COMPUTING BUDGET ALLOCATION FOR SIMULATION EXPERIMENTS WITH DIFFERENT SYSTEM STRUCTURES

Chun-Hung Chen
Yu Yuan

Hsiao-Chang Chen

Dept. of Systems Engineering
University of Pennsylvania
Philadelphia, PA 19104-6315, U.S.A.

Enver Yücesan

INSEAD
Technology Management Area
Fontainebleau, FRANCE

Liyi Dai

Dept. of Systems Science and
Mathematics
Washington University
St. Louis, MO 63130 U.S.A.

ABSTRACT

Simulation plays a vital role in analyzing discrete-event systems, particularly, in comparing alternative system designs with a view to optimize system performance. Using simulation to analyze complex systems, however, can be both prohibitively expensive and time consuming. We present effective algorithms to intelligently allocate computing budget for discrete-event simulation experiments with different system structures. These algorithms dynamically determine the best simulation lengths for all simulation experiments and thus significantly reduce the total computation cost for a desired confidence level. Numerical illustrations are included. We also compare our algorithms with our earlier approach in which different system structures are not considered. Numerical testing shows that we can further improve simulation efficiency.

1 INTRODUCTION

Simulation plays a central role in designing and efficiently managing large man-made systems such as communication networks, traffic systems, and manufacturing facilities since closed-form analytical solutions are scarce for such problems. Unfortunately, simulation can be both expensive and time consuming. Suppose we want to compare k different systems (competing designs or alternative operating policies). We conduct N simulation replications for each of the k designs. Therefore, we need kN simulation replications. Simulation results become more accurate as N increases. If the accuracy requirement is high (N is not small), and if the total number of designs in a decision problem is large (k is large), then kN can be very large, which may easily make total simulation cost prohibitively high and preclude the feasibility of simulation for ranking and selection problems.

The effective reduction of computation costs while obtaining a good decision is therefore crucial. Dudewicz and Dalal (1975) develop a two-stage procedure for selecting the best design or a design that is very close to the best system. In the first stage, all systems are simulated through n_0 replications. Based on the results of the first stage, the number of additional simulation replications for each design in the second stage is estimated in order to reach the desired confidence level. Rinott (1978) presents an alternative way to estimate the required number of simulation replications in the second stage. Many researchers have extended this idea to more general ranking and selection settings in conjunction with new developments. Chiu (1974), Gupta and Panchapakesan (1979), Matejcik and Nelson (1993, 1995), Bechhofer et al. (1995), and Hsu (1996) present methods based on the classical statistical model adopting a frequentist view. Berger (1980), Berger and Deely (1994), Bernardo and Smith (1994), Gupta and Berger (1988), and Chick (1997), on the other hand, use a Bayesian framework for constructing ranking and selection procedures.

Chen et al. (1996) present a new approach to improve simulation efficiency in ranking and selection. The underlying ideas are as follows. Intuitively, to ensure the correct selection of the best design, a larger portion of the computing budget should be allocated to those designs that are critical in the process of identifying the best design. In other words, a larger number of replications must be conducted with those critical designs in order to reduce estimator variance. On the other hand, limited computational effort should be expanded on non-critical designs that have little effect on identifying the best design even if they have large variances. In doing so, less computational effort is spent on simulating non-critical designs and more computational effort is spent on simulating critical designs; hence, the overall simulation efficiency is improved. Ideally, we want to optimally choose the number of simulation replications for all

designs to maximize simulation efficiency with a given computing budget. In fact, this question is equivalent to optimally decide which designs will receive additional computing budget for continuing the simulation further or to find an optimal way to reach an optimal design. Chen et al. (1997) present an approach incorporating a simple steepest-descent method (Luenberger 1984). Chen et al. (1998) compare the approach given in Chen et al. (1996) with Rinott's and Dudewicz's two-stage procedures and show that the new budget allocation is significantly faster than two-stage procedures.

