Proceedings of the 2008 Winter Simulation Conference
S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, J. W. Fowler eds.

A NEW PERSPECTIVE ON FEASIBILITY DETERMINATION

Roberto Szechtman

Department of Operations Research
1411 Cunningham Rd.
Naval Postgraduate School
Monterey, CA 93943, U.S.A.

Enver Yücesan

Technology and Operations Management Area
INSEAD
Boulevard de Constance
77305 Fontainebleau, France

ABSTRACT

We consider the problem of feasibility determination in a stochastic setting. In particular, we wish to determine whether a system belongs to a given set Γ based on a performance measure estimated through Monte Carlo simulation. Our contribution is two-fold: (i) we characterize fractional allocations that are asymptotically optimal; and (ii) we provide an easily implementable algorithm, rooted in stochastic approximation theory, that results in sampling allocations that provably achieve in the limit the same performance as the optimal allocations. The finite-time behavior of the algorithm is also illustrated on two small examples.

1 BACKGROUND

We consider the problem of feasibility determination in a stochastic setting. In particular, we wish to determine whether a system belongs to a given set Γ based on a performance measure estimated through Monte Carlo simulation. While it is an interesting problem in its own right, feasibility determination has recently attracted much attention from researchers in ranking and selection (R&S) whose objective is to select the best system in the presence of a stochastic constraint.

Ranking and selection (R&S) techniques are statistical methods developed to select the best system, or a subset of systems from among a set of alternative system designs. R&S via simulation is particularly appealing as it combines modeling flexibility of simulation with the efficiency of statistical techniques for effective decision making. Furthermore, it is relatively straightforward to satisfy the underlying technical assumptions of these techniques in simulation experiments, which also allow for multi-stage sampling as required by some R&S methods.

Due to randomness in output data, comparing a number of simulated systems requires care. Suppose we conduct n simulation replications for each of r designs. Therefore, we need rn simulation replications. Simulation results become more precise as n increases. If the precision requirement is high (n is not small), and if the total number of designs in a decision problem is large (r is large), then rn can be very large, which may easily render the total simulation cost prohibitively high and preclude the feasibility of using simulation for R&S problems. The effective employment of the simulation budget in the course of obtaining a good decision is therefore crucial.

The rich literature on R&S has been developing along two principal axes. The optimal computing budget allocation (OCBA) approach has its roots in Dudewicz and Dalal (1975) who developed a two-stage procedure for selecting the best design or a design that is very close to the best system. Rinott (1978) presents an alternative way to compute the required number of simulation replications in the second stage. This idea has been ultimately extended to fully sequential algorithms. Paulson (1964), Gupta and Panchapakesan (1979), Nelson and Matejcik (1995), Matejcik and Nelson (1995), Hsu (1996), and Kim and Nelson (2006) present methods based on the classical statistical model adopting a frequentist view. On the other hand, the value of information (VIP) approach, exemplified by Berger (1985), Bernardo and Smith (1994), Gupta and Miescke (1996), Chick and Inoue (2001), and Chen et al. (2000) uses a Bayesian framework for managing the trade-off between the consequences of an immediate decision and the cost of additional sampling.

The overwhelming majority of the R&S research focuses on a single unconstrained performance measure (Kim and Nelson 2006). Most applications, however, either have multiple performance measures (Butler, Morrice, and Mullarkey 2001) or face some constraints on the primary performance measure (Andradottir, Goldsman, and Kim 2005). In a manufacturing setting, for example, one might be simultaneously interested in selecting the scheduling policy that maximizes throughput, with a physical limit on the work in process or a service constraint on lead times. While some of these measures might be correlated, others might simply conflict with one another. The literature dealing with

978-1-4244-2708-6/08/$25.00 ©2008 IEEE

multiple performance measures or stochastic constraints has been sparse.

In the case with multiple performance measures, only a partial ordering may be possible, i.e., these populations may not be ordered in terms of a vector-valued parameter. In such cases, one typically defines a real-valued function of the parameters and use this function to rank the populations. Gupta et al. (1973) propose several such scalar functions. Andijani (1998) uses the analytic hierarchical process (AHP) and Butler et al. (2001) use the multiple attribute utility theory (MAUT) to construct a scalar function by assigning weights to the performance measures. Kim and Lin (2001) use a *max-min* approach whereby they maximize the poorest performing criterion. Lee et al. (2006) and Lee et al. (2007) seek to identify a Pareto set of non-dominated designs.

R&S in the presence of stochastic constraints has recently started receiving some attention. In a setting where R&S is based on a primary performance measure subject to the feasibility of a (possibly correlated) secondary performance measure, Andradottir et al. (2005) propose a two-phase approach : phase I identifies feasible systems while phase II selects the best among them. With the objective of accelerating the first phase, Batur and Kim (2005) have recently introduced procedures for finding a set of feasible or near-feasible systems. The present article also is concerned with efficient feasibility determination. Unlike the previous work that is based on the indifference zone perspective, we use large deviations theory to minimize the expected number of incorrect determinations. We characterize the optimal employment of the simulation budget to sampling from each system, and present a stochastic approximation algorithm that yields a budget allocation that provably converges to the optimal one. It is worth pointing out that we are not the first ones using the large deviations framework for R&S. Glynn and Juneja (2004) have used this framework not only to determine an optimal budget allocation strategy, but to also demonstrate the deteriorating effect of violating the normality assumptions, which are very commonly used in the R&S literature. Our notation closely follows theirs. Like them, we characterize the optimal allocation of the computational budget. However, we also provide an algorithm to achieve the optimal sampling allocations.

