



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.

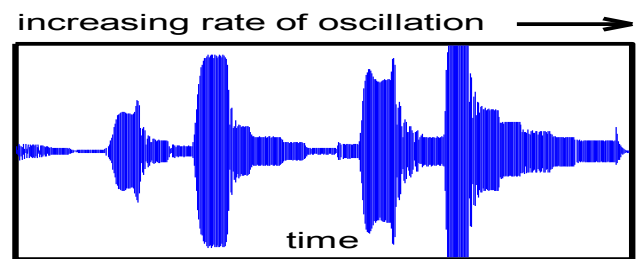


Fig 2 - Simple Plate Response

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 its 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 mass 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?



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 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 loses 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 possible 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 6c - Phase




Fig 6d - Real




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 8b – 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 8c – 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 contains 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.

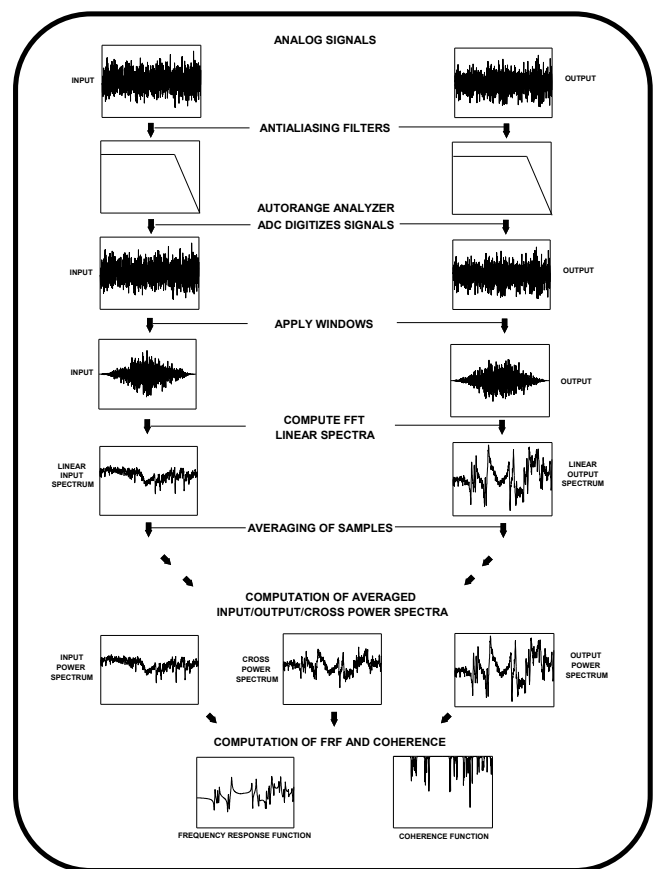


Fig 11 - Anatomy of an FFT Analyzer

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.



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 in 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.



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.

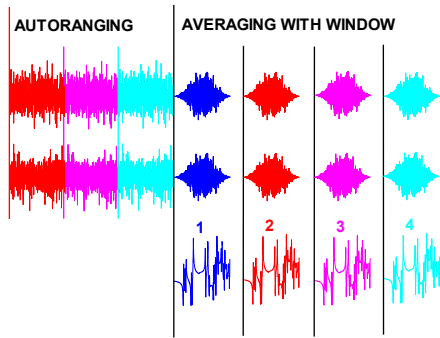


Fig 14 - Shaker Testing – Excitation Considerations
Random Excitation with Hanning window

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.

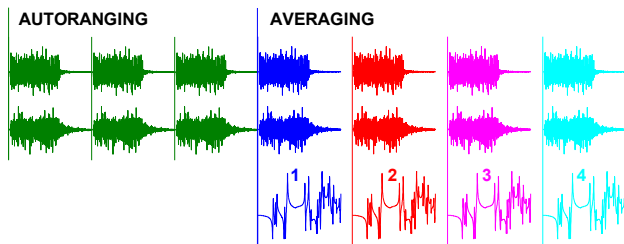


Fig 15 - Burst Random Excitation Without a Window

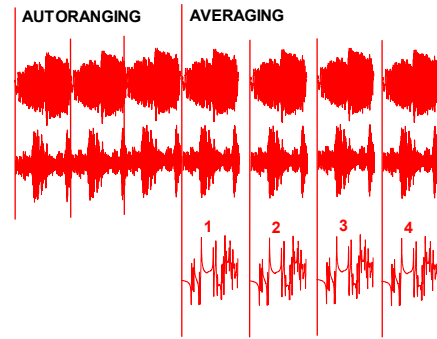



Fig 16 - Sine Chirp Excitation Without a Window

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 windows 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 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

