

Spectrum Analyzer Measurements and Noise

Application Note 1303

Measuring Noise and Noise-like Digital Communications Signals with a Spectrum Analyzer




Table of contents

Part I: Noise measurements	3
Introduction	3
Simple noise—Baseband, Real, Gaussian.....	3
Bandpassed noise—I and Q	3
Measuring the power of noise with an envelope detector	6
Logarithmic processing	7
Measuring the power of noise with a log-envelope scale.....	8
Equivalent noise bandwidth	8
The noise marker.....	9
<i>Spectrum analyzers and envelope detectors</i>	<i>10</i>
<i>Cautions when measuring noise with spectrum analyzers.....</i>	<i>12</i>
Part II: Measurements of noise-like signals	14
The noise-like nature of digital signals	14
Channel-power measurements	14
Adjacent-Channel Power (ACP).....	16
Carrier power.....	16
<i>Peak-detected noise (and TDMA ACP measurements)</i>	<i>18</i>
Part III: Averaging and the noisiness of noise measurements	19
Variance and averaging.....	19
Averaging a number of computed results.....	20
Swept versus FFT analysis	20
Zero span.....	20
The standard deviation of measurement noise.....	21
Examples.....	22
Part IV: Compensation for instrumentation noise	23
CW signals and log versus power detection	23
Power-detection measurements and noise subtraction	24
<i>Log scale ideal for CW measurements.....</i>	<i>25</i>
Bibliography	27
Glossary of terms.....	28

Part I: Noise measurements

Introduction


Noise. It is the classical limitation of electronics. In measurements, noise and distortions limit the dynamic range of test results.

In this four-part paper, the characteristics of noise and its direct measurement are discussed in Part I. Part II contains a discussion of the measurement of noise-like signals exemplified by digital CDMA and TDMA signals. Part III discusses using averaging techniques to reduce noise. Part IV is about compensating for the noise in instrumentation while measuring CW (sinusoidal) and noise-like signals.

Simple noise—Baseband, Real, Gaussian

Noise occurs due to the random motion of electrons. The number of electrons involved is large, and their motions are independent. Therefore, the variation in the rate of current flow takes on a bell-shaped curve known as the Gaussian Probability Density Function (PDF) in accordance with the central limit theorem from statistics. The Gaussian PDF is shown in Figure 1.

Figure 1. The Gaussian PDF is maximum at zero current and falls off away from zero, as shown (rotated 90 degrees) on the left. A typical noise waveform is shown on the right.



The Gaussian PDF explains some of the characteristics of a noise signal seen on a baseband instrument such as an oscilloscope. The baseband signal is a real signal; it has no imaginary components.


Bandpassed noise—I and Q

In RF design work and when using spectrum analyzers, we usually deal with signals within a passband, such as a communications channel or the resolution bandwidth (RBW, the bandwidth of the final IF) of a spectrum analyzer. Noise in this bandwidth still has a Gaussian PDF, but few RF instruments display PDF-related metrics.

Instead, we deal with a signal's magnitude and phase (polar coordinates) or I/Q components. The latter are the in-phase (I) and quadrature (Q) parts of a signal, or the real and imaginary components of a rectangular-coordinate representation of a signal. Basic (scalar) spectrum analyzers measure only the magnitude of a signal. We are interested in the characteristics of the magnitude of a noise signal.

We can consider the noise within a passband as being made of independent I and Q components, each with Gaussian PDFs. Figure 2 shows samples of I and Q components of noise represented in the I/Q plane. The signal in the passband is actually given by the sum of the I magnitude, v_I , multiplied by a cosine wave (at the center frequency of the passband) and the Q magnitude, v_Q , multiplied by a sine wave. But we can discuss just the I and Q components without the complications of the sine/cosine waves.

Figure 2: Bandpassed noise has a Gaussian PDF independently in both its I and Q components.




Spectrum analyzers respond to the magnitude of the signal within their RBW passband. The magnitude, or envelope, of a signal represented by an I/Q pair is given by:

$$v_{env} = \sqrt{(v_I^2 + v_Q^2)}$$

Graphically, the envelope is the length of the vector from the origin to the I/Q pair. It is instructive to draw circles of evenly spaced constant-amplitude envelopes on the samples of I/Q pairs as shown in Figure 3.

Figure 3: Samples of I/Q pairs shown with evenly spaced constant-amplitude envelope circles.




If one were to count the number of samples within each annular ring in Figure 3, we would see that the area near zero volts does not have the highest count of samples. Even though the density of samples is highest there, this area is smaller than any of the other rings.

The count within each ring constitutes a histogram of the distribution of the envelope. If the width of the rings were reduced and expressed as the “count” per unit of ring width, the limit becomes a continuous function instead of a histogram. This continuous function is the PDF of the envelope of bandpassed noise. It is a Rayleigh distribution in the envelope voltage, v , that depends on the sigma of the signal; for $v \geq 0$:

$$\text{PDF}(v) = \left(\frac{v}{\sigma^2}\right) \exp\left(-\frac{1}{2}\left(\frac{v}{\sigma}\right)^2\right)$$

The Rayleigh distribution is shown in Figure 4.

Figure 4: The PDF of the voltage of the envelope of a noise signal is a Rayleigh distribution. The PDF is zero at zero volts, even though the PDFs of the individual I and Q components are maximum at zero volts. It is maximum for $v = \text{sigma}$.



Measuring the power of noise with an envelope detector

The power of the noise is the parameter we usually want to measure with a spectrum analyzer. The power is the “heating value” of the signal. Mathematically, it is the average of v^2/R , where R is the impedance of the signal and v is its instantaneous voltage.

At first glance, we might like to find the average envelope voltage and square it, then divide by R . But finding the *square of the average* is not the same as finding the *average of the square*. In fact, there is a consistent under-measurement of noise from squaring the average instead of averaging the square; this under-measurement is 1.05 dB.

