1

Introduction to Experimental Modal Analysis: A Simple Non-mathematical Presentation

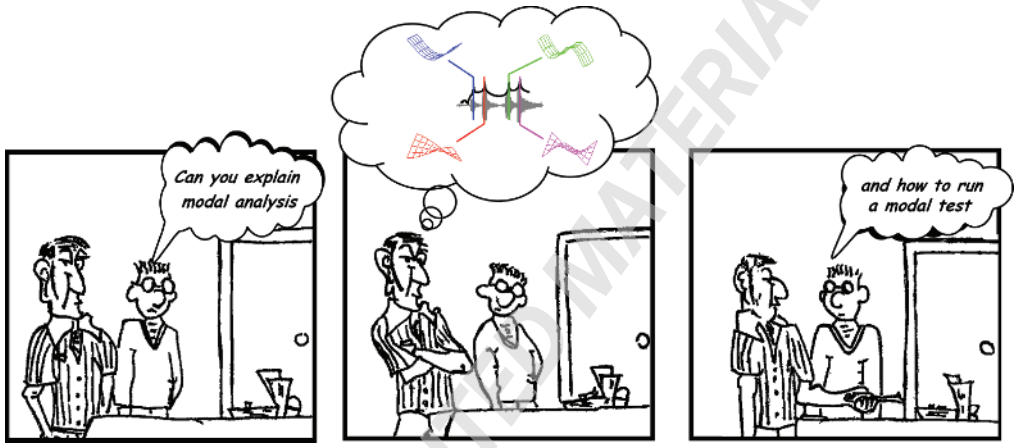

COPYRIGHTED MATERIAL

All structures and systems have operating conditions that cause them to respond due to these excitations. The loadings are generally not just static. Just about any structure is exposed to both static and dynamic loads, and it is the dynamic loads that are of concern to a structural dynamics or vibrations engineer. These excitations cause responses that may not be acceptable for the intended operation of the structure. When this is the case, the engineer must determine what if, anything, can be done to minimize or eliminate the undesirable response in the structure. Sometimes this can be very difficult if the cause of the unwanted response is unknown.

Now structural dynamics is the study of the response of a system to applied loads. These loads can cause responses at different frequencies depending on the dynamic characteristics of the structure. These dynamic characteristics are the frequency, damping and mode shapes. Each of the modes of the structure may contribute in varying degrees to the response of the system and it is sometimes very difficult to understand how the structure responds from the total response of all the modes of the system. So looking at the complete picture may not provide an insight as to how to fix a particular problem. This is where modal analysis comes in.

Modal analysis is the study of the dynamic character of a system that is defined independently from the loads applied to the system and the response of the system. Each of the modes of a system has a certain frequency, with a particular damping, and, most importantly, the characteristic deformation that the structure will undergo given an excitation at its natural frequency. This deformation is related to the mode shape characteristic for the particular mode. Modal analysis, by itself, can only identify the characteristics and not the actual physical deformations. The actual response and physical deformation can only be identified if loads are known

Modal Testing: A Practitioner's Guide, First Edition. Peter Avitabile.
© 2018 John Wiley & Sons Ltd. Published 2018 by John Wiley & Sons Ltd.
Companion Website: www.wiley.com/go/avitabile/modal-testing

and applied to the structure. This is sometimes confusing to many people, but let's put it in perspective with a simpler case.

Let's consider a cantilever beam. Now the beam can be described in terms of its characteristics. These might be the length, width, weight, density, Young's modulus, cross-sectional area, and moment of inertia. But given these characteristics, the deformation of the beam cannot be identified, nor can it be determined if the beam is going to fail in a particular application. This can only be done if the load is known: loads must be identified to determine the deformation, stress, or strain. But once loads are identified, and then the displacement, stress and strain can be determined. But even at this point, the usefulness of the beam for a particular application cannot be determined until the design specification is identified. This specification will identify the relevant design criteria (such as allowable deflection, allowable stress, and allowable strain) and then an engineering judgment can be made as to the suitability of the cantilever beam for the intended use.

Well, modal analysis falls into this same situation. The frequency, damping, and mode shapes are just characteristics of a structure. But whether or not these are good or bad cannot be stated until the intended application is identified, loads identified, and design specification identified. So modal analysis, by itself, is not sufficient to decide if a structure is acceptable or not; the loads and design specification must be identified. (But it is important to point out that in solving many vibration problems, there is sometimes very little understood about the actual loading and often there is no relevant specification available; this is the reality of real-world engineering.)

However, understanding the modal characteristics of a structure can be very useful when performing a structural dynamic analysis. Depending on how the structural dynamic analysis is performed, the underlying modal characteristics may be used for the determination of the response, which helps in gaining an understanding of which modes, how many modes, and to what degree the modes all contribute to the response of the system. Suffice it to say that modal analysis is a very important part of gaining an understanding of a structural dynamic system.

Figure 1.1 shows a computer cabinet, which has responses to a variety of inputs: disk drive inputs, fan inputs and of course any external inputs that "excite" the system. The response comprises the response to all of these individual excitations. The structural dynamics analysis is the study of how the computer cabinet responds to all of these inputs. The time input force shown may be a combination of rotating inputs as well as random inputs. The output time response is due to all those inputs. But the inputs and output responses are not easy to interpret in the time

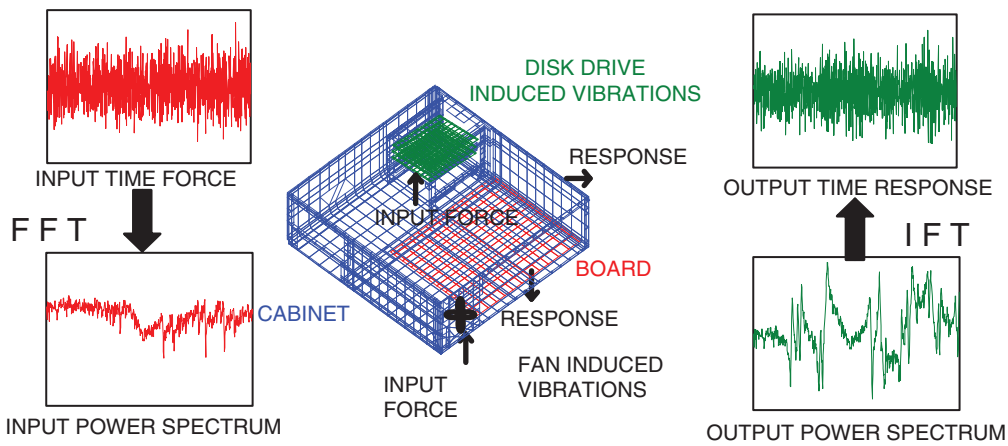

Figure 1.1 Structural dynamics vs modal analysis.

domain. But once they have been transformed to the frequency domain, there is a better picture of the energy distribution in the input force as well as the output response. Clearly, there are some frequencies that seem to have larger responses in the output frequency spectrum. Now if an experimental modal test was performed, those high frequency response peaks are likely associated with the modes of the system. So having modal information helps the designer to understand how the structure might respond to various frequency excitations – be it a discrete or a broadband response.

Now that the input–output scenario has been shown and described, it is useful to discuss a slightly simpler structure that is subject to some input excitation. Figure 1.2 shows a simple plate structure that has a random input excitation. And the output time response is also random in nature. From the time domain signal, there are no hints as to how or why the structure responds in the way it does. However, if the input is transformed to the frequency domain then there is a much clearer picture of the input force excitation. In the frequency domain, the modes of the system (the natural frequencies, damping, and mode shapes) act just like band-pass filters. Each mode "knows" exactly how to amplify and attenuate the input excitation on a frequency basis. And each mode has a separate effect on the input, but all the responses from each filter (each mode) are added together to determine the overall response. This combined response gives hints about where the response is high and generally corresponds to where the modes of the system lie. But in this output response spectrum, all of the modes are not equally excited because the input force spectrum does not have equal energy at all frequencies. So the response is strongly affected by the variation of the input force spectrum. But overall, the modes of the system can be seen as very important indicators as to where the response may be large (if there is significant input at that frequency). So a signal flow diagram provides a very good insight into why the modes of the system are critical pieces of information that need to be clearly understood.

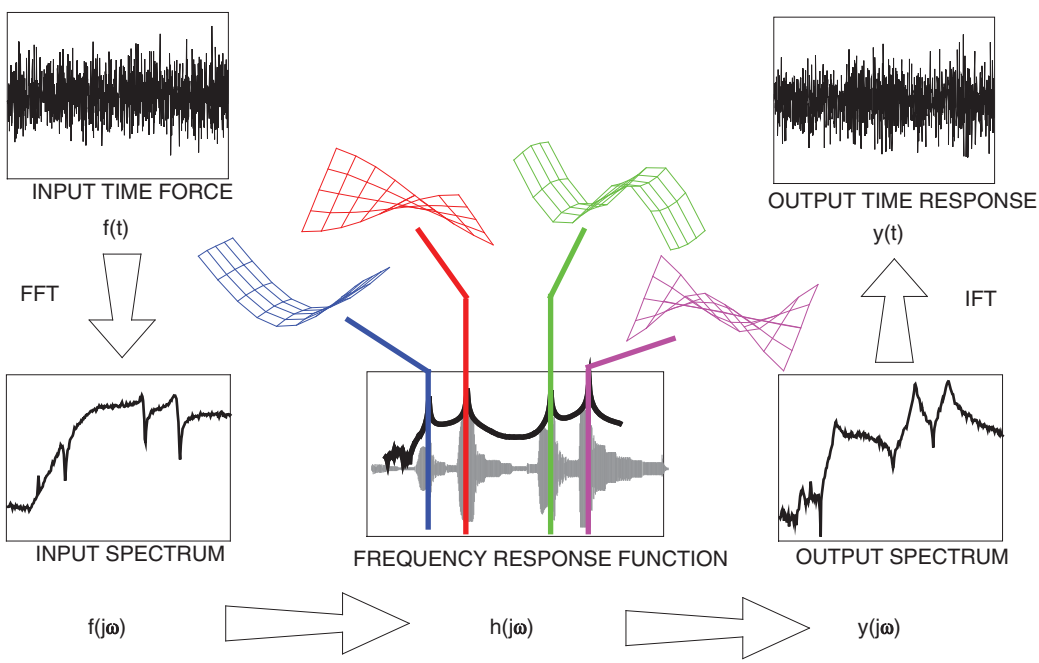

Figure 1.2 Signal flow diagram showing modal filtering of input resulting in output.

So now let's move on to the subject at hand, namely modal analysis. This seems to be an important part of the puzzle: frequencies and mode shapes appear to be central to understanding any structural dynamic problem.

Often times, people ask simple questions regarding modal analysis and how to run a modal test. Mostly, it is impossible to describe the process simply and some of the basic underlying theory needs to be addressed in order to fully explain some of the concepts involved. However, sometimes the theory is just a little too much to handle, although some of the concepts can be described without a rigorous mathematical treatment. This chapter will attempt to explain some concepts about how structures vibrate and to introduce aspects of modal analysis that are used to solve structural dynamic problems. The intent is to simply explain how structures vibrate from a non-mathematical perspective. This chapter serves to introduce some very basic material, which is then expanded on in later chapters.

With that being said, let's start with the first question that is usually asked in regards to experimental modal analysis.

1.1 Could you Explain Modal Analysis to Me?

In a nutshell, we could say that modal analysis is a process whereby we describe a structure in terms of its natural characteristics or "dynamic properties", namely the frequency, damping, and mode shapes. Well that's a mouthful so let's explain what it means. Without getting too technical, modal analysis can be very simply introduced 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. While the structure of a plate is very simple compared to more complicated everyday structures that are evaluated, it can be used to explain the basic underlying theory and concepts very easily.

Let's consider a freely supported flat plate, as shown in Figure 1.3. Let's apply a constant force to one corner of the plate. We usually think of a force in a static sense, with the force causing some static deformation of the plate. But here the force applied varies in a sinusoidal fashion. Let's consider a fixed frequency of oscillation of the constant force: we will change the rate of oscillation but the peak force will always be the same value. We will 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 (see Figure 1.4). There will be increases as well as decreases in amplitude at different points as we sweep up from low frequency to high frequency over time. This seems very odd: 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 because we are applying the same peak force all the time; only the rate of oscillation is changed.

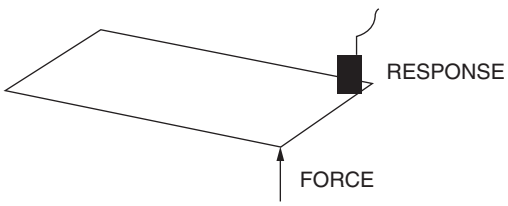

Figure 1.3 Simple plate excitation–response model.

