

Noise figure

Many receiver manufacturers specify the performance of their receivers in terms of noise figure, rather than sensitivity. As we shall see, the two can be equated. A spectrum analyzer is a receiver, and we shall examine noise figure on the basis of a sinusoidal input.

Noise figure can be defined as the degradation of signal-to-noise ratio as a signal passes through a device, a spectrum analyzer in our case. We can express noise figure as:

$$F = \frac{S_i/N_i}{S_o/N_o}$$

where F = noise figure as power ratio (also known as noise factor)
 S_i = input signal power
 N_i = true input noise power
 S_o = output signal power
 N_o = output noise power

If we examine this expression, we can simplify it for our spectrum analyzer. First of all, the output signal is the input signal times the gain of the analyzer. Second, the gain of our analyzer is unity because the signal level at the output (indicated on the display) is the same as the level at the input (input connector). So our expression, after substitution, cancellation, and rearrangement, becomes:

$$F = N_o/N_i$$

This expression tells us that all we need to do to determine the noise figure is compare the noise level as read on the display to the true (not the effective) noise level at the input connector. Noise figure is usually expressed in terms of dB, or:

$$NF = 10 \log(F) = 10 \log(N_o) - 10 \log(N_i).$$

We use the true noise level at the input, rather than the effective noise level, because our input signal-to-noise ratio was based on the true noise. As we saw earlier, when the input is terminated in 50 ohms, the kTB noise level at room temperature in a 1 Hz bandwidth is -174 dBm.

We know that the displayed level of noise on the analyzer changes with bandwidth. So all we need to do to determine the noise figure of our spectrum analyzer is to measure the noise power in some bandwidth, calculate the noise power that we would have measured in a 1 Hz bandwidth using $10 \log(BW_2/BW_1)$, and compare that to -174 dBm.

For example, if we measured -110 dBm in a 10 kHz resolution bandwidth, we would get:

$$\begin{aligned} NF &= [\text{measured noise in dBm}] - 10 \log(\text{RBW}/1) - kTB_{B=1 \text{ Hz}} \\ &= -110 \text{ dBm} - 10 \log(10,000/1) - (-174 \text{ dBm}) \\ &= -110 - 40 + 174 \\ &= 24 \text{ dB} \end{aligned}$$

4. This may not always be precisely true for a given analyzer because of the way resolution bandwidth filter sections and gain are distributed in the IF chain.

Noise figure is independent of bandwidth⁴. Had we selected a different resolution bandwidth, our results would have been exactly the same. For example, had we chosen a 1 kHz resolution bandwidth, the measured noise would have been -120 dBm and $10 \log(\text{RBW}/1)$ would have been 30. Combining all terms would have given $-120 - 30 + 174 = 24 \text{ dB}$, the same noise figure as above.

The 24 dB noise figure in our example tells us that a sinusoidal signal must be 24 dB above kTB to be equal to the displayed average noise level on this particular analyzer. Thus we can use noise figure to determine the DANL for a given bandwidth or to compare DANLs of different analyzers on the same bandwidth.⁵

Preamplifiers

One reason for introducing noise figure is that it helps us determine how much benefit we can derive from the use of a preamplifier. A 24 dB noise figure, while good for a spectrum analyzer, is not so good for a dedicated receiver. However, by placing an appropriate preamplifier in front of the spectrum analyzer, we can obtain a system (preamplifier/spectrum analyzer) noise figure that is lower than that of the spectrum analyzer alone. To the extent that we lower the noise figure, we also improve the system sensitivity.

When we introduced noise figure in the previous discussion, we did so on the basis of a sinusoidal input signal. We can examine the benefits of a preamplifier on the same basis. However, a preamplifier also amplifies noise, and this output noise can be higher than the effective input noise of the analyzer. As we shall see in the “Noise as a signal” section later in this chapter, a spectrum analyzer using log power averaging displays a random noise signal 2.5 dB below its actual value. As we explore preamplifiers, we shall account for this 2.5 dB factor where appropriate.

Rather than develop a lot of formulas to see what benefit we get from a preamplifier, let us look at two extreme cases and see when each might apply. First, if the noise power out of the preamplifier (in a bandwidth equal to that of the spectrum analyzer) is at least 15 dB higher than the DANL of the spectrum analyzer, then the noise figure of the system is approximately that of the preamplifier less 2.5 dB. How can we tell if this is the case? Simply connect the preamplifier to the analyzer and note what happens to the noise on the display. If it goes up 15 dB or more, we have fulfilled this requirement.

On the other hand, if the noise power out of the preamplifier (again, in the same bandwidth as that of the spectrum analyzer) is 10 dB or more lower than the displayed average noise level on the analyzer, then the noise figure of the system is that of the spectrum analyzer less the gain of the preamplifier. Again we can test by inspection. Connect the preamplifier to the analyzer; if the displayed noise does not change, we have fulfilled the requirement.

But testing by experiment means that we have the equipment at hand. We do not need to worry about numbers. We simply connect the preamplifier to the analyzer, note the average displayed noise level, and subtract the gain of the preamplifier. Then we have the sensitivity of the system.

What we really want is to know ahead of time what a preamplifier will do for us. We can state the two cases above as follows:

$$\begin{array}{l} \text{If } NF_{pre} + G_{pre} \geq NF_{sa} + 15 \text{ dB,} \\ \text{Then } NF_{sys} = NF_{pre} - 2.5 \text{ dB} \end{array}$$

And

$$\begin{array}{l} \text{If } NF_{pre} + G_{pre} \leq NF_{sa} - 10 \text{ dB,} \\ \text{Then } NF_{sys} = NF_{sa} - G_{pre} \end{array}$$

5. The noise figure computed in this manner cannot be directly compared to that of a receiver because the “measured noise” term in the equation understates the actual noise by 2.5 dB. See the section titled “Noise as a signal” later in this chapter.

Using these expressions, we'll see how a preamplifier affects our sensitivity. Assume that our spectrum analyzer has a noise figure of 24 dB and the preamplifier has a gain of 36 dB and a noise figure of 8 dB. All we need to do is to compare the gain plus noise figure of the preamplifier to the noise figure of the spectrum analyzer. The gain plus noise figure of the preamplifier is 44 dB, more than 15 dB higher than the noise figure of the spectrum analyzer, so the noise figure of the preamplifier/spectrum-analyzer combination is that of the preamplifier less 2.5 dB, or 5.5 dB. In a 10 kHz resolution bandwidth, our preamplifier/analyzer system has a sensitivity of:

$$\begin{aligned} kTB_{B=1} + 10 \log(\text{RBW}/1) + \text{NF}_{\text{sys}} &= -174 + 40 + 5.5 \\ &= -128.5 \text{ dBm} \end{aligned}$$

This is an improvement of 18.5 dB over the -110 dBm noise floor without the preamplifier.

There might, however, be a drawback to using this preamplifier, depending upon our ultimate measurement objective. If we want the best sensitivity but no loss of measurement range, then this preamplifier is not the right choice. Figure 5-4 illustrates this point. A spectrum analyzer with a 24 dB noise figure will have an average displayed noise level of -110 dBm in a 10 kHz resolution bandwidth. If the 1 dB compression point⁶ for that analyzer is 0 dBm, the measurement range is 110 dB. When we connect the preamplifier, we must reduce the maximum input to the system by the gain of the preamplifier to -36 dBm. However, when we connect the preamplifier, the displayed average noise level will rise by about 17.5 dB because the noise power out of the preamplifier is that much higher than the analyzer's own noise floor, even after accounting for the 2.5 dB factor. It is from this higher noise level that we now subtract the gain of the preamplifier. With the preamplifier in place, our measurement range is 92.5 dB, 17.5 dB less than without the preamplifier. The loss in measurement range equals the change in the displayed noise when the preamplifier is connected.

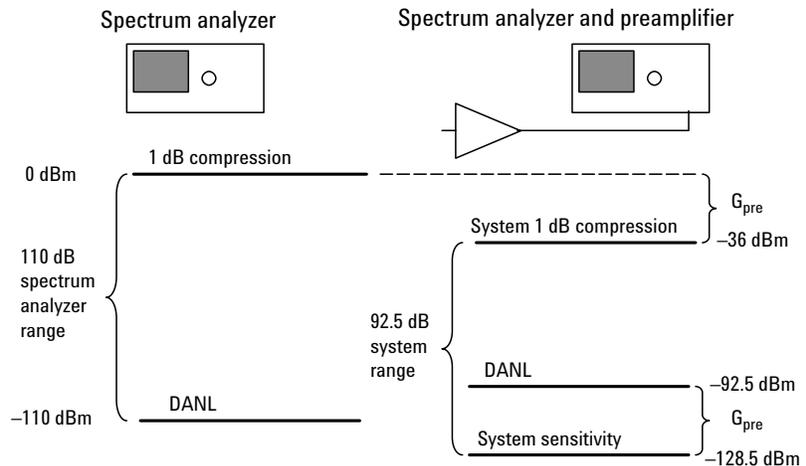


Figure 5-4. If displayed noise goes up when a preamplifier is connected, measurement range is diminished by the amount the noise changes

6. See the section titled "Mixer compression" in Chapter 6.

Finding a preamplifier that will give us better sensitivity without costing us measurement range dictates that we must meet the second of the above criteria; that is, the sum of its gain and noise figure must be at least 10 dB less than the noise figure of the spectrum analyzer. In this case the displayed noise floor will not change noticeably when we connect the preamplifier, so although we shift the whole measurement range down by the gain of the preamplifier, we end up with the same overall range that we started with.

To choose the correct preamplifier, we must look at our measurement needs. If we want absolutely the best sensitivity and are not concerned about measurement range, we would choose a high-gain, low-noise-figure preamplifier so that our system would take on the noise figure of the preamplifier, less 2.5 dB. If we want better sensitivity but cannot afford to give up any measurement range, we must choose a lower-gain preamplifier.

Interestingly enough, we can use the input attenuator of the spectrum analyzer to effectively degrade the noise figure (or reduce the gain of the preamplifier, if you prefer). For example, if we need slightly better sensitivity but cannot afford to give up any measurement range, we can use the above preamplifier with 30 dB of RF input attenuation on the spectrum analyzer. This attenuation increases the noise figure of the analyzer from 24 to 54 dB. Now the gain plus noise figure of the preamplifier (36 + 8) is 10 dB less than the noise figure of the analyzer, and we have met the conditions of the second criterion above. The noise figure of the system is now:

$$\begin{aligned} NF_{\text{sys}} &= NF_{\text{SA}} - G_{\text{PRE}} \\ &= 54 \text{ dB} - 36 \text{ dB} \\ &= 18 \text{ dB} \end{aligned}$$

This represents a 6 dB improvement over the noise figure of the analyzer alone with 0 dB of input attenuation. So we have improved sensitivity by 6 dB and given up virtually no measurement range.

Of course, there are preamplifiers that fall in between the extremes. Figure 5-5 enables us to determine system noise figure from a knowledge of the noise figures of the spectrum analyzer and preamplifier and the gain of the amplifier. We enter the graph of Figure 5-5 by determining $NF_{\text{PRE}} + G_{\text{PRE}} - NF_{\text{SA}}$. If the value is less than zero, we find the corresponding point on the dashed curve and read system noise figure as the left ordinate in terms of dB above $NF_{\text{SA}} - G_{\text{PRE}}$. If $NF_{\text{PRE}} + G_{\text{PRE}} - NF_{\text{SA}}$ is a positive value, we find the corresponding point on the solid curve and read system noise figure as the right ordinate in terms of dB above NF_{PRE} .

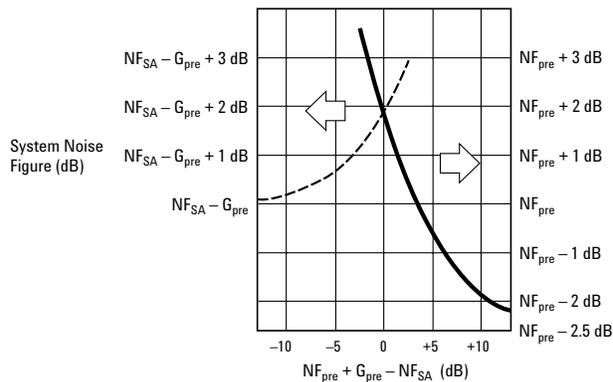


Figure 5-5. System noise figure for sinusoidal signals

Let's first test the two previous extreme cases.

As $NF_{PRE} + G_{PRE} - NF_{SA}$ becomes less than -10 dB, we find that system noise figure asymptotically approaches $NF_{SA} - G_{PRE}$. As the value becomes greater than $+15$ dB, system noise figure asymptotically approaches NF_{PRE} less 2.5 dB. Next, let's try two numerical examples. Above, we determined that the noise figure of our analyzer is 24 dB. What would the system noise figure be if we add an Agilent 8447D, a preamplifier with a noise figure of about 8 dB and a gain of 26 dB? First, $NF_{PRE} + G_{PRE} - NF_{SA}$ is $+10$ dB. From the graph of Figure 5-5 we find a system noise figure of about $NF_{PRE} - 1.8$ dB, or about $8 - 1.8 = 6.2$ dB. The graph accounts for the 2.5 dB factor. On the other hand, if the gain of the preamplifier is just 10 dB, then $NF_{PRE} + G_{PRE} - NF_{SA}$ is -6 dB. This time the graph indicates a system noise figure of $NF_{SA} - G_{PRE} + 0.6$ dB, or $24 - 10 + 0.6 = 14.6$ dB⁷. (We did not introduce the 2.5 dB factor previously when we determined the noise figure of the analyzer alone because we read the measured noise directly from the display. The displayed noise included the 2.5 dB factor.)

Many modern spectrum analyzers have optional built-in preamplifiers available. Compared to external preamplifiers, built-in preamplifiers simplify measurement setups and eliminate the need for additional cabling. Measuring signal amplitude is much more convenient with a built-in preamplifier, because the preamplifier/spectrum analyzer combination is calibrated as a system, and amplitude values displayed on screen are already corrected for proper readout. With an external preamplifier, you must correct the spectrum analyzer reading with a reference level offset equal to the preamp gain. Most modern spectrum analyzers allow you to enter the gain value of the external preamplifier from the front panel. The analyzer then applies this gain offset to the displayed reference level value, so that you can directly view corrected measurements on the display.

Noise as a signal

So far, we have focused on the noise generated within the measurement system (analyzer or analyzer/preamplifier). We described how the measurement system's displayed average noise level limits the overall sensitivity. However, random noise is sometimes the signal that we want to measure. Because of the nature of noise, the superheterodyne spectrum analyzer indicates a value that is lower than the actual value of the noise. Let's see why this is so and how we can correct for it.

By random noise, we mean a signal whose instantaneous amplitude has a Gaussian distribution versus time, as shown in Figure 5-6. For example, thermal or Johnson noise has this characteristic. Such a signal has no discrete spectral components, so we cannot select some particular component and measure it to get an indication of signal strength. In fact, we must define what we mean by signal strength. If we sample the signal at an arbitrary instant, we could theoretically get any amplitude value. We need some measure that expresses the noise level averaged over time. Power, which is of course proportionate to rms voltage, satisfies that requirement.

7. For more details on noise figure, see Agilent *Application Note 57-1, Fundamentals of RF and Microwave Noise Figure Measurements*, literature number 5952-8255E.

We have already seen that both video filtering and video averaging reduce the peak-to-peak fluctuations of a signal and can give us a steady value. We must equate this value to either power or rms voltage. The rms value of a Gaussian distribution equals its standard deviation, σ .

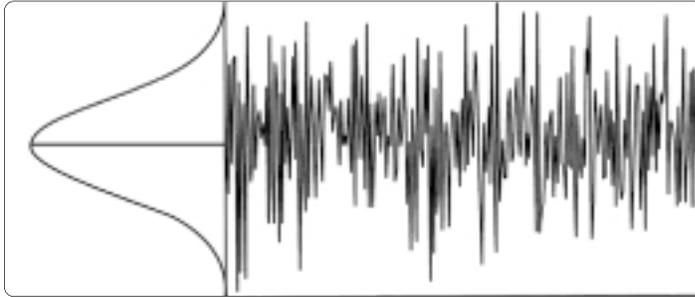


Figure 5-6. Random noise has a Gaussian amplitude distribution

Let's start with our analyzer in the linear display mode. The Gaussian noise at the input is band limited as it passes through the IF chain, and its envelope takes on a Rayleigh distribution (Figure 5-7). The noise that we see on our analyzer display, the output of the envelope detector, is the Rayleigh distributed envelope of the input noise signal. To get a steady value, the mean value, we use video filtering or averaging. The mean value of a Rayleigh distribution is 1.253σ .

But our analyzer is a peak-responding voltmeter calibrated to indicate the rms value of a sine wave. To convert from peak to rms, our analyzer scales its readout by 0.707 (-3 dB). The mean value of the Rayleigh-distributed noise is scaled by the same factor, giving us a reading that is 0.886σ (1.05 dB below σ). To equate the mean value displayed by the analyzer to the rms voltage of the input noise signal, then, we must account for the error in the displayed value. Note, however, that the error is not an ambiguity; it is a constant error that we can correct for by adding 1.05 dB to the displayed value.

In most spectrum analyzers, the display scale (log or linear in voltage) controls the scale on which the noise distribution is averaged with either the VBW filter or with trace averaging. Normally, we use our analyzer in the log display mode, and this mode adds to the error in our noise measurement. The gain of a log amplifier is a function of signal amplitude, so the higher noise values are not amplified as much as the lower values. As a result, the output of the envelope detector is a skewed Rayleigh distribution, and the mean value that we get from video filtering or averaging is another 1.45 dB lower. In the log mode, then, the mean or average noise is displayed 2.5 dB too low. Again, this error is not an ambiguity, and we can correct for it⁸.

8. In the ESA and PSA Series, the averaging can be set to video, voltage, or power (rms), independent of display scale. When using power averaging, no correction is needed, since the average rms level is determined by the square of the magnitude of the signal, not by the log or envelope of the voltage.

This is the 2.5 dB factor that we accounted for in the previous preamplifier discussion, whenever the noise power out of the preamplifier was approximately equal to or greater than the analyzer's own noise.

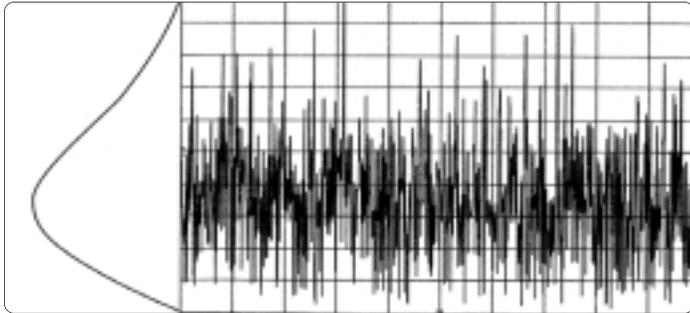


Figure 5-7. The envelope of band-limited Gaussian noise has a Rayleigh distribution

Another factor that affects noise measurements is the bandwidth in which the measurement is made. We have seen how changing resolution bandwidth affects the displayed level of the analyzer's internally generated noise. Bandwidth affects external noise signals in the same way. To compare measurements made on different analyzers, we must know the bandwidths used in each case.

Not only does the 3 dB (or 6 dB) bandwidth of the analyzer affect the measured noise level, the shape of the resolution filter also plays a role. To make comparisons possible, we define a standard noise-power bandwidth: the width of a rectangular filter that passes the same noise power as our analyzer's filter. For the near-Gaussian filters in Agilent analyzers, the equivalent noise-power bandwidth is about 1.05 to 1.13 times the 3 dB bandwidth, depending on bandwidth selectivity. For example, a 10 kHz resolution bandwidth filter has a noise-power bandwidth in the range of 10.5 to 11.3 kHz.

