

HANDLING STOCHASTIC CONSTRAINTS IN DISCRETE OPTIMIZATION VIA SIMULATION

Chuljin Park
Seong-Hee Kim

Georgia Institute of Technology
765 Ferst Dr NW
Atlanta, GA 30332, USA

ABSTRACT

We consider a discrete optimization via simulation problem with stochastic constraints on secondary performance measures where both objective and secondary performance measures need to be estimated by simulation. To solve the problem, we present a method called penalty function with memory (PFM), which determines a penalty value for a solution based on history of feasibility check on the solution. PFM converts a DOvS problem with stochastic constraints into a series of new optimization problems without stochastic constraints so that an existing DOvS algorithm can be applied to solve the new problem.

1. INTRODUCTION

Simulation techniques are useful to analyze performance of a complex or large-scale system with various random factors and often used in an optimization problem where performance measures in the problem cannot be analytically or numerically calculated. Optimization via Simulation (OvS) algorithms are to find an optimal or a good solution when the objective function needs to be evaluated through simulation. An OvS algorithm generates solutions from the search space, estimates the performance measures of each sampled solution, moves to a promising solution or region and repeats these steps until some stopping criteria are satisfied. Andradóttir (1998) and Fu (2006) provide a thorough review of OvS research and practice for both discrete and continuous decision variables. In this paper, we focus on discrete OvS (DOvS).

For a DOvS problem, a number of algorithms have been presented. The nested partitions (NP) method due to Shi and Ólafsson (2000) is one of globally convergent DOvS algorithms and works for both continuous and discrete decision variables. Pichitlamken and Nelson (2003) developed a DOvS algorithm using the NP framework. Their algorithm uses sample averages as estimates of performance measures and is proven to be globally convergent. The convergent optimization via most-promising-area stochastic search (COMPASS) (Hong and Nelson 2006) is a framework for discrete optimization to find the set of local optima. Although these algorithms are shown to perform well for many stochastic optimization problems, they are not directly applicable to a DOvS problem with stochastic constraints.

In deterministic optimization, the penalty function is one of popular methods to handle deterministic constraints. The penalty function consists of a penalty parameter and a measure of violation of the constraints. The penalty function is added to the objective function, which converts the original optimization problem with deterministic constraints into a series of unconstrained problems whose solutions converge to the solution of the original problem (Nocedal and Wright 1999). One may want to adopt the idea of the penalty function to a DOvS problem with stochastic constraints. Li et al. (2009) proposed an idea similar to the penalty function, in which the penalty parameter converges to infinity and the measure of violation of a constraint is estimated by sample means of secondary performance measures. Unfortunately, the method

works for only one stochastic constraint with a strict inequality.

In this paper, we propose a new method, namely the penalty function with memory (PFM) to handle multiple stochastic constraints with general inequalities in DOvS. A measure of violation of the constraints is determined by sample means as in Li et al. (2009). However, PFM differs from the method due to Li et al. (2009) in the sense that the penalty parameter in Li et al. (2009) diverges to infinity for all solutions while PFM determines the value of the penalty parameter based on history of feasibility check of a visited solution and either converges to 0 for a feasible solution or diverges to ∞ (if minimization) for an infeasible solution as the number of observations obtained for the solution goes to infinity.

This paper is organized as follows: Section 2 defines our problem, details notations, and provides a common structure of existing DOvS algorithms. Section 3 presents a general form of PFM and discusses convergence properties of the method. Section 4 gives an example algorithm that combines a version of NP with PFM. Experimental results are discussed in Section 5, followed by concluding remarks in Section 6.

2. BACKGROUND

In this section, we define problem and detail notation and assumptions used throughout the paper. We also review a common structure of existing DOvS algorithms.

2.1 Problem

Let $\mathbf{x} = (x_1, \dots, x_d)$ represent a solution and Θ represent the whole decision variable space which is a bounded and finite set in \mathbb{R}^d . Let $G_i(\mathbf{x})$ represent the i th observation corresponding to the primary performance measure. Similarly, for all $\ell = 1, 2, \dots, m$, $H_{\ell i}(\mathbf{x})$ represents the i th observation corresponding to the secondary performance measures on the ℓ th constraint. We assume that for a given \mathbf{x} , $G_i(\mathbf{x})$ are independent and identically distributed (iid) random variables for $i = 1, 2, \dots$ and observed only through simulation. For $\ell = 1, 2, \dots, m$, $H_{\ell i}(\mathbf{x})$ are also iid and must be generated via simulation. It is possible that $G_i(\mathbf{x})$ and $H_{\ell i}(\mathbf{x})$ are correlated. The expectation and variance of the primary performance measure are $\mathbf{E}[G_i(\mathbf{x})]$ and $\mathbf{Var}[G_i(\mathbf{x})]$, respectively. Also, for all $\ell = 1, 2, \dots, m$, the expectation and variance of the ℓ th secondary performance measure are $\mathbf{E}[H_{\ell i}]$ and $\mathbf{Var}[H_{\ell i}(\mathbf{x})]$, respectively. We assume that the following assumption holds for a DOvS problem in consideration.

