Proceedings of the 2013 Winter Simulation Conference
R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, eds.

A PROCEDURE TO SELECT THE BEST SUBSET AMONG SIMULATED SYSTEMS USING ECONOMIC OPPORTUNITY COST

Franco Chingcuanco

Dept. of Civil and Environmental Engineering
Engineering Systems Division
Massachusetts Institute of Technology
Cambridge, MA 02139, USA

Carolina Osorio

Dept. of Civil and Environmental Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139, USA

ABSTRACT

We consider subset selection problems in ranking and selection with tight computational budgets. We develop a new procedure that selects the best m out of k stochastic systems. Previous approaches have focused on individually separating out the top m from all the systems being considered. We reformulate the problem by casting all m-sized subsets of the k systems as the alternatives of the selection problem. This reformulation enables our derivation to follow along traditional ranking and selection frameworks. In particular, we extend the value of information procedure to subset selection. Furthermore, unlike previous subset selection efforts, we use an expected opportunity cost (EOC) loss function as evidence for correct selection. In minimizing the EOC, we consider both deriving an asymptotic allocation rule as well as approximately solving the underlying optimization problem. Experiments show the advantage of our approach for tests with small computational budgets.

1 INTRODUCTION

The best m of k simulated systems are to be identified, where best is defined as the smallest output mean. This subset selection problem falls under ranking and selection (R&S), which is a branch of statistics that uses statistical sampling to infer the best out of the k unknown competing systems. R&S procedures can be classified as either frequentist or Bayesian. Frequentist approaches, such as indifference zone (IZ) methods (Kim and Nelson 2006), determine the additional replications required to maximize the probability of correct selection based on statistically conservative assumptions (Chick and Inoue 2001). On the other hand, Bayesian formulations like the value of information procedure (VIP) (Chick and Inoue 2001) and the optimal computing budget allocation (OCBA) (Chen 1996) use the posterior distributions of the unknown means to provide evidence for correct selection. VIP allocates additional samples to maximize the expected value of information (EVI) from those samples, while OCBA approximates how additional replications change the parameter uncertainty associated with each system (Chick 2006).

The IZ, VIP and OCBA methods use loss functions to describe the loss associated with an incorrect selection. Note that loss functions are estimates of the penalty if a selection is wrong, whereas the EVI is a framework that measures the value additional observations bring in reducing the uncertainty associated with each system under consideration, before a decision is made. Two-stage and sequential R&S formulations exist. In the former, first-stage observations are used to inform the allocation of replications for the second-stage. The latter formulation is a natural extension of two-stage procedures where all observations up until stage l serve as the first-stage observations to inform stage $l+1$. By using information from all earlier stages, sequential approaches could lead to improved sampling efficiency (Chick and Inoue 2001). We focus on sequential formulations and refer to current- and next-stages to describe any consecutive stages l

and $l+1$ in our selection procedure. See Kim and Nelson (2006), Chick (2006) and Branke et al. (2007) for comprehensive reviews on these topics.

Most R&S procedures focus on identifying a single system with the best performance, even though one of the main papers R&S originates from is the seminal work of Gupta (1965) on subset selection (Kim and Nelson 2006). Gupta presents a single-stage procedure that seeks to identify a subset of random size that contains the best system. Sullivan and Wilson (1989) extend this work to a two-stage IZ procedure that selects a subset of at most size m. However, the primary motivation for both approaches is to screen inferior systems out, leaving the selected subset for further scrutiny (Chen et al. 2008). As Chen et al. (2008) note, the first attempt to identify the best m systems was not until Koenig and Law (1985). Recently, Chen et al. (2008) apply the latest developments in R&S to subset selection problems by extending OCBA to choosing the best m systems. The procedure, called OCBA-m, allocates a limited computational budget across the k systems to maximize the probability of correct selection (PCS) that the top-m systems are chosen. Zhang et al. (2012) present an improved version of this algorithm, called OCBA-m+, while LaPorte et al. (2012) extend OCBA for subset selection under very small computing budgets. Ryzhov and Powell (2009) have also recently developed a subset selection algorithm for online problems. In their procedure, entire subsets are treated as the individual alternatives of the selection problem, and a covariance structure is imposed on the priors that capture the correlations among overlapping alternatives. This design means that replications are allocated across entire subsets and not individual systems.

In this paper, we propose a new procedure to select the best subset of size m across k simulated systems. Our overall purpose is to integrate this subset selection procedure with a recently developed simulation-based optimization (SO) framework for urban transportation optimization by Osorio and Bierlaire (forthcoming). Since this SO framework operates under very tight computational budgets, our subset selection procedure focuses on problems with few systems and very small budgets. There are three notable distinctions between our approach and previous contributions. First, we formulate our problem as finding a subset of m elements that have the smallest collective performance out of all the possible subsets across the k simulated systems. Previous approaches, such as OCBA-m, attempt to identify and separate out the best m out of the k systems, whereas we consider entire subsets as our alternatives. Unlike Ryzhov and Powell (2009) however, we can allocate our replications across individual systems and are not limited to allocating across entire subsets. Second, rather than maximizing the PCS as in OCBA-m, we use an opportunity cost loss function and minimize the expected opportunity cost of incorrect selection. Third, we allocate additional samples using Bayesian decision theory to maximize the allocation's expected value of information, and hence we build on the work of Chick and Inoue (2001) by extending VIP to subset selection problems. As such, we call our procedure VIP-m.