If we use $10 \log(BW_2/BW_1)$ to adjust the displayed noise level to what we would have measured in a noise-power bandwidth of the same numeric value as our 3 dB bandwidth, we find that the adjustment varies from:

$$10 \log(10,000/10,500) = -0.21 \text{ dB}$$

to

$$10 \log(10,000/11,300) = -0.53 \text{ dB}$$

In other words, if we subtract something between 0.21 and 0.53 dB from the indicated noise level, we shall have the noise level in a noise-power bandwidth that is convenient for computations. For the following examples below, we will use 0.5 dB as a reasonable compromise for the bandwidth correction⁹.

9. ESA Series analyzers calibrate each RBW during the IF alignment routine to determine the noise power bandwidth. The PSA Series analyzers specify noise power bandwidth accuracy to within 1% (± 0.044 dB).

Let's consider the various correction factors to calculate the total correction for each averaging mode:

Linear (voltage) averaging:

Rayleigh distribution (linear mode):	1.05 dB
3 dB/noise power bandwidths:	<u>-50 dB</u>
Total correction:	0.55 dB

Log averaging:

Logged Rayleigh distribution:	2.50 dB
3 dB/noise power bandwidths:	<u>-50 dB</u>
Total correction:	2.00 dB

Power (rms voltage) averaging:

Power distribution:	0.00 dB
3 dB/noise power bandwidths:	<u>-50 dB</u>
Total correction:	-50 dB

Many of today's microprocessor-controlled analyzers allow us to activate a noise marker. When we do so, the microprocessor switches the analyzer into the power (rms) averaging mode, computes the mean value of a number of display points about the marker¹⁰, normalizes and corrects the value to a 1 Hz noise-power bandwidth, and displays the normalized value.

The analyzer does the hard part. It is easy to convert the noise-marker value to other bandwidths. For example, if we want to know the total noise in a 4 MHz communication channel, we add $10 \log(4,000,000/1)$, or 66 dB to the noise-marker value¹¹.

Preamplifier for noise measurements

Since noise signals are typically low-level signals, we often need a preamplifier to have sufficient sensitivity to measure them. However, we must recalculate sensitivity of our analyzer first. We previously defined sensitivity as the level of a sinusoidal signal that is equal to the displayed average noise floor. Since the analyzer is calibrated to show the proper amplitude of a sinusoid, no correction for the signal was needed. But noise is displayed 2.5 dB too low, so an input noise signal must be 2.5 dB above the analyzer's displayed noise floor to be at the same level by the time it reaches the display. The input and internal noise signals add to raise the displayed noise by 3 dB, a factor of two in power. So we can define the noise figure of our analyzer for a noise signal as:

$$NF_{SA(N)} = (\text{noise floor})_{dBm/RBW} - 10 \log(RBW/1) - kTB_{B=1} + 2.5 \text{ dB}$$

If we use the same noise floor that we used previously, -110 dBm in a 10 kHz resolution bandwidth, we get:

$$NF_{SA(N)} = -110 \text{ dBm} - 10 \log(10,000/1) - (-174 \text{ dBm}) + 2.5 \text{ dB} = 26.5 \text{ dB}$$

As was the case for a sinusoidal signal, $NF_{SA(N)}$ is independent of resolution bandwidth and tells us how far above kTB a noise signal must be to be equal to the noise floor of our analyzer.

10. For example, the ESA and PSA Series compute the mean over half a division, regardless of the number of display points.

11. Most modern spectrum analyzers make this calculation even easier with the Channel Power function. The user enters the integration bandwidth of the channel and centers the signal on the analyzer display. The Channel Power function then calculates the total signal power in the channel.

When we add a preamplifier to our analyzer, the system noise figure and sensitivity improve. However, we have accounted for the 2.5 dB factor in our definition of $NF_{SA(N)}$, so the graph of system noise figure becomes that of Figure 5-8. We determine system noise figure for noise the same way that we did previously for a sinusoidal signal.

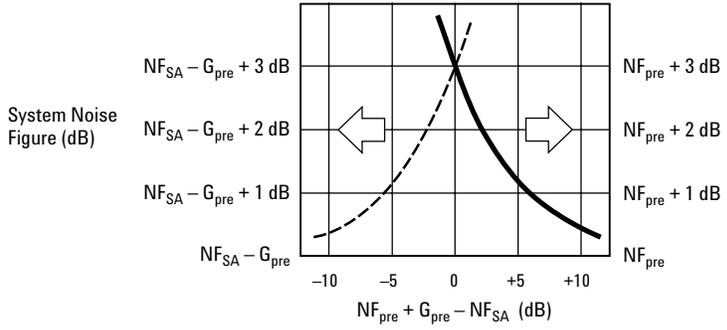


Figure 5-8. System noise figure for noise signals

Chapter 6

Dynamic Range

Definition

Dynamic range is generally thought of as the ability of an analyzer to measure harmonically related signals and the interaction of two or more signals; for example, to measure second- or third-harmonic distortion or third-order intermodulation. In dealing with such measurements, remember that the input mixer of a spectrum analyzer is a non-linear device, so it always generates distortion of its own. The mixer is non-linear for a reason. It must be nonlinear to translate an input signal to the desired IF. But the unwanted distortion products generated in the mixer fall at the same frequencies as the distortion products we wish to measure on the input signal.

So we might define dynamic range in this way: it is the ratio, expressed in dB, of the largest to the smallest signals simultaneously present at the input of the spectrum analyzer that allows measurement of the smaller signal to a given degree of uncertainty.

Notice that accuracy of the measurement is part of the definition. We shall see how both internally generated noise and distortion affect accuracy in the following examples.

Dynamic range versus internal distortion

To determine dynamic range versus distortion, we must first determine just how our input mixer behaves. Most analyzers, particularly those utilizing harmonic mixing to extend their tuning range¹, use diode mixers. (Other types of mixers would behave similarly.) The current through an ideal diode can be expressed as:

$$i = I_s(e^{qv/kT} - 1)$$

where I_s = the diode's saturation current
 q = electron charge (1.60×10^{-19} C)
 v = instantaneous voltage
 k = Boltzmann's constant (1.38×10^{-23} joule/°K)
 T = temperature in degrees Kelvin

We can expand this expression into a power series:

$$i = I_s(k_1 v + k_2 v^2 + k_3 v^3 + \dots)$$

where $k_1 = q/kT$
 $k_2 = k_1^2/2!$
 $k_3 = k_1^3/3!$, etc.

Let's now apply two signals to the mixer. One will be the input signal that we wish to analyze; the other, the local oscillator signal necessary to create the IF:

$$v = V_{LO} \sin(\omega_{LO}t) + V_1 \sin(\omega_1t)$$

If we go through the mathematics, we arrive at the desired mixing product that, with the correct LO frequency, equals the IF:

$$k_2 V_{LO} V_1 \cos[(\omega_{LO} - \omega_1)t]$$

A $k_2 V_{LO} V_1 \cos[(\omega_{LO} + \omega_1)t]$ term is also generated, but in our discussion of the tuning equation, we found that we want the LO to be above the IF, so $(\omega_{LO} + \omega_1)$ is also always above the IF.

1. See Chapter 7, "Extending the Frequency Range."

With a constant LO level, the mixer output is linearly related to the input signal level. For all practical purposes, this is true as long as the input signal is more than 15 to 20 dB below the level of the LO. There are also terms involving harmonics of the input signal:

$$\begin{aligned} & (3k_3/4)V_{LO}V_1^2 \sin(\omega_{LO} - 2\omega_1)t, \\ & (k_4/8)V_{LO}V_1^3 \sin(\omega_{LO} - 3\omega_1)t, \text{ etc.} \end{aligned}$$

These terms tell us that dynamic range due to internal distortion is a function of the input signal level at the input mixer. Let's see how this works, using as our definition of dynamic range, the difference in dB between the fundamental tone and the internally generated distortion.

The argument of the sine in the first term includes $2\omega_1$, so it represents the second harmonic of the input signal. The level of this second harmonic is a function of the square of the voltage of the fundamental, V_1^2 . This fact tells us that for every dB that we drop the level of the fundamental at the input mixer, the internally generated second harmonic drops by 2 dB. See Figure 6-1. The second term includes $3\omega_1$, the third harmonic, and the cube of the input-signal voltage, V_1^3 . So a 1 dB change in the fundamental at the input mixer changes the internally generated third harmonic by 3 dB.

Distortion is often described by its order. The order can be determined by noting the coefficient associated with the signal frequency or the exponent associated with the signal amplitude. Thus second-harmonic distortion is second order and third harmonic distortion is third order. The order also indicates the change in internally generated distortion relative to the change in the fundamental tone that created it.

Now let us add a second input signal:

$$v = V_{LO} \sin(\omega_{LO} t) + V_1 \sin(\omega_1 t) + V_2 \sin(\omega_2 t)$$

This time when we go through the math to find internally generated distortion, in addition to harmonic distortion, we get:

$$\begin{aligned} & (k_4/8)V_{LO}V_1^2V_2 \cos[\omega_{LO} - (2\omega_1 - \omega_2)]t, \\ & (k_4/8)V_{LO}V_1V_2^2 \cos[\omega_{LO} - (2\omega_2 - \omega_1)]t, \text{ etc.} \end{aligned}$$

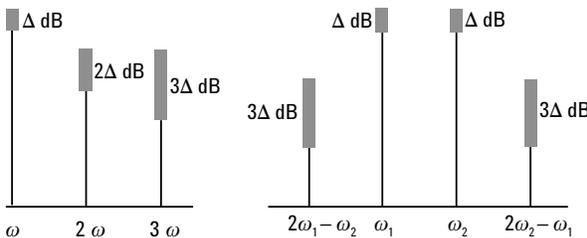


Figure 6-1. Changing the level of fundamental tones at the mixer

These represent intermodulation distortion, the interaction of the two input signals with each other. The lower distortion product, $2\omega_1 - \omega_2$, falls below ω_1 by a frequency equal to the difference between the two fundamental tones, $\omega_2 - \omega_1$. The higher distortion product, $2\omega_2 - \omega_1$, falls above ω_2 by the same frequency. See Figure 6-1.

Once again, dynamic range is a function of the level at the input mixer. The internally generated distortion changes as the product of V_1^2 and V_2 in the first case, of V_1 and V_2^2 in the second. If V_1 and V_2 have the same amplitude, the usual case when testing for distortion, we can treat their products as cubed terms (V_1^3 or V_2^3). Thus, for every dB that we simultaneously change the level of the two input signals, there is a 3 dB change in the distortion components, as shown in Figure 6-1.

This is the same degree of change that we see for third harmonic distortion in Figure 6-1. And in fact, this too, is third-order distortion. In this case, we can determine the degree of distortion by summing the coefficients of ω_1 and ω_2 (e.g., $2\omega_1 - 1\omega_2$ yields $2 + 1 = 3$) or the exponents of V_1 and V_2 .

All this says that dynamic range depends upon the signal level at the mixer. How do we know what level we need at the mixer for a particular measurement? Most analyzer data sheets include graphs to tell us how dynamic range varies. However, if no graph is provided, we can draw our own².

We do need a starting point, and this we must get from the data sheet. We shall look at second-order distortion first. Let's assume the data sheet says that second-harmonic distortion is 75 dB down for a signal -40 dBm at the mixer. Because distortion is a relative measurement, and, at least for the moment, we are calling our dynamic range the difference in dB between fundamental tone or tones and the internally generated distortion, we have our starting point. Internally generated second-order distortion is 75 dB down, so we can measure distortion down 75 dB. We plot that point on a graph whose axes are labeled distortion (dBc) versus level at the mixer (level at the input connector minus the input-attenuator setting). See Figure 6-2. What happens if the level at the mixer drops to -50 dBm? As noted in Figure 6-1, for every dB change in the level of the fundamental at the mixer there is a 2 dB change in the internally generated second harmonic. But for measurement purposes, we are only interested in the relative change, that is, in what happened to our measurement range. In this case, for every dB that the fundamental changes at the mixer, our measurement range also changes by 1 dB. In our second-harmonic example, then, when the level at the mixer changes from -40 to -50 dBm, the internal distortion, and thus our measurement range, changes from -75 to -85 dBc. In fact, these points fall on a line with a slope of 1 that describes the dynamic range for any input level at the mixer.

2. For more information on how to construct a dynamic range chart, see the Agilent *PSA Performance Spectrum Analyzer Series Product Note, Optimizing Dynamic Range for Distortion Measurements*, literature number 5980-3079EN.

We can construct a similar line for third-order distortion. For example, a data sheet might say third-order distortion is -85 dBc for a level of -30 dBm at this mixer. Again, this is our starting point, and we would plot the point shown in Figure 6-2. If we now drop the level at the mixer to -40 dBm, what happens? Referring again to Figure 6-1, we see that both third-harmonic distortion and third-order intermodulation distortion fall by 3 dB for every dB that the fundamental tone or tones fall. Again it is the difference that is important. If the level at the mixer changes from -30 to -40 dBm, the difference between fundamental tone or tones and internally generated distortion changes by 20 dB. So the internal distortion is -105 dBc. These two points fall on a line having a slope of 2 , giving us the third-order performance for any level at the mixer.

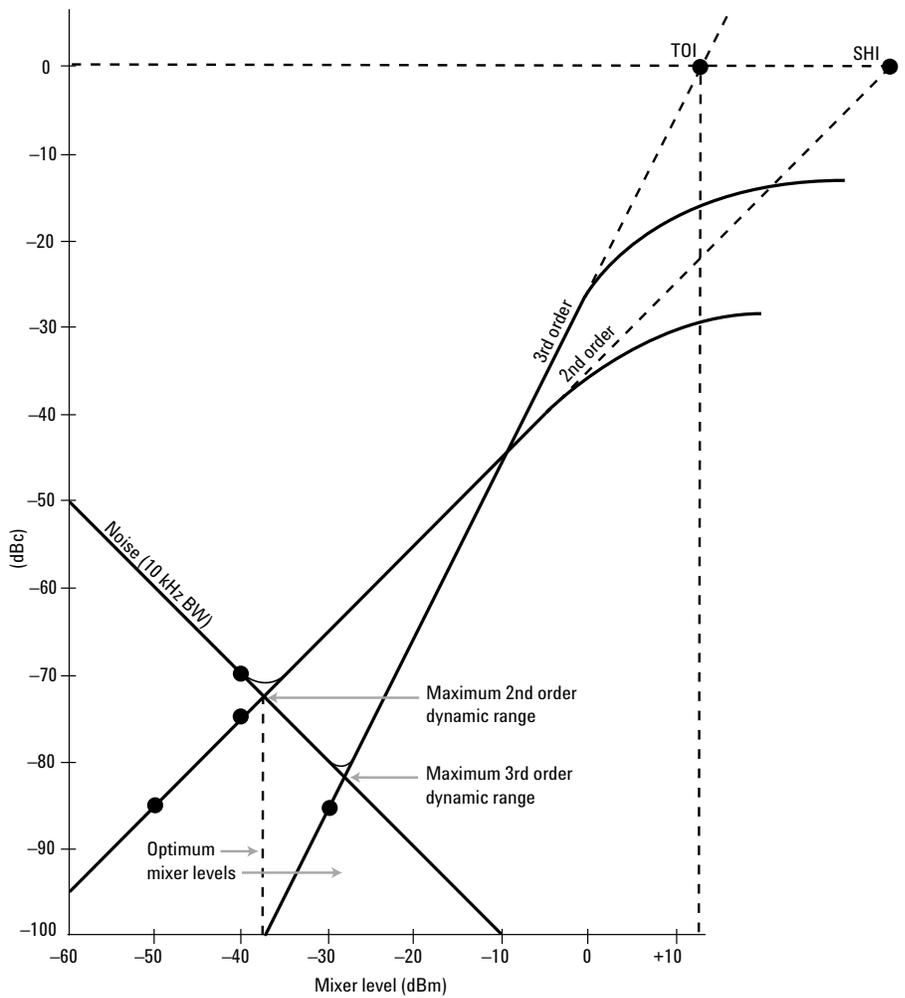


Figure 6-2. Dynamic range versus distortion and noise

Sometimes third-order performance is given as TOI (third-order intercept). This is the mixer level at which the internally generated third-order distortion would be equal to the fundamental(s), or 0 dBc. This situation cannot be realized in practice because the mixer would be well into saturation. However, from a mathematical standpoint, TOI is a perfectly good data point because we know the slope of the line. So even with TOI as a starting point, we can still determine the degree of internally generated distortion at a given mixer level.

We can calculate TOI from data sheet information. Because third-order dynamic range changes 2 dB for every dB change in the level of the fundamental tone(s) at the mixer, we get TOI by subtracting half of the specified dynamic range in dBc from the level of the fundamental(s):

$$\text{TOI} = A_{\text{fund}} - d/2$$

where A_{fund} = level of the fundamental in dBm
 d = difference in dBc between fundamental and distortion

Using the values from the previous discussion:

$$\text{TOI} = -30 \text{ dBm} - (-85 \text{ dBc})/2 = +12.5 \text{ dBm}$$

Attenuator test

Understanding the distortion graph is important, but we can use a simple test to determine whether displayed distortion components are true input signals or internally generated signals. Change the input attenuator. If the displayed value of the distortion components remains the same, the components are part of the input signal. If the displayed value changes, the distortion components are generated internally or are the sum of external and internally generated signals. We continue changing the attenuator until the displayed distortion does not change and then complete the measurement.

Noise

There is another constraint on dynamic range, and that is the noise floor of our spectrum analyzer. Going back to our definition of dynamic range as the ratio of the largest to the smallest signal that we can measure, the average noise of our spectrum analyzer puts the limit on the smaller signal. So dynamic range versus noise becomes signal-to-noise ratio in which the signal is the fundamental whose distortion we wish to measure.

We can easily plot noise on our dynamic range chart. For example, suppose that the data sheet for our spectrum analyzer specifies a displayed average noise level of -110 dBm in a 10 kHz resolution bandwidth. If our signal fundamental has a level of -40 dBm at the mixer, it is 70 dB above the average noise, so we have a 70 dB signal-to-noise ratio. For every dB that we reduce the signal level at the mixer, we lose 1 dB of signal-to-noise ratio. Our noise curve is a straight line having a slope of -1, as shown in Figure 6-2.

If we ignore measurement accuracy considerations for a moment, the best dynamic range will occur at the intersection of the appropriate distortion curve and the noise curve. Figure 6-2 tells us that our maximum dynamic range for second-order distortion is 72.5 dB; for third-order distortion, 81.7 dB. In practice, the intersection of the noise and distortion graphs is not a sharply defined point, because noise adds to the CW-like distortion products, reducing dynamic range by 2 dB when using the log power scale with log scale averaging.

Figure 6-2 shows the dynamic range for one resolution bandwidth. We certainly can improve dynamic range by narrowing the resolution bandwidth, but there is not a one-to-one correspondence between the lowered noise floor and the improvement in dynamic range. For second-order distortion, the improvement is one half the change in the noise floor; for third-order distortion, two-thirds the change in the noise floor. See Figure 6-3.