The remainder of the paper is organized as follows. Section 2 introduces the necessary notation along with some preliminary ideas from the large deviations theory. Section 3 introduces our algorithm for feasibility determination. A numerical illustration is presented in Section 4. Section 5 concludes the paper.

2 OPTIMAL ASYMPTOTIC ALLOCATION

We consider r systems, each with unknown performance measure $\mu_1,\ldots,\mu_r \in R$. Given a set $\Gamma=[\gamma,\infty)$, we wish to employ Monte Carlo simulation to determine for each system i whether $\mu_i \in \Gamma$ or not. We assume the simulationist can obtain i.i.d. replicates of the random variable X_i to form unbiased estimators $\bar{X}_i(n_i)=n_i^{-1}\sum_{k=1}^{n_i} X_{i,k}$ of μ_i, for $i=1,\ldots,r$.

The simulation budget n is allocated in order to minimize the expected number of incorrect determinations. If, without loss of generality, we assume that $\mu_1,\ldots,\mu_a \in \Gamma$ and $\mu_{a+1},\ldots,\mu_r \notin \Gamma$, then the feasibility determination problem is to

$$\min_{p_1,\ldots,p_r\in\mathscr{M}} g_n(p_1,\ldots,p_r)$$

where

$$g_n(p_1,\ldots,p_r)=\sum_{i=1}^{a} P(\bar{X}_i(p_i n)\notin\Gamma)+\sum_{i=a+1}^{r} P(\bar{X}_i(p_i n)\in\Gamma),$$

and $\mathscr{M}=\{p_i\geq 0:\sum_i p_i\leq 1\}$. The p_i's represent the fraction of the simulationist's budget that is allocated to sampling from each system, where for simplicity we assume that each system has the same per-sample cost.

Our contribution is two-fold: (i) we characterize fractional allocations $p_1^*,\ldots,p_r^*$ that are optimal (in log scale) as $n\to\infty$; and (ii), we provide an easily implementable algorithm, rooted in stochastic approximation theory, that results in sampling allocations that provably achieve the same performance (in log scale) as the optimal allocations in the limit as $n\to\infty$.

Let $I_i(\cdot)$ be the large deviations rate function

$$I_i(x)=\sup_{\theta\in R}\{\theta x-\log M_i(\theta)\}, \tag{1}$$

where

$$M_i(\theta)=E\exp(\theta X_i)$$

is the moment generating function, for each system $i=1,\ldots,r$. Define $\mathscr{D}_i=\{\theta\in R:\log M_i(\theta)<\infty\}$, and $\mathscr{D}_i^o$ the interior of $\mathscr{D}_i$. Our main assumptions are that:

A1. The performance measures do not lie exactly at the boundary: $\mu_i\neq\gamma$ for all $i=1,\ldots,r$.

A2. For each system, there exists $\theta_i^*\in\mathscr{D}_i^o$ such that $(\log M_i(\theta_i^*))'=\gamma$.

Assumption A1 ensures that the rate functions $I_i(\cdot)$ evaluated at the boundary point γ are positive, so that no system requires all the simulation budget. Assumption A2 is standard in situations where the underlying distributions are light tailed, and implies that for each system there exists $\theta_i^*\in R$ (unique) such that $I_i(\gamma)=\theta_i^*\gamma-M_i(\theta_i^*)$. It can be shown that $\theta_i^*<0$ for $i\in\{1,\ldots,a\}$ and $\theta_i^*>0$ for $i\in\{a+1,\ldots,r\}$.

We are now ready to state the main results of this section.

Proposition 1 *Suppose Assumption A2 holds. Then*

$$\lim_{n\to\infty}\frac{1}{n}\log g_n(p_1,\ldots,p_r) = -\min_i p_i I_i(\gamma). \qquad (2)$$

Proof: The proof consists of two parts. For $p_i > 0$ and θ scalar

$$\frac{1}{n}\log E\left(\exp(n\theta\bar{X}_i(p_i n))\right) = p_i \log M_i(\theta/p_i).$$

Moreover,

$$\begin{aligned}\sup_{\theta\in R}\{\gamma\theta - p_i\log M_i(\theta/p_i)\} &= p_i \sup_{\theta\in R}\{\gamma\theta/p_i - \log M_i(\theta/p_i)\}\\ &= p_i I_i(\gamma).\end{aligned}$$

Hence, the Gärtner-Ellis Theorem implies that $1/n\log P(\bar{X}_i(np_i) \notin \Gamma) \to -p_i I_i(\gamma)$ for $i = 1,\ldots,a$, and $1/n\log P(\bar{X}_i(np_i) \in \Gamma) \to -p_i I_i(\gamma)$ for $i = a+1,\ldots,r$, as $n\to\infty$.

