Proceedings of the 2014 Winter Simulation Conference
A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, eds.

A PENALTY FUNCTION APPROACH FOR SIMULATION OPTIMIZATION WITH STOCHASTIC CONSTRAINTS

Liujia Hu
Sigrún Andradóttir

H. Milton Stewart School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332, USA

ABSTRACT

This paper is concerned with continuous simulation optimization problems with stochastic constraints. Thus both the objective function and constraints need to be estimated via simulation. We propose an Adaptive Search with Discarding and Penalization (ASDP) method for solving this problem. ASDP utilizes the penalty function approach from deterministic optimization to convert the original problem into a series of simulation optimization problems without stochastic constraints. We present conditions under which the ASDP algorithm converges almost surely, and conduct numerical studies aimed at assessing its efficiency.

1 INTRODUCTION

This paper develops a provably convergent algorithm for solving optimization problems involving continuous decision variables, constraints, and uncertainties. More specifically, in the following, we consider the benchmark problem:

$$\begin{aligned} &\underset{\theta\in\Theta}{\text{Maximize}} \quad f(\theta) = E[h(\theta, X(\omega))] \\ &\text{subject to} \quad g_j(\theta) = E[u_j(\theta, Y_j(\omega))] \le b_j,\ j \in \mathscr{C}, \end{aligned} \tag{1}$$

where the feasible region Θ is a subset of $\mathbb{R}^s$, E denotes the mathematical expectation operation, $\mathscr{C}$ is a finite set of indexes, $X(\omega)$ and $Y_j(\omega)$, $j \in \mathscr{C}$, are random elements defined on some probability space $(\Omega, \Sigma, \mathbb{P})$, and h, u_j, $j \in \mathscr{C}$, are deterministic functions. We allow the feasible region Θ to be uncountable and assume $f(\theta)$ and $g_j(\theta)$ at any $\theta \in \Theta$ and $j \in \mathscr{C}$ cannot be evaluated exactly. Thus, if there are deterministic constraints, we assume they are incorporated in the feasible region Θ. Let f^* be the optimal objective value of the optimization problem (1) and $\bar{f} = \sup_{\theta\in\Theta} f(\theta)$, so that $f^* \le \bar{f}$. Assume $\bar{f} < \infty$.

Several simulation optimization techniques have been developed to solve the problem (1) when Θ is discrete and no stochastic constraints are present; a thoroughly review can be found in Fu (2002). When Θ is discrete, and stochastic constraints are involved in (1), the Optimal Computing Budget Allocation (OCBA) approach can be used to maximize the probability of selecting the best system, see, e.g., Pujowidianto et al. (2009) and Hunter and Pasupathy (2013). Moreover, Li, Sava, and Xie (2009) and Park and Kim (2011) proposed random search methods where stochastic constraints are taken into account in an augmented performance function via a nonnegative penalty factor.

When Θ is continuous and there are no stochastic constraints in (1), a few simulation optimization algorithms that do not use gradient information have been proposed. For example, Baumert and Smith (2002) proposed a deterministic shrinking ball method based on pure random search and Yakowitz and Lugosi (1990) developed a global random search with resampling method. Moreover, Norkin, Pflug, and Ruszczyński (1998) proposed a stochastic branch and bound method, Rubinstein and Kroese (2004) developed the cross entropy approach, Hu, Fu, and Marcus (2008) proposed stochastic model reference

978-1-4799-7486-3/14/$31.00 ©2014 IEEE

adaptive search, Andradóttir and Prudius (2010) proposed adaptive search with resampling, and Hu and Andradóttir (2014a) proposed adaptive search with resampling and discarding.