Assumption 1 For all $\mathbf{x} \in \Theta$ and $\ell = 1, 2, \dots, m$,

$$\mathbf{E}[|G_i(\mathbf{x})|] < \infty, \mathbf{Var}[G_i(\mathbf{x})] < \infty, \mathbf{E}[|H_{\ell i}(\mathbf{x})|] < \infty \text{ and } \mathbf{Var}[H_{\ell i}(\mathbf{x})] < \infty.$$

Then our DOvS problem with stochastic constraints is defined as follows:

$$\begin{aligned} & \operatorname{argmin}_{\mathbf{x} \in \Theta} \mathbf{E}[G_i(\mathbf{x})], \\ & \text{subject to } \mathbf{E}[H_{\ell i}(\mathbf{x})] \geq q_{\ell}, \ell = 1, 2, \dots, m. \end{aligned} \tag{1}$$

2.2 Existing DOvS Algorithms

Existing DOvS algorithms share a general structure of the random search framework (Andradóttir 2006) as follows:

- Step 1: Set search iteration counter $k = 1$. Choose an initial sampling strategy.
- Step 2: Sample solutions and take Δn number of additional observations for each sampled solution.
- Step 3: Update estimates of performance measures.

Step 4: Update the sampling strategy for next iteration, set $k \leftarrow k + 1$, and repeat Steps 2 and 3 until stopping criteria are satisfied.

In general, it is required that the sampling strategy guarantees that each solution has non-zero probability of being sampled at any iteration k to ensure that a solution receives an infinite number of observations as k goes to infinity. PFM is designed for an algorithm that follows this structure.

Additional notation is defined below:

\mathcal{D} := an existing DOvS algorithm designed for a DOvS problem without any stochastic constraint;
 x_o^b := the solution to Problem 1;
 $v_k(\mathbf{x})$:= the number of visits up to iteration k for \mathbf{x} ;
 $n_r(\mathbf{x})$:= the total number of observations obtained up to the r th visit for \mathbf{x} ;
 $n_{v_k}(\mathbf{x}) := n_{v_k(\mathbf{x})}(\mathbf{x})$:= the total number of observations obtained up to iteration k for \mathbf{x} ;
 $\lambda_\ell^r(\mathbf{x})$:= penalty parameter of the ℓ th constraint at the r th visit for \mathbf{x} ;
 $\lambda_\ell^{v_k}(\mathbf{x}) := \lambda_\ell^{v_k(\mathbf{x})}(\mathbf{x})$:= penalty parameter of the ℓ th constraint up to iteration k for \mathbf{x} ;
 \mathcal{F} := an index set of all stochastic constraints such as $\mathcal{F} := \{1, 2, \dots, m\}$;
 $\mathcal{F}_{S(\mathbf{x})} := \{\ell \mid \mathbf{E}[H_\ell(\mathbf{x})] > q_\ell, \ell \in \mathcal{F}\}$, an index set of stochastic constraints for which \mathbf{x} is strictly feasible;
 $\mathcal{F}_{A(\mathbf{x})} := \{\ell \mid \mathbf{E}[H_\ell(\mathbf{x})] = q_\ell, \ell \in \mathcal{F}\}$, an index set of active stochastic constraints for \mathbf{x} ;
 $\mathcal{F}_{I(\mathbf{x})} := \{\ell \mid \mathbf{E}[H_\ell(\mathbf{x})] < q_\ell, \ell \in \mathcal{F}\}$, an index set of stochastic constraints for which \mathbf{x} is infeasible;
 $\bar{G}_k(\mathbf{x}) := \frac{1}{n_{v_k}(\mathbf{x})} \sum_{i=1}^{n_{v_k}(\mathbf{x})} G_i(\mathbf{x})$, cumulative sample mean of the objective up to iteration k ;
 $\bar{H}_{\ell k}(\mathbf{x}) := \frac{1}{n_{v_k}(\mathbf{x})} \sum_{i=1}^{n_{v_k}(\mathbf{x})} H_{\ell i}(\mathbf{x})$, $\ell = 1, 2, \dots, m$, cumulative sample mean of the ℓ th secondary performance measure up to iteration k .

Algorithm \mathcal{D} will be combined with PFM to solve Problem 1. In this paper, we assume that \mathcal{D} satisfies the following assumption:

Assumption 2 The sampling strategy of \mathcal{D} guarantees that for any $\mathbf{x} \in \Theta$,

$$\mathbf{P}[\lim_{k \rightarrow \infty} v_k(\mathbf{x}) = \infty] = 1 \quad \text{and} \quad \mathbf{P}[\lim_{k \rightarrow \infty} n_{v_k}(\mathbf{x}) = \infty] = 1.$$

NP-based algorithms due to Pichitlamken and Nelson (2003) and random search methods due to (Andradóttir 1999) are example algorithms that satisfy Assumption 2.

