Proceedings of the 1999 Winter Simulation Conference
P. A. Farrington, H. B. Nembhard, D. T. Sturrock, and G. W. Evans, eds.

AN ASYMPTOTIC ALLOCATION FOR SIMULTANEOUS SIMULATION EXPERIMENTS

Hsiao-Chang Chen
SynQuest Inc.
Clifton Park, NY 12065, U.S.A.

Chun-Hung Chen
Jianwu Lin
Dept. of Systems Engineering
University of Pennsylvania
Philadelphia, PA 19104, U.S.A.

Enver Yücesan
INSEAD
Technology Management Area
Fontainebleau, FRANCE

ABSTRACT

In this paper, we consider the allocation of a fixed total number of simulation replications among competing design alternatives in order to (i) identify the best simulated design, (ii) intelligently determine the best simulation run lengths for all simulation experiments, and (iii) significantly reduce the total computation cost. An asymptotically optimal allocation rule for maximizing a lower bound of the probability of correct selection is presented. Moreover, we illustrate the efficiency of our method with a series of generic numerical experiments. The simulation cost is significantly reduced with our sequential approach.

1 INTRODUCTION

In order to appropriately design large man-made systems such as communication networks, traffic systems, and automated manufacturing facilities, it is often necessary to apply extensive simulation to study their performance since no closed-form analytical solutions exist for these complex models. Unfortunately, using simulation to solve such problems can be both computationally expensive and time consuming due to their massive search space, evolution in time according to complex man-made rules, and the influence of random occurrences. Suppose we want to compare k different designs based on the data obtained when random samples are drawn from each of the k designs. If the performance of each system is completely unknown, we would conduct N simulation replications for each of the k designs. Therefore, we need a total of kN simulation replications. If the accuracy requirement is high (N is not small), and if the total number of designs in the selection problem is not small (k is large), the total simulation cost might become extremely high, therefore precluding the feasibility of simulation. The effective reduction of computation cost while obtaining a good decision is therefore a very important topic in simulation.

In general, there are two principal approaches in solving the ranking and selection problem:

(a) Specify the number of simulation replications first and hope that the resulting precision is satisfactory, or
(b) Specify the desired precision first and hope that the number of replications required to achieve this precision is not unacceptably large.

The well-known two-stage sampling procedure developed by Stein (1945) exemplifies the first approach. Bechhofer et al. (1954) prove that Stein's procedure satisfies a user-defined requirement for correct selection. Dunnett and Sobel (1954) develop an exact analysis for the probability of correct selection. Dudewicz and Dalal (1975) address the problem of selecting the normal population with the largest (mean) performance with unknown variances that are not necessarily common. They develop a two-stage procedure for selecting the best design, or a design that is very close to the best. At the first stage, all designs are simulated for n_0 replications. If the evidence is insufficient, based on the results obtained from the first stage, additional simulation replications are prescribed for each design in order to reach the desired confidence level. Rinott (1978) presents an alternative way to estimate the number of simulation replications required at the second stage.

As examples of the second approach, Tong and Wetzell (1984) and Futschik and Pflug (1997) develop adaptive procedures for sequentially allocating a fixed simulation budget. Tong and Wetzell use the results in Bechhofer (1969) to ensure that the best design obtains a desired proportion of the total number of replications, while developing a new rule to allocate the remaining replications to the other designs. Futschik and Pflug (1997), on the other hand, formulate a specific closed-form expression for an expected loss function. Therefore, an approximate solution for the objective function is given by applying nonlinear programming and a regularization technique.

These approaches have been extended to more general ranking and selection problems in conjunction with new developments. Chiu (1974), Gupta and Panchapakesan (1979), Matejcik and Nelson (1995), Bechhofer et al. (1995), and Hsu (1996) present methods based on the classical statistical model adopting a frequentist view. On the other hand, Berger and Deely (1994), Bernardo and Smith (1994), Chick (1997a, 1997b), and Chick and Inoue (1998) use a Bayesian framework for constructing ranking and selection procedures.

