MODAL SPACE

(in our own little world)

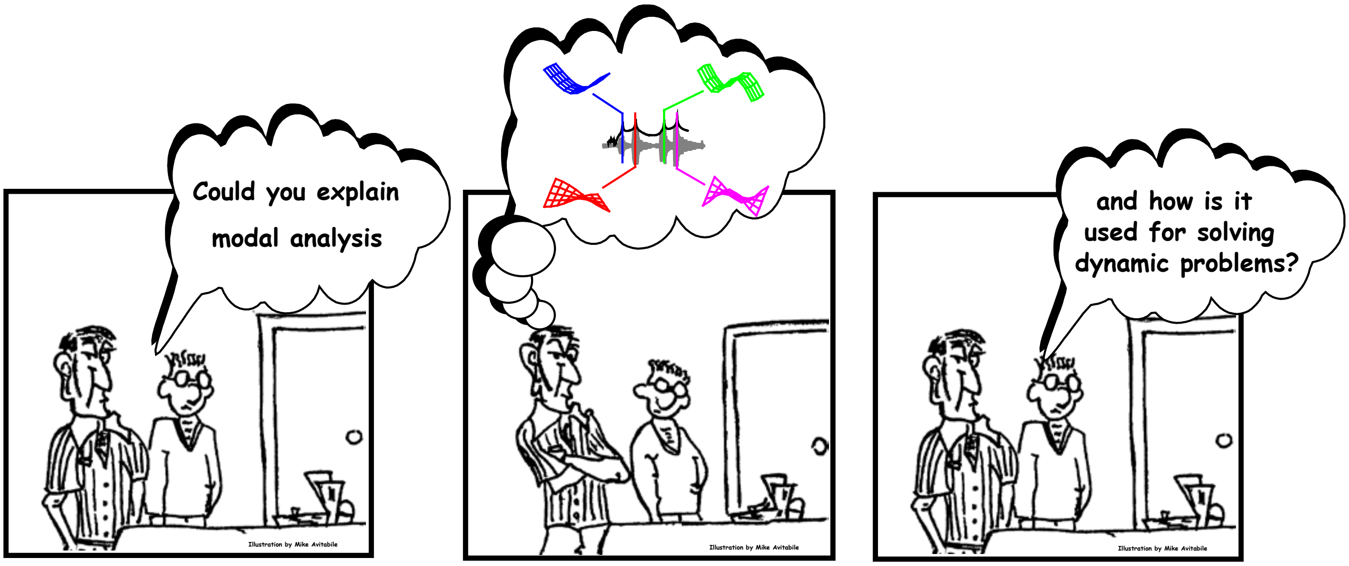

This document contains a collection of articles relating to modal analysis
Experimental Modal Analysis – A Simple Non-Mathematical Overview
(published by Sound & Vibration magazine in January 2001)
Modal Space – Back to Basics (1998 to 2014)
(published in SEM Experimental Techniques)

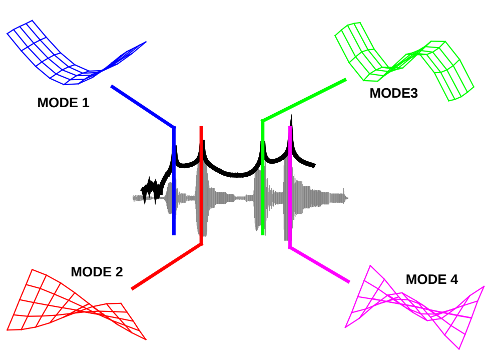

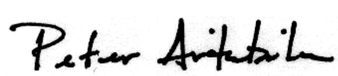

Dr. Peter Avitabile
Professor, Mechanical Engineering
Structural Dynamics and Acoustics Systems Laboratory
University of Massachusetts Lowell
One University Avenue
Lowell, Massachusetts 01854
peter_avitabile@uml.edu
http://sdasl.uml.edu
978-934-3176

EXPERIMENTAL MODAL ANALYSIS

(A Simple Non-Mathematical Presentation)

Peter Avitabile
Modal Analysis and Controls Laboratory
Mechanical Engineering Department
University of Massachusetts Lowell
Lowell, Massachusetts USA

Preface

Often times, people ask some simple questions regarding modal analysis and how structures vibrate. Most times, it is impossible to describe this simply and some of the basic underlying theory needs to be addressed in order to fully explain some of these concepts.

However, many times the theory is just a little too much to handle and some of the concepts can be described without a rigorous mathematical treatment. This document will attempt to explain some concepts about how structures vibrate and the use of some of the tools to solve structural dynamic problems. The intent of this document is to simply identify how structures vibrate from a non-mathematical perspective.

With this being said. Let's start with the first question that is usually asked.

Could you explain modal analysis for me?

In a nutshell, we could say that modal analysis is a process whereby we describe a structure in terms of its natural characteristics which are the frequency, damping and mode shapes - it's dynamic properties. Well that's a mouthful so let's explain what that means. Without getting too technical, I often explain modal analysis in terms of the modes of vibration of a simple plate. This explanation is usually useful for engineers who are new to vibrations and modal analysis.

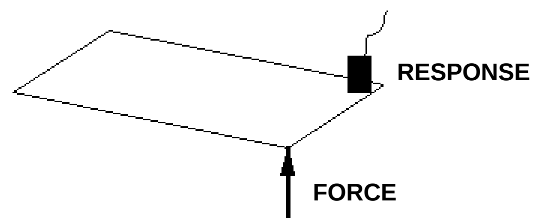

Fig 1 - Simple Plate Excitation/Response Model

Let's consider a freely supported flat plate (Fig 1). Let's apply a constant force to one corner of the plate. We usually think of a force in a static sense which would cause some static deformation in the plate. But here what I would like to do is to apply a force that varies in a sinusoidal fashion. Let's consider a fixed frequency of oscillation of the constant force. We will change the rate of oscillation of the frequency but the peak force will always be the same value - only the rate of oscillation of the force will change. We will also measure the response of the plate due to the excitation with an accelerometer attached to one corner of the plate.

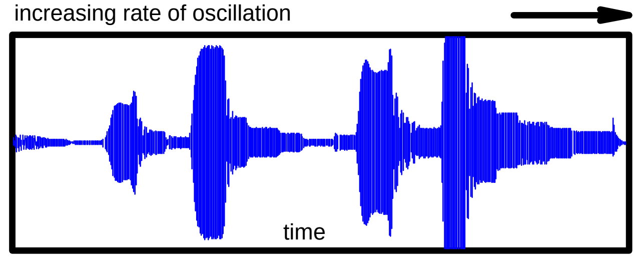

Fig 2 - Simple Plate Response

Now if we measure the response on the plate we will notice that the amplitude changes as we change the rate of oscillation of the input force (Fig 2). There will be increases as well as decreases in amplitude at different points as we sweep up in time. ***This seems very odd*** since we are applying a constant force to the system yet the amplitude varies depending on the rate of oscillation of the input force. But this is exactly what happens - the response amplifies as we apply a force with a rate of oscillation that gets closer and closer to the natural frequency (or resonant frequency) of the system and reaches a maximum when the rate of oscillation is at the resonant frequency of the system. When you think about it, that's pretty amazing since I am applying the same peak force all the time - only the rate of oscillation is changing!

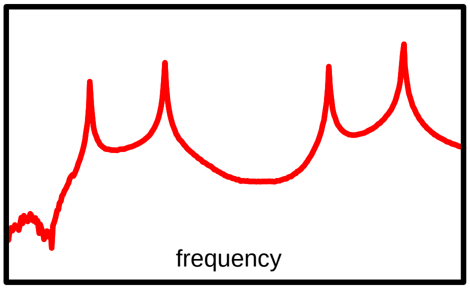

Fig 3 - Simple Plate Frequency Response Function

This time data provides very useful information. But if we take the time data and transform it to the frequency domain using the Fast Fourier Transform then we can compute something called the frequency response function (Fig 3). Now there are some very interesting items to note. We see that there are peaks in this function which occur at the resonant frequencies of the system. And we notice that these peaks occur at frequencies where the time response was observed to have maximum response corresponding to the rate of oscillation of the input excitation.

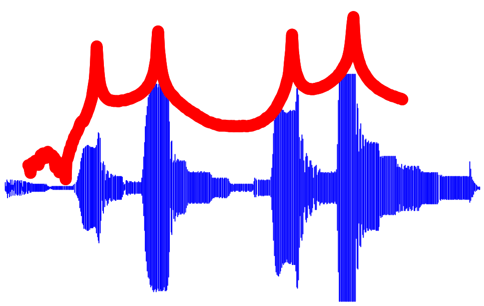

Fig 4 - Overlay of Time and Frequency Response Function

Now if we overlay the time trace with the frequency trace what we will notice is that the frequency of oscillation at the time at which the time trace reaches it's maximum value corresponds to the frequency where peaks in the frequency response function reach a maximum (Fig 4). So you can see that we can use either the time trace to determine the frequency at which maximum amplitude increases occur or the frequency response function to determine where these natural frequencies occur. Clearly the frequency response function is easier to evaluate.

Now most people are amazed at how the structure has these natural characteristics. Well, what's more amazing is that the deformation patterns at these natural frequencies also take on a variety of different shapes depending on which frequency is used for the excitation force.

Now let's see what happens to the deformation pattern on the structure at each one of these natural frequencies. Let's place 45 evenly distributed accelerometers on the plate and measure the amplitude of the response of the plate with different excitation frequencies. If we were to dwell at each one of the frequencies - each one of the natural frequencies - we would see a deformation pattern that exists in the structure (Fig 5). The figure shows the deformation patterns that will result when the excitation coincides with one of the natural frequencies of the system. We see that when we dwell at the first natural frequency, there is a first bending deformation pattern in the plate shown in blue (mode 1). When we dwell at the second natural frequency, there is a first twisting deformation pattern in the plate shown in red (mode 2). When we dwell at the third and fourth natural frequencies, the second bending and second twisting deformation patterns are seen in green (mode 3) and magenta (mode 4), respectively. These deformation patterns are referred to as the mode shapes of the structure. (That's not actually perfectly correct from a pure mathematical standpoint but for the simple discussion here, these deformation patterns are very close to the mode shapes, from a practical standpoint.)

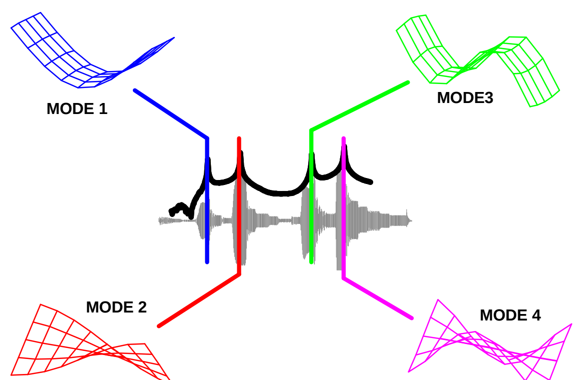

Fig 5 - Simple Plate Sine Dwell Response

Now these natural frequencies and mode shapes occur in all structures that we design. Basically, there are characteristics that depend on the weight and stiffness of my structure which determine where these natural frequencies and mode shapes will exist. As a design engineer, I need to identify these frequencies and know how they might affect the response of my structure when a force excites the structure. Understanding the mode shape and how the structure will vibrate when excited helps the design engineer to design better structures. Now there is much more to it all but this is just a very simple explanation of modal analysis.

So, basically, modal analysis is the study of the natural characteristics of structures. Understanding both the natural frequency and mode shape helps to design my structural system for noise and vibration applications. We use modal analysis to help design all types of structures including automotive structures, aircraft structures, spacecraft, computers, tennis rackets, golf clubs, ... the list just goes on and on.

Now we have introduced this measurement called a frequency response function but exactly what is it?

Just what are these measurements that are called FRFs?

The frequency response function is very simply the ratio of the output response of a structure due to an applied force. We measure both the applied force and the response of the structure due to the applied force simultaneously. (The response can be measured as displacement, velocity or acceleration.) Now the measured time data is transformed from the time domain to the frequency domain using a Fast Fourier Transform algorithm found in any signal processing analyzer and computer software packages.

Due to this transformation, the functions end up being complex valued numbers; the functions contain real and imaginary components or magnitude and phase components to describe the function. So let's take a look at what some of the functions might look like and try to determine how modal data can be extracted from these measured functions.

Let's first evaluate a simple beam with only 3 measurement locations (Fig 6). We see the beam below with 3 measurement locations and 3 mode shapes. There are 3 possible places forces can be applied and 3 possible places where the response can be measured. This means that there are a total of 9 possible *complex valued* frequency response functions that could be acquired; the frequency response functions are usually described with subscripts to denote the input and output locations as $h_{out,in}$ (or with respect to typical matrix notation this would be $h_{row,column}$)

The figure shows the magnitude, phase, real and imaginary parts of the frequency response function matrix. (Of course, I am assuming that we remember that a complex number is made up of a real and imaginary part which can be easily converted to magnitude and phase. Since the frequency response is a complex number, we can look at any and all of the parts that can describe the frequency response function.)

Now let's take a look at each of the measurements and make some remarks on some of the individual measurements that could be made.

First let's drive the beam with a force from an impact at the tip of the beam at point 3 and measure the response of the beam at the same location (Fig 7). This measurement is referred to as h_{33}. This is a special measurement referred to as a drive point measurement. Some important characteristics of a drive point measurement are

- all resonances (peaks) are separated by anti-resonances
- the phase looses 180 degrees of phase as we pass over a resonance and gains 180 degrees of phase as we pass over an anti-resonance
- the peaks in the imaginary part of the frequency response function must all point in the same direction

So as I continue and take a measurement by moving the impact force to point 2 and measuring the response at point 3 and then moving the impact force on to point 1 to acquire two more measurements as shown. (And of course I could continue on to collect any or all of the additional input-output combinations.)

So now we have some idea about the measurements that we could possibly acquire. One important item to note is that the frequency response function matrix is symmetric. This is due to the fact that the mass, damping and stiffness matrices that describe the system are symmetric. So we can see that $h_{ij} = h_{ji}$ - this is called reciprocity. So we don't need to actually measure all the terms of the frequency response function matrix.

One question that always seems to arise is whether or not it is necessary to measure all of the possible input-output combinations and why is it possibly to obtain mode shapes from only one row or column of the frequency response function matrix.

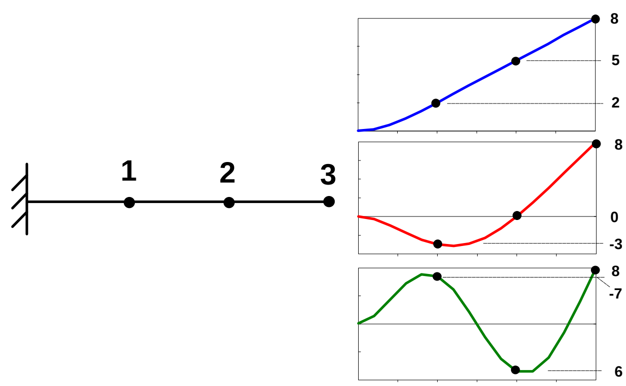

Fig 6a – Beam 3 DOF Model

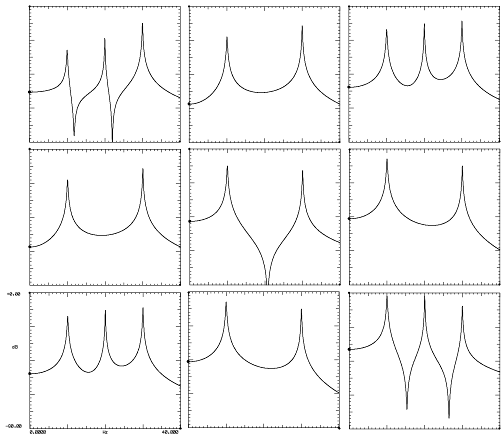

Fig 6b - Magnitude

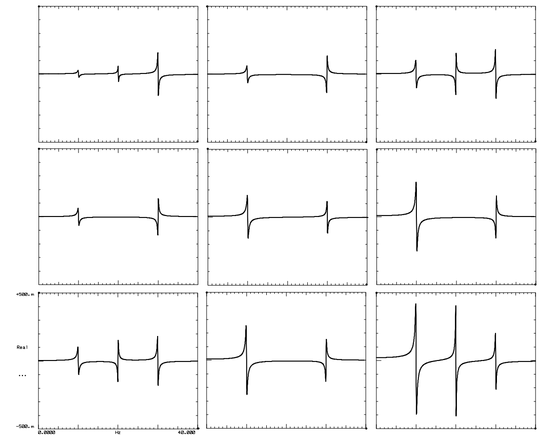

Fig 6d - Real

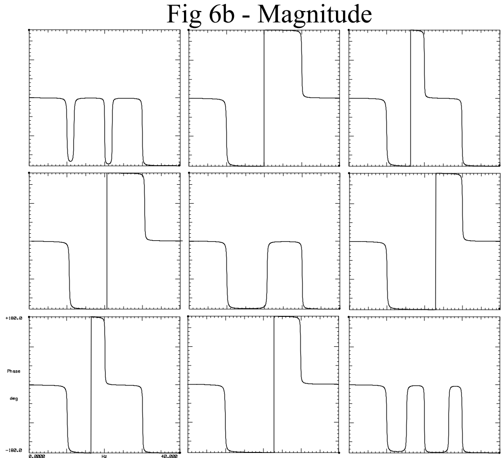

Fig 6c - Phase

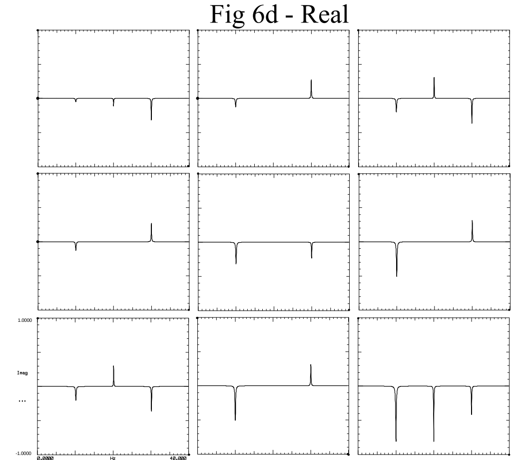

Fig 6e - Imaginary

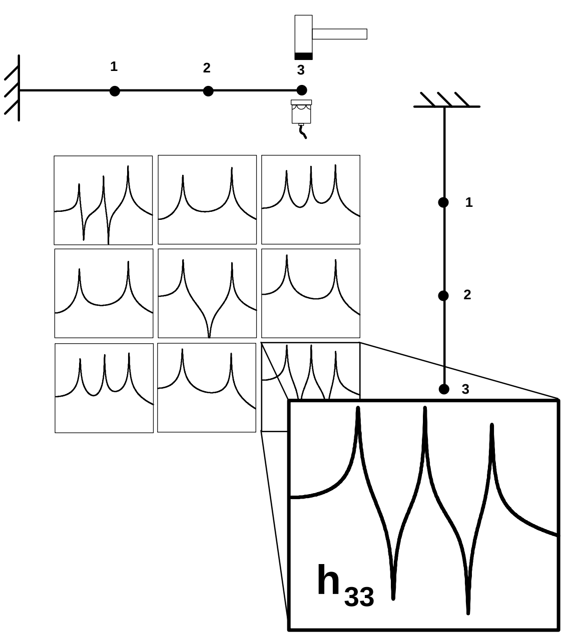

Fig 7a – Drive Point FRF for Reference 3

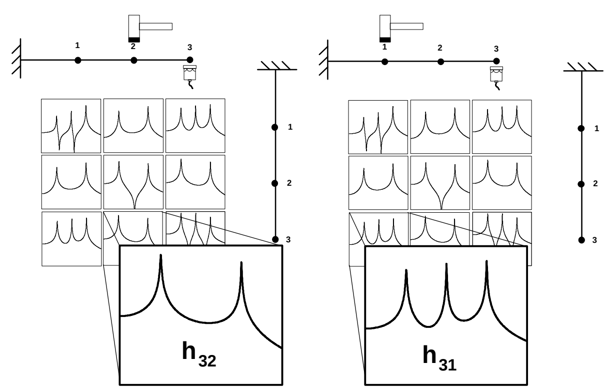

Fig 7 – Cross FRF s for Reference 3

Why is only one row or column of the FRF matrix needed?

It is very important for us to understand how we arrive at mode shapes from the various measurements that are available in the frequency response function matrix. Without getting mathematical, let's discuss this.

Let's just take a look at the third row of the frequency response function matrix and concentrate on the first mode. If I look at the peak amplitude of the imaginary part of the frequency response function, I can easily see that the first mode shape for mode 1 can be seen (Fig 8a). So it seems fairly straightforward to extract the mode shape from measured data. A quick and dirty approach is just to measure the peak amplitude of the frequency response function for a number of different measurement points.

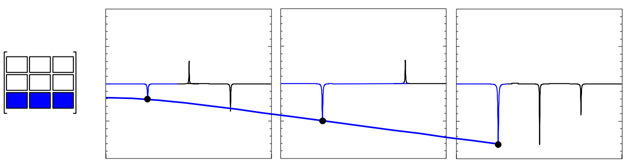

Fig 8a – Mode 1 from Third Row of FRF Matrix

Now look at the second row of the frequency response function matrix and concentrate on the first mode (Fig 8b). If I look at the peak amplitude of the imaginary part of the frequency response function, I can easily see that the first mode shape for mode 1 can be seen from this row also.

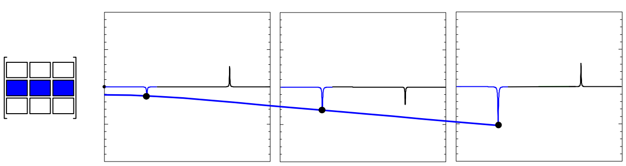

Fig 8a – Mode 1 from Second Row of FRF Matrix

We could also look at the first row of the frequency response function matrix and see the same shape. This is a very simple pictorial representation of what the theory indicates. We can use any row to describe the mode shape of the system. So it is very obvious that the measurements contain information pertaining to the mode shapes of the system.

Let's now take a look at the third row again and concentrate on mode 2 now (Fig 8c). Again if I look at the peak amplitude of the imaginary part of the frequency response function, I can easily see the second mode shape for mode 2 can be seen.

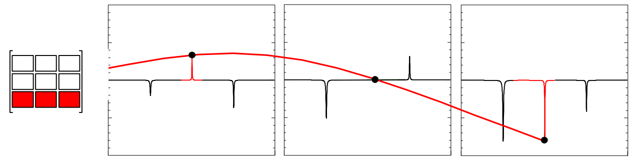

Fig 8a – Mode 2 from Third Row of FRF Matrix

And if I look at the second row of the frequency response function matrix and concentrate on the second mode, I will be a little surprised because there is no amplitude for the second mode (Fig 8d). I wasn't expecting this but if we look at the mode shape for the second mode then we can quickly see that this is a node point for mode 2. The reference point is located at the node of the mode.

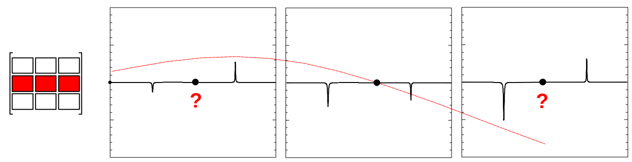

Fig 8d – Mode 2 from Second Row of FRF Matrix

So this points out one very important aspect of modal analysis and experimental measurements. The reference point cannot be located at the node of a mode otherwise that mode will not be seen in the frequency response function measurements and the mode cannot be obtained.

Now we have only used 3 measurement points to describe the modes for this simple beam. If we add more input-output measurement locations then the mode shapes can be seen more clearly as shown in Figure 9. The figure shows 15 measured frequency response functions and the 3 measurement points used in the discussion above are highlighted. This figure shows the 15 frequency response functions in a waterfall style plot. Using this type of plot, it is much easier to see that the mode shapes can be determined by looking at the peaks of the imaginary part of the frequency response function.

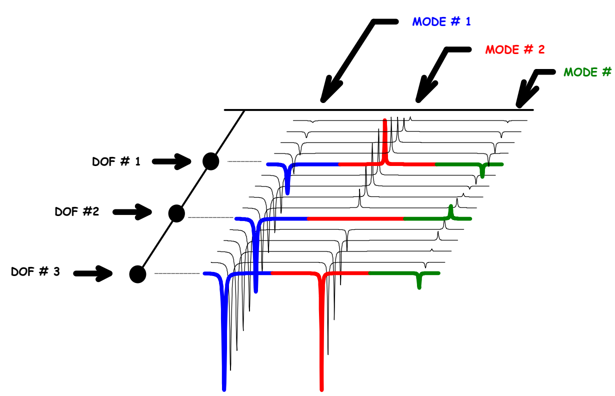

Fig 9 - Waterfall Plot of Beam Frequency Response Functions

Now the measurements that we have discussed thus far have been obtained from an impact testing consideration. What if the measured frequency response functions come from a shaker test?

What's the difference between a shaker test and an impact test?

From a theoretical standpoint, it doesn't matter whether the measured frequency response functions come from a shaker test or an impact test. Figures 10a and 10b show the measurements that are obtained from an impact test and a shaker test. *An impact test generally results in measuring one of the rows of the frequency response function matrix whereas the shaker test generally results in measuring one of the columns of the frequency response function matrix.* Since the system matrices describing the system are square symmetric, then reciprocity is true. For the case shown, the third row is exactly the same as the third column, for instance.

Theoretically, there is no difference between a shaker test and an impact test. That is, from a theoretical standpoint! If I can apply pure forces to a structure without any interaction between the applied force and the structure and I can measure response with a massless transducer that has no effect on the structure - then this is true. But what if this is not the case?

Now let's think about performing the test from a practical standpoint. The point is that shakers and response transducers generally do have an effect on the structure during the modal test. The main item to remember is that the structure under test is not just the structure for which you would like to obtain modal data. It is the structure plus everything involved in the acquisition of the data - the structure suspension, the mass of the mounted transducers, the potential stiffening effects of the shaker/stinger arrangement, etc. So while theory tells me that there shouldn't be any difference between the impact test results and the shaker test results, often there will be differences due to the practical aspects of collecting data.

The most obvious difference will occur from the roving of accelerometers during a shaker test. The weight of the accelerometer may be extremely small relative to the total weight of the whole structure, but its weight may be quite large relative to the effective weight of different parts of the structure. This is accentuated in multi-channel systems where many accelerometers are moved around the structure in order to acquire all the measurements. This can be a problem especially on light-weight structures. One way to correct this problem is to mount all of the accelerometers on the structure even though only a few are measured at a time. Another way is to add dummy accelerometer masses at locations not being measured; this will eliminate the roving mass effect.

Another difference that can result is due to the shaker/stinger effects. Basically, the modes of the structure may be affected by the mass and stiffness effects of the shaker attachment. While we try to minimize these effects, they may exist. The purpose of the stinger is to uncouple the effects of the shaker from the structure. However, on many structures, the effects of the shaker attachment may be significant. Since an impact test does not suffer from these problems, different results may be obtained. So while theory says that there is no difference between a shaker test and an impact test, there are some very basic practical aspects that may cause some differences.

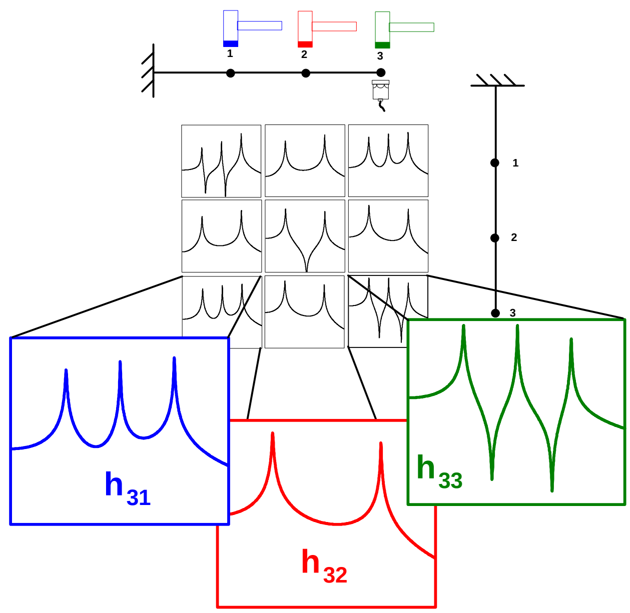

Fig 10a - Roving Impact Test Scenario

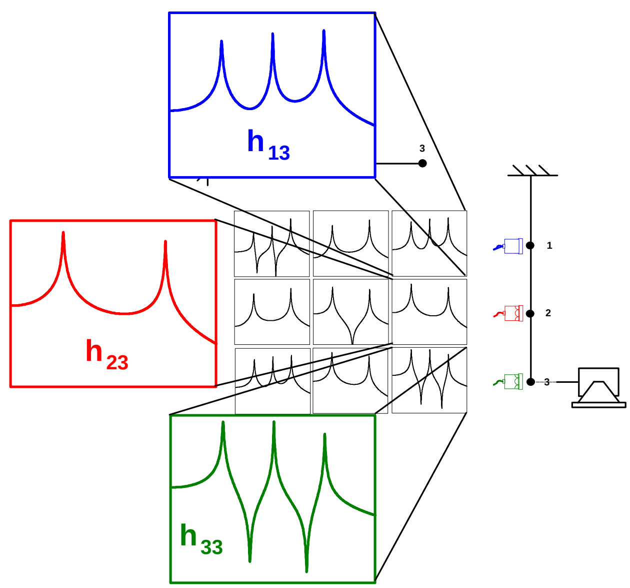

Fig 10b - Roving Response Test Scenario

What measurements do I actually make to compute the FRF?

The most important measurement that is needed for experimental modal analysis is the frequency response function. Very simply stated, this is the ratio of the output response to the input excitation force. This measurement is typically acquired using a dedicated instrument such as an FFT (Fast Fourier Transform) analyzer or a data acquisition system with software that performs the FFT.

Let's briefly discuss some of the basic steps that occur in the acquisition of data to obtain the FRF. First, there are analog signals that are obtained from our measuring devices. These analog signals must be filtered to assure that there is no aliasing of higher frequencies into the analysis frequency range. This is usually done through the use of a set of analog filters on the front-end of the analyzer called anti-aliasing filters. Their function is to remove any high frequency signals that may exist in the signal.

The next step is to digitize the analog signal to form a digital representation of the actual signal. This is done by the analog to digital converter that is called the ADC. Typically this digitization process will use 10, 12 or 16 bit converters; the more bits available, the better the resolution possible in the digitized signal. Some of the major concerns lie in the sampling and quantization errors that could potentially creep into the digitized approximation. Sampling rate controls the resolution in the time and frequency representation of the signals. Quantization is associated with the accuracy of magnitude of the captured signal. Both sampling and quantization can cause some errors in the measured data but are not nearly as significant and devastating as the worst of all the signal processing errors – leakage!

Leakage occurs from the transformation of time data to the frequency domain using the Fast Fourier Transform (FFT). The Fourier Transform process requires that the sampled data consist of a complete representation of the data for all time or contain a periodic repetition of the measured data. When this is satisfied, then the Fourier Transform produces a proper representation of the data in the frequency domain. However, when this is not the case, then leakage will cause a serious distortion of the data in the frequency domain. In order to minimize the distortion due to leakage, weighting functions called windows are used to cause the sampled data to appear to better satisfy the periodicity requirement of the FFT. While windows greatly reduces the leakage effect, it cannot be completely removed.

Once the data is sampled, then the FFT is computed to form linear spectra of the input excitation and output response. Typically, averaging is performed on power spectra obtained from the linear spectra. The main averaged spectra computed are the input power spectrum, the output power spectrum and the cross spectrum between the output and input signals.

These functions are averaged and used to compute two important functions that are used for modal data acquisition – the frequency response function (FRF) and the coherence. The coherence function is used as a data quality assessment tool which identifies how much of the output signal is related to the measured input signal. The FRF contains information regarding the system frequency and damping and a collection of FRFs contain information regarding the mode shape of the system at the measured locations. This is the most important measurement related to experimental modal analysis. An overview of these steps described is shown in Figure 11.

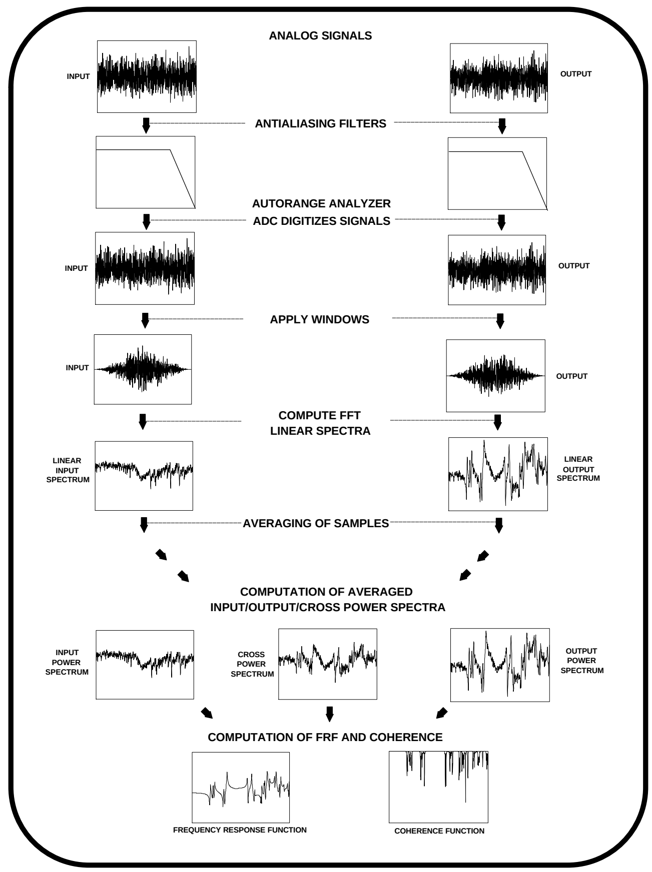

Fig 11 - Anatomy of an FFT Analyzer

Of course, there are many important aspects of measurement acquisition, averaging techniques to reduce noise and so on, that are beyond the scope of this presentation. Any good reference on digital signal processing will provide assistance in this area. Now the input excitation needs to be discussed next. Basically, there are two commonly used types of excitation for experimental modal analysis – impact excitation and shaker excitation.

Now let's consider some of the testing considerations when performing an impact test.

What are the biggest things to think about when impact testing?

There are many important considerations when performing impact testing. Only two of the most critical items will be mentioned here; a detailed explanation of all the aspects pertaining to impact testing is far beyond the scope of this paper.

First, the selection of the hammer tip can have a significant effect on the measurement acquired. The input excitation frequency range is controlled mainly by the hardness of the tip selected. The harder the tip, the wider the frequency range that is excited by the excitation force. The tip needs to be selected such that all the modes of interest are excited by the impact force over the frequency range to be considered. If too soft a tip is selected, then all the modes will not be excited adequately in order to obtain a good measurement as seen Figure 12a. The input power spectrum does not excite all of the frequency range shown as evidenced by the rolloff of the power spectrum; the coherence is also seen to deteriorate as well as the frequency response function over the second half of the frequency range.

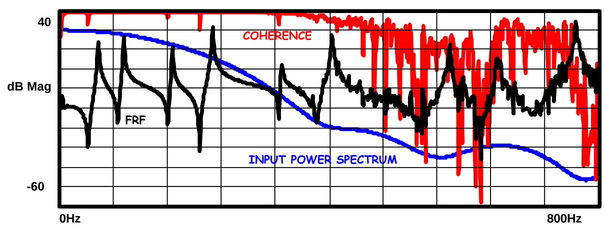

Fig 12a - Hammer Tip Not Sufficient to Excite All Modes

Typically, we strive to have a fairly good and relatively flat input excitation forcing function. The frequency response function is measured much better as evidenced by the much improved coherence function. When performing impact testing, care must be exercised to select the proper tip so that all the modes are excited well and a good frequency response measurement is obtained.

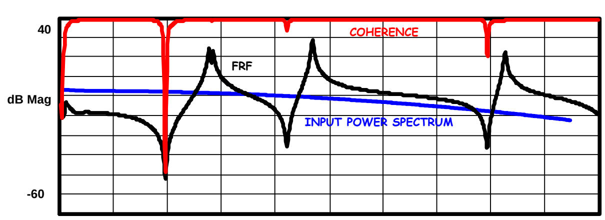

Fig 12b - Hammer Tip Adequate to Excite All Modes

The second most important aspect of impact testing relates to the use of an impact window for the response transducer. Generally for lightly damped structures, the response of the structure due to the impact excitation will not die down to zero by the end of the sample interval. When this is the case, the transformed data will suffer significantly from a digital signal processing effect referred to as leakage.

In order to minimize leakage, a weighting function referred to as a window is applied to the measured data. This window is used to force the data to better satisfy the periodicity requirements of the Fourier Transform process, thereby minimizing the distortion effects of leakage. For impact excitation, the most common window used on the response transducer measurement is the exponentially decaying window. The implementation of the window to minimize leakage is shown in Figure 13.

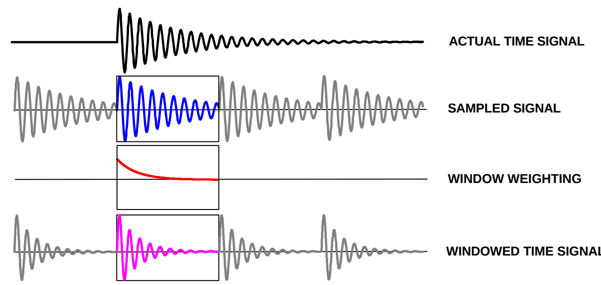

Fig 13 - Exponential Window to Minimize Leakage Effects

Windows cause some distortion of the data themselves and should be avoided whenever possible. For impact measurements, two possible items to always consider are the selection of a narrower bandwidth for measurements and to increase the number of spectral lines of resolution. Both of these signal processing parameters have the effect of increasing the amount of time required to acquire a measurement. These will both tend to reduce the need for the use of an exponential window and should always be considered to reduce the effects of leakage.

Now let's consider some of the testing considerations when performing a shaker test.

What are the biggest things to think about when shaker testing?

Again, there are many important items to consider when performing shaker testing but the most important of those center around the effects of excitation signals that minimize the need for windows or eliminate the need for windows altogether. There are many other considerations when performing shaker testing but a detailed explanation of all of these is far beyond the scope of this presentation.

One of the more common excitation techniques still used today is random excitation due to its ease of implementation. However, due to the nature of this excitation signal, leakage is a critical concern and the use of a Hanning window is commonly employed. This leakage effect is serious and causes distortion of the measured frequency response function even when windows are used. A typical random excitation signal with a Hanning window is shown in Figure 14. As seen in the figure, the Hanning window weighting function helps to make the sampled signal appear to better satisfy the periodicity requirement of the FFT process, thereby minimizing the potential distortion effects of leakage.

While this improves the distortion of the FRF due to leakage, the window will never totally remove these effects; the measurements will still contain some distortion effects due to leakage.

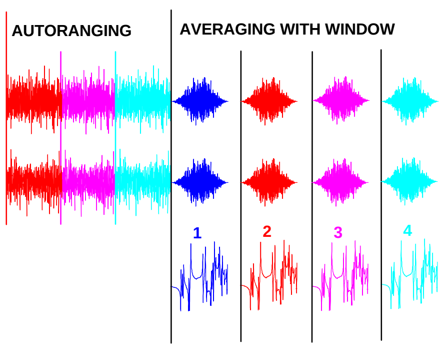

Fig 14 - Shaker Testing – Excitation Considerations
Random Excitation with Hanning window

Two very common excitation signals widely used today are burst random and sine chirp. Both of these excitations have a special characteristic that do not require the need for windows to be applied to the data since the signal are inherently leakage free in almost all testing situations. These excitations are relatively simple to employ and are commonly found on most signal analyzers available today. These two signals are shown schematically in Figure 15 and 16.

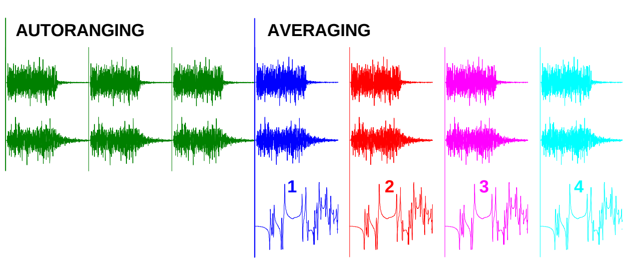

Fig 15 - Burst Random Excitation Without a Window

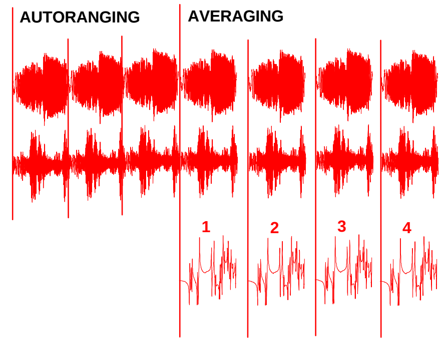

Fig 16 - Sine Chirp Excitation Without a Window

In the case of burst random, the periodicity requirement of the FFT process is satisfied due to the fact that the entire transient excitation and response are captured in one sample interval of the FFT. For the sine chirp excitation, the repetition of the signal over the sample interval satisfies the periodicity requirement of the FFT process. While other excitation signals also exist, these are the most common excitation signals used in modal testing today.

So now we have a better idea how to make some measurements.

Windows are, in many measurement situations, a necessary evil. While I would rather not have to use any windows at all, the alternative of leakage is definitely not acceptable either. As discussed above, there are a variety of excitation methods that can be successfully employed which will provide leakage free measurements and do not require the use of any window. However, there are many times, especially when performing field testing and collecting operating data, where the use of windows is necessary. So what are the most common windows typically used.

Basically, in a nutshell, the most common widows used today are the Uniform, Hanning, Flat Top and Force/Exponential windows. Rather than detail all the windows, let’s just simply state when each are used for experimental modal testing.

The Uniform Window (which is also referred to as the Rectangular Window, Boxcar or No Window) is basically a unity gain weighting function that is applied to all the digitized data points in one sample or record of data. This window is applied to data where the entire signal is captured in one sample or record of data or when the data is guaranteed to satisfy the periodicity requirement of the FFT process. This window can be used for impact testing where the input and response signals are totally observed in one sample of collected data. This window is also used when performing shaker excitation tests with signals such as burst random, sine chirp, pseudo-random and digital stepped sine; all of these signals generally satisfy the periodicity requirement of the FFT process.

The Hanning window is basically a cosine shaped weighting function (bell shaped) that forces the beginning and end of the sample interval to be heavily weighted to zero. This is useful for signals that generally do not satisfy the periodicity requirement of the FFT process. Random excitations and general field signals usually fall into this category and require the use of a window such as the Hanning window.

The Flat Top window is most useful for sinusoidal signals that do not satisfy the periodicity requirement of the FFT process. Most often this window is used for calibration purposes more than anything else in experimental modal analysis.

The force and exponential windows are typically used when performing impact excitation for acquiring FRFs. Basically, the force window is a unity gain window that acts over a portion of the sample interval where the impulsive excitation occurs. The exponential window is used when the response signal does not die out within the sample interval. The exponential window is applied to force the response to better satisfy the periodicity requirement of the FFT process.

Each of the windows has an effect of the frequency representation of the data. In general, the windows will cause a degradation in the accuracy of the peak amplitude of the function and will appear to have more damping than what really exists in the actual measurement. While these errors are not totally desirable, they are far more acceptable than the significant distortion that can result from leakage.

So how do I get mode shapes from the plate FRFs?

So now that we have discussed various aspects of acquiring measurements, let's go back to the plate structure previously discussed and take several measurements on the structure. Let's consider 6 measurement locations on the plate. Now there are 6 possible places where forces can be applied to the plate and 6 possible places where we can measure the response on the plate. This means that there are a total of 36 possible input output measurements that could be made. The frequency response function describes how the force applied to the plate causes the plate to respond. If we applied a force to point 1 and measured the response at point 6, then the transfer relation between 1 and 6 describes how the system will behave (Figure 17).

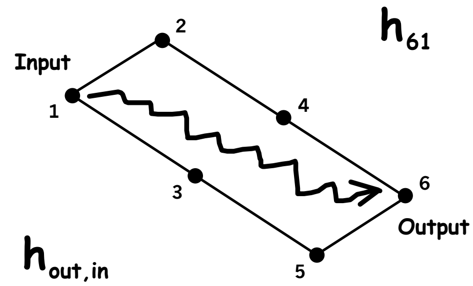

Fig 17 - Input-Output Measurement Locations

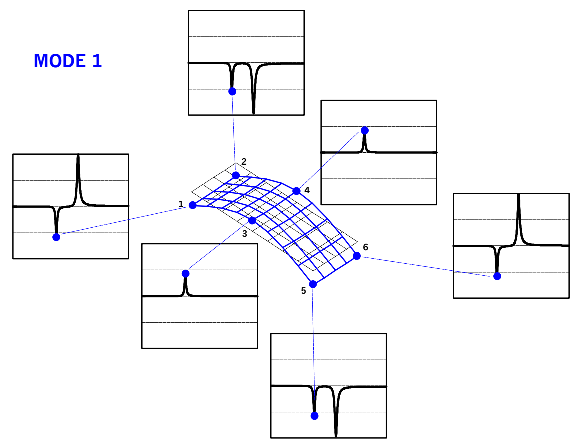

Fig 18 - Plate Mode Shapes For Mode 1- Peak Pick of FRF

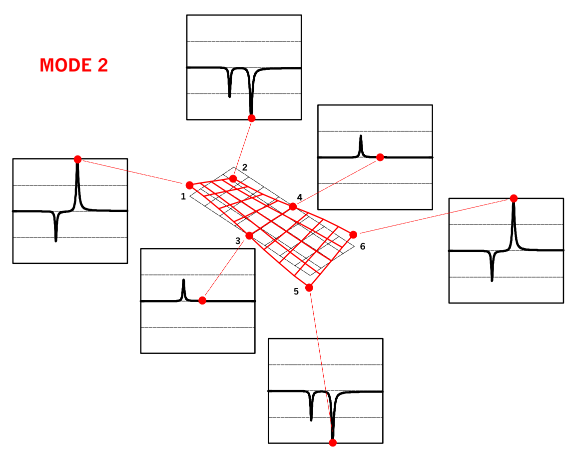

Fig 19 - Plate Mode Shapes For Mode 2 - Peak Pick of FRF

While the technique shown above is adequate for very simple structures, mathematical algorithms are typically used to estimate the modal characteristics from measured data. The modal parameter estimation phase, which is often referred to as curvefitting, is implemented using computer software to simplify the extraction process. The basic parameters that are extracted from the measurements are the frequency, damping and mode shapes – the dynamic characteristics. The measured FRF is basically broken down into many single DOF systems as shown in Figure 20.

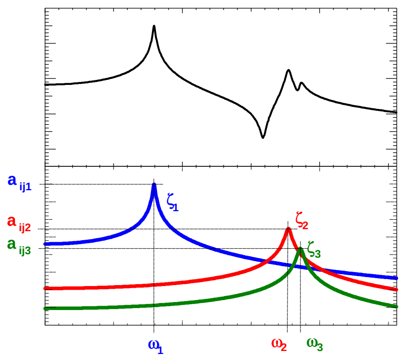

Fig 20 - Breakdown of a Frequency Response Function

These curvefitting techniques use a variety of different methods to extract data. Some techniques employ time domain data while others use frequency domain data. The most common methods employ multiple mode analytical models but at times very simple single mode methods will produce reasonably good results for most engineering analyses (Figure 21). Basically, all of the estimation algorithms attempt to break down measured data into the principal components that make up the measured data – namely the frequency, damping and mode shapes.

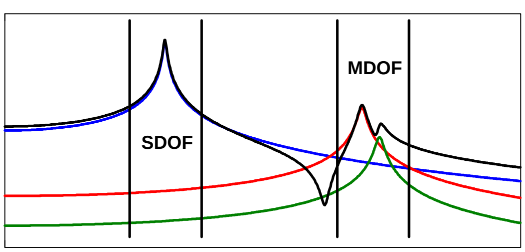

Fig 21 - Curvefitting Different Bands using Different Methods

The key inputs that the analyst must specify are the band over which data is extracted, the number of modes contained in the data and the inclusion of residual compensation terms for the estimation algorithm. This is schematically shown in Figure 22.

Much more could be said concerning the estimation of modal parameters from measured data, the tools available for deciphering data and the validation of the extracted model but a detailed explanation is far beyond the scope of this paper.

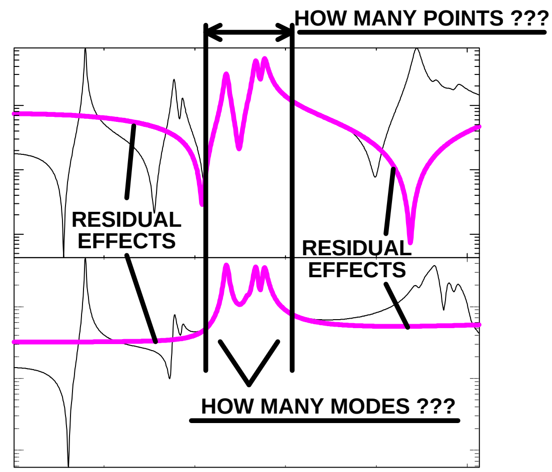

Fig 22 - Curvefitting a Typical FRF

All structures respond to externally applied forces. But many times the forces are not known or cannot be measured easily. We could still measure the response of a structural system even though the forces may not be measured. So the next question that is often asked concerns operating data.

What is operating data?

We first need to recognize that the system responds to the forces that are applied to the system (whether or not I can measure them). So for explanation purposes, let's assume for now that we know what the forces are. While the forces are actually applied in the time domain, there are some important mathematical advantages to describing the forces and response in the frequency domain. For a structure which is exposed to an arbitrary input excitation, the response can be computed using the frequency response function multiplied by the input forcing function. This is very simply shown in Figure 23.

The excitation shown is a random excitation that excites all frequencies. The most important thing to note is that the frequency response function acts as a filter on the input force which results in some output response. The excitation shown causes all the modes to be activated and therefore, the response is, in general, the linear superposition of all the modes that are activated by the input excitation. Now what would happen if the excitation did not contain all frequencies but rather only excited one particular frequency (which is normally what we are concerned about when evaluating most operating conditions).

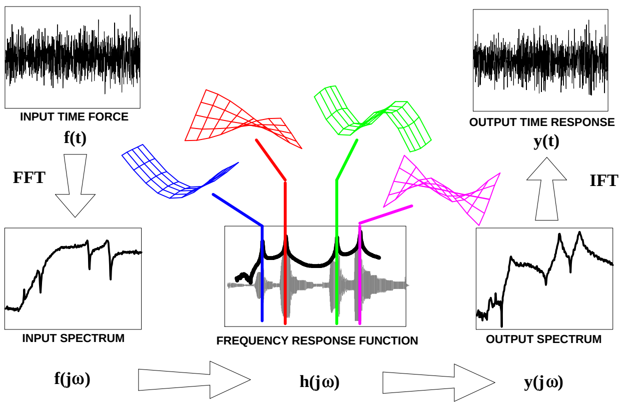

Fig 23 - Schematic Overviewing the Input-Output Structural Response Problem

To illustrate this, let's use the simple plate that we just discussed. Let's assume that there is some operating condition that exists for the system; a fixed frequency operating unbalance will be considered to be the excitation. It seems reasonable to use the same set of accelerometers that were on the plate to measure the response of the system. If we acquire data, we may see something that looks like the deformation pattern shown in Figure 24. Looking at this deformation, it is not very clear why the structure responds this way or what to do to change the response. Why does the plate behave in such a complicated fashion anyway??? This doesn't appear to be anything like any of the mode shapes that we measured before.

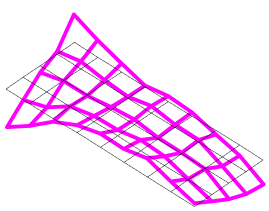
Fig 24 - Measured Operating Displacements

In order to understand this, let's take that plate and apply a simple sinusoidal input at one corner of the plate. For the example here, we are only going to consider the response of the plate assuming that there are only 2 modes that are activated by the input excitation. (Of course there are more modes, but let's keep it simple to start.) Now we realize that the key to determining the response is the frequency response function between the input and output locations. Also, we need to remember that when we collect operating data, we don't measure the input force on the system and we don't measure the system frequency response function - we only measure the response of the system.

First let's excite the system with a sinusoid that is right at the first natural frequency of the plate structure. The response of the system for one frequency response function is shown in Figure 25. So even though we excite the system at only one frequency, we know that the frequency response function is the filter that determines how the structure will respond. We can see that the frequency response function is made up of a contribution of both mode 1 and mode 2. We can also see that the majority of the response, whether it be in the time or frequency domain, is dominated by mode 1. Now if we were to measure the response only at that one frequency and measure the response at many points on the structure, then the operating deflection pattern would look very much like mode 1 - but there is a small contribution due to mode 2. Remember that with operating data, we never measure the input force or the frequency response function - we only measure the output response. So that the deformations that are measured are the *actual response* of the structure due to the input excitation - whatever it may be.

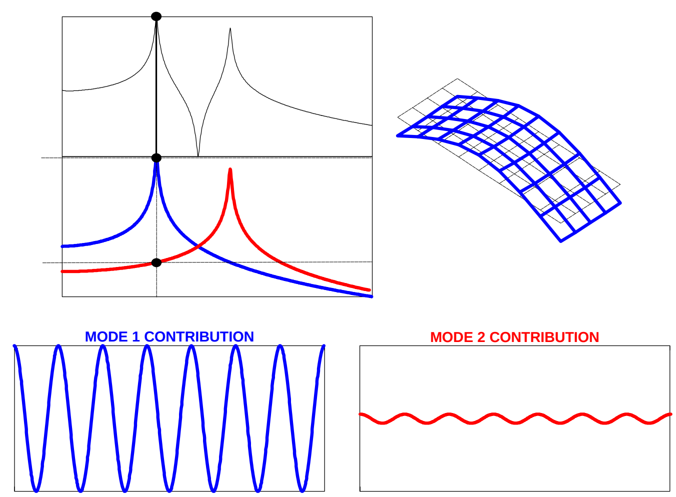

Fig 25 - Excitation Close to Mode 1

When we measure frequency response functions and estimate modal parameters, we actually determine the contribution to the total frequency response function solely due to the effects of mode 1 acting alone, as shown in blue, and mode 2 acting alone, as shown in red, and so on for all the other modes of the system. Notice that with operating data, we only look at the response of the structure at one particular frequency - which is the linear combination of all the modes that contribute to the total response of the system. So we can now see that the operating deflection pattern will look very much like the first mode shape if the excitation primarily excites mode one.

Now let's excite the system right at the second natural frequency. Figure 26 below shows the same information as just discussed for mode 1. But now we see that we primarily excite the second mode of the system. Again, we must realize that the response looks like mode 2 - but there is a small contribution due to mode 1.

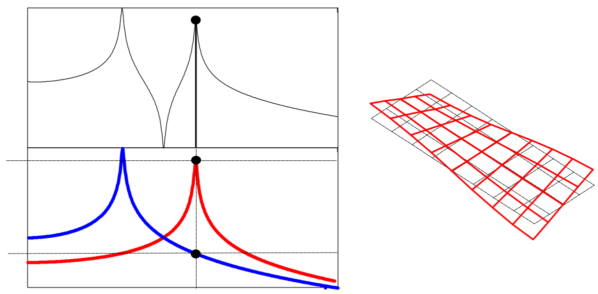

Fig 26 - Excitation Close to Mode 2

But what happens when we excite the system away from a resonant frequency. Let's excite the system at a frequency midway between mode 1 and mode 2. Now here is where we see the real difference between modal data and operating data. The next figure shows the deformation shape of the structure.

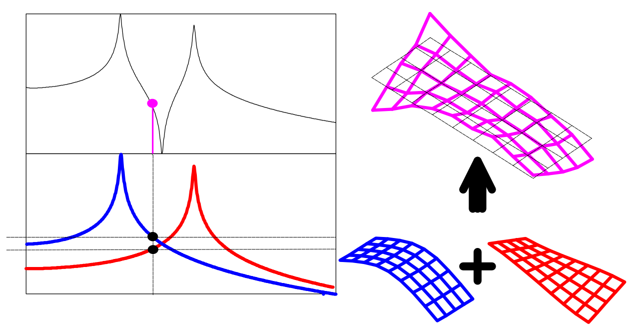

Fig 27 - Excitation Somewhere Between Mode 1 and Mode 2

At first glance, it appears that the deformation doesn't look like anything that we recognize. But if we look at the deformation pattern long enough, we can actually see a little bit of first bending and a little bit of first torsion in the deformation. So the operating data is primarily some combination of the first and second mode shapes. (Yes, there will actually be other modes but primarily mode 1 and 2 will be the major participants in the response of the system.)

Now, we have discussed all of this by understanding the frequency response function contribution on a mode by mode basis. When we actually collect operating data, we don't collect frequency response functions but rather we collect output spectrums. If we looked at those, it would not have been very clear at to why the operating data looked like mode shapes. Figure 28 shows a resulting output spectrum that would be measured at one location on the plate structure. Now the input applied to the structure is much broader in frequency and many modes are excited. But, by understanding how each of the modes contributes to the operating data, it is much easier to see how the modes all contribute to the total response of the system.

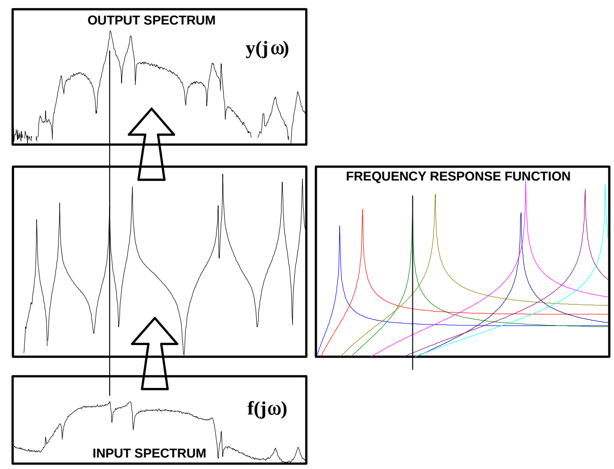

Fig 28 - Broadband Plate Excitation

So actually, there is a big difference between operating deflections and mode shapes - *we can now see that the modes shapes are summed together in some linear fashion to form the operating deflection patterns*. But typically we are interested in the total deformation or response of the system. Why do I even want to bother to collect modal data? It seems like a lot of work to acquire measurements and extract data.

So what good is modal data?

Modal data is an extremely useful piece of information that can assist in the design of almost any structure. The understanding and visualization of mode shapes is invaluable in the design process. It helps to identify areas of weakness in the design or areas where improvement is needed. The development of a modal model is useful for simulation and design studies. One of these studies is structural dynamic modification.

This is a mathematical process which uses modal data (frequency, damping and mode shapes) to determine the effects of changes in the system characteristics due to physical structural changes. These calculations can be performed without actually having to physically modify the actual structure until a suitable set of design changes is achieved. A schematic of this is shown in Figure 29. There is much more that could be discussed concerning structural dynamic modification but space limitations restrict this.

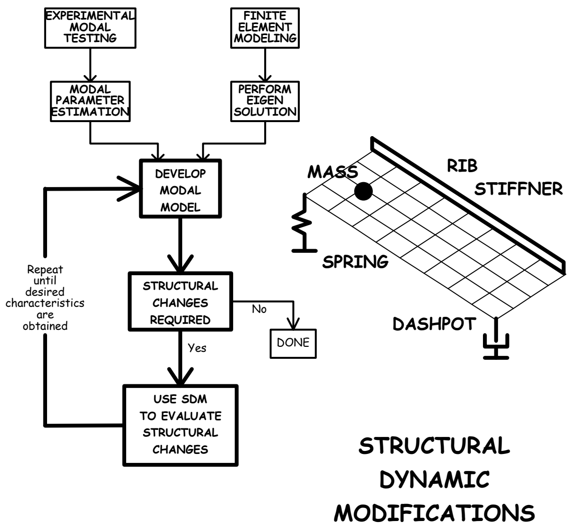

Fig 29 – Schematic of the SDM Process

In addition to structural dynamic modification studies, other simulations can be performed such as force response simulation to predict system response due to applied forces. And another very important aspect of modal testing is the correlation and correction of an analytical model such as a finite element model. These are a few of the more important aspects relating to the use of a modal model which are schematically shown in Figure 30.

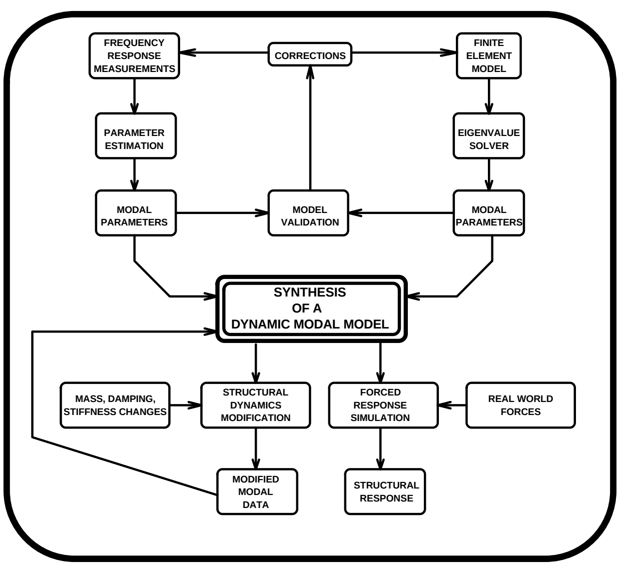

Fig 30 – Overall Dynamic Modeling Process

And one of the final questions that is often asked is which test is best to perform

Of course with tight schedules and budgets, do I really need to collect both modal data and operating data? This is always difficult to answer but it is always better to have both whenever possible. If only one of the two is available, then many times some engineering decisions may be made without full knowledge of the system characteristics. To summarize, let's point out the differences between each of the data sets.

Modal data requires that the force is measured in order to determine the frequency response function and resulting modal parameters. Only modal data will give the true principal characteristics of the system. In addition, structural dynamic modifications and forced response can only be studied using modal data (operating data cannot be used for these types of studies). Also correlation with a finite element model is best performed using modal data. But of course it needs to be clearly stated that modal data alone does not identify whether a structure is adequate for an intended service or application since modal data is independent of the forces on the system.

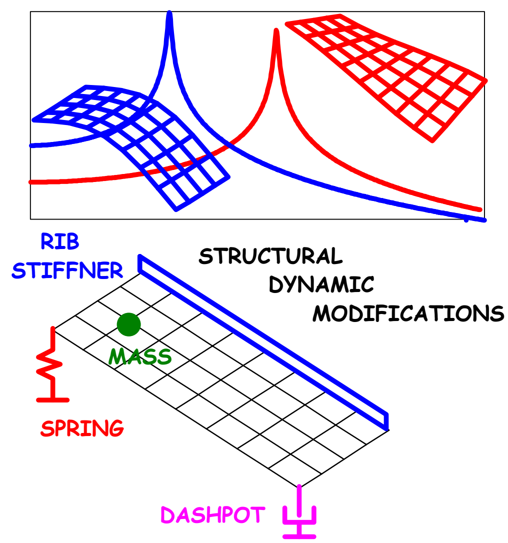

Fig 31 - Modal Model Characteristics

Operating data on the other hand is an actual depiction of how the structure behaves in service. This is extremely useful information. However, many times the operating shapes are confusing and do not necessarily provide clear guidance as to how to solve or correct an operating problem (and modification and response tools cannot be utilized on operating data).

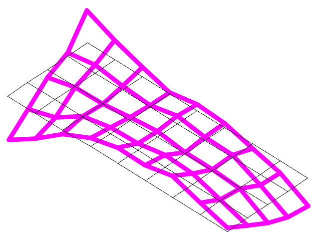

Fig 32 - Operating Data Characteristics

The best situation exists when both operating data and modal data are used in conjunction to solve structural dynamic problems.

Summary

Some simple explanations were used to describe structural vibration and the use of some of the available tools for the solution of structural dynamic problems. This was all achieved without the use of any detailed mathematical relationships. In order to better understand more of the details of the data presented here, a theoretical treatment of this material is necessary.

References

1) Lecture Notes, Peter Avitabile, Modal Analysis I & II, University of Massachusetts Lowell
2) Seminar Presentation Notes, Peter Avitabile
3) The Modal Handbook, A Multimedia Computer Based Training and Reference Guide, Dynamic Decisions, Merrimack, NH, info@dynamicdecisions.com

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Could you explain modal analysis for me?
Well...it will take a little bit but here's one that anyone can understand.

You're not the first one to ask me to explain modal analysis is simple terms so anyone can understand it. In a nutshell, we could say that modal analysis is a process whereby we describe a structure in terms of it's natural characteristics which are the frequency, damping and mode shapes - it's dynamic properties. Well that's a mouthful so let's explain what that means. Without getting too technical, I often explain modal analysis in terms of the modes of vibration of a simple plate. This explanation is usually useful for engineers who are new to vibrations and modal analysis.

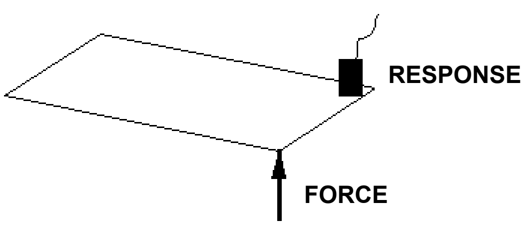

Let’s consider a freely supported flat plate. Let's apply a constant force to one corner of the plate. We usually think of a force in a static sense which would cause some static deformation in the plate. But here what I would like to do is to apply a force that varies in a sinusoidal fashion. Let's consider a fixed frequency of oscillation of the constant force. We will change the rate of oscillation of the frequency but the peak force will always be the same value - only the rate of oscillation of the force will change. We will also measure the response of the plate due to the excitation with an accelerometer attached to one corner of the plate.

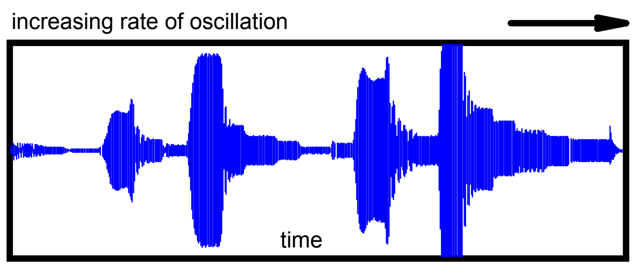

Now if we measure the response on the plate we will notice that the amplitude changes as we change the rate of oscillation of the input force. There will be increases as well as decreases in amplitude at different points as we sweep up in time. ***This seems very odd*** since we are applying a constant force to the system yet the amplitude varies depending on the rate of oscillation of the input force. But this is exactly what happens - the response amplifies as we apply a force with a rate of oscillation that gets closer and closer to the natural frequency (or resonant frequency) of the system and reaches a maximum when the rate of oscillation is at the resonant frequency of the system. When you think about it, that's pretty amazing since I am applying the same peak force all the time - only the rate of oscillation is changing!

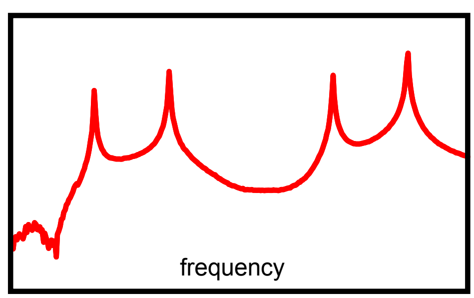

Copyright 1998
Pete Avitabile

This time data provides very useful information. But if we take the time data and transform it to the frequency domain using the Fast Fourier Transform then we can compute something called the frequency response function. Now there are some very interesting items to note. We see that there are peaks in this function which occur at the resonant frequencies of the system. Now we notice that these peaks occur at frequencies where the time response was observed to have maximum response corresponding to the rate of oscillation of the input excitation.

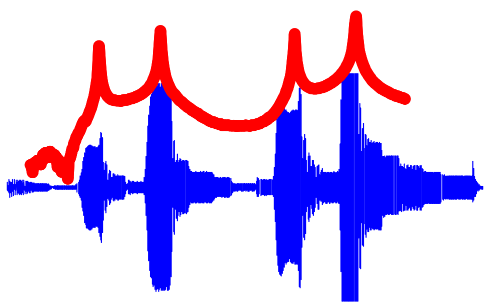

Now if we overlay the time trace with the frequency trace what we will notice is that the frequency of oscillation at the time at which the time trace reaches it's maximum value corresponds to the frequency where peaks in the frequency response function reach a maximum. So you can see that we can use either the time trace to determine the frequency at which maximum amplitude increases occur or the frequency response function to determine where these natural frequencies occur. Clearly the frequency response function is easier to evaluate.

You thought it was pretty amazing how the structure has these natural characteristics. Well, the deformation patterns at these natural frequencies also take on a variety of different shapes depending on which frequency is used for the excitation force.

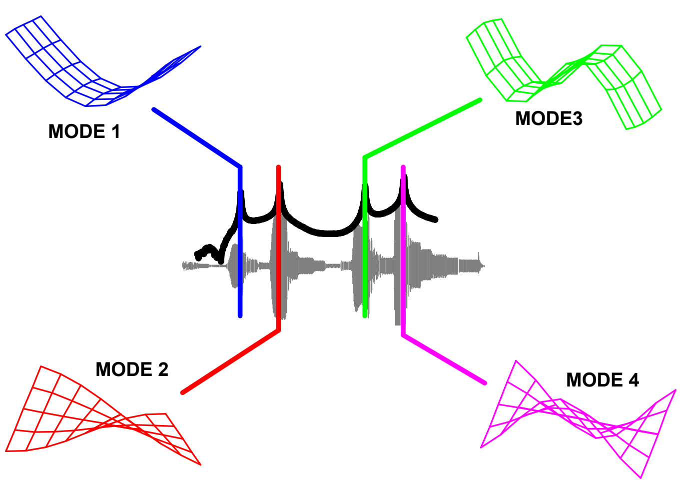

Now let's see what happens to the deformation pattern on the structure at each one of these natural frequencies. Let's place 45 evenly distributed accelerometers on the plate and measure the amplitude of the response of the plate with different excitation frequencies. If we were to dwell at each one of the frequencies - each one of the natural frequencies - we would see a deformation pattern that exists in the structure. The figure shows the deformation patterns that will result when the excitation coincides with one of the natural frequencies of the system. We see that when we dwell at the first natural frequency, there is a first bending deformation pattern in the plate shown in blue. When we dwell at the second natural frequency, there is a first twisting deformation pattern in the plate shown in red. When we dwell at the third and fourth natural frequencies, the second bending and second twisting deformation patterns are seen in green and magenta, respectively. These deformation patterns are referred to as the mode shapes of the structure. (That's not actually perfectly correct from a pure mathematical standpoint but for the simple discussion here, these deformation patterns are very close to the mode shapes, for all practical purposes.)

Now these natural frequencies and mode shapes occur in all structures that we design. Basically, there are characteristics that depend on the weight and stiffness of my structure which determine where these natural frequencies and mode shapes will exist. As a design engineer, I need to identify these frequencies and know how they might affect the response of my structure when a force excites the structure. Understanding the mode shape and how the structure will vibrate when excited helps the design engineer to design better structures. Now there is much more to it all but this is just a very simple explanation of modal analysis.

Now we can better understand what modal analysis is all about - it is the study of the natural characteristics of structures. Both the natural frequency and mode shape (which depends on the mass and stiffness distributions in my structure) are used to help design my structural system for noise and vibration applications. We use modal analysis to help design all types of structures including automotive structures, aircraft structures, spacecraft, computers, tennis rackets, golf clubs, ... the list just goes on and on.

I hope this very brief introduction helps to explain what modal analysis is all about. I know I explained modal analysis to my Mom using the example above and I think for the first time she actually understood what I actually do. Since then, she has been heard explaining it to her friends using a variety of words closely resembling *modal analysis*, of which the best one was the time she referred to it as *noodle analysis* ... but that's another story!

Copyright 1998
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Could you explain the difference between time domain, frequency domain and modal space?
I hear it all the time but I'm not sure what's the difference.
There's a lot to explain but let's start with something simple.

This question gets asked often. There's a lot of different aspects relating to this so let's start with a simple explanation without using too much math and explain all of this with a simple schematic. Let's use the figure to discuss all these different aspects of the time domain, frequency domain, modal space and physical space. Now there are a lot of parts to discuss in the figure, so let's take them in pieces - one at a time - and then summarize everything at the end. You might also want to remember the discussion we had before when you asked me about what modal analysis was all about ("Could you explain modal analysis for me?") to help with the discussion here.

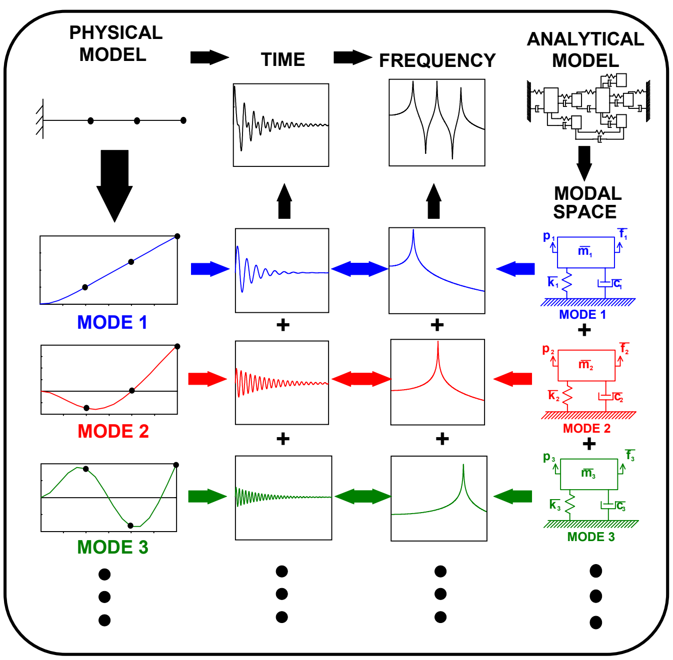

First, let's consider a simple cantilever beam and imagine that the beam is excited by a pulse at the tip of the beam. The response at the tip of the beam will contain the response of all the modes of the system (shown in the black time response plot); notice that there appears to be response at several different frequencies. This time response at the tip of the beam can be converted to the frequency domain by performing a Fourier Transform of the time signal. There is a significant amount of math that goes along with this process but it is a common transformation that we perform all the time. The frequency domain representation of this converted time signal is often referred to as the frequency response function, or FRF for short (shown in the black frequency plot); notice that there are peaks in this plot which correspond to the natural frequencies of the system.

Before we discuss the time and frequency plots any further, let's talk about the physical model in the upper left part of the figure. We know that the cantilever beam will have many natural frequencies of vibration. At each of these natural frequencies, the structural deformation will take on a very definite pattern, called a mode shape, as described previously [1]. For this beam, we see that there is a first bending mode shown in blue, a second bending mode shown in red and a third bending mode shown in green. Of course, there are also other higher modes not shown and we will only discuss the first three modes here but it could easily be extended to higher modes.

Now the physical beam could also be evaluated using an analytical lumped mass model or finite element model (shown in black) in the upper right part of the figure. This model will generally be evaluated using some set of equations where there is an interrelationship, or coupling, between the different points, or degrees of freedom (dof), used to model the structure. This means that if you pull on one of the dofs in the model, the other dofs are also affected and also move. This coupling means that

Copyright 1998
Pete Avitabile

the equations are more complicated in order to determine how the system behaves. As the number of equations used to describe the system get larger and larger, the complication in the equations becomes more involved. We often use matrices to help organize all of the equations of motion describing how the system behaves which looks like

$$[M]\{\ddot{x}\} + [C]\{\dot{x}\} + [K]\{x\} = \{F(t)\}$$

where [M], [C], [K] are the mass, damping and stiffness matrices respectively, along with the corresponding acceleration, velocity and displacement and the force applied to the system. Usually the mass is a diagonal matrix and the damping and stiffness matrices are symmetric with off-diagonal terms indicating the degree of coupling between the different equations or dofs describing the system. The size of the matrices is dependent on the number of equations that we use to describe our system. Mathematically, we perform something called an eigensolution and use the modal transformation equation to convert these coupled equations into a set of uncoupled single dof systems described by diagonal matrices of modal mass, modal damping and modal stiffness in a new coordinate system called modal space described as

$$\begin{bmatrix} \diagdown & & \\ & \overline{M} & \\ & & \diagdown \end{bmatrix}\{\ddot{p}\} + \begin{bmatrix} \diagdown & & \\ & \overline{C} & \\ & & \diagdown \end{bmatrix}\{\dot{p}\} + \begin{bmatrix} \diagdown & & \\ & \overline{K} & \\ & & \diagdown \end{bmatrix}\{p\} = [U]^T\{F\}$$

So we can see that the transformation from physical space to modal space using the modal transformation equation is a process whereby we convert a complicated set of coupled physical equations into a set of simple uncoupled single dof systems. And we see in the figure that the analytical model can be broken down into a set of single dof systems where the single dof describing mode 1 is shown in blue, mode 2 is shown in red and mode 3 is shown in green. Modal space allows us to describe the system easily using simple single dof systems.

Now let's go back to the time and frequency responses shown in black. We know that the total response can be obtained from the contribution of each of the modes. The total response shown in black comes from the summation of the effects of the response of the model shown in blue for mode 1, red for mode 2 and green for mode 3. This applies whether I describe the system in the time domain or the frequency domain. Each domain is equivalent and just presents the data from a different viewpoint. It's a lot like money - as I go from country to country, the money in each country looks different but it's really the same thing. So we can see that the total time response is made up of the part of the time response due to the contribution of the time response of mode 1 shown in blue, mode 2 in red and mode 3 in green. We can also see that the total FRF is made up of the part of the FRF due to the contribution of the FRF of mode 1 shown in blue, mode 2 in red and mode 3 in green. (We have only shown the magnitude part of the FRF here; this function is actually complex which is correctly displayed using both magnitude and phase or real and imaginary parts of the FRF).

Since we can break the analytical model up into a set of single dof systems, we could determine the FRF for each of the single dof systems as shown with mode 1 in blue, mode 2 in red, and mode 3 in green. We could also determine the time response for each of these single dof systems through a closed form solution for the response of a single dof system due to the pulse input or we could simply inverse Fourier Transform the FRF for each of the single dof systems. We could also measure the response of the beam at the tip due to the pulse and filter the response of each of the modes of the system, and we would see the response of each of the modes of the system with mode 1 shown in blue, mode 2 in red and mode 3 in green. (Of course, I'm simplifying a lot of theory here so we can understand the concepts.)

Now that we have pulled apart all the pieces of the figure, I think it should be much clearer that there is really no difference between the time domain, frequency domain, modal space and physical space. Each domain is just a convenient way for presenting or viewing data. However, sometimes one domain is much easier to see things than another domain. For instance, the total time response does not clearly identify how many modes there are contributing to the response of the beam. But the total FRF in the frequency domain is much clearer in showing how many modes are activated and the frequency of each of the modes. So often we transform from one domain to another domain simply because the data is much easier to interpret.

While there is a lot more to it all, I hope this simple schematic and explanation helps to put everything in better perspective. Think about it and if you have any more questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

Illustration by Mike Avitabile

Is there any difference between a modal test with a shaker excitation or impact excitation?
Well ... that's a good question.
The answer is yes and no.

This is another question that gets asked often. There's a lot of different aspects relating to this. Let's start with some basics to understand why it is so difficult to answer this question as either yes or no. A few simple equations are needed to help explain this.

First, we have to remember that any system can be described by it's equation of motion. Basically, the equation is simply the force balance of mass times acceleration plus damping times velocity plus stiffness times displacement which is equal to the applied force. For a number of reasons, it is easier to work with this equation in the Laplace domain. By taking the Laplace transform of the equation of motion, we can write

$$\Big[[M]s^2+[C]s+[K]\Big]\{X(s)\}=\{F(s)\}\Rightarrow[B(s)]\{X(s)\}=\{F(s)\}$$

We use matrices to help organize all of the equations. Remember that [M], [C], [K] are the mass, damping and stiffness matrices respectively. It is very important to note that these matrices are symmetric. Therefore, the system matrix, [B(s)], is also symmetric. The system transfer function is the inverse of the system matrix given by

$$[B(s)]^{-1}=[H(s)]=\frac{\mathrm{Adj}[B(s)]}{\det[B(s)]}=\frac{[A(s)]}{\det[B(s)]}$$

And, of course, you remember that the frequency response function that we measure during a modal test is nothing more than the system transfer function evaluated along the frequency axis. Most of the time, we write the frequency response function in partial fraction form, for convenience, as

$$[H(s)]_{s=j\omega}=[H(j\omega)]=\sum_{k=1}^{m}\frac{[A_k]}{(j\omega-p_k)}+\frac{[A_k^*]}{(j\omega-p_k^*)}$$

and an individual term can be written as

$$h_{ij}(j\omega)=\sum_{k=1}^{m}\frac{a_{ijk}}{(j\omega-p_k)}+\frac{a_{ijk}^*}{(j\omega-p_k^*)}$$

So why did I bother writing out all these equations? That's because there are some very important things to note in these equations relative to your question. Remember that [B(s)] and [H(s)] are symmetric since [M], [C], and [K] are symmetric. That means that [H(jω)] is also symmetric. This implies that $h_{ij}=h_{ji}$ which is called reciprocity. This means that you can measure the FRF by impacting point 'i' and measuring the response at point 'j' and get exactly the same FRF as impacting point 'j' and measuring the response at point 'i'. This is what is meant by reciprocity.

Now, let's consider an impact test situation for a simple beam with three measurement locations. There are a total of nine possible input-output FRFs that could be measured. But for this case, let's put our accelerometer at point 3 and make FRF measurements by impacting the beam at point 1, 2, and 3. We call point 3 the reference location since it is the same response point for each of the measurements that I make. Since the hammer is roving from one point to another point, the FRFs that are measured come from one row of the FRF matrix, the last row of the matrix.

Impact Test Measurements

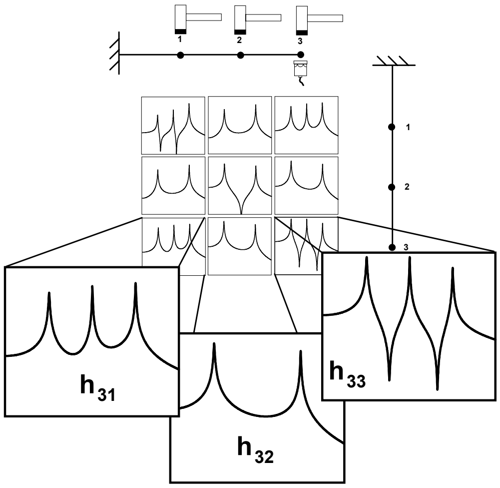

Shaker Test Measurements

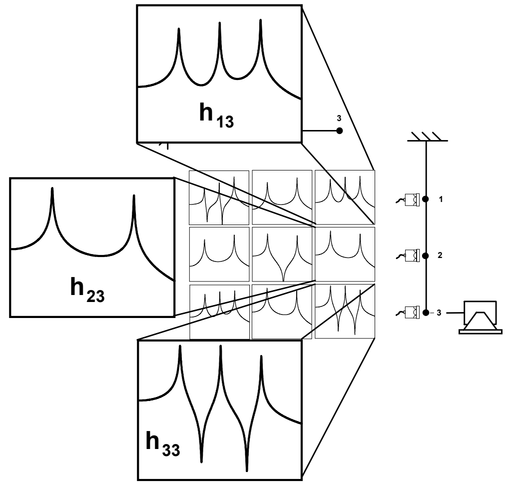

Before we talk about anything else, let's discuss the same set of measurements from a shaker test. Let's place our shaker at point 3 and make FRF measurements by roving the accelerometer to point 1, 2, and 3 on the beam; note that point 3 is still the reference location since the force is applied to the same point for each measurement. Now that the force is stationary, the FRFs that are measured come from one column of the FRF matrix, the last column of the matrix.

If I look at the measurements taken, I'll notice that h_{13} from the shaker test is *exactly the same* as h_{31} from the impact test. Also notice that h_{23} from the shaker test is *exactly the same* as h_{32} from the impact test. Well, this is what reciprocity is all about. So, from a theoretical standpoint, it doesn't matter whether I collect data from a shaker test or an impact test. The data is exactly the same - from a theoretical standpoint. If fact, there is no reason why the impact test can't be performed by impacting the same point on the structure and roving the accelerometer around to all the different measurement locations. I could draw the same analogy for the shaker test also. We could keep the response accelerometer at the same location and move the shaker from point to point (but I don't know anyone who wants to run a test that way!) The point is that from a theoretical standpoint, it doesn't matter how the data is collected as long as the input-output characteristics are obtained.

So the answer is that there is no difference between a shaker test and an impact test. That is, from a theoretical standpoint! If I can apply pure forces to a structure without any interaction between the applied force and the structure and I can measure response with a massless transducer that has no effect on the structure - then this is true. But what if this is not the case.

Now let's think about performing the test from a practical standpoint. The point is that shakers and response transducers generally do have an effect on the structure during the modal test. The main item to remember is that the structure under test is not just the structure that you would like to obtain modal data. It is the structure plus everything involved in the acquisition of the data - the structure suspension, the mass of the mounted transducers, the potential stiffening effects of the shaker/stinger arrangement, etc. So while theory tells me that there shouldn't be any difference between the impact test results and the shaker test results, often there will be differences due to the practical aspects of collecting data.

The most obvious difference will occur from the roving of accelerometers during a shaker test. The weight of the accelerometer may be extremely small relative to the total weight of the whole structure, but it's weight may be quite large relative to the effective weight of different parts of the structure. This is accentuated in multi-channel systems where many accelerometers are moved around the structure in order to acquire all the measurements. This can be a problem especially on light weight structures. One way to correct this problem is to mount all of the accelerometers on the structure even though only a few are measured at a time. Another way is to add dummy accelerometer masses at locations not being measured; this will eliminate the roving mass effect.

Another difference that can result is due to the shaker/stinger effects. Basically, the modes of the structure may be affected by the mass and stiffness effects of the shaker attachment. While we try to minimize these effects, they may exist. The purpose of the stinger is to divorce the effects of the shaker from the structure. However, on many structures, the effects of the shaker attachment may be significant. Since an in impact test does not suffer from these problems, different results may be obtained.

So while theory says that there is no difference between a shaker test and an impact test, there are some very basic practical aspects that may cause some differences. I hope this clears up this question.

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

Illustration by Mike Avitabile

Is there a difference between a roving hammer and roving accelerometer test?
Well ... it depends
Let's explain what the differences could be.

Basically, there is no difference between a roving hammer and roving accelerometer modal test. This is true providing the same measurements are collected. Let me explain by discussing this seemingly simple but tricky fine point about a modal test.

Back when we performed a modal test with a 2 channel analyzer, it was fairly straightforward to perform an impact test. Usually, the hammer roved around the structure with a stationary accelerometer. Typically, we impacted the structure at every point in the x, y, and z directions to obtain FRFs relative to the reference location of the stationary accelerometer. But when we started using multichannel analyzers to perform the same test, there are some slight differences that need to be addressed. Let's consider an impact test for the 9 points shown on the structure. Let's also assume that I have an impact hammer and a tri-axial accelerometer with a 4 channel FFT analyzer or acquisition system.

MODAL TEST CONFIGURATION

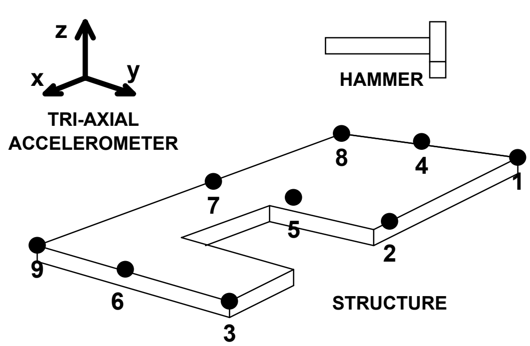

One way to run the test is to place the tri-axial accelerometer at a fixed location and impact, in one direction, at all 9 points. We would then obtain 27 FRFs for the structure. Another way to run the test is to impact at one point, in one direction, and have the tri-axial accelerometer rove to all 9 points. Again we would collect 27 FRFs. So in both cases, we measure 27 FRFs by impacting in only one direction.

But are the two tests the same? At first glance, you would think that both test setups should produce the same results. In order to confirm whether this is true or not, let's step through the measurement process and list out what measurements are actually being made for each test setup.

Test Setup #1

Let's say that we want to run a modal test shown in setup #1. In this test, the tri-axial accelerometer is stationary at point 9 and measures x, y, and z outputs. The input hammer force is applied in the z direction only and roves to each of the 9 points shown.

Now let's list each of the FRFs that will be collected from this test setup. When we impact point 1 in the z direction, the response is measured at 9x, 9y, and 9z. So the FRFs measured are 9x/1z, 9y/1z, 9z/1z for the first measurement made. Next we impact point 2 in the z direction and the response is measured at 9x, 9y, and 9z. This set of FRFs are 9x/2z, 9y/2z, 9z/2z. We can continue on here but I think you get the hang of it. But what did we actually measure? Let's arrange all of these measurements in the FRF matrix to see what we have.

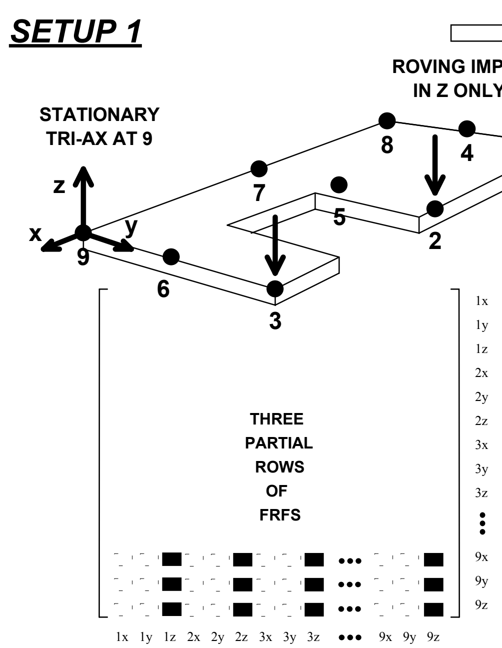

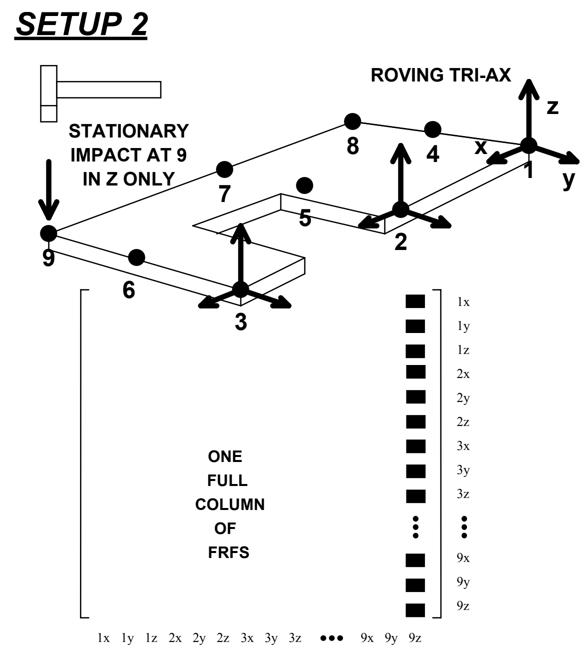

When we take a close look at the FRF matrix, we notice that we have measured only parts of three different rows of this matrix. So we only have three partial descriptions of the characteristic of the system. But in each of the partial descriptions, we can only see the characteristic information in the z direction. This would be fine if there was only motion in the z direction. But what if there was significant motion in the z direction when the structure is excited in the x direction? We have only measured response due to excitation in the z direction!

Test Setup #2

Now let's say that we also want to run the modal test shown in setup #2. In this test, the hammer impacts only in the z direction at point 9. The tri-axial accelerometer roves to each of the 9 points shown for this test, measuring the x, y, and z directions.

Let's list each of the FRFs that will be collected from this test setup. When we impact point 9 in the z direction, the response is measured at 1x, 1y, and 1z. So the FRFs measured are 1x/9z, 1y/9z, 1z/9z for the first measurement made. Next we move the accelerometer to point 2 and the response is measured at 2x, 2y, and 2z. This set of FRFs are 2x/9z, 2y/9z, 2z/9z. So what did we actually measure? Again, let's arrange all of these measurements in the FRF matrix to see what we have.

Now we notice that we have measured one complete column of the FRF matrix. Now we can describe the response of the system in a more complete sense. We have now measured enough FRFs that we can describe the response of the system for all points. Of course, I'm assuming that the reference location at point 9 in the z direction is not the node of a mode!

So what should I do?

So while it appeared on the surface that both tests were the same, there actually is a difference!!! So how could I change these test setups so that the same data is measured. Well, there are two ways. First, Setup #1 could be changed as follows. Instead of using a tri-axial accelerometer, we could use a single uniaxial accelerometer to acquire data at 9z, for instance. But the difference would be that the impact excitation needs to be applied in the x, y and z directions. Then the data collected would be a row of the FRF matrix with 9z as the reference. This is exactly the same data as collected in Setup#2 provided that reciprocity holds true.

The other way to make sure that the same data is collected is as follows. In Setup #1, the impact hammer needs to be used to excite the x, y, and z direction. So the roving hammer needs to impact in all three directions. In Setup #2, the stationary impact at point 9 would need to be used to excite the structure in all three directions. Both tests would then produce 3 complete rows or columns of the FRF matrix.

Now you still may be a little confused by this. I know it's not easy to comprehend the first time you hear it. The best way to convince yourself is to write out all the FRF measurements that you intend to collect to assure that at least one complete row or one complete column of the FRF matrix is acquired.

I hope this simple explanation helps to clear up your question. You need to carefully think about the measurements you are going to make. Remember what I always say: "Thinking is *not* optional!" If you have any more questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Should I always use a hard tip for impact testing . . .
so the input spectrum is flat over all frequencies?
Well . . . too hard a tip may cause problems.

For some reason, everyone thinks that the input spectrum should be flat over the whole frequency range of interest when performing an impact test. But what do we mean by "flat" anyway. Well, it would be better to say that the input spectrum should be "reasonably flat" over all frequencies with "no significant drop-outs or zeros" in the frequency spectrum. So what does that mean.

Basically, we want the input spectrum to have sufficient, fairly even excitation over the frequency range of concern. If the input spectrum were to completely drop off to zero, then the structure would not be excited at that frequency which is not desirable. I use the words "reasonably flat" to allow for some engineering judgment as to what is acceptable.

Of course, many times people don't like engineers to use judgment, so they identify specific criteria or limits to force the situation to be controlled. At times, specifications have been written with specific criteria such as "the input spectrum should roll-off no more than 3 dB over the FFT analysis frequency range". This is a very specific requirement which does not allow the engineer to think. It just forces him to follow a rule without thinking. A criteria like this one may force a poor measurement to be made. But if we don't have to think (or are not allowed to think) then inappropriate measurements could be acquired.

Now you asked about using a hard tip for all your impact tests. I'll answer that in a minute but first let's discuss some basics about the selection of hammer tips for an impact test. First of all, let's remember that the input force spectrum exerted on the structure is a combination of the stiffness of the hammer/tip as well as the stiffness of the structure. Basically the input power spectrum is controlled by the length of time of the impact pulse. A long pulse in the time domain, results in a short or narrow frequency spectrum. A short pulse in the time domain, results in a wide frequency spectrum.

Let's look at some cases and see what this means from a measurement standpoint. (In all the following figures, black is the FRF, blue is the input spectrum and red is the coherence).

Now let's use a very soft tip to excite a structure over an 800 Hz frequency range. As shown in Figure 1, we see that the input power spectrum (blue) has some significant roll-off of the spectrum past 400 Hz. We also notice that the coherence (red) starts to drop off significantly after 400 Hz and the FRF (black) does not look particularly good past 400 Hz. The problem here is that there is not enough excitation at higher frequencies to cause the structure to respond. If there is not much input, then there is not much output. Then none of the measured output is due to the measured input and the FRF as well as the coherence are not acceptable.

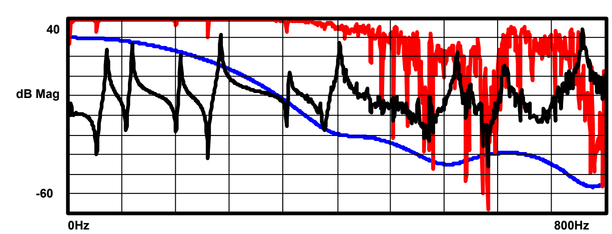

Figure 1 - Very Soft Tip

Now let's use a very hard tip to excite a structure over a 200 Hz frequency range. As shown in Figure 2, we see that the input power spectrum (blue) is extremely flat over all frequencies of interest. We also notice that the coherence (red) is not particularly good for this measurement. The problem here is that there is too much excitation at higher frequencies causing all the modes of the structure to respond. (We'll discuss this further in a moment.)

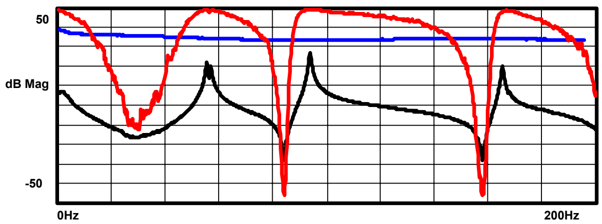

Figure 2 - Very Hard Tip

Now let's use a medium hardness tip to excite a structure over an 200 Hz frequency range such that the input force spectrum does not drop off significantly by the end of the frequency range of interest. As shown in Figure 3, we see that the input power spectrum (blue) rolls off by 10 to 20 dB by 200 Hz. We also notice that the coherence (red) looks especially good at all frequencies over the 200 Hz band with the exception of anti-resonances. The drop off of the coherence is fully acceptable at these frequencies since the structure is non-resonant (anti-resonant) at these frequencies. This means that there is no response to measure so the coherence is expected to drop here. This is a good measurement.

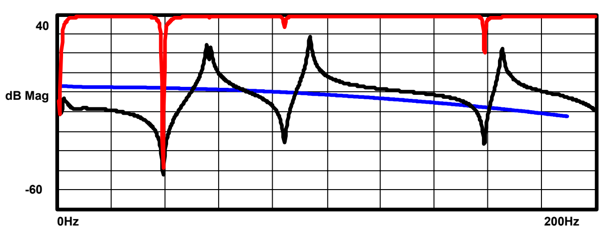

Figure 3 - The Right Tip

Notice that the input spectrum is not perfectly flat as you suggested it should be. In fact, when the input is almost perfectly flat as shown in Figure 2, the measurement is not as good. Let's explain why this happens. Consider the measurement shown in Figure 4. This measurement was taken over a 400 Hz bandwidth. The hammer tip used had approximately 20 dB rolloff over the 400 Hz band which is probably acceptable for this measurement.

Now let's say that I wanted to only measure to 128 Hz and that I wanted to impose a restriction that the input spectrum could not roll off more than 3 dB. Well look at Figure 4 with the 128 Hz bandwidth specified. The input force spectrum rolls off approximately 2 to 3 dB over this 128 frequency band. So the measurement should be acceptable. But what you have to realize is that while the analysis frequency band is only 128 Hz, the response of the structure is based on the energy imparted to the structure. So the structure responds well past 128 Hz because the input force excites all of those modes - *even though I might not be interested in those frequencies.*

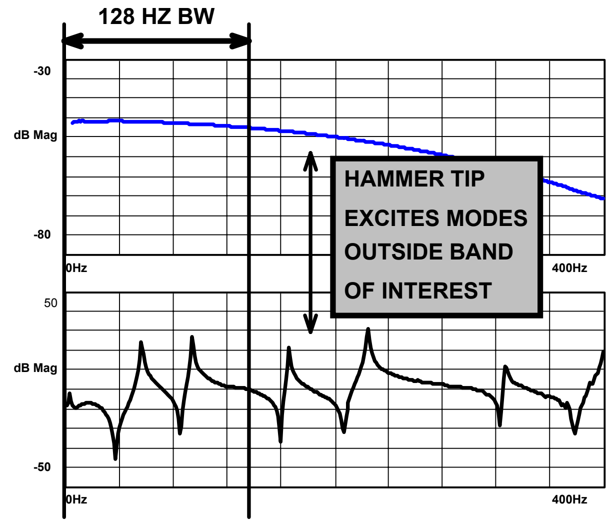

Figure 4 - Exciting Modes Outside the Band of Interest

The accelerometer, mounted on the structure, measures all that response and outputs a voltage which is input to the analyzer. Just doing a quick eyeball of the total area under the curve of the FRF, it appears that only one-third of the energy is associated with the bandwidth of interest. The rest of the energy is associated with something that I'm not interested in measuring. But the accelerometer senses that energy! The ADC on your analyzer may need to be setup such that an overload does not occur due to the total response of the structure.

If the signal is not analog filtered before it reaches the analyzer, then the ADC may need to set excessively high to avoid a potential overload. *Remember, most of the energy of the signal is probably outside the 128 Hz bandwidth of interest!!!* This results in a quantization problem in the ADC. This can easily be corrected through the use of an impact tip that does not needlessly excite modes outside the bandwidth of interest.

So now you can see why I don't like to use a hard tip all the time for impact testing. Sure it gives a good flat input force spectrum. The problem is that it excites more modes than desired and may cause a poor measurement. Think about it and if you have any more questions about modal analysis, just ask me.

Copyright 1998
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Which shaker excitation is best? Is there any difference?
Well ... that's a good question
Let's talk about the different techniques.

Let's discuss the most commonly used excitation techniques for modal analysis today. These are random, burst random, sine chirp and digital stepped sine. But before we discuss the excitation techniques themselves, there are a few basics that we need to discuss first. Let's try to categorize the different techniques and explain when to use which technique. First of all, let's break up the excitations into deterministic and non-deterministic (or random) excitations.

Deterministic signals are those that can be described at any point in time by a mathematical function - they can be *determined*. Typical signals of this type are sinusoidal in nature, such as sine chirp and digital stepped sine. Random signals, on the other hand, can not be described by a mathematical function but are rather described by their statistical characteristics. Typical signals of this type are random and burst random.

In general, we use deterministic signals on linear systems. We also use deterministic signals to determine if a system is linear by performing a linearity check. We use random signals to average slight nonlinearities in a system due to things such as rattles. If we have a structure that has gross nonlinearities, then we need to stop and think just how useful the results of a *linear modal analysis* will be. But understanding the difference between these two categories helps in deciding which technique will provide the best measurement. Depending on the system being tested, you may want to *document the linearity* of the system under test, or you may want to *linearize* any slight nonlinearities that exist..

Now first, let's consider a random excitation. Random is used quite widely for general vibration testing today. But it is not considered one of the best techniques for acquiring FRF measurements for modal testing (although it is still often used). The random nature of the signal excites the structure with varying amplitude and phase as averages are collected. This tends to average any slight nonlinearities that may exist in the structure. While this is a benefit, the signal nevers satisfies the periodicity requirement of the FFT process. Therefore, leakage is a tremendous problem. Even with a Hanning window applied, the resulting FRFs will always suffer from leakage; the peak amplitude will be affected and there will be an appearance of more damping in the structure due to the leakage and windowing effects. A typical measurement sequence is shown in Figure 1. The resulting FRF and COH are shown in Figure 2. Notice how the coherence dips at the resonances of the system; this is a characteristic of random excitation.

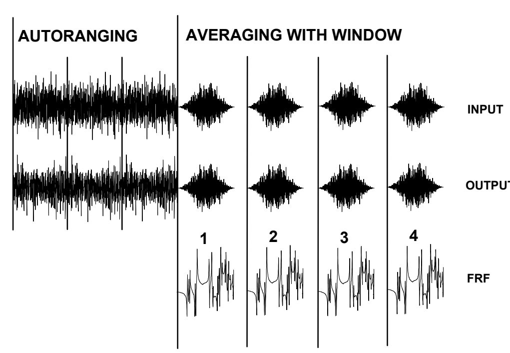

Figure 1 - Typical Random Measurement Sequence

Now, let's consider a burst random excitation. The only difference is that the random signal is only used during a portion of the data capture. If a pretrigger delay is also used, then the signal is totally observed within one sample interval.

Copyright 1998
Pete Avitabile

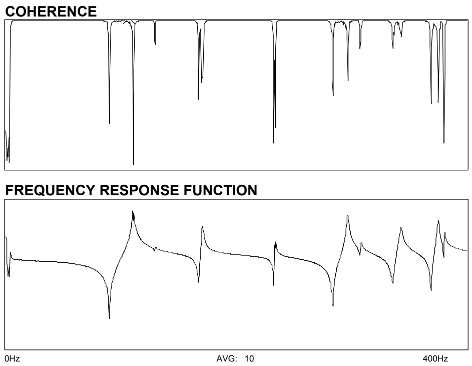

Figure 2 - Random Excitation w/Hanning

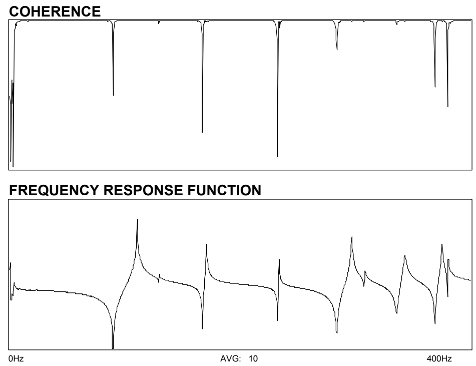

Figure 4 - Burst Random Excitation

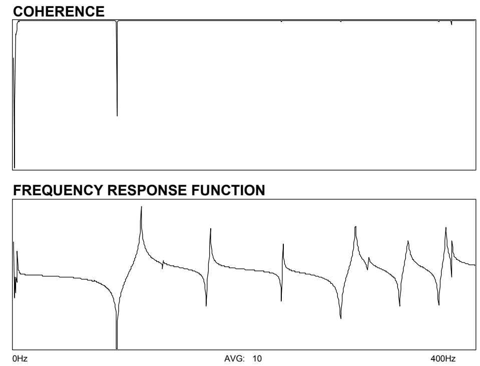

Figure 5 - Sine Chirp Excitation

Therefore, the signal satisfies the periodicity requirement of the FFT process. This means that no leakage will occur and no window is needed. Of course, both the input and responsesignals need to satisfy this requirement. This is easily done for most structures. This signal is well suited for averaging out slight nonlinearities that may be found in the measurement. A typical time measurement is shown in Figure 3. Notice that the excitation is terminated such that the response signal also decays to zero within the sample interval. The resulting FRF and COH are shown in Figure 4. Notice the improvement in the measurement and coherence when compared to Figure 2. The peaks are much sharper and better defined; the coherence is especially good at the resonances.

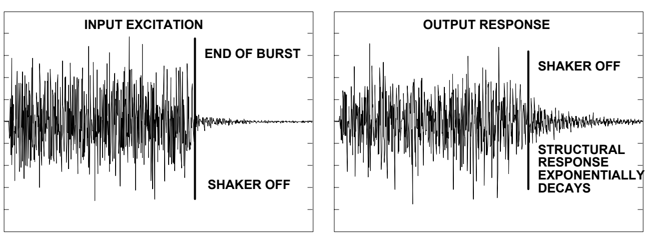

Figure 3 - Typical Random Measurement Sequence

Now, the sine chirp is a fast sweep from low to high frequency within one sample interval of the analyzer. The signal repeats and therefore satisfies the periodicity requirement of the FFT process. This means that no leakage will occur and no window is needed. Of course, the signal must be played continuously so that the structure steady state response is achieved. The resulting FRF and COH are shown in Figure 5. The measurement is very similar to the results from the burst random test. By changing the input force level applied to the system, linearity checks can be easily made using this excitation technique.

Finally, the digital stepped sine technique requires that a single frequency, coincident with an analyzer spectral line, is used to excite the system. Since the signal is guaranteed to be periodic with regards to the FFT process, no leakage occurs and no windows are necessary. Since it is not broadband in nature, this technique is the slowest of all techniques because each spectral line is evaluated individually. However, it is excellent for documenting nonlinearities and is likely to produce the best measurement of all the excitation techniques above.

When comparing the techniques, the burst random and sine chirp will produce similar results if the system is linear. In general, the random measurement will always suffer from leakage and the quality of the measurement will suffer when using this technique. To illustrate the degradation of the measurement when using random excitation, Figure 6 compares random and burst random with an expanded look around the first resonant peak of the system. The random signal contains alot of variance and the peak is distorted at resonance (where the coherence is known to dip). In fact, there almost appears to be two modes at that frequency; this is due to the distortion of leakage. The burst random measurement is clean and sharp. Clearly, the burst random measurement is the better of the two measurements.

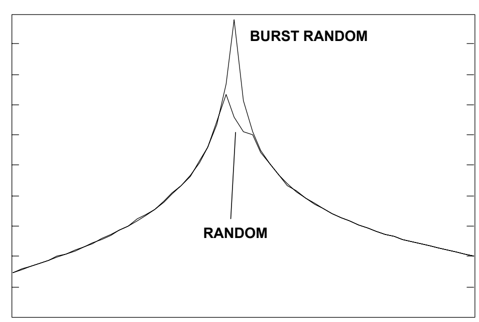

Figure 6 - FRF for Burst Random and Random

We could spend alot more time discussing all the details of each of the techniques (as well as others not mentioned) but there isn't enough time right now to cover everything. Maybe another time we can discuss each of the techniques in more detail. But this quick overview should give you what you need to know. If you have any more questions on modal analysis, just ask me.

Copyright 1998
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

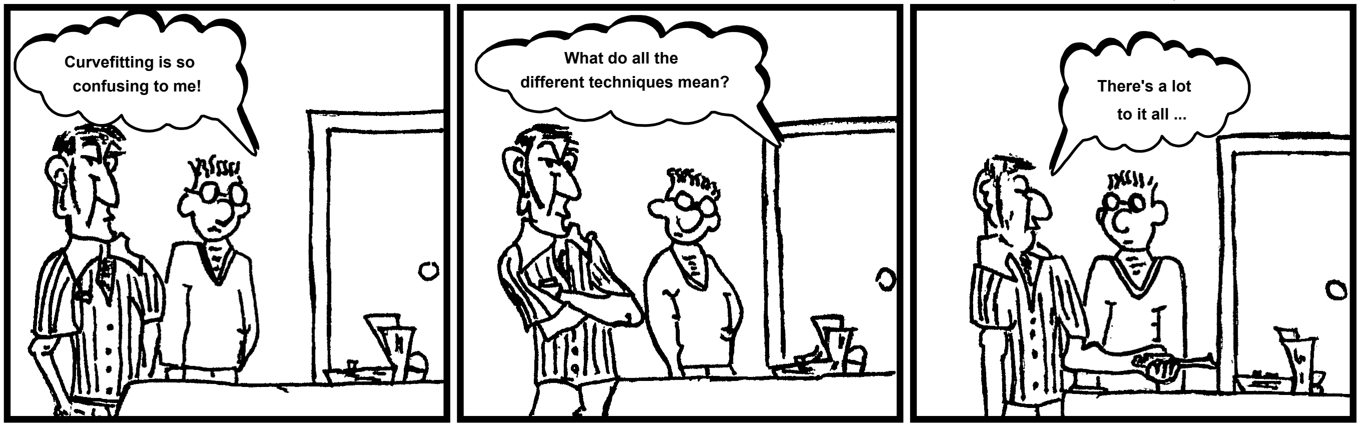

Illustration by Mike Avitabile

Curvefitting is so confusing to me!
What do all the different techniques mean?
There's a lot to it all ...

Curvefitting is probably the most difficult part of the whole experimental modal analysis process for most people. Actually, its better to refer to it as *modal parameter estimation*. But that's a mouthful - so we usually just call it *curvefitting*. But we are actually trying to extract modal parameters (frequency, damping and mode shapes) from measured data. Let's discuss a few general items first.

Basically, we need to describe the system in terms of it's modes of vibration. For example, the three mode system shown in Fig 1 can be described by the following frequency domain representation of the system as

$$h_{ij}(j\omega) = \sum_{k=1}^{3} \frac{a_{ijk}}{(j\omega - p_k)} + \frac{a_{ijk}^*}{(j\omega - p_k^*)})$$

or broken down into the contribution of each of the modes as

$$h_{ij}(j\omega) = \frac{a_{ij1}}{(j\omega - p_1)} + \frac{a_{ij1}^*}{(j\omega - p_1^*)} + \frac{a_{ij2}}{(j\omega - p_2)} + \frac{a_{ij2}^*}{(j\omega - p_2^*)} + \frac{a_{ij3}}{(j\omega - p_3)} + \frac{a_{ij3}^*}{(j\omega - p_3^*)}$$

Now as you start to look at this measurement, some quick thoughts come to mind. How many data points should I use? What should the order of the model be? Are there any effects from modes outside the band of the curvefitter? Does the same technique need to be applied to all the modes? When do I use a SDOF vs. a MDOF technique? Should I use a time or frequency domain curvefitter? (And the most important thing that should come to mind is - Oh how I wished I listened in modal class that day instead of going out to party!)

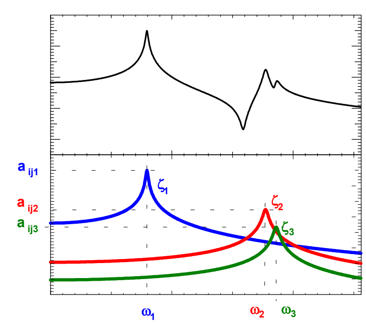

Figure 1 - Three mode system

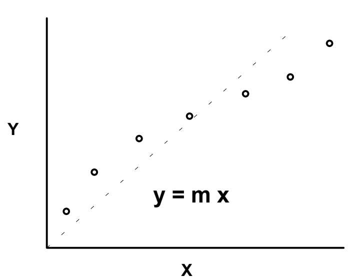

Force data to pass thru zero
Figure 2a

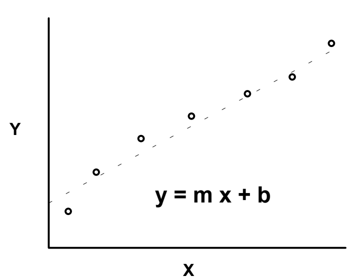

Allow for compensation
Figure 2b

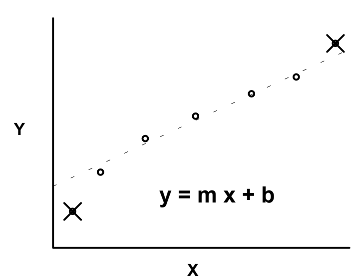

Use only part of the data
Figure 3b

Copyright 1999
Pete Avitabile

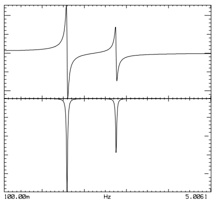

Figure 4a

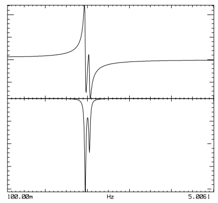

Figure 4b

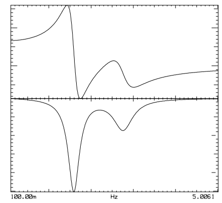

Figure 4c

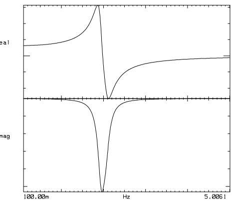

Figure 4d

First let's consider the some simple force gage calibration data in Fig 2. Now if the force gage should read zero at zero load, then Fig 2a represents the best straight line fit of the data - but that fit doesn't look very good. But what if the force gage had a preload. Then it may be necessary to allow for some compensation as shown in Fig 2b. And what if some of the measured data was outside the useful range of the force transducer. Possibly only a portion of the data should be used as shown in Fig 2c. And who said that the force gage was linear with a first order approximation of y=mx+b? I could possibly envision a cubic function that would better describe the measured data. For some reason everyone understands this force gage example but have a hard time realizing that my measured FRF has the same characteristics. Basically the analyst must decide on the order of the model, the amount of data to use and the need for residual compensation as shown in Fig 3. The basic equation to address this measurement is

$$\left[H(s)\right] = \text{lower residuals} + \sum_{k=i}^{j} \frac{\left[A_k\right]}{\left(s - s_k\right)} + \frac{\left[A_k^*\right]}{\left(s - s_k^*\right)} \quad + \text{upper residuals}$$

Basically, I select a band of modes to fit, specify the order of the model and decide on inclusion of residual terms.

Now I need to know when to use a SDOF or MDOF technique. What I need to know is how much modal overlap exists from one mode to the next. Fig 4 shows a variety of different situations for a two DOF system. Fig 2a shows modes that are well separate with very light damping. These types of modes can be approximated with a SDOF fit. Fig 2b shows modes that are closely spaced with very light damping.

