



Analog Dialogue

Volume 51, Number 3, 2017

Your Engineering Resource for Innovative Design

- 5 MEMS Vibration Monitoring: From Acceleration to Velocity
- 17 Transmit LO Leakage (LOL)—An Issue of Zero-IF That Isn't Making People Laugh Out Loud
- 22 Ultrawideband Digital Predistortion (DPD): The Rewards (Power and Performance) and Challenges of Implementation in Cable Distribution Systems
- 31 Improved DAC Phase Noise Measurements Enable Ultralow Phase Noise DDS Applications
- 36 Mirror, Mirror on the Wall—Understanding Image Rejection and Its Impact on Desired Signals
- 41 Wireless Current Sense Circuit Floats with Sense Resistor
- 45 GaN Breaks Barriers—RF Power Amplifiers Go Wide and High
- 48 Intelligent Video Analytics at the Edge of IoT
- 51 Uncompromising Linearity from the LTC2185 and ADA4927-1

10 Massive MIMO and Beamforming: The Signal Processing Behind the 5G Buzzwords



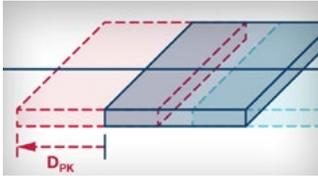
50
1967 to 2017



Visit analogdialogue.com

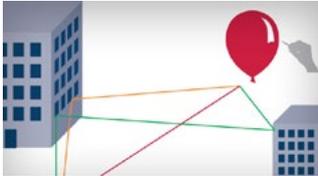
**ANALOG
DEVICES**
AHEAD OF WHAT'S POSSIBLE™

In This Issue



5 MEMS Vibration Monitoring: From Acceleration to Velocity

MEMS accelerometers have finally reached a point where they are able to measure vibration on a broad set of machine platforms. Recent advances in their capability, along with the many advantages that MEMS accelerometers already had over more traditional vibration sensors (size, weight, cost, shock immunity, ease of use), are motivating the use of MEMS accelerometers in an emerging class of condition-based monitoring (CBM) systems. As a result, many CBM system architects, developers, and even their customers are giving consideration to these types of sensors for the first time.



10 Massive MIMO and Beamforming: The Signal Processing Behind the 5G Buzzwords

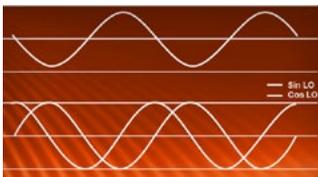
Our thirst for high speed mobile data is insatiable. As we saturate the available RF spectrum in dense urban environments, it's becoming apparent that there's a need to increase the efficiency of how we transmit and receive data from wireless base stations.

Base stations consisting of large numbers of antennas that simultaneously communicate with multiple spatially separated user terminals over the same frequency resource and exploit multipath propagation are one option to achieve this efficiency saving. This technology is often referred to as massive MIMO (multiple-input, multiple-output).



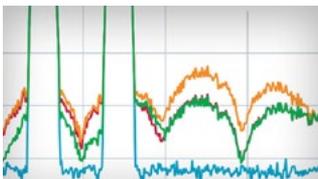
15 Orthogonal Perspectives

I am using a MEMS inertial measurement unit (IMU) in a self-balancing guidance control system for a personal transportation platform. Can I expect a consumer targeted IMU to eliminate all misalignment errors between each sensor if all of the core sensor elements are on a single piece of silicon?



17 Transmit LO Leakage (LOL)—An Issue of Zero-IF That Isn't Making People Laugh Out Loud

Transmit LOL refers to transmit local oscillator leakage. If you are working with RF mixers and have not corrected transmit LOL, this article explains why LOL correction is so important, especially in zero-IF architecture.



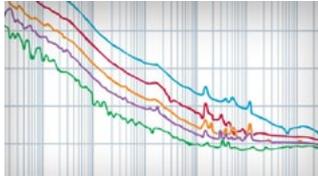
22 Ultrawideband Digital Predistortion (DPD): The Rewards (Power and Performance) and Challenges of Implementation in Cable Distribution Systems

Digital predistortion (DPD) could enhance the efficiency of power amplifiers in our cable infrastructure and more. Time to improve efficiency for the next generation of cable technologies.



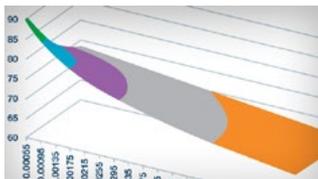
29 Burned by Low Power? When Lower Current Consumption Can Get You Into Trouble

What happens if you are replacing one part with a better (less power consuming) one and forget to check the power supply specifications? This RAQ comes out of a customer experience discussing some unexpected regulator phenomena.



31 Improved DAC Phase Noise Measurements Enable Ultralow Phase Noise DDS Applications

How can a good high speed or radar system be improved further? Phase noise is one of the critical issues to take care of. Not surprisingly, surrounding components and test setup improvements could make the difference. Peter Delos and Jarrett Liner—system engineers of instrumentation, aerospace, and defense—discuss these topics.



36 Mirror, Mirror on the Wall—Understanding Image Rejection and Its Impact on Desired Signals

Patrick Wiers introduces you to image rejection. Despite the title, it has nothing to do with what you see in a mirror. Images in the frequency domain should be rejected. The questions are: how to do that with a zero-IF approach? Without the possibility of active or passive filtering? And what is the right choice between the wideband and narrow-band transceivers offered?



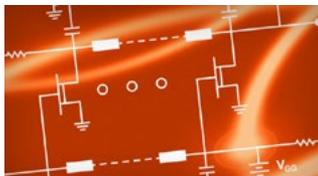
41 Wireless Current Sense Circuit Floats with Sense Resistor

In our premiere LTC article, Kris Lokere introduces a wireless current sense circuit that floats with sense resistors. An interesting concept— isolation without using traditional optocouplers or inductive or capacitive barriers? Curious? The secret is called SmartMesh® IP network. You will also see how well Analog Devices and Linear Technology's products fit together in a combined solution.



43 Accelerometer Tilt Measure Over Temperature and in the Presence of Vibration

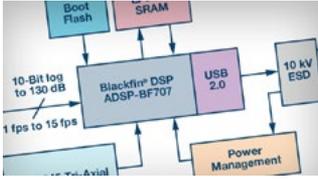
As you know, we live on the Earth with $1g$ acceleration toward the planet. An accelerometer with its sensing element oriented in this direction will measure $1g$. Tilt the sensor 90° and the output becomes $0g$. Everything between the tilt is related to the sine and cosine of the angle. But add in the factor of temperature and vibration, and are you still able to measure it with good enough accuracy? See what Christopher Murphy has to say.



45 GaN Breaks Barriers—RF Power Amplifiers Go Wide and High

Keith Benson discusses new approaches with GaN in power amplifier designs. He looks at products that demonstrate the possibilities of GaN technology, including covering wide bandwidths and providing high power and efficiency.

In This Issue



48 Intelligent Video Analytics at the Edge of IoT

Tyle Jesiel introduces us to SNAP sensors working in combination with our Blackfin® technology—a mix of logarithmic imaging and node analytics in the Internet of Things (IoT). By bringing in-depth video processing and analytics to the node, it will be possible to dramatically reduce the amount of data transmission into the cloud.



51 Uncompromising Linearity from the LTC2185 and ADA4927-1

I hope you had a chance to read our article from the August issue featuring products from Linear Technology and Analog Devices. We continue the series this month with the combination of a precision op amp with great linearity driving a precision 16-bit ADC, which compliments the linear requirements for excellent ac performance applications.



53 A Low Power, Low Cost, Differential Input to a Single-Ended Output Amplifier

A low power, low cost, differential input to a single-ended output amplifier from Jordyn Rombola and Chan Trau is the topic of this RAQ. This article asks: How often do you need a ground referenced signal after a differential transmission of that signal? Using a differential to a single-ended application, the differential amplifier rejects the common-mode voltage and the remaining voltage is amplified and presented on the amplifier output as a single-ended voltage.



Bernhard Siegel, Editor

Bernhard became editor of *Analog Dialogue* in March 2017, when the preceding editor, Jim Surber, decided to retire. Bernhard has been with Analog Devices for over 25 years, starting at the ADI

Munich office in Germany.

Bernhard has worked in various engineering roles including sales, field applications, and product engineering, as well as in technical support and marketing roles.

Residing near Munich, Germany, Bernhard enjoys spending time with his family and playing trombone and euphonium in both a brass band and a symphony orchestra.

You can reach him at bernhard.siegel@analog.com.



Analog Dialogue

Analog Dialogue is a technical magazine created and published by Analog Devices. It provides in-depth design related information on products, applications, technology, software, and system solutions for analog, digital, and mixed-signal processing. Published continuously for 50 years—starting in 1967—it is produced as a monthly online edition and as a printable quarterly journal featuring article collections. For history buffs, the *Analog Dialogue* archive includes all issues, starting with Volume 1, Number 1, and four special anniversary editions. To access articles, the archive, the journal, design resources, and to subscribe, visit the *Analog Dialogue* homepage, analogdialogue.com.

MEMS Vibration Monitoring: From Acceleration to Velocity

By Mark Looney

Share on   

Introduction

MEMS accelerometers have finally reached a point where they are able to measure vibration on a broad set of machine platforms. Recent advances in their capability, along with the many advantages that MEMS accelerometers already had over more traditional vibration sensors (size, weight, cost, shock immunity, ease of use), are motivating the use of MEMS accelerometers in an emerging class of condition-based monitoring (CBM) systems. As a result, many CBM system architects, developers, and even their customers are giving consideration to these types of sensors for the first time. Quite often, they are faced with the problem of quickly learning how to evaluate the capability of MEMS accelerometers to measure the most important vibration attributes on their machine platforms. This might seem difficult at first, as MEMS accelerometer data sheets often express the most important performance attributes in terms that these developers may not be familiar with. For example, many are familiar with quantifying vibration in terms of linear velocity (mm/s), while most MEMS accelerometer data sheets express their performance metrics in terms of gravity-referenced acceleration (g). Fortunately, there are some simple techniques for making this translation from acceleration to velocity and for estimating the influence that key accelerometer behaviors (frequency response, measurement range, noise density) will have on important system-level criteria (bandwidth, flatness, peak vibration, resolution).

Basic Vibration Attributes

This process starts with a review of linear vibration from an inertial motion point of view. Within this context, vibration is a mechanical oscillation that has zero mean displacement. For those who don't want their machines to be moving across the factory floor, zero mean displacement is pretty important! The value of the core sensor in a vibration sensing node will be directly related to how well it can represent the most important attributes of a machine's vibration. In order to start assessing the capability of a specific MEMS accelerometer in this capacity, it's important to start with a basic understanding of vibration from an inertial motion point of view. Figure 1 provides a physical illustration of a vibration motion profile, where the gray box represents the middle point, the blue image represents the peak displacement in one direction, and the red image represents the peak displacement in the other direction. Equation 1 provides a mathematical model that describes the instantaneous acceleration of the rectangular object when it is vibrating at one frequency (f_v), at a magnitude of A_{rms} .

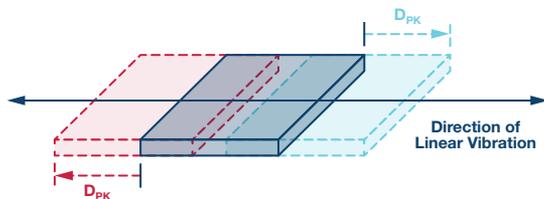


Figure 1. Simple linear vibration motion.

$$a(t) = A_{PK} \times \sin(\omega_v \times t)$$

$$A_{rms} = \frac{A_{PK}}{\sqrt{2}} \tag{1}$$

$$f_v = \frac{\omega_v}{2\pi}$$

In most CBM applications, the vibration on a machine platform is often going to have more complex spectral signature than the model in Equation 1, but this model provides a nice starting point in the discovery process, as it identifies two common vibration attributes that CBM systems often track: magnitude and frequency. This approach is also useful in translating key behaviors into terms of linear velocity as well (more on that later). Figure 2 provides a spectral view of two different types of vibration profiles. The first profile (see the blue lines in Figure 2) has a constant magnitude across its frequency range, which is between f_1 and f_6 . The second profile (see the green lines in Figure 2) has peaks in its magnitude at four different frequencies: f_2 , f_3 , f_4 , and f_5 .

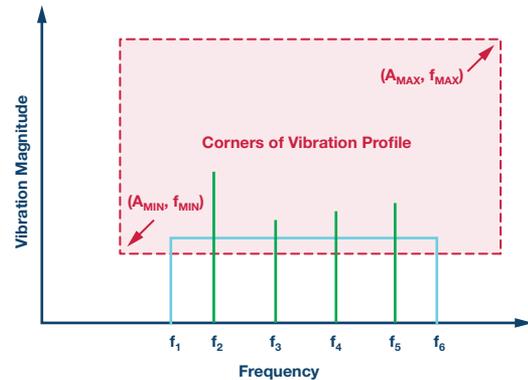


Figure 2. CM vibration profile examples.

System Requirements

Measurement range, frequency range (bandwidth), and resolution are three common attributes that often quantify the capability of a vibration sensing node. The red dashed lines in Figure 2 illustrate these attributes through a rectangular box that is bound by the minimum frequency (f_{MIN}), maximum frequency (f_{MAX}), minimum magnitude (A_{MIN}), and maximum magnitude (A_{MAX}). When considering a MEMS accelerometer for the role of the core sensor in a vibration sensing node, system architects will likely want to analyze its frequency response, measurement range, and noise behaviors fairly early in their design cycle. There are simple techniques for evaluating each of these accelerometer behaviors to predict the accelerometer's suitability for a given set of requirements. Obviously, system architects will eventually need to validate these estimates through actual validation and qualification, but even those efforts will value the expectation that comes from early analysis and predication of the accelerometer's capabilities.

Frequency Response

Equation 2 presents a simple, first-order model that describes a MEMS accelerometer's response (y) to linear acceleration (a) in the time domain. In this relationship, the bias (b) represents the value of the sensor's output when it is experiencing zero linear vibration (or any type of linear acceleration). The scale factor (K_A) represents the amount of change in the MEMS accelerometer's response (y), with respect to the change in linear acceleration (a).

$$y(t) = K_A \times a(t) + b \quad (2)$$

The frequency response of a sensor describes the value of the scale factor (K_A), with respect to frequency. In a MEMS accelerometer, the frequency response has two primary contributors: (1) response of its mechanical structure and (2) the response of the filtering in its signal chain. Equation 3 presents a generic, second-order model that presents an approximation for the mechanical portion of a MEMS accelerometer's response to frequency. In this model, f_0 represents the resonant frequency and Q represents the quality factor.

$$H_M(s) = \frac{\omega_0^2}{s^2 + \left(\frac{\omega_0}{Q}\right) \times s + \omega_0^2} \quad (3)$$

$$f_0 = \frac{\omega_0^2}{2\pi}$$

The contribution from the signal chain will often depend on the filtering that the application requires. Some MEMS accelerometers use a single-pole, low-pass filter to help lower the gain of the response at the resonant frequency. Equation 4 offers a generic model for the frequency response associated with this type of filter (H_{SC}). In this type of filter model, the cutoff frequency (f_c) represents the frequency at which the magnitude of the output signal is lower than its input signal by a factor of $\sqrt{2}$.

$$H_{SC}(s) = \frac{\omega_C}{s + \omega_C} \quad (4)$$

$$\omega_C = 2\pi f_C$$

Equation 5 combines the contributions of the mechanical structure (H_M) and the signal chain (H_{SC}).

$$H_T(s) = H_M(s) \times H_{SC}(s)$$

$$H_T(s) = \frac{\omega_0^2}{s^2 + \left(\frac{\omega_0}{Q}\right) \times s + \omega_0^2} \times \frac{\omega_C}{s + \omega_C} \quad (5)$$

Figure 3 provides a direct application of this model to predict the frequency response of the ADXL356 (x-axis). This model assumes a nominal resonant frequency of 5500 Hz, a Q of 17, and the use of a single-pole, low-pass filter that has a cutoff frequency of 1500 Hz. Note that Equation 5 and Figure 4 only describe the sensor's response. This model does not include consideration of the manner in which the accelerometer is coupled to the platform that it is monitoring.

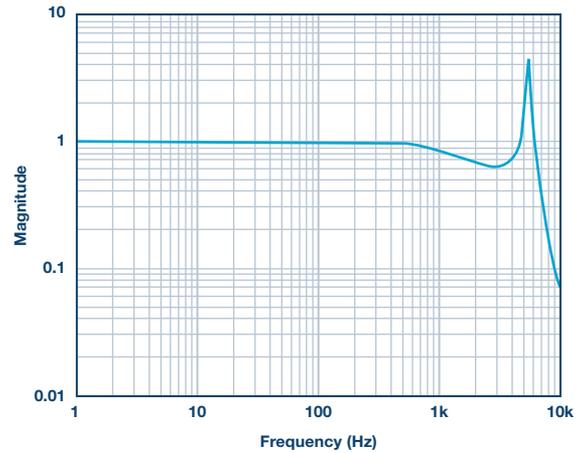


Figure 3. ADXL356 frequency response.

Bandwidth vs. Flatness

In signal chains that leverage a single-pole, low-pass filter (like the one in Equation 4) to establish their frequency response, their bandwidth specification often identifies the frequency at which its output signal is delivering 50% of the power of the input signal. In more complex responses, such as the third-order model from Equation 5 and Figure 3, bandwidth specifications will often come with a corresponding specification for the flatness attribute. The flatness attribute describes the change in the scale factor over the frequency range (bandwidth). Using the ADXL356's simulation from Figure 3 and Equation 5, the flatness at 1000 Hz is approximately 17% and at 2000 Hz, the flatness is ~40%.

While many applications will need to limit the bandwidth that they can use due to their flatness (accuracy) requirements, there are cases where this may not be as concerning. For example, some applications may be more focused on tracking relative changes over time, rather than absolute accuracy. Another example could come from those who will leverage digital postprocessing techniques to remove the ripple over the frequency ranges that they are most interested in. In these cases, the repeatability and stability of the response is often more important than the flatness of the response over a given frequency range.

Measurement Range

The measurement range metric for a MEMS accelerometer represents the maximum linear acceleration that the sensor can track in its output signal. At some linear acceleration level that is beyond the measurement range rating, the output signal of the sensor will saturate. When this happens it introduces significant distortion and makes it very difficult (if not impossible) to extract useful information from the measurements. Therefore, it is important to make sure that a MEMS accelerometer will support the peak acceleration levels (see A_{MAX} in Figure 2).

Note that the measurement range will have a dependence on frequency since the mechanical response of the sensor introduces some gain to the response, with the peak of the gain response happening at the resonant frequency. In the case of the simulated response for the ADXL356 (see Figure 3), the gain peaks at approximately 4 \times , which reduces the measurement range from $\pm 40 g$ to $\pm 10 g$. Equation 6 offers an analytical approach to predicting this same number, using Equation 5 as a starting point:

$$\begin{aligned}
A_{MAX}(5500 \text{ Hz}) &= \frac{A_{MAX}(0 \text{ Hz})}{H_A(5500 \text{ Hz})} \\
A_{MAX}(5500 \text{ Hz}) &= \frac{\pm 40 \text{ g}}{4} \\
A_{MAX}(5500 \text{ Hz}) &= \pm 10 \text{ g}
\end{aligned} \tag{6}$$

The large change in scale factor and reduction in measurement range are two reasons why most CBM systems will want to confine the maximum frequency of their vibration exposure to levels that are well below the resonant frequency of the sensor.

Resolution

“The resolution of an instrument may be defined as the smallest change in the environment that causes a detectable change in the indication of the instrument.”¹ In a vibration sensing node, noise in the acceleration measurement will have a direct influence on its ability to detect changes in vibration (aka resolution). Therefore, noise behaviors are an important consideration for those who are considering a MEMS accelerometer to detect small changes in the vibration on their machine platforms. Equation 7 provides a simple relationship for quantifying the impact that a MEMS accelerometer’s noise will have on its ability to resolve small change in vibration. In this model, the sensor’s output signal (y_M) is equal to the sum of its noise (a_N) and the vibration that it is experiencing (a_V). Since there will be no correlation between the noise (a_N) and the vibration (a_V), the magnitude of the sensor’s output signal ($|y_M|$) will be equal to the root sum square (RSS) combination of the noise magnitude ($|a_N|$) and the vibration’s magnitude ($|a_V|$).

$$\begin{aligned}
y_M(t) &= a_N(t) + a_V(t) \\
|y_M| &= \sqrt{|a_V|^2 + |a_N|^2}
\end{aligned} \tag{7}$$

So, what level of vibration is required to overcome the noise burden in the measurement and create an observable response in the sensor’s output signal? Quantifying the vibration level in terms of the noise level can help explore this question in an analytical manner. Equation 8 establishes this relationship through ratio (K_{VN}) and then derives a relationship to predict the level of change in the sensor’s output, in terms of that ratio:

$$\begin{aligned}
|a_V| &= K_{VN} \times |a_N| \\
|y_M| &= \sqrt{(K_{VN} \times |a_N|)^2 + |a_N|^2} \\
|y_M| &= \sqrt{(K_{VN}^2 + 1) \times |a_N|^2} \\
\frac{|y_M|}{|a_N|} &= \sqrt{(K_{VN}^2 + 1)}
\end{aligned} \tag{8}$$

Table 1 provides some numerical examples of this relationship to help illustrate the increase in the sensor’s output measurement, with respect to the ratio (K_{VN}) of the vibration and noise magnitudes. For simplicity, the remainder of this discussion will assume that the total noise in the sensor’s measurement will establish its resolution. From Table 1, this relates to the case where K_{VN} is equal to one, which is when the vibration magnitude is equal to the noise magnitude. When that happens, the magnitude in the sensor’s output will increase by 42% over its output magnitude when there is zero vibration. Note that each application may need to consider what level of increase will be observable in their system in order to establish a relevant definition for resolution in that situation.

Table 1. Sensors Response to Vibration/Noise

K_{VN}	$ y_M / a_N $	Increase %
0	1	0
0.25	10.3	3
0.5	1.12	12
1	1.41	41
2	2.23	123

Predicting Sensor Noise

Figure 4 presents a simplified signal chain of a vibration sensing node that will use a MEMS accelerometer. In most cases, the low-pass filter provides some support for antialiasing, while the digital processing will provide more defined boundaries in the frequency response. In general, these digital filters will seek to preserve the signal content that represents the real vibration, while minimizing the influence of out of band noise. Therefore, the digital processing will often be the most influential part of the system to consider when estimating the noise bandwidth. This type of processing can come in the form of time-domain techniques, such as a band-pass filter or through spectral techniques, such as the fast Fourier transform (FFT).

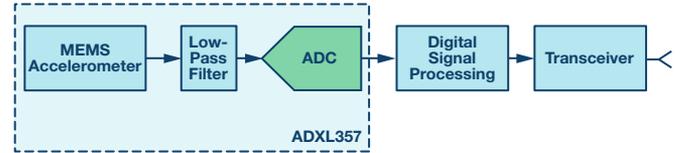


Figure 4. Signal chain of a vibration sensing node.

Equation 9 provides a simple relationship for estimating the total noise in a MEMS accelerometer’s measurement (A_{NOISE}), using its noise density (ϕ_{ND}) and the noise bandwidth (f_{NBW}) associated with the signal chain.

$$A_{NOISE} = \phi_{ND} \times \sqrt{f_{NBW}} \tag{9}$$

Using the relationship in Equation 9, we can estimate that when using a filter that has a noise bandwidth of 100 Hz on the ADXL357 (noise density = $80 \mu\text{g}/\sqrt{\text{Hz}}$), total noise will be 0.8 mg (rms).

Vibration in Terms of Velocity

Some CBM applications need to evaluate core accelerometer behaviors (range, bandwidth, noise) in terms of linear velocity. One method for making this translation starts with the simple model from Figure 1 and the same assumptions that produced the model in Equation 1: linear motion, single frequency, and zero mean displacement. Equation 10 expresses this model through a mathematical relationship for the instantaneous velocity (v_V) of the object in Figure 1. The magnitude of this velocity, expressed in terms of root mean square (rms), is equal to the peak velocity, divided by the square root of 2.

$$\begin{aligned}
v_V(t) &= V_{pk} \times \sin(2\pi f_V t) \\
V_{rms} &= \frac{V_{pk}}{\sqrt{2}}
\end{aligned} \tag{10}$$

Equation 11 takes the derivative of this relationship to produce a relationship for the instantaneous acceleration of the object in Figure 1:

$$\begin{aligned}
 a_V(t) &= d \frac{v_V(t)}{dt} \\
 a_V(t) &= d \frac{V_{pk} \times \sin(2 \times \pi \times f_V \times t)}{dt} \\
 a_V(t) &= 2 \times \pi \times f_V \times V_{pk} \times \cos(2 \times \pi \times f_V \times t)
 \end{aligned} \tag{11}$$

Starting with the peak value of the acceleration model from Equation 11, Equation 12 derives a new formula that relates the acceleration magnitude (A_{rms}) to the velocity magnitude (V_{rms}) and vibration frequency (f_V).

$$\begin{aligned}
 A_{rms} &= \frac{2 \times \pi \times f_V \times V_{pk}}{\sqrt{2}} \\
 A_{rms} &= 2 \times \pi \times f_V \times V_{rms} \\
 A_{PK} &= 2 \times \pi \times f_V \times V_{pk}
 \end{aligned} \tag{12}$$

Case Study

Let's bring this all together with a case study of the ADXL357, which expresses its range (peak) and resolution for a vibration frequency range of 1 Hz to 1000 Hz, in terms of linear velocity. Figure 5 provides graphical definition of several attributes that will contribute to this case study, starting with a plot of the ADXL357's noise density over the frequency range of 1 Hz to 1000 Hz. For the sake of simplicity in this discussion, all of the computations in this particular case study will assume that the noise density is constant ($\varphi_{ND} = 80 \mu g/\sqrt{Hz}$) over the entire frequency range. The red spectral plot in Figure 5 represents the spectral response of a band-pass filter and the green vertical line represents the spectral response of a single frequency (f_V) vibration, which is useful in developing velocity-based estimates of resolution and range.

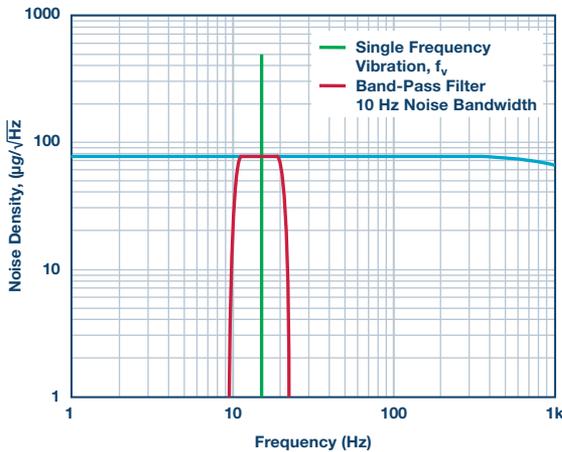


Figure 5. Case study noise density and filtering.

The first step in this process uses Equation 9 to estimate the noise (A_{NOISE}) that comes from four different noise bandwidths (f_{NBW}): 1 Hz, 10 Hz, 100 Hz, and 1000 Hz. Table 2 presents these results in terms of two different units of measure for linear acceleration: g and mm/s^2 . The use of g is fairly common in most MEMS accelerometer specifications tables, while vibration metrics are not often available in these terms. Fortunately, the relationship between g and mm/s^2 is fairly well known and is available in Equation 13.

$$1 g = 9.81 \frac{m}{s^2} \tag{13}$$

Table 2. Sensors Response to Vibration/Noise

f_{NBW} (Hz)	A_{NOISE}	
	(mg)	(mm/s ²)
1	0.08	0.78
10	0.25	2.48
100	0.80	7.84
1000	2.5	24.8

The next step in this case study rearranges the relationship in Equation 12 to derive a simple formula (see Equation 14) for translating the total noise estimates (from Table 2) into terms of linear velocity (V_{RES} , V_{PEAK}). In addition to offering the general form of this relationship, Equation 14 also offers one specific example, using the noise bandwidth of 10 Hz (and the acceleration noise of 2.48 mm/s^2 , from Table 2). The four dashed lines in Figure 6 represent the velocity resolution for all four noise bandwidths, with respect to the vibration frequency (f_V).

$$\begin{aligned}
 V_{RES}(f_{NBW}) &= \frac{A_{NOISE}(f_{NBW})}{2 \times \pi \times f_V} \\
 V_{RES}(10 \text{ Hz}) &= \frac{A_{NOISE}(10 \text{ Hz})}{2 \times \pi \times f_V} \\
 V_{RES}(10 \text{ Hz}) &= \frac{2.48 \frac{mm}{s^2}}{2 \times \pi \times f_V}
 \end{aligned} \tag{14}$$

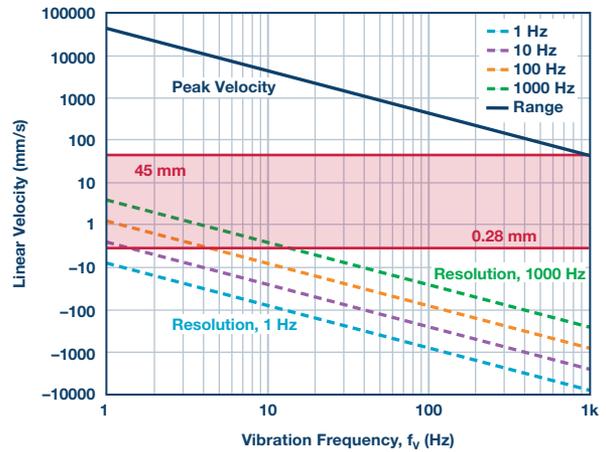


Figure 6. Peak and resolution vs. vibration frequency.

In addition to presenting the resolution for each bandwidth, Figure 6 also provides a solid blue line that represents the peak vibration levels (linear velocity) with respect to frequency. This comes from the relationship in Equation 15, which starts with the same general form as Equation 14, but instead of using the noise in the numerator it uses maximum acceleration that the ADXL357 can support. Note that the $\sqrt{2}$ factor in the numerator scales this maximum acceleration to reflect the rms level, assuming a single-frequency vibration model.

$$\begin{aligned}
 V_{RANGE} &= \frac{A_{RANGE}}{2 \times \pi \times f_V} \\
 V_{RANGE} &= \frac{1}{\sqrt{2}} \times \frac{\pm 40 g}{2 \times \pi \times f_V} \times \frac{9810 \frac{mm}{s^2}}{1 g}
 \end{aligned} \tag{15}$$

Finally, the red box represents how to apply this information to system-level requirements. The minimum (0.28 mm/s) and maximum (45 mm/s) velocity levels from this red box come from some of the classification levels in a common industry standard for machine vibration: ISO-10816-1. Overlaying the requirements on the range and resolution plots for the ADXL357 provides a quick method for making simple observations, such as:

- ▶ The worst case for the measurement range is at the highest frequency, where the $\pm 40 g$ range of the ADXL357 appears capable of measuring a very large portion of the vibration profiles associated with ISO-10816-1.
- ▶ When processing the ADXL357's output signal with a filter that has a noise bandwidth of 10 Hz filter, the ADXL357 appears capable of resolving the lowest vibration level from ISO-10816-1 (0.28 mm/s) across the frequency range of 1.5 Hz to 1000 Hz.
- ▶ When processing the ADXL357's output signal with a filter that has a noise bandwidth of 1 Hz filter, the ADXL357 appears capable of resolving the lowest vibration level from ISO-10816-1 across the entire 1 Hz to 1000 Hz frequency range.

