A BAYESIAN APPROACH TO THE DESIGN OF SIMULATION EXPERIMENTS

Mark A. Turnquist
Northwestern University

Joseph M. Sussman
Massachusetts Institute of Technology

INTRODUCTION

Simulation is inherently an experimental modelling tool. One obtains results from simulation models by operating them, or in other terms, by conducting experiments with them. The intrinsic experimental nature of simulation (and its resulting costliness as a modelling methodology) implies that in order to utilize simulation effectively, it is important to make the fullest possible use of the tools of experimental design and statistical analysis.

Effective design of simulation experiments requires consideration of a number of issues, including:

- the context of the experiments (i.e., deterministic or stochastic model, continuous or discrete domain of experimentation, etc.);
- the nature of the decisions to be made, based on the outcome of the experiments;
- the special experimental capabilities of simulation models.

The approach to experimental design described here is oriented toward stochastic models, and to situations involving discrete experimental options, as opposed to experiments in a continuous space. A Bayesian approach has several advantages over previous methods for dealing with problems of this type. As preface to a discussion of this Bayesian approach, the following three sections clarify the special experimental capabilities of simulation, assumptions regarding the context of the experiments and nature of resulting decisions, and previous research in this area.

SIMULATION AND SYSTEMS PLANNING

The purpose of simulation experiments is to acquire additional understanding about the system under study in order to improve the decisions made by the analyst. Effective design of experiments must thus recognize the nature of the decisions to be made, in order to assess the value of the information which might be obtained through any of several possible experiments.

Simulation may often be used very effectively in the initial planning stage of a project, when several basic alternative system configurations are being evaluated and compared. The product of the initial planning stage is the selection of a basic system configuration to undergo detailed design. The decision which must be made by the analyst is thus a selection of that alternative which is "best," in some sense, from among a set of possible alternatives. The decisions made at this stage are extremely important because they strongly influence the final outcome of the project. No amount of detailed design work can overcome the effects of an incorrect decision with regard to basic system configuration.

The cost of an incorrect decision at the initial planning stage may be considered to be an opportunity cost associated with selection of an inferior alternative. That is, if alternative "i" is selected, subjected to detailed design and implemented, a certain measure of system performance will be achieved. If, however, another alternative, say "j," is really the best alternative, selection and implementation of "j" would have resulted in a superior measure of performance. The cost of the incorrect decision in this case is the opportunity loss, or difference between what is achieved and what might have been achieved.

If one is designing experiments to compare alternatives, the potential value of any possible experiment is thus dependent upon the cost of an incorrect decision, and should be incorporated explicitly into the experimental design process. This is a major aspect of the approach described here. A second major aspect of this approach is that it effectively utilizes the capabilities of simulation as an experimental environment. Simulation does provide very special opportunities to the experimenter, and the nature of these opportunities must be fully understood.

SPECIAL EXPERIMENTAL CAPABILITIES IN SIMULATION

The experimental environment provided by simulation is a unique one because of its controllability and reproducability.

If one is constructing a set of experiments to investigate the effect of varying a certain input variable, the ability to control all other variables is very important. In this way, the effect of the change in which the analyst is interested can be isolated, and not obscured by the effects of other input variables.

The reproducability of the simulation experimental

environment stems from the "pseudo-random" nature of the random elements in simulation models. Because of this "pseudo-randomness," parallel experiments under the control of the same random-number streams can be performed. This gives the analyst opportunities to reduce the variance of some of the sample statistics of interest (or to use less experimental effort to obtain a given level of confidence in these model outputs). Such techniques are typically referred to as "variance reduction" methods. Among the better-known of these are antithetic variates and common random-number streams. Considerable prior work has been done in the analysis of these techniques by Page [11], Fishman [2], Kleijnen [4], Moy [9], and others.