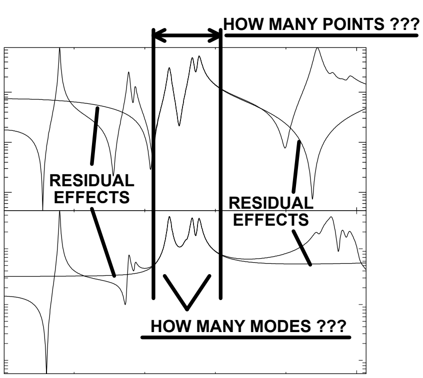

Figure 3 - Schematic of Analyst Curvefitting Decisions

There is some overlap from one mode to the next which may not be correctly compensated with a SDOF fit. It is likely that a MDOF fit may need to be employed for these two modes. Fig 4c shows well separated modes but damping causes some overlap which may also require a MDOF fit. But for both of these last two cases, you may try a SDOF fit for comparison with the MDOF fit. Fig 4d shows modes that are closely spaced with heavy damping. A MDOF fit would be needed for this case.

The last thing to consider is whether to use a time or frequency domain technique. The mathematical relationship is basically the same - it just looks different. Many times we write a relationship in a given form because there is some mathematical *gimmick* that makes the equation easier to solve or more efficient from a computational standpoint. But, in essence, both domains are equivalent. However, many times we tend to use the time domain techniques for lightly damped systems and the frequency domain techniques for heavily damped systems.

If I now look at Fig 5, what would I think would be appropriate for estimating parameters for this measurement. Well it is probably allright to use a SDOF for that first peak. But modes 2 and 3 are too closely spaced to use a SDOF, so most likely a MDOF technique would be used for these modes. Another thing to realize is that the cursors don't need to overlap or cover the whole frequency band. Remember that we are trying to extract parameters that identify the frequency, damping and residue for the system for each of the modes.

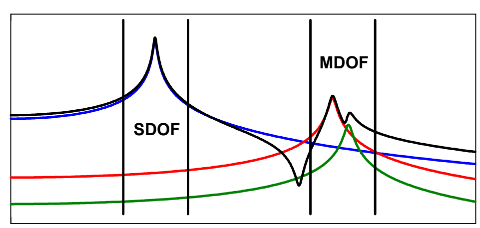

Figure 5 - Possible Curvefit Bands

We really need to spend a lot more time discussing all the details of each of the techniques but there isn't enough time right now to cover everything. But this quick overview should give you an idea of some of the concepts involved. Think about what we have discussed and maybe another time we can discuss each of the techniques in more detail. If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

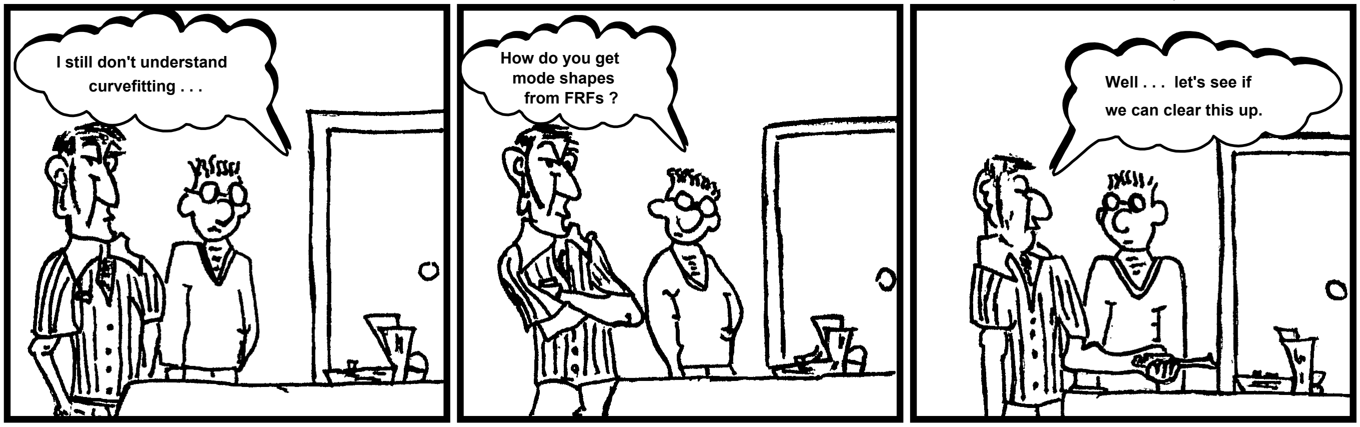

Illustration by Mike Avitabile

I still don't understand curvefitting ...
How do you get mode shapes from FRFs?
Well . . . let's see if we can clear this up.

Modal parameter estimation (commonly referred to as curvefitting) is probably, by far, the hardest part of experimental modal analysis for most people to understand. I know I can write out all the equations to explain this. But I will probably bore you to death. Not only do I have to write out all the equations relating to the modal parameter estimation process, I also have to show the equations relating the residue to the mode shape. And, of course, the concept of a residue is another abstract concept. (Oh, how I wished they had called it a mode shape rather than a residue since this only confuses everyone.)

Last time (Feb 1999), we talked about the curvefitting model and the basic equation we use for estimating parameters, of which one form is

$$\left[H(s)\right] = \text{lower residuals} + \sum_{k=i}^{j} \frac{\left[A_k\right]}{\left(s-s_k\right)} + \frac{\left[A_k^*\right]}{\left(s-s_k^*\right)} \quad + \text{upper residuals}$$

Now those terms in the matrix, [A], are the residues which are obtained from the curvefitting process; we also get the poles, or frequency and damping, from the denominator of the equation. Now these residues can be shown to be related to the mode shapes. Without going through all the steps, the resulting relationship is shown below (with some terms expanded)

$$\left[A(s)\right]_k = q_k \left\{u_k\right\} \left\{u_k\right\}^T$$

$$\begin{bmatrix} a_{11k} & a_{12k} & a_{13k} & \cdots \\ a_{21k} & a_{22k} & a_{23k} & \cdots \\ a_{31k} & a_{32k} & a_{33k} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} = q_k \begin{bmatrix} u_{1k}u_{1k} & u_{1k}u_{2k} & u_{1k}u_{3k} & \cdots \\ u_{2k}u_{1k} & u_{2k}u_{2k} & u_{2k}u_{3k} & \cdots \\ u_{3k}u_{1k} & u_{3k}u_{2k} & u_{3k}u_{3k} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

And if we were to look at each of the columns we would see the mode shape is contained in the column with some scalar multipliers; we would also see that due to reciprocity, the rows also contain the mode shapes. If we were to look at one column, such as the first column, then we would see

$$\begin{Bmatrix} a_{11k} \\ a_{21k} \\ a_{31k} \\ \vdots \end{Bmatrix} = q_k u_{1k} \begin{Bmatrix} u_{1k} \\ u_{2k} \\ u_{3k} \\ \vdots \end{Bmatrix}$$

The residues are, therefore, nothing more than the mode shape multiplied by a scalar which is the value of the mode shape at the reference location, u, and the scaling constant, q. (The q scale constant allows for mode shapes to be represented with different scale constants (unit modal mass, unit length, etc.)

Great, so here are some equations that you may or may not fully understand or appreciate. Maybe a better way to explain the concept is through some simple pictures. Let's go back to that simple plate that we discussed some time ago (Feb 1998) and explain very simply how we can get mode shapes from measurements (then maybe you'll appreciate what the math is doing for us).

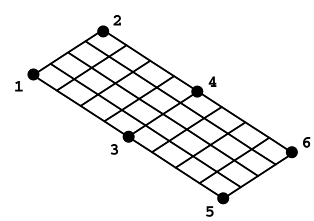

Copyright 1999
Pete Avitabile

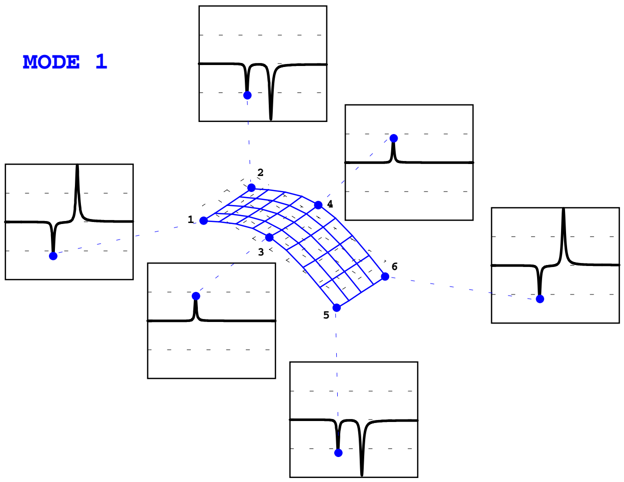

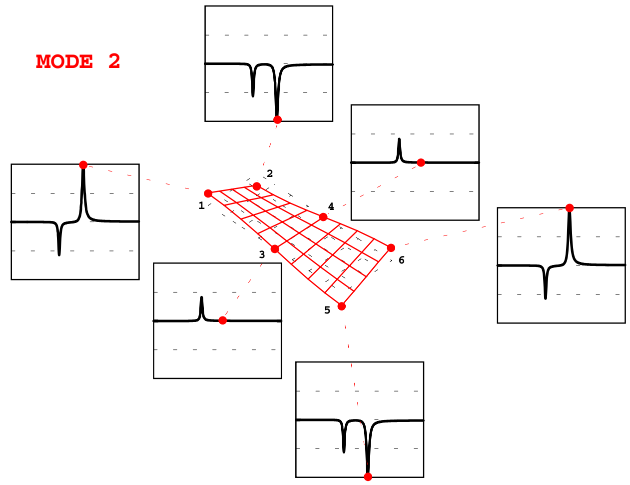

Now let's take some measurements on the plate so that we get a total of 6 FRFs - at the 4 corners and at the 2 mid-points. We want to be able to determine what the first two mode shapes look like from these measurements. Now we could look at the log magnitude of the FRFs but this is not very useful since all the peaks would be positive in this plot.

A more informative plot is the imaginary part of the FRF. This shows both amplitude and, most importantly, the direction of the response. Without getting into all of the technical math, we know that the peak amplitude of the imaginary part of the FRF is directly related to the residue (and the residue is related to the mode shape). This approximate equation is shown below

$$h(j\omega)\,\Big|_{\omega\to\omega_n} = \frac{a_1}{(j\omega_n + \sigma - j\omega_d\,)} + \frac{a_1^*}{(j\omega_n + \sigma + j\omega_d\,)}$$

$$a_1 = \sigma\, h(j\omega)\,\Big|_{\omega\to\omega_n}$$

This very simplistic approach to determining mode shapes is commonly referred to as peak picking since we are picking the peak of the FRF. Now let's look at some of the peaks for each of the measurements at each of the points. (In all of the plots shown, the amplitude of the scale ranges from minus one to plus one and the dashed line is one-half. In addition, the frequency axis has been removed.) Now let's just concentrate on mode 1 first and then go on to mode 2.

Look at the FRF for mode 1 for point 1. Notice that this amplitude is 0.5 and it is negative. If we look at point 2, then we see that the amplitude is also 0.5 and it is also negative. This means that point 1 and 2 are moving with the same amplitude and in the same direction for mode 1. If we look at point 5 and 6, we see the same thing as point 1 and 2. So we can see that points 1, 2, 5, and 6 are all moving with the same amplitude in the same direction.

A important point to make here is that if I only measure these four points, then it would appear to me that the mode shape of the plate would be a rigid body mode (all four points moving together with equal amplitude). This is a common problem encountered when too few points are used to describe the mode shape of a system.

Now look at point 3 for mode 1. Notice that since amplitude is 0.5 but that it is positive. Then same can be said for point 4. So we see that point 3 and 4 have the same amplitude and move in the same direction together. But we also notice that points 3 and 4 are moving in the opposite direction from the rest of the points. Now, while we haven't measured more than 6 points, we start to see that the plate is deflecting into a pattern that is plate bending in characteristic. If we measured more points, then we would see a much better defined mode shape.

Now if we look at mode 2 we can step through and look at all the points and what we will see is that point 1 and 2 have the same amplitude but now they are moving in opposite directions. The same is true for points 5 and 6. But we notice that points 1 and 5 are also moving in opposite directions; the same is true for points 2 and 6. So we see that there is some type of twisting or torsional type deformation pattern for mode 2. If we look at points 3 and 4, we notice that these points have zero value. This is because points 3 and 4 are node points for the torsional mode of the plate. Again, adding more points better defines the shape.

So now we can see that the peaks of the imaginary part of the FRF are directly related to the mode shape of the plate for each of the modal peaks. Without going through all the math, the residues are terms that are extracted from the curvefitting process and these residues are directly related to the modes shapes of the plate. This was shown pictorially to keep things simple.

I hope that this helps to clear up the mystery as to how we get mode shapes from FRFs. Think about it and if you have any more questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

Illustration by Mike Avitabile

What's the difference between operating deflection shapes and mode shapes? sometimes they look the same to me! Well . . . let's describe the differences.

This is a common stumbling point for many people. This is partly due to the words that we use. I would much rather call the data we receive from an operating condition, an *operating deflection pattern,* rather than use the word *shape*. But unfortunately, I can't change the nomenclature at this point.

Let's first recall how a structure responds, in general, due to any excitation

$$h(j\omega) \times f(j\omega) = y(j\omega) \qquad (1)$$

Of course, we realize that the input forcing function is actually applied in the time domain but we represent it in the frequency domain; also the response actually occurs in the time domain but it can also be represented it in the frequency domain.

So for a structure which is exposed to an arbitrary input excitation, the response can be computed using the frequency response function multiplied by the input forcing function. This is very simply shown in the schematic in Figure 1.

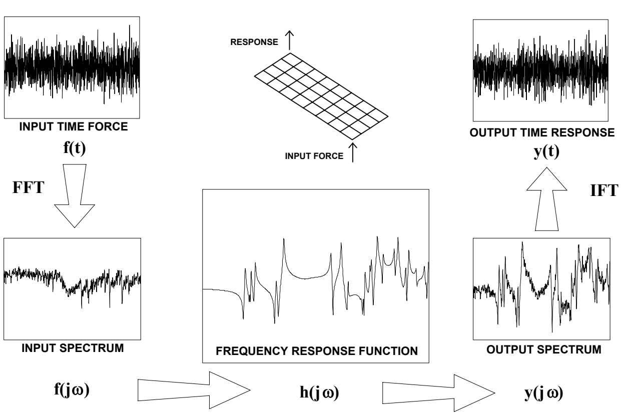

Figure 1 - Schematic Overviewing the Input-Output Structural Response Problem

The excitation shown is a random excitation that excites all frequencies. The most important thing to note is that the frequency response function acts as a filter on the input force which results in some output response. The excitation shown causes all the modes to be activated and therefore, the response is, in general, the linear superposition of all the modes that are activated by the input excitation. Now what would happen if the excitation did not contain all frequencies but rather only excited one particular frequency (which is normally what we are concerned about when evaluating operating conditions).

Let's consider a simple plate that is excited by an input force that is sinusoidal in nature. And let's also assume that the force is applied at one corner of the plate. For the example here, we are only going to consider the response of the plate assuming that there are only 2 modes that are activated by the input excitation. (Of course there are more modes, but let's keep it simple to start.) Now from figure 1 and equation 1 we realize that the key to determining the response is the FRF between the input and output locations. Also, we need to remember that when we collect operating data, we don't measure the input force on the system and we don't measure the system FRF - we only measure the response of the system.

First let's excite the system with a sinusoid that is right at the first natural frequency of the plate structure. The response of the system for one FRF is shown in Figure 2. So even though we excite the system at only one frequency, we know that the FRF is the filter that determines how the structure will respond. We can see that the FRF is made up of a contribution of both mode 1 and mode 2. We can also see that the majority of the response, whether it be in the time or frequency domain, is dominated by mode 1. Now if we were to measure the response only at that one frequency and measure the response at many points on the structure, then the operating deflection pattern would look very much like mode 1 - but there is a small contribution due to mode 2. Remember that with operating data, we never measure the input force or the FRF - we only measure the output response. So that the deformations that are measured are the *actual response* of the structure due to the input excitation - whatever it may be.

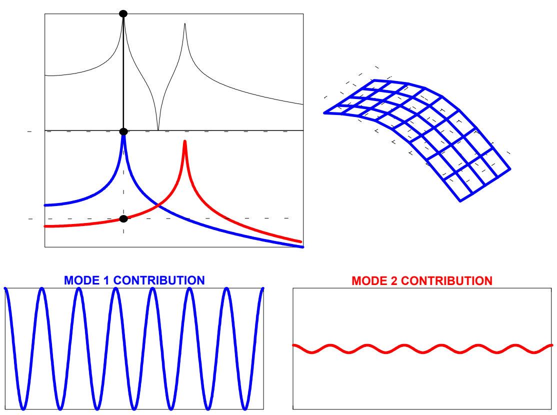

Figure 2 - Excitation Close to Mode 1

When we measure FRFs and estimate modal parameters, we actually determine the contribution to the total FRF soley due to the effects of mode 1 acting alone, as shown in blue, and mode 2 acting alone, as shown in red, and so on for all the other modes of the system. Notice that with operating data, we only look at the response of the structure at one particular frequency - which is the linear combination of all the modes that contribute to the total response of the system. So we can now see that the operating deflection pattern will look very much like the first mode shape if the excitation primarily excites mode one.

Now let's excite the system right at the second natural frequency. Figure 3 shows the same information as just discussed for mode 1. But now we see that we primarily excite the second mode of the system. Again, we must realize that the response looks like mode 2 - but there is a small contribution due to mode 1.

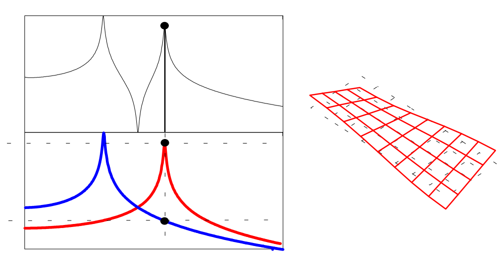

Figure 3 - Excitation Close to Mode 2

But what happens when we excite the system away from a resonant frequency. Let's excite the system at a frequency midway between mode 1 and mode 2. Now here is where we see the real difference between modal data and operating data. Figure 4 shows the deformation shape of the structure. At first glance, it appears that the deformation doesn't look like anything that we recognize. But if we look at the deformation pattern long enough, we can actually see a little bit of first bending and a little bit of first torsion in the deformation. So the operating data is primarily some combination of the first and second mode shapes. (Yes, there will actually be other modes but primarily mode 1 and 2 will be the major participants in the response of the system.)

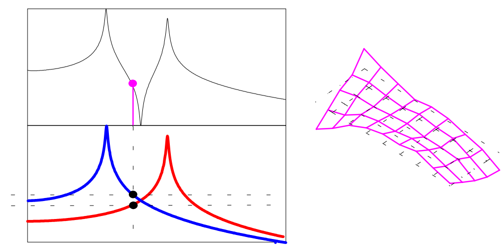

Figure 4 - Excitation Somewhere Between Mode 1 and Mode 2

Now, we have discussed all of this by understanding the FRF contribution on a mode by mode basis. When we actually collect operating data, we don't collect FRFs but rather we collect output spectrums. If we looked at those, it would not have been very clear at to why the operating data looked like mode shapes. Figure 5 shows a resulting output spectrum that would be measured at one location on the plate structure. Now the input applied to the structure is much broader in frequency and many modes are excited. But, by understanding how each of the modes contributes to the operating data, it is much easier to see how the modes all contribute to the total response of the system. So actually, there is a big difference between operating deflections and mode shapes - *we can now see that the modes shapes are summed together in some linear fashion to form the operating deflection patterns*. I hope that this helps to clear up the mystery as to the differences between operating deflection patterns and mode shapes.

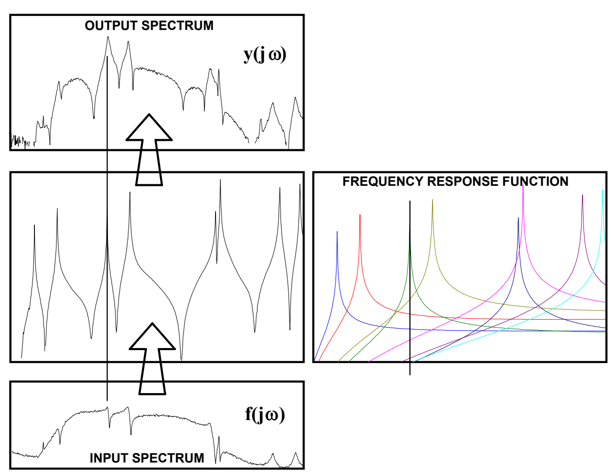

Figure 5 - Broadband Plate Excitation

Think about it and if you have any more questions about modal analysis, just ask me.

Copyright 1999
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Are you sure you can get mode shapes from one row or column of the H matrix?
Sure! Let's walk through an example.

Let's use the beam that we have discussed before as an example. For this beam, we considered three measurement points. There are a total of nine possible input-output FRFs that can be measured. Remember we discussed that these measurements can be obtained from either shaker or impact testing. So that we have some numbers to discuss, the beam mode shape values are shown in Figure 1. (The values will be kept simple for discussion purposes.)

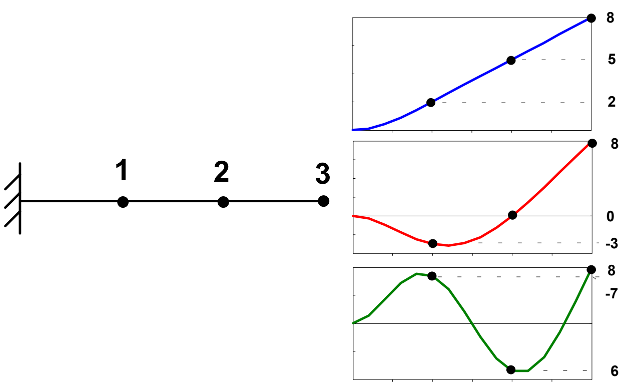

Figure 1

Awhile ago we described how the peak amplitude of the imaginary part of the FRF is directly related to the residue (which is directly related to the mode shape). In fact, we said that the residue was approximated by

$$a_1 = \sigma\, h(j\omega)\, \big|_{\omega \to \omega_n}$$

and that the individual values of the mode shape can be obtained from

$$a_{ijk} = q_k u_{ik} u_{jk}$$

Now let's plot the FRF matrix for this beam with three measurement points. I could show any one of the different parts of the FRF, but it turns out that the imaginary part of the FRF is most informative for this discussion since it shows both magnitude and direction; all the plots have the same -10 to +10 scale. This is shown in Figure 2.

Now let's use the third row of measurements to determine the mode shape for mode 1; this implies that point three is the reference location. Now if I were to pick the peak of the FRF for mode 1, the amplitudes are proportional to the shape of the cantilever beam first mode as seen in Figure 3.

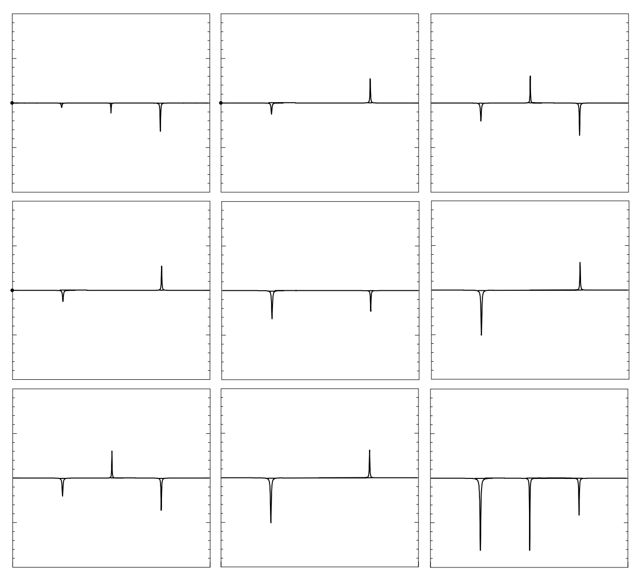

Figure 2

Copyright 1999
Pete Avitabile

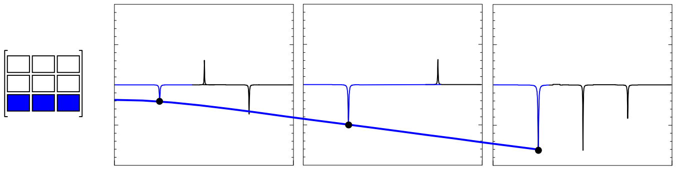

Figure 3

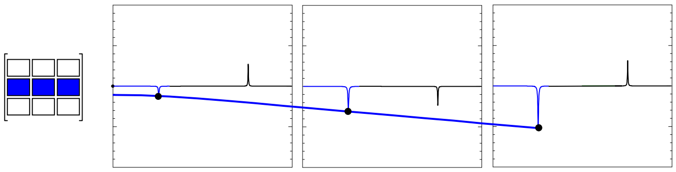

Figure 4

If you look at the values of the amplitude for mode 1 for points 1, 2 and 3, you will see that they are -2, -5, -8, respectively. These values are the values of the mode shape shown in Figure 1. (Notice that I have arbitrarily scaled the values to maintain an easy interpretation of the data. Also notice that the shape could be either plus or minus since the "shape" is the same.)

Now let's use the second row of the FRF matrix. If I pick the peak of the FRF for mode 1, the amplitudes are again proportional to the shape of the cantilever beam first mode as seen in Figure 4.

If you look at the values of the amplitude for mode 1 for points 1, 2 and 3, you will see that they are approximately -1.2, -3.13, -5, respectively. At first glance, these values look different but we can notice that the "ratio" or "shape" is exactly the same as the previous case.

In fact, if I scale the values of the mode shape from the third row by the ratio of the value of the mode shape at reference point 2 (5.0) to the value of the mode shape at reference 3 (8.0), then I will get the mode shape listed above for the 2nd row of the FRF matrix [2 (5/8)=1.2, 5 (5/8)=3.13, 8 (5/8)=5]. This is exactly what I expect to get based on the theory relating mode shapes to residues, so I'm actually not surprised. (We could also look at the first row of the FRF matrix and arrive at the same results.)

So we can see that we can get the mode shape of the beam from any row of the FRF matrix. If we remember that the reciprocity holds true, then we know that the rows and columns contain the same information. So now I can also see that the mode shape in every column of the FRF matrix. So this is why we say that you can use any row or column of the FRF matrix to estimate the mode shape. Of course, I can write out all the equations to show this but the pictorial description is sufficient (and I know how you hate it when I start writing equations!)

Now let's look at mode 2 and use the third row of the FRF matrix. If I pick the peak of the FRF for mode 2, the amplitudes are proportional to the shapes of the cantilever beam second mode as seen in Figure 5.

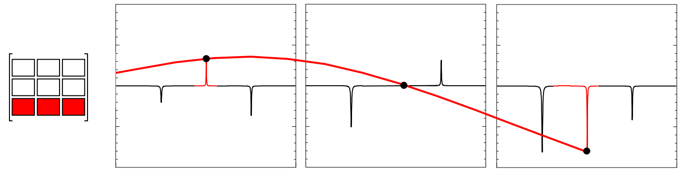

Figure 4

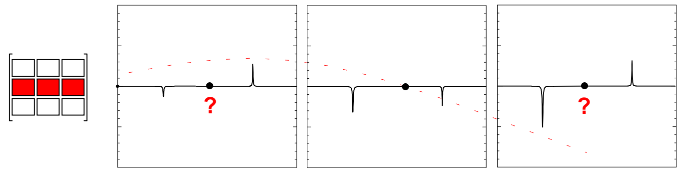

Figure 6

If you look at the values of the amplitude for points 1, 2 and 3, you will see that they are 3, 0, -8, respectively. They are the values of the mode shape shown in Figure 1.

But when I look at the second row of the FRF matrix for mode 2, there is no information pertaining to mode 2. How could this happen? Well, the value of the mode shape for mode 2 at point 2 is zero - its the node of the mode. Anytime we use an input location or response location that is located at a node point (zero shape value) then we will not be able to see the mode from that reference location.

One last picture may help to put it all together for you. Figure 7 shows a waterfall plot of the imaginary part of 15 measurements taken on the beam; the three measurements corresponding to the ones in Figure 2 are shown in color. In this plot, the information pertaining to mode 1 is shown in blue, mode 2 in red and mode 3 in green. We see that the mode shapes can be obtained from the peak of the imaginary part of the FRF. From these plots we can see the first, second and third bending shapes for the cantilever beam.

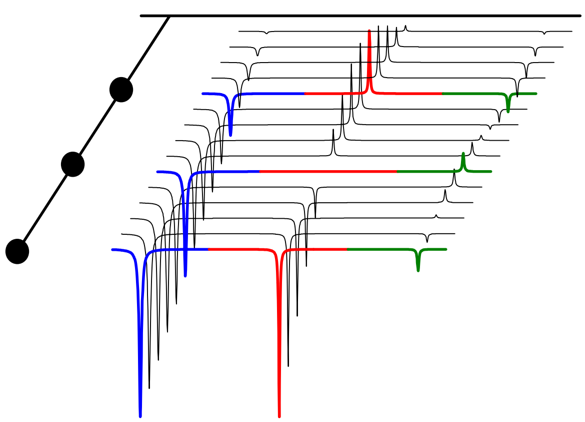

Figure 7

So, in conclusion, we can say that you can use any row or any column of the FRF matrix to estimate any mode of the system, provided that the reference is not located at the node of a mode. I hope this answers your question. If you have any other questions about modal analysis, just ask me.

Copyright 1999
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

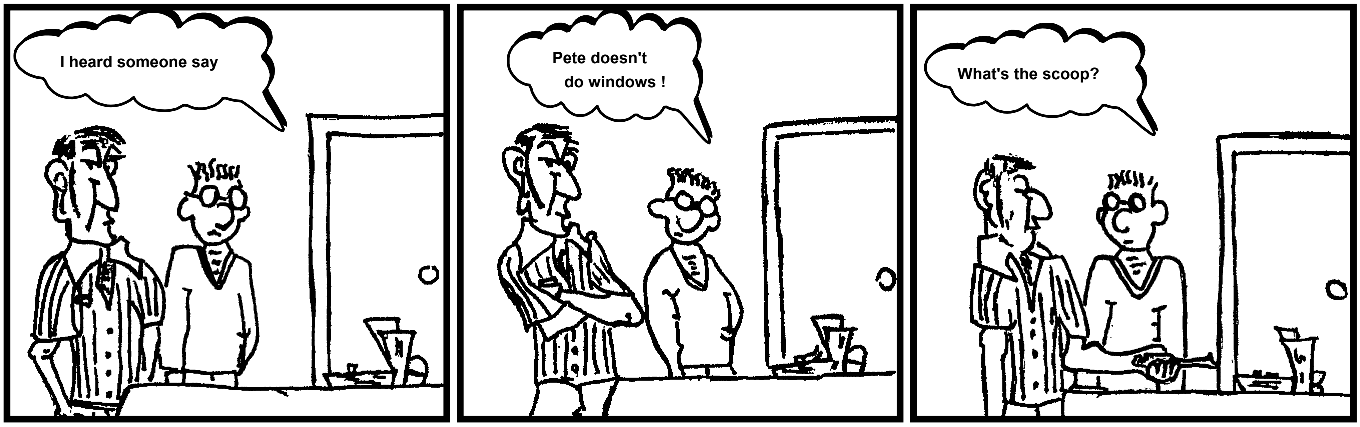

Illustration by Mike Avitabile

I heard someone say Pete doesn't do windows! What's the scoop?

Well ... that's right. But you have to let me qualify that statement. Of course there are many data acquisition situations where it is a necessity to use windows. But almost all of the time when performing a modal test, the input excitation can be selected such that the use of windows can be eliminated. Let's first understand why acquisition of certain types of data can be distorted by the digitization and sampling process, what needs to be done to minimize the distortion, and how to work around the acquisition problem through the selection of specialized test excitation techniques.

First let's remember that the Fourier Transform is defined from $-\infty$ to $+\infty$ but that we only acquire data over a very short time interval. As long as we can reconstruct the data, for all time, from the very small sample we measure, then there is no problem.

Figure 1 shows a simple sine wave, sampled for one time record, with the reconstruction of the time signal from the sample. Figure 1 also shows the FFT of this sampled signal. The time signal is expressed in the frequency domain as one discrete spectral line as expected. This happened because we captured an integer number of cycles of the sine wave in one record or sample of the data - in which case we say that the signal is periodic with respect to the sample interval.

But what if this is not the case. Figure 2 shows this situation. As before, we see the signal, the sample, the reconstructed signal and the FFT of the signal. Notice that the reconstructed signal contains a discontinuity that clearly did not exist in the original signal. The FFT of this signal is far from being a single spectral line as expected. Due to the sampling distortion, the frequency representation is smeared over the whole frequency bandwidth. This very serious error is called leakage and is by far the most serious digital signal processing error that is encountered.
But why does this happen? The original signal was a simple sine wave. How did the frequency representation get so distorted? There's an easy explanation for this. The sampled data does ***not*** contain an integer number of cycles or repetitions of the signal.

Let's stop and recall some simple things we learned about Fourier series. If we start with a simple sine wave, we know that it is a trivial task to describe that signal with a Fourier series. It is basically just one term of the Fourier series which is a sine wave at ω with some amplitude A_0. But do you remember what the series expansion was for a signal such as a rectangular series of pulses? Well, I don't want to expand on all of this right now but I think you would remember that it was a series of sinusoids at different frequencies with different amplitudes. In fact for the rectangular pulse, there were many terms in the series required in order to approximate that signal. That happened because the shape of the discontinuous rectangular pulse doesn't look like a nice smooth sine wave.

Now if I look back at the sampled sine wave in Figure 2, I can now see that by not capturing an integer number of cycles of the signal I have distorted the signal such that it *appears* to have a discontinuous nature at the end of the sample interval. This explains why the FFT is smeared over the frequency bandwidth. Basically, there are many terms needed in order to approximate this *apparently discontinuous* signal.

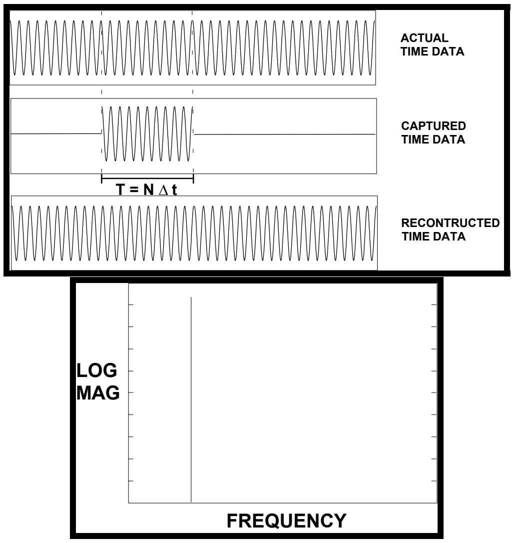

Figure 1

In order to minimize this error (and notice that I said *minimize* and not eliminate), we use weighting functions called *windows*. Basically we apply a weighting function to make the signal *appear* to better satisfy the periodicity requirements of the FFT process. Figure 3 shows a windowed time history.

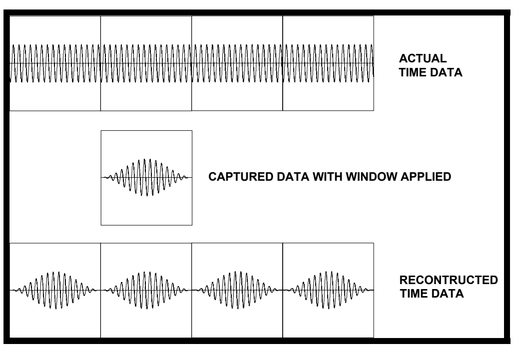

Figure 3

The most common windows for modal testing today are the Rectangular window, the Hanning window, and the Flat Top window for shaker testing and the Force/Exponential window for impact testing. The main thing to understand right from the start is that ***all windows distort data!*** Without going into all the detail, windows always distort the peak amplitude measured and always give the appearance of more damping then what actually exists in the measured FRF - two very important properties that we try to estimate from measured functions. The amplitudes are distorted as much as 36% for the Rectangular window and 16% for the Hanning window. The effects of these windows is best seen in the Frequency domain representation of the weighting function. All windows have a characteristic shape that identifies the amount of amplitude distortion possible, the damping effects introduced and the amount of smearing of information possible.
Figure 4 shows the Rectangular, Hanning and Flat Top windows frequency representation. Sometime soon we will discuss what

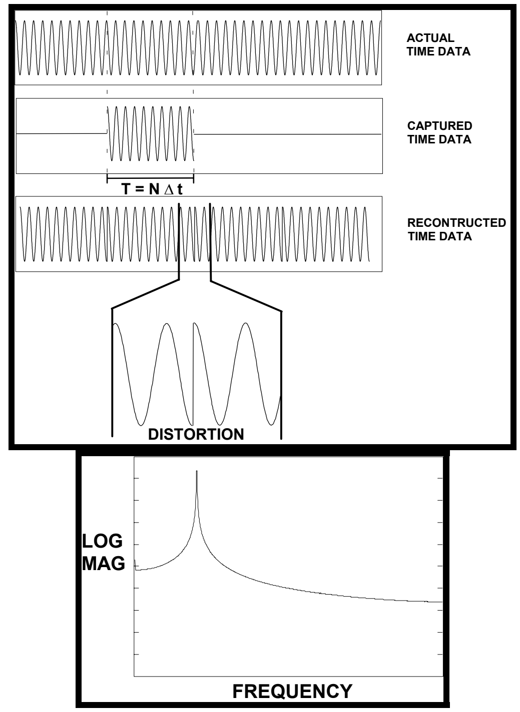

Figure 2

these curve more but for right now, I'm happy if you just understand that the windows, while a necessary evil in some measurement situations, distort data.

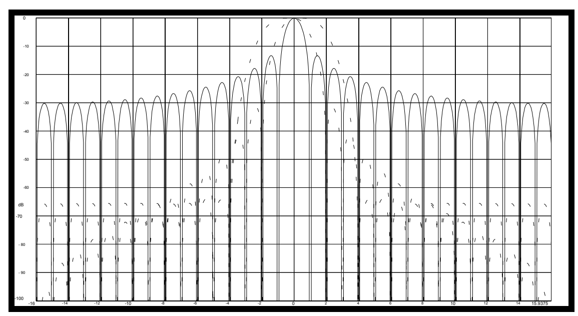
Figure 4

So how do I get around not using windows on measured FRFs for a modal test? Basically, I try to satisfy Fourier's request - "either sample a repetition of the data or completely observe the signal in one sample of data". If you think about it, signals such as pseudo-random, burst random, sine chirp, and digital stepped sine all satisfy this requirement under most conditions and therefore are leakage free and do not require the use of a window. Maybe we can discuss the particulars about each of the windows another time, but this short explanation should suffice for now.

Now I hope you understand why I don't like to use windows and I will avoid the use of windows at all costs - but every once and a while, I have no other choice. (Especially at home, where I can never get out of "doing windows"!) If you have any other questions about modal analysis, just ask me.

Copyright 1999
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

I'm still overwhelmed by all this modal stuff
Laplace, Fourier, FRFs, and all that!
Can you put the big picture together for me?

Sure ... sometimes it helps to stand back and look at everything from a complete picture. I have a figure that I have used for many years now to help people see things more clearly. I call it "The Big Picture". Let's just look at this picture and discuss all the pieces individually.

First let's start with an analytical representation such as the finite element model shown. Basically, we use the FEM to approximate a lumped mass system that is interconnected by springs to represent the physical system. Since the analytical approximation is described in terms of a force balance for each mass that is described in the system, we end up with one equation for each mass (or degree of freedom) used to approximate the system. Since we need many small little finite elements to accurately describe the system, I end up with many equation and unknowns. Right away, it becomes convenient to describe all these equations using matrices. Now once I have assembled all these equations, a mathematical routine called an eigensolution is used to represent the system in simpler terms - the system's frequencies and mode shapes. This is what we do in the finite element process.

Well, without getting into all the details, I can take those same equations and transform them into the Laplace domain. (No - we don't convert to the Laplace domain to make your life miserable - we do it to make some of the equations easier to handle. Please believe me on this one!) Now in the Laplace domain, we have, [B(s)], the system equation and its inverse, [Hs)], the system transfer function. Now we know that this inverse is the adjoint of the system matrix (or the cofactors of the system matrix) divided by the determinant of the system matrix. This inverse is described in all vibrations text books (usually in Appendix A).

So big deal! What's that mean to you! Well, it turns out that the adjoint matrix contains the modal vectors and we call this the Residue Matrix.. The determinant of [B(s)] contains the roots, or poles of the system. Well, this is the same basic information that is obtained from the analytical model. So we could determine the system dynamic characteristics from either the analytical model or from the Laplace domain representation - they both will give the same results.

Now another important relationship is the Frequency Response Function, FRF. This is the system transfer function evaluated along the $j\omega$ axis. The FRF is actually a matrix of terms, $[H(j\omega)]$. Well, since we are dealing with a matrix, it is convenient to identify input-output measurements with a subscript. So a particular output response at point 'i' due to an input force at point 'j' is called $h_{ij}(j\omega)$.

Now remember that the system transfer function has been defined up to this point from mass, damping and stiffness quantities. This function can be computed or *synthesized* for any input-output combination over any frequency band desired. So if we wanted, we could synthesize several FRFs that make up either one full row or one full column of the FRF matrix if needed or desired as shown in the figure.

Now what we need to realize is that those FRFs that were generated (synthesized) contain information relative to the system characteristics. Remember that the FRFs can be generated from residues and poles. And that the residues are directly related to the mode shapes and the poles are the frequency and damping of the system.

Copyright 1999
Pete Avitabile

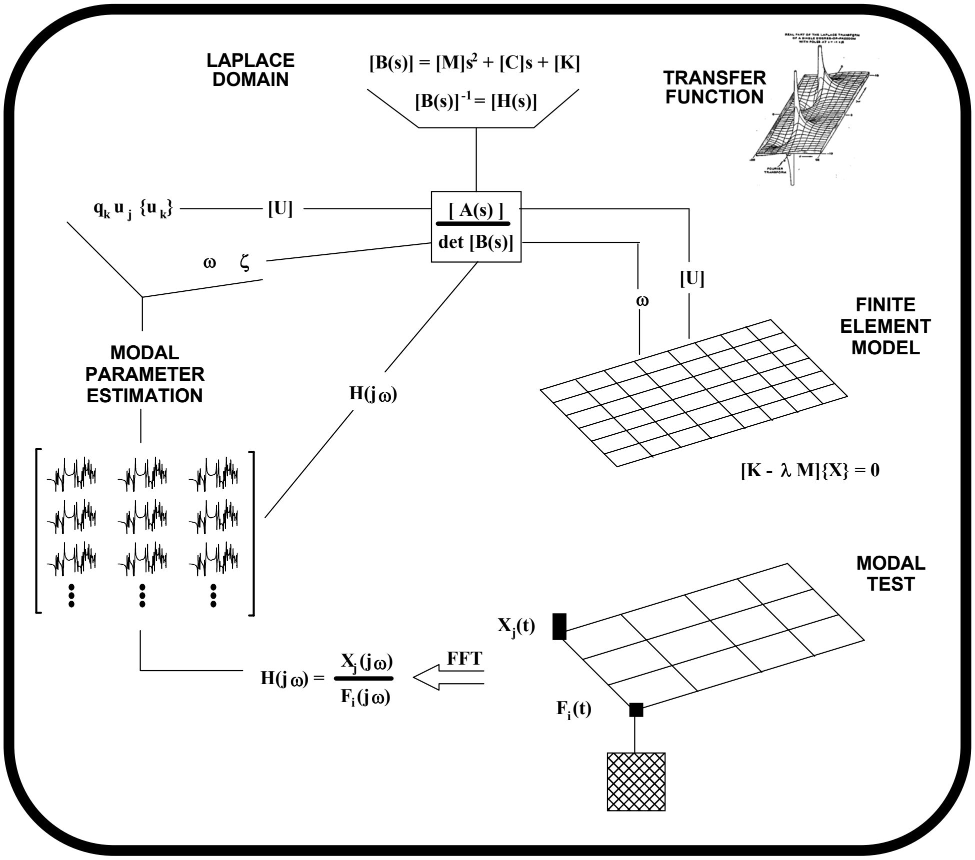

So the parameters that make up the FRFs, are the parameters that we wish to extract from the FRFs. This is what modal parameter estimation is all about. Basically, we use the FRFs in a mathematical algorithm to extract the generic information that makes up the FRFs - the frequency, damping and mode shapes. We often refer to this process as curvefitting. The basic information that is extracted is the mode shapes which are related to information contained in the adjoint matrix or residue matrix and the poles which relate to information in the determinant of the system matrix.

This pretty much summarizes the process - except one important thing needs to be addressed. Up until now we have only discussed using the mass, damping and stiffness approximations to compute system characteristics from the finite element model or from the Laplace domain representation of the system. Both these approaches use approximations of the physical parameters of mass, damping and stiffness to describe the system and so they will both provide the same basic information. If there were some other way to estimate those FRFs without assuming physical properties then I could employ the modal parameter estimation techniques to extract the desired information.

This is where modal testing comes in. Basically, my structure is excited with some measured force. The response of the system due to the applied force is measured along with the force. Now this time data is transformed to the frequency domain using the FFT and basically a ratio of output response to input force is computed to form an approximation of the FRF.

There are many implications of making these measurements which involve digital signal processing concepts which are much too involved to discuss in detail right now (but I think you get the idea where I'm going with all this).

So we could measure one input-output FRF based on this approach. If we used a shaker to excite the structure and move the accelerometer to many points then we could measure a column of the FRF matrix. (If we collected the data using impact techniques then we would measure on row of the FRF matrix). So the big advantage of making measurements is that I measure the response of the system due to the applied force - I don't ever make any assumptions as to the mass, damping and stiffness of the system - and I avoid any erroneous approximations I may make. Of course, I need to make sure that I make very good measurements otherwise I will distort my system characteristics.

So I hope this clears some things up for you. If you have any other questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

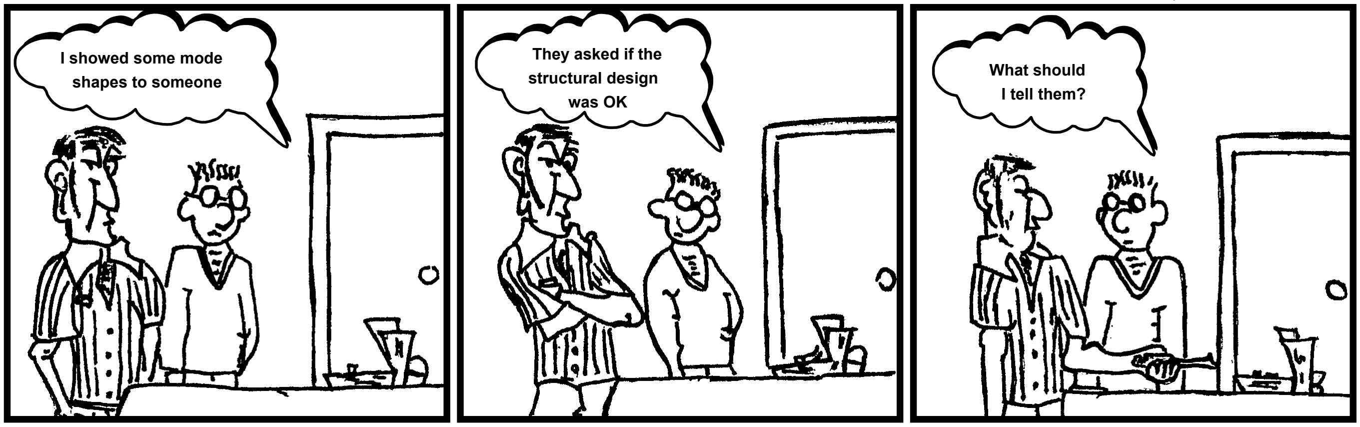

Illustration by Mike Avitabile

I showed some mode shapes to someone
They asked me if the structural design was ok
What should I tell them ?

If I could have a dollar for every time I have heard that question, I'd be rich! The basic answer is you just don't have enough information to answer that question. People who ask that question have no idea what they are asking about. You have to be very diplomatic in telling them that the question is a silly one to ask.

One of the reasons why they are apt to ask the question is because you probably showed them an animation and their impression is that the structure is deforming (since they see the deflections on the computer screen). Of course, you know that this is only a characteristic shape that the structure will undergo when subjected to a force that excites that mode. Sometimes I've been known to say "well let's increase the amplitude of the animation and see if we get the structure to break on the screen". (Of course, this is ridiculous!!! This can't happen.) I use this statement to start to explain what shapes are all about. Animation is only a mechanism to understand how the structure may deform if that mode is excited by the forcing function.

One of the key points here is that we need to know the applied force. For some reason, people forget that we need a force applied to the system to get a response. The physical equation of motion is

$$[M]\{\ddot{x}\} + [C]\{\dot{x}\} + [K]\{x\} = \{F(t)\}$$

and the equivalent modal space representation is

$$\begin{bmatrix} \diagdown & & \\ & \bar{M} & \\ & & \diagdown \end{bmatrix}\{\ddot{p}\} + \begin{bmatrix} \diagdown & & \\ & \bar{C} & \\ & & \diagdown \end{bmatrix}\{\dot{p}\} + \begin{bmatrix} \diagdown & & \\ & \bar{K} & \\ & & \diagdown \end{bmatrix}\{p\} = [U]^T\{F\}$$

Notice that there is a force on the right hand side of this equation. When we solve for the characteristic equation of the system, we assume that there is no force on the right hand side. This is how we obtain the dynamic characteristics of the system. One way to look at it is that the modes of the system are nothing more than a very elaborate set of filters which have the ability to amplify and attenuate an input signal on a frequency basis. If we just look at the filters themselves, can we make any assessment whether the filters are good or bad for a particular application? Of course not! All we can say is that the filters have some characteristics which relate to a center frequency, rolloff and some gain settings as seen in Figure 1.

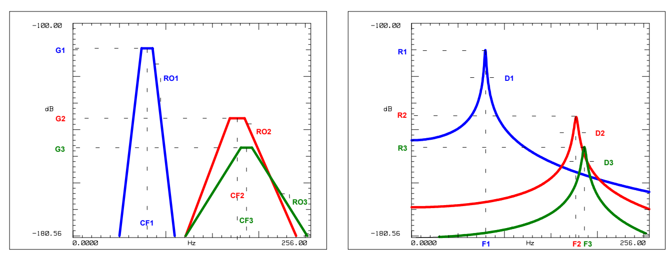

Figure 1

Well ... the dynamic characteristics of a structural system are quite the same. We can identify each mode (each filter) as having a natural frequency (center frequency), damping (rolloff) and residue/mode shape (gain). We need to very clearly understand that the mode shapes are only characteristics and we cannot determine the goodness or badness of a mode unless we know the forcing function - that is, the right hand side of the equation.

As another example, let's say we wanted to determine the stiffness of a cantilever beam. Well, we could go out in the lab and apply a force to the tip of the cantilever beam and measure the resulting displacement. We know that we could determine the stiffness as K = F / X. Now this stiffness is an important

Copyright 2000
Pete Avitabile

parameter or characteristic of the beam. But once I determine the stiffness, do I know if the beam will fail or not? Of course not! I would need to know the actual force that was applied to the beam - wouldn't I? You see, in the test lab we applied an arbitrary force and measured the displacement due to that force in order to determine the character of the beam. Someone needs to identify the actual real world force before I can compute the actual displacement. And then I need to have some specification defined as to how to assess the acceptability of the structure due to the design or real world forces - which brings me to another important point.

One thing that people often forget is that once the mode shapes are obtained and a dynamic design force is specified, the response can be computed, but someone needs to identify a specification defining what is acceptable and unacceptable for the response. This, at times, can be one of the most frustrating parts of the structural dynamic response modeling process. The responses can be computed but no one has defined what the level of acceptance is. Many times this very important detail is overlooked in the process of extracting pretty animated mode shapes. Then everyone asks... how much deflection is acceptable, how long will the component life be, does it "feel" good, is the response too noisy, etc.

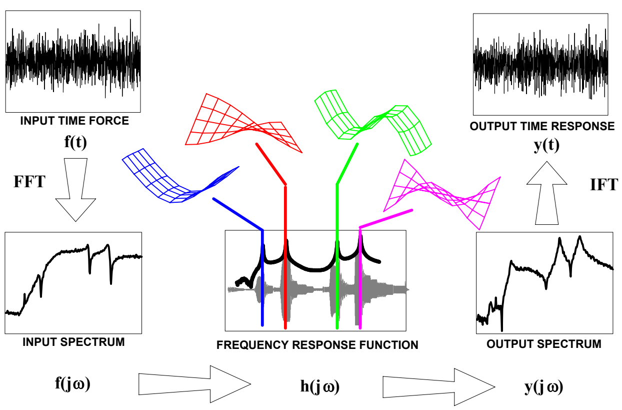

Figure 2

So now that we have discussed a few of these things, let's go back to our plate example that we have discussed before concerning different aspects of modal analysis. Figure 2 shows a schematic of a typical forced vibration problem. There is some force which is applied in the time domain. Well, this time signal is very confusing so it helps to identify some important characteristics of this force if it is transformed to the frequency domain using the FFT process. Now I know that this force is multiplied times the frequency response function in order to get the output of the system. That output could then be transformed back to the time domain if desired. Well, the important point to make here is that the FRF is multiplied by the force spectrum.

That means that the input force spectrum is amplified and attenuated by this multiplication. The FRF controls how this force is amplified and attenuated on a frequency basis. In the figure above, the FRF appears to have contribution for all four modes shown. That assumes that the applied force and response location exists at a point where there is participation of each of the four modes of the system.

But what if the force was applied at a location of a node of a mode. Let's say that the force was applied along the symmetry line along the length of the plate. Then, the applied force would not excite any of the torsional modes from that location; then we say that those modes don't participate in the response of the plate due to that force. The same is true for the response location. So we can see that both the input and output locations will have an effect on the response of the system. (In fact, the mode shape amplitudes have a strong influence on how much a particular mode contributes to the overall response.)

While we could say that certain modes may not participate in the response of the system, that does not imply that those modes don't exist - they just are not needed to compute the response of the system. But the modes still exist - they define the dynamic characteristics of the system. Depending on the location of the applied force and the point where response needs to be measured (as well as the frequency content of the signal), will determine how the structure responds. Some modes may be more dominant in the level of response and others may be less dominant in the response - again depending on the particular input-output location selected. But all the modes exist - they just may not all be activated on a uniform basis.

So what we need to remember is that a modal test only defines the character of the system. We apply an arbitrary force which is measured along with the response of the system due to the applied force. This enables us to determine the dynamic characteristics of the system - the frequency, damping and mode shapes. These are only characteristics of the system. We display the mode shape (animate them) to better understand how the structure may deform if a force is applied to the system that excites one or more modes of the system. Remember, modal analysis doesn't use the force on the right hand side of the equation - the mode shapes are independent of the force.

Now I hope you understand why you can't answer the question that you asked. If you have any other questions about modal analysis, just ask me.

Copyright 2000
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

I ran one test with an x-excitation and can see some modes
and another test with a y-excitation and see some different modes
could I use an oblique angle instead ?

Well, that's a very good question. Its one that comes up often in terms of running modal tests with a shaker excitation. Of course it is totally acceptable to run one test with a shaker at some oblique angle to the structure. But the only thing we need to be careful about is to assure that we don't select the reference point at the node of a mode. Let's talk about this a little more.

Let's start this discussion with a simple structure that has mode shapes that are very directional in nature. Now just what do I mean by that. That means that the response of the structure is primarily in one direction with very little or no response in the other directions for a given mode of the structure. Yet another mode of the structure may have response in a different direction than the first mode with little or no response in the other directions.

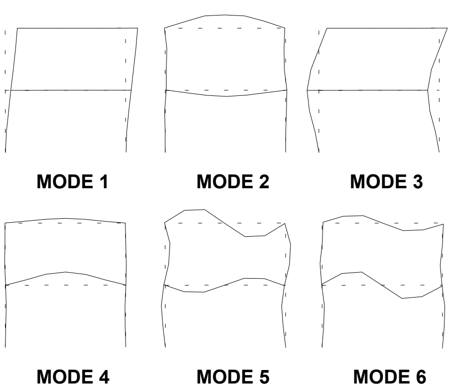

Figure 1

To illustrate this, let's consider the very simple frame shown in figure 1. We see that mode 1 of the structure has motion primarily in the horizontal direction with very little response in the vertical direction. However, mode 2 of the structure has motion primarily in the vertical direction with very little motion in the horizontal direction. We can also see that mode 3 and mode 4 follow the same trend. Mode 5 and mode 6 have motion in both the horizontal and vertical directions with the vertical direction being slightly more predominant.

If we look at a drive point measurement in the vertical direction (as shown in figure 2) over the bandwidth of the first six modes of the structure, we notice that there are only 2 peaks that are visible in the measured frequency response function. Yet we know that there are 6 modes in this frequency range. And if we took a drive point measurement in the horizontal direction we would also notice only 4 peaks. But upon closer examination of the measurement, we would notice that the first two frequencies of each of the measurements is different.

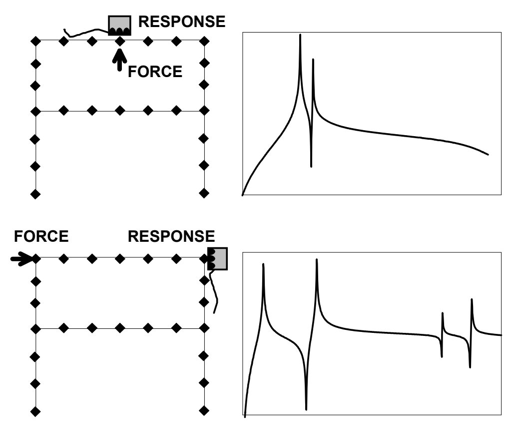

Figure 2

Copyright 2000
Pete Avitabile

So now we can see that the modes cannot be seen in every measurement. That directly implies that if we were to select either one of the two measurement points shown as a reference point then clearly we would not see all the modes of the structure. But just why does that happen.

Let's recall the equation for the frequency response function

$$h_{ij}(j\omega) = \sum_{k=1}^{m} \frac{a_{ijk}}{(j\omega - p_k)} + \frac{a_{ijk}^*}{(j\omega - p_k^*)}$$

which (for illustration) is expanded for 3 modes

$$h_{ij}(j\omega) = \frac{a_{ij1}}{(j\omega - p_1)} + \frac{a_{ij1}^*}{(j\omega - p_1^*)} + \frac{a_{ij2}}{(j\omega - p_2)} + \frac{a_{ij2}^*}{(j\omega - p_2^*)} + \frac{a_{ij3}}{(j\omega - p_3)} + \frac{a_{ij3}^*}{(j\omega - p_3^*)}$$

This equation is described by the residues (in the numerator) and the poles (in the denominator) for each of the modes of the system in the formation of the frequency response function. We also need to remember that the residues are directly related to the mode shapes (and a scaling factor) as

$$a_{ijk} = q_k u_{ik} u_{jk}$$

so that the frequency response function can be written either in terms of residues or mode shapes. When written as a mode shape, then it becomes very clear that if the value of the mode shape at the reference point is zero (or almost zero) then that mode will not be seen in the frequency response function. So the trick to performing a good modal test setup is to always select a reference point where all of the modes can be seen all the time from that reference point. But sometimes this is easier said than done, especially when I am not sure what the expected modes of the system are going to be. (Its always easy being a Monday morning quarterback.)

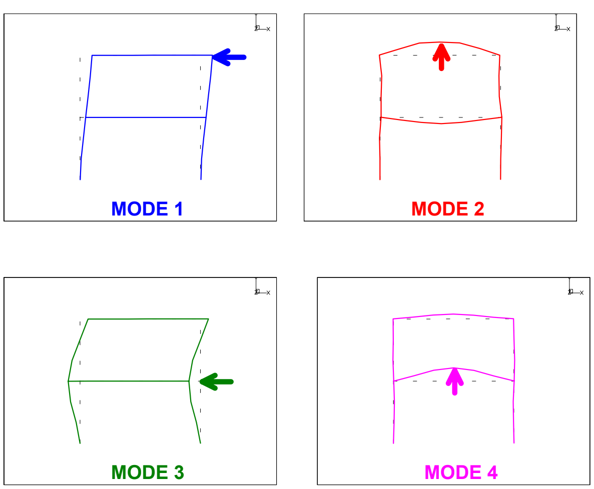

So now let's look at the first 4 modes separately. Then the optimum reference location is easily seen to be the point on the structure where the mode shape value is largest. But I quickly realize that the point is different for each of the modes. The trick is to select one point where each of the modes can be observed "reasonably well". The point where most people get hung up is in trying to think in terms of our simple rectangular coordinate system - with x, y and z directions.

What would happen if I tried to apply a force to the structure at a point that has a 45 degree angle to my global rectangular coordinate system. Would the reference point shown be suitable to be able to measure all the modes of interest?

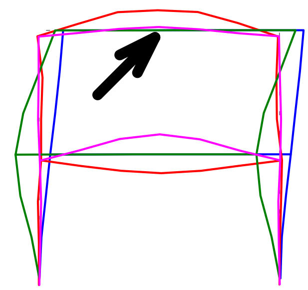

The best way to answer that question is to look at the equation that forms the frequency response function. We quickly notice that the equation can be written in terms of mode shapes. When we consider the first 4 modes of the system, we notice that each mode has a component of response in this 45 degree angle. This means that this reference point would be suitable for measuring the first 4 modes of the system. However, if I take a closer look and consider modes 5 and 6, then I will quickly discover that these two higher modes will not be measured at all from this reference point. This is because, the center point on the cross member is a node point for both modes 5 and 6 of the structure.

So we could pick any point on the structure including an angle relative to the global coordinate system. The only requirement is that the mode shape should have a significant value in relation to its mode shape; if the reference point selected is a node of a mode then I will not see that mode in the measured response. One last thing to quickly mention is that in order to obtain valid "scaled mode shapes", a drive point measurement is necessary. (We will talk about mode shape scaling in a future article.) This means that the response must be measured ***at the same point*** and ***in the same direction*** as the applied force - this also applies to a reference measurement which is taken at an angle to the global coordinate system.

I hope this explanation helps you to understand that you can pick any angle for the reference - just as long as its not the node of a mode. If you have any other questions about modal analysis, just ask me.

Copyright 2000
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

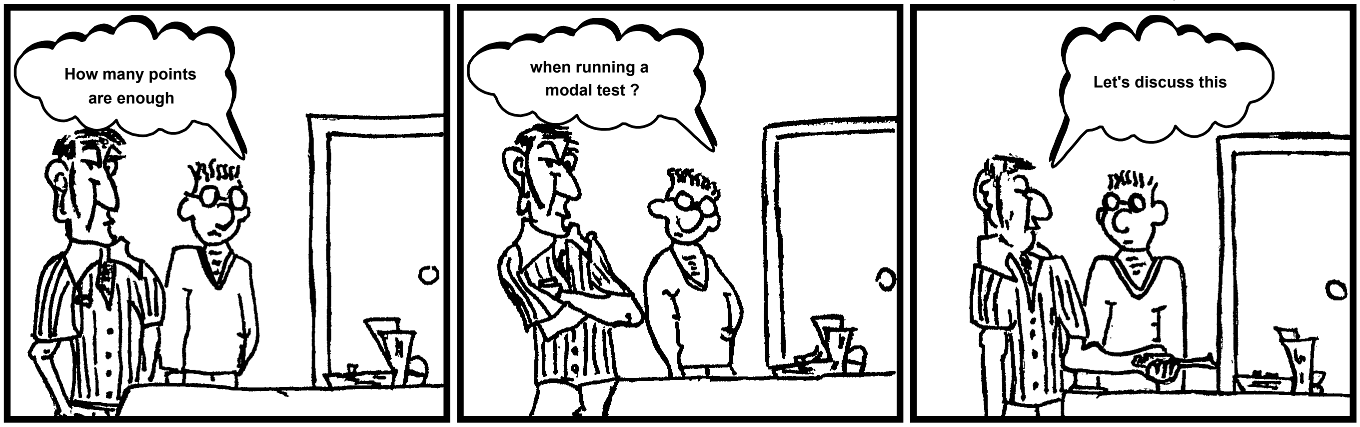

Illustration by Mike Avitabile

How many points are enough
when running a modal test ?
Let's discuss this

I expected that eventually you would get around to asking me that question. Its another one that I get asked all the time. Basically the simplest answer is that you need to measure a sufficient number of points so that you can uniquely describe the mode shape. This answer may not be completely obvious. We need to talk about this a little more.

Let's start with a simple structure that we have discussed before. The simple plate structure.

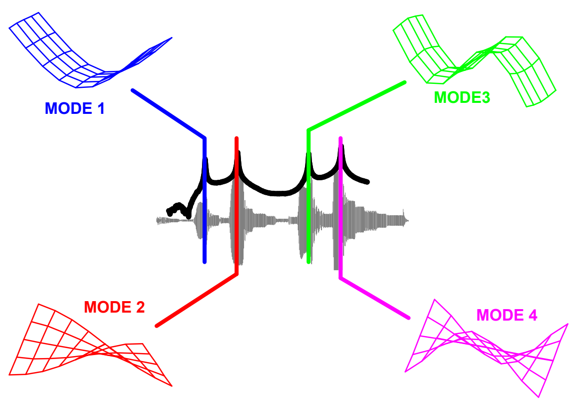

Now from the mode shapes shown we can see that there are sufficient number of points to describe the mode shape for each mode. But there are a total of 45 measurement locations on this plate.

Now let's consider only 5 points along one edge of the plate to illustrate some important points. If I look at mode 1 and mode 3, I quickly realize that there are not enough points to adequately describe the differences between the two modes. And also considering mode 2 and mode 4, the same conclusion can be drawn.

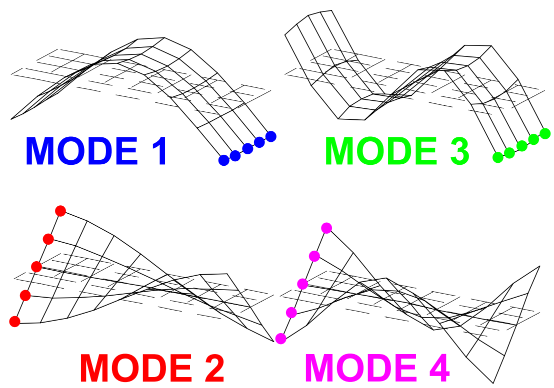

But I think we would all realize that only 5 points will not be sufficient to adequately describe the mode shape. Would it be possible to measure the mode reasonably well with only 15 measurement points? Well, most likely - but it is heavily dependent on ***where*** the 15 points are located. Let's consider 15 points - but I am going to pick the points to illustrate a point.

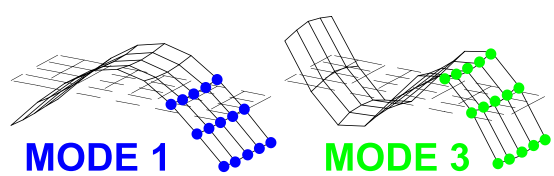

If I look at these 15 points then it is very hard to distinguish between mode 1 and mode 3. For all practical purposes, the mode shapes look almost exactly the same.

Now let's consider that we only took measurements along the front and back edges of the plate. You would be very hard

Copyright 2000
Pete Avitabile

pressed to tell the difference between the first rigid body mode and the first flexural mode.

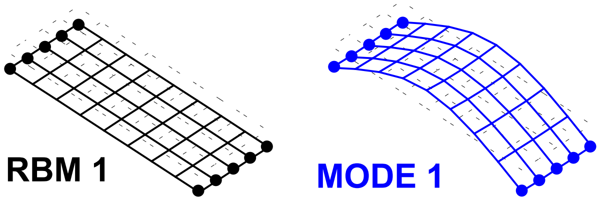

So from all of these simple examples above, it becomes obvious that we need a distribution of points located ***appropriately*** such that each mode shape can be uniquely distinguished. If I am only interested in characterizing mode 1 and mode 2, then possibly I could get a fairly good description with only 6 points as shown but fewer points than that would be difficult especially if we needed to distinguish the flexible modes from the rigid body modes.

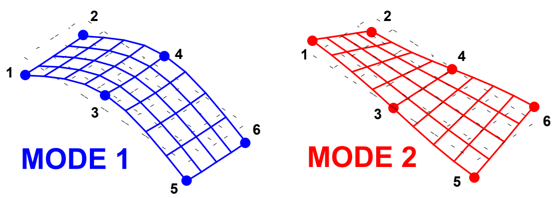

Now let's consider another example - the simple frame shown. Suppose that the only surfaces that are accessible are the three exterior surfaces. If I only collected measurements on those surfaces and did not have any measurements on the interior surface, then the description of the mode shapes may not be sufficient to uniquely describe the mode shapes. Just consider mode 2 and mode 4. In one mode, the two cross beams are out of phase with each other and for the other mode they are in phase with each other. The same is true for mode 5 and mode 6. If there are no measurements available on the interior surface then it is very difficult to distinguish the mode shapes for these modes. This is a common testing problem that occurs in many modal tests. Too few points are used to describe the mode shape due to inaccessibility of all the significant modally active portions of the structure.

Another common problem encountered in performing a modal test is the reluctance to measure adjacent portions of a structure. A typical comment that will be made is that we are only interested in a portion of a structure that we have responsibility for. We are not interested in the rest of the structure because it doesn't fall under our jurisdiction. To illustrate the problem with this statement we can also use the simple frame again.

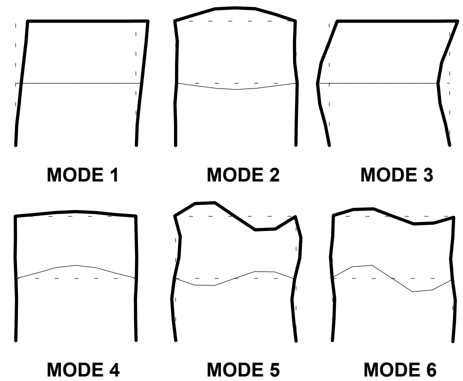

But this time measurements and modal data are only collected for the interior surface of the structure. We can quickly see that some of the mode shape information is strongly controlled by the exterior of the structure. If we don't measure enough information to fully describe the mode shape, then it may be very difficult to determine what the cause of the problem is when we blindly limit the data we look at.

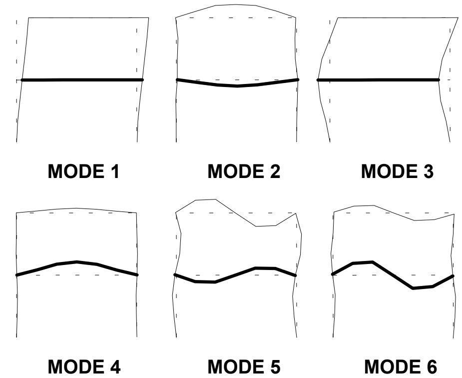

A good example of this brings to mind a recent modal test on a troublesome torsional vibration shaker system. The folks interested in solving their problem were only interested in the fixture and test article that came under their responsibility. Their impression of the shaker system and supporting structure was that they really didn't care what the outside world was doing. But in fact, the rest of the structure was actually responsible for the trouble that they were experiencing. The local response on the fixture and test article was largely due to some major global modes of the system. Without measuring this information, you would not have sufficient data to correct or understand the problem.

Now I hope you have a better understanding of how many points are needed for a modal test - a sufficient number to uniquely identify the mode. If you have any other questions about modal analysis, just ask me.

Copyright 2000
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Someone told me SDM will never work
Because you don't have all the mode shapes of the system
Well, let’s talk about this.

Structural Dynamic Modification (SDM) became a popular tool in the early 80's. Due to some misunderstandings of the technique, some erroneous results could be obtained. But given the right circumstances, SDM is a very powerful tool to help the design engineer make very good design decisions. First, let's briefly recall the technique and show how the technique can be sensitive to its biggest problem - modal truncation.

Basically, SDM is an analytical tool that uses modal data (either analytical or experimental) to estimate how the system dynamic characteristics will change when basic changes in the mass damping, stiffness of the system are investigated. Note that only modal data (frequency, damping and mode shapes) are used for the prediction - the original FEM or test data need not be modified to explore these changes. However, once a set of desired changes are obtained, then it is strongly recommended to re-run the modified FEM or re-test the modified test article.

The physical system equations can be developed and the eigensolution obtained. The modal representation can be obtained from either an analytical model or from test data. The modal representation of a physical system in modal space is given by

$$\begin{bmatrix} \ddots & & \\ & \overline{M} & \\ & & \ddots \end{bmatrix}\{\ddot{p}\}+\begin{bmatrix} \ddots & & \\ & \overline{C} & \\ & & \ddots \end{bmatrix}\{\dot{p}\}+\begin{bmatrix} \ddots & & \\ & \overline{K} & \\ & & \ddots \end{bmatrix}\{p\}=[U]^T\{F\}$$

Now changes to the physical system mass, ΔM, damping, ΔC, and stiffness, ΔK, can be represented in modal space (through the modal transformation equation) as

$$\left[\Delta\overline{M}\right]=[U]^T[\Delta M][U];\left[\Delta\overline{C}\right]=[U]^T[\Delta C][U];\left[\Delta\overline{K}\right]=[U]^T[\Delta K][U]$$

Assuming a proportionally damped system, an eigensolution can be obtained for the modified system. One important part of this solution is the computation of the final physical modes of the system from

$$[U_2]=[U_{12}][U_1]$$

which implies that the final modified modes of the modified system are made up from linear combinations of the unmodified modes of the original system. It is this important equation that we will use to show the effects of truncation of the predicted results.

Let's consider a simple example of a free free beam which we will use to make two simple structural changes - a simple support and a cantilever beam. We will modify the structure using two springs to ground and perform the SDM equations to obtain the modified frequencies and mode shapes. The original unmodified frequencies and resulting modified frequencies are shown in the Table 1 (Note: The frequencies identified in italics are an approximation of the constraint modes of the system and are beyond the scope of this discussion).

Notice that the simple support produces very accurate modified modes using only the first 5 modes of the original unmodified system whereas the cantilever beam does not. Why does the simple support do so well and the cantilever does not? The answer lies in how the mode shapes are formed from the original system modes.

Copyright 2000
Pete Avitabile

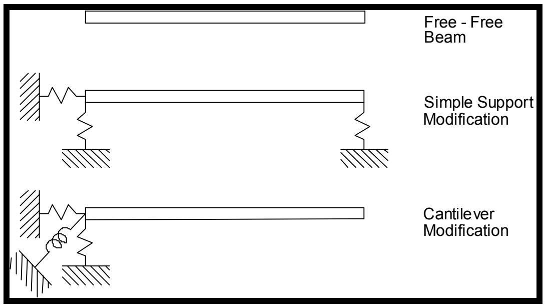

Fig 1 - Models Evaluated

		Simple Support		Cantilever	
#	Free	Ref.	SDM	Ref.	SDM
1	0.	71.9	72.0	21.6	24.8
2	0.	285.7	288.4	139.3	162.8
3	128.	636.5	646.0	396.1	476.0
4	367.	1114.9	*9108.3*	781.8	*1274.5*
5	738.	1706.3	*9593.6*	1292.0	*9437.8*

Table 1 - Frequencies of Different Systems

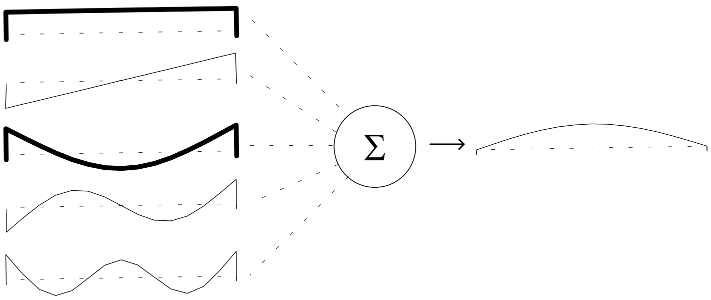
Fig 2 - First Simple Support Mode

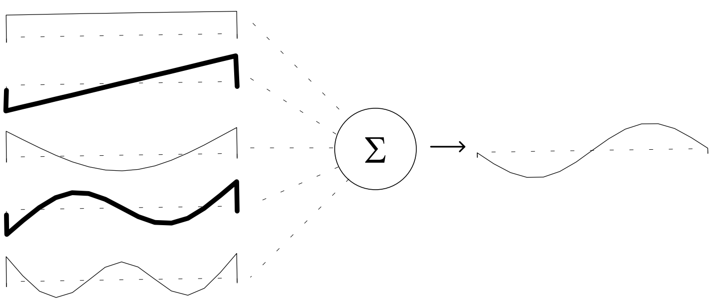
Fig 3 - First Simple Support Mode

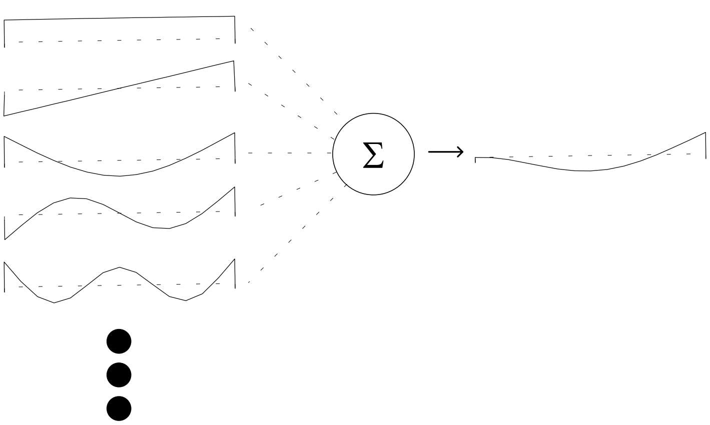
Fig 4 - Second Cantilever Support Mode

The simple support modified modes are easily made up from linear combinations of the unmodified modes of the original system. When we look at Figure 2, we notice that mode 1 and 3 are the most significant contributors to the first final modified mode for the simple support beam. And when we look at Figure 3, we notice that modes 2 and 4 are the most significant contributors to the second final modified mode for the simple support beam.

But when I consider the modes of the cantilever beam modification, there is a significant contribution from all 5 modes of the unmodified system. In fact, many more modes are needed to improve the accuracy of this cantilever predicted modes. (Note: Mode 2 is shown in Figure 4 for the cantilever)

It turns out that the simple support can be easily made from the available linear combinations of the 5 free-free modes of the original system whereas the cantilever can not! So that fact that all the modes are not available (modal truncation) is not always a problem. The real problem is that the final modified modes must be able to be formulated from the original unmodified modes.

Another important item to note is that the rigid body modes of the free-free beam are very important to the accurate prediction of the modified modes. If the rigid body modes are not available, then the predicted modes will be in error. This is an important consideration for the development of the experimental modal database since, often times, rigid body modes are not acquired as part of the test. However, it can be easily seen that the rigid body modes are very important to the success of the modification, even for the case of the simple support modification.

The bottom line is that in order to compute an accurate modified model using SDM, the final modes must always be made up of linear combinations of the unmodified modes. If this is possible, then good results can be obtained. If not, then errors will result due to modal truncation.

Without getting into all the detailed equations, some simple graphics were used to illustrate how SDM uses the unmodified modes of the system to obtain estimates of the modified modes of the system. I hope that this helps you to understand how SDM could be affected by modal truncation. If you have any other questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

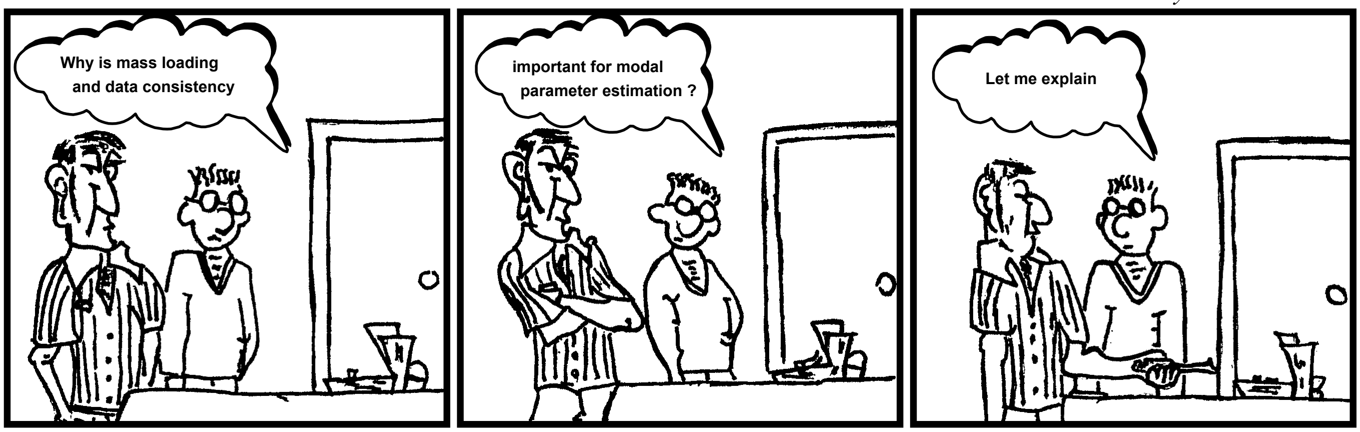

Illustration by Mike Avitabile

Why is mass loading and data consistency important for modal parameter estimation? Let me explain

This is other good example where people can get confused when performing modal parameter estimation. All too often when the curvefitting results are confusing or appear distorted, the effects will be blamed on noise or nonlinearities. This is often a blanket statement that many people use when they don't understand or can't explain something easily. Let's look at why data consistency is important and what effects mass loading will have.

The first thing to recall is that the model we use to fit data comes from a linear, symmetric set of equations where the poles (frequency and damping) are defined in terms of global quantities and reciprocity is assumed to be inherent in the formulation of the equations. Now as long as our data fits that model then everything is OK. But how does my testing and data acquisition have an effect.

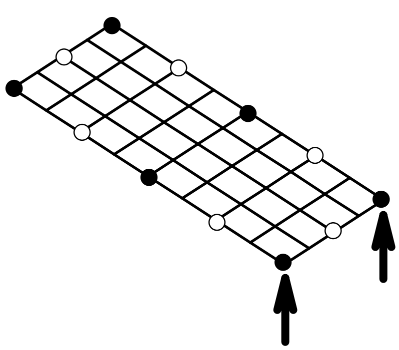

Fig 1 - MIMO Test Setup with 2 Sets of Points

Let's consider a simple plate test setup that is driven by two shakers for a MIMO test with an 8 channel data acquisition system. Now I'll acquire FRFs using good measurement techniques to assure the best possible measurements are obtained for the 6 accelerometers mounted on the plate shown in Figure 1 (the solid fill points are for the first test and the other points are associated with the second test and are obtained by roving the accelerometers on the structure).

The mode indicator function is shown in Figure 2 and the stability diagram is shown in Figure 3. The poles are extracted for the first two modes only (for illustration purposes). The stability diagram shows these two poles very clearly. Notice that as the order of the model increases, the poles are clearly identified (overlaid on the summation function). Once the poles are extracted, then the residues or mode shapes are obtained to provide modal data associated with these 6 measurement points; a typical curvefit is shown in Figure 4.

However, this first set of data only consists of 6 measurement points. In order to better define the mode shapes, more measurement points are needed.

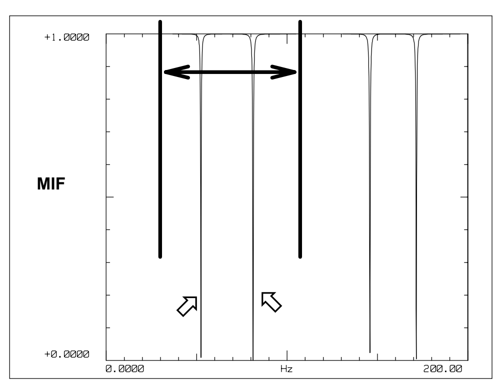

Fig 2 - MIF for Data from First Test

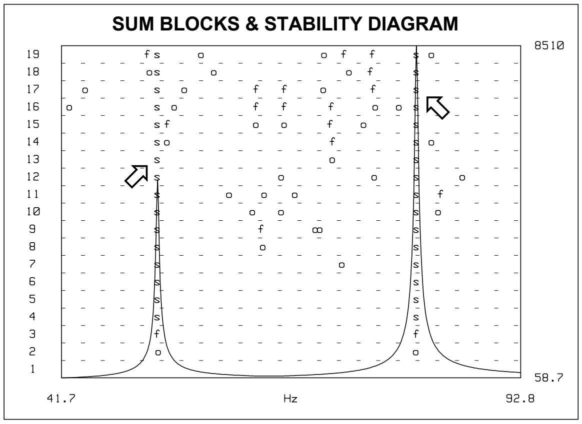

Fig 3 - Stability Diagram for Data from First Test

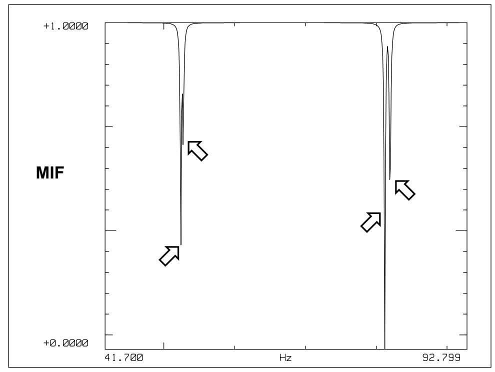

Fig 5 - MIF for Test 1 & 2 Combined

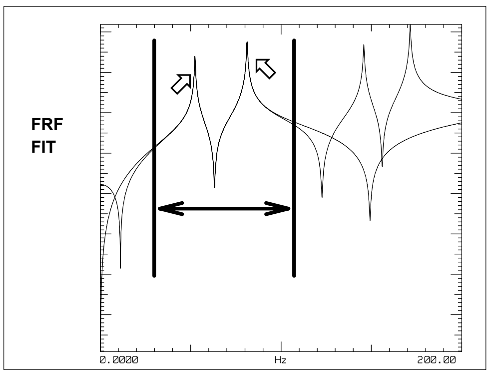

Fig 4 - Typical Curvefit from First Test

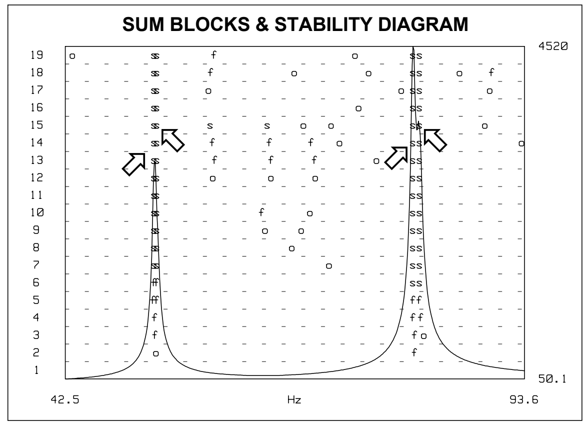

Fig 6 - Stability Diagram for Test 1 & 2 Combined

For the additional points, the accelerometers are relocated to the measurement points shown (non-filled points) and a second set of MIMO measurements were collected. Again, poles are extracted *using just this second set* of measurement points and a stability diagram obtained. Again the poles are clearly identified and mode shapes associated with these 6 points identified. (These results are not shown here but are similar to the first case.) *But the two sets of data were evaluated separately to estimate the poles and residues.*

Now let's combine the two data sets together and evaluate the data. The mode indicator function and stability diagram are computed again. Now instead of 2 distinct peaks as we saw earlier in the MIF, there are now 4 distinct peaks over the same band (Figure 5). The estimation of the poles for the same frequency band (Figure 6) used earlier now shows 4 modes instead of 2!!! How could this possibly be? The plate didn't change - did it?

Well, the plate didn't change - ***but the test setup sure did!*** The roving accelerometers have a mass effect that caused the modes to shift slightly. So when all the data is processed simultaneously, some of the measurements indicate the poles at a certain frequency and the other measurements indicate the poles at a different frequency.

So which is correct? It is likely that neither is correct. That's because the test setup had an effect on the measured modes of the system. The question is which poles are the correct ones to be used for the modal parameter estimation process. Well, you really can't identify a global set of poles for all the measurements since they are not "global" for all the measurements. Actually, the correct way to extract parameters in this case is to collect a "consistent" set of data by eliminating the mass loading effect by mounting all the instrumentation on the structure (or adding dummy masses) for the duration of the test. This will provide more "consistent" data which conforms to the model being used to fit the data. Of course, it is very important to point out that we have modified the structure due to the addition of all of the masses. But at least all the data will be consistent and will not distort the modal parameter estimation process due to mass loading effects.

Of course, real world structures have all kinds of measurement problems with respect to noise, linearity, time variability, etc. The modal parameter estimation process is complicated enough. Don't complicate the process further by letting simple items such as mass loading distort your data. I hope this helps to answer your question as to why mass loading and data consistency is so important. If you have any other questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

Illustration by Mike Avitabile

I hear about SVD all the time
Could you explain it simply to me?
Sure ...

I'm surprised you haven't asked this question sooner. SVD, singular valued decomposition, is probably one of the most important linear algebra tools that we use today to solve many of our structural dynamic problems. First let's present the mathematical formulation of SVD and some of its variations and then describe where it is commonly used in experimental modal analysis. Of course, I will try to explain the use of SVD and its use rather than give a detailed mathematical development.

First we have to realize that we are going to be dealing with matrices here. (I know you all shudder when we say matrices - but as I have said before "Matrices are your friends!".) So let's assume that we have some matrix [A] that is a n x n square matrix. The basic SVD equation is

$$[A]=[U][S][V]^T$$

Now this formulation looks pretty simple but let's expand out some of these terms to see the real power of SVD

$$[A]=\left[\{u_1\} \quad \{u_2\} \quad \{u_3\} \quad \cdots\right]\begin{bmatrix} s_1 & & & \\ & s_2 & & \\ & & s_3 & \\ & & & \ddots \end{bmatrix}\left[\{v_1\}^T \quad \{v_2\}^T \quad \{v_3\}^T \quad \cdots\right]$$

The expansion of this gives

$$[A]=\{u_1\}s_1\{v_1\}^T+\{u_2\}s_2\{v_2\}^T+\{u_3\}s_3\{v_3\}^T+\quad\cdots$$

Now that's pretty incredible because it implies that the matrix A is made up of a set of vectors and singular values that describe the matrix.

We could also say that there are parts of the matrix A that are comprised of other matrices who are very simply described as one vector and a corresponding eigenvalue. So the SVD really has the ability to determine the "principal pieces" that comprise the matrix. This also implies that the rank of the matrix can be determined. So let's try a few numbers here to see what this means.

Let's start with a simple vector with an eigenvalue to illustrate the basic SVD equation. Let's define a vector with a singular value as

$$u_1=\begin{Bmatrix}1\\2\\3\end{Bmatrix};\ s_1=1$$

So the matrix A can be found by simply multiplying out these terms to be

$$[A]_1=\begin{bmatrix}1 & 2 & 3\\ 2 & 4 & 6\\ 3 & 6 & 9\end{bmatrix}$$

So this is pretty neat because I started with a vector and I formed a matrix. Now this matrix is clearly a 3x3 matrix in size, but what can I say about its rank? Well, if I look at the different rows of the matrix, I can very quickly see that row two and three are linearly related to row 1. That means that while I have a 3x3 matrix, there is only one linearly independent piece of information that makes up this matrix. (Of course, we know that this is true since we made the matrix from one vector). We would then say that this matrix has a rank of 1 - because there is only one linearly independent piece of information that makes up this matrix.

Copyright 2000
Pete Avitabile

Now let's consider another simple vector with an eigenvalue as

$$u_2 = \begin{Bmatrix} 1 \\ 1 \\ -1 \end{Bmatrix}; \; s_2 = 1$$

So the matrix A can be found by simply multiplying out these terms to be

$$[A]_2 = \begin{bmatrix} 1 & 1 & -1 \\ 1 & 1 & -1 \\ -1 & -1 & 1 \end{bmatrix}$$

Again I make all the same comments about this matrix as I did for the first matrix we looked at. The rank of this matrix is 1 because it is made up from one linearly independent piece of information.

Now let's consider a general matrix as

$$[A]_3 = \begin{bmatrix} 2 & 3 & 2 \\ 3 & 5 & 5 \\ 2 & 5 & 10 \end{bmatrix}$$

Now this matrix is a 3x3 but it is not clear to me what its rank is. The simpliest way to determine this is to do an SVD on this matrix. The resulting decomposition is

$$[A] = \begin{bmatrix} \begin{Bmatrix} 1 \\ 2 \\ 3 \end{Bmatrix} & \begin{Bmatrix} 1 \\ 1 \\ -1 \end{Bmatrix} & \begin{Bmatrix} 0 \\ 0 \\ 0 \end{Bmatrix} \end{bmatrix} \begin{bmatrix} 1 & & \\ & 1 & \\ & & 0 \end{bmatrix} \begin{bmatrix} \{1 & 2 & 3\} \\ \{1 & 1 & -1\} \\ \{0 & 0 & 0\} \end{bmatrix}$$

So the beauty of SVD is that I can write the matrix A in terms of the linearly independent pieces that make up the matrix. This can be expressed in summation form as

$$[A] = \begin{Bmatrix} 1 \\ 2 \\ 3 \end{Bmatrix} 1 \{1 \;\; 2 \;\; 3\} + \begin{Bmatrix} 1 \\ 1 \\ -1 \end{Bmatrix} 1 \{1 \;\; 1 \;\; -1\} + \begin{Bmatrix} 0 \\ 0 \\ 0 \end{Bmatrix} 0 \{0 \;\; 0 \;\; 0\}$$

So I think this helps to explain the basic principles of SVD. But now I need to discuss some of the applications where SVD is commonly used. (There are many different applications for SVD but only a few specific ones related to experimental testing issues are addressed.)

One application of SVD is for the collection of MIMO data for an experimental modal test. While the data acquisition system may generate forcing functions for all of the MIMO shakers that are uncorrelated (linearly unrelated), the actual shaker force excitation may not be completely uncorrelated for each of the shakers due to the interaction of the shakers with the structure.

The linear independence of the input spectrum matrix needs to be checked. During the acquisition of MIMO data, the Gxx matrix of the shakers can be used to perform what is commonly called a principal component analysis. This technique decomposes the Gxx matrix using SVD and then plots the singular values for each of the inputs on a frequency basis. If the shakers are all linearly independent, then there will be a significant singular value at all frequencies for each of the independent inputs. This is shown in Figure 1.

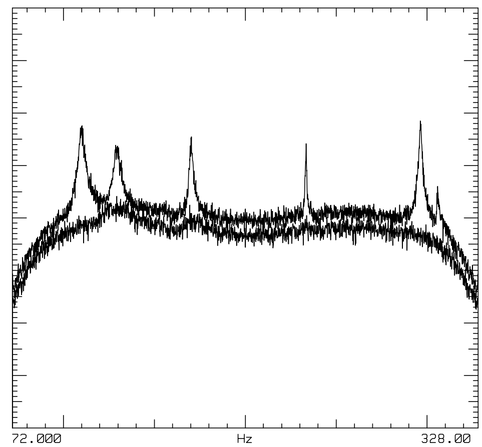

Figure 1 - 2 Shaker MIMO SVD

Another application is used in modal parameter estimation. The FRF matrix from several different references can be decomposed using SVD to determine where there are roots (or modes) of the system. This decomposition is the basis of the CMIF modal parameter estimation approach. A plot of the significant singular values of this SVD will provide plots which will indicate where the modes of the system are located. A typical plots of this is shown in Figure 2 for a system that has repeated roots.

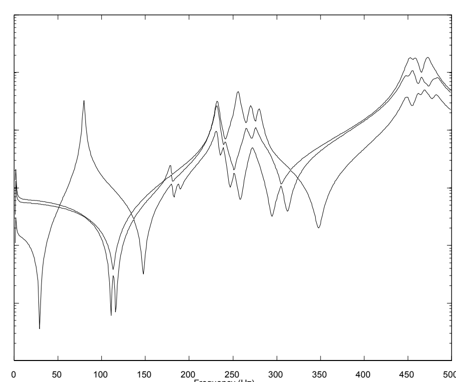

Figure 2

While there are many more applications of SVD, I hope that these few examples help you better understand the technique. If you have any other questions about modal analysis, just ask me.

Copyright 2000
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

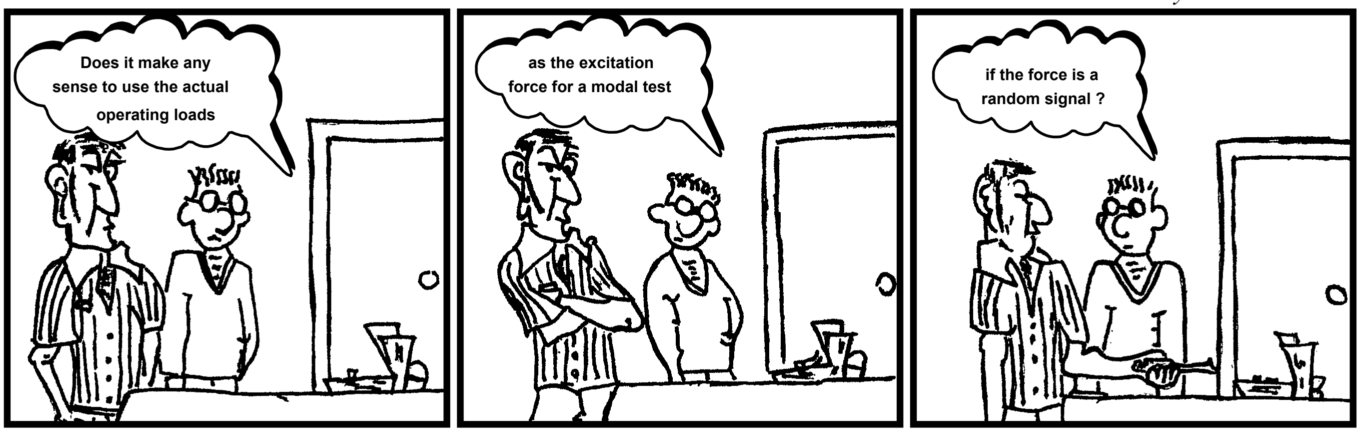

Illustration by Mike Avitabile

Does it make any sense to use the actual operating loads
as the excitation force for a modal test
if the force is a random signal ?

The answer to this question is not an easy one. There are many aspects related to this question that we need to discuss in order to fully understand the answer.

The use of a random operating excitation may seem to be an excellent idea, but the bottom line is that the modal parameters that are extracted are not likely to be nearly as good as those obtained from a modal test where the excitation is one of the more traditional excitation techniques. Let's discuss this to see where some of the pitfalls exist. In order to understand all of the implications, there have been some other modal questions that have been asked and answered that will help shed light on this question (SEM ET V23 No1, V23 No4, V23 No6).

Let's recall that an experimental modal test is typically performed to extract the underlying modal parameters of the structure - that is, the frequency, damping and mode shapes. Accurate measured frequency response functions are needed in order to extract these parameters. Typically, we go to extreme lengths to excite the structure with very specialized excitations to minimize, and ultimately eliminate, leakage and other signal processing errors that can possibly result. Remember that any signal processing errors that do result, distort the measured frequency response and manifest themselves as less accurate modal parameters.

As a general rule, random signals do not provide the best excitation for the development of accurate frequency response functions. Random excitation techniques are notorious for causing leakage in the measured spectra. Even with the use of windows, the measured frequency response functions are distorted when compared to other leakage-free measurement techniques (ie, burst random, sine chirp, digital stepped sine).

A comparison of a frequency response function from a random excitation and a burst random excitation is shown in Figure 1. It is very clear in the measurement that the burst random, leakage free measurement is far superior to the random measurement. (While not shown, the coherence is also far superior.)

To go one step further, the extracted modal parameters from the random excitation will also be distorted, and in many cases, there actually appears to be two peaks as seen in the measurement. This is a typical effect seen in frequency response functions measured using random excitation. Leakage is a serious concern and windows are necessary to minimize leakage. The whole purpose for the development of specialized functions for modal testing is to provide highly accurate frequency response functions which do not require the use of any windows and provide leakage free measurements for the accurate extraction of modal parameters.

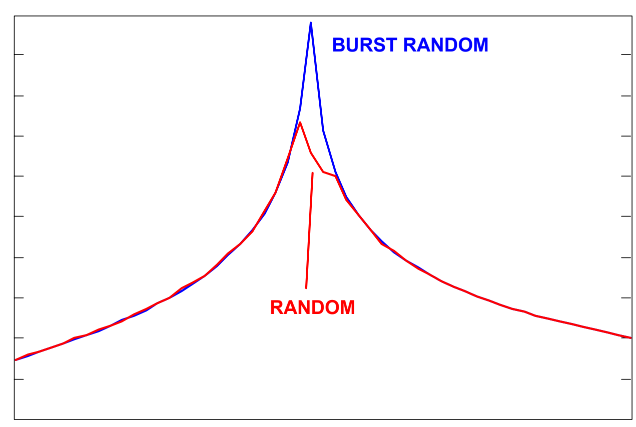

Figure 1 - FRF for Burst Random and Random

Copyright 2001
Pete Avitabile

So what would ever posses anyone to perform a modal test using an operating random excitation. Well, if the actual force was used to excite the structure, then the response will be similar to the actual response in service. This response will be an accurate depiction of the actual in-service deformations that will be seen in the structure. But then the response that is measured is more appropriate for use in an operating deflection analysis - but not an experimental modal survey!

Figure 2 shows a schematic of the response of a structure due to an arbitrary input excitation. There are several aspects of this figure that will give greater insight into the question at hand. The forcing function is broadband, but has a very distinct profile which is not flat, thereby exciting all of the modes with different excitation levels.

First, and foremost, notice that the frequency response function is nothing more than an bandpass filter which amplifies and attenuates the input force excitation as a function of frequency. What would happen if the estimation of this frequency response was tainted or distorted by the digital signal processing procedure (ie, digitization, quantization, leakage, windows, FRF method, etc.) ??? Well, of course, there would be an effect on the computed response! The goal of a modal test is to extract the accurate dynamic system characteristics.

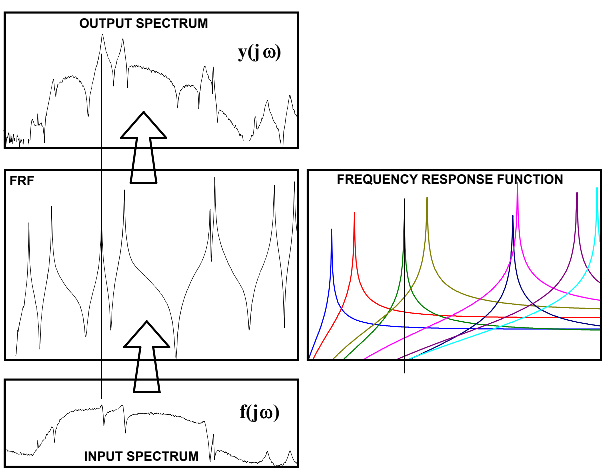

Figure 2 - Typical Input-Output Situation

Second, the level of the force spectrum over the frequency band has a direct effect on the response of the system. Figure 3 very clearly shows that the response has significant variation over the frequency band. Since the ADC maximum setting is determined by the total spectrum, there will be a wide variation in the accuracy of the measured function. In fact, the lower response spectral components will have a much larger effect due to quantization errors associated with the analog to digital conversion process. This is particularly true when looking at the response of mode 1 and mode 3. Notice that mode 1 shows very little response due to the extremely low input excitation; mode 1 response will be very small and may be affected by noise.

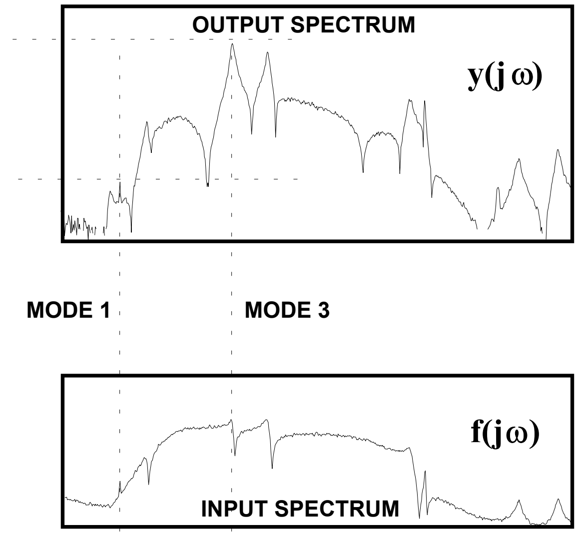

Figure 3 - Signal Level Differences

Third, remember that if a random signal is used, then a Hanning window must be applied otherwise the measured signals will contain significant leakage. In any event, the measured frequency response function will be affected by the window and leakage that does result. The measured function will not be of the best quality and the extracted modal parameters will suffer from these signal processing effects.

Fourth, the measured frequency response function will have errors associated with poor excitation signal strength over some frequency regions, leakage and window errors due to the random nature of the signal type, frequency response function errors as seen in the coherence associated with leakage especially at the resonant peaks, and modal parameter estimation errors due to poor estimated frequency response functions used for the modal parameter estimation process.

So in the big picture of the development of a modal model from measured functions, the best excitation techniques will provide the best representation of the modal parameters of the structure. This will not necessarily occur using an operating random spectrum. Once a modal model is developed, then the actual response of the structure can be determined, if necessary, using the measured frequency response function as pictorially shown in Figure 2. But in order for accurate response to be computed, an accurate modal model from accurate frequency response functions is of paramount importance.

Now, there is tremendous merit in performing an operating test using operating excitations. However, this is not necessarily the best way to estimate frequency response functions for use in the development of a modal model. Now I hope you understand the problems associated with running an experimental modal test with an operating excitation. If you have any other questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

When impact testing, can the use of the exponential window cause any problems?
Let's discuss this

The exponential window can cause some problems if not used properly. If an excessive amount of damping is needed to minimize the effects of leakage, then you run the risk of missing closely spaced modes. There are a few examples to show the relative to the use of the window and what can happen if care is not exercised in using the exponential window.

First of all, let me clearly state that in many impact testing situations the use of an exponential window is necessary. However, before any window is applied, it is advisable to try alternate approaches to minimize the leakage in the measurement. Increasing the number of spectral lines or halving the bandwidth are two things that should always be investigated prior to using a damping window. Both of these items will essentially increase the total time for the collected data. This can often help by allowing the response of the system to naturally decay before the end of the sample period. If this can be accomplished, then the use of the exponential window may not be necessary.

However, if the response still does not decay by the end of the sample period, then an exponential window may be necessary. The use of the window should not be employed until these first two items (mentioned above) are checked as possible ways to minimize the leakage problem. The arbitrary use of the exponential window without first looking at the time response is not recommended as the first step in the measurement process. Let's look at this through the use of a simple example.

A very simple, lightly damped structure was subjected to an impact test. The signal processing parameters were selected for a 400 Hz bandwidth which resulted in a 1.0 second time window. Since the structure was expected to have a response that would not decay by the end of the sample interval, an exponential window was applied such that the windowed response would decay to a reasonably small value by the end of the sample interval thereby minimizing the effects of leakage. The impact excitation, windowed exponential response and the FRF are shown in Figure 1. On the surface, this measurement looks acceptable. [Note that the input spectrum (not shown) was reasonably flat over the entire frequency range thereby allowing sufficient excitation of the structure. Also note that the coherence (not shown) was also considered very acceptable.]

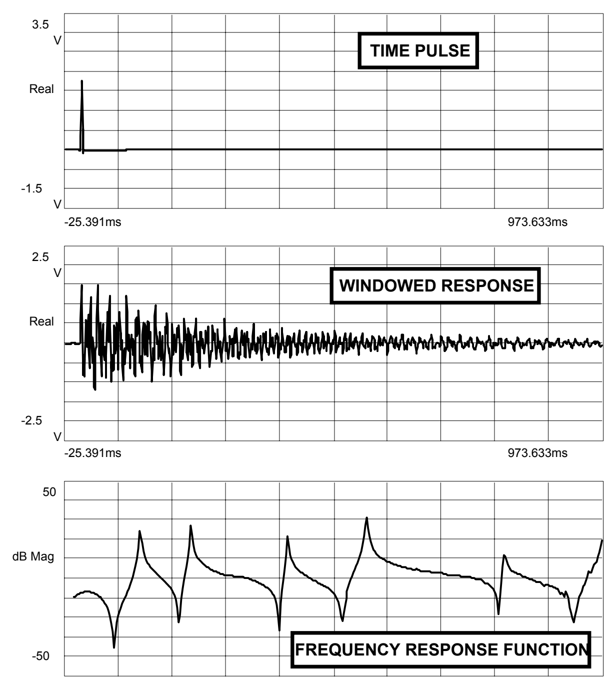

Figure 1 - FRF with slightly too much damping

Copyright 2001
Pete Avitabile

From all aspects, this measurement appears very acceptable. But we need to look at this measurement in more depth. First, let's consider the same measurement but add significantly more damping to the response signal. Figure 2 shows the same data but with a significantly larger value of damping used for the exponential window. The FRF that results from the impact measurement of this signal clearly has significantly more damping than that shown in the FRF of Figure 1. The peaks of the FRF show this effect; notice that the peaks are much wider due to the excessive use of the damping window.

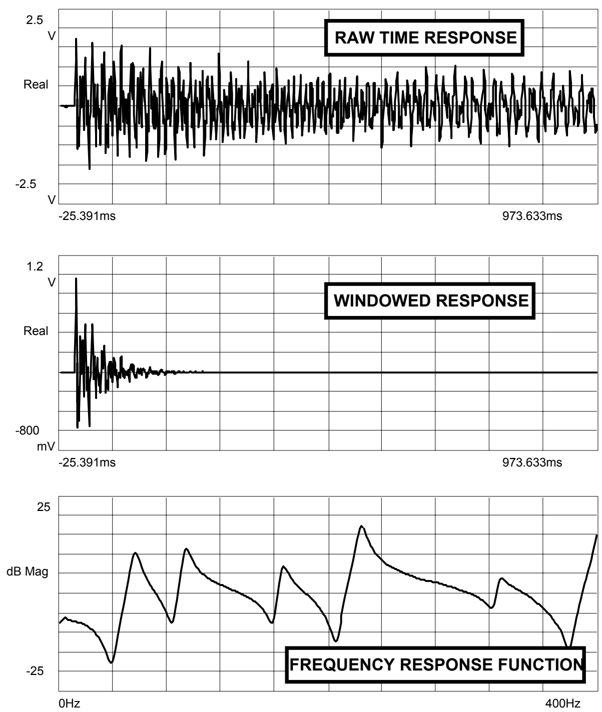

Figure 2 - FRF with too much damping

Now let's look yet a little deeper into this measurement and try some alternate signal processing parameters. In order to minimize the use of the damping window, either the bandwidth can be shortened or the number of spectral/time lines of resolution can be increased. Both of these changes result in an increase in the total time necessary to collect the sample of data. If the total time is increased, then there is less need for a significant amount of damping window to be applied to the collected time data.

Figure 3 shows a doubling of the number of spectral/time lines of resolution. The time sample was increased from 1.0 second to 2.0 seconds. While an exponential window was still necessary to minimize leakage, the overall damping effect that was added to the measurement is far less than that used for the measurements shown in Figure 1 and Figure 2 above.

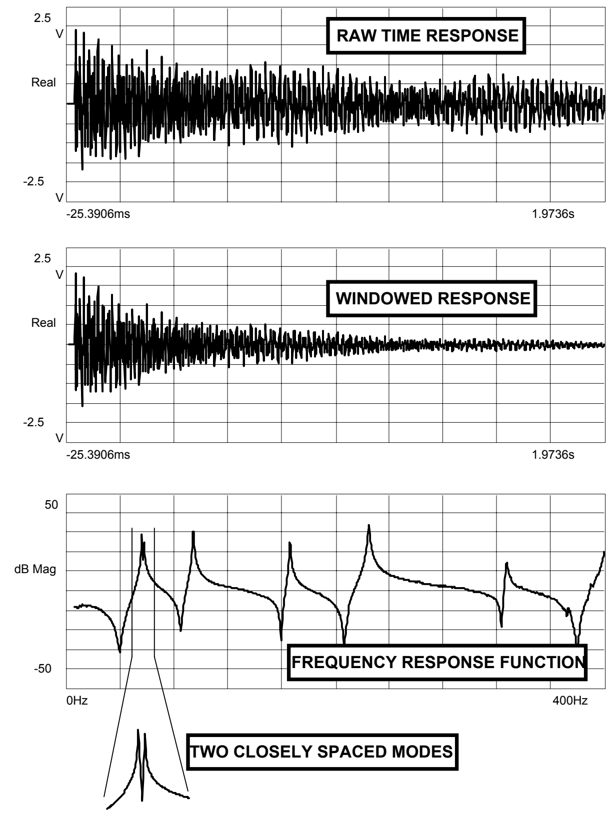

Figure 3 - FRF with increased time/spectral resolution

The most important item to notice in Figure 3 is that what appeared to be one mode at the first peak in the FRF actually turns out to be two very closely spaced modes of the structure. The use of the damping window in Figure 1 and 2 resulted in an FRF that appeared to have only one mode at the first peak in the FRF. The use of the damping window caused these two distinct modes to appear as only one peak in the FRF.

While the damping window was necessary to minimize the leakage, the window distorted the actual FRF in Figure 1 and 2 such that it was difficult to observe that two peaks existed at this frequency. The use of the exponential window, while necessary for digital signal processing considerations, can cause some significant difficulties when evaluating structures with light damping and closely spaced modes as seem in this example.

Now, I hope you can see some of the effects of the exponential window in this example. While an exponential window may be necessary to minimize the effects of leakage, the use of the window may also hide or distort the modes in the measurement. It is extremely important to be very careful when using the exponential window when performing an impact test. If you have any other questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Is there any benefit to using multiple references?
I thought only one reference was needed
Let's discuss this.

This is a very good question. Its one that comes up often in terms of estimating modal parameters from test data. From modal analysis theory, we can easily show that only one reference is necessary in order to determine all of the modes of a system - at least from a theoretical standpoint! While theoretically this is true, from a practical standpoint, there is a strong need to have multiple references in many cases. Before we can understand this, let's take a look at some basic concepts that will help illustrate some of the problems that we might encounter.

Let's start this discussion with a simple structure that has mode shapes that are very directional in nature. We have used this structure before in other discussions (May/Jun 2000, Vol. 24, No. 3). The structure is shown in Figure 1 along with the first several modes.

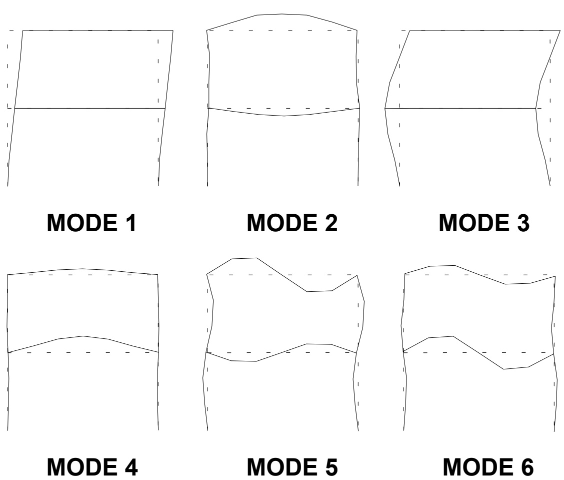

Figure 1

If we look at a reference point in the vertical direction (as shown in figure 2) over the bandwidth of the first six modes of the structure, we notice that there are only 2 peaks that are visible in the measured frequency response function. Yet we know that there are 6 modes in this frequency range. And if we took a reference point in the horizontal direction, we would also notice only 4 peaks. But upon closer examination of the measurement, we would notice that the first two frequencies of each of the measurements is different.

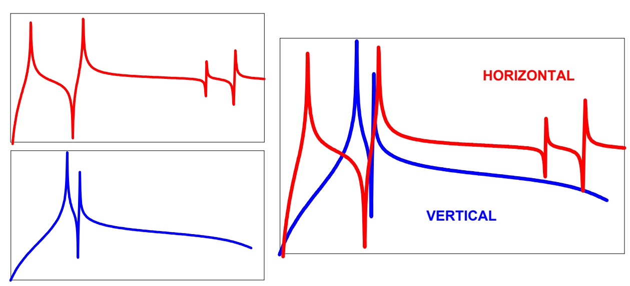

Figure 2

Let's recall the equation for the frequency response function

$$h_{ij}(j\omega) = \sum_{k=1}^{m} \frac{a_{ijk}}{(j\omega - p_k)} + \frac{a_{ijk}^*}{(j\omega - p_k^*)}$$

Basically, this equation is described by the residues (in the numerator) and the poles (in the denominator) for each of the modes of the system. We must remember that this frequency response function can be written for any one of the input-output combinations of interest. Now the interesting part of this equation is that while the residues change depending on which input-output combination is acquired, the poles do not change.

Copyright 2001
Pete Avitabile

This implies that the poles of the system are global. They are independent of the particular input-output point. However, the residues do, in fact, change.

Now when a modal test is performed, typically all of the measurements are acquired relative to a particular reference. The reference location is typically either the fixed excitation location when performing a shaker excitation or the stationary accelerometer location when performing an impact test. So the measurements acquired will contain residues, relative to a particular reference as

$$\lfloor a_{11} \quad a_{12} \quad a_{13} \quad \cdots \rfloor_k$$

In this case, the reference is "1" since all of the residues are related to that DOF. The residues are a_{11}, a_{12}, a_{13} and so on. (Note that the "k" subscript is used to denote a particular mode of the system.)