To further reduce the overall computation cost, Chen (1995) formulates the procedure for selecting the best design as another optimization problem. The idea in Chen's formulation is as follows: Intuitively, some inferior designs can be discarded during the early stages of simulation. As the simulation proceeds, on the basis of the estimated values of the parameters, a decision is made on which designs can be eliminated from further consideration as higher simulation accuracy for the remaining designs is obtained. This procedure is repeated until a desired assurance of correct selection is obtained. Little effort is therefore expanded on simulating inferior designs, reducing the overall simulation time. Ideally, we want to optimally allocate the number of simulation replications to each design in order to minimize the total simulation cost while obtaining the desired confidence level. In fact, this question is equivalent to optimally decide which designs will receive an additional computing budget for continuing the simulation or to find an optimal selection procedure to identify the best design.

Chen et al. (1996) provide an approach to solve such an optimization problem. They use Chernoff bounds to estimate the gradient information and then apply the modified steepest ascent method to solve this optimization problem. In a follow-up work, Chen et al. (1997) obtain the gradient information through finite differencing and then apply the steepest-ascent method to solve the budget allocation problem. Numerical results show that these approaches significantly reduce the overall simulation cost, while achieving the desired confidence level.

Instead of using the nonlinear programming methods described above, we opt for a decision-theoretic approach. Thus, we propose an asymptotically optimal allocation rule, Optimal Computing Budget Allocation (OCBA), for maximizing the lower bound of the probability of correct selection subject to a fixed total number of simulation replications.

The paper is organized as follows: In the next section, we formulate the optimal computing budget allocation (OCBA) problem and discuss the major issues in solving this optimization problem. Since our approach is based on the Bayesian framework, we also provide a brief discussion of that framework for completeness. Section 3 presents the techniques for solving OCBA. The performance of these techniques is illustrated with a series of generic numerical experiments in Section 4. Section 5 concludes the paper.

2 FORMAL STATEMENT OF THE PROBLEM

The principle goal is to select the best of k alternative system designs. Without loss of generality, we consider minimization problems in this paper; thus, the "best" design means the design with the smallest mean performance, μ_i. We assume that the competing designs have *known* variances that are not necessarily common. We further assume that the computing budget is limited and the number of competing designs is large. Denote by

- k : the total number of designs,
- X_{ij}: the j-th *i.i.d.* sample of the performance measure from design i,
- $\mathbf{X}_i$: the vector representing the simulation output for design i; $\mathbf{X}_i = \{X_{ij} : j=1,2,..,N_i \}$,
- N_i : the number of simulation replications for design i,
- $\bar{\mu}_i$: the sample mean performance measure for design i, i.e., $\bar{\mu}_i = \frac{1}{N_i}\sum_{j=1}^{N_i} X_{ij}$,
- μ_i: the mean performance measure; $\mu_i = E(X_{ij})$,
- σ_i^2 : the variance for design i,
- b: the design having the smallest sample mean performance measure, i.e., $\bar{\mu}_b \le \min_i \bar{\mu}_i$,
- s: the design having the second smallest sample mean performance measure i.e., $\bar{\mu}_b \le \bar{\mu}_s \le \min_{i \ne b} \bar{\mu}_i$,
- $\delta_{j,i} \equiv \bar{\mu}_j - \bar{\mu}_i$.

Note that when N_i's are large, $\bar{\mu}_i$ can be a good approximation for μ_i, since, according to the law of large numbers, $P\{\lim_{N_i \to \infty} \bar{\mu}_i = \mu_i\} = 1$. However, given the fact that we can conduct only a finite number of simulation replications, $\bar{\mu}_i$ is simply an approximation to μ_i. Since an approximation is used to select the best design, it is desirable to state the confidence in our decision, typically expressed as the probability of *correct selection*. In this paper, we define *correct selection* as the event that a design with the smallest sample performance measure (i.e., design b) *is* actually the best design. In the remainder of this paper, we let "CS" denote "correct selection."

