1981 Winter Simulation Conference Proceedings
T.I. Oren, C.M. Delfosse, C.M. Shub (Eds.)

SIMULATION SENSITIVITY ANALYSIS: A FREQUENCY DOMAIN APPROACH

Lee W. Schruben
and
Vincent James Cogliano
Cornell University
School of Operations Research and Industrial Engineering
Ithaca, New York 14853

A procedure is presented for assessing the sensitivity of a discrete event digital simulation model to the values assumed for its input parameters. A frequency domain approach is taken where the input parameters oscillate throughout a run of the model. Parameter sensitivities are indicated by changes in the frequency spectrum of the simulation response. The spectrum can be used to identify a regression model for the simulated response surface. Several continuous parameters may be screened in a single run.

1. INTRODUCTION

In a simulation study one is often interested in how the simulation output, or response, depends on the input parameters. Parameters are fixed but unknown quantities affecting the system environment. For example, the average demand rate is a parameter of a simulated inventory system. The system's sensitivity to a particular parameter influences the amount of effort that should be spent in estimating the parameter and in experimenting with the simulation model for different values of the parameter. In this paper we present an inexpensive experimental method for characterizing the sensitivity of the response to the input parameters.

The experimental unit in a simulation study has traditionally been a single run of a computer program. For each run the parameter values are set according to an experimental design and remain fixed throughout the run. When the number of input parameters is large, classical experimental designs require a large number of runs. The result is that classical experimental designs can be prohibitively expensive for assessing the sensitivity of the simulation response to the parameters, particularly when one does not know whether some of the parameters are important at all. Hillier and Lieberman (1980) summarize this unfortunate situation and give an indication of the current state of practice in simulation sensitivity analysis:

> Therefore it is very expensive to conduct a sensitivity analysis of the parameter values assumed by the model. The only possible way would be to conduct a new series of simulation runs with different parameter values, which would tend to provide relatively little information at a relatively high cost.

Research for this paper has been supported by the Office of Naval Research Grant No. N00014-81-K-0037.

We propose a method that increases the amount of information that can be obtained from each simulation run by shifting the analysis into the frequency domain. In the frequency domain, the experimental unit is a frequency band within a run, and each run contains a large number of experimental units. During each run, each parameter is assigned to one of the frequencies and is varied according to a sinusoidal oscillation at its assigned frequency. This is possible in simulation studies because one can design a simulation model to give the experimenter complete control over the parameter values, including the ability to change those values during a simulation run. If the response is sensitive to a particular parameter, then the modulation of that parameter will tend to induce predictable oscillations in the response series. If the response is insensitive to a parameter, then the modulation of that parameter will not alter the response. Spectral analysis can be used to identify the induced oscillations in the response series and attribute their cause to particular parameters.

The simulation response surface acts to amplify parameter oscillations. Consider the simple situation pictured in Fig. 1. The response is a linear function of one parameter over the experimental region. If the parameter oscillates during the run at the frequency ω, the response will display an oscillation at the same frequency amplified by the slope of the response surface.

