Vibration Testing of Small Satellites

This series of papers provides a tutorial along with guidelines and recommendations for vibration testing of small satellites. Our aim with these papers is to help you (a) ensure the test meets its objectives in demonstrating flight worthiness and (b) avoid test failures, whether associated with a design deficiency or with excessive loading during test. Addressed are sine-burst testing, random vibration testing, and low-level diagnostic sine sweeps. Although much of the guidance provided in this series applies to CubeSats, the series is primarily aimed at satellites in the 50 – 500 lb (23 – 230 kg) range. Most of the guidance applies to larger satellites as well if they will be tested on a shaker.

The plan is for this series to include seven parts, each of which will be released when completed:

1.	Introduction to Vibration Testing	(released April 11, 2014; last revised July 19, 2017)
2.	Test Configuration, Fixtures, and Instrumentation	(released April 11, 2014; last revised July 19, 2017)
3.	Low-level Sine-Sweep Testing	(released May 13, 2015; last revised July 19, 2017)
4.	Sine-Burst Testing	(released April 28, 2017; last revised July 19, 2017)
5.	Random Vibration Testing	(released April 7, 2016; last revised July 19, 2017)
6.	Notching and Force Limiting	(released May 13, 2015; last revised July 19, 2017)
7.	Designing a Small Satellite to Pass the Vibration Test	(yet to be released)

The most recent versions of these papers are available for free download at

http://instarengineering.com/vibration_testing_of_small_satellites.html.

Part 3: Low-Level Sine-Sweep Testing

Tom Sarafin, Poti Doukas, Lenny Demchak, and Mike Browning
Instar Engineering and Consulting, Inc.
Contact: tom.sarafin@instarengineering.com

Rev. A, July 19, 2017

Part 1 of this series introduces low-level sine-sweep tests and why we do them for small satellites. Here we explore low-level sine sweeps in greater detail to help you pick the frequency range, level, and sweep rate to best meet your objectives. We also aim to help you establish test criteria associated with comparing pre-test and post-test sine-sweep data and to interpret any differences you find in such comparisons. Finally, we address how to derive damping from the test data.

In a sine sweep, the shaker introduces sinusoidal acceleration at gradually increasing (or decreasing, less commonly) frequency. Input and response are plotted across the selected frequency range. If the test article has only one mode of vibration excited within that range, the response plot looks much like the transmissibility plot for a mass on a spring, shown in Fig. 1-2 in Part 1. The presence of multiple response modes essentially causes superposition of multiple transmissibility plots.

Figure 3-1 shows an example of acceleration data from sine-sweep testing of a small satellite. Note how, at low frequencies, well below the first mode, response is slightly higher than the input acceleration. As frequency increases, the gain increases, just as it does for a mass on a spring. Two closely spaced rocking modes[1] (154 – 165 Hz) are followed by isolation above 300 Hz until, at 500 Hz, another significant mode appears.