There exists a large literature on selecting the best design based on the classical statistical model. Goldsman and Nelson (1994) provide an excellent survey of ranking, selection, and multiple comparison techniques for selecting the best system (e.g., Gupta and Panchapakesan 1979, Kleijnen 1987, Goldsman, Nelson, and Schmeiser 1991, and Law and Kelton 1991). In addition, Bechhofer, Santner, and Goldsman (1995) give a systematic and more detailed discussion on this issue. These approaches are mainly suitable for problems having a small number of competing designs. However, for large-scale industrial problems, the number of designs can rapidly grow extremely large.

We adopt the Bayesian framework for constructing an efficient approach to ranking and selection problems (Chen 1996, Chick and Inoue 1998). Under a Bayesian model, we assume that the simulation output, X_{ij}, has a normal distribution with mean μ_i and known variance σ_i^2. After the simulation is performed, a posterior distribution of μ_i, $P\{\ \mu_i \mid X_{ij}, j = 1, ..., N_i\ \}$, can be constructed based on two pieces of information: (i) the prior knowledge on the system's performance, and (ii) the simulation output. Thus, the probability of correctly selecting the best design can be defined by

$$\begin{aligned} P\{\text{CS}\} &= P\{\text{design } b \text{ is actually the best design}\} \\ &= P\{\mu_b < \mu_i,\ i \neq b \mid \mathbf{X}_i,\ i=1,2,..,k\ \}. \end{aligned} \tag{1}$$

To simplify the notation used, we rewrite Eq. (1) as $P\{\hat{\mu}_b < \hat{\mu}_i, i \neq b\}$, where $\hat{\mu}_i$ denotes the random variable whose probability distribution is the posterior distribution for design *i*.

As indicated earlier, we consider the case where the variance σ_i^2 is known. Under this assumption, the unknown mean μ_i has the conjugate normal prior distribution $N(\mu_0, v_0)$. Furthermore, the posterior distribution of any simulation output still follows the normal distribution. Thus, the posterior distribution of μ_i is

$$\hat{\mu}_i \sim N\left(\frac{\sigma_i^2\mu_0 + N_i v_0^2 \bar{\mu}_i}{\sigma_i^2 + N_i v_0^2}, \frac{\sigma_i^2 v_0^2}{\sigma_i^2 + N_i v_0^2}\right)$$
$$\text{for } i = 1,2,\ldots,k$$

Suppose that the performance of any design is completely unknown before conducting the simulation. In that case, De Groot (1970) suggests a procedure whereby a prior distribution is found by taking the parameter of the conjugate prior distribution to some limiting value. In particular, we can consider v_0 as an extremely large number. Chen et al. (1999a) then show that the posterior distribution of μ_i is given by

$$\hat{\mu}_i \sim N\left(\frac{1}{N_i}\sum_{j=1}^{N_i} X_{ij}, \frac{\sigma_i^2}{N_i}\right).$$

After the simulation is performed, σ_i^2 can be approximated by S_i^2, the sample variance. We further estimate $P\{\text{CS}\}$ using Monte Carlo simulation.

The goal is to maximize $P\{\text{CS}\}$, with a given number of simulation replications, *T*. Since the simulation budget is restricted, it is necessary to develop an allocation rule for N_i in such a way that the information about the best design is maximized. In other words, simulation replications should be allocated in a way that provides as much information as possible for the identification of the best design.