We also need to remember that the residues are directly related to the mode shapes (and a scaling factor) as

$$a_{ijk} = q_k u_{ik} u_{jk}$$

This means that the residues are actually directly related to the mode shapes of the system as

$$\begin{Bmatrix} a_{11k} \\ a_{21k} \\ a_{31k} \\ \vdots \end{Bmatrix} = q_k \begin{Bmatrix} u_{1k} u_{1k} \\ u_{2k} u_{1k} \\ u_{3k} u_{1k} \\ \vdots \end{Bmatrix} = q_k u_{1k} \begin{Bmatrix} u_{1k} \\ u_{2k} \\ u_{3k} \\ \vdots \end{Bmatrix}$$

Notice that the reference DOF at point 1 can be factored out since it is common to all of the measurements. In doing this, it becomes very clear that the reference DOF carries a tremendous amount of weight regarding the magnitude of the residue; this is directly related to the magnitude of the frequency response function. If the reference point is associated with a very small mode shape response location on the structure for a particular mode, then the magnitude of the frequency response function will also be very small for that mode. On the other hand, if the reference point is associated with a very large mode shape response location then the magnitude of the frequency response function will be very large.

Of course, we can then also see that if the reference location is located at a DOF where the mode shape value is very large for one mode and very small for another mode, then the amplitude of the frequency response function will have the same attributes. This is a common problem in performing any modal test. We always try to locate the accelerometer at a location where all of the modes can be observed with the same strength across the desired frequencies of interest. However, this is often very difficult and, in many cases, almost impossible.

However, we can use some of the redundancy in the frequency response matrix to help with this situation. If we look at some of the terms of this matrix, then there are some interesting things to note. The residue matrix is shown along with some of the terms expanded for reference.

$$[A(s)]_k = q_k \{u_k\} \{u_k\}^T$$

$$\begin{bmatrix} a_{11k} & a_{12k} & a_{13k} & \cdots \\ a_{21k} & a_{22k} & a_{23k} & \cdots \\ a_{31k} & a_{32k} & a_{33k} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} = q_k \begin{bmatrix} u_{1k}u_{1k} & u_{1k}u_{2k} & u_{1k}u_{3k} & \cdots \\ u_{2k}u_{1k} & u_{2k}u_{2k} & u_{2k}u_{3k} & \cdots \\ u_{3k}u_{1k} & u_{3k}u_{2k} & u_{3k}u_{3k} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

Notice that there is redundancy in this matrix. Each column contains information that is related to the kth mode shape of the system times the reference DOF. (Also note that due to symmetry, the rows contain the same information.) This very important fact is the reason why many modal parameter estimation algorithms utilize multiple reference data from a modal test. Each of the references contains the same basic information that is only scaled by the reference DOF for a particular mode. Therefore, this redundant information can be extracted and used in the curvefitting process.

More importantly, if there is one reference that does not excite a particular mode very well (ie, the reference is located close to the node of a mode for that mode), then there are other references that may be much better reference locations for the determination of that mode. So using multiple references minimizes the need to be absolutely certain that all of the modes of the system can be reasonably well excited from only one reference location. The modal parameter estimation process uses weighting terms, called modal participation factors, in order to utilize all of the referenced data to extract valid modal parameters. So the use of multiple referenced data is a tremendous help in determining modal parameters. The use of redundant data allows for the selection of several references, each of which may be very good for several modes, but not all the modes, of the system. However, using multiple references allows the adequate description of all the modes from the combination of references. This way, many references gives the best possible chance to adequately determine all of the modes of the system. This may not be totally possible using only one reference - even though theoretically, it is possible!

I hope this explanation helps you to understand why multiple references are useful even though they are not theoretically necessary. If you have any other questions about modal analysis, just ask me.

Copyright 2001
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

What are some of the most important things to consider when impact testing?
Let's discuss this.

This is a very good question. The most important considerations can be broken down into those that are impact related and those that are response related. The excitation concerns are numerous. Only issues pertaining to hammer tip, trigger delay and double impacts are discussed here. However, other issues related to overload/underload of the analog to digital converter, poor utilization of the digitizer, and difficulties with testing nonlinear structures, are some additional concerns (but are not addressed in this article). The response concerns lie with the signal decay and the need for windows to minimize leakage. Let's first discuss the excitation issues and then the response issues.

First of all, the hammer tip is largely responsible for the frequency spectrum that is excited. In general, the harder the tip, the wider the frequency range that is exicted. The hammer tips typically used range from rubber to metal on the extremes with various intermediate tips such as the soft plastic, hard plastic, eraser, etc. Each of these tips are designed to have a certain amount of elastic deformation during impact. The total time duration of the tip impact is directly related to the corresponding frequency range that is excited. Generally, the shorter the length of the time pulse, the wider the frequency range that is excited (Figure 1 shows some typical tips).

While this is generally the case, often the local flexibility of the structure can play an important part of the total time of the impact and therefore can have an effect on the force spectrum imparted to the structure. You may have noticed this when testing structures that have dramatically varying stiffnesses throughout the structure. When impacting a stiff region, one input frequency spectrum is observed and a much different, narrower frequency range is excited when impacting more flexible regions. (The published impact tip frequency spectrum provided by the hammer manufacturer does not include any of the local structure flexibility effects.) Be careful !!!

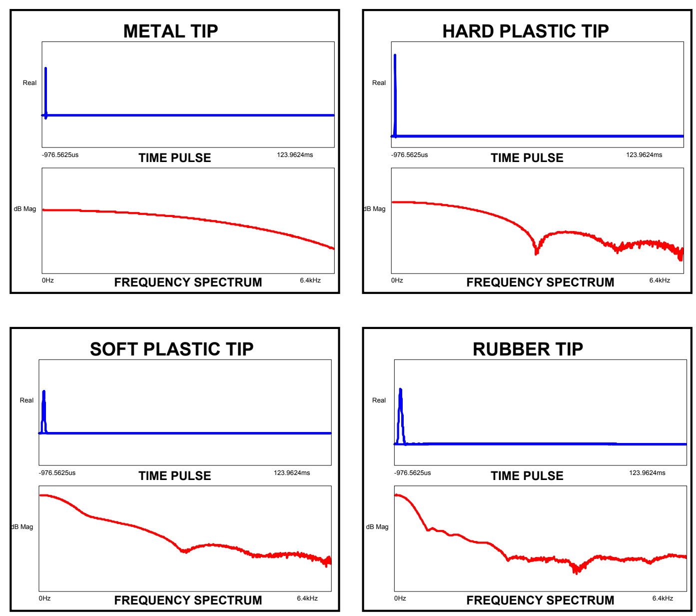

Figure 1

Another somewhat important test consideration for impact testing relates to the need for pretrigger delay. Since the measurement initiation is controlled by the leading slope of the pulse, part of the pulse is lost and the resulting spectrum is distorted unless some pretrigger delay is specified. The effects of spectrum distortion is shown in Figure 2. The red pulse and resulting frequency spectrum are clearly different than the correct blue pulse and resulting frequency spectrum with pretrigger delay. This can cause some input frequency spectrum distortion which will have an effect on the computed FRF.

Copyright 2001
Pete Avitabile

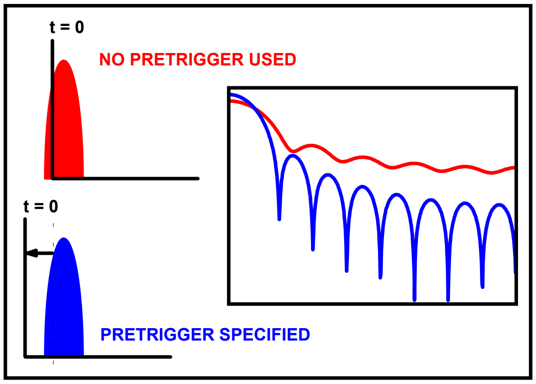

Figure 2

Typically, a pretrigger delay of 1% to 5% of the time window is sufficient to eliminate this effect. Care must be exercised when specifying this delay since some FFT analyzers use a plus (+) delay for pretrigger delay while others use a minus (-) delay for pretrigger delay. This causes a totally incorrect frequency spectrum if not applied correctly. Check your time pulse to assure that the entire pulse is captured in the time signal. In addition, some analyzers use a percentage of the block whereas others use an absolute time value in seconds. If absolute time is used then this can cause problems especially when changing bandwidths during test setup. Be careful !!!

Another annoying impact testing problem is the double impact. The double impact generally causes a non-uniform, non-flat input force spectrum. Two typical double impacts are shown in Figure 3. The "ripple" in the spectrum is not desirable especially if the force spectrum dips substantially. A drop of 30 dB or more is cause for concern especially if it occurs at a resonant peak - and it often does.

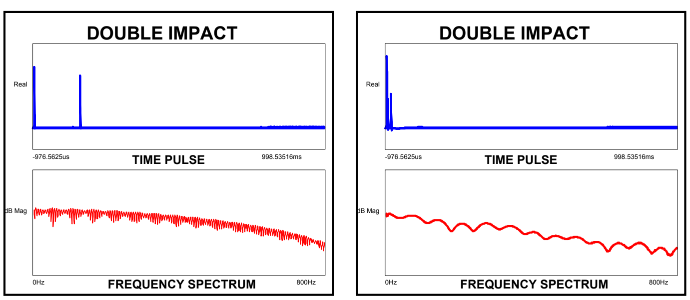

Figure 3

The reason for the double impact is generally from two possibilities. First, many double impacts occur due to new or inexperienced impact testers. It takes some time to get accustomed to swinging the hammer - it is a much different technique than driving nails! But even with experience, sometimes a double impact is unavoidable. Often, with lightly damped structures, the response of the structure is so fast that the hammer can not move away from the structure due to the response of the structure. In these cases, double impacts are unavoidable. The problem is that often the impact spectrum will have significant drop out at the major resonances of the structure. This can produce undesirable effects and must be avoided. One possible technique to overcome the double impact problem is to use the principle of reciprocity. The impact and response locations can be swapped thereby eliminating the double impact problem. This can often solve the problem but many times mass loading effects can become an separate issue.

The last major concern relates to the response and the need for the exponential window. The response of the system may not decay to zero within the sample interval of the FFT. When this is the case, then leakage can occur unless a window is used. The most appropriate window is the exponential window but should only be used when necessary. Many times the window is not necessary if the signal naturally decays within the sample interval.

Often, the data acquisition system can be setup to allow this to happen. Two signal processing parameters should always be explored before using the window. The bandwidth selection can be changed which has a direct effect on the total time required to capture data. If you halve the bandwidth, you double the time sample. Another approach to increase the total time of the sample interval is to change the total number of samples for the acquisition. Both of these two signal processing parameters allow more time data to be collected and should always be explored prior to the use of the window. However, a window may still be required to minimize the effects of leakage if the signal does not die out by the end of the sample interval. Figure 4 shows one time signal in blue that requires an exponential window whereas the red time signal and sample interval does not require a window (or at least substantially reduces the need for the window).

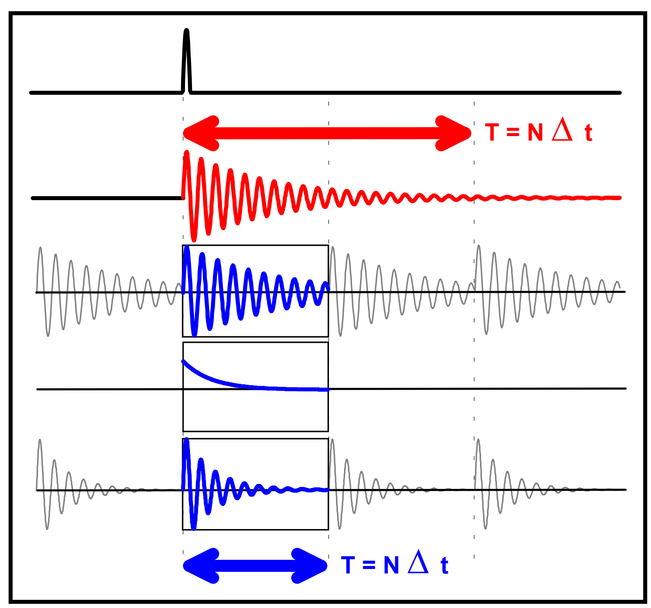

Figure 4

I hope this helps to explain some of the more important concerns when impact testing. If you have any other questions about modal analysis, just ask me.

Copyright 2001
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

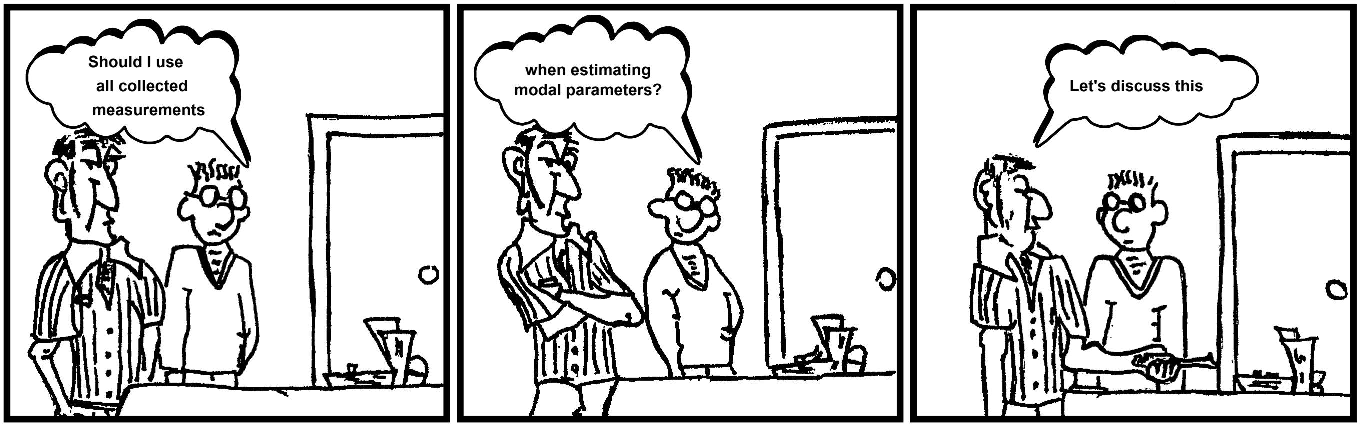

Illustration by Mike Avitabile

Should I use all collected measurements when estimating modal parameters?
Let's discuss this.

This is a very good question. There is no reason to not include all the data collected providing that the data is well measured and consistently related. Providing that there is good dynamic range, with accurate sensitive transducers and all modes are well excited from all reference points and at all the response locations, then, of course, all the data can be used for estimating modal parameters.

But as I said that mouthful of requirements, I could tell from the expression on your face that it is ***highly unlikely*** that ***all*** your measurements meet that requirement. In the past quarter century, I know that I have ***never*** had that happen in any test I have conducted or been associated with - so join the club! What I just described is a measurement situation that will likely only occur with an analytical model with infinite dynamic range and infinite frequency resolution. The real fact is that from a practical standpoint, this will probably never happen. So let's discuss the reality of the situation and discuss some practical approaches to minimize some of the measurement shortcomings.

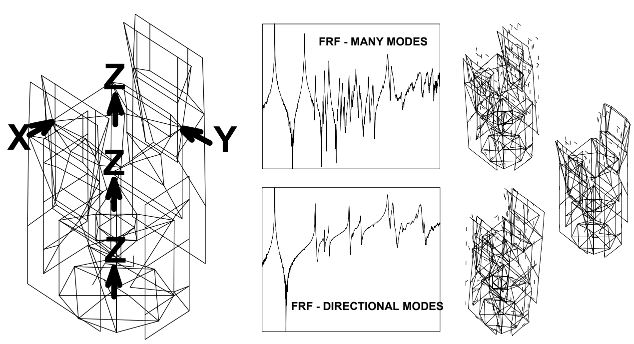

Figure 1 - Aerospace Structure with FRFs and Several Modes

As an example of a common measurement problem, I will use a test that was run many years ago on an aerospace structure that had very directional modes as well as numerous local modes. The structure is shown in Figure 1 along with some typical FRFs. Notice that the lower FRF only shows a few modes but the upper FRF shows all the modes of the structure. Actually, the problem isn't just an aerospace problem but a general problem that can be seen in many structures we test. In fact, the measurements shown are typical of those that could be from almost any structure subjected to modal testing.

The particular structure shown had several bending and torsional lower order modes followed by many local modes with bending, torsion, in-phase, out-of-phase types of modes for the panels and peripheral equipment on the structure. The actual structure was tested using 5 independent shaker excitations (three vertical and two separate horizontal directions).

The first mode of the structure consisted of bending in the x-direction with almost no response in the y-direction. Obviously, the shaker in the x-direction can do a very good job of exciting the x-direction modes but the shaker in the y-direction does not excite the structure in the x-direction very well at all. So the measurements obtained from the y shaker are obviously going to be very poor due to the lack of participation of the first mode in the y-direction.

On the other hand, the second mode of the structure consisted of bending in the y-direction with almost no response in the x-direction. Here the opposite is true from that just discussed. The y shaker can do a very good job of exciting the structure in the y-direction but the shaker in the x-direction cannot excite the structure in the y-direction. But both shakers can do a very good job of exciting the torsional mode from both shaker locations. This directly implies that all of the measurements

Copyright 2001
Pete Avitabile

will not be measured with the same degree of accuracy for each mode.

During the MIMO excitation with 5 shakers, all of the FRFs are collected simultaneously but clearly not all of the modes are excited equally from each of the shaker locations. This is a physical reality of most test structures that is typically impossible to overcome. So how can all of this data be efficiently and accurately processed.

Most modal parameter estimation performed today, generally utilizes a two step process. First, the poles are estimated and then the residues or mode shapes are computed (once global poles have been extracted). With this in mind, the poles of the system do not need to be estimated using ***all*** the measurements collected. The poles can be estimated using only a subset of measured functions that best describe the poles of interest. Once the global poles have been estimated, then the residues or mode shapes can be extracted using all the measurement DOFs. (It is also not necessary to estimated residues for all references, especially if the references do not sufficiently excite all the modes). The selection of particular FRFs for the extraction of poles is schematically shown in Figure 2.

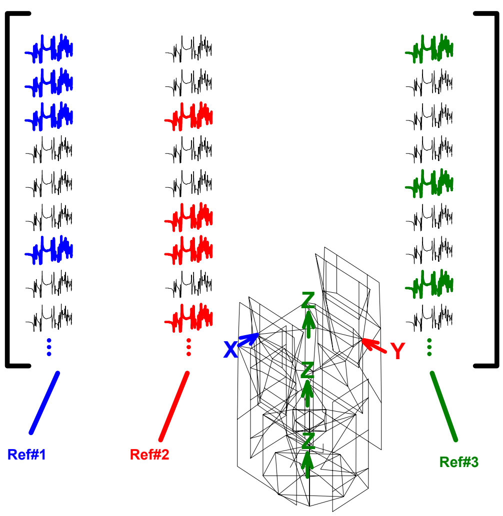

Figure 2 - Schematic Depiction of Measurement Selection

In the example discussed, the first x-bending mode was estimated using only the x-response location from the x-excitation location. Only the y-response locations were used for the y-excitation location for the y bending modes. But both x and y excitations with the x and y responses were used for the torsional mode. Notice that the z-direction excitation and response were ***not*** used for the estimation of any of these poles. This is because the z-excitation locations have a very hard time exciting either the x or y direction modes efficiently. While these references/excitations are necessary for the excitation of some of the higher frequency modes, these vertical excitations are not very good for the excitation of the lower order x and y direction modes. But, of course, once the poles are estimated, then the residues or mode shapes are estimated using all the measurements in the x, y and z directions - but only using the x and y shaker excitations for the x and y lower order modes.

During the modal parameter estimation process, extreme care needs to be exercised to extract the best possible poles to describe the system characteristics. However, many of the measurements and often times all of the references are not optimum for all the modes of the system. As an example, a large telescope structure was recently tested with 4 reference excitation locations. Clearly, the references were not all optimum for all the lower order directional modes of the structure. As a first pass on evaluating the data, all the FRFs from all the reference locations were used to extract poles and residues for the structure. Once parameters were selected, a synthesized FRF was generated and compared to the actual acquired measurement as part of the validation process. The synthesized and measured FRF are shown in Figure 3a. Please carefully note that this is ***not*** a good comparison of the measured and synthesized FRFs. However, after a very careful evaluation of the data and careful selection of measurements to extract the poles of the system (followed by residue extraction), a far better model was obtained. This is confirmed by the comparison of the synthesize and measured FRF shown in Figure 3b. Of course, this approach requires significant effort but the modal parameters are generally greatly improved.

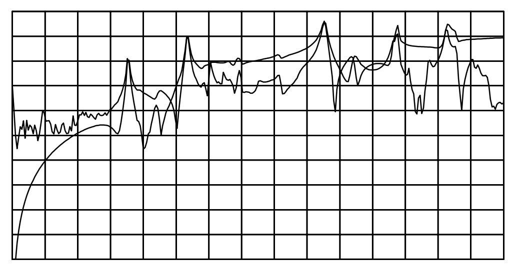

Figure 3a - Poor Extraction and Synthesis of FRFs

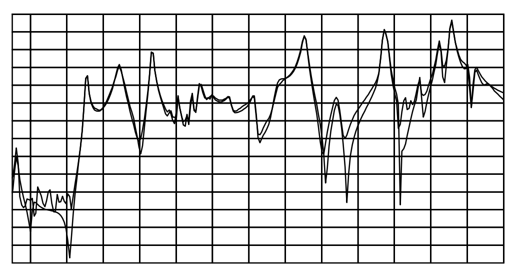

Figure 3b - Good Extraction and Synthesis of FRFs

I hope this explanation helps you to understand why it may not be necessary (or actually detrimental) to use all of the measured FRFs when extracting modal parameters. A careful selection of the best measured FRFs will generally produce much better global poles of the system for the modal parameter estimation process. If you have any other questions about modal analysis, just ask me.

Copyright 2001
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

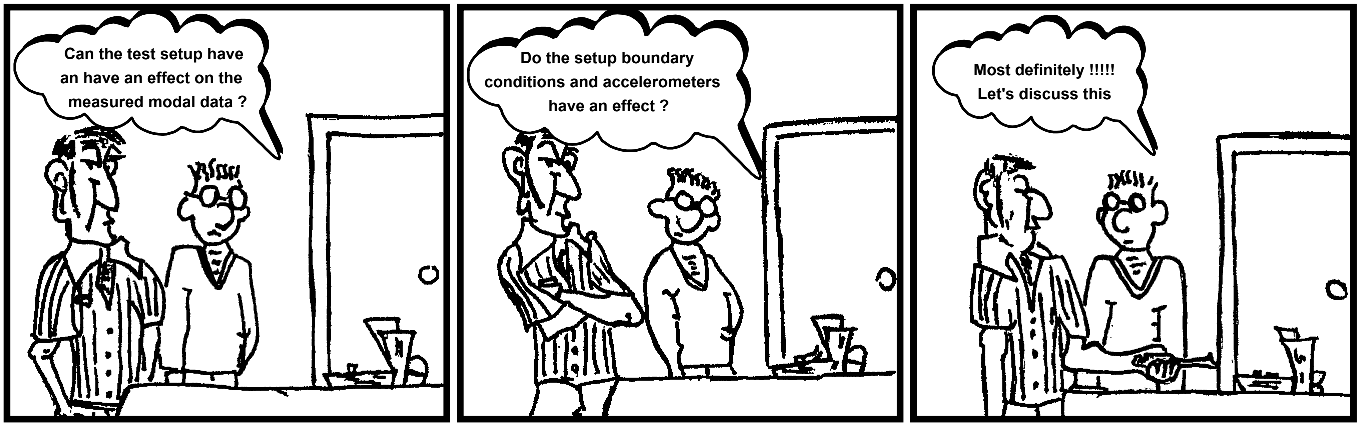

Illustration by Mike Avitabile

Can the test setup have an effect on the measured modal data ?
Do the setup boundary conditions and accelerometers have an effect?
Most definitely !!!!!   Let's discuss this

There is no doubt that the test setup and instrumentation may have an effect on the measured data. This is especially true when testing items such as disk drives, turbine blades, cabinets, computer boards and other small lightweight structures.

While it may be obvious to a seasoned test engineer that the test setup and instrumentation may have an effect on the results of a modal test, this may not necessarily be obvious to the new test engineer to modal testing. (I recently read a report of an experimental modal test on a light weight structure where, after many different tests and analyses were performed, it was "revealed" that the accelerometer mass had an effect on the natural frequency measured on the test article. So it is definitely worthwhile to discuss this further.)

From a practical standpoint, it is straightforward to realize that the instrumentation that is applied to the structure during the test is a direct addition of mass to the structure. However, many times I am shocked that most test engineers new to modal testing just don't realize this fact. For some reason, the instrumentation is perceived as non-intrusive. But, in fact, the instrumentation that is mounted on the structure can, in many instances, have an effect on the measured frequency response functions. From a theoretical standpoint, the natural frequency is related to the square root of the ratio of stiffness to mass. So it stands to reason that if the mass of an accelerometer is "added" to the structure to make a measurement, then the natural frequency will be lowered. Obviously, the larger the accelerometer mass, the more pronounced and obvious the shift of the frequency. And, of course, the size of the test article will have an effect on this. If an accelerometer is added to a large massive structure, such as a bridge or building, then the effects of the accelerometer are likely negligible. But, as the size and mass of the structure under test becomes smaller, then the effect of the accelerometer mass becomes much more important.

It is also very important to note that the mass of the structure is not necessarily the entire mass of the structure but rather the effective mass of the "modally" active portion of the structure. For instance, consider the modal test of a large computer rack with disk drives and computer boards. The mass of the accelerometer on the main structural portions of the rack may not pose any problems. However, the weight of the accelerometer on a cabinet panel or on a computer board or on the armature of the disk drive may have a significant effect on the measured frequencies. Often people get confused by thinking that the mass of the accelerometer is related to the total mass of the structure. This is not the case. It is the mass of the accelerometer relative to the mass of the modally active portion of the structure which may be vastly different than the total mass of the entire structure.

The best way to illustrate the accelerometer mass effect is go down to the lab and take a measurement. To illustrate the mass loading effect, a lightweight disk drive bracket was used for measurement purposes. This was a rectangular structure approximately 5in x 3in x 2in high used for mounting some older disk drives. (Now these measurements are not my pride and joy, but they will clearly illustrate the point.)

A very lightweight, a reasonably lightweight and heavier accelerometer were attached to an open span on the side of the bracket. Three separate impact tests were performed to obtain typical measurements. Only impact excitation and accelerometer response were measured in the x-direction to obtain the drive point measurement shown. Two measurements for the extreme mass cases are shown in Figure 1. The arrows depict two frequencies, for example, that were measured to be 260 Hz (red) and 271 Hz (blue); the third intermediate measurement with a frequency of 266 Hz is not shown. For this frequency, there exists significant difference. So the mass of the accelerometer can have a significant effect; the higher frequencies are effected by an even greater degree.

Another important note is that the two lower amplitude frequencies shown do not appear to be significantly affected by the mass of the accelerometers. These two modes are either y-direction or z-direction predominant in their response. Since the drive point measurement was only obtained in the x-direction, then the mass of the accelerometer is essentially located at the node of the mode for the lower amplitude frequencies and therefore has negligible effect.

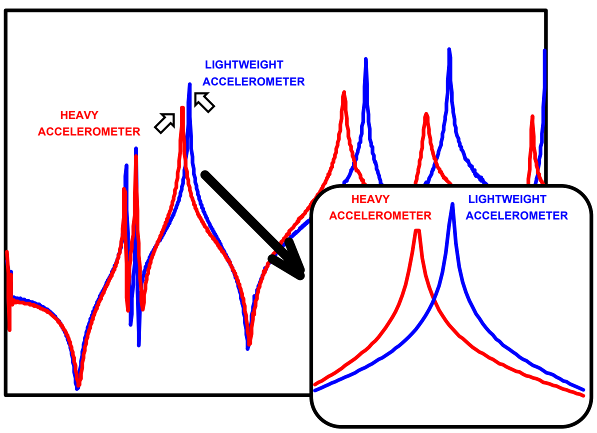

Figure 1 - FRF with Two Different Accelerometer Masses

Therefore, two important items result from this simple test that was performed. First, the mass of the accelerometer has an effect. This must be true since the equation defining the natural frequency of any system involves both the stiffness and mass. Second, the location of the mass of the accelerometer will also have an effect. If the mass is located at the node of a mode (point of zero amplitude), then the mass addition will have no effect *on that mode*. If the mass of the accelerometer is located at the anti-node of a mode (point of maximum amplitude), then it will have the largest effect *on that mode*.

Of course, mass loading effects can pose problems especially if accurate frequency measurements are required. The use of noncontacting (or less intrusive) measuring devices can be used to measure the natural frequencies. For instance, a laser device can be used to obtain high quality FRF measurements without causing any mass loading effects on the structure. However, these devices are typically very expensive and not found in every lab. Other measurements can also be made using eddy current probes or strain gages with reasonably good results. However, these are not always as convenient to apply to the structure under test.

How can the effects of mass loading be identified? Well, the easiest way is to mount two accelerometers at the same location on the structure. Take one FRF measurement with both accelerometers and a second measurement with only one accelerometer. This will quickly identify whether or not the mass loading will be an issue. Then some corrective measures need to be taken if this poses a problem. (Further discussion of this is beyond the scope of this article and will be addressed at some future point in time.)

But there is another item that I feel is just as important as the accelerometer mass loading effect that *is almost always overlooked*. Many modal tests are conducted in a "free-free" condition. Actually, there is no way we can do this here on earth. At best, we can simulate something that is reasonably close to unconstrained (free-free).

Several tests were performed with different mechanisms for supporting the bracket. Three FRFs are shown in Figure 2. The bracket was supported on thick foam (green), airbag packing material (red) and hung from rubberbands (blue). (The same frequency range as used for the mass loading discussion will be addressed here.) The frequencies depicted ranged in frequency from 266 Hz to 272 Hz (and the amplitude is substantially different). This is just about the same amount of frequency variation observed with the accelerometer mass effects!!! So when everyone gets all upset about mass loading effects but don't even consider the support mechanism for the structure, I laugh to myself (and then provide some helpful thoughts to consider). Clearly, the support variation is as critical as the accelerometer mass loading effects. In many cases, the support mechanism effects are much more important than the mass loading effects - *so be careful !!!*

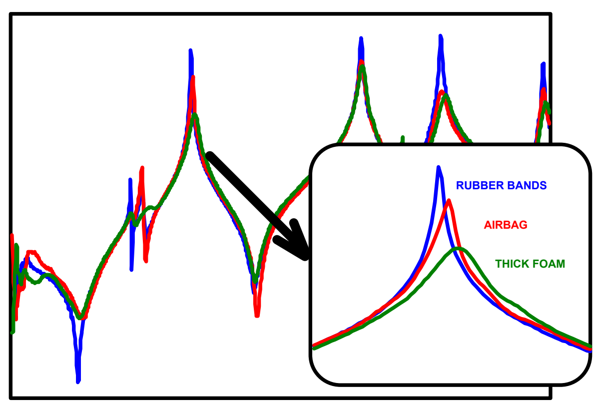

Figure 2 - FRF with Three Different Supports

Now these measurements were just made using what I had available in the lab to quickly show you this problem. It didn't take much effort at all to illustrate the problem of mass loading. But also realize that the support system used to hold the structure in a "free-free" condition is equally important. I hope this helps to answer your question concerning test setup. If you have any other questions about modal analysis, just ask me.

Copyright 2001
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

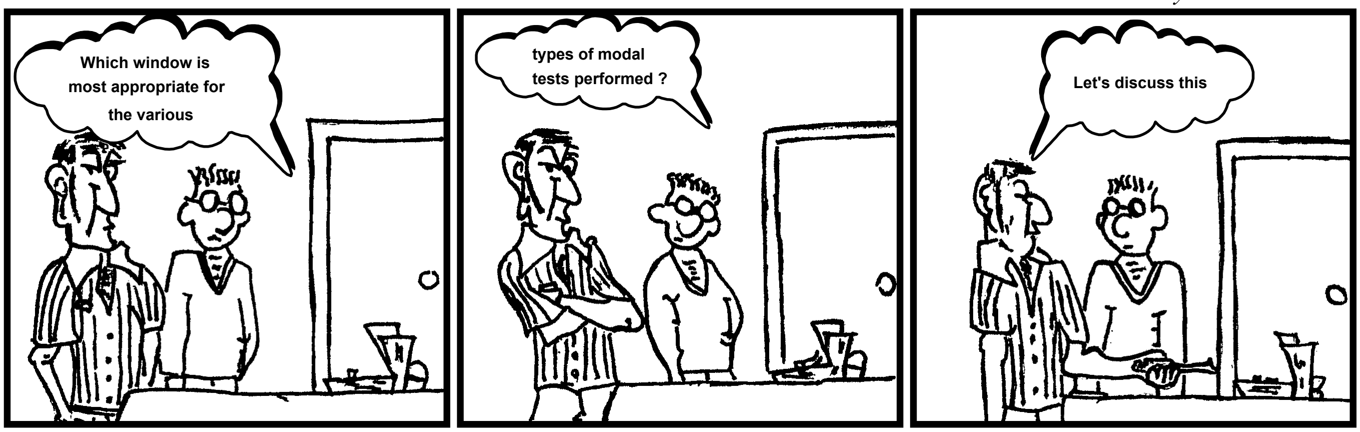

Illustration by Mike Avitabile

Which window is most appropriate for the various types of modal tests performed?
Let's discuss this

This is a good question. Let's review some of the commonly used excitation techniques and associated windows. Actually, the Fourier transformation using the FFT contains some constraints that must be considered first. It is these requirements that will help shed some light on why certain excitations are used and what windows are most appropriate.

First let's remember that the Fourier Transform is defined from $-\infty$ to $+\infty$. As long as the entire transient is measured ***or*** a repetition of the signal is captured, then the requirements of the FFT are satisfied. If this is not true, then there will be serious consequences from the most important signal processing problem called leakage. Windows are weighting functions that are used to *minimize* the effects of leakage - the effects of leakage can never be eliminated. This is really the problem that needs to be addressed. So with this basic fact, let's discuss the different types of excitation signals used for experimental modal testing and explain the windows typically employed for these excitations.

Impact testing is a very common testing technique that is often used for modal testing. Impact testing always causes some type of transient response that is the summation of damped exponential sine waves. This being the case, the entire transient event can be captured such that the FFT requirements can be met and leakage will not be a problem. But for most structures and especially, lightly damped structures, the exponentially decaying response often does not decay sufficiently within the sample record of captured data. This then implies that the FFT requirement may not be satisified. In these cases, an exponential window is typically applied to the data, thereby, weighting the data to better satisify the FFT requirement. Figure 1 shows an impact time pulse along with the raw time response and the exponentially windowed time response.

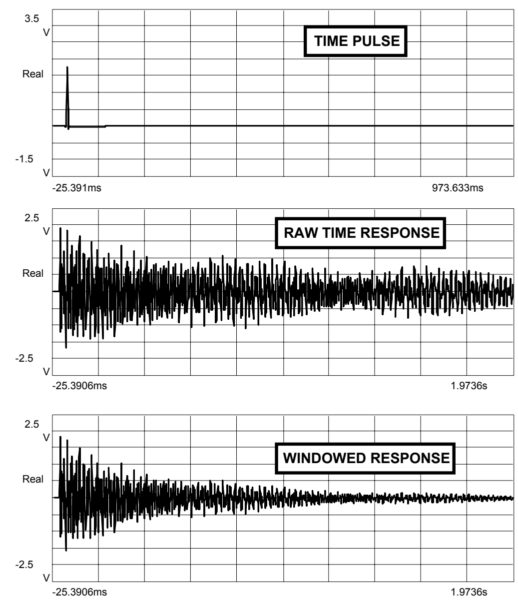

Figure 1 - Impact Excitation and Response

The exponentially windowed response has been weighted such that this signal better satisfies the requirement of the FFT process. The entire signal appears to have been captured - but at the price of the window function. Another alternative to applying the exponential window is to either adjust the bandwidth of the measurement to allow for more captured time data ***or*** increase the the total number of samples which has the direct effect of acquiring more time data. In any event, if the signal does not decay essentially to zero by the end of the sample period, then the exponential window may be necessary in order to minimize the effects of leakage.

In many data acquisition systems, there is also a force window that can be applied for the impact portion of the excitation. This force window is used to eliminate the effects of noise that may

Copyright 2002
Pete Avitabile

be present on the hammer excitation channel. Typically, this window is set to approximately 10% of the sample window such that the impact pulse is located within this unity gain window - that balance of the time record is weighted to zero. The force window may not always be necessary but is available on almost all data acquisition systems. It is very important to note that this window should ***never*** be used to try to remove the effects of a double impact that may occur during impact testing. Use of the impact force window to remove the effects of the second impact, resulting from a double impact, will seriously distort the input force spectrum.

Now that impact excitation has been addressed, let's discuss the window considerations for common shaker excitations used for experimental modal testing. The first most common one is random excitation. The problem with a random excitation is that there will ***never*** be a repetition of the signal within the sample interval. Therefore, a window will be required to minimize the effects of leakage.

The most common window used for random excitations is the Hanning window. But it must be pointed out that the use of a window, any window for that matter, will have an effect on the measured data - but the use of the window is a necessary evil in order to reduce the effects of leakage. Remember that the effects of leakage can only be minimized through the use of a window - it will never be eliminated. All windows will always have the effect of measuring amplitudes that are reduced from the true amplitude and, generally, have the effect of appearing to have more damping then what actually exists. A typical input-output measurement resulting from random excitation is shown in Figure 2 along with the application of the Hanning window on both the measured input and response channels. The use of the Hanning window can cause amplitude distortions of as much as 16%. Of course, this is much better than the distortion due to leakage if no window were applied.

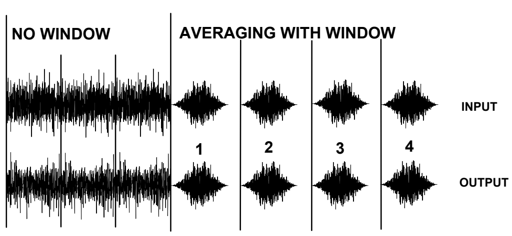

Figure 2 - Typical Random Measurement Sequence

Due to the problems associated with leakage and the measurement distortion through the use of windows, other excitations were developed that were specifically intended to eliminate leakage and the need for windows. Excitations such as pseudo random, periodic random, burst random, sine chirp and digital stepped sine were all developed. (All of these excitations will be discussed in some future article). The most commonly used excitation for modal testing today is burst random and will be discussed here.

The basic problem is that unless the entire transient is captured or a repetition of the data is captured, then there will be leakage. In one way or another all of the specialized excitation techniques attempt to satisify this requirement of the FFT. If this is satisfied, then there is no leakage and therefore, no need to apply any window weighting function. In the case of burst random, the excitation is applied to the structure in a manner whereby the excitation signal starts and stops within the sample interval. This directly implies that the basic requirement of the FFT process is completely satisfied; there is no leakage associated with this signal and no window weighting functions are required. Typically, a burst of 50% to 80% for the sample is customary and can be specified by the user. Now there is no leakage associated with the input excitation signal but some additional consideration must be given to the response channels.

The response of the structure does not stop instantaneously when the shaker excitation is terminated. Generally, there is some exponential decay that is seen to exist on the response channels after the shaker excitation is terminated. (In fact, there is also some measured force that is seen on the excitation channel after the shaker signal is terminated; this is part of the input that must be measured as part of the input forcing function.) A typical input-output burst random signal is shown in Figure 3.

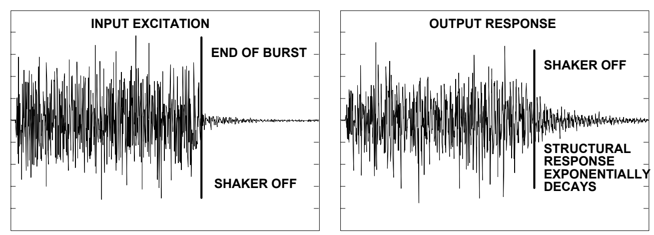

Figure 3 - Typical Burst Random Measurement Sequence

As long as the measured response decays essentialy to zero by the end of the sample period, then the entire signal is captured and there is no need to apply a window weighting function. However, if this is not the case then some adjustments are required. In order to have the entire transient be captured, then either the length of the excitation burst can be reduced, the bandwidth adjusted to provide more time data or more lines of resolution provided which essentially lengthens the captured time sample. All of these will generally help to assure that the entire response of the structure is captured within the sample of data collected. Generally, the use of windows for this excitation technique is not required. In fact the purpose of this excitation technique is to eliminate the use of any weighting functions. This will then provide a leakage free measurement that satisifies the periodicity requirement of the FFT process.

Now, I hope you understand which windows are most appropriate for these most commonly used experimental modal analysis excitation techniques. (Other excitations will be discussed in a future article.) If you have any other questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

What's the difference between a complex mode and a real normal mode?
There's a lot to explain but let's start with some simple examples.

Now that's a question that comes up often and gets many people confused. So let's discuss this in a little detail to explain the differences. Unfortunately, we are going to have to include a little math and some theory here to help explain this.

Let's start with an undamped set of equations and proceed on to a damped case with proportional and then non-proportional damping. It is here where the differences will become apparent. A simple example will be used to illustrate some points here. The equations describing a general system can be written as

$$[M]\{\ddot{x}\}+[C]\{\dot{x}\}+[K]\{x\}=\{F\}$$

where [M], [C], [K] are the mass, damping and stiffness matrices respectively, along with the corresponding acceleration, velocity, displacement and force.

The transformation to modal space will yield

$$\begin{bmatrix}\diagdown & & \\ & \overline{M} & \\ & & \diagdown\end{bmatrix}\{\ddot{p}\}+\begin{bmatrix}\diagdown & & \\ & \overline{C} & \\ & & \diagdown\end{bmatrix}\{\dot{p}\}+\begin{bmatrix}\diagdown & & \\ & \overline{K} & \\ & & \diagdown\end{bmatrix}\{p\}=[U]^T\{F\}$$

with diagonal matrices of modal mass, modal stiffness and, under certain conditions, modal damping. The mode shapes will uncouple the mass and stiffness matrices and for certain specific types of damping, these mode shapes will also uncouple the damping matrix. In order to understand some of these conditions, a simple example will be shown.

For the example here, the matrices will be defined with

$$[M]=\begin{bmatrix}2&0\\0&3\end{bmatrix}\ ;\ [K]=\begin{bmatrix}2&-1\\-1&1\end{bmatrix}$$

$$[C_0]=\begin{bmatrix}0&0\\0&0\end{bmatrix}\ ;\ [C_P]=\begin{bmatrix}0.4&-0.1\\-0.1&0.4\end{bmatrix}\ ;\ [C_N]=\begin{bmatrix}0.4&-0.1\\-0.1&0.1\end{bmatrix}$$

First, the undamped case is considered. The mass, [M], and stiffness, [K], will be used with the $[C_0]$ matrix. The eigensolution of this set of matrices will yield frequencies, residues and shapes as:

$$\lambda_1=0+0.3737j;\quad a_1=\begin{Bmatrix}0+0.1230j\\0+0.2116j\end{Bmatrix};\quad u_1=\begin{Bmatrix}0.2459\\0.4232\end{Bmatrix}=\begin{Bmatrix}1\\1.72\end{Bmatrix}$$

$$\lambda_2=0+1.0926j;\quad a_2=\begin{Bmatrix}0+0.1868j\\0-0.0724j\end{Bmatrix};\quad u_2=\begin{Bmatrix}0.3735\\-0.1447\end{Bmatrix}=\begin{Bmatrix}1\\-0.387\end{Bmatrix}$$

Notice that the mode shape is a sign valued (+ or -) real number. The first mode has both DOFs with the same sign indicating that both of these DOFs are in phase with each other differing only in magnitude. The second mode has both DOF with differing signs indicating that both DOFs are out of phase with each other and have differing magnitudes.

Now let's consider the second case with damping which is proportional to either the mass and/or stiffness of the system. The damping here is $[C_P]$ to be used with the [M] and [K]. The eigensolution of this set of matrices will yield frequencies, residues and shapes as:

$$\lambda_1=-0.0579+0.3693j\quad a_1=\begin{Bmatrix}0+0.1244j\\0+0.2141j\end{Bmatrix}\quad u_1=\begin{Bmatrix}0.2488\\0.4282\end{Bmatrix}=\begin{Bmatrix}1\\1.72\end{Bmatrix}$$

$$\lambda_2=-0.1097+1.0871j\quad a_2=\begin{Bmatrix}0+0.1877j\\0-0.0727j\end{Bmatrix}\quad u_2=\begin{Bmatrix}0.3754\\-0.1455\end{Bmatrix}=\begin{Bmatrix}1\\-0.387\end{Bmatrix}$$

Copyright 2002
Pete Avitabile

Notice that the eigensolution yields the same mode shapes as the undamped case. This is due to the fact that the damping is proportional to the mass and /or stiffness of the system. This results in modes that are referred to as "real normal modes". So it is clear that the mode shapes for the undamped and proportionally damped cases are exactly the same.

Now let's consider the third case with damping which is not proportional to either the mass and/or stiffness of the system. The damping here is $[C_N]$ to be used with the [M] and [K]. The eigensolution of this set of matrices will yield frequencies, residues and shapes as:

$$\lambda_1 = -0.0162 + 0.3736j \quad a_1 = \begin{Bmatrix} -0.0071 + 0.1288j \\ -0.0048 + 0.2116j \end{Bmatrix} \quad u_1 = \begin{Bmatrix} 0.2456 + 0.0143j \\ 0.4232 + 0.0095j \end{Bmatrix}$$

$$\lambda_2 = -0.1005 + 1.0872j \quad a_2 = \begin{Bmatrix} -0.0071 + 0.1885j \\ -0.0048 - 0.0726j \end{Bmatrix} \quad u_2 = \begin{Bmatrix} -0.3771 + 0.0142j \\ 0.1451 + 0.0096j \end{Bmatrix}$$

Now for this case, the mode shapes are seen to be different than the previous cases. First of all, the mode shapes are complex valued. Upon closer inspection of these shapes, it can be seen that the relative phasing between each DOF for each of the modes is NOT either totally in-phase or out-of-phase. This results in modes which are described as "complex modes". This is very different from the two previous cases. This will typically occur when the damping for the system is not related to the mass and/or stiffness of the system and is referred to as non-proportional damping. In order to perform the eigen solution, a slightly different form is used where the equations are cast in state space in order to perform the solution.

Basically all the equations get more complicated when considering complex modes. Some simple statements between a real normal mode and a complex mode can be summarized as follows:

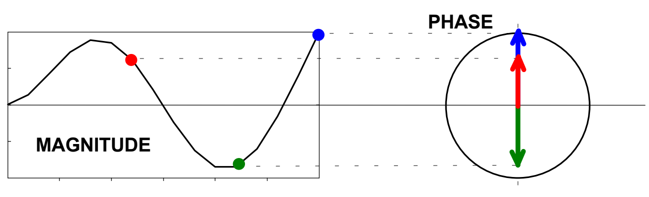

Figure 1a - Proportional (Real Normal) Mode Schematic

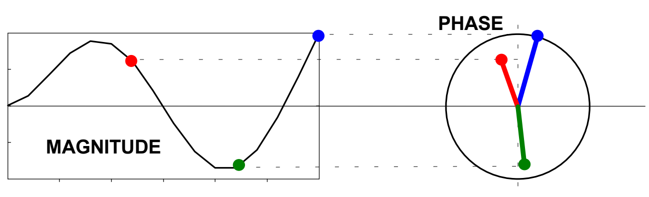

Figure 1b - Non-Proportional (Complex) Mode Schematic

REAL NORMAL MODE

Some characteristics of a real normal mode are:

1. The mode shape is described by a standing wave which has the presence of a fixed stationary node point
2. All points pass through their maxima and minima at the same instant in time
3. All points pass through zero at the same instant in time
4. The mode shape can be described as a sign valued, real number
5. All points are either totally in-phase or out-of-phase with any other point on the structure
6. The mode shapes from the undamped case are the same as the proportionally damped case. These shapes uncouple the [M],[C], and [K]

COMPLEX MODE

Some characteristics of a complex mode are:

1. The mode shape is described by a traveling wave and appears to have a moving node point on the structure
2. All points do not pass through their maxima at the same instant in time - points appear to lag behind other points
3. All points do not pass through zero at the same instant in time
4. The mode shape can not be described by real valued numbers - the shapes are complex valued
5. The different DOFs will have some general phase relationship that will not necessarily be in-phase or 180 degrees out-of-phase with other DOF
6. The mode shapes from the undamped case will not uncouple the damping matrix

In order to further visualize some of these statements. A simple mode shape is plotted for a real normal mode and a complex mode for one of the modes of a cantilever beam. In the real normal mode (Figure 1a), the relative phasing between the DOF is either totally in phase (as in the case of the blue and red DOF) or totally out of phase by 180 degrees (as in the case of the green DOF relative to the blue and red DOF). A complex mode does not have this simple phasing relationship and the mode shape must be described by both amplitude and phase, or real and imaginary components (Figure 1b). The plots in Figure 1 are intended to visualize this relationship of the phase.

Now it is very important to point out that phasing can be seen in FRF measurements all the time. Sometimes this may be an indication of complex mode behavior, but be careful to jump conclusions. The data acquisition, instrumentation, signal processing, FFT, and modal parameter estimation are all stages that can distort a measurement and force a mode shape to "appear" as if it is complex.

While there is a lot more to it all, I hope this simple explanation helps to put everything in better perspective. Think about it and if you have any more questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

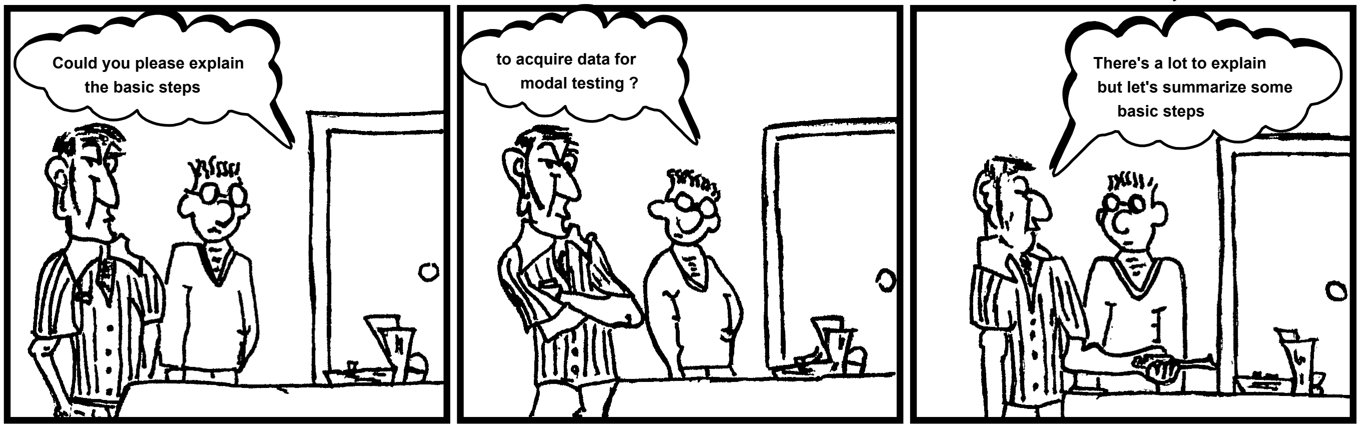

Illustration by Mike Avitabile

Could you please explain the basic steps to acquire data for modal testing?
There's a lot to explain but let's summarize some basic steps.

The basic data acquisition process has several steps. These steps for impact and shaker testing are slightly different and will be discussed separately.

For impact testing, the excitation is applied to the structure using a hammer or some type of impacting device where the force transmitted to the structure is measured. The response of the structure is often measured using an accelerometer but sometimes a laser or other measuring transducer is used. Typically, the force used to excite the structure is measured on the lowest channel of the data acquisition system. Although this is not a requirement on many data acquisition systems today, many test engineers still follow this practice. The response signal(s) is measured on the remaining channel(s) (depending on whether a dual or multichannel system is being used). In order to start the data acquisition measurement, typically the measurement is started from a trigger from the impact device. Some minimal voltage must be specified in order for the data acquisition system to start the measurement process. A trigger level of 10% to 20% of the maximum voltage of the measured force is a good value to use for most tests performed. In many data acquisition systems, a pretrigger delay is specified to capture the entire transient of the impact device. By using the pretrigger delay feature, none of the impact force pulse is omitted.

The transducer data is collected and this data is always passed through a low pass, analog filter before any digitization is performed. This is done mainly to filter out high frequencies that are not of interest and to prevent aliasing from occurring. These analog filters, often referred to as anti-aliasing filters, remove high frequencies which might otherwise contaminate the measured frequencies.

This data is then passed into the analog to digital converter (ADC) where the data is sampled and converted into digital form. Two concerns exist at this point. The data must be sampled at a rate so as to adequately characterize the time data for conversion to the frequency domain. Generally, the data must be sampled at least twice as fast as the highest frequency of interest for transformation to the frequency representation. If time data processing is needed to evaluate time characteristics of the structure, then sampling should be performed at least 10 to 20 times faster than the maximum frequency of interest in order to adequately interpret the system characteristics. In order to properly characterize the amplitude of the signal correctly, the ADC must be set to an appropriate voltage level to characterize the signal. If this is not done properly, then quantization errors in the measured signal may pose a problem. In many data acquisition systems, a feature referred to as autoranging assists in setting appropriate voltage levels for all of the data acquisition channels. The ADC levels can also be set manually but care must be exercised to assure that the acquisition channels are set properly. Otherwise, the signal may suffer from quantization error if the levels are set to high or from clipping and overloads is the level is set too low.

At this point, the digital data describing the impact and response is available in raw digitized form. Depending on the character of the actual time signals, windows may need to be applied in order to minimize any leakage that may otherwise result. Leakage will occur during impact testing if the entire transient is not captured during the acquired sample of data. If there is significant noise on the impact channel, then a force window may be used to minimize this if necessary. If the response signal does not decay sufficiently to zero by the end of the sample interval, then an exponential window may be necessary to avoid distortion of the signal due to the Fourier transformation process.

Copyright 2002
Pete Avitabile

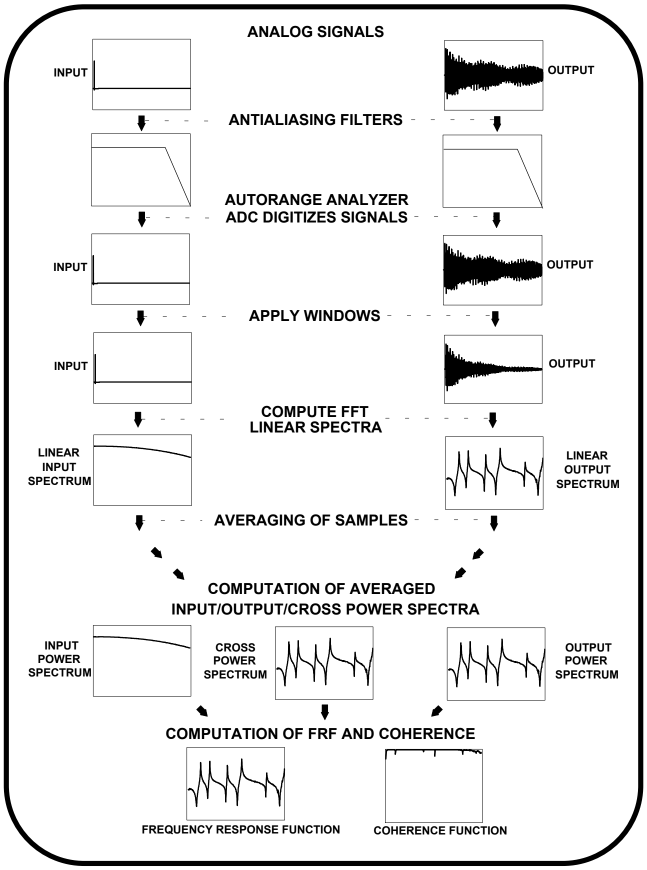

Figure 1 - Impact Test Flow Diagram

Before an exponential window is applied, two additional signal processing features should be used. In order to allow more time for the response signal to naturally decay to zero, the bandwidth of the measurement can be reduced or the number of spectral lines increased; both of these parameters will ultimately lengthen the total time required for the measurement. This will allow more time for the response signal to naturally decay thereby minimizing the need for application of a significant damping window.

For shaker testing, the excitation of the shaker is typically measured at the lowest channel of the acquisition system (again while this may not be required, many follow this practice). The response transducers are measured in the remaining channels of the data acquisition system. Depending on the excitation used, triggering will vary. For random excitations, a "free run" mode is typically used. However, other excitations (such as burst random, sine chirp, etc.) will start from a signal source or force trigger. In addition, sometimes a pretrigger delay is specified for burst random excitation.

For many shaker excitations used, no window is applied since these signals usually have "special" characteristics that are employed in order to provide leakage free measurements that satisfy the FFT requirements. However, if any arbitrary signal such as random excitation is employed, then a window such as a Hanning window, is used to minimize the leakage effects.

For either impact or shaker testing, the time captured data must be transformed to the frequency domain using the FFT and the transform algorithm. The FFT provides the linear Fourier spectrum of the input and output(s) signals. (Note that these functions are complex valued functions.) This then provides the input spectrum (Gxx), output spectrum (Gyy) and cross spectrum (Gyx). These three spectra are then averaged using all the individual data records collected. Once the Gxx, Gyx and Gyy are obtained then the frequency response function and coherence are computed. While different forms of the frequency response function are available, H_1 is the most popular form of the frequency response function employed in the majority of single input modal testing performed today. Figures 1 and 2 depict the measurement process for impact and shaker testing, respectively.

While frequency response functions are the only measurements required for development of an experimental modal model, many times the auto- and cross-spectra along with the coherence are saved as part of the dataset. (With the abundant availability of disk drive storage, there is no reason to not save all of the measurements!)

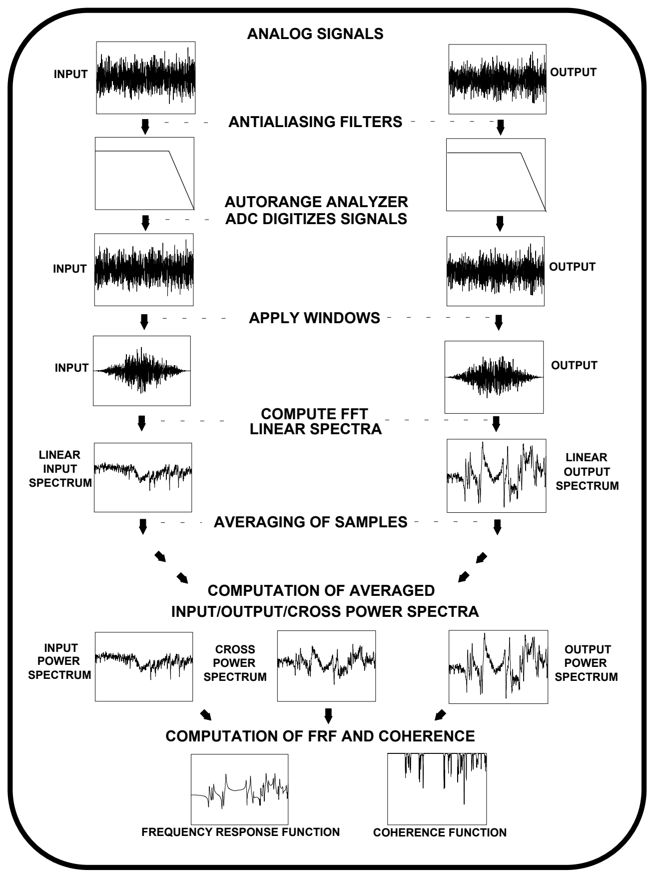

Figure 2 - Random Shaker Test Flow Diagram

Obviously, there is much more that could be discussed but I hope this helps to explain some of the basic steps in the overall measurement process for experimental modal testing. If you have any more questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Is there any real advantage to MIMO testing?
Why not just use SISO and then move the shaker?
Let's talk about the differences.

Multiple input multiple output (MIMO) testing has many advantages when compared to data collected from single input single output (SISO) testing. The energy from multiple shakers allows the structure to be more uniformly excited throughout the entire structure and thus allows for the development of better frequency response functions (FRF). When only using a single shaker, the measurements obtained are generally not as good as those obtained from multiple shaker excitations, especially when considering larger structures. With single shaker methods, many times it is difficult to get a reasonably good level of excitation throughout the entire structure.

Another important factor is the effect of the shaker setup on the test article to be measured. With single shaker testing methods, the shaker system must be setup multiple times in order to obtain multiple reference data necessary for polyreference curvefitting techniques. Many times the test setup may have an effect on the measured FRFs. When multiple reference data is collected with a single shaker, this can be a serious concern. The data that is collected with MIMO testing is generally more consistent when compared to equivalent data collected SISO testing.

In order to see some of the differences, some data was collected on a structure using both single shaker and multiple shaker excitation techniques. Four different sets of data were collected:

- SISO with random excitation and a Hanning window
- SISO with burst random excitation and no window
- MIMO with random excitation and a Hanning window
- MIMO with burst random excitation and no window

In all cases, a reciprocal FRF measurement between the two excitation locations was collected. The accelerometers were permanently mounted to the structure for all testing performed to minimize any mass loading effects that might otherwise occur. Only the shakers were connected and disconnected between the various measurements obtained. (The shakers were actually placed at the measurement locations and then connected or disconnected as necessary to minimize the effects of shaker setup problems.) The FRF measurements for all four cases are shown in Figure 1. Each graph contains two reciprocal FRFs - Hij and Hji.

On first glance, it appears that all four techniques provide similar data. The FRFs appear reasonably good. However, upon closer inspection of each of the techniques, differences will be easily seen in the reciprocal measurements (Figure 2).

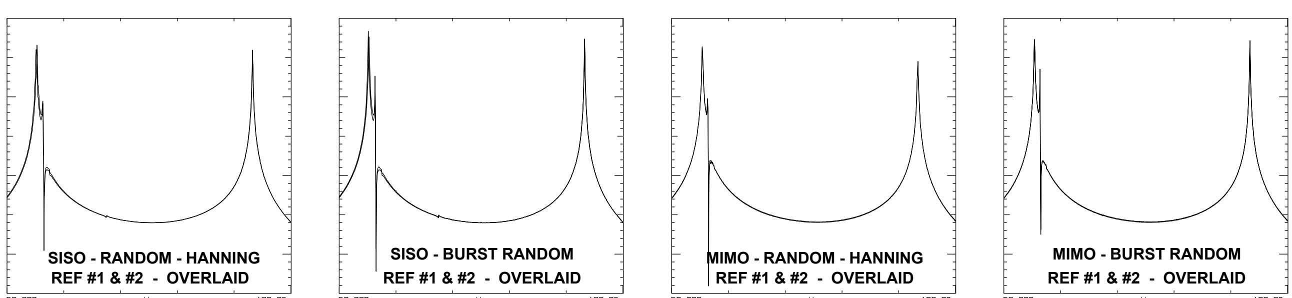

Figure 1 - Reciprocal FRFs for All Data Sets Compared over the 50 to 100 Hz range

Copyright 2002
Pete Avitabile

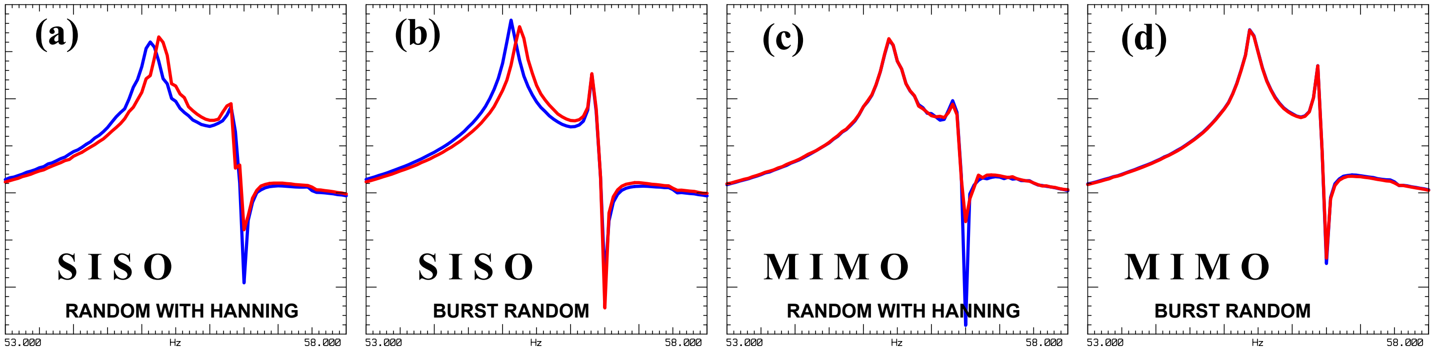

Figure 2 - Reciprocal FRFs for All Data Sets Compared over the 53 to 58 Hz range

First, look at the FRF measurement using SISO with random excitation shown in Figure 2(a). Note that the two curves do not line up very well especially at the first resonant peak shown. Also notice that the random excitation exhibits significant variance on the FRF measurement even with the Hanning window applied; this is due to the leakage effects of the random excitation process which are not completely removed through the use of the Hanning window. The shift in the peak of the FRF is directly attributed to the shaker setup and the stiffness effects of the stinger attaching the shaker to the structure. The test setup clearly has an effect on the measured FRF data.

Next, look at the FRF measurement using SISO with burst random excitation shown in Figure 2(b). Notice that the two peaks still do not line up very well at the first resonant peak. However, the burst random excitation provides a much better measurement when compared to the random excitation with the Hanning window applied. The burst random excitation generally provides a much better overall measurement since no window is necessary for the collected data. This is due to the fact that the measurement satisfies the periodicity requirements of the FFT process. The random excitation will generally need many more averages in order to reduce the variance on the measured FRF data and will still not produce as nice a measurement as the burst random excitation.

Now, look at the FRF measurement using MIMO with random excitation shown in Figure 2(c). The peaks at the first resonance compare much better than in the previous case. However, there is still significant variance of the FRF in general and significant differences exist mainly at the antiresonance. Since both shakers are mounted on the structure at the same time, the effects of the mounting of the shaker are similar for the duration of the test. Therefore, the resonant peak is not affected by the setup configuration.

Finally, look at the FRF measurement with the MIMO with burst random excitation shown in Figure 2(d). This reciprocal measurement is almost identically the same. Obviously, this measurement is the best of all the measurements considered. The burst random MIMO excitation has the best overall characteristics - the measurement does not need a window to be applied since there is no leakage and has the consistency necessary in the reciprocal FRF since the shakers are mounted simultaneously on the structure for the duration of the test.

So from the SISO and MIMO data evaluated, it is clear that the MIMO data produces more consistently related data. The SISO measurements clearly showed differences when evaluating the resonant peaks. In addition, the burst random excitation produced much better results when compared to the random excitation with the Hanning window applied.

Another interesting point is shown in Figure 3. Notice that the two SISO and MIMO all produce different results. This clearly shows that the test setup has an effect on the measured FRFs. Of course, we realize that the peaks of the FRF may be affected by the stinger stiffness attaching the shaker to the structure, but at least with the MIMO configuration, the peaks are consistently related. (This can be a very critical point especially when performing multiple reference modal parameter estimation techniques to extract the modes of the system.)

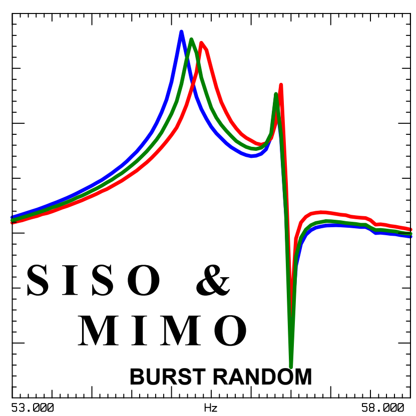

Figure 3 - Comparison of SISO & MIMO FRFs

I hope this example helps to show the advantages of MIMO testing techniques. If you have any more questions on modal analysis, just ask me.

Copyright 2002
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

What are some of the major mistakes people make in modal tests so I don't make the same ones.
Let's talk about some of them.

Well, there are many things I have seen over many years that are big mistakes that people have made during modal testing and analysis. Some of them are actually pretty funny and some of them are fairly serious. So let me discuss some of what I refer to as the "modal testing funnies" that I have observed over the past several decades of modal testing.

Impact testing and coherence - Many, many years ago (as well as just a few months ago) I observed a modal test where the engineers would only take one average for each measurement location. When asked why, they very clearly stated that since the coherence was one, why bother making any more measurements - how much better could they get ! I have seen many fall into this trap. The coherence is a function that can only be evaluated for averaged measurements. With only one average, there is no variation relative to the only measurement made. So, therefore the coherence can only be one - but is definitely not an indication that the measurement is acceptable. You can only use the coherence function to evaluate the variation on a set of averaged functions. Averaging is required!

Reference location at the node of a mode - In order to conduct a modal survey, the reference location (stationary response in a roving impact test and stationary input location for impact or shaker testing) must not be located at the node of a mode. The node point is a location of zero response. If there is no response, then how can a good FRF be obtained. I remember many years ago, a group was testing a large cantilevered type of structure and they had just procured some very expensive low frequency accelerometers. Concerned that the device might fall from the structure, the reference accelerometer was mounted at the base of the cantilevered structure - but of course there is no response to measure so the FRFs were very poor to say the least. Another test of a rib-panel cabinet structure was instrumented with all the measurement locations at the majority of the stiff rib intersections and no accelerometers on any of the panels portions of the structure. For this particular structure, it turned out that the modes of interest were mainly related to the panel responses. All the accelerometer locations were essentially at nodes of all the modes. Unfortunately, someone worked on this over several months and acquired many sets of useless data before asking for guidance and help. The measurements need to be at locations where there is response to measure. Avoid nodes of modes!

Mass loading effects - Now mass loading effects have been discussed in other articles. The mass effect of instrumentation can cause an effect on the measured FRFs and give a misrepresentation of the system natural frequencies. Care must be exercised to determine what mass loading effects, if any, exist. This may be a pronounced effect or relatively insignificant. But I have a mass loading story that is hard to beat. Back in the early 80s, I was involved in impedance testing of isolation components for submarine propulsion systems. New designs were being evaluated and impedance testing was being performed over a 5 KHz range with a new approach at the time called digital stepped sine; each test took approximately 3 hours to conduct. With many measurements for each configuration, many proposed prototypes to evaluate and only a dual channel analyzer, testing was performed 24-7 to complete the tight scheduled program. The configuration involved a set of 2 ton concrete blocks supported on air bags and separated by the isolation system. There was one particular co-op student (about 6'5" and about 350lb) who would run some of the third shift tests. Three months after the program ended, the analysts evaluating data called with some questions regarding several data sets that appeared to have a 400 lb difference in mass that appeared part way through the frequency range. Their comment sent a chill up my spine as I remembered coming into the shop assembly area where the tests were being conducted every morning and remembering seeing a pillow and blanket on the concrete blocks. After some lengthy discussion, it was determined that once the test was started, the co-op student would climb up onto the airbag supported concrete block to take a nap while the test was underway !!! You can imagine how hard it was to explain that unique testing problem. It turned out that since a very good logbook was kept for all the tests, it was easy to determine which tests the co-op was involved in and additional measurements were reviewed and some determined to be inappropriate for use. No one likes to keep an accurate logbook but this is one case where that information was invaluable. Now that is a mass loading effect that I don't think will be easy for anyone to top. Watch out for mass loading!

Double impact and force windows - Many times double impacts are unavoidable while performing an impact test. There have been several times where I have seen test engineers use the force window to eliminate the effects of the second impact of the double impact. Their thought process was that since I really didn't want to have a double impact, why not just zero out the effects of the second unwanted impact. At first thought, this might seem reasonable, but the reality is that the structure actually did see the double impact. Therefore, it is totally incorrect to apply the force window to try to analytically remove the effects of the second impact. The structure actually responds due to the actual double impact applied. Never try to remove the effects of a double impact through the use of the force window.

Averaging using a different point for each average - This is one you have to laugh at in complete disbelief. I witnessed a test eons ago, where the resulting mode shapes of a fairly simple structure appeared to be nothing more than shear gibberish. After checking some of the basic things that could be wrong, a lengthy discussion was held to determine why the coherence of the measurement could be so bad for a simple structure when 25 averages had been acquired. It turns out that there were 25 measurement points to be acquired. The test engineer took 25 averages for each measurement point for the modal test. Unfortunately, each of the individual measurements that made up the summation of all 25 measurements were obtained through an impact test where each individual point on the structure was impacted at a different location for the total 25 measurements. Basically, each averaged FRF came from the 25 averages from impacting ALL of the individual 25 points on the structure. This one averaged function was then labeled for one of the test structure measurement points. This process was repeated 25 times until all the measurement points were measured. Obviously, this is completely incorrect. An averaged FRF for a given point on the structure must come from a measured function where the SAME point is impacted for each of the averages. I still can't believe that one.

Coordinate systems and point/direction information - The proper identification of point and direction identification for a modal test is a fairly simple process - but sometimes errors result from incorrect specification of this information. A simple remedy is to clearly mark the coordinate system at the test setup. I usually put down tape on the floor with the x, y and z directions clearly labeled. Many people laugh behind my back at this routine but I never mess up point direction information for a modal test and many others often do. So I guess it is "he who laughs last, laughs best" and I am still laughing at some of the mishaps that I have seen. I remember one modal test of an engine block on its mount system. The test was used to basically identify the rigid body modes of the engine on the mount system - these are typically very low frequency modes. One day I received a phone call with a problem where the modal test revealed a flexible mode of the engine block in the 10 Hz frequency range. This was highly unlikely and I questioned the proper identification of the point and direction information for all of the measurement locations. The people at the test lab very abruptly stated that they were experienced modal test engineers and knew exactly how to identify this basic information. Not wanting to ruffle their feathers, I asked for the data to review. Upon close inspection of their data, it became very clear and blatantly obvious that one face of the engine block had all of the X direction transducers mislabeled 180 out of phase (basically they were pointing in the opposite direction as that specified in the modal software package). Once the phase was corrected for the measurement points of concern, the engine block rigid body modes appeared exactly as expected. Point and direction information is a fairly simple straightforward process - care needs to be exercised in this important step of the geometry generation.

Finite element models aren't always correct - Now that's a loaded statement. Many people have heard the statement that "Everyone believes the test except the test engineer and no one believes the model except the analyst." A large satellite structure was tested about ten years ago where a great deal of care was employed to accurately identify the shaker locations for a modal test with a very elaborate pre-test analysis using the finite element model. The model identified several shakers along the length of this long cantilevered structure attached to a huge seismic mass - but the two horizontal shakers were only set up in one of the directions perpendicular to the length of the cantilevered satellite structure. When questioned on the omission of exciters in the other perpendicular direction, the analysts firmly responded that there was no need for the exciters in the other horizontal direction since an extensive pre-test analysis was performed and ALL the modes of the structure would clearly be excited by the selected excitation directions. Well, it turned out that the model was not perfectly correct (and actually had many lumped mass elements incorrectly defined and located in the model). Therefore, the pre-test analysis was biased by the errors in the finite element model and therefore provided inaccurate information. The selection of reference locations is not an easy task. A finite element model, if available, is a great tool to assist in the selection of references. But care should be taken to not put too much faith and confidence in a model that has not been verified (which was actually the point of the test under way).

Bottom line - The bottom line for almost any modal test is that you need to carefully think about each step of the measurement process to assure that correct FRFs are obtained . The worst situation occurs when people stop thinking about what they are trying to do. I get upset when I walk in to a lab and see measurements being made with no understanding of what is being measured. Many times the response of the concerned people is that "this is the way we have always done it and it must be right because we have been doing it this way for years". It is all right to take measurements following a set procedure but it is imperative that everyone understands the logic and reasoning behind the approach and methodology used for the acquisition of FRF data. The first thing that should always be done is ***question assumptions*** to assure that everyone knows why things are done a certain way. The next thing I always say is that ***thinking is not optional!*** This is not push button technology like a hamburger joint where the choices are simple - burger, fries, soda, $4.52 please. Modal testing has come a long way in the past 25 years but we are still not to the fully automated modal test just yet.

I hope some of these stories have brought a smile to your face but just make sure you don't make the same mistakes! If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Is it better to collect averaged FRF data for a modal test?
Or collect time data and process it afterwards?
Let's talk about the differences.

Both approaches are acceptable providing that good data is collected but time data offers many more advantages. Let's discuss some of the different aspects of collected averaged data vs. time data streamed to disk.

In the old days, the option of streaming to disk was not possible. Typically, time data was collected and immediately averaged to obtain FRF data. Generally, computer memory and disk drive capacities were very small and this necessitated the immediate processing of time data. (Actually, in the very early days of modal testing, it was very rare that all the spectra would be saved and many times only the FRF was saved - you had to think twice about anything that was saved due to the expense of storage devices.)

This was the typical mode of operation for most modal tests performed. As the data was collected, the averaged input, output and cross spectrum was available for review along with the FRF and coherence. With this approach, there was immediate information available to assure that adequate measurements were obtained. The measurements could be scrutinized after each set of measurements were collected and if necessary, additional averages could be collected to obtain improved measurements or determine what might be causing poor measurements. As each set of measurements were collected, this data review continued for each set of measurements obtained. If any problems occurred during any of the measurements acquired, there was immediate feedback through review of the FRF and coherence as to the adequacy of each measurement obtained.

Depending on the application, at times, data was collected and recorded on magnetic tape in the field at the test site. This data was then brought back to the laboratory for processing to obtained averaged FRF data.

However, the use of magnetic tape and the associated tape recording equipment, at times, introduced a wide array of different issues that could possibly contaminate some, if not all, of the data collected. While this introduced problems of its own, the advantage of having time data enabled further processing following the completion of the test. Sampling parameters could be studied to determine various signal processing effects since time data was available. This enables the test engineer to gain further insight into various aspects of the data collected. If only averaged FRF data is available, none of the additional processing is possible.

Today, it is very common to obtain time data that is directly streamed to disk. (This largely due to the availability of inexpensive large capacity disk drives.) This data is collected and then processed after the completion of the test. There is no doubt that time data is by far the best data to collect today. With time data, the same processing still needs to be performed (as is done with averaged FRF data) on the data (Figure 1). However, the time data is always available for additional processing if needed. Additional signal processing scenarios can be investigated if desired or needed. This is not possible with averaged FRF data; once the data is collected there is very little additional processing that can be performed since time data is not available once the data is processed.

Based on all the statements above, it appears fairly obvious that time data is the best data to collect. There doesn't seem to be any reason to collect anything but time data followed by whatever spectral processing is needed. In this way, any subsequent processing can be investigated and explored with the time streamed data. Once time data is available, new concepts and processing can be performed at a future data as the technology progresses. If averaged data were collected, then future processing could not be explored.

Copyright 2002
Pete Avitabile

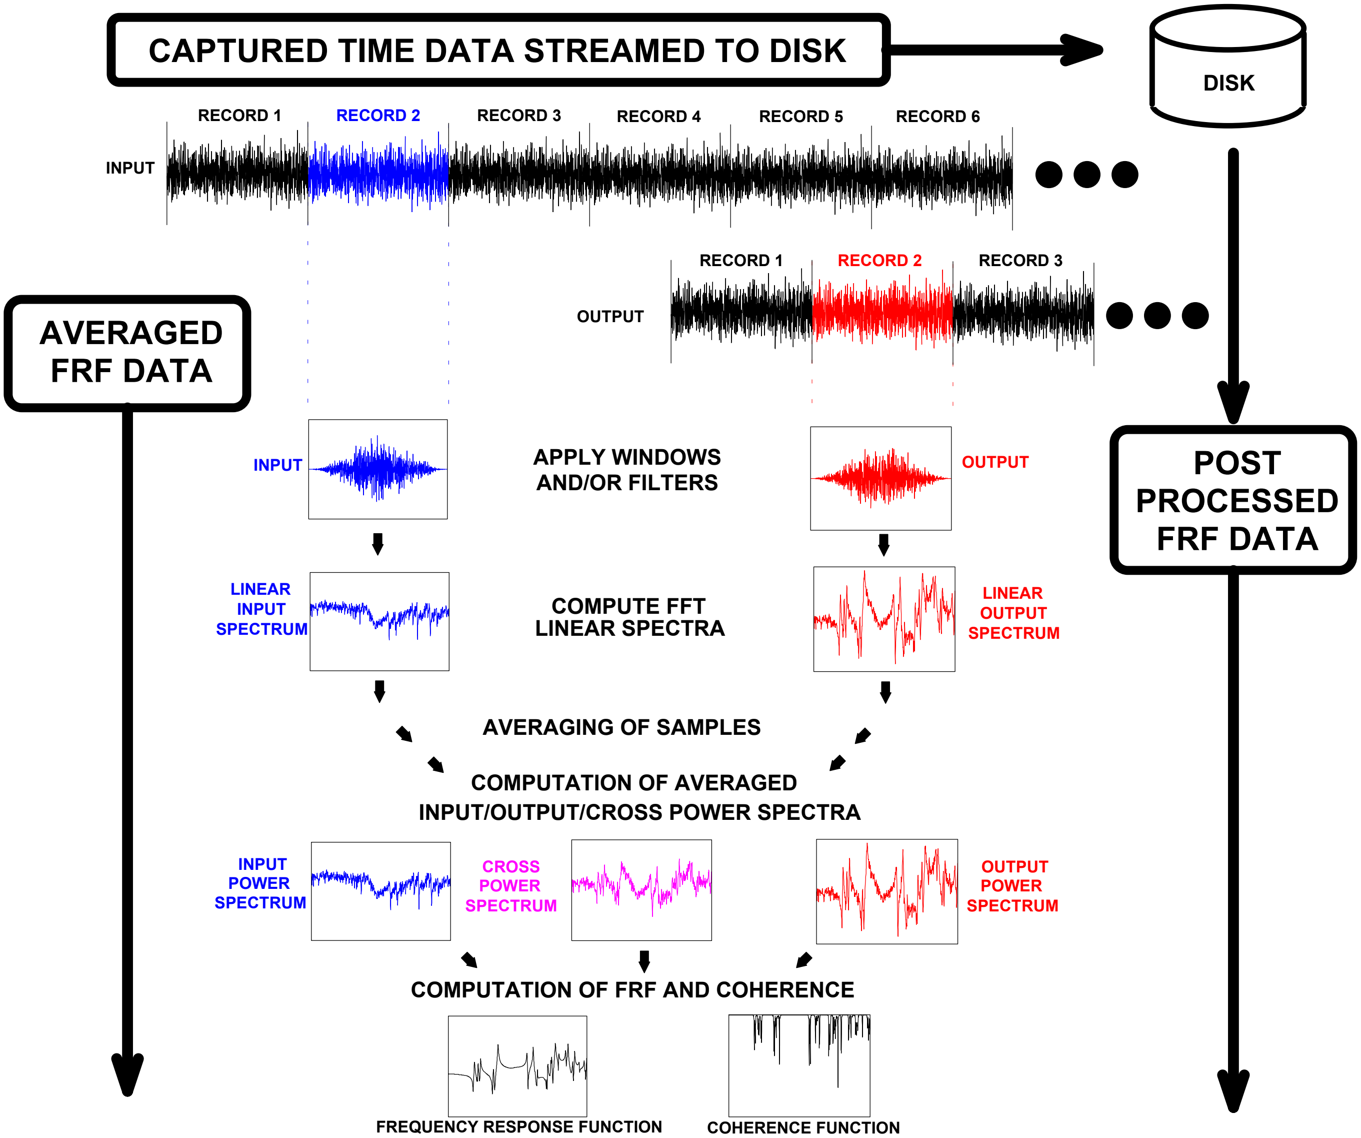

Figure 1 - Schematic Overview of Time Streamed Data to Disk vs. Averaged FRF Data

However, whenever any testing is performed, it is extremely dangerous to only collect time data. Now that may sound like a contradiction after everything that was discussed above but just hear me out. Whenever time data is collected there is no immediate spectral data that provides any information as to the adequacy of the data. All we know is that we have collected time data streamed to disk.

But how good will the processed FRFs be once the time data is processed? Has sufficient data been collected to obtain good FRFs with acceptable coherence functions? And many other statements can be raised here as to the adequacy of the time data collected. You will not know how good the data is until it is processed. You would hate to come back from an expensive field test only to find out that all the data is unacceptable!

So the rule is - collect all the time data you want but you had better process some typical data sets in the field to assure that the data collected will be acceptable. There is no substitute for viewing an FRF and coherence!!!

I hope discussion helps with time streamed data vs. averaged FRF data. If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Why is calibration and mode shape scaling important?
And does it make a difference?
Let's talk about this

Calibration and mode shape scaling are two important items for the development of an accurate dynamic model that would be used for other structural dynamic studies. Some of these would be simulation and prediction, modification, correlation, to name a few. While there may be some instances when calibration and scaling may not be critical, I will always recommend that they are done since this may be the only data ever acquired. First let's discuss calibration and then discuss scaling.

Calibration of the whole acquisition system is very important. Back in the early days of 2 channel FFT analyzers, many times we may have stepped around calibration when performing troubleshooting or quick investigative tests since we were only interested in the ratio of output to input - so the exact units may not have been critical. This may have been tolerable since we may have only been interested in general shapes of the structure. But as soon as the use of the modal data for simulation, prediction, etc. was needed, then a fully calibrated model was necessary. An accurate calibration was required when the dynamic model is used for other structural dynamic studies.

As larger channel FFT analyzers became available, the use of several accelerometers (possibly with different sensitivities) required at least some nominal calibration value to be used. If not, then different regions of the structure could possibly show relative differences in shape which could cause confusion in understanding the mode shapes.

So what calibration should be performed? Well a complete calibration is always best. This would involve a complete calibration of an entire acquisition channel as a unit - the accelerometer, signal conditioner, ADC channel together. While calibration of each individual piece is often acceptable, the calibration of the complete system together is preferred. There are a variety of types of calibration. Accelerometers can be calibrated relative to some well-maintained reference accelerometer that is traceable to a source. This can be performed in the lab using a piggyback arrangement for the test accelerometer to the reference accelerometer. Or the test accelerometer can be calibrated through drop test with some known mass. Another common calibration utilizes an excitation through a force gage mounted to a known mass with an accelerometer. This ensures that the ratio of force to acceleration using the equation of motion with the known mass. (The individual transducers can be identified if one of them is known). The most accurate way to perform this calibration is through the acquisition channels to be used for the test.

And while many calibration service companies provide calibration in fixed increments (ie, 50, 100, 200, 500, 1000, ...), this only provides information at those discrete frequencies. The better way of calibrating is to perform broadband input excitation over the frequency range of interest.

Now that calibrated modal data has been addressed, we need to discuss mode shape scaling. Yes - I know that the shape is the relative motion between points. But there is a scaling that needs to be preserved. That is, the relationship between the modal mass, modal damping and modal stiffness. The shapes can be anything but the relationship between the shapes and physical quantities is very important. The shapes can be scaled to anything you'd like - but the most common is *unit modal mass* scaling (but others such as unit length, largest value equals one are also common). The most important item is that the shapes are scaled to some quantity that is identified for future reference. Scaling can be a critical item with respect to further

Copyright 2003
Pete Avitabile

use of the dynamic model for simulation, prediction, correlation, etc.
In order to address scaling, a simple 3 dof system will be used.

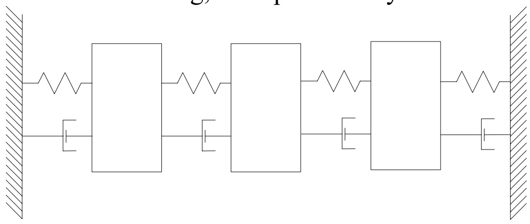

The equation of motion and specific system values are:

$$[M]\{\ddot{x}\}+[C]\{\dot{x}\}+[K]\{x\}=[F(t)]$$

$$\begin{bmatrix}1 & & \\ & 1 & \\ & & 1\end{bmatrix}\begin{Bmatrix}\ddot{x}_1\\ \ddot{x}_2\\ \ddot{x}_3\end{Bmatrix}+\begin{bmatrix}0.2 & -0.1 & \\ -0.1 & 0.2 & -0.1\\ & -0.1 & 0.2\end{bmatrix}\begin{Bmatrix}\dot{x}_1\\ \dot{x}_2\\ \dot{x}_3\end{Bmatrix}+\begin{bmatrix}20000 & -10000 & \\ -10000 & 20000 & -10000\\ & -10000 & 20000\end{bmatrix}\begin{Bmatrix}x_1\\ x_2\\ x_3\end{Bmatrix}=\begin{Bmatrix}f_1\\ f_2\\ f_3\end{Bmatrix}$$

and the corresponding eigensolution yields:

$$\left[[K]-\lambda[M]\right]\{X\}=\{0\}$$

$$\left[\Omega^2\right]=\begin{bmatrix}5858 & & \\ & 20000 & \\ & & 34142\end{bmatrix}\;;\;[U]=\left[\{u_1\}\;\{u_2\}\;\{u_3\}\right]=\left[\begin{Bmatrix}0.500\\ 0.707\\ 0.500\end{Bmatrix}\begin{Bmatrix}0.707\\ 0\\ -0.707\end{Bmatrix}\begin{Bmatrix}-0.500\\ 0.707\\ -0.500\end{Bmatrix}\right]$$

Now the first mode of the system will be the only one addressed.
The frequency, damping and complex pole for mode 1 are:
Frequency 12.18Hz Damping 0.038%
Complex pole -0.029 ± j 76.537 rad/sec

Let's recall that the poles and residues are the values that describe the FRF measured. For mode 1, this is

$$h(j\omega)=h(s)\Big|_{s=j\omega}=\frac{a_1}{(j\omega-p_1)}+\frac{a_1^*}{(j\omega-p_1^*)}$$

Now let's also recall that the residues are directly related to the mode shapes of the system from

$$\left[A(s)\right]_k=q_k\,\{u_k\}\,\{u_k\}^T$$

which can be expanded as

$$\begin{bmatrix}a_{11k} & a_{12k} & a_{13k} & \cdots\\ a_{21k} & a_{22k} & a_{23k} & \cdots\\ a_{31k} & a_{32k} & a_{33k} & \cdots\\ \vdots & \vdots & \vdots & \ddots\end{bmatrix}=q_k\begin{bmatrix}u_{1k}u_{1k} & u_{1k}u_{2k} & u_{1k}u_{3k} & \cdots\\ u_{2k}u_{1k} & u_{2k}u_{2k} & u_{2k}u_{3k} & \cdots\\ u_{3k}u_{1k} & u_{3k}u_{2k} & u_{3k}u_{3k} & \cdots\\ \vdots & \vdots & \vdots & \ddots\end{bmatrix}$$

The scaling constant 'q' is a very important term in this equation. While there are many different types of scaling that may be used, the most common one is *unit modal mass* scaling. (This is also a very common scaling used in finite element modeling software packages). With this, the system parameters of modal mass, modal damping and modal stiffness are defined as:

modal mass $\quad \bar{m}_k=\dfrac{1}{q_k\bar{\omega}_k}$

modal damping $\quad \bar{c}_k=2\sigma_k\bar{m}_k$

modal stiffness $\quad \bar{k}_k=(\sigma_k^2+\bar{\omega}_k^2)\bar{m}_k$

Now if we consider the first column of these equations, then the residues can be related to the mode shapes using

$$\begin{Bmatrix}a_{11k}\\ a_{21k}\\ a_{31k}\\ \vdots\end{Bmatrix}=q_k u_{1k}\begin{Bmatrix}u_{1k}\\ u_{2k}\\ u_{3k}\\ \vdots\end{Bmatrix}$$

We notice that there is a scale factor 'q' which is important in this equation. This scaling constant helps to preserve the proper scaling relationship between the mode shapes and the system modal mass, modal damping and modal stiffness. Notice that if we take a measurement such as h_{31} that a_{31}=q u_3u_1 (also h_{21} that a_{21}=q u_2u_1 and so on). For each of these equations, there is one extracted value of the residue (from the curvefitting process) but two values of the mode shape. So the best that can be said about the mode shape is that there is a 'relative' motion between the various points. This relative motion using the residues can be animated and provides a wealth of knowledge. However, if the drive point measurement is considered, then we see that h_{11} provides a_{11}=q u_1u_1 and it is this equation that can be used to solve for u_1 which is then used to scale all the other terms that were measured.

Every now and then I will hear someone say that there is no need to scale the mode shapes and that there is no need to take a drive point measurement. While this is true to in order to visually observe the mode shapes, without any scale factor, this modal information cannot be used for any further analytical manipulation using this data. The scaled modal data is required for any further analyses such as structural modification, forced response, prediction, simulations, correlation, etc.

Since the modes in this example are real normal modes, the residues are complex valued but will only have an imaginary part of the residue. In order to simply the numbers, the residue will be converted to a real valued expression using r = 2 j a - (note that this is a common representation of the residue in many commercially available modal analysis packages).

The values of the residues 'r' for this example for mode 1 are
dof1 = (0.003266 ± j 0.0)
dof2 = (0.004619 ± j 0.0)
dof3 = (0.003266 ± j 0.0) .
Then the relationship of the residues to the mode shapes with the scaling factor are given as:

$$\begin{Bmatrix}0.32664E-2\\ 0.46194E-2\\ 0.32664E-2\end{Bmatrix}=\begin{Bmatrix}r_{11}\\ r_{21}\\ r_{31}\end{Bmatrix}^{(1)}=q_1u_1\begin{Bmatrix}u_1\\ u_2\\ u_3\end{Bmatrix}^{(1)}=\frac{1}{76.537}(0.500)\begin{Bmatrix}0.500\\ 0.707\\ 0.500\end{Bmatrix}$$

So it can be seen that the is a definite scaling relationship that does exist and the mode shapes ***must*** be scaled using the drive point measurement in order to accomplish this. If you have any other questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Is it really necessary to reject a double impact?
Are they really a problem?
Let's talk about this.

I know that many people will say that a double impact is totally unacceptable. It is not the optimum condition for collecting impact data for a modal test. However, under certain circumstances, it may be reasonable to accept a measurement that has resulted from a double impact. Let's discuss this problem that may arise when a double impact occurs and explain how to determine if the measurement is acceptable or not.

First, let me state that I would like to avoid double impact measurements at all costs. It is not a desirable situation but at times it is unavoidable. The real concern should be the adequacy of the frequency response measurement which is really the deciding factor. Just recently, I was involved in a modal test where double impact measurements were a concern. The engineer involved in the test was quite firm in his position that no double impact measurements would ever be acceptable. (In fact, the engineer quoted reputable sources as to his position on this subject). When asked why double impacts were unacceptable, the engineer responded with a comment that any good test engineer knows not to accept a double impact measurement! This is good advice but the engineer really didn't understand what were the limiting factors in this type of measurement situation; he only knew not to accept the measurement. Of course , if the measurement is not good as evidenced by the coherence and poor input forcing function spectrum, then the double impact measurement may well not be acceptable.

Unfortunately, since double impacts were a problem, the engineer picked a measurement location that avoided double impacts but resulted in an extremely poor measurement overall. The measured FRF was much worse than the measurement that resulted from the double impact. In order to illustrate some points, let's take a look at some of the acquired measurements and explain some of the problems, pitfalls and things to consider when faced with this problematic measurement situation.

While the actual structure under test is not shown, the system can be simply depicted by the schematic shown in Figure 1. The cantilevered plate-like structure is very responsive and prone to double impacts during testing. Two measurement locations were considered - the end of the cantilevered where double impacts are likely to occur and a location on the plate closer to the cantilevered end where double impacts are avoided. (Note that in all frequency plots, a dB scale was used with 100 dB of dynamic range for plotting the input spectrum and frequency response function; the coherence is plotted on top of the frequency measurements for ease of interpretation with a range of 0 to 1).

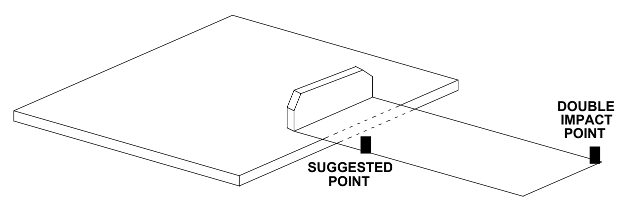

Figure 1 - Structure with Measurement Locations

The engineer wanted to avoid the "double impact location" and identified a "suggested point" where there was no double impact observed. The measurement for the "suggested point" (impact time history, input force spectrum, frequency response function and coherence) are shown in Figure 2.

The first thing to notice is that the impact and force spectrum appear to be very good. The force pulse contains one pulse and the resulting frequency spectrum is reasonably flat over the entire frequency range with less than 10 dB rolloff over the entire frequency range.

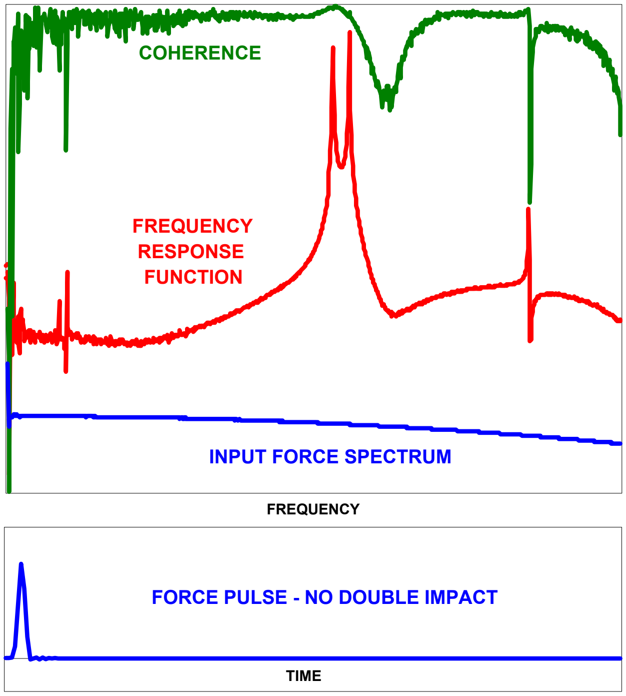

Figure 2 - Measurement with No Double Impact

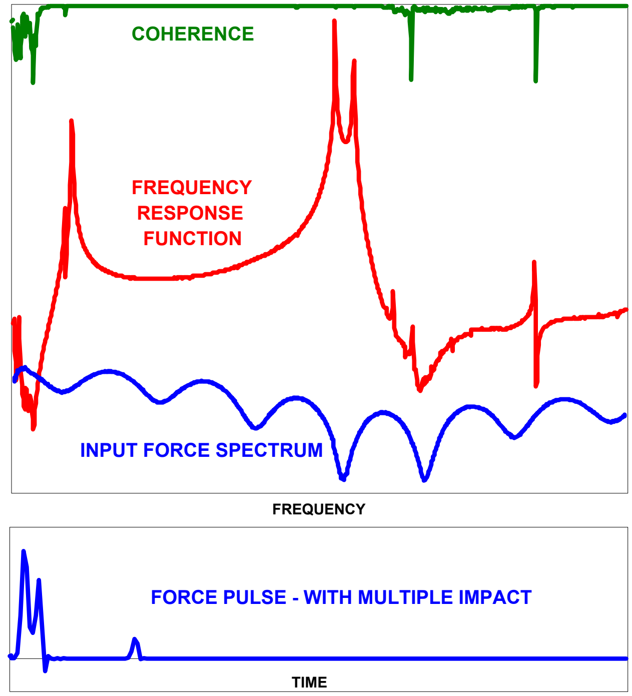

Figure 3 - Measurement with Multiple Impacts

The resulting frequency response function and coherence are also presented in the same plot. In general, the frequency response function is not particularly good as evidenced by the poor coherence at many frequencies over the spectrum. Notice that the coherence is very poor at low frequencies and there appears to be 2 peaks that are not excited very well and appear to be in the noise of the measurement.

Clearly, there is no double impact but the measurement adequacy in my mind is very poor - and most would agree that this is not a very good frequency response measurement. In fact, the engineer attempted to defend this poor measurement stating that the structure is very complicated, with many joints and possible nonlinear behavior. (I wish I had a dollar for every time I have heard that statement ! Nonlinearity and joints and damping - oh my!).

Now let's consider the measurement where the double impact, actually multiple impacts, are observed. Now this measurement clearly has multiple impacts on the input force excitation. The input spectrum is not flat and has some variance over the frequency spectrum. The actual variation is between 20 and 25 dB over the spectrum. Of course, I agree that I would avoid this particular measurement but the frequency response measurement and coherence are actually very good.

The frequency response function is relatively good and the peaks in the measurement are well defined especially at the lower frequency range (where previously the peaks were not measured well and were contaminated with noise). There are also two additional peaks in the higher frequency range that were not even observed in the previous measurement. Actually if I hadn't made all this fuss concerning the double impact and just showed the frequency response function and coherence, most people would have accepted the measurement without any questions. (It also can be stated that the previous measurement would not have been considered acceptable if only the frequency response function and coherence were shown).

So what do we need to be concerned about? If the input spectrum has significant drop out at any particular frequency, then the measurement may not be adequate. But before we can make any assessment, the input force spectrum, frequency response function and coherence need to be reviewed and evaluated. We cannot make a blanket statement that double impact measurements are unacceptable.

I agree that I will avoid double impacts at all costs to be safe - but we have to realize that the double impact itself is not necessarily a problem if the input force spectrum, frequency response spectrum and coherence are all acceptable. If you have any more questions on modal analysis, just ask me.

Copyright 2003
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Someone told me that you must have multiple references to identify pseudo repeated roots.
Let's talk about this.

First, let me say that repeated roots are generally rarely found in the majority of structures we test. But recently, I have seen many tests with what are referred to as "pseudo repeated roots". This refers to modes with two or more frequencies that are present in one discrete Δf spectral line of the measured data. Therefore, the observed peak is comprised of the sum of two or more modes in the response function measured.

In order to see the peaks as distinct individual peaks, finer frequency resolution is needed. Of course, this may not be feasible. Many times to obtain this resolution would require excessively long time blocks which may not be easy to obtain. But what happens if sufficient resolution is not available. Typically, most people would state that if multiple roots exist, then multiple references must be acquired otherwise the roots cannot be extracted. Now this is a very strong statement and I do not necessarily agree. The modal parameter estimation algorithms (curvefitters) are very robust and there is really no reason to think that the algorithms are deficient when it comes to extracting pseudo-repeated roots (two examples will be shown). Multiple references are actually not always necessary in these situations.

First, let me make sure that there are no misrepresentations here. I fully advocate using multiple reference data collection when conducting an experimental modal test. With multiple channel acquisition systems commonly available, it is easy to acquire multiple reference data. Typically with a 4 or 8 channel data acquisition system, multiple accelerometers can be placed at a variety of different locations as references for an impact test (MRIT - multiple reference impact test). Then, the impact data can be collected at all the measurement locations with the multiple stationary acceleometer locations. (Of course with a shaker excitation modal test, multiple reference data can be more work if multiple shakers are not available. But, typically multiple reference data can be collected.) But the question really is this - Is multiple reference data required to extract multiple roots?

As far as I am concerned, the answer is not necessarily so. If the data collected is good measured data, then the frequency response function is the sum of all the modes and the modal parameter estimation algorithms can extract multiple roots accurately from the measured data. Two structures are evaluated to show the results of extraction of multiple roots (two distinct modes within one Δf of the analysis spectral resolution) from measured data. In both cases, there appeared as if only one root existed as indicated by the available mode indicator tools (SUM, MIF, etc). In one structure, only a handful of cross directional FRFs from a total of over 100 measurements showed some indication of two frequencies. In the other structure, none of the measurements revealed the fact that the first two "peaks" in the FRF contained multiple roots at each peak!

The first structure was a prototype composite spar from a wing structure. The geometry typified a taperd beam (with a type of I shape) which had no geometric symmetry. A typical impact measured FRF is shown in Figure 1 with a photo of the structure.

The measurement is reasonably good with some noise seen on the measurement over the frequency range. A frequency domain polynomial curvefit algorithm was used to extract mode shapes. (Actually, three different commercially available modal packages were used with essentially identical results). The second peak in the function actually contains two separate roots. The curvefitter extracted the two roots at the second peak with no difficulty whatsoever even though there appeared to be only one peak present. Clearly, the modes extracted show two well defined mode shapes shown in Figure 2. Even though only one

Copyright 2003
Pete Avitabile

reference was available to extract modal parameters, two modes were very successfully extracted!

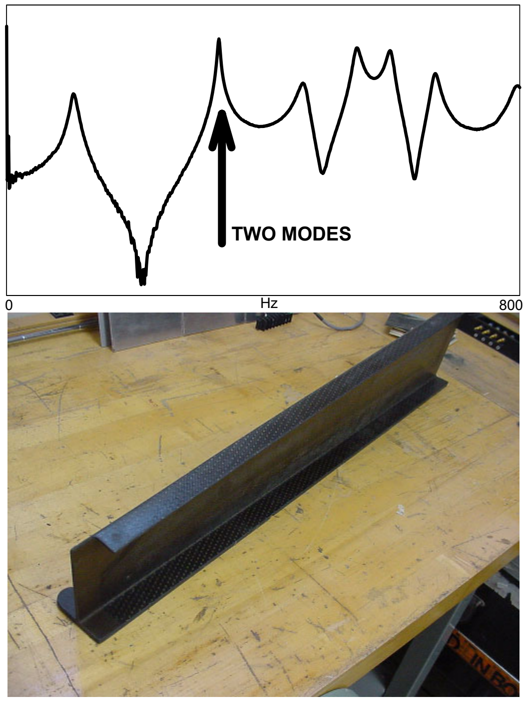

Figure 1 - Typical FRF and Photo of Structure

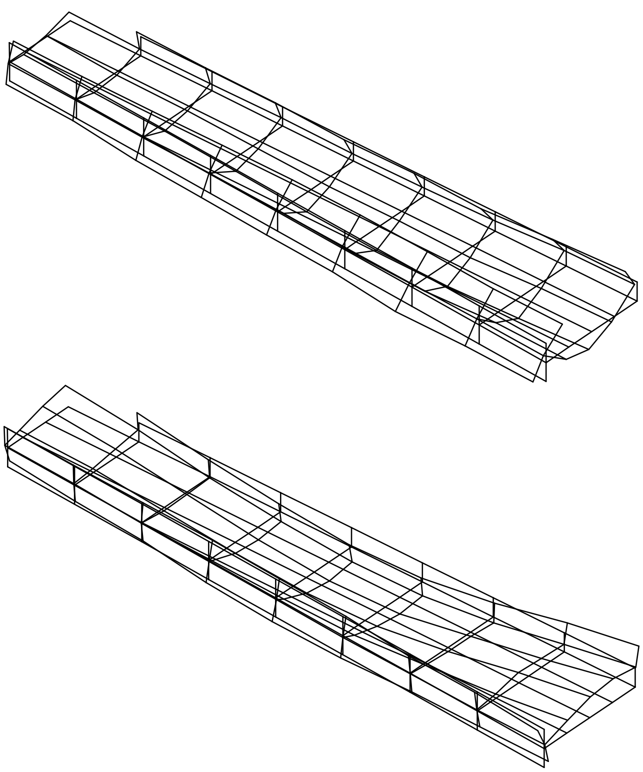

Figure 2 - Pseudo-Repeated Mode Shapes at Second Peak

The second structure was a simple magnesium shaker table slip plate. The geometry was such that multiple roots were not actually anticipated. A typical impact measured FRF is shown in Figure 3 with a photo of the structure.

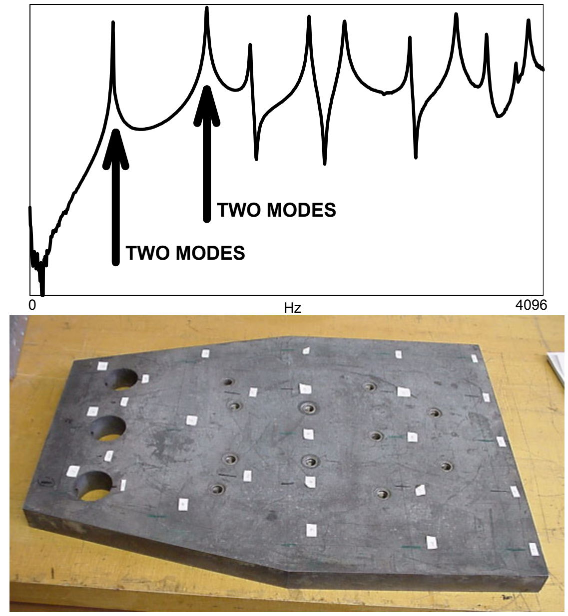

Figure 3 - Typical FRF and Photo of Structure

The first peak appeared to be a single mode for all of the FRFs obtained; to some degree the second peak had similar characteristics. The mode indicator tools also indicate that only one mode is expected. However, estimating parameters assuming only one mode does not provide the "expected" mode shapes. Upon refitting the data with two modes at each peak reveals the mode shapes expected; the shapes are shown in Figure 4.

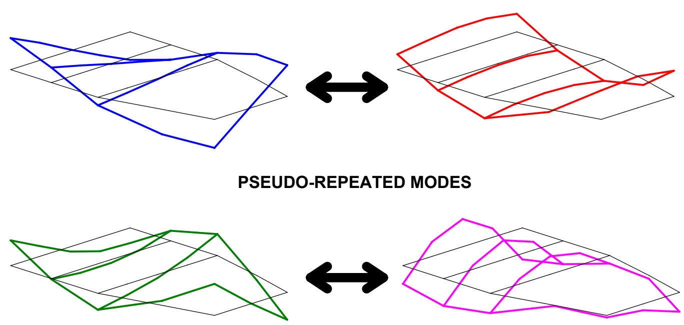

Figure 4 - Pseudo-Repeated Mode Shapes

Now in both cases, the mode indicator tools did not necessarily indicate multiple roots. User intervention was required in order to extract the multiple roots. Also notice that the multiple roots were very successfully extracted with single reference data - multiple reference data was not required! Please understand that multiple reference data is extremely useful - but is not always required in order to extract multiple pseudo-repeated roots.

I hope this answers your question. If you have any more questions on modal analysis, just ask me.

Copyright 2003
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

I ran a modal test on a portion of a structure of concern and many modes look the same!
What did I do wrong?
Let's talk about this.

This is another common problem that I often see in experimental modal testing. Many times only a portion of the structure or component of the system is of interest to you or your company. So immediately you focus on only that portion of the overall structure since that is your area of responsibility or concern. This seems reasonable especially since you may not want to test the entire system. (Ahhh - if life could be so simple and easy!).

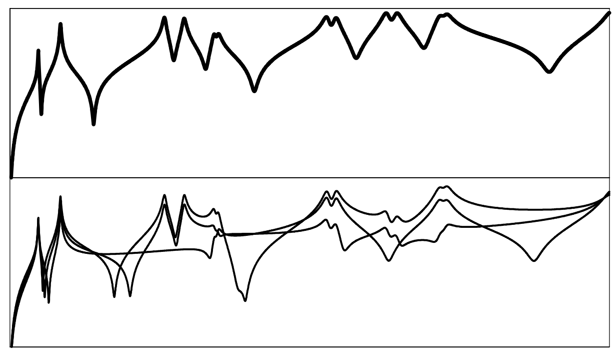

Figure 1 - Drive Point FRF and Typical Cross FRFs

Unfortunately many times this may not be possible. Most times there is significant dynamic coupling between different components in the system or different portions of the system. It is not always possible to just measure the portion of the structure of interest to you. Of course, you can certainly measure <u>only</u> the portion of the structure of concern, but many times there may be significant dynamic interaction between the various components of subsystems in the system. If measurements for an experimental modal test are only collected over a portion of the structure, then the mode shapes may be confusing since the entire mode shape over the whole structure is not known. It is as if you have put blinders on your view to only look at a portion of the structure - this can leave the user fairly confused. Many times people will comment on this type of test data that there are two first bending modes or two first torsional plate modes, etc. Obviously this is entirely not possible! There cannot be two first modes of the system. But from your limited vantage point when only a portion of the structure is measured, it certainly appears as if there are two very similar modes.

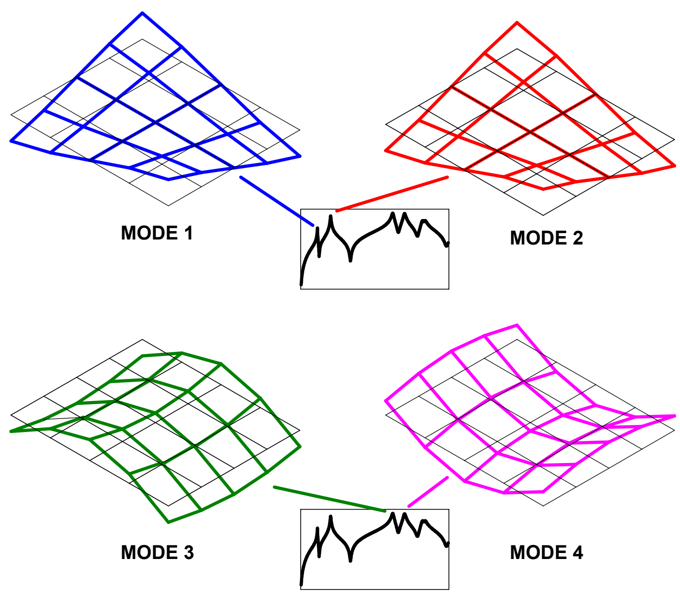

Figure 2 - Mode Shapes of Upper Plate Component

Recently, I saw a modal test of a frame type structure that had various platforms at different levels. An experimental modal test was performed on one platform surface since some important equipment was mounted on one particular platform.

Copyright 2003
Pete Avitabile

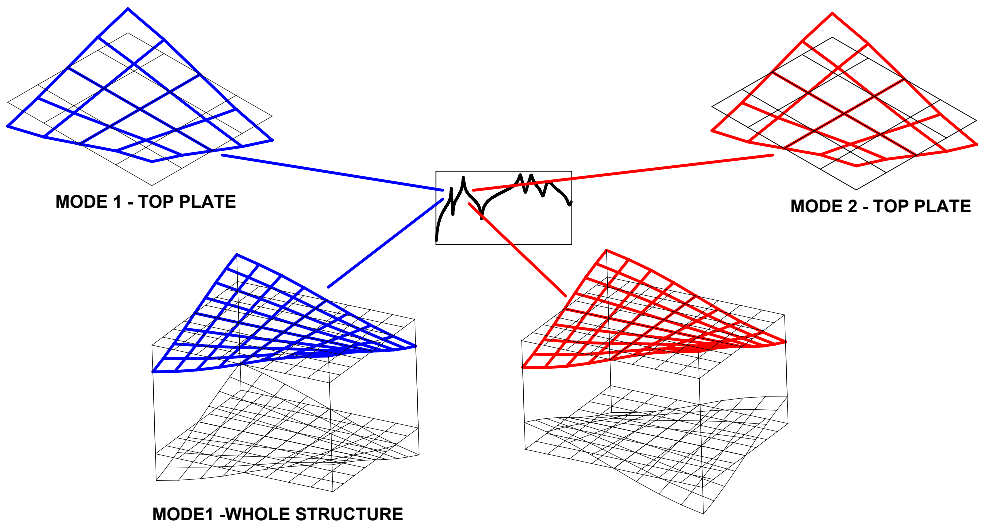

Figure 3 - Mode Shapes of Upper Plate Component along with Entire Structure

Since data was only collected only on that one platform, it appeared as if there were many similar (and almost exactly the same) modes. The problem clearly is that only a portion of the entire mode shape was identified and does not provide a clear understanding of the entire mode shape.

As an example of this problem, a simple two plate structure separated by support columns is used to illustrate typical FRFs and resulting mode shapes from an experimental modal analysis. In one case, only one plate was measured and in another case both plates were measured (but only vertical motion was considered to simplify the explanation of this problem). A typical drive point FRF on the upper plate along with some other typical cross FRFs are shown in Figure 1.

Upon reducing all the data (considering only the upper plate of the structure, the experimental mode shapes revealed two very similar torsion and two very similar bending modes for the first four peaks seen in the FRF. These shapes are shown in Figure 2. Obviously, this is not possible - but with the limited number of measurements on just a portion of the structure, this is entirely possible.

In order to better understand the 'actual' mode shapes of the structure, a more extensive array of points were used to describe the FRF matrix. These FRFs were used to determine the mode shapes of the entire system and are shown in Figure 3 for the first two peaks in the FRF. Clearly, the addition of the extra points clarifies the actual mode shapes of the upper plate in relation to the rest of the structure. This phenomena happens often in many structures when testing is performed for only a small portion of an entire structure or system.

In this simple example, it is clear that the mode shapes for the entire system must be obtained otherwise some confusion may exist. However, this also occurs many times when complicated structures are tested where access to the entire structure is not possible. This might happen with internal components that are not easily accessible for testing. These internal components may have significant modal energy related to one or modes of the structure or system. In these cases, only a portion of the entire mode shape is acquired since it is not possible to instrument interior portions of the structure. Just imagine in the two plate example if the lower plate were not visible or covered by some exterior shroud or covering. If the lower plate were not accessible, the measurements may only reveal the portion of the mode shapes related to the upper plate. In this case the same problem will exist.