2 NOTATIONS, ASSUMPTIONS & FORMULATION

2.1 Notation

The following notation in used in this paper:

- W_i: the unknown mean of system i
- σ_i^2: the unknown variance of system i
- λ_i: the unknown precision of system i ($\lambda_i = 1/\sigma_i^2$)
- x_{ij}: the jth independent and identically distributed observation or realization of performance measure for system i
- n_i: the number of simulation replications at the current stage for system i
- $\mathbf{x_i}$: the vector of simulation outputs for system i, $x_i = \{x_{ij} : j = 1,2\ldots n_i\}$
- $\bar{x}_i$: the sample mean of system i for samples observed up until the current stage, $\bar{x}_i = (1/n_i)\sum_{j=1}^{n_i} x_{ij}$

- $\hat{\sigma}_i^2$: the sample variance of system i, $\hat{\sigma}_i^2 = (1/(n_i-1))\sum_{j=1}^{n_i}(x_{ij}-\bar{x}_i)^2$
- r_i: the number of next-stage simulation replications to be allocated to system i
- Z_i: the predicted overall sample mean of system i prior to observing next-stage samples, $Z_i = (1/(n_i+r_i))\sum_{j=1}^{n_i+r_i} x_{ij}$
- $\tilde{W}_a$: the unknown mean of subset a, defined as the sum of the performance measures of all systems i in subset a, $\tilde{W}_a = \sum_{i\in a} W_i$
- $\tilde{\sigma}_a^2$: the variance of subset a, which is $\tilde{\sigma}_a^2 = \sum_{i\in a}\sigma_i^2$ for all independent systems in a
- J_{ab}: the difference in the unknown means of two subsets a and b, $J_{ab} = \tilde{W}_a - \tilde{W}_b$

2.2 Posterior Distribution of Systems and Subsets

Assume that the simulation output x_{ij} for system i $(i = 1,2\ldots k)$ and run j $(j = 1,\ldots n)$ is normally distributed with unknown mean W_i and unknown variance σ_i^2. Assume that the x_{ij}'s are jointly independent. We adopt a Bayesian framework and use information from an initial stage of sampling $(\mathbf{x_i})$ to infer the values of the unknown means W_i and precisions λ_i (the inverse of variances), treating them as random variables. Following Chick and Inoue (2001), we consider a noninformative prior distribution, which leads to a normal-gamma posterior for the joint distribution of W_i and λ_i after an initial n_i samples are taken (DeGroot 1970, Chapter 9.6). The marginal posterior distribution of the unknown mean is then a three-parameter Student-t distribution with $n_i - 1$ degrees of freedom

$$W_i|\mathbf{x}_i \sim St\left(\bar{x}_i, \frac{n_i}{\hat{\sigma}_i^2}, n_i - 1\right).$$

$W_i|\mathbf{x}_i$ gives the predicted performance of system i, but what we want is to identify the best subset of size m out of the k simulated systems. We define the performance of a subset as the collective performance of its elements (i.e., the sum of the elements in that subset). Note that out of the k simulated systems, there are $s = \binom{k}{m}$ possible subsets of size m, where $\binom{k}{m}$ is the binomial coefficient. For each subset a, let $\tilde{W}_a$ represent its unknown mean and $\tilde{\sigma}_a^2$ its unknown variance. As each system i is assumed to be normally distributed and independent of all of the other systems, then

$$\tilde{W}_a = \sum_{i\in a} W_i \quad \textit{and} \quad \tilde{\sigma}_a^2 = \sum_{i\in a}\sigma_i^2.$$

Since the posteriors of the individual W_i's are Student-t distributed, there is no known tractable expression for the distribution of their sum $\tilde{W}_a$ when the cardinality of a is greater than two and when the individual W_i's have arbitrary degrees of freedom. For our purposes, we instead use a normal density for the posteriors of the individual systems $W_i|\mathbf{x}_i \sim N\left(\bar{x}_i, \frac{\hat{\sigma}_i^2}{n_i}\right)$ by approximating $\sigma_i^2 \approx \hat{\sigma}_i^2$, similar to OCBA (Chen 1996). For independent systems, this allows the posterior density of the sum $\tilde{W}_a = \sum_{i\in a} W_i$ to be expressed as

$$\tilde{W}_a|\mathbf{x}_a \sim N\left(\sum_{i\in a}\bar{x}_i, \sum_{i\in a}\frac{\hat{\sigma}_i^2}{n_i}\right)$$

where $\mathbf{x}_a$ is the data observed for systems that belong in subset a.

2.3 Posterior Distribution Prior to Next-Stage Sampling

After an initial stage of sampling, the posterior distribution of the unknown means and variances of each subset is used as the prior distribution for the next stage. The goal is to determine the allocation rule

$r = (r_1, r_2, \ldots r_k)$ across the k systems that maximizes the probability that the best subset is selected after all next-stage replications have been run. Note that while we want to identify the best subset, we allocate the simulation runs across the individual systems and not across entire subsets.

