SAMPLE-PATH SOLUTION OF STOCHASTIC VARIATIONAL INEQUALITIES, WITH APPLICATIONS TO OPTION PRICING

Gül Gürkan
 A. Yonca Özge
 Stephen M. Robinson

Department of Industrial Engineering
 University of Wisconsin–Madison
 1513 University Avenue
 Madison, WI 53706-1539, U.S.A.

ABSTRACT

This paper shows how to apply a variant of sample-path optimization to solve stochastic variational inequalities, including as a special case finding a zero of a gradient. We give a new set of sufficient conditions for almost-sure convergence of the method, and exhibit bounds on the error of the resulting approximate solution. We also illustrate the application of this method by using it to price an American call option on a dividend-paying stock.

1 INTRODUCTION

This paper shows how to use the technique of sample-path optimization (SPO) to solve stochastic variational inequalities, which include as special cases systems of nonlinear equations. This capability extends the range of application of SPO, since in some important cases the functions to be optimized are difficult to deal with, but their gradients can be approximated. In this section we review the SPO technique and give an example of a case in which its usual form is difficult to apply. The remainder of the paper then shows how to adapt the method to variational inequalities, and gives a numerical example illustrating an application of the new method.

To describe the usual form of SPO, we introduce an extended-real-valued stochastic process $\{L_n(x) | n = 1, 2, \dots\}$ where $x \in \mathbb{R}^k$. For each n and x , $L_n(x)$ is a random variable defined on the common probability space $(\Omega, \mathcal{F}, P)$; it takes values in $\mathbb{R} \cup \{+\infty\} \cup \{-\infty\}$, though we generally exclude $+\infty$ by requiring the function to be *proper* (for maximization): that is, never $+\infty$ and not everywhere $-\infty$. We can use the extended value $-\infty$ to model constraints on x , by setting $L_n(x) = -\infty$ for infeasible values of x . In the following we write $L_n(\omega, x)$ when we want to emphasize dependence of $L_n(x)$ on the sample point ω .

We assume the existence of a deterministic func-

tion L_∞ such that the L_n converge pointwise to L_∞ with probability one. A convenient interpretation is to regard the $L_n(x)$ as estimates of $L_\infty(x)$ computed by a simulation run of length n . Such convergence is often justified by regeneration theorems in the case of dynamic systems, or by the strong law of large numbers in the case of static systems. We wish to find (approximately) a constrained maximizer of $L_\infty(x)$, under the condition that one cannot observe L_∞ but only the L_n .

The method in its pure form is very simple: we fix a large simulation length n and the sample point ω (representing the random number streams used in the simulation), compute a maximizer x_n^* of the resulting deterministic function $L_n(\omega, \cdot)$, and take this point as an estimate of a maximizer of L_∞ .

This form of the method was proposed in Plambeck *et al.* (1993, 1996) for use with infinitesimal perturbation analysis (IPA) gradient estimates; the key point is that IPA – when it applies – gives *exact* gradients of the L_n , so that one can apply the powerful technology of constrained deterministic optimization. Convergence of the general method is analyzed in Robinson (1996). Similar ideas were proposed by Rubinstein and Shapiro (1993) for use with the likelihood-ratio (LR) method, and these methods are also closely related to the retrospective optimization proposals of Healy and Schruben (1991) and to M-estimation and other techniques. Robinson (1996) gives a brief survey of these and other ideas similar to SPO that have appeared in the literature; a summary of the method's properties can be found in Gürkan *et al.* (1994). Gürkan (1996) and Gürkan and Özge (1996) show how SPO can be applied to the buffer allocation problem in certain tandem production lines. Plambeck *et al.* (1996) and Gürkan (1996) report extensive numerical results on the performance of the method on fairly large systems with various constraints on the variables.

In these experiments the method performed at least

as well as the stochastic approximation method or its single-run optimization variant (Meketon (1983), Suri and Leung (1989), Leung (1990)) in cases to which the latter two methods were applicable. It was also effective for problems with inequality constraints, in which stochastic approximation can have difficulty in maintaining feasibility; see Appendix F of Plambeck *et al.* (1996) for an example. In addition, the method does not require a predetermined choice of step size, thereby avoiding a potential difficulty of the stochastic approximation approach; see Fu and Healy (1992) and L'Ecuyer *et al.* (1994).