This happens many times with structures where all significant portions of the structure are not available for instrumentation or where some disassembly is required to gain access to all the pertinent areas of the structure for modal testing. In these cases the same problem exists. So it is very important to be careful when testing structures where interior components or subsystems are not readily accessible for testing - there may be significant regions of the structure that contain critical information that identify the modal character of the system.

I hope this answers your question. If you have any more questions on modal analysis, just ask me.

Copyright 2003
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

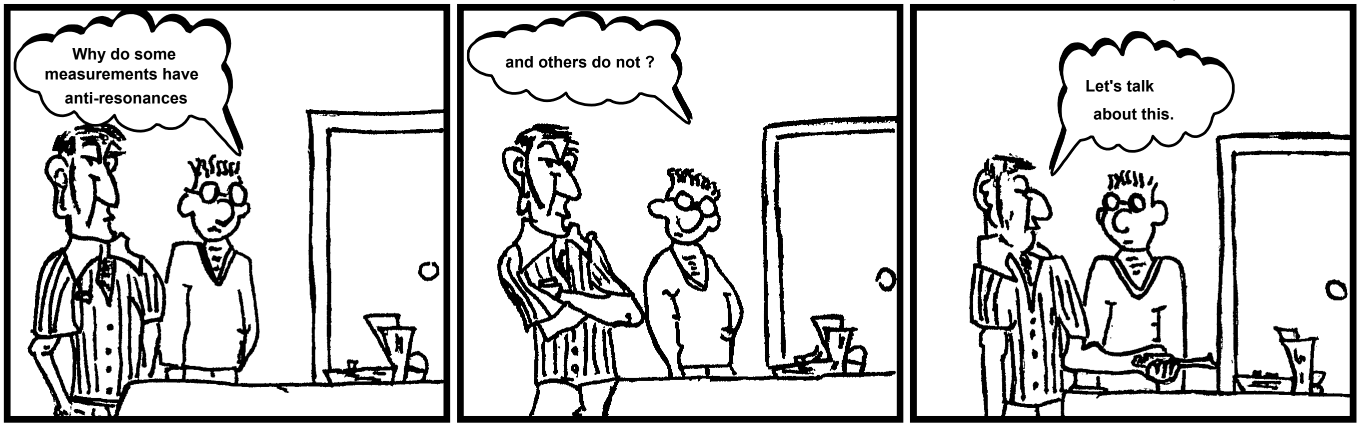

Illustration by Mike Avitabile

Why do some measurements have anti-resonances and others do not?
Let's talk about this.

This is a good question. You are absolutely correct - some measurements have anti-resonances and others do not. But why does this happen. Let's first discuss some properties of a particular measurement called a drive point measurement and then extend this discussion to explain how anti-resonances occur in a measurement.

Let's first explain a drive point measurement. A drive point measurement is one where the input force and output response are made at the same point and in the same direction. A typical drive point measurement is shown in Figure 1.

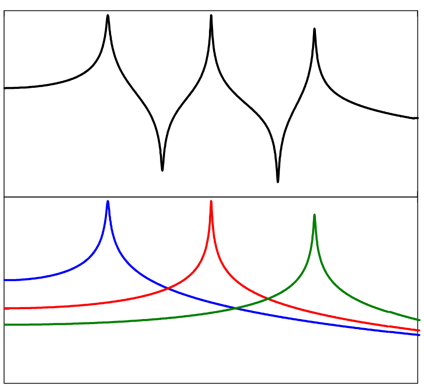

MAGNITUDE

$$h_{ij}(j\omega) = \frac{a_{ij1}}{(j\omega - p_1)} + \frac{a_{ij1}^*}{(j\omega - p_1^*)} + \frac{a_{ij2}}{(j\omega - p_2)} + \frac{a_{ij2}^*}{(j\omega - p_2^*)} + \frac{a_{ij3}}{(j\omega - p_3)} + \frac{a_{ij3}^*}{(j\omega - p_3^*)}$$

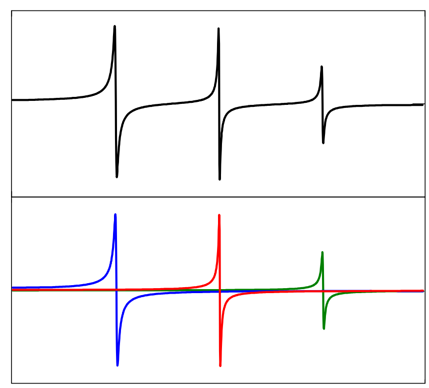

REAL

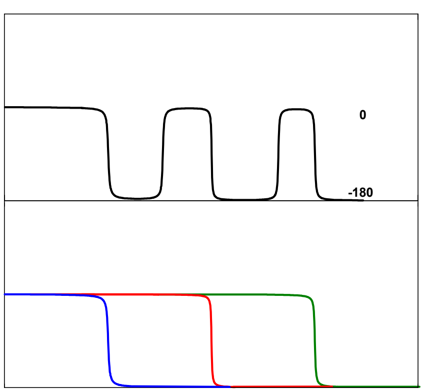

PHASE

$$h_{ij}(j\omega) = \frac{q_1 u_{i1} u_{j1}}{(j\omega - p_1)} + \frac{q_1 u_{i1} u_{j1}{}^*}{(j\omega - p_1^*)} + \frac{q_2 u_{i2} u_{j2}}{(j\omega - p_2)} + \frac{q_2 u_{i2} u_{j2}{}^*}{(j\omega - p_2^*)} + \frac{q_3 u_{i3} u_{j3}}{(j\omega - p_3)} + \frac{q_3 u_{i3} u_{j3}{}^*}{(j\omega - p_3^*)}$$

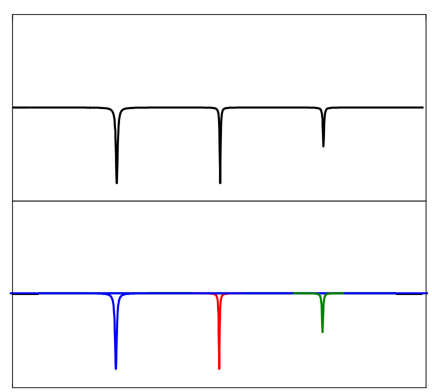

IMAGINARY

Figure 1 - Drive Point FRF (Magnitude, Phase, Real, Imaginary)

Copyright 2003
Pete Avitabile

For a driving point measurement several items can be noted:

- all resonances are separated by antiresonances as seen in the magnitude plot
- the phase looses 180 degrees of phase passing over a resonance and gain 180 degrees of phase passing over an antiresonance
- the peaks in the imaginary part of the FRF must all point in the same direction

The drive point measurement can be viewed as a summation of all the modes or as the contribution due to each mode. As seen in the four plots in Figure 1, the upper plot contains the summation due to all the modes and the lower plot shows the contribution due to each mode. For the first three modes shown, the frequency response function is made up of the sum of each of the single degree of freedom oscillators describing each mode of the system. For reference, recall that the frequency response function equation can be written as either residues or mode shapes as shown in Figure 1.

Now that the drive point measurement is understood, several other items can be discussed. For instance, the imaginary part of the frequency response function must all have the same direction and in this condition an anti-resonance exists between each mode. This is due to the fact that the magnitude of the FRF of mode 1 and mode 2 is equal at the anti-resonant frequency. But at this frequency, while the magnitudes are equal, the phase is 180 degrees out of phase with each other. This implies that the sum of mode 1 and mode 2 are equal and opposite. Therefore the function trends towards zero. (There is actually a contribution from other modes that is generally very small when the modes are far spaced as shown.)

Now this implies that when the imaginary part of each mode has an opposite sign, the phase is not necessarily out of phase - and then the modes add and an anti-resonance does not result. So each measurement can have anti-resonances or no anti-resonances (saddles) depending on the direction of the imaginary part of the frequency response function. When the imaginary part of the frequency response function for sequential modes have the same direction, then an anti-resonance will exist between those two modes. When the imaginary part of sequential modes have different signs or directions, then a saddle exists between those two modes.

Actually, the direction (or sign) of the function is directly related to the mode shapes of the system. As seen in Figure 1, the frequency response function can be written in the form of residues. But the residues can be expressed in terms of the mode shapes of the system. When written as mode shapes, the directional sign of the residue can be clearly seen as a result of the mode shape of the system. Figure 2 shows the measurements for a simple 3 DOF system. Upon reviewing each of the individual FRF measurements, the phase relationship and occurrence of anti-resonances and saddles in the frequency response function can now be better understood.

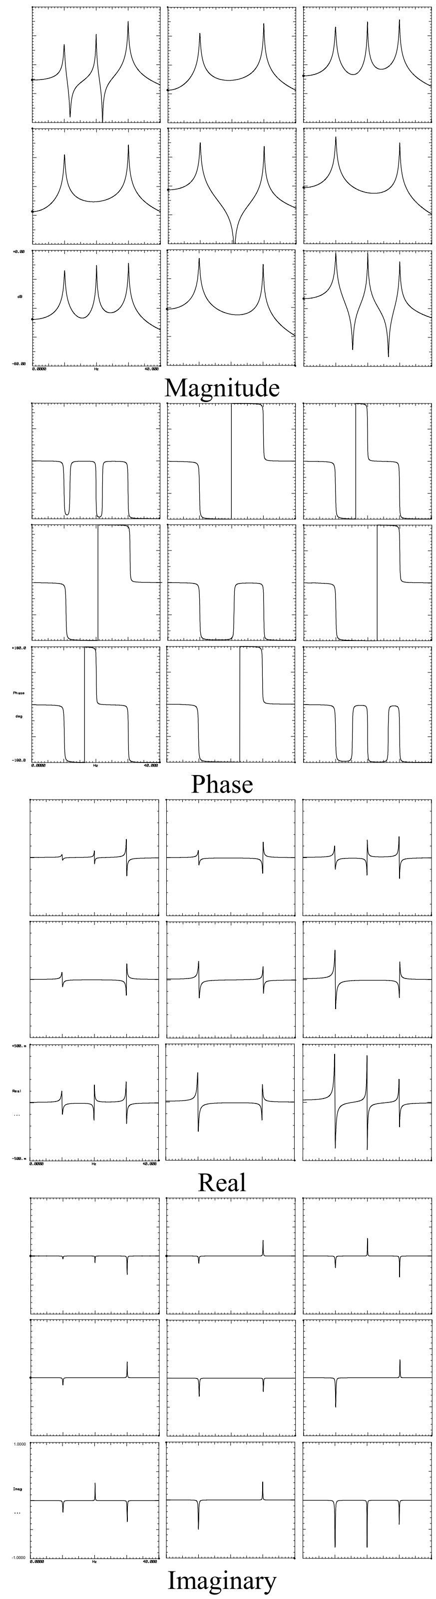

Figure 2 - FRF Matrix for a 3 DOF System

I hope this explanation answers your question. If you have any more questions on modal analysis, just ask me.

Copyright 2003
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

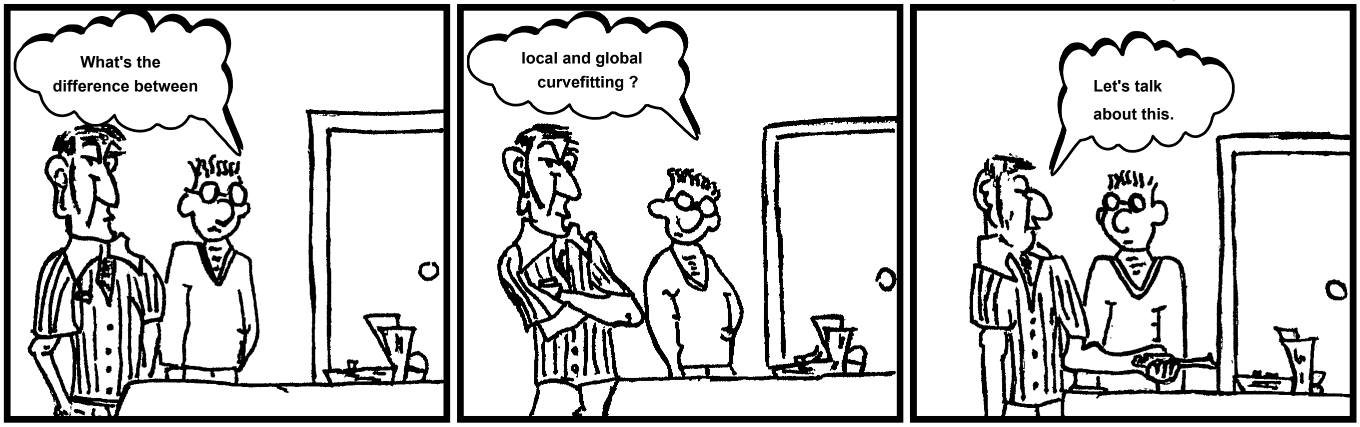

Illustration by Mike Avitabile

What's the difference between local and global curvefitting ??
Let's talk about this.

This is a good question. In order to explain this, a few quick equations are needed followed by a simple example that will illustrate the differences. Let's recall the frequency response function which is

$$h_{ij}(j\omega) = \sum_{k=1}^{m} \frac{a_{ijk}}{(j\omega - p_k)} + \frac{a_{ijk}^*}{(j\omega - p_k^*)}$$

There are numerous 'ij' (output-input) combinations; a matrix of possible FRFs is illustrated in Figure 1.

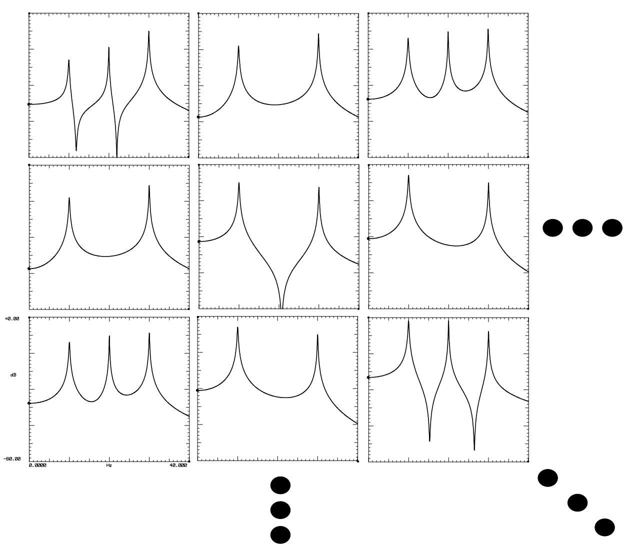

Figure 1 - FRF Matrix for a Multiple DOF System

Now each FRF is defined by poles and residues. The FRF is different from one measurement to the next because the residue is different. This is true since the mode shapes are related to the residues as

$$a_{ijk} = q_k u_{ik} u_{jk}$$

But it is very important to note that the denominator of the FRF is constant and does not change from one measurement to the next. Since the pole does not change from one measurement to the next, then it is said that the pole is a "global" property of the system. This means that while the residue changes from one measurement to the next, as expected, the pole does not change - at least theoretically! But in real measurements, this may not necessarily be the case. In actuality, the pole may shift from one measurement to the next. This can cause a problem.

To understand this, consider data to be fit with a straight line as shown in Figure 2. Now, if only two points are selected (blue) different from another set of points (red), there can be dramatic differences in the slope and y-intercept computed from the two sets of points. In other words, there are differences and inconsistencies in the slope and y-intercept depending on which data is used to extract parameters. When all the data is used together in a least squares fashion, then the "best" overall estimate of the slope and y-intercept results.

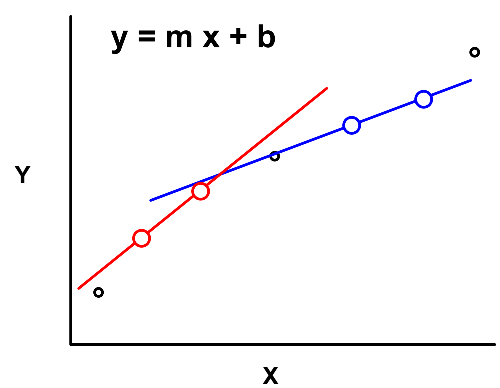

Figure 2 - Illustration on Parameter Variation

The same effect can be seen in the extraction of modal parameters from measured FRFs if each FRF is evaluated

Copyright 2003
Pete Avitabile

independently from every other FRF. Depending on which FRF is used, there may be differences in the estimated pole - but the theory indicates that this should not happen. However, this is exactly what happens when real measurements are used to extract modal parameters when each measurement is considered independently from each other. This is referred to as "local" curvefitting. In order to circumvent this problem, all of the measurements are used together, as one set, to find the best pole in a least squares fashion, to describe the best "global" representation of the pole. Once the pole is estimated, the residues are then estimated with the "global" estimate of the pole used in the modal parameter estimation equations. This is a two step process where the best "global" pole of the system is estimated first, followed by the estimation of residues with the estimated "global" pole of the system locked to a fixed value regardless of what each measurement may indicate. This is global curvefitting.

To illustrate the differences in local and global curvefitting, FRF measurements on a simple planar frame are used. Several FRFs are shown in Figure 3. There are 5 distinct modes in the band shown. Notice that the top two FRFs show all the peaks for each of the modes of the frame but that the lower two FRFs do not contain peaks at each one of the frequencies shown in the upper two plots. (This is due to the fact that some of the measurements are located at nodes of some of the modes.)

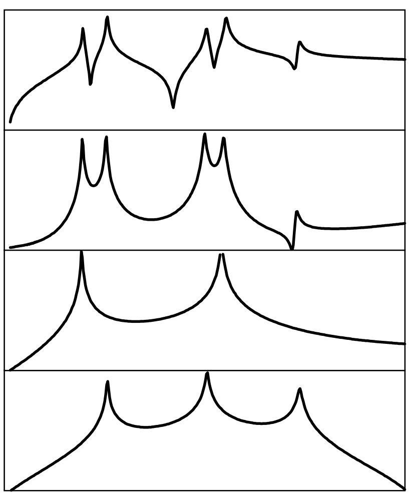

Figure 3 - Several FRFs from the Planar Frame

Now if there is no peak in a particular measurement, then how can pole values be estimated? This poses a serious problem and it is exactly these situations that the local curvefitting breaks down. If local curvefitting is performed on this type of data, then the estimated modal parameters may contain poorly extracted values from the individual FRF local curvefitting approach since the pole is estimated poorly. A local curvefitting technique was used to estimate modal parameters for the planar frame structure. The modes shapes are shown in Figure 4.

Notice that there are several locations in the mode shape where the data appears to be inconsistent from the expected mode shape. The modes shapes are distorted. It turns out that these points correspond to nodes of modes of the structure. (This is a well-known problem with local curvefitting.)

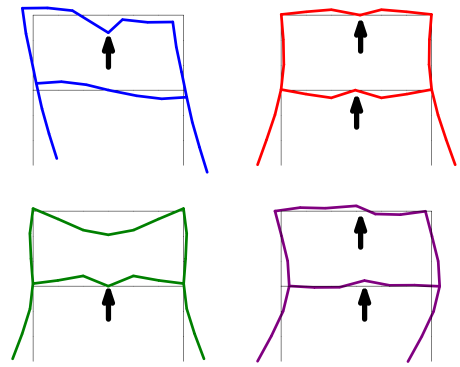

Figure 4 - Distorted Mode Shapes from Local Curvefitting

The same set of FRFs was used for global curvefitting. First, the best global pole of the system was estimated and then the residues were estimated in a second pass with the global pole used for all the FRF measurements when estimating residues. The mode shapes are shown in Figure 5. Notice that these mode shapes are the expected shapes of the planar frame structure.

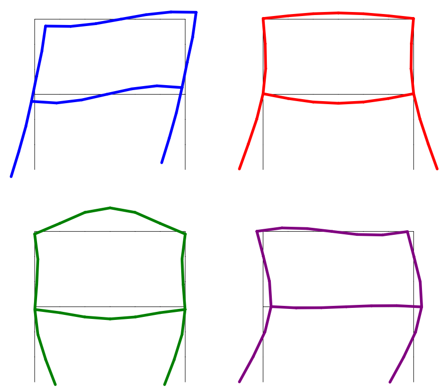

Figure 5 - Correct Mode Shapes from Global Curvefitting

Now from this example, it is clear that global curvefitting produces superior results. However, when collecting data, care must be exercised to assure that the data satisfies the requirement of global curvefitting - the modes must be global in all of the measurements collected! If the data is inconsistent, then errors may result in the estimation process. Care must be exercised to collect FRF data that satisfies the global nature necessary for the global data reduction process.

I hope this clears up your question. If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

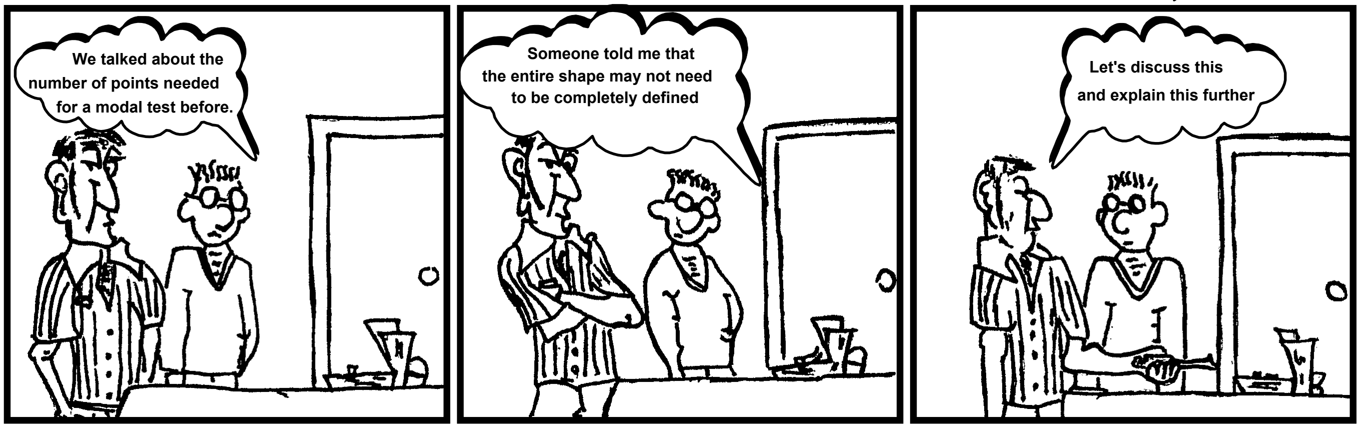

Illustration by Mike Avitabile

We talked about the number of points needed for a modal test before.
Someone told me that the entire shape may not need to be completely defined.
Let's discuss this and explain this further.

The last discussion we had regarding the total number of points for an experimental modal test was directed towards the adequate description for the mode shape. But in order to define a good dynamic model, the requirement can be different. Let's first restate what was identified earlier regarding the mode shape definition and then proceed on to identify the points needed to identify a good dynamic model.

Let's start with a simple structure that we have discussed before. The simple plate structure.

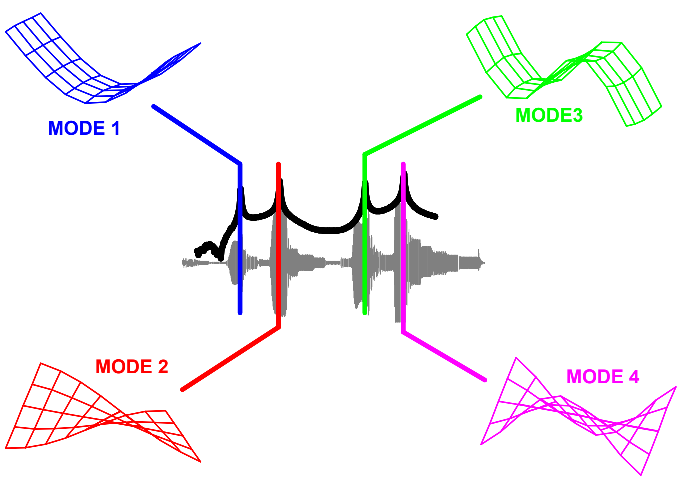

For this plate, there are many possible measurement locations that can be used to describe the dynamic characteristics of the structure. If the experimental modal analysis is being conducted to correlate with a finite element model, then it is necessary to have a reasonable number of well distributed measurement dofs in order to have sufficient spatial distribution for comparison to the model. It is also important in order to visualize the shape.

So let's consider a finite element model that has 45 node points with 32 plate elements. Let's also consider an experimental modal test with 15 evenly distributed measurement locations. A comparison of the finite element node points and experimental measurement points to describe the plate are shown for the first two free-free flexible modes of the plate.

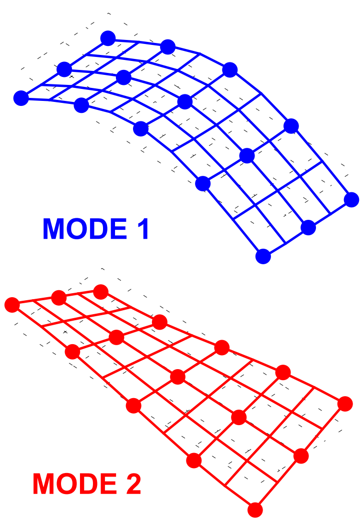

For this comparison, there is a sufficient distribution of points such that the modes can be uniquely defined for the correlation of the finite element model and experimental modal model. If the model is to be correlated using the Modal Assurance Criteria and using Pseudo Orthogonality Checks, then there is sufficient information to perform a valid analysis. However, a response model may not need the same distribution of points.

Copyright 2004
Pete Avitabile

Consider a structure where there are two separate points where forces are applied to the system. Also consider that there are only three points on the structure where critical response needs to be measured. This is pictured below with the input forces in blue and response measurements in red.

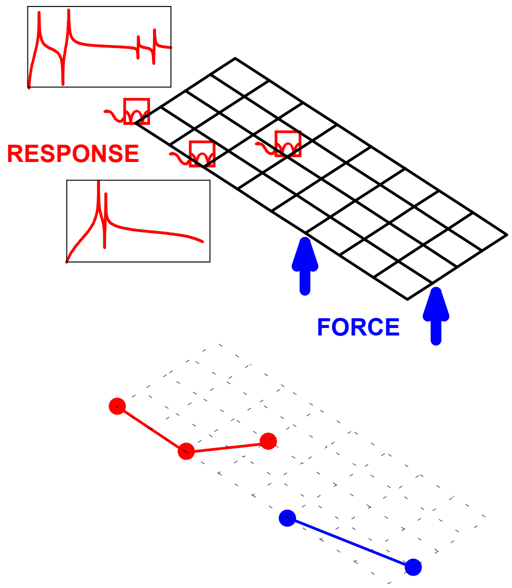

The point-to-point frequency response function describes the characteristic output response due to a unit sinusoidal input force. This is computed for a particular input-output measurement on the structure. A particular frequency response measurement has no relationship to other measurements on the structure. If only one frequency response function is of concern, it doesn't matter whether two, three or one hundred additional measurements are made - the individual frequency response measurement for a particular input-output measurement does not get better or worse as additional measurements are added.
(This is not true of the finite element approach where adding additional nodes and elements generally has a direct effect on the results.)

If the only requirement is to determine what the dynamic response of the system is due to the applied forces, then the only points that are required for the experimental modal model are the three red points and two blue points pictured. These points completely describe the response of the system at these points. Of course, the points shown may not adequately describe the deformation sufficiently to determine exactly how the load is distributed through the system - but this limited set of points is sufficient to define the dynamic model.

Now let's continue on and consider that the system may need to be modified with structural changes to the system. The figure below depicts the structure with a rib stiffener modification that could possibly be proposed to modify the structure. In order to perform any of these structural dynamic modifications, a set of points associated with all of the dofs used for the structural change need to be included. These are shown in green in the figure.

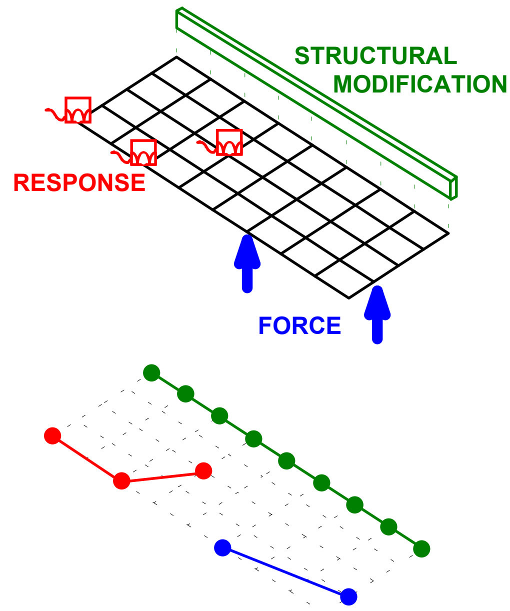

So these two models show a completely different set of measurement dofs that are needed for each of the different models described. For the response model and structural dynamic modification model, a significantly different set of dofs are needed in order to develop the proper dynamic model to describe the system.

So ..... the set of required points can be stated as follows.
In order to have an adequate dynamic model, there must be points to describe

1. all dofs where forces are applied to the system
2. all dofs where response needs to be measured
3. all dofs where structural modifications are considered

Any additional dofs included in the model are included for your viewing pleasure only!!! Additional points are included in the experimental modal test for you to better visualize the characteristic shapes of the system (or for correlation of a model if that is the purpose of the test). This minimum set of dofs is all that is required to develop an appropriate dynamic model.

Now I hope you have a better understanding of how many points are needed for a modal test - it varies depending on the ultimate use of the dynamic model. If you have any other questions about modal analysis, just ask me.

Copyright 2004
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Do I need to have an accelerometer mounted in the X, Y and Z directions to do a modal test?
Well ... let's discuss this.

This is an item that often causes confusion for many people. There is some preconceived notion that there must be an accelerometer mounted in each of the three principal directions in order to acquire data for a modal test. Well, it turns out that this is not necessary but in some tests it may be strongly advised or even required. But many times people think that you can't get three dimensional mode shapes unless you have accelerometers in all three directions.

The basic equation we use for estimating parameters can be written in one form as

$$\left[H(s)\right] = \text{lower residuals} + \sum_{k=i}^{j} \frac{\left[A_k\right]}{\left(s - s_k\right)} + \frac{\left[A_k^*\right]}{\left(s - s_k^*\right)} \quad + \text{upper residuals}$$

The terms in the matrix, [A], are the residues which are obtained from the curvefitting process; we also get the poles, or frequency and damping, from the denominator of the equation. But these residues are directly related to the mode shapes from

$$\left[A(s)\right]_k = q_k \left\{u_k\right\}\left\{u_k\right\}^T$$

This relationship between the residues and the mode shapes holds the answer to the question posed. Let's expand that equation to look at some of the terms that are found in each term of the matrix.

$$\begin{bmatrix} a_{11k} & a_{12k} & a_{13k} & \cdots \\ a_{21k} & a_{22k} & a_{23k} & \cdots \\ a_{31k} & a_{32k} & a_{33k} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} = q_k \begin{bmatrix} u_{1k}u_{1k} & u_{1k}u_{2k} & u_{1k}u_{3k} & \cdots \\ u_{2k}u_{1k} & u_{2k}u_{2k} & u_{2k}u_{3k} & \cdots \\ u_{3k}u_{1k} & u_{3k}u_{2k} & u_{3k}u_{3k} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

And if we were to look at each of the columns we would see the mode shape is contained in the column with some scalar multipliers; we would also see that due to reciprocity, the rows also contain the mode shapes. If we were to look at one column, such as the first column, then we would see

$$\begin{Bmatrix} a_{11k} \\ a_{21k} \\ a_{31k} \\ \vdots \end{Bmatrix} = q_k u_{1k} \begin{Bmatrix} u_{1k} \\ u_{2k} \\ u_{3k} \\ \vdots \end{Bmatrix}$$

So the value of the mode shape that is factored out of the equation is called the "reference" DOF. In other words, all of the measurements are affected by the value of this reference DOF. If this DOF is zero (at the node of the mode) then no matter how many measurements are made, that particular mode will not be observed from the measured data.

It is this basic equation that really contains the answer to the question raised. As long as the reference DOF has a non-zero value for each of the modes of interest, then the frequency response function will have residue associated with that input-output relationship. As long as the mode shape in the X, Y and Z direction has a value associated relative to the reference DOF, then the mode shape(s) can be observed from measurements made relative to that reference DOF. It's that simple!

Now let's use a simple structure to illustrate this point. A simple L-shaped bracket will be used for discussion purposes and illustration of the reference DOF and its relationship to all of the measured DOFs for the experimental modal survey. For the discussion, the reference point where the accelerometer is to be mounted will be shown in black in the following figures and the various impact locations will be shown in blue, red and green for distinction between the different points.

Copyright 2004
Pete Avitabile

For this first mode of this bracket, the reference accelerometer could possibly be located at the upper corner of the structure in the x-direction. Notice that if the structure is impacted at the upper corners of the bracket in the x-direction (red or green) or on the lower corner in the z-direction (blue), the structure has significant response at all these locations. This implies that if the structure is impacted in the x-direction (red or green) in the upper corner, there is response in the x-direction at the reference point. And if the structure is impacted on the lower corner (blue) in the z-direction, there is response at the reference point in the x-direction. So this mode can easily be seen from the selected reference location.

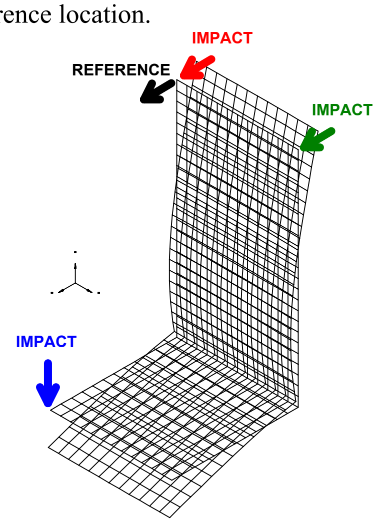

For the second mode of the structure, the reference accelerometer could possibly be located at the lower corner on the structure in the z-direction. If the structure is impacted in the z-direction at the same point (blue), the structure has significant response at this point for this mode. But also notice that if the structure is impacted at the upper corner (red and green) in the x-direction, there is response at the reference accelerometer location in the z-direction. So this reference is good for this mode.

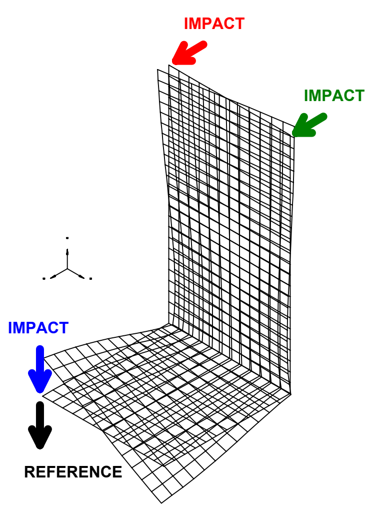

And if the third mode is considered with the reference accelerometer located at the lower corner in the z-direction, all three impact locations on the structure have significant response at this point for this mode.

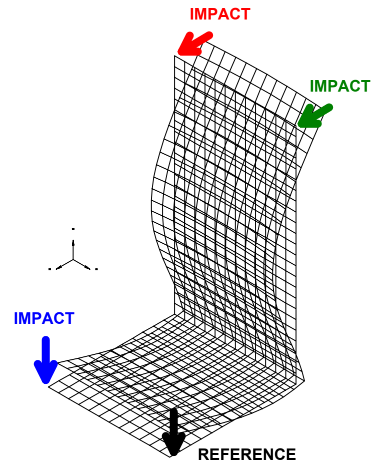

Now the real question is if there is ONE reference location that can be selected that will adequately capture the dynamic characteristics of the structure for all the modes of interest. For this case, it seems reasonable that the lower corner of the structure in the z-direction is sufficient to observe all of the modes for this case.

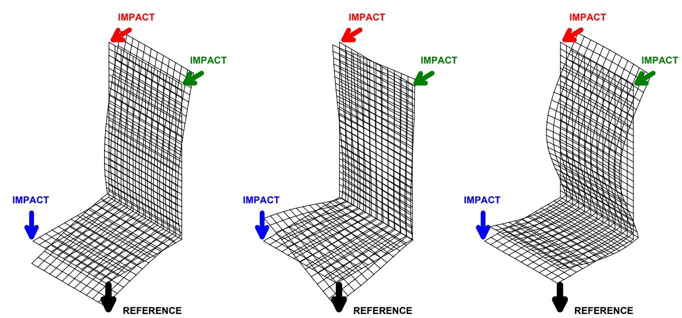

Now we can see that the single reference point is sufficient to adequacy observe all the modes of interest - and only one direction is necessary to accomplish this. Of course, if more references are used this is totally acceptable and is definitely a better way to test the structure - but these extra references are not necessarily needed in order to extract modes shapes that are three dimensional in nature.

I hope that this helps to clear up the misconception regarding the need to measure in three separate reference directions for a modal test. Think about it and if you have any more questions about modal analysis, just ask me.

Copyright 2004
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

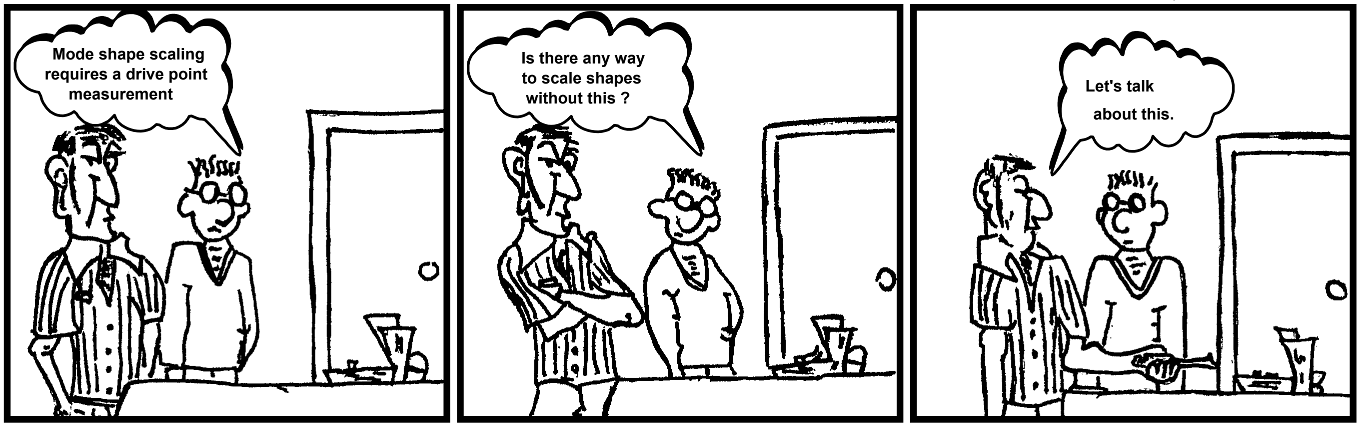

Illustration by Mike Avitabile

Mode shape scaling requires a drive point measurement.
Is there any way to scale shapes without this?
Let's talk about this

Mode shape scaling is an important item for the development of an accurate dynamic model that would be used for other structural dynamic studies. Some of these would be simulation and prediction, modification, and correlation, to name a few. While there may be some instances when scaling may not be critical, I will always recommend that this is done since this may be the only data ever acquired. Generally, a drive point measurement is required to scale mode shapes. However, there are alternate ways to collect measurements and obtain scaled mode shapes without a drive point measurement. Let's discuss this.

Recall that the poles and residues are the values that describe the measured frequency response function and can be written as

$$[H(s)] = \begin{matrix}\text{lower}\\ \text{residuals}\end{matrix} + \sum_{k=i}^{j} \frac{[A_k]}{(s-s_k)} + \frac{[A_k^*]}{(s-s_k^*)} + \begin{matrix}\text{upper}\\ \text{residuals}\end{matrix}$$

Now these residues can be shown to be related to the mode shapes. Without going through all the steps, the resulting relationship for the 'k' mode of the system can be written (with some terms expanded) as

$$[A(s)]_k = q_k \{u_k\}\{u_k\}^T$$

$$\begin{bmatrix} a_{11k} & a_{12k} & a_{13k} & \cdots \\ a_{21k} & a_{22k} & a_{23k} & \cdots \\ a_{31k} & a_{32k} & a_{33k} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} = q_k \begin{bmatrix} u_{1k}u_{1k} & u_{1k}u_{2k} & u_{1k}u_{3k} & \cdots \\ u_{2k}u_{1k} & u_{2k}u_{2k} & u_{2k}u_{3k} & \cdots \\ u_{3k}u_{1k} & u_{3k}u_{2k} & u_{3k}u_{3k} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

Now if we consider the 'r' column of these equations, then the residues can be related to the mode shapes using

$$\begin{Bmatrix} a_{1r} \\ a_{2r} \\ a_{3r} \\ \vdots \\ a_{rr} \\ \vdots \end{Bmatrix} = q\ \ u_r \begin{Bmatrix} u_1 \\ u_2 \\ u_3 \\ \vdots \\ u_r \\ \vdots \end{Bmatrix}$$

So for each measurement, a relationship between the residue and mode shape can be obtained as

$$\begin{aligned} h_{1r} &\Rightarrow a_{1r} = u_1 u_r \\ h_{2r} &\Rightarrow a_{2r} = u_2 u_r \\ h_{3r} &\Rightarrow a_{3r} = u_3 u_r \\ &\vdots \end{aligned}$$

but notice that there are more unknowns than equations and it doesn't matter how many extra equations are added to the list. The shapes cannot be determined unless one particular measurement is included - the drive point measurement which is given as

$$h_{rr} \Rightarrow a_{rr} = u_r u_r$$

With the drive point measurement, then the mode shape at the reference location can be obtained - thereby allowing all the other mode shape coefficients to be determined.

Copyright 2004
Pete Avitabile

But what happens if a drive point measurement is not available or is very difficult to obtain. Is there any other way that the mode shapes can be scaled using other measurements that could possibly be made? Well, it turns out that the answer to this is YES. Let's describe a set of measurements that will enable an equivalent representation of the drive point scaling measurement.

Let's consider some terms in a frequency response matrix at arbitrary locations as shown. The 'r' subscript is the reference and the 'o', 'p', 'q', 's' and 't' are arbitrary measurements in that matrix. Most of the measurements are made relative to the 'r' reference but one measurement is not. We are assuming that the drive point measurement, h_{rr} , has not been measured but is shown in the matrix for illustration purposes. There are three particular measurements of interest that are needed to write some simple equations (these are shown with a double bar underscore in the matrix).

$$\begin{bmatrix} & & & h_{or} & \\ \rightarrow & \underline{\underline{h_{pq}}} & & \underline{\underline{h_{pr}}} & \leftarrow \\ & & & \underline{\underline{h_{qr}}} & \leftarrow \\ & & & \vdots & \\ & & \Rightarrow & h_{rr} & \Leftarrow \\ & & & h_{sr} & \\ & & & h_{tr} & \end{bmatrix}$$

Recall that we can write the residue - mode shape relationship for a particular mode and for a particular measurement as

$$(1) \quad h_{pq} \Rightarrow a_{pq} = u_p u_q$$

$$(2) \quad h_{pr} \Rightarrow a_{pr} = u_p u_r$$

$$(3) \quad h_{qr} \Rightarrow a_{qr} = u_q u_r$$

(Note: For brevity, the scaling coefficient has been dropped) Three specific measurements have been selected here to illustrate the development of an alternate scaling mechanism.

Now, the first equation can be rewritten as

$$u_p = \frac{a_{pq}}{u_q}$$

and substituted into the second equation to give

$$u_r = \frac{a_{pr}}{a_{pq}} u_q$$

The third equation can be rewritten as

$$u_q = \frac{a_{qr}}{u_r}$$

and substituted into the modified second equation to give

$$u_r = \frac{a_{pr}}{a_{pq}} \frac{a_{qr}}{u_r}$$

And then, rearranging terms, gives the drive point equivalent as

$$u_r^{\,2} = \frac{a_{pr}\, a_{qr}}{a_{pq}}$$

I know that I usually don't have this many equations to explain things but this only involved a few simple manipulations to reveal an alternate mechanism to obtain the mode shape coefficient for the reference degree of freedom. Remember that the drive point measurement was not used to obtain the mode shape coefficient for the reference point.

At times this can become a very useful approach especially when there is no access for a drive point measurement or it is inconvenient to obtain the drive point measurement. While I haven't used this often, it does come in handy when performing impact measurements and it is difficult to get the impact device into an area of the structure where access is restricted. It is also useful during shaker testing when it is difficult to make a drive point measurement.

I hope this clarifies your question regarding mode shape scaling and the need for a drive point measurement. If you have any other questions about modal analysis, just ask me.

Copyright 2004
Pete Avitabile

 by Pete Avitabile

Illustration by Mike Avitabile

I find overload and underload of the digitizer and range setting difficult to understand.
Let's discuss this.

This is a very good item to discuss. Previously we had discussed issues in impact testing related to hammer tip, trigger delay and double impacts. There are other issues related to overload/underload of the analog to digital converter, poor utilization of the digitizer, and difficulties with testing nonlinear structures that are additional concerns. In order to try to explain the issue of range settings and their impact (no pun intended) on the analog to digital converter (ADC) setting, several typical measurements discussed previously will be used.

Figure 1 shows an impact measurement where the input force excitation does not adequately excite the entire frequency range of interest. Approximately half of the frequency range does not see sufficient force input to excite the structure and therefore both the input and output signals are very low over this frequency range. But at the lower frequencies, the input force signal is strong as is the response signal.

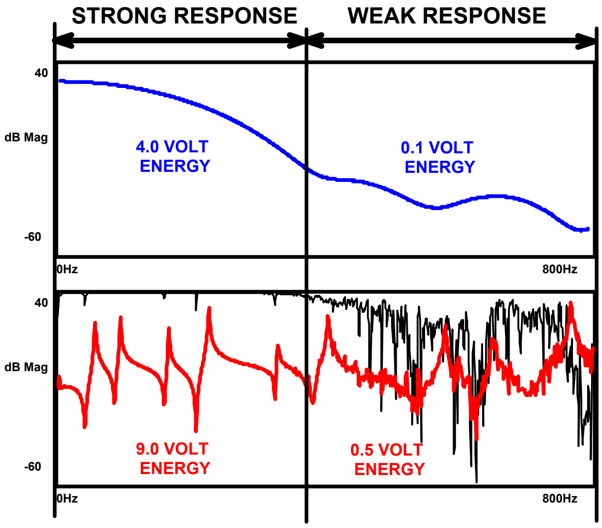

Figure 1 - Hammer Input Only Excites a Portion of Spectrum

The low frequency components dominate the energy of the signal. Now the real issue is how the energy is distributed over the frequency range (which is essentially an assessment of the area of the curve for discussion purposes). For the sake of argument, let's say that the input force spectrum (blue) has approximately 4Volts in the low frequency range and 0.1Volts in the higher frequency range and has a total of 4.1Volts. Let's also say that the response (red) is distributed as 9Volts and 0.5Volts in the low and high frequency ranges, respectively. Clearly, the total voltage is dominated by the low frequency range. In this case, the input channel and output channel may be set to 5 volts and 10 volts for possible ranges on each channel.

So what does this mean. Let's use a simple sine wave to explain resolution. For illustration purposes, a simple 6 bit ADC is set to full range and then set to a lower range to clearly show how the digitizer can affect the amplitude measured. (Please note that all values are approximate and rounded off for illustration purposes). Figure 2 shows a sine wave with 1.5 V peak amplitude and an ADC set to a maximum of 10 Volts. Figure 2 only shows the portion of the ADC which contains the signal. Notice that the resolution is poor and that the actual amplitude of the sine wave is not identified correctly due to quantization error. This would result if the ADC range setting was set much larger than the actual signal to be measured (in this case the full range of the ADC is 10V).

Now if the ADC range is set to 2.0 Volts as shown in Figure 3, the resolution of the signal is much better. This is because all of the dynamic range of the ADC is dedicated to the signal of interest (the ADC is set to 2.0 Volts to measure the 1.5 Volt signal).

Copyright 2004
Pete Avitabile

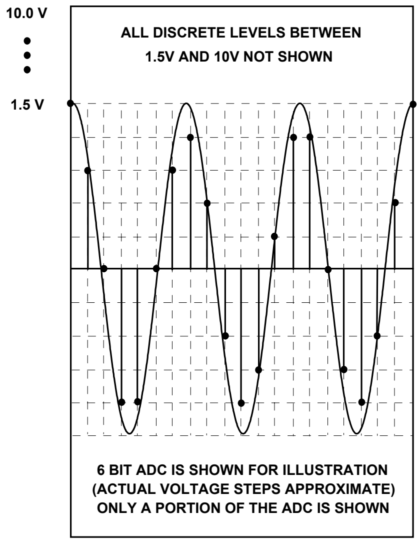

Figure 2 - Sine Wave with Poor Resolution

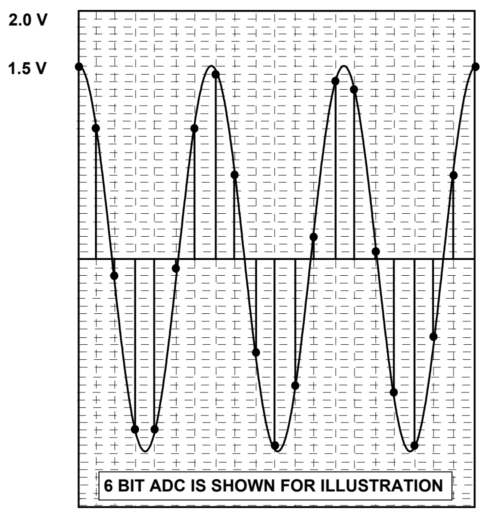

Figure 3 - Sine Wave with Good Resolution

Now let's consider one more case where there are two sine waves at different frequencies and different amplitudes to be measured simultaneously. The same 2.0 Volt range is used. In Figure 4, it is very clear that the larger of the two sine waves dominates the ADC setting. However, it is very important to note that the smaller of the two signals will suffer from quantization error more than the larger signal. This is very common in frequency measurements that are made on structural systems. There is no way to avoid it. But imagine how much worse the quantization error would be if the ADC were set to a 10 Volt max range.

Now that we have some idea about range setting for a simple sine wave, we can better understand the problem with the measurement in Figure 1. The higher frequencies are not well excited and there is little response of the higher modes. The measurement at the higher frequencies suffers from quantization errors. This problem in Figure 1 at the higher frequencies is analogous to the problems cited in Figure 4.

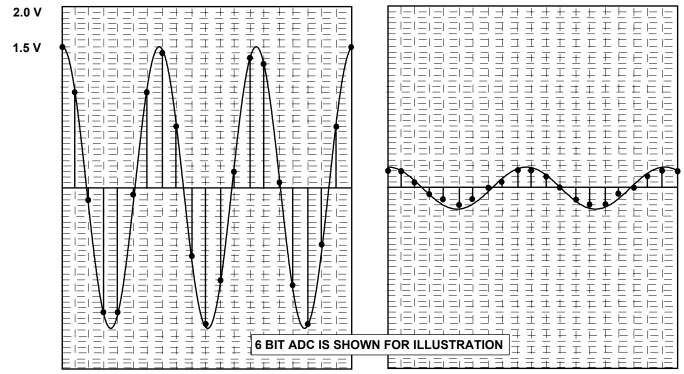

Figure 4 - Two Sine Waves with Possible Resolution Problems

Now let's consider one more case where the impact to the structure excites modes well beyond the frequency range of interest (128 Hz) as seen in Figure 5. For the sake of argument, an assumed energy distribution between the desired (lower) frequency range and the higher (excited but outside the range of interest) frequency range. Notice that the transducers will measure the entire response (energy) of the system even though only a portion of that energy is actually used in the digitization for the frequency information. What happens is that the ADC must be set much higher than actually needed since the total voltage from the transducers is actually heavily affected by the higher frequencies. This implies that the ADC will be set much higher than needed to accommodate these higher frequencies - the result is that the lower frequencies will suffer from quantization errors.

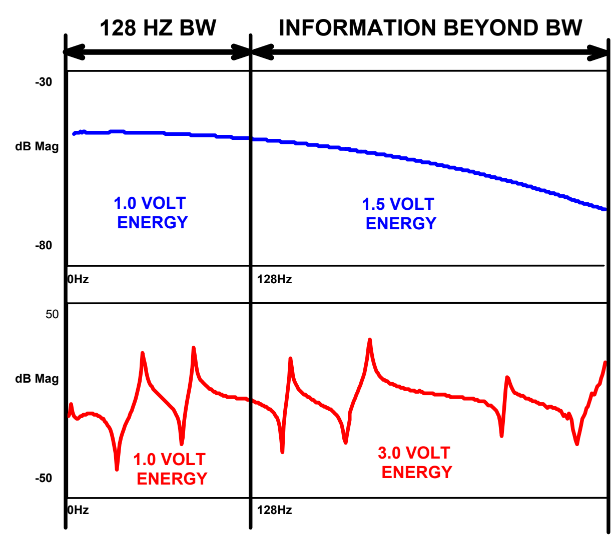

Figure 5 - Exciting Modes Outside the Band of Interest

I have tried to answer your question about digitizer settings in a round-about way using a typical measurement that might be collected for a structural system. I hope this helps to explain this important part of taking measurements. If you have any other questions about modal analysis, just ask me.

Copyright 2004
Pete Avitabile

Illustration by Mike Avitabile

Do mode shapes of a plate have any particular predetermined order?
Let's discuss this.

This is a question that comes up more often than you would ever imagine. This needs to be discussed and some examples will be given to explain and clarify this often confused item.

Many times I have heard people say that modes of a structure always first start with torsion. While it may be true for their particular application that the modes that are typically seen in their structural configuration may start with torsion, there is no predetermined rule as to the order or sequencing of modes.

For instance, many people often think that modes of a plate must start with torsion – but there is no mathematical reason for this to occur. It just may have been that all the plate structures that they have seen in the past have had a first mode that was a torsional mode. Of course once someone sees a torsional mode as the first mode of a structure several times, then forevermore, *all plate like structure must have a first mode which is torsion*. ***This is not true at all***.

I remember an instance many years ago when an analysis was performed on a new stiffened body-in-white configuration. The structural engineers had spent a significant amount of time concentrating on designing a structure that had a significant increase in the first flexible modes of the car frame. Prior to that time, the modes of a configuration of this type *always* started first with torsion (T) and then followed by bending (B). When the first analytical models were analyzed, the first mode of the car frame turned out to be bending rather than torsion. There was incredible confusion concerning this since up until this time the first mode was *always* torsion – almost as if it was etched in stone (as the 11th commandment!)

No one believed the model since this appeared to be totally contradictory to what everyone believed to be the "*way things were meant to be*". But there is really no basis for things to be that way. The model is a distribution of mass and stiffness that results in an eigenproblem that yields frequencies and mode shapes which satisfy the force balance equation. If the model is prepared correctly then the solution will identify the frequencies and mode sequence that satisfies the mathematical problem. (Of course, there may have been errors in the model but that's a totally different story.)

The simple fact is that the frequency and mode shape sequence is due to the mass and stiffness distribution of the structural configuration and not due to anything else.

In order to illustrate the mode sequence arrangement possibilities, three different plate configurations with different length to width aspect ratios were generated and finite element solutions were obtained for each. These are shown in the figure with the arrangement from lowest to highest mode from top to bottom. The modes are further indicated with a B for bending along the longer length of the plate, B2 for bending along the shorter length of the plate, and T for torsion about the symmetry axis. For the three different plates analyzed, there is no specific ordering of the mode shape sequence. Each of the plates has a different combination as seen in the figure.

And as long as we are on the subject of mode shapes, the question to ask is if the bending along the longer length of the plate (B) will always occur at a lower frequency than the bending along the shorter length of the plate (B2)? Now before you too quickly just answer that question, stop and think about it for a minute……..

Is this a trick question? What do I need to think about before I answer that question? What are the material properties? And are they the same along the long and short length?

Copyright 2004
Pete Avitabile

If the material is homogeneous, then the bending mode along the long length (B) will occur before the bending mode along the short length (B2). But what if the material is a reinforced carbon fiber composite where the stiffening fibers run along the longer length of the plate? Then there is a possibility that the plate will be much stiffer along that length. Then it is also possible that the frequency of the bending along the shorter length (B2) may occur before the frequency along the longer length (B).

(Oh my, Toto – I am not sure we are in Kansas anymore!). Obviously, the bottom line here is that you really need to think about this possibility. It is a very possible reality!

I have tried to answer your question about mode order for a plate. Obvoiusly, this holds true for any structural configuration that has characteristic bending and torsion modes – not just a plate configuration. If you have any other questions about modal analysis, just ask me.

B-T-T-B-B2

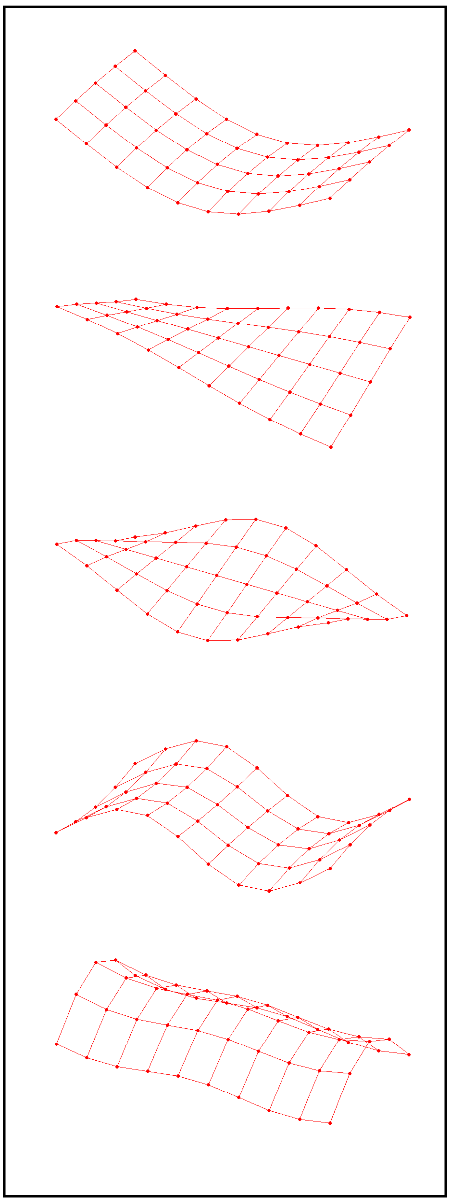

B-T-B-T-B

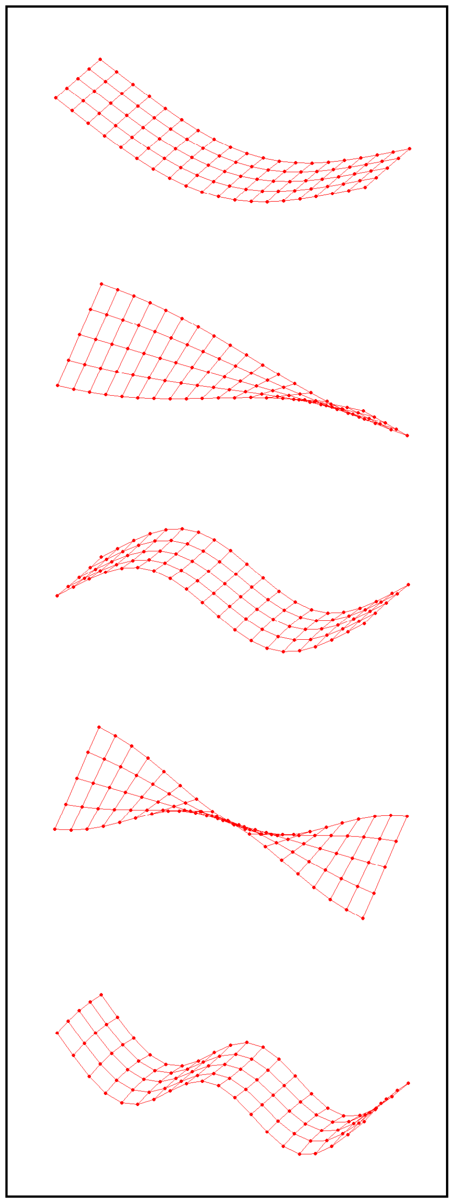

T-B-B2-T-T

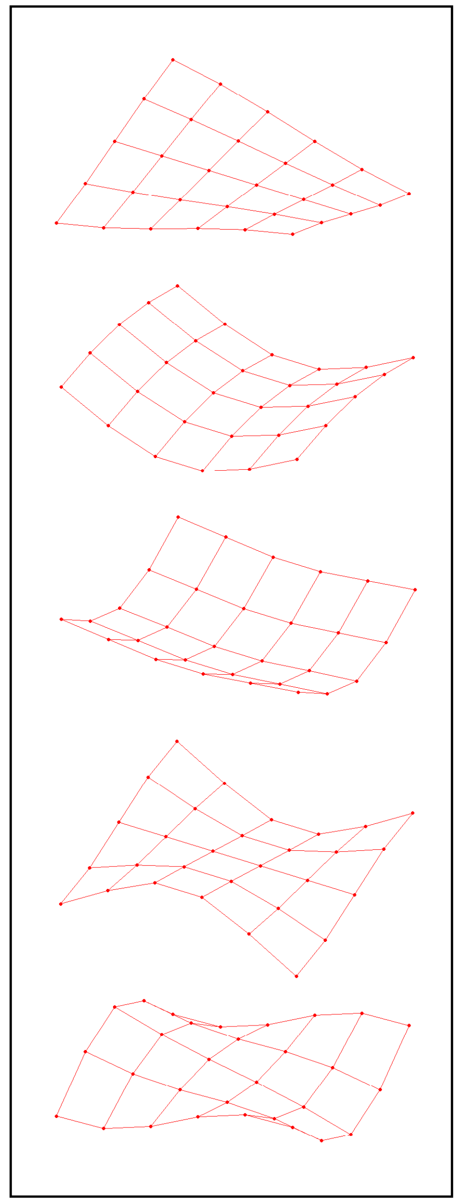

Copyright 2004
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Once I have set up a good measurement, is there any reason to watch the time and frequency results for every FRF? Let's discuss this.

This question brings up a very important topic. Once a test has been set up and care has been taken to make sure that the measurement is good, you should always monitor all of the measurements made on the system. This is a must!

I have seen some people take a great deal of care measuring the impact force and response, measure the input spectrum, FRF and coherence for just one point and then just disregard monitoring all the points for the rest of the test. The general feeling is that once the drive point measurement is made, the force spectrum is checked and the coherence is acceptable, then the test should proceed without any major difficulty. The problem is that just because one point seems to be very good, doesn't necessarily mean that all of the points will be measured the same. I have seen many tests where measurements made on various parts of the structure have had very different measured characteristics than might be expected.

So let's start with a typical measurement scenario and identify what could possibly go wrong if attention is not given to every measurement made during the test. A bracket that was in the lab was used to acquire some measurements. Obviously, the time and frequency data should be reviewed from some measurement points. Typically, the drive point measurement may be acquired as a starting point. For the structure under consideration, an impact excitation was used and the time input force and time response are shown in Figure 1. The response seems to almost decay to zero by the end of the sample interval so possibly a light exponential damping window could be used to minimize the slight amount of leakage that might occur.

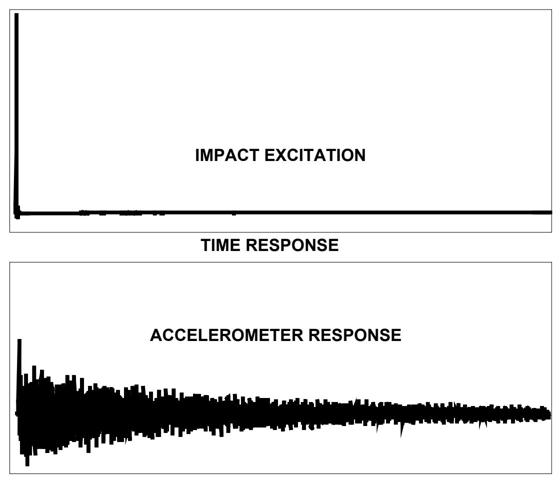

Figure 1: Force Excitation and Accelerometer Response

Next the input force spectrum is checked to make sure that a sufficient amount of force is applied to the system over the frequency range of interest; usually this input spectrum should be reasonably flat over the spectrum with approximately 10 to 15 to possibly 30 dB roll-off over the desired frequency range. (Notice that I said the "desired" frequency range which may not be the entire spectrum measured.) The coherence is checked to make sure that there is reasonably good causal relationship between the measured input force and the output response to assure that a good measurement is made. And of course, the FRF is checked for peaks in the measurement indicating modes of the system. These are shown in Figure 2 and this measurement looks very good. In addition to the magnitude of the FRF, it is a very good idea to also check the complex parts of the FRF.

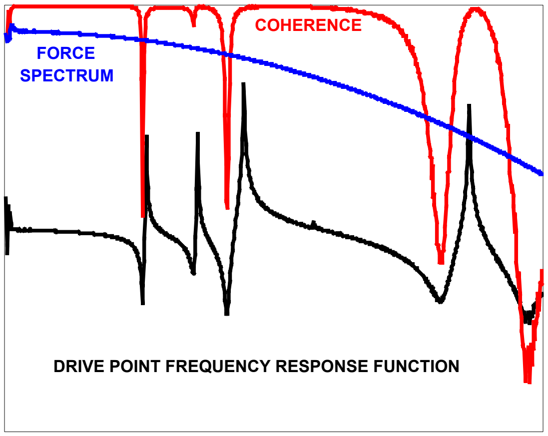

Figure 2: FRF, Input Force Spectrum, Coherence

The real and imaginary parts of the FRF should be inspected to make sure that the measurement looks as expected. Figure 3 shows a good representation of this.

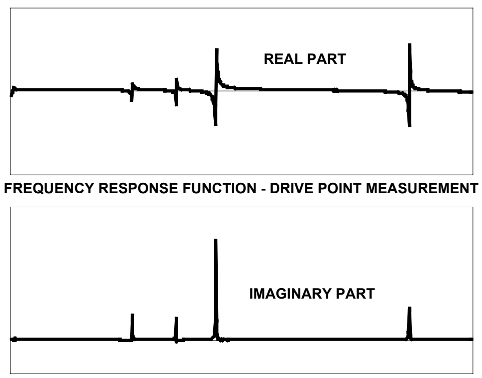

Figure 3: Real/Imaginary FRF – Drive Point Measurement

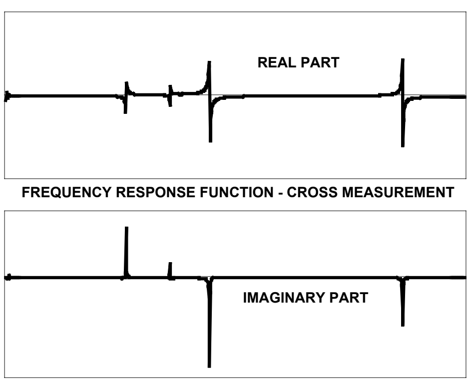

Figure 4: Real/Imaginary FRF – Cross Measurement

However, a note of caution. Many times people only look at the drive point measurement at the start of a test. While this is a critical measurement of the system especially for mode shape scaling considerations, it is not the best measurement to look at all the time. For instance, the peaks of the imaginary part of the FRF will always have the same phase relationship. But if two modes are very close to each other than it is sometimes very difficult to determine how many modes really exist in the data. Many times it is better to check one of the cross measurements as shown in Figure 4. Notice that all of the peaks in the imaginary part of the FRF do NOT have the same phase relationship. This is very useful in determining peaks of closely spaced modes and should always be taken during preliminary testing setup. So once this is done, is there any real need to continuously monitor the time and frequency data for all of the measurement locations. The measurement shown in Figure 5 will help to show why constant monitoring is needed.

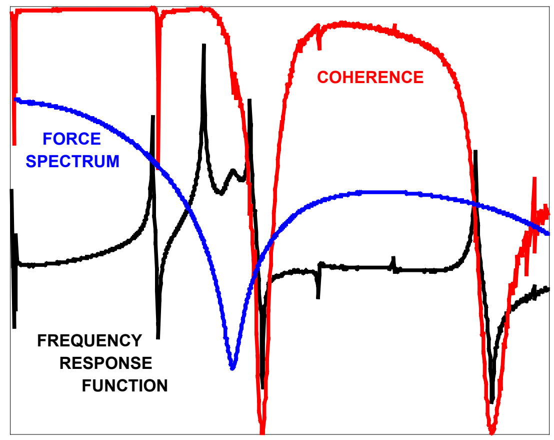

Figure 5: FRF, Input Force Spectrum, Coherence

If just the FRF and coherence were measured, then part of the picture is lost. With just the FRF and coherence, it may be seen that the measurement was poor and possibly blamed on nonlinearities, noise, complex damping and a host of other well-known problems that can affect data. But the real culprit for this measurement is none of those. The input spectrum which was reasonably flat for all the other measurements has a much different force spectrum than in earlier measurements. For this particular structure, impacting at certain locations, there is a dramatic change in the local compliance in the structure and it is very difficult to maintain a fairly uniform input spectrum. It is just a "quirk" of this structure but can happen on any structure. So if you aren't going to check every measurement during the test, make sure that each measurement is saved for every point measured – and that includes all the parts of the measurement not just the FRF and coherence. Because you can see from this case, the input spectrum had important information that was critical to interpreting the measurement in Figure 5.

I hope I have shed a little more light on different aspects of running a test and possible items that need to be addressed. If you have any other questions about modal analysis, just ask me.

Copyright 2004
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

I know that certain shaker excitations have different characteristics – but which is the best to use?
Let's discuss this.

Let's discuss the most commonly used excitation techniques for modal analysis today. These are random, pseudo random, burst random, sine chirp and digital stepped sine. These will just be briefly reviewed (because they have been covered before), but the more important issue is do they all provide the same results all the time. There is not always a straightforward answer to this so we will discuss some issues to consider when performing shaker testing. The main excitation techniques utilized are shown in Figure 1 for reference.

Random is still used at times today, even though leakage and windows cause some distortion of the signals acquired. The time signal is shown in different colors. (This is mainly done to highlight the fact that each measured time signal is different from one record to the next.) Since each signal sample is different than every other signal, the system is excited with different spectral characteristic for each record of data collected. If the system has slight non-linearities, then the system will respond differently for each record of data – the averaged data will then reflect the best linear description of the system in the presence of these slight nonlinearities. This type of excitation is very useful to minimize or smooth data that is subjected to noise or rattles or other measurement contaminants. But leakage and windows tend to distort the measurement so this is not the optimal excitation technique.

Pseudo random is really nothing more than a set frequency spectral lines over a frequency band of interest that are inverse transformed to the time domain to create an excitation signal. Since the excitation is basically sinusoidal in nature, the effects of leakage are non-existent providing the system is excited sufficiently long enough that steady state response is achieved. This proves to be a very useful excitation. However, because the signal is repetitious (notice that the excitation color is identical from one record to the next), the system will respond in a deterministic fashion. This will not average any slight nonlinearities or rattles that may exist in the system. Pseudo random works very well with structures that are fairly linear in character.

Burst random excitation was developed as a very good excitation technique which combines the advantages of both random and pseudo-random excitation. The signal is random from one record to the next (notice the different colors from one record to the next) and nonlinearities are averaged in the process. Since the signal is completely observed in one sample interval, there is no leakage and windows are not needed. The only concern is to make sure that both the input AND output response are totally observed within the sample record of data.

Sine chirp excitation is a fast swept sine that is completely observed within the sample interval. The effects of leakage are non-existent providing that steady state response is achieved. This excitation is very similar in advantages/disadvantages to pseudo random. One additional advantage is that the level of force can be controlled and can be used to identify the nonlinear character of the system.

Digital stepped sine excitation is yet another very useful excitation technique. This is very similar to pseudo random except that only one frequency is excited at a time. But one significant difference lies in the improved description of the signal amplitude. Broadband techniques (those discussed above) require that the analog to digital converter (ADC) be set to capture all the energy over the entire spectrum. But the frequency character may have a wide variation in amplitude over the frequency spectrum. This is not an issue with digital stepped sine since all of the energy of the excitation/response is dedicated to one single spectral line in the frequency spectrum. Therefore, quantization error is not an issue for this excitation.

Copyright 2005
Pete Avitabile

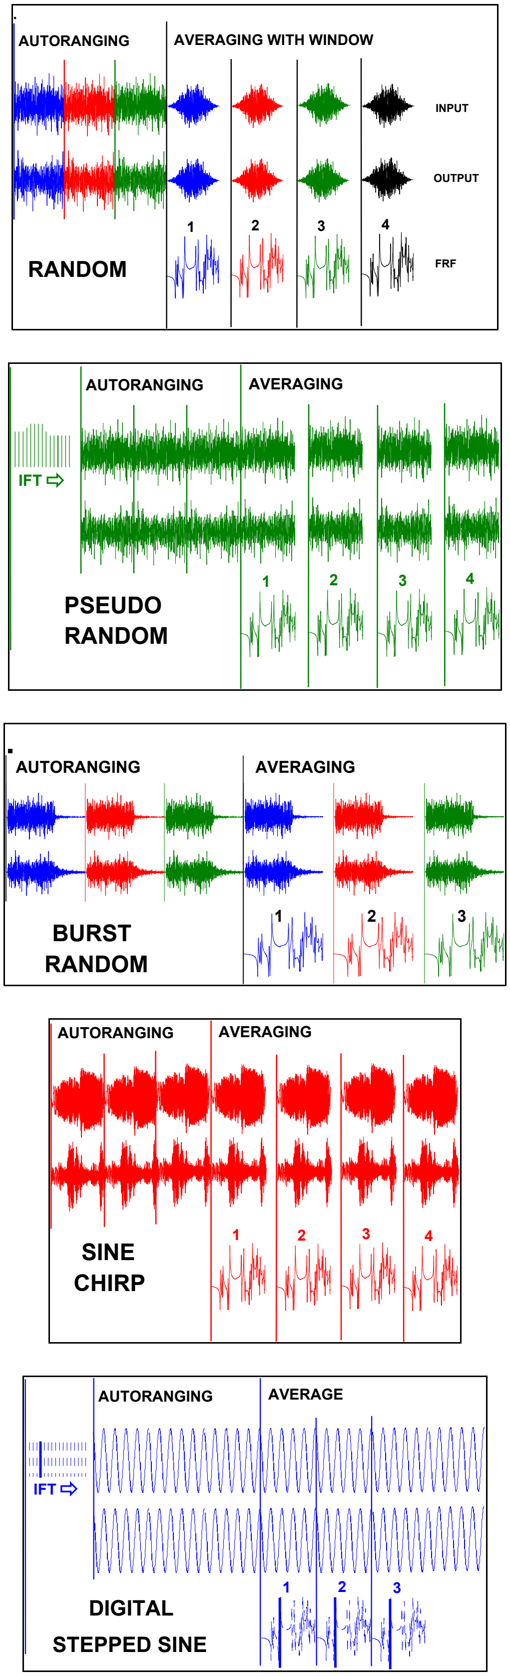

Figure 1 - Typical Shaker Excitations Employed

So it seems pretty straightforward what to do. Well, the reality is that things aren't always so simple. So now that we have categorized all the excitation techniques, let's talk about what some of the issues are that might arise. In general, over the past several decades I have generally found that burst random usually works best overall. But I have also used sine chirp on many occasions when the structures were fairly linear in nature. When I have needed extremely high resolution FRFs, then I have used digital stepped sine. Once or twice, I have used either pseudo-random and random. So let me explain some of the times I have used other signals and explain why.

One structure I tested many years ago was a very lightly damped system. It turned out that burst random was not very effective. The system was so lightly damped that the response could not be totally observed within one sample interval of the time record – even with the burst set to less than 5% of the time window. Fortunately, this system was fairly linear so that pseudo-random excitation was employed (but sine chirp could have also been effectively employed for this structure).

When a structure does have some nonlinear character, then it may be desirable to perform the test at a level that is comparable to the in-situ conditions. Sine chirp proves to be a very good excitation for this type of test. So why not use digital stepped sine – well for this particular test there were not a sufficient number of acquisition channels available to make the test feasible.

I guess the best thing to realize is that there is always going to be some situation which may make one of the available excitation techniques provide a better measurement than the other excitation techniques. Each of the techniques needs to be compared to determine which is the best. Don't just rely on one technique because it has proven to be acceptable in the past.

But today, with the large channel count systems that are more common for modal testing, my recommendation would be to utilize all of the excitation techniques. Today there is sufficient disk storage that this shouldn't be a concern. Since it takes a good deal of time to set up a large number of accelerometers on the structure, why not run all the different excitation techniques – even digital stepped sine which takes much more time than the broadband based techniques. If you have spent 3 or 4 days setting up a large test, do you think anyone will care if you take a few hours and collect all the data possible? I don't think there will be objections. At least then you have all the data.

I hope that I have answered your question regarding the different excitation techniques. If you have any more questions on modal analysis, just ask me.

Copyright 2005
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

When I perform impact testing, the input spectrum looks distorted – do you think my FFT analyzer has a problem? Let's take a measurement and see what's going on.

The only way to figure out what is going on is to do a little trouble shooting. But of course we have to have some idea of what the correct answer should be. Well … for an impact test, the force in the time domain should be just a simple pulse over a very short duration. And the resulting frequency spectrum should be a relatively flat input profile over some frequency range. The width of the input spectrum is directly related to the time of the pulse. Basically, the shorter the pulse in the time domain, the wider the frequency range that is excited. And conversely, the wider the pulse in the time domain, the narrower the frequency range that is excited. At least this is what we expect to get. So now let's go down to the lab and take a few measurements with your analyzer to see if this is what we get. (In this article, I will try to replicate what was actually observed during a test at a particular lab. But I cannot reproduce exactly what occurred due to FFT hardware differences.) For this test, a typical impact hammer is being used and the structure under test is just a typical structure with an accelerometer mounted on it.

The impact force is shown in Figure 1 with the time response (upper two traces in black). Now the time signal looks reasonable. The impact signal is a sharp pulse with relatively flat zero signal over the entire time. In terms of the response signal, the damped response of the system appears to be reasonable but a window may be necessary to minimize leakage.

The input spectrum (shown in blue in the lower trace and labeled as distorted force spectrum) is also fairly flat with little roll-off to the input spectrum as expected. But one thing to notice is that there is a significant spike in the spectrum at very low frequency which does not look correct. This is definitely not expected and is probably the reason you have asked a question regarding this measurement. This needs to be investigated further.

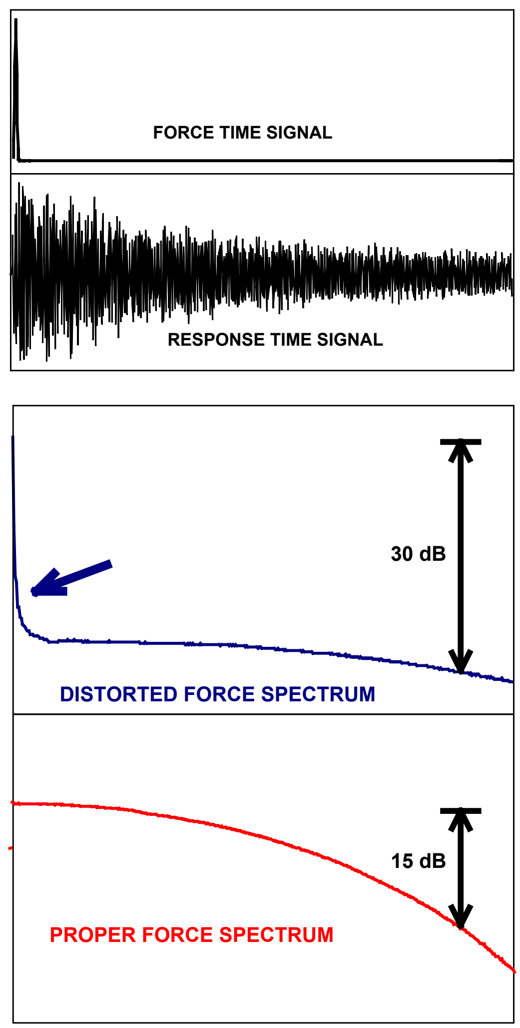

Figure 1 – Impact Time Force, Exponential Accelerometer Response and Force Spectrum (Proper and Distorted)

Now let's return to the time signals and do a little investigating on the measured signals. The first thing to check is if the input transducers have a suitable signal and whether the electronics may be a problem (ie, check for transducer saturation along with checking cables, batteries, etc). Usually the FFT analyzer can be setup into some type of time sampling mode to view the time signals as would be done on an oscilloscope. As long as there appears to be a suitable signal, then you can proceed to the FFT representation of the signal.

The time signal looked adequate … but a question needs to be asked if any windows have been applied to the measured data. Based on the analyzer window settings for this particular analyzer, there *does not appear* to be any windows applied to the force input signal. But the response signal may need a window. However, many times the labels for the force and response windows are confusing on many FFT analyzers. Most times the labels make sense but many times the analyzer software uses labels that may not clearly identify the window being used.

In this particular case, the input signal was labeled as not needing any window – the label on the force channel indicated something similar to "**response only window**". Now that label is confusing to me. The label might mean that "only a window is applied to the response channel" – or does it mean something else? Of course, we should open up the user's manual and read the paragraphs related to this particular window configuration. But many times, a typical user may think he knows what all the buttons/labels mean since he is familiar with other FFT analyzer equipment used in the past and not think that reading the user manual is really necessary. On the other hand, sometimes the user manuals are just as confusing. Often times I am more confused by the user manual information than when I started.

So, typically, I resort to a very simple "hunt and peck" approach. That is, select one parameter to be changed while others remain constant. In the case of the impact hammer, the first check is to make sure that the force gage is not damaged – swapping another force hammer is the simplest thing to check along with cabling and signal conditioners.

The next check is the FFT acquisition system setup. The analyzer window for the force channel was set to "**response only window".** The response window can also be changed – but our first impression is that this has nothing to do with the input force spectrum. Originally, the response window had been setup as a Hanning window which is obviously incorrect – this may have been set as a default setting for some other tests such as processing random type signals. The window should either be no window or the exponential window. As a first check, the rectangular window was selected (and should be done first even if an exponential window is ultimately required). Now the response signal appears to be the same as previously observed.

Now this is not possible since the Hanning window had been applied previously. Well, it turns out that the original signals were observed as the "un-windowed" signals which is the default setting for many analyzers – specific setting have to be selected in order to view the windowed signal on certain analyzers. While the response signal is obviously affected by the application of the window, the interesting observation is that the ***force spectrum on the input channel also appears to be affected by the window on the response channel!*** The force spectrum is shown in Figure 1 in the lower trace (shown in red and labeled as proper force spectrum). How could the force spectrum on the input channel change when the response window is changed?

Is the analyzer broken? Should you report this flaw to the manufacturer? Well – actually no! On many FFT analyzers, the response window is actually applied to the response signal as well as the force signal in order to properly compute the ratio of the output to input (this is commonly done to avoid the distortion that results from applying the window to just the output response – there are theoretical issues that justify this but we don't need to discuss those here). The important item here is that the response window is applied to both the input and output channels on many (but not all) FFT analyzers. So now we can see why that original force spectrum was distorted – now it is clear what "**response only window**" meant. The original FFT analyzer was setup with a default setting using the Hanning window – on the response channel. Even though the only channel being evaluated was the force channel, the window effect on the response channel was equally important.

You need to be extremely careful when performing any testing using FFT analyzers. You need to be very sure you understand how the analyzer operates – many analyzers have many different features and not all of them are the same. To be sure, make some simple checks like were done here to assure yourself that the measurements are acquired in a proper sense.

I hope that I have answered your question regarding this impact measurement problem. If you have any more questions on modal analysis, just ask me.

(Author Note: The events described in this paper have been replicated using different analyzer hardware to show what was essentially observed during this interesting measurement. It is not a statement concerning the hardware used – but rather a caution that test engineers need to carefully understand the functionality of their particular FFT analyzer hardware.)

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Why does my stability diagram of a component on a system show modes that the Sum or MIF don't show?
Let's look at some measurements and see what's going on.

There must be more to this problem than what was stated by your question. My guess is that you are performing a modal test on a structure but you are not measuring all the significant modally active portions of the system during the modal test.

Now what do I mean by that. Well, there are many times when a modal test is performed and there is only an interest in a portion of the structure and nobody wants you to spend any time testing more than what you actually need to do or are being funded to do. This happens all the time in real lab environments. Let's say for instance that you are trying to solve a vibration problem on the floor board of an automotive structure. Now your first thought might be that you don't want to make measurements on the exhaust system of the car – since you are only interested in the floor board.

Of course when you only make measurements on the floor board the rest of the automotive structure is not divorced from the measurements made on the floor board. That means that the measurements made see the response of the entire system. Now granted that most of the measurements on the floor board will be primarily due to the response of the floor board. But there will also be effects of other portions of the system that will be observed in the measurements. Their response may not be strong but it will be there. So a measurement on the floor board will also have the effects of other portions of the structure such as the exhaust system, seating system, etc. It is impossible to completely separate out the response of these other systems. Unless of course the floor board is cut out of the structure and tested separately. But then the floor board modes do not have the same boundary conditions as assembled in the system so the modes of this separate test may not provide the required insight into the response of the system.

This is a problem that can be seen in many experimental modal tests that are conducted on just about any type of structure. It could be floor boards in an automotive structure. Or maybe fuselage modes of a shell like structure when the only area of interest is the wing modes for flutter studies. Or it could be ….. well I just don't have enough room to list all the possible scenarios. But rest assured, it is a prevalent problem in just about every structure that could be subjected to modal testing.

So in order to illustrate what could happen, I went down to the lab and used one of my existing structures that has lots of local modes, a few global modes and some nonlinear behavior due to joint problems – just a typical structure that I often use to illustrate these types of problems you have described. The structure is shown in Figure 1 and is setup for some shaker testing. This structure consists of a very stiff outer frame and a very flexible panel structure that is held in with a clip arrangement. Notice that the shakers are set up for testing only the outer frame of the structure and in the initial test, accelerometers are only located on this outer structure; initially there are no measurements on the panel structure since it is not of immediate concern (or so it is assumed).

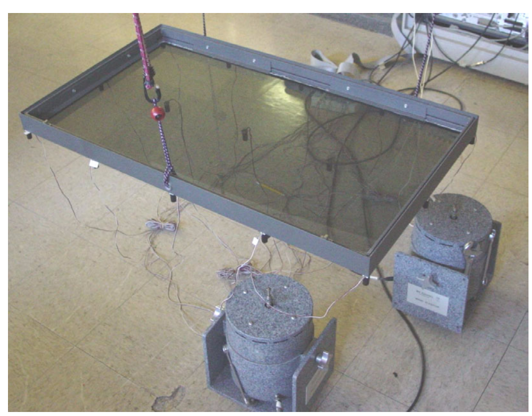

Figure 1 – Ribstiffened Panel Structure

Now Figure 2 contains a drive point measurement from one of the shaker reference locations. Notice that there are three well defined peaks and some other characteristics that are not well defined.

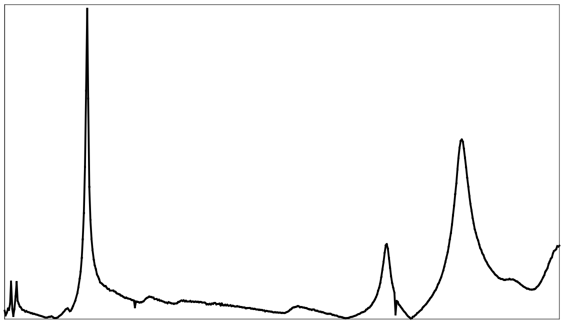

Figure 2 – Drive Point FRF on Frame Structure

Now using just the measurements on the frame structure, a stability diagram is developed as shown in Figure 3. Notice that there are many more than three fairly well stabilized poles in that plot. The SUM function and MIF function show the three peaks very well but the balance of the peaks are not shown very clearly at all. So this stabilization diagram appears to be identifying many more modes than what appear to be interpreted from the SUM and MIF functions.

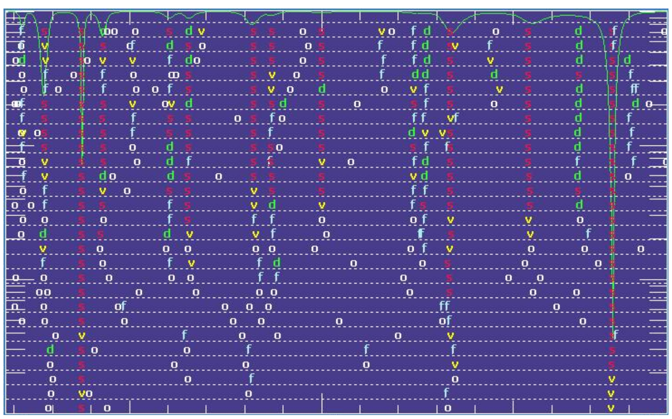

Figure 3 – Stability Diagram and MIF from Frame FRFs Alone

Now the problem is that the structure has many more modes that what can be easily seen on the frame portion of the structure. That panel has many modes that have very little contribution to the response of the frame portion of the structure but their effects can be seen in the measurements taken only on the frame. That is to say that the poles of the system can be seen in the stability diagram even though the SUM and MIF do not show those peaks very well at all.

Now let's take a set of measurements that includes the panel portion of the structure. A drive point FRF on the panel is shown in Figure 4. Notice that there are many more peaks in this measurement than seen in the previous drive point FRF in Figure 2. (Note: these measurements have significant effects of joint slop and have nonlinear characteristics but this structure is very good for illustrating the local modes effects of concern in this discussion).

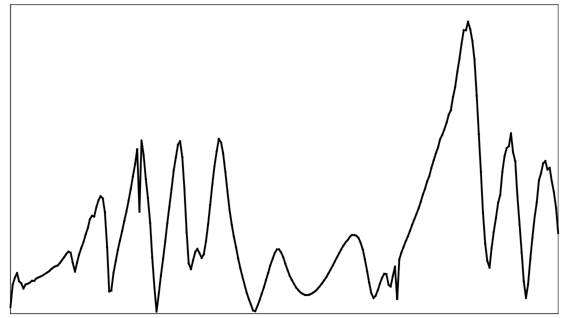

Figure 4 – Drive Point FRF on Panel Structure

Now using all the measurements on the outer fame as well as the panel, the stability diagram in conjunction with the SUM and MIF function shown in Figure 5 appear to present a much clearer picture of all the so called extra modes from the stability diagram of Figure 3 shown previously.

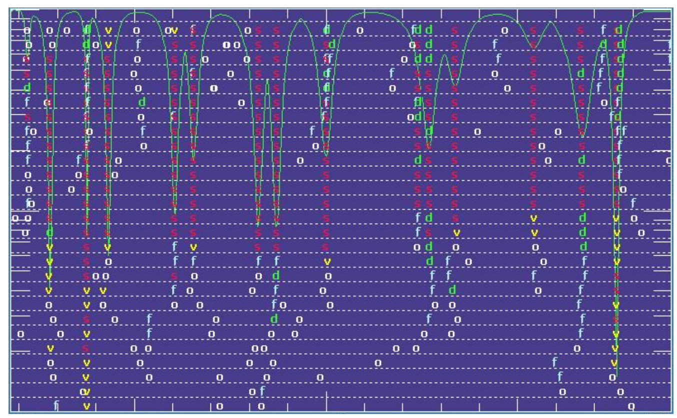

Figure 5 – Stability Diagram and MIF using all FRFs

So now it becomes much clearer as to how many modes are in the band and why the previous stability diagram did not provide useful information. In order to use the tools to interpret the measured data, it is imperative that a sufficient set of FRFs be acquired to adequately describe all the dynamics of the structure – not just the portion of the structure of immediate concern. I see this situation occur very often in industry. A problem arises regarding a portion of a system (or contractually you are only obligated to a portion of the system) and measurements are taken on just that portion of the system. The data obtained only contains a piece of the puzzle. The use of the mode identification tools can become confusing if only a subset of data is acquired. Many times more than just your "region of interest" may need to be measured in order to understand the complete dynamics of the system.

I hope that I have answered your question regarding this interpretation of the stability diagram. If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Sometimes my impact force is very smooth just as expected but often it looks like it is oscillating – Why is that?
Let's look at some measurements and talk about this.

Many times when performing impact testing, the force pulse appears to very regular shaped with a pulse that resembles a half-sine wave. The event starts at zero, followed by the pulse and returns to zero for the duration of the measurement event.

However, many times the force pulse seems to oscillate about zero after the initial half-sine pulse. The real question is why does this happen, should it occur, is it possibly a double impact and should a window be used to minimize the effect of this.

Well … there is a lot to answer here and I may not be able to cover it all in this one article. This problem is referred to as "filter ring". Let's start out with some simple measurements to show this problem that is often seen. Just by taking a few sample measurements, the effect can be observed and hopefully better understood with some simple examples and illustrations.

This is a problem that can be seen on many FFT analyzers. For the measurements and discussion here, I am going to use a general BRAND XYZ FFT analyzer. A typical measurement will be made on a typical structure with a impact force hammer and response accelerometer. However, only the force input will be discussed here. Some of the force pulses will be very regular shaped just as we would expect to see in a textbook case. But other measurements will have a force pulse that has an oscillation to the end of the time pulse as if it is the response of a simple single degree of freedom system. This problem is often refered to as "filter ring". It is due to the fact that the analog anti-aliasing filters on the front end of the analog to digital converter (ADC) may show some response due to their own natural frequencies that are possibly excited due to the force pulse. This is actually what occurs. The force pulse will excite different frequency ranges depending on the tip that is used to excite the structure as is well understood by everyone.

But here is the problem. Depending on what frequency range (bandwidth) is selected, this filter ring may or may not be noticeable on your analyzer. Now on the surface this doesn't sem reasonable until the actual inside working of the FFT analyzer is considered. Usually, the FFT manufacturers have different sets of anti-aliasing filters – one for low frequency work and one for high frequency work. Typically, if you are measuring lower frequency ranges, the lower frequency filter is employed. If a soft impact tip is used then this will not significantly cause any filter ring. But if a slightly harder tip is used, then the upper frequency range of the hammer excitation may excite the low frequency analog anti-aliasing filter. The filter gets excited and has a dynamic response characteristic which manifests itself on the force pulse as this filter ring.

So let's take some measurements to illustrate this filter ring characteristic and see how setting different frequency bandwidths may have an effect on the filter ring observed. An impact hammer will be used with four different tips over two different frequency ranges. The hammer tips will consist of a very soft red air capsule, a medium blue plastic tip, a harder white plastic tip and a metal tip. In each case, the hammer is used to impact a structure to acquire a time trace. In one set of measurements, the frequency bandwidth is set at 400 Hz and in the second set of measurements the bandwidth is set to 1600 Hz. The two figures on the next page show the results of the different impacts over the two frequency bands. The tips range from softest to hardest from top to bottom.

Notice that the 400 Hz bandwidth has significantly more filter ring as the hammer tips go from softer to harder. This is because the harder tip excites a wider frequency range and has a great possibility of exciting the low frequency analog anti-aliasing filter.

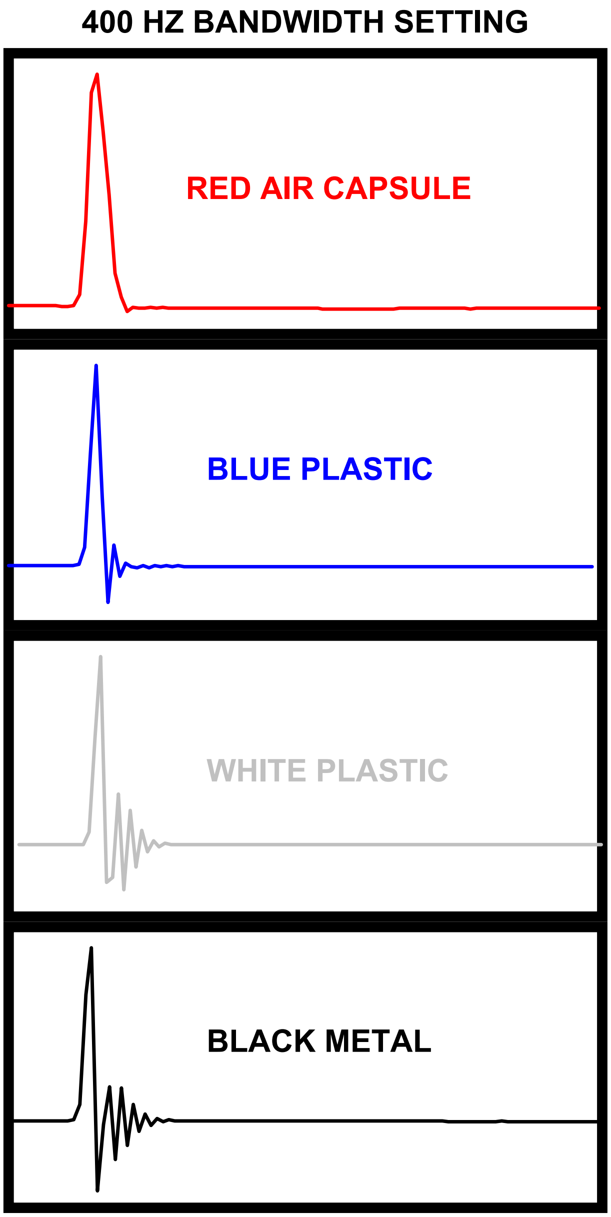

Figure 1 – Impact with 400 Hz Bandwidth

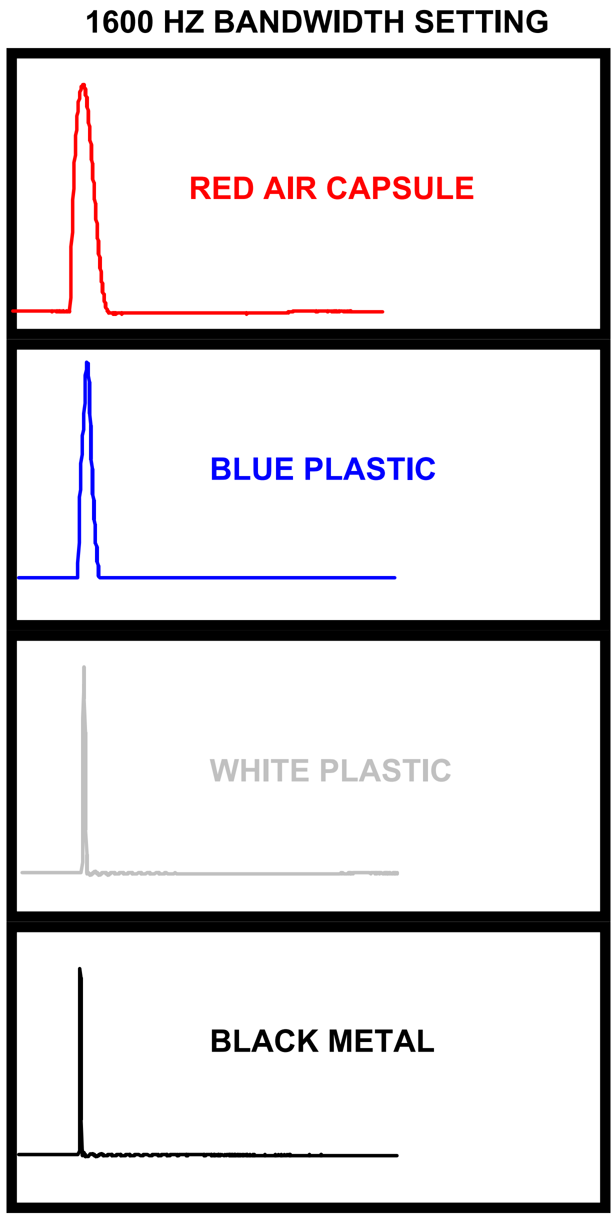

Figure 2 – Impact with 1600 Hz Bandwidth

Comparing the 400 Hz bandwidth to the 1600 Hz bandwidth, there is a noticeable change in the filter ring – there is hardly any ring at all for the 1600 Hz bandwidth. And the only difference was the selection of the bandwidth.

On this particular FFT analyzer, the two sets of anti-aliasing filters are used depending on which bandwidth is selected. Clearly, the filter ring is much more obvious when the harder tip is used over the lower frequency range. This is because the harder tip has significantly more energy at the higher frequencies which excites the filter dynamic characteristics. Notice how the softer tip doesn't excite this filter ring very much at all.

Generally, a softer tip is a better selection to assure that the filter ring does not occur. If there is filter ring then it makes sense to select a higher frequency range so that the filter ring is minimized. Then it is not a serious issue and the problem is resolved.

I hope that this little discussion has shed some light on this problem regarding filter ring observed on the force time history. If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

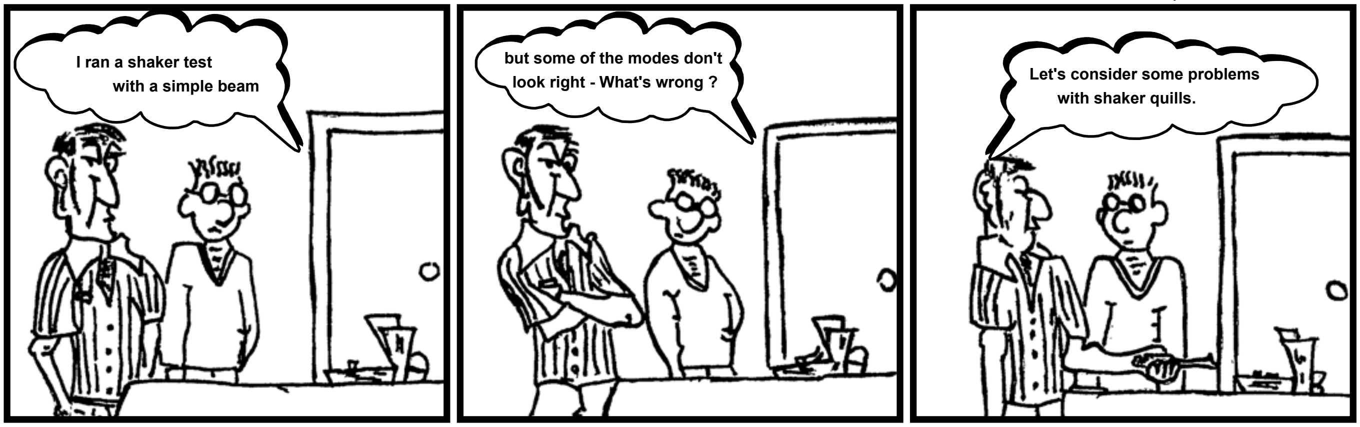

Illustration by Mike Avitabile

I ran a shaker test with a simple beam but some of the modes don't look right – What's wrong?
Let's consider some problems with shaker quills.

Shaker testing for experimental modal analysis can pose some special difficulties if care is not taken in setting up the shaker and attachment device commonly called a "quill" or "stinger". Typically a system is setup as shown in Figure 1. The idea of the stinger is to allow for axial motion to be imparted into the structure which is measured by the force gage for simple compression and tension type loads.

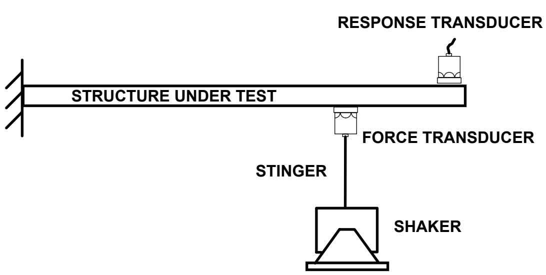

Figure 1 – Typical Shaker Setup

The purpose of the stinger is to allow for loads in the direction of excitation but to minimize the lateral loads that may be imparted into the system. Essentially, a free body diagram concept allows us to know the force imparted into the structure at the attachment point. Therefore, all the dynamic effects of the shaker system and stinger are not included in the dynamic characterization of the structure under test. At least that's what is happening from a theoretical standpoint. Of course this assumes that the stinger has essentially no lateral stiffness and does not have any significant contribution to the overall dynamic characterization of the system. This is extremely important because the force gage only measures the axial load applied – if there is any other loads (lateral or moment) that occur, the force gage does not measure them.

The measurement that was made is described next. (This measurement was received from an outside source). A relatively flexible beam was set up for testing with a shaker similar to that shown in Figure 1. However, the stinger was relatively short and there was a possibility that the rotational stiffness of the stinger may affect the beam flexible modes.

So now let's take a look at some of the measurements that were made. Figure 2 shows an FRF measurement that was taken with a shaker system attached to the structure with a stinger that was possibly too short. This then caused the rotational stiffness of the stinger to be more pronounced – especially relative to the flexible beam that was being measured. A modal test revealed that there was a classical 1^{st} and 2^{nd} bending mode for the first two peaks as expected. However, the next two peaks revealed essentially the same classical 3^{rd} bending mode of the beam. FRF measurements were obtained only for the structure under test but none on the stinger.

Subsequent tests (and additional measurements on the stinger itself) revealed that the two peaks were actually the result of a tuned absorber effect. The stinger was actually in phase with the structure mode shape at 3^{rd} peak of the FRF and out of phase with the structure motion at the 4^{th} peak of the FRF.

The force gage only accounts for the axial motion imparted by the shaker excitation – there is no measurement of the rotational effects associated with the beam rotary stiffness introduced by the stinger in the test setup. But the stinger actually looks like a rotational spring relative to the beam at the attachment point.

In order to confirm the observation, a longer stinger was utilized in a second test of the structure. The longer stinger effectively minimizes the effect of the rotation stiffness imparted to the structure under test.

Figure 3 shows the FRF with the longer stinger attached. It is clear that the FRF is much cleaner and follow the expected pattern of beam like mode response. A brief modal survey was conducted and the first three peaks correspond to the first three classical mode shapes for a cantilever beam.

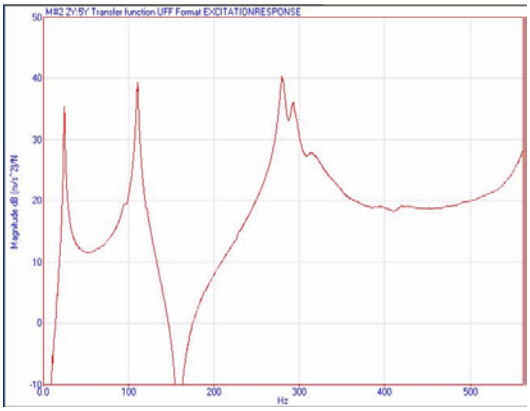

Figure 2 – FRF with Short Stinger

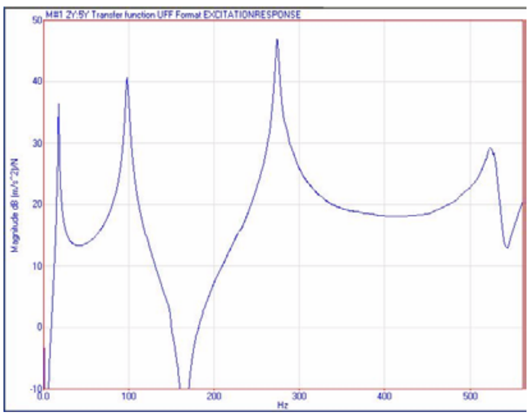

Figure 3 – FRF with Longer Stinger

Clearly, the first two modes see a shifting of frequency due to the different stinger configurations. This can be due to a variety of reasons which might be mass loading effect, stinger effect, different test setup, etc. (These are measurements that were provided by an outside source so I can not be sure of the actual test setup – but the effect is very clear). The third peak is significantly different. There is a splitting of the main peak as is typically seen in tuned absorber applications – there is also a significant reduction in the overall amplitude of the measured response (as is seen in tuned absorber theory).

Figure 4 shows the expected shape that would result if this stinger acted as a tuned absorber to the measurement system. (Again these measurements were provided from an outside source and are used to illustrate the effect that is expected to exist here). Obviously the rotational effects of the stinger at the attachement point on the structure will be more pronounced as the stinger is shortened. If the stinger happens to have the same frequency as one of the modes of the main structure, then the coupling would definitely produce FRFs as shown in Figure 2.

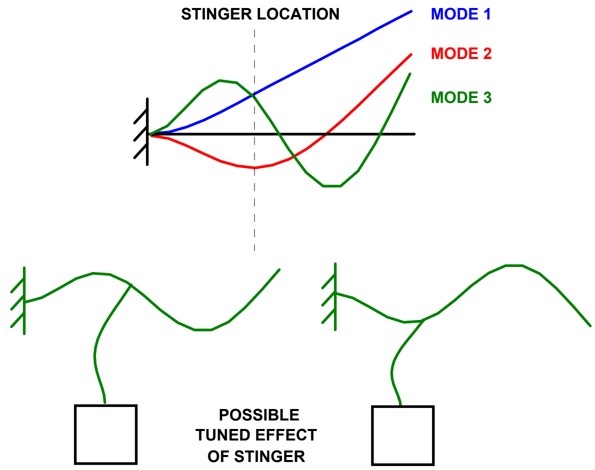

Figure 4 – Stinger Tuned Absorber Effect
(Note: Shapes not to scale; shapes sketched to show expected effect of stinger rotational stiffness coupled to main structure)

Clearly, the effect of the shaker stinger length plays a very important role in the measurement of accurate FRF measurements. If the stinger is too short then there is a general stiffening effect that can be seen in the measured response function. For this particular case, there is a general tuned absorber effect that can be easily seen. This tuned absorber effect may not occur in every stinger application but was observed in this particular measurement setup.

Figure 5 shows an overlay of the two FRF measurements acquired – one with the short stinger and one with the longer stinger. Comparing the two measurements shows significant differences on all the modes of the system measured.

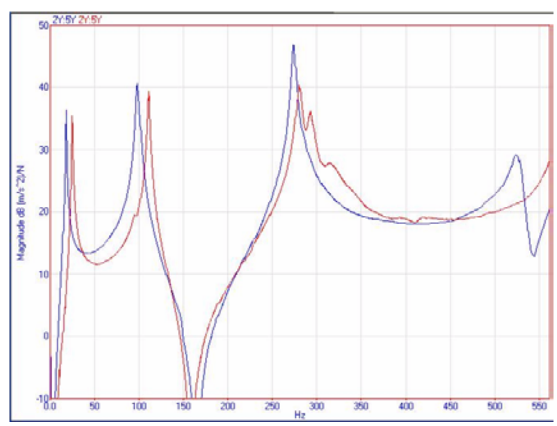

Figure 4 – Comparison FRF

I hope that this little discussion has shed some light on this problem regarding shaker stinger setup. If you have any more questions on modal analysis, just ask me.

Copyright 2005
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Should the measurement bandwidth match the frequency range of interest for impact testing?
Let's discuss this to see why they may not need to match.

Now here is a question that may appear simple on the surface. But as we discuss it, you may realize that there are some alternate issues that may make you think differently regarding this problem. On the surface, it would appear that the measurement to be made should be over the bandwidth of interest.

Obviously, if the bandwidth is narrower, then the higher modes of interest may not be observed. And, of course, if the bandwidth is too large then there will be response of higher modes that may not be of interest. But the real point is this – is the latter case undesirable or can a wider frequency range be selected and get an equivalent or better measurement? Hmmm…. maybe this needs to be discussed and evaluated some more before a final call is made here.

Let's consider a simple measurement on a typical structure where the first two or three modes are of interest. These first three modes are expected to exist over an 800 Hz bandwidth. A typical measurement over that 800 Hz bandwidth with 800 lines of resolution can be seen in Figure 1.

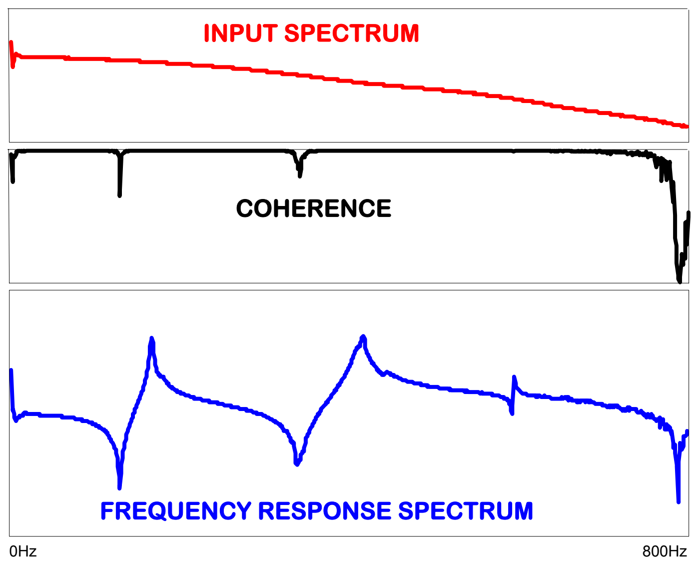

Figure 1 – Input Spectrum, Coherence and Frequency Response Function over an 800 Hz Bandwidth

In general, the measurement looks reasonably good. The frequency response shows the desired peaks well and the measurement appears acceptable. The input spectrum shows reasonably flat input over all frequencies with approximately 20 dB roll off over the frequency range. The coherence is reasonably good at most frequencies in the range of interest. (While difficult to see in the plot, there is some minor drop off of the coherence over the frequency range even at the resonances but likely to be acceptable for most engineers' use.)

So what could possibly be wrong with the measurement? Let's take a look at the time signal associated with the response of the system.

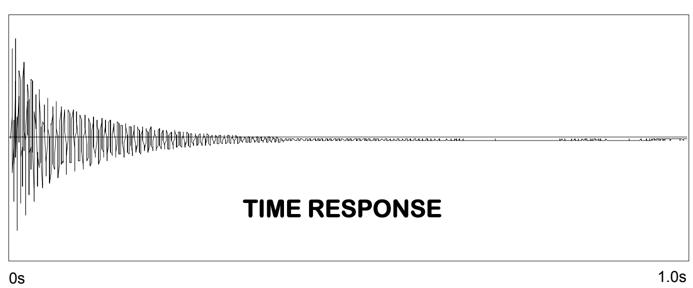

Figure 2 – Time Response Output for 800 Hz Bandwidth

Now the time response is noted to be fairly well diminished within one-quarter of the time record.

Copyright 2005
Pete Avitabile

So is this a problem. On the surface – no. But the question I really want to ask is this. Could I make a better measurement? And how would I do that?

Look at the time response in Figure 2. There is a very strong possibility that any noise on the response channel might be significant in the measured frequency response function. (For the case shown in this example, there was not any appreciable amount of noise. But if there were, then the frequency response function would be degraded as well as the coherence.)

Let's consider a different frequency bandwidth for the measurement. For the next measurement, let's try to optimize the time response to be a significant signal for the majority of the time record or block of data collected. If the frequency range is quadrupled, then the time record length will be one fourth of the original time record. This signal is shown in Figure 3. Notice that the time response now fills the majority of the time record.

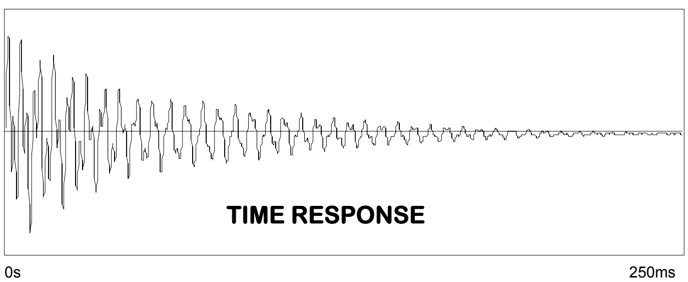

Figure 3 – Time Response Output for 3200 Hz Bandwidth

Now also take a look at the resulting input spectrum, coherence and frequency response function shown in Figure 4.

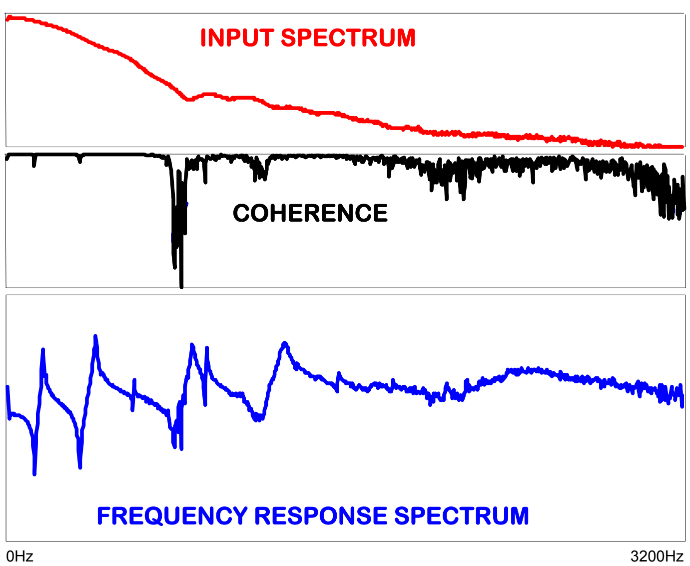

Figure 4 – Input Spectrum, Coherence and Frequency Response Function over a 3200 Hz Bandwidth

The measurement at first glance does not look very good over the entire frequency range. BUT over the range of interest for the first several modes, the measurement is actually very good. (Again, while difficult to see in the plot, the coherence is actually as good as Figure 2, if not better overall.)

