David O. Smallwood
Thomas L. Paez
Sandia National Laboratories
Albuquerque New Mexico 87185

A Frequency Domain Method for the Generation of Partially Coherent Normal Stationary Time Domain Signals

A procedure for generating vectors of time domain signals that are partially coherent in a prescribed manner is described. The procedure starts with the spectral density matrix, $[G_{xx}(f)]$, *that relates pairs of elements of the vector random process* $\{\{x(t)\}, -\infty < t < \infty\}$. *The spectral density matrix is decomposed into the form*

$$[G_{xx}(f)] = [U(f)][S(f)][U(f)]'$$

where $[U(f)]$ *is a matrix of complex frequency response functions, and* $[S(f)]$ *is a diagonal matrix of real functions that can vary with frequency.*

The factors of the spectral density matrix, $[U(f)]$ *and* $[S(f)]$, *are then used to generate a frame of random data in the frequency domain. The data is transformed into the time domain using an inverse FFT to generate a frame of data in the time domain. Successive frames of data are then windowed, overlapped, and added to form a vector of normal stationary sampled time histories,* $\{x(t)\}$, *of arbitrary length.* © *1993 John Wiley & Sons, Inc.**

INTRODUCTION

The generation of realizations of a vector of random processes, which are partially coherent in a prescribed manner, is of interest in testing and dynamic analysis of structures exposed to a variety of natural environments described as stochastic processes. Examples include an airplane or missile exposed to turbulent flow, the response of a ship or sea platform to a confused sea, or the response of a building to wind or an earthquake.

Several techniques have been developed to generate these time histories. ARMA (auto regressive moving average) models have been used since about 1970 [Spanos and Mignolet, 1989]. Time Data Corp. developed and patented (U.S. Patent 3,848,115) a frequency domain process for

*This article is a US Government work and, as such, is in the public domain in the United States of America.

This work was supported by the United States Department of Energy under Contract DE-AC04-76DP00789.

Requests for reprints should be sent to David O. Smallwood.

Shock and Vibration, Vol. 1, No. 1, pp. 45–53 (1993)

CCC 1070-9622/93/01045-09

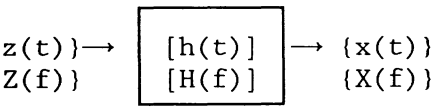

FIGURE 1 Multiple input/multiple output problem.

a single time history. This method was employed in the development of single-channel random vibration controllers in the early 1970s. The first publication detailing this method was by Tebbs and Hunter in 1974. The method was extended to vectors of time histories by Smallwood in 1978.

This article will discuss a generalization of the frequency-domain methods discussed above.

THEORY

Consider the general multiple input/multiple output linear problem as illustrated in Fig. 1. A vector of inputs, $\{z(t)\}$ excites a linear system resulting in a vector of outputs, $\{x(t)\}$. The matrix $[H(f)]$ describes the frequency response functions between all pairs of inputs and outputs. The Fourier transforms of the inputs and outputs are related through the equation

$$\{X(f)\} = [H(f)]\{Z(f)\} \tag{1}$$

where the upper case variable is a function of frequency, f, and is the Fourier transform of the corresponding lower case variable that is a function of time, t.

It is well known that the cross-spectral density matrix of the output and the input are related through the equation [Dodds and Robson, 1975; Bendat and Piersol, 1986]

$$[G_{xx}(f)] = [H(f)][G_{zz}(f)][H[(f)]' \tag{2}$$

where []′ is the conjugate transpose of a matrix.

For the special case where the inputs, $\{z(t)\}$, are independent, the spectral density of the inputs, $[G_{zz}(f)]$, will be diagonal, with real elements. The diagonal elements will be the auto-spectral densities of each process, and the cross-spectral densities, the off-diagonal terms, will all be zero.

Now consider the decomposition of the cross-spectral density of the outputs, $[G_{xx}(f)]$, into the form

$$[G_{xx}(f)] = [U(f)][S(f)][U(f)]'. \tag{3}$$

$[U(f)]$ is a matrix of complex factors that are functions of frequency. $[S(f)]$ is a diagonal matrix of real functions of frequency.