Existing sufficient conditions for the convergence of SPO can be found in Robinson (1996). A brief summary of these conditions is as follows: we require L_∞ to be a proper deterministic function with a nonempty, compact set of maximizers, and the L_n to be (with probability one) proper, upper semicontinuous functions such that the sequence $\{-L_n\}$ *epiconverges* to $-L_\infty$. Epiconvergence is a type of convergence often used in optimization, which is independent of the usual pointwise convergence. Kall (1986) gives a good review of epiconvergence and its connections to other ideas of convergence. Under the conditions just mentioned, Robinson (1996) shows that with probability one, for large n the set M_n of maximizers of L_n will be nonempty and compact, and the distance from *any* point of M_n to *some* point of the set M_∞ of maximizers of L_∞ will be small. Thus, for large enough n , if we maximize L_n we are guaranteed to be close to a maximizer of L_∞ .

Unfortunately, for some problems found in applications it is difficult or impossible to maximize L_n . For example, Figure 1 shows the graph of L_{50} for an option-pricing problem described further in Section 3. This function is an average of step functions, and even though it would look better if we used more than 50 replications, it would still be discontinuous and extremely unpleasant to maximize. On the other hand, Figure 2 shows (also for 50 replications) the graph of a function f_{50} that is quite smooth. The f_n , which we shall show in Section 3 how to compute, have the property that f_n almost surely converges pointwise, as $n \rightarrow \infty$, to the *derivative* dL_∞ of L_∞ .

Therefore, in this case a reasonable strategy seems to be to try to find a zero of f_n for large n , in the hope that it might be close to a zero of dL_∞ , which under suitable conditions would be a maximizer of L_∞ . The theory that we shall present in Section 2 shows when this procedure can be justified. Thus, by using this new theory we can apply SPO to this problem, whereas the version requiring maximization of the L_n would have been very difficult or impossible to apply.

Figure 1: Sample option value

The rest of this paper has three numbered sections. Section 2 briefly introduces stochastic variational inequalities, and builds a theoretical framework to justify applying SPO to solve them. Particular cases covered by this theory include the first-order optimality conditions for nonlinear-programming problems, and this in turn includes unconstrained optimization by computing a critical point of the gradient. In Section 3 we illustrate the latter case by applying the method to price an American call option on a dividend-paying stock. We summarize what we have done in Section 4, and then conclude with acknowledgments and references.

2 VARIATIONAL INEQUALITIES AND SIMULATION

In this section we introduce the variational inequality problem, briefly review its significance, and indicate

Figure 2: Sample gradient

how in a very special case this problem reduces to that of solving systems of nonlinear equations. Then we show how this problem can arise in a stochastic context, and we extend the SPO approach to provide a simple algorithm to solve such problems. Finally, we state a theorem giving conditions under which this approach is justified.

The ingredients of the (deterministic) variational inequality problem are a closed convex set C (in general, a subset of a Hilbert space, but here a subset of $\mathbb{R}^k$) and a function f from an open set Φ meeting C to $\mathbb{R}^k$. The problem is to find a point $x_0 \in C \cap \Phi$, if one exists, satisfying

$$\text{For each } x \in C, \quad \langle x - x_0, f(x_0) \rangle \geq 0, \quad (1)$$

where $\langle y, z \rangle$ denotes the inner product of y and z . Geometrically, (1) means that $f(x_0)$ is an inward normal to C at x_0 .

The problem (1) models a very large number of equilibrium phenomena in economics, physics, and operations research. A survey of some of these can be found in Harker and Pang (1990). In problems arising from applications the convex set C is frequently polyhedral. As a simple example, consider the problem of solving systems of k nonlinear equations in k unknowns: to put this in the form (1) we need only take $C = \mathbb{R}^k$, so that the requirement is to find $x_0 \in \Phi$ such that $f(x_0) = 0$. Another very commonly occurring special case arises from the fact that the first-order necessary optimality conditions for a nonlinear-programming problem with continuously differentiable objective and constraint functions can be written in the form (1).

However, not all variational inequality problems arise from optimization, just as not all systems of nonlinear equations do. For example, in models of economic equilibrium, the lack of certain symmetry properties results in a model that is said to be *non-integrable*; in such a case, it is not possible to find the equilibrium prices and quantities by substituting an associated optimization problem for the variational inequality. For discussion of an actual model of this type that was heavily used in policy analysis, see Hogan (1975). The theory that we shall develop does not require any symmetry properties, so it applies to non-integrable models. In fact, one of the possible future applications that we have in mind is the solution of stochastic economic equilibrium models involving expectations or steady-state functions; in such cases one might hope that simulation together with gradient-estimation techniques could provide an effective alternative to discrete scenario representations of uncertainty, with their associated data-management problems.