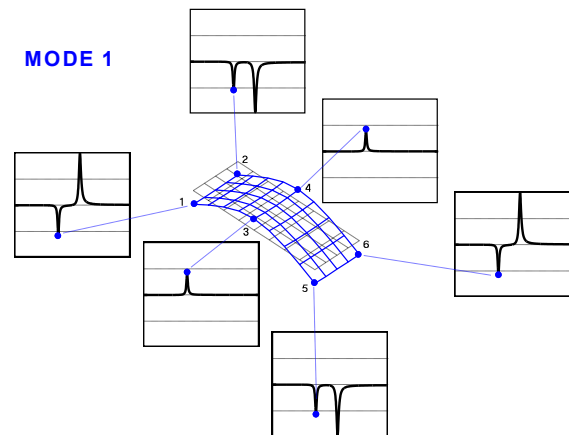


Fig 18 - Plate Mode Shapes for Mode 1- Peak Pick of FRF

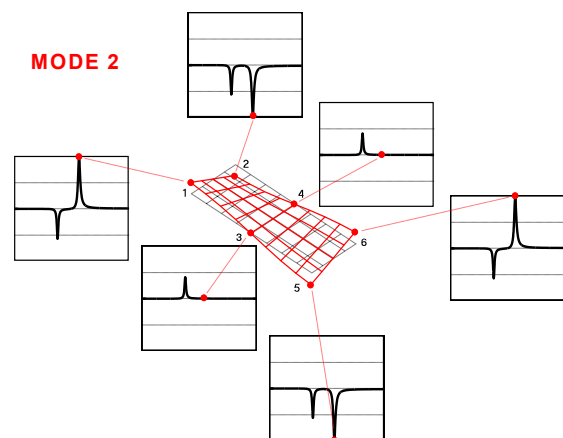


Fig 19 - Plate Mode Shapes for Mode 2 - Peak Pick of FRF

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.

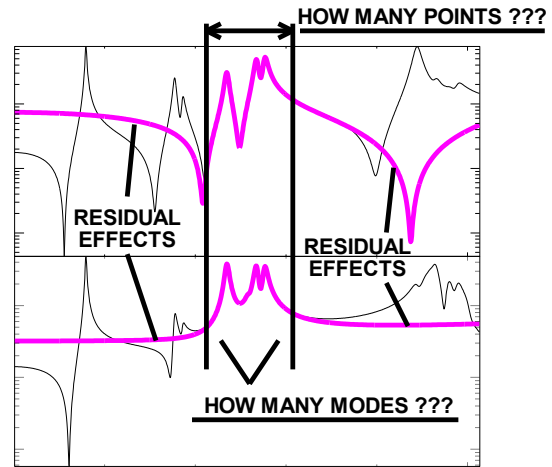



Fig 22 - Curvefitting a Typical FRF

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.



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.

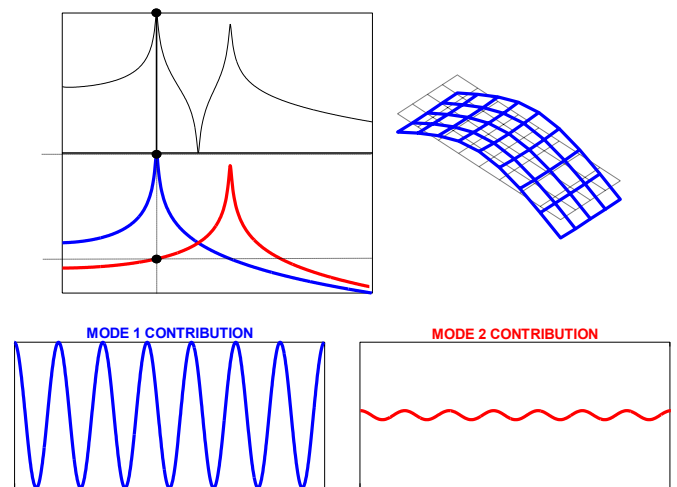


Fig 25 - Excitation Close to Mode 1

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. Figure 27 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.

