Proceedings of the 2015 Winter Simulation Conference
L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, eds.

OPTIMAL SAMPLING LAWS FOR BI-OBJECTIVE SIMULATION OPTIMIZATION ON FINITE SETS

Susan R. Hunter

School of Industrial Engineering
Purdue University
315 N. Grant Street
West Lafayette, IN 47907, USA

Guy Feldman

Department of Statistics
Purdue University
250 N. University Street
West Lafayette, IN 47907, USA

ABSTRACT

We consider the bi-objective simulation optimization (SO) problem on finite sets, that is, an optimization problem where for each "system," the two objective functions are estimated as output from a Monte Carlo simulation. The solution to this bi-objective SO problem is a set of non-dominated systems, also called the Pareto set. In this context, we derive the large deviations rate function for the rate of decay of the probability of a misclassification event as a function of the proportion of sample allocated to each competing system. Notably, we account for the presence of dependence between the estimates of each system's performance on the two objectives. The asymptotically optimal allocation maximizes the rate of decay of the probability of misclassification and is the solution to a concave maximization problem.

1 INTRODUCTION

The simulation optimization (SO) problem is an optimization problem in which the objective(s) and constraint(s) are estimated as output from a Monte Carlo simulation. The literature on SO on finite sets, broadly called Ranking and Selection (R&S), is particularly rich in the context of a single objective (see, e.g., Kim and Nelson 2006 for an overview). By comparison, the body of work on R&S in the presence of multiple performance measures is new, and is developing along the classic lines of literature for a single objective — class $\mathcal{P}$ procedures, which provide a finite-time guarantee on solution quality, and class $\mathcal{L}$ procedures, which provide no finite-time guarantee on solution quality, and instead focus on providing efficiency (Pasupathy and Ghosh 2013). Recent work on R&S in the context of multiple performance measures, one of which is used as the objective function and the rest as constraints, includes a class $\mathcal{P}$ procedure by Andradóttir and Kim (2010) and class $\mathcal{L}$ procedures by Lee et al. (2012), Hunter and Pasupathy (2013), and Pasupathy et al. (2014).

R&S with multiple competing objectives has arguably seen less development to date than stochastically constrained R&S, since we are not aware of any published class $\mathcal{P}$ procedures in this area. However the current standard among class $\mathcal{L}$ procedures for efficient sampling in this context is Multi-objective Optimal Computing Budget Allocation (MOCBA) (Lee et al. 2010). Under the assumption of normally distributed simulation output, Lee et al. (2010) provide a heuristic sampling framework for the case of simultaneous multiple objectives, developed along the lines of the popular Optimal Computing Budget Allocation (OCBA) framework for a single objective by Chen et al. (2000). This procedure does not explicitly account for correlation between the objectives. (See Butler, Morrice, and Mullarkey 2001 for a utility function approach to multi-objective R&S; work on multi-objective SO on integer-ordered or continuous sets includes Ryu et al. 2009, Kim and Ryu 2011b, Kim and Ryu 2011a, Huang and Zabinsky 2014, Li et al. 2015).

We provide mathematical results underlying a competing class $\mathcal{L}$ procedure for R&S in the case of two objectives; our results specifically account for dependence between the objective vector estimates for a single system. (For corresponding work in the more-than-two-objectives case, see Feldman, Hunter, and Pasupathy 2015.) In the bi-objective case, the solution to the R&S problem is a Pareto set, that is, a set of systems that are non-dominated in the objectives. A misclassification event occurs if a truly Pareto system is falsely estimated as non-Pareto or a truly non-Pareto system is falsely estimated as Pareto. To derive a competing procedure, we employ a large deviations (LD) analysis to find the rate of decay of the probability of a misclassification event as a function of the proportional simulation budget allocated to each of the competing systems. We then characterize the asymptotically optimal sampling allocation as the solution to a concave maximization problem that maximizes the rate of decay of a misclassification event. Our work is preceded by Glynn and Juneja (2004), Szechtman and Yücesan (2008), Hunter (2011), Hunter and Pasupathy (2013), Pasupathy et al. (2014), who likewise use an LD analysis to derive an asymptotically optimal sampling allocation in the context of a single objective, a single constraint, and a single objective with multiple constraints, respectively.