3. PENALTY FUNCTION WITH MEMORY

In this section, we present a new method to handle stochastic constraints, namely PFM. The method consists of the penalty parameter of the ℓ th constraint, $\lambda_\ell^{v_k}(\mathbf{x})$, and a measure of violation of the ℓ th constraint, $\max(0, q_\ell - \bar{H}_{\ell k}(\mathbf{x}))$. The penalty parameter throughout iteration k becomes a sequence of real-valued numbers which we call a penalty sequence. The sequences $\lambda_\ell^{v_k}(\mathbf{x})$ should satisfy the following assumption.

Assumption 3 If $\ell \in \mathcal{F}_{S(\mathbf{x})}$, $\mathbf{P}[\lim_{k \rightarrow \infty} \lambda_\ell^{v_k}(\mathbf{x}) = 0] = 1$.
 Else if $\ell \in \mathcal{F}_{A(\mathbf{x})}$, $\lim_{k \rightarrow \infty} \mathbf{P}[\lambda_\ell^{v_k}(\mathbf{x}) = 0] = 1$.
 Else, $\mathbf{P}[\lim_{k \rightarrow \infty} \lambda_\ell^{v_k}(\mathbf{x}) = \infty] = 1$.

Assumption 3 implies that the penalty sequence for the ℓ th constraint converges to 0 with probability 1 (w.p. 1) if \mathbf{x} is strictly feasible with respect to the constraint but diverges to ∞ w.p. 1 if \mathbf{x} is infeasible for the constraint. If \mathbf{x} is active for the constraint, the sequence converges to 0 in probability. For an

active constraint, there is always non zero probability that $\bar{H}_{\ell k}(\mathbf{x}) < q_\ell$. Thus the convergence occurs only in probability rather than w.p. 1.

With the penalty sequences that satisfy Assumption 3, a new objective function with PFM at search iteration k , $Z_k(\mathbf{x})$, is defined as

$$Z_k(\mathbf{x}) = \bar{G}_k(\mathbf{x}) + \sum_{\ell \in \mathcal{F}} \left[\lambda_\ell^{v_k}(\mathbf{x}) \times \max\{0, q_\ell - \bar{H}_{\ell k}(\mathbf{x})\} \right]. \quad (2)$$

We apply \mathcal{D} to solve $\min_{\mathbf{x} \in \Theta} \mathbf{E}[Z_k(\mathbf{x})]$. As $k \rightarrow \infty$, we want to guarantee that $\lim_{k \rightarrow \infty} \mathbf{E}[Z_k(\mathbf{x}_k^b)] = \mathbf{E}[G_i(\mathbf{x}_o^b)]$ where \mathbf{x}_k^b is the solution to $\min_{\mathbf{x} \in \Theta} \mathbf{E}[Z_k(\mathbf{x})]$ and \mathbf{x}_o^b is the solution to Problem 1. If \mathcal{D} satisfies Assumption 2, $n_{v_k}(\mathbf{x})$ goes to infinity as k goes to infinity and $\bar{G}_k(\mathbf{x})$ and $\bar{H}_{\ell k}(\mathbf{x})$ converge to $\mathbf{E}[G_i(\mathbf{x})]$ and $\mathbf{E}[H_{\ell i}(\mathbf{x})]$, respectively, by the strong law of large numbers (SLLN). Then it is clear that $Z_k(\mathbf{x})$ converges to $\mathbf{E}[G_i(\mathbf{x})]$ for a feasible solution but diverges to infinity for an infeasible solution if Assumption 3 holds. The following theorem is proved in Park and Kim (2011).

Theorem 1 Suppose that Assumptions 1, 2, and 3 hold. Then if \mathbf{x} is strictly feasible (i.e., $\mathcal{F}_{S(\mathbf{x})} = \mathcal{F}$), $\mathbf{P}[\lim_{k \rightarrow \infty} Z_k(\mathbf{x}) = \mathbf{E}[G_i(\mathbf{x})]] = 1$.
 Else if \mathbf{x} is feasible with any active constraint (i.e., $\mathcal{F}_{A(\mathbf{x})} \neq \emptyset$ and $\mathcal{F}_{I(\mathbf{x})} = \emptyset$), $\lim_{k \rightarrow \infty} \mathbf{P}[Z_k(\mathbf{x}) = \mathbf{E}[G_i(\mathbf{x})]] = 1$.
 Else, $\mathbf{P}[\lim_{k \rightarrow \infty} Z_k(\mathbf{x}) = \infty] = 1$.

For Theorem 1, it is important that the penalty sequences $\lambda_\ell^{v_k}(\mathbf{x})$ satisfy Assumption 3. We give an example of such penalty sequences, namely multiple penalty sequence with constants (PS_m):

$$\lambda_\ell^{v_k}(\mathbf{x}) = \begin{cases} \lambda_\ell^{v_{k-1}}(\mathbf{x}) \times \theta_a, & \text{if } \bar{H}_{\ell k}(\mathbf{x}) < q_\ell; \\ \lambda_\ell^{v_{k-1}}(\mathbf{x}) \times \theta_d, & \text{if } \bar{H}_{\ell k}(\mathbf{x}) \geq q_\ell, \end{cases}$$

where $\lambda_\ell^0(\mathbf{x}) = \lambda_\ell^0$ is an initial penalty constant for constraint ℓ and θ_a is an appreciation factor and θ_d is a depreciation factor such that $\theta_a > 1$, $0 < \theta_d < 1$ and $\theta_a \theta_d < 1$.

