A JMP Script to Enhance Spectral Density Analysis in JMP's Time Series Analysis Platform

Presented at the

JMP Discovery Summit
September 2014
Cary, NC

Jerry Fish
President, Experistats LLC
experistats@gmail.com
859-433-6231

Contents

Abstract

JMP's Time Series Analysis platform contains an optional Spectral Density Analysis for studying repeating signals. Spectral Analysis is often used in engineering, such as electrical, vibration, and acoustic signal analysis.

The author will demonstrate a JMP Script which adds important functionality to the existing Time Series/Spectral Density offering, including:

- A brief review of Fourier Series and Frequency Analysis
- Spectrum display
- Autospectrum display
- Waterfall plots
- Subsample Averaging (to help clean up non-repeating signal noise)
- Overlapping (useful if frequencies are varying with time)
- Windowing (to clean up periodicity that does not have an integer number of cycles per frame)
- Transfer Functions (allowing frequency domain analysis of input/output effects)
- Damping Ratio estimates (based on peak widths)

Introduction

Opportunities to investigate the behavior of dynamic systems abound. The investigation methods used to analyze these systems range from those that help to reveal the fundamental behavior of the system, to those that allow forecasting and prediction of future behavior. The only requirements for these analyses are that 1) data exists to be analyzed, and 2) the data points are collected at a constant interval.

JMP's Time Series Analysis platform offers a number of techniques to develop models to forecast and/or predict future behavior of dynamic systems. Among these techniques is one called Spectral Density. Spectral Density is used to determine periodicity in a time series signal. This is important when fitting Seasonal ARIMA models, but it can also be important outside of ARIMA, such as for understanding the fundamental dynamics of a given system.

While JMP's Spectral Density tool is useful in pulling out simple periodicity in a signal, the general field of Spectral Analysis can be used in a wide variety of ways. This paper documents a new JMP script that has been developed to exploit some of these methods.

Spectral Analysis Basics

Time Domain Analysis of Periodic Signals

Spectral analysis can be accomplished in many ways. If the signal being studied has relatively little noise, and only one or two clear frequencies, spectral analysis might be as simple as looking at the signal vs. amplitude plot and measuring distance between adjacent peaks to give a direct measurement of the period of a cycle, as in Figure 1. (The frequency is obviously the reciprocal of the period). Accuracy can be improved by averaging over multiple cycles.

When noise begins to contaminate the signal, the characteristic frequencies may become more difficult to discern (Figure 2).

Instead of analyzing the signal in Figure 2 in the time domain, it is might be wise to study the signal in the frequency domain. This is done by transforming the data using a Fourier Series.

The General Fourier Series

The Fourier Series is named after the French mathematician Jean Baptiste Joseph Fourier (1768-1830) (1). The Fourier Series allows a time signal to be decomposed into a sum of sine waves. If enough sine waves are included in the sum, the initial time signal can be reproduced exactly.

Mathematically expressed, if we have a function f(t), it can be expressed as follows (2):

$$f(t) = A + \sum_{n=1}^{\infty} B_n \sin(nt - \varphi_n) \quad \text{Equation 1}$$

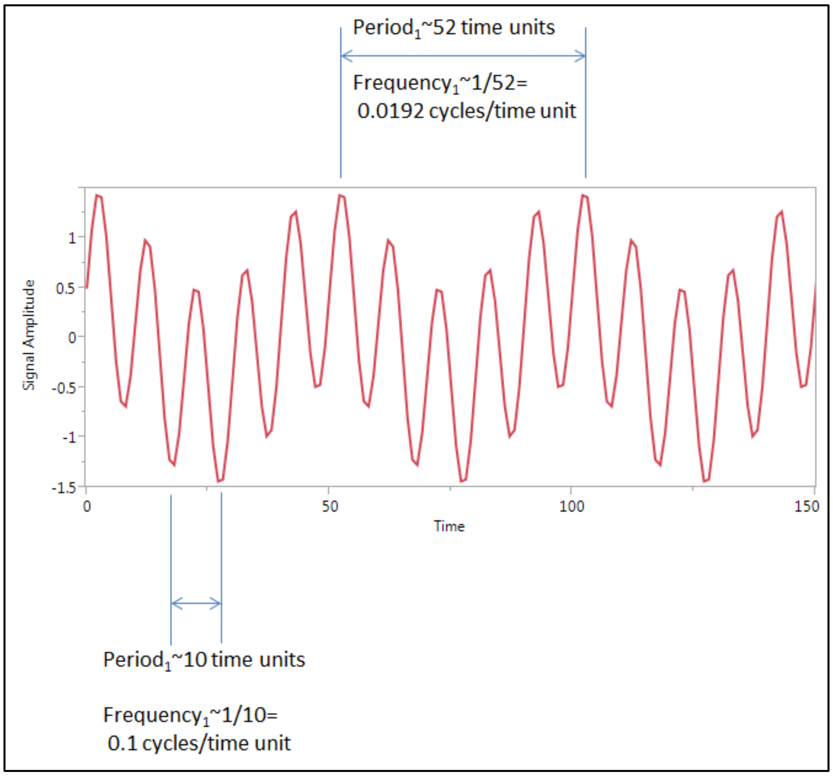

Figure 1: Determining frequencies from time plots.

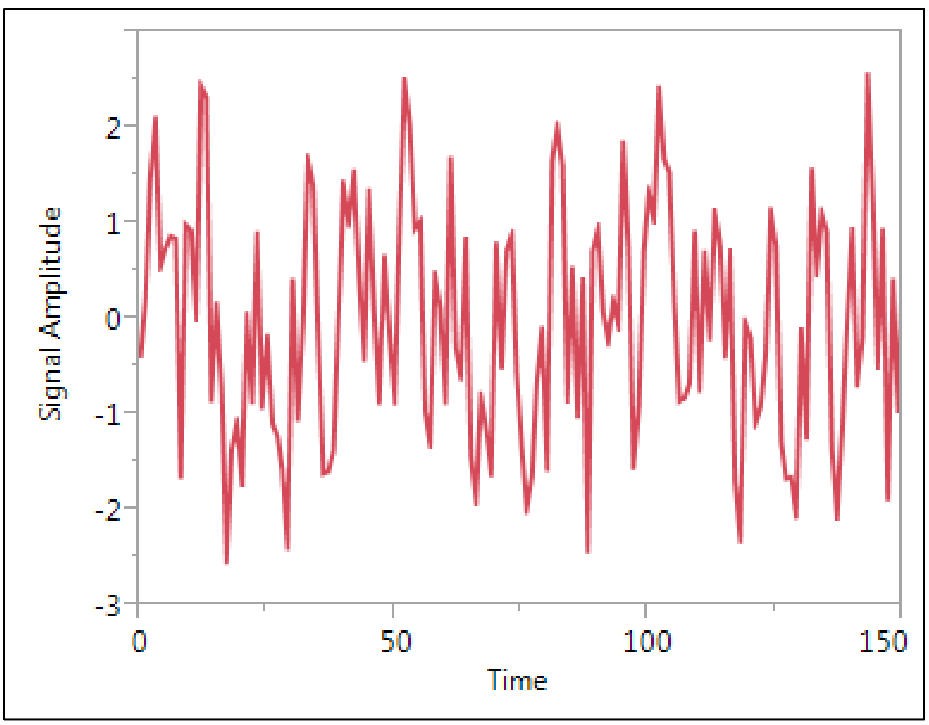

Figure 2: Signal contaminated with noise. Note that frequencies are not as clearly discerned as in Figure 1.

where:

A is a constant (the average of the time signal)

n is an integer

B_n is a constant associated with each sinusoidal frequency

t is time

ϕ_n is a phase angle

An assumption of the Fourier Series applies if the signal is of finite length. The finite length original signal assumed to have repeated exactly in the past, and will continue to repeat exactly in the future. This will become important in later sections.

The Continuous Frequency Spectrum

The Frequency Spectrum is a plot of all of the A and B_i coefficients from Equation 2 against their respective frequencies.

As an example, let us assume that our time signal is composed of a single sine wave of the form:

$$f(t) = 2sin(5\pi t - \pi/4) \qquad \text{Equation 2}$$

where t is in seconds. The time domain plot of this signal is shown in Figure 3.

In this case (comparing to Equation 2), A=0, and there is only one value of B (B=2). The only phase value is $\phi=\pi/4$. The frequency is 5π radians per second, or 2.5 cycles/second. A continuous frequency spectrum would then look like Figure 4.

Figure 4 indicates that there is a single frequency in the original time series, with frequency=2.5 Hz, and 0-pk Amplitude of 2 units.

In addition, the phase angles ϕ_i can also be plotted against frequency (Figure 5). We know that the frequency is 2.5 Hz, so we read off of the phase plot that the phase angle is 45 degrees. (Note that the amplitude at all other frequencies is zero, so the phase is moot. By definition, we give those phase angles a value of 0 degrees.)

If we have a discretely sampled signal (both a fixed time interval and a fixed length of time for the entire sample), then the frequency spectrum also becomes discrete. The next section begins to explore those aspects.

The Discrete Time Series

If a signal is sampled at even time intervals[1], and a known number of samples are collected, then the series is called a Discrete Time Series. (3) The sampling rate and the sample size are both chosen by the user, based on the dynamic content of the signal and other considerations (Figure 6).

[1] "Time" is only one aspect of sampling intervals. "Distance" can also be used as an interval. For example, certain printing processes have light/dark artifacts on the printed output that are periodically spaced down the page. One could create a "distance series", where the darkness of the printed image was measured in even distance increments.

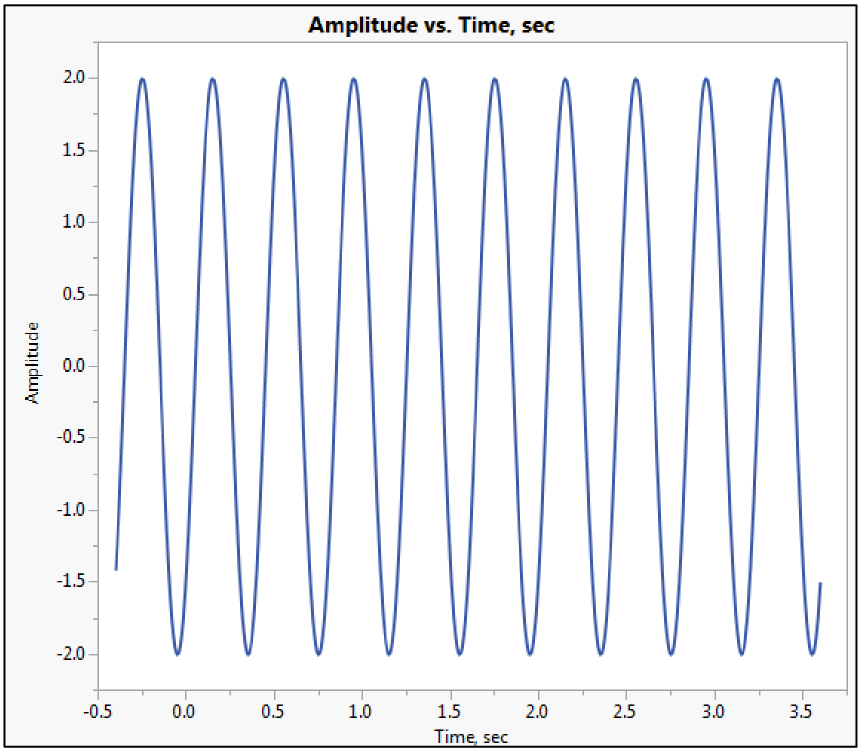

Figure 3: Continuous Time Series Plot of 5 Hz Sine Wave with 45 Degree Phase Shift

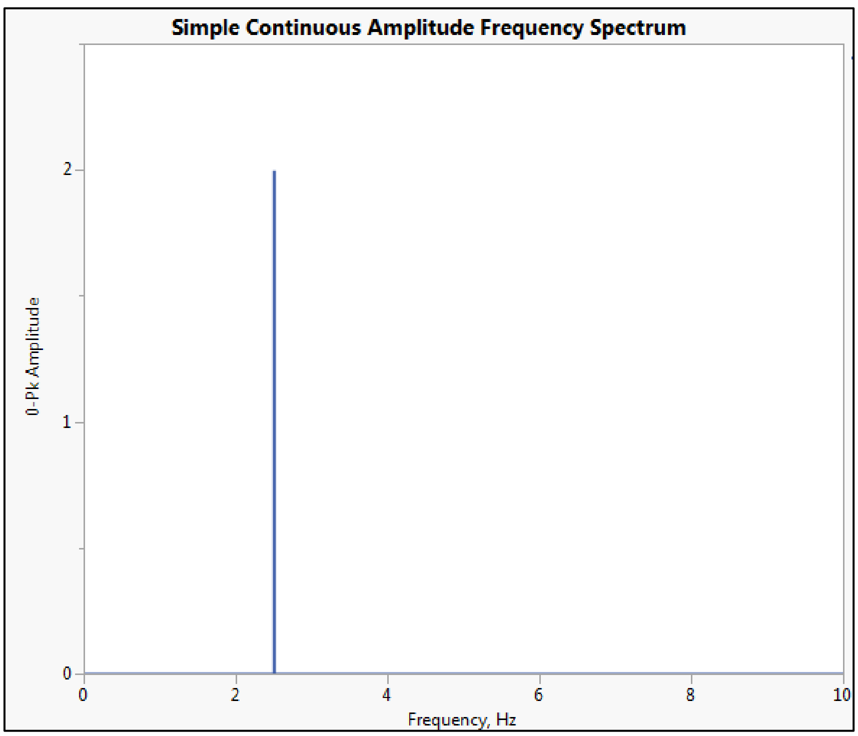

Figure 4: Continuous Frequency Spectrum of Simple Sine Wave

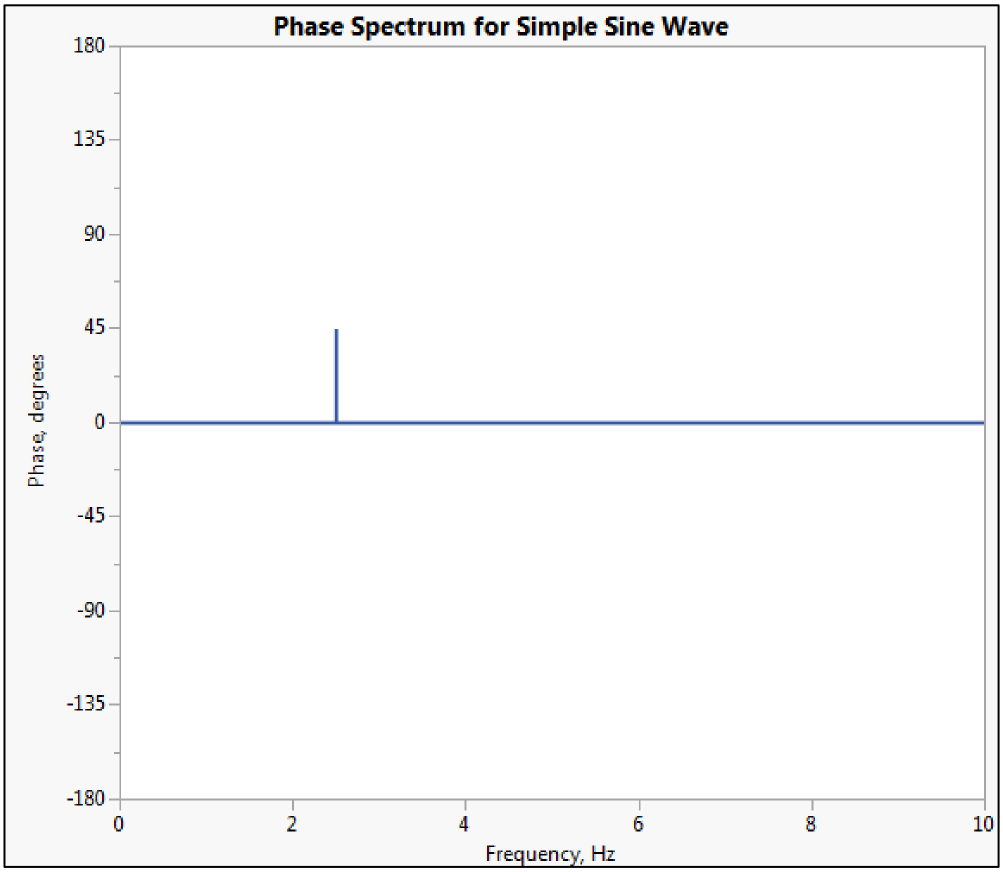

Figure 5: Phase Angle vs. Frequency Plot for Continuous Frequency Spectrum

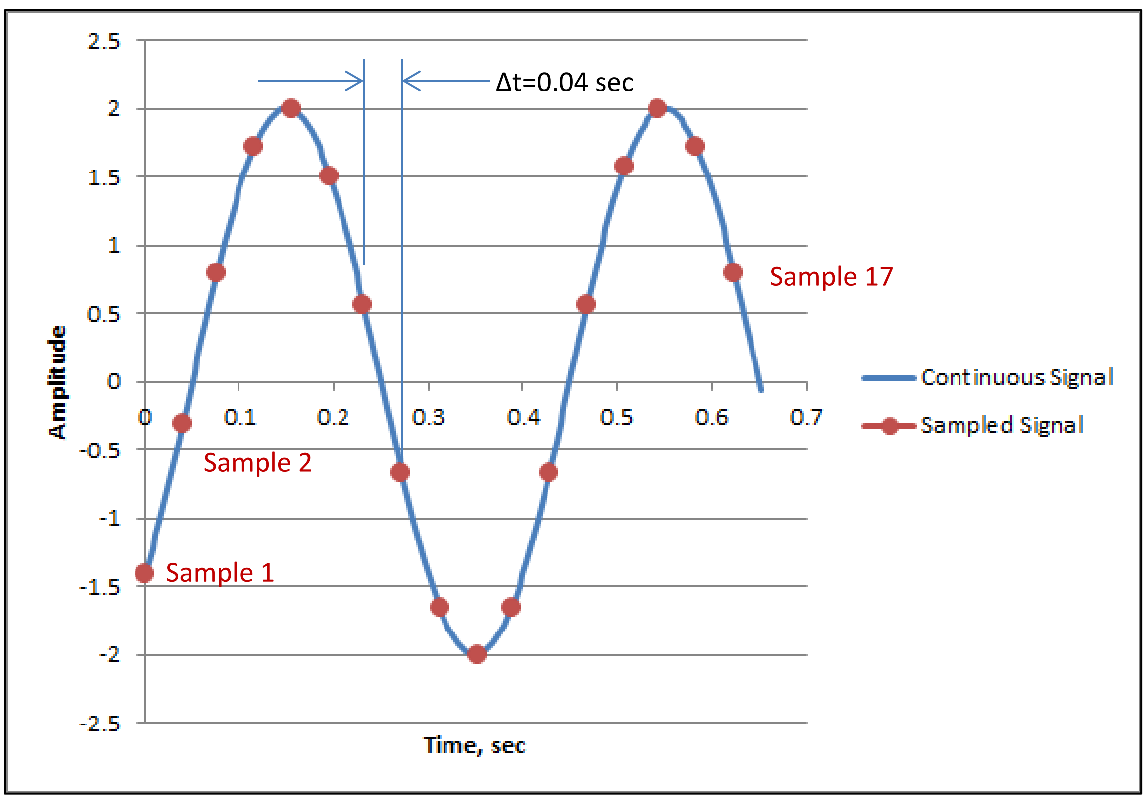

Figure 6: A continuous time signal, sampled at even increments of 0.04 seconds. A total of 17 samples are shown.

The Discrete Frequency Spectrum

In a previous section, we discussed a time signal that was continuous, which resulted in a frequency spectrum that was also continuous. For a discretely-sampled time series (Figure 6), a Discrete Fourier Series may be fit to the discrete time series data, resulting in a discrete frequency spectrum.

The Discrete Fourier Series (4)

If the time signal consists of N data points sampled at a regular interval Δt, the Fourier Series can be rewritten as:

$$f(t) = A + \sum_{i=1}^{N/2} B_i sin\left(\frac{2\pi i}{N\Delta t}t - \varphi_i\right) = A + \sum_{i=1}^{N/2} B_i sin(\omega_i t - \varphi_i) \qquad \text{Equation 3}$$

where:

A is a constant (the average of the time series)

i is an integer

N is the number of time samples in the time series

Δt is the sampling interval

t is time

$\omega_i = \frac{2\pi i}{N\Delta t}$ = the i^{th} sinusoidal frequency (radians/time unit)

B_i is a constant associated with each sinusoidal frequency

ϕ_i is the phase angle (radians) associated with each sinusoidal frequency

The Discrete Frequency Spectrum

The frequency spectrum that results from the discrete fourier series now has a finite number of frequencies for which amplitude and phase data exists. For the above sine wave, the resulting spectrum might look like that shown in Figure 7. Note that there are discrete frequency points in this spectrum. In other words, there are only certain frequencies that are fit to the original time series.