The asymptotically optimal sampling allocation we present cannot be implemented as written, since we assume prior knowledge of the rate functions. However, our allocation can easily be incorporated into a sequential sampling framework, as in Hunter (2011), Hunter and Pasupathy (2013), Pasupathy et al. (2014), Hunter and McClosky (2015). Indeed, Hunter and McClosky (2015) presents an implementable version of the proposed allocation for the special case in which the simulation output is Gaussian and the two objectives are the mean and variance of the underlying Gaussian distribution; therefore the two objective function estimators are independent. A numerical comparison of the approach with MOCBA in the context of plant breeding is also included in Hunter and McClosky (2015). However, in more general applications, solving for the optimal allocation requires solving a potentially burdensome bi-level optimization problem. While this approach is successful in Hunter and McClosky (2015), we ideally seek a framework that is easier to implement. Therefore we view this work as the first step in deriving easily implementable Sampling Criteria for Optimization using Rate Estimators (SCORE) allocations (Pasupathy et al. 2014) in the bi-objective context — that is, allocations that are provably optimal, in a certain rigorous sense, both as the sampling budget tends to infinity and as the number of systems tends to infinity.

1.1 Problem Statement

Consider a finite set of r systems, each with unknown objective values $g_i \in \mathbb{R}$ and $h_i \in \mathbb{R}$ for $i = 1,2,\ldots,r$. We wish to select the set of systems that are non-dominated in both objectives, where system k *dominates* system i, written $k \preceq i$, if $g_k \leq g_i$ and $h_k < h_i$, or $g_k < g_i$ and $h_k \leq h_i$. That is, we consider

$$\text{Problem } P: \quad \text{Find } \arg\min_{i \in \{1,2,\ldots,r\}} \; (g_i, h_i),$$

where g_i and h_i are expectations, and estimates of g_i and h_i are observed together through simulation as sample means. The solution to Problem P is the Pareto set of non-dominated systems

$$\mathcal{P} := \{\text{systems } i : \nexists \text{ system } k \in \{1,2,\ldots,r\} \text{ such that } k \preceq i\},$$

which we assume is unique.

Let $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \ldots, \alpha_r)$ be a vector containing the proportion of the sampling budget allocated to each system, where $\sum_{i=1}^{r} \alpha_i = 1$ and $\alpha_i \geq 0$ for all systems $i = 1,\ldots,r$. After the sampling budget has been expended, a misclassification event occurs if a Pareto system has been falsely estimated as non-Pareto, or non-Pareto system has been falsely estimated as Pareto. The probability of a misclassification event tends to zero as the sampling budget tends to infinity. Then we ask, what proportional allocation of the sampling budget $\boldsymbol{\alpha}$ maximizes the rate of decay of the probability of misclassification as the sampling budget tends to infinity?

1.2 Assumptions

To estimate the unknown quantities g_i and h_i, we assume we may obtain replicates of the random variables (G_i, H_i) from each system. We also assume the following, where $i \leq r$ is shorthand for $i = 1, 2, \ldots, r$.

Assumption 1 We assume the random variables (G_i, H_i) are mutually independent for all $i \leq r$.

That is, we develop a model to guide sampling that does not specifically account for correlation between systems, such as the correlation that would arise with the use of common random numbers (CRN). Note that we *do* account for correlation between objectives, since we *do not require* that G_i and H_i be independent for a particular system i. We also require the following technical assumption which is standard in optimal allocation literature, since it ensures all systems in the Pareto set are distinguishable on each objective with a finite sample size.

Assumption 2 We assume $g_i \neq g_k$ and $h_i \neq h_k$ for all $i \in \mathcal{P}$, $k = 1, 2, \ldots, r$, $i \neq k$.

Since we employ a large deviations (LD) analysis, we require the following Assumptions 3 and 4, included here for completeness. We refer the reader to Dembo and Zeitouni (1998) for further explanation; we note that these assumptions are similar to those required in Glynn and Juneja (2004), Hunter and Pasupathy (2013), and Pasupathy et al. (2014). First, we define the required notation.