Conclusion

MEMS accelerometers are coming of age as vibration sensors and they are playing a key role in what appears to be a perfect storm of technology convergence in CBM systems for modern factories. New solutions in sensing, connectivity, storage, analytics, and security are all coming together to provide factory managers with a fully integrated system of vibration observation and process feedback control. While it is easy to get lost in the excitement of all of this amazing technology advancement, someone still needs to understand how to relate these sensor measurements to real-world conditions and the implications that they represent. CBM developers and their customers will be able to draw value from these simple techniques and insights, which provide an approach for translating MEMS performance specifications into their impact on key system-level criteria using familiar units of measure.

References

1. Gerald C. Gill and Paul L. Hexter. "IEEE Transactions on Geoscience Electronics." *IEEE*, Vol. 11, Issue. 2, April, 1973.

Mark Looney [mark.looney@analog.com] is an applications engineer at Analog Devices in Greensboro, North Carolina. Since joining ADI in 1998, he has accumulated experience in high performance inertial MEMS technology, sensor-signal processing, high speed analog-to-digital converters, and dc-to-dc power conversion. He earned a B.S. (1994) and M.S. (1995) degree in electrical engineering from the University of Nevada, Reno, and has published numerous technical articles. Prior to joining ADI, he helped start IMATS, a vehicle electronics and traffic solutions company, and worked as a design engineer for Interpoint Corporation.



Mark Looney

Also by this Author:

[Designing for Low Noise Feedback Control with MEMS Gyroscopes](#)
Volume 50, Number 2

Massive MIMO and Beamforming: The Signal Processing Behind the 5G Buzzwords

By Claire Masterson

Share on   

Introduction

Our thirst for high speed mobile data is insatiable. As we saturate the available RF spectrum in dense urban environments, it's becoming apparent that there's a need to increase the efficiency of how we transmit and receive data from wireless base stations.

Base stations consisting of large numbers of antennas that simultaneously communicate with multiple spatially separated user terminals over the same frequency resource and exploit multipath propagation are one option to achieve this efficiency saving. This technology is often referred to as massive MIMO (multiple-input, multiple-output). You may have heard massive MIMO described as beamforming with a large number of antennas. But this raises the question ... what is *beamforming*?

Beamforming vs. Massive MIMO

Beamforming is a word that means different things to different people. Beamforming is the ability to adapt the radiation pattern of the antenna array to a particular scenario. In the cellular communications space, many people think of beamforming as steering a lobe of power in a particular direction toward a user, as shown in Figure 1. Relative amplitude and phase shifts are applied to each antenna element to allow for the output signals from the antenna array to coherently add together for a particular transmit/receive angle and destructively cancel each other out for other signals. The spatial environment that the array and user are in is not generally considered. This is indeed beamforming, but is just one specific implementation of it.

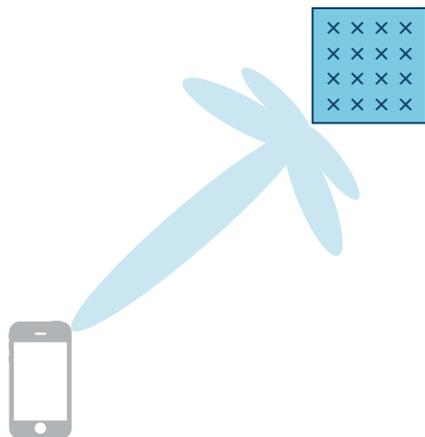


Figure 1. Traditional beamforming.

Massive MIMO can be considered a form of beamforming in the more general sense of the term, but is quite removed from the traditional form. *Massive* simply refers to the large number of antennas in the base station antenna array. *MIMO* refers to the fact that multiple spatially separated users are catered for by the antenna array in the same time and frequency resource. *Massive MIMO* also acknowledges that in real-world systems, data transmitted between an antenna and a user terminal—and vice versa—undergoes filtering from the surrounding environment. The signal may be reflected off buildings and other obstacles, and these reflections will have an associated delay, attenuation, and direction of arrival, as shown in Figure 2. There may not even be a direct line of sight between the antenna and the user terminal. It turns out that these nondirect transmission paths can be harnessed as a power for good.

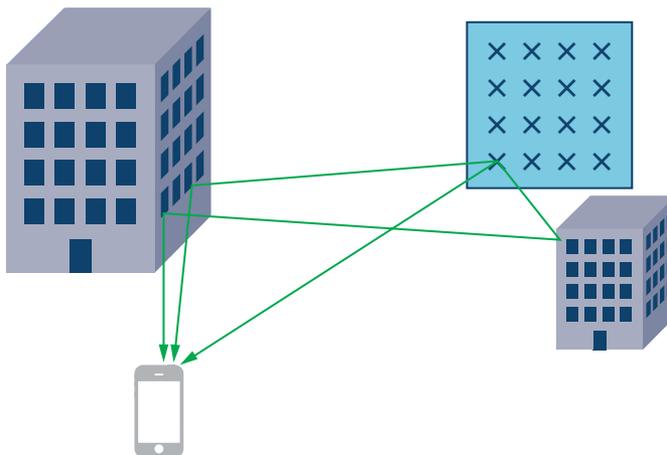


Figure 2. Multipath environment between antenna array and user.

In order to take advantage of the multiple paths, the spatial channel between antenna elements and user terminals needs to be characterized. In literature, this response is generally referred to as channel state information (CSI). This CSI is effectively a collection of the spatial transfer functions between each antenna and each user terminal. This spatial information is gathered in a matrix (H), as shown in Figure 3. The next section looks at the concept of CSI and how it is collected in more detail. The CSI is used to digitally encode and decode the data transmitted from and received by the antenna array.

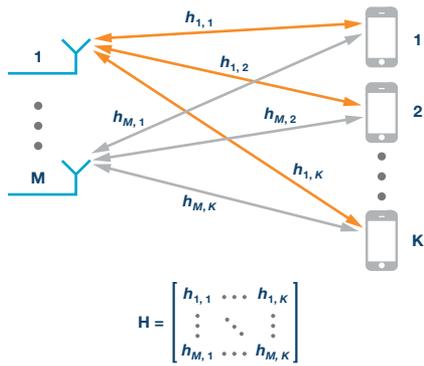


Figure 3. Channel state information needed to characterize a massive MIMO system.

Characterizing the Spatial Channel Between Base Station and User

An interesting analogy is to consider a balloon being popped at one location and the sound of this pop, or impulse, being recorded at another, as shown in Figure 4. The sound recorded at the microphone position is a spatial impulse response that contains information unique to the particular position of both the balloon and the microphone in the surrounding environment. The sound that is reflected off obstacles is attenuated and delayed compared to the direct path.

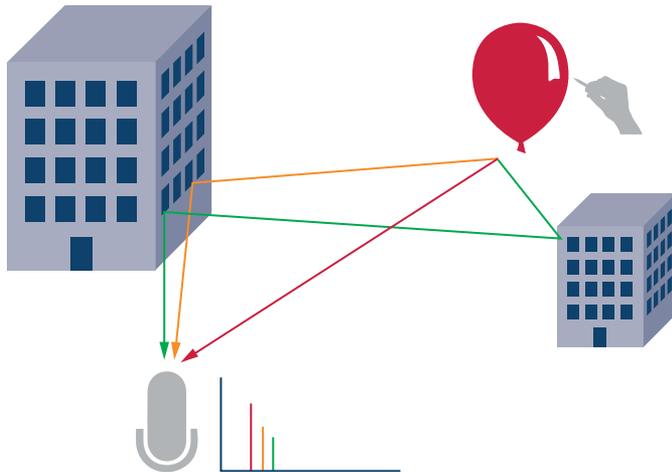


Figure 4. Audio analogy to demonstrate spatial characterization of a channel.

If we expand the analogy to compare to the antenna array/user terminal case, we need more balloons, as seen in Figure 5. Note that in order to characterize the channel between each balloon and the microphone, we need to burst each balloon at a separate time so the microphone doesn't record the reflections for different balloons overlapping. The other direction also needs to be characterized, as shown in Figure 6. In this instance, all the recordings can be done simultaneously when the balloon is popped at the user terminal position. This is clearly a lot less time consuming!

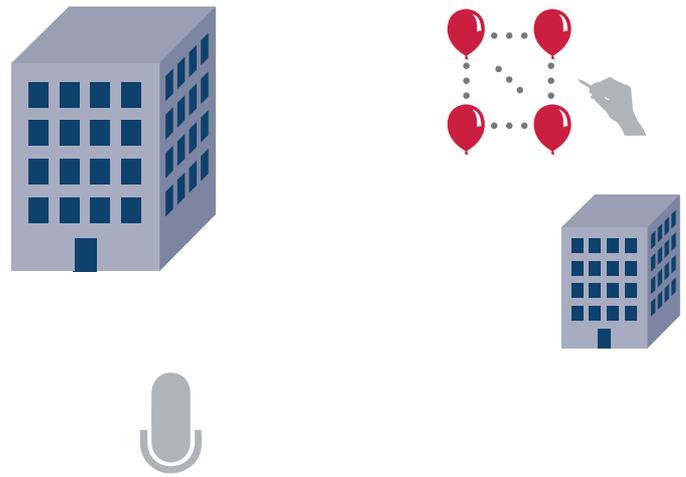


Figure 5. Audio analogy to downlink channel characterization.

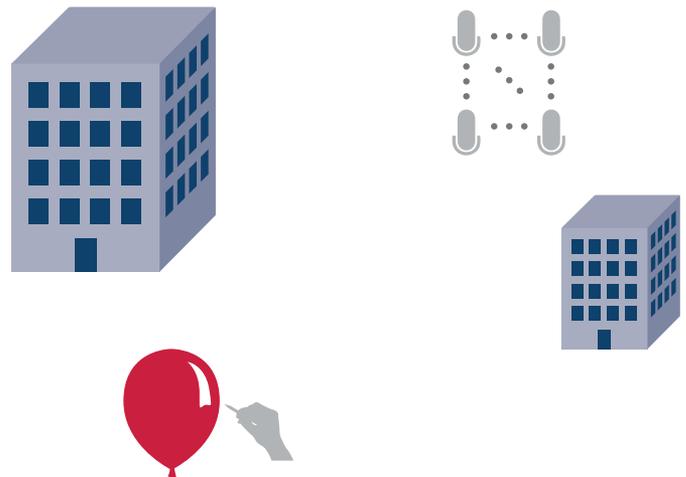


Figure 6. Audio analogy to uplink channel characterization.

In the RF space, pilot signals are used for characterizing the spatial channels. The over-the-air transmission channels between antennas and user terminals are reciprocal, meaning the channel is the same in both directions. This is contingent on the system operating in time division duplex (TDD) mode as opposed to frequency division duplex (FDD) mode. In TDD mode, uplink and downlink transmissions use the same frequency resource. The reciprocity assumption means the channel only needs to be characterized in one direction. The uplink channel is the obvious choice, as just one pilot signal needs to be sent from the user terminal and is received by all antenna elements. The complexity of the channel estimation is proportional to the number of user terminals, not the number of antennas in the array. This is of critical importance given the user terminals may be moving, and hence the channel estimation will need to be performed frequently. Another significant advantage of uplink-based characterization means that all the heavy duty channel estimation and signal processing is done at the base station, and not at the user end.

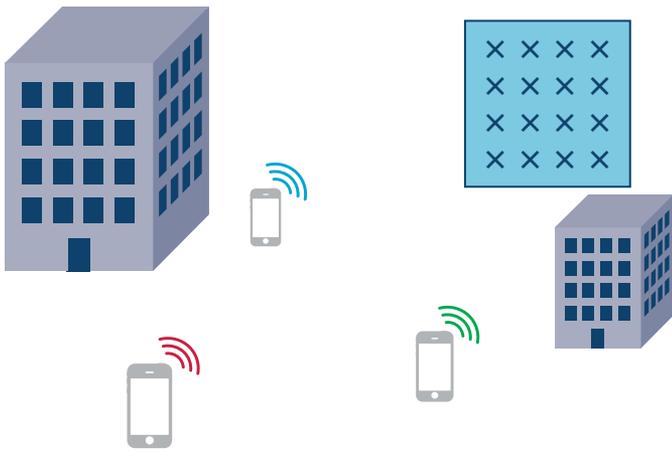


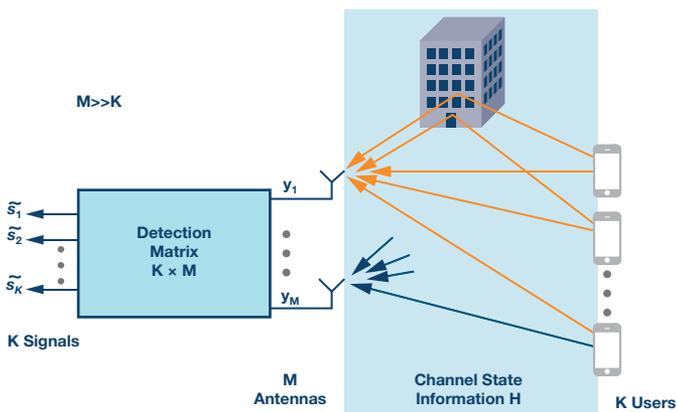
Figure 7. Each user terminal transmits orthogonal pilot symbol.

So now that the concept of collecting CSI has been established, how is this information applied to data signals to allow for spatial multiplexing? Filtering is designed based on the CSI to precode the data transmitted from the antenna array so that multipath signals will coherently add at the user terminals position. Such filtering can also be used to linearly combine the data received by the antenna array RF paths so that the data streams from different users can be detected. The following section addresses this in more detail.

The Signal Processing that Enables Massive MIMO

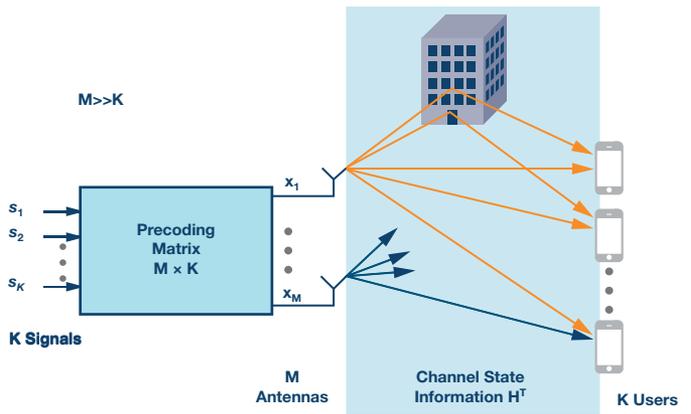
In the previous section we've described how the CSI (denoted by the matrix H) is estimated. Detection and precoding matrices are calculated based on H . There are a number of methods for calculating these matrices. This article focuses on linear schemes. Examples of linear precoding/detection methods are maximum ratio (MR), zero forcing (ZF), and minimum mean-square error (MMSE). Full derivations of the precoding/detection filters from the CSI are not provided in this article, but the criteria they optimize for, as well as the advantages and disadvantages of each method are discussed. A more detailed treatment of these topics can be found in the references at the end of this article.^{1, 2, 3}

Figure 8 and Figure 9 give a description of how the signal processing works in the uplink and downlink respectively for the three linear methods previously mentioned. For precoding there may also be some scaling matrix to normalize the power across the array that has been omitted for simplicity.



Detection Type	
Maximum Ratio (MR)	$\tilde{s} = H^H y$
Zero Forcing (ZF)	$\tilde{s} = (H^H H)^{-1} H^H y$
NMSE or RZF	$\tilde{s} = (H^H H + \beta I)^{-1} H^H y$

Figure 8. Uplink signal processing. H denotes the conjugate transpose.



Precoding Type	
Maximum Ratio (MR)	$x = H^T s$
Zero Forcing (ZF)	$x = H^T (H^T H)^{-1} s$
NMSE or RZF	$x = H^T (H^T H + \beta I)^{-1} s$

Figure 9. Downlink signal processing. T denotes the transpose. * denotes the conjugate.

Maximum ratio filtering, as the name suggests, aims to maximize the signal-to-noise ratio (SNR). It is the simplest approach from a signal processing viewpoint, as the detection/precoding matrix is just the conjugate transpose or conjugate of the CSI matrix, H . The big downside of this method is that the interuser interference is ignored.

Zero forcing precoding attempts to address the interuser interference problem by designing the optimization criteria to minimize for it. The detection/precoding matrix is the pseudoinverse of the CSI matrix. Calculating the pseudoinverse is more computationally expensive than the complex conjugate as in the MR case. However, by focusing so intently on minimizing the interference, the received power at the user suffers.

MMSE tries to strike a balance between getting the most signal amplification and reducing the interference. This holistic view comes with signal processing complexity as a price tag. The MMSE approach introduces a regularization term to the optimization—denoted as β in Figures 8 and 9—that allows for a balance to be found between the noise covariance and the transmit power. It is sometimes also referred to in literature as regularized zero forcing (RZF).

This is not an exhaustive list of precoding/detection techniques, but gives an overview of the main linear approaches. There are also nonlinear signal processing techniques such as dirty paper coding and successive interference cancellation that can be applied to this problem. These offer optimal capacity but are very complex to implement. The linear approaches described above are generally sufficient for massive MIMO, where the number of antennas gets large. The choice of a precoding/detection technique will depend on the computational resources, the number of antennas, the number of users, and the diversity of the particular environment the system is in. For large antenna arrays where the number of antennas is significantly greater than the number of users, the maximum ratio approach may well be sufficient.

The Practical Obstacles Real-World Systems Present to Massive MIMO

When massive MIMO is implemented in a real-world scenario, there are further practical considerations to be taken into account. Consider an antenna array with 32 transmit (Tx) and 32 receive (Rx) channels operating in the 3.5 GHz band as an example. There are 64 RF signal chains to be put in place and the spacing between the antennas is approximately 4.2 cm

given the operating frequency. That's a lot of hardware to pack into a small space. It also means there is a lot of power being dissipated, which brings inevitable temperature concerns. Analog Devices' integrated transceivers offer a highly effective solution to many of these issues. The AD9371 will be discussed in more detail in the next section.

Previously in this article, the application of reciprocity to the system to drastically cut the channel estimation and signal processing overheads were discussed. Figure 10 shows the downlink channel in a real-world system. It is split into three components; the over-the-air channel (H), the hardware response of the base station transmit RF paths (T_{BS}), and the hardware response of the user receive RF paths (R_{UE}). The uplink is the opposite of this with R_{BS} characterizing the base station receive hardware RF paths and T_{UE} characterizing the users transmit hardware RF paths. While the reciprocity assumption holds for the over the air interface, it does not for the hardware paths. The RF signal chains introduce inaccuracies into the system due to mismatched traces, poor synchronization between the RF paths, and temperature-related phase drift.

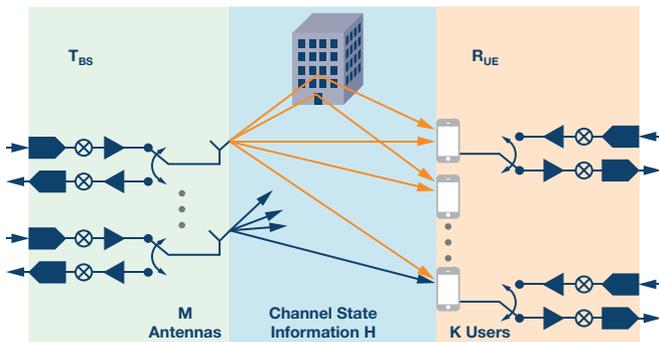


Figure 10. Real-world downlink channel.

Using a common synchronized reference clock for all LO (local oscillator) PLLs in the RF paths and synchronized SYSREFs for the baseband digital JESD204B signals will help address latency concerns between the RF paths. However, there will still be some arbitrary phase mismatch between the RF paths at system startup. Temperature-related phase drift contributes further to this issue and it is clear that calibration is required in the field when the system is initialized and periodically thereafter. Calibration allows for the advantages of reciprocity, such as maintaining the signal processing complexity at the base station and uplink only channel characterization to be kept. It can generally be simplified so that only the base station RF paths (T_{BS} and R_{BS}) need to be considered.

There are a number of approaches to calibrating these systems. One is to use a reference antenna positioned carefully in front of the antenna array to calibrate both the receive and transmit RF channels. It's questionable whether having an antenna placed in front of the array in this way is suited to practical base station calibration in the field. Another is to use mutual coupling between the existing antennas in the array as the calibration mechanism. This may well be feasible. The most straight forward approach is probably to add passive coupling paths just before the antennas in the base station. This adds more complexity in the hardware domain, but should provide a robust calibration mechanism. To fully calibrate the system a signal is sent from one designated calibration transmit channel, which is received by all RF receive paths through the passive coupled connection. Each transmit RF path then sends a signal in sequence that is picked up at the passive coupling point before each antenna, relayed back to a combiner, and then to a designated calibration receive path. Temperature related effects are generally slow to change, so this calibration does not have to be performed very frequently, unlike the channel characterization.

Analog Devices' Transceivers and Massive MIMO

Analog Devices' range of integrated transceiver products are particularly suited to applications where there is a high density of RF signal chains required. AD9371 features 2 transmit paths, 2 receive paths, and an observation receiver, as well as three fractional-N PLLs for RF LO generation in a 12 mm × 12 mm package. This unrivaled level of integration enables manufacturers to create complex systems in a timely and cost-effective manner.

A possible system implementation featuring multiple AD9371 transceivers is shown in Figure 11. This is a 32 transmit, 32 receive system with 16 AD9371 transceivers. Three AD9528 clock generators provide the PLL reference clocks and JESD204B SYSREFs to the system. The AD9528 is a 2-stage PLL with 14 LVDS/HSTL outputs and an integrated JESD204B SYSREF generator for multiple device synchronization. The AD9528s are arranged in a fanout buffer configuration with one acting as a master device with some of its outputs used to drive the clock inputs and the SYSREF inputs of the slave devices. A possible passive calibration mechanism is included—shown in green and orange—where a dedicated transmit and receive channel are used to calibrate all the receive and transmit signal paths through a splitter/combiner, as discussed in the previous section.

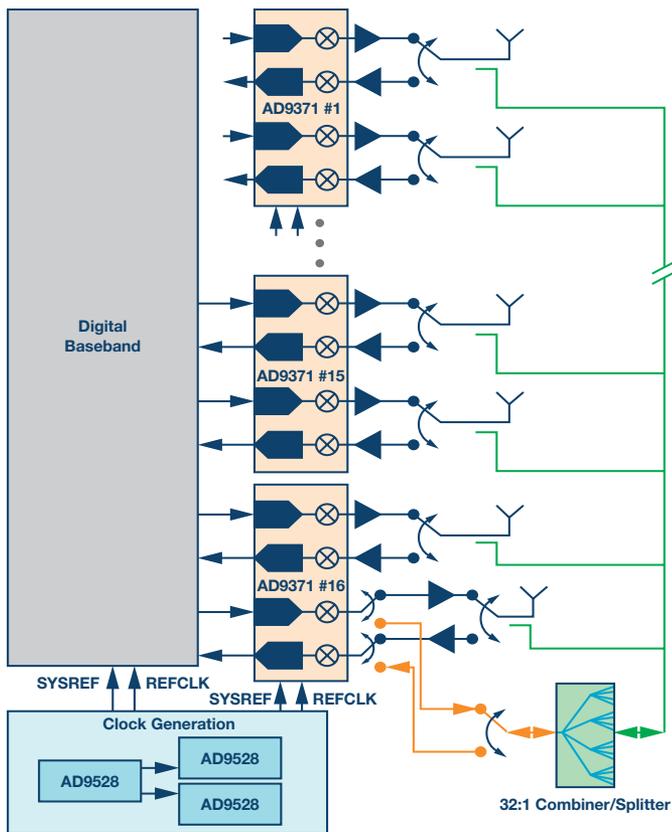


Figure 11. Block diagram of 32 Tx, 32 Rx massive MIMO radio head featuring Analog Devices' AD9371 transceivers.

Conclusion

Massive MIMO spatial multiplexing has the potential to become a game changing technology in the cellular communications space, allowing for increased cellular capacity and efficiency in high traffic urban areas. The diversity that multipath propagation introduces is exploited to allow for data transfer between a base station and multiple users in the same time and frequency resource. Due to reciprocity of the channel between the base station antennas and the users, all the signal processing complexity can be kept at the base station, and the channel characterization can be done in the uplink. Analog Devices' RadioVerse™ family of integrated transceiver products allow for a high density of RF paths in a small space, so they are well suited to massive MIMO applications.

References

1. Xiang Gao. *Massive MIMO in Real Propagation Environments*. Lund University, 2016.
2. Michael Joham, Josef A. Nossek, and Wolfgang Utschick. "Linear Transmit Processing in MIMO Communications Systems." *IEEE Transactions on Signal Processing*, Vol. 53, Issue 8, Aug, 2005.
3. Hien Quoc Ngo. *Massive MIMO: Fundamentals and System Design*. Linköping University, 2015.

Claire Masterson [claire.masterson@analog.com] is a systems applications engineer in the Communication Systems Team at Analog Devices Limerick working on systems implementation, software development, and algorithm development. Claire received a BAI and PhD from Trinity College in Dublin and joined ADI in 2011 after graduating.



Claire Masterson

Rarely Asked Questions—Issue 142

Orthogonal Perspectives

By Ian Beavers

Share on   



Question:

I am using a MEMS inertial measurement unit (IMU) in a self-balancing guidance control system for a personal transportation platform. Can I expect a consumer targeted IMU to eliminate all misalignment errors between each sensor if all of the core sensor elements are on a single piece of silicon?

Answer:

No, this is generally not a safe expectation for your design. Industrial grade IMUs, which use robust discrete sensors with optimal packaging and calibration, offer much better alignment precision than consumer-targeted IMUs residing on a single piece of silicon.

Consumer targeted and industrial targeted IMUs tend to specify axis alignment behaviors differently. Consumer IMUs typically lump all misalignment errors into a single cross-axis sensitivity specification. Industrial targeted IMUs, such as the recently released [ADIS16490](#), specify alignment precision more directly using two different specifications: axis-to-axis misalignment error and axis-to-package misalignment error. The axis-to-package misalignment error describes how well the alignment in each axis relates to mechanical features within the IMU package. Axis-to-axis misalignment error describes how well the alignment of each accelerometer and gyroscope axis fits into the ideal case of mutual orthogonality. This is why axis-to-axis misalignment error is also commonly known as orthogonal error.

We can see the mathematical relationship between cross-axis sensitivity (CAS) and axis-to-axis misalignment error (A2A_MAE) as described below:

$$\text{CAS} = \sin(\text{A2A_MAE})$$

$$\text{A2A_MAE} = \text{asin}(\text{CAS})$$

The effect of nonorthogonality occurs between sensor axes, across sensors, or from package misalignment between the sensors and the enclosure. On an industrial targeted IMU, these specifications are fully described in the data sheet after factory calibration. For discrete components, the cross-axis sensitivity specification does not account for assembly variances to a PCB.

Ideally, multiple axes within gyroscopes and accelerometers are mutually orthogonal to each other. However, it is a common misconception that since a multi-axis gyroscope or accelerometer can be designed within one discrete MEMS component that each of the axes are perfectly orthogonal at 90° to one another. Although all inertial sensors in these devices are on a single piece of silicon, inherent errors introduced from fabrication and manufacturing variances can still accumulate an orthogonal error. The resulting equivalent alignment precision is actually not very impressive when compared to fully calibrated, industrial targeted IMUs.

A quick survey of consumer targeted devices reveals that cross-axis sensitivity is often in the range of 1% to 5%. Using the above relationship, that results in equivalent axis-to-axis misalignment errors of 0.57° to 2.87°. However, it could also be defined in units of milliradian, equal to 0.057°. Industrial grade IMUs will typically be much more precise. We can also use this relationship to translate the axis-to-axis misalignment error of an industrial targeted IMU of 0.018° into an equivalent cross-axis sensitivity of 0.031%.

$$\text{CAS} = \sin(\text{A2A_MAE}) = \sin(0.018^\circ) = 0.00031 = 0.031\%$$

Despite the apparent disadvantage of not having all inertial sensors on one piece of silicon, the ADIS16489 industrial grade IMU still offers ~32× better performance than the best consumer devices.

To understand the effect of nonorthogonal errors, let's assume that one accelerometer axis is pointed perfectly upward and the device is exactly level. The accelerometer on this z-axis is ideally measuring the total impact of gravity. If the other two axes were perfectly orthogonal, they would not measure any vector of gravity. However, if there is a nonorthogonality error, these two other horizontal axes would measure some portion of the gravity

vector. For example, if a device offers a cross-axis sensitivity of 1%, its equivalent response to gravity will be 10 mg. This equates to an equivalent alignment error of 0.6°. Conversely, if the first axis is not orthogonal to the level frame, it will measure less than the complete gravity vector.

Orthogonality errors are particularly stable components of the total error from an accelerometer. They may therefore yield to corrections based on a one time calibration. To determine the orthogonality error of accelerometer axis pairs, the static response of each axis to gravity is measured as the accelerometer is rotated through the space of all possible 90° orientations. This can be done using either a precision gimbal mount or on a known orthogonal surface.

It can be a challenging proposition to effectively calibrate out the orthogonal errors across the full operating conditions after mounting components onto a PCB. Inertial calibration requires observation of each sensor response, while the devices are experiencing well-controlled motion profiles. These types of motion profiles often require highly specialized equipment and

expertise to operate effectively over time. In contrast to an industrial targeted IMU that is already precalibrated for mounting, each mounted consumer MEMS device on a PCB would need to be calibrated against the other sensors, environmental performance, and temperature.

Performance from an industrial IMU, composed of 3 gyroscope axes and 3 accelerometer axes, leverages a calibration step within manufacturing after discrete components are mounted on a mini-PCB in a rugged module. This single factory calibration identifies and compensates not only the non-orthogonality of the MEMS devices themselves, but also for any assembly related skew. This minimizes the errors associated with variances from assembly, cross-axis error, and temperature. The ADIS16489 provides factory calibration that minimizes axis alignment errors within platform stabilization, navigation, or robotics applications. With a digital tri-axis gyroscope and a tri-axis accelerometer, the ADIS16489 offers a mere $\pm 0.018^\circ$ axis-to-axis gyroscope misalignment error and $\pm 0.035^\circ$ accelerometer axis-to-axis error. In addition to the high performance sensor parameters, the ADIS16489 also provides a parylene coating as a moisture barrier for its internal circuitry.

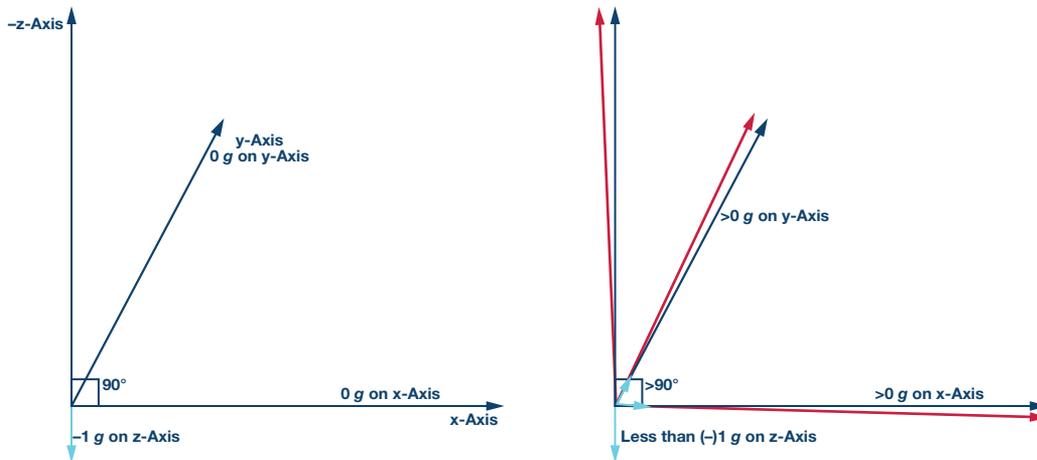


Figure 1. An ideal 3-axis orthogonal case on the left reflects the true impact of a vector. A nonorthogonal error allows leakage of rotation or force to be seen across all axes.

