Proceedings of the 2007 Winter Simulation Conference
S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, eds.

NEW GREEDY MYOPIC AND EXISTING ASYMPTOTIC SEQUENTIAL SELECTION PROCEDURES: PRELIMINARY EMPIRICAL RESULTS

Stephen E. Chick

INSEAD
Technology and Operations Management Area
Boulevard de Constance
F-77305 Fontainebleau, FRANCE

Jürgen Branke, Christian Schmidt

Institute AIFB
University of Karlsruhe (TH)
D-76128 Karlsruhe, GERMANY

ABSTRACT

Statistical selection procedures can identify the best of a finite set of alternatives, where "best" is defined in terms of the unknown expected value of each alternative's simulation output. One effective Bayesian approach allocates samples sequentially to maximize an approximation to the expected value of information (EVI) from those samples. That existing approach uses both asymptotic and probabilistic approximations. This paper presents new EVI sampling allocations that avoid most of those approximations, but that entail sequential myopic sampling from a single alternative per stage of sampling. We compare the new and old approaches empirically. In some scenarios (a small, fixed total number of samples, few systems to be compared), the new greedy myopic procedures are better than the original asymptotic variants. In other scenarios (with adaptive stopping rules, medium or large number of systems, high required probability of correct selection), the original asymptotic allocations perform better.

1 OVERVIEW

Selection procedures are intended to select the best of a finite set of alternatives, where best is determined with respect to the largest mean, and the mean must be inferred via statistical sampling (Bechhofer et al. 1995, Kim and Nelson 2006). Selection procedures help to select the best of a finite set of alternative actions whose effects are evaluated with simulation. Selection procedures have attracted interest in combination with tools like evolutionary algorithms (Branke and Schmidt 2004, Schmidt et al. 2006), and discrete optimization via simulation (Boesel et al. 2003).

Since inference about the unknown mean performance of each system is estimated with stochastic simulation output, it is not possible to guarantee that the best alternative is selected with probability 1 in finite time. An ability to minimize the expected penalty for incorrect selections with a limited number of samples is desired. The expected penalty is typically measured by the probability of incorrect selection (PICS), or the expected opportunity cost (EOC) associated with potentially selecting a system that is not the best. Several frequentist and Bayesian formulations have been used to motivate and derive selection procedures.

Branke et al. (2007b) compare several of those formulations for fully sequential sampling procedures. A fully sequential procedure allocates one sample at a time to one of the simulated alternatives, runs a replication of the simulation for that alternative, updates the statistics for that system, and iterates until the selection procedure stops and identifies a system to select. They indicate that specific instances of two of those approaches, the expected value of information approach (VIP) (Chick and Inoue 2001) and the optimal computing budget allocation approach (OCBA) (Chen 1996, He et al. 2007) perform quite well, when used with particular adaptive stopping rules. The Online Companion to Branke et al. (2007b) also provides a theoretical explanation for why the VIP and OCBA approaches, when used to minimize the EOC objective function for sampling, perform similarly in numerical experiments.

The articles that derived the original VIP and OCBA procedures (Chick and Inoue 2001, Chen 1996, He et al. 2007) make use of asymptotic approximations (in the number of samples), approximations to the distribution of the difference of two variables with t distributions (e.g. Welch's approximation for the so-called Behrens-Fisher problem), and approximations to bound the probability of certain events (such as Bonferroni's or Slepian's inequality).

This paper presents new *small-sample* allocations with the VIP approach (derived in Branke, Chick, and Schmidt 2007a), and presents numerical comparisons of the new small-sample procedures with the original VIP approach. The derivation of the new allocations avoids an asymptotic approximation, the Behren's-Fisher problem (no Welch approximation) and the need to use Bonferroni's or Slepian's inequality. Because those approximations are averted, one might expect the new small-sample VIP procedures to perform better than their original large-sample counterparts.

We show empirically that this is the case only in certain situations. As expected, the new small-sample allocations work particularly well for a small total expected number of samples, and also for the case of a fixed total number of samples (as opposed to a flexible stopping rule) with few systems to be compared. The original allocations function better when used in combination with an adaptive stopping rule, and for high values of the probability of correct selection. This sheds some insight as to when, and how much, these various approximations degrade the performance of the selection procedures. The fact that one approach seems to do better than the other in specific types of situations leaves open the potential to combine the approaches to develop an even more effective procedure.

2 OLD AND NEW SELECTION PROCEDURES

We first describe the problem, assumptions and notation. Section 2.2 describes measures of the evidence of correct selection and stopping rules that influence the efficiency of the procedures. Section 2.3 recalls the original VIP procedures, in order to provide context for the new procedures in Section 2.4. The new procedures are derived elsewhere (Branke, Chick, and Schmidt 2007a).

2.1 Setup, Assumptions and Notation

The best of k simulated systems is to be identified, where 'best' means the largest output mean. Let X_{ij} be a random variable whose realization x_{ij} is the output of the jth simulation replication of system i, for $i = 1, \ldots, k$ and $j = 1, 2, \ldots$. Let w_i and σ_i^2 be the unknown mean and variance of simulated system i, and let $w_{[1]} \leq w_{[2]} \leq \ldots \leq w_{[k]}$ be the ordered means. In practice, the ordering $[\cdot]$ is unknown, and the best system, system $[k]$, is to be identified with simulation. The procedures considered below are derived from the assumption that simulation output is independent and normally distributed, *conditional* on w_i and σ_i^2,

$$\{X_{ij} : j = 1, 2, \ldots\} \stackrel{iid}{\sim} \texttt{Normal}\left(w_i, \sigma_i^2\right), \text{ for } i = 1, \ldots, k.$$

Although the normality assumption is not always valid, it is often possible to batch a number of outputs so that normality is approximately satisfied. Vectors are written as $\mathbf{w} = (w_1, \ldots, w_k)$ and $\boldsymbol{\sigma}^2 = (\sigma_1^2, \ldots, \sigma_k^2)$. A problem instance (*configuration*) is denoted by $\chi = (\mathbf{w}, \boldsymbol{\sigma}^2)$.

