REALTIME
Update
Hewlett-Packard Spring 1994

A Newsletter for Noise, Vibration and Electromechanical Test Professionals

Also in this issue:

What's New	5
Teaching Bridges to Fight Back	6
Improving the Accuracy of Modal Measurements	7
Realtime Answers	10
Measuring Floor Vibration	11

The Fine Art of Order Tracking

by Neil Kirtley, LMS International

Ever try to locate a noise or rattle in your car? As soon as you think you've found the source and start to fix it, a new sound pops up somewhere else. Change speed even slightly, and a new problem appears.

Troubleshooting noise and vibration in cars is challenging for several reasons. First, a car is a dynamically complex structure with countless vibration/acoustic modes. Second, the large number of operating forces (such as engine vibration and power-train noise) are equally complex and also time variant. As each harmonic (or order) from the revolving components sweeps through the structure, it momentarily stimulates each resonance and then moves on. Third, the resulting forces can travel through many separate transmission paths and can be completely modified before reaching the occupants. The forces can also interfere with each other, further complicating the measurement situation.

Moreover, the test specialist's goal may not only be to reduce the overall level of noise but to shape the sound in some way to create a particular noise profile (an important task with sports cars, for instance, since drivers want them to sound just right). In either case, the first step is to get a global overview of the noise and vibration situation by acquiring a series of spectra at various speed intervals.

HEWLETT PACKARD

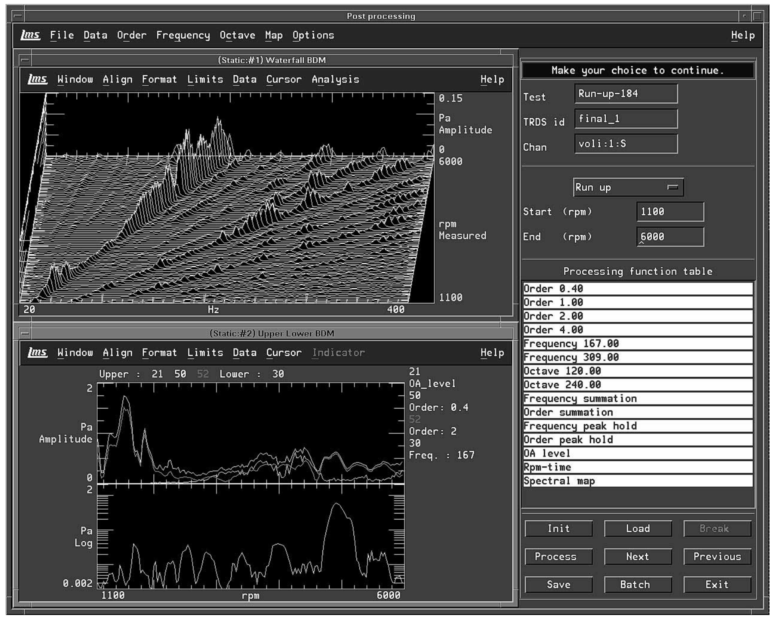

Figure 1:
Using a spectral map to get a global overview and to synthesize order tracks.

Getting to Know the Car

Acquiring a series of noise or vibration spectra at various speed intervals will give you a global picture of the car's behavior. Figure 1 shows a spectral map acquired and displayed by the LMS CADA-X system running with HP 3565S measurement hardware. The map display is assembled by taking measurements at fixed RPM intervals (such as every 50 RPM), using the tachometer channel to control the acquisition process. Structural resonances show up as peaks in the map that do not change frequency with changes in RPM; the forces generated by rotating elements show up as a characteristic "fanning out" of the orders.

You can interpret the orders more easily by remapping the X-axis into the order domain, which keeps the orders parallel on the display. To get a better look at low-level rotational phenomena, you can also use a summation diagram to enhance the order-related components through repeated addition while suppressing peaks from the structural resonances. Another useful display tool is the geometry animation you see in figure 2, which shows the operating deflection shapes (or vibration pattern) of the car as it changes speed. You can think of this as the modern "digital camera" equivalent of the traditional analog stroboscope.

Remapping a spectral map to the order domain does have some limitations, however. First, even at relatively low RPM rates, higher orders fall off the top of the map and can't be processed. Second, accuracy is limited, particularly on interchannel phase, which is vital if you need to resolve noise transmission paths. Third, this approach can miss data at medium and high slew rates.

Keeping up with Digital Order Tracking

The solution to all three of these shortcomings is to acquire data directly in the order domain by tying the system's sampling rate to the engine's rotational rate. LMS CADA-X uses the HP computed order tracking algorithm, which starts by sampling the data at a very high rate and then digitally reprocesses the samples to produce an effective sampling rate that is linked to rotational speed. Computed order tracking also sidesteps the limitations of the analog filters traditionally used in order tracking.

One of the primary benefits of digital order tracking is speed. The data processing overhead is reduced because only the preselected orders are calculated. Performance of 25 measurements per second is typical. In addition, the results offer the phase accuracy needed for precise noise path analysis.

