# About FFT Spectrum Analyzers

*Application Note #1*

## What is an FFT Spectrum Analyzer?

FFT Spectrum Analyzers, such as the SR760, SR770, SR780 and SR785, take a time varying input signal, like you would see on an oscilloscope trace, and compute its frequency spectrum.

Fourier's theorem states that any waveform in the time domain can be represented by the weighted sum of sines and cosines. The FFT spectrum analyzer samples the input signal, computes the magnitude of its sine and cosine components, and displays the spectrum of these measured frequency components.

## Why Look at a Signal's Spectrum?

For one thing, some measurements which are very hard in the time domain are very easy in the frequency domain. Consider the measurement of harmonic distortion. It's hard to quantify the distortion of a sine wave by looking at the signal on an oscilloscope. When the same signal is displayed on a spectrum analyzer, the harmonic frequencies and amplitudes are displayed with amazing clarity. Another example is noise analysis. Looking at an amplifier's output noise on an oscilloscope basically measures just the total noise amplitude. On a spectrum analyzer, the noise as a function of frequency is displayed. It may be that the amplifier has a problem only over certain frequency ranges. In the time domain it would be very hard to tell.

Many of these measurements were once done using analog spectrum analyzers. In simple terms, an analog filter was used to isolate frequencies of interest. The signal power which passed through the filter was measured to determine the signal strength in certain frequency bands. By tuning the filters and repeating the measurements, a spectrum could be obtained.

## The FFT Analyzer

An FFT spectrum analyzer works in an entirely different way. The input signal is digitized at a high sampling rate, similar to a digitizing oscilloscope. Nyquist's theorem says that as long as the sampling rate is greater than twice the highest frequency component of the signal, the sampled data will accurately represent the input signal. In the SR7xx (SR760, SR770, SR780 or SR785), sampling occurs at 256 kHz. To make sure that Nyquist's theorem is satisfied, the input signal passes through an analog filter which attenuates all frequency components above 156 kHz by 90 dB. This is the anti-aliasing filter. The resulting digital time record is then mathematically transformed into a frequency spectrum using an algorithm known as the Fast Fourier Transform, or FFT. The FFT is simply a clever set of operations which implements Fourier's theorem. The resulting spectrum shows the frequency components of the input signal.

Now here's the interesting part. The original digital time record comes from discrete samples taken at the sampling rate. The corresponding FFT yields a spectrum with discrete frequency samples. In fact, the spectrum has half as many frequency points as there are time points. (Remember Nyquist's theorem.) Suppose that you take 1024 samples at 256 kHz. It takes 4 ms to take this time record. The FFT of this record yields 512 frequency points—but over what frequency range? The highest frequency will be determined by the period of two time samples or 128 kHz. The lowest frequency is just the period of the entire record or 1/(4 ms) or 250 Hz. Everything below 250 Hz is considered to be DC. The output spectrum thus represents the frequency range from DC to 128 kHz, with points every 250 Hz.

## Advantages of FFT Analyzers

The advantage of this technique is its speed. Because FFT spectrum analyzers measure all frequency components at the same time, the technique offers the possibility of being hundreds of times faster than traditional analog spectrum analyzers. In the case of a 100 kHz span and 400 resolvable frequency bins, the entire spectrum takes only 4 ms to measure. To measure the signal with higher resolution, the time record is increased. But again, all frequencies are examined simultaneously providing an enormous speed advantage.

In order to realize the speed advantages of this technique we need to do high speed calculations. And, in order to avoid sacrificing dynamic range, we need high-resolution ADCs. SRS spectrum analyzers have the processing power and front-end resolution needed to realize the theoretical benefits of FFT spectrum analyzers.