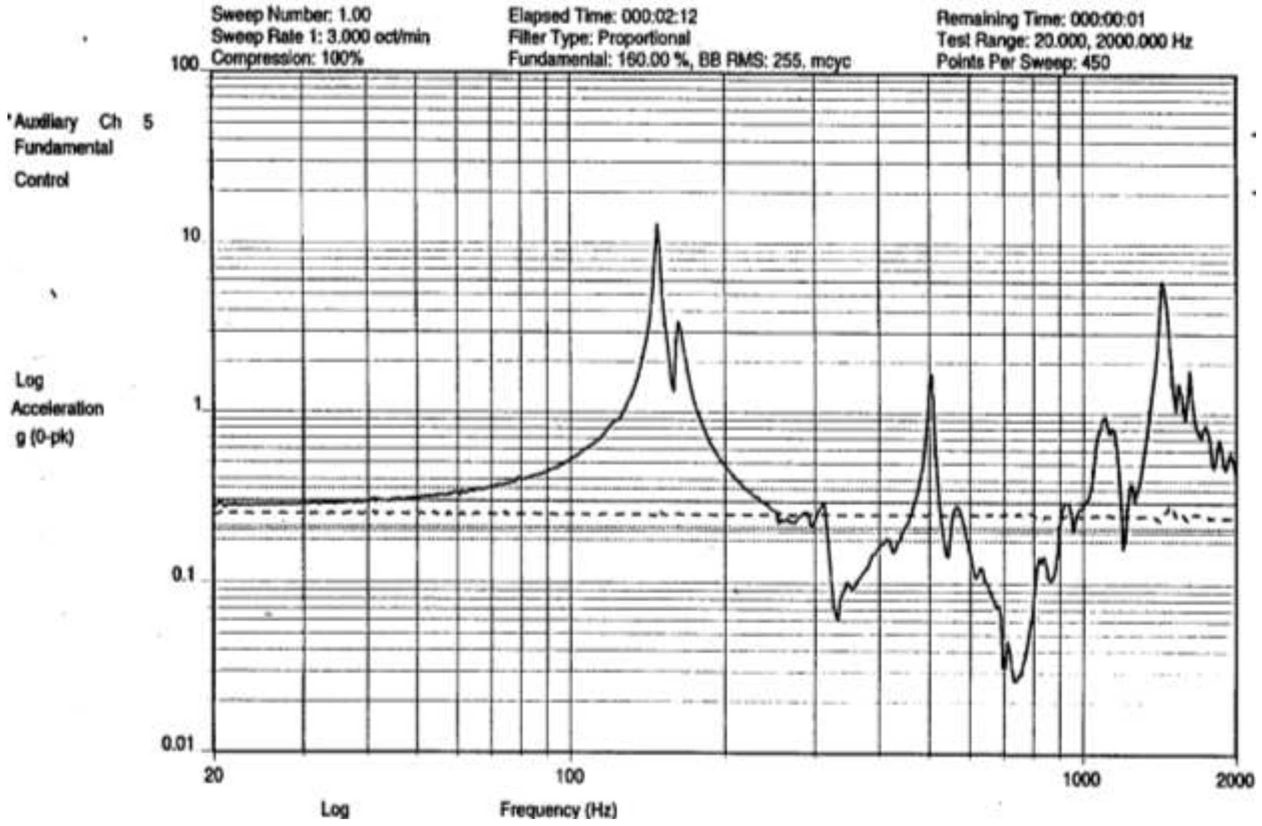

Fig. 3-1. Example Sine-sweep Response Data. This response plot is from a lateral test of a 43-lb satellite run at 0.25g swept-sine input (dashed line). The accelerometer was placed at the top of the structure, and response was measured in the same axis as the input. The first and highest peak (154 Hz) represents the fundamental lateral mode, which was rocking of the structure on the separation mechanism. The second peak at approximately 165 Hz is from rocking about an axis 90 degrees from the fundamental rocking axis. The presence of both peaks indicates the rocking modes are about axes that are skewed from the satellite's coordinate system.

Objectives

---

[1] In these two modes, the structure rocks laterally about two axes that are ninety degrees apart; the axes of rotation are skewed from the test axis so that both modes are excited by lateral motion in the test axis.

The objective of a low-level sine sweep is to acquire response data that helps us understand the test-article's modes of vibration, primarily natural frequencies but also damping and mode shapes. We can use this information to

- compare with applicable requirements, such as a lower limit on fundamental frequency;
- correlate a finite element model; or
- assess health of the test article by comparing response plots from pre-test and post-test sine sweeps (tests that are done before and after a random vibration or sine-burst test).

Sweep Rate

Sweep rate—the rate at which frequency increases for input vibration—is traditionally in the units of octaves per minute. Each octave is a doubling of frequency. For example, a rate of 2 oct/min means it takes one minute for the frequency to double twice, or to increase by a factor of four (two squared). Four octaves is a factor of sixteen (two to the fourth power).

The main reason to sweep at a higher rate is to decrease time in test. It takes twice as long to sweep across a given frequency range at 2 oct/min as it does at 4 oct/min. To sweep from 20 – 2000 Hz, a rate of 2 oct/min takes about 3 minutes and 20 seconds, whereas 4 oct/min saves half that time. If you do the sweep, say, three times per axis—once before random vibe, once after random vibe, and once after sine burst—the time savings is negligible between these two sweep rates. However, if you end up doing repeated diagnostic sine sweeps as part of anomaly investigation, the time savings adds up.

