Proceedings of the 2012 Winter Simulation Conference
C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, and A. M. Uhrmacher, eds.

CLOSED-FORM SAMPLING LAWS FOR STOCHASTICALLY CONSTRAINED SIMULATION OPTIMIZATION ON LARGE FINITE SETS

Nugroho Artadi Pujowidianto

Industrial and Systems Engineering
National University of Singapore, SINGAPORE

Susan R. Hunter

Operations Research and Information Engineering
Cornell University, Ithaca, NY, USA

Raghu Pasupathy

Industrial and Systems Engineering
Virginia Tech, Blacksburg, VA, USA

Loo Hay Lee

Industrial and Systems Engineering
National University of Singapore, SINGAPORE

Chun-Hung Chen

Department of Electrical Engineering
National Taiwan University, Taipei, TAIWAN

ABSTRACT

Consider the context of constrained simulation optimization (SO), that is, optimization problems where the objective function and constraints are known through a Monte Carlo simulation, with corresponding estimators possibly dependent. We identify the nature of sampling plans that characterize efficient algorithms, particularly in large countable spaces. We show that in a certain asymptotic sense, the optimal sampling characterization, that is, the sampling budget for each system that guarantees optimal convergence rates, depends on a single easily estimable quantity called the score. This result provides a useful and easily implementable sampling allocation that approximates the optimal allocation, which is otherwise intractable due to it being the solution to a difficult bilevel optimization problem. Our results point to a simple sequential algorithm for efficiently solving large-scale constrained simulation optimization problems on finite sets.

1 INTRODUCTION

The constrained Simulation Optimization (SO) problem is a nonlinear optimization problem in which the objective and constraint functions can only be observed with error through a Monte Carlo simulation. The functions involved in the optimization are implicitly expressed through the simulation, as opposed to more traditional settings in which an explicit expression of the functions is required. The flexibility of incorporating complex stochastic function structures into optimization problems has resulted in the widespread adoption of SO formulations and consequent attention among researchers.

SO problems are broadly categorized by the nature of the feasible region and the type of solution required. Depending on the nature of the feasible region, SO problems are generally considered either categorical, integer-ordered, or continuous, with problems falling in more than one of these categories being considered mixed SO problems. SO problems in each of the integer-ordered and continuous categories can require global or local solutions. For examples of SO problems in each of these categories, visit the library of SO problems at www.simopt.org (Pasupathy and Henderson 2006; Pasupathy and Henderson 2011).

In this paper, we consider the stochastically constrained SO problem on categorical or finite spaces. This variation involves identifying the best system (or design) from a finite population of systems, as

978-1-4673-4781-5/12/$31.00 ©2012 IEEE

measured by an estimable objective function, from among those systems that are feasible, as measured by a set of estimable constraint functions. This formulation subsumes the unconstrained version of the categorical SO problem. Unlike the unconstrained version, research on the constrained version is still in its infancy. Attempts at solution have been relatively few and very recent; entry points to work in this area include Andradóttir and Kim (2010), Hunter and Pasupathy (2012), and Lee et al. (2012). Among this work on the stochastically constrained SO problem, we know of no work that focuses on optimal allocation to efficiently solve problems on *large* finite sets.

To provide a better sense of the specific questions we answer, consider the following setting. Suppose simulation runs are allocated across the available systems according to an allocation scheme on the design space. After expending a certain amount of the simulation budget, the system with the largest observed objective function estimate among those estimated to be feasible is chosen as the best system. The estimated best system may or may not coincide with the true best system, thereby giving rise to the notion of a false selection event, which is the event that the estimated best system is not the true best system. The probability of false selection ($P\{FS\}$) is the probability of observing a false selection event.

Our questions in this paper relate to the behavior of $P\{FS\}$ as a function of the simulation budget and its allocation across systems, with a special emphasis on settings where the design space is large and while making no independence assumptions between the objective and constraint estimators for a system. Specifically, we ask:

- Q.1 What is the optimal simulation budget allocation across designs, that is, what is the nature of the budget allocation that maximizes the rate of decay of $P\{FS\}$ to zero?
- Q.2 What is the nature of the optimal budget allocation scheme as the number of systems within the design space becomes large?
- Q.3 Can the answer to Q.2 be used to construct an easily implementable algorithmic scheme to solve large-scale stochastically constrained finite SO problems?

Our answer to question Q.1 appears in Section 3 and is a relatively simple extension of recent work by Hunter (2011), Hunter and Pasupathy (2012), and Hunter et al. (2012), and is in general based on the seminal work by Glynn and Juneja (2004). We answer question Q.2 in Section 4, where we demonstrate that the optimal allocation in the proposed setting reduces to a form that is remarkably simple in structure and intuition. Specifically, we show that as the number of systems becomes large, the optimal simulation budget allocated to any suboptimal system i is inversely proportional to a suboptimality/infeasibility measure of a system that we call the *score*. Not surprisingly, the score for system i depends only on the random variables inherent to the system i and best feasible system. Furthermore, the score has an expression that seems easily estimable when the distributions driving the observations from each system are known or assumed. For example, when the observations corresponding to the constraint and objective functions from each system are independently normal, the score for a system is the sum across its optimality gap and infeasibility gaps for violated constraints, each measured in standard deviation units. More generally, calculating the score amounts to minimizing a strictly convex function with box constraints.