Consider the posterior distribution of system i. Let $\bar{y}_i$ be the sample average of system i from only the next-stage observations (i.e., $\bar{y}_i$ only includes next-stage observations after an additional r_i's are allocated and observed and does not include samples from previous stages). After observing the next-stage samples, the overall sample average for i becomes (Chick et al. 2010)

$$z_i = E[W_i|\mathbf{x}'_i] = \frac{n_i\bar{x}_i + r_i\bar{y}_i}{n_i + r_i}$$

where $\mathbf{x}'_i$ is the overall data observed for system i, including next-stage samples. Prior to observing the next-stage samples, $\bar{Y}_i$ is a random variable. Following the work of Chick and Inoue (2001), it can be shown that the distribution of Z_i is then (DeGroot 1970, Chapter 11.9)

$$Z_i|\mathbf{X}'_i \sim N\left(\bar{x}_i, \frac{\hat{\sigma}_i^2}{n_i}\frac{r_i}{(n_i+r_i)}\right).$$

That is, the posterior distribution of system i after an initial round of sampling but prior to observing the next-stage samples is distributed around the current-stage sample mean $\bar{x}_i$, with a precision that is scaled by $(n_i + r_i)/r_i$. Following the approach above, the posterior distribution of subset a after r_i runs are allocated to system i ($i \in a$) is then

$$\tilde{Z}_a|\mathbf{X}'_a \sim N\left(\sum_{i\in a}\bar{x}_i, \sum_{i\in a}\frac{\hat{\sigma}_i^2}{n_i}\frac{r_i}{(n_i+r_i)}\right)$$

where $\mathbf{X}'_a$ is the overall data to be observed for systems that belong in subset a (including current- and next-stage samples). For an opportunity cost loss function (described below), we are interested in the posterior distribution of the random variable $J_{a_1a_2} = \tilde{Z}_{a_1} - \tilde{Z}_{a_2}$ for any two subsets a_1 and a_2. For independent systems, this posterior distribution is then

$$J_{a_1a_2}|\mathbf{X}'_a \sim N\left(\sum_{i\in a_1 i\notin a_2}\bar{x}_i - \sum_{j\notin a_1 j\in a_2}\bar{x}_j, \left[\sum_{i\in a_1 i\notin a_2}\frac{\hat{\sigma}_i^2}{n_i}\frac{r_i}{(n_i+r_i)} + \sum_{j\notin a_1 j\in a_2}\frac{\hat{\sigma}_j^2}{n_j}\frac{r_j}{(n_j+r_j)}\right]\right).$$

Note that the elements included in the posterior distribution of $\tilde{Z}_{a_1} - \tilde{Z}_{a_2}$ are only those that are unique across a_1 and a_2. This can be easily shown, and the result intuitively makes sense since when the difference $\tilde{Z}_{a_1} - \tilde{Z}_{a_2}$ is taken, the common elements within a_1 and a_2 cancel each other out.

3 EVIDENCE FOR CORRECT SELECTION OF THE BEST SUBSET

3.1 Loss Functions

At a given stage, we are faced with a total of s subsets, each with m elements. We want to allocate a limited computational budget across the k simulated systems to find the best subset with the smallest collective performance measure. Denote by b^* the subset with the true best performance, and denote the selected subset by b, which is the subset that has been observed to have the best performance at the current stage. In the R&S literature, the loss function $L(d, \omega)$ is used to provide evidence for correct selection. $L(d, \omega)$ describes the loss when a decision d is chosen given the current state ω. For the case at hand, a correct selection occurs when the selected subset b is the true best subset b^*. Branke et al. (2007) discuss two loss functions that are applied in both frequentist and Bayesian procedures:

1. the zero-one loss function $L_{0-1}(b, \omega) = I(\tilde{w}_b \neq \tilde{w}_{b^*})$ which equals 1 if the true best subset is not correctly selected and 0 otherwise;

2. and the opportunity cost loss function $L_{OC}(b,\omega) = \tilde{w}_b - \tilde{w}_{b^*}$ which is 0 if the true best is selected, and is otherwise the difference between the true best and the selected system.

For Bayesian R&S procedures, the probability of correct selection given the data seen so far ($\mathbf{x}$) is then (Branke et al. 2007)

$$PCS_{Bayes} = 1 - E[L_{0-1}(b,\mathbf{W})|\mathbf{x}] = P(\tilde{W}_b \neq \tilde{W}_{b^*}|\mathbf{x}).$$

An alternate measure of evidence for correct selection is the expected opportunity cost, which makes better sense when measuring output that have economic value (e.g., financial, time, man-hours, etc.)

$$EOC_{Bayes} = E[L_{OC}(b,\mathbf{W})|\mathbf{x}] = E[\tilde{W}_b - \tilde{W}_{b^*}|\mathbf{x}].$$

The original OCBA and OCBA-*m* procedures were formulated to maximize the PCS, while Chick and Inoue (2001) formulate the VIP by using both opportunity cost and zero-one loss functions. Since our intended application is for urban transportation problems where the performance measures can be naturally translated into economic metrics, we formulate our approach using an opportunity cost loss function. However, we do note that using L_{0-1} can be derived in an analogous manner.