In the second part we obtain lower and upper bounds for $(1/n)\log g_n(p_1,\ldots,p_r)$. Suppose that $p_{i^*}I_{i^*} = \min_i p_i I_i(\gamma)$. Then, since $g_n(p_1,\ldots,p_r) \geq P(\bar{X}_{i^*}(np_{i^*}) \notin \Gamma)$ if $i^* \in \{1,\ldots,a\}$ and $g_n(p_1,\ldots,p_r) \geq P(\bar{X}_{i^*}(np_{i^*}) \in \Gamma)$ if $i^* \in \{a+1,\ldots,r\}$, we obtain the lower bound $\liminf_{n\to\infty}(1/n)\log g_n(p_1,\ldots,p_r) \geq -p_{i^*}I_{i^*}(\gamma)$. For the upper bound,

$$\begin{aligned}\frac{1}{n}\log g_n(p_1,\ldots,p_r) \leq \frac{1}{n}\max\{&\log(rP(\bar{X}_1(np_1)\notin\Gamma)),\ldots,\\ &\log(rP(\bar{X}_a(np_a)\notin\Gamma)),\log(rP(\bar{X}_1(np_{a+1})\in\Gamma)),\ldots,\\ &\log(rP(\bar{X}_1(np_r)\in\Gamma))\}.\end{aligned}$$

Taking limits,

$$\limsup_{n\to\infty}\frac{1}{n}\log g_n(p_1,\ldots,p_r) \leq -p_{i^*}I_{i^*}(\gamma).$$

$\otimes$

The best achievable exponential decay rate for $g_n(\cdot)$ is described next.

Proposition 2 *Suppose assumptions A1 and A2 hold. Then*

$$\lim_{n\to\infty}\frac{1}{n}\log g_n(p_1^*,\ldots,p_r^*) = -\frac{1}{\sum_{j=1}^r I_j^{-1}(\gamma)}$$

where

$$p_i^* = \frac{I_i^{-1}(\gamma)}{\sum_{j=1}^r I_j^{-1}(\gamma)}. \qquad (3)$$

For any other $(p_1,\ldots,p_r) \in \mathscr{M}$ we have

$$\liminf_{n\to\infty}\frac{1}{n}\log g_n(p_1,\ldots,p_r) \geq -\frac{1}{\sum_{j=1}^r I_j^{-1}(\gamma)}.$$

Proof: The first statement follows immediately by evaluating Eq. (2) with p_i^*. For the second statement we need to show that $p_1^*,\ldots,p_r^*$ maximize the exponential decay rate of $g_n(\cdot)$, given in Eq.(2). Hence we consider the problem $\max_{p\in\mathscr{M}}\min_i p_i I_i(\gamma)$, which is equivalent to the linear problem

$$\begin{aligned}&\max \eta\\ \text{st}\quad &\eta - p_i I_i(\gamma) \leq 0 \text{ for } i = 1,\ldots,r\\ &\sum_{j=1}^r p_j \leq 1\\ &p_i \geq 0 \text{ for } i = 1,\ldots,r\end{aligned}$$

The dual of this problem is

$$\begin{aligned}&\min \pi\\ \text{st}\quad &\pi - q_i I_i(\gamma) \geq 0 \text{ for } i = 1,\ldots,r\\ &\sum_{j=1}^r q_j \geq 1\\ &\pi, q_i \geq 0 \text{ for } i = 1,\ldots,r\end{aligned}$$

The variables $p_i = q_i = I_i^{-1}(\gamma)/\sum_{j=1}^r I_j^{-1}(\gamma)$, $\eta = \pi = 1/\sum_{j=1}^r I_j^{-1}(\gamma)$ are primal and dual feasible and $\eta = \pi$, so that p_i^* is primal optimal by weak duality. $\otimes$

We view p_i^* as the optimal allocation scheme, meaning that no other allocation achieves a higher exponential decay rate for $g_n(\cdot)$ as the sampling budget goes to infinity.

Example 1 *Suppose for every system we have $X_i \sim N(\mu_i,\sigma_i^2)$, with $\sigma_i^2 > 0$. For each system i (see Exercise 2.2.23 of Dembo and Zeitouni 1998),*

$$I_i(\gamma) = \frac{1}{2}\left(\frac{\mu_i-\gamma}{\sigma_i}\right)^2.$$

Hence Proposition 2, specialized to the Normal case, becomes

$$p_i^* = \frac{\sigma_i^2/(\mu_i-\gamma)^2}{\sum_{j=1}^r \sigma_j^2/(\mu_j-\gamma)^2}. \qquad (4)$$

In other words, when the underlying distributions are normal, each system is sampled in proportion to the square of the distance of its scaled performance measure to the boundary point γ.