In this paper, we will extend the results in Chen et al. (1998) to simulation experiments with different system structures. Namely, the simulation cost per replication varies from one system to another. This is true for many applications, in which some design alternatives are much more complex than others. As a result, the simulation costs for those complex systems are higher than others. We will present a budget allocation approach for the simulation experiments with different system structures. Numerical experiments show that the presented approach is more efficient than the one in Chen et al. (1997) which doesn't consider the impact of different system structures.

In the next section, we formulate the optimal computing budget allocation problem and discuss the major issues in solving this optimization problem. A Bayesian model for our budget allocation approaches is described in Section 3. In Section 4, we present an effective method of estimating sensitivity information and a sequential framework for dealing with budget allocation problems. Section 5 contains three numerical experiments. The results show that our approach obtains the desired confidence level while significantly reducing computation cost. Section 6 concludes this paper.

2 PROBLEM FORMULATION

Suppose that our goal is to select a design associated with the smallest mean performance measure among k alternative designs with unknown and possibly unequal variances. Without loss of generality, we define the "best" system as the one with the smallest mean performance measure in this paper. Further assume that the simulation output is independent from replication to replication. Note that we consider terminating (finite-horizon) simulations in this paper. Our approach is equally applicable to steady-state simulations where we need N independent samples rather than N independent simulation replications. In that case, the batch means method (Schmeiser 1982) can be applied to ensure the independence of the samples. The computing budget is limited. Denote by

- X_{ij} : the j -th *i.i.d.* sample of the performance measure from design i ,
- N_i : the number of simulation replications for design i ,

- c_i : the computation cost per replication for design i ,
- $\mathbf{X}_i$: the vector representing the simulation output for design i ; $\mathbf{X}_i = \{X_{ij} : j=1, 2, \dots, N_i\}$,
- $\bar{\mu}_i$: the sample average of the simulation output for design i ; $\bar{\mu}_i(N_i) = \frac{1}{N_i} \sum_{j=1}^{N_i} X_{ij}$,
- μ_i : the unknown mean performance measure; $\mu_i = E[X_{ij}]$,
- b : the design with the smallest sample mean performance; $b = \arg \min_i \{ \bar{\mu}_i \}$.

In ranking and selection problems, while the design with the smallest sample mean (design b) is usually picked, design b is not necessarily the one with the smallest unknown mean performance. *Correct selection* is therefore defined as the event that design b is actually the best design (i.e., with the smallest population performance). In the remainder of this paper, let "CS" denote *Correct Selection*. Thus,

$$P\{CS\} = P\{a = b\} \\ = P\{ \bar{\mu}_a(N_a) < \bar{\mu}_i(N_i), \text{ for all } i \neq a \}.$$

If the simulations are performed on a sequential computer, the computation cost is $c_1N_1 + c_2N_2 + \dots + c_kN_k$. To reduce the simulation efforts, we might want to minimize the total simulation cost for a desired $P\{CS\}$:

$$\min_{N_1, \dots, N_k} \{ c_1N_1 + c_2N_2 + \dots + c_kN_k \} \\ \text{s.t. } P\{CS\} \geq P^*. \quad (1.a)$$

where P^* is a user-defined confidence level requirement. Or, alternatively, we wish to maximize $P\{CS\}$ by utilizing a limited computing budget B :

$$\max_{N_1, \dots, N_k} P\{CS\} \\ \text{s.t. } c_1N_1 + c_2N_2 + \dots + c_kN_k = B. \quad (1.b)$$

Some difficulties in solving (1) include the following:

- There is no closed-form expression for the confidence level $P\{CS\}$.
- $P\{CS\}$ can not be computed before conducting the simulations if the mean and variance of each design are unknown.
- Most optimization techniques require sensitivity information.

Due to these problems, solving (1) can be very difficult, especially when k is not small. Since the purpose of solving (1) is to improve $P\{CS\}$ with a limited computing budget, we need a relatively fast and inexpensive way of solving (1) *during* the simulation experiment. Otherwise, the additional cost of solving (1)

will eliminate the benefits of the optimal computing budget allocation. These problems are addressed in Sections 3 and 4. We first introduce a Bayesian model to estimate $P\{CS\}$.