In the same way, the discrete phase plot also consists of individual points (Figure 8).

How Discrete Time Series Sampling Affects the Discrete Frequency Spectrum

These two parameters (the time interval between samples Δt, and the total number of sampled points N) determine the general parameters that describe the frequency spectrum.

Maximum Frequency (a.k.a. Nyquist Frequency)

Considering the signal shown in Figure9, there are enough samples that we get good fidelity of reproduction of the original signal, even if we connect the points with straight lines as shown in Figure 9.

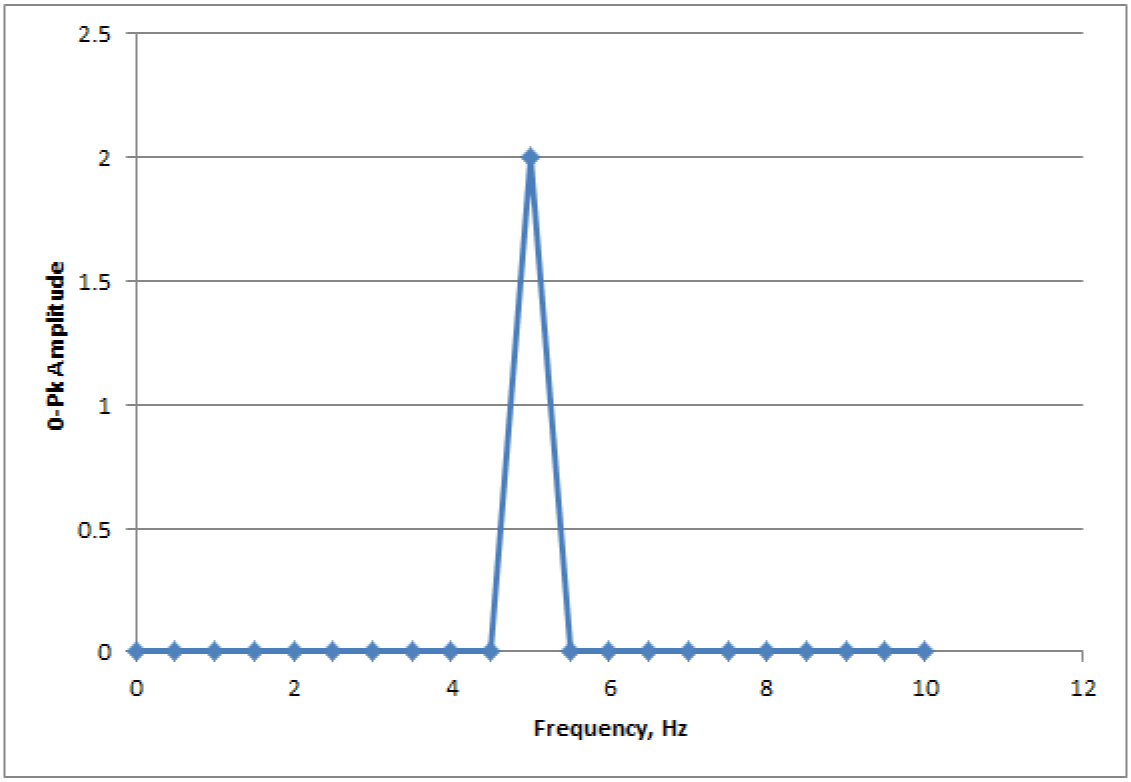

Figure 7: Example of a Discrete Frequency Spectrum for a Simple Sine Wave

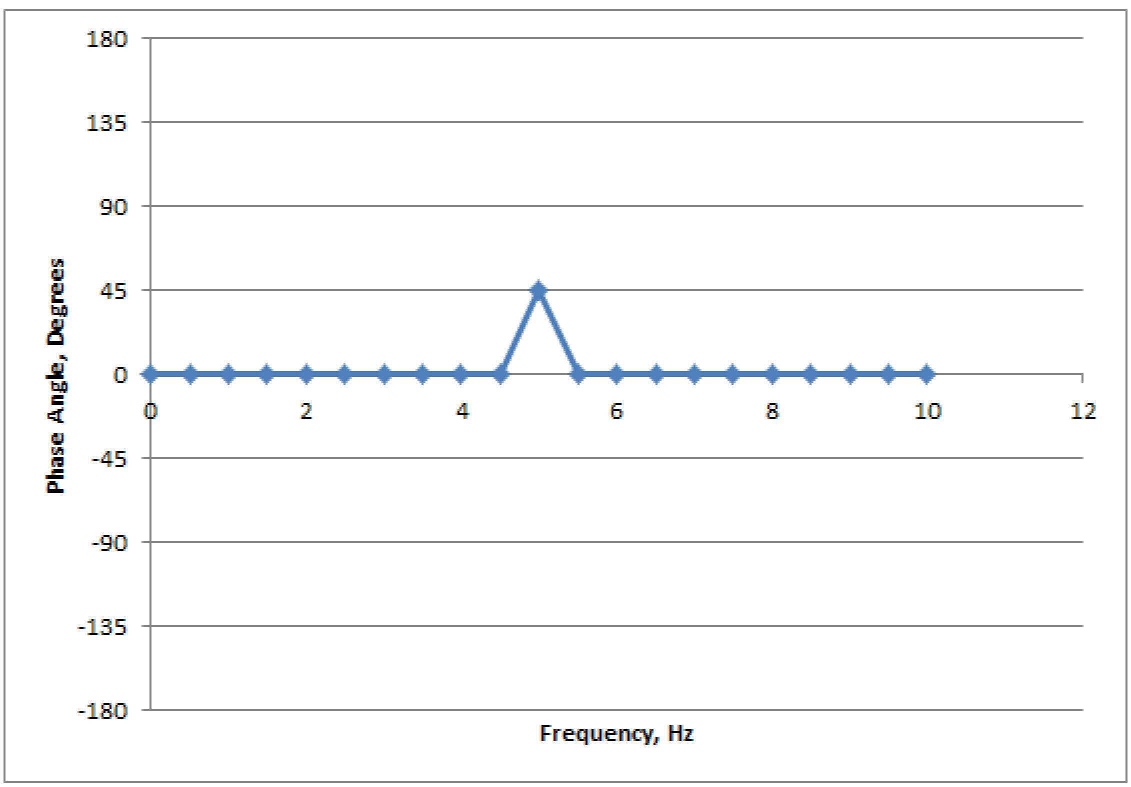

Figure 8: Example of a Discrete Frequency Phase Plot

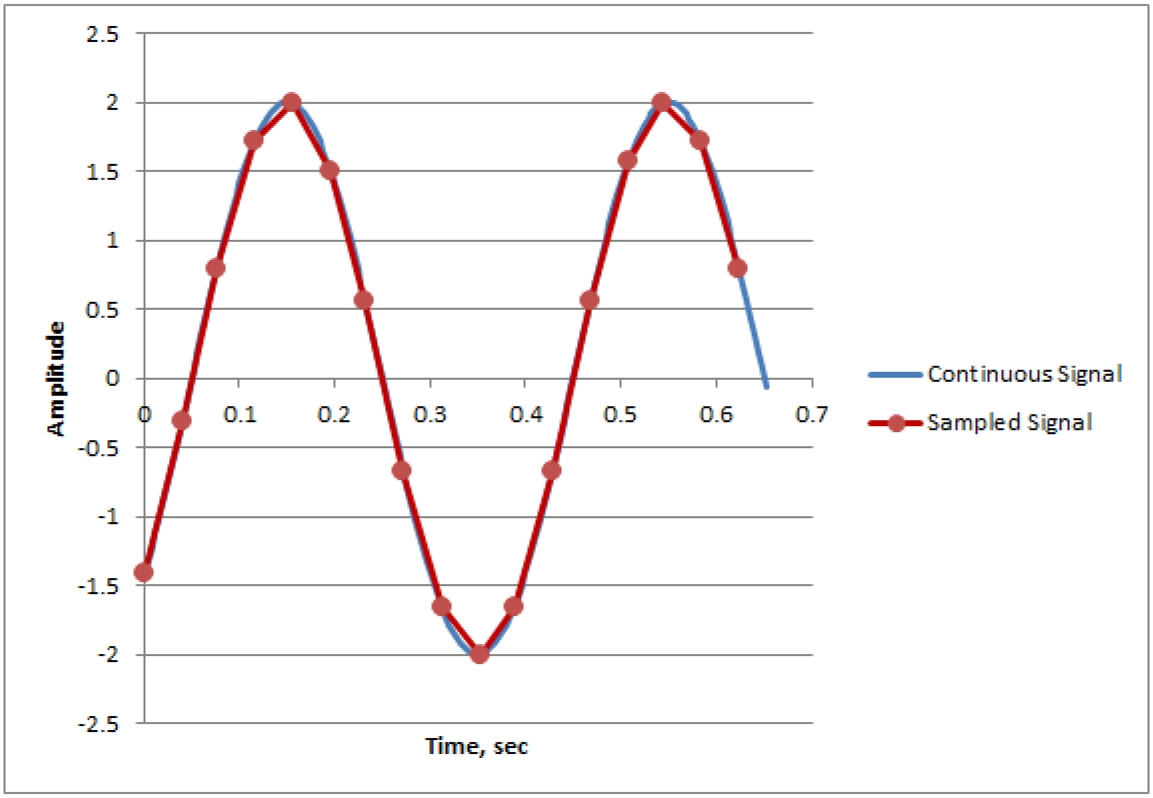

Figure 9: Waveform Sampled at Many Samples per Cycle

If we sample the same sinusoidal signal over a longer sampling interval, we begin to lose some fidelity of the original signal, but the periodic nature of the signal is still intact (Figure 10). Even though the original signal frequency can still be seen in the sampled signal, now we begin to see a clear impact on the amplitude of the sampled signal as compared to the original signal.

If we then choose a sampling interval that coincides with exactly half of the period of the sine wave, we might get the result in Figure 11.

If we continue to increase the sampling interval to, say, 0.75 of the original signal sinusoidal period, we get the result shown in Figure 12. Now not only is the amplitude of the sampled signal affected, but the frequency has also changed. This phenomenon is known as "aliasing". (5)

It turns out that the maximum frequency that can be accurately detected in a sampled signal is 0.5 times the sampling frequency. This maximum frequency is known as the Nyquist Frequency (6), after Harry Nyquist (1889-1976), an electrical engineer specializing in communications technology. (7) The mathematical relationship between this maximum frequency and the sampling interval Δt is:

$$f_{max} = f_{Nyquist} = \frac{1}{2\Delta t} \qquad \text{Equation 4}$$

In practice, a low-pass analog or digital filter is often used to filter out any frequencies that are higher than the Nyquist frequency. Filtering (as well as aliasing and folding) are beyond the scope of this paper.

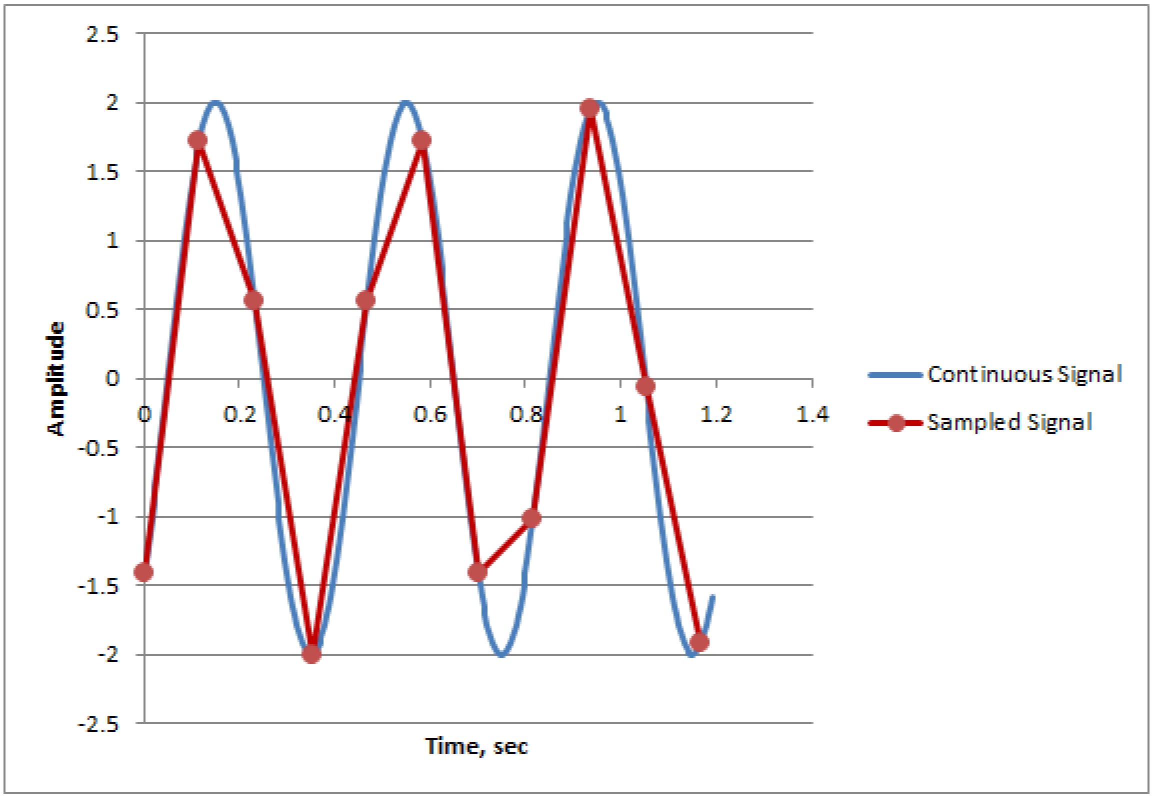

Figure 10: Signal sampled at about 3 samples per cycle

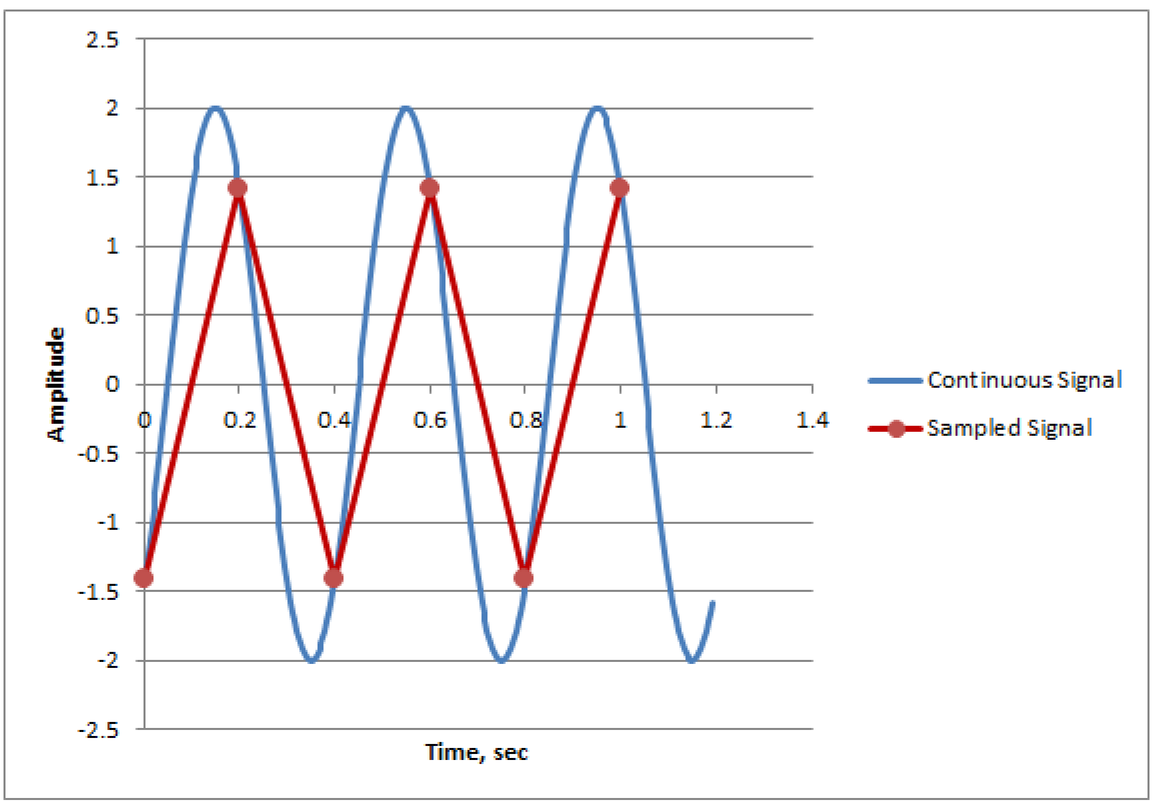

Figure 11: Signal Sampled at Exactly 2 Samples per Cycle

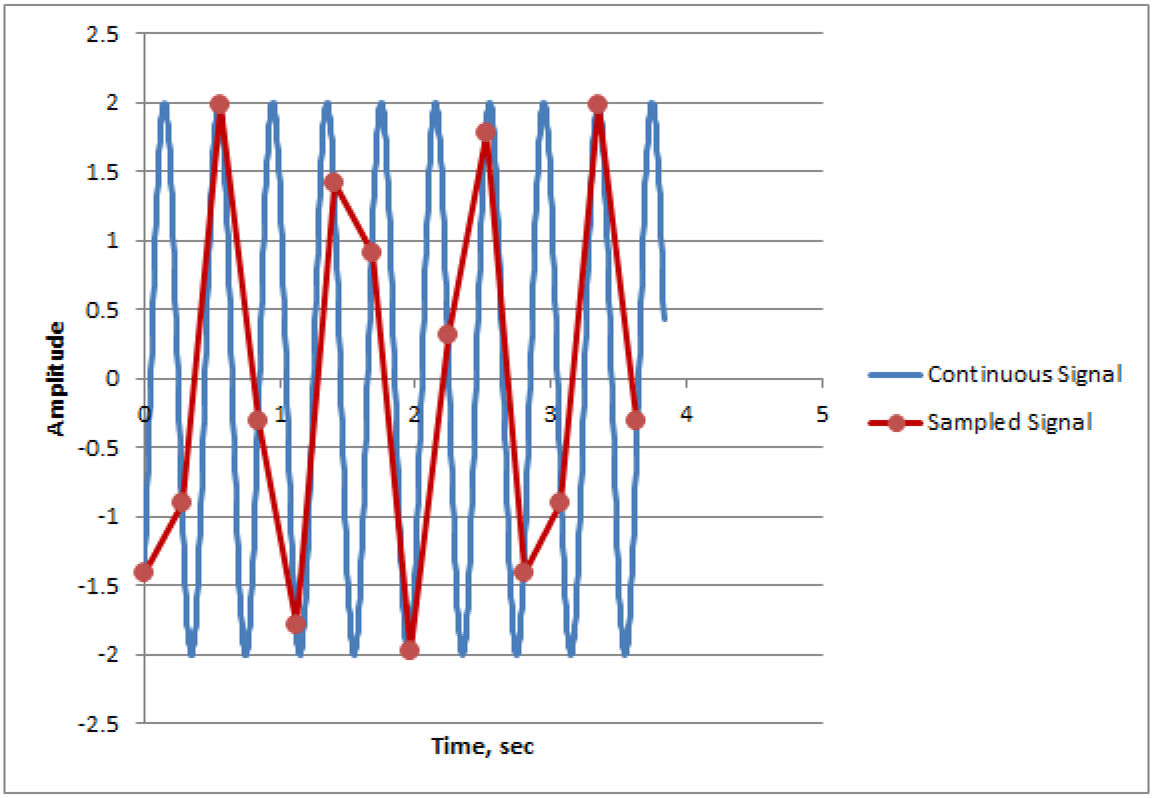

Figure 12: Signal Sampled at 0.75 Samples per Cycle. (Note that signal frequency is not represented accurately in sampled series.)

Number of Spectral Lines

Based on Discrete Fourier Series theory, it happens that the total number of frequencies that comprise the discrete frequency spectrum is exactly half the number of samples collected, or:

$$N_{lines} = \frac{N}{2} \qquad \text{Equation 5}$$

If N is odd, we round down to the nearest even integer to find Nlines.