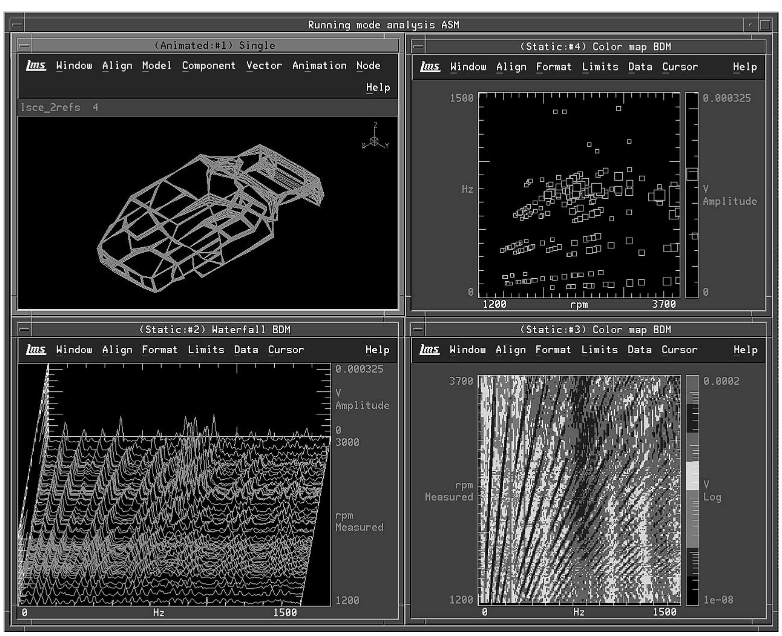

Figure 2:
A geometry animation lets you see the vibration pattern at any point in the run-up.

Staying on Track with the Kalman Filter

Digital order tracking with traditional FFT processing can handle a wide variety of rotational noise and vibration problems, but it can run into trouble when the measurement involves very high slew rates. Some common examples are engine ignition, clutch engagement, and pass-by noise analysis. Critical time-variant phenomena related to nonlinearities or load-dependent system dynamics may show up only under these conditions of rapid changes in speed.

The FFT's assumption of an essentially periodic signal structure and the unavoidable trade-off between time and frequency resolution can hinder its ability to generate accurate results under these conditions. Other time/frequency analysis approaches, such as wavelet analysis, may offer some improvement but are still constrained by the trade-off between time and frequency resolution.

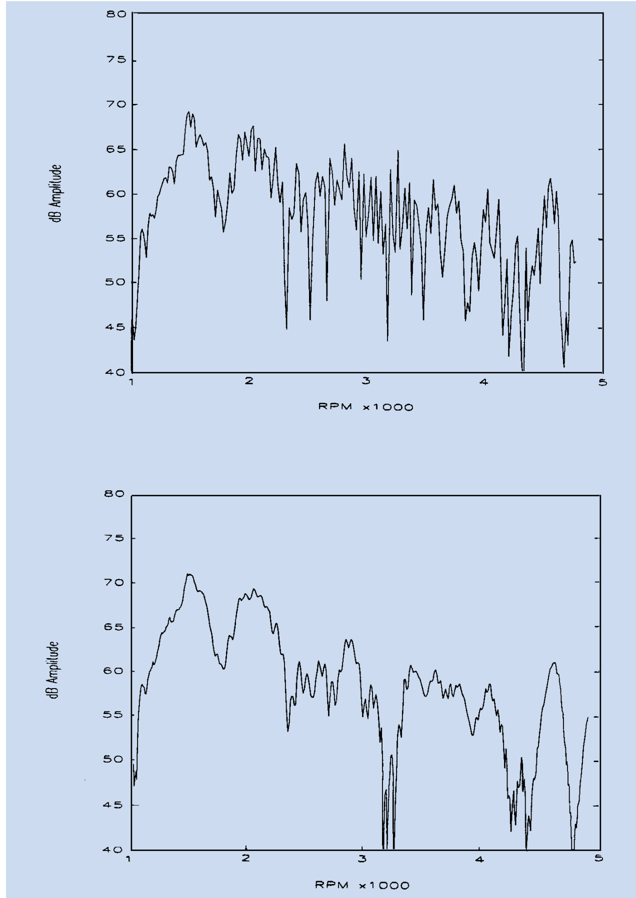

Figure 3:
The FFT results (top trace) and Kalman filter results (bottom trace) are in close agreement, but the Kalman approach has a clear edge in resolution and dynamic range.

Extreme slew-rate situations call for an analysis model that can adapt as the system under test changes. The adaptive filtering technique known as Kalman filtering has proven to be an ideal solution for these tests. The Kalman filter is designed to track signals that have a known structure (that is, known to have order-related components) and that are buried in noise and signal components of a different and unknown structure. LMS has recently implemented a Kalman filter module in the LMS CADA-X analysis system.