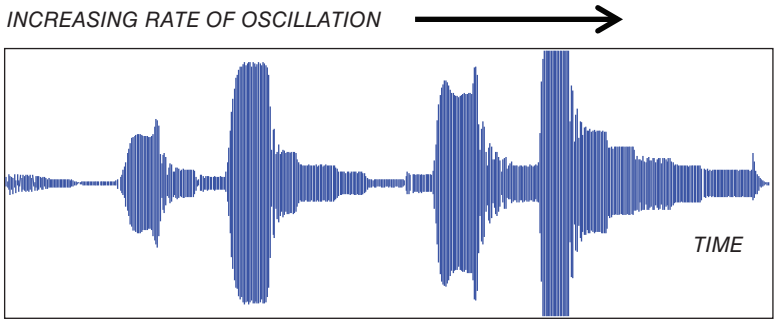

Figure 1.4 Simple plate response due to sinusoidal sweep excitation.

Figure 1.5 Simple plate frequency response function.

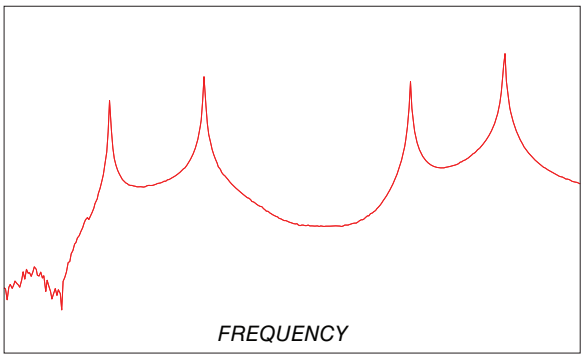

The accelerometer time response data in Figure 1.4 provides very useful information. But if we take the time data and transform it to the frequency domain using the fast Fourier transform (FFT) then we can compute something called the frequency response function (Figure 1.5). Now there are some very interesting aspects to note in this graph. We see that there are four peaks in this function, which occur at the resonant frequencies of the system, and we notice that they occur at the frequencies at which the time response was observed to have its maximum response corresponding to the rate of oscillation of the input excitation.

Now if we overlay the time trace and the frequency trace, what we will notice is that the frequencies at which the time trace reaches its maximum values correspond to the frequencies at which the peaks in the frequency response function reach their maxima (Figure 1.6). So you can see that we can use either the time trace to determine the frequencies 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. And it is important to note that while the sine sweep is very easy to evaluate, a random time signal would not be easy to interpret at all. And it is this frequency response that is widely used in measurements describing responses of structural systems.

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. Identifying and understanding these patterns (or what are called *mode shapes*) is critically important when designing a structure or solving a dynamic response problem. But with only one measurement location, the actual deformation pattern cannot be identified.

Now let's see what happens to the deformation pattern of the plate structure at each one of these natural frequencies. Let's evenly distribute 45 accelerometers on the plate and measure

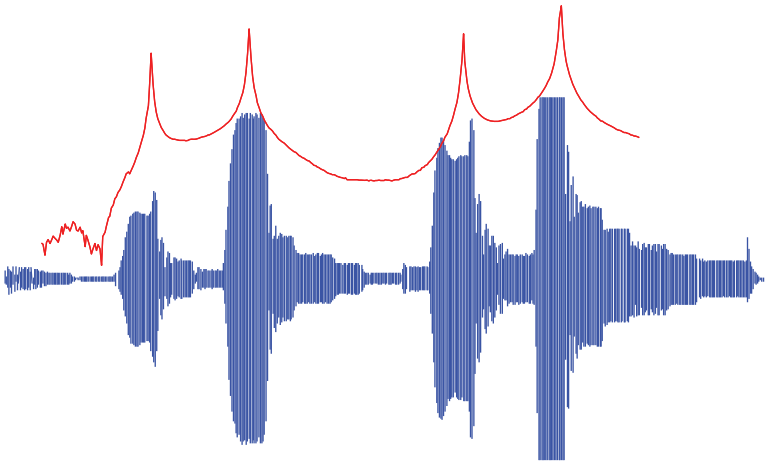

Figure 1.6 Overlay of time and frequency response functions for the simple plate structure.

the response of the plate at different excitation frequencies. If we were to dwell at each one of the four natural frequencies, we would see the deformation patterns that exist in the structure shown in Figure 1.7. The structure will have a very specific deformation pattern depending on the resonant frequency at which we dwell while we measure the response. 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 (mode 1, blue). When we dwell at the second natural frequency, there is a first *twisting* deformation pattern (mode 2, red). At the third and fourth natural frequencies, the second bending and second twisting deformation patterns are seen (mode 3, green and mode 4, magenta, respectively). These deformation patterns are referred to

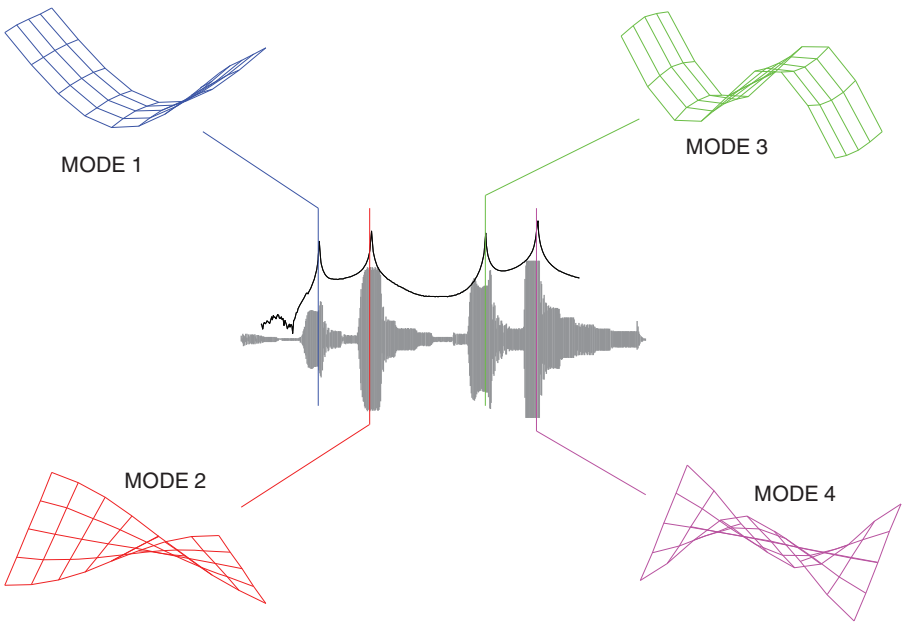

Figure 1.7 Simple plate sine dwell response.

as the *mode shapes* of the structure, although it should be noted that from a pure mathematical standpoint this is not strictly true. However, for the simple discussion here and from a practical standpoint, these deformation patterns are very close to the mode shapes.

Natural frequencies and mode shapes occur in all structures: the mass and stiffness of the structure determines where these natural frequencies and mode shapes will exist. As a design engineer, you need to identify these frequencies and know how they might affect the response of my structure when a force is applied. Understanding the mode shape and how the structure will vibrate when excited helps the design engineer to design better structures. And these frequencies and mode shapes are also critical for test engineers trying to troubleshoot operational problems.

Now there is much more to it all but this is just a very simple explanation of modal analysis. But in order to try to put it into a simpler context, there are two analogies that I commonly use to help people understand what is really needed; see Box 1.1.

Box 1.1 Analogy to help explain modal analysis

Example 1: We all know that there are many, many ingredients that are needed for a multitude of different recipes that we might find in a cookbook. But each recipe might only use a very small subset of ingredients that are needed for each recipe in the cookbook. And each of the identified ingredients is only added in certain proportions for each recipe. Well the frequencies and mode shapes for a structure behave in a very similar manner. If there is a certain loading condition on a structure, then there may be only a particular set of modes that "participate" in the response of the structure. Some modes may participate much more than other modes depending on the particular loading applied to the structure. And under a completely different loading condition, different modes may participate for that particular circumstance. And for the two loading conditions, there may be some common modes that are needed for both loading conditions but all the same modes may not be excited for each of the loading scenarios. Plus there will be differences in the modes that participate especially when the excitation has either low frequency content or high frequency content. So the structure will have many modes that describe its response and there are certain sets of modes that are necessary for the description of the response for each of the different loading scenarios. This is analogous to the cookbook which has many recipes but each recipe has a different set of ingredients for the individual recipes.

Example 2: We know that we may have a large 100-piece orchestra that is needed to play a variety of different scores. Each of the scores requires a different set of instruments that will participate at different times during the score with different intensities. And some scores do not necessarily use all of the instruments in the orchestra for each score. So each individual instrument will participate to a different degree for each of the particular scores to be played. And each instrument has a useful frequency band that will be heard from each instrument. Again these instruments are very much like all of the modes that make up the dynamic characteristic of a structure; each mode has a particular frequency slice and will participate to varying degrees and intensities depending on the particular loading that is applied. One additional item to note is that when we listen to a well-tuned orchestra, the music will be very good and easy to listen to. But all it takes is one member of the orchestra to play out of tune and then the entire score may not sound very good at all. If we just look at the orchestra as a whole unit, it may be very difficult to identify what is wrong and how to fix it. But if we look at each one of the members of the orchestra separately then it is easier to identify where the problem may be and then easy to correct. Well it turns out that structural response is very similar. If the response is acceptable then there is no need to look at all the contributors. But if the structural response is poor then it is very hard to determine how to fix the problem unless we can determine what each mode is and how it contributes. This is where modal analysis has its greatest strength and contribution to design. The structure can be evaluated in terms of its frequencies and mode shapes to identify how each mode behaves and how it contributes to the overall response of the system.

In essence, modal analysis is the study of the natural characteristics of structures. Understanding both the natural frequencies and mode shapes helps in designing a structural system for noise and vibration applications. We use modal analysis in the design of many types of structures, including automotive structures, aircraft structures, spacecraft, computers, tennis rackets, golf clubs…the list goes on and on. There is more detail in later chapters.

Now in the discussion above, we have introduced this measurement called a frequency response function, but exactly what is a frequency response function or what we commonly call an FRF?

1.2 Just what are these Measurements called FRFs?

In Section 1.1, we introduced the frequency response function, but exactly what *is* this? The frequency response function (FRF) 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 a displacement, a velocity or an acceleration.) Now the measured time data is transformed from the time domain to the frequency domain using an FFT algorithm, which can be found in any signal processing analyzer and in many 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 three measurement locations and three mode shapes (Figure 1.8). There are three possible places forces can be applied, and three possible places where the response can be measured. For this example, we will have as many modes as degrees of freedom (DOFs), but in a typical application we have many more measurements than modes. For this example, there is a total of nine possible complex valued frequency response functions. The frequency response functions are usually described with subscripts to denote the input and output locations: as $h_{out,in}$ (or in normal matrix notation $h_{row,column}$).

Figure 1.8b–e shows the magnitude, phase, real, and imaginary parts of the frequency response function matrix. Note that a complex number is made up of a real and imaginary part that can be easily converted to a magnitude and phase; because the frequency response is a complex number, we can look at any and all of the parts that can describe the frequency response function. More detail on this subject is left for subsequent chapters. Now let's take a look at each of the measurements.

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 at point 3 shown in Figure 1.9a. This measurement is denoted as h_{33}. This is a special measurement, which is 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° as we pass over a resonance and gains 180° as we pass over an anti-resonance
- the peaks in the imaginary part of the frequency response function must all be in the same direction.