Frequency Resolution

The frequency resolution of the spectrum is the spacing of the spectral lines in the spectrum. Using Equations XX and YY above, we have:

$$f_{res} = \frac{f_{max}}{N_{lines}} = \frac{1/(2\Delta t)}{N/2} = \frac{1}{N\Delta t} \qquad \text{Equation 6}$$

All of the above relationships are summarized in the following table:

Table 1: Relationships Between Time Sampling Parameters and Discrete Frequency Spectrum Parameters

		Affected Parameters			
		sampfreq	Nyquist Frequency	Number of Spectral Lines	Spectral Resolution
Sampling Parameters	Δt	Δt is inversely proportional to sampling frequency: $f_s=1/\Delta t$	Δt sets the maximum frequency that can be analyzed: $fmax=f_s/2=1/(2\Delta t)$		Δt is inversely proportional to spectral resolution. If you want a smaller frequency increment between spectral lines, increase Δt: $\Delta f=f_{max}/N_{lines} = 1/(N\Delta t)$
	N			N is twice the number of lines in the spectrum: $N_{lines} = N/2$	N is inversely proportional to spectral resolution. If you want a smaller frequency increment between spectral lines, increase N: $\Delta f=f_{max}/N_{lines} = 1/(N\Delta t)$

What's an "FFT"?

FFT stands for Fast Fourier Transform. (8) It manipulates a discrete time series to produce a discrete frequency spectrum, just like described above. However, an FFT is a special numeric algorithm that performs the math in a very fast fashion, by taking advantage of a specific constraint on the discrete time series. That constraint requires that the number of samples N in the discrete time series be a power of 2 (e.g. 32, 64, 128, …). Lengths of time series other than these are not allowed. This constraint also limits the number of spectral lines in the frequency spectrum to a power of 2. The FFT is generally employed by most frequency analyzers on the market today.

JMP chooses to use a different algorithm in computing their Spectral Density function in the Time Series Analysis platform. (9) The function used by JMP has no limit on the number of samples in the time series, so there is no constraint on the number of spectral lines that can exist in the resulting frequency spectrum (other than Nlines = N/2, as described previously). This is a more flexible algorithm, though it is computationally slower than the FFT.

The JMP Script described below generally conforms to the JMP-style algorithm.

The Frequency Analysis Script for JMP

The Frequency Analysis Script was written to take advantage of the theory outlined above, allowing a more robust analysis of periodic signals. The following sections describe the operation, functions, and outputs of the script.

Input Data Description

As with most JMP operations, the process begins by the user building a data table. For the Frequency Analysis Script, this means building up to 3 columns of data. The columns may be in any order. These columns include Response Data (required), Time Data (optional), and Forcing Function Data (optional).

Response data

Response data is the measured value of the signal of interest, sampled at discrete time increments. The data is entered sequentially, beginning in Row 1 and continuing through the last sample at Row N. This column is a requirement.

Time data

Time data is an optional column. If entered, it contains the values of time where the response data was collected. It also has N rows.

Note that if the Time Data column is omitted, the user has an option of specifying the time data (in terms of sampling increment or sampling frequency) later in the Frequency Analysis Script inputs.

Forcing Function data

The Forcing Function is also an optional data column. If entered, transfer functions will be calculated between the response data and the forcing function data. Again, this column contains N rows.

Units

The user can optionally specify Units for Response, Time, and Forcing Function data. This is a normal data table function within JMP. Simply right-click on the column header, select Column Properties, and then select Units. Type in the units in the appropriate box. (See Figure 13.) If Units are not specified, the Script assigns the default values shown in Table 2 to each of the columns.

Starting the Frequency Analysis Script

The script can be installed as an Add-In like any other script, or it can be opened and run from the Script Editing window. Refer to the JMP Scripting Guide for further instruction.

Selecting Columns Window

The Frequency Analysis Script uses a dialog that is very similar to other dialogs that start various JMP platforms. As shown in Figure 14, the user can specify the three inputs described above: Response series, Forcing Function series, and the Time series. Cast the columns into their roles, then click on OK to proceed.

Spectral Analysis Setup Window

Once the input columns have been specified, a new window appears allowing the user to set up the frequency analysis parameters for this analysis. Figure 15 shows an example of this window, the aspects of which will be described below.

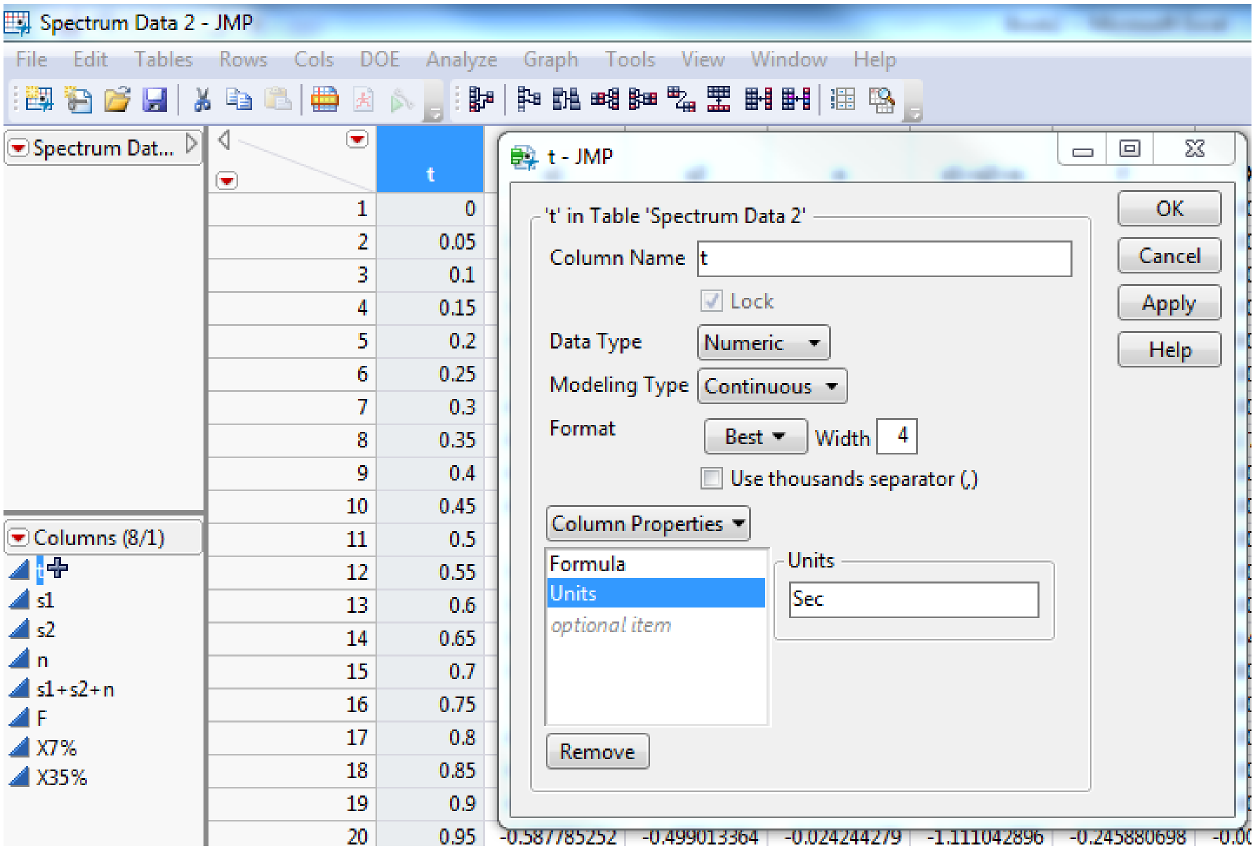

Figure 13: Specifying Units for a Column Within a JMP Data Table

Column	Default Units
Time	“Time Units”
Response	“Response Units”
Forcing Function	“Forcing Function Units”

Table 2: Default Data Units Assumed by Frequency Analysis Script

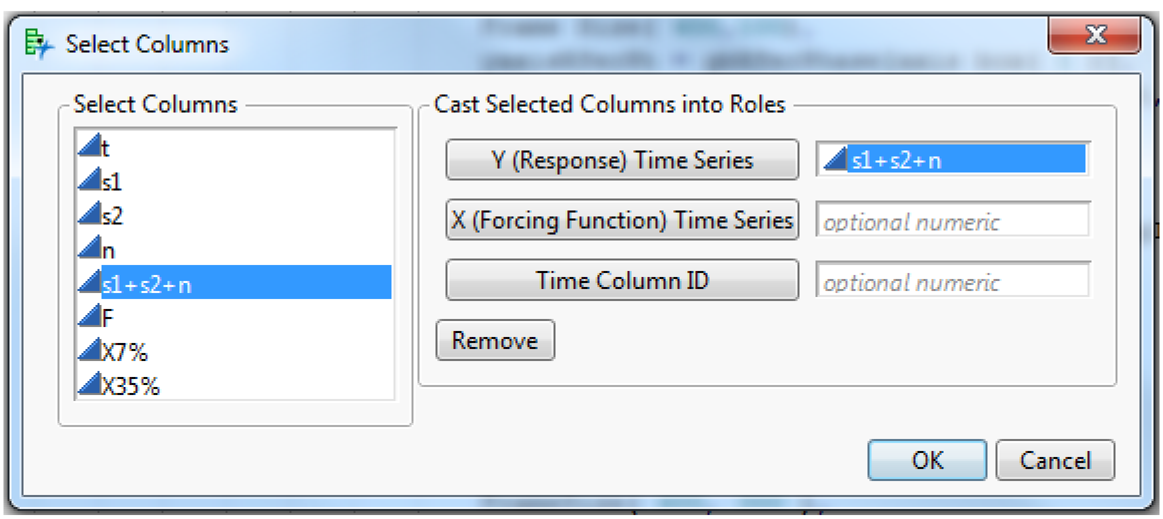

Figure 14: Specifying Columns for the Frequency Analysis Script. In this example, only a Response column is specified. The other columns are optional.

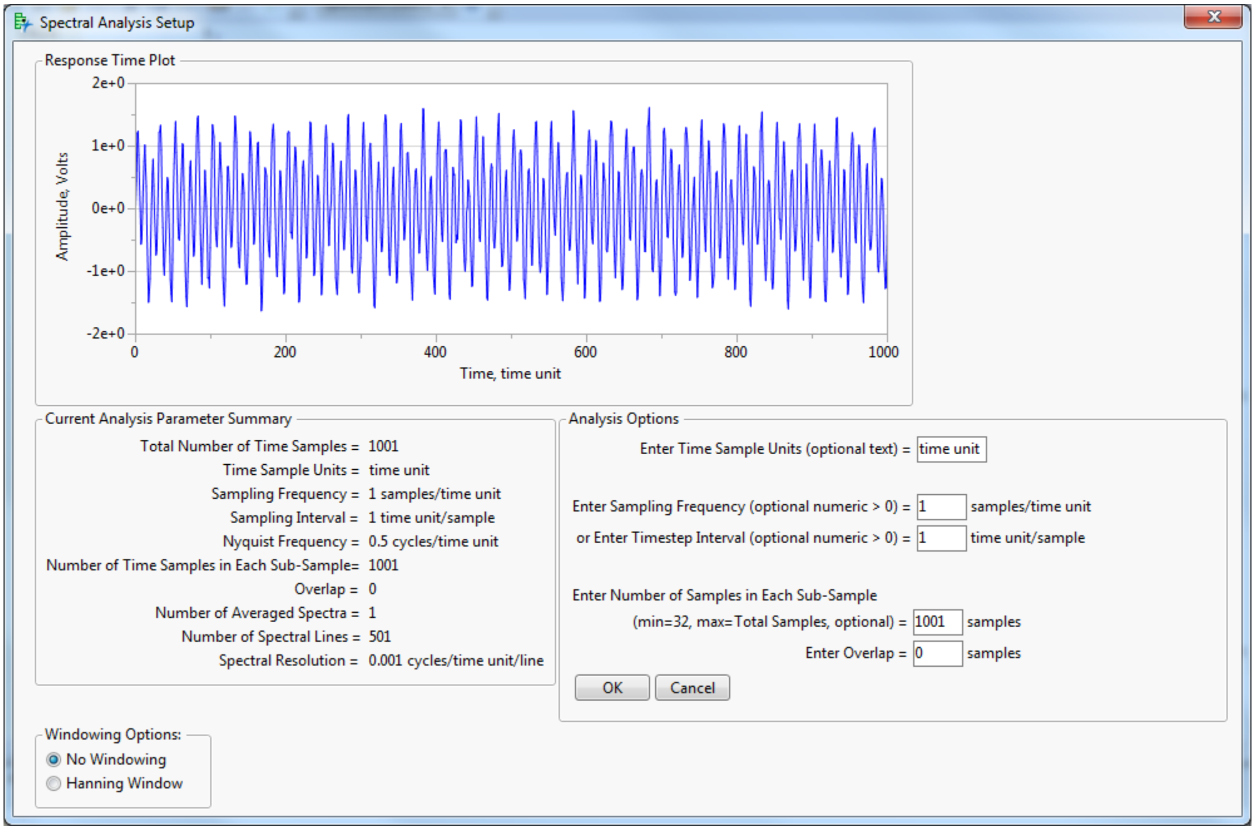

Figure 15: The Spectral Analysis Setup Window

Response Time Plot

At the top of the Spectral Analysis Setup is the "Response Time Plot". This shows the sampled response signal, plotted in the time domain.

Note that the Amplitude (y-axis) has units of "Volts" in Figure 15, since Volts were specified as units for the response column.

Also note that the Time axis currently contains "time unit" as the units. The Time column was not specified in this example, so at this point the script does not know what the time increments or units are.

Current Analysis Parameter Summary Panel

To the left and below the Response Time Plot is the Current Analysis Parameter Summary Panel. This panel summarizes all of the parameters for the upcoming analysis, based on current inputs. All of these parameters have been described previously, with the exception of "Number of Time Samples in Each Sub-Sample", and "Overlap". These will be discussed later in this report.

Analysis Options Panel

The user has the ability to change analysis parameters in the Analysis Options panel. The first three options are described below. The last two options ("Number of Time Samples in Each Sub-Sample", and "Overlap") will be described later in this report.

Enter Time Sample Units (optional text)
The user has the option to change time units in this text box. Note that this option only appears if the time column was not specified, or if time column was specified but time units were not attached to the column.

Once this text box is populated, all references to time units are immediately updated in the Response Time Plot, the Current Analysis Parameter Summary Panel, and the Analysis Options Panel.

Enter Sampling Frequency (optional numeric >0)
The sampling frequency can be entered here. (Alternatively, the Timestep Interval can be entered immediately below this option.)   This is the number of samples collected per "time unit". Note that this option only appears if the time column was not specified. Also note that the sampling frequency must be a number greater than zero.

Once this numeric box is populated, all references regarding sampling frequency (or time intervals) are immediately updated in the Response Time Plot, the Current Analysis Parameter Summary Panel, and the Analysis Options Panel.

Enter Timestep Interval (optional numeric >0)
The time between samples can be entered here. (Alternatively, the Sampling Frequency can be entered immediately above this option.)   Note that this option only appears if the time column was not specified. Also note that the Timestep Interval must be a number greater than zero.

Once this numeric box is populated, all references regarding time intervals (or sampling frequency) are immediately updated in the Response Time Plot, the Current Analysis Parameter Summary Panel, and the Analysis Options Panel.

Windowing Options
Windowing options are used to help "clean up" the response spectrum, the forcing function spectrum, and the transfer function. Windowing will be described later in this report.

Spectral Analysis Plots Window
For the purposes of explanation in this section, a simple sine wave is used as the example. The sine wave consists of:

- Amplitude=0.5
- Units=none (default is "Response Units")
- Frequency=0.8 Hz
- Phase Angle = 0 degrees
- Sampling Frequency = 20 samples/sec (or sampling interval = 0.05 sec)
- Nsamples=1001 (or total length of sample=(1001-1)*0.05=50 seconds)

After clicking "OK" in the Spectral Analysis Setup window, a progress bar window appears as calculations are performed. Once complete, the progress window disappears, and the Spectral Analysis Plots window appears (see Figure16).

Since a Forcing Function was not selected, there are no Forcing Function or Transfer Function spectra to be plotted. Hence, the only options are to show the Response Spectrum and the Response Autospectrum.

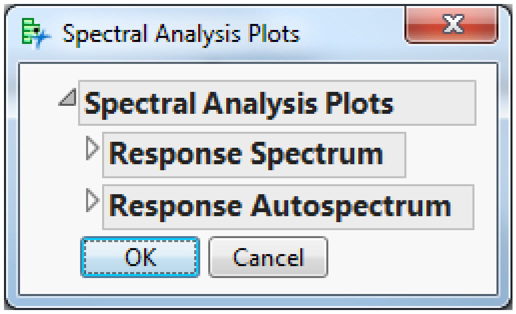

Figure 16: Spectral Analysis Plots Window. Note that since a Forcing Function series was not specified, only the Response Spectrum and Autospectrum are available and shown.

Response Spectrum

The Response Spectrum is accessed by expanding the drop-down triangle. This results in the plot shown in Figure 17. There are several things to note about this plot:

Amplitude Units

Amplitude units are in 0-Peak Response Units. Response Units are presented according to whether units were defined with the Response Input column. (In this sample, response units were not defined, so the script uses the default "Response Units".

Phase Units

Phase units are always expressed in Degrees.

Frequency Units

The frequency axis shows frequency in Cycles/Time Unit. In this example, we specified time units as "Sec", so it is displayed accordingly, both here and in the Text Readouts (see below).

Red Cursors

Red cursors are displayed in the Amplitude and Phase plots. The vertical cursor responds to the position of the mouse in the Amplitude vs. Frequency plot. It moves to the mouse position either by clicking or by clicking and dragging. The horizontal cursor and blue circle reposition based on the amplitude at the frequency of the vertical cursor. The vertical cursor in the Phase vs. Frequency plot tracks the vertical cursor in the lower plot.

Text Readouts

The text printed in the Amplitude vs. Frequency plot is updated each time the cursor is moved. Displayed values include:

- Frequency: The frequency value where the vertical cursor is located, in Cycles/Time Unit.
- Period: The reciprocal of Frequency, in Time Units/cycle.
- Ampl: The 0-Pk amplitude at the vertical cursor position, in Response Units.
- Phase: The phase angle at the vertical cursor position, in degrees.

Dynamic Frequency Scale

The frequency scale is dynamic, as with any other JMP plot. The user can grab the axis and stretch it according to his/her needs. This is particularly useful if there are many lines in the frequency spectrum, making it difficult to position the vertical cursor exactly at a particular peak. Simply stretch out the scale so the peak has a wider definition, and then place the cursor.

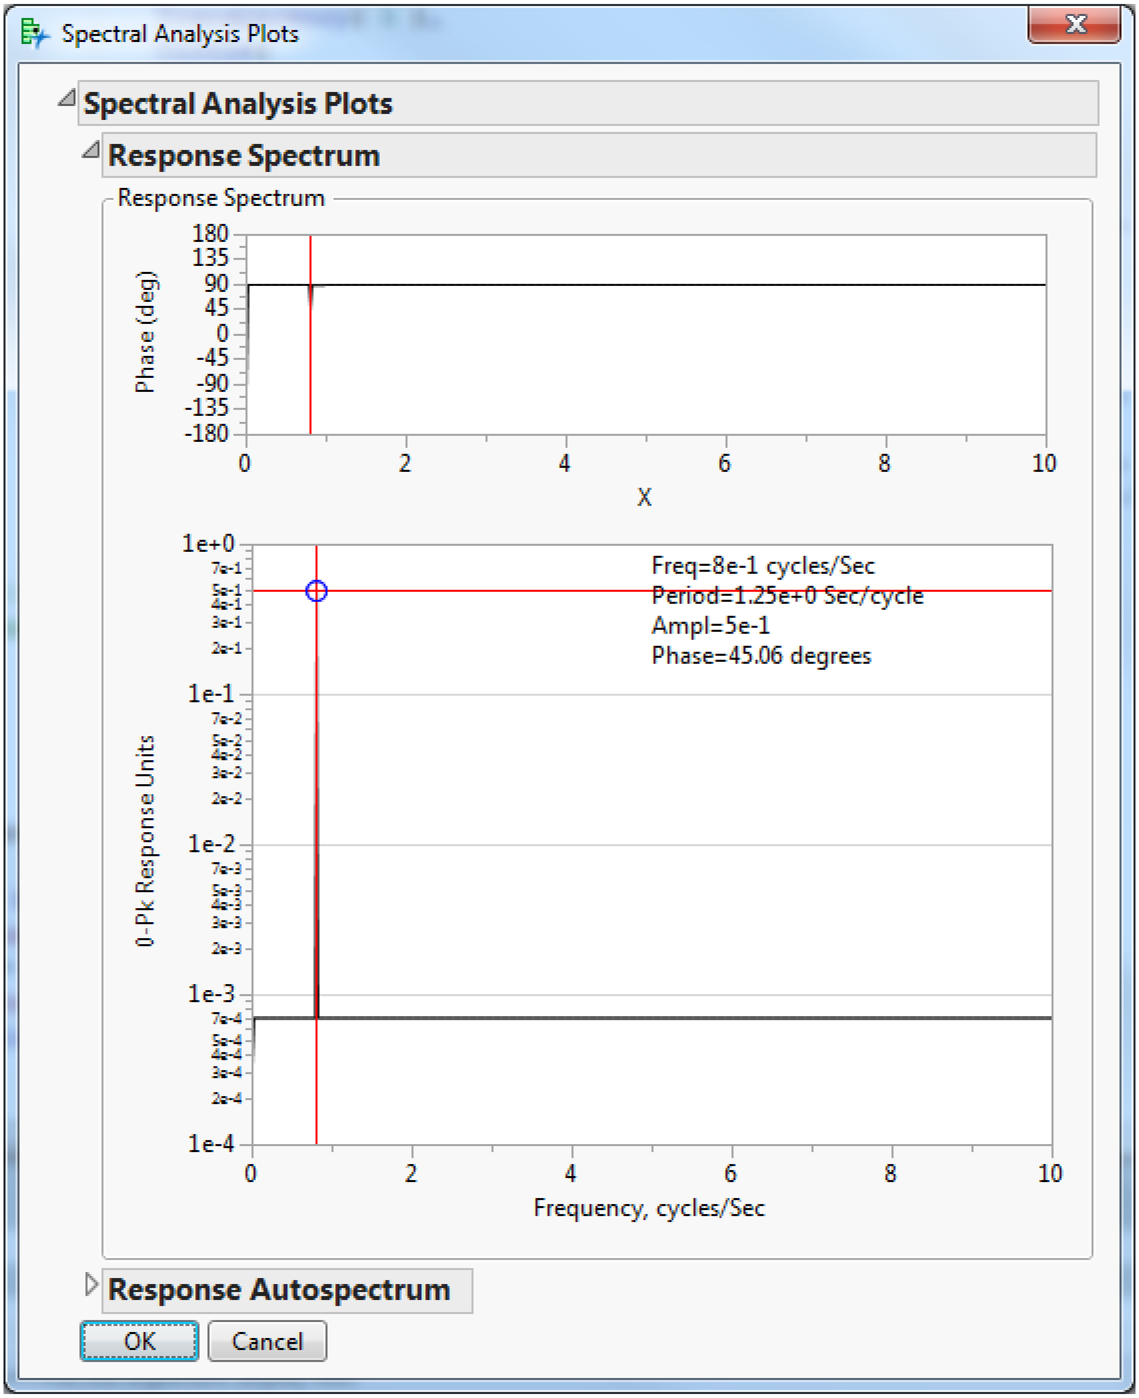

Figure 17: Example of Response Spectrum Plot

Response Autospectrum

The Response AutoSpectrum is accessed by expanding the drop-down triangle in the Spectral Analysis Plots window. This results in the plot shown in Figure 18. In addition to the notes regarding the Response Spectrum plot, please not the following:

Display Response Spectrum and Response Autospectrum Simultaneously

These two plots can be opened simultaneously as desired.