The Kalman filter's ability to track noisy target signals makes it crucial to estimate instantaneous rotational speed precisely. Otherwise, the filter could end up tracking the wrong signal. The LMS solution is to sample the tachometer signal at the same rate as the measurement channels and to perform a least-squares fit of a cubic spline to the estimated period.

The tracking characteristics of these Kalman filters have been optimized to accurately track signals with rapidly changing amplitude and phase, even with extreme slew rates. In addition, because the analysis procedure is a true filtering process, it yields an order amplitude for every sample of the run-up.

The Kalman filter provides another unique advantage for noise and vibration analysis: the ability to extract the time history of a particular

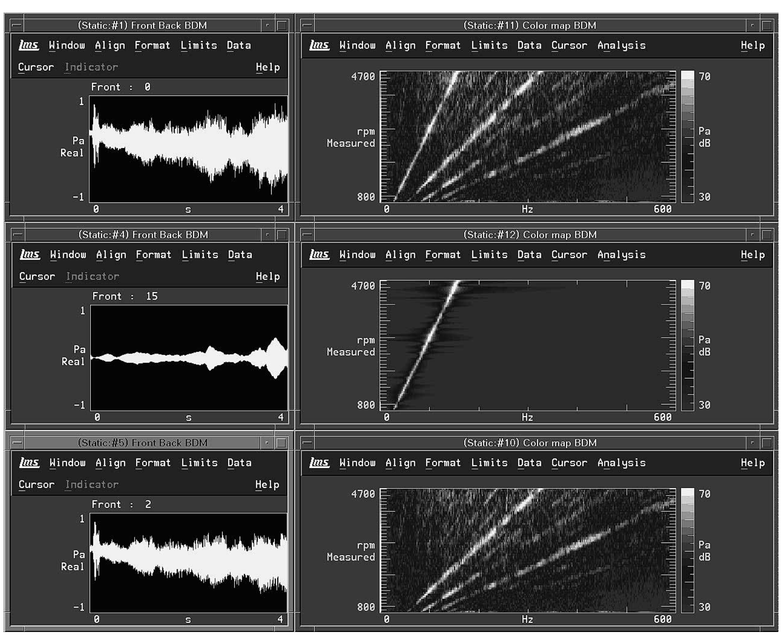

Figure 4:
The top window shows the original signal with the color mapped waterfall, the center window shows the extracted 2nd order, and the bottom window shows the result of subtracting the 2nd order from the original signal.

order from the overall signal. Because it doesn't introduce phase distortion, the Kalman filter can generate error-free time histories for each order. One key application of this feature is isolating the effects of particular orders by subtracting them from the overall signal. To see whether a specific order is contributing to your noise problem, you simply extract its time history, subtract that from the overall signal, then play the result back through a digital to analog converter. This provides a subjective check of the effect that suspect orders are having on noise levels. It also lets you perform "what-if" tests by listening to the simulated improvement before you actually modify the structure in an attempt to remove the effects of one or more orders.

Figure 3 shows a test in which we removed the 2nd order from the overall signal by isolating it and then subtracting it. You can see in the bottom trace that the region around the 2nd order still shows the noise floor caused by aperiodic effects, but the order itself has been removed.

Comparing FFT and Kalman Filtering

To illustrate the advantage of Kalman filtering in high-slew-rate situations, we ran a four-cylinder engine from 1000 RPM to 4900 RPM in roughly 13 seconds. We used a microphone as the response transducer and extracted the 4th order for comparison (see figure 4). The top trace shows the results using the FFT approach, and the bottom trace shows the Kalman filter results. The results are similar, but the Kalman filter has a clear edge in both resolution and dynamic range. ■

Check 1 on the Reply Card to receive the latest information on LMS products.

Neil Kirtley is with LMS International, an HP solutions partner based in Leuven, Belgium. LMS offices in the U.S. may be reached at 313-952-5664 (Detroit), 714-891-4205 (Los Angeles), or 201-227-8722 (New York).

Software Enhancements for the HP 3566A/67A

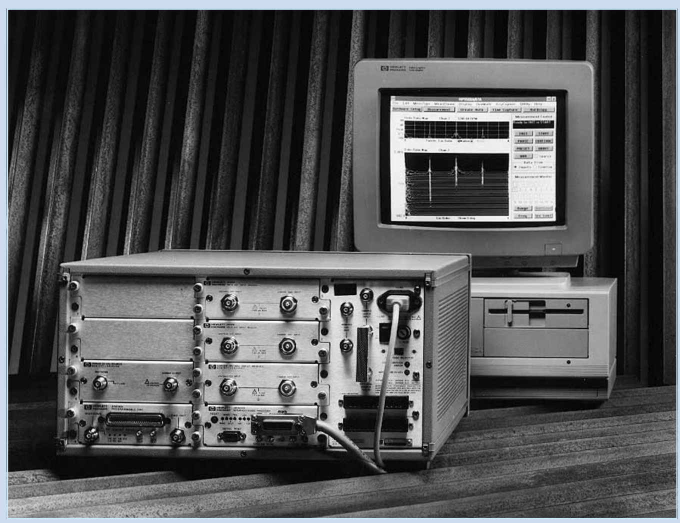

The HP 3566A/67A PC spectrum/network analyzers now offer an even more attractive solution for engine runup and rundown tests and multi-channel measurements in general. With software version A.03.02, these Microsoft Windows-based systems include some key enhancements to computed order tracking and support for up to 48 simultaneous measurement channels.