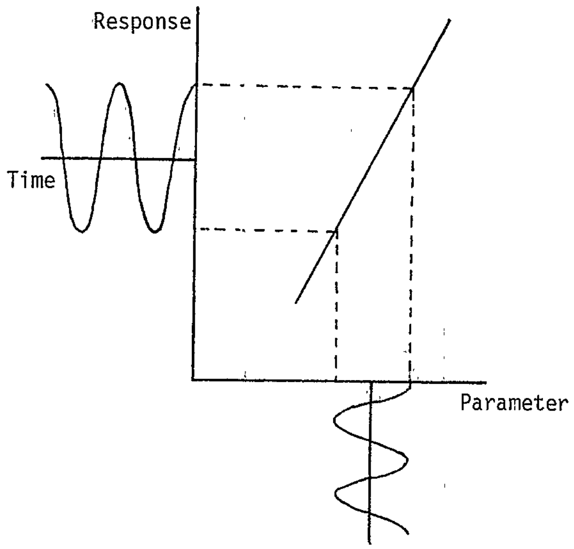

Figure 1
Parameter Oscillations Causing Response Oscillations

Driving a system with sinusoidal oscillations is common in areas such as electrical engineering and seismology. To the authors' knowledge, this technique has not been applied to the analysis of discrete event digital simulation models. Our development differs from that found in the electrical engineering literature in that we make use of elementary statistical experimental design concepts. We presume the reader has some familiarity with the spectral analysis of time series and the statistical design of experiments. Chatfield (1975) gives a sufficient background in spectral analysis in Chapters 6 and 7 of his book, and Kleijnen (1975) gives an overview of experimental design for simulation studies in Chapter 4.

2. TECHNICAL DEVELOPMENT

The objective of the simulation study is to identify a polynomial regression model which expresses the response as a function of the input parameters; that is,

$$E[y] = \sum_{j=1}^{g} a_j t_j$$

where y is the simulation response,

a_j is the regression coefficient corresponding to the term t_j,

t_j is a term in the regression model; it is a product of powers of the parameters such as $t_j = x_1$ or $t_j = x_1^2 x_3$, and

g is the number of terms in the regression model.

We propose to identify the proper terms to include in the regression model by driving the parameters with sinusoidal oscillations and analyzing the spectrum of the simulation response.

To estimate the response spectrum, the response must be observed at equally-spaced intervals. Let $y_1,\ldots,y_n$ be the response series. The spectrum is estimated by computing

$$\hat{f}(\omega) = \frac{1}{\pi} \sum_{k=-m}^{m} \lambda_k c_k \cos \omega k$$

using the sample auto-covariance function

$$c_k = \begin{cases} \frac{1}{n} \sum_{t=1}^{n-k} (y_t - \bar{y})(y_{t+k} - \bar{y}) & \text{if } k \geq 0 \\ c_{-k} & \text{if } k < 0 \end{cases}$$

and the spectral window

$$\lambda_k = \frac{1}{2}(1 + \cos \pi k/m)$$

$$m = \frac{1}{3} n.$$

The spectral window should be chosen to smooth the spectrum, balancing resolution against the precision of the estimates. We have used the Tukey window whose size is one-third of the number of observations for reasons that will be discussed later.

There are two major advantages in shifting the analysis into the frequency domain. The first is that independent frequency bands of the spectrum provide the vehicle for screening several parameters in the same run. The spectrum of the response decomposes the variability of the response into components at each particular frequency. A peak in the spectrum indicates a large contribution to the variability of the response at the corresponding frequency. For a linear system with a uniform gain (gain will be discussed later), driving the parameters with sinusoidal oscillations induces sinusoidal oscillations at the same frequency in the response. The response oscillations are amplified by the slope of the response surface for each parameter. Suppose the parameters $x_1,\ldots,x_p$ are driven by sinusoidal oscillations at the driving frequencies $\omega_1,\ldots,\omega_p$, respectively. The spectral estimates $\hat{f}(\omega_1),\ldots,\hat{f}(\omega_p)$ at the driving frequencies indicate the relative importance of these frequencies in the response, from which we can make inferences about the relative importance of the parameters oscillating at these frequencies. Because the frequency domain approach provides a large number of frequency bands in one run, several parameters may be screened in one run.

The second advantage in shifting the analysis into the frequency domain is that non-linear effects, such as powers and product interactions of the parameters, can be detected with no additional experimentation. When the parameters are driven by sinusoidal oscillations, it turns out that powers and products of the parameters are driven by compound oscillations which are the sums of sinusoidal oscillations. The frequencies present in a compound oscillation are called the indicator frequencies for the power or product term. Each term in the response model has a unique set of indicator frequencies. The presence of a particular term is indicated by increases in the spectrum at its indicator frequencies. Indicator frequencies can be determined by examining trigonometric identities. For example, suppose the parameter x_i is driven by a sinusoidal oscillation at the frequency ω_i,

$$x_{it} = \alpha_i \cos \omega_i t.$$

The quadratic term x_i^2 can be expressed

$$x_{it}^2 = \frac{1}{2} \alpha_i^2 (1 + \cos 2\omega_i t).$$

Hence the presence of the term x_i^2 is indicated by an increase in the spectrum at the frequency $2\omega_i$. For products suppose

