



Log Swept Sine Script Notes

Version 130b6 October 2010

Ian Heaton, Applications Engineer, Prism Media Products Ltd.



Table of Contents

Introduction.....	3
Overview.....	3
Features.....	3
Pre-requisites	4
Installation.....	5
File List.....	5
Un-installing the scripts.....	5
Connections.....	6
Connection Notes:.....	7
Running The Script.....	8
Scripted Interface Notes.....	8
Scripted User Interfaces.....	8
Scripted User Traces	9
Getting Started.....	9
Setup.....	10
Sweep Settings.....	10
Measurement frequency range (from start frequency to end frequency in Hz).....	10
Generator Level.....	10
FFT Size.....	11
Sample Rate.....	11
The "Apply" Button.....	11
Sensitivity Range.....	11
Microphone Polarity.....	12
Results Section Settings.....	12
Distortion.....	12
THD.....	12
Acoustic Phase.....	12
Impedance (magnitude and phase).....	12
Smoothing.....	13
Calibration.....	14
Microphone Calibration.....	14
Automatic Microphone calibration:	14
Manual Microphone Calibration	15
Impedance Calibration	16
File Management.....	17
Background.....	17
Practical File Management.....	17
Create a new file.....	17
Open a file.....	17
Save a file.....	17
Remove a file.....	17
File Export.....	17
Technical Notes.....	18
Log Swept Sine Stimulus.....	18
Derivation of an Impulse Response.....	18
FFT Windowing and anechoic measurements.....	20
Distortion Harmonics.....	21
Distortion Normalisation.....	21
Distortion Normalisation Summary:.....	22



Introduction

These notes are to accompany the dScope script “LSS Automation 130b4.dss” – an automation script to configure and manage the dScope Series III for measurements acoustic transducers.

Overview

The dScope Series III is capable of measuring a wide range of parameters directly from its user interface, however, for some tasks this tends to be quite low level and requires a lot of configuration. This automation script is intended to manage many of the settings and configuration options of dScope in order to make a range of measurements appropriate to acoustic transducer testing. It also goes further than this by managing the data generated by dScope (graphs etc.) and by calculating other parameters from these data. It does this by using a scripted interface to set and control dScope settings and automate some of the more difficult procedures so as to make microphone measurements more readily accessible.

The script here is intended to measure an electrical to acoustic transducer, such as a loudspeaker, by using the Log Swept Sine sweep (Farina) Method.

Features

This script is intended to measure the following:

- **Frequency Response**
 - Quasi-anechoic response derived from windowed impulse response
 - Measured from the impulse response of the fundamental (excluding distortion components)
 - Raw FFT spectrum or smoothed response
 - In calibrated units of dBSPL
- **Distortion Harmonics vs. Frequency**
 - Harmonics 2 up to 8 can be plotted against frequency
 - Optional frequency normalisation
 - Raw FFT spectrum or smoothed
- **THD trace**
 - Power sum of individual distortion harmonics
- **Sensitivity**
 - Calculated over a user defined range of frequencies in dBSPL 1W/1M
- **Polarity**
 - Transducer polarity as determined from the impulse response
- **Acoustic Phase**
 - Acoustic phase in degrees with the phase shift due to the time-of-flight removed
- **Impedance Magnitude and phase**
 - Electrical Impedance magnitude and phase against frequency
 - Derived values of minimum impedance magnitude and resonant frequency



Pre-requisites

The script requires dScope Series III software version 1.30a on a computer platform supported by this software (Windows 98, 2000, XP or Vista 32). In addition, the following will also be required:

- an audio power amplifier
- a series resistor and calibration resistor for impedance measurements
- a microphone pre-amplifier
 - dScope does not supply phantom power, and a microphone pre-amplifier is usually necessary to power the microphone and increase the level.
- a calibrator / pistonphone for calibrating the microphone (optional).