When Θ is continuous and stochastic constraints are involved in (1), the sample average approximation approach can handle the stochastic constraints, see, for example, Dentcheva and Ruszczynski (2003) and Pagnoncelli, Ahmed, and Shapiro (2009). However, to the best of our knowledge, very few papers have addressed provably convergent *random search* algorithms to solve the benchmark problem (1). The main advantage of random search over sample average approximation is that it assumes little structure (of the objective and constraint functions). In this paper, we propose a provably convergent algorithm called Adaptive Search with Discarding and Penalization (ASDP) to solve (1). We use a sequence of positive real numbers to dynamically penalize the sample points that appear to be infeasible or whose feasibility is ambiguous (meaning that they appear to be very close to the boundary of the feasible region), as well as two other sequences of non-negative real numbers to discard points that are likely to be infeasible and/or inferior.

The remainder of this article is organized as follows. In Section 2, we present our ASDP algorithm. In Section 3, we prove its almost sure convergence. In Section 4, we provide a numerical study. In Section 5, we summarize the main contributions of this article.

2 THE ASDP ALGORITHM

In this section, we present our algorithm for solving continuous simulation optimization problems with stochastic constraints. We start by introducing some notation. For all $\theta \in \Theta$ and $k \in \mathbb{N}$, let $N_k(\theta)$ be the number of observations of the objective function $f(\theta)$ as well as the constraint functions $g_j(\theta)$ by the end of iteration k, and let $S_k(\theta)$ be the sum of these $N_k(\theta)$ observations of $f(\theta)$ and $S_{j,k}(\theta)$ be the sum of these $N_k(\theta)$ observations of $g_j(\theta)$ for all $j \in \mathscr{C}$. Also, for all $\theta \in \Theta$, $j \in \mathscr{C}$, and $k \in \mathbb{N}$, let $\hat{f}_k(\theta) = S_k(\theta)/N_k(\theta)$ and $\hat{g}_{j,k}(\theta) = S_{j,k}(\theta)/N_k(\theta)$. In addition, a sequence $\{a_k\}$ is said to be $O(k^n)$ for some $n \in \mathbb{R}$ if there exists a $C_1 \in \mathbb{R}^+$ such that $0 \leq a_k \leq C_1 k^n$ for all $k \in \mathbb{N}$. A sequence $\{a_k\}$ is said to be $\Phi(k^n)$ for some $n \in \mathbb{R}$ if there exists a $C_2 \in \mathbb{R}^+$ such that $a_k \geq C_2 k^n$ for all $k \in \mathbb{N}$. A sequence $\{a_k\}$ is said to be $\Omega(k^n)$ for some $n \in \mathbb{R}$ if it is both $O(k^n)$ and $\Phi(k^n)$.

Let $\{M(i)\}_{i=1}^{\infty}$ be a strictly increasing sequence of positive integers with $M(1) = 1$. In our ASDP Algorithm, we alternate between adaptively sampling from the feasible region (if the current iteration number is equal to some element in the sequence $\{M(i)\}_{i=1}^{\infty}$), or resampling a previously sampled point (otherwise). After a new point has been sampled, we decide whether or not to accept the newly sampled point (because it appears promising), ensure that we have collected enough objective function observations at each sampled point under consideration, and update the estimate of the optimal solution. Then those points exhibiting inferior qualities are discarded.

Due to the randomness involved in the constraints in problem (1), a penalty addressing the feasibility of estimated constraints is added to the estimate of the objective function as follows. Let $\{\lambda_i\}_{i=1}^{\infty}$, $\{\xi_i\}_{i=1}^{\infty}$ be two sequences of positive real numbers. For all $\theta \in \Theta$ and $i \in \mathbb{N}^+$, at iteration $M(i)$, define

$$F_i(\theta) = \hat{f}_{M(i)}(\theta) - \lambda_i G_i(\theta, \xi_i),$$

where $G_i(\theta, \xi_i) = \mathbb{1}_{\{\sum_{j \in \mathscr{C}} \mathbb{1}_{\{\hat{g}_{j,M(i)}(\theta) > b_j - \xi_i\}} \geq 1\}}$ and $\mathbb{1}_A$ is 1 if event A is true, 0 otherwise. Note that the penalty $G_i(\theta, \cdot)$ is positive when the current point θ either appears to be infeasible or shows ambiguity between feasible and infeasible (meaning that θ appears to be feasible but is very close to the boundary). Here we use the sequence $\{\xi_i\}_{i=1}^{\infty}$ to control our criteria to test feasibility, and $\{\lambda_i\}_{i=1}^{\infty}$ to control the scale of penalty when $G_i(\theta, \xi_i)$ is positive.

