Analog Dialogue

Volume 56, Number 1, 2022

Your Engineering Resource for Innovative Design

Visit *analogdialogue.com*

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

In This Issue

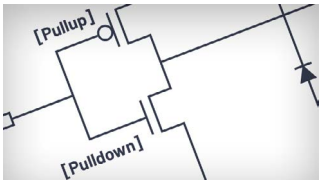

Bernhard Siegel,
Editor in Chief

Bernhard became editor in chief of *Analog Dialogue* in March 2017. He has been with Analog Devices for over 30 years, starting at the ADI Munich office in Germany. In his current role as the editor in chief, he is responsible for the worldwide technical article program within Analog Devices.

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.

AnalogDialogue

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 over 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* home page, analogdialogue.com.

Optical Liquid Analysis Prototyping Platform Lights the Pathway to Ubiquitous Sensing

Sydney Wells, Applications Engineer, and
Scott Hunt, Systems Applications Engineer

Abstract

Monitoring the environment in real time is critical to improving global sustainability. Having the ability to quickly analyze a sample and identify a problem is key to a fast resolution with minimal impact to the ecosystem. This drive toward ubiquitous, real-time sensing has shifted requirements for liquid sensors to require smaller size, higher robustness, and lower power while still delivering high quality results. As the industry advances, intelligent platforms for on-the-go sensing are needed. These platforms need to be highly versatile, capable of satisfying unique requirements for a wide range of applications from environmental waters to process control. This article will introduce a portable, real-time sensing solution and prototyping platform for rapid liquid sensing.

A Common Liquid Analysis Technique

Various methods exist to test liquids, the purpose being to measure the concentration of an unknown parameter in a sample, such as pH, fluorescence, or turbidity. A popular method is to evaluate liquids optically, as it is noninvasive and provides stable and accurate results. Precision optical liquid measurements require mixed domain knowledge in electronics, optics, and chemistry. In simple terms, the analysis begins with a sample that is exposed to light from a source such as an LED. After interacting with the sample, the resulting light is processed by a photodiode. That measured response is plotted against the measured responses of a set of standard samples of known concentrations. This is known as a calibration curve. Using the calibration curve, the unknown value can be determined. This describes the general laboratory method for analytical measurements, but to meet the needs of ubiquitous sensing it must be scaled to different analytes and to measurement techniques, as well as fit into a small form factor, all of which increases the complexity of design and evaluation.

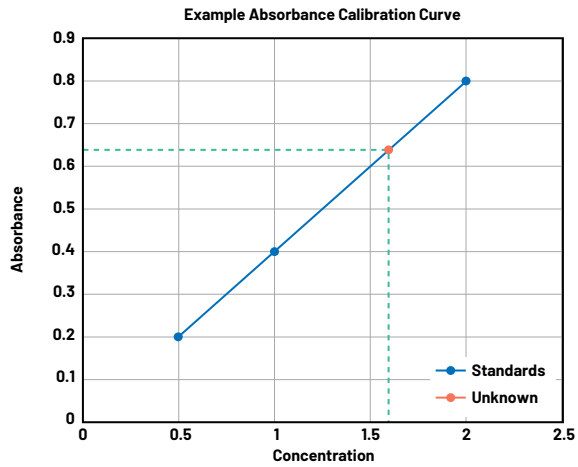

Figure 1. An example of a calibration curve for absorbance.

Modular ADI Solution for Rapid Liquid Measurement

The ADPD4101 from Analog Devices is an optical analog front end (AFE) capable of driving LEDs and synchronously receiving and processing signals from photodiodes to make high precision optical measurements. The ADPD4101 is highly configurable, featuring a high optical signal-to-noise ratio of up to 100 dB and high ambient light rejection provided by on-chip synchronous detection methods, allowing it to be used without an optically dark enclosure in many cases.

The CN0503 reference design was created to enable rapid prototyping of liquid analysis measurements with the ADPD4101. The CN0503 features the ADPD4101 as its core product, but adds up to four modular optical paths as well as measurement firmware and application software targeted for liquid analysis. The CN0503 interfaces directly with the ADICUP3029 board, which manages the measurement

routine and data flow. The ADICUP3029 board can be connected directly to a laptop to view results in the evaluation GUI. The CN0503 can measure fluorescence, turbidity, absorbance, and colorimetry. A sample is prepared in a cuvette and placed in the 3D-printed cuvette holder, which houses the optics, including a lens and beam splitter. The cuvette holder slots into the appropriate optical path for plug-and-play measurement. In addition, the LED and photodiode cards can be switched out for even more customization.

To demonstrate creating calibration curves and measuring unknowns with the CN0503, measurements of pH, turbidity, and fluorescence will be shown. The evaluation GUI was used to take measurements to create calibration curves. The noise value and limit of detection (LOD) were calculated to determine the lowest concentration feasible to be detected by the CN0503 for each example.

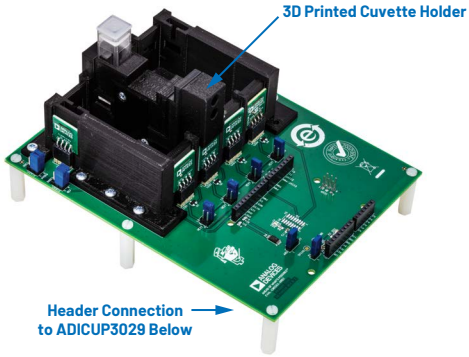

Figure 2. The CN0503 evaluation board.

Measuring pH with Absorbance

Absorbance Background

Absorbance involves determining the concentration of a known solute in a solution based on how much light is absorbed at a particular wavelength. The concentration is proportional to the absorbance, per the Beer-Lambert Law. Many colorless analytes can be measured by adding a color-changing reagent. This example is measuring pH, one of the most common parameters measured across many industries from water quality to wastewater treatment. Absorbance measurements are used for many other parameters, including dissolved oxygen/biological oxygen demand, nitrates, ammonia, and chlorine.

Optics

The optical path configuration for absorbance measurements is shown in Figure 3. With the CN0503, absorbance measurements can be made in any optical path (1 through 4). The incident beam is directed at the beam splitter where a reference photodiode samples the intensity of the beam. The remaining power is directed through the sample. Taking the ratio of sample to reference light removes variation and noise of the LED source, and synchronous pulse and receive windows provide ambient light rejection.

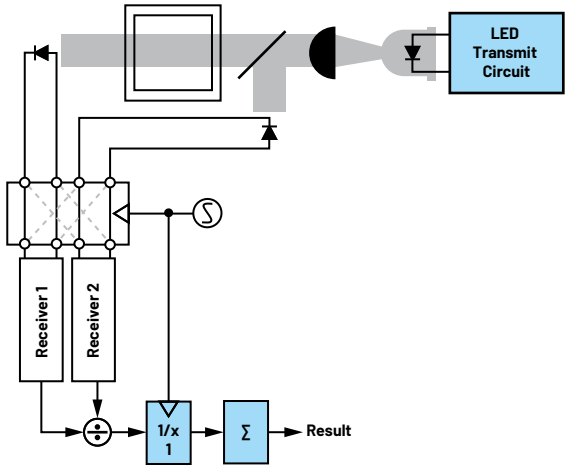

Figure 3. An optical light path for measuring absorbance.

Setup Equipment

- CN0503 evaluation board
- EVAL-ADICUP3029 evaluation board
- API pH Test & Adjuster Kit
- pH standards

Figure 4. The CN0503 taking a pH measurement.

In this experiment, a color indicator (bromothymol blue) was added to prepared solutions with different pH values. The solutions were transferred into cuvettes and tested at two different wavelengths, 430 nm and 615 nm, where the indicator shows absorption changes vs. pH. The CN0503 makes this easy; the two different wavelength LED cards can be inserted into optical path 2 and optical path 3. Then the cuvette holder is simply moved into different paths for the different measurements.

Results

Using the CN0503 evaluation GUI, measurement results from both optical paths were easily exported into Excel. The resulting calibration curves for the two different wavelengths are shown in Figure 5 and Figure 6.

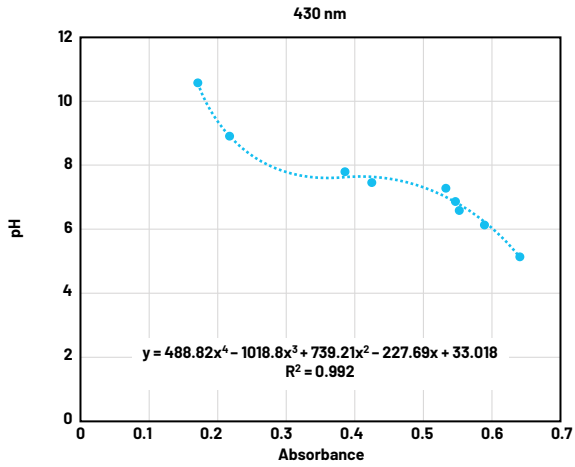

Figure 5. Absorbance calibration curve of pH at 430 nm.

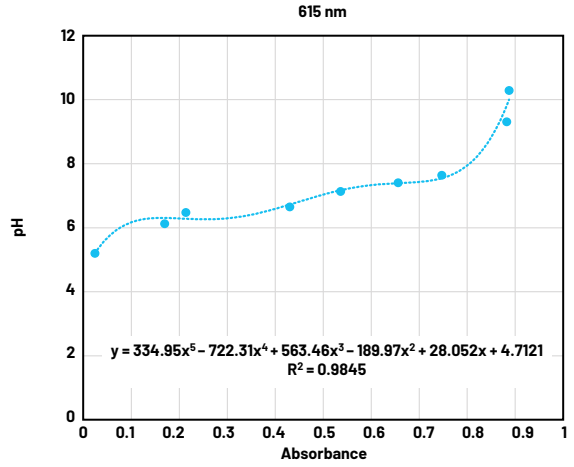

Figure 6. Absorbance calibration curve of pH at 615 nm.

In each case, the pH vs. absorbance was plotted to create the calibration curve. The **Add Trendline** option is then used to get an equation for the curve. The equations are then used to determine concentrations of unknown samples. The sensor output is the x-variable, and the resulting y-value is the pH. This can be done by hand; however, the CN0503 can also be used for this. The firmware implements two fifth-order polynomials, INS1 and INS2. Once the polynomials are stored, the INS1 or INS2 mode can be selected so that measurement results are reported directly in the desired unit—in this case, pH. This makes it simple to quickly get a result for an unknown sample.

To get the noise value, two different data points were chosen for each wavelength: one lower pH value and one higher pH value. Two points were used because the curve fit was not linear in this case. The standard deviation of a set of repeated measurements for each point, reported as the noise value in Table 1, describes the precision of the measurement, excluding variations in sample preparation.

Table 1. pH Measurement Noise Values

	6.1 pH Sample		7.5 pH Sample	
	430 nm	615 nm	430 nm	615 nm
RMS Noise Value (pH)	0.002098	0.000183	8.18994×10^{-7}	0.000165

LOD is typically determined by measuring noise at low concentration and multiplying by 3 to give a confidence interval of 99.7%. Because pH is a logarithmic scale, pH 7 was chosen as the number to check the LOD, as shown in Table 2.

Table 2. pH Measurement Limit of Detection

	7 pH Sample	
	430 nm	615 nm
Limit of Detection (pH)	0.001099	0.001456

Measuring Turbidity

Turbidity Background

Turbidity measurements of a liquid sample use the light scattering property of particles suspended in the liquid. Ultimately, it is a measure of the relative clarity of a liquid. The amount of light scattered and the scattering angle differ based on particle size, concentration, and the wavelength of incident light. Measuring turbidity is conducted in many industries, including water quality and life sciences. Besides general turbidity, the CN0503 could be used to determine algae growth by measuring optical density.

Optics

Figure 7 shows the optical path for turbidity measurement with a 90° or 180° detector. With the CN0503, turbidity measurements can be made in optical path 1 or 4 only, as a 90° detector is needed. There are various measurement configurations and standards for turbidity. This example demonstrates a modified version of EPA Method 180.1, calibrated and reported in nephelometric turbidity units (NTU).

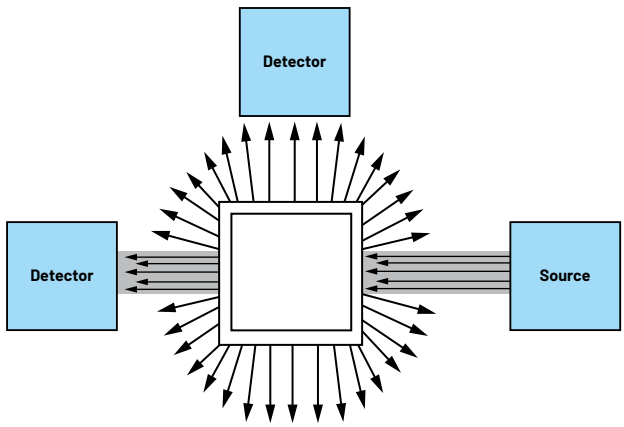

Figure 7. Optical path for turbidity measurement.

Setup Equipment

- CN0503 evaluation board
- EVAL-ADICUP3029 evaluation board
- Hanna Instruments® turbidity standard calibration set

For this experiment, optical path 4 was used with the 530 nm LED board inserted for testing.

Figure 8. Turbidity calibration standards.

Results

Using the CN0503 evaluation GUI, measurement results were exported into Excel. The resulting calibration curve is shown in Figure 9.

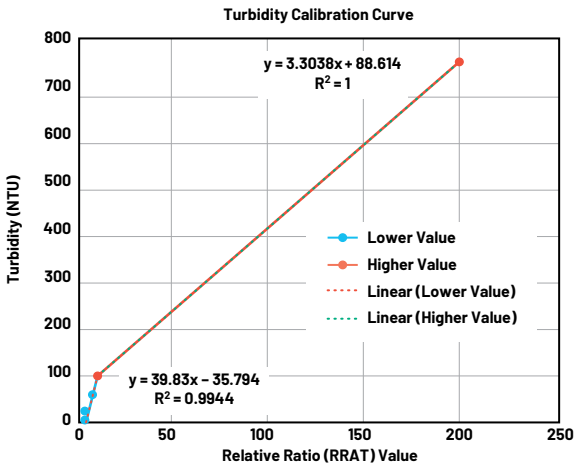

Figure 9. Calibration curve for turbidity.

The response curve was split into two sections because the 90° scattering measurement is less responsive to high turbidities. One section represents lower turbidity (0 NTU to 100 NTU), and the other higher turbidity (100 NTU to 750 NTU). Two linear fits were then made to each section. Even though there are now two equation values, the CN0503 can still be used to quickly show resulting NTU values. This is because each optical path can store its own equation values in INS1 and INS2. One important note is that INS1 and INS2 are dependent. The result of the first equation, INS1, is the input variable for the second equation, INS2. Once the equation values are stored, INS1 can be used to measure lower-turbidity samples, and INS2 for higher.

To get the noise value, we choose a data point to take the standard deviation of repeated measurements. The standard deviation is the noise value. One data point is chosen near the bottom of the range because the equation fit is linear.

Table 3. Turbidity Measurement Noise Value

12 NTU
RMS Noise Value (NTU)
0.282474

To determine the LOD, the noise value is measured for a blank or low concentration sample, then multiplied by 3 to represent a 99.7% confidence interval.

Table 4. Turbidity Measurement Limit of Detection

Blank Sample
Limit of Detection (NTU)
0.69204

Measuring Fluorescence with Spinach Solutions

Fluorescence Background

When light is shined into a sample containing fluorescent molecules, the electrons move into a higher energy state and then lose some of that energy before emitting light at a longer wavelength. The fluorescence emission is chemically specific and can be used to identify the presence and amount of specific molecules in a medium. In this example, fluorescing chlorophyll was demonstrated by using spinach leaves. Among many applications, fluorescence measurements are common in biological assays, dissolved oxygen, chemical oxygen demand, and detecting if pasteurization has been successful in milk.

Optics

The optical path configuration for fluorescence measurements is shown in Figure 10. With the CN0503, absorbance measurements can be made in optical path 1 or 4 only, which is because of the 90° detector. Usually, a fluorescence detector is positioned at 90° from the incident light and a monochromatic or long-pass filter is used to increase isolation between the excitation and emitted light. Fluorescence is a very sensitive low level measurement and is subject to interference, so the reference detector and the synchronous detection methods are used to reduce error sources.

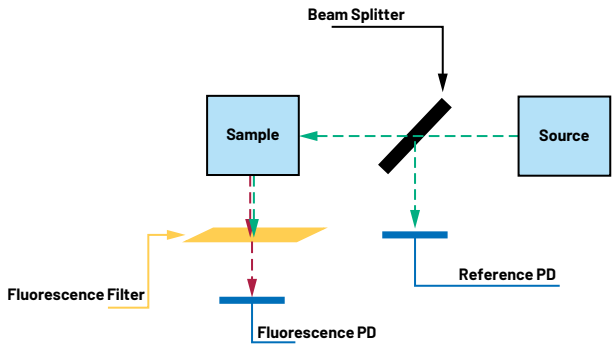

Figure 10. Optical path for fluorescence measurement.

Setup Equipment

- CN0503 evaluation board
- EVAL-ADICUP3029 evaluation board
- Spinach solution

For this experiment, a spinach solution was created by blending spinach leaves with water. This was then filtered and kept as the stock solution. The stock solution was then diluted to result in samples of different percent spinach solution. These were used as standards to create a curve of percent spinach solution through fluorescence. Optical path 1 was used along with the 365 nm LED card and the long-pass filter inserted.

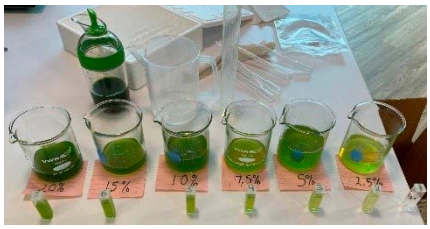

Figure 11. Chlorophyll samples made with spinach.

Results

The calibration curve for percent spinach solution is shown in Figure 12.

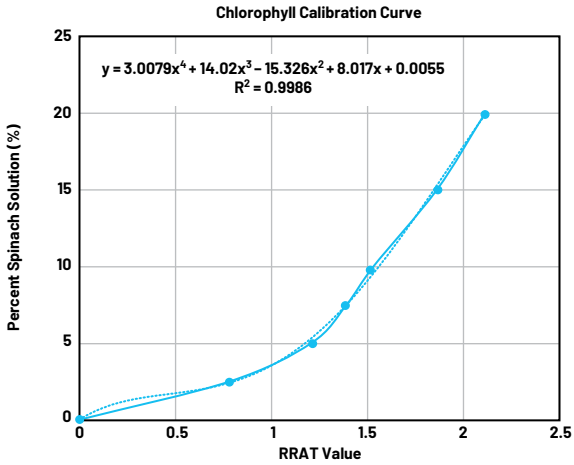

Figure 12. Calibration curve for percent spinach solution.

The trendline equation for this calibration curve can be stored so that the results are directly reported as a percentage by the CN0503.

To get the noise value, two different data points were picked: one near the bottom of the range and one near the top because the curve fit was not linear. The noise is given by the standard deviation of a set of repeated measurements for each point, as shown in Table 5.

Table 5. Fluorescence Measurement Noise Value

	7.5% Spinach Sample	20% Spinach Sample
RMS Noise Value (% Spinach)	0.0616	0.1158

To determine the LOD, the noise value is measured for a blank or low concentration sample, then multiplied by 3 to represent a 99.7% confidence interval.

Table 6. Fluorescence Measurement Limit of Detection

Blank Sample
Limit of Detection (% Spinach)
0.1621

Conclusion

Prototyping complex optical liquid analysis measurements is a challenge that requires careful consideration of how chemistry, optics, and electronics interact to produce a precise result. Integrated AFE products like the ADPD4101 pave the way for higher performance optical liquid sensing in smaller spaces. The CN0503 builds on the ADPD4101, including optical design, firmware, and software for an easy to use and highly customizable rapid prototyping platform that is capable of taking accurate optical measurements of liquid parameters, including absorbance, colorimetry, turbidity, and fluorescence.

References

"HI98703-11 Turbidity Calibration Standards." Hanna Instruments, Inc.

Optical Platform: Turbidity Measurement Demo. Analog Devices, Inc.

About the Author

Sydney Wells is currently in the field development rotational program at ADI, training for the role of field applications engineer. She earned her bachelor's degree in electrical engineering from the University of Connecticut in 2020. Sydney has previously worked on power converter evaluations, manufacturing automation, and supplier development engineering. Her current interests are in power and instrumentation. She can be reached at sydney.wells@analog.com.

About the Author

Scott Hunt is a system applications engineer specializing in scientific instruments in the Instrumentation Market Group of Analog Devices in Wilmington, MA. Scott joined Analog Devices in 2011 as a product applications engineer for high performance integrated precision amplifiers, such as instrumentation amplifiers, before transferring into the instrumentation group in 2016. He has a bachelor's degree in electrical and computer systems engineering from Rensselaer Polytechnic Institute. Scott was honored as a recipient of the 2015 Outstanding Technical Writing award and the 2015 Outstanding Program Support award from Analog Devices. He can be reached at scott.hunt@analog.com.

How to Accelerate Peripheral Monitoring in Low Power Wearables with DMA

Brandon Hurst, Hardware and Embedded Firmware Engineer

Abstract

This article explains the use cases, advantages, and disadvantages of utilizing direct memory access (DMA) in embedded systems programming. The article describes how DMA interacts with peripheral and memory modules for more efficient operation of CPUs. The article will also introduce the reader to different DMA bus access architectures, and the advantages of each.

One task that is common for embedded systems is managing external input. Managing input can put a lot of unnecessary computational strain on the processor, causing longer periods in active power modes and slow response times. For optimizing power, preserving quick responses to events, and managing large continuous data transfers, a microcontroller with direct memory access (DMA) may offer the best solution.

Direct Memory Access (DMA)

In system applications involving peripherals, there are many points at which a microprocessor can become bottlenecked. For instance, when managing an ADC that is constantly sending data, a processor can be interrupted so often that it struggles to accomplish other tasks. DMA is a method of moving data and minimizing processor involvement in large or fast data transactions. You can think of the DMA controller as a coprocessor whose sole purpose is to interact with memory and peripherals. This allows the main processor to successfully manage a greedy peripheral, focus on another task, or even go to sleep and conserve power while data transactions happen in the background. For example, on Arm® architectures, a DMA module can operate during LP2 (sleep) or LP3 (run) modes. This can give a distinct advantage in applications that require extended battery life, such as wearable sensor hubs and smart watches.

Advantages and Drawbacks

DMA is useful in many digital systems, and sometimes it is even required to manage large amounts of bus traffic. It has been used in network cards, graphics cards, and even some of the original IBM PCs. That being said, incorporating DMA into a design does have some trade-offs.

Table 1. Advantages of Using DMA

Advantages of Using DMA	
CPU Time	DMA minimizes the need for processor execution and interrupts, decreasing the required CPU time for data transactions.
Power Consumption	Using DMA can yield opportunities to minimize power consumption if it allows the processor to sleep during DMA transfers.
Parallel Operation	Depending on the architectural details of the system bus, the processor may be able to execute other operations while peripheral transactions are taking place.

Table 2. Disadvantages of Using DMA

Disadvantages of Using DMA	
Cost	Incorporating a system with DMA requires a DMA controller, and this can make a system more expensive.
Complexity	While DMA can reduce the frequency of interrupts, it can increase the size and complexity of application firmware.
Platform Dependence	DMA controllers have differing internal architectures between and within manufacturers and can have different behavior depending on their native bus access schemes.
Cache Incoherency	DMA transactions can cause logical errors to occur by writing to a cached layer of the memory hierarchy. This can be solved by using cache-coherent system architectures or by invalidating cache storage upon DMA completion.

Bus Access and CPU Cycles

While DMA controllers can be incredibly effective at conserving power or speeding up embedded systems, their implementation is not heavily standardized. There are multiple schemes for making sure that internal bus access is not granted simultaneously with the CPU. The goal of the bus access scheme is primarily to avoid concurrent access to the same memory locations, which can lead to cache incoherency and logical errors. A single DMA controller will usually be configured to employ one of these schemes, since different hardware or firmware control may be required to use each of them. The bus access schemes used by most DMA controllers are burst, cycle-stealing, and transparent DMA.

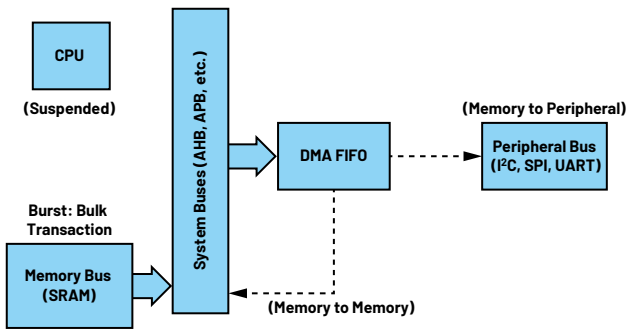

Figure 1. An architectural diagram of burst DMA during DMA operations.

Burst DMA occurs through infrequent large bursts, where the DMA controller sends as much data to the destination buffer as the buffer can hold. The DMA controller blocks CPU operation for a very short period to move a large chunk of memory, and then relinquishes the bus back to the main CPU, repeating until the transfer is complete. Burst DMA is generally considered the fastest type.

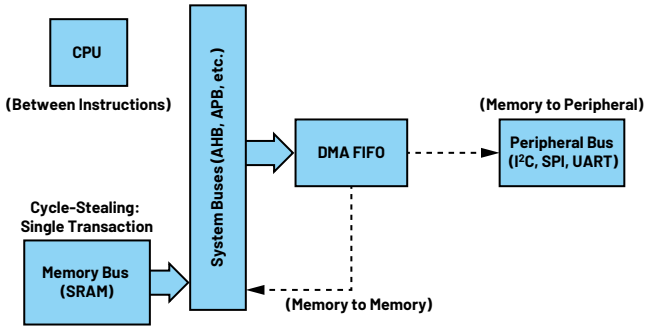

Figure 2. Cycle-stealing DMA during DMA operations occurs between two CPU cycles.

Conversely, single byte transfer or cycle-stealing DMA takes a cue from the CPU and only carries out operations between CPU instructions. It inserts a single operation between two CPU cycles, and thus is in effect "stealing" CPU time. Due to the limitation of executing one operation at a time, it is generally slower than burst DMA.

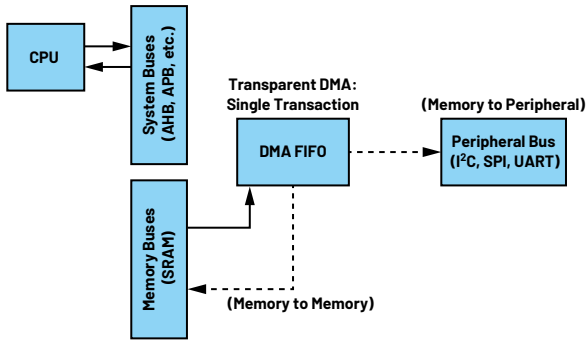

Figure 3. Transparent DMA during DMA operations occurs while the processor works on tasks that do not access the data or address buses.

Transparent DMA can only execute a single operation at a time, but it must also wait for the processor to execute instructions in which it yields access to the desired data or address buses. Extra logic is required to verify this access restriction, and this type of DMA is generally the slowest. Transparent DMA may be advantageous in applications where one has extra processing to do that does not require access to the memory buses. The advantage in this case would be the elimination of throttling the CPU, since the processor does not have to stop operating completely.

Table 3. Summary of DMA Types and Their Pros/Cons

Type of DMA	Pros	Cons
Burst DMA	Fastest type of DMA	Relatively long periods of CPU idle time
Cycle-Stealing DMA	CPU is not idle for long contiguous periods	Slower than burst DMA
Transparent DMA	No throttling of CPU use needed	Slowest form of DMA

Example of a Burst DMA Architecture

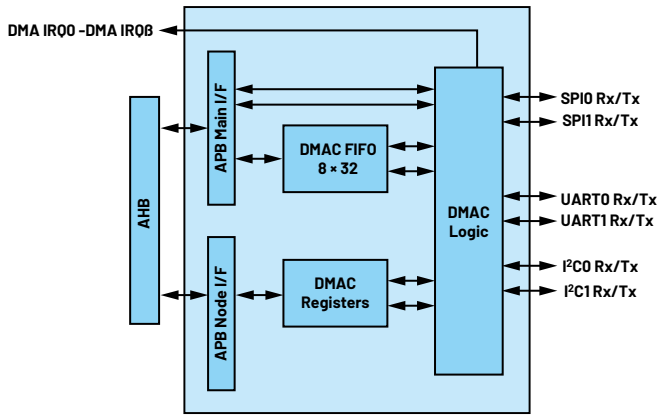

Figure 4. An architectural diagram of the DMA controller on the MAX32660.

An example of a burst DMA controller can be found on the MAX32660 (see Figure 4). The upper path corresponds to data flow, and the lower path represents control/status flow between the advanced high performance bus (AHB) and the DMA logic. The DMA controller can behave as a buffer interface between the AHB and memory or peripheral modules, depending on how it is configured. DMA logic sits between the DMA buffer and each peripheral to independently manage each unique peripheral bus during transactions. A DMA operation can move up to 32 bytes at a time, provided the source/destination buffers can contain this much data. The buffer can hold up to 16 MB and is configurable to transmit or receive I²C, SPI, I²S, and UART in addition to internal memory transfers. Programming DMA control may vary slightly between protocols, but the peripheral transactions are managed exclusively by the DMA controller. An arbiter module controls the bus access restrictions between the four DMA channels and the CPU, granting requests according to a priority system.

Modern DMA Options

In summary, DMA is a critical feature for modern embedded systems that manage an abundance of sensors and require high throughput, efficiency, and low power operation. It behaves like a coprocessor dedicated exclusively to memory and peripheral bus transactions.

Using DMA is imperative for many applications to minimize power consumption and lighten processor loads. For example, health and wearable devices handle large amounts of data throughput, but they also must conserve as much battery charge as possible, all while handling sensitive data. Analog Devices offers fast burst DMA architectures on microcontrollers well-equipped for low power wearable designs, such as the MAX32660 and MAX32670. In addition, DARWIN Arm microcontrollers such as the MAX32666 are built for wearable and IoT applications with integrated Bluetooth® 5. These devices have two 8-channel burst DMA controllers with integrated support for event-based transactions. They even feature best-in-class security hardware with a secure bootloader and trust protection unit (TPU) for accelerating ECDSA, SHA-2, and AES encryption. From the early IBM PCs to network cards, and now to secure, low power wearable and IoT devices, DMA is an essential feature of modern digital systems.

About the Author

Brandon Hurst is a hardware and embedded firmware engineer working with the Training and Technical Services Group for Maxim Integrated, now part of Analog Devices. He graduated with a bachelor's in electrical engineering from Cal Poly, San Luis Obispo and joined Maxim in January 2021. Previously, Brandon did an internship with both Maxim's TTS team and the Product Safety Engineering team at Apple, Inc. He can be reached at brandon.hurst@analog.com.

Current Sensing with PMBus Digital Power System Managers—Part 2

Michael Peters, Senior Applications Engineer

Abstract

Part 2 of this article series describes how to measure current on high voltage or negative supply rails and how to set configuration registers for the I_{MON} sensing method. This article covers the accuracy considerations of measuring current and provides instructions on programming devices using LTpowerPlay®. In Part 1, we covered the general concepts of current sensing including the various methods and circuit topologies.

Beyond the Limits

The LTC297x devices have limits to the voltage applied to the sense pins, both V_{SENSE} and I_{SENSE}. The limit is 6 V. For much of the following, we will be discussing most of the LTC297x family, the exception to this being the LTC2971, which has limits to ±60 V. For supply rails that are greater than 6 V or for a negative supply, an indirect method of sensing voltage across an inductor or sense resistor must be devised.

Resistor Divider

For supply voltages that are higher than the maximum voltage rating of I_{SENSE} pins, one may be tempted to use two voltage dividers. This seems like a reasonable idea until you calculate the error of the divided down "signal." A voltage divider is placed on each side of the sense element. The "output" is taken from each of the dividers and fed into the LTC297x sense pins. If the top-to-bottom resistor ratios match each other, the goal of accurately dividing down the HV signals is achieved. The rail voltage is divided enough to keep the LTC297x inputs within their limits, and the divided output develops a proportional voltage that may be measured by the LTC297x. However, the required resistor tolerance makes this approach impractical. Moreover, the more the voltage is divided, the greater the error will be. For example, if just one of the resistors has an error of just 0.1%, the result is a fixed offset error. The gain error contributes very little, and the offset error dominates.

For example, assume you need to measure the output current of a 12 V supply. The supply is capable of 2 A, and a 10 mΩ shunt resistor (R_{SNS}) is placed in the output path. The shunt will produce a 20 mV signal at full load. A divide-by-3 circuit is an appropriate choice, and the top and bottom resistors are chosen to be 2 kΩ and 1 kΩ, respectively. This places the common-mode voltage of the I_{SENSE} pins at 4 V. The use of relatively low values is meant to keep the source impedance low, as seen by the LTC297x device to reduce leakage current induced errors from the Thevenin equivalent resistance of the divider.