$$x_{ht} = \alpha_h \cos \omega_h t \quad \text{and}$$

$$x_{it} = \alpha_i \cos \omega_i t.$$

Then

$$x_{ht} x_{it} = \frac{1}{2} \alpha_h \alpha_i (\cos(\omega_h + \omega_i)t + \cos(\omega_h - \omega_i)t),$$

so that the presence of the term $x_h x_i$ is indicated by increases in the spectrum at both $\omega_h + \omega_i$ and $\omega_h - \omega_i$. In general, the indicator frequencies for the power x_i^k are $k\omega_i$, $(k-2)\omega_i$, $(k-4)\omega_i$, etc., and the indicator frequencies for a product are all possible sums and differences of the driving frequencies of the factors in the product. By looking for increases in the spectrum at the indicator frequencies of powers and products, non-linear terms can be detected with no additional experimentation.

There are a few considerations to keep in mind when using the frequency domain approach to identify the important terms in a model of the response. The first is that the response spectrum can become quite crowded with indicator frequencies. For example, if screening four parameters for linear effects, quadratic effects, and two-way product interactions, there are twenty indicator frequencies to consider, $\{\omega_i, 2\omega_i, \omega_i \pm \omega_j$ for $i,j = 1,\ldots,4\}$. To ensure that effects at these frequencies can be distinguished from one another, a high-resolution spectrum should be used. The spectral window provides control over the resolution of the spectrum, with larger window sizes giving estimates with higher resolution and lower precision. Chatfield (1975) recommends that the window size be between one-third and one-twentieth of the number of observations. Because resolution is important in this approach, we recommend a window size from the higher end of this range. An added benefit of a high-resolution spectrum is that spectral estimates at different frequencies, especially non-neighboring frequencies, are more nearly independent.

Another consideration is the phenomenon of aliasing. When the spectrum is estimated from observations taken at equally-spaced intervals, the highest frequency that can be detected is one-half cycle per observation, known as the Nyquist frequency. Frequencies above the Nyquist frequency appear as if they were other frequencies below the Nyquist frequency. Aliasing has the effect of "folding" the entire range of frequencies onto the interval between zero and one-half; its effect is determined by the following rules:

(1) Frequencies $\omega \in [0, 1/2]$ are unchanged.

(2) Frequencies $\omega \in [1/2, 1]$ appear at $1-\omega$.

(3) Frequencies $\omega > 1$ appear at the fractional part of ω; that is, at $\omega - \lfloor \omega \rfloor$.

(4) Frequencies $\omega < 0$ appear at $|\omega|$.

All frequencies must be interpreted in terms of their aliases. For example, suppose $\omega_i = .4$ is the driving frequency of the parameter x_i. The indicator frequency of the quadratic term x_i^2 is $2\omega_i = .8$. Under rule (2) this frequency appears in the spectrum at the alias frequency .2. Hence a quadratic effect would be indicated by an increase in the spectrum at the frequency .2.

Another consideration is the system gain. The system gain prescribes how the system amplifies or suppresses input oscillations at each frequency. Systems sometimes have the effect of filtering out particular types of variability. A low-pass filter, such as computing an exponentially-smoothed average, suppresses high-frequency fluctuations so only low-frequency variability appears in the spectrum. A high-pass filter, such as computing a period-by-period change, suppresses low-frequency cycles so only high-frequency variability appears in the spectrum. In general, the system gain can greatly affect the spectrum, raising it at some frequencies and lowering it at others. If it lowers the spectrum at one of the indicator frequencies, it may cause one to miss identifying an important term in the response model. One way to compensate for system gain is to estimate it. This requires additional simulation runs and is rather complicated, considering there may be an interaction between parameter and frequency--the system may act as a low-pass filter for one parameter and a high-pass filter for another. Instead of estimating system gain we recommend treating gain as an unknown nuisance factor and blocking on it. A few simulation runs could be made changing the driving frequencies of each parameter between runs. For example, a second run could assign a high driving frequency to each parameter that had a low driving frequency in the first run and a low driving frequency to each parameter that had a high driving frequency in the first run. This way there is a smaller likelihood that an important term can be suppressed by a low system gain on every run. We have found blocking to be an effective and efficient technique for dealing with system gain.

