Proceedings of the 1986 Winter Simulation Conference
J. Wilson, J. Henriksen, S. Roberts (eds.)

SIMULATION OPTIMIZATION USING FREQUENCY DOMAIN METHODS

Lee Schruben

School of O.R.I.E.
Cornell University
Ithaca, NY 14853

ABSTRACT

Simulation programs can be quite complex, sometimes involving a great number of input factors and parameters. Conventional experiments where each setting of the input values requires a separate simulation run can be quite time consuming and expensive; such experiments will be referred to as "run oriented" simulation experiments. An alternative approach is to allow the input variables in a simulation to vary according to specific patterns during a single run. Various output spectra can then be analyzed to gain information about the simulation; such experiments will be refered to as "frequency-domain" simulation experiments. This technique was initially developed primarily to aid in factor screening and to perform a global sensitivity analysis of the input parameters in a simulation. It has since been developed into a method for identifying a meta-model for the simulation response surface.

In this paper an overview of frequency domain simulation experiments is first presented. A new technique for including discrete factors in frequency domain experiments will also be discussed. Recently frequency domain optimality criteria have been discovered that can be used to identify local optima in the simulation response. The focus of this paper is on these optimality criteria. A brief discussion of how optimization algorithms might be designed using frequency domain simulation experiments is presented. This last topic is the subject of current research in frequency domain simulation experiments.

1. INTRODUCTION:

In frequency domain simulation experiments, input variables are changed according to specific patterns during a run of the program. The patterns used in such experiments might include sinusoidal oscillations and rectangular oscillations. The only requirement is that these patterns of variation form a complete orthogonal basis for the simulation output series when it is viewed as a vector in n dimensions where n is the number of observations collected during a run. The analogy is drawn between the simulation program and a "black box" electrical system such as an amplifier. This system has an input port for each input variable. The oscillations of each input variable are an application of power to this system.

A key element in the analysis of the output of such experiments is the assumption that the output response can be modeled as a polynomial function of the input variables. Such a function is often referred to as a meta-model of the simulation response. See Kleijnen, et.al, 1979 and Kleijnen, 1980 for discussions of the usefulness of meta-models in simulation.

Specifically, if Y represents the simulation response and $x_1, x_2, \ldots, x_p$ are the inputs, then the response is given by,

$$E[Y] = \beta_0 + \sum_{j=1}^{q} \beta_j t_j \qquad (1)$$

where

Y is the response

$t_1, \ldots, t_q$ are possible terms in the polynomial model; they are products of non-negative integral powers of continuous parameters $x_1, \ldots, x_p$.

$\beta_0, \ldots, \beta_q$ are real-valued coefficients.

Initally the experiment specifys a prospective polynomial response surface model, say, all terms of order less than some number, k. An experimental technique for identifying the functional form of the meta-model in (1) is presented in Schruben and Cogliano, 1985. That methodology was developed for continuous input variables only. Extensions to discrete input variables will be discussed later in this paper. We next summarize this procedure.

1.1. Meta-Model Identification Procedure:

Briefly, frequency domain meta-model identification experiments are run as follows.

(1) Select a range of interest,

$$\{(x_1, \ldots, x_p) | L_i \le x_i \le U_i\},$$

for each continuous input factor. Here L_i is a lower limit and U_i is an upper limit on the values of factor x_i. The larger the experimental region, the more power there will be in detecting input factor effects. However, just as in dealing with electrical systems one must take care that the simulation does not "blow a fuse". That is, the values taken on by the input variables do not put the simulation in an unstable region for too much of the time.

(2) Select p+1 driving frequencies, ω_i, i = 1,2,...,p, for each continuous input variable in the simulation. These frequencies are between 0 and 1/2 cycles per output observation. Partial confounding of indicator frequencies may occur if driving frequencies are selected carelessly. Independent spectrum estimates for different indicator frequencies are desirable, so indicator frequencies should be as widely separated as possible in the interval (0,1/2]. The frequency selection problem depends on the number of input factors and the list of terms in the prospective response model. The problem of selecting driving frequencies can be formulated as a mixed integer linear program [Cogliano, 1982]. However, frequency selection is not critical. Term indicator frequencies and their aliases should be at least one bandwidth apart. The selection of bandwidth is under the control of the experimenter. If bandwidth is decreased the same estimator precision can be obtained by increasing the run length. Determine the indicator frequencies for each term in the prospective polynomial response model. These frequencies are in the sets given by,

