ACPR on Cassini: Direct Measurement for Spread Spectrum Signals

Adjacent Channel Power Ratio (ACPR) is one of the many ways to measure distortion in active circuits. It is often used with spread spectrum signals. It is a measurement of the amount of signal, or power, in the frequencies adjacent to the primary signal. ACPR is defined as the ratio of the RMS power in the adjacent frequency channel (or offset) to the RMS power in the transmitted frequency channel. ACPR is essentially the spread-spectrum counterpart to two-tone intermodulation measurements. The difference is that, instead of just two tones, a spread-spectrum signal contains many components that all intermodulate with each other. When this signal passes through active circuitry, these interactions create small sidebands that spill into adjacent frequency channels. The more severe the intermodulation distortion, the stronger these adjacent-channel sidebands become.

ACPR is measured by measuring the amplitude of the carrier over its channel bandwidth, then measuring the amplitude of the power in the adjacent channel and calculating the ratio of the two measurements. This presents several considerations, as we will see. Note that we will use WCDMA as our example, but these comments can apply to any spread spectrum signal.

Modulation Pattern
The modulation is typically a pseudo-random signal, structured to be as close to uniform as possible. The pseudo-random pattern must be chosen carefully for optimum performance.

In one example, the user noticed an occasional measurement that was slightly different, with the carrier measuring about 0.25 dB higher than normal. The user found that the pattern, in time domain, had small amplitude bumps. The Roos tester is fast enough that its measurement would occasionally occur completely inside the peak of those bumps, more careful pattern design might have eliminated this, but a more expedient solution was to slow the measurement to encompass more time, averaging out the bumps. Although it slowed the measurement, this was a viable solution.

Patterns should be designed for clock rates of 1MHz, 10MHz, 20MHz, 32MHz, or 40MHz. If not, the pattern must be re-sampled to one of those rates. The pattern must be less than 8M complex pairs to properly fit into the RI8508C's pattern memory.

Once created, the pattern can be loaded into a Roos RI8508C signal generator. For more information on creating and importing patterns, refer to Importing a Waveform File into Cassini.

Carrier Measurement
Measuring the carrier is straightforward. It can be measured with a wide filter, as the tiny sidebands added in have negligible effect. This assumes, of course, that there are no other active channels leaking in. In a typical test environment, that can be controlled.

There are two considerations that must be remembered when measuring the carrier:

1) Absolute power accuracy is important, because the Device Under Test (DUT) is typically operating in the compressed non-linear region. A small error in absolute power can make a significant difference in ACPR.

2) Your receiver must have enough bandwidth to encompass the entire channel. For example, in WCDMA , the channel bandwidth is usually 5 MHz, with the middle 3.84 MHz used for data. The Roos wideband filter is 5 MHz, which is enough, some instruments measure in a narrower bandwidth and extrapolate, which is possible if the pattern is trusted to be flat across its frequency bandwidth.

Spurious-free Dynamic Range and Dynamic Range are Different

Dynamic range is simply the largest signal that a receiver can measure minus the smallest signal that same receiver can measure with no interference from other signals. If the receiver has variable attenuation or gain, those settings can be independently set for each measurement.

Spurious-free dynamic range is the ability to measure a small signal in the presence of a large interfering signal. Spurious-free dynamic range is always smaller than dynamic range because the receiver's front end must be adjusted to a lower sensitivity, so that the large signal doesn't cause distortion of the small signal.

When measuring ACPR, the adjacent channel sideband is a small signal in the presence of the main channel’s large signal. Therefore, the spurious-free dynamic range of a receiver must be considered.

On a typical Roos Cassini, the dynamic range varies with multiple factors but is typically about 140 dB. The spurious-free dynamic range, if fully optimized, can be about 60 dB, but external factors typically limit it to about 55 dB.

