REALTIME
Update
Hewlett-Packard Winter 1994

A Newsletter for Noise, Vibration and Electromechanical Test Professionals

Managing a Large-Scale Modal Test

by Abraham Frydman, Dean Li, Stuart Shelley, and Jeff Poland

Keeping the data and details straight in any modal test can be a challenge. Add several configurations of a large test article and a multiple-input, multiple-output (MIMO) setup with hundreds of response points, and you have a sizable management task on your hands.

We recently had the opportunity to test our management skills at the U.S. Army's Aberdeen Proving Grounds in Maryland. Our assignment was to collect detailed modal data for correlating with a finite element (FE) model. The test articles were two M113 Armored Personnel Carriers, one a standard issue aluminum version and the other a special composite version.

To characterize the vehicle and its components fully, we tested two configurations of each version. The "basic" configuration was the bare hull. With no nonlinear joints, rubber elements, or other complications, we knew this version would exhibit linear, lightly damped behavior. This submodel would give us a good look at the dynamics of the vehicle shell and serve as the foundation for our complete system model.

Also in this issue:

HEWLETT PACKARD

In the "tight" configuration, we added all components that we judged had a tight fit with the hull. The goal here was to keep things as linear as possible while still characterizing a more complete vehicle. The elements we added included the engine, hatch, back ramp, and final drive casings.

Setting up the Test

We needed to achieve empirical boundary conditions that were as close as possible to the FE assumptions, so we supported the vehicles on four Firestone Airmount Isolators (see photo on page 1). This established a quasi-free boundary condition in which the supports did not significantly affect the vehicles' elastic modes.

Four MB Dynamics Model 50 electromagnetic shakers provided the excitation. We positioned one vertically at each corner and used stingers to prevent lateral and torque loads from being applied to the structure. A PCB Model 208A02 load cell measured the input stimulus from each shaker.

The complexity of the vehicle geometry dictated quite a few response measurements. We used roughly 230 accelerometers on the basic configurations and 350 on the tight configurations; the setup involved three types of PCB Flexcel accelerometers. Since the vehicles are fairly rigid along the fore-aft axis, there was little response in this direction. Therefore, the measurements were primarily biaxial in the lateral and vertical directions.

Managing the Data

Keeping track of two or three hundred measurement points is no simple task. Fortunately, the University of Cincinnati's Structural Dynamics Research Laboratory (UC-SDRL) has a great deal of experience with similar tests. UC-SDRL has also created custom hardware and software to help calibrate accelerometers and organize data. The UC-SDRL array calibrator allowed us to calibrate 128 transducers simultaneously. This system creates a computer file that links calibration values with transducer serial numbers, helping to ensure accurate measurements once we moved to the test location.

User programs written for the CADA-X data acquisition software from LMS (running on an HP Model 380 workstation) managed the measurement channels, transducer calibrations, and mounting point locations and directions. A pair of 128-channel PCB Data Harvesters provided signal conditioning for the transducers. We switched the Data Harvester outputs into an HP 3565S measurement system 64 channels at a time. This required four acquisition cycles for the basic configurations and six cycles for the tight configurations (see figure 1).

Signal Processing Considerations

Covering the relevant frequency range on the basic configurations required two complete measurement passes, once at 0 to 128 Hz and again at 100 to 228 Hz. We measured the tight configurations once at the 0 to 128 Hz span. Approximately 100 averages were used to estimate each frequency response function.

We fed the shakers with a burst random signal, using a 50% burst cycle for the basic configurations and a 70% cycle for the tight configurations. Because the very lightly damped response of the basic configurations didn't decay by the end of each time record, we forced the response to zero by applying an exponential window to these measurements.

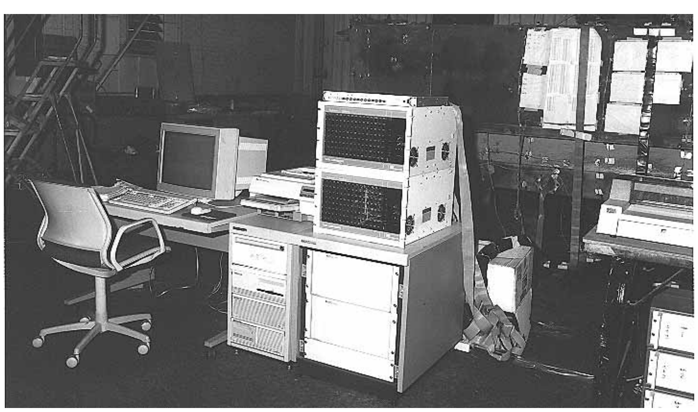

Figure 1: Collecting and Managing Data Our test setup included HP 3565S measurement hardware, PCB Data Harvesters and acclerometers, and LMS CADA-X software.