Ian Beavers [ian.beavers@analog.com] is a product engineering manager for the Automation Energy and Sensors Team at Analog Devices (Greensboro, NC). He has worked for the company since 1999. Ian has over 19 years of experience in the semiconductor industry. Ian earned a bachelor's degree in electrical engineering from North Carolina State University and an M.B.A. from the University of North Carolina at Greensboro.



Ian Beavers

Also by this Author:

[The Case of the Misguided Gyro](#)

Volume 51, Number 1

Transmit LO Leakage (LOL)—An Issue of Zero-IF That Isn't Making People Laugh Out Loud

By Dave Frizelle

Share on   

Introduction

There are several major advantages to zero-IF architecture. However, there are also some challenges that need to be overcome. Transmit local oscillator leakage (referred to as transmit LOL) is one such challenge. Uncorrected, transmit LOL will produce an unwanted emission within the desired transmission, potentially breaking system specifications. This article discusses the issue of transmit LOL and examines the techniques used to eliminate it, as implemented in ADI's RadioVerse transceiver family (which includes the AD9371; see [ADI RadioVerse Website](#) for more details). If transmit LOL can be reduced to a low enough level that it no longer causes system or performance issues, perhaps people can learn to laugh out loud about LOL!

What Is LOL?

An RF mixer has two input ports and one output port, as shown in Figure 1. The ideal mixer would produce an output that is the product of the two inputs. In frequency terms, the output should be $F_{IN} + F_{LO}$ and $F_{IN} - F_{LO}$, nothing else. If either input is undriven there will be no output.

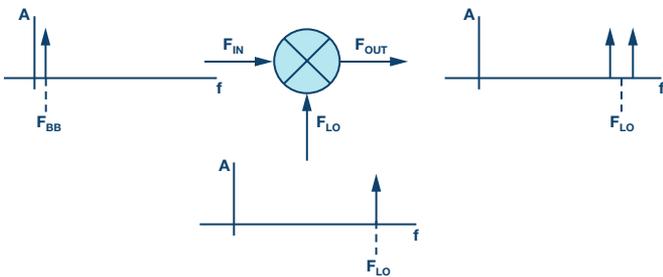


Figure 1. Ideal mixer.

In Figure 1, F_{IN} is set to F_{BB} with a baseband frequency of 1 MHz and F_{LO} is set to F_{LO} with a local oscillator frequency of 500 MHz. If the mixer were ideal it would produce an output that comprises two tones: one at 499 MHz and one at 501 MHz. However, as shown in Figure 2, a real-world mixer will also produce some energy at F_{BB} and F_{LO} . The energy at F_{BB} can be ignored because it is far away from the desired output and will be filtered out by the RF components located after the mixer output. Regardless of the energy at F_{BB} , the energy at F_{LO} can be a problem. It is very close to or within the desired output signal and difficult or impossible to remove

by filtering, since the filtering would also filter the desired signal. It is this unwanted energy at F_{LO} that is referred to as LOL. The local oscillator (LO), which drives the mixer, has leaked to the mixer's output port. There are also other paths for the LO to leak to the system output, such as through power supplies or across the silicon itself. Regardless of how the LO leaks out, it can be referred to as LOL.

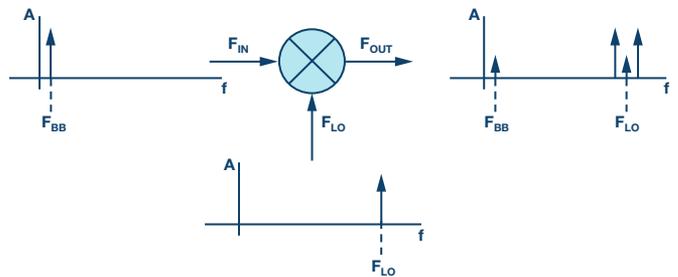


Figure 2. Real-world mixer.

In a real-IF architecture where only one sideband is to be transmitted, it is possible to resolve LOL by using RF filtering. In contrast, in a zero-IF architecture where both sidebands are to be transmitted, the LOL sits in the middle of the desired output and presents a more difficult challenge (see Figure 3). Conventional filtering is no longer an option, because any filtering that would remove the LOL would also remove portions of the wanted transmission. Therefore, other techniques must be used to eliminate it. Otherwise, it will likely end up becoming an unwanted emission within the overall desired transmission.

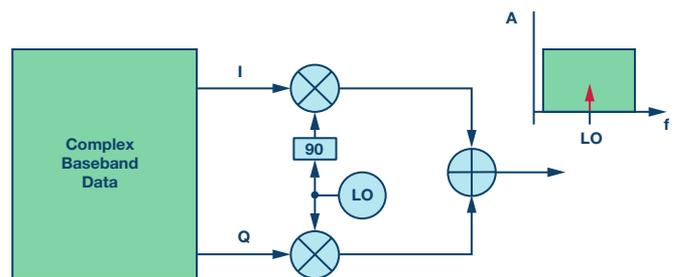


Figure 3. Unwanted energy at F_{LO} shown in red. This unwanted energy at F_{LO} is called LOL.

Eliminating LO Leakage (also Called LOL Correction)

The elimination of LOL is achieved by generating a signal that is equal in amplitude but opposite in phase to the LOL, thus cancelling it, as shown in Figure 4. Assuming we know the exact amplitude and phase of the LOL, the cancellation signal can be generated by applying dc offsets to the transmitter's inputs.

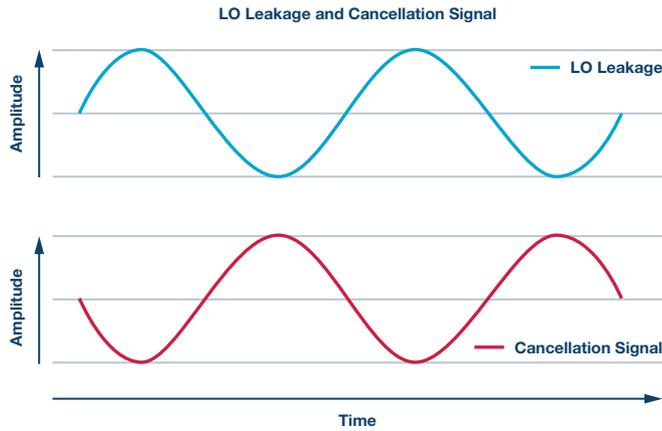


Figure 4. LO leakage and cancellation signals.

Generation of the Cancellation Signal

The complex mixer architecture lends itself well to the generation of the cancellation signal. Because quadrature signals at the LO frequency exist in the mixer (they are at the heart of how the complex mixer works),¹ they allow for the generation of a signal at the LO frequency with any phase and amplitude.

The quadrature signals that drive the complex mixer can be described as $\sin(\text{LO})$ and $\cos(\text{LO})$ —these are orthogonal signals at the LO frequency that drive the two mixers. To generate the cancellation signal, these orthogonal signals are added together with different weights. In mathematical terms, we can produce an output that is $I \times \sin(\text{LO}) + Q \times \cos(\text{LO})$. By applying different signed values in place of I and Q , the resulting sum will be at the LO frequency and can have any desired amplitude and phase. Examples are shown in Figure 5.

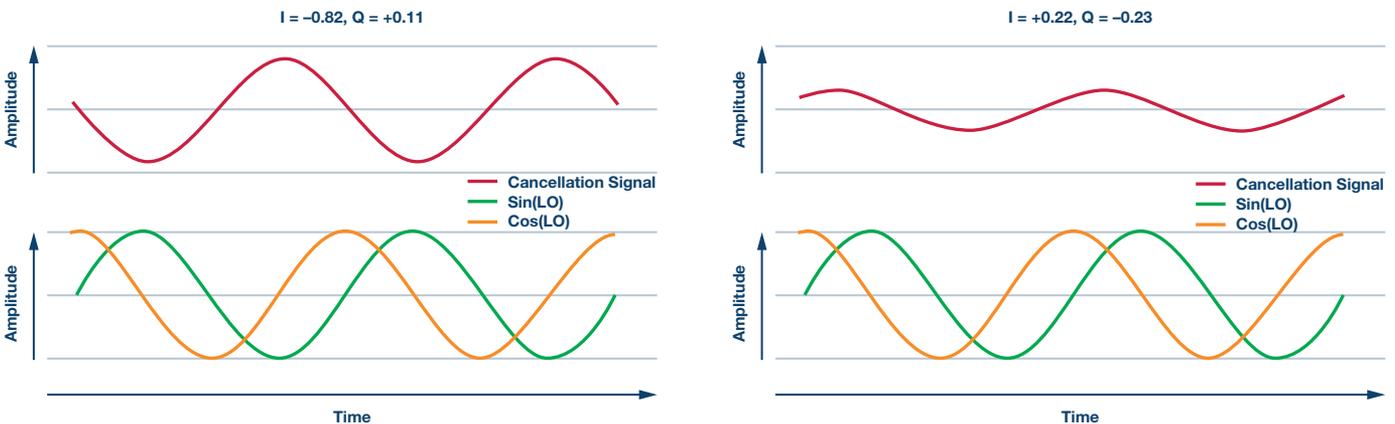


Figure 5. Examples of any phase and any amplitude cancellation signal being generated.

The desired transmission signal will need to be applied to the transmitter's inputs. By applying a dc bias to the transmission data, the output of the mixer will contain both the desired transmission signal and also the desired LOL cancellation signal. The intentionally generated cancellation signal will combine with the unwanted LOL and they will cancel, leaving only the desired transmission signal.

Observing the Transmit LOL

The transmit LOL is observed using an observation receiver, as shown in Figure 6. In this example, the observation receiver uses the same LO as the transmitter, so any transmit energy at the LO frequency will appear as dc at the output of the observation receiver.

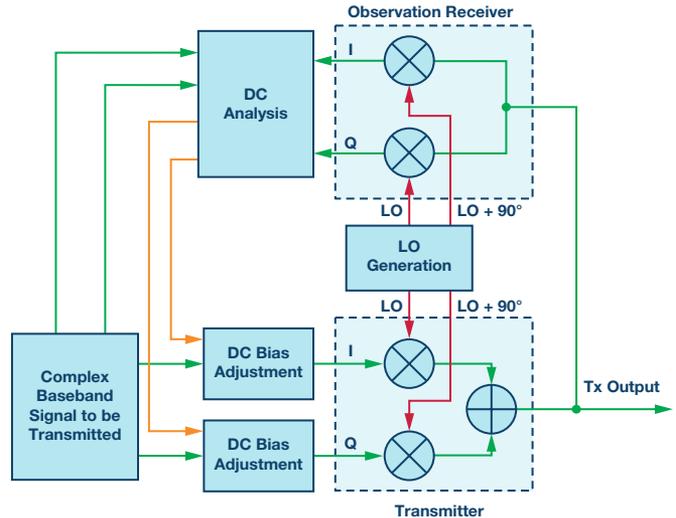


Figure 6. Basic concepts of observation and correction of Tx_{LO} leakage.

The approach shown in Figure 6 has an inherent weakness: by using the same LO to transmit and observe, transmit LOL will appear as dc in the observation receiver's output. The observation receiver itself will have some amount of dc due to component mismatch in the circuit, so the total dc output by the observation receiver will be the sum of the transmit LOL

and the native dc offset that exists in the observation receiver. There are ways to overcome this issue, but a better approach is to use a different LO frequency for observation, thereby separating the native dc in the observation path from the transmit LOL separating result. This is shown in Figure 7 below.

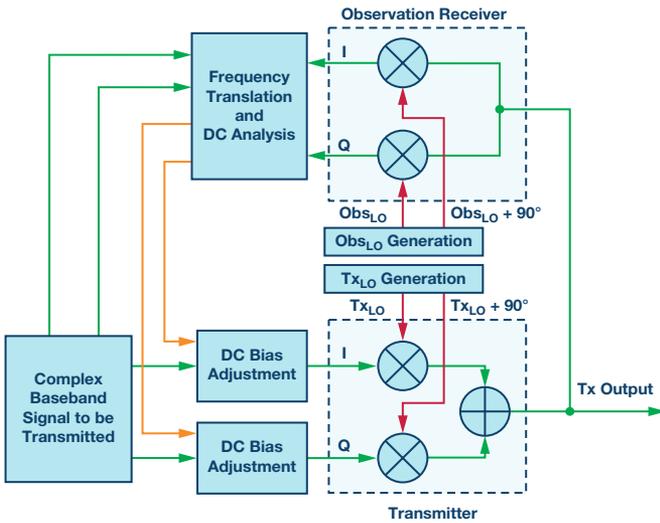


Figure 7. Using different LOs to transmit and observe.

Because the transmission is being observed using a frequency other than transmit LO, energy at the transmit LO frequency will not appear at dc in the observation receiver. Instead, it will appear as a baseband tone whose frequency is equal to the difference between the transmit LO and the observation LO. DC native in the observation path will still appear at dc, so there will be total separation of observation dc and transmit LOL measurement results. Figure 8 illustrates this concept using single-mixer architecture for simplicity. The input to the transmitter is transmit LOL. Frequency shifting is done after the observation receiver to move the transmit LOL observed energy to dc.

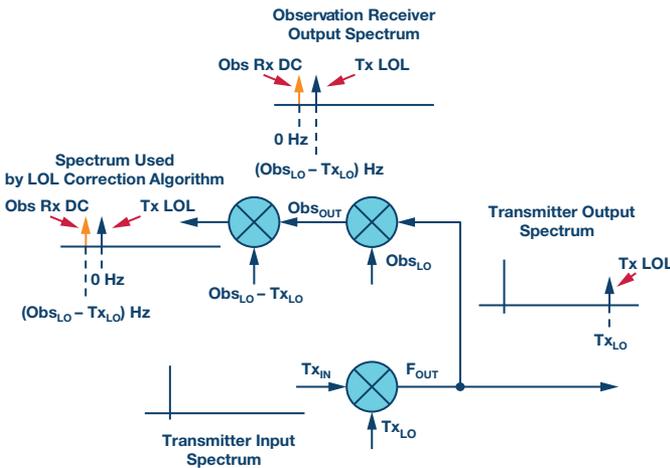


Figure 8. Separating observation receiver dc from Tx LOL.

Finding the Necessary Correction Values

The required correction values are found by taking the observation receiver's output, dividing it by the transfer function from transmit input to observation receiver output, and comparing this result to the intended transmission. The transfer function in question is shown in Figure 9.

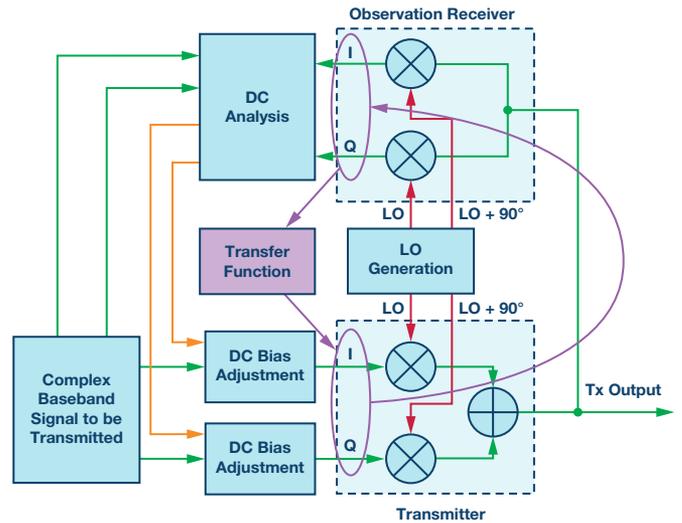


Figure 9. The transfer function from transmitter input to observation receiver output.

The transfer function from the transmitter baseband input to the observation receiver baseband output is comprised of two components: amplitude scaling and phase rotation. Each is explained independently in more detail in the following sections.

Figure 10 shows that the amplitude of the transmit signal reported by the observation receiver may not represent the actual amplitude of the transmit signal being transmitted if the loopback path from transmit output to observation receiver input has gain or attenuation in the path, or if the gain of the transmitter circuit were different from the gain of the observation receiver circuit.

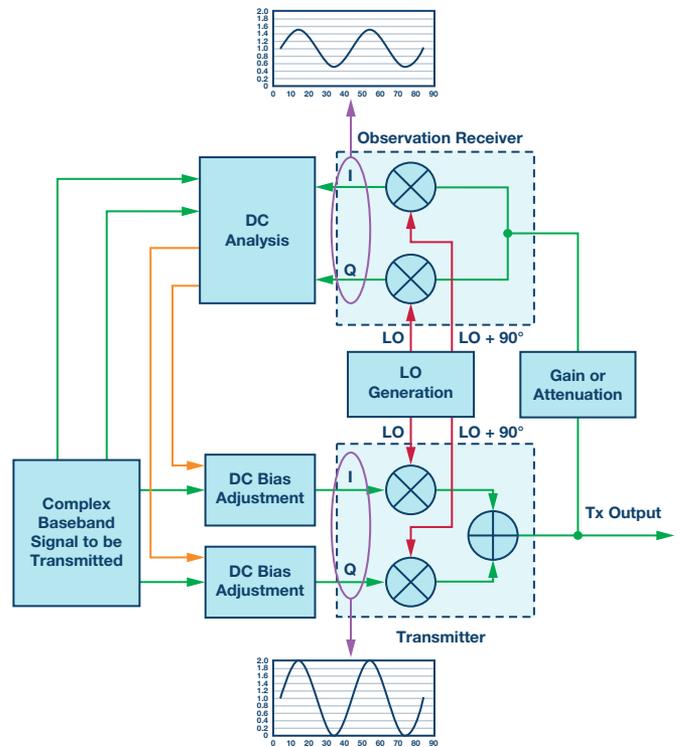


Figure 10. Amplitude scaling due to attenuation in the loopback path.

Now let's consider phase rotation. It is important to realize that signals do not travel instantaneously from point A to point B. For example, signals travel through copper at approximately half the speed of light, meaning that a 3 GHz signal travelling along a copper strip has a wavelength of approximately 5 cm. This means that if the copper strip is probed with multiple oscilloscope probes spaced a few centimeters apart, the oscilloscope will show multiple signals that are out of phase with each other. Figure 11 illustrates this principle, showing three scope probes that are spaced out along a copper strip. The signal seen at each point is at a frequency of 3 GHz, but there is a phase difference between the three signals.

Note that moving a single scope probe down the copper strip would not show this effect, as the scope would always trigger at 0° phase. It is only by using multiple probes that the relationship between distance and phase can be observed.

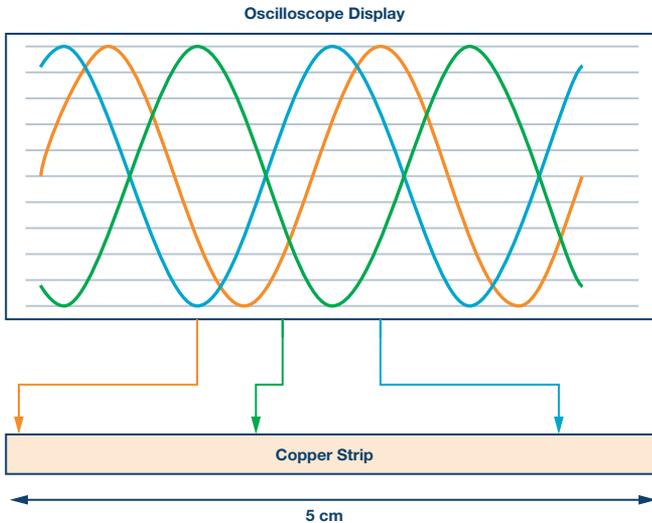


Figure 11. The relationship of distance and phase, a 5 cm trace, a 3 GHz signal, and probe points at 0 cm, 2 cm, and 4 cm.

Just as there is a phase change along the copper strip, there will be a phase change from transmitter input to observation receiver output, as shown in Figure 12. It is essential that the LOL correction algorithm knows how much phase rotation has occurred for it to compute the correct correction values.

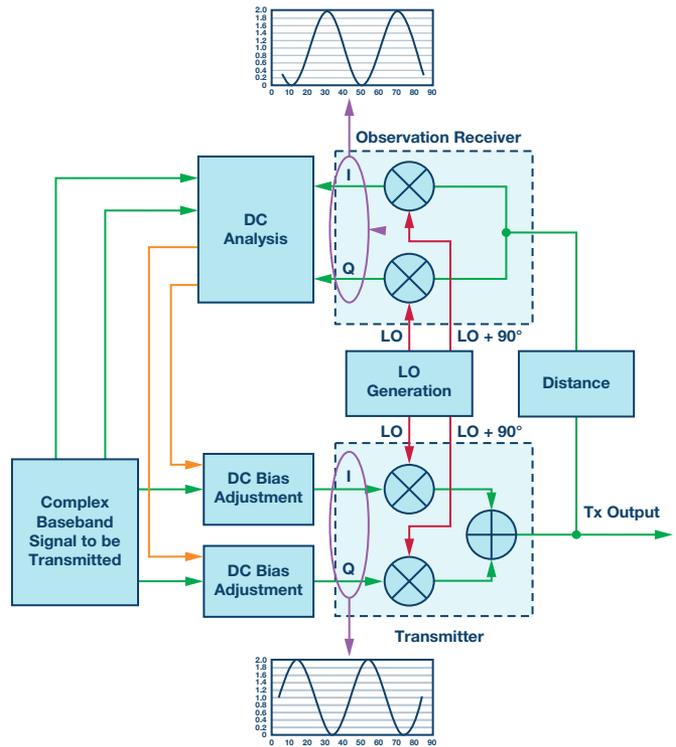


Figure 12. Phase rotation due to physical distance in the loopback path.

Determining the Transfer Function from Transmitter Input to Observation Receiver Output

The transfer function shown in Figure 13 may be learned by applying an input to the transmitter and comparing it to the output from the observation receiver. However, some points need to be kept in mind. If a static (dc) signal is applied to the transmitter input, it will produce an output at the transmit LO frequency and the transmit LOL will combine with it. This will prevent the transfer function from being learned correctly. It should also be noted that the transmit output may be connected to an antenna, so intentionally applying signals to the transmitter input may not be allowed.

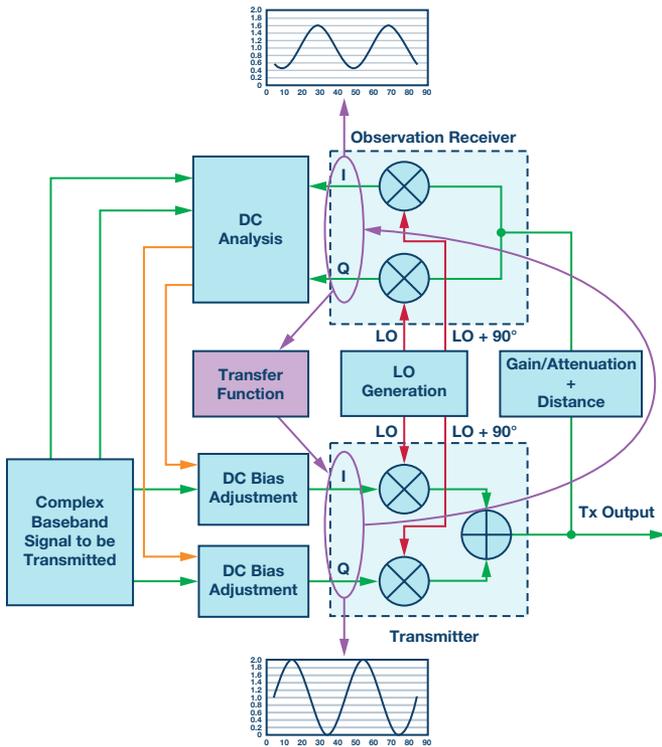


Figure 13. Determining the transfer function from transmitter input to observation receiver output.

To overcome these challenges, the ADI transceivers use an algorithm that applies a low level dc offset to the transmitted signal. The offset level is adjusted periodically and these perturbations will show up in the observation receiver's outputs. The algorithm then analyzes the deltas in input values compared to the deltas in observed values, as outlined in Table 1. In this example there is no user signal being transmitted, but the method still holds in the presence of user signal.

Table 1. The Deltas in Input Value Compared to the Deltas in Observed Value

	Tx Input Signal	Tx Output Port	Observation Receiver Output
Case 1	DC offset 1	$T_{X_{LO}} 1 + Tx\ LOL$	$(T_{X_{LO}} 1 + Tx\ LOL) \times$ transfer function
Case 2	DC offset 2	$T_{X_{LO}} 2 + Tx\ LOL$	$(T_{X_{LO}} 2 + Tx\ LOL) \times$ transfer function

By performing a subtraction of the two cases, the constant transmit LOL is eliminated from the equation and the transfer function can be learned. The number of cases can be expanded to more than two, giving many independent results that can be averaged to increase the accuracy.

Summary

The LOL correction algorithm will learn the transfer function from transmit input to observation receiver output. It will then take the observation receiver's output and divide it by the transfer function to refer it to the transmitter's input. By comparing the dc levels in the intended transmission to the dc levels in the observed transmission, the transmit LOL will be determined. Finally, the algorithm will compute the necessary correction values to eliminate the transmit LOL and apply them as a dc bias to the desired transmission data.

This article provides an overview of one aspect of the algorithms used in ADI's RadioVerse transceivers. For a broader understanding on the concepts of zero-IF and algorithms, see this article on complex RF mixers.¹

References

- David Frizelle and Frank Kearney. "Complex RF Mixers, Zero-IF Architecture, and Advanced Algorithms: The Black Magic in Next-Generation SDR Transceivers." *Analog Dialogue*, Vol. 51, February, 2017.

Dave Frizelle [david.frizelle@analog.com] works as an applications manager in the Transceiver Product Group at Analog Devices Limerick, supporting the integrated transceiver family of products. He has worked at ADI since graduation in 1998. His previous engineering roles include six years working in Japan and Korea supporting the development and design-in of ADI components into advanced consumer products.



Dave Frizelle

Also by this Author:

[Complex RF Mixers, Zero-IF Architecture, and Advanced Algorithms: The Black Magic in Next-Generation SDR Transceivers](#)

Volume 51, Number 1

Ultrawideband Digital Predistortion (DPD): The Rewards (Power and Performance) and Challenges of Implementation in Cable Distribution Systems

By Patrick Pratt and Frank Kearney

Share on   

Introduction

The first cable systems in the U.S. started to appear in the early '50s. Even with the rapid changes in technology and distribution methods, cable has maintained a prominent position as a conduit for the distribution of data. New technologies have layered themselves on the existing cable network. This article focuses on one aspect of that evolution—power amplifier (PA) digital predistortion (DPD). It's a term that many involved in cellular system networks will be familiar with. Transitioning the technology to cable brings substantial benefits in terms of power efficiency and performance. With these benefits come substantial challenges; this article dives deep into some of these challenges and provides an overview as to how they may be solved.

Understanding the Requirements

When power amplifiers are operated in their nonlinear region, their output becomes distorted. The distortion can affect the in-band performance and may also result in unwanted signal spilling over into adjacent channels. The spill-over effect is particularly important in wireless cellular applications, and adjacent channel leakage ratio—or ACLR, as it is termed—is tightly specified and controlled. One of the prominent control techniques is digitally shaping or predistorting the signal before it gets to the power amplifier so that the nonlinearities in the PA are cancelled.

The cable environment is very different. Firstly, it can be regarded as a closed environment; what happens in the cable stays in the cable! The operator owns and controls the complete spectrum. Out-of-band (OOB) distortions are not a major concern. In-band distortions are, however, of critical importance. The service providers have to ensure the highest quality in-band transmission conduit so that they can leverage the maximum data throughputs. One of the ways that they ensure this is by running the cable power amplifier strictly within its linear region. The trade-off for this mode of operation is very poor power efficiencies.

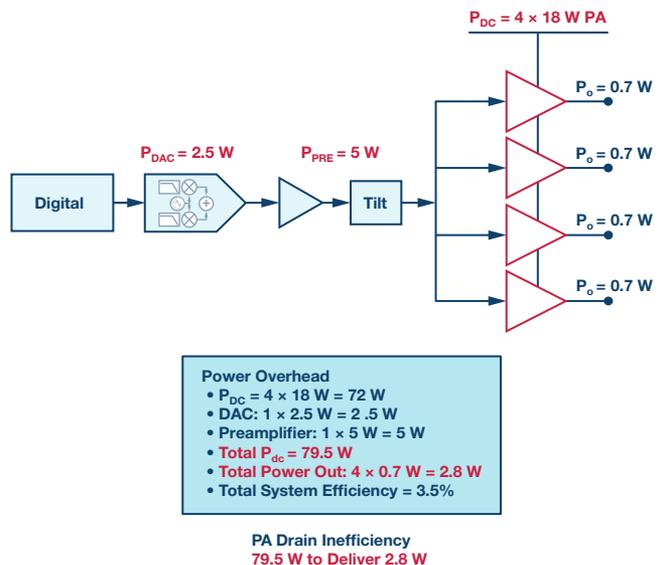


Figure 1. Power efficiencies in cable power amp drivers.

Figure 1 provides an overview of a typical cable application. Although the system consumes nearly 80 W of power, just 2.8 W of signal power is delivered. The power amplifiers are very low efficiency Class A architectures. The maximum instantaneous peak efficiency can be calculated to be 50% (when the signal envelope is at maximum, assuming inductive loading). If the PA is to operate entirely in its linear region, then taking into account the very high peak to average ratio of the cable signals (typically 14 dB) means that the amplifiers need to operate on average 14 dB below the start of compression, hence ensuring that no signal compression occurs even at the peaks of the signal. There is a direct correlation between the back-off and the amplifier operating efficiency. As the amplifier is backed off 14 dB to accommodate the full range of cable signals, the operating efficiency will reduce by $10^{-14/10}$. Hence, the operating efficiency drops from its theoretical max of 50% to $10^{-14/10} \times 50\% = 2\%$. Figure 2 provides an overview.