Order Tracking Enhancements

Automobile and engine manufacturers who need to measure noise, vibration, and ride harshness during runup and rundown tests will appreciate the new capabilities in the HP 35636A order tracking software used with the HP 3566A/67A. These new features can boost productivity and accuracy in both R&D and production testing.

The new RPM tracking measurement displays vibration or noise as a function of engine RPM during runup and rundown tests. Digital filtering techniques model the analog tracking filters traditionally used to make this measurement. Versatile display choices let you plot vibration in 20 orders, 5 fixed frequency bands, or 10 1/3-octave bands.

Other improvements to the HP 35636A software include averaging of multiple runup/rundown sessions to decrease noise and a new "track peak order" result that displays the most dominant order.

Support for 48-Channel Measurements

Version A.03.02 of the HP 3566A/67A software also expands channel capacity from 16 to 48 for most measurements. The extra channels can reduce test time and provide better characterization of mechanical structures during modal or general vibration tests. ■

For more details on these new capabilities, check 2 on the Reply Card.

Teaching Bridges to Fight Back with Active Vibration Control

by Stuart J. Shelley, A. Emin Aktan, David L. Brown, and Randall J. Allemang, University of Cincinnati

The next time you drive across a bridge, imagine the beating that poor structure takes day after day. The forces from water, wind, and seismic commotion, in addition to the relentless pounding from traffic, are enough to give any bridge a case of the shakes.

Infrastructure engineers continually develop new designs and materials to help structures withstand these shock and vibration forces. The challenge, of course, is either to minimize the vibrations that a structure is exposed to or to keep structural resonances out of the frequency range of those vibrations. Doing either or both extends the lifespan of the structure and increases the quality of service it delivers. However, sometimes there are limits to what can be done with static structures, even with advanced designs and materials.

A promising answer for many structures is active vibration control. With this technique, the structure no longer just sits there; it takes steps to counter the vibration with some vibration of its own. You may have heard of noise cancellation techniques, in which equipment noise is canceled when summed with a copy of the same noise signal that has been phase shifted 180 degrees. Active vibration control works on the same principle.

Putting Theory to the Test

With support from the National Science Foundation, the Structural Dynamics Research Laboratory and the Cincinnati Infrastructure Institute (both affiliated with the University of Cincinnati) recently put active vibration control to the test. The structure in question was a decommissioned 250-foot steel truss bridge outside Columbus, Ohio (see figure 1). Our goals were to test active control in a realistic setting and to try a new adaptive modal filter-based control algorithm that takes advantage of a large number of response signals simultaneously.

An active control system consists of (1) sensors to measure vibration on the structure, (2) a computer-based measurement system to evaluate the vibration and generate the necessary compensating control force, and (3) force actuators to apply the compensation vibration to the structure. We used 25 PCB Model 393C seismic accelerometers, an HP 3565S multichannel measurement system, and a 1000 lbf 32-inch stroke electromagnetic mass actuator from Aura Systems. The control algorithm was written with the HP 3565S's Toolkit package.

A vibration control cycle consisted of measuring one time sample from each of the 25 channels, calculating the appropriate control force, and then sending that signal through the digital to analog converter module to the Aura actuator. Based on our study of the bridge's resonant frequencies, we designed the control force to combat global vertical bending modes at 3.4 Hz, 5.5 Hz, 5.7 Hz, 9.4 Hz, and 9.8 Hz. We repeated the control cycle 64 times per second.

Bringing the Bridge Under Control

Our challenge was to reduce the vertical vibration caused by trucks passing over the bridge at highway speeds. Since the bridge was out of service it was not possible to use an actual passing truck to stimulate the bridge. A disturbance actuator did the job for us, hitting the bridge with the same level of vibration that a truck would generate at approximately 55 MPH.

With the bridge subjected to the simulated truck vibration, we unleashed our active control system. The improvements were dramatic. Figure 2 shows the frequency response function measured near midspan, both with and without our control system in operation. You can see that the system attenuated the 3.4 Hz mode by 20 dB and significantly improved the other controlled modes as well. No attempt was made to control the 5.2 Hz mode because our control actuator was positioned at a node point on this mode. (The FRF measurements were made with an HP 35665A dual-channel dynamic signal analyzer using a 1000 lbf random disturbance force from a servohydraulic reaction mass actuator. The measurements were made at the driving point.)

Figure 1:
The Ohio Department of Transportation provided access to the Big Darby Creek bridge on the outskirts of Columbus.

Figure 2:
The active vibration control system dramatically reduced vibration on the bridge deck.