Let n_i be the number of replications for system i run so far. Let $\bar{x}_i = \sum_{j=1}^{n_i} x_{ij}/n_i$ be the sample mean and $\hat{\sigma}_i^2 = \sum_{j=1}^{n_i}(x_{ij} - \bar{x}_i)^2/(n_i - 1)$ be the sample variance. Let $\bar{x}_{(1)} \leq \bar{x}_{(2)} \leq \ldots \leq \bar{x}_{(k)}$ be the ordering of the sample means based on all replications seen so far. Equality occurs with probability 0 in contexts of interest here. The quantities n_i, $\bar{x}_i, \hat{\sigma}_i^2$ and (i) are updated as more replications are observed.

Each selection procedure generates estimates $\hat{w}_i$ of w_i, for $i = 1, \ldots, k$. For the procedures studied here, $\hat{w}_i = \bar{x}_i$, and a correct selection occurs when the selected system, system $\mathcal{D}$, is the best system, $[k]$. Here, system $\mathcal{D} = (k)$ is selected as best.

If T_ν is a random variable with standard Student t distribution with ν degrees of freedom, we denote the distribution of $\mu + \frac{1}{\sqrt{\kappa}} T_\nu$ by $\texttt{St}\,(\mu, \kappa, \nu)$ (Bernardo and Smith 1994). If $\nu > 2$ the variance is $\kappa^{-1}\nu/(\nu - 2)$. If $\kappa = \infty$ or 1/0, then $\texttt{St}\,(\mu, \kappa, \nu)$ is a point mass at μ. Denote the cumulative distribution function (cdf) of the standard t distribution $(\mu = 0, \kappa = 1)$ by $\Phi_\nu()$ and probability density function (pdf) by $\phi_\nu()$. The standardized EOC function is

$$\Psi_\nu[s] \stackrel{\text{def}}{=} \int_{u=s}^{\infty} (u - s)\phi_\nu(u)du = \frac{\nu + s^2}{\nu - 1}\phi_\nu(s) - s\Phi_\nu(-s).$$

2.2 Evidence for Correct Selection

The measures of effectiveness and some stopping rules for the procedures are defined in terms of loss functions. The zero-one loss function, $\mathcal{L}_{0-1}(\mathcal{D}, \mathbf{w}) = \mathbb{1}\left\{w_\mathcal{D} \neq w_{[k]}\right\}$, where the indicator function $\mathbb{1}\{\cdot\}$ equals 1 if its argument is true (here, the best system is not correctly selected), and is 0 otherwise. The opportunity cost $\mathcal{L}_{oc}(\mathcal{D}, \mathbf{w}) = w_{[k]} - w_\mathcal{D}$ is 0 if the best system is correctly selected, and is otherwise the difference between the best and selected system.

There are several frequentist measures that describe a procedure's ability to identify the actual best system. The frequentist probability of correct selection ($\mathrm{PCS}_{\mathrm{IZ}}$) is the probability that the mean of the system selected as best, system $\mathcal{D}$, equals the mean of the system with the highest mean, system $[k]$, conditional on the problem instance. (The subscript $_{\mathrm{IZ}}$ denotes indifference zone, a frequentist approach to selection procedures.) The probability is defined with respect to the simulation output X_{ij} generated by the procedure (the realizations x_{ij} determine $\mathcal{D}$),

$$\mathrm{PCS}_{\mathrm{IZ}}(\chi) \stackrel{\text{def}}{=} 1 - \mathrm{E}\left[\mathcal{L}_{0-1}(\mathcal{D}, \mathbf{w}) \mid \chi\right] = \Pr\left(w_\mathcal{D} = w_{[k]} \mid \chi\right).$$

We denote the corresponding probability of *incorrect* selection as $\mathrm{PICS}_{\mathrm{IZ}}(\chi) = 1 - \mathrm{PCS}_{\mathrm{IZ}}(\chi)$.

Another measure of selection quality is the frequentist opportunity cost of a potentially incorrect selection,

$$\mathrm{EOC}_{\mathrm{IZ}}(\chi) \stackrel{\text{def}}{=} \mathrm{E}\left[\mathcal{L}_{oc}(\mathcal{D}, \mathbf{w}) \mid \chi\right] = \mathrm{E}\left[w_{[k]} - w_\mathcal{D} \mid \chi\right].$$

At times, the configuration itself is sampled randomly prior to the application of a selection procedure. We then average over the sampling distribution of the configurations,

$$\mathrm{EOC_{IZ}} = \mathrm{E}[\mathrm{EOC_{IZ}}(\chi)] \text{ or } \mathrm{PCS_{IZ}} = \mathrm{E}[\mathrm{PCS_{IZ}}(\chi)].$$

When the mean performance of each system is unknown, a different set of measures is required to describe the evidence for correct selection. Bayesian procedures assume that parameters whose values are unknown are random variables (such as the unknown means $\mathbf{W}$), and use the posterior distributions of the unknown parameters to measure the quality of a selection. Given the data $\mathcal{E}$ seen so far, two measures of selection quality are

$$\begin{aligned}
\mathrm{PCS_{Bayes}} &\stackrel{\text{def}}{=} 1 - \mathrm{E}\left[\mathcal{L}_{0-1}(\mathcal{D}, \mathbf{W}) \mid \mathcal{E}\right] \\
&= \Pr\left(W_{\mathcal{D}} \geq \max_{i \neq \mathcal{D}} W_i \mid \mathcal{E}\right) \\
\mathrm{EOC_{Bayes}} &\stackrel{\text{def}}{=} \mathrm{E}\left[\mathcal{L}_{oc}(\mathcal{D}, \mathbf{W}) \mid \mathcal{E}\right] \\
&= \mathrm{E}\left[\max_{i=1,2,\ldots,k} W_i - W_{\mathcal{D}} \mid \mathcal{E}\right],
\end{aligned}$$

the expectation taken over $\mathcal{D}$ (a function of the random X_{ij}) and the posterior distribution of $\mathbf{W}$ given $\mathcal{E}$.