3.2 Objective Function

Let $b(\mathbf{x}'_r)$ be the selected subset that has the best performance after all next-stage simulation outputs ($\mathbf{x}'_r$) have been observed. The subscript $\mathbf{r}$ is used as a reminder than the next-stage outputs are dependent on the simulation allocations $\mathbf{r}$ across the k systems. Given $\mathbf{x}'_r$ and a loss function L, the expected loss is then $E_{\mathbf{W}|\mathbf{x}'_r}[L(b(\mathbf{x}'_r),\mathbf{W})|\mathbf{x}'_r]$ (Chick and Inoue 2001). Since $\mathbf{r}$ is chosen before the next stage of sampling occurs, $\mathbf{X}'_r$ is random. Hence, nested expectations across $\mathbf{X}'_r$ are taken, (i.e., $E_{\mathbf{W}|\mathbf{x}'_r}[L(b(\mathbf{x}'_r),\mathbf{W})|\mathbf{x}'_r] = E_{\mathbf{X}'_r}[E_{\mathbf{W}|\mathbf{x}'_r}[L(b(\mathbf{x}'_r),\mathbf{W})|\mathbf{X}'_r]]$).

We use the objective function proposed in Chick and Inoue (2001). Assuming unit replication costs c for all systems, the objective is to minimize the expected loss when the subset with the best overall sample mean is selected subject to a replication budget constraint u

$$\min_r \rho(\mathbf{r}) = c\mathbf{r}^T + E_{\mathbf{X}'_r}[E_{\mathbf{W}|\mathbf{x}'_r}[L(b(\mathbf{x}'_r),\mathbf{W})|\mathbf{X}'_r]]) \tag{1}$$

$$s.t. \sum_{i=1}^{k} cr_i = u$$

$$r_i \in Z^+, i = 1,2\ldots k$$

The allocation rule that minimizes $\rho(\mathbf{r})$ is the vector $\mathbf{r}$ that minimizes a nested expectation, the inner expectation corresponding to the expected loss after $\mathbf{X}'_r = \mathbf{x}'_r$ is observed, the outer expectation averaging over $\mathbf{X}'_r$ (Chick 2006).

3.3 Approximate Objective Function

A general technique for determining an expression for $E_{\mathbf{X}'_r}[E_{\mathbf{W}|\mathbf{x}'_r}[L(b(\mathbf{x}'_r),\mathbf{W})|\mathbf{X}'_r]]$ is to consider an auxiliary loss function L^* that has the same optimal decision as the original L. Chick and Inoue (2001) derive a lower bound for $\rho(\mathbf{r})$ using this approach, which we adopt for our formulation. Consider the modified loss function

$$L^*_{OC}(b(\mathbf{X}'_r),\mathbf{w}) = L_{OC}(b(\mathbf{X}'_r),\mathbf{w}) - L_{OC}(b^*,\mathbf{w})$$

Adding $L_{OC}(b^*,\mathbf{w})$ to the original loss function does not change the optimal answer (DeGroot 1970, Chapter 8.4). An approximation to $E[L^*_{OC}(b(\mathbf{X}'_r),\mathbf{W})|\mathbf{X}'_r]$ is obtained using a Bonferroni-type bound (Chick and Inoue 2001, see Appendix for the full proof). The use of the auxiliary loss function and Bonferroni-type bound lead to the following approximate objective function

$$\min_{r} \rho^*(\mathbf{r}) = c\mathbf{r}^T + E[\min_{a} \tilde{w}_{b^*} - \tilde{w}_a] - \sum_{a \neq b} \lambda_{ba}^{-1/2} \Psi[\lambda_{ba}^{1/2} d_{ba}] \tag{2}$$

where

$$d_{ba} = \left(\sum_{j \in b, j \notin a} \bar{x}_j - \sum_{i \in a, j \notin b} \bar{x}_i \right)$$

$$\lambda_{ba}^{-1} = \left(\sum_{j \in b j \notin a} \frac{\hat{\sigma}_j^2}{n_j} \frac{r_j}{(n_j + r_j)} + \sum_{i \in a i \notin b} \frac{\hat{\sigma}_i^2}{n_j} \frac{r_i}{(n_j + r_j)} \right)$$

and $\Psi[s]$ is the standard normal loss function, defined as $\Psi[s] = \int_s^\infty (t-s)\phi(t)dt = \phi(s) - s(1-\Phi(s))$ (Chick 2006, Nahmias 2000, p. 262). Refer to Chingcuanco and Osorio (2013) for the full derivation.