## Dual-Channel FFT Analyzers

One of the most common applications of FFT spectrum analyzers is to measure the transfer function of a mechanical or electrical system. A transfer function is the ratio of the output spectrum to the input spectrum. Single-channel analyzers, such as the SR760, cannot measure transfer functions. Single-channel analyzers with integrated sources, such as the SR770, can measure transfer functions but only by assuming that the input spectrum of the system is equal to the spectrum of the integrated source. In general, to measure a general transfer function, a two-channel analyzer (such as the SR785) is required. One channel measures the spectrum of the input, the other measures the spectrum of the output, and the analyzer performs a complex division to extract the magnitude and phase of the transfer function. Because the input spectrum is actually measured and divided out, you're not limited to using a predetermined signal as the input to the system under test—any signal will do.

## Frequency Spans

Before continuing, a couple of points about frequency span need clarification. We just described how we arrived at a DC to 128 kHz frequency span using a 4 ms time record. Because the signal passes through an anti-aliasing filter at the input, the entire frequency span is not useable. The filter has a flat response from DC to 100 kHz and then rolls off steeply

from 100 kHz to 156 kHz. No filter can make a 90 dB transition instantly. The range between 100 kHz and 128 kHz is therefore not useable, and the actual displayed frequency span stops at 100 kHz. There is also a frequency bin labeled 0 Hz (or DC). This bin actually covers the range from 0 Hz to 250 Hz (the lowest measurable frequency) and contains the signal components whose period is longer than the time record (not only DC). So the final displayed spectrum contains 400 frequency bins. The first covers 0 to 250 Hz, the second 250 to 500 Hz, and the 400th covers 99.75 to 100.0 kHz.

### Spans Less Than 100 kHz

The length of the time record determines the frequency span and resolution of our spectrum. What happens if we make the time record 8 ms (twice as long)? Well, we ought to get 2048 time points (sampling at 256 kHz) yielding a spectrum from DC to 100 kHz with 125 Hz resolution containing 800 points. But the SR7xx places some limitations on this. One is memory. If we keep increasing the time record we will need to store more and more points. (1 Hz resolution would require 256 k values.) Another limitation is processing time. The time it takes to calculate an FFT with more points increases more than linearly.

To overcome this problem, the analyzer digitally filters and decimates the incoming data samples (at 256 kHz) to limit the bandwidth and reduce the number of points in the FFT. This is similar to the anti-aliasing filter at the input except the digital filter's cutoff frequency can be changed. In the case of the 8 ms record, the filter reduces the bandwidth to 64 kHz with a filter cutoff of 50 kHz (the filter rolls off between 50 kHz and 64 kHz). Remember that Nyquist only requires samples at twice the frequency of the highest signal frequency. Thus, the digital filter only has to output points at 128 kHz, or half of the input rate (256 kHz). The net result is the digital filter outputs a time record of 1024 points, effectively sampled at 128 kHz, to make up an 8 ms record. The FFT processor operates on a constant number of points, and the resulting FFT will yield 400 points from DC to 50 kHz. The resolution or linewidth is 125 Hz.

This process of doubling the time record and halving the span can be repeated by using multiple stages of digital filtering. The SR7xx can process spectra with a span of only 191 mHz with a time record of 2098 seconds if you have the patience. However, this filtering process only yields baseband measurements (frequency spans which start at DC).

### Starting the Span Somewhere Other Than DC

In addition to choosing the span and resolution of the spectrum, we may want the span to start at frequencies other than DC. It would be nice to center a narrow span around any frequency below 100 kHz. Using digital filtering alone requires that every span start at DC. We need to frequency shift, or heterodyne, the input signal. Multiplying the incoming signal by a complex sine will frequency shift the signal. The resulting spectrum is shifted by the frequency of the complex sine. If we incorporate heterodyning with our digital filtering, we can shift any frequency span so that it

starts at DC. The resulting FFT yields a spectrum offset by the heterodyne frequency.

Heterodyning allows the analyzer to compute zoomed spectra (spans which start at frequencies other than DC). The digital filter processor can filter and heterodyne the input in real time to provide the appropriate filtered time record at all spans and center frequencies. Because the digital signal processors in the SR7xx are so fast, you won't notice any calculation time while taking spectra. All the signal processing calculations, heterodyning, digital filtering and Fourier transforming are done in less time than it takes to acquire the data. So the SR7xx can take spectra seamlessly, i.e. there is no dead time between one time record and the next.