Installation

The files are supplied in a zip archive with a script installer which copies files to the required locations. This script reads the dScope installation folder location from the registry, then installs the necessary files, creating new folders if and where necessary.

To install the files, unzip the contents of the zip file to an empty folder (location is not important) and double click on the file "install test files.vbs". This may get stopped by anti-virus software as it uses the registry to get the path to the dScope folder. If so, you will need to tell the anti-virus software to allow the script to run. Once it has run, it generates a text file called "report.txt" which lists the actions taken by the script, including where it has installed files. The files in this instance are:

File List

File Name	Description	File Location*
LSS Automation 130b4.dss	Main dScope automation script	\scripts\automation\
LSS default.dsc	Default dScope configuration	\configurations\
LSS Test Notes 130b4.pdf	This document	\scripts\automation\
LSS default.xml	Script configuration file	\configurations\LSS\

*Location is relative to the dScope installation folder

Un-installing the scripts

There is no un-installer. To remove the scripts, you will need to delete the files installed as listed in the installation report. These are also listed above.

The process of installing the files makes no changes to the registry.

Connections

This script is set up and configured for analogue output and analogue input only. In order to make acoustic measurements, the connections below should be used:




Figure 1) Basic acoustic measurement connections

In order to measure impedance and phase, an additional series resistor and a connection to the second channel of the dScope is required as below:




Figure 2) Full measurement connections



Connection Notes:

- The dScope Series III analogue inputs do not have phantom power and so most condenser or electret microphones are going to need an external phantom power supply.
- The use of a microphone pre-amplifier will help to get the levels more appropriate for the dScope inputs and improve the measurements.
- The choice of series resistor is a bit of a trade-off between measurement performance and influence on the damping of the loudspeaker: Too small a value results in a very low voltage drop across the resistor and results in noisy, unrepeatable measurements. Too large a resistor results in a test circuit that will have an undesirable effect on the damping of the loudspeaker. In practice a 0.1 Ohm resistor is a common compromise.




Running The Script

Once installed, the script is run from the dScope user interface by clicking on “Automation” followed by “Run Script” and navigating to the script “LSS Automation 130b1.dss”. Alternatively, you can click on the “run script” icon on the dScope toolbar and select the script from the list.

To make the script easier to run repeatedly, it can be set up as a button on the dScope user bar. See the dScope help topic entitled “Customize User bar dialogue box” for more information.

If the currently loaded configuration is set up to work with the log swept sine script, the script will keep the configuration, retrieve its settings and proceed as normal. Note that if the configuration has been changed since it was last saved it is not re-loaded and this may cause problems if the script and the configuration are no longer in sync. If the current configuration isn't set up to work with the script, the script will load a default configuration.

Running the script presents the user interface shown below:



This contains a series of buttons to perform a range of functions described later and a status field at the top which gives feedback about the current status of the script. Below these are a set of reading windows which display measured parameters after a test has been performed. Not all measurements are always active, in particular the bottom four parameters depend on the impedance being measured and will therefore only be updated if the “Plot impedance” option is selected. Note also that the accuracy of these measurements may depend on the impedance calibration. At the bottom are some buttons for managing different configurations.

Scripted Interface Notes

In order to make the most of this interface, it's worth taking the time to understand what this interface is doing and how it works.

Scripted User Interfaces

This script's user interface shown above is not part of the dScope user interface – it is part of an “Automation Script” which accesses dScope's methods and properties through its automation interface. It is written in VBScript and uses a DLL called “ScriptDlg” (Script Dialogue) to create the user interface.

The script can set dScope parameters and read them using dScope's ActiveX controls, but it is not dynamically linked. This means that changing a parameter in the dScope interface does not automatically change it in the script. The script has to actively check for changes – something it does not do continuously. It checks some parameters when a test is run or when some other action such as calibration is performed.



