



# Swept Sine Chirps for Measuring Impulse Response

Ian H. Chan  
Design Engineer

Stanford Research Systems, Inc.

*Log-sine chirp and variable speed chirp are two very useful test signals for measuring frequency response and impulse response. When generating pink spectra, these signals possess crest factors more than 6dB better than maximum-length sequence. In addition, log-sine chirp separates distortion products from the linear response, enabling distortion-free impulse response measurements, and variable speed chirp offers flexibility because its frequency content can be customized while still maintaining a low crest factor.*

## 1. Introduction

Impulse response and, equivalently, frequency response measurements are fundamental to characterizing any audio device or audio environment. In principle, any stimulus signal that provides energy throughout the frequency range of interest can be used to make these measurements. In practice however, the choice of stimulus signal has important implications for the signal-to-noise ratio (SNR), distortion, and speed of the audio measurements. We describe two signals that are generated synchronously with FFT analyzers (chirp signals) that offer great SNR and distortion properties. They are the **log-sine chirp** and the **variable speed chirp**. The log-sine chirp has a naturally useful pink spectrum, and the unusual ability to *separate* non-linear (distortion) responses from the linear response [1,2]. The utility of variable speed chirp comes from its ability to reproduce an *arbitrary* target spectrum, all the while maintaining a low crest factor. Because these signals mimic sines that are swept in time, they are known generically as swept sine chirps.

## 2. Why Swept Sine Chirps?

Most users are probably familiar with measuring frequency response at discrete frequencies. A sine signal is generated at one frequency, the response is measured at that frequency, and then the signal is changed to another frequency. Such measurements have very high signal-to-noise ratios because all the energy of the signal at any point in time is concentrated at one frequency. However, it can only manage measurement rates of a handful of frequencies per second at best. This technique is best suited to making measurements where very high SNR is needed, like acoustic measurements in noisy environments, or when measuring very low level signals, like distortion or filter stop-band performance. In contrast, broadband stimulus signals excite many frequencies all at once. A 32k sample signal, for example, generated at a sample rate of 64kHz can excite 16,000 different frequencies in only half a second. This results in much faster measurement rates, and while energy is more spread out than with a sine, in many situations the SNR is more than sufficient to enable good measurements of low level signals. We will show several such measurements in Section 5.

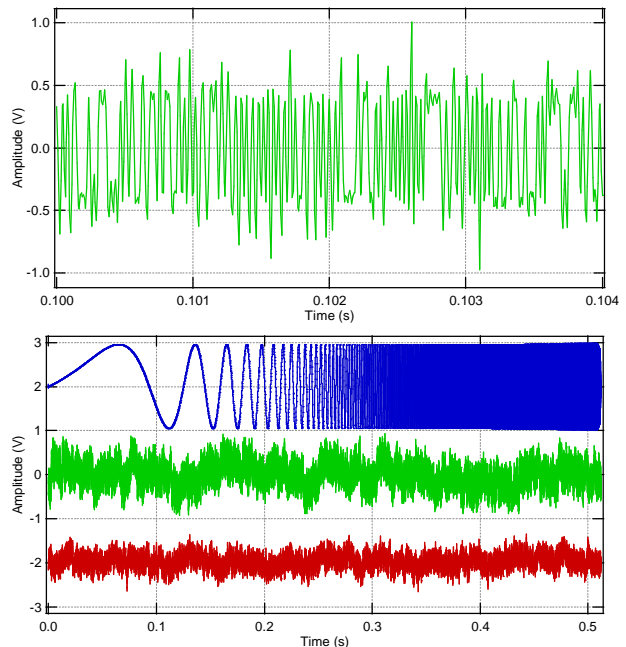


Figure 1. a) Close-up of an MLS signal showing the large excursions due to sudden transitions inherent in the signal. Crest factor is about 8dB instead of the theoretical 0dB. b) Three signals with pink spectra. From top to bottom, log-sine chirp, filtered MLS, and filtered Gaussian noise. The crest factor worsens from top to bottom. All signals have a peak amplitude of 1V. Signals offset for clarity.