However, time savings usually does not balance the disadvantages of sweeping at 4 oct/min. We tend to interpret a peak response during a sine-sweep test as if it were data from a sine dwell at the natural frequency, corresponding to the peak in the transmissibility curve shown in Fig. 1-2. In such a dwell, the peak response occurs at resonance, when equal amounts of energy are coming in from the shaker and leaving the structure because of damping. Recall from Part 1 that, for a sinusoidally base-driven mass on a spring, Q is the peak response acceleration divided by the peak input acceleration when at resonance, and $Q \approx 1/(2\zeta)$, where ζ is the damping ratio. But true resonance is not achieved in a sine sweep, and the difference can be quite significant when the sweep rate is high and damping is low (high Q). Figure 3-2 shows, for a base-driven mass on a spring, how the sweep rate affects how the ***observed Q*** (ratio of peak response to peak input) differs from actual Q.

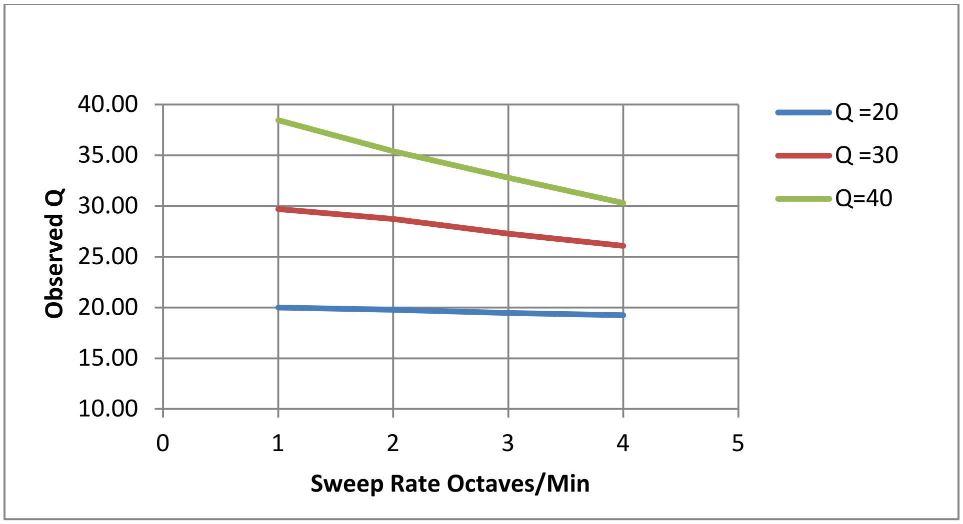

Fig. 3-2. How Sweep Rate Affects the Difference Between Observed *Q* and Actual *Q*.

Another effect of a high sweep rate is that the peak response is not as repeatable from one test to the next. When comparing pre- and post-test sine-sweep data, we want to be sure any differences in response plots are the result of changes within the structure, not variation from test to test, and a slower sweep rate causes less variation. In addition, when sweeping at a high rate we are more likely to see differences in response plots if we sweep forward, with increasing frequency (more commonly the case), or backward, with decreasing frequency. In some cases, a mode that will show up in a low-rate sweep will not be detected in a high-rate sweep.

When *Q* is below about 40, which is typically the case for a small satellite's high-mass modes, a good compromise in sweep rate is 3 oct/min. Data tends to be quite repeatable in such a case. However, to determine damping for model correlation, we'll need to be careful in how we interpret the data, as discussed later in this paper.

Frequency Range

Because random vibration tests in the space industry traditionally go from 20 – 2000 Hz, the same frequency range is often used for low-level sine sweeps. Sine-sweep data, however, is most meaningful for the first few modes. We tend not to pay much attention to higher-order modes because they tend to have less effective mass, and small changes in the structural assembly (e.g., slight loss in fastener preload) can make big differences in the pre- and post-test data. So we may choose to cut off the sweep at, say, 1000 Hz instead of going all the way to 2000 Hz. Doing so, though, won't save much test time. It takes just as long to increase the frequency from 20 Hz to 40 Hz as it takes to go from 1000 Hz to 2000 Hz.

