# AnalogDialogue

Volume 52, Number 4, 2018

Your Engineering Resource for Innovative Design



Visit *analogdialogue.com*

ANALOG DEVICES

AHEAD OF WHAT'S POSSIBLE™

# In This Issue

Bernhard Siegel, Editor

Bernhard became editor of *Analog Dialogue* in March 2017. He has been with Analog Devices for over 25 years, starting at the ADI Munich office in Germany. In his current role as the chief technical editor, 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* homepage, *analogdialogue.com*.

# Use of Integrated Passives in Micromodule SIPs

By **Mark Murphy** and **Pat McGuinness**



## Introduction

Integrated passives are nothing new in our business—they're long established and well understood. In fact, ADI produced such components for the market in the past. When standalone discrete passives or even integrated passive networks are included as part of a chipset, then routing parasitics, device compatibility, and board assembly considerations will require careful design management. While integrated passives continue to have a strong presence in the industry, their most significant value begins to be realized only when they are included in system in package applications.

A new integrated passives technology initiative (*i*Passives™) began at ADI several years ago. The aim of this initiative was to deliver passive components, diodes, resistors, inductors, and capacitors that could enable ADI to encompass more of the signal chain while overcoming the past limitations and complexities of the existing approaches of adopting passive components. The demand from ADI's customer base for more complete solutions in spatially efficient footprints also drove these developments. From the designer's perspective, *i*Passives can be viewed as a flexible design tool that enables the design of system solutions with best-in-class performance and robustness in exceptionally short development cycle times. ADI has many signal conditioning ICs whose superior performance is enabled by the unique silicon fabrication processes that we have at our disposal. Without the complexity of having to develop highly complex integrated processes, ADI can leverage the diversity of its existing offering to produce plug and play systems that have exceptional performance characteristics. The integrated passives technology is being used to tie this all together in highly customizable networks and packaged up via system in package technology to create µModule® devices that are fully qualified, tested, and characterized. Systems that were previously board-level solutions can now be reduced to what appears as a single device. From our customer perspective, they now get complete solutions with excellent out of the box performance, short development cycles, and cost savings—all in very compact packages.

## Passives Technology

So, let's go back briefly to the basics and remind ourselves of what a passive component is. Passive components are unpowered devices that give a relatively simple relationship between current and voltage. These components are resistors, capacitors, inductors, transformers (which are effectively coupled inductors), and diodes. Sometimes the current-voltage relationship is very simple, like in a resistor where the current is linearly dependent on voltage. For a diode there is a direct relationship between current and voltage, but the relationship is exponential. In inductors and capacitors, the relationship has a current and voltage transient dependency. Tables 1 shows the formulae that define these relationships for four of the basic passive components:

**Table 1. Basic Formulae for Primary Passive Components**

| Discrete Element | Equation | Symbol | |
|---|---|---|---|
| Resistor | $V = I \times R$ |  | V = Voltage<br>I = Current<br>t = Time<br>R = Resistance in ohms<br>C = Capacitance in farads<br>L = Coil inductance in henrys<br>$I_S$ = Diode saturation current<br>$V_T$ = Thermal voltage<br>h = Diode ideality factor |
| Capacitor | $I = C\frac{dV}{dt}$ |  | |
| Inductor | $V = L\frac{dI}{dt}$ |  | |
| Diode | $I = I_s\left(e^{\frac{V}{hV_T}} - 1\right)$ |  | |

Passive devices can be used individually, can be connected in series or parallel, and are essential components in analog signal processing (RLC for amplification, attenuation, coupling, tuning, and filtering), in digital signal processing (pull-up, pull-down, and impedance matching resistors), in EMI suppression (LC noise suppression), and in power management (R for current sensing and limiting, LC for energy accumulation).

## Limitations of Discrete Components

Historically, passive components have been discrete, meaning they are manufactured separately and connected in a circuit using conductive wires or tracks on a printed circuit board (PCB). Over time they have evolved along three paths: smaller size, lower cost, and higher performance. The evolution is now mature and optimized, but the size of the footprint and profile means that discrete passive components are always limiting the effort to reduce overall solution area and volume. Passives generally constitute over 80% of the bill of materials in an application, occupy about 60% of the area, and make up about 20% of the overall component spend. These factors compound to produce very complex inventory control and storage challenges.

By their very nature, discrete devices are individually processed components. While there may be ways of ensuring that components can be selected from certain process batches, there is still a high degree of uniqueness to every single component. This, however, is a significant disadvantage when it comes to needing very well matched components. For devices that are meant to match, uniqueness and differences between components contribute to errors that degrade time zero circuit performance. Additionally, this performance degradation invariably gets worse across the operating temperature and lifetime of the circuit.

Another drawback to discrete passive devices is that assembly and wiring of individual components takes time, and simply occupies a large space. The elements are connected using a soldering process, generally through-hole or surface-mount technology (SMT) assembly. Through-hole is the older assembly technology where leaded parts are inserted into holes on a PCB, any excess in lead length is clinched and cut back, and a molten wave of solder connects the leads of the devices to the PCB interconnect tracks. Surface-mount assembly has enabled the development of smaller passive components; in this case, a landing pattern is etched on a PCB, solder paste is used to cover the patterns, and then SMT components are positioned using a pick and place machine. The PCB is then run through a soldering reflow process, where the solder liquefies and establishes electrical connections, and, when cooled, the solder solidifies and mechanically affixes the SMT components to the PCB. The primary problem with both assembly technologies is that the soldering processes can be very unreliable and, in an industry where defect targets are in parts per million, this is becoming more and more of a concern. Several factors are important in ensuring solder reliability: the actual composition of the solder (which is generally lead free now and thus less reliable), mechanical stability during the solder reflow process (mechanical vibrations can render a dry solder joint), purity of the solder (any contaminants adversely affect solder reliability), and the time and temperature in the solder reflow process. How quickly solder is heated, what the actual temperature and uniformity of temperature is, and for how long the solder is heated are critical. Any variation here can lead to damage to landing pads or through-holes, or it can induce mechanical stress on components that will lead to failure over time.

Another limitation in employing passive components on PCBs due to them being spread apart is that trace lengths need to be long. This can introduce unaccounted for parasitic elements that can limit performance and the repeatability of results. Typically, PCB traces have in the order of 1 nH/mm length and capacitance that depends on the width of trace and proximity to its neighbors. PCB trace tolerances contribute to variability in parasitics, so not only are the parasitics disruptive, but they are also unpredictable. Tighter tolerances on PCB boards cost more.

Passive devices also present a lot of potential contact points to the outside world, points where ESD events can occur, via manual or machine handling. Again, this has consequences and risks in overall reliability and robustness.

## Advantages of Integrated Passives

Before delving into the advantages integrated passives have over their discrete passive counterparts, let us first outline the origins of integrated passives. Integrated circuits now contain many transistors (millions in fact) wired together with very well defined metal interconnections. Special processes have been developed for analog type applications like DACs and ADCs that contain portfolios of passive components such as resistors and capacitors in addition to the transistors. To achieve the performances required in these precision analog applications, very high quality passive components have been developed. It is these high quality passive components that are used to build integrated passives. Just as integrated circuits contain many transistors, integrated passives can contain many high quality passive components packed into a very small area. Like integrated circuits, integrated passives are fabricated on large area substrates (wafers) where multiple passive networks are produced at the same time.

One of the most compelling advantages integrated passives have to offer over discrete passives is the precise matching that can be achieved with them. When integrated passive networks are fabricated, all components within a network are manufactured at the same time, under the same conditions, with the same material set and, because of network compactness, essentially in the same place. Passive components manufactured in this way have a much better chance of being very well matched than discrete passive alternatives. To illustrate this let's say we have an application that requires two matched resistors. These resistors are fabricated on circular substrates such as silicon wafers, as illustrated in Figure 1. Due to slight process variations like resistive film thickness, chemical properties of the film, contact resistance, etc., there will be some level of resistance variation within a batch and even more variation across multiple batches. In the example in Figure 1, dark green indicates the resistance is on the high side of the tolerance range and yellow indicates the resistance is on the low side of the tolerance range. For standard discrete devices, there is the possibility that each of the two resistors may come from different fabrication batches as indicated by the two separate resistors drawn in red. The full tolerance range of the process may be observable between the two discrete resistors and hence the matching will not be very good. With special ordering restrictions it may be possible to have the two discrete resistors chosen from the same batch as indicated by the two separate resistors drawn in blue. The tolerance range within just one batch may be observable between the two resistors. While the matching between these resistors will be better than the random discrete case, there is still scope for some level of mismatch. Finally, with integrated passives, the two resistors come from the same die, as indicated by the resistors drawn in black in Figure 1. The tolerance range within one die is the only range that will be observable between the two resistors. Hence, matching between the two resistors will be excellent. Additional techniques of using cross quad layouts and other means can be employed to tighten the spread between the two resistors even further, which culminates in extremely well matched components. Matching between integrated passive components is not only much better than discrete passives at time zero, but will also track better across temperature, mechanical stress, and lifetime due to their manufacture having been so well coupled.

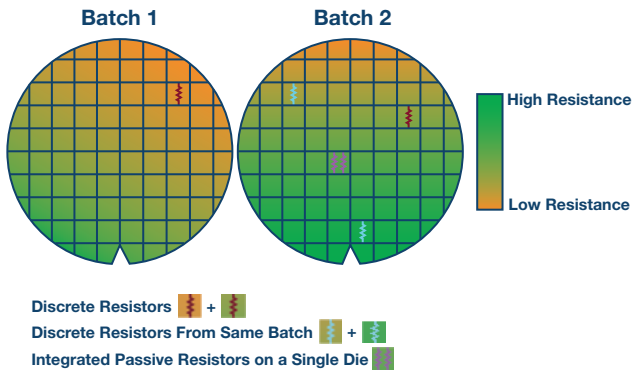


*Figure 1. Matching of discrete vs. passive resistors.*

The individual components within integrated passives are placed closely together (within microns in fact) and, because of this, the interconnect parasitics such as trace resistance and inductance can be kept to an absolute minimum. On PCBs, interconnect parasitics can be variable due to trace tolerances and component placement tolerances. With integrated passives, interconnect tolerances and component placement tolerances are extremely tight due to the photolithographic processes employed in their manufacture. On integrated passives, not only are the parasitics very small, but the few that are there are very predictable and, hence, can be reliably accounted for.

The miniaturization of passive networks through integrated passives has the very straightforward benefit of making circuit boards smaller. This leads directly to reducing circuit board costs and to allowing more and more functionality and performance to be packed into a smaller footprint. Building systems with high channel counts becomes much more practical when using integrated passives.

Another significant advantage of integrated passives is the robustness of the complete wiring network around them. Rather than needing lots of soldered connections, integrated passives are essentially forged together in one complete unit, sealed up with glass, and then further protected with a robust plastic encapsulant. Within integrated passive networks, issues with dry solder joints, corrosion, or misplaced components are not present.