4 SOLUTION TO THE OPTIMIZATION PROBLEM

4.1 Optimal Allocation Rule

One approach to solve the optimization problem posed above is to derive an optimal allocation rule that allocates the next-stage budget u to the different systems. Following Chick and Inoue (2001), we optimize $\sum_{a \neq b} \lambda_{ba}^{-1/2} \Psi[\lambda_{ba}^{1/2} d_{ba}]$ by relaxing the integrality constraints and taking partial derivatives with respect to the r_i's subject to a budget constraint $\sum_{i=1}^{k} r_i = u$. Since λ_{ba} is itself a function of the next-stage samples, a large-sample approximation is made so that $\lambda_{ba} \to \lambda_{ba*}$ where

$$\lambda_{ba*}^{-1} = \left(\sum_{j \in b j \notin a} \frac{\hat{\sigma}_j^2}{n_j} + \sum_{i \in a i \notin b} \frac{\hat{\sigma}_i^2}{n_j} \right). \tag{3}$$

and does not depend on next-stage samples, similar to Chick and Inoue (2001). The optimal allocation rule is then

$$r_i = \frac{\left[u + \sum_{j=1}^{k} n_j \right]}{\left[\sum_{j=1}^{k} \left(\frac{\hat{\sigma}_j^2 \eta_j}{\hat{\sigma}_i^2 \eta_i} \right)^{1/2} \right]} - n_i \tag{4}$$

where

$$\eta_i = \sum_{a'} \left(\frac{1}{2} \lambda_{ba*}^{1/2} \phi\left(\lambda_{ba*}^{1/2} d_{ba} \right) \right) \begin{array}{l} \quad a = 1, 2 \ldots s \\ a' : i \in a \quad if \quad i \notin b \\ a' : i \notin a \quad if \quad i \in b \end{array} \tag{5}$$

See Chingcuanco and Osorio (2013) for the details of the derivation. Note that η_i is a constant that does not depend on r_i. If i is in the best subset b, then it will be canceled out of all the subsets being compared against it. Hence, the summation for η_i is taken over all subsets a that do not contain i ($a' : i \notin a$). If i is not part of the best subset b, then simply sum over all subsets that contain i ($a' : i \in a$). It is worth mentioning that the allocation rule above is similar to the one obtained by Chick and Inoue (2001), with the only significant difference being how η_i is defined. Theirs is reproduced below for the case of normal posteriors to maintain consistency with our formulation

$$\eta_i = \left(\frac{1}{2} \lambda_{ik*}^{1/2} \phi\left(\lambda_{ik*}^{1/2} d_{ik} \right) \right) \quad [i] \neq [k]$$

$$\eta_i = \sum_{j=1}^{k-1} \eta_j \quad [i] = [k].$$

In our formulation, η_i is a sum of constants across all subsets that contain system i that capture how these subsets compare to the best subset b, while in the work of Chick and Inoue (2001) η_i pertains to how an individual system i compares against the current best. In addition, their η_i for the currently best system is a summation of k terms for the k systems, while we may have different number of terms across all η_i's depending on: 1) whether or not a system i belongs to the best subset b; and 2) the size of m. The reason for this difference is that in our formulation, systems common to any two subsets being compared cancel out each other.

4.2 VIP-m Analytical Allocation Procedure

We now describe our proposed subset selection procedure. Consider k systems in which the best m are to be selected. The description below assumes that best is defined as the smallest means. Set an initial sample size n_0, a next-stage budget u and a total budget t. Set $n_i = 0$ for all i.

1. Take n_0 replications for each of the k systems and set $n_i = n_i + n_0$.
2. Compute the current-stage sample means $\bar{x}_i = (1/n_i) \sum_{j=1}^{n_i} x_{ij}$ and sample variances $\hat{\sigma}_i^2 = \frac{1}{(n_i-1)} \sum_{j=1}^{n_i} (x_{ij} - \bar{x}_i)^2$ for the k systems.
3. Compute the current-stage order statistics for the k systems $\bar{x}_{[1]} \leq \bar{x}_{[2]} \leq \ldots \leq \bar{x}_{[m]} \leq \bar{x}_{[m+1]} \leq \ldots \bar{x}_{[k-1]} \leq \bar{x}_{[k]}$.
   (a) Take the smallest m systems and set this as the best subset b.
   (b) Out of the k systems, generate the remaining $s = \binom{k}{m}$ subsets for consideration.
4. While $t > 0$
   (a) Initialize the systems considered for additional replications $S \leftarrow 1, \ldots, k$.
   (b) Compute the next-stage replications for all systems using equations (3) to (5).
   (c) Remove any $r_i < 0$ from S and return to 4b.
   (d) Round the r_i's and take an additional r_i replications for each system i.
   (e) Set $n_i = n_i + r_i$.
   (f) Update sample and order statistics, as well as the best subset b and those subsets in s.
   (g) Reduce the total budget t by u.
5. Return the m elements with the smallest means.

Other stopping criteria besides $t > 0$ (e.g., meeting a certain threshold for the probability of correct selection or economic opportunity cost) could also be used (Branke et al. 2007).