Antithetic sampling is a variance reduction technique which is useful in the estimation of the mean response from a simulation model for a particular set of input conditions. The basic motivation for antithetic sampling can be illustrated by a simple example in which two simulation runs are performed, and two samples, Y_1 and Y_2, are drawn from the distribution of the model output, Y.

The expected value of Y may be estimated by the sample mean:

$$\bar{Y} = \frac{Y_1 + Y_2}{2} \tag{1}$$

and in general, the variance of $\bar{Y}$ is given by:

$$V(\bar{Y}) = \frac{V(Y)}{2} + \frac{\text{cov}(Y_1, Y_2)}{2} . \tag{2}$$

If the samples are independent, $\text{cov}(Y_1, Y_2) = 0$, and the result is the familiar one:

$$V(\bar{Y}) = \frac{V(Y)}{2} . \tag{3}$$

However, if $\text{cov}(Y_1, Y_2)$ is negative, the variance of $\bar{Y}$ will be reduced, implying a more efficient estimate of E(Y) than could be achieved through independent sampling.

Antithetic sampling is one of several ways in which the analyst can attempt to induce negative values of $\text{cov}(Y_1, Y_2)$. The various methods available have been surveyed by Moy [9], and more recently by McGrath and Irving [8]. The attractiveness of antithetics stems from the fact that it is one of the easiest techniques to apply, and at the same time is one of the most effective.

The application of antithetic sampling involves the performance of pairs of simulation runs. In the first run, a stream of random numbers, RN_1, RN_2,..., is transformed into samples of the various inputs to the simulation model. The result of this simulation is a sample, Y, from the output distribution, f(Y). For the second run, the so-called "antithetic" random-number stream, $1-RN_1$, $1-RN_2$, ..., is used, resulting in output samples which are negatively correlated with those from the first run, and producing an output sample, Y', also from f(Y), which is (one hopes) negatively correlated with the sample, Y, from the first run.

Common random numbers is a technique which may be used for variance reduction in the comparison of alternative system configurations, or options. To illustrate the motivation for this technique, consider a simple example with two options, denoted λ_1 and λ_2. Comparison of these two options can be done by defining the difference in response, Z:

$$Z = (Y|\lambda_1) - (Y|\lambda_2). \tag{4}$$

Mean value comparison then involves the expected value of Z:

$$\begin{aligned} E(Z) &= E[(Y|\lambda_1) - (Y|\lambda_2)] \\ &= E(Y|\lambda_1) - E(Y|\lambda_2). \end{aligned} \tag{5}$$

This mean difference can be estimated by the sample mean difference, $\bar{Z}$:

$$\bar{Z} = (\bar{Y}|\lambda_1) - (\bar{Y}|\lambda_2), \tag{6}$$

with variance:

$$\begin{aligned} V(\bar{Z}) &= V(\bar{Y}|\lambda_1) + V(\bar{Y}|\lambda_2) \\ &\quad - 2\text{cov}(\bar{Y}|\lambda_1, \bar{Y}|\lambda_2). \end{aligned} \tag{7}$$

Clearly, if $\text{cov}(\bar{Y}|\lambda_1, \bar{Y}|\lambda_2)$ can be made positive, a variance reduction in Z will be achieved. Positive correlation between these sample means can be induced by constructing pairs of positively correlated samples through the use of common random numbers. By using a common random-number stream, RN_1, RN_2, ..., to sample the inputs for each option, the inputs are the same, and the outputs, $Y|\lambda_1$ and $Y|\lambda_2$, are (one hopes) positively correlated.

In light of these unique capabilities of the experimental environment available in simulation, it is clear that the traditional statistical tools for experimental design must be modified and extended if they are to be used most effectively for the design of simulation experiments. The following section describes previous work in development of experimental design procedures for simulation, and a new approach based on statistical decision theory is then presented.