An additional benefit of integrated passive networks being so well sealed up is that the number of exposed nodes in a system is much reduced. Hence, the possibility of systems being damaged by accidental shorting or electrostatic discharge (ESD) events are significantly diminished.

Maintaining and controlling the stock of components for any circuit board assembly is quite a complex task. Integrated passives having multiple passive components within one device greatly unburdens the customers' bill of materials, which results in the cost of ownership going down. Customers receive integrated passive networks fully tested and proven good. This means that the final board build yields are improved and this not only enables further cost savings but also enhances supply chain predictability.

## Use of Integrated Passives by ADI (*i*Passives)

As alluded to earlier, high quality passive devices have been core to the circuit performances achieved by many of ADI's products over many years. During that time, the range and quality of the passive devices has grown and the integrated passives portfolio now contains a significant number of components. The integrated passives process is modular, which means that the processing steps needed to produce a certain type of passive device need only be performed if that particular component is needed. An *i*Passives network can essentially be built up with just the processing complexity required—no more and no less. As illustrated in Figure 2, there are a number of passive building blocks to choose from and constructing an integrated passives network can be as simple as piecing together the required components.

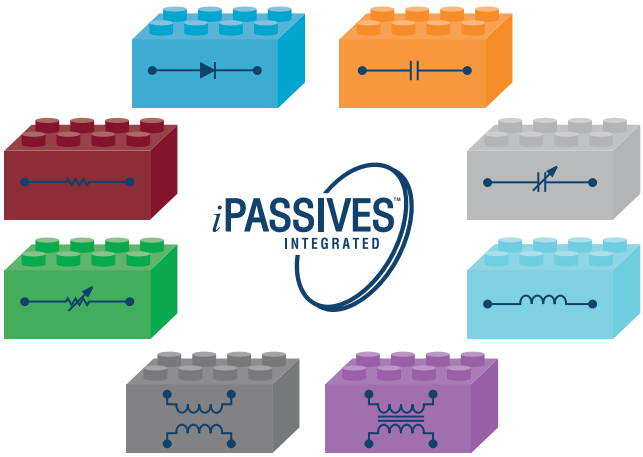


*Figure 2. Building blocks for iPassives.*

As outlined earlier in this article, integrated passives hold many advantages over discrete passives. ADI has brought these advantages a step further by employing their use in µModule devices. These modules leverage off the capabilities of a diverse range of integrated circuits. These circuits are manufactured through tailored processes that provide enhanced performance and that are not achievable from any one single process. ADI is using *i*Passives to tie these integrated circuits together and in doing so are building complete precision signal chains within a single device. The two µModule device examples in Figure 3 contain data converters, amplifiers, and other components, and bringing them together are passive gain and filtering networks that are built with integrated passives.

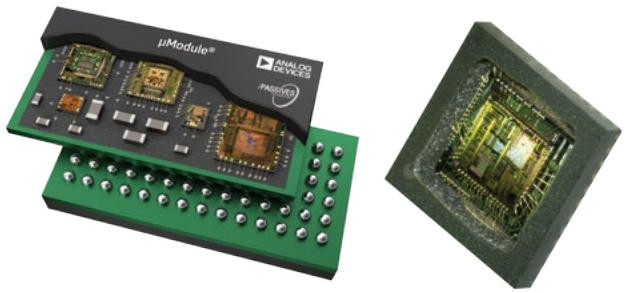


*Figure 3. µModule product examples using iPassives.*

ADI produces precision signal conditioning systems that are highly customizable. By adopting a reusability approach from a huge portfolio of field proven ICs and combining this approach with the versatility of *i*Passives, both the development cycle times and costs are dropping significantly. This decision offers immense benefits to customers who themselves can get to market faster and more efficiently with state-of-the-art performance.

## Conclusion

At first glance, the use of integrated passives may appear only incrementally more advantageous over more established approaches. The advantages, however, turn out to be significant and *i*Passives as employed by ADI are redefining not just what can be done, but also what speeds, costs, and sizes are beneficial to customers.


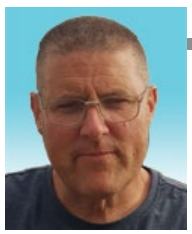

**Mark Murphy**

Mark Murphy [mark.murphy@analog.com] is a product marketing engineer with the Precision Converter Group in Limerick Ireland. He holds a B.S.E.E. from Merrimack College and an M.B.A. from the University of Limerick.

**Pat McGuinness**

Pat McGuinness [pat.mcguinness@analog.com] is a product engineer at Analog Devices with responsibility for the use of integrated passive components in system in package applications. Since joining ADI in 1994, Pat has worked on product yield improvement activities, circuit design, device modeling, and new process development. Throughout his time at ADI he has designed precision converters, switches, sensors, and passive networks. Pat received his bachelor's degree in electrical and microelectronic engineering from University College Cork.

# A Large Current Source with High Accuracy and Fast Settling

By **Nick Jiang**

Share on

Voltage controlled current sources (VCCSs) are widely used in many areas, like medical machine and industrial automation. The dc precision, ac performance, and drive capability of a VCCS are highly important in these applications. This article analyzes the enhanced Howland current source (EHCS) circuit's limitations and shows how to improve it with a composite amplifier topology to implement a ±500 mA current source with high precision and fast settling.

## Enhanced Howland Current Source

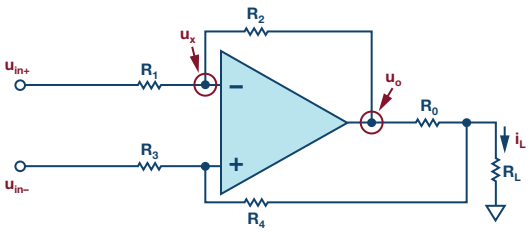


*Figure 1. Howland current source circuit.*

Figure 1 shows the traditional Howland current source (HCS) circuit, while Equation 1 shows how the output current can be calculated. The output current will be constant if R2 is large enough.

$$i_L R_0 \left[1 + \frac{R_1}{R_2} + R_L\left(\frac{1}{R_2} + \frac{1 - \frac{R_1}{R_2} \times \frac{R_4}{R_3}}{R_0}\right)\right] = u_{in+}\left(\frac{R_0}{R_2} + \frac{R_4}{R_3} + 1\right) - u_{in-}\left(\frac{R_4}{R_3} + \frac{R_1}{R_2} \times \frac{R_4}{R_3}\right) \quad (1)$$

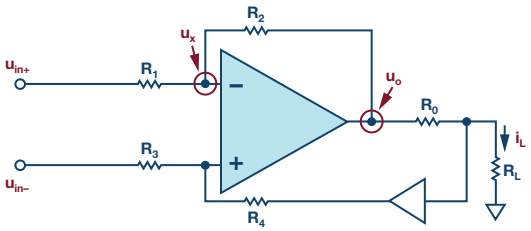


*Figure 2. Enhanced Howland current source circuit.*

While a large R2 will reduce the speed and precision of the circuit, inserting a buffer into the feedback route to form an enhanced Howland current source will eliminate this, as shown in Figure 2. All the current flows through $R_0$ is through into $R_L$. The output current is calculated with Equation 2.

$$i_L R_0 \left[1 + \frac{R_1}{R_2} + R_L\left(1 - \frac{R_1}{R_2} \times \frac{R_4}{R_3}\right)\right] = u_{in+}\left(\frac{R_4}{R_3} + 1\right) - u_{in-}\left(\frac{R_4}{R_3} + \frac{R_1}{R_2} \times \frac{R_4}{R_3}\right) \quad (2)$$

If R1/R2 = R3/R4 = k, the equation is changed to Equation 3. The output current is independent of the load and only controlled by the input voltage. It's an ideal VCCS.

$$i_L = \frac{u_{in+} - u_{in-}}{kR_0} \quad (3)$$

## Performance Analysis

Equation 3 is based on the ideal system. Figure 3 shows the dc error analysis model of the EHCS. $V_{OS}$ and IB+/IB– are the input offset voltage and bias current of the main amplifier. $V_{OSbuf}$ and $I_{Bbuf}$ are the input offset voltage and bias current of the buffer. The total output error can be calculated by Equation 4.

$$\begin{aligned} I_O\left[R_0 + R_L\left(1 - \frac{R_1 + R_2}{R_1} \times \frac{R_3}{R_3 + R_4}\right)\right] = {} & \left(\frac{R_2}{R_1} + 1\right) \times V_{os} + \left(\frac{R_2}{R_1} + 1\right) \times \\ & \frac{R_3 R_4}{R_3 + R_4} \times I_{B+} - \left(\frac{R_2}{R_1} + 1\right) \times \\ & \frac{R_1 R_2}{R_1 + R_2} \times I_{B-} + \left(\frac{R_2}{R_1} + 1\right) \times \\ & \frac{R_3}{R_3 + R_4} \times V_{OSbuf} - I_{Bbuf} R_0 \end{aligned} \quad (4)$$

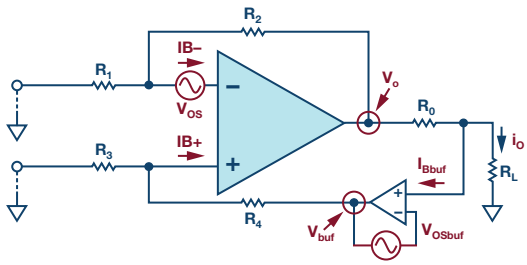


*Figure 3. Offset voltage calculation.*

Ignore the mismatch from the gain resistors, and consider $R_1/R_2 = R_3/R_4 = k$, $R_1//R_2 = R_3//R_4$. The output offset current depends on the amplifiers' offset and bias current, as shown in Equation 5.

$$I_O = \left(\frac{1}{k} + 1\right) \times \frac{1}{R_0} \times V_{os} + \left(\frac{1}{k} + 1\right) \times \frac{1}{R_0} \times \frac{R_1 R_2}{R_1 + R_2} \times I_{os} + \frac{1}{R_0} \times V_{OSbuf} - I_{Bbuf} \quad (5)$$

Taking the mismatch of $R_1/R_2$ and $R_3/R_4$ into consideration, $R_L$ will influence the output offset current. The worst relative error is shown in Equation 6. The error depends on $R_L/R_0$ and k. A smaller load resistor and higher k will decrease the offset error.

$$Max\ Relative\ Error\ of\ I_O = \frac{R_L}{R_0} \times \frac{2\Delta_k}{k(k+1)} \quad (6)$$

We can also calculate the temperature drift of the circuit, which comes from amplifiers and resistors. Amplifiers' offset voltage and bias current change with the work temperature. For most CMOS input amplifiers, the bias current doubles for every increase of 10°C. The drift of resistors changes a lot with different types. For example, carbon composition units' TC is approximately 1500 ppm/°C, while metal film and bulk metal resistors' TC can be 1 ppm/°C.