After the decomposition is accomplished, the factor $[U(f)]$ of the decomposition is associated with the frequency-response functions, $[H(f)]$ of Eq. (2) and the diagonal factor, $[S(f)]$ is associated with the auto-spectral density of the independent inputs, $[G_{zz}(f)]$, in Eq. (2). The frequency-response functions, $[H(f)]$, associated with the factors $[U(f)]$ do not need to have any physical significance, they are merely used as a model of the system as will be seen later.

It is important to realize that any decomposition of the cross-spectral density matrix that yields the form of Eq. (3) will be acceptable. Several methods of decomposition that have been or could be used are described below. Each method has particular advantages in certain applications. Independent noise sources are then generated in the frequency domain, coupled with the established frequency-response functions, transformed to the time domain to generate the partially coherent realizations. These steps are outlined in Fig. 2. Each step will now be discussed in detail.

SPECTRAL DECOMPOSITIONS

Cholesky Decomposition (CD)

The first method used to decompose the spectral-density matrix was Cholesky decomposition [Smallwood, 1978; Dodds and Robson, 1975]. Using CD the assumption is made that the cross-spectral density matrix of the independent inputs is white with unity amplitude, hence

$$[G_{zz}(f)] = [S(f)] = [I] \tag{4}$$

where $[I]$ is the identity matrix, and Eq. (3) reduces to

$$[G_{xx}(f)] = [U(f)][U(f)]' \tag{5}$$

where, $[U(f)]$ is a lower triangular matrix. The diagonal elements are real and the off-diagonal elements are complex.

By writing out the terms for the elements of $[G_{xx}(f)]$ in Eq. (5) it is straightforward to derive recursive relations for the elements in $[U(f)]$ from previously derived elements of $[U(f)]$ and elements of $[G_{xx}(f)]$. The recursive formulas for

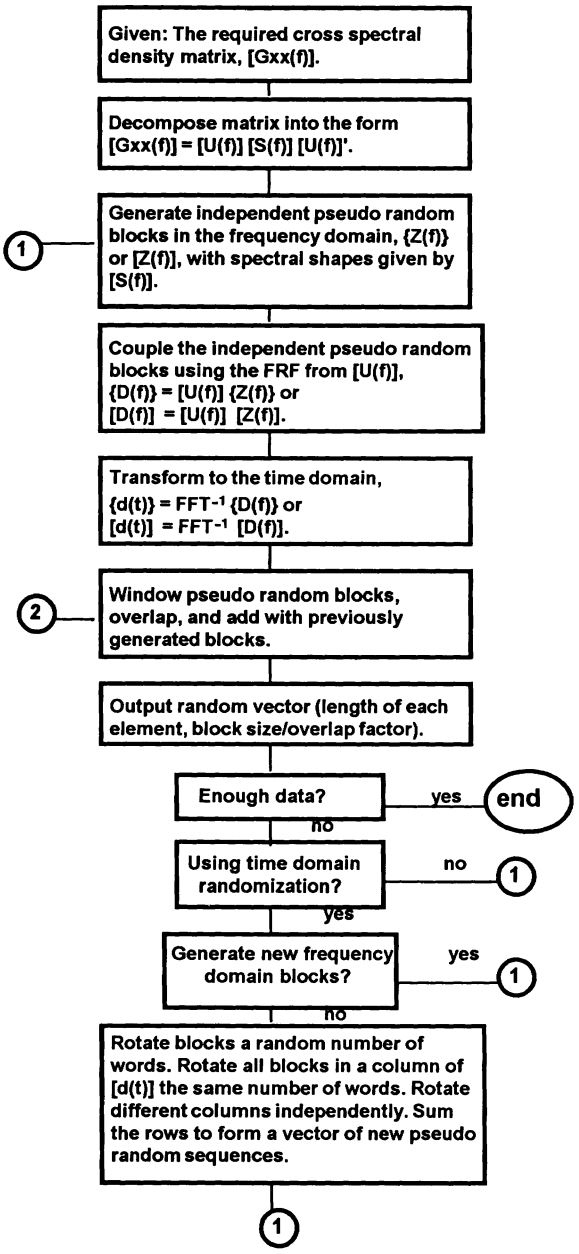

FIGURE 2 Frequency domain method.