We are concerned here with stochastic variational inequalities: that is, problems of the form (1) in which the function f is an (unobservable) expectation or steady-state function, and in which we can only observe f approximately by computing a sequence $\{f_n\}$ of functions (e.g., by simulation) with the property that the f_n almost surely converge to f in a sense that will be made precise below. We shall show how one can use the sample-path technique on this problem when the set C is polyhedral.

Our technique will be first to convert the variational inequality (1) into an equivalent equation which, however, will involve a function that is continuous but is generally nonsmooth even if f is smooth. Then we introduce some technical terminology that will permit us to state the main theorem. That theorem gives conditions under which solutions of variational inequalities involving the approximate functions f_n will converge to solutions of the problem that we really want to solve. We conclude this section by commenting on some special cases, one of which involves the numerical example that we present in Section 3.

To convert (1) into an equivalent equation, we introduce the *normal map* induced by f and C , defined by $f_C(z) = f(\Pi_C(z)) + z - \Pi_C(z)$, where Π_C is the Euclidean projector on C . The function f_C is then well defined on $\Pi_C^{-1}(\Phi)$, and it is clearly continuous if f is. Further, it is easy to show that if x_0 satisfies (1) then $z_0 = x_0 - f(x_0)$ satisfies

$$f_C(z_0) = 0, \quad (2)$$

and that if (2) holds then $x_0 = \Pi_C(z_0)$ satisfies (1). Therefore for purposes of computation (1) and (2) are equivalent; we shall work with (2). For additional information and further references on normal maps, see e.g. Robinson (1995).

The new form of SPO works in the setting of a vector-valued stochastic process $\{f_n(x) \mid n = 1, 2, \dots\}$, and a vector-valued function $f_\infty(x)$, where the parameter x takes values in $\mathbb{R}^k$. Again, for all $n \geq 1$ and all $x \in \mathbb{R}^k$, the random variables $f_n(x)$ are defined on a common probability space $(\Omega, \mathcal{F}, P)$. The aim is to find a point x_0 satisfying (1) for $f = f_\infty$, provided that one exists. We proceed by fixing a large simulation length n and the sample point ω , finding a point $x_n^*(\omega)$ solving (1) for $f = f_n$ (again, provided that one exists) and taking this point as an estimate of x_0 . We shall give conditions below that ensure the existence of these quantities.

In the particular special case of an unconstrained simulation optimization problem, C would be $\mathbb{R}^k$ and f_∞ would be the gradient of the limit function L_∞ .

By finding a zero of f_∞ (a critical point of L_∞), we are solving the first-order necessary optimality conditions, which is the approach taken by most deterministic optimization algorithms. Of course, such a critical point may or may not be an maximizer, and additional work is generally necessary to determine whether the appropriate second-order conditions hold. In the unconstrained case this amounts to checking for negative definiteness of the Hessian at the critical point. In some cases, notably when the problem is unconstrained and the objective function is known to be locally concave near the critical point, no second-order analysis is needed.

It turns out that *continuous convergence* plays a crucial role in the convergence analysis of this method. This notion of convergence is equivalent to uniform convergence to a continuous limit on compact sets. For an elementary treatment of the relationship between different types of functional convergence, see Kall (1986), and for a comprehensive treatment of continuous convergence and related issues, see Rockafellar and Wets (1996).

Definition 1 A sequence f_n of extended-real-valued functions defined on $\mathbb{R}^k$ converges continuously to an extended-real-valued function f_∞ defined on $\mathbb{R}^k$ (written $f_n \xrightarrow{c} f_\infty$) if for any $x \in \mathbb{R}^k$ and any sequence $\{x_n\}$ converging to x , one has $f_n(x_n) \rightarrow f_\infty(x)$. A sequence of functions from $\mathbb{R}^k$ into $\mathbb{R}^m$ converges continuously if each of the m component functions does so.

The fact motivating our use of continuous convergence is that if $f_n \xrightarrow{c} f_\infty$ and $x_n \rightarrow x$, and if for each n , x_n solves (1) with $f = f_n$, then x solves (1) with $f = f_\infty$. This is very easy to check directly from (1). Thus, if the problems with $f = f_n$ have solutions, and if those solutions converge, then their limit will solve the problem with $f = f_\infty$.