### Table 1. Precision Amplifiers Parameters

| Devices | $V_{os}$ Max (μV) | IB Max (pA) | GPB (MHz) | Slew Rate (V/μs) | Isc (mA) |
|---|---|---|---|---|---|
| ADA4522 | 5 | 150 | 3 | 1.3 | 22 |
| ADA4077 | 25 | 1500 | 4 | 1 | 22 |
| LTC2057HV | 4 | 120 | 2 | 1.2 | 26 |
| LT1012 | 25 | 100 | 1 | 0.2 | 13 |

Choosing a precision amplifier is good for the dc accuracy of the output current. However, there are many limitations in the precision amplifier selection. The drive capability and ac performance are not good enough. Table 1 lists some common precision amplifiers. We want to build a ±500 mA current source with 1 μs settling time. For a current source we would need high drive capability. For a current source with additional high settling time we need good ac performance. In general, precision amplifiers do not provide that specification combination as the slew rate and bandwidth are not good enough. This requires choosing from a few other amplifiers.

## EHCS Implementation

ADA4870 is a high speed, high voltage, and high drive capacity amplifier. It can supply 10 V to 40 V with 1.2 A output current limitation. Its bandwidth is over 52 MHz for a large signal and the slew rate is up to 2500 V/μs. All these specifications make it the right fit for fast settling and a large current source. Figure 4 shows an EHCS circuit based on the ADA4870 that generates a ±500 mA output current source by 10 V input.

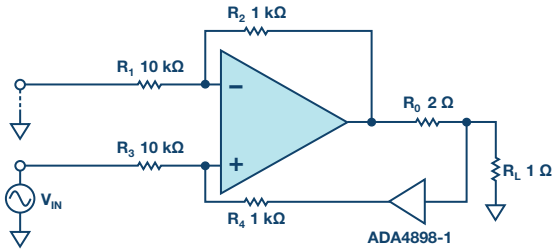


*Figure 4. EHCS circuit based on ADA4870.*

In ac specifications, we are more care about settling time, slew rate, bandwidth, and noise. The settling time is about 60 ns and bandwidth is about 18 MHz as shown in Figure 5. The output current slew rate can be calculated by measuring the slope of the rising and falling stage. The positive and negative slew rate are +25 A/μs and –25 A/μs. The noise performance is shown in the output noise density curve. It's about 24 nV/√Hz at 1 kHz.

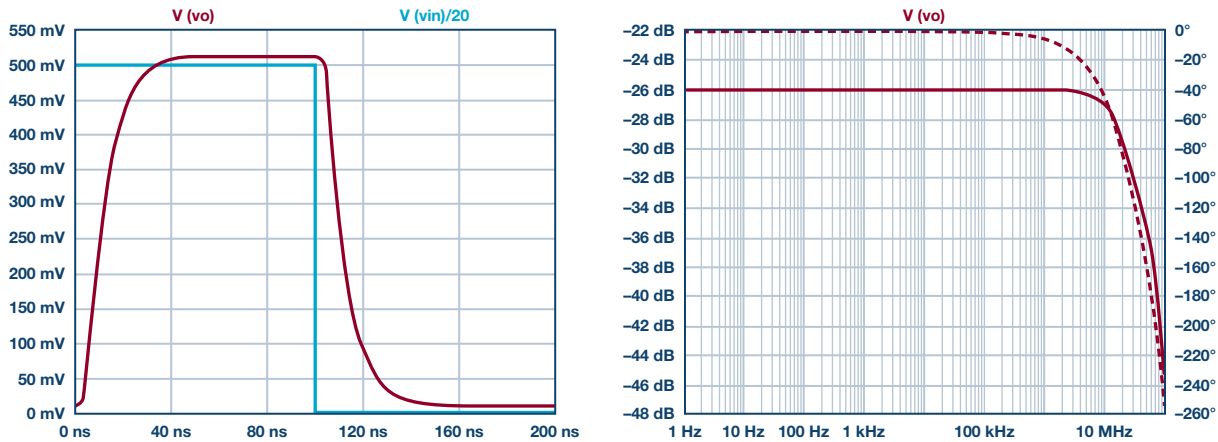


*Figure 5. Settling time and frequency response of an EHCS based on ADA4870*

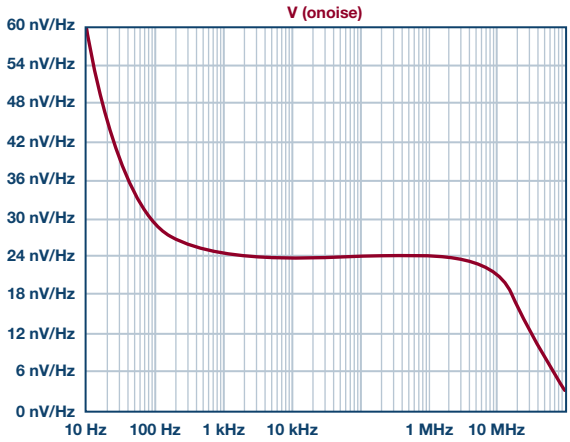


*Figure 6. Output noise density curve of EHCS based on ADA4870.*

Due to large input offset voltage and bias current, the dc precision is not good in this circuit. Table 2 shows different dc error sources and contribution. The main dc error comes from the $V_{os}$ and $I_B$ of ADA4870. The typical output current offset is about 11.06 mA, which is about 2.21% range error referring to 500 mA full range.

### Table 2. The DC Error of EHCS Based on ADA4870

| Error Source | Parameters (Typ) | Error Output (mA) | Percentage |
|---|---|---|---|
| IB | –12 µA | 6.00 | 54.2% |
| IB+ | +9 µA | 4.50 | 40.7% |
| $V_{OS}$ | 1 mV | 0.55 | 5.0% |
| $I_{Bbuf}$ | –0.1 µA | 0.00 | 0.0% |
| $V_{OSbuf}$ | 0.02 mV | 0.01 | 0.1% |
| Total | | 11.06 | 100% |

## Composite Amplifier Technology

High drive amplifiers like ADA4870's dc parameters limit output current accuracy, and high precision amplifiers don't have enough speed. Here we can combine all these qualities into one circuit with composite amplifier technology. Figure 7 shows the composite amplifier enhanced Howland current source (CAEHCS) that is formed by ADA4870 and ADA4898-2.

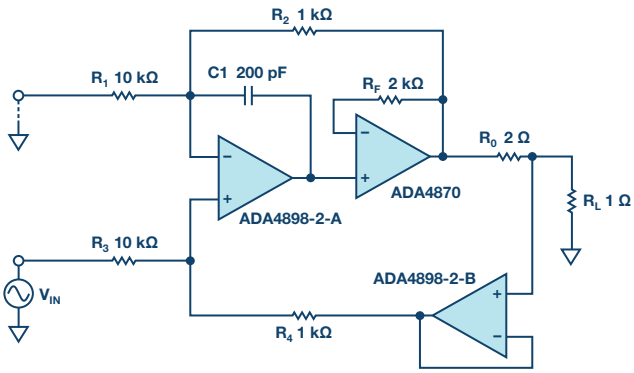


*Figure 7. EHCS circuit with composite amplifier.*

ADA4898-2 is chosen to form the composite amplifier for its excellent ac and dc performance. Its –3 dB bandwidth is 63 MHz. The settling time to 0.1% with a 5 V output step is 90 ns and the slew rate is up to 55 V/µs. It has ultralow noise, too. The voltage noise density is 0.9 nV/√Hz and current noise density is 2.4 pA/√Hz. As for dc specifications, it performs well, too. The typical input offset voltage is 20 µV with 1 µV/°C temperature drift. The bias current is 0.1 µA. Table 3 shows the dc error of the CAEHCS. The output current offset is decreased to 0.121 mA, which means the range error is lower than 0.03%.

### Table 3. The DC Error of CAEHCS Based on ADA4898

| Error Source | Parameters (Typ) | Error Output (mA) | Percentage |
|---|---|---|---|
| IB– | –0.1 µA | 0.050 | 41.3% |
| IB+ | +0.1 µA | 0.0050 | 41.3% |
| $V_{OS}$ | 20 mV | 0.011 | 9.1% |
| $I_{Bbuf}$ | –0.1 µA | 0.000 | 0.1% |
| $V_{OSbuf}$ | 20 µV | 0.01 | 8.2% |
| Total | | 0.121 | 100% |

The ac performance of the CAEHCS is shown in Table 4. The settling time and bandwidth are lower than EHCS due to the loop delay of the composite amplifier. CAEHCS output noise is much lower than EHCS output noise due to low current noise of ADA4898-2. As specified in the data sheet, the ADA4870's invert input current noise density is 47 pA/√Hz. With several kΩ resistors, it will generate much higher noise than the voltage noise (2.1 nV/√Hz). While the input current noise density of the CAEHCS is 2.4 pA/√Hz. It will generate much lower output noise.

### Table 4. The AC Specification of the CAEHCS

| Parameter | CAEHCS | EHCS |
|---|---|---|
| Settling time (ns) | 200 | 60 |
| Slew rate (A/µs) | 7.7 | 25 |
| Bandwidth (MHz) | 6 | 18 |
| Output Noise Density at 1 kHz (nV/√Hz) | 4 | 24 |

Above all, CAEHCS has greatly improved the dc accuracy of the VCCS with comparable drive capacity and ac performance. Besides, there are many selections of the composite amplifiers for different requirements. Table 5 shows the performance of different amplifiers in the CAEHCS circuit. LT6275 is the best in ac performance. Its settling time can be within 100 ns, and the slew rate is up to 15 A/µs. Zero-drift amplifiers like ADA4522-2 are suitable for high precision applications that have about 0.002 mA output current offset error.

### Table 5. Selection of Main Amplifier in CAEHCS

| Main Amplifier | EHCS | CAEHCS |
|---|---|---|
| ADA4898 | Good | Good |
| LT6275 | Good | Excellent |
| ADA4522 | Excellent | Not good |

## Test Results

The performance of the EHCS and CATHCS based on ADA4898 are shown in Table 6 and Figure 8.