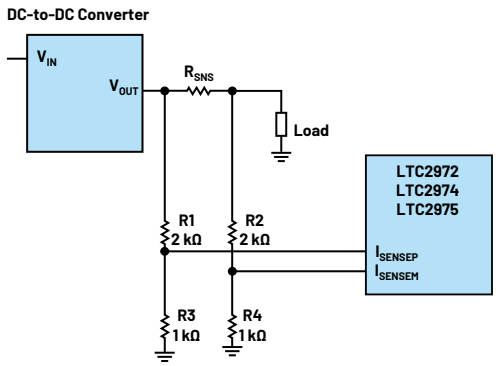

Figure 1. A resistor divider for current sense produces a large error.

Assume a no-load condition and all resistors are perfect. Each divider midpoint will be 4.00 V and the delta V will be zero. Therefore, the LTC297x will have a READ_IOUT value of 0.000 A. However, if the resistance of one of the 2 kΩ components is 0.1% high (2002 Ω), the delta V will be 2.665 mV. But remember that full scale is 20 mV/3 or 6.667 mV, as seen by the I_{SENSE} pins. The 2.665 mV reading translates to an output current of 0.4 A. This is 40% of the expected full-scale reading! As mentioned earlier, the error introduced is an offset error, not a gain error. It is a large error, nonetheless. This approach is too sensitive to resistor tolerance, and we must look for another solution.

High-Side Sense Amplifier

Since the LTC2972/LTC2974/LTC2975 have a 6 V limit to the I_{SENSE} pins, the solution to this problem is to use a high-side current sense amplifier (CSA) for level translation. The LT6100/LTC6101 are popular for fixed/user-selectable gain. The accuracy is much better than with divider resistors.

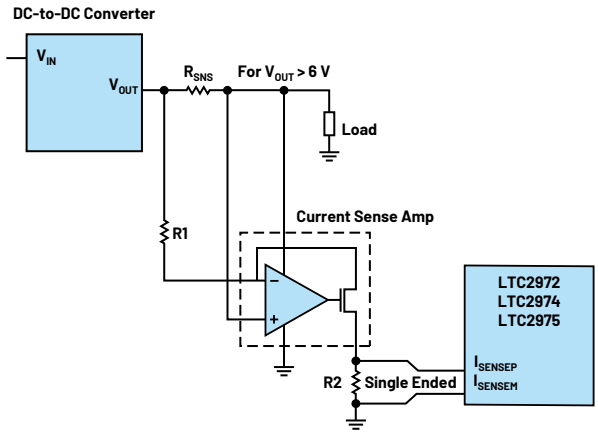

Figure 2. A current sense amplifier used as a level translator.

Here are the relevant equations and conditions:

V_{OUT} of CSA = $I_{LOAD} \times R_{SNS} \times (R2 / R1)$

Set IOUT_CAL_GAIN = $R_{SNS} \times (R2 / R1)$

Keep $V_{ISENSEP} < \pm 170$ mV

LTC2971 for Current Sensing on High Voltage Rails

For high voltage rails, use an LTC2971 (2-channel DPSM) to directly sense current and voltages up to 60 V. The LTC2971 is available in four different ordering options. The LTC2971-1 supports 60 V sensing on one channel and –60 V on the other channel. The LTC2971-2 supports –60 V on both channels, and the LTC2971-3 option supports 60 V and 1.8 V. The LTC2971 supports 60 V sensing on both channels. The direct connection to the IOUT_SNS pins avoids external CSAs, which would increase cost, add board space, and introduce error. The LTC2971 current measurement accuracy is 0.6% of the READ_IOUT reading.

Table 1. LTC2971 Ordering Options

Option	CH0	CH1
LTC2971	0 V to 60 V	0 V to 60 V
LTC2971-1	0 V to 60 V	0 V to –60 V
LTC2971-2	0 V to –60 V	0 V to –60 V
LTC2971-3	0 V to 60 V	0 V to 1.8 V

Low-Side Current Sense

Low-side current sensing may be an option in some cases. A sense resistor can be placed on the low side of the load and the I_{SENSE} pins tied across the resistor. This allows the I_{SENSE} pins to have a common-mode voltage near GND. For supply voltages greater than 6 V, it may suit your application well. This is a good solution for measuring current on virtually any power supply rail, including high voltage rails. Selecting the R_{SENSE} value is a compromise between obtaining a large enough signal for good accuracy vs. a low enough resistance that does not produce a significant IR drop, causing the output voltage to drop as seen by the load—that is, poor load regulation. Figure 3 shows the feedback resistors and the Kelvin sense connections for V_{SENSE}. Kelvin sense is a term used to describe the connections made to the sense element, which do not include voltage drops.

Care should be taken to establish a return current path for the sense resistor. Many high density boards are designed with many layers of ground pour, thus allowing return current to flow through multiple paths. The use of a shunt resistor gives you the ability to force the return current through this element, thereby allowing Kelvin sense connections across the element and tied back to the I_{SENSE} pins of the PSM device.

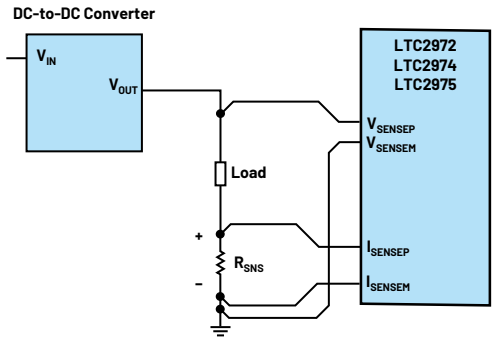

Figure 3. Low-side sensing solves high voltage current sensing but has drawbacks.

Current Sense on a Negative Rail

There are a few different ways to monitor the output current of a negative supply. The simplest solution is to use a low-side CSA, such as the LTC6105. Figure 4 shows the inputs tied across a shunt, and the CSA is powered from VDD33 of the PSM and the low side of the negative rail itself. The output is a single-ended signal that can be wired either to the I_{SENSE} or V_{SENSE} pins of the PSM.

If the CSA is wired to the I_{SENSE} pins, then set IOUT_CAL_GAIN to $R_{SNS} \times GAIN_{CSA}$. For example, if the shunt resistance is 10 mΩ and the CSA gain is 10, then set IOUT_CAL_GAIN to 100. The IOUT_CAL_GAIN units are milliohms.

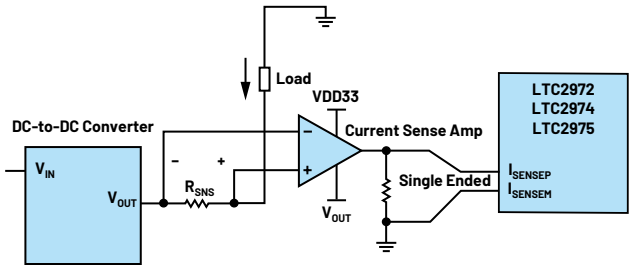

Figure 4. Sensing current with a CSA (LTC6105).

Using the LTC2971-1 or LTC2971-2 is the simplest solution for monitoring output current on a negative supply. These are 2-channel devices, and both channels of the LTC2971-2 can natively sense current on rails to –60 V. The LTC2971-1 is capable of sensing current on negative rails on Channel 1 only.

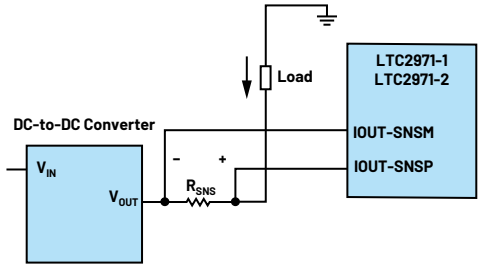

Figure 5. Sensing current on a negative rail without external components.

Note: The LTC2971's READ_VOUT value is L16 format and is unsigned. The value displayed in the GUI for the negative rail voltages is inverted.

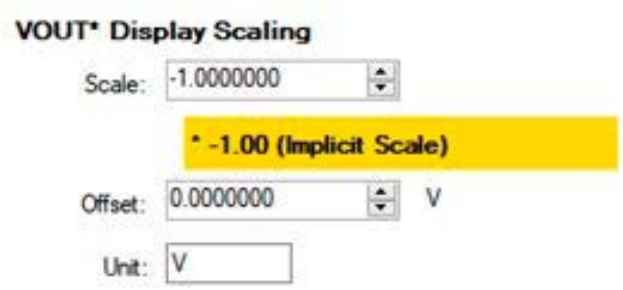

Figure 6. The LTpowerPlay setup tab for LTC2971-1 Channel 1 and both LTC2971-2 channels.

I_{MON} Examples

Current-driven I_{MON} pins allow the user to choose a resistor value that sets the current sensing gain and maximum voltage. The PSM device measures the voltage difference between the I_{SENSEP} and I_{SENSEM} pins, and sensing gain needs to be set using MFR_IOUT_CAL_GAIN, which is similar to shunt sensing.

The LT3081 LDO regulator has an I_{MON} pin that can used as an example. The LT3081 I_{MON} current is the load current divided by 5000. Suppose a 2 kΩ resistor is used. The I_{MON} pin voltage per amp of load current is:

$V_{IMON} = (I_{LOAD} / 5000) \times 2000\ \Omega = 0.4\ V/A$

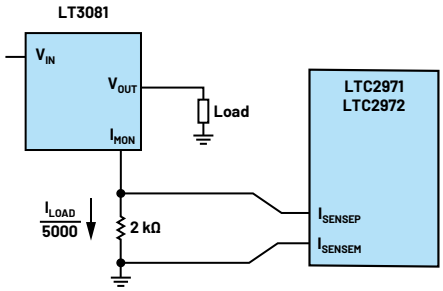

Figure 7. Using the LT3081 I_{MON} pin.

If the load current is 2 A, the I_{MON} voltage will be 0.8 V. From this equation, we can see that the I_{MON} voltage can be made more sensitive to load current by simply increasing the I_{MON} resistor value. If we do so, the maximum voltage (full load) may well be >1 V. The PSM device's I_{SENSE} pins need to accommodate this large excursion. For the LTC2974/LTC2975, this violates the differential voltage, which is limited to ±170 mV. Fortunately, the LTC2971 and LTC2972 have a configuration bit that, when imon_sense is set, places the current sense circuits into a mode that allows sensing single-ended voltages up to 6 V.

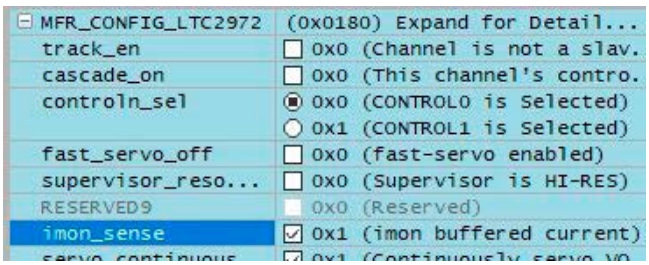

Figure 8. MFR_CONFIG imon_sense bit.

The configuration commands must be set according to the hardware we've selected. In this example, IOUT_CAL_GAIN should be set to 400 (0.4 V/A). The units are milliohms. The other current-related commands may have default values if there is no temperature coefficient or thermal time constant that may affect the READ_IOUT value. The default values for MFR_IOUT_CAL_GAIN_TC, MFR_IOUT_CAL_GAIN_TAU_INV, and MFR_IOUT_CAL_GAIN_THETA are set to zero.

The I_{MON} pin of the LT7101 buck regulator is an example of a pin with a voltage-driven output. The output also has an offset voltage. That is, under no load conditions, the I_{MON} pin sits at 0.4 V. At first this may seem problematic since the differential voltage limit is ±170 mV. However, the LTC2972/LTC2971 PSM devices can sense this type of I_{MON} pin and allow a much larger differential signal on the I_{SENSE} pins. Let's work through a real example.

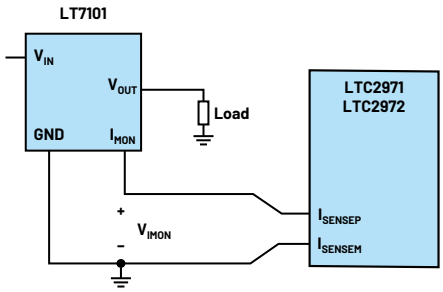

Figure 9. Using the LT7101 I_{MON} pin.

An LTC2971/LTC2972 can be connected to a LT7101 by connecting the LTC297x I_{SENSEM} pin to ground and connecting the I_{SENSEP} pin to the I_{MON} pin. The command values can be calculated as follows:

Start with the READ_IOUT equation,

$$READ\_IOUT = \frac{V_{IOUT\_SNSP} - V_{IOUT\_SNSM}}{IOUT\_CAL\_GAIN \times T_{CORRECTION}} + IOUT\_CAL\_OFFSET$$

Rewrite the equation solving for IOUT_CAL_GAIN:

$$IOUT\_CAL\_GAIN = \frac{V_{IOUT\_SNSP} - V_{IOUT\_SNSM}}{READ\_IOUT - IOUT\_CAL\_OFFSET}$$

assuming $T_{CORRECTION} = 1$.

The LT7101 data sheet provides I_{MON} voltage levels for 1 A and 0.25 A load currents, 1.21 V and 0.603 V respectively. Therefore, the IOUT_CAL_GAIN value is:

$$\frac{1.21\ V - 0.603\ V}{1\ A - 0.25\ A} = 810\ m\Omega$$

The IOUT_CAL_OFFSET is:

$$\left(\frac{0.603\ V}{0.81\ \Omega}\right) - 0.25\ A = -0.494\ A$$

The IOUT_CAL_OFFSET is a negative value because we need to reduce the READ_IOUT value. You may find that the calculated register values need to be changed to better correlate the measured load current with the READ_IOUT readings. This involves adding a calibration step. Force a known load current, compare the READ_IOUT value against the expected value, and write the adjusted value to IOUT_CAL_GAIN and/or IOUT_CAL_OFFSET. In general, I_{MON} accuracy on many regulators is not as precise as sense resistors for current measurements, but calibrating the current measurement will greatly improve the accuracy.

Accuracy

A current measurement is only as accurate as the sum of its parts. In most systems, accuracy is important in the mid to high end of the load current range. Some demand good accuracy under light load conditions, which means the signal in the sensing chain is very small. We can break down the accuracy parts into four

categories: the sense element, board layout, the amplifier, and the sense measuring circuit.

Before covering accuracy in more detail, the term TUE needs to be defined. The total unadjusted error or TUE is a specification listed in every LTC297x data sheet. There are TUE specs for voltage and current measurements. The TUE is the combined error contributed by the PSM device's internal reference, gain, and offset errors in the buffers and amplifiers in the path from the V_{SENSE} or I_{SENSE} pins to the digital portion of the chip. The TUE is the worst-case error as a percentage of the READ_IOUT or READ_VOUT reading across all process variation and temperature. This removes the burden of calculating the individual contributors from the chip, such as V_{REF} error and ADC errors. External components—CSAs and associated resistors, shunt resistors, inductor DCR, I_{MON} current—contribute their own error and must be considered in the overall error budget.

As stated earlier, the best accuracy comes from a resistive sense element that is placed in the output path. The R_{SENSE} tolerance is commonly specified as 1%. These are inexpensive and easy to find. Values typically range from 0.5 mΩ to tens of mΩ. To determine the value, one must consider the current range of interest and the accuracy desired at each end of the range. As current flows through the R_{SENSE}, a small voltage is developed across the element, a delta V. It is this signal that needs to be measured and converted via Ohm's Law into a current. One would like to obtain a large enough signal for good accuracy under light load conditions; however, a large IR drop will occur under heavy loads and could have a negative effect on the performance of the supply. We are assuming that the regulator's feedback is taken from the load itself, that the sense points are wired across the load. This accounts for any voltage drop in the output path, both high side and GND return path. The R_{SENSE} is located inside the regulator's feedback loop. Any PCB copper in the layout that contributes to IR loss is included.

The following is an example that covers accuracy. Suppose a power supply is capable of 10 A max, and we need good accuracy down to 100 mA. At full load, it is recommended to keep the IR drop to <50 mV. If the sense resistor is placed within the feedback loop, you can afford to create a larger sense voltage. The downside to a large signal is the power loss in the sense element. This is the basic trade-off in selecting a resistor value. The R_{SENSE} value is calculated from the voltage sense at full load current—for this example, 50 mV/10 A or 5 mΩ. Suppose we choose a 5 mΩ sense resistor that has a tolerance of 1%.

The accuracy achieved will be 1% (resistor tolerance) + 0.3% (TUE from the data sheet) or 1.3%, for LTC2972/LTC2974/LTC2975 input sense voltages >20 mV, which translates to load currents greater than 4 A. For sense levels <20 mV, the TUE is specified as ±60 µV. For a load current of 100 mA, the signal produced is 0.1 A × 0.005 Ω or 500 µV. The error is much larger under light load conditions at ±12% (60 µV/500 µV), which is dominated by the TUE, and the resistor tolerance has little influence over the accuracy. In absolute terms, it amounts to only ±12 mA of error. TUE accounts for internal reference and ADC errors. Choosing a tighter tolerance sense resistor yields much better accuracy.

Table 2. Example I_{SENSE} Accuracy Calculations

Load Current	Sense Voltage	LTC2972/ LTC2974/ LTC2975 TUE	Resistor Tolerance	READ_IOUT Accuracy	READ_IOUT Accuracy
100 mA	500 µV	60 µV	1.0%	±13%	±13 mA
100 mA	500 µV	60 µV	0.1%	±12.1%	±12 mA
10 A	50 mV	0.3%	1.0%	±1.3%	±130 mA
10 A	50 mV	0.3%	0.1%	±0.4%	±40 mA

The previous discussion applies to most of the LTC297x family for supply rails <6 V, where the LTC2972/LTC2974/LTC2975 I_{SENSE} pins can be tied directly across the sense element, avoiding the need for an external CSA. If the supply rail is >6 V, a CSA is necessary for most of the PSM manager family. The LTC2971 is the exception and allows the direct connection of I_{SENSE} pins up to ±60 V. The LTC2971's TUE is 0.6%, double that of the LTC2972/LTC2974/LTC2975; however, the IOUT_SNS pins may be connected directly to sense resistors on supply voltages up to ±60 V.

When using the LTC2977/LTC2979/LTC2980/LTM2987 to measure output current on a supply voltage >6 V, it is possible to use a CSA single-ended output to drive the V_{SENSE} pins. Any channel may be used, and the adc_hires bit should remain at its default setting of 0. The output current measurement is read from the READ_VOUT register and must be translated from volts to amps. It is important to realize that you have larger dynamic range at the V_{SENSE} pin than the 170 mV limitation on the I_{SENSE} pins of the LTC2974/LTC2975. This allows the CSA gain to be set higher to generate a larger sense voltage since the V_{SENSEP} pin can be driven to 6 V. The parameter to be considered is the CSA's input offset voltage V_{OS}. It is the V_{OS} that is multiplied by the gain that sets the output error of the CSA. If the V_{OS} is 85 µV (LTC6101) and the gain set to 100, the output error could be as much as 8.5 mV. The TUE of the V_{SENSE} pins <1 V is 2.5 mV and >1 V is 0.25%. The CSA gain should be set low in order to minimize the output error, yet large enough to take advantage of the large signal range of the V_{SENSE} pin. The error contributed by the CSA is fixed mV error for a given gain setting. The error in the translated output current value is shown in the last column. Table 8 illustrates an example. The R_{SENSE} is 5 mΩ.

Table 3. LTC2977/LTC2979/LTC2980/LTM2987 Accuracy Calculations for adc_hires = 0 with an External CSA

Load Current	Sense Voltage	CSA Gain	V_{SENSE}	LTC297x TUE	CSA Error	READ_ VOUT Error	Translated Output Current
100 mA	500 µV	20	10 mV	25%	17%	±42%	±42 mA
100 mA	500 µV	100	50 mV	5%	17%	±22%	±22 mA
10 A	50 mV	20	1 V	0.25%	0.17%	±0.42%	±42 mA
10 A	50 mV	100	5 V	0.25%	0.17%	±0.42%	±42 mA

This illustrates that external CSAs provide reasonably good accuracy for large sense voltages but introduce more error under low sense level conditions.

An accurate current measurement is made possible by creating enough sense voltage or signal. It is the delta V from the sense element that needs to be large enough to overcome any noise and errors that are introduced by the chip and other sources such as layout. Make an estimate of the signal-to-noise ratio (SNR) by first deciding how important light load accuracy is to you. One can calculate an optimum value by considering the lowest sense voltage that produces an acceptable accuracy divided by the lowest value of current in the range to be measured.

To achieve the highest accuracy, it is best to create the largest signal and minimize component/layout errors. That is, use a large R_{SENSE} value and use tight tolerance resistors. You also may consider calibrating the current readback value. Apply a known load current and observe the READ_IOUT value. Adjust the IOUT_CAL_GAIN value to minimize the error in the readback value. Store any changed values to the chip's EEPROM by issuing a STORE_USER_ALL command.

Shunt Resistor Sensing Accuracy

The upside of the shunt resistor method is that it is more accurate than the inductor DCR method because a shunt resistor value is typically accurate to 1% or better. The temperature coefficient is quite low compared to inductor DCR. However,

even though you can purchase very tight tolerance resistors, this may be overridden by layout and solder issues.

The downside to the shunt resistor method is that it is lossy due to IR drop. This generates heat and creates a voltage drop in the output path. As mentioned before, the IR drop is largely mitigated by placing the sense resistor inside the feedback loop, which allows the regulator's loop to reduce the voltage drop to a negligible level.

The Rcm resistors need to be the same value due to LTC297x differential input current that will cause a differential error voltage. Mismatched Rcm resistors introduce an error solely due to filter component tolerances. Generally, keep these resistor values less than 1 kΩ.

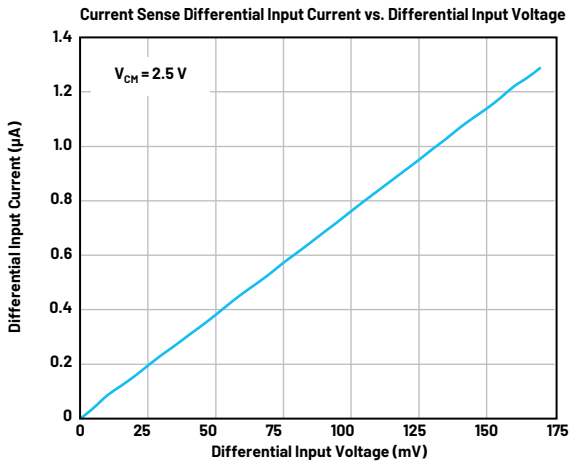

Figure 10. I_{SENSE} pin current.

Layout

Whether you plan to use a discrete sense resistor or inductor DCR to measure current, the layout becomes important under high load conditions. This has the most potential for IR drop in the solder connection and sensing connections can be impactful. It is best to avoid sense connections that are made to a pad that include an IR drop between the sensing points. If you compare the layouts in Figure 11, the example that shows connections to the inside of the pads will have little or no IR drop since those areas of the pad experience little or no current flow. The layout labeled "fair" suffers from IR drops due to the location of the sense point (side of pad), which is partially in the current path.

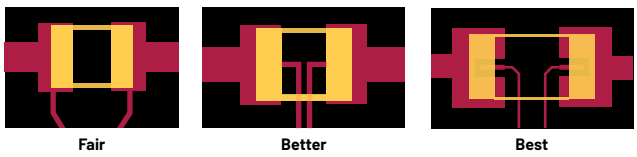

Figure 11. Layout proposals for shunt resistors.

There are 4-terminal sense resistors available on the market. They offer two terminals for the main current path and two terminals for the Kelvin sense connections. For applications requiring good accuracy for currents greater than 20 A or so, there are 4-terminal metal alloy sense resistors that have values down to 100 µΩ. Some manufacturers specify the higher value resistors with tighter tolerance than the low value resistors, so you have a basic trade-off here—use 1 mΩ at 0.1% vs. 400 µΩ at 0.5%.

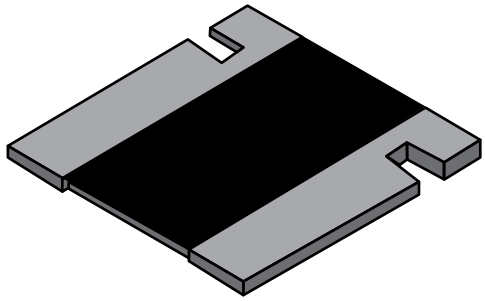

Figure 12. A 4-terminal shunt resistor.

Please refer to "Optimize High-Current Sensing Accuracy by Improving Pad Layout of Low-Value Shunt Resistors" for more specifics on accuracy when laying out sense resistors.

Using the LTC2977/LTC2979/LTC2980/LTM2987 to Measure Output Current

The LTC2977/LTC2979/LTC2980/LTM2987 devices have a limited ability to measure current. They can be configured to measure current on odd-numbered channels: channels 1, 3, 5, and 7. To configure for current measurement, the channel must be set to hi-res mode (MFR_CONFIG_LTC2977, bit9). This allows the V_{SENSEM} pin to be tied to a common-mode voltage up to 6 V. The V_{SENSEP} and V_{SENSEM} pins may be connected across an inductor (DCR) or resistive sense (R_{SNS}) element.

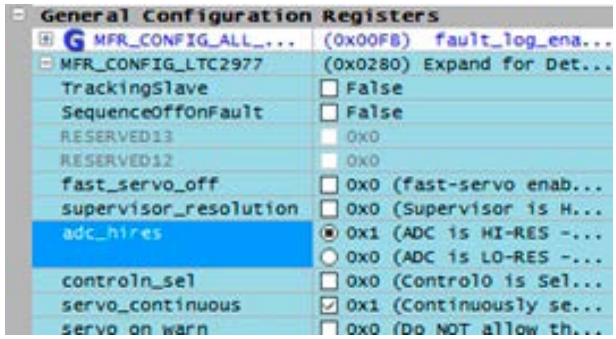

Figure 13. MFR_CONFIG adc_hires bit.

Even-numbered channels do not support this feature, and the V_{SENSEM} pin (channels 0, 2, 4, and 6) must remain within ±100 mV of GND.

In this mode, the only function that this channel provides is telemetry readback of the current. Setting the adc_hires bit disables the VOUT_EN pin and disables all fault responses. Essentially it forces the channel to the "off" state as far as the LTC2977 is concerned, and it only reads back the voltage in mV across the sense element.

The LTC2977/LTC2979/LTC2980/LTM2987 devices do not have a convenient READ_IOUT register or a register to store the DCR or R_{SNS} value. Instead, you use the READ_VOUT command to get raw differential voltage readings. The system host needs to calculate the current based on this reading divided by the sense resistor value. Note that these values are given in L11 format, not L16 format. The units are millivolts. If a system host or FPGA/CPU is used to read current, it must perform the math to convert a millivolt value to a value in milliamps or amps. The application note AN135 covers example code to convert L11 hex to a floating-point value.

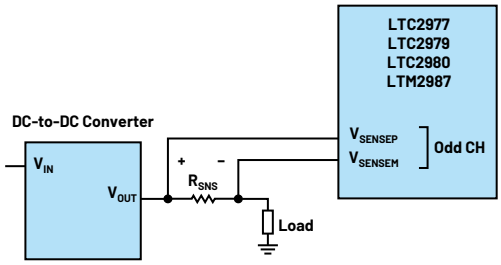

Figure 14. V_{SENSE} pins used for differential current sensing.

LTpowerPlay has a feature that conveniently translates this mV reading to a current readback value in mA. There is a scale factor that can be used to generate an adjusted value in the READ_VOUT register. This is accessed by clicking the **Setup** tab in the **Config** window.

The value entered into the **VOUT Display Scaling** box should be equal to $1/R_{SNS}$. If an external CSA is used, one needs to set the scale factor to $1/(GAIN_{CSA}/R_{SNS})$. There is a **Display Units** field that can be changed from volts to amps by replacing V with A. These changes allow the readout to display a computed current that is consistent with the actual current based on the sense resistance in the circuit. For example, if the R_{SNS} is 10 mΩ (0.01 Ω), the **VOUT Display Scaling** is 100. The READ_VOUT register will now report a value in mA that reflects 100 mA for every mV that is measured by the chip. In this example, a 592 mA load was applied to a supply rail with R_{SNS} of 10 mΩ, and the chip was measuring 5.92 mV. Note: the scaling/offset values under **Setup** are not saved to the device's NVM, but they are saved to the .proj file.

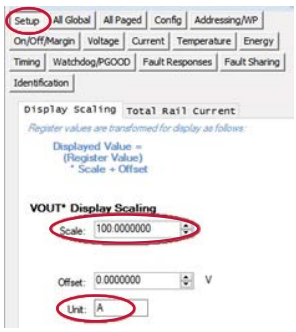

Figure 15. VOUT Display Scaling in the Setup tab.

Telemetry -- Output Voltage (V)	
MFR_VOUT_PEAK_LTC	+0.000 A
READ_VOUT	+592.041 mA
MFR_VOUT_MIN_LTC	+0.000 A

Figure 16. READ_VOUT telemetry shows scaled value and units in mA.

Since the differential voltage (VSENSEPn – VSENSEMn) is limited to ±170 mV, the sense element must be chosen so that the IR drop does not exceed this limit. The common-mode voltage of these pins is allowed up to 6 V. For example, if the current is expected to be in the 3 A range, a sense resistor of 50 mΩ provides 150 mV of voltage to the ADC and allows excursions to 3.4 A. This is great for the accuracy because you have a large signal, but 150 mV is a fairly large IR drop in the output path. This is a basic trade-off to be made when deciding between current measurement accuracy and IR drop in the output. One should always close the feedback loop at the load. This allows the regulator/servo to adjust to the proper output voltage. Refer to the LTC2977 data sheet for further details.

As an example, one of the odd-numbered channels has been allocated to measure output current. Channel 7 measures the I_{OUT} of Channel 6, a 3.0 V supply.

READ_VOUT (All Pages in System)	
U0:0 - LTC2977	1.2004 V
U0:1	1.5001 V
U0:2	1.8002 V
U0:3	1.9995 V
U0:4	2.5002 V
U0:5	2.7003 V
U0:6	2.9995 V
U0:7	+453.125 mA

Figure 17. READ_VOUT translated to mA (Channel 7).

When odd-numbered channels are configured to ADC hi-res mode, the VOUT_EN pin cannot be used and the supervisor function is disabled; hence, there is no possibility to quickly detect an overcurrent condition. However, it is possible to supervise current on any channel (in ADC low-res mode) if you use a CSA and output a single-ended signal to the V_{SENSEP} pin. This dedicates a voltage channel to supervising the output of a CSA. The propagation delay will be determined by the sum of the delay through the CSA, the delay introduced by the PSM device, and any delays that passive components (that is, RC) may introduce. The PSM delay depends on the configuration, whether the fault response is set to immediate off or deglitched off and delay count setting.

OC/UC Fault Supervision

Protecting loads from an overcurrent condition may be desirable on rails that power high value loads. Output current supervisors are built into the LTC2974/LTC2975. The dedicated hardware allows the user to configure a channel to shut down if the supervisor detects an overcurrent or undercurrent condition. These devices have both voltage and current supervision, which means that the channel will be shut down if the output voltage or output current goes outside the user-defined limits. The voltage supervisor and current supervisor are combined into the VOUT_EN logic internally. Table 1 in Part 1 of this article summarizes this feature for all PSM managers.

Output Current	
IOUT_OC_FAULT_LIMIT	10.000 A
IOUT_OC_WARN_LIMIT	5.000 A
IOUT_UC_FAULT_LIMIT	-1.000 A
Fault Responses -- Output Current	
⊞ IOUT_UC_FAULT_RESPONSE	(0x00) Ignore
⊞ IOUT_OC_FAULT_RESPONSE	(0xBF) Deglitched Off,Retry

Figure 18. I_{OUT} OC/UC fault/warn limits.

The fault supervisors are sampled comparators that have a user-adjustable threshold. The comparator is sampled every 12.1 µs and allows the user to deglitch output noise based on user-defined settings. The supervisor will trip only when the fault condition has been present for a consecutive number of times, or what is called the delay_count. This is essentially a time-based filter. The delay_count can be set up to 7, which provides an 84 µs deglitch response to an OC event. This allows a narrow glitch to go undetected while declaring a fault for wider pulses. Any RC filters that are inserted between the load and manager add an additional delay. The filter reduces the amplitude of any glitches but adds a delay to the response time of the supervisor. The data sheet suggests a time constant that is a tenth of the switching frequency yet not so long that the delay through the filter is much longer than the supervisor response time. For quiet supplies that need a fast OC response, one may choose 200 Ω/10 nF or 2 µs delay. For noisy supplies, an RC of 1 kΩ/0.1 µF yields a delay of 100 µs. While this may seem like a lot of delay, it is much faster than an ADC reading, which could be ~100 ms.

Undercurrent supervision will detect a low current or reverse current condition in the output. A low current condition is typical of light loads, and a UC fault

might not be desirable. However, the measured output current value includes negative values. While undercurrent supervision is not typically used, it may be used to detect reverse current conditions by setting IOUT_UC_FAULT_LIMIT to a negative value. To disable UC fault detection, set the IOUT_UC_FAULT_RESPONSE to ignore and set the IOUT_UC_FAULT_LIMIT to a large negative value. The default setting is –1 A.

While the LTC2971/LTC2972 do not have OC fault detection, the devices have an OC warning feature that will pull ALERTB low based on the ADC output current measurement. Warnings will pull ALERTB low and update the STATUS_IOUT register. The ADC-based reading will result in a slower response and is meant to be used as status indicators via a hardware pin and PMBus® register. It is possible to tie ALERTB to a CONTROL pin to shut down the channel. Or a microcontroller can respond to ALERTB by declaring an interrupt and drive a CONTROL pin or issue a PMBus command to shut off the channel. The downside to tying ALERTB to CONTROL is that any warning or fault will shut down the channel.

