Proceedings of the 1989 Winter Simulation Conference
E.A. MacNair, K.J. Musselman, P. Heidelberger (eds.)

SIMULATION SENSITIVITY ANALYSIS USING FREQUENCY DOMAIN EXPERIMENTS

Douglas J. Morrice
Lee W. Schruben
School of O.R.I.E.
Cornell University
Ithaca, NY 14853

ABSTRACT

In this paper we review frequency domain experiments (FDE) for simulation sensitivity analysis and contrast it with conventional methodology. With very few runs of the model (often one or two runs is sufficient), FDE can be used to gain information on which factors significantly influence the performance of the simulated system. Two examples are given which illustrate some of the fundamental differences between frequency domain and conventional run-oriented analysis.

1. SIMULATION SENSITIVITY ANALYSIS

To run a simulation program the experimenter typically must choose values for a large number of input variables. Broadly speaking, the goal of a sensitivity analysis experiment is to determine how the values of these input variables influence the simulation output. The objective of a particular sensitivity analysis experiment can range from merely classifying the input variables to estimating system response gradients.

1.1 Classification of Input Variables

Input variables can be classified as discrete (including qualitative variables) or as continuous. Continuous variables are further classified by whether or not the output (response) is differentiable, continuous, or neither with respect to these variables.

In traditional experimental design literature, input variables are further classified as being factors or parameters of the system. Param. are considered not to be under the direct contro. of the experimenter whereas factor values are considered to be part of the system design. In a simulated system all input variables can be controlled by the experimenter. Nevertheless, it is still useful to think of these variables as being either parameters or factors depending on whether they can be easily controlled in the real system being simulated. An example of a system parameter would be the demand rate for a service or product. An example of a system factor would be the number of machines to place in a factory cell or the number of servers in a service facility.

The estimation of parameters usually involves collection and analysis of real-world data which can be a very expensive part of a simulation study. The determination of factor settings for a good system design requires extensive experimentation with the simulation and can also be expensive.

1.2 Factor Screening

The qualitative classification of the input variables as being either important or unimportant is commonly called screening. The objective of screening is to simplify a model by identifying which factors should be varied in an experimental design or which input parameters need to be estimated accurately from real data and which do not. The most popular screening approach is to run a traditional screening design (e.g. group screening or fractional factorial designs) and apply conventional regression (including ANOVA) analysis.

Efficient screening of factors can be beneficial as a preliminary step in system design so that optimization can be done over a smaller dimension factor space. A qualitative screening of unimportant parameters can also help in controlling data collection costs.

1.3 Metamodeling

An intermediate experimental objective might be to fit a regression model for the simulation output variables as functions of the values of input variables. Such experiments are sometimes referred to as metamodeling experiments. Both conventional screening methods and metamodeling experiments are discussed extensively in [Kleijnen, 1987]. Using ANOVA-like techniques qualitative factors can be included in such experiments.

1.4 Response Optimization

A more quantitative sensitivity experiment involves response gradient estimation. Available methodologies for such experiments are necessarily restricted to scalar responses and continuously differentiable quantitative input variables. The objective of such experiments is typically in system design where one seeks optimal values of input factors. Stochastic optimization with respect to more than a very few input factors is known to be a very difficult task. If in addition some of the factors are discrete or qualitative then optimization without some preliminary factor screening is virtually hopeless.

Response Surface Methodology (RSM) is a traditional approach to response optimization [Myers, 1971]. Unfortunately, RSM requires a large number of simulation runs and is supported by very restrictive assumptions on system behavior.

A developing technique for simulation optimization experiments is Perturbation Analysis [Ho, et.al, 1979]. This method is known to give efficient gradient estimates for certain simple systems and empirically appears to be useful even in cases where the resulting gradient estimates are known to be biased. There is also a potential for perturbation analysis of discrete factors but this methodology needs more development before it can be strongly endorsed. When possible an appropriate gradient estimation algorithm such as infinitesimal perturbation analysis should be used.

1.5 The Domain of FDE

The appropriate domain for frequency domain simulation experiments as presently developed is in screening. Most input variables (whether they be continuous or discrete, quantitative or qualitative, factors or parameters) can be included in an FDE. Where the cost of a simulation run is considerably more expensive than the marginal cost of a single observation of an output variable during a run, then FDE can be recommended as a first screening step in developing an overall experimental design for a simulation study.