3 A BAYESIAN MODEL

To solve the budget allocation problem in (1), $P\{CS\}$ must be estimated. There exists a large literature on assessing $P\{CS\}$ based on classical statistical model. Goldsman and Nelson (1994) provide an excellent survey on available approaches to ranking, selection, and multiple comparisons (e.g., Goldsman, Nelson, and Schmeiser 1991, Gupta and Panchapakesan 1979). In addition, Bechhofer et al. (1995) provide a systematic and more detailed discussion on this issue.

In this paper, a Bayesian model (Bernardo and Smith, 1984) is used to develop an effective approach for solving the budget allocation problem in (1). Under a Bayesian model, we assume that the simulation output X_{ij} has a normal distribution with unknown mean μ_i . The unequal variance σ_i^2 can be known or unknown. After the simulation is performed, a posterior distribution of μ_i , $p(\mu_i | \mathbf{X}_i)$, can be constructed based on the prior knowledge on the system's performance and the simulation output. The probability of correctly selecting the best design can then be estimated by

$$P\{CS\} = p(\mu_b < \mu_i, i \neq b \mid \mathbf{X}_i, i=1,2,\dots,k).$$

First, consider the case where the variance σ_i^2 is known. Further assume that the unknown mean μ_i has the conjugate normal prior distribution $N(\eta_i, \nu^2)$. Under this assumption, the posterior distribution for any simulation output still belongs to the normal family. In particular, the posterior distribution of μ_i is

$$p(\mu_i | \mathbf{X}_i) \sim N\left(\frac{\sigma_i^2 \eta_i + N_i \nu^2 \bar{\mu}_i}{\sigma_i^2 + N_i \nu^2}, \frac{\sigma_i^2 \nu^2}{\sigma_i^2 + N_i \nu^2}\right).$$

Suppose that the performance of any design is unknown before conducting the simulation and that ν is an extremely large number. Then the posterior distribution of μ_i is given by

$$p(\mu_i | \mathbf{X}_i) \sim N\left(\bar{\mu}_i, \frac{\sigma_i^2}{N_i}\right).$$

If the variance σ_i^2 is unknown, a t -distribution model can be used (Inoue and Chick 1998, Chen et al. 1998). Then the posterior distribution of μ_i has a t distribution with

$$\begin{aligned} \text{mean} &= \bar{\mu}_i / N_i, \\ \text{precision} &= N_i / S_i^2, \\ \text{degrees of freedom} &= N_i - 1. \end{aligned}$$

After the simulation is performed, $\bar{\mu}_i$ and S_i^2 can be calculated, and $P\{CS\}$ can be estimated using a Monte Carlo simulation regardless of whether a normal distribution or a t distribution model is adopted. However, estimating $P\{CS\}$ via Monte Carlo simulation is time-consuming. Our budget allocation approach, as will be presented later, needs to estimate $P\{CS\}$ several times. Since the purpose of budget allocation is to improve simulation efficiency, we need a relatively fast and inexpensive way of estimating $P\{CS\}$ within the budget allocation procedure. Efficiency is more crucial than estimation accuracy in this setting. From Chen (1996) we have

$$P\{CS\} \geq \prod_{i=1, i \neq b}^k p(\mu_b < \mu_i \mid \mathbf{X}_b, \mathbf{X}_i) \equiv \text{APCS}. \quad (2)$$

We refer to the lower bound of the correct selection probability in (2) as the *Approximate Probability of Correct Selection (APCS)*. APCS can be computed very easily and quickly; no extra Monte Carlo simulation is needed. Numerical testing in Chen (1996) shows that APCS provides a reasonably good approximation to $P\{CS\}$. Furthermore, sensitivity information on the confidence level with respect to the number of simulation replications, N_i , which is central to the allocation of the computing budget, can be easily obtained. APCS is therefore used to approximate $P\{CS\}$ within the budget allocation procedure.