So the bottom line here is that the second measurement may actually be the preferred measurement depending on the coherence of the measured response. The trick to this measurement is that the hammer tip should be selected to excite ***only the frequency range of interest*** – NOT the entire bandwidth of the FFT analyzer. In this way, a good measurement can be obtained for the modes of interest.

I have run across this issue several times in a variety of different measurement situations. Generally, people are bewildered why this measurement might be acceptable but as I discuss this measurement problem it becomes apparent that the overall measurement can actually be better than the narrow bandwidth.

A specific example relates to some measurements taken a few years ago on a surveillance pod for an aircraft structure. The initial measurements over the narrow specific frequency range were very noisy since the response died very quickly in the measurement time record. Selecting a wider frequency range, where the response signal was significant over the entire time record, actually produced a much better measurement for the modes of interest. And again, the force hammer tip was selected to excite only the modes of interest and not the entire frequency range of the FFT analysis process. A typical measurement from that structure is shown in Figure 5. (Unfortunately, the narrow band FRF measurement is not available for comparison but was a much poorer measurement overall.)

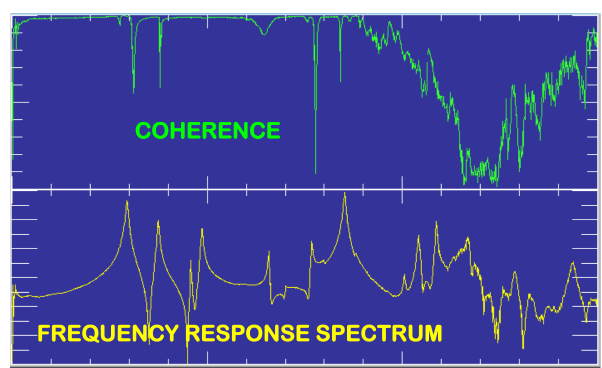

Figure 5 – FRF/Coherence for POD Measurement

I hope that this little discussion has shed some light on alternate ways to improve a measurement. In either case, judgment needs to be made to determine which measurement is the best overall before proceeding with a specific course of action. If you have any more questions on modal analysis, just ask me.

Copyright 2005
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

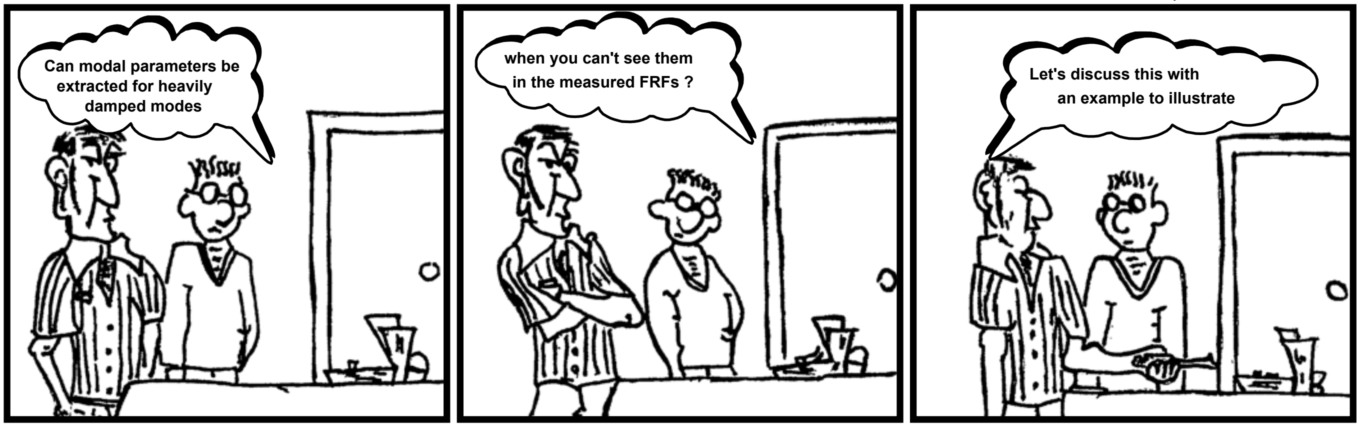

Illustration by Mike Avitabile

Can modal parameters be extracted for heavily damped modes when you can't see them in the measured FRFs? Let's discuss this with an example to illustrate.

Now this is a question that I have heard many times over the years. The answer is bittersweet in many respects. Of course you can extract heavily damped modes from FRFs! But you need to know there is a root in the FRF and make sure that you make a good measurement so that the root can be extracted. Let's elaborate on this with an example to help show that the parameter estimation algorithms are very robust and well-suited to extract heavily damped roots.

To illustrate this point, I am going to refer to a simple model that we have used for many years. It is a very simple 2 DOF mass-spring-dashpot system that has non-proportional damping. The equation of motion and mass, damping and stiffness are defined as

$$[M]\{\ddot{x}\} + [C]\{\dot{x}\} + [K]\{x\} = \{F(t)\}$$

$$M = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}; C = \begin{bmatrix} 0 & 0 \\ 0 & 646.225 \end{bmatrix}; K = \begin{bmatrix} 428400 & -132900 \\ -132900 & 532800 \end{bmatrix}$$

These matrices can be used to extract the complex solution (frequency, damping and mode shapes). In addition, the frequency response functions can be synthesized to simulate a set of collected data. In this case, only one reference, for the first DOF will be used to generate a row of the frequency response matrix as shown in Figure 1. Now it is very clear that only one mode is observed in the peak amplitude of the magnitude of the frequency response function. If we only looked at the magnitude of the function then it would appear that there is only one mode in the system. But if we also looked at the phase, then there is an indication that maybe there is something other than a single mode in the band of interest. (In fact for this case, there are definitely two modes in this band.)

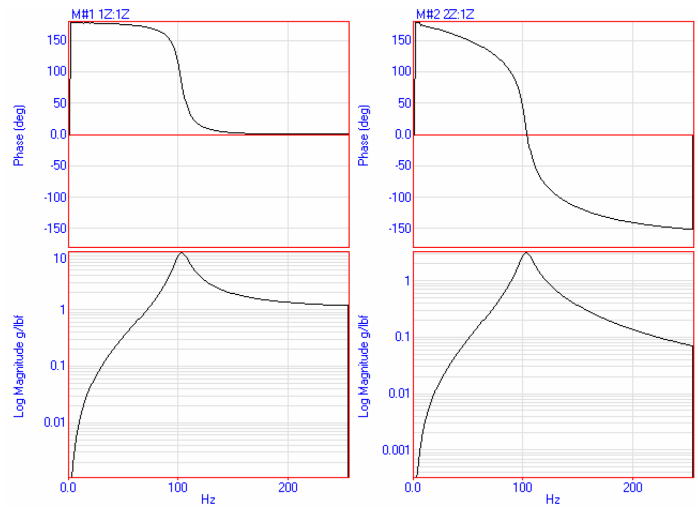

Figure 1 – H_{11} and H_{12} Frequency Response Functions

But the big question is "Can the modal parameter estimation algorithms extract reasonable (accurate) values for the residues of the system?". As a user, the proper order model must be identified for the extraction of residues. If only one mode is requested, then obviously only one mode will be estimated – and it might provide marginal estimates for one of the roots. And if the model is overspecified with too many modes, then the results may be equally distorted - possibly there will be a reasonable estimate for one or two of the modes but they are also likely to be poorly estimated.

Now if the proper order model is specified, will the correct modal parameters be estimated? Using a 2 DOF model fit with an orthogonal polynomial estimation algorithm (popular in many commercially available software packages), the poles extracted are reported in Table 1.

Table 1 – Poles Extracted from Orthogonal Polynomial

Mode	Frequency (Hz)	Damping (%)
1	103	5.31
2	103	40.7

These are actually the poles that would be obtained from a complex eigensolution for the non-proportioanl system matrices presented above. From this table, you can see that there are two roots at the same frequency with one at 5% of critical damping and another at 40% of critical damping which is very heavy damping. Now proceeding on, the residues can be extracted in a similar fashion. For H11 and H12, the residues are shown in Table 2 and 3, respectively. These residues match well with the analytical residues that were determined from the analytical model used to generate the frequency response functions.

Table 2 – H_{11} Residues from Orthogonal Polynomial

Mode	Frequency (Hz)	Damping (%)	Res Mag (g/lbf-sec)	Res Phs (deg)
1	103	5.31	745	183
2	103	40.7	108	21.8

Table 3 – H_{12} Residues from Orthogonal Polynomial

Mode	Frequency (Hz)	Damping (%)	Res Mag (g/lbf-sec)	Res Phs (deg)
1	103	5.31	257	107
2	103	40.7	309	306

Now from this simple example, it is clear that the modes can be extracted from a frequency response function and it is independent on the damping of the system - whether it be lightly damped or heavily damped and whether it be proportional or non-proportionally damped system.

Let's be very clear here… The modal parameter estimation algorithms are very capable curvefitters that are commonly used today in almost all commercial software packages. The problem does not lie with the curvefitter as much as it lies with the measurement and the engineer using the software.

Obviously, the engineer needs to have some indication that there are a certain number of roots in a given frequency band – either from the mode indicator tools or from apriori knowledge that there are a certain number of roots in a given band. While the curvefitters are generally robust, what I find many times is that the mode indicator tools, at times, do not provide a clear, concise indication of the number of roots in a band. And even more often that that, there is generally a poor set of measurements that have been collected that are not considered adequate to extract modal parameters. The problem most times lies with the actual measurements. Generally, they are not of sufficient quality to extract accurate modal parameters. That is the plain and simple fact in most cases!

Another case is a structure with relatively heavy damping and pseudo-repeated roots. The mode indicator tools in Figure 2 clearly help to identify the number of modes present for the roots extracted and shown in Table 4. This case is included here to show that all the tools must be used together to assist in the modal parameter estimation process. These measurements are reasonably good which leads to success in the extraction of modal parameters.

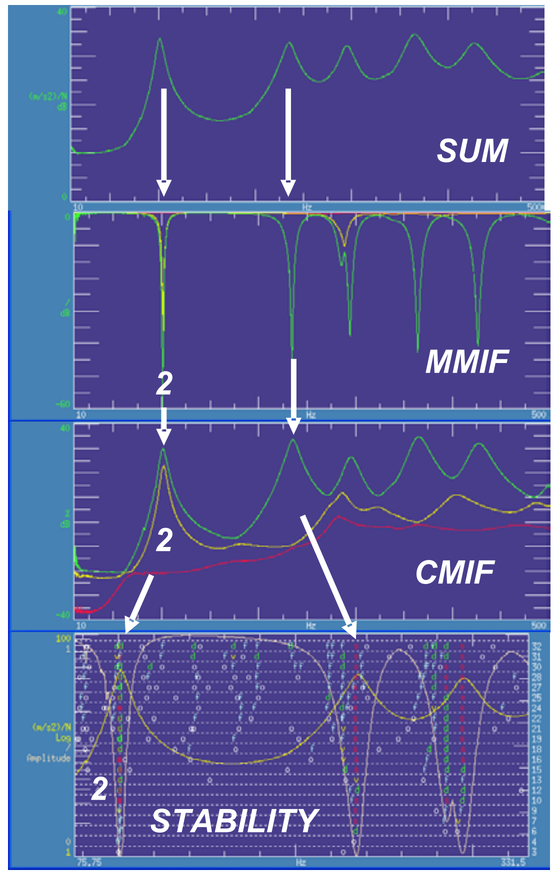

Figure 2 – SUM, MMIF, CMIF, Stability

Table 4 – Time Domain Polyreference Extraction

No.	Frequency	Damping	Stab.	DOFs	Stored
1	101.08 Hz	3.95 %	s	0	no
2	101.62 Hz	3.29 %	s	0	no
3	234.55 Hz	3.20 %	s	0	no
4	285.65 Hz	2.88 %	s	0	no
5	294.90 Hz	2.90 %	s	0	no

I hope that this little discussion has shed some light on estimation of modal parameters with heavy damping and pseudo- repeated roots. The modal parameter estimation algorithms are capable of extracting these roots provided good measurements are provided. But good measurements are the critical key. If you have any more questions on modal analysis, just ask me.

Copyright 2006
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

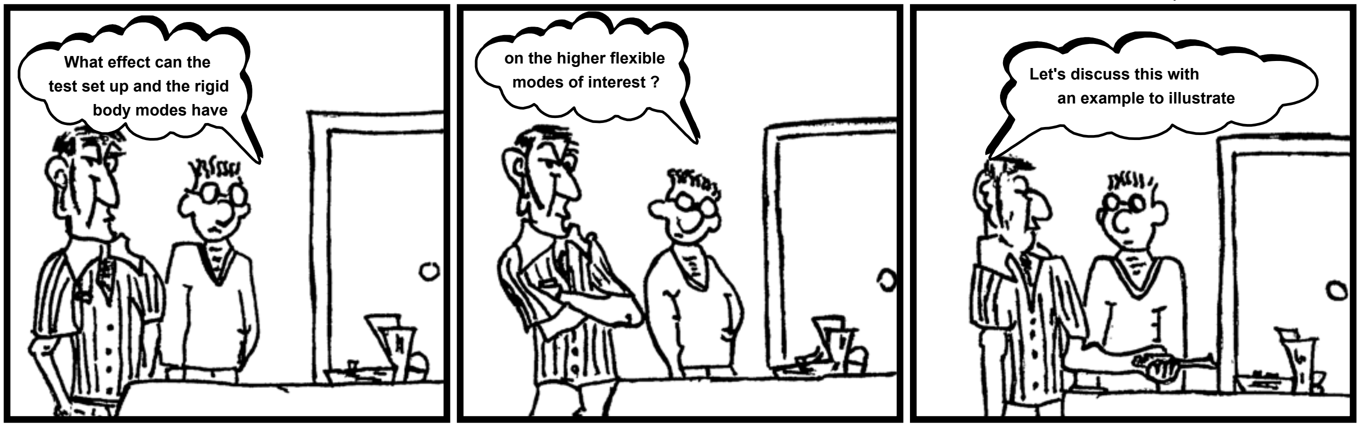

Illustration by Mike Avitabile

What effect can the test set up and rigid body modes have on the higher flexible modes of interest?
Let's discuss this with an example to illustrate.

This particular question comes up very often. In this particular case, the question was posed relative to ground vibration testing of an aircraft. The concern lies with the fact that if a different support configuration is used, how or will this effect the measured flexible modes of the system.

Now there are some very important questions to be answered here. All of them may be more than can be answered in one article but at least some concepts can be presented and some possible ways to better understand the problem can be presented.

In order to do this, I want to show some data that was recently collected in the lab for some composite plate specimens that were subjected to impact testing. The main purpose of the testing was to determine the damping of the composite material using a newer material formulation and compare these results to commercially available composite resins that are typically used.

One of the first things that was done was to subject the first prototype plate to a number of different proposed test set up configurations to determine if the test setup would have a significant effect on the results obtained. Since the plates were of a very lightweight construction, many different set up configurations were explored. Only four different configurations will be shown here to illustrate some of the differences that could possibly result. The composite plate was supported on a very soft elastic system and subject to impact testing using the multiple reference impact technique using three reference accelerometers. The impact test was conducted with the four different support configurations as shown in Figure 1 along with a photo of one of the test set ups. A typical frequency response function measured for one of the configurations is shown in Figure 2 (just for reference).

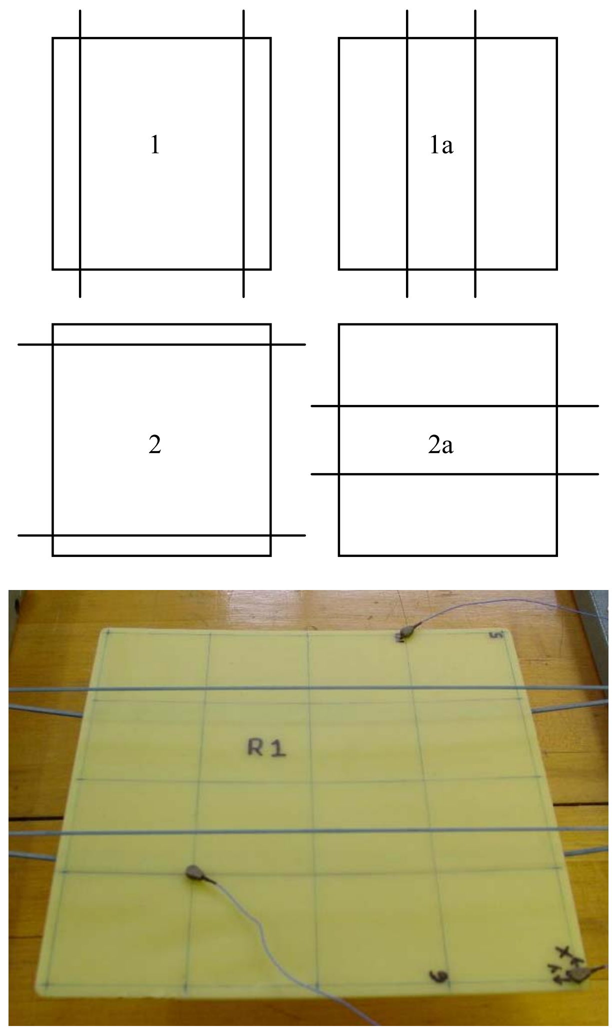

Figure 1 – Schematic of Four Different Test Support Configurations for One Composite Plate Sample

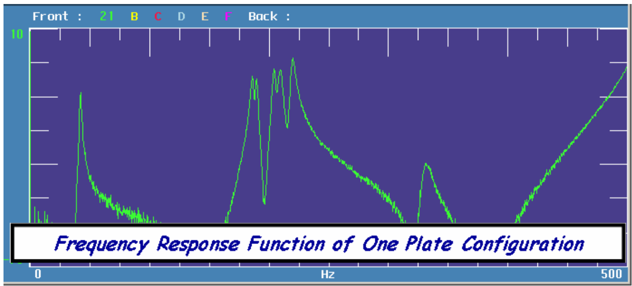

Figure 2 – Typical FRF Measured on Composite Plate

The data was reduced using normal modal extraction procedures and the results for the first four modes are shown in the following tables. The results all seem to be fairly consistent with the exception of the first mode of the structure. There is a definite difference between the different set up configurations. The frequency for most modes varies less than 1% except for the first mode which shows up to 5% variation on frequency results. (Now we can argue about fiber orientation and other factors but the bottom line is that there are differences.)

Test 1 Results (Outboard Supports)

Mode	Frequency (Hz)	Damping (%)
1	43.89	2.15
2	188.48	0.92
3	203.24	0.71
4	207.88	1.03

Test 1a Results (Inboard Supports)

Mode	Frequency (Hz)	Damping (%)
1	42.03	2.33
2	188.02	0.91
3	204.00	0.81
4	209.72	1.09

Test 2 Results (Outboard Supports Rotated 90 deg)

Mode	Frequency (Hz)	Damping (%)
1	43.97	2.19
2	188.51	0.91
3	203.01	0.75
4	209.73	1.05

Test 2a Results (Inboard Supports Rotated 90 deg)

Mode	Frequency (Hz)	Damping (%)
1	42.11	2.32
2	188.47	0.92
3	203.92	0.82
4	209.88	1.07

It is important to note that the rigid body modes are significantly lower than the first flexible mode of the system. (It is difficult to see from the measurement shown but the rigid body modes are close to 1 Hz.) That means that there is much more than a 10:1 ratio between the rigid body modes and first flexible mode of the system. But notice that the first flexible mode is definitely affected by the support configuration.

Everyone always says that as long as there is greater than a 10:1 ratio then there is no effect between the rigid body modes and the flexible modes of the system. But that really depends on what you agree is “close enough”. In this case, if you are willing to accept a 5% variation in frequency, then maybe we could agree that there is “essentially” no effect of the rigid body modes on the flexible modes of this system with the 40:1 ratio between the rigid body modes and flexible modes. But you need to check this and you need to have people agree that this is acceptable. It really depends on how accurate you need your data to be. This will always vary from case to case and industry to industry and test configuration to test configuration.

In the case of these composite plates, there were many modal tests performed and all the results were carefully compared. And not just frequency was compared. Mode shapes were also compared to determine the variance that might be observed from the sets of data collected. You need to check both frequency and mode shapes.

The data collected must be interrogated to determine how the frequencies and mode shapes will vary due to these different test configurations. Maybe there is very little difference in mode shape which may be the parameter of interest. Or maybe the frequency is a sensitive parameter for the design under evaluation. This really depends on the application at hand.

So what should you do? Well… if there is an analytical model available, then it is a very easy task to investigate the effects of different boundary condition on both the frequencies and mode shapes. Each configuration can be easily evaluated using correlation tools to determine the effect on the reported frequencies and mode shapes using vector correlation tools commonly available. This can be done prior to running the actual test to determine what effects, if any, might be observed. In this way, some evaluation can be made as to the expected variation in modal characteristics. Analyses can be performed to determine how these changes in characteristics may effect the final system response. If the effects are significant, then the effects of the test support condition need to be carefully evaluated. But if the results of the system response are not significantly different then the effects of the test support condition can be considered to be not as critical. But someone needs to make this evaluation. Rules of Thumb that are used are exactly that – they must be evaluated in more depth and should not be blindly followed. And remember that a change in the stiffness of the test support ***must*** have an effect on all the frequencies. If you add stiffness the frequencies must shift upwards – the question is how much do the frequencies shift and is it of importance or is it measurable.

I hope that this little discussion has shed some light on the effects of test set up on the frequencies and mode shapes. You need to evaluate this carefully. If you have any more questions on modal analysis, just ask me *by Pete Avitabile*

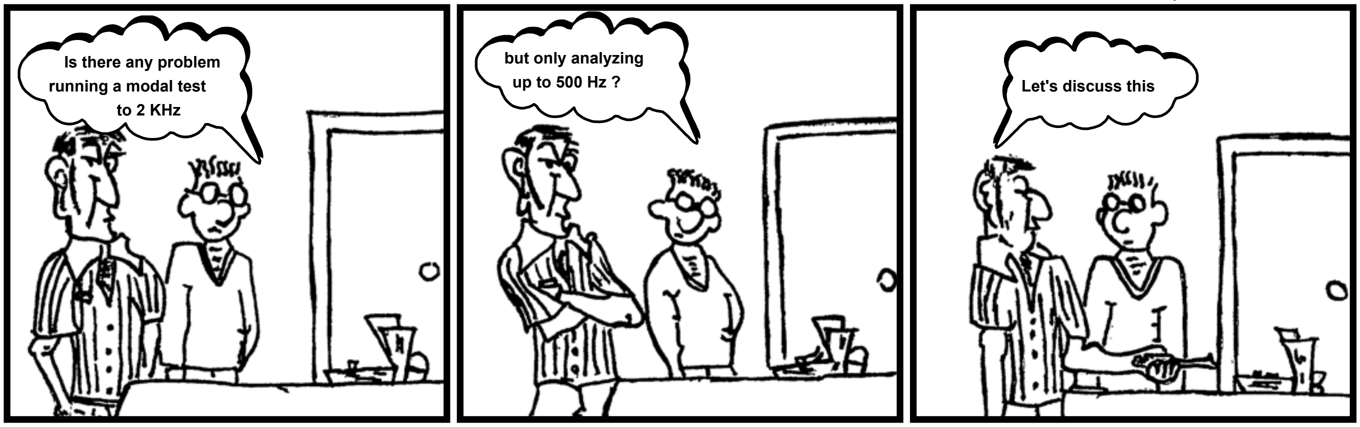

Illustration by Mike Avitabile

Is there any problem running a modal test with 2 KHz excitation but only analyzing up to 500 Hz ?
Let's discuss this.

Now this is an interesting question. There are several issues to be discussed relative to this. The more important question is maybe why one would want to run a test in this fashion in the first place and then discuss some of the issues that might have an effect on the overall measurement and then possibly some alternate things to consider.

So let's consider a measurement as shown in Figure 1. As the question was posed, the measurement would be acquired over a 2KHz range but the only range up to 500 Hz is to be analyzed.

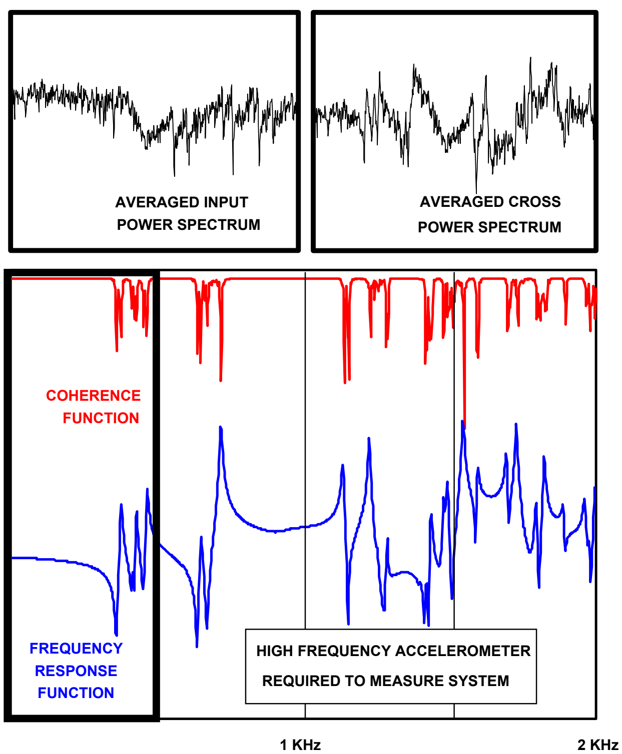

Figure 1 – Measurement over 2KHz with 500 Hz Concern

There is really no right or wrong answer here but I have some strong feelings regarding the adequacy of this measurement as shown. Without some very specific details, I really don't want to make this measurement as requested. Looking at the input power spectrum, cross power spectrum, frequency response function and coherence, there is definitely excitation and response to 2 KHz. There appears to be considerably higher response levels in the higher frequency range as well as many more modes of the system. This measurement looks acceptable overall but is it really the best possible measurement over the 500 Hz frequency range of interest?

The first thing to maybe consider is why is there only a need to extract model information up to 500 Hz when the excitations are a much higher frequency. Well, the analysis or design to be considered may only involve lower order frequencies. It may be that the model to be developed is only needed to address response up to 200 or 400 Hz and there is no need to consider the contribution of higher frequencies for the aspects of the design to be considered. That implies that the higher modes do not significantly participate in the overall response of the system and can be excluded from the analysis.

So if this is the case then the excitation need not extend to a high frequency to extract the measurements and model to describe the system dynamics appropriately. But possibly the excitation may have come from an operating condition where the input excitation is broadband and excites this wide frequency range. But because it is an operating condition, it may be considered a better excitation than an artificially generated excitation – but this is definitely debatable.

But there may also be a dual purpose need for the test. While you may only be concerned for frequencies up to 500 Hz for your analysis, there may be others that need to use and analyze the data for other applications up to 2 KHz. This is always a

Copyright 2006
Pete Avitabile

problem when one test is to be used for multiple purposes and analyses. This is not the optimum way to conduct a test but may be used purely in consideration of time aspects when a test article is not available for long duration or is an expensive piece of hardware on a tight production schedule. In any event, there may be multiple reasons for this type of test scenario.

But what might be the issues that might affect the overall measurement. Well, there needs to be some consideration to the transducers used to acquire the measurement. If the excitation extends to well beyond 500 Hz (and up to 2KHz) then the transducers selected must be suitable for responses at this high frequency range. Of course, this implies that the accelerometers selected should be suitable for high frequency and, as such, may not be as sensitive at lower frequencies than an accelerometer that is selected specifically for a lower frequency range. So the issue that is of concern is the appropriate selection of transducer that is going to provide a suitable measurement below 500 Hz while not being overloaded or saturated by the higher frequency excitation. This can cause an inappropriate transducer selection

As another issue, the excitation up to 2KHz will cause high frequency response that may not be of interest or may excite other problems (such as nonlinearities) that might contaminate the overall measurement. My preference would be to measure only the frequency range of interest as shown in Figure 2.

It seems much wiser to limit the excitation used with a low pass filter and not ever excite the higher frequency modes of the system. This would then possibly allow the use of more

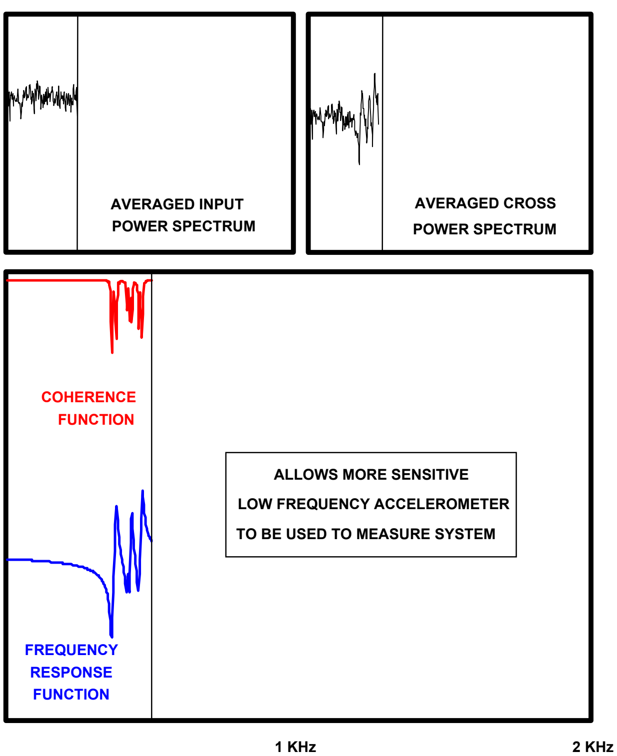

Figure 2 – Frequency Excitation to 500 Hz

sensitive lower frequency accelerometers that would provide a much better measurement overall. This also allows for a better utilization of the analog to digital converter in the acquisition system. But the bottom line is that the instrumentation and their associated signal conditioning must also be considered. Unnecessary loading of the transducer makes no sense at all. Why excite and measure something that isn't of concern?

But looking at the measurement, there may be some concern as to the contribution of the modes just beyond 500Hz and up to 1 KHz. If they are not measured then at some time in the future, there may be a reason or need to evaluate beyond what was required today. And looking at that next band in Figure 3, you can see that there is definitely some dominant modes that may be of interest (if not today, then maybe tomorrow). So you see that often there is not a clear cut answer as to what frequency range might be appropriate. But one thing is clear – the transducers selected for making the measurements are very sensitive to the actual frequency range to be tested and this needs to be well thought out before a test is conducted.

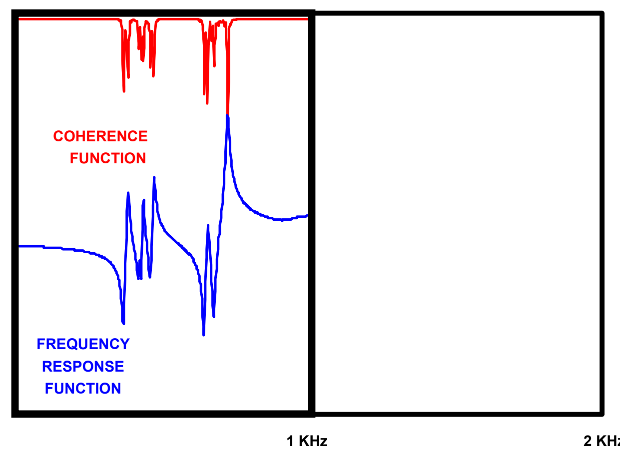

Figure 3 – Measured Response to 1 KHz

So what if I am forced to run a test with a 2 KHz excitation but only analyze to 500 Hz. It might be best to run a test with 2KHz excitation and a second test with 500 Hz excitation. Both measurements should provide equivalent information if all the issues identified above have been properly addressed. And if I am forced to excite the structure to 2 KHz, then I would run both tests and analyze both sets of data to see if there are any significant differences. Of course, this still would imply that the instrumentation would have to be suitable for both frequency ranges and therefore may not be optimum for the lower frequency range.

I hope that this little discussion has shed some light on the effects of acquiring data well beyond the actual frequency range of interest. It can be done, if required, but there may be some issues related to selecting transducers that appropriately measure the actual frequency band of interest accurately. You need to evaluate this carefully. If you have any more questions on modal analysis, just ask me.

Copyright 2006
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

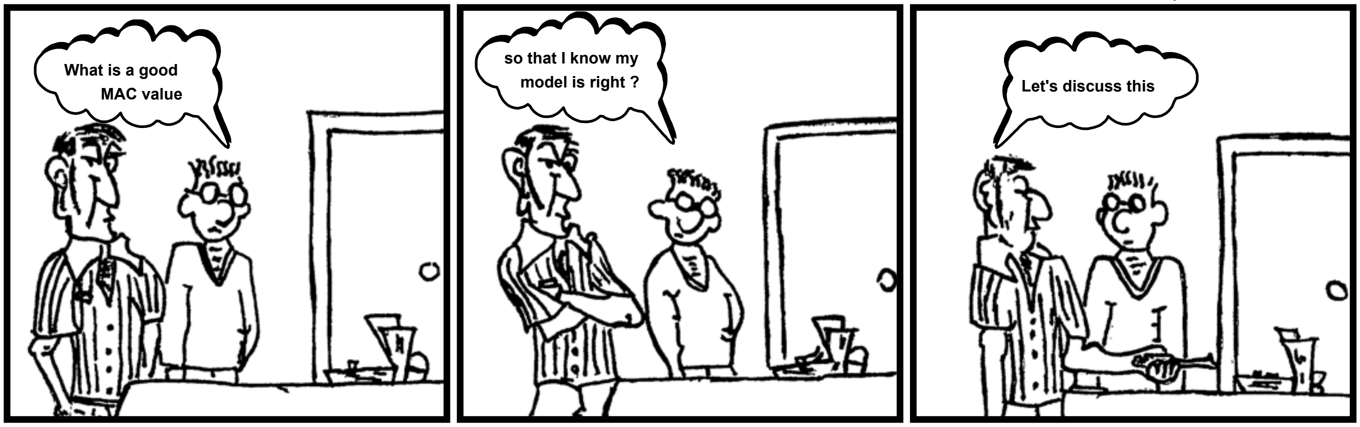

Illustration by Mike Avitabile

What is a good MAC value so I know my model is right ?
Let's discuss this.

Now this is a question which needs a lot of discussion. Many people are often confused about MAC and the other correlation tools that are commonly used. There are a few issues to be discussed in order to clarify some misconceptions.

For purposes of discussion, let's assume that we have an analytical model and experimental data that has close to perfect vector correlation viewed from the MAC (Modal Assurance Criteria) and POC (Pseudo-Orthogonality Check); both approach the desired unity matrix. But, while the vectors correlated well, let's assume that the frequency correlation is not quite as good and assume that there is a 10% frequency variation for the first mode and only a 1% variation for the second mode. So what does this correlation mean then.

To help with this discussion, let's look at the response of a simple plate that has been discussed in several previous Modal Space articles to explain some simple concepts. Figure 1 shows the response of the plate due to a random excitation as the input excitation and corresponding output random response due to that input. Also shown is the frequency respresentation of that input-output phenomena.

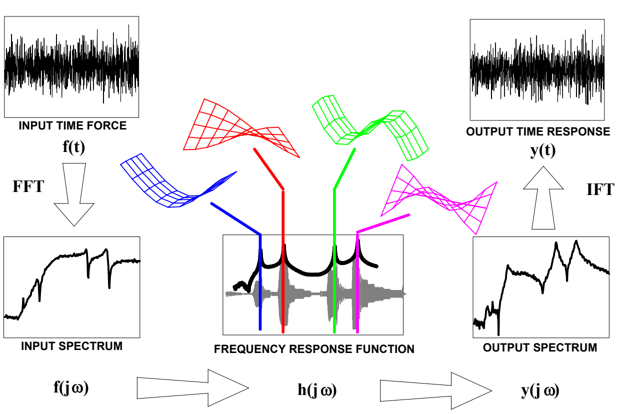

Figure 1 –Overall Response with Random Input Schematic

The FRF and impulse response is nothing more than a filter applied to the input excitation. The FRF is also shown with each of the corresponding mode shapes at each of the resonant frequencies. So we see that the frequency value as well as the mode shape is important for identifying the response of the system. While the shapes are correct, the frequency difference is also important.

If the frequency value is not correct then the response will vary depending on how the input spectrum varies. In this case the second mode frequency is very accurate and the input spectrum is fairly flat over the region of the second mode so the slight frequency variation only causes slight change in the response of the system.

However, for the first mode there is a 10% variation of the frequency. For this mode, there is a significant variation of the input frequency spectrum in this frequency range. So the variation of the frequency is more important for this mode than the second mode.

So it starts to become fairly obvious that the MAC is only an indicator of the vector correlation. But that only identifies if the vectors are correlated. It doesn't provide any information as to the suitability of the model to accurately predict the response of the system. But how does the vector affect the response. Well, the best way to understand the vector effect on the response is to look at the basic equation of motion.

The physical response of the system is

$$[M]\{\ddot{x}\} + [C]\{\dot{x}\} + [K]\{x\} = \{F(t)\}$$

where [M], [C], [K] are the mass, damping and stiffness matrices respectively, along with the corresponding

Copyright 2006
Pete Avitabile

acceleration, velocity and displacement and the force applied to the system. This can be written in modal space as

$$\begin{bmatrix} \ddots & & \\ & \overline{M} & \\ & & \ddots \end{bmatrix}\{\ddot{p}\} + \begin{bmatrix} \ddots & & \\ & \overline{C} & \\ & & \ddots \end{bmatrix}\{\dot{p}\} + \begin{bmatrix} \ddots & & \\ & \overline{K} & \\ & & \ddots \end{bmatrix}\{p\} = [U]^T\{F\}$$

where the diagonal matrices are the modal mass, modal damping and modal stiffness along with the modal acceleration, modal velocity, and modal displacement. The right hand side of the equation has the modal force. Notice the the force is projected to modal space using the transpose of the modal vectors. So the mode shapes are important for the identification of the modal characteristics as well as the appropriation of the physical force to each of the modal oscillators.

If the mode shape varies then the distribution of load and response will vary. So we have to think about how we are going to use the model and more importantly we need to identify what types of loads will be applied and what response is critical to the overall performance of the system. With a random excitation that is broadband and fairly uniform in nature, some of these effects will generally be small.

Now to continue on with another example, Figure 2 shows a sinusoidal excitation with some harmonic components to the driving frequency. Notice that the driving frequency is NOT at one of the resonances of the system. But what if the model frequency was wrong? and the excitation was actually aligned to the first mode? Then there would be more response than predicted.

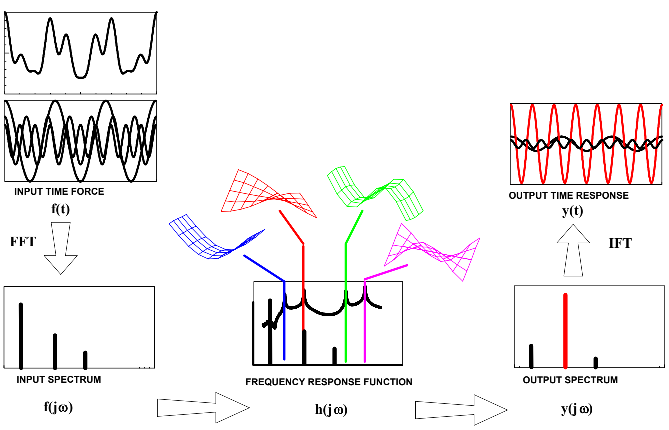

Figure 2 – Overall Response with Sine Input Schematic

And on the other hand, what would happen if the second mode frequency was wrong in the model? Notice that a harmonic of the driving frequency is aligned to the second mode of the system. The response would be predicted wrong.

So we have to start to think about what the MAC values mean in regards to the entire model and the response of that model. The MAC (and POC) helps us to identify how accurate the shape is. But we also need to think about the frequency correlation as well and the forcing function.

So when a correlation is performed, it is important to obtain the best correlation possible. But what does that really mean? There needs to be some assessment of the response of the model due to all the design loadings anticipated. Then someone needs to determine what variability can exist in the model and what effect that has on the computed response. It is then and only then that I can determine how much the frequencies and vectors can vary once someone has defined the acceptable variability in the model.

What we need to realize is that no model is ever perfect. Every model will have variation. And design loadings will have some variation also (if we consider real loading conditions). So before we can ever define levels of acceptable correlation, someone needs to define what is acceptable in terms of the overall system level response. If this isn't done then the levels specified for the frequency correlation and MAC/POC correlation are meaningless. If they are arbitrarily selected, then they may not really be good overall indicators as to how accurate the model prediction may be.

In certain applications there may a very strict requirement that the first and second modes MUST have very accurate frequency correlation as well as shape correlation if the loadings are very sensitive to the frequency of the signal. This is true in applications that involve rotating equipment where the specific operating speeds are critical to the overall response of the system. It may be more critical to have the frequency accurate in some instances and have less stringent requirements on the mode shape correlation. But in other applications where the inputs are uniform broadband excitations, then the frequency correlation may not be as critical and the shape correlation is more important.

This can not be simply identified in a fixed correlation specification. The simple fact is that the correlation and levels of acceptance need to be identified as a result of a detailed analysis of the system in question due to the specific loadings anticipated. Without this important evaluation, then the levels of correlation do not have any practical relevance.

Of course, it would be nice if the models developed satisfy some basic levels of correlation but that does not imply that the model will necessarily produce accurate results if these generic correlation levels are achieved.

I hope that this little discussion has shed some light on the correlation process and why the specific correlation values achieved must be used with an understanding that there needs to be some relationship to the response of the system. If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Someone told me that operating modal analysis produces better results and that damping is much more realistic ? Now this is something that needs discussion.

Now this is a topic that I have seen which causes some confusion among many people. These techniques are very powerful and can have very good results but ….. there are several issues that need to be clearly identified when using these tools that often slip by very quietly but can have very serious consequences if not understood. Let's discuss some of the critical items and issues of concern.

Over the past several years there have been many techniques developed which can reduce data from operating systems. These techniques have been previously referred to as "output only systems" or more recently as "operating modal characteristics". The critical feature of this type of analysis is that the input forces need not be measured in order to reduce the measured data to extract deformation characteristics. This is its biggest benefit ….. and also one of its downfalls. While there is no need to meaure the input, there is also no guarantee that the input exciting the system actual causes response of all the desired system characteristics. This can lead to definition of a system model that does not totally identify all the system characteristics – only those characteristics excited by the unmeasured/unknown force are estimated.

Figure 1 is a schematic we have used before to illustrate the input-output problem for a structural dynamic system. In output only systems, the output response is the only item measured. The assumption is that the input force is generally broadband and excites a frequency band that defines the operating characteristics of the system. However, in Figure 1, the input force (which is not measured) clearly does not excite all the low frequency modes of the system which may be critical to the definition of the dynamic characteristics of the system overall.

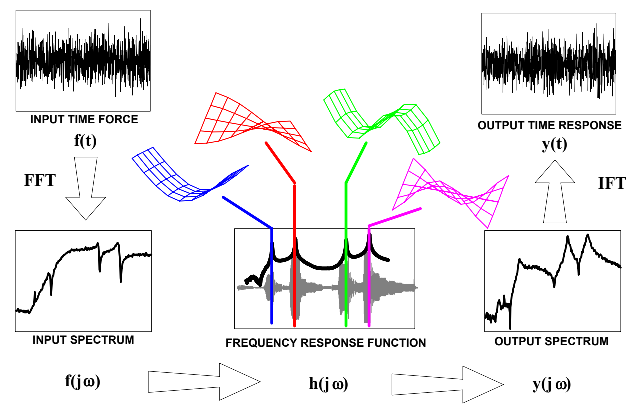

Figure 1 –Overall Response with Random Input Schematic

Well, this may not be a problem as long as this force is truly representative of the actual forcing function and there is no other possibility that the lower frequency modes may be excited by other operating forces. But the problem with output only systems is that you never really know what the excitation force is and if all the modes of the systems have been adequately excited for extraction of a model that completely describes the system characteristics.

So let's just realize that the forcing function is a concern and that it must be broadband in order to extract all the dynamic characteristics adequately. Provided that this is achieved then the modes that represent the system can likely be extracted. But another issue of concern is that there needs to be some way to scale the operating modal data to be used for any further dynamic simulations, correlations to a finite element model, forced response simulation or other dynamic analyses that require scaled mode shapes.

While there has been some research in this area, there still needs to be more work devoted to developing useful techniques that cover a broad range of situations and can provide scaled modes. Hopefully, future efforts will provide these tools.

So one additional critical item that needs to be discussed is related to the estimation of the pole of the system. While the frequency can be estimated relatively easily, the damping is not as simple at all. Many times I have heard people state that the damping obtained from an operating modal analysis is much more accurate that a traditional modal test. While this may be very true of systems that have nonlinear characteristics or bearings or other complicated construction features, the fact is that just about all of the operating extraction algorithms that are available ***all predict damping that appears to be higher than what actually exists*** in a linear time invariant (LTI) system.

To illustrate the fact that output only data reduction schemes always produce higher damping, even on an LTI system, results from two models will be presented here – one case is a pure analytical development of simulated operating data and the other is an actual experimental set up on a system which is extremely linear and for all practical purposes is an LTI system.

For the first case, let's assume that we can start with a linear time invariant system with analytically determined frequency and damping values. The damping will be specified to be 2% for this study. An analytically derived random signal can be applied to drive the LTI system and the output response can be computed. From this time data, a simulated set of data can be used to mimic the output only measurement process. This data can then be processed to extract system characteristics.

This analytical simulation was performed and the starting system characteristics along with the characteristics extracted from the simulated operating data produced the results in Table 1. While the frequencies and shapes are very good approximations of the LTI system, notice that the calculated damping from the output only system is much higher than the starting system. Obviously this is a result of the extraction process; the estimated damping from the output only system is higher than the original damping that was specified for the LTI system.

Table 1: Analytical Model - Prescribed 2% Critical Damping

Original Analytical Model		Random Operating Response	
Freq (Hz)	Damping	Freq (Hz)	Damping
9.1	2.0%	9.1	3.8%
32.5	2.0%	32.6	2.3%
60.3	2.0%	60.3	2.3%

In the second case, an experimental system was setup with a snowboard with no bindings or attachments to the board; this system is extremely linear with none of the typical joints or interactions of components that might cause the system to be nonlinear or appear to have more damping.

A traditional experimental modal test was performed first to estimate the system characteristics. Then a time stream of response (due to arbitrarily tapping the snowboard) was used to provide the simulated operating data. This output only data was processed and system characteristics were extracted.

The results of the traditional experimental modal test with the output only results are shown in Table 2. While the frequencies and mode shapes are very accurate, the damping estimates are not comparable at all. (NOTE: In both cases studies here, commonly used commercial extraction algorithms were employed to estimate parameters to produce the results)

Table 2: Comparison of Traditional Experimental Modal with Output Only Results for a Snowboard Configuration

Experimental Modal Results		Output Only System Response	
Freq (Hz)	Damping	Freq (Hz)	Damping
18.1	0.70%	18.2	2.44%
38.9	0.44%	38.8	1.73%
42.1	0.65%	42.2	1.71%
62.4	0.44%	62.0	2.0%
66.8	0.70%	67.5	2.0%

Now I know that I have presented only two cases here. But there have been many tests (and analyses) on numerous configurations that have substantiated this claim over the years. This may not always be the case but it appears to be true for almost all the cases I have seen thus far. So the most important item to note here is that, in general, output only systems tend to always predict much higher damping than what really exists – even in an LTI system. So please be very careful using the results from these operating modal analyses because the damping predicted may be higher than what really exists in the actual system.

As time progresses, these algorithms will improve and hopefully, they will provide more realistic results as time progresses. But in the meantime be careful using those results. I hope that this little discussion has shed some light on operating modal analysis (or output only systems). If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

Illustration by Mike Avitabile

What is MRIT? I hear people talk about it for impact testing.
Let's talk about this testing technique.

MRIT, or Multiple Reference Impact Technique, has been around for many years now. It became popular when multichannel FFT analyzers became more affordable and more commonly available in everyday experimental modal analysis testing. Let's first start with some simple concepts related to single input single output systems and then move on to a deeper understanding on the information in the FRF matrix. This will lead us to understand why we might be interested in MRIT as a testing technique for the development of multi-referenced data.

In the old days, most people only had a two channel FFT analyzer at best. (You know… we had to walk uphill to school, both ways, in the snow and rain, with no boots or rain coats!). We collected FRFs for one input output location at a time. Then another measurement was taken. Now depending on whether it was an impact test or shaker test would determine the reference location.

In a shaker test, the force measurement was the reference and the accelerometer was "roved" around the structure to different locations. (Obviously it was easier to move the accelerometer rather than the shaker.) Once all the measurements were acquired, a column of the FRF matrix was obtained. The particular column that was measured was determined by the location of the force measurement on the structure.

But in regards to impact testing, possibly the hammer could "rove" while the accelerometer was kept in the same location. In this case, the accelerometer was the reference and a row of the FRF matrix was obtained. Again the particular row is determined by the location of the accelerometer on the structure. (But there is also the possibility that the hammer could be held stationary and the accelerometer would "rove" around the structure).

In any event, the stationary measurement was called the "reference" because it was the same for every input output measurement acquired. Figure 1 shows a typical column from the FRF matrix for a shaker test (or stationary hammer test) in blue and a typical row of the FRF matrix for an impact test (where the hammer roves around the structure) in red.

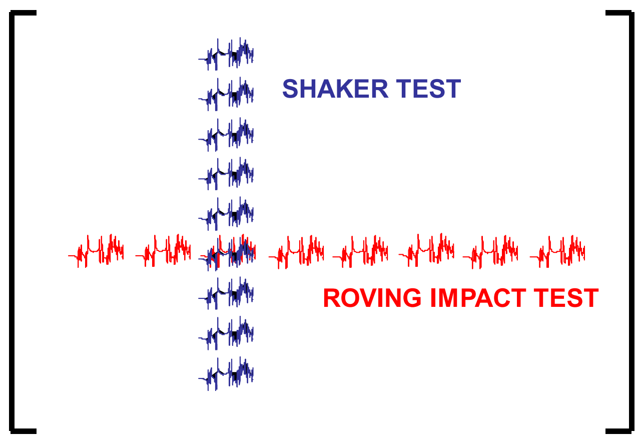

Figure 1 – Typical Row/Column Measured in FRF Matrix

OK – so now we know the old days. Because only one FRF was measured at a time, it was fairly simple to conduct a modal test. But the most critical aspect of the test was the appropriate selection of the reference location. This has been discussed several times before but it is clear that the reference location must be able to measure the mode shape for all the modes of interest from that reference location. The mode shape is related to the residues as

Copyright 2006
Pete Avitabile

$$\begin{Bmatrix} a_{11k} \\ a_{21k} \\ a_{31k} \\ \vdots \end{Bmatrix} = q_k u_{1k} \begin{Bmatrix} u_{1k} \\ u_{2k} \\ u_{3k} \\ \vdots \end{Bmatrix}$$

for one particular reference location. This corresponds to one column of the residue matrix. (Remember that the residue matrix is symmetric so this can also be written to address a row of the residue matrix.) If the reference location is close to the node of a mode for one or more modes, then the measured FRFs will not provide the best information for extraction of the modal parameters. Therefore, this reference selection is critical. However, if more than one row or column of the FRF matrix is collected then redundant information is available. So as discussed several times before, the entire residue matrix is defined as

$$\begin{bmatrix} a_{11k} & a_{12k} & a_{13k} & \cdots \\ a_{21k} & a_{22k} & a_{23k} & \cdots \\ a_{31k} & a_{32k} & a_{33k} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} = q_k \begin{bmatrix} u_{1k}u_{1k} & u_{1k}u_{2k} & u_{1k}u_{3k} & \cdots \\ u_{2k}u_{1k} & u_{2k}u_{2k} & u_{2k}u_{3k} & \cdots \\ u_{3k}u_{1k} & u_{3k}u_{2k} & u_{3k}u_{3k} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

The collection of multiple rows or columns of the FRF matrix is therefore very desirable. The multiple reference modal parameter estimation algorithms take advantage of this redundant information to get the best possible modal parameters from the redundant multiple references. Now I said "redundant" several times to emphasize this important fact. But sometimes these extra references may not be optimal for all the modes if that were the only reference. This is really the reason why multiple references are often used. Just in case one of the references is not located at an optimal location, there will be other references that will contain better information.

Now we understand that it is good to have more than one reference for the estimation of modal parameters. So as multiple channel FFT analyzers became more commonplace, the ability to collect simultaneous sets of references from multiple locations became very possible.

Thus the birth of ***Multiple Reference Impact Testing***. Generally, this can be done by placing multiple accelerometers at various locations on the structure that are expected to be reasonably good references for most of the modes of the structure. So, for example, if a four channel FFT was utilized, then one channel would be used for the force hammer and the remaining three channels would be used for a reference accelerometer. And contrary to popular belief, this does not have to be a triaxial accelerometer at one point on the structure – it is probably better to use three separate single axis accelerometers located at three different locations (and they don't have to be located one in x, one in the y and one in the z axis!).

Using this strategy, then each time a set of averages are acquired, there would be three different FRFs, in three different rows of the FRF matrix. As the hammer roves from one point to another, three additional FRFs would be acquired and as all impact locations were completed, then three separate rows of the FRF matrix would be acquired as seen in Figure 2. This data collection process is referred to as Multiple Reference Impact Testing.

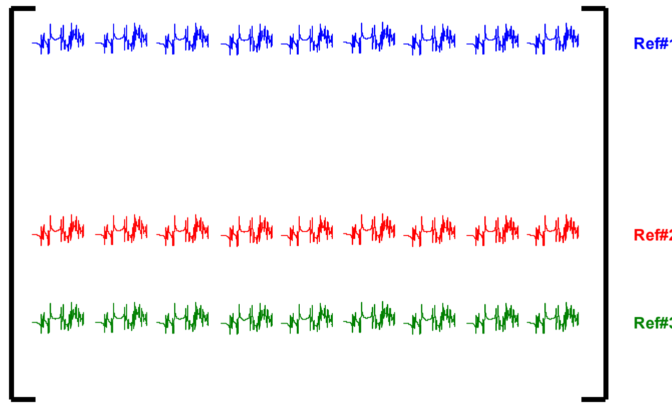

Figure 2 – Multiple Rows Measured in FRF Matrix

One variation of this MRIT occurs when a large multichannel system is used to measure all the accelerometer responses simultaneously. If only one location is impacted then one complete column of the FRF matrix is measured similar to the shaker test in Figure 1. Of course, if we would continue and impact a few different locations, then multiple columns of the FRF matrix would be obtained as seen in Figure 3.

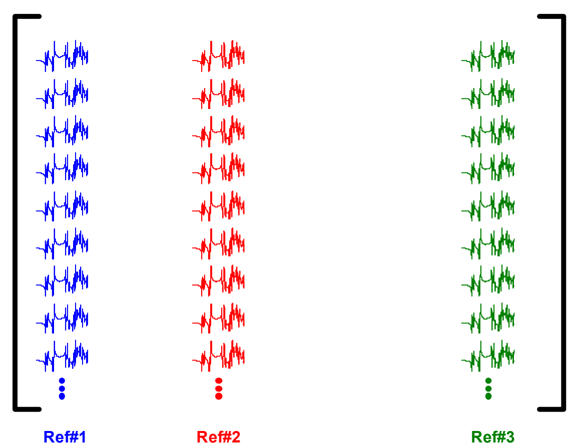

Figure 3 – Multiple Columns Measured in FRF Matrix

In both cases described, multiple reference data is obtained from the MRIT approach. This is a very good way to collect multiple referenced data. If a multiple channel FFT is available, I can't imagine not performing a MRIT test. It doesn't take any more time and multiple reference data results which is very useful.

If you have any more questions on modal analysis, just ask me.

Copyright 2006
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

What is the difference between all the mode indicator functions? What do they all do ?
Let's discuss this.

This is a good question. The indicator functions are very useful. There are several different mode indicator functions that are routinely used in experimental modal analysis when data is reduced. Let's talk about each of the most common tools to show their strengths and weaknesses and how to interpret data from each.

Of course, the measured frequency response function (FRF) can be viewed also but with only one FRF it may very difficult to identify how many modes exist. This is a problem because all of the modes may not be active in the particular FRF measured. The modes may be directional and from one measurement all the modes may not be easily observed. This might also be especially true of the drive point measurement where all the peaks will have the same phase; two very closely spaced modes may be very difficult to observe. So to assist in the process of pole selection, many different tools have been developed over the years. The main tools used are:

- SUM – Summation Function
- MIF – Mode Indicator Function
- MMIF – Multivariate MIF
- CMIF – Complex Mode Indicator Function
- Stability Diagram

So let's discuss each of these. For an example structure, a simple plate will be used as shown in Figure 1. But this plate has some closely spaced modes which will tax all of the mode indicator tools. The plate is subjected to MIMO testing with 2 shaker reference points and 15 accelerometer locations.

The first tool discussed is the Summation Function, SUM. This is a very simple formulation. Basically, it is the sum of all of the FRFs measured (or sometimes only a subset of all the FRFs is used). The SUM will reach a peak in the region of a mode of the system. The idea is that if all the FRFs are considered then all of the modes will be seen in the majority of the measurements. As more and more FRFs are included, there is a greater chance that all of the modes will be seen in the collection of FRFs summed together. This is obviously better than one particular measurement where all the modes may not be present.

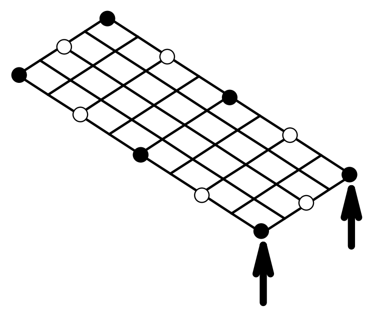

Figure 1 - Plate Test Setup with 2 References

A SUM function for all the measured response functions is shown in Figure 2. The SUM function will identify modes reasonably well especially if the modes are well separated. In the figure, there are five peaks observed which indicates that there are at least five modes in the frequency band shown. Another important feature of the SUM function is that each of the peaks is generally fairly wide and if closely spaced modes exist, then this may not show all of the modes well.

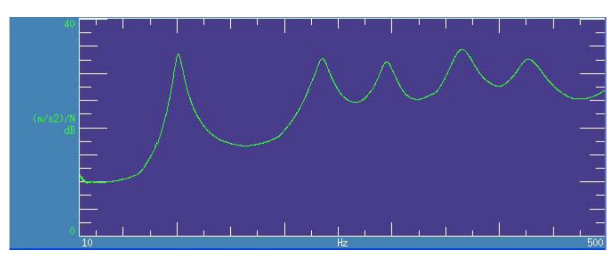

Figure 2 - SUM for 2 References and 15 Accelerometers

While the SUM function is useful, it is not always very clear when modes are closely spaced. The original Mode Indicator Function (MIF) was formulated to provide a better tool for

Copyright 2007
Pete Avitabile

identifying closely spaced modes. Basically the mathematical formulation of the MIF is that the real part of the FRF is divided by the magnitude of the FRF. Because the real part rapidly passes through zero at resonance, the MIF generally tends to have a much more abrupt change across a mode. The real part of the FRF will be zero at resonance and therefore the MIF will drop to a minimum in the region of a mode. An extension of the MIF is the Multivariate MIF (MMIF) which is an extended formulation of MIF for multiple referenced FRF data. The MMIF follows the same basic description of a single MIF. The big advantage is that multiple referenced data will have multiple MIFs (one for each reference) and can detect repeated roots. This is shown in Figure 3.

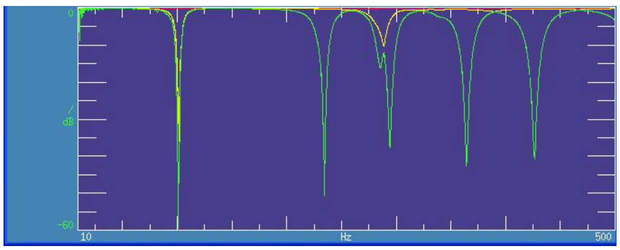

Figure 3 - MMIF for 2 References and 15 Accelerometers

If the first MIF drops, then there is an indication that there is a pole of the system. Every one of the drops in Figure 3 for the MIF1 (shown in green) indicates a mode of the system. Notice that there are six dips in the function – one more than was observed in the SUM function. Clearly, there is one mode that is closely spaced around 300 Hz which was not clearly identified in the SUM function.

Now if the second MIF also drops at the same frequency as the first MIF, then there is an indication that there is a repeated (or pseudo-repeated root). Clearly, the second MIF in Figure 3 (shown in yellow) indicates that there is a repeated root at the first dip in the MIF close to 100 Hz. (Note that the SUM only indicated one mode in this range.) However, the other small dip in the second MIF close to 300 Hz is not an indication of a mode because the second MIF does not dip at the same frequency as the first MIF. In order to have an indication of two roots both MIFs must dip at the same frequency.

The MMIF is a much more accurate tool for indication of modes. However, the assumption is that real part of the FRF is zero at resonance. If the measurements have some distortion or if there is some phasal information in the measurements (associated with non-real normal or complex modes) then the MMIF may not be able to accurately depict the modes accurately.

The Complex Mode Indicator Function (CMIF) is a better tool if this is the case. The CMIF is based on a singular valued decomposition of the FRF matrix to determine all the principal modes that are observed in the set of measurements. The plot of the singular values also helps to identify poles of the system. The CMIF will peak where maximum values exist indicating poles of the system. There will be one CMIF curve for each reference. Figure 4 shows the CMIF.

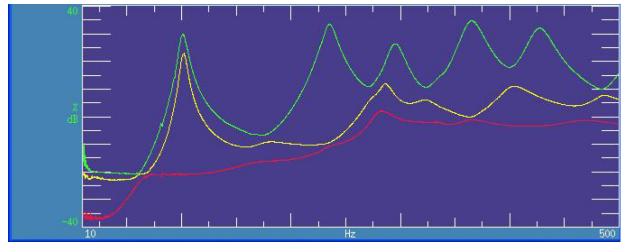

Figure 4 - CMIF for 3 References and 15 Accelerometers

Clearly, the two CMIF curves peak close to 100 Hz indicating that there are two peaks at that frequency. In the 300 Hz frequency range, there is an indication that there are two (or possibly three) modes in that range. The CMIF function provides some additional insight into the number of poles in the frequency band of interest.

All of the tools assist in the selection of poles during the extraction process. The last tool is the Stability Diagram, SD. The basic philosophy is that poles that are extracted from increasing order mathematical model will repeat as the order is increased if the pole is a global characteristic of the system. Other indications of roots will not maintain consistent indication as the order of the model is increased. A plot of these characteristics when a pole migrates to a stable configuration provides yet additional insight into the poles of the system. Figure 5 shows a stability diagram over a narrower frequency range than previously shown. Notice that there is an indication of a repeated root near 100 Hz and another pair of roots close to 300 Hz. (Discussion on details of the stability diagram will be discussed in a future article.) So this confirms the findings from the MMIF and CMIF.

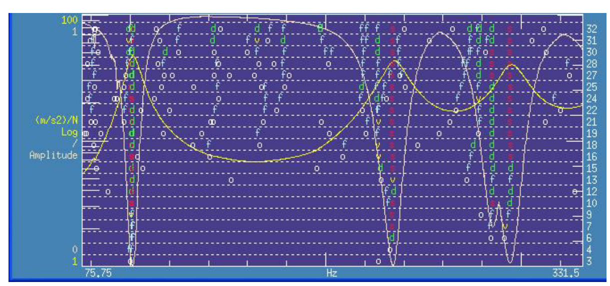

Figure 5 – Stability Diagram for FRF Data

There is a lot more that can be discussed but the majority of the mode indicator tools are explained in this article. If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

Illustration by Mike Avitabile

How do you select the reference location for a modal test? What needs to be considered? Let's discuss this to see how to think about this.

Now the selection of the reference location is one of the more important steps of performing an experimental modal test. If the reference(s) are selected poorly, then there is a strong possibility that one or more modes of the system may be represented poorly or, in the worst case, not at all. Many times the references are selected with a priori knowledge if similar structures have been tested many times in the past. In these cases, the selection is much easier. But when the structure is unique and has no previous history, then the selection of the reference can be much more complicated. Obviously, experience is a very strong plus in these situations. And it may be that there is an analytical model that may assist in this selection of the reference. So let's discuss some basics and describe some things to consider when selecting the reference location(s).

The first thing to really show is the basic equation that dominates the selection of the reference. As I always say to all my students, "Remember … the most important answer to almost all of your modal questions is very simply $\boldsymbol{u_i}\,\boldsymbol{u_j}$". Of course the students all make fun of me for saying this over and over but then they realize that most of their modal questions are often answered with this very statement! So what do I mean by this. Recall that the residue matrix is given by

$$\begin{bmatrix} a_{11k} & a_{12k} & a_{13k} & \cdots \\ a_{21k} & a_{22k} & a_{23k} & \cdots \\ a_{31k} & a_{32k} & a_{33k} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} = q_k \begin{bmatrix} u_{1k}u_{1k} & u_{1k}u_{2k} & u_{1k}u_{3k} & \cdots \\ u_{2k}u_{1k} & u_{2k}u_{2k} & u_{2k}u_{3k} & \cdots \\ u_{3k}u_{1k} & u_{3k}u_{2k} & u_{3k}u_{3k} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

But we do not collect data for all of these input output combinations (and theory tells us that we do not need to measure all of them either). So there needs to be a very careful selection of which rows or columns are measured. If we look at one column then we can write

$$\begin{Bmatrix} a_{11k} \\ a_{21k} \\ a_{31k} \\ \vdots \end{Bmatrix} = q_k u_{1k} \begin{Bmatrix} u_{1k} \\ u_{2k} \\ u_{3k} \\ \vdots \end{Bmatrix}$$

Obviously, the value of the mode shape at the reference location must be significant for all the modes to be measured. If this is done then the FRFs measured will have strong response of the modes of the system. But if the value of the mode shape at the reference location is not significant for one or more modes of the system, then the FRFs may not contain strong response for all the modes of the system. This will make the modal parameter estimation process more difficult.

So if an analytical model is available, then the mode shapes can be reviewed to select optimal reference locations. One simple tool that is often used is the drive point residue. Basically, this is an assessment of the mode shape represented as a residue

$$a_{iik} = q_k u_{ik} u_{ik}$$

This is a common tool used in preliminary assessment usually called a Pre-Test Analysis. Of course there are other tools such as Mode Shape Summation, MODMAC, Effective Independence, along with others that are beyond what can be discussed here. But what if a finite element model is not available or (let me say this quietly) what if the model is not correct. So we need to be able to select the references without any previous knowledge or assistance from an analytical model.

So often times, an experimental test is setup and the first thing that is done is to make sample measurements to determine how many modes might exist in the structure. At times, the drive

Copyright 2007
Pete Avitabile

point FRFs are reviewed – possibly the imaginary part of the FRF is viewed. Unfortunately, this is probably not one of the better measurements to view. This is because closely spaced modes may be difficult to identify since all of the peaks of the imaginary part of the FRF will have the same positive or negative going peaks in the function. Actually the non-drive point measurements are better because the values of the amplitude can be both positive and negative enabling a better chance to identify closely spaced modes. For example, two measurements are shown in Figure 1. The upper trace is a drive point FRF and it is impossible to identify that there are two modes at the first peak. The lower trace is a cross measurement and it is more obvious that there are two modes at that frequency. So you can see that a drive point measurement is useful but may also be deceiving in that the strength of each mode is not apparent in the measurement.

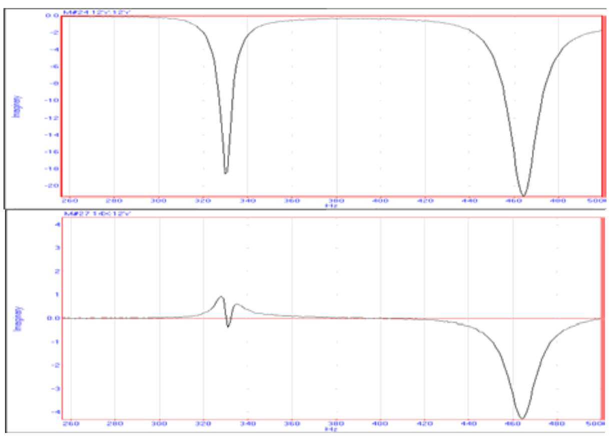

Figure 1 – Drive Point Vs Cross FRF with Close Modes

As a modal test is set up, often, a random sampling of FRFs is made with an educated guess as to what might be reasonable references. This random selection is shown in Figure 2 with the selected measurements shown in different colors.

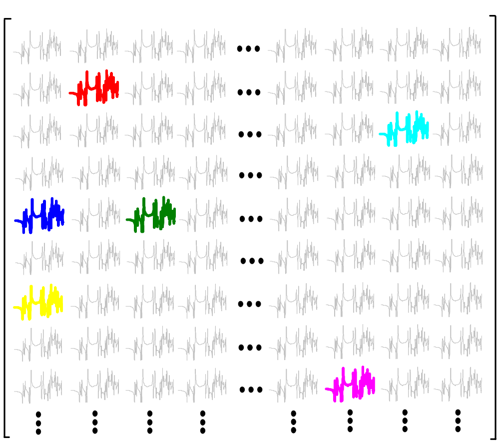

Figure 2 – Random Selection of FRFs for Test Setup

The FRFs are reviewed and identification of peaks in the FRFs are noted from one measurement to the next. If all the peaks are the same and no additional peaks are obtained, then the references might be reasonably selected from those measurements made. Unfortunately all the measurements are made in a somewhat random fashion. There is also a very strong possibility that critical modes may be missed with this procedure. (I have seen even the best of test engineers occasionally miss major modes of a structure)

Another possibility to identify potential references is to obtain a small set of FRFs at all of the potenetial candidate reference locations. This set of FRFs is shown schematically in Figure 3. An SVD is then performed on this matrix. By evaluating the SVD of submatrices of this original matrix (ie, removing individual references in a controlled fashion), an evaluation of the number of significant modes can be determined. If the same number of significant modes are obtained, then the reference removed was not a critical reference for the identification of the modes of the system. However, if fewer significant modes are identified then the reference removed was an important reference for those modes no longer observed and should be retained as a reference for the modal test.

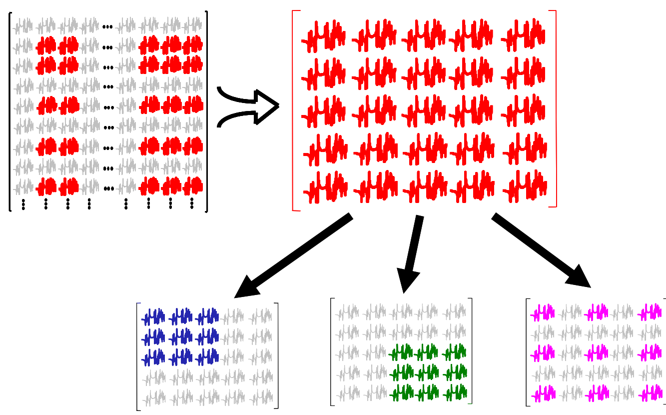

Figure 3 – Systematic Selection of FRF Submatrices for SVD

So while it is common practice to take a handful of randomly selected FRFs to identify a potential reference location, a possible alternate approach that utilizes a mathematical approach to perform an SVD with a set of FRFs may be a much more rigorous mechanism for identifying potential references. This approach, commonly called the Test Reference Identification Procedure (TRIP), offers a technique for reference determination. This is especially useful when no analytical model is available or when there is skepticism as to the accuracy of the finite element model used for the Pre-Test analysis.

The real trick here is to pick a reasonable $\boldsymbol{u_i}\,\boldsymbol{u_j}$ term such that the value of the mode shape at the reference location is a significant value. This will then cause the FRFs to have significant peaks that allow for adequate measurements to be made. Of course, you have to have an idea of what the mode shapes of the system are in order to achieve this. A finite element model or apriori knowledge is very beneficial to accomplish this.

If you have any more questions on modal analysis, just ask me

Copyright 2007
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

Illustration by Mike Avitabile

How do you interpret the stability diagram? And how do data points affect the fit?
There are some concepts here that are important to discuss.

The parameter estimation process is a very important part of the extraction of a model (poles and residues). Usually this is broken down into two parts – the extraction of the poles in the first step and then the estimation of the residues in the second step. The stability diagram is a tool that is used in the development of the extraction of the pole from the data. Let's discuss the estimation of poles and the use of the stability diagram. A few simple examples are included here to drive home the point of critical issues in the estimation process.

Let's assume that we have a set of data as shown in Figure 1. As a starting point, a third order fit will be assumed to describe the phenomena well. In general, the fit is reasonable as evidenced by the R^2 coefficient which is large. But when the variance tolerance is included (dotted lines), there is a fair amount of variation possible. One point is clearly seen as an outlier to the fit of the data. If this outlier point is removed from the data set as seen as in Figure 2, then the R^2 coefficient increases. So from the set of data shown here, it becomes very clear that the data quality is very important to the extraction of a valid set of parameters. It is of paramount importance to have good quality data for the estimation process.

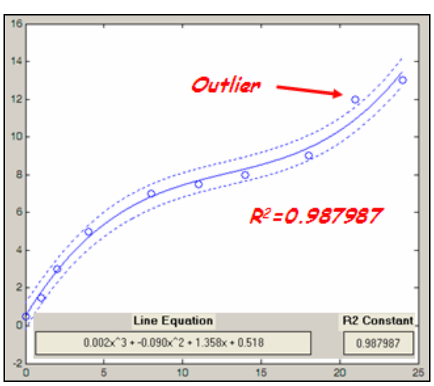

Figure 1 – Fit of Data with Obvious Outlier

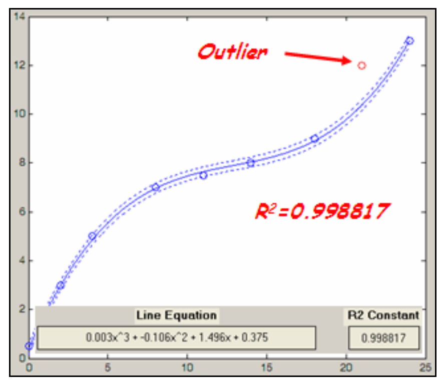

Figure 2 – Fit of Data with Outlier Removed

From this simple example, it is clear that good data is important. Now consider the data set shown in Figure 3. This is a very simple set of data that appears to have a very simple first order characteristic. Let's study the estimated parameters as the order of the model is increased.

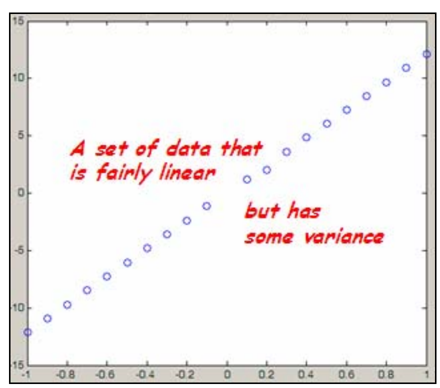

Figure 3 – Set of Fairly Linear Data

Copyright 2007
Pete Avitabile

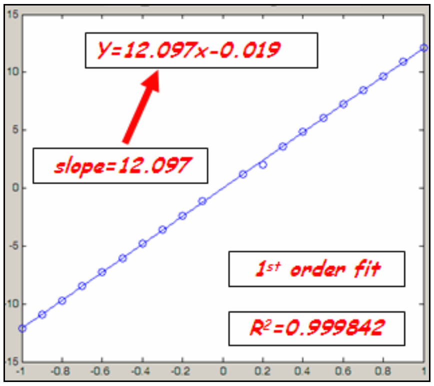

Figure 4a – First Order Fit of Data

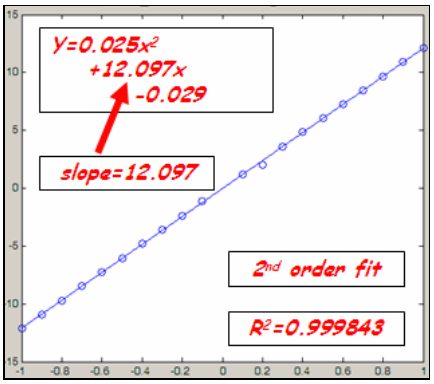

Figure 4b – Second Order Fit of Data

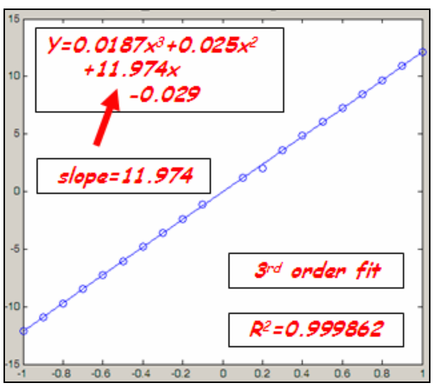

Figure 4c – Third Order Fit of Data

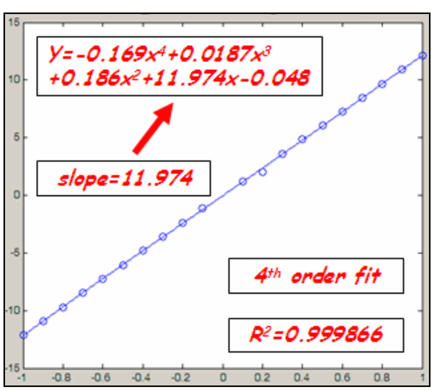

Figure 4d – Fourth Order Fit of Data

The plots in Figure 4 show the progression of the estimation of the slope as the order of the model is increased from first order to fourth order. In Figure 4a, the first order fit produces a slope of 12.097 with a very good R^2 value. Now as the order of the model is increased to second order, the slope is still 12.097 with a good R^2 value. So increasing the order of the model to second order has not produced a change in the estimation of the slope. Of course, the higher order terms are basically making adjustments to account for the variance on the measured data.

As the order of the model is increased to third order, the slope is 11.974 which is very close to the slope previously computed from the first order and second order models. In fact, the slope is only 1% different. So we could argue that the slope is basically the same and has not changed significantly from the previous estimates. And as the order model is further increased to a fourth order model, the slope is again estimated to be 11.974 which is unchanged.

So after this process is complete, the general consensus would be that the parameter of the slope of the data is approximately 12.0 and that very little change occurs as the order of the model is increased. Also note that it doesn't matter which order model I use because to within the tolerance of 1%, all orders produce essentially the same slope!

This simple example really provides an understanding of exactly what goes on behind the scenes in the development of the stability diagram. As the order of the model is increased, there will be estimates of poles. If the pole estimated only changes very slightly from one order model to the next, then the software will provide a flag (or indicator) to help understand if the pole has reached some "stable value" within some specified tolerance. (These tolerances might be set set to 1% on frequency and 5% on damping to identify pole stabilization.) There are usually some indicators that will be provided superimposed on a SUM, MMIF or CMIF plot. A typical stability plot is shown in Figure 5 for reference. The stability diagram helps to identify which poles are "consistent" or stable as the order of the model is increased.

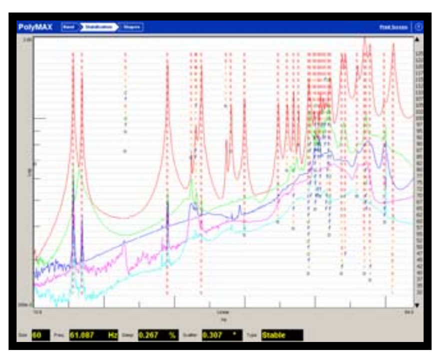

Figure 5 – Typical Stabilization Diagram.

If you have any more questions on modal analysis, just ask me.

Copyright 2007
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

The Stability Diagram has poles that are not indicated by the MMIF or CMIF. Are they really modes? There are some concepts here that are important to discuss.

Now this is a problem that can very possibly occur and needs some discussion in order to sort out what is happening in this situation. We have discussed this before when the problem pertained to not measuring a significant portion of the structure. But in this example, we are going to see that even if we measure appropriately, there are additional issues that must be addressed.

For the example, I am going to use the same plate structure previously used that has two very closely spaced modes in order to show some situations that are possible. We just recently discussed all the mode indicator tools so their use is understood. This plate will be evaluated with several different references to illustrate some points.

For the first case, the plate will be evaluated with more references than needed in order to show the modes of interest. Three references will be used on the plate as shown in Figure 1.

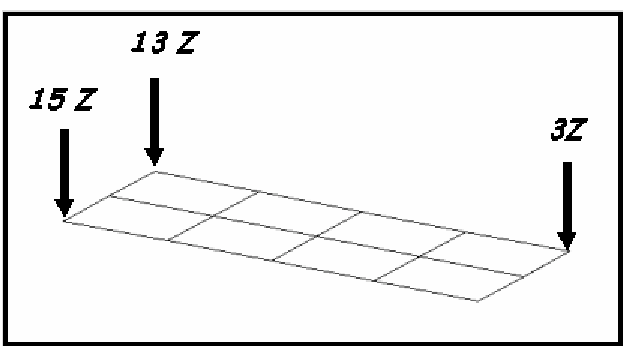

Figure 1 – Plate with Three References Identified

A typical SUM block (upper) and three drive point FRFs (lower) are shown in Figure 2 for reference. Now using these three references, the MMIF and CMIF both show that there are two closely spaced modes at that first frequency around 100 Hz as seen in Figure 3; only the CMIF is shown.

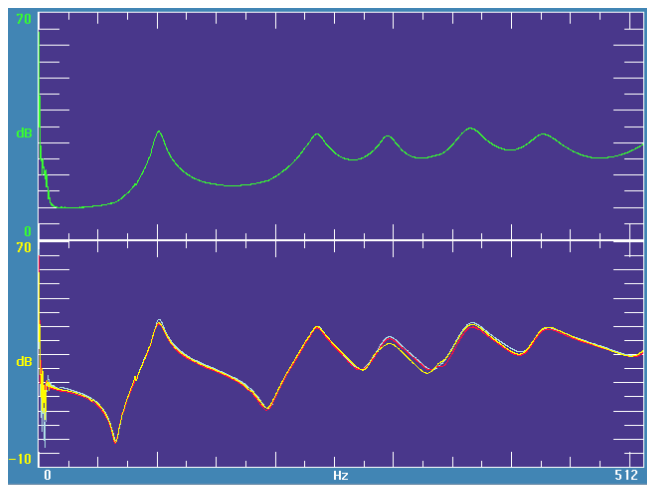

Figure 2 – SUM (upper) and FRFs (lower) for Plate Structure

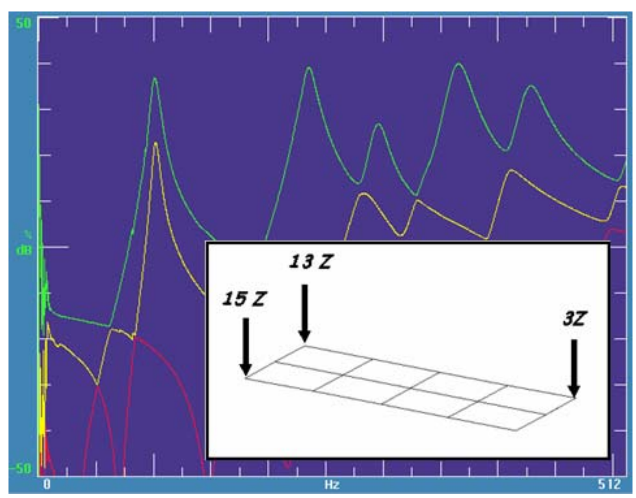

Figure 3 – MMIF and CMIF for Three references

The stability diagram very clearly shows that there are two modes present in that frequency range as shown in Figure 4.

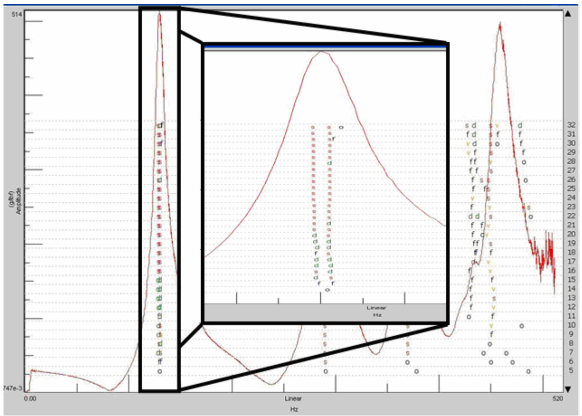

Figure 4 – Stability Diagram with Three References

The mode shapes corresponding to this frequency range are bending and torsion as indicated in Figure 5. These two modes occur at almost the same frequency and while not perfectly repeated, they do occur so close that they are referred to as “pseudo-repeated” roots.

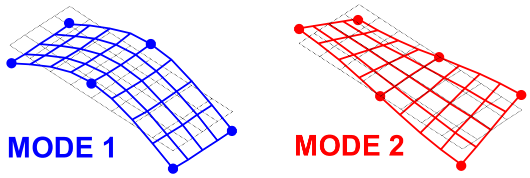

Figure 5 – Bending and Torsion Modes of the Plate

Now Figure 6 shows the CMIF with different combinations of only two of the three original references. Notice that the two references adjacent to each other in Figure 6a and 6b both show two modes in that frequency range but that the two references at opposite corners in Figure 6c do not. (Note that only the CMIF is shown for brevity but the MMIF which is not shown confirms the same results seen with CMIF. Also note that the stability diagram is essentially the same as Figure 4 using any two of the references shown in Figure 6.)

So why does this happen? Why do the MMIF and CMIF not clearly show the modes all the time? In order to answer this, the modes shapes of the structure must be discussed relative to the reference location.

Consider the reference in Figure 6a. Notice that these two corners of the plate both have negative shapes values for bending while the torsion has a negative and positive value for shape. The same is also true for the references in Figure 6b. But when looking at the mode shapes at the reference locations in Figure 6c, something different happens.

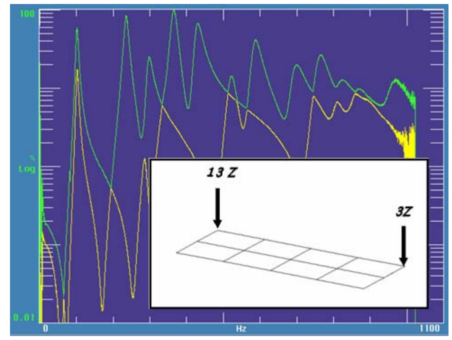

(a) – Reference 3Z and 13Z

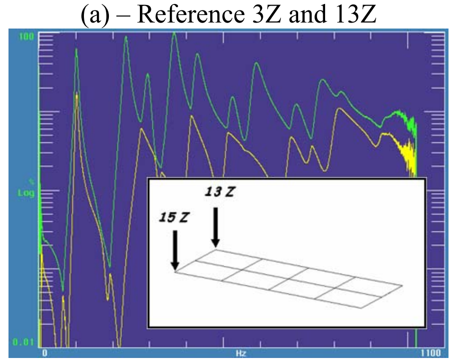

(b) – Reference 13Z and 15Z

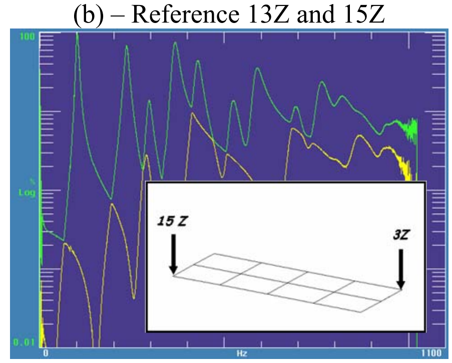

(c) – Reference 3Z and 15Z

Figure 6 – CMIF for Different References

In this case, the mode shape for the first mode has the same sign and direction - and the mode shape for the second mode also has the same sign and direction. Whether they are plus or minus is not important. What is important is that the points have the same phase. There is no way to distinguish the difference between mode 1 and mode 2 from the reference location in Figure 6c. But the references in Figure 6a and 6b can distinguish the difference in the mode shape because of the phase information at those reference locations.

So it is not enough to have two references on the structure in order to identify pseudo-repeated roots. The references must provide an independent view of the mode of the system from the reference location in order to distinguish the modes. But the stability diagram can still identify the fact that there are two roots at that frequency. So from this example, it is possible to have roots in the stability diagram that may not been seen in the MMIF and CMIF.

If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

Illustration by Mike Avitabile

Bolted joints are common in structures. Can the frequency change significantly due to joints?
Now there are several things to discuss here.

Bolted joints are common in many applications. This connection mechanism is found just about everywhere and may be a significant contributor to the frequencies of the system. Now if the structure is assembled very carefully and in a repeatable fashion, then the structure frequency may have relatively small variation in the structure's frequency. However, if the joint is not assembled in a consistent fashion, then there may be considerable difference that may cause a significant variation in the frequency of the structure. Obviously, there will be many variations possible depending on the joint configurations used in each particular application.

While there may be analytical models that may be developed to study the effects of the joint configuration, these models will have many of their own assumptions which may cause variation. For instance, the element type, mesh density, joint configuration and actual connection configuration will all contribute to the variation that may be seen. In fact, a detailed study of some of these parameters shows that there are many issues to be understood.

Rather than discuss all of the analytical modeling issues that may need to be addressed, an actual configuration of a portal frame with bolted joints will be used to show some of the frequency variations that may result due to bolted joint configurations. This portal frame has been used for many different studies including effects of bolted joint arrangements (and has been used in the Los Alamos Dynamics Summer School program for a variety of different studies). The portal frame used for this study is shown in Figure 1. The structure will be tested with a normal well assembled joint and then the structure will be assembled with very deliberate joint mis-orientation to show the change in the frequency of the structure.

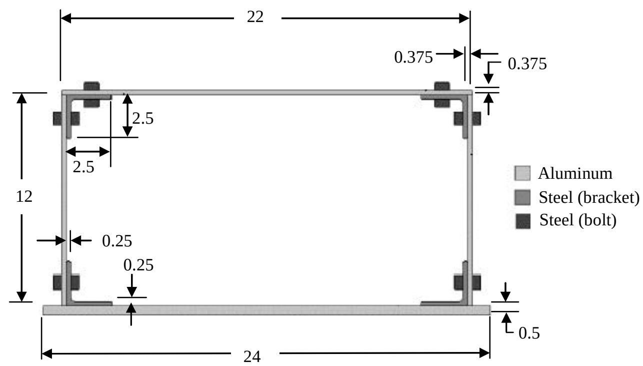

Notes:
All dimensions in inches
Depth into plane = 2 in., except base plate = 6 in.
All four brackets are identical with thickness of 0.25 in.
The two sides are identical with thickness 0.375 in.

Figure 1 – Portal Frame Configuration

Generally, there is care in the development of any joint in a structural system. But what if the structure assembly is not properly performed or if there is some manufacturing variation that causes difficulty in assembling the structure. In addition to the normal assembly of the joint, two cases are considered here. One case allows for a misalignment in the angle bracket and another case considers a shim to force a misalignment. The three different configurations are shown in Figure 2. A properly mated assembly is shown in Figure 2a, a sloppy assembly with misalignment is shown in Figure 2b and a shim assembly is shown in Figure 2c.

(a) (b) (c)

Figure 2 – Three Joint Configurations Studied

Now each of these configurations were assembled to determine the frequencies for the first three modes of the structure for comparison. A typical drive point measurement on the upper beam of the portal frame in the vertical direction was made for comparison. Figure 3 shows the original measurement for the properly mated assembly (black - top), a sloppy assembly due to misalignment (red - middle) and the shim assembly (blue - bottom). Even with the naked eye, there are observed differences in the peaks of the frequency response functions for the three different configurations. Generally, the amplitudes are very similar but the frequencies are definitely different.

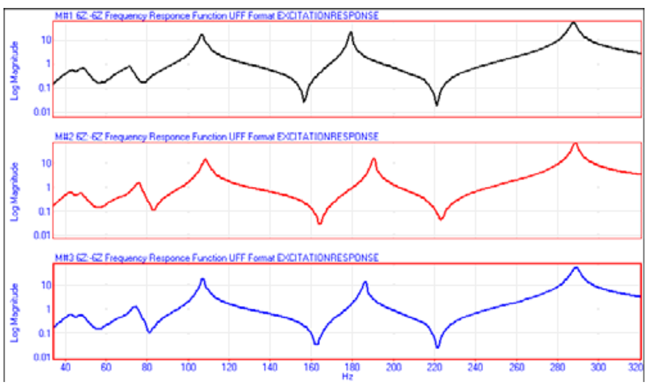

Figure 3 – Frequency Response Functions for Normal Assembled (upper), Misaligned (middle), and Shim (lower)

For each of the frequency response measurements shown in Figure 3, modal parameters were estimated using a frequency domain polynomial approach. The resulting frequencies and damping for each of the configurations are shown in Table 1.

Table 1 –Frequencies/Damping for Three Configurations

Normally Mated Assembly

Mode	Frequency (Hz)	Damping (% Critical)
1	71.7	2.65
2	106.	1.08
3	179.	0.334

Sloppy/Misaligned Assembly

Mode	Frequency (Hz)	Damping (% Critical)
1	75.7	2.15
2	108.	1.07
3	190.	0.364

Shim Assembly

Mode	Frequency (Hz)	Damping (% Critical)
1	74.4	2.53
2	107.	0.843
3	186.	0.425

So from this data, it is very easy to see that an improperly assembled joint can have a change in the frequencies of the structure. This needs to be very carefully evaluated to understand the variation that may result from the assembly process if care is not exercised in the manufacturing/assembly process.

In this quick study, only a few configurations were shown to illustrate what could happen with just a few very simple alterations to the joint assembly configuration. The results show differences and need to be carefully assessed and evaluated.

Bolted joints pose very significant effects that may need to be evaluated in much greater detail than that shown here. But I hope that these simple cases illustrate some of the variation that may result. If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Sometimes the mode shapes appear to be rotated from what is expected. Are the modes wrong? What's up? Now this is something that needs to be discussed.

Now this is a topic that comes up often. So it is going to need some discussion. I have seen people often get confused about the mode shapes for a system. Often times people have a preconceived notion as to what the results "should be". When the mode shapes appear different than that expected, then you might think that the modes are wrong.

Most often when modes of a structure are very closely spaced the mode shapes that satisfy the system can be linear combinations of each other. Therefore, the shapes might be rotated from what you might have expected. The only real requirement is that the modes of the system are orthogonal with respect to the system mass and stiffness matrices. Each of the modes of the system are unique from each other.

In order to help describe this, a simple geometry example will be used along with a simple beam and plate to describe what can be happening here.

A simple x-y coordinate system to describe a rectangular shaped area is shown on the left in Figure 1. Now we selected the coordinate system arbitrarily and aligned the x and y to the side walls in the lower left hand corner. That just happens to be convenient. And then all of our dimensions are easy to understand.

But what if I have the irregular shaped area shown in the right in Figure 1. Now the selection of the reference coordinate system can be selected in several locations and no one location appears to be better than another. With the reference selected at the upper most corner, then the description of the original reference location in the lower left corner is described differently because of the coordinate system selected.

This just means that the description of any point in the area will be described differently. But the point in the area will not change.

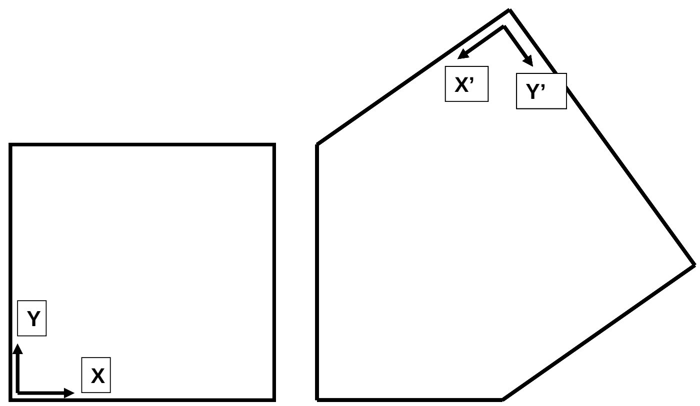

Figure 1 – Rectangular and Irregular Shaped Area

So using this simple geometry example will help set the stage for the discussion of a general mode shape that may be described in different ways depending on how the coordinate system is selected. Figure 2 and 3 show a description of the rigid body modes for a simple planar beam structure.

In the first case shown in Figure 2, the first two rigid body modes of the beam consist of a classic bounce mode and a rocking mode that occurs about the geometric center. This is exactly what everyone would expect those two modes to be. And if this were to occur, not one would question this at all.

But for the second case shown in Figure 3, the first two rigid body modes have a slightly different appearance. At first glance, most people would say that those rigid body modes were not correct. And that statement would only be made because it wasn't what you were expecting. You will notice that one mode is mainly bounce but has a little bit of rocking and that the other mode is mainly rocking but not about the geometric center.

While they may not look like what you would expect (or like) to see, these modes are perfectly correct. Because they are essentially at the same frequency, any linear combinations of these modes form a linearly independent set of vectors that are orthogonal with respect to the system mass and stiffness matrices.

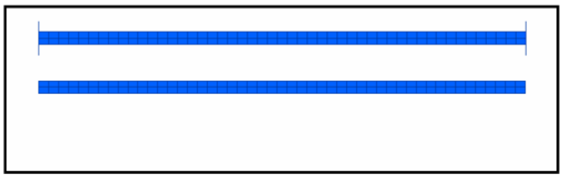

(a) Bounce Rigid Body Mode

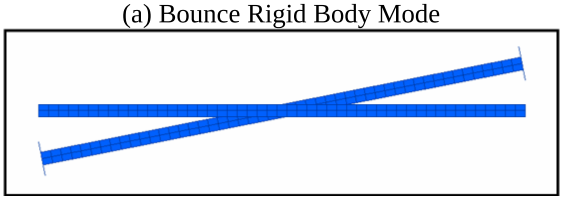

(b) Rocking Rigid Body Mode About Center
Figure 2 – Rigid Body Modes Geometric Center

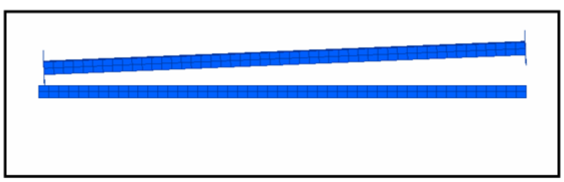

(a) Bounce Rigid Body Mode with Some Rotation

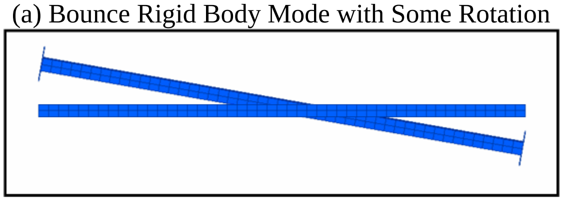

(b) Rocking Mode Offset From Center
Figure 3 – Rigid Body Modes Not About Geometric Center

This can also happen with the flexible modes of the system when the frequencies are repeated or pseudo-repeated. Figure 4 shows a set of modes that are pseudo-repeated – they occur at essentially the same frequency. These modes are seen as first bending and first torsion as expected. But these same modes are also seen in Figure 5 but they do not appear as simple bending and simple torsion. But these modes just have a different coordinate system to describe them. As long as the modes represent an orthogonal set of vectors then they are mathematically correct. They just may not be what you would expect to see.

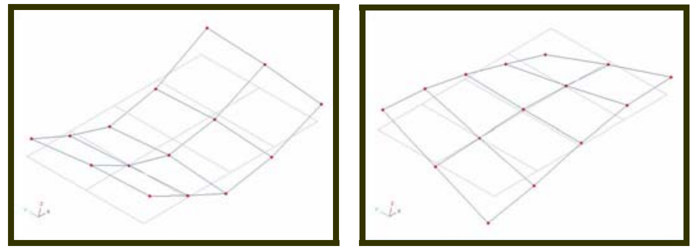

Figure 4 – Pure Bending and Pure Torsion Modes

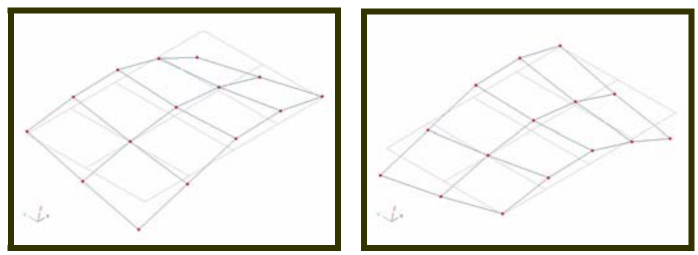

Figure 5 – Modes with Mixing of Bending and Torsion

This issue occurs with structures that have double symmetry and when either repeated roots or pseudo repeated roots occur. Another time it can happen is when using different numerical solution algorithms. Because the solution will typically iterate to a set of solution vectors, there is no reason why the vectors should converge towards a particular reference coordinate system. Actually the beam solutions shown in Figure 2 and 3 were obtained using two different finite element eigensolution approaches – one solution just happened to converge to the modes the way we would have expected them to occur whereas the other solution scheme did not. The modes in Figure 4 and 5 were obtained from actual test data on a structure that is known to have pseudo-repeated roots.

I hope this clarifies your confusion about modes shapes and their possible orientations. If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

A reference DOF is needed for all modal software systems. Why is the FFT analyzer reference different? This is another area of confusion.

OK – so this is another area of confusion that people often stumble upon and it causes grief when trying to perform a modal survey and obtain results. Usually this is often encountered when data is collected separately from the modal analysis software system using an independent FFT analyzer. But it can also happen when collecting data in any modal analysis system.

Before we get into specifics, let's just review a few basic items to try to put this problem in perspective. The problem really stems from the fact that the FFT analyzer is not used for just modal testing and the generation of mode shapes using a modal analysis software system. The FFT analyzer is a general purpose instrument that is used to generate frequency response functions for just about anything. The measurements can be general purpose signals for circuit analysis, acoustic measurements, transmissibility measurements, etc.

The concept of the analyzer is that measurements are made on two or more channels. The measurement ratio of output to input is typically what people are trying to obtain. That being said, the ratio of two signals is typically of interest.

Now let's say that a measurement is being made where the output voltage of a circuit (filter) needs to be obtained relative to the voltage applied to the circuit as seen in Figure 1 for instance. (Notice that these measurements are general and not necessarily the typical force and acceleration that are often obtained for a modal test.) So I could put the input voltage into channel 1 of the FFT analyzer and the output voltage of the circuit into channel 2 of the FFT analyzer. Or I could swap those two channels because it really doesn't matter.

But what does matter is that the frequency response measurement desired is that of the output voltage of the circuit "relative to" the applied input voltage. So in terms of the FFT analyzer, the output voltage spectrum is measured relative to a "reference signal" which is the input voltage.

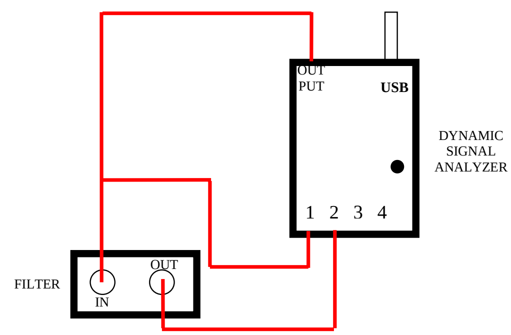

Figure 1 – Typical Input/Output Measurement Setup

For the FFT analyzer, the reference channel will depend on which channel was used to measure the input reference voltage – whether it be channel 1 or 2 or whatever. So the frequency response function measured might look something like that shown in Figure 2 where the output voltage of the filter is measured relative to the input signal.

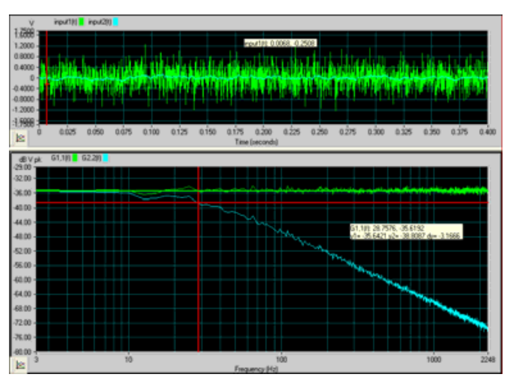

Figure 2 – Typical Input/Output FRF Measurement

Copyright 2008
Pete Avitabile

So the reference for the FFT analyzer is related to the measurements of output to input. Now what is meant by reference for an experimental modal test? Generally, the reference is the item in the measurement that doesn't change. The typical measurements made for a modal test will depend on whether an impact test or shaker test is performed. The terms of the FRF matrix obtained for each are shown in Figure 3.

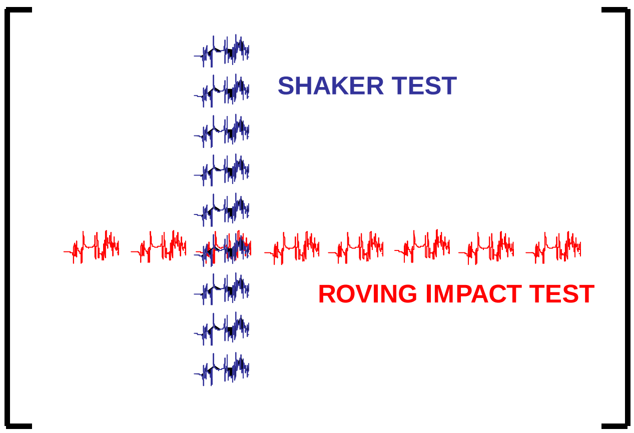

Figure 3 – FRF Matrix for Shaker and Impact Excitations

For a shaker test, the force is applied at the same location and the response is measured for all the measurement points desired. This is consistent with the FFT reference nomenclature.

But for a roving impact test, the hammer moves from one point to another but the accelerometer remains at a fixed location – so the accelerometer is called the "reference" for modal testing purposes but the FRF measured relates the output acceleration relative to the input force. So it is right here that the problem arises. The word "reference" means different things to different applications.

So depending on what FFT analyzer is used, there may be a procedure or recommended file naming convention that may need to be used in order to not "confuse" the modal software in regards to this "reference" notation.

Now of course everyone realizes that the measurements in the FRF matrix are reciprocal and that this reference notation is just an administrative procedure that needs to be addressed when data is transferred from an FFT analyzer to your particular modal analysis software package. But it is a frustrating administrative procedure that must be addressed in order to be able to use the measurements obtained from the FFT analyzer in the modal software package.

Once the proper procedure is identified, then it is a simple matter to document it so everyone realizes what needs to be done. I know that over the years I have generated many different schematics to remind myself of what procedure I needed to follow that particular day for the particular FFT analyzer that I was using that day. A typical (old) schematic is shown in Figure 4 where a file naming convention was used to clarify which measurement was which.

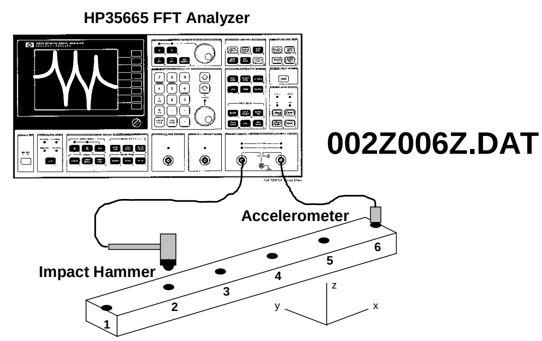

Figure 4 – Filename Convention Schematic

Now depending on the software package that you currently use, there are a variety of ways that the data may need to be handled in order to properly identify the "modal reference" for the measurements. Now this may get handled by a "swap" or "switch" command to obtain the proper "modal reference".

Or it may be required that the measurements be written out to a universal file format with a specific organization of the measurements – or the manner in which the measurements are imported into the software may handle this issue. This may be through the identification of whether the measurements come from a "roving" hammer or "stationary" hammer reference location.

But how you need to handle this will depend on your particular FFT analyzer and modal software system used. But I can guarantee you that each of the modal software vendors all have to face the same issue and will have specific procedures to handle this commonly encountered problem. I guess the amazing thing about this is that it is very clear what the problem is and it should be a very simple fix for the software to handle this problem but all the different software packages require this one obvious step to be performed.

I hope that this clears up the confusion on references – whether it be on the FFT analyzer or modal software package. If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Should more residuals be used to improve a curvefit? The results look better when many extra terms are used. This is an area that needs to be discussed..

When using frequency domain curvefitting techniques, many software packages allow the incorporation of extra terms in the polynomial in order to account for out of band effects. This is very useful in order to obtain accurate modal parameters. However, the user can specify many additional extra terms in order to improve the fit of the data. While this may "look" better, it is questionable where or not the parameters are actually better. So let's discuss the basic underlying equation and concept behind using residuals for modal parameter estimation. The basic frequency response equation can be written as

$$[H(s)] = \sum_{k=1}^{m} \frac{[A_k]}{(s-s_k)} + \frac{[A_k^*]}{(s-s_k^*)}$$

Now if we only write this equation over a band somewhere in the middle of the frequency response function, then there will be three different terms – one for the terms below the band of interest, the band of interest and one for the terms above the bands of interest. This is written as

$$[H(s)] = \sum_{\text{terms}}^{\text{lower}} \frac{[A_k]}{(s-s_k)} + \frac{[A_k^*]}{(s-s_k^*)}$$

$$\sum_{k=i}^{j} \frac{[A_k]}{(s-s_k)} + \frac{[A_k^*]}{(s-s_k^*)}$$

$$\sum_{\text{terms}}^{\text{upper}} \frac{[A_k]}{(s-s_k)} + \frac{[A_k^*]}{(s-s_k^*)}$$

And we often write this equation with only the modes of interest, over the band of interest, and apply extra terms called residuals to compensate for out of band effects and is written as

$$[H(s)] = LR + \sum_{k=i}^{j} \frac{[A_k]}{(s-s_k)} + \frac{[A_k^*]}{(s-s_k^*)} + UR$$

A typical frequency response function illustrating this is shown in Figure 1.

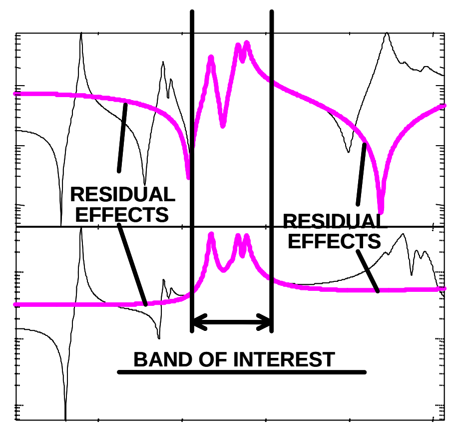

Figure 1 – Illustration of FRF with Band of Interest

In order to describe the residual terms, it is advantageous to look at the single degree of freedom displacement frequency response function. Figure 2 shows that frequencies below that of the resonant frequency are basically described by a dominant stiffness term and that the frequencies above that of the resonant frequency are basically described by a dominant mass term. It is this basic fact that allows the frequency response function to be written with the band of interest along with a lower residual term (LR) and an upper residual term (UR). Usually 4 extra residual terms in a polynomial curvefitter are sufficient in order to approximate these terms.

So now let's use a measurement to illustrate what happens when residual terms are overspecified to extract parameters. A simple 6 DOF model with a band of four modes bounded by two dominant modes will be used.

Now a curvefit for the four modes in the middle of the band is performed using the typical residual terms in most polynomial curvefitters (4 extra terms) and the fit is seen in Figure 3. Notice that the fit is reasonable but it doesn't fit the data well over all over all frequencies – at least from a visual perspective. Because the fit only used 4 extra residual terms, the next curvefit performed uses 10 extra residual terms and is seen in Figure 4. Now this fit appears better overall – from a visual standpoint anyway. And just to illustrate a point, the fit is also done with a simple SDOF shown in Figure 5.

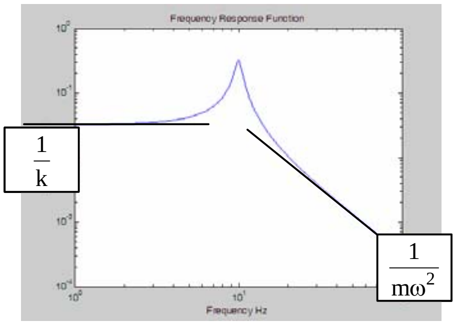

Figure 2 – Single DOF System with Residual Terms

But in order to really evaluate these fits, the extracted data needs to be compared to the actual parameters that were used to develop the frequency response functions. Table 1 lists the frequencies, damping and residues for the four modes along with the parameters extracted from both curvefit approaches.

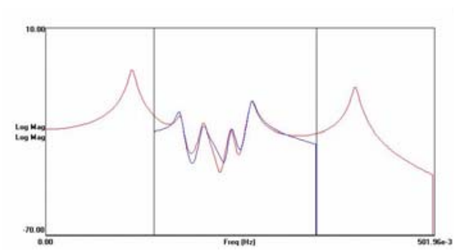

Figure 3 – Curvefit with 4 Residual Terms

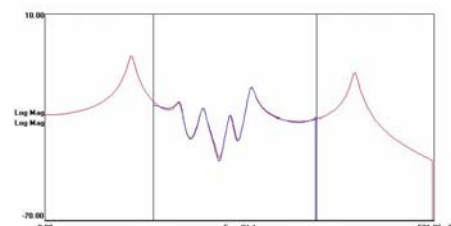

Figure 4 – Curvefit with 10 Residual Terms

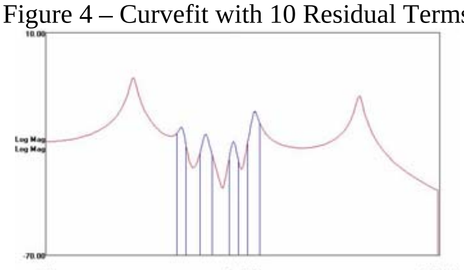

Figure 5 – Curvefit with SDOF Polynomial

Once the data in Table 1 is evaluated and assessed, it becomes clear that the addition of extra residual terms does not improve the parameter estimation overall and actually might degrade the results somewhat. Also note that the SDOF produces the best results overall. This brings up the point that the modal parameter estimation process is about extracting reasonable parameters to describe the system characteristics – and not necessarily about making curves that overlay on top of each other. In all the years of estimating modal parameters, it has become very clear that the overspecification of residual terms is only trying to compensate for noise or imperfections in the frequency response functions obtained. The overspecification of residual terms is not considered to be the reasonable approach for extracting modal parameters. The default residual terms specified in most commercially available software packages are reasonable for most curvefitting applcations. If many extra residual terms are needed to fit measured frequency response functions to "look better", then it is likely that the measured functions are contaminated with noise or imperfections and better measurements are likely needed.

Table 1 –Frequencies/Damping/Residues

Mode	Frequency Hz	Damping % Critical	Residue Value
Exact Analytical Results			
1	0.173	2.46	311
2	0.203	1.95	233
3	0.239	1.55	159
4	0.265	1.49	595
4 Modes Extracted With 4 Residual terms			
1	0.173	2.17	349
2	0.202	2.22	223
3	0.239	1.65	149
4	0.265	1.51	596
4 Modes Extracted With 10 Residual terms			
1	0.173	2.66	348
2	0.203	1.91	231
3	0.238	1.43	137
4	0.265	1.50	584
4 Modes Extracted With SDOF Approach			
1	0.173	2.30	314
2	0.203	1.96	234
3	0.239	1.68	159
4	0.265	1.50	594

I hope that these simple cases illustrate some important points regarding modal parameter estimation. Overspecifying residual terms is not the preferred approach for extracting accurate parameters. If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

Illustration by Mike Avitabile

My coherence is better in some measurements than others when impact testing. Am I doing something wrong? There are definitely some issues to discuss here.

OK – so this yet another area of measurement quality that needs to be discussed. Impact testing is by far the most common and most popular of the approaches for obtaining frequency response functions for the description of a structural system. The impact test is a very economical approach for frequency response testing. In addition, impact testing is very easy to setup and is extremely portable for field testing. Due to the ease with which measurements can be made, impact testing is widely used in many industries and applications.

But there are a wide range of issues that need to be recognized when performing impact testing. Some of these relate to double impacts, pre-trigger delay, high peak voltages compared to overall RMS level of the signal, nonlinear systems, etc. Some of these are commonly cited “areas of concern” when impact testing. These often become the stated reasons why impact test results may have coherence values that are not as acceptable as may be desired. But these may not be the only reasons – one very important consideration that I would like to discuss in this article is the effect of impact location on the resulting frequency response function and its coherence.

When performing impact testing, the input impact location can have a very significant effect on the resulting frequency response function. And this can be seen in the coherence function measured for each set of averaged data. First, let’s take a set of measurements where care is exercised in the impact location during the test to show a very good high quality measurement. Then some “less than perfect” impact measurements will be made on the same structure to show the degradation of the coherence.

The structure is a very simple structure with what are expected to be some very good measurements. A typical impact measurement is going to be made for the frequency response function at the drive point on the structure. Sampling parameters are selected such that the input force and response acceleration are totally observed signals within one sample record of data. This eliminates the need for any window functions on the input or output signals measured.

The measured frequency response (lower trace) and coherence (upper trace) are shown in Figure 1. Notice that the frequency response function appears to be a very good measurement and the coherence is very good for this measurement. The coherence for most of the frequency range is extremely close to one. The coherence has a slight dip in antiresonant regions but is not a problem for this measurement. (Note that drops in the coherence in antiresonant regions are expected due to the fact that the structure has no response at these frequencies and therefore the response of the system is not coherently related to the measured input signal.)