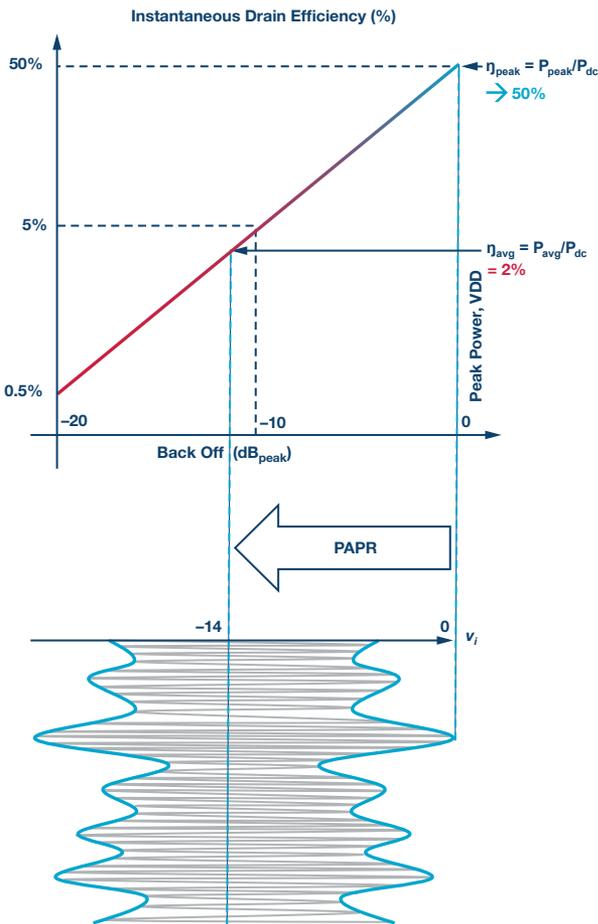


Figure 2. High peak to average ratios push a back-off operational mode and a dramatic decrease in efficiency.

In summary, power efficiency is the major issue. The lost power has a cost implication, but, just as importantly, it also uses up a scarce resource within the cable distribution system. As cable operators add more features and services, they require more processing, and the power for that processing may be constrained within existing power budgets. If wasted power can be retrieved from the PA inefficiencies, then it can be reallocated to those new functions.

The proposed solution to the PA inefficiency is digital predistortion. It's a method universally adopted and employed right across the wireless cellular industry. DPD allows the user to operate the PA in a more efficient, but more nonlinear region, and then pre-emptively correct for the distortions in the digital domain before the data is sent to the PA. DPD is essentially shaping the data before it gets to the PA to counteract the distortions the PA will produce, and hence extend the linear range of the PA, as shown in Figure 3. That extended linear range can be used to support higher quality processing, deliver lower modulation error rates (MER),¹ or allow the PA to run at a reduced bias setting—thus saving power. Although DPD has been widely used in wireless cellular infrastructures, implementing DPD in a cable environment has unique and challenging requirements.

As shown in Figure 4, actual operating efficiencies for the cable application sit at approximately 3.5%! Implementing DPD results in the power requirements of the system dropping from 80 W to 61 W—a power saving of 19 W, which is a 24% reduction. Previously each PA required 17.5 W of power; now that drops to 12.8 W.

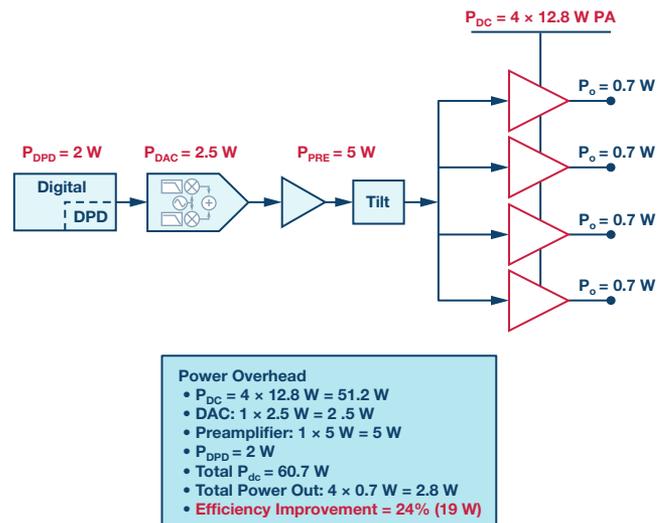


Figure 4. Overview of power saving through DPD implementation.

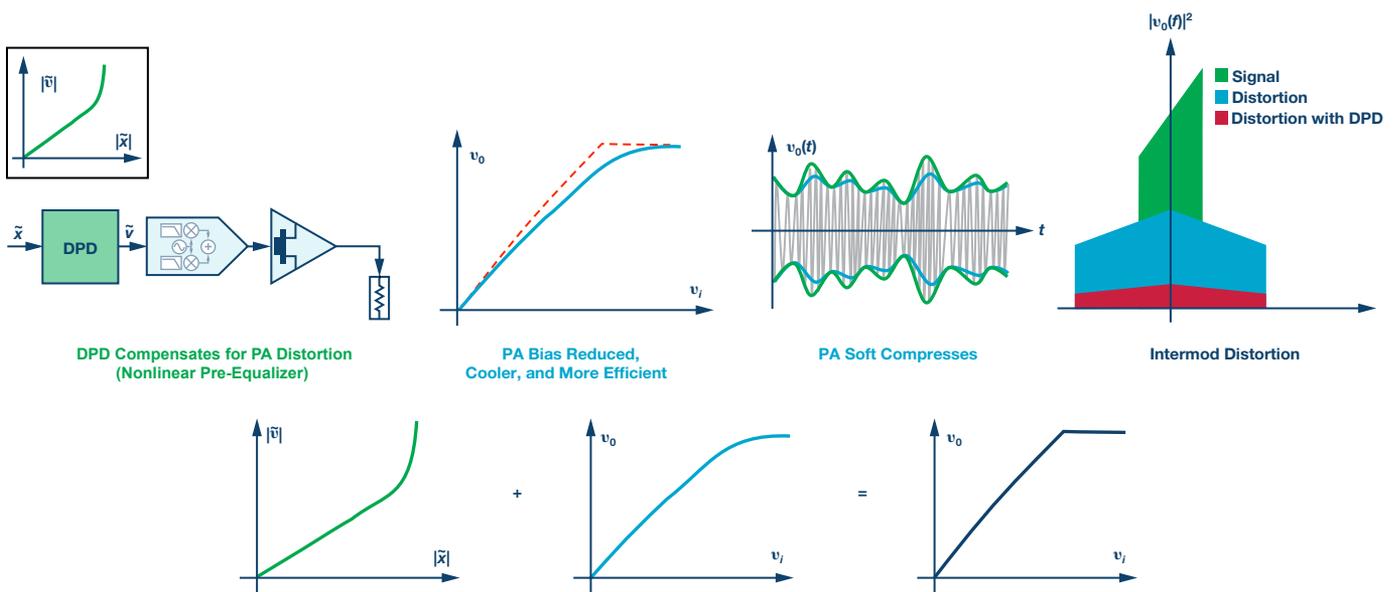


Figure 3. Digital predistortion overview.

The Challenges of Implementation

The value of DPD is clear but the cable application poses many unique challenges to its implementation. Those technical challenges have to be met within the confines of the available resources. For example, the solution itself must be power efficient, as there is little value in optimizing the PA efficiency if the power saved gets shifted into powering the solution. Likewise, the digital processing resources need to be appropriate so that they can reside efficiently within the current FPGA architectures. A very large/complex algorithm with nonstandard hardware requirements and extensive architectural change is unlikely to get adapted.

Ultrawide Bandwidth

Perhaps the most prominent difference between the cable application and that of the wireless cellular environment is the bandwidth of operation. In cable, there is approximately 1.2 GHz of bandwidth to linearize. The wide bandwidth challenge is compounded by the fact that the spectrum starts just 54 MHz from dc and that the signal bandwidth is larger than the channel center frequency. We must remember that the power saving is going to come by driving the PA into its nonlinear region of operation; that delivers better efficiency, but at the cost of generating nonlinear products. The DPD has to cancel the nonlinearities created by the PA, with a particular focus on those that fall back in-band of the wanted signal. That poses a unique challenge in the cable application.

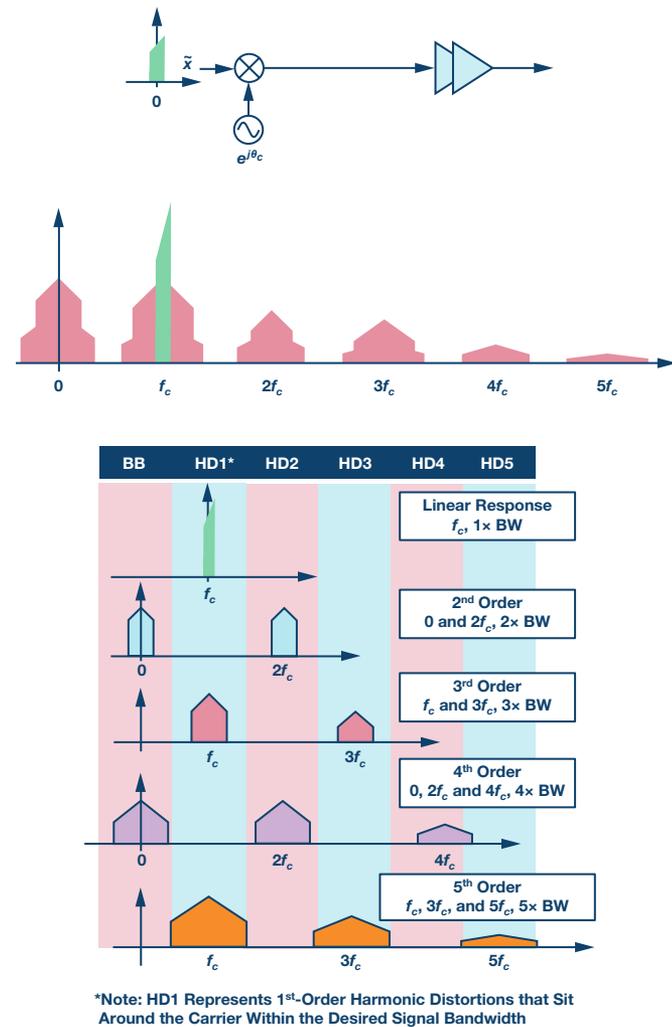


Figure 5. Harmonic distortion terms in conventional narrow-band interpretation.

Figure 5 provides an overview of the wideband harmonic distortion terms we might expect for a conventional, narrow-band (narrow-band is defined later in this section) upconverted baseband signal that goes through a nonlinear amplification stage. The nonlinear PA output is typically described by a power series expression, such as the Volterra series having the form of

$$y(n) = \sum_{m_1=0}^M \dots \sum_{m_k=0}^M h_{k,m_1,\dots,m_k} \prod_{q=1}^k x(n-m_q)$$

which can be understood as a generalization of the Taylor power series to include memory effects. Of fundamental importance to note is that each nonlinear term ($k = 1, 2, \dots, K$) generates multiple harmonic distortion (HD) products. For example, the 5th order has 3 terms: 5th order at 1st harmonic, 5th order at 3rd harmonic, and 5th order at 5th harmonic. Note also that harmonic bandwidth is a multiple of its order; for example, 3rd-order terms are three times wider than the stimulus bandwidth.

In cable, it is not so much the large signal bandwidth, but its positioning on the spectrum (just 54 MHz from dc), that poses a particular challenge to DPD. Harmonic distortion occurs in all nonlinear systems; the focus of cable DPD is the harmonic distortion that fall in band. Looking at Figure 5, it can be seen that in conventional narrow-band applications, the focus will be the 3rd- and 5th-order harmonics. Although others are created, they fall out of the band of interest and can be removed through conventional filtering. We can define wideband and narrow-band applications by their fractional bandwidth where the fractional bandwidth is defined as

$$\frac{(f_n - f_1)}{f_c}$$

(f_n = highest frequency, f_1 = lowest frequency, and f_c = center frequency). When the fractional bandwidth goes above 1, the application can be deemed wideband. Most cellular applications have fractional bandwidth of 0.5 or lower. Therefore, their HD behavior adheres to the characteristics shown in Figure 6.

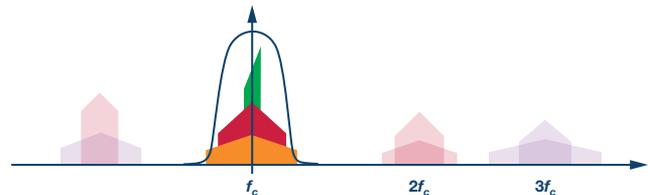


Figure 6. Narrow-band simplification; only products around the 1st harmonic need be considered.

For such narrow-band systems, only in-band distortion around the 1st harmonic needs be cancelled by the DPD since a band-pass filter can be employed to remove all other products. Note also, as no even order products fall in band, the DPD needs only treat odd order terms.

In the cable application we can approximate $f_n \sim 1200$ MHz, $f_1 \sim 50$ MHz, and $f_c \sim 575$ MHz, thus giving us a fractional bandwidth of 2. To determine what minimum HD order needs to be corrected, the equation

$$f_1 K_{\min} < f_n$$

(K_{\min} is the lowest nonlinear order to be considered) can be used, or numerically $50 \text{ MHz} \times 2 = 100 \text{ MHz}$, which is less than 1200 MHz—hence the 2nd-order HD falls well within the band of operation and must be corrected. Therefore, if the decision is made to operate the cable PA outside of its very safe and linear operation, the resultant harmonic distortion will be as indicated in Figure 7.

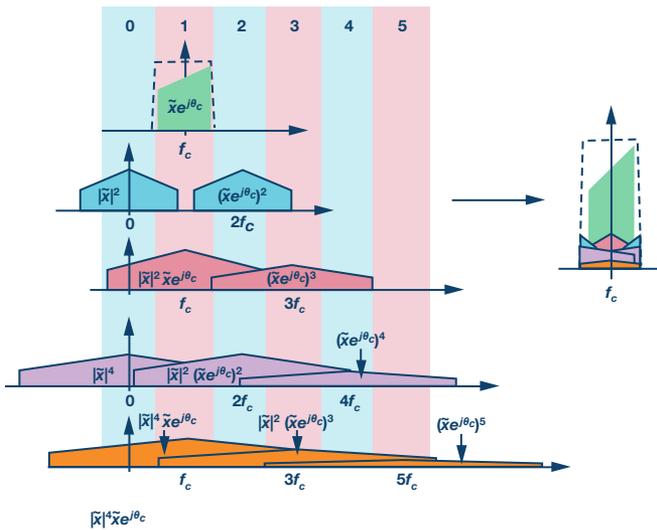


Figure 7. The effects of broadband harmonic distortion in wideband cable applications.

In contrast to wireless cellular where only the odd order harmonics were of concern, within the cable application both even and odd terms fall in band, creating multiple, overlapping zones of distortion. This has some serious implications for the complexity and sophistication of any DPD solution, as the algorithm has to move past simple narrow-band assumptions. The DPD solution must accommodate terms for each order of harmonic distortion.

In a narrow-band system, the even order terms can be ignored and the odd order produce 1 term each within the band of interest. DPD in the cable application has to concern itself with both odd and even order harmonic distortions and it also must consider that each order can have multiple overlapping in-band elements.

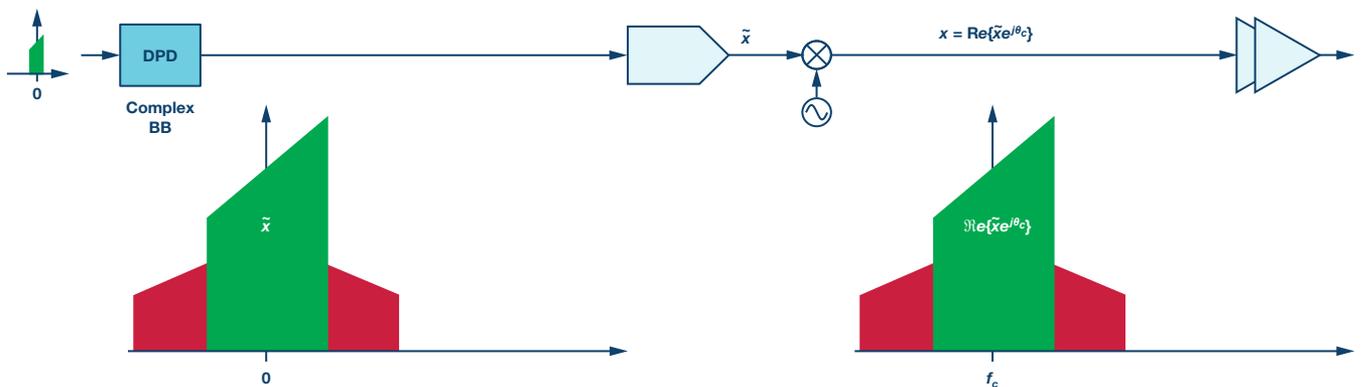
Positioning the Harmonic Distortion Corrections

Considering a conventional narrow-band DPD solution where processing is done at complex baseband, we are principally concerned with harmonic distortions that sit symmetrically around the carrier. In wideband cable systems, although that symmetry is maintained for those terms located around the 1st harmonic, that symmetry no longer holds for the higher harmonic products.

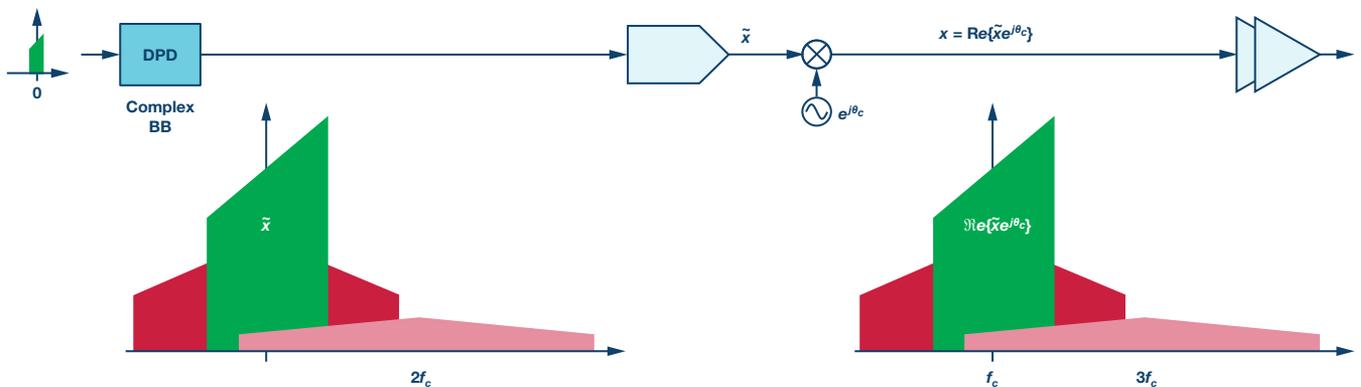
As shown in Figure 8, conventional narrow-band DPD is done at complex baseband. In these instances, only the 1st harmonic products fall in band such that their baseband representation translate directly to RF. When we consider wideband cable DPD, the higher harmonic distortions must be frequency offset so that baseband representation after upconversion is positioned correctly in the real RF spectrum.

Loop Bandwidth Limitations

Closed-loop DPD systems employ a transmission and an observation path. In an idealized model, neither of the paths would be bandwidth restricted and both would be wide enough to pass all DPD terms; that is, both in-band and out-of-band terms are passed.



(a) Conventional Narrow-Band DPD Processing Done at Complex Baseband.



(b) Wideband Cable DPD, OOB HD Must be Frequency Offset to Allow for RF Upconversion.

Figure 8. Note the frequency offset requirements in the complex baseband processing for wideband DPD.

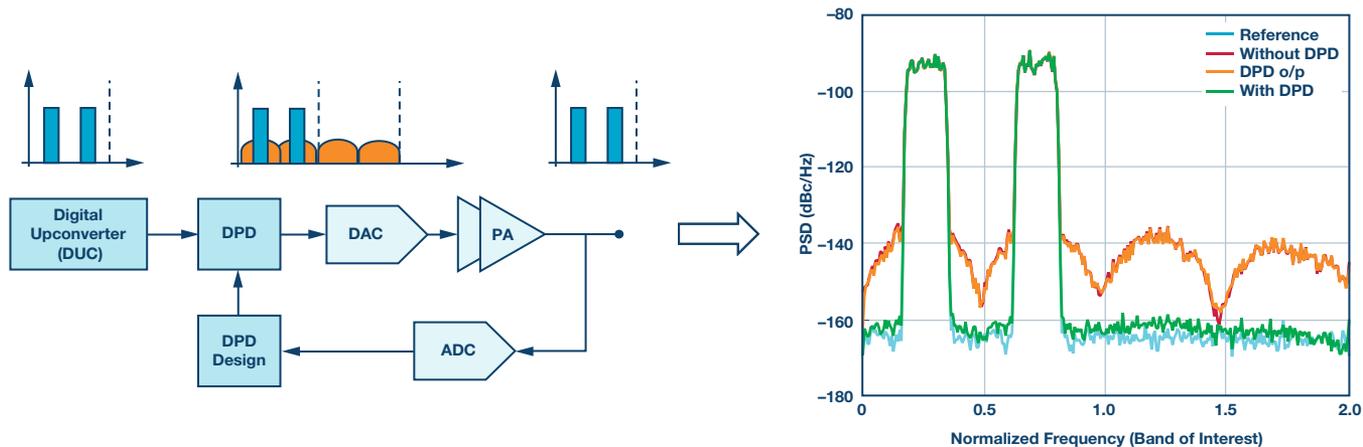


Figure 9. Idealized DPD implementation with no bandwidth restrictions.

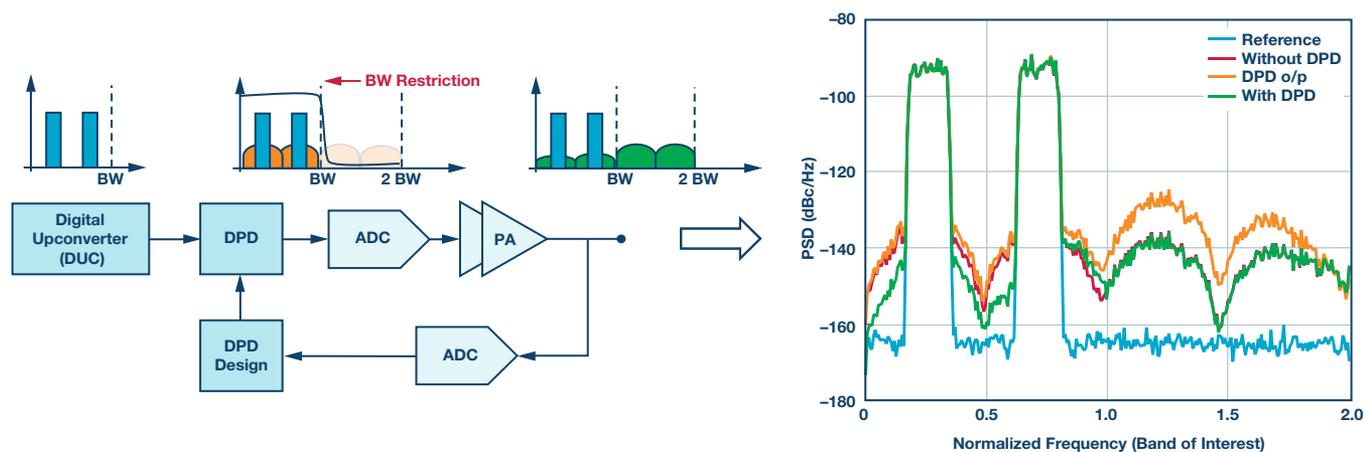


Figure 10. Decreased performance of DPD as bandwidth limitations in the signal path limit the OOB terms.

Figure 9 provides an overview of a DPD implementation. In an ideal situation, the path from the digital upconverter (DUC) via the DPD to the DAC and through the PA will have no bandwidth restrictions. Likewise, the ADC on the observation path will digitize the full bandwidth (note that for illustration purposes we are showing a 2× bandwidth signal path; in some wireless cellular applications that may extend to 3× to 5×). The ideal implementation has the DPD produce terms both in-band and out of band that totally cancel the distortion introduced by the PA. It is important to note that for accurate cancellation, terms are created well outside the bandwidth of the signal of interest.

In a practical implementation, the signal path has bandwidth limitations that modify the DPD performance from the ideal implementation.

In the cable application, the bandwidth limitations can come from a variety of sources: The JESD link between the FPGA and the DAC, the DAC anti-imaging filter, and the PA input matching. The most notable effect of these limitations is the OOB performance. As can be observed on the simulation shown in Figure 10, the DPD fails to correct the OOB distortions. In cable, where OOB distortions contribute to in-band performance degradation, this can be of particular significance; bandwidth restrictions in the signal path can, and do, affect the in-band performance.

The cable environment is unique in that the operator owns the entire spectrum. Emissions that fall out of the band of interest (54 MHz to 1218 MHz) are in a portion of the spectrum not used by anyone else and are also subject to decay due to the inherent cable losses at high frequency. The observation path need only be concerned with monitoring what is happening within the band of operation.

An important distinction needs to be made here; emissions that fall out of band are not of concern, but those that are generated out of band and extend back down into band are. Hence, although OOB emissions are not of concern, the terms that create them are. The implementation is very different to the wireless cellular application where the observation bandwidth requirement is typically 3× to 5× that of the band of operation. In cable, the focus is in-band performance, and OOB terms need only be considered regarding their effect on the in-band performance.

Cable DPD only needs to correct for in-band products: 54 MHz to 1218 MHz for DOCSIS 3. The DPD generates 2nd, 3rd, ... cancellation terms. Although we only need to correct over the cable bandwidth, within the DPD actuator these terms extend over a wider bandwidth (for example, the 3rd order extends out to 3× 1218 MHz). To preserve stability of a conventional DPD adaptation algorithm, these OOB terms should be preserved around the loop. Any filtering of the DPD terms tends to destabilize the adaptation algorithm. In a cable system there is band limiting and thus a conventional algorithm can fail.

DPD and Cable Tilt Compensation

As with all other mediums of transmission, cable introduces attenuation. In general, this attenuation can be considered as functions of the cable quality, the cable run distance, and the frequency of transmission. If a relatively uniform received signal strength is to be achieved at the receiving end of the cable, right across the spectrum of operation, then pre-emphasis (tilt) must be added at the transmit side. The tilt can be regarded as an inverse transfer function of the cable. It applies pre-emphasis, or shaping, that is proportional to the frequency of transmission.

The shaping is achieved by a low power passive analog equalizer known as a tilt compensator, located just prior to the power amplifier. Little or no attenuation is applied at the high frequencies, whereas the maximum attenuation is applied to the lower frequencies. Signals at the output of the tilt compensator can have level variations of up to 22 dB across the spectrum of operation.

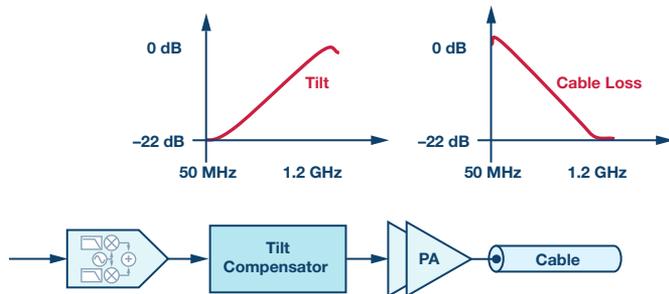


Figure 11. Tilt compensator implementation.

The tilt compensator shapes the signal, and that shaping profile is maintained as the signal is processed through the PA. Conventional DPD implementation would see the shaping as an impairment and try to correct for it as DPD is an (nonlinear) equalizer. It might seem reasonable to suggest that if the inverse to the tilt was added to the observation path that it would mitigate the effect. However, this is not true. Because the PA is nonlinear, commutativity doesn't hold—or, in other words,

$$PA\{T\} \neq T\{PA\}$$

(PA is the model of the power amplifier and T is the model of the tilt compensator).

For optimal operation, the DPD processing block needs an explicit knowledge of the signal that will present itself at the input to the PA. In a cable DPD application, the tilt compensation must be maintained while simultaneously having the DPD algorithm model the PA. This presents some very unique and difficult challenges. We need a low cost, stable solution that doesn't equalize out the tilt. While the nature of the solution can't be disclosed in this article, ADI has found an innovative solution to this problem, which may be described in detail in future publications.

DPD and Cable PA Architecture

As shown in Figure 4, the typical cable application will have the output from one DAC split and supplied to four separate PAs. To get maximum power savings, the DPD needs to be implemented on all of those PAs. One possible solution could be the implementation of four independent DPD and DAC blocks. The solution works but the efficiencies are reduced and the system implementation costs increased. The additional hardware has a dollar and a power cost.

Not all PAs are created equal and, although process matching (during manufacturing) may deliver units that have similar personalities, differences will persist and may grow larger with aging, temperature, and supply variations. Having said that, using one PA as the master and developing an optimized DPD for it, which is then applied to other PAs, does deliver a system performance benefit as shown from the simulated results in Figure 12.

The plots to the left-hand side indicate the PA performance with no DPD applied. The nonlinear operational mode causes distortions and this is reflected in the MER¹ performance, having a range of 37 dBc to 42 dBc. Closed-loop DPD is applied by observing the output from the master PA; the green plot to the right-hand side of the diagram shows the enhanced performance. The DPD has corrected for the PA distortions and the result is that the overall performance has been shifted to deliver a MER ranging from 65 dBc to 67 dBc. The remaining plots in the middle show the performance of the slave PAs—that is, the PAs that have been corrected based on the master PA. As can be seen, implementing closed-loop DPD by just observing one PA has benefited the performance of all the PAs. However, the performance of the slave PAs continues to have points of operation that will fail. The span of performance on the slave PAs ranges from 38 dBc to 67 dBc. The wide range itself is not the concern, but a portion of that range falls below the acceptable operational threshold (typically 45 dBc for cable).

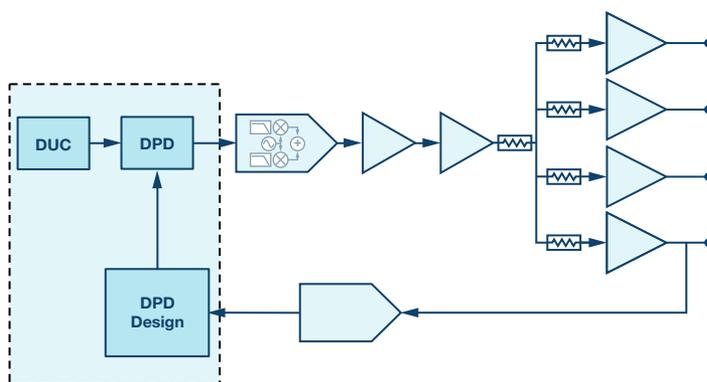
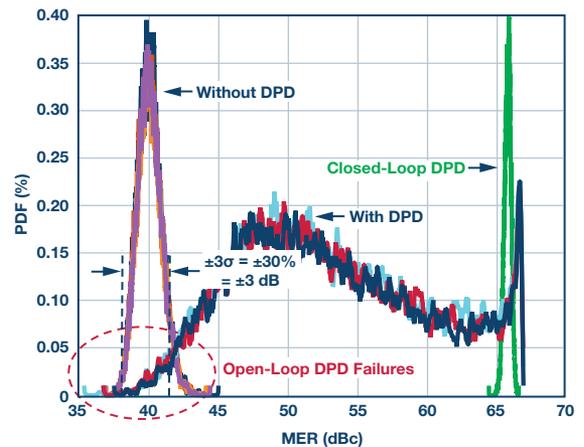


Figure 12. Single DPD with multiple PAs (simulated results).



The unique system architecture in cable provides an extra challenge for DPD. Optimized performance requires closed-loop DPD implementation. However, conventional thinking would suggest that to do so in cable would require additional hardware in each of the PA paths. An optimal solution needs to deliver the enhancements of closed-loop DPD to each PA but without the additional hardware costs.

Solving the Challenges with SMART Algorithms

As described earlier in this article, cable DPD presents the designer with particularly unique and difficult challenges. The challenges have to be resolved, but within the constraints of power and hardware so that the advantages are not eroded; there is little value in saving PA power if that power is used in additional DACs or the FPGA. Likewise, the power saving must be balanced with a hardware cost. ADI solves the challenge using a combination of high performance analog signal processing, coupled with advanced algorithm implementation.