The practical implication of this is that you may not be able to get away with changing settings in the dScope interface directly and expect the scripted results to still work. In general, if you can make a change in the script interface or on the dScope, you should use the script interface as this keeps the two in sync. Changing units in the dScope interface may also have unpredictable effects on the results obtained from the script since the script reads some parameters from the dScope in the currently selected units and will be expecting the returned results to be in a particular unit.

Scripted User Traces

This script makes extensive use of “user traces” which are traces that are created using the dScope automation interface. Each trace is created with a number of points, particular units, ranges etc. When the script configuration is changed, if the change means that more or less traces are required or the traces must have a different number of points, different units, or different ranges, the script removes all the user traces from the interface and re-creates them with the new settings. This has the side effect that any copies of user traces that are created are also removed. For this reason, if you need to keep copies of traces and need to change settings such as the number of FFT points, fractional octave spacing, etc, you should save your traces to file first as any copy traces will be removed.

Things that cause the traces to be regenerated in this instance include:

- Changing the number of harmonics that are displayed
- Enabling or disabling the smoothing of traces
- Changing the FFT buffer size
- Changing the Sample Rate

Once the script has created the user traces, it avoids scaling them and turning them on and off. In most cases you can change the trace range that is visible manually and it will not change it back.

Since the traces created by the script are user traces (and not dScope native traces), you can not change the units – they are fixed when the script is created.

Getting Started

Once you have connected the dScope and other equipment as described in the “connections” section above, you will need to configure the test to get the measurements you require.

When you first run the script, it will load a default configuration. This has a generator level of 200mV so any connecting amplifier and speakers should be set to an appropriate gain level before you run a test to avoid damage. If the level is OK, you can immediately press the “Run Test” button and make a measurement. It is unlikely however that this measurement will be useful “out of the box” as it will not be calibrated and will not be set up to work with your amplifier, microphone etc. For that, it will need a bit of setting up. Read on!



Setup

Pressing the “Setup” button brings up a separate window which is used to set up the dScope for acoustic measurements. Clicking the OK button will apply the settings but not save them. Settings changed here are NOT saved with the configuration until the configuration is saved using the “Save” or “Save As” buttons on the script interface.

The setup dialogue is shown below:

Sweep Settings

Within the top section of the setup dialogue are the following settings:

Measurement frequency range (from start frequency to end frequency in Hz)

- This setting is for the frequency range over which the measurements are plotted.
- It does not change the frequency range of the stimulus which is set at 1Hz to the Nyquist frequency of the currently set sample rate.

Generator Level

- This is the RMS voltage level of the Log Swept Sine stimulus at the analogue output of dScope with the output unloaded.
- Since the output will be going through an amplifier, this is not the level seen at the loudspeaker terminals.
- If you change the generator level, you should use the “Apply” button to allow the dScope to set appropriate gain ranges. Likewise, if you change power amplifier or microphone pre-amplifier gain settings, running the “Apply” routine allows the script to optimise the gain settings.



FFT Size

- The size of the FFT buffer applies to both the generator wavetable and the analyzer sample buffer and has a direct impact on the duration of the sweep and the frequency resolution of the resulting measurements.
- The size is given in samples and the available values are powers of 2 (eg, 2^n where n is an integer between 12 and 17 inclusive).
- Changing the FFT Size means that the stimulus must be re-calculated. If this is done, you are required to press the “Apply” button in order that the dScope can check the settings and configure the analysis correctly

Sample Rate

- This determines both the dScope generator and analyzer digital audio sampling rate in kHz. Only two rates are supported 48kHz and 96kHz.
- 96kHz gives faster tests, with a higher upper frequency limit, but with reduced frequency resolution
- 48kHz gives slower test times and lower upper frequency limits, but with increased frequency resolution.
- Changing the sample rate means that the stimulus must be re-calculated. If this is done, you are required to press the “Apply” button in order that the dScope can check the settings and configure the analysis correctly.