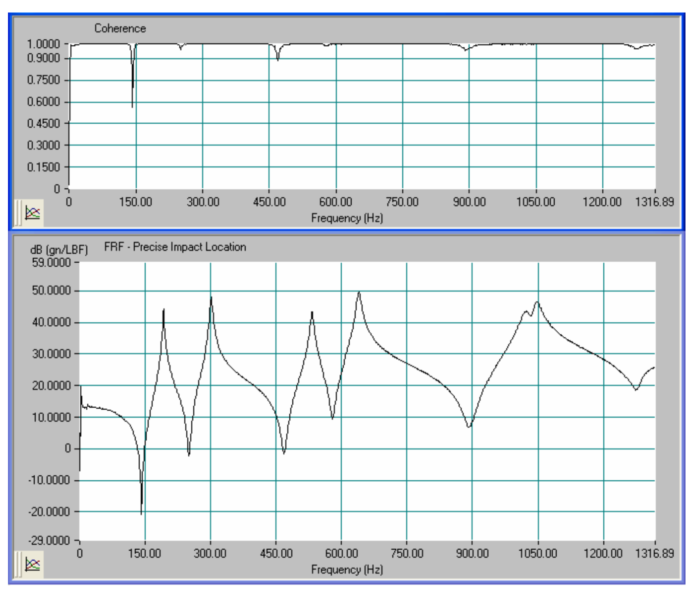

Figure 1 – FRF & Coherence for a Well Controlled Impact

Now in that first measurement, extreme care was exercised to assure that each average was the result of an impact at the same location in the same direction. This is a very important concern when impact testing.

To illustrate what happens where this care is not exercised, a measurement is made where each average is intentionally made within a region that is very close to the desired input location but there is some slight variation in the actual input location. With the same number of averages, the frequency response (lower trace) and coherence (upper trace) are shown in Figure 2 for this measurement where there is some variation for each impact location. While the frequency response function looks reasonable, the coherence is seen to have some significant degradation across the entire frequency range. While the coherence is acceptable in the immediate region of the peaks of the frequency response function, overall the coherence is poor.

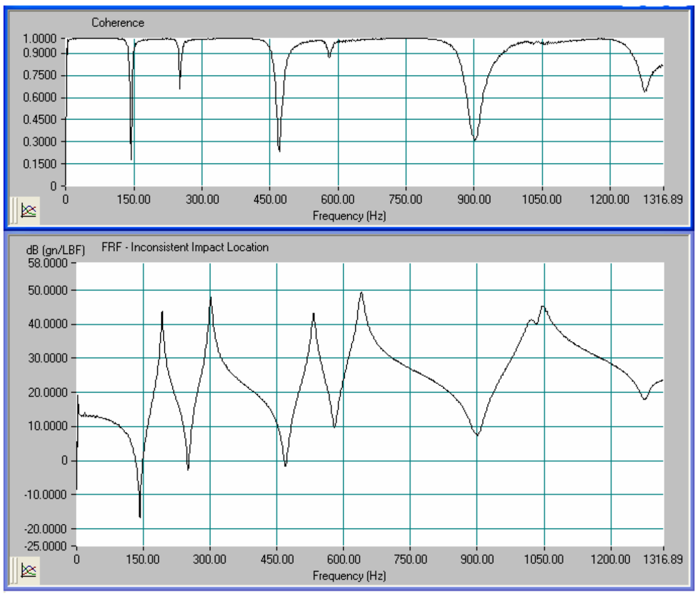

Figure 2 – FRF & Coherence for a Poorly Controlled Impact

The coherence is most significantly affected in the antiresonant regions of the frequency response function. This is due to the fact that while resonances are global characteristics of a system, the antiresonant regions are absolutely not global in character at all. The antiresonant regions are highly dependent on the particular input-output measurement location. Because care was not exercised during the impact test to assure than all impacts were made at the same location, the antiresonant region changes for each input output measurement that makes up the total average for the measurement. Therefore, from one measurement location to the next there is no consistency in the measurement and therefore the coherence reflects this.

One additional set of averages was made where the impact point was kept the same but the angle of the impact excitation was allowed to vary during each of the averages. The frequency response (lower trace) and coherence (upper trace) are shown in Figure 3 for this measurement. Similar to the previous case, the coherence is also degraded. There is also a lack of consistency in the antiresonant regions for this measurement.

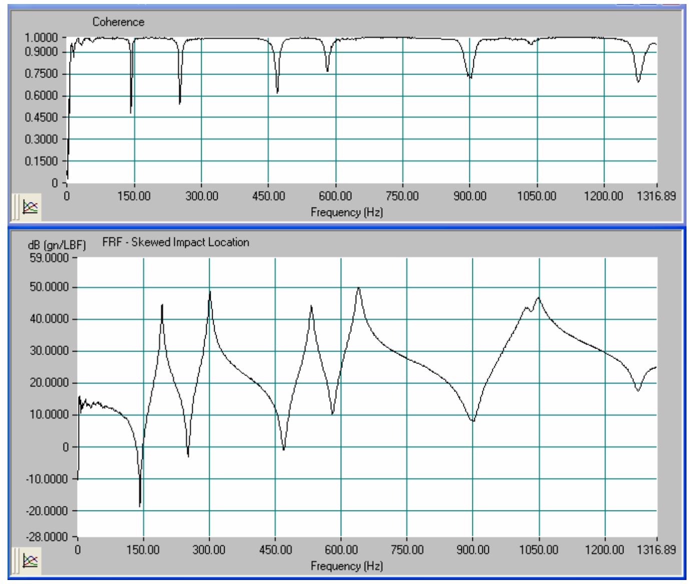

Figure 3 – FRF & Coherence for a Skewed Impact

For both of the cases shown in Figure 2 and 3, the coherence is not nearly as good as the measurement shown in Figure 1. This is due to the inconsistency of the impact location – whether it be not impacting the same location for each measurement or for not maintaining a consistent strike angle for each measurement. Both cases clearly show a degradation of the measurement coherence. A very well controlled, precise impact excitation needs to be maintained for each average that makes up the complete measurement.

These cases are presented here because this is a very common problem during impact testing. This is especially true when the test lasts for a long period of time for measuring many locations. Generally, as time goes on it is very easy to become bored and not maintain the consistent impact during the entire test. This is also very common when the impact locations are at inconvenient locations around, on top or underneath the test structure. When climbing all around the structure (and often in very unnatural positions), it is very easy to not maintain a consistent impact for all averages making up a measurement.

So as a word of caution when impact testing… be very sure to impact the same point, in the same direction, for each of the averages that make up the frequency response function to assure that an overall acceptable coherence is obtained for all measurements.

I hope that this clears up the concerns about possible coherence degradation when impact testing. While there are many more items that could affect coherence, this is one that has an effect. If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

Illustration by Mike Avitabile

My accelerometer is not overloaded but my measurement is terrible. What could be wrong? Some discussion of this is needed here.

OK – there can be many things that might cause this problem. The measurements can be contaminated by a variety of sources. Many different types of problems may be encountered in different situations. But in this particular case you have a very strange problem from the measurement that was provided. At first glance, the structure seems to be one that can be tested with little problem.

Let's start with a different structure and recreate the measurement problem that actually existed in your measurement system. For the structure here, a simple plate was instrumented with an accelerometer and subjected to impact testing. Three different cases will be shown to show what could have happened with the measurement.

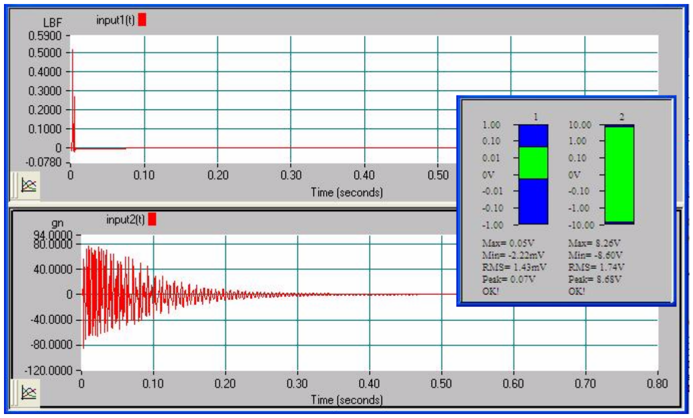

Figure 1 – Excitation (top) and Response (bottom) with Sensitive Accelerometer and Exponential Window for Case 1

Case 1 – Sensitive Accelerometer with Exponential Window

In the first measurement, an impact excitation was used. A very sensitive accelerometer was used and because leakage may be a problem, an exponential window was used for this measurement. Figure 1 shows the input excitation and the response from the accelerometer. Also shown in Figure 1 are the ADC range settings that resulted from the measurement. The measurement looks reasonable and there doesn't appear to be any problem with the time measurement.

However, looking at the frequency response function and the coherence in Figure 2, the measurement looks terrible indeed. The measurement has no real useful information anywhere in the frequency range shown. Clearly, this measurement is not good at all.

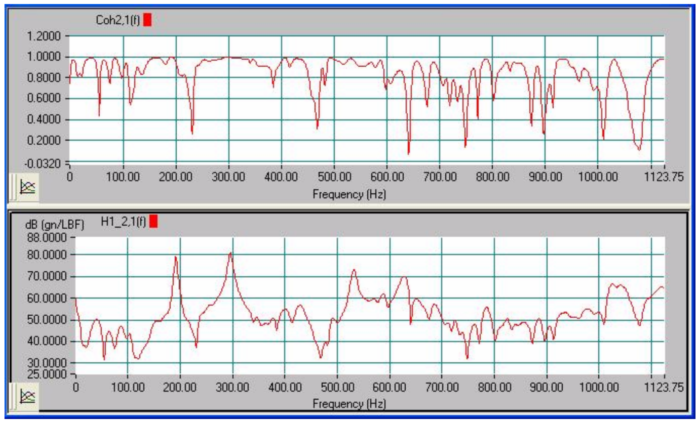

Figure 2 – FRF (bottom) & Coherence (top) with Sensitive Accelerometer and Exponential Window for Case 1

Copyright 2008
Pete Avitabile

Case 2 – Sensitive Accelerometer with No Window

In the second measurement, an impact excitation was used again but no window was applied to the response window to see if there was any additional information that could be seen.

Figure 3 shows the input excitation and the response from the accelerometer. Also shown in Figure 3 are the ADC range settings that resulted from the measurement. There doesn't appear to be any overload with the time measurement.

Again, looking at the frequency response function and the coherence in Figure 4, the measure still looks terrible.

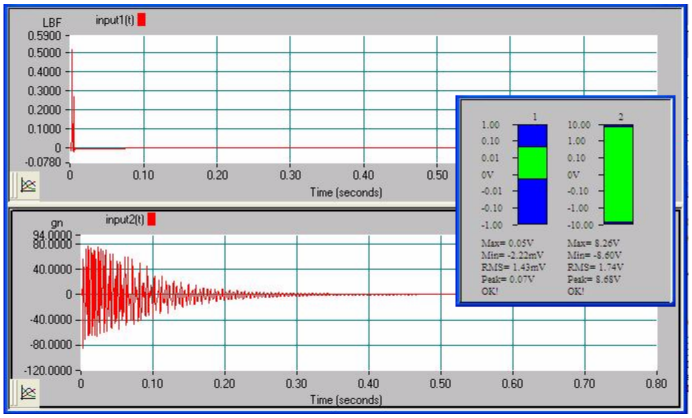

Figure 3 – Excitation (top) and Response (bottom) with Sensitive Accelerometer and Exponential Window for Case 2

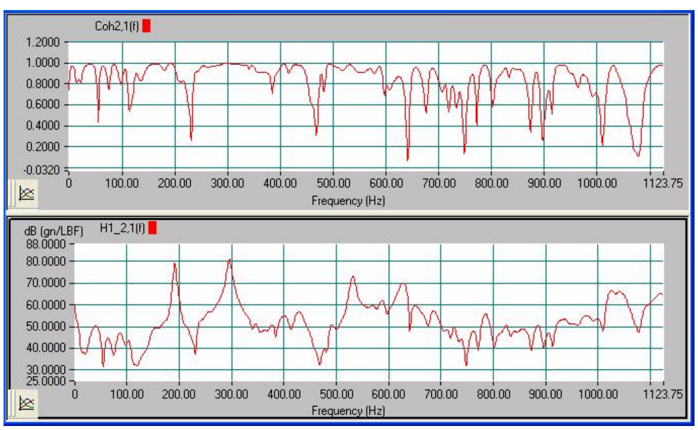

Figure 4 – FRF (bottom) & Coherence (top) with Sensitive Accelerometer and Exponential Window for Case 2

But looking at the time trace, the response does not appear to be what would be expected for a second order exponentially decaying system. What has actually occurred here is the accelerometer response was so large that it saturated the accelerometer response causing it to respond in a nonlinear fashion. During the first 0.05 seconds of time response, the system does not appear to respond in an exponential fashion. But the interesting part is that the total accelerometer voltage output was not greater than 10 volts and therefore did not overload the ADC of the acquisition system!

Case 3 – Less Sensitive Accelerometer with No Window

In the third measurement, an impact excitation was used again but no window was applied and a less sensitive accelerometer was used for the measurement. Now the time response in Figure 5 and frequency response in Figure 6 looks like what was expected.

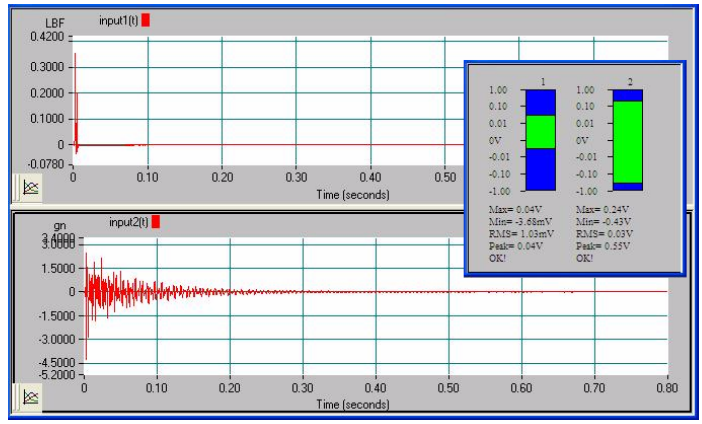

Figure 5 – Excitation (top) and Response (bottom) with Sensitive Accelerometer and Exponential Window for Case 3

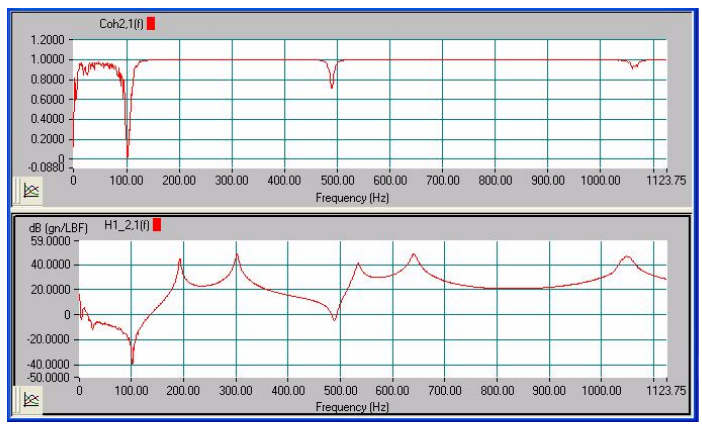

Figure 6 – FRF (bottom) & Coherence (top) with Sensitive Accelerometer and Exponential Window for Case 3

The problem in this case was that too sensitive an accelerometer was used for the impact test. While the FFT analyzer ADC did not overload, the accelerometer was saturated by the large response; this caused a response that was far different from the damped exponential response expected. So it is very important to look at all the various pieces of the time and frequency measurements made.

I hope that this sheds additional light onto this measurement problem. You not only have to worry about the measurement system but also the transducers used to make the measurement. If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

Illustration by Mike Avitabile

Double impacts are undesirable. What about multiple impacts?
Ahhh… Now this is something that we have to discuss.

We have discussed double impacts before and have shown that while they are undesirable, they may be unavoidable in many cases. In fact, previously we showed that the double impact measurement wasn't necessarily as bad as most people profess. Of course, the overall measurement, including the frequency response and the coherence, must be checked along with the averaged spectrums for the measurement.

Now the question here is really if multiple impacts can be used as an excitation technique and if there is any problem using a measurement made from multiple impacts.

This is actually a very good question and needs to be thought through carefully. An impact measurement typically is the result of a single impact; the response due to that impact is generally a damped exponentially decaying response.

Now if we were to consider an arbitrary input force, then that signal can be thought of as a series of impulses added together spaced delta t seconds apart in order to characterize the input. In fact, this is the way that arbitrary signals are handled in any vibrations text book – the solution method is called the superposition method, or convolution integral, or Duhamel's integral – and is used to compute arbitrary response of any system.

In this case, the series of pulses will be applied to the structure. But some care needs to be used here. The impulses should be applied in a very incoherent fashion in terms of their timing and spacing. The pulses should also not be applied for the entire sample period. They should be applied for a portion of the sample interval, 50% to 75% for instance. But it is also important for the response to be totally observed within the sample interval so that no leakage will occur.

In this way, all the requirements of the Fourier transform are satisfied. In fact, the signal will start to approach a broad band excitation with characteristics similar to that of a random signal like a burst random.

A simple structure is used to illustrate the technique. Due to the responsive nature of the structure, double impact measurements are unavoidable but they are not serious enough so as to corrupt the measurement overall.

In the first case, a single impact measurement is applied – or least the intent is to apply a single impact. Figure 1 shows the time signals for the impact and response. Figure 2 shows the input power spectrum with the frequency response. Figure 3 shows the frequency response function along with the coherence. Overall the measurement is good but the effects of double impact are seen in the input time excitation and the input spectrum noted by a varying input spectrum. The variation of the input spectrum is small enough so as to not distort the overall measurement for the system as evidenced by the coherence.

In the second case, a series of impact measurements were applied to the structure. Figure 4 shows the time signals for the impact and response. Figure 5 shows the input power spectrum with the frequency response. Figure 6 shows the frequency response function along with the coherence. While multiple impacts were applied, the overall measurement is very good. The resulting frequency response and coherence are very good.

I hope that this shows that multiple impact excitations can in fact be used to excite the structure and measure good overall response functions. If you have any more questions on modal analysis, just ask me.

Copyright 2008
Pete Avitabile

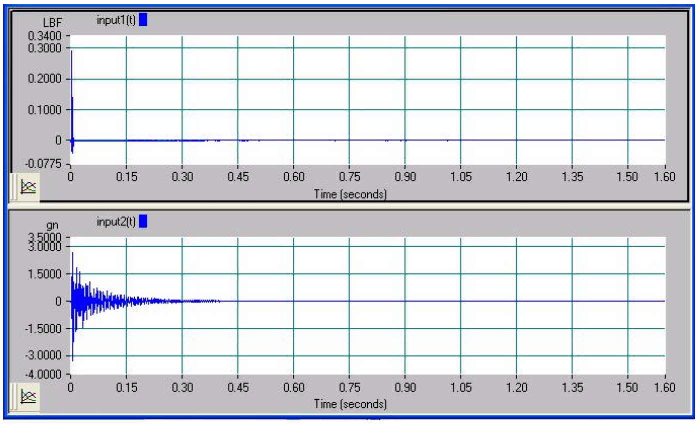

Figure 1 – Excitation (top) and Response (bottom)
with Single Impact Excitation for Case 1

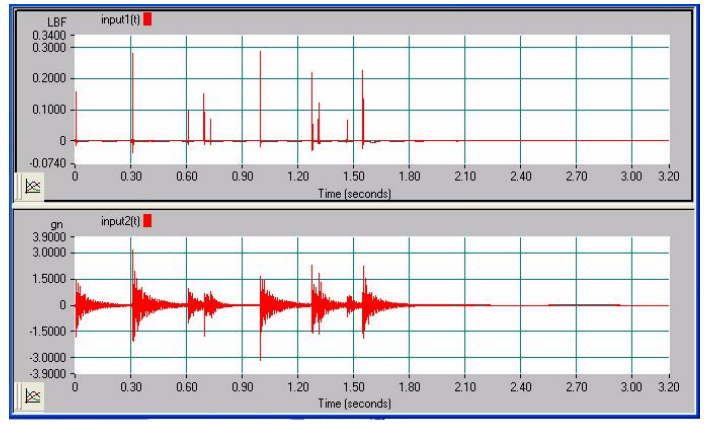

Figure 4 – Excitation (top) and Response (bottom)
with Multiple Impact Excitation for Case 2

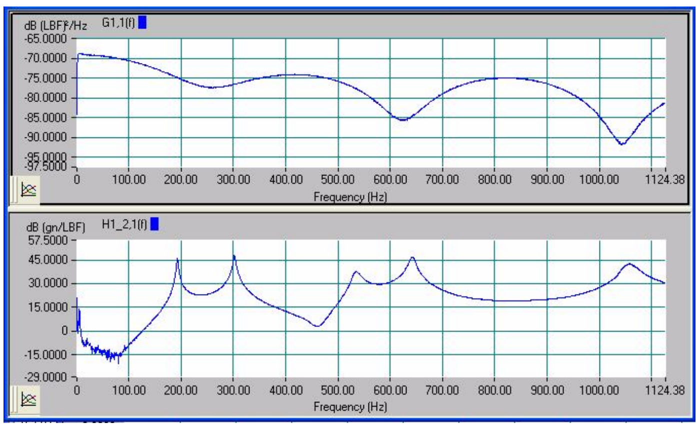

Figure 2 – FRF (bottom) & Input Power (top)
with Single Impact Excitation for Case 1

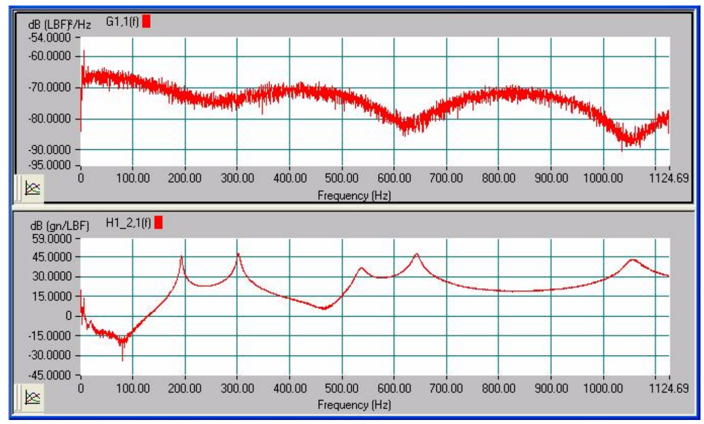

Figure 5 – FRF (bottom) & Input Power (top)
with Multiple Impact Excitation for Case 2

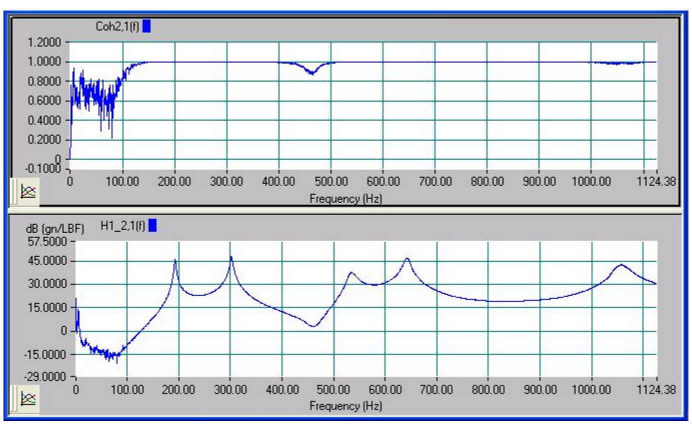

Figure 2 – FRF (bottom) & Coherence (top)
with Single Impact Excitation for Case 1

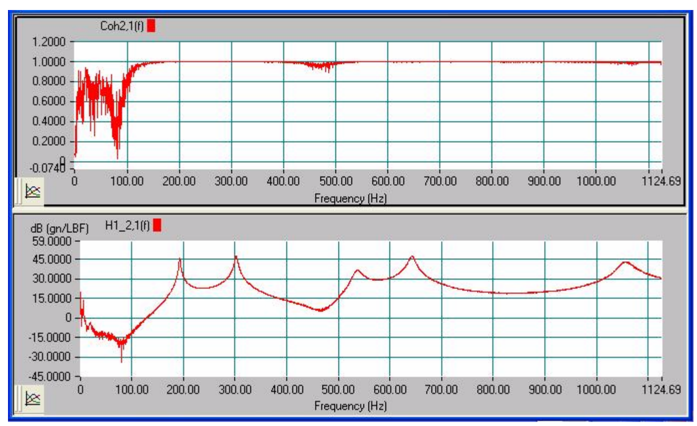

Figure 6 – FRF (bottom) & Coherence (top)
with Multiple Impact Excitation for Case 2

Copyright 2008
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Can you describe reciprocity? It just doesn't make sense to me.
This is something that often confuses people.

Alright – let's discuss the reciprocity of measurements when doing modal testing. This is a very important item when doing modal tests. People say the words but sometimes they really don't believe it – mainly because when we take measurements there are many reasons why the actual measurement may not satisfy the theoretical requirement of reciprocity.

Let's first simply state what reciprocity is. Figure 1 shows a structure where an input-output measurement is to be made at point "i" and point "j". Now in one measurement the force is applied at point "i" and the response is measured at point "j". And in the second measurement, the force is measured at point "j" and the response is measured at point "i". From the principle of reciprocity, the h_{ij} must equal h_{ji}

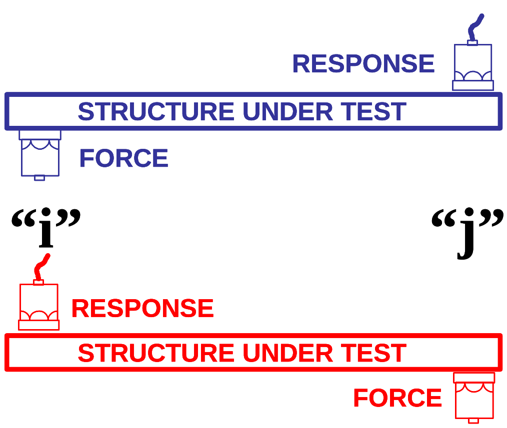

Figure 1 – Schematic for Reciprocity Measurement

From the complete set of measurements possible, Figure 2 shows a frequency response matrix where one row and one column are measured. Several reciprocal measurements are highlighted in that matrix for reference.

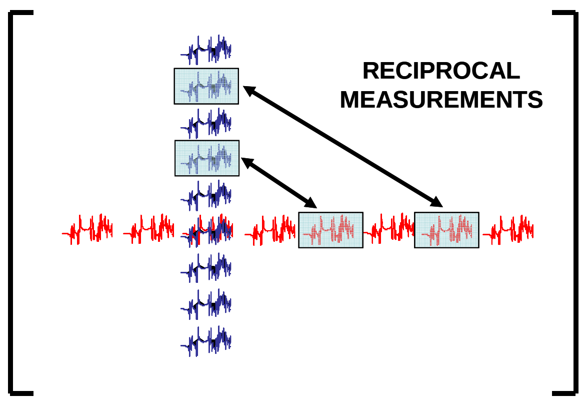

Figure 2 – FRF Matrix Showing Reciprocal Measurements

Of course, the first time you try to explain reciprocity to someone who is not familiar with this concept, it always seems to raise an eyebrow. So let's try to show where reciprocity comes from in the basic equations describing the system. (Some theory will have to be presented here to show reciprocity)

First let's realize that we start from an equation of motion written in matrix form for a multiple degree of freedom system as:

$$[M]\{\ddot{x}\} + [C]\{\dot{x}\} + [K]\{x\} = \{F(t)\}$$

Now the important point to note here is that these matrices are square symmetric (for a structural system). This immediately implies that the "ij" and "ji" terms of the matrix are the same.

Now let's write the equation of motion in the Laplace domain for that physical equation of motion written above. Assuming initial conditions are zero, then this is:

$$\left[[M]s^2 + [C]s + [K]\right]\{X(s)\} = \{F(s)\}$$

Copyright 2008
Pete Avitabile

Of course we have to realize that each of the terms of this matrix is also square symmetric. From this Laplace equation of motion, the system matrix [B(s)] and its inverse, the system transfer function [H(s)], is also square symmetric. This is:

$$\left[B(s)\right]^{-1} = \left[H(s)\right] = \frac{\text{Adj}\left[B(s)\right]}{\det\left[B(s)\right]} = \frac{\left[A(s)\right]}{\det\left[B(s)\right]}$$

Now with some manipulation, the system transfer function can be written in partial fraction form as the summation of all the individual modes of the system. This is:

$$\left[H(s)\right] = \sum_{k=1}^{m} \frac{\left[A_k\right]}{\left(s - p_k\right)} + \frac{\left[A_k^*\right]}{\left(s - p_k^*\right)}$$

The frequency response function is the system transfer function evaluated at s-jω and is given as:

$$\left[H(s)\right]_{s=j\omega} = \left[H(j\omega)\right] = \sum_{k=1}^{m} \frac{\left[A_k\right]}{\left(j\omega - p_k\right)} + \frac{\left[A_k^*\right]}{\left(j\omega - p_k^*\right)}$$

Now it is important to remember that the residue matrix [A(s)] is also square symmetric because all the matrices used to ultimately form it were square symmetric.

Now a single "ij" measurement can be written as:

$$h(s)_{ij}\Big|_{s=j\omega} = h_{ij}(j\omega) = \sum_{k=1}^{m} \frac{a_{ijk}}{(j\omega - p_k)} + \frac{a_{ijk}^*}{(j\omega - p_k^*)}$$

and expanding the first three mode terms for this gives:

$$\begin{aligned} h_{ij}(j\omega) = {} & \frac{a_{ij1}}{(j\omega - p_1)} + \frac{a_{ij1}^*}{(j\omega - p_1^*)} \\ & + \frac{a_{ij2}}{(j\omega - p_2)} + \frac{a_{ij2}^*}{(j\omega - p_2^*)} \\ & + \frac{a_{ij3}}{(j\omega - p_3)} + \frac{a_{ij3}^*}{(j\omega - p_3^*)} \end{aligned}$$

But in this form it is not clearly obvious that reciprocity exists. So the residue form of the equation needs to be extended. Recall that the residue matrix for the kth mode of the system can be obtained from singular value decomposition and written as:

$$\left[A(s)\right]_k = q_k \{u_k\}\{u_k\}^T$$

or expanded as:

$$\begin{bmatrix} a_{11k} & a_{12k} & a_{13k} & \cdots \\ a_{21k} & a_{22k} & a_{23k} & \cdots \\ a_{31k} & a_{32k} & a_{33k} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} = q_k \begin{bmatrix} u_{1k}u_{1k} & u_{1k}u_{2k} & u_{1k}u_{3k} & \cdots \\ u_{2k}u_{1k} & u_{2k}u_{2k} & u_{2k}u_{3k} & \cdots \\ u_{3k}u_{1k} & u_{3k}u_{2k} & u_{3k}u_{3k} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

and in this form is simply written as:

$$h(s)_{ij}\Big|_{s=j\omega} = h_{ij}(j\omega) = \sum_{k=1}^{m} \frac{q_k u_{ik} u_{ij}}{(j\omega - p_k)} + \frac{q_k^* u_{ik}^* u_{jk}^*}{(j\omega - p_k^*)}$$

and can be expanded for the first three mode terms as:

$$\begin{aligned} h_{ij}(j\omega) = {} & \frac{q_1 u_{i1} u_{j1}}{(j\omega - p_1)} + \frac{q_1 u_{i1} u_{j1}^*}{(j\omega - p_1^*)} \\ & + \frac{q_2 u_{i2} u_{j2}}{(j\omega - p_2)} + \frac{q_2 u_{i2} u_{j2}^*}{(j\omega - p_2^*)} \\ & + \frac{q_3 u_{i3} u_{j3}}{(j\omega - p_3)} + \frac{q_3 u_{i3} u_{j3}^*}{(j\omega - p_3^*)} \end{aligned}$$

Now in this form the reciprocity can be very easily seen. This is because the residue is nothing more than the value of the mode shape at the ith degree of freedom times the value of the mode shape at the jth degree of freedom (plus a few other terms that are constant). This implies that it doesn't matter whether we measure force at point "i" or point "j" when we measure the response at the other point – the product of the mode shape values at the input and output location is the same. So reciprocity must happen for this case. As an example, a simple 3x3 FRF matrix is shown in Figure 3 for magnitude plots. Clearly, the reciprocity can be seen in the matrix. Note that the real, imaginary and phase are also symmetric but not shown here for brevity.

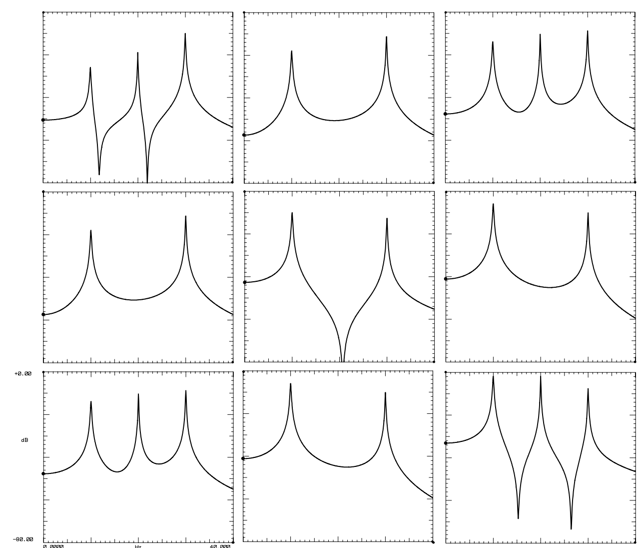

Figure 3 – Magnitude FRF Matrix

Of course this is a theoretical presentation and must hold true. However, measurements may not be so cooperative all the time. This will be discussed at some future point in time.

I hope that this discussion clears up any confusion as to why reciprocity must hold true. While there were some theoretical equations presented, they are necessary in order to show that reciprocity must hold true. If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Do we really need an accurate updated model? What are the effects if it is not perfect?
I have some very good examples to illustrate this.

Model updating is an important step in the development of an accurate system model. If any of the components of the system are not modeled correctly, then the overall system characteristics will not be accurate. Of course someone has to define what is acceptable and unacceptable in terms of response for the system. But that is a different item to address.

What I want to talk about here is the effect of a component frequency in relation to the system response. Components are typically much easier to model and update than the overall system. In fact, the biggest problems in developing a system model are the boundary conditions and the interaction of components in a system model. So the thing that I want to address here is the relationship of components with each other in a system model. Of course there are many ways to write all of those relationships. What I want to do here is identify the relationship from a very simple representation of the various systems. First the simple single degree of freedom tuned absorber will be considered and then the mode complicated multiple degree of freedom system will be addressed.

So the first thing that has to be discussed is the simple representation of a component in terms of its mass and stiffness, or its modal characteristics or its frequency response characteristic. Figure 1 shows a conceptualization of a finite element model of a component which is described in terms of its mass and stiffness. But in this form, it is not easy to interpret how the various modes affect the component overall. An eigensolution of the component reduces the complicated mass and stiffness into more simplistic set of single degree of freedom systems which are linearly independent and orthogonal with each other. The lower portion of Figure 1 shows the component as a set of single degree of freedom systems as well as a set of frequency responses for each of the modes of the system. So the component can be best understood if the modal characteristics are understood.

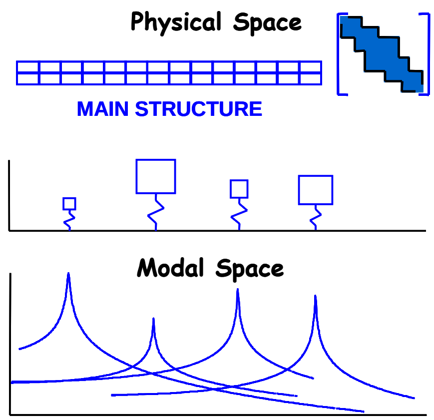

Figure 1 – Physical and Modal Representation for a Component

Now what would happen to the component if another spring-mass system were attached (but for now let's assume that it is not aligned with any frequency of the system)? Figure 2 adds that spring-mass system and the effects are very minimal on the original component modes. But what if the spring-mass system is "tuned" to a particular mode?

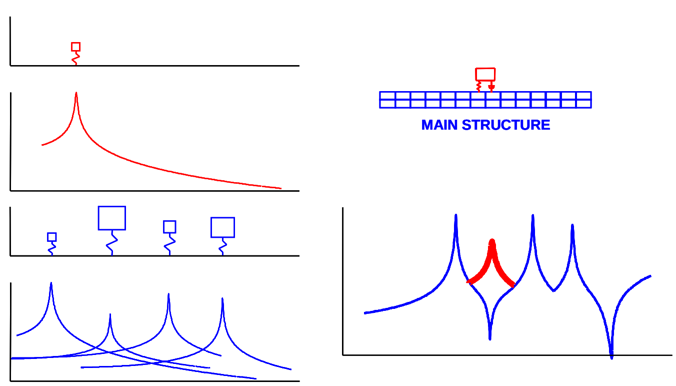

Figure 2 – Component with Untuned Spring-Mass System

Copyright 2009
Pete Avitabile

Figure 3 shows the effects on the frequency response if the spring-mass system is coincident with one of the modes and Figure 4 shows the effects if two spring-mass systems are added to the component. These effects are exactly what are expected for a tuned mass-spring absorber. There is a dynamic interaction between the added spring-mass system and the component modes of the system.

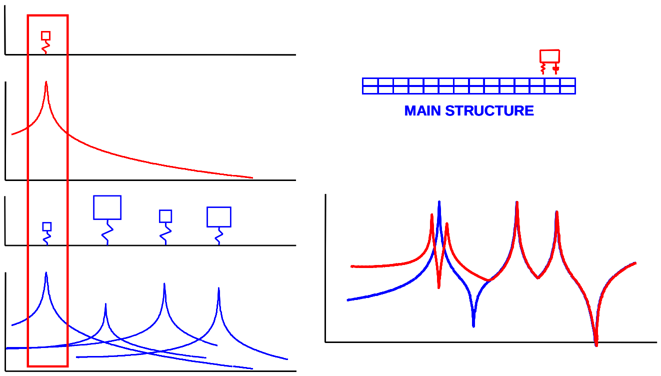

Figure 3 – Component with Tuned Spring-Mass System

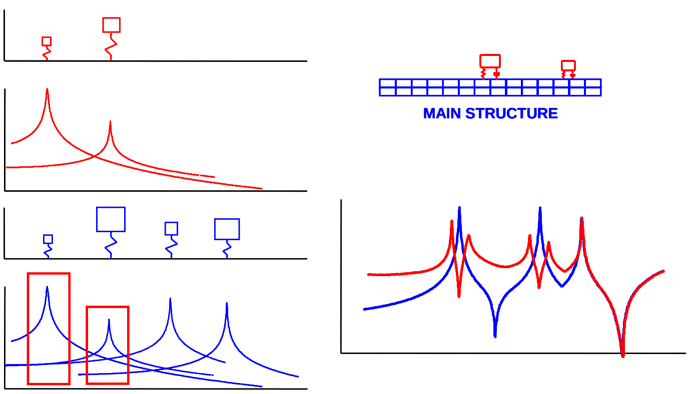

Figure 4 – Component with Two Tuned Spring-Mass Systems

So from these two schematics in Figure 3 and 4, it is very clear that if the frequencies of the spring-mass system are specifically selected to be coincident with one of the modes of the system, there will be a dramatic change in the dynamic characteristics of the component. However, if the frequencies of the spring-mass system are not selected correctly (as in Figure 2) then there will be no significant dynamic coupling between the added spring-mass system and the component.

Now that we have that concept in place, let's consider the coupling of two components to form a system model. But instead of considering the two components as mass and stiffness matrices, it is much more advantageous to consider them as either a collection of modal mass/spring systems in modal space or as a set of single degree of freedom response functions.

Figure 5 shows this representation for Component A and Component B. In order for the proper coupling of the two components to occur, the modes of each of the components must be properly specified. In Figure 5, the modes of Component A are not close to the modes in Component B, and therefore, there will not be significant coupling between the two components.

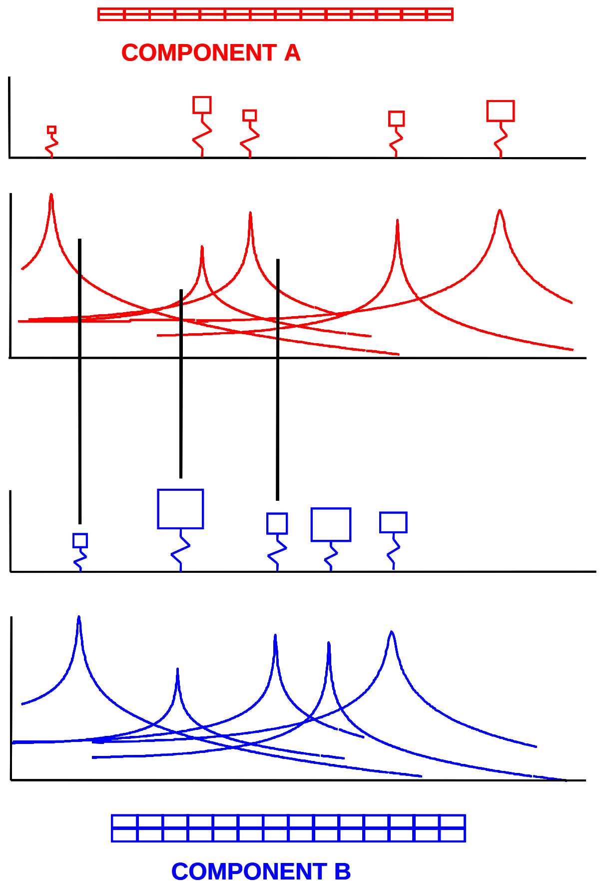

Figure 5 – Physical Coupling of Two Components

But what if the modes of either one of the components are not correct and the actual modes are much closer to each other? Then there should be a significant amount of dynamic coupling between the two components wherever the modes of both are aligned. The coupling between the two components is heavily dependent on the relative relationship of the frequencies of the two systems. Therefore, it is imperative that the modes of each component be properly identified so that the correct dynamic interaction exists in the assembled system model.

I hope that this helps to show why the modes of each component must be identified correctly. The component modes must be updated to properly identify the dynamic characteristics of the component. In just considering the mass and stiffness matrices, it is not apparent why the modes must be identified correctly. By representing the component in the modal or frequency domain, the need for updating the component models is much more obvious. If you have any more questions on modal analysis, just ask me.

Copyright 2009
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

Illustration by Mike Avitabile

I made a stiffness modification to a free-free system. The flexible modes shifted down! What's up?
Now this needs to be discussed.

Alright – that's a pretty bold statement that will turn almost anyone's head. I think we need to first describe what actually happened in this particular case. But rather than using the specific data originally presented, a simple beam can be used to describe what happened in this case.

The way this problem was described was that a free-free system was tested and then the system was constrained to fix or constrain the corners of the system. When the actual modification was performed to constrain the free-free system, the modes obtained were lower than the flexible modes of the unconstrained system.

Of course, the first thing that everyone stated was that if you increase the stiffness of any system, then the modes must shift upwards because

$$\omega_{initial}=\sqrt{\frac{k}{m}}$$

and if the stiffness is increased then

$$\omega_{modified}=\sqrt{\frac{k+k_{added}}{m}}$$

So it stands to reason that the frequencies must shift upwards – and the fact that the frequencies were lower just doesn't make sense.

So let's start with a simple beam that was analyzed and tested in a free-free condition. The first several free-free modes were 164 Hz, 452 Hz and 888 Hz. The unconstrained modes of the planar beam are shown in Figure 1 for reference.

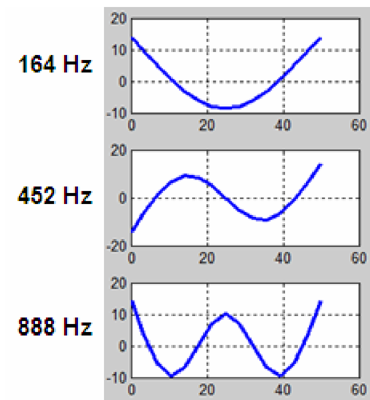

Figure 1 – Flexible Modes of Free-Free Beam

Then the simple beam was tested in a pinned (or constrained) condition. The first several free-free modes were 72 Hz, 288 Hz and 647 Hz. The constrained modes of the planar beam are shown in Figure 2 for reference. Clearly, the modes did not shift up as expected.

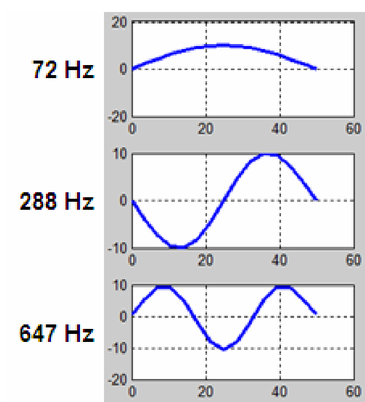

Figure 2 – Pinned Modes of Free-Free Beam

Copyright 2009
Pete Avitabile

So exactly what happened here???

This is a very simple problem. But I have heard this many times over many years and inevitably the same problem exists.

Typically, people are concerned about the flexible modes of the system because those are the modes that generally cause all the noise and vibration problems that we encounter. But those are not the only modes that are needed to describe the entire system.

The basic problem here is that everyone forgot that an unconstrained system has flexible modes AND the rigid body modes. Now most times people don't measure the rigid body modes in test and they don't include them as part of the flexible modes measured during a modal test. And from an analytical standpoint, many times the eigensolution is performed but either a shifted eigenproblem is solved or only the flexible modes are obtained.

While the rigid body modes exist, many times people ignore them – mainly because they are not the source of the noise and vibration problems that are of concern. So Figure 3 shows the set of modes for the beam to more correctly include the rigid body modes as well as the flexible modes. So now once we realize that the first modes are actually at 0 Hz from the analytical model or a very low frequency from a test, then the intuition that adding stiffness shifts the modes upwards makes proper sense. Table 1 shows the frequencies of the free-free beam with rigid body modes along with the pinned modes

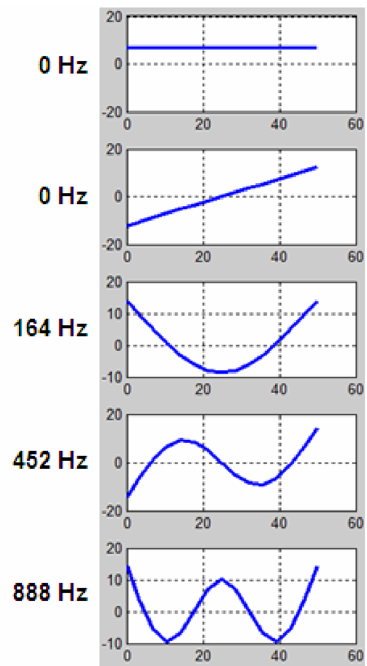

Figure 3 – Flexible Modes of Free-Free Beam with Rigid Body Modes

Table 1: Free-Free and Constrained Modes for Beam

Mode	Free	Constrained
1	0.	72.
2	0.	288.
3	164.	647.
4	452.	
5	888.	

So the basic problem is that the rigid body modes can't be ignored; they are a part of the total description for the beam. Notice now that all the frequencies in Table 1 do shift upwards as the stiffness is increased.

One way to easily prove this to yourself is to make a simple free-free beam model. The next model to develop is the beam with two very soft springs. Then make subsequent beam models where the spring stiffness is increased until ultimately the spring is so stiff that it is an approximation of a pinned end condition.

Along the way, it would be very beneficial to look at the mode shapes. When the springs supporting the beam are very soft, then the mode shape for the beam looks very much like a rigid body.

But as the stiffness of the support beam increases, the frequencies will increase and the mode shapes will start to migrate from rigid body modes to modes that have some rigid body mode component but also start to develop some more flexible attributes.

When the support spring stiffness gets larger and larger, the rigid body mode characteristic will diminish as the flexible characteristic becomes more pronounced. Ultimately, the rigid body characteristic will disappear and the flexible characteristic will completely dominate the mode shape.

This little exercise will then clearly show that the rigid body modes are critically needed to describe the modes of the system.

I hope this simple example clears up any misconceptions that you may have had. If you have any more questions on modal analysis, just ask me.

Copyright 2009
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

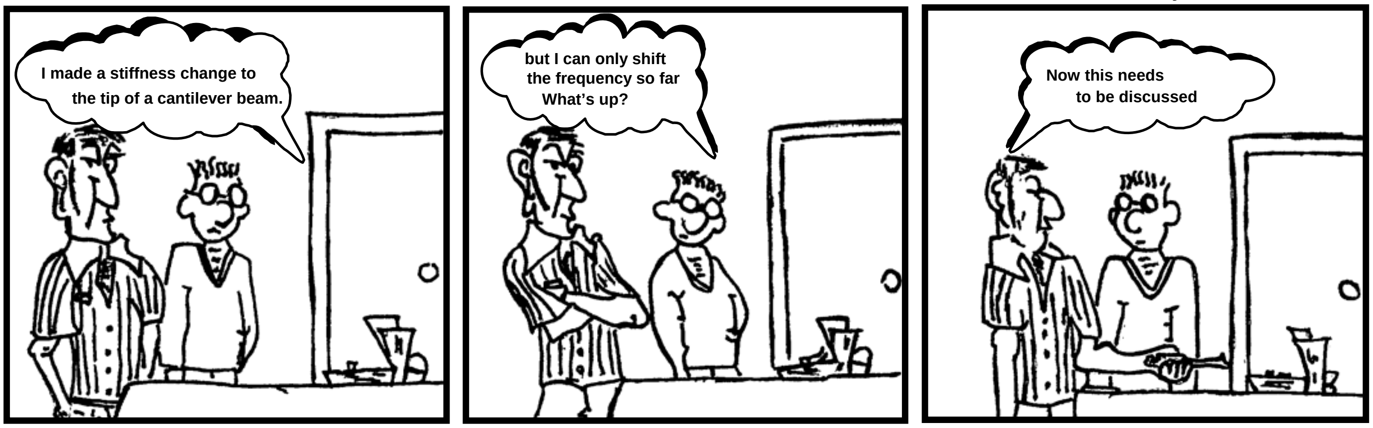

Illustration by Mike Avitabile

I made a stiffness change to the tip of a cantilever beam but I can only shift the frequency so far. What's up? Now this needs to be discussed.

OK. This is another one of those problems that I see many people get confused about. Let's start with a simple cantilever beam and explain some basic properties that are inherent in the system.

First, let's start with a simple finite element model to investigate the effects of stiffness at the tip of the cantilever beam. Figure 1 shows the cantilever beam along with the cantilever beam with a spring at the tip and the cantilever beam with the end pinned. A finite element model of the beam will be used to lend some insight into what happens when the spring at the tip of the beam is varied from low stiffness to high stiffness.

Table 1 shows the first three modes of the cantilever beam and then the change in frequency as the stiffness is increased along with the final pinned result if the spring was infinitely stiff. It is very important to notice that as the spring stiffness is increased, the final frequencies converge towards the final result where the cantilever is pinned at the tip.

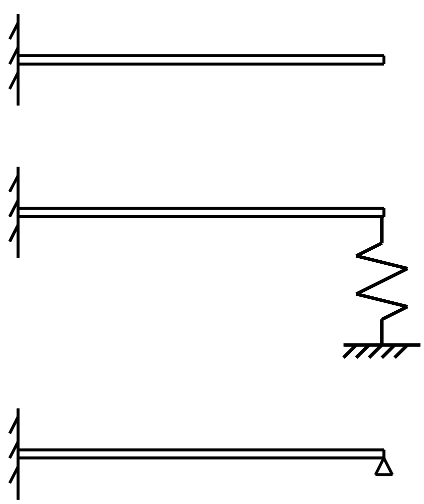

Figure 1 – Cantilever Beam, Cantilever Beam with Tip Spring, and Cantilever with Tip Pinned

So this implies that the no matter how much stiffness you add at the end of the cantilever beam, the frequency can only shift so far and then any additional increases in stiffness have very little effect at all – it is a point of diminishing returns.

Now let's further consider the simple cantilever beam and let's look at the tip response. The frequency response function is shown in Figure 2 with a drive point measurement at the tip of the beam where the stiffness is to be added to the beam.

So now let's look at the frequency response function and discuss the different parts of this function. At the natural frequencies, there is a peak in the function. Basically, this is a region in frequency where it takes very little force to cause large response. At the resonant frequency it appears that the structure has no apparent stiffness.

Table 1: Cantilever Beam Frequencies with Various Tip Conditions (Free, Spring, Pinned)

Condition	Mode 1	Mode 2	Mode 3
Cantilever	58.	363.	1017.
K=1E1	68.	365.	1018.
K=1E2	123.	382.	1024.
K=1E3	224.	546.	1092.
K=1E4	251.	787.	1544.
K=1E5	254.	820.	1705.
K=1E6	254.	823.	1718.
Pinned	254.	823.	1720.

Copyright 2009 Pete Avitabile

Now at the antiresonances, this is a region in frequency where it takes excessive force and there is very little to essentially no response. At the antiresonant frequency, it appears that the structure is infinitely stiff. That is to say that at the antiresonant frequencies, there is no displacement and it appears that the cantilever is pinned at that point at that antiresonant frequency.

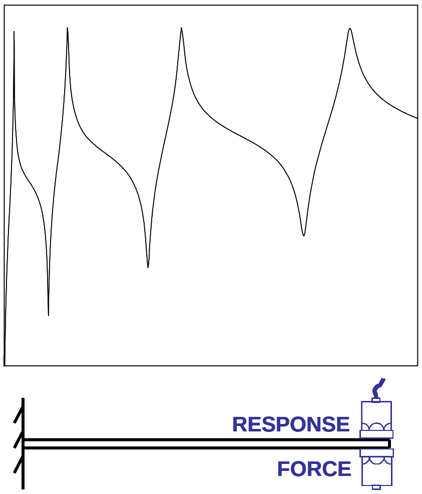

Figure 2 – Drive Point FRF Measurement for the Tip of the Cantilever Beam

Now if there would be a change in stiffness at the tip of the cantilever beam, then there would be a shift in the peaks of this function. If stiffness is added to the tip of the beam then the peaks will shift upward. This is shown in Figure 3.

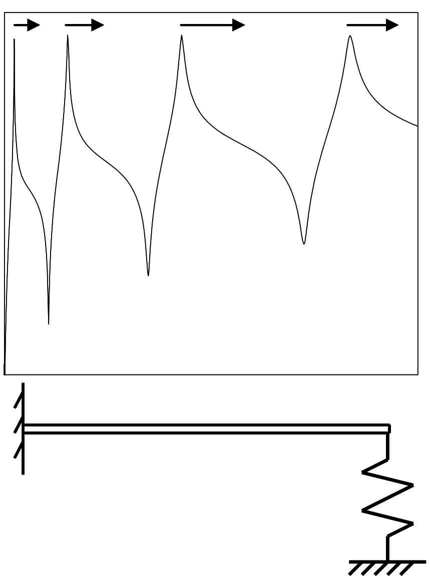

Figure 3 – Shift of Frequencies Due to Spring

But as the stiffness is increased, there will be some limit to the shift in the frequency of the system. Now if we realize that the antiresonance is actually the frequency at which the cantilever beam tip displacement is zero, then it is obvious that this is the frequency where the beam appears to be pinned at the tip. This is shown in Figure 4. From that schematic it is easy to realize that the peaks of the unconstrained cantilever beam can never shift past the antiresonances of the cantilever beam because this is essentially the cantilever constrained at the tip which is the pinned condition.

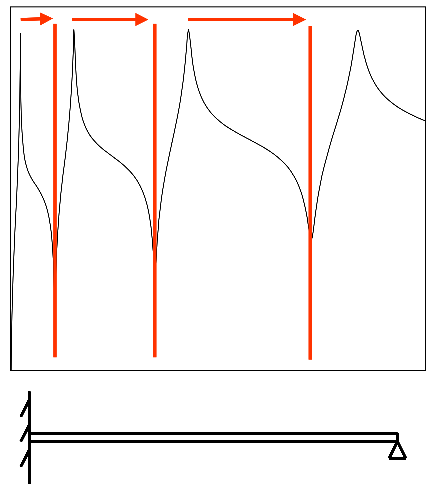

Figure 4 – Maximum Shift of Frequencies with Constraint

So from this discussion, it should be clear that the cantilever beam frequencies can only shift so far when a spring is considered at the tip of the beam. Further, we can actually identify how far those frequencies can shift by looking at the antiresonances at the tip of the unconstrained cantilever beam.

I hope that this discussion clears up the mystery as to why the frequencies can only shift so far before there is no further change in the frequencies. The best way to prove it to yourself is to make a simple finite element model and check out the results. If you have any more questions on modal analysis, just ask me.

Copyright 2009
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

If I run a shaker test with the input oblique to the global coordinate system, how do I decompose the force into the specific components in each direction? Wait – we need to discuss this before you take data.

Well it turns out that you don’t really need to break the force up into components in the global coordinate system. There is an easier way to account for this. But first let’s discuss a few basics before we get to how we are going to help fix the problem you described.

The first thing we need to understand is that even if you could break the input force up into two separate inputs in the global coordinate system, you wouldn’t actually have two separate independent inputs; the two inputs are linearly related to the one independent input applied to the structure. So even if you could break it up into components there would be no advantage to doing that. But let’s think about how we got ourselves into even thinking that we needed to decompose the force into separate components.

Let's start this discussion with a simple structure that has mode shapes that are very directional in nature as seen in Figure 1.

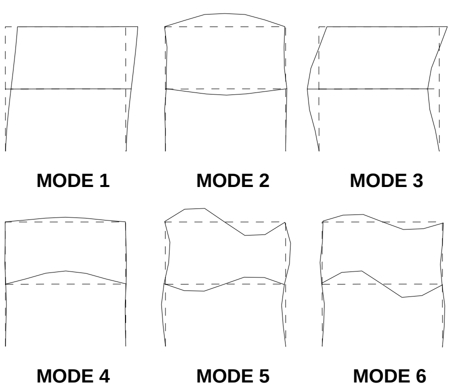

Figure 1 – Mode Shapes With Very Directional Character

Now just what do I mean by directional modes. That means that the response of the structure is primarily in one direction with very little or no response in the other directions for a given mode of the structure. Yet another mode of the structure may have response in a different direction than the first mode with little or no response in the other directions. We can see this in Mode 1 and Mode 3 in Figure 1; notice that the mode shapes are basically in the horizontal direction with very little motion in the vertical direction. But if we look at Mode 2, Mode 4, Mode 5, and Mode 6, we see that the main motion in the shape is in the vertical direction with little motion in the horizontal direction.

So if we wanted to pick a reference point on the structure for the modal test, then it isn’t easy to do if I restrict myself to either the horizontal direction (X) or the vertical direction (Y). So maybe I would need to have some reference that is oblique to the global coordinate system that I selected for the set of measurements. For the purpose of this discussion, let’s assume that I am only interested in the first four modes of the system. First let’s write the equations assuming that I will have a reference in one modal test for the x-direction modes and then a reference in a second modal test for the y-direction modes. (Eventually, we will write the equations for an oblique reference to show how to select a reference that is suitable for all the modes in one modal test.)

The most important thing to discuss now is the importance of the drive point measurement and how it relates to the equations describing the residues and mode shapes. Let's recall the equation for the frequency response function

$$h_{ij}(j\omega) = \sum_{k=1}^{m} \frac{a_{ijk}}{(j\omega - p_k)} + \frac{a_{ijk}^{*}}{(j\omega - p_k^{*})}$$

We need to remember that the residues are directly related to the mode shapes (and the q scaling factor) for a particular measured degree of freedom as

$$a_{ijk} = q_k u_{ik} u_{jk}$$

or for the whole set of measurements in matrix form as

$$\begin{bmatrix} a_{11k} & a_{12k} & a_{13k} & \cdots \\ a_{21k} & a_{22k} & a_{23k} & \cdots \\ a_{31k} & a_{32k} & a_{33k} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} = q_k \begin{bmatrix} u_{1k}u_{1k} & u_{1k}u_{2k} & u_{1k}u_{3k} & \cdots \\ u_{2k}u_{1k} & u_{2k}u_{2k} & u_{2k}u_{3k} & \cdots \\ u_{3k}u_{1k} & u_{3k}u_{2k} & u_{3k}u_{3k} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

So if we picked a particular reference such as 7x and measured at 24 points in the x and y directions, then the set of data would be written for a particular mode as

$$\begin{Bmatrix} a_{1x7x} \\ a_{1y7x} \\ a_{2x7x} \\ a_{2y7x} \\ a_{3x7x} \\ \vdots \\ a_{7x7x} \\ \vdots \\ a_{24x7x} \\ a_{24y7x} \end{Bmatrix} = q\ u_{7x} \begin{Bmatrix} u_{1x} \\ u_{1y} \\ u_{2x} \\ u_{2y} \\ u_{3x} \\ \vdots \\ u_{7x} \\ \vdots \\ u_{24x} \\ u_{24y} \end{Bmatrix}$$

and then we would see that the drive point measurement at 7x would be the measurement needed to scale the residues to get scaled mode shapes using

$$a_{7x7x} = q\ u_{7x} u_{7x}$$

But we have to remember that from the reference in the x-direction (7x), only Mode 1 and Mode 3 can be easily measued because these modes are in the x-diection whereas Mode 2 and Mode 4 are modes in the y-direction and cannot be easily measured if the reference is in the x-direction.

In order to measure Mode 2 and Mode 4, a reference in the y-direction is necessary. Of course, a second test is needed to accomplish this. If a reference is selected at point 20Y for instance, then the equation would be written relative to that reference and the drive point at 20Y would be used to obtain a scaled mode shape as discussed above and is

$$\begin{Bmatrix} a_{1x20y} \\ a_{1y20y} \\ a_{2x20y} \\ a_{2y20y} \\ a_{3x20y} \\ \vdots \\ a_{20y20y} \\ \vdots \\ a_{24x20y} \\ a_{24y20y} \end{Bmatrix} = q\ u_{20y} \begin{Bmatrix} u_{1x} \\ u_{1y} \\ u_{2x} \\ u_{2y} \\ u_{3x} \\ \vdots \\ u_{20y} \\ \vdots \\ u_{24x} \\ u_{24y} \end{Bmatrix} \quad \text{and} \quad a_{20y20y} = q\ u_{20y} u_{20y}$$

But this requires that a modal test be run twice with two different references. Another approach would be to select an additional point on the structure at some oblique angle at an arbitrary point 99s for instance where a drive point measurement of force and acceleration is made. The input to the structure is shown for illustration in Figure 2. This reference can be any point on the structure at any oblique angle as long as that location is suitable to excite all the modes of interest.

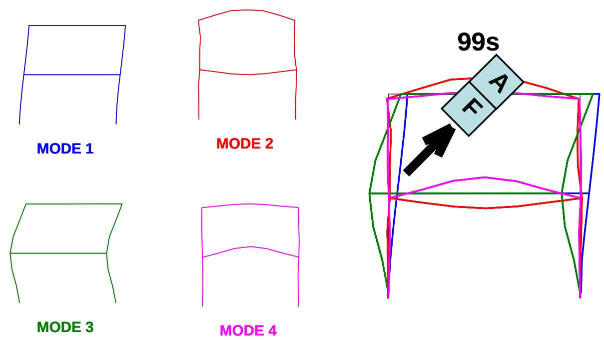

Figure 2 – Modal Test with Oblique Input Excitation

With this set of measurements, the set of equations relative to reference at point 99s and the drive point measurement would be

$$\begin{Bmatrix} a_{1x99s} \\ a_{1y99s} \\ a_{2x99s} \\ a_{2y99s} \\ a_{3x99s} \\ a_{3y99s} \\ a_{4x99s} \\ a_{4y99s} \\ \vdots \\ a_{99s99s} \end{Bmatrix} = q\ u_{99s} \begin{Bmatrix} u_{1x} \\ u_{1y} \\ u_{2x} \\ u_{2y} \\ u_{3x} \\ u_{3y} \\ u_{4x} \\ u_{4y} \\ \vdots \\ u_{99s} \end{Bmatrix} \quad \text{and} \quad a_{99s99s} = q\ u_{99s} u_{99s}$$

So once the mode shapes are scaled using the drive point measurement, then there really is no need to include the reference point in the description of the mode shapes. This happens to be a very convenient way to obtain scaled mode shapes without ever needing to include the oblique drive point measurement in the actual geometry description of the mode shapes. But of course it is critical to remember that the oblique reference location must be a good location where all of the modes can be observed from that one reference location for this to work.

I hope this explanation helps you to understand that you can pick any angle for the reference - just as long as its not the node of a mode. And you can use this oblique reference location as a drive point for scaling the modes of the system. If you have any other questions about modal analysis, just ask me.

Copyright 2009
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD by Pete Avitabile

Illustration by Mike Avitabile

When the transfer function is evaluated along the frequency axis, the damping is zero.
Does this mean there is no damping in the system. Let's clarify this confusing point.

Well I find that this becomes a confusing point for many people so let's try to talk about it and explain what is actually happening with this. So I will discuss a few items along the way here as part of the discussion.

First, let's write the system transfer function in partial fraction form

$$h(s) = \frac{a_1}{(s - p_1)} + \frac{a_1^*}{(s - p_1^*)}$$

and realize that the roots or poles of this function for an underdamped system can be written as

$$s_{1,2} = -\zeta\omega_n \pm \sqrt{(\zeta\omega_n)^2 - \omega_n^2} = -\sigma \pm j\omega_d$$

Because the function is complex, the roots will be a function of two variables, σ and ω , which are the real and imaginary parts of the root. The numerator is called the residue of the system transfer function (and is so named because it comes from the Residue Theorum used to evaluate the function).

Now when we plot this function, the plot is going to map a surface because the function is defined by two independent variables, namely σ and ω . So if we hold σ constant and vary ω and then incrementally change σ and recompute the range of ω there will be a matrix of complex numbers that are generated. Because the numbers are complex, we can make a plot of the real and imaginary parts separately but we could also plot the magnitude and phase for the function. In any event, this surface can be plotted in any one of these forms to describe the system transfer function.

This is shown in Figure 1 (from Vibrant Technology webpage). We can discuss each of the pieces of the system transfer function but I really want to concentrate on the magnitude of this function for the discussion here. (But we need to always remember that this is a complex valued function that has real and imaginary, or magnitude and phase, to describe the total function.)

So when we say that we evaluate the function at σ =0, we aren't really saying that the damping is zero but rather that the function is evaluated along the $j\omega$ axis.

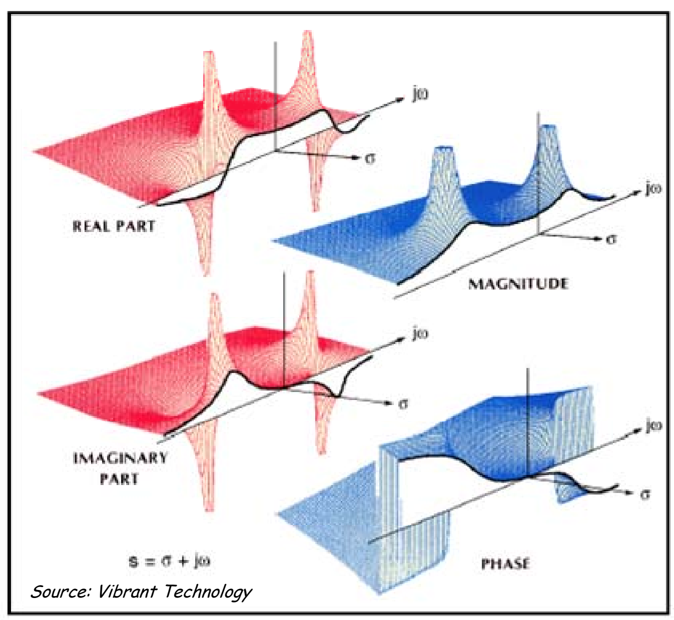

Figure 1 – System Transfer Function

Copyright 2009
Pete Avitabile

Now if we write this equation evaluated this way then we can write the frequency response function as

$$h(j\omega) = h(s)\Big|_{s=j\omega} = \frac{a_1}{(j\omega - p_1)} + \frac{a_1^*}{(j\omega - p_1^*)}$$

And if we were to look at the magnitude of the system transfer function evaluated along the jω axis, and project the face of that cut along that axis we would see the plot shown in Figure 2 that is projected from that slice. This what we measure in the FFT analyzer - the frequency response function. And we can see that there is only one independent variable ω used to describe that function. We would also notice that we only have a line now rather than the surface described for the system transfer function.

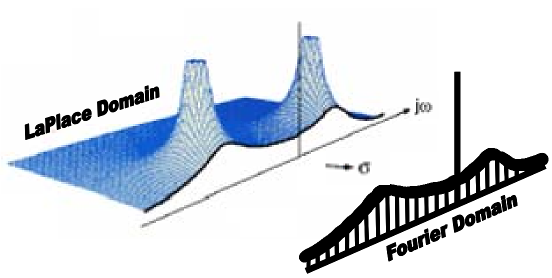

Figure 2 – System Transfer Function (Magnitude) with Frequency Resonse Function

So now we have a handle on where this frequency response function comes from. Now we want to describe the splane and the system transfer function surface. Well, the surface looks like a tent with two poles so I want to use this as an analogy with a wedding with a seating arrangement under the tent. We know that there are two sides to the seating arrangement – the bride and groom (the pole and conjugate of the pole). Now you could be seated on either side depending on which side of the wedding party you are with.

Let's say that you are with the bride's side of the wedding and you are seated in the first row-second seat. Now when you sit down and look up you will notice that the tent is a certain height above your head (the magnitude at that particular value of σ and ω). You will also notice that there is a mirror image seat on the groom's side (conjugate) at the first row-second seat; and the height above that seat is the same in terms of its magnitude.

But let's say that someone else was seated at the second row-third seat on the bride's side. Now at that point, the height of the tent is much higher than the first case. And of course, there is also a mirror image seat on the groom's side which has the same height.

So each of these seats has a particulat tent height above the seat location. That height maps the surface of the tent. But you notice that there is one seat on both sides that corresponds to where the pole is located; these are analgous to the roots which are complex conjugate pairs. Notice that no one can sit there and no one can really tell what the magnitude of the tent is at that location because it is undefined; the magnitude of the system transfer function can not be determined at the pole (root) of the tent because it is undefined at that location.

So this tent analogy is a pretty good description of the system transfer function. The value of the function is determined by the location in the seating plan. The amplitude changes depending on location. The system transfer function is undefined at the poles or roots of the system; that is where the poles of the tent are located and no one can sit at that location to determine the amplitude or height of the tent. (We use the Residue Theorum to evaluate this.)

And of course we all know that the first row is the most important row. In fact, that is where Mr Fourier resides – right along the jω axis – which is the slice we took out of the system transfer function.

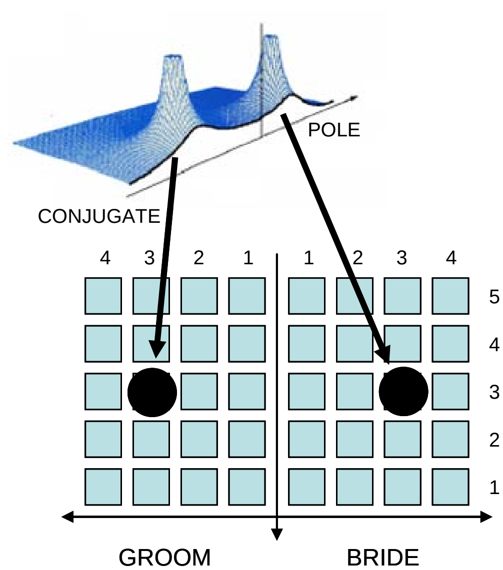

Figure 3 – The S-Plane Representation

So I hope these simple examples shed a little more light on the system transfer function and the frequency response function and how they are related to each other. One last note regarding a recent wedding where the best man was asked to introduce the bridal party. He did just fine up until he introduced (for the first time) the groom and bride as John and Angela – it was also the last time because the bride's name wasn't Angela! But that's another story.

If you have any other questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

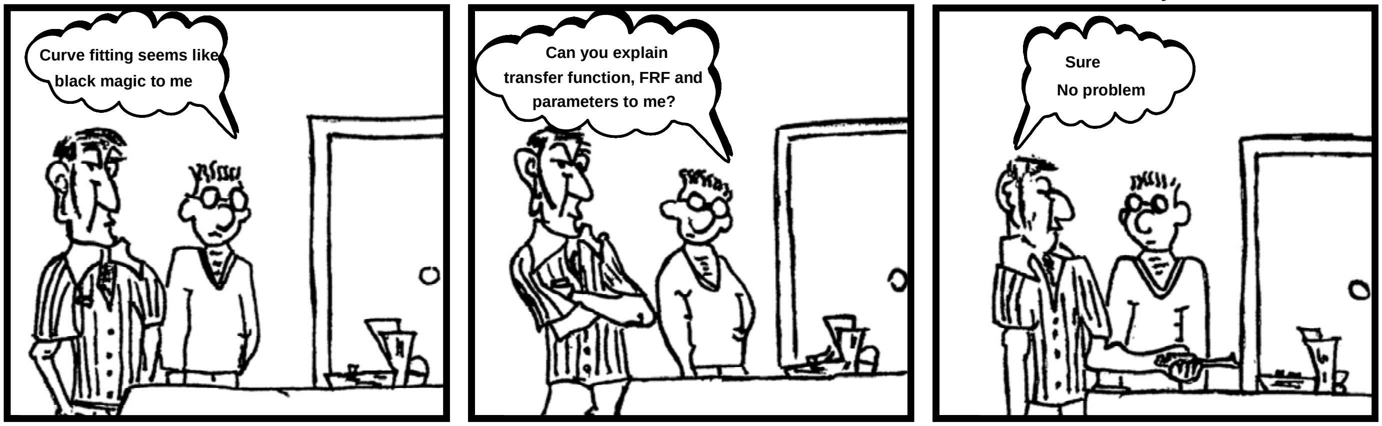

Illustration by Mike Avitabile

Curvefitting still seems like black magic to me. Can you explain transfer function, FRF and parameters to me?
Sure – no problem.

Well, curvefitting might look like black magic at first but I want to make a few simple analogies to help you understand that it is really fairly straight-forward and with the example I will show is very simple indeed.

The last time we discussed some related information regarding the system transfer function and the frequency response function (FRF). We wrote the system transfer function in partial fraction form for a single degree of freedom system as

$$h(s) = \frac{a_1}{(s-p_1)} + \frac{a_1^*}{(s-p_1^*)}$$

and we also wrote the frequency response equation as

$$h(j\omega) = h(s)\,\Big|_{s=j\omega} = \frac{a_1}{(j\omega-p_1)} + \frac{a_1^*}{(j\omega-p_1^*)}$$

Now if we look at these two equations we notice in the first case the independent variable is “s” and in the second equation it is “ω” and that the “h” depends on these values. But I also notice that there are two constants, or parameters that are “a” the residue and “p” the pole. So these are the parameters that define “h” given some value “ω”; we call these modal parameters.

Now if we look at the system transfer function or a piece of the system transfer function called the frequency response function, we need to realize that the surface of the system transfer function as well as the curve of the frequency response function are defined by only two parameters for the single degree of freedom system, namely the pole “p” and residue “a”. So looking at Figure 1, we need to realize that only two parameters define that surface and line – pretty amazing.

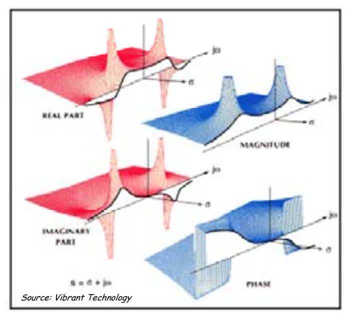

$$h(s) = \frac{a_1}{(s-p_1)} + \frac{a_1^*}{(s-p_1^*)}$$

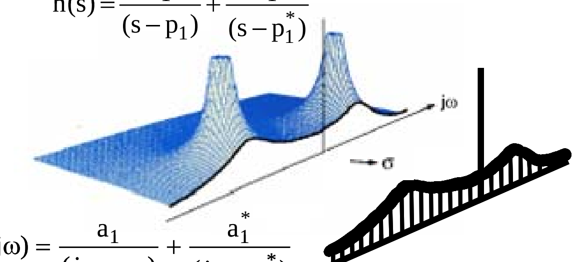

$$h(j\omega) = \frac{a_1}{(j\omega-p_1)} + \frac{a_1^*}{(j\omega-p_1^*)}$$

Figure 1 – System Transfer Function and FRF

Now let’s take a step back to something a little simpler and more commonly understood. Let’s look at a very simple straight line fit of of some measured data. We are going to perform a least squares error minimization for the data presented in Figure 2. Now we know we can fit any line to the data but for this set of data it seems that a first order fit makes the most sense. Of course the model we are going to use is

$$y = mx + b$$

and there are two parameters that define the line, namely the slope and y intercept.

Copyright 2009
Pete Avitabile

So for instance, in Figure 2 the resulting least squares fit of the data resulted in two parameters with a slope of 12.097 and a y-intercept of -0.019. Also realize that this data was obtained from a set of measured data that had some variance and that the least squares regression analysis identified the best parameters to represent this data with these two parameters of slope and y-intercept.

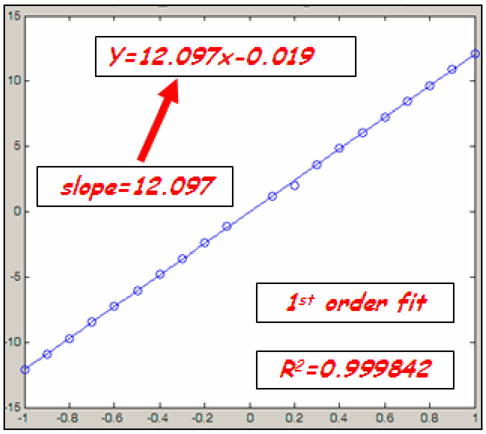

Figure 2 Example of Simple Straight Line Fit

So if we were to apply this same logic to a single degree of freedom frequency response function then I would fit a first order model of the form of a frequency response function as written above to the data presented in Figure 3. And if you looked at this schematic it would be very easy to see that there is a set of data and curvefit from which two parameters are obtatained, namely the pole and residue.

It is really the same as the straight line fit except that the data is complex and the line is a little more complicated. But in principle, it is the same methodology. We measure data at discrete data points as complex data and then fit a line of the frequency response function to the data to find the parameters that best describe the data in a least squares fashion.

Now of course the data in Figure 3 is for a single degree of freedom system. This same approach can be extended to a higher order function as shown in Figure 4. So in this way we can fit multiple modes (or essentially a higher order polynomial) to the data described by the discrete complex data measured from the frequency response function. And all the same arguments relating to the estimation of modal parameters can be made again here with the data in Figure 4.

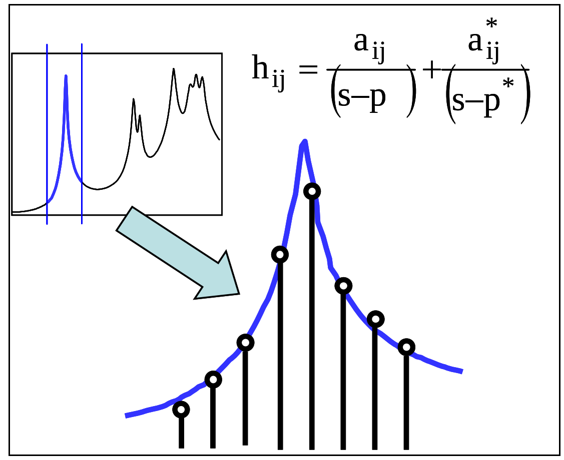

Figure 3 – Conceptual SDOF Curvefit

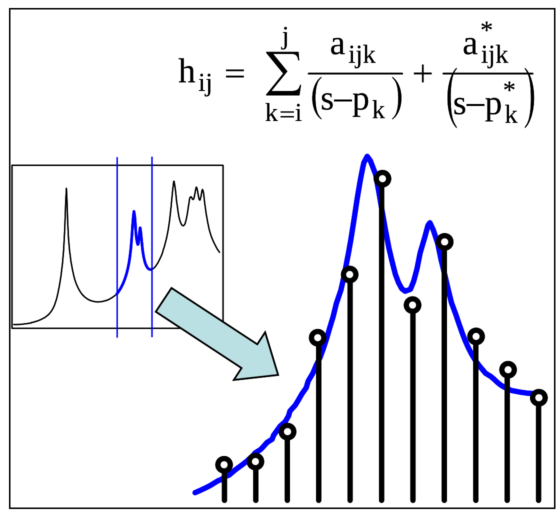

Figure 4 – Conceptual MDOF Curvefit

So if you accept the procedure that you always perform for the simple straight line fit, then you have to agree that the same procedure is applied for the modal parameter estimation process (but of course the data is complex and the line is slightly more complicated). Essentially in both cases, parameters are extracted, in a least squares fashion, that describe the function.

So there really isn't any black magic at all to the curvefitting process. It is really the same process that we all perform with simpler straight line regression analyses. Modal parameter estimation is just an extention of simpler curvefitting of data. If you have any other questions about modal analysis, just ask me.

Copyright 2009
Pete Avitabile

Illustration by Mike Avitabile

Is there any effect in impact testing when hitting harder or softer with the different tips? This needs some discussion.

Well there can definitely be some significant effects depending on which tips you are using and how hard you are hitting the structure. Actually with imact testing, you really need to have the hammer do the work. This is not like driving nails in a wood framing operation. What I mean by this is that the hammer is held in your hand and your wrist is used like a pin joint as you swing the hammer. I have seen some people perform impact tests and they are swinging the hammer excessively which is not really the way to do this. (And I have seen some hammer tips that look like they have been through a nuclear detonantion they are so damaged.)

From what I have seen, there are many different hammer tips that are often supplied with the hammer kits commercially available. In particular, I want to concentrate on a few tips that I have seen that have significantly different force spectrums during various impacts with increasing levels of force applied to a structure. These tips need to be used with caution. The air capsule and the plastic cap on top of the hard plastic tip will be compared to the hard plastic tip to show some overlooked effects of these tips when used with different force levels.

In all cases, I will impact a very large, massive, steel block and I will keep the applied force level to a softer level, a medium level and a harder level of force applied to the block. In the first case, the air capsule will be evaluated, then the plastic cap on top of the hard plastic tips followed by the hard plastic tip. The results of these impacts (time pulse and input spectrum) are shown in the attached figure.

Air Capsule – This tip shows a dramatic change in the time pulse and resulting force spectrum depending on force level applied. Notice that the frequency spectrum excited with an attenuation of 20 dB changes significantly (and the hardest hit excites a much wider range overall). So that if you are performing an impact test and every hit for an average has a different level of force, the spectrum excited by that impact is significantly different each time. This could have a significant effect in the coherence in the higher frequency ranges.

Plastic Cap on Top of Hard Plastic Tip – This tip can also exhibit the same type of behavior in many cases. For this particular test, the plastic cover was slightly longer than the hard plastic tip so that there was effectively a small air pocket included. Again depending on the level of excitation applied there may be a significantly different input force spectrum/frequency range excited.

Hard Plastic Tip – Notice that this tip shows relatively little variation in the spectrum force characteristics over the frequency spectrum excited. There is same small variation but relative to the previous two tips, it is relatively small. So the frequency range excited with this tip will be relatively constant even with relatively different impact levels applied.

This effect is very important when testing structures where the frequency range to be excited is critical. For the first two tips, the frequency range excited is very dependent on the level of excitation used for the test. Care must be exercised to assure that a fairly consistent force strike is applied with every impact, for every average, for every measurement. This is not so easy to do all the time. So use care when using some of those special impact tips as part of your impact hammer kit.

I hope this explanation helps you to understand that you need to be very careful when performing impact tests. If you have any other questions about modal analysis, just ask me.

Copyright 2010
Pete Avitabile

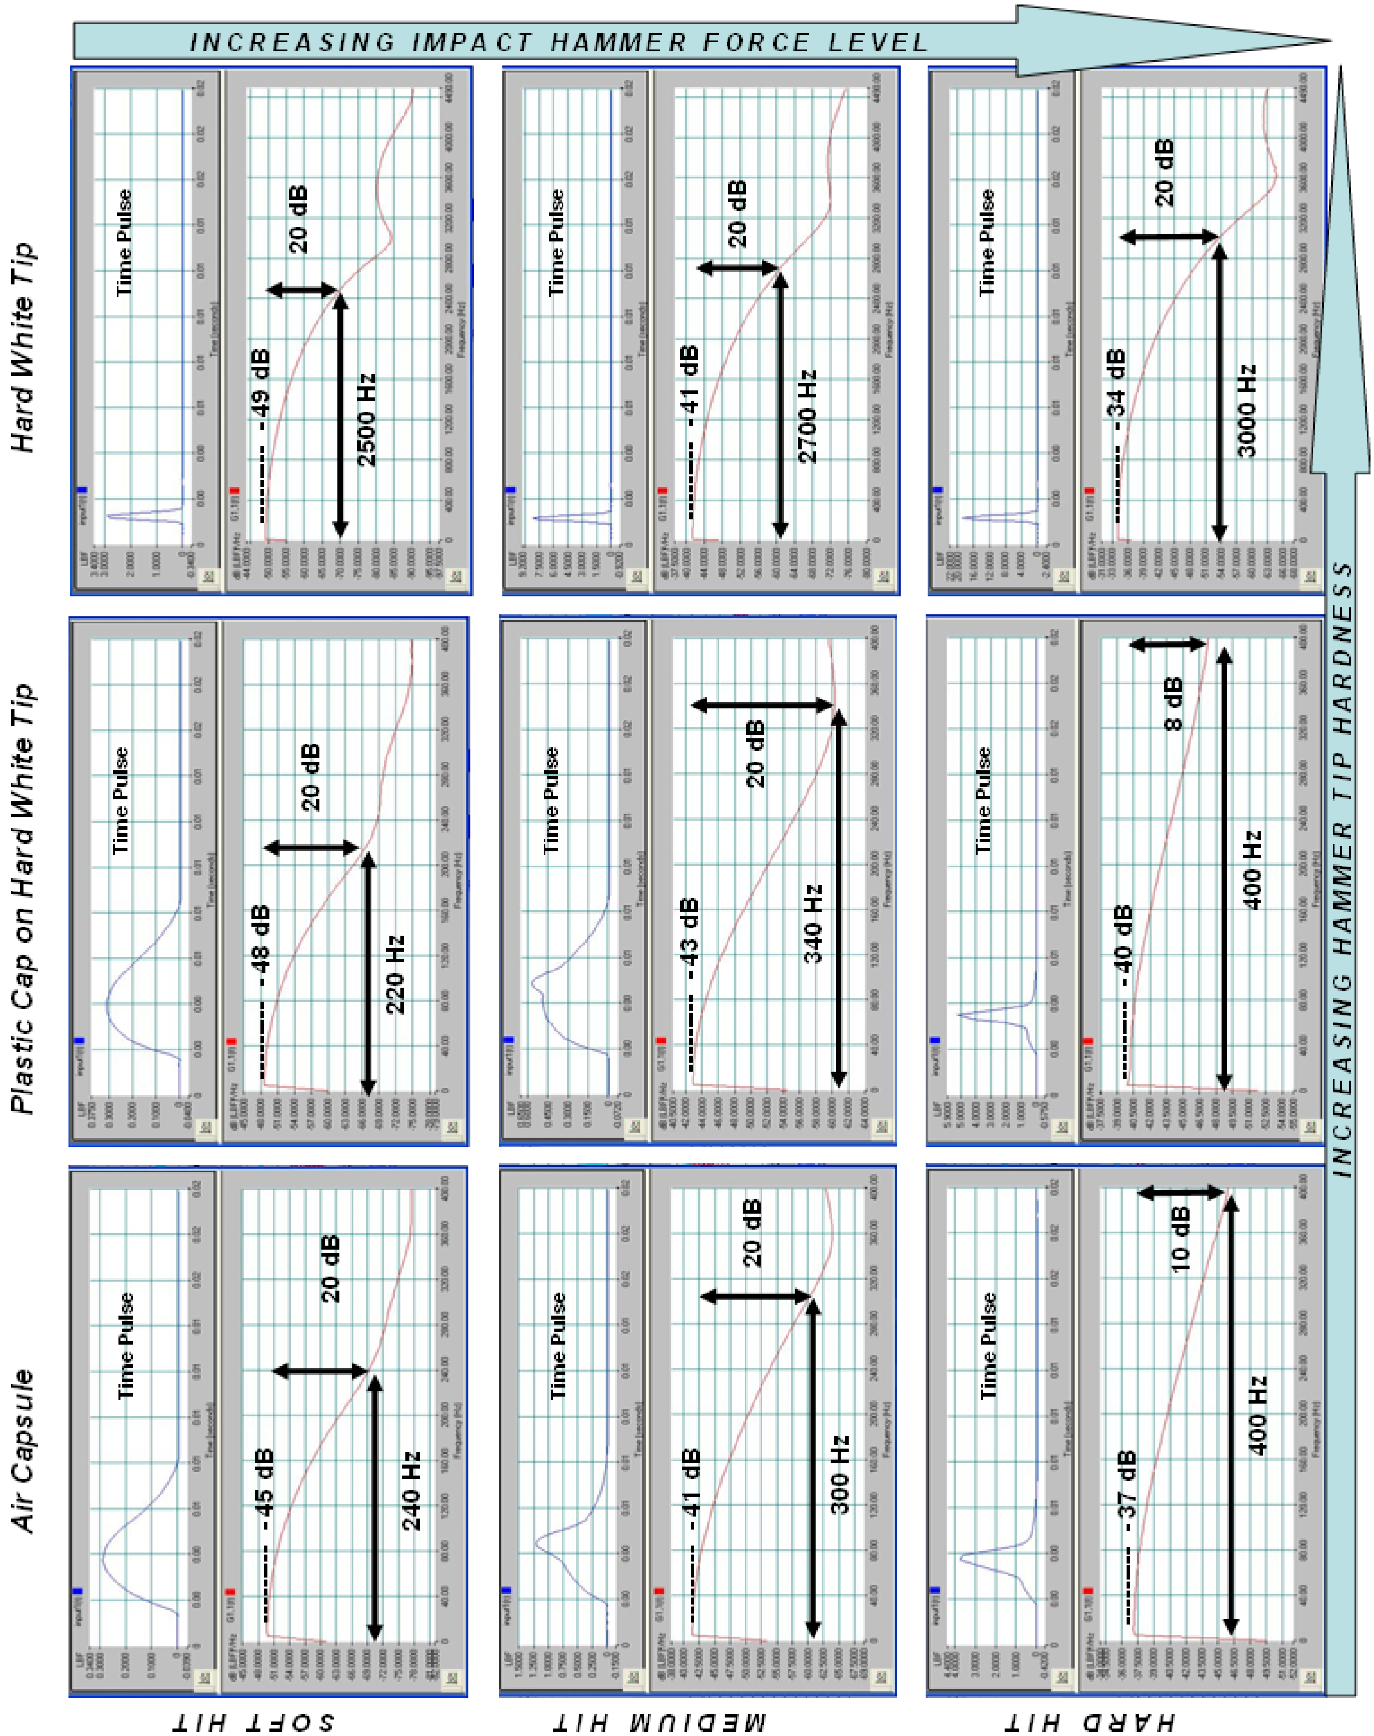

Copyright 2010
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD by Pete Avitabile

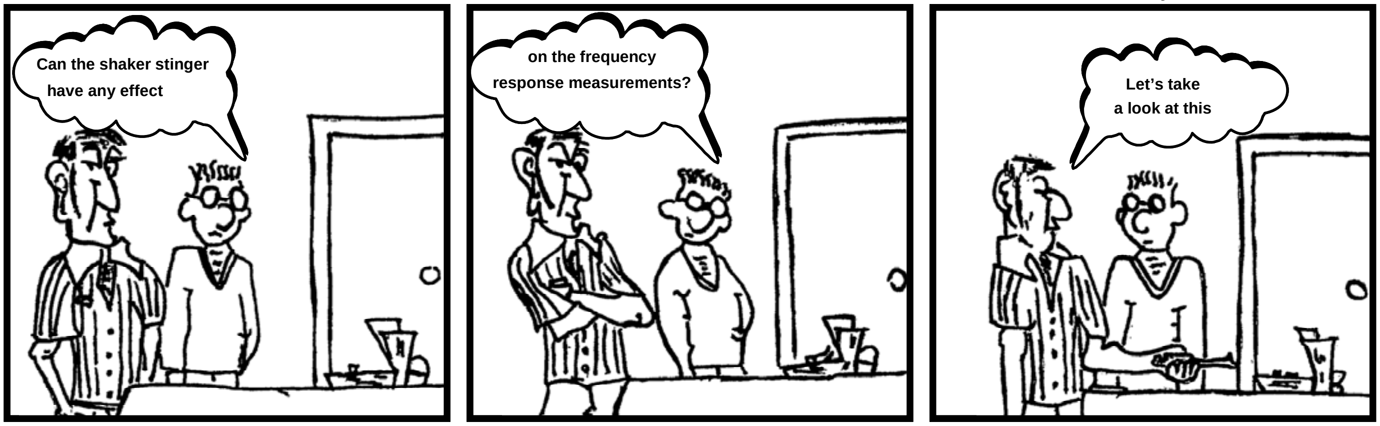

Illustration by Mike Avitabile

Can the shaker stinger have any effect on the FRF measurements? Let's take a look at this.

When performing shaker tests, there should be concern when connecting the shaker to the structure under test. Many times I have found that there are misconceptions in regards to this type of shsker testing when used for the development of frequency response functions for a modal test. Often times people are familiar with base excitation shaker testing when used for qualification of equipment due to some simulation of actual loads or for conducting qualification testing in accordance with some specification. Modal testing is a little different than these types of qualifications tests.

Rather than directly conneting the test structure to the shaker armature, there is a long slender rod, referred to as a quill or stinger, that is placed between the shaker and the test structure. The intent of this stinger is to provide axial excitation along the length of the rod to the structure while providing very little lateral stiffness to the structure under test. Figure 1 shows a typical shaker test set up configuration for illustration.

Now I could spend a lot of time to discuss all of the different situations that could possibly exist but there is not enough room here to discuss all the possible scenarios. At best, I can probably illustrate some typical testing situations that must be considered and show some possible frequency response function measurement distortions that could result from an inappropriate test set up with the shaker and the stinger.

Case A: Let's consider the effect of bending of the shaker stinger during a test. Remember the intent is to provide only input excitation along the length of the stinger and to minimize any bending of the stinger. Two thngs can happen when the stinger bends. The stinger can impart a rotational load which is not measured by the force gage; remember that the force gage expects to see only uniform compressive or tensile loads and any moment will distort the force gage reading as well as impart a moment into the structure that is not measured as a rotational load. Also, the stinger can introduce rotational stiffness into the structure under test which is not really part of the structure's dynamic characteristics. Figure 2 shows a measurement where the shaker excitation was applied at different elevations to a simple structure. Clearly, there is an effect on the measured frequency response functions. This should be checked at the preliminary stages of test set up.

Case B: Now consider the situation where the alignment of the shaker relative to the test structure is not set up properly. This is similar to the case above in that the is a bending contribution that the shaker imposes on the measurement set up. Figure 3 shows measurement distortions resulting from shaker misalignment. Care should be exercised in aligning the shaker.

Case C: Another consideration is related to the length of the shaker quill in the test set up. Again the measurements in Figure 4 show that there can be differences in the measurements obtained. Preliminary testing should be performed to understand if this effect is important or not.

Case D: Finally, the type of stinger used can also have an effect on the measurements obtained. Figure 5 shows the effects resulting from different stingers used and should be checked at the start of testing.

Copyright 2010
Pete Avitabile

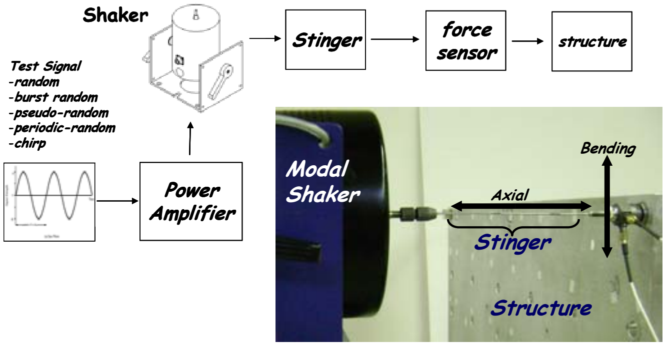

Figure 1 – Typical Modal Test Shaker Set Up Arrangement

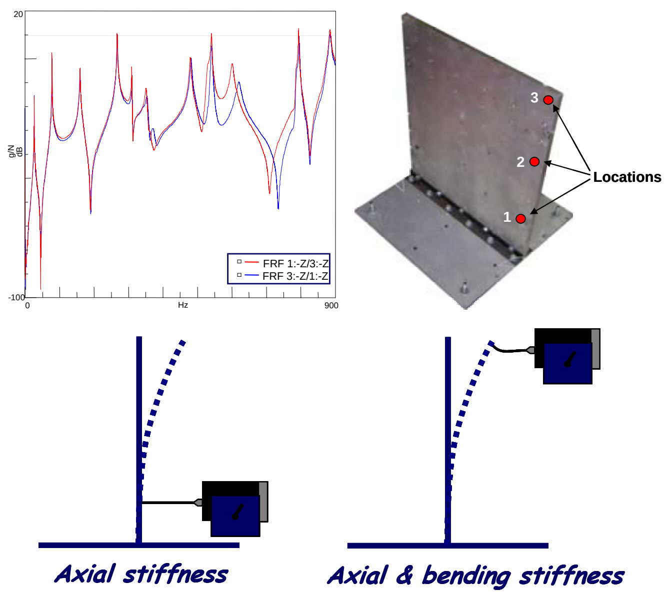

Figure 2 – Stinger Bending Effects

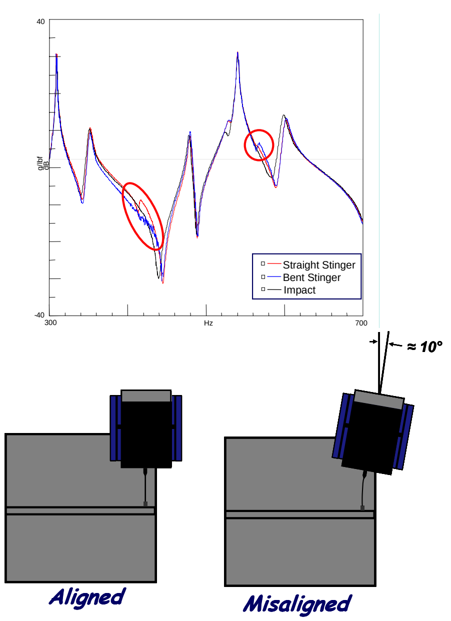

Figure 3 – Stinger Misalignment Effects

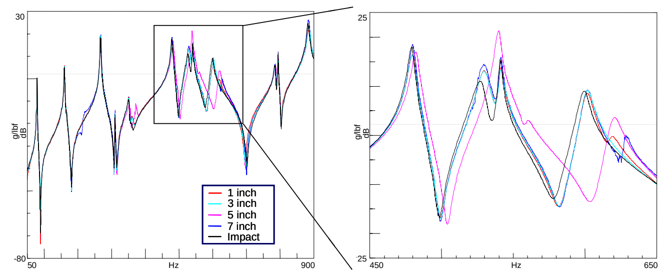

Figure 4 Stinger Length Effects

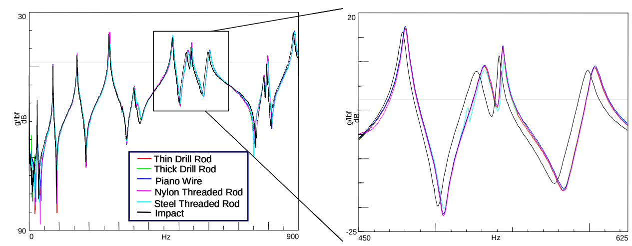

Figure 5 Stinger Type Effects

So in all the cases shown there is clearly an effect of the stinger on the results obtained. Whether it be the position of the shaker/stinger on the structure, the skewed alignment to the structure, the length of the stinger or type of the stinger, there can be an effect of the bending of the stinger that affects the measured frequency response function.

Care must be exercised in the test set up. Unfortuantely there is not one clear cut answer as to which stinger configuration will provide the optimal results. This is heavily dependent on the structure under test and the frequency range of interest. However, it is very important to try different scenarios to assure yourself that the best possible frequency response measurements are being acquired for the configuration eventually utilized for the test.

I hope this explanation helps you to understand that you need to be very careful when performing shaker tests and that the stinger can have a significant effect on the overall results.. If you have any other questions about modal analysis, just ask me.

by Pete Avitabile

Illustration by Mike Avitabile

I am still confused about how the SDOF in modal space is related to the physical response?
This needs some discussion.

Well - this is a concept that is actually very simple but does need some explaining to make sure it is comprehended properly. First let's start with a few summary equations that we have presented several times before in previous articles. Of course the equation of motion in matrix form is the starting point

$$[M]\{\ddot{x}\}+[C]\{\dot{x}\}+[K]\{x\}=\{F(t)\}$$

This coupled set of matrix equations is then uncoupled by performing an eigensolution. The modal transformation equation is obtained from the set of modal vectors obtained from the eigensolution. The physical coordinate $\{x\}$ is related to the modal coordinate $\{p\}$ using the collection of modal vectors $[U]$

$$\{x\}=[U]\{p\}=\{u_1\}p_1+\{u_2\}p_2+\{u_3\}p_3+\cdots$$
$$\text{with } [U]=\left[\{u_1\}\quad\{u_2\}\quad\{u_3\}\quad\cdots\right]$$

Substituting this into the physical equation and premultiplying by the transpose of the projection operator [U] will result in a very simple diagonal set of equations in modal space where every equation (modal oscillator) is orthogonal and linearly independent (uncoupled) from each other and is given as

$$\begin{bmatrix}\overline{m}_1 & & \\ & \overline{m}_2 & \\ & & \ddots\end{bmatrix}\begin{Bmatrix}\ddot{p}_1\\ \ddot{p}_2\\ \vdots\end{Bmatrix}+\begin{bmatrix}\overline{c}_1 & & \\ & \overline{c}_2 & \\ & & \ddots\end{bmatrix}\begin{Bmatrix}\dot{p}_1\\ \dot{p}_2\\ \vdots\end{Bmatrix}+\begin{bmatrix}\overline{k}_1 & & \\ & \overline{k}_2 & \\ & & \ddots\end{bmatrix}\begin{Bmatrix}p_1\\ p_2\\ \vdots\end{Bmatrix}=\begin{Bmatrix}\{u_1\}^T\{F\}\\ \{u_2\}^T\{F\}\\ \vdots\end{Bmatrix}$$

There are several important things that need to be noted about this equation. And the mode shape matrix [U] has a lot to do with this.

First is that every equation contains only one variable to describe each equation – the modal displacement for each particular mode. Second is that every equation is uncoupled from every other equation. Third is that each equation is basically a very simple single degree of freedom (SDOF) system. Fourth is that the right hand side of the equation identifies the force that is appropriated to the modal oscillator from the physical force applied to the physical system. Figure 1 shows a schematic of a multiple degree of freedom system (MDOF) where coupling between degrees of freedom exist in the physical model and the resulting equivalent set of SDOF systems representing the modal system in modal space.

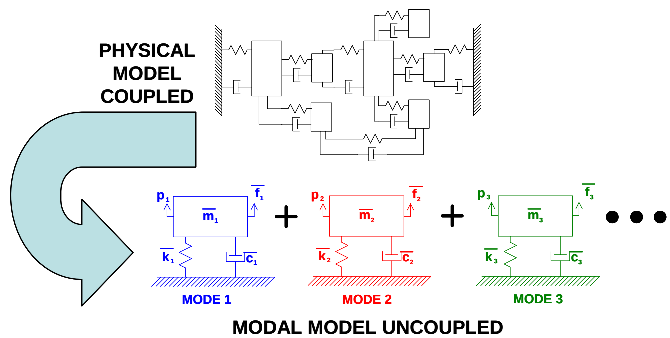

Fig 1 – MDOF System Schematic and SDOF Equivalent

So if we write out any one equation from the modal space form and use an "i" subscript the "ith" equation we would get

$$\overline{m}_i\ddot{p}_i+\overline{c}_i\dot{p}_i+\overline{k}_i p_i=\{u_i\}^T\{F\}=\overline{f}_i$$

Copyright 2010
Pete Avitabile

So with this simple SDOF equation we can calculate the response due to any force applied on the equivalent system. Of course we can see that the right hand side of the equation identifies how much of the force is appropriated from physical space to the equivalent modal system through the mode shape. With this force, then the response for the equivalent system can be identified. This response can be simply found from any Vibrations textbook – usually this is one of the first four chapters in most textbooks for free response, forced sinusoidal response or arbitrary input response. For sake of the discussion here, let's assume that an impact is applied at one point on the physical system.

Now that physical force will be appropriated to each of the modal DOF in modal space. So if we look at the first mode then we could calculate the impulse response for the SDOF describing mode 1. This SDOF response is then distributed over all the physical DOF using the modal transformation equation; this essentially scales the SDOF response to all the physical DOF using each value of the first mode shape at each individual DOF. This is schematically shown in Figure 2 (for just a few DOF to illustrate the concept). Now this only provides the part of the response of the physical system that is related to the contribution that mode 1 makes over all the physical DOF in the system; the portion of the response related to mode 1 is shown in blue.

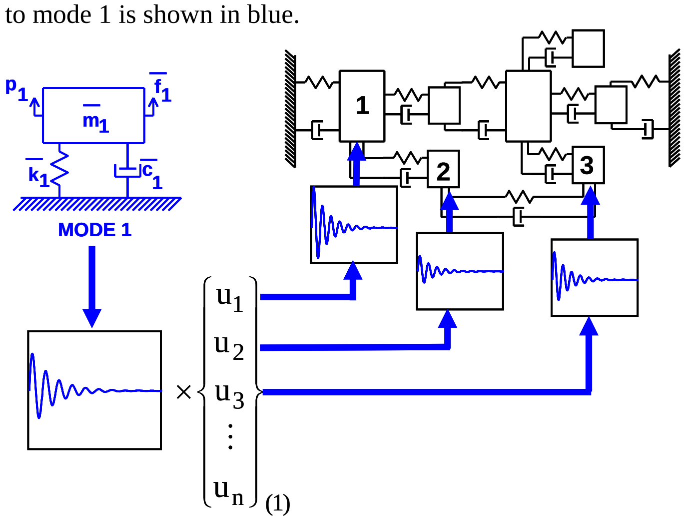

Fig 2 – Schematic Response for Mode 1 Contribution

This is not the entire physical response of the system – it is just the portion of the response that is related to mode 1. Now the contribution of the other modes needs to also be included. If we look at the second mode then we could calculate the impulse response for the SDOF describing mode 2 with the force that is appropriated to mode 2. Again this SDOF response for mode 2 needs to be distributed over all the physical DOF using the second mode shape; this only represents the portion of the response that is related to the second mode of the system. This is shown in Figure 3 (for just a few DOF to illustrate the concept); the portion of the response related to mode 2 is shown in red.

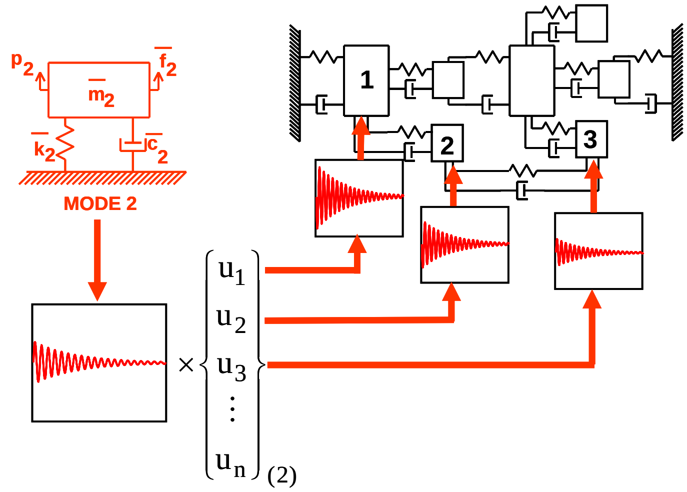

Fig 3 – Schematic Response for Mode 2 Contribution

This process is then continued for all the modes that contribute to the total response of the physical system. Of course you have to include all the modes that have a contribution to the overall response otherwise some of the solution is lost. The entire process is best seen in Figure 4. This figure shows the physical equation and the modal transformation equation which allows the coupled physical system to be written as a set of equivalent SDOF systems in modal space with the equivalent modal force applied on all the modal oscillators in modal space. It is important to realize that each mode is linearly independent from every other mode but that the total response is made up of the linear combination of the response of all the modes that participate in the response of the system.

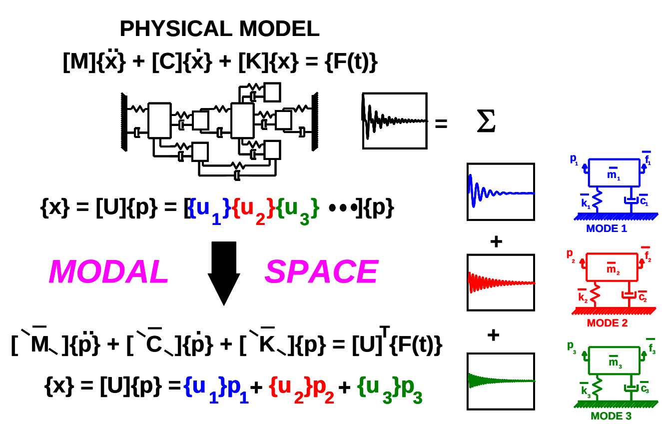

Fig 4 – Overview of the Modal Space Representation

I hope this explanation helps you to understand how the SDOF response is characterized in physical space from the modal space response. If you have any other questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

I hear mode participation all the time. What does that really mean? OK – let's discuss this.

So people always talk about mode participation but maybe it really isn't clear what they are referring to when they talk about it. So let's discuss what this concept is all about and put it in terms that might make more sense.

But to put it in context, let's write the equation of motion

$$[M]\{\ddot{x}\}+[C]\{\dot{x}\}+[K]\{x\}=\{F(t)\}$$

and recognize that the modal transformation is obtained from the eigen solution with the physical coordinate {x} is related to the modal coordinate {p} using the collection of modal vectors [U]

$$\{x\}=[U]\{p\}=\{u_1\}p_1+\{u_2\}p_2+\{u_3\}p_3+\cdots$$
$$\text{with } [U]=[\{u_1\} \quad \{u_2\} \quad \{u_3\} \quad \cdots]$$

and then further remember that the modal space equation is

$$\begin{bmatrix}\overline{m}_1 & & \\ & \overline{m}_2 & \\ & & \ddots\end{bmatrix}\begin{Bmatrix}\ddot{p}_1\\ \ddot{p}_2\\ \vdots\end{Bmatrix}+\begin{bmatrix}\overline{c}_1 & & \\ & \overline{c}_2 & \\ & & \ddots\end{bmatrix}\begin{Bmatrix}\dot{p}_1\\ \dot{p}_2\\ \vdots\end{Bmatrix}+\begin{bmatrix}\overline{k}_1 & & \\ & \overline{k}_2 & \\ & & \ddots\end{bmatrix}\begin{Bmatrix}p_1\\ p_2\\ \vdots\end{Bmatrix}=\begin{Bmatrix}\{u_1\}^T\{F\}\\ \{u_2\}^T\{F\}\\ \vdots\end{Bmatrix}$$

So the most important thing to understand is that the right hand side of this equation has the transpose of the mode shape times the physical force vector that is applied on the structure. So if you looked at a particular mode of interest, then you would see that the mode shape values have a strong effect on how much of that physical force is allocated or approporiated to the particular mode of interest.

What I mean by that is that if the value of the mode shape is large associated with the particular degree of freedom where the force is applied then that particular mode will get a lot of force appropriated to it in modal space. On the other hand if the value of the mode shape is small then there will be much less force appropriated to that particular mode in modal space. And if the value of the mode shape is zero then there will be no force allocated to that mode in modal space – this means that this particular mode has no contribution to the response because it sees no force applied on that particular modal oscillator in modal space.

So the modal transformation equation identifies how to uncouple all the coupled set of physical equations and it also identifies how much of the physical force is allocated for each of the modal oscillators in modal space. The larger the force that is applied to a particular modal oscillator, the large the response (in general) and the more that particular mode contributes or participates in the total response of the system.

So let's try a simple 3 dof system and see what force gets appropriated to modal space. Here is the model, with the equation of motion in physical space

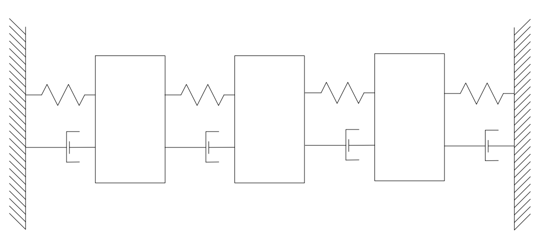

$$\begin{bmatrix}1 & & \\ & 1 & \\ & & 1\end{bmatrix}\begin{Bmatrix}\ddot{x}_1\\ \ddot{x}_2\\ \ddot{x}_3\end{Bmatrix}+\begin{bmatrix}0.2 & -0.1 & \\ -0.1 & 0.2 & -0.1\\ & -0.1 & 0.2\end{bmatrix}\begin{Bmatrix}\dot{x}_1\\ \dot{x}_2\\ \dot{x}_3\end{Bmatrix}+\begin{bmatrix}20000 & -10000 & \\ -10000 & 20000 & -10000\\ & -10000 & 20000\end{bmatrix}\begin{Bmatrix}x_1\\ x_2\\ x_3\end{Bmatrix}=\begin{Bmatrix}f_1\\ f_2\\ f_3\end{Bmatrix}$$

Fig 1 – Simple 3DOF Model and Equation of Motion

Copyright 2010
Pete Avitabile

Now an eigensolution for this will result in frequencies and mode shapes as

$$\left[\Omega^2\right]=\begin{bmatrix}5858 & & \\ & 20000 & \\ & & 34142\end{bmatrix}$$

$$[U]=\left[\{u_1\}\quad\{u_2\}\quad\{u_3\}\right]=\left[\begin{Bmatrix}0.500\\0.707\\0.500\end{Bmatrix}\quad\begin{Bmatrix}0.707\\0\\-0.707\end{Bmatrix}\quad\begin{Bmatrix}-0.500\\0.707\\-0.500\end{Bmatrix}\right]$$

Now let’s consider two different forcing functions.

$$F=\begin{Bmatrix}f_1\\0\\0\end{Bmatrix}\quad \text{and} \quad F=\begin{Bmatrix}0\\f_2\\0\end{Bmatrix}$$

For the first case with just f_1 applied, the force on mode 1 would be $0.5*f_1$, the force on mode 2 would be $0.707*f_1$ and the force on mode 3 would be $-0.5*f_1$. Now there are different allocations of the physical force on each of the modal oscillators which is controlled by the value of the mode shape associated with the degree of freedom where the force is applied.

Now for the second case with just f_2 applied, the forces on each of the modal oscillators would be $0.707*f_2$, $0.0*f_2$ and $0.707*f_2$ for each of the three modes. Notice that mode 2 sees no force because the value of the mode shape is zero for the degree of freedom associated with the force for mode 2. So we can say that mode 2 does not participate in the response of the system. Its modal participation is zero. But that doesn’t mean that mode 2 doesn’t exist – it just means that it doesn’t contribute to the response of the system for this particular loading scenario (but it certainly has contribution for the first loading condition).

So let me try to give a little example to explain this a little better. Let’s say that you are a cook in a restaurant. And imagine that you have a lot of different recipes that you might make. You also have a lot of ingredients and spices that could be used in all the different recipes that you make. Here is my question. Do you use all of your spices in all of your recipes with equal amounts of all the spices. No. You use varying amounts of spices in each recipe. And in some recipes there are many spices that are never even used. What I am trying to say is that you don’t use all of your spices all the time. Only certain spices “participate” in each recipe and to varying amounts.

Now if I am cooking some French dish then there are certain spices that will be more predominant then if I was making a Japanese dish. But my spice cabinet still has all the spices I could possibly need for all the different types of dishes that I may cook. But that doesn’t mean that I use every spice I have just because they are in the cabinet. And one particular spice isn’t in every recipe (except if you are cooking Italian, then garlic goes in everything!) But I think you get the idea now.

BASIL AND GARLIC MINESTRONE
3 ounces pancetta, finely chopped, 3 cloves garlic, 1 cup olive oil, pinch of oregano, 1/2 cup Italian tomatoes, 1 tsp chicken bouillon, 1 tablespoon salt, pinch of pepper, 10 ounces penna pasta, 3 tablespoons fresh basil, 1 cup Parmesan cheese

Another good example would be that of an orchestra. There are many instruments that are available in the orchestra. Every possible score won’t use all the possible instruments all the time. In fact as the particular score proceeds there will be varying contributions from each of the instruments. Sometimes the horns will be dominant and sometimes the strings will be the strongest instrument. And as the score progresses, each of the instruments will participate to varying degrees depending on the particular musical arrangement. And sometimes some of the scores will not need any contribution from certain instruments (like the guy in the back with the triangle). But the orchestra always consists of all the members of the orchestra – but all of the instruments have varying participation for each of the different scores that the orchestra plays.