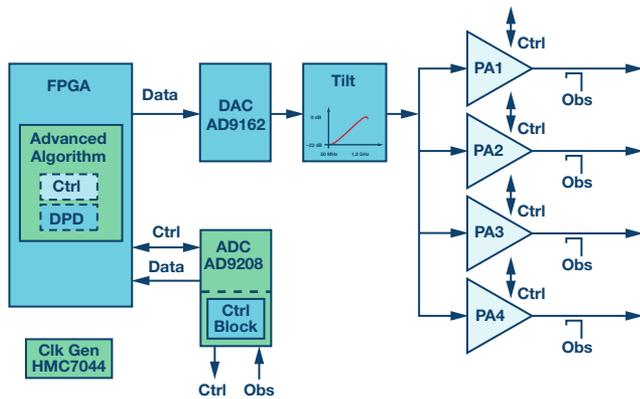


Figure 13. Cable DPD implementation using advanced converters and SMART algorithms.

A high level overview of the ADI implementation is shown in Figure 13. The solution can be regarded as having three key elements: The use of advanced converter and clocking products, an architecture that supports a comprehensive signal chain monitoring/control, and, lastly, an advanced DPD algorithm that can utilize the former knowledge to deliver optimal performance.

The algorithm sits at the heart of the solution. It uses its extensive knowledge of the signal being processed and the transfer function of the signal path to shape the output while simultaneously adjusting the dynamic control of some aspects of the signal path. The dynamic system solution means not only does the system designer have the capability to obtain

substantial power saving, but that those power savings can be directly traded off against performance. The algorithm is such that once the user defines the MER¹ performance level to which the system must operate, the system tuning is implemented so that performance is achieved across all the outputs. It's important to note that the algorithm also ensures that the performance threshold is achieved while maintaining optimal power usage at each PA; no PA takes more power than required to achieve the target performance.

The previous paragraph provides an overview of the solution implementation. The specifics of the algorithm itself are ADI proprietary IP and beyond the scope of this article. The SMART algorithm has the ability of learning the system path and then changing both the nature of the data being transmitted through the path and the characteristics of the path itself to deliver optimal results. We define optimal results as maintaining the quality of the MER while simultaneously lowering the power requirement.

The path characteristics, along with the nature of the transmitted signal, are in constant flux. The algorithm has the self-learning capability to deal with that dynamic adaptability. What's more, the adaption happens while the system is live without interrupting or distorting the transmitted streams.

Conclusion

The cable environment continues to be an important infrastructure for the delivery of data services. As the technology evolves, so too do the pressures on spectrum and power efficiencies. The next generation of developments is calling for ever increasing demands and pushing higher order modulation schemes and better power efficiencies. These enhancements must come without affecting system performance (MER), and while DPD provides one possible avenue of enablement, its implementation in the cable application poses unique and difficult challenges. ADI has developed a total system solution to address these challenges. That solution encompasses silicon (DACs, ADCs, and clocks), PA control, and advanced algorithms. The combination of all three technologies provides users an adaptable solution where they can easily trade off power and performance requirements with minimal compromise. This software-defined solution also supports a painless transition to the next generation of cable technologies that are expected to incorporate full duplex (FD) and envelope tracking (ET).

Note 1: Modulation error rate is a measure of the quality of modulation. It represents the difference between the target symbol vector and the transmitted symbol vector. $MER = 10\log(\text{average signal power}/\text{average error power})$. It can be regarded as a measure of accuracy of the symbol placement within the constellation.

Patrick Pratt [patrick.pratt@analog.com] is a senior research scientist heading up Analog Devices cable DPD algorithm development. His career spans over 30 years and encompasses algorithm research and development activity, both in private organizations and academic institutions. Patrick holds a Ph.D. in electronic engineering from Cork Institute of Technology.



Patrick Pratt

Frank Kearney [frank.kearney@analog.com] works as algorithm development manager in the Communication Systems Engineering Team at Analog Devices Limerick. He has worked at ADI since graduation in 1988. He recently returned from China where he held the position of senior applications manager for the Systems Engineering Team in Asia Pacific. He is currently a doctoral candidate at University College Dublin.



Frank Kearney

Also by this Author:

[Passive Intermodulation \(PIM\) Effects in Base Stations: Understanding the Challenges and Solutions](#)

Volume 51, Number 1

Rarely Asked Questions—Issue 143

Burned by Low Power? When Lower Current Consumption Can Get You Into Trouble

By **Abhinay Patil**

Share on   



Question:

I changed one of my devices to a newer and better part with lower current consumption. Nothing is working anymore and even the new part got “burned.” Do you have any explanation?

Answer:

Linear regulators are fairly simple devices and don’t have too many challenges. With that said, you can still get into trouble with them occasionally.

In my role as a field applications engineer, there are times when customers request me to recommend a drop in replacement for a part they are using from another supplier. In many of these cases, the decision to replace a part is made by the customer’s production/purchase team and the original circuit designer may not be aware of the change. The decision process is fairly simple and the alternative part should have the same functionality, package, and pin configuration, as well as equal or better electrical specifications when compared to the device being replaced. Once all of these boxes are checked, you provide the necessary comparison data to the component engineer and the new component is added to the bill of materials as a second source alternative. After doing that, you should be all set. That is, until the product that was working faithfully with the older component starts failing on the production line after changing over to the alternative. What went wrong?!

I was involved in one such case where we followed the previously mentioned process and got one of our isolated RS-485 transceivers designed in as a second source to another supplier’s device in a customer’s design. The devices were form, fit, and functionality compatible, and our part had better electrical specifications. The customer went ahead and ordered a large quantity of this part from us and there seemed no reason for anything

to go wrong. However, the customer reported that the new RS-485 transceiver started failing on the production test bed. Since nothing else in the design had been changed, it came down to something being wrong with the new device in the picture.

Upon further investigation, we found that the linear regulator powering the bus side of the transceiver was not regulating down to 5 V as expected, but instead was going up to a much higher voltage. We needed to do a careful review and comparison of the data sheets for the older transceiver and the replacement one, as well as the data sheet of the linear regulator to figure out what was going wrong.

Better is a qualitative term and depends on the parameter in question. For example, when it comes to speed/CMRR/PSRR, higher is better; when it comes to offset voltage/drift, lower is better; and you don’t need great engineering insight to know that when it comes to power consumption, lower is always better. Or is it? In this particular case, it turned out it wasn’t. The older transceiver consumed 15 mA (typically) on the bus side in idle state, while the new device consumed only 2 mA (maximum). It was no wonder that the new device looked better on paper. Unfortunately, the linear regulator didn’t seem to agree and went berserk.

As I mentioned at the start of this article, linear regulators are fairly simple devices that don’t have too many demands. Their one peculiar requirement is that they need a minimum load current to function properly. If this condition is not met, the regulator stops regulating correctly and the output voltage goes out of bounds. The situation could also get worse if the input voltage to the regulator is much higher than the desired output voltage. Read this [excellent article](#) for more details on this topic.¹

Many modern linear regulators take special care in their design to address this and pose no such problems. Some of the older devices (like the one in the customer’s design) don’t, and, hence, need extra precautions during the system design. In some cases, the feedback resistor network for the adjustable output LDO takes care of the minimum load current. Unfortunately, one could run into this issue inadvertently by deciding to supersize these resistors while keeping the ratio the same. There could also be another scenario in which the device being powered by the LDO meets the load requirement during normal operation, but not while it’s in standby. These are all potential pitfalls to watch out for, so be sure to read the LDO data sheet carefully. If there’s a minimum load current requirement, it would usually say so in some way. Here are a couple of examples:

Load Regulation	10 mA ≤ I _{OUT} ≤ I _{MAX} ⁽¹⁾⁽³⁾	T _J = 25°C	0.1%	0.3%
		Over Full Operating Temperature Range	0.3%	1%
I _{MIN}	Minimum Load Current ⁽³⁾	–40°C to +125°C	1	mA

Figure 1. An example of a minimum load current from a data sheet.

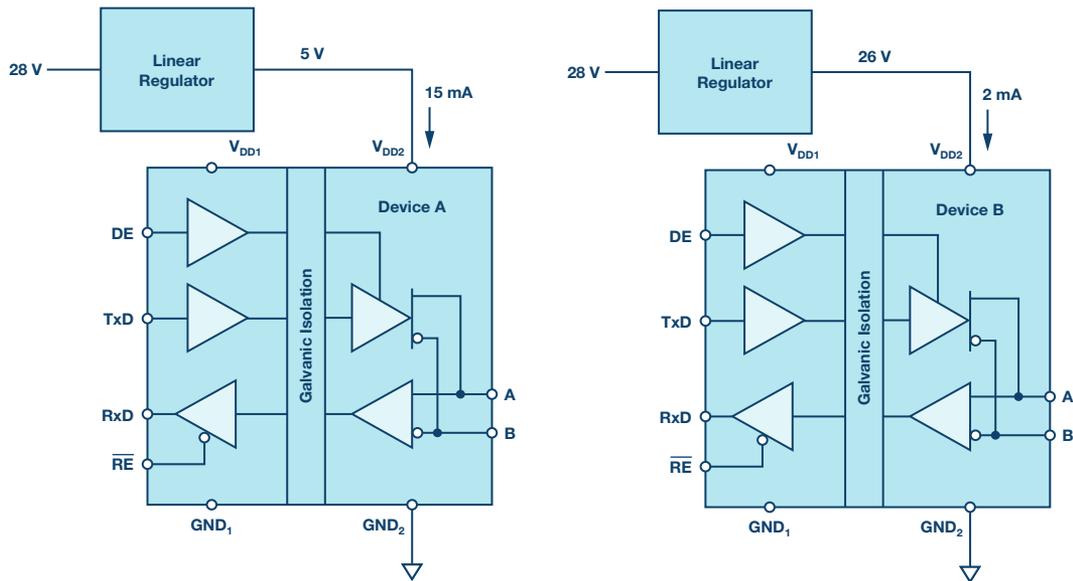


Figure 2. Left schematic: the regulator operating correctly with the older device (minimum load current requirement satisfied). Right schematic: the unstable regulator with the new device (load current not sufficient).

Going back to our story—once the root cause of the problem was understood, it was fairly simple to fix. All we had to do was add a bleeder resistor on the regulator output to draw the minimum load current. While it would have been easy for an unreasonable customer to blame it on our part, this particular customer could see the funny side of this story and was able to appreciate the technical learning from this episode.

It all ended well like any fairy tale—a little blood was spilled, but everyone lived happily ever after.

References

1. Luca Vassalli. “Minimum Load Current Operation—Zero-Load Operation.” *Analog Dialogue*, Vol 48, September, 2014.

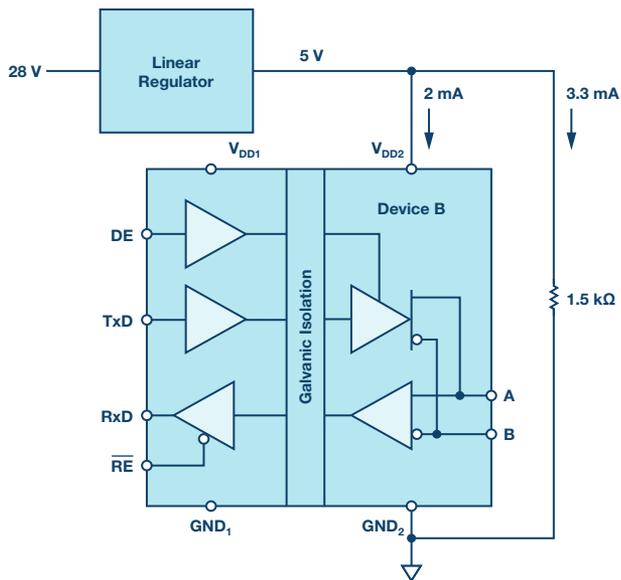


Figure 3. The problem resolved after adding a bleeder resistor to draw minimum required load current.

Abhinay Patil [Abhinay.patil@analog.com] joined Analog Devices India in 2003 and is currently working as a field applications manager in Key Accounts, based in Bangalore. He holds a bachelor of engineering degree in electronics and communications.



Abhinay Patil

Improved DAC Phase Noise Measurements Enable Ultralow Phase Noise DDS Applications

By Peter Delos and Jarrett Liner

Share on   

Introduction

In radar applications, phase noise is a critical performance metric for systems requiring high clutter attenuation. Phase noise is a concern for all radio systems, but radar in particular can require phase noise performance at frequency offsets much closer to the carrier than a communication system.

System designers in these high performance systems will choose ultralow phase noise oscillators and the objective of the signal chains, from a noise perspective, is to add minimal degradation to the oscillator phase noise profile. This requires residual or additive phase noise measurements of the varied components in the signal chain.

Recent product releases of high speed digital-to-analog converters (DACs) are extremely attractive for both waveform generation and frequency creation for any LOs needed in frequency conversion stages. The radar objectives, however, challenge the DAC phase noise performance.

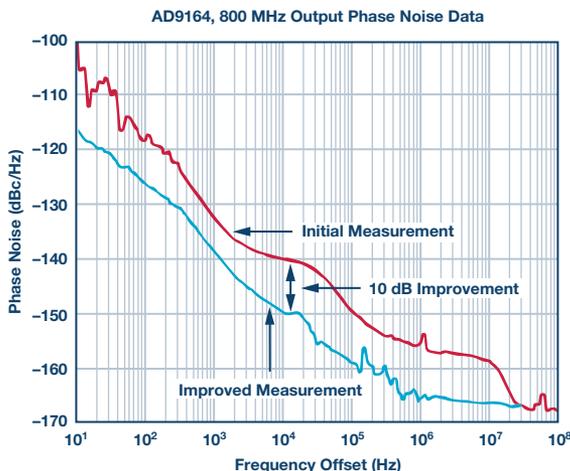


Figure 1. AD9164 phase noise improvements.

In this article we show measured improvements of over 10 dB at 10 kHz offsets using the AD9164 DAC. Figure 1 illustrates the improvement and we will discuss how the results were achieved through a combination of both power supply regulator selection and test setup improvements.

Phase Noise Definition

Phase noise is a measure of the deviation in the zero crossing of a periodic signal. Consider a cosine wave with phase fluctuations

$$x(t) = \cos(2\pi ft + \Phi(t))$$

f = instantaneous frequency

$\Phi(t)$ = randomly fluctuating phase in radians

Phase noise is determined from the power spectral density of the phase variations

$$S(f) = \frac{\Phi(t)_{RMS}^2}{BW} \text{ with units of } \frac{\text{rad}^2}{\text{Hz}}$$

In linear terms, the single sided phase noise is defined as

$$L(f) = \frac{S(f)}{2}$$

Phase noise is normally expressed in units of dBc/Hz from $10\log(L(f))$. Phase noise data is then plotted at offset frequencies relative to the RF carrier.

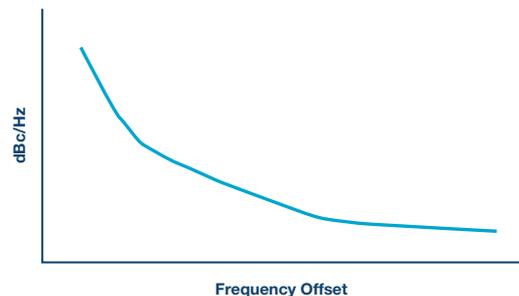


Figure 2. Phase noise plot method.

An important further definition of phase noise is absolute phase noise vs. residual phase noise. Absolute phase noise is the total phase noise measured in the system. Residual phase noise is the additive phase noise of the device under test. This distinction becomes critical in the test setups and in the process of determining component level phase noise contributions in a system.

DAC/DDS Phase Noise Measurement Methods

The figures in this section illustrate DDS phase noise test setups. For DAC phase noise measurements, it is assumed the DAC is used as part of a direct digital synthesizer (DDS) subsystem. A DDS is implemented with a digital sinewave pattern to a DAC that could be in a monolithic IC or an FPGA or ASIC communicating to a DAC. In modern DDS design, the digital phase errors can be made much less than the DAC errors, and DDS phase noise measurements are typically limited by the DAC performance.

The simplest and most common test setup is shown in Figure 3. A clock source is used for the DDS and the DDS output is fed to a cross correlation type phase noise analyzer. This is easy to implement since only a single DDS is required. However, with this test setup there is no method to extract the oscillator contribution to show only the DDS phase noise.

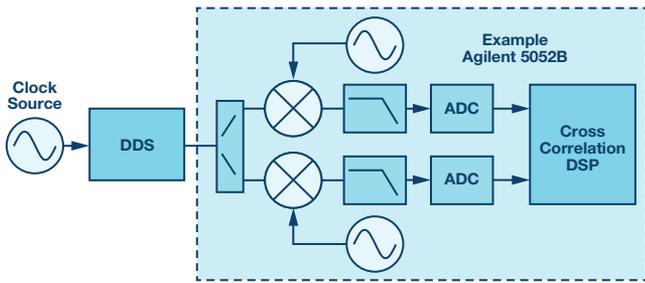


Figure 3. Absolute phase noise DDS test setup includes both DAC and oscillator noise.

Figure 4 shows two common methods to remove the oscillator phase noise from the measurement, providing a residual noise measurement. The drawback of the measurements is that additional DACs are required in the test setup. However, the benefit is a much better indicator of the DAC phase noise contribution that can be applied in system-level analysis budgets.

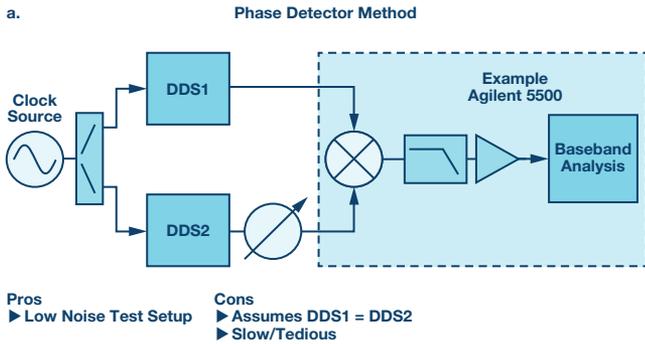


Figure 4a. DDS residual phase noise measurement using the phase detector method.

Figure 4a shows the phase detector method. In this case, two DACs are used and the oscillator contribution subtracts from both DUTs in the downconversion to dc.

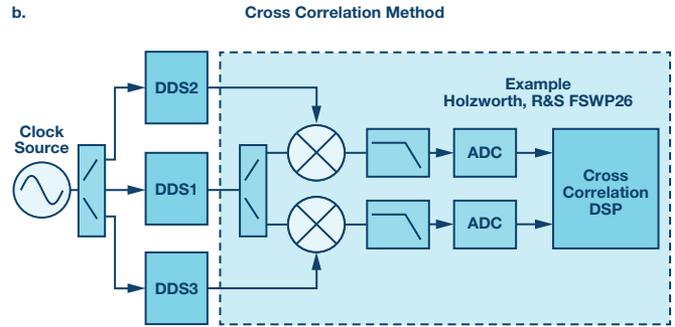


Figure 4b. DDS residual phase noise measurement using the cross correlation method.

Figure 4b shows a method using cross correlation phase noise analysis. In this case, DDS2 and DDS3 are used to translate the clock contribution to the LO ports of the measurement, their contribution is removed in the cross correlation algorithms, and the DDS1 residual phase noise is obtained in the measurement.

Power Supply Noise Contributions

In low noise analog and RF design, power supply noise is a well-known factor to consider. Power supply ripple that is periodic modulates onto the RF carrier and creates spurs on the RF carrier at frequency offsets equal to the ripple frequency. Regulator 1/f noise modulates onto the RF carrier also and contributes to the phase noise profile. Figure 5 illustrates the principles.

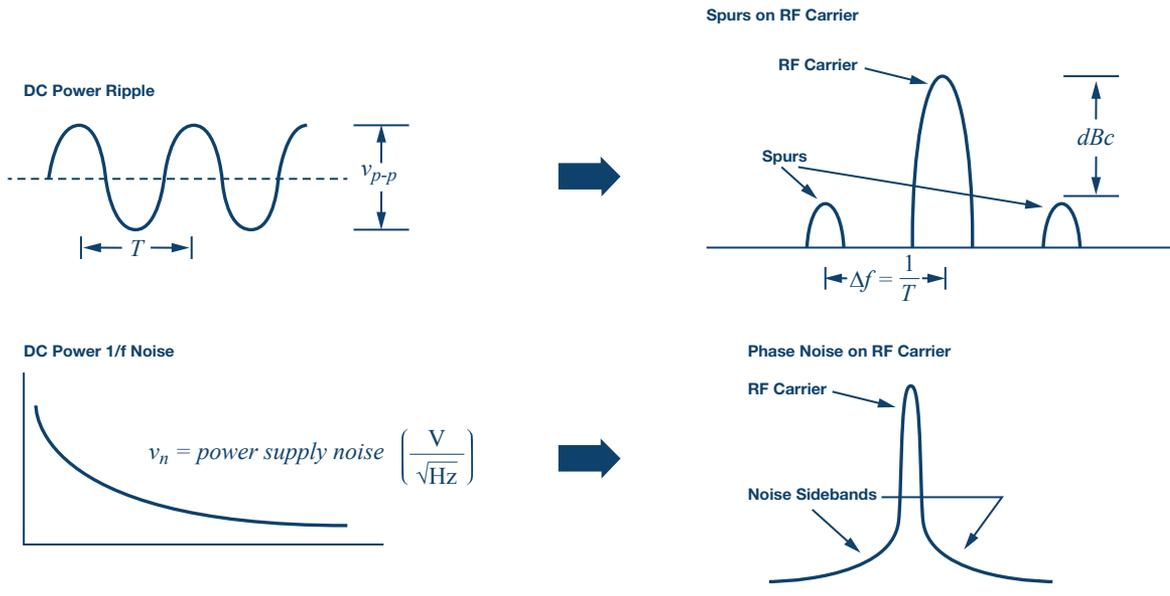


Figure 5. Power supply imperfections modulated on to the RF carrier.

Measured Results

During investigation of the true DAC phase noise performance, both the test setups and the regulator noise performance were considered.

The initial DAC evaluation board included the ADP1740 regulator for the analog and clock voltages. Noise spectral densities were compared with recently released ultralow noise regulators and the ADM7155 was chosen. Figure 6 illustrates the comparison of the noise densities as shown in the

product data sheets. The power supply modification was merely to use the ADM7155 for both the AD9164 clock (data sheet pins VDD12_CLK) and the analog voltage (data sheet pins VDD12A).

Next, test setup options were considered for residual phase noise measurements. The cross correlation method was chosen with the Rohde & Schwarz FSWP primarily out of availability and convenience. The test setup used is shown in Figure 7.

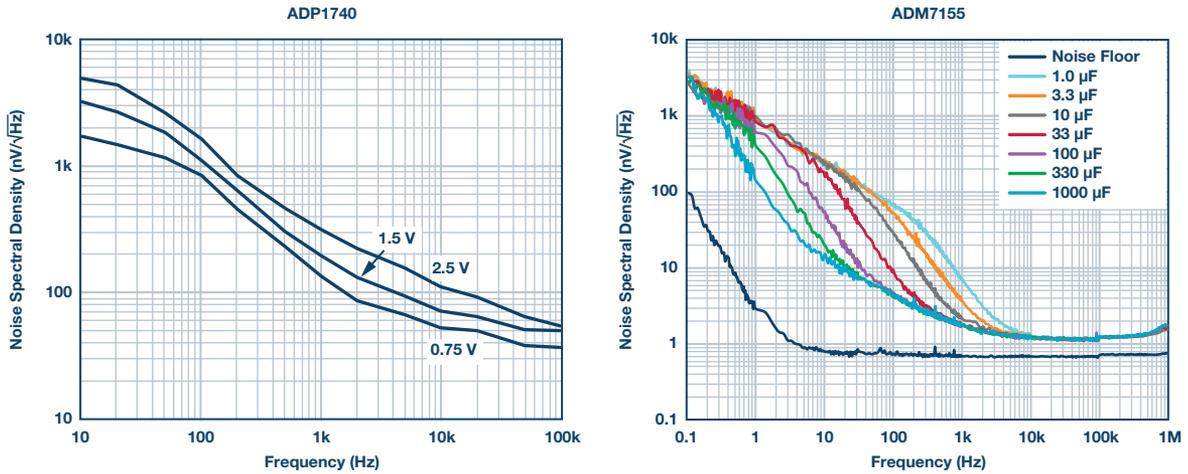


Figure 6. Regulator noise density comparison. Note the Y-axis units—the ADM7155 is an order of magnitude improved.

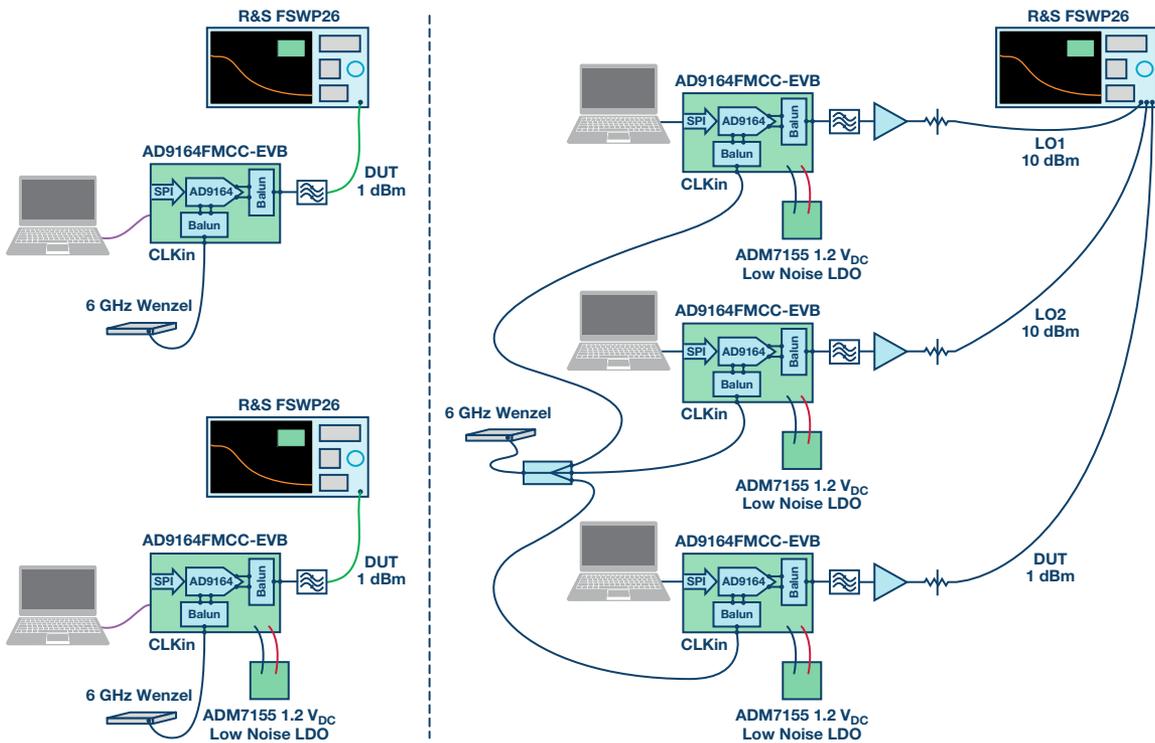
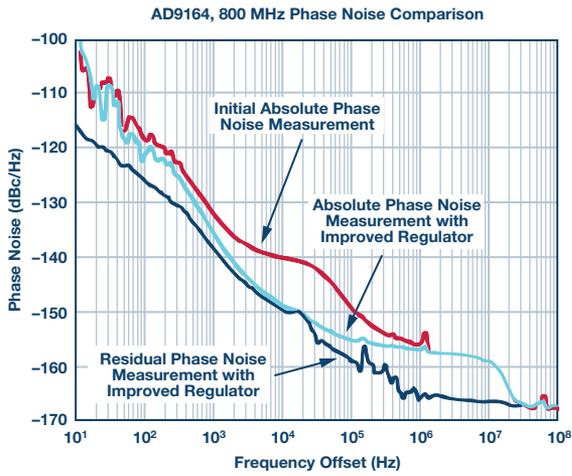


Figure 7. Test setups for AD9164 phase noise measurements.



- Improvements**
- ▶ <1 kHz
 - Test Setup Limited
 - ▶ 1 kHz to 100 kHz
 - Regular Limited
 - ▶ 100 kHz to 1 MHz
 - Test Setup Limited

Figure 8. AD9164 800 MHz output phase noise comparisons.

Figure 8 is a measurement of three cases. The initial evaluation board measurement taken with an absolute phase noise approach is shown as the red curve. The light blue curve is also an absolute measurement but with the regulator improvement. The dark blue curve is a residual phase noise measurement that also includes the regulator improvement.

The measurement indicates three general regions of limitations in the initial measurement that were not obvious in the beginning of the

investigation. Frequencies below 1 kHz were limited by the close in noise of the clock source. Frequencies from 1 kHz to 100 kHz were limited by the regulator selection. Frequencies above 100 kHz were limited by the clock source. The sharp drop off above 10 MHz is the clock source contribution as the clock used was a multiplied crystal oscillator to create 6 GHz and the roll-off is from the RF filters used in the multiplication stages.

Residual phase noise measurements with the regulator improvement were taken at additional DAC frequencies and several are summarized in Figure 9. The modifications were duplicated on several evaluation boards and all cases showed the same improved results.

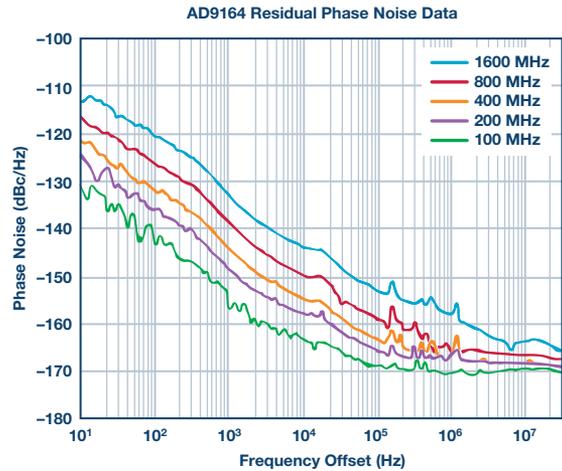


Figure 9. AD9164 residual phase noise measurements with low noise regulator improvement.

Table 1. Family of Regulators with State-of-the-Art Noise Density Performance

Part Number	V _{IN} Min (V)	V _{IN} Max (V)	V _{OUT} Options or Adj Range (V)	I _{OUT} (mA)	PSRR @ 1 kHz	PSRR @ 1 MHz (μV rms) ¹	RMS Noise 100 Hz to 100 kHz (nV/√Hz)	Noise Spectral Density 100 kHz (nV/√Hz)	Dropout @ Rated I _{OUT} Typ (mV)	Total Accuracy Max (±%)	Package
ADM7150	4.5	16	Fixed: 1.5 to 5.0	800	94	62	1	2	600	2	3 mm × 3 mm, 8-lead LFCSP, 8-lead SOIC
ADM7151	4.5	16	Adjustable: 1.5 to 5.1	800	94	62	1	2	600	2	3 mm × 3 mm, 8-lead LFCSP, 8-lead SOIC
ADM7154 New	2.3	5.5	Fixed: 1.2 to 3.3	600	90	58	1	1.2	120	2	3 mm × 3 mm, 8-lead LFCSP, 8-lead SOIC
ADM7155 New	2.3	5.5	Adjustable 1.2 to 3.3	600	90	58	1	1.2	120	2	3 mm × 3 mm, 8-lead LFCSP, 8-lead SOIC

¹ Noise independent at fixed output voltage.