If simulations are performed on a sequential computer with the assumption that the simulation cost for each replication is roughly the same across different designs, the total computation cost can be approximated by $N_1 + N_2 + \cdots + N_k$. In the case where the simulation cost depends on the particular designs under consideration, our approach can be appropriately modified to handle such a case (Chen et al. 1998). Thus, under this environment, our problem can be formulated as

$$\begin{aligned} &\max_{N_1,\cdots,N_k} \ P\{\text{CS}\} \\ &\quad s.t.\ N_1 + N_2 + \cdots + N_k = T. \end{aligned} \tag{2}$$

The solution to (2) is complicated by the following restrictions:

1. $N_1, N_2, ..., N_k$ are integers and the number of combinations for $N_1, N_2, ..., N_k$ is large even for moderate *k*.
2. There is no closed-form expression for the confidence level, $P\{\text{CS}\}$.
3. Since the optimal allocation depends on uncertain parameters, $P\{\text{CS}\}$ can be computed *only after* exhausting the total simulation budget.

Due to these difficulties, obtaining an exact solution for (2) becomes virtually impossible, especially when *k* is extremely large. Since the purpose of solving (2) is simply to enhance $P\{\text{CS}\}$, we should not expand too much effort solving (2) during the overall ranking and selection process. Otherwise, the additional cost of solving (2) may overwhelm the benefits of optimal computing budget allocation. Hence, we need to find an inexpensive method that can solve (2) efficiently, even if this means obtaining a sub-optimal solution. In other words, efficiency is more crucial than optimality in this setting. In the next section, an asymptotic allocation rule with respect to the number of simulation replications, N_i, is presented.

3 OPTIMAL COMPUTING BUDGET ALLOCATION (OCBA)

The problem is the optimal allocation of the simulation replications to each competing design to maximize $P\{\text{CS}\}$. As discussed in the previous section, this allocation problem is difficult, since a change in any one of the N_i's always modulates the correlation structure of the underlying multivariate normal distribution, making it virtually impossible to compute $P\{\text{CS}\}$ exactly.

Let Y_i be a random variable (i = 1,2,…k). According to the Bonferroni inequality, $P\{\bigcap_{i=1}^{k}(Y_i < 0)\} \geq 1 - \sum_{i=1}^{k}[1 - P\{(Y_i < 0)\}]$. We replace Y_i by the random variable $(\hat{\mu}_b - \hat{\mu}_i)$ to provide a lower bound for the probability of correct selection. That is,

$$\begin{aligned} & P\{\text{CS}\} \\ &= P\{\bigcap_{\substack{i=1 \\ i\neq b}}^{k}(\hat{\mu}_b - \hat{\mu}_i < 0)\} \\ &\geq 1 - \sum_{i=1, i\neq b}^{k}[1 - P\{\hat{\mu}_b - \hat{\mu}_i < 0\}] \\ &= 1 - \sum_{i=1, i\neq b}^{k} P\{\hat{\mu}_b > \hat{\mu}_i\} \\ &\equiv APCS. \end{aligned}$$

We refer to this lower bound on the correct selection probability as the *Approximate Probability of Correct Selection (APCS). APCS* can be computed very easily and quickly; no extra Monte Carlo simulation is needed. We therefore use *APCS* to approximate $P\{\text{CS}\}$ as the simulation experiment proceeds. Chen et al. (1999b) approximate $P\{\text{CS}\}$ using the Chernoff bounds (Ross 1994) and offer an asymptotically solution, which is summarized in the following theorem.

Theorem 1. Given a total number of simulation replications T to be allocated to k competing designs whose performance is depicted by random variables with means $\mu_1, \mu_2, \ldots, \mu_k$, and finite variances $\sigma_1^2, \sigma_2^2, \ldots, \sigma_k^2$, respectively, the Approximate Probability of Correct Selection (*APCS*) can be asymptotically maximized when

$$\text{(1)} \quad \frac{N_b}{N_s} \rightarrow \frac{\sigma_b}{\sigma_s}\left[\sum_{\substack{i=1 \\ i\neq b}}^{k}\left(\frac{\delta_{b,s}^2}{\delta_{b,i}^2}\right)\right]^{1/2},$$