PREVIOUS RESEARCH

The classical statistical framework for designing experiments to compare several alternatives is provided by multiple ranking procedures. These procedures are discussed in detail by Bechhofer [1], and their application to simulation has been studied by Naylor, et al. [10], Kleijnen, et al. [6], and Schmidt, et al. [13]. The major drawback of multiple ranking procedures is that they make rather inefficient use of the simulation environment. A basic assumption of these procedures is that all samples are independent, and thus the use of variance reduction techniques, such as antithetics and common random numbers, is specifically excluded.

In recent years, groundbreaking work has been done by Fishman [3], Kleijnen [5] and others in the development of experimental design methods which are explicitly simulation-oriented and are able to make effective use of various reduction techniques. To date, however, the available methods are limited to pairwise comparisons and are based on a design objective of minimizing the variance of the estimated difference in performance between the two alternatives.

This minimum-variance criterion serves as a surrogate for the objective of minimizing the risk of incorrect decision by the analyst. It only indirectly addresses the decision of which alternative is the best, since there is no inclusion of the mean difference.

The approach described here is an evolutionary step in improving methods for the design of simulation experiments. It incorporates both the capability to handle more than two alternatives at a time, and explicit recognition of the costs of incorrect decisions into the experimental design framework, while retaining the ability to exploit the special experimental capabilities of simulation.

A DECISION-THEORETIC APPROACH TO EXPERIMENTAL DESIGN

In order to provide explicit consideration of the consequences of decisions to be made on the basis of experimental results, a decision-theoretic approach to the design of these experiments has been adopted. The formulation is based on previous work in decision theory and experimental design by Raiffa and Schlaifer [12], and Lin [7].

The basis of this approach to experimental design is the multivariate Normal-Normal probability model. The following discussion presents a brief description of the method. For a more complete presentation, the interested reader is referred to Turnquist [14].

If r alternatives or options are being considered and compared, the r mean responses may be written as a vector:

$$\underline{\mu} = [\mu_1, \mu_2, \ldots, \mu_r] \tag{8}$$

The decision problem faced by the analyst is then to identify from among this set the option which produces the largest (or smallest) mean response. The maximization and minimization problems are equivalent mathematically, so the development here will be phrased in terms of selecting the largest mean value.

The analyst may express his uncertainty about the unknown values of the elements of $\underline{\mu}$ in the form of a multivariate probability density function, $f(\underline{\mu})$. In the Normal-Normal model, $f(\underline{\mu})$ is a multivariate normal density function with mean vector

$$\underline{m} = [m_1, m_2, \ldots, m_r] \tag{9}$$

and covariance matrix

$$\underline{\underline{S}} = \begin{bmatrix} s_{11} & s_{12} & \cdots & s_{1r} \\ s_{21} & & & \\ \vdots & & & \\ s_{r1} & \cdot & \cdot & s_{rr} \end{bmatrix} \tag{10}$$

where s_{ij} represents the covariance of μ_i and μ_j.

If the analyst decides that μ_i is the largest mean, when in fact the largest is μ_j, the penalty, or loss, incurred as a result of the decision is assumed to be:

$$L(\mu_j, \mu_i) = \mu_j - \mu_i. \tag{11}$$

This linear loss function is the formal representation of the notion of opportunity loss discussed earlier. Clearly, if the correct decision is made (i.e., i=j), the loss is zero. For any incorrect decision, a penalty is incurred.

Given the representation of the analyst's state of uncertainty, $f(\underline{\mu})$, and the loss structure associated with incorrect decisions, $L(\mu_j, \mu_i)$, the optimal Bayes decision (i.e., minimizing expected loss) is simply to select the option corresponding to the largest element of the vector, $\underline{m}$. This decision has associated with it a certain level of risk, or expected loss.

In an attempt to reduce the risk associated with making a decision, the analyst may try to obtain more information by conducting experiments. Experimental observations are taken on the output of the simulation model, a random variable, Y. A single observation (say, on option i) may be written in the form

$$y = \mu_i + w \tag{12}$$

where μ_i is the mean response for option i and w is a random disturbance about the mean, due to the stochastic nature of the simulation model. The Normal-Normal probability model requires that the random disturbance be normally distributed with mean zero.