Well the same is true of the response of any structural system. The total response of the system is made up from linear combinations of a subset of the total number of modes that possibly exist in the system. Not all modes contribute to the response for every forcing condition that might exist. Only certain modes may contribute substantially and some others may contribute a little bit and yet some others may not have any contribution whatsoever. This amount of contribution of all the modes changes depending on what loading condition is considered. So certain modes have different modal participations depending on what loading scenario is considered.

But the important thing to understand is that the mode shape plays an important part in determining the amount of force that is appropriated to the modal oscillator in modal space along with the distribution of force on the physical system. So understanding the modes shapes also helps to identify the force applied to all the modes of the system.

I hope this explanation helps you to understand modal participation a little better. If you have any other questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Hey – I ran a test and got some extra modes I didn't expect. Does the overall test set up have any effect? This really needs some discussion

The test set up can have a significant effect on the overall test results in some circumstances. You must be very careful with this. From what you explained to me, it seems that in this case the test conditions related to the test set up changed during the course of the test and had an effect on the modes that were extracted.

I guess the first thing to do is to recreate some of the data that led you to believe that there may have been some problems with the test that was conducted. When you first showed me the original stability diagram, you indicated that you didn't expect that there would be several frequencies with very close and almost repeated roots. The stability diagram in Figure 1 had what appeared to be very close frequencies.

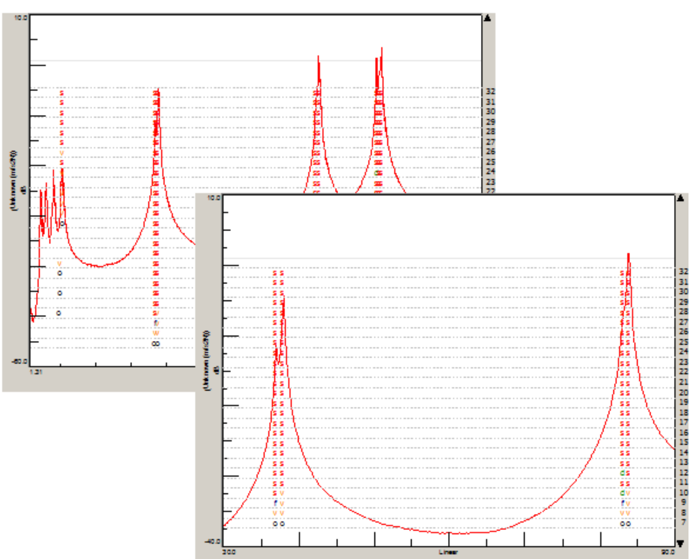

Figure 1 – Stability Diagram over 130 Hz Band and over 30 to 90 Hz Band

It is very clear that there are multiple roots at each of the frequencies. But this was not expected for this particular component. As we discussed the data you mentioned that there were multiple reference accelerometers but that the data from each reference was collected at different times and was not collected all at the same time.

This fact alone starts to lead me to believe that maybe there was a change in the overall test set up between the first test and the second test. Taking a closer look at that stability diagram especially in the range of 30 to 90 Hz in Figure 1 shows that there are indeed multiple poles at 36.96 Hz and 37.96 Hz with 3Δf spacing and multiple poles at 83.08 Hz and 83.8 Hz with 2Δf spacing. But the question is – are they really separate modes or is it a problem with the test set up?

In order to sort this out, we probably need to look at the data in a little more detail and try to sort out what may be happening with this data. Well the first thing to do is to look at the data from each of the references separately. When we do this what we see is that each of the individual references sees only one pole at each of the frequencies and that each of the references predicts that frequency at different values for each of the separate references. This seems to indicate that there is some type of a shift in all the modes of the system from the time the first test was conducted and when the second test was conducted. This is clearly seen in Figure 2 for the 30-90 Hz band.

In order to further confirm this, the rigid body mode frequencies were also evaluated. Now here is where there are some more obvious differences. When evaluating the two references separately, clearly there is a significant shift or change in the rigid body modes of the system as seen in Figure 3.

Copyright 2010
Pete Avitabile

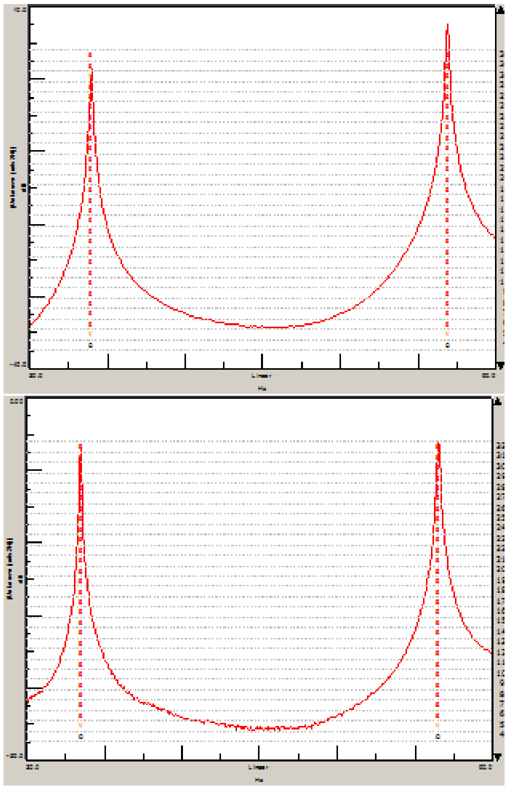

Fig 2 – Stability Diagram for Each Reference Separately

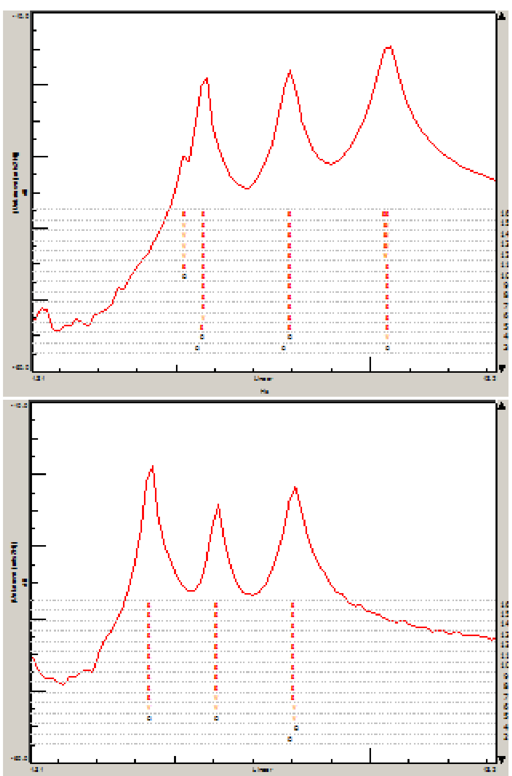

Fig 3 – Stability Diagram for Rigid Body Frequency Range for Each Reference Separately

So the change in the rigid body modes from the two different tests shows a dramatic change in the frequencies. From one reference set of data the rigid body modes observed were 4.3, 6.1and 8.1Hz and from the other reference they were 5.7, 7.9 and 10.4 Hz which is a significant shift in these frequencies from one test to another.

After some detective work on the test set up, the cause of the problem was likely due to the fact the air pressure in the structure support system was not maintained at the same pressure for both tests and caused the support stiffness to change significantly between the two tests. This is further highlighted when the sets of data are overlaid and compared between the two references as shown in Figure 4. Note that the blue FRFs are related to one reference set of data collected on one day and the red FRFs are related to the other reference collected on a different day. This highlights the obvious inconsistency in the two data sets that were collected on different days resulting in obvious differences.

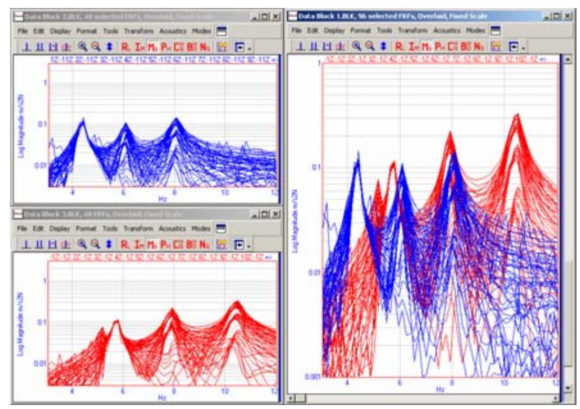

Fig 4 – Comparison of the FRFs from Two References

So now it is very clear that there are differences from the test set up which resulted in not only the rigid body modes shifting but also the flexible modes were significantly affected by the slight change in the test set up conditions.

Of course we could argue about which are the "real" set of frequencies for the structure but we really don't know which are correct because the test set up does have an effect on the frequencies observed. What is more important to realize is that it is imperative to make sure that all the data is collected in a consistent fashion. This would have been best accomplished by collecting all of the data at the same time using the multiple reference impact test technique. This would have prevented the inconsistency that resulted from conducting two separate tests on two separate days in what ultimately appears to be in two essentially different configurations.

I hope this explanation helps you to understand the need to collect data in a consistent manner. If you have any other questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD by Pete Avitabile

Illustration by Mike Avitabile

Will the support mechanism have any effect on FRFs? Does bungie cord vs. fishing line make any difference?
It really depends - Let's discuss this.

Well, we have discussed the effects of the test set up before. But maybe we need to shed a little more light on this.

For your specific problem, there was some concern as to whether or not the damping may be affected by the way the structure is supported – for instance, whether we use an elastic cord or maybe fishing wire as suggested. In order to see what kind of an effect this may have, let's test a simple structure and see what effects may result from several different mechanisms to support the structure for modal testing.

For this test, we used a simple plate that was basically hung horizontally from the four corners of the plate using long flexible elastics (rubber bands) and nylon cord (fishing wire) in one set of tests. But then in a second set of tests, the plate was hung vertically and only supported from two corners in a pendulum type fashion. In all cases, the plate response was measured with an accelerometer that was fixed on the plate for all testing and an impact excitation was used to provide the input to measure the frequency response function.

As far as signal processing parameters are concerned, care was taken to assure that the time sample was long enough so that the response signal was essentially zero by the end of the time sample; this then ensured that the FFT was not affected by leakage and no weighting function (windows) were applied to the measurement.

The plate along with several different support configurations are shown in Figure 1. The impact hammer and accelerometer are also shown for reference. Rather than discuss the results for each of the individual tests performed, the results will be discussed for all four tests conducted.

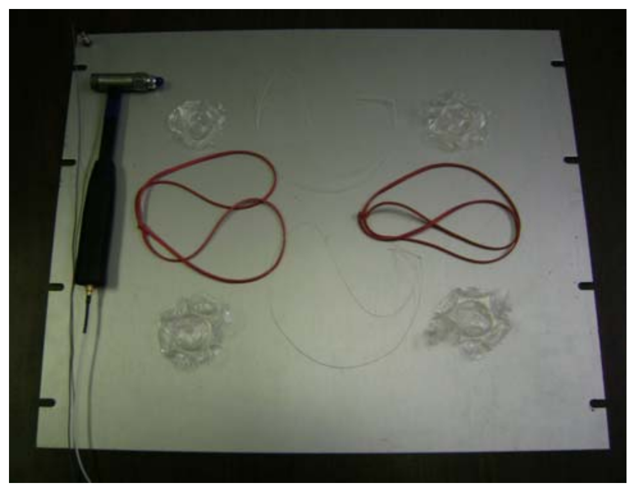

Figure 1- Plate with elastic cord and fishing wire
(as well as some bubble wrap pieces for other testing)

Figure 2 shows an overlay of the measured frequency response functions as well as the results of the frequency and damping obtained from the rational fraction polynomial curvefit (using the MEscope software). Only the first five modes were considered

At first glance, it is very obvious that there is some difference between the four different frequency response functions obtained from the four different tests performed.

In the case of the plate testing performed with the plate hung in a vertical orientation (in a pendulous configuration), there seems to be little difference in the measured frequency response functions; the blue frequency response function is with the elastic rubber bands and the green frequency response function is with the fishing line.

The frequency values are very similar for these two tests and the damping values are also similar but do show some differences. For this configuration, the results are not exactly the same but reasonably close to the same values (at least for the purposes of what is being presented here).

Now when the plate is supported horizontally, a different situation exists. For the elastic rubber band supports (red), the frequencies are very similar but the damping values are somewhat different with the values appearing to be higher than those for the vertically hung plate.

And then when looking at the fishing line with the plate hung horizontally (black), there is a very clear increase in the natural frequency for all the modes investigated. The damping values are also higher than any other configuration investigated.

So from these very simple, quick tests that were run, there is definitely a difference in the resulting frequency response functions and extracted parameters depending on how the test is set up.

I hope this explanation helps you to understand that the test set up can have a dramtic effect on some of the critical parameters of interest typically obtained from a modal test. If you have any other questions about modal analysis, just ask me.

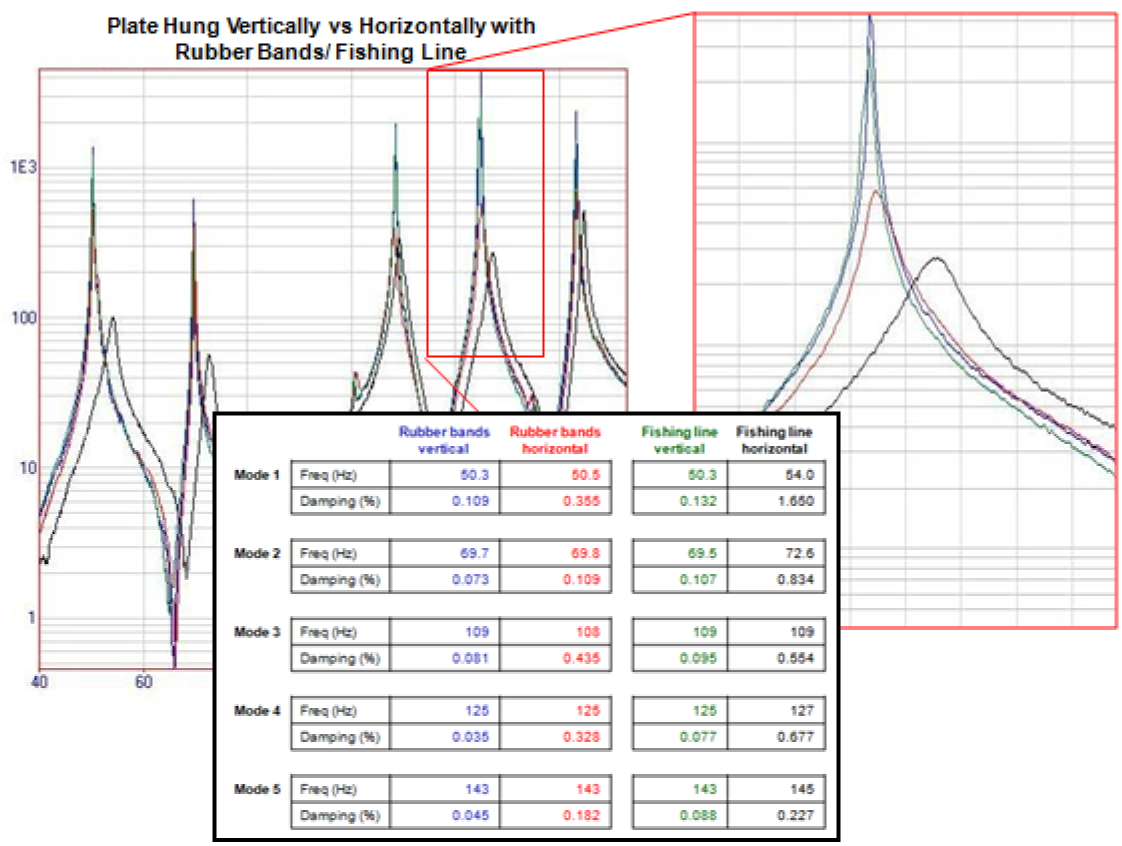

		Rubber bands vertical	Rubber bands horizontal	Fishing line vertical	Fishing line horizontal
Mode 1	Freq (Hz)	50.3	50.5	50.3	54.0
	Damping (%)	0.109	0.355	0.132	1.650
Mode 2	Freq (Hz)	69.7	69.8	69.5	72.6
	Damping (%)	0.073	0.109	0.107	0.834
Mode 3	Freq (Hz)	109	108	109	109
	Damping (%)	0.081	0.435	0.095	0.554
Mode 4	Freq (Hz)	125	125	125	127
	Damping (%)	0.035	0.328	0.077	0.677
Mode 5	Freq (Hz)	143	143	143	145
	Damping (%)	0.045	0.182	0.088	0.227

Figure 2 – Comparison of FRFs for two different test set ups with two different support mechanisms along with frequency and damping extracted from the measured frequency response functions

Copyright 2010
Pete Avitabile

Illustration by Mike Avitabile

Does the structure need to come to rest between impact measurements? Doesn't the damping window take care of that? This is important - Let's discuss this.

So now let's talk a little bit about what kind of problems can result from the measurement you described. The measurement you made was on a very lightly damped structure and in order to prevent leakage most likely an exponential window is needed. From what you described, the measurement is likely to look like what is shown in Figure 1.

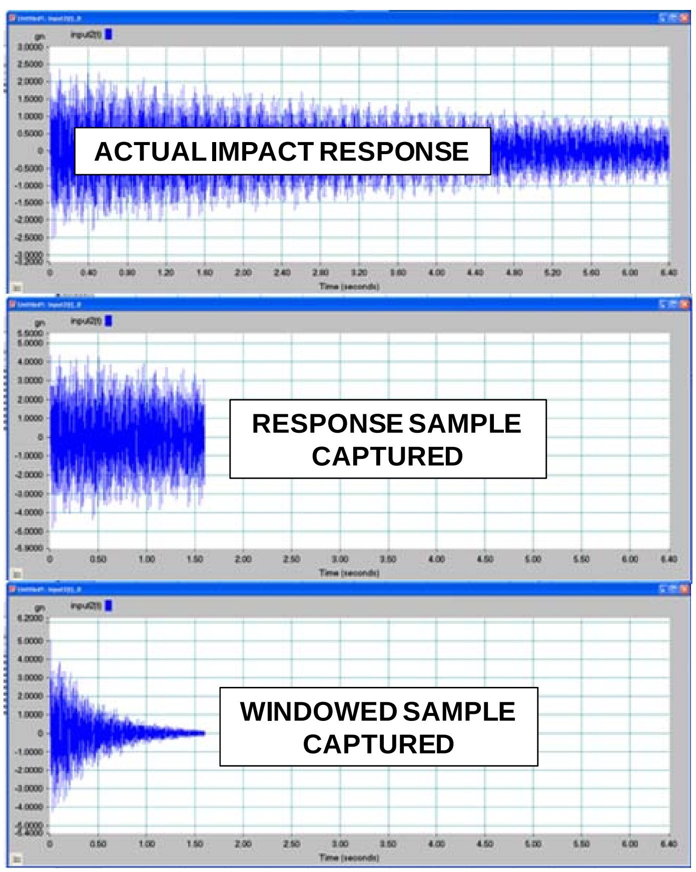

Figure 1 – Impact Response for One Sample

Now the upper trace shows the time response for a much longer sample than what you used for acquisition. The middle trace is what was actually captured from the FFT for the T seconds of data collected. And the lower trace is the time response with the window applied to the output response. So up until this point everything looks reasonably fine.

From what you described, the averaging was performed by impacting the structure and measuring the response for a series of many averages. A sample of these averages is shown in Figure 2. As far as you were concerned, the window was applied and the response was measured and averaging was performed to obtain the data described.

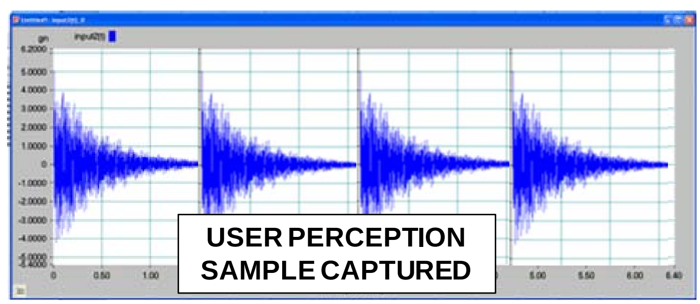

Figure 2 – User Perception of Impact Averaged Response

However, the frequency response function (FRF) that resulted (shown in Figure 3) did not look very good overall and the coherence was not very good either. In addition, this drive point FRF lacks the typical measurement characteristics that is expected with strong resonant and anti-resonant frequencies.

Copyright 2011 Pete Avitabile

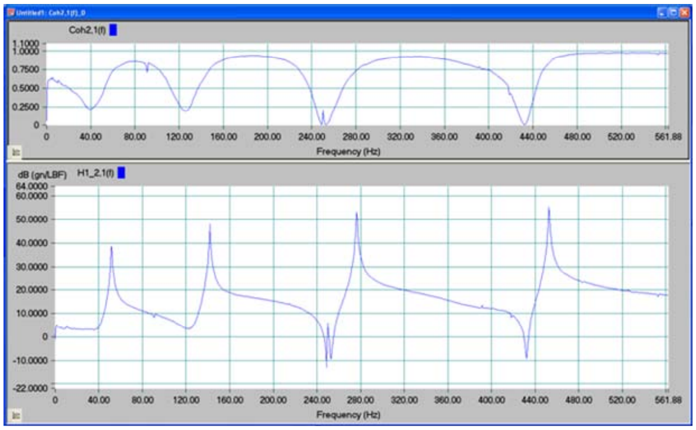

Figure 3 – FRF and Coherence from Initial Measurements

So what could possibly have gone wrong here. To understand what happened, we need to go back to the formulation of the system transfer function. When we write the equation of motion and perform the Laplace Transform we get

$$\left(ms^2 + cs + k\right)X(s) = f(s) + (ms + c)x_0 + m\dot{x}_0$$

and we get fairly comfortable writing the system transfer function as

$$H(s) = \frac{X(s)}{F(s)} = \frac{1}{ms^2 + cs + k}$$

but in order to do that we have to realize that the extra terms on the right hand side of the equation have been eliminated. It turns out that these are the initial conditions for the transformation.

So ignoring those terms assumes that the initial conditions are zero. But the problem is that the way that the original measurement was acquired, the structure's response in between each individual impact was assumed to be zero. While a damping window was applied to the data and it looks like the response has been decayed to zero, that is only with respect to the software used to acquire the data.

Actually what most likely happened is that the measurements were taken in close succession and the actual response of the structure never actually died out before the next sample was taken. This is schematically shown in Figure 4. So what happens is the response of the second average is contaminated by the remaining response of the first impact. And the third average is contaminated by the remaining response of the first and second average. And this continues for all the averages you took. So basically the measured response of each average (after the first average) is not the result of the impact excitation for that particular average and it is the response due to other than the measured force for that particular average. So that is why the coherence is so poor.

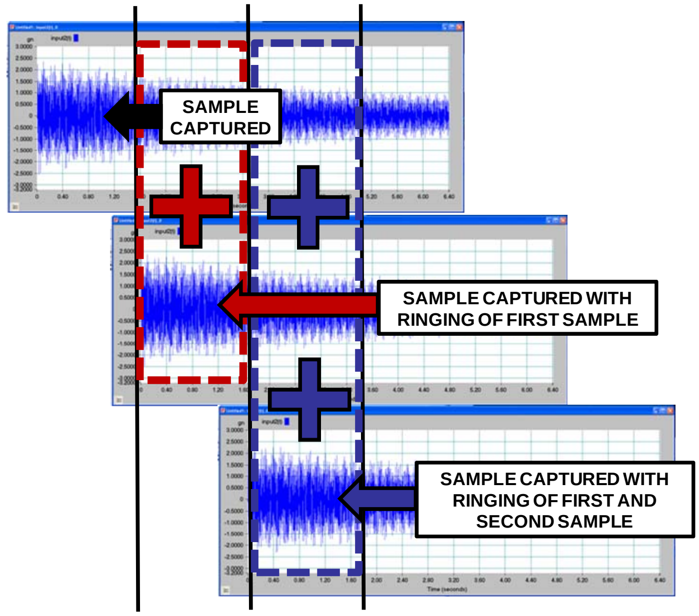

Figure 4 Impact Response from Structure Standpoint

To confirm that this is the case, another measurement was made where sufficient time was given to allow the structure to return to a steady state (no response) condition. The resulting FRF and coherence is shown in Figure 5 and it is very clear that this measurement is far superior to the one shown in Figure 3.

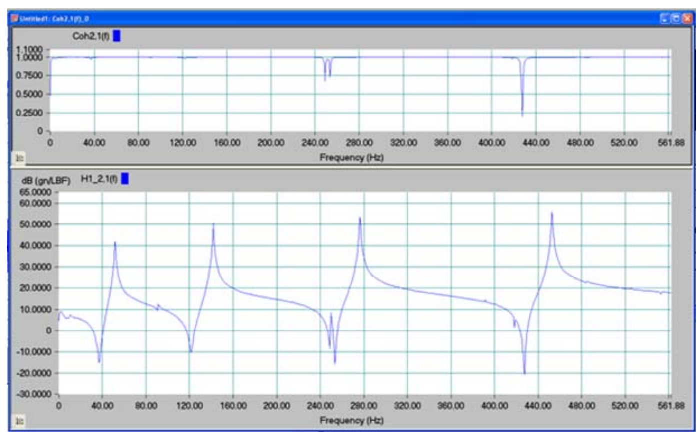

Figure 5 – Good FRF and Coherence from Proper Technique

I hope this explanation helps you to understand that the formation of the FRF is subject to some of the assumptions made in the formulation of the system transfer function, namely, that the initial conditions are assumed to be zero. Once these restrictions are observed, then proper measurements can be acquired. If you have any other questions about modal analysis, just ask me.

Copyright 2011
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Why can't I run a modal test with one big shaker and just "crank up the signal"?
That isn't a good idea - Let's discuss this.

OK – so we need to talk about a few things here. Many times people who get involved in modal testing sometimes come from the "vibration qualification world" and have a completely different mentality compared to the "modal world".

In vibration qualification testing, a large shaker is used and a test article is normally hard mounted to the top surface of the armature and then some base excitation is applied and usually monitored by controlling some prescribed acceleration. The device under test (DUT) is normally subjected to some operating environment, generic spectrum or some excessive environment to determine if the equipment is suitable for the intended service. A typical test schematic is shown in Figure 1 showing the shaker system and a test article mounted to a fixture and expander head which is all attached to the shaker armature.

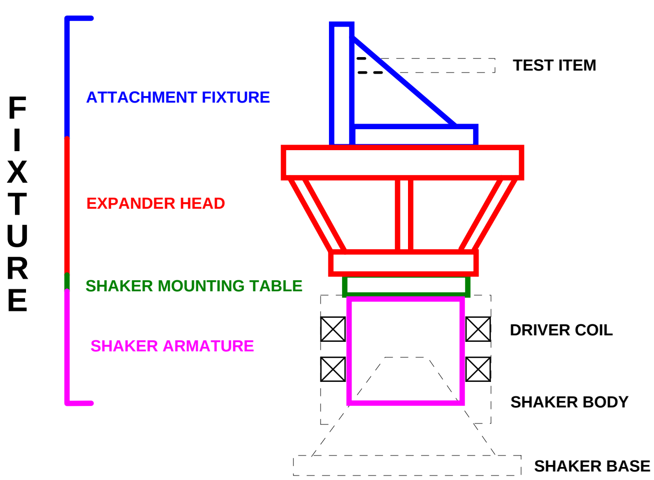

Figure 1 – Typical Shaker Qualification Test Set Up

This is a completely different test than what we try to do with modal testing. In modal testing, the shaker is attached to the structure with a long rod commonly called a stinger or quill. The force imparted to the structure is measured with a force gage or an impedance head mounted on the structure side of the shaker exciter set up. This is shown in Figure 2.

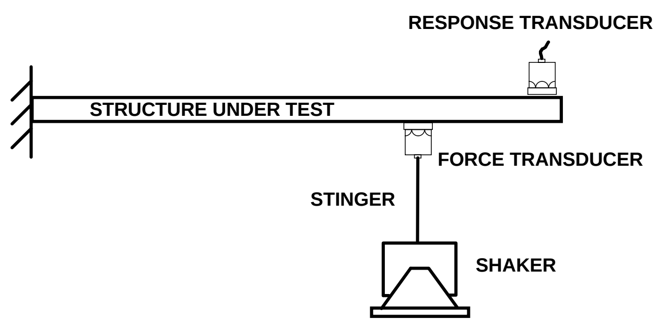

Figure 2 – Typical Modal Test Set Up

In modal testing, the intent is to use lower levels of excitation and identify system characteristics – the test is not intended to provide operating level input excitations. In fact, if higher levels are used then sometimes nonlinear characteristics of the structure are excited and the overall measurement becomes distorted and not particularly useful for modal parameter estimation.

Now of course it also depends on what kind of structure you are testing. If it is a very simple component of a larger system and the component itself is fairly linear then there is no problem using a single shaker with an appropriate force level specified.

Copyright 2011
Pete Avitabile

But when the structure becomes more complicated with many components assembled together to form a system, then the ability to provide a force excitation to measure all the locations on the structure to identify the mode shapes can become more difficult. This can then be compounded when the various components are attached with mounting devices to isolate all the components from each other. The problem becomes that it is very hard to provide an adequate excitation from one shaker location and be able to make adequate FRF measurements at all the response points to be measured. Then it becomes necessary to “crank up the signal” to be able to get measurable vibration at all the response locations. When this is done, then it is very likely that nonlinearities will be excited and then the overall measurement will be degraded.

I have been involved in many tests where this is the case. Just recently, a test on a large propulsion system had an isolation system that was intended to isolated all the components for vibration transmission considerations. The actual data can’t be shown but a laboratory structure with several components attached through an isolation system was used to illustrate the problem with using just one shaker to excite the system.

The laboratory structure is shown in Figure 3 with three plate components attached with isolators to a larger frame structure.

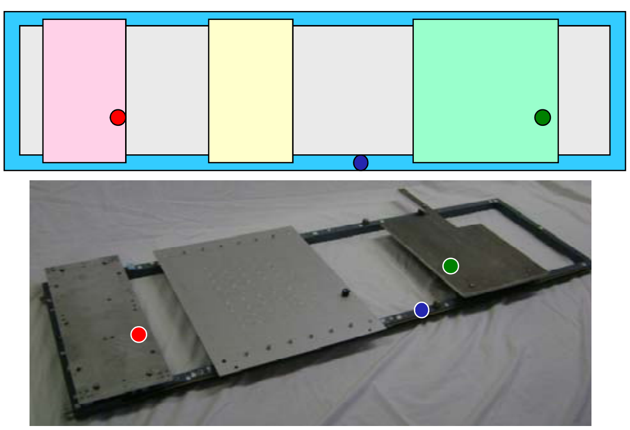

Figure 3 – Laboratory Structure with Isolated Components

A single shaker was attached on the main frame and FRF measurements were made. In addition, a three shaker MIMO test was also conducted to compare the measurements obtained. Figure 4 shows a typical drive point measurement (on the main frame in this case). The FRF in red is related to the SISO test. Also shown in the figure is the same FRF (black) obtained from the three shaker MIMO test that was conducted with much lower overall excitation applied to the structure.

In looking at the FRF, it is very clear that the SISO FRF obtained is not the same quality as the MIMO FRF obtained using lower overall shaker excitation levels; this is especially true when looking at the coherence. A cross measurement that is even poorer quality is shown in Figure 5 and again the FRF and coherence is seen to be much worse from the SISO test.

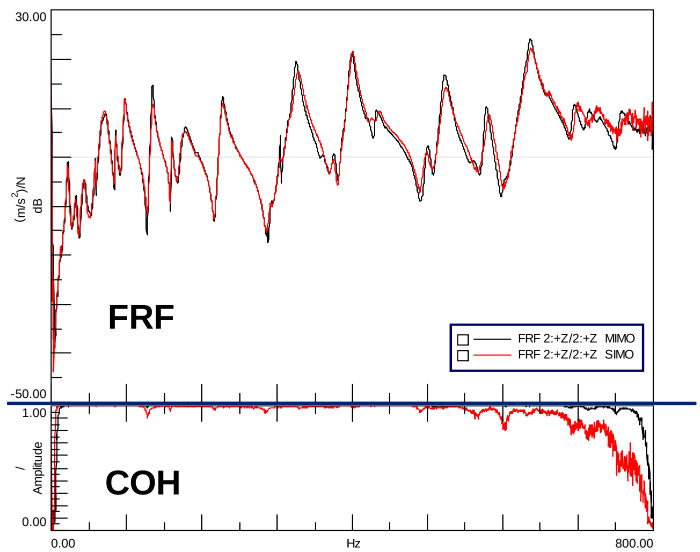

Figure 4 – SISO vs MIMO FRF Drive Point Measurement

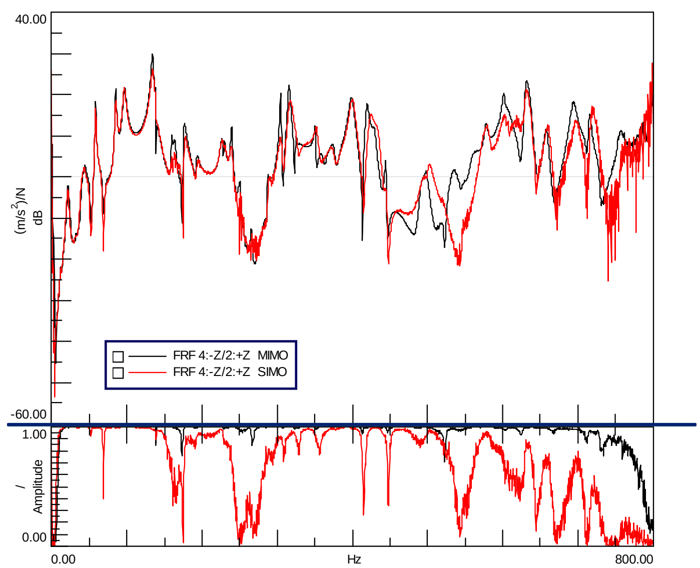

Figure 5 – SISO vs MIMO FRF Cross Measurement

Now this is a very short discussion but in the next article this will be expanded to further illustrate that using one shaker with higher level of excitation is not preferred to the multiple shaker excitation approach with lower excitation levels.

I hope this explanation helps you to understand that using one shaker with the “signal cranked up” will not provide a good measurement overall. If you have any other questions about modal analysis, just ask me.

Copyright 2011
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD by Pete Avitabile

Illustration by Mike Avitabile

Is there really a difference running a modal test with MIMO as opposed to SISO?
You bet there is - Let's discuss this.

Well I hear this question a lot. I guess mainly it is due to the fact that additional hardware and software needs to be purchased to do MIMO testing. And so it needs to be justified that MIMO is really much better than SISO.

So last time we discussed the fact that using a single shaker and "cranking up the signal" could likely excite nonlinearities in the structure and that the overall FRF would likely be affected by this. So from that data, it is very obvious that the single shaker test may not provide the best set of FRFs for modal parameter estimation.

Another approach that I often see people try is to use one shaker but then move the shaker to all the different locations for the desired number of references. On the surface this may seem to be a useable solution but there are limitations to this approach. The first problem is as we already discussed – the level of force with one shaker will need to be much higher in order to get adequate response at all the measurement locations in the structure.

Now a single shaker may work for structures that are not very complicated with many components and substructures that are attached in a manner to minimize the flow of energy through the subsystems. The situation is much different when the components are isolated from each other. In these situations it is very hard to get adequate response throughout the structure with just one excitation source. In these cases, multiple references are needed.

So let's discuss the difficulty with the data collected from a single shaker that is moved to the different reference locations to collect the multiple referenced FRF data. Unfortunately, many of the tests and data sets that I have seen are not available for public release. So instead, a simpler structure that contains all of the features typically seen in complicated structures with components and subsystems that are mounted to minimize the flow of energy in the structure (isolated) was assembled in the lab.

The laboratory structure is shown in Figure 1. This structure was assembled with 3 components mounted to a frame. Each of the components was mounted with a very soft mount, an intermediate mount and a very hard mount. Now the main frame and the attachments do have some of the typical "pesky" rattles and noise that plague the collection of FRF data; no attempt was made to minimize any of these noise sources and in fact they are welcome to illustrate a typical structure measurement.

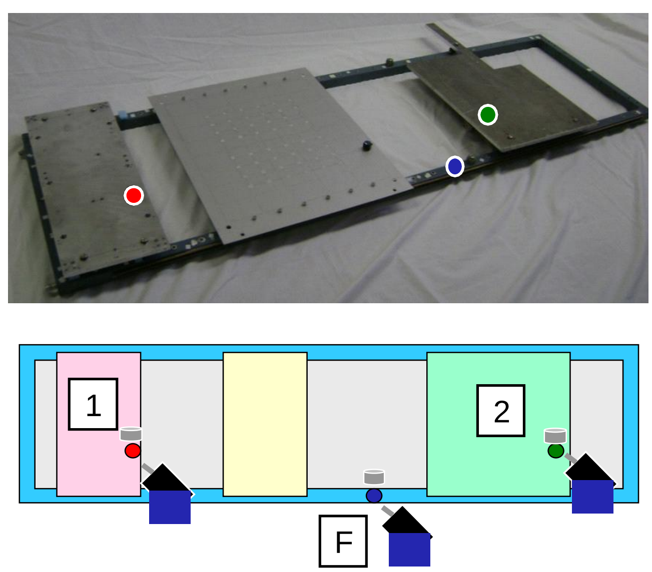

Figure 1 – Laboratory Structure with Isolated Components

The structure was tested in many different configurations and only a few of them are presented here to show the problem with the FRFs collected with single shaker set up and with a multiple

Copyright 2011
Pete Avitabile

shaker set up.The three shaker reference locations are shown in Figure 1.

Now separate tests were run with each of the individual shakers used to collect FRF data from the structure as well as a multiple reference MIMO set of data. However, in order to make the best possible measurements, the individual SISO shaker tests needed more force excitation level to make suitable measurements; the MIMO configuration needed lower force levels in order to make accepatable FRF measurements.

In order to compare all the measurements, several FRFs were compared. In all FRFs the reference was made to the shaker mounted on the frame; the other references could be used and yielded essentially the same results as those presented next. In Figures 2, 3 and 4, the FRF in red was obtained from the SISO test and the FRF in black was obtained from the MIMO test. Two measurements are shown from the frame to the attached components and one of the measurements was a drive point on the frame itself.

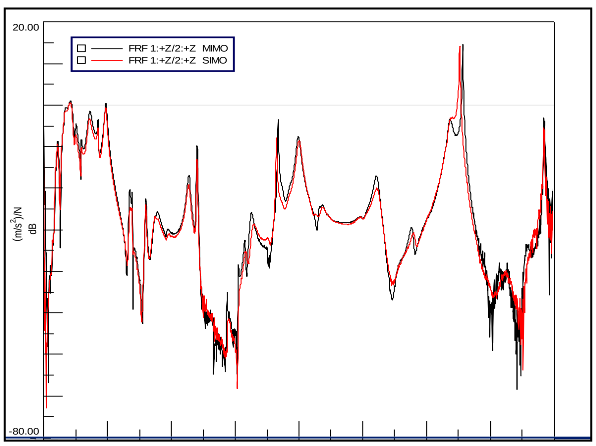

Figure 2 – FRF Component (1) to Frame (F) Reference

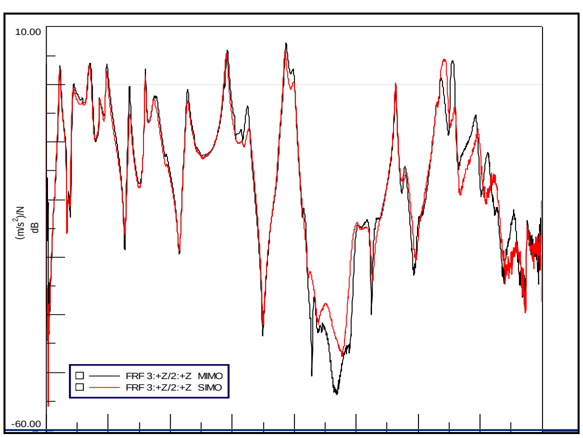

Figure 3 – FRF Component (2) to Frame (F) Reference

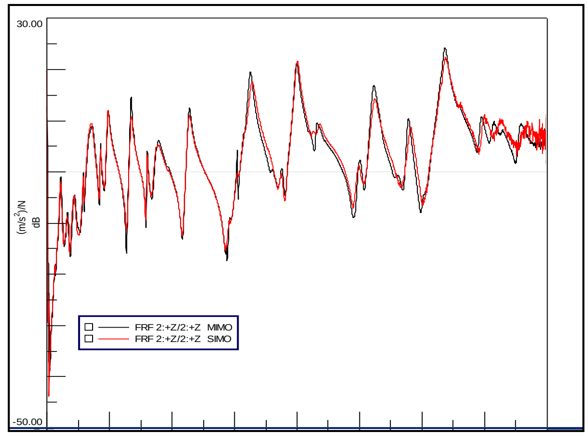

Figure 4 – FRF Frame (F) to Frame (F) Reference

So at first glance, the data in Figures 2, 3 and 4 don't look terribly different and I know that many people might actually say that data is just fine. But if you start poking around and looking more closely at some of the reciprocal FRFs, then it becomes very clear that the peaks of the FRFs from the SISO tests don't line up with each of the different SISO tests that were conducted. This then causes a discrepancy or inconsistency between the different data sets. A few of these are shown in Figure 5.

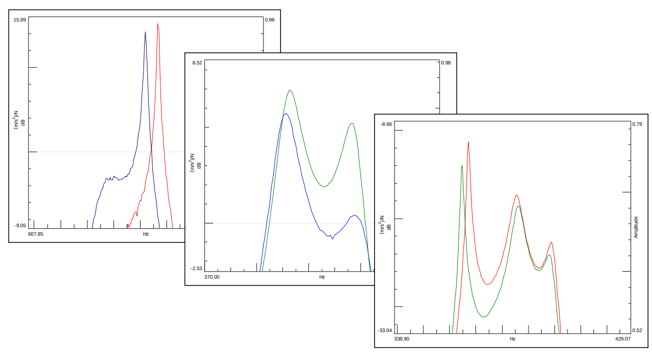

Figure 5 – Close Up of Several FRFs Showing Inconsistency

The bottom line of all of this is that the reciprocity between the different data sets is not satisfied! This will have a significant effect when modal parameters are extracted (and will be discussed in the next article).

I hope this explanation helps you to understand that using one shaker at different locations does not necessarily provide the best data. MIMO tests are needed in order to provide more consistently related FRF data from multiple references. If you have any other questions about modal analysis, just ask me.

Copyright 2011
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD by Pete Avitabile

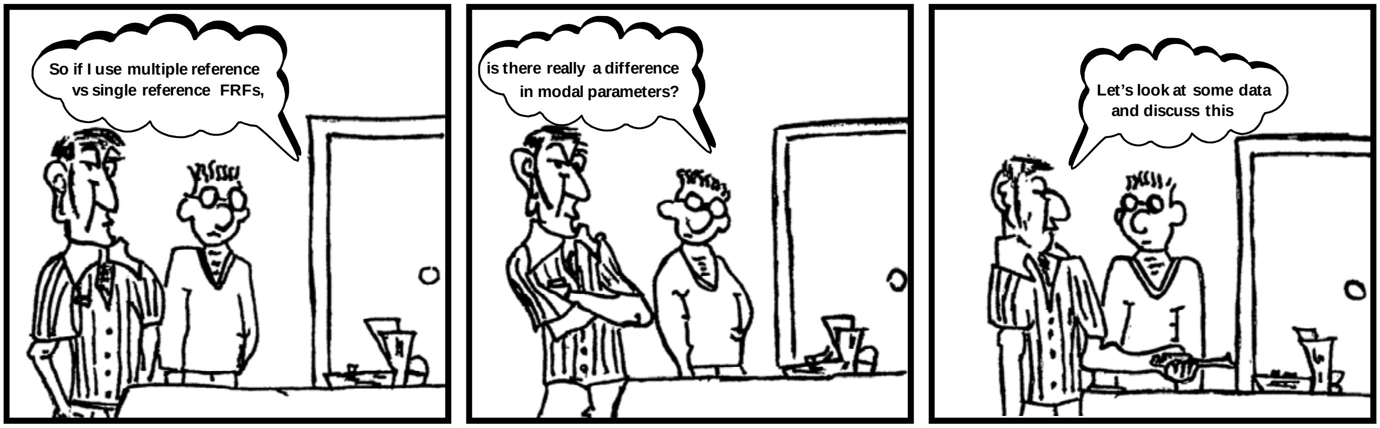

Illustration by Mike Avitabile

So if I use multiple reference vs single reference FRFs, is there really a difference in the modal parameters?
Let's look at some data and discuss this.

Well from a purely theoretical standpoint, you are able to extract modal parameters from any reference location as long as it is not at the node of a mode. But of course, the theory is perfect and we need to consider the practicality of the measurements we can make on any real structure.

In the last two articles, several aspects of the measurements were discussed. Overall, the FRF measurements are always much better overall when the data is collected simultaneously in a MIMO test. If a single shaker is used two issues arise that tend to provide FRFs that are not of the best quality for modal parameter estimation.

In one case, a single shaker needs to have a higher excitation level in order to make adequate measurements but this invariably cause nonlinearities to be excited and generally tends to increase the variance and the FRF measurements are not as good as one would like.

The second issue noted was that when multiple referenced data was formed from single reference tests, generally the FRFs are likely to not be related in a consistent fashion and the FRF peaks may show some slight variation in frequency. While the structure may be time invariant, the test set up can have an effect on the measured FRFs when the tests are obtained from separate tests. And then another variability can result in the fact that all the data is collected at different times and there may be slight environment changes that could compound this problem.

In order to have some continuity with the two previous articles, the test data for this discussion will be the same data previously used. Of course we noted that there were some shifting of the frequencies for some of the modes and that the reciprocity was not satisfied for all the SISO data that was collected and used to form the multiple reference data set.

The laboratory structure is schematically shown in Figure 1. Three reference sets of data were collected using SISO methodology in three separate tests; data was also collected for all three references simultaneously using a MIMO methodology.

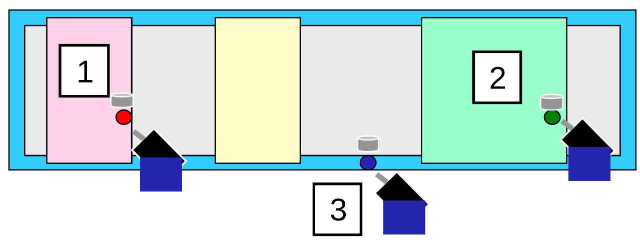

Figure 1 – Laboratory Structure with Isolated Components

So the previous articles discussed some of the measurement issues. Here what I want to do is to process some of this data to show some of the difficulties in identifying modal parameters. In all cases the stability diagram will be used to show how some of the variance on the data will present challenges for identification of the system poles.

So the first thing to try is to take all three separate SISO test FRFs and form one set of multiple reference data for processing. (And please note that I am not calling this MIMO data because it was actually all collected separately.) The first step in the modal parameter estimation process is to identify the system poles. This is usually done using the stability diagram with an overlay of one of the mode indicator function; for the plots here, the CMIF is used in all cases.

Figure 2 shows the stability diagram for this case. While this diagram may be acceptable to many, there is definitely some variation in the system poles and there is not a good strong

Copyright 2011
Pete Avitabile

stable pole identified for every one of the system poles. (As we reprocess this data, the improvement in the stability diagram will be seen.)

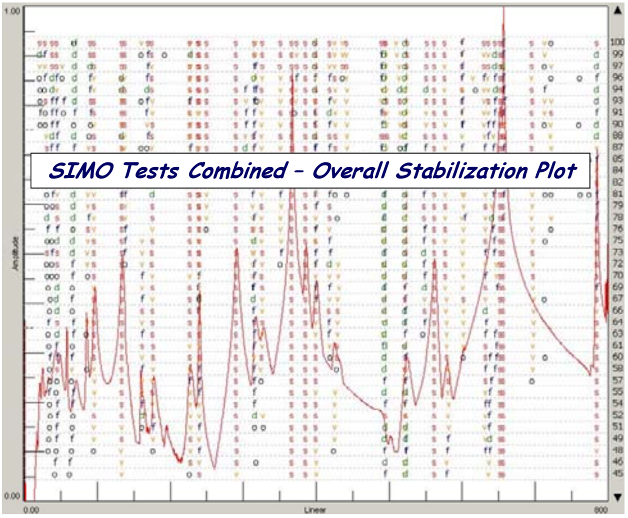

Figure 2 – Stability Diagram for Combined SISO FRFs

But before we look at the MIMO data set, let's look at the individual SISO data sets alone. Figure 3 shows the three separate SISO test data sets processed individually before being combined into one multiple referenced data set. The thing that is very obvious is that the stability diagram for each of the separate test cases produces very consistent stable system poles. There is no question what the system pole is when the data looks this good.

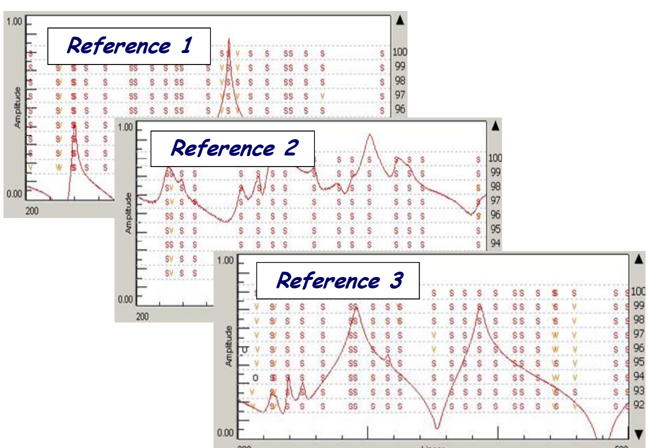

Figure 3 - Stability Diagram for Three Separate SISO Tests

So why are the individual data sets (Figure 3) so obvious as to what the system poles are and is not as clear cut (Figure 2) when all the data sets are combined? Well remember that each of the individual SISO data sets were collected consistently for each of the individual SISO tests. Even in light of some of the noise and nonlinearities that were discussed in the previous two articles, the identification of the system pole is not difficult at all. But when all the individual SISO data sets are combined, there is no guarantee that the data will be consistently related between the three different SISO tests that were performed. And in fact, the shifting of the peak of the FRF measurements was pointed out in the previous article. This shifting was noted in several measurements such as the reciprocal FRFs. So the main culprit here is the fact that the data was collected in three separate tests and the data was not necessarily guaranteed to be consistently related. This is why the stability diagram becomes a little more difficult to interpret and the system pole identification is not as straight-forward when this happens.

To confirm this, the MIMO data set (where all the data is collected simultaneously and in a consistent fashion) is used to generate a stability diagram. This is shown in Figure 4. This stability diagram is much better than the one shown in Figure 2. Of course there are some frequencies that are still not perfect – but this is much better than the previous scenario where the data was collected separately and the consistency of the data could not be guaranteed.

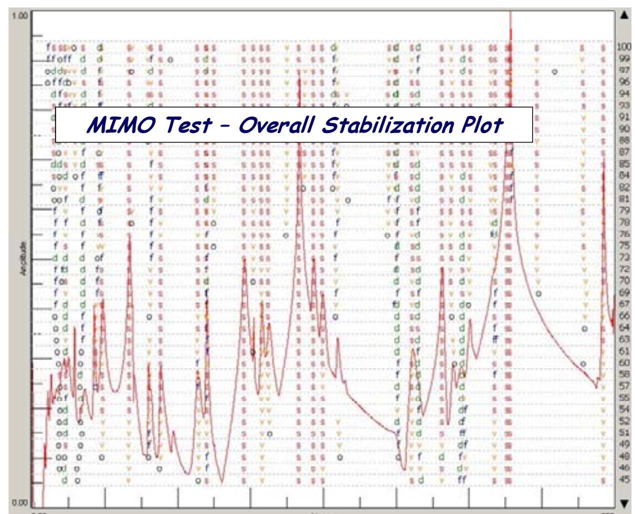

Figure 4 – Stability Diagram for MIMO FRFs

So the real problem here lies with the data. The FRFs must be collected in a consistent manner. The SISO test can not provide data with this consistency but the MIMO test generally does due to the nature by which data is collected.

I hope this helps to explain why collecting data in a consistent fashion is very important for any multiple reference test. Collecting data from separate single input reference tests may not provide the best data for modal parameter estimation. MIMO tests are needed in order to provide more consistently related FRF data from multiple references. If you have any other questions about modal analysis, just ask me.

 by Pete Avitabile

Illustration by Mike Avitabile

Someone told me that you can't accept a FRF when the input spectrum has more than a 20 dB rolloff
Let's discuss this and consider difficulties.

Well this is a very touchy topic with many people. I remember back when some people claimed that there could be no more than 1 dB rolloff on the input spectrum. Well this was a very harsh criterion and in fact this actually excited many modes well outside the band of interest and could potentially saturate the accelerometers thereby making a poor measurement.

Now let's understand why we even try to make rules to live by in modal testing. Many times there may be some tests where we may want to provide some guidance as to typical ways to conduct the test. This is intended to protect us from making measurements that may not be particularly useful in some testing scenarios.

But the problem is that some of these "suggested rules" get interpreted as if they are cast in stone as if they were the Ten Commandments. And maybe at the time the "suggested rules" were made might have been back 20 or more years ago when instrumentation was not as good as it is today and back when 12 bit acquisition systems were very commonplace. But maybe those rules are not as critically needed today with much better instrumentation and 24 bit acquisition systems commonly used.

So while I think "suggested practices" are clearly needed, I also think that we need to realize that they are suggested and we need to understand how to interpret if the measurement is useful or not.

So to illustrate this, a simple plate structure was tested with an impact excitation technique. Two tests were performed. One test with a harder tip with an input spectrum with a 10 dB rolloff over the frequency range of interest. The second test was with a softer tip with 30 to 35 dB rolloff - approximately 10 dB rollover over the first third of the spectrum, approximately 25 dB rolloff over the next third of the spectrum with the remaining rolloff over the last third of the spectrum.

The hard tip and soft tip input force spectrum are shown in Figure 1.

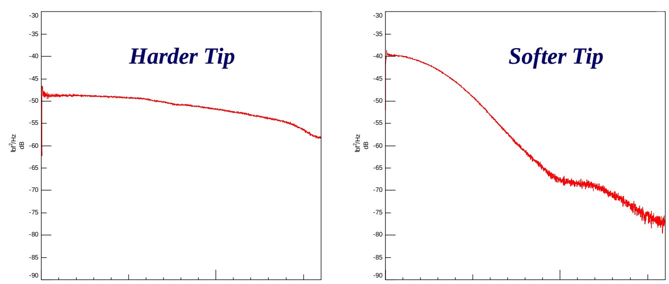

Figure 1 – Comparison of Hard Tip and Soft Tip Force Spectrum

The drive point FRF for the modal test with the harder hammer tip is shown in Figure 2 and the drive point FRF for the test with the softer tip is shown in Figure 3. Now clearly, the FRF with the harder tip is overall a much better measurement as evidenced by the coherence. One thing to notice in the FRF with the softer tip is that the measurement at the higher frequency shows some variance on the FRF overall and there is a slight degradation of the coherence at the higher frequencies.

Now we have to ask ourselves exactly why are we taking the measurements and performing the modal test. Sometimes tests are performed to obtain very high quality measurements for very specific applications. But sometimes measurements are made to get a general understanding of the generic characteristic shapes for the structure and maybe do not need to have the same high quality as some other tests that we may need to perform.

Think of it like buying lumber for a home building project. We don't always need knot free wood for the entire project. Sometimes wood of a lower quality is more than adequate for the project undertaken.

Copyright 2011
Pete Avitabile

Now I would always like to take high quality measurements all the time but sometimes the cost involved in doing that makes the test prohibitively expensive. So let's see just how good or bad these measurements are. Modal parameters were estimated from both sets of measurements. The generic mode shapes are shown in Figure 4 for reference. A MAC was also computed for the two sets of mode shapes and is shown in Table 1.

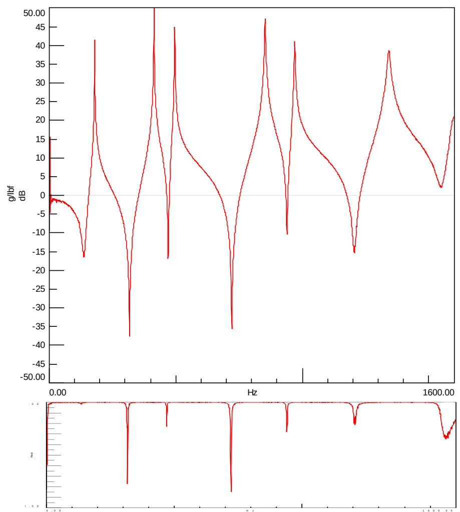

Figure 2 – FRF and Coherence for Hard Tip

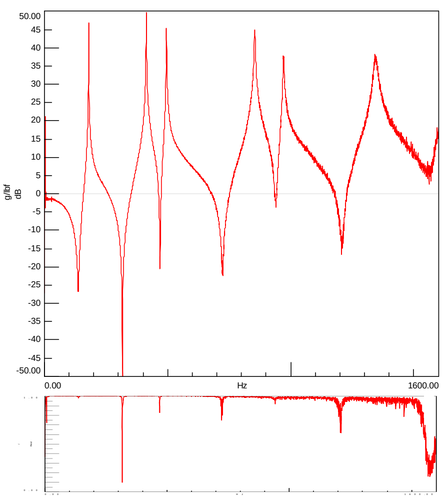

Figure 3 – FRF and Coherence for Soft Tip

Now the mode shapes are seen to be essentially the same from both tests. So the FRF measurements seem to be adequate for the simple assessment of mode shapes for the structure.

Now I am not advocating that this type of input force spectrum rolloff is acceptable but sometimes there is still useful information that can be obtained from data. So while we have "suggested rules" that doesn't necessarily mean that the data is not useful. But we do need to be careful as to how we collect the data and interpret the results.

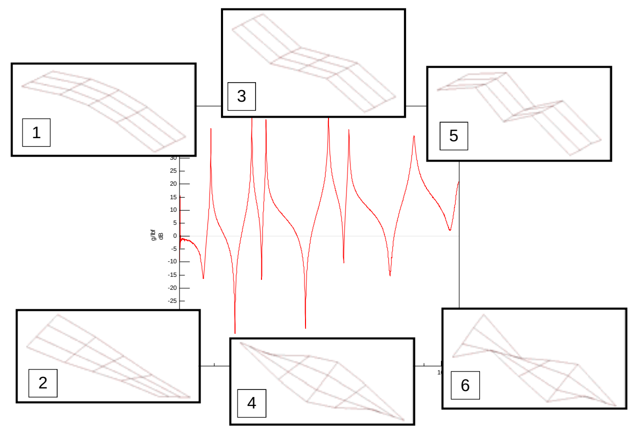

Figure 4 – Mode Shapes for Structure

Table 1 – MAC for Two Modal Tests Performed

Frequency	*179.3 Hz*	*413.5 Hz*	*495.1 Hz*	*853.7 Hz*	*970.6 Hz*	*1345.2 Hz*
179.3 Hz	*100*	*0.006*	*0.152*	*0.048*	*32.868*	*0.006*
413.5 Hz	*0.006*	*100*	*0.015*	*0.123*	*0.002*	*9.974*
495.1 Hz	*0.152*	*0.015*	*100*	*0.001*	*0.165*	*0.075*
853.6 Hz	*0.048*	*0.124*	*0.001*	*100*	*0*	*0.179*
970.6 Hz	*32.873*	*0.002*	*0.165*	*0*	*100*	*0*
1345.2 Hz	*0.006*	*9.975*	*0.075*	*0.179*	*0*	*100*

I hope this helps to explain that sometimes we have "suggested rules" but that sometimes we can still use information beyond the typical acceptable range of useful data. If you have any other questions about modal analysis, just ask me.

Copyright 2011
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Sometimes when I have a double impact, I switch hammer tips to eliminate it. Is that OK?
Let's take some measurements to see what impact this has.

So we have talked about double impacts before but this is a different scenario. On the surface it sounds like this might be a way to mitigate the double impact but there may be some ramifications as a result of that. So let's take some measurements on the same structure we discussed in the last article to see what impact this has (no pun intended).

Last time we were discussing the rolloff of the hammer and we showed that the rolloff itself didn't significantly degrade the resulting mode shapes of the system but that there was some degradation of the FRFs measured as expected.

Now during that original test we were fairly careful to avoid any double impacts (with the harder tip). But we have gone back to that same structure and acquired some additional measurements and made sure that some of the measurements were acquired with double impacts. And in fact we took another whole set of data and specifically made sure that every one of the FRFs acquired came from impact excitation where double impacts were applied.

For reference, the typical input force spectrum for a single impact and a double impact is shown in Figure 1.

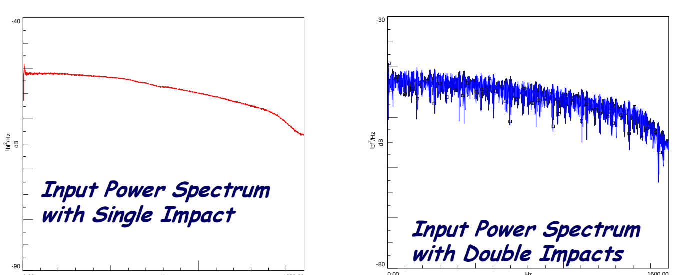

Figure 1 – Comparison of Single and Double Force Spectrum

While the double impact shows variation of the input force spectrum over the entire frequency band, it is important to note that there are no serious drops in the input spectrum which would be the major concern. And for reference, Figure 2 shows the typical mode shapes for the structure we are testing.

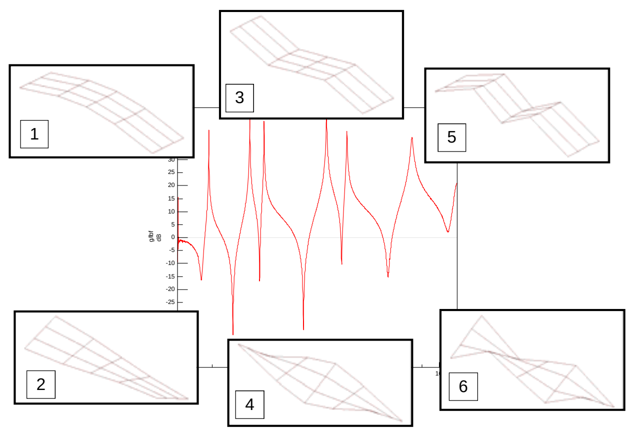

Figure 2 – Mode Shapes for Structure

Now what I am going to do is use the data set with the harder tip and no double impacts as the reference for the comparisons that we will consider here. And I am going to acquire some measurements in locations of the structure where double impacts could possibly occur and use the softer tip to acquire those measurements. (Just to make sure I document this properly, the outer 10 FRFs of the structure are measured with the harder tip and the inner 10 FRFs are measured with the softer tip.)

For comparison, two FRFs from each hammer tip are shown in Figure 3 and Figure 4 for the harder tip and the softer tip, respectively.

Copyright 2011
Pete Avitabile

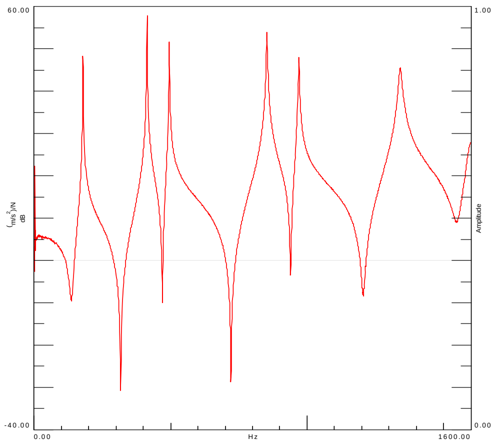

Figure 3 – Typical FRF for Harder Impact Tip

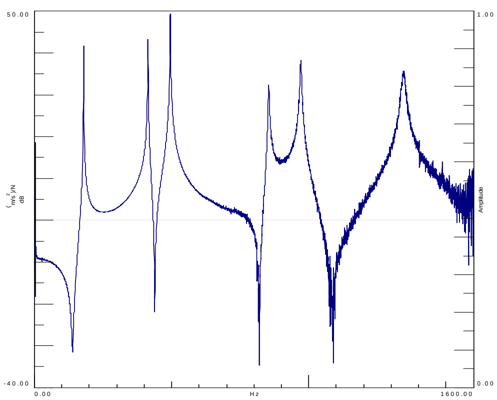

Figure 4 – Typical FRF for Softer Impact Tip

Now for the first comparison, the MAC was computed for the reference modal data set and the "hybrid" set of modal data where some of the measurements were made with the harder tip and some were made with the softer tip; the original idea was to minimize the double impact with the softer tip. The MAC is shown in Table 1 for this case.

Table 1 – MAC for Reference Test and Hybrid Data Set

Frequency	*179.270 Hz*	*413.356 Hz*	*495.121 Hz*	*852.661 Hz*	*970.418 Hz*	*1341.456 Hz*
179.304 Hz	*98.547*	*0.207*	*0.048*	*0.17*	*30.453*	*0.114*
413.501 Hz	*0.052*	*98.088*	*0.007*	*0.253*	*0.149*	*10.311*
495.105 Hz	*0.114*	*0.189*	*99.798*	*0.144*	*0.173*	*0.216*
853.646 Hz	*0.107*	*0.573*	*0.002*	*97.825*	*0.121*	*0.31*
970.634 Hz	*33.247*	*0.144*	*0.09*	*0.082*	*95.881*	*0.126*
1345.196 Hz	*0.122*	*9.725*	*0.07*	*0.431*	*0.132*	*97.921*

Notice that the MAC for the diagonal terms ranges from about 95 to 99 for the corresponding modes; the off-diagonal terms are not as critical to this evaluation because spatial aliasing is the main difficulty with such a limited set of data points.

But remember from the last article, when we compared all the harder tip modal data set with the softer tip modal test there was essentially no difference between the modes. So what has happened here?

Basically, as we switched the tip on the hammer we had an effective change in the input spectrum which essentially changed the calibration for the hammer. Because all the measurements were not collected with the same hammer tip, there is a bias on some of the measurements relative to the balance of the measurements. This means that we have created an imbalance in the scaling of the FRFs. So this directly implies that we really shouldn't switch the hammer tip in the middle of the test or else there can be a bias on the FRFs collected – unless if we calibrate to normalize that effect in the data acquired.

Now let's take this just one step further and use another set of data. While I am not an advocate of using double impact data, we have shown in the past that sometimes we might need to collect data with double impacts and maybe that data is not horrible to use – *as long as we use care to make sure that all the data seems reasonable with good coherence.* Now I am going to use the data set where all the FRFs were measured with some type of double impact but all FRFs were acquired with the same hard tip for all measurements.

Now another MAC was computed for the reference modal data and the modal data with some type of double impact at all measurement points. The MAC is shown in Table 2 for this case. Now notice that MAC for all the diagonal terms are all above 99. So this shows that the data was actually very good overall and the FRFs collected with double impacts are actually better than the data where we tried to minimize the double impact by using a softer tip at a subset of locations on the structure. I guess you would never expect that result but it makes sense if you consider that the double impact data was collected with a somewhat consistent input excitation whereas the "hybrid" data set was not.

Table 2 – MAC for Reference Test and Double Impact Test

Frequency	*179.454 Hz*	*414.166 Hz*	*495.463 Hz*	*855.208 Hz*	*972.122 Hz*	*1346.707 Hz*
179.304 Hz	*99.634*	*0.014*	*0.085*	*0.093*	*33.183*	*0.024*
413.501 Hz	*0.024*	*99.823*	*0.004*	*0.137*	*0*	*12.293*
495.105 Hz	*0.039*	*0.036*	*99.906*	*0.034*	*0.093*	*0.058*
853.646 Hz	*0.1*	*0.175*	*0*	*99.475*	*0.065*	*0.341*
970.634 Hz	*33.476*	*0.01*	*0.117*	*0.072*	*99.579*	*0.051*
1345.196 Hz	*0.018*	*11.365*	*0.06*	*0.216*	*0.009*	*99.292*

I hope this helps to illustrate that double impacts are maybe not as bad as you would have guessed. And the switching of the impact tip during the middle of the test, without accounting for the effective change in the input force spectrum, changes the calibration and needs to be considered. If you have any other questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

I took a measurement one day and the next day the measurement looked different. How can that be?
Let's take a look at this measurement and try to see what's happening.

So we have talked about a lot of different measurement problems in the past but this one is really one that needs some attention. Of course there are always differences between measurements from one day to the next or hour to hour or even week to week; these are typically small differences that we need to live with when we test structures – this is normal variability.

But the measurement you showed me is quite different than what we normally see as variation we would expect to see. In this particular measurement, there was very small variation in the lower frequency range and then there was a very significant and dramatic difference in the higher frequency range.

So let's step through this particular measurement that was made and discuss what happened. The main item that we will home in on is that the fixturing for the test may have played a very significant effect in the measurement for the system.

The measurement that was made was for a small wind turbine blade. The blade was to be tested in a "built-in" or "clamped" condition. The blade itself weighs less that two pounds and is attached to an 800 pound optical table. Now the optical table is certainly large enough to be able to adequately simulate a built-in condition for the blade. In fact, an analytical model was available for the structure and the anticipated built in modes were available from the model. While the analytical model is never perfect and has approximation that are made when the model is developed, a model for this type of simple structure is expected to be reasonably accurate given the simple geometry for the configuration.

So the first frequency response function was made with the blade attached to the optical table. Normal impact measurement methodologies were employed for the impact and response measurements with an FFT analyzer. The blade was attached to the optical table with a solid mounting block and some threaded rods to attach the block (as a spacer) between the turbine blade and the optical table; this was necessary due to the curvature of the blade geometry to allow for clearance between the cantilevered blade and optical table.

Overall the measurement looked acceptable and the coherence for the measurement was also acceptable. The frequency response function is shown in Figure 1 with a picture of the test set up. The low frequency portion of the measurement was not of interest and is hidden behind the picture of the test set up.

So the first measurement taken appeared to be acceptable. Now this measurement was taken on a Friday afternoon followed by additional testing to be performed the following week. As good practice on Monday morning, the measurement was repeated prior to the balance of testing to be performed. This second measurement also looked acceptable overall. The frequency response is shown in Fugure 2 with a picture of the test set up. Again the low frequency portion of the measurement was not of interest and is hidden behind the picture of the test set up.

But the measurement taken on Monday morning did not look the same as the measurement taken on Friday which had several people scratching their heads. Figure 3 shows the comparison of two measurements overlaid.

Fortunately someone recognized these dramatic differences and stopped testing to determine what could have possibly caused such a dramatic difference in the measurement. Figure 3 clearly shows that the lower frequency modes are essentially the same with normal variation which is to be expected. But the higher frequency range had a dramatic difference in the measurement, almost as if it was from a completely different structure.

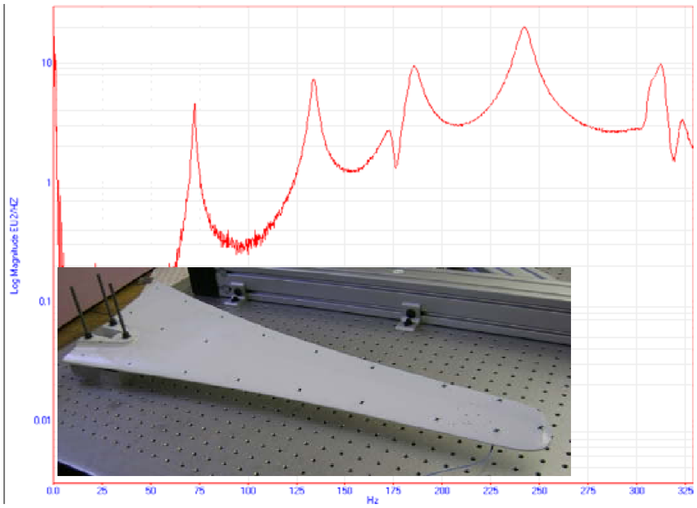

Figure 1 – FRF and Test Set Up - friday

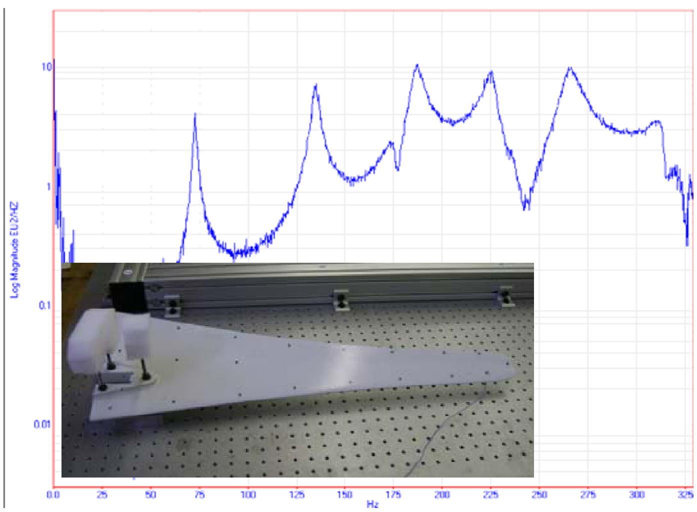

Figure 2 – FRF and Test Set Up - Monday

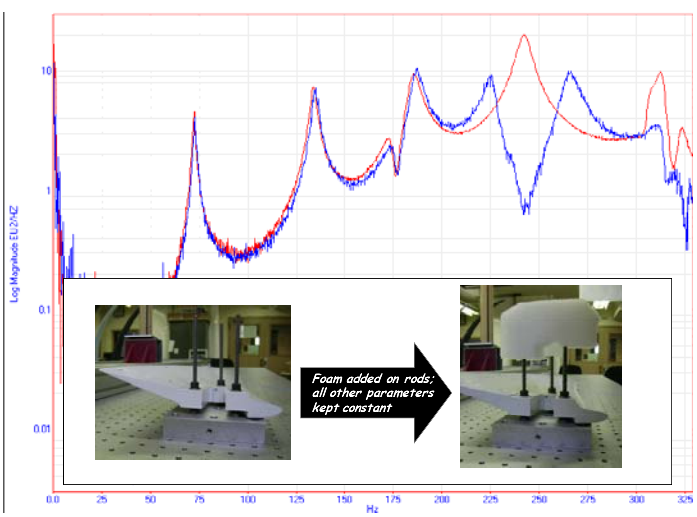

Figure 3 – Comparison of Friday and Monday Measurements

Obviously the first thoughts are directed to the fixturing that attached the blade to the optical table. But inspection of the joint and tightening and reassembly of the joint did not make any appreciable difference.

So when I first saw this measurement, I immediately suspected that there was likely something different in the two tests that might be attributed to some sort of tuned absorber effect. The reason why I jumped to this conclusion is because there is a general shifting of frequencies in the lower frequency range that could result from a tuned absorber effect.

So after a little detective work and asking some general questions as to what could be different between the two tests, some thoughts emerged as to what could be the problem. After a few more questions and some close interrogation of the pictures of the test set up, there is one thing that caught my eye.

Well as it turns out, the lab manager was very concerned about safety and the three threaded rods that protrude from the blade and block attaching it to the optical table could cause some injury if someone wasn't careful around the test set up. So as a precaution, the lab manager had some soft foam pushed onto the long threaded rod. (Now this foam is basically the soft packing foam that you find in packaging of electronics and is extremely light and really has no structural effect whatsoever.) BUT…

That foam does have a very, very small amount of mass and at the end of the long threaded rod has the effect of changing the cantilevered mode of the threaded rod. And oddly enough, the effect of shifting that cantilevered rod frequency just happened to coincide with one of the natural frequencies of the blade. This then caused the two modes to split in a very traditional tuned absorber effect. And this effect was very dramatic for sure.

And if you don't believe that this could possibly happen and that this was in fact the cause of the problem, then please rest assured that the structure was tested about five more times with and without the small piece of packing foam attached to the threaded rod and both sets of measurements from Friday and Monday were replicated each time the test was conducted.

So the bottom line here is that fixturing in any dynamic test can have an effect on the measured data. In this case the effect observed was very dramatic for the fifth mode of the blade. Care must be exercise in all fixturing for all dynamic tests performed. You never know what can happen! If you have any other questions about modal analysis, just ask me.

Illustration by Mike Avitabile

We compare tests to fully built-in models all the time. Can you really simulate built-in down in the lab? Let's take a look at this to understand this problem.

So this has been a constant item for discussion for as long as I can remember – and while I might forget certain things as I get older, this topic is one I remember very well. It seems to constantly pop up as an item for discussion all the time. Of course, this happens because the analytical world is quite different than the experimental world.

In analytical modeling, we can always very simply identify a boundary condition to apply to our model. We can make it totally free-free if we want. But of course the real world can be much different than the analytical world (which is filled with "assumptions" that may not be possible in the real world).

But at least with a free-free boundary condition we can often do a fairly good job of approximating that condition. In fact, we often like to test this way because then there is very little interaction with the test fixturing and related set ups for conditions that are other than free-free. (Remember in the last article, how a very seemingly insignificant change to the test set up ended up having a very significant change to the test results.)

The problem is that many times we would like to validate our analytical models with a constrained boundary condition at the attachment points to our component. From an analytical standpoint, this is very desireable. But from a practical testing standpoint, this can be very messy from a variety of different perspectives. There are all kinds of issues related to mounting surface flatness, bolting preloads, etc. that are of concern.

But one item that is always one that can be misunderstood is, eaxtly how stiff is stiff and how massive is massive in regards to creating that so called built-in condition. Many people will try to design "an infinitely stiff" support frame or test fixture. But we all know that any structure will have resonant frequencies, it is just a matter of where they occur and what effect they may have on the ability to actually create a built-in condition. The test fixture may be adequate for the first few lower order flexible modes but eventually there are fixture resonances that may interact with the test article.

Another way to simulate a built-in condition is to provide a large seismic mass. This generally tends to be a better mechanism to achieve a built-in boundary condition but often times people don't realize exactly how much mass is actually required to achieve this condition. Often times you will hear people state that the seismic mass needs to be 10 times larger than the mass of the test object. For some reason, people think that the 10:1 ratio is the answer for all problems. But what we forget is that these "rules of thumb" evolved back in the early days when all we had was a slide rule for calculations. And with that level of accuracy maybe that 10:1 rule was a good guesstimate. But now with all the sophisticated models we can evaluate today, we really should rethink that 10:1 rule.

In order to illustrate this, a simple example for a long beam like structure will be used to explain this. We have recently performed some free-free testing for flapwise modes for a 9m turbine blade and this is an excellent article to discuss in regards to a mass loaded interface to simulate a built-in condition.

The 9m blade was tested free-free and a very coarse beam finite element model was utilized to model the turbine blade for flapwise motion alone for the first few modes. The intent was to use that free-free model and then apply a "perfect analytical built-in boundary condition" and compare it to a variety of different modes to study the effect of the amount of mass needed to actually anchor the blade to ground.

The free-free test and analytical modes are shown in Table 1 for reference. The analytical model and test performed are not described here for brevity.

Table 1 - Comparison of Model and Test Results Free-Free for a 9 m Wind Turbine Blade

Mode	Model (Hz)	Test (Hz)	MAC
1	7.84	7.76	99.85
2	18.5	21.26	98.28
3	34.52	31.34	98.85

These results are considered reasonable considering the coarseness of the rough beam finite element model.

Now the analytical model can be used to identify the "perfect" built-in boundary condition and will be used as a reference. Another analytical model will also be used to compare the effects of adding a very large "seismic mass" to the root of the blade model. This model, with different mass conditions will be compared to the "perfect" reference model. The frequencies and shapes will be compared to show the effects of the amount of mass that is needed to achieve this constrained condition. And it is very important to note that the mass is not just a lumped mass; the mass has rotational effects and it is these rotational mass effects that are the most important ones for the development of a built-in simulation for a long overhung structure such as the turbine blade. For reference the turbine blade weighs on the order of 400 lbs. One approximation of a seismic mass used a 66" x 72" and 24" thick steel plate that weighed approximately 22,000 lbs or roughly 55 times heavier than the turbine blade.

Several models were developed with various ratios of the lumped mass, designated as M, and the rotary mass, designated as MR and were compared to the "perfect" reference model, for both frequency and shape of the resulting model. Table 2 summarizes the results considering just the first mode of the turbine blade and Figure 1 shows the shape comparison for the different cases shown in Table 2.

First notice that the model with just M and MR as the seismic mass approximation does not replicate the frequency very well; realize that the anchor is over 50 times the weight of the blade. Also notice that the curvature of the mode shape does not match the "perfect" reference model well. Now doubling the lumped mass, 2*M, with the same rotary mass, MR, shows some improvement but still has differences. And if the inertia is doubled and lumped mass kept the same, the results are approximately the same. The curvatures are improved but there are still differences. Notice that the rotary mass of the seismic anchor is the most critical item to cause the frequencies and shapes to match much better. If the rotary inertia is increased by an order of magnitude, then the frequency better correlates and the curvature of the shape also starts to match very well.

So the bottom line here is that the seismic mass needs to be very large in order to approximate a built-in condition and that the rotary mass effect is much more important due to the large overhung effects of the wind turbine blade. Obviously, a different structure with a lower center of gravity will have different results but this model certainly shows the importance of identifying the proper characteristics necessary for the seismic anchor for overhung structures.

Table 2 - Comparison of Different Seismic Anchor Approximations and "Perfect" Built-in Condition for a 9 m Wind Turbine Blade

Properties of Anchor		Freq. of Mode 1 (Hz)	
Mass	Inertia	With Anchor	Truly built in
M	MR	5.29	4.36
2*M	MR	4.87	4.36
M	2*MR	4.87	4.36
M	10*MR	4.47	4.36

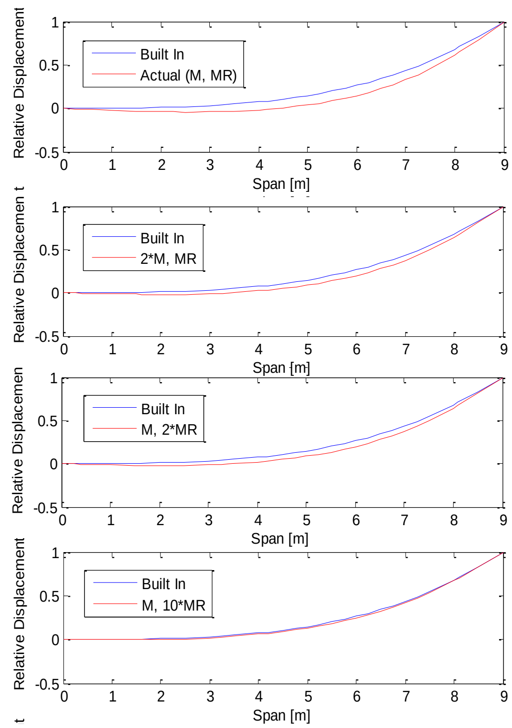

Figure 1 – Comparison of Mode Shape Curvature

But at the end of the day you have to decide how much variation from the real built-in condition is acceptable and how much deviation in the actual mode shape is acceptable for the intent of the test that is being performed. The problem is that often times people haven't considered these issues in any depth and therefore do not have clear statements as to what deviation may be acceptable. If you have any other questions about modal analysis, just ask me.

Copyright 2012
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

If you don't excite a structure at its natural frequency, then how can you know what it is?
This is an important item to discuss.

Well, this is an area where I find people often get confused. Many times I hear people say that they have to tweak and tune the excitation frequency so that they get the excitation right at the natural frequency otherwise the frequency will not be identified properly. I also hear people say that the excitation method must have broadband energy at all frequencies otherwise the system will not be excited properly.

There is a misconception that the frequency of excitation must be *exactly* at the natural frequency otherwise the results are not valid. Well this is not really a problem with the way we do frequency response testing and how we extract parameters from measured data to estimate the frequency and damping for a system.

So let's discuss some of this and help you to understand why we really don't need to excite a structure *exactly* at its natural frequency when we run a test - but we do have to make sure that the data we collect is good data because there is no substitute for good data.

Now let's take a step back to something a little simpler and more commonly understood. Let's look at a very simple straight line fit of some measured data. We are going to perform a least squares error minimization for the data presented in Figure 1. Now we know we can fit any line to the data but for this set of data it seems that a first order fit makes the most sense. Of course the model we are going to use is

$$y = mx + b$$

and there are two parameters that define the line, namely the slope and y intercept.

The data and fit of the data is shown in Figure 1. Now let's look very closely at the data. We know that we can compute the slope, but did we ever really measure the slope? Not really – we measured data and then fit a mathematical function to that data which is then used to obtain the slope – but technically speaking we did not measure the slope. And let's also look at the y intercept. If you look closely you will see that we never measured data that directly obtained the y-intercept. But certainly we would say that we could obtain the y-intercept but we really never actually measured it, did we? And let's take this just one step further to see if we could obtain the value of the y function for a given value of x, let's say 0.707. But we never actually measured the function for that particular value of x – but we would say that we know what it would be from the function we fit. I think we are all comfortable we these statements and the values that we obtain – right?

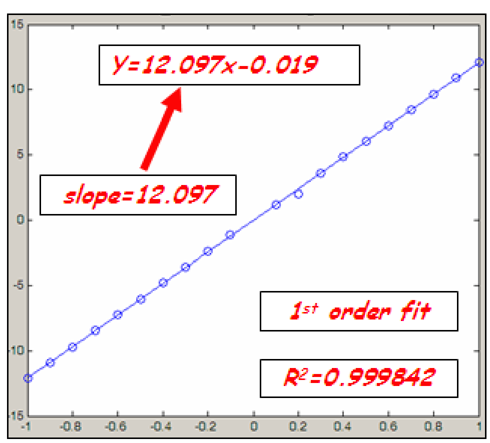

Figure 1: Example of Simple Straight Line Fit

Copyright 2012
Pete Avitabile

So now let's apply this logic to the mathematical function that is used to represent a single degree of freedom frequency response function. We can write the system transfer function in partial fraction form for a single degree of freedom system as

$$h(s) = \frac{a_1}{(s - p_1)} + \frac{a_1^*}{(s - p_1^*)}$$

and we can also write the frequency response equation as

$$h(j\omega) = h(s)\Big|_{s=j\omega} = \frac{a_1}{(j\omega - p_1)} + \frac{a_1^*}{(j\omega - p_1^*)}$$

First we need to realize that this function is written as a function of frequency and it contains two constants which are the pole, p, and the residue, r; these are the two parameters that we need to extract (just like we did for the slope and y-intercept for the straight line). Now we can evaluate this function at a series of frequencies spaced delta f apart. This is shown in Figure 2. These data points are the set of data we collect when a frequency response measurement is made. We can fit a line to the data where the line is the complex frequency response function and extract the parameters of interest.

I can make the same argument that was made for the straight line here for the frequency resonse function. We are going to fit a line to the measured data and extract several key parameters that define that line. These are namely, the pole and the residue. And just like we argued for the straight line I do not necessarily need a data point *exactly* at the natural frequency to obtain an estimate of the pole or residue. As long as I have good data that represents the function, I can fit this simple single degree of freedom model to the measured data to extract the parameters of interest which are the poles and residues.

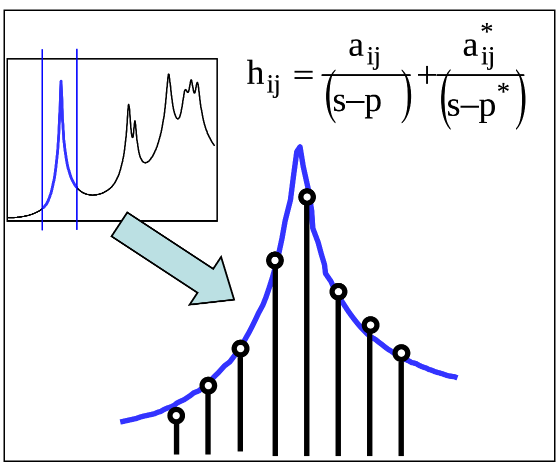

Figure 2 – Conceptual SDOF Curvefit

Now of course I can extend this from a single degree of freedom system to a multiple degree of freedom system and the problem just gets mathematically more complicated – but it is the same process overall. A multiple degree of freedom frequency response measurement is shown in Figure 3.

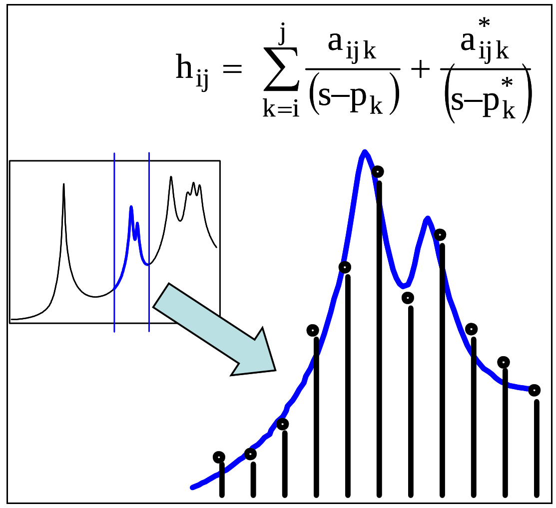

Figure 3 – Conceptual MDOF Curvefit

Data is obtained at selected frequency spacing and there is no need to make sure that the frequency data points coincide *exactly* with the precise natural frequencies for the modes obtained from the parameter estimation process. But surely I have to remind everyone that the data that is collected must be obtained using the best of measurement methodologies to assure that there is little variance between the data collected and the line that is used to describe the data and for the extraction of the poles and residues.

So I hope that this little explanation helps to clarify the fact that the measured data does not necessarily need to lie *exactly* on the precise natural frequency for the system. The process of modal parameter estimation (which is really nothing more than a very elaborate least squares error minimization process) where parameters, namely the poles and residues, are extracted. If you have any other questions about modal analysis, just ask me.

Copyright 2012
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

How free does a test need to be? Does it really matter that much?
This is an important item to discuss.

Alright – now here is an item that I think everyone gets confused about. It really stems from the fact that we all are familiar with rigid body dynamics and the concepts surrounding that. But we don't let go of those concepts easily.

What I mean is that rigid body dynamics is a good approximation when the body is described by a center of mass concept and all the points on the structure can be described in terms of that one point. What we then have is a point in space that is the only point we need to describe the entire motion of the structure associated with that one point.

And that is a pretty powerful statement. It means that every point on the geometry of the structure can be defined completely by that one point on the structure. Therefore, if this was the case then we would say that we have a rigid body.

So now what do we mean when we say we have a free-free system and we have rigid body modes that describe that system. That implies that we have no constraint whatsoever to ground and that the structure is essentially floating freely in space. If that is the case, then we would have rigid body modes describing the six independent ways that the body can move in space. There would be three separate translation motions in the three principle directions as well as three separate rotations in the three separate directions. Of course we have to realize that the six independent motions could possibly be comprised of linear combinations of each other as another possibility. Just because we tend to think in x,y,z directions doesn't mean that the rigid body motion needs to be isolated to those three directions – any linear combinations are also valid.

OK – so now we have this rigid body mode concept down. Now let's talk about a simple beam that we might possibly model with a finite element model. Let's assume to start that the beam is a uniform cross section and uniform weight distribution so there is nothing fancy about this beam. To further simplify the discussion we will only consider planar motion but there is no reason we couldn't extend it to six degrees of freedom to be general.

So let's first describe the first few modes of this planar system. Figure 1 shows the first four modes of the planar beam system. Notice that the first two modes are the rigid body modes and that the next two modes are the flexible modes of the system.

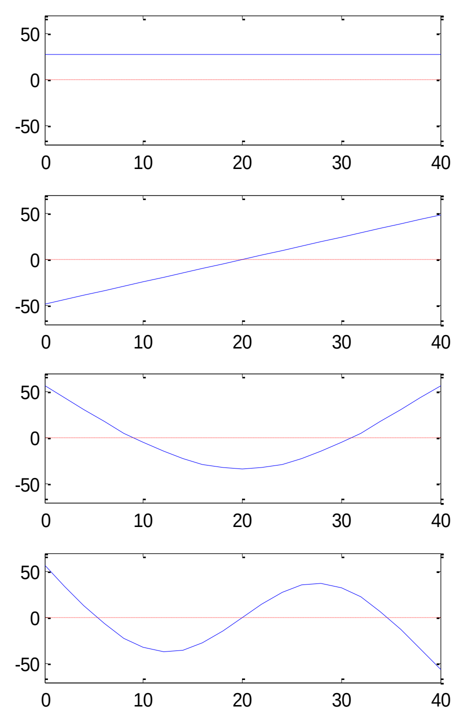

Figure 1 –Modes of Free-Free Beam

Copyright 2012
Pete Avitabile

Notice that the first rigid body mode is a bounce mode with up and down motion and that the second rigid body mode is a rocking mode about the geometric center of the beam. This is what is expected for the free-free modes of a fully unconstrained beam structure.

Now let's consider that the beam really can't float in space unconstrained when we test it in the lab. And let's consider a range of spring stiffnesses to apply to the two ends of te beam. And let's further let the stiffnesses range from close to zero all the way up to a very high stiffness approaching a pinned condition or perfectly constrained. This is schematically shown in Figure 2.

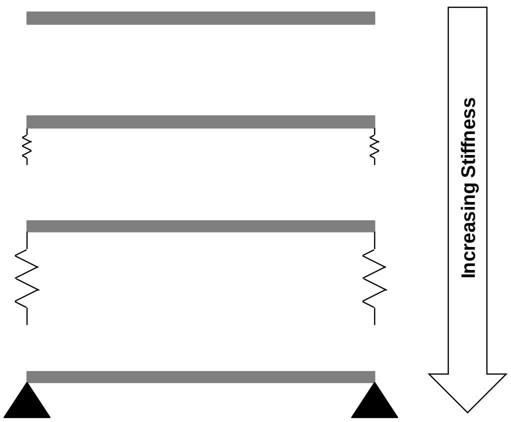

Figure 2 – Elastic Support for Beam

Now to make this simple, we are just going to look at how the first mode changes as we add increasing spring stiffness to the ends of the beam. We are going to look at what happens to the mode shape as the stiffness increases. This is shown in Figure 3 with the mode shape shown from top to bottom with increasing stiffness.

The first shape plotted is the free-free beam first mode shape. So as we increase the stiffness at the end of the beam, the natural frequency will shift upwards because there was an increase in stiffness as expected. So if we added just a little bit of stiffness the mode shape may not change appreciably. And we will notice that in the second plot from the top that the mode shape is still very similar to a rigid body mode but that there is slight amount of curvature in the beam. As we increase the stiffness we see that in the third plot that the shape doesn't really look like a perfect rigid body mode and that the shape is starting to take on more of a curvature like the first flexible mode of the system. By the time we increase the stiffness even more, the fourth and fifth plots don't really resemble a rigid body mode any longer and basically the mode shape really resembles a flexible mode with just a tiny bit of rigid body motion.

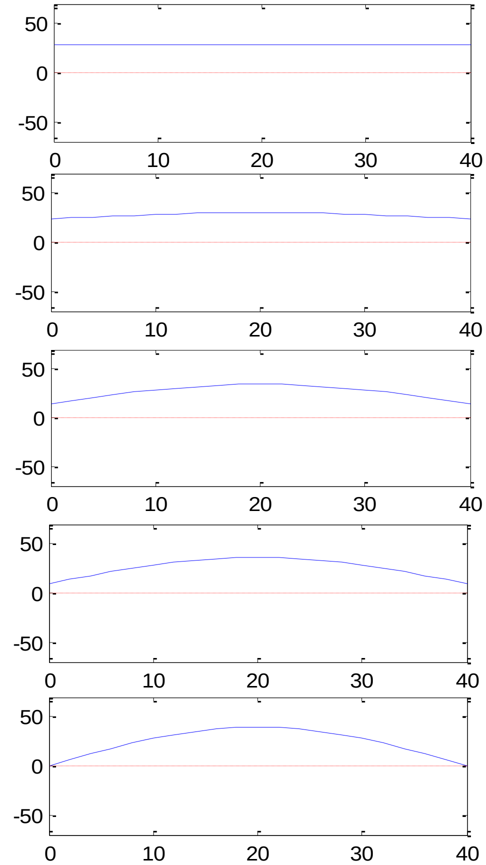

Figure 3 – Progression of Mode 1 Shape

So now the story is basically told. The rigid body mode is only truly a rigid body mode when it is completely free-free. Once any amount of stiffness is added to the ends of the beam, then the mode starts to change from a rigid body type mode to a flexible type mode and the proportions of rigid body and flexible mode is heavily dependent on the amount of stiffness added as well as the stiffness of the structure itself.

This means that when we measure any structure in the lab in the so called free-free state, the actual rigid body mode obtained will always have some of the flexible modes included and is really not a perfectly rigid body mode. Depending on how the test is set up and and how stiff the free-free suspension is will have a direct effect on just how rigid those rigid body modes are. I hope this simple explanation clears up any misconceptions that you may have had. If you have any more questions on modal analysis, just ask me.

Copyright 2012
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD *by Pete Avitabile*

Illustration by Mike Avitabile

So what really is a drive point FRF? Do you have to impact exactly at the same point?
Let's take a look at this.

Drive point measurements have always raised questions in regards to experimental modal tests. There are several things that need to be considered when conducting a test especially for this measurement. So it is very important to discuss it.

The drive point measurement is a very important measurement to be made as part of an experimental modal test. The drive point frequency response function is a measurement where both the input force and response are measured on a structure at the same point and in the same direction. Now a few things need to be considered when we discuss this type of measurement.

For sure, it is very difficult to actually hit the structure at the same location where you are simultaneously measuring the response, so there are some practical implications that need to be considered. I have seen some cases where accelerometer casings appear to have been subjected to physical impacts in order to try to take this drive point frequency response function. Now this is definitely not recommended as the way to take this measurement. So we need to think about how to make the measurement and how to consider the implications of the practicality of actually taking this measurement.

So obviously we need to try to achieve the desired result as closely possible without actually impacting right on the accelerometer itself. So one way to achieve that would be to measure on the opposite side of the structure. If the cross section is very stiff or a solid cross section then this would appear to be a possible way to achieve that result. The only difference would be that the phase of the measurement would need to be considered so that if the positive sensing direction of the accelerometer was 180 degrees opposite to the desired measurement then the phase would need to be corrected. And in just about every modal software package available, the software allows for the phase to be included with the specification of the measurement being in either the "plus" direction or in the "minus" direction. So that is not really a problem (but we will discuss one difficulty in a few moments).

The other way to achieve the drive point measurement is to impact alongside the accelerometer when making the measurement. Now this is not truly a drive point measurement but if the structure is very large, then this is not a problem. So if I were to take a measurement on a big wind turbine blade, then the effects of this small difference in the location of the impact would be essentially insignificant. But if I were to take the same drive point measurement on a much smaller structure such as a disk drive or jet engine turbine blade then the size of the structure relative to the small difference in the actual geometric location of the accelerometer and the actual impact location may have a fairly significant change in the drive point measurement in that case.

The effect is going to be very dependent on the change in the value of the mode shape over that very small distance. If the mode shape doesn't change very much then the difference in the actual drive point measurement and the acquired drive point measurement may be essentially insignificant. But as the structure starts to get smaller or higher modes are considered, then the effects of the actual change in the mode shape can have a much bigger impact (no pun intended). This can really all be related back to the equation describing the frequency response function written in terms of mode shapes for a single mode approximation can be given as

$$h(j\omega) = h(s)\Big|_{s=j\omega} = \frac{(q\,u_i\,u_j)}{(j\omega - p_1)} + \frac{(q\,u_i\,u_j)^*}{(j\omega - p_1^*)}$$

Obviously if the value of the mode shape between point "i" and "j" is extremely small then the change in the actual measured frequency response function and the drive point measurement will be very small. So it is all dependent on size and the change in the mode shape over the very small distance of the accelerometer and impact location.

But let's consider one additional case that might be a more common problem that needs to be addressed. Many times a measurement will be made and the accelerometer is located on the opposite side of the structure for convenience. If the structure is a solid cross section or it is very stiff then it would seem reasonable to make that measurement in that manner. Or it might not be possible due to space constraints. In any event, a simple tubular beam cross section will be used to show some additional concerns that need to be considered. The beam cross section is shown in Figure 1 with two small teardrop accelerometers mounted on the structure along with a schematic to the right with a red accelerometer shown as the true drive point measurement and the blue accelerometer shown as the approximation of the drive point measurement that might typically be acquired. Obviously the measurement here can be made because the FRF drive point measurement is made at the end of the beam where access is available; but if this measurement was needed at an interior location then this measurement of the true drive point measurement could not be possible. (For reference, this an aluminum beam approximately 60 inches long with a 1 inch by 2 inch cross section with a 3/16 inch wall thickness.)

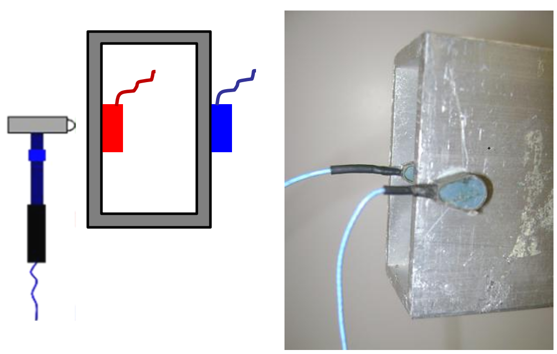

Figure 1 –Schematic of the Beam Measurement

Now an impact measurement was taken over a 4000 Hz range and also zoomed in over a 1100 Hz range to more clearly see the difference in the frequency response function. Figure 2 shows the imaginary part of the frequency response function and the two traces (red for the true drive point frequency response and the blue for the approximate frequency response function) are overlaid for comparison. Essentially there is no difference in the imaginary part of the function. Remember that the imaginary part of the frequency response will be a peak when the real part is a zero for a proportionately damped system with well spaced modes. Figure 2 seems to indicate that there is essentially no difference at all and would lead you to believe that there is no error in this measurement.

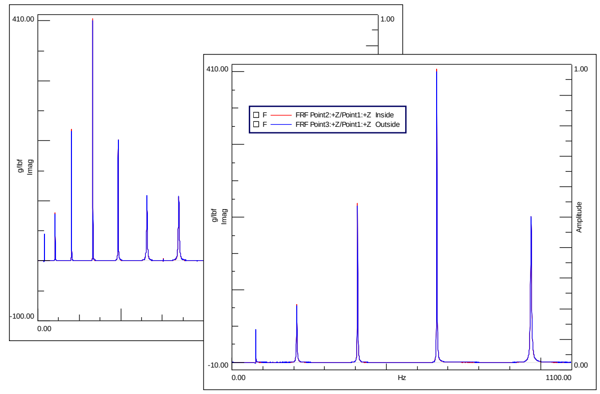

Figure 2 – Imaginary Part of the FRF

However, if we look at the magnitude of the frequency response function we see something that indicates a different story. Notice that the anti-resonances do not line up between the two measurements. This is directly related to a phase difference between the two measurements. So while the magnitudes line up properly, the phase between the two measurements shows a significant difference. Yet upon first looking at this simple beam section, the lower order modes would be expected to be relatively unaffected by the difference between measuring the exact drive point and the approximation of the drive point measurement, especially for the lower order modes. But it is clearly seen that there is a difference. (And just to be sure there was no instrumentation issues the measurement was repeated with both accelerometers mounted on top of each other and the measurement was essentially identical.)

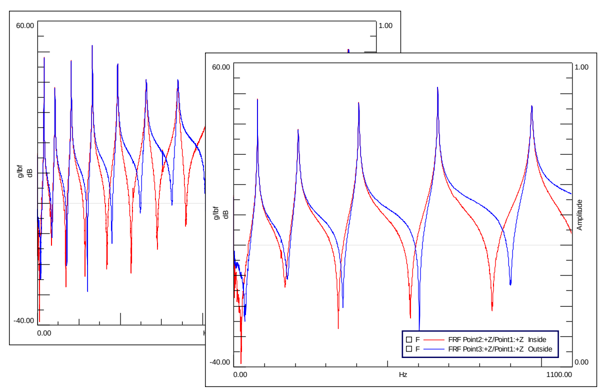

Figure 3 – Magnitude Part of the FRF

While the amplitudes would likely give a good representation of the mode shape, the more important item to observe is that if these FRFs were used for any frequency based substructuring type applications, then that phase/anti-resonance issue would cause difficulties in numerically processing any inconsistent data that might be collected at different measurement points. I hope you have a better appreciation of drive point measurements now. If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD by Pete Avitabile

Illustration by Mike Avitabile

We updated a composite plate model but the properties seem to not be physical.
Can the model be used for response studies? Let's take a look at this to understand this problem.

So this is an area where people often get confused. We make analytical models as an approximation of the characteristics of a system. Many times we use modeling approximations or equivalencies that help to obtain the right overall characteristic for the model. But it doesn't necessarily mean that the model portrays the actual physical property as we may expect to see it as if we were to get it from a material properties table.

Actually I have a very good case that may be a good one to present related to the modeling of a composite plate that was modeled with some radically different modeling strategies. I really don't need to go into all the details of the modeling philosophy and strategy deployed but I will only concentrate on the confusion of the material properties that resulted from the model updating performed.

A finite element model of a wind turbine composite section was developed. The physical structure had a balsa core and a 5 ply resin fiber layer (0-90° warp-weft architecture) on each side of the balsa. The finite element model used a unique modeling strategy to capture the resin/fiber composite with a plate and beam formulation to capture the shear and bending characteristics of the fiber embedded in the resin which is modeled as a plate. This modeling scenario had been used in the past but a prototype was also fabricated to perform some modal tests to validate the model. The testing was performed on a 3 ft x 3ft plate in a free-free condition as well as in 4 separate configurations with each side of the plate clamped to a 500 lb block to form a constrained end condition. The material properties were provided from the materials group and were identified as being accurate for the identification of its characteristics. The finite element model was developed with these properties identified as supplied by the materials group. A free-free modal test was performed and used to study the model adequacy.

The free-frre correlation of the first dozen modes produced very good MAC values but the frequency had a very consistent shift in frequency for all the modes of the system. A model updating study produced a very significant change in the balsa material properties to cause the frequency difference to be minimized; the basic premise was that the resin and fiber material properties that were provided were correct. But the balsa which is geometrically located at the neutral axis of the plate needed to have a tremendous change in stiffness in order to accomplish this shift in frequency; the updated balsa properties for Young's Moduls essentially needed to be that of steel. While this may seem unrealistic from a practical standpoint, the reality is that the finite element approximation made with the composite materials used required this in order to achieve the proper stiffness to represent the modal characteristics properly. The results of the original free-free correlation and the updated free-free correlation is shown in Table 1.

In order to confirm that the model was a reasonable approximation of the system even with the unrealistic balsa properties identified, the plate was tested in 4 perturbed conditions where it was clamped on each of the 4 sides of the plate. The finite element model *with the updated properties* was correlated to each of these configurations and produced comparable results to the free-free update models; only limited results are shown to keep the article brief. Table 2 shows the results of the correlation of the finite element model *with the updated properties* for the stiffer and softer direction of the composite plate clamped to a 500 lb anchor in the lab; the model was made with the anchor included to best represent the clamped arrangement. Clearly, the update parameters are suitable to predict a significant change in the boundary conditions made to the free-free composite plate structure.

Table 1: Correlation of Free-Free Composite Plate Before Model Updating (left) and Correlation After Model Updating (right)

#	FEA Hz	EMA Hz	MAC
1	22.13	49.90	99.9
2	48.82	84.34	99.3
3	65.20	129.94	99.7
4	101.74	138.49	98.7
5	109.56	166.55	99.8
6	128.46	230.26	99.6
7	135.95	252.14	99.3
8	142.08	267.53	99.3
9	194.58	386.93	96.6
10	204.93	359.29	97.6

#	FEA Hz	EMA Hz	MAC
1	49.64	49.90	100
2	87.19	84.34	99.8
3	131.23	129.94	99.9
4	134.28	138.49	99.6
5	163.65	166.55	99.9
6	229.83	230.26	99.8
7	249.63	252.11	99.8
8	266.60	267.53	99.8
9	355.74	359.29	99.5
10	381.52	382.79	97.1

Table 2: Correlation of the Update Finite Element Model to Two Different Tests with Significantly Perturbed Boundary Conditions with the plate Clamped on One Edge

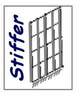

#	FEA Hz	EMA Hz	MAC
1	20.06	19.19	99.1
2	32.86	31.98	99.1
3	99.69	98.19	98.9
4	124.61	127.25	97.8
5	133.35	135.64	98.2
6	199.99	204.16	99.4
7	241.21	241.29	99.2
8	281.52	298.25	97.5
9	307.01	317.41	94.4
10	355.27	354.81	98.1

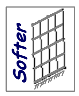

#	FEA Hz	EMA Hz	MAC
1	13.93	12.68	99.2
2	30.20	29.54	99.6
3	82.12	84.58	94.7
4	108.58	109.27	99.6
5	143.58	147.37	99.8
6	203.78	208.57	99.0
7	224.81	229.96	97.6
8	236.02	240.43	99.1
9	329.40	338.66	99.0
10	360.70	363.54	98.0

So many people might ask, how can you have such different material properties for the balsa approaching those of steel and expect the model to be reasonable. Well that is an excellent question and is one that needs to be discussed. I have a great example that I think you will very quickly accept.

We know that an I-beam gets its significant Area Moment of Inertia from the flanges being offset from the neutral axis. But I could reduce that down to an equivalent rectangular cross section to give me the same effective stiffness as seen in the upper portion of Figure 1. I have certainly captured the stiffness properly to get the right deflection of the system. But if you looked at the cross section you would say that it doesn't really look like an I beam.

Now in the composite plate model where we changed the balsa property, the effective composite materials we used did not really represent the true stiffness of the material. Because the composite fibers are outboard of the neutral axis of the balsa, their effect in defining the stiffness of the system is critical. In the first model we used "whimpy" material properties so to speak. So the only way that the model could reflect the difference in stiffness is by adjusting the Young's Modulus of the only element left in the model that could change – that is the balsa. So the balsa had to be very stiff – almost on steroids so to speak. While we may not believe that modulus from a physical standpoint, as far as the model was concerned, the balsa needed to be that stiff. If you look at the overall EI of the balsa and resin and fibers *as a complete unit*, then the overall stiffness is represented correctly as well as the mass distributed correctly and we predicted well over a dozen modes correctly – and in 2 different perturbed boundary conditions as well.

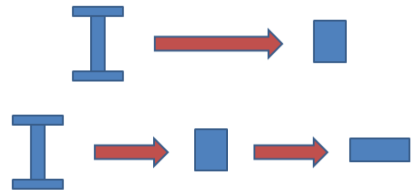

Figure 1: Schematic – Equivalent Section Properties

And let me take it one step further and look at the lower sketch in Figure 1. Let's say for some reason we couldn't let the rectangular approximation be as high as it is and we needed to make it half as thick. In order to do that I would need to change the stiffness somehow because the I about the weaker axis would not be as stiff as the thicker section. Because the I has a t^3 term, we would have to increase the effective stiffness of that thinner section by adjusting the EI term to account for the change. So the E would need to be adjusted by 2^3 to make the overall cross section have the same effective stiffness. But the real E of the material would not be that value –the model would need to adjust the material to compensate for the change in physical dimension.

So the bottom line here is that we take everything into consideration in terms of the overall mass and stiffness distributions to cause the system to have the right overall effective representations so that the ***response*** is measured properly and the structure has the right overall weight as well as the right stiffness such that if you put a static force on it you get the right displacement.

And as a side note, eventually the material properties were re-evaluated with updated material testing methodologies and correlation to the finite element model was significantly improved with very acceptable frequency and shape correlations. Now the material properties of all materials are more in line with what we may have expected. But the bottom line is that any of the updated models could have been used for a proper estimation of system characteristics when only considering response of the system. If you have any other questions about modal analysis, just ask me.

Copyright 2012
Pete Avitabile

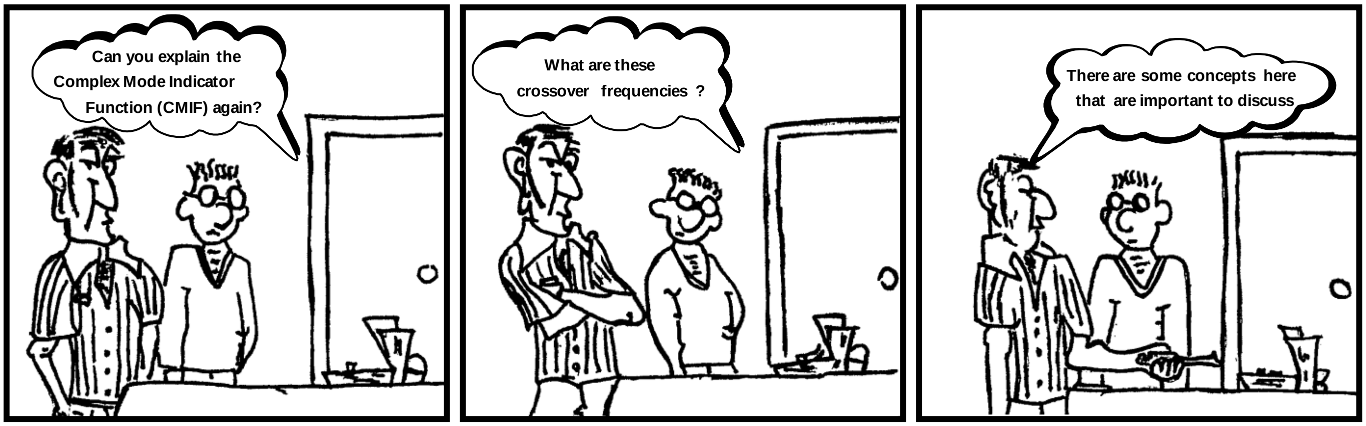

Illustration by Mike Avitabile

Can you explain the Complex Mode Indicator Function (CMIF) again? What are these crossover frequencies? There are some concepts here that are important to discuss.

Now I know I have heard many people comment on this as a point of confusion. But it really is not that hard to explain and I have examples to put it in perspective.

First let's get some of the messy math out of the way but that is needed to start the conversation. Basically, the Complex Mode Indicator Function (CMIF) uses the frequency response matrix and performs a singular value decomposition (SVD) to identify how many "significant" eigenvalues exist in each individual spectral line of the frequency response function (FRF) matrix.

WOW – now that is surely a mouthful to say the least. Let's present the equation and then try to pull it apart and make some sense of it all. If we collected a set of FRF data with multiple references, then I would have a matrix of FRFs. And I could write that equation using SVD as

$$[H]=[U][S][V]^T$$

And I could also write it down in expanded form to see some of the important pieces as

$$[H]=[\{u_1\}\quad\{u_2\}\quad\{u_3\}\quad\cdots]\begin{bmatrix} s_1 & & & \\ & s_2 & & \\ & & s_3 & \\ & & & \ddots \end{bmatrix}\begin{bmatrix} \{v_1\}^T \\ \{v_2\}^T \\ \{v_3\}^T \\ \vdots \end{bmatrix}$$

Now we can see that there is a matrix [U] and matrix [V] which contain eigenvectors (left and right hand vectors to be specific) and a diagonal matrix of scalar values [S] called singular values; an earlier article about ten years ago describes this in much more detail. But of course the most important part of the SVD is that this can be written as the sum of each of the individual pieces that make up this FRF matrix.

Now when we perfom an SVD and plot the singular values, we will get as many curves as there are references for all of the spectral lines in the band considered. So if there are three references then there will be three singular values, namely, S_1, S_2, S_3 , which will exist for all the spectral lines in the band considered. Now if we plot these, then there will be three separate curves over that frequency band; the singular values are shown in the diagonal [S] matrix in the second equation. It is these three lines that are of interest for CMIF (but the eigenvectors will also play a small roll as will be explained shortly). A set of CMIF curves are shown in Figure 1. The upper set of curves shows one form of the CMIF and the lower set of curves shows another form of the CMIF.