$$\text{(2)} \quad \frac{N_i}{N_s} \rightarrow \left(\frac{\sigma_i/\delta_{b,i}}{\sigma_s/\delta_{b,s}}\right)^2 \text{ for } i = 1, \ldots, k \text{ and } i \neq s \neq b,$$

where N_i is the number of replications allocated to design i, $\delta_{b,i} = \bar{\mu}_b - \bar{\mu}_i$ and $\bar{\mu}_b \leq \bar{\mu}_s \leq \min_{i\neq b\neq s} \bar{\mu}_i$. #

We now present a cost-effective sequential approach based on OCBA to select the best design from k alternatives within a given computing budget. Initially n_0 simulation replications for each of k designs are conducted to obtain some information about the performance of each design. As simulation proceeds, the means and variances of each design are estimated from the data already collected up to that stage. According to these estimated values, an incremental computing budget, Δ, is sequentially allocated. Ideally each new replication should bring us closer to the optimal solution. Namely, a decision is made on which designs should be eliminated from further consideration and how much computational effort should be invested for the next round of the simulation experiment. This procedure is continued until the total budget T is exhausted.

4 NUMERICAL TESTING

In this section, we will test our OCBA and compare it with several different allocation procedures by performing a series of numerical experiments.

4.1 Different Allocation Procedures

In addition to the OCBA we present in this paper, we test two more procedures and compare their performances. We briefly summarize the compared allocation procedures as follows.

4.1.1 Equal Allocation

This is the more straightforward way to conduct simulation experiments and has been well applied. The computing budget is equally allocated to all simulated designs. Namely, $N_i = T/k$ for each i. The performance of equal allocation will serve as a benchmark for comparison.

4.1.2 Two-Stage Rinott Procedure

The two-stage procedure of Rinott (1978) has been widely applied. Unlike the OCBA approach, the two-stage procedures are developed based on the classical statistical model. See Bechhofer et al. (1995) for a systematic discussion of the two-stage procedures. In the first stage, all designs are simulated for n_0 samples. Based on the sample variance estimate (S_i^2) obtained from the first stage, the number of additional simulation samples for each design in the second stage is determined:

$$N_i = \max(n_0, \lceil (S_i^2 h^2/d^2 \rceil), \text{ for } i = 1, 2, \ldots, k,$$

where $\lceil \bullet \rceil$ is the integer "round-up" function, d is the indifference zone, h is a constant which solves Rinott's integral (h can also be found from the tables in Wilcox 1984). In short, the computing budget is allocated proportional to the estimated sample variances. The major drawback is that only the information of variances is used when determining the simulation allocation, while our OCBA utilizes the information of both means and variances. As a result, the performance of Rinott procedure is not as good as our OCBA.

4.2 Numerical Experiments

4.2.1 Experiment 1. Normal Distribution

There are ten design alternatives. Suppose $X_{ij} \sim N(i, 6^2)$, i = 1, 2, .., 10. We want to find a design with the minimum mean. It is obvious that design 1 is the actual best design. In the numerical experiment, we compare the convergence of $P\{CS\}$ for different allocation procedures. Furthermore, 10,000 independent experiments are performed to estimate $P\{CS\}$. In all the numerical illustrations, we estimate $P\{CS\}$ by counting the number of times 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.

Figure 1 shows the test results using OCBA and other two different procedures given in section 4.1. All procedures obtain a higher $P\{CS\}$ as the available computing budget increases. However, OCBA achieves a higher $P\{CS\}$ than other procedures do with a same amount of computing budget. In particular, we indicate the computation costs in order to have $P\{CS\}$ = 99% for different procedures in Figure 1.

It is worth noting that Rinott procedure does not perform much better than the simple equal allocation. This is because Rinott Procedure determines the number of simulation samples for all designs using only the information of sample variances. On the hand, Rinott procedure is much slower than the our OCBA. This is because when determining budget allocation, OCBA exploits the information of both sample means and variances. The sample means can provide the valuable information of relative differences across the design space.