finding the elements of $[U(f)]$ in terms of the cross-spectral density matrix are given by

$$U_{kk} = \left[G_{kk} - \sum_{i=1}^{k-1} U_{ki}U_{ki}^*\right]^{1/2} \qquad k = 1, \ldots, N \quad (6)$$

$$U_{jk} = \left[G_{jk}^* - \sum_{i=1}^{k-1} U_{ki}U_{ji}^*\right] \Big/ U_{kk} \qquad j = k + 1, N \quad (7)$$

where * is the complex conjugate.

The sum is not performed for $k = 1$.

Cholesky Decomposition without Square Roots (CDw/oSQRT)

A modified form of Cholesky decomposition can be performed without taking the square roots in Eq. (6) [Lawson and Hanson, 1974]. This is the method used by Bendat and Piersol [1986] in analyzing multiple input data. The spectral-density matrix is decomposed in the form of Eq. (3), where $[U(f)]$ is a lower triangular matrix with unity on the diagonal, and $[S(f)]$ is a diagonal matrix whose amplitudes determine the auto-spectral densities of the independent noise sources. The recursive formulas for finding the elements of $[U(f)]$ and $[S(f)]$ are

$$S_{kk} = \left[G_{kk} - \sum_{i=1}^{k-1} U_{ki}S_{ii}U_{ki}^*\right] \qquad k = 1, \ldots, N \quad (8)$$

$$U_{kk} = 1 \qquad k = 1, \ldots, N$$

$$U_{jk} = \left[G_{jk}^* - \sum_{i=1}^{k-1} U_{ki}S_{ii}U_{ji}^*\right] \Big/ S_{kk} \qquad j = k + 1, \ldots N. \quad (9)$$

As before the sum is not performed when $k = 1$.

Cholesky Decomposition with Zeros (CDwZ)

A major weakness of the above two methods of Cholesky decomposition is that the methods will fail if the rank of the cross-spectral density matrix is less than the size of the matrix. The rank of the cross-spectral density matrix can be a function of frequency, complicating the problem. This can easily happen if the number of underlying sources (which may be unknown) is less than the number of measurement points. For example, consider a structure with response measured at N points. Assume the structure is excited by M independent random sources, where $M < N$. The cross-spectral density matrix of the responses will be of rank M. Another example would be the measurement of a set of responses on a structure dominated by a single mode of vibration at a particular frequency. The rank of

the cross-spectral density matrix at that frequency can be almost one. For these examples Cholesky decomposition will fail. In these cases some of the diagonal elements will be zero (or negative due to noise) and Eq. (7) will fail when the divide by zero is attempted. For CDw/oSQRT diagonal elements in $[S(f)]$ will be zero and the divide by zero in Eq. (9) will fail. This shortcoming can be overcome by setting all the elements in the corresponding column to zero when a zero or negative value is found for a diagonal element in the CD method. For the CDw/oSQRT method the corresponding element of $[S(f)]$ is set to zero, and all the corresponding off-diagonal elements in a column of $[U(f)]$ are set to zero.

The above three methods of Cholesky decomposition provide a very useful tool to interpret random data, and have proved very useful in both random vibration testing and the analysis of random data.

Singular Value Decomposition (SVD)

SVD is a robust method for decomposing matrices in the form

$$[A] = [U][S][V]' \tag{10}$$

where $[S]$ is a diagonal matrix with nonnegative real elements in decreasing order. $[U]$ and $[V]$ are complex, the columns forming orthogonal vectors. For the special case of a Hermitian matrix, $[A]$ (a spectral density matrix is Hermitian, meaning $A_{ij} = A_{ji}^*$), Eq. (10) reduces to Eq. (3), because $[U] = [V]$.

The algorithms are too complicated to give in this paper, but numerous references are available [Lawson and Hanson, 1974; MATLAB, 1993; Press, Flannery, Teukolsky, and Vetterling, 1986].