The family of ultralow noise regulators with similar noise density is shown in Table 1. As demonstrated, the impact on the DAC phase noise is significant and these are also recommended for consideration for any areas of the RF system requiring optimum phase noise performance.

Summary

A phase noise review was provided for the fundamental definition, absolute vs. residual phase noise, DAC phase noise measurement test setups, and regulator noise contributions.

DAC phase noise improvements were demonstrated for including both residual phase noise test methods and optimum regulator selection. The end result is the AD9164 now is an enabler for ultralow phase noise, DDS-based applications when the analog voltages and clock voltages are powered from the Analog Devices family of low noise regulators.

References

Bergeron, Jarrah. "Analyzing and Managing the Impact of Supply Noise and Clock Jitter on High Speed DAC Phase Noise." *Analog Dialogue*, Vol. 51, 2017.

Calosso, Claudio E., Yannick Gruson, and Enrico Rubiola. "Phase Noise and Amplitude Noise in DDS." IEEE Frequency Control Symposium, 2012.

Jayamohan, Umesh. "Powering GPS or RF Sampling ADCs; Switcher vs LDO." *Analog Dialogue*, Vol. 50, 2016.

"Product Note 11729B-1, Phase Noise Characterization of Microwave Oscillators: Phase Detector Method." Agilent, May, 2007.

Reeder, Rob. "Designing Power Supplies for High Speed ADC." Analog Devices, Inc., 2012

Walls, Warren F. "Cross-Correlation Phase Noise Measurements." IEEE Frequency Control Symposium, 1992.

Peter Delos [peter.delos@analog.com] is a technical lead at Analog Devices, Inc., in the Aerospace and Defense Group. He received his B.S.E.E. from Virginia Tech in 1990 and M.S.E.E. from NJIT in 2004. He worked in the Naval Nuclear Power program from 1990 to 1997. This work included completion of the Naval Nuclear Power School Officer's Program, and work as an instructor in a Naval submarine facility and as a lead electrical field engineer on the Seawolf class Submarines in Groton, CT. In 1997, he accepted a position with Lockheed Martin in Moorestown, NJ, and began a prolific career developing receivers/exciters and synthesizers for multiple radar and EW programs. This experience encompassed architecture definition, detailed design, rapid prototypes, manufacturing coverage, field installations, and coordination among many engineering disciplines. This work led the migration of phased array receiver/exciter electronics from centralized architectures to on-array digital beamforming systems.



Peter Delos

Jarrett Liner [jarrett.liner@analog.com] is an RF systems application engineer with Analog Devices, Inc., in the Aerospace and Defense Group in Greensboro, NC. He has significant experience in the area of RF system and component design.

Formerly, Jarrett was an applications engineer for GaN on SiC amplifiers for the military and aerospace sector. His prior experience also includes design and test of RF IC WLAN power amplifier and front-end modules for 13 years. He served 6 years in the United States Navy as an electronics technician. Jarrett received his B.S.E.E. from North Carolina Agricultural and Technical State University located in Greensboro, NC, in 2004.

When Jarrett isn't simulating circuit solutions or taking data in the lab, he might be found mountain biking, teaching cycle class at the gym, running, or chasing his four kids around the yard.



Jarrett Liner

Mirror, Mirror on the Wall—Understanding Image Rejection and Its Impact on Desired Signals

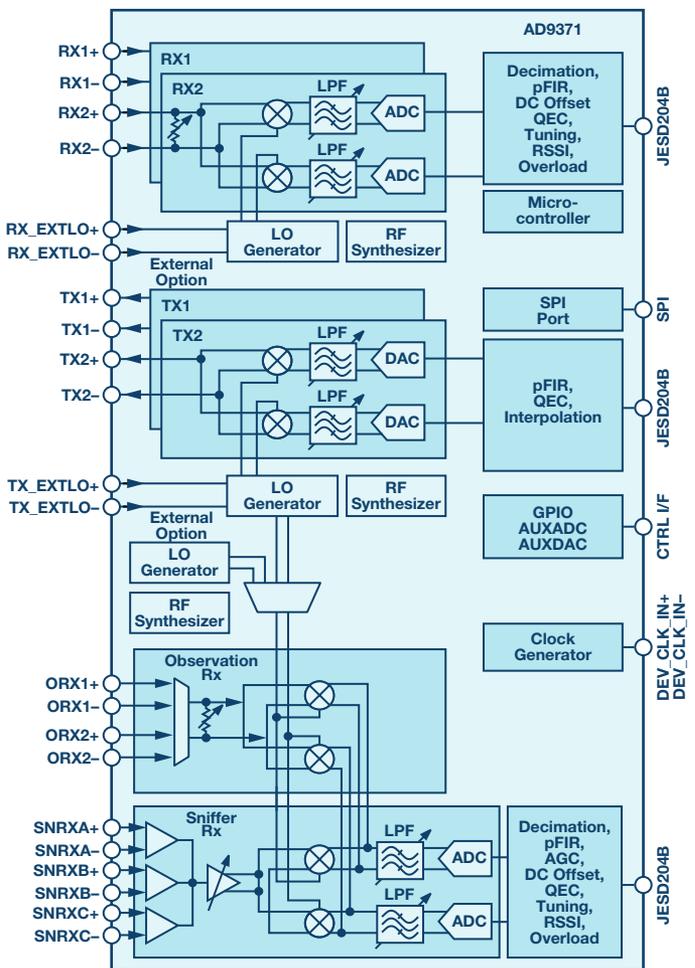
By Patrick Wiers

Share on   

The [AD9361](#) and [AD9371](#) RadioVerse wideband transceiver families both offer an unparalleled level of integration, as well as a plethora of features and many user-selectable options. The two families have significantly different levels of performance in several key areas and power consumption is also quite different between the two product families. Image rejection is one of the areas of performance that differs between the two families. This article explores where images come from, what they are, and how they impact overall system performance. Armed with that information, customers can make informed decisions and select the right transceiver for the application.

Image Rejection Basics

Both the AD9361 and AD9371 families use zero intermediate frequency (aka, zero-IF or ZIF) architectures, which results in a very high level of integration and significantly reduces the number of frequency-dependent components in the system. As shown in the block diagram of the AD9371 in Figure 1, both the main receive and transmit signal paths use one complex mixer stage to translate between the radio frequency (RF), which is centered around the local oscillator (LO) frequency, and baseband, which is centered around dc. For an excellent starting point to understanding the complex mixers used in ZIF transceivers, refer to the article on complex RF mixers referenced at the end of this article.¹



Notes
1. For JESD204B Pins, See Figure 4.

Figure 1. Block diagram of the RadioVerse AD9371 transceiver.

While providing many advantages due to such high integration, ZIF radios bring challenges as well. A complex mixer has in-phase (I) and quadrature-phase (Q) signals. Any mismatch of phase or amplitude of those signals degrades the summing and cancellation that occurs when the upconverted I and Q signals are combined. This is described in the article referenced above. Imperfect cancellation results in an inverted copy of the desired signal appearing on the opposite side of the local oscillator (LO) frequency from the desired signal when transmitting. This copy is referred to as an image and has a reduced amplitude compared to the desired signal. Similarly, an inverted copy of the desired signal appears on the opposite side of dc from the desired signal when receiving. In some other architectures (for example, super heterodyne), this image can be filtered at an intermediate stage. The primary advantage of the ZIF architecture is the removal of those filters and the intermediate mixer stage, but this necessitates very good I and Q balance to reduce the image amplitude to an acceptable level.

The simplified diagram of a receive signal path in Figure 2 shows where these mismatches occur with the mismatches designated by ΔA , Δf_c , and $\Delta\phi$. Only one path shows the added factors since it is the imbalance between the signal paths that creates the image, not the absolute gains and phases of the signal paths. Therefore, it is mathematically correct to show all of the imbalance factors in just one of the paths. The complex mixer shown in Figure 2 is also known as a quadrature mixer since the two LO signals provided to the mixers are orthogonal to each other.

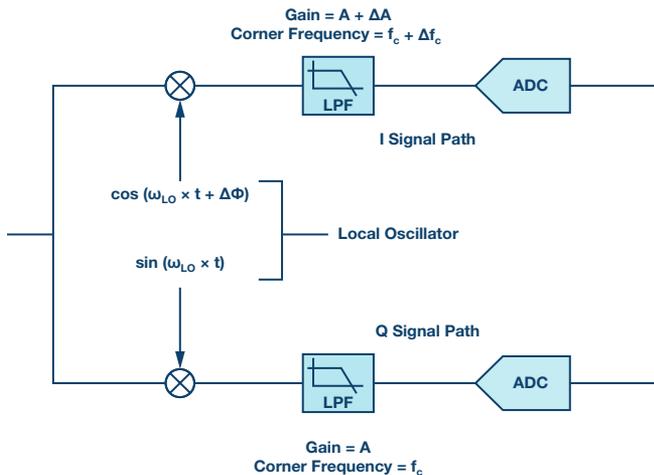


Figure 2. Simplified quadrature receiver signal path showing impairments.

Figure 3 shows an example using a single-tone or continuous wave (CW) desired signal and the resulting unwanted CW image. The desired signal is downconverted to frequency ω_c . If the quadrature balance is not perfect, the image will appear at frequency $-\omega_c$. The image rejection ratio (IRR) is the difference in dB between the desired signal and the unwanted image signal. Reducing the quadrature balance is known as quadrature error correction (QEC).

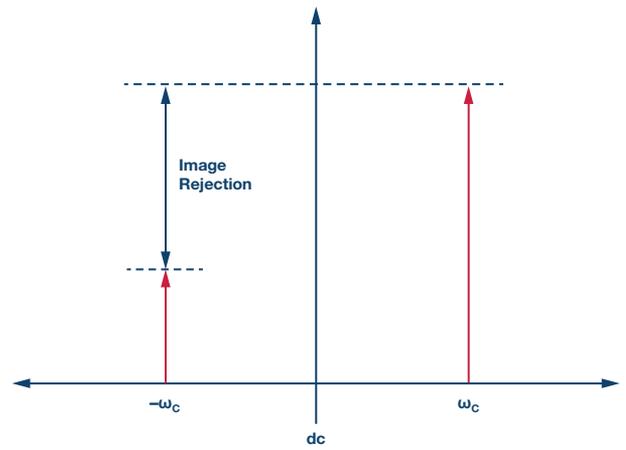


Figure 3. Single-tone desired signal and undesired image.

The image amplitude is related to the gain and phase mismatch by the following equation:

$$Image\ Rejection\ (dB) = 10 \times \log_{10} \frac{1 + 2 \times \sqrt{\Delta} \times \cos\theta + \Delta}{1 - 2 \times \sqrt{\Delta} \times \cos\theta + \Delta}$$

Where:

Δ = the amplitude imbalance in dB (a value of one is ideal)

θ = the phase error in degrees (a value of zero is ideal)

Equation 1 results in a two-dimensional matrix since image rejection will be degraded by both input variables independently. A portion of this matrix is shown in Figure 4, where the axis across the page is amplitude imbalance, the axis into the page is phase imbalance, and the vertical axis is the image rejection in dB. For example, if a 0.00195 amplitude error exists and the system needs to achieve an image rejection of 76 dB, the phase error must be better than 0.01286°. Even within a single integrated circuit device, it is very difficult to reach an image rejection better than 50 dB by controlling all of the factors that affect the I and Q matching. Achieving the 76 dB image rejection typically achieved with the AD9371 requires digital algorithms to manipulate analog path variables and to apply correction in the digital domain.

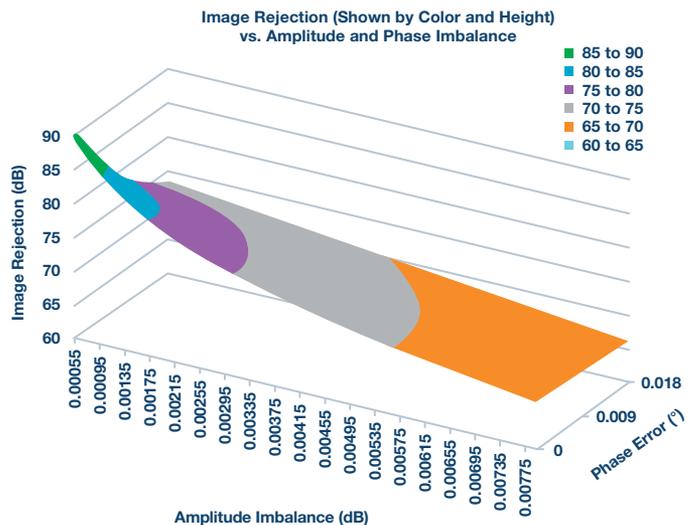


Figure 4. Image rejection in dB vs. amplitude imbalance in dB and phase imbalance in degrees.

The Effects of Images on Desired Signals

Figure 5 shows a simplified drawing of a single carrier case with the waveform centered at dc after downconversion. An example of this waveform would be a single instance of a 20 MHz LTE downlink OFDM signal. As seen in Figure 5, a portion of the desired signal on the negative side will have an image on the positive side and vice versa. In the single carrier case centered on dc, the image lands within (or on top of) the desired signal and corrupts the desired signal.

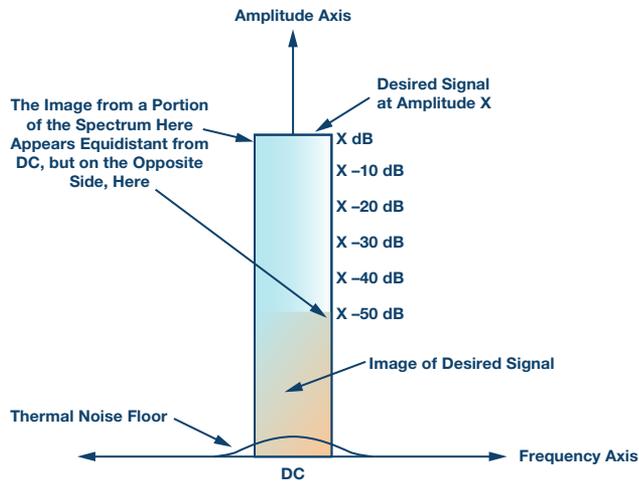


Figure 5. Single modulated carrier with undesired image.

When receiving and then demodulating a signal, there are several impairments that will be present. Thermal noise, which increases the noise floor of the receive signal path, is one example. Images add noise as well if they land within the desired signal. If the total of all noise sources is too high, then it will not be possible to demodulate the signal. The thermal noise floor shown in single carrier and multicarrier figures is an example and in these discussions is ignored as a contributor.

When using the AD9361's internal LO that is suitably supplied with a reference clock source with the recommended performance, the AD9361 will achieve approximately -40 dB EVM when not noise-floor limited. The EVM is limited to -40 dB by the phase noise of the RF PLL. The AD9361 image rejection performance of approximately 50 dBc means that in the single carrier case shown in Figure 5, the image alone will only degrade EVM by approximately 0.5 dB. Such low EVM degradation means that the transceiver will typically not be a limiting factor with modulation schemes such as 64-QAM and even higher. In this single-carrier case, the image is always smaller than the desired signal by approximately 50 dB and as shown in Figure 5.

Figure 6 shows a multicarrier example. Here, the desired signals are offset from dc after downconversion.

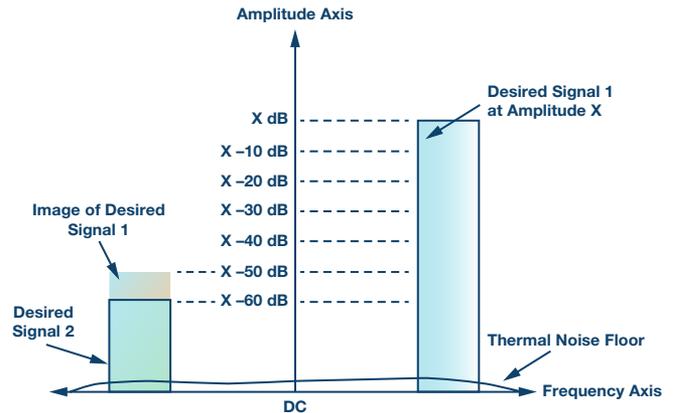


Figure 6. Multicarrier modulated signal with undesired image from Signal 1 corrupting Signal 2.

The image of each desired signal reflects through dc and shows up on the opposite side of the spectrum. In this example, two desired signals have been downconverted to the same offset from dc, desired Signal 1 on the positive side, and desired Signal 2 on the negative side. Note that the amplitude of desired Signal 2 is 60 dB below the amplitude of desired Signal 1. Two carriers with differing amplitudes is a common occurrence in multicarrier cases that arises when the signals from two mobile stations encounter different amounts of path loss when traveling to the same base station. This can occur when the two stations are different distances from the base station or one station transmits through or around different objects than the other station.

The amplitude of desired Signal 2 is 10 dB lower than the image from desired Signal 1. This represents a signal-to-noise ratio for desired Signal 2 of -10 dB. Demodulation, even using the simplest modulation techniques, would be extremely difficult if not impossible. Clearly, better image rejection performance is needed to handle these scenarios.

Figure 7 shows the same scenario but with receive image rejection performance typical of the AD9371.

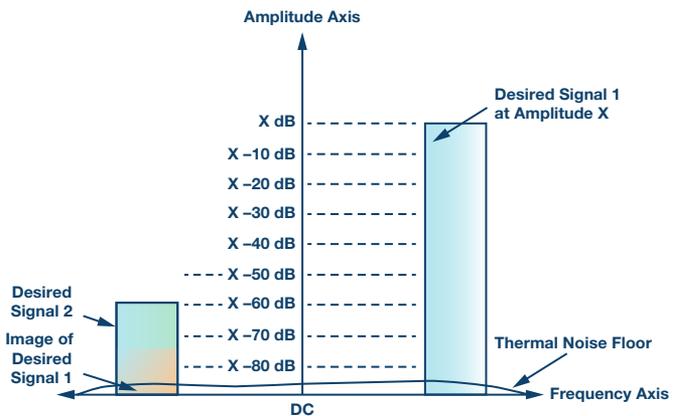


Figure 7. Multicarrier modulated signal with undesired image from Signal 1 and below Signal 2.

The image from desired Signal 1 now lies 15 dB below the amplitude of desired Signal 2. This provides 15 dB of signal noise ratio, which is enough to demodulate desired Signal 2 using a variety of modulation schemes.

Techniques Reducing Quadrature Imbalance in the AD9361 and the AD9371

Both the AD9361 and the AD9371 have optimized analog signal and LO paths that inherently reduce quadrature imbalance. As mentioned above, however, there is a limit to what good silicon can achieve. Digital correction can improve image rejection by several orders of magnitude.

The AD9361 receiver quadrature calibration uses an algorithm that analyzes the entire received spectrum of data to create an average correction across the entire bandwidth. For single-carrier use cases and relatively narrow bandwidth such as 20 MHz, this correction yields good image rejection across the bandwidth of interest. This is referred to as a frequency-independent algorithm. The algorithm operates on received data and updates in real time.

The AD9371 runs receive image rejection calibrations both during initialization by injecting test tones and during operation using actual received data. These more advanced calibrations adjust for frequency-dependent imbalance as well as frequency-independent imbalance. The algorithm updates in real time. The more advanced algorithm and circuitry that implements the corrections yield performance approximately 25 dB better than the AD9361 across the occupied signal bandwidth.

This article demonstrates the origination and effects of quadrature imbalance using the receive signal path, but a ZIF transceiver must overcome the same issues in the transmit signal path as well. The output of the transmitter includes both the desired signals as well as their images when the signal path or the LO paths are imbalanced.

For the transmit signal path, the AD9361 uses an initialization calibration to reduce quadrature imbalance below that provided by the optimized hardware design. The calibration uses a CW tone placed at a single frequency and at a single attenuation setting. The algorithm typically results in an image approximately 50 dB below the power of the desired signal. Another way to write this would be -50 dBc (decibels below carrier). Operation over temperature, across a wide bandwidth, or at different attenuation settings can affect the image level.

The AD9371 uses an initial transmit path calibration using many internally generated tones spread across the desired signal bandwidth and it determines the correction factors across multiple transmit attenuation settings. During operation, a transmit signal path tracking calibration uses real transmitted data and periodically updates the correction factors during operation. The AD9371 achieves approximately 15 dB better image rejection compared to the AD9361 and it provides this benefit over temperature, attenuation, and across the occupied signal bandwidth.

A Simplified Specific Example

Using all that has been covered in this article so far, let's perform a thought experiment and assume that we are building a system in which there is a central base station and multiple client devices. To simplify the example, the system will be operating away from objects such as buildings that cause multipath. The base station is expected to communicate with the client devices that can be spread throughout an area with a radius of 100 meters, as shown in Figure 8.

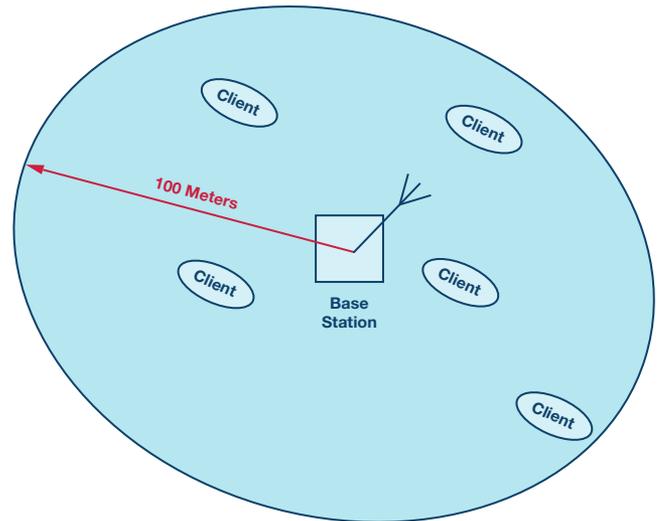


Figure 8. Imaginary cell coverage area showing base station and client stations.

The system will use multiple, simultaneously transmitted, 6 MHz wide carriers across a total bandwidth of 18 MHz. Let us say that in this system, one client unit could be very close to the base station—say, 0.3 meters—while the farthest could of course be 100 meters away. The free space path loss difference between the two would be approximately 50 dB. Let us also assume that the base station baseband processor can measure the received power and then tell the client to increase or decrease its transmit power by up to 10 dB. A nearby client would decrease by 10 dB while a client at the far end of the range would transmit at full power. This reduces the received power at the base station by 10 dB, resulting in an overall potential difference of 40 dB, as shown in Figure 9. The two displayed carriers represent the worst-case situation described above. For clarity, the optional carrier that could reside between the two desired signals has been omitted.

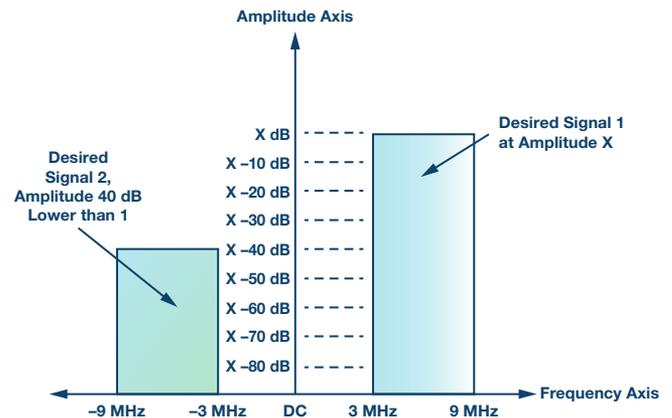


Figure 9. Multicarrier modulated signal example.

In this system, we will assume that the same transceiver will be used for both base station and client. If the AD9361 is used, the transmitted image can be approximately 50 dB below the amplitude of the desired signal. The receiver will also add a similar amount of image power. The two quadrature imbalances combine to create an image approximately 47 dB below the desired signal.

If the AD9371 is used for both ends of the link, then the transmitted image would typically be 65 dB down and the receiver would contribute an image 75 dB down below the desired signal. Adding those two up results in a total image of approximately 64.5 dB below the desired signal. Figure 10 shows both results.

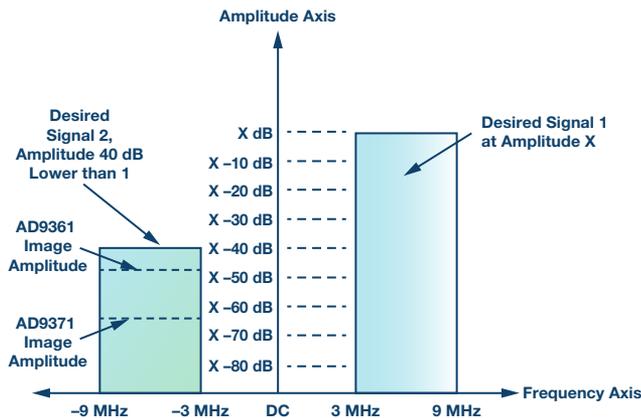


Figure 10. Multicarrier modulated signal example with image amplitudes for the AD9361 and AD9371.

In this simplified example, we are considering only the effects of the image and ignoring influences on SNR such as thermal noise, phase noise, and nonlinearities. Here, the AD9361 case results in an SNR of approximately 7 dB while the AD9371 case results in approximately 24.5 dB of SNR in this example. If complex modulation schemes such as 64-QAM are used in this system, then the total system SNR requirement would probably make the AD9371 the best choice. If simpler modulation schemes are

used, such as QPSK, the AD9361 may be more than adequate. Techniques used in the baseband processor would determine the actual system SNR needed to demodulate the signal. And, of course, to move from this thought experiment to a real system, the effects previously ignored such as thermal noise would have to be considered.

Conclusion

The previous plots and description of the quadrature correction algorithms in the two transceivers focused on the receive signal path. The effects of undesired images also apply to the transmit path for the same reasons. A transmission image that lands on top of a smaller carrier is equally troublesome for the station receiving the signals.

The section describing the techniques used by the transceivers to reduce image level shows the quantitative differences achieved by the two different families of devices. Then the specific example above takes us through a system design and narrows the design decision down to a few concise questions about the SNR required to demodulate the received signal. While the AD9371 family will always have better image performance than the AD9361 family, the higher power consumption of the AD9371 family and the use of a high speed serial interface require system engineers to look at all aspects of the design and find the best solution for their application.

References

1. Frizelle, Dave and Frank Kearney. "Complex RF Mixers, Zero-IF Architecture, and Advanced Algorithms: The Black Magic in Next-Generation SDR Transceivers." *Analog Dialogue*, Vol. 51, 2017.

Patrick Wiers [patrick.wiers@analog.com] is the applications manager for the Transceiver Products Group at Analog Devices Greensboro. He joined ADI in 2007, starting as an applications engineer for transceivers. Prior to ADI, he worked on board level circuit design for defense and audio products and also spent five years as an FAE for software-defined radio and WLAN products.



Patrick Wiers

Wireless Current Sense Circuit Floats with Sense Resistor

By Kris Lokere

Share on   

Introduction

Measuring the current that flows through a sense resistor seems easy. Amplify the voltage, read it with an ADC, and now you know what the current is. However, it gets more difficult if the sense resistor sits at a voltage that is very different from system ground. The typical solutions that bridge that voltage differ in either the analog or digital domain. But here is a different approach—wireless.

Analog current sense ICs are compact solutions. However, the voltage difference that they can withstand is limited by semiconductor processes. It is difficult to find devices that are rated for more than 100 V, and these circuits often lose accuracy if the sense resistor's common-mode voltage changes quickly or swings both above and below system ground.

Digital isolation techniques (magnetic or optical) are a bit more bulky, but work without loss of accuracy and can typically withstand thousands of volts. These circuits need an isolated power supply, but that can sometimes be integrated in the isolator component. If the sense resistor is physically separated from the main system, then you also may need to run long wires or cables.

A wireless current sense circuit overcomes many of these limitations. By allowing the entire circuit to float with the common mode of the sense resistor and transmitting the measured data wirelessly over the air, there are no voltage limitations at all. The sense resistor can be located anywhere without the need to run cables. If the circuit is very low power, then you don't even need an isolated power supply and can instead run for many years on a small battery.

Design Overview

Figure 1 shows the block diagram of the design. The current sense circuit is based on the [LTC2063](#) chopper-stabilized op amp to amplify the voltage drop across a sense resistor. The micropower SAR ADC [AD7988](#) digitizes the value and reports the result via an SPI interface. The [LTP5901-IPM](#) is the radio module that contains not only the radio, but also the networking firmware needed to automatically form an IP-based mesh network. In addition, the LTP5901-IPM has a built-in microprocessor that reads the AD7988 ADC SPI port. The [LTC3335](#) is a low power, dc-to-dc power supply that converts the battery voltage to a constant output voltage. The LTC3335 also includes a coulomb counter that reports the cumulative charge pulled from the battery.

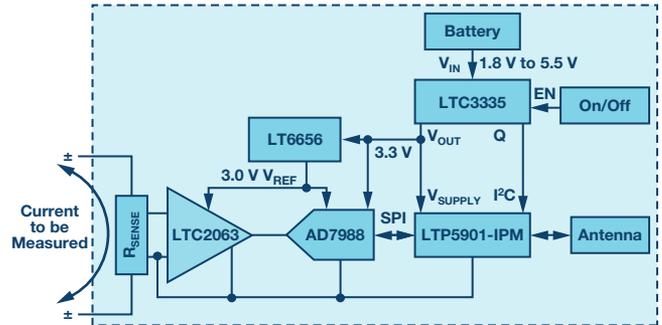
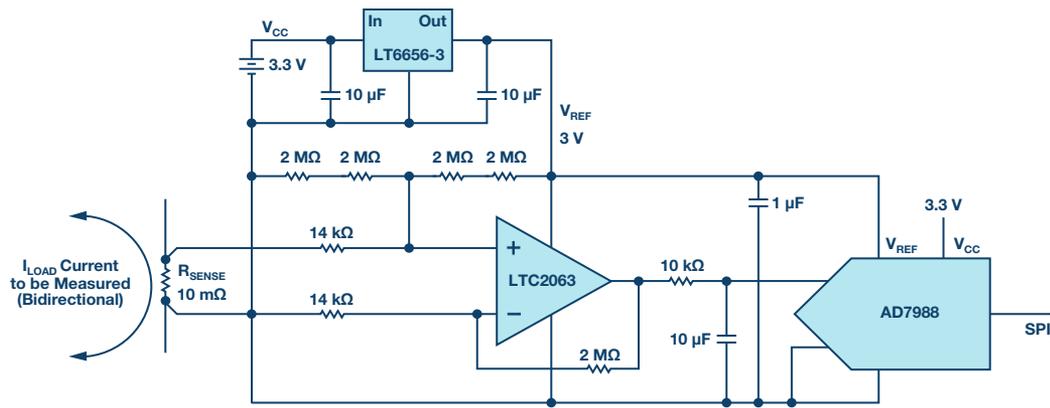


Figure 1. A low power, wireless current sense circuit is formed by a low power chopper op amp to amplify the sense voltage. It is digitized using a low power ADC and reference and connected to a SmartMesh IP wireless radio module. A low power, dc-to-dc converter conditions the battery and also keeps track of the charge drawn from the battery.