Results

In any modal test, it's good practice to check the data periodically. This guideline is doubly important with expensive and time-consuming tests. When you're spending thousands of dollars to conduct a test, you don't want to discover bad data the day after you've dismantled everything. Quick checks can help uncover problems with transducer mounting or calibration, unexpected nonlinearities in the structure, faulty or improperly configured data acquisition channels, and other potential glitches.

We used three techniques to look for problems. The first is based on Maxwell's reciprocity theorem, which states that a measurement result obtained by exciting at point A and measuring at point B should be identical to the result obtained by exciting at point B and measuring at point A. (The theorem applies only to systems that are not under closed loop control and that have no actively rotating components, which applied in our case.)

The second quality check was a linearity test to see how the vehicles responded to different stimulus levels. We checked once at 50% of the normal level and again at 200% of normal. The final check was a cursory curve fitting analysis. We used the modal analysis module in the LMS CADA-X software to see whether the mode shapes made sense, based on what we knew about the geometry.

Thanks to careful planning and some equipment designed for high-channel count measurements, we met the goal of providing accurate data for a precise modal model. Paying attention to details and keeping everything organized are the secrets to achieving high-quality results. ■

Abraham Frydman and Dean Li are with the U.S. Army Research Laboratory and can be reached at 301-394-2804. Dr. Stuart Shelley is with the University of Cincinnati and can be reached at 513-556-2720. Jeff Poland is with Vold Solutions and can be reached at 513-871-4136.

Getting Maximum Accuracy In Vibration Measurements

by John Jensen
Hewlett-Packard Company

Today's dynamic signal analyzers can provide remarkably accurate sound and vibration measurements. To get high-quality measurements every time, however, you have to play an active role in the process. Transducers, signal conditioning equipment, and the analyzer make up a measurement chain. The better you understand how the chain fits together, the more accurate your measurements will be.

Even if you aren't testing to meet certain specifications (in which case the need for accuracy is obvious), calibrated accuracy can still be vital. If you're trending machine vibration over time, for instance, calibrating to a known accuracy standard enables you to compare current and historical measurements with confidence.

Using Engineering Units

Whether you're measuring acceleration, velocity, or displacement, the input signal is simply a voltage level as measured by the analyzer. To calibrate the measurement and display the results in your desired units (such as gs for acceleration or mm/s for velocity), you need to use the feature called engineering units (EUs). This is simply a multiplier that tells the analyzer to scale the measurement results according to the units you've selected. The scaling factor is determined by the individual transducer and any signal conditioning equipment you may be using.

To scale a measurement in EUs, you enter a volts/EU scaling factor using the appropriate softkey menu. Some analyzers, such as the HP 35670A, provide a CAL AT MARKER capability. This lets you measure a known calibration signal and use that value to scale the EUs. In addition, you can enter an EU label on most analyzers so that the display, markers, and hardcopy are presented in gs, mm/s, or whatever units you've selected.

Making accurate, calibrated EU measurements involves three steps: determining the transducer's sensitivity, selecting a reference level, and running a "reality check" on the results.

Step 1: Determine the Transducer's Sensitivity

For force transducers and accelerometers with integrated charge amplifiers, simply obtain the V/EU value from the manufacturer's data sheet and enter that number into the analyzer's EU conversion field.

Some displacement transducers, such as LVDTs (Linear Voltage Differential Transformers) and capacity and eddy current proximity probes, are scaled in mV per millimeter or mV per milli-inch. If you enter these scale factors directly, the analyzer might produce strange display scales such as "milli mils." To avoid this, enter the scale factor in volts/meter or volts/inch. The analyzer will then create the proper scaling in micrometers or micro-inches.