Let the vector of sample means after n samples be $(\bar{G}_i(n), \bar{H}_i(n)) := (\frac{1}{n}\sum_{k=1}^{n} G_{ik}, \frac{1}{n}\sum_{k=1}^{n} H_{ik})$ for all $i = 1, 2, \ldots, r$. We define $(\hat{G}_i, \hat{H}_i) := (\bar{G}(\alpha_i n), \bar{H}(\alpha_i n))$ as the estimators of g_i and h_i after scaling the total sample size n by $\alpha_i > 0$, the proportional sample allocation to system i. Since our analysis is asymptotic, we ignore issues relating to the fact that $\alpha_i n$ is not an integer. Let $\Lambda^{(n)}_{(G_i,H_i)}(\boldsymbol{\theta}) = \log \mathrm{E}[e^{\langle \boldsymbol{\theta}, (\bar{G}_i(n), \bar{H}_i(n)) \rangle}]$, be the cumulant generating function of $(\bar{G}_i(n), \bar{H}_i(n))$, where $\boldsymbol{\theta} \in \mathbb{R}^2$ and $\langle \cdot, \cdot \rangle$ denotes the dot product. Let the effective domain of a function $f(\cdot)$ be denoted $\mathcal{D}_f = \{x : f(x) < \infty\}$, and its interior $\mathcal{D}^\circ_f$. Let $\nabla f(\boldsymbol{x})$ be the gradient of f with respect to $\boldsymbol{x}$. We make the following standard assumption.

Assumption 3 For each system $i = 1, 2, \ldots, r$,

(i) the limit $\Lambda_{(G_i,H_i)}(\boldsymbol{\theta}) = \lim_{n \to \infty} \frac{1}{n} \Lambda^{(n)}_{(G_i,H_i)}(n\boldsymbol{\theta})$ exists as an extended real number for all $\boldsymbol{\theta} \in \mathbb{R}^2$;

(ii) the origin belongs to the interior of $\mathcal{D}_{\Lambda_{(G_i,H_i)}}$;

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

(iv) $\Lambda_{(G_i,H_i)}(\boldsymbol{\theta})$ is steep, that is, for any sequence $\{\boldsymbol{\theta}_n\} \in \mathcal{D}_{\Lambda_{(G_i,H_i)}}$ converging to a boundary point of $\mathcal{D}_{\Lambda_{(G_i,H_i)}}$, $\lim_{n \to \infty} |\nabla \Lambda_{(G_i,H_i)}(\theta_n)| = \infty$.

Assumption 3 implies that by the Gärtner-Ellis theorem, the probability measure governing $(\bar{G}_i(n), \bar{H}_i(n))$ obeys the large deviations principle (LDP) with good, strictly convex rate function

$$I_i(x_i, y_i) = \sup_{\theta \in \mathbb{R}^2} \{\langle \boldsymbol{\theta}, (x_i, y_i) \rangle - \Lambda_{(G_i,H_i)}(\boldsymbol{\theta})\}$$

(Dembo and Zeitouni 1998, p.44). Let

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

and let $\mathcal{F}^c_d$ denote the closure of the convex hull of the set $\{(g_i, h_i) : (g_i, h_i) \in \mathbb{R}^2, i \in \{1, 2, \ldots, r\}\}$.

Assumption 4 The closure of the convex hull of all points $(g_i, h_i) \in \mathbb{R}^2$ is a subset of the intersection of the interiors of the effective domains of the rate functions $I_i(x_i, y_i)$ for all $i = 1, 2, \ldots, r$, that is, $\mathcal{F}^c_d \subset \cap_{i=1}^{r} \mathcal{F}^\circ_{(G_i,H_i)}$.

Henceforth, for ease of notation, all vectors are column vectors. For brevity, we omit proofs of the results.

2 RATE FUNCTION DERIVATION

Consider a procedure to estimate the Pareto set that consists of expending some amount of simulation budget n to estimate the objective values for each system, and then returning the estimated set of non-dominated or Pareto systems to the user. Let us define the estimated Pareto set as

$$\widehat{\mathcal{P}} = \{\text{systems } i : \nexists \text{ system } k \in \{1,2,\ldots,r\} \text{ such that } k \,\hat{\preceq}\, i\},$$

where $k \,\hat{\preceq}\, i$ if and only if $\hat{G}_k \leq \hat{G}_i$ and $\hat{H}_k < \hat{H}_i$, or $\hat{G}_k < \hat{G}_i$ and $\hat{H}_k \leq \hat{H}_i$. Recall that $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \ldots, \alpha_r)$ is the proportional allocation of the total simulation budget n to the systems $i \leq r$, where for now, we ignore that $n\alpha_i$ is not necessarily an integer. To derive an efficient sampling method, in this section, we construct the rate function corresponding to the rate of decay of the probability of misclassification, as a function of the proportion of the sample given to each system, $\boldsymbol{\alpha}$.