SVD is a very powerful, robust method and seldom fails. The values on the diagonal of $[S]$ are called the singular values of the matrix. If the rank of the matrix is less than the dimension size of the matrix, SVD will return zeros, or values that are on the order of the computer roundoff, for the singular values.

SVD also provides a consistent way to handle errors in the matrix. Errors of magnitude ε in a rank deficient matrix will produce singular values of magnitude ε. Loosely, a singular value of magnitude ε means that changes of magnitude ε to the original matrix could produce a zero singular value. This implies that if an estimate of the errors in a matrix is known, then singular values of less than this error are not meaningful. A useful tool is to set these singular values to zero.

When the SVD is interpreted as Eq. (2), the singular values give the number and strength of the underlying independent sources in the data. The vectors in $[U]$ give the contribution of each source to each location where data was taken.

However, SVD does not provide as good an intuitive model to analyze random data as does CD [Smallwood, 1978; Dodds and Robson, 1975; Bendat and Piersol, 1986]. A useful compromise might be to use SVD to analyze the noise and the number of independent sources in the data. The noise would be removed by setting nonmeaningful singular values to zero. The spectral-density matrix would be reformed with the noise removed using Eq. (3). The modified spectral-density matrix would be decomposed using CDwZ to provide the intuitive model for analyzing the data. This needs to be investigated further.

GENERATION OF INDEPENDENT RANDOM SOURCES IN THE FREQUENCY DOMAIN

In order to use the decomposition of the spectral density into the form of Eq. (3) with the model provided by Fig. 1 and Eq. (2), frames or blocks of independent random sources need to be generated in the frequency domain with a specified spectral density. Two methods are available to accomplish this and are discussed below.

Specified Spectral Amplitude and Random Phase (RP)

The simplest and first method used is to specify the magnitude of the spectral lines of a discrete Fourier transform (DFT) by taking the square root of the values from the decomposition, $[S(f)]$. The values are normalized to account for the bandwidth between the lines. This will be a line spectrum. The phase angles, ϕ, are then randomly selected for each line, uniformly distributed from 0 to 2π.

$$Z(f) = S^{1/2}(f)e^{i\phi} \tag{11}$$

where $Z(f)$ is the element corresponding to a generic element, $S(f)$, from the matrix $[S(f)]$.

To improve execution time, some quite crude

methods have been used successfully. For example, several vibration control systems simply pick the sign of the real and imaginary parts of each line at random. This requires only two random bits/spectral line. This randomly picks one of four phase angles. To avoid the time to compute the random numbers, bits are selected from a precomputed and stored table of random bits. A better method would be to use a well-tested random number generator. One frame of data of specified amplitudes and random phases is needed for each independent noise source required.

White Noise Method (WN)

White noise, $n(i)$, can be generated in the time domain by generating a sequence of L Gaussian independent random samples where L is the length of the DFT, the block size. The discrete Fourier transform (usually the fast Fourier Transform, FFT) is then computed for this sequence. This gives one frame of independent data in the frequency domain. Using this method both the amplitude and phase of the frame will have the proper distributions. If required, the data can then be colored by multiplying by the square root of the desired spectral density.

$$Z(f) = \text{FFT}[n(i)]S^{1/2}(f). \tag{12}$$

A limitation of the random-phase method for narrow-band processes will be discussed later.

GENERATION OF PARTIALLY COHERENT NOISE SOURCES IN THE FREQUENCY DOMAIN

The independent noise sources generated in the previous step are then multiplied by frequency response functions determined by associating the factors $[U(f)]$ from the decomposition with the model in Fig. 1. Two forms can be used. First, the independent sources from the previous step can be considered a vector. This leads to a vector of coupled sources in the frequency domain, $\{D)(f)\}$, as shown in Fig. 3.

$$\{D(f)\} = [U(f)]\{Z(f)\}. \tag{13}$$

Second, if time-domain randomization (to be described later) is used, the independent sources

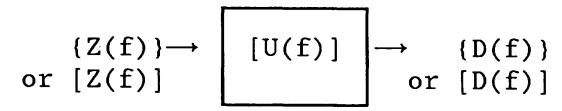

FIGURE 3 Independent noise sources, coupled to give partially coherent sources.

are written as a diagonal matrix, resulting in a square matrix of coupled noise sources, $[D(f)]$,

$$[D(f)] = [U(f)][Z(f)] \tag{14}$$

where $[Z(f)]$ is now a diagonal matrix with the elements of $\{Z(f)\}$ on the diagonal.