The average envelope voltage is given by integrating the product of the envelope voltage and the probability that the envelope takes on that voltage. This probability is the Rayleigh PDF, so:

$$\bar{v} = \int_0^{\infty} v \text{PDF}(v) dv = \sigma \sqrt{\frac{\pi}{2}}$$

The average power of the signal is given by an analogous expression with v^2/R in place of the “ v ” part:

$$\bar{p} = \int_0^{\infty} \left(\frac{v^2}{R}\right) \text{PDF}(v) dv = \frac{2\sigma^2}{R}$$

We can compare the true power, from the average power integral, with the voltage-envelope-detected estimate of v^2/R and find the ratio to be 1.05 dB, independent of σ and R .

$$10 \log\left(\frac{\bar{v}^2/R}{\bar{p}}\right) = 10 \log\left(\frac{\pi}{4}\right) = -1.05 \text{ dB}$$

Thus, if we were to measure noise with a spectrum analyzer using voltage-envelope detection (the “linear” scale) and averaging, an additional 1.05 dB would need to be added to the result to compensate for averaging voltage instead of voltage-squared.

Logarithmic processing

Spectrum Analyzers are most commonly used in their logarithmic (“log”) display mode, in which the vertical axis is calibrated in decibels. Let us look again at our PDF for the voltage envelope of a noise signal, but let’s mark the x -axis with points equally spaced on a decibel scale, in this case with 1 dB spacing. See Figure 5. The area under the curve between markings is the probability that the log of the envelope voltage will be within that 1 dB interval. Figure 6 represents the continuous PDF of a logged signal which we predict from the areas in Figure 5.

Figure 5: The PDF of the voltage envelope of noise is graphed. 1-dB spaced marks on the x -axis shows how the probability density would be different on a log scale. Where the decibel markings are dense, the probability that the noise will fall between adjacent marks is reduced.





Figure 6: The PDF of logged noise is about 30 dB wide and tilted toward the high end.



Measuring the power of noise with a log-envelope scale

When a spectrum analyzer is in a log (dB) display mode, averaging of the results can occur in numerous ways. Multiple traces can be averaged, the envelope can be averaged by the action of the video filter, or the noise marker (more on this below) averages results across the x -axis.

When we express the average power of the noise in decibels, we compute a logarithm of that average power. When we average the output of the log scale of a spectrum analyzer, we compute the average of the log. The *log of the average* is not equal to the *average of the log*. If we go through the same kinds of computations that we did comparing average voltage envelopes with average power envelopes, we find that log processing causes an under-response to noise of 2.51 dB, rather than 1.05 dB.¹

The log amplification acts as a compressor for large noise peaks; a peak of ten times the average level is only 10 dB higher. Instantaneous near-zero envelopes, on the other hand, contain no power but are expanded toward negative infinity decibels. The combination of these two aspects of the logarithmic curve cause noise power to be underestimated.

Equivalent noise bandwidth

Before discussing the measurement of noise with a spectrum analyzer “noise marker,” it is necessary to understand the RBW filter of a spectrum analyzer.


The ideal RBW has a flat passband and infinite attenuation outside that passband. But it must also have good time domain performance so that it behaves well when signals sweep through the passband. Most spectrum analyzers use four-pole synchronously tuned filters for their RBW filters. We can plot the power gain (the square of the voltage gain) of the RBW filter versus frequency as shown in Figure 7. The response of the filter to noise of flat power spectral density will be the same as the response of a rectangular filter with the same maximum gain and the same area under their curves. The width of such a rectangular filter is the “equivalent noise bandwidth” of the RBW filter. The noise density at the input to the RBW filter is given by the output power divided by the equivalent noise bandwidth.

1. Most authors on this subject artificially state that this factor is due to 1.05 dB from envelope detection and another 1.45 dB from logarithmic amplification, reasoning that the signal is first voltage-envelope detected, then logarithmically amplified. But if we were to measure the voltage-squared envelope (in other words, the power envelope, which would cause zero error instead of 1.05 dB) and then log it, we would still find a 2.51 dB under-response. Therefore, there is no real point in separating the 2.51 dB into two pieces.

The ratio of the equivalent noise bandwidth to the -3 dB bandwidth (the “name” of the RBW is usually its -3 dB BW) is given by the following table:

Filter type	Application	NBW/ -3 dB BW
4-pole sync	Most SAs analog	1.128 (0.52 dB)
5-pole sync	Some SAs analog	1.111 (0.46 dB)
Typical FFT	FFT-based SAs	1.05 (0.23 dB)

Figure 7: The power gain versus frequency of an RBW filter can be modeled by a rectangular filter with the same area and peak level, and a width of the “equivalent noise bandwidth.”



The noise marker

As discussed above, the measured level at the output of a spectrum analyzer must be manipulated in order to represent the input spectral noise density we wish to measure. This manipulation involves three factors, which may be added in decibel units:

1. Under-response due to voltage envelope detection (add 1.05 dB) or log-scale response (add 2.51 dB).
2. Over-response due to the ratio of the equivalent noise bandwidth to the -3 dB bandwidth (subtract typically 0.52 dB).
3. Normalization to a 1 Hz bandwidth (subtract 10 times the log of the RBW, where the RBW is given in units of Hz).

A further operation of the noise marker in HP spectrum analyzers is to average 32 measurement cells centered around the marker location in order to reduce the variance of the result.

The final result of these computations is a measure of the noise density, the noise in a theoretical ideal 1 Hz bandwidth. The units are typically dBm/Hz.

Spectrum analyzers and envelope detectors

A simplified block diagram of a spectrum analyzer is shown in Figure A.




Figure A: Simplified spectrum analyzer block diagram

The envelope detector/logarithmic amplifier block is shown configured as they are used in the HP 8560 E-Series spectrum analyzers. Although the order of these two circuits can be reversed, the important concept to recognize is that an IF signal goes into this block and a baseband signal (referred to as the “video” signal because it was used to deflect the electron beam in the original analog spectrum analyzers) comes out.