Next, let Θ_i be the set of solutions sampled and accepted by the end of iteration $M(i)$ without discarding already accepted points. Let Θ_i^* denote the set of solutions sampled, accepted, and not discarded by the end of iteration $M(i)$. Let Θ_i^+ be the set of solutions sampled and accepted by iteration $M(i)$, and not discarded prior to the discarding procedure in iteration $M(i)$. Our ASDP algorithm focuses on finding the optimum $\theta_i^* \in \arg\max_{\theta \in \Theta_i^*} F_i(\theta)$, and its pseudo-code is given in Algorithm 1.

Algorithm 1 Adaptive Search with Discarding and Penalization (ASDP).

Select $c > 0$, $\{\lambda_i\}_{i=1}^{\infty}$, $\{\xi_i\}_{i=1}^{\infty}$, and $\{\delta_i\}_{i=1}^{\infty}$, three sequences of positive real numbers, $\{K(i)\}_{i=1}^{\infty}$, a nondecreasing sequence of positive integers with $K(i) = \Phi(i^c)$, a sampling strategy, a resampling strategy, and an acceptance criterion. Let $\eta_i = 3\xi_i$, $\forall i \in \mathbb{N}^+$, $\Theta_0^* = \emptyset$, $i = 1$, and $k = 0$.
while Stopping criterion is not satisfied **do**
  Let $k = k+1$
  if $k = M(i)$ **then**
    Sample a solution θ_i from Θ using the sampling strategy
    Based on the acceptance criterion, decide whether to include θ_i in the set Θ_i^+, so that $\Theta_i^+ \in \{\Theta_{i-1}^*, \Theta_{i-1}^* \cup \{\theta_i\}\}$, and update $N_k(\theta_i)$, $S_k(\theta_i)$, and $S_{j,k}(\theta_i)$ if needed
    For each $\theta \in \Theta_i^+$, if $N_k(\theta) < K(i)$, obtain $K(i) - N_k(\theta)$ additional observations of $f(\theta)$ and $g_j(\theta)$ $(j \in \mathscr{C})$, and update $N_k(\theta)$, $S_k(\theta)$, and $S_{j,k}(\theta)$ accordingly
    Let $\Theta_i^* = \Theta_i^+$
    Select an estimate of the current best solution $\theta_i^* \in \arg\max_{\theta \in \Theta_i^*} F_i(\theta)$
    if $\hat{g}_{j,k}(\theta_i^*) \leq b_j - \eta_i, \forall j \in \mathscr{C}$ **then**
      For each $\theta \in \Theta_i^*$, if $F_i(\theta_i^*) - F_i(\theta) > \delta_i$, remove θ from Θ_i^* and update $\Theta_i^* = \Theta_i^* \setminus \{\theta\}$
    end if
    Let $i = i+1$
  else
    Sample a solution θ from Θ_{i-1}^* using the resampling strategy
    Obtain additional estimates of $f(\theta)$, $g_j(\theta)$ $(j \in \mathscr{C})$, and update $N_k(\theta)$, $S_k(\theta)$, and $S_{j,k}(\theta)$
  end if
end while
Return θ_{i-1}^* as an estimate of the optimal solution.