We can then move the impact force to point 2 while continuing to measure the response at point 3, and then move the impact force to point 1, still measuring the response at point 3. This gives two more measurements (Figure 1.9b). Notice that all measurements are made

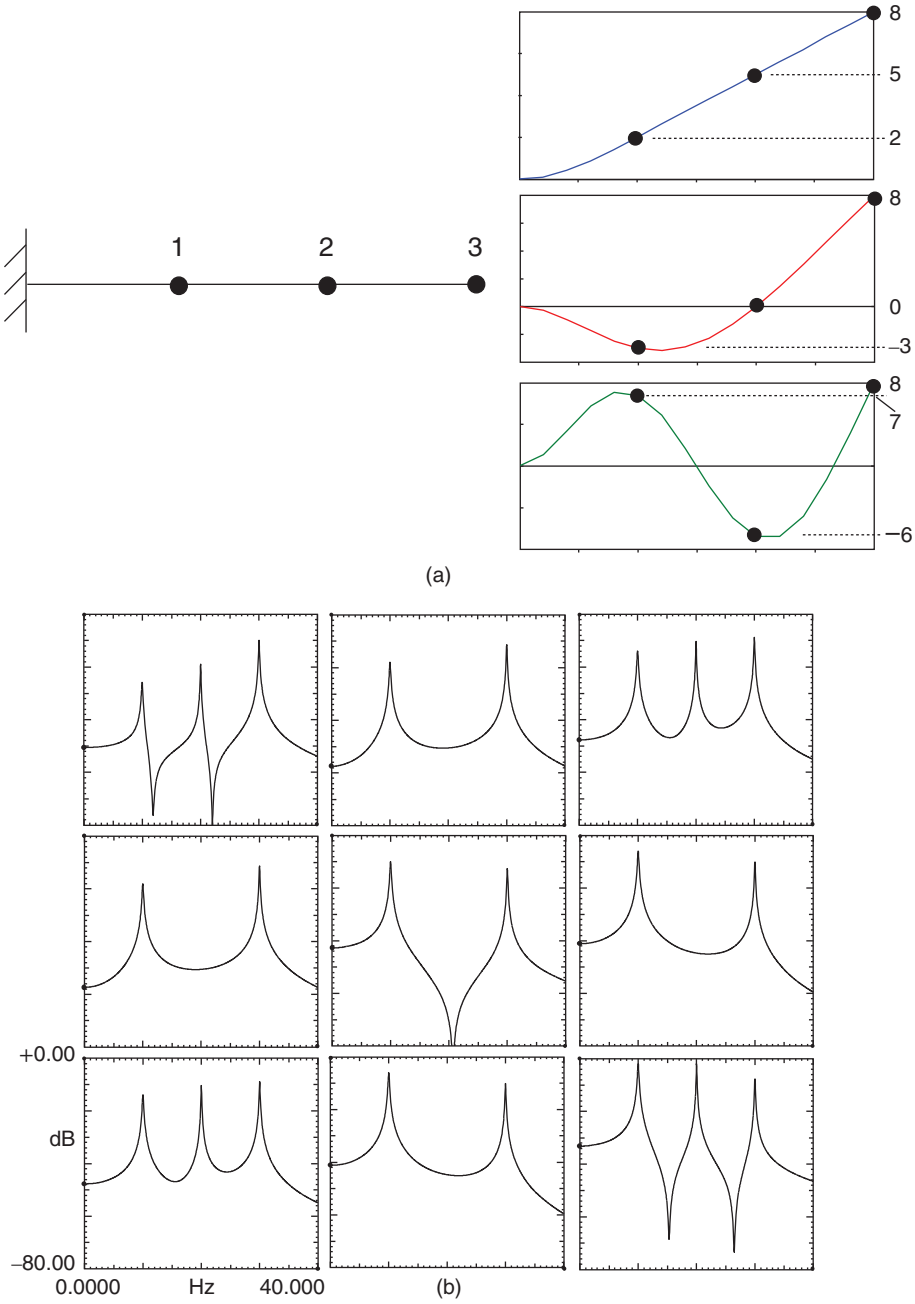

Figure 1.8 3 DOF beam: (a) model for input–output frequency response function matrix and (b) magnitude, (c) phase, (d) real, and (e) imaginary portions of the frequency response matrix.

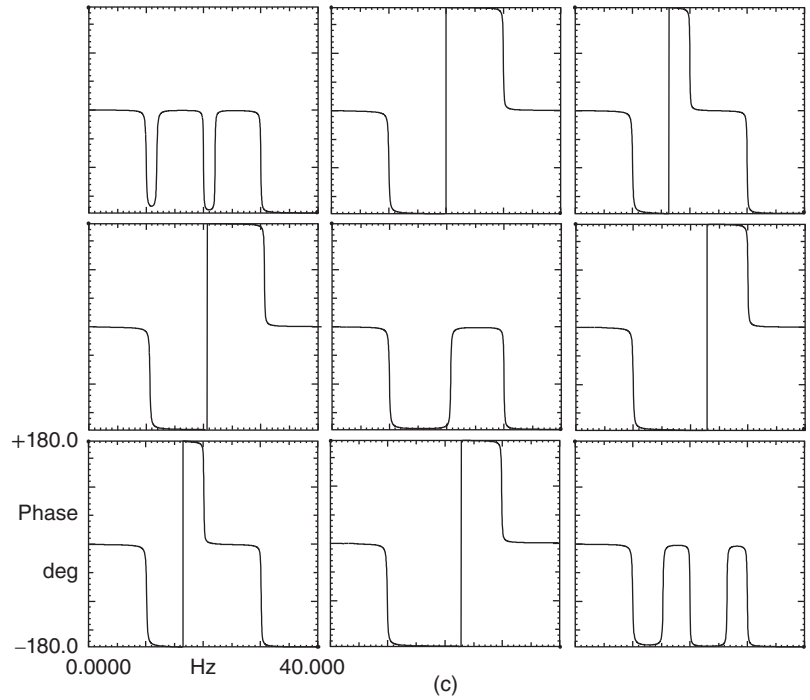

(c)

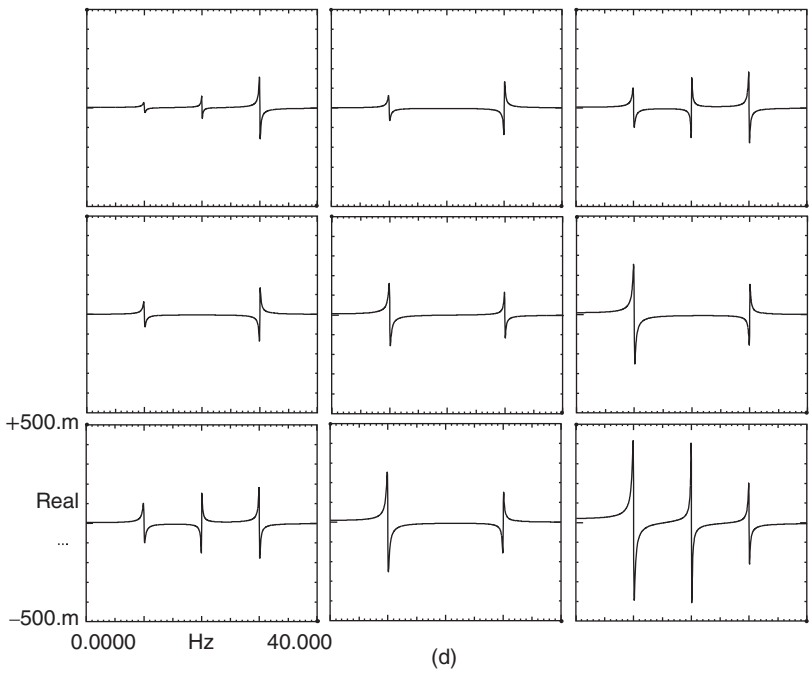

(d)

Figure 1.8 (*Continued*)

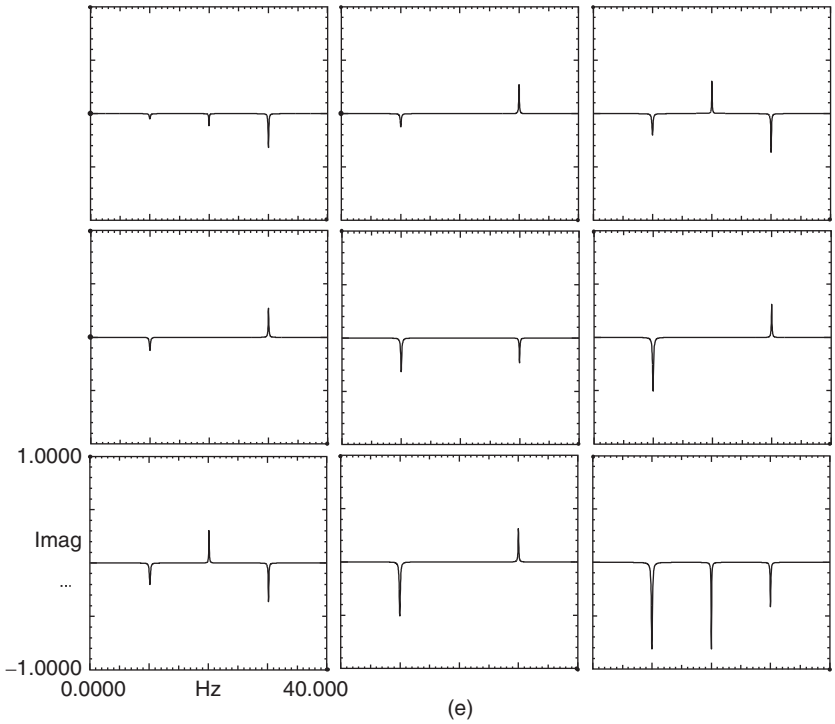

Figure 1.8 (*Continued*)

relative to point 3 and this is commonly called a "reference". Of course, it would be possible to also collect some or all of the remaining input–output combinations. So now we have some idea about the measurements that we *could* possibly acquire. However, in this case, we have taken measurements for just one row of the frequency response function matrix, which is the last row of the matrix of possible terms. It is important to note that the frequency response function matrix is symmetric. This is because the mass, damping, and stiffness matrices that describe the system are symmetric. We can therefore see that $h_{ij} = h_{ji}$; this characteristic is called "reciprocity". It means that we don't need to measure all the terms of the frequency response function matrix; many can be determined using the reciprocity characteristic.

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 it is possible to obtain mode shapes from only one row or column of the frequency response function matrix.

1.2.1 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 measurements in the frequency response function matrix. Without getting mathematical just yet, let's discuss this; the math will come later in the theoretical development of the equations.

Let's just take a look at the third row of the frequency response function matrix and concentrate on the first mode (in blue). Examining the peak amplitude of the imaginary part of the frequency response function for each FRF, the first mode shape for mode 1 is as shown in Figure 1.10a. It therefore seems fairly straightforward to extract the mode shape from measured data. A quick approach is just to measure the peak amplitude of the frequency response

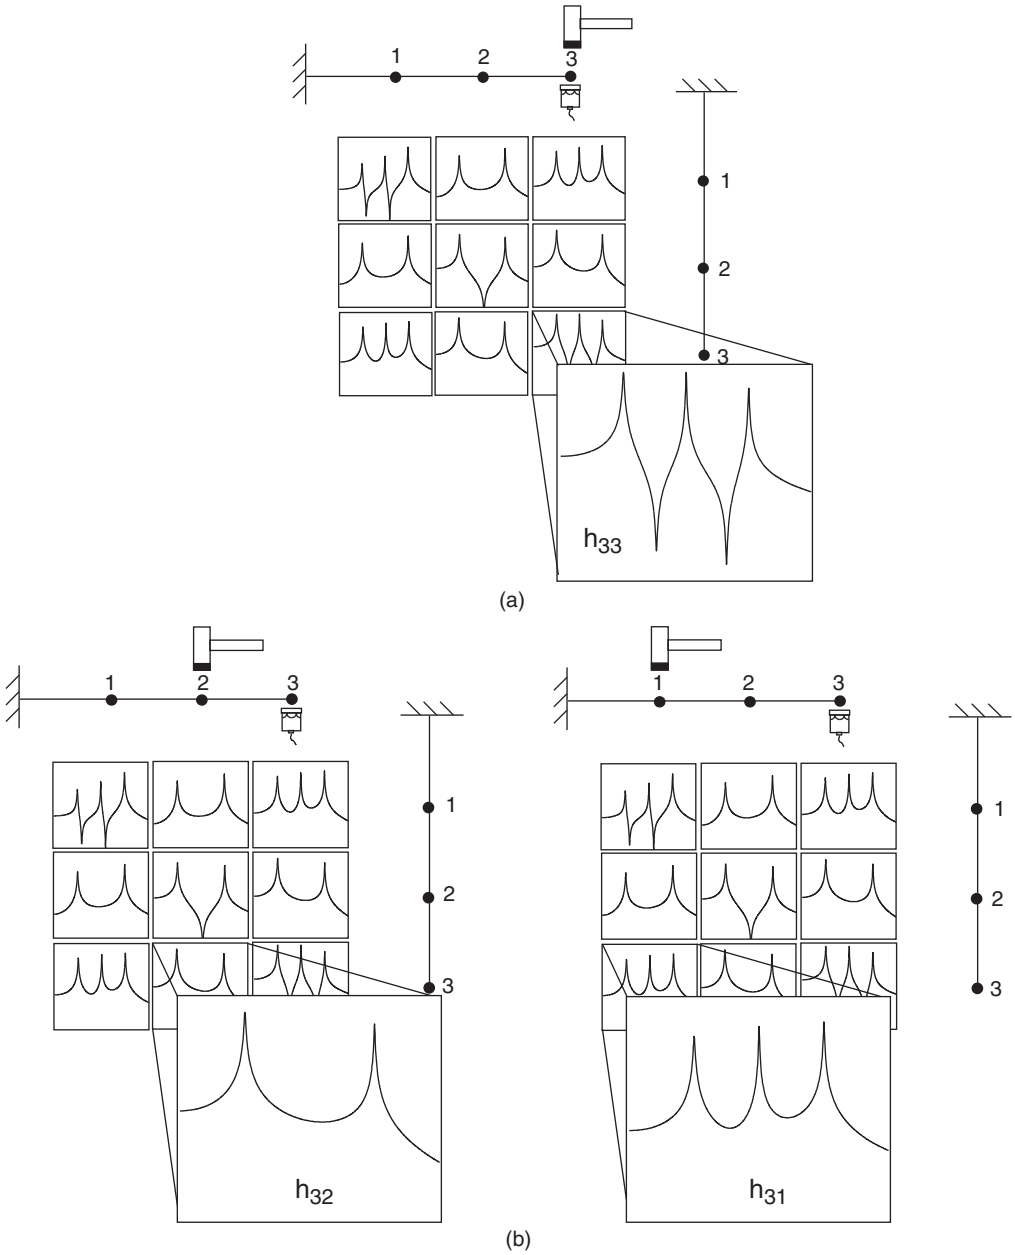

Figure 1.9 3DOF beam: (a) drive point FRF (magnitude) for reference 3; (b) cross FRFs (magnitude) for Reference 3.

function for a number of different measurement points. Clearly, the first bending deformation pattern for the first mode (in blue) is seen from these amplitudes at each of the three points.