When using T_{SENSE} pins on the LTC2971/LTC2972/LTC2974/LTC2975 for compensating inductor DCR, the temperature may be used to shut down channels, which is another form of supervision. The overtemperature fault, warning limits, and fault response may be adjusted to suit the application on a channel-by-channel basis. That is, it can be used to shut down an individual channel and is not a global (entire chip) setting.

Current Readback L11 Format

The hex value being read back from the PSM device is formatted as L11. Whether you are reading the READ_VOUT register on an LTC2977 (ADC hi-res mode) or reading the READ_IOUT register on an LTC2975/LTC2974/LTC2972/LTC2971, the L11 format is a signed value that has a 5-bit exponent and 11-bit mantissa.

The L11 format supports polarity of the current measurement. Because it is a signed format, this allows the READ_IIN and READ_IOUT registers to provide this information to the system host about the direction of the current flow. The LTC2974/LTC2975 have undercurrent thresholds for the output current. A negative value could be useful to shut down a channel that sinks too much current, a reverse current.

There is a particular point to make about the L11 format, and that is the granularity. The LTC2971/LTC2972/LTC2974/LTC2975 data sheets show a table that lists the granularity of the READ_IOUT value across a wide range of currents. There is an inherent granularity that is due to the L11 hex format and is not the device's ADC or any other hardware limitations. The table also lists MFR_READ_IOUT granularity for comparison. The MFR_READ_IOUT value is a custom format and provides improved resolution with 2.5 mA granularity above 2 A. It is limited to ±81.92 A. If the board's host CPU/FPGA need to convert L11 to floating point, it can issue reads to either register. The READ_IOUT register has better resolution for currents below 2 A and does not have the 81.92 A limit, but the MFR_READ_IOUT values will resolve to the nearest 2.5 mA value.

Programming PSM Devices and LTpowerPlay

As with the entire family of LTC297x devices, programming the PSM device and successfully powering your hardware for the first time can be very rewarding. The use of LTpowerPlay is your easiest path. LTpowerPlay is a free download and runs on Windows®. The software has a built-in programming utility that takes the configuration data you have saved and writes it to the device's EEPROM. After power cycling, the chip automatically loads its RAM from the EEPROM and is ready to run autonomously.

Whether you're new to LTpowerPlay or already a power user, learn how to configure, design, evaluate, diagnose, and debug using the LTpowerPlay Software-Based Power Supply Configuration and Debug Tool. If you do not plan to use LTpowerPlay for programming or to provide telemetry, downloading Linduino C code examples is an alternative solution. The code examples are provided in the LTSketchbook zip file.

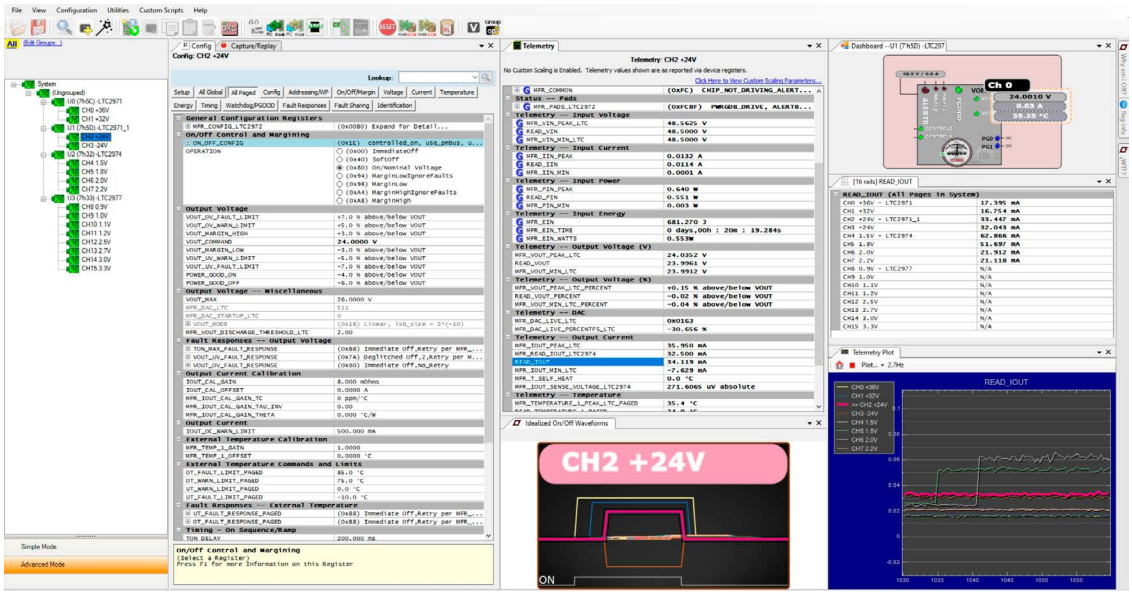

Figure 19. LTpowerPlay is a powerful, Windows-based development environment supporting Analog Devices' digital power system management (PSM) products.

Checklist to create .proj file:

- Ensure each PSM device has a unique address on PMBus (hardware strapping).
- Set IOUT_CAL_GAIN on each output channel.
  - This is the R_{SENSE}, inductor DCR, or calculated I_{MON} value.
- Set IIN_CAL_GAIN on each device that measures input supply current (LTC2971/LTC2972/LTC2975).
- Set temperature-related configuration (for example, MFR_IOUT_CAL_GAIN_TC, MFR_IOUT_CAL_GAIN_TAU_INV, MFR_IOUT_CAL_GAIN_THETA).
- Set IOUT_OC_FAULT_LIMIT and IOUT_OC_FAULT_RESPONSE (LTC2974/LTC2975).
- Set IOUT_UC_FAULT_LIMIT and IOUT_UC_FAULT_RESPONSE (LTC2974/LTC2975).
- Tip: Use the Configuration Wizard in LTpowerPlay to ease file generation.

Summary

ADI's DPSM LTC297x devices are mixed-signal PMBus ICs that can measure and supervise power supply currents. Various sensing methods have been presented—resistor shunt, inductor DCR, and I_{MON} are among them. The current measurement capabilities add to the feature set of the family by providing another level of protection in the form of OC/UC fault supervision. These devices bring to any power supply the ability to monitor, supervise, and measure voltages and currents. These features are very desirable for supply rails of high value. The LTC297x offers the ability to configure the device's PMBus registers, which adds flexibility to make changes to your board design at any point in the design phase, even after the board is deployed in the field.

About the Author

Michael Peters is a senior applications engineer for power system management devices at Analog Devices. He has more than 30 years of experience in analog and digital circuits, including working on memory devices at previous companies. He received his B.S.E.E. degree from the University of Michigan, Ann Arbor, Michigan. He can be reached at michael.peters@analog.com.

RAQ Issue 197: RF Demystified—Understanding Wave Reflections

Anton Patyuchenko, Field Applications Engineer

Question:

Wave reflections—why are they important to understand in RF designs?

Answer:

This article presents a short discussion for non-RF engineers on the terminology associated with one of the key properties inherent to RF designs: wave reflections. The key difference between ordinary circuits operating at low frequencies and circuits designed for RF frequencies is their electrical size. RF designs can be of many wavelengths in size, leading to the variation of voltages and currents in magnitude and phase over the physical dimensions of their elements. This gives RF circuits a number of fundamental properties[1] that underly the core principles used for their design and analysis.

Basic Concepts and Terminology

Consider a transmission line—for example, a coaxial cable or a microstrip—terminated in an arbitrary load, and define the wave quantities a and b as shown in Figure 1.

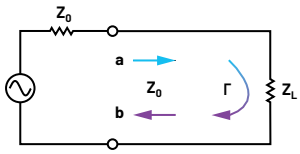

Figure 1. Transmission line with a matched signal source terminated in a single-port load.

These wave quantities are complex amplitudes of the voltage waves incident on and reflected from this load. We can use these quantities now to define the voltage reflection coefficient, Γ, which describes the ratio of the complex amplitude of the reflected wave to that of the incident wave:

$$\Gamma = \frac{b}{a} \tag{1}$$

The reflection coefficient can also be expressed using the characteristic impedance of the transmission line Z_0 and the complex input impedance of the load Z_L as:

$$\Gamma = \frac{Z_L - Z_0}{Z_L + Z_0} \tag{2}$$

RF engineering typically relies on $Z_0 = 50\ \Omega$, which is a compromise between signal attenuation and power handling capacity that can be achieved with coaxial transmission lines. However, in some applications—for instance, in broadcasting systems where RF signals need to be transmitted over longer distances—$Z_0 = 75\ \Omega$ is a more common choice allowing for lower cable losses.

Regardless of what that value of the characteristic impedance is, if the load impedance is the same ($Z_L = Z_0$), this load is said to be matched to the transmission line. It should be noted that this condition is only valid when the signal source is matched to the transmission line, as shown in Figure 1, which we assume in this article. In this case, we will not have any reflected waves ($\Gamma = 0$) and the load will receive the maximum power from the signal source, while in the case of total reflection ($|\Gamma| = 1$), no power is delivered to the load at all.

If the load is mismatched ($Z_L \neq Z_0$), it will not receive all the incident power. The corresponding "loss" in power is known as return loss (RL), which can be related to the magnitude of the reflection coefficient using the following equation:

$$RL = -20 \log(|\Gamma|)\ \text{dB} \tag{3}$$

Return loss describes a ratio of the power incident on the load to the power reflected back from it. Return loss is always a non-negative quantity that indicates how well the load is matched to the impedance of the network "seen" at the load toward the source.

If the load is mismatched, the presence of the reflected wave leads to standing waves, resulting in a non-constant voltage magnitude that changes with position along the line. The measure used to quantify this impedance mismatch of the line is called standing wave ratio (SWR) and defined as:

$$SWR = \frac{1 + |\Gamma|}{1 - |\Gamma|} \quad (4)$$

Since SWR is often interpreted in terms of maximum and minimum voltages, this quantity is also known as voltage standing wave ratio (VSWR). SWR is a real number that can take values in the range from 1 to infinity, where SWR = 1 implies a matched load.

Conclusion

RF circuits have a number of fundamental properties that distinguish them from ordinary circuits. Design and analysis of microwave circuits require the use of extended concepts to resolve the problems of practical interest. This article introduced and discussed some of the key concepts and terminology related to one of the main properties of RF systems: wave reflections.

Analog Devices provides the broadest portfolio of RF integrated circuits in the industry coupled with deep system design expertise to address the most demanding requirements across a wide variety of RF applications. Moreover, ADI offers the entire ecosystem including design tools, simulation models, reference designs, rapid prototyping platforms, and a discussion forum to support RF engineers and ease the development process of their target applications.

References

[1] M. S. Gupta. "What Is RF?" *IEEE Microwave Magazine*, Vol. 2, No. 4, December 2001.

Hiebel, Michael. *Fundamentals of Vector Network Analysis*. Rohde & Schwarz, 2007.

Pozar, David M. *Microwave Engineering*, 4th Edition. Wiley, 2011.

About the Author

Anton Patyuchenko received his Master of Science in microwave engineering from the Technical University of Munich in 2007. Following his graduation, Anton worked as a scientist at German Aerospace Center (DLR). He joined Analog Devices as a field applications engineer in 2015 and is currently providing field applications support to strategic and key customers of Analog Devices specializing in RF applications. He can be reached at anton.patyuchenko@analog.com.

IBIS Modeling—Part 3: How to Achieve a Quality Level 3 IBIS Model Through Bench Measurement

Christine C. Bernal, Product Applications Engineer,
Janchris Espinoza, Product Applications Engineer, and
Aprille Hernandez-Loyola, Product Applications Engineer

Abstract

IBIS models are commonly generated through design circuit simulations. However, there are some cases when the design files are obsolete, unavailable, or only available in an unworkable schematic file format due to old, released parts. This article aims to provide a high level procedure for generating IBIS models via bench measurement using an actual unit—from data extraction to model validation. A dedicated test fixture that minimizes impedance mismatches, which can arise from parasitic traces, was used in data gathering to manage signal integrity constraints and ensure a reliable IBIS model. It was then validated through simulation and bench measurement, making it compliant to Quality Level 3 of the IBIS model.

Introduction

The Input/Output Buffer Information Specification (IBIS) is a behavioral model that is gaining worldwide popularity as a standard format to generate device models. The accuracy of the device model depends on the quality of the IBIS models offered by the industry. Therefore, providing quality, reliable IBIS models for signal integrity simulations is a strong commitment to the customer.

One way of generating IBIS models is through simulations; however, there are some cases when the design files are not available, making it impossible to generate IBIS models from simulation results. In this case, generating IBIS models via bench measurement is a solution to this gap that can offer a high quality and more realistic device behavioral model. Figure 1 shows the complete stage in generating IBIS models through bench measurement. Using the actual silicon, the device's receiver and driver buffer behaviors are extracted to represent the current vs. voltage (I-V) data and the voltage vs. time (V-t) data.

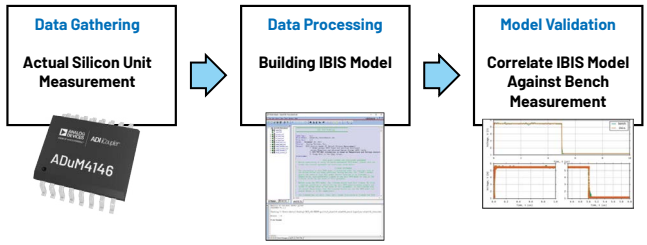

Figure 1. IBIS model through bench measurement generation process.

The model will then be validated against an actual bench setup with complete loading conditions. This procedure provides a Quality Level 2b IBIS model. To achieve the higher Quality Level 3 model, the generated IBIS model will also be validated against the device's transistor-level design, also with the recommended loading conditions.

To characterize the quality, the IBIS Quality Task Group has formulated a quality control (QC) process using five QC stages. They developed a checklist to define different quality levels, as shown in Table 1.

Table 1. Quality Levels in the IBIS Quality Specification

Quality Level	Description
0	Passes IBISCHK
1	Complete and correct as defined in checklist documentation
2a	Correlation with simulation
2b	Correlation with actual silicon measurement
3	All of the above

The quality levels presented in Table 1 provide a standard for IBIS model quality, which varies from vendor to vendor.[1] Having a standard for IBIS model accuracy

will ensure customers that they are getting accurate and reliable models. The higher the quality level of the model is, the more accurate its data since higher quality levels require more validation processes.

Based on the book *Semiconductor Modeling: For Simulating Signal, Power, and Electromagnetic Integrity* by Roy Leventhal and Lynne Green,[2] there are five recognized quality levels of the IBIS correctness checklist.

Quality Level 0—Passes IBISCHK

A Quality Level 0 requirement should at least pass the IBIS parser. IBISCHK must yield zero errors, and all warnings must be explained if they cannot be eliminated. Ideally, there should be no warnings, but it is recognized that some warnings cannot be removed. The "Error," "Warning," and "Note" messages from the parser check serve as guides for the IBIS model maker to identify errors and easily correct them. See Figure 2 for the IBIS model parser check.

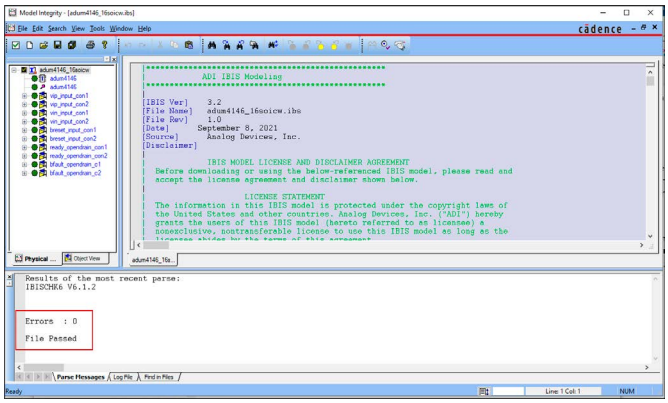

Figure 2. IBIS model passes IBISCHK.

Quality Level 1—Complete and Correct as Defined in Checklist Documentation

A Quality Level 1 IBIS model passes Quality Level 0 with an additional check for correctness and completeness of a basic simulation test. It includes the correctly defined package parasitic, pin configuration, and load parameters. Ramp rate and typical/minimum/maximum values must be in accordance with the device specifications. The detailed requirements under Quality Level 1 found here can also be used as a reference.

Quality Level 2a—Correlation with Simulation

Quality Level 2a compares the performance of an IBIS model to the device's transistor-level design. The IBIS model's performance when connected to a load is correlated against the device's transistor-level design with the same load. The results from the two simulation setups are then compared and checked if the model passes Quality Level 2a. Details are discussed in the section "Validation and Results."

Quality Level 2b—Correlation with Actual Silicon Measurement

Quality Level 2b compares the performance of an IBIS model against the device's actual unit. Like Quality Level 2a, the same load must be connected to the two setups during correlation. The model will pass as a Quality Level 2b based on the correlation results. Details will be discussed in the section "Validation and Results."

Quality Level 3—Correlation of Transistor-Level Simulation and IBIS Bench Measurement

Quality Level 3 specifies that the IBIS model is validated against the transistor-level design and the actual unit. For the model to pass Quality Level 3, it must pass the correlation for both quality levels 2a and 2b. On top of that, the model must pass the IBIS parser test (Quality Level 0) and satisfy the IBIS quality checklist (Quality Level 1). Details will be discussed in the section "Validation and Results."

The Use Case

In this article, the case under study is the ADuM4146, an isolated gate driver specifically optimized for driving silicon carbide (SiC) MOSFETs. The ADuM4146 has three input pins (VIP, VIN, and $\overline{\text{RESET}}$) and two open-drain pins (READY and $\overline{\text{FAULT}}$), but this article will only discuss one pin for each buffer type. It is because the procedure for building and validating the IBIS models for pins with similar buffer types is the same. The VIP pin will be used as the use case for input buffer, and $\overline{\text{FAULT}}$ will be used as the use case for open-drain buffer.

It is important to note that although similar buffer types have the same IBIS modeling procedure and validation, it does not necessarily mean that they have the same IBIS data. The article only discusses one pin per buffer type to simplify the explanation of building the IBIS model and validation process.

The ADuM4146 has a standard small outline wide-body package (SOIC_W), which is represented as a resistor, inductor, and capacitor (RLC) parasitic in the validation process. The package RLC values were extracted through simulations by a package engineer. The dedicated printed circuit board (PCB) has a similar case

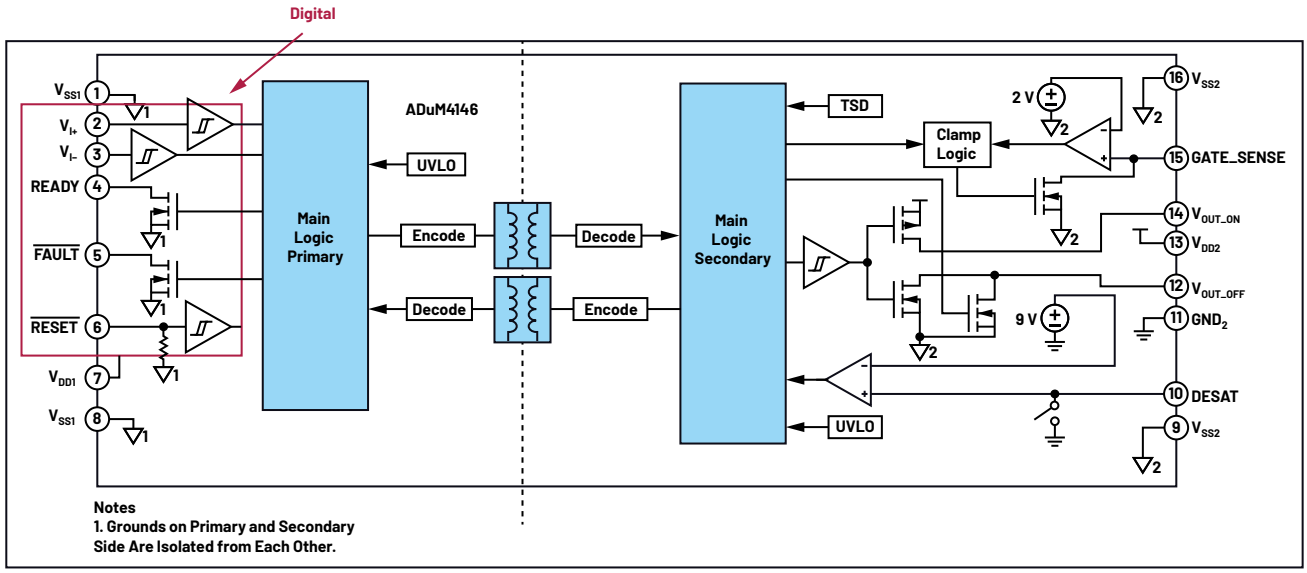

Figure 3. ADuM4146 functional block diagram.

to the package parasitic: it is represented by the RLC parasitic, and the values were extracted by a PCB engineer.

Table 2 shows the ADuM4146 pin configuration and the corresponding buffer type for each pin. This information will be used in the [Pin] keyword of the IBIS model.

Table 2. ADuM4146 Pinouts and Their Corresponding Buffer Type

[Pin]	Signal Name	Model_Name
1	VSS1	GND
2	VIP	vip_input
3	VIN	vin_input
4	READY	ready_opendrain
5	$\overline{\text{FAULT}}$	bfault_opendrain
6	$\overline{\text{RESET}}$	breset_input
7	VDD1	POWER
8	VSS1	GND
9	VSS2	GND
10	DESAT	NC
11	GND2	GND
12	VOUT_OFF	NC
13	VDD2	POWER
14	VOUT_ON	NC
15	GATE_SENSE	NC
16	VSS2	GND

IBIS Bench Measurement Procedure

Gathering data through bench measurement may be affected by different external factors. These factors should be compensated to achieve correlation and provide quality models.

To minimize the effect of external factors, the device under test (DUT) is placed on a dedicated fixture, which aims to reduce unwanted capacitance that may cause inaccuracy to the measured device behaviors, as shown in Figure 4. Parasitic capacitance is a significant problem in actual silicon measurement and is often the factor limiting the operating frequency and bandwidth of a device model.

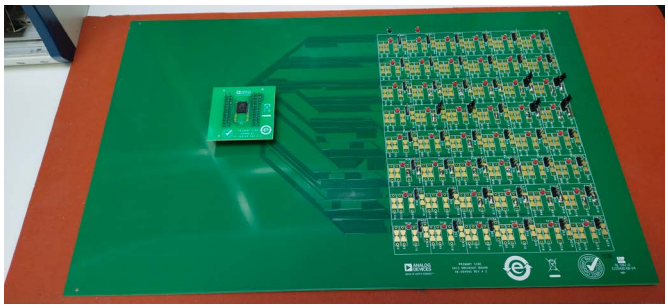

Figure 4. Dedicated fixture for IBIS bench measurement.

Steps on generating an IBIS model through bench measurement:

Prepare the Setup

Table 3 shows the IBIS pre-modeling stage requirements for bench measurement, and Table 4 shows the different model types and model components that define the buffer behavior. The model types are discussed in detail in articles "IBIS Modeling—Part 1: Why IBIS Modeling Is Critical to the Success of Your Design"[3] and "IBIS Modeling—Part 2: Why and How to Create Your Own IBIS Model."[4] You may also refer to the *IBIS Modeling Cookbook*.[5]

Table 3. IBIS Bench Measurement Pre-Modeling Stage

Requirements	Contents
Device Under Test (DUT)/Sample Units	Provide tested good units
Adaptor Board	Define the package type of the device
RLC Package Parasitic Values	Provide the bonding diagram of the device
Product Data Sheet Specification	Consider the following: ▶ Logic supply voltage range ▶ Digital supply range (if applicable) ▶ Pin configurations ▶ Operating temperature range ▶ Logic high/low input voltage range ▶ Logic high/low output voltage range ▶ Timing test loads and characteristics ▶ Theory of operation

Table 4. IBIS Bench Measurement Buffer Type Selection

Model Type	Model Components
Input	▶ [Power Clamp] ▶ [Ground Clamp]
Output 3-State, Output 2-State, I/O	▶ [Power Clamp] ▶ [Ground Clamp] ▶ [Pullup] ▶ [Pulldown] ▶ [Rising Vddref] ▶ [Falling Vddref] ▶ [Rising Gndref] ▶ [Falling Gndref]
Open_drain, I/O_open_drain	▶ [Power Clamp] ▶ [Ground Clamp] ▶ [Pulldown] ▶ [Rising Vddref] ▶ [Falling Vddref]
Open_source, I/O_open_source	▶ [Power Clamp] ▶ [Ground Clamp] ▶ [Pullup] ▶ [Rising Gndref] ▶ [Falling Gndref]

Bench Setup

Understanding how the device operates is essential in data gathering for IBIS models. As shown in Figure 1, this is the first stage and this is done by extracting the I-V data and the V-t data. Both are represented in tabular form.

I-V data includes the ESD clamp behavior and the driver strength, while the V-t data indicates the transition from low state to high state and vice versa. The switching behavior is measured with a load connected to the output pin, equivalent to the value that the output buffer will drive. Nevertheless, the usual load value is 50 Ω to represent the typical transmission line impedance.

For I-V measurement, a programmable power supply capable of sinking and sourcing current and a curve tracer are used to sweep the voltage and gather the current behavior of the buffer. The data is recommended to be taken in the voltage range of $-V_{DD}$ to $2 \times V_{DD}$, and the typical, minimum, and maximum corners. V-t measurement requires the use of the oscilloscope with appropriate bandwidth and a low capacitance probe.

The DUT is mounted on the dedicated fixture and is to be tested under varying temperature conditions with the use of a temperature forcing system to capture the minimum, typical, and maximum performance. In this case, the minimum (weakest drive strength, slowest edge) data is taken at 125°C, and the maximum (strongest drive strength, fastest edge) data is taken at –40°C.

Bench Data Extraction

Once verified that the bench setup is ready, the process of gathering the required I-V and V-t data can begin. Output and I/O buffers demand both I-V tables and rise/fall data, while input buffers only demand I-V tables.

- I-V (Current vs. Voltage) Data Measurement

The I-V curve measurements cover the four IBIS keywords—[Pullup] and [Pulldown] represent the I-V behavior of the pull-up component when driving high and the pull-down component when driving low, while [Power Clamp] and [GND Clamp] represent the I-V behavior of the ESD protection diodes during the high impedance state.

To measure the I-V characteristics, mount the component on the dedicated board and connect the power and ground pins to the power supply. Prepare the temperature forcing system, adjust to the desired temperature, and wait for it to stabilize. Sweep the voltage within the recommended range, then, using the curve tracer, measure the current of the required buffer.

The positive node of the sweeping device for pull-up and power clamp data should be connected to the supply voltage and the negative node should be connected to the pin, while the sweeping device for pull-down and ground clamp data are referenced to ground. Extrapolation may be required at times when the curve tracer is not able to sweep the entire range.

Figure 5 shows the bench setup for input buffer (V_{I+}) I-V ground clamp measurement, while Figure 6 shows its measured behavior. The ground clamp circuit is triggered when the input goes below ground resulting in a negative current, approaching and settling at zero. The input pin (VIP) does not have a power clamp component, so its model will not have power clamp data.

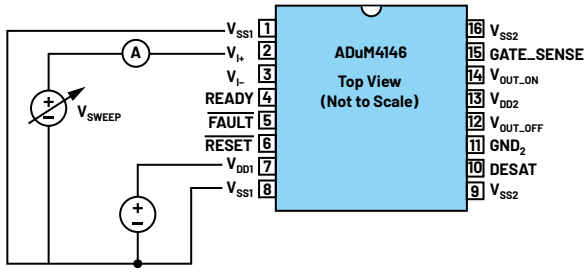

Figure 5. ADuM4146 bench setup for I-V clamp measurement.

Figure 6. ADuM4146 input buffer bench measured ground clamp.

The same method is implemented to the ESD clamp, pull-up, and pull-down data of the output buffer. In this case, however, the ADuM4146 READY and $\overline{\text{FAULT}}$ pins are open-drain buffers; thus, they do not have a pull-up component and they only require pull-down data.

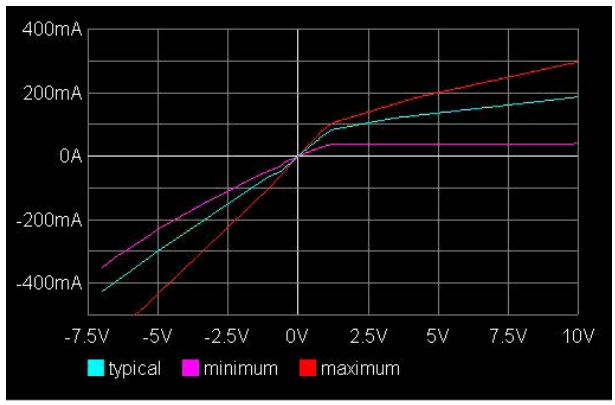

Figure 7. ADuM4146 open-drain buffer pull-down result.

Figure 7 shows the ADuM4146 open-drain buffer's pull-down data result. The pull-down curve starts from a negative current, then crosses zero going to the positive quadrant, also in the range of $-V_{DD}$ to $2 \times V_{DD}$.

- Buffer Capacitance (C_comp) Extraction

According to the *IBIS Modeling Cookbook for IBIS Version 4.0*, "The total die capacitance of each pad, or the C_comp parameter, is the capacitance seen when looking from the pad into the buffer for a fully placed and routed buffer design, exclusive of package effects."[5] One way to obtain the C_comp value is by using the following equation.

$$C\_comp = C_{IN} - C_{pkg}$$

Where:

C_{IN} = device input capacitance

C_{pkg} = device package capacitance

▶ V-t (Output Voltage vs. Time) Data Measurement

The V-t curve measurements also cover four IBIS keywords—[Rising Vddref] and [Falling Vddref] pertain to the transitions from low-to-high and high-to-low with load referenced to the supply, while [Rising Gndref] and [Falling Gndref] pertain to the transitions from low-to-high and high-to-low with load referenced to the ground. Related to these is the keyword [Ramp], which defines the transition rate when changing from one state to another, taken at 20% to 80% of the waveform.

Measuring the rise and fall time data requires the use of the oscilloscope on the buffer driving the required load. In this case, a 50 Ω resistor is used to represent the transmission line impedance. For the open-drain type, connect the load to the buffer and to the supply voltage to measure the switching behavior with reference to VDD1. Be sure to stabilize the temperature as required using the temperature forcing system to capture the minimum, typical, and maximum range. Figure 8 shows ADuM4146 actual bench setup for READY and $\overline{\text{FAULT}}$ pins' switching behaviors. Given that ADuM4146 digital output pins are open drain, only the rising and falling behaviors referenced to the supply voltage are required.

Figures 9 and 10 show the captured rising and falling waveforms for the $\overline{\text{FAULT}}$ pin both in transistor-level simulation and actual silicon measurement. Both setups use the same loading conditions of 50 Ω connected to VDD1, across typical, minimum, and maximum corners.

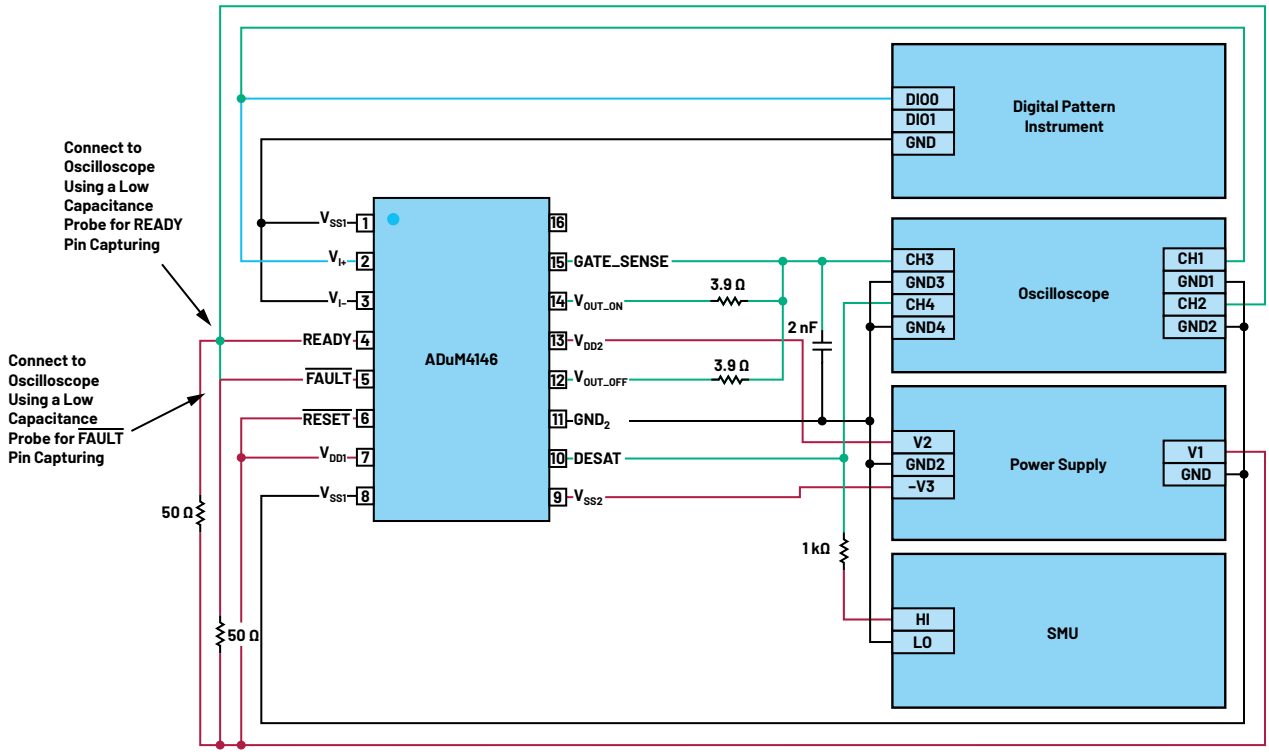

Figure 8. ADuM4146 bench setup for READY/$\overline{\text{FAULT}}$ switching behaviors.