### Measurement Basics

An FFT spectrum is a complex quantity. This is because each frequency component has a phase relative to the start of the time record. (Alternately, you may wish to think of the input signal being composed of sines and cosines.) If there is no triggering, the phase is random and we generally look at the magnitude of the spectrum. If we use a synchronous trigger, each frequency component has a well defined phase.

### Spectrum

The spectrum is the basic measurement of an FFT analyzer. It is simply the complex FFT. Normally, the magnitude of the spectrum is displayed. The magnitude is the square root of the FFT times its complex conjugate. (Square root of the sum of the real (sine) part squared and the imaginary (cosine) part squared.) The magnitude is a real quantity and represents the total signal amplitude in each frequency bin, independent of phase.

If there is phase information in the spectrum, i.e. the time record is triggered in phase with some component of the signal, then the real (cosine) or imaginary (sine) part or the phase may be displayed. The phase is simply the arctangent of the ratio of the imaginary and real parts of each frequency component. The phase is always relative to the start of the triggered time record.

### Power Spectral Density (PSD)

The PSD is the magnitude of the spectrum normalized to a 1 Hz bandwidth. This measurement approximates what the spectrum would look like if each frequency component were really a 1 Hz wide piece of the spectrum at each frequency bin.

What good is this? When measuring broadband signals (such as noise), the amplitude of the spectrum changes with the frequency span. This is because the linewidth changes, so the frequency bins have a different noise bandwidth. The PSD, on the other hand, normalizes all measurements to a 1 Hz bandwidth, and the noise spectrum becomes independent of the span. This allows measurements with different spans to be compared. If the noise is Gaussian in nature, the amount of noise amplitude in other bandwidths may be approximated by scaling the PSD measurement by the square root of the

bandwidth. Thus, the PSD is displayed in units of V/√Hz or dBV/√Hz.

Since the PSD uses the magnitude of the spectrum, the PSD is a real quantity. There is no real or imaginary part, or phase.

### Time Record

The time record measurement displays the filtered and decimated (depending on the span) data points before the FFT is taken. In the SR760 and SR770, this information is available only at full span. In the SR780 and SR785, time records can be displayed at all spans. For baseband spans (spans that start at DC), the time record is a real quantity. For non-baseband spans, the heterodyning discussed earlier transforms the time record into a complex quantity which can be somewhat difficult to interpret.

### Two-Channel Measurements

As we discussed earlier, two-channel analyzers (such as the SR780 and SR785) offer additional measurements such as transfer function, cross-spectrum, coherence and orbit. These measurements, which only apply to the SR780 and SR785, are discussed below.

### Transfer Function

The transfer function is the ratio of the spectrum of channel 2 to the spectrum of channel 1. For the transfer function to be valid, the input spectrum must have amplitude at all frequencies over which the transfer function is to be measured. For this reason, broadband sources (such as noise, or periodic chirps) are often used as inputs for transfer function measurements.

### Cross Spectrum

The cross spectrum is defined as:

$$\text{cross spectrum} = \text{FFT2} \times \text{conj(FFT1)}$$

The cross spectrum is a complex quantity which contains magnitude and phase information. The phase is the relative phase between the two channels. The magnitude is simply the product of the magnitudes of the two spectra. Frequencies where signals are present in both spectra will have large components in the cross-spectrum.

### Orbit

The orbit is simply a two dimensional display of the time record of channel 1 vs. the time record of channel 2. The orbit display is similar to an oscilloscope displaying a "Lissajous" figure.

### Coherence

Coherence measures the percentage of power in channel 2 which is caused by (phase coherent with) power in the input channel. Coherence is a unitless quantity which varies from 0 to 1. If the coherence is 1, all the power of the output signal is due to the input signal. If the coherence is 0, the input and output are completely random with respect to one another. Coherence is related to signal-to-noise ratio (S/N) by the formula:

$$\text{S/N} = \gamma^2/(1-\gamma^2)$$

where $\gamma^2$ is the traditional notation for coherence.