So the measurement in Figure 1.10a was taken with the accelerometer stationary at point 3 as we impacted all three points. But if we had the accelerometer positioned at point 2 then we would collect the data in row two of the matrix as we impacted each of the points. Now look at the second row of the frequency response function matrix and concentrate on the first

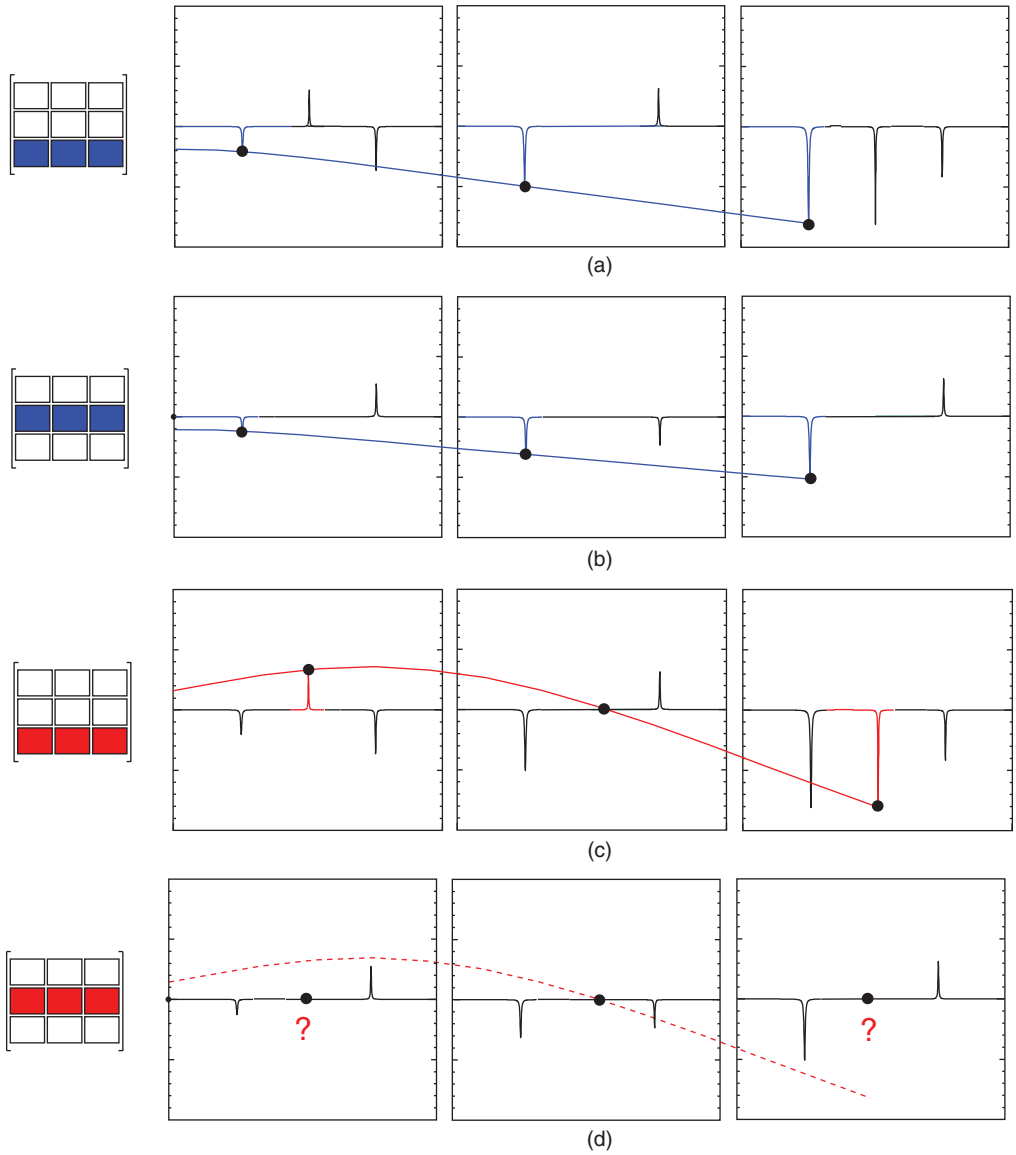

Figure 1.10 3DOF beam: (a) Mode 1 from third row of frequency response matrix; (b) Mode 1 from second row of frequency response matrix; (c) Mode 2 from third row of frequency response matrix; (d) Mode 2 from second row of frequency response matrix.

mode (Figure 1.10b). The first mode shape of mode 1 (in blue) can easily be seen in the peak amplitudes of the imaginary part of the frequency response function; mode 1 can be seen from this row also. Again, the deformation pattern for the first mode (in blue) is clearly seen from these amplitudes at each of the three points. But one important thing to note is that all of the amplitudes are lower from the second row when compared to the third row of the frequency response function 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, in red, (Figure 1.10c). Again if I look at the peak amplitude of the imaginary part of the frequency response function, I can easily see the second bending mode shape for mode 2 (in red) in the peak amplitudes of the imaginary part. Clearly, the deformation pattern for the second mode (in red) is seen from these amplitudes at each of the three points.

Now if I look at the second row of the frequency response function matrix (as we did for mode 1) and concentrate on the second mode, I will be a little surprised because there is no amplitude for the second mode shown in Figure 10d. 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 and the peak amplitude of the imaginary part of the frequency response function is zero. The second mode cannot be observed from this reference location on the beam (this is actually expected).

So this reveals a very important aspect of modal analysis and the experimental measurements we collect during a modal test. 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. Later, in our theoretical analysis, this will be very clear.

Now we have only used three 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. Figure 1.11 shows 15 measured frequency response functions and the three 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. The theory involved is explained in other chapters. In the figure, mode 1 is in blue, mode 2 in red, and mode 3 in green; this color scheme will be used throughout this book.

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?

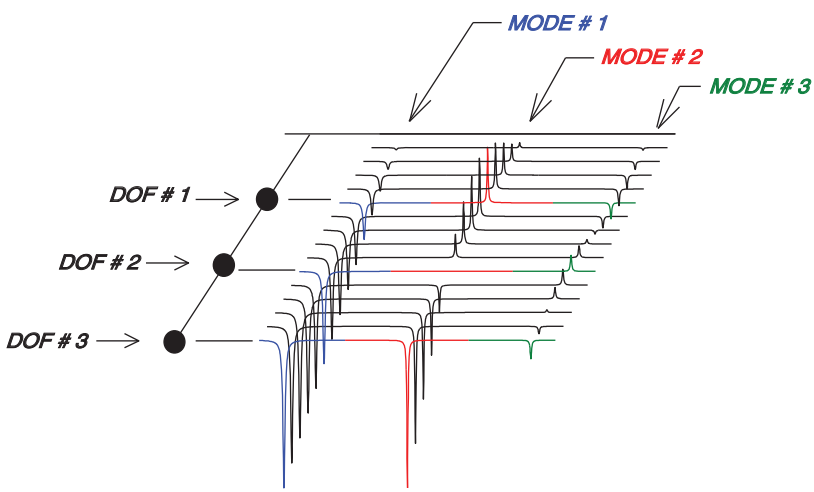

Figure 1.11 Waterfall plot of 3DOF beam frequency response functions.

1.3 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 1.12a and b show the measurements that are obtained from an impact test and a shaker test of the 3DOF beam. A roving

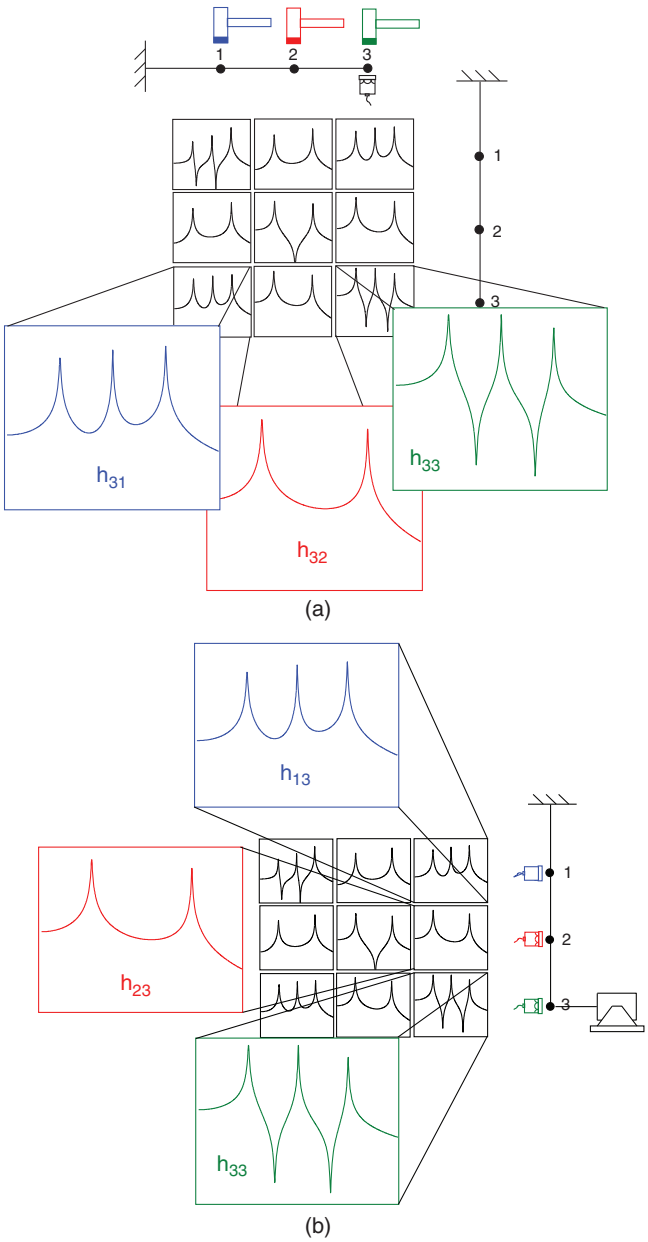

Figure 1.12 Test scenarios for frequency response matrix for the 3DOF beam. (a) roving impact hammer test, (b) shaker test.

impact hammer test generally results in completing 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. Because the matrices describing the system are square symmetric, the reciprocity characteristic is true. It is very important to note that the stationary accelerometer reference will be related to a row of the frequency response matrix whereas a stationary force reference will be related to a column of the frequency response matrix. For the case shown in Figure 1.12, 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. If pure forces can be applied to a structure without any interaction between the applied force and the structure and the response is measured with a massless transducer that has no effect on the structure, then this is true. But what if it is not the case?

Well 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 point 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, and so on. So while theory tells us 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; later in the text there will be examples to illustrate this.

The most obvious difference will occur as a result of 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, particularly 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 lightweight structures. One way to correct for this problem is to mount all of the accelerometers on the structure even though only a few are used at a time. Another way is to add dummy accelerometer masses at locations that are not being measured; this will minimize the roving mass effect.

Another difference is due to shaker/stinger effects. The modes of the structure may be affected by the mass and stiffness effects of the shaker attachment, which may exist even if efforts are made to minimize them. 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. Because 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. The subsequent chapters have more detail and depth on this.

1.3.1 What Measurements do we 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 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 frequency response function. First, there are analog signals that are obtained from the measuring devices. These analog signals must be filtered to ensure 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. These are called anti-aliasing filters. Their function is to remove any high frequency signals that may exist in the signal. Good, high-quality phase-matched anti-aliasing filters contribute to the higher costs of FFT analyzers compared to the much cheaper analog-to-digital converters (ADCs).

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 12, 16 or 24 bit converters; the more bits available, the better the resolution in the digitized signal. Some of the major concerns lie in the sampling and quantization errors that could potentially creep into the digitized approximation. The sampling rate controls the resolution in the time and frequency representation of the signals. Quantization is associated with the accuracy of the magnitude of the captured signal. Both sampling and quantization can cause errors in the measured data but are not nearly as significant and devastating as the worst of all the signal processing errors, namely leakage.

Leakage occurs in the transformation of time data to the frequency domain using the FFT. The FFT process requires that the sampled data consist of a complete representation of the data for all time or that it contain a periodic repetition of the measured data. When this is satisfied, then the FFT produces a true 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 make the sampled data better satisfy the periodicity requirement of the FFT. While the use of windows greatly reduces the leakage effect, it cannot be completely removed. There are many windows that can be used for general signal processing, but there are only a few that are commonly used for experimental modal testing. These windows are discussed in more depth in the signal processing chapter.