FRAME OF TIME DOMAIN DATA

The vector or matrix of $D(f)$ is then transformed to the time domain using an inverse FFT. The result is blocks of data that are one period of periodic data commonly called periodic or pseudo random. If time-domain randomization is not used this will be a vector of time histories, each one block-length long. If time-domain randomization is used the result will be a matrix of time histories each one block-length long. At this point the spectrum is a line spectrum.

Time Domain Randomization

If a circular shift of the frame of periodic random data is performed, a linear phase shift is introduced into all the frequency components. If the amount of the shift is chosen at random, a random linear phase is introduced into the frequency components. Introducing a random circular shift of the blocks of periodic random data, in place of generating new blocks of periodic random data through the steps previously outlined, is known as time-domain randomization.

This method of randomizing the data has been used very effectively in random-vibration control systems.

GENERATION OF CONTINUOUS SAMPLED TIME HISTORIES

The blocks of data are then windowed. The windowing converts the line spectrum into a continuous spectrum. Leakage between adjacent lines will randomize the amplitudes if the random

phase (RP) method was used. The blocks of data are overlapped and added with previously generated blocks of data as outlined in Fig. 4. This further randomizes the amplitudes and phases of the spectrum. If a window and overlap were not used, serious leakage problems would occur because the frames of data would not be continuous at the frame boundaries. Each successive frame is one period of a different periodic function. The goal of the windowing and overlapping is to remove the discontinuities thereby reducing leakage, and hence improve sidelobe performance, randomize the amplitudes and phases, but maintain a stationary output.

The output will be stationary if the sum of the squares of the overlapped windows at all points in time is a constant, and the periodic random data in the frames are independent.

A short discussion of appropriate windows used is discussed below.

Half-Sine or Half-Cosine Windows

The window, $w(i)$, in discrete form, is given by

$$\begin{aligned} w(i) &= \cos(\pi i/L) \quad |i| \leq L/2 \quad \text{or} \\ w(i) &= \sin(\pi i/L) \quad 0 \leq i < L. \end{aligned} \tag{15}$$

This window is used in many vibration-control systems. The characteristics are detailed by Harris [1978]. The window is stationary with an overlap factor of $P \geq 2$. The overlap factor, P, is defined as the ratio of the block length, L, and the number of samples, N, that a block is delayed with respect to the previous block in and overlap-and-add operation. In this article P is restricted to integers.

Potter Windows

Hewlett Packard used a window that is a slight variation of the half-sine window, designed by Potter, in their vibration control system, that gave slightly better performance than the half-sine window. The window is given by

$$w(i) = \left[0.5 - 0.5625 \cos\left(\frac{2\pi i}{L-1}\right) + 0.0625 \cos\left(\frac{6\pi i}{L-1}\right)\right]^{1/2} \quad 0 \leq i < L. \tag{16}$$