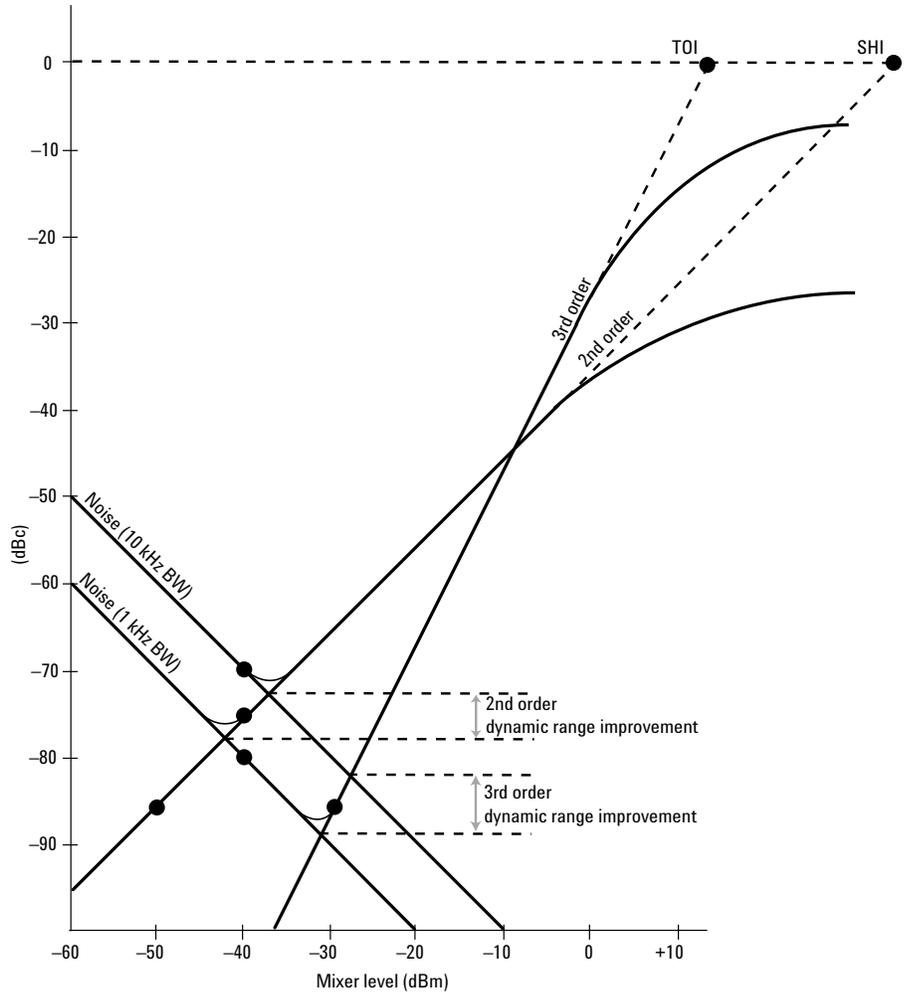


Figure 6-3. Reducing resolution bandwidth improves dynamic range

The final factor in dynamic range is the phase noise on our spectrum analyzer LO, and this affects only third-order distortion measurements. For example, suppose we are making a two-tone, third-order distortion measurement on an amplifier, and our test tones are separated by 10 kHz. The third-order distortion components will also be separated from the test tones by 10 kHz. For this measurement we might find ourselves using a 1 kHz resolution bandwidth. Referring to Figure 6-3 and allowing for a 10 dB decrease in the noise curve, we would find a maximum dynamic range of about 88 dB. Suppose however, that our phase noise at a 10 kHz offset is only -80 dBc. Then 80 dB becomes the ultimate limit of dynamic range for this measurement, as shown in Figure 6-4.

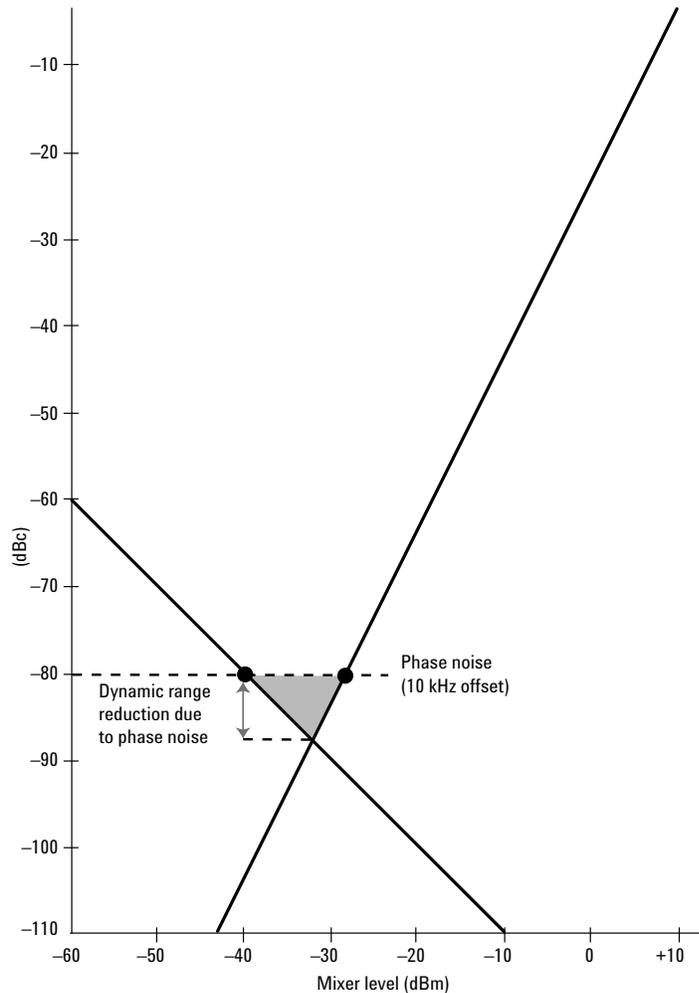


Figure 6-4. Phase noise can limit third-order intermodulation tests

In summary, the dynamic range of a spectrum analyzer is limited by three factors: the distortion performance of the input mixer, the broadband noise floor (sensitivity) of the system, and the phase noise of the local oscillator.

Dynamic range versus measurement uncertainty

In our previous discussion of amplitude accuracy, we included only those items listed in Table 4-1, plus mismatch. We did not cover the possibility of an internally generated distortion product (a sinusoid) being at the same frequency as an external signal that we wished to measure. However, internally generated distortion components fall at exactly the same frequencies as the distortion components we wish to measure on external signals. The problem is that we have no way of knowing the phase relationship between the external and internal signals. So we can only determine a potential range of uncertainty:

$$\text{Uncertainty (in dB)} = 20 \log(1 \pm 10^{d/20})$$

where d = difference in dB between the larger and smaller sinusoid (a negative number)

See Figure 6-5. For example, if we set up conditions such that the internally generated distortion is equal in amplitude to the distortion on the incoming signal, the error in the measurement could range from +6 dB (the two signals exactly in phase) to -infinity (the two signals exactly out of phase and so canceling). Such uncertainty is unacceptable in most cases. If we put a limit of ± 1 dB on the measurement uncertainty, Figure 6-5 shows us that the internally generated distortion product must be about 18 dB below the distortion product that we wish to measure. To draw dynamic range curves for second- and third-order measurements with no more than 1 dB of measurement error, we must then offset the curves of Figure 6-2 by 18 dB as shown in Figure 6-6.

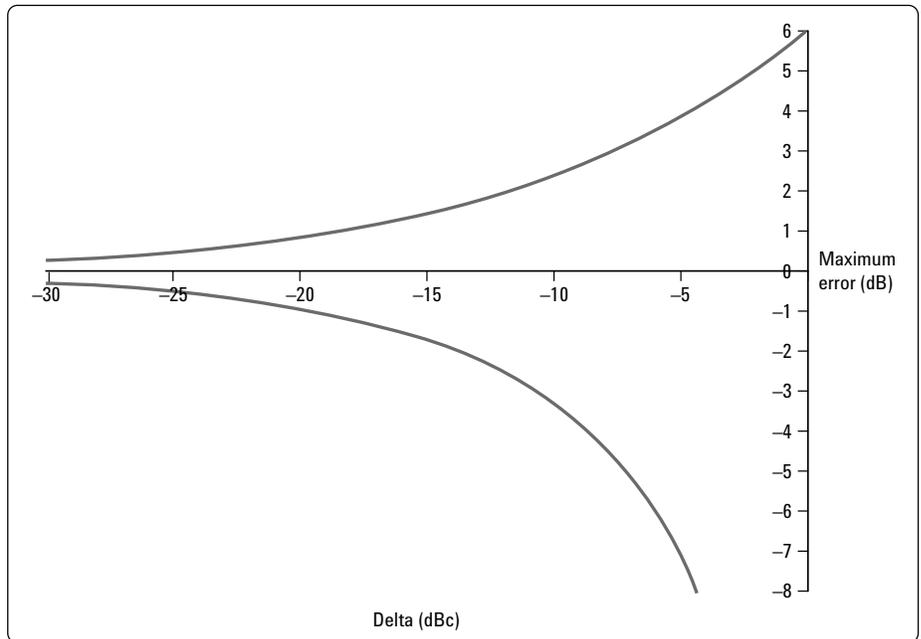


Figure 6-5. Uncertainty versus difference in amplitude between two sinusoids at the same frequency

Next, let's look at uncertainty due to low signal-to-noise ratio. The distortion components we wish to measure are, we hope, low-level signals, and often they are at or very close to the noise level of our spectrum analyzer. In such cases, we often use the video filter to make these low-level signals more discernable. Figure 6-7 shows the error in displayed signal level as a function of displayed signal-to-noise for a typical spectrum analyzer. Note that the error is only in one direction, so we could correct for it. However, we usually do not. So for our dynamic range measurement, let's accept a 0.3 dB error due to noise and offset the noise curve in our dynamic range chart by 5 dB as shown in Figure 6-6. Where the distortion and noise curves intersect, the maximum error possible would be less than 1.3 dB.

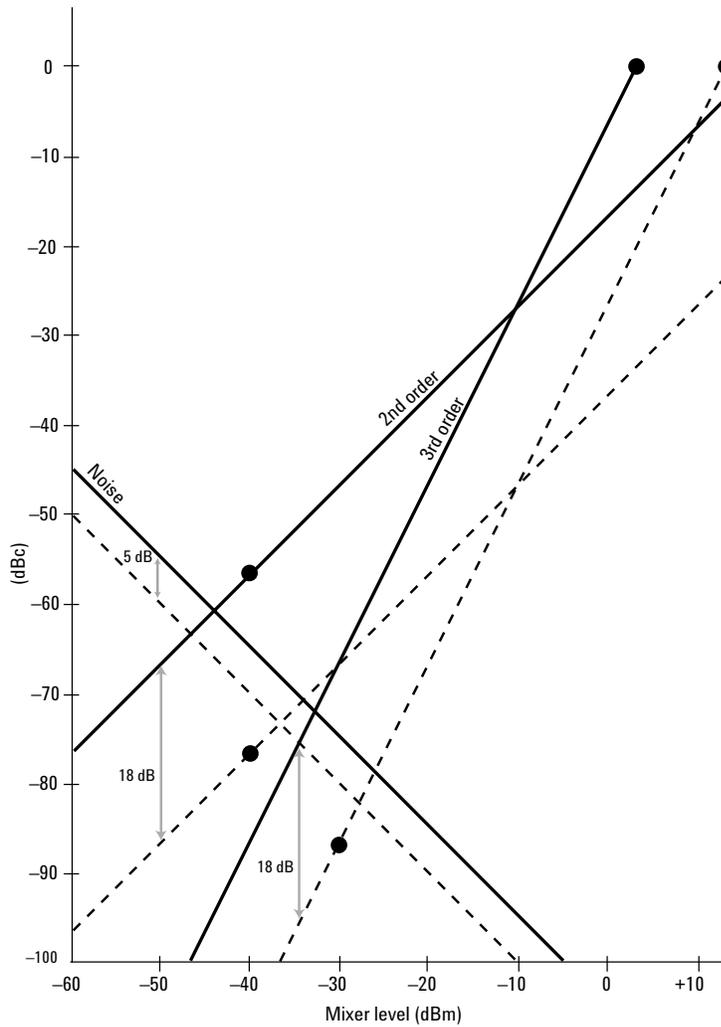


Figure 6-6. Dynamic range for 1.3 dB maximum error

Let's see what happened to our dynamic range as a result of our concern with measurement error. As Figure 6-6 shows, second-order-distortion dynamic range changes from 72.5 to 61 dB, a change of 11.5 dB. This is one half the total offsets for the two curves (18 dB for distortion; 5 dB for noise). Third-order distortion changes from 81.7 dB to about 72.7 dB for a change of about 9 dB. In this case, the change is one third of the 18 dB offset for the distortion curve plus two thirds of the 5 dB offset for the noise curve.

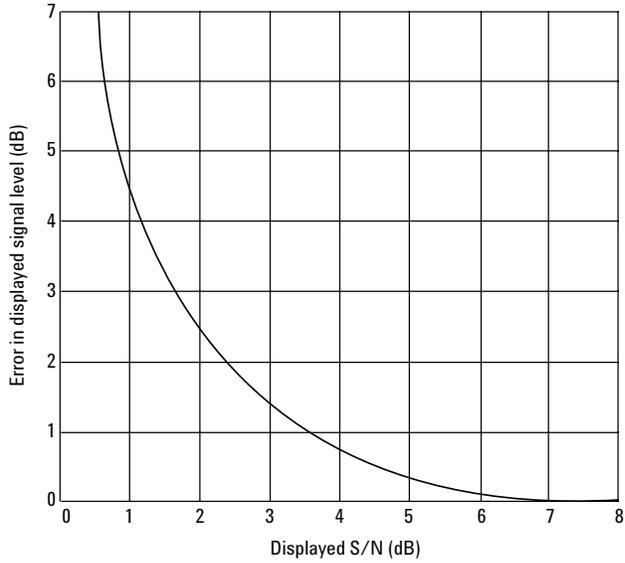


Figure 6-7. Error in displayed signal amplitude due to noise

Gain compression

In our discussion of dynamic range, we did not concern ourselves with how accurately the larger tone is displayed, even on a relative basis. As we raise the level of a sinusoidal input signal, eventually the level at the input mixer becomes so high that the desired output mixing product no longer changes linearly with respect to the input signal. The mixer is in saturation, and the displayed signal amplitude is too low. Saturation is gradual rather than sudden. To help us stay away from the saturation condition, the 1-dB compression point is normally specified. Typically, this gain compression occurs at a mixer level in the range of -5 to $+5$ dBm. Thus we can determine what input attenuator setting to use for accurate measurement of high-level signals³. Spectrum analyzers with a digital IF will display an “IF Overload” message when the ADC is over-ranged.

Actually, there are three different methods of evaluating compression. A traditional method, called CW compression, measures the change in gain of a device (amplifier or mixer or system) as the input signal power is swept upward. This method is the one just described. Note that the CW compression point is considerably higher than the levels for the fundamentals indicated previously for even moderate dynamic range. So we were correct in not concerning ourselves with the possibility of compression of the larger signal(s).

A second method, called two-tone compression, measures the change in system gain for a small signal while the power of a larger signal is swept upward. Two-tone compression applies to the measurement of multiple CW signals, such as sidebands and independent signals. The threshold of compression of this method is usually a few dB lower than that of the CW method. This is the method used by Agilent Technologies to specify spectrum analyzer gain compression.

A final method, called pulse compression, measures the change in system gain to a narrow (broadband) RF pulse while the power of the pulse is swept upward. When measuring pulses, we often use a resolution bandwidth much narrower than the bandwidth of the pulse, so our analyzer displays the signal level well below the peak pulse power. As a result, we could be unaware of the fact that the total signal power is above the mixer compression threshold. A high threshold improves signal-to-noise ratio for high-power, ultra-narrow or widely chirped pulses. The threshold is about 12 dB higher than for two-tone compression in the Agilent 8560EC Series spectrum analyzers. Nevertheless, because different compression mechanisms affect CW, two-tone, and pulse compression differently, any of the compression thresholds can be lower than any other.

Display range and measurement range

There are two additional ranges that are often confused with dynamic range: display range and measurement range. Display range, often called display dynamic range, refers to the calibrated amplitude range of the spectrum analyzer display. For example, a display with ten divisions would seem to have a 100 dB display range when we select 10 dB per division. This is certainly true for modern analyzers with digital IF circuitry, such as the Agilent PSA Series. It is also true for the Agilent ESA-E Series when using the narrow (10 to 300 Hz) digital resolution bandwidths. However, spectrum analyzers with analog IF sections typically are only calibrated for the first 85 or 90 dB below the reference level. In this case, the bottom line of the graticule represents signal amplitudes of zero, so the bottom portion of the display covers the range from -85 or -90 dB to infinity, relative to the reference level.

3. Many analyzers internally control the combined settings of the input attenuator and IF gain so that a CW signal as high as the compression level at the input mixer creates a deflection above the top line of the graticule. Thus we cannot make incorrect measurements on CW signals inadvertently.

The range of the log amplifier can be another limitation for spectrum analyzers with analog IF circuitry. For example, ESA-L Series spectrum analyzers use an 85 dB log amplifier. Thus, only measurements that are within 85 dB below the reference level are calibrated.

The question is, can the full display range be used? From the previous discussion of dynamic range, we know that the answer is generally yes. In fact, dynamic range often exceeds display range or log amplifier range. To bring the smaller signals into the calibrated area of the display, we must increase IF gain. But in so doing, we may move the larger signals off the top of the display, above the reference level. Some Agilent analyzers, such as the PSA Series, allow measurements of signals above the reference level without affecting the accuracy with which the smaller signals are displayed. This is shown in Figure 6-8. So we can indeed take advantage of the full dynamic range of an analyzer even when the dynamic range exceeds the display range. In Figure 6-8, even though the reference level has changed from -8 dBm to -53 dBm, driving the signal far above the top of the screen, the marker readout remains unchanged.

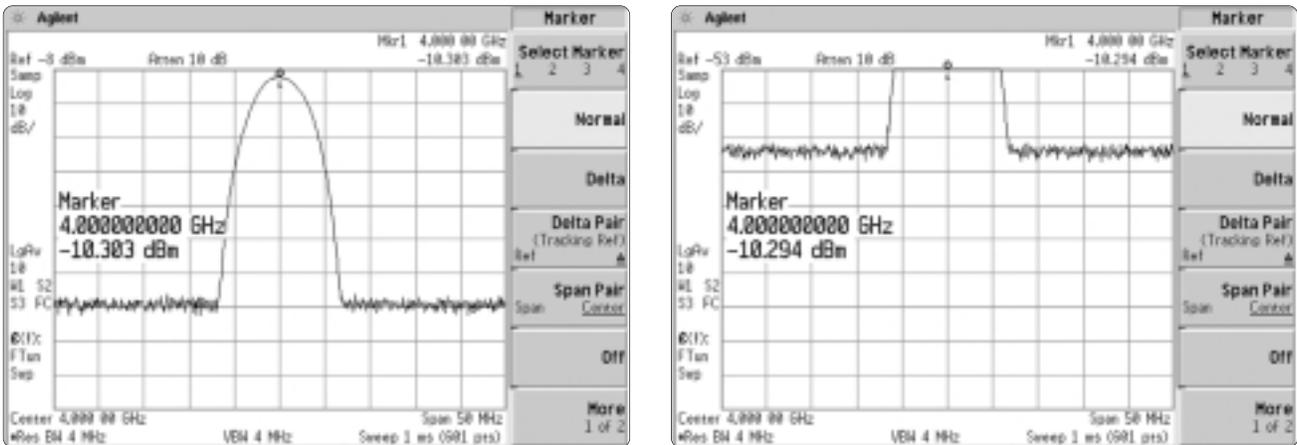


Figure 6-8. Display range and measurement range on the PSA Series

Measurement range is the ratio of the largest to the smallest signal that can be measured under any circumstances. The maximum safe input level, typically +30 dBm (1 watt) for most analyzers, determines the upper limit. These analyzers have input attenuators settable to 60 or 70 dB, so we can reduce +30 dBm signals to levels well below the compression point of the input mixer and measure them accurately. The displayed average noise level sets the other end of the range. Depending on the minimum resolution bandwidth of the particular analyzer and whether or not a preamplifier is being used, DANL typically ranges from -115 to -170 dBm. Measurement range, then, can vary from 145 to 200 dB. Of course, we cannot view a -170 dBm signal while a +30 dBm signal is also present at the input.

Adjacent channel power measurements

TOI, SOI, 1 dB gain compression, and DANL are all classic measures of spectrum analyzer performance. However, with the tremendous growth of digital communication systems, other measures of dynamic range have become increasingly important. For example, adjacent channel power (ACP) measurements are often done in CDMA-based communication systems to determine how much signal energy leaks or “spills over” into adjacent or alternate channels located above and below a carrier. An example ACP measurement is shown in Figure 6-9.