### Table 6. Comparison of EHCS vs. CAEHCS

| Main Amplifier | | EHCS | CAEHCS |
|---|---|---|---|
| DC Parameters | Output current offset (mA) | 10.9 | 0.2 |
| AC Parameters | Settling time (ns) | 100 | 100 |
| | Slew rate (A/µs) | 22.2 | 12.6 |
| | Bandwidth (MHz) | 18 | 8 |

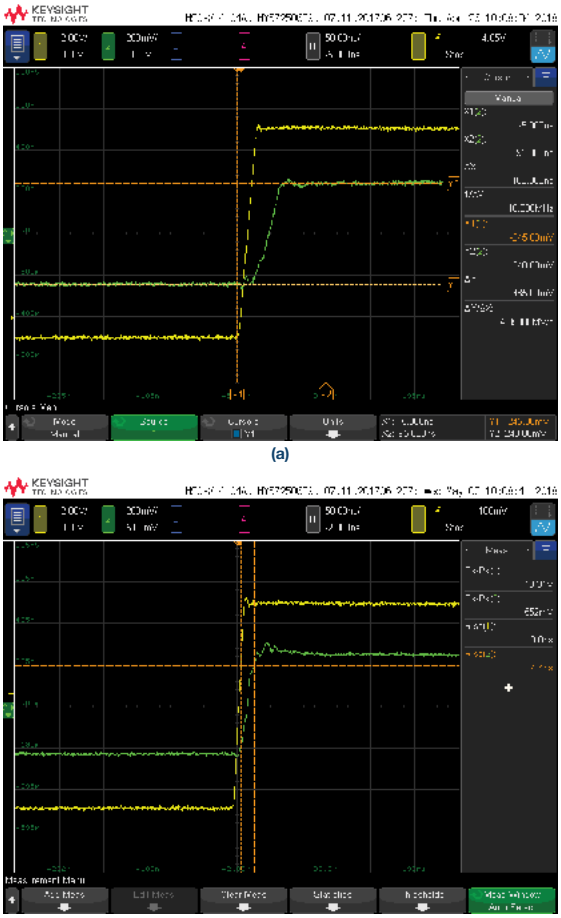

*Figure 8. The settling time of the ADA4898-2 (CH1-Input, CH2-Output).*

The CAEHCS circuit has much better dc specifications than an EHCS circuit. Its output current offset is 0.2 mA, while the EHCS circuit's output current offset is 10.9 mA. The CAEHCS circuit has good ac specifications, too. The settling time both are about 100 ns. The bandwidth of the EHCS circuit is 18 MHz, and the CAEHCS circuit is 8 MHz.

The performance of the CAEHCS based on the ADA4522-2 and LT6275 is shown in Table 7. The output offset error of ADA4522-2 version is lower to 0.04 mA. The settling time of LT6275 version is about 60 ns and the output current slew rate is up to 16.6 A/µs, which is shown in Figure 9.

### Table 7. Test Results of the Different Main Amplifier in CAEHCS

| Main Amplifier | Ios (mA) | Settling time (ns) | Slew Rate (A/µs) | Bandwidth (MHz) |
|---|---|---|---|---|
| ADA4898 | 0.2 | 100 | 12.6 | 10 |
| LT6275 | 0.8 | 60 | 16.6 | 11 |
| ADA4522 | 0.04 | 1000 | 0.4 | 1.2 |

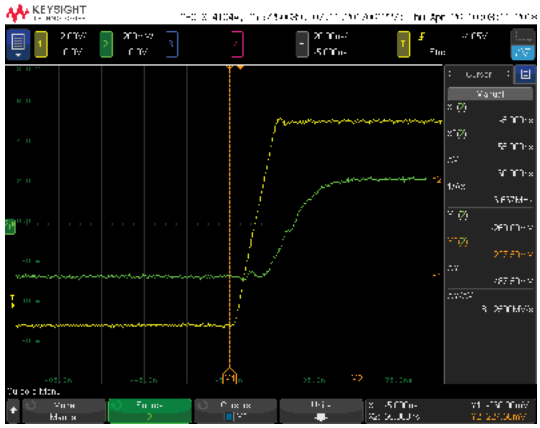

*Figure 9. The settling time of the LT6275 (CH1-Input, CH2-Output).*

## Thermal Consideration

The output current of the VCCS can be several hundred milliamperes. The whole power dissipation can be several watts. If output efficiency is bad, the temperature of the part will rise rapidly. The thermal resistance ($\theta_{JA}$) of ADA4870 without a head sink can be 15.95°C/W. The temperature rising can be calculated by using Equation 7.

$$T_{rise} = \theta_{JA} \times P \quad (7)$$

The value of $R_0$ will influence the power dissipation of ADA4870. Table 8 shows the temperature rise with different $R_0$ selected at a ±20 V supply. The temperature rising will decrease greatly when larger $R_0$ is used. Therefore, larger $R_0$ is recommended to decrease the temperature rise.

### Table 8. ADA4870's Power Dissipation and Temperature Rise vs. RO (Io = 500 mA)

| RL/Ω | Power Dissipation (W) | | Temperature Rise (°C) | |
|---|---|---|---|---|
| | $R_0$ = 2 Ω | $R_0$ = 10 Ω | $R_0$ = 2 Ω | $R_0$ = 10 Ω |
| 1 | 6.92 | 4.92 | 110.4 | 78.5 |
| 5 | 5.92 | 3.92 | 94.5 | 62.6 |
| 10 | 4.67 | 2.67 | 74.6 | 42.7 |

## Conclusion

The CAEHCS circuit that combines a high drive amplifier and a high precision amplifier can provide excellent ac and dc performance with large output capacity in VCCS applications. ADA4870 combined with ADA4898, LT6275, and ADA4522 are recommended for use in this circuit.


Nick Jiang is a product applications engineer in the Analog Devices Linear and Precision Technology Group, and he is based in Beijing. He graduated from Xi'an Jiaotong University, with a bachelor's and master's degree in electrical engineering.

Nick Jiang

# Modeling and Control for a Current-Mode Buck Converter with a Secondary LC Filter

By **Ricky Yang**

Share on

## Introduction

Modern signal processing system design utilizing ADCs, PLLs, and RF transceivers demands lower power consumption and higher system performance. Selecting proper power supplies for those noise sensitive devices is always the pain point for system designers. There is always a trade-off between high efficiency and high performance.

Traditionally, LDO regulators are often used to power those noise sensitive devices. LDO regulators reject the low frequency noise that often presents in system power supplies and they provide clean power to ADCs, PLLs, or RF transceivers. But LDO regulators usually have low efficiency, especially in systems where LDO regulators must regulate down from a power rail several volts above their output voltage. In this kind of situation, LDO regulators typically offer 30% to 50% efficiency, while using switching regulators can reach 90% or even higher efficiency.

Switching regulators are more efficient than LDO regulators, but they are too noisy to directly power ADCs or PLLs without significant performance degradation. One of the noise sources of switching regulators is the output ripple, which can appear as distinct tones or spurs in ADCs output spectrum. To avoid degrading the signal-to-noise ratio (SNR) and spurious-free dynamic range (SFDR), minimizing the output ripple and output noise of switching regulators can be very important.

In order to maintain high efficiency and high system performance at the same time, it is often desirable to add a secondary LC filter ($L_2$ and $C_2$) to the output of switching regulator to reduce the ripple and noise, as Figure 1 shows. However, the two-stage LC output filter has associated disadvantages, too. The power stage transfer function is ideally modeled as a fourth-order system that can easily be unstable. If the sample data effect of a current loop[1] is also taken into consideration, the complete control-to-output transfer function is shown to be fifth order. An alternate solution is to sense the output voltage from the primary LC filter ($L_1$ and $C_1$) point to stabilize the system. However, applying this approach results in poor output voltage regulation due to the large voltage drops on the secondary LC filter when the load current is heavy, which is not acceptable in some applications.

In this article, a new hybrid feedback method is proposed to provide adequate stability margin and maintain the output accuracy over all the load condition in the application where switching regulators with secondary LC filters are used to provide high efficiency, high performance power supplies to ADCs, PLLs, or RF transceivers.

There has been some published research work on the dc-to-dc converter with a secondary LC output filter.[2–5] Specifically, the articles "Control Loop Design for Two-Stage DC-to-DC Converters with Low Voltage/High Current Output" and "Comparative Evaluation of Multiloop Control Schemes for a High Bandwidth AC Power Source with a Two-Stage LC Output Filter" discuss the modeling and control of a two-stage voltage-mode converter, which can't be directly applied to a current-mode converter. A current-mode converter with a secondary LC filter has been analyzed and modeled in "Secondary LC Filter Analysis and Design Techniques for Current-Mode-Controlled Converters" and "Three-Loop Control for Multimodule Converter Systems." However, both articles have an assumption that the secondary inductor has a much smaller inductance value than the primary inductor, which is not always eligible in real applications.

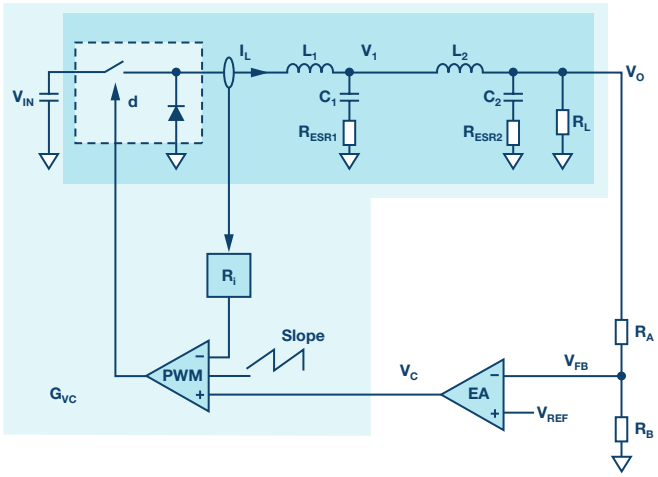


*Figure 1. Circuit diagram of a current-mode buck converter with a secondary LC filter.*

The outline of this article is as follows:

The small signal modeling of a buck converter with a secondary LC filter is analyzed. A new fifth-order control-to-output transfer function is presented, which is very accurate regardless of the peripheral inductor and capacitor parameters.

A new hybrid feedback method is proposed to provide adequate stability margins while maintaining good dc accuracy of the output voltage. The limitation of the feedback parameters has been analyzed for the first time, which can provide basic criteria for practical design.

Based on the power stage small signal model and new hybrid feedback method, the compensation network is designed. The stability of the closed-loop transfer function is evaluated using a Nyquist plot.

A simple design example is presented based on the power management product ADP5014. With the secondary LC filter, the output noise of ADP5014 in a high frequency range is even better than an LDO regulator.

Appendix I and Appendix II present necessary small signal transfer function for power stage and feedback network respectively.

## Small Signal Modeling of Power Stage

Figure 2 shows the small signal block diagram for Figure 1. The control loop is composed by inner current loop and outer voltage loop. The sample data coefficient $H_e(s)$ in the current loop refers to the model proposed by Raymond B. Ridley in "A New, Continuous-Time Model for Current-Mode Control." Note that in the simplified small signal block diagram in Figure 2, the input voltage disturbance and load current disturbance are assumed to be zero since transfer functions related with input voltage and load current will not be discussed in this article.