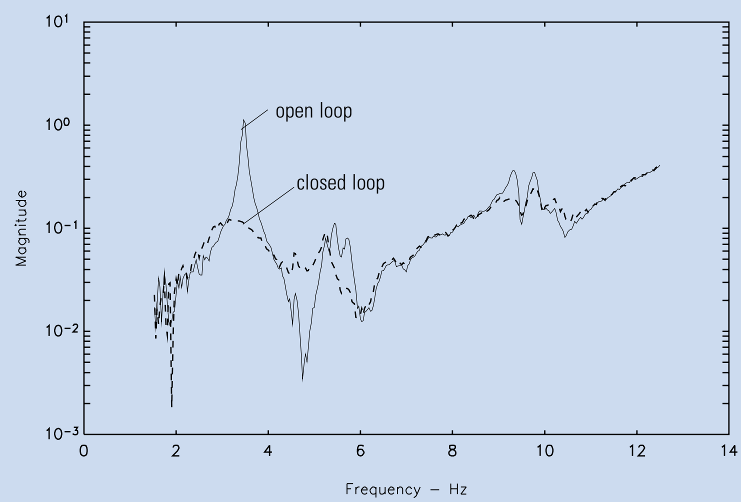

We also evaluated the control system's performance with two other methods, hitting the bridge at midspan with 10,000 lbf vertical step-release inputs and applying a 1000 lbf sinusoidal disturbance, also at midspan. Without our control system in operation, the sinusoidal vibration made it impossible to stand on the bridge for more than a few minutes without experiencing motion sickness. The response amplitudes were as high as 0.2 $in_{p\text{-}p}$. Activating the control system cut this vibration by 75%, making it possible to stand on the bridge without much discomfort.

Structures Don't Have to Be Helpless

The active control system met the project challenge of helping the bridge cope with traffic vibration, and it demonstrated the viability of this technique in a realistic setting. The technique is not limited to bridges, however. Buildings face many of the same vibration problems as bridges, with a few of their own special issues (such as machinery and manufacturing installations). In fact, designers in Japan are equipping a dozen existing or planned buildings with active control systems.

We've moved on to our next vibration control challenge: a large dance floor. If we can keep that from shaking during a high-energy dance, we should be able to control just about anything. ■

Stuart Shelley, David Brown, and Randall Allemang are with the Structural Dynamics Research Laboratory at the University of Cincinnati; they can be reached at 513-556-3535. Emin Aktan is with the Cincinnati Infrastructure Institute, also at the University of Cincinnati, and can be reached at 513-556-3689.

Realtime Basics:

Improving the Accuracy of Modal Measurements

by Ken Singleton
Eastman Chemical Company

Successful modal tests require a solid understanding of structural dynamics and measurement techniques, along with a healthy dose of hands-on judgment. The more structures you can test and analyze, the better you'll be able to reconcile theory (what you'd expect to see) with reality (what you actually observe in your tests). You can also take advantage of other people's experience through the numerous articles and publications available on structural theory and test techniques.

This article shares five key points of accumulated wisdom: understanding the assumptions that apply to all modal tests, selecting measurement points, using impact hammers, setting up force and exponential windows, and keeping measurements clean by avoiding hammer bounce.

Important Assumptions in Modal Testing

Modal testing assumes a number of things about the structure you're testing and your measurement setup. Of course, it's unlikely one would ever encounter a situation in which every one of these assumptions holds true, but understanding them provides more insight for setting up tests and interpreting the results:

1. The structure will not start to vibrate until a known, controlled excitation is applied.

2. The induced vibrations will die out after the excitation is removed.

3. The dynamics of the structure will not change during the test. (Consider, for instance, how the effects of mass loading from transducers may violate this assumption.)

4. The structure exhibits symmetry (Maxwell's reciprocity theorem). In other words, a stimulus applied at point A with the response measured at point B will yield the same results as the identical stimulus applied at point B and measured at point A.

5. Only one mode exists at each frequency. (This assumption in particular is invalid for many structures.)

6. Modes are global, not local. In other words, the entire structure vibrates when the mode is excited, not just a portion or localized section.

7. The structure responds in a linear manner. (A common concern with this assumption is making sure there are no loose connections between structural components, loose bolts, or other structural discontinuities.)

Again, you can't expect to encounter a perfect situation with every modal test, but knowing how each structure might deviate from the assumptions will help you understand the behavior you're measuring.

Selecting Measurement Points

One of the most important opportunities for exercising good judgment is in selecting the points on the structure where you'll make your measurements. The optimal number and location of measurement points depends on both the structure itself and the purpose of your test. The more points you choose, naturally, the more time you'll need to set up and collect and analyze data (and you may need additional equipment as well). On the other hand, consider the effect of having too few measurement points. Figure 1 demonstrates the misleading modes that can appear when you don't use enough points to accurately characterize mode shapes.

Figure 1:
Using too few measurement points can cause spatial aliasing, which produces false mode shape indications.