The “Apply” Button

The “Apply” button does a lot more than just change the settings. It runs through an automated sequence which:

- re-calculates the generator stimulus if needed,
- checks that the analyzer levels are set appropriately
- Measures the delay in the system and compensates by adjusting the generator sequence
- Places the FFT window into edit mode to allow you to adjust the position of the FFT window.

The “Apply” button should be used whenever:

- The generator voltage is changed
- The FFT size is changed
- The sample rate is changed
- The amplifier gain is changed
- The microphone or microphone pre-amplifier gain is changed
- The distance between the speaker and microphone is increased or processing is introduced that may cause additional delay in the system.
 - A drop in high frequency response may be a symptom of the capture buffer not containing the end of the sweep. Try using the “Apply” button to allow the dScope to work out how much delay to allow for.

Sensitivity Range

- The sensitivity range is the range of frequencies over which the calculation of sensitivity is performed. Set both the start and end to the same frequency (e.g. 1kHz) to measure at a specific point.



Microphone Polarity

- The microphone polarity radio buttons allow the script to compensate for measurement microphones that output a voltage that is either proportional or inversely proportional to pressure.
- This is used in the determination of the polarity of the transducer under test.

Results Section Settings

This section allows you to select which results to plot.

Distortion

- This is used to set how many distortion harmonics are measured.
- The value set is the highest harmonic that will be measured.
- The distortion can be measured in dB or percent.
- The distortion trace frequency range will typically be smaller than the fundamental for two reasons:
 - The upper frequency limit of the trace will be progressively limited by the bandwidth of the measurement
 - The lower frequency will be progressively limited by the size of the FFT window that can be applied over the harmonic impulse without including adjacent harmonics. (The choice of FFT window follows the selection used for the fundamental)
- The resulting distortion traces can be plotted “frequency normalised” .
 - When the “Normalised” checkbox is checked, the distortion harmonics displayed are shown in the conventional manner whereby the harmonics are calculated relative to the level of the fundamental that caused them and plotted at the frequency of the fundamental
 - When the “Normalised” checkbox is NOT checked, the distortion harmonics are calculated relative to the fundamental at the same frequency as the actual frequency of the distortion component and plotted at that frequency.
 - For more details see the notes on distortion normalisation in the notes section at the end of this document.

THD

- Total Harmonic Distortion (not THD+N as noise is largely removed)
- This is a trace generated using a power sum of the distortion harmonics that are currently being calculated.
- Relativity and normalisation follows the selection in the distortion harmonics section.
- This can be displayed in dB or percent

Acoustic Phase

- The dScope natively only measures phase +/-180 degrees. The script attempts to remove the phase due to the time of flight and unwrap the phase plot and may sometimes get this wrong.

Impedance (magnitude and phase)

- This can only be measured when the system is set up to measure impedance and the impedance has been calibrated. See the next section for more about this.
- Impedance is calculated by measuring the voltage across a series resistor in the presence of a known calibration resistor and storing a calibration coefficient and phase



reference for each frequency of the FFT. When the calibration resistor is replaced by the speaker under test, the script can work out the impedance magnitude of the speaker at each frequency and the phase difference and plot them.

- The calibration is specific to each configuration and will change if any of the following change:
 - The generator voltage is changed
 - The FFT size or window is changed (calibration gets automatically disabled)
 - The sample rate is changed
 - The amplifier gain is changed
 - The series resistor is changed.
- In general it is a good idea to re-calibrate the impedance regularly.

Smoothing

An FFT measurement inherently has data points at linear frequency intervals, but quite often we would like to see the data smoothed in octaves or fractions of an octave. The smoothing option processes the FFT data points to extract only those near fractional octave centre frequencies and apply smoothing to them. At low frequencies with small FFT buffers the spacing of the frequency points becomes insufficient to provide separate FFT points for each fractional octave point and the matching between the two becomes increasingly inaccurate. At some point the same FFT point will be mapped to more than one fractional octave. In this case the data will still be plotted, but there will be a message at the bottom of the configuration menu to indicate what frequency the smoothing algorithm had separate FFT points from.