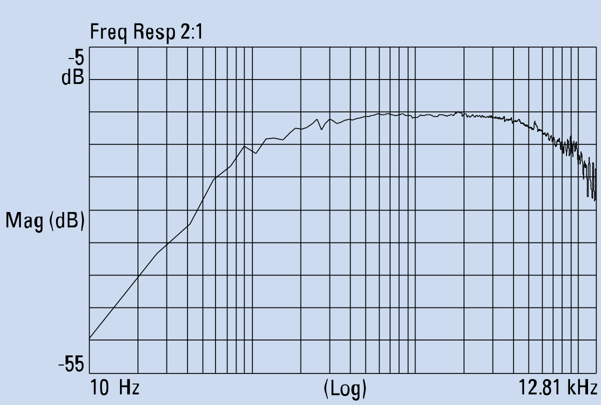

Figure 1: The Risk of Bypassing Charge Amplifiers
Measurement quality suffers if you bypass the charge amplifier when using charge-coupled transducers. Here you see the roll-off in the transducer's frequency response, which will be left uncorrected if you don't use a charge amplifier.

Obtaining a V/EU scaling is more complex with transducers that use external charge amplifiers (partly because you need to account for the charge-to-voltage conversion and partly because there are no industry-wide standards for charge amplifiers). Start with the sensor's output, which is given in pC (picoCoulomb) per g (or some other unit). Dial this value into a calibrated 10-turn potentiometer, and then select a gain or full-scale range value, such as 300 g. If the range is in g's full scale, you need to get the full-scale output voltage from the charge amplifier's manual.

Common full-scale output ranges from charge amplifiers are 5 V and 10 V. A 10 V full-scale output range and a 300 g full-scale range yield a scaling factor of 33.3 mV/g. (The scaling on these charge amplifiers was intended to simplify the visual conversion of volts to gs when using oscillographic recorders and oscilloscopes.) If you're using a charge amplifier that's already scaled in mV/g, simply enter that value into your analyzer and skip this step.

Sometimes charge-coupled accelerometers are connected directly to the DSA without a charge amplifier. With a handheld shaker calibrator, you can easily enter the scaling factor with the analyzer's CAL AT MARKER feature. These calibrators typically generate a 1 g signal at 159.2 Hz (a signal chosen to yield a round number for displacement, $20\ \mu m_{pk\text{-}pk}$.)

However, using a charge-coupled accelerometer without a charge amplifier will not produce a fully calibrated result, even after you've used the handheld calibrator. This setup provides a general indication of vibration levels, but since the measurement is calibrated only at the calibrator's 159.2 Hz frequency, this technique is risky. (Working without the charge amplifier presents two major accuracy problems. First, the measurement is more vulnerable to noise from nearby power lines and other sources. Second, without the charge amplifier, the transducer no longer yields a flat frequency response (see figure 1). You have no assurance that a 1 g signal at 10 Hz will produce the same voltage as a 1 g signal at the calibrated frequency of 159.2 Hz.)

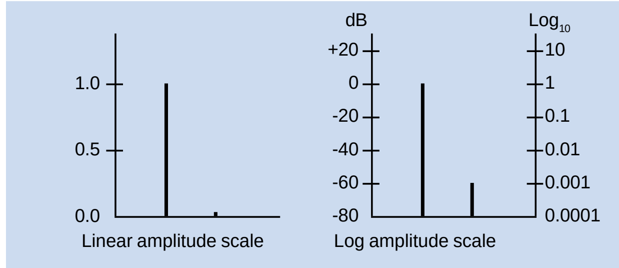

Figure 2: Benefits of the dB Scale
Displaying measurements with the log (dB) scale makes it much easier to see small and large signals simultaneously.

Step 2: Select a Vibration Reference Level

Your next consideration is the vibration reference level. If you want to use a log amplitude scaling to display your data, you'll need this reference to define the decibel (dB) units. Even if you plan to use linear scaling, you'll need to choose the units to display the measurement properly.

Most DSA users prefer log (dB) scaling because it lets them see large and small signals simultaneously, a frequent need in vibration testing (see figure 2). The dB scale is convenient for calculations as well, since figuring the distance from a measured point to your reference level is a simple subtraction problem. You can also calculate absolute levels from a dB value and the reference level value:

$$dB = 20 \log \frac{G_m}{G_{ref}}$$

where G_m is the measured level and G_{ref} is the reference level.