Signal Chain

The LTC2063 is an ultra low power, chopper-stabilized op amp. With a maximum supply current of 2 μA , it is uniquely suited for use in battery powered applications. Because the offset voltage is less than 10 μV , it can measure very small voltage drops without loss of accuracy. Figure 2 shows the LTC2063 configured to gain up and level shift the voltage across a 10 m Ω sense resistor. The gain is chosen so that the ± 10 mV full-scale input at the sense resistor (corresponding to ± 1 A of current) maps to a near full-scale range at the output, centered around 1.5 V. This amplified signal is fed into a 16-bit SAR ADC. The AD7988 was chosen for its very low standby current and good dc accuracy. At low sample rates, the ADC automatically shuts down in between conversions, resulting in an average current consumption of as little as 10 μA at 1 kSPS. The LT6656 biases the amplifier, the level-shift resistors, and the ADC's reference input. The LT6656 voltage reference consumes less than 1 μA and can drive up to 5 mA loads with low dropout, making it easy to output a precise 3 V, even when powered from the 3.3 V system supply.

There are three roughly equal sources to offset error in this signal chain, together contributing about 0.5% relative to a ± 10 mV full-scale input. They are the offset voltage of the LTC2063 and AD7988, as well as mismatch in the level-shift resistors (0.1% resistors are recommended). A one-point calibration step could largely eliminate that offset. Gain error is generally dominated by inaccuracies in available sense resistors, which tend to be worse than the 0.05%, 10 ppm/ $^{\circ}\text{C}$ specifications of the LT6656 voltage reference.



(Some Supply Bypass Capacitors Omitted for Clarity)

Figure 2. The current sense circuitry floats with the sense resistor voltage. The LTC2063 chopper op amp amplifies the sense voltage and biases it mid-rail for the AD7988 ADC. The LT6656-3 provides the precision 3 V reference.

Power Management

The LTC3335 is a nanopower buck-boost converter with an integrated coulomb counter. It is configured to provide a regulated 3.3 V output from an input supply between 1.8 V and 5.5 V. This allows the circuit to be powered by two alkaline primary battery cells. For duty-cycled wireless applications, the load current can easily vary from 1 μ A to 20 mA, depending on whether the radio is in active or sleep mode. The LTC3335 has a quiescent current of just 680 nA at no load, which keeps the entire circuit operating at very low power when the radio and signal chain are in sleep mode. Still, the LTC3335 can output as much as 50 mA, which easily provides enough power during radio transmit/receive and for a variety of signal chain circuits.

The LTC3335 also has a handy, built-in coulomb counter. When switching, it keeps track of the total charge that it draws from the battery. This information can be read out using an I²C interface and can then be used as a predictor for when it may become time to replace the batteries.

Wireless Networking

The LTP5901-IPM is a complete wireless radio module that includes the radio transceiver, embedded microprocessor, and SmartMesh IP networking software. The LTP5901-IPM performs two functions in this application: wireless networking and housekeeping functions (processes). When multiple SmartMesh IP motes are powered up in the vicinity of a network manager, the motes automatically recognize each other and form a wireless mesh network. The entire network is automatically time-synchronized, which means that each radio is only powered on during very short, specific time intervals. As a result, each node can function not only as a source of sensor information, but also as a routing node to relay data from other nodes toward the manager. This creates a highly reliable, low power mesh network where multiple paths are available from each node to the manager, even though all nodes, including the routing nodes, operate on very low power.

The LTP5901-IPM includes an ARM[®] Cortex[®]-M3 microprocessor core that runs the networking software. In addition, users may write application firmware to perform tasks specific to the user application. In this example, the microprocessor inside the LTP5901-IPM reads the SPI port of the current measurement ADC (AD7988) and reads the I²C port of the coulomb counter (LTC3335). The microprocessor can also put the chopper op amp (LTC2063) in shutdown mode, further reducing its current consumption

from 2 μ A to 200 nA. This provides additional power savings in use models with extremely long intervals between measurements.

Overall Power Consumption

The total power consumption of the complete application circuit depends on various factors, including how often the signal chain takes a reading and how the nodes are configured in the network. Typical power consumption for a mote reporting once per second is less than 5 μ A for the measurement circuit and can be 40 μ A for the wireless radio, allowing years of operation on small batteries.



Figure 3. A complete wireless current sense circuit is implemented on a small PCB. The only physical connections are the banana jacks for the current to be measured. The wireless radio module is shown on the right. The circuit is powered from two AAA batteries connected on the back of the board.

Conclusion

Combining Linear Technology and Analog Devices signal chain, power management, and wireless networking products enables the design of a truly wireless current sense circuit. Figure 3 shows an example implementation. The new ultra low power LTC2063 chopper op amp can accurately read small voltage drops across a sense resistor. The entire circuit, including the micropower ADC and voltage reference, floats with the common mode of the sense resistor. The nanopower LTC3335 switcher can power the circuit for years from a small battery, while reporting cumulative battery usage with its built-in coulomb counter. The LTP5901-IPM wireless module manages the entire application and automatically connects to a highly reliable SmartMesh IP network.

Kris Lokere [kris.lokere@analog.com] is a strategic applications manager for signal chain products, joining ADI as part of the merger with Linear Technology. Kris enjoys architecting systems that combine technologies from multiple product lines. In the past 20 years, Kris has designed op amps, built engineering teams, and managed product line strategy. Kris holds multiple patents and received an M.S.E.E. degree from Katholieke Universiteit Leuven and an M.B.A. from Babson College.



Kris Lokere

Rarely Asked Questions—Issue 144

Accelerometer Tilt Measure over Temperature and in the Presence of Vibration

By Chris Murphy

Share on   



Question:

My consumer grade accelerometer can theoretically measure $<1^\circ$ of tilt. Will this still be possible over temperature and in the presence of vibration?

Answer:

Most likely, the answer is no. Questions around definitive tilt accuracy values are always difficult to answer, as many environmental factors need to be accounted for when it comes to MEMS sensor performance. Typically, consumer grade accelerometers struggle to detect less than 1° of tilt in dynamic environments. To show this, a general-purpose, consumer grade accelerometer is compared with a next-generation, low noise, low drift, and low power MEMS accelerometer. The comparison looks at a number of error sources present in tilt applications and what errors can be compensated or removed.

Errors such as zero- g bias accuracy, zero- g bias shift due to soldering, zero- g bias shift due to PCB enclosure alignment, zero- g bias tempco, sensitivity accuracy and tempco, nonlinearity, and cross-axis sensitivity are observable and can be reduced through the postassembly calibration processes. Other error terms such as hysteresis, zero- g bias shift over life, sensitivity shift over life, zero- g shift due to humidity, and PCB bend and twist due to temperature variations over time can't be addressed in cali-

bration or else they require some level of in-situ servicing to be reduced. For this comparison it is assumed that cross-axis sensitivity, nonlinearity, and sensitivity are compensated, as they require much less effort to minimize compared to tempco offset drift and vibration rectification.

Table 1 shows an estimate of the consumer grade ADXL345 accelerometer's ideal performance specifications and the corresponding tilt errors. When trying to achieve the best possible tilt accuracy, it is imperative to apply some form of temperature stabilization or compensation. For this example, a constant temperature of 25°C is assumed. The largest error contributors that can't be fully compensated out are offset over temperature, bias drift, and noise. Bandwidth can be lowered to reduce the noise as inclination applications typically require bandwidths below 1 kHz.

Table 1. ADXL345 Error Source Estimates

Sensor Parameter	Exemplary Performance	Condition/Note	Typical g	Application Error Tilt $^\circ$
Noise	x/y axis 290 $\mu\text{g}/\sqrt{\text{Hz}}$	Bandwidth at 6.25 Hz	0.9 mg	0.05 $^\circ$
Bias drift	Allan deviation	Short-term (for example, 10 days)	1 mg	0.057 $^\circ$
Initial offset	35 mg	No compensation	35 mg	2 $^\circ$
		With compensation	0 mg	0 $^\circ$
Error	No compensation	6.25 Hz bandwidth	36.9 mg	2.1 $^\circ$
Error	With compensation	6.25 Hz bandwidth	1.9 mg	0.1 $^\circ$

Table 2 shows the same criteria for the ADXL355. Short-term bias values were estimated from the root Allan variance plots in the ADXL355 data sheet. At 25°C , the compensated tilt accuracy is 0.1° for the general-purpose ADXL345 and for the industrial grade ADXL355 the tilt accuracy is 0.005° . Comparing the ADXL345 and ADXL355, it can be seen that larger error contributors like noise have been reduced significantly from 0.05° to 0.0045° and that bias drift has been reduced from 0.057° to 0.00057° . This shows the massive leap forward in MEMS capacitive accelerometer performance in terms of noise, temperature coefficient, offset, and bias drift, enabling much higher levels of inclination accuracy under dynamic conditions.

Table 2. ADXL355 Error Source Estimates

Sensor Parameter	Exemplary Performance	Condition/Note	Typical <i>g</i>	Application Error Tilt °
Noise	25 $\mu\text{g}/\sqrt{\text{Hz}}$	Bandwidth at 6.25 Hz	78 μg	0.0045°
Bias drift	Allan deviation	X/Y axis short-term (for example, 10 days)	<10 μg	0.00057°
Initial offset	25 mg	No compensation	25 mg	1.43°
		With compensation	0 mg	0°
Total error	No compensation	6.25 Hz bandwidth	25 mg	1.43°
Total error	With compensation	6.25 Hz bandwidth	88 μg	0.005°

The importance of selecting a higher grade accelerometer is crucial in achieving required performance, especially if your application demands <1° tilt accuracy. Application accuracy can vary depending on application conditions (large temperature fluctuations, vibration) and sensor selection (consumer grade vs. industrial or tactical grade). In this case, the ADXL345 will require extensive compensation and calibration effort to achieve <1° tilt accuracy, adding to the overall system effort and cost. Depending on the magnitude of vibrations in the end environment and temperature range, this may not even be possible. Over 25°C to 85°C, the tempco offset drift is 1.375°—already exceeding the requirement for less than 1° of tilt accuracy.

$$0.4 \frac{\text{mg}}{\text{C}} \times \frac{1^\circ}{17.45 \text{ mg}} \times (85^\circ\text{C} - 25^\circ\text{C}) = 1.375^\circ$$

For the ADXL355 the tempco offset drift from 25°C to 85°C is

$$0.15 \frac{\text{mg}}{\text{C}} \times \frac{1^\circ}{17.45 \text{ mg}} \times (85^\circ\text{C} - 25^\circ\text{C}) = 0.5^\circ$$

Vibration rectification error (VRE), as shown in Table 3, is the offset error introduced when accelerometers are exposed to broadband vibration. When an accelerometer is exposed to vibrations, VRE contributes significant error in tilt measurements when compared to 0 *g* offset over temperature and noise contributions. This is one of the key reasons it is left off data sheets, as it can very easily overshadow other key specifications.

Chris Murphy [christopher.murphy@analog.com] is an applications engineer with the European Centralized Applications Center, based in Dublin, Ireland. He has worked for Analog Devices since 2012, providing design support on motor control and industrial automation products. He holds a master's of engineering in electronics by research and a bachelor's of engineering in computer engineering.



Chris Murphy

Table 3. Errors Shown in Degrees of Tilt

	Maximum Tilt Error 0 <i>g</i> Offset vs. Temperature(°/°C)	Noise Density (°/√Hz)	Vibration Rectification (°/ <i>g</i> ² rms)
ADXL354	0.0085	0.0011	0.023'
ADXL355	0.0085	0.0014	0.023'

¹±2 *g* range in a 1 *g* orientation offset due to 2.5 *g* rms vibration.

In environments with higher magnitude vibrations, it is essential to have a higher *g* range accelerometer in order to minimize clipping leading to offsets. Table 4 shows the ADXL35x family of accelerometers and their corresponding *g* range and bandwidths.

Table 4. ADXL354/ADXL355/ADXL356/ADXL357 Measurement Ranges

Part	Measurement Range (<i>g</i>)	Bandwidth (kHz)
ADXL354B	±2, ±4	1
ADXL354C	±2, ±8	1
ADXL355B	±2, ±4, ±8	1.5
ADXL356B	±10, ±20	1.5
ADXL356C	±10, ±40	1.5
ADXL357B	±10.24, ±20.48, ±40.96	1

Choosing a part from the ADXL35x family for tilt applications will provide high stability and repeatability, and this family is robust to temperature fluctuation and broadband vibration, as well as requires less compensation and calibration compared to a lower cost accelerometer. The hermetic package helps ensure that the end product conforms to its repeatability and stability specifications long after it leaves the factory. By providing repeatable tilt measurement under all conditions, Analog Devices' next-generation accelerometers enable minimal tilt error without extensive calibration in harsh environments, as well as minimizing the need for post-deployment calibration.

GaN Breaks Barriers—RF Power Amplifiers Go Wide and High

By Keith Benson

Share on   

Introduction

The increasing demand for higher data rates in telecommunications and higher resolution in industrial systems is pushing the frequency of operation higher for the electronics that support them. Many of these systems operate over a wide frequency spectrum and further increased bandwidth requirements are a common request for new designs. Across many of these systems there is a push to use one signal chain for all frequency bands. Advancements in semiconductor technology have led to breakthroughs in capability for high power and wideband amplifiers. An area that was once dominated by traveling wave tubes has begun to cede ground to semiconductor devices thanks to the GaN revolution that is sweeping the industry and enabling MMICs that generate >1 W of power over many decades of bandwidth. As shorter gate length GaAs and GaN transistors become available, coupled with improved circuit design techniques, new devices are becoming available that can perform comfortably to millimeter wave frequencies, opening new applications that were hard to contemplate a decade ago. This article will briefly describe the state of the semiconductor technology that is enabling these developments, circuit design considerations to achieve optimum performance, and examples of both GaAs and GaN wideband power amplifiers (PAs) that demonstrate today's technology.

Many wireless electronic systems operate over wide frequency ranges. In the military industry, there are radar bands from a few hundred MHz to many GHz. There are electronic warfare and electronic countermeasure systems required to work over very wide bandwidth. Threats can come at a variety of frequencies, such as MHz to 20 GHz, or even higher frequencies today. As more electronics become available at higher frequencies, the need for higher frequency electronic warfare systems will proliferate. In telecommunications, base stations operate from 450 MHz to ~3.5 GHz and continue to increase as the need for more bandwidth continues. Satellite communications systems operates from mainly C-band to Ka-band. Instrumentation used to measure these different electronics needs to work over all required frequencies to be universally accepted. As a result, the systems engineer faces challenges trying to design electronics to cover the entire frequency range. Given the possibility of having one signal chain cover the entire frequency range, most systems engineers and procurement folks would be very excited. There are many advantages to having one signal chain cover the entire frequency range, including simpler design, faster time to market, less component inventory to manage, and more. The challenge with the one signal chain approach is always related to the performance degradation that comes with a wide-band solution vs. a narrow-band solution. At the heart of this challenge is the power amplifier, which commonly has superior performance in terms of power and efficiency when tuned over a narrow bandwidth.

Semiconductor Technology

In years past, traveling wave tube (TWT) amplifiers have dominated higher power electronics as the output power amplifier stage in many of these systems. There are some nice attributes to TWTs, including capability of kW's of power, operation over octaves or even multiple octaves of bandwidth,

high efficiency in back off condition, and good stability over temperature. There are some drawbacks to TWTs that include poor long-term reliability, lower efficiency, and they need very high voltage to operate (~1 kV or higher). Given the long-term reliability of semiconductor ICs, there has been a push toward these electronics for many years, starting with GaAs. When possible, many systems engineers have worked to combine multiple GaAs ICs to generate large output power. Entire companies have been created based entirely on combining technology and doing it efficiently. There are many different types of combining technologies, such as spatial combining, corporate combining, etc. These combining techniques all suffer from the same fate—combining has loss and, ideally, you would not have to use these combining techniques. This motivates us to use high power electronics to start the design. The easiest way to increase the RF power from a power amplifier is to increase the voltage, which has made gallium nitride transistor technologies so attractive. If we compare the various semiconductor process technology, we can see how the power generally increases with high operating voltage IC technology. Silicon germanium (SiGe) technology uses a relatively low operating voltage of 2 V to 3 V, but is very attractive for its integration benefits. GaAs has been used widely for power amplifiers for many years in microwave frequencies and has operating voltages of 5 V to 7 V. Silicon LDMOS technology operating at 28 V has been used for many years in telecommunications, but it is mainly useful below 4 GHz, so it's not as widely used in broadband applications. The emergence of GaN technology operating at 28 V to 50 V on a low loss, high thermal conductivity substrate like silicon carbide (SiC) has opened up a range of new possibilities. Today, GaN on silicon technology is limited to operation below 6 GHz. The RF losses associated with the silicon substrate and its lower thermal conductivity compared to SiC compromises the gain, efficiency, and power as the frequency increases. Figure 1 shows a comparison of various semiconductor technologies and how they compare to each other.

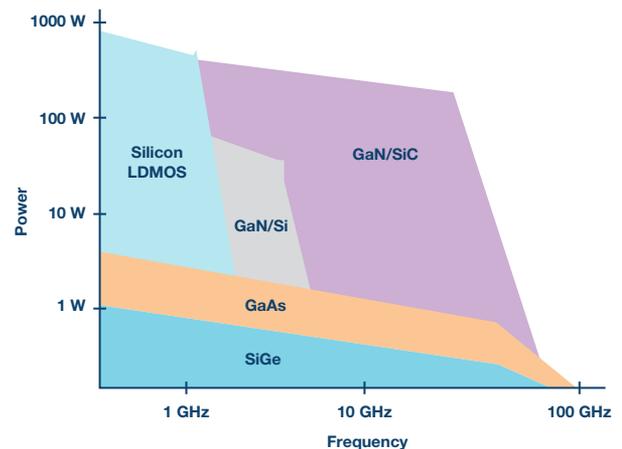


Figure 1. A process technology comparison of microwave frequency range power electronics.

The emergence of GaN technology has created an industry shift away from TWT amplifiers and a move toward GaN amplifiers as the output stage of many of these systems. The driver amplifier in many of these systems is still commonly GaAs, as much of this technology already exists and continues to be improved. Next we will look at how to use circuit design to extract as much power, bandwidth, and efficiency out of these wideband power amplifiers. Certainly GaN-based designs are capable of higher output powers than GaAs-based designs, and the design considerations are largely the same.

Design Considerations

There are different topologies and design considerations for the IC designer to use when choosing how to start a design to optimize power, efficiency, and bandwidth. The most common type of monolithic amplifier design is a multistage, common-source, transistor-based design, also known as a cascade amplifier design. Here the gain multiplies from each stage, leading to high gain and allowing us to increase output transistor sizes in order to increase the RF power. GaN offers benefits here because we are able to greatly simplify the output combiners, reducing loss, and thereby improving efficiency, as well as shrinking the die size, as shown in Figure 2. As a result, we are able to achieve wider bandwidths and improve performance. A less obvious benefit of going to GaN devices from GaAs is to achieve a given RF power level, perhaps 4 W—the transistor size will be less, which will result in higher gain per stage. It will lead to fewer stages per design and ultimately higher efficiency. The challenge with this cascade amplifier technique is that it is difficult to achieve bandwidths over an octave without significantly compromising the power and efficiency, even with the help of GaN technology.

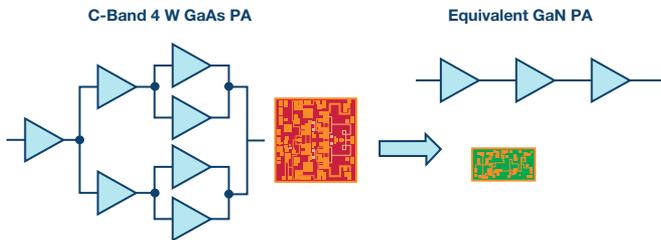


Figure 2. Comparison of multistage GaAs PA vs. an equivalent GaN PA.

Lange Coupler

One approach to achieve wide bandwidth design is to implement a balanced design with Lange couplers on the RF input and output, shown in Figure 3. Here the return loss is ultimately dependent on the coupler design, as it becomes easier to optimize the gain and power response over frequency without also needing to optimize the return loss. Even while using Lange couplers, it becomes more difficult to achieve bandwidths over an octave, but they do offer a very nice return loss for the design.

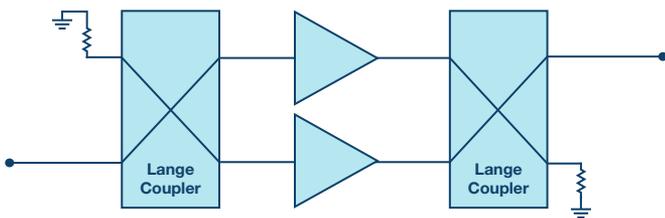


Figure 3. Balanced amplifier using Lange couplers.

Distributed Amplifier

The next topology to consider is the distributed power amplifier shown in Figure 4. The benefit of a distributed power amplifier is accomplished by incorporating the parasitic effects of the transistor into the matching networks between devices. The input and output capacitances of the device can be combined with the gate and drain line inductance, respec-

tively, to make the transmission lines virtually transparent, excluding transmission line loss. By doing this, the gain of the amplifier should only be limited by the transconductance of the device and not the capacitive parasitics associated with the device. This only happens if the signal traveling down the gate line is in phase with the signal traveling down the drain line, so that each transistor's output voltage adds in phase with the previous transistors output. The signal traveling to the output will constructively interfere so that the signal grows along the drain line. Any reverse waves will destructively interfere since these signals will not be in phase. The gate line termination is included to absorb any signals that are not coupled to the gates of the transistors. The drain line termination is included to absorb any reverse traveling waves that could destructively interfere with the output signal and improve the return loss at low frequencies. As a result, multiple decades of bandwidth are able to be realized from kHz to many GHz. This topology is popular when more than an octave of bandwidth is needed and there are some nice benefits, such as flat gain, good return loss, high power, etc. An illustration of a distributed amplifier is shown in Figure 4.

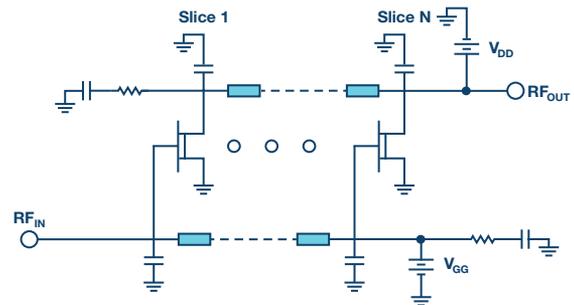


Figure 4. A simplified block diagram of a distributed amplifier.

One challenge here with distributed amplifiers is that the power capability is dictated by the voltage applied to the device. Since there is no narrow-band tuning capability, you are essentially providing a 50 Ω impedance to the transistor or close to it. When we consider Equation 1, the average power of a power amplifier, R_L or optimum load resistance, essentially becomes 50 Ω. Therefore, the achievable output power is set by the voltage applied to the amplifier such that if we want to increase the output power, we need to increase the voltage applied to the amplifier.

$$P_{OUT} = \frac{V_{DD}^2}{2 \times R_L}$$

This is where GaN becomes very helpful, as we can quickly go from a 5 V supply voltage with GaAs to a 28 V supply voltage in GaN, and the achievable power goes from 0.25 W to almost 8 W simply by changing from GaAs to GaN technology. There will be other considerations to look at such as the gate length of the process available in GaN and if they can achieve the gain you need at the high frequency end of the band. As time progresses, more of these GaN processes become available.

The fixed R_L of 50 Ω for distributed amplifiers is different compared to the cascaded amplifier where we change the resistance value presented to the transistor by matching networks to optimize the power from the amplifier. There is a benefit in optimizing the resistance value presented to the transistor with cascade amplifiers in that it can improve the RF power. Theoretically, we can continue increasing the transistor periphery size to continue increasing the RF power, but there are practical limitations to this such as complexity, die size, and combining loss. The matching networks also tend to limit the bandwidth, as they become difficult to provide an optimum impedance over wide frequencies. In the distributed power amplifier there are only transmission lines whose purpose is to have the signals constructively interfere along the amplifier, rather than matching networks. There are additional techniques to further improve the power in distributed amplifiers, such as using a cascode amplifier topology to further increase the voltage supply to the amplifier.

Results

We have shown that there are various techniques and semiconductor technologies that offer trade-offs in providing optimum power, efficiency, and bandwidth. Each of these different topologies and technologies will likely have a place in the semiconductor world, as they each provide benefits, which is why they have survived to this day. Here we'll focus on a few results that we believe show what is possible with these technologies today to achieve high power, efficiency, and bandwidth.

Today's Product Capability

We'll look at a GaAs-based, distributed power amplifier operating from dc to 30 GHz, which is a product released from Analog Devices, [HMC994A](#). This part is interesting because it covers many decades of bandwidth, lots of different applications, and achieves high power and efficiency. The performance is shown in Figure 5. Here we see a saturated output power covering MHz to 30 GHz with over 1 W of power and a power added efficiency (PAE) of 25% nominal. This particular product also has a strong third-order intercept (TOI) performance of 38 dBm nominal. This result shows that with GaAs-based designs we are able to achieve an efficiency that is close to what is possible with many narrow-band power amplifier designs. The positive gain slope with frequency, high PAE, wideband power performance, and strong return loss make the HMC994A an interesting product.

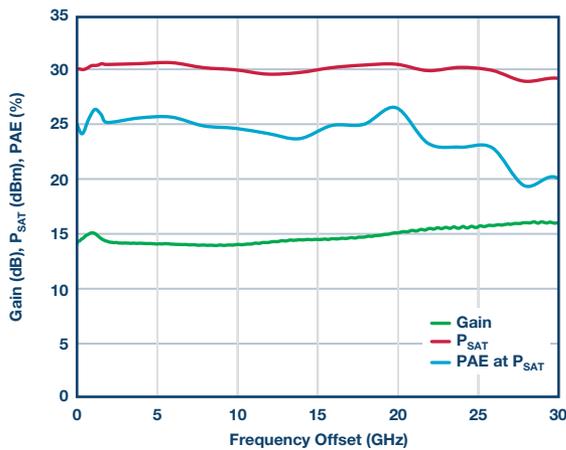


Figure 5. HMC994A gain, power, and PAE vs. frequency.

It's also interesting to see what is achievable with GaN-based technology. Analog Devices offers a standard product, [HMC8205BF10](#), which is based in GaN and combines high power, efficiency, and bandwidth. This product operates from a 50 V power supply and provides 35 W of RF power at 35% nominal efficiency with ~20 dB of power gain covering over a decade of bandwidth. In this case, a single IC is able to provide roughly 10× more power compared to similar approaches in GaAs. In years past, this would have required a complicated combining scheme of GaAs die that would not have been able to reach the same efficiency. This product demonstrates what is possible with GaN technology covering wide bandwidths and providing high power and efficiency, as shown in Figure 6. It also shows progress in high power electronic packaging technology, as this part is housed in a flange package capable of supporting continuous wave (CW) signals needed for many military applications.

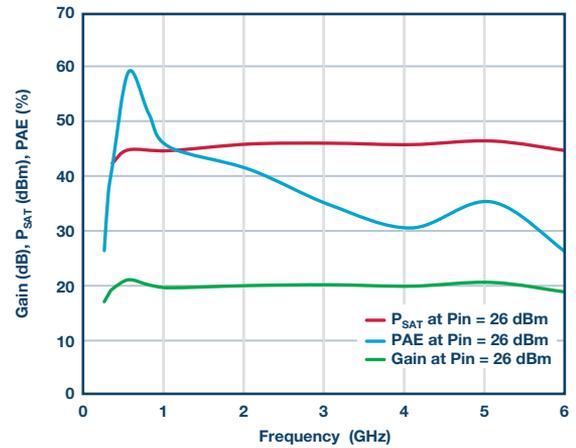


Figure 6. HMC8205BF10 power gain, P_{SAT}, and PAE vs. frequency.

Summary

The emergence of new semiconductor materials like GaN have opened the possibilities to reach higher power levels covering wide bandwidths. Shorter, gate length GaAs devices have extended frequency ranges from 20 GHz to 40 GHz and beyond. The reliability of these devices is shown in literature to exceed 1 million hours, making them ubiquitous for modern day electronic systems. We expect the trends of higher frequencies and wider bandwidth to continue into the future.

Keith Benson [keith.benson@analog.com] graduated from the University of Massachusetts, Amherst with a B.S.E.E. in 2002 and from the University of California, Santa Barbara with an M.S.E.E. in 2004. He spent much of his early career at Hittite Microwave designing ICs used in RF wireless electronics. He then moved to manage a team of IC design engineers focusing on wireless communication links. In 2014, Hittite Microwave was acquired by Analog Devices where he became a product line director for RF/microwave amplifiers and phased array ICs. Keith currently holds three U.S. patents related to novel amplifier techniques.



Keith Benson

Intelligent Video Analytics at the Edge of IoT

By Tyler Jesiel

Share on   

A Mix of Logarithmic Imaging and Node Analytics

Video analytics applications operating within the Internet of Things (IoT) can be improved by using node analytics and logarithmic imagers. Video analytics applications attempt to take advantage of the abundant information in the everyday world for several reasons. These reasons range from facial recognition to everyday surveillance, but the majority are centered on predictive and behavioral analysis. The information gathered in these applications can be processed extensively at a higher level by cloud computing. However, in-depth processing has its limitations and can be improved in many ways by adding node analytics and logarithmic imagers to the mix.

By adding node analytics to the mix, data analysis can be improved by eliminating communication with the cloud. Cloud computing requires two, if not three orders of magnitude of more bandwidth than node analytics applications. As a result, node analytics consumes much less computational power and latency is decreased. Densely populated markets, chaotic traffic sections, and city parking spaces are some of the intricate atmospheres that can be sensed for predictive and behavioral analysis using node analytics. High level processing of these environments performed in the cloud can advance business strategies, divert general traffic flow, and improve the efficacy of government-managed parking spaces. However, by implementing lower level software at the sensor node instead of performing analysis in the cloud—latency, bandwidth, security, and power can be improved in these scenarios.

In addition to intelligence at the node, adding logarithmic imagers to the mix can strengthen these systems by providing advantages in areas where conventional imagers fall short. Logarithmic imagers provide a higher dynamic range for image processing, in addition to reduced dependence on luminosity changes. For example, shadows, reflections, abrupt changes in light, and high contrast scenarios are the areas in which a logarithmic imager can outperform a conventional imager. The resolution of these issues in video applications can bolster data capture, which strengthens the analytics at the node. By improving data capture, the overall video analytics application can be significantly improved.

The improvements offered by node analytics and logarithmic imagers can help solve problems in video analytics applications operating within the IoT. Some of the engineering difficulties common to IoT applications are security, decision making latency, data bandwidth, and computational power. These engineering problems are greatly reduced with the elimination of data transmission, which is why at-the-node analytics is attractive for IoT applications. In video analytics applications, limited

contrast and dependence on luminosity is a common struggle. This struggle is almost eliminated with logarithmic imagers, which is key for video analytics applications. Overall, video analytics applications operating within IoT are enhanced with node analytics and logarithmic imagers.

Intelligence at the Edge

By processing the data based upon expected visual events, the measured data can be transformed quickly into appropriate actions with little or no data communicated to a cloud server. This quick analysis of video data, instead of transmitting to the cloud, localizes the decision making process and improves the latency in the system. The latency in decision making is not only reduced significantly, but the security is enhanced by eliminating the transmission of data that would normally introduce the risk of interception.