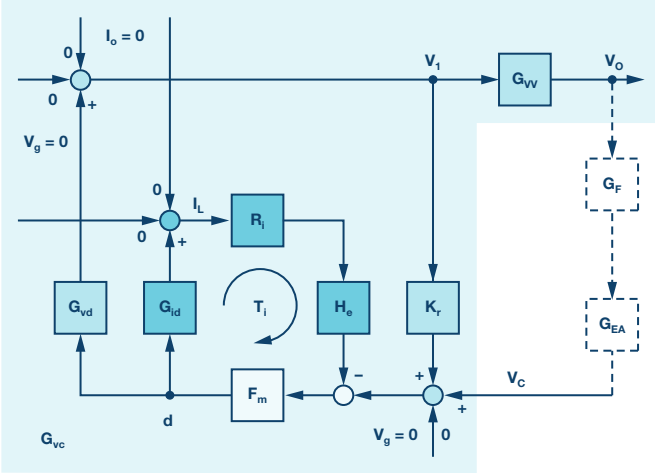


*Figure 2. A small signal block diagram of a current-mode buck converter with a secondary LC filter.*

## Buck Converter Example

The new small signal model is demonstrated with a current-mode buck converter with the following parameters:

- $V_g = 5$ V
- $V_o = 2$ V
- $L_1 = 0.8$ µH
- $L_2 = 0.22$ µH
- $C_1 = 47$ µF
- $C_2 = 3\times 47$ µF
- $R_{ESR1} = 2$ mΩ
- $R_{ESR2} = 2$ mΩ
- $R_L = 1$ Ω
- $R_i = 0.1$ Ω
- $T_s = 0.833$ µs

## Current-Loop Gain

The first transfer function of interest is the current-loop gain measured at the output of the duty-cycle modulator. The resulting current loop transfer function (see Equation 16 in Appendix I) exhibits a fourth-order system with two pairs of complex conjugate poles, which results in two system resonances ($\omega_1$ and $\omega_2$). Both of these two resonance frequency are determined by $L_1$, $L_2$, $C_1$, and $C_2$. A domain zero is contributed by load resistor $R_L$, $C_1$, and $C_2$. One pair of complex conjugate zeros ($\omega_3$) is determined by $L_2$, $C_1$, and $C_2$. Besides, the sample data coefficient $H_e(s)$ in the current loop will introduce a complex pair of right half plane (RHP) zeros at half of the switching frequency.

Compared with the conventional current-mode buck converter without the secondary LC filter, the new current-loop gain has one more pair of complex conjugate poles and one more pair of complex conjugate zeros, which locate very close to each other.

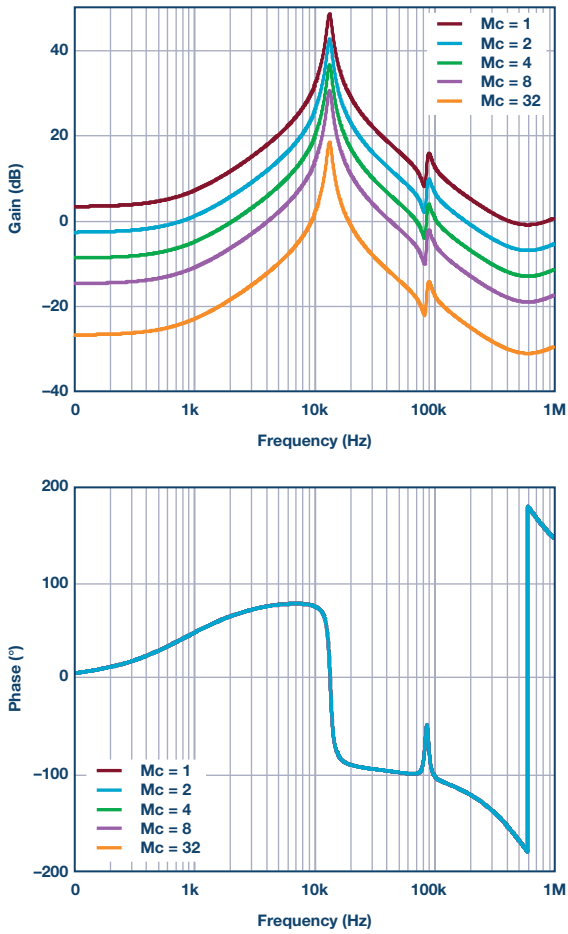


*Figure 3. Buck converter current-loop gain.*

Figure 3 shows a plot of the current-loop gain with different values of external ramp. For the case without external slope compensation ($M_c = 1$), it can be seen that there is very little phase margin in the current loop, which may lead to subharmonic oscillation. With added external slope compensation, the shape of the gain and phase curves do not change, but the amplitude of the gain will decrease and phase margin will increase.

## Control-to-Output Gain

A new control-to-output transfer function is created when the current loop is closed. The resulting control-to-output transfer function (see Equation 19 in Appendix I) exhibits a fifth-order system with one domain pole ($\omega_p$) and two pairs of complex conjugate poles ($\omega_l$ and $\omega_h$). The domain pole is mainly determined by load resistor $R_L$, $C_1$, and $C_2$. The lower frequency pair of conjugate poles is determined by $L_2$, $C_1$, and $C_2$, while the higher frequency pair of conjugate poles locates at half of the switching frequency. Additionally, two zeros are contributed by the *ESR* of $C_1$ and *ESR* of $C_2$, respectively.

Figure 4 shows a plot of the control-to-output loop gain with different values of external ramp. Compared with the conventional current-mode buck converter, there is one more pair complex conjugate poles ($\omega_l$) in the control-to-output gain of current-mode buck converter with a secondary LC filter. The additional resonant poles will give up to 180° additional phase delay. The phase margin drops dramatically, and it can make the system unstable even with Type III compensation. Besides, Figure 4 clearly shows the transition from current-mode to voltage-mode control as the slope compensation is increased.

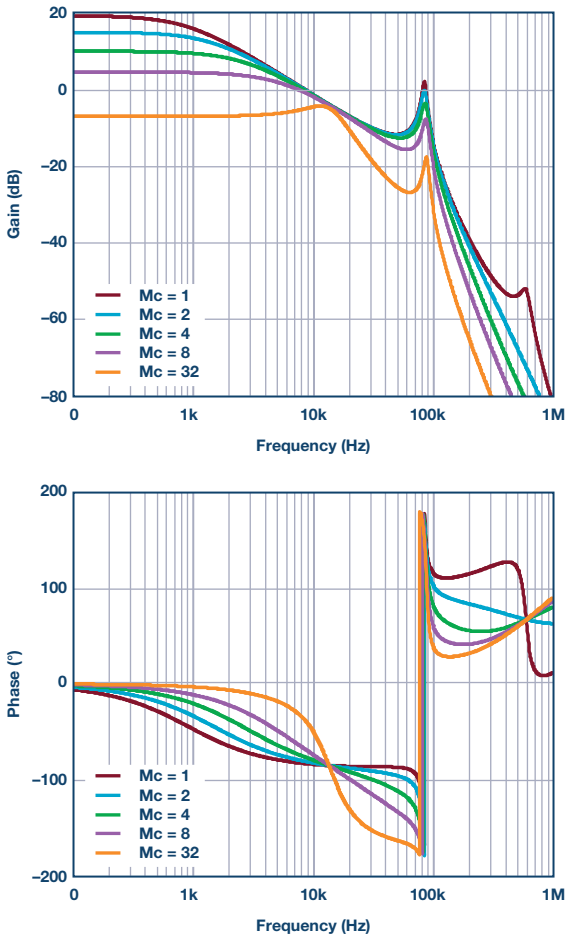


*Figure 4. Control-to-output transfer function for a buck converter*

## The Hybrid Feedback Method

This article presents a new hybrid feedback structure, as Figure 5(a) shows. The idea of hybrid feedback is to stabilize the control loop by using an additional capacitor feedback from the primary LC filter. The outer voltage feedback from the output through resistor divider is defined as the remote voltage feedback and the inner voltage feedback though capacitor $C_F$ will be referred to as the local voltage feedback hereafter. The remote feedback and local feedback carry different information on the frequency domain. Specifically, the remote feedback senses the low frequency signal to provide good dc regulation of the output, while the local feedback senses the high frequency signal to provide good ac stability for the system. Figure 5(b) shows the simplified small signal block diagram for Figure 5(a).

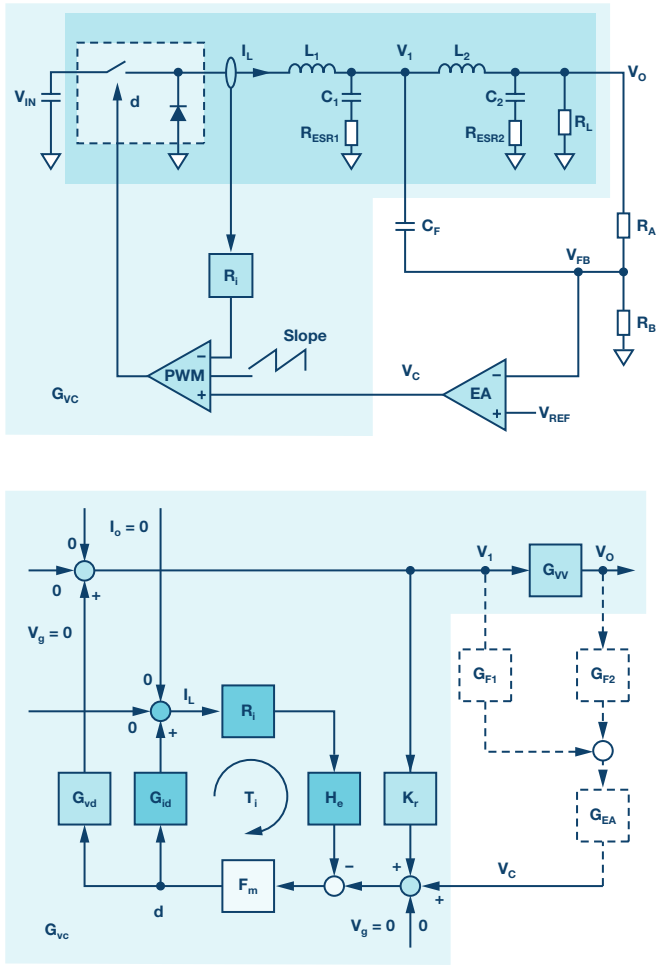


*Figure 5. A current-mode buck converter with a proposed hybrid feedback method, showing (a) a circuit diagram and (b) a small signal model.*

## The Feedback Network Transfer Function

The resulting equivalent transfer function (see Equation 31 and Equation 32 in Appendix II) of a hybrid feedback structure differs significantly from the transfer function of conventional resistor divider feedback. The new hybrid feedback transfer function has more zeros than poles, and the additional zeros will lead to 180° phase ahead at the resonant frequency determined by $L_2$ and $C_2$. Therefore, with the hybrid feedback method, the additional phase delay in control-to-output transfer function will be compensated for by the additional zeros in the feedback transfer function, which will facilitate the compensation design based on the complete control-to-feedback transfer function.

## The Limitation of Feedback Parameters

Apart from those parameters in the power stage, there are two more parameters in the feedback transfer function. Parameter β (see Equation 30 in Appendix II) is the output voltage magnification ratio, which is already well-known. However, the parameter α is a brand new concept.