Theorem 2 Under Assumptions 1 and 2, PS_m satisfies Assumption 3.

We explain intuitively why Theorem 2 should hold. See Park and Kim (2011) for its formal proof. If there is any evidence that \mathbf{x} is infeasible with respect to constraint ℓ at iteration k (i.e., $\bar{H}_{\ell k}(\mathbf{x}) < q_\ell$ for any ℓ), then PS_m multiplies θ_a , a constant greater than 1, to the previous penalty parameter of the constraint. Otherwise, θ_d , a constant between 0 and 1, is multiplied. As iteration k grows, SLLN applies and PS_m keeps multiplying θ_a to the penalty parameter of any infeasible constraint. As results, the penalty sequence of the infeasible constraint diverges to infinity as $k \rightarrow \infty$. Similarly, for any strictly feasible constraint of \mathbf{x} , the penalty sequence of the constraint keeps receiving the depreciation factor θ_d as $k \rightarrow \infty$ and thus converges to 0. For any active constraint, PS_m multiplies θ_a or θ_d with equal probability and the penalty sequence converges to 0 in probability as $k \rightarrow \infty$ because $\theta_a \theta_d < 1$.

4. EXAMPLE ALGORITHM

In this section, we take a version of NP due to Pichtlamken and Nelson (2003) as a DOvS algorithm and combine it with PFM that uses PS_m . NP focuses on \mathcal{R}_k called the most promising region and spends more computational efforts in \mathcal{R}_k . A complement set of \mathcal{R}_k , $\Theta \setminus \mathcal{R}_k$ is called the surrounding region. NP systematically partitions \mathcal{R}_k into several subregions. Let ω represent the number of subregions. Then

$\mathcal{R}_k(j)$, $j = 1, 2, \dots, \omega$, are mutually disjoint subsets of \mathcal{R}_k such that the union of $\mathcal{R}_k(j)$, for $j = 1, 2, \dots, \omega$, is equal to \mathcal{R}_k .

NP samples and assesses solutions from $\mathcal{R}_k(j)$ for $j = 1, 2, \dots, \omega$ and $\Theta \setminus \mathcal{R}_k$ and compares all solutions visited so far. Let $\tau_k(j)$, $j = 1, 2, \dots, \omega$, be the number of sampled solutions from subregion j at iteration k and τ_k be the total number of sampled solutions at iteration k (i.e., $\sum_{j=1}^{\omega} \tau_k(j) = \tau_k$, if $\mathcal{R}_k = \Theta$). If a sampled solution is visited for the first time, n_0 observations are taken and otherwise, we take Δn observations for the solution. Let \mathcal{S}_k define the set of solutions sampled at iteration k and \mathcal{V}_k define the set of all solutions visited up to iteration k . The sample best among solutions in \mathcal{V}_k is denoted as $\hat{\mathbf{x}}_k^*$. If $\hat{\mathbf{x}}_k^*$ is in one of $\mathcal{R}_k(j)$, $j = 1, 2, \dots, \omega$, the subregion will be the next promising region, \mathcal{R}_{k+1} . Otherwise, Θ will be set to \mathcal{R}_{k+1} . The steps of NP+PFM are provided below.

Algorithm : NP + PFM

Step 0. Initialization:

- Choose a sampling strategy and a partitioning scheme.
- Set $k = 1$, $\mathcal{R}_k = \Theta$, and $\mathcal{V}_0 = \emptyset$.
- Sample an initial solution, $\hat{\mathbf{x}}_0^*$ randomly from Θ .
- Select constants ω , τ_k , Δn , $\lambda_\rho^0(\mathbf{x})$, θ_a and θ_d .

Step 1. Partitioning:

- Partition \mathcal{R}_k into ω disjoint subregions, $\mathcal{R}_k(1), \mathcal{R}_k(2), \dots, \mathcal{R}_k(\omega)$. If \mathcal{R}_k is a singleton, set $\mathcal{R}_k(1) = \mathcal{R}_k$ and $\mathcal{R}_k(2) = \dots = \mathcal{R}_k(\omega) = \emptyset$.
- Set $\mathcal{R}_k(\omega + 1) = \Theta \setminus \mathcal{R}_k$ which denotes the surrounding region.

Step 2. Sampling Solutions:

- From each region $\mathcal{R}_k(j)$, $j = 1, 2, \dots, \omega + 1$, sample $\tau_k(j)$ solutions using the chosen sampling strategy. Always sample $\hat{\mathbf{x}}_{k-1}^*$ so that $\hat{\mathbf{x}}_{k-1}^* \in \mathcal{S}_k$.
- Include all sampled solutions \mathbf{x} into \mathcal{S}_k .
- If $\mathbf{x} \notin \mathcal{V}_k$ for any $\mathbf{x} \in \mathcal{S}_k$, then $\mathcal{V}_k = \mathcal{V}_k \cup \{\mathbf{x}\}$.