Example 2 *Assume $X_i \sim Bernoulli(\mu_i)$, $0 < \mu_i < 1$, for $i = 1,\ldots,r$, and let $0 < \gamma < 1$. The rate functions are*

(Exercise 2.2.23 of Dembo and Zeitouni 1998),

$$I_i(\gamma) = \gamma \log\left(\frac{\gamma}{\mu_i}\right) + (1-\gamma)\log\left(\frac{1-\gamma}{1-\mu_i}\right),$$

so that Proposition 2 results in

$$p_i^* = \frac{\left[\gamma\log\left(\frac{\gamma}{\mu_i}\right) + (1-\gamma)\log\left(\frac{1-\gamma}{1-\mu_i}\right)\right]^{-1}}{\sum_{j=1}^r \left[\gamma\log\left(\frac{\gamma}{\mu_j}\right) + (1-\gamma)\log\left(\frac{1-\gamma}{1-\mu_j}\right)\right]^{-1}}. \quad (5)$$

While Proposition 2 characterizes the optimal sampling allocation, the rate functions $I_i(\gamma)$ are generally unknown and need to be estimated. However, estimating each rate function involves solving a root finding problem, as can be seen in Eq. (1) by solving for the derivative with respect to θ equal to zero. This suggests that the computational cost of sequentially estimating the rate functions dominates all other costs when running the simulation, and therefore is not a promising course of action. In the next section we present a stochastic approximation algorithm that overcomes these issues, and leads to fractional allocations that converge almost surely to the optimal p_i^* allocations.

3 STOCHASTIC APPROXIMATION ALGORITHM

Initially we warm-up each system with $n_0 > 0$ replicates and compute the sample averages $\tilde{X}_{i,0} = n_0^{-1}\sum_{k=1}^{n_0} X_{i,k}$. Let

$$\begin{aligned} I_{i,0} &= \sup_{\theta\in R}\left\{\theta\gamma - \log\left(\frac{1}{n_0}\sum_{k=1}^{n_0}\exp(\theta X_{i,k})\right)\right\} \\ &= \theta_{i,0}\gamma - \log\left(M_{i,0}\right), \end{aligned}$$

where $\theta_{i,0}$ satisfies the root problem

$$\gamma = \frac{\sum_{k=1}^{n_0} X_{i,k}\exp(\theta_{i,0}X_{i,k})}{\sum_{k=1}^{n_0}\exp(\theta_{i,0}X_{i,k})},$$

and the sample moment generating functions $M_{i,0}$ evaluated at $\theta_{i,0}$ are given by

$$M_{i,0} = \frac{1}{n_0}\sum_{k=1}^{n_0}\exp(\theta_{i,0}X_{i,k}).$$

Finally, we initialize the sample sizes, $\lambda_{i,0} = n_0$ for $i = 1,\ldots,r$.

Algorithm.

1. Initialize $n = 0$.
2. Generate a replicate α from the p.m.f. $I_{i,n}^{-1}/\sum_{j=1}^r I_{j,n}^{-1}$, for $i = 1,\ldots,r$.
3. Update sample sizes: $\lambda_{\alpha,n+1} = \lambda_{\alpha,n} + 1$, and $\lambda_{i,n+1} = \lambda_{i,n}$ for $i \neq \alpha$.
4. Generate a replicate from system α, (abusing notation, say) X_{n+1}.
5. Update $\tilde{X}_{\alpha,n+1}$, $M_{\alpha,n+1}$, $\theta_{\alpha,n+1}$, and $I_{\alpha,n+1}$:

$$\tilde{X}_{\alpha,n+1} = \tilde{X}_{\alpha,n} + \frac{1}{\lambda_{\alpha,n+1}}(X_{n+1} - \tilde{X}_{\alpha,n}),$$

$$M_{\alpha,n+1} = M_{\alpha,n} + \frac{1}{\lambda_{\alpha,n+1}}(\exp(\theta_{\alpha,n}X_{n+1}) - M_{\alpha,n}),$$

$$\begin{aligned} \theta_{\alpha,n+1} &= \theta_{\alpha,n} \\ &\quad - \frac{1}{\lambda_{\alpha,n+1}}(X_{n+1}-\gamma)\exp\left(\theta_{\alpha,n}(X_{n+1}-\gamma)\right), \end{aligned}$$

and

$$I_{\alpha,n+1} = \theta_{\alpha,n+1}\gamma - \log(M_{\alpha,n+1}),$$

if $\theta_{\alpha,n+1}\gamma - \log(M_{\alpha,n+1}) > 0$, and $I_{\alpha,n+1} = \min_j I_{j,n}$ otherwise. (This ensures that $I_{\alpha,n+1} > 0$ always.) For $i \neq \alpha$, set $\tilde{X}_{i,n+1} = \tilde{X}_{i,n}$, $M_{i,n+1} = M_{i,n}$, $\theta_{i,n+1} = \theta_{i,n}$, and $I_{i,n+1} = I_{i,n}$.

6. Increase $n \leftarrow n+1$ and go back to 2.

To ensure that each system is sampled infinitely often, let (ν_n) be an increasing sequence such that $\nu_n \to \infty$ and $n^{-1}\sum_{k=1}^n J(\nu_k \leq n) \to 0$, where $J(\cdot)$ is the indicator function. We sample from all r systems at stages $\nu_1, \nu_2, \ldots$, and update the parameters according to steps 3 and 5 of the algorithm.