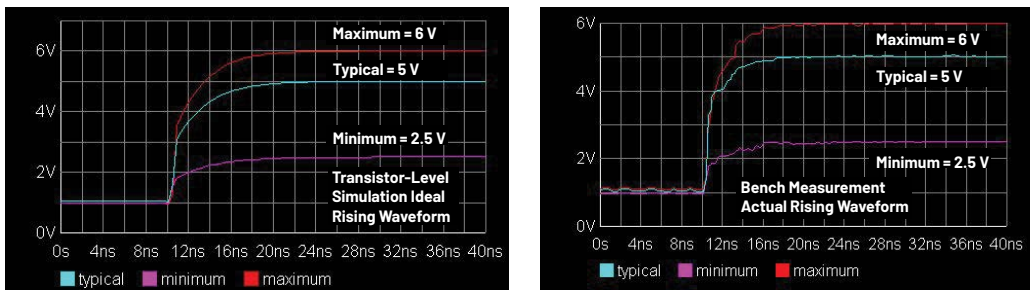

Figure 9. ADuM4146 $\overline{\text{FAULT}}$ pin rising waveform at VDD1 reference.

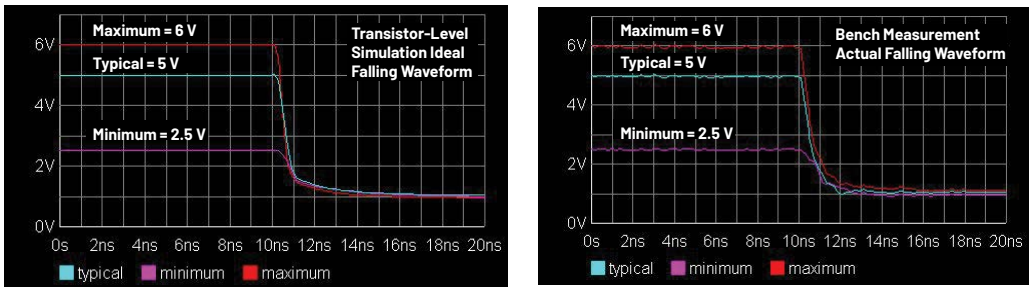

Figure 10. ADuM4146 $\overline{FAULT}$ pin falling waveform at VDD1 reference.

Building the IBIS Model

The next stage in creating an IBIS model is processing the gathered data and building the model itself. In this stage, the raw data tables are inserted in an IBIS text format following the necessary keywords and including the device parameters. The detailed process for this is discussed in the article "IBIS Modeling—Part 1: Why IBIS Modeling Is Critical to the Success of Your Design."[3]

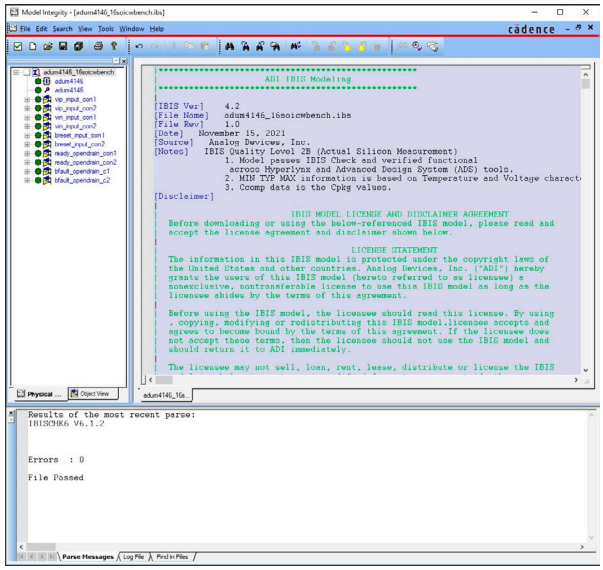

Figure 11. ADuM4146 IBIS model generated from bench measurement.

Figure 11 shows the ADuM4146 IBIS model generated from bench measurement. The model should pass the IBIS parser, which includes basic checks such as matching between the I-V and V-t tables and reviewing the monotonicity of table data. All errors, warnings, and notes should be completely resolved before continuing with the validation process. In addition, the model should satisfy the IBIS quality checklist.

Validation and Results

The validation process of this article will follow the one presented in the second article of this series, "IBIS Modeling—Part 2: Why and How to Create Your Own IBIS Model."[4] More details regarding the validation process of an IBIS model are discussed there.

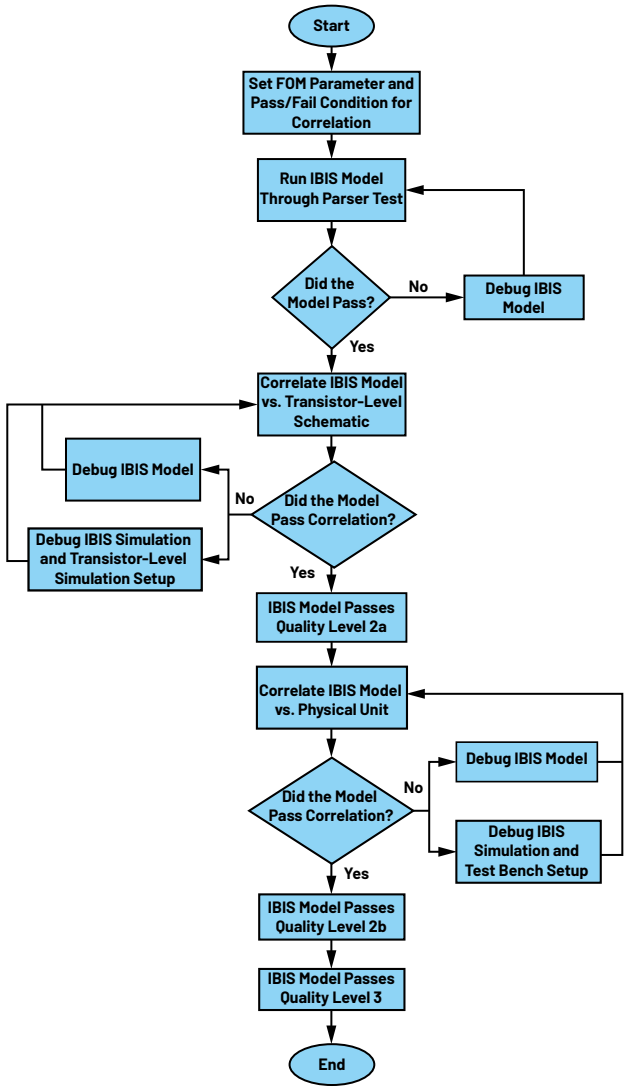

Figure 12. IBIS model Quality Level 3 validation process flowchart.

The model must first pass the parser test, which can be checked using a software with integrated IBISCHK or using the open-source executable code from ibis.org. After passing the parser test, the model must then be correlated against either its transistor-level schematic or the actual silicon unit. Since this article aims to achieve the Quality Level 3 model, ADuM4146's IBIS model will be correlated against

both its transistor-level schematic and actual unit. The figure of merit (FOM) value will be set to determine whether the IBIS model will pass both correlations. In this case, the FOM value for both correlations must be greater than or equal to 95% to pass the Quality Level 3 IBIS model validation. Figure 12 presents a flowchart diagram of the validation process an IBIS model must go through to pass Quality Level 3.

The area under the curve metric will be used to compute the FOM values of both correlations. The same loading conditions must be placed on the two sets of correlation. During validation, it is advisable to follow the loading conditions indicated in the data sheet to test the device in its normal operation.

To correctly validate the IBIS model against a reference—for example, IBIS vs. bench measurement correlation—the PCB traces that the signal will go through in the bench measurement setup must be added to the IBIS simulation setup.

Below are the two conditions performed to achieve the Quality Level 3 IBIS model.

IBIS Quality Level 2a Validation

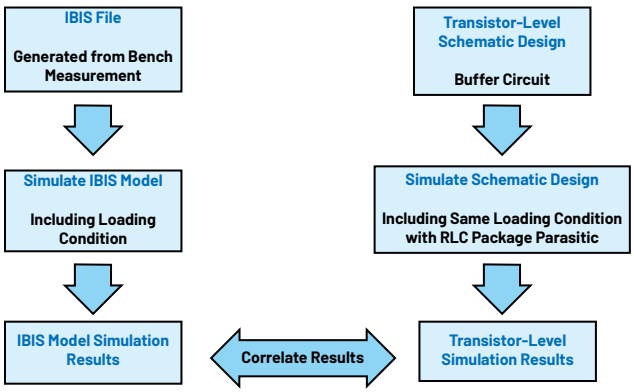

Figure 13. IBIS model Quality Level 2a validation process.

Figure 13 shows the IBIS model Quality Level 2a validation process. This correlation process intends to assess the degree to which the IBIS model data will result in simulations that match the transistor-level simulation results. Figure 14 shows the ADuM4146's IBIS model simulation setup for both its input and open-drain buffers with loading conditions.

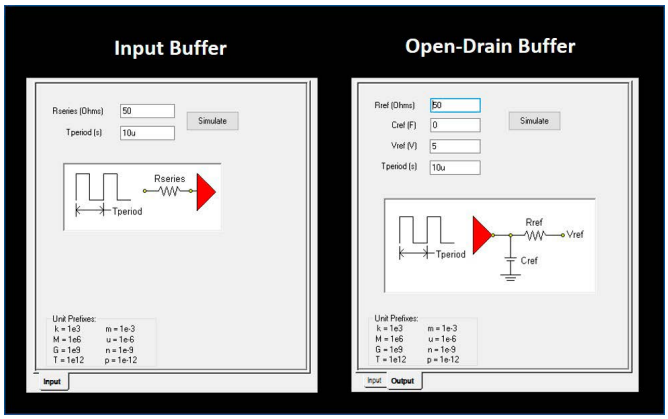

Figure 14. ADuM4146 input and open-drain buffer simulation setup.

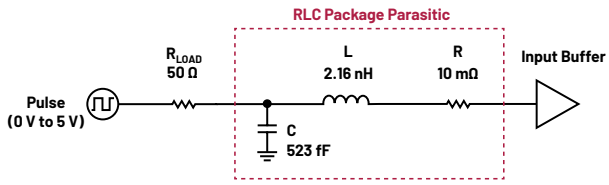

Figure 15. ADuM4146 transistor-level design simulation setup with loading conditions (input buffer).

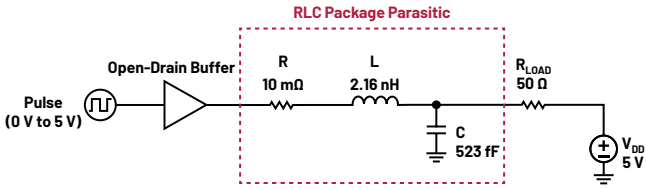

Figure 16. ADuM4146 transistor-level design simulation setup with loading conditions (open-drain buffer).

Figures 15 and 16 show the transistor-level design simulation setup with loading conditions for input and open-drain buffers, respectively. The package RLC values of the device are added in between the buffer and the load to replicate the package parasitic in the IBIS setup.

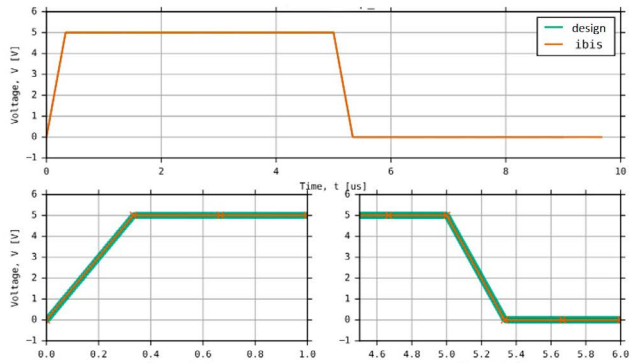

Figure 17. Transistor-level design vs. IBIS model validation results (input buffer).

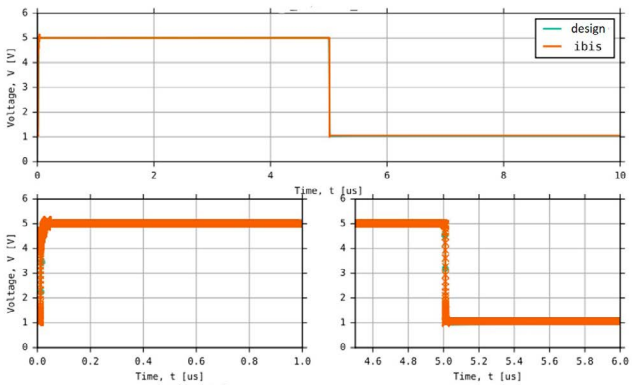

Figure 18. Transistor-level design vs. IBIS model validation results (open-drain buffer).

Figures 17 and 18 show the correlation results of both input and open-drain buffers, respectively, when running the IBIS model with standard load and comparing the results against a transistor-level reference simulation using the same load. A 50 Ω resistor is used as a load for the IBIS vs. transistor-level correlation setups of the open-drain buffer. A transient analysis is performed for both setups with a 10 µs pulse input.

Table 5 shows the computed FOM values for the two buffer models when correlated against their transistor-level schematic. Since both buffer models have FOM values greater than 95%, the IBIS model passes Quality Level 2a.

Table 5. Quality Level 2a Validation FOM Values of Input and Open-Drain Buffers

Buffer Model	FOM
Input	99.99%
Open Drain	99.68%

IBIS Quality Level 2b Validation

IBIS Quality Level 2b requires the model to be correlated against bench measurement, so factors that may affect the performance of bench measurement need to be considered. The main challenge when performing bench measurement is signal attenuation, which is mostly caused by trace parasitics. When measuring data from the actual unit, it is best to use a dedicated board with a low capacitance probe to reduce the effects of trace parasitics as much as possible. In this case, the IBIS bench dedicated board was a solution for signal integrity issues, reducing attenuation caused by unwanted signals that might get introduced to the signal of interest. Figure 19 shows the validation process for IBIS Quality Level 2b.

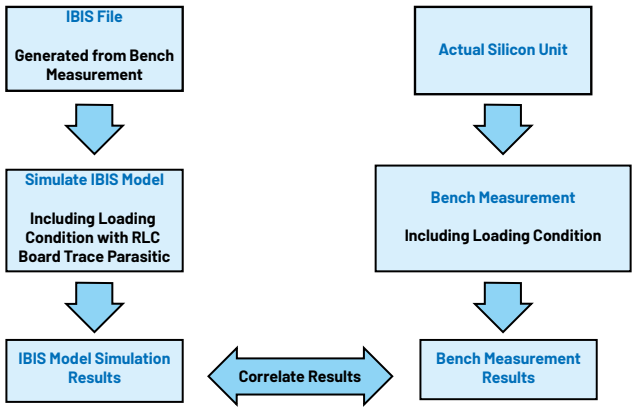

Figure 19. IBIS model Quality Level 2b validation process.

The main objective in the IBIS model correlation is to obtain results that are as close as possible to the reference. In capturing rise/fall time data in the oscilloscope, it is best to use a probe with extremely low loading to reduce signal attenuation. The errors introduced by the probe and instrument combination can provide a significant contribution to the signal of interest. According to Tektronix, "Special filtering techniques and proper selection of tools to de-embed the measurement system's effects on the signal, displaying edge times, and other signal characteristics are key factors to consider when measuring the actual silicon performance."[6]

Figures 20 and 21 show the simulation setups of the IBIS models with input and open-drain buffers, respectively, considering the loading conditions. The RLC values connected in series to the buffers are the parasitic values from the board traces. It is important to consider their effect on the model's performance when adding any loading by the fixture to replicate the lab setup.

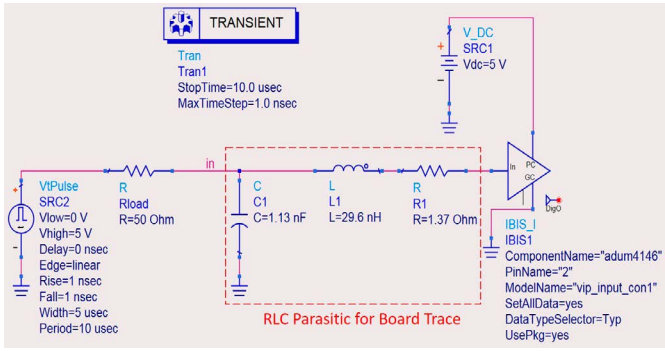

Figure 20. Actual IBIS simulation setup with loading conditions (input buffer).

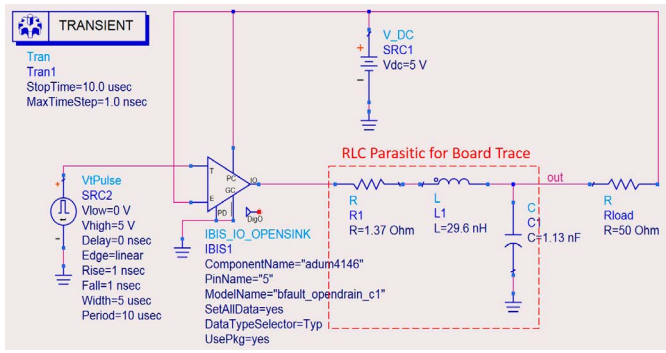

Figure 21. Actual IBIS simulation setup with loading conditions (open-drain buffer).

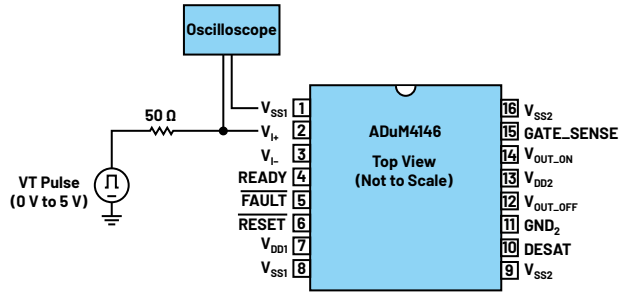

Figure 22. Bench setup with loading conditions (input buffer).

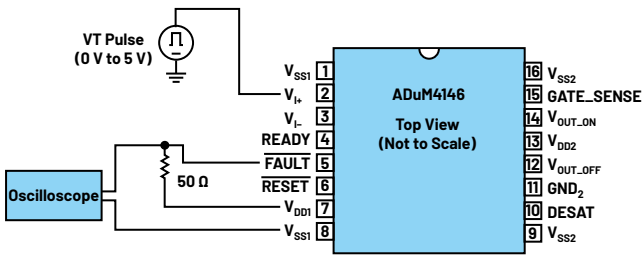

Figure 23. Bench setup with loading conditions (open-drain buffer).

Figures 22 and 23 present a diagram representation of the bench setups with loading conditions for input and open-drain buffers, respectively. A 5 V pulse signal is used to drive the open-drain buffer, which is connected to a 50 Ω load. Correlation results for IBIS simulation and bench measurement of input and open-drain buffers are shown in Figure 24 and Figure 25, respectively.

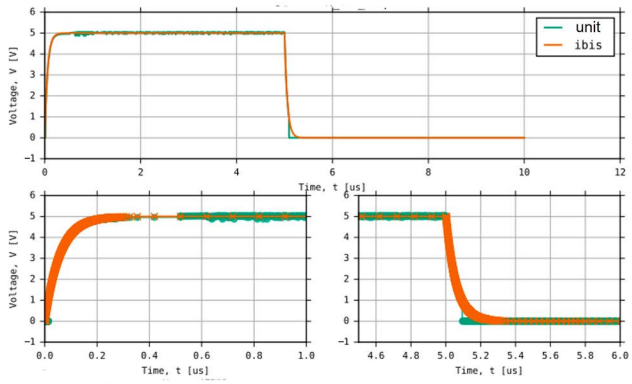

Figure 24. Actual silicon unit vs. IBIS model validation results (input buffer).

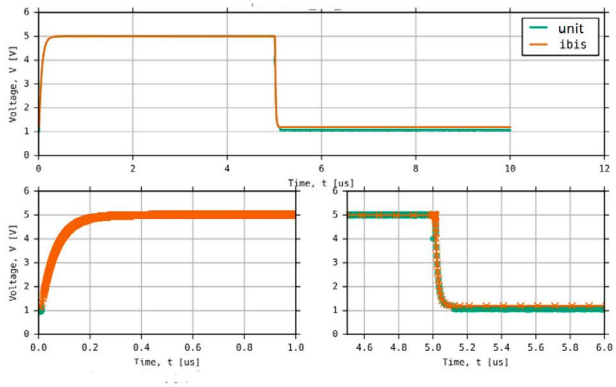

Figure 25. Actual silicon unit vs. IBIS model validation results (open-drain buffer).

Table 6 shows the FOM values of input and open-drain buffers when correlated against the actual silicon bench measurements. The FOM values are greater than 95%, which means that the IBIS models of the two buffers pass Quality Level 2b. Since the model passes Quality Level 2a and Quality Level 2b, it can now be considered as a Quality Level 3 IBIS model.

Table 6. Quality Level 2b Validation FOM Values of Input and Open-Drain Buffers

Buffer Model	FOM
Input	99.23%
Open Drain	98.52%

Conclusion and Takeaway

Extracting the data necessary for hardware-to-model correlation is one of the most challenging procedures in building a quality IBIS model through bench measurement. By careful attention to detail and understanding the behavior of the I/O circuit, it is possible to achieve a close correlation between the lab measurements and IBIS simulation results. Eliminating as much attenuation as possible is key to having a high FOM value in correlation. Considering this, it is advisable to use a dedicated test fixture with well-matched equipment and accessories to ensure the integrity of the signal.

It is also important to keep in mind that in correlation, the IBIS model and reference setups must be identical in terms of the traces that the signal will go through. This decreases the error caused in the correlation—thus, increasing the FOM value.

Having an IBIS Quality Level 3 model is an advantage for both semiconductor vendors and customers. This ensures having a higher accuracy level of the model when it was validated from pre-silicon to actual silicon measurements.

Acknowledgements

We would like to thank the ADI design engineers, ADGT test development engineer, and ISO team for all the support to complete this project. Also, we appreciate the support of the ADGT system integration manager for the project sponsorship.

References

[1]Mercedes Casamayor. "AN-715 Application Note—A First Approach to IBIS Models: What They Are and How They Are Generated." Analog Devices, Inc., 2004.

[2]Roy Leventhal and Lynne Green. *Semiconductor Modeling: For Simulating Signal, Power, and Electromagnetic Integrity*. Springer, 2006.

[3]Jermaine Lim and Keith Francisco-Tapan. "IBIS Modeling—Part 1: Why IBIS Modeling Is Critical to the Success of Your Design." *Analog Dialogue*, Vol. 55, No. 3, September 2021.

[4]Roylnd Aquino, Francis Ian Calubag, and Janchris Espinoza. "IBIS Modeling—Part 2: Why and How to Create Your Own IBIS Model." *Analog Dialogue*, Vol. 55, No. 4, October 2021.

[5]Michael Mirmak, John Angulo, Ian Dodd, Lynne Green, Syed Huq, Arpad Muranyi, and Bob Ross. *IBIS Modeling Cookbook for IBIS Version 4.0*. The IBIS Open Forum, September 2005.

[6]"The Basics of Serial Data Compliance and Validation Measurements." Tektronix, March 2010.

IBIS Quality Specification—Revision 1.0. IBIS Quality Committee, November 2004.

I/O Buffer Accuracy Handbook. IBIS Open Forum, April 2000.

Oscilloscope Fundamentals. Tektronix, 2009.

About the Author

Christine C. Bernal joined Analog Devices in June 2007 as a product engineer. She worked on developing IBIS models for various ADI products in 2016 under the New Technology Integration Team. Christine graduated with a Master of Science in electronics and communications engineering M.S.E.C.E., majoring in microelectronics, at Mapua University in Manila in 2015. She can be reached at christine.bernal@analog.com.

About the Author

Janchris Espinoza is a product application engineer at Analog Devices under the New Technology Integration Team. His work focuses on IBIS modeling and simulations for ADI products. He worked as an intern at ADI in 2019 under the Analog Garage team and officially joined ADI in September 2020. He graduated from De La Salle University with a bachelor's degree in electronics engineering in February 2020. He can be reached at janchris.espinoza@analog.com.

About the Author

Aprille Arjhilynne Hernandez-Loyola joined Analog Devices in August 2015. She is a product applications engineer under the New Technology Integration Team. Her work focuses on modeling and circuit simulations, particularly LTspice® and IBIS for ADI products. She graduated from De La Salle University-Dasmariñas with a bachelor's degree in electronics and communications engineering. She can be reached at aprille.hernandez@analog.com.

In the New Era of Wireless Battery Management Systems (wBMS), Security Takes the Spotlight

Lei Poo, Director of System Architecture

The full benefits of wBMS technology can only be achieved if a system security can be assured from process to product.

The challenges identified in early conversations with electric vehicle (EV) OEMs about the technological and business benefits of wireless battery management systems (wBMS) seemed daunting, but the rewards are too promising to ignore. Wireless connectivity's many inherent advantages over wired/cabled architectures have already been proven in countless commercial applications, and BMS was another clear-cut candidate for cord cutting.

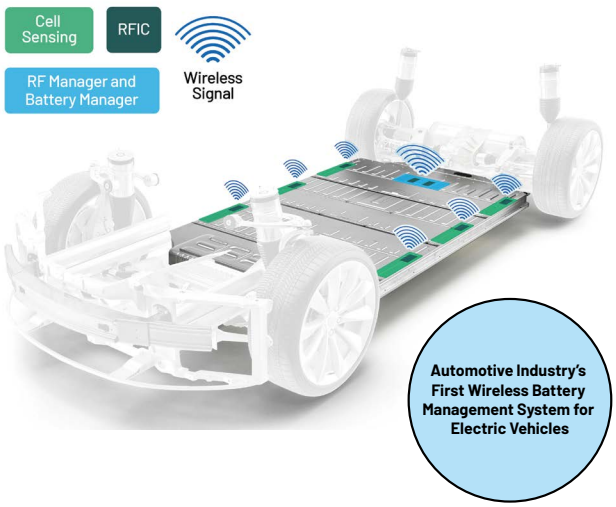

Figure 1. An electric vehicle using a wireless battery management system (wBMS).

The prospect of a more lightweight, modular, and compact EV battery pack—finally liberated from its cumbersome communication wiring harness—has been roundly embraced. By eliminating up to 90% of the pack wiring, and 15% of the pack volume, the entire vehicle's design and footprint can be streamlined significantly, as can its bill of materials (BOM) cost, development complexity, and associated manual installation/maintenance labor.

What's more, a single wireless battery design could be readily scaled across an OEM's entire EV fleet, precluding extensive and costly pack harness redesigns for each make and model. With wBMS, OEMs can freely modify their car frame designs without worrying about having to reroute extensive BMS wiring within the battery pack.

Taking a longer view, continued reductions in vehicle weight and battery pack size will be essential for extending EV driving ranges in the coming years. wBMS technology will therefore play an instrumental part in helping OEMs boost their range capabilities, and in so doing, help overcome consumers' long-lingering EV range anxiety.

This not only promises to spur greater overall EV market adoption, but also it gives OEMs the opportunity to leapfrog into EV market leadership positions on the strength of their driving range claims. This will remain a major differentiator among EV OEMs going forward. Further details about the advantages and market analysis can be found in "Electric Vehicle Wireless Battery Management Revolution Has Begun and the ROI Potential Is Huge."[1]

A New Security Standard

There are numerous challenges that need to be overcome to achieve the promises that wBMS offers. Wireless communications used in wBMS need to be sufficiently robust to interference when the car is driven, and the system must be safe under all conditions. But robust and safe design alone may not suffice against a determined attacker—this is where system security comes into play.

Sources of interference change depending on where the car is driving (for example, city vs. rural area) as well as whether someone is using another wireless device in the car that operates in the same frequency band. Reflections within the battery pack can also degrade performance depending on the material used for the pack housing the battery cells. There is a significant possibility that the

wBMS signal could fluctuate, potentially disrupting communication under natural conditions, let alone in the face of a malicious actor.

If the wBMS communication was somehow interrupted, the car can revert to a "safe mode" with reduced performance to allow the driver to act, or, with complete loss of wBMS communications, come to a safe stop. This can be accomplished by proper safety design, which takes into account all possible failure modes in the system and implements end-to-end safety mechanisms that address random failures of components.

But safety design does not consider the possibility of malicious actors who might exploit the system to their advantage, which may include taking control of the car remotely. This possibility was demonstrated by researchers on a moving vehicle during the 2016 Black Hat conference, using remote access via the vehicle's gateway. Thus, wireless robustness and fail-safe design are not sufficient; they need to be accompanied by security. The Black Hat demonstration was a valuable lesson, showing that future wireless systems in cars need to be designed in such a way that they cannot be leveraged as another remote entry point. In contrast, conventional wired battery packs offer no remote access, so to gain access to battery data, a hacker would need physical access to a high voltage environment in a vehicle.

Additional security challenges can arise throughout the EV battery life cycle, as shown in Figure 2. At Analog Devices, Inc. (ADI), our approach to designing wBMS focuses on understanding the different stages an EV battery goes through from birth to factory, to deployment and maintenance, and finally to either the next life or end-of-life. These use cases define the various functions a wBMS must support. For example, preventing unauthorized remote access is one consideration during EV deployment, but more flexible access is needed during manufacturing. Another example is in serviceability, where right-to-repair laws require a way for car owners to fix issues that stem from the battery cells or the associated wBMS. This means that a legitimate way of updating software in the wBMS must be supported and that the update mechanism should not compromise the safety of the vehicle when it leaves the service center.

Also, EV batteries are sometimes redeployed to the energy sector when they no longer meet EV performance standards. This requires secure ownership transfer of the EV battery from its first life to the next life. Since the batteries are devices without built-in intelligence, it falls upon the accompanying wBMS to enforce the proper security policies that best serve the EV battery life cycle. Secure erasure of first life secrets is needed before transitioning to a second life.

ADI anticipated these concerns and addressed them in accordance with our own core design principles that place a premium value—and exhaustive scrutiny—on maintaining and enhancing security integrities from process to product. In parallel, the ISO/SAE 21434[2] standard on "Road Vehicles: Cybersecurity Engineering" that has been in development over the last three years was officially released in August 2021. It defines a similar exhaustive, end-to-end process framework, with four levels of Cybersecurity Assurance. Automotive OEMs and suppliers are rated on a scale from 1 to 4, with 4 denoting the highest level of conformance (see Figure 3).

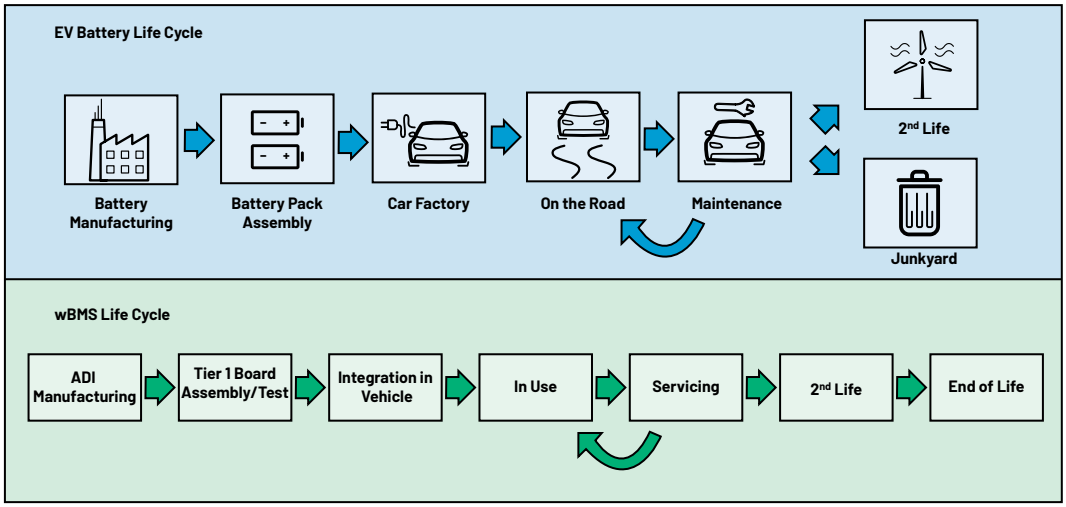

Figure 2. EV battery life cycle and its associated wBMS life cycle.