Only the most valuable information needs to be connected beyond the node and into the cloud for predictive or behavioral processing. This optimized partitioning of data maximizes cloud value, because full bandwidth of a video analytics frame is typically not needed. Most of the visual data from frame to frame is static on a fixed mounted camera that can be filtered at the node. Edge node video analytics can provide many filtered interpretations to differentiate expected object types: car, truck, bicycle, human, animal, etc. This decimation reduces data bandwidth and associated computational power that otherwise would be required within a cloud server to analyze the full frame rate of video data transmitted downstream. This reduction in bandwidth can be a two or three orders of magnitude improvement over cloud computing applications, which is a key improvement offered by node analytics.

Logarithmic Imaging

Video analytics applications can be improved further by addressing common concerns related to conventional imagers with the substitution of logarithmic imaging. Most conventional imagers are linear and use pixels that generate a voltage that is a linear function of light, which can result in a limited contrast. Linear imagers also utilize a uniform exposure phase that limits their dynamic range to the pixel exposure time within the frame rate. Lastly, conventional imager contrast is dependent on luminosity, which can introduce reflection-related contrast issues. These common problems are eliminated with the substitution of a logarithmic imager by using pixels that generate a voltage that is a logarithmic function of light. The differences between a traditional imager and a logarithmic imager can be observed in Figures 1, 2, and 3.



Figure 1. Conventional image exposure.



Figure 2. ADIS17002 logarithmic image exposure.



Figure 3. ADIS17002 logarithmic image exposure with built-in edge detection.

Some conventional imagers struggle with contrast-related difficulties that hinder users from fully capturing their target environment. These contrast issues stem from the linear properties of the voltage generation within each pixel. The voltage that is generated within a linear imaging pixel is directly proportional to the amount of photons striking it; as result, the dynamic range is limited compared to its logarithmic counterpart. The reduced contrast associated with these linear imagers is a result of the reduced dynamic range. This reduced contrast can corrupt analytics at the sensor node in IoT applications, ultimately affecting the overall system performance. The logarithmic imager provides for a wider range of light levels and, thus, increased contrast due to the pixel voltage being logarithmically generated. However, this increased contrast results in higher sensitivity to light, which can be an undesired effect in some applications. Alternatively, this increased sensitivity to luminosity can be an advantage—it all depends on the application.

Video capture with conventional imagers can be further obstructed by reflections in sunny or bright environments. For example, facial recognition in vehicles can become increasingly difficult when reflections are present on the windshield. This obstruction of video capture can negatively impact video analytics by introducing errors in the system or by the loss of crucial data. These reflections are introduced because linear imager contrast between pixels is dependent on luminosity; as a result, reflections are more prominent. This dependence on luminosity can be observed in Equation 1. Alternatively, logarithmic imager contrast is independent of luminosity, due to its natural logarithmic properties, which helps mitigate reflections or abrupt changes in light. The independence of luminosity for logarithmic imagers can be observed by Equation 2.

$$\text{Linear Image Contrast}_{AB} = \text{Pixel}_A I - \text{Pixel}_B I = I(\text{Pixel}_A - \text{Pixel}_B) \quad (1)$$

$$\begin{aligned} \text{Logarithmic Image Contrast}_{AB} &= \log(\text{Pixel}_A I) - \log(\text{Pixel}_B I) = \\ \log\left(\frac{\text{Pixel}_A I}{\text{Pixel}_B I}\right) &= \log\left(\frac{\text{Pixel}_A}{\text{Pixel}_B}\right) = \log(\text{Pixel}_A) - \log(\text{Pixel}_B) \end{aligned} \quad (2)$$

Beyond Individual Components

Analog Devices is going beyond individual components in order to offer platform-level solutions; these solutions help customers rapidly deploy proven intelligent solutions with improved performance at a lower system cost. Smart applications start with reliable and accurate data that is achieved by ADI's advanced sensing and measuring capabilities. In addition, Analog Devices is partnering up with customers to develop unique, system-level solutions that approach the whole problem. One of these solutions, which is capable of quarter video graphics array (QVGA) imaging analytics, is the ADIS1700x.

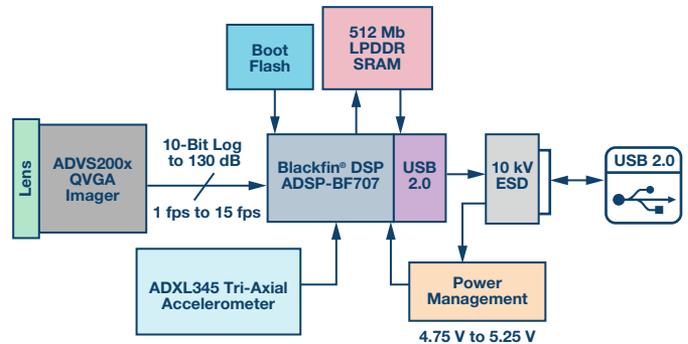


Figure 4. Functional block diagram.

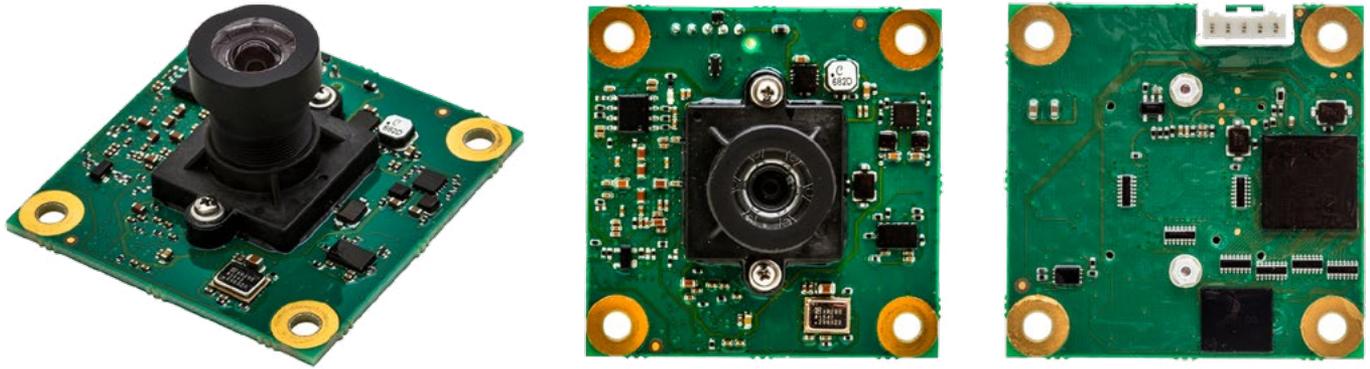


Figure 5. ADIS17002 at angle (left), above at the lens side of the board (center), and back side (right).

The ADIS1700x is a QVGA analytics imager module in a small form factor with logarithmic sensitivity, combined with digital signal processing that optimizes video performance. The module utilizes a low power Blackfin processor to offer the ability to perform analytics at the node, in addition to an accelerometer for image stabilization, tilt, and impact detection. It also utilizes built-in edge detection for tracking and counting object motion. Each $14\ \mu\text{m} \times 14\ \mu\text{m}$ pixel has a unique exposure phase, unlike conventional imagers. Conformal coating for outdoor operation make this a perfect module for mass deployment that allows the creation of emerging smart city and building applications. The [ADIS17001](#) provides a 110° field of view (FOV) lens, while the ADIS17002 provides a 67° FOV lens. These two options provide for a diverse range of target applications, which include parking spot monitoring, parking violation enforcement, traffic queue detection, and industrial analytics.

Overall, video applications in IoT can be significantly improved with node analytics and logarithmic imagers, which is the approach Analog Devices is taking with the release of the ADIS1700x. Node analytics, instead of cloud computing, can benefit IoT applications moving forward. Logarithmic imagers have benefits that its counterpart cannot match, improving IoT applications even further. In summary, video analytics applications operating within the IoT combined with node analytics and logarithmic imagers makes for a robust, system-level solution.

Tyler Jesiel [tyler.jesiel@analog.com] is a product engineer in the Autonomous Inertial MEMS Group at Analog Devices, primarily working on system in package (SiP) solutions and inertial measurement units (IMUs). His experience includes characterization, design verification, and test automation. He received his B.S.E.E. from the University of Michigan, Ann Arbor, and he is actively pursuing an M.S.E.E. at North Carolina State University.



Tyler Jesiel

Uncompromising Linearity from the LTC2185 and ADA4927-1

By Clarence Mayott

Share on   

The **LTC2185** is a 125 MSPS, 16-bit ADC with excellent noise and linearity performance while only consuming 185 mW per channel. It is ideal for demanding, low power applications that require excellent ac performance. A high performance ADC like the LTC2185 requires a high performance amplifier driving it to maintain its excellent performance. The **ADA4927-1** delivers the linearity performance required by the LTC2185 while only consuming 215 mW. The well designed package of the ADA4927-1 allows for a simple layout that reduces parasitic capacitance in the feedback path that can erode the phase margin of the amplifier. This combination of ADC and driver allows for excellent performance from 62.5 MHz to 125 MHz a region, where other high speed amplifiers are lacking.

The LTC2185 is a 2-channel, simultaneous sampling parallel ADC that offers a choice of full-rate CMOS or double data rate (DDR) CMOS/LVDS digital outputs. Pin-compatible speed grade options include 25 MSPS, 40 MSPS, 65 MSPS, 80 MSPS, and 105 MSPS with approximate power dissipation of just 1.5 mW/MSPS per channel. It includes popular features such as the digital output randomizer and alternate bit polarity (ABP) mode that minimize digital feedback when using parallel CMOS outputs. Analog full power bandwidth of 550 MHz and ultralow jitter of 0.07 ps rms allow under-sampling of IF frequencies with excellent noise performance. To maintain this level of performance, the LTC2185 needs to be driven with an appropriate amplifier like the ADA4927-1.

The ADA4927 is a high speed, differential current feedback amplifier. Fabricated on Analog Devices' silicon-germanium process, the ADA4927-1 has excellent distortion and an input voltage noise of only 1.3 nV/√Hz. This allows it to drive high speed ADCs like the LTC2185. The gain of the ADA4927-1 is set with external feedback resistors located next to the input pins. By keeping the feedback pins and input pins close on the package, the ADA4927-1 provides a clean layout and minimizes the parasitic capacitance in the feedback network. This makes the ADA4927-1 an ideal choice for driving high performance ADCs, like the LTC2185, from dc to 125 MHz.

Figure 1 shows a schematic of the ADA4927-1 driving the LTC2185. The corresponding layout is shown in Figure 2. The feedback pins on the ADA4927-1 are adjacent to the input pins, which minimizes the parasitic capacitance of the feedback node and improves the phase margin of the amplifier. It also simplifies the layout by making it possible to place feed-

back resistors directly across the two pins and not having additional trace length in the feedback path. There is a simple filter between the amplifier and ADC that reduces the wideband noise of the amplifier and improves the SNR of the system. This filter also attenuates the sampling glitches from the ADC before they reach the amplifier. This helps keep the output network of the ADA4927 from oscillating in response to these glitches. This filter network can be modified to accommodate a wide range of input bandwidth requirements.

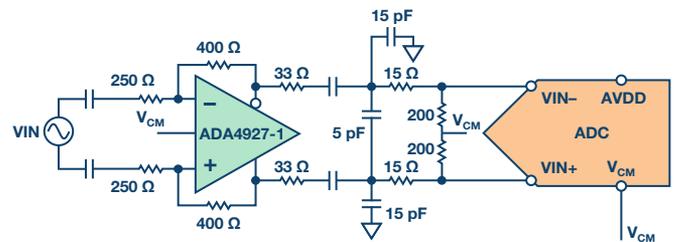


Figure 1. Schematic showing an ADA4927-1 driving one channel of the LTC2185.

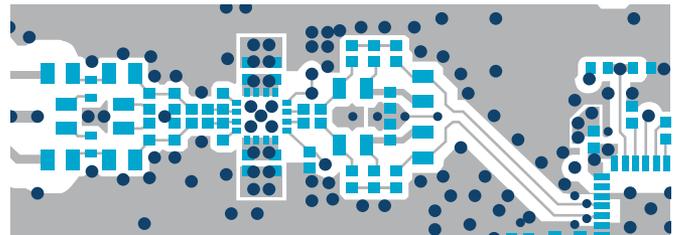


Figure 2. Layout showing an ADA4927-1 driving one channel of the LTC2185.

Figure 3 and Figure 4 show the SNR and SFDR of the LTC2185 and ADA4927-1 combination. The SFDR stays above 67 dB out to 125 MHz while the SNR is better than 63 dB to the same frequency. This combination only consumes 250 mW. With a sample rate of 125 MSPS, this combination provides good performance through the entire second Nyquist zone where other amplifiers begin to have poor linearity.

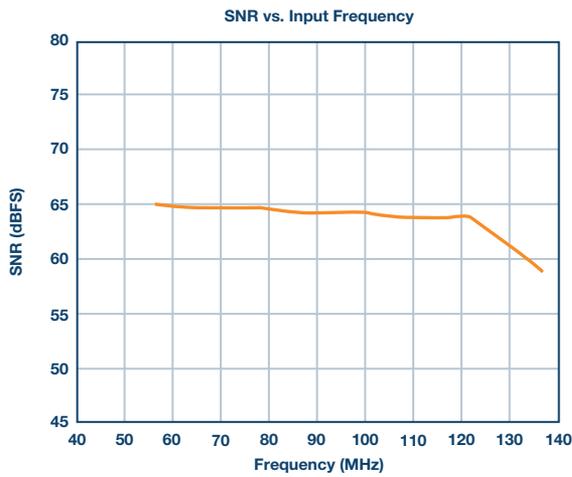


Figure 3. SNR of the LTC2185 driven with the ADA4927-1.

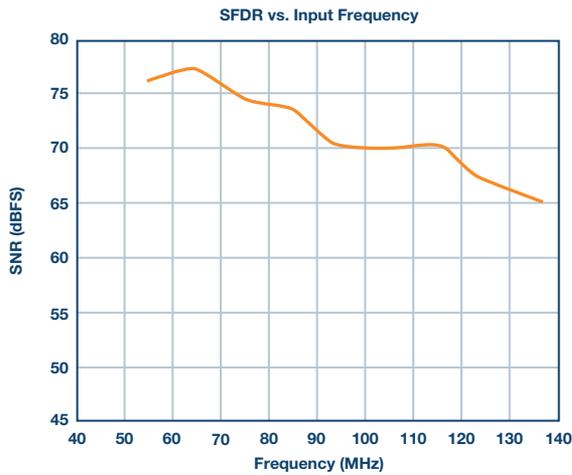


Figure 4. SFDR of the LTC2185 driven with the ADA4927-1.

Using the ADA4927-1 to drive the LTC2185 provides excellent linearity while keeping the power consumption low. The fact that the ADA4927-1 stays very linear out to 125 MHz allows this ADC amplifier combination to be used in demanding communication and medical applications that require the use of the second Nyquist zone of the LTC2185. The pin out of the ADA4927-1 and filter design minimize the complexity of the layout while maintaining excellent performance on a low power budget.

Clarence Mayott [clarence.mayott@analog.com] is a mixed-signal application section leader with over 10 years of experience at Linear Technology.

Clarence designed demo boards with complete signal chains combining amplifiers and ADC combinations to help end customers evaluate systems more easily. He also designed companion boards, including clock and signal source boards, to help facilitate the evaluation of high speed ADC demo boards. Clarence manages the continued development of PScope, the software used for various pipeline and SAR ADCs.

With the release of the LTC2000, Clarence has expanded his knowledge base to include high speed DACs and waveform generation, in addition to high speed ADCs. As an application section leader, he oversees the continued development of LTDACGen—a new software tool for generating complex waveforms for high speed DACs.

Clarence holds an M.S. degree in electrical engineering from Santa Clara University and a B.S. degree in electrical engineering from California Polytechnic State University, San Luis Obispo.



Clarence Mayott

Rarely Asked Questions—Issue 145

A Low Power, Low Cost, Differential Input to a Single-Ended Output Amplifier

By Chau Tran and Jordyn Rombola

Share on   



Question:

How do I make a low cost, low power, differential input into a single-ended output amplifier?

Answer:

Introduction

In many applications, there are requirements of low power, high performance differential amplifiers to convert small differential signals to a readable ground referenced output signal. Input voltages at two inputs usually share a large common-mode voltage. The differential amplifier rejects the common-mode voltage and the remaining voltage is amplified and presented on the amplifier output as single-ended voltage. The rejected voltage can either be ac or dc, and this common-mode voltage is typically larger than the differential input voltage. The effectiveness of the rejection diminishes as the frequency of the common-mode voltage increases. Amplifiers within the same package have better matching, the same parasitic capacitance, and no external wiring is needed. Therefore, a high performance, high bandwidth dual amplifier performs better with frequency than discrete amplifiers.

A simple solution is to use dual precision amplifiers with a resistive gain network, as shown in Figure 1. This circuit shows a simple way of converting a differential input to a single-ended output with adjustable gain. The gain of the system can be set by Equation 1:

$$V_{OUT} = Gain \times (V_{IN1} - V_{IN2}) \quad (1)$$

Where gain = $RF/1 \text{ k}\Omega$ and $(V_{IN1} - V_{IN2})$ is the differential input voltage

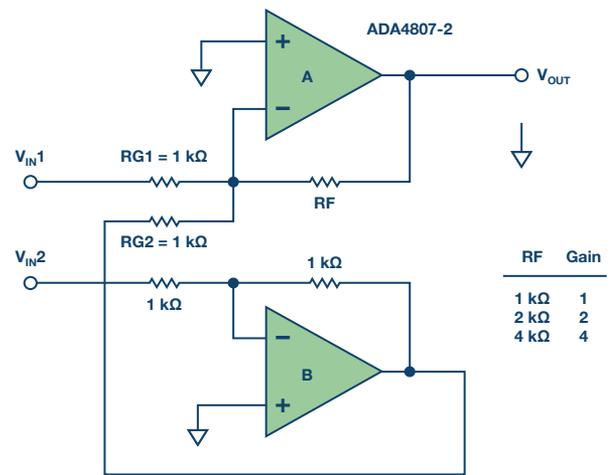


Figure 1. The differential input, single-ended output amplifier.

Generally, this technique provides a more stable reading when EMI or RFI is present and, therefore, it is recommended when noise is a problem. This is especially true when measuring thermocouple, strain gage, and bridge type pressure sensor inputs, since they produce very small signals in a noisy environment.

This circuit can offer a performance improvement over single-ended inputs by not only measuring the difference in voltage between the positive and negative terminals of the sensor, but also providing common-mode rejection with some system gain. Furthermore, the sensor ground can be different from the analog ground. An output voltage referenced to ground is important in many applications. The accuracy of the system depends on the tolerance of the network resistors.

The circuit can convert a differential input to a single-ended output with an adjustable gain. The gain of the system can be set by the ratio of RF and $RG1$ with the assumption that $RG2 = RG1$ and amplifier B has a gain of -1 .

For example, the ADA4807-2, a 180 MHz dual amplifier, can be built as an inverting amplifier for this application and this circuit provides lower noise. With a low quiescent current of 1000 μA per amplifier, the circuit is ideal for low power, high resolution data conversion systems.

The input common mode can go beyond the supply voltages. The output is rail to rail, which makes it useful in the presence of a large common-mode signal or a large output voltage applications. For example, the data acquisi-

tion board has an ADC that accepts a single-ended input of 0 V to 5 V. However, the signal source happens to be differential voltage generated from a sensor bridge where one terminal swings positive while the other swings negative in response to pressure in the presence of common noise.

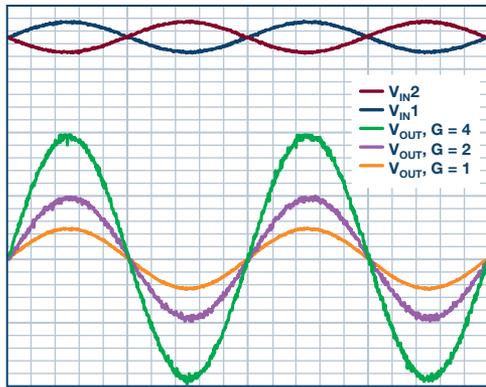


Figure 2. The performance of the simple differential to single-ended amplifier.

The plot in Figure 2 was taken by applying a differential input voltage and varying the gain of the circuit. The value of R_F sets the gain of the system. As one can observe, the plot shows the system gain of 1, 2, and 4, with the differential input voltage of 1 V p-p at 1 kHz.

The circuit is useful for measuring a small difference between two large voltages. For example, consider a solution with a simple 1% accuracy of monitoring a typical Wheatstone bridge circuit that is excited by 3 V/GND in a 3 V battery-powered system. Using 1% resistors or better will allow for the level of accuracy needed, and the circuit will reject any common mode and amplify the attenuated bridge signal by the gain that the circuit is set for. If driving an ADC, some level-shifting will need to be applied to get the output signal in the range of 0 V to 5 V.

This circuit offers a combination of excellent distortion with low quiescent current. The 2 op amp solution results in lower system cost and the use of a differential amplifier yields better performance.

Chau Tran [chau.tran@analog.com] joined Analog Devices in 1984 and works in the Integrated Amplifier Products (IAP) Group in Wilmington, MA. In 1990, he graduated with an M.S.E.E. degree from Tufts University. Chau holds more than 10 patents and has authored more than 10 technical articles.



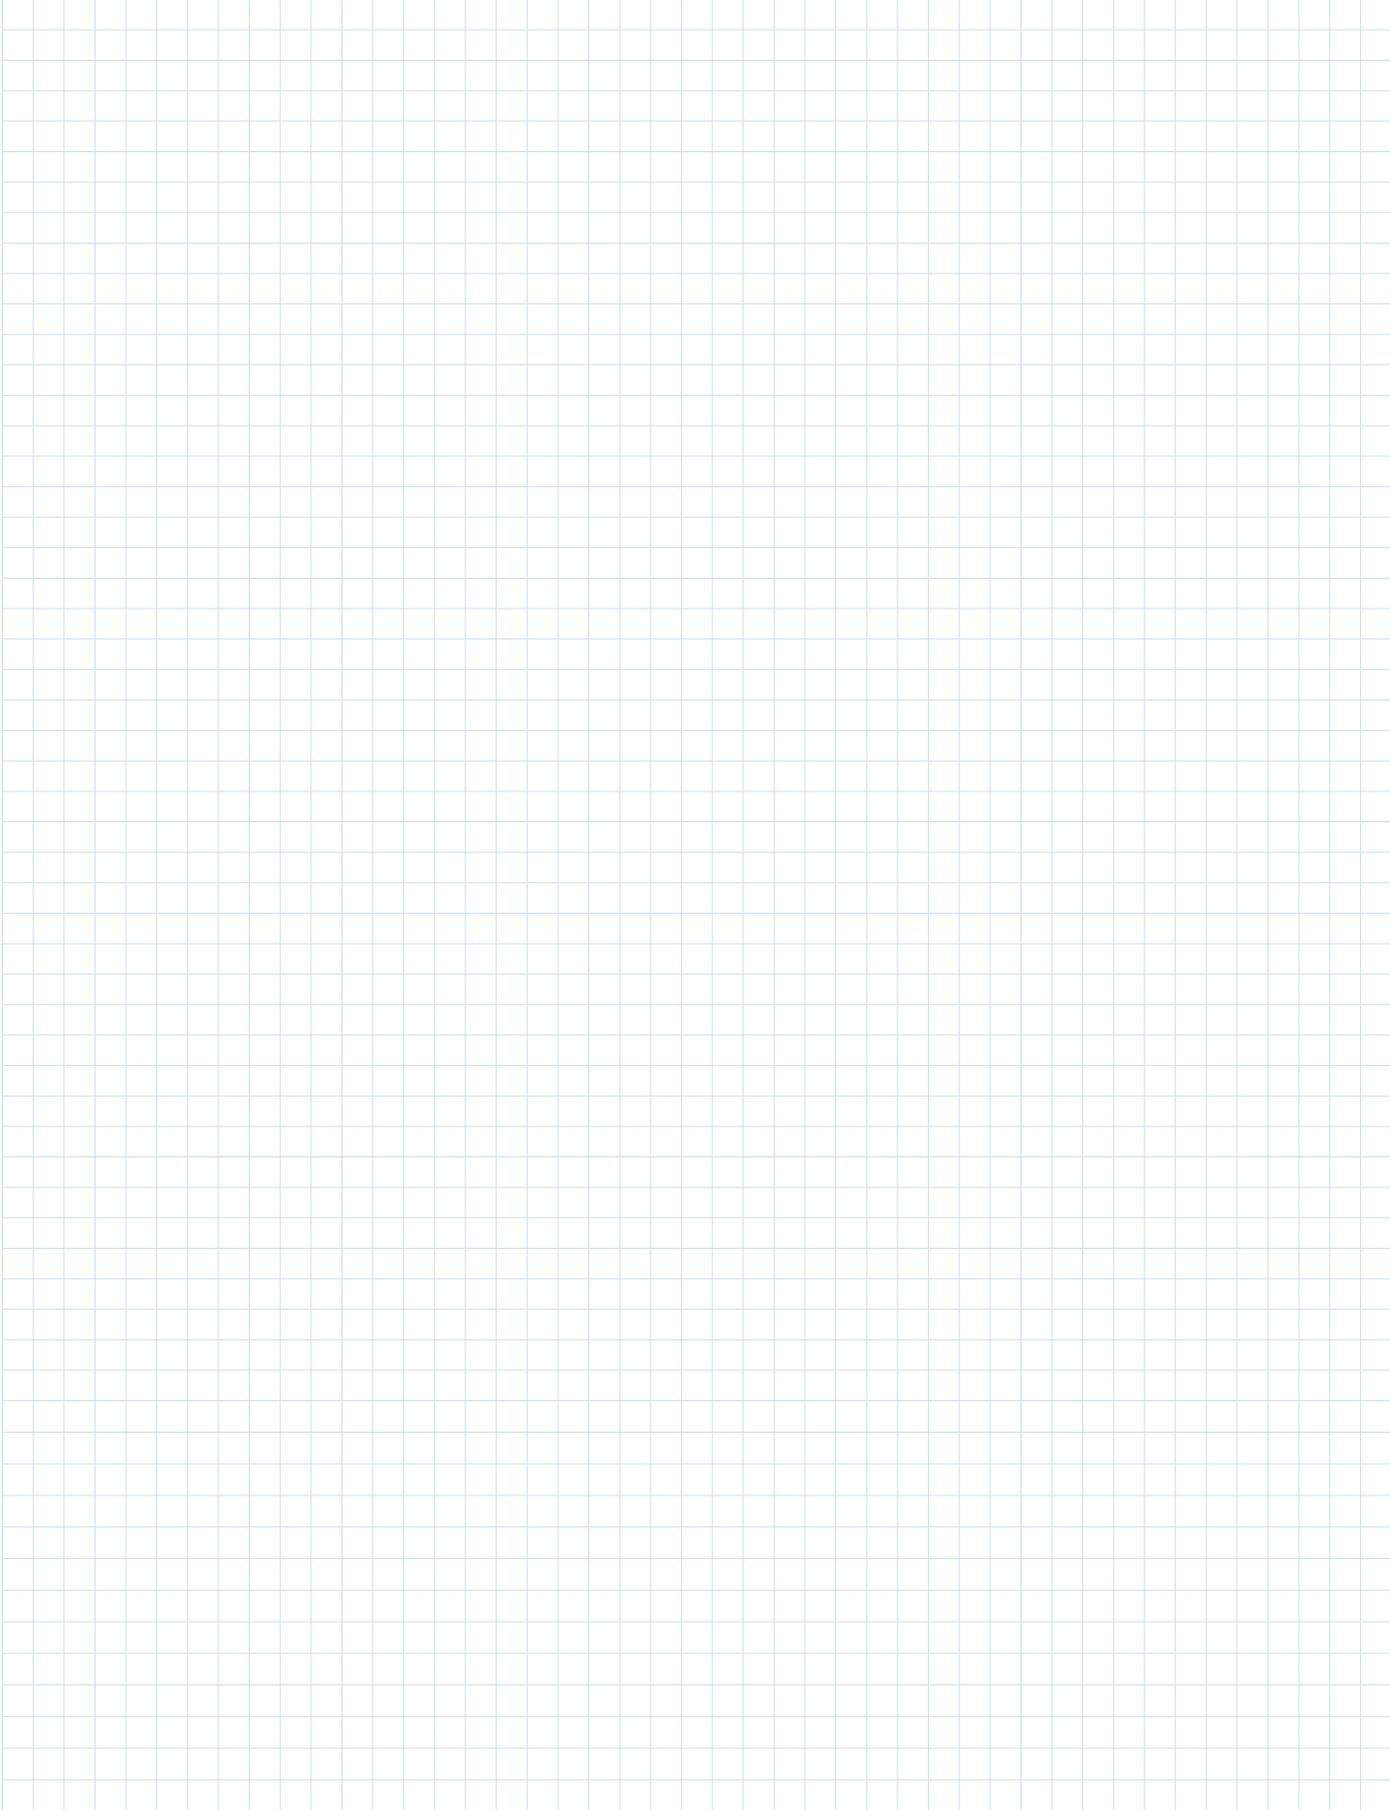
Chau Tran

Jordyn Rombola [jordyn.rombola@analog.com] is a product engineer in the Linear and Precision Technology (LPT) Group here at ADI. She joined Analog Devices in January 2014 after finishing her bachelor's degree in electrical and computer engineering from Worcester Polytechnic Institute (WPI).



Jordyn Rombola

Notes



Stay Connected with *Analog Dialogue*

Your Engineering Resource for Innovative Design

Be the first to receive the latest technology, applications, and systems-level insights needed to solve real-world design challenges by subscribing to *Analog Dialogue*. And be sure to tell a friend! analogdialogue.com



Like *Analog Dialogue* to discover the latest articles and author insights.
facebook.com/analogdialogue



EngineerZone® Support Community

Engage with the ADI technology experts in the community to ask your tough design questions, browse our rich knowledge base, or hear from our authors in the *Analog Dialogue* blog. ez.analog.com, ezchina.analog.com (China)



Facebook

[facebook.com/
AnalogDevicesInc](https://facebook.com/AnalogDevicesInc)



@ADI_News

[twitter.com/
ADI_News](https://twitter.com/ADI_News)



LinkedIn

[linkedin.com/
company/analog-devices](https://linkedin.com/company/analog-devices)



Sina Weibo (China)

[weibo.com/
analogdevices](https://weibo.com/analogdevices)



WeChat (China)



Analog Devices, Inc. Worldwide Headquarters

Analog Devices, Inc.
One Technology Way
P.O. Box 9106
Norwood, MA 02062-9106
U.S.A.
Tel: 781.329.4700
(800.262.5643, U.S.A. only)
Fax: 781.461.3113

Analog Devices, Inc. Europe Headquarters

Analog Devices GmbH
Ott-Aicher-Str. 60-64
80907 München
Germany
Tel: 49.89.76903.0
Fax: 49.89.76903.157

Analog Devices, Inc. Japan Headquarters

Analog Devices, KK
New Pier Takeshiba
South Tower Building
1-16-1 Kaigan, Minato-ku,
Tokyo, 105-6891
Japan
Tel: 813.5402.8200
Fax: 813.5402.1064

Analog Devices, Inc. Asia Pacific Headquarters

Analog Devices
5F, Sandhill Plaza
2290 Zuchongzhi Road
Zhangjiang Hi-Tech Park
Pudong New District
Shanghai, China 201203
Tel: 86.21.2320.8000
Fax: 86.21.2320.8222

©2017 Analog Devices, Inc. All rights reserved. Trademarks and registered trademarks are the property of their respective owners. Ahead of What's Possible is a trademark of Analog Devices. MO2000-0-10/17

analog.com



AHEAD OF WHAT'S POSSIBLE™