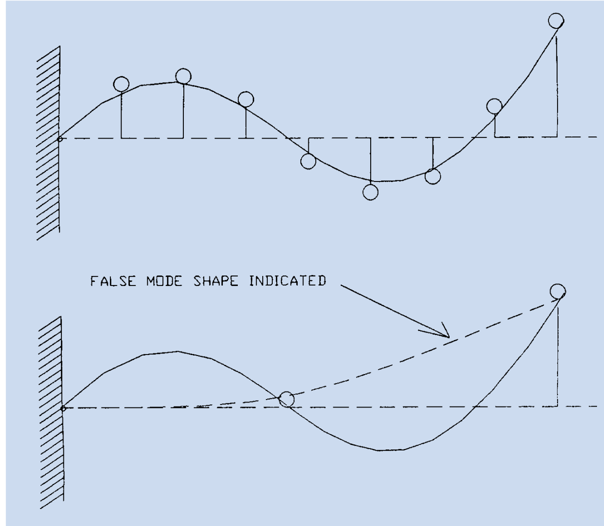

Understanding Impact Forces

Impact testing using instrumented hammers is a popular modal test technique. The equipment is relatively inexpensive and easy to transport. Most impact tests are easy to set up, and the results are accurate enough for a wide variety of test scenarios.

However, consider the nature of the impact force that your hammer delivers to the structure under test. A hammer impact is a broadband signal, which means it delivers energy across a wide frequency span. (In contrast, an electromagnetic shaker driven by bandlimited random noise or a sinusoidal signal contains the energy inside the frequency span you've chosen.) Depending on your structure and the measurement span you've selected on your analyzer, the hammer impact might excite out-of-band modes. You won't see these modes in your measurement results, but they can affect the frequency response functions inside your frequency span.

Another aspect of hammer impacts to consider is the nature of the hammer tip. As you can see in figure 2, hard, medium, and soft tips generate different types of impact pulses and therefore have different frequency profiles. For instance, a hard tip generates higher frequency components that may affect your measurement, but it provides a much more even distribution of energy across the spectrum. On the other hand, soft tips provide more energy at lower frequencies but roll off much faster. A good rule of thumb when choosing the hammer tip is to make sure that the response doesn't roll off by more than 10 to 20 dB at the maximum frequency of interest.

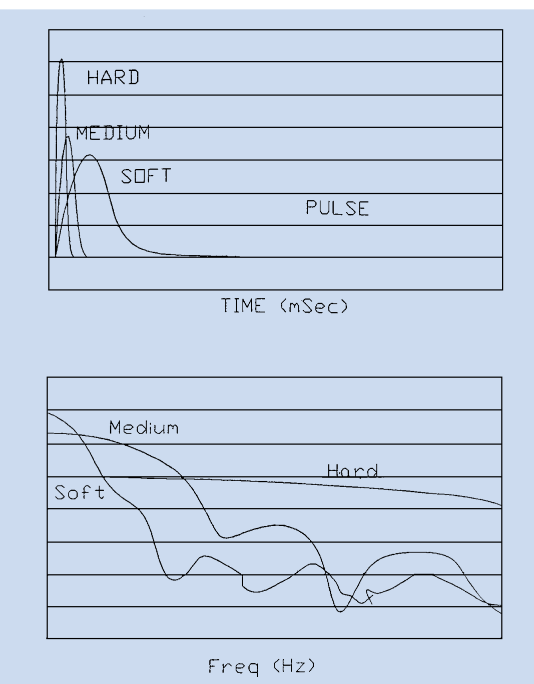

Figure 2:
The shape of the impact pulse and its frequency profile depend on the hammer tip you've chosen.

Setting Up Windows

Your analyzer provides force and exponential windows to help you deal with three characteristic issues in impact testing. First, the duration of the energy pulse generated by the hammer impact is typically much shorter than the time record over which you're collecting measurement data. In other words, the analyzer is still collecting input data when the pulse is long gone. To avoid contamination from noise during this part of the time record, you should apply the force window on the stimulus channel. The force windows are typically set at 10% of the length of the time record, meaning that the stimulus channel will be safe from noise during the last 90% of the record.

The second problem can show up on the response channel when you're measuring lightly damped structures in which the response doesn't decay to zero by the end of the time record. Because the FFT process assumes periodicity from record to record, a response that is still "live" at the end of the record will create leakage across the frequency spectrum. The solution is to force the response to zero by applying the exponential window, as you see in figure 3.