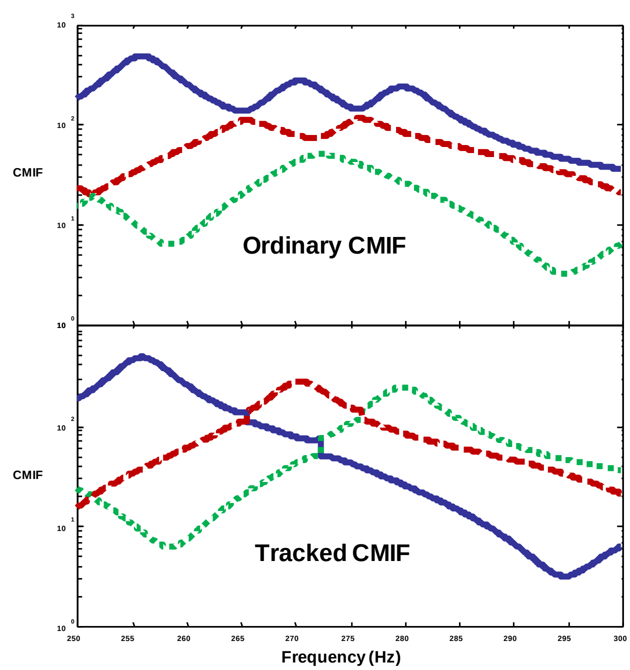

Figure 1 – CMIF (upper) and Tracked CMIF (lower)

Copyright 2013 Pete Avitabile

Now in the CMIF in Figure 1, notice that the blue curve is larger than the red and green curves in the upper plot. To interpret CMIF, wherever there is a peak in the CMIF, there will be an indication of a mode at that frequency; in the upper plot there are three peaks in the blue curve so there are three modes identified by the CMIF in this case. If there is a second peak ***at the same frequency as the first peak*** then there is an indication that there are two modes at that frequency – but the second curve ***must*** peak at the ***same*** frequency as the first peak, otherwise it is not another mode.

So what are the peaks in the red and green curve and why aren't they indicators of modes? Well these are a result of these crossover frequencies you asked about. It all comes down to what you want to track from the singular values. Do you want to track the biggest singular value or do you want to track the vector related to the singular value. The lower plot in Figure 1 tracks the vector associated with the singular value rather than the largest singular value. So you can see that when we track the blue vector, it peaks at a lower frequency and then steadily declines. We also see that the red vector starts off small and then peaks in the middle of the frequency band and then steadily declines. And then you can see that the green line starts off very small and then eventually peaks as the blue and red decline. So it makes a difference whether I track the largest singular value (upper plot) or the vector related to the singular value (lower plot) – it just depends on what you want to look at in the CMIF.

If this still isn't clear, I have a good example. If you happen to like the horse races, then I can relate the horse race to the SVD for the CMIF. Just before the horse race starts, all of the horses are lined up at the starting gate and they are all at the same point. But as soon as the race starts, different horses will end up in different positions and that will change as the race progresses. But do you want to track who is ahead or do you want to track your horse – it is just a matter of preference (and depends on whether you are betting to win, place or show, etc). Figure 2 shows a schematic of the race at different points during the race. Clearly, at different points during the race, a different horse is in the lead. In terms of who wins, you want to track who is ahead and that changes during the race. But you might also want to track your favorite horse but he may not be in the lead.

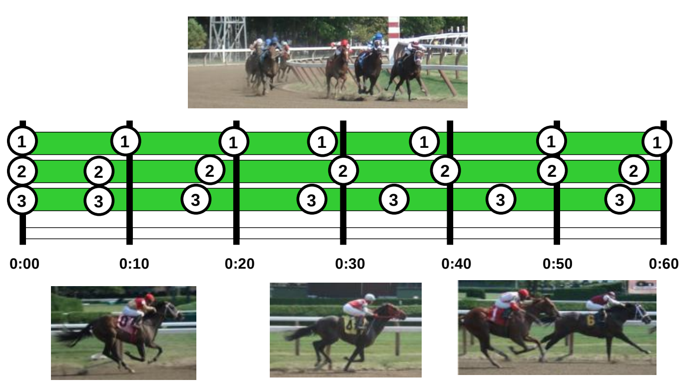

Figure 2 – CMIF/SVD and Horse Racing

In the first example, there was only one prominent mode at three different frequencies in Figure 1. But what would happen if there were multiple modes at the same frequency. Well then the CMIF would show one or more of the singular values peak at the same frequency in the CMIF plot. Figure 3 shows exactly this case where there are actually three modes at the same frequency at the first peak in the expanded view of the CMIF plot followed by three separate peaks higher in frequency – so in that expanded band there are actually six separate modes indicated by the CMIF plot. (That would be analogous to three horses nose to nose and if they were at the finish line, there would have to be a photo finish to determine the winner.)

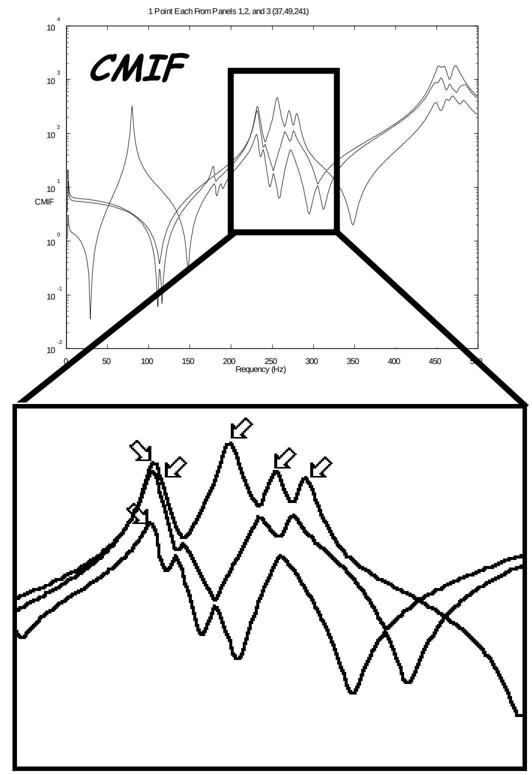

Figure3 – CMIF with Multiple Modes

So I hope this explanation helps to clear up any confusion related to the CMIF and the interpretation of the curves as well as why the crossover frequencies exist.

And for some reason, it always seems that I never pick the right horse no matter which way I track them. If you have any more questions on modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD by Pete Avitabile

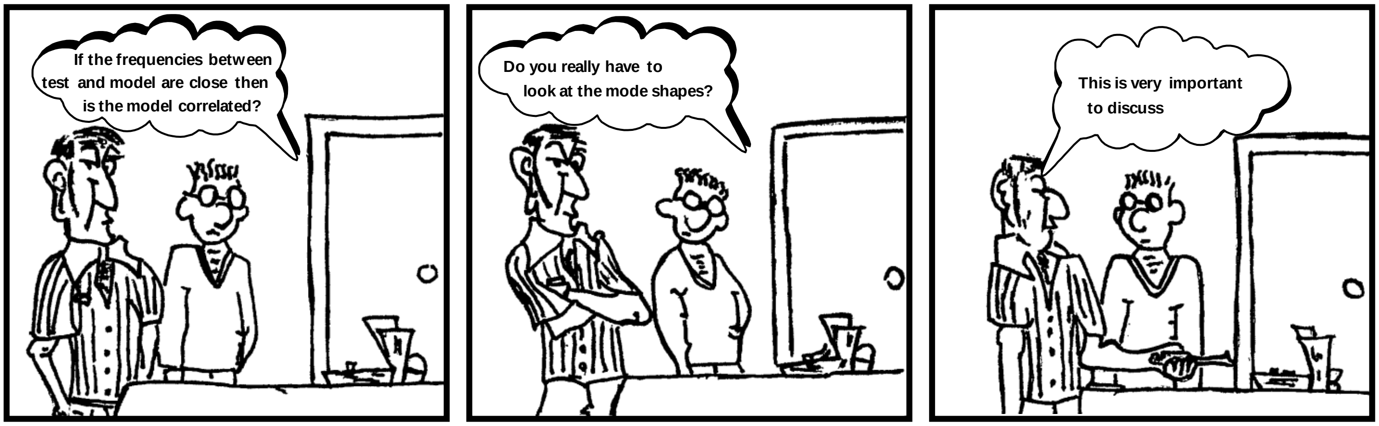

Illustration by Mike Avitabile

If the frequencies between test and model are close, then is the model correlated?
Do you really have to look at mode shapes? This is very important to discuss.

So this is yet another area where people often get confused. So many times people develop finite element models and there is a desire to make sure that the model is reasonable. Often times experimental modal tests are performed specifically for this sole purpose – to provide a sanity check for the model.

The finite element model is developed from many different assumptions and there are many areas where there is question as to:

- how the structure was modeled
- what material properties were used
- how the joints and connections were modeled

… and the list goes on and on.

This is because the finite element model is an approximation and it is a modeling tool that we use to assure that a design is reasonable for the intended use in its particular application. The model contains hundreds of approximations all of which may be reasonable for most designs. Fortunately, in design, we build in factors of safety and stress limits and other criteria to compensate for that which we don't know or understand.

Many times in the finite element model there are "penalty factors" or "knock-down factors" that are applied to the model because we just are not sure that we believe the properties that we use for the model possibly due to the manufacturing process used or because of fabrication techniques that may impose loads and degrade the general properties of the structure, etc.

The finite element models are approximations. We use the models to build in a "comfort factor" in the systems we design and build to have a greater confidence in our design. But the bottom line is that the finite element models we build are not perfect by any means. They are hopefully good approximations of the systems we try to dynamically characterize but … often times the approximations are forgotten in the development of the model.

For instance, everyone builds models that are very complicated and at times the complication is where the focus of concern may be directed. But many times simple questions as to what is the Young's Modulus or density of the material will raise an eyebrow as a possible source or error. Yet many times no one ever weighs the test article to see if the model weight and actual part weight are the same. And Young's Modulus is always just accepted as the published value with no thought of the variance that might be expected and how it should be checked.

And often times the CAD model is used for the model geometry generation without any real regard for what the actual geometry might be and how it may change the actual frequencies and mode shapes. One very important case is the flatness of panels that are included in a finite element model. The model may have the panel modeled as flat but the actual warpage of the plate may have a strong effect on the overall frequency prediction. Figure 1 shows some scans on panels that were stated to be flat and modeled as such in the model, but these deviations played a very important part of the frequency and shape determination.

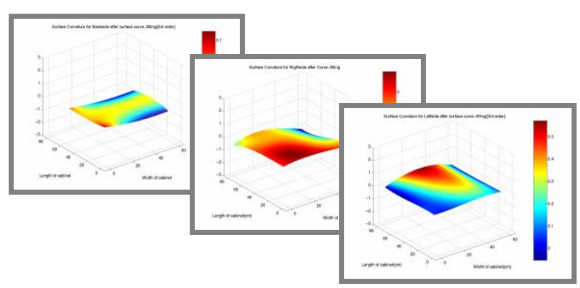

Figure 1: Distorted Geometry of "Assumed" Flat Panels

Copyright 2013
Pete Avitabile

In this particular case, the frequencies of the model matched many of the tested frequencies but the shapes showed essentially no correlation at all – so just matching the frequencies does not necessarily mean that the model is OK.

One particular difficulty that always causes problems is when the experimental modal test is performed with a clamped or built in condition. Where the finite element model can easily predict a built in or clamped condition, it is very hard to accomplish this in an experimental modal test set up – but people try to do tests this way all the time.

One case involved a composite plate that was developed with some new fiber and the characteristics for the composite plate model were being questioned. The plate was set up in a fixture to try to achieve a built in condition even though it was clearly known that the boundary condition may be difficult to achieve.

Unfortunately, the analytical group assured everyone that the fixture was more than adequate. And they clearly stated that it was "rigid", "stiff" and "more that adequate to simulate a built in condition for test".

Now the only thing that can be said here in terms of the test is that "It is what it is" and whatever is measured will identify the reality of the test set up. The first several modes of the composite plate were measured and a correlation study with the finite element model was performed. However, remember that the original finite element model fiber characteristics were not clearly known and the purpose of the test was to help define the fiber characteristics. Of course, the first correlation performed showed significant differences between the test and the model. This was anticipated to be largely due to the fiber unknowns. The frequencies showed quite a bit of difference but the shape correlation was considered reasonable as seen in Table 1. The three mode shape pairs are shown for reference in Figure 2.

Obviously, the model had some fiber characteristics that did not represent the true stiffness of the fiber in the composite arrangement. So with some minor tweaking of the fiber characteristics, the model was adjusted and the resulting frequencies were shown to be greatly improved. The model frequencies were then identified and were shown to be much improved as seen in Table 2.

But if you look at Table 2 you will notice that the MAC values are not shown. Due to time and budget constraints, no additional correlation studies were performed with the "adjusted" finite element model. Everyone felt that because the frequency comparison was greatly improved that there was no need to further validate the model – the frequencies are close so "end of story" so to speak.

Table 1: Correlation of Composite Plate Model and Test Data

	FEA (Hz)	EMA (Hz)	%Diff	MAC(%)
1	81.3	117.4	30.8	99.5
2	165.7	213.9	22.5	89.9
3	165.7	232.8	28.8	80.6

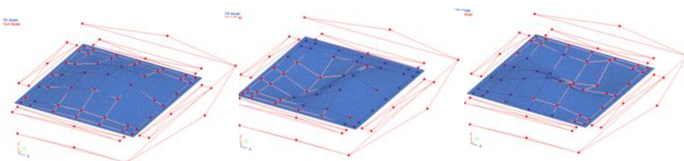

Figure 2: First Set of Correlated Mode Pairs

Table 2: Correlation of "Adjusted" Composite Plate Model and Test

	FEA (Hz)	EMA (Hz)	%Diff
1	113.7	117.4	3.1
2	233.7	213.9	9.2
3	233.7	232.8	0.4

But some time later several additional correlations were performed and as it turned out, the shape correlation had degraded significantly from the original correlation identified with MAC values for mode 2 and 3 much lower than before. While the frequencies appeared to be very close, everyone was happy to use the updated properties of the adjusted model. As it turned out there was significant effect of the boundary condition of the so-called built in or clamped condition. The frequencies had indeed gotten closer to the test frequencies but the fiber parameters of the model that were adjusted had a significant effect on the shapes of the first three modes – but that was never checked as part of the adjustment of the model. In fact, the boundary condition of the fixture was not as stiff as anticipated by the analyst and this had a big effect on the shape correlation.

So the bottom line is that the correlation clearly needs to compare frequencies and mode shapes. If only the frequencies are considered then the model can be adjusted (or maybe distorted) in just about any way to achieve "matched" frequencies. It is critically important that the shapes be evaluated as part of the correlation process to further justify and substantiate the correlation of the model to the test data.

There have been numerous instances to substantiate this and it is always recommended that the correlation of the model to the test include frequency comparison as well as shape comparison. The MAC is a first step to the shape correlation process but orthogonality checks are also needed for the validation of the model. (Some concerns of using only MAC as a vector correlation tool will be discussed in a future article.) If you have any other questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

Illustration by Mike Avitabile

Do you really need to measure FRFs? Or is a Transmissibility OK? We need to discuss this.

So there are many times that transmissibility is made for measurements in many different situations. This might be due to the fact that the data is collected during a shaker qualification test where the test article is mounted on a big shaker and all of the "device under test" accelerometers are measured relative to the base acceleration input to the test article.

Or it may be that the measurements are made on equipment in operation and the force cannot be measured and only response measurements with accelerometers are available. This is common when flight tests, vehicle test, suspension tests or similar tests are performed. This might be the only data that is available. But there are some slight differences that need to be noted. And it is also important to make sure that we are all using the same nomenclature when we use all these fancy words; sometimes I find that the words mean different things in different industries so it is always important to check the definitions are understood.

So let's make some simple definitions to explain some of the differences in all the measurements we may possibly make. If we make the following definitions,

- x(t) - time domain input to the system
- y(t) - time domain output to the system
- Sx(f) - linear Fourier spectrum of x(t)
- Sy(f) - linear Fourier spectrum of y(t)
- H(f) - system transfer function
- h(t) - system impulse response

then Fig 1 shows the input-output schematic for linear spectra.

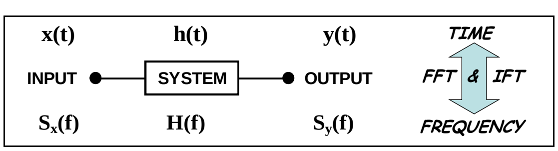

Figure 1 – Linear Spectra

And if we make these additional definitions,

- Rxx(t) - autocorrelation of the input signal x(t)
- Ryy(t) - autocorrelation of the output signal y(t)
- Ryx(t) - cross correlation of y(t) and x(t)

- Gxx(f) - autopower spectrum of x(t) $G_{xx}(f)=S_x(f)\bullet S_x^*(f)$
- Gyy(f) - autopower spectrum of y(t) $G_{yy}(f)=S_y(f)\bullet S_y^*(f)$
- Gyx(f) - cross power spectrum of y(t) and x(t) $G_{yx}(f)=S_y(f)\bullet S_x^*(f)$

then Fig 2 shows the input-output schematic for power spectra.

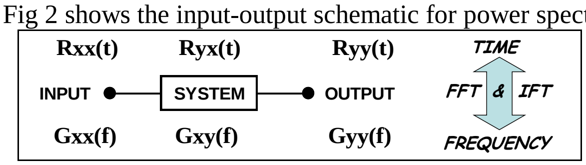

Figure 2 – Power Spectra

So now that we have these equations defined, let's identify some measurements we typically make and understand how they are used to compute things such as the FRF and Transmissibility for instance.

The first thing to notice is that Linear Spectra are complex valued functions that have both magnitude and phase – so S_x and S_y are complex linear spectra. But their companion power spectra, G_{xx} and G_{yy}, are not complex valued but are real valued, magnitude only measurements. This is very important because they have no phase information associated with them. But notice that the cross spectrum, G_{yx}, is a complex valued measurement that has both magnitude and phase.

So let's proceed and identify the FRF and Transmissibility. The FRF is the cross power spectrum divided by the input power spectrum whereas the Transmissibility is just the ratio of the output spectrum divided by the input spectrum. These are given as:

Copyright 2013
Pete Avitabile

$$H=\frac{S_y \bullet S_x^*}{S_x \bullet S_x^*}=\frac{G_{yx}}{G_{xx}} \qquad TR=\frac{S_y \bullet S_y^*}{S_x \bullet S_x^*}=\frac{G_{yy}}{G_{xx}}$$

So while we generally say that both measure the output relative to the input, there is a big difference – the FRF is a complex function with both magnitude and phase whereas the Transmissibility (TR) is just the ratio of the magnitudes; this is very different because of the lack of phase. But there is one very important difference. Generally the FRF is measured with a reference to a measured force whereas the TR has no force measured as it is typically obtained. This is very important when the data is needed for development of a calibrated model for model validation, structural dynamic modification, system model development, and forced response studies – a measured force is needed to calibrate the model so to speak.

So now that we have all the definitions out of the way, let's look at some measurements for a FRF and TR to show a few differences. A simple free-free beam will be used for some typical measurements. In the first measurement, a drive point FRF is made with an impact hammer and accelerometer. Figure 3 shows the Log Mag and phase for the drive point measurement; notice that the function is complex as shown in the figure. Figure 4 shows the Log Mag and phase for a cross measurement from one end of the beam to the other end of the beam; this measurement is also complex valued.

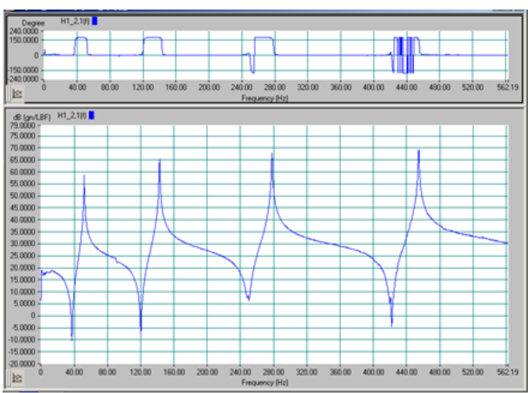

Figure 3 – FRF Drive Point Measurement on Beam

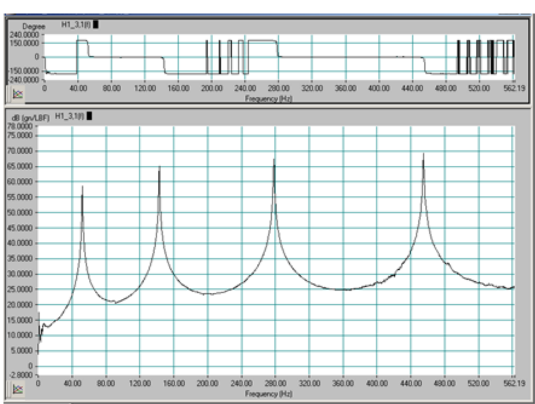

Figure 4 – FRF Cross Measurement on Beam

Figure 5 shows the autopower spectrum of the accelerometer at the drive point measurement location and Figure 6 shows the autopower spectrum of the accelerometer at the cross measurement location.

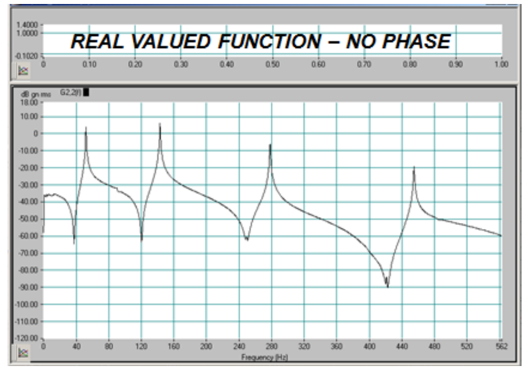

Figure 5 Autopower Spectrum of the Drive Point Location

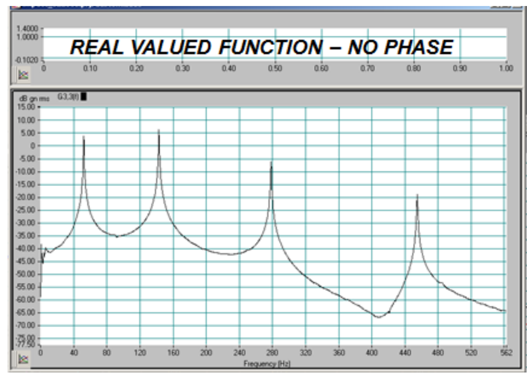

Figure 6 Autopower Spectrum of the Cross Point Location

Now both of the spectrum have some similar features when compared to the FRF measurements shown in Figure 3 and 4. But there is a very important piece of imformation missing which is the phase of the measurement. The power spectra are real valued functions but do not have any phase information. So when relating the magnitudes to each other there is no phase information that can be obtained from the measurement.

In order to have any directional information, there needs to be a complex measurement obtained so that phase is included. Now don't get me wrong here because the transmissibility can be very useful in many cases where no other measurement is possible. But we just have to make sure that we realize that there is some critical information needed if we want to understand the mode shapes of the structure. If you have any other questions about modal analysis, just ask me.

Illustration by Mike Avitabile

If we perform a roving impact hammer test and impact many points, is there any possibility of missing a mode? Well… you need to be careful where you place the accelerometer.

I am glad you asked this question because it is a very important consideration when perfoming a modal test. Let's discuss a few things in regards to your question.

Now a roving hammer with a stationary accelerometer is one way to run an impact test that is commonly used. The other way an impact test can be performed is by keeping the hammer stationary and moving the accelerometer. Both are acceptable ways to run this test and because of reciprocity, there really isn't a difference from a theoretical standpoint. In fact if you consider the measurements made you will fill one row of the FRF matrix when you have a roving hammer and you will fill one column of the FRF matrix if you have a roving response transducer. This is shown schematically in Figure 1 along with the reciprocal measurements.

Anytime you run a modal test, you always have to be careful to avoid having the reference located at the node of a mode. This is the most important consideration.

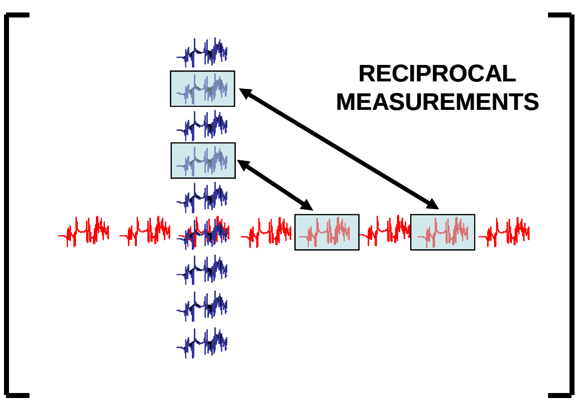

Figure 1 – FRF Matrix Showing Roving Impact (Red Row) and FRF Matrix Showing Roving Response (Blue Column)

In order to understand this, a few very basic equations describing the FRF equation need to be presented. Now a single FRF "ij" measurement can be written in summation form in terms of residues as

$$h(s)_{ij}\Big|_{s=j\omega} = h_{ij}(j\omega) = \sum_{k=1}^{m} \frac{a_{ijk}}{(j\omega - p_k)} + \frac{a_{ijk}^*}{(j\omega - p_k^*)}$$

But in this residue form of the equation, it is not so easy to realize what the residues imply. But if I write this equation using the residues expressed as mode shapes then

$$h(s)_{ij}\Big|_{s=j\omega} = h_{ij}(j\omega) = \sum_{k=1}^{m} \frac{q_k u_{ik} u_{ij}}{(j\omega - p_k)} + \frac{q_k^* u_{ik}^* u_{jk}^*}{(j\omega - p_k^*)}$$

The resulting FRF for this equation might look like the FRF shown in Figure 2 where the two forms of the FRF formulation are shown (and colored in blue and red for clarity of the individual mode contribution for each term of the FRF summation); in Figure 2, the equations above have been expanded for the first two terms of the summation to show the contribution that each mode makes to the total FRF. The important thing to realize is that the FRF is made up of the summation of each of the individual modes.

When we write this equation in terms of the mode shapes, it becomes very clear how the mode shapes of the structure have a strong influence on the amplitude of the FRF for a particular "ij" term. The residue is basically developed from a scaling coefficient, q, and the ***value of the mode shape at the output response location*** times the ***value of the mode shape at the input excitation location***. With that said, I think it becomes very clear that if either the output location or input location mode shape value is zero (that is, located at the node of the mode), then there will be no amplitude for that particular mode.

Copyright 2013 Pete Avitabile

$$h_{ij}(j\omega)=\frac{q_1u_{i1}u_{j1}}{(j\omega-p_1)}+\frac{q_1^*u_{i1}^*u_{j1}^*}{(j\omega-p_1^*)}$$

$$+\frac{q_2u_{i2}u_{j2}}{(j\omega-p_2)}+\frac{q_2^*u_{i2}^*u_{j2}^*}{(j\omega-p_2^*)}+\cdots$$

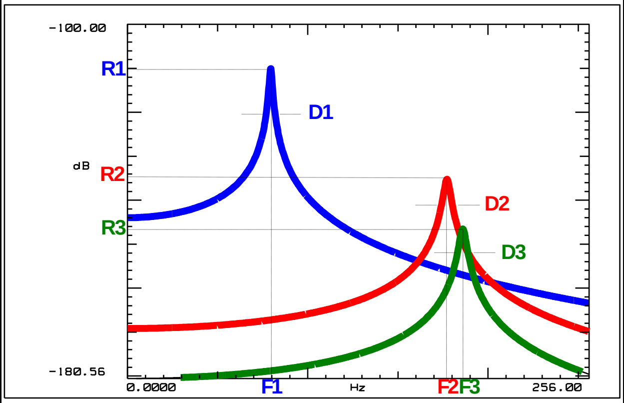

$$h_{ij}(j\omega)=\frac{a_{ij1}}{(j\omega-p_1)}+\frac{a_{ij1}^*}{(j\omega-p_1^*)}$$

$$+\frac{a_{ij2}}{(j\omega-p_2)}+\frac{a_{ij2}^*}{(j\omega-p_2^*)}+\cdots$$

Figure 2 – FRF Summation

So now that we realize this, it is also clear that any time we measure at the node of a mode then there will be no apparent amplitude in the FRF measurement related to that mode – and it doesn't matter if it is the excitation or response location.

And I should also point out that if the reference is located very close to the node of the mode then the amplitude of the FRF will be very low for that particular mode. Actually the 3rd mode shown in green in Figure 2 is a prime example of that. That amplitude is low compared to the 1st and 2nd mode because the value of the mode shape for the input and/or output location is much smaller than that for the 1st and 2nd mode; that's why the amplitude is much lower.

But as far as the original question you asked, what really matters is the reference location - whether it be the stationary hammer as a reference or the stationary accelerometer as a reference. If the reference is located at the node of a mode then there will be no apparent response in the FRF for that particular mode.

Now in order to conduct a good modal test, we need to have a pretty good idea of what the mode shapes are for all the modes of interest so that a proper reference can be selected.

But many times in more complicated structures or structures with very directional modal characteristics, it may be very hard if not impossible to select one location where it is easy to see all the modes from that one reference location. That is why so many times, we conduct modal tests with several references. This way we have the ability to see all the modes from a collection of different reference locations.

Often times, we will use 3 accelerometers when we have a 4 channel acquisition system (or 7 accelerometers if we have an 8 channel acquisition system), and perform a roving hammer impact modal test with stationary accelerometers at different locations on the structure. That way we have a much better chance to make sure that we can see all the modes from all the different reference locations.

Hopefully we can pick 3 (or 7 locations) from which all the references are not located on the nodes of modes for all the references. In fact you might think that it is almost impossible to pick that many references and have all of them be on the node of a mode all simultaneously.

Well as luck might have it, there was one modal test I saw where there were 9 references used with a roving impact hammer. Your first thought might be how could you possibly have a problem with missing a mode with 9 references. And wouldn't you know it, that all 9 accelerometers all happened to be all located on the node of one of the modes of the structure. Figure 3 shows this very unbelievable test that was run where, with 9 references, one of the modes of the structure was missed.

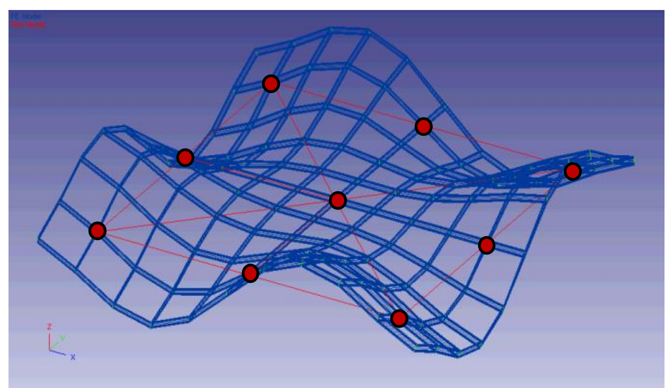

Figure 3 – Plate Mode Missed with 9 References

I hope that this helps to explain that you need to be very careful when identifying the reference for a modal test. If you have any other questions about modal analysis, just ask me.

Copyright 2013
Pete Avitabile

Illustration by Mike Avitabile

So now I have another question…if the hammer impacts a node, then there is no response elsewhere? That seems incorrect. OK… let's discuss this further.

So last time we discussed how the roving hammer test is a great way to run a modal test but that you needed to be very careful to not place the reference accelerometer at the node of a mode. We also discussed that it didn't matter how many points we impacted if the reference transducer was at the node of a mode then we would not see that mode.

So in answering the question I also stated that the same would be true for a stationary hammer test with a roving accelerometer. You still needed to make sure that the stationary hammer input at the reference location was not at the node of a mode – or else you would not see any response from that mode.

After we finished discussing it you seemed to not be comfortable with the fact that you could hit the structure at a node point and the structure would not move at that point and you mentioned it was counter-intuitive to you. So let me give a few more examples to try to get you more comfortable with this fact.

I will use a simple free-free beam and a simple plate to try to drive home some additional points for you. Let's first recall the FRF equation we wrote last time which can be in terms of the residues or mode shapes. But in this residue form of the equation, it is not so clear so I prefer the mode shape form of the equation because it is there that you can clearly see the effect of the mode shape on the peaks in the FRF for each of the modes.

$$h(s)_{ij}\Big|_{s=j\omega} = h_{ij}(j\omega) = \sum_{k=1}^{m} \frac{q_k u_{ik} u_{ij}}{(j\omega - p_k)} + \frac{q_k^* u_{ik}^* u_{jk}^*}{(j\omega - p_k^*)}$$

This equation is very clear in that the amplitude of the FRF is very much controlled by the ***value of the mode shape at the output response location*** times the ***value of the mode shape at the input excitation location***.

So now let's consider a simple free-free beam. A series of measurements will be made at 15 equispaced points along the length of the beam. The measurements will be made and then the FRFs will be plotted in a waterfall plot so the shape can be clearly seen. And in order to see this, it is very important that we use the imaginary part of the FRF to map the mode shape. We may want to recall that the imaginary part of the FRF will be a peak at the natural frequency while the real part of the FRF will be a zero; this is true only for displacement and acceleration measurements. Now FRF measurements were made over a wide frequency band but I want to zoom in to just around the resonant frequency for the first and second mode for this free-free beam.

Figure 1 (blue) shows the waterfall plot of the imaginary part of the FRF for all 15 measurments with a frequency band around the 45 Hz first mode. Figure 2 (red) shows the same plot but with a frequency band around the 140 Hz second mode. In both plots the peak of the imaginary part of the FRF is circled for all the measurements made on the beam. In Figure 1 (blue), it is very clear that the shape is that of the first free-free flexible mode of the beam whereas Figure 2 (red) very clearly shows the second free-free flexible mode of the beam. The most important thing to note right now is that the amplitude of the imaginary part of the FRF for mode 1 changes sign from positive to negative back to positive as you traverse down the length of the beam. At some point it crosses zero. At this location, there is no response for that particular mode. And it doesn't matter if it identified with a hammer input or accelerometer output.
The FRF will be zero at that point for that mode.

Copyright 2013 Pete Avitabile

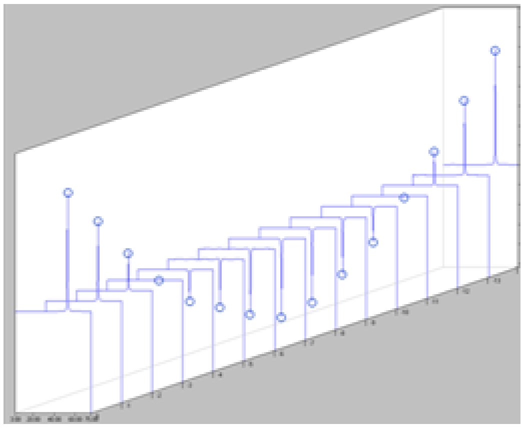

Figure 1 Waterfall plot of the Imaginary Part of the FRF showing the First Mode Shape

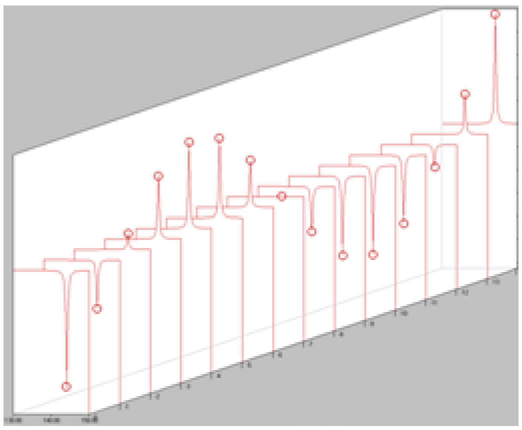

Figure 2 Waterfall plot of the Imaginary Part of the FRF showing the Second Mode Shape

So the node of the mode is a special location where there is no output for that mode. And it doesn't matter if you measure with an accelerometer or impact with a force hammer – there is no response because the value of the mode shape at either the output response point or at the input excitation point is zero and therefore there will be no peak amplitude for that mode. But there may be some response at that point due to the other modes of the structure which may not be related to the node of those other modes.

So in the case of the roving hammer, if the accelerometer is located at the node of a mode then it doesn't matter how many points are used for the hammer excitation, there will be no contribution to the response due to that particular mode.

And the converse is also true. If you have a stationary input that is located at the node of a mode, it doesn't matter how many points are measured with the response accelerometers, there will be no contribution to the response due to that particular mode.

So now let's extend it from the beam to a plate to see the same effect. A rectangular plate has been used in previous articles and is used here for this example. Figure 3 shows 6 FRF measurements around the perimeter of the plate where the first peak is related to the first bending mode of the plate (blue) and Figure 4 shows the same information but highlights the second mode of the plate (red).

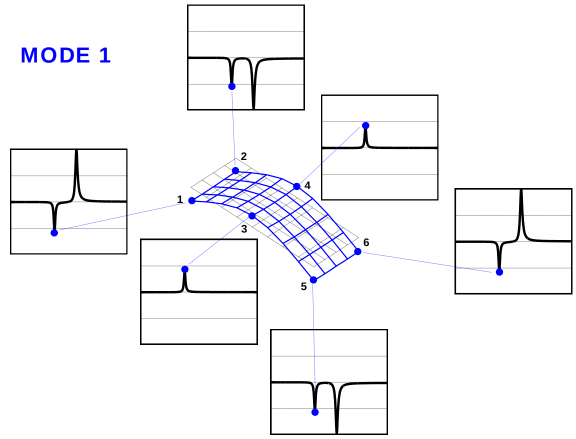

Figure 3 Plate Bending Mode

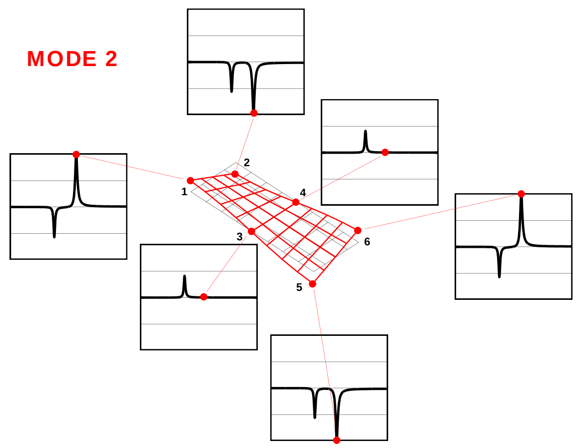

Figure 4 Plate Torsion Mode

So in looking at Figure 3 and 4, the same statements can be made regarding the node of the mode and the impact or response location. The mode shape can be simply described by plotting the imaginary part of the FRF. The function has positive and negative values depending of the mode shape of the structure. There must be a zero crossing at some point which corresponds to the node of the mode – this implies a point of zero response for that particular mode. And the response is zero at the node of the mode whether the measurement is related to the response accelerometer location of the hammer impact location – but just for that particular mode.

I hope that this helps to further explain these questions you had. If you have any other questions about modal analysis, just ask me.

Copyright 2013
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD

by Pete Avitabile

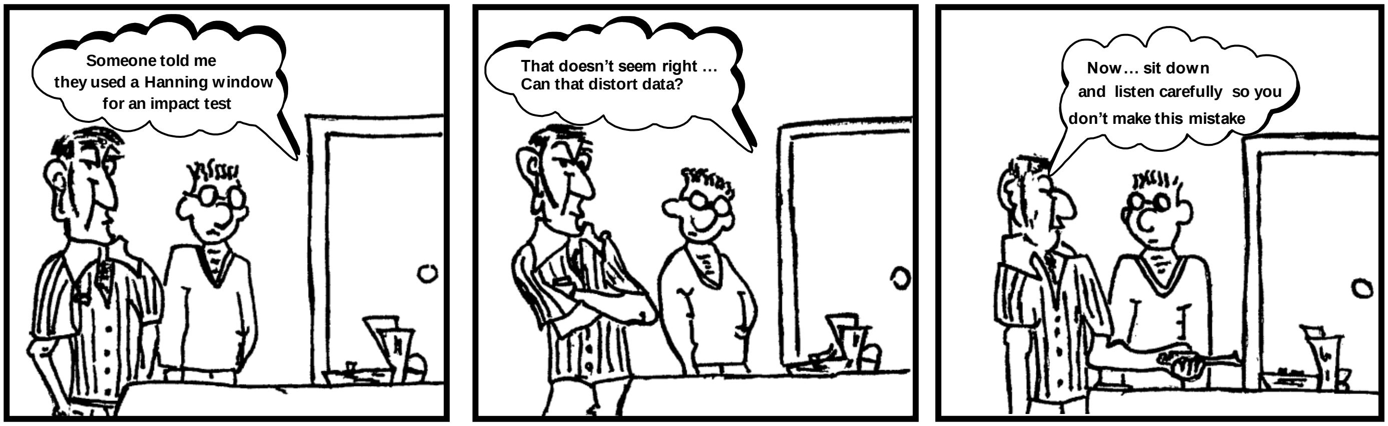

Illustration by Mike Avitabile

Someone told me they used a Hanning window for an impact test. That doesn't seem right…Can that distort data? Now… sit down and listen carefully so you don't make this mistake.

So…where do I start here??? First, let's just try to understand how things like this happen. Many times people take measurements in very difficult situations where the measurement conditions are not optimal to say the least. There are many instances where measurements are not the pretty textbook figures that we all wish we had for our measured frequency response functions.

There are plenty of situations where the system is in a very noisy environment, or the measurement transducers are not optimal, or the excitation is not sufficient to provide a measurable response, etc. And these are just some of the issues we face taking measurements. And please let's not forget that there may be nonlinearities (our arch-enemy) and complicated damping mechanisms (our arch-enemy's best friend) that all compound the measurement situation.

And because we have these types of difficulties and because we see them so often, we come to expect that all of our measurements are always going to have all of these difficulties. And then it becomes the rule, rather than the exception, that we come to expect this is just the way a measurement should be all the time.

But is that really so? Do all of our measurements really have such poor qualities all the time? Or is it just that we have become complacent and assume that's the way it should be?

So let's start by looking at the measurement that you provided in Figure 1. Oh my…..that is a really bad looking measurement. And at first glance I know that all of us could argue that this is from a nonlinear system, with a complicated damping mechanism, with a noisy environment with transducers that are the best that can be used to obtain this measurement and so on.

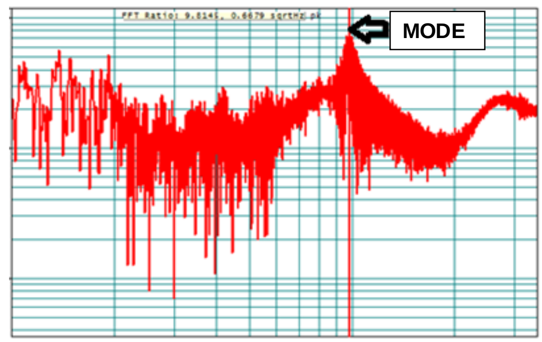

Figure 1: An example of a really bad measurement

But the real question is maybe why is this measurement so poor? Is it really from a nonlinear system? Is there really a complicated damping mechanism? Is there really a noisy environment? Are the transducers really that bad?

Or are all these just easy excuses that we can say because we really don't know or have just become accustomed to such poor measurements and assume "that's just the way it is".

I can't really comment on the measurement above other than you mentioned that it was from an impulsive excitation and that the measurement was made with a Hanning window applied – because that's the way it has traditionally been done. So the question is if this is really the right way to make this measurement.

So let's procede with a typical impact measurement on a general structure; actually this is a composite rib stiffened spar type structure that has been used before in other measurement situations. So the first thing we will do is to make a measurement with what would be the appropriate measurement

Copyright 2013
Pete Avitabile

dignal processing parameters and then retake the measurement but use a Hanning window to show dramatic differences.

This first frequency response measurement (FRF) is shown in Figure 2 along with the coherence and the input excitation and output time responses. Now for this particular configuration, there really isn't any need to use a window on the input or output because the measurement is completely observable within the sample interval and satisfies the periodicity requirement of the fourier transform process. Notice that the coherence is very good and the FRF is also very good for this measurement. ***And just to be clear if any window was to be applied then it would be the exponential window for the response.***

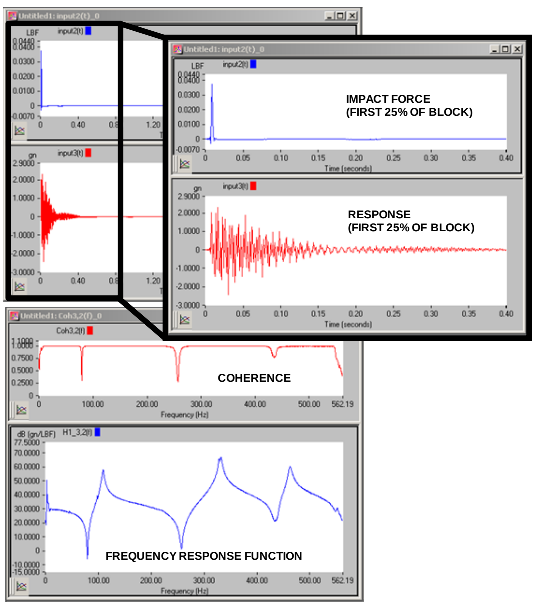

Figure 2: An example of an impact measurement with appropriate signal processing parameters applied

Now… let's make a measurement and apply a Hanning window on the measurement which is shown on the left column in Figure 3. ***Now please make sure you understand this is <u>not</u> the way to take this measurement but I am going to show exactly how bad this measurement can be.*** The input excitation and time response are the same but you can see that the FRF and coherence for this measurement are terrible – and terrible is an understatement of how bad this measurement actually is.

But what is confusing is that the time signals really don't look terrible by any means. Well what you have to realize is that the signals that are shown are the raw measured data and do not show the effects of the windows on the data. So it is not really showing how the data has been affected by the Hanning window applied in the time domain – but certainly the frequency domain shows a dramatic degradation of the measured FRF and coherence.

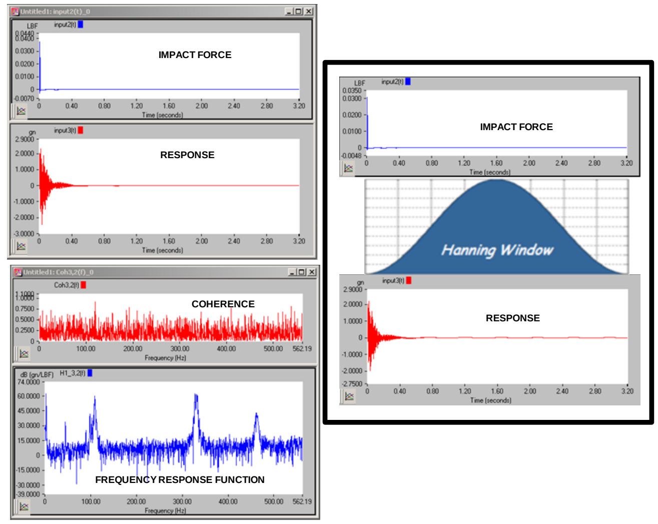

Figure 3: An example of an impact measurement with totally inappropriate signal processing parameters applied

In order to understand what happens when the Hanning window is applied, the right column in Figure 3 shows the two time signals with the Hanning window also displayed with the window appropriately scaled to visually compare with the measurements made. Now if you look closely at Figure 3 you will quickly see that the Hanning window is going to seriously attenuate the beginning of the time record of the input excitation and output response – and in fact will essentially weight all the important information regarding the transient response to zero thereby leaving a measurement which is essentially a measurement of the noise in the system. This now becomes very clear why the FRF and coherence in Figure 3 are so poor – the measurement has been essentially reduced to noise.

But if you weren't paying attention or have been misguided to think that this is just the way measurements typically look, then you might think that this was "the best measurement that can be obtained under the circumstances". But the reality of the situation here is that the signal processing parameters to process the data have been totally, incorrectly specified and the measurement has been completely distorted by this.

I understand that there are measurement situations where there are difficulties due to all the reasons mentioned **<u>but</u>** that does not give you the right to inappropriately process the data and cause additional errors "because you think it doesn't make a difference". In the case shown here, all the distortion was due to incorrectly processing the data and, in so doing, totally good data was converted into terrible, ugly measurements that are not acceptable under any circumstances.

I hope that this helps to explain the questions you had. If you have any other questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD   by Pete Avitabile

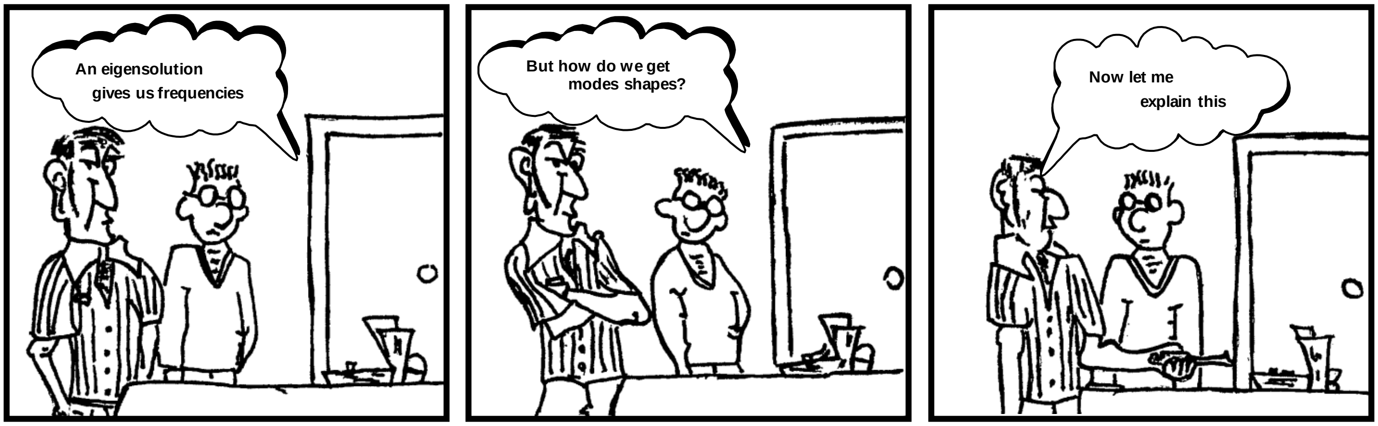

Illustration by Mike Avitabile

An eigensolution gives us frequencies. But how do we get mode shapes?
Now let me explain this.

OK. I think the first thing that we have to say is that the eigensolution actually gets us both the frequencies and mode shapes. The mathematical process of the eigensolution can be performed a number of different ways. There are what are called direct techniques and indirect techniques for the solution.

For smaller matrices, the direct techniques decompose the set of equations to get all of the eigenvalues and eigenvectors. Techniques such as Jacobi, Givens and Householder are common methods that are used.

But when the matrices get larger, like those of the large finite element model that are generally developed today, then an indirect technique is used where only a few of the lower order modes are obtained. Techniques such as Subspace Iteration, Simultaneous Vector Iteration and Lanczos are some of these indirect techniques that are used.

But I really don't want to make this article a math class or really get into the details of the solution sequence. So let's discuss the eigensolution and what we are attempting to do when we find the frequencies and mode shapes. I want to explain it so it makes sense to you.

So let's write the eigensolution in general form.

$$\left[[K]-\lambda[M]\right]\{x\}=\{0\} \qquad (1)$$

The first thing that I want to say is that the eigenvalues can be found from the determinant of the matrix. Well, that determinant will really be nothing more than a high order polynomial whose roots are the eigenvalues. Now numerically those can be obtained from any root solving algorithm such as Secant Method or Newton-Rapson Method as a few well known approaches.

So the eigen equation and a typical polynomial that may result is shown in Figure 1. The function zero crossings are the roots where the polynomial is zero.

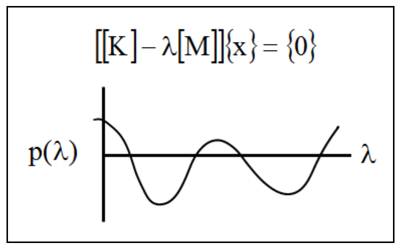

Figure 1: Graphical Representation of Roots of Determinant

Now that gives us the frequencies of the set of equations and the next step is to determine the mode shapes. Well, if you take the first eigenvalue, $\lambda = \omega_1^2$, and substitute it into the eigensolution equation, then you can solve for the $\{x_1\}$ vector because you know [M], [K], and ω_1^2. The solution for that vector is straightforward using any decomposition scheme such as Crout-Doolittle, Cholesky, LDL decomposition to name several well known popular approaches.

So that $\{x_1\}$ vector is actually the mode shape for that particular frequency that was used to solve the set of equations. Figure 2 shows this schematically for the first free-free mode for a simple beam; note that blue is used to identify this as the first mode of the system. And if you follow through with the equation in Figure 2 you will notice that the elastic forces are equal to the inertial forces in the way they are written. We could also say the the beam is in dynamic equilibrium at that frequency which is ω_1^2 . And if you looked at the system from an energy perspective you could see why there are node points where the system oscillates about those points and there is equal positive and negative parts of the shape to keep it in equilibrium.

Copyright 2014
Pete Avitabile

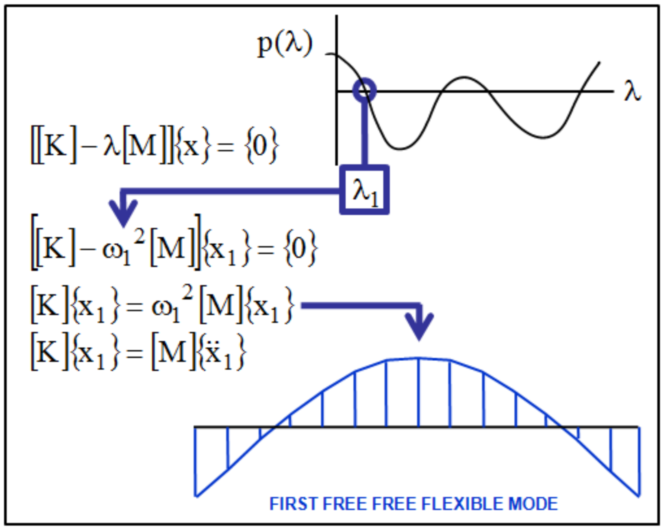

Figure 2: Schematic for eigensolution for mode 1

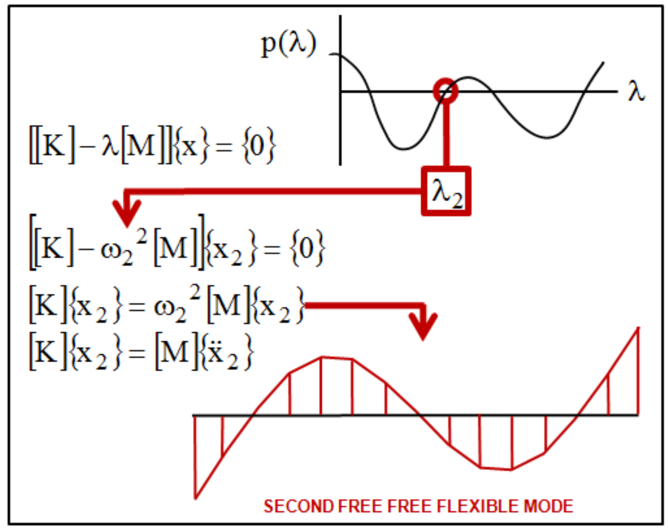

Figure 3: Schematic for eigensolution for mode 2

Of course now we need to do the same thing for the second frequency. If you now take the second eigenvalue, $\lambda = \omega_2^2$, and substitute it into the eigensolution equation, then you can solve for the $\{x_2\}$ vector because you know [M], [K], and ω_2^2. Now the $\{x_2\}$ vector is actually the mode shape for the second frequency. Figure 3 shows this schematically for the second free-free mode for a simple beam; note that red is now used to identify this as the second mode of the system. Again let's follow through with the equation in Figure 3 in red you will notice that the elastic forces are equal to the inertial forces in the way they are written. We could also say the the beam is in dynamic equilibrium but now at the frequency which is ω_2^2 . And just like we did with mode 1, we will see that the node points are locations where the system oscillates about those points and there is equal positive and negative parts of the shape to keep it in equilibrium.

We then continue this process for all the modes of interest. Of course the way I explained it may not be the way the different solution algorithms actually decompose the matrices and obtain the final answer. But the way I have explained it will probably give you a much better overall idea how the frequencies and mode shapes come from the system set of equations.

So it is important to realize that the eigensolution is used to obtain what is called the eigenpair – that is, the frequency and the vector associated with the eigen-equation. This is in fact the mode shape.

Now another thing to realize is that the mode shapes are linearly independent and the mode shapes are also orthogonal with respect to the mass and stiffness matrices. This is a by-product of the eigensolution. This is a very important fact that is often used when we check our finite element model with measured experimental data. We perform a type of orthogonality check, often called a pseudo-orthogonality check, to compare the measured experimental vectors with those obtained from the eigensolution.

I hope that this helps to explain the questions you had. If you have any other questions about modal analysis, just ask me.

Copyright 2014
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD by Pete Avitabile

Illustration by Mike Avitabile

Do you need to mount triaxial accelerometers at all locations? That can result in more channels needed and more cost. Let's discuss this and think about it.

Well, let's talk about this as well as a few other things that relate to how your structure is instrumented.

First, let me say that triaxial accelerometers are very useful in many, many applications. They allow for a very compact package to be used to monitor all three directions from one physical mounted transducer on a structure. Yes, I do use them but as it turns out I really don't use them all the time and there are many cases where I absolutely will not use triaxial accelerometers and we will discuss some of the reasons why.

First of all, we all realize that we can "make" a triaxial accelerometer by mounting three separate accelerometers onto one mounting block. Now of course this is not as elegant as a triaxial accelerometer, but it is one simple and economical way to accomplish this. And of course it also means that you can buy three separate accelerometers which are just about the same price as the triaxial accelerometer. But remember that there is one different distinction.

When I mount the triaxial accelerometer on my structure under test and I really only need to measure one direction, then I have wasted two accelerometers for all practical purposes. And if someone else needs to make some measurements, you have three accelerometers tied up for each measurement location whether or not you use all three. Now if you had three separate accelerometers then you wouldn't be tying up all the accelerometer inventory! Now this may sound silly but when you don't have a lot of accelerometers and all of them are triaxial, then you have tied up a lot of instrumentation if you really only needed a single axis accelerometer. I have seen some laboratories that have bought all triaxial accelerometers and when there are multiple tests to be run, all the instrumentation is tied up on one test.

OK. So now let's discuss a few more things. Let me first start with a simple free-free beam test. (You know that all of us "educators in academia" all test beams all the time.) So if we want to test a simple beam and find the modes in transverse bending in only one direction, then we would have a test set up with something like that shown in Figure 1 where there are 15 measurement locations along the length of the beam.

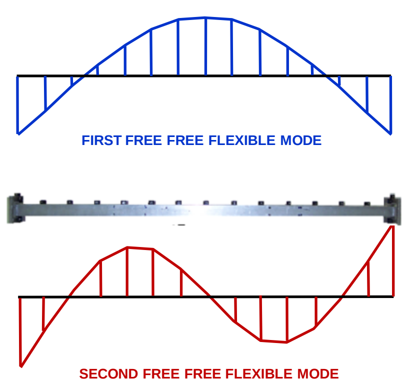

Figure 1: Schematic Planar Beam Modal Test

Now if all I had was traxial accelerometers, I would be tying up 45 measurement transducers and really only needing 15 of those for the measurement at hand. Now of course you could argue that I might need to also test the other planar beam bending direction too and would need another 15 accelerometers for that. But I still would have 15 measurement transducers that are not utilized if I really didn't need the axial direction too.

And again you would probably say that this is an academic situation and you may really need all those triaxial accelerometers for a typical application. So I will agree but let me show a few cases where you might want to rethink this.

Copyright 2014 Pete Avitabile

Recently we have done quite a few big wind turbine blade modal tests where the main interest is the bending in two directions referred to as the flapwise and edgewise modes of the anchored wind turbine blade. (And so realize this is nothing more than a "really big" beam for all practical purposes.) Figure 2 shows the schematic for a 9 meter wind turbine blade test with some accelerometer configurations. Notice that there are only measurements in 2 directions (x and y) because the axial direction is really not of interest. This test was run with a very portable 8 channel system with 7 accelerometers and one hammer. When the test was run the first set of measurements were made with 7 accelerometers at 7 points but all in the x-direction. Then the accelerometers were all reoriented to the y-direction for the second set of measurements. Eventually the accelerometers were all roved to all the points of interest. Now one advantage of using single axis accelerometers here was that all the cables remained attached to the accelerometer and DAQ as they were reoriented and then roved to all points. In this way, there was never a concern that there were any cable swaps resulting in a mismatch between accelerometer location or direction. Had all triaxial accelerometers been used then there is a much greater possibility of getting cabling screwed up.

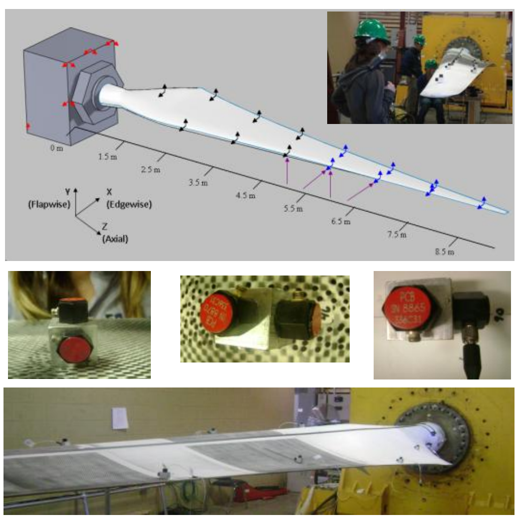

Figure 2: Schematic for 9 Meter Wind Turbine Blade Test

Another wind turbine blade test was performed for a turbine blade that was in the 50 meter long range. This test also was only really interested in the flapwise and edgewise modes of the blade but several argued that it may be necessary to also measure the axial direction too. Figure 3 shows the blade test with cabling configuration and expected mode shapes for the test along with a related measurement.

But the axial direction is very stiff compared to the two flap and edge motions and the displacement is very small. Now I will say that in this test we actually did mount triaxial accelerometers just in case we finally needed to measure all three directions but fortunately many realized that there was very little to measure in the axial direction. But there was another very important concern that many never really consider.

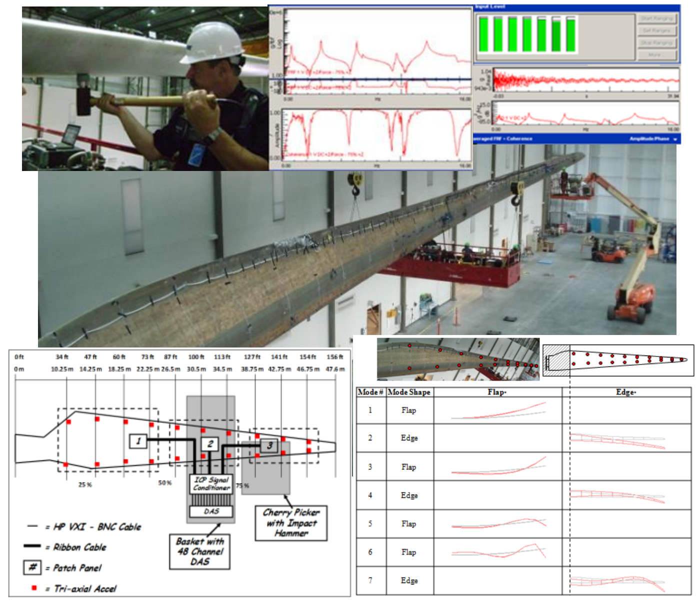

Figure 3: Schematic for Large Wind Turbine Blade Test

The flap and edge motion is large and an accelerometer sensitivity of 100 mv/g is very suitable for the motion in these two flexible directions. However, the motion is very small in the axial direction and a sensitivity of 1V/g or higher is necessary in order to make a good measurement. The problem with a triaxial accelerometer is that the sensitivity in all three directions are nominally the same – so the measurement in the axial direction with a triaxial accelerometer with 100 mv/g would be plagued by noise and poor signal strength and for all practical purposes would not provide a suitable measurement at all!

So here is where I end this article with the very clear statement that for this last test scenario, I would be much better off with three separate accelerometers with sensitivities that are suitable for the motion to be measured for the test of the large wind turbine blade. A triaxial accelerometer would not be the wise choice for this test with everything considered.

But I will point out that we did mount triaxial accelerometers for this test but it was mainly to allow us to pre-cable the entire blade before it was hoisted up on the test stand for the test. And yes I did have one accelerometer direction that we never measured during the test. And in fact we only measured two directions with flap and edge for test and never wired up the axial accelerometer channel for any of the measurement locations. And if I had wired up all three directions of each triaxial accelerometer, then I would have needed more DAQ channels than what were available on my acquisition system used. But I used the triaxial accelerometer "just in case"!!!

I hope that this helps to explain the questions you had. If you have any other questions about modal analysis, just ask me.

Copyright 2014
Pete Avitabile

 by Pete Avitabile

Illustration by Mike Avitabile

Does gravity come into play when testing to find frequencies? Does orientation make a difference? This is something that needs to be discussed.

There are several different things to discuss in regards to this. There are several different things that may come into play when performing a modal test and we need to talk about this for sure.

Generally, from a theoretical standpoint, the effects of gravity will not have an effect on the frequencies and mode shapes for most of the situations we face. That is because the equations we write for the definition of the system are written about a static equilibrium standpoint and the effects of gravity are not necessarily an influence (but in a moment I will discuss practical situations where this is not necessarily true).

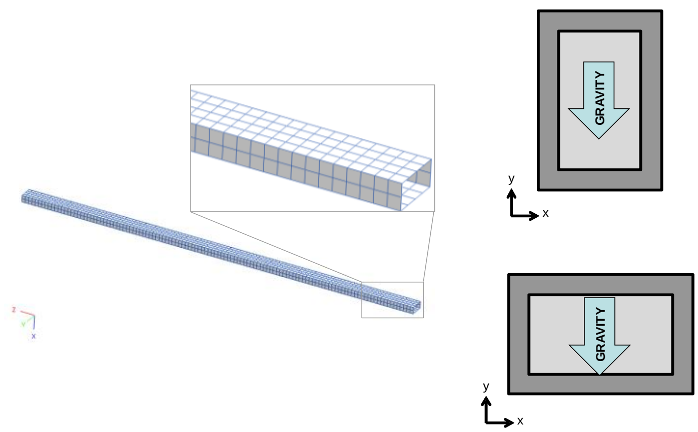

Figure 1: Schematic of Beam for Modal Test

Let's consider a simple beam as shown in Figure 1.   Now when we make the finite element model of this beam, it really doesn't matter which way we orient the beam relative to gravity and there generally will not be any difference in the frequencies if we chose the cross section orientation on the top right or the cross section orientation shown on the lower right. The frequencies computed will be the same because gravity is not considered and we are assuming that the structural configuration is evaluated about the static equilibrium point and that *there is essentially no significant deformation of the structure due to the effects of gravity*. (At least that's what we are assuming when we make the finite element model.)

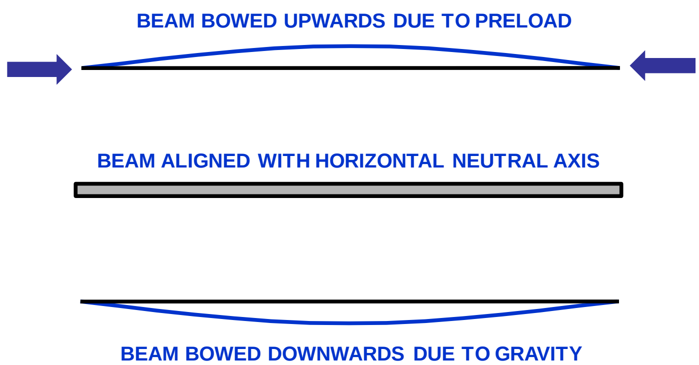

Figure 2: Schematic of Beam in Several Configurations

But it is that last statement that needs some additional considerations. Figure 2 shows several configurations that may need to have some additional discussion. The first configuration is the one shown in the middle section in Figure 2. Here the beam is assumed to be oriented along the neutral axis and there is no significant deflection of the beam. In this case the beam orientation really makes no difference at all. The tested frequencies will not be affected by the orientation of the beam relative to gravity.

Copyright 2014
Pete Avitabile

But now let’s consider the configuration in the upper portion in Figure 2. In this case an axial load has been applied to cause the beam to bow upwards. When this happens then there really is an effect on the stiffness of the beam. The stiffness is increased because of the upward bow of the beam. And we know that this is expected to be true. The arched configuration is actually stiffer than the flat configuration. Just consider any bridge span and the girders are always slightly arched because that is a stiffer configuration. So if a compressive preload is applied and the beam deflects upward, then this arched configuration is slightly stiffer than the nominal undeformed configuration.

So the lower configuration in Figure 2 shows a deformation due to the gravity load. It stands to reason that the deflection due to gravity will have an effect on the stiffness especially for very flimsy lightweight configurations such as wind turbine blades. Now if the beam is rotated and the stiffer cross section is now taking the load due to gravity, then the deflection will be significant lower and the effects of gravity are minimized significantly.

So theoretically, the gravity load does not create an effect because the assumption is that the deflection is small and the gravity effects are insignificant. But if the static deflection is more pronounced due to the flexible nature of the structure then the assumption may not be a reasonable one to make. Then the orientation of the beam can have a significant effect.

So when might this become a concern. Well, for large wind turbine blades, the structure is very flexible and the orientation of the blade cross section (flap vs edge) can have a significant effect on the effective stiffness of the blade due to the orientation of the blade with respect to gravity.

So Figure 3 shows two wind turbine blade configurations and both are sensitive to the orientation of the blade with respect to gravity. The turbine blade is strongly affected by the orientation with regards to gravity and the natural frequencies will be affected by the orientation. Looking at the blade it is very obvious that there is a difference in the two orientations. The flap (weaker axis) direction is much more sensitive to the effects of dead weight than the much stiffer edge direction.

But there is another consideration that many often overlook. The dead weight loading causes a deflection in the structure as expected. But that dead weight deflection may cause some significant loads on some of the internal members – these loads may cause enough deflection in the spars and ribstiffening internal structures that they are deflected into a configuration that is much different than the nominal dimensions on the design drawings. These deflections will basically cause the internal stiffening members to have a much different stiffness that that of the nominal dimensions based on the design drawing dimensions. This is very similar to the preloaded beam shown in the upper portion of Figure 2 which is much different than the nominal dimensions.

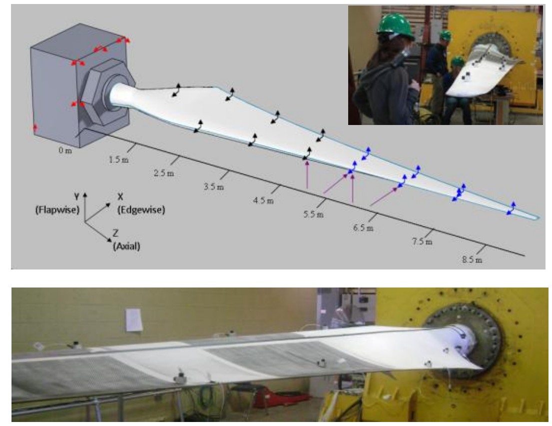

Figure 3: Schematic for 9 Meter Wind Turbine Blade Test

The ribstiffened airfoil panel configuration in Figure 4 is a very good example of a cross section where this may be of concern; the sketch on the right shows a simple wing configuration where the left wing (blue) has essentially no deformation due to gravity but the right wing (red) shows significant deflection due to gravity. For this thin, flimsy panel configuration (red), the dead weight or structural loading can cause deflections which may warp the rib stiffened interior structural panels – and this loading may cause significant deflection to result in a geometry that is no longer following the nominal dimensions identified on the CAD drawings. Therefore the stiffness of these internal ribstiffening members may be very sensitive to the orientation of the test structure when considering gravity loading.

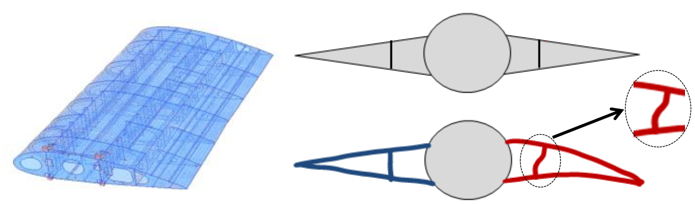

Figure 4: Ribstiffened Panel Airfoil Configuration

So usually we don’t have to consider the effects of gravity unless these effects cause significant deflections and seriously change the geometry defining the finite element model.

I hope that this helps to explain the questions you had. If you have any other questions about modal analysis, just ask me.

Copyright 2014
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD by Pete Avitabile

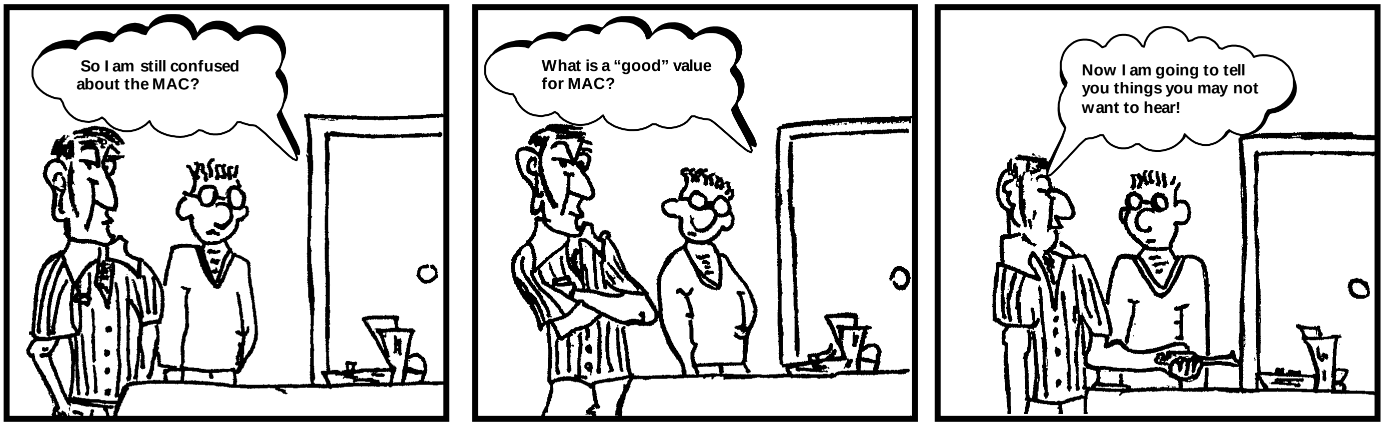

Illustration by Mike Avitabile

So I am still confused about the MAC? What is a "good" value for MAC?
Now I am going to tell you things you may not want to hear!

This is a subject that I find many people have a hard time understanding and accepting. Many would like to act like an ostrich and "stick your head in the ground" and hope that the problem will go away. Well…I am about to show some things here that many people have a hard time accepting. But first let me just start with a few words concerning correlation and orthogonality because there is a big difference between the two and many would like to think that the Modal Assurance Criteria (MAC) is the same as the orthogonality check – but in fact they are dramatically different.

As a formality, the two equations describing the MAC and Orthogonality, respectively, are:

$$MAC_{ij}=\frac{[\{a_i\}^T\{b_j\}]^2}{[\{a_i\}^T\{a_i\}][\{b_j\}^T\{b_j\}]}$$

$$ORT_{ij}=\{u_i\}^T\big[M\big]\{j_j\}^T$$

And the first thing to point out is that the MAC is really nothing more than a vector dot product that is scaled such that the values will range between 0 and 1. And if the value of the MAC is close to 0.0 then we say that there is little correlation between the two vectors. And if the value of MAC approaches 1.0 then the two vectors are very similar. But you notice that the word *orthogonal* was never used. Only the word *similar* was used.

Now the orthogonality is a mathematical property that results from the eigensolution of the mass and stiffness matrices that describe the system. A by product of the eigensolution is that the vectors are *"linearly independent"* and the vectors are *"orthogonal with respect to the mass and stiffness matrices simultaneously"*. So the orthogonality is a property that is guaranteed as a result of the eigensolution. The MAC has no such guarantees with that calculation.

The orthogonality check is a much more rigorous check that is performed and often times it is mandated as part of the certification process in the aerospace and military applications. The analytical/finite element model mode shapes are often compared to the measured experimental vectors from test. The governing bodies have mandated that the mass orthogonality of similar vectors must be greater that 90% or 95% and that different vectors must have values no higher than 5% to 10%. That is to say that the diagonal terms of the mass orthogonality matrix must be greater than 90% and all the off-diagonal terms must be lower than 10%. Now in these industries, the MAC is not typically used for the validation of the model because the orthogonality is a better correlation identifier.

Now what about other industries. Well there really isn't a mandate or governing body so many times companies or industries have "good practices" that are generally followed. The MAC really resulted from a test environment where test engineers wanted to identify if the measured shapes from one test to the next were similar, or one prototype was similar to a production configurations or … (on and on with many different ways that we can use the MAC).

So why didn't they use the orthogonality check. Well, remember the MAC started from the testing guys and back 30-40 years ago, the test guys didn't have access to a mass matrix; only a few of the analytical engineers had access to very primitive finite element modeling tools. Plus the MAC was an easy calculation to perform. Let's face it, they didn't have the mass matrix and didn't want to bother with the much more intensive mass orthogonality check. The test guys were just trying to get some simple comparisons.

But then everything "grew up" and all of a sudden we had people using the MAC to correlate the finite element mode shapes with the measured test data. And all of a sudden we had people "correlating" models using MAC and then they started to use some of the same general criteria that were developed with orthogonality checks.

But often times I will hear people saying that they will accept MAC values for correlated vectors with values that are lower than 90% and sometimes accept correlation values as low as 80% because "we are working with real structures and not simple academic examples". Well I am not so sure that I agree with that mentality because there is no mass matrix involved in the MAC calculation.

I want to show two examples to show some MAC values that result from modes that are clearly not similar at all. The first case is to compare the rocking rigid body mode of a free free beam with the first cantilever mode. Figure 1 shows that there is about 60% similarity indicated by MAC.

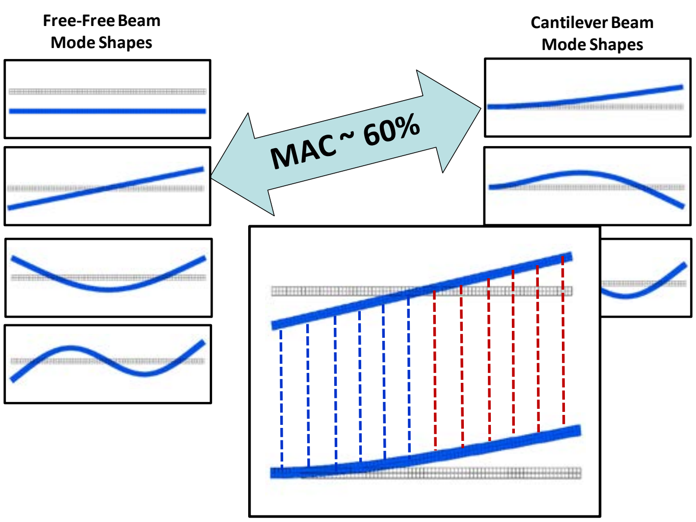

Figure 1: Mode Shapes for Free-Free Beam (left) and Cantilever Beam (right)

Good Golly Miss Molly! These two vectors have nothing to do with each other but MAC shows 60% correlation. How could that possibly happen. Well if you look lower right inset in Figure 1, you will see that the values of the cantilever mode shape from the base to midspan of the beam are very small and their contribution in the MAC calculation are very small. And if you look at the mode shape from the midspan to the tip of the beam, the values are much larger and with the naked eye you can see that the rocking rigid body mode looks quite similar. And that is what the MAC is indicating. But we know these two modes are not similar at all. Now you may argue that I have no right to compare these modes, but I did just to illustrate what a 60% MAC indicates. How bizarre is that !!!!!

Now let's proceed on with two cantilever beams – one with a uniform mass distribution and one with an additional lumped mass at the center of the beam (20% of the weight of beam). Now if I looked at the MAC between Mode 1 and Mode 2 with the uniform mass distribution, the MAC will show little correlation as expected. But when the MAC is performed on the beam with the additional lumped mass, the MAC between Mode 1 and Mode 2 is almost 80%. *Missy Molly cannot believe this at all.* But if you look at the uniform beam modes in the upper right portion of Figure 2 and compare them to the modes in the lower right portion in Figure 2, you start to see the same type of issue as discussed in the first case. The mode shape values close to the built in end have very little contribution to the MAC. And the values of the shape towards the end of the beam for Mode 1 and Mode 2 look very similar. And that is what MAC is indicating again. But we know that these two modes are orthogonal to each other but the MAC completely breaks down here. However, the mass orthogonality clearly identifies the vector status because the proper mass distribution was included in the orthogonality check. The MAC has no way to account for the uneven mass distribution for this case.

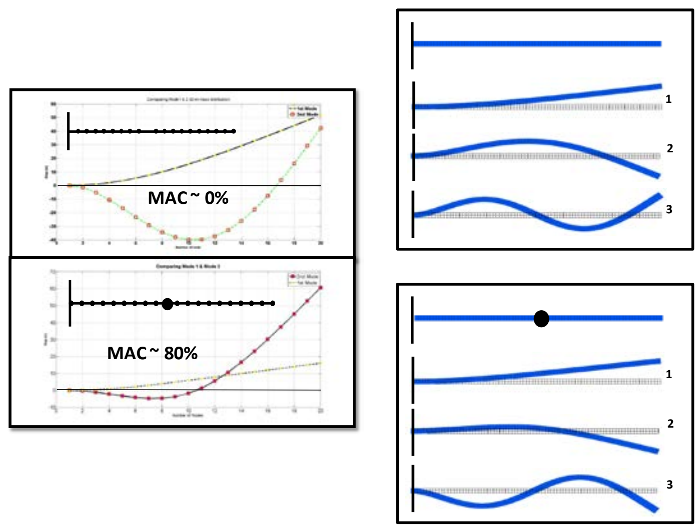

Figure 2: Cantilever Beam Mode 1 and Mode 2 with Uniform Mass Distribution and Center Lumped Mass

So I hope you can see that the MAC values can be deceiving. The MAC has no mass matrix. That is its biggest strength ***<u>and</u>*** its biggest downfall. Now don't get me wrong…I use MAC all the time to help with sorting out model and test results. But I am very leary when the supposedly "correlated modes" have MAC values that do not have very strong indicators.

I hope that this helps to explain the questions you had. We will likely talk more about MAC in a future article. If you have any other questions about modal analysis, just ask me.

MODAL SPACE - IN OUR OWN LITTLE WORLD by Pete Avitabile

Illustration by Mike Avitabile

So what is the best way to make a free-free test set up … because nothing is really free-free.
Alright … let's discuss some non-traditional ways to do this.

Alright … so many people always ask me what is the right way to set up for a free-free test. Well there is no right way but certainly there can be very poor ways to do this.

The most common way is to use bungie cords (or something similar depending on the weight of the structure). I have seen missiles supported from bungie cords as well as large wind turbine blades. I have seen airbag systems deployed in many different instances. Well, I could go on and on with all the different ways we could do this. But the bottom line is that you have to make sure that the boundary conditions are not intrusive on the system under test. The boundary condition should have little effect on the flexible modes of the system. When this is done, then we can say that the test set up has very little effect on the flexible modes of interest in the test structure.

But we actually need to check that to make sure that the test set up does not have an effect on the flexible modes of the system. We need to set the structure up with one set of support locations and then retest the structure with a different set of support locations or change the stiffness of the support at the support location (possibly by using twice as many bungie cords or changing the pressure in the air bag support system for instance). If the flexible modes of the system do not change appreciably then the support condition likely has little effect. But there might still be some effect from the boundary condition and it needs to be carefully checked.

So as an example, I have two structures that were recently tested and some very non-traditional boundary conditions were used. The first is a smaller lighter weight structure that has some very closely spaced frame bending and torsion modes whereas the second structure is a much heavier anchor plate used for some shock response spectrum testing work.

The support for the first structure was actually inspired from a phone conversation with a close colleague where he had mentioned in class that you could use almost anything for an isolation system. A student asked what extremes could be taken and he quickly, as a funny remark, said "I don't care if you use marshmallows if you want". Well hearing that I decided to test one of our standard lab structures with various sized marshmallows; very small mini-marshmallows and very large jumbo marshmallows were used to perform a modal test for our frame structure in the lab. This particular frame is designed to have the first bending and first torsion mode to be very, very close in frequency to the point of being almost repeated.

So the first test (Test #1) was set up with 4 jumbo marshmallows located at the four mid-section of each leg of the frame which corresponds to the node points for the torsion mode. The second test (Test #2) was set up with 10 mini marshmallows distributed around the frame. These two tests were performed and the first thing that was noticed was that the bending and torsion modes were swapped depending on which of these first two tests were used. So a third test (Test #3) was set up where the jumbo marshmallows were located at the corners of the frame.

The rigid body modes were definitely affected by the arrangement of the marshmallows. But it is important to note that the flexible modes also showed a little frequency difference in each of the different configurations. So the boundary condition does have a little effect on the flexible modes of the system. But more importantly, the sequencing of the bending and torsion modes occurred differently in Test #1 and Test #2. So it is very important to realize that the support condition may have an important effect on the frequencies of the modes as well as the organization of the modes. Notice that Test #2 and Test #3 however, have the same organization of the mode sequencing for these two tests. So not only do we need to be cautious about

Copyright 2014
Pete Avitabile

the shifting of frequencies, we also need to be concerned about the organization of the modes due to the test set up. Figure 1 shows the results of the first two modes for the three different test set up configurations along with the photo of the structure with marshmallow support and typical drive point measurement for each configuration.

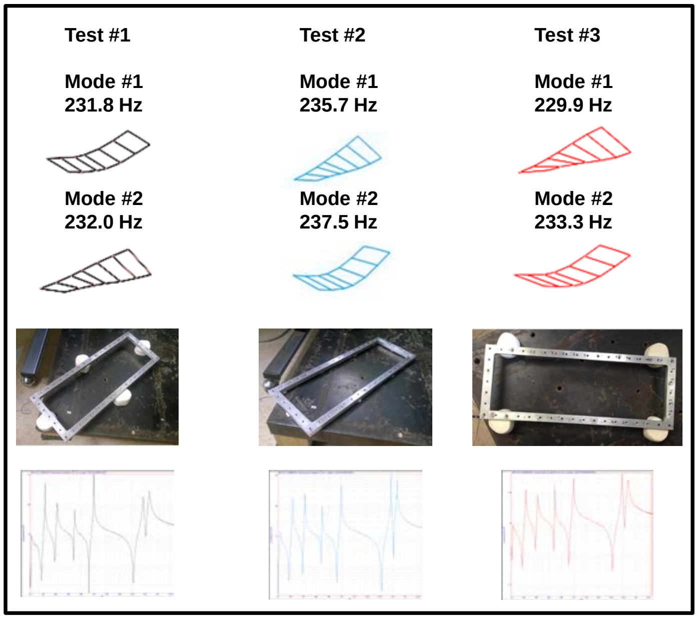

Figure 1: Results of Frame on Marshmallow Support

Now the second structure is a shock response spectrum plate structure that is available in the lab. The structure is mounted up onto an air-piston floatation system. But before the air-piston system was available, a very quick modal test was needed to validate the model and make some preliminary shock predictions. Without the air-pistons to support the plate, a very crude floating support was devised. Now at a university, money is always limited so a pratical, economical support needed to be provided. After some long thought one day, a brilliant idea came upon me. That handy old toilet plunger seemed to be a very good possibility for the support of the shock plate.

The hardware store was quite surprised when I appeared at the cash register with 6 toilet plungers. Our 250lb shock plate was tested with two configurations – one with 3 plungers located at the locations of the air-pistons and one with 6 plungers.

And the results of this test were very good. The rigid body modes in both configurations were very good and the flexible modes were similar as seen in Figure 2 (left) for the three plunger configuration and in Figure 2 (right) for the six plunger configuration. And the results were so good that a recent visit from a European colleague sparked the question where could he buy some of these plungers on E-Bay (to which I replied to just go down to your local hardware store and buy some brand new ones – they really only cost about $5 per plunger and were a bargain compared to some of the more expensive configurations that people have concocted).

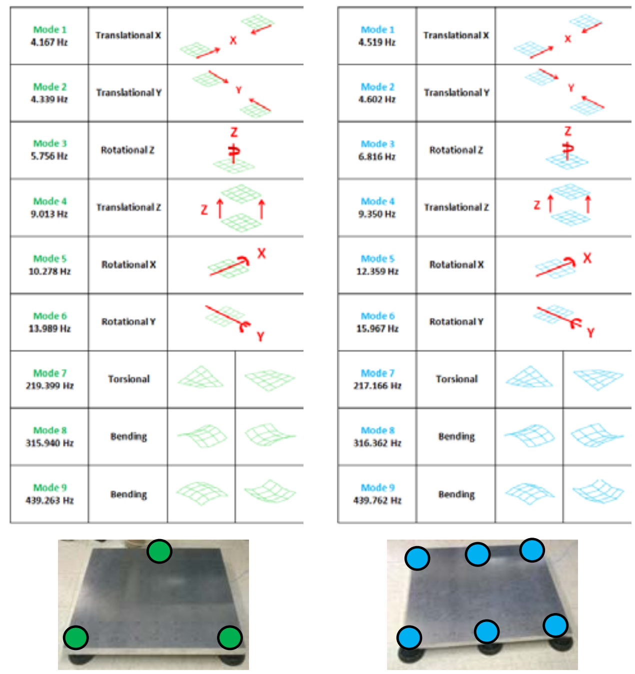

Figure 2: Results of Anchor Shock Plate on Toilet Plungers

So you can see that there can be a number of different and very simple mechanisms to create the free-free test configuration. But you do need to be mindful of the fact that the boundary condition may cause some shifting of the frequency that may be important to the further use of the data for subsequent analyses and that the boundary condition may have an effect on the organization of the different modes of the system as was seen in the first test arrangement for the frame structure.

I hope that this helps to shed some light on the questions you had. You can see that from marshmallows to toilet plungers, many different support conditions can be used to accomplish the support for the system under test – you just need to be careful and check your measurements. If you have any other questions about modal analysis, just ask me.

Copyright 2014
Pete Avitabile

MODAL SPACE - IN OUR OWN LITTLE WORLD by Pete Avitabile

Illustration by Mike Avitabile

So before you end this Modal Space series … can you provide some last pieces of advice?
Why sure … So here is my Top Ten list of important items.

Well, if you look back over all the articles, there is plenty of information that will help you to be careful when you are involved in a modal analysis project – whether it be analytical or experimental. But surely there are some really important items that you need to be aware of when working in this area. While there are many items that could be listed, I am going to make a Top Ten list (because ten seems like a good number to pick). So let's count down these items that I have selected as some of the "top" things to watch out for.

So in David Letterman style "Here we go"

Number 10 …
Why are you performing this test
Why ask why?
Well that is because it is the most important question to ask…..

So let me elaborate on this a little bit here. Many times we conduct tests because someone believes that the test will solve some problem or that it is a test that someone "thinks" will solve a problem.

I have no problem performing a test but many times people really do not realize what the test may or may not provide. That is the reason for me to always ask "Why" do you want to run this test. Is there an operating problem? What additional items are expected from the test? What frequency range is really of interest? How many modes are really of concern? And on and on. So it is really important to find out as much as you can before you run the test to make sure everyone is all "on the same page" in terms of what the test will provide.

And I say it that way because I have seen many, many instances where people have "claimed" to understand the test and are adamant about they want from the test and have been very clear as to what they wanted. But then once the test results are provided, then there are questions as to why the test does not answer the questions of interest. And sometimes the disconnect occurs because sometimes the words we use may mean different things to different people. So generally I always ask very specifically what people want to know and I very specifically ask what they mean (with an explanation) by each of the things that they have requested from the test.

And as an example I remember a group of young engineers in the automotive industry want to "learn" how to do modal testing and how to "correlate" to a finite element model for a simple brake rotor configuration. All the right questions were asked and it seemed like a very good effort to try to understand the very basic material and learn how to take some baby steps to understanding what is necessary before undertaking a much more complicated system. OK – so it seemed like all the right discussions were made and everything considered.

But before the project started, this goup of young engineers wanted to make a presentation to their management as to what they were about to undertake – again a very good thing to do to get everyone to "buy into" the project. Everything still seemd to be going smoothy until they introduced the project in this way.

"Hello everyone. This will be a project that will perform testing on a brake rotor to correlate to a finite element model. The results of this project will solve our brake squeal problem". And that was the first time they ever mentioned brake squeal. So suffice it to say, the squeal problem and what was originally discussed, were completely disconnected.

So why ask why? That is exactly why!!!

Number 9 …

Selection of appropriate test points

Often times I see people start a modal test and they get all wound up selecting all the points for measuring and make an elaborate geometry file and get all the coordinates lined up – but they haven't taken a single measurement.

Before you go head over heels making a geometry, go out and make a measurement first. Actually make a few measurements. Check different measurement locations and in different directions. This is critical especially if you really don't know what all of the modes of the system might be. It doesn't make any sense to select all the points until you have some idea what all the modes might be for the system.

Often times, the points you think you need to measure may not actually be the best locations depending on the modes of the system. Somehow, in my mind, I think that the FRF will tell you so much about the structure and frequencies that you really need to worry about that first.

Then maybe take just a handful of measurements to make sure you really do know what the mode shapes might be for the structure. Once you are sure you know what all the mode shapes might be, then you can select many more measurements but with the understanding of what the shapes might be. Too often I have seen people identify 100 to 150 points, run the modal test, curvefit the data, and then all sit looking at the mode shape only to realize that they placed all their measurements on a portion of the structure that really has very little to do with the modes of interest for the structure.

Also make sure that your reference location for your FRF measurements is at a location(s) where you are sure that you can see most if not all of the modes. Certainly if all the modes cannot be seen then it is imperative that additional references be used. When performing impact testing it is always advisable to use as many references as possible.

If you have a 4 channel system then you should have one channel for the hammer and three references on the structure. They don't have to be oriented into each of the three directions – X, Y, Z. But you want to make sure that they are all located to see as many of the modes of the system as possible.

If you have an 8 channel system then use 7 references if you are doing a roving impact test. You might think it is overkill but it really doesn't take much additional effort to collect the data and it never hurts to have more data.

And you think with 7 references you would get all the modes – well most times you would think so. But I can recall one test on a large symmetric composite plate structure where 9 reference accelerometers were used to run the test. But as it turned out one of the higher modes was missed because all of the 9 accelerometers ended up located at the nodes of this higher mode. Who would ever guess you could be that unlucky. (I recommended that this guy never go gamble in Las Vegas because his luck was obviously bad.) The higher order mode and the measurement locations for the 9 accelerometers are shown in the Figure 1.

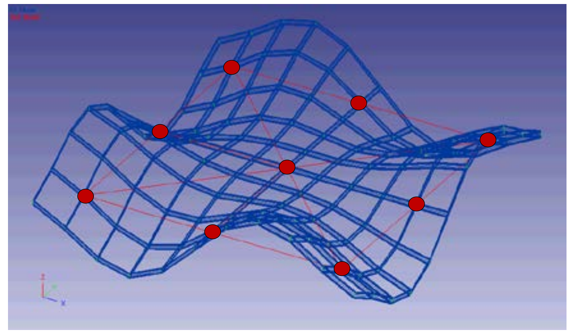

Figure 1: Nine measurement locations unfortunately all located at the nodes of this particular mode

Number 8 …

Hammer tip selection

Now selecting the proper hammer tip can sometimes be confusing to the novice. Basically what you want to do is make sure that you select a hammer tip that will excite a frequency range similar to the range of frequencies that will be excited when the structure is in service. Of course that means that you have to have some idea what frequency range is really important. I remember back many years ago when we started doing some modal testing on baseball bats, there was a very long discussion as to what would be the best tip to use. I explained that you needed to have a hammer tip that would excite a similar range of frequencies as those excited by the actual ball hitting the bat. The next day when I arrived in the lab the students had taken a baseball and put a 10-32 tapped stud into the baseball and then screwed that onto the hammer. Of course, this was a brilliant idea because it is as close as we can get to the actual impact scenario for the ball hitting the bat.

But you also have to remember that the hammer tip is not the only thing that controls the input force spectrum. The local flexibility of the structure can also play a critical role in the actual force spectrum imparted into the structure for the modal test. So you really need to look at this closely. And by the way, you can take those published curves you get from the hammer manufacturer and just put them aside because those are all generated by impacting a massive, stiff steel block which is never what we actually have when we perform a modal test.

Another critical item in impact testing that is often not taken seriously is that the hammer must impact the structure ***consistently*** with the same point impacted in the same direction for every measurement. If this is not done then the FRF will have some variability between each measurement which will result in reduced coherence. On a large structure this may not be hard to do. However, on a smaller structure this can be difficult.

One test for a golf club head utilized a unique tripod/hammer configuration to consistently impact the same point in the same direction for every measurement as seen in Figure 2.

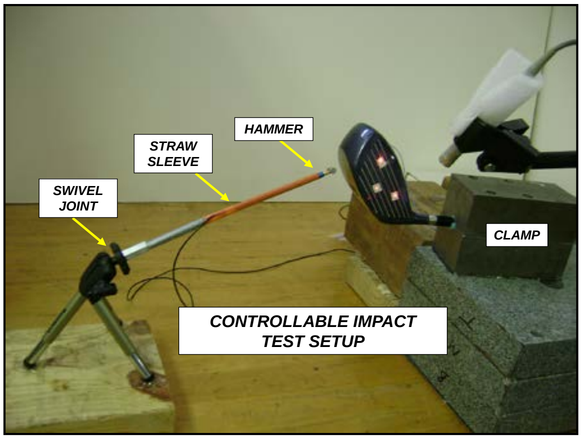

Figure 2: Impact hammer test configuration

Number 7 …
How free does it need to be
Well there have been a few articles on this subject. The most important thing to realize is that your test article is actually your structure plus all the instrumentation and support condition. Your finite element model of your structure can be modelled as free but the reality is that there are soft springs that really need to be included in the model to properly account for the support system for your structure along with all the instrumentation added. Many times this does not affect the overall test but in many cases this is actually very important to include in your analysis of the structure.

But what you really want is for the rigid body modes of your structure to be reasonably well separated from the flexible modes of the structure and have little modal overlap or coupling between the rigid body modes and the flexible modes. While this is very easy to say, often times this is not so easy to achieve. Most times I recommend that the finite element model include the effects of the support structure in the model so as to get a clear understanding as to how the test set up might interact with the test article. While the finite element model may not be perfect, the model is a great way to study the effects of stiffness changes in the support structure and the corresponding effect on the flexible modes of the system overall.

But if there is no model available, then this needs to be checked when the test is set up to identify exactly what the interaction might be for the test configuration. This might take some extra effort but it is a critical part of the test set up that needs to be documented and identified.

So a test where this was of concern was when missiles are tested. It is very hard to get them into a free-free condition. So the best we can do is to test the missile hung from a gantry and perform the test with the missle supported at the nodal locations for the first flexible mode; then the support condition is not very intrusive because it is supported at the node of the mode. Figure 3a shows a typical missile configuration with Dilbert performing the impact test here; Figure 3b shows a smaller missile undergoing shaker modal testing.

Figure 3a: Impact test on missile hung at locations that are close to the nodes of the first bending modes

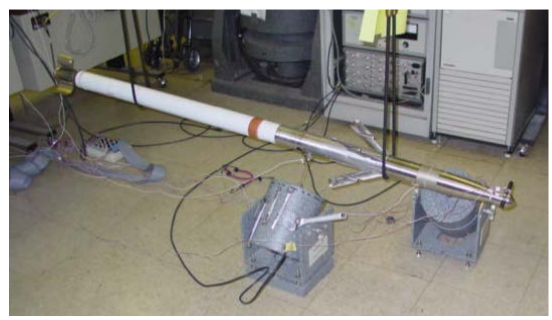

Figure 3b: Shaker test on missile hung at locations that are close to the nodes of the first bending modes

Number 6 …
Some other common blunders
There are always some of the most simple things that often get overlooked. These are the simple sanity checks to make sure that everything is set up properly.

Make sure all your cables are good and have not been crimped or bent. Make sure all the connectors are tightly connected. Often spurious signals, especially with the impact hammer, may be the result of a loose cable connection.

Of course make sure all your signal conditioners are turned on. And make sure that you understand if your transducers are either voltage or ICP. I have seen many tests run where the ICP transducers were set as voltage transducers and the measurements are essentially useless. Of course you would have expected that the measurements would not look good but if you go into the measurement process assuming that you have a very complicated, nonlinear, heavily damped system then you are expecting your measurements to not look good.

Of course if your measurement system is not set up properly your measurements won't look good and therefore you may think that this is the best you can do – even though your measurements are terribly wrong.

You also have to realize that if you only own one hammer that does not mean that it is useful for ALL the tests you plan to conduct. I have seen people violently wailing away on a large structure with an impact hammer that is clearly too small to excite the structure and insufficient to conduct the test. (And believe me I have seen some hammer tips that appear that they have been exposed to a nuclear explosion they are so severly battered to death.) Get an appropriate sized hammer to conduct the test you need to perform rather than try to use an inappropriate hammer for the test.

Another important consideration is in regards to the size of the accelerometer that is used. Mass loading can be a very important consideration. There have been many articles written to understand these effects. This needs to be addressed and documented. Just because it is the smallest accelerometer that you own does not mean that the mass loading is not of concern. And it is not just the mass of the accelerometer relative to the total mass of your test structure, it is the mass relative to the effective mass of the structure where it is mounted. An accelerometer weight at a very stiff/massive location on a structure is different than that same accelerometer mounted on a thin lightweight flimsy panel in the same structure.

And one more important item is that you need to make sure that you have not saturated your transducers in which case they will not be able to provide useful measurements. I have been party to tests where people have bought very sensitive transducers because they think they are "better" but then only to find out that their structure is very responsive and the response saturates the transducer.

Number 5 …

<u>Double impacts</u>

Now we really do want to avoid double impacts if at all possible. But there will be many instances where we just can't avoid them. So try your best to impact with single impacts. But if you do have a double impact, then the thing to do is look at the input power spectrum of the force hammer. As long as the force spectrum is reasonably flat and there is no significant dropout in the force spectrum and the FRF/Coherence looks good, then most likely the measurement will be adequate for the test to identify the frequencies and mode shapes.

But of course you can ask how flat does the force spectrum need to be and how much of a drop in the force is tolerable. And these are good questions to ask. I would rather not see the force spectrum drop more than 5 to 10 db but as long as the coherence is good then the FRF may be acceptable for a measurement.

I know some people might argue and say that much of a drop is totally unacceptable. But if you look back in some of the articles we have shown that the frequencies and mode shapes were actually very acceptable when comparing a test with no double impacts and a test with several or even quite a few double impacts. But you still need to be very careful to make sure that the data is useful.

And just for the record, there were a few articles that discussed double impacts and one article where multiple impacts were intentionally applied to the structure for a "burst impact" excitation test. While that was shown on an academic structure, over the past year we actually tested a large radio telescope and a large (50 meter plus) wind turbine blade and very cleary showed that the multiple impact technique provided far superior results. The measurement in Figure 4a/4b shows the result of an FRF measurement on a very large wind turbine blade with the coherence. The first measurement (4a) is made with a single impact and clearly the variance on the FRF measurement and the coherence show that the measurement is contaminated with noise. But the next measurement (4b) shows the result for the multiple impact and it is very obvious that the FRF and coherence are dramatically improved with the multiple impact test technique used. Of course you need to be careful to make sure that the entire input and output are observed within one sample interval of the FFT time window, but if that is done then the measurement can be very much improved.

Number 4 …

<u>Windows</u>

I am sorry to say that as far as I am concerned, no window is a good window – any window distorts data – windows are a necessary evil. These are strong statements that I live by.

Do everything possible to assure that your input signal and response signal are either periodic in the sample window or entirely captured within the sample interval. If you can do this, then you don't need to use any window.

When performing impact testing, always try to change the acquisition parameters such that the signal can be completely observed in one sample interval of the measurement process. If this can be done then there will not be any leakage and a window is not needed. Figure 5 shows how simply changing the sample time, the need for a window can be eliminated.

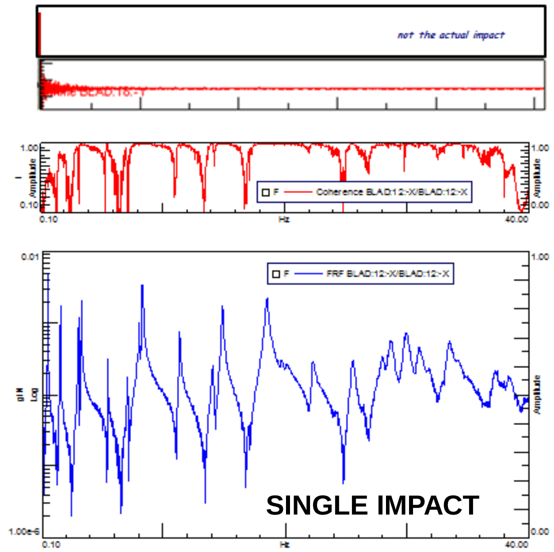

Figure 4a: Single impact FRF for large wind turbine blade

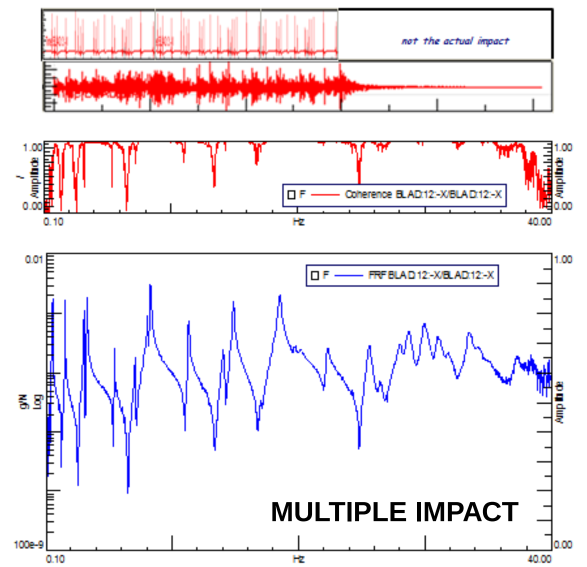

Figure 4b: Multiple impact FRF for large wind turbine blade

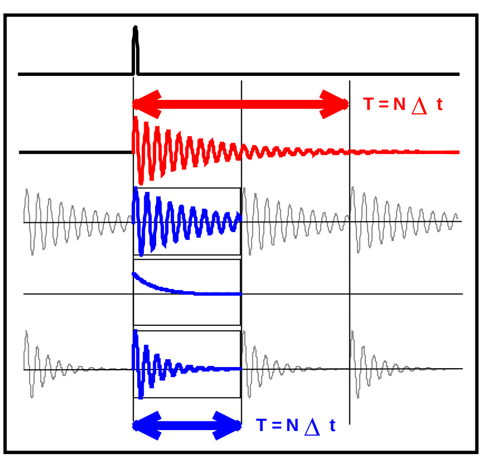

Figure 5: Window required for shorter time sample in blue can be eliminated by changing the time sample in red

And actually the same is true for shaker testing. But in this case we try to create a sample of data which is completely measured within one sample of collected data (the same as was done in the impact just described). OR, in shaker testing the other option is to create an excitation signal that forms a response that repeats; if this can be done then the system will get to steady state response and then the Fourier transform will be satisfied and leakage will not be a concern and a window will not be needed.

In shaker testing many signals will create this situation and are used often in shaker testing. These signals are specialized for modal testing – pseudo-random, random transient, burst random and sine chirp are all signals that were created specifically for this type of modal testing. Figure 6 shows the most commonly used burst random excitation which provides an excitation that starts and ends within one sample interval of the time sample for the FFT and therefore does not need a window because there is no leakage of concern. And providing that the response also starts and ends within the time sample then a window is not needed on the response either. So this excitation has no leakage and therefore no windows are required.

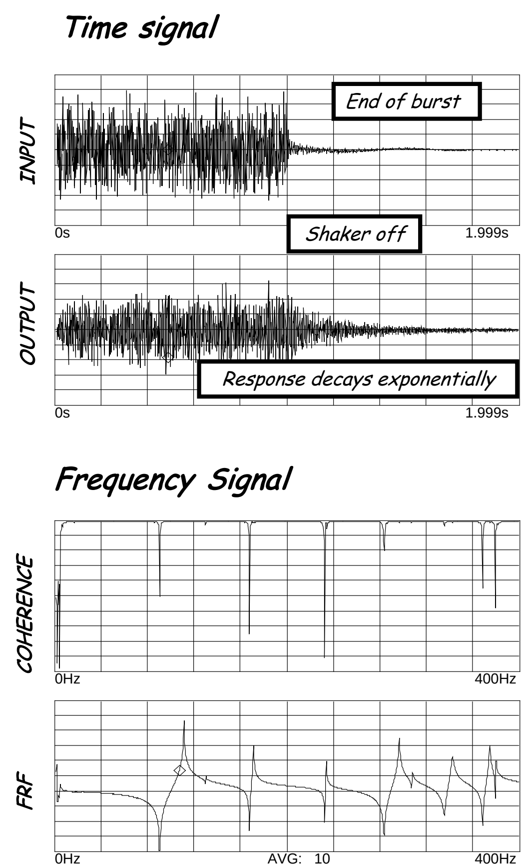

Figure 6: Example of shaker excitation (burst random) which provides a leakage free FRF measurement

Copyright 2014
Pete Avitabile

Number 3 …

Modal impact test set up ritual

So every time I set up to perform an impact test, there is a ritual that I usually go through to make sure that I can make the best possible FRF measurements. There isn't a specific set of steps that I take every time I do this but there are certainly key things that I do every time I make a measurement. Of course, I am talking about taking a measurement on something that I have never tested before or something that is completely new to me. (If it is a structure that I test every day then maybe some of these steps will not be needed because I have apriori information which gives me a good understanding of what is expected).

So when I start a measurement I never take anything for granted and I start with a measurement with a frequency bandwidth which is higher than the frequency range that everyone "believes" is the frequency range of interest. I then use a hammer tip to excite the structure over this range of interest and I always check the input power force spectrum applied to the structure under test. Of course, while I make this first measurement I may need to adjust the voltage level for the hammer input as well as the accelerometer responses. This may need to be done manually unless if your acquisition system has provision to "autorange" all of the response levels. Of course at this point I may need to change the hammer tip to excite the appropriate frequency range of interest and then check to make sure that all the proper response ranges are still appropriate as the different hammer tips are studied.

Once we have a good input excitation then we will start to look at the response and FRF and coherence. But the first thing to do is to look at the response decay to see if the entire response can be captured within one time sample of the measurement. If this is satisfied, then we do not need to apply a window. If it is not satisfied then we may want to consider a longer time window. If this is not possible then we may need to apply a window, which in this case would be an exponentially decaying window.

Once this is done then we would want to take several averages to look at the FRF and coherenece. If this is an acceptable measurement then the next step would be to change the hammer tip to excite a slightly lower frequency range – remember that when I started this process, I selected a higher frequency range than what may have been prescribed for the test. So this is a good opportunity to make sure that the hammer tip is actually exciting the frequency range of interest (because the frequency range is still set for the higher frequency range). But now that less input force is being applied to the structure, then it is important to make sure that all the voltage ranges are still set properly, that the damping window if originally used is still necessary along with other parameters set for the initial set of tests. Once this is all checked then a measurement would be made to assess the FRF and coherence.

Following this then the frequency range of the FFT analyzer could be changed to the lower frequency range associated with the actual softer hammer tip excitation range of the last measurement. And again all the same parameters would need to be checked to make sure that an appropriate level is set a good measurement is obtained.

So for the measurement process I just described you can see that all of the parameters need to be checked each time I change each and every one of the individual items that can change. Remember that I have the ability to change the bandwidth of the measurement, the number of spectral lines, the hammer tip and the use of windows, if needed. All of these need to be considered when making the measurement. And I keep changing all these parameters until I am happy with the measurement that has been made. At this point I would start to collect sets of measurements for the experimental modal test.

Number 2 …

Ui times Uj

Now this is probably the biggest item to consider. So what does this mean. Well let's write an equation down to explain what this means. The FRF can be written in terms of residues or in terms of mode shapes (and has been used in many different articles in this series) as seen in Figure 7.

$$\mathbf{h}_{ij}(j\omega) = \frac{\mathbf{q}_1\mathbf{u}_{i1}\mathbf{u}_{j1}}{(j\omega - \mathbf{p}_1)} + \frac{\mathbf{q}_1^*\mathbf{u}_{i1}^*\mathbf{u}_{j1}^*}{(j\omega - \mathbf{p}_1^*)} + \frac{\mathbf{q}_2\mathbf{u}_{i2}\mathbf{u}_{j2}}{(j\omega - \mathbf{p}_2)} + \frac{\mathbf{q}_2^*\mathbf{u}_{i2}^*\mathbf{u}_{j2}^*}{(j\omega - \mathbf{p}_2^*)} + \cdots$$

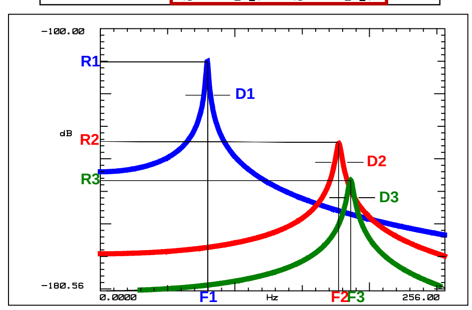

$$\mathbf{h}_{ij}(j\omega) = \frac{\mathbf{a}_{ij1}}{(j\omega - \mathbf{p}_1)} + \frac{\mathbf{a}_{ij1}^*}{(j\omega - \mathbf{p}_1^*)} + \frac{\mathbf{a}_{ij2}}{(j\omega - \mathbf{p}_2)} + \frac{\mathbf{a}_{ij2}^*}{(j\omega - \mathbf{p}_2^*)} + \cdots$$

Figure 7: Frequency Response Function written on a mode by mode basis using the residue formulation and the mode shape formulation

The lower equation is the common way that it is normally written in most of the literature. This is useful but only if you really understand what a residue is. The upper equation is actualy the same equation but with the residues expressed in terms of mode shape information. Specifically the residue (which is directly related to the amplitude of the frequency response measurement) is related to the value of the mode shape at the input excitation location times the value of the mode shape at the output response location for a particular mode of interest and will determine the amplitude of the frequency response function for tht particular mode; and of course the effects of all the modes are the linear summation of all the modes of the system.

So what does this tell me? Basically it gives a very clear definition of the peak amplitude of the FRF is related to the values of the mode shape for a particular mode at the input-output location.

Often times people will ask me why the amplitude of a particular mode is very low for a particular measurement. Well…this equation tells me that for that particular mode either the input excitation or output response (or both) is a very small value and probably close to the node of a mode. If you want to see that mode with a more pronounced peak in the FRF then you really need to change the input and/or output location to be at a place where the mode shape values are much larger and away from the node points.

And actually if you want to conduct a test and select good locations for measurements, then you really need to look to see where the mode shapes are large for each of the modes of the system. The finite element model is a very good tool to use to help decide where to place all of your transducers. While the model may not be perfect, certainly it is a reasonable representation of your structure under test.

And I think if you look at a good number of all the articles in this series, you will find that this is a theme for many of the articles written. Firmly understanding this principle will be a great asset to your understanding of many questions that arise in the conduct of an experimental modal test.

Actually the students in the lab have a list of their top ten things from their perspective of what I always say…

The Rules of Modal
(at least from a modal student's perspective)

1. The Big Dog has great advice.
2. $u_i \bullet u_j$
3. $u_i \bullet u_j$
4. See rules 2 and 3.
5. Don't ask the Big Dog a question unless you want more project work.
6. Don't ask Big Dog a question if you plan to leave within 30 minutes.
7. In one breath, you must be able to say, "The magnitude of a complex number is the square root of the sum of the squares of the real and imaginary parts."
8. In the same breath, you must also be able to say, "No window is needed provided that it meets the periodicity requirement of the Fourier transform process."
9. Document everything.
10. And then document that.

And you can see that #2 and #3 and the follow up #4 pretty much say that this is one of the critical rules of modal and that it is likely the answer to many of your questions.

Number 1…
Thinking is not optional

OK so now let's talk about the "Numero Uno" item. And that is to realize that you really need to think about what you are doing all the time when you are doing testing or analysis. None of this is mundane and thinking is required. This is not like you are working at Burger King where everything is all so very clearly defined. Burger, fries, coke…push the button and the price is determined without you needing to think at all.

Once you stop thinking and just blindly follow a set of rules then you are likely to fall into the hand of the Modal Monster and your results may not be useful if you have encountered any problems that really required your attention and some thinking to realize what may have happened with your measurement.

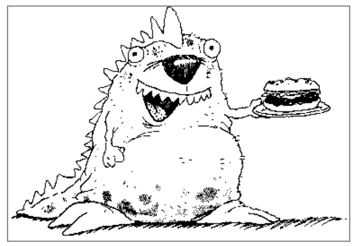

Don't let the Modal Monster rule you – understand what you are doing – think always – question assumptions – be vigilant when you are making measurements and conducting modal tests. For sure go back and read all the articles. There are many important issues that may help answer some of your questions and concerns.

I hope that this last bit of advice helps many of you. If you have any other questions about modal analysis, just ask me.

Author Commentary: This will be the last article that will be published by SEM in Experimental Techniques; this series has existed for 17 years. I hope that the information has been useful to all those people working in the analytical and experimental modal analysis area.

While the series will end in Experimental Techniques, the Modal Space articles will continue to be published on the web at the Structural Dynamics and Acoustic Systems webpage http://sdasl.uml.edu and more specifically at http://sdasl.uml.edu/umlspace/mspace.html

Thank you for all the comments, questions, emails, and support over this duration of this series in Experimental Techniques.
I will continue to publish more information as time goes on until all your questions have been answered.