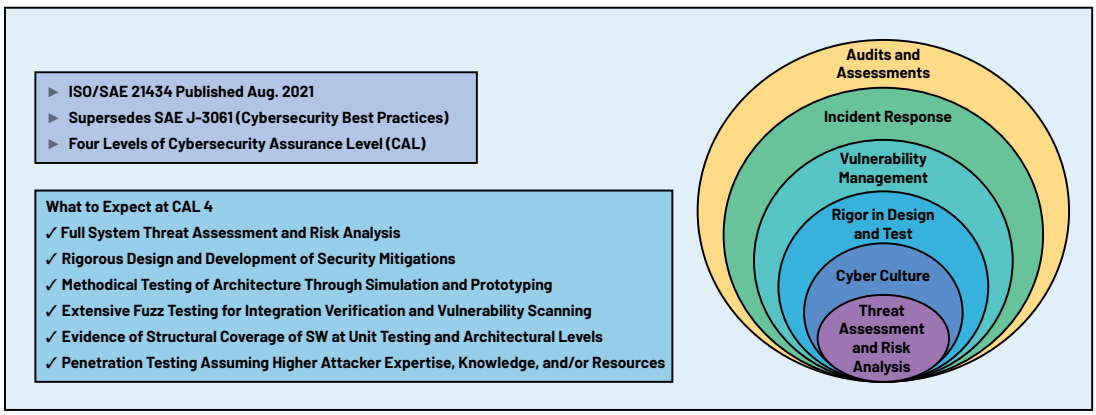

Figure 3. ISO/SAE 21434 framework and CAL 4 expectations.

ADI's approach for wBMS resonates with ISO/SAE 21434 for the highest level of examination and rigor that is needed for secure product development in the automotive industry. To this end, ADI engaged with TÜV-Nord, a well-known trusted certification lab, to assess our internal development policies and processes. This resulted in our policies and processes being vetted to fully comply with the new standard ISO 21434, as shown in Figure 4.

TÜV NORD

ZERTIFIKAT
CERTIFICATE

Hiermit wird bescheinigt, dass die Firma / *This is to certify, that the company*

Analog Devices Inc.
One Analog Way
Wilmington, MA 01887
USA

ein Managementsystem konform zu den Anforderungen der ISO/SAE 21434 betreibt und innerhalb der Laufzeit des Zertifikats von 3 Jahren auf Konformität überwacht wird.
operates a management system conforming to the requirements of ISO / SAE 21434 and within the the duration of the certificate of 3 years is monitored for conformity.

Road vehicles – Cybersecurity engineering
gemäß / *according to* ISO/SAE 21434:2021-08

Geltungsbereich: *Scope of application:* **ADI1381 Development Process for Product Security**
Automotive Business Unit (all locations)

Geprüft nach: *Compliant in accordance with:* **ISO/SAE 21434:2021-08**

Registrier-Nr. / *Registration No.* 44 331 13082139
Prüfbericht Nr. / Test Report No. TNM_8118984937-D06-CSA-R_V1.1
Aktenzeichen / *File reference* 8003030776

Gültigkeit / *Validity*
von / *from* 2021-10-30
bis / *until* 2024-10-29

Zertifizierungsstelle
der TÜV NORD CERT GmbH

Essen, 2021-10-30

TÜV NORD CERT GmbH Am TÜV 1 45307 Essen www.tuev-nord-cert.de

Figure 4. Certificate from TÜV-Nord.

Rigorous Scrutiny from Device to Network

Following our systematic process in the product design of wBMS, a threat assessment and risk analysis (TARA) was conducted to map out the threat landscape based on how the customer intends to use the product. By understanding what the system does, and the various ways it will be used over its lifetime, we can determine what key assets need protection and from what potential threats.

There are several choices of TARA techniques, including the well-known Microsoft STRIDE method, which attempts to model threats by considering the six threats abbreviated by the word STRIDE: Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, and Elevation of privilege. We can then apply this to the different interfaces of the components that make up the wBMS system, as shown in Figure 5. These interfaces are natural stops along the data and control flow paths where a potential attacker may gain unauthorized access to the assets of the system. Here, by role playing the attacker and asking ourselves how applicable each threat is on each interface and why, we can map out possible attack paths and determine how likely the threat is to happen and how severe the consequences may be if it succeeds. We then repeat this thought process across the different life cycle stages as the threat likelihood and impact may vary depending on the environment the product is in (for example, warehouse vs. deployment). This information will point to the need for certain countermeasures.

Take for example the wireless channel between the wireless cell monitors and wBMS manager during deployment, as shown in Figure 5. If the asset is the data from the wireless cell monitors, and the concern is the leakage of the data values to an eavesdropper, then we may want to encrypt the data when it passes through the wireless channel. If our concern is about the data being tampered with as it passes through the channel, then we may want to protect the data with a data integrity mechanism, such as a message integrity code. If the concern is about identifying where the data came from, then we will need a method for authenticating the wireless cell monitors to the wBMS manager.

Going through this exercise allows us to identify the key security objectives of a wBMS system, as shown in Figure 6. These goals will require some mechanisms to be implemented.

Very often, the question of "how far do we go in selecting the mechanisms to achieve a particular security goal?" is asked. If more countermeasures are added, it would almost certainly improve the overall security posture of the product, but at great cost, and possibly add unnecessary inconvenience to the end consumer using the product. A common strategy is to mitigate the most likely threats that are the easiest to deploy. More sophisticated attacks that tend to target assets of higher value will likely require stronger security countermeasures, but these may be extremely unlikely to occur and hence yield low returns if implemented.

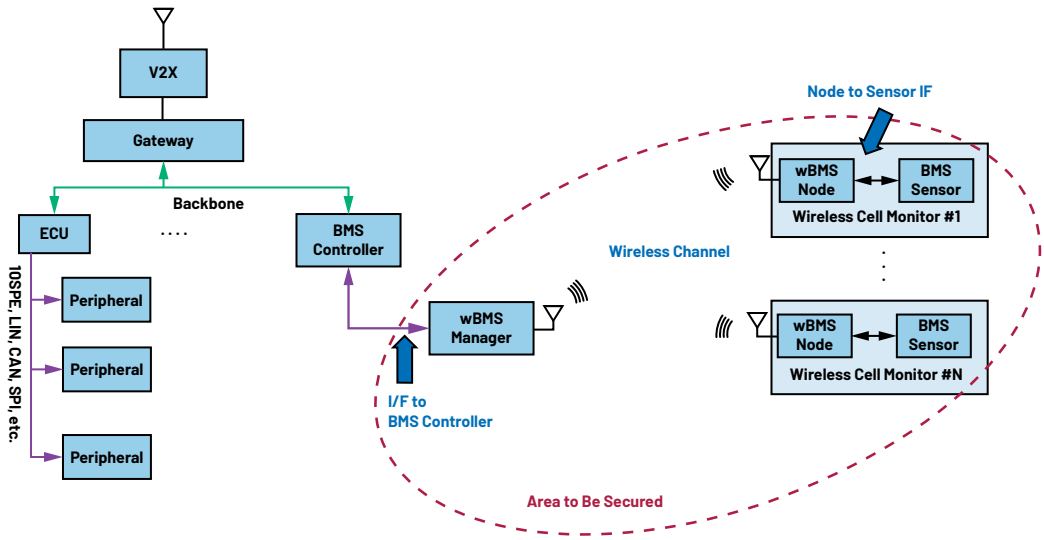

Figure 5. Threat surface considerations for wBMS.

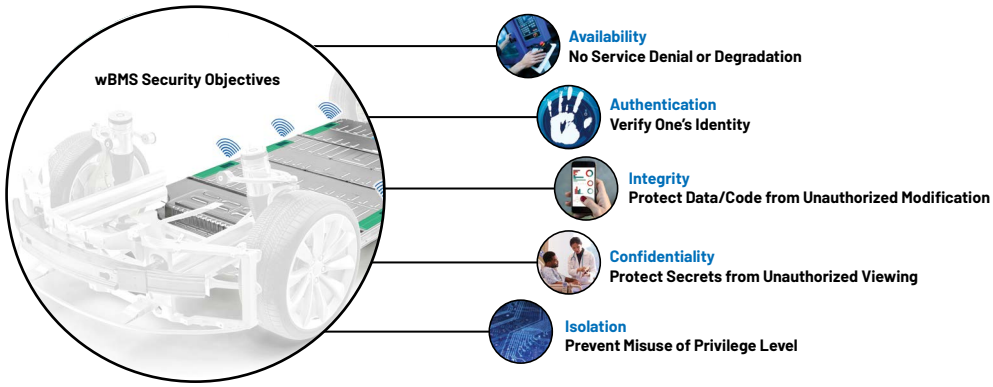

Figure 6. Security objectives of wBMS.

For example, in wBMS, physical tampering of the IC components to gain access to battery data measurements while the vehicle is driving down the road is extremely unlikely, since one would need a well-trained mechanic with deep EV battery knowledge to do gymnastics on the car parts while the car is in motion. An attacker in real life would likely attempt an easier path if one exists. A common type of attack on networked systems is a denial of service (DoS) attack—depriving the user of product utility. You could create a portable wireless jammer to try to interfere with wBMS function (hard), but you can also let the air out of the tires (easy).

This step of reconciling the risks with a set of appropriate mitigations is known as risk analysis. By weighing the impact and likelihood of the relevant threats before and after suitable countermeasures have been introduced, we can determine if the residual risks have been reasonably minimized. The end result is one where security features are incorporated only because they are needed and at a cost level acceptable to the customer.

The TARA for wBMS points to two important aspects of wBMS security: device-level security and wireless network security.

The first rule of any secure system is "keep your keys secret!" This means both on the devices and in our worldwide manufacturing operations. ADI's wBMS device security takes into account the hardware, ICs, and low level software on the ICs, and ensures that the system is able to boot securely from immutable memory into a trusted platform for code to run. All software code is authenticated prior to execution, and any in-field software update requires authorization by pre-installed credentials. Rollback to a previous (and possibly vulnerable) version of the software is prohibited after the system is deployed in the vehicle. Additionally, debug ports are locked once the system is deployed, thereby eliminating the possibility of unauthorized backdoor access into the system.

Network security is targeted at protecting the over-the-air communication between a wBMS cell monitor node and the network manager within a battery pack enclosure. Security starts at network joining where membership is checked for all of the participating nodes. This prevents random nodes from joining the network even if they happen to be physically close by. Mutual authentication of the nodes to the network manager at the application layer will further secure the wireless communication channel, making it impossible for a man-in-the-middle attacker to masquerade as a legitimate node to the manager or vice versa. Furthermore, to ensure that only the intended recipient can access the data, AES-based encryption is used to scramble the data cryptographically, preventing information leakage to any potential eavesdroppers.

Securing the Keys

As with all secure systems, the heart of security is a set of cryptographic algorithms and keys. ADI's wBMS follow NIST-approved guidelines, which means choosing algorithms and key sizes that align to a minimum security strength of 128 bits that are suitable for data-at-rest protection (for example, AES-128, SHA-256, EC-256) and uses algorithms found in well-tested wireless communications standards such as IEEE 802.15.4.

The keys used in device security are typically installed during ADI's manufacturing process and never leave the IC devices. These keys, which are used to ensure system security, are in turn protected physically by the IC devices both in use and at rest from unauthorized access. A hierarchical key framework then protects all application-level keys by saving them as encrypted blobs in nonvolatile memory, including the ones used in network security.

To facilitate mutual authentication of the nodes in the network, ADI's wBMS has provisioned a unique public-private key pair and a signed public-key certificate into each wBMS node during manufacturing. The signed certificate allows a node to verify that it is talking to another legitimate ADI node and valid network member, while the unique public-private key pair is used by the node in a key agreement scheme to establish a secure communication channel with another node or with the BMS controller. One benefit of this approach is easier wBMS installation without needing a secure installation environment, as the nodes are programmed to automatically handle network security after deployment.

In contrast, past schemes that use pre-shared keys to establish secure channels often required a secure installation environment and an installer to manually program the key value for the communication end points. To simplify and lower the cost of handling the key distribution problem, assigning a default common network key for all nodes in the network was often the shortcut many took. This often resulted in a hard lesson learned when a break-one, break-all disaster occurred.

As the OEM production scales, having the ability to leverage the same wBMS with varying number of wireless nodes across different EV platforms and to install at

different manufacturing or servicing sites that are necessarily secure, we lean in favor of the distributed key methodology that simplifies the overall key management complexity.

Conclusion

The full benefits of wBMS technology can only be achieved if security can be assured from device to network, and over the lifetime of the EV battery. Security, in this light, requires a system-level design philosophy, encompassing both process and product.

ADI anticipated the core cybersecurity concerns addressed by the ISO/SAE 21434 standard during its draft period and embraced them within our own wBMS design and development ethos. We are proud to be one of the first technology vendors to achieve ISO/SAE 21434 compliance on our policies and processes, and are currently undergoing certification for wBMS technology to the highest Cybersecurity Assurance Level.

References

[1]Shane O'Mahony. "Electric Vehicle Wireless Battery Management Revolution Has Begun and the ROI Potential Is Huge." Analog Devices, Inc., November 2021.

[2]ISO/SAE 21434:2021 - Road Vehicles. ISO, 2021.

About the Author

Lei Poo is director of system architecture in the E-Mobility Group within the Automotive Business Unit at Analog Devices and currently manages the Systems Architecture Team, which is responsible for designing wireless battery management systems (wBMS). She previously led ADI's Security Architecture and Platforms Team to establish the internal, secure product development process and now builds HW embedded security into ADI's emerging silicon products for Industrial Ethernet and wBMS. Prior to joining ADI, Lei was with NXP, Broadcom, and Marvell, where she was an embedded systems and security architect who designed secure chip/controller solutions for smartcards/smartphones, set-top boxes, and secure disk drives. Lei received a Ph.D. in electrical engineering from Stanford University in 2005 and holds 20 U.S. patents in the areas of HW embedded security, systems, and algorithms. She can be reached at lei.poo@analog.com.

Honey, Where Is My Power Cord?

Thong Huynh, Application Engineering Director

Introduction

A more common modern-day question would probably be "Honey, where is my battery charger?" The reduction in cost and improvement in battery performance, especially Li-Ion based, at the turn of the century has fueled a steady growth of battery-powered energy storage and portable equipment. Also, supercapacitors (aka ultracapacitors) are increasingly finding usage in a variety of applications due to their unique characteristics. The lead-acid battery, a 150-year-old technology, is still popularly used in cars, wheelchairs, scooters, golf carts, and uninterruptable power supply (UPS) systems. These energy storage devices must be recharged once their energy has been depleted. The worldwide charging IC shipment was 1.16 billion units in 2019 and is expected to grow to 1.72 billion units in 2024 with a healthy 8.6% annual growth rate. The respective revenue was $518.1 billion and $735.4 billion with a 7.3% CAGR. Figure 1 shows this trend, according to OMDIA's "Power IC Market Tracker - 2019."[1]

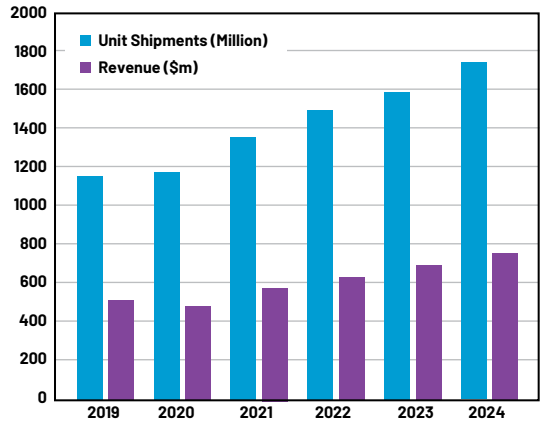

Figure 1. The world market for charging ICs.

The demand for more power, along with longer range or run time, dictates an increase in the voltages used in energy storage devices. For example, Li-Ion battery stacks have gone from one or two cells to multiple (up to 12) cells used in robots, drones, power tools, and a host of other things. A 12-cell Li-Ion battery stack provides 50.4 V maximum voltage. A 12-cell battery would last 12 times longer than a 1-cell battery at the same current rating. Alternatively, 12 batteries can be connected in parallel for higher power, but this method would increase the current 12×. Higher current causes more conduction losses, so paralleling batteries is not preferred.

Industrial systems such as emergency lighting with battery backup, UPS backup power, and HVAC use a 24 V_{DC} power source—that is, a 24 V battery is used to back up these systems. The 24 V_{DC} power source, however, can rise to 60 V peak voltage during transient conditions, according to IEC 61131-2 and IEC 60664-1 standards.

In either situation, the equipment requires charger solutions that can accommodate higher battery voltage and withstand higher input voltage during transient events.

Charger Basics

There are many charger topologies. The linear charger drops the voltage difference between the power source and the battery through a power switch. This type of charger is the least efficient, since it dissipates a lot of power across the power switch when the voltage difference between the power source and the battery is large. The boost charger boosts the voltage from the power source to the battery voltage. This topology requires the power source voltage to be lower than the battery voltage. The buck charger steps down the voltage from the power source and requires the power source voltage to be higher than the battery voltage. The buck-boost charger can charge the battery with a power source voltage that is either higher or lower voltage than the battery voltage. This topology requires four power switches (compared to two for the buck) and generally is not as efficient.

The synchronous rectification buck charger is the most efficient and is the focus of this article. Figure 2 shows a generic synchronous rectification buck charger circuit. Most buck chargers today operate at a relatively low voltage. Many are rated at only 28 V input with some at 40 V. Allowing ±10% input voltage regulation and a 2 V drop across the buck charger, a 28 V-rated charger can only practically charge a 5S Li-Ion battery stack (maximum). We will examine a new family of 60 V input charger ICs that allow higher voltage charging—up to 52 V battery voltage (or a 12-cell Li-Ion stack)—and that can withstand a 65 V input voltage transient.

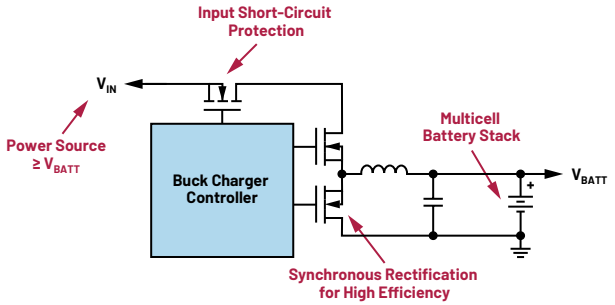

Figure 2. Generic synchronous rectification buck charger.

The standby current on a charger should be low to save energy. Energy Star® assigns five stars to mobile phone chargers and other small chargers that draw 30 mW or less on standby. One star goes to chargers with 300 mW or more, and there are other ratings for everything else in between. Energy Star aims to reduce current consumption of personal chargers that are mostly left plugged in when not in use. There are over 1 billion such chargers connected to the grid globally at any given time.

Even though the lead-acid battery, Li-Ion-based battery, and supercapacitor are all energy storage devices, they have very distinct charging/discharging characteristics. We will examine these characteristics and discuss a charging solution for each of them. A good battery charger provides battery performance and durability, especially when charging under adverse conditions.

Lead-Acid Battery Charger

Lead-acid is the oldest rechargeable battery in existence and was invented by the French physician Gaston Planté in 1859.[2] One hundred and fifty years later, it is still popularly used in cars, wheelchairs, scooters, electric bikes, golf carts, and UPS systems.

The lead-acid battery must be charged slowly. Typical charge time is 8 to 16 hours. The battery must always be stored in a charged state, and a periodic fully saturated charge is essential to prevent sulfation. It is common practice to charge lead-acid batteries to 70% in about 8 hours, and another 8 hours to do the all-important absorption charge. A partial charge is fine provided the lead-acid occasionally receives a fully saturated charge to prevent sulfation. Leaving the battery on float charge for a prolonged time does not cause damage.

Finding the ideal charge voltage limit is critical. A high voltage (above 2.45 V/cell) produces good battery performance but shortens service life due to grid corrosion on the positive plate. A low voltage limit is subject to sulfation on the negative plate. Temperature also affects the cell voltage with a typical –5 mV/°C (0.028 V per cell for every 10°F).[3] A good charger must compensate for this temperature coefficient to avoid overcharge of the battery when hot or undercharge when cold.

As an example, the MAX17702 (see Figure 3) is a complete lead-acid battery charger controller designed to operate over an input voltage range of 4.5 V to 60 V. The device offers a high efficiency (over 97%), high voltage, synchronous buck solution to charge 12 V/24 V/48 V lead-acid battery stacks. Figures 4a and 4b show its charging cycle and charging efficiency.

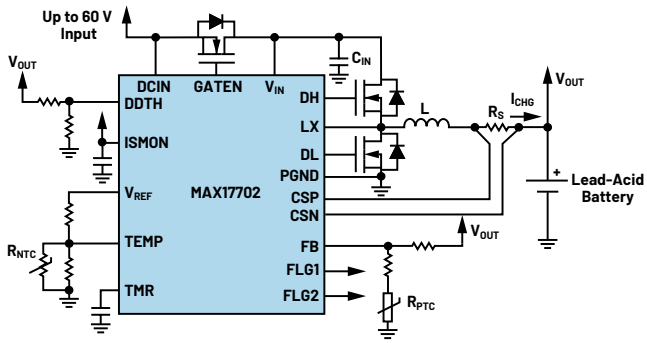

Figure 3. High voltage lead-acid battery charger controller.

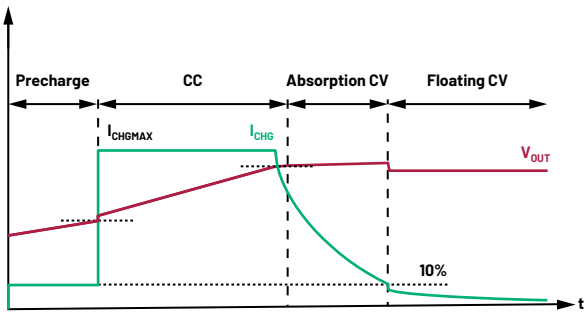

Figure 4a. MAX17702 lead-acid charging cycle.

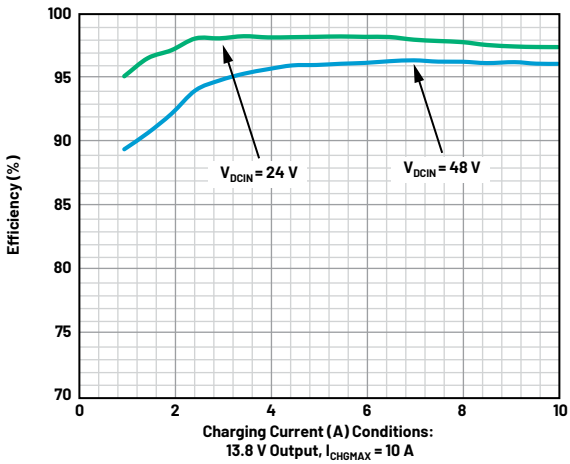

Figure 4b. MAX17702 charging efficiency.

The lead-acid battery has low energy densities, making it unsuitable for portable devices. This is where a lithium-based battery comes into play.

Li-Ion Battery Charger

Li-Ion is the universally accepted battery for portable applications, heavy industries, electric powertrains, and satellites due to its light weight and high energy density.

Li-Ion is a low maintenance battery. The battery has no memory and does not need exercising (deliberate full discharge) to keep it in good shape. But it needs protection circuits, both built-in inside the battery pack as well as in the charger to prevent short circuit, overcharge, thermal runaway, and overdischarge. If a Li-Ion battery has dwelled below 1.5 V/cell for a week or longer, dendrites may have developed that could compromise safety.

To prevent overdischarge, the built-in battery protection circuit puts the battery into a sleep condition. This happens when storing the battery in a discharged state in which self-discharge brings the voltage to the cutoff point. A regular charger treats such a battery as unserviceable, and the pack is often discarded. An advanced Li-Ion charger includes a wake-up feature, or "precharge," to allow recharging if a Li-Ion battery has fallen asleep due to overdischarge. In precharge mode, the charger applies a small charge current to safely raise the voltage to between 2.2 V/cell and 2.9 V/cell to activate the protection circuit, at which point a normal charge commences.

During normal charge, the Li-Ion charger operates on constant current constant voltage (CCCV). The charge current is constant, and the voltage is capped when it reaches a set limit. Reaching the voltage limit, the battery saturates; the current drops until the battery can no longer accept further charge and charging terminates. Each battery has its own low current threshold.

Li-Ion batteries should always stay cool on charge. Li-Ion cannot absorb overcharge. Thus, it is very important to monitor the battery temperature and its charging voltage to assure battery health and safety. A good charger must include these features.

Figure 5 shows an example of an advanced Li-Ion battery charger. The MAX17703 is a high efficiency, high voltage, synchronous, step-down charger controller designed to operate over a wide input voltage range of 4.5 V to 60 V. The device offers a complete charging solution for up to 12 Li-Ion cell stacks.

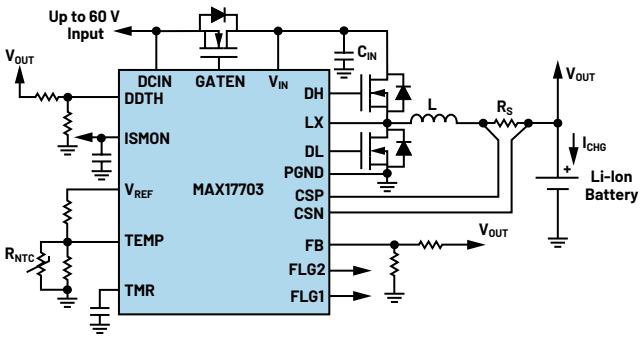

Figure 5. Advanced, high voltage Li-Ion battery charger circuit.

The device offers accurate CCCV charging current/voltage at ±4% and ±1%, respectively. The charger enters a top-up-charge state when the charging current reduces to the taper-current threshold and then exits charging after a taper-timer period elapse. The charger initiates a recharge cycle when the output voltage falls below the recharge threshold voltage. This is a nice feature to keep the battery fully charged, if left in the charging cradle for a long period, without using too much power and to comply with Energy Star requirements. The device can detect and precondition deeply discharged batteries, waking them up with the precharge feature. For added protection, the device senses the battery temperature and allows charging only when within the temperature range. There is also an input short-circuit protection feature, which prevents discharging of the battery when the input is accidentally short circuited. Figure 6 illustrates the MAX17703's charging cycle.

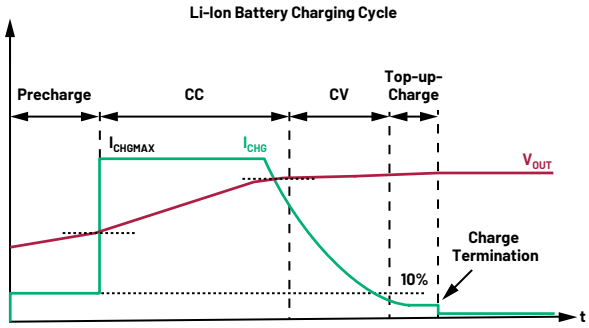

Figure 6. MAX17703 Li-Ion battery charging cycle.

Supercapacitor Charger

Supercapacitors are increasingly finding usage in a variety of applications, thanks to their unique advantages over batteries. Supercapacitors function on electrostatic principles with no chemical reactions, avoiding the lifetime issues associated with chemical storage of batteries. Their high durability allows for millions of charge/discharge cycles with lifetimes up to 20 years, one order of magnitude above batteries. Their low impedance enables fast charge and discharge in a matter of seconds. This, in conjunction with their moderate ability to hold charge over long periods of time, makes supercapacitors ideal for applications requiring short charge and discharge cycles. They are also used in parallel with batteries, in applications where instantaneous peaks of power delivery are necessary during load transitions.

Supercapacitor short-charge and discharge cycles require chargers to handle high currents and work smoothly in constant current (CC) mode during a charge, which may start at 0 V, and in constant voltage (CV) mode once the final output value is achieved. In high voltage applications, many supercapacitors are connected in series, requiring chargers to manage high input and output voltage.

The MAX17701 (see Figure 7) is a high efficiency, high voltage, synchronous, step-down supercapacitor charger controller designed for high current charging and operates over an input voltage range (V_{DCIN}) of 4.5 V to 60 V. The output voltage is programmable from 1.25 V up to (V_{DCIN} – 4 V). The device uses an external N-MOSFET to provide an input supply-side OR'ing function, preventing supercapacitor discharge back to the input. Figure 8 illustrates the simplistic, but high current charging profile.

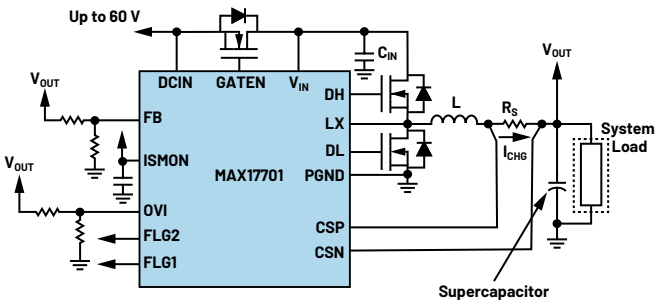

Figure 7. Supercapacitor high voltage, high current charger.

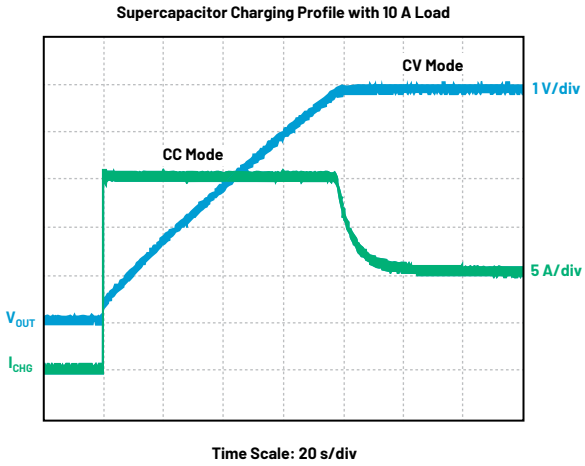

Figure 8. MAX17701 supercapacitor charging profile.

Conclusion

The use of battery-powered energy storage and portable equipment has grown steadily. The demand for more power, along with longer range or run time, dictates an increase in the voltages used in their battery stacks. Applications in industrial systems using a 24 V_{DC} power supply can see 60 V peak voltage during transient conditions. Legacy charger solutions are mostly limited to 28 V input. Newer charger solutions from Analog Devices enable higher battery stack voltage and higher charging efficiency, thanks to the high voltage, synchronous buck charging topology.

Lead-acid batteries, lithium-based batteries, and supercapacitors are all energy storage devices that have very distinct charging/discharging characteristics and require a dedicated charger for an optimum charging solution. An advanced battery charger also provides adequate protection to give battery performance and durability, especially when charging under adverse conditions. These are also addressed in newer charger solutions.

References

[1]Kevin Anderson. "Power IC Market Tracker - 2019." OMDIA, September 2020.

[2]"Advancements in Lead Acid." Battery University, July 2016.

[3]US 31DC XC2 - Data Sheet. U.S. Battery, 2019.

About the Author

Anthony T. Huynh (aka Thong Anthony Huynh) was a principal member of technical staff (MTS), applications engineering, at Maxim Integrated (now part of Analog Devices). He has more than 20 years of experience designing and defining isolated and nonisolated switching power supplies and power management products. At ADI, he has defined more than 100 power management products including DC-to-DC converters, hot swap controllers, Power over Ethernet, and various system-protection ICs adopted by the world's leading manufacturers.

Anthony holds four U.S. patents in power electronics and has written several public articles and application notes in this area. He has a B.S. in electrical engineering from Oregon State University and has completed all coursework for an M.S. in electrical engineering at Portland State University, where he also taught a power electronics class as an adjunct instructor.

RAQ Issue 198: Why AMR Sensors Are a Great Option for High Precision Position Measurements!

Benjamin Reiss, Field Applications Engineer

Question:

What is the best type of sensor to measure the position of moving and rotating elements?

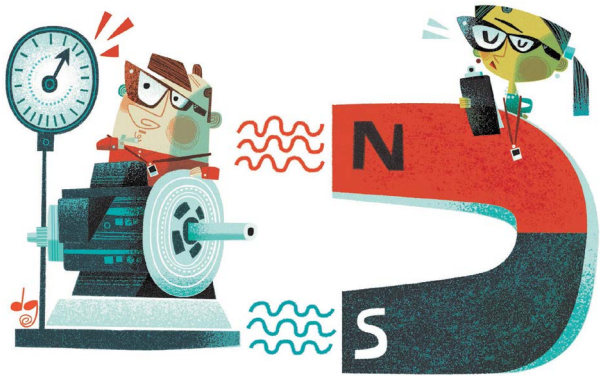

Answer:

Try an anisotropic magnetoresistive (AMR) sensor, such as the ADA4570.

Nowadays, there are many different ways to measure the position of moving and rotating elements. Apart from optical encoders, Hall sensors, and resolvers, magnetic sensors based on the magnetoresistive effect are used. There is a giant magnetoresistive (GMR) and a tunnel magnetoresistive (TMR) effect as well as an anisotropic magnetoresistive (AMR) effect, which this article delves further into.

The ADA4570 AMR sensor from Analog Devices makes use of the characteristic of ferromagnetic materials in which the electrical resistance depends on the direction of magnetization—a phenomenon that was discovered by William Thomson (Lord Kelvin) around 1851:

$$R = R0 + \Delta R \times \cos(2 \times \alpha) \quad (1)$$

Where α is the angle between the direction of magnetization and the direction of current flow.

For an optimal sensor response in linear displacement measurements, the sensor is placed such that the magnet and the sensor are in the same plane and the center of the magnet is in line with the center of the sensor. Because an AMR sensor cannot distinguish between north and south poles, the position of the magnet cannot be changed.

For rotating elements, so-called off-shaft or end-of-shaft configurations are common. In the example with the off-shaft configuration (see Figure 1), the sensor sine/cosine outputs repeat the absolute information for every pole—for example, 45° for a 4-pole pair magnet.