We present next a theorem, for the case in which C is polyhedral, giving sufficient conditions for these approximate solutions to exist and to converge. In order to state this theorem we need a technical property that generalizes nonsingularity. We explain it very briefly here, and then mention some important special cases in which the property is easier to understand. For additional examples in more general cases we refer to Robinson (1995).

For this explanation we need some definitions. First, the normal cone $N_C(x)$ of C at x is the set

$$\{y^* \mid \text{for each } c \in C, \quad \langle y^*, c - x \rangle \leq 0\}$$

provided that $x \in C$, and it is empty otherwise. If $x \in C$ then the *tangent cone* of C at x , written $T_C(x)$,

is the polar of $N_C(x)$: that is, the set of all y such that $\langle y^*, y \rangle \leq 0$ for each $y^* \in N_C(x)$. Second, the *critical cone* defined by C and a given point $z \in \mathbb{R}^k$ is

$$K(z) = T_C(\Pi_C(z)) \cap \{y^* \in \mathbb{R}^k \mid \langle y^*, z - \Pi_C(z) \rangle = 0\}.$$

Now fix any z and write $K = K(z)$. As K is polyhedral it has only finitely many faces; for each nonempty face F the normal cone of K takes a constant value, say N_F , on the relative interior of F . Then the set $\sigma_F = F + N_F$ is a nonempty polyhedral convex set of dimension k in $\mathbb{R}^k$. The collection $\mathcal{N}_K$ of all these σ_F for nonempty faces F of K is called the *normal manifold* of K . In each of these σ_F the projector Π_K reduces to an affine map (generally different for different σ_F). If M is a linear transformation from $\mathbb{R}^k$ to $\mathbb{R}^k$, then we say that the normal map M_K is *coherently oriented* if in each σ_F the determinant of the affine map obtained by restricting M_K to σ_F has the same nonzero sign.

As a simple illustration of this property, we can consider the case in which K happens to be a subspace (the “strict complementary slackness” situation in the optimization literature). Then the coherent orientation requirement reduces to nonsingularity of the section of M in K . In particular, if $C = \mathbb{R}^k$ (the case of nonlinear equations), then $K(z) = \mathbb{R}^k$ for each $z \in \mathbb{R}^k$, and then $\mathcal{N}_{K(z)}$ has only one cell, namely $\mathbb{R}^k$ itself. Then M_K is coherently oriented exactly when M is nonsingular. In general, the coherent orientation condition is a way of generalizing nonsingularity to the case of a nontrivial set C .

Theorem 1 Let Φ be an open subset of $\mathbb{R}^k$ and let C be a polyhedral convex set in $\mathbb{R}^k$. Let x_0 be a point of Φ , and suppose f_∞ is a function from Φ to $\mathbb{R}^k$. Let $\{f_n \mid n = 1, 2, \dots\}$ be random functions from Φ to $\mathbb{R}^k$ such that for all $x \in \Phi$ and all finite n the random variables $f_n(x)$ are defined on a common probability space $(\Omega, \mathcal{F}, P)$. Let $z_0 = x_0 - f_\infty(x_0)$ and assume the following:

- With probability one, each f_n for $n = 1, 2, \dots$ is continuous and $f_n \xrightarrow{c} f_\infty$.
- x_0 solves the variational inequality (1) defined by f_∞ and C .
- f_∞ has a strong Fréchet derivative $df_\infty(x_0)$ at x_0 , and $df_\infty(x_0)_{K(z_0)}$ is coherently oriented.

Then there exist a positive number λ , a compact subset $C_0 \subset C \cap \Phi$ containing x_0 , a neighborhood $V \subset \Phi$ of x_0 , and a set $\Delta \subset \Omega$ of measure zero, with the following properties: for $n = 1, 2, \dots$ and $\omega \in \Omega$ let

$$\xi_n(\omega) := \sup_{x \in C_0} \|f_n(\omega, x) - f_\infty(x)\|,$$

and

$$X_n(\omega) := \{x \in C \cap V \mid \text{for each } c \in C, \langle f_n(\omega, x), c - x \rangle \geq 0\}.$$