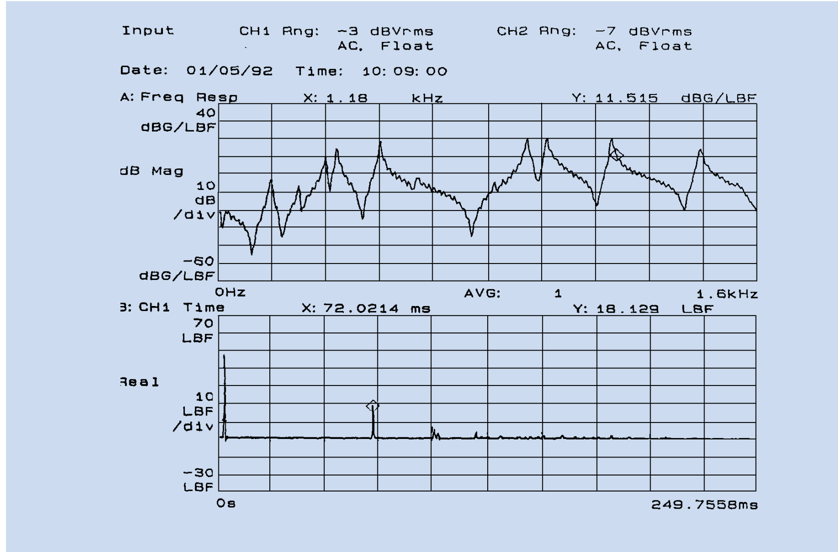

The third potential concern is a response channel issue as well. The risk of noise contamination exists in heavily damped structures that decay long before the time record ends (not unlike the noise problem on the stimulus channel.) The exponential window can help here as well by causing the analyzer to ignore any activity (namely noise) that appears on the response channel after the response has died away.

Getting Clean Hammer Hits

One more challenge with hammer testing is trying to ensure clean, consistent impacts. A very lightly damped structure may vibrate so much after initial impact that it strikes the hammer head before you pull the hammer away. Failing to hit the structure quickly and crisply can have the same effect.

The force window keeps multiple impacts out of the stimulus side of the frequency response calculation, but the energy is already in the structure and will contaminate the response. You can see the effect this has in figure 4, where the response curve is jagged and uneven. To avoid this problem, always use manual preview mode on the analyzer. This lets you view each response to see whether your hit was clean, and you reject any measurement that looks suspect.

Learning More

Experience is the best teacher for all of us in the modal analysis field, and it's also helpful to review as many conference proceedings, journal articles, and manufacturer's application notes as possible. Every test situation is unique, and the more we know, the better equipped we are to handle each new scenario. ■

For a copy of the references used in developing this article, check 8 on the Reply Card.

Ken Singleton is with the Machinery Analysis Group/Plant Engineering at the Tennessee Eastman Division of Eastman Chemical Company in Kingsport, Tennessee. He can be reached at 615-229-3053.

Figure 4:
You can see how hammer bounce affects the frequency response (top trace); the multiple hits made the measurement look rough and ragged.

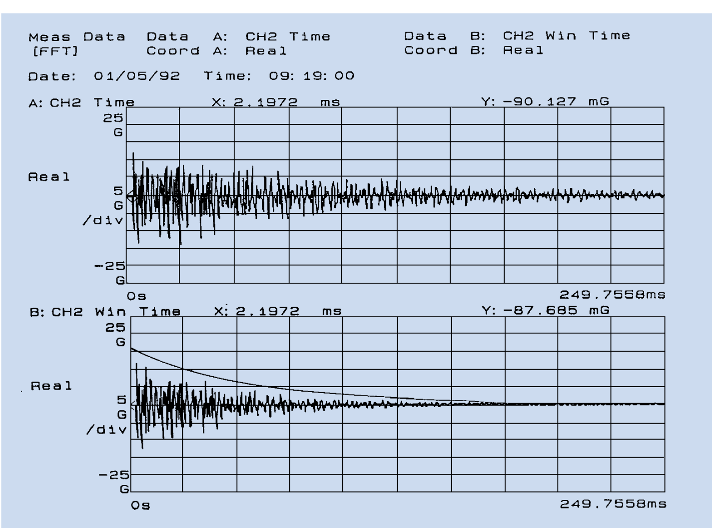

Figure 3:
The top trace shows the unwindowed response; the bottom trace shows how the exponential window drives the response toward zero by the end of the time record.

REALTIME Update

Issue: Spring '94
Realtime Update is a quarterly newsletter published by Hewlett-Packard Company. Please address comments, special requests, and other correspondence to:

Rick Van Ness, Editor
Realtime Update
8600 Soper Hill Road
Everett, WA 98205-1298 USA.
Fax: (206) 335-2828
Tel: (206) 335-2654

Contributing Editors
Dick Bingham
Steve Bye
Paul Gallagher
John Jensen

Manuscript Editing, Design and Production
The Dovel Group
Desktop Connexion, Inc.

Subscription and Fulfillment
Louise Williamson
HP Direct Marketing Organization

Realtime Answers

Question: I'm using an HP 35665A and would like to make a power spectrum measurement in which all the traces are displayed in a waterfall. I want to trigger the spectrum measurements at random times, adding a trace to the waterfall whenever trigger conditions are met. The result would display frequency on the X-axis, amplitude on the Y-axis, and time on the Z-axis. How can I do this?

Answer: Start by setting up your HP 35665A for external trigger, with arming set to Time Step Arm. Make sure the arm time step size is smaller than the smallest trigger interval you'll encounter. For example, if you will trigger as frequently as every second, be sure the trigger is armed in less than a second. This ensures that the analyzer is always armed and ready.