Approximations in the form of bounds on the above losses are useful to derive sampling allocations and to define stopping rules. *Slepian's inequality* and the *Bonferroni inequality* (e.g., Kim and Nelson 2006) imply that the posterior evidence that system (k) is best satisfies

$$\begin{aligned}
\mathrm{PCS_{Bayes}} &\geq \prod_{j:(j)\neq(k)} \Pr\left(W_{(k)} > W_{(j)} \mid \mathcal{E}\right) \\
&\approx \prod_{j:(j)\neq(k)} \Phi_{\nu_{(j)(k)}}(d^*_{jk}) \stackrel{\text{def}}{=} \mathrm{PCS_{Slep}} \\
\mathrm{PCS_{Bayes}} &\geq 1 - \sum_{j:j\neq(k)} \Pr\left(W_j > W_{(k)} \mid \mathcal{E}\right) \\
&\approx 1 - \sum_{j:j\neq(k)} \Phi_{\nu_{(j)(k)}}(-d^*_{jk}) \stackrel{\text{def}}{=} \mathrm{PCS_{Bonf}}
\end{aligned}$$

where d^*_{jk} is the normalized distance for systems (j) and (k), and $\nu_{(j)(k)}$ comes from *Welch's approximation* for the difference $W_{(k)} - W_{(j)}$ of two shifted and scaled t random variables (Law and Kelton 2000, p. 559),

$$\begin{aligned}
d^*_{jk} &= d_{(j)(k)} \lambda_{jk}^{1/2}, \text{ with} \\
d_{(j)(k)} &= \bar{x}_{(k)} - \bar{x}_{(j)}, \\
\lambda_{jk}^{-1} &= \frac{\hat{\sigma}^2_{(j)}}{n_{(j)}} + \frac{\hat{\sigma}^2_{(k)}}{n_{(k)}}, \text{ and} \\
\nu_{(j)(k)} &= \frac{[\hat{\sigma}^2_{(j)}/n_{(j)} + \hat{\sigma}^2_{(k)}/n_{(k)}]^2}{\frac{[\hat{\sigma}^2_{(j)}/n_{(j)}]^2}{(n_{(j)}-1)} + \frac{[\hat{\sigma}^2_{(k)}/n_{(k)}]^2}{(n_{(k)}-1)}}.
\end{aligned}$$

The Bayesian analogs to approximate the probability of correct selection and for the expected opportunity cost of a potentially incorrect selection are

$$\begin{aligned}
\mathrm{PCS_{Slep}} &= \prod_{j:(j)\neq(k)} \Phi_{\nu_{(j)(k)}}(\lambda_{jk}^{1/2}(d_{(j)(k)})) \\
\mathrm{EOC_{Bayes}} &\leq \sum_{j:(j)\neq(k)} \int_{w=0}^{\infty} w\, f_{(j)(k)}(w)\, dw \\
&\approx \sum_{j:(j)\neq(k)} \lambda_{jk}^{-1/2} \Psi_{\nu_{(j)(k)}}\left[d^*_{jk}\right] \stackrel{\text{def}}{=} \mathrm{EOC_{Bonf}},
\end{aligned}$$

where $f_{(j)(k)}(\cdot)$ is the posterior pdf for the difference $W_{(j)} - W_{(k)}$ given $\mathcal{E}$. That difference has approximately a $\mathtt{St}\left(-d_{(j)(k)}, \lambda_{jk}, \nu_{(j)(k)}\right)$ distribution.

The values of $\mathrm{EOC_{Bonf}}$ and $\mathrm{PCS_{Slep}}$ are used in the *stopping rules* of the sampling procedures below.

1. Sequential ($\mathcal{S}$): Repeat sampling while $\sum_{i=1}^{k} n_i < B$ for some specified total budget B.
2. Probability of correct selection ($\mathrm{PCS_{Slep}}$): Repeat while $\mathrm{PCS_{Slep}} < 1 - \alpha^*$ for a specified probability target $1 - \alpha^*$.
3. Expected opportunity cost ($\mathrm{EOC_{Bonf}}$): Repeat while $\mathrm{EOC_{Bonf}} > \beta^*$, for a specified EOC target β^*.

2.3 Value of Information Procedure (VIP)

VIPs allocate samples to each alternative in order to maximize the EVI of those samples. Some balance the cost of sampling with the EVI, and some maximize EVI subject to a sampling budget constraint (Chick and Inoue 2001). Procedures 0-1($\mathcal{S}$) and $\mathcal{LL}$($\mathcal{S}$) are sequential variations of those procedures that are designed to iteratively improve $\mathrm{PCS_{Bonf}}$ and $\mathrm{EOC_{Bonf}}$, respectively. Those procedures allocate τ replications per stage until a total of B replications are run. We recall those procedures in order to serve as a basis for comparison with the new procedures in Section 2.4.

Procedure 0-1.

1. Specify a first-stage sample size $n_0 > 2$, and a total number of samples $\tau > 0$ to allocate per subsequent stage. Specify stopping rule parameters.
2. Run independent replications $X_{i1}, \ldots, X_{in_0}$, and initialize the number of replications $n_i \leftarrow n_0$ run so far for each system, $i = 1, \ldots, k$.
3. Determine the sample statistics $\bar{x}_i$ and $\hat{\sigma}^2_i$, and the order statistics, so that $\bar{x}_{(1)} \leq \ldots \leq \bar{x}_{(k)}$.
4. WHILE stopping rule not satisfied DO another stage:
   (a) Initialize the set of systems considered for additional replications, $\mathcal{S} \leftarrow \{1, \ldots, k\}$.