Adjacent Channel Measurement

Measuring the adjacent channel is where the most significant challenges arise. There are 3 main challenges:

1) To a receiver, spread spectrum signals look like noise. This means that vector noise reduction techniques, such as averaging complex signals to extract a signal from the noise, don't work. Any receiver must be configured to act as a 'scalar-only' receiver. On Roos equipment, the complex detector can be used with scalar averaging, but typically the RMS detector is simpler.

2) The system's receive front end is always being hit with the carrier power, regardless of whether we are measuring the adjacent channel or not. This means that the receive RF path must always be set so that the carrier does not compress the front end (or, worse, damage it). IF gain, which occurs after the IF filter, can be adjusted as necessary but the receiver's front-end sensitivity cannot be optimized for the adjacent channel. Note that this ability to measure a small signal in the presence of a large signal is called 'spurious-free dynamic range'.

   Also, you must account for the Peak-to-Average Ratio (PAR) of the signal to keep the peaks from compressing the front end. Note that the tester's overdrive warnings can't be fully trusted here, because they're based on the RMS power, not the peak. It's possible to have an RMS power below the warning level but still have peaks above it. You could simply drop the power into the receiver by the PAR amount and be safe, but for optimum sensitivity, it's usually possible to allow a little more power than that amount. Allowing the very highest peaks to compress a little typically has insignificant effect on the measurement.

3) Some standards specify to use a specific bandwidth (typically 30 kHz) to measure in the middle of the adjacent channel. Others, like WCDMA, specify measuring over the entire channel. This is a problem because it's impossible to filter out the large carrier to measure the small adjacent channel sideband. Any filter wide enough to encompass the adjacent channel won't sufficiently

reject the carrier. The most effective technique is to use a narrow filter, step it along over the adjacent channel and sum the powers.

Measurement Technique for a WCDMA Signal

1) With the test setup optimized for the carrier power with the RMS detector and 5 MHz IF bandwidth, measure the carrier power.

2) Change to a narrow filter. The 30 kHz filter will work but we prefer the 7 kHz filter as it has sharper walls and better rejects the carrier. Start 0.5 MHz into the adjacent channel and measure. Then step and measure. Continue until you are 0.5 MHz from the edge of the adjacent channel. This gives a measurement span of 3-4 MHz, but if you are still getting a little of the carrier signal power leaking into the measurement, start a little further away from the carrier.

3) Add the adjacent power measurements together, apply a multiplication factor, and ratio this to the carrier measurement.

About the Step Size

The idea is to choose a step size that minimizes measurement time without missing major portions of the pattern. In a perfect world, a perfect step size would precisely match the filter bandwidth. Then you could simply add all the measured powers and get the total ACP. However, it is virtually impossible to choose a step size that perfectly matches the filter bandwidth. If the step is too small, then some power is 'double measured'. If the step size is too large, then some power is not measured.

However, if the pattern is reasonably consistent, then you don't have to measure every frequency point. You can take larger steps, add the powers, then apply a multiplication factor to account for the missed bandwidth.

There's a little bit of trial and error to pick the optimum step size. A smaller step size (measurements closer together) gives more repeatable values but takes more time. However, too large of a step size can reduce repeatability.

We found that approximately 50 steps provide solid repeatable values, although some users report acceptable results with as few as 20 steps.

About the Multiplication Factor

The multiplication factor is because we're not equally measuring every frequency. The exact value of the multiplication factor is difficult to precisely calculate, primarily because the filter's noise bandwidth is not precisely known. See the sidebar showing why noise bandwidth is always wider than CW bandwidth below. Also, the exact pseudo-random pattern can have some effect, depending on which frequency points are measured.

The simplest technique is to measure a few known ACPR's and choose the multiplication factor to match their known values. Some might say that this technique is cheating, however we're not creating a 'fudge factor'. Rather than try to calculate based on the unknowable noise bandwidth, we're doing solid repeatable measurements and empirically determining the correct factor. In other words, we're calibrating the measurement, knowing precisely why the calibration is needed.

