Proceedings of the 2001 Winter Simulation Conference
B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, eds.

CONSTRAINED OPTIMIZATION OVER DISCRETE SETS VIA SPSA WITH APPLICATION TO NON-SEPARABLE RESOURCE ALLOCATION

James E. Whitney, II
Latasha I. Solomon

Department of Electrical and Computer Engineering
Morgan State University
5200 Perring Parkway
Baltimore, MD 21251, U.S.A.

Stacy D. Hill

Johns Hopkins University
Applied Physics Laboratory
11100 Johns Hopkins Rd.
Laurel, MD 20723, U.S.A.

ABSTRACT

This paper presents a version of the Simultaneous Perturbation Stochastic Approximation (SPSA) algorithm for optimizing non-separable functions over discrete sets under given constraints. The primary motivation for discrete SPSA is to solve a class of resource allocation problems wherein the goal is to distribute a finite number of discrete resources to finitely many users in such a way as to optimize a specified objective function. The basic algorithm and the application of the algorithm to the optimal resource allocation problem is discussed and simulation results are presented which illustrate its performance.

1 INTRODUCTION

In many practical non-linear, high-dimensional optimization problems, the relationship between the problem parameters will be sufficiently complex such that the objective function is not easily described in closed, analytic form. It is then necessary to use some iterative, numeric technique to find the objective function optimizer. Additionally, in the constrained discrete resource allocation problem, the number of resources available may differ from the solution which optimizes the objective function. This paper considers the application of a modified version of *simultaneous perturbation stochastic approximation* (SPSA) for such difficult optimization problems. The SPSA method has been found to be a fast and computationally efficient technique for optimizing large dimensional problems when the objective function is sufficiently smooth yet analytically unavailable or difficult to obtain in closed form. The method is essentially a Kiefer-Wolfowitz stochastic approximation scheme that relies on a very efficient approximation of the objective function gradient.

2 BASIC SPSA ALGORITHM

The focus of the present paper is the discrete version of SPSA. Before discussing the procedure, a brief review of the basic SPSA algorithm (Spall 1992) for continuous parameter optimization will be given. A modified version of the algorithm for constrained optimization will be discussed in Sections 3 and 4.

Consider the problem of minimizing a real-valued function, $L(\theta)$, defined on an open domain D in p-dimensional Euclidean space, R^p. The function $L(\theta)$ is assumed to be at least three-times differentiable and have a unique minimizer in D. Although the exact form of L is not assumed known, it is assumed, minimally, that noisy measurements of the function are available. In particular, assume that measurements $M(n,\theta)$ are available, where

$$M(n,\theta) = L(\theta) + \varepsilon_n(\theta) \quad (1)$$

and ε_n is a zero-mean measurement noise process. The SPSA algorithm uses two measurements of L at iteration k to form a gradient estimate:

$$M_k^+(\theta) = L(\theta + c\Delta_k) + \varepsilon_k(\theta + c\Delta_k) \quad (2)$$
$$M_k^-(\theta) = L(\theta - c\Delta_k) + \varepsilon_k(\theta - c\Delta_k). \quad (3)$$

The SPSA gradient estimate of the i-th component of the gradient at the k-th iteration is

$$H_i(k,\theta) = \frac{\left(M_k^+(\theta) - M_k^-(\theta)\right)}{2c\Delta_{ki}}. \quad (4)$$

The parameter c is a positive-valued step size, and Δ_k is a random perturbation vector; $\Delta_k = \left(\Delta_{k1}, \ldots, \Delta_{kp}\right)^T$. The

Δ_{ki}'s form an i.i.d. sequence of ±1 Bernoulli random variables and are assumed to be independent of the measurement noise process.

3 PENALTY FUNCTION

The algorithm presented in Section 2 is for unconstrained optimization. This section will discuss a modification of the algorithm for use on constrained optimization problems. The modification utilized here relies on the method of penalty functions. A penalty function of the following form is considered

$$P(\theta) = \sum_{j=1}^{s} w_j p(q_j(\theta)) \tag{5}$$

where, w_j are the positive scalar weights and, $p(\bullet)$ is a real-valued function on R. The penalty function is implemented to transform the original constrained optimization problem into an unconstrained optimization problem given by

$$\underset{\theta}{\text{opt}}\, L(\theta) + rP(\theta) \tag{6}$$

where, r is the penalty parameter and is a positive real number. r is typically chosen larger than some threshold $\bar{r}$, which depends on both $L(\theta)$ and $P(\theta)$. Because precise information on $L(\bullet)$ is unavailable, the optimum value of r is also unknown. Therefore, r is allowed to slowly increase, forming a sequence of unconstrained optimization problems defined by