This algorithm provides fractional allocations that converge a.s. to the optimal allocations, as the next theorem demonstrates.

Theorem 1 *Under assumptions A1 and A2,*

$$\frac{\lambda_{i,n}}{n} \to p_i^*,$$

almost surely as $n \to \infty$.

Proof outline: Suppose momentarily that $\theta_{i,n} \to \theta_i^*$ a.s. Let $p_{i,n} = \lambda_{i,n}/n$, so that step 3 of the algorithm can be expressed as $p_{i,n+1} = p_{i,n} + (J(\alpha_n = i) - p_{i,n})/(n+1)$, where α_n is the nth replicate of α generated in step 2 of the algorithm, and $J(\cdot)$ is the indicator function. The recursion for $p_{i,n+1}$ can be re-written as

$$\begin{aligned} p_{i,n+1} = p_{i,n} + \frac{1}{n+1}(p_i^* - p_{i,n}) + \frac{1}{n+1}(J(\alpha_n = i) - q_{i,n}) \\ + \frac{1}{n+1}(q_{i,n} - p_i^*), \quad (6) \end{aligned}$$

where $q_{i,n} = I_{i,n}^{-1} / \sum_{j=1}^{r} I_{j,n}^{-1}$. Since $E[J(\alpha_n = i) | I_{j,0}, \alpha_k, X_k, j = 1,\ldots,r, k = 0,\ldots,n] = q_{i,n}$, the sequence $\{J(\alpha_n = i) - q_{i,n}\}$ is a martingale difference with respect to the sequence of σ-algebras generated by $\{I_{j,0}, \theta_{j,0}, \alpha_k, X_k, j = 1,\ldots,r, k = 0,\ldots,n\}$. Additionally,

$$\sum_k \frac{|q_{i,k} - p_i^*|}{k+1} \leq \sum_k \frac{J(|q_{i,k} - p_i^*| > 1/(1+k))}{k+1} + \sum_k \frac{1}{(k+1)^2}, \quad (7)$$

almost surely. The set $\{\sum_k J(|q_{i,k} - p_i^*| > 1/(1+k))/(k+1) = \infty\} \subseteq \cup_{\ell=0}^{\infty} \{\sum_k J(|q_{i,k} - p_i^*| > 1/(1+\ell))/(k+1) = \infty\}$, and $P(\sum_k J(|q_{i,k} - p_i^*| > 1/(1+\ell))/(k+1) = \infty) = 0$ for each ℓ since $\theta_{i,n} \to \theta_i^*$ a.s. implies $q_{i,n} \to p_i^*$ a.s. Therefore $P(\sum_k J(|q_{i,k} - p_i^*| > 1/(1+k))/(k+1) = \infty) = 0$, and we conclude from Eq. (7), that $\sum_k |q_{i,k} - p_i^*|/(k+1) < \infty$ a.s.

Hence, all the assumptions of Theorem 5.2.1 of Kushner and Yin (2003) are satisfied. Moreover, the ordinary differential equations $p_i' = p_i^* - p_i$, $i = 1,\ldots,r$, have a unique globally asymptotically stable point, p_i^*, so that Theorem 5.2.1 of Kushner and Yin (2003) results in $p_{i,n} \to p_i^*$ a.s. (The special instances of $n = \nu_k, k = 1,2,\ldots$, can be handled by introducing an extra term in Eq. (6), and do not affect the asymptotic behavior of $p_{i,n}$.)

Regarding the asymptotic behavior of $\theta_{i,n}$, our algorithm guarantees that $\lambda_{i,n} \to \infty$ w.p.1, so that we can focus on the epochs when a system is sampled. Re-write the recursion for $\theta_{\alpha,n+1}$ as

$$\theta_{\alpha,n+1} = \theta_{\alpha,n} + \frac{1}{\lambda_{\alpha,n+1}} \times \left(\bar{g}_\alpha(\theta_{\alpha,n}) - (X_{n+1} - \gamma) \exp\left(\theta_{\alpha,n}(X_{n+1} - \gamma)\right) \right) - \frac{\bar{g}_\alpha(\theta_{\alpha,n})}{\lambda_{\alpha,n+1}},$$

where $\bar{g}_i(\theta_i) = E[(X_i - \gamma)\exp(\theta_i(X_i - \gamma))]$. Since $\log M_i(\theta_i)$ is strictly convex and $(\log M_i(\theta_i^*))' = \gamma$, we have $\bar{g}_i(\theta_i) < 0$ for $\theta_i < \theta_i^*$, $\bar{g}_i(\theta_i) > 0$ for $\theta_i > \theta_i^*$, and $\bar{g}_i(\theta_i) = 0$ for $\theta_i = \theta_i^*$.