The feedback parameter α (see Equation 29 in Appendix II) can be adjusted to understand the behavior of the feedback transfer function. Figure 6 exhibits the change trend of the zeros in the feedback transfer when α is decreased. It clearly shows that one pair of conjugate zeros will be pushed from left half plane (LHP) to right half plane (RHP) with decreased α.

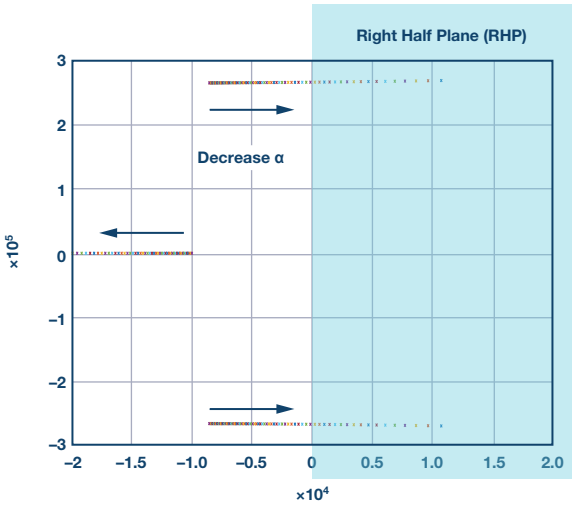


*Figure 6. Effect of feedback parameter α on the zeros of the feedback network.*

Figure 7 is a plot of the feedback transfer function with a different α. It shows that when α is decreased to $10^{-6}$ (for example: $R_A$ = 10k, $C_F$ = 1 nF), the transfer function of the feedback network will exhibit 180° phase delay, which means the complex zeros have become RHP zeros. The feedback transfer function has been simplified to a new form (see Equation 33 in Appendix II). To keep the zeros in the LHP, the parameter α should always meet the following condition:

$$\alpha > \frac{L_2 \times C_2}{\frac{L_2}{R_L} + R_{ESR2} \times C_2} \qquad \text{(Formula 1)}$$

Formula 1 gives a minimum limitation basis for feedback parameter α. As long as the condition is satisfied, the control system will be easily stable. However, since $R_A$ and $C_F$ will work as an RC filter of output voltage change during a load transient, the load transient performance will be degraded with a very big α. So α should not be too large. In practical design, the parameter α is recommended to be 20% to ~30% bigger than the minimum limitation value.

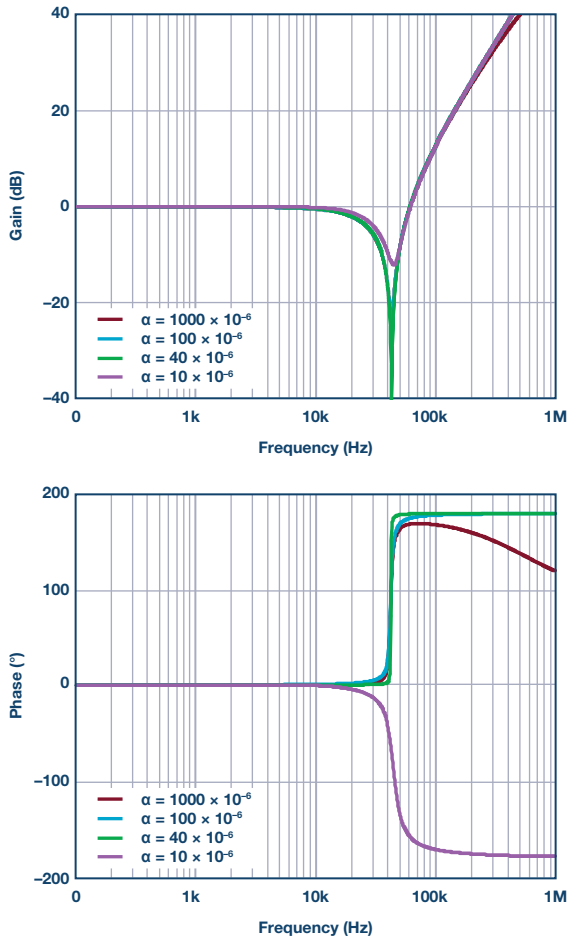


*Figure 7. The transfer function of a proposed hybrid feedback network with a different parameter α.*

## Loop Compensation Design

### Design the compensation

The control-to-feedback transfer function GP(s) can be derived by the product of the control-to-output transfer function $G_{vc}(s)$ and the feedback transfer function $G_{FB}(s)$. The compensation transfer function $G_C(s)$ is designed to have one zero and one pole. The asymptotic Bode plots of the control-to-feedback and compensation transfer function, as well as closed-loop transfer function $T_V(s)$, are shown in Figure 8. The following procedures show how to design the compensation transfer function.

Determine the cross frequency ($f_c$). Since the bandwidth is limit by $f_{z1}$, choosing an $f_c$ smaller than $f_{z1}$ is recommended.

Calculate the gain of $G_P(s)$ at fc, then the middle frequency band gain of $G_C(s)$should be the opposite number of $G_P(s)$.

Place the compensation zero at the domain pole ($f_{p1}$) of the power stage.

Place the compensation pole at the zero ($f_{z2}$) caused by the ESR of output capacitor $C_1$.

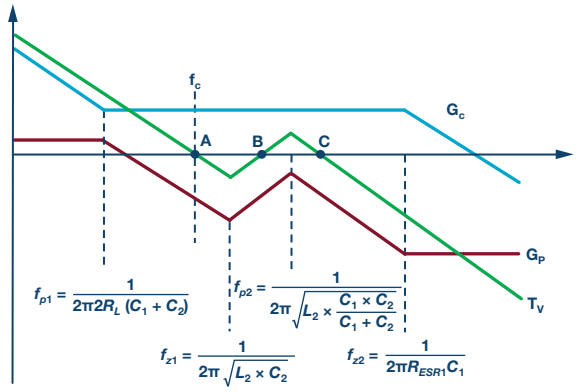


*Figure 8. A loop gain design based on a proposed control-to-output and hybrid feedback transfer function.*

## Using a Nyquist Plot to Analyze Stability

According to Figure 8, the closed-loop transfer function $T_V(s)$ has crossed 0 dB three times. The Nyquist plot is used to analyze the stability of closed-loop transfer function, as Figure 9 shows. Since the plot is far away from (–1, j0), the closed loop is stable and has adequate phase margin. Note that the points A, B, and C in the Nyquist plot correspond to the points A, B, and C in the Bode plot.