Note that APCS in (2) can be computed easily regardless of whether the posterior is normally or t distributed. For notational simplicity and the ease of discussion, we will only discuss the case of normal distribution in the rest of the paper, which is the case when the variances are known. If the variances are unknown, the discussions are still valid except that t distribution model is used.

4 SENSITIVITY ESTIMATION AND ASEQUENTIAL APPROACH

We now present a cost-effective method to estimate the sensitivity information and a sequential approach to resolve the difficulties in solving (1). Before conducting the simulation, neither APCS nor a good way to allocate the simulation budget is known. Therefore, all designs are initially simulated with n_0 replications, yielding the following posterior distribution for the performance of design i :

$$N\left(\frac{1}{n_0} \sum_{j=1}^{n_0} X_{ij}, \frac{\sigma_i^2}{n_0}\right).$$

We use this distribution to obtain the necessary sensitivity information. Let Δ_i be the additional computing budget allocated to design i ; more specifically, Δ_i is a non-negative integer denoting the number of additional simulation replications allocated to design i . We are

interested in assessing how the *APCS* would be affected if an additional simulation budget of Δ_i is allocated to design i . Note that this assessment must be done before actually conducting Δ_i simulation replications on design i . A preposterior distribution (Berger 1980 and DeGroot 1970) can provide a means for such a purpose. However, a Monte Carlo simulation is needed even to estimate *APCS*, which is too expensive in our OCBA framework. In this paper, we adopt a simple and fast approximation scheme. We know that if we conduct additional Δ_i replications on design i , the posterior distribution for design i is

$$N\left(\frac{1}{n_0 + \Delta_i} \sum_{j=1}^{n_0 + \Delta_i} X_{ij}, \frac{\sigma_i^2}{n_0 + \Delta_i}\right).$$

Assume that Δ_i is relatively small that $\frac{1}{n_0} \sum_{j=1}^{n_0} X_{ij}$ is a good approximation to $\frac{1}{n_0 + \Delta_i} \sum_{j=1}^{n_0 + \Delta_i} X_{ij}$. A good approximation to the preposterior distribution is

$$N\left(\frac{1}{n_0} \sum_{j=1}^{n_0} X_{ij}, \frac{\sigma_i^2}{n_0 + \Delta_i}\right), \text{ for design } i. \quad (3)$$

The estimated *APCS* can then be calculated by plugging (3) into the *APCS* formula in (2). We refer to this approximation as *EAPCS*. (*Estimated Approximate Probability of Correct Selection*). *EPACS* can be computed very easily.

Similarly, when the simulation is executed for $(N_1, N_2, \dots, N_{i-1}, N_i, N_{i+1}, \dots, N_k)$ replications, we can also use the available information to estimate how *APCS* will change if design i is given an additional budget Δ_i , i.e., *EAPCS*($N_1, N_2, \dots, N_{i-1}, N_i + \Delta_i, N_{i+1}, \dots, N_k$), by using an approximated posterior distribution $N\left(\frac{1}{N_i} \sum_{j=1}^{N_i} X_{ij}, \frac{\sigma_i^2}{N_i + \Delta_i}\right)$ for design i .

More specifically, when the simulation is executed for $(N_1, N_2, \dots, N_{i-1}, N_i, N_{i+1}, \dots, N_k)$ replications,

$$\begin{aligned} & \text{EAPCS}(N_1, N_2, \dots, N_{i-1}, N_i + \Delta_i, N_{i+1}, \dots, N_k) \\ &= P(\mu_b < \tilde{\mu}_i \mid \mathbf{X}_b, \mathbf{X}_i) \cdot \prod_{j=1, j \neq b, j \neq i}^k P(\mu_b < \mu_j \mid \mathbf{X}_b, \mathbf{X}_j), \\ & \text{where } \tilde{\mu}_i \sim N\left(\frac{1}{N_i} \sum_{j=1}^{N_i} X_{ij}, \frac{\sigma_i^2}{N_i + \Delta_i}\right), \end{aligned} \quad (4)$$