Level

Another parameter we'll need to select for a sine sweep is ***level***, or amplitude of peak input acceleration. We want to select a level that is safe for the hardware yet high enough to give us meaningful data.

Many people use 0.25g for small satellites, but that level is more suited to spacecraft in the 50-lb class than the 500-lb class. 0.25g may sound harmless, but it actually can cause high structural loads in the test article, depending on its size. If the *Q* is 20 for the satellite's fundamental mode in a given axis, response to 0.25g swept sinusoidal input can be as high as 5g. The primary structure of a 500-lb satellite might be designed for a limit quasi-static load of about 10g, so a 0.25g sine sweep may cause stresses equal to 50%

of the limit stresses. Given all the loading cycles from "low-level" sine-sweep testing in the course of a vibration test program, a stress equal to half the limit stress can introduce unwanted and unnecessary fatigue damage to the structural materials.

The 43-lb satellite tested to produce the plot in Fig. 3-1 clearly had a Q higher than 20, as the response at the top of the satellite for the 154-Hz rocking mode hit a peak of over 12g for an input of 0.25g. (For a mode such as this, deriving Q from test data is not as simple as dividing response by input, even if the input is a dwell rather than a sweep; we need to account for mode shape, as discussed later in this paper.) This 43-lb satellite probably was designed for a limit quasi-static load of 25g and also had high margins of safety. Thus, 0.25g as an input level was acceptable for sine-sweep testing of this satellite.

A better level for satellites weighing more than about 100 lb—or smaller ones with high *Q* and less-robust design margins—is 0.1g. Most test equipment can control such low-level sinusoidal input quite well. If this is true for your test (check with the selected test lab!), the only downside to such a low level is that it won't stress the joints as much, so the *Q* measured during the sweep may be significantly higher than it is for higher-level loading. To keep the difference relatively small so that higher-level responses are more predictable while avoiding or minimizing fatigue damage, you'll probably want to increase the level above 0.1g for satellites weighing less than 100 lb.

Comparing Pre- and Post-test Sine-sweep Data

It's important to assess the health of the structure after completing a high-level random vibration or sine-burst test. The most useful tool for doing so is the low-level sine sweep (or low-level random vibration test, as discussed below). If nothing has changed in the structure, the response spectrum (often referred to as a "response signature") from pretest and post-test sine sweeps should overlay perfectly. In other words, if the pretest and post-test signatures differ, something has changed in the structure. Changes associated with the first few modes of vibration (lowest in frequency), which typically have the most effective mass, are of the most concern.

When one of the first few modes of vibration changes after completion of a high-level test, the frequency nearly always drops, and so does the peak response to a given input acceleration. The most common reasons for such changes are (a) loss of preload in threaded fasteners, (b) localized yielding in joints, and (c) fatigue cracks or full rupture.

a. Preload loss—Any loss of preload can cause a reduction in stiffness; the joint doesn't have to be loose, with all preload gone. Fasteners lose preload for several reasons, including relative rotation of the bolt and the nut or threaded insert, local yielding of the bolt or the clamped parts (e.g., washers) in the joint, and temperature change.

- Relative rotation of the threaded parts is mostly the result of cyclic lateral motion at the bolt threads and is most common when the joint slips back and forth within the limits of the fastener clearance holes. Most locking features do not prevent such loss of preload (ref. NASA-STD-5020 [1]).
- Local yielding causes a loss of preload when the total stress in the bolt or a clamped part during the high-level test exceeds the material's elastic limit. (Recall that yielding does not begin at the yield stress; it begins at the elastic limit, which is nearly the same as the proportional limit for most metals.) This is most often the case when the preload by itself stresses the material near the elastic limit; then the applied load during the test adds additional stress.

- Unless the materials making up the bolt, the washers, and the other clamped parts have the same coefficient of thermal expansion, a change of temperature causes a change in preload. Because temperatures of the structural materials typically don't change much between the pretest and post-test sine sweeps, this effect is not usually the reason for any loss in preload.