In general, a set of observations from n experiments can be written in matrix form as follows:

$$\underline{y}^{(n)} = \underline{\underline{A}}^{(n)} \underline{\mu}^T + \underline{w}^{(n)} \tag{13}$$

where

- $\underline{y}^{(n)}$ = n x 1 vector of experimental observations
- $\underline{\underline{A}}^{(n)}$ = n x r coefficient matrix relating each element of $\underline{y}^{(n)}$ to a specific option, k
- $\underline{w}^{(n)}$ = n x 1 vector of random errors, distributed multivariate normal with a zero mean vector and covariance matrix denoted $\underline{\underline{\Sigma}}^{(n)}$.

The $\underline{\underline{A}}^{(n)}$ matrix is composed of zeros and ones. Each row corresponds to an experiment, and has r-1 zeros and a single 1 in the k^{th} position (if the experiment is run on the k^{th} option).

The elements of the covariance matrix, $\underline{\underline{\Sigma}}^{(n)}$, are

assumed known. Both this known variance assumption and the normality assumption for the errors, $\underline{w}^{(n)}$, are likely to be violated often in practice. Because of this, an important aspect of the research has been sensitivity testing for the effects of various types of violations. These sensitivity tests, described by Turnquist [14], indicate that the method exhibits considerable robustness relative to these assumptions.

Given equations (9), (10) and (13), the process of incorporating new information into the prior distribution, $f(\underline{\mu})$, is very straightforward, as shown by Raiffa and Schlaifer [12]. The covariance matrix and mean vector of the posterior distribution (including the data) are as shown in equations (14) and (15):

$$\underline{\underline{SP}}^{-1} = \underline{\underline{S}}^{-1} + \underline{\underline{A}}^{(n)T}[\underline{\underline{\Sigma}}^{(n)}]^{-1}\underline{\underline{A}}^{(n)} \qquad (14)$$

$$\underline{mp} = \underline{m} + \underline{\underline{SP}}\,\underline{\underline{A}}^{(n)T}[\underline{\underline{\Sigma}}^{(n)}]^{-1}(\underline{y}^{(n)} - \underline{\underline{A}}^{(n)}\underline{m}^{T}) \qquad (15)$$

where $\underline{\underline{SP}}$ = posterior covariance matrix

$\underline{m}$ = posterior mean vector.

Equations (14) and (15) tell us how to incorporate the data we gather into the representation of the state of our knowledge about the system under study. We can now focus on the experimental design issues —that is, deciding what data to gather.

In this decision-theoretic framework, the problem of designing the $(n+1)^{st}$ experiment becomes one of specifying the $(n+1)^{st}$ rows to add to the $\underline{\underline{A}}^{(n)}$ and $\underline{\underline{\Sigma}}^{(n)}$ matrices. These rows define the option on which an experiment is to be run, and the correlation structure with previous experiments.

The basic criterion for designing optimal experiments is the Expected Value of Sample Information (EVSI). This is a measure of the expected reduction in risk which would result from the availability of the new information associated with an experimental outcome. EVSI is a function of the experiment being planned, and explicitly incorporates the "cost of being wrong" in the selection of the best system alternative.

In general, one compares EVSI, as a measure of the benefit from an experiment, with the cost of performing that experiment, to obtain a measure of net benefit for each possible experiment. The experiment with the largest net benefit is then selected as the next one to be performed. After the performance of each experiment, the updating equations, (14) and (15), are applied, and the next experiment is planned. When the EVSI criterion suggests that no further experiments are worthwhile, the process terminates.

The method for locating the optimal next experiment utilizes the structure due to the "pairwise" nature of antithetic sampling and the "r-tuple" nature of common random numbers. Antithetic sampling is applied by generating pairs of negatively correlated observations with each option. Common random numbers are applied by making a run under each option with the same stream of random numbers.

