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

Some extensions and limitations of frequency domain experiments

Arnold H. Buss
The John M. Olin School of Business
Washington University
St. Louis, MO 63130, U.S.A.

ABSTRACT

We present three extensions of the frequency domain approach proposed by Schruben and Cogliano (1987). The first is the assignment of multiple frequencies to input factors. The frequency selection in this case is nearly identical to that of one frequency per factor. The second extension is the use of the time series of batch means to flatten the noise spectrum and make the identification of peaks in the output spectrum easier. Finally, we show that using common random numbers (as suggested in Schruben and Cogliano, 1987) for the signal and noise runs do indeed tend to decrease the peaks in the noise spectrum and also facilitate factor identification.

We discuss several limitations to the use of frequency domain methodology as it currently exists. One limitation is the case in which a factor-frequency interaction is present. Although none of the simulation models currently being examined seems to have this property, it is one present in certain chaotic dynamical systems. These systems demonstrate qualitatively different behavior when oscillated. Another limitation involves the time index. For complicated systems it may not be feasible to define a single "clock" with which to oscillate parameters. Current work (Jacobson, *et al.*, 1988) shows that certain models may use the global time clock, but it is not clear that there is always a solution.

1. MULTIPLE FREQUENCY ASSIGNMENTS

In this section we will first discuss the relationship between Frequency Domain Experiments (FDE's) and Response Surface Analysis (RSM). A FDE may be viewed as a dynamic version of a RSM as follows. The RSM approach consists of hypothesizing a meta-model which is polynomial in the factors (Box and Draper). The experimenter specifies a region of interest and performs experiments within that region by specifying the various factor levels and running the simulation for each combination of factor level. The model is fitted by a polynomial regression on the resulting data (Kleijnen, 1987).

The FDE approach also begins by hypothesizing a response meta-model which is polynomial in the factors, but the method of experimentation differs. Rather than one run at each factor value, in a FDE the factors are oscillated throughout the run; this constitutes the *signal* run. A second run is often performed in which each factor is held constant at its nominal value. This is simply a conventional run, but is referred to as a *noise* run. The output power spectrum is computed for each run and the ratio of the spectrum at each frequency is obtained, the *signal to noise ratio*. The experimenter looks for peaks in this ratio, either visually or statistically. Relatively large peaks at certain frequencies are indicators that corresponding combinations of factors are influential in the simulation. Thus, factors may be screened in an efficient manner.

To perform FDE's, the experimenter makes an assumption in addition to those for RSM: oscillations of the factors at their respective *driving frequencies* result in peaks in the output spectrum at sums and differences of certain multiples of those frequencies, the *term indicator frequencies*. These multiples do not exceed the order of the hypothesized polynomial the experimenter assumes for RSM. Some models which exhibit this behavior include the Dynamic Response Surface Model (Schruben, Heath, and Buss; Sanchez and Buss), the Hammerstein model, and the Weiner model (Jacobson, 1987). In this paper we will only consider models of order 2; that is, we assume that the response surface may be approximated by a second order polynomial. This is commonly used for RSM (Box and Draper, 1987). FDE's should thus be viewed as a complement to RSM, at the very least as a relatively efficient method of factor screening (Schruben and Cogliano, 1987).

We will now consider assigning more than one driving frequency to each factor. The problem of selecting driving frequencies has been solved to optimality for up to seven factors in Jacobson, Buss, and Schruben (1986). They also show that driving frequencies may always be selected so that there is no confounding amongst term indicator frequencies. We discuss confounding below. The frequency selection problem for multiple driving frequencies is no more difficult than for single driving frequencies. Each factor is simply treated as being several factors for the purpose of assigning frequencies. The total number of frequencies to be assigned is the sum of the cardinality of the frequencies assigned to each factor. The frequency selection problem is solved using that number as the number of factors. Observe, however, that the interpretation of the term indicator frequencies (linear, quadratic, interaction, etc.) is different in this case than the one frequency per factor approach. For example, all quadratic and interaction term indicator frequencies for those frequencies associated with a single factor are associated with the quadratic term for that factor. Also, there are many more interaction terms: if each of two factors are assigned s and r frequencies, respectively, then there are $2rs$ term indicator frequencies representing interaction between those two variables. An attempt should be made to spread out the frequencies for each factor by assigning it both high and low frequencies. If the system acts in a way that certain frequencies do not affect the output (*e.g.* low pass filters) then using more than one frequency, increases the chance that peaks in the output spectrum will actually be observed when the corresponding terms are in fact in the model.