Once the data has been sampled, the FFT is computed, so as to form linear spectra of the input excitation and output responses. 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 and the coherence. The coherence function is used as a data quality assessment tool, by determining how much of the output signal is related to the measured input signal. The frequency response function contains information regarding the system frequency and damping and a collection of frequency response functions contains information regarding the mode shape of the system at the measured locations. This is the most important measurement in experimental modal analysis. An overview of these steps is shown in Figure 1.13, which shows the steps in the measurement process from an input to a system to the resulting output from the system. These are the acquisition of the signal, filtering to prevent aliasing, digitization of the signal, windowing (if needed), the FFT to form linear spectra, followed by the averaging of power spectra and the computation of the frequency response function and coherence.

Of course, there are many important aspects to measurement acquisition – averaging techniques to reduce noise and so on – that are discussed more in the signal processing chapters. Any good reference on digital signal processing will also provide assistance in this area. Now the input excitation needs to be discussed next. 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.

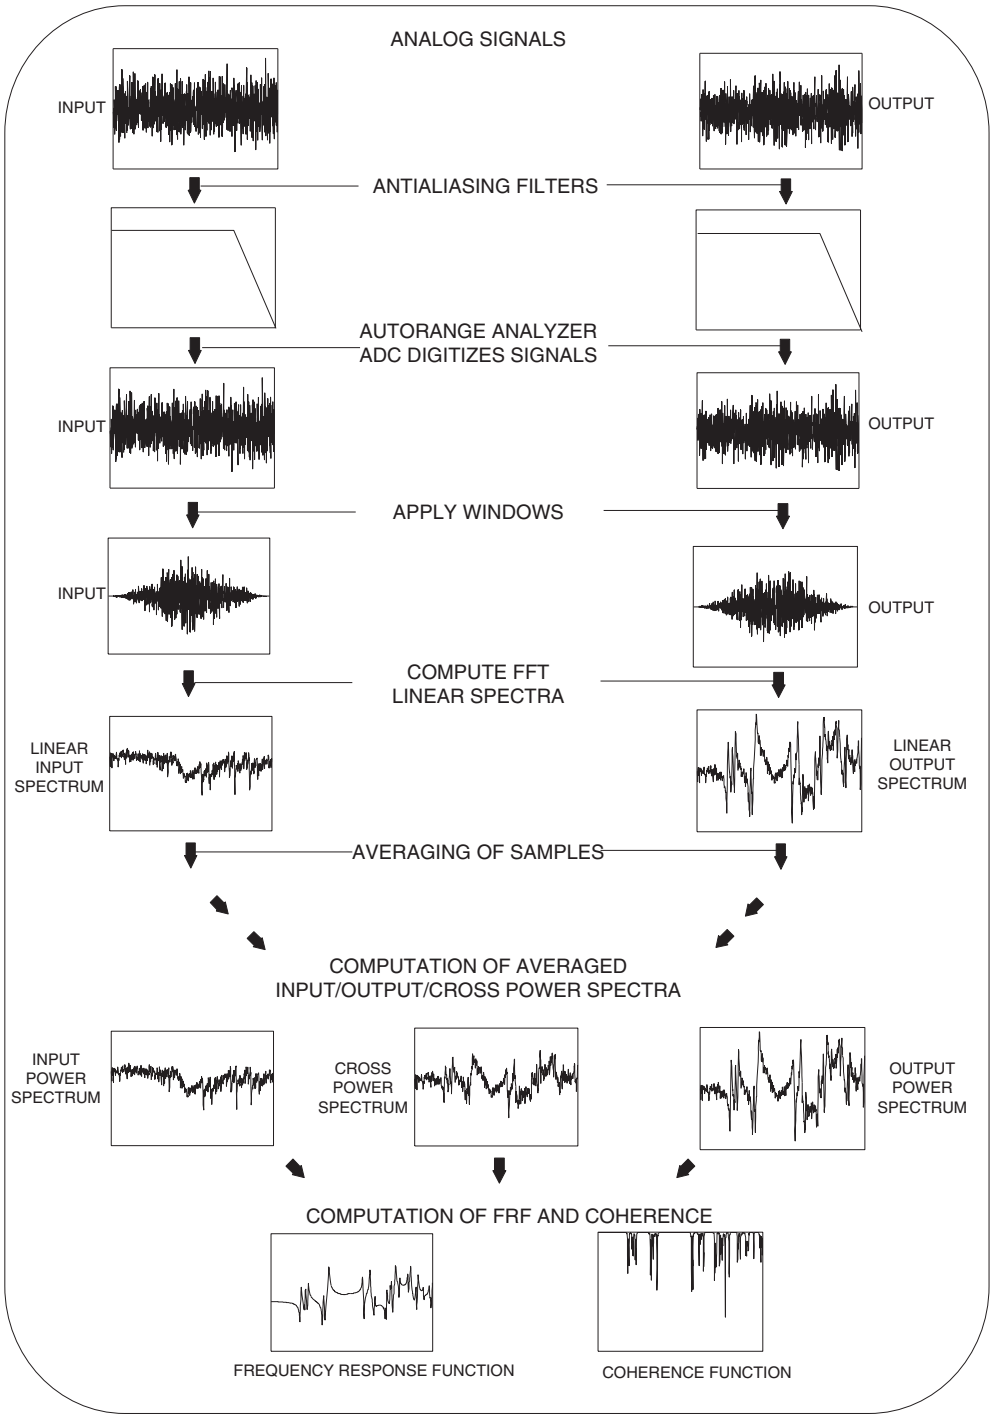

Figure 1.13 Anatomy of an FFT analyzer for typical experimental modal test measurements.

1.4 What's the Most Important Thing 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 of impact testing is presented in the impact chapter and the applications of impact testing are described throughout this book.

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 softer the tip, the narrower the frequency range excited. This is a general statement, but there can be exceptions to this rule and these are discussed in detail later. But basically, 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 not all the modes will be adequately excited, preventing a good measurement from being obtained (Figure 1.14a). The input power spectrum does not excite all of the frequency range, as evidenced by the roll off 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. However, this measurement may not be a problem if we are only interested in the first half of the frequency spectrum shown. Typically, we strive to have a fairly good and relatively flat input excitation forcing function, as seen in Figure 1.14b. The frequency response function is measured much better, as evidenced by the much improved coherence function. When performing impact testing,

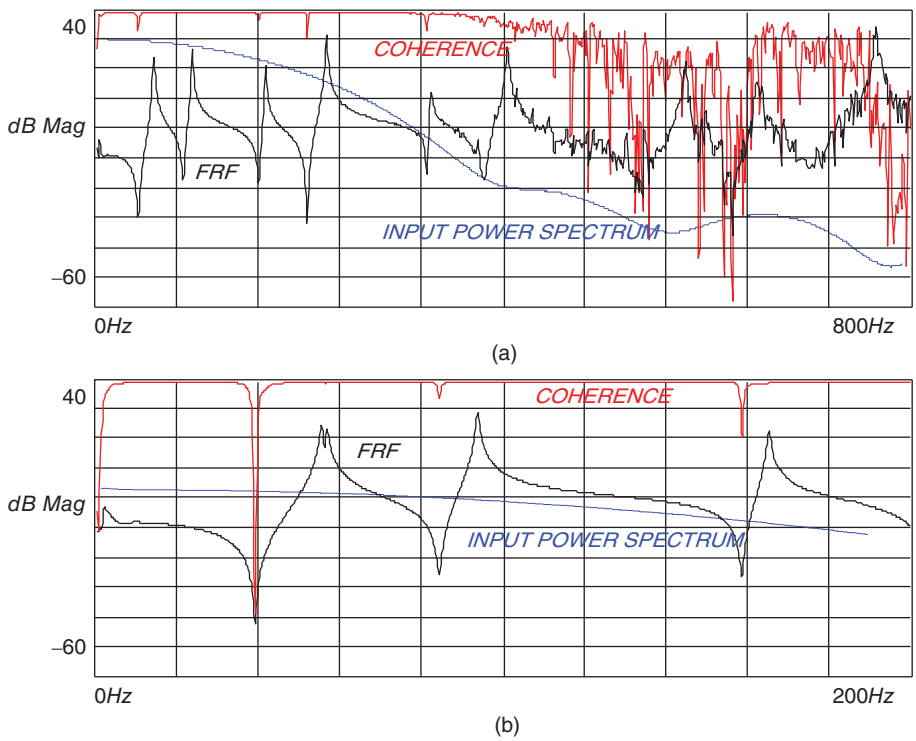

Figure 1.14 Hammer tip choice: (a) hammer tip insufficient to excite all modes; (b) hammer tip adequate to excite all modes.

care must be taken to select a suitable tip, so that all the modes are excited well and a good frequency response measurement is obtained. It typically requires a few iterations in order to achieve acceptable results for a particular measurement scenario.

The second most important aspect of impact testing is the use of a window for the response transducer. Generally, for lightly damped structures, the response 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.

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 FFT 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 this window to minimize leakage is shown in Figure 1.15. The exponential window is one way to minimize leakage, but there are also others, as discussed in the impact excitation chapter and the applications sections.

Windows themselves cause some distortion of the data and should be avoided whenever possible. For impact measurements, two alternative ways 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. Impact testing is covered extensively in the chapter on impact testing and applications.

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

1.5 What's the Most Important Thing to Think about when Shaker Testing?

As with impact testing, there are many important considerations when performing shaker testing, but the most important is the use of excitation signals that minimize the need for windows or eliminate the need for windows altogether. There are many other important considerations, however, a detailed explanation of these is contained in the chapter on shaker testing.

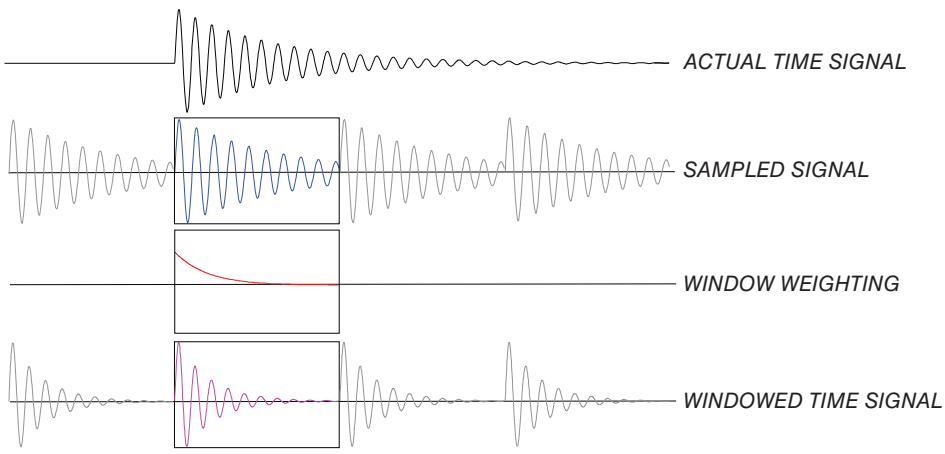

Figure 1.15 Exponentially decaying window to minimize leakage effects.

One of the more common excitation techniques used today is random excitation. This is due to its ease of implementation, but because of the nature of this excitation signal, leakage is a critical concern and a Hanning window is commonly employed. However, the leakage effect when random excitation is used 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 1.16. 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 reduces the distortion of the frequency response function due to leakage, the window will never totally remove these effects; the measurements will still contain some distortion effects due to leakage.

Two very common excitation signals are known as burst random and sine chirp. Both have the special characteristic that they do not require windows to be applied to the data because the signals are inherently leakage free in almost all testing situations. They are relatively simple to employ and are found on most signal analyzers available today. The two signals are shown schematically in Figures 1.17 and 1.18.

In the case of burst random excitation, the periodicity requirement of the FFT process is satisfied because 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. While other excitation signals also exist, these are the most common ones in use in modal testing today. Much more detail on shaker testing and excitation signals is discussed in the chapter related to shaker testing and applications. So now we have a better idea how to make some measurements.