Unlike some sensitivity methodologies, FDE methods are not invasive; the simulation code need not be altered to run these experiments. However, the model is run very differently and the output analyzed differently.

2. FREQUENCY DOMAIN EXPERIMENTS

2.1 Background

We present here only a brief summary of FDE. Further details on implementation of the methodology can be found in [Schruben and Cogliano, 1987], and a discussion of theoretical issues is given in [Schruben, Heath, and Buss, 1988]. A detailed application of the method can be found in [Sanchez and Schruben, 1986].

In a traditional simulation experiment the values of the input variables are fixed for each run of the model. The experimenter is willing to change the values of these variables only when making a different simulation run. In an FDE, input variables are oscillated sinusoidally at different frequencies (called *driving frequencies*) throughout the same simulation run; hence the moniker "frequency domain experiment".

While persons might be quite comfortable with running a system in the traditional manner, they might find changing input settings during a run quite disconcerting (see [Kleijnen, 1987], page 242). Of course, the very concept of "running a system", while familiar, is nonetheless abstract and the definition of what constitutes a run often arbitrary.

The philosophical issue here is not the willingness to turn a system control knob while experimenting with a system. Both traditional experiments and FDE involve changing system variables. The question is whether one must hypothetically "stop and restart" the system to make such a change. In conducting an FDE the simulation code is not changed, we simply do not stop a "run" to change the value of an input variable.

There is no question that a system will behave differently when its control knobs are twisted while the system is running; the practical issue is what information can be gained in this manner. Fixed variable, simultaneous experiments have been inherited from fields such as medicine and agriculture where there are no practical alternatives; time compression is not an option. Where varying an input during an experiment has been possible, it has found ready application. The fact is that in many real-world systems FDE are more common than fixed variable experiments simply because they provide more information (examples of FDE abound in electrical engineering, seismology, nuclear engineering, etc).

After running a frequency domain experiment, the series of output observations is regressed against sinusoids at the various input driving frequencies. Such a regression is called harmonic analysis. Amplification of a particular driving frequency by the system is an indication that the response is sensitive to the corresponding input variable. If a driving frequency is not found in the output process then the response is presumed to be insensitive to the corresponding input factor.

While not critical in a screening experiment, different regression terms (products of integer powers of the input variables) will result in different sets of frequencies being present in the output. The presence of a particular term in the response regression is indicated in the output by the presence of an oscillation at a corresponding set of *term indicator frequencies*. For example, if the response is sensitive to the product of two input variables (called an interaction term in conventional regression modeling) then there should be frequency amplification in output at the sum and difference of the two factor driving frequencies. Consult the reference by [Jacobson, et. al., 1987] on selecting driving frequencies for FDE for further information on term indicator frequencies.

In summary: for FDE a simulation code is not changed, it is "run" differently from a traditional experiment. Like traditional analysis, regression is used for FDE. However, a different basis is used; the output series is viewed in a different coordinate system. The key idea is that rather than runs of the simulation being the experimental unit, frequency bands or "tones" in the output series are the experimental unit. There are many usable frequency bands in the output of a single run which provides FDE with a potential advantage.

There have been some recent changes in the methodology not mentioned in the earlier references. One change is that least squares regression of the output series on sinusoid independent variables is used; that is, harmonic analysis is used rather than spectral estimation. In addition, the sinusoids are added to input random variables rather than to input parameters. These two changes, while making theoretical treatment of the method more straightforward, have a fairly minor impact on the practical applications.

2.2. Implementation Issues

2.2.1. Noise

There is a fundamental difference between how random noise is treated in frequency domain experiments and in conventional RSM or metamodeling regression. In the conventional methodologies randomness from all sources is "pooled" into the sum of squares for noise. This is used to compute the denominator for an F test for term significance.