### Correlation

The SR780 and SR785 analyzers also compute auto and cross-correlation. Correlation is a time domain measurement which is defined as follows:

$$\text{Auto Correlation} = \int x^*(t)x(t-\tau)dt$$

$$\text{Cross Correlation} = \int x^*(t)y(t-\tau)dt$$

where x and y are the channel 1 and channel 2 input signals and the integrals are over all time. It is clear that the auto correlation at a time t is a measure of how much overlap a signal has with a delayed-by-t version of itself, and the cross-correlation is a measure of how much overlap a signal has with a delayed-by-t version of the other channel. Although correlation is a time-domain measurement, the SR780 and SR785 use frequency-domain techniques to compute it in order to make the calculation faster.

### Spectrum

The most common measurement is the spectrum and the most useful display is the log magnitude. The log magnitude display graphs the magnitude of the spectrum on a logarithmic scale using dBV as units.

Why is the log magnitude display useful? Remember that the SR7xx has a dynamic range of about 90 dB below full scale. Imagine what something 0.01 % of full scale would look like on a linear scale. If we wanted it to be 1 centimeter high on the graph, the top of the graph would be 100 meters above the bottom. It turns out that the log display is both easy to understand and clearly shows features which have very different amplitudes.

Of course, the analyzers are also capable of showing the magnitude on a linear scale. The real and imaginary parts are always displayed on a linear scale. This avoids the problem of taking the log of negative voltages.

### Phase

In general, phase measurements are only used when the analyzer is triggered. The phase is relative to the start of the time record.

The phase is displayed in degrees or radians on a linear scale from −180 to +180 degrees. There is no phase unwrap on the

SR760 and SR770. The SR780 and SR785 can display unwrapped phase which is very useful, for instance, in displaying the phase of filter transfer functions which may vary over hundreds or even thousands of degrees.

The phase of a particular frequency bin is set to zero if neither the real nor imaginary part of the FFT is greater than 0.012 % of full scale (−78 dB below f.s.). This avoids the messy phase display associated with the noise floor. (Remember, even if a signal is small, its phase extends over the full 360 degrees.)

### Watch Out For Phase Errors

The FFT measurement can be thought of as N band pass filters, each centered on a frequency bin. The signal within each filter shows up as the amplitude of each bin. If a signal's frequency is between bins, the filters act to attenuate the signal a little bit. This results in a small amplitude error. The phase error, on the other hand, can be quite large. Because these filters are very steep and selective, they introduce very large phase shifts for signals not exactly on a frequency bin.

On full span, this is generally not a problem. The bins are 250 Hz apart, and most synthesized sources have no problem generating a signal right on a frequency bin. But when the span is narrowed, the bins move much closer together and it becomes very hard to place a signal exactly on a frequency bin.

### Windowing

What is windowing? Let's go back to the time record. What happens if a signal is not exactly periodic within the time record? We said that its amplitude is divided into multiple, adjacent frequency bins. This is true but it's actually a bit worse than that. If the time record does not start and stop with the same data value, the signal can actually smear across the entire spectrum. This smearing will also change wildly between records because the amount of mismatch between the starting value and ending value changes with each record.

Windows are functions defined across the time record which are periodic in the time record. They start and stop at zero and are smooth functions in between. When the time record is windowed, its points are multiplied by the window function, time-bin by time-bin, and the resulting time record is by definition periodic. It may not be identical from record to record, but it will be periodic (zero at each end).

### In the Frequency Domain

In the frequency domain a window acts like a filter. The amplitude of each frequency bin is determined by centering this filter on each bin and measuring how much of the signal falls within the filter. If the filter is narrow, only frequencies near the bin will contribute to the bin. A narrow filter is called a selective window—it selects a small range of frequencies around each bin. However, since the filter is narrow, it falls off from center rapidly. This means that even frequencies close to the bin may be attenuated somewhat. If the filter is wide, frequencies far from the bin will contribute to the bin

amplitude, but those close by will not be attenuated significantly.

The net result of windowing is to reduce the amount of smearing in the spectrum from signals not exactly periodic with the time record. The different types of windows trade off selectivity, amplitude accuracy and noise floor.