In many systems the spectrum may show such a distinct increase at some of the indicator frequencies that a visual inspection is sufficient to identify the important terms in the response model. Rigorous statistical analysis of the spectral results is also possible. Significance tests for changes in the spectrum at particular frequencies can be based on the fact that $[8n\ \hat{f}(\omega)]/[3m\ f(\omega)]$ has an asymptotic chi-square distribution with $8n/3m$ degrees of freedom. For example, suppose ω is an indicator frequency of a particular term. Let $\hat{f}_s(\omega)$ be the spectral estimate from a signal run in which the parameters are driven by sinusoidal oscillations, and let $\hat{f}_n(\omega)$ be the spectral estimate from an independent noise run in which the parameters are held fixed. If the term has no effect in the response model and if driving the parameters with sinusoids changes the spectrum only at the indicator frequencies, then the true spectral values $f_s(\omega)$ and $f_n(\omega)$ are equal. Hence the spectral ratio $\hat{f}_s(\omega)/\hat{f}_n(\omega)$ has an F distribution with $8n/3m$ and $8n/3m$ degrees of freedom and can be used to test for the presence of the term in the response model.

3. DESCRIPTION OF THE PROCEDURE

The following steps outline a procedure for screening parameter sensitivities in the frequency domain.

(1) Design the simulation model so the parameters can be driven with sinusoidal oscillations.

(2) Decide what types of terms will be screened; for example, linear effects only or linear effects and interactions. The procedure will select a subset of these terms for inclusion in the response model.

(3) Choose driving frequencies so that the indicator frequencies of the terms are distinguishable from one another.

(4) Make a signal run, driving the parameters with sinusoidal oscillations. Estimate the response spectrum, and look for peaks in the spectrum.

(5) Make a noise run, holding the parameters fixed. Determine the indicator frequencies at which the signal spectrum shows an increase. This determination may be made by either a visual inspection or a statistical analysis.

(6) If it is suspected that there may be effects from a non-uniform system gain, make additional runs changing the driving frequencies assigned to the parameters. The ultimate experimental objective should guide the need for additional runs. If the objective is to determine whether interactions are present to guide the choice of a suitable design for further experimentation, one signal run showing peaks at some interactions will suffice. If the objective is to eliminate unimportant variables from further consideration, additional runs may be needed to be sure that an important term has not been suppressed by a low system gain at its indicator frequency.

4. AN EXAMPLE

To illustrate how spectral analysis can be used to identify a response model, we will use the frequency domain approach to analyze a "black box" linear system. The system response is a function of four input parameters and a noise process:

$$
\begin{aligned}
y_t &= x_{1(t-5)} \\
&+ x_{1(t-5)}^2 \\
&+ \tfrac{1}{3}(x_{2(t-1)}+x_{2(t-2)}+x_{2(t-3)}) \\
&+ \tfrac{1}{3}(x_{2(t-1)}x_{3(t-1)}+x_{2(t-2)}x_{3(t-2)}+x_{2(t-3)}x_{3(t-3)}) \\
&+ a_t
\end{aligned}
$$

where y_t is the response at time t,

x_{it} is the value of parameter i at time t,

a_t is the auto-regressive noise process $a_t = .6a_{t-1} + .8e_t$, and

e_t is a standard normal white noise process.

The response contains x_1 and x_1^2 lagged five periods, three-period moving averages of x_2 and x_2x_3, and some auto-regressive noise. The response is independent of x_4. An ideal screening procedure should identify x_1, x_1^2, x_2, and x_2x_3 as the important terms in the response model.