In the end-of-shaft configuration (see Figure 1), the sensor is situated below a rotating dipole magnet; here, the north and south poles form a uniform field above the center of the magnet. The sensor is positioned in such a way that the magnetic field and the element to be measured are in the same plane. A typical application is rotor position measurement and control in brushless DC motors. In the case of AMR sensors, which are characteristically 180° angle sensors, the motor must be an even pole pair-motor; motors with an odd number of pole pairs require 360° information for commutation. Compared with conventional Hall sensors, which are also used in motor control, AMR sensors such as the ADA4570 and the ADA4571 have a higher precision. They also reduce the torque ripple and provide true power on absolute position information after startup or an idle state irrespective of the motor position.

The AMR technology from Analog Devices measures the angle by means of two Wheatstone bridges with one rotated by 45° with respect to the other (see Figure 2).

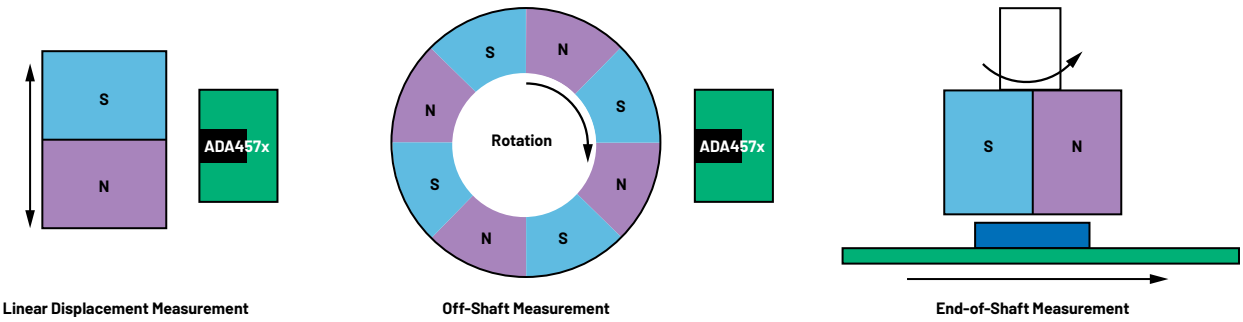

Figure 1. Linear displacement measurement (left), off-shaft measurement (middle), and end-of-shaft measurement (right).

The angle is calculated via sine and cosine functions and represents the orientation from 0° to 180° in relation to the sensor (ADA4570).

$$\alpha = \arctan2\ (V_{SIN} / V_{COS}) / 2 \quad (2)$$

In AMR sensors, differentiation is made between electrical angles and mechanical angles. Due to the working principle of AMR sensors and the above-described 45° angle between the Wheatstone bridges, the absolute angle can be measured by means of Equation 2 over 180° mechanical rotation. The electrical period repeats twice over the 360° rotation for a dipole magnet. Given that AMR sensors work in saturation, the absolute field strength is irrelevant with a certain minimum magnetic field strength present allowing for robust systems when working with strong magnets.

Apart from optical sensors, Hall sensors, and resolvers, magnetic sensors provide a further elegant solution for measuring position with high precision and robustness in many different applications. Analog Devices offers a number of possibilities for this, such as the ADA4570, ADA4571, and ADA4571-2 where redundancy is needed. To find out more about this topic, please see the references listed at the end of this article.

References

Guyol, Robert. "AN-1314 Application Note: AMR Angle Sensors." Analog Devices, Inc., October 2014.

Nicholl, Enda. "Dual AMR Motor Position Sensor for Safety Critical Applications." *Analog Dialogue*, Vol. 53, No. 4, November 2019.

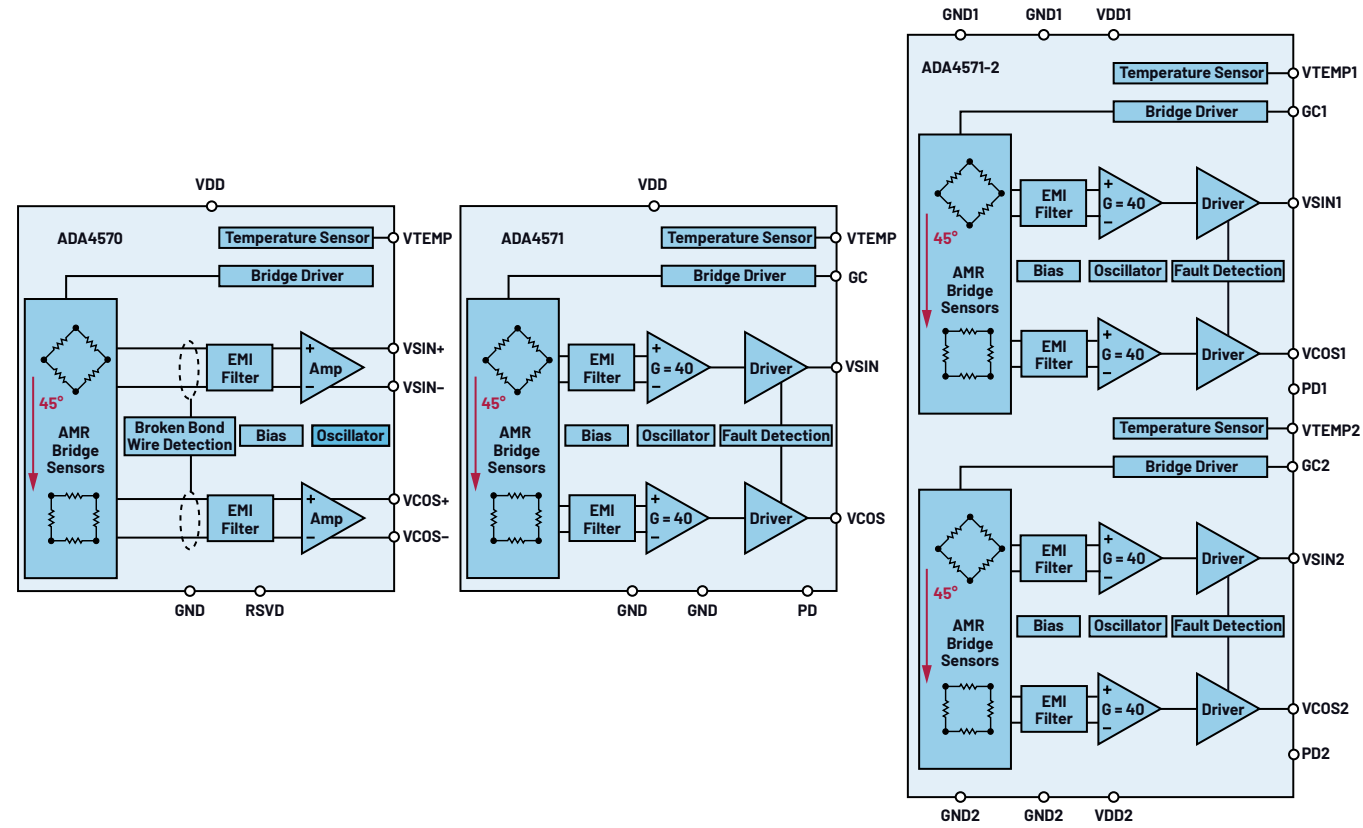

Figure 2. ADA4570, ADA4571, and ADA4571-2 AMR configurations with Wheatstone bridges at a 45° angle to each other.

About the Author

Benjamin Reiss has been working at Analog Devices in Munich, Germany since April 2017. He graduated in 2016 from the Friedrich-Alexander University in Erlangen with a master's degree in nanotechnology. After completion of the trainee program at Analog Devices, he joined the regional team as a field applications engineer supporting several broad market accounts. He can be reached at benjamin.reiss@analog.com.

Transceiver with Scalable Power and Performance: A Solution to Mission Critical Communications

Michelle Tan, Product Applications Engineer

Abstract

This article discusses the ADRV9001, Analog Devices' newest generation software-defined radio (SDR) transceiver monolithic integrated circuit (IC) designed to provide scalable power and performance for many satellite, military, land mobile, utility infrastructure, and cellular mission critical communications. It first introduces three user-defined power saving options within the ADRV9001 from the component level, channel level, and IC system level. Then it further discusses a unique system feature called monitor mode, which not only saves power for the ADRV9001 itself but also enables the user to reduce the baseband integrated circuit (BBIC) operating power so that the best overall system power saving target can be achieved. This article also demonstrates power consumption savings for each power saving option and provides a detailed explanation of the associated performance trade-offs. By thoroughly understanding those trade-offs, the best system power saving strategy can be determined to achieve the optimal power consumption with satisfactory system performance.

Introduction

The ADRV9001 is part of a highly agile, user-configurable, new generation SDR IC transceiver family. It provides state-of-the-art RF performance with a set of advanced system features such as multichip synchronization (MCS), digital predistortion (DPD), dynamic profile switching (DPS), and fast frequency hopping (FFH). The IC supports both frequency division duplex (FDD) and time division duplex (TDD) operations with a wide range of RF frequency from 30 MHz to 6 GHz, which covers the ultrahigh frequency (UHF); very high frequency (VHF); industrial, scientific, and medical (ISM); and cellular frequency bands. It can handle both narrow-band (down to 12 kHz) and wideband (up to 40 MHz) signals with an almost continuous sample rate from 24 kSPS to 61.44 MSPS.

Empowered with all these capabilities, it is ideally suited as a platform for use across many different mission critical applications. Several general-purpose system on modules (SOMs) have already been developed from ADI partners, such as Alciom, Epiq Solutions, NextGen RF Design, and Vanteon Wireless Solutions. The SOM products target mission critical communications, including industrial automation and advanced metering applications. They share the same attributes of achieving a perfect balance between performance, power, size, and cost made possible by this IC. Figure 1 showcases the major power saving options provided at the component, channel, and system levels. Note: different transceiver variants in the ADRV9001 family could have different numbers of channels and different system features, which are simplified in Figure 1.

As shown in Figure 1, the component-level power saving options, highlighted in purple, mainly involve components such as the analog-to-digital converter (ADC), RF PLL, baseband (BB) PLL, analog transmit low-pass filter (Tx LPF), and receive low-pass filter (Rx LPF). Unlike most traditional transceivers, the ADRV9001 provides a pair of high performance (HP) and low power (LP) ADCs for both I and Q datapaths that users can select. For each component, multiple power saving options are provided. The channel-level power saving options are highlighted in red for a pair of transmit and receive channels. This is designed specifically for TDD applications, in which the transmit and receive operations are time multiplexing with each other. Therefore, while one channel is operating, the other channel is idle, which could be powered down. Different levels of channel power saving schemes are provided by requiring different wake-up times to resume operation. The system-level power saving options are highlighted in green; these can be employed to achieve more power savings for some applications expecting longer periods of inactivity, such as digital mobile radio (DMR) handset systems.[1]

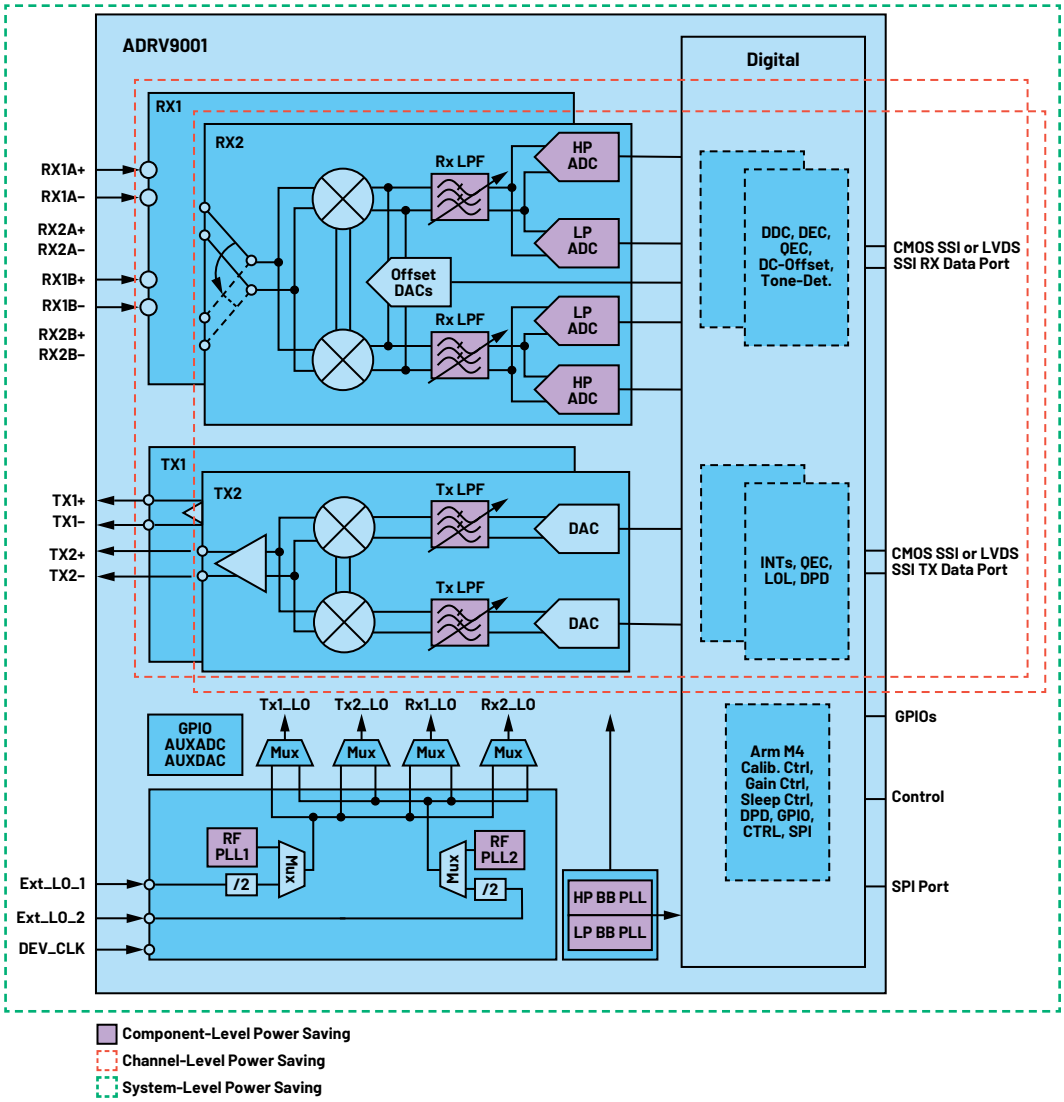

Figure 1. A high level diagram of the ADRV9001 power saving options at three different levels.

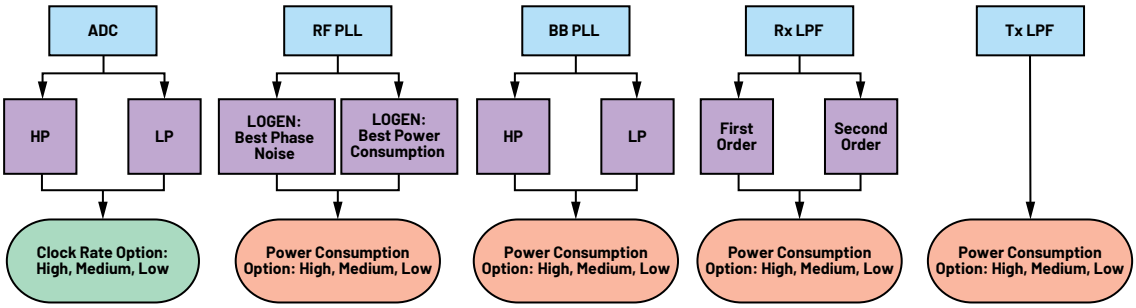

Figure 2. ADRV9001 component-level power saving options.

Besides all those power saving options, the transceiver features a monitor mode that allows both the ADRV9001 and the BBIC to go to sleep during the system idle time period. During sleep, the ADRV9001 can periodically wake up one receive channel to perform signal detection. Therefore, it could offload signal detection responsibility from the BBIC and allow it to sleep through the entire idle time period to achieve the best overall system power saving target.

In the following sections, all the power saving options and the monitor mode will be discussed in depth. By thoroughly understanding the associated performance trade-offs, a design engineer can explore all the potential power saving possibilities to keep the power consumption under control while guaranteeing a satisfactory system performance.

Component-Level Power Saving

Component-level power saving can be easily achieved by configuring individual hardware components through application programming interfaces (APIs) provided by the software development kit (SDK) during the device initialization stage. Figure 2 presents the major hardware components that offer multiple power saving options, including the ADC, RF PLL, BB PLL, receive LPF, and transmit LPF. To properly configure those components, it is crucial to understand the performance trade-offs.

The ADRV9001 provides an option to select between the HP ADC and the LP ADC. The HP ADC is based on continuous-time sigma-delta (CTSD) architecture and is 5 bits wide. The LP ADC is based on voltage-controlled oscillator (VCO) architecture and is 16 bits wide. The HP and LP ADCs provide a similar performance of dynamic range (full scale to thermal noise) but a different performance in linearity.[2] Figure 3 compares the input third-order intercept point (IIP3) and input second-order intercept point (IIP2) performance of the HP ADC and the LP ADC. It is measured with two continuous wave (CW) tones (with 1 MHz frequency spacing) using a wideband profile under room temperature and maximum receiver gain. Note: the x-axis stands for the baseband frequency for the first tone (lower frequency), and the second tone frequency is 1 MHz higher than the first tone.

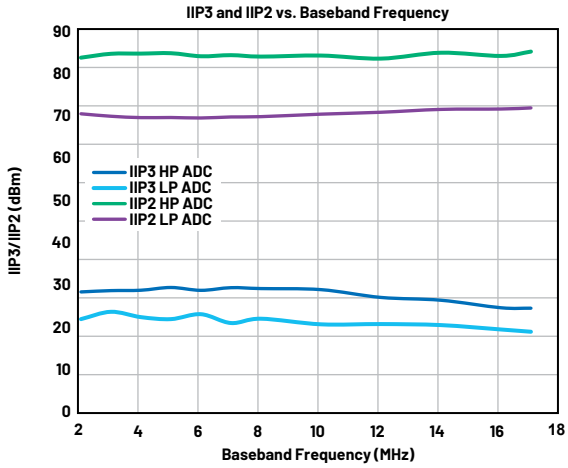

Figure 3. ADRV9001 linearity performance comparison between its HP ADC and its LP ADC.

As shown in Figure 3, both the HP ADC and the LP ADC demonstrate a high linearity performance. The HP ADC can achieve about 12 dB better performance with IIP2 and 6 dB better performance with IIP3 than the LP ADC by consuming more power. For both HP ADC and LP ADC, the user could further choose a high, medium, or low ADC sample rate. Choosing a higher sample rate improves the noise performance; in addition, it mitigates the requirement on transition band sharpness in the antialiasing filter design at the expense of consuming more power to process data at a faster rate.

The transceiver contains two RF PLLs, each driving its own local oscillator (LO) generator. Two options of LO generators are provided to achieve the best phase noise performance or the best power consumption performance. The best power consumption mode consumes less power by slightly sacrificing the phase noise performance. Note: the best phase noise performance option is only available for an LO frequency less than 1 GHz. For each mode, three different power consumption options are provided with different LO output swing levels. A higher swing level results in better phase noise performance but consumes more power.

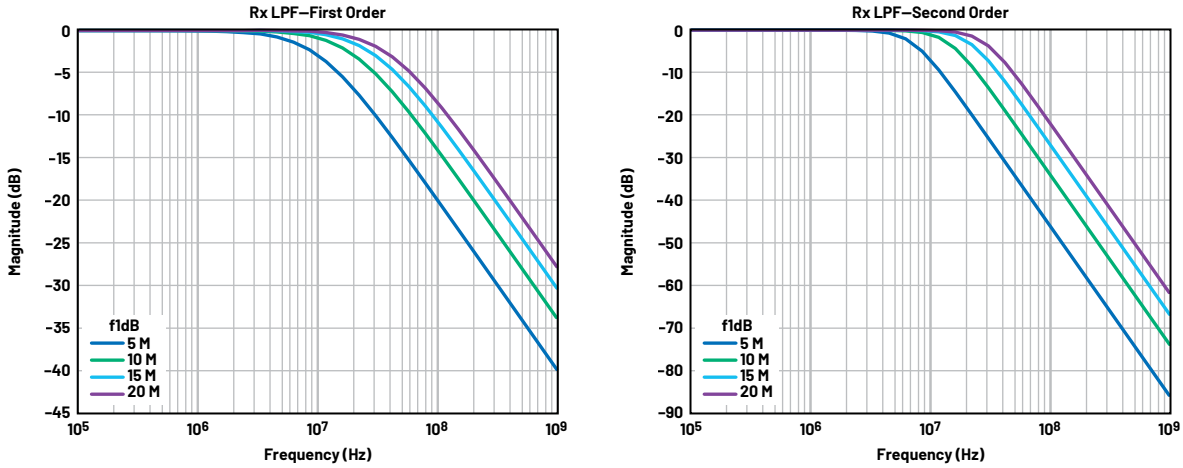

Figure 4. First-order and second-order Rx LPF frequency response at different LPF f1dB configurations.

The BB PLL generates all baseband and data port related clocks. Similar to the ADC, both HP BB PLL and LP BB PLL options are provided. The HP BB PLL has a programmable frequency range of 7.2 GHz to 8.8 GHz, while the LP BB PLL has a programmable range of 3.3 GHz to 5 GHz. The HP BB PLL offers greater flexibility in generating clocks to support a wider range of sample rates. When the signal sample rate is greater than 53.33 MHz, the HP CLK PLL must be used. The LP BB PLL has a limitation in supporting certain sample rates but consumes less power.

The receive LPF attenuates out-of-band signals by supporting a variable bandwidth from 5 MHz to 50 MHz. It also converts the baseband signal current to voltage. It could be configured in transimpedance amplifier (TIA) mode as a first-order single-pole filter, or in bi-quad (BIQ) mode as a second-order filter with two complex poles in the transfer function. While the in-band performance of both modes is similar, the second-order BIQ mode achieves additional out-of-band attenuation compared with the first-order TIA mode. Figure 4 compares the simulated frequency response at different f1dB configurations between these two filters. The selection of the second-order LPF consumes more power than the first-order mode. In addition to that, the in-band noise of the second-order LPF is around 2.5 dB higher than the first-order LPF. For both first-order and second-order modes, the user can further select three different power consumption levels at high, medium, or low by sacrificing the noise and linearity performance.

The transmit LPF is a second-order Butterworth filter used to attenuate the sampling images of the digital-to-analog converter (DAC). It also converts the current from the DAC to a voltage and reconstructs the analog spectrum by low-pass filtering the output. Similar to the receive LPF, it provides three options of power consumption levels at high, medium, or low at a cost of linearity performance.

Usually, the best performance could be achieved by configuring all the components at its highest power consumption option. For an FDD 1T1R LTE 20 MHz profile, by configuring the highest power consumption option when both transmit and receive channels are active, the total power consumption of the ADRV9001 is measured at about 1800 mW. Note: even with the same configurations, measurement results might vary depending on hardware and temperature. Table 1 presents the amount of power saving achieved through configuring different power saving options. In this 1T1R LTE 20 MHz profile, both Receive Channel 1 and Transmit Channel 1 are enabled and the LO is configured at 900 MHz. Note: the numbers in each row of Table 1 show the relative amount of power saving in mW by only enabling this single power saving option. For example, using HP ADC with Medium Clock Rate saves about 72 mW compared against the highest power consumption around 1800 mW with all the highest power consumption options enabled.

Table 1. ADRV9001 Component-Level Power Saving Measurement

ADC		Power Saving (mW)
HP	Medium Clock Rate	–72
	Low Clock Rate	–41
LP	High Clock Rate	–100
	Medium Clock Rate	–177
	Low Clock Rate	–158
RF PLL		
Best Phase Noise	Medium Power Consumption	–44
	Low Power Consumption	–84
Best Power Consumption	High Power Consumption	–50
	Medium Power Consumption	–80
	Low Power Consumption	–108
BB PLL		
HP	Medium Power Consumption	–5
	Low Power Consumption	–10
LP	High Power Consumption	–45
	Medium Power Consumption	–47
	Low Power Consumption	–49
Rx LPF		
Second Order	Medium Power Consumption	–26
	Low Power Consumption	–40
First Order	High Power Consumption	–77
	Medium Power Consumption	–101
	Low Power Consumption	–116
Tx LPF		
Medium Power Consumption		–29
Low Power Consumption		–47

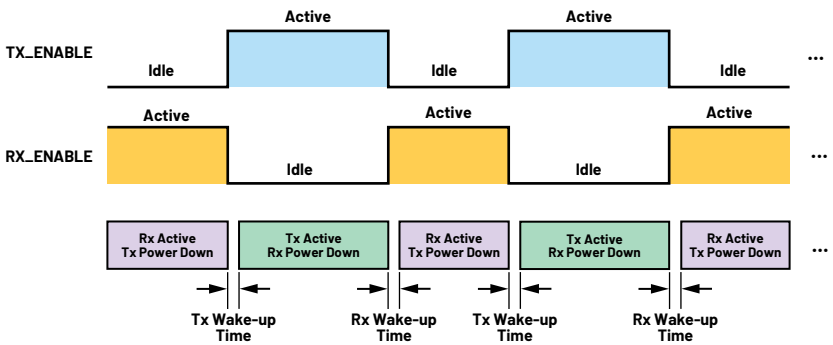

Figure 5. Channel-level power saving in a general TDD operation.

Based on Table 1, if an application has a relaxed performance requirement, by selecting the lowest power consumption option for each component, a total power saving of about 480 mW could possibly be achieved with this profile. Note: the component-level power saving options are mainly static, which means once configured during the device initialization stage, they could not be changed dynamically on-the-fly. One exception is the selection between the HP ADC or the LP ADC, which allows changing on-the-fly through an API command.

Another static power saving option worth mentioning is related to the configuration of one of its power domains. The ADRV9001 requires five different power supply domains: 1 V digital (VDD_1P0), 1.8 V digital (VDD_1P8), 1 V analog (VDDA_1P0), 1.3 V analog (VDDA_1P3), and 1.8 V analog (VDDA_1P8). Among them, the VDDA_1P0, which is used to power all the transmit and receive channel LO circuits, is optional. This domain can be powered using internal low dropout (LDO) regulators, which generate the 1 V required. Alternatively, it can be powered externally by bypassing some of the ADRV9001 internal LDO regulators, which is desirable to achieve more power saving by turning off the LDO regulators and applying a higher efficiency external power source.[3] Note: all the measurements performed in this article use internal LDO regulators to power up the VDDA_1P0 power domain.

Channel-Level Power Saving

Different from the static component-level power saving, channel-level power saving is dynamic. It is designed for TDD operations specifically. As shown in Figure 5, in TDD, transmit and receive operations are time multiplexing with each other. While one channel is active, the other channel is idle; therefore, it could be powered down to reduce power consumption. Different from component-level power saving, it does not result in any performance penalty by powering down the idle channel, but it takes more time to wake up to resume the normal operation.

One method to power up and down channels is to use the channel enable signal (TX_ENBALE/RX_ENABLE) rising and falling edge, respectively. As shown in Figure 5, the channel being powered down starts to wake up at the corresponding enable signal rising edge and it takes some time to become fully operational. If more channel components are powered down, then more wake-up time is required. The user should evaluate if the required wake-up time can satisfy the transmit and receive channel transition timing requirement in their TDD applications.

Three different modes of channel-level power saving are provided: Mode 0, Mode 1, and Mode 2. Each higher mode powers down additional channel-associated components by requiring a longer wake-up time. Table 2 summarizes these three modes along with the required approximate wake-up time at different RF PLL calibration modes and RF PLL reference clock rates.

Table 2. Channel-Level Power Saving Modes and the Required Wake-Up Time

Channel Components Powered Down		Mode 0	Mode 1	Mode 2
Transmit	Analog and Digital Data Path	X	X	X
	Tx Internal PLLs		X	X
	Tx PLL LDOs and Channel LDOs			X
Receive	Analog and Digital Data Path	X	X	X
	Rx Internal PLLs		X	X
	Rx PLL LDOs and Channel LDOs			X
Wake-Up Time at Different Configurations		**Mode 0**	**Mode 1**	**Mode 2**
Approximate Power-Up Time (µs) with RF PLL Normal Calibration Mode and Different RF PLL REF CLK Rates	RF PLL REF CLK = 30 MHz	4.5	350	500
	RF PLL REF CLK = 50 MHz		180	380
	RF PLL REF CLK = 100 MHz		170	370
Approximate Power-Up Time (µs) with RF PLL Fast Calibration Mode and Different RF PLL REF CLK Rates	RF PLL REF CLK = 30 MHz		100	300
	RF PLL REF CLK = 50 MHz		60	260
	RF PLL REF CLK = 100 MHz		40	240

As shown in Table 2, a higher channel-level power saving mode powers down additional channel components at the expense of longer wake-up time. By default, the channel power saving Mode 0 is always enabled if the user does not configure other modes. It powers down analog and digital datapath components such as mixers, converters, filters, etc. when the channel is idle. In Mode 0, only the RX_ENABLE and TX_ENABLE signals can be employed to trigger the power-up and

power-down. The wake-up time is short—around 4.5 µs. The channel power saving Mode 1 further powers down the channel's internal PLL. When the PLL is powering up, recalibration is mandatory, so the PLL wake-up time includes PLL power-up time and PLL calibration time. The ADRV9001 provides two PLL calibration modes: normal mode and fast mode. The fast mode does not guarantee a lock over the entire temperature range as the normal mode does, but it is more suitable when the channel stays at a particular frequency for a short period of time. As shown in Table 2, fast mode takes less calibration time than the normal mode; therefore, the PLL can wake up more quickly. In addition, a higher RF PLL reference clock rate also reduces the PLL calibration time. The channel-level power saving Mode 2 further powers down PLL LDO regulators and channel LDO regulators. It adds a fixed amount of wake-up time to turn on the LDO regulators. Note: the measurement of wake-up time displayed in Table 2 is performed at the ADRV9001 standard system clock rate of 184.32 MHz. When a custom profile with arbitrary sample rate is used, the system clock rate could change, which scales the PLL power-up time accordingly (lower system clock rate will increase the required PLL power-up time). The user could retrieve the system clock information from the ADRV9001 transceiver evaluation software (TES).

Modes 1 and 2 can be triggered by the RX_ENABLE and TX_ENABLE signal rising edge the same as Mode 0. In the case that a pair of transmit and receive channels shares the same internal PLL and its LDO regulators, the power saving achieved by modes 1 and 2 is limited when one channel is active, since the PLL and its LDO regulators must be powered up. Higher power saving can be achieved when both channels are idle. Different from Mode 0, Mode 1 and Mode 2 can also be triggered by a pre-assigned digital general-purpose input/output (DGPIO) pin. However, one DGPIO pin powers up and down both channels. Therefore, the DGPIO pin method can only be used when both transmit and receive channels are idle.

Figure 6 shows an example of using a DGPIO pin to trigger power saving Mode 1 or Mode 2. In this example, the entire TDD time period is divided into multiple time frames and each of them consists of four time slots. The first one is a transmit time slot, followed by two idle time slots, and the last one is a receive time slot. By default, Mode 0 is always enabled, which powers down the idle channel. However, during the idle time slots 2 and 3, both transmit and receive channels are idle; therefore, the DGPIO pin method could be used to trigger power saving Mode 1 or Mode 2, which achieves additional power saving than Mode 0 only.

It is important to emphasize that the DGPIO pin method should always trigger higher channel-level power saving modes than RX_ENABLE and TX_ENABLE signals as in the example shown in Figure 6. The DGPIO pin method helps to achieve more power saving in the scenario when Mode 1 and Mode 2 could not be triggered by RX_ENABLE and TX_ENABLE signals due to insufficient transmit and receive channel transition time.

In some TDD applications, one channel might be initialized but not used for a long period of time. In that case, an API command to power down the unused channel similar to Mode 2 (powering down its datapath, PLL, and LDO regulators) is provided for the user. This moves the unused channel to the hibernate state. Before the channel starts to operate, the user could power it up by using another API command. This ensures the best channel-level power saving for the unused channel is achieved. More discussions about channel/system states will be presented in later sections.

To demonstrate the power saving achieved through three different channel-level power saving modes, a DMR profile with 24 kSPS is employed. In DMR handset systems, the battery life is one of the key factors to decide the user experience. After powering up, the DMR handset is switched among three different states: transmit, receive, and idle. A typical cycle case is denoted as 5-5-90, which means the handset spends approximately 5% of the time on transmit, 5% of the time on receive, and 90% of the time on idle. Usually, the battery life data with the 5-5-90 cycle case needs to be published in the DMR handset data sheet as an important system parameter.[1]

Since power consumption is critical for DMR applications, the best power saving options are adopted at the component level. In addition, for a pair of transmit and receive channels, only one PLL is employed. Since the ADRV9001 receiver uses intermediate frequency (IF) mode and the transmitter uses zero-IF mode, the PLL is retuned when one channel is switching to the other channel. Figure 7

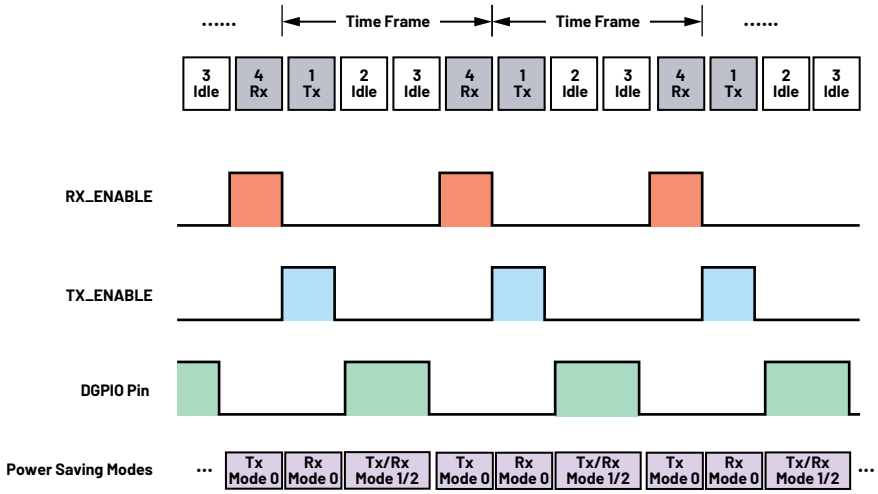

Figure 6. An example of using DGPIO to trigger channel-level power saving using Mode 1 or Mode 2.