A figure of merit that distinguishes different broadband stimulus signals is the crest factor, the ratio of the peak to RMS level of the signal. A signal with a low crest factor contains greater energy than a high crest factor signal with the same peak amplitude, so a low crest factor is desirable. Maximum-length sequence (MLS) theoretically fits the bill because it has a mathematical crest factor of 0dB, the lowest crest factor possible. However, in practice, the sharp transitions and bandwidth-limited reproduction of the signal result in a crest factor of about 8dB (Fig. 1a). Filtering MLS to obtain a more useful pink spectrum further increases the crest factor to 11-12dB. Noise is even worse. Gaussian noise has a crest factor of about 12dB (white spectrum), which increases to 14dB when pink-filtered.<sup>1</sup> On the other hand, log-sine chirp has a measured crest factor of just 4dB (Fig. 1b), and has a naturally pink spectrum. The crest factor of variable speed chirp is similarly low, measuring 5dB for a pink target spectrum. These crest factors are 6-8dB better than that of pink-filtered MLS. That is, MLS needs to be played more than twice as loud as these chirps, or averaged more than four times as long at the same volume, for the same signal-to-noise ratio.

### 3. Generating Swept Sine Chirps

Compounding the low crest factor advantage are the unique properties of log-sine chirp to remove distortion, and variable speed chirp to produce an arbitrary spectrum. To understand how these properties come about requires a knowledge of how these signals are constructed. First up is the log-sine chirp. The log-sine chirp is essentially a sine wave whose frequency increases exponentially with time (e.g. doubles in frequency every 10ms). This is encapsulated by [2]

$$x(t) = \sin\left(\frac{2\pi f_1 T}{\ln(\frac{f_2}{f_1})} \left[ \exp\left(\frac{\ln(\frac{f_2}{f_1})t}{T}\right) - 1 \right]\right) \quad (1)$$

where  $f_1$  is the starting frequency,  $f_2$  the ending frequency, and  $T$  the duration of the chirp. This signal is shown in Figure 2. The explanation of its special property will come in Section 4, when the signal is analyzed.

The variable speed chirp's special property comes from the simple idea of using the *speed* of the sweep to control the frequency response [3]. The greater the desired response, the slower the sweep through that frequency (Fig. 3). To generate a variable speed chirp, it is easiest to go into the frequency domain. This entails specifying both the magnitude and phase of the signal, and then doing an inverse-FFT to obtain the desired time-domain signal. The magnitude of a variable speed chirp is simply that of the desired target frequency response  $|H_{user}(f)|$ . The phase is a little trickier to specify. What we have to do is first specify the *group delay* of the signal  $t_G$ , and then work out the phase from the group delay. The group delay for variable speed chirp is [3]

$$t_G(f) = t_G(f - df) + C|H_{user}(f)|^2 \quad (2)$$

where

$$C = \frac{t_G(f_2) - t_G(f_1)}{\sum_{f=0}^{f_s/2} |H_{user}(f)|^2} \quad (3)$$




Figure 2. a) Power spectrum of a log-sine chirp signal. It is pink except at the lowest frequencies. b) Time record of a log-sine chirp signal. The frequency of the signal increases exponentially before repeating itself.

<sup>1</sup> True Gaussian noise has an *infinite* crest factor; the excursion of the noise here was limited to  $\pm 4s$ .

That is, the group delay at  $f$  is the group delay at the previous frequency bin, plus an amount dependent on the magnitude-squared of the target response. The group delay at the first frequency bin is  $t_G(f_1)$ .

The starting and ending group delays,  $t_G(f_1)$  and  $t_G(f_2)$ , represent the start and stop times of the sweep respectively, and are specified by the user. They must fall within the time interval of the chirp  $T$

$$0 \leq t_G(f_1) < t_G(f_2) \leq T. \quad (4)$$

Because the signal is generated in the frequency domain, the actual start and stop times will leak over a little in the time domain. Depending on your requirements, you may want to start the sweep a little after 0, and stop the sweep a little before  $T$ . Recalling that group delay  $t_G(f) = -\frac{d\phi(f)}{2\pi df}$ , phase (in radians) can be obtained by integrating the group delay

$$\phi(f) = -2\pi \int t_G(f) df. \quad (5)$$

This is best done numerically.