From the implementation standpoint of Q.3, when solving constrained SO problems with large finite spaces, our insight from answering Q.2 points to a solution scheme with three repeating steps: estimate the score, update the optimal simulation allocation across systems to be in inverse proportion to the estimated scores, and then select designs on which to execute the simulation according to the updated allocation scheme. As we demonstrate, this procedure results in a simple sequential algorithm that asymptotically achieves the optimal budget allocation scheme, while reliably solving "large" problems with known or assumed distributions.

2 PROBLEM SETTING AND FORMULATION

In this section, we outline a formal problem statement, notational conventions, and assumptions.

2.1 Problem Statement

The problem statement we consider here is identical to that in Hunter and Pasupathy (2012). Unlike Hunter and Pasupathy (2012), however, the investigation in this paper pertains to the setting where the number of systems r is large.

> **Problem Statement:** Consider a finite set $i = 1,2,\ldots,r$ of systems, each with an unknown objective value $h_i \in \mathbb{R}$ and unknown constraint values $g_{ij} \in \mathbb{R}$, $j = 1,2,\ldots,s$ and $i = 1,2,\ldots,r$. Given constants $\gamma_j \in \mathbb{R}$, $j = 1,2,\ldots,s$, we wish to select the system with the lowest objective value h_i, subject to the constraints $g_{ij} \leq \gamma_j$. That is, we consider
>
> $$\text{Problem } P: \text{ Find } \arg\min_{i=1,2,\ldots,r} \quad h_i$$
> $$\text{s.t.} \quad g_{ij} \leq \gamma_j, \text{ for all } j = 1,2,\ldots,s,$$
>
> where h_i and g_{ij} are expectations, estimates of h_i and g_{ij} are observed together through simulation as sample means, and a unique solution to Problem P is assumed to exist.
>
> Let $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \ldots, \alpha_r)$ be a vector denoting the proportion of the total sampling budget given to each system, so that $\sum_{i=1}^{r} \alpha_i = 1$ and $\alpha_i \geq 0$ for all $i = 1,2,\ldots,r$. Let the system having the smallest estimated objective value among the estimated-feasible systems be selected as the estimated solution to Problem P. Then we ask, what vector of proportions $\boldsymbol{\alpha}$ maximizes the rate of decay of the probability that this procedure returns a suboptimal solution to Problem P?

2.2 Notational Conventions

Where it is reasonable to do so, we generally use upper case letters for random variables, lower case letters for fixed quantities, bold type for vectors, and script letters for sets. For brevity, we write $i \leq r$ and $j \leq s$ to indicate $i = 1,2\ldots,r$ and $j = 1,2\ldots,s$. For vectors $\boldsymbol{x} = (x_1, x_2, \ldots, x_m)$ and $\boldsymbol{y} = (y_1, y_2, \ldots, y_m)$, the notation $\boldsymbol{x} \leq \boldsymbol{y}$ means $x_i \leq y_i$ for all $i \leq m$. Throughout the paper, we let system 1 denote the best feasible system, that is, the system with the smallest value of h_i that satisfies the constraints $g_{ij} \leq \gamma_j$ for all $j \leq s$.

2.3 Assumptions

This paper follows from the general theory for constrained simulation optimization with correlation between the objective and constraint estimators outlined in Hunter et al. (2012). To this end, we require the same assumptions as those required in Hunter et al. (2012). First, to estimate the unknown quantities h_i and $\boldsymbol{g}_i = (g_{i1}, g_{i2}, \ldots, g_{is})$, we assume we may obtain replicates of the output random variables $(H_i, \boldsymbol{G}_i)$ from each system, where each system is simulated independently of the others.

Assumption 1 The systems are simulated independently of each other, that is, the random vectors $(H_i, \boldsymbol{G}_i)$ are mutually independent for all $i \leq r$.

We also require the assumption that no system lies exactly on a constraint, and that no system has exactly the same objective function value as that of the best feasible system, system 1. This assumption is standard in literature that seeks an optimal sampling allocation since it ensures that two values may be distinguished with finite simulation budget.

Assumption 2 We assume $h_i \neq h_1$ for all $i \leq r$ and $g_{ij} \neq \gamma_j$ for all $i \leq r, j \leq s$.

Since this paper builds directly from the theory derived in Hunter et al. (2012), the following two assumptions, standard in literature using large deviations theory, are required. Since our focus in this paper is to derive a broad sampling law for a large number of systems, we replicate these assumptions for completeness and refer the reader to Dembo and Zeitouni (1998) for further explanation. We first define the required notation.

Let random variables from the simulation be replicates of $(H_i, \boldsymbol{G}_i) = (H_i, G_{i1}, G_{i2}, \ldots, G_{is})$ for each $i \leq r$. Then define the sample mean after t observations as

$$(\bar{H}_i(t), \bar{\boldsymbol{G}}_i(t)) = (\bar{H}_i(t), \bar{G}_{i1}(t), \ldots, \bar{G}_{is}(t)) = \left(\tfrac{1}{t}\textstyle\sum_{k=1}^{t} H_{ik}, \tfrac{1}{t}\sum_{k=1}^{t} G_{i1k}, \ldots, \tfrac{1}{t}\sum_{k=1}^{t} G_{isk}\right).$$

