WHITE PAPER

Spectrum Analyzer Basics

Fundamentals of swept-tuned, superheterodyne spectrum analyzers

KEYSIGHT

Introducing the Spectrum Analyzer

A simple way to describe a spectrum analyzer is a frequency-selective, peak-responding voltmeter calibrated to display the RMS value of a sine wave. Although a spectrum analyzer displays power, it is a power meter. A signal analyzer can calculate power when a sine wave's peak or average value is known and the resistance across the value measured is known. This white paper describes the basic signal analyzer and how it has evolved considering today's digital technology and digital signal processing that provides many more capabilities.

Contents

Introducing the Spectrum Analyzer 2

What is a Spectrum Analyzer? 4

Frequency Domain vs. Time Domain 5

Types of Spectrum / Signal Analyzers 7

How Does a Spectrum Analyzer Work? 7

Spectrum Analyzer Block Diagram / Components and Their Functions 26

Summary 28

What is a Spectrum Analyzer?

As you analyze various signals with a voltmeter, oscilloscope, or other instrument, have you ever wondered if a spectrum analyzer can provide additional valuable information that may help with your application? Perhaps you are not sure what a spectrum analyzer does or the signal analysis advantages it offers. This white paper describes the spectrum analyzer, providing helpful advice and recommendations for fundamental spectrum analyzer capabilities and settings.

Let's start with what a spectrum analyzer is. A spectrum analyzer decomposes time-varying signals into a collection of sinewaves displaying a frequency spectrum (power versus frequency). Technical experts use individual sinewave components and measurements to troubleshoot, assess quality, or deepen their understanding of complex signals. A spectrum analyzer is an essential test instrument for RF, wireless communication, radar, and related applications.

Applications that use spectrum analyzers

We can look at various spectrum analyzer applications to better understand what a spectrum analyzer is. Spectrum analyzers measure frequency information for signals in applications like spectrum monitoring, noise and distortion, radiofrequency devices, EMI testing for device compliance, and telecommunications analysis of occupied bandwidth and sources of interference.

Spectrum monitoring is a critical spectrum analyzer application. Government regulatory agencies allocate different frequencies for various radio services like television, radio, mobile phones, police, emergency communications, and many more. It is critical that each of these services operates in its assigned frequency and stays within the allocated channel bandwidth. Transmitters must operate at closely spaced adjacent frequencies. Other components in these communications systems, such as power amplifiers, require performance checks to ensure their energy does not affect adjacent signals or cause interference. Wireless communications are keenly interested in finding and resolving out-of-band and spurious emissions. One example is analyzing a cellular radio system, looking for harmonics of the carrier signal to ensure it does not interfere with other systems operating at the same frequencies as the harmonics. Additionally, spectrum analysis can look for distortion of a message modulated onto a carrier signal.

Electromagnetic interference (EMI) is unwanted radiated emissions that can impair other systems and cause adverse effects on performance. Designers and manufacturers of electrical and electronic products must test emissions levels versus frequency according to regulations set by various government agencies and industry standards.

Engineers often test their active circuits and devices for excess noise. Engineers use noise figure and signal-to-noise ratio (SNR) measurements to characterize a device's performance to ensure no excessive noise generation.

Frequency Domain vs. Time Domain

It may be more familiar to think of signal events as they occur in time or within the time domain. An oscilloscope measures signal information and electrical events in the time domain.

We can refer to the Fourier theory when comparing the time domain to the frequency domain. The Fourier theory explains that any time-domain electrical phenomenon includes one or more sine waves of appropriate frequency, amplitude, and phase. In other words, the Fourier transform is a mathematical function that separates a waveform (function of time) into the frequencies that comprise it, resulting in a function of frequency. Now, we use a spectrum analyzer to visualize electronic signals in the frequency domain. The frequency domain measurements indicate how much energy is present at each frequency.

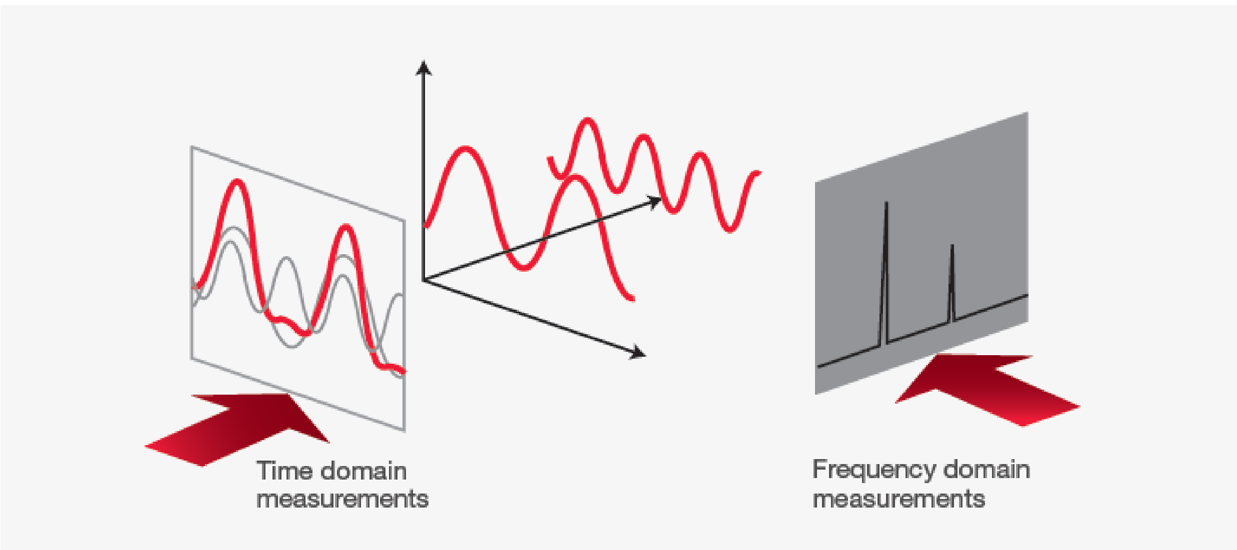

Figure 1. Relationship between the time and frequency domains

Stepping through a Fourier transform example, we can begin with the waveform signal that appears below:

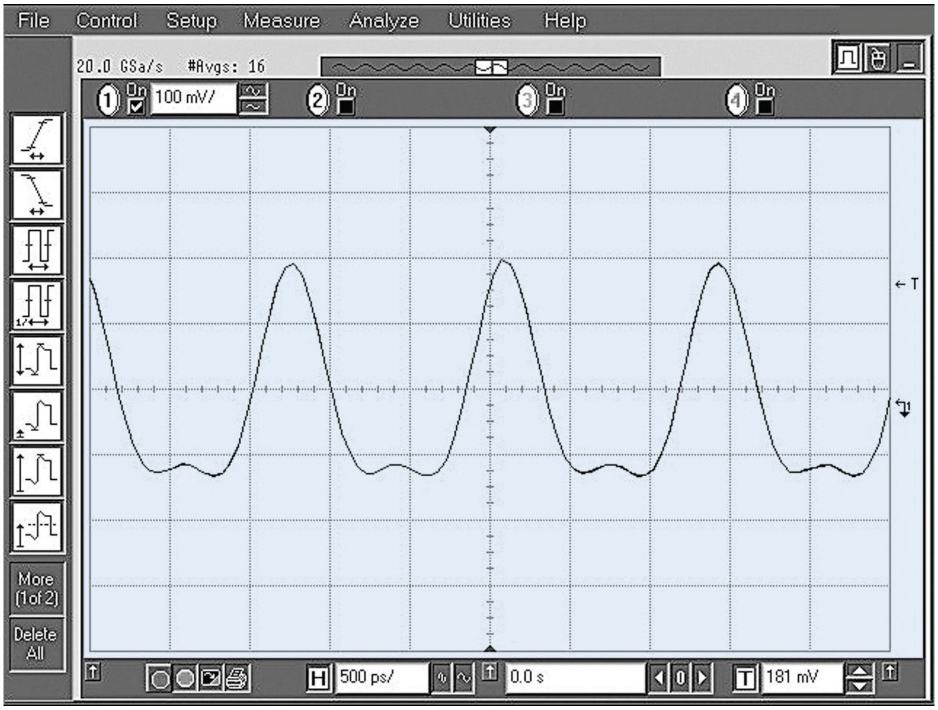

Figure 2. A complex time-domain signal

We can use filtering to separate the waveform into sinusoidal signals or spectral components for independent analysis. Amplitude and phase measurements help to characterize each sine wave. But suppose the signal we intend to analyze is periodic. In that case, Fourier says the essential sine waves appear individually in the frequency domain represented by 1/T, where T is the period of the signal[1].