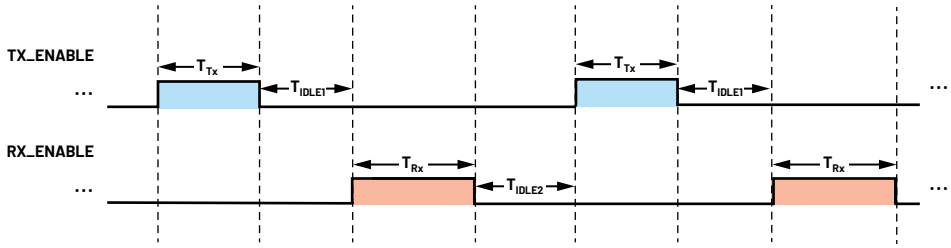

Figure 7. A general DMR TDD timing configuration for power consumption measurement using channel-level power saving modes.

describes a general TDD timing configuration. T_{TX} and T_{RX} stand for the transmit and receive active time, respectively. T_{IDLE1} and T_{IDLE2} stand for the idle time. For simplicity, wake-up time is not indicated since in general it is much shorter compared with the channel active and idle time; therefore, it is insignificant in power consumption calculation.

Table 3 presents the power consumption measured during T_{TX}, T_{RX}, and idle time (T_{IDLE1} /T_{IDLE2}) with the channel-level power saving modes 0, 1, and 2. In this measurement, the LO is configured at 900 MHz.

Table 3. Power Consumption During Different Time Periods for a TDD DMR Profile Using Channel-Level Power Saving Modes 0, 1, and 2

Channel-Level Power Saving Mode	Power Consumption (mW)		
	P_{TX} (Transmit Only)	P_{RX} (Receive Only)	P_{IDLE} (Idle)
Mode 0	580	525	368
Mode 1	580	509	205
Mode 2	580	502	173

By knowing the power consumption during different time periods, the average power consumption could be further calculated as:

$$P_{AVG} = P_{TX} \times \left(\frac{T_{TX}}{T_{TOTAL}}\right) + P_{RX} \times \left(\frac{T_{RX}}{T_{TOTAL}}\right) + P_{IDLE} \times \left(\frac{T_{IDLE1} + T_{IDLE2}}{T_{TOTAL}}\right) \quad (1)$$

$$T_{TOTAL} = T_{TX} + T_{RX} + T_{IDLE1} + T_{IDLE2}$$

Considering the typical 5-5-90 DMR use case, the average power consumption by using Mode 2 can be calculated as 580 × 5% + 502 × 5% + 173 × 90%, which is about 210 mW.

As shown in Table 3, Mode 1 and Mode 2 save more power during the idle time period since the PLL and its related LDO regulators can be powered down. But during the channel active time (either transmit or receive), the PLL and its LDO regulator can't be powered down since they are shared between both channels; therefore, the power saving is very limited by only powering down idle channel related components such as the channel LDO regulators.

System-Level Power Saving

As discussed in the previous section, channel-level power saving modes power down channel-associated components such as the datapath, RF PLL, and LDO regulators. In the case both transmit and receive channels are idle, such as in the scenario described in Figure 6, the system-level components could be further powered down to achieve additional power saving. Those system-level components include clock PLL, converter LDO regulators, clock PLL LDO regulators, and the Arm® processor and its memories. Similar to the channel-level power saving modes, three system-level power saving modes are provided, with the higher number modes powering down additional system components, which are summarized in Table 4.

Table 4. System-Level Power Saving Modes and the Required Wake-Up Times

Channel and System Components Powered Down		Mode 3	Mode 4	Mode 5
Tx	Analog and Digital Data Path	X	X	X
	Tx Internal PLLs	X	X	X
	PLL LDOs and Tx LDOs		X	X
Rx	Analog and Digital Data Path	X	X	X
	Rx Internal PLLs	X	X	X
	PLL LDOs and Rx LDOs		X	X
System	CLK PLL	X	X	X
	Converter LDOs and CLK PLL LDOs		X	X
	Arm + Memories			X
Wake-Up Time at Different Configurations		**Mode 3**	**Mode 4**	**Mode 5**
Approximate Power-Up Time (µs)		250	650	3200

As shown in Table 4, Mode 3 powers down the CLK PLL in addition to Mode 1, Mode 4 powers down the CLK PLL, converter LDO regulators and CLK PLL LDO regulators in addition to Mode 2. Mode 5 further powers down the Arm device and its memories in addition to Mode 4. Similarly, powering down more components causes longer wake-up time. In Mode 5, it takes approximately 3.2 ms to power up all the components.

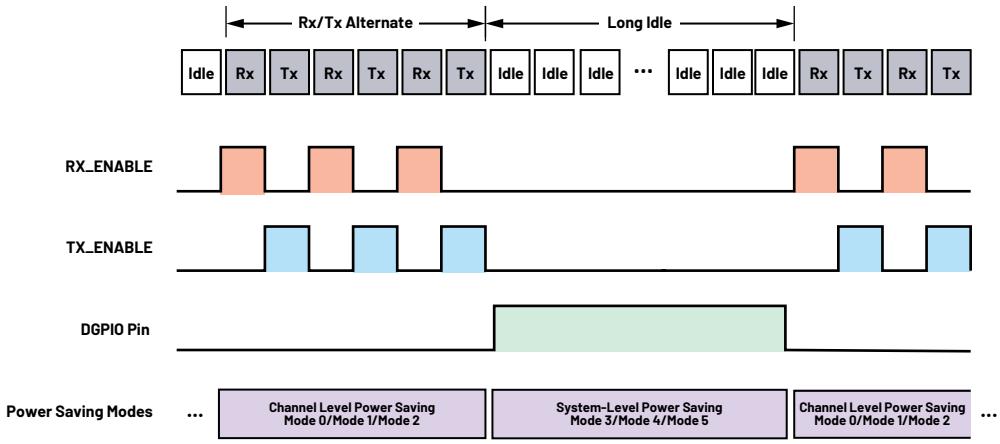

Figure 8. An example of using combined channel-level and system-level power saving.

Different from the channel-level power saving, the system-level power saving must be triggered by a DGPIO pin. Figure 8 shows an example of how to use a combined channel-level power saving and system-level power saving during different time periods of a TDD operation to achieve better power saving.

In this example, during the time period transmit and receive operations are alternate, users can select the highest possible channel power saving mode by using RX_ENABLE and TX_ENABLE signals. During the long idle time period when no channel is operating, the user can employ a DGPIO pin to trigger the highest system-level power saving mode, which allows to power down additional system components. This helps to achieve the better power saving compared with the channel-level power saving only. Similar to the DGPIO pin method in the channel-level power saving Mode 1 and Mode 2, the DPGIO pin method in the system-level power saving can only be employed when both TX_ENABLE and RX_ENABLE signals are low.

Table 5 presents the power consumption for the DMR use case shown in Figure 7 by using power saving Mode 2 when one channel is active, and three different system-level power saving modes when both channels are idle.

Table 5. Power Consumption During Different Time Periods of a TDD DMR Profile Using Channel-Level Power Saving Mode 2 and System-Level Power Saving Modes 3, 4, and 5

System-Level Power Saving Mode (For Idle Only)	Power Consumption (mW)		
	P_{TX} (Transmit Only, Mode 2)	P_{RX} (Receive Only, Mode 2)	P_{IDLE} (Idle)
Mode 3	580	502	100
Mode 4	580	502	65
Mode 5	580	502	35

Compared with Table 3, during the idle time period, it is clear that using the System-Level power saving modes can save more power. For the same 5-5-90 DMR use case, the average power consumption using Mode 5 is further lowered, which can be calculated as 580 × 5% + 502 × 5% + 35 × 90% = 86 mW.

Monitor Mode

In the previous sections, three different levels of power saving options were discussed. To achieve the best power saving in a system, reducing the power consumption for only the ADRV9001 might not be sufficient. Ideally, during the long idle time period, the best power saving of the entire system is achieved when all the major components can be powered down. To achieve this goal, a monitor mode is provided that allows both the ADRV9001 and the BBIC to go into deep sleep during the entire idle time period—except for one receive channel, which could optionally wake up to perform signal detection periodically. When a valid signal is found, the ADRV9001 wakes up the BBIC immediately. This offloads the signal detection responsibility from the BBIC and allows it (and possibly other circuitry in the system controlled by the BBIC) to sleep during the entire idle time period to achieve the highest overall system power saving.

Figure 9 shows a simplified state diagram of the ADRV9001 and how it transitions between normal operation mode and monitor mode.

As shown in Figure 9, with normal operation mode, after the ADRV9001 is powered up, it automatically goes to the standby state, during which the user can configure component-level power saving options. The standby state will switch to the calibrated state if the initialization is successful. As mentioned earlier, in this state, unused channels (although initialized) can be powered down using an API command to move from the calibrate ready substate to hibernate substate. From the calibrated state, the radio on command further primes the channels to prepare for transmit and receive operations and all channels are switched to the primed ready substate. Note: this substate is equivalent to the default

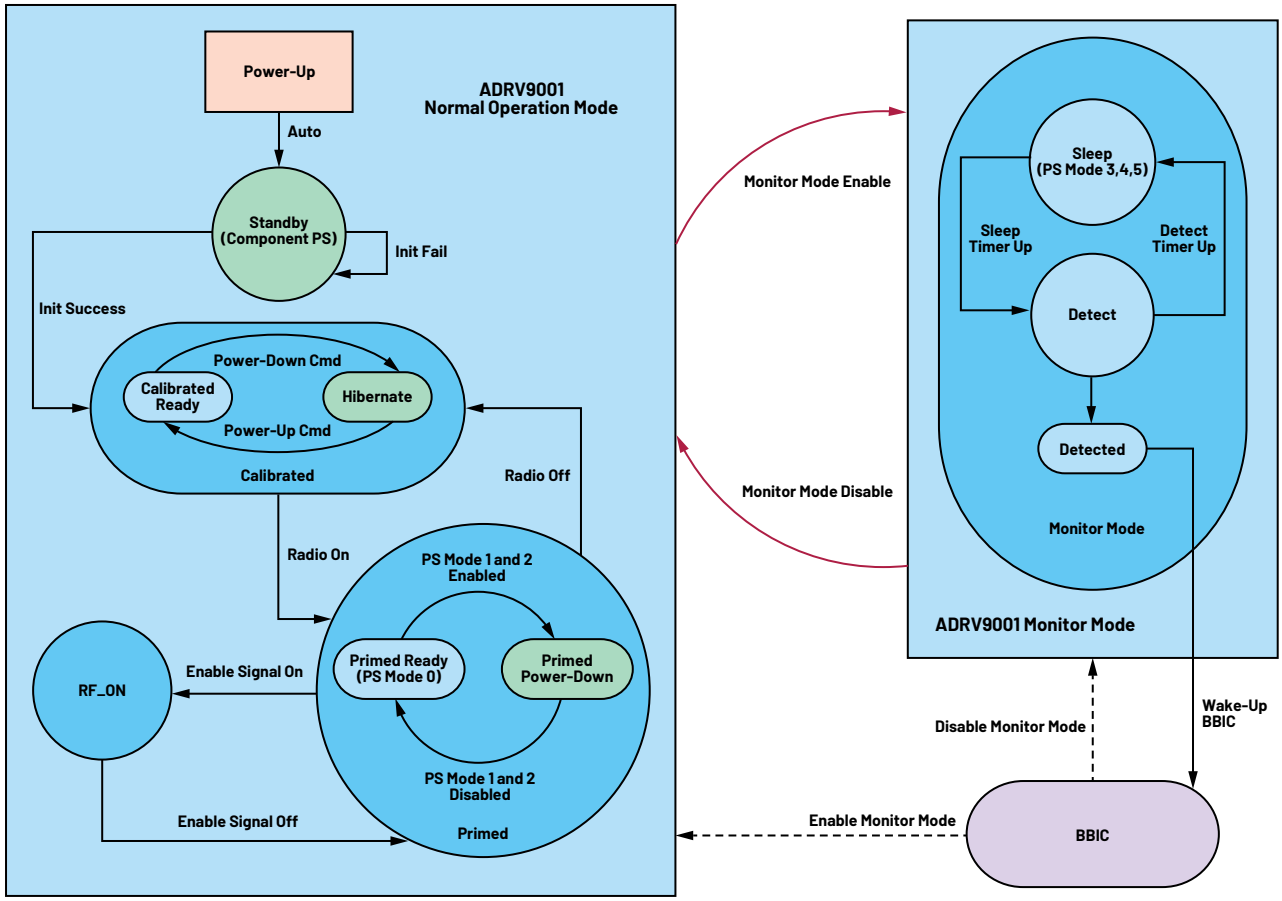

Figure 9. State diagram of the ADRV9001 in normal operation mode and monitor mode.

channel-level power saving Mode 0. When the channel enable signal is on, channels are further moved to RF_ON state to start operation. As discussed earlier, during a TDD operation, idle channels can be powered down using channel-level power saving modes. If using power saving Mode 0, it moves the idle channel from RF_ON state to the primed ready substate. If using power saving Mode 1 or Mode 2, it moves the idle channel from RF_ON to primed power-down substate.

The transition from normal operation mode to monitor mode is initiated by the BBIC after it detects the start of a long idle time period. In monitor mode, the BBIC employs system-level power saving Mode 3, Mode 4, or Mode 5 based on the configuration set by the BBIC. Both the ADRV9001 and BBIC go to sleep—except one ADRV9001 receive channel could optionally wake up to perform signal detection periodically. When a valid signal is found, the ADRV9001 will wake up the BBIC and the BBIC will further disable monitor mode to resume normal operation.

As shown in Figure 9, monitor mode consists of three different states: sleep, detect, and detected. The sleep and detect cycles are controlled through timers. When the time is up, one state will transition to another state if no valid signal is detected. The BBIC determines the timer and which state monitor mode should start with. If a valid signal is detected during the detect state, the ADRV9001 will transition to detected state immediately and wake up the BBIC. The BBIC then disables monitor mode, and the ADRV9001 switches back to normal operation mode. The start of monitor mode is triggered by a DGPIO pin as in the system-level power saving mode, since fundamentally these two are very similar except that monitor mode incorporates a signal detection capability. As a matter of fact, the ADRV9001 can dynamically switch between system-level power saving mode and monitor mode through an API command.

Figure 10 describes the detailed timing events happening during monitor mode for both the ADRV9001 and the BBIC. When the monitor mode DGPIO pin is asserted by the BBIC, the BBIC will start sleep and the ADRV9001 will wait for a configurable initial delay before going to sleep-detect pattern by using the configured timers. The ADRV9001 can perform signal detection during the initial delay to make sure that no signal is present before going to sleep. The sleep-detect pattern of the ADRV9001 continues to go on until a valid signal is detected. The ADRV9001 then wakes up the BBIC and starts to buffer the valid receive data to make sure the BBIC will not lose any valid data during the sleep. After the BBIC fully wakes up, it enables the receive channel to first retrieve all the buffered data at a preconfigured higher interface data rate. Then it further disables monitor mode to resume normal operation. Note: the BBIC can set the detect timer to be 0 so that the ADRV9001 will not perform any signal detection, and instead the BBIC will perform signal detection and terminate monitor mode by de-asserting the DGPIO pin at any time when a valid signal is found.

The ADRV9001 provides multiple signal detection methods to accommodate different radio standards, including receive signal strength indicator (RSSI), synchronization (SYNC), and fast Fourier transform (FFT). The RSSI method compares the receive signal level with a threshold to determine a valid signal so it can be used for any type of radio standards. SYNC method detects specific synchronization signal patterns defined by the DMR standard. FFT method is only applicable for standards using FSK modulation schemes. Therefore, there is no limitation to use monitor mode to other standards besides the DMR.

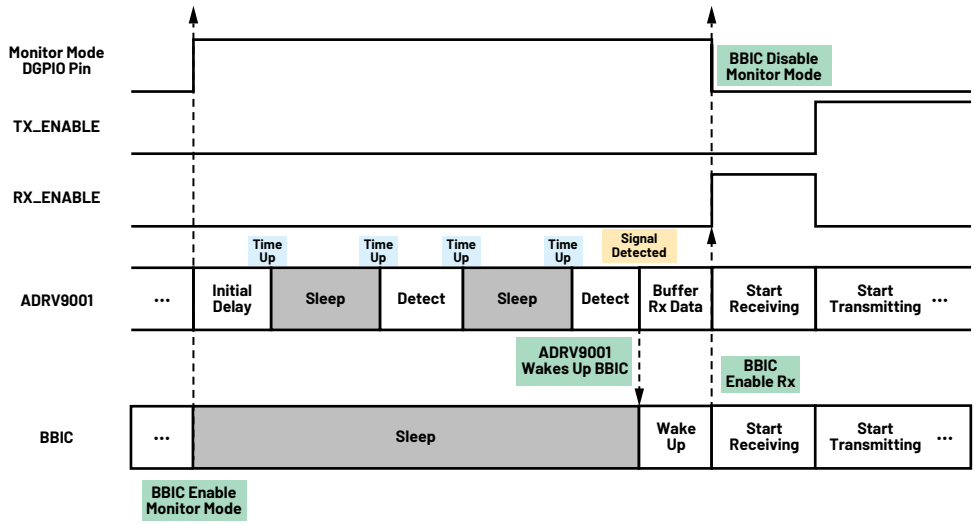

Figure 10. Timing events of ADRV9001 and BBIC during monitor mode.

Table 6 presents the power consumption during the sleep state and detect state utilizing different system-level power saving modes in monitor mode during the idle time period for the DMR use case shown in Figure 7.

Table 6. Power Consumption of a TDD DMR Profile During Sleep and Detect States Using System-Level Power Saving Modes 3, 4, and 5

System-Level Power Saving Mode	Power Consumption (mW)	
	Sleep	Detect
Mode 3	100	240
Mode 4	65	240
Mode 5	35	225

Depending on the timer configuration for sleep and detect states, the average power consumption during monitor mode could be determined. Although the ADRV9001 spends more power performing detection in detect state than sleep state, it allows the BBIC to sleep through the entire idle time period, which could result in higher overall system power saving.

Power Consumption Evaluation Through TES

All the power consumption measurements presented in this article are performed through the ADRV9001 TES with the ADRV9001 evaluation board (EVB). More information about TES and EVB can be found on the ADRV9002 product page. Both Xilinx® ZC706 and ZCU102 FPGA boards are supported by TES.[3] All the power saving options including monitor mode could be configured in TES, as shown in Figure 11.

The self-explanatory power saving configuration pages are very easy to use. To help users further evaluate power consumption, the ADRV9001 EVB is equipped with a power monitor chip to monitor and measure the power consumption in real time. Detailed power consumption at different power domains can be displayed in TES at 30 second intervals, as shown in Figure 12, which is a powerful visual tool to evaluate the power performance on-the-fly at different channel states. Good measurement accuracy can be achieved to be within ±2.5% error tolerance.

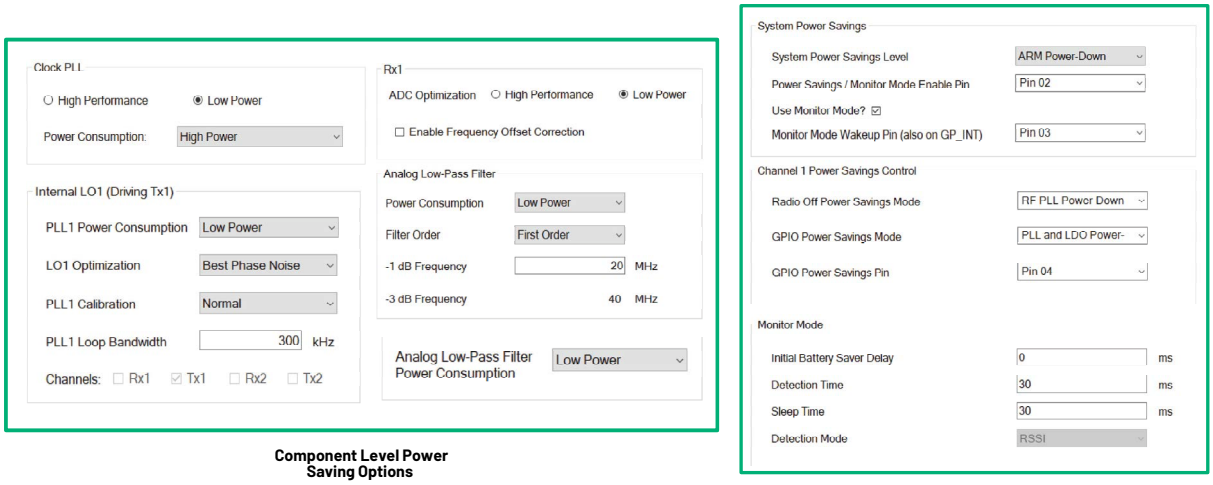

Figure 11. Power saving options and power monitor mode configuration in TES.

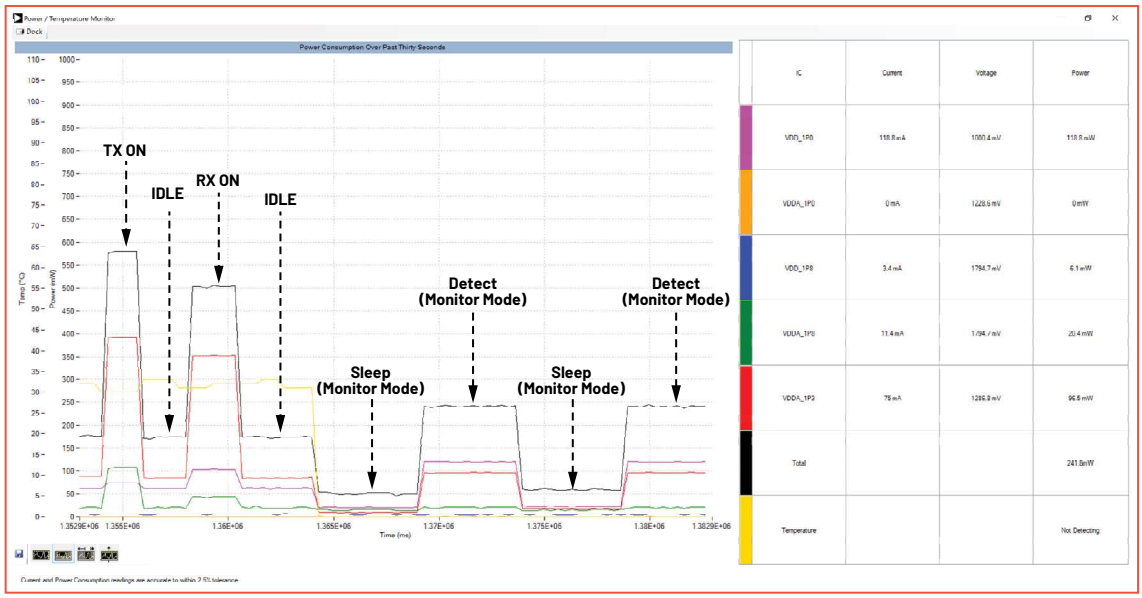

Figure 12. Power consumption real-time display using TES.

Conclusion

As discussed in this article, empowered with a variety of power saving options at the component, channel, and system levels, as well as the monitor mode, the ADRV9001 transceiver family is able to achieve a scalable power and performance for many mission critical applications. Understanding the associated performance trade-offs for each power saving option is crucial to determine the best system power saving strategy. All power saving options can be thoroughly evaluated through the ADRV9001 TES and EVB by a powerful real-time display of power consumptions on all power domains with a high accuracy.

References

[1] "Two-Way Radios and Battery Life." Hytera Europe, December 2016.

[2] ADRV9002 Dual Narrow-Band and Wideband RF Transceiver Data Sheet. Analog Devices, Inc., April 2021.

[3] ADRV9001 System Development User Guide. Analog Devices, Inc., October 2021.

About the Author

Mizhou (Michelle) Tan is a product applications engineer with Analog Devices. She has supported the design and development of RF transceiver products and applications for about 3 years. Prior to joining ADI, she received her B.S. degree and M.S degree in electrical engineering from Sichuan University in China and her Ph.D. degree in electrical and computer engineering from New Jersey Institute of Technology in 2004. After that, she worked as an algorithm, system, and software engineer at Agere Systems, LSI Logic, and Intel® Corp. from 2004 to 2018. She has published more than 15 papers in technical conferences and journals and owns nine issued patents in the wireless communication and digital signal processing area. She can be reached at mizhou.tan@analog.com.

How to Effectively Compare the Performance of CMOS Switches with Solid-State Relays

Stephen Nugent, Product Applications Engineer

Abstract

Off capacitance between the source and drain, $C_{DS(OFF)}$, is a measure of the ability of an off switch to block a signal on the source from coupling to the drain. It is a common specification seen in solid-state relays (also known as PhotoMOS®, OptoMOS®, photorelays, or MOSFET relays), and it is often referred to as output capacitance, C_{OUT}, in solid-state relay data sheets. CMOS switches do not usually include this specification, but the off-isolation spec is a different method to characterize the same phenomenon—that is, the amount of a signal that is presented to the source of an off switch that couples to the drain. This article will discuss how to derive C_{OUT} from off isolation and how this can be used to compare the performance of solid-state relays and CMOS switches more effectively. This is important as CMOS switches are a fit for many applications where solid-state relays are used, such as switching DC and high speed AC signals.

How to Derive $C_{DS(OFF)}$ from Off Isolation

Figure 1 shows the off isolation vs. frequency typical performance plot of the ADG5412. This plot shows that as the frequency of the signal on the source increases, the off isolation decreases.

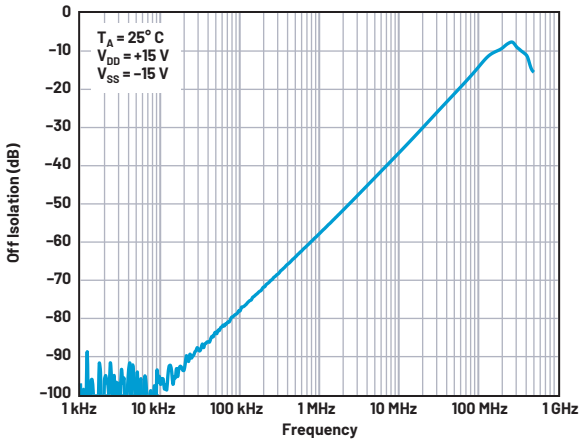

Figure 1. ADG5412 off isolation vs. frequency, ±15 V dual supply.

This means that more of the signal present on the source will appear on the drain of an off switch as the frequency of the signal increases. This is not surprising when you investigate the equivalent circuit for a switch in the off condition, as shown by the test circuit in Figure 2. When a switch is open there is a parasitic capacitance between source and drain, shown as $C_{DS(OFF)}$ in the figure. This parasitic capacitance enables high frequency signals to pass, and to characterize this is the purpose of the off-isolation plot.

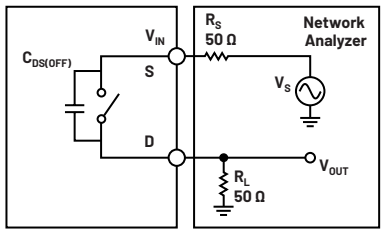

Figure 2. Off-isolation measurement test circuit.

Off isolation is calculated by taking V_S and V_{OUT} from the Figure 2 test circuit and inserting them into the following equation:

$$Off\ Isolation = 20 \log_{10} \frac{V_{OUT}}{V_S}$$

Using the results of the off-isolation plot combined with the equivalent circuit of an open switch, $C_{DS(OFF)}$ can be calculated in a CMOS switch. First if we consider the off-switch channel and the load we can equivalate the circuit to a high-pass filter, as shown in Figure 3.

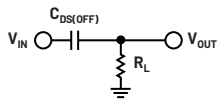

Figure 3. $C_{DS(OFF)}$ and R_L high-pass filter.

The transfer function of the circuit shown can then by derived by:

$$\frac{V_{OUT}(s)}{V_{IN}(s)} = \frac{R_L}{R_L + \frac{1}{C_{DS(OFF)}s}} = \frac{R_L C_{DS(OFF)}s}{1 + R_L C_{DS(OFF)}s}$$

Next is to consider the source voltage, V_S, and its impedance from Figure 2. The source impedance, R_S, is 50 Ω and this matches the load impedance, R_L, of 50 Ω. If we assume the ideal case where $C_{DS(OFF)}$ is a short circuit, then V_S is double V_{IN}, as the impedances are equal. This means when the transfer function is calculated in relation to V_S, the overall transfer function is doubled.

Therefore, the transfer function of the whole system is:

$$H(s) = \frac{V_{OUT}(s)}{V_S(s)} = \frac{2R_L C_{DS(OFF)}s}{1 + R_L C_{DS(OFF)}s}$$

This transfer function can then be substituted into the off-isolation equation to give:

$$Off_{iso} = 20 \log_{10}\left(\frac{2R_L C_{DS(OFF)}s}{1 + R_L C_{DS(OFF)}s}\right)$$

$$= 20 \log_{10}\left(\frac{2j2\pi f R_L C_{DS(OFF)}}{1 + j2\pi f R_L C_{DS(OFF)}}\right) \text{ as } s = j2\pi f$$

This equation can then be rearranged to make $C_{DS(OFF)}$ the subject:

$$C_{DS(OFF)} = \frac{1}{2\pi f R_L\left(\frac{2}{10^{\,offiso/20}} - 1\right)}$$

This means that if we know R_L, the frequency of the input signal, f, and the off-isolation specification value in dB, the $C_{DS(OFF)}$ can be calculated. These values can be found in the data sheet of switch or multiplexer products in the Analog Devices portfolio. The following example will outline how it can be done.

$C_{DS(OFF)}$ Calculation Example

The SPI controlled, quad SPST switch, the ADGS1612, will be used in this example. The off-isolation specification of the ADGS1612 is −65 dB, and this can be read from Table 1 of the data sheet. From the test conditions section of the off-isolation specification, R_L is given as 50 Ω and the signal frequency, f, is stated as 100 kHz. By putting these values into the $C_{DS(OFF)}$ equation, the capacitance value can be calculated.

$$C_{DS(OFF)} = \frac{1}{2\pi f R_L\left(\frac{2}{10^{\,offiso/20}} - 1\right)}$$

$$= \frac{1}{2\pi(100 \times 10^3)(50)\left(\frac{2}{10^{-65/20}} - 1\right)} \approx 9 \text{ pF}$$

Note, the measurement circuit for off isolation for switches and multiplexers may contain an additional 50 Ω termination before the source pin of the switch channel, as shown in Figure 4. The $C_{DS(OFF)}$ equation can still be used with off-isolation specifications that were measured in this way. However, 6 dB has to be added to the off-isolation specification from the data sheet when a 50 Ω termination was used at the source pin before using it in the $C_{DS(OFF)}$ equation. This is to compensate for the fact that the 50 Ω termination at the source decreases the voltage by half, which is equivalent to −6 dB.

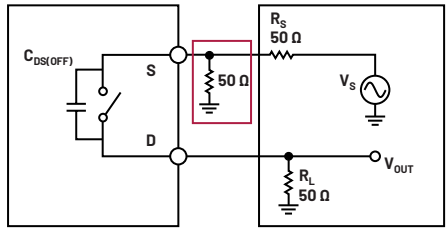

Figure 4. Off-isolation test circuit with a 50 Ω termination on the source.

CMOS Switches vs. Solid-State Relays

Table 1 shows the $C_{DS(OFF)}$ values for a selection of switch products in the Analog Devices portfolio. The ADG54xx and ADG52xx families can handle signal voltages with up to a 44 V swing, while the ADG14xx and ADG12xx families can pass signal voltages with up to a 33 V swing. This comparable signal ranges to 30 V and 40 V solid-state relays. The last column on the table also shows how the $C_{DS(OFF)}$ can be used in conjunction with the switch on resistance to calculate the R_{ON}, $C_{DS(OFF)}$ product, which is used as an order of merit in solid-state relays. The R_{ON}, $C_{DS(OFF)}$ product indicates how little a switch will attenuate a signal when it is on and combined with how well a switch can block high speed signals when the switch is off. The table shows that the ADG1412 has an R_{ON}, C_{OFF} product of less than 5, which is extremely competitive to solid-state relays on the market.