Step 3. Estimating the Promising Index: For each $\mathbf{x} \in \mathcal{S}_k$,

- If \mathbf{x} has never been visited before, take n_0 observations and set $n_1(\mathbf{x}) = n_0$. Otherwise, take Δn observations and set $n_{v_k}(\mathbf{x}) = n_{v_{k-1}}(\mathbf{x}) + \Delta n$.
- Update $Z_k(\mathbf{x})$ as defined in (2).
- Select $\hat{\mathbf{x}}_k^*$ such that $\hat{\mathbf{x}}_k^* \equiv \operatorname{argmin}_{\mathbf{x} \in \mathcal{V}_k} Z_k(\mathbf{x})$.

Step 4. Selecting the Most Promising Region and Backtracking:

- Determine j^* such that $\hat{\mathbf{x}}_k^* \in \mathcal{R}_k(j^*)$.
- If $\mathcal{R}_k(j^*) \subset \mathcal{R}_k$, then $\mathcal{R}_{k+1} = \mathcal{R}_k(j^*)$. Otherwise, $\mathcal{R}_{k+1} = \Theta$.
- Set $k = k + 1$.

Step 5. Stopping Rule If the stopping rule is satisfied, then stop and return $\hat{\mathbf{x}}_k^*$ as the best solution. Otherwise repeat Steps 1 through 5.

For various sampling strategies and partitioning schemes, see Pichitlamken and Nelson (2003) or Shi and Ólafsson (2008). The global convergence is achieved when k goes to infinity but, in practice, the algorithm should terminate with finite search iterations. Popular stopping criteria discussed in Hong and Nelson (2006) include (i) when all computational budget is consumed or a certain number of search iterations is made, (ii) all solutions are visited so that the algorithm had enough search iterations or (iii) the identity of the sample best $\hat{\mathbf{x}}_k^*$ does not change for many iterations. Another example of stopping

criteria is to stop when event E_1 occurs for n_E times consecutively, where

$$E_1 := \{\hat{\mathbf{x}}_k^* = \hat{\mathbf{x}}_{k-1}^*, |Z_k(\hat{\mathbf{x}}_k^*) - Z_{k-1}(\hat{\mathbf{x}}_{k-1}^*)| < \epsilon, \mathcal{R}(k) \text{ is a singleton}\}$$

for a small positive constant ϵ . A decision maker needs to choose ϵ and n_E .

5. EXPERIMENTS

In this section, we test our method NP+PFM on a DOvS problem based on the Goldstein-Price function (Goldstein and Price 1971) and compare the performance of NP+PFM with a method due to Li et al. (2009).

The Goldstein-Price function is one of the famous deterministic and continuous optimization problems with a 2-dimensional quadratic function defined by,

$$g(\mathbf{x}) = \left\{ 1 + (x_1 + x_2 + 1)^2 \cdot (19 - 14x_1 + 3x_1^2 - 14x_2 + 6x_1x_2 + 3x_2^2) \right\} \\ \times \left\{ 30 + (2x_1 - 3x_2)^2 \cdot (18 - 32x_1 + 12x_1^2 + 48x_2 - 36x_1x_2 + 27x_2^2) \right\}.$$

Let $\phi_i(\mathbf{x})$ and $\psi_{\ell i}(\mathbf{x})$, $\ell = 1, 2, \dots, m$ be iid normal random variables with mean zero and standard deviations $\gamma(\mathbf{x})$ and $\sigma_{\ell}(\mathbf{x})$. We define $G_i(\mathbf{x}) = g(\mathbf{x}) + \phi_i$ and $H_{\ell i}(\mathbf{x}) = a_{\ell}x_1 + b_{\ell}x_2 + \psi_{\ell i}$ and want to minimize $\mathbf{E}[G_i(\mathbf{x})]$ with constraints $\mathbf{E}[H_{\ell i}(\mathbf{x})] \geq q_{\ell}$. We set $\Theta = \{-2.50, -2.49, \dots, 1.99, 2.00\}^2$ which is a two-dimensional discretized set in $[-2.50, 2.00]^2$. The function $g(\mathbf{x})$ has four local minima and the global minimum at $(0, -1)$. In Θ , the largest and smallest values of $g(\mathbf{x})$ are 1,015,685 and 3, respectively. We take $\gamma(\mathbf{x}) = 0.15g(\mathbf{x})$ and $\sigma_{\ell}(\mathbf{x}) = 0.15(a_{\ell}x_1 + b_{\ell}x_2)$.