Notice that there is a second set of detectors in the block diagram: the peak/pt/sample hardware of what is normally called the “detector mode” of a spectrum analyzer. These “display detectors” are not relevant to this discussion, and should not be confused with the envelope detector.

The salient features of the envelope detector are two:

1. The output voltage is proportional to the input voltage envelope.
2. The bandwidth for following envelope variations is large compared to the widest RBW.

Figure B: Detectors: a) half-wave, b) full-wave implemented as a “product detector”, c) peak. Practical implementations usually have their gain terms implemented elsewhere, and implement buffering after the filters that remove the residual IF carrier and harmonics. The peak detector must be cleared; leakage through a resistor or a switch with appropriate timing are possible clearing mechanisms.

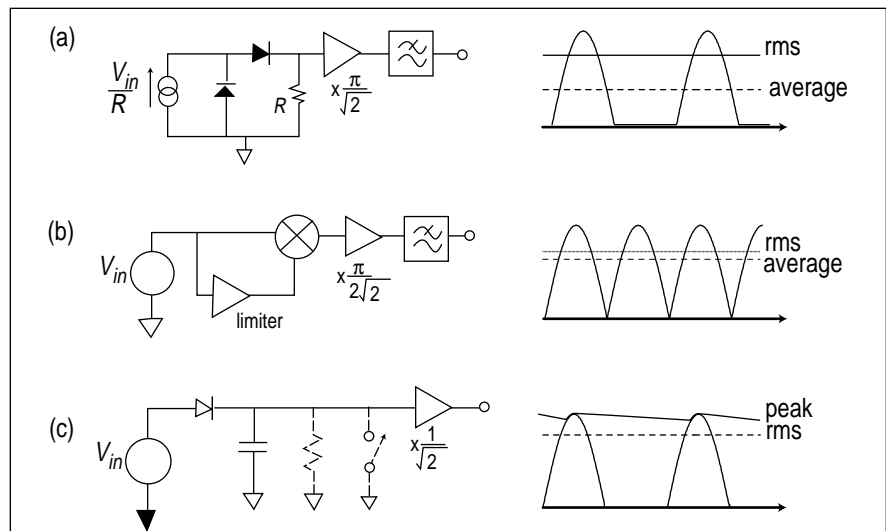



Figure B shows envelope detectors and their associated waveforms in (a) and (b). Notice that the gain required to make the average output voltage equal to the r.m.s. voltage of a sinusoidal input is different for the different topologies.

Some authors on this topic have stated that “an envelope detector is a peak detector.” After all, an idealized detector that responds to the peak of each cycle of IF energy independently makes an easy conceptual model of ideal behavior. But real peak detectors do not reset on each IF cycle. Figure B, part c shows a typical peak detector with its gain calibration factor. It is called a peak detector because its response is proportional to the peak voltage of the signal. If the signal is CW, a peak detector and an envelope detector act identically.

But if the signal has variations in its envelope, the envelope detector with the shown LPF (low pass filter) will follow those variations with the linear, time-domain characteristics of the filter; the peak detector will follow non-linearly, subject to its maximum negative-going dv/dt limit, as demonstrated in Figure C. The non-linearity will make for unpredictable behavior for signals with noise-like statistical variations.

Figure C: An envelope detector will follow the envelope of the shown signal, albeit with the delay and filtering action of the LPF used to remove the carrier harmonics. A peak detector is subject to negative slew limits, as demonstrated by the dashed line it will follow across a response pit. This drawing is done for the case in which the logarithmic amplification precedes the envelope detection, opposite to Figure A; in this case, the pits of the envelope are especially sharp.



A peak detector may act like an envelope detector in the limit as its resistive load dominates and the capacitive load is minimized. But practically, the non-ideal voltage drop across the diodes and the heavy required resistive load make this topology unsuitable for envelope detection. All spectrum analyzers use envelope detectors, some are just misnamed.


Cautions when measuring noise with spectrum analyzers

There are three ways in which noise measurements can look perfectly reasonable on the screen of a spectrum analyzer, yet be significantly in error.

Caution 1, input mixer level. A noise-like signal of very high amplitude can overdrive the front end of a spectrum analyzer while the displayed signal is within the normal display range. This problem is possible whenever the bandwidth of the noise-like signal is much wider than the RBW. The power within the RBW will be lower than the total power by about ten decibels times the log of the ratio of the signal bandwidth to the RBW. For example, an IS-95 CDMA signal with a 1.23 MHz bandwidth is 31 dB larger than the power in a 1 kHz RBW. If the indicated power with the 1 kHz RBW is -20 dBm at the input mixer (i.e., after the input attenuator), then the mixer is seeing about $+11$ dBm. Most spectrum analyzers are specified for -10 dBm CW signals at their input mixer; the level below which mixer compression is specified to be under 1 dB for CW signals is usually 5 dB or more above this -10 dBm. The mixer behavior with Gaussian noise is not guaranteed, especially because its peak-to-average ratio is much higher than that of CW signals.

Keeping the mixer power below -10 dBm is a good practice that is unlikely to allow significant mixer nonlinearity. Thus, caution #1 is: Keep the total power at the input mixer at or below -10 dBm.

Figure D: In its center, this graph shows three curves: the ideal log amp behavior, that of a log amp that clips at its maximum and minimum extremes, and the average response to noise subject to that clipping. The lower right plot shows, on expanded scales, the error in average noise response due to clipping at the positive extreme. The average level should be kept 7 dB below the clipping level for an error level below 0.1 dB. The upper left plot shows, with an expanded vertical scale, the corresponding error for clipping against the bottom of the scale. The average level must be kept 14 dB above the clipping level for an error below 0.1 dB.