According to Fourier's theory, when we shift from the time domain to the frequency domain or from the frequency domain to the time domain, the signal evaluation should be overall time over infinity. However, in practice, we always use a finite time period when making a measurement. Making measurements in a finite bandwidth produces acceptable results, capturing most of the signal energy. When you perform a Fourier transformation on frequency domain data, the phase of the individual components is critical. For example, if you do not preserve the phase, a square wave that transforms to the frequency domain and back again, could turn into a sawtooth wave.

[1] If the time signal occurs only once, then T is infinite, and the frequency representation is a continuum of sine waves.

Types of Spectrum / Signal Analyzers

Initial swept-tuned, superheterodyne analyzers measure only amplitude but gradual technology advances required additional measurements. For example, more complex communication systems made phase measurements more critical, and newer spectrum analyzers preserved the signal's phase and amplitude information. Additionally, smaller circuitry has enabled rugged, portable spectrum analyzers. Today's spectrum analyzers digitize the signal after one or more stages of frequency conversion. As a result of more technological advances and new measurement needs, there are various types of analyzers:

- Analog analyzers (see Basic Spectrum Analyzers)
- Vector analyzers (see Vector Signal Analysis Basics application note)
- Fast Fourier transform (FFT) analyzers (see Keysight 35670A FFT Dynamic Signal Analyzer)
- Analyzers that combine these capabilities (see Keysight X-Series Signal Analyzers)
- Ruggedized and portable analyzers (see FieldFox Handheld Spectrum Analyzers)
- Modular analyzers (see Modular PXI Signal Analyzers)

How Does a Spectrum Analyzer Work?

While today's technology makes it possible to replace many analog circuits with modern digital implementations, it is helpful to understand classic spectrum analyzer architecture and how it works. Using examples in this section, we will describe and demonstrate spectrum analyzer fundamentals and their simplest, most important configuration elements and measurements.

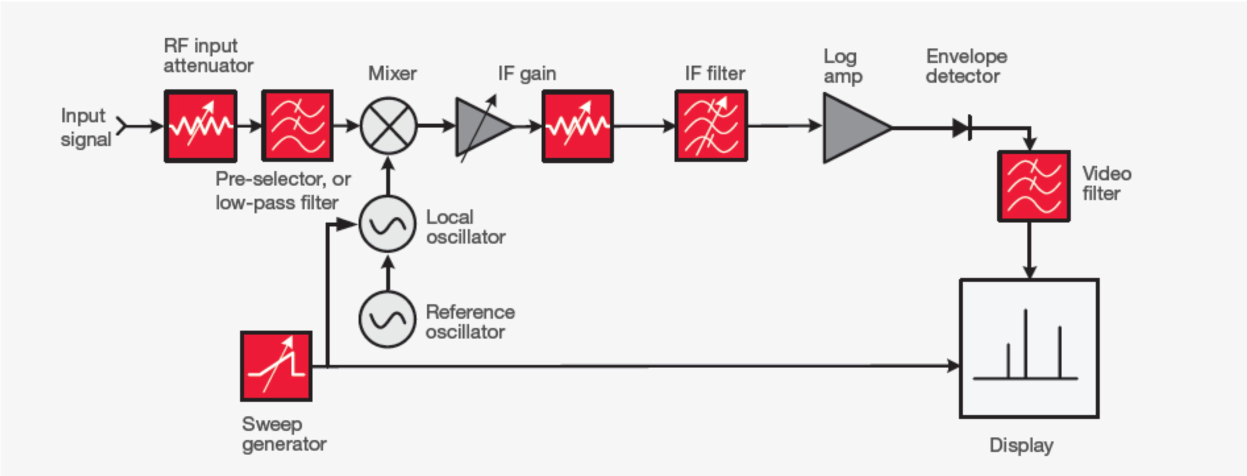

Figure 3. Block diagram of a classic superheterodyne spectrum analyzer

Spectrum analyzer display and settings

The output of a spectrum analyzer is an X-Y trace on a display. The display appears on a grid (graticule) with 10 main horizontal divisions and 10 main vertical divisions.

Horizontal axis

The horizontal axis is linearly calibrated in frequency that increases from left to right. By setting the frequency, you can determine the absolute frequency of any signal displayed and the relative frequency difference between two signals. There are two ways you can set the frequency:

1. Set the center frequency and span. The controls are independent, so changing the center frequency does not alter the frequency span.
    a. First, adjust the frequency at the centerline of the graticule with the center frequency control.
    b. Then, adjust the frequency range (span) across the full 10 divisions with the frequency span control.
2. Set the start and stop frequencies.

Vertical axis

You can calibrate the vertical axis in amplitude. Choose a linear scale calibrated in volts or a logarithmic scale calibrated in dB. A log scale is used more often because it has a much wider usable range. Additionally, the log scale allows signals as far apart in amplitude as 70 to 100 dB (voltage ratios of 3,200 to 100,000 and power ratios of 10,000,000 to 10,000,000,000) to display simultaneously.

Figure 4. Typical spectrum analyzer display with control settings

Alternatively, the linear scale is usable for signals that differ by no more than 20 to 30 dB (voltage ratios of 10 to 32). In either case, the top line of the graticule, the reference level, is set to an absolute value using calibration techniques and uses the scaling per division to assign values to other locations on the graticule. This enables us to measure either the signal's absolute value or the relative amplitude difference between any two signals. The display includes the scale calibration, frequency, and amplitude, as pictured in Figure 4.

Tuning the analyzer

Knowing how to tune a spectrum analyzer to a desired frequency range is essential. Tuning relates to the intermediate frequency (IF) filter's center frequency, the local oscillator's (LO) frequency range, and the frequency range that passes through the low-pass filter to reach the mixer. We must select the best LO frequency and an IF to achieve an optimal tuning range for our analyzer.

Example 1: Let's say we want a tuning range from 0 to 3.6 GHz. First, we need to select the IF. We can try 1 GHz since it is within our designated tuning range. The mixer output also includes the original input signals, so an input signal at 1 GHz would give us a constant mixer output at the IF. The 1 GHz signal would pass through the system and give us a constant amplitude response on the display regardless of the LO tuning. Unfortunately, the result would be a hole in the frequency range where we could not analyze signals because the amplitude response would be independent of the LO frequency. A 1 GHz IF will not work.

Example 2: Let's choose an IF that is above the highest frequency in our tuning range — 0 Hz to 3.6 GHz. Using the Keysight X-Series signal analyzer, we can tune to 3.6 GHz. The first LO frequency range is 3.8 to 8.7 GHz; for this example, the IF is 5.1 GHz. If we start the LO frequency at the IF filter frequency (LO minus IF = 0 Hz) and tune it upward from that point to 3.6 GHz above the IF, we can cover the tuning range with the LO minus the IF mixing product. Using this information, we can generate a tuning equation:

$F_{sig} = f_{LO} - f_{IF}$

Where f_{sig} = signal frequency

f_{LO} = local oscillator frequency

f_{IF} = intermediate frequency

To determine the LO frequency needed to tune the analyzer to a low-, mid-, or high-frequency signal (let's use 1 kHz, 1.5 GHz, and 3 GHz), restate the tuning equation in terms of f_{LO}:

$f_{LO} = f_{sig} + f_{IF}$

Then, apply the numbers for the signal and IF in the tuning equation:

f_{LO} = 1 kHz + 5.1 GHz = 5.100001 GHz

f_{LO} = 1.5 GHz + 5.1 GHz = 6.6 GHz

f_{LO} = 3 GHz + 5.1 GHz = 8.1 GHz

Figure 6. demonstrates a spectrum analyzer's tuning and shows that f_{LO} is not high enough to cause the $F_{LO}-F_{sig}$ mixing product to fall in the IF passband, so there is no response. By adjusting the ramp generator to tune the LO higher, the mixing product will be in the IF passband at some point on the ramp (sweep), and you will see a response.

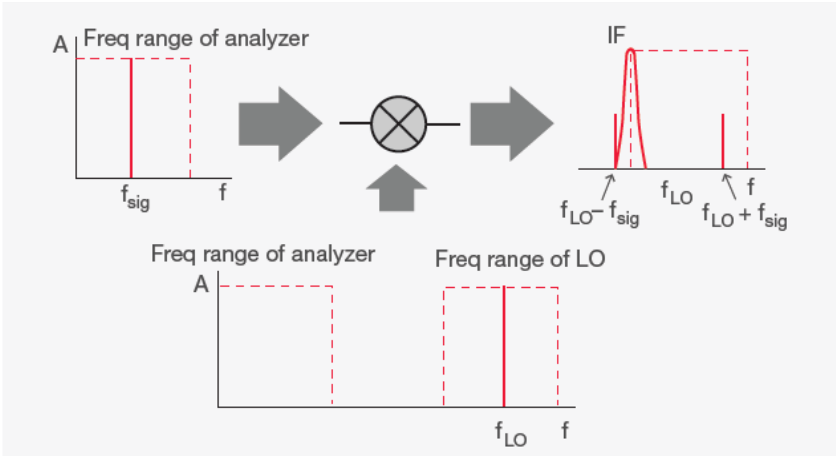

Figure 6. Tune the LO to f_{IF}+f_{sig} to produce a displayed response

The ramp generator controls both the horizontal position of the trace on the display and the LO frequency so we can calibrate the display's horizontal axis in terms of the input signal frequency.