Torque striping is a good technique for determining whether the threaded parts are rotating relative to each other. After applying torque to the nut or bolt head at assembly, we paint a stripe on the nut or bolt head that continues onto the neighboring parts. (See Fig. 3-3.) After the high-level vibration test, we visually check to see if the stripe is broken. Of course, this technique works only if we have visual access to the fastener. For flight hardware, be sure to use paint with acceptable outgassing properties.

Fig. 3-3. Torque Striping. The stripe is painted after the nut is torqued. After vibration testing, a break in the stripe between the nut and the end of the bolt indicates relative rotation between these parts, hence preload loss. For a screw going into a blind hole, the stripe must continue onto the clamped structural part, as shown above for the nut. (Image from americanprecisionassembly.com)

As noted above, local yielding of the bolt or the clamped parts also can cause preload loss. In this event, preload is lost even if a torque stripe remains unbroken.

To check whether preload loss is the cause of a change in pre- and post-test sine-sweep data, retorque fasteners and then repeat the sine sweep. If preload loss is the cause of the modal change, the response plot should closely match the pretest response plot. Again, this is an option only for accessible fasteners.

b. Localized yielding—Yielding is often thought to cause natural frequencies to drop, but this is not usually the case in a random vibration test. If a random, high-level loading cycle causes a material to yield, that material still has the same modulus of elasticity as it had originally, so stiffness of the structure is unchanged for lower-level loading cycles. Stiffness decreases only temporarily at and above the stress that causes yielding.

At low stress levels, which is the case for a low-level sine sweep, the stiffness of a structure, hence the natural frequency for a given mode, changes as a result of prior material yielding only if the yielding changes the load path. This can be the case for certain locations of yielding with fastened joints. Fastener yielding causing preload loss is one example, as noted above. Figure 3-4 shows another example. Another example is a preloaded bolt that yields under applied tensile load. Such yielding causes a loss of preload, hence less clamped material and joint that is less stiff. This effect is normally not significant unless the joint loses a large percentage of its preload.

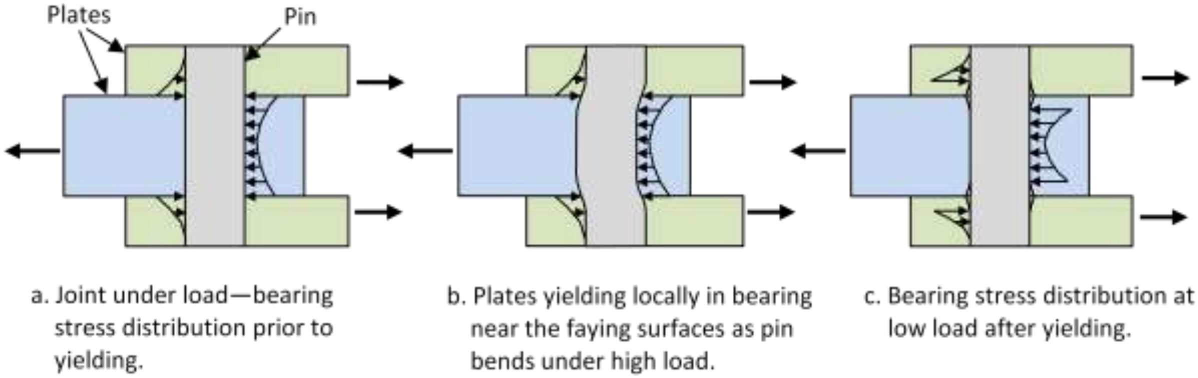

a. Joint under load—bearing stress distribution prior to yielding.

b. Plates yielding locally in bearing near the faying surfaces as pin bends under high load.

c. Bearing stress distribution at low load after yielding.

Fig. 3-4. Example of Local Yielding that Reduces Stiffness. In this lug-and-clevis joint, local yielding of the plates in bearing under high applied load causes a change in the load path and different distribution of bearing stress. As a result, the pin is more flexible in bending.

c. Fatigue Cracks and Rupture—When a ductile material fails during a random vibration test, the failure is almost always from ***fatigue***: crack formation and growth as a result of cyclic stress. A single random cycle of high stress causes the material to yield before it ruptures, thus absorbing energy and limiting the stress such that it doesn't reach the rupture stress. Cyclic high levels of stress, however, can cause a crack to form and then grow in the region of highest stress.