Caution 2, overdriving the log amp. Often, the level displayed has been heavily averaged using trace averaging or a video bandwidth (VBW) much smaller than the RBW. In such a case, instantaneous noise peaks are well above the displayed average level. If the level is high enough that the log amp has significant errors for these peak levels, the average result will be in error. Figure D shows the error due to overdriving the log amp in the lower right corner, based on a model that has the log amp clipping at the top of its range. Typically, log amps are still close to ideal for a few dB above their specified top, making the error model conservative. But it is possible for a log amp to switch from log mode to linear (voltage) behavior at high levels, in which case larger (and of opposite sign) errors to those computed by the model are possible. Therefore, caution #2 is: Keep the displayed average log level at least 7 dB below the maximum calibrated level of the log amp.

Caution 3, underdriving the log amp. The opposite of the overdriven log amp problem is the underdriven log amp problem. With a clipping model for the log amp, the results in the upper left corner of Figure D were obtained. The caution #3 is: Keep the displayed average log level at least 14 dB above the minimum calibrated level of the log amp.


Part II: Measurements of noise-like signals

In Part I, we discussed the characteristics of noise and its measurement. In this part, we'll discuss three different measurements of digitally modulated signals, after showing why they are very much like noise.

The noise-like nature of digital signals

Digitally modulated signals are created by clocking a DAC with the symbols (a group of bits simultaneously transmitted), then passing the DAC output through a pre-modulation filter (to reduce the transmitted bandwidth), then modulating the carrier with the filtered signal. See Figure 8. The resulting signal is obviously not noise-like if the digital signal is a simple pattern. It also does not have a noise-like distribution if the bandwidth of observation is wide enough for the discrete nature of the DAC outputs to significantly affect the distribution of amplitudes.

Figure 8: A simplified model for the generation of digital communications signals.



But, under many circumstances, especially test conditions, the digital signal bits are random. And, as exemplified by the “channel power” measurements discussed below, the observation bandwidth is narrow. If the digital update period (the reciprocal of the symbol rate) is less than one-fifth the duration of the majority of the impulse response of the resolution bandwidth filter, the signal within the RBW is approximately Gaussian according to the central limit theorem.

A typical example is IS-95 CDMA. Performing spectrum analysis, such as the adjacent-channel power ratio (ACPR) test, is usually done using the 30 kHz RBW to observe the signal. This bandwidth is only one-fortieth of the symbol clock (1.23 Msymbols/s), so the signal in the RBW is the sum of the impulse responses to about forty pseudorandom digital bits. A Gaussian PDF is an excellent approximation to the PDF of this signal.

Channel-power measurements

Most modern spectrum analyzers allow the measurement of the power within a frequency range, called the channel bandwidth. The displayed result comes from the computation:

$$P_{ch} = \left(\frac{B_s}{B_n}\right) \left(\frac{1}{N}\right) \sum_{i=n1}^{n2} 10^{(p_i/10)}$$


p_{ch} is the power in the channel, B_s is the specified bandwidth (also known as the channel bandwidth), B_n is the equivalent noise bandwidth of the RBW used, N is the number of data points in the summation, p_i is the sample of the power in measurement cell i in dB units (if p_i is in dBm, p_{ch} is in milliwatts). $n1$ and $n2$ are the end-points for the index i within the channel bandwidth, thus $N = (n2 - n1) + 1$.

The computation works excellently for CW signals, such as from sinusoidal modulation. The computation is a power-summing computation. Because the computation changes the input data points to a power scale before summing, there is no need to compensate for the difference between the log of the average and the average of the log as explained in Part I of this article series, even if the signal has a noise-like PDF (probability density function). But if the signal starts with noise-like statistics but is averaged in decibel form (typically with a VBW filter on the log scale) before the power summation, some of 2.51 dB under-response explained in Part I will be incurred. If we are certain that the signal is of noise-like statistics, and we fully average the signal before performing the summation, we can add 2.51 dB to the result and have an accurate measurement. Furthermore, the averaging reduces the variance of the result.

But if we don't know the statistics of the signal, the best measurement technique is to do no averaging before power summation. Using a $\text{VBW} \geq 3\text{RBW}$ is required for insignificant averaging, and is thus recommended. But the bandwidth of the video signal is not as obvious as it appears. In order to not peak-bias the measurement, the "sample" detector must be used. Spectrum analyzers have lower effective video bandwidths in sample detection than they do in peak detection mode, because of the limitations of the sample-and-hold circuit that precedes the A/D converter. Examples include the HP 8560E-Series spectrum analyzer family with 450 kHz effective sample-mode video bandwidth, and 800 kHz bandwidth in the HP 8590E-Series spectrum analyzer family.

Figure 9 shows the experimentally determined relationship between the VBW:RBW ratio and the under-response of the partially averaged logarithmically processed noise signal.

Figure 9: For $\text{VBW} \geq 3\text{RBW}$, the averaging effect of the VBW filter does not significantly affect power-detection accuracy.



Adjacent-Channel Power (ACP)

There are many standards for the measurement of ACP with a spectrum analyzer. The issues involved in most ACP measurements is covered in detail in an earlier article in *Microwaves & RF*, May, 1992, "Make Adjacent-Channel Power Measurements." A survey of other standards is available in "Adjacent Channel Power Measurements in the Digital Wireless Era" in *Microwave Journal*, July, 1994.

For digitally modulated signals, ACP and channel-power measurements are similar, except ACP is easier. ACP is usually the ratio of the power in the main channel to the power in an adjacent channel. If the modulation is digital, the main channel will have noise-like statistics. Whether the signals in the adjacent channel are due to broadband noise, phase noise, or intermodulation of noise-like signals in the main channel, the adjacent channel will have noise-like statistics. A spurious signal in the adjacent channel is most likely modulated to appear noise-like, too, but a CW-like tone is a possibility.

If the main and adjacent channels are both noise-like, then their ratio will be accurately measured regardless of whether their true power or log-averaged power (or any partially averaged result between these extremes) is measured. Thus, unless discrete CW tones are found in the signals, ACP is not subject to the cautions regarding VBW and other averaging noted in the section on channel power above.