For the tuning example in Figure 6., let's adjust the input frequency to 9.0GHz. As the LO tunes through its 3.8 to 8.7 GHz range, it reaches 3.9 GHz, which is IF away from the 9.0 GHz input signal. At this frequency, we have a mixing product that is equal to the IF creating a response on the display, or:

$F_{sig} = f_{LO} + f_{IF}$

The equation indicates that the spectrum analyzer's architecture could also support a tuning range of 8.9 to 13.8 GHz for in-range signals that reach the mixer.

Follow these guidelines for tuning a single-band RF spectrum analyzer:

- Choose an IF above the highest frequency of the tuning range.
- Make the LO tunable from the IF to the IF and the upper limit of the tuning range.
- Include a low-pass filter in front of the mixer that cuts off below the IF.

Tuning for more closely spaced signals may require a spectrum analyzer with IF bandwidths as narrow as 1 kHz, 10 Hz, or even 1 Hz. Filters this narrow are challenging to achieve at 5.1 GHz, so we must add two to four mixing stages to down-convert from the first IF. Figure 7 shows an IF chain in a typical spectrum analyzer. The figure only shows passive filters, although the implementation includes amplification in the narrower IF stages. Depending on the analyzer's design, the final IF may contain other components like logarithmic amplifiers or analog-to-digital converters.

Most RF spectrum analyzers permit an LO frequency as low as, and even below the first IF. In this case, the LO appears at the mixer output because there is finite isolation between the LO and IF ports of the mixer. When the LO equals the IF, the system processes the LO signal that appears as a response on the display as if it were a 0 Hz input signal. The LO feedthrough response can mask very low-frequency signals, so not all analyzers allow the display range to include 0 Hz.

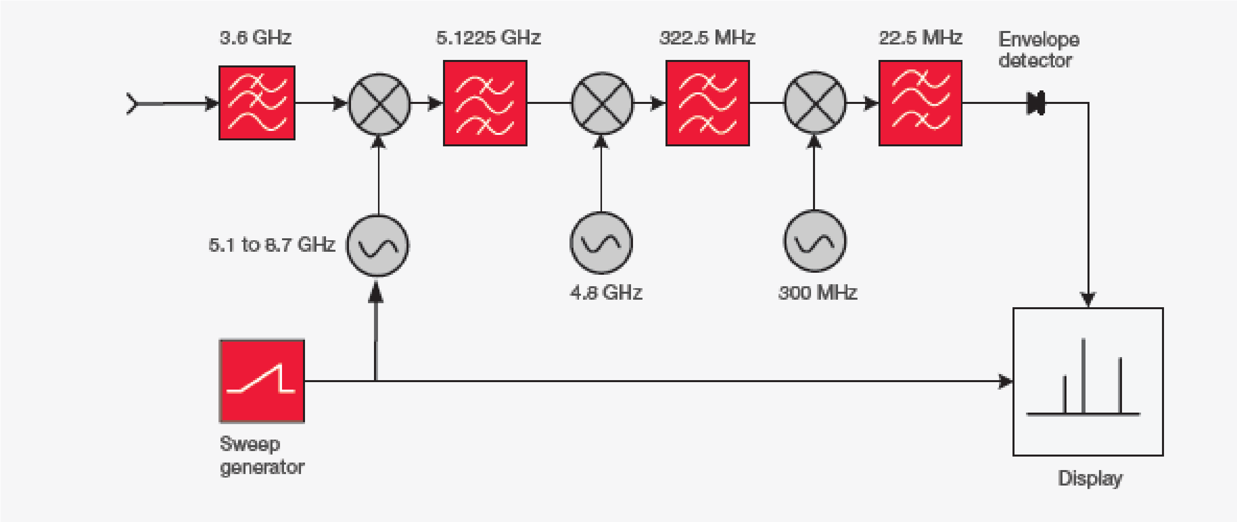

Figure 7. Most spectrum analyzers use two to four mixing stages to reach the final IF

For tuning this spectrum analyzer, use the following formula:

$F_{sig} = f_{LO1} - (f_{LO2} + f_{LO3} + f_{LOfinalIF})$

In addition, make note to simplify the tuning equation by using just the first IF results in the same answer:

$f_{LO2} + f_{LO3} + f_{LOfinalIF}$ = 4.8 GHz +300 MHz + 22.5 MHz

= 5.1225 GHz, the first IF

Analog filters

Frequency resolution is a spectrum analyzer's ability to separate two input sinusoidal signals into distinct responses. Fourier says a sinewave signal only has energy at one frequency, so there should be no resolution issues. Two signals, regardless of how close they are in frequency, should appear as two lines on the spectrum analyzer's display. These signals must be far enough apart on the display, or the signal traces may fall on top of each other and appear as a single response. Use the spectrum analyzer's selectable resolution IF filters and select one narrow enough to resolve closely spaced signals. A closer look at a superheterodyne receiver demonstrates why displayed signal responses have definite widths.

A mixer's output includes the two original signals (input and LO) and the sum and difference products. A bandpass filter determines the intermediate frequency, selects the desired mixing product, and rejects other signals. We know that the input signal remains fixed, and since the LO is swept, the products from the mixer are also swept. During a sweep, if a mixing product sweeps past the IF, the bandpass filter's characteristic shape is traced on the display, as shown in Figure 8. The narrowest filter in the chain determines the overall displayed bandwidth. In Figure 7, the narrowest filter and resulting bandwidth is 22.5 MHz IF.

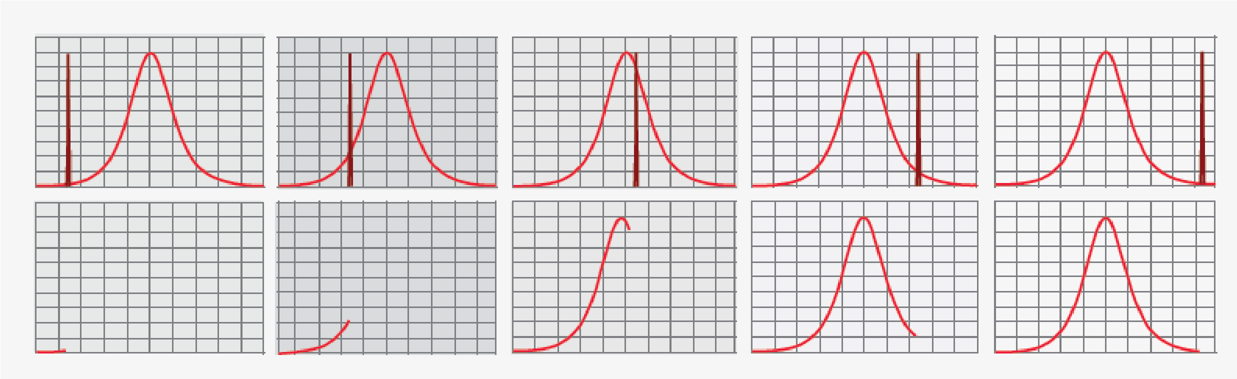

Figure 8. As a mixing product sweeps past the IF, the filter shape traces appear on the display

It is possible to resolve two very close signals by separating them using the 3-dB bandwidth of a selected IF filter. The IF filter's 3-dB bandwidth value determines how close together equal-amplitude sinusoids can be and still resolve. In the case of the example in Figure 9, you see an approximate 3-dB dip in the trace between the two signal peaks. Although the signals could be even closer before merging completely, using the IF filter's 3-dB bandwidth rule-of-thumb is a suitable method for resolving equal amplitude signals.