For each $\omega \notin \Delta$ there is then a finite integer N_ω such that for each $n \geq N_\omega$ the set $X_n(\omega)$ is a nonempty, compact subset of the ball

$$B(x_0, \lambda \xi_n(\omega)) = \{x \mid \|x - x_0\| \leq \lambda \xi_n(\omega)\}.$$

In the rest of this section we discuss some work in the literature related to Theorem 1 and point out similarities and differences. First, King and Rockafellar (1993) gave various convergence and asymptotic properties for solutions to stochastic generalized equations, of which the variational inequalities that we treat here are a particular case. Among other things they proved under certain regularity conditions that the limit superior of a localization of the perturbed solution set was nonempty and contained in the solution set of the limit problem; this corresponds in our case to a proof that the solutions of the problems involving the f_n have cluster points, and that each such point solves the limit problem. They also showed under additional assumptions that the solutions of the perturbed problems converged in probability to the solution of the limit problem, with the speed of convergence dependent on the goodness of the approximation. Our conditions are in some ways more restrictive than theirs, but we are able to establish almost-sure convergence and to bound the error of the approximate solutions by a constant multiple of the error in the approximation. Under considerably stronger assumptions than ours, they also obtained asymptotic results about the distribution of errors, a question which we do not treat at all here.

Shapiro and Wardi (1994, 1995) used conditions somewhat similar to those of Theorem 1 (in the special case $C = \mathbb{R}^k$) to prove the convergence of stochastic gradient descent algorithms. However, there are significant differences: for example, they made no assumption that the limiting function had a nonsingular derivative, but they derived no bounds on the resulting solutions. Also, they assumed that the functions involved were gradients or elements of generalized gradients. As we have indicated above, we have in mind applications in which it is important that the functions *not* be restricted to be gradients.

Shapiro (1995) investigates the convergence of approximate solutions of quite general problems by introducing a “stopping criterion” to measure the quality of approximation. However, he assumes the existence of solutions to the approximating problems, whereas under the hypotheses of Theorem 1 we can

prove the existence of such solutions as well as bounding their distance from the true solution. On the other hand, Shapiro (1995) also gives results on the asymptotic distribution of solutions, as did King and Rockafellar (1993). We do not give such results.

As we have shown in Theorem 1, we can find approximate solutions of stochastic variational inequalities by a variant of sample-path optimization, provided that a generalized nonsingularity condition holds. Moreover, as $n \rightarrow \infty$ the distance of the approximate solutions from an exact solution is bounded above by a constant multiple of the uniform norm of $f_n - f_\infty$ on a compact set containing the true solution. In the next section we illustrate an application of this general result.

3 APPLICATION: OPTION PRICING

In this section we present a small numerical example illustrating how to apply SPO to the pricing of an American call option on a dividend-paying stock. Such an option is a derivative security giving the holder the right to *exercise* the option (by buying an underlying asset) at the *exercise price* K at any time up to the *expiration date* T . Under certain assumptions about markets, including the absence of transaction costs, it is never optimal to exercise an American call option prematurely (*i.e.*, before T) unless the stock pays a dividend during the life of the option contract (see Hull (1993) and Stoll and Whaley (1993)). In such a case, in order to obtain the right to receive the dividend it may be optimal to exercise the option just prior to an *ex-dividend date* t_1 , provided that the stock price then exceeds a *threshold price*. Note that the choice of a threshold price is under the control of the option holder. We can value the option by finding its greatest expected return over all possible choices of threshold price.

Early examples of using Monte Carlo simulation in option pricing can be found in Boyle (1977), Hull and White (1987), and Johnson and Shanno (1987). In these studies, simulation was employed as a “black box” without providing any sensitivity information. More recently, Fu and Hu (1995) and Broadie and Glasserman (1996) have improved this approach by using gradient estimation methods. Welch and Chen (1991) formulated the pricing of an American call option on a dividend-paying stock as an optimization problem. Fu and Hu (1995) adapted this viewpoint and applied stochastic approximation to value this option.

Following the notation and discussion in Fu and Hu (1995), we define the variables that will be used

in the rest of this section as follows :

S_0 = the initial stock price,

S_t = the stock price at time t ,

r = the annualized riskless interest rate (compounded continuously),

σ = volatility of the stock price process,

D = the dividend amount,

J_T = the net present value of the return on the option.