In FDE, randomness is not pooled but kept separate for each potential term in the regression. We use periodogram estimators at Fourier frequencies that are not term indicator frequencies for noise estimation. In this way, each term has its own numerator *and* denominator for tests of significance. These may be term indicator frequences from a second independent run of the simulation where all or some of the inputs are held fixed (called a noise run) or they may be periodogram components at neighboring frequencies (within a "band-width" of an indicator frequency) from the same run.

Another effective method for reducing the impact of random noise in the output series is to make the amplitudes of the driving oscillations as large as possible. See [Jacobson, 1989] for a discussion of amplitude selection issues.

2.2.2. Serially dependent output: Gain

Serial dependence (system memory) causes some of the most difficult problems with conventional time-domain simulation analysis (e.g. initialization bias, interval estimator coverage failure). It is not too surprising that serial dependency also causes problems in the frequency domain. The frequency domain manifestation of system memory is called gain and it is perhaps the most troublesome problem in FDE. Gain is the phenomenon where some input frequencies will be naturally filtered out by the simulation model whereas others might be naturally amplified.

In typical simulations of queueing systems there is a tendency for positive serial correlation between observations in the output series. In the frequency domain this positive memory makes the system in the frequency domain behave as a low-pass filter. Low frequencies, if present in the input, are amplified by the system; high frequencies are filtered out. The impact of system gain on FDE has several interesting aspects; however, from a practical viewpoint it means that driving frequencies must necessarily be chosen at the low end of the frequency spectrum. The use of low driving frequencies is the only effective means developed so far for dealing with gain in FDE; Slowly oscillating inputs tend to work well where fast oscillations do not. The price of using only low driving frequencies is that the usable frequency bands are much narrower resulting in an increased run length. (See [Schruben and Cogliano, 1987] for a limited discussion of this issue.) This is satisfactory when the cost of increasing the run length is not significant compared with other costs of a simulation study.

The tables in [Jacobson, Buss and Schruben, 1987] can be used to choose driving frequencies across the whole visible spectrum (0 to .5 cycles per observation). A slight modification of the algorithm in this reference can be used to select driving frequencies across just a lower portion of the visible spectrum.

Filtering techniques which have been stunningly successful in other frequency domain applications (i.e. Kalman filtering in image processing) remain to be tried in FDE. The application of filtering, a frequency domain counterpart to variance reduction techniques via dependency induction, is a current topic of research by the authors.

2.2.3. Aliasing

In discretely sampled output there is also the problem of frequencies higher than .5 cycle per observation being aliased (folded) onto lower frequencies. These aliased frequencies might be confused with lower frequency term indicator frequencies for other variables. The potential of aliasing (not really a serious problem in this application since potential aliases are known) is reduced if low frequencies are used.

2.2.4. Indexing

When running a frequency domain experiment the input oscillations and the output observations must have a common index. It is important that each output observation have an index that is common to all oscillations that might contribute to the value of the observation.

At first glance, simulation clock time is an obvious choice for an index; however, this index does not work as one might expect (see [Jacobson, Morrice, and Schruben, 1988]). In open queueing simulations the use of the customer arrival sequence number as an index is quite effective. The oscillations are then made in units of cycles per customer. If the customers are shuffled (say by feedback) the output is sorted before harmonic analysis. In closed queueing networks, assigning a new identification number to a customer each time it passes a data collection point also works; this requires a little care [Sargent and Som, 1988]. In a factory simulation where part assembly may occur then the output product number (order number) is used as the common index for all parts on the bill of material for that product [Yucesan, 1988]. In all these cases, frequencies are in units of cycles per output observation (see [Sanchez and Schruben, 1986] for a detailed illustration).

2.2.5. Qualitative or Discrete Factors

In the above referenced example by Sanchez and Schruben there were discrete or qualitative input factors such as the probability distributions used or the number of discrete resources in the system. In FDE such factors can be included in the same experiment as continuous quantitative factors. Including all types of input in the same FDE run is one of the strengths of the approach. In [Schruben, 1986] a technique for including discrete factors in a simulation is presented. The trick is to randomize over discrete sets using an oscillating randomization scheme (e.g. a coin flip with an oscillating probability of coming up heads). This works largely due to FDE not pooling system noise. [Sanchez and Sanchez, 1989] present two enrichments of this randomization technique for multiple valued factors.