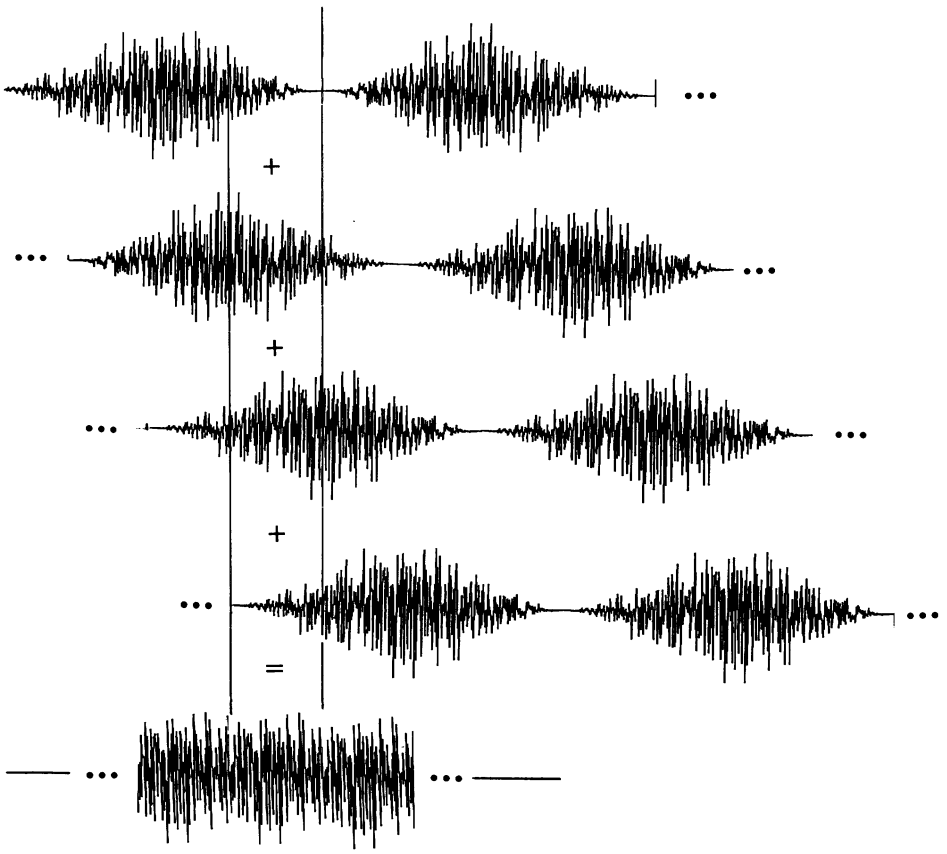

FIGURE 4 Overlap and add operation to generate a continuous output sequence. Overlap factor = 4, Hanning window.

This window was designed for better sidelobe reduction than the cosine window, but still is stationary with an overlap factor of 2. The half-sine and Potter window were used with overlap factor of 2.

A very useful class of windows in the discrete form is given by

$$w(i) = \frac{1}{L}\sum_{k=0}^{K} a_k \cos(2\pi ki/L) \quad |i| \le L/2. \quad (17)$$

For notational convenience the window is centered around $i = 0$. In actual use the window is usually defined for $0 \le i < L$.

Rectangular or Boxcar Windows

For $K = 0$, and $a_0 = 1$, a rectangular window results. This window is not used in this application, but is included for completeness.

Hanning Windows

For $K = 1$, $a_0 = 1/2$, and $a_1 = 1/2$, a Hanning window results. A Hanning window will be stationary with on overlap factor of ≥ 3.

Three- and Four-Term Windows

A set of windows using three and four terms ($K = 2$ or 3) are described by Nuttall [1981]. These windows have very good sidelobe characteristics. For example, Nuttall's minimum three-term window ($a_0 = 0.4243801$, $a_1 = 0.4973406$, $a_2 = 0.0782793$) has a highest sidelobe of -71.48 dB. And the minimum four-term window ($a_0 = 0.3635819$, $a_1 = 0.4891775$, $a_2 = 0.1365995$, $a_3 = 0.0106411$) has a highest sidelobe of -96.17 dB.

Windows of the above type will be stationary if the overlap factor, P, is $\ge 2K$. Some of these windows are stationary for an overlap factor of $P \ge 2K - 1$. The windows are summarized in Table 1. The maximum stationarity error is defined as follows.

Let max = maximum of the sqrt(sum of the squares of the overlapped windows); min = minimum of the sqrt(sum of the squares of the overlapped windows). Then the maximum stationarity error = 1 − (min/max). The peak/rms value (the number in parentheses in Table 1) will be defined later.

As can been seen from Table 1, if exact stationarity is not required, one of the better windows can be used with an overlap factor of 4. An overlap factor of 3 can be used with the Hanning window, but this is not convenient when frames that are powers of 2 long are used with an FFT.