(b) For each (i) in $\mathcal{S}\backslash\{(k)\}$: If $(k) \in \mathcal{S}$ then set $\lambda_{ik}^{-1} \leftarrow \hat{\sigma}_{(i)}^2/n_{(i)} + \hat{\sigma}_{(k)}^2/n_{(k)}$, and set $\nu_{(i)(k)}$ with Welch's approximation. If $(k) \notin \mathcal{S}$ then set $\lambda_{ik} \leftarrow n_{(i)}/\hat{\sigma}_{(i)}^2$ and $\nu_{(i)(k)} \leftarrow n_{(i)} - 1$.

(c) Tentatively allocate a total of τ replications to systems $(i) \in \mathcal{S}$ (set $\tau_{(j)} \leftarrow 0$ for $(j) \notin \mathcal{S}$):

$$\tau_{(i)} \leftarrow \frac{(\tau + \sum_{j\in\mathcal{S}} n_j)(\hat{\sigma}_{(i)}^2 \gamma_{(i)})^{\frac{1}{2}}}{\sum_{j\in\mathcal{S}} (\hat{\sigma}_j^2 \gamma_j)^{\frac{1}{2}}} - n_{(i)},$$

where

$$\gamma_{(i)} \leftarrow \begin{cases} \lambda_{ik} d_{ik}^* \phi_{\nu_{(i)(k)}}(d_{ik}^*) & \text{for } (i) \neq (k) \\ \sum_{(j)\in\mathcal{S}\backslash\{(k)\}} \gamma_{(j)} & \text{for } (i) = (k). \end{cases}$$

(d) If any $\tau_i < 0$ then fix the nonnegativity constraint violation: remove (i) from $\mathcal{S}$ for each (i) such that $\tau_{(i)} \leq 0$, and go to Step 4b. Otherwise, round the τ_i so that $\sum_{i=1}^k \tau_i = \tau$ and go to Step 4e.

(e) Run τ_i additional replications for system i, for $i = 1, \ldots, k$. Update sample statistics $n_i \leftarrow n_i + \tau_i$; $\bar{x}_i$; $\hat{\sigma}_i^2$, and the order statistics, so $\bar{x}_{(1)} \leq \ldots \leq \bar{x}_{(k)}$.

5. Select the system with the best estimated mean, $\mathcal{D} = (k)$.

The formulas in Steps 4b-4c are derived from optimality conditions to improve a Bonferroni-like bound on the EVI for asymptotically large τ (Chick and Inoue 2001). Depending on the stopping rule used, the resulting procedures are named 0-1($\mathcal{S}$), 0-1(PCS$_{\text{Slep}}$), 0-1(EOC$_{\text{Bonf}}$).

Procedure $\mathcal{LL}$ is a variant of 0-1 where sampling allocations seek to minimize EOC$_{\text{Bonf}}$.

Procedure $\mathcal{LL}$. Same as Procedure 0-1, except set $\gamma_{(i)}$ in Step 4c to

$$\gamma_{(i)} \leftarrow \begin{cases} \lambda_{ik}^{1/2} \frac{\nu_{(i)(k)} + (d_{ik}^*)^2}{\nu_{(i)(k)} - 1} \phi_{\nu_{(i)(k)}}(d_{ik}^*) & \text{for } (i) \neq (k) \\ \sum_{(j)\in\mathcal{S}\backslash\{(k)\}} \gamma_{(j)} & \text{for } (i) = (k). \end{cases}$$

2.4 New Small-Sample Procedures

Procedures 0-1 and $\mathcal{LL}$ allocate additional replications using an EVI approximation based on asymptotically *large* number of replications (τ) per stage. A performance improvement in the procedures *might* be obtained by better approximating EVI when there are a *small* number of replications per stage that are all run for one system.

The derivation of Procedure $\mathcal{LL}_1$, a small-sample procedure that is derived in (Branke, Chick, and Schmidt 2007a) and that is introduced below, avoids the asymptotic approximation in the derivation of $\mathcal{LL}$, as well as the Bonferroni and Welch approximations. The reasons it can do so are that:

- the predictive distribution of the sample mean to be computed for the single system which receives replications is known (and is a t distribution),
- only one system is sampled, therefore the Behrens-Fisher problem is avoided (no Welch approximation for the difference of t random variables),
- the system that would be selected as best after a single stage of simulation is either going to be the current best, or the current second best system (so only one comparison need be made, so the Bonferroni inequality is not needed),
- the EVI can be computed exactly, given the above assumptions, without resort to asymptotic approximations.

The procedure's name is distinguished from its large-sample counterpart by the subscript $_1$ (one system gets all replications in a given stage). In spite of the ability to compute the EVI precisely for a given single stage, the application of the resulting allocation sequentially leads to a greedy myopic policy. This new approach therefore may be suboptimal when applied sequentially.

The new procedures use the following variables, which are determined by the predictive distributions of the sample means of the different alternatives, given that additional samples will be, but have not yet been taken.

$$\begin{aligned} d_{\{jk\}}^* &= \lambda_{\{jk\}}^{1/2} d_{(j)(k)} \\ \lambda_{\{jk\}}^{-1} &= \left(\frac{\tau_{(k)} \hat{\sigma}_{(k)}^2}{n_{(k)}(n_{(k)} + \tau_{(k)})} + \frac{\tau_{(j)} \hat{\sigma}_{(j)}^2}{n_{(j)}(n_{(j)} + \tau_{(j)})} \right) \end{aligned} \tag{1}$$

Procedure $\mathcal{LL}_1$. Same as Procedure 0-1, except replace Steps 4a-4d by:

(a) For each $i \in \{1, 2, \ldots, k\}$, see if allocating to (i) is best:

   (i) Tentatively set $\tau_{(i)} \leftarrow \tau$ and $\tau_\ell \leftarrow 0$ for all $\ell \neq (i)$; set $\lambda_{\{jk\}}^{-1}$, $d_{\{jk\}}^*$ with Equation (1) for all j.