Table 1. $C_{DS(OFF)}$ for a Selection of SPST × 4 Switches in the Analog Devices Portfolio

	Max Supply Voltage	Off Isolation	CDS(OFF)	On Resistance	R × C
ADG5412	±22 V, +40 V	−78 dB @ 100 kHz	4 pF	9.8 Ω	39.2
ADG5212	±22 V, +40 V	−80 dB @ 1 MHz	0.32 pF	160 Ω	51.2
ADG1412	±16.5 V, +16.5 V	−80 dB @ 100 kHz	3.2 pF	1.5 Ω	4.8
ADG1212	±16.5 V, +16.5 V	−80 dB @ 1 MHz	0.32 pF	120 Ω	38.4

There are also many advantages of CMOS switches compared to solid-state relays. These include:

- Easier to Drive Switch Logic

The typical digital input current for most Analog Devices' CMOS switches is 1 nA, while the recommended forward current for the diode in solid-state relays is 5 mA. This means that CMOS switches can be easily controlled directly by the GPIOs on microcontrollers.

- Faster Switching Speeds

The ADG1412 has a typical turn-on time of 100 ns compared to solid-state relays, which have turn-on times in the region of hundreds of milliseconds.

- More Switches per Package

For example, the ADGS1414D has eight switch channels with 1.5 Ω on resistance and 5 pF $C_{DS(OFF)}$ in a 5 mm × 4 mm package. That is one switch per 2.5 mm^2 of package area.

Conclusion

The ability of a switch to block signals when in the off state is key. In solid-state relays, the C_{OFF} specification is a measure of the capacitance across the switch, which allows the coupling of signals from the input to the output of the closed switch. In CMOS switches this capacitor is not directly measured; however, the effect of this capacitor is measured through the off-isolation specification. The off-isolation value in dB, the frequency of the input signal, and the load resistance can be used to determine the $C_{DS(OFF)}$ by deriving the transfer function of an open switch. The $C_{DS(OFF)}$ is important in comparing CMOS switches to the C_{OUT} specification of solid-state relays. Furthermore, $C_{DS(OFF)}$ can also be used to calculate the R_{ON}, $C_{DS(OFF)}$ product, which is an order of merit that is used to show the overall off isolation and signal lost performance of a switch. This allows for direct competition between CMOS switches and solid-state relays when selecting a switch for an application. CMOS switches also have many benefits over solid-state relays—namely, easier to drive switch logic, faster switching speeds, and the ability to have more switches in a package.

About the Author

Stephen Nugent is an applications engineer in the Automatic Test Equipment Group at Analog Devices. He graduated from Queen's University Belfast in 2014 with a master's degree in electrical and electronic engineering. Stephen joined Analog Devices in 2014 supporting the analog switch and multiplexer product portfolio. He can be reached at stephen.nugent@analog.com.

Level-Setting DAC Calibration for ATE Pin Electronics

Minhaaz Shaik, Product Applications Engineer

Abstract

This article provides the methodology to calibrate digital-to-analog converters (DACs) specifically for pin electronic drivers, comparators, load, PMU, and DPS. DACs have nonlinear properties such as differential nonlinearity (DNL) and integral nonlinearity (INL), which can be minimized with the use of gain and offset adjustments. This article describes how to make those corrections for improved level-setting performance.

Introduction

Automated test equipment (ATE) describes testing apparatuses designed to perform a single or sequence of tests on one device or multiple devices at a time. Different types of ATE tests electronics, hardware, and semiconductor devices. Timing devices, DACs, ADCs, multiplexers, relays, and switches are the supporting blocks in the tester or ATE system. These pin electronic devices can deliver signals and power with precise voltages and currents. These precision signals are configured by the level-setting DACs. In the ATE portfolio, some pin electronic devices have calibration registers, and some calibration settings are stored off-chip. This article describes the DACs' function, errors, and calibration via gain and offset adjustments.

Digital-to-Analog Converter (DAC)

A DAC is a type of data converter that converts digital inputs to corresponding analog output levels. An N-bit DAC can support 2^N output levels. A higher number of bits corresponds to a higher DAC output resolution.

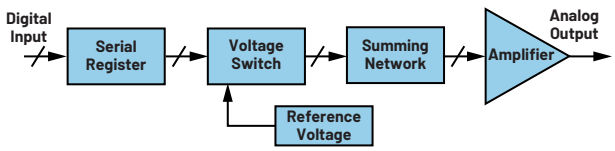

Figure 1. A digital-to-analog converter (DAC) block diagram.

First, the N-bit digital input is provided to a DAC serial register. The voltage switch and resistor summing network converts the digital inputs to analog output levels. The transfer characteristics of the DAC plot are shown in Figure 2. For a 3-bit DAC, 2^3 digital input yields eight analog output levels.

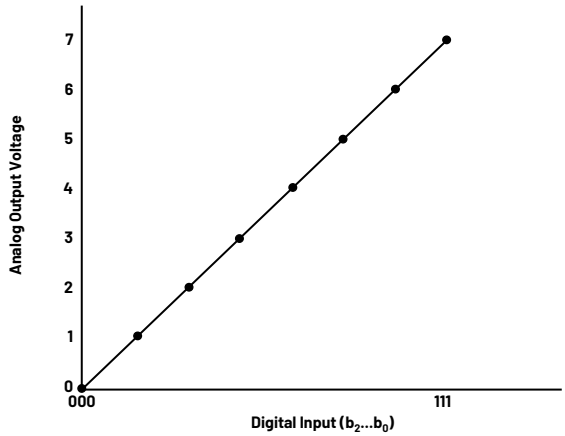

Figure 2. Ideal transfer function of a 3-bit DAC.

DAC Errors

In the real world, converters are not ideal. Because of the variance in resistance values, interpolation, and sampling, the DAC transfer function will not be a straight line, or linear. These errors are namely referred to as differential nonlinearity (DNL) and integral nonlinearity (INL). DNL is the maximum deviation of the output levels from ideal step sizes. It is derived from the difference between two successive output voltage levels. INL is the maximum deviation of the input/output characteristic from the ideal transfer function. With the gain and offset corrections, the INL errors can be reduced.

The INL in Figure 3 shows the deviation between actual transfer function and ideal transfer function. The gain error of the DAC indicates how well the slope of the linear approximation of the actual transfer function matches the slope of the ideal transfer function. Adjusting the gain will affect the angle of the linear approximation when graphed. The offset error is the difference between the measured value and chosen desired zero-offset point. Adjusting the offset will shift the entire linear approximation up or down accordingly. The INL of a single code is the sum of both gain error and offset error at any given point. After calibration, the transfer function can be a line drawn between end points once the gain and offset errors have been minimized.

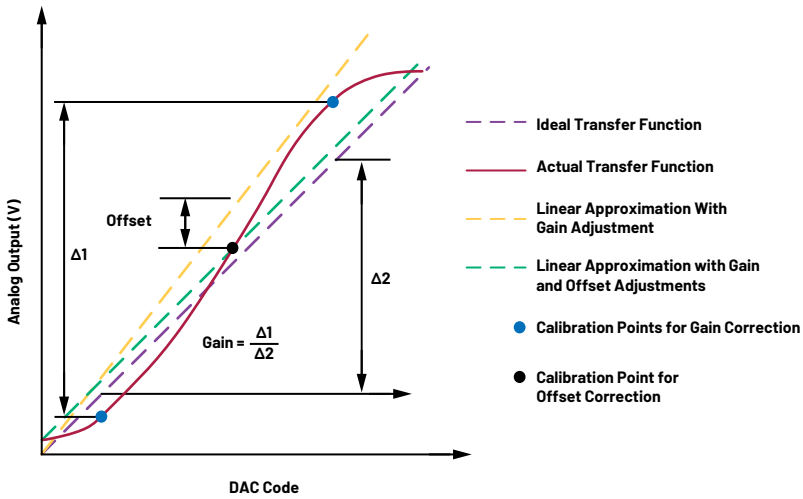

Figure 3. INL error transfer function.

Calibration Routine

The user can establish a calibration routine to reduce DAC nonlinearities using gain and offset corrections. The following procedure explains the step-by-step process of an example calibration routine.

For an N-bit DAC:

$$Maximum\ Code\ (MC) = (2^N - 1)$$

$$Voltage\ Range\ (V_{RANGE}) = Maximum\ DAC\ Output\ Voltage\ (V_{MAX}) - Minimum\ DAC\ Output\ Voltage\ (V_{MIN}) = 4 \times V_{REF}$$

$$DAC\ Input\ Code\ (Without\ Calibration) = \left(\frac{MC}{V_{RANGE}}\right) \times (V_{OUT} - V_{MIN})$$

- Gain correction (GC):

DACs tend to become less linear at the lowest and highest binary values. Therefore, it is recommended to choose calibration points within 5% to 10% in between the outer binary values or EC table recommended calibration points. For the following calculation, we assume 5% calibration points.

- Set the DAC input to 5% above the lowest binary value. Calculate the expected voltage output and record it as IDEAL1. Measure the output voltage and record it as MEAS1.
- Set the DAC input 5% below the highest binary value. Calculate and record IDEAL2. Measure the output voltage and record it as MEAS2.
- $$GC = \frac{MEAS2 - MEAS1}{IDEAL2 - IDEAL1}$$

$$DAC\ Input\ Code\ (with\ Gain\ Correction)$$

- $$= \left(\frac{MC}{V_{RANGE}}\right) \times (V_{OUT} - V_{MIN}) \times \frac{1}{GC}$$

- Offset correction (OC):

The desired zero-offset point varies by application. The user should define the best value based on their application. Some users may prefer to use zero volts to get an exact ground reference point. Some users prefer to use the midpoint of their operating range to minimize the overall INL error.

- Apply the gain correction of the DAC to the slope of the voltage-to-code equation to establish unity gain.
- Choose the desired zero-offset voltage point and record it as IDEAL3. Calculate the code using your updated voltage-to-code equation. Program your calculated code, then measure the output voltage and record it as MEAS3.
- $$OC = MEAS3 - IDEAL3$$
- $$DAC\ Input\ Code\ (with\ Gain\ and\ Offset\ Correction) = \left(\frac{MC}{V_{RANGE}}\right) \times (V_{OUT} - V_{MIN} - OC) \times \frac{1}{GC}$$

Example 1

Consider the MAX32007, an octal DCL with integrated level-setting DACs and PMU switches. The MAX32007 has internal DACs for level-setting VDH, VDL, VDT/VCOM, VCH, VCL, VCPH, and VCPL. These DACs do not have internal calibration registers. To calibrate the DACs, follow this procedure:

- Power up the MAX32007 evaluation (EV) kit by following the instructions in the EV kit data sheet.
- Connect the SMB connectors DATAOA and NTRMOA to 1.2 V.
- Connect the SMB connectors NDATAOA and TRMOA to ground through a 50 Ω terminator.
- Connect the EV kit to a Windows® 10 PC through a USB cable. Open the MAX32007 EV kit software (GUI).
- $$VDH\ DAC\ Resolution = N = 14$$

$$Maximum\ Code = (MC) = 2^N - 1 = 16383$$

$$Voltage\ Range\ (V_{RANGE}) = Maximum\ DAC\ Output\ Voltage\ (V_{MAX}) - Minimum\ DAC\ Output\ Voltage\ (V_{MIN}) = 7.5 - (-2.5) = 10$$

$$VDH\ DAC\ Input\ Code\ (Without\ Calibration) = \left(\frac{MC}{V_{RANGE}}\right) \times (V_{OUT} - V_{MIN})$$

Apply the DAC voltage levels and driver settings as shown in Figure 4. Note that the lowest operating VDH DAC value is –1.5 V, the highest operating value is 4.5 V; in this case, the zero-offset point value is 1.5 V.

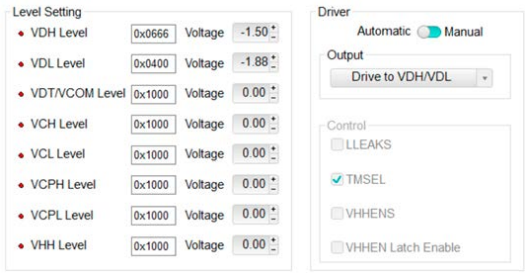

Figure 4. DAC-level setup of the MAX32007 using eval board software.

- Apply VDH = –1.5 V and measure the output voltage value.
- Apply VDH = 4.5 V and measure the output voltage value.
- Gain correction = Difference between measure output voltage values/difference between ideal values. For example, (4.501 – (–1.497)) / (4.5 – (–1.5)) = 0.999667.
- After applying gain correction,

$$VDH\ DAC\ Input\ Code\ (with\ Gain\ Correction) = \left(\frac{MC}{V_{RANGE}}\right) \times (V_{OUT} - V_{MIN}) \times \frac{1}{GC}$$

To apply gain correction, open **Menu → Options → Calibration**, as shown in Figure 5.

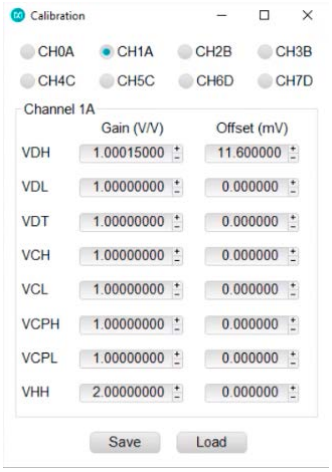

Figure 5. Calibration menu of the MAX32007 DAC.

- Apply VDH = 1.5 V (with gain correction code) and measure the output voltage value.
- Offset correction = Measure output value – Ideal value. For example, (1.502 – 1.5) = 0.002.
- After applying gain and offset correction,

$$VDH\ DAC\ Input\ Code\ (with\ Gain\ and\ Offset\ Correction) = \left(\frac{MC}{V_{RANGE}}\right) \times (V_{OUT} - V_{MIN} + OC) \times \frac{1}{GC}$$

Example 2

Consider the MAX9979, a dual DCL with integrated level-setting DACs and PMU. The MAX9979 has internal DACs for level-setting VDH, VDL, VDT, VCH, VCL, VCPH, VCPL, VCOM, VLDH, VLDL, VIN, VIOS, CLAMPHI/VHH, and CLAMPLO. These DACs have internal calibration registers. In Example 1, the DAC input codes are adjusted to minimize the INL error. In Example 2, the DAC input code remains the same and the calibration registers adjust the output stage buffer to minimize the INL errors, as depicted in Figure 6. To calibrate the DACs, use the following procedure:

- Power up the MAX9979 EV kit by following the instructions in the EV kit data sheet.
- Connect the SMB connectors DATAOA and NTRMOA to 1.2 V.
- Connect the SMB connectors NDATAOA and TRMOA to ground through the 50 Ω terminator.
- Connect the EV kit to a Windows 10 PC through a USB cable. Open the MAX9979 EV kit software (GUI).
- $VDH\ DAC\ Resolution = N = 16$

  $Maximum\ Code = (MC) = 2^{16} - 1 = 65535$

  $Voltage\ Range\ (V_{RANGE}) = Maximum\ DAC\ Output\ Voltage\ (V_{MAX}) - Minimum\ DAC\ Output\ Voltage\ (V_{MIN}) = 7.5 - (-2.5) = 10$

  $VDH\ DAC\ Input\ Code\ (Without\ Gain\ Correction) = \left(\frac{MC}{V_{RANGE}}\right) \times (V_{OUT} - V_{MIN})$
- Apply the DAC voltage levels and driver settings as shown in Figure 7. Note that the VDH DAC lowest recommended value is –1.5 V, the highest recommended value is 4.5 V, while the zero-offset point value is at 1.5 V.

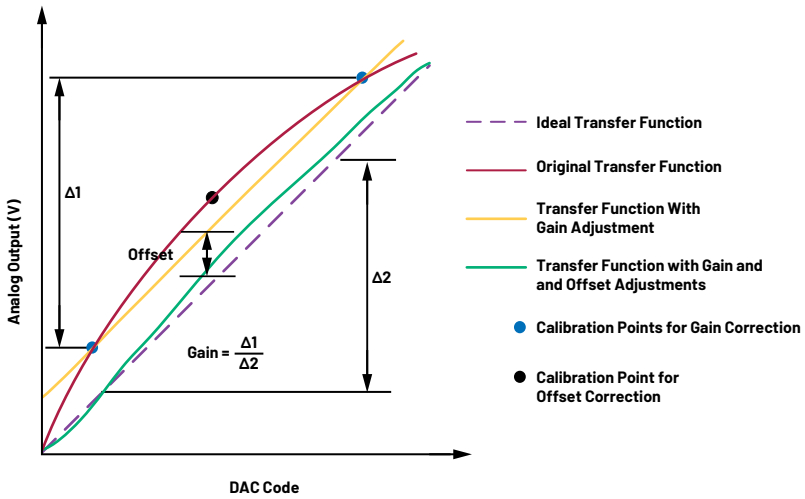

Figure 6. INL error correction for DACs with calibration registers.

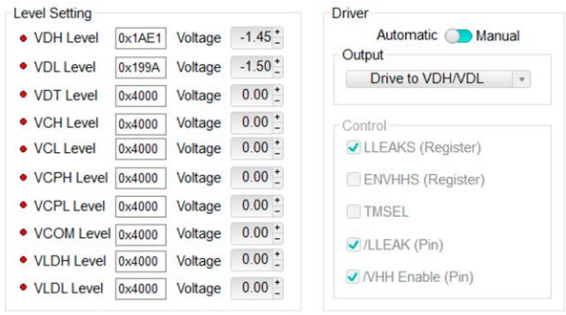

Figure 7. DAC-level setup of the MAX9979 using the eval board software.

- Apply VDH = –1.45 V and measure the output voltage value.
- Apply VDH = 6.5 V and measure the output voltage value.
- Gain correction = Difference between measure output voltage values/ difference between ideal values. For example, (6.501 V – (–1.455 V))/(6.5 V – (–1.45 V)) = 1.0007 V.
- After applying gain correction,

$$VDH\ DAC\ Input\ Code\ (with\ Gain\ Correction) = \left(\frac{MC}{V_{RANGE}}\right) \times (V_{OUT} - V_{MIN}) \times \frac{1}{GC}$$

- Apply VDH = 1.5 V (with gain correction code) and measure the output voltage value.
- Offset correction = Measure output value – Ideal value. For example (1.502 – 1.5) = 0.002.
- After applying gain and offset correction,

$$VDH\ DAC\ Input\ Code\ (with\ Gain\ and\ Offset\ Correction) = \left(\frac{MC}{V_{RANGE}}\right) \times (V_{OUT} - V_{MIN} + OC) \times \frac{1}{GC}$$

Note: gain and offset corrections can be applied on **Menu → Options → Change → Calibration**, as shown in Figure 8. Conversion of gain and offset corrections to gain and offset codes are given in the MAX9979 data sheet.

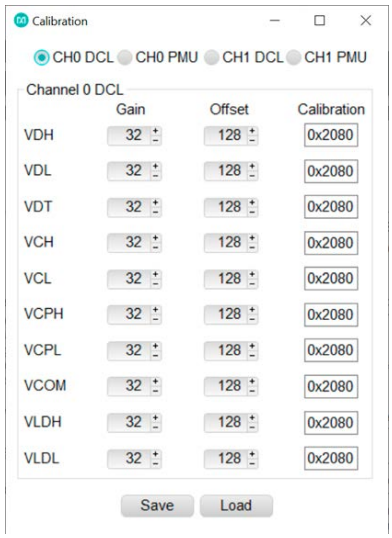

Figure 8. Calibration register setup for the MAX9979.

About the Author

Minhaaz Shaik is a member of the technical staff at Analog Devices with more than five years of experience as an applications/ systems engineer in the analog and mixed-signal domain. Minhaaz mainly focuses on product lines such as automated test equipment (ATE) pin electronics, ADCs, DACs, and supervisory and interface ICs. She is a highly skilled professional in electronic system design, lab evaluation, automation, customer support, and technical writing, and has significant knowledge in SPICE simulations and circuit design. She can be reached at minhaaz.shaik@analog.com.

RAQ Issue 199: Three Compact Solutions for High Step-Down Voltage Ratios

Olivier Guillemant, Central Applications Engineer

Question:

What are some methods for achieving a compact design under high step-down voltage ratios?

Answer:

This article will address why the nonisolated DC-to-DC buck converter (referred to simply as buck converter in this article) is facing serious challenges to downconverting high DC input voltages to very low output voltages at high output current. Three different approaches will be presented for downconverting steep voltage ratios while keeping a small form factor.

Introduction

System designers can be faced with the challenge of downconverting high DC input voltages to very low output voltages at high output current (such as 60 V down to 3.3 V at 3.5 A), while maintaining high efficiency, small form factor, and simple design.

Combining high input-to-output voltage difference with high current automatically excludes the linear regulator due to the excessive power dissipation. Consequently, the designer must opt for a switching topology under these conditions. However, even with such topologies, it is still challenging to implement a design that is sufficiently compact for space-restricted applications.

Challenges Faced by DC-to-DC Buck Converters

One candidate for high step-down ratios is the buck converter because it is the topology of choice when having to step down an input voltage to a lower output voltage (such as V_{IN} = 12 V down to V_{OUT} = 3.3 V) in an efficient way, with a significant amount of current while also using a small footprint. However, there are conditions under which the buck converter faces serious challenges to keep its output voltage regulated. To understand these challenges, we must remember that the simplified duty cycle (D) of a buck converter operating in continuous conduction mode (CCM) is:

$$D = \frac{V_{OUT}}{V_{IN}} \quad (1)$$

Now, the duty cycle also relates to the switching frequency (f_{SW}) in the following way, where the on-time (t_{ON}) is the duration over which the control FET stays on during each switching period (T):

$$D = \frac{t_{ON}}{T} = t_{ON} \times f_{SW} \quad (2)$$

Combining Equation 1 and Equation 2 shows how t_{ON} is influenced by the step-down voltage ratio and f_{SW}:

$$t_{ON} = \frac{V_{OUT}}{V_{IN}} \times \frac{1}{f_{SW}} \quad (3)$$

Equation 3 tells us that the on-time decreases when the input-to-output voltage ratio (V_{IN}/V_{OUT}) and/or f_{SW} increase. This means that the buck converter must be able to operate with very low on-time to regulate the output voltage in CCM under high V_{IN}/V_{OUT} ratio, and it becomes even more challenging with a high f_{SW}.

Let's consider an application with $V_{IN(MAX)}$ = 60 V, V_{OUT} = 3.3 V at $I_{OUT(MAX)}$ = 3.5 A. When required, we shall use values from the LT8641 data sheet because a solution with the LT8641 will be provided in a later section. The required minimum on-time ($t_{ON(MIN)}$) corresponds to the highest input voltage ($V_{IN(MAX)}$). In order to assess this $t_{ON(MIN)}$, it is advised to make Equation 3 more accurate. By including $V_{SW(BOT)}$ and $V_{SW(TOP)}$, the voltage drops for the two power MOSFETs of the buck converter, and replacing V_{IN} with $V_{IN(MAX)}$ we obtain:

$$t_{ON(MIN)} = \frac{V_{OUT} + V_{SW(BOT)}}{V_{IN(MAX)} - V_{SW(TOP)} + V_{SW(BOT)}} \times \frac{1}{f_{SW}} \quad (4)$$

Using Equation 4 with $V_{IN(MAX)}$, f_{SW} = 1 MHz, we obtain a $t_{ON(MIN)}$ of 61 ns. For $V_{SW(BOT)}$ and $V_{SW(TOP)}$, we made use of the values provided for $R_{DS(ON)(BOT)}$ and $R_{DS(ON)(TOP)}$ in the LT8641 data sheet, knowing as well that $V_{SW(BOT)} = R_{DS(ON)(BOT)} \times I_{OUT(MAX)}$ and $V_{SW(TOP)} = R_{DS(ON)(TOP)} \times I_{OUT(MAX)}$.

Buck converters can rarely guarantee a $t_{ON(MIN)}$ with the short value of 61 ns obtained above; therefore, the system designer is forced to search for alternative topologies. There are three possible solutions for high step-down voltage ratios.

Three Compact Solutions for $V_{IN(MAX)}$ = 60 V, V_{OUT} = 3.3 V at $I_{OUT(MAX)}$ = 3.5 A

Solution 1: Using the LT3748 Non-opto Flyback

The first option consists of using an isolated topology, where the transformer performs most of the downconversion thanks to its N:1 turn ratio. For that matter, Analog Devices offers flyback controllers such as the LT3748 that do not require a third transformer winding or opto-isolator, making the design simpler and compact. The LT3748 solution for our conditions is presented in Figure 1.

Even though the LT3748 solution simplifies the design and saves space compared with a standard flyback design, a transformer is still required. For applications where isolation between input and output sides is not required, it is preferred to avoid this component, which adds complexity and increases the form factor vs. a nonisolated solution.

Solution 2: Using the LTM8073 and LTM4624 µModule Devices

As an alternative, the designer can downconvert in two steps. To achieve a reduced component count of only 10, two µModule® devices and eight external components can be used, as demonstrated in Figure 2. Moreover, the two µModule devices already integrate their respective power inductor, sparing the system engineer a design task that is rarely straightforward. The LTM8073 and LTM4624 both come in BGA packages, with respective dimensions of 9 mm × 6.25 mm × 3.32 mm and 6.25 mm × 6.25 mm × 5.01 mm (L × W × H), providing a solution with a small form factor.

Since the LTM4624 exhibits an efficiency of 89% under these conditions, the LTM8073 supplies at most 1.1 A to the input of the LTM4624. Given that the LTM8073 can provide up to 3 A of output current, it can be used to supply other circuit rails. It is with this purpose in mind that we selected 12 V as the intermediary voltage (V_{INT}) in Figure 2.

Despite avoiding the usage of a transformer, some designers might be reluctant to implement a solution that requires two separate buck converters, especially if no intermediary voltage is required to supply other rails.

Solution 3: Using the LT8641 Buck Converter

Consequently, in many cases, using a single buck converter would be preferred because it provides the optimal solution to combine system efficiency, a small footprint, and design simplicity. But did we not just demonstrate that buck converters cannot cope with high V_{IN}/V_{OUT} combined with high f_{SW}?

This statement might apply to most buck converters, but not to all of them. The ADI portfolio includes buck converters such as the LT8641, which is specified with a very short minimum on-time of 35 ns typical (50 ns max) over the full operating temperature range. Those specifications are safely below the required minimum on-time of 61 ns previously calculated, providing us with a third possible compact solution. Figure 3 shows how simple the LT8641 circuit can be.

It is also worth noting that the LT8641 solution can be the most efficient of the three. Indeed, if efficiency must be further optimized compared with Figure 3, we can decrease f_{SW} and select a bigger inductor size.

Although f_{SW} can also be decreased with Solution 2, the integration of the power inductors does not offer the flexibility to increase the efficiency beyond a certain point. Moreover, the use of two consecutive downconversion stages has a small negative impact on the efficiency.

In the case of Solution 1, the efficiency will be very high for a flyback design, thanks to the operation in boundary mode and to all components removed with the no-optical feedback design. However, the efficiency cannot be fully optimized because there is a limited number of transformers to select from, as opposed to the broad portfolio of inductors available for Solution 3.

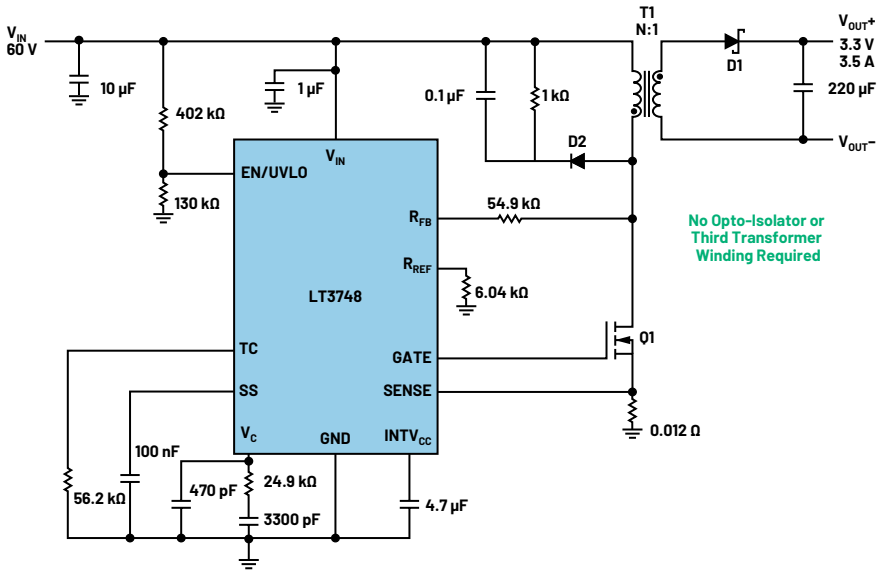

Figure 1. A circuit solution with the LT3748 downconverting 60 V input to 3.3 V output.

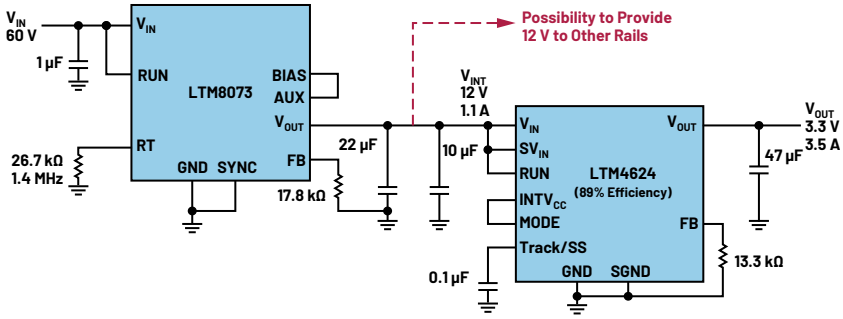

Figure 2. A circuit solution with the LTM8073 and LTM4624, downconverting 60 V input to 3.3 V output.

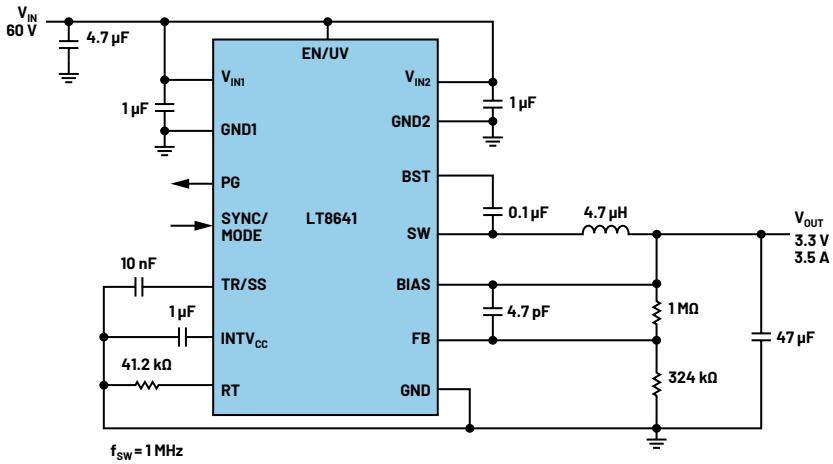

Figure 3. A circuit solution with the LT8641 downconverting 60 V input to 3.3 V output.

An Alternative Way to Check Whether LT8641 Fulfills Requirements

In most applications, the only adjustable parameter in Equation 4 is the switching frequency. Consequently, we reformulate Equation 4 to assess the maximum permitted f_{SW} for the LT8641 under given conditions. By doing this, we obtain Equation 5, which is also provided on page 16 of the LT8641 data sheet.

$$f_{SW(MAX)} = \frac{V_{OUT} + V_{SW(BOT)}}{t_{ON(MIN)} \times (V_{IN(MAX)} - V_{SW(TOP)} + V_{SW(BOT)})} \quad (5)$$

Let's use this equation with the following example: V_{IN} = 48 V, V_{OUT} = 3.3 V, $I_{OUT(MAX)}$ = 1.5 A, f_{SW} = 2 MHz. An input voltage of 48 V is commonly found in automotive and industrial applications. By inserting those conditions in Equation 5, we obtain:

$$f_{SW(MAX)} = \frac{5\ \text{V} + 0.0825\ \text{V}}{50\ \text{ns} \times (48\ \text{V} - 0.1575\ \text{V} + 0.0825\ \text{V})} = 2.12\ \text{MHz} \quad (6)$$

Therefore, under the provided application conditions, the LT8641 would operate safely with f_{SW} set as high as 2.12 MHz, confirming that the LT8641 is a good choice for this application.

Conclusion

Three different methods were presented to achieve a compact design under high step-down voltage ratios. The LT3748 flyback solution does not require a bulky opto-isolator and is recommended for designs where isolation is necessary between input and output sides. The second method, which involves implementing the LTM8073 and LTM4624 µModule devices, is of particular interest when the designer is hesitant to select the optimal inductor for the application and/or when an additional intermediary rail must be supplied. The third method, a design based on the LT8641 buck converter, offers the most compact and simplest solution when the sole requirement is the steep voltage downconversion.

About the Author

Olivier Guillemant is a central applications engineer at Analog Devices in Munich, Germany. He provides design support for the Power by Linear portfolio for European broad market customers. He has held various power application positions since 2000 and joined ADI in 2021. He received his M.Sc. in electronics and telecommunications from University of Lille, France. He can be reached at olivier.guillemant@analog.com.

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

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, ez.analog.com/cn (China)

Facebook
facebook.com/AnalogDevicesInc

@ADI_News
twitter.com/ADI_News

LinkedIn
linkedin.com/company/analog-devices

Sina Weibo (China)
weibo.com/analogdevices

WeChat (China)

Analog Devices, Inc.
Worldwide Headquarters

Analog Devices, Inc.
One Analog Way
P.O. Box 9106
Wilmington, MA 01887
U.S.A.
Tel: (781) 935-5565
((800) 262-5643, U.S.A. only)
Fax: 781.461.3113

Analog Devices, Inc.
Europe Headquarters

Analog Devices GmbH
Otl-Aicher-Str. 60-64
80807 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

M02000-5/22

analog.com