Calibration

Microphone Calibration

In order for the dScope to display measurements in acoustic units (dB SPL re 20 μ Pa), it must be told the sensitivity of the microphone and pre-amplifier that is driving it. There are two ways to achieve this: automatically and manually.


Automatic Microphone calibration:

In order to calibrate the microphone and pre-amplifier automatically, a simple calibration routine must be run. To do this, you will need an appropriate acoustic calibrator for the microphone you are using. To run the calibration, from the script user interface, click the “calibration” button. The dialogue below will appear:

A screenshot of a software dialog box titled "LSS Tests - Calibration". The dialog has a blue title bar with standard Windows window controls. It contains two main sections. The first section, "Calibration Menu", has a radio button selected for "Use mic calibration:". Below this, it shows the last calibration date and time: "Last calibrated 19/02/2010 12:19:17 using 94dB SPL calibrator resulting in 10.1mV at the dScope input." There are input fields for "Calibrate using:" (114) and "dB SPL at" (250) Hz, with a "Calibrate" button to the right. The second section, "Use mic manufacturer's data:", has an unselected radio button. It includes input fields for "Mic sensitivity" (12) mV/Pa and "Pre-amp gain" (0) dB gives. Below this is an "Impedance calibration" section with an unselected radio button. It shows the last calibration date and time: "Last calibrated 01/04/2010 13:47:53 using 3 Ohm Series resistor and 32 Ohm calibration resistor." There are input fields for "Series resistor" (0.1) Ohms and "Calibration resistor" (4.7) Ohms, with a "Calibrate" button to the right. At the bottom are "OK" and "Cancel" buttons.

Selecting the “Use mic calibration” radio button allows you to use an automated microphone calibration routine. To use this, enter the output level of the microphone calibrator you are using and its frequency. The frequency is simply used to check that the signal is received correctly during the calibration process.

When you are ready, click “Calibrate”. You will be prompted to attach and turn on the calibrator as shown on the next page / below.



Allow plenty of time for the calibrator to settle and stabilize. When this has happened, press “OK”. The dScope will make a measurement of the incoming signal and, in conjunction with the level set in the user interface, will work out and store the conversion necessary to display the acoustic sound pressure level in units of dB SPL. The values acquired in this way are saved when the calibration is performed and are kept with the configuration and not globally.

Manual Microphone Calibration

If you do not have access to an appropriate microphone calibrator for your microphone, you can enter the manufacturer's sensitivity values manually, as well as the gain of any pre-amplifier you are using. To do this, from the script interface, click “Calibration” and then select the “Use mic manufacturer's data” radio button

Enter the microphone sensitivity in mV/Pa and the pre-amplifier gain in dB. Note that the pre-amplifier gain on the calibration page is NOT set in the dScope “pre-amplifier gain” field in the Signal Analyzer panel. This is because the dScope uses the pre-amplifier gain to work out the level out of the




microphone, and that is not what we want in this instance. See the dScope Operation Manual help topic on the “Signal Analyzer” in the “References” section to understand how dScope uses the pre-amplifier gain parameter. The settings you enter are applied when the “OK” button is pressed and are stored with the configuration not globally.

Impedance Calibration

In order to make impedance measurements, it is necessary to first calibrate the setup that is being used. In the calibration menu are edit fields to enter the series resistor you are using and the calibration resistor.

- The choice of series resistor is a bit of a trade-off between measurement performance and influence on the damping of the loudspeaker: Too small a value results in a very low voltage drop across the resistor and results in noisy, unrepeatable measurements. Too large a resistor results in a test circuit that will have an undesirable effect on the damping of the loudspeaker. In practice a 0.1Ohm resistor is a common compromise.
- The calibration resistor should be of a similar value to the nominal impedance of the drive unit you are measuring.
- Ideally you should avoid inductive wire wound resistors and use thick film types of sufficient power handling that they will not heat up or be damaged by the currents passing through them.