3 CONVERGENCE ANALYSIS

In this section, we present our main convergence result for Algorithm 1. Let $\Theta_{\mathscr{F}} = \{\theta \in \Theta | g_j(\theta) \leq b_j, \forall j \in \mathscr{C}\}$. For each $\varepsilon, \Delta \in \mathbb{R}$, define $\Theta_\varepsilon = \{\theta \in \Theta | f(\theta) \geq f^* - \varepsilon\}$ and $\Theta_{\mathscr{F},\Delta} = \{\theta \in \Theta | g_j(\theta) \leq b_j - \Delta, \forall j \in \mathscr{C}\}$, and let $\Theta_{\varepsilon,\Delta} = \Theta_\varepsilon \cap \Theta_{\mathscr{F},\Delta}$. Define $\theta \in \Theta$ to be a "near-optimal" point with respect to ε if $\theta \in \Theta_\varepsilon$. Let $f_n(\theta)$ be the estimate of $f(\theta)$ obtained from a sample average of n independent observations of $f(\theta)$, $g_{j,n}(\theta)$ be the estimate of $g_j(\theta)$ obtained from a sample average of n independent observations of $g_j(\theta)$, where $j \in \mathscr{C}$. We need the following assumptions.

Assumption 1 For each $\theta \in \Theta$, we can generate independent and unbiased observations $\{h(\theta, X_k(\omega)\}$ of $f(\theta)$, and $\{u_j(\theta, Y_{j,k}(\omega)\}$ of $g_j(\theta)$ for each $j \in \mathscr{C}$. Moreover, there exist $l, w \in \mathbb{N}\backslash\{0,1\}$ and $R \in \mathbb{R}^+$ such that $E[(h(\theta, X_k(\omega)) - f(\theta))^{2l}] \leq R$ and $E[(u_j(\theta, Y_{j,k}(\omega)) - g_j(\theta))^{2w}] \leq R$ for $j \in \mathscr{C}$, $\theta \in \Theta$, and $k \in \mathbb{N}^+$.

Assumption 2 The random elements used for estimating the objective function and constraints values (e.g., in steps 7 and 16 of ASDP) are independent of the random elements used in the execution of algorithmic decisions (e.g., in steps 5 and 15 of ASDP).

Assumption 3 For each $\varepsilon > 0$, there exists $\Delta(\varepsilon) > 0$ such that $P(\theta_i \in \Theta_i \cap \Theta_{\varepsilon,\Delta(\varepsilon)}, i.o.) = 1$, where *i.o.* stands for "infinitely often."

Assumption 1 imposes the finiteness of moments for the random variables under consideration in this paper. Assumption 2 is an assumption about implementation that can always be satisfied and allows for the use of common random numbers to estimate the objective function and constraints values at different solutions. Assumption 3 imposes restrictions on the optimization problem (1), namely that there exist "enough" near-optimal points with respect to ε in the interior of the feasible region (e.g., objective function with an isolated optimal solution would violate Assumption 3). We show how Assumption 3 can be verified in Hu and Andradóttir (2014b). Now we present our main theorem in this paper; the detailed proof is also provided in Hu and Andradóttir (2014b).

Theorem 1 Suppose Assumptions 1, 2, and 3 hold. Choose $\{\lambda_i\}_{i=1}^\infty$ such that $\liminf_{i\to\infty} \lambda_i > \bar{f} - f^*$. Let $\delta_i = \Omega(i^{-\gamma_1})$ and $\xi_i = \Omega(i^{-\gamma_2})$, where $\gamma_2 > 0$. If $c(l-1) > 3$, $c(l-1) - 2\gamma_1 l > 3$, and $c(w-1) - 2\gamma_2 w > 3$, then $f(\theta_i^*) \to f^*$ and $\theta_i^* \in \Theta_{\mathscr{F}}$ almost surely (a.s.) as $i \to \infty$.

Note that Theorem 1 not only guarantees almost surely convergence (i.e., that $f(\theta_i^*) \to f^*$ a.s. as $i \to \infty$), but we also have $\theta_i^* \in \Theta_{\mathscr{F}}$ a.s. as $i \to \infty$. The latter ensures that the estimate of the optimal solution converges from inside the feasible region.