#### 4. Analysis of Swept Sine Chirps

Swept-sine chirps are analyzed using two-channel FFT techniques (Fig. 4) to determine the frequency response of the device under test (DUT),

$$H_{DUT}(f) = \frac{Y(f)}{X(f)} \quad (6)$$

where  $Y(f)$  is the FFT of the input channel, and  $X(f)$  is the FFT of the reference channel. Division by the reference channel response (assuming it is non-zero) cancels out both the magnitude *and* phase irregularities present in the test signal.<sup>2</sup> This automatically accounts for any intended or unintended non-flatness in the stimulus signal, as well as zeros out any group delay present. The impulse response of the DUT is then the inverse-FFT of  $H_{DUT}(f)$ .

Now we can see how the log-sine chirp can create a distortion-free impulse response. The group delay of a log-sine chirp (which will be removed), is

$$t_G(f) = T \frac{\ln(\frac{f}{f_1})}{\ln(\frac{f_2}{f_1})}. \quad (7)$$

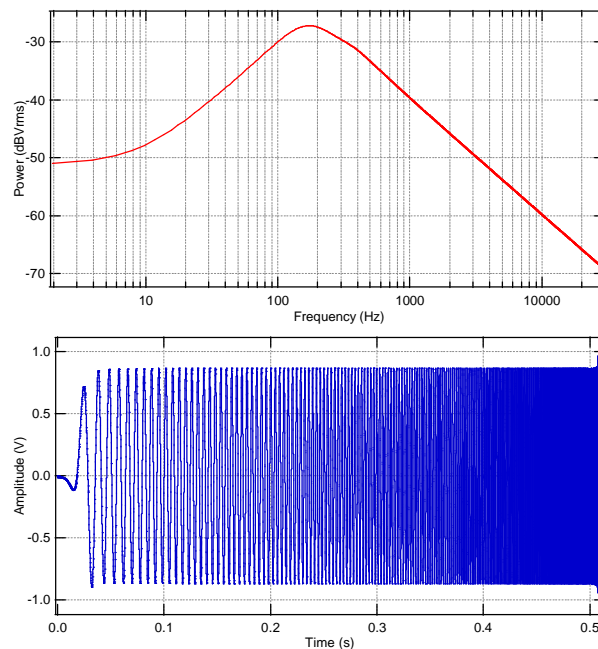


Figure 3. a) Power spectrum of a variable speed chirp signal. In this example, the target EQ was that of USASI noise. b) Time record of the same variable speed chirp signal. The signal dwells at frequencies that have large emphasis, and speeds through frequencies with small emphasis, to achieve a low crest factor (4.4dB in this case).

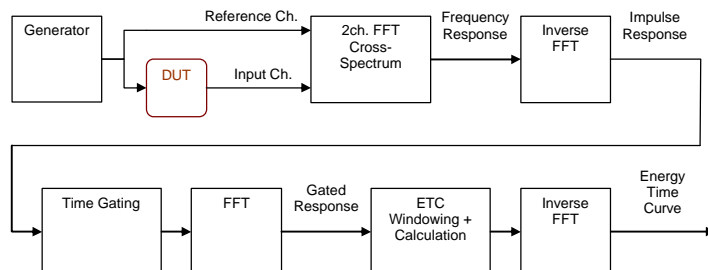


Figure 4. Diagram illustrating signal flow in a 2-channel FFT measurement, such as found in the SR1 Audio Analyzer employed in this paper. Time gating and energy-time curve (ETC) computation are typically used in acoustic measurements.

<sup>2</sup> It is also important that the FFT sees a *consistent* stimulus spectrum, especially if there are delays between reference and input channels. Using a noise-like stimulus that has inconsistent shot-to-shot spectra can result in unreliable measurements.

This represents the *time* at which the (fundamental) frequency  $f$  is produced. The time at which the  $N^{th}$  harmonic,  $f_N$ , is produced is the time when the fundamental is  $\frac{f_N}{N}$ . Thus, the group delay of the harmonic is

$$t_{GN}(f_N) = T \frac{\ln(\frac{f_N}{Nf_1})}{\ln(\frac{f_2}{f_1})} \tag{8}$$

Following the two-channel FFT analysis, the group delay of the fundamental is zeroed out, meaning that the time delay at each frequency is subtracted according to (7). The group delay of the harmonics thus becomes

$$\Delta t_{GN}(f_N) = t_{GN}(f_N) - t_G(f_N) = -T \frac{\ln(N)}{\ln(\frac{f_2}{f_1})} \tag{9}$$

Note that it only depends on the *order* of the harmonic,  $N$ , but not at all on the frequency,  $f_N$ ! So after analysis, all frequencies arising from a particular harmonic order arrive *at the same time*, creating an impulse response that precedes the linear impulse response by a time  $|\Delta t_{GN}|$  (Fig. 5a). This is quite a result, and is a property unique to the log-sine chirp.