One disadvantage is the fact that in order to stay within the appropriate region of the response surface, the amplitudes for each term must be reduced roughly in proportion to the number of frequencies it is assigned. However, this may be outweighed by the information obtained about gain. In particular, since each factor will typically have term indicator frequencies at high, medium, and low frequencies, the signal and noise run may be substituted for the Latin square design proposed in Schruben and Cogliano (1987). Of course, since there will be many more frequencies, the distance between term indicator frequencies will be much smaller, and longer runs will be required to get good spectral estimates. On the other hand, due to the greater number of term indicator frequencies for each term we may allow some confounding to occur in the design, particularly amongst interaction terms.

Confounding occurs when distinct terms in the meta-model have the same indicator frequencies. For a given order polynomial meta-model, we may always choose frequencies so that *all* term indicator frequencies are unconfounded. If the order of the polynomial is greater than hypothesized then, as with RSM, there will be model mis-specification. In a FDE this will result in peaks for higher order term indicator frequencies being aliased to term indicator frequencies corresponding to lower order terms. Thus, there will be confounding. This cannot be detected by a FDE, since it is assumed *a priori* that the hypothesized order of the polynomial is large enough. Of course, there is a similar difficulty with RSM, in which most of the analysis deals with a quadratic (second order) model (Box and Draper, 1987).

1.1 Example

We will illustrate multiple driving frequencies with an M/M/1 queue with factors the arrival rate and the service rate. The experiment was run for 50,000 customers and the values as shown in Table 1. Observe that each factor has a high and a low driving frequency. The resulting signal to noise ratio is shown in Figure 1.

Table 1: M/M/1 Queue Parameters

Factor	Nominal Value	Amplitudes	Frequencies
Arrival Rate	0.6	0.35 0.20	0.04 0.22
Service Rate	1.0	0.40 0.25	0.05 0.36

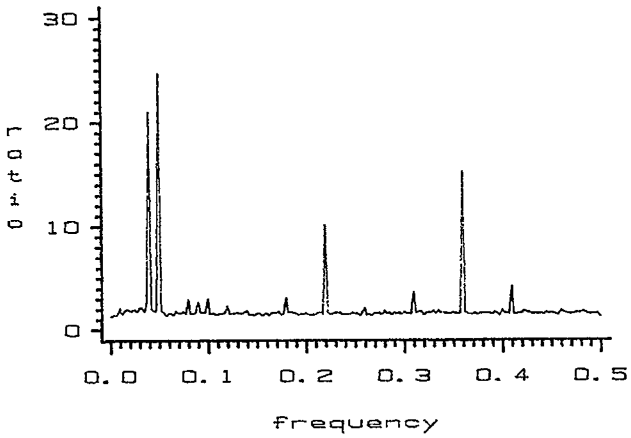

Figure 1: Signal/Noise Ratio for M/M/1 Queue

2. BATCH MEANS

Batch means are a way of reducing serial correlation inherent in simulation experiments (Fishman, 1978). One would therefore expect that using batch means for FDE's would result in a flatter noise spectrum and reduce the effect of gain and frequency-pass qualities of the simulation model. Sargent and his co-workers have extensive experience with the use of batch means in queueing and computer network models (Sargent, 1988). In this section we will show that the output spectra of FDE's are extremely sensitive to the exact value of the batch size.

We will be making the following assumptions for the FDE:

1) The noise run is a (second order) stationary process with covariance function $R_N(n)$.

2) The signal run is the sum of two uncorrelated processes: one having covariance function $R_N(n)$ (the same as for the noise run), and the other having *periodic* covariance function $R_S(n)$ with finitely many frequencies.

We first consider a steady state output $\{X_n\}$ with covariance function $R(n)$ (of either signal or noise form) and define the batch means Y_j as usual: $Y_j = (X_{(j-1)b+1} + \ldots + X_{jb})/b, \quad j = 1, \ldots, m$, where b is the batch size, m is the number of batches, and $n = mb$. Let $\bar{R}_b(k)$ be the covariance between batches j and $j+k$ (observe that $\bar{R}_b(k)$ is independent of j by stationarity). Then

$$\bar{R}_b(k) = \frac{1}{b} \sum_{s=-(b-1)}^{b-1} \left(1 - \frac{s}{b}\right) R(kb+s). \quad (1)$$

Furthermore, if $\sum_{n=-\infty}^{\infty} |R(n)| < \infty$ then (Law and Carson, 1982) $\bar{R}_b(k) \to \delta_{k0}$, where δ_{ij} is the Kronecker delta function:

$$\delta_{ij} = \begin{cases} 1, & \text{if } i = j \\ 0, & \text{if } i \neq j. \end{cases}$$