Among these variables S_t and J_T are random and the rest are deterministic. We assume that after the ex-dividend date the stock price drops by the dividend amount: i.e., $S_{t_1^+} = S_{t_1^-} - D$. The sample option value for the American call is then given by

$$J_T = e^{-rT} [I_{\{S_{t_1^-} > s\}} (S_{t_1^-} - K)e^{r(T-t_1)} + I_{\{S_{t_1^-} \leq s\}} (S_T - K)_+] \quad (3)$$

where I_A is the indicator function of the set A in the sense of probability theory, and $(\cdot)_+ = \max\{\cdot, 0\}$ (the “plus function”). We assume that except for the downward jump at the ex-dividend date the stock price changes continuously according to a function $h(\omega; S, t, r, \sigma)$ which gives the stock price at time t net of the present value of the dividend, given the stock price S at time 0 and the random vector ω . In this work we use the particular form of h that arises when the stock price fluctuations are described by the lognormal distribution, namely

$$h(\omega; S, t, r, \sigma) = Se^{(r-\sigma^2/2)t + \sigma\omega\sqrt{t}}, \quad (4)$$

where ω is a standard $N(0, 1)$ random variable. For $N(0, 1)$ random variables ω_1 and ω_2 we have

$$S_{t_1^-} = h(\omega_1; S_0 - De^{-rt_1}, t_1, r, \sigma) + D$$

and

$$S_T = h(\omega_2; S_{t_1^+}, T - t_1, r, \sigma).$$

We construct the functions L_n by drawing n samples from the standard normal distribution and averaging the resulting values of J_T . For a fixed, finite set of samples, L_n is therefore a deterministic step function of the threshold price s . As can be seen in Figure 1, the L_n are unsuitable for optimization. We therefore applied the general approach of Section 2 (with $C = \mathbb{R}$) by using an unbiased gradient estimator for $dE[J_T]/ds =: f_\infty$ provided in Fu and Hu (1995) to construct approximating functions f_n . Figure 2 shows a representative of f_{50} , which is quite well behaved. We next discuss briefly the construction of the f_n and some of their properties.

In Shi (1995) a discontinuous perturbation analysis (DPA) gradient estimator is derived for general

step functions of one variable, and in Fu and Hu (1995) a smoothed perturbation analysis (SPA) gradient estimator is derived specifically for the gradient of $E[J_T]$ with respect to various parameters of the option. Each of these approaches results in the same unbiased gradient estimator for $dE[J_T]/ds$, which is shown in (5). Denoting by p the probability density function of ω and setting $\tau = T - t_1$, we have

$$\left(\frac{dJ_T}{ds}\right)_{PA} = e^{-rT} \frac{\partial h^{-1}(y^*)}{\partial s} p(h^{-1}(y^*)) [(h(\omega; s - D, \tau, r, \sigma) - K)_+ - (s - K)e^{r\tau}] \quad (5)$$

where $y^* = (s - D; S_0 - De^{-rt_1}, t_1, r, \sigma)$ and the shorthand notation $h^{-1}(y^*)$ stands for

$$h(\cdot; S_0 - De^{-rt_1}, t_1, r, \sigma)^{-1}(s - D).$$

For $y > 0$ and the choice of h given by (4), we have

$$h(\cdot; S, t, r, \sigma)^{-1}(y) = \frac{1}{\sigma\sqrt{t}} [\ln \frac{y}{S} - (r - \sigma^2/2)t],$$

while for $s > D$ the partial derivative in (5) is

$$\frac{\partial h^{-1}(y^*)}{\partial s} = \frac{1}{(s - D)\sigma\sqrt{t_1}}.$$

For $s \in (D, +\infty)$, the expression in (5) is therefore a continuous function of the threshold price s .

We constructed the approximating functions f_n by averaging the values of $(dJ_T/ds)_{PA}$. One can show that $f_n \xrightarrow{C} f_\infty$ and that f_∞ is continuously F-differentiable. Hence if there is an s^* with $dE[J_T(s^*)]/ds = 0$ and $d^2E[J_T(s^*)]/ds^2 < 0$, then Theorem 1 shows that for sufficiently large sample sizes n , the f_n will have zeros and these zeros will be close to a locally unique maximizer s^* .

In this example we used for the function h the particular form (4) arising from the geometric Brownian motion, since this is a popular stock price process in the finance community. Other forms could also be used, provided that they satisfy conditions to ensure the unbiasedness of the gradient estimate; see Fu and Hu (1995) and Shi (1995).