The nature of the test usually dictates the reference level to use. Military and government agencies, industry associations, and various international test standards organizations often specify certain reference levels. Mil. Std. 740B, for instance, specifies 1 $\mu m/s^2$. In other words, 1 $\mu m/s^2$ input signal will result in a 0 dB reading on the analyzer. A 1 g input (9.81 m/s^2) will produce a 140 dB reading.

With so many references in common use, you need to make sure you're using the correct one and that you've scaled your measurements accordingly. Table 1 lists other common reference levels.

In today's competitive markets, vibration data must be compatible with your customers' needs, whether your customer is an external company or someone inside your organization. Some may want the data scaled in English or metric units (or both) in dB relative to any of the common reference levels. At one time, converting English to metric required changing the V/EU value from English to metric and remeasuring the raw data. Now, DSAs such as HP 35670A have the conversion capability built in, eliminating the need to remeasure.

The conversion routines provide acceleration-to-velocity-to-displacement conversions for spectral measurements and acceleration-to-velocity conversions for time domain measurements.

Step 3: Run a Reality Check

Running a reality check will help ensure you're getting reasonable results. The handheld 1 g shaker calibrator can help you verify that everything in the measurement chain is functioning properly. However, there are a couple of cautions to consider. First, your knowledge of the calibrator's sensitivity is likely to be less accurate than your knowledge of the transducers' sensitivities. In other words, the calibrator probably won't uncover small errors.

Second, an accelerometer can have a fractured crystal and still produce a signal level that looks reasonable but is in fact incorrect. So use the handheld calibrator to look for substantial differences (more than 5%) between expected and measured values.

If something does look out of line, cross-check the calibrator output with an accelerometer that you know to be good. Then recheck the charge amplifier settings, the accelerometer cabling, and the accelerometer itself.

Table 1

English units

g_{ref}	1 g dB level	1 mg dB level
1 g=32.2 ft/s^2	0 dB	-120 dB
386.4 in/s^2	0 dB	-120 dB
1 μg	120 dB	0 dB
1 in/s^2	51.7 dB	-68.26 dB

Metric units

g_{ref}	1 g dB level	1 mg dB level
9.81 m/s^2	0 dB	-120 dB
10 m/s^2	-0.1667 dB	-120.167 dB
10 $\mu m/s^2$	119.8 dB	-0.17 dB
1 $\mu m/s^2$	139.8 dB	19.6 dB

Also, look at the time domain displays of the calibrator and accelerometer signals. Loose connections or defective cables produce noise spikes around the peaks of the sine wave. Figure 3 is an example of the signal from an accelerometer with a loose connector.

During your reality check and any subsequent troubleshooting, use the transducer's calibrated sensitivity value as your primary calibration standard. Then use the handheld calibrator to verify that all equipment in the measurement chain is working within acceptable limits.

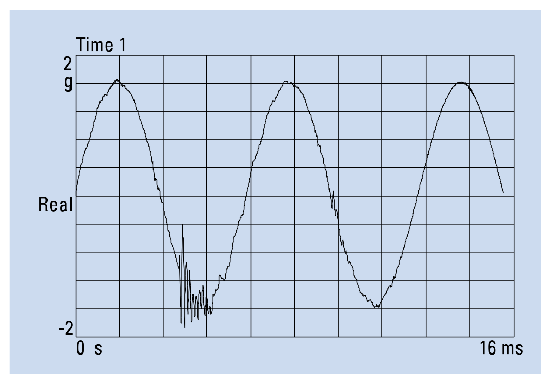

Figure 3: Searching for Setup Problems
Looking at the time domain helps uncover setup problems such as loose accelerometer cables.

One final note: Always recheck the calibration of the accelerometer at the end of the data-gathering session and save the pretest and posttest calibration measurements. This will give you a calibration audit trail in case anyone questions the results later. ■

John Jensen is a senior application engineer in Hewlett-Packard's Mountain View, California, office. He can be reached at 415-694-3399.

Realtime Basics:

Extracting Small Signals from Noise

by Anthony Nash
Charles M. Salter Associates, Inc.

Wouldn't it be great if every signal you wanted to measure jumped right out of the background noise? You could get a clear look and always be confident that you were measuring the signal of interest. Unfortunately, life isn't always so cooperative. Signals are sometimes buried in random noise, making accurate measurements a big challenge. In our tests of massive structures, we frequently encounter this problem.