   (ii) Compute EVI$_{LL,(i)}$ as

$$\begin{cases} \lambda_{\{ik\}}^{-1/2} \Psi_{n_{(i)}-1}\left[d_{\{ik\}}^*\right] & \text{if } (i) \neq (k) \\ \lambda_{\{k-1,k\}}^{-1/2} \Psi_{n_{(k)}-1}\left[d_{\{k-1,k\}}^*\right] & \text{if } (i) = (k). \end{cases}$$

(b) Set $\tau_{(i)} \leftarrow \tau$ for the system that maximizes EVI$_{LL,(i)}$, and $\tau_\ell \leftarrow 0$ for the others.

The derivation of Procedure 0-1_1, another small-sample procedure that is derived in (Branke, Chick, and Schmidt 2007a), similarly avoids one of the two asymptotic approximations in the derivation of 0-1, as well as the Bonferroni and Welch approximations.

Procedure 0-1_1. Same as Procedure $\mathcal{LL}_1$, except the EVI is approximated with respect to the expected 0-1 loss,

$$\mathrm{EVI}_{0-1,(i)} = \begin{cases} \Phi_{n_{(i)}-1}(-d^*_{\{ik\}}) & \text{if } (i) \neq (k) \\ \Phi_{n_{(k)}-1}(-d^*_{\{k-1,k\}}) & \text{if } (i) = (k). \end{cases}$$

3 EMPIRICAL RESULTS

A detailed comparison of different selection procedures, including $\mathcal{KN}$++ (Goldsman et al. 2002), $\mathcal{OCBA}$ variations, and the original VIP procedures, can be found in Branke et al. (2007b). One main result of that paper was that $\mathcal{OCBA}_{\mathrm{LL}}$ and $\mathcal{LL}$ behave rather similar, and are among the best procedures over a large variety of different test cases.

In this follow-up study, we compare the small-sample and asymptotic VIP procedures. The test setup is similar to the one used in Branke et al. (2007b), which assessed different problem configurations, like monotone decreasing means, slippage configurations, and random problem instances. Many different parameter settings have been tested. We report here on only a typical subset of these runs due to space limitations, and focus on the main qualitative conclusions.

We report here on the efficiency of the procedures, which is defined in terms of the average number of samples needed to reach a given frequentist probability of correct selection (or expected opportunity cost, depending on the objective). All results reported are averages of 100,000 applications of the selection procedure. The initial number of evaluations for each system has been set to $n_0 = 6$, as this yielded good results in Branke et al. (2007b).

3.1 Monotone Decreasing Means

In a *monotone decreasing means (MDM)* configuration the means are equally spaced with distance δ. The independent outputs have a $\texttt{Normal}\left(w_i, \sigma^2\right)$ distribution,

$$X_{ij} \sim \texttt{Normal}\left(-(i-1)\delta, \sigma^2\right).$$

Figure 1 compares the asymptotic VIP procedures $\mathcal{LL}$ and 0-1 to the small-sample counterparts $\mathcal{LL}_1$ and 0-1_1, together with a very efficient adaptive stopping rule, $\mathrm{EOC}_{\mathrm{Bonf}}$. The small-sample procedures slightly outperform the original versions if only a few samples are taken (which corresponds to relatively high levels of empirical $\mathrm{PICS}_{\mathrm{IZ}}$). This is the scenario that the small-sample procedures have been designed for, and the removal of the Bonferroni and Welch approximations seems to actually pay off.

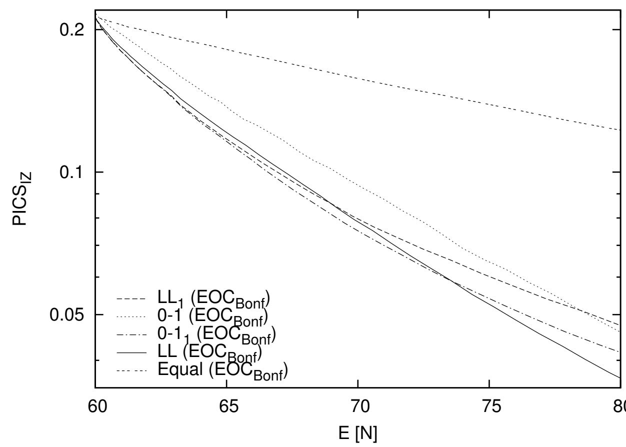

Figure 1: $\mathrm{PCS}_{\mathrm{IZ}}$ efficiency for the original and new allocations with the $\mathrm{EOC}_{\mathrm{Bonf}}$ stopping rule, and a small mean number of additional samples (MDM, $k = 10$, $\delta = 0.5$).

The situation changes completely with more demanding levels of $\mathrm{PICS}_{\mathrm{IZ}}$, as shown in Figure 2. In that case, the asymptotic procedures outperform the small-sample procedures by a wide margin, and the difference grows with increasing $\mathrm{E}[N]$. Still, note that the small-sample procedures are much better than the naïve Equal allocation.

Figure 3 shows the output for the same configuration, but with a fixed sampling budget ($\mathcal{S}$ stopping rule) as opposed to an adaptive stopping rule. All allocations become less efficient with the fixed sampling budget, as expected (compare Figures 2 and 3). The asymptotic allocations suffer more than the small-sample variants. In particular, 0-1($\mathcal{S}$) is the worst procedure for this experiment. The two small-sample procedures perform similar and somewhere between $\mathcal{LL}(\mathcal{S})$ and 0-1($\mathcal{S}$), even for large $\mathrm{E}[N]$ (which corresponds to low levels of $\mathrm{PICS}_{\mathrm{IZ}}$).