The batches are thus asymptotically uncorrelated. The corresponding property in the frequency domain is that the batch spectrum, the Fourier transform of the covariance of the sequence of batch means, becomes flat. Let $h(\omega)$ be the generalized spectral density for $\{X_n\}$ and $\bar{h}_b(\omega)$ the corresponding spectral density for $\{Y_k\}$, the sequence of batch means. We have

$$b\bar{h}_b(\omega) = \sum_{s=-(b-1)}^{b-1} \left(1 - \frac{|s|}{b}\right) \sum_{k=-\infty}^{\infty} e^{2\pi i \omega k} R(kb+s). \quad (2)$$

Thus, if $R(n)$ is absolutely summable, then $b\bar{h}_b(\omega) \to h(0)$ as $b \to \infty$, for $0.0 \leq \omega \leq 0.5$. Equivalently, $\bar{h}_b(\omega)/\bar{h}_b(0) \to 1$ (Heidelberger and Welsh, 1981).

The situation is quite different in the presence of discrete spectra. Indeed, if $R_S(n)$ is not identically zero, then $\sum_{n=-\infty}^{\infty} |R(n)| = \infty$. For simplicity, consider a covariance consisting of a single frequency term at ω_0: $R(n) = \cos 2\pi\omega_0 n$. Then by Equation (2) we have:

$$b\bar{h}_b(\omega) = \sum_{s=-(b-1)}^{b-1} \left(1 - \frac{|s|}{b}\right) \cos 2\pi\omega_0 s \cdot \sum_{k=-\infty}^{\infty} e^{2\pi i \omega k} cos 2\pi\omega_0 bk$$

$$= F_b(\omega_0)\delta(\omega - \omega_0 b), \quad (3)$$

in which $F_N(\omega) = \sin^2(\pi\omega N)/N sin^2(\pi\omega)$ is the Fejer kernel and k is the Dirac delta function. In the above we have added the "power" at $+\omega$ and $-\omega$ to obtain the one-sided power spectrum. The contribution of the oscillation of the covariance function at frequency ω_0 contributes power at $b\omega_0$, but the Fejer kernel tends to decrease the power for $b\omega_0 > 1$. Now, $b\omega_0$ will be aliased into the interval [0.0,0.5], so the exact location of this peak depends on b. Let $\omega_0 = r/d$ where r and d are relatively prime integers. We are confined to rational frequencies, not only because of the finite precision of the computer, but because optimal frequencies will be rational and every known algorithm for determining a non-confounding set of driving frequencies results in rational frequencies (Jacobson, Buss, and Schruben, 1987). Choosing batch size b a multiple of d, the frequency's denominator, will result in all term indicator frequencies being aliased to the origin, which is clearly undesirable. However, since we have complete control over both batch size and frequency selection, we can always choose the driving frequnecies to avoid this difficulty. For example, choosing $b = kr + 1$ for integer k results in $b\omega_0$ being aliased to ω_0. In general, different batch sizes lead to peaks in the output spectrum at any multiple of $1/r$. The difficulty with this is that the Fejer kernel is extremely small for batch sizes greater than the denominator of the frequencies.

It is clear that batch means should be used for FDE's with great caution, since the term indicator frequencies do get shifted from their unbatched locations. Furthermore, the presence of the Fejer kernel in the batch spectrum at non-zero values could suppress peaks which would otherwise be observed. This would be an example of the "false negative," a situation in which a peak *should* be observed since the corresponding term is truly in the model, but the peak is not observed (see Morrice and Schruben, 1988 for further discussion of the false negative problem). However, since batching may be viewed as a low pass linear filter (Heidelberger and Welsh, 1981), it seems reasonable that careful design of the FDE should avoid these difficulties.

One solution is to divide the original driving frequencies by the batch size b. Clearly the term indicator frequencies from the batch spectrum will then be the original term indicator frequencies, since each will be multiplied by b. For even moderate batch sizes, the effect will be to have all parameters oscillating at very low frequencies. However, batching has the effect of increasing the frequencies, and this may always be done so that the resulting batch term indicator frequencies are those which the experimenter would have chosen in the output had not been batched.

Let us examine the impact of choosing frequencies in this manner on the output spectra. For a single frequency ω_0 as above we have

$$b\bar{h}_b(\omega_0) = F_b(\omega_0/b)\delta(\omega - \omega_0) \tag{4}$$

so that the peak is now at ω_0. Now observe that

$$F_b(\omega_0) = \frac{\sin^2 \pi\omega_0}{b \sin^2 \pi\omega_0/b}$$

$$= b\left(\frac{\sin^2 \pi\omega_0}{\pi\omega_0}\right)^2 \left(\frac{\pi\omega_0/b}{\sin \pi\omega_0/b}\right)^2.$$

For large batch sizes b the rightmost term is nearly unity, so that $F_b(\omega_0/b) = O(b)$ as $b \to \infty$. Therefore, we would expect the signal to noise ratio corresponding to batch size b to increase, at term indicator frequencies for terms present in the model, about on the order of b. For small frequencies the increase is nearly b, while for high frequencies (near 0.5) the increase is about $b(2/\pi)^2$.