Amplitude Units

Amplitude units are in 0-Peak Response Units *Squared*. Response Units are presented according to whether units were defined with the Response Input column. (In this sample, response units were not defined, so the script uses the default "Response Units".

Phase Plot

There is no phase plot with the AutoSpectrum.

Frequency Units

Same as Response Spectrum.

Red Cursors

Same operation as Response Spectrum. Note that if the Response Spectrum plot is also opened, the cursors are linked between the two plots.

Text Readouts

The text printed in the Amplitude vs. Frequency plot is updated each time the cursor is moved. Displayed values include:

- Frequency: The frequency value where the vertical cursor is located, in Cycles/Time Unit.
- Period: The reciprocal of Frequency, in Time Units/cycle.
- Ampl: The 0-Pk amplitude *squared* at the vertical cursor position, in Response Units.

Dynamic Frequency Scale

The frequency scale is dynamic, as with any other JMP plot. The user can grab the axis and stretch it according to his/her needs. This is particularly useful if there are many lines in the frequency spectrum, making it difficult to position the vertical cursor exactly at a particular peak. Simply stretch out the scale so the peak has a wider definition, and then place the cursor.

An Example Including a Forcing Function

A 1 degree-of-freedom spring-mass-damper system (10) will be used as an example to show the features of the Forcing Function Spectrum, Forcing Function Autospectrum, and Transfer Function plots in the Frequency Analysis Script.

The 1 DOF Spring-Mass-Damper System

The 1 DOF spring-mass-damper system used in this example is shown in Figure 19. Parameters used for this particular example were:

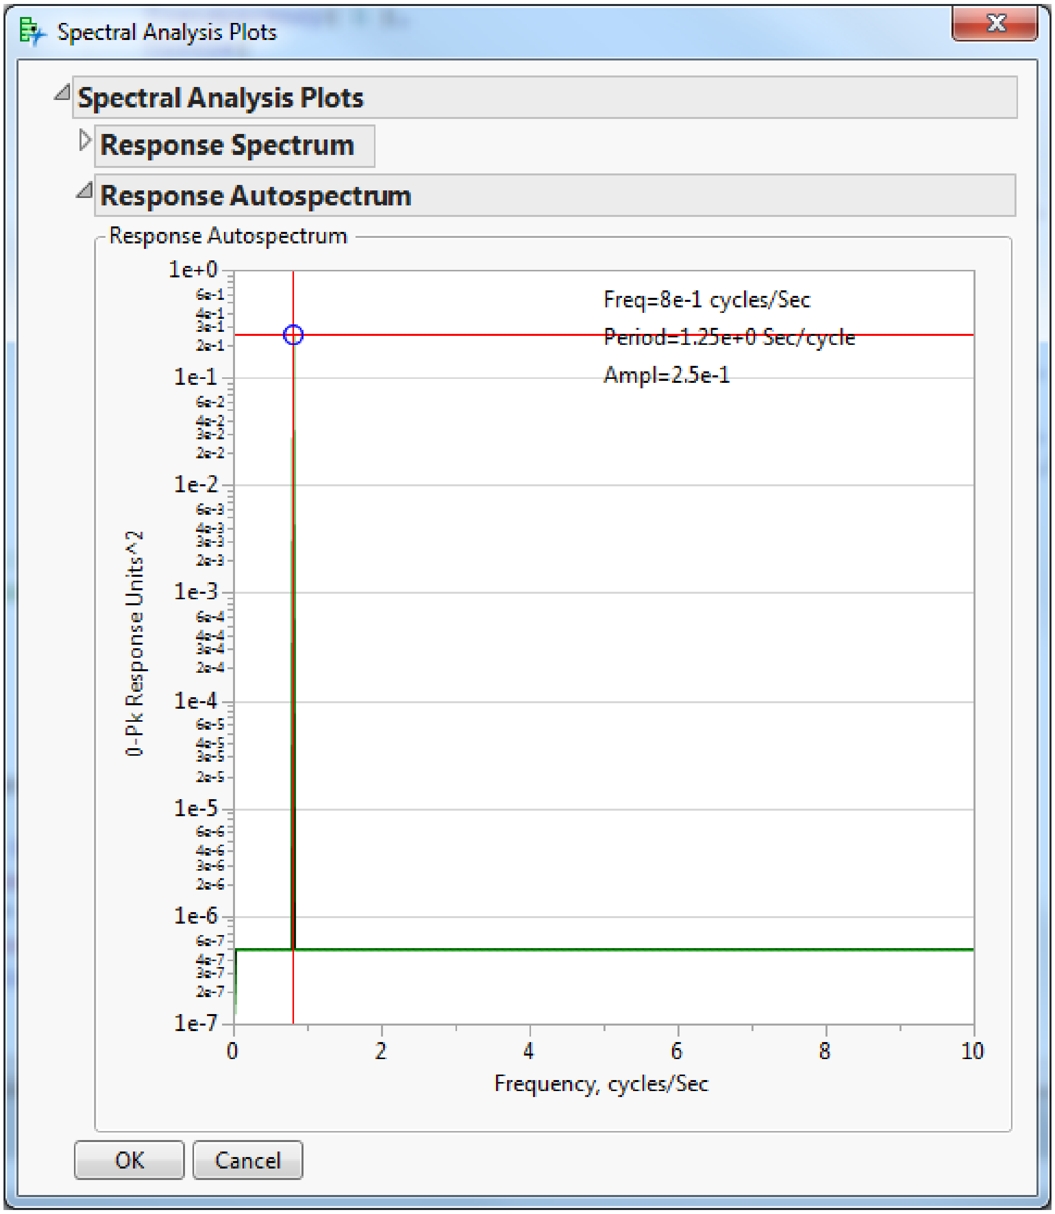

Figure 18: Example of Response Autospectrum

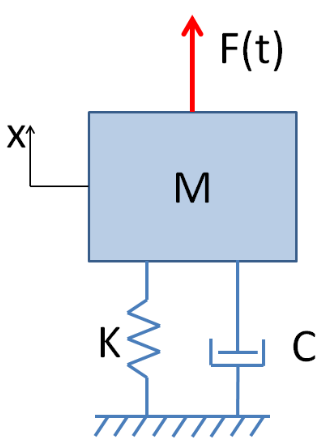

Figure 19: 1 DOF Spring-Mass-Damper System

- M = 1kg
- K = 500N/m
- C = 1 N/(m/s)
- F(t) = Random Noise, -0.5N to +0.5N, uniform distribution
- X = motion of mass (response), measured in meters.

From basic vibration equations, we have:

$$\boldsymbol{natural\ frequency = \omega_n = \sqrt{\frac{K}{M}} = \sqrt{\frac{500}{1}} = 70.7107 rad/sec}$$

$$\boldsymbol{natural\ frequency = f_n = \frac{\omega_n}{2\pi} = \sqrt{\frac{500}{1}} = 11.254\ Hz}$$

$$\boldsymbol{critical\ damping = c_{crit} = \sqrt{4KM}} = \sqrt{4 * 500 * 1} = \boldsymbol{14.142\ N/(\frac{m}{s})}$$

$$\boldsymbol{damping\ ratio = \frac{c}{c_{crit}} = \frac{1}{14.142} = 0.0707 = 7.07\%}$$

$$\boldsymbol{damped\ natural\ frequency = f_d = f_n\sqrt{1-(damping\ ratio)^2} = 11.254 * \sqrt{1-0.0707^2} = 11.226\ Hz}$$

A numerical simulation provided the response X for the input function F. The resulting forcing function and response were then sampled at 0.015 second intervals, and 1000 samples were collected (total of14.985 seconds of data).

These data were copied into a JMP data table, and the Frequency Analysis Script was used to analyze the data. A sample of the data table is shown in Figure 20. All three columns were input to the script, and all 3 columns had units pre-defined and attached to the columns prior to running the script.

Spectral Analysis Setup Window for 1DOF Example

Figure 21 shows the forcing function and response signals, and the various setup parameters in the Spectral Analysis Setup Window.

Spectral Analysis Plots Window for 1 DOF Example

Since we specified a Forcing Function in the inputs for this example, the Spectral Analysis Plots window has additional options available to us, as shown in Figure 22.

Response Spectrum

The Response Spectrum plot functions as described above. The resulting plot for the example data is shown in Figure 23. Note that there appears to be a peak at 11.14 Hz, but that the spectrum is rather noisy in appearance.

Response Autospectrum

The Response Spectrum plot functions as described above. The resulting plot for the example data is shown in Figure 24.