Sketch of the proof of Theorem 1: Let $\bar{A}$ denote the complement of any set A. Fix $0 < \varepsilon < \liminf_{i\to\infty}[\lambda_i - (\bar{f} - f^*)]/2$. To prove the algorithm converges from inside $\Theta_{\mathscr{F}}$ almost surely, it suffices to show that (a) $P(\Theta_i^* \cap \Theta_{\varepsilon/2,2\xi_i} = \emptyset, i.o.) = 0$ and (b) $P(\theta_i^* \in \bar{\Theta}_\varepsilon \cup \bar{\Theta}_{\mathscr{F}}, \Theta_i^* \cap \Theta_{\varepsilon/2,2\xi_i} \neq \emptyset, i.o.) = 0$. Note that (a) ensures that all near-optimal interior points are not discarded infinitely often, and (b) ensures that the algorithm does a good job with estimation so that the estimate θ_i^* of the optimal solution is selected well when near-optimal interior points are available.

4 NUMERICAL ANALYSIS

In this section, we conduct a numerical analysis aimed at investigating how stochastic constraints affect the performance of ASDP; refer to Hu and Andradóttir (2014b) for additional numerical examples.

More specifically, consider (1) with

$$f(\theta) = -\theta^2 + 100,$$
$$g(\theta) = \theta \leq b.$$

The feasible region is $\Theta = [-10, 10]$. The global maximum is 100 at $\theta = 0$, the range of $f(\theta)$ on Θ is $[0, 100]$, and we consider $b \in \{0, 5\}$. Let $h(\theta, X(\omega)) = f(\theta) + X(\omega)$ and $u(\theta, Y(\omega)) = g(\theta) + Y(\omega)$ for all $\theta \in \Theta$, with $X(\omega)$ being $\mathscr{N}(0, 100)$ and $Y(\omega)$ being $\mathscr{N}(0, 10)$. Here $\mathscr{N}(\mu, \sigma^2)$ denotes the normal distribution with mean μ and variance σ^2.

We use balanced exploration and exploitation to sample new points, as in Andradóttir and Prudius (2009). Explicitly, in iteration $k = M(i)$, with probability $p > 0$, a new solution is sampled uniformly from the whole feasible set Θ, and with probability $1 - p$, a new solution is sampled uniformly from $N(\theta_{i-1}^*)$, where

$$N(\theta) = \{\theta' \in \Theta : |\theta - \theta'| \leq r\}$$

for all $\theta \in \Theta$ (the first point is sampled uniformly from Θ). Here r denotes the radius of the "local" neighborhood. We accept every newly sampled point.

We use the same resampling strategy as Andradóttir and Prudius (2010) and Hu and Andradóttir (2014a). For any $x \in \mathbb{R}$, $\lceil x \rceil$ denotes the smallest integer not less than x and $\lfloor x \rfloor$ denotes the largest integer not greater than x. Let $M(i) = \lfloor i^v \rfloor$, and note that $m_k = \lfloor k^{1/v} \rfloor$ is the number of points sampled by the end of iteration k. Then, a point $\theta \in \Theta_{m_k}^*$ is resampled in iteration k with probability

$$p_k(\theta) = \frac{\exp\left(F_{m_k}(\theta)/T_{m_k}\right)}{\sum_{\theta' \in \Theta_{m_k}^*} \exp\left(F_{m_k}(\theta')/T_{m_k}\right)},$$

where $T_i = T/\log\left(M(i)+1\right)$ with $T > 0$. This resampling procedure puts more weight on the points that have better estimated objective function values.

Finally, we choose $\delta_i = \frac{D_1}{i^{\gamma_1}}$ and $\xi_i = \frac{D_2}{i^{\gamma_2}}$, where $\gamma_2 > 0$. Choose $\gamma_1 = \gamma_2 = 0.2$, $D_1 = 10$, and $D_2 = \sqrt{10}$. Let $T = 1$, $r = 0.2$, $v = 1.1$, $\lambda_i = i^p$, and $K(i) = \lceil i^c \rceil$, where $p = c = 0.5$.