We use $(\hat{H}_i, \hat{\boldsymbol{G}}_i) \equiv (\bar{H}_i(\alpha_i t), \bar{\boldsymbol{G}}_i(\alpha_i t))$ as shorthand for the estimator of $(h_i, \boldsymbol{g}_i)$ when system i receives $\alpha_i > 0$ proportion of the total sampling budget t. For simplicity, we ignore that $\alpha_i t$ is not necessarily an integer. Let the cumulant generating functions of $\bar{H}_i(t)$, $\bar{G}_{ij}(t)$, and $(\bar{H}_i(t), \bar{\boldsymbol{G}}_i(t))$ be $\Lambda^{(t)}_{H_i}(\theta) = \log \mathrm{E}\left[e^{\theta \bar{H}_i(t)}\right]$, $\Lambda^{(t)}_{G_{ij}}(\theta) = \log \mathrm{E}\left[e^{\theta \bar{G}_{ij}(t)}\right]$, and $\Lambda^{(t)}_{(H_i, \boldsymbol{G}_i)}(\boldsymbol{\theta}) = \log \mathrm{E}\left[e^{\langle \boldsymbol{\theta}, (\bar{H}_i(t), \bar{\boldsymbol{G}}_i(t))\rangle}\right]$, respectively, where $\theta \in \mathbb{R}$, $\boldsymbol{\theta} \in \mathbb{R}^{s+1}$, and $\langle \cdot, \cdot \rangle$ denotes the dot product. Let the effective domain of a function $f(\cdot)$ be denoted by $\mathcal{D}_f = \{x : f(x) < \infty\}$ and its interior by $\mathcal{D}_f^\circ$. Let $\nabla f(\boldsymbol{x})$ be the gradient of f with respect to $\boldsymbol{x}$, and $f'(x)$ the derivative of f with respect to x.

Assumption 3 Let the following hold for each $i \leq r$ and $j \leq s$:

(1) the limit $\Lambda_{(H_i, \boldsymbol{G}_i)}(\boldsymbol{\theta}) = \lim_{t\to\infty} \frac{1}{t} \Lambda^{(t)}_{(H_i, \boldsymbol{G}_i)}(t\boldsymbol{\theta})$ exists as an extended real number for all $\boldsymbol{\theta} \in \mathbb{R}^{s+1}$, where we denote $\Lambda_{H_i}(\theta) = \lim_{t\to\infty} \frac{1}{t} \Lambda^{(t)}_{H_i}(t\theta)$ and $\Lambda_{G_{ij}}(\theta) = \lim_{t\to\infty} \frac{1}{t} \Lambda^{(t)}_{G_{ij}}(t\theta)$ for all $\theta \in \mathbb{R}$;

(2) the origin belongs to the interior of $\mathcal{D}_{\Lambda_{(H_i, \boldsymbol{G}_i)}}$, that is, $0 \in \mathcal{D}^\circ_{\Lambda_{(H_i, \boldsymbol{G}_i)}}$;

(3) $\Lambda_{(H_i, \boldsymbol{G}_i)}(\boldsymbol{\theta})$ is strictly convex and C^∞ on $\mathcal{D}^\circ_{\Lambda_{(H_i, \boldsymbol{G}_i)}}$;

(4) $\Lambda_{(H_i, \boldsymbol{G}_i)}(\boldsymbol{\theta})$ is steep, that is, for any sequence $\{\boldsymbol{\theta}(t)\} \in \mathcal{D}_{\Lambda_{(H_i, \boldsymbol{G}_i)}}$ converging to a boundary point of $\mathcal{D}_{\Lambda_{(H_i, \boldsymbol{G}_i)}}$, then $\lim_{t\to\infty} |\nabla \Lambda_{(H_i, \boldsymbol{G}_i)}(\boldsymbol{\theta}(t))| = \infty$.

Under Assumption 3, the large deviations principle (LDP) holds for the estimators $\bar{H}_i(t)$, $\bar{G}_{ij}(t)$, and $(\bar{H}_i(t), \bar{\boldsymbol{G}}_i(t))$ with strictly convex rate functions $I_i(x) = \sup_{\theta\in\mathbb{R}}\{\theta x - \Lambda_{H_i}(\theta)\}$, $J_{ij}(y) = \sup_{\theta\in\mathbb{R}}\{\theta y - \Lambda_{G_{ij}}(\theta)\}$, and $I_i(x, \boldsymbol{y}) = \sup_{\boldsymbol{\theta}\in\mathbb{R}^{s+1}}\{\langle \boldsymbol{\theta}, (x, \boldsymbol{y})\rangle - \Lambda_{(H_i, \boldsymbol{G}_i)}(\boldsymbol{\theta})\}$, respectively. Let $\boldsymbol{\gamma} = (\gamma_1, \ldots, \gamma_s)$,

$$(x, \boldsymbol{y}) \in \mathcal{F}^\circ_{(H_i, \boldsymbol{G}_i)} = \mathrm{int}\{\nabla \Lambda_{(H_i, \boldsymbol{G}_i)}(\boldsymbol{\theta}) : \boldsymbol{\theta} \in \mathcal{D}^\circ_{\Lambda_{(H_i, \boldsymbol{G}_i)}}\},$$

and let $\mathcal{F}_d^c$ denote the closure of the convex hull of the set of points $\{(h_i, \boldsymbol{\gamma}) : (h_i, \boldsymbol{\gamma}) \in \mathbb{R}^{s+1}\}$.

Assumption 4 The closure of the convex hull of all points $(h_i, \boldsymbol{\gamma}) \in \mathbb{R}^{s+1}$ is a subset of the intersection of the interiors of the effective domains of the rate functions $I_i(x, \boldsymbol{y})$ for all $i \leq r$, that is, $\mathcal{F}_d^c \subset \cap_{i=1}^r \mathcal{F}^\circ_{(H_i, \boldsymbol{G}_i)}$.