The SR7xx offers several types of window functions including Uniform (none), Flattop, Hanning, Blackman-Harris and Kaiser.

### Uniform

The uniform window is actually no window at all. The time record is used with no weighting. A signal will appear as narrow as a single bin if its frequency is exactly equal to a frequency bin. (It is exactly periodic within the time record.) If its frequency is between bins, it will affect every bin of the spectrum. These two cases also have a great deal of amplitude variation between them (up to 4 dB).

In general, this window is only useful when looking at transients which do not fill the entire time record.

### Hanning

The Hanning window is the most commonly used window. It has an amplitude variation of about 1.5 dB (for signals between bins) and provides reasonable selectivity. Its filter rolloff is not particularly steep. As a result, the Hanning window can limit the performance of the analyzer when looking at signals close together in frequency and very different in amplitude.

### Flattop

The Flattop window improves on the amplitude accuracy of the Hanning window. Its between-bin amplitude variation is about 0.02 dB. However, the selectivity is a little worse. Unlike the Hanning, the Flattop window has a wide pass band and very steep rolloff on either side. Thus, signals appear wide but do not leak across the whole spectrum.

### Blackman-Harris

The Blackman-Harris window is a very good window to use with SRS FFT analyzers. It has better amplitude accuracy (about 0.7 dB) than the Hanning, very good selectivity, and the fastest filter rolloff. The filter is steep and narrow and reaches a lower attenuation than the other windows. This allows signals close together in frequency to be distinguished, even when their amplitudes are very different.

### Kaiser

The Kaiser window, which is available on the SR780 and SR785 only, combines excellent selectivity and reasonable accuracy (about 0.8 dB for signals between exact bins). The Kaiser window has the lowest side-lobes and the least broadening for non-bin frequencies. Because of these

properties, it is the best window to use for measurements requiring a large dynamic range. On the SR760 and SR770, the Blackman-Harris window is the best large dynamic range window.

## Averaging

The SR7xx analyzers supports several types of averaging. In general, averaging many spectra together improves the accuracy and repeatability of measurements.

### RMS Averaging

RMS averaging computes the weighted mean of the sum of the squared magnitudes (FFT times its complex conjugate). The weighting is either linear or exponential.

RMS averaging reduces fluctuations in the data but does not reduce the actual noise floor. With a sufficient number of averages, a very good approximation of the actual random noise floor can be displayed.

Since rms averaging involves magnitudes only, displaying the real or imaginary part, or phase, of an rms average has no meaning. The rms average has no phase information.

### Vector Averaging

Vector averaging averages the complex FFT spectrum. (The real part is averaged separately from the imaginary part.) This can reduce the noise floor for random signals since they are not phase coherent from time record to time record.

Vector averaging requires a trigger. The signal of interest must be both periodic and phase synchronous with the trigger. Otherwise, the real and imaginary parts of the signal will not add in phase, and instead will cancel randomly.

With vector averaging, the real and imaginary parts (as well as phase displays) are correctly averaged and displayed. This is because the complex information is preserved.

### Peak Hold

Peak Hold is not really averaging. Instead, the new spectral magnitudes are compared to the previous data, and if the new data is larger, the new data is stored. This is done on a frequency bin-by-bin basis. The resulting display shows the peak magnitudes which occurred in the previous group of spectra.

Peak Hold detects the peaks in the spectral magnitudes and only applies to Spectrum, PSD and Octave Analysis measurements. However, the peak magnitude values are stored in the original complex form. If the real or imaginary part (or phase) is being displayed for spectrum measurements, the display shows the real or imaginary part (or phase) of the complex peak value.

### Linear Averaging

Linear averaging combines N (number of averages) spectra with equal weighting in either RMS, Vector or Peak Hold fashion. When the number of averages has been completed, the analyzer stops and a beep is sounded. When linear averaging is in progress, the number of averages completed is continuously displayed below the averaging indicator at the bottom of the screen.

Auto ranging is temporarily disabled when a linear average is in progress. Be sure that you don't change the input range manually. Changing the range during a linear average invalidates the results.