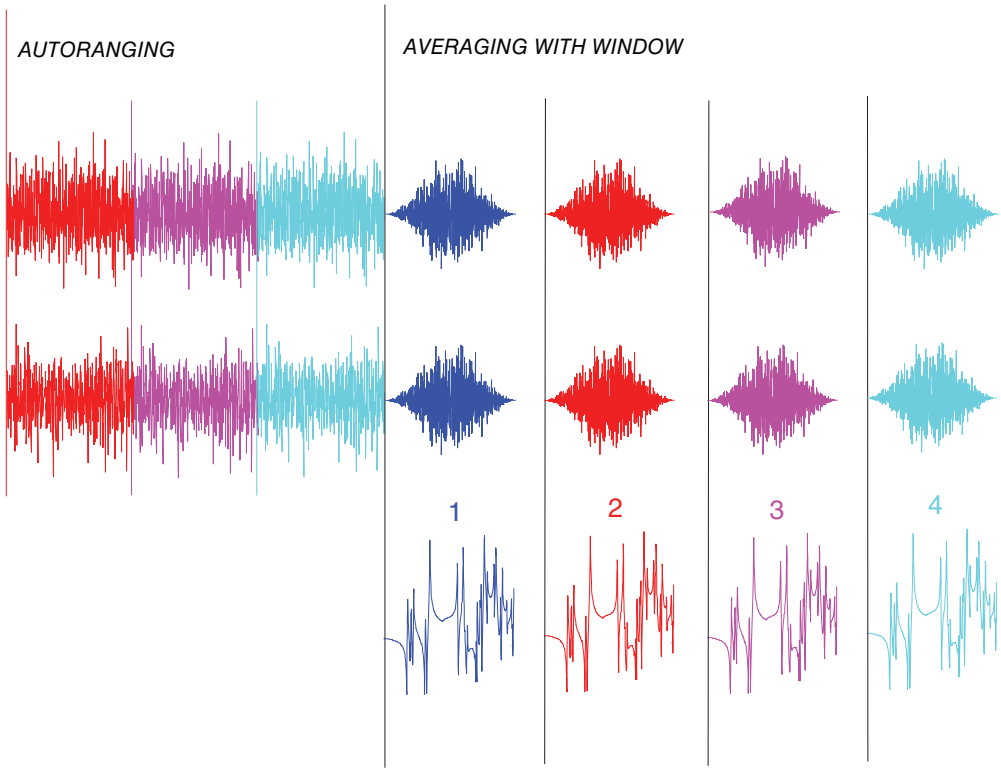

Figure 1.16 Shaker testing: random excitation with Hanning window.

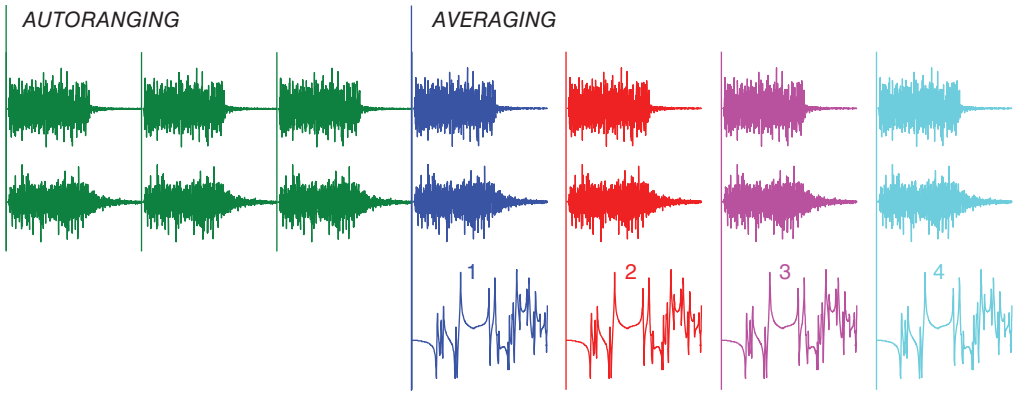

Figure 1.17 Burst random excitation without a window.

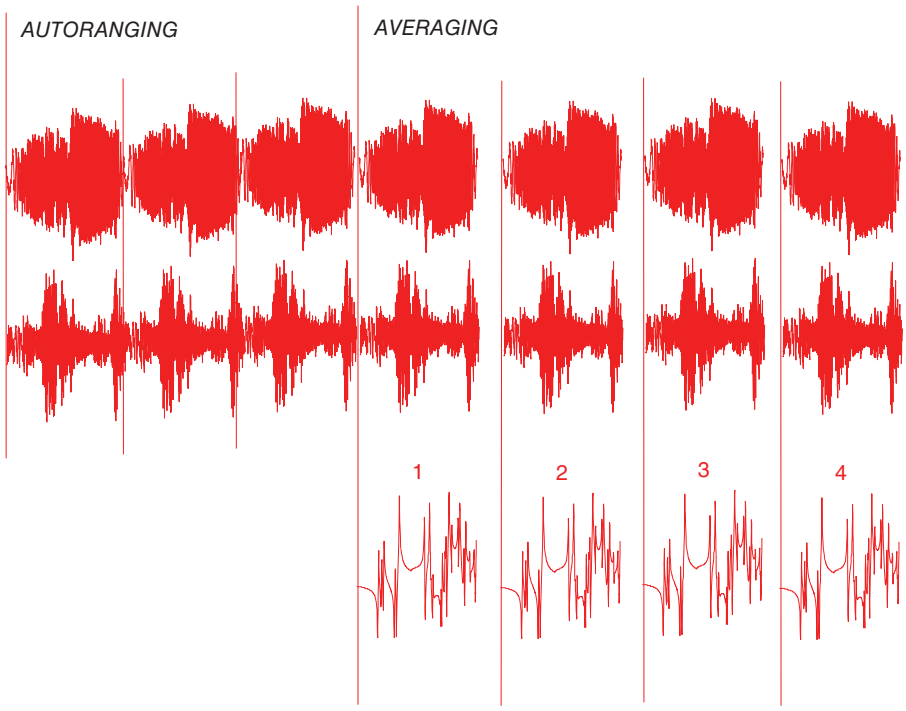

Figure 1.18 Sine chirp excitation without a window.

1.6 Tell me More About Windows; They Seem Pretty Important!

Windows are, in many measurement situations, a necessary evil. While it would be preferable to use no 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 used to provide leakage free measurements and which 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 used? Basically, the most common windows used today are the uniform, Hanning, flat top and force/exponential windows. Rather than detail all these, let's just simply state when each is used for experimental modal testing. More discussion on windows is contained in the signal processing and measurements chapters.

The uniform window (which is also referred to as the rectangular window, boxcar or no window) is a unity-gain weighting function that is applied to all the digitized data points in one sample or record of data. It 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, for example, burst random, sine chirp, pseudo-random, or digital stepped sine excitation signals, all of which generally satisfy the periodicity requirement of the FFT process.

The Hanning window is a cosine shaped weighting function (that is, bell-shaped) that forces the beginning and end of the sample interval to be heavily weighted to zero. It 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 therefore 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 or when there is a constant speed excitation.

The force and exponential windows are typically used when using impact excitation. The force window is a unity-gain window that acts over a portion of the sample interval in which 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 on 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 suggest more damping than really exists. While these errors are not totally desirable, they are far more acceptable than the significant distortion that results from leakage. The chapter on signal processing and measurement definition has much more in depth material related to this.

1.7 So how do we get Mode Shapes from the Plate FRFs?

So now that we have discussed various aspects of acquiring measurements, let's take several measurements on the plate structure discussed at the start of this chapter. Let there be six measurement locations on the plate. Now there are six possible places where forces can be applied to the plate and six possible places where we can measure the responses. This means that there are a total of 36 possible input–output measurements that could be taken. The frequency response function describes how the force applied to the plate causes the plate to respond. If we apply a force to point 1 and measured the response at point 6, then the transfer relation between 1 and 6 describes how the system behaves for this measurement (Figure 1.19).

Now if we were to take six measurements on the plate, then we could go through the same simple peak picking process as we did for the beam in Section 1.2. If we evaluate the first mode and look at the imaginary part of the frequency response function and evaluate the peaks for mode 1 (in blue), as shown in Figure 1.20, then we would see that points 1 and 2 both have −1 amplitude, as do points 5 and 6, but points 3 and 4 have an equal +1 amplitude. From these six points it can be seen that this is really a first bending mode of the plate. Of course, all 45

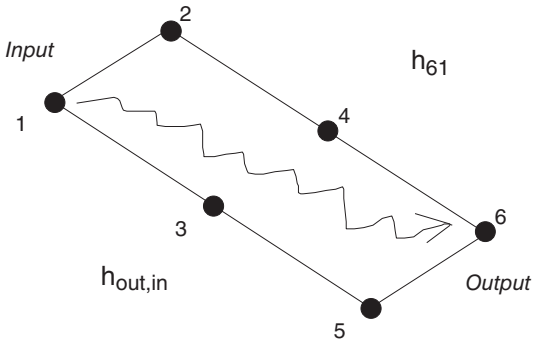

Figure 1.19 Input–output measurement locations.

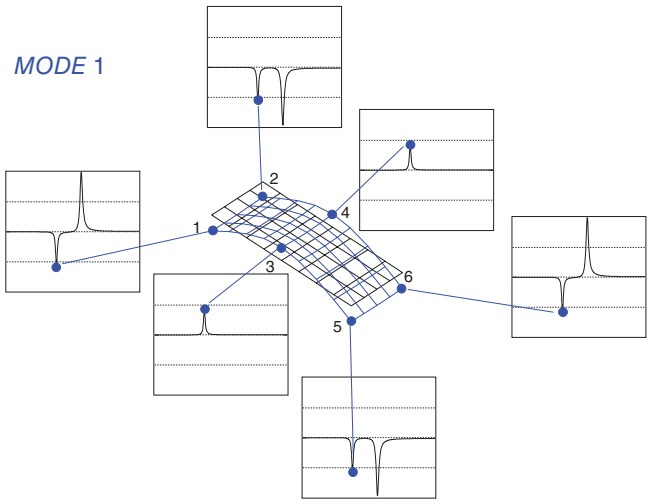

Figure 1.20 Plate mode shapes for Mode 1: peak pick of FRF.

points are shown in this figure, revealing the resulting mode shape for mode 1; showing all the measurements would be cumbersome and would clutter the figure, rendering it useless. But clearly the trend of the measurements would map out the first bending mode of the plate (in blue).

If we then move on to the second mode in red (Figure 1.21), we could also look at the peak amplitude of the imaginary part of the frequency response function. Point 1 is shown to have +2 units of amplitude, and point 2 has −2 units of amplitude. At the other end of the plate, point 5 has −2 units of amplitude and point 6 has +2 units of amplitude. This suggests that this is a torsional mode of the plate, which is confirmed when all 45 points are shown. However, it is very important to note that points 3 and 4 have no amplitude, which implies that these are node points, as indicated by the mode shape.

While the peak pick technique 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, known collectively as the dynamic characteristics. The measured frequency response function is broken down into many single DOF (SDOF) systems, as shown in Figure 1.22. The upper portion of the figure shows the magnitude of a typical frequency response function and the lower portion shows

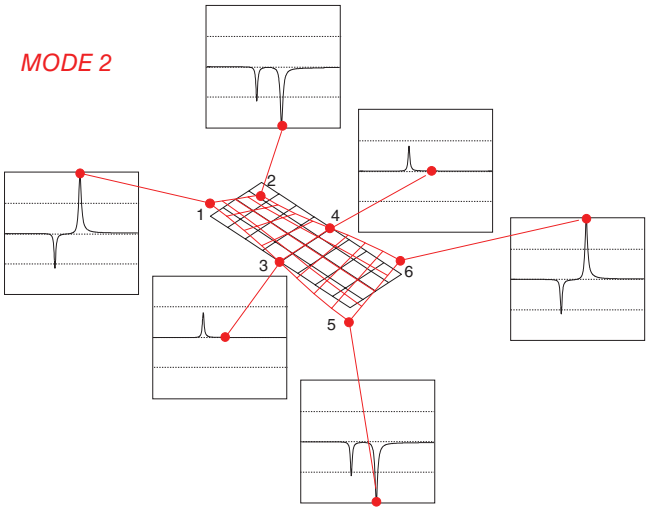

Figure 1.21 Plate mode shapes for Mode 2: peak pick of FRF.

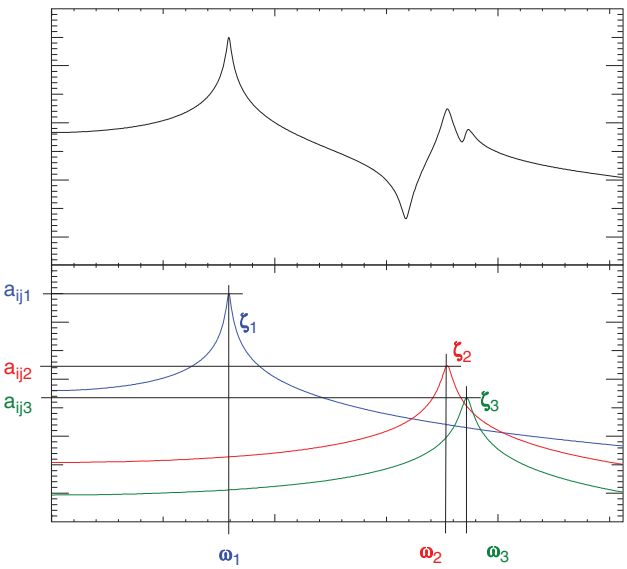

Figure 1.22 Breakdown of a frequency response function.

the magnitude of the frequency response function of each of the individual single degree of freedom (SDOF) systems that form the frequency response function in the upper plot. In fact, the modal parameter estimation process is intended to decompose the measured frequency response function into the basic fundamental or primary parts that make up the measurement. These are the frequency, damping, and mode shape for the structure. The main job for the structural test engineer is to acquire measurements that characterize the overall combined response and to convert or decompose all the data acquired into the primary information that is embedded in it.