The influence of the stopping rule is even more apparent in Figure 4, which compares three different stopping rules for $\mathcal{LL}$ and $\mathcal{LL}_1$. For the $\mathcal{LL}$ stopping rule, $\mathrm{EOC}_{\mathrm{Bonf}}$ is clearly the most efficient, followed by $\mathrm{PCS}_{\mathrm{Slep}}$ and $\mathcal{S}$. The influence is rather large: For example, to reach a $\mathrm{PCS}_{\mathrm{IZ}}$ of 0.005, $\mathcal{LL}$ requires approximately 115, 125, or 172 samples, depending on stopping rule—the $\mathcal{S}$ stopping rule is much worse than the other two. For $\mathcal{LL}_1$, the ranking is the same for higher acceptable $\mathrm{PICS}_{\mathrm{IZ}}$, but the influence is smaller. For the above example of a $\mathrm{PCS}_{\mathrm{IZ}}$ of 0.005, the required numbers of samples are approximately 162, 190, and 195 for the three stopping rules. For very low values of $\mathrm{PICS}_{\mathrm{IZ}}$, the stopping rule with a fixed budget becomes even better

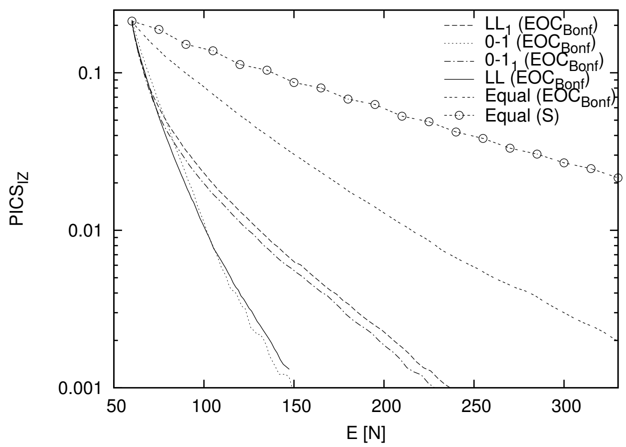

Figure 2: $\mathrm{PCS}_{\mathrm{IZ}}$ efficiency for the original and new allocations with the $\mathrm{EOC}_{\mathrm{Bonf}}$ stopping rule, and a large mean number of additional samples (MDM, $k = 10$, $\delta = 0.5$).

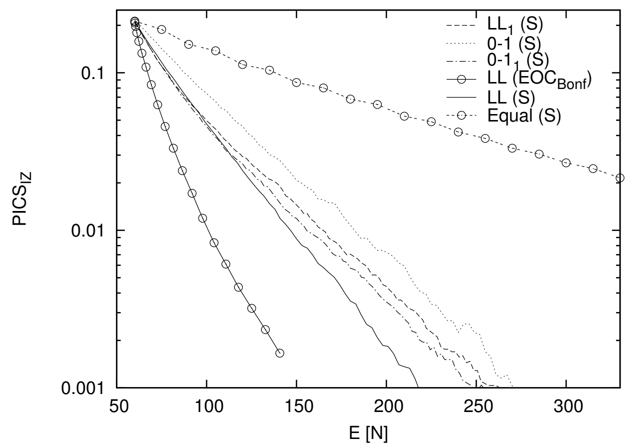

Figure 3: $\mathrm{PCS}_{\mathrm{IZ}}$ efficiency for the original and new small-sample VIPs with $\mathcal{S}$ stopping rule (MDM, $k = 10$, $\delta = 0.5$).

than the $\mathrm{PCS}_{\mathrm{Slep}}$ stopping rule (rightmost two lines); the corresponding efficiency line shows less curvature.

Figure 5 shows the influence of the number of systems, k, on the relative performance of $\mathcal{LL}(\mathrm{EOC}_{\mathrm{Bonf}})$ and $\mathcal{LL}_1(\mathrm{EOC}_{\mathrm{Bonf}})$. For larger k (e.g., $k = 20$ in the figure), the gap between $\mathcal{LL}(\mathrm{EOC}_{\mathrm{Bonf}})$ and $\mathcal{LL}_1(\mathrm{EOC}_{\mathrm{Bonf}})$ is larger for a given low desired $\mathrm{PICS}_{\mathrm{IZ}}$ (e.g., $\mathcal{LL}_1(\mathrm{EOC}_{\mathrm{Bonf}})$ requires approximately 60% more samples than $\mathcal{LL}(\mathrm{EOC}_{\mathrm{Bonf}})$ to reach a $\mathrm{PICS}_{\mathrm{IZ}}$ of 0.003). The relative performance of the small-sample procedure improves with decreasing k, until it is actually slightly better for $k = 2$ (in this case, to reach a $\mathrm{PICS}_{\mathrm{IZ}}$ of 0.003, it requires approximately 5% fewer samples on average than the asymptotic counterpart). This is achieved in combination with the $\mathrm{EOC}_{\mathrm{Bonf}}$ stopping rule, which is generally less favorable for the small-sample procedures as shown above.

3.2 Slippage Configuration

In a *slippage configuration (SC)* the means of all systems except the best are tied for second best. We use the parameter δ to describe the configurations of the independent outputs with $\texttt{Normal}\left(w_i, \sigma^2\right)$ distribution,

$$
\begin{aligned}
X_{1j} &\sim \texttt{Normal}\left(0, \sigma^2\right) \\
X_{ij} &\sim \texttt{Normal}\left(-\delta, \sigma^2\right) \text{ for } i = 2, \ldots, k
\end{aligned}
$$

The above MDM example with $k = 2$ is actually also a SC. More SC results with $k = 5$ are shown in Figure 6. Again, in combination with the $\mathrm{EOC}_{\mathrm{Bonf}}$ stopping rule, $\mathcal{LL}$ is more efficient than $\mathcal{LL}_1$, while in combination with the $\mathcal{S}$ stopping rule, the opposite is true.