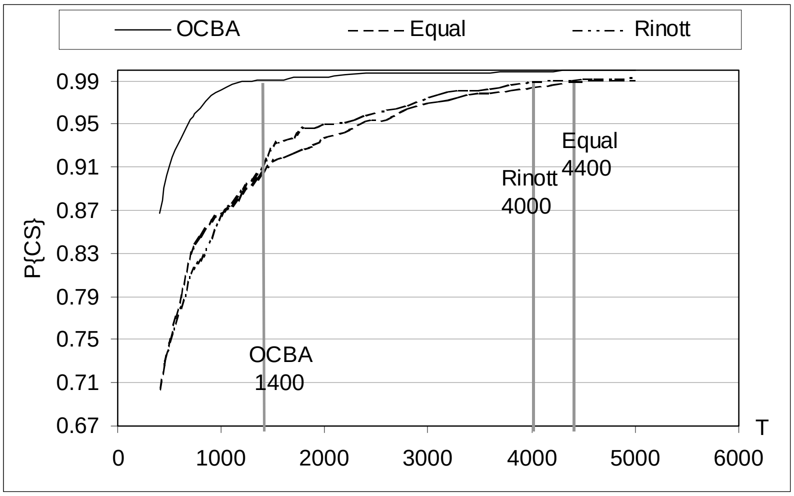

Figure 1: $P\{CS\}$ vs. the computation budget T for experiment 1. The computation cost for obtaining $P\{CS\}$=99% using different allocation procedures are indicated.

4.2.2 Experiment 2. Uniform Distribution

We consider a non-normal distribution for the performance measure: $X_{ij} \sim$ Uniform (i-10.5, i+10.5), i = 1, 2, .., 10. The endpoints of the uniform distribution are chosen such that the corresponding variance is close to that in experiment 1. Again, we want to find a design with the minimum mean and design 1 is the actual best design. All other settings are identical to experiment 1. Figure 2 contains the simulation results for all compared allocation procedures. We can see that the relative performances of different procedures are very similar with what we see in experiment 1. OCBA is much faster than Rinott and equal allocation.

4.2.3 Experiment 3. Larger Variance

This is a variation of experiment 1. All settings are preserved except that the variances for each design is doubled. Namely, $L(\theta_i, \xi)) \sim N(i, 2\cdot 6^2)$, i = 1, 2, .., 10. Figure 3 contains the simulation results for the compared three allocation procedures. We can see that the relative performances of different procedures are very similar with what we see in previous experiments, except that bigger computing budgets are needed in order to obtain a same $P\{CS\}$, due to larger variance. Once again, OCBA is much faster than the other two allocation procedures.

5 CONCLUSIONS

We have shown that the algorithm based on the OCBA technique is a powerful tool for selecting the best design out of k (simulated) alternatives. The great advantage of the OCBA technique is that the algorithm dynamically determines the simulation lengths for all simulationexperiments and thus significantly improves simulation efficiency with a given computing budget. Numerical results show that algorithms based on OCBA can indeed significantly improve simulation efficiency.

ACKNOWLEDGEMENTS

This work has been supported in part by NSF under grant DMI-9732173, by Sandia National Laboratories under contract BD-0618, and by the University of Pennsylvania Research Foundation.