In the context of our estimation procedure, a misclassification (MC) event can occur in two ways. First, misclassification by exclusion (MCE) is the event in which a truly Pareto system is falsely excluded from the estimated Pareto set by being estimated as dominated by another system, be it Pareto or non-Pareto. Second, misclassification by inclusion (MCI) is the event in which a truly non-Pareto system is falsely included in the estimated Pareto set by being estimated as non-dominated. That is,

$$\mathrm{MCE} := \underbrace{\bigcup_{i \in \mathcal{P}} \bigcup_{k \leq r,\, k \neq i} (\hat{G}_k \leq \hat{G}_i) \cap (\hat{H}_k \leq \hat{H}_i)}_{\text{some } i \in \mathcal{P} \text{ dominated by some } k} \quad \text{and} \quad \mathrm{MCI} := \underbrace{\bigcup_{j \in \mathcal{P}^c} \bigcap_{k \leq r,\, k \neq j} (\hat{G}_j \leq \hat{G}_k) \cup (\hat{H}_j \leq \hat{H}_k)}_{\text{some } j \in \mathcal{P}^c \text{ not dominated by any } k}.$$

A straightforward way of writing the MC event is $\mathrm{MC} := \mathrm{MCE} \cup \mathrm{MCI}$, and hence the probability of an MC event is $P\{\mathrm{MC}\} = P\{\mathrm{MCE} \cup \mathrm{MCI}\}$. However, this probabilistic statement is difficult to analyze, since there is dependence in the MCI term. We now reformulate the MC term for easier analysis, using what we call *phantom* Pareto systems.

2.1 Misclassification event reformulation

Let us label the true Pareto systems by the ordering of their means as

$$g_1 < g_2 < \ldots < g_{p-1} < g_p < g_{p+1} := \infty \quad \text{and} \quad h_0 := \infty > h_1 > h_2 > \ldots > h_{p-1} > h_p,$$

where p is the cardinality of $\mathcal{P}$. Then the true Pareto systems are at the coordinates (g_ℓ, h_ℓ), for $\ell = 1, \ldots, p$, where ℓ uniquely indexes the Pareto systems. To reformulate the misclassification event, we introduce the notion of *phantom* Pareto systems. The *phantom* Pareto systems are constructed by taking the g coordinate of the $(\ell+1)$th Pareto system with the h coordinate of the ℓth Pareto system. More formally, the true *phantom* Pareto systems are at the coordinates $(g_{\ell+1}, h_\ell)$ for $\ell = 0, 1, \ldots, p$, where we place phantom systems at (g_1, ∞) and (∞, h_p). There are a total of $p+1$ phantom systems. Assuming the true Pareto set is known, Figure 1 displays the Pareto systems and the phantom Pareto systems. The phantom Pareto systems allow us to rewrite the MCI term by considering the event in which non-Pareto systems falsely dominate the phantom systems, resulting in an MCI event or MCE event.

Since the locations of the Pareto systems are unknown, the phantom Pareto systems must be estimated. Define $\hat{G}_{[\ell]}$ and $\hat{H}_{[\ell]}$ as the ℓth order statistics of the estimated objectives of the true Pareto set. That is,

$$\hat{G}_{[1]} < \ldots < \hat{G}_{[p-1]} < \hat{G}_{[p]} < \hat{G}_{[p+1]} := \infty \quad \text{and} \quad \hat{H}_{[0]} := \infty > \hat{H}_{[1]} > \hat{H}_{[2]} > \ldots > \hat{H}_{[p]},$$

where the definitions of $\hat{G}_{[p+1]}$ and $\hat{H}_{[0]}$ hold for all n. Now the *estimated phantom systems for the true Pareto set* are at coordinates $(\hat{G}_{[\ell+1]}, \hat{H}_{[\ell]})$ for $\ell = 0, 1, \ldots, p$. Define misclassification by dominating a phantom system (MCI_{ph}) as

$$\mathrm{MCI}_{ph} := \underbrace{\bigcup_{j \in \mathcal{P}^c} \cup_{\ell=0}^{p} (\hat{G}_j \leq \hat{G}_{[\ell+1]}) \cap (\hat{H}_j \leq \hat{H}_{[\ell]})}_{\text{some } j \in \mathcal{P}^c \text{ dominates some phantom system}},$$

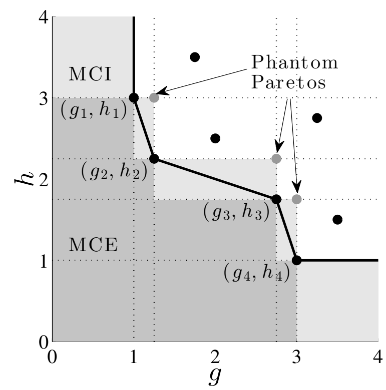

Figure 1: Suppose the true Pareto set is known. Then to be falsely estimated as Pareto without dominating any of the true Pareto systems, the non-Pareto systems must be falsely estimated as being in the light-gray MCI "region," which implies that they dominate a phantom Pareto system. Non-Pareto systems falsely estimated in the dark-gray region result in an MCE event.

and rewrite the misclassification event as $\mathrm{MC}_{ph} := \mathrm{MCE} \cup \mathrm{MCI}_{ph}$. The following Theorem 1 states the equivalence of the probability of an MC event and the probability of an MC_{ph} event. This result was first stated and proved in Hunter and McClosky (2015) in a slightly different context, but the result still holds in our context.