On the one hand, $\Delta_1, \Delta_2, \dots, \Delta_k$ should not be too large to obtain a good estimate of the *EAPCS*. On the other hand, a sufficient number of additional replications should be allocated to the competing designs to maximize $P\{\text{CS}\}$. To resolve this trade-off, we develop the following sequential algorithm for determining the number of simulation

replications. Instead of finding the best $N_1, N_2, \dots, N_k$ at the beginning of a simulation experiment, we sequentially add computing budget and determine the number of simulation replications such that the improvement of *EAPCS* is maximized. In doing so, $\Delta_1, \Delta_2, \dots, \Delta_k$ won't be too large.

Let $\Delta \equiv \sum_{i=1}^k c_i \Delta_i$. More specifically, Δ is a one-time

incremental computing budget in our sequential algorithm. We sequentially increase the computing budget by Δ until the total budget B is exhausted. At step l , $l = 1, 2, \dots$, we therefore solve

$$\begin{aligned} & \max_{\Delta_1, \dots, \Delta_k} \text{EAPCS}(N_1^l + \Delta_1^l, N_2^l + \Delta_2^l, \dots, N_k^l + \Delta_k^l), \\ & \text{s.t. } \sum_{i=1}^k c_i \Delta_i^l = \Delta \text{ and } \Delta_i^l \geq 0 \text{ for all } i. \end{aligned} \quad (5)$$

In summary, we have the following algorithm for the budget allocation problem in (1.a):

A Sequential Algorithm for Optimal Computing Budget Allocation (OCBA)

- Step 0.** Perform n_0 simulation replications for each of the k designs,
 $l \leftarrow 0$,
 $N_1^l = N_2^l = \dots = N_k^l = n_0$,
- Step 1.** If $\text{APCS}(N_1^l, N_2^l, \dots, N_k^l) \geq P^*$, stop, otherwise, go to Step 2.
- Step 2.** Solve (5).
- Step 3.** Perform additional Δ_i^l simulation replications for design i , $i = 1, \dots, k$.
 $N_i^{l+1} = N_i^l + \Delta_i^l$, for $i = 1, \dots, k$.
 $l \leftarrow l + 1$.
 update $\frac{1}{N_i^l} \sum_{j=1}^{N_i^l} X_{ij}$ (and S_i^2 if the variance is unknown), for $i = 1, \dots, k$.
 go to step 1.

At iteration l , our sequential algorithm determines the additional number of simulation replications, Δ_i^l , for each design. Since our goal is to achieve the highest $P\{\text{CS}\}$, we choose Δ_i^l by solving (5) where *EAPCS* is maximized. Once Δ_i^l is determined for each design, we execute Δ_i^l additional simulation replications on design i . This procedure is repeated until the desired confidence level P^{**} is attained or the total simulation budget is exhausted. In Section 6 we provide further guidelines for the selection of n_0 and Δ .

We present a simple greedy approach, where, m most promising designs are chosen and the computing budget is

distributed equally among them (each design has Δ/m). Since we want to maximize the *EAPCS*, promising designs can be defined as those designs that have the highest potential to increase the *APCS* if they are simulated further. For each design, we calculate the anticipated increase in the *APCS*, i.e., *EAPCS*, if a computing budget of Δ/m is allocated to it. Then a design is chosen if its *EAPCS* increase is among the top- m . Note that the simulation costs are different for different designs. The numbers of simulation replications in each iteration for different designs are different either.

Step 0. Choose appropriate m and Δ . Let $\tau_i = \frac{\Delta}{mc_i}$.

Step 1. For $i=1, \dots, n$, calculate

$$D_i \equiv EAPCS(N_1^l, N_2^l, \dots, N_{i-1}^l, N_i^l + \tau_i, N_{i+1}^l, \dots, N_n^l) - APCS(N_1^l, N_2^l, \dots, N_{i-1}^l, N_i^l, N_{i+1}^l, \dots, N_n^l).$$

Step 2. Find the set $S(m) \equiv \{i : D_i \text{ is among the highest } m\}$

Step 3. $\Delta_i^l = \tau_i$, for all $i \in S(m)$; otherwise, $\Delta_i^l = 0$.