We now present the results of a small numerical experiment. The data used are taken from Stoll and Whaley (1993): $S_0 = 50.0$, $K = 50.0$, $T = 90.0$, $t_1 = 60.0$, $D = 2.0$, $r = 0.10$, $\sigma = 0.30$. This is a problem with only one dividend; hence one can use the analytical formula to find the optimal threshold price and the option value. Stoll and Whaley (1993) found these to be $s^* = 51.82$ and $E[J_T] = 2.93$.

The same problem is considered in Fu and Hu (1995); they applied stochastic approximation and reported that “convergence within a penny of the actual

option value is achieved within 1000 simulations.” Fu and Hu (1995) used 9 replications, which means that their total simulation budget was 9,000.

We applied SPO by finding a zero of f_n using the nonsmooth Newton method of Qi and Sun (1993). Results are presented in Table 1. The true solution is found within one penny with a simulation budget of 7,000 (two iterations, each with a sample size of 3,500). The table also gives results for larger sample sizes n . In addition, we used a very large sample size of 500,000 and found the exact solution (reported above). In each case the method converged in 2 iterations. Thus, on this problem SPO performed at least comparably to stochastic approximation.

Table 1: Results of SPO for Option Pricing

n	Option value	Threshold price	# Iter
35E2	2.937	51.850	2
50E2	2.931	51.803	2
10E4	2.939	51.828	2

4 CONCLUSIONS

In this paper we have shown how to use a simulation-based method called sample-path optimization (SPO) for solving stochastic variational inequalities. We presented a new set of sufficient conditions for the convergence of the method, and by applying these to the special case of nonlinear equations we showed how the method can be applied to the pricing of certain American call options. As an illustration we presented the results of a small numerical experiment. On that small problem SPO performed at least as well as stochastic approximation. However, the SPO technique appears to be applicable to a considerably larger class of problems than is stochastic approximation. Therefore, the sufficient conditions that we have given here may help in the solution of problems that are difficult to handle with current techniques.

ACKNOWLEDGMENT

The research reported here was sponsored by the U. S. Army Research Office under grant number DAAH04-95-1-0149, by the Air Force Office of Scientific Research, Air Force Materiel Command, USAF, under grant number F49620-95-1-0222, and by the North Atlantic Treaty Organization under grant number CRG 950360.

REFERENCES

- Boyle, P. P. 1977. Options: a Monte Carlo approach. *Journal of Financial Economics* 4:323–338.
- Broadie, M., and P. Glasserman. 1996. Estimating security price derivatives using simulation. *Management Science* 42:269–285.
- Fu, M. C., and K. Healy. 1992. Simulation optimization of (s, S) inventory systems. In *Proceedings of the 1992 Winter Simulation Conference*, ed. J. J. Swain, D. Goldsman, R. C. Crain and J. R. Wilson, 506–514. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Fu, M. C., and J.-Q. Hu. 1995. Sensitivity Analysis for Monte Carlo Simulation of Option Pricing. *Probability in the Engineering and Informational Sciences* 9:417–449.
- Gürkan, G. 1996. Performance optimization in simulation: Sample-path optimization of buffer allocations in tandem lines. Ph.D. Dissertation, Department of Industrial Engineering, University of Wisconsin-Madison, Madison, WI.
- Gürkan, G., and A. Y. Özge. 1996. Sample-path optimization of buffer allocations in a tandem queue – Part I: Theoretical issues. Preprint, Department of Industrial Engineering, University of Wisconsin-Madison, Madison, WI.
- Gürkan, G., A. Y. Özge, and S. M. Robinson. 1994. Sample-path optimization in simulation. In *Proceedings of the 1994 Winter Simulation Conference*, ed. J. D. Tew, S. Manivannan, D. A. Sadowski and A. F. Seila, 247–254. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Harker, P. T., and J.-S. Pang. 1990. Finite-dimensional variational inequality and nonlinear complementarity problems: A survey of theory, algorithms, and applications. *Mathematical Programming* 48:161–220.
- Healy, K., and L. W. Schruben. 1991. Retrospective simulation response optimization. In *Proceedings of the 1991 Winter Simulation Conference*, ed. B. L. Nelson, W. D. Kelton, and G. M. Clark, 901–906. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Hogan, W. W. 1975. Energy policy models for Project Independence. *Computers & Operations Research* 2:251–271.
- Hull, J. C. 1993. *Options, Futures, and Other Derivative Securities*, 2nd ed. Englewood Cliffs, NJ: Prentice-Hall.
- Hull, J. C., and A. White. 1987. The pricing of options on assets with stochastic volatilities. *Journal of Finance* 42:281–300.