This covariance structure can be illustrated as shown in Figure 1, for the case in which there are just two options (r = 2).

FIGURE 1

Covariance Structure of Experiments for Two-Option Problem (r = 2)

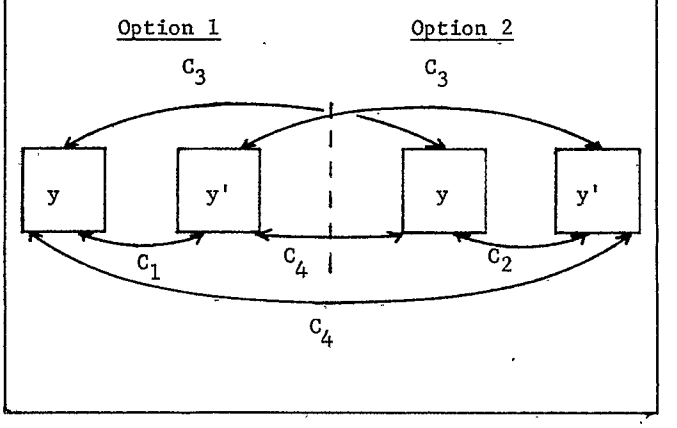

The covariances indicated in Figure 1 are as follows:

- C_1 = negative covariance due to antithetic sampling on option 1;
- C_2 = negative covariance due to antithetic sampling on option 2;
- C_3 = positive covariance due to common random numbers between options 1 and 2;
- C_4 = negative covariance due to simultaneous use of antithetics and common random numbers.

In general, we can visualize a structure of "2r-tuples" of correlated experiments.

This structure implies that the experimental design problem can be viewed as a process of filling "slots" in a tableau, as shown in Figure 2. Each row of the tableau comprises a "2r-tuple" of correlated experimental possibilities, and the rows are all independent from one another.

FIGURE 2

Experimental Design Tableau

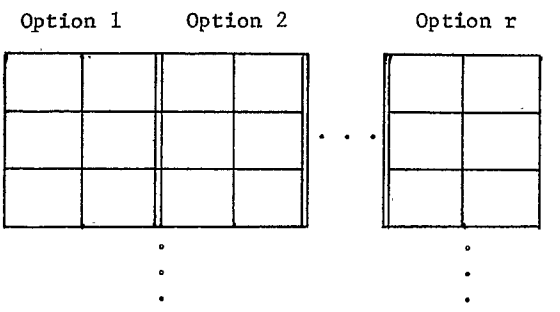

Before experimentation is begun, the tableau can be visualized as being completely empty; and as the experimental program proceeds, various boxes are filled by running the particular experiment which that box connotes. The experimental design process thus amounts to searching over unfilled boxes in the tableau, each of which corresponds to an experiment, and selecting the experiment which results in the largest expected net benefit.

If all available experiments at a given stage are of equal cost, the optimal experiment is clearly the one with maximum value of EVSI. However, if the available experiments are of unequal cost, or when the expected informational value of the optimal experiment may not be sufficient to cover its cost, a measure of net benefit must be constructed for each experiment.

If the units of EVSI are monetary, and the benefits measured by EVSI are accruing to the same group as the cost of the experiment, the simple procedure of computing the Expected Net Gain of Sampling (ENGS) is adequate:

$$ENGS(\epsilon) = EVSI(\epsilon) - C(\epsilon) \qquad (16)$$

where $C(\epsilon)$ is the cost of the experiment, ϵ.

However, the units of EVSI are often not monetary, and in many cases the benefits and costs of experimentation may be accruing to different people. In these situations, the comparison of benefits and costs must be less direct.

In order to understand the dimensions of the benefit/cost trade-off, it is important to recognize that the units of EVSI are always the same as the units of the simulation output. The job of the analyst is then to weigh the value of information, expressed in these units, against the cost of experimentation, expressed in monetary terms, and decide which, if any, experiments to perform.