We note one final consequence of batch means flattening the spectrum. The purpose of performing signal and noise runs and taking the spectral ratio is to mitigate effects of gain, namely properties of the system which tend to enhance or suppress certain frequencies. If the noise spectrum were truly flat then we would only have to perform the signal run. Since batch means tend to flatten the spectrum if to periodicities are present, while enhancing the peaks of frequencies corresponding to the discrete spectrum, we envision the possibility of dispensing with the noise run.

2.1 Example

We will illustrate the impact of batching on the M/M/1 queue with the arrival and service rates the parameters of interest. In each case the nominal values of arrival and service rates were 0.6 and 1.0, respectively and the (batched) driving frequencies were 0.04 and 0.4, respectively. The settings are given in Table 2 below.

Figure 2a shows the low pass nature of the queue, while Figures 3a and 4a demonstrate the gradual flattening of the spectrum with increasing batch size. As these three are all the noise runs (no oscillation of parameters), this replicates behavior previously observed (*e.g.* Heidelberger and Welsh, 1981).

Turning to the signal spectra, we observe increasingly higher peaks at 0.04 and at 0.40 as the batch size increases (Figures 2b, 3b, and 4b, respectively), while the width of the peaks increases due to the correspondingly fewer observations. This effect is more pronounced for the lower frequency than the higher one. As batch size increases, the second order effects become more apparent. The square of the arrival rate appears as a peak at 0.08 and the interaction between arrival and service rates as small but distinct peaks at 0.36 and 0.44. The signal/noise ratios (Figures 2c, 3c, and 4c) show the increasing heights of the peaks as well as the increase in order of magnitude described above.

Table 2: Frequencies and Batch Sizes for the M/M/1 Queue

Run	Batch Size	# Batches	ω_1 (Arrival)	ω_2 (Service)
1	1	50,000	0.04	0.40
2	10	5,000	0.004	0.040
3	100	500	0.0004	0.0040