3 CHARACTERIZATION OF THE OPTIMAL BUDGET ALLOCATION

Recall that our problem context is Problem *P* (see Section 2.1), and our solution context involves three steps: sample from each of the designs to obtain objective function and constraint function estimators; estimate the feasible set of systems by observing their constraint function estimators; and estimate the optimal system from the estimated feasible set as that system having the smallest estimated objective function value. In this section, we rigorously characterize the optimal allocation as the allocation that minimizes the probability that the system returned as the "solution" at the end of some sampling effort t is not the true best system.

We build upon the characterization of the optimal budget allocation for general distributions in the presence of correlation between the objective and constraint estimators that was formally derived in Hunter et al. (2012). Hunter et al. (2012) characterize the optimal allocation as the solution to a concave maximization problem. We replicate the key results here, and then further characterize the solution to the concave maximization problem in terms of its Karush-Kuhn-Tucker (KKT) conditions (Boyd and Vandenberghe 2004). For brevity, results are presented without proof.

Recall that t is the computing budget, $\alpha_i \in [0, 1]$ is the fraction of the simulation budget devoted to system i, $\hat{H}_i = (\alpha_i t)^{-1} \sum_{k=1}^{\alpha_i t} H_{ik}$, and $\hat{G}_{ij} = (\alpha_i t)^{-1} \sum_{k=1}^{\alpha_i t} G_{ijk}$. From Hunter et al. (2012), the probability of

selecting a system other than the best feasible system, system 1, is

$$P\{FS\} = P\{\underbrace{[\cup_{j=1}^{s} \hat{G}_{1j} > \gamma_j]}_{\substack{\text{system 1}\\ \text{estimated}\\ \text{infeasible}}} \cup_{i=2}^{r} [\underbrace{(\cap_{j=1}^{s} \hat{G}_{ij} \leq \gamma_j)}_{\substack{\text{system i}\\ \text{estimated}\\ \text{feasible}}} \cap \underbrace{(\hat{H}_i \leq \hat{H}_1)}_{\substack{\text{system i}\\ \text{“beats”}\\ \text{system 1}}}]\}$$

and the rate function of the $P\{FS\}$ is

$$-\lim_{t\to\infty} \frac{1}{t} \log P\{FS\} = \min\left(\min_{j\leq s} \alpha_1 J_{1j}(\gamma_j), \min_{2\leq i\leq r} R_i(\alpha_1, \alpha_i) \right), \tag{1}$$

where $\alpha_1 J_{1j}(\gamma_j)$ is the rate function for the probability that system 1 is classified infeasible due to the violation of the jth constraint, and $R_i(\alpha_1, \alpha_i) = \inf_{x_i \leq x_{1i}, \boldsymbol{y_i} \leq \boldsymbol{\gamma}} (\alpha_1 I_1(x_{1i}) + \alpha_i I_i(x_i, \boldsymbol{y_i}))$ is the rate function for the probability that system i is estimated feasible *and* system i has a "better" estimated objective function value that system 1. We use the double subscript on x_{1i} to indicate that this variable is the variable corresponding to system 1 in the ith problem $R_i(\alpha_1, \alpha_i)$.

We are interested in identifying the allocation $\boldsymbol{\alpha}$ that maximizes the rate of decay in (1). This problem can be formally stated as

$$\max \quad \min\left(\min_{j\leq s} \alpha_1 J_{1j}(\gamma_j), \min_{2\leq i\leq r} R_i(\alpha_1, \alpha_i) \right) \quad s.t. \quad \sum_{i=1}^{r} \alpha_i = 1, \quad \alpha_i \geq 0.$$

We may equivalently write this problem as

$$\begin{aligned} \text{Problem } Q: \qquad \max \quad & z \quad s.t. \\ & \alpha_1 J_{1j}(\gamma_j) \geq z, \quad j \leq s \\ & R_i(\alpha_1, \alpha_i) \geq z, \quad i = 2, \ldots r, \\ & \sum_{i=1}^{r} \alpha_i = 1, \ \alpha_i \geq 0, \end{aligned}$$

where for each $i = 2, \ldots, r$, the values of $R_i(\alpha_1, \alpha_i)$ are obtained by solving

$$\text{Problem } R_i: \qquad \min \quad \alpha_1 I_1(x_{1i}) + \alpha_i I_i(x_i, \boldsymbol{y_i}) \quad s.t. \quad x_i \leq x_{1i}, \quad \boldsymbol{y_i} \leq \boldsymbol{\gamma}.$$

As a matter of notation, we distinguish Problem R_i as an optimization problem in $(x_{1i}, x_i, \boldsymbol{y_i})$, and $R_i(\alpha_1, \alpha_i)$ as the value of its objective function at optimality. By Hunter et al. (2012), Problem R_i is a strictly convex minimization problem with a unique optimal solution. Further, Problem Q is a concave maximization problem to which the optimal solution exists, and the solution is strictly positive, that is, $\boldsymbol{\alpha}^* = (\alpha_1^*, \alpha_2^*, \ldots, \alpha_r^*) > \mathbf{0}$ and hence all systems receive a nonzero fraction of the sampling budget at optimality.