The good news is that today's dynamic signal analyzers can help. Assisted by HP engineers and their application literature, we have devised an averaging and windowing technique that extracts those elusive signals from the noise.

Time Averaging

The HP 35670A and other DSAs give you a choice of averaging methods, each with a specific goal. With the right measurement setup, *time averaging* can actually reduce the noise level in your measurements. The result is a clearer picture of the signals you're trying to analyze.

To use time averaging in a vibration response measurement, you need a repetitive stimulus signal that is synchronized to the collection of the time records used in the average. Since the deterministic part of the response signal (the part driven by the known stimulus signal) is repeated from record to record, it is reinforced as each new time record is added to the average.

In contrast, random noise in the measurement does not repeat from one time record to the next and tends to cancel as the averaging progresses. With each new time record, the signal-to-noise ratio increases. Time averaging reduces random noise by 10 log N, where N is the number of time averages.

The Uniform Window

Noise reduction isn't the only benefit of time averaging in this case. Another bonus is the potential for improving the quality of the frequency filter skirts in the FFT computation.

Each FFT record is a "snapshot" that represents a finite Fourier series based on an assumed infinite repetition of the same time record. This assumption is valid for a given signal only if it oscillates an integral number of periods within each time record (see figure 1). When the signal does not fit precisely into the time record, you have the potential for leakage, which shows up as misleading signal artifacts in the frequency spectrum.

The general tactic to minimize leakage is to use a Hanning or flat top window. This deemphasizes the discontinuity at the record boundaries. These windows are compromises, however, because you trade the unpredictable error of leakage for the predictable error introduced by the window shapes.

It's possible to avoid this compromise in some cases. If you control the stimulus frequency so that an integral number of periods fall within the time record, the window is unnecessary because you have met the FFT's periodicity and continuity assumptions.

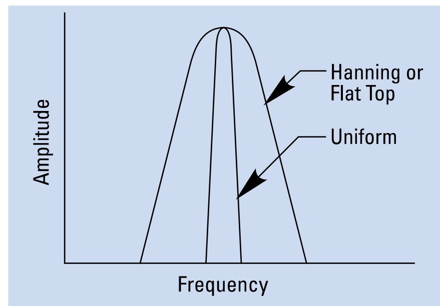

Figure 2: Benefits of the Uniform Window
The uniform window offers much narrower FFT filter skirts.

You can now use the uniform window (no window at all, actually) and still avoid leakage. The benefit of this technique is that the FFT filter skirts become very narrow, effectively improving the filter's selectivity (see figure 2).

How can you find a precise signal frequency that has an integral number of periods within the time record? You can calculate it based on the time record length, find it experimentally, or use the instrument's marker features (as explained in the sidebar).

Putting It to the Test

We recently put this averaging and window technique to work measuring floor displacement in a mid-rise, reinforced concrete building. We positioned an inertial vibration exciter horizontally on the fifth floor of the building (APS Dynamics Model 113). A sensitive accelerometer recorded the lateral motion of the same floor (on a north-south axis in this case). Figure 3 shows two measured responses generated by the same sinusoidal input force at 0.875 Hz. We used this very low excitation

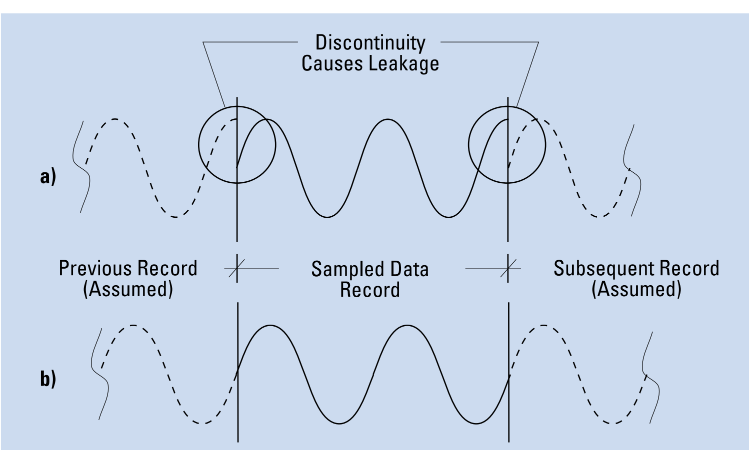

Figure 1: Periodicity and the Need for Windows
If the signal isn't exactly periodic in the time record, you'll have to apply a Hanning or flat top window to avoid leakage.