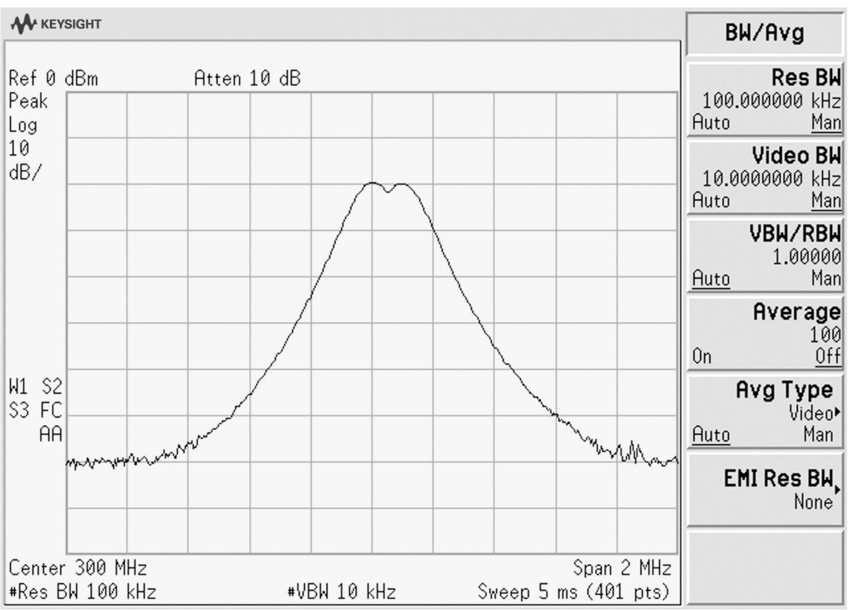

Figure 9. It is possible to resolve two equal-amplitude sinusoids separated by a selected IF filter's 3-dB bandwidth

It is more common to have sinusoidal signals that are not equal in amplitude and where the lower-amplitude signal hides in the skirt of the higher-amplitude signal's trace. In the Figure 10 example, the top trace looks like a single signal but consists of a 300 MHz (0 dBm) and 300.005 MHz (-30 dBm) signal. The lower trace demonstrates the resulting trace after removing the 300 MHz signal.

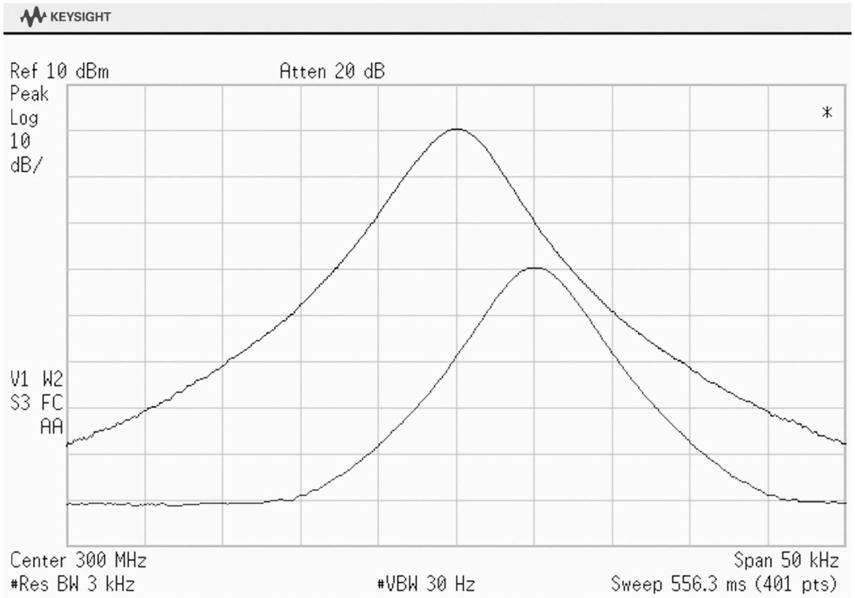

Figure 10. A low-amplitude signal may hide under the skirt of a higher-amplitude signal

Resolution filters also include a bandwidth selectivity specification, referred to as "selectivity" or "shape factor." A filter's bandwidth selectivity helps to determine the resolving power for unequal sinusoidal signals. For Keysight analyzers, bandwidth selectivity is the ratio of the 60 dB bandwidth to the 3 dB bandwidth, as shown in Figure 11. For this example, the analyzer is a four-pole, synchronously tuned design with a nearly Gaussian shape and a bandwidth selectivity of approximately 12.7:1. Analyzers with digital filters can achieve even better bandwidth selectivity.

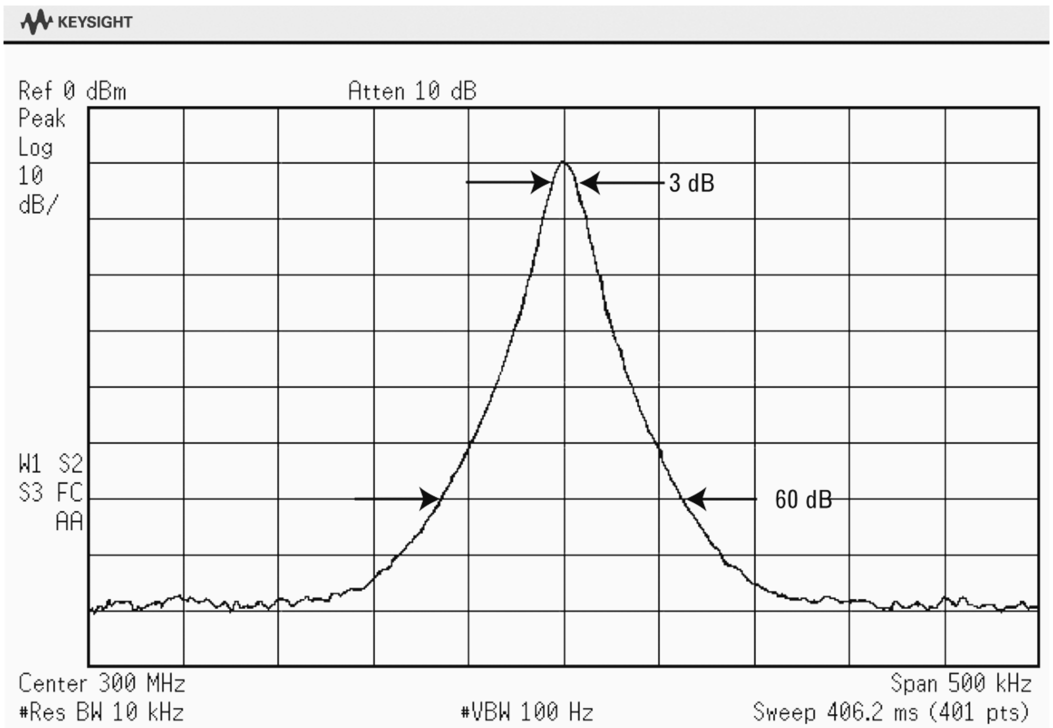

Figure 11. Bandwidth selectivity ratio of 60 dB to 3 dB bandwidths

For example, what resolution filter's bandwidth should we choose to resolve signals that differ by 4 kHz and 30 dB, assuming 12.7:1 bandwidth selectivity?

We need to consider the frequency difference from the filter center frequency to the skirt rather than the full bandwidth because the larger signal may be rejected when the analyzer is tuned to the smaller signal. To determine how far down the filter skirt is at a given offset, we use the following equation:

$$H(\Delta f) = -10(N) \log_{10} [(\Delta f/f_0)^2 + 1]$$

Where $H(\Delta f)$ is the filter skirt rejection in dB, N is the number of filter poles, Δf is the frequency offset from the center in Hz, and f_0 is:

$$\frac{RBW}{2\sqrt{2^{1/N}-1}}$$

For our example, N = 4 and Δf = 4000. Let's begin by trying the 3-kHz RBW filter. First, we compute f_0:

$$f0 = \frac{3000}{2\sqrt{2^{1/4}-1}} = 3448.44$$

Now, we can determine the filter rejection at a 4-kHz offset:

H(4000) = –10(4) log10 [(4000/3448.44)2 + 1] = −14.8 dB

This is not enough to enable us to see the smaller signal. Let's determine H(Δf) again using a 1-kHz filter:

$$f0 = \frac{1000}{2\sqrt{2^{1/4}-1}} = 1149.48$$

This allows us to calculate the filter rejection:

H(4000) = –10(4) log10[(4000/1149.48)2 + 1] = −44.7 dB

The 1 kHz resolution bandwidth filter successfully resolved the smaller signal, as shown in Figure 12.

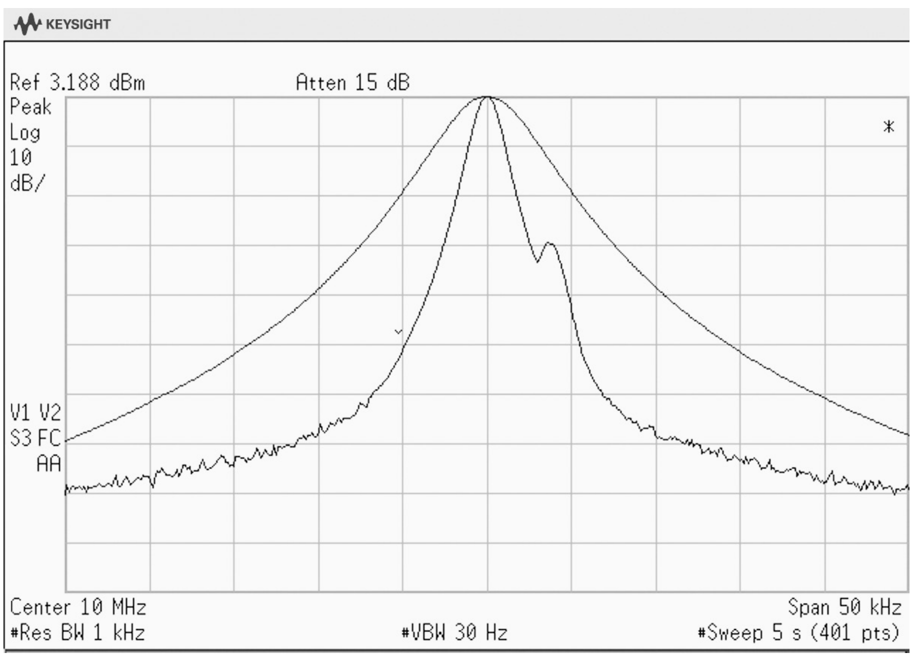

Figure 12. A 3 kHz filter (top trace) does not resolve the smaller signal. However, a 1 kHz filter (bottom trace) resolves the smaller signal.