When the instrument triggers, it will update the waterfall display. Each trigger will add one more trace to the waterfall, and Time Step Arm mode keeps track of the time each trace is added. Use the Waterfall Marker "Trace Select" to scroll through the waterfall display picking out each trace and time. The time will be relative to the initial trigger. The display will show traces at regular intervals along the Z-axis, even if the actual trigger intervals were not regular. The marker will read out the correct time values. (Note that this technique is not applicable to 1/3 octave measurements.)

Question: During a recent measurement with my HP 3569A, I turned on the Hold Min mode. When my measurement was completed I noticed that the final Minimum Y values were larger than some of the values I'd seen during the measurement. Why did Hold Min end up with larger values than some I'd seen while the measurement was in progress?

Answer: When Hold Min is active, the HP 3569A displays the lowest average value in each frequency bin. Assume you set the HP 3569A to take ten measurements, with two-second integration times. The result will be ten spectra, and the process will take twenty seconds. Hold Min selects the minimum value from one of the ten measurements.

If the HP 3569A is set to display incremental updates, you will see a new display four times a second. Because your integration time is two seconds, you'll see eight display updates. The minimum value will be the lowest *average* value for any completed measurement. During the incremental updates you may see a lower value, but the final value after the two seconds of integration is the only valid data point.

Question: Using the HP 3560A and SDF, I can transfer measurement data to a PC; can I also transfer set-up state information?

Answer: First select either HP PRINT or HP PLOT for the HP 3560A's display type. (HP PRINT will give you an ASCII file of the screen data; HP PLOT will give you an HP-GL file that you can import into word processing programs.) Next, set the analyzer to display the Status or Measurement state. With the analyzer connected to your computer's serial port, run the SDF utility DOWNLOAD. The computer will indicate that it is "Waiting...". When you see this message, hit the Shift-Print Screen key sequence on the HP 3560A. The analyzer will then send the file to the PC. Note that this technique works with the HP 3569A as well.

The HP 3560A manual explains how to set up the RS-232 interface correctly. The HP SDF Utilities manual (included with the HP 3560A) explains how to use the DOWNLOAD utility. ■

Ask Us *Your* Question

If you have a question about the operation, performance, or programming of any Hewlett-Packard signal analyzer, we're here to help. See page 12 for the most convenient way to reach us.

Measuring Low-Level Floor Vibration

by John Jensen

As semiconductor manufacturers push deeper and deeper into submicron geometries, floor vibration becomes an increasingly vital concern. Measuring this vibration is one of the more challenging applications of dynamic signal analysis. On one hand, a measurement result that is higher than the actual vibration level could lead to expensive seismic isolation work that really isn't necessary. On the other hand, a measurement that underestimates the actual level could mean that sensitive fabrication equipment won't operate to specification, and the cause will go undetected.

Because the acceptable level of vibration is so low, these measurements often push the sensitivity limits of both the analyzer and the transducer. Three key precautions will reduce errors and improve the measurement results.

The Analyzer's Noise Floor

The first step is making sure the analyzer's own noise floor is low enough. How low is low? A good rule of thumb is to choose an analyzer whose noise floor is 10 times lower than the vibration level you need to characterize. For instance, if you need to verify that the vibration doesn't exceed 16 μG_{rms}, make sure the analyzer's noise floor is no higher than 1.6 μG_{rms} inside the relevant frequency range. You can check this by connecting a resistor across the analyzer's input channel and measuring the power spectrum with the lowest input range selected.

The Transducer's Noise Floor

Similarly, you need to measure the transducer's noise floor to make sure it is low enough. This step can be a challenge because the transducer may sense ambient vibration from air conditioning units, street traffic, foot traffic, everyday seismic activity, and other sources. The only way to ensure an accurate evaluation is to mechanically isolate the transducer (by suspending it from a tripod, for instance) in a "quiet" room that is situated on or near bedrock.

The Analyzer's Input Range

Once you're set up and making measurements, you must take care to maintain the correct range on the analyzer's input channels. For instance, if you activate "autorange up only" on the analyzer and then touch or move the transducer, you might cause the input range to ratchet up. As the input range moves up, so does the noise floor. To avoid inadvertently measuring on a higher input range, mount the accelerometer with its power supply on and then set the analyzer to its minimum input range. If the overload light goes out within a couple of minutes, you're ready to go. If it doesn't, increase the input range one notch and try again. ■

This article was adapted from "Low-Level Floor Vibration Measurements," a tutorial paper that contains more-detailed test steps and a number of example measurements. To receive a copy, check 9 on the Reply Card.

John Jensen is a Senior Applications Consultant in Hewlett-Packard's Mountain View, California, office (415-694-3399). John wishes to thank Wayne Vogen of Vibration Engineering Consultants (510-339-8719) for pointing out the need for this tutorial, offering his insights and advice, and providing the quiet floor for making these measurements.