See It for Yourself

You can easily demonstrate the effects of time averaging and the uniform window. These steps apply to the HP 35670A, but you can adapt them to any DSA that has the necessary features.

1) Connect a cable between the source output and one input channel

2) Set the analyzer state as follows:

- FFT mode
- Frequency span as desired
- Measured data: power spectrum
- Uniform window
- Averaging type: (See below)
- Number of averages: 1000
- Averaging off
- Source on, Sine
- Adjust input range and source level to be compatible
- Trigger on the input channel connected to the source
- Trigger level: approximately 10% of source amplitude

3) Move the marker to any desired frequency

4) With the source menu displayed, select fixed sine then press the **Marker Value** hard key (this sets the sine output to the displayed frequency you've selected with the marker)

5) If the instrument is paused, press the **Start** key

You should now observe a very narrow filter skirt. Now select the Hanning or flat top window and compare the results. Also try entering a slightly different source frequency using the numerical key pad. The filter skirt should widen in both cases. (You may need to expand the x-axis scale to view the skirt width.)

Next, select RMS averaging and repeat steps 3 through 5. The noise in adjacent frequency bins should stabilize at some constant value. Finally, while maintaining the same source frequency, select the uniform window with time averaging. The noise in adjacent frequency bins should now gradually decrease as the time records accumulate.

frequency to quantify the building's stiffness or "spring-like" characteristics.

The dashed line in figure 3 is the floor displacement measured with a Hanning window and asynchronous RMS averaging (the primary alternative to time averaging). The coherence for this measurement was only 0.7 at 0.875 Hz, indicating a signal-to-noise ratio close to 0 dB.

Now compare the RMS averaging technique with the time averaging/ uniform window result. Despite the increased variance of noise amplitudes (inherent in time averaging), the time averaging improved the signal-to-noise ratio throughout the spectrum by about 14 dB. It also reduced noise at the excitation frequency. You can see that the true floor displacement measured with time averaging is 2.6 dB less than indicated by the noisier RMS averaging technique. (The coherence function is not available with time averaging, but you can assess the quality of a frequency response measurement by observing random variations in the phase angle as the time averages accumulate. Moreover, the response amplitude should continue decreasing until it stabilizes at the true value.)

The peak at -157 dB represents a displacement of 14 nanometers (roughly 2% of the wavelength of light), so you can see why noise reduction in our measurements is so crucial. Without the time averaging/ uniform window measurement technique, we'd be stuck with a questionable result. ■

Anthony Nash, PE, is with Charles M. Salter Associates in San Francisco. He can be reached at 415-397-0442.

For more information on the HP 35670A, check 9 on the Reply Card.

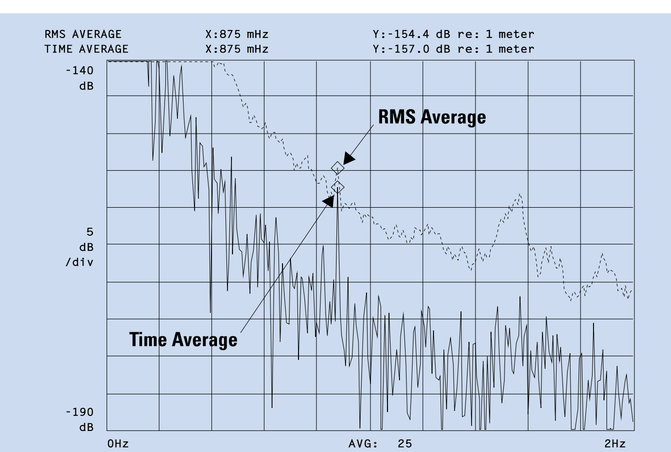

Figure 3: Comparing RMS and Time Averaging
Time averaging with the uniform window (solid line) provided a more accurate result than RMS averaging with the Hanning window (dashed line).