Digital filters

An advantage of digital filters in spectrum analyzers is that they dramatically improve bandwidth selectivity. Keysight's PSA and X-Series signal analyzers implement the resolution bandwidths digitally. Other analyzers, like the Keysight ESA-E Series, use analog filters for wider bandwidths and digital filters for bandwidths of 300 Hz and below.

Residual FM

Determine the minimum usable resolution bandwidth from the instability and residual FM of the analyzer's LOs, especially the first LO. Today's analyzers dramatically improve residual FM compared to earlier models, which were unable to resolve bandwidths narrower than 1 kHz due to an undetermined cause of display instability. Keysight analyzers include:

- High-performance X-Series signal analyzers with a residual FM of nominally 0.25 Hz
- PSA Series analyzers' residual FM is 1 to 4 Hz
- ESA Series analyzers' residual FM is 2 to 8 Hz

These residual FM signal analyzer examples enable bandwidths as low as 1 Hz. Any observed instability is due to the incoming signal.

Phase noise

Phase noise, or sideband noise, is ultimately due to the LOs in the spectrum analyzer. While we may not be able to see frequency jitter of a spectrum analyzer LO system, we can observe a manifestation of the LO frequency or phase instability. The initial LO instability transfers to the spectrum analyzer's mixing products due to the LO and input signals. This means the LO phase noise modulation sidebands appear around any spectral component on the display that is far enough above the broadband noise floor of the system, as shown in Figure 13. The amplitude difference between a displayed spectral component and the phase noise is a function of the LO stability. A stable LO will result in lower phase noise. The amplitude difference is also a function of the resolution bandwidth. Reducing the resolution bandwidth by a factor of 10 causes the displayed phase noise lever to decrease by 10 dB.

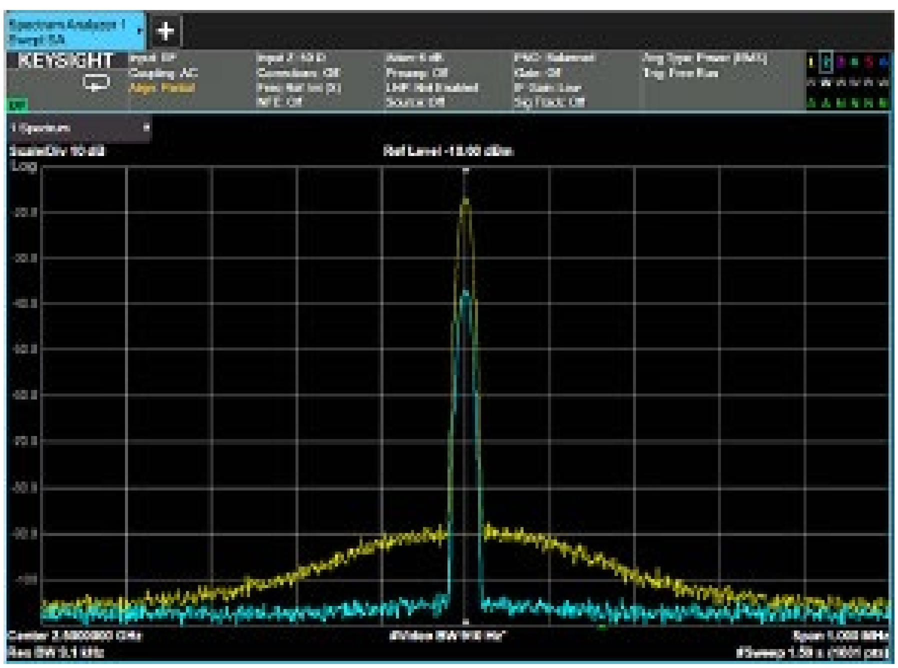

Figure 13. Phase noise displays only when a signal appears far enough above the system noise floor

The shape of the phase noise spectrum relates directly to the analyzer design, specifically the complexity of the phase-lock loops used to stabilize the LO. Phase noise is specified in terms of dBc, dB relative to a carrier, and normalized to a 1 Hz noise power bandwidth. Additionally, we can represent phase noise at a specific frequency offset or present a curve that shows the phase noise characteristics over a range of offsets.

You will typically see a spectrum analyzer's inherent phase noise just in narrower resolution filters when it obscures the lower skirts of the filters. Unfortunately, digital filters do not change this inherent phase noise. The phase noise can hide under the filter skirt for wider filters, as in the case of the two unequal sinusoids discussed earlier.

Spectrum analyzers enable you to select different LO stabilization modes to optimize the phase noise for different measurement conditions. For example, the high-performance X-Series signal analyzers offer three LO stabilization modes:

- Optimize phase noise for frequency offsets < 140 kHz from the carrier.
  - Optimizes LO phase noise for the area close to the carrier at the expense of phase noise beyond 140-kHz offset.
- Optimize phase noise for frequency offsets > 160 kHz from the carrier.
  - Optimizes phase noise for offsets above 160 kHz away from the carrier.
- Optimize LO for fast tuning.
  - LO behavior compromises phase noise at all offsets from the carrier below approximately 2 MHz, minimizes measurement time, and allows the maximum measurement throughput when changing the center frequency or span.

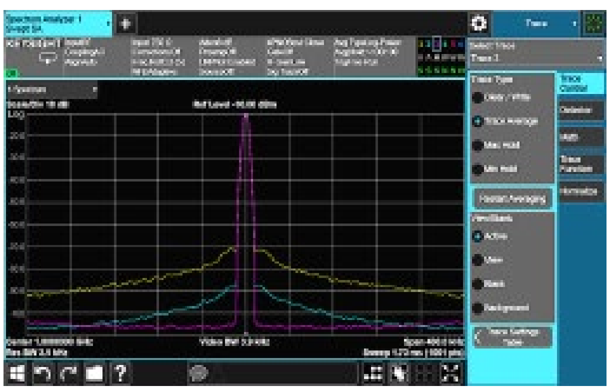

Figure 14a. Optimization of phase noise performance for different measurement conditions

Figure 14b. Detail of a 140 kHz carrier offset region

Spectrum analyzers, such as the high-performance X-Series signal analyzer, have an auto mode for phase noise optimization that automatically sets the instrument to optimize speed or dynamic range for various operating conditions.

For example:

- Fast tuning mode
  - when the span is >44.44 MHz or the RBW is >1.9 MHz
- Best close-in phase noise
  - when center frequency < 195 kHz
  - or when center frequency ≥ 1 MHz and span ≤ 1.3 MHz and RBW ≤ 75 kHz
- Best wide-offset phase noise
  - if the previous conditions are unmet

In any case, phase noise is the main limitation of a spectrum analyzer's ability to resolve signals of unequal amplitude. Figure 15 shows two resolved signals based on the 3 dB bandwidth and selectivity. Unfortunately, the phase noise covers the smaller signal.