Theorem 1 (Hunter and McClosky 2015) $P\{\mathrm{MC}\} = P\{\mathrm{MC}_{ph}\}$.

Since Theorem 1 holds, henceforth we use the notation $P\{\mathrm{MC}\}$ to refer to the probability of a misclassification event.

2.2 Rate Function Derivation

Recall that our goal is to to identify the sampling allocation vector $\boldsymbol{\alpha}$ that maximizes the rate of decay of the $P\{\mathrm{MC}\}$. Then letting $\ell b = \max(P\{\mathrm{MCE}\}, P\{\mathrm{MCI}_{ph}\})$, it follows that $\ell b \leq P\{\mathrm{MC}\} \leq 2\ell b$, which implies

$$-\lim_{n\to\infty} \frac{1}{n} \log P\{\mathrm{MC}\} = \min\left(-\lim_{n\to\infty} \frac{1}{n} \log P\{\mathrm{MCE}\}, -\lim_{n\to\infty} \frac{1}{n} \log P\{\mathrm{MCI}_{ph}\}\right), \tag{1}$$

assuming the limit exists. (Also, recall that we have assumed $\alpha_k > 0$ for all $k = 1, \ldots, r$, so that each system has nonzero sample for the derivation of the rate functions.)

2.2.1 Rate of Decay of $P\{\mathrm{MCE}\}$

Since the rate function corresponding to the MCE term in equation (1) is most straightforward, we analyze it first in the following Lemma 1. For brevity, for all $i \in \mathcal{P}, k \leq r, k \neq i$, define the rate function

$$R_i(\alpha_i, \alpha_k) := \inf_{x_k \leq x_i,\, y_k \leq y_i} \alpha_i I_i(x_i, y_i) + \alpha_k I_k(x_k, y_k).$$

Lemma 1 The rate of decay of $P\{\mathrm{MCE}\}$ is

$$-\lim_{n\to\infty} \frac{1}{n} \log P\{\mathrm{MCE}\} = \min_{i\in\mathcal{P}} \min_{k\leq r, k\neq i} R_i(\alpha_i, \alpha_k).$$

Lemma 1 states that the rate of decay of $P\{\mathrm{MCE}\}$ is determined by the slowest rate function for the probability that a Pareto system is falsely dominated by some other system. We note that a similar result appears in Li (2012).

2.2.2 Rate of Decay of $P\{\text{MCI}_{ph}\}$

Now, consider the term corresponding to MCI_{ph} in (1). To begin, when considering the MCI_{ph} event, note that many arrangements of the true Pareto systems can occur. In addition to some system j dominating a phantom Pareto system, the Pareto systems themselves may be estimated out of "order;" hence the need for order statistics in the statement of MCI_{ph}.

To handle details of the ordering of Pareto systems, we require additional notation. Recall that we labeled the Pareto systems from $1,2,\ldots,p$. Let $\mathcal{O}=\{(1,1),(2,2),\ldots,(p,p)\}$ be an ordered list denoting the positions of the true Pareto set on each objective, where we "count" from left to right and top to bottom according to the ordering in Figure 1. For example, Pareto system 1 is in "position 1" on objective g (smallest) and "position 1" on objective h (largest), corresponding to $(1,1)$, and Pareto system 2 is in position 2 on objective g (2nd smallest) and position 2 on objective h (2nd largest), corresponding to $(2,2)$, and so on. Let $\hat{\mathcal{O}}$ denote the ordered list of estimated positions of the true Pareto set. Then the first element in the list $\hat{\mathcal{O}}$ will be $(2,6)$ if Pareto system 1 is estimated as being in the position 2 on objective g and position 6 on objective h. Let $\mathcal{S}$ denote a (fixed) realized instance of $\hat{\mathcal{O}}$, where $\mathcal{S}$ is an ordered set of p elements of the form $\{(x_1,y_1),(x_2,y_2),\ldots,(x_p,y_p)\}$ where the x and y coordinates are separately drawn without replacement from the set of Pareto indices $\{1,2,\ldots,p\}$. Define MCI_{ph} *without* order statistics as