$$S_i = \{a_{ij}\omega_i, (a_{ij} - 2)\omega_i, \ldots, -a_{ij}\omega_i\}, \quad i = 1, \ldots, q.$$

for each term involving the a_{ij} th power of the single variable x_i in the prospective response model given by (1). For terms involving the interactions of powers of more than one input variable, the indicator frequencies are given as the direct sum of the sets, S for each term in the interaction. Compute the minimum spacing, b, between these frequencies evaluated at their principle aliases. The value of b will be the band width necessary for the spectrum estimators. Some effort should be made to select driving frequencies that maximize this band width. The problem of Maximizing the minimum spacing between term indicator frequencies is rather difficult. A method that involves nesting two univariable integer optimization programs has been developed [Jacobson and Schruben, 1986]. This method, while heuristic, has resulted in the optimal selection of driving frequencies each time it has been tried and the optimal solution is known. The problem of selecting driving frequencies for frequency domain simulation experiments is a current topic of research underway.

(3) Choose a window size of m and a run length of n observations such that $m \geq 4/3b$ if using the Tukey spectral window and $m \geq 1.86/b$ if using the Parzen sepctral window in step (5) below. Also it is recommended in [Chatfield, 1975] that $20m \geq n \geq 3m$. See the book by Chatfield for an introductory discussion of window size selection (note that Chatfield revised his recommendation as to window truncation point in the second edition of his book...the selection of the truncation point m is an art form; try several values). I further recommend that n include at least 10 full cycles of the lowest term indicator frequency.

(4) Run p+1 independently seeded replications of the simulation program using a Latin square design to assign input factors to driving frequencies for each run as described in Schruben and Cogliano. For each run the input factors oscillate according to

$$x_i(t) = \tfrac{1}{2}(U_i + L_i) + \tfrac{1}{2}(U_i - L_i)\cos 2\pi\omega_i t.$$

(5) Compute the sample spectrum

$$\hat{f}(\omega) = \sum_{k=-m}^{m} \lambda_k c_k \cos(2\pi\omega k) \qquad 0 \leq |\omega| \leq \tfrac{1}{2}$$

for each response series. The λ_k are the weightings for a particular spectral window. The Parzen and the Tukey windows have both been successfully used for frequency domain experiments with very little difference in the results.

(6) Compute the spectrum ratios

$$F(\omega_i) = \frac{\frac{1}{\nu r}\sum_{j=1}^{r} \hat{f}_j(\omega_i)}{\frac{1}{\nu s}\sum_{j=1}^{s} \hat{f}_j^c(\omega_i)}$$

The spectral estimators with the superscript, c, are for frequencies that are not changed during a particular run. The c is for "control" and this spectrum estimator measures the ambient noise in the response at a particular frequency. The spectral ratios above will have an approximate F distribution if the term, t_j, has a zero coefficient in the meta-model. This can be used to compute the observed significance, denoted as $p(\omega_i)$, of each term in the meta-model. This significance level can be computed for several such "signal to noise" ratios in the Latin Square experiment. Finally, these observed significance levels are combined into an overall significance level for each term in the prospective meta-model. Fisher's method of computing an overall significance level is used by the author but other methods might be considered (see Rosenthal, 1978).

$$p^* = -2\sum_{i=1}^{k} \log p(\omega_i),$$

This overall significance level, p^* is used to evaluate each term for inclusion in the prospective simulation response surface model.

1.2 Frequency Domain Experiments with Discrete Variables:

Three techniques for dealing with discrete variables in frequency domain experiments. The most straightforward technique would be to use the discrete driving functions given by the Walsh basis to vary the values of these parameters during each run [Sanchez and Schruben, 1985]. This approach has the problem that the Walsh spectrum is not invariant to time lags. Since most simulation responses to changes in input variables are delayed this makes the direct application of Walsh analysis of little use in simulation experiments. A second technique would be to use a conventional run-oriented design for the discrete factors and embed a frequency domain experiement involving the continuous factors in each run of such a design. This approach would not take advantage of the main benefit of frequency domain simulation experiments; frequency bands rather than simulation runs are the experimental units.