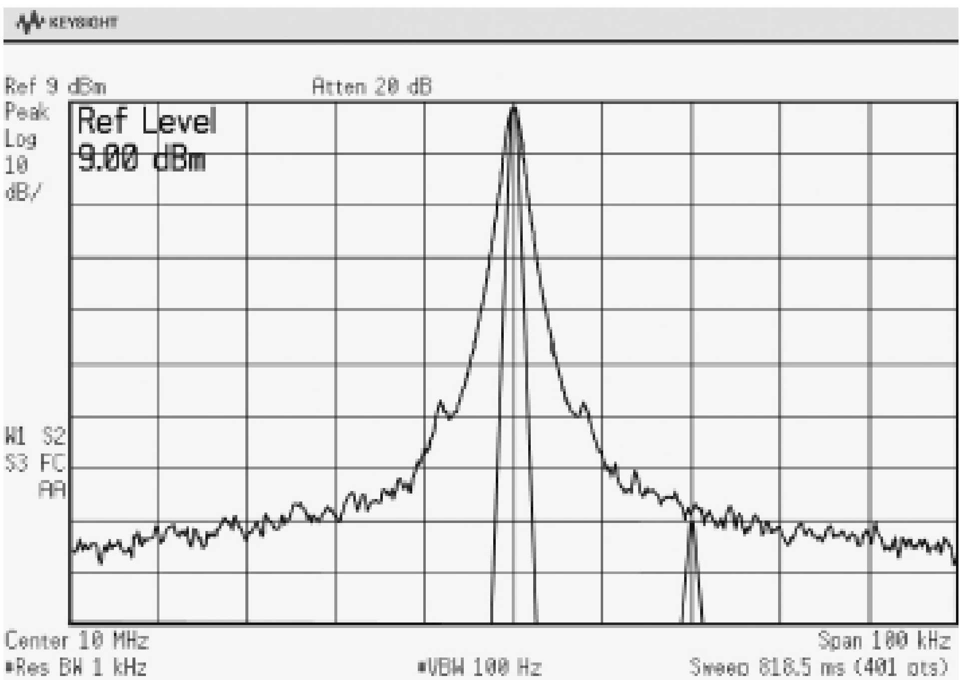

Figure 15. Phase noise prevents the resolution of unequal signals

Sweep time

Analog resolution filters

If a spectrum analyzer's resolution stood on its own, the narrowest possible resolution IF filter would be the best choice. However, resolution impacts sweep time, which affects how long measurement completion takes.

Resolution is a concern because the IF filters are band-limited circuits that require finite times to charge and discharge. If the mixing products sweep through the IF filters too quickly, the displayed amplitude will be less than it should be, as shown in Figure 16. The length of time a mixing product stays in the IF filter pass band is directly proportional to bandwidth and inversely proportional to the sweep in Hz per unit time, or:

Time in pass band = RBW/(Span/ST) = (RBW)(ST)/Span

Where RBW = resolution bandwidth and ST = sweep time

Alternatively, the filter's rise time is inversely proportional to its bandwidth. If we include a constant of proportionality, k, then:

Rise time = k/RBW

If we make the terms equal and solve for sweep time:

k/RBW = (RBW)(ST)/Span

or $ST = k(Span)/RBW^2$

The k value is between two and three for the synchronously timed, near-Gaussian filters used in many analog analyzers.

It is most important to remember that a change in resolution dramatically affects sweep time. The sweep time of older analog analyzers was affected by a factor of about 10 with each step in resolution. Keysight X-Series signal analyzers provide a bandwidth step of 10% for a better compromise between span, resolution, and sweep time.

Spectrum analyzers automatically couple sweep time to the span and resolution bandwidth settings. Sweep time adjusts automatically to maintain a calibrated display. However, manual override of the automated settings is possible. Additionally, if you select a sweep time outside the available range, the analyzer indicates the display is uncalibrated with a "Meas Uncal" message in the upper right of the graticule.

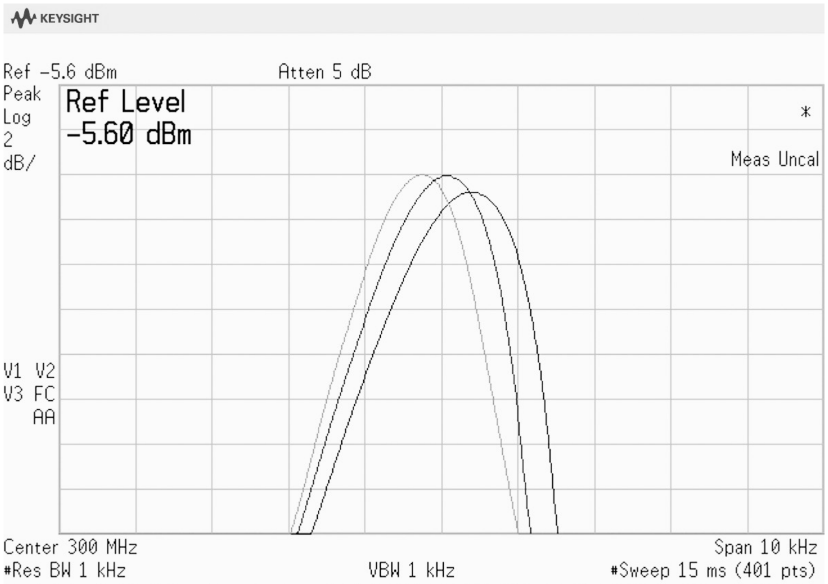

Figure 16. When a spectrum analyzer's sweep time is too fast, there can be a drop in displayed amplitude and a shift in the indicated frequency

Digital resolution filters

Digital resolution filters have a different effect on sweep time than analog filters. For swept analysis, the speed of digitally implemented filters and no additional processing can provide a 2X to 4X improvement.

Keysight X-Series signal analyzers with Option FS1 correct effects of sweeping too fast for resolution bandwidths between about 3 kHz and 100 kHz. The correction enables sweep times to be milliseconds rather than many seconds based on the settings. Referring to Figure 17, the sweep time without correction would be 79.8 seconds. Option FS1 enables a sweep time of just 1.506 seconds. Sweep times are very short for the widest resolution bandwidths, as demonstrated using the formula with k=2 on a span of 1 GHz and an RBW of 1 MHz where the sweep time is 2 ms.

Keysight's X-Series use fast Fourier transforms (FFTs) to process the data and achieve shorter sweep times for narrower resolution bandwidths. This is possible by using frequency blocks to analyze signals. For example, if a frequency block is 1 kHz, and we select a 10 Hz resolution bandwidth, the analyzer simultaneously processes the data in each 1 kHz block through 100 contiguous 10 Hz filters. The fast sweep times make it easy to see the benefits of digital filters.

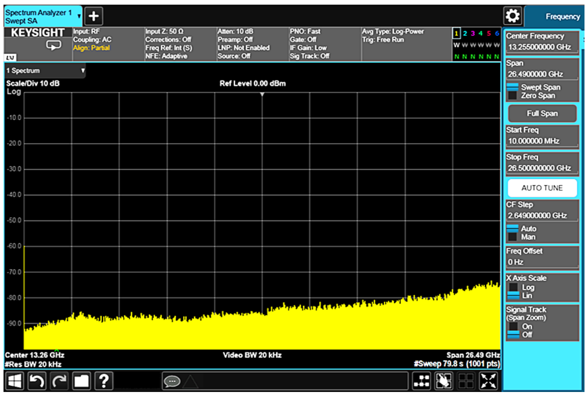

Figure 17a. Full span sweep speed, RBW of 20 kHz without Option FS1

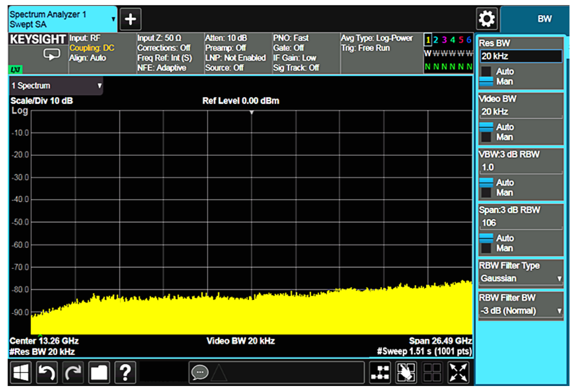

Figure 17b. Full span sweep speed, RBW of 20 kHz with Option FS1

Display detector types

Each data point of a displayed signal must represent what occurs over some frequency range and time interval. The calculation of each data interval extracts the desired information from the input signal, stores it in memory, and writes it to the display. This process enables great flexibility, resulting in six different detector types: sample, positive peak (peak), negative peak, normal, average, and quasi-peak.

When gathering data, detectors save data in a bucket, as shown in Figure 18. The bucket contains data from a frequency span and time frame.

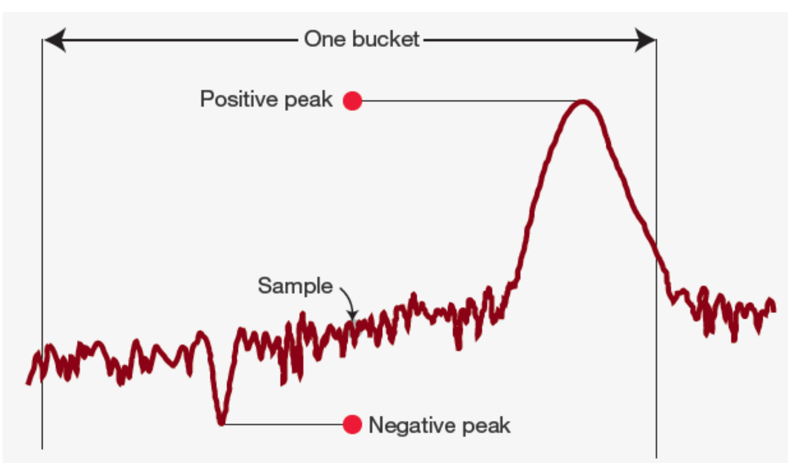

Figure 18. The trace point saved in memory is based on the detector type of algorithm

Equations define the type of applied detector:

Frequency: bucket width = span/(trace points – 1)

Time: bucket width = sweep time/(trace points – 1)

Spectrum analyzer sampling rates are different for various instruments. However, decreasing the span or increasing the sweep time will result in better accuracy because the number of samples per bucket is greater.