AN EXAMPLE APPLICATION

As an example of the application of this experimental design method, consider the following situation. Suppose we are operating a system which may be modeled as a single-server queuing system with a FIFO queue. Customers arrive in Poisson fashion, with mean inter-arrival time of ten minutes. Service time is exponentially distributed with mean nine minutes.

Because of the degree of congestion in this system, we are considering three alternatives for increasing the service capacity of the system:

1. Replace the current server with a new server which works twice as fast —i.e., has exponential service times with mean 4.5 minutes;
2. Replace the current server by a single Erlang-2 server with mean service time of 5.2 minutes;
3. Add another server of the same type as currently in use, to make a two-parallel-server system, in which customers in the single FIFO queue enter the first available server.

Further suppose that these are equal-cost alternatives, and that the criterion for evaluation is a weighted combination of the mean total time (wait plus service) of customers in the system, and the probability of having a very long wait time. Specifically, let the criterion be:

$$Y = E(W_t) + 100\ P(W_q>30) \qquad (16)$$

where W_t = total system time (minutes)

W_q = wait time in queue (minutes).

Note the weighting coefficient on $P(W_q>30)$ has an implicit unit of "minutes" in order to eliminate dimensionality problems.

Note also that this problem may be solved analytically, so that in practice, simulation would not be used. However, it does provide a convenient example problem.

The required inputs from the analyst are a prior distribution and a set of variance and covariance values. The prior distribution is specified as "non-informative," with the following mean vector, $\underline{m}$, and covariance matrix, $\underline{S}$:

$$\underline{m} = [30\ 30\ 30]$$

$$\underline{S} = \begin{bmatrix} 10^6 & .5 \times 10^6 & .5 \times 10^6 \\ .5 \times 10^6 & 10^6 & .5 \times 10^6 \\ .5 \times 10^6 & .5 \times 10^6 & 10^6 \end{bmatrix}$$

This prior distribution is "non-informative" because of the very large variances (diagonal elements of $\underline{S}$).

This "three-option" problem requires estimates of three variances and nine covariances (a simple extension of the structure illustrated in Figure 1). The values used are as follows:

$V_1 = 60$	$C_1 = -30$	$C_5 = -30$
$V_2 = 60$	$C_2 = -30$	$C_6 = 15$
$V_3 = 60$	$C_3 = 15$	$C_7 = -10$
	$C_4 = -10$	$C_8 = 15$
		$C_9 = -10$

The argument leading to the selection of these values will not be given here. The interested reader may refer to Turnquist [14].

A single experiment, or simulation run, is defined to be the servicing of 200 customers, beginning with the system empty. The experimental design results in the performance of 34 simulations, arranged as follows (using the tableau notation):

Option 1		Option 2		Option 3	
2	5	1	4	3	6
9	11	7	8		
10	12	14	13	16	15
17	18	19	20		
21	22	23	24		
25	27	28	26		
29	31	33	32	34	30

The criterion for terminating experimentation is that the maximum value of EVSI for any possible experiment is less than .01 (minutes). This represents a judgment on the part of the analyst that an experiment which is expected to yield less information than this is not worth the cost of running. Note that the units of EVSI are the units of the response, Y, and the analyst is required to make the trade-off between this value and the cost of an experiment.

It is clear that the bulk of the experimental effort is directed toward option 1 (the single faster exponential server) and option 2 (the Erlang-2 server). A few experiments with option 3 (the parallel-server system) suffice to eliminate it from contention. The final predicted values of Y for the three systems are indicated in Table 1, along with the true values (from analytic solution) for comparison.

TABLE 1

Predicted Performance of the Three Options

	True Value	Predicted Value
Y_1	9.33	8.72
Y_2	10.505	9.18
Y_3	12.37	13.84

APPLICATIONS TO MORE COMPLEX MODELS

Although it will not be discussed here, a major case study is described by Turnquist [14]. This case study involves the application of this experimental deisgn method to a problem of comparing alternative routing algorithms for Dial-A-Ride, a demand-responsive transit system.