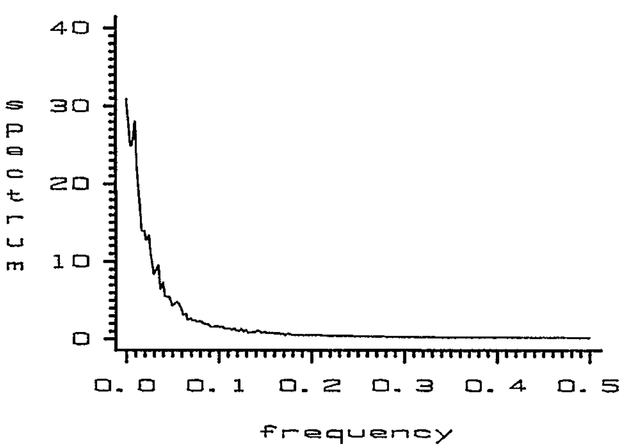

Figure 2a: Noise Run for M/M/1 Queue
Batch Size = 1

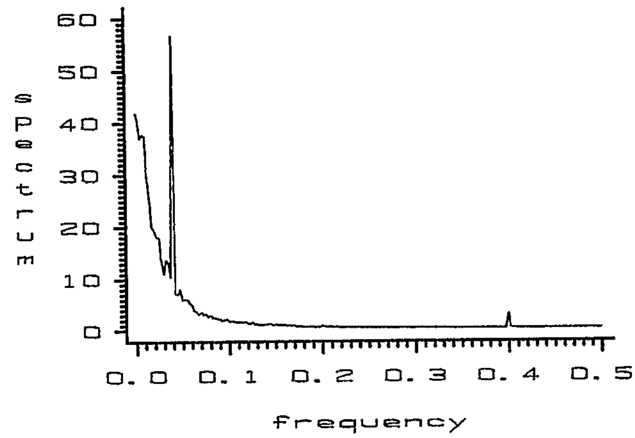

Figure 2b: Signal Run for M/M/1 Queue
Batch Size = 1

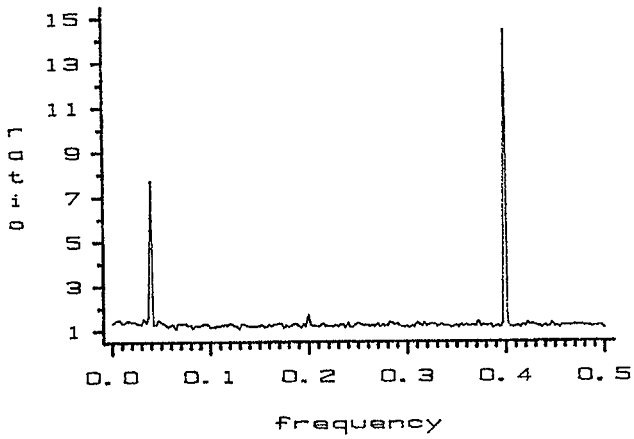

Figure 2c: Signal/Noise Ratio for M/M/1 Queue
Batch Size = 1

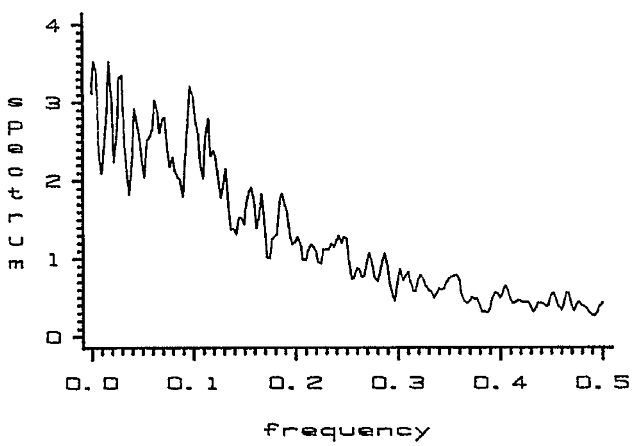

Figure 3a: Noise Run for M/M/1 Queue
Batch Size = 10

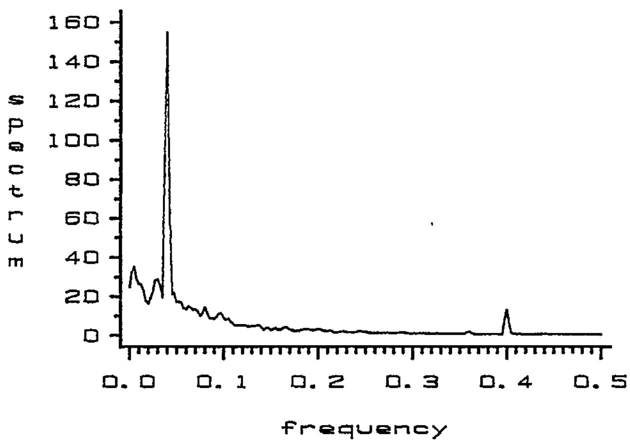

Figure 3b: Signal Run for M/M/1 Queue

Batch Size = 10

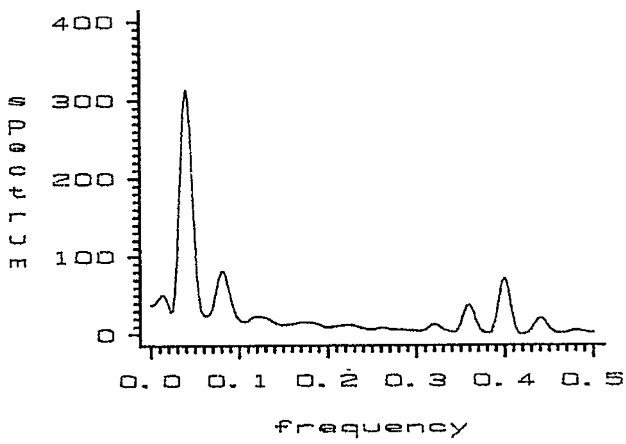

Figure 4b: Signal Run for M/M/1 Queue

Batch Size = 100

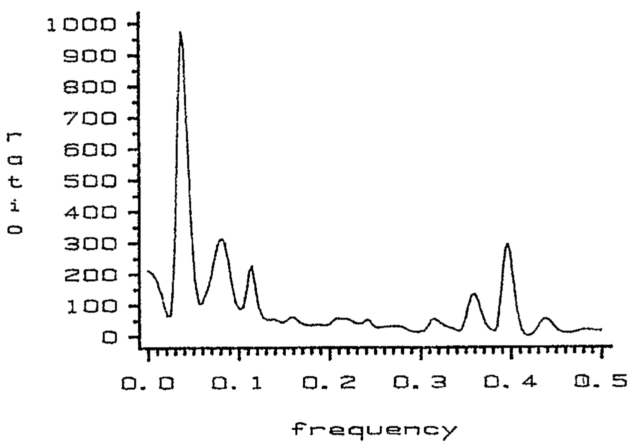

Figure 3c: Signal/Noise Ratio for M/M/1 Queue

Batch Size = 10

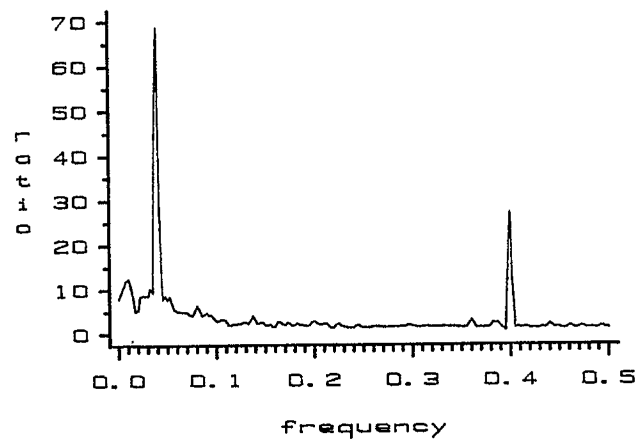

Figure 4c: Signal/Noise Ratio for M/M/1 Queue

Batch Size = 100

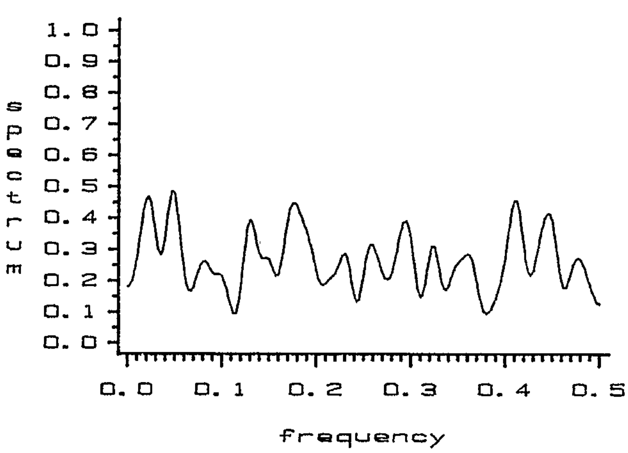

Figure 4a: Noise Run for M/M/1 Queue

Batch Size = 100

3. COMMON RANDOM NUMBERS FOR SIGNAL AND NOISE RUNS

In this section we investigate the effects of using common random numbers for the signal and noise runs. This was suggested by Schruben and Cogliano (1987). The usual F test under the null hypothesis of no significant frequency response is not appropriate, even under the assumption that the spectral estimators are asymptotically chi-square. This is because the F distribution is based on ratios of *independent* chi-square variables.

