximation algorithm, *Proc. Winter Simulation Conference*:509-517.

Hill, S.D. and M. C. Fu. 1994. Optimizing Discrete Event Systems with the Simultaneous Perturbation Stochastic Approximation Algorithm, *Proc. 33rd IEEE Conf. Dec. Contr.* to appear.

Spall, J.C. 1992. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation, *IEEE Trans. Aut. Contr.* **37**:332-341.

Spall, J.C. and J.A. Cristion. 1994. Nonlinear Adaptive Control Using Neural Networks: Estimation with a Smoothed Simultaneous Perturbation Gradient Approximation, *Statistica Sinica* **4**:1-27.

AUTHORS BIOGRAPHIES

STACY D. HILL is a Systems Engineer on the Senior Professional Staff of the Johns Hopkins University Applied Physics Laboratory. He received the D.Sc. degree in control systems engineering and applied mathematics from Washington University in 1983. He has been project and technical lead in developing, testing, and applying statistical estimation techniques and software, and has authored and co-authored journal articles on estimation theory and techniques, including an article in the cross-disciplinary book *Bayesian Analysis of Time Series and Dynamic Models*. His present research interests include optimization and simulation of discrete event systems, with application to transportation problems. Dr. Hill is a member of **IEEE**, **ASA**, and **Sigma Xi**.

MICHAEL C. FU is an Associate Professor of Management Science & Statistics in the College of Business and Management, with a joint appointment in the Institute for Systems Research, at the University of Maryland at College Park. He received an S.B. degree 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, inventory control, and the pricing of financial derivatives. He teaches courses in applied probability, stochastic processes, discrete-event simulation, and operations management. He is a member of **IEEE** and **ORSA**.