A general method for incorporating discrete valued variables in a frequency domain experiment that takes full advantage of the frequency domain approach is to simply map a continuous variable into the discrete valued variable. Then the continuous variable is considered as the input to the simulation; that is, this new variable is oscillated with a sinusoidal driving frequency. There are several examples of techniques of making this continuous variable to discrete variable mapping. One example of such a mapping is to pick the values for a discrete variable by inverting a uniform discrete cumulative distribution function. The argument of the inverse uniform distribution function would be a continuous variable between zero and one that takes on values according to a sinosoidal function during the simulation. This is just a deterministic version of the common technique of generating the value of a random varible by taking the inverse of a distribution function with a uniform pseudorandom number as its argument.

A stochastic mapping that could be used for converting a discrete input variable to a continuous variable would be to consider the value of each discrete variable as determined by an indicator function. The probabilities for the indicator function are then oscillated according to sinusoidal driving functions. The actual value of the discrete variable used in the simulation program is randomly drawn from a distribution with oscillating probabilities. These continuous probabilities are now the inputs to the simulation. This second technique of selecting the values of discrete variables randomly according to oscillating probabilities has been found to work quite well is the limited testing so far.

Regardless of the method for varying discrete valued variables during a simulaiton run there remains a serious issue that the experimenter must address. How are discrete variables to be changed during a run if their values alter the logic of the simulation? For example, if a queueing simulation is being run and the number of servers is being oscillated, what happens to customers of a server that is removed? The best approach to this problem found so far is to let the system change "naturally" when values of such variables are changed. In the queueing example just mentioned the logic of a server breakdown (or coffee break) is followed when a server is removed. This may mean that the current customer (if any) is served or this may mean that the current customer may have service terminated. The server may also simply hang out a "this line closed" sign depending on the natural server break behavior in the real system being simulated.

2. USING FREQUENCY DOPMAIN EXPERIMENTS IN RESPONSE OPTIMIZATION:

There are several ways that frequency domain experiments can be useful in simulation optimization. Fore example: in classical Response Surface Methodology frequency domain experiments can be used to assess lack of fit by a linear model and indicate when phase II of RSM (quadratic model fitting) should be started.

Here we discuss a more direct use of frequency domain experiments in simulation optimization. A criterion for local optimality in the frequency domain is first presented and illustrated. This is followed by some approaches to optimization in the frequency domain.

2.1 A Criterion for Optimality:

When a frequency domain experiment is centered in a small region about a local optimum, the spectral amplification at indicator frequencies of all terms involving the unit power of the input variables disappear. Intuitively, this is because the first partial derivatives of the response function are zero. These partial derivatives for unit powers of variables are directly proportional to the coefficients of linear terms in the response surface polynomial meta-model.

To illustrate: consider the response surface given by the function given by

$$E[Y] = -2x_1^2 + 5x_2^2 - 4x_1x_2 + 96x_1 + 48x_2 - 960.$$

This function was used in [Smith, 1976] as a test function for an automated RSM procedure for optimizing a simulation response. In Figure 1 we have the spectrum of the output of a single run at the initial (non-optimal) point of the RSM procedure in Smith's experiment. The frequency assigned to the factor x_1 was .07 cycles per observation and the frequency assigned to factor x_2 was .28 cycles per observation. Note that in Figure 1 spikes at frequencies corresponding to the linear terms dominate the spectrum.

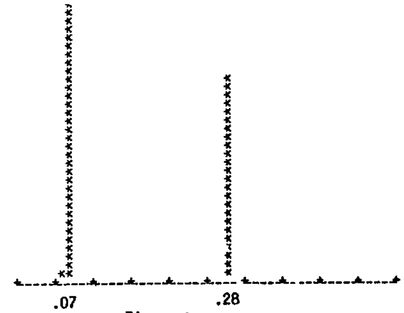

Figure 1

Now consider the same experiment but with the oscillations centered at the optimal values of x_1 and x_2. The specturm of a single frequency domain run centered at the optimum is given in Figure 2. Note the shift in the spectrum. All linear terms have no spectral amplification at the optimum as expected.