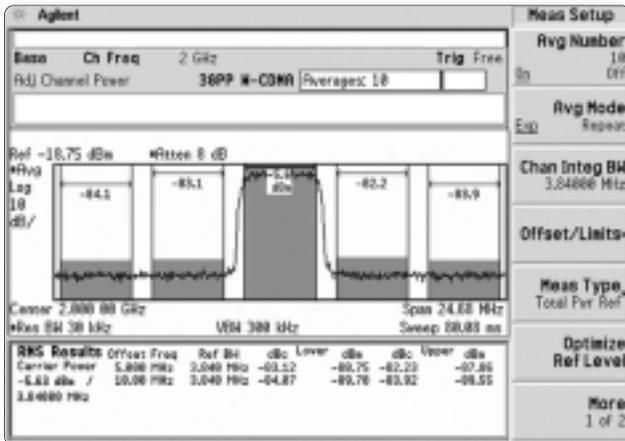


Figure 6-9. Adjacent channel power measurement using PSA Series

Note the relative amplitude difference between the carrier power and the adjacent and alternate channels. Up to six channels on either side of the carrier can be measured at a time.

Typically, we are most interested in the relative difference between the signal power in the main channel and the signal power in the adjacent or alternate channel. Depending on the particular communication standard, these measurements are often described as “adjacent channel power ratio” (ACPR) or “adjacent channel leakage ratio” (ACLR) tests. Because digitally modulated signals, as well as the distortion they generate, are very noise-like in nature, the industry standards typically define a channel bandwidth over which the signal power is integrated.

In order to accurately measure ACP performance of a device under test (DUT), such as a power amplifier, the spectrum analyzer must have better ACP performance than the device being tested. Therefore, spectrum analyzer ACPR dynamic range has become a key performance measure for digital communication systems.

Chapter 7 Extending the Frequency Range

As more wireless services continue to be introduced and deployed, the available spectrum becomes more and more crowded. Therefore, there has been an ongoing trend toward developing new products and services at higher frequencies. In addition, new microwave technologies continue to evolve, driving the need for more measurement capability in the microwave bands. Spectrum analyzer designers have responded by developing instruments capable of directly tuning up to 50 GHz using a coaxial input. Even higher frequencies can be measured using external mixing techniques. This chapter describes the techniques used to enable tuning the spectrum analyzer to such high frequencies.

Internal harmonic mixing

In Chapter 2, we described a single-range spectrum analyzer that tunes to 3 GHz. Now we wish to tune higher in frequency. The most practical way to achieve such an extended range is to use harmonic mixing.

But let us take one step at a time. In developing our tuning equation in Chapter 2, we found that we needed the low-pass filter of Figure 2-1 to prevent higher-frequency signals from reaching the mixer. The result was a uniquely responding, single band analyzer that tuned to 3 GHz. Now we wish to observe and measure higher-frequency signals, so we must remove the low-pass filter.

Other factors that we explored in developing the tuning equation were the choice of LO and intermediate frequencies. We decided that the IF should not be within the band of interest because it created a hole in our tuning range in which we could not make measurements. So we chose 3.9 GHz, moving the IF above the highest tuning frequency of interest (3 GHz). Since our new tuning range will be above 3 GHz, it seems logical to move the new IF to a frequency below 3 GHz. A typical first IF for these higher frequency ranges in Agilent spectrum analyzers is 321.4 MHz. We shall use this frequency in our examples. In summary, for the low band, up to 3 GHz, our first IF is 3.9 GHz. For the upper frequency bands, we switch to a first IF of 321.4 MHz. Note that in Figure 7-1 the second IF is already 321.4 MHz, so all we need to do when we wish to tune to the higher ranges is bypass the first IF.

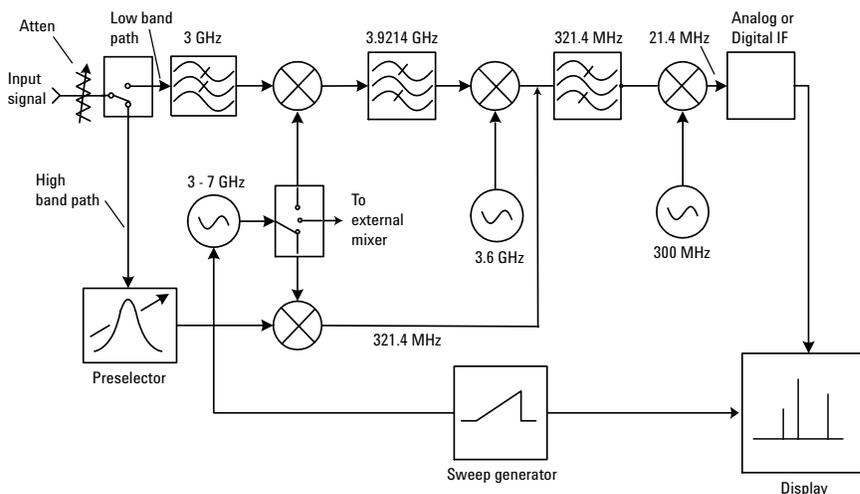


Figure 7-1. Switching arrangement for low band and high bands

In Chapter 2, we used a mathematical approach to conclude that we needed a low-pass filter. As we shall see, things become more complex in the situation here, so we shall use a graphical approach as an easier method to see what is happening. The low band is the simpler case, so we shall start with that. In all of our graphs, we shall plot the LO frequency along the horizontal axis and signal frequency along the vertical axis, as shown in Figure 7-2. We know we get a mixing product equal to the IF (and therefore a response on the display) whenever the input signal differs from the LO by the IF. Therefore, we can determine the frequency to which the analyzer is tuned simply by adding the IF to, or subtracting it from, the LO frequency. To determine our tuning range, then, we start by plotting the LO frequency against the signal frequency axis as shown by the dashed line in Figure 7-2. Subtracting the IF from the dashed line gives us a tuning range of 0 to 3 GHz, the range that we developed in Chapter 2. Note that this line in Figure 7-2 is labeled “1-” to indicate fundamental mixing and the use of the minus sign in the tuning equation. We can use the graph to determine what LO frequency is required to receive a particular signal or to what signal the analyzer is tuned for a given LO frequency. To display a 1 GHz signal, the LO must be tuned to 4.9 GHz. For an LO frequency of 6 GHz, the spectrum analyzer is tuned to receive a signal frequency of 2.1 GHz. In our text, we shall round off the first IF to one decimal place; the true IF, 3.9214 GHz, is shown on the block diagram.

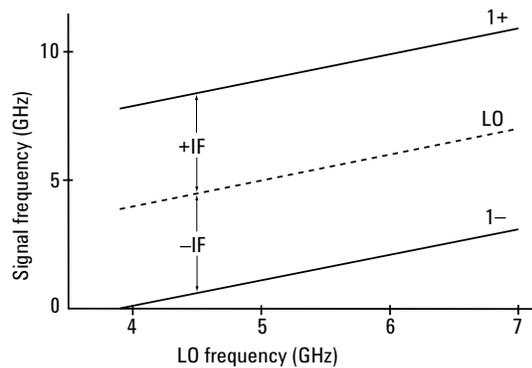


Figure 7-2. Tuning curves for fundamental mixing in the low band, high IF case

Now let's add the other fundamental-mixing band by adding the IF to the LO line in Figure 7-2. This gives us the solid upper line, labeled 1+, that indicates a tuning range from 7.8 to 10.9 GHz. Note that for a given LO frequency, the two frequencies to which the analyzer is tuned are separated by twice the IF. Assuming we have a low-pass filter at the input while measuring signals in the low band, we shall not be bothered by signals in the 1+ frequency range.

Next let's see to what extent harmonic mixing complicates the situation. Harmonic mixing comes about because the LO provides a high-level drive signal to the mixer for efficient mixing, and since the mixer is a non-linear device, it generates harmonics of the LO signal. Incoming signals can mix against LO harmonics, just as well as the fundamental, and any mixing product that equals the IF produces a response on the display. In other words, our tuning (mixing) equation now becomes:

$$f_{\text{sig}} = n f_{\text{LO}} \pm f_{\text{IF}}$$

where $n = \text{LO harmonic}$
 (Other parameters remain the same as previously discussed)

Let's add second-harmonic mixing to our graph in Figure 7-3 and see to what extent this complicates our measurement procedure. As before, we shall first plot the LO frequency against the signal frequency axis. Multiplying the LO frequency by two yields the upper dashed line of Figure 7-3. As we did for fundamental mixing, we simply subtract the IF (3.9 GHz) from and add it to the LO second-harmonic curve to produce the 2^- and 2^+ tuning ranges. Since neither of these overlap the desired 1^- tuning range, we can again argue that they do not really complicate the measurement process. In other words, signals in the 1^- tuning range produce unique, unambiguous responses on our analyzer display. The same low-pass filter used in the fundamental mixing case works equally well for eliminating responses created in the harmonic mixing case.

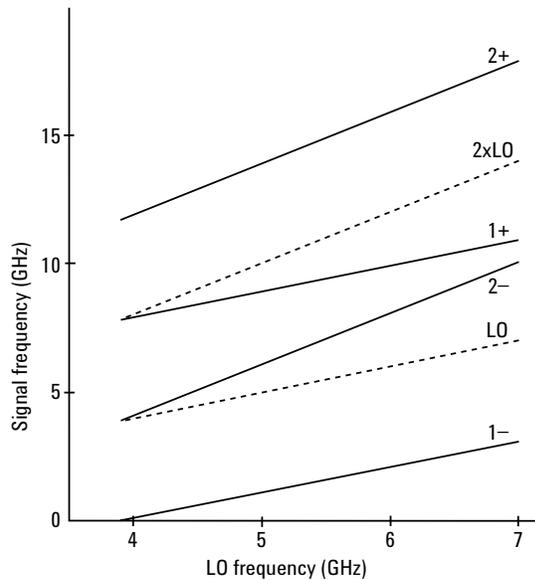


Figure 7-3. Signals in the "1 minus" frequency range produce single, unambiguous responses in the low band, high IF case

The situation is considerably different for the high band, low IF case. As before, we shall start by plotting the LO fundamental against the signal-frequency axis and then add and subtract the IF, producing the results shown in Figure 7-4. Note that the 1^- and 1^+ tuning ranges are much closer together, and in fact overlap, because the IF is a much lower frequency, 321.4 MHz in this case. Does the close spacing of the tuning ranges complicate the measurement process? Yes and no. First of all, our system can be calibrated for only one tuning range at a time. In this case, we would choose the 1^- tuning to give us a low-end frequency of about 2.7 GHz, so that we have some overlap with the 3 GHz upper end of our low band tuning range. So what are we likely to see on the display? If we enter the graph at an LO frequency of 5 GHz, we find that there are two possible signal frequencies that would give us responses at the same point on the display: 4.7 and 5.3 GHz (rounding the numbers again). On the other hand, if we enter the signal frequency axis at 5.3 GHz, we find that in addition to the 1^+ response at an LO frequency of 5 GHz, we could also get a 1^- response. This would occur if we allowed the LO to sweep as high as 5.6 GHz, twice the IF above 5 GHz. Also, if we entered the signal frequency graph at 4.7 GHz, we would find a 1^+ response at an LO frequency of about 4.4 GHz (twice the IF below 5 GHz) in addition to the 1^- response at an LO frequency of 5 GHz. Thus, for every desired response on the 1^- tuning line, there will be a second response located twice the IF frequency below it. These pairs of responses are known as *multiple responses*.

With this type of mixing arrangement, it is possible for signals at different frequencies to produce responses at the same point on the display, that is, at the same LO frequency. As we can see from Figure 7-4, input signals at 4.7 and 5.3 GHz both produce a response at the IF frequency when the LO frequency is set to 5 GHz. These signals are known as *image frequencies*, and are also separated by twice the IF frequency.

Clearly, we need some mechanism to differentiate between responses generated on the 1^- tuning curve for which our analyzer is calibrated, and those produced on the 1^+ tuning curve. However, before we look at signal identification solutions, let's add harmonic-mixing curves to 26.5 GHz and see if there are any additional factors that we must consider in the signal identification process. Figure 7-5 shows tuning curves up to the fourth harmonic of the LO.

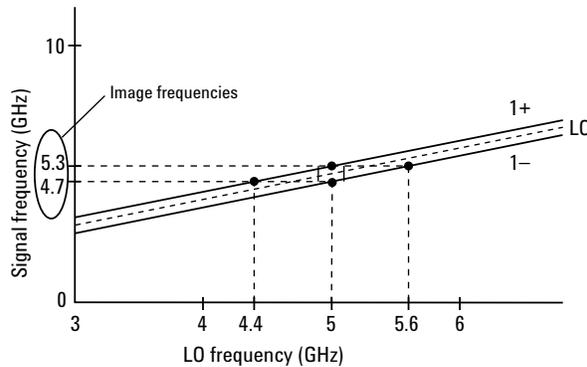


Figure 7-4. Tuning curves for fundamental mixing in the high band, low IF case

In examining Figure 7-5, we find some additional complications. The spectrum analyzer is set up to operate in several tuning bands. Depending on the frequency to which the analyzer is tuned, the analyzer display is frequency calibrated for a specific LO harmonic. For example, in the 6.2 to 13.2 GHz input frequency range, the spectrum analyzer is calibrated for the 2^- tuning curve. Suppose we have an 11 GHz signal present at the input. As the LO sweeps, the signal will produce IF responses with the 3^+ , 3^- , 2^+ and 2^- tuning curves. The desired response of the 2^- tuning curve occurs when the LO frequency satisfies the tuning equation:

$$11 \text{ GHz} = 2 f_{LO} - 0.3$$

$$f_{LO} = 5.65 \text{ GHz}$$

Similarly, we can calculate that the response from the 2^+ tuning curve occurs when $f_{LO} = 5.35$ GHz, resulting in a displayed signal that appears to be at 10.4 GHz.

The displayed signals created by the responses to the 3^+ and 3^- tuning curves are known as in-band multiple responses. Because they occur when the LO is tuned to 3.57 GHz and 3.77 GHz, they will produce false responses on the display that appear to be genuine signals at 6.84 GHz and 7.24 GHz.

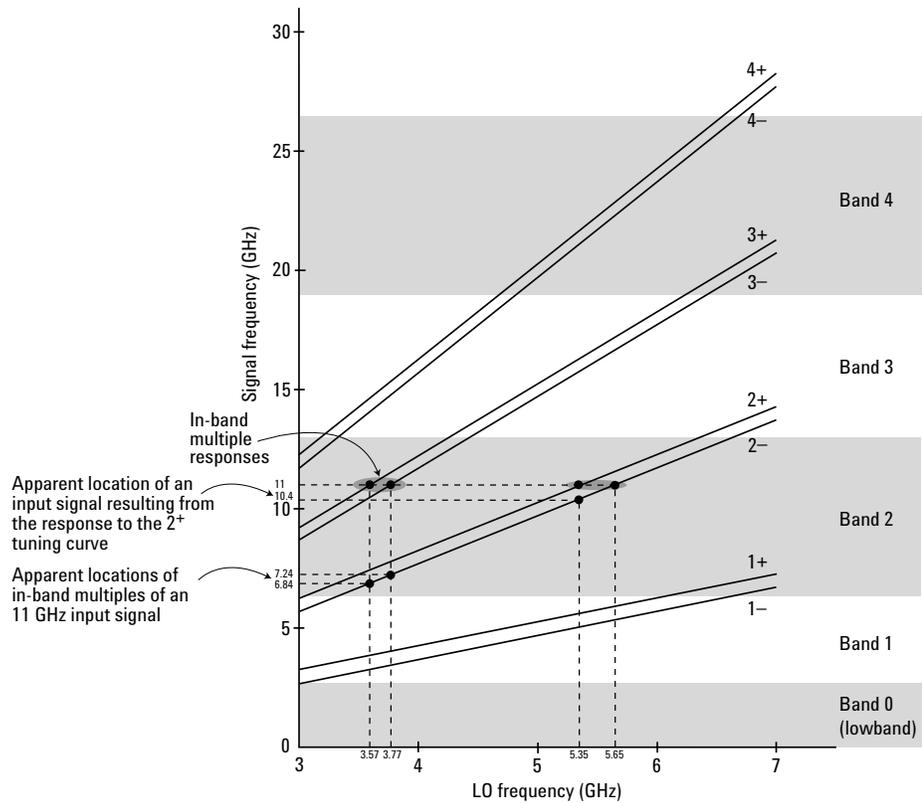


Figure 7-5. Tuning curves up to 4th harmonic of LO showing in-band multiple responses to an 11 GHz input signal.

Other situations can create out-of-band multiple responses. For example, suppose we are looking at a 5 GHz signal in band 1 that has a significant third harmonic at 15 GHz (band 3). In addition to the expected multiple pair caused by the 5 GHz signal on the 1⁺ and 1⁻ tuning curves, we also get responses generated by the 15 GHz signal on the 4⁺, 4⁻, 3⁺, and 3⁻ tuning curves. Since these responses occur when the LO is tuned to 3.675, 3.825, 4.9, and 5.1 GHz respectively, the display will show signals that appear to be located at 3.375, 3.525, 4.6, and 4.8 GHz. This is shown in Figure 7-6.

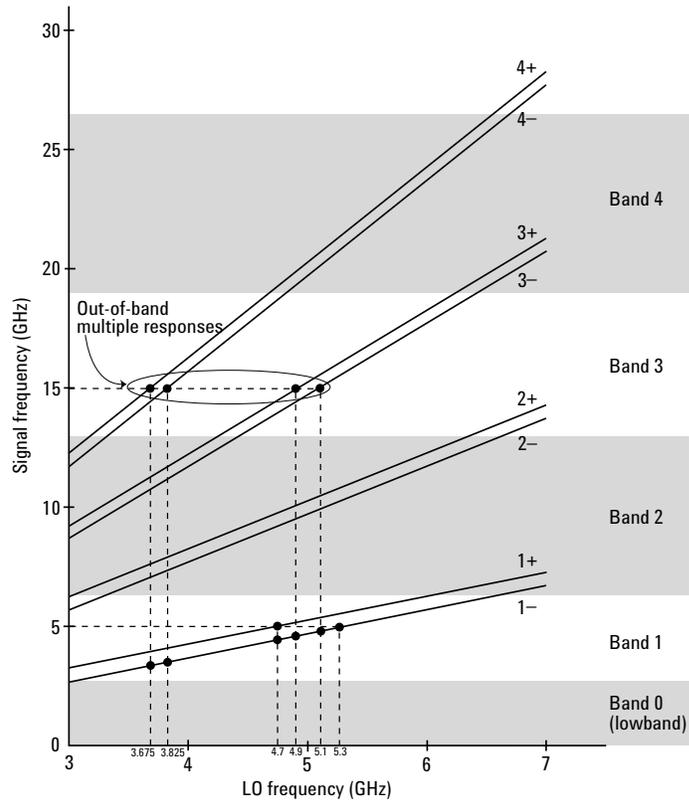


Figure 7-6. Out-of-band multiple responses in band 1 as a result of a signal in band 3

Multiple responses generally always come in pairs¹, with a “plus” mixing product and a “minus” mixing product. When we use the correct harmonic mixing number for a given tuning band, the responses will be separated by 2 times f_{IF} . Because the slope of each pair of tuning curves increases linearly with the harmonic number N , the multiple pairs caused by any other harmonic mixing number appear to be separated by:

$$2f_{IF} (N_c/N_A)$$

where N_c = the correct harmonic number for the desired tuning band
 N_A = the actual harmonic number generating the multiple pair

1. Often referred to as an “image pair.” This is inaccurate terminology, since images are actually two or more real signals present at the spectrum analyzer input that produce an IF response at the same LO frequency.