$$\underset{\theta \in R^d}{\text{opt}}\, L(\theta) + r_n P(\theta)\,. \tag{7}$$

The penalized cost function is then given by

$$L_n(\theta) = L(\theta) + r_n(\theta)\,. \tag{8}$$

(Wang and Spall 1999) give the necessary conditions such that $\theta_n^* = \theta^*$. The modified parameter update equation is then given by

$$\theta_{n+1} = \theta_n - a_n \hat{g}_n(\theta_n) - a_n r_n \nabla P(\theta_n) \tag{9}$$

where, $\hat{g}_n(\theta_n)$ is an estimate of the gradient of $L(\bullet)$ at θ_n, $\{r_n\}$is an increasing sequence of positive scalars with $\lim_{n\to\infty} r_n = \infty$, $\nabla P(\theta)$ is the gradient of $P(\theta)$ at θ, and $\{a_n\}$ is a positive scalar sequence satisfying $a_n \to 0$ and $\sum_{n=1}^{\infty} a_n = \infty$. The gradient estimate, $\hat{g}_n(\theta_n)$, is obtained as before, from two noisy measurements of the cost function , i.e., Eqs. (2) and (3), and the SPSA gradient estimate of the i-th component of the gradient at the k-th iteration is given by Eq. (4).

4 DISCRETE SPSA

The discrete version of SPSA requires some modification of $H(k+1, \hat{\theta}_k)$ to ensure that the iterates satisfy the constraints placed on θ. Herein, several possible modifications are investigated. To simplify the treatment, it is assumed that the domain of interest is the grid of points in R^p with discrete-valued coordinates.

In Method 1, the estimate of θ is constrained by constraining the gradient estimate $H(k,\theta)$ to discrete values. At each iteration of the algorithm, the discrete solution is formed by

$$\hat{\theta}_{k+1} = \hat{\theta}_k - a \cdot round\left\{H(k+1, \hat{\theta}_k)\right\} \tag{10}$$

Upon convergence, the final values of the parameter vector, $\hat{\theta}$, are given by $\hat{\theta}_{final} = round\{\hat{\theta}_k\}$.

Method 2 is constrained in a manner similar to that of Method 1, with the exception that in this case, the gain constant a is also included in the rounding operation, i.e.,

$$\hat{\theta}_{k+1} = \hat{\theta}_k - round\left\{aH(k+1, \hat{\theta}_k)\right\} \tag{11}$$

Upon convergence, the final values of the parameter vector, $\hat{\theta}$, are already discrete valued.

In the final method, Method 3, the entire parameter estimate is rounded either up or down to the nearest discrete value, i.e.,

$$\hat{\theta}_{k+1} = round\left\{\hat{\theta}_k - aH(k+1, \hat{\theta}_k)\right\}, \tag{12}$$

on each iteration.

5 RESOURCE ALLOCATION

In the resource allocation problem, let there be M user classes and n types of resources allocated over these M user classes. There are N_i resources of type i and each resource of type i is assumed identical. The number of resources of type i that are allocated to user class j is denoted as θ_{ij}. The entire set of θ_{ij} is denoted by Θ. These relationships are illustrated in Table 1.

Table 1: User Class vs. Resource Type

User Class

$$\Theta = \text{ResourceType}\begin{bmatrix} \theta_{11} & \theta_{12} & \theta_{1M} \\ \theta_{21} & & \theta_{2M} \\ & & \\ \theta_{n1} & & \theta_{nM} \end{bmatrix}$$

The authors previously looked at the use of SPSA on resource allocation problems with separable objective functions (Gerencsér, Hill, Vágó 1999, Whitney and Hill 2001), i.e., functions of the form

$$L(\Theta) = \sum_{j=1}^{M} L_j(\theta_j) \qquad (13)$$

where, $L_j(\theta_j)$ is the individual cost incurred by class j, and $\theta_j = \theta_{1j}, \theta_{2j}, \ldots, \theta_{nj}$. In this case, the discrete, multiple constrained resource allocation problem is

$$\min L(\Theta)$$

subject to

$$\sum_{j=1}^{M} \theta_{ij} = N_i, \theta_{ij} \geq 0, 1 \leq i \leq n\,, \qquad (14)$$

where, N_i is the total number of resources of type i, j is the class of resource, θ_{ij}'s are non-negative integers, and M is the number of user classes.

