Proceedings of the 1988 Winter Simulation Conference
M. Abrams, P. Haigh, and J. Comfort (eds.)

Designing computer simulation experiments

W. David Kelton
Department of Management Sciences
Carlson School of Management
University of Minnesota
Minneapolis, Minnesota 55455

ABSTRACT

This tutorial focuses on that part of a simulation study concerning setting the various specific model parameters and the experimental conditions under which the model will be exercised. Other issues in setting up and designing simulation experiments, such as variance reduction, ranking and selection, and optimization, are also mentioned. The focus is on careful choice of such parameters beforehand, with an eye toward the statistical analysis of the simulation's results. Output analysis is not treated per se, being covered in another tutorial in this conference (by Gordon Clark), but the design and analysis activities must be done hand in hand.

1. INTRODUCTION AND SCOPE

Modeling and coding a simulation is in itself a daunting task, often consuming vast resources. Thus, a sigh of relief is certainly in order when a model is finally built, coded, verified (debugged), and validated in some way. However, the point of doing all of this work is so that the model can be <u>used</u> in an organized, scientific way to learn about the system being simulated. Such use of a simulation model can be thought of as being composed of two parts: <u>design</u> and <u>analysis</u>.

The <u>design</u> phase is concerned with setting the various parameters of the model, and deciding on the conditions under which the model will be run. For instance, in a manufacturing study we would need to decide which specific scenarios to consider, and also the length and number of simulation runs of each scenario. Thus, the design phase must be done before the main runs of the simulation are carried out.

The <u>analysis</u> phase is concerned with interpreting and using the output from the simulation runs to infer something about the system. In the manufacturing example, this could include building a confidence interval for the expected daily throughput of the system. The analysis that is possible (and appropriate) depends on what was decided in the design phase; for example, we would be able to make inferences on only those scenarios considered, and the accuracy and precision of our inferences could depend on the length and number of runs we decided to make.

Thus, the design and analysis phases of a simulation study must be done hand in hand. This tutorial will focus on the design phase, and another tutorial in this conference (by Gordon Clark) treats the analysis phase. However, since the design must look forward to the analysis, some aspects of the latter will have to be mentioned here.

Generally, our focus will be on dynamic, stochastic simulation models, for this probably includes the preponderance of industrial simulation applications, and in addition poses the most difficult design (and analysis) problems. However, many of the ideas to be discussed here could be used as well for deterministic or static models. For example, it may be that a deterministic model of a large system is itself very expensive to run, and we thus must take care in choosing which system variants we can evaluate directly with the model, to allow for inferring about other variants which are not possible to evaluate. Also, in estimating a complex multidimensional integral by (static) Monte Carlo simulation, we would have to decide on the number of independent replications to make in order to achieve the desired level of precision in the estimate.

The tutorial will cover methods for setting those parameters that are tied to the <u>model</u> (in Section 2), and then discuss setting the parameters having to do with the <u>experimental conditions</u> (Section 3). Section 4 then touches on some other issues that must be considered in the design phase.

2. SETTING MODEL FACTORS

In most simulation models there are <u>parameters</u> or, in classical statistical experimental design parlance, <u>factors</u>, that must be set to define the model itself. To take an example which will be carried through the tutorial, consider a small manufacturing facility consisting of some number m of identical machines operating in parallel, all fed by a single queue of parts. For this system, two model parameters of potential interest come immediately to mind:

(1) The number of parallel machines, m

(2) The order of service of the parts in queue, being either First-In-First-Out (FIFO) or Shortest-Job-First (SJF)

In this case, the possibilities for the second factor are simply given, and are qualitative, i.e., not numerical. The first factor (m), however, is less clearly defined and we must come up with a way to specify it numerically.

In addition to the above two factors, we must specify the manner in which parts arrive, and how they are processed on a machine. Typically, we specify a probability distribution to describe the arrivals, in terms of the duration of the times between successive arrivals, called interarrival times. Also, we would have to describe the service times by some probability distribution. Thus, we have two more model factors:

(3) The interarrival time distribution

(4) The service time distribution

While perhaps less graphic than the first two factors, the choice of these particular distributions (including both their form as well as their parameters) is quite important, having a potentially major impact on the model's behavior and results. The distribution-specification activity, then, must be regarded as part of the design of the simulation experiment having to do with the model-based parameters. Methods for selecting distributions from which to sample are discussed in most simulation texts, e.g., Banks and Carson (1984), Bratley, Fox, and Schrage (1987), Fishman (1973, 1978), or Law and Kelton (1982). In the case of specifying a multivariate input distribution, a comprehensive reference is Johnson (1987b). For discussions of distribution specification in past meetings of this conference, see Johnson (1987a) and Kelton (1984).