For the implementation of NP, we use the bisection scheme for partitioning, which bisects the possible range of each variable x_q , $q = 1, 2, \dots, d$ of \mathbf{x} in \mathcal{R}_k . For example, if $\mathcal{R}_k = [0, 2]^2$, the bisection method would create four subregions: $[0, 1] \times [0, 1]$, $[1, 2] \times [0, 1]$, $[0, 1] \times [1, 2]$ and $[1, 2] \times [1, 2]$. We set $n_0 = 10$, $\Delta n = 10$, $\tau_k = 16$, and $\omega = 4$. For sampling solutions, we use the following sampling strategy: If $\mathcal{R}_k = \Theta$, $\tau_k(j) = 4$ for all $j = 1, 2, 3, 4$. Else if $|\mathcal{R}_k| > 1$, $\tau_k(j) = \min(|\mathcal{R}_k|, 3)$ for all $j = 1, 2, 3, 4$ and $\tau_k(5) = \tau_k - \sum_{j=1}^4 \tau_k(j)$. Else, $\tau_k(1) = 1$, $\tau_k(j) = 0$, for $j = 2, 3, 4$ and $\tau_k(5) = \tau_k - 1$. If $|\mathcal{R}_k| \neq 1$, we employ the uniform sampling, which generates $\tau_k(j)$ solutions randomly from $R_k(j)$ for each j . Otherwise (i.e., R_k is a singleton), we randomly generate seven solutions among eight perturbed solutions from the solution in \mathcal{R}_k , which is called the local search sampling (Shi and Ólafsson 2008), and use the uniform sampling to generate eight solutions randomly in the surrounding region.

Li et al. (2009) provide a penalty function for a DOvS problem with one stochastic constraint, which we call augmented cost function (ACF). Their penalty function is defined as

$$\sum_{\ell \in \mathcal{F}} \alpha_{\ell}^k \times \max\{0, q_{\ell} - \bar{H}_{\ell k}(\mathbf{x})\},$$

which is a straightforward extension to multiple constraints. We take α_{ℓ}^k similar to the one used in their online companion:

$$\alpha_{\ell}^k = \begin{cases} \frac{e^k}{\min_{\mathbf{x} \in \Upsilon_k} (q_{\ell} - \bar{H}_{\ell k}(\mathbf{x}))}, & \text{if } \Upsilon_k \neq \emptyset; \\ 10^6, & \text{otherwise,} \end{cases}$$

where $\Upsilon_k \equiv \{\mathbf{x} | q_{\ell} > \bar{H}_{\ell k}(\mathbf{x}) \text{ and } \mathbf{x} \in \mathcal{V}_k\}$. For NP+ACF, we use $Z'_k(\mathbf{x}) = \bar{G}_k(\mathbf{x}) + \sum_{\ell \in \mathcal{F}} \alpha_{\ell}^k \times \max\{0, q_{\ell} - \bar{H}_{\ell k}(\mathbf{x})\}$ as a new objective function.

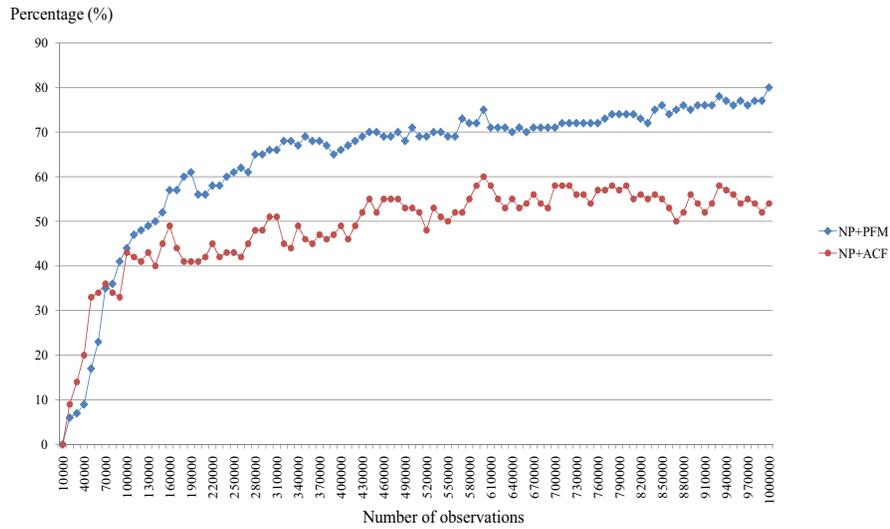


Figure 1: Percentage (%) of time that $\hat{\mathbf{x}}_k^* = \mathbf{x}_o^b$ under a single constraint.

For PFM, we employ PS_m and set $\theta_a = \sqrt{2}$, $\theta_d = \frac{1}{2\theta_a}$, and $\lambda_\ell^0(\mathbf{x}) = 10^6$. See Park and Kim (2011) for a discussion on parameter choices in PS_m .

We first consider a single stochastic constraint and then add one more constraint. We make 100 macro replications and each macro replication stops when the total number of observations reaches a million. We report (i) percentage of time that the sample best at search iteration k , $\hat{\mathbf{x}}_k^*$ was equal to the true best optimal \mathbf{x}_o^b and (ii) average estimated objective value $\bar{G}_k(\hat{\mathbf{x}}_k^*)$ over 100 macro replications.