From the analysis in Pujowidianto et al. (2012), under some regularity conditions, for large enough r, the optimal allocation $\boldsymbol{\alpha}^*$ is obtained as the solution to the KKT system

$$R_i(\alpha_1^*, \alpha_i^*) = R_k(\alpha_1^*, \alpha_k^*) \text{ for all } i, k \neq 1, \tag{2}$$

$$\sum_{i=2}^{r} \frac{\partial R_i(\alpha_1^*, \alpha_i^*)/\partial \alpha_1}{\partial R_i(\alpha_1^*, \alpha_i^*)/\partial \alpha_i} = 1. \tag{3}$$

We derive an explicit expression for the summands in equation (3) in the following Lemma 1.

Lemma 1 For a system i, the ith term in the summand of equation (3) is given by

$$\frac{\partial R_i(\alpha_1,\alpha_i)/\partial\alpha_1}{\partial R_i(\alpha_1,\alpha_i)/\partial\alpha_i} = \frac{I_1(x_{1i}^*)}{I_i(x_i^*,\boldsymbol{y}_{\boldsymbol{i}}^*)}, \tag{4}$$

where $(x_{1i}^*, x_i^*, \boldsymbol{y}_{\boldsymbol{i}}^*)$ is the unique optimal solution to Problem R_i.

Let $\lambda_{ix} \geq 0$ and $\lambda_{ij} \leq 0, j \leq s$ be the Lagrange multipliers associated with the constraints in Problem R_i, where $\boldsymbol{\lambda_i} = (\lambda_{ix}, \lambda_{i1}, \ldots, \lambda_{is})$. We also define the following sets which are functions of these Lagrange multipliers and the optimal solution to Problem R_i:

$$\begin{aligned}
\mathcal{C}_I^{i*} &= \{j : \lambda_{ij} < 0 \text{ and } y_{ij}^* = \gamma_j\};\\
\mathcal{C}_F^{i*} &= \{j : \lambda_{ij} = 0 \text{ and } y_{ij}^* \leq \gamma_j\};\\
\Gamma^* &= \{i : \lambda_{ix} > 0, x_{1i}^* = x_i^* \text{ and } \mathcal{C}_I^{i*} \text{ empty}, i \neq 1\};\\
\mathcal{S}_b^* &= \{i : \lambda_{ix} = 0, x_i^* \leq x_{1i}^* \text{ and } \mathcal{C}_I^{i*} \text{ nonempty}, i \neq 1\};\\
\mathcal{S}_w^* &= \{i : \lambda_{ix} > 0, x_{1i}^* = x_i^* \text{ and } \mathcal{C}_I^{i*} \text{ nonempty}, i \neq 1\}.
\end{aligned}$$

The sets $\mathcal{C}_I^{i*}$ and $\mathcal{C}_F^{i*}$ form a partitioning of the set of constraints $\{1,2,\ldots,s\}$ for each design i, and the sets $\Gamma^*, \mathcal{S}_b^*$, and $\mathcal{S}_w^*$ form a partition of the design space $\{1,2,\ldots,r\}$. For example, it is seen that when the objective function and constraint estimators are mutually independent, the sets $\mathcal{C}_I^{i*}$ and $\mathcal{C}_F^{i*}$ are the sets of constraints on which system i is infeasible and feasible, respectively. Likewise, under mutual independence, the sets $\Gamma^*, \mathcal{S}_b^*$, and $\mathcal{S}_w^*$ correspond to the set of truly feasible designs, the set of truly infeasible designs that are better than system 1 in objective function value, and the set of truly infeasible designs that are worse than system 1 in objective function value, respectively.

From equations (2) and (3), the terms of the simplified summand in equation (4) of Lemma 1 can be further simplified to

$$R_i(\alpha_1^*, \alpha_i^*) = R_k(\alpha_1^*, \alpha_k^*) \text{ for all } i,k \neq 1, \tag{5}$$

$$\sum_{i \in \Gamma^* \cup \mathcal{S}_w^*} \frac{I_1(x_{1i}^*)}{I_i(x_i^*, \boldsymbol{y}_{\boldsymbol{i}}^*)} = 1. \tag{6}$$

We note that the summand in equation (6) contains nonzero terms only for systems in $\Gamma^* \cup \mathcal{S}_w^*$ since the rate functions for systems in $\mathcal{S}_b^*$ do not depend on the rate function corresponding to system 1, $I_1(\cdot)$.

Since the rate functions involved in (5) and (6) are unknown and cumbersome to estimate, Problem Q is usually impractical to solve in all generality. However as we demonstrate in the sections that follow, the KKT conditions for Problem Q become remarkably easier to solve under certain conditions — most notably when the number of systems r tends to infinity. This limiting approximation forms the basis of our proposed solution to Problem Q.

4 A LIMITING APPROXIMATION TO THE OPTIMAL BUDGET ALLOCATION

With a view toward efficiently solving Problem Q, this section proposes a "closed-form" limiting approximation to the solution of the KKT system for Problem Q, obtained as a certain asymptotic limit. Specifically, it is shown in Section 4.2 that under certain conditions, the fraction of the budget that should be devoted to each of the suboptimal systems is inversely proportional to an easily-expressed penalty measure that we call the *score*. In the following Section 4.1, we detail some important properties of the summands appearing in the KKT system for Problem Q before we proceed to the main results.

4.1 Some Key Properties

First, we show through Lemma 2 that the summands appearing in the KKT system for Problem Q are within a positive finite constant from each other.