The class of allocation problems considered in the current work involve non-separable functions of the resource allocation parameters, θ_{ij}. The particular objective function may be represented mathematically as follows (Einbu 1984):

$$\max L = \sum_{k=1}^{K} r_k \left(\sum_{j=1}^{J} e_{jk}\theta_{jk} \right), \qquad (15)$$

subject to constraints

$$\sum_{k=1}^{K} \theta_{jk} \leq h_j; j = 1,\ldots,J\,, \qquad (16)$$

$$\theta_{jk} \geq 0;\, j = 1,\ldots,J;\, k = 1,\ldots,K\,, \qquad (17)$$

where, J, is the number of resource types, K, the number of activities, θ_{jk}, the (unknown) quantity of resource j, allocated to activity k, h_j, the available quantity of resource j, r_k, the return function of activity k, and e_{jk}, the effectiveness of resource j when allocated to activity k; it is assumed that $e_{jk} \geq 0$ for at least one k for every j. It is further assumed that r_k is continuously differentiable, strictly concave, and strictly increasing in both the positive and negative domains. The return functions were of the form

$$r_k\left(\sum_{j=1}^{J} e_{jk}\theta_{jk} \right) = a_k \left[1 - \exp\left(- \sum_{j=1}^{J} e_{jk}\theta_{jk} \right) \right], \qquad (18)$$

$k = 1, \ldots, K$. The matrix and vector elements e_{jk}, a_k, and h_j were sampled from uniform distributions over the (continuous) intervals [0.2,1], [5,10], and [1,3] respectively.

In the current problem, $L(\theta)$ is concave, so the objective is to distribute the resources in such a way that $L(\theta)$ is maximized.

6 SIMULATION

Simulation results are presented for a cost function of the form

$$L = \sum_{k=1}^{K} r_k \left(\sum_{j=1}^{J} e_{jk}\theta_{jk} \right). \qquad (19)$$

In this example, there are three resource types, J. Resource type '1', has an available quantity of 6, resource type '2', has an available quantity of 3, resource type '3'has an available quantity of 4, and the number of activities, a, is 3. The problem is then the optimum distribution of these available resources, i.e., θ^*_{jk}, subject to constraints

$$\begin{array}{ccccccc} \theta^*_{11} & + & \theta^*_{12} & + & \theta^*_{13} & \leq 6 \\ \theta^*_{21} & + & \theta^*_{22} & + & \theta^*_{23} & \leq 3\,. \\ \theta^*_{31} & + & \theta^*_{32} & + & \theta^*_{33} & \leq 4 \end{array} \qquad (20)$$

7 RESULTS

The results for the three different methods of parameter constraint are illustrated in Figures 1-3. These figures illustrate the cost function maximization for the three methods as a function of the instantaneous estimate of the parameter vector, θ, at iteration k. The parameter vector, $\hat{\theta}$, was initialized to a zero matrix, and the efficiencies were arbitrarily set to '1'.

Figure 1 illustrates the result when the parameters are constrained using Method 1. In Method 1, the gradient estimate is truncated, but during the iterations, $\hat{\theta}_k$, is allowed to take on non-discrete values. When the specified number of iterations are run, or, the convergence criteria is met, the final $\hat{\theta}_k$ is then truncated.

Figure 2 illustrates the result when the parameters are constrained using Method 2. In Method 2, on each iteration the product of the gain factor, *a*, and the gradient estimate, $g(\theta)$, is rounded.

Figure3 illustrates the result when the parameters are constrained utilizing Method 3; in this method the entire quantity $\left\{\hat{\theta}_k - aH(k+1,\hat{\theta}_k)\right\}$ is truncated.