One way to specify the model-based parameters, such as the four outlined above, is to use traditional methods from classical statistical experimental design, such as those treated by Box, Hunter, and Hunter (1978). For example, we could use a two-level factorial design as a starting point, in which case we must specify two "values" for each parameter, resulting in 16 possible combinations (in our four-factor example); it may not be necessary to consider all 16 possible cases, if we can find an appropriate fractional factorial design. Then the prescribed combinations of factors would be simulated, and the results judged by traditional experimental design methods, e.g., analysis of contrasts. In the above example, there may be additional factors if we wanted to consider the various parameters of the interarrival and service time distributions individually. Again, in a previous meetings of this conference this class of techniques, in its specific application to simulation, was discussed by Kelton (1986).

If we were interested in a more elaborate exploration of the effects of the model's parameters on the output measures, we could use a more complete (and expensive) design, and view the output measures at the different design combinations as functions of the input factors (which, after all, they are, but related by the complex simulation model itself). If all the factors are numerical, we could actually try to fit a regression model, in which the output would be assumed to depend on the input in some relatively simple algebraic fashion. The result is often called a <u>metamodel</u> in the simulation literature, since it is a "model" of the simulation model itself. The metamodel is then used as a proxy for the full simulation model, perhaps to optimize some output measure as a function of the model parameters. This is the basic idea behind using the response surface methodology approach to optimization of simulation models.

In setting the model factors, it is seldom clear just what values they should be given. Typically this can only be done with some knowledge of the system itself. For example, in a simple two-level factorial design, we generally want to choose the levels to be extreme in some sense, which clearly calls for an element of judgement by someone familiar with the actual system being simulated.

3. SETTING EXPERIMENTAL FACTORS

In running a simulation experiment, we must not only decide on the model-based parameters discussed in Section 2, but also on the nature of the runs themselves. Here it appears that, instead of being motivated by concepts from classical statistical experimental design, we will be driven by the nature of the model's definition (especially its time frame of operation) and by the methods to be used to analyze the simulation output. Importantly, these decisions must be made before the simulation is run, rather than as an afterthought.

In a dynamic simulation, we must obviously specify the length of the simulation. In our manufacturing example, this could be defined in terms of some fixed amount of simulated time (e.g., 8 hours), or by some event (e.g., 5000 parts have been produced). Stopping rules of this type are thus really a part of the model itself, in that the output measures being estimated are dependent upon how the simulation is stopped.

On the other hand, we may want to estimate a long-run, or <u>steady-state</u> quantity, such as the long-run expected number of parts produced per hour, after the initial conditions of the simulation have worn off. In this case, we would in principle run the simulation forever; this being impractical (no matter how good or cheap microcomputers may get) we must stop the simulation at some point, and the form of this stopping rule can be a highly complex issue, tied intimately to the method to be used for output analysis.

Depending on the way the simulation is to be stopped, we may in addition face the issue of how many independent runs, or <u>replications</u> of the simulation (of each scenario) to make. An unfortunately widespread answer to this question is "one," i.e., only a single run is made. In a stochastic simulation, the output values are random, so that simply running the simulation once is just as dangerous as flipping a coin once and concluding that both faces are the same as the one that happened to come up.

Although actually an issue of output analysis, the decision on the length and number of runs must be made before the simulations are done, so are part of the design phase as well. One way to determine the answers to these questions is to let the simulation decide along the way, giving rise to sequential procedures for determining the length and number of runs; see Law (1983) for a survey of such methods. However, we still must specify ahead of time the parameters that will drive the sequential procedure, such as the maximally acceptable width of a confidence interval.

4. OTHER ISSUES

In this section we mention examples of other concerns that arise when getting ready to run a course of simulation experimentation.

4.1 Variance Reduction

In many stochastic simulations, the opportunity arises to run the simulation in something other than the obvious, straightforward method, and in return get estimates that are more precise (of lower variance) than would otherwise have been the case. Again, the decision to use a variance reduction technique must be made before the simulation is run, since it in general affects the way the model is actually coded, or the kinds of quantities that are observed during the simulation. Thus, planning to use (or not to use) a variance reduction technique is part of the design phase. For a recent tutorial on variance reduction techniques in an earlier meeting of this conference, see Nelson (1987).

4.2 Ranking and Selection

In most simulation projects we are interested in more than just a single variant of the system. In our manufacturing example, we would probably want to compare the effect of the different queue disciplines, as well as the effect of the number m of machines. Here we must choose which system variants are of interest, and using the classical statistical experimental design methods for setting the model-oriented factors discussed in Section 2 will help us make such a decision. Again, this is a pre-run decision, and should be considered a part of the design. Goldsman (1987) treats a version of this problem when there are two factors involved.