5 NUMERICAL TESTING

This test case is a G/G/1 queue in which the objective is to select a design with minimum expected waiting time over a set of 10 competing designs ($k = 10$). All designs have the same interarrival time uniformly distributed over $[0.1, 1.9]$. Service time in design i is uniformly distributed over $[0.1, 1.3 + 0.05i]$, $i = 1, 2, \dots, 10$. We want to find a design with minimum average waiting time for customers served within the first 10 time units (terminating simulation). Since a higher service rate results in shorter waiting times, design 1 is the best design. In the numerical experiment, we compare the computation costs and the actual convergence probabilities $P\{CS\}$ for different approaches. Various parameter settings are used for each algorithm.

We set $\Delta = 12$ and $n_0 = 10$ in this example. To avoid spending too much time in solving (5), we allow only a small number of iterations when applying the steepest-descent method. 10,000 independent experiments are performed so that the average computation cost and $P\{CS\}$ can be estimated. For all compared algorithms in this paper, we estimate the $P\{CS\}$ by counting the number of times in which we successfully find the true best design (design 1 in this example) in those 10,000 independent experiments. $P\{CS\}$ is then obtained by dividing this number by 10,000, representing the correct selection frequency. Different confidence level requirements are also tested. Three examples for different simulation costs are tested.

Tables 1, 3, and 5 contain the test results using our OCBA algorithms presented in this paper for the three examples. On the other hand, tables 2, 4 and 6 include the test results using our earlier approach (Chen et al. 1997), which doesn't consider different simulation cost structure.

Example 1.

The simulation costs of all designs are

$$[c_1, c_2, c_3, \dots, c_{10}] = [4.0, 3.0, 2.0, 1.0, 1.0, 1.0, 1.0, 1.0, 1.0, 1.0].$$

Table 1: Consider different simulation structure. Average total number of simulation replications and $P\{CS\}$ by OCBA. We compare different numbers of promising designs we pick at each iteration for further simulation.

P^*	$m = 1$		$m = 2$	
	Total # of	$P\{CS\}$	Total # of	$P\{CS\}$
	Replication		Replication	
	s		s	
60%	453.9	72.1%	405.9	71.5%
80%	815.8	82.5%	818.6	80.8%
90%	1348.6	95.2%	1181.2	96.3%
95%	1550.6	96.7%	1435.8	98.7%

Table 2: Without considering different simulation structure. Average total number of simulation replications and $P\{CS\}$ by OCBA. We compare different numbers of promising designs we pick at each iteration for further simulation.

P^*	$m = 1$		$m = 2$	
	Total # of	$P\{CS\}$	Total # of	$P\{CS\}$
	Replication		Replication	
	s		s	
60%	473.2	71.2%	418.9	66.8%
80%	857.3	89.5%	853.6	88.0%
90%	1506.5	91.1%	1342.4	95.6%
95%	1952.6	96.0%	1667.9	97.4%

Example 2.

The simulation costs of all designs are

$$[c_1, c_2, c_3, \dots, c_{10}] = [3.0, 3.0, 2.0, 1.0, 3.0, 2.0, 1.0, 1.0, 1.0, 1.0].$$

Table 3: Consider different simulation structure.

P^*	$m = 1$		$m = 2$	
	Total # of	$P\{CS\}$	Total # of	$P\{CS\}$
	Replication		Replication	
	s		s	
60%	483.70	78.2%	468.93	69.8%
80%	838.02	88.6%	873.47	89.5%
90%	1378.64	90.9%	1367.45	96.2%
95%	1804.58	99.1%	1830.52	98.4%

Table 4: Without considering different simulation structure.