4.3 VIP-m Numerical Allocation Procedure

Alternatively, it is also possible to minimize the objective function numerically which avoids the large-sample approximation made in (3). For the numerical allocation procedure, the posterior variances are approximated as $\frac{\hat{\sigma}_i^2}{n_i+r_i}$ similar to OCBA (Chen 1996). The full derivation of the procedure is available in Chingcuanco and Osorio (2013). When the replication costs are identical and when the integrality constraints are relaxed, the numerical optimization problem simplifies to

$$\min_r \sum_{a \neq b} \tilde{\lambda}_{ba}^{-1/2} \Psi[\tilde{\lambda}_{ba}^{1/2} d_{ba}] \tag{6}$$

$$s.t. \sum_{i=1}^{k} r_i = u$$

$$r_i \geq 0 \quad i = 1,2\ldots k$$

$$d_{ba} = \left(\sum_{j\in b, j\notin a} \bar{x}_j - \sum_{i\in a, j\notin b} \bar{x}_i \right)$$

$$\tilde{\lambda}_{ba*}^{-1} = \left(\sum_{j\in b j\notin a} \frac{\hat{\sigma}_j^2}{n_j + r_j} + \sum_{i\in a i\notin b} \frac{\hat{\sigma}_i^2}{n_i + r_i} \right).$$

The numerical allocation procedure is similar to the algorithm presented in Subsection 4.2, except that steps 4a to 4c are replaced with solving the optimization problem posed above.

5 NUMERICAL EXPERIMENTS

5.1 Description of Procedures Tested and Configurations

This section compares the performance of the proposed subset selection procedures against other subset selection methods under different testing configurations. The following *procedures* are evaluated:

1. The proposed analytical allocation rule described in Subsection 4.2.
2. The proposed numerical optimization procedure described in Subsection 4.3.
3. The OCBA-*m* (Chen et al. 2008) allocation procedure

$$\frac{n_1 + r_1}{\hat{\sigma}_1^2/\delta_1} = \frac{n_2 + r_2}{\hat{\sigma}_2^2/\delta_2} \ldots = \frac{n_k + r_k}{\hat{\sigma}_k^2/\delta_k}$$

   where $\delta_i = \bar{x}_i - c$ and c is a parameter derived from the order statistics. The code of the OCBA-*m* procedure was obtained from its authors for use in the numerical experiments.
4. The uniform allocation rule, i.e., $r_i = u/k$ for all i.
5. The proportional to variance allocation rule

$$\frac{n_1 + r_1}{\hat{\sigma}_1^2} = \frac{n_2 + r_2}{\hat{\sigma}_2^2} \ldots = \frac{n_k + r_k}{\hat{\sigma}_k^2}.$$

In all experiments, 6 systems are considered $(k = 6)$, where the best 3 elements $(m = 3)$ are to be identified. Three different testing *configurations* are used and described below:

1. Systems have equal variance and are distributed according to $N(i,6)$ $i = 1,2\ldots 6$.
2. Systems have variances that are increasing with their means, where systems are distributed as $N(i,i^2)$ $i = 1,2\ldots 6$.
3. Systems have variances that are decreasing with their means, where systems are distributed as $N(i,(k-i+1)^2)$ $i = 1,2\ldots 6$.

5.2 Results of Numerical Experiments

In comparing the procedures, two metrics are used and are calculated across 20,000 experiments $(N = 20,000)$. The estimated probability of correct selection (PCS) of a procedure is the fraction of time it correctly identifies the best *m*-systems across the *N* experiments. However, the PCS metric treats all incorrect selections equally and does not capture the magnitudes of the errors. Hence, economic opportunity cost (EOC) is also used to measure performance and is calculated as $EOC = \sum_{j=1}^{m} (\tilde{w}_b - \tilde{w}_{b^*})$. When the selected best is not the true best system $(b \neq b^*)$, the EOC is positive and it is zero otherwise. All tests were run with very small budgets: t was increased from 24 to 60 and the systems were initially sampled 3 times

($n_0 = 3$). The next-stage budget was set equal to the number of systems, i.e., $u = k = 6$. The results of the experiments are displayed in Figures 1 to 3 below.

Figure 1 shows that the analytical and numerical VIP-*m* procedures outperform the other three methods in the constant variance test. The relative improvements of the two VIP-*m* procedures improve as the budget is increased. As shown in the PCS and EOC graphs, OCBA-*m* has a higher PCS than uniform allocation but loses some of its advantage for higher budgets when the EOC is measured. When the systems have variances that increase with the mean (Figure 2), the two VIP-*m* procedures maintain their advantage over the other three methods but the improvement of OCBA-*m* over the uniform and proportional to variance allocations are diminished. Figure 3 illustrates this as well for the case where systems have decreasing variance. While the results are preliminary, they show the promise of our proposed subset selection procedure for problems with few systems and tight computational budgets.