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

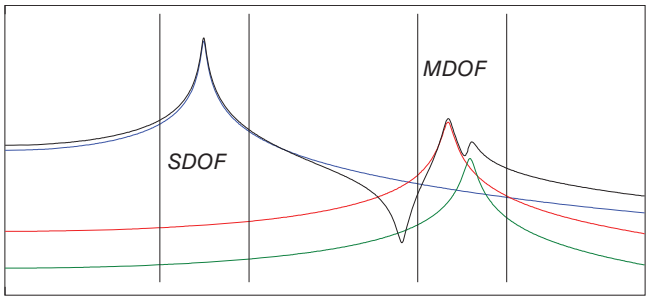

Figure 1.23 Curvefitting different bands using different methods.

common methods employ multiple mode analytical models but at times very simple single mode methods will produce reasonably good results, as shown in Figure 1.23. Again, it is very clear that the individual modes form the entire frequency response function measurement. 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.

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 1.24. While the measurement to obtain the frequency response functions is a critical and sometimes a time consuming process, the extraction of modal parameters can also be tedious, especially when good measurements are not acquired. Generally, frequency response function data is fit over bands of the frequency range, in which several modes may be grouped. The identification of how many modes are

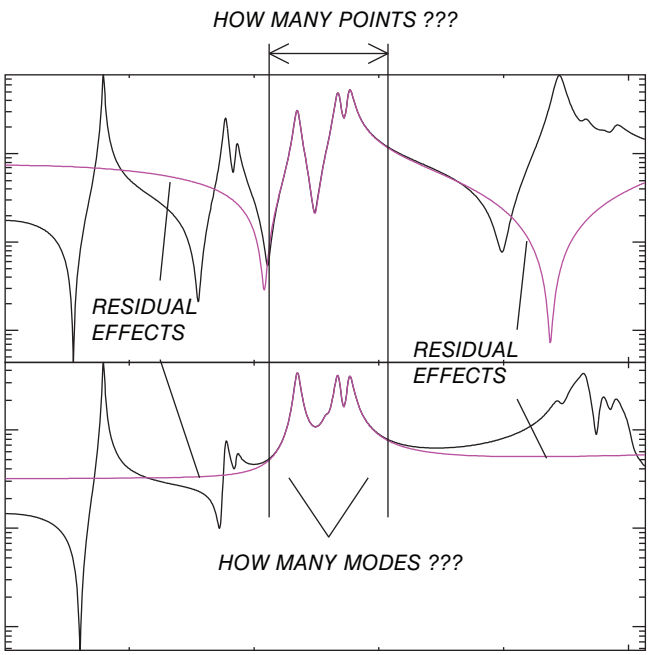

Figure 1.24 Curvefitting a typical FRF.

present in the band can be difficult in some situations, especially when the modes are extremely close together in frequency. Certain tools exist to help with this mode number determination.

Much more needs to be said concerning the estimation of modal parameters from measured data, the tools available for deciphering data, and the validation of the extracted model. A detailed explanation is contained in the chapter related to modal parameter estimation.

1.8 Modal Data and Operating Data

All structures respond to externally applied forces. But often 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.

1.8.1 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 they can be measured). Assume 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 that 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 1.25. 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 that 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.

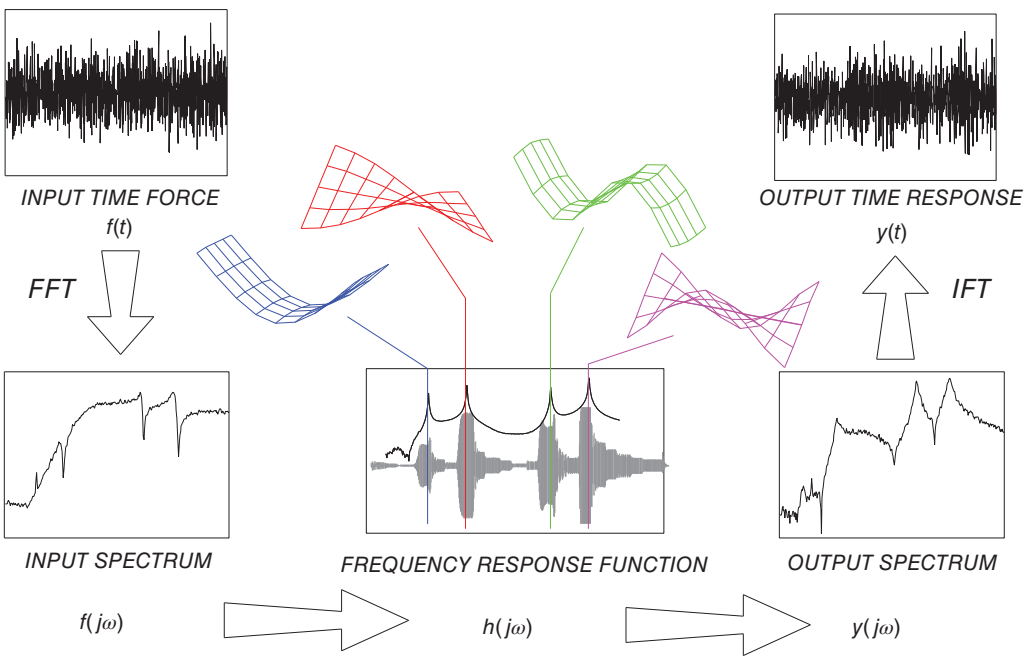

Figure 1.25 Schematic overview of the input–output structural response problem.

The excitation shown is a random excitation and the response is also random in nature, as seen in the time data in Figure 1.25. However, looking at this time data, it is very hard to determine what is dominating the response. Now if the response level was acceptable then there really wouldn't be a need to further evaluate the response. But if the response was unacceptable, then it would be very advantageous to look at what frequency or frequencies were causing the objectionable response. So by taking the FFT of the input and response signals, there is a much clearer picture as to what mode or modes may be causing the objectionable response. And looking at the output power spectrum, there are two peaks that seem to be much more dominant in the overall response. The actual deformation pattern would also be of interest, in order to understand what is happening.

Now what would happen if the excitation did not contain all frequencies but rather only excited one particular frequency (which is common in the case of constant rotating speeds, for instance)? In this case, it is very beneficial to understand the deformation pattern due to the forced response.

To illustrate this, let's use the simple plate again. Let's assume that there is some operating condition that exists for the system; a fixed frequency operating imbalance 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 1.26. 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? It doesn't appear to be anything like any of the mode shapes that we measured before.

In order to understand this, let's take the 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 only two modes are activated by the input. Of course there are more modes, but let's keep it simple to start with. 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 sinusoidal signal that is very close to the first natural frequency of the plate structure. The response of the system for one frequency response function is shown in Figure 1.27. 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 from both mode

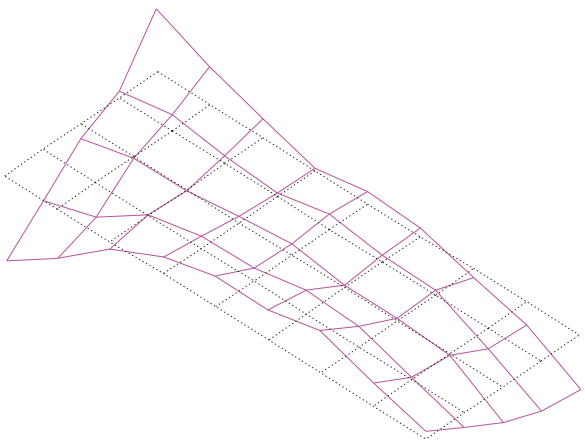

Figure 1.26 Measured operating displacements.

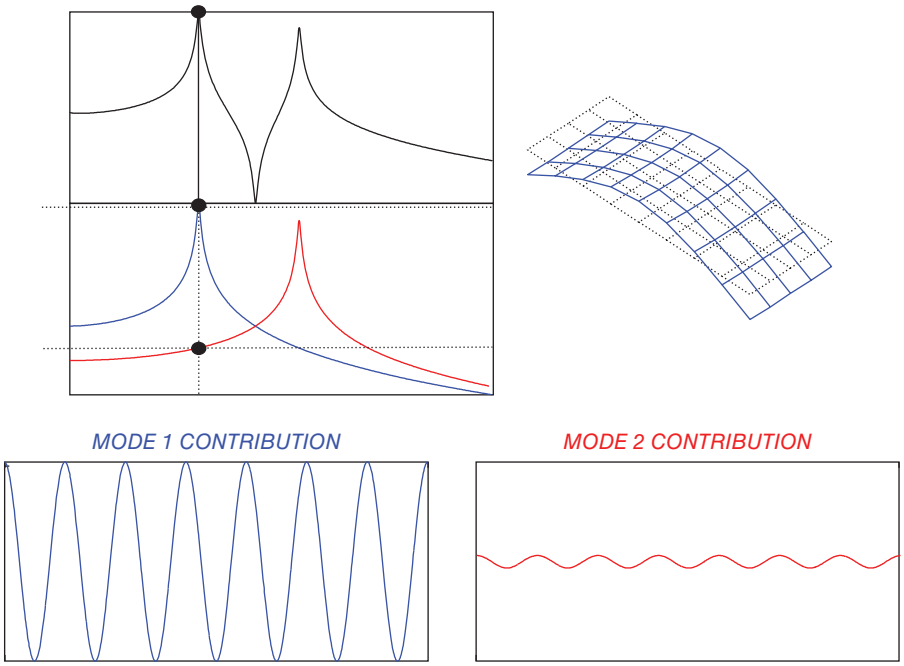

Figure 1.27 Excitation close to mode 1.

1 and mode 2. We can also see that the majority of the response, whether it is 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, with 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; the deformations that are measured are the *actual response* of the structure due to the input excitation, whatever it may be.

When we measure frequency response functions and estimate the modal parameters, we actually determine the contribution to the total frequency response function solely due to the effects of mode 1mode 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. We can now see that the operating deflection pattern will look very much like the first mode shape if the excitation primarily excites mode 1.

Now let's excite the system very close to the second natural frequency. Figure 1.28 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 from mode 1.

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. Here we see the real difference between modal data and operating data. Figure 1.29 shows the deformation shape of the structure.

At first glance, it appears that the deformation doesn't look like anything that we recognize, as displayed in the upper right corner of Figure 1.29. But if we look at the deformation pattern

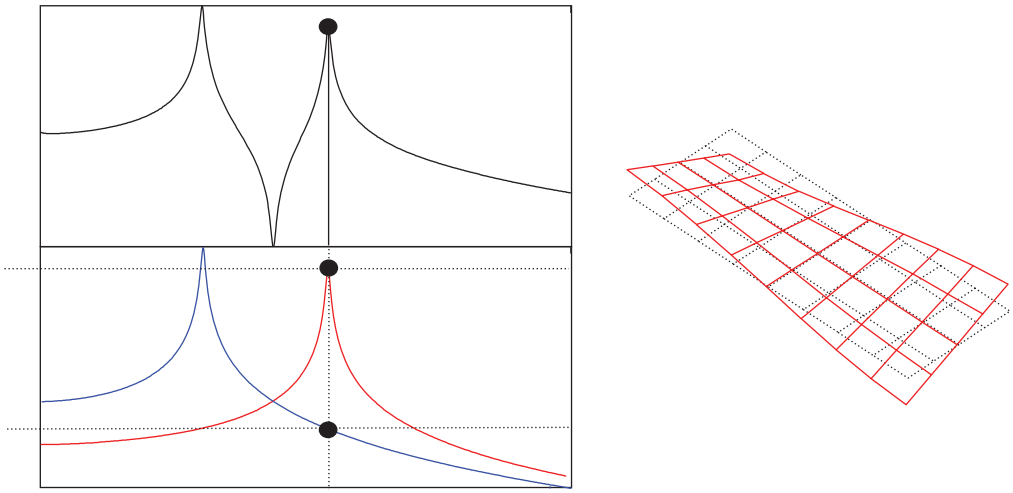

Figure 1.28 Excitation close to mode 2.