3. CONTRASTING FDE WITH CONVENTIONAL REGRESSION

The supporting assumptions for least squares regression and ANOVA are basically that the samples are independent and identically distributed (iid) Normal. Sufficient conditions to insure that FDE will work include that the system can be modeled as a time-invariant linear filter that transforms the input to the output. A Hammerstein-like model of the simulation has been developed that has many of the characteristics of discrete event simulations [Sanchez and Buss, 1987]. Empirical evidence suggests that this model is a good approximation to the behavior of many discrete event simulations for low driving frequencies (We mention again that the lower the input frequency the better the performance of FDE but the longer the required run length...no free lunch here!). A more complete analysis of the relative power of conventional least squares and FDE is presented in section 4 of this paper. Here we examine two empirical studies from [Sargent and Som, 1988].

3.1. An M/M/1 Queue

Frequency domain experiments with an M/M/1 queueing system were studied extensively by [Sargent and Som, 1988]. The response of customer waiting time, W, to changes in the mean service time, S, was examined. They attempted to contrast ordinary least squares (OLS) with FDE; however, they failed to apply the two different methodologies to the same system. In their study, the authors applied ordinary least squares (OLS) to the deterministic equation of the mean system response and applied FDE to the simulated stochastic system. Thus their experiments do not provide any basis for comparison.

When we applied standard OLS and FDE to the same simulation, the factor screening results were essentially the same. The regression analysis is summarized in Table 1. It is based on 20 runs of 6000 data points where 1000 data points where truncated on each run in order to reduce the initial transient effect. The R^2 and R^2-adjusted statistics are used to compare the fit of each regression model. $X = S - 7$ represents the shifted mean service time, i.e., $S \in [5,9]$ is shifted to $X \in [-2,2]$. Based on these statistics, the linear and quadratic terms are important, the cubic term is questionable, and the fourth and fifth order terms add nothing to the model.

The FDE results are summarized in Table 2. They are based on two runs (a signal run and a control run - see [Schruben and Cogliano, 1987]) of length 65536 each (the run length of 2^{16} facilitates the use of the Fast Fourier Transform) with no truncation for initial transient effect. During the signal run, the input factor, X, is varied according to

$$X(t) = 2.0 \cos (2 \pi (t / 1024))$$

The signal-to-noise ratio is constructed from the ratio of two periodogram ordinates (the numerator from the signal run and the denominator from the control or noise run). Under certain regularity assumptions and the assumption that the terms in the model have no effect on the waiting time, this ratio has an asymptotic F distribution with (2,2) degrees of freedom (see [Brillinger, 1981], Chapter 5). Only the signal-to-noise ratio values associated with the linear and quadratic terms are significant when compared with the 95th percentile of an $F_{2,2}$ distribution, which is equal to 19.00.

Table 1: Waiting Time Regression Results for M/M/1 Queue Using OLS.

Metamodel	R^2	R^2- adj
$W = 31.2 + 15.8\ X$	0.835	0.826
$W = 22.6 + 5.86\ X^2$	0.134	0.088
$W = 22.6 + 15.8\ X + 5.86\ X^2$	0.968	0.965
$W = 22.6 + 11.2\ X + 5.86\ X^2 + 1.75\ X^3$	0.982	0.978
$W = 23.1 + 11.2\ X + 4.77\ X^2 + 1.75\ X^3 + 0.290\ X^4$	0.982	0.978
$W = 23.1 + 12.4\ X + 4.77\ X^2 + 0.43\ X^3 + 0.290\ X^4 + 0.273\ X^5$	0.982	0.977

Table 2: FDE Results for Waiting Time in M/M/1 Queue

Term	Signal-to-Noise Ratio
X	1712.7
X^2	96.733
X^3	10.443
X^4	2.0103
X^5	0.83664

Here FDE and least squares behaved the same for the identification of significant factors; with our model implementation the FDE was much less expensive. Least squares does allow estimation of model parameters whereas FDE as applied here only permitted factor screening; so something was gained by the extra expense of the least squares experiment.