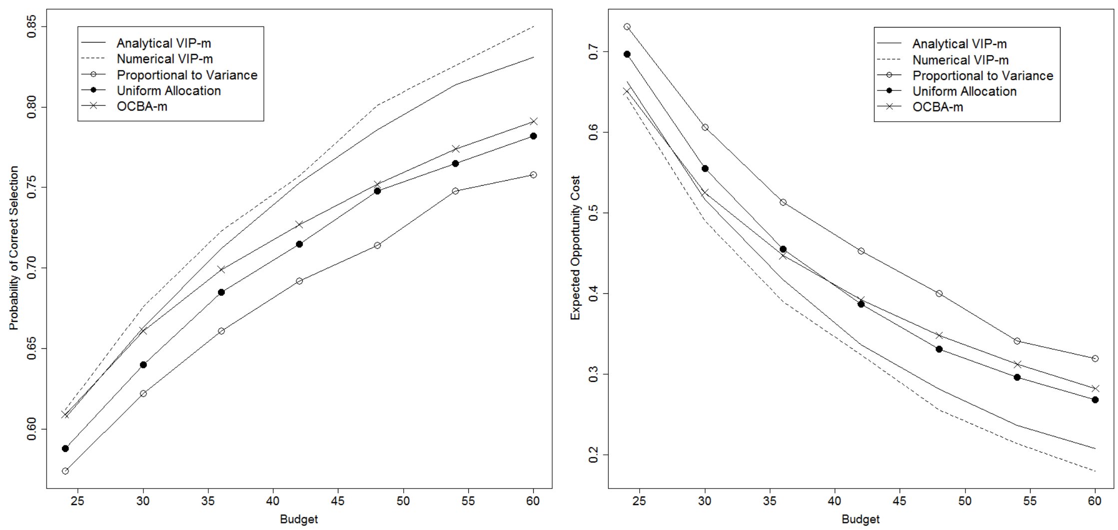

Figure 1: PCS and EOC for Test 1 (constant variance).

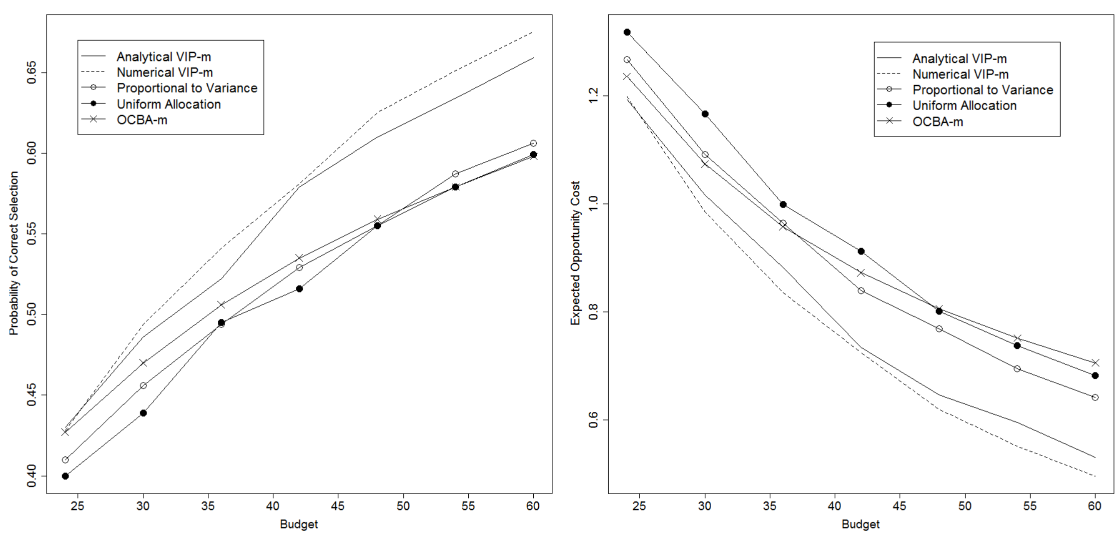

Figure 2: PCS and EOC for Test 2 (variance increasing with mean).

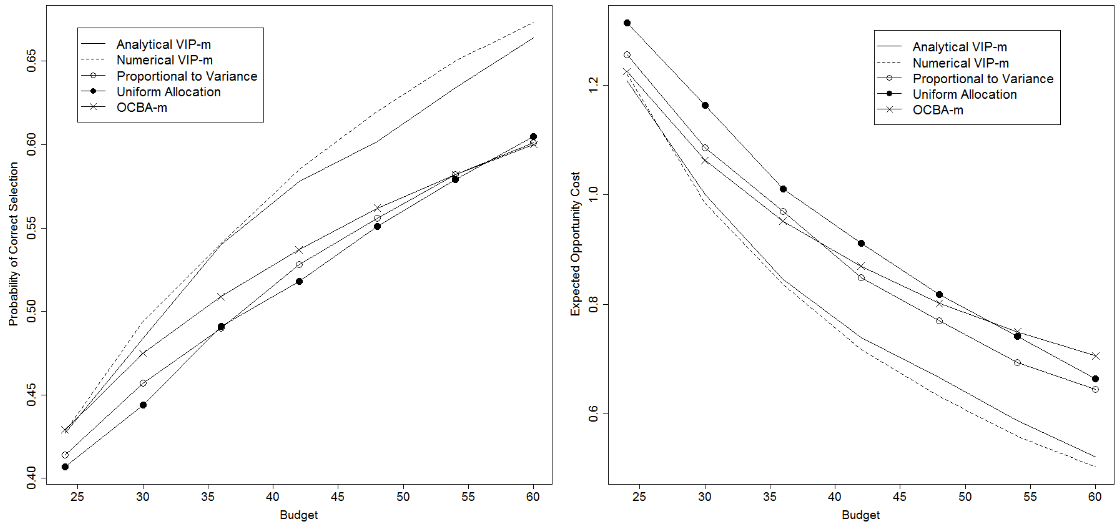

Figure 3: PCS and EOC for Test 3 (variance decreasing with mean).