### 5. Examples of Swept Sine Measurements

For the measurements made here, we employed our new SR1 Audio Analyzer. This analyzer includes both log-sine chirp and variable speed chirp generators, as well as a two-channel FFT analyzer (actually, it has two). In Figures 5 and 6, we use log-sine chirp to measure the behavior of an elliptical low-pass filter with a 6kHz cut-off frequency. Figure 5a shows the impulse response of the filter with the harmonic impulses neatly separated in time. In this example,  $\frac{f_2}{f_1} = 4095$  and  $T = 128\text{ms}$ , and all the harmonic impulses are seen to arrive at their expected times. By time-gating the impulse response to

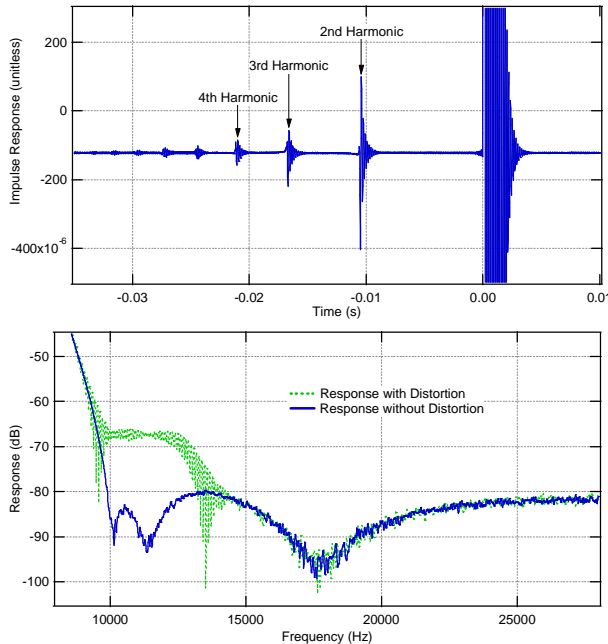


Figure 5. a) Measured electrical impulse response of an elliptical low-pass filter with log-sine chirp. The main response (which is distortion-free) occurs after  $t=0$ . The non-linear responses consist of peaks that *precede* the main response, with the higher-order responses occurring earlier. b) By gating the impulse response, we can examine the stop-band frequency response with distortion products (green), and without (blue). The intrusion of 2<sup>nd</sup> harmonic energy at -65dB up to twice the low-pass frequency is clear.




Figure 6. a) Second harmonic distortion of the elliptical low-pass filter, as measured using a log-sine chirp (blue), and a conventional stepped sine sweep (red). Both measurements made referred to input. The log-sine chirp data tracks the conventional measurement very well down to about -100dB. b) Phase response of the second harmonic with the delay removed, as measured using a log-sine chirp.

include or exclude the distortion impulses, we obtain the stop-band frequency response of the DUT with or without distortion (Fig. 5b). The blue curve is the distortion-free frequency response of the filter, while the green curve includes the distortion components. The difference between the two curves is entirely because of distortion (predominantly second harmonic, as evidenced by the  $\sim 12\text{kHz}$  roll-off, which is twice the filter cut-off frequency).

Now, each harmonic impulse response in Figure 5a captures the *full behavior* of the distortion product (amplitude and phase), and may be analyzed just like the linear response. In figure 6, the second harmonic impulse response is time-gated and analyzed (between  $-14\text{ms}$  and  $-4\text{ms}$ , with 5% raised cosine windowing applied to both ends).<sup>3</sup> The admittedly unusual second harmonic distortion response tracks the results of a conventional stepped sine measurement closely, so this verifies that the distortion measured with log-sine chirp is accurate. Phase response of the second harmonic is shown in Figure 6b, after the constant delay has been removed. These measurements were averaged for less than a second (128ms chirp averaged 4 times).