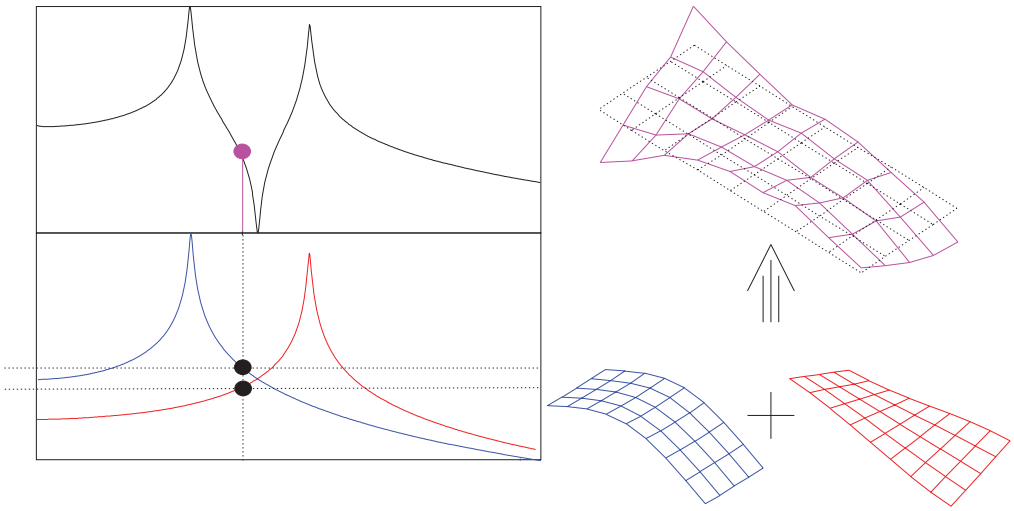

Figure 1.29 Excitation somewhere between mode 1 and mode 2.

long enough, we can actually see a little bit of the first bending and a little bit of first torsion patterns in the deformation. So the operating data is primarily a combination of the first and second mode shapes. (Yes, there will actually be other modes, but primarily modes 1 and 2 are 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, however, we don't collect frequency response functions but rather we collect output spectrums, from which it will not be very clear why the operating data looked like the mode shapes. Figure 1.30 shows an output spectrum 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. However, by understanding how each of the modes contributes to the operating data, it is much easier to see how each one contributes to the total response of the system.

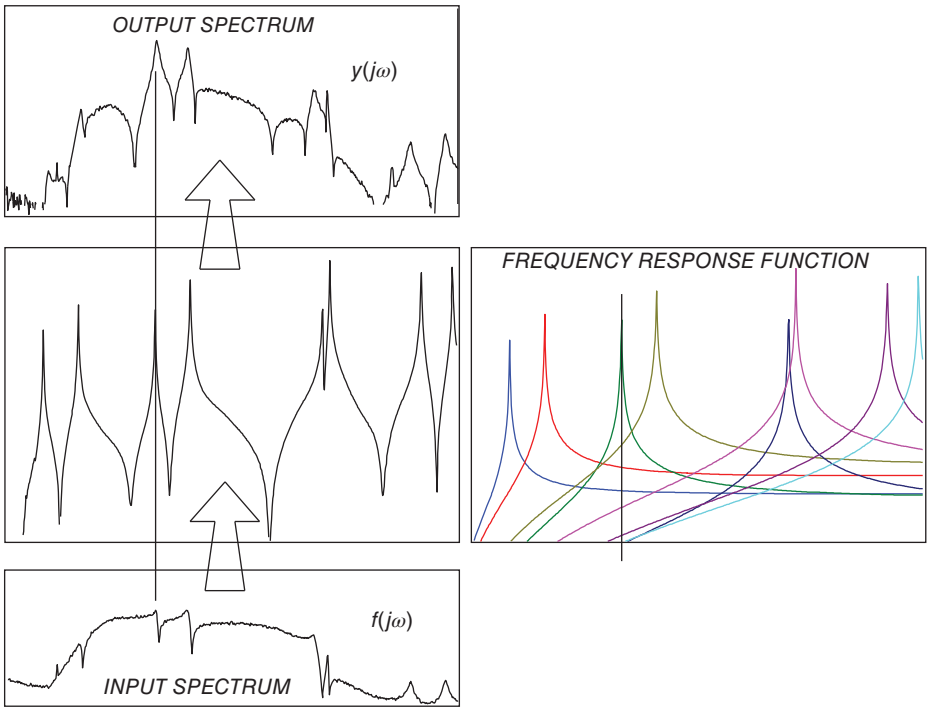

Figure 1.30 Broadband plate excitation.

So 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. However, typically we are interested in the total deformation or response of the system. Why even bother to collect modal data? It seems like a lot of work to acquire measurements and extract data.

1.8.2 So what Good is Modal Data?

Modal data is extremely useful in the design of almost any structure; the understanding and visualization of mode shapes is invaluable in the design process, helping to identify areas of weakness in the design. The development of a modal model is useful for simulation and design studies. One of these studies is called *structural dynamic modification* (SDM). This is a mathematical process that 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 having to *physically* modify the actual structure until a suitable set of design changes is achieved. A schematic of the process is shown in Figure 1.31. There is much more that could be discussed concerning structural dynamic modification but is beyond the scope of this book.

In addition to structural dynamic modification studies, other simulations can be performed, for example force response simulations to predict system responses 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 of the use of a modal model, which are schematically shown in Figure 1.32.

And one of the final questions that are often asked is which test is best to perform – a modal test or an operating test.

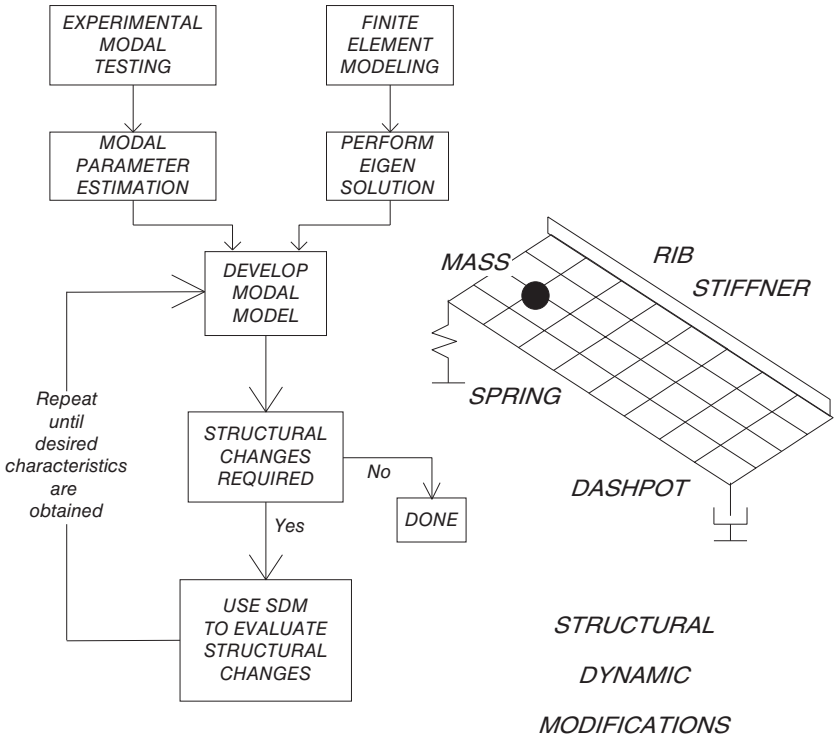

Figure 1.31 Schematic of the SDM process.

1.8.3 So Should I Collect Modal Data or Operating Data?

One of the other questions that is often asked is whether it is better to perform a modal test or an operating test. With tight schedules and budgets, it is debatable whether it is necessary 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 some engineering decisions may have to be made without full knowledge of the system characteristics. To summarize, let's point out the differences between each of the datasets.

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 responses can only be studied using modal data (operating data cannot be used for these types of studies); Figure 1.33 shows this idea schematically. Moreover, 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 because modal data is independent of the forces on the system.

Operating data, on the other hand, is an actual depiction of how the structure behaves in service, as shown in Figure 1.34. This is extremely useful information, but the operating shapes are often confusing and do not necessarily provide clear guidance as to how to solve or correct an operating problem (and structural dynamic modification and response tools cannot be utilized on operating data).

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

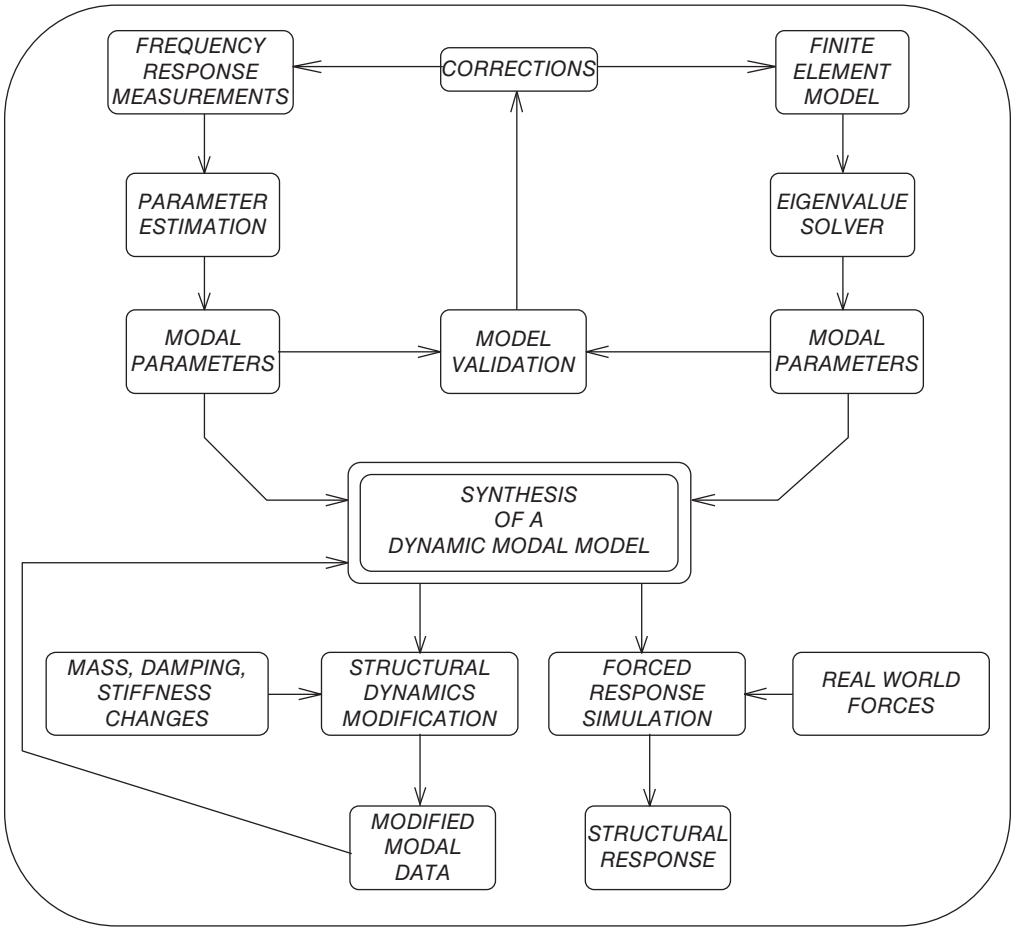

Figure 1.32 Overall dynamic modeling process.

Figure 1.33 Modal model characteristics.

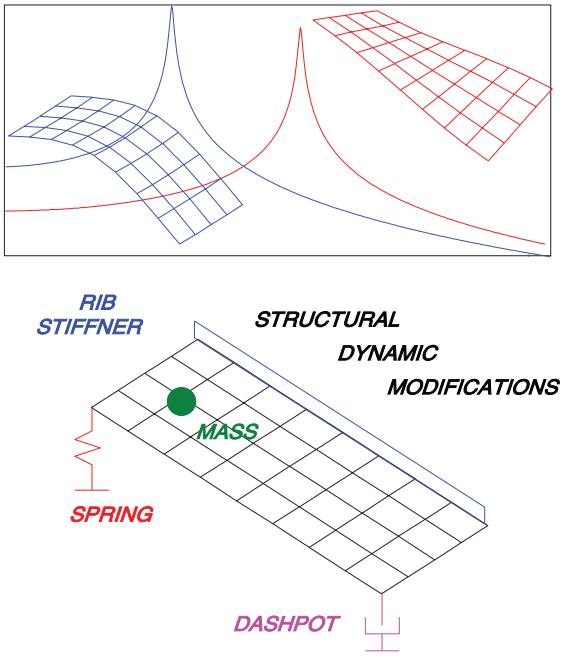

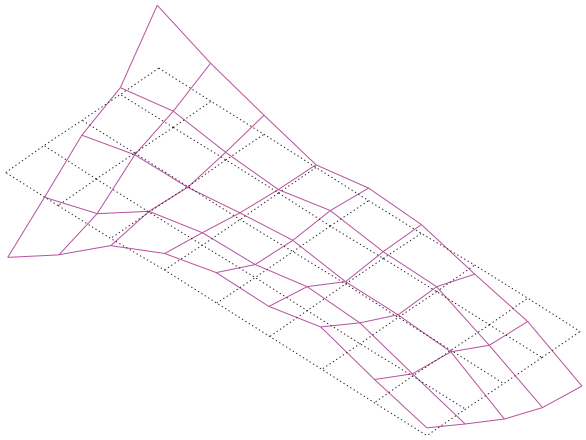

Figure 1.34 Operating data characteristics.

1.9 Closing Remarks

Simple approaches have been 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 mathematics. However, in order to better understand the details of the data presented here, a theoretical treatment of this material is necessary. The next chapter provides some of the basic theoretical material upon which experimental modal analysis is based. This is then followed by chapters on signal processing, measurement definitions, excitation techniques and parameter estimation concepts. The remainder of the book then addresses some very practical aspects of experimental modal analysis and gives many hints and tips as to how to perform a modal test to acquire good frequency response functions and then extract the modal parameters.