6 CONCLUSION

We propose a new procedure, VIP-*m*, that selects the best subset of size *m* across *k* simulated systems. Building on the work of Chick and Inoue (2001), our approach adopts an EVI framework and uses an expected opportunity cost loss function. We focus on problems with few systems and very small budgets, and experiments show the advantage of our approach under these conditions. In the tests above, two consistent trends are observed. First, the VIP-*m* perform just as well or better compared to the other methods under all configurations. Note that all tests were limited to very small computational budget, which means that the advantage of the VIP-*m* may be diminished when the budgets and initial sample sizes are increased. Second, numerically solving the VIP-*m* (equation 6) yields only a very small margin over its analytic counterpart (equation 4). However, the numerical formulation takes on average about 5 times longer to solve. Hence, the marginal gain in performance may not outweigh the additional computational burden required when solving the problem numerically. Finally, at present, the VIP-*m* formulation is limited to a small number of systems because the number of subsets *s* of size *m* grows exponentially as the number of systems *k* increase. However, it is possible to circumvent this problem by employing smarter subset generation procedures (e.g., ignoring subsets that are clearly inferior).

There are ongoing experiments to further understand the performance of the VIP-*m* procedure when various parameters are changed, in particular: the initial sample sizes (n_0); the overall computational budget (t); the number of systems (k); and the next-stage allocation budget (u). Besides this, current efforts are devoted to applying the proposed technique to urban transportation problems. Specifically, we are designing how VIP-*m* could be integrated into the simulation-based optimization framework of Osorio and Bierlaire (forthcoming) as a step towards developing traffic-responsive control strategies.

REFERENCES

Branke, J., S. E. Chick, and C. Schmidt. 2007. "Selecting a Selection Procedure". *Management Science* 53:1916–1932.

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

Chen, C.-H., D. He, M. Fu, and L. H. Lee. 2008. "Efficient Simulation Budget Allocation for Selecting an Optimal Subset". *INFORMS Journal on Computing* 20:579–595.

Chick, S. E. 2006. “Subjective Probability and Bayesian Methodology”. In *Handbook in OR & MS, Vol. 13*. Elsevier.

Chick, S. E., J. Branke, and C. Schmidt. 2010. “Sequential Sampling to Myopically Maximize the Expected Value of Information”. *INFORMS Journal on Computing* 22:71–80.

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

Chingcuanco, F., and C. Osorio. 2013. “Development of a multi-instance metamodel-based simulation-optimization procedure for traffic responsive control”. Technical report, Massachusetts Institute of Technology.

DeGroot, M. H. 1970. *Optimal Statistical Decisions*. McGraw-Hill.

Gupta, S. S. 1965. “On some multiple decision (ranking and selection) rules”. *Technometrics* 7:225–245.

Kim, S.-H., and B. L. Nelson. 2006. “Selecting the Best System”. In *Handbook in OR & MS, Vol. 13*. Elsevier.

Koenig, L. W., and A. M. Law. 1985. “A procedure for selecting a subset of size m containing the l best of k independent normal populations”. *Communications in Statistics - Simulation and Computation* 14:719–734.

LaPorte, G. J., J. Branke, and C.-H. Chen. 2012. “Optimal Computing Budget Allocation for Small Computing Budgets”. In *Proceedings of the 2012 Winter Simulation Conference*.

Nahmias, S. 2000. *Production and Operations Analysis, 4th Edition*. McGraw-Hill.

Osorio, C., and M. Bierlaire. (forthcoming). “A simulation-based optimization framework for urban transportation problem”. *Operations Research*.

Ryzhov, I. O., and W. Powell. 2009. “The knowledge gradient algorithm for online subset selection”. In *IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning*.

Sullivan, D. W., and J. R. Wilson. 1989. “Restricted subset selection procedures for simulation”. *Operations Research* 37:52–71.

Zhang, S., L. H. Lee, E. P. Chew, C.-H. Chen, and H.-Y. Jen. 2012. “An Improved Simulation Budget Allocation Procedure to Efficiently Select the Optimal Subset of Many Alternatives”. In *2012 IEEE International Conference on Automation Science and Engineering (CASE)*.

ACKNOWLEDGMENTS

We would like to thank Chun-Hung Chen, Donghai He and Zhang Si for providing the OCBA-m code. We also appreciate Stephen Chick’s insights in the early stages of the work. Comments from the reviewers have also been helpful.

AUTHOR BIOGRAPHIES

FRANCO CHINGCUANCO is a candidate for dual Master’s degrees in Transportation and in Technology & Policy from the Massachusetts Institute of Technology (MIT). He is interested in developing models and algorithms to inform the planning and operations of urban transportation systems. His email address is fchingcu@mit.edu.

CAROLINA OSORIO is an Assistant Professor of Civil and Environmental Engineering (CEE) at the Massachusetts Institute of Technology (MIT). She received a Ph.D in Mathematics from from EPFL in 2010. She develops operations research techniques to inform the design and operations of transportation systems. She is interested in the techniques that combine ideas from the fields of probability theory, simulation, simulation-based optimization, derivative-free optimization, nonlinear optimization, statistics, traffic control and traffic flow theory. Her email address is osorioc@mit.edu.