REALTIME Update

Issue: Winter '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

Editor
Rick Van Ness

Contributing Editors
Paul Gallagher
John Jensen
Joe Megaw

Manuscript Editing, Design and Production
The Dovel Group

Subscription and Fulfillment
Kathy Minor, Bunny Spaulding,
HP Direct Marketing Organization

Realtime Answers

Question: I am using an HP 35665A to do real-time 1/3 octave measurements. I then transfer the data to my PC. When I use the SDF utilities (especially the VIEWDATA program) to examine the data, I get peculiar center frequencies in my 1/3 octave bins. Is something wrong?

Answer: Nothing is wrong with your analyzer or with VIEWDATA; what you're seeing is a result of the way the 1/3 octave standards were written. The preferred 1/3 octave center frequencies are defined by two standards: ANSI S1.6-1984 and ANSI Z17.1-1973. These standards contain the equations for the computation of these frequencies. The resulting center frequencies are rounded off to become what we call "nominal" frequencies. The HP analyzers that offer real-time 1/3 octave (HP 3567A, 3569A, 35665A, and 35670A) display these nominal center frequencies. However, the analyzers use the true center frequencies from the ANSI equations for all internal calculations.

When you store measurement data, the analyzers store the actual data, including the true center frequencies. Transferring the data to a PC and examining it with VIEWDATA will give you a display of the true center frequencies. The displayed result meets all standards but no longer shows the rounded ANSI nominal values.

Question: I am using an HP 3567A with the optional internal 420 MByte SCSI disk. When I do a time capture I get better real-time rates if I capture to RAM. Is there any way to capture signals to RAM and then copy them to the internal SCSI disk?

Answer: The SCSI disk option to the HP 3566/67A provides access to many storage devices (magnetic disks, optical drives, tape drives, etc.) The HP 3566/67A internal RAM offers the fastest capture rates but is limited to a maximum of 16 MBytes. Combining time capture to RAM with backup to a SCSI device offers the best of both worlds.

The technique for doing this with the HP 3566/67A:

Click **Time Capture**
Click **Capture Mode**
Click **Time Capture to Memory/DOS**

This sets up the analyzer to do a time capture to internal RAM with a subsequent file transfer to the system disk on the host PC. Follow the appropriate instructions to do your signal acquisition. When you have successfully completed the time capture:

Click **File**
Click **Save Time Capture to DOS**

Follow the instructions to name and save your time capture on the PC system disk. Now you are ready to back up the file to your SCSI device.

Click **Capture Mode**
Click **Time Capture to SCSI Disk**
Click **File**
Click **Recall Time Capture from DOS**

Follow the instructions to identify the file you wish to recall. The file will be recalled to memory and copied to the SCSI device.

Question: What are the latest revisions for all of HP's dynamic signal analyzers?

Answer: Some of the DSAs have user-installable firmware updates; others do not. Here are the analyzers with user-installable updates, along with the latest revision numbers:

HP 3560A: Firmware revision A.00.02. You can see the revision number after Preset or at power on.

HP 3566/67A: Software revision A.03.02. You can see the revision number by clicking on **File** and then clicking on **About HP 3566/67A** on the host PC.

HP 3569A: Firmware revision A.00.03. You can see the revision number at Preset or at power-on.

HP 35665A: Firmware revision A.01.08. You can see the revision number by pressing **[System Utility]** and then the softkey [S/N VERSION].

HP 35670A: Firmware revision A.00.01. You can see the revision number by pressing **[System Utility]**, the softkey [MORE] and then the softkey [S/N VERSION].

There are no firmware updates for HP 3561A, 3562A or 3563A.

If you use the SDF Utilities, make sure you have the latest revision (B.02.00). You can find the revision number at any command by using a /U switch or by examining the contents of the REVID file on your SDF disks.

If you need a firmware update or a software revision, contact your local HP sales representative or contact us at *Realtime Update* (see below for contact information). Several of these firmware revisions are available free of charge. ■

How To Reach Us

Rick Van Ness, editor:
206-335-2654

Rick Van Ness
Editor, Realtime Update
Hewlett-Packard
8600 Soper Hill Road
Everett, WA 98205-1298 U.S.A.

Fax line: 206-335-2828

Internet address:
rick_vanness@hpa100.desk.hp.com

For more information on the HP products mentioned in this issue, please call 800-452-4844 ext. 7920