When you have set the values for the resistors in use, connect the calibration resistor in place of the test speaker as shown below:



Press the “Calibrate” button and follow the instructions on screen. Calibration is quick and easy, just requiring one sweep in order that the dScope can capture and save the required coefficients.

Calibration values are saved to file when the calibration is performed. The interface will show when it was last calibrated and what the resistors in use at the time were. The calibration is specific to each configuration and will change if any of the following change:

- The generator voltage is changed
- The FFT size is changed (calibration gets automatically disabled)
- The sample rate is changed
- The amplifier gain is changed
- The series resistor is changed.



File Management

Background

This automation script makes use of parameters that are not stored in the normal dScope configuration and therefore must be managed by the script. It is important to use the script file management controls or the parameters controlled by the script will not be saved. When you save a configuration from a script, it saves both the dScope configuration and a separate script settings file.

Practical File Management

Create a new file

The easiest way to create a new file for use with the Log Swept Sine script is to use the “Save As” button to create a copy of a known configuration with a new name. This stores the original dScope configuration and copies of the script settings and impedance calibration files (if present) with the new name. All the scripts and configurations are stored in the dScope “configurations” folder and its sub-folders. It is not possible to save to other locations.

Open a file

Use the script “Open” button to select from the configurations that have been set up to work with this script.

Save a file


Save a configuration with its script settings using the “Save” button on the script interface. While saving a configuration with the dScope controls will save the configuration, it won't save the script settings and this can allow the two to get out of sync.

Remove a file

To remove a file that you no longer want, use the “Open” dialogue to list the script files, select the file that you want to remove and click “Remove”. This removes not only the configuration, but also the script settings file and any calibration files.

File Export

After a measurement has been made, the data can be exported to file using the “Export” button on the main user interface. This can export either the acoustic data (magnitude and phase) or the impedance data (magnitude and phase) to a variety of file formats for import into other programs such as Excel or loudspeaker design packages. Files are exported to the dScope “Graph Exports” folder.





Technical Notes


Log Swept Sine Stimulus

The Log Swept Sine stimulus consists of a sine wave whose frequency is swept logarithmically (in a particular rate of octaves per second) from low to high frequency. For example, if the sweep is set up to run at 1 octave per second, every second the frequency of the sweep would double (an octave being a doubling of frequency). Since distortion harmonics occur at multiples of the fundamental, with a log swept sine, as the fundamental frequency increases at a rate of octaves per second, so the harmonics will also increase at the same rate of octaves per second. This is only true of logarithmically swept sines and has important implications for deriving the distortion harmonics as discussed later.

Another by-product of the logarithmic sweep rate is that the signal spends disproportionately long at low frequencies, and shorter at high frequencies. This gives the log swept sine a pink weighted energy spectrum.

Derivation of an Impulse Response

In the screen shot below, the upper trace is an analyzer buffer capture of a log swept sine as played through a reverb unit. The lower trace is the original generator stimulus fed directly to the analyzer.




If we were to perform a cross correlation between the two buffers, we would get the impulse response of the system. The cross correlation can be thought of as taking the two waveforms as shown above and wrapping them around a cylinder so that they are one above the other. You then cut the cylinder in half horizontally so that the top trace can rotate separately from the bottom trace. Starting with the two traces aligned, you then rate them for how similar they are as you rotate them. For most positions of the two traces they will be dis-similar, but at one point they will line up quite well and have fairly similar




shapes. If we make the “rating of how similar they are” a bit more scientific, at each rotational point we can multiply all the points on each trace with the points from the other trace that are lined up with it and add them all together. It turns out that for most rotational steps this sum will result in a very low number as the positive and negative values tend to cancel each other out. At one point, however, this calculation will result in a clear spike or impulse. If the polarity of both waveforms is the same, this will be a positive going impulse, if not, it will be negative going. In practice, dScope does not perform a cross-correlation, but does the calculations in the frequency domain using FFT and inverse FFT calculations which is computationally faster and gives the same results.