Can we conclude from this discussion that a harmonic mixing spectrum analyzer is not practical? Not necessarily. In cases where the signal frequency is known, we can tune to the signal directly, knowing that the analyzer will select the appropriate mixing mode for which it is calibrated. In controlled environments with only one or two signals, it is usually easy to distinguish the real signal from the image and multiple responses. However, there are many cases in which we have no idea how many signals are involved or what their frequencies might be. For example, we could be searching for unknown spurious signals, conducting site surveillance tests as part of a frequency-monitoring program, or performing EMI tests to measure unwanted device emissions. In all these cases, we could be looking for totally unknown signals in a potentially crowded spectral environment. Having to perform some form of identification routine on each and every response would make measurement time intolerably long.

Fortunately, there is a way to essentially eliminate image and multiple responses through a process of prefiltering the signal. This technique is called preselection.

Preselection

What form must our preselection take? Referring back to Figure 7-4, assume that we have two signals at 4.7 and 5.3 GHz present at the input of our analyzer. If we were particularly interested in one, we could use a band-pass filter to allow that signal into the analyzer and reject the other. However, the fixed filter does not eliminate multiple responses; so if the spectrum is crowded, there is still potential for confusion. More important, perhaps, is the restriction that a fixed filter puts on the flexibility of the analyzer. If we are doing broadband testing, we certainly do not want to be continually forced to change band-pass filters.

The solution is a tunable filter configured in such a way that it automatically tracks the frequency of the appropriate mixing mode. Figure 7-7 shows the effect of such a preselector. Here we take advantage of the fact that our superheterodyne spectrum analyzer is not a real-time analyzer; that is, it tunes to only one frequency at a time. The dashed lines in Figure 7-7 represent the bandwidth of the tracking preselector. Signals beyond the dashed lines are rejected. Let's continue with our previous example of 4.7 and 5.3 GHz signals present at the analyzer input. If we set a center frequency of 5 GHz and a span of 2 GHz, let's see what happens as the analyzer tunes across this range. As the LO sweeps past 4.4 GHz (the frequency at which it could mix with the 4.7 GHz input signal on its 1^+ mixing mode), the preselector is tuned to 4.1 GHz and therefore rejects the 4.7 GHz signal. Since the input signal does not reach the mixer, no mixing occurs, and no response appears on the display. As the LO sweeps past 5 GHz, the preselector allows the 4.7 GHz signal to reach the mixer, and we see the appropriate response on the display. The 5.3 GHz image signal is rejected, so it creates no mixing product to interact with the mixing product from the 4.7 GHz signal and cause a false display. Finally, as the LO sweeps past 5.6 GHz, the preselector allows the 5.3 GHz signal to reach the mixer, and we see it properly displayed. Note in Figure 7-7 that nowhere do the various mixing modes intersect. So as long as the preselector bandwidth is narrow enough (it typically varies from about 35 MHz at low frequencies to 80 MHz at high frequencies) it will greatly attenuate all image and multiple responses.

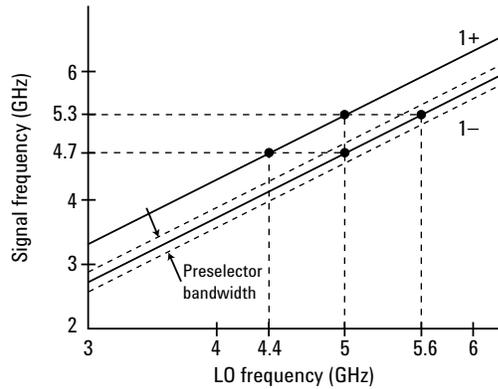


Figure 7-7. Preselection; dashed lines represent bandwidth of tracking preselector

The word eliminate may be a little strong. Preselectors do not have infinite rejection. Something in the 70 to 80 dB range is more likely. So if we are looking for very low-level signals in the presence of very high-level signals, we might see low-level images or multiples of the high-level signals. What about the low band? Most tracking preselectors use YIG technology, and YIG filters do not operate well at low frequencies. Fortunately, there is a simple solution. Figure 7-3 shows that no other mixing mode overlaps the 1^- mixing mode in the low frequency, high IF case. So a simple low-pass filter attenuates both image and multiple responses. Figure 7-8 shows the input architecture of a typical microwave spectrum analyzer.

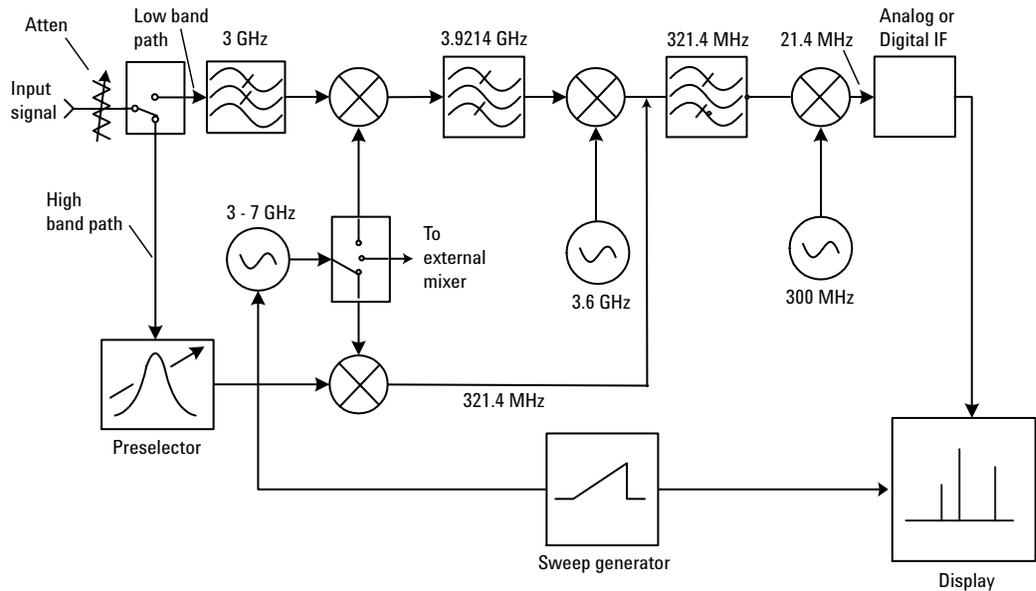


Figure 7-8. Front-end architecture of a typical preselected spectrum analyzer

Amplitude calibration

So far, we have looked at how a harmonic mixing spectrum analyzer responds to various input frequencies. What about amplitude?

The conversion loss of a mixer is a function of harmonic number, and the loss goes up as the harmonic number goes up. This means that signals of equal amplitude would appear at different levels on the display if they involved different mixing modes. To preserve amplitude calibration, then, something must be done. In Agilent spectrum analyzers, the IF gain is changed. The increased conversion loss at higher LO harmonics causes a loss of sensitivity just as if we had increased the input attenuator. And since the IF gain change occurs after the conversion loss, the gain change is reflected by a corresponding change in the displayed noise level. So we can determine analyzer sensitivity on the harmonic-mixing ranges by noting the average displayed noise level just as we did on fundamental mixing.

In older spectrum analyzers, the increase in displayed average noise level with each harmonic band was very noticeable. More recent models of Agilent spectrum analyzers use a double-balanced, image-enhanced harmonic mixer to minimize the increased conversion loss when using higher harmonics. Thus, the “stair step” effect on DANL has been replaced by a gentle sloping increase with higher frequencies. This can be seen in Figure 7-9.

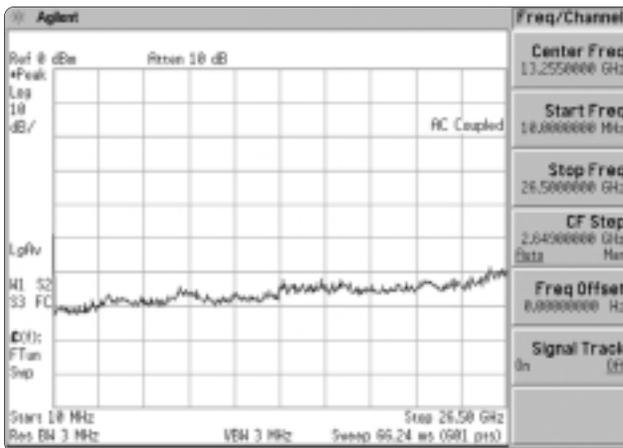


Figure 7-9. Rise in noise floor indicates changes in sensitivity with changes in LO harmonic used

Phase noise

In Chapter 2, we noted that instability of an analyzer LO appears as phase noise around signals that are displayed far enough above the noise floor. We also noted that this phase noise can impose a limit on our ability to measure closely spaced signals that differ in amplitude. The level of the phase noise indicates the angular, or frequency, deviation of the LO. What happens to phase noise when a harmonic of the LO is used in the mixing process? Relative to fundamental mixing, phase noise (in decibels) increases by:

$$20 \log(N),$$

where N = harmonic of the LO

For example, suppose that the LO fundamental has a peak-to-peak deviation of 10 Hz. The second harmonic then has a 20 Hz peak-to-peak deviation; the third harmonic, 30 Hz; and so on. Since the phase noise indicates the signal (noise in this case) producing the modulation, the level of the phase noise must be higher to produce greater deviation. When the degree of modulation is very small, as in the situation here, the amplitude of the modulation side bands is directly proportional to the deviation of the carrier (LO). If the deviation doubles, the level of the side bands must also double in voltage; that is, increase by 6 dB or $20 \log(2)$. As a result, the ability of our analyzer to measure closely spaced signals that are unequal in amplitude decreases as higher harmonics of the LO are used for mixing. Figure 7-10 shows the difference in phase noise between fundamental mixing of a 5 GHz signal and fourth-harmonic mixing of a 20 GHz signal.

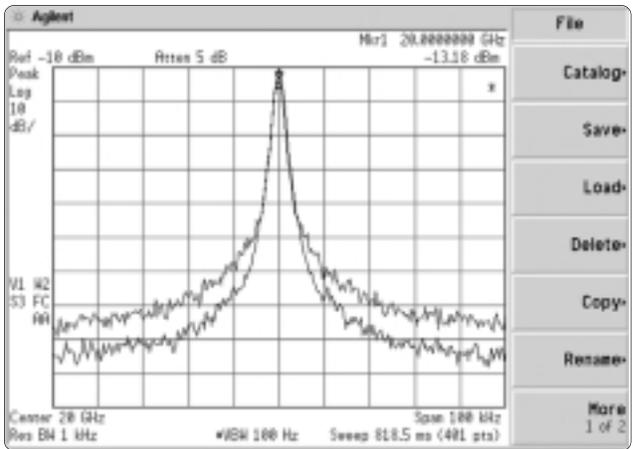


Figure 7-10. Phase noise levels for fundamental and 4th harmonic mixing

Improved dynamic range

A preselector improves dynamic range if the signals in question have sufficient frequency separation. The discussion of dynamic range in Chapter 6 assumed that both the large and small signals were always present at the mixer and that their amplitudes did not change during the course of the measurement. But as we have seen, if signals are far enough apart, a preselector allows one to reach the mixer while rejecting the others. For example, if we were to test a microwave oscillator for harmonics, a preselector would reject the fundamental when we tuned the analyzer to one of the harmonics.

Let's look at the dynamic range of a second-harmonic test of a 3 GHz oscillator. Using the example from Chapter 6, suppose that a -40 dBm signal at the mixer produces a second harmonic product of -75 dBc. We also know, from our discussion, that for every dB the level of the fundamental changes at the mixer, measurement range also changes by 1 dB. The second-harmonic distortion curve is shown in Figure 7-11. For this example, we shall assume plenty of power from the oscillator and set the input attenuator so that when we measure the oscillator fundamental, the level at the mixer is -10 dBm, below the 1 dB compression point.

From the graph, we see that a -10 dBm signal at the mixer produces a second-harmonic distortion component of -45 dBc. Now we tune the analyzer to the 6 GHz second harmonic. If the preselector has 70 dB rejection, the fundamental at the mixer has dropped to -80 dBm. Figure 7-11 indicates that for a signal of -80 dBm at the mixer, the internally generated distortion is -115 dBc, meaning 115 dB below the new fundamental level of -80 dBm. This puts the absolute level of the harmonic at -195 dBm. So the difference between the fundamental we tuned to and the internally generated second harmonic we tuned to is 185 dB! Clearly, for harmonic distortion, dynamic range is limited on the low-level (harmonic) end only by the noise floor (sensitivity) of the analyzer.

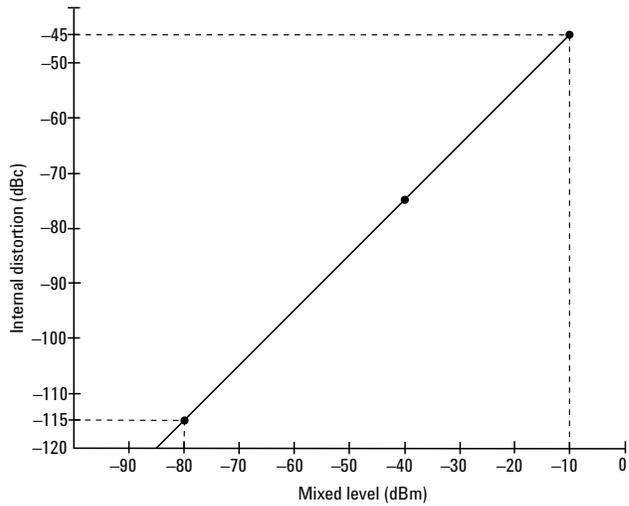


Figure 7-11. Second-order distortion graph

What about the upper, high-level end? When measuring the oscillator fundamental, we must limit power at the mixer to get an accurate reading of the level. We can use either internal or external attenuation to limit the level of the fundamental at the mixer to something less than the 1 dB compression point. However, since the preselector highly attenuates the fundamental when we are tuned to the second harmonic, we can remove some attenuation if we need better sensitivity to measure the harmonic. A fundamental level of $+20$ dBm at the preselector should not affect our ability to measure the harmonic.

Any improvement in dynamic range for third-order intermodulation measurements depends upon separation of the test tones versus preselector bandwidth. As we noted, typical preselector bandwidth is about 35 MHz at the low end and 80 MHz at the high end. As a conservative figure, we might use 18 dB per octave of bandwidth roll off of a typical YIG preselector filter beyond the 3 dB point. So to determine the improvement in dynamic range, we must determine to what extent each of the fundamental tones is attenuated and how that affects internally generated distortion. From the expressions in Chapter 6 for third-order intermodulation, we have:

$$(k_4/8)V_{LO}V_1^2V_2 \cos[\omega_{LO} - (2\omega_1 - \omega_2)]t$$

and

$$(k_4/8)V_{LO}V_1V_2^2 \cos[\omega_{LO} - (2\omega_2 - \omega_1)]t$$

Looking at these expressions, we see that the amplitude of the lower distortion component ($2\omega_1 - \omega_2$) varies as the square of V_1 and linearly with V_2 . On the other side, the amplitude of the upper distortion component ($2\omega_2 - \omega_1$) varies linearly with V_1 and as the square of V_2 . However, depending on the signal frequencies and separation, the preselector may not attenuate the two fundamental tones equally.

Consider the situation shown in Figure 7-12 in which we are tuned to the lower distortion component and the two fundamental tones are separated by half the preselector bandwidth. In this case, the lower-frequency test tone lies at the edge of the preselector pass band and is attenuated 3 dB. The upper test tone lies above the lower distortion component by an amount equal to the full preselector bandwidth. It is attenuated approximately 21 dB. Since we are tuned to the lower distortion component, internally generated distortion at this frequency drops by a factor of two relative to the attenuation of V_1 (2 times 3 dB = 6 dB) and equally as fast as the attenuation of V_2 (21 dB). The improvement in dynamic range is the sum of 6 dB + 21 dB, or 27 dB. As in the case of second harmonic distortion, the noise floor of the analyzer must be considered, too. For very closely spaced test tones, the preselector provides no improvement, and we determine dynamic range as if the preselector was not there.

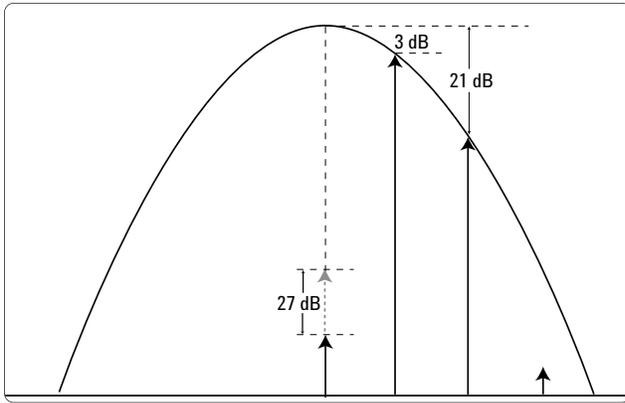


Figure 7-12. Improved third-order intermodulation distortion; test tone separation is significant, relative to preselector bandwidth

The discussion of dynamic range in Chapter 6 applies to the low-pass-filtered low band. The only exceptions occur when a particular harmonic of a low band signal falls within the preselected range. For example, if we measure the second harmonic of a 2.5 GHz fundamental, we get the benefit of the preselector when we tune to the 5 GHz harmonic.

Pluses and minuses of preselection

We have seen the pluses of preselection: simpler analyzer operation, uncluttered displays, improved dynamic range, and wide spans. But there are some minuses, relative to an unpreselected analyzer, as well.

First of all, the preselector has insertion loss, typically 6 to 8 dB. This loss comes prior to the first stage of gain, so system sensitivity is degraded by the full loss. In addition, when a preselector is connected directly to a mixer, the interaction of the mismatch of the preselector with that of the input mixer can cause a degradation of frequency response. Proper calibration techniques must be used to compensate for this ripple. Another approach to minimize this interaction would be to insert a matching pad (fixed attenuator) or isolator between the preselector and mixer. In this case, sensitivity would be degraded by the full value of the pad or isolator.

Some spectrum analyzer architectures eliminate the need for the matching pad or isolator. As the electrical length between the preselector and mixer increases, the rate of change of phase of the reflected and re-reflected signals becomes more rapid for a given change in input frequency. The result is a more exaggerated ripple effect on flatness. Architectures such as those used in the ESA Series and PSA Series include the mixer diodes as an integral part of the preselector/mixer assembly. In such an assembly, there is minimal electrical length between the preselector and mixer. This architecture thus removes the ripple effect on frequency response and improves sensitivity by eliminating the matching pad or isolator.

Even aside from its interaction with the mixer, a preselector causes some degradation of frequency response. The preselector filter pass band is never perfectly flat, but rather exhibits a certain amount of ripple. In most configurations, the tuning ramp for the preselector and local oscillator come from the same source, but there is no feedback mechanism to ensure that the preselector exactly tracks the tuning of the analyzer. Another source of post-tuning drift is the self-heating caused by current flowing in the preselector circuitry. The center of the preselector passband will depend on its temperature and temperature gradients. These will depend on the history of the preselector tuning. As a result, the best flatness is obtained by centering the preselector at each signal. The centering function is typically built into the spectrum analyzer firmware and selected either by a front panel key in manual measurement applications, or programmatically in automated test systems. When activated, the centering function adjusts the preselector tuning DAC to center the preselector pass band on the signal. The frequency response specification for most microwave analyzers only applies after centering the preselector, and it is generally a best practice to perform this function (to mitigate the effects of post-tuning drift) before making amplitude measurements of microwave signals.

External harmonic mixing

We have discussed tuning to higher frequencies within the spectrum analyzer. For internal harmonic mixing, the ESA and PSA spectrum analyzers use the second harmonic ($N=2^-$) to tune to 13.2 GHz, and the fourth harmonic ($N=4^-$) to tune to 26.5 GHz. However, what if you want to test outside the upper frequency range of the spectrum analyzer? Some spectrum analyzers provide the ability to bypass the internal first mixer and preselector and use an external mixer to enable the spectrum analyzer to make high frequency measurements². For external mixing we can use higher harmonics of the 1st LO. Typically, a spectrum analyzer that supports external mixing has two additional connectors on the front panel. An LO OUT port routes the internal first LO signal to the external mixer, which uses the higher harmonics to mix with the high frequency signals. The external mixer's IF output connects to the analyzer's IF IN port. As long as the external mixer uses the same IF frequency as the spectrum analyzer, the signal can be processed and displayed internally, just like any signal that came from the internal first mixer. Figure 7-13 illustrates the block diagram of an external mixer used in conjunction with a spectrum analyzer.