3.2. A tandem queue

A more interesting example, also from the Sargent and Som reference cited above, is a tandem queueing system consisting of an M/M/1 queue feeding into an M/D/∞ queue. Again Sargent and Som do not apply OLS and FDE to the same system (OLS calculations were made with a deterministic function and FDE experiments were run with a stochasitic simulation). Quite expectedly, they get different results. However, this time when we ran the experiments using the same simulation we still got different results! In this example, FDE indicates that the total customer waiting time is strongly influenced by the deterministic service time; OLS analysis concludes that this service time is not an important factor.

Sargent and Som incorrectly conclude that the least squares results are right while the FDE results are wrong. Clearly, the service time of the second M/D/∞ system has a strong influence on customer waiting time. Any change in this service time changes not only the average customer waiting time but the entire sample path of all customer waiting times by exactly the same amount; this factor is a perfect control! The failure of OLS to detect the importance of the second service time is due to its pooling of mean-square-error from all sources as mentioned above. Random noise used to compute the denominator in the OLS F-statistics comes primarily from the first queue. This randomness from the first queue is large enough to mask the response to the second queue. Changes in the service time for the second queue do not "explain" very much of this pooled variation.

When the experimental range for the service rates for the first M/M/1 queue is narrowed, the R^2 value for the M/D/∞ queue approaches 1. In fact, the OLS R^2 assigned to the service time for the second queue can be changed almost arbitrarily by changing the range of values used for the service rate in the first queue. FDE does not pool random noise so the influence of the M/D/∞ queue is clearly seen in the analysis. When regression is correctly used, conclusions concerning service rates for the first M/M/1 queue for OLS and FDE are again essentially the same as the second queue contributes no noise to the system.

4.0. POWER COMPARISON: FACTORIAL vs FDE

Consider a conventional screening experiment using a fractional factorial design where we are only interested in testing for main effects of K input factors. An appropriate model for the two level factorial experiment is,

$$Y(s) = \sum_{i=1}^{K} A_i X_i(s) + \varepsilon(s), \quad s = 1,..,M, \tag{1}$$

where

Y(s) is the output response for run s,

K is the number of input factors,

A_i is the coefficient for the ith input factor,

$X_i(s)$ is the level of input factor i on run s,

$\{\varepsilon(s)\}_{s=1}^{M}$ are iid $N(0,\sigma^2)$ random variables,

and M is the total number of runs.

The model for FDE is,

$$Y(t) = \sum_{i=1}^{K} A_i \cos(\omega_i t + \varphi_i) + \varepsilon(t), \quad t = 1,..,N$$

where

K and A_i are defined as in (1),

Y(t) is the output response for observation t,

ω_i is the frequency (in radians per unit of t) of input factor i,

$\{\varphi_i\}_{i=1}^{K}$ are iid $U(-\pi,\pi)$,

$\{\varepsilon(t)\}_{t=1}^{N}$ are iid $N(0,\sigma^2)$ random variables,

and N is the total number of observations.

Suppose that the cost of producing one additional observation within a simulation run is one unit and the set-up cost for each simulation run, in the same unit of measure, is C (C includes costs of program loading and any necessary compilation or truncation for initialization bias, etc.). Then the total cost, C_1, associated with the two level factorial experiment is,

$$C_1 = M (C + 1) \tag{2}$$

because each one of the M runs of the simulation has C units of set-up cost and generates one observation. The total cost, C_2, for the FDE is,

$$C_2 = C + N \tag{3}$$

since the set-up cost is incurred once and all N observations are generated from one run of the simulation. In this analysis, we make only one run for the FDE test because term indicator frequencies are used to estimate the main effects. Their neighboring frequencies, which are not used as term indicator frequencies, are used to estimate the error.

Equating C_1 and C_2 in (2) and (3) yields,

$$N = M (C + 1) - C \tag{4}$$

where (4) provides the number of observations for a FDE that equates its total cost with that of the factorial experiment.

Under iid normal assumptions, it is easy to construct F-tests to test for the significance of the input factor in either model (for factorial experiments see [Box and Draper, 1987], Chapters 4 and 5, and for harmonic analysis see [Hartley, 1949] or [Anderson, 1971], Chapter 4) and compute their power. Figure 1 is a graph of power versus C for both tests (which have level of significance 0.01). Here M is fixed at 32 and K equals 25 (a 2^{25-20} fractionated design of resolution III is used).