The sequence $\{\bar{g}_\alpha(\theta_{\alpha,n}) - (X_{n+1} - \gamma)\exp\left(\theta_{\alpha,n}(X_{n+1} - \gamma)\right)\}$ is a martingale difference with respect to the sequence of σ-algebras generated by $\{I_{j,0}, \theta_{j,0}, \alpha_k, X_k, j = 1,\ldots,r, k = 0,\ldots,n\}$. Since $\lambda_{i,n} \to \infty$ a.s. for each system i, Theorem 5.2.1 of Kushner and Yin (2003) applies, meaning that $\theta_{i,n}$ converges a.s. to the unique globally asymptotically stable point of the ODE

$$\theta_i' = -\bar{g}_i(\theta_i),$$

which is given by θ_i^*, for $i = 1,\ldots,r$. $\otimes$

Remark. $p_i^* I_i(\gamma) = p_j^* I_j(\gamma)$ suggests that step 2 of the algorithm could be replaced by $\alpha = \arg\min_i \{\lambda_{i,n} I_{i,n}\}$. An argument similar to the one employed in the proof of Theorem 1 shows that this approach also results in asymptotically optimal fractional allocations.

The next result shows that the expected number of incorrectly determined systems produced by the stochastic approximation algorithm approaches 0 at the best possible rate.

Theorem 2 *Suppose that assumptions A1 and A2 hold. Then,*

$$\frac{1}{n} \log \left(\sum_{i=1}^{a} P(\tilde{X}_{i,n} \notin \Gamma) + \sum_{i=a+1}^{r} P(\tilde{X}_{i,n} \in \Gamma) \right) \to -\frac{1}{\sum_{j=1}^{r} I_j^{-1}(\gamma)},$$

as $n \to \infty$.

Proof outline: The proof resembles the proof of Proposition 1. First of all, since for each system the samples are i.i.d., we obtain

$$\frac{1}{n} \log \left(E \exp(n\theta \tilde{X}_{i,n}) \right) = \frac{1}{n} \log E \left[M_i \left(\theta \frac{n}{\lambda_{i,n}} \right)^{\lambda_{i,n}} \right],$$

where the expectation is with respect to the probability measure of the random variable $\lambda_{i,n}$. Assumptions A1 and A2, Theorem 1, and dominated convergence lead to

$$M_i \left(\theta \frac{n}{\lambda_{i,n}} \right)^{\lambda_{i,n}/n} \to M_i \left(\frac{\theta}{p_i^*} \right)^{p_i^*} \quad (8)$$

a.s. By Jensen's inequality

$$\left(E \left[M_i \left(\theta \frac{n}{\lambda_{i,n}} \right)^{\lambda_{i,n}/n} \right] \right)^n \leq E \left[M_i \left(\theta \frac{n}{\lambda_{i,n}} \right)^{\lambda_{i,n}} \right],$$

so that

$$p_i^* \log M_i \left(\frac{\theta}{p_i^*} \right) \leq \liminf_{n \to \infty} \frac{1}{n} \log \left(E \exp(n\theta \tilde{X}_{i,n}) \right). \quad (9)$$

By Eq. (8), given $\varepsilon > 0$,

$$M_i \left(\theta \frac{n}{\lambda_{i,n}} \right)^{\lambda_{i,n}/n} \leq M_i \left(\frac{\theta}{p_i^*} \right)^{p_i^* + \varepsilon}$$

a.s. for all n sufficiently large, which results in

$$\limsup_{n \to \infty} \frac{1}{n} \log \left(E \exp(n\theta \tilde{X}_{i,n}) \right) \leq (p_i^* + \varepsilon) \log M_i \left(\frac{\theta}{p_i^*} \right).$$

Sending $\varepsilon \to 0$, together with Eq. (9), yields

$$\frac{1}{n} \log \left(E \exp(n\theta \tilde{X}_{i,n}) \right) \to p_i^* \log M_i \left(\frac{\theta}{p_i^*} \right).$$

By the Gärtner-Ellis theorem, we conclude that $1/n \log P(\tilde{X}_{i,n} \notin \Gamma) \to -p_i^* I_i(\gamma)$ for $i = 1,\ldots,a$, and $1/n \log P(\tilde{X}_{i,n} \in \Gamma) \to -p_i^* I_i(\gamma)$ for $i = a+1,\ldots,r$, as $n \to \infty$. The rest of the proof is similar to the second part of the proof of Proposition 1, and is thus omitted. $\otimes$

4 NUMERICAL EXAMPLES

We now illustrate our stochastic approximation algorithm through the two examples introduced in Section 2, where the performance metric of interest is distributed according to Normal and Bernoulli distributions, respectively. In addition to providing us with an opportunity to examine the finite-time performance of our algorithm, these two examples also illustrate the impact of deviations from the assumption of normality, which is quite common in the R&S literature. All the results are the average of 100 independent mega-replications. The results, which reflect the allocation of the computing budget as a function of the number of iterations of the algorithm, lead to two key observations: the efficacy of sequential algorithms and the importance of adequate initialization of such algorithms. The efficacy of these algorithms is particularly valuable in critical settings, i.e., those settings where the system performance is close to the boundary of the feasible region.