To begin the frequency domain approach, we decided to limit our screening to first- and second-order effects; namely, linear terms, quadratic terms, and two-way product interactions. We assigned driving frequencies to the parameters so that these effects can be distinguished from one another. One possible set of driving frequencies is .06, .20, .29, and .39. The indicator frequencies associated with each potential first- and second-order term are summarized in Fig. 2. A signal run was made in which each parameter oscillated at its driving frequency with a mean of 0 and an amplitude of 2. The response was observed for three hundred periods, and the estimate of the signal spectrum $\hat{f}_s(\omega)$ was computed.

The signal spectrum showed distinct peaks at the frequencies .06, .09, .12, .20, and .49, which correspond to the terms x_1, x_1^2, x_2, and x_2x_3. To test the significance of these peaks, we made an independent noise run, in which the parameters were held fixed at their mean levels. A different random number stream was used to ensure that the two runs were independent. Again the response was observed for three hundred periods, and the estimate of the noise spectrum $\hat{f}_n(\omega)$ was computed. We computed the signal-to-noise ratio $\hat{f}_s(\omega)/\hat{f}_n(\omega)$. Under the assumption that the true spectral values $f_s(\omega)$ and $f_n(\omega)$ would be equal

Term	Indicator Frequency		Formula for Indicator Frequency
x_1	.06		ω_1
x_2	.20		ω_2
x_3	.29		ω_3
x_4	.39		ω_4
x_1^2	.12		$2\omega_1$
x_2^2	.40		$2\omega_2$
x_3^2	.42		$2\omega_3$
x_4^2	.22		$2\omega_4$
x_1x_2	.14	.26	$\omega_1 \pm \omega_2$
x_1x_3	.23	.35	$\omega_1 \pm \omega_3$
x_1x_4	.33	.45	$\omega_1 \pm \omega_4$
x_2x_3	.09	.49	$\omega_2 \pm \omega_3$
x_2x_4	.19	.41	$\omega_2 \pm \omega_4$
x_3x_4	.10	.32	$\omega_3 \pm \omega_4$

Figure 2
Indicator Frequencies for
Potential Model Terms

if the parameters have no effect, the ratio has an F distribution with 8 and 8 degrees of freedom. The signal-to-noise ratio is graphed in Fig. 3. Note that the spectrum at the indicator frequencies of the terms x_1, x_1^2, x_2, and x_2x_3 is significantly higher in the signal run. We have correctly identified the response model.

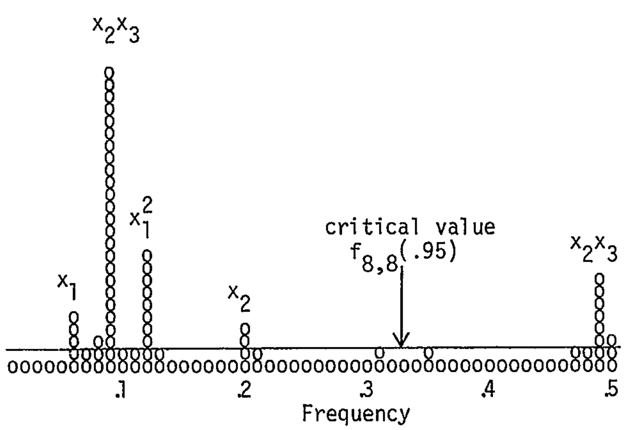

Figure 3
Spectral Ratio Showing Important Terms

5. CONCLUSION

Our experiences with using the frequency domain to characterize parameter sensitivities and identify a response model have been most satisfactory. The use of frequency bands as experimental units increases the amount of information that can be obtained from each run. Several continuous parameters, along with their powers and interactions, may be screened in one or two runs. Practitioners with complicated systems and tight budgets may wish to use this approach to be able to quickly focus their attention on a few important parameters. Researchers may wish to explore further the possibilities offered by the frequency domain.

REFERENCES

Chatfield, C. (1975), The Analysis of Time Series: Theory and Practice, Chapman and Hall, London, 263 p.

Hillier, F.S. and G.J. Lieberman (1980), Introduction to Operations Research, Third Edition, Holden-Day, San Francisco, 829 p.

Kleijnen, J.P.C. (1975), Statistical Techniques in Simulation, Part II, Marcel Dekker, New York, 775 p.