5.1 Single Stochastic Constraint

We set the stochastic constraint as

$$\mathbf{E}[-x_1 - x_2 + \psi_{1i}] \geq 1.5.$$

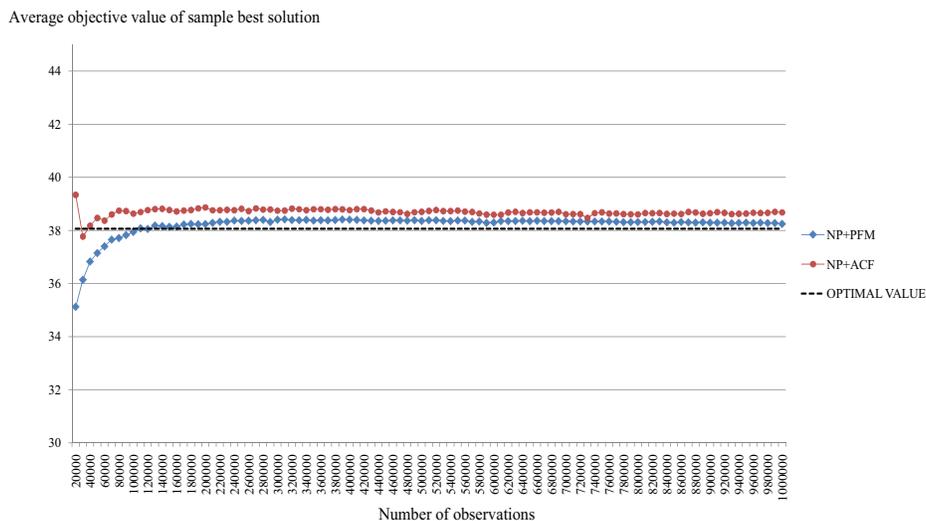


Figure 2: Average estimated objective value of $\hat{\mathbf{x}}_k^*$ under a single constraint.

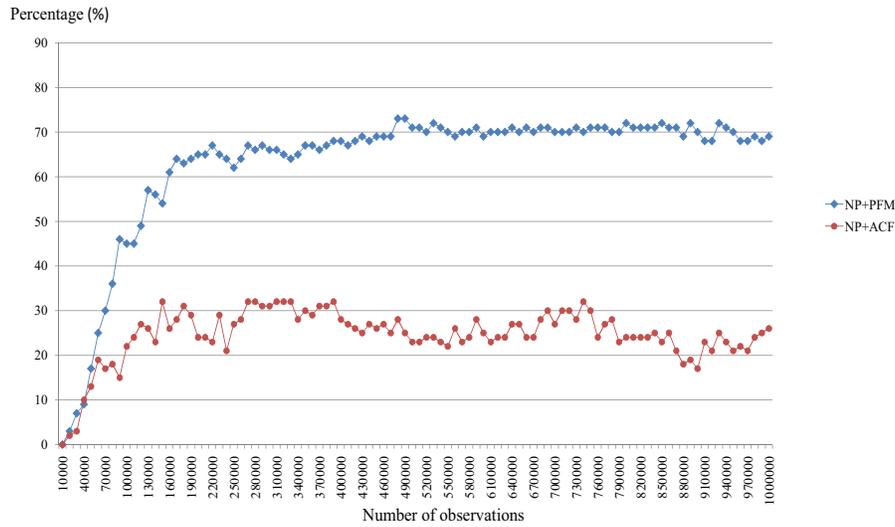


Figure 3: Percentage (%) of time that $\hat{\mathbf{x}}_k^* = \mathbf{x}_o^b$ under two constraints.

With this constraint, the true optimal solution (\mathbf{x}_o^b) is $(-0.3, -1.2)$ and the stochastic constraint becomes active at \mathbf{x}_o^b . The true objective value at \mathbf{x}_o^b is exactly $g(\mathbf{x}_o^b) = 38.0625$. Figure 1 shows the percentage of time that $\hat{\mathbf{x}}_k^* = \mathbf{x}_o^b$ over 100 macro replications. The percentage in NP+PFM goes up to 80% while NP+ACF achieves only up to 55%. As shown in Figure 2, NP+PFM obtains better average estimated objective values than NP+ACF.

5.2 Two Stochastic Constraints

Now we consider two stochastic constraints:

$$\mathbf{E}[-x_1 - x_2 + \psi_{1i}] \geq 1.5, \quad \text{and} \quad \mathbf{E}[x_1 - x_2 + \psi_{2i}] \geq 0.9.$$