But some ACP standards call for the measurement of absolute power, rather than a power ratio. In such cases, the cautions about VBW and other averaging do apply.

Carrier power

Burst carriers, such as those used in TDMA mobile stations, are measured differently than continuous carriers. The power of the transmitter during the time it is on is known as the "carrier power."

Carrier power is measured with the spectrum analyzer in "zero span." In this mode, the LO of the analyzer does not sweep, thus the span swept is zero. The display then shows amplitude normally on the y axis, and time on the x axis. If we set the RBW large compared to the bandwidth of the burst signal, then all the display points include all the power in the channel. The carrier power is computed simply by averaging the power of all the signals that represent the times when the burst is on. Depending on the modulation type, this is often considered to be any point within 20 dB of the highest registered amplitude. (A trigger and gated spectrum analysis may be used if the carrier power is to be measured over a specified portion of a burst-RF signal.)

Using a wide RBW for the carrier-power measurement means that the signal will not have noise-like statistics. It will not have CW-like statistics, either, so it is still wise to set the VBW as wide as possible. But let's consider some examples to see if the sample-mode bandwidths of spectrum analyzers are a problem.

For PDC, NADC and TETRA, the symbol rates are under 25 kb/s, so a VBW set to maximum will work excellently. It will also work well for PHS and GSM, with symbol rates of 380 and 270 kb/s. For IS-95 CDMA, with a modulation rate of 1.2288 MHz, we could anticipate a problem with the 450 and 800 kHz effective video bandwidths discussed in the section on channel power above. Experimentally, an instrument with an 800 kHz sample-mode bandwidth experienced a 0.2 dB error, and one with 450 kHz BW a 0.6 dB error with an OQPSK (mobile) burst signal.

Peak-detected noise (and TDMA ACP measurements)

TDMA (time-division multiple access, or burst-RF) systems are usually measured with peak detectors, in order that the burst “off” events are not shown on the screen of the spectrum analyzer, potentially distracting the user. Examples include ACP measurements for PDC (Personal Digital Cellular) by two different methods, PHS (Personal Handiphone System) and NADC (North American Dual-mode Cellular). Noise is also often peak detected in the measurement of rotating media, such as hard disk drives and VCRs.


The peak of noise will exceed its power average by an amount that increases (on average) with the length of time over which the peak is observed. A combination of analysis, approximation and experimentation leads to this equation for v_{pk} , the ratio of the average power of peak measurements to the average power of sampled measurements:

$$v_{pk} = [10 \text{ dB}] \log_{10} [L \ln(2\pi \tau BW_i + e)]$$

Tau (τ) is the observation period, usually given by either the length of an RF burst, or by the spectrum analyzer sweep time divided by the number of cells in a sweep. BW_i is the “impulse bandwidth” of the RBW filter, which is 1.62 times the -3 dB BW for the four-pole synchronously tuned filter used in most spectrum analyzers. Note that v_{pk} is a “power average” result; the average of the log of the ratio will be different.

The graph in Figure E shows a comparison of this equation with some experimental results. The fit of the experimental results would be even better if 10.7 dB were used in place of 10 dB in the equation above, even though analysis does not support such a change.

Figure E: The peak-detected response to noise increases with the observation time.



Part III: Averaging and the noisiness of noise measurements

The results of measuring noise-like signals are, not surprisingly, noisy. Reducing this noisiness is accomplished by three types of averaging:

- increasing the averaging within each measurement cell of a spectrum analyzer by reducing the VBW.
- increasing the averaging within a computed result like channel power by increasing the number of measurement cells contributing to the result.
- averaging a number of computed results.

Variance and averaging

The variance of a result is defined as the square of its standard deviation, therefore it is symbolically σ^2 . The variance is inversely proportional to the number of independent results averaged, thus when N results are combined, the variance of the final result is σ^2/N .

The variance of a channel-power result computed from N independent measurement cells is likewise σ^2/N where σ is the variance of a single measurement cell. But this σ^2 is a very interesting parameter.

If we were to measure the standard deviation of logged envelope noise, we would find that the σ is 5.57 dB. Thus, the σ of a channel-power measurement that averaged log data over, for example, 100 measurement cells would be 0.56 dB ($5.6/\sqrt{100}$). But averaging log data not only causes the aforementioned 2.51 dB under-response, it also has a higher than desired variance. Those not-rare-enough negative spikes of envelope, such as -30 dB, add significantly to the variance of the log average even though they represent very little power. The variance of a power measurement made by averaging power is lower than that made by averaging the log of power by a factor of 1.64.

Thus, the σ of a channel-power measurement is lower than that of a log-averaged measurement by a factor of the square root of this 1.64:

$$\sigma_{noise} = 4.35 \text{ dB}/\sqrt{N} \quad [\text{power averaging}]$$

$$\sigma_{noise} = 5.57 \text{ dB}/\sqrt{N} \quad [\text{log processing}]$$

Averaging a number of computed results

If we average individual channel-power measurements to get a lower-variance final estimate, we do not have to convert dB-format answers to absolute power to get the advantages of avoiding log averaging. The individual measurements, being the results of many measurement cells summed together, no longer have a distribution like the “logged Rayleigh” but rather look Gaussian. Also, their distribution is sufficiently narrow that the log (dB) scale is linear enough to be a good approximation of the power scale. Thus, we can dB-average our intermediate results.

Swept versus FFT analysis

In the above discussion, we have assumed that the variance reduced by a factor of N was of independent results. This independence is typically the case in swept-spectrum analyzers, due to the time required to sweep from one measurement cell to the next under typical conditions of span, RBW and sweep time. FFT analyzers will usually have many fewer independent points in a measurement across a channel bandwidth, reducing, but not eliminating, their theoretical speed advantage for true noise signals.