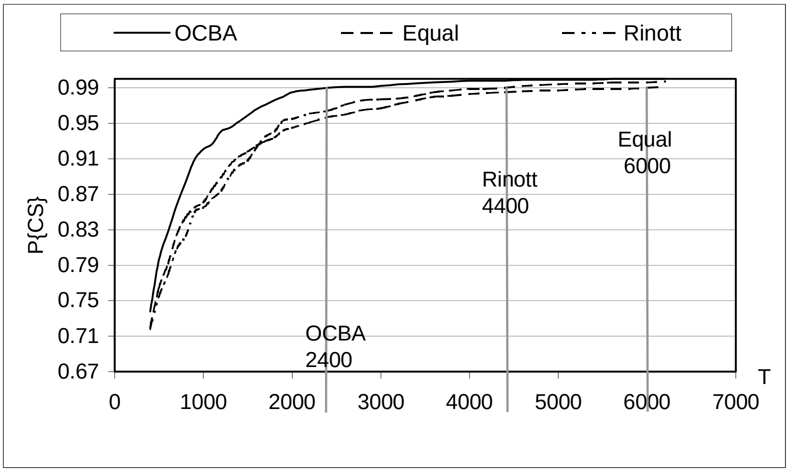

Figure 2: $P\{CS\}$ vs. the computation budget T for experiment 2. The computation costs for obtaining $P\{CS\}$=99% using different allocation procedures are indicated.

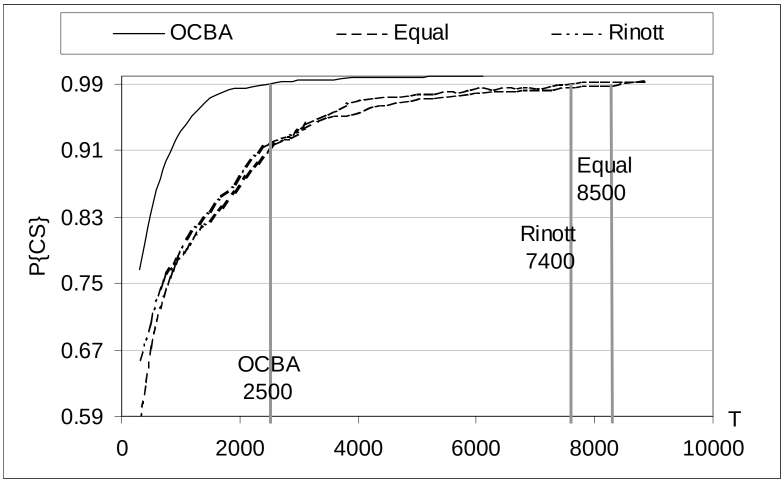

Figure 3: $P\{CS\}$ vs. the computation budget T for experiment 3. The computation costs for obtaining $P\{CS\}$=99% using different allocation procedures are indicated.

REFERENCES

Bechhofer, R. E., C. W. Dunnett, and M. Sobel 1954. A two-sample multiple decision procedure for ranking means of normal populations with a common unknown variance. *Biometrika*, 41, 170-176.

Bechhofer, R. E. 1969. Optimal allocation of observations when comparing several treatments with a control. In Multivariate Analysis II, P. R. Krishnaiah, ed., Academic Press, Inc. 463-473. Berger

Bechhofer 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., 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*. John Wiley.

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. 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., 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.

Chen, H. C., C. H. Chen, L. Dai, E. Yücesan, 1997. New Development of Optimal Computing Budget Allocation for Discrete Event Simulation. *Proceedings of the 1997 Winter Simulation Conference*, 334-341.

Chen, C. H., E. Yücesan, Y. Yuan, H. C. Chen, and L. Dai,. 1998b. "Computing Budget Allocation for Simulation Experiments with Different System Structures," to appear in the *Proceedings of the 1998 Winter Simulation Conference*.

Chen, C. H., S. D. Wu, and L. Dai. 1999a. "Ordinal Comparison of Heuristic Algorithms Using Stochastic Optimization," *IEEE Transactions on Robotics and Automation*, 15: 44-56.

Chen, H. C., C. H. Chen, and E. Yücesan. 1999b. "Computing Efforts Allocation for Ordinal Optimization and Discrete Event Simulation," To appear in *IEEE Transactions on Automatic Control*.