Table 1. Summary of Windows

	Overlap Factor				
Window	2	3	4	6	8
Cosine	0[a]	0	0	0	0
	(2)[b]	(2.3)	(2.6)	(2.8)	(3.6)
Hanning	0.29	0	0	0	0
	(1.4)	(2.0)	(2.3)	(2.8)	(3.3)
Minimum 3-term	0.51	0.12	0.01	0	0
	(1.4)	(1.8)	(2.2)	(2.7)	(3.1)
Minimum 4-term	0.69	0.27	0.05	0	0
	(1.4)	(1.6)	(2.0)	(2.5)	(2.8)

[a] Maximum stationarity error.
[b] Peak/rms one nonzero line.

Required Matrix Formulation for Random-Periodic Coupled Signal Blocks when Time-Domain Randomization Is Used

When time-domain randomization is used, the phase shift caused by the random rotations must be the same for all signals that originate from a common independent noise source. Otherwise, the desired coherence will be destroyed. But signals originating from different noise sources must be randomized independently. By writing the blocks of data in a matrix form, this task can be accomplished. The columns will contain blocks of data from a common noise source, and the rows indicate the blocks that must be summed to form a vector of coupled signals [Smallwood, 1983].

Limits when Using the RP Method and Time-Domain Randomization

The RP method depends on leakage between adjacent lines introduced by the window, and further randomization introduced by the overlap-and-add operations to generate a final output with Gaussian characteristics. The Central Limit Theorem guarantees this will happen if the spectra are broadband. If the spectra are sufficiently narrowband only an approximate Gaussian distribution will result.

Consider the limiting case where the spectral density is zero except for one spectral line. In this case, if the RP method is used, the frame of

data will be a sine with a fixed amplitude and a randomly varying phase. Consider a peak amplitude of one for the sine, a cosine window and an overlap factor of 2. The output will be stationary, with an rms value of 0.707. The largest peak that could be observed is for the two sine waves in the overlapped frames to be in phase at the maximum point of overlap. This will give a peak of (1, the amplitude of the sine) × (2, the number of sines) × (0.707, the magnitude of the window at the point of maximum overlap) = 1.414. The ratio of the maximum peak that can be observed relative to the rms is then 1.414/0.707 = 2. Thus the largest peak that could be observed is twice the rms, not a particularly good result. This value is tabulated for other windows and overlap factors in Table 1. As can be seen the peak/rms values improve as the overlap factor is increased.

If two lines are nonzero and equal, and all other lines are zero, the peak increases by a factor of 2, and the rms increases by a factor sqrt(2). Therefore, the peak to rms increases by a factor of sqrt(2) or 1.41, and the values in Table 1 must be multiplied by this factor. For this case of two lines, the cosine window and an overlap factor of 2 gives a peak/rms ratio of 2.82, a better result. For m lines the increase in the ratio of peak to rms over one line is the square root of m. Thus a cosine window, with an overlap factor of 2, and four nonzero lines gives a peak/rms value of 4.0. Ten lines will give possible peaks of 6.3 times the rms. Possible peaks of 4 times the rms will give a distribution that is very nearly Gaussian. This results in the recommendation that if the RP method is used to generate blocks of independent noise, at least four lines in the frequency domain should be of the same order of magnitude. This restriction does not apply if the WN method is used to generate the independent blocks of noise.

ALTERNATE METHOD FOR GENERATION OF CONTINUOUS SAMPLED TIME HISTORIES

Continuous time histories can also be generated with a generalization of the overlap-add method of high speed convolution outlined by Gold and Rader [1969].