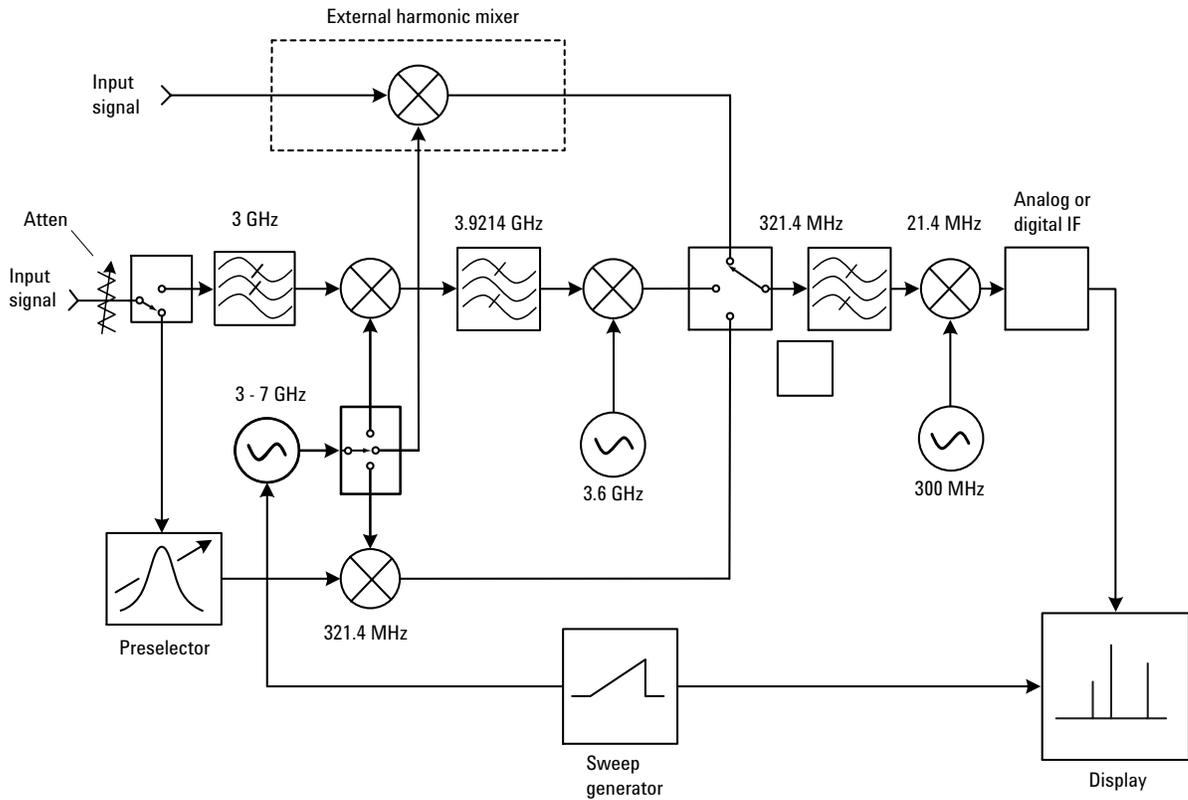


Figure 7-13. Spectrum analyzer and external mixer block diagram

2. For more information on external mixing, see Agilent Application Note 1485, *External Waveguide Mixing and Millimeter Wave Measurements with Agilent PSA Spectrum Analyzers*, literature number 5988-9414EN.

Table 7-1 shows the harmonic mixing modes used by the ESA and PSA at various millimeter wave bands. You choose the mixer depending on the frequency range you need. Typically, these are standard waveguide bands. There are two kinds of external harmonic mixers; those with preselection and those without. Agilent offers unpreselected mixers in six frequency bands: 18 to 26.5 GHz, 26.5 to 40 GHz, 33 to 50 GHz, 40 to 60 GHz, 50 to 75 GHz, and 75 to 110 GHz. Agilent also offers four preselected mixers up to 75 GHz. Above 110 GHz, mixers are available from other commercial manufacturers for operation up to 325 GHz.

Some external mixers from other manufacturers require a bias current to set the mixer diodes to the proper operating point. The ESA and PSA spectrum analyzers can provide up to ± 10 mA of DC current through the IF OUT port to provide this bias and keep the measurement setup as simple as possible.

Table 7-1. Harmonic mixing modes used by ESA-E and PSA Series with external mixers

Band	Harmonic mixing mode (N ^a)	
	Preselected	Unpreselected
K (18.0 to 26.5 GHz)	n/a	6 ⁻
A (26.5 to 40.0 GHz)	8 ⁺	8 ⁻
Q (33.0 to 50.0 GHz)	10 ⁺	10 ⁻
U (40.0 to 60.0 GHz)	10 ⁺	10 ⁻
V (50.0 to 75.0 GHz)	14 ⁺	14 ⁻
E (60.0 to 90.0 GHz)	n/a	16 ⁻
W (75.0 to 110.0 GHz)	n/a	18 ⁻
F (90.0 to 140.0 GHz)	n/a	20 ⁻
D (110.0 to 170.0 GHz)	n/a	24 ⁻
G (140.0 to 220.0 GHz)	n/a	32 ⁻
Y (170.0 to 260.0 GHz)	n/a	38 ⁻
J (220.0 to 325.0 GHz)	n/a	46 ⁻

Whether performing harmonic mixing with an internal or an external mixer, the issues are similar. The LO and its harmonics mix not only with the RF input signal, but any other signal that may be present at the input as well. This produces mixing products that can be processed through the IF just like any other valid signals. There are two ways to deal with these unwanted signals. A preselector designed into the external mixer will offer you the same type of tunable filter, as in the spectrum analyzer, for the frequency band of interest. Figure 7-14 shows a spectrum analyzer and an external mixer with internal preselection. The benefits and drawbacks of a preselected external mixer are very similar to those for the preselector inside the spectrum analyzer. The most significant drawback of preselected mixers is the increased insertion loss due to the filter, resulting in lower sensitivity for the measurement. Preselected mixers are also significantly more expensive than unpreselected mixers. For these reasons, another way to deal with these unwanted signals has been designed into the spectrum analyzer. This function is called “signal identification.”

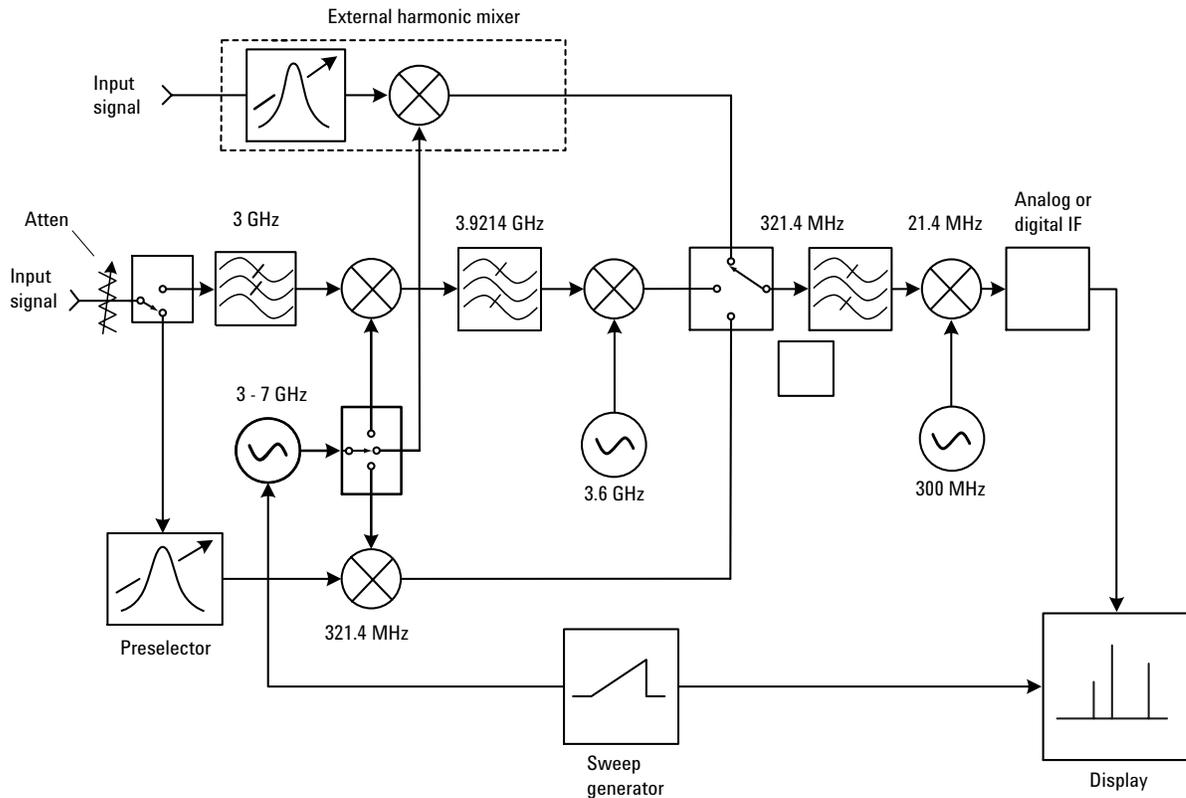


Figure 7-14. Block diagram of spectrum analyzer and external mixer with built-in preselector

Signal identification

Even when using an unpreselected mixer in a controlled situation, there are times when we must contend with unknown signals. In such cases, it is quite possible that the particular response we have tuned onto the display has been generated on an LO harmonic or mixing mode other than the one for which the display is calibrated. So our analyzer must have some way to tell us whether or not the display is calibrated for the signal response in question. For the purposes of this example, assume that we are using an Agilent 11970V 50 to 75 GHz unpreselected mixer, which uses the 14^- mixing mode. A portion of this millimeter band can be seen in Figure 7-15.

The Agilent E4407B ESA-E spectrum analyzer offers two different identification methods: *Image shift* and *Image suppress*. We shall first consider the image shift method. Looking at Figure 7-16, let's assume that we have tuned the analyzer to a frequency of 58.5 GHz. The 14^{th} harmonic of the LO produces a pair of responses, where the 14^- mixing product appears on screen at the correct frequency of 58.5 GHz, while the 14^+ mixing product produces a response with an indicated frequency of 57.8572 GHz, which is 2 times f_{IF} below the real response. Since the ESA has an IF frequency of 321.4 MHz, the pair of responses is separated by 642.8 MHz.

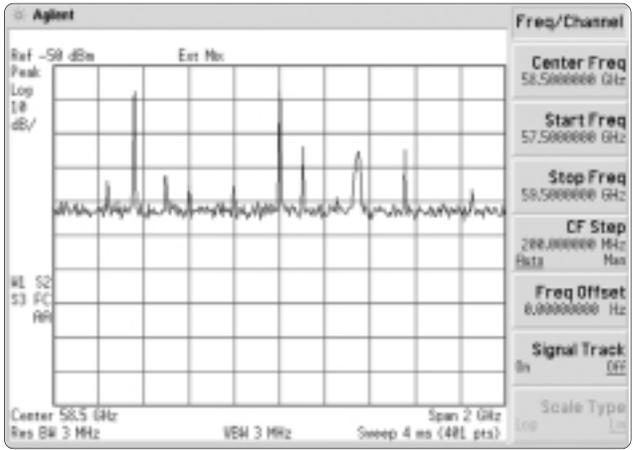


Figure 7-15. Which ones are the real signals?

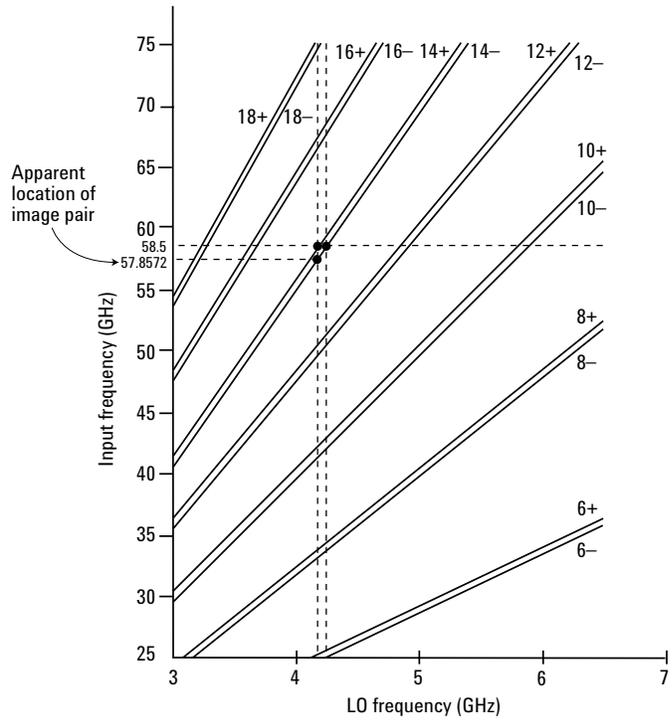


Figure 7-16. Harmonic tuning lines for the E4407B ESA-E spectrum analyzer

Let's assume that we have some idea of the characteristics of our signal, but we do not know its exact frequency. How do we determine which is the real signal? The image-shift process retunes the LO fundamental frequency by an amount equal to $2f_{IF}/N$. This causes the Nth harmonic to shift by $2f_{IF}$. If we are tuned to a real signal, its corresponding pair will now appear at the same position on screen that the real signal occupied in the first sweep. If we are tuned to another multiple pair created by some other incorrect harmonic, the signal will appear to shift in frequency on the display. The ESA spectrum analyzer shifts the LO on alternate sweeps, creating the two displays shown in Figures 7-17a and 7-17b. In Figure 7-17a, the real signal (the 14^- mixing product) is tuned to the center of the screen. Figure 7-17b shows how the image shift function moves the corresponding pair (the 14^+ mixing product) to the center of the screen.

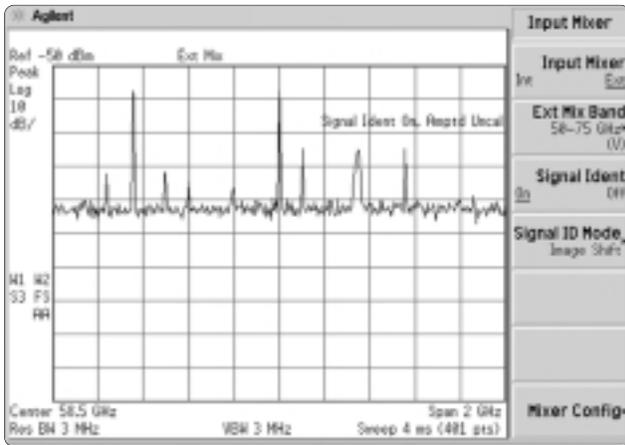


Figure 7-17a. 14^- centered

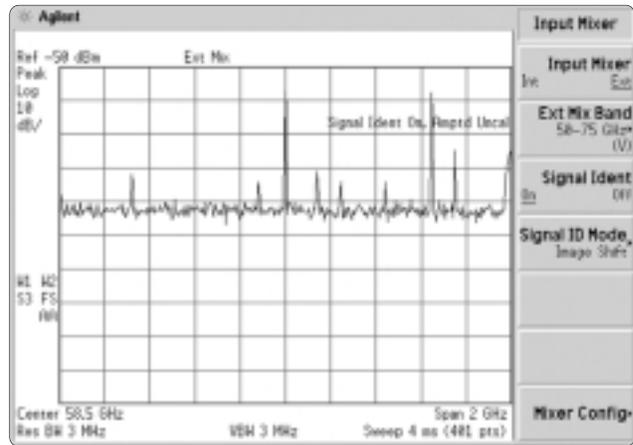


Figure 7-17b. 14^+ centered

Figure 7-17. Alternate sweeps taken with the image shift function

Let's examine the second method of signal identification, called image suppression. In this mode, two sweeps are taken using the MIN HOLD function, which saves the smaller value of each display point, or bucket, from the two sweeps. The first sweep is done using normal LO tuning values. The second sweep offsets the LO fundamental frequency by $2f_{IF}/N$. As we saw in the first signal ID method, the image product generated by the correct harmonic will land at the same point on the display as the real signal did on the first sweep. Therefore, the trace retains a high amplitude value. Any false response that shifts in frequency will have its trace data replaced by a lower value. Thus, all image and incorrect multiple responses will appear as noise. This is shown in Figure 7-18.

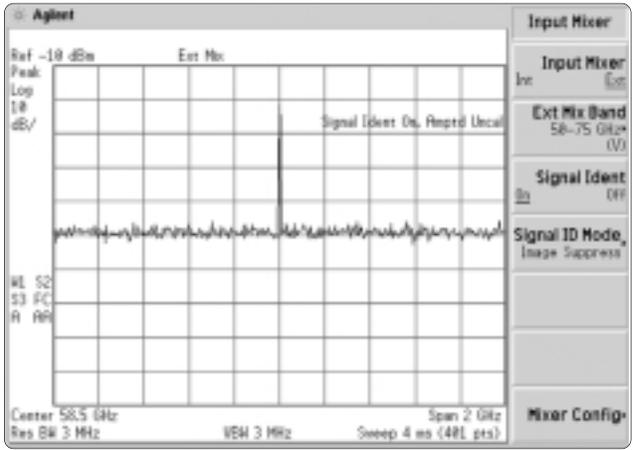


Figure 7-18. The image suppress function displays only real signals

Note that both signal identification methods are used for identifying correct frequencies only. You should not attempt to make amplitude measurements while the signal identification function is turned on. Note that in both Figures 7-17 and 7-18, an on-screen message alerts the user to this fact. Once we have identified the real signal of interest, we turn off the signal ID function and zoom in on it by reducing the span. We can then measure the signal's amplitude and frequency. See Figure 7-19.

To make an accurate amplitude measurement, it is very important that you first enter the calibration data for your external mixer. This data is normally supplied by the mixer manufacturer, and is typically a table of mixer conversion loss, in dB, at a number of frequency points across the band. This data is entered into the ESA's amplitude correction table. This table is accessed by pressing the [AMPLITUDE] key, then pressing the {More}, {Corrections}, {Other} and {Edit} softkeys. After entering the conversion loss values, apply the corrections with the {Correction On} softkey. The spectrum analyzer reference level is now calibrated for signals at the input to the external mixer. If you have other loss or gain elements between the signal source and the mixer, such as antennas, cables, or preamplifiers, the frequency responses of these elements should be entered into the amplitude correction table as well.