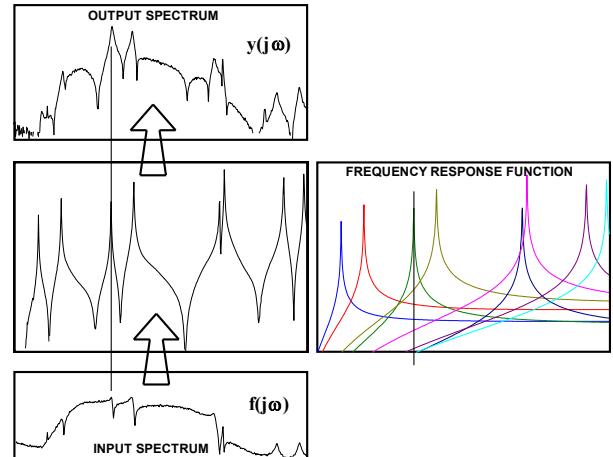


Fig 28 - Broadband Plate Excitation

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.



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.



So should I collect modal data or operating data?

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.

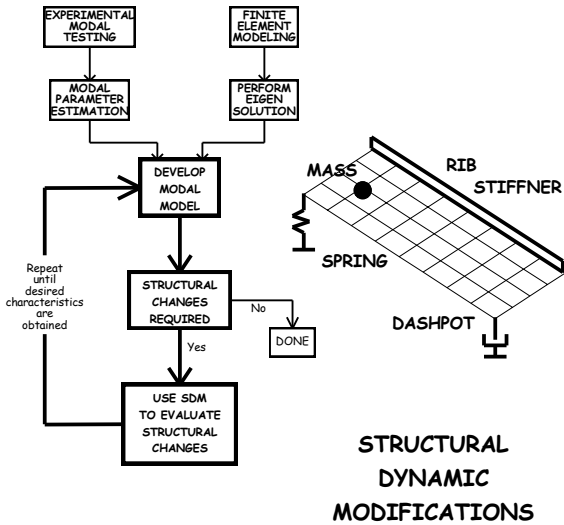


Fig 29 – Schematic of the SDM Process

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.

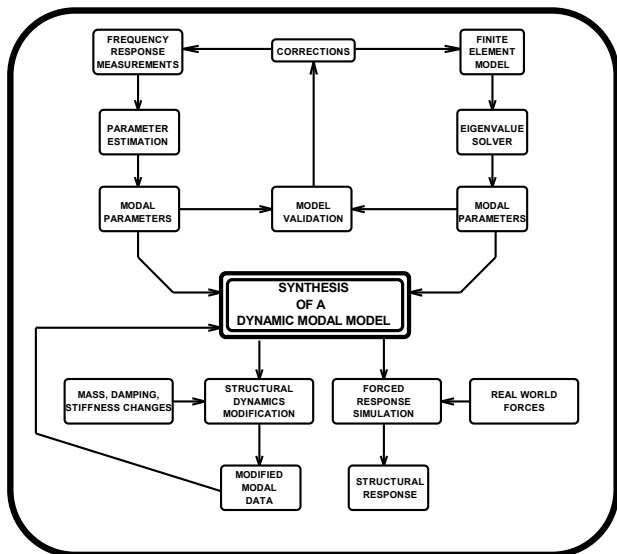


Fig 30 – Overall Dynamic Modeling Process

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

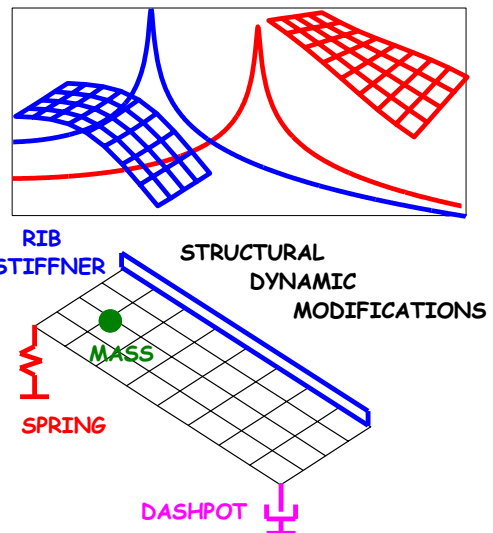


Fig 31 - Modal Model Characteristics

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

February 2025 - Some additional editorial notes by the author.
This document was originally written in 2000. So much has evolved in just 25 years. But the essence of modal analysis is still captured in this article as written.

Additional resources and information can be found in the Modal Space article series written for Experimental Techniques (1998-2014)
<https://www.uml.edu/research/sdasl/education/modal-space.aspx>