1. The process is started by taking the inverse FFT of the coupling matrix $[U(f)]$, giving the matrix $[u(t)]$. Each element, $u(t)$, in the resulting matrix, $[u(t)]$, is then a time history N points long. These terms can be thought of as the impulse responses of the coupling terms. The impulse responses are not necessarily realizable. For example, the diagonal terms have a real FFT and are therefore even.
2. The elements, $u(t)$, are circular rotated in the time domain $N/2$ points. This puts time zero in the middle of the frame. This will introduce a time lag of $N/2$ points in the output, but this is of no practical significance.
3. The elements are then windowed with your favorite window to reduce any leakage caused by truncation of the impulse response.
4. The elements, $u(t)$, are then extended with N zeros, giving blocks $2N$ long. This is done to prevent circular convolution errors later.
5. The elements, $u(t)$, are then transformed back into the frequency domain with an FFT to give an extended version of $[U(f)]$. Call this matrix $[U_e(f)]$.
6. Each element in a vector composed of frames of independent random data in the time domain, $\{z(t)\}$, is then extended with N zeros, and transformed into the frequency domain with an FFT giving, $\{Z_e(f)\}$.
7. The matrix vector product, $[U_e(f)]\{Z_e(f)\}$ is then formed, and transformed into the time domain, giving $\{d_e(t)\}$. Note that the sequences for each of the elements in $\{d_e(t)\}$ are $2N$ points long.
8. The first half of each element of $\{d_e(T)\}$ is summed with the last half of the sequences formed from the previous loop. Zeros are used to initialize the process. This gives an output sequence of N points for each element of the output vector. The last half is saved for the next loop.
9. The process is repeated by looping back to 6., with a new sequence $\{z(t)\}$, for as long as desired. Note that the elements of the coupling matrix, $[u_e(t)]$, and hence $[U_e(f)]$, need to be formed only once.

The result is a continuous output sequence with no discontinuities at the frame boundaries, and no leakage caused by the circular convolution. Frequency resolution and dynamic range are limited only by the size of the frame, N, and

the window used to shape the impulse response functions. The sequences will be stationary, independent of the window used, and always require a 50% overlap (an overlap factor of 2). Time domain randomization can be used to generate the new sequences, $\{z(t)\}$, but is not recommended.

CONCLUSIONS

The method provides a very robust stable method for producing partially coherent sequences. Sequences corresponding to any realizable cross-spectral density matrix can be easily generated. No complicated filter designs are required. Frequency resolution is limited only by the sample rate of the data and the frame size selected. In principle, any desired frequency resolution can be achieved. The selection of the frame size and window give excellent spectral characteristics, which are fully characterized in shape and dynamic range. The selection of the window and the overlap parameter can guarantee a stationary output. The method is quite efficient. Sequences of 2048 points for three partially coherent signals have been computed in significantly less than 0.2 seconds using an array processor.

REFERENCES

Bendat, J. S., and Piersol, A. G., 1986, *Random Data, Analysis and Measurement Procedures,* 2nd ed., Wiley, New York.

Dodds, C. J., and Robson, J. D., "Partial coherence in Multivariate Random Processes," 1975, *Journal of Sound and Vibration,* Vol. 42, pp. 243–249.

Gold, and Rader, 1969, *Digital Processing of Signals,* Chap. 7, McGraw–Hill, New York.

Harris, F. J., January 1978, "On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform," *Proceedings of the IEEE,* Vol. 66.

Lawson, C. L., and Hanson, R. J., 1974, *Solving Least Squares Problems,* Prentice–Hall, Englewood, NJ.

MATLAB Users Guide, Mathworks, S. Natick, MA.

Nuttall, A. H., February 1981, "Some Windows with Very Good Sidelobe Behavior," *IEEE Transactions on Acoustics, Speech, and Signal Processing,* Vol. ASSP-29.

Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T., 1986, *Numerical Recipes, The Art of Scientific Computing,* Cambridge, New York.

Smallwood, D. O., April 1978, "Multiple Shaker Random Control with Cross Coupling," Institute of Environmental Sciences 1978 Proceedings, pp. 341–347.

Smallwood, D. O., September 1983, "A Random Vibration Control System for Testing a Single Test Item with a Multiple Input Control System," *1982 SAE Transactions.*

Spanos, P. D., and Mignolet, M. P., November 1989, "ARMA Monte Carlo Simulation in Probabilistic Structural Analysis," *The Shock and Vibration Digest,* Vol. 21, pp. 3–14.

Tebbs, J. D., and Hunter, N. F., April 1974, "Digital Control of Random Vibration Tests Using a Sigma V Computer," Institute of Environmental Sciences 1974 Proceedings, pp. 36–43.