Lemma 2 There exists $c > 0$ such that

$$\frac{I_1(x^*_{1i})}{I_i(x^*_i, \boldsymbol{y}^*_{\boldsymbol{i}})} \leq c \frac{I_1(x^*_{1k})}{I_k(x^*_k, \boldsymbol{y}^*_{\boldsymbol{k}})} \text{ for all } i,k \in \Gamma^* \cup \mathcal{S}^*_w.$$

Lemma 2 essentially states that neither $I_1(x^*_{1i}), i \in \Gamma^* \cup \mathcal{S}^*_w$, nor $I_i(x_i, \boldsymbol{y}^*_{\boldsymbol{i}})$ is zero, an assertion that is intuitively clear upon noting that the sampling proportions $\alpha^*_i > 0$ and the rate functions $I_1(\cdot), I_i(\cdot,\cdot)$ are strictly convex.

We now demonstrate two key assertions in the following Theorem 3. First, we state that at optimality, each of the summands appearing in (6) of the KKT system for Problem Q tends to zero as the number of systems in $\Gamma^* \cup \mathcal{S}^*_w$ tends to infinity. The second assertion uses the first to demonstrate that, as the number of systems in $\Gamma^* \cup \mathcal{S}^*_w$ tends to infinity, the optimal budget fractions are such that the true best design receives "far more" of the simulation budget than any of the suboptimal designs.

Theorem 3 As $|\Gamma^* \cup \mathcal{S}^*_w| \to \infty$, the following hold.

(i) $\dfrac{I_1(x^*_{1i})}{I_i(x^*_i, \boldsymbol{y}^*_{\boldsymbol{i}})} \to 0$ for all $i \in \Gamma^* \cup \mathcal{S}^*_w$.

(ii) $\alpha^*_i / \alpha^*_1 \to 0$ for all $i \in \Gamma^* \cup \mathcal{S}^*_w$.

Theorem 3, particularly through assertion (ii), conveys an important message. As the number of systems in $\Gamma^* \cup \mathcal{S}^*_w$ tends to infinity, optimality dictates that the fraction of the budget given to the optimal design, system 1, far exceed any of the fraction given to suboptimal designs in $\Gamma^* \cup \mathcal{S}^*_w$. This result makes intuitive sense if one thinks of each of the suboptimal systems as individually attempting to "beat the best design" by inducing a false selection event. Optimality dictates that the best design receive far more sample than these competitors in a bid to minimize the probability of occurrence of the most likely of the numerous false selection events, made possible by the assumption $|\Gamma^* \cup \mathcal{S}^*_w| \to \infty$.

Consider the assumption that $|\Gamma^* \cup \mathcal{S}^*_w| \to \infty$ in the context where the objective function and constraint estimators for each system are mutually independent. In such a context, the assumption implies that the collective cardinality of the set of truly feasible systems, Γ, and the set of truly infeasible systems that are "worse" than the best design, $\mathcal{S}_w$, tends to infinity. In the more general context, the interpretation becomes slightly more nuanced. A sufficient condition that guarantees that the assumption holds in the general context is that the cardinality of the set of truly feasible systems Γ tends to infinity. The conditions on the supremum norm of the variances and means being bounded is to avoid pathological cases where the number of designs tends to infinity, but the effect of the designs being added is diminishing in the sense that they do not compete in any real way to increase the overall probability of false selection.

4.2 The Limiting Approximation and A Sequential Algorithm

We now have the machinery required to characterize the allocation given to the suboptimal systems as $|\Gamma^* \cup \mathcal{S}^*_w| \to \infty$. Theorem 4 asserts that as $|\Gamma^* \cup \mathcal{S}^*_w| \to \infty$, the ratio of the rate $R_i(\alpha^*_1, \alpha^*_i)$ to the optimal fraction α^*_i for the ith system tends to the minimum value attained by the rate function $I_i(x_i, \boldsymbol{y}_{\boldsymbol{i}})$ in the box $x_i \leq h_1, \boldsymbol{y}_{\boldsymbol{i}} \leq \boldsymbol{\gamma}$.

Theorem 4 As $|\Gamma^* \cup \mathcal{S}^*_w| \to \infty$,

$$\frac{R_i(\alpha^*_1, \alpha^*_i)}{\alpha^*_i} = \frac{R_i(\alpha^*_i)}{\alpha^*_i} = \inf_{x_i \leq h_1, \boldsymbol{y}_{\boldsymbol{i}} \leq \boldsymbol{\gamma}} I_i(x_i, \boldsymbol{y}_{\boldsymbol{i}}) \text{ for all } i = 2, \ldots, r.$$

Theorem 4 combined with the fact that the KKT system for Problem Q in (5) dictates equating $R_i(\alpha^*_1, \alpha^*_i) = R_k(\alpha^*_1, \alpha^*_k)$ for $i,k \in \{2,3,\ldots,r\}$ provides a basis for budget allocation. Theorem 5 provides an expression for this allocation among suboptimal systems through the scores $\mathbb{S}_i, i = 2,3,\ldots,r$ for the various systems.

Theorem 5 As $|\Gamma^* \cup \mathcal{S}_w^*| \to \infty$, the allocation to non-best systems $i = 2, \ldots, r$ is determined by

$$\frac{\alpha_i^*}{\alpha_k^*} \to \frac{\mathbb{S}_k}{\mathbb{S}_i} = \frac{\inf_{x_k \le h_1, \boldsymbol{y_k} \le \boldsymbol{\gamma}} I_k(x_k, \boldsymbol{y_k})}{\inf_{x_i \le h_1, \boldsymbol{y_i} \le \boldsymbol{\gamma}} I_i(x_i, \boldsymbol{y_i})},$$

where the *score* $\mathbb{S}_i$ for (suboptimal) system i as

$$\mathbb{S}_i := \inf_{x_i \le h_1, \boldsymbol{y_i} \le \boldsymbol{\gamma}} I_i(x_i, \boldsymbol{y_i}) \text{ for all } i = 2, \ldots, r.$$

Theorem 5 shows that when the number of systems is large, the allocation to the suboptimal systems becomes inversely proportional to the score of each suboptimal system. This result is useful in practice, as we demonstrate through Algorithm 1 — a sequential algorithm that uses the notion of the score.