As a fatigue crack grows, natural frequencies can drop, especially for the modes that stress the material most near the crack. Figure 3-5 shows a comparison of pre- and post-test sine-sweep data for a small-satellite engineering model; a fatigue crack caused the difference.

It's important to understand whether a significant drop in natural frequency is the result of fastener issues (preload loss or local yielding, as noted above) or a growing fatigue crack. With continued loading cycles—e.g., during launch—a fatigue crack can continue to grow until the part ruptures, sometimes catastrophically. If retorquing the fasteners doesn't increase the natural frequencies approximately to the original frequencies, post-test inspection for cracks is prudent.

Sometimes a structural part fully ruptures without causing catastrophic failure of the structural assembly. In other words, the test article continues to withstand the test environment after one part ruptures. This may be the case for a ***statically indeterminate*** (***redundant***) structure, which is one with multiple load paths. Such a part failure usually causes natural frequencies to drop and mode shapes to change. Even though the structural assembly may still withstand the environment, we probably would conclude it did not pass the test. Debonding between plies in a composite laminate is an example of material rupture (rupture of the adhesive) that may not propagate to full catastrophic rupture of the structure and, in some cases, may not be considered failure to pass the test, depending on the consequences of debonding.

As noted above, for parts made of ductile materials, full rupture in a random vibration test is usually the result of fatigue. For brittle materials, however, rupture is caused by the highest load—a strength failure. (Such a failure is also referred to as a ***single-passage failure***.) A brittle material can rupture from a random high spike of load during a random vibration test or from the peak acceleration during a sine-burst test. Ductile materials also can rupture from a single cycle of dynamic load, but normally this is the case only for failure modes in which the weak cross section is of short length such that yielding exhibits little plastic displacement. Otherwise, yielding may prevent some of the mass from seeing the full acceleration. As a result, the structure may not fail, but we may not consider the test to be adequate. Part 4 of this series addresses this topic for sine-burst testing.