Refer to Figure 18, which demonstrates sample, peak, and negative peak detection.

Sample detection

Sample is a data point's instantaneous level at the center of each bucket. Drawing vectors between the sample points of each bucket results in a continuous signal display, as shown in Figure 19. Comparing the displays in Figure 19 shows that more trace points provide a better analog signal replication. The number of display points can vary for different analyzers. For example, the number of frequency domain traces of a Keysight X-Series analyzer is selectable from 1 point to 40,001 points.

The sample detection mode is good at indicating random noise but not very good for analyzing sinusoidal signals. Take the example of looking at a 100 MHz comb using an X-Series signal analyzer with a span from 0 to 26.5 GHz. Even with 1,001 display points, each display point represents a span (bucket) of 26.5 MHz, which is much wider than the maximum 8 MHz resolution bandwidth. Display of the signal's true amplitude would only occur if the mixing product happened to be at the IF center. Therefore, sample detection does not capture all the signals or reflect the signal's true peak values. From this example, we learn that the sample detection mode can produce erroneous results when the resolution bandwidth is narrower than the sample interval (bucket width).

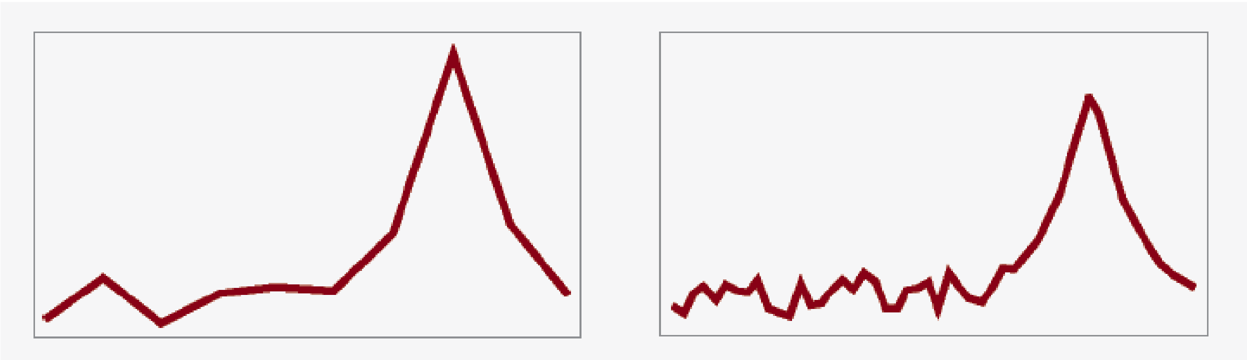

Figure 19. A sample display mode with fewer points (left) and more points (right) produces a better replication of the analog signal

Peak detection

Peak detection helps to ensure the representation of a signal's true amplitude by displaying the maximum value for each bucket. Peak is the default mode for many spectrum analyzers because it prevents missed signals regardless of the resolution bandwidth and bucket width ratio. However, unlike sample mode, peak detection does not provide a good indication of random noise because it only displays the maximum value in each bucket. Most spectrum analyzers with a primary peak detection, generally offer sample detection as an alternative.

Negative peak detection

Negative peak detection displays a signal created by the minimum value of each bucket and is available in most spectrum analyzers. Negative peak detection is not as common as other detection modes but is helpful when differentiating continuous waveform (CW) signals from impulsive signals in EMC testing.

Normal detection

Normal detection provides a better visual display of random noise than peak detection and does not miss signals like the sample mode. When a signal both rises and falls (determined by positive peak and negative peak detectors), normal detection will classify the signal as noise. When noise occurs, an odd-numbered data point displays the maximum value encountered within the bucket. An even-numbered data point displays the minimum value within the bucket. As its mixing product sweeps past the IF filter, a sinusoidal signal appears only to rise or fall as it approaches the filter's center frequency. The positive-peak and negative-peak detectors' sense and amplitude change in one direction, and the maximum value in each bucket gets output to the display, as shown in Figure 20.

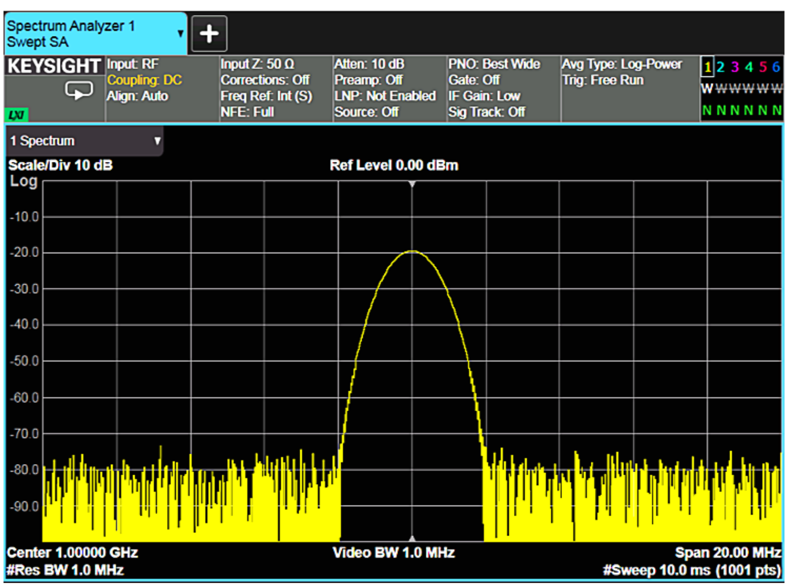

Figure 20. Normal detection displays maximum values in buckets where the signal only rises or falls

Problems can occur when using normal detection if the resolution bandwidth is narrow relative to the bucket. For example, if the signal rises and falls within an odd-numbered bucket, the analyzer will plot the maximum peak value, which is what you want to display. Alternatively, if the signal rises and falls within an even-numbered bucket, the analyzer only plots the minimum value it encounters. The minimum value may differ from the true peak value depending on the resolution bandwidth and the bucket width ratio. In an extreme case where the bucket is much wider than the resolution bandwidth, the minimum and maximum difference within the bucket is similar to the difference between the peak signal value and noise, as shown in Figure 21, bucket 6. This example shows, for an odd-numbered bucket, a comparison of the previous bucket (6) to the current bucket (7) with a resulting display of the greater of the two values in the current bucket (7).

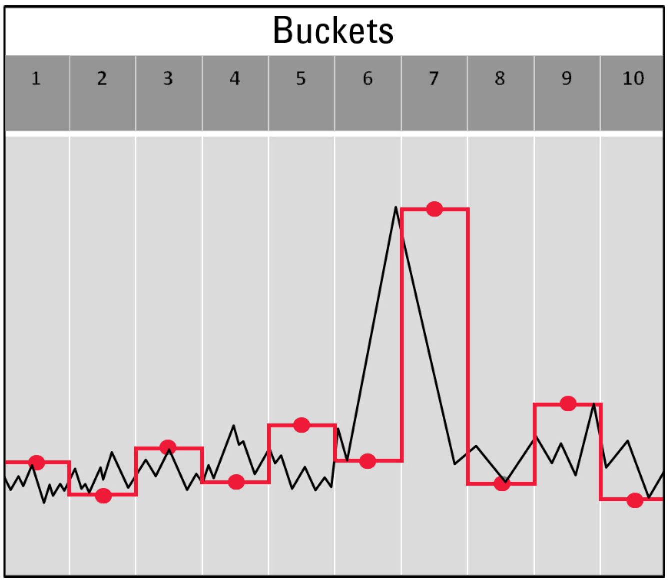

Figure 21. The normal detection algorithm selects the trace point

Signal within a bucket	Even-numbered bucket	Odd-numbered bucket
Rises and falls	Displays the minimum (negative peak) value in the bucket. Remembers the maximum peak.	Displays the maximum (positive peak) value determined by comparing the previous bucket peak (remembered) with the current peak.
Only rises or only falls	Displays the peak	Displays the peak

Table 1. Normal detection algorithm

The normal detection algorithm in Table 1 may display a maximum value (peak) one data point too far to the right, but the offset is typically a small percentage of the span. Some spectrum analyzers compensate for this potential effect by moving the LO start and stop frequencies.

Another normal detection error occurs when one peak appears as two peaks, as seen in Figure 22. The display outlines the two peaks using peak detection with a wider RBW.

From the detection error, we learn that peak detection is best for locating CW signals well of the noise. Sample detection is ideal for looking at noise, while normal detection is best for viewing signals and noise.