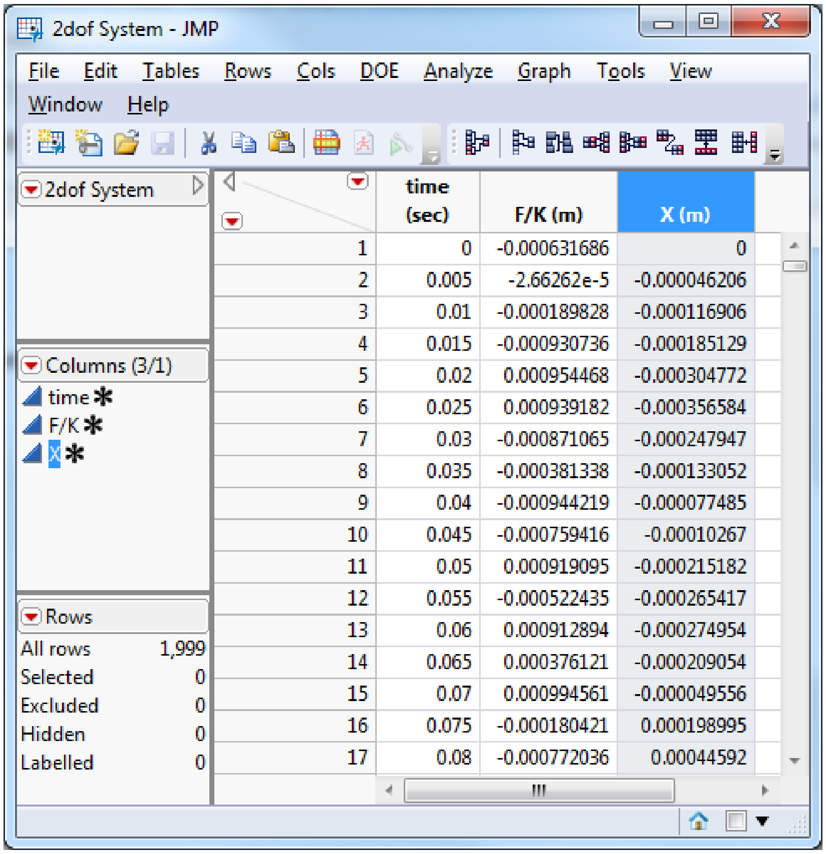

Figure 20: Sample Data for 1DOF Example

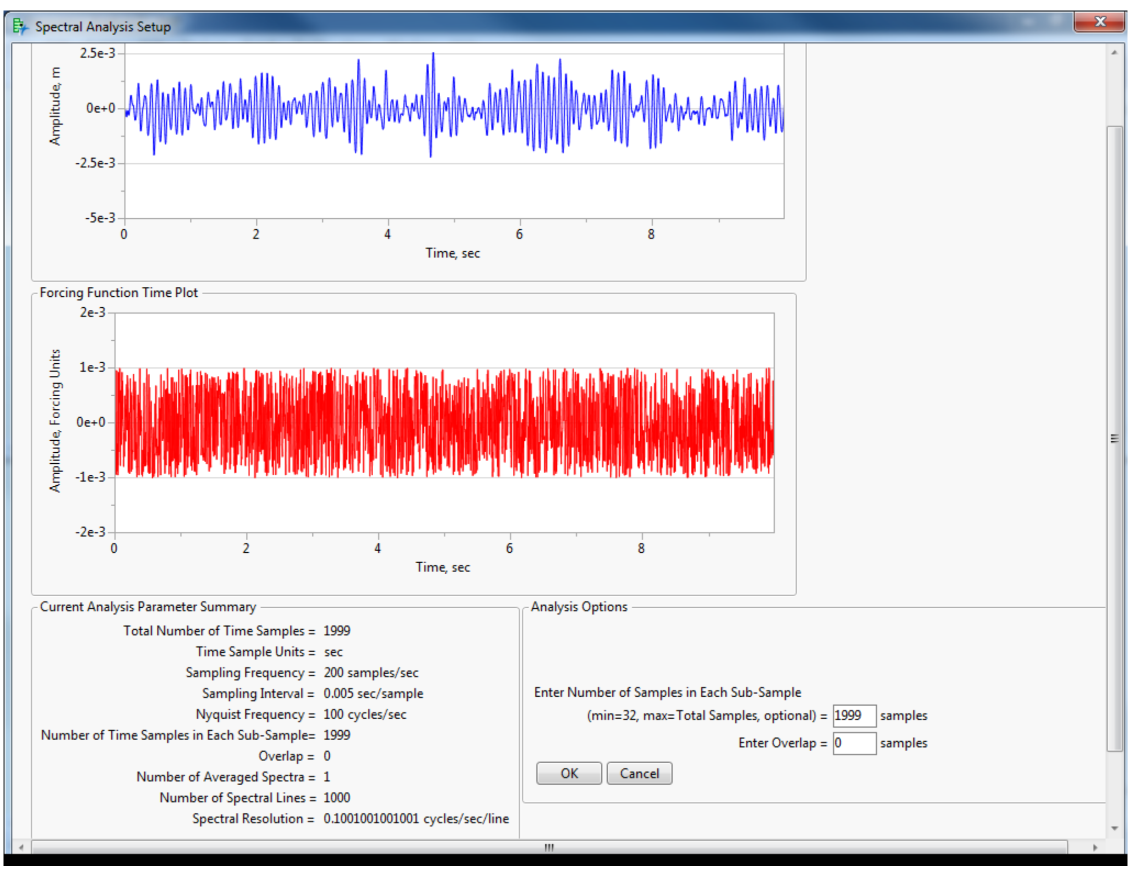

Figure 21: Spectral Analysis Setup for 1DOF System Example

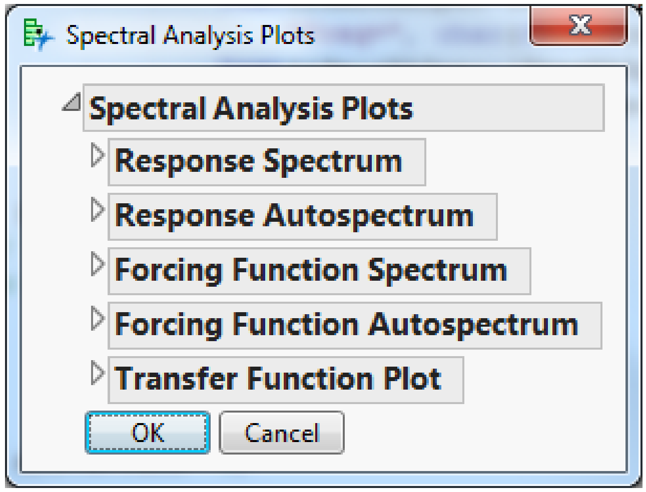

Figure 22: Spectral Analysis Plots Window When Forcing Function Is Defined

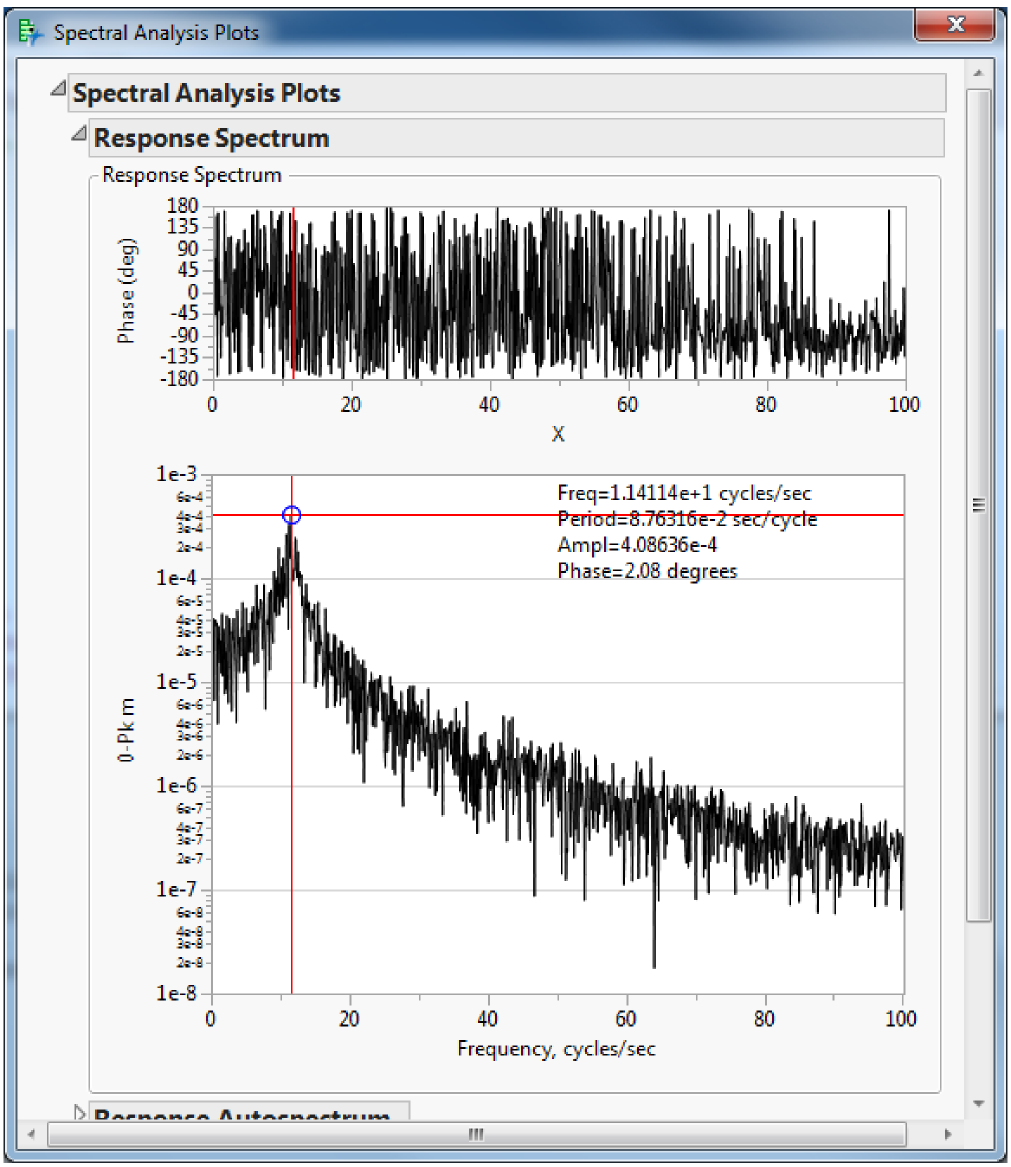

Figure 23: Response Spectrum for 1DOF Example

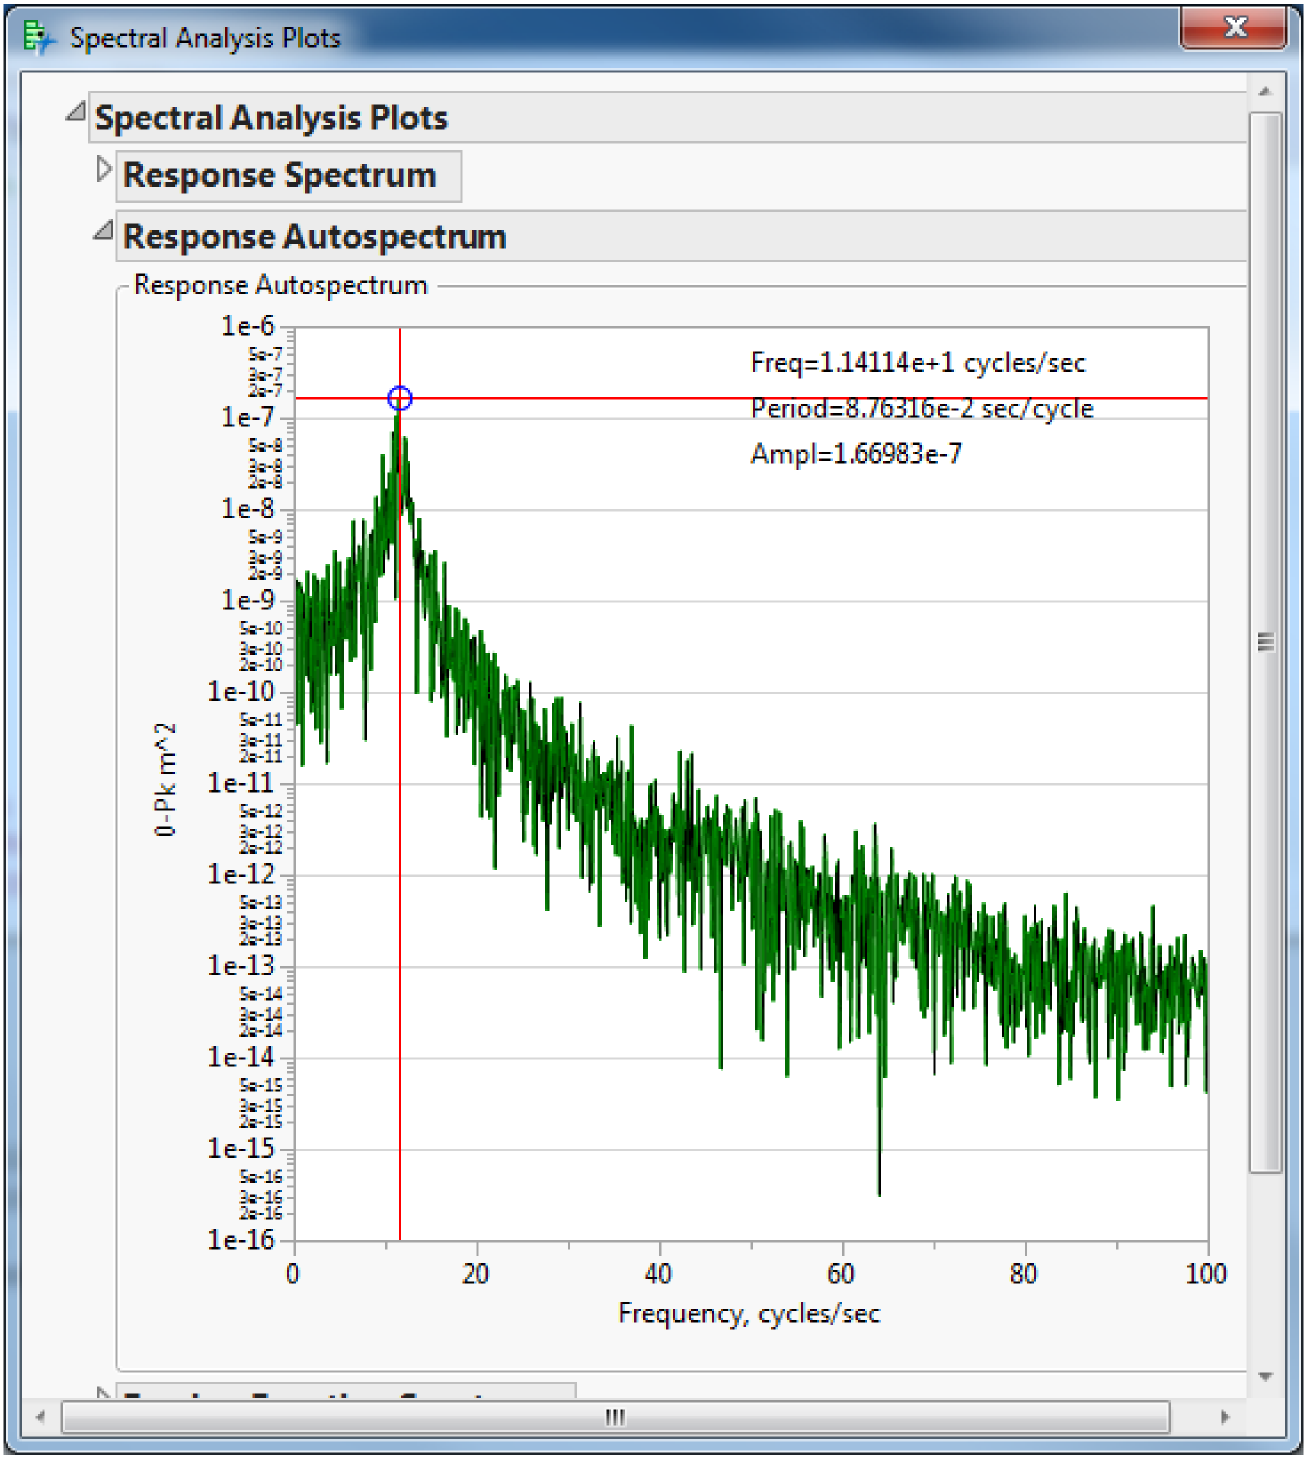

Figure 24: Response Autospectrum for 1DOF Example

Forcing Function Spectrum

The Forcing Function Spectrum for the 1DOF example is shown in Figure 25. The plot controls are identical to the Response Spectrum controls. Note that the input Forcing Function was defined to be random (white) noise, which appears as a relatively flat (albeit noisy) frequency spectrum.

Forcing Function Autospectrum

The Forcing Function Autospectrum for the 1DOF example is shown in Figure 26. The plot controls are identical to the Response Autospectrum controls.

Transfer Function Plot

The Transfer Function (also known as a Frequency Response Function) Plot represents the relationship of the Response to the Input. It can be thought of as:

> **The transfer function is the frequency-by-frequency behavior of a response signal to a sinusoidal input function of unit amplitude.**

The Transfer Function for the 1DOF example is shown in Figure 27.

The Transfer Function Plot works much the same way as the Response Spectrum and Forcing Function Spectrum plots, in that the cursor controls work the same, and there is an amplitude and a phase plot. However, Damping is also reported. This is the Damping Ratio for the system.

Damping Ratio

The Damping Ratio is determined from the "half power points" around the peak in the transfer function. From the user selected cursor position, the software looks for the points to the left (lower frequency) and right (higher frequency) around the peak that have amplitudes equal to half of the peak amplitude. The damping ratio is then determined from the following expression:

$$measured\ damping\ ratio = \frac{f_{peak}}{\Delta f}$$

where Δf is the distance (in frequency units) between the half-power points.