Let $N_k = \sum_{\theta \in \tilde{\Theta}_{m_k}} N_k(\theta)$ be the total number of objective function evaluations by the end of iteration k, where $\tilde{\Theta}_{m_k}$ denotes the set of sampled points by the end of iteration k. Let $N = 10,000$ be the simulation budget. The performance of the ASDP algorithm is averaged over 100 independent replications. Its performance is documented by plotting 100 pairs (x, y), where $x \in \{0.01N, 0.02N, \ldots, N\}$ and y is the average objective function value at the estimated optimal solution after x objective function observations have been collected.

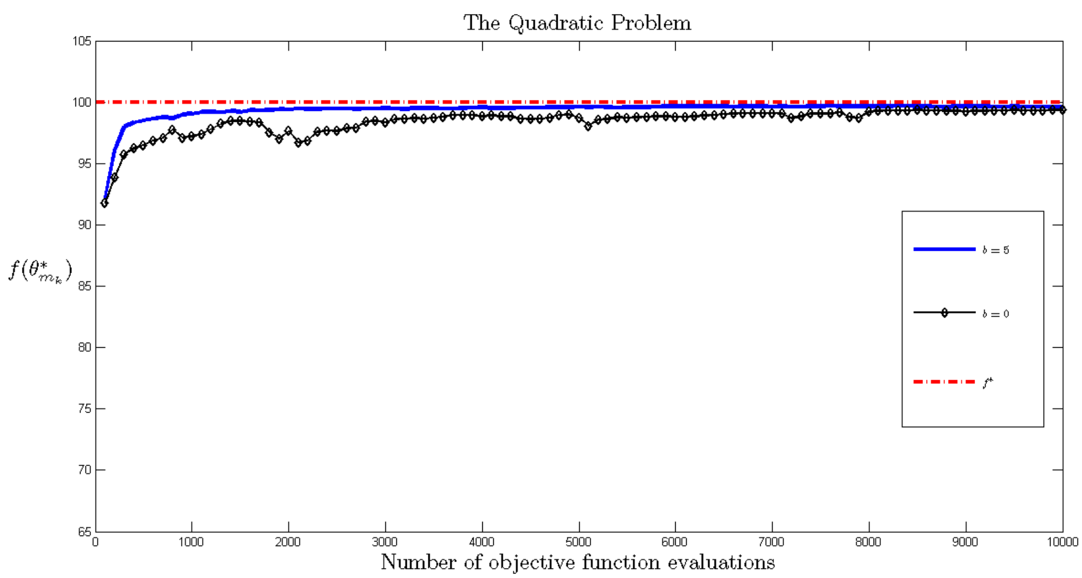

Figure 1: Performance of ASDP under different constraints.

Figure 1 shows the empirical performance of the ASDP method for $b \in \{0, 5\}$. The horizontal line labeled f^* denotes the optimal objective function value under $b \in \{0, 5\}$. From Figure 1, it is clear that the ASDP algorithm converges for both choices of b. Moreover, although the optimal solution is the same for $b \in \{0,5\}$, the convergence rate is slower for $b = 0$ due to the difficulty of ensuring convergence from inside the feasible region when the optimal solution is on the boundary of the feasible region.

5 CONCLUSION

In this paper, we propose and analyze a novel penalty function based random search algorithm, called Adaptive Search with Discarding and Penalization (ASDP), for continuous simulation optimization with stochastic constraints. We show that ASDP is an almost surely convergent algorithm, that the estimate of the optimal solution converges from inside the feasible region, and provide numerical results for a simple quadratic problem.

ACKNOWLEDGMENTS

We thank Xinyu Min for his assistance with developing a fast C++ code for our ASDP algorithm. The second author was supported by the National Science Foundation under grant CMMI-0856600.

REFERENCES

Andradóttir, S., and A. A. Prudius. 2009. "Balanced Explorative and Exploitative Search with Estimation for Simulation Optimization". *INFORMS Journal on Computing* 21:193–208.