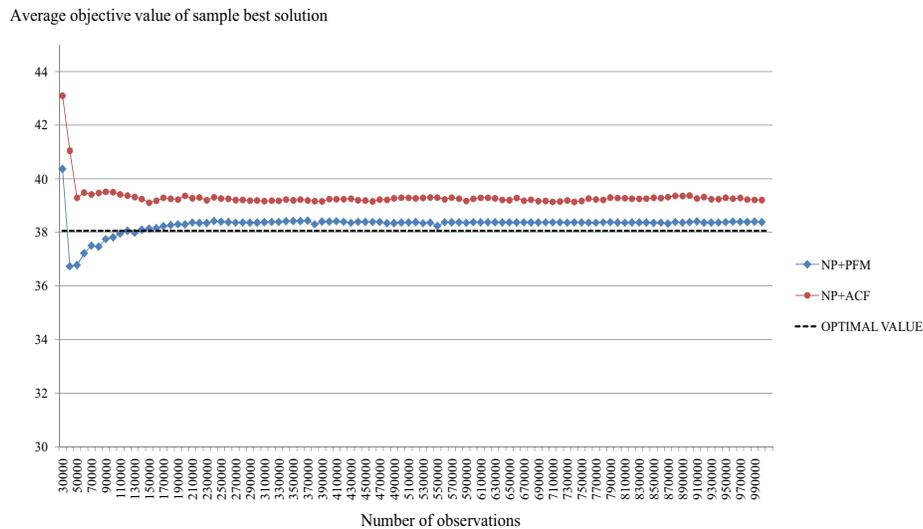


Figure 4: Average objective value of $\hat{\mathbf{x}}_k^*$ under two stochastic constraints.

The true optimal feasible solution \mathbf{x}_o^b is $(-0.3, -1.2)$ which has both constraints as active constraints. Its objective function $g(\mathbf{x}_o^b)$ is 38.0625. Figure 3 shows that NP+PFM returns the true best up to 70% of time as the number of observations increases while NP+ACF shows only 20% ~ 30% convergence. The convergence percentage is slightly lower than the single constraint case but it is expected because this problem with two stochastic constraints is more difficult. NP+PFM still gives a better estimated objective value than NP+ACF as shown in Figure 4.

6. CONCLUSION

In this paper, we present PFM that replaces a DOvS problem with stochastic constraints into a new DOvS problem without stochastic constraints. PFM determines penalty based on history of feasibility check for each solution. We discuss convergence properties of PFM and an example of penalty sequences for PFM. Our experimental results show that the method works well. There exist more sophisticated statistical methods for feasibility check and selection of the most promising region as presented in, for example, Batur and Kim (2010) and Andradóttir and Kim (2010). PFM can be combined with them to improve better performance, which is a topic of current ongoing work.

ACKNOWLEDGMENTS

This work is supported by NSF grant number CMMI-0644837.

REFERENCES

- Andradóttir, S. 1998. "Simulation optimization". In *Handbook of Simulation*, edited by J. Banks. New York: Wiley.
- Andradóttir, S. 1999. "Accelerating the convergence of random search methods for discrete stochastic optimization". *ACM Transactions on Modeling and Computer Simulation* 9:349–380.
- Andradóttir, S. 2006. "An overview of simulation optimization via random search". In *Handbooks in Operations Research and Management Science: Simulation*, edited by S. G. Henderson and B. L. Nelson. Oxford: Elsevier Science.
- Andradóttir, S., and S.-H. Kim. 2010. "Fully sequential procedures for comparing constrained systems via simulation". *Naval Research Logistics* 57 (5): 403–421.
- Batur, D., and S.-H. Kim. 2010. "Finding feasible systems in the presence of constraints on multiple performance measures". *ACM Transactions on Modeling and Computer Simulation* 20 (13): 1–26.
- Fu, M. C. 2006. "Optimization for simulation: Theory vs. practice". *Operations Research* 14:192–215.
- Goldstein, A. A., and J. F. Price. 1971. "On descent from local minima". *Mathematics of Computation* 25 (115): 569–574.
- Hong, L., and B. L. Nelson. 2006. "Discrete optimization via simulation using COMPASS". *Operations Research* 54:115–129.
- Li, J., S. 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* 54:2900–2904.
- Nocedal, J., and S. Wright. 1999. *Numerical optimization*. New York: Springer.
- Park, C., and S.-H. Kim. 2011. "Penalty function with memory for discrete optimization via simulation with stochastic constraints". Technical report, H. Milton Stewart School of Industrial Systems and Engineering, Georgia Institute of Technology, Atlanta, Georgia.
- Pichitlamken, J., and B. L. Nelson. 2003. "A combined procedure for optimization via simulation". *ACM Transactions on Modeling and Computer Simulation* 13:155–179.
- Shi, L., and S. Ólafsson. 2000. "Nested partitions method for stochastic optimization". *Methodology and Computing in Applied Probability* 2:271–291.
- Shi, L., and S. Ólafsson. 2008. *Nested Partitions Method, Theory and Applications*. New York: Springer.

AUTHOR BIOGRAPHIES

CHULJIN PARK is a Ph.D. student in H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology. He received his B.S. in Mechanical Engineering from Korea Advanced Institute of Science and Technology (KAIST) in 2007. His research interests include optimization via simulation with stochastic constraints and its applications to environmental management. His e-mail and web addresses are cpark41@gatech.edu and <http://www.isye.gatech.edu/~cpark41>, respectively.

SEONG-HEE KIM is an Associate Professor in H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology. She received her Ph.D. in Industrial Engineering and Management Sciences from Northwestern University in 2001. Her research interests include ranking and selection, optimization via simulation, quality control, simulation output analysis, and applications of simulation methods to environmental management. She is a member of INFORMS, and serves as an associate editor in the simulation area of *Operations Research* and the OR/simulation area of *The American Statistician*. Her e-mail and web addresses are skim@isye.gatech.edu and <http://www.isye.gatech.edu/~skim>, respectively.