Example 3 *Suppose we have five systems with* $X_i \sim N(\mu_i, \sigma_i^2)$, *with* $\mu_i = [9.51, 9.45, 9.40, 9.55, 9.60]$ *and* $\sigma_i^2 = 1$. *Let the boundary point* $\gamma = 9.50$. *As a starting point, we will use Eq. (4), and sample from each system in proportion to the square of the distance of its scaled performance measure to* γ. *In this case, the optimal budget allocation for the five systems is given by* p_i^* *= [0.9091, 0.0364, 0.0091, 0.0364, 0.0091]. Figure 1 shows that convergence is rather rapid, that is, within 100000 iterations, the optimal allocation of the computing budget is achieved along with the correct classification of all the systems. Note that the algorithm has been initialized with only 100 samples from each system.*

Given the fact that the rate functions are generally unknown and need to be estimated, we next deploy our stochastic approximation algorithm to classify five systems with $X_i \sim N(\mu_i, \sigma_i^2)$, *with* $\mu_i = [9.20, 8.50, 9.00, 9.80, 10.01]$ *and* $\sigma_i^2 = 1$. *Let the boundary point* $\gamma = 9.50$. *In this case, the optimal budget allocation for the five systems is given by* p_i^* *= [0.3577, 0.0322, 0.1288, 0.3577, 0.1238]. Figure 2 shows that convergence has slowed down due to the effort required in estimating the rate functions; that is, it now takes over 1000000 iterations for the stochastic approximation algorithm to achieve the optimal allocation of the computing budget along with the correct classification of all the systems. Note that the algorithm has been initialized with 50000 samples from each system.*

Example 4 *Assume* $X_i \sim Bernoulli(\mu_i)$, *with* $\mu_i = [0.92, 0.85, 0.90, 0.98, 0.88]$. *Let the boundary point* $\gamma =$ 0.95. *Based on Eq. (5), the optimal budget allocation for the five systems is given by* p_i^* *= [0.4492, 0.0618, 0.1878, 0.1927, 0.1084]. Figure 3 shows that, within 10000 iterations, our stochastic approximation algorithm achieves the optimal allocation of the computing budget along with the correct classification of all the systems. One must recall, however, that appropriate initialization of the algorithm is crucial. The above performance is obtained after collecting 50000 observations to initialize the algorithm.*

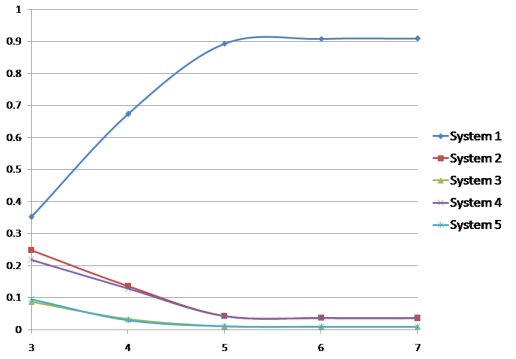

Figure 1: Convergence of the budget allocation as a function of log(iterations) for the Normal case with the exact rate function

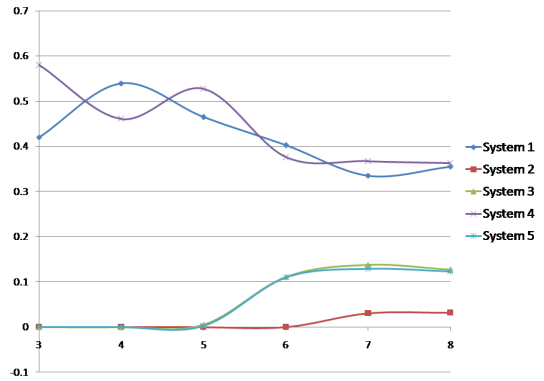

Figure 2: Convergence of the budget allocation as a function of log(iterations) for the Normal case with the estimated rate function

5 CONCLUDING REMARKS

We address the problem of feasibility determination in a stochastic setting. More specifically, we wish to determine whether a system belongs to a given set Γ based on a perfor-

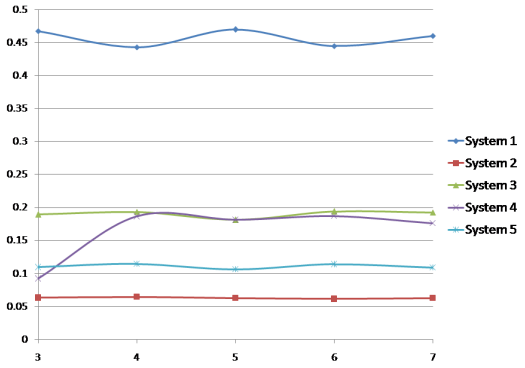

Figure 3: Convergence of the budget allocation as a function of log(iterations) for the Bernoulli case with the estimated rate function

mance measure estimated through Monte Carlo simulation. Our contribution is two-fold: (i) we characterize fractional allocations that are asymptotically optimal; and (ii) we provide an easily implementable algorithm, rooted in stochastic approximation theory, that results in sampling allocations that provably achieve in the limit the same performance as the optimal allocations. The finite-time behavior of the algorithm, illustrated on two small examples, is quite promising. Current research is aimed at generalizing the approach to settings where the feasible region is d-dimensional with $d > 1$.