Algorithm 1 A Sequential Algorithm

Require: Number of pilot samples $\delta_0 > 0$; number of samples between allocation vector updates $\delta > 0$; and a minimum-sample vector $\boldsymbol{\varepsilon} > 0$.

1: Initialize: collect δ_0 samples from each system $i \le r$.
2: Initialize: total simulation effort $n = r\delta_0$, effort for each system $n_i = \delta_0$.
3: Update the objective and constraint estimators $(\hat{H}_i(n_i), \hat{\boldsymbol{G}}_{\boldsymbol{i}}(n_i))$, the rate function estimator $\hat{I}_{i,n}(x_i, \boldsymbol{y_i})$, and the score estimator $\hat{\mathbb{S}}_{i,n}$ for all $i \le r$.
4: **if** no systems are estimated feasible **then**
5: Set $\hat{\boldsymbol{\alpha}}_n^* = (1/r, 1/r, \ldots, 1/r)$.
6: **else**
7: Update $\hat{1}(n)$, the estimated system 1, and its allocation $\hat{\alpha}_{1,n}^*$.
8: Set $\hat{\alpha}_{i,n}^* = (\sum_{k=1}^r \hat{\mathbb{S}}_k^{-1})^{-1} \times \hat{\mathbb{S}}_i^{-1} \times (1 - \hat{\alpha}_{1,n}^*)$ for all systems $i \ge 2$.
9: **end if**
10: Collect one sample at each system $X_k, k = 1, 2, \ldots, \delta$, where the X_k's are iid random variates with probability mass function $\hat{\boldsymbol{\alpha}}_n^*$ and support $\{1, 2, \ldots, r\}$. Update $n_{X_k} = n_{X_k} + 1$.
11: Set $n = n + \delta$ and update $\bar{\boldsymbol{\alpha}}_{\boldsymbol{n}} = \{n_1/n, n_2/n, \ldots, n_r/n\}$.
12: **if** $\bar{\boldsymbol{\alpha}}_n > \boldsymbol{\varepsilon}$ **then**
13: Set $\delta^+ = 0$.
14: **else**
15: Collect one sample from each system in the set of systems receiving insufficient sample $\mathcal{I}_n$.
16: Update $n_i = n_i + 1$ for all $i \in \mathcal{I}_n$. Let $\delta^+ = |\mathcal{I}_n|$.
17: **end if**
18: Set $n = n + \delta^+$ and go to step 3.

The essence of Algorithm 1 is straightforward. The algorithm evolves in stages by collecting a fixed number of simulation observations from systems chosen strategically at the beginning of each stage, updating the relevant estimators, and then proceeding to the next stage to begin the process over again. Specifically, at the beginning of each stage, $\delta > 0$ observations are obtained from systems chosen with probabilities in accordance with the prevailing estimated optimal fractions $\hat{\boldsymbol{\alpha}}_n^* = \{\hat{\alpha}_{1,n}^*, \hat{\alpha}_{2,n}^*, \ldots, \hat{\alpha}_{r,n}^*\}$, where n represents the expended number of simulation calls. The observations are then used to update the estimated scores $\hat{\mathbb{S}}_{i,n}$ for systems $i \ge 2$, and the estimated best solution $\hat{1}(n)$. The iterative process continues by using the updated scores to modify the estimated optimal fractions $\hat{\boldsymbol{\alpha}}_n^*$, which will in turn be used as the system choice probabilities in the subsequent stage.

While we have characterized the relative optimal allocations for the suboptimal systems $i \ge 2$ through the scores, this says nothing about what fraction of the budget should be allocated to the best system.

Theorem 3 is not of direct relevance in this context since only tells us that the fraction received by best system should far exceed that of any of the other suboptimal systems in the limit. This question is under investigation, and a heuristic based on the rate at which α_i^*/α_1^* tends to zero seems reasonable.

CONCLUDING REMARKS

The question of efficiently identifying the best among a finite set of competing systems in the presence of "stochastic" objective and constraint functions seems to be an important SO variation about which little is currently known. As we have shown, it is possible to pose this question as an optimization problem having a certain measure of simulation efficiency as the objective function, and the fraction of the simulation budget given to the various systems involved as the decision variables. Solving this optimization problem, however, turns out to be difficult in general because of its bilevel nesting structure. This intractability, in combination with an interest in solving large-scale problems, has inspired our investigation into the existence of "near optimal" allocations that are easily identifiable.

Our analysis has revealed two interesting facts: (i) as the number of systems in contention grows, the solution to the bilevel nested optimization problem becomes "closed-form" when expressed in terms of a single measure, called the score, encompassing the infeasibility and suboptimality of each system; and (ii) the score for each system is itself easy to estimate in many common scenarios. These two facts lead to a sequential algorithm that seems to efficiently solve large-scale constrained SO problems on finite sets. For example, preliminary numerical experiments in the multivariate normal context reveal that the outlined algorithm reliably solves constrained SO problems with many thousands of systems within seconds on a laptop computer (with 2.66 GHz Intel Core 2 Duo processor and 8GB of memory).