For digital communications signals, FFT analyzers have an even greater speed advantage than their throughput predicts. Consider a constant-envelope modulation, such as used in GSM cellular phones. When measured with a sweeping analyzer, with an RBW much narrower than the symbol rate, the spectrum looks noise-like. But in an FFT span wider than the spectral width of the signal, the total power looks constant, so channel power measurements will have very low variance.

Zero span


A zero-span measurement of carrier power is made with a wide RBW, so the independence of data points is determined by the symbol rate of the digital modulation. Data points spaced by a time greater than the symbol rate will be almost completely independent.

Zero span is sometimes used for other noise and noise-like measurements where the noise bandwidth is much greater than the RBW, such as in the measurement of power spectral density. For example, some companies specify IS-95 CDMA ACPR measurements that are spot-frequency power spectral density specifications; zero span can be used to speed this kind of measurement.

The standard deviation of measurement noise

Figure 10 summarizes the standard deviation of the measurement of noise. The figure represents the standard deviation of the measurement of a noise-like signal using a spectrum analyzer in zero span, averaging the results across the entire screen width, using the log scale. t_{INT} is the integration time (sweep time). The curve is also useful for swept spectrum measurements, such as channel-power measurements. There are three regions to the curve.

Figure 10: Noise measurement standard deviation for log-response spectrum analysis depends on the sweep-time/RBW product, the VBW/RBW ratio, and the number of display cells.



The left region applies whenever the integration time is short compared to the rate of change of the noise envelope. As discussed above, without VBW filtering, the σ is 5.6 dB. When video filtering is applied, the standard deviation is improved by a factor. That factor is the square root of the ratio of the two noise bandwidths: that of the video bandwidth, to that of the detected envelope of the noise. The detected envelope of the noise has half the noise bandwidth of the undetected noise. For the four-pole synchronously tuned filters typical of most spectrum analyzers, the detected envelope has a noise bandwidth of $(\frac{1}{2}) \times 1.128$ times the RBW. The noise bandwidth of a single-pole VBW filter is $\pi/2$ times its bandwidth. Gathering terms together yields the equation:

$$\sigma = (9.3 \text{ dB}) \sqrt{VBW/RBW}$$

The middle region applies whenever the envelope of the noise can move significantly during the integration time, but not so rapidly that individual sample points become uncorrelated. In this case, the integration behaves as a noise filter with frequency response of $\text{sinc}(\pi t_{INT})$ and an equivalent noise bandwidth of $1/(2t_{INT})$. The total noise should then be 5.6 dB times the square root of the ratio of the noise bandwidth of the integration process to the noise bandwidth of the detected envelope, giving

$$5.2 \text{ dB}/\sqrt{t_{INT} \text{ RBW}}$$

In the right region, the sweep time of the spectrum analyzer is so long that the individual measurement cells are independent of each other. In this case, the standard deviation is reduced from that of the left-side case (the sigma of an individual sample) by the square root of the number of measurement cells in a sweep.

The noise measurement sigma graph should be multiplied by a factor of about 0.8 if the noise power is filtered and averaged, instead of the log power being so processed. (Sigma goes as the square root of the variance, which improves by the cited 1.64 factor.) This factor applies to channel-power and ACP measurements, but does not apply to VBW-filtered measurements by any current-generation spectrum analyzers.

Examples

Let's use the curve in Figure 10 in two examples. In the measurement of CDMA ACPR, we can power-average a 400-point zero-span trace for a frame (20.2 ms) in the specified 30 kHz bandwidth. Power averaging requires $\text{VBW} > \text{RBW}$. For these conditions, we find $t_{INT} \text{ RBW} = 606$, and we approach the right-side asymptote of $5.6 \text{ dB}/\sqrt{400 \text{ points}}$ or 0.28 dB. But we are power averaging, so we multiply by 0.8 to get $\text{sigma} = 0.22 \text{ dB}$.

In a second example, we are measuring noise in an adjacent channel in which the noise spectrum is flat. Let's use a 600-point analyzer with a span of 100 kHz and a channel BW of 25 kHz, giving 150 points in our channel. Let's use an RBW of 300 Hz and a VBW = 10 Hz; this narrow VBW will prevent power detection and lead to about a 2.3 dB under-response (see Figure 9) for which we must manually correct. The sweep time will be 84 s, or 21 s within the channel. $t_{INT} \text{ RBW} = 6300$; if the center of Figure 10 applied, sigma would be 0.066 dB. Checking the right asymptote, it works out to be 0.083 dB, so this is our final predicted standard deviation. If the noise in the adjacent channel is not flat, the averaging will effectively extend over many fewer samples and less time, giving a higher standard deviation.

Part IV: Compensation for instrumentation noise

In Parts I, II and III, we discussed the measurement of noise and noise-like signals respectively. In this part, we'll discuss measuring CW and noise-like signals in the presence of instrumentation noise. We'll see why averaging the output of a logarithmic amplifier is optimum for CW measurements, and we'll review compensation formulas for removing known noise levels from noise-plus-signal measurements.


CW signals and log versus power detection

When measuring a single CW tone in the presence of noise, and using power detection, the level measured is equal to the sum of the power of the CW tone and the power of the noise within the RBW filter. Thus, we could improve the accuracy of a measurement by measuring the CW tone first (let's call this the "S+N" or signal-plus-noise), then disconnect the signal to make the "N" measurement. The difference between the two, with both measurements in power units (for example, milliwatts, not dBm) would be the signal power.

But measuring with a log scale and video filtering or video averaging results in unexpectedly good results. As described in Part I, the noise will be measured lower than a CW signal with equal power within the RBW by 2.5 dB. But to first order, the noise doesn't even affect the S+N measurement! See "Log Scale Ideal for CW Measurements" later in this section.

Figure 11 demonstrates the improvement in CW measurement accuracy when using log averaging versus power averaging.

Figure 11: Log averaging improves the measurement of CW signals when their amplitude is near that of the noise. (a) shows a noise-free signal. (b) shows an averaged trace with power-scale averaging and noise power 1 dB below signal power; the noise-induced error is 2.5 dB. (c) shows the effect with log-scale averaging—the noise falls 2.5 dB and the noise-induced error falls to only 0.6 dB.