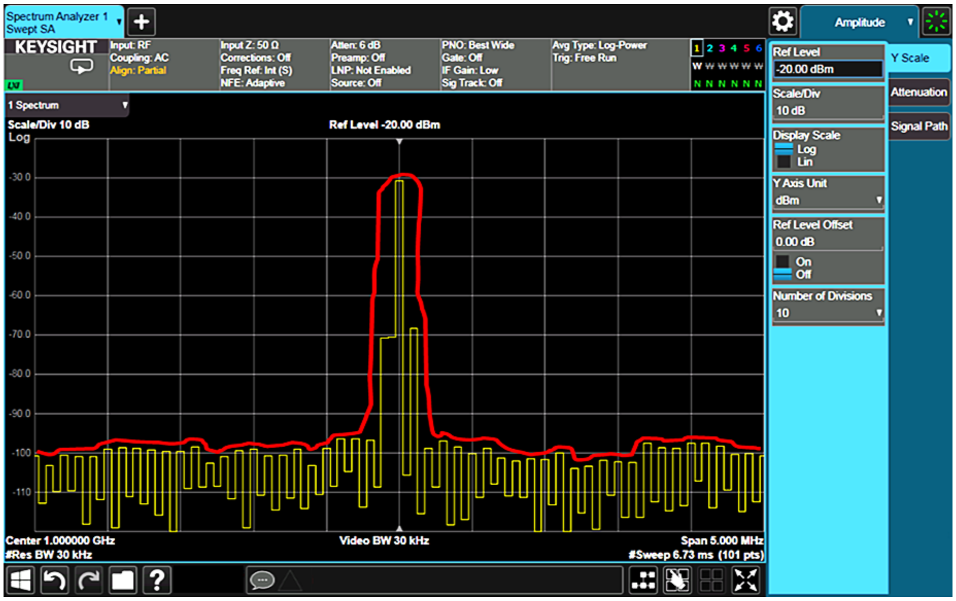

Figure 22. Normal detection can show two peaks when only one peak exists

Average detection

Average detection uses all the data values that collect within the bucket's time and frequency interval rather than taking samples and dismissing data that does not represent a positive or negative peak. In addition, average detection includes different options:

- Power (RMS) averaging: Computes RMS levels by taking the square root of the average of the squares of the voltage data measured during the bucket interval. The analyzer then squares this voltage and divides it by the input impedance of the spectrum analyzer, which is usually 50 ohms. This method is best for measuring the power of complex signals.
- Voltage averaging: Averages the linear voltage data of the envelope signal measured during the bucket interval. Voltage averaging is a good general-purpose average detector often used in EMI testing for narrowband signal measurement.
- Log-power (video) averaging: Averages the logarithmic amplitude values in dB of the envelope signal measured during the bucket interval. This method works best for observing sinusoidal signals near noise.

The advantage of average detection over sample detection relates specifically to power determination. For example, sample detection requires multiple sweeps to collect enough data points to provide accurate average power information. Alternatively, average detection captures all the power information available and has a lower variance result for the same measurement time.

Spectrum Analyzer Block Diagram / Components and Their Functions

Below is a simplified block diagram of a superheterodyne spectrum analyzer. Heterodyne means to combine two nearly equal high frequencies to produce a lower frequency. "Super" refers to superaudio frequencies or frequencies above the audio range. The block diagram shows an input signal passing through an attenuator, then through a low-pass filter to a mixer, where it combines with a signal from the LO.

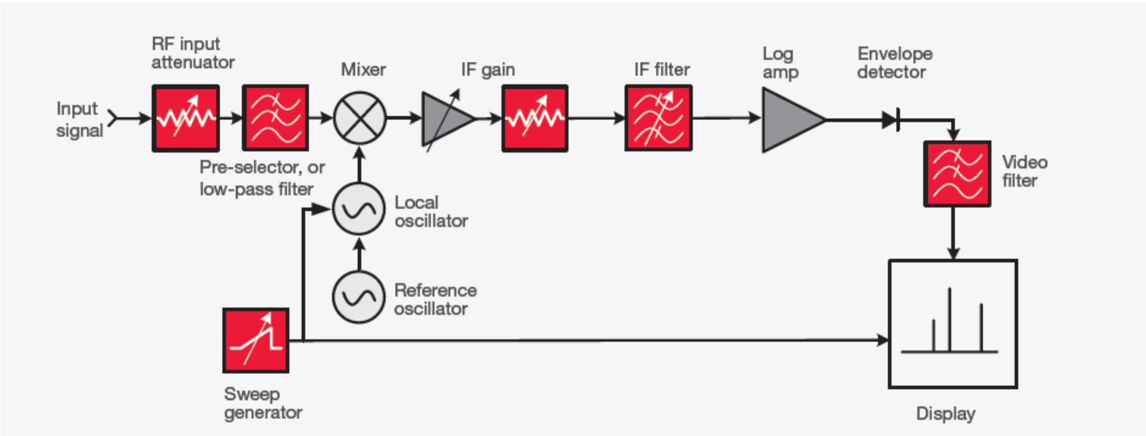

Figure 23. Block diagram of a classic superheterodyne spectrum analyzer

RF attenuator

The spectrum analyzer begins with the RF input attenuator that ensures the signal enters the mixer at the optimal level to prevent overload, gain compression, and distortion. This protective attenuation is typically set automatically, based on the reference level. However, manual attenuation is selectable in steps of 10, 5, 2, or 1 dB. The attenuator circuit below in Figure 24 is an example, providing a maximum attenuation of 70 dB in increments of 2dB.

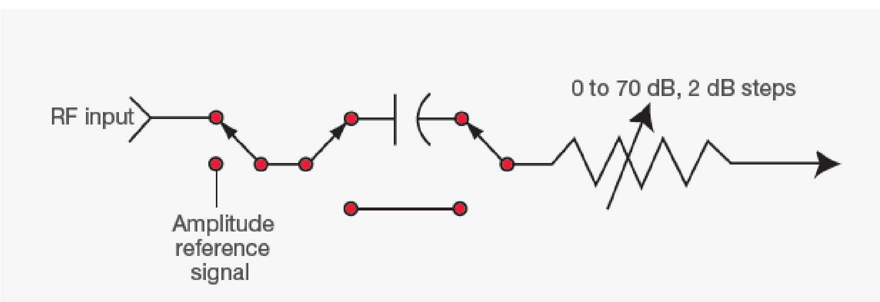

Figure 24. RF input attenuator circuitry

The blocking capacitor helps to prevent damage to the analyzer from a DC signal or a DC offset of a viewed signal. Conversely, it also attenuates low-frequency signals and increases the minimum useable start frequency of the analyzer to 9 kHz, 100 kHz, or 10 MHz, depending on the analyzer.

In some analyzers, an amplitude reference signal input is available to provide a precise frequency and amplitude signal for the analyzer to use for periodical self-calibration.

Low-pass filter or preselector

The low-pass filter blocks high-frequency signals from reaching the mixer, preventing out-of-band signals from mixing with the LO and creating unwanted responses on the display. Microwave spectrum analyzers replace the low-pass filter with a preselector, which is a tunable filter that rejects all frequencies except those we want to view.

Mixer

Because the mixer is a non-linear device, its output includes the two original signals, plus their harmonics and the sums and differences of the original frequencies and their harmonics. Any mixed signals are process when within the IF filter pass band(amplified and compressed on a logarithmic scale). The envelope detector rectifies the signal, which then filters through the low-pass filter and displays. A ramp generator creates the horizontal movement across the display from left to right. The ramp also tunes the LO so its frequency change is in proportion to the ramp voltage.

Variable gain amplifier

The variable gain amplifier follows the mixer in the spectrum analyzer architecture in Figure 37. The gain amplifier adjusts the signal's vertical position on the display without affecting the signal level at the input mixer. When we change the IF gain, the reference level value adjusts accordingly to preserve the correct value for displayed signals. The input attenuator and IF gain couple to keep the reference level from changing when we change the input attenuator. Changing the input attenuation will automatically change the IF gain to offset the effect of the change in input attenuation to keep the signal in a constant display position.

Summary

A signal analyzer is the instrument of choice for signal analysis that goes beyond the time domain and into the frequency domain. The sine waves that spectrum analyzers display as frequency spectrum can assist technical experts as they troubleshoot, assess a device or system's quality, or simply help them better understand complex signals. Depending on the types of signals you need to analyze, there are several specialized analyzers designed for the application. The spectrum analyzer fundamentals described in this white paper will help get you started making measurements and analyzing your application's signals.

KEYSIGHT

Keysight enables innovators to push the boundaries of engineering by quickly solving design, emulation, and test challenges to create the best product experiences. Start your innovation journey at www.keysight.com.

This information is subject to change without notice. © Keysight Technologies, 2023, Published in USA, October 31, 2023, 3123-1730.EN