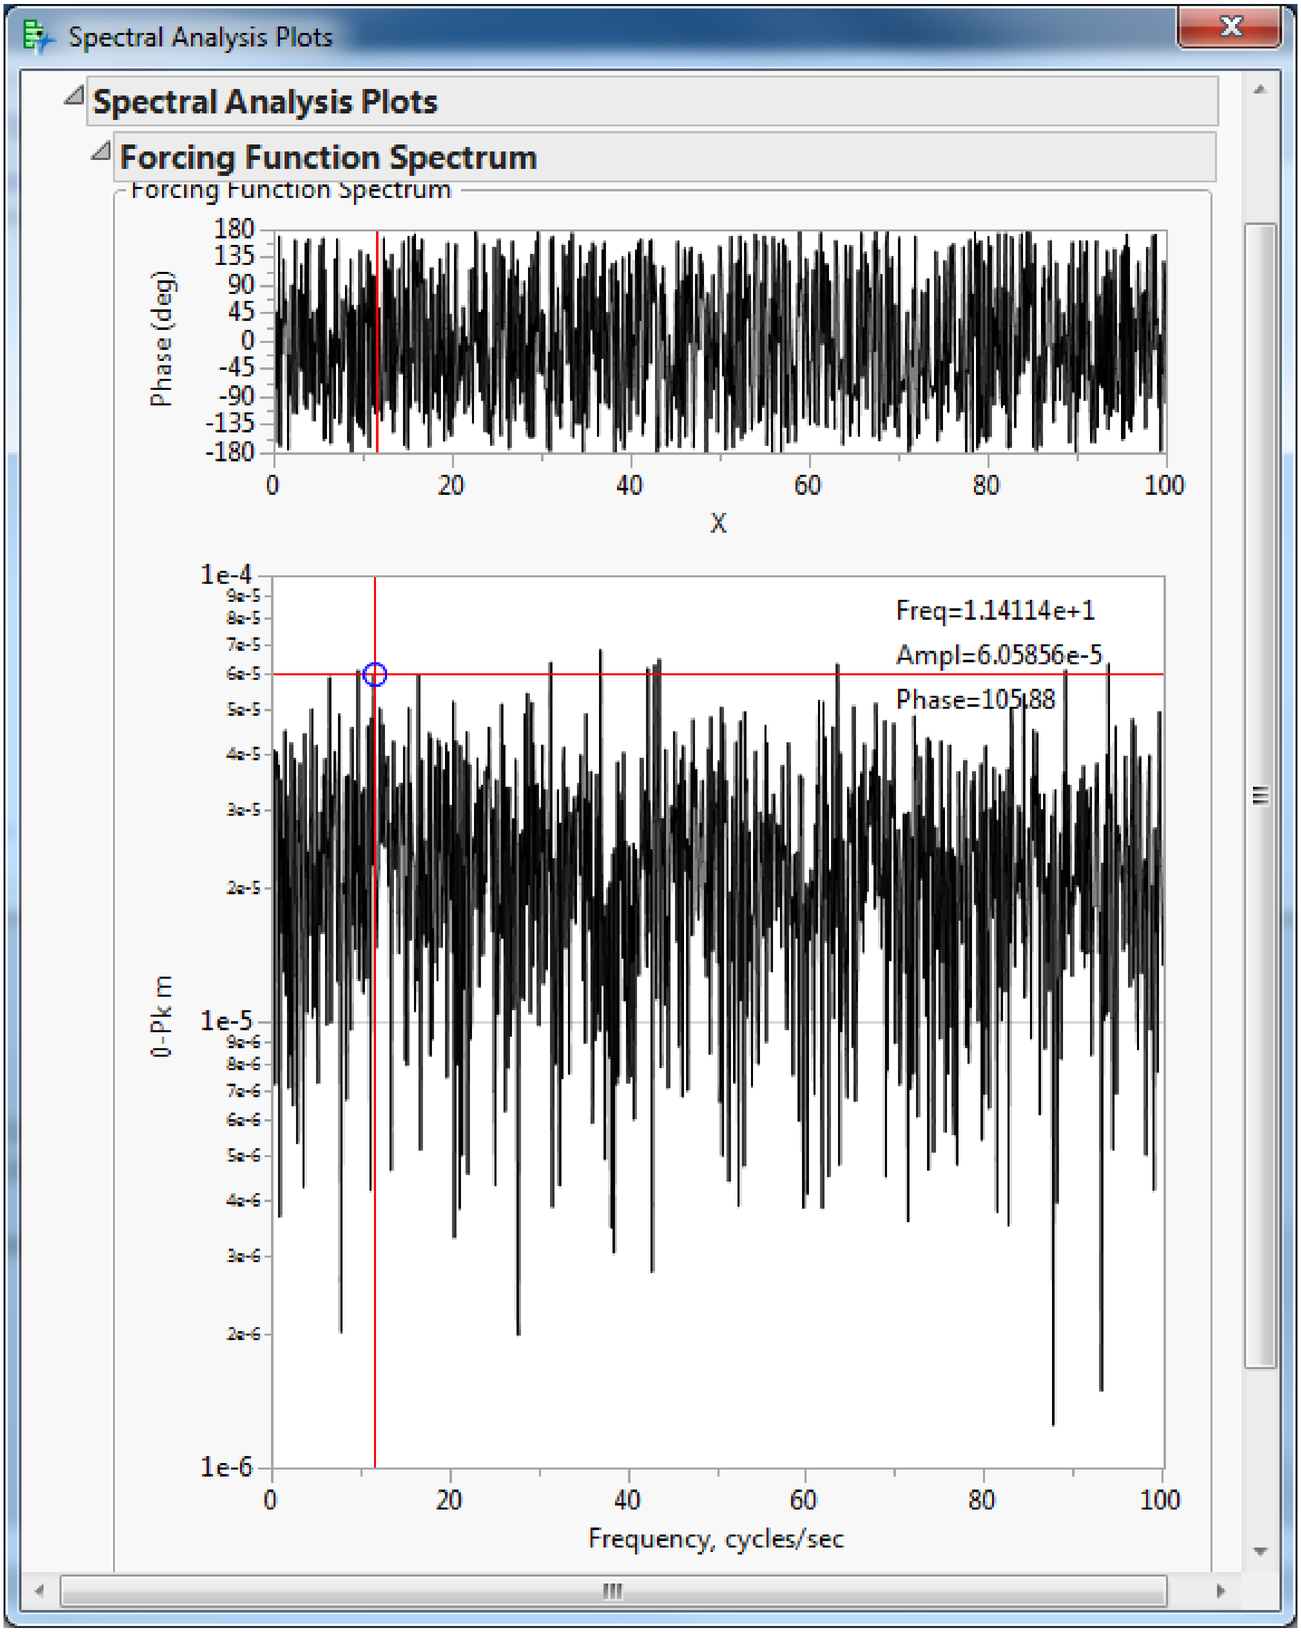

Figure 25: Forcing Function Plot for 1DOF Example

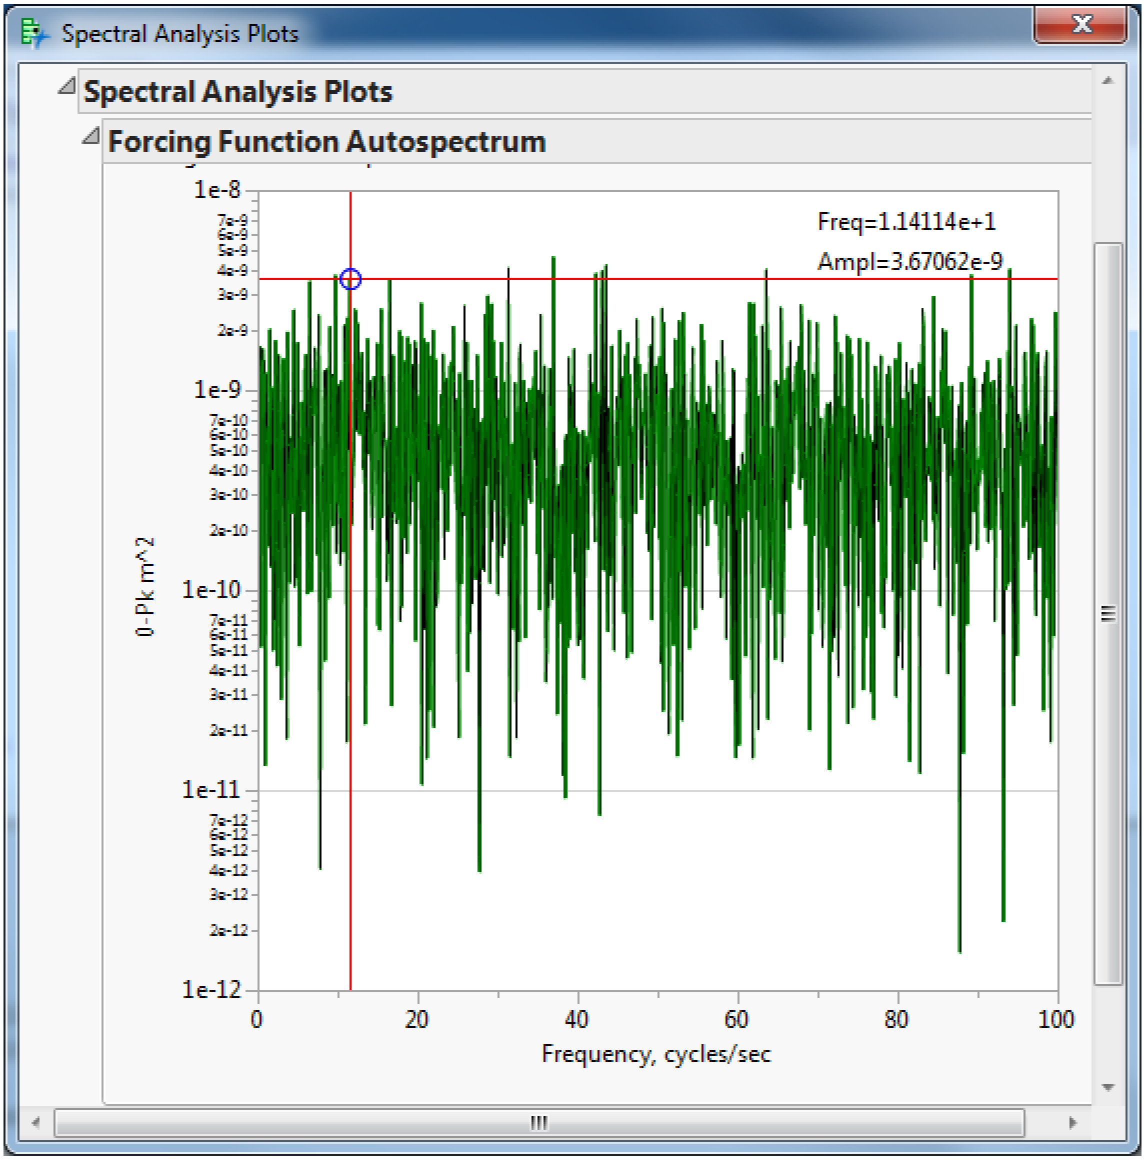

Figure 26: Forcing Function Autospectrum for 1DOF System

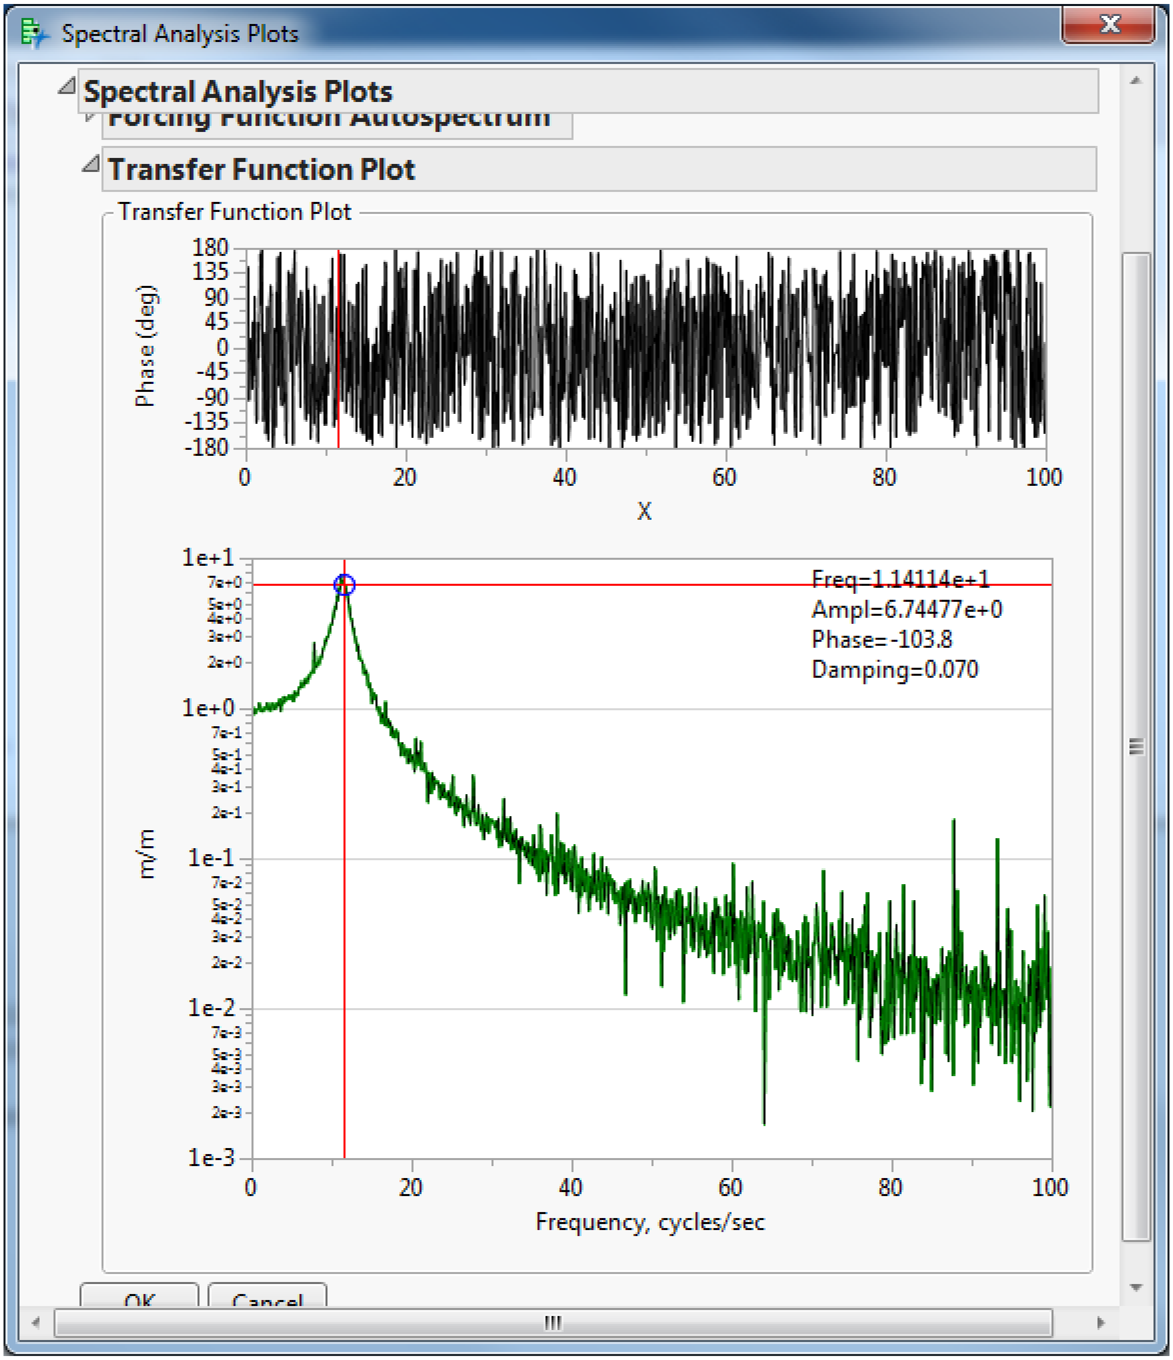

Figure 27: Transfer Function for 1DOF System

Additional Functionality of the Frequency Analysis Script

As shown in Figure 15 and discussed previously, there are three additional options available in the Frequency Analysis Script which may help in signal analysis. These options, as well as a resulting differences in the Spectrum and Autospectrum plots, and the new Waterfall Plot, are described below.

Number of Samples in Each Sub-Sample

The total number of samples in the data set is defined by the number of rows in the data table. However, there may be times when the user wishes to “split up” the large sample into sub-samples. This can help in smoothing out noisy signals, or selecting smaller samples to show the effects of frequency shifts in a waterfall plot.

Figure 28 shows a typical response time plot for a time series with 2000 data points. If the user chooses to subdivide the sample into, say, 180 sample sub-series, the Frequency Analysis Script would subdivide the original time series into 11 adjacent and sequential sub-samples, comprising 11*180=1980 data points. (The last 20 data points are ignored.)

Spectral analysis is then performed on each of the sub-samples. Note that since the number of points in each sub-sample is smaller than the original sample, frequency resolution will be reduced accordingly.

Once the computations are complete, the user can select from a new set of Spectral Analysis Plots (see Figure 29). While the first two of these plots has already been discussed, there is new functionality available in the plots as discussed below.

Response Spectrum

Figure 30 shows the Response Spectrum plot that results from specifying sub-samples and/or overlap. As discussed above, the frequency resolution of the spectrum has been adjusted according to the number of samples in each sub-sample. However, there is now a slider below the plot that allows the user to quickly index between sub-sample spectra to observe differences over time.

Response Autospectrum

The Response Autospectrumnow shows two green lines, one above and one below the center black line. (See Figure 31.)

The center black line represents the average amplitude at each frequency of all of the sub-samples.

The two green lines represent the upper and lower 95% confidence bounds on the average amplitude at each frequency in the spectrum. These lines give an idea of the repeatability of the spectral amplitudes over time.

Waterfall Plot

If SubSamples are specified (and/or if Overlap is specified, as described later), then the Response Waterfall Plot option is available to the user. A typical waterfall plot is shown in Figure 31.

The Waterfall plot shows all of the individual spectral amplitudes for the individual spectra in the left-side plot. The “Y Scale” input box allows the user to adjust the scale factors on the spectral amplitudes. (Y-Scale = 1 by default. Depending on the amplitude of your particular signal, you may have to increase or decrease this value to give reasonable plots.) A vertical cursor can be dragged back and forth across the plot with the mouse to identify specific frequencies.

To the right side of the waterfall plot is a second plot. This plot shows the (max – min)/2 value for each subsample. In this case, amplitude is plotted on the x axis, and subsample number is plotted on the Y axis. The red horizontal cursor can be dragged up and down the plot to identify where maximum excursions occurred in the time signal.

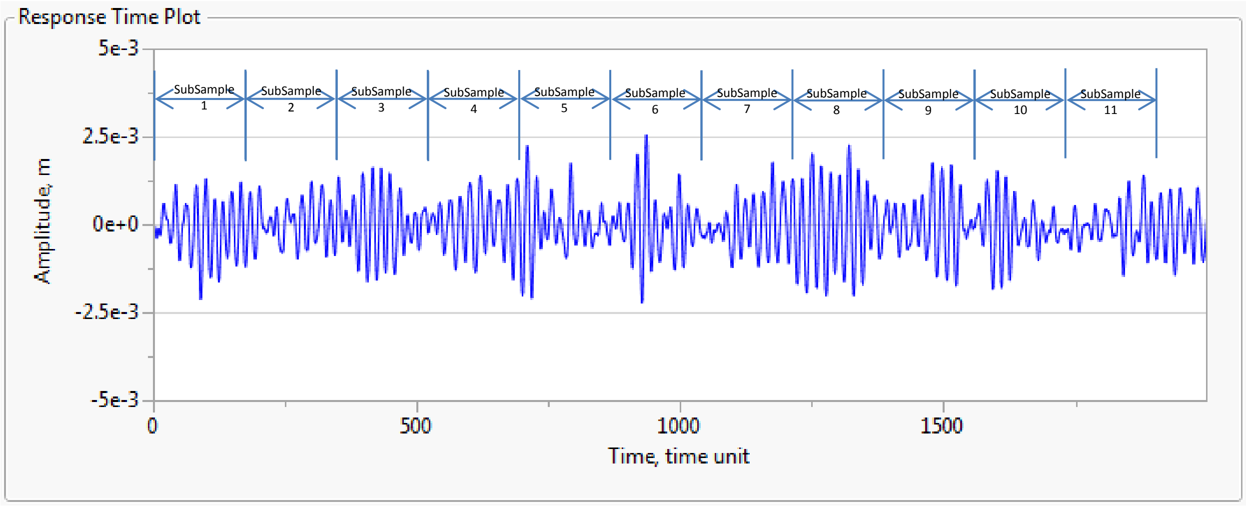

Figure 28: Response Time Plot showing how Sub-Samples are defined.

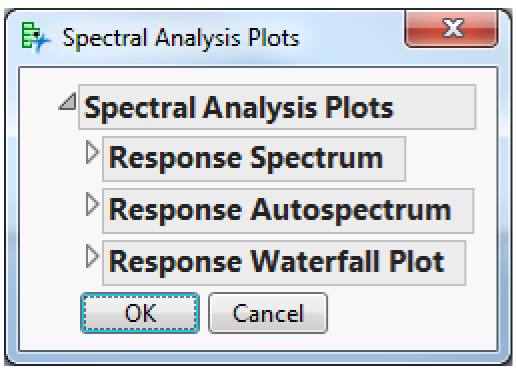

Figure 29: Spectral Analysis Plot Window Options when Sub-Samples and/or Overlap Are Specified

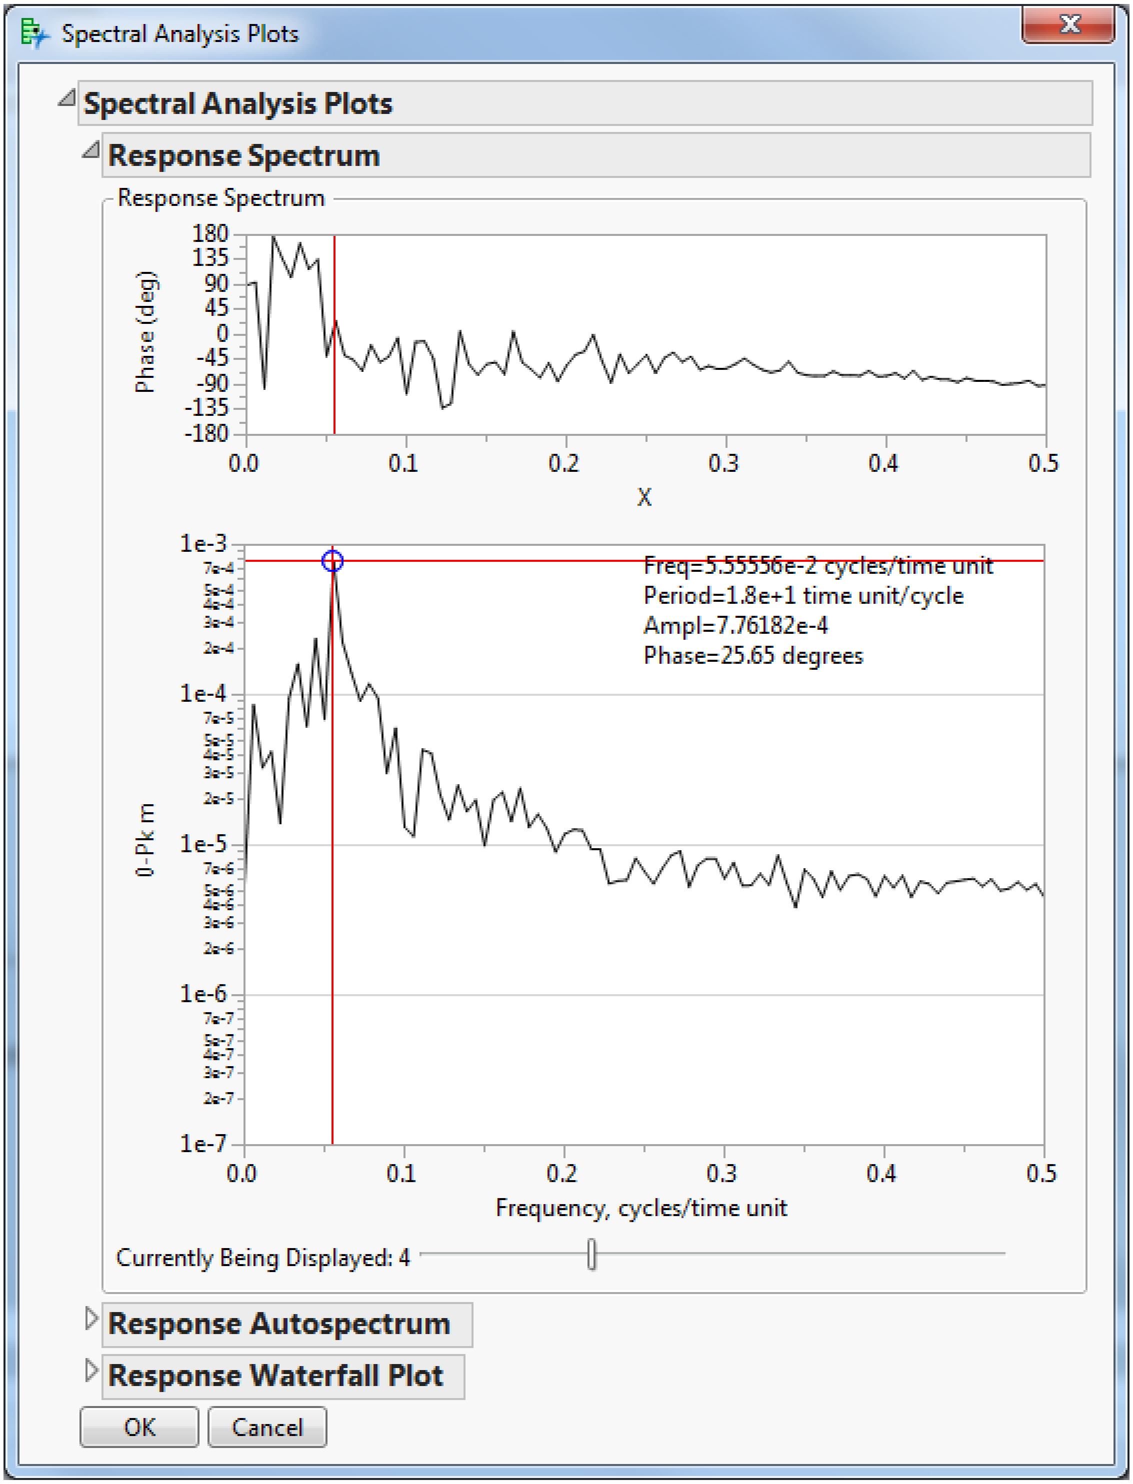

Figure 30: Response Spectrum Plot When SubSamples and/or Overlap Have Been Specified

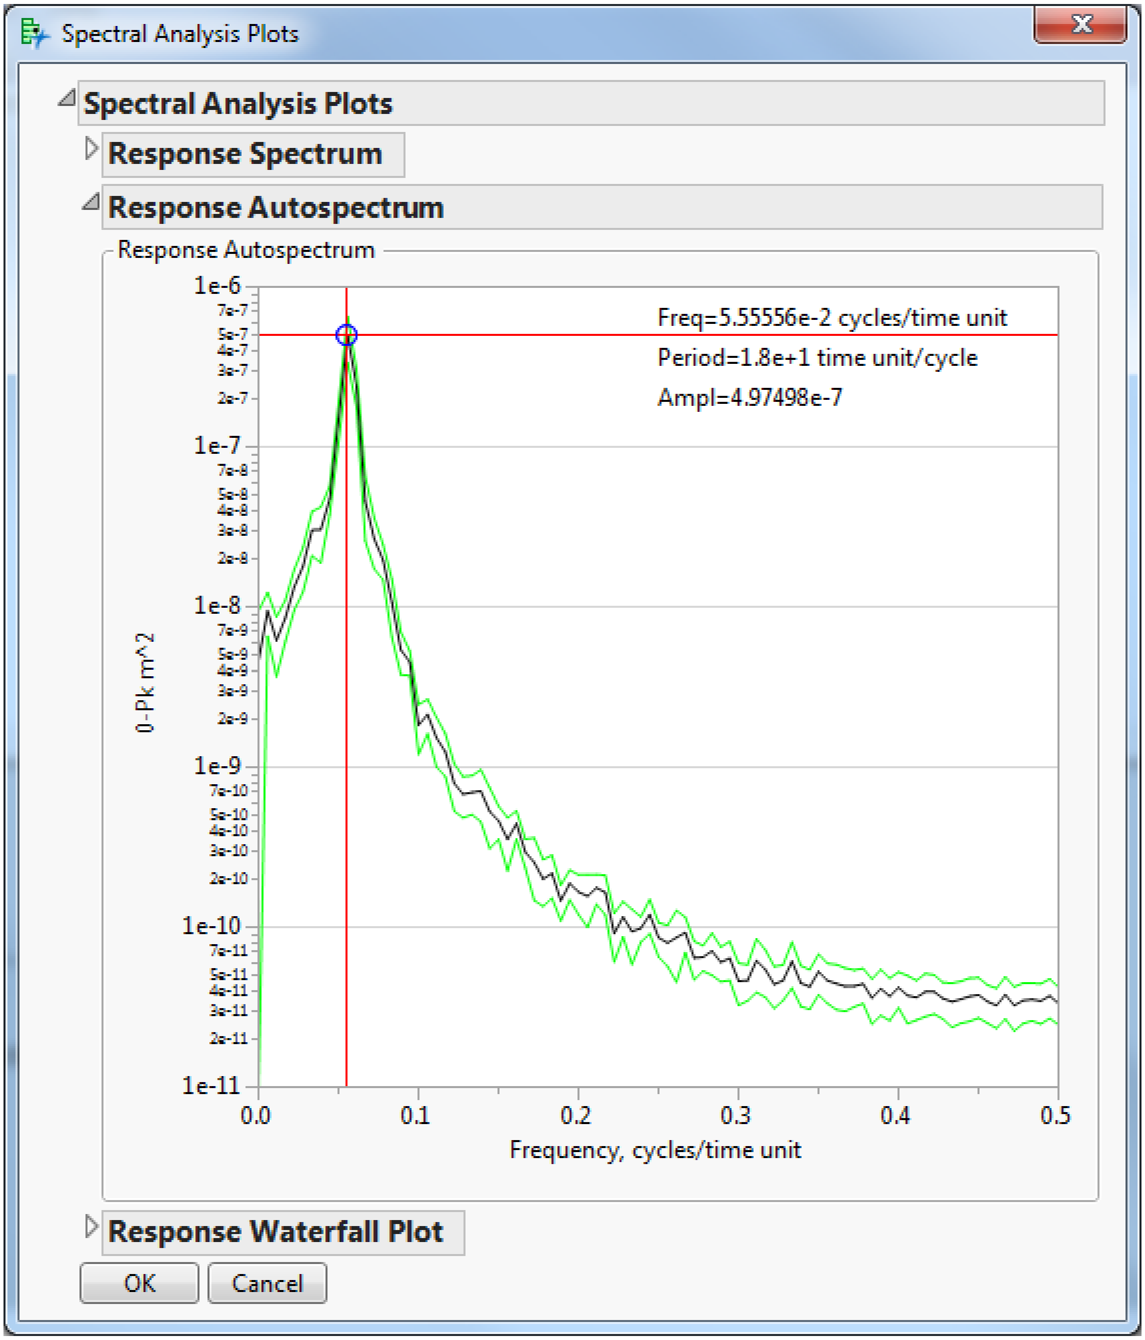

Figure 31: Response Autospectrum Plot When Sub-Samples and/or Overlap Are Specified