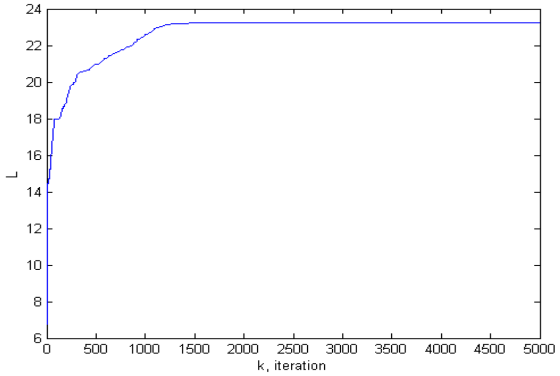

Figure 1: Convergence Rate of $L(\theta)$, Method 1

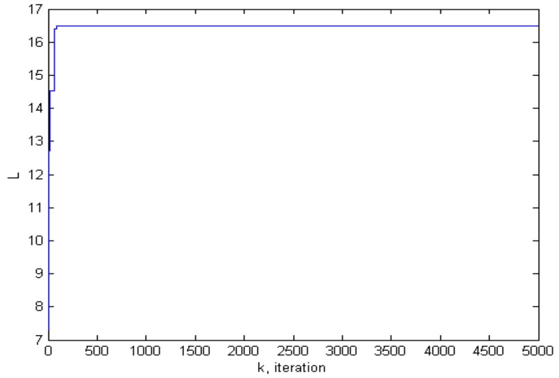

Figure 2: Convergence Rate of $L(\theta)$, Method 2

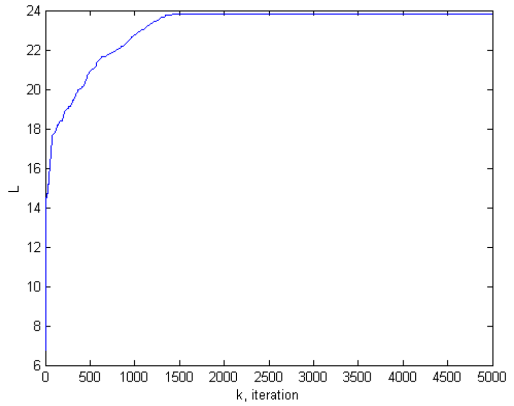

Figure 3: Convergence Rate of $L(\theta)$, Method 3

For Method 1, the resulting resource allocation matrix is, $X = \begin{bmatrix} 2 & 3 & 1 \\ 2 & 0 & 1 \\ 2 & 1 & 1 \end{bmatrix}$, and the optimized value of the objective function, *L*, is 23.2532.

For Method 2, the resulting resource allocation matrix is, $X = \begin{bmatrix} 0 & 0 & 4 \\ 2 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$, and the optimized value of the objective function, *L*, is 16.5010.

For Method 3, the resulting resource allocation matrix is, $X = \begin{bmatrix} 1 & 4 & 1 \\ 2 & 0 & 1 \\ 2 & 2 & 0 \end{bmatrix}$, and the optimized value of the objective function, *L*, is 23.8321.

8 SUMMARY

This work has illustrated a method of implementing the simultaneous perturbation stochastic approximation (SPSA) algorithm for discrete-valued constrained optimization problems. A penalty function was used to implement the resource allocation constraint criteria. Three different methods were illustrated for truncating the parameter estimates. In this application, Methods 1 and 3 arrive at very similar optimized objective function values, as compared to Method 2. Additionally, Methods 1 and 3 utilize all of the available resources of each available type. One possibility in the cause of the difference in the resource allocation distribution is the method used for parameter truncation. This will be examined in future work.

ACKNOWLEDGMENTS

This work was supported by a research contract with The Johns Hopkins University-Applied Physics Laboratory.

REFERENCES

Einbu, J. M. "A finite method for the solution of a multi-resource allocation problem with concave return functions", *Mathematics of Operations Research*, Vol. 9, No.2, pp. 232-243, May 1984.

Gerencsér, L., Hill, S. D. and Vágó, Z.. "Optimization over discrete sets via SPSA," P*roceedings of the Conference on Decision and Control, CDC 38,* pp. 1791-1794, 1999.

Spall, J. C. "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation," *IEEE Transactions on Automatic Control*, vol. 37, pp. 332-341, 1992.

Wang, I-J. and Spall, J.C. "A constrained simultaneous perturbation stochastic approximation algorithm based on penalty functions", *Proceedings of the American Control Conference*, pp. 393-399, San Diego, California, June 1999.

Whitney, J. and Hill, S. D. "Constrained optimization over discrete sets via spsa with application to resource allocation", *Proceeding of the 35th Annual Conference on Information Sciences and Systems*, pp. 515-518, March 2001.

AUTHOR BIOGRAPHIES

JAMES E. WHITNEY is an Assistant Professor of Electrical Engineering at Morgan State University. He received his Ph.D. from Marquette University in 1995. Dr. Whitney's research and engineering efforts have focused on digital and statistical signal processing methods. Current research involves investigation of signal processing algorithms for incorporation into smart sensors and instrumentation. His email and web addresses are `<whitney@eng.morgan.edu>` and `<www.eng.morgan.edu/~sigsens>`

LATASHA I. SOLOMON is a graduate of Morgan State University where she recently received a Master of Engineering Degree. Currently, she conducts research in the Signal and Sensor (SIGSENS) Laboratory at Morgan State University. Her research interests include image, signal, and statistical signal processing. Her email address is `<lsnvg@eng.morgan.edu.>`.

STACY D. HILL received his B.S. and M.S. degrees from Howard University in 1975 and 1977 respectively. He received his D.Sc. degree in control systems engineering and applied mathematics from Washington University in 1983. He is currently on the Senior Professional Staff of The Johns Hopkins University-Applied Physics Laboratory where he has been project and technical lead in developing, testing, and aplying statistical techniques and software, and has led systems analysis and modeling projects. He has published papers on diverse topics in statistics and engineering, including subjects such as simulation, optimization, and parameter estimation. His email address is `<Stacy.Hill@jhuapl.edu>`