$$\text{MCI}^*_{ph} := \bigcup_{j\in\mathcal{P}^c} \cup_{\ell=0}^{p} (\hat{G}_j \le \hat{G}_{\ell+1}) \cap (\hat{H}_j \le \hat{H}_\ell),$$

where $\hat{G}_{p+1} := \infty, \hat{H}_0 := \infty$ for all n. The following Lemma 2 breaks down the $P\{\text{MCI}_{ph}\}$ rate into rates that explicitly account for the Pareto set "ordering," and then shows that the only rate that can be the unique minimum when determining the overall rate of decay of $P\{\text{MC}\}$ is the rate corresponding to the probability that MCI^*_{ph} occurs and the Pareto ordering is estimated correctly.

Lemma 2 The rate of decay of $P\{\text{MC}\}$ is

$$-\lim_{n\to\infty}\frac{1}{n}\log P\{\text{MC}\} = \min\left(-\lim_{n\to\infty}\frac{1}{n}\log P\{\text{MCE}\}, -\lim_{n\to\infty}\frac{1}{n}\log P\{\text{MCI}^*_{ph}\cap\hat{\mathcal{O}}=\mathcal{O}\}\right). \tag{2}$$

Loosely speaking, the result in Lemma 2 makes sense: the event that a false inclusion occurs *and* the true Pareto set is estimated "out of order" is, intuitively, less likely than the event that a false inclusion occurs while the true Pareto set is estimated "in order." The proof proceeds by bounding the rate of decay of all other events below by the rate of decay of $P\{\text{MCE}\}$. In the following Theorem 2, we present the overall rate of decay of $P\{\text{MC}\}$, where for brevity in Theorem 2 and the remainder of the paper, for all $j\in\mathcal{P}^c, \ell=0,\ldots,p$, define the rate function

$$R_j(\alpha_j,\alpha_\ell,\alpha_{\ell+1}) := \begin{cases} \inf\limits_{x_j\le x_1} & \alpha_j I_j(x_j,y_j)+\alpha_1 I_1(x_1,y_1) & \text{if } \ell=0 \\ \inf\limits_{x_j\le x_{\ell+1},\, y_j\le y_\ell} & \alpha_j I_j(x_j,y_j)+\alpha_\ell I_\ell(x_\ell,y_\ell)+\alpha_{\ell+1}I_{\ell+1}(x_{\ell+1},y_{\ell+1}) & \text{if } \ell\in\{1,\ldots,p-1\} \\ \inf\limits_{y_j\le y_p} & \alpha_j I_j(x_j,y_j)+\alpha_p I_p(x_p,y_p) & \text{if } \ell=p, \end{cases}$$

where $\alpha_0 := 1$ and $\alpha_{p+1} := 1$. This rate function corresponds to the rate of decay of the probability of a subset of events contained in $\text{MCI}^*_{ph}\cap\hat{\mathcal{O}}=\mathcal{O}$ (from Lemma 2). Then Theorem 2 follows.

Theorem 2 The rate of decay of the probability of misclassification is

$$-\lim_{n\to\infty}\frac{1}{n}\log P\{\text{MC}\} = \min\left(\min_{i\in\mathcal{P}}\min_{k\in\mathcal{P},k\ne i} R_i(\alpha_i,\alpha_k), \min_{j\in\mathcal{P}^c}\min_{\ell=0,\ldots,p} R_j(\alpha_j,\alpha_\ell,\alpha_{\ell+1})\right).$$

Theorem 2 states that the overall rate of decay of $P\{\text{MC}\}$ is the minimum of two rates. The first rate corresponds to the minimum pairwise rate of decay of the probability that Pareto systems falsely exclude each other (MCE), and the second rate corresponds to the minimum "pairwise" rate of decay of the probability that a non-Pareto system falsely dominates a phantom Pareto system (MCI_{ph}). Note that the rate of decay of the probability that a Pareto system is falsely excluded by a non-Pareto system is accounted for in the second rate corresponding to MCI_{ph}, since a non-Pareto system falsely excludes a Pareto system if and only if it also falsely excludes a phantom Pareto system.