We will focus on a single frequency ω_0 and begin by assuming that the spectral estimators are asymptotically chi-square with ν degrees of freedom, where ν is the "equivalent" degrees of freedom and depends on the exact type of window used (Priestley, 1981). We further assume that the chi-square variables are obtained from the sums of squares of "equivalent" standard normal random variables. We will initially assume that the signal and noise run are of identical length, so ν is the same for each. Thus, under the null hypothesis of no factor effect, the outputs from the signal and noise runs are considered to be transformed into $\{X_1, \ldots, X_\nu\}$ and $\{Y_1, \ldots, Y_\nu\}$, respectively, where the X_i's and Y_i's are two sets of ν independent standard normal random variables. The spectral estimators at ω_0 are asymptotically proportional to $\sum_{i=1}^{\nu} X_i^2$ and $\sum_{i=1}^{\nu} Y_i^2$ for the signal and noise runs, respectively. Finally, we assume that the effect of the common random numbers is to induce correlation ρ between X_i and Y_i, identical for each i, but the pairs (X_i, Y_i) are independent and jointly normal.

Under these assumptions, the joint distribution of the spectral estimators for the signal and noise runs are proportional to the following bivariate gamma density (Johnson and Kotz, 1971):

$$f(x,y) = \frac{(\sqrt{xy})^{(\nu-2)/2} e^{-(x+y)/2(1-\rho^2)}}{4(1-\rho^2)\Gamma(\nu/2)}. \tag{5}$$