ACKNOWLEDGMENTS

This work has been completed while the second author was visiting the Operations Research Department at the Naval Postgraduate School. He wishes to thank his co-author and the members of the Department for the stimulating research environment.

REFERENCES

Andijani, A. A. 1998. A multi-criterion approach for kanban allocations. *Omega* 26:483–493.

Andradottir, S., D. Goldsman, and S.-H. Kim. 2005. Finding the best in the presence of stochastic constraints. In *Proceedings of the 2005 Winter Simulation Conference*, ed. M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, 732–738. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Batur, D., and S.-H. Kim. 2005. Finding the set of feasible systems when the number of systems or constraints is large. In *Proceedings of the 2005 Winter Simulation Conference*, ed. M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, 692–698. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Berger, J. O. 1985. *Statistical decision theory and bayesian analysis*. 2nd ed. New York: Springer-Verlag.

Bernardo, J. M., and A. F. M. Smith. 1994. *Bayesian theory*. New York: Wiley Series in Probability and Mathematical Statistics.

Butler, J., D. J. Morrice, and P. W. Mullarkey. 2001. A multiple attribute utility theory approach to ranking and selection. *Management Science* 47:800–816.

Chen, H. C., C. H. Chen, and E. Yücesan. 2000. Computing efforts allocation for ordinal optimization and discrete event simulation. *IEEE Transactions on Automatic Control* 45:960–964.

Chick, S., and K. Inoue. 2001. New two-stage and sequential procedures for selecting the best simulated system. *Operations Research* 49:1609–1624.

Dembo, A., and O. Zeitouni. 1998. *Large deviations techniques and applications*. 2nd ed. New York: Springer-Verlag.

Dudewicz, E. J., and S. R. Dalal. 1975. Allocation of observations in ranking and selection with unequal variances. *Sankhya* B37:28–78.

Glynn, P., and S. Juneja. 2004. A large deviations perspective on ordinal optimization. In *Proceedings of the Winter Simulation Conference*, ed. R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 192–199. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Gupta, S. S., and K. J. Miescke. 1996. Bayesian look ahead one-stage sampling allocations for selection of the best population. *Journal of Statistical Planning and Inference* 54:229–244.

Gupta, S. S., K. Nagel, and S. Panchapakesan. 1973. On the order statistics from equally correlated normal random variables. *Biometrika* 60:403–413.

Gupta, S. S., and S. Panchapakesan. 1979. *Multiple decision procedures: Theory and methodology of selecting and ranking populations*. New York: Wiley.

Hsu, J. C. 1996. *Multiple comparisons: Theory and methods*. New York: Chapman & Hall.

Kim, K., and D. Lin. 2001. Simultaneous optimization of mechanical properties of steel by maximizing exponential desirability functions. *J. Roy. Statist. Soc. (Series C)* 49:1609–1624.

Kim, S.-H., and B. L. Nelson. 2006. Selecting the best: Theory and method. In *Handbooks in Operations Research and Management Science*, ed. S. G. Henderson and B. L. Nelson. Oxford: Elsevier.

Kushner, H. J., and G. G. Yin. 2003. *Stochastic approximation and recursive algorithms and applications*. 2nd ed. New York: Springer-Verlag.

Lee, L. H., E. P. Chew, and S. Tang. 2006. Integration of statistical selection with search mechanism for solving multi-objective simulation optimization problems. In *Proceedings of the 2006 Winter Simulation Conference*, ed. L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, 294–303. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Lee, L. H., E. P. Chew, and S. Tang. 2007. Finding the pareto set for multi-objective simulation models by minimization of expected opportunity cost. In *Proceedings of the 2007 Winter Simulation Conference*, ed. S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, 513–521. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Matejcik, F. J., and B. L. Nelson. 1995. Two-stage multiple comparisons with the best for computer simulation. *Operations Research* 43:633–640.

Nelson, B. L., and F. J. Matejcik. 1995. Using common random numbers for indifference-zone selection and multiple comparisons in simulation. *Management Science* 41:1935–1945.

Paulson, E. 1964. A sequential procedure for selecting the population with the largest mean from k normal populations. *Annals of Mathematical Statistics* 35:174–180.

Rinott, Y. 1978. On two-stage selection procedures and related probability-inequalities. *Communications in Statistics* A7:799–811.

AUTHOR BIOGRAPHIES

ROBERTO SZECHTMAN received his Ph.D. from Stanford University, after which he joined Oracle Corp. (2001-03). Currently, he is Assistant Professor in the Operations Research Department at the Naval Postgraduate School. His research interests include applied probability and military operations research. His email address is <rszechtm@nps.edu>.

ENVER YÜCESAN is a Professor of Operations Management at INSEAD in Fontainebleau, France. He holds a BS degree in Industrial Engineering from Purdue University and an MS and PhD in Operations Research from Cornell University. His research interests include ranking and selection, stochastic optimization, and supply chain management. He has also served as the Proceedings co-editor for WSC'02.