2.3 OPTIMAL ALLOCATION STRATEGY

Since $R_i(\alpha_i,\alpha_k)$ and $R_j(\alpha_j,\alpha_\ell,\alpha_{\ell+1})$ are concave functions of (α_i,α_k) and $(\alpha_j,\alpha_\ell,\alpha_{\ell+1})$, respectively, for all $i,k\in\mathcal{P},\ j\in\mathcal{P}^c$, and $\ell=0,1,\ldots,p$, the asymptotically optimal sample allocation can be expressed as the solution to the concave maximization problem

$$\begin{aligned}\text{Problem } Q: \quad & \text{maximize } z \text{ s.t.}\\ & R_i(\alpha_i,\alpha_k)\ge z \text{ for all } i,k\in\mathcal{P},k\ne i,\\ & R_j(\alpha_j,\alpha_\ell,\alpha_{\ell+1})\ge z \text{ for all } j\in\mathcal{P}^c,\ell=0,\ldots,p,\\ & \textstyle\sum_{i=1}^r \alpha_i=1,\ \alpha_i\ge 0 \text{ for all } i\le r,\end{aligned}$$

where, for some $i,k\in\mathcal{P},\ k\ne i$ and a given value of (α_i,α_k), the value of $R_i(\alpha_i,\alpha_k)$ is obtained by solving

$$\text{Problem } R_{ik}^{\mathcal{P}}: \quad \text{minimize } \alpha_i I_i(x_i,y_i)+\alpha_k I_k(x_k,y_k) \text{ s.t. } x_k\le x_i,\ y_k\le y_i,$$

and for some $j\in\mathcal{P}^c,\ell\in\{0,\ldots,p\}$, and a given value of $(\alpha_j,\alpha_\ell,\alpha_{\ell+1})$, the value of $R_j(\alpha_j,\alpha_\ell,\alpha_{\ell+1})$ is obtained by solving

$$\begin{aligned}\text{Problem } R_{j\ell}: \quad & \text{minimize } \alpha_j I_j(x_j,y_j)+\alpha_\ell I_\ell(x_\ell,y_\ell)\mathbb{I}_{[\ell\ne 0]}+\alpha_{\ell+1}I_{\ell+1}(x_{\ell+1},y_{\ell+1})\mathbb{I}_{[\ell\ne p]}\\ & \text{s.t. } (x_j-x_{\ell+1})\mathbb{I}_{[\ell\ne p]}\le 0,\ (y_j-y_\ell)\mathbb{I}_{[\ell\ne 0]}\le 0.\end{aligned}$$

Recall that Problems $R_{ik}^{\mathcal{P}}$ and $R_{j\ell}$ are strictly convex minimization problems for which the KKT conditions are necessary and sufficient.

Problem Q has $p\times(p-1)$ constraints corresponding to controlling the rate of decay of $P\{\mathrm{MCE}\}$ and $(r-p)\times(p+1)$ constraints corresponding to controlling the rate of decay of $P\{\mathrm{MCI}_{ph}\}$. Slater's condition (Boyd and Vandenberghe 2004) holds for this problem, so the KKT conditions are necessary and sufficient for global optimality.

We remind the reader that to solve Problem Q and obtain the asymptotically optimal sampling allocation, we would have to know the entire rate function. Since the rate function is unknown, the optimal allocation cannot be implemented as written, but must be incorporated into some type of sequential sampling framework, as in Hunter and Pasupathy (2013), Pasupathy et al. (2014), Hunter and McClosky (2015). Broadly, the philosophy in these papers is to assume knowledge of the distributional family, e.g., normal, and use plug-in estimators for the distributional parameters. While this approach is somewhat theoretically controversial (Glynn and Juneja 2011), it is numerically successful; in particular, see Pasupathy et al. (2014), Hunter and McClosky (2015).

3 CONCLUDING REMARKS

We have characterized the asymptotically optimal sampling allocation for bi-objective R&S problems, fully accounting for dependence between the objectives. Since the computational complexity of Problem Q increases with the number of systems, ongoing work includes obtaining an easily-implementable SCORE allocation for the case when the number of systems is large.

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.

Butler, J. C., D. J. Morrice, and P. Mullarkey. 2001. "A Multiple Attribute Utility Theory Approach to Ranking and Selection". *Management Science* 47 (6): 800–816.

Chen, C.-H., J. Lin, E. Yücesan, and S. E. Chick. 2000. "Simulation Budget Allocation for Further Enhancing the Efficiency of Ordinal Optimization". *Discrete Event Dynamic Systems* 10 (3): 251–270.