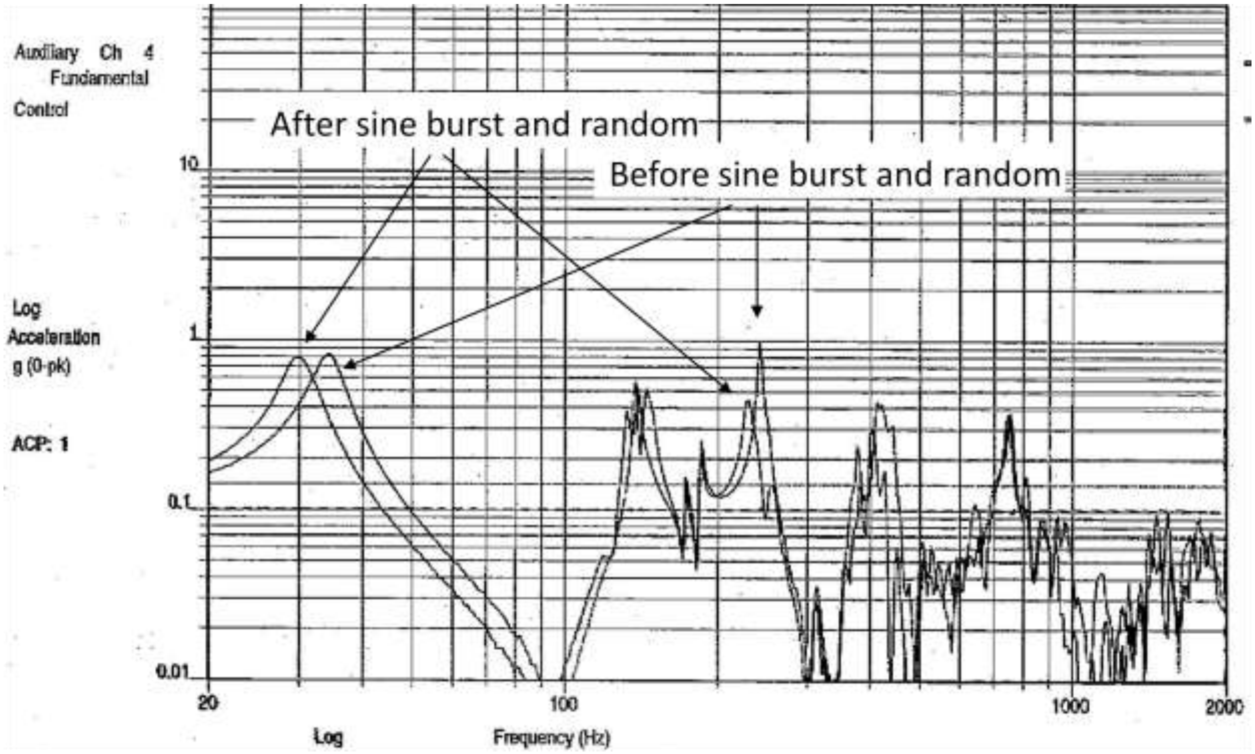

Fig. 3-5. Example Comparison of Pre- and Post-test Sine-sweep Data. The first mode, a lateral rocking mode, started out at 34 Hz and dropped to 30 Hz after the sine-burst and random vibration tests. Post-test inspection revealed a fatigue crack near the mounting base.

Establishing Criteria for Comparisons of Pre- and Post-test Sine-sweep Data

Small changes between pre- and post-test sine-sweep data, such as from localized yielding in fastened joints or slight loss of fastener preload, are unavoidable in many structures and usually are not detrimental. Prior to the test, we should establish criteria associated with the sine-sweep comparison for determining whether investigation is needed. Such criteria should depend on type of structure and risk tolerance of the program. Table 3-1 lists suggested criteria.

Such criteria should apply only to the modes showing the highest response acceleration in a given test, using the accelerometer channels that show the highest response for those modes. Changes in low-frequency, high-mass modes imply something has changed in the primary structure, whereas changes in higher-order modes imply something has changes in secondary structures or their mounting interfaces.

TABLE 3-5. Suggested Criteria for Comparison of Pre- and Post-test Sine-sweep Data. These criteria apply to the fundamental mode; less stringent criteria should be used for higher-order modes. See text.

Type of Structure	Stop the test and investigate the anomaly …	
	if the first natural frequency drops by more than …	or if the response peak drops by more than …
"Loose"—a structure with many joints using threaded fasteners and oversized fastener holes	5%	30%
"Tight"—a structure with few or no mechanical joints or with joints welded, bonded, or pinned (with interference-fit shear pins)	2%	15%

As an example, consider the test data shown in Fig. 3-1. This structure was all metallic, with structural parts bolted together. The bolted joints were relatively lightly loaded, with high margins of safety and high preload, which means we would expect little if any slipping of joints within the fastener clearance holes. The spacecraft structure was mounted on its separation mechanism (band-clamp type), which was bolted to the test fixture. Because of the separation mechanism, we would consider this test article a "loose" structure; separation mechanisms tend to have stiffness that is nonlinear and not repeatable. For this structure and this accelerometer channel, we would apply the following criteria for stopping the test for investigation:

- First mode (154 Hz): 5% change in natural frequency or 30% drop in response peak
- Second mode (162 Hz): no criteria (off-axis mode that is not well pronounced—small changes in a structure can cause two closely spaced modes such as these to merge into one)
- 500-Hz mode: 5% change in frequency or 50% drop in response peak
- No criteria for the other, higher-order modes

We recommend you formalize such criteria by documenting them in the test plan and adding procedural steps to assess the criteria and document conclusions after each post-test sine sweep. Doing so adds discipline to the test process. It is far easier to investigate anomalies while the test article is still in test configuration than it is after tearing down the setup!

A key point here is that sine-sweep comparisons are not always sufficient for judging whether a satellite passes its vibration test. Final confirmation requires post-test performance testing. This is particularly the case for alignment-critical systems, but it applies for other aspects of performance as well. The satellite should receive full testing afterwards to ensure it still functions properly.

Reference

1. NASA-STD-5020. "Requirements for Threaded Fastening Systems in Spaceflight Hardware." March 12, 2012.