To compensate S+N measurements on a log scale for higher-order effects and very high noise levels, use this equation where all terms are in dB units:

$$power_{cw} = power_{S+N} - 10.42 \times 10^{-0.333(\delta SN)}$$

$power_{S+N}$ is the observed power of the signal with noise. δSN is the decibel difference between the S+N and N-only measurements. With this compensation, noise-induced errors are under 0.25 dB even for signals as small as 9 dB below the interfering noise. Of course, in such a situation, the repeatability becomes a more important concern than the average error. But excellent results can be obtained with adequate averaging. And the process of averaging and compensating, when done on a log scale, converges on the result much faster than when done in a power-detecting environment.

Power-detection measurements and noise subtraction

If the signal to be measured has the same statistical distribution as the instrumentation noise, in other words, if the signal is noise-like, then the sum of the signal and instrumentation noise will be a simple power sum:

$$power_{S+N} = power_S + power_N \quad [mW]$$

Note that the units of all variables must be power units such as milliwatts, and not log units like dBm, nor voltage units like mV. Note also that this equation applies even if $power_S$ and $power_N$ are measured with log averaging.

The power equation also applies when the signal and the noise have different statistics (CW and Gaussian respectively) but power detection is used. The power equation would never apply if the signal and the noise were correlated, either in-phase adding or subtracting. But that will never be the case with noise.

Therefore, simply enough, we can subtract the measured noise power from any power-detected result to get improved accuracy. Results of interest are the channel-power, ACP and carrier-power measurements described in Part II. The equation would be:

$$power_S = power_{S+N} - power_N \quad [mW]$$


Care should be exercised that the measurement setups for $power_{S+N}$ and $power_N$ are as similar as possible.

Log scale ideal for CW measurements

If one were to “design” a scale (such as power, voltage, log power, or an arbitrary polynomial) to have response to signal-plus-noise that is independent of small amounts of noise, one could end up designing the log scale.

Consider a signal having unity amplitude and arbitrary phase, as in Figure G. Consider noise with an amplitude much less than unity, r.m.s., with random phase. Let us break the noise into components that are in-phase and in-quadrature with the signal. Both of these components will have Gaussian PDFs, but for this simplified explanation, we can consider them to have values of $\pm x$, where $x \ll 1$.

Figure G: Noise components can be projected into in-phase and quadrature parts with respect to a signal of unity amplitude and arbitrary phase.




The average response to the signal plus the quadrature noise component is the response to a signal of magnitude

$$\sqrt{1+x^2}$$

The average response to the signal plus in-phase noise will be lower than the response to a signal without noise if the chosen scale is compressive. For example, let x be ± 0.1 and the scale be logarithmic. The response for $x = +0.1$ is $\log(1.1)$; for $x = -0.1$ is $\log(0.9)$. The mean of these two is 0.0022, also expressible as $\log(0.9950)$. The mean response to the quadrature components is $\log(\sqrt{2}(1 + (0.1)^2))$, or $\log(1.0050)$. Thus, the log scale has an average deviation for in-phase noise that is equal and opposite to the deviation for quadrature noise. To first order, the log scale is noise-immune. Thus, an analyzer that averages (for example, by video filtering) the response of a log amp to the sum of a CW signal and a noise signal has no first-order dependence on the noise signal.

Figure H shows the average error due to noise addition for signals measured on the log scale and, for comparison, for signals measured on a power scale.

Figure H: CW signals measured on a logarithmic scale show very little effect due to the addition of noise signals.



Bibliography

1. Nutting, Larry. Cellular and PCS TDMA Transmitter Testing with a Spectrum Analyzer. Hewlett-Packard Wireless Symposium, February, 1992.
2. Gorin, Joe. Make Adjacent Channel Power Measurements, *Microwaves & RF*, May 1992, pp 137-143.
3. Cutler, Robert. Power Measurements on Digitally Modulated Signals. Hewlett-Packard Wireless Communications Symposium, 1994.
4. Ballo, David and Gorin, Joe. Adjacent Channel Power Measurements in the Digital Wireless Era, *Microwave Journal*, July 1994, pp 74-89.
5. Peterson, Blake. Spectrum Analysis Basics. Hewlett-Packard Application Note 150, literature part number 5952-0292, November 1, 1989.
6. Moulthrop, Andrew A. and Muha, Michael S. Accurate Measurement of Signals Close to the Noise Floor on a Spectrum Analyzer, *IEEE Transactions on Microwave Theory and Techniques*, November 1991, pp. 1182-1885.

Glossary of terms

ACP: See Adjacent Channel Power.

ACPR: Adjacent Channel Power Ratio. See Adjacent-Channel Power; ACPR is always a ratio, whereas ACP may be an absolute power.

Adjacent Channel Power: The power from a modulated communications channel that leaks into an adjacent channel. This leakage is usually specified as a ratio to the power in the main channel, but is sometimes an absolute power.

Averaging: A mathematical process to reduce the variation in a measurement by summing the data points from multiple measurements and dividing by the number of points summed.

Burst: A signal that has been turned on and off. Typically, the on time is long enough for many communications bits to be transmitted, and the on/off cycle time is short enough that the associated delay is not distracting to telephone users.

Carrier Power: The average power in a burst carrier during the time it is on.

CDMA: Code Division Multiple Access or a communications standard (such as cdmaOne (R)) that uses CDMA. In CDMA modulation, data bits are xored with a code sequence, increasing their bandwidth. But multiple users can share a channel when they use different codes, and a receiver can separate them using those codes.

Channel Bandwidth: The bandwidth over which power is measured. This is usually the bandwidth in which almost all of the power of a signal is contained.

Channel Power: The power contained within a channel bandwidth.

Clipping: Limiting a signal such that it never exceeds some threshold.