### Exponential Averaging

Exponential averaging weights new data more than old data. Averaging takes place according to the formula,

New Average = (New Spectrum • 1/N) + (Old Average) • (N−1)/N

where N is the number of averages.

Exponential averages "grow" for approximately the first 5N spectra until the steady state values are reached. Once in steady-state, further changes in the spectra are detected only if they last sufficiently long. Make sure that the number of averages is not so large as to eliminate the changes in the data that might be important.

### Real-Time Bandwidth and Overlap Processing

What is real-time bandwidth? Simply stated, it is the frequency span whose corresponding time record exceeds the time it takes to compute the spectrum. At this span and below, it is possible to compute the spectra for every time record with no loss of data. The spectra are computed in "real time". At larger spans, some data samples will be lost while the FFT computations are in progress.

For all frequency spans, the SR7xx can compute the FFT in less time than it takes to acquire the time record. Thus, the real-time bandwidth of the SR7xx is 100 kHz. This includes the real-time digital filtering and heterodyning, the FFT processing, and averaging calculations. The SR7xx employs two digital signal processors to accomplish this. The first collects the input samples, filters and heterodynes them, and stores a time record. The second computes the FFT and averages the spectra. Since both processors are working simultaneously, no data is ever lost.

The SR780 and SR785 accomplish high-speed processing with a single, advanced-technology, floating-point DSP chip.

### Averaging Speed

How can you take advantage of this? Consider averaging. Other analyzers typically have a real-time bandwidth of around 4 kHz. This means that even though the time record at 100 kHz span is only 4 ms, the "effective" time record is

25 times longer due to processing overhead. An analyzer with 4 kHz of real-time bandwidth can only process about 10 spectra a second. When averaging is on, this usually slows down to about 5 spectra per second. At this rate it takes a few minutes to do 500 averages.

The SR7xx, on the other hand, has a real-time bandwidth of 100 kHz. At a 100 kHz span, the analyzer is capable of processing 250 spectra per second. In fact, this is so fast that the display can not be updated for each new spectra. The display only updates about 6 times a second. However, when averaging is on, all of the computed spectra will contribute to the average. The time it takes to complete 500 averages is only a few seconds. (Instead of a few minutes!)

## Overlap

What about narrow spans where the time record is long compared to the processing time? The analyzer computes one FFT per time record and can wait until the next time record is complete before computing the next FFT. The update rate would be no faster than one spectra per time record. With narrow spans, this could be quite slow.

And what is the processor doing while it waits? Nothing. With overlap processing, the analyzer does not wait for the next complete time record before computing the next FFT. Instead, it uses data from the previous time record, as well as data from the current time record, to compute the next FFT. This speeds up the processing rate. Remember, most window functions are zero at the start and end of the time record. Thus, the points at the ends of the time record do not contribute much to the FFT. With overlap, these points are "re-used" and appear as middle points in other time records. This is why overlap effectively speeds up averaging and smooths out window variations.

Typically, time records with 50 % overlap provide almost as much noise reduction as non-overlapping time records when rms averaging is used. When rms averaging narrow spans, measurement time can be reduced by a factor of two.

## Overlap Percentage

The amount of overlap is specified as a percentage of the time record. 0 % is no overlap, and 99.8 % is the maximum (511 out of 512 samples re-used). The maximum overlap is determined by the amount of time it takes to calculate an FFT and the length of the time record, and thus varies according to the span.

The SR760/SR770 always try to use the maximum amount of overlap possible. This keeps the display updating as fast as possible. Whenever a new frequency span is selected, the overlap is set to the maximum possible value for that span. If less overlap is desired, use the average menu to enter a smaller value. On the widest spans (25, 50 and 100 kHz), no overlap is allowed.

The SR780 and SR785 use a slightly different system for specifying the overlap. The overlap entered by the user is the "requested overlap". The instrument attempts to make the

actual overlap as close as possible to the requested overlap. The SR780 and SR785 compute and display the actual overlap so that it is obvious when it differs from the requested overlap.