When C = 0, the two level factorial test uniformly dominates the frequency domain test in terms of power. However, as C increases the factorial experiment becomes more costly relative to the FDE. For the same total cost (see the relationship in (4)), many more observations (and degrees of freedom) can be generated for the FDE, thus increasing the power of the frequency domain test. Therefore, in this example and under the stated assumptions, the frequency domain method gives superior power when the set-up cost of each run is higher than the cost of four or five observations. For example, if we were simulating a queueing system and observing customer waiting times then FDE have better power if the set-up cost

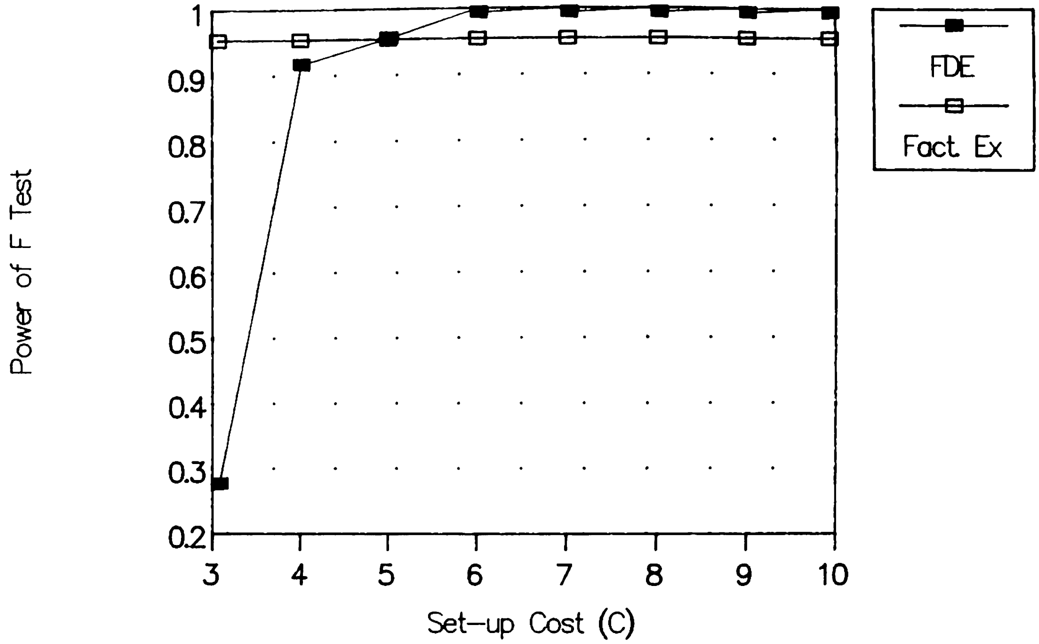

Figure 1: Power versus Cost Curves for Fractional Factorial and Frequency Domain Experiments.

per run is greater then the cost of simulating five customers.

It is important to note that the assumption of independent error is more realistic for the factorial experiment than it is for FDE; observations from the same run are typically not independent. This assumption might bias the power comparison in favor of FDE. Positive serially correlated error will reduce the power of the FDE F-tests. On the other hand, it may also result in the need for more initial data truncation and increase run set-up costs (which favors the use of FDE).

Therefore, the analysis in this section should be viewed as providing guidelines for <u>not using</u> FDE (i.e., when run set-up is nearly free). The power of FDE tests with correlated error is a topic of current study.

5. CLOSING COMMENTS

When the marginal cost of each observation is small compared to the cost of a simulation run, the running of a preliminary FDE can be an effective and efficient method of initial factor screening. Isolation of important qualitative and quantitative input variables makes conventional experiments, data collection, parameter estimation, and system optimization significantly easier. Screening can also be used to detect programming errors in an otherwise working simulation code. As with any factor screening methodology, the user is cautioned against drawing quantitative conclusions with this methodology.

ACKNOWLEDGEMENTS

This research was partially supported by a Cornell National Supercomputer Facility Grant. The authors wish to thank Professor Robert Sargent for many interesting and challenging discussions of FDE.