Here, $I_\mu(z)$ is the modified Bessel function of the second type of order μ. The density in Equation (5) may be written as a mixture of univariate gamma densities:

$$f(x,y) = \sum_{k=0}^{\infty} \eta(k;\rho^2,\nu/2)\gamma(x;\nu/2+k,1/[2(1-\rho^2)]) \cdot \gamma(y;\nu/2+k,1/[2(1-\rho^2)])$$

in which

$$\eta(k;p,N) = \binom{N+1}{N-1} p^k (1-p)^N, \quad k = 0, \ldots$$

is the negative binomial mass function and

$$\gamma(x;\alpha,\lambda) = \frac{\lambda^\alpha e^{-\lambda x}}{\Gamma(\alpha)}$$

is the gamma density function. Thus, given K (a negative binomial random variable), the joint distribution of the spectral estimators (under the null hypothesis) are proportional to independent gamma random variables with parameters $\nu/2 + K$ and $1/[2(1-\rho^2)]$. We may therefore write the density of the signal to noise ratio as a mixture of F distributions:

$$f(z) = \sum_{k=0}^{\infty} \eta(k;\rho^2,\nu/2) f_{\nu+2k,\nu+2k}(z)$$

in which $f_{\nu+2k,\nu+2k}(z)$ is the density of an F distribution with $\nu + 2k$ degrees of freedom in both the numerator and denominator. This density was apparently first studied by Bose(1938).

Using the moments of the F distribution, the mean and variance of this distribution are

$$\mu = \frac{\nu - 2\rho^2}{\nu - 2} \qquad (\nu > 2)$$

and

$$\sigma^2 = \frac{4\nu(nu-1)(1-\rho^2)}{(\nu-2)^2(\nu-4)} - \frac{4\rho^2(1-\rho^2)(5\nu-8)}{(\nu-2)^2(\nu-2)} \qquad (\nu > 4)$$

Thus, using common random numbers should result in smaller signal/noise ratios under the null hypothesis. When the term for which the indicator frequency in question is in fact *not* present in the model, then using common random numbers should result in smaller peaks, as desired. Thus, using common random numbers is desirable when we are particularly interested in detecting the *absence* of factors in the model.

The distribution of the signal to noise ratio depends on the unknown parameter, ρ, which must be estimated. Since this is the cross-correlation (with no lag) between the signal and noise runs, the usual estimator is asymptotically unbiased and consistent. Therefore, with large samples sizes there should be no difficulty replacing ρ with its estimator $\hat{\rho}$. The exact distribution appears to be extremely complicated.

Suppose the equivalent degrees of freedom for the signal run are greater than those for the denominator. If we assume that the extra degrees of freedom come from more normal random variables independent of the noise run, we have the following distribution of the signal to noise ratio for a given frequency under the null hypothesis:

$$f(x,y) = \sum_{k=0}^{\infty}\sum_{h=0}^{\infty} \eta(k;\rho^2,\nu/2)\eta(h;\rho^2,\nu/2) \cdot \gamma(x;\nu/2+\mu/2+k+h, 1/[2(1-\rho^2)] \cdot \gamma(x;\nu/2+k, 1/[2(1-\rho^2)]$$

in which there are $\nu + \mu$ degrees of freedom in the numerator and ν in the denominator.

There does not appear to be a nice representation of the spectral ratio under these conditions when the alternative hypotheses (existence of a peak in the signal spectrum) is true. One could approximate it by replacing the F distributions with non-central F distributions corresponding to the numerator in the negative binomial mixture above.

4. LIMITATIONS OF FREQUENCY DOMAIN EXPERIMENTS

In this section we will discuss some limitations to the use of Frequency Domain Methodology. Although it appears to work very well for a number of models, Frequency Domain Methodology is still relatively untested. Its good and bad properties are not known to the extent of conventional models. This in itself has been used as a basis for dismissing FDE's as a useful tool for analysis. However, all established techniques were once new and untested, so we do not feel that this justifies the denial of benefits of FDE's. The experience of the author, as well as other workers in the field of frequency domain methodology, has been then FDE's work extremely well, even when perhaps they shouldn't. This suggests a certain robustness to the Frequency Domain approach. However, there appear to be situations in which we might expect FDE's to perform badly. These situations should be those in which interaction of the frequencies of oscillation and the factors themselves is present.

The issue of validity and robustness is far from settled. There are several papers which discuss the kinds of models for which FDE's will perform well (Schruben, Heath, and Buss, 1988; Jacobson, 1988; Sanchez and Buss, 1987). These models include static models as special cases: regression, RSM, etc. Each of these models has the difficulty of the respective analysis invalidated by model mis-specification.