P^*	$m = 1$		$m = 2$	
	Total # of	$P\{CS\}$	Total # of	$P\{CS\}$
	Replication		Replication	
	s		s	
60%	499.70	66.3%	480.24	71.4%
80%	850.66	83.4%	1013.98	80.8%
90%	1445.60	94.5%	1472.47	91.1%
95%	1942.86	98.8%	1973.82	99.1%

Example 3.

The simulation costs of all designs are

$$[c_1, c_2, c_3, \dots, c_{10}] \\ = [1.0, 1.0, 2.0, 1.0, 1.0, 2.0, 1.0, 3.0, 2.0, 4.0].$$

Table 5: Consider different simulation structure.

P^*	$m = 1$		$m = 2$	
	Total # of	$P\{CS\}$	Total # of	$P\{CS\}$
	Replication		Replication	
	s		s	
60%	415.66	78.5%	392.25	71.5%
80%	563.79	85.4%	571.72	90.5%
90%	826.05	95.6%	768.53	97.0%
95%	964.23	98.2%	991.34	97.4%

Table 6: Without considering different simulation structure.

P^*	$m = 1$		$m = 2$	
	Total # of	$P\{CS\}$	Total # of	$P\{CS\}$
	Replication		Replication	
	s		s	
60%	426.44	72.5%	430.28	75.5%
80%	644.14	89.6%	614.22	91.3%
90%	909.17	96.3%	890.11	95.5%
95%	1142.74	96.8%	1130.06	96.5%

6 CONCLUSIONS

From the above numerical testing, we observe that under different simulation cost structure, our newly developed algorithm can effectively improve simulation efficiency. The efficiency improvement is particularly significant when the confidence requirement is high.

ACKNOWLEDGMENTS

The authors would like to thank Professors Stephen Chick and Barry Nelson for their helpful suggestions and valuable comments, which significantly improved this

work. Also, this work was supported in part by the National Science Foundation under grants DMI-9732173 and ECS-962479 and the University of Pennsylvania Research foundation

REFERENCES

- Bechofer R.E., T.J. Santner, and D.M. Goldsman. 1995. *Design and Analysis of Experiments for Statistical Selection, Screening, and Multiple Comparisons*. John Wiley & Sons, Inc.
- Berger, J. O. 1980. *Statistical Decision Theory, Foundations, Concepts, and Methods*. Springer Verlag.
- Berger, J. O., and J. Deely. 1994. "A Bayesian Approach to Ranking and Selection of Related Means with Alternative to AOV Methodology," *Journal of American Statistics Association*. 83, 364-373.
- Bernardo, J.M., and A.F.M. Smith. 1994. *Bayesian Theory*. Wiley.
- 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 8, 1227-1231.
- Chen, C. H. 1995. An Effective Approach to Smartly Allocate Computing Budget for Discrete Event Simulation. *Proceedings of the 34th IEEE Conference on Decision and Control*, 2598-2605.
- Chen, C. H., L. Dai, H. S. Chen, and E. Yücesan. 1998. "Efficient Computation of Optimal Budget Allocation for Discrete Event Simulation Experiment," submitted to *Operations Research*.
- Chen, H. C., C. H. Chen, L. Dai, and E. Yücesan. 1997. "New Development of Optimal Computing Budget Allocation For Discrete Event Simulation," *Proceedings of the 1997 Winter Simulation Conference*.
- Chen, C. H., H. C. Chen, and L. Dai. 1996. A Gradient Approach for Smartly Allocating Computing Budget for Discrete Event Simulation. *Proceedings of the 1996 Winter Simulation Conference*, 398-405.
- Chick, S. E. 1997. Selecting The Best System: a Decision-Theoretic Approach. *Proceedings of the 1997 Winter Simulation Conference*, 326-333.
- Chiu, W.K., 1974. The Ranking of Means of Normal Populations for a Generalized Selection Goal. *Biometrika*, 61, 579-584.
- DeGroot. 1970. *Optimal Statistical Decisions*. McGraw-Hill Book Co., New York.
- Dudewicz, E. J. and S.R. Dalal. 1975. Allocation of Observations in Ranking and Selection with Unequal Variances. *Sankhya*, B37:28-78.
- Goldsman, D., B. L. Nelson, and B. Schmeiser. 1991. Methods for Selecting the Best System. *Proceedings of the 1991 Winter Simulation Conference*, 177-186.