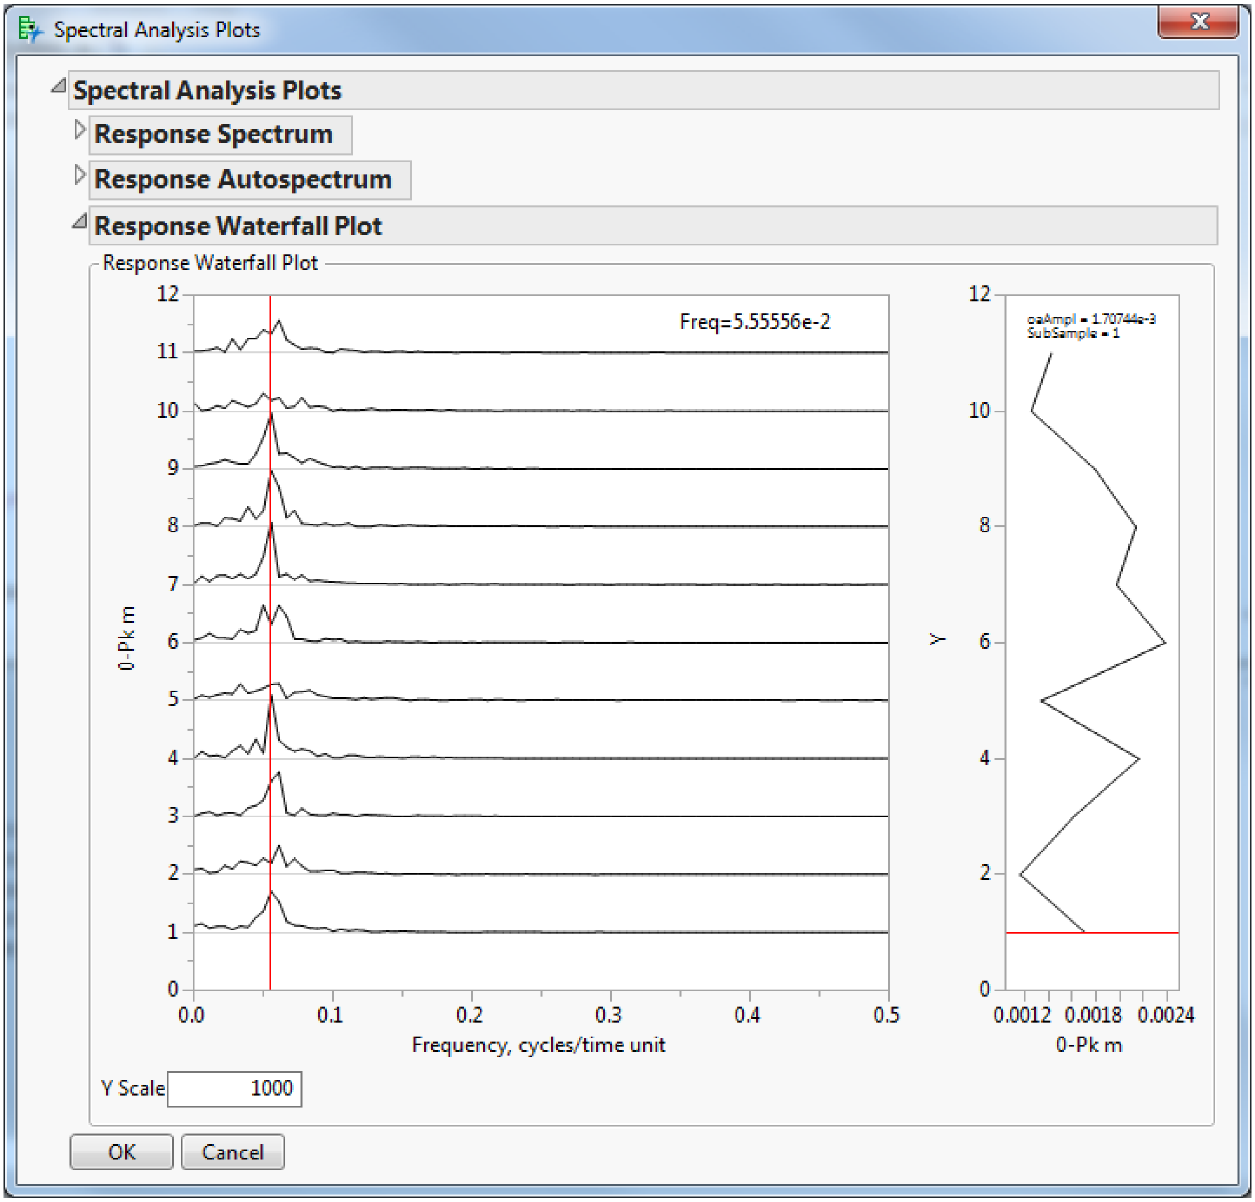

Figure 32: Response Waterfall Plot

Forcing Function Plots

If Subsamples (and/or Overlap) have been identified in the setup window, and a Forcing Function has also been entered, the same plots and functions will be available for Forcing Function data as were described previously for Response data.

Overlap

Overlap is defined as the number of samples that sub-samples can overlap each other. For example, if we look at the same signal shown in Figure 28, we might choose to have the 180 point subsamples overlap each other by 30 samples. This is depicted graphically in Figure 33. Note that the spectral resolution will not change from the previous example (since we are still using N=180 for the subsample size), but we now have 13 total subsamples available because of the overlap parameter.

The Response Spectrum, Response Autospectrum, and Response Waterfall Plot all work as described above.

Hanning Window

Windowing is a signal conditioning technique that helps to reduce noise in a spectrum caused by various artifacts. One of these windows is called the Hanning Window, named after Julius von Hann (1839-1921), an Austrian physicist and meteorologist.

The Hanning Window is used to reduce spectral noise due to "leakage". Recall the earlier discussion that the Discrete Fourier Series assumes that the signal repeats in time. Consider a 1 Hz sine wave that is sampled at 0.1 samples per second, and a total of 13 samples are collected. The resulting signal is shown in Figure 34. Also shown in Figure 34 are the imaginary repeating series preceding and following the measured series.

If we were to use the Frequency Analysis Script (without windowing) to analyze such a signal (where the length of the time series does not equal an integer number of periods of a perfect sine wave), we might get a Response Spectrum plot like that shown in Figure 35. Note that in Figure 35, instead of being a single spectral line, the shape of the spectrum now has smooth "curves" descending from either side of the peak. This phenomenon is called "leakage", and affects both the amplitude and the phase of the resulting spectrum.

One way to avoid this problem is to multiply the original time signal by a function that forces the endpoints of the time signal to zero. By doing this, the time series is guaranteed to repeat.

In fact, the Hanning Window does just that. The Hanning function is a "raised Cosine" function which has a value of zero at the endpoints and unity in the center.

Consider the time series data shown in Figure 36. (This is the same data used in calculating the spectrum shown in Figure 35). Shown are the original signal (N=193 data points), and the assumed repeating series preceding and following the actual series. Note the discontinuities at the beginning and end of the series.

Figure 37 shows the result of applying the Hanning Window to the original time series. Note that the beginning and end points of the new series are now zero.

Figure 38 shows that same series of Figure 37, along with the assumed preceding and following series. Now there are no discontinuities in the signal.

Figure 39 shows the resulting spectrum. While the spectrum is still not a perfect "spike" at the input frequency, the side slopes are greatly attenuated. (Compare to Figure 35.)

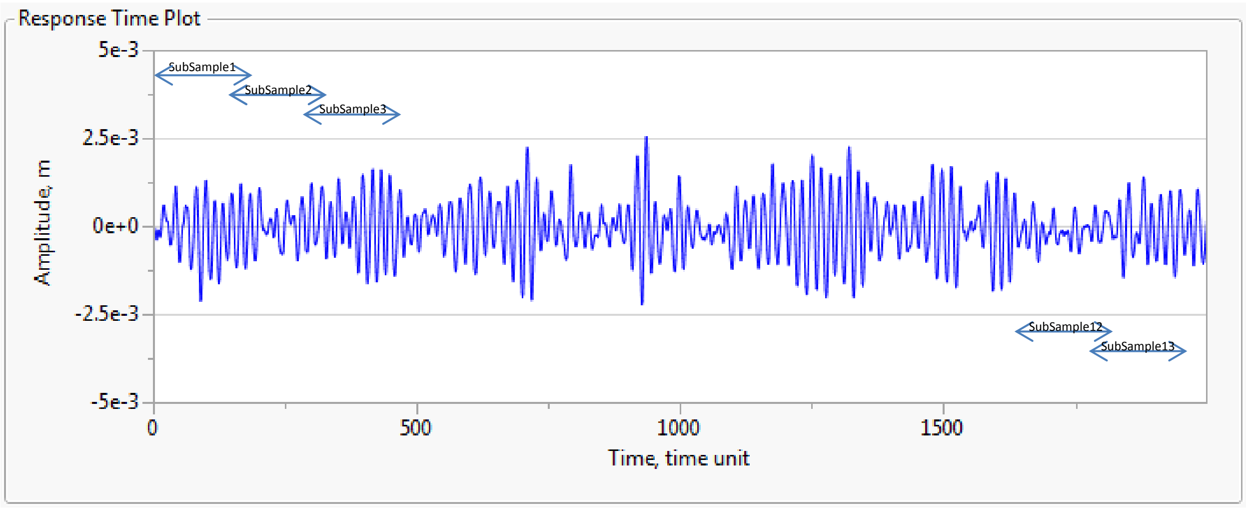

Figure 33: Demonstration of Overlap of Sub-Samples

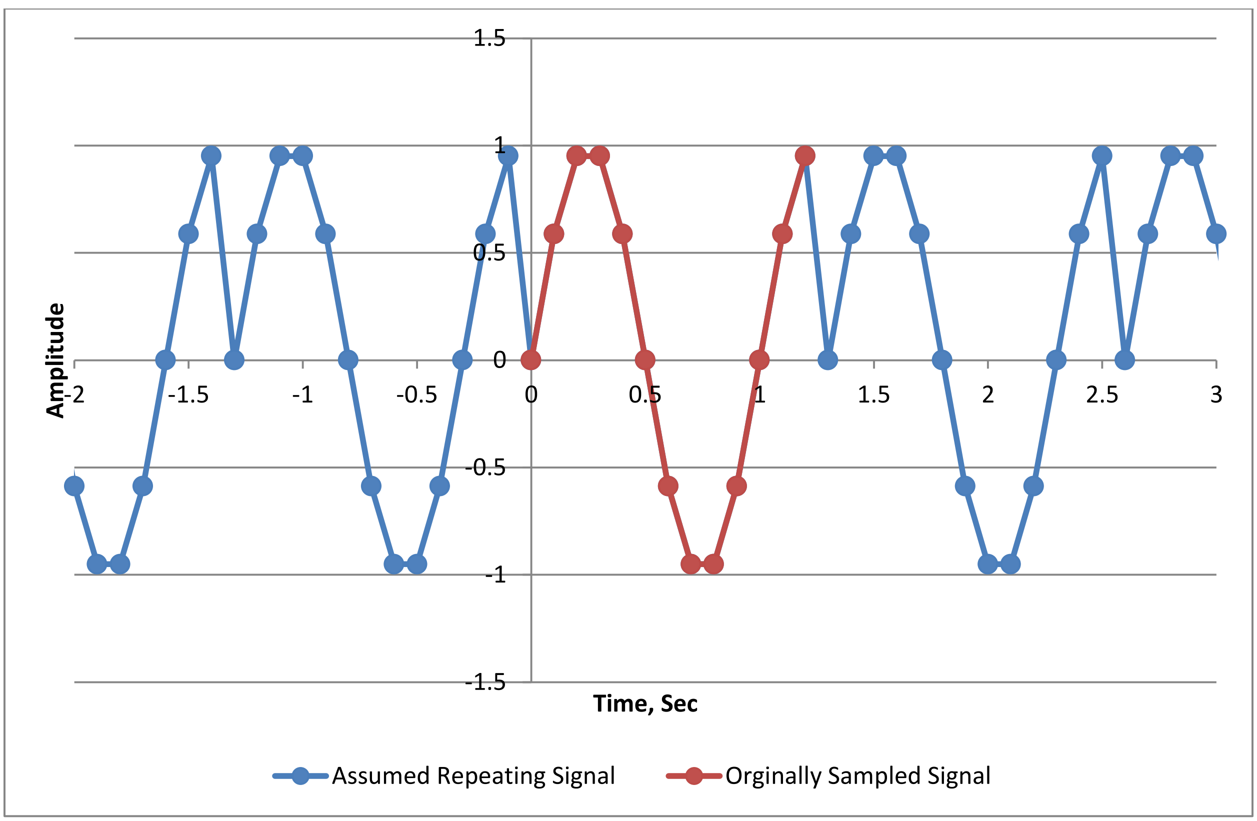

Figure 34: Demonstration of Discontinuities at Boundaries of Sine Wave if Not Matched with Sample Window Length

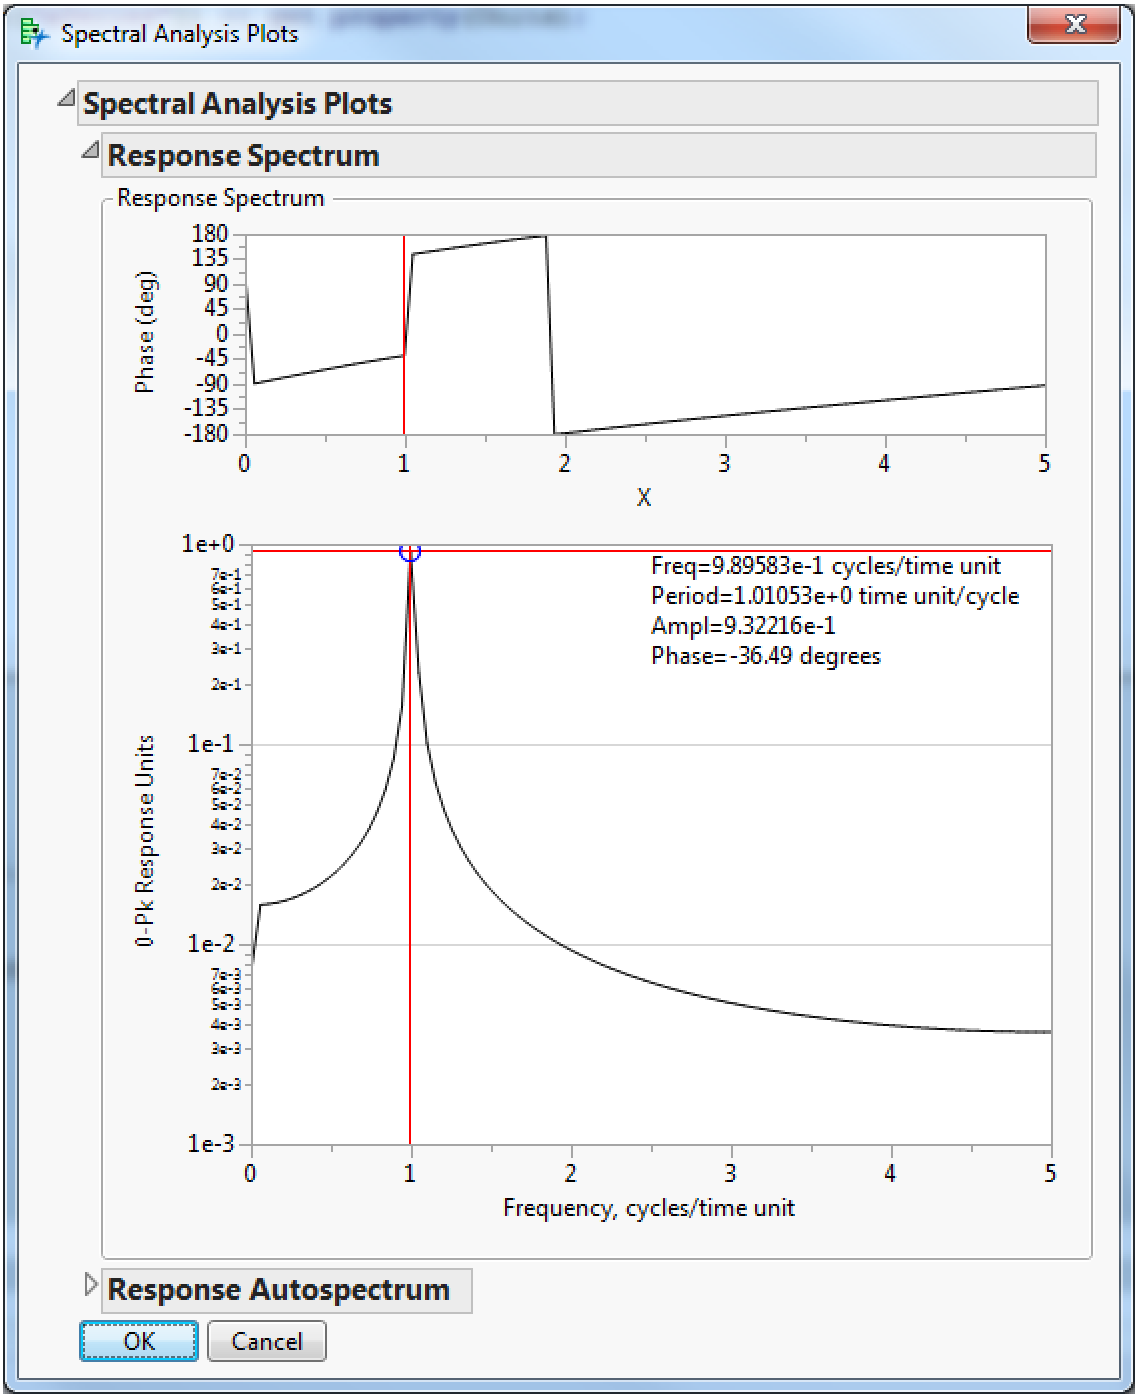

Figure 35: Example of "Leakage" when Length of Time Series Is Not an Integer Multiple of Sine Wave

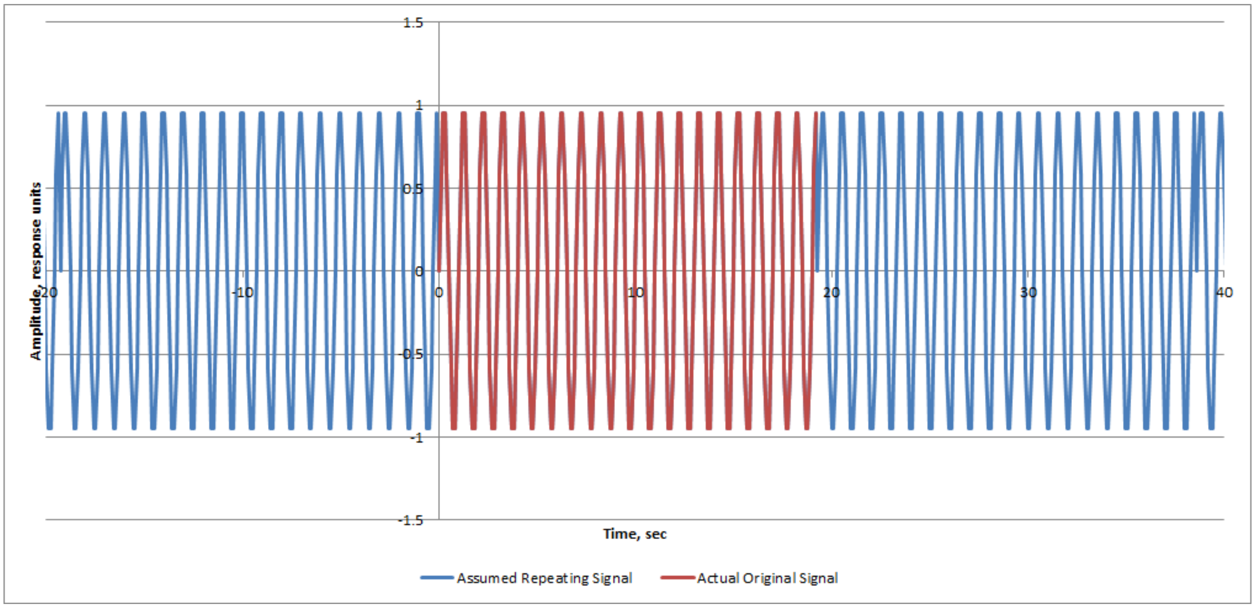

Figure 36: Time Series Data for Simple Sine Wave with Non-Integer Number of Cycles in Sample