The results of this case study indicate that the approach is very useful, and can lead to substantially more effective experimentation with complex simulation models.

CONCLUSIONS

This discussion of experimental design has had two major goals. First, it has attempted to focus attention on the unique experimental capabilities offered by the simulation environment. The controllability and reproducibility of experimental conditions in this environment require that experimental design tools be modified and extended to make the fullest possible use of techniques such as antithetic variates and common random numbers.

The second goal of this discussion has been to introduce a more complete methodological framework for the design of simulation experiments than has previously been available. This framework is oriented toward the special experimental capabilities of simulation in the planning of complex systems and the nature of decisions to be made based on the experimental results.

It is the opinion of the authors that the decision-theoretic approach described in this paper represents progress toward both of these goals. It is the hope of the authors that this work will stimulate greater interest in the problems of experimental design for simulation models, and that even more progress will be made as a result of this interest.

BIBLIOGRAPHY

1. Bechhofer, R.E., "A Sequential Multiple Decision Procedure for Selecting the Best One of Several Normal Populations with a Common Unknown Variance, and Its Use with Various Experimental Designs," BIOMETRICS, 14 (1958), pp. 408-429.
2. Fishman, G.S., "Variance Reduction in Simulation Studies," JOURNAL OF STATISTICAL COMPUTATION AND SIMULATION, 1:2 (April 1972), pp. 173-182.
3. Fishman, G.S., "The Allocation of Computer Time in Comparing Simulation Experiments," OPERATIONS RESEARCH, 16:2 (1968), pp. 280-295.
4. Kleijnen, J.P.C., "Variance Reduction Techniques in Simulation," Ph.D. dissertation, Katholieke Hogeschool, Tilburg, Netherlands, January 1971.
5. Kleijnen, J.P.C., "Antithetic Variates, Common Random Numbers and Optimal Computer Time Allocation in Simulation," mimeograph, February 1972.
6. Kleijnen, J.P.C., T.H. Naylor and T.C. Seaks, "The Use of Multiple Ranking Procedures to Analyze Simulations of Management Systems: A Tutorial," MANAGEMENT SCIENCE, 18:6 (February 1972), pp. B-245.
7. Lin, C.Y., "A Decision Theoretic Approach to the Design and Analysis of Industrial Experiments," Ph.D dissertation, Sloan School of Management, M.I.T., June 1968.
8. McGrath, E.J., and D.C. Irving, "Applications of Variance Reduction to Large Scale Simulation Problems," COMPUTERS AND OPERATIONS RESEARCH, 1:2 (August 1974), pp. 283-312.
9. Moy, W.A., "Monte Carlo Techniques: Practical," in Taylor, T.H., ed., THE DESIGN OF COMPUTER SIMULATION EXPERIMENTS, Duke University Press, Durham, N.C. 1969.
10. Naylor, T.H., K. Werty and T.H. Wonnacott, "Methods for Analyzing Data from Computer Simulation Experiments," COMMUNICATIONS OF THE ACM, 10 (November 1967), pp. 703-710.
11. Page, E.S., "On Monte Carlo Methods in Congestion Problems: II. Simulation of Queuing Systems," OPERATIONS RESEARCH, 13:2 (March-April 1965), pp. 300-305.
12. Raiffa, H., and R. Schlaifer, APPLIED STATISTICAL DECISION THEORY, The M.I.T. Press, Cambridge, Mass., 1961.
13. Schmidt, J.W., R.E. Taylor and G.K. Bennett, "A Heuristic Elimination Method for Optimum Selection of Alternatives Through Simulation," SIMULATION, 21:5 (November 1973), pp. 129-134.
14. Turnquist, M.A., "A Bayesian Approach to Simulation —Specific Experimental Design with Application to Modelling Transportation Systems," Ph.D. dissertation, Department of Civil Engineering, M.I.T., May 1975.