4.3 Optimization

To go a step beyond ranking and selection, we might be ambitious enough to attempt to optimize some measure of system performance. For example, we could place a profit measure on the operation of our manufacturing facility accounting for both the number m of machines (with higher m meaning lower profit) as well as the production (which would have a positive cost coefficient), and then try to find that combination of m and the queue discipline that would maximize profit. In general, true optimization of simulation models is difficult due to their stochastic nature. However, optimizing is a very good result if it can be attained, and we must design for it ahead of time. This can use classical response surface methods such as those discussed by Box, Hunter, and Hunter (1978), or perhaps more simulation-specific techniques discussed by Meketon (1987).

5. CONCLUSIONS

Careful design of simulation experiments, before the runs are made (at a cost which is frequently considerable), is ignored only at great peril. Speaking from personal experience, there are few moments in life as wrenching as realizing after the runs have been made (and the money spent) that they should have been done differently, and that this could have been foreseen. There appear to be two classes of design parameters, those tied to the model (Section 2) and those tied to the experimental conditions and the output analysis technique to be used (Section 3). The attempt in this tutorial has been to identify the nature of these parameters, and to discuss and make reference to techniques for setting them that has been found useful in the simulation context.

REFERENCES

Banks, J. and J.S. Carson II (1984). Discrete-Event System Simulation. Prentice-Hall, Englewood Cliffs, New Jersey.

Box, G.E.P., Hunter, W.G., and Hunter. J.S. (1978). Statistics for Experimenters: An Introduction to Design, Data Analysis, and Model Building. Wiley, New York.

Bratley, P., B.L. Fox, and L.E. Schrage (1987). A Guide to Simulation, 2nd ed. Springer-Verlag, New York.

Fishman, G. (1973). Concepts and Methods in Discrete Event Digital Simulation. Wiley, New York.

Fishman, G. (1978). Principles of Discrete Event Simulation. Wiley, New York.

Goldsman, D. (1987). Ranking and selection tutorial: 2-factor normal means procedures. In: Proceedings of the 1987 Winter Simulation Conference (A. Thesen, H. Grant, and W.D. Kelton, eds.). Institute of Electrical and Electronics Engineers, San Francisco, 52-57.

Johnson, M.E. (1987a). Modeling and generating input processes. In: Proceedings of the 1987 Winter Simulation Conference (A. Thesen, H. Grant, and W.D. Kelton, eds.). Institute of Electrical and Electronics Engineers, San Francisco, 30-32.

Johnson, M.E. (1987b). Multivariate Statistical Simulation. Wiley, New York.

Kelton, W.D. (1984). Input data collection and analysis. In: Proceedings of the 1984 Winter Simulation Conference (S. Sheppard, U. Pooch, and D. Pegden, eds.). Institute of Electrical and Electronics Engineers, San Francisco, 91-95.

Kelton, W.D. (1986). Statistical design and analysis. In: Proceedings of the 1986 Winter Simulation Conference (J. Wilson, J. Henriksen, and S. Roberts, eds.). Institute of Electrical and Electronics Engineers, San Francisco, 45-51.

Law, A.M. (1983). Statistical analysis of simulation output data. Operations Research 31, 983-1029.

Law, A.M. and W.D. Kelton (1982). Simulation Modeling and Analysis. McGraw-Hill, New York.

Meketon, M. (1987). Optimization in simulation: a survey of recent results. In: Proceedings of the 1987 Winter Simulation Conference (A. Thesen, H. Grant, and W.D. Kelton, eds.). Institute of Electrical and Electronics Engineers, San Francisco, 58-67.

Nelson, B.L. (1987). Variance reduction for simulation practitioners. In: Proceedings of the 1987 Winter Simulation Conference (A. Thesen, H. Grant, and W.D. Kelton, eds.). Institute of Electrical and Electronics Engineers, San Francisco, 43-51.

AUTHOR'S BIOGRAPHY

W. David Kelton is an Associate Professor in the Management Sciences Department at the Carlson School of Management, University of Minnesota, in Minneapolis. He received a B.A. in Mathematics from the University of Wisconsin-Madison, an M.S. in Mathematics from Ohio University, and M.S. and Ph.D. degrees in Industrial Engineering from the University of Wisconsin-Madison. His research interests are in simulation methodology, stochastic modeling and estimation, and quality control. He is a member of ORSA, TIMS, ASA, and SCS. He served as Program Chair for the 1987 Winter Simulation Conference.

W. David Kelton
Department of Management Sciences
Carlson School of Management
University of Minnesota
Minneapolis, Minnesota 55455
612/624-8503
dkelton@umnacvx.bitnet