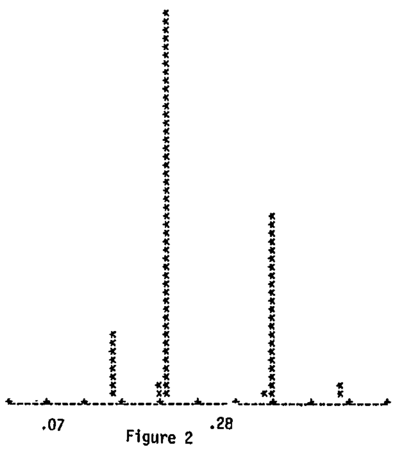

Figure 2

2.2 Optimization Techniques:

Research is still underway on optimization techniques using frequency domain methods at the time of the writing of this paper. Adaptations of classical non-linear algorithms to the frequency domain as well as completely new approachs are under study. The latest results in this area will be reported at the conference.

3. SUMMARY AND SUGGESTIONS FOR RESEARCH:

In this paper we have reviewed the application of frequency domain methods to simulation experiments. We have also presented a new approsch for including discrete valued factors in such experiments as well as a new frequency domain criterion for local optimality of a simulation response surface. Using frequency domain methods in simulation experiments is a relatively new idea that holds much promise. The adaptation of these experiments in simulation optimization is a particularly interesting area of research.

The author wishes to acknowledge the help of Sheldon Jacobson who ran the experiments presented in this paper.

REFERENCES

Chatfield, C. (1984). The Analysis of Time Series: An Introduction (3rd.Ed.). Chapman and Hall.

Cogliano, V.J. (1982). "Sensitivity Analysis and Model Identification in Simulation Studies", Ph.D. dissertation, School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY.

Jacobson, S., and Lee Schruben, (1986) "An Algorithm for Selecting Driving Frequencies for Frequency Domain Simulation Experiments", Technical Report, S.O.R.I.E., Cornell University, Ithaca, NY.

Kleijnen, J.P.C., A.J. van den BURG, and R.Th. van der Ham (1979). "Generalization of Simulation Results: Practicality of Statistical Methods", European Journal of Operations Research Vol. 3, No. 1. pp. 50-64.

Kleijnen, J.P.C. (1980). "Experimentation with Models: Statistical Design and Analysis Techniques", proceedings for "Simulation '80", Interlaken, Switzerland, June 1980.

Rosenthal, R. (1978). "Combining Results of Independent Studies", Psychological Bulletin, Vol. 85, pp. 185-193.

Sanchez, P.J., and L. W. Schruben (1985) "Significant Factor Identification using Discrete Spectral Methods", Tech. Rpt. 654, SORIE, Cornell University, NY.

Schruben, L.W. and V.J. Cogliano (1981). "Simulation Sensitivity Analysis: a Frequency Domain Approach", Proceedings of the 1981 Winter Simulation Conference, Atlanta, Georgia.

Schruben, L. and V. J. Cogliano, (1985). "An Experimental Procedure for Simulation Response Surface Model Identification", Tech. Rpt. 669, SORIE, Cornell University, Ithaca, N.Y. 14853.

Smith, D. E., (1976), "Optimization of a Computer Simulation Response", Technical Report 106-3, Desmatics Inc. State College, PA. 16801.

AUTHOR'S BIOGRAPHY

LEE SCHRUBEN is on the faculty of the School of Operations Research and Industrial Engineering at Cornell University. He received his undergraduate degree in engineering from Cornell University and a Masters degree from the University of North Carolina. His Ph.D. is from Yale University. Prior to coming to Cornell he was the Associate Director of the Health Systems Research Division and an Assistant Professor of Pharmacology in the Medical School at the University of Florida. Before going to graduate school he was a manufacturing systems engineer with the Emerson Electric Co. in St. Louis, Mo. His research interests are in the statistical design and analysis of large scale simulation experiments. His consulting activities have been primarily in the area of manufacturing systems simulation. He is currently the chairman of the TIMS College on Simulation and is also a member of ASA, ORSA, and SCS where he serves on several editorial boards.

Prof. Lee Schruben
School of O.R.&I.E.
Cornell University
Ithaca, NY 14853
ph: (607) 255-9133
mail: FX9J@CORNELLA