ACKNOWLEDGMENTS

Susan R. Hunter was supported in part by National Science Foundation Grants CMMI 0758441 and CMMI 0800688. Raghu Pasupathy acknowledges the financial support from the Office of Naval Research via contract N00014–08–1–0066. Chun-Hung Chen acknowledges support by the Department of Energy under Award DE–SC0002223, NIH under Grant 1R21DK088368–01, and National Science Council of Taiwan under Award NSC–100–2218–E–002–027–MY3.

REFERENCES

Andradóttir, S., and S.-H. Kim. 2010. "Fully Sequential Procedures for Comparing Constrained Systems via Simulation". *Naval Research Logistics* 57 (5): 403–421.

Boyd, S., and L. Vandenberghe. 2004. *Convex Optimization*. New York: Cambridge University Press.

Dembo, A., and O. Zeitouni. 1998. *Large Deviations Techniques and Applications*. 2nd ed. New York: Springer.

Glynn, P. W., and S. Juneja. 2004, December. "A large deviations perspective on ordinal optimization". In *Proceedings of the 2004 Winter Simulation Conference*, edited by R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 577–585. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Hunter, S. R. 2011. *Sampling Laws for Stochastically Constrained Simulation Optimization on Finite Sets*. Ph. D. thesis, Virginia Polytechnic Institute and State University.

Hunter, S. R., and R. Pasupathy. 2012. "Optimal sampling laws for stochastically constrained simulation optimization on finite sets". To appear in *INFORMS Journal on Computing*.

Hunter, S. R., N. A. Pujowidianto, R. Pasupathy, L. H. Lee, and C. H. Chen. 2012. "Constrained and Correlated Simulation Optimization on Finite Sets". Under Review.

Lee, L. H., N. A. Pujowidianto, L.-W. Li, C.-H. Chen, and C. M. Yap. 2012. "Approximate Simulation Budget Allocation for Selecting the Best Design in the Presence of Stochastic Constraints". To appear in *IEEE Transactions on Automatic Control*.

Pasupathy, R., and S. G. Henderson. 2006, December. "A Testbed of Simulation-Optimization Problems". In *Proceedings of the 2006 Winter Simulation Conference*, edited by L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, 255–263. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Pasupathy, R., and S. G. Henderson. 2011, December. "SimOpt: A Library of Simulation Optimization Problems". In *Proceedings of the 2011 Winter Simulation Conference*, edited by S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, and M. Fu, 4080–4090. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Pujowidianto, N. A., S. R. Hunter, R. Pasupathy, L. H. Lee, and C.-H. Chen. 2012. "On approximating optimal sampling laws for stochastically constrained simulation optimization on large finite sets". Under Review.

AUTHOR BIOGRAPHIES

NUGROHO A. PUJOWIDIANTO is a Ph.D. candidate in the Department of Industrial and Systems Engineering, National University of Singapore. He teaches Design and Technology and Computer Applications at Pioneer Secondary School, Singapore. He received his B.Eng. (Mechanical Engineering) degree from Nanyang Technological University in 2006. His research interests include simulation optimization and its application in health care. His email address is nugroho@nus.edu.sg.

SUSAN R. HUNTER is a postdoctoral associate in the School of Operations Research and Information Engineering at Cornell University. Her research interests include Monte Carlo methods and simulation optimization. Her email address is hunter@cornell.edu, and her webpage is http://people.orie.cornell.edu/srh227.

RAGHU PASUPATHY is an associate professor in the Grado Department of Industrial and Systems Engineering at Virginia Tech. His research interests lie broadly in Monte Carlo methods with a specific focus on simulation optimization. He is a member of INFORMS, IIE, and ASA, and serves as an associate editor for *Operations Research*, *ACM TOMACS* and *INFORMS Journal on Computing*. His email address is pasupath@vt.edu and his web page is https://filebox.vt.edu/users/pasupath/pasupath.htm.

LOO HAY LEE is an Associate Professor and Deputy Head (Graduate Studies and Research) in the Department of Industrial and Systems Engineering, National University of Singapore. He received his B.S. (Electrical Engineering) degree from the National Taiwan University in 1992 and his Ph.D. degree in 1997 from Harvard University. He is currently a senior member of IEEE, a member of ORSS and INFORMS. He is the associate editor for IIE Transactions, Flexible Services and Manufacturing Journal, the Asia Pacific Journal of Operational Research, the co-editor for Journal of Simulation and is a member in the advisory board for OR Spectrum. His research interests include simulation-based optimization, and maritime Logistics. His email address is iseleelh@nus.edu.sg.

CHUN-HUNG CHEN is a Professor in the Department of Electrical Engineering and Molecular Imaging Center at National Taiwan University. He received his Ph.D. from Harvard University in 1994. His research interests are mainly in development of very efficient methodology for simulation and optimization and its applications. Dr. Chen has served as Co-Editor of the Proceedings of the 2002 Winter Simulation Conference and Program Co-Chair for 2007 Informs Simulation Society Workshop. He is currently an associate editor of IEEE Transactions on Automatic Control, area editor of Journal of Simulation Modeling Practice and Theory, associate editor of International Journal of Simulation and Process Modeling, and simulation department editor for IIE Transactions. His email address is cchen9@cc.ee.ntu.edu.tw.