CW: Carrier Wave or Continuous Wave. A sinusoidal signal without modulation.

Digital: Signals that can take on only a prescribed list of values, such as 0 and 1.

Display detector: That circuit in a spectrum analyzer that converts a continuous-time signal into sampled data points for displaying. The bandwidth of the continuous-time signal often exceeds the sample rate of the display, so display detectors implement rules, such as peak detection, for sampling.

DAC: Digital to Analog Converter.

Envelope Detector: The circuit that derives an instantaneous estimate of the magnitude (in volts) of the IF (intermediate frequency) signal. The magnitude is often called the envelope.

Equivalent Noise Bandwidth: The width of an ideal filter with the same average gain to a white noise signal as the described filter. The ideal filter has the same gain as the maximum gain of the described filter across the equivalent noise bandwidth, and zero gain outside that bandwidth.

Gaussian and Gaussian PDF: A bell-shaped PDF which is typical of complex random processes. It is characterized by its mean (center) and sigma (width).

I and Q: In-phase and Quadrature parts of a complex signal. I and Q, like x and y, are rectangular coordinates; alternatively, a complex signal can be described by its magnitude and phase, also known as polar coordinates.

Linear scale: The vertical display of a spectrum analyzer in which the y axis is linearly proportional to the voltage envelope of the signal.

NADC: North American Dual mode (or Digital) Cellular. A communications system standard, designed for North American use, characterized by TDMA digital modulation.

Near-noise Correction: The action of subtracting the measured amount of instrumentation noise power from the total system noise power to calculate that part from the device under test.

Noise Bandwidth: See Equivalent Noise Bandwidth.

Noise Density: The amount of noise within a defined bandwidth, usually normalized to 1 Hz.

Noise Marker: A feature of spectrum analyzers that allows the user to read out the results in one region of a trace based on the assumption that the signal is noise-like. The marker reads out the noise density that would cause the indicated level.

OQPSK: Offset Quadrature-Phase Shift Keying. A digital modulation technique in which symbols (two bits) are represented by one of four phases. The set of four phases is offset by 45 degrees on alternate symbols.

PDC: Personal Digital Cellular (originally called Japanese Digital Cellular). A cellular radio standard much like NADC, originally designed for use in Japan.

PDF: See probability density function.

Peak Detect: Measure the highest response within an observation period.

PHS: Personal Handy-Phone. A communications standard for cordless phones.

Power Detection: A measurement technique in which the response is proportional to the power in the signal, or proportional to the square of the voltage.

Power Spectral Density: The power within each unit of frequency, usually normalized to 1 Hz.

Probability Density Function: A mathematical function that describes the probability that a variable can take on any particular x-axis value. The PDF is a continuous version of a histogram.

Q: See I and Q.

Rayleigh: A well-known PDF which is zero at $x=0$ and approaches zero as x approaches infinity.

RBW filter: The resolution bandwidth filter of a spectrum analyzer. This is the filter whose selectivity determines the analyzer's ability to resolve (indicate separately) closely-spaced signals.

Reference Bandwidth: See specified bandwidth.

RF: Radio Frequency. Frequencies that are used for radio communications.

Sigma: The symbol and name for standard deviation.

Sinc: A mathematical function. $\text{Sinc}(x) = (\sin(x))/x$.

Specified Bandwidth: The channel bandwidth specified in a standard measurement technique.

Standard Deviation: A measure of the width of the distribution of a random variable.

Synchronously Tuned Filter: The filter alignment most commonly used in analog spectrum analyzers. A sync-tuned filter has all its poles in the same place. It has an excellent tradeoff between selectivity and time-domain performance (delay and step-response settling).

Symbol: A combination of bits (often two) that are transmitted simultaneously.

Symbol Rate: The rate at which symbols are transmitted.

TDMA: Time Division Multiple Access. A method of sharing a communications channel by assigning separate time slots to individual users.

TETRA: Trans-European Trunked Radio. A communications system standard.

Variance: A measure of the width of a distribution, equal to the square of the standard deviation.

VBW Filter: The Video Bandwidth filter, a low-pass filter that smoothes the output of the detected IF signal, or the log of that detected signal.

Zero Span: A mode of a spectrum analyzer in which the local oscillator does not sweep. Thus, the display represents amplitude versus time, instead of amplitude versus frequency. This is sometimes called fixed-tuned mode.

For more information about Hewlett-Packard test and measurement products, applications, services, and for a current sales office listing, visit our web site, <http://www.hp.com/go/tmdir>. You can also contact one of the following centers and ask for a test and measurement sales representative.

United States:

Hewlett-Packard Company
Test and Measurement Call Center
P.O. Box 4026
Englewood, CO 80155-4026
1 800 452 4844

Canada:

Hewlett-Packard Canada Ltd.
5150 Spectrum Way
Mississauga, Ontario L4W 5G1
(905) 206 4725

Europe:

Hewlett-Packard
European Marketing Centre
P.O. Box 999
1180 AZ Amstelveen
The Netherlands
(31 20) 547 9900

Japan:

Hewlett-Packard Japan Ltd.
Measurement Assistance Center
9-1, Takakura-Cho, Hachioji-Shi,
Tokyo 192, Japan
Tel: (81) 426 56 7832
Fax: (81) 426 56 7840

Latin America:

Hewlett-Packard
Latin American Region Headquarters
5200 Blue Lagoon Drive, 9th Floor
Miami, Florida 33126, U.S.A.
Tel: (305) 267-4245
(305) 267-4220
Fax: (305) 267-4288

Australia/New Zealand:

Hewlett-Packard Australia Ltd.
31-41 Joseph Street
Blackburn, Victoria 3130, Australia
1 800 629 485

Asia Pacific:

Hewlett-Packard Asia Pacific Ltd.
17-21/F Shell Tower, Times Square,
1 Matheson Street, Causeway Bay,
Hong Kong
Tel: (852) 2599 7777
Fax: (852) 2506 9285