Chick, S. E., 1997a. Bayesian Analysis for Simulation Input and Output. *Proceedings of the 1997 Winter Simulation Conference*, 253-260.

Chick, S. E., 1997b. Selecting The Best system: A decision-Theoretic Approach. *Proceedings of the 1997 Winter Simulation Conference*, 326-333.

Chick, S. E. and Inoue, K., 1998, “How to More Effectively Allocate Computing Resources to Identify The Best Simulated Alternative” working paper.

Chiu, W. K., 1974. Selecting the *n* populations with Largest Means From *k* normal Populations with Unknown Variances. *Australian Journal of Statistics*, 16, 144-147.

De Groot, M. H. 1970. *Optimal Statistical Decisions*. McGraw-Hill, Inc.

Dudewicz, E. J. and S. R. Dalal, 1975. Allocation of Observations in Ranking and Selection with Unequal Variances. *Sankhya*, B37:28-78.

Dunnett, C. W. and M. Sobel, 1954. A Bivariate Generalization of Student's *t* distribution, with Tables for Certain Special Cases, *Biometrika*, 41, 153-169.

Futschik, A. and G. C. Pflug, 1997. Optimal Allocation of Simulation Experiments in Discrete Stochastic Optimization and Approximative Algorithms. *European Journal of Operational Research*, 245-260.

Goldsman, G., B. L. Nelson, and B. Schmeiser, 1991. Methods for Selecting the Best System. *Proceedings of the 1991 Winter Simulation Conference*, 177-186.

Goldsman, G., 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 S. Panchapakesan, 1979. *Multiple Decision Procedures: Theory and Methodology of Selecting and Ranking Populations*. John Wiley.

Gupta, S. S. and J. O. Berger, 1988. *Statistical Decision Theory and Related Topics IV*. Springer Verlag.

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*.

Kleijnen, J.P.C., 1987. *Statistical Tools for Simulation Practitioners*. Dekker.

Law, A. M. and W. D. Kelton, 1991. *Simulation Modeling & Analysis*. McGraw-Hill, Inc.

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

Patsis, N. T., C. H. Chen, and M. E. Larson. 1997. "SIMD Parallel Discrete Event Dynamic System Simulation," *IEEE Transactions on Control Systems Technology*, 5:30-41.

Rinott, Y., 1978. On Two-stage Selection Procedures and Related Probability Inequalities. *Communications in Statistics* A7, 799-811.

Ross, S. 1994. *A First Course in Probability*, Prentice Hall Inc.

Stein, C., 1945. A two-sample test for linear hypothesis whose power is independent of the variance. *The Annals of Mathematical Statistics*, 16, 243-258.

Tong, Y. L. and D. E. Wetzell, 1984. Allocation of Observations for Selecting The Best Normal Population. 213-224. *Design of Experiments : Ranking & Selection*, 213-224.

AUTHOR BIOGRAPHIES

HSIAO-CHANG CHEN is currently a senior engineer at SynQuest Inc. in Atlanta, Georgia. He received a B.S. degree in Mathematics and Computer Science from the Eastern Michigan University, an M.S. degree in Systems Science and Mathematics from Washington University, MO, and a Ph.D. degree in Systems Engineering from University of Pennsylvania. His research interests include developing efficient approaches for discrete event simulation, stochastic optimization, and using simulation in the planning and design of supply chain.

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.

JIANWU LIN is a Ph.D. candidate at the Systems Engineering Department, University of Pennsylvania. He received a B.E. degree in Biomedical Engineering and a dual degree in Industrial Engineering from Tsinghua University in 1996. And he received a M.S. in Biomedical Engineering from Tsinghua University in 1998. He is working on developing efficient approaches for discrete event simulation.

ENVER YÜCESAN is a Professor of Operations Research at INSEAD in Fontainebleau, FRANCE. He holds a BSIE degree from Pudue 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. His research interests include web-based simulation, systems design and optimization, and supply chain management.