REFERENCES:

Anderson, T.W. (1971). *The Statistical Analysis of Time Series*. John Wiley & Sons, New York.

Box, G.E.P. and Draper, N.R. (1987). *Empirical Model-Building and Response Surfaces*. John Wiley and Sons, New York.

Brillinger, D.R. (1981). *Time Series: Data Analysis and Theory*, Second Edition. Holden-Day, Inc., San Francisco.

Hartley, H.O. (1949). Tests of Significance in Harmonic Analysis. *Biometrica*, Volume 36, 194-201.

Ho, Y.C., Eyler M.A., and Chien, T.T. (1979). A Gradient Technique for General Buffer Storage Design in a Production Line. *International Journal of Production Research*, Volume 17, Number 3, 557-580.

Jacobson, S.H. (1989). Oscillation Amplitude Considerations in Frequency Domain Experiments. In *Proceedings of the 1989 Winter Simulation Conference* (E.A. MacNair, ed.), Washington, D.C.

Jacobson, S.H., Buss, A.H., and Schruben, L.W. (1987). Driving Frequency Selection for Frequency Domain Simulation Experiments. (Submitted for Publication).

Jacobson, S.H., Morrice, D.J., and Schruben, L.W. (1988). The Global Simulation Clock as the Frequency Domain Experiments Index. In *Proceedings of the 1988 Winter Simulation Conference* (M.A. Abrams, ed.), San Diego, California, 558-563.

Kleijnen, J.P.C. (1987). *Statistical Tools for Simulation Practioners*. Marcel Dekker, New York.

Myers, R.H. (1971). *Response Surface Methodology*. Boston: Allyn and Bacon. (Reprinted by Edwards Bros., Ann Arbor, Michigan.)

Sanchez, P.J. and Buss, A.H. (1987). A Model for Frequency Domain Experiments. In *Proceedings of the 1987 Winter Simulation Conference* (A. Thesen, H. Grant, and W.D. Kelton, eds.), Atlanta, Georgia, 424-427.

Sanchez, P.J. and Sanchez, S.M. (1989). Spectral Design of Experiments for Discrete-Valued Factors. Working Paper #89-003, Department of Systems & Industrial Engineering, University of Arizona, Tucson, Arizona.

Sanchez, P.J. and Schruben, L.W. (1986). Simulation Factor Screening Using Frequency Domain Methods: An Illustrative Example. (Submitted for Publication).

Sargent, R.T. and Som, T.K. (1988). Personal Communication.

Sargent, R.T. and Som, T.K. (1988). A Different View of Frequency Domain Experiments. Working Paper #88-01, Simulation Research Group, Syracuse University, Syracuse, New York.

Schruben, L.W. (1986). Simulation Optimization Using Frequency Domain Methods. In *Proceedings of the 1986 Winter Simulation Conference* (J.R. Wilson, J.O. Henriksen, and S.D. Roberts, eds.), Washington, D.C., 366-369.

Schruben, L.W. and Cogliano, V.J. (1987). An Experimental Procedure for Simulation Response Surface Model Identification. *Communications of the Association for Computing Machinery*, Volume 30, Number 8, 716-730.

Schruben L.W., Heath D.C., and Buss, A.H. (1988). A Dynamic Response Surface Model for Frequency Domain Simulation Experiments. Technical Report 808, School of Operations Research and Industrial Engineering, Cornell University, Ithaca, New York.

Yucesan, E. (1988). Personal Communication.

AUTHORS' BIOGRAPHIES:

DOUGLAS J. MORRICE is a Ph.D. student in the School of Operations Research and Industrial Engineering at Cornell University. He holds a Bachelors degree in Operations Research from Carleton University in Ottawa, Canada and an M.S. degree in O.R. from Cornell University. His research interests include simulation output analysis, spectral methods, and manufacturing systems.

Douglas J. Morrice
S.O.R.I.E.
Cornell University
Ithaca, N.Y. 14853
(607) 255-1211.

LEE W. 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. His research interests are in statistical design and analysis of large scale simulation experiments. His consulting activities have been primarily focused in the area of manufacturing systems simulation.

Lee W. Schruben
S.O.R.I.E.
Cornell University
Ithaca, N.Y. 14853
(607) 255-9133.