The number of milliseconds by which the traces have to be rotated before they line up gives the position of the impulse and corresponds to the delay in the system. In the above example it happens that the top trace is about 2.3ms later than the bottom one: the latency through the digital delay is 2.3ms.

In addition, any reflections of the waveform (caused by delay taps within the reverb or by echoes in a real world system) appear as copies of the swept sine waveform at a lower level and delayed relative to the original. These are very hard to spot in the above diagram because they are superimposed on the original waveform. When the cross-correlation is performed, however, these show up as distinct impulses after the original impulse. These correspond to the impulse responses of each echo or reflection. The impulse response resulting from the two buffers of data shown above is shown below:



There is a very strong correlation near the beginning of the buffer which is the direct sound through the reverb. The early reflections appear as distinct impulses at periods some time after the initial impulse, and there follows a reverberant “tail” made up of many reflections and reflections of reflections overlaid on each other and decaying with time. This is seen more clearly if we zoom in:




In the screen shot above, we have zoomed the Y axis so that most of the initial impulse is off the top of the trace window, but it has allowed us to see the reflections quite clearly. There is a cursor on the impulse itself and we have placed another cursor at the first impulse and are displaying the delay between the two in milliseconds. This indicates that the first reflection in this reverb occurs at 18.15ms after the initial impulse. Given that sound travels at roughly 340m/s in air, this is equivalent to a path length difference of more than 6m. This reverb is emulating a fairly large space with the first boundary at least 3m away.

FFT Windowing and anechoic measurements

From Fourier theory, a perfect impulse contains all frequencies with equal amplitude. If we perform a Fourier transform on a perfect impulse, we should get a flat line. In this instance, the impulse represents the transfer function of the system under test. By performing an FFT on the initial impulse we can find the frequency response of the system.

We achieve this in practice by “windowing” out just the initial impulse and performing an FFT on the data that has been weighted by this function whose purpose is to progressively remove any data we do not want to appear in our frequency response. This is shown below:



Here the green line is the impulse response, the blue is the FFT window function and the red is the resulting FFT derived frequency response. Note that the FFT window ends just before the first reflection so as to make the measurement “anechoic” (literally “without echoes”). Similarly we could actually window the first reflection (or any reflection) and view its frequency response in isolation if we were interested in the absorption properties of partition walls for example.

Distortion Harmonics

We noted earlier that it is a special property of the Log Swept Sine test signal that the distortion harmonics sweep at the same rate as the fundamental. This has an interesting effect when it comes to the derivation of the impulse response. As far as the cross-correlation is concerned, the distortion “sweeps” are higher frequency versions of the fundamental, and actually correlate quite well with the higher end of the original sweep – in fact, they appear to be waveforms that started before the original impulse and were already at a higher frequency when the the main impulse came along. They appear in the buffer as small impulses before the main impulse. Since the buffers are “wrapped” around as far as the maths is concerned, they will usually be displayed at the end of the dScope sample buffer. Placing the FFT window over these distortion impulses results in the frequency content of the harmonics. The position of the distortion impulses relative to the fundamental can be calculated from the sweep rate and the sample rate. This is done in the script and the FFT window placed automatically.

Distortion Normalisation

There are a couple further steps to be performed when plotting the distortion harmonics. Normally when we plot the 2nd and 3rd order harmonics, we plot them at the frequency of the fundamental (ie, even though the 2nd harmonic of a 1kHz tone is at 2kHz, we would plot it at 1kHz and call it the 2nd harmonic of a 1kHz tone.) The direct FFT of the harmonic impulses in the impulse response doesn't do this by default – it plots it directly at the actual frequency of the harmonic. The script gives us the option to take this data and correct it so that it is displayed at the frequency of the fundamental (using the “Normalise” checkbox) as is conventional. In this case the distortion harmonic traces originally had data all the way up to the Nyquist frequency, but since they are shifted down and plotted against the fundamental from



which they are derived, the distortion harmonic traces get progressively shorter as the order of the harmonic increases as you would expect.