- Johnson, H., and D. Shanno. 1987. Option pricing when the variance is changing. *Journal of Financial and Quantitative Analysis* 22:143–151.
- Kall, P. 1986. Approximation to optimization problems: An elementary review. *Mathematics of Operations Research* 11:9–18.
- King, A. J., and R. T. Rockafellar. 1993. Asymptotic theory for solutions in statistical estimation and stochastic programming. *Mathematics of Operations Research* 18:148–162.
- L'Ecuyer, P., N. Giroux, and P. W. Glynn. 1994. Stochastic optimization by simulation: numerical experiments with the $M/M/1$ queue in steady-state. *Management Science* 40:1245–1261.
- Leung, Y. T. 1990. Single-run optimization of discrete-event simulations. Ph.D. Dissertation, Department of Industrial Engineering, University of Wisconsin–Madison, Madison, Wisconsin.
- Qi, L., and Sun, J. 1993. A nonsmooth version of Newton's method. *Mathematical Programming* 58:353–367.
- Meketon, M. 1983. A tutorial on optimization in simulations. Unpublished tutorial presented at the 1983 Winter Simulation Conference.
- Plambeck, E. L., B.-R. Fu, S. M. Robinson, and R. Suri. 1993. Throughput optimization in tandem production lines via nonsmooth programming. In *Proceedings of 1993 Summer Computer Simulation Conference*, ed. J. Schoen, 70–75. San Diego, CA: Society for Computer Simulation.
- Plambeck, E. L., B.-R. Fu, S. M. Robinson, and R. Suri. 1996. Sample-path optimization of convex stochastic performance functions. Accepted by *Mathematical Programming*.
- Robinson, S. M. 1995. Sensitivity analysis of variational inequalities by normal-map techniques. In *Variational Inequalities and Network Equilibrium Problems*, ed. F. Giannessi and A. Maugeri, 257–269. New York: Plenum Press.
- Robinson, S. M. 1996. Analysis of sample-path optimization. Accepted by *Mathematics of Operations Research*.
- Rockafellar, R. T., and R. J-B Wets. 1996. *Variational Analysis*. Berlin: Springer-Verlag (forthcoming).
- Rubinstein, R. Y., and A. Shapiro. 1993. *Discrete Event Systems: Sensitivity Analysis and Stochastic Optimization by the Score Function Method*. Chichester and New York: John Wiley & Sons.
- Shapiro, A. 1995. Simulation based optimization – Convergence analysis and statistical inference. To appear in *Stochastic Models*.
- Shapiro, A., and Y. Wardi. 1994. Convergence analysis of stochastic algorithms. To appear in *Mathematics of Operations Research*.
- Shapiro, A., and Y. Wardi. 1995. Convergence analysis of gradient descent stochastic algorithms. Preprint, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA.
- Shi, L. 1995. Discontinuous Perturbation Analysis of Discrete Event Dynamic Systems. Preprint, Department of Industrial Engineering, University of Wisconsin - Madison, Madison, WI.
- Stoll, H. R., and R. E. Whaley. 1993. *Futures and Options*. Cincinnati: South-Western Publishing Co.
- Suri, R., and Y. T. Leung. 1989. Single-run optimization of discrete event simulations – an empirical study using the $M/M/1$ queue. *IIE Transactions* 34:35–49.
- Welch, R. L., and D. M. Chen. 1991. Static optimization of American contingent claims. *Advances in Futures and Options Research* 5:175–184.

AUTHOR BIOGRAPHIES

GÜL GÜRKAN is Research Assistant in the Department of Industrial Engineering, University of Wisconsin–Madison. Her research interests include stochastic and deterministic optimization, simulation, and stochastic processes. She is a member of INFORMS.

A. YONCA ÖZGE is Research Assistant in the Department of Industrial Engineering, University of Wisconsin–Madison. Her research interests include deterministic optimization, stochastic optimization, and simulation. She is a member of INFORMS.

STEPHEN M. ROBINSON is Professor of Industrial Engineering and Computer Sciences at the University of Wisconsin–Madison. His research interests are in quantitative methods in managerial economics, especially deterministic and stochastic optimization and methods to support decision under uncertainty. He is a member of INFORMS and of the Mathematical Programming Society, as well as of SIAM and IIE.