3.3 Random Problem Instances

Random problem instances (RPI) are more realistic in the sense that problems faced in practice typically are not the SC or MDM configuration. The RPI experiment here samples configurations χ from normal-inverse gamma family. If $S \sim \texttt{InvGamma}\,(\alpha, \beta)$, then $\mathrm{E}[S] = \beta/(\alpha - 1)$ and $S^{-1} \sim \texttt{Gamma}\,(\alpha, \beta)$ with $\mathrm{E}[S^{-1}] = \alpha\beta^{-1}$ and $\mathrm{Var}[S^{-1}] = \alpha\beta^{-2}$. A random χ is generated by sampling the σ_i^2 independently, then sampling the W_i conditionally independent, given σ_i^2,

$$
\begin{aligned}
p(\sigma_i^2) &\sim \texttt{InvGamma}\,(\alpha, \beta) \\
p(W_i \,|\, \sigma_i^2) &\sim \texttt{Normal}\left(\mu_0, \sigma_i^2/\eta\right).
\end{aligned}
$$

Increasing η makes the means more similar. We set $\beta = \alpha - 1 > 0$ to standardize the mean of the variances to be 1. Increasing α reduces the variability in the variances. The noninformative prior distributions used for standard VIP and OCBA derivations correspond to $\eta \to 0$.

A typical result for the RPI problem instances is shown in Figure 7. It is consistent with the previous observations—the asymptotic variants are better in combination with the $\mathrm{EOC}_{\mathrm{Bonf}}$ stopping rule, while the small-sample procedures are at least competitive in combination with the $\mathcal{S}$ stopping rule or for a small number of additional samples.

4 DISCUSSION AND CONCLUSION

The choice of the selection procedure and its parameters can have a tremendous effect on the effort spent to select the best system, and the probabilistic evidence for making a correct selection.

The new small-sample VIP procedures avoid the Bonferroni and Welch approximations in their derivation, along

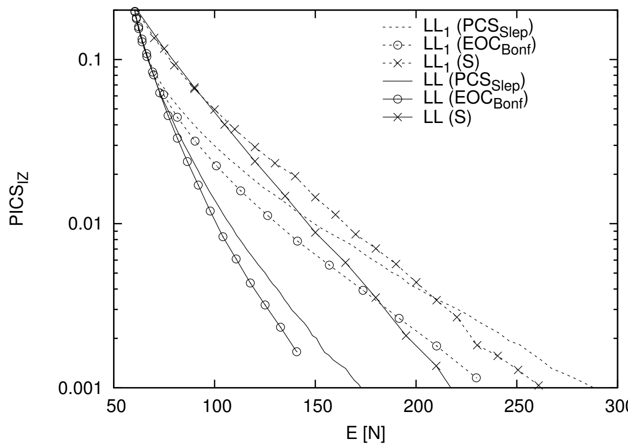

Figure 4: Comparison of different stopping rules for $\mathcal{LL}$ and $\mathcal{LL}_1$ (MDM, $k = 10$, $\delta = 0.5$).

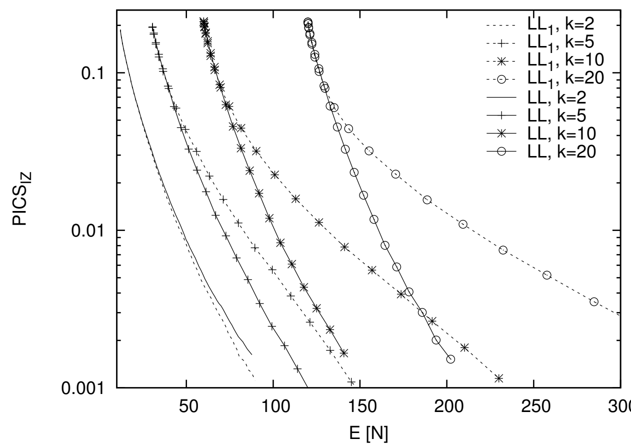

Figure 5: Influence of the number of systems, k, on $\mathcal{LL}(\mathrm{EOC}_{\mathrm{Bonf}})$ and $\mathcal{LL}_1(\mathrm{EOC}_{\mathrm{Bonf}})$ (MDM, $k = 10$, $\delta = 0.5$).

with an asymptotic approximation. That provides a potential net benefit relative to previous VIP derivations, which assumed a large-sample asymptotic and other approximations. The small-sample procedures suffer, however, from a myopic allocation that presumes that a selection will be made after one allocation. Repeatedly applying that allocation greedily, in order to obtain a sequential procedure, results in a sub-optimality. Naturally, this sub-optimality is more significant if the number of samples allocated before selection is large (i.e., if low numbers of $\mathrm{PICS}_{\mathrm{IZ}}$ are sought).

Overall, the empirical results that are presented above show that the small-sample procedures are particularly competitive if either the number of additional samples allocated is small, a fixed budget is used as stopping rule (as opposed to a flexible stopping rule such as $\mathrm{EOC}_{\mathrm{Bonf}}$), or the number of systems is small. The PCS-based procedures 0-1 and $0\text{-}1_1$ were almost always inferior the EOC-based allocations. Here, the small-sample variant $0\text{-}1_1$ yielded a greater improvement relative to 0-1, as compared with the relative difference in performance between $\mathcal{LL}$ and $\mathcal{LL}_1$.

One explanation for the still satisfactory performance of the asymptotic procedures may be that the approximations for the original VIP allocations are not significant, given that the allocations are rounded to integers in order to run an integer number of replications at each stage. The asymptotic allocations are designed to sample for long run behavior, and therefore perform well when more stringent levels of evidence are required.

Because there appear to be specific scenarios where the small-sample EVI procedures can be more effective than the original procedures, there may be a potential to adaptively combine the two allocation techniques in order to obtain an even more efficient selection procedure.

REFERENCES

Bechhofer, R. E., T. J. Santner, and D. M. Goldsman. 1995. *Design and Analysis for Statistical Selection, Screening, and Multiple Comparisons*. New York: John Wiley & Sons, Inc.

Bernardo, J. M. and A. F. M. Smith. 1994. *Bayesian Theory*. Chichester, UK: Wiley.