However, there should be cases in which FDE's simply do not work due to gross violation of the dynamic assumptions. A class of models from nonlinear dynamics, the so-called chaotic systems, exhibit qualitatively different behavior when excited by periodic functions (Guckenheimer and Holmes, 1983). Two well-studied examples of such systems are the Duffing oscillator and the Van der Pol Oscillator. These arise from relatively simple nonlinear models of physical phenomena. Although the researchers in that field do not think of them as such, we may consider such models to be FDE's. The unforced system corresponds to the noise run and the forced system (essentially replacing a nominal value of zero by a sine wave) corresponds to the signal run. The unforced system may decay to a stable value or to a periodic limit cycle. The forced system, however, can exhibit either periodic or "chaotic" behavior. The spectrum of periodic systems consists of peaks at the appropriate frequencies, while the chaotic spectrum is typically wild, resembling more a white noise spectrum than a periodic one. Evidently the type of behavior (periodic or chaotic) depends on the frequency and the amplitude of the forcing sinusoid. the behavior of the oscillated system depends on the frequency. If the experimenter is so unfortunate to attempt a FDE on such a system, the conclusions could be radically different when different frequencies are used.

The second limitation concerns the time index. There are two types of simulations for which a consistent time index may *always* be selected: synchronous, or "time slice" simulations and FIFO simulations, such as the M/M/1 queue. Somewhat more complicated systems have a defineable time index, but in some cases one has to resort to a re-ordering of the output (Som, *et al.*, 1987). In some circumstances one need not re-order the output, but instead may determine effects of "shuffling" on the term indicator frequencies (Jacobson, *et al.*, 1988) and use the simulation time clock as the index. For extremely compliacted systems, however, with multiple entities and vast amounts of shuffling within it is not apparent that a consistent time index may always be found. It is a limitation which Frequency Domain Methodology must overcome if it is to prove a useful tool for analyzing large simulation models. If we can understand the circumstances under which FDE's will *not* work, we may understand better which circumstances under which they *will* work.

5. SUMMARY AND CONCLUSIONS

We have presented three extensions to Frequency Domain Methodology as originally proposed by Schruben and Cogliano (1987): use of multiple frequencies for each factor, use of the batch means time series for spectral estimation, and the use of common random numbers for signal and noise runs. The initial results for each of these are encouraging. Further study is needed to demonstrate the utility of these techniques. We have discussed some validity issues concerning Frequency Domain Methodology itself. It would appear that factor-frequency interaction is the most damaging characteristic of a system to have if the experimenter wishes to perform FDE's on it. Chaotic dynamical systems are one class of models for which this occurs.

REFERENCES

Bose, S. (1935) On the Distribution of the Ratio of Variances of Two Samples Drawn from a Given Bivariate Normal Population, *Sankhya*, **2**, 65-72.

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

Bratley, P., Fox, B., and Schrage, L. (1987) *A Guide to Simulation, Second Edition*, Springer-Verlag, New York, NY.

Fishman, G. (1978) Grouping Observations in Digital Simulation, *Management Science*, **24**, 510-521.

Guckenheimer, J. and Holmes, P. (1983) *Nonlinear Oscillations, Dynamical Systems, and Bifurcation of Vector Fields*, Springer-Verlag, New York, NY.

Heidelberger, P. and Welsh, P. (1981) A Spectral Method for Confidence Interval Generation and Run Control Length, *Communications of the Association for Computing Machinery*, **24**, 233-245.

Jacobson, S. (1988) PhD Thesis, School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY.

Jacobson, S., Buss, A., and Schruben, L. (1986) Driving Frequency Selection for Frequency Domain Simulation Experiments, Technical Report No. 714, School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY.

Jacobson, S. Morrice, D., and Schruben, L. (1988) The Global Simulation Clock as the Frequency Domain Experiment Index, *Proceedings of the 1988 Winter Simulation Conference* (M. Abrams, ed.), San Diego, CA.

Johnson, N. and Kotz, S. (1971) *Continuous Multivariate Distributions*, John Wiley, New York, NY.

Kleijnen, J. (1987) *Statistical Tools for Simulation Practitioners*, Marcel Dekker, New York, NY.

Law, A. and Carson, J. (1979) A Sequential Procedure for Determining the Length of a Steady-State Simulation, *Operations Research*, **27**, 1011-1025.

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

Sargent, R. (1988) Personal Communication.

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

Schruben, L., Heath, D., and Buss, A. (1988) A Dynamic Response Surface Model for Frequency Domain Simulation Experiments, Technical Report No. 808, School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY.

Som, T., Sargent, R., and Schruben, L. (1987) Frequency Domain Metamodelling of a Feedback Queue, *Proceedings of the 1987 Winter Simulation Conference* (A. Thesen, H. Grant, and W. D. Kelton, eds.), Atlanta, GA.

AUTHOR'S BIOGRAPHY

ARNOLD BUSS is Assistant Professor of Management Science at The John M. Olin School of Business. He received his PhD in Operations Research from Cornell University. His research interests include bizarre methods of simulation analysis, stochastic models, and dynamic cement.

Arnold Buss
The John M. Olin School of Business
Campus Box 1133
Washington University
One Brookings Drive
St. Louis, MO 63930