- Goldsman, D., and B. L. Nelson. 1994. Ranking, Selection, and Multiple Comparison in Computer Simulation. *Proceedings of the 1994 Winter Simulation Conference*, 192-199.
- Gupta, S. S. and J.O. Berger. 1988. *Statistical Decision Theory and Related Topics IV*. Springer Verlag.
- Gupta, S. S. and S. Panchapakesan. 1979. *Multiple Decision Procedures: Theory and Methodology of Selecting and Ranking Populations*. John Wiley.
- Hsu, J. C. 1996. *Multiple Comparisons Theory and Methods*. Chapman & Hall.
- Inoue, K., and S. E. Chick. 1998. "Comparison of Bayesian and Frequentist Assessments of Uncertainty for Selecting the Best System," to appear in the *Proceedings of the 1998 Winter Simulation Conference*.
- Luenberger, D. G., 1984. *Linear and Nonlinear Programming*. Addison-Wesley.
- Matejcek, F. J. and B.L. Nelson. 1993. Simultaneous Ranking, Selection and Multiple Comparisons for Simulation. *Proceeding of the 1993 Winter Simulation Conference*, 386-391.
- Matejcek, F. J. and B. L. Nelson. 1995. "Two-stage multiple comparisons with the best for computer simulation." *Operations Research*, 43, 633-640.
- Rinott, Y. 1978. On Two-stage Selection Procedures and Related Probability Inequalities. *Communications in Statistics A7*, 799-811.
- Schmeiser, B. 1982. Batch Size Effects in the Analysis of Simulation Output. *Operations Research*, 30:3, 556-567.

AUTHOR BIOGRAPHIES

CHUN-HUNG CHEN is an Assistant Professor of Systems Engineering at the University of Pennsylvania, Philadelphia, PA. He received his Ph.D. degree in Simulation and Decision from Harvard University in 1994. His research interests cover a wide range of areas in Monte Carlo simulation, optimal control, stochastic decision processes, ordinal optimization, perturbation analysis, and their applications to manufacturing systems. Dr. Chen won the 1994 Harvard University Eliahu I. Jury Award for the best thesis in the field of control. He is also one of the recipients of the 1992 MasPar Parallel Computer Challenge Award.

ENVER YÜCESAN is an associate Professor of Operations Research at INSEAD in Fontainebleau, FRANCE. He holds a BSIE degree from Purdue University, and an MS and a Ph.D. both in OR, from Cornell University. The work described in this paper has been initiated while he was visiting the Department of Systems Engineering at the University of Pennsylvania.

YU YUAN is a Ph.D. candidate at the Systems Engineering Department, University of Pennsylvania. He received a B.E. degree in Electrical Engineering and an M.S. degree in Control and Automation Engineering from Tsinghua University in 1994 and 1997, respectively. He is working on developing efficient approaches for discrete event simulation and optimization.

HSIAO-CHANG CHEN is a Ph.D. candidate at the Systems Engineering Department, University of Pennsylvania. He received a B.S. degree in Mathematics and Computer Science from the Eastern Michigan University in 1992, and he received an M.S. degree in Systems Science and Mathematics from Washington University, St. Louis in 1994. He is working on developing efficient approaches for discrete event simulation.

LIYI DAI is an assistant professor in the Department of Systems Science and Mathematics at Washington University, MO. He received the M.S. degree from the Institute of Systems Science, Academia Sinica, Beijing, China, in 1986, and the Ph.D. degree from Harvard University in 1993. His research interests include discrete event dynamic systems, simulation, stochastic optimization, communication systems, and singular systems. He has coauthored over 30 papers in various journals and is the author of *Singular Control Systems* (Berlin: Springer-Verlag, 1989). Dr. Dai is listed in *Who's Who* among Asian Americans and is a recipient of the NSF CAREER award.