Andradóttir, S., and A. A. Prudius. 2010. "Adaptive Random Search for Continuous Simulation Optimization". *Naval Research Logistics* 57:853–604.

Baumert, S., and R. L. Smith. 2002. "Pure Random Search for Noisy Objective Functions". *Technical Report 01-03, University of Michigan*.

Dentcheva, D., and A. Ruszczynski. 2003. "Optimization with Stochastic Dominance Constraints". *SIAM Journal on Optimization* 14:548–566.

Fu, M. C. 2002. "Optimization for Simulation: Theory vs. Practice". *INFORMS Journal on Computing* 14:192–215.

Hu, J., M. C. Fu, and S. I. Marcus. 2008. "A Model Reference Adaptive Search Method for Stochastic Global Optimization". *Communications in Information and Systems* 8 (3): 245–276.

Hu, L., and S. Andradóttir. 2014a. "Adaptive Random Search with Discarding for Continuous Simulation-Based Optimization". *Under revision*.

Hu, L., and S. Andradóttir. 2014b. "Simulation-Based Continuous Optimization with Stochastic Constraints". *Working paper*.

Hunter, S. R., and R. Pasupathy. 2013. "Optimal Sampling Laws for Stochastically Constrained Simulation Optimization on Finite Sets". *INFORMS Journal on Computing* 25:527–542.

Li, J., A. Sava, and X. Xie. 2009. "Simulation-Based Discrete Optimization of Stochastic Discrete Event Systems Subject to Non Closed-Form Constraints". *IEEE Transactions on Automatic Control* 54 (12): 2900–2904.

Norkin, V. I., G. C. Pflug, and A. Ruszczyński. 1998. "A Branch and Bound Method for Stochastic Global Optimization". *Mathematical Programming* 83:425–450.

Pagnoncelli, B. K., S. Ahmed, and A. Shapiro. 2009. "Sample Average Approximation Method for Chance Constrained Programming: Theory and Applications". *SIAM Journal on Optimization* 142 (2): 399–416.

Park, C., and S. Kim. 2011. "Handing Stochastic Constraints in Discrete Optimization via Simulation". In *Proceedings of the Winter Simulation Conference*, edited by S. Jain, R. Creasey, J. Himmelspach, K. White, and M. Fu, 4217–4226. Phoenix, Arizona.

Pujowidianto, N. A., L. H. Lee, C.-H. Chen, and C. M. Yap. 2009. "Optimal Computing Budget Allocation for Constrained Optimization". In *Proceedings of the Winter Simulation Conference*, edited by A. Dunkin, R. G. Ingalls, E. Yucesan, M. Rossetti, and B. J. R. Hill, 584–589. Austin, Texas.

Rubinstein, R. Y., and D. P. Kroese. 2004. *The Cross-Entropy Method*. Springer, New York, NY.

Yakowitz, A., and E. Lugosi. 1990. “Random Search in the Presence of Noise, with Application to Machine Learning”. *SIAM Journal on Scientific and Statistical Computing* 11 (4): 702–712.

AUTHOR BIOGRAPHIES

LIUJIA HU is a PhD candidate in the H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology. He received his M.S. in Operations Research from Columbia University, and B.A. in Economics and B.S. in Mathematics from Wuhan University in China. His e-mail address is lhu9@gatech.edu.

SIGRÚN ANDRADÓTTIR is a Professor in the H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology. Prior to joining Georgia Tech, she was an Assistant Professor and later an Associate Professor in the Departments of Industrial Engineering, Mathematics, and Computer Sciences at the University of Wisconsin – Madison. She received her Ph.D. in Operations Research from Stanford University in 1990. Her research interests include simulation, applied probability, and stochastic optimization. She is a member of INFORMS and served as Editor of the Proceedings of the 1997 Winter Simulation Conference. She was the Simulation Area Editor of *Operations Research Letters* from 2002 to 2008, and has served as Associate Editor for various journals. Her e-mail address is sa@gatech.edu, and her web page is http://www.isye.gatech.edu/faculty/sa.