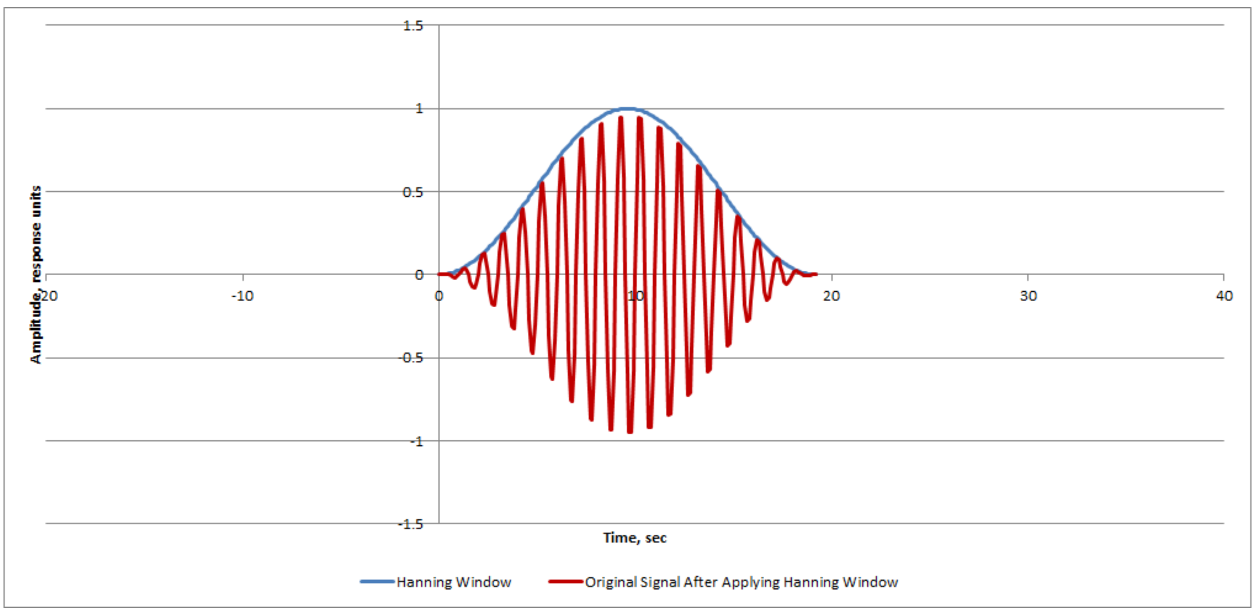

Figure 37: Hanning Window applied to time series of Figure 36

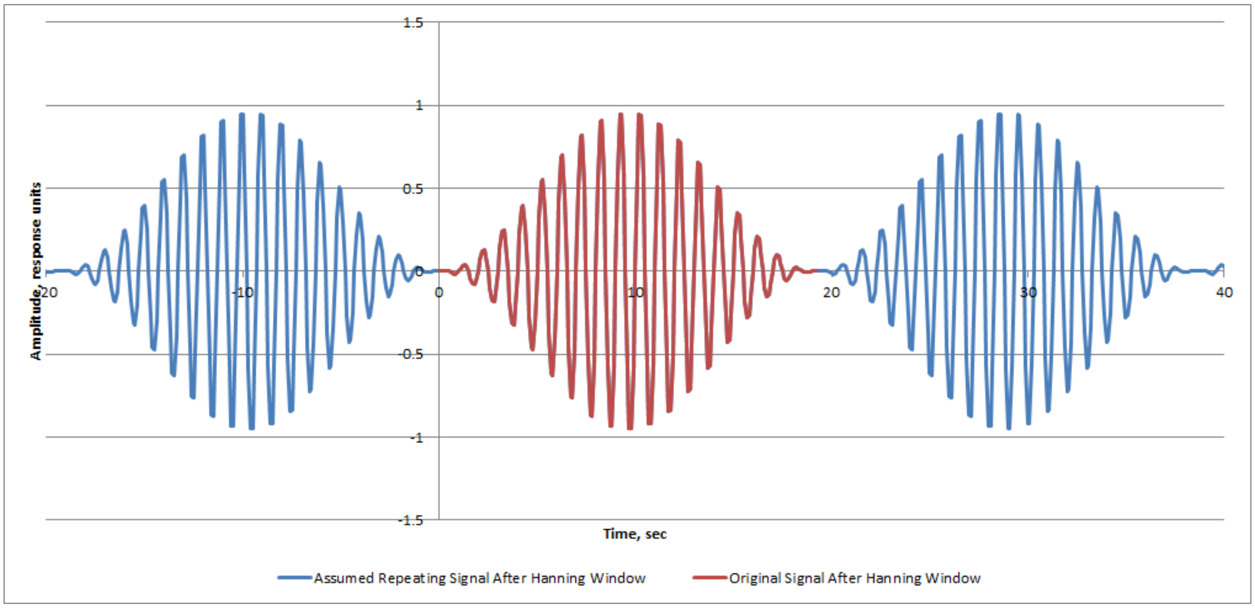

Figure 38: Time Series of Fig 36 After Hanning, with Assumed Preceding and Following Series

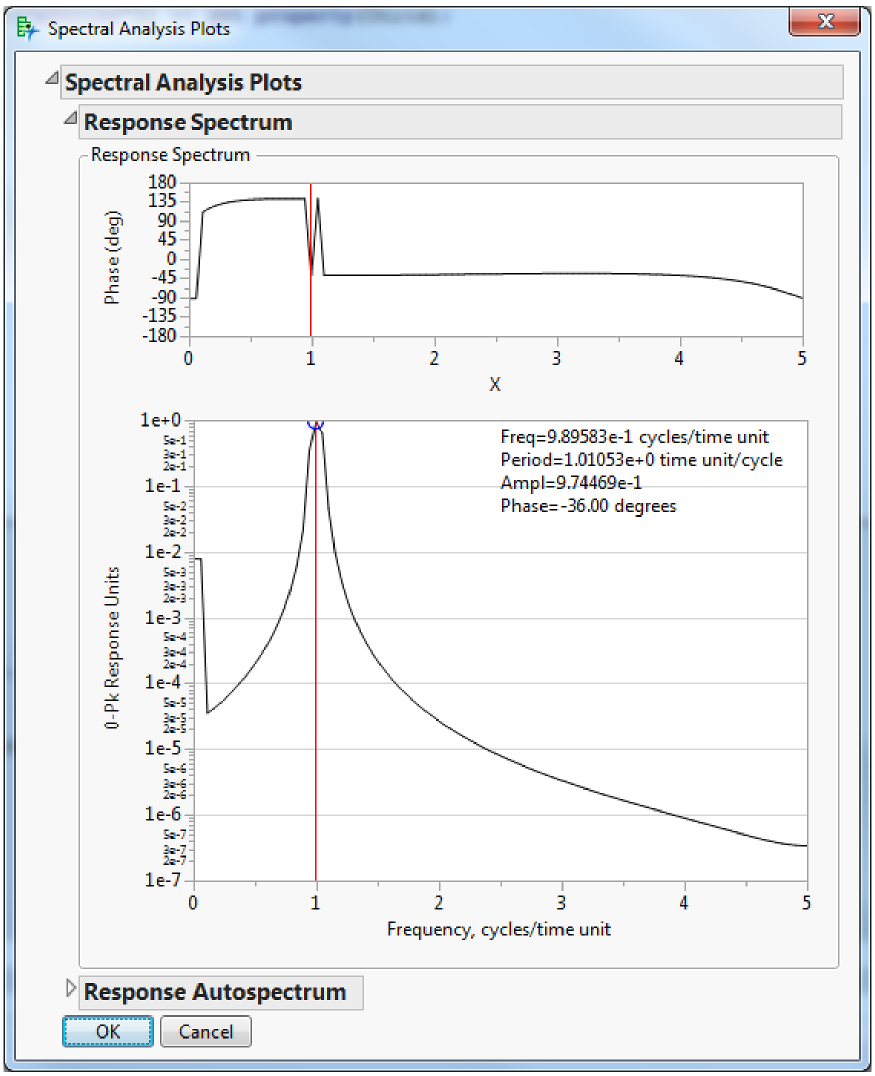

Figure 39: Spectrum Resulting from Time Series in Figure 37 (after applying Hanning Window)

Example Analysis of a Time Varying Signal

As mentioned previously, Waterfall Plots can be used to study time varying signals.

To demonstrate this feature, two (human) whistles were recorded independently, and the two signals combined (added) to produce the signal to be analyzed:

- The first whistle was a constant tone.
- The second was a time varying tone (started at low frequency, swept to high frequency, then swept back to low frequency.
- Signals were recorded at 8000 samples per second.
- A total of 1.5 seconds of data was collected (a total of 12000 data points).

Figure 40 shows the Spectral Analysis Setup window, showing the original time signal and the parameters used in the analysis. Note that we requested 500 point subsamples, with an overlap of 250 points, resulting in a total of 47 frequency spectra.

Figure 50 shows the resulting waterfall plot. There is a clear tone at a constant frequency of 1824 Hz. The time varying signal also stands out, starting at 1200 Hz, climbing to 2320 Hz (shown by the red cursor line in Figure 50), and then falling back to 960 Hz at the end of the sample. In the plot to the right, we see that the peak overall signal amplitude occurred within the 9th subsample.

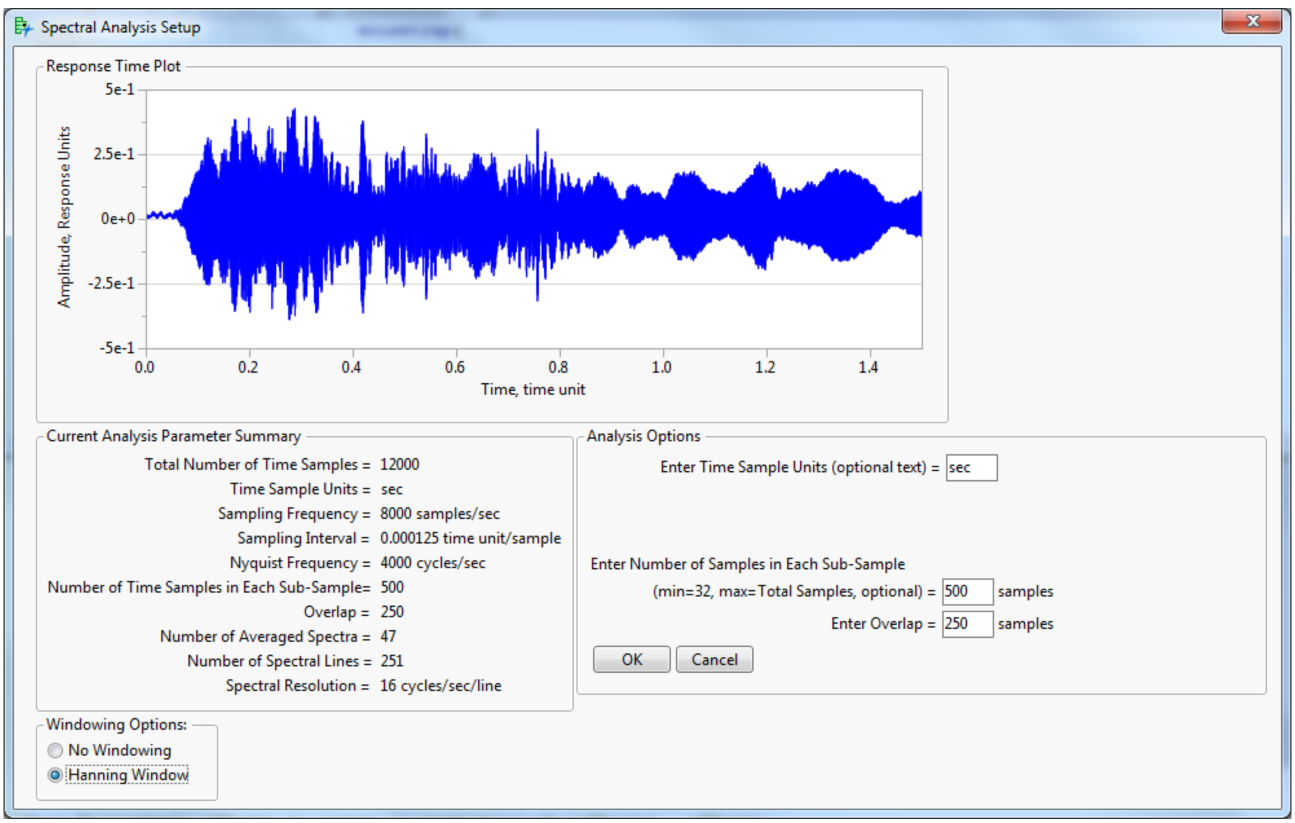

Figure 40: Spectral Analysis Setup for Frequency Varying Input Signal

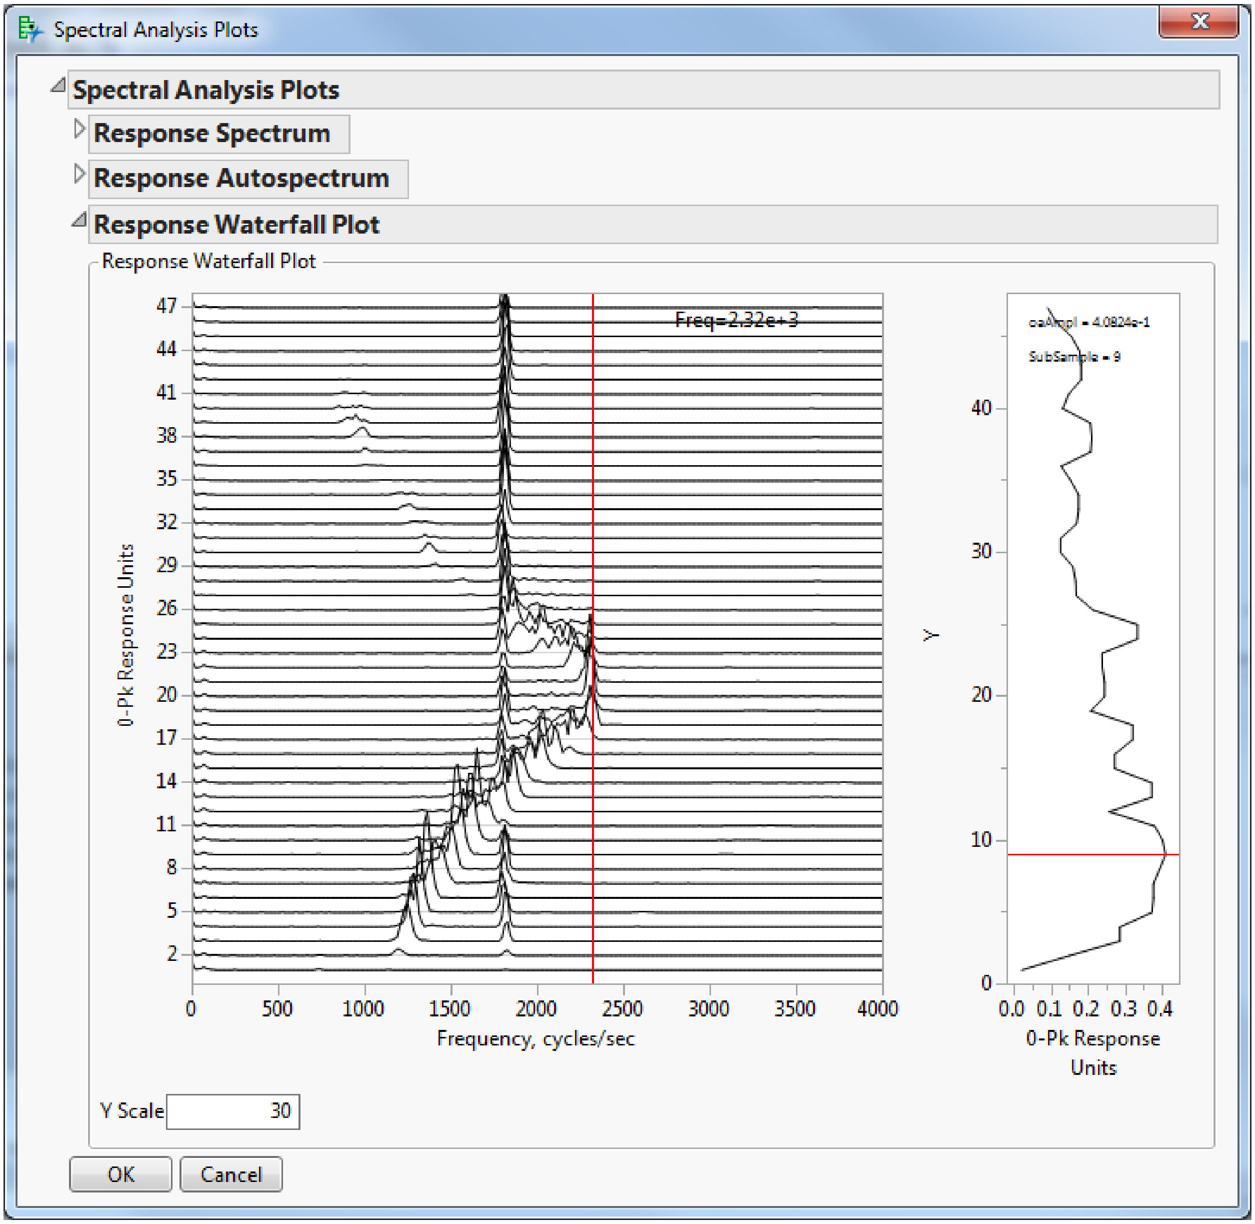

Figure 41: Waterfall Plot for Frequency Varying Input Signal

Future Enhancement of the Frequency Analysis Script

There are several features and enhancements that are planned for the Frequency Analysis Script. These are:

Amplitude Axis Options

The current script displays spectrum and autospectrum amplitudes in 0-Pk units. Often the user may need to display in RMS or dB units, and may wish to display the amplitude axis and frequency axis with a linear scale.

Additionally, a simple way of normalizing amplitude is to divide amplitude values by the spectral resolution. This allows amplitude comparisons between frequency spectra that may have different numbers of spectral lines.

These options will be offered to the user from a popup menu.

A-Weighting Function

If the input signal is from a calibrated microphone, then the user may wish to apply an A-Weighting Filter to the results. A-Weighting is an approximation of how the human ear hears different frequencies. Below 1000 Hz, our ears perceive that the observed sound pressure level is less than what it actually is. Similarly, above 4000 Hz our ears tend to attenuate the incoming signal. Conversely, between 1000-4000 Hz, our ears are most sensitive, and they actually perceive noise in this range as slightly louder than it actually is.

An A-Weighting function will be added that approximates this human hearing curve, and will be selectable by the user.

Coherence Function

The Autospectrum and Transfer Function plots available in the script currently have lower and upper 95% confidence limits plotted for each frequency (assuming multiple spectra have been averaged). For Transfer Functions, another way to look at how well the output function is correlated to the input function is via the Coherence Function. This frequency-dependent function is akin to the normal R-Squared function, resulting in a value between 0 (no correlation) and 1 (perfect correlation) at each frequency. This will be included as a standard display with the Transfer Function plot.

Power Between Cursors

Sometimes the analyst wishes to know the power under a peak in the autospectrum, but the peak is wider than a single spectral line. For this, we will add a "power between cursors" function, where the user can select two cursor locations surrounding the peak and the script will report the power between the two positions.

Periodogram

A Periodogram is simply frequency spectrum whose frequency axis (cycles/time unit) has been converted to 1/frequency (time units/cycle). This will be included as a standard plot among the current plots.

Additional Windowing Options

There are a number of additional windowing functions besides the Hanning Window that could be added to the script. These will be added as the need arises.

Save data to table

Like most other platforms in JMP, there will be an option to save spectral data to a JMP data table for additional processing or plotting at the user's discretion.

Collect data from external instruments

JMP Scripts have the capability of talking to various instruments to enhance data collection. For instance, a digital oscilloscope might have a USB output that could be collected to the user's PC, so that data could be seamlessly exported from the scope directly to the JMP data table. Since all instruments differ in their communication protocols, this is left as a customization to the current script. Please contact the author of this report if you have interest in this functionality.

Add an "About" Button

This button provides contact info for support, version number of the script being run, etc.

Use Column Names on Plots

Instead of "Response" and "Forcing Function", use the actual column names on the plot labels.

Conclusion

A new Frequency Analysis Script has been written for JMP. The script has been described and demonstrated. Future development features have been outlined.

For additional information about this script, please contact the author

Jerry Fish
President, Experistats LLC
experistats@gmail.com
859-433-6231

About the Author

Jerry Fish is owner and President of Experistats LLC. He is a trained Lean Six Sigma Black Belt, and has many years of experience in statistics and DOE while working for Caterpillar Tractor, Cummins Engine, GE Aircraft Engines, and Lexmark International. He and his wife live in Lexington, KY.

Works Cited

1. Fourier Series. *Wikipedia.* [Online] 8 8, 2014. http://en.wikipedia.org/wiki/Fourier_series.

2. Weisstein, Eric W. "Fourier Series." From MathWorld--A Wolfram Web Resource. http://mathworld.wolfram.com/FourierSeries.html. *Wolfram Mathworld.* [Online]

3. Discrete-Time Signal. *Wikipedia.* [Online] 6 30, 2014. http://en.wikipedia.org/wiki/Discrete-time_signal.

4. Discrete Fourier Series. *Wikipedia.* [Online] 8 12, 2014. http://en.wikipedia.org/wiki/Discrete_Fourier_series.

5. Aliasing. *Wikipedia.* [Online] 8 4, 2014. http://en.wikipedia.org/wiki/Aliasing.

6. Nyquist Frequency. *Wikipedia.* [Online] 8 7, 2014. http://en.wikipedia.org/wiki/Nyquist_frequency.

7. Harry Nyquist. *Wikipedia.* [Online] 8 14, 2014. http://en.wikipedia.org/wiki/Harry_Nyquist.

8. Fast Fourier Transform. *Wikipedia.* [Online] 8 12, 2014. http://en.wikipedia.org/wiki/Fast_Fourier_transform.

9. **JMP, a Business Unit of SAS.** Specialized Models, Version 11. Cary, NC : SAS Publishing, 2013.

10. **Robert F. Steidel, Jr.** *An Introduction to Mechanical Vibrations.* New York : John Wiley & Sons, Inc., 1971. pp. 205-246.