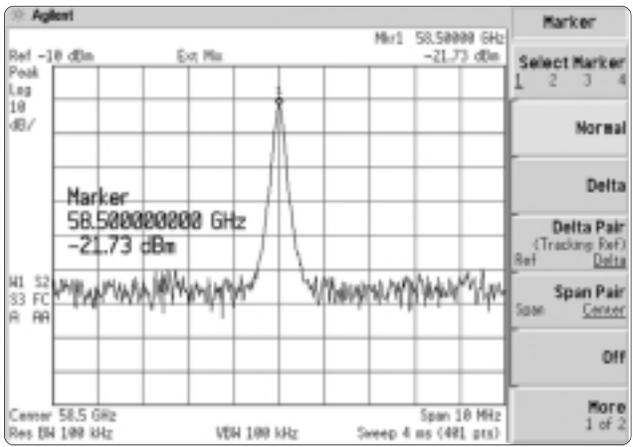


Figure 7-19. Measurement of positively identified signal

Chapter 8

Modern Spectrum Analyzers

In previous chapters of this application note, we have looked at the fundamental architecture of spectrum analyzers and basic considerations for making frequency-domain measurements. On a practical level, modern spectrum analyzers must also handle many other tasks to help you accomplish your measurement requirements. These tasks include:

- Providing application-specific measurements, such as adjacent channel power (ACP), noise figure, and phase noise
- Providing digital modulation analysis measurements defined by industry or regulatory standards, such as GSM, cdma2000, 802.11, or *Bluetooth*
- Performing vector signal analysis
- Saving data
- Printing data
- Transferring data, via an I/O bus, to a computer
- Offering remote control and operation over GPIB, LAN, or the Internet
- Allowing you to update instrument firmware to add new features and capabilities, as well as to repair defects
- Making provisions for self-calibration, troubleshooting, diagnostics, and repair
- Recognizing and operating with optional hardware and/or firmware to add new capabilities

Application-specific measurements

In addition to measuring general signal characteristics like frequency and amplitude, you often need to make specific measurements of certain signal parameters. Examples include channel power measurements and adjacent channel power (ACP) measurements, which were previously described in Chapter 6. Many spectrum analyzers now have these built-in functions available. You simply specify the channel bandwidth and spacing, then press a button to activate the automatic measurement.

The complementary cumulative distribution function (CCDF), showing power statistics, is another measurement capability increasingly found in modern spectrum analyzers. This is shown in Figure 8-1. CCDF measurements provide statistical information showing the percent of time the instantaneous power of the signal exceeds the average power by a certain number of dB. This information is important in power amplifier design, for example, where it is important to handle instantaneous signal peaks with minimum distortion while minimizing cost, weight, and power consumption of the device.

Other examples of built-in measurement functions include occupied bandwidth, TOI and harmonic distortion, and spurious emissions measurements. The instrument settings, such as center frequency, span, and resolution bandwidth, for these measurements depend on the specific radio standard to which the device is being tested. Most modern spectrum analyzers have these instrument settings stored in memory so that you can select the desired radio standard (GSM/EDGE, cdma2000, W-CDMA, 802.11a/b/g, and so on) to properly make the measurements.

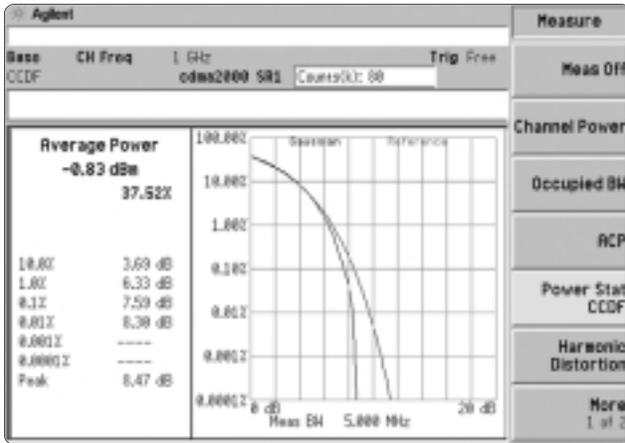


Figure 8-1. CCDF measurement

RF designers are often concerned with the noise figure of their devices, as this directly affects the sensitivity of receivers and other systems. Some spectrum analyzers, such as the PSA Series and ESA-E Series models, have optional noise figure measurement capabilities available. This option provides control for the noise source needed to drive the input of the device under test (DUT), as well as firmware to automate the measurement process and display the results. Figure 8-2 shows a typical measurement result, showing DUT noise figure (upper trace) and gain (lower trace) as a function of frequency. For more information on noise figure measurements using a spectrum analyzer, see Agilent *Application Note 1439, Measuring Noise Figure with a Spectrum Analyzer*, literature number 5988-8571EN.

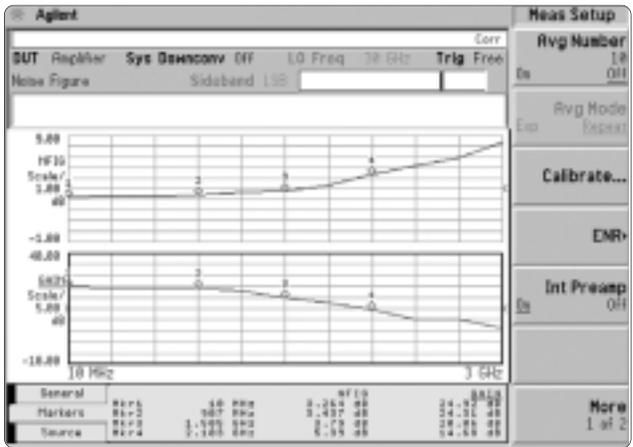


Figure 8-2. Noise figure measurement

Similarly, phase noise is a common measure of oscillator performance. In digitally modulated communication systems, phase noise can negatively impact bit error rates. Phase noise can also degrade the ability of Doppler radar systems to capture the return pulses from targets. Many Agilent spectrum analyzers, including the ESA, PSA, and 8560 Series offer optional phase noise measurement capabilities. These options provide firmware to control the measurement and display the phase noise as a function of frequency offset from the carrier, as shown in Figure 8-3.

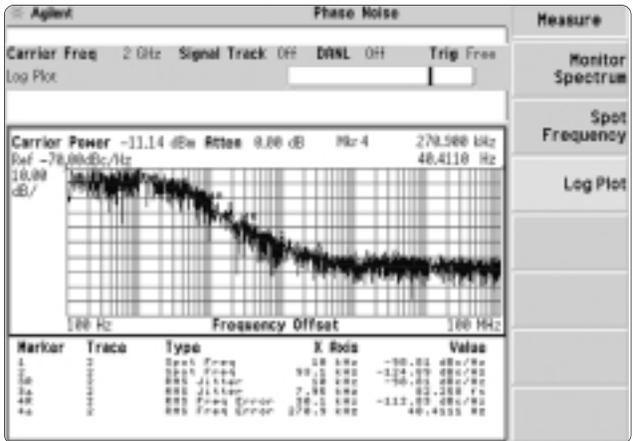


Figure 8-3. Phase Noise measurement

Digital modulation analysis

The common wireless communication systems used throughout the world today all have prescribed measurement techniques defined by standards-development organizations and governmental regulatory bodies. Optional measurement personalities are commonly available on spectrum analyzers to perform the key tests defined for a particular communication format. For example, if we need to test a transmitter to the Bluetooth wireless communication standard, we must measure parameters such as:

- Average/peak output power
- Modulation characteristics
- Initial carrier frequency tolerance
- Carrier frequency drift
- Monitor band/channel
- Modulation overview
- Output spectrum
- 20 dB bandwidth
- Adjacent channel power

These measurements are available on the Agilent ESA-E Series spectrum analyzer with appropriate options. For more information on *Bluetooth* measurements, please refer to Agilent *Application Note 1333, Performing Bluetooth RF Measurements Today*, literature number 5968-7746E. Other communication standards-based measurement personalities available on the ESA-E Series include cdmaOne and GSM/GPRS/EDGE.

Measurement capabilities for a wide variety of wireless communications standards are also available for the PSA Series spectrum analyzers. Optional measurement personalities include:

- GSM/EDGE
- W-CDMA
- HSDPA
- cdma2000
- 1xEV-DO
- 1xEV-DV
- cdmaOne
- NADC and PDC
- TD-SCDMA

Figure 8-4 illustrates an error vector magnitude (EVM) measurement performed on a GSM/EDGE signal. This test helps diagnose modulation or amplification distortions that lead to bit errors in the receiver.

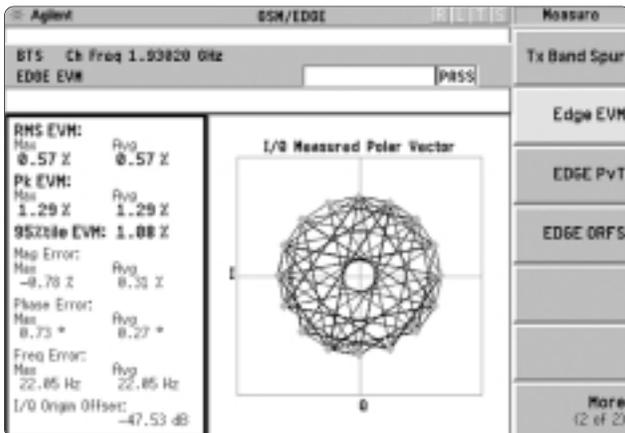


Figure 8-4. EVM measurement results and constellation display

Not all digital communication systems are based on well-defined industry standards. Engineers working on non-standard proprietary systems or the early stages of proposed industry-standard formats need more flexibility to analyze vector-modulated signals under varying conditions. This can be accomplished in two ways. First, modulation analysis personalities are available on a number of spectrum analyzers. Alternatively, more extensive analysis can be done with software running on an external computer. For example, the Agilent 89600 Series vector signal analysis software can be used with either the ESA or PSA Series spectrum analyzers to provide flexible vector signal analysis. In this case, the spectrum analyzer acts as an RF downconverter and digitizer. The software communicates with the spectrum analyzer over a GPIB or LAN connection and transfers IQ data to the computer, where it performs the vector signal analysis. Measurement settings, such as modulation type, symbol rate, filtering, triggering, and record length, can be varied as necessary for the particular signal being analyzed.

Saving and printing data

After making a measurement, we normally want to keep a record of the test data. We might simply want to make a quick printout of the instrument display. Depending on the particular analyzer and printer model, we might use the parallel, RS-232, or GPIB ports to connect the two units.

Very often, we may want to save measurement data as a file, either in the spectrum analyzer's internal memory or on a mass-storage device such as a floppy disk. In this case, there are several different kinds of data we may wish to save. This could include:

- **An image of the display** - Preferably in a popular file format, such as bitmap, .GIF, or Windows metafile.
- **Trace data** - Saved as X-Y data pairs representing frequency and amplitude points on the screen. The number of data pairs can vary. Modern spectrum analyzers such as the ESA and PSA Series allow you to select the desired display resolution by setting a minimum of 2 up to a maximum of 8192 display points on the screen. This data format is well suited for transfer to a spreadsheet program on a computer.
- **Instrument state** - To keep a record of the spectrum analyzer settings, such as center frequency, span, reference level, and so on, used in the measurement. This is useful when documenting test setups used for making measurements. Consistent test setups are essential for maintaining repeatable measurements over time.

Most Agilent spectrum analyzers come with a copy of Agilent's *IntwiLink* software. This software lets you transfer instrument settings and trace data directly to a Microsoft® Excel spreadsheet or Word document.

Data transfer and remote instrument control

In 1977, Agilent Technologies (part of Hewlett-Packard at that time) introduced the world's first GPIB-controllable spectrum analyzer, the 8568A. The GPIB interface (also known as HP-IB or IEEE-488) made it possible to control all major functions of the analyzer and transfer trace data to an external computer. This innovation paved the way for a wide variety of automated spectrum analyzer measurements that were faster and more repeatable than manual measurements. By transferring the raw data to a computer, it could be saved on disk, analyzed, corrected, and operated on in a variety of ways.

Today, automated test and measurement equipment has become the norm, and nearly all modern spectrum analyzers come with a variety of standard interfaces. The most common one remains GPIB, but in recent years, Ethernet LAN connectivity has become increasingly popular, as it can provide high data transfer rates over long distances and integrates easily into networked environments such as a factory floor. Other standard interfaces used widely in the computer industry are likely to become available on spectrum analyzers in the future to simplify connectivity between instrument and computer.

A variety of commercial software products are available to control spectrum analyzers remotely over an I/O bus. You can also write your own software to control spectrum analyzers in a number of different ways. One method is to directly send programming commands to the instrument. Older spectrum analyzers typically used proprietary command sets, but newer instruments, such as Agilent's ESA and PSA spectrum analyzers, use industry-standard SCPI (standard commands for programmable instrumentation) commands. A more common method is to use standard software drivers, such as *VXIplug&play* drivers, which enable higher-level functional commands to the instrument without the need for detailed knowledge of the SCPI commands. Most recently, a new generation of language-independent instrument drivers, known as "interchangeable virtual instrument," or IVI-COM drivers, has become available for the ESA and PSA families. The IVI-COM drivers are based on the Microsoft Component Object Model standard and work in a variety of PC application development environments, such as the Agilent T&M Programmers Toolkit and Microsoft's Visual Studio .NET.

Some applications require that you control the spectrum analyzer and collect measurement data from a very long distance. For example, you may want to monitor satellite signals from a central control room, collecting data from remote tracking stations located hundreds or even thousands of kilometers away from the central site. The ESA and PSA Series spectrum analyzers have software options available to control these units, capture screen images, and transfer trace data over the Internet using a standard Web browser.

Firmware updates

Modern spectrum analyzers have much more software inside them than do instruments from just a few years ago. As new features are added to the software and defects repaired, it becomes highly desirable to update the spectrum analyzer's firmware to take advantage of the improved performance.

The latest revisions of spectrum analyzer firmware can be found on the Agilent Technologies website. This firmware can be downloaded to a file on a local computer. A common method to transfer new firmware into the spectrum analyzer is to copy the firmware onto several floppy disks that are then inserted into the spectrum analyzer's floppy disk drive. Some models, such as the PSA Series, allow you to transfer the new firmware directly into the spectrum analyzer using the Ethernet LAN port.

It is a good practice to periodically check your spectrum analyzer model's Web page to see if updated firmware is available.

Calibration, troubleshooting, diagnostics, and repair

Spectrum analyzers must be periodically calibrated to insure that the instrument performance meets all published specifications. Typically, this is done once a year. However, in between these annual calibrations, the spectrum analyzer must be aligned periodically to compensate for thermal drift and aging effects. Modern spectrum analyzers such as the ESA and PSA Series have built-in alignment routines that operate when the instrument is first turned on, and during retrace at predetermined intervals, or if the internal temperature of the instrument changes. These alignment routines continuously adjust the instrument to maintain specified performance. In the past, spectrum analyzers normally had to be turned on in a stable temperature environment for at least thirty minutes in order for the instrument to meet its published specifications. The auto-alignment capability makes it possible for the ESA and PSA spectrum analyzers to meet published specifications within five minutes.

Modern spectrum analyzers usually have a service menu available. In this area, you can perform useful diagnostic functions, such as a test of the front panel keys. You can also display more details of the alignment process, as well as a list of all optional hardware and measurement personalities installed in the instrument. When a spectrum analyzer is upgraded with a new measurement personality, Agilent provides a unique license key tied to the serial number of the instrument. This license key is entered through the front panel keypad to activate the measurement capabilities of the personality.

Summary

The objective of this application note is to provide a broad survey of basic spectrum analyzer concepts. However, you may wish to learn more about many other topics related to spectrum analysis. An excellent place to start is to visit the Agilent Technologies Web site at www.Agilent.com and search for “spectrum analyzer.”



Agilent Technologies Signal Analysis Division would like to dedicate this application note to Blake Peterson, who recently retired after more than 46 years of outstanding service in engineering applications and technical education for Agilent customers and employees. One of Blake's many accomplishments includes being the author of the previous editions of Application Note 150. To our friend and mentor, we wish you all the best for a happy and fulfilling retirement!

Glossary of Terms

Absolute amplitude accuracy: The uncertainty of an amplitude measurement in absolute terms, either volts or power. Includes relative uncertainties (see Relative amplitude accuracy) plus calibrator uncertainty. For improved accuracy, some spectrum analyzers have frequency response specified relative to the calibrator as well as relative to the mid-point between peak-to-peak extremes.

ACPR: Adjacent channel power ratio is a measure of how much signal energy from one communication channel spills over, or leaks into an adjacent channel. This is an important metric in digital communication components and systems, as too much leakage will cause interference on adjacent channels. It is sometimes also described as ACLR, or adjacent channel leakage ratio.

Amplitude accuracy: The uncertainty of an amplitude measurement. It can be expressed either as an absolute term or relative to another reference point.

Amplitude reference signal: A signal of precise frequency and amplitude that the analyzer uses for self-calibration.

Analog display: The technique in which analog signal information (from the envelope detector) is written directly to the display, typically implemented on a cathode ray tube (CRT). Analog displays were once the standard method of displaying information on a spectrum analyzer. However, modern spectrum analyzers no longer use this technique, but instead, use digital displays.

Average detection: A method of detection that sums power across a frequency interval. It is often used for measuring complex, digitally modulated signals and other types of signals with noise-like characteristics. Modern Agilent spectrum analyzers typically offer three types of average detection: power (rms) averaging, which measures the true average power over a bucket interval; voltage averaging, which measures the average voltage data over a bucket interval; and log-power (video) averaging, which measures the logarithmic amplitude in dB of the envelope of the signal during the bucket interval.

Average noise level: See Displayed average noise level.

Bandwidth selectivity: A measure of an analyzer's ability to resolve signals unequal in amplitude. Also called shape factor, bandwidth selectivity is the ratio of the 60 dB bandwidth to the 3 dB bandwidth for a given resolution (IF) filter. For some analyzers, the 6 dB bandwidth is used in lieu of the 3 dB bandwidth. In either case, bandwidth selectivity tells us how steep the filter skirts are.

Blocking capacitor: A filter that keeps unwanted low frequency signals (including DC) from damaging circuitry. A blocking capacitor limits the lowest frequency that can be measured accurately.

CDMA: Code division multiple access is a method of digital communication in which multiple communication streams are orthogonally coded, enabling them to share a common frequency channel. It is a popular technique used in a number of widely used mobile communication systems.

Constellation diagram: A display type commonly used when analyzing digitally modulated signals in which the detected symbol points are plotted on an IQ graph.

Delta marker: A mode in which a fixed, reference marker has been established and a second, active marker is available that we can place anywhere on the displayed trace. A read out indicates the relative frequency separation and amplitude difference between the reference marker and the active marker.

Digital display: A technique in which digitized trace information, stored in memory, is displayed on the screen. The displayed trace is a series of points designed to present a continuous looking trace. While the default number of display points varies between different models, most modern spectrum analyzers allow the user to choose the desired resolution by controlling the number of points displayed. The display is refreshed (rewritten from data in memory) at a flicker-free rate; the data in memory is updated at the sweep rate. Nearly all modern spectrum analyzers have digital flat panel LCD displays, rather than CRT-based analog displays that were used in earlier analyzers.

Display detector mode: The manner in which the signal information is processed prior to being displayed on screen. See Neg peak, Pos peak, Normal and Sample.

Digital IF: An architecture found in modern spectrum analyzers in which the signal is digitized soon after it has been downconverted from an RF frequency to an intermediate frequency (IF). At that point, all further signal processing is done using digital signal processing (DSP) techniques.

Display dynamic range: The maximum dynamic range for which both the larger and smaller signal may be viewed simultaneously on the spectrum analyzer display. For analyzers with a maximum logarithmic display of 10 dB/div, the actual dynamic range (see Dynamic range) may be greater than the display dynamic range.

Display scale fidelity: The uncertainty in measuring relative differences in amplitude on a spectrum analyzer. The logarithmic and linear IF amplifiers found in analyzers with analog IF sections never have perfect logarithmic or linear responses, and thus introduce uncertainty. Modern analyzers with digital IF sections have significantly better display scale fidelity.

Display range: The calibrated range of the display for the particular display mode and scale factor. See Linear and Log display and Scale factor.

Displayed average noise level: The noise level as seen on the analyzer's display after setting the video bandwidth narrow enough to reduce the peak-to-peak noise fluctuations such that the displayed noise is essentially seen as a straight line. Usually refers to the analyzer's own internally generated noise as a measure of sensitivity and is typically specified in dBm under conditions of minimum resolution bandwidth and minimum input attenuation.

Drift: The very slow (relative to sweep time) change of signal position on the display as a result of a change in LO frequency versus sweep voltage. The primary sources of drift are the temperature stability and aging rate of the frequency reference in the spectrum analyzer.

Dynamic range: The ratio, in dB, between the largest and smallest signals simultaneously present at the spectrum analyzer input that can be measured to a given degree of accuracy. Dynamic range generally refers to measurement of distortion or intermodulation products.