Noise Bandwidth is Always Wider than CW bandwidth

When you apply a CW signal to a filter, the output power is dependent on where you are in the filter, full power in the middle, decreasing power at the sides. The bandwidth is the frequency range between the points where the power has dropped by 3 dB.

When you apply noise to a filter, the output power is the total integrated noise power through the filter. Since it's white noise, the concept of a 3 dB bandwidth for noise doesn't apply. You get the same power anywhere you tune.

Accepted practice for noise is to use 'equivalent noise bandwidth', which is usually shortened to just 'noise bandwidth'.

Equivalent noise bandwidth is defined as the bandwidth of a theoretically perfect brick wall filter that passes the same total integrated power as your filter. In other words, if you have a perfectly rectangular filter that passes the same noise power as your filter, they have the same noise bandwidth.

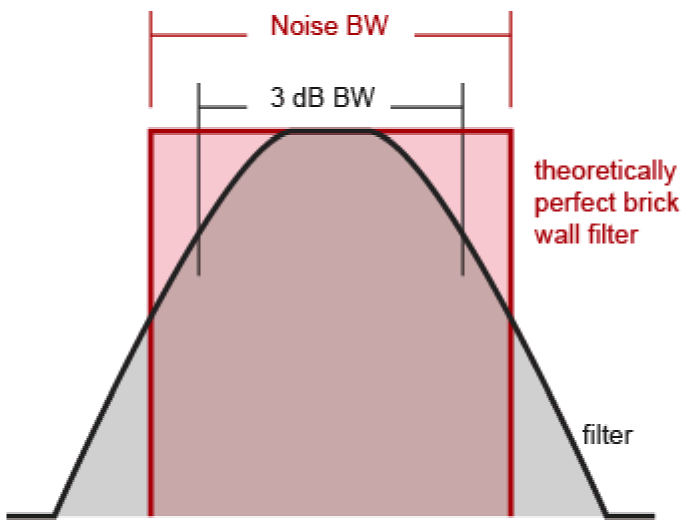

Noise bandwidth is ALWAYS wider than 3 dB bandwidth, because there's still filter left beyond the 3 dB points. The more poles the filter has, the steeper the sides, and the closer the noise bandwidth is to the 3 dB bandwidth. Since a perfectly rectangular filter is not possible, the noise bandwidth is always wider than the 3 dB bandwidth.

With regards to the filter in the Roos tester, we call it a 7 kHz filter, and its 3 dB bandwidth is approximately 7 kHz. But its noise bandwidth is closer to 10 kHz. That's why we use 10 kHz in our calculations.

Other Measurement Techniques

Roos uses a direct measurement approach, using filters and frequency steps to define the ACPR. There are, however, other ways to measure.

Fast Fourier Transform (FFT)

One technique measures several samples in time domain, then uses an FFT to convert to frequency domain, then calculate the ACPR.

This is a similar direct approach, where the signal is modulated and measured, only the measurement is done in time domain instead of frequency domain and then converted.

With the FFT approach there are trade-offs between step size, speed of each measurement and the total measurement time. To get temporal resolution, measurements must be made very fast. To get frequency resolution, the measurements must be made for a relatively long time. However, if properly chosen a good compromise of resolution, speed, and range can usually be achieved.

For direct approaches, Roos prefers to measure directly in the frequency domain and avoid sampling issues that can occur when using Fourier Transforms to convert domains.

Model-Based

There are many measurements that can be performed to measure distortion. ACPR is simply one way. If the fundamental distortion characteristics of a device can be accurately measured, then any of the distortion measurements (ACPR, IP3, P1dB, etc.) can be calculated from those measurements, if a proper model exists.

This is the concept behind model-based measurements. Roos has done significant research into this and has been able to get excellent correlation without having to modulate the signal. The key is accurately measuring complex distortion (AM to AM and AM to PM) through a device, which Roos testers excel at.

For more information on model-based ACPR reference documentation:
ACPR on Cassini, Model Based (using WCDMA ,EDGE, and WiMAX examples)