Next we demonstrate using the variable speed chirp to make an acoustic measurement of a loudspeaker. Being able to tailor the frequency response of the test signal is often very useful. In this case, we chose a target EQ of USASI noise (Fig. 2a), which resembles the frequency spectrum of program material. We also chose it because power falls off below  $100\text{Hz}$ , and due to time-gating, we did not expect meaningful data below a few hundred hertz anyway. The measurement setup used was similar to Figure 4, except that the driving amplifier output was fed into the reference channel of the SR1 Audio Analyzer (Fig. 7). Doing so removes the frequency response of the driving amplifier, leaving the measured response as that of the 2-way speaker (DUT) and the calibrated microphone. Power to the speaker was set at  $2V_{\text{rms}}$  ( $1\text{W}$  into  $4\Omega$ ), and the measurement microphone was placed about 10 feet away.

Figure 8a shows part of the impulse response measured with the variable speed chirp. The main response begins at about  $9.5\text{ms}$ , and the first echo follows about  $2.5\text{ms}$  later. Figure 8b shows the quasi-anechoic response in blue (gated from  $8\text{ms}$  to  $12\text{ms}$ , with 5% raised-cosine window at both ends) overlaid with the ungated frequency response in green. The time-gated frequency response is, of course, much smoother and more meaningful than the ungated response due to the exclusion of echoes. The gated impulse response and energy-time curve (ETC) are shown in Figure 9. The ETC was computed using a half-Hann window [4], and is an indication of the energy response of the loudspeaker. Measurements here were averaged over about 2 seconds ( $512\text{ms}$  long chirp averaged 4 times).




Figure 7. Modified signal flow, where the output of the power amplifier is fed into the reference channel. This removes the amplitude and phase imperfections of the amplifier from the measurement of the loudspeaker.




Figure 8. a) Impulse response of a speaker measured with a variable speed chirp that had a target EQ of USASI noise to simulate program material. The impulse response starts at about  $9.5\text{ms}$  due to the distance between the mic and the speaker. The first echo (reflection from floor) follows about  $2.5\text{ms}$  later. b) The raw frequency response of the speaker, including the echoes, is shown in green. The trace in blue is the gated, quasi-anechoic response, which shows a much smoother, meaningful response. The gating was applied between  $8\text{ms}$  and  $12\text{ms}$ , with a 5% raised-cosine window at both ends. Due to the first echo arriving just  $2.5\text{ms}$  later, the response below about  $400\text{Hz}$  is not accurate.

<sup>3</sup> When analyzing harmonic impulses, remember to divide the frequency axis by  $N$ . e.g. the response at  $20\text{kHz}$  for a second harmonic impulse was generated by a fundamental at  $10\text{kHz}$ .

## 6. Conclusions

Both the log-sine chirp and variable speed chirp signals are powerful additions to the toolbox of the professional audio engineer. These signals can be used to make measurements very quickly compared to traditional stepped sine sweeps, and have crest factors significantly better than that of MLS. The log-sine chirp also has the unique advantage of being able to separate distortion response from linear response, while the variable speed chirp is able to generate a customized frequency spectrum at a low crest factor. The advantages of these signals were demonstrated in real-world test situations using the Stanford Research Systems SR1 Audio Analyzer.




Figure 9. Energy-time curve of the speaker under test (red) together with the corresponding gated impulse response (blue), as measured using a variable speed chirp. The target frequency spectrum of the chirp was USASI noise. A half-Hann window was used to calculate the ETC.

## 7. References

- [1] A. Farina, "Simultaneous measurement of impulse response and distortion with a swept sine technique," *presented at the 108<sup>th</sup> AES Convention*, Paris, France, February 2000.
- [2] T. Kite, "Measurement of audio equipment with log-swept sine chirps," *presented at the 117<sup>th</sup> AES Convention*, San Francisco, October 2004.
- [3] S. Müller and P. Massarini, "Transfer-Function Measurement with Sweeps," *J. Audio Eng. Soc.*, vol. 49, pp. 443-471, June 2001.
- [4] J. Vanderkooy and S. P. Lipshitz, "Uses and Abuses of the Energy-Time Curve," *presented at the 87<sup>th</sup> AES Convention*, New York, October 1989.