## Octave Analysis

The magnitude of the normal spectrum measures the amplitudes within equally divided frequency bins. Octave analysis computes the spectral amplitude in logarithmic frequency bands whose widths are proportional to their center frequencies. The bands are arranged in octaves with either 1, 3 or 12 bands per octave (1/1, 1/3 or 1/12 octave analysis). Octave analysis measures spectral power closer to the way people perceive sound: in octaves.

The actual method used to calculate octave measurements differs for each of the analyzers. In the SR780 and SR785, the input data passes into a bank of parallel digital filters. The filter center frequencies and shapes are determined by the type of octave analysis (1/1, 1/3 or 1/12 octave) and comply with ANSI s1-11-1986, Order 3, Type 1-D. The output of each filter is rms averaged to compute the power and displayed as a bar-type graph. This is a real-time measurement of the power within each band and is the only available octave measurement. Since the bands are spaced logarithmically, octave displays always have a logarithmic x-axis.

## Band Center Frequencies

The center frequency of each band is calculated according to ANSI standard S1.11 (1986). The shape of each band is a third-order Butterworth filter whose bandwidth is either a full, 1/3 or 1/12 octave. The full octave bands have band centers at:

$$\text{Center Frequency} = 1 \text{ kHz} \times 2^n$$

The 1/3 octave bands have center frequencies given by:

$$\text{Center Frequency} = 1 \text{ kHz} \times 2^{((n-30)/3)}$$

Finally, the SR780 and SR785 only can calculate octave power in 1/12 octave bins whose center frequencies are at:

$$\text{Center Freqeuncy} = 1 \text{ kHz} \times 2^{1/24} \times 2^{n/12}$$

## Swept-Sine Measurements

The SR780 and SR785 contain an additional measurement mode, the swept-sine mode, which is useful for making measurements with high dynamic range. A swept-sine measurement is basically a sine sweep which steps through a specified sequence of frequency points. At each point, the source maintains a constant frequency, and the inputs measure only signals at this frequency. After each point has been measured, the source moves on to the next point in the sequence. Unlike the FFT, which measures many frequencies at once, swept-sine measures one frequency at a time. As we'll see, this technique is somewhat slower but leads to increases in dynamic range.

Transfer functions can be measured using the FFT mode or the swept-sine mode. However, if the transfer function has a large variation within the measurement span, the FFT may not be the best measurement technique. It's limitation comes from the nature of the chirp source that must be used. The FFT simultaneously measures the response at all frequencies within the span . Thus, the source must contain energy at all of the measured frequencies. In the time record, the frequency components in the source add up, and the peak source amplitude within the time record generally exceeds the amplitude of each frequency component by about 30 dB. Since the input range must be set to accommodate the amplitude peak, each component is measured at −30 dB relative to full scale. This effectively reduces the dynamic range of the measurement by about 30 dB! If the transfer function has a variation from 0 to −100 dB within the measurement span, each bin of the FFT must measure signals from −30 dBfs to −130 dBfs. Even with a large number of vector averages, this proves difficult—especially with large measurement spans.

Swept-sine measurements, on the other hand, can optimize the measurement at each frequency point. Since the source is a sine wave, all of the source energy is concentrated at a single frequency, eliminating the 30 dB chirp dynamic range penalty. In addition, if the transfer response drops to −100 dBV, the input range of channel 2 can auto range to −50 dBV and maintain almost 100 dB of signal-to-noise. In fact, simply optimizing the input range at each frequency can extend the dynamic range of the measurement to beyond 140 dB.

For transfer functions with both gain and attenuation, the source amplitude can be optimized at each frequency. Reducing the source level at frequencies where there is gain prevents overloads, and increasing the amplitude where there is attenuation preserves signal-to-noise. To optimize the measurement time of sweeps covering orders of magnitude in frequency, the detection bandwidth can be set as a function of frequency. More time can be spent at lower frequencies and less time at higher frequencies. In addition, frequency points can be skipped in regions where the response does not change significantly from point to point. This speeds measurements of narrow response functions.