Also, the spacing of the distortion harmonic impulses restricts the size of FFT window that can be placed over them without overlapping adjacent distortion harmonics. The size of the FFT window that can be placed over the harmonic determines the low frequency resolution of the FFT calculations and this is reflected in the starting frequency of the distortion traces. The higher the harmonic, the closer the impulse spacing, the smaller the FFT window, and the higher the low frequency cut-off of the measurement.

In addition, since we have performed a measurement of the fundamental over frequency, if we select the normalise option we can actually plot the level of each distortion harmonic in dB relative to the level of the fundamental at the frequency from which it came. This means that if there was a particularly high resonance or dip in the response of the transducer fundamental, this would not directly cause a corresponding peak or dip in the distortion response as it is plotted relative to this peak or dip. If the "Normalise" option is not selected, the harmonics are calculated relative to the level of the fundamental at the same frequency. This is summarised below.

Distortion Normalisation Summary:

1) Normalise checkbox checked: (default)

- This is compatible with the standard way of making distortion measurements using a sine wave whereby the harmonics are calculated relative to the level of the fundamental that caused them and plotted at the frequency of that fundamental (i.e., with a 1kHz fundamental, the 2nd harmonic is at 2kHz, but is plotted at 1kHz relative to the level of the fundamental at 1kHz.)
- This answers the question "If I play a tone at a certain frequency, what is the level of the distortion harmonics relative to the tone?".
- This takes into consideration the frequency response of the EUT as the amplitude of the distortion components is expressed relative to the level of the fundamental that caused it.
- If the measurement system does not have a flat amplitude response over the range from the fundamental to the highest harmonic of interest, it will not be able to make an accurate distortion measurement. This needs to be corrected for using pre-weighting.

2) Normalised checkbox NOT checked

- The distortion harmonics are plotted at the frequency at which they occur relative to the level of the fundamental at that same frequency (eg, a 2kHz 2nd harmonic distortion is calculated relative to the fundamental at 2kHz and plotted at 2kHz, despite the fact that it was caused by a fundamental at 1kHz.) This is a particularly strange way to do things because it effectively plots the distortion harmonics relative to a signal that has nothing much to do with causing it. It is included here for compatibility with other systems.
- This is perhaps easier to understand in the context of broadband noise in the presence of resonant effects that take place after the sound generation (such as horns) as it helps to answer the question "Given a broadband input, what proportion of the sound at a given frequency is made up of distortion harmonics?".
- This does not take the frequency response of the EUT into consideration and will give unusual results with devices with non-flat responses.
- This does compensate for the frequency response of the measurement system and any resonant / horn effects to some degree as any influence that the horn or system has will affect the fundamental and harmonics equally at any given frequency and since the harmonics are plotted relative to the fundamental at the same frequency (ie, a 2kHz distortion harmonic is calculated relative to the fundamental at 2kHz and shown on the trace at 2kHz) any anomaly will have no effect.



- This needs to be treated with some care as it is not really a distortion measurement in the normal sense because of the way the relativity is applied – take the situation where you measure a two way loudspeaker with a crossover frequency of 3kHz. If we consider the point on the 2nd harmonic trace at 4kHz in the case where the “Normalise” box is not checked, this will be showing the level of the 2nd harmonic distortion measured at 4kHz plotted relative to the fundamental at 4kHz. Bearing in mind that the 2nd harmonic at 4kHz was generated by a fundamental frequency of 2kHz (ie, below the crossover frequency) it was actually caused primarily by the mid-woofer, but we are plotting it relative to the level at 4kHz which will be generated primarily by the tweeter. Although it's not really a conventional “distortion” trace It may still be useful to help diagnose and compensate for problems caused by resonances in the system that occur after the transducer such as horn, port or room resonances.