Boesel, J., B. L. Nelson, and S.-H. Kim. 2003. Using ranking and selection to ‘clean up’ after simulation optimization. *Operations Research 51*, 814–825.

Branke, J., S. Chick, and C. Schmidt. 2007a. Greedy versus asymptotic sequential selection procedures. Technical report, INSEAD, Technology and Operations Management Area working paper.

Branke, J., S. Chick, and C. Schmidt. 2007b. Selecting a selection procedure. *Management Science 53*, to appear.

Branke, J. and C. Schmidt. 2004. Sequential sampling in noisy environments. In X. Yao et al. (Ed.), *Parallel Problem Solving from Nature*, Volume 3242 of *LNCS*, pp. 202–211. Springer.

Chen, C.-H. 1996. A lower bound for the correct subset-selection probability and its application to discrete event simulations. *IEEE Transactions on Automatic Control 41*(8), 1227–1231.

Chick, S. E. and K. Inoue. 2001. New two-stage and sequential procedures for selecting the best simulated system. *Operations Research 49*(5), 732–743.

Goldsman, D., S.-H. Kim, W. S. Marshall, and B. L. Nelson. 2002. Ranking and selection for steady-state simulation: Procedures and perspectives. *INFORMS Journal on Computing 14*(1), 2–19.

He, D., S. E. Chick, and C.-H. Chen. 2007. The opportunity cost and OCBA selection procedures in ordinal opti-

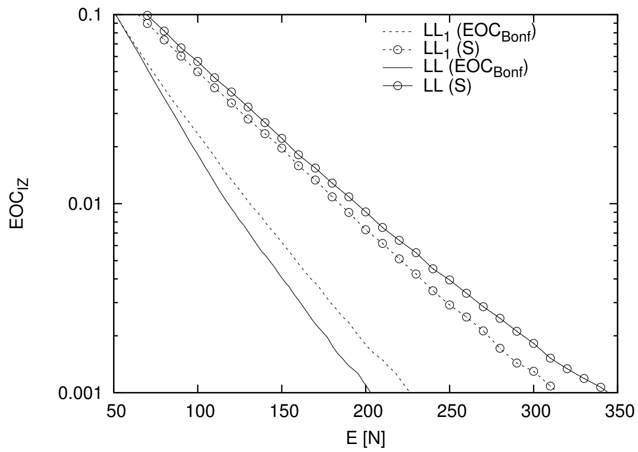

Figure 6: Influence of the stopping rule on $\mathcal{LL}(\text{EOC}_{\text{Bonf}})$ and $\mathcal{LL}_1(\text{EOC}_{\text{Bonf}})$ (SC, $k = 5$, $\delta = 0.5$).

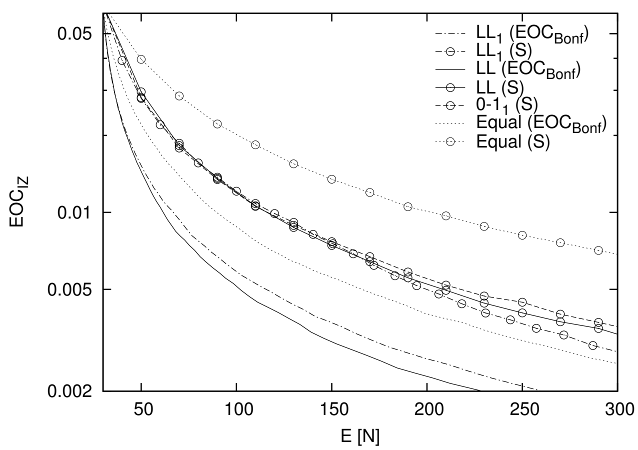

Figure 7: Efficiency of different allocation procedures for RPI (RPI, $k = 5$, $\eta = 2$, $b = 100$).

mization for a fixed number of alternative systems. *IEEE Trans. Systems, Machines, Cybernetics C*, to appear.

Kim, S.-H. and B. L. Nelson. 2006. Selecting the best system. In S. G. Henderson and B. L. Nelson (Eds.), *Handbook in Operations Research and Management Science: Simulation*. Elsevier.

Law, A. M. and W. D. Kelton. 2000. *Simulation Modeling & Analysis* (3rd ed.). New York: McGraw-Hill.

Schmidt, C., J. Branke, and S. Chick. 2006. Integrating techniques from statistical ranking into evolutionary algorithms. In *Applications of Evolutionary Computation*, Volume 3907 of *LNCS*, pp. 753–762. Springer.

AUTHOR BIOGRAPHIES

STEPHEN E. CHICK is a Professor of Technology and Operations Management at INSEAD. He has worked in the automotive and software sectors prior to joining academia, and now teaches operations with applications in manufacturing and services, particularly the health care sector. He enjoys Bayesian statistics, stochastic models, and simulation. His email address is <stephen.chick@insead.edu>, and his web page is <faculty.insead.edu/chick/>.

JÜRGEN BRANKE is a Research Associate of the Institute for Applied Computer Science and Formal Description Methods (AIFB) at the University of Karlsruhe, Germany. His main research area is nature-inspired optimization, with a special focus on optimization in the presence of uncertainties, including noisy or dynamically changing environments. Other research interests include multi-objective optimization, parallelization, and agent-based modelling. His email address is <jbr@aifb.uni-karlsruhe.de>, and his web page is <www.aifb.uni-karlsruhe.de/~jbr/>.

CHRISTIAN SCHMIDT is a scientist at Locom GmbH (Karlsruhe), a strategic logistics consulting company. His main duty is to develop and implement new optimization heuristics to solve complex, practical logistics problems. His main research interest is simulation-based optimization using nature-inspired heuristics. Last year, Dr. Schmidt finished his PhD on “Evolutionary Computation in Stochastic Environments” at the University of Karlsruhe, Germany. His email address is <christian.schmidt@locom.de>.