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

Feldman, G., S. R. Hunter, and R. Pasupathy. 2015. "Multi-Objective Simulation Optimization on Finite Sets: Optimal Allocation via Scalarization". In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti. Piscataway, NJ: Institute of Electrical and Electronics Engineers, Inc.

Glynn, P. W., and S. Juneja. 2004. "A Large Deviations Perspective on Ordinal Optimization". In *Proc. of the 2004 Winter Simulation Conference*, edited by R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 577–585. Piscataway, NJ: Institute of Electrical and Electronics Engineers, Inc.

Glynn, P. W., and S. Juneja. 2011. "Ordinal Optimization: A Nonparametric Framework". In *Proc. of the 2011 Winter Simulation Conference*, edited by S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, and M. Fu. Piscataway, NJ: Institute of Electrical and Electronics Engineers, Inc.

Huang, H., and Z. B. Zabinsky. 2014. "Multiple Objective Probabilistic Branch and Bound for Pareto Optimal Approximation". In *Proc. of the 2014 Winter Simulation Conference*, edited by A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, 3916–3927. Piscataway, NJ: 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 B. McClosky. 2015. "Maximizing Quantitative Traits in the Mating Design Problem via Simulation-Based Pareto Estimation". *Under 3rd review for IIE Transactions*.

Hunter, S. R., and R. Pasupathy. 2013. "Optimal Sampling Laws for Stochastically Constrained Simulation Optimization on Finite Sets". *INFORMS Journal on Computing* 25 (3): 527–542.

Kim, S., and J. Ryu. 2011a. "The Sample Average Approximation Method for Multi-Objective Stochastic Optimization". In *Proc. of the 2011 Winter Simulation Conference*, edited by S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, and M. Fu, 4026–4037. Piscataway, NJ: Institute of Electrical and Electronics Engineers, Inc.

Kim, S., and J. Ryu. 2011b. "A Trust-Region Algorithm for Bi-Objective Stochastic Optimization". *Procedia Computer Science* 4:1422–1430.

Kim, S.-H., and B. L. Nelson. 2006. "Selecting the Best System". In *Simulation*, edited by S. G. Henderson and B. L. Nelson, Handbooks in Operations Research and Management Science, Volume 13, 501–534. Elsevier.

Lee, L. H., E. P. Chew, S. Teng, and D. Goldsman. 2010. "Finding the Non-Dominated Pareto Set for Multi-Objective Simulation Models". *IIE Transactions* 42:656–674.

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". *IEEE Transactions on Automatic Control* 57 (11): 2940–2945.

Li, H., L. H. Lee, E. P. Chew, and P. Lendermann. 2015. "MO-COMPASS: A Fast Convergent Search Algorithm for Multi-Objective Discrete Optimization via Simulation". *IIE Transactions*.

Li, J. 2012. *Optimal Computing Budget Allocation for Multi-Objective Simulation Optimization*. Ph. D. thesis, National University of Singapore.

Pasupathy, R., and S. Ghosh. 2013. "Simulation Optimization: A Concise Overview and Implementation Guide". In *TutORials in Operations Research*, edited by H. Topaloglu, Chapter 7, 122–150. INFORMS.

Pasupathy, R., S. R. Hunter, N. A. Pujowidianto, L. H. Lee, and C. Chen. 2014. "Stochastically Constrained Ranking and Selection via SCORE". *ACM Transactions on Modeling and Computer Simulation* 25 (1): 1–26.

Ryu, J., S. Kim, and H. Wan. 2009. "Pareto Front Approximation with Adaptive Weighted Sum Method in Multi-Objective Simulation Optimization". In *Proc. of the 2009 Winter Simulation Conference*, edited by M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin, and R. G. Ingalls, 623–633: Institute of Electrical and Electronics Engineers, Inc.

Szechtman, R., and E. Yücesan. 2008. "A New Perspective on Feasibility Determination". In *Proc. of the 2008 Winter Simulation Conference*, edited by S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, and J. W. Fowler, 273–280. Piscataway, NJ: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHIES

SUSAN R. HUNTER is an assistant professor in the School of Industrial Engineering at Purdue University. Her research interests include Monte Carlo methods and simulation optimization, especially in the presence of multiple performance measures. Her e-mail address is susanhunter@purdue.edu, and her website is http://web.ics.purdue.edu/~hunter63/.

GUY FELDMAN is a Ph.D. student in the Department of Statistics at Purdue University. His research interests include simulation optimization and machine learning. His email address is gfeldman@purdue.edu. His website is http://www.stat.purdue.edu/~gfeldman/.