Envelope detector: A circuit element whose output follows the envelope, but not the instantaneous variation, of its input signal. In a superheterodyne spectrum analyzer, the input to the envelope detector comes from the final IF, and the output is a video signal. When we put our analyzer in zero span, the envelope detector demodulates the input signal, and we can observe the modulating signal as a function of time on the display.

Error vector magnitude (EVM): A quality metric in digital communication systems. EVM is the magnitude of the vector difference at a given instant in time between the ideal reference signal and the measured signal.

External mixer: An independent mixer, usually with a waveguide input port, used to extend the frequency range of those spectrum analyzers designed to utilize external mixers. The analyzer provides the LO signal and, if needed, mixer bias. Mixing products are returned to the analyzer's IF input.

FFT (fast Fourier transform): A mathematical operation performed on a time-domain signal to yield the individual spectral components that constitute the signal. See Spectrum.

Flatness: See Frequency response.

Frequency accuracy: The uncertainty with which the frequency of a signal or spectral component is indicated, either in an absolute sense or relative to some other signal or spectral component. Absolute and relative frequency accuracies are specified independently.

Frequency range: The minimum to maximum frequencies over which a spectrum analyzer can tune. While the maximum frequency is generally thought of in terms of an analyzer's coaxial input, the range of many microwave analyzers can be extended through use of external waveguide mixers.

Frequency resolution: The ability of a spectrum analyzer to separate closely spaced spectral components and display them individually. Resolution of equal amplitude components is determined by resolution bandwidth. The ability to resolve unequal amplitude signals is a function of both resolution bandwidth and bandwidth selectivity.

Frequency response: Variation in the displayed amplitude of a signal as a function of frequency (flatness). Typically specified in terms of \pm dB relative to the value midway between the extremes. Also may be specified relative to the calibrator signal.

Frequency span: The frequency range represented by the horizontal axis of the display. Generally, frequency span is given as the total span across the full display. Some earlier analyzers indicate frequency span (scan width) on a per-division basis.

Frequency stability: A general phrase that covers both short- and long-term LO instability. The sweep ramp that tunes the LO also determines where a signal should appear on the display. Any long term variation in LO frequency (drift) with respect to the sweep ramp causes a signal to slowly shift its horizontal position on the display. Shorter term LO instability can appear as random FM or phase noise on an otherwise stable signal.

Full span: For most modern spectrum analyzers, full span means a frequency span that covers the entire tuning range of the analyzer. These analyzers include single band RF analyzers and microwave analyzers such as the ESA and PSA Series that use a solid-state switch to switch between the low and preselected ranges.

NOTE: On some earlier spectrum analyzers, full span referred to a sub-range. For example, with the Agilent 8566B, a microwave spectrum analyzer that used a mechanical switch to switch between the low and preselected ranges, full span referred to either the low, non-preselected range or the high, preselected range.

Gain compression: That signal level at the input mixer of a spectrum analyzer at which the displayed amplitude of the signal is a specified number of dB too low due just to mixer saturation. The signal level is generally specified for 1 dB compression, and is usually between +3 and -10 dBm, depending on the model of spectrum analyzer.

GSM: The global system for mobile communication is a widely used digital standard for mobile communication. It is a TDMA-based system in which multiple communication streams are interleaved in time, enabling them to share a common frequency channel.

Harmonic distortion: Unwanted frequency components added to a signal as the result of the nonlinear behavior of the device (e.g. mixer, amplifier) through which the signal passes. These unwanted components are harmonically related to the original signal.

Harmonic mixing: The utilization of the LO harmonics generated in a mixer to extend the tuning range of a spectrum analyzer beyond the range achievable using just the LO fundamental.

IF gain/IF attenuation: Adjusts the vertical position of signals on the display without affecting the signal level at the input mixer. When changed, the value of the reference level is changed accordingly.

IF feedthrough: A raising of the baseline trace on the display due to an input signal at the intermediate frequency passing through the input mixer. Generally, this is a potential problem only on non-preselected spectrum analyzers. The entire trace is raised because the signal is always at the IF, i.e. mixing with the LO is not required.

Image frequencies: Two or more real signals present at the spectrum analyzer input that produce an IF response at the same LO frequency. Because the mixing products all occur at the same LO and IF frequencies, it is impossible to distinguish between them.

Image response: A displayed signal that is actually twice the IF away from the frequency indicated by the spectrum analyzer. For each harmonic of the LO, there is an image pair, one below and one above the LO frequency by the IF. Images usually appear only on non-preselected spectrum analyzers.

Incidental FM: Unwanted frequency modulation on the output of a device (signal source, amplifier) caused by (incidental to) some other form of modulation, e.g. amplitude modulation.

Input attenuator: A step attenuator between the input connector and first mixer of a spectrum analyzer. Also called the RF attenuator. The input attenuator is used to adjust level of the signal incident upon the first mixer. The attenuator is used to prevent gain compression due to high-level and/or broadband signals and to set dynamic range by controlling the degree of internally generated distortion. In some analyzers, the vertical position of displayed signals is changed when the input attenuator setting is changed, so the reference level is also changed accordingly. In modern Agilent analyzers, the IF gain is changed to compensate for input attenuator changes, so signals remain stationary on the display, and the reference level is not changed.

Input impedance: The terminating impedance that the analyzer presents to the signal source. The nominal impedance for RF and microwave analyzers is usually 50 ohms. For some systems, e.g. cable TV, 75 ohms is standard. The degree of mismatch between the nominal and actual input impedance is given in terms of VSWR (voltage standing wave ratio).

Intermodulation distortion: Unwanted frequency components resulting from the interaction of two or more spectral components passing through a device with non-linear behavior (e.g. mixer, amplifier). The unwanted components are related to the fundamental components by sums and differences of the fundamentals and various harmonics, e.g. $f_1 \pm f_2$, $2f_1 \pm f_2$, $2f_2 \pm f_1$, $3f_1 \pm 2f_2$, and so forth.

Linear display: The display mode in which vertical deflection on the display is directly proportional to the voltage of the input signal. The bottom line of the graticule represents 0 V, and the top line, the reference level, some non-zero value that depends upon the particular spectrum analyzer. On most modern analyzers, we select the reference level, and the scale factor becomes the reference level value divided by the number of graticule divisions. Although the display is linear, modern analyzers allow reference level and marker values to be indicated in dBm, dBmV, dBuV, and in some cases, watts as well as volts.

LO emission or feedout: The emergence of the LO signal from the input of a spectrum analyzer. The level can be greater than 0 dBm on non-preselected spectrum analyzers but is usually less than -70 dBm on preselected analyzers.

LO feedthrough: The response on the display when a spectrum analyzer is tuned to 0 Hz, i.e. when the LO is tuned to the IF. The LO feedthrough can be used as a 0-Hz marker, and there is no frequency error.

Log display: The display mode in which vertical deflection on the display is a logarithmic function of the voltage of the input signal. We set the display calibration by selecting the value of the top line of the graticule, the reference level, and scale factor in dB/div. On Agilent analyzers, the bottom line of the graticule represents zero volts for scale factors of 10 dB/div or more, so the bottom division is not calibrated in these cases. Modern analyzers allow reference level and marker values to be indicated in dBm, dBmV, dBuV, volts, and in some cases, watts. Earlier analyzers usually offered only one choice of units, and dBm was the usual choice.

Marker: A visible indicator that we can place anywhere along the displayed signal trace. A read out indicates the absolute value of both the frequency and amplitude of the trace at the marked point. The amplitude value is given in the currently selected units. Also see Delta marker and Noise marker.

Measurement range: The ratio, expressed in dB, of the maximum signal level that can be measured (usually the maximum safe input level) to the lowest achievable average noise level. This ratio is almost always much greater than can be realized in a single measurement. See Dynamic range.

Mixing mode: A description of the particular circumstance that creates a given response on a spectrum analyzer. The mixing mode, e.g. 1⁺, indicates the harmonic of the LO used in the mixing process and whether the input signal is above (+) or below (-) that harmonic.

Multiple responses: Two or more responses on a spectrum analyzer display from a single input signal. Multiple responses occur only when mixing modes overlap and the LO is swept over a wide enough range to allow the input signal to mix on more than one mixing mode. Normally not encountered in analyzers with preselectors.

Negative peak: The display detection mode in which each displayed point indicates the minimum value of the video signal for that part of the frequency span and/or time interval represented by the point.

Noise figure: The ratio, usually expressed in dB, of the signal-to-noise ratio at the input of a device (mixer, amplifier) to the signal-to-noise ratio at the output of the device.

Noise marker: A marker whose value indicates the noise level in a 1 Hz noise power bandwidth. When the noise marker is selected, the sample display detection mode is activated, the values of a number of consecutive trace points (the number depends upon the analyzer) about the marker are averaged, and this average value is normalized to an equivalent value in a 1 Hz noise power bandwidth. The normalization process accounts for detection and bandwidth plus the effect of the log amplifier when we select the log display mode.

Noise sidebands: Modulation sidebands that indicate the short-term instability of the LO (primarily the first LO) system of a spectrum analyzer. The modulating signal is noise, in the LO circuit itself and/or in the LO stabilizing circuit, and the sidebands comprise a noise spectrum. The mixing process transfers any LO instability to the mixing products, so the noise sidebands appear on any spectral component displayed on the analyzer far enough above the broadband noise floor. Because the sidebands are noise, their level relative to a spectral component is a function of resolution bandwidth. Noise sidebands are typically specified in terms of dBc/Hz (amplitude in a 1 Hz bandwidth relative to the carrier) at a given offset from the carrier, the carrier being a spectral component viewed on the display.

Phase noise: See Noise sidebands.

Positive peak: The display detection mode in which each displayed point indicates the maximum value of the video signal for that part of the frequency span and/or time interval represented by the point.

Preamplifier: An external, low noise-figure amplifier that improves system (preamplifier/spectrum analyzer) sensitivity over that of the analyzer itself.

Preselector: A tunable bandpass filter that precedes the input mixer of a spectrum analyzer and tracks the appropriate mixing mode. Preselectors are typically used only above 2 GHz. They essentially eliminate multiple and image responses and, for certain signal conditions, improve dynamic range.

Quasi-peak detector (QPD): A type of detector whose output is a function of both signal amplitude as well as pulse repetition rate. The QPD gives higher weighting to signals with higher pulse repetition rates. In the limit, a QPD will exhibit the same amplitude as a peak detector when measuring a signal with a constant amplitude (CW) signal.

Raster display: A TV-like display in which the image is formed by scanning the electron beam rapidly across and slowly down the display face and gating the beam on as appropriate. The scanning rates are fast enough to produce a flicker-free display. Also see Vector display and Sweep time.

Reference level: The calibrated vertical position on the display used as a reference for amplitude measurements. The reference level position is normally the top line of the graticule.

Relative amplitude accuracy: The uncertainty of an amplitude measurement in which the amplitude of one signal is compared to the amplitude of another regardless of the absolute amplitude of either. Distortion measurements are relative measurements. Contributors to uncertainty include frequency response and display fidelity and changes of input attenuation, IF gain, scale factor, and resolution bandwidth.

Residual FM: The inherent short-term frequency instability of an oscillator in the absence of any other modulation. In the case of a spectrum analyzer, we usually expand the definition to include the case in which the LO is swept. Residual FM is usually specified in peak-to-peak values because they are most easily measured on the display, if visible at all.

Residual responses: Discrete responses seen on a spectrum analyzer display with no input signal present.

Resolution: See Frequency resolution.

Resolution bandwidth: The width of the resolution bandwidth (IF) filter of a spectrum analyzer at some level below the minimum insertion loss point (maximum deflection point on the display). For Agilent analyzers, the 3 dB bandwidth is specified; for some others, it is the 6 dB bandwidth.

Rosenfell: The display detection mode in which the value displayed at each point is based upon whether or not the video signal both rose and fell during the frequency and/or time interval represented by the point. If the video signal only rose or only fell, the maximum value is displayed. If the video signal did both rise and fall, then the maximum value during the interval is displayed by odd-numbered points, the minimum value, by even-numbered points. To prevent the loss of a signal that occurs only in an even-numbered interval, the maximum value during this interval is preserved, and in the next (odd-numbered) interval, the displayed value is the greater of either the value carried over or the maximum that occurs in the current interval.

Sample: The display detection mode in which the value displayed at each point is the instantaneous value of the video signal at the end of the frequency span and/or time interval represented by the point.

Scale factor: The per-division calibration of the vertical axis of the display.

Sensitivity: The level of the smallest sinusoid that can be observed on a spectrum analyzer, usually under optimized conditions of minimum resolution bandwidth, 0 dB RF input attenuation, and minimum video bandwidth. Agilent defines sensitivity as the displayed average noise level. A sinusoid at that level will appear to be about 2 dB above the noise.

Shape factor: See Bandwidth selectivity.

Signal identification: A routine, either manual or automatic, that indicates whether or not a particular response on the spectrum analyzer's display is from the mixing mode for which the display is calibrated. If automatic, the routine may change the analyzer's tuning to show the signal on the correct mixing mode, or it may tell us the signal's frequency and give us the option of ignoring the signal or having the analyzer tune itself properly for the signal. Generally not needed on preselected analyzers.

Span accuracy: The uncertainty of the indicated frequency separation of any two signals on the display.

Spectral purity: See Noise sidebands.

Spectral component: One of the sine waves comprising a spectrum.

Spectrum: An array of sine waves of differing frequencies and amplitudes and properly related with respect to phase that, taken as a whole, constitute a particular time-domain signal.

Spectrum analyzer: A device that effectively performs a Fourier transform and displays the individual spectral components (sine waves) that constitute a time-domain signal. Phase may or may not be preserved, depending upon the analyzer type and design.

Spurious responses: The improper responses that appear on a spectrum analyzer display as a result of the input signal. Internally generated distortion products are spurious responses, as are image and multiple responses.

Sweep time: The time to tune the LO across the selected span. Sweep time does not include the dead time between the completion of one sweep and the start of the next. In zero span, the spectrum analyzer's LO is fixed, so the horizontal axis of the display is calibrated in time only. In non-zero spans, the horizontal axis is calibrated in both frequency and time, and sweep time is usually a function of frequency span, resolution bandwidth, and video bandwidth.

Time gating: A method of controlling the frequency sweep of the spectrum analyzer based on the characteristics of the signal being measured. It is often useful when analyzing pulsed or burst modulated signals; time multiplexed signals, as well as intermittent signals.

TDMA: Time division multiple access is a digital communication method in which multiple communication streams are interleaved in time, enabling them to share a common frequency channel.

Units: Dimensions of the measured quantities. Units usually refer to amplitude quantities because they can be changed. In modern spectrum analyzers, available units are dBm (dB relative to 1 milliwatt dissipated in the nominal input impedance of the analyzer), dBmV (dB relative to 1 millivolt), dBuV (dB relative to 1 microvolt), volts, and in some analyzers, watts. In Agilent analyzers, we can specify any units in both log and linear displays.

Vector diagram: A display type commonly used when analyzing digitally modulated signals. It is similar to a constellation display, except that in addition to the detected symbol points, the instantaneous power levels during state transitions are also plotted on an IQ graph.

Vector display: A display type used in earlier spectrum analyzer designs, in which the electron beam was directed so that the image (trace, graticule, annotation) was written directly on the CRT face, not created from a series of dots as in the raster displays commonly used today.

Video: In a spectrum analyzer, a term describing the output of the envelope detector. The frequency range extends from 0 Hz to a frequency typically well beyond the widest resolution bandwidth available in the analyzer. However, the ultimate bandwidth of the video chain is determined by the setting of the video filter.

Video amplifier: A post-detection, DC-coupled amplifier that drives the vertical deflection plates of the CRT. See Video bandwidth and Video filter.

Video average: A digital averaging of a spectrum analyzer's trace information. The averaging is done at each point of the display independently and is completed over the number of sweeps selected by the user. The averaging algorithm applies a weighting factor ($1/n$, where n is the number of the current sweep) to the amplitude value of a given point on the current sweep, applies another weighting factor $[(n - 1)/n]$ to the previously stored average, and combines the two for a current average. After the designated number of sweeps are completed, the weighting factors remain constant, and the display becomes a running average.

Video bandwidth: The cutoff frequency (3 dB point) of an adjustable low pass filter in the video circuit. When the video bandwidth is equal to or less than the resolution bandwidth, the video circuit cannot fully respond to the more rapid fluctuations of the output of the envelope detector. The result is a smoothing of the trace, i.e. a reduction in the peak-to-peak excursion of broadband signals such as noise and pulsed RF when viewed in the broadband mode. The degree of averaging or smoothing is a function of the ratio of the video bandwidth to the resolution bandwidth.

Video filter: A post-detection, low-pass filter that determines the bandwidth of the video amplifier. Used to average or smooth a trace. See Video bandwidth.

Zero span: That case in which a spectrum analyzer's LO remains fixed at a given frequency so the analyzer becomes a fixed-tuned receiver. The bandwidth of the receiver is that of the resolution (IF) bandwidth. Signal amplitude variations are displayed as a function of time. To avoid any loss of signal information, the resolution bandwidth must be as wide as the signal bandwidth. To avoid any smoothing, the video bandwidth must be set wider than the resolution bandwidth.

Agilent Technologies' Test and Measurement Support, Services, and Assistance

Agilent Technologies aims to maximize the value you receive, while minimizing your risk and problems. We strive to ensure that you get the test and measurement capabilities you paid for and obtain the support you need. Our extensive support resources and services can help you choose the right Agilent products for your applications and apply them successfully. Every instrument and system we sell has a global warranty. Support is available for at least five years beyond the production life of the product. Two concepts underlie Agilent's overall support policy: "Our Promise" and "Your Advantage."

Our Promise

Our Promise means your Agilent test and measurement equipment will meet its advertised performance and functionality. When you are choosing new equipment, we will help you with product information, including realistic performance specifications and practical recommendations from experienced test engineers. When you use Agilent equipment, we can verify that it works properly, help with product operation, and provide basic measurement assistance for the use of specified capabilities, at no extra cost upon request. Many self-help tools are available.

Your Advantage

Your Advantage means that Agilent offers a wide range of additional expert test and measurement services, which you can purchase according to your unique technical and business needs. Solve problems efficiently and gain a competitive edge by contracting with us for calibration, extra-cost upgrades, out-of-warranty repairs, and on-site education and training, as well as design, system integration, project management, and other professional engineering services. Experienced Agilent engineers and technicians worldwide can help you maximize your productivity, optimize the return on investment of your Agilent instruments and systems, and obtain dependable measurement accuracy for the life of those products.



Agilent Email Updates

www.agilent.com/find/emailupdates

Get the latest information on the products and applications you select.

Agilent T&M Software and Connectivity

Agilent's Test and Measurement software and connectivity products, solutions and developer network allows you to take time out of connecting your instruments to your computer with tools based on PC standards, so you can focus on your tasks, not on your connections. Visit www.agilent.com/find/connectivity for more information.

By internet, phone, or fax, get assistance with all your test & measurement needs

Phone or Fax

United States:

(tel) 800 829 4444

Canada:

(tel) 877 894 4414

(fax) 905 282 6495

China:

(tel) 800 810 0189

(fax) 800 820 2816

Europe:

(tel) (31 20) 547 2323

(fax) (31 20) 547 2390

Japan:

(tel) (81) 426 56 7832

(fax) (81) 426 56 7840

Korea:

(tel) (82 2) 2004 5004

(fax) (82 2) 2004 5115

Latin America:

(tel) (305) 269 7500

(fax) (305) 269 7599

Taiwan:

(tel) 0800 047 866

(fax) 0800 286 331

Other Asia Pacific Countries:

(tel) (65) 6375 8100

(fax) (65) 6836 0252

Email: tm_asia@agilent.com

Online Assistance:

www.agilent.com/find/assist

Product specifications and descriptions in this document subject to change without notice.

© Agilent Technologies, Inc. 2004

Printed in U.S.A., April 27, 2004

5952-0292



Agilent Technologies