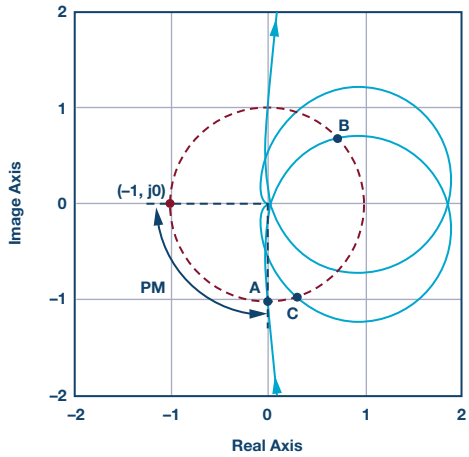


*Figure 9. Nyquist plot of the closed-loop transfer function.*

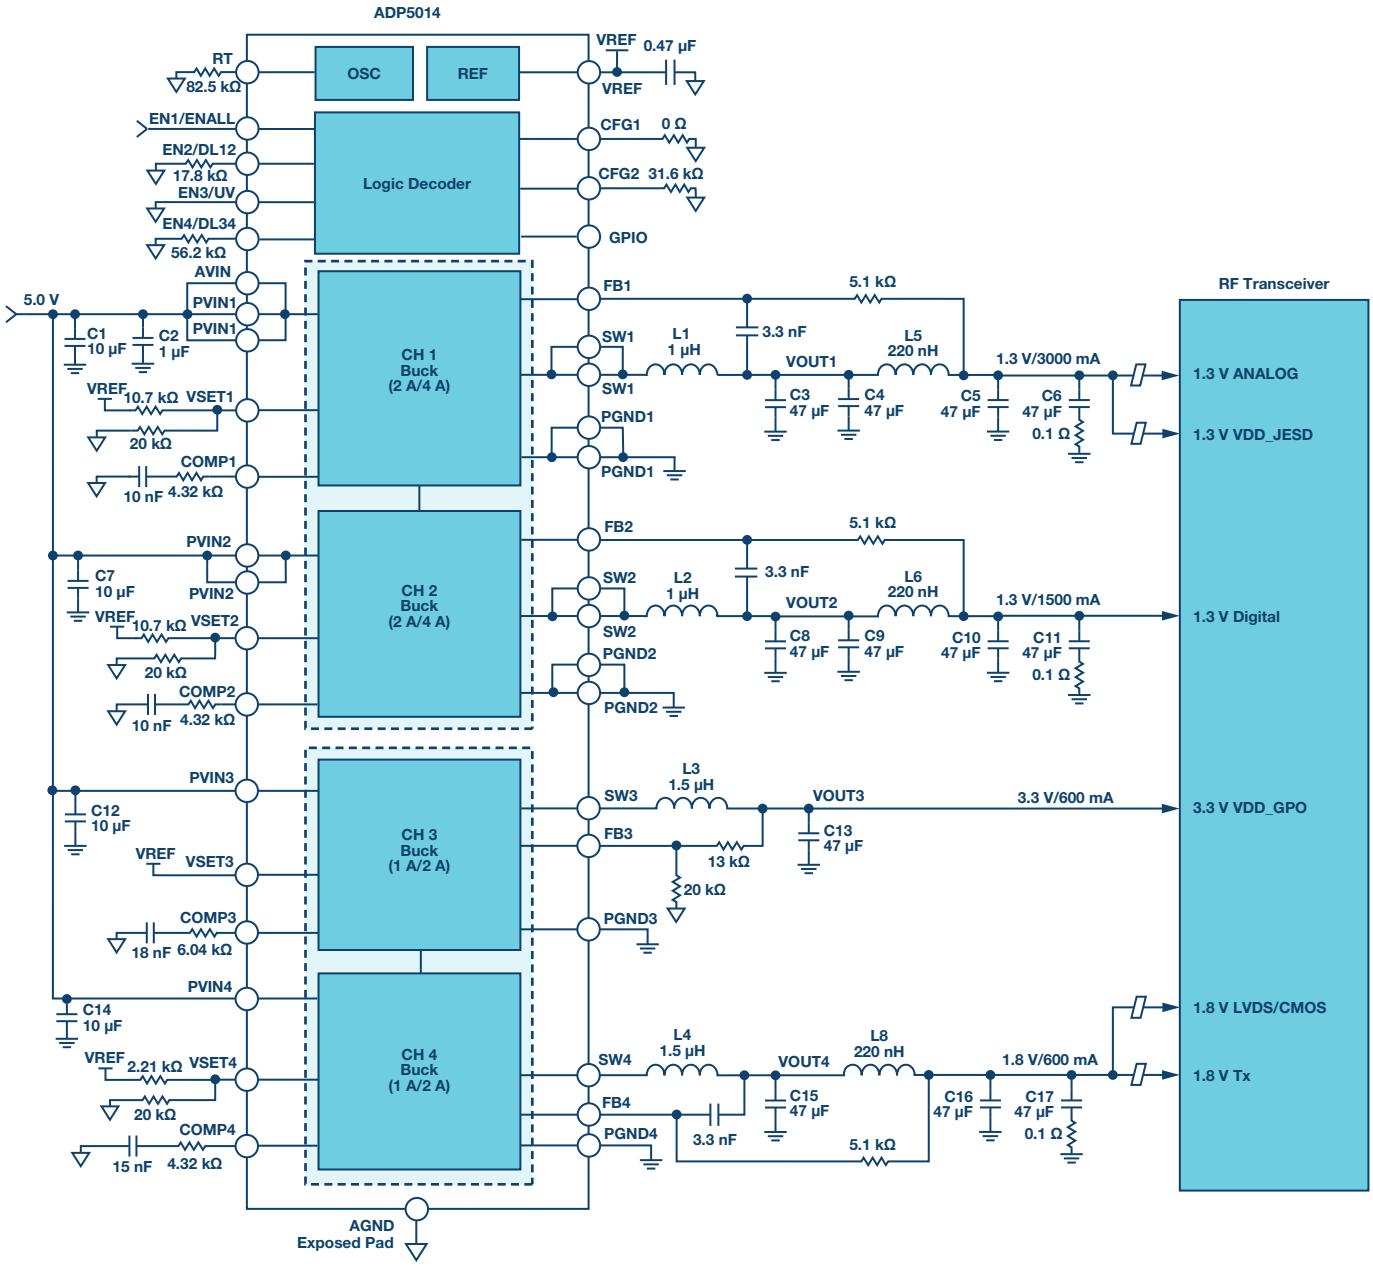


*Figure 10. RF transceiver powered up by ADP5014 with a secondary LC filter.*

## Design Example

The ADP5014 optimizes many analog blocks to achieve lower output noise at a low frequency range. The unit-gain voltage reference structure also makes its output noise independent from the output voltage setting when $V_{OUT}$ setting is less than the $V_{REF}$ voltage. A secondary LC filter is added to attenuate the output noise at a high frequency range, especially for fundamental switching ripple and its harmonic. Figure 10 shows the design details.

Figure 11 shows the ADP5014 noise spectral density measurement from a 10 Hz to 10 MHz frequency range and integrated rms noise from a 10 Hz to 1 MHz frequency range, compared to the ADP1740s as another traditional, 2 A, low noise LDO regulator. The output noise of ADP5014 in the high frequency range is even better than ADP1740's.

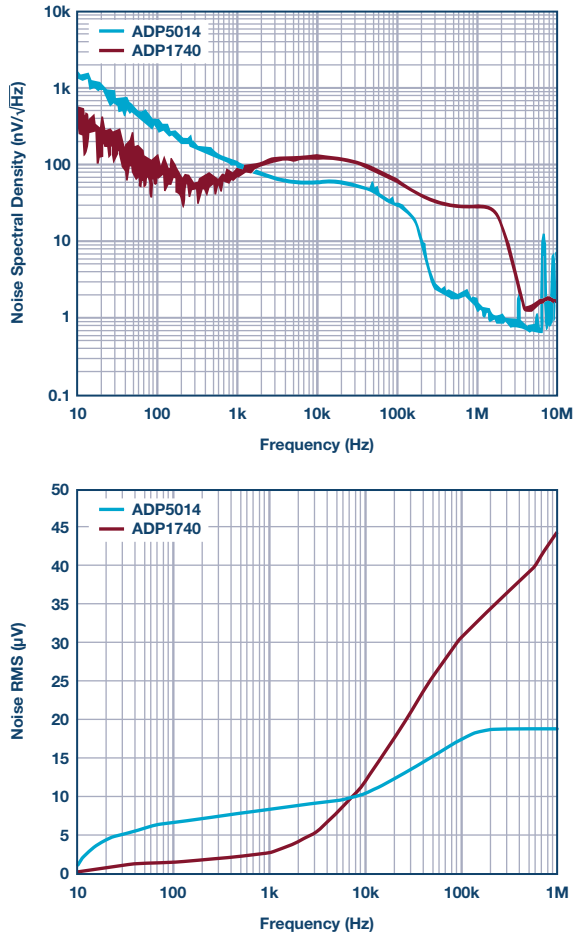


*Figure 11. A comparison of output noise performance between ADP5014 and ADP1740 showing (a) noise spectral density and (b) integrated rms noise.*

## Conclusion

This article presents a general analysis framework for modeling and control for a current-mode buck converter with a secondary-stage LC output filter. The accurate control-to-output transfer function is discussed. A new hybrid feedback structure is proposed and the feedback parameter limitation is deduced.

The design example shows that a switching regulator with secondary LC filter and hybrid feedback method can provide a clean, stable power supply that is competitive with, or even better than, an LDO regulator.

Modeling and control in this article concentrates on a current-mode buck converter, but the methods described here can be applied to voltage-mode buck converters as well.

## Appendix 1

The power stage transfer functions in Figure 2 are as follows.

$$G_{id} = \frac{i_L}{d} = \frac{V_g}{R_L} \times \frac{[1 + R_L \times (C_1 + C_2) \times S] \times \left(1 + \frac{S}{Q_3 \times \omega_3} + \frac{S^2}{\omega_3{}^2}\right)}{\left(1 + \frac{S}{Q_1 \times \omega_1} + \frac{S^2}{\omega_1{}^2}\right) \times \left(1 + \frac{S}{Q_2 \times \omega 2} + \frac{S^2}{\omega_2{}^2}\right)} \quad (1)$$

$$G_{vd} = \frac{v_1}{d} = \frac{V_g \times (1 + R_{ESR1} \times C_1 \times S) \times \left(1 + \frac{S}{Q_4 \times \omega_4} + \frac{S^2}{\omega_4{}^2}\right)}{\left(1 + \frac{S}{Q_1 \times \omega_1} + \frac{S^2}{\omega_1{}^2}\right) \times \left(1 + \frac{S}{Q_2 \times \omega 2} + \frac{S^2}{\omega_2{}^2}\right)} \quad (2)$$

$$G_{vv} = \frac{v_o}{v_1} = \frac{1 + R_{ESR2} \times C_2 \times S}{1 + \frac{S}{Q_4 \times \omega_4} + \frac{S^2}{\omega_4{}^2}} \quad (3)$$

where:

$$\omega_1 = \frac{1}{\sqrt{L_1 \times (C_1 + C_2) + L_2 \times C_2}} \quad (4)$$

$$Q_1 = \frac{1}{\omega_1 \times \left(\frac{L_1}{R_L} + \frac{L_2}{R_L} + R_{ESR1} \times C_1 + R_{ESR2} \times C_2\right)} \quad (5)$$

$$\omega_2 = \frac{1}{\sqrt{\frac{L_1 \times L_2 \times C_1 \times C_2}{L_1 \times (C_1 + C_2) + L_2 \times C_2}}} \quad (6)$$

$$Q_2 = \frac{1}{\omega_2 \times \left(\frac{(L_1 + L_2) \times C_1 \times C_2 \times R_{ESR1} + \frac{L_1 \times L_2 \times C_1}{R_L} + L_1 \times C_1 \times C_2 \times R_{ESR2}}{L_1 \times C_1 + L_1 \times C_2 + L_2 \times C_2}\right)} \quad (7)$$

$$\omega_3 = \frac{1}{\sqrt{L_2 \times \frac{C_1 \times C_2}{C_1 + C_2}}} \quad (8)$$

$$Q_3 = \frac{1}{\omega_3 \times \frac{L_2 \times C_1 + R_{ESR1} \times R_L \times C_1 \times C_2 + R_{ESR2} \times R_L \times C_1 \times C_2}{R_L \times (C_1 + C_2)}} \quad (9)$$

$$\omega_4 = \frac{1}{\sqrt{L_2 \times C_2}} \quad (10)$$

$$Q_4 = \frac{1}{\omega_4 \times \left(\frac{L_2}{R_L} + R_{ESR2} \times C_2\right)} \quad (11)$$

where: $L_1$ is the primary inductance.

$C_1$ is the primary capacitance.

$R_{ESR1}$ is the equivalent series resistance of the primary capacitor.

$L_2$ is the secondary inductance.

$C_2$ is the secondary capacitance.

$R_{ESR2}$ is the equivalent series resistance of the secondary capacitor. $R_L$ is the load resistance.

The gain blocks in the current loop are as follows.

$$F_m = \frac{1}{m_c \times S_n \times T_s} \quad (12)$$

$$H_e = 1 - \frac{T_s}{2} \times S + \frac{S}{\omega_h{}^2} \quad (13)$$

where:

$$m_c = 1 + \frac{S_e}{S_n} \quad (14)$$

$$\omega_h = \frac{\pi}{T_s} \quad (15)$$

where: $R_i$ is the equivalent current sense resistance

$S_e$ is the sawtooth ramp of slope compensation

$S_n$ is on-time slope of the current sense waveform

$T_s$ is the switching period

The current-loop gain is

$$T_i(S) = F_m \times R_i \times G_{id} \times H_e = \frac{L_1}{R_L \times D' \times m_c \times T_s} \times \frac{[1 + R_L \times (C_1 + C_2) \times S] \times \left(1 + \frac{S}{Q_3 \times \omega_3} + \frac{S^2}{\omega_3{}^2}\right)}{\left(1 + \frac{S}{Q_1 \times \omega_1} + \frac{S^2}{\omega_1{}^2}\right) \times \left(1 + \frac{S}{Q_2 \times \omega 2} + \frac{S^2}{\omega_2{}^2}\right)} \times H_e \quad (16)$$

where:

$$D' = 1 - D \quad (17)$$

where:

D is the duty cycle

According to Figure 2, the gain block $k_r$ is given by

$$k_r = \frac{R_i \times T_s}{2 \times L_1} \quad (18)$$

The control-to-output transfer function is

$$G_{vc} = \frac{v_o}{v_c} = \frac{R_L}{R_i} \times \frac{1}{1 + \frac{R_L \times T_s}{L_1} \times (mc \times D' - 0.5)} \times F_l(s) \times F_h(s) \quad (19)$$

where:

$$F_l(s) = \frac{(1 + R_{ESR1} \times C_1 \times S) \times (1 + R_{ESR2} \times C_2 \times S)}{\left(1 + \frac{S}{\omega_p}\right) \times \left(1 + \frac{S}{Q_1 \times \omega_1} + \frac{S^2}{\omega_1{}^2}\right)} \quad (20)$$

$$F_h(s) = \frac{1}{1 + \frac{S}{Q_h \times \omega_h} + \frac{S^2}{\omega_h{}^2}} \quad (21)$$

$$\omega_p = \frac{1 + \frac{R_L \times T_S}{L_1} \times (mc \times D' - 0.5)}{R_L \times (C_1 + C_2)} \quad (22)$$

$$\omega_1 = \frac{1}{\sqrt{L_2 \times \frac{C_1 \times C_2}{C_1 + C_2}}} \quad (23)$$

$$Q_1 = \frac{1}{\omega_1 \times \frac{L_2 \times C_1 + (R_{ESR1} + R_{ESR2}) \times R_L \times C_1 \times C_2}{R_L \times (C_1 + C_2)}} \quad (24)$$

$$Q_h = \frac{1}{\pi \times [m_c \times (1 - D) - 0.5]} \quad (25)$$

## Appendix II

In Figure 5, the local feedback and remote feedback transfer function are

$$G_{F1} = \frac{R_A \times C_F \times S}{1 + \frac{R_A}{R_B} + R_A \times C_F \times S} \quad (26)$$

$$G_{F2} = \frac{1}{1 + \frac{R_A}{R_B} + R_A \times C_F \times S} \quad (27)$$

According to Equation 1 through Equation 27, the control-to-feedback transfer function is given by

$$G_P = \frac{v_{FB}}{v_C} = \frac{1 + (\alpha + R_{ESR2} \times C_2) \times S + \alpha \times \left(\frac{L_2}{R_L} + R_{ESR2} \times C_2\right) \times S^2 + \alpha \times L_2 \times C_2 \times S^3}{(\beta + \alpha \times S) \times (1 + R_{ESR2} \times C_2 \times S)} \times G_{VC} \quad (28)$$

where:

$$\alpha = R_A \times C_F \quad (29)$$

$$\beta = 1 + \frac{R_A}{R_B} \quad (30)$$

where: $R_A$ is the top resistor of feedback resistor divider

$R_B$ is the bottom resistor of feedback resistor divider

$C_F$ is the local feedback capacitor

The equivalent feedback network transfer function is

$$G_{FB} = \frac{v_{FB}}{v_o} = \frac{1 + (\alpha + R_{ESR2} \times C_2) \times S + \alpha \times \left(\frac{L_2}{R_L} + R_{ESR2} \times C_2\right) \times S^2 + \alpha \times L_2 \times C_2 \times S^3}{(\beta + \alpha \times S) \times (1 + R_{ESR2} \times C_2 \times S)} \quad (31)$$

The approximate feedback transfer function is

$$G_{FB\_appr} = \frac{1 + \alpha \times S}{(\beta + \alpha \times S) \times (1 + R_{ESR2} \times C_2 \times S)} \times \Delta \quad (32)$$

where:

$$\Delta = 1 + \left(\frac{L_2}{R_L} + R_{ESR2} \times C - \frac{L_2 \times C_2}{\alpha}\right) \times S + L_2 \times C_2 \times S^2 \quad (33)$$

In typical low noise applications, the unit-gain voltage reference structure is usually applied, so parameter β will be equal to 1. Then the feedback transfer function is

$$G_{FB} = \frac{\Delta}{1 + R_{ESR2} \times C_2 \times S} \quad (34)$$

## References


[1]Raymond B. Ridley. "A New, Continuous-Time Model for Current-Mode Control." *IEEE Transactions on Power Electronics*, Vol. 6, No. 2, 1991.

[2]Julie Yixuan Zhu and Brad Lehman. "Control Loop Design for Two-Stage DC-to-DC Converters with Low Voltage/High Current Output." *IEEE Transactions on Power Electronics*, Vol 20, No. 1, 2005.

[3]Patricio Cortes, David O. Boillat, Hans Ertl, and Johann W. Kolar. "Comparative Evaluation of Multiloop Control Schemes for a High Bandwidth AC Power Source with a Two-Stage LC Output Filter." International Conference on Renewable Energy Research and Applications, IEEE, 2013.

[4]Raymond B. Ridley. "Secondary LC Filter Analysis and Design Techniques for Current-Mode-Controlled Converters." *IEEE Transactions on Power Electronics*, Vol. 3, No. 4, 1988.

[5]Byungcho Choi, Bo H. Ch, Fred C. Lee, and Raymond B. Ridley. "Three-Loop Control for Multimodule Converter Systems." *Power Electronics IEEE Transactions on Power Electronics*, Vol. 8, No. 4, 1993.



Ricky Yang [ricky.yang@analog.com] joined Analog Devices in April 2008 and works as an applications engineer for the Power Management Group based in Shanghai, China. He graduated from Shanghai Jiaotong University with a bachelor's degree in information and controlling engineering and a master's degree in power electronics. He has 10 years of experience in application work for various switching regulators and related system design.


Ricky Yang

# Rarely Asked Questions—Issue 158 Driving a Unipolar Gate Driver in a Bipolar Way

By **Ryan Schnell**

Share on

## Question:

Do you need a specialized gate driver to deliver positive and negative voltages?

## Answer:

No, you can adapt a unipolar gate driver to drive in a bipolar manner.

If a positive and negative gate drive is required for a particular power device, circuit designers don't need to search for a special gate driver that handles bipolar operation specifically. Use this simple trick to make a unipolar gate driver deliver bipolar voltages!

When driving medium to high powered MOSFETs and IGBTs, there is a risk of a Miller effect induced turn-on when a high rate of change of voltage is seen across the power device. Current is injected onto the gate of a power device through the gate-to-drain or gate-to-collector capacitance. If the current injection is large enough to bring the gate voltage above the device threshold voltage, parasitic turn-on can be observed resulting in lower efficiency or even device failure.

The Miller effect can be mitigated through use of a very low impedance path from the power device gate to the source or drain, or by driving the gate to a negative voltage with respect to the source or drain. The goal of the Miller effect turn-on mitigation techniques is to keep the gate voltage below a desired threshold when a current spike through the Miller capacitance occurs.

Certain power device types even require a negative voltage to be fully off, necessitating some kind of negative voltage drive coming from the gate driver. Device manufacturers that recommend negative gate drive voltage include standard silicon MOSFETs, IGBTs, SiC, and GaN devices.

There is a wide selection of isolated gate drivers that operate on a unipolar power supply on the secondary side (the side driving the power device), but much fewer gate driver devices that allow for explicit bipolar voltage drives. One method to overcome this lack of negative gate drive devices is to offset the gate driver from the power device, thereby creating a negative gate drive relative to the source or drain of the power device, while the gate driver IC still only sees a unipolar supply. Unipolar and bipolar gate drive waveform examples are shown in Figure 1.

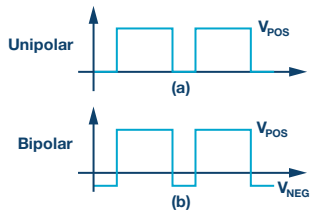


*Figure 1.(a) Unipolar and (b) bipolar gate drive waveforms.*

A schematic with ideal voltage sources pictured is shown in Figure 2. In this example, the driver IC is powered by a voltage equal to the sum of $V_1$ and $V_2$, while the gate of the MOSFET is driven to a $+V_1$ in the ON state and a $–V_2$ in the OFF state, relative to the MOSFET source node. Note that in this example, both voltage sources are decoupled with individual capacitors. The effective decoupling seen by the gate driver IC is the series combination of the capacitors, which is less than the value of each individual capacitor. Additional decoupling can be added between $V_{DD}$ and GND if desired, but it is important to keep $C_1$ and $C_2$ as the capacitors provide low impedance paths for the gate current during turn-on and turn-off separately.

Isolated gate driver ICs often come with an undervoltage lockout (UVLO) to prevent a power device from being driven weakly if the gate driver is being driven with too low of a gate voltage. When driving a unipolar gate driver as shown in Figure 2, care must be taken with the expected operation of the UVLO as the UVLO is usually referenced to the ground of the gate driver. Consider a case where $V_1 = 15$ V, $V_2 = 9$ V, and the gate driver UVLO is around 11 V, which is common for IGBT operation. If $V_1$ were to drop more than 4 V, the UVLO would not trigger, but the IGBT would be driven under 11 V during the ON time, thereby underdriving the IGBT.

Creating two separate voltage sources for this purpose can be accomplished by using two isolated power supplies, but cost is often a concern for this approach. If a flyback topology is used, multiple winding taps could be used to obtain multiple voltages relatively easily.

There are isolated voltage source modules that can provide isolated power, and some manufacturers are selecting voltages conducive to power device voltages. One example is RECOM, with devices such as the IGBT targeted product line that produces an isolated +15 V and –9 V rail.

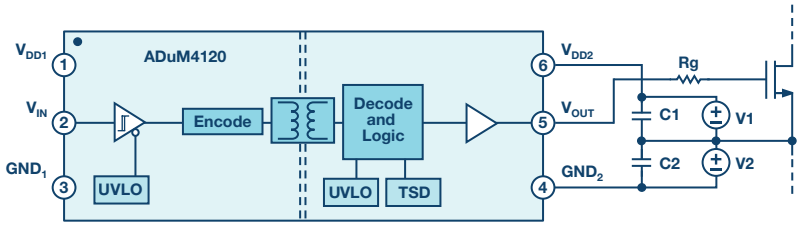


*Figure 2. Example bipolar supply setup.*

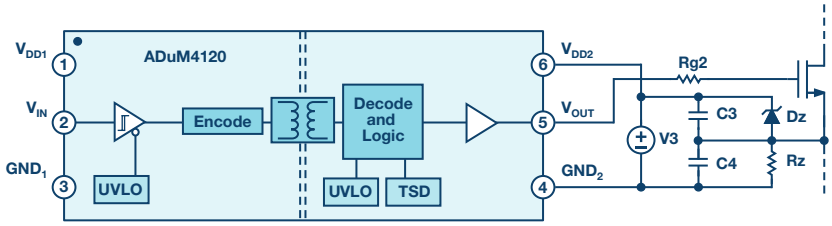


*Figure 3. Zener diode example.*

For such a large voltage swing, the gate driver must be able to withstand a larger range than the range at which other devices were targeted. Two gate drivers that work well with these voltages are ADI's ADuM4135 and ADuM4136 IGBT gate drivers with *i*Coupler® technology, which have a recommended voltage range that allows up to 30 V. Both provide a dedicated ground pin on the output side, allowing the driver UVLO to be referenced against the positive supply rail. The ADuM4135 also includes an integrated Miller Clamp, which can further help suppress the Miller induced turn-on gate voltage bump.

A simple method for creating a bipolar supply from a single voltage source is to create a second voltage source using a biased Zener diode. Although gate drivers provide high currents during the turn-on and turn-off of a power device, the average current actually needed from the power supply is relatively low—often in the tens of mA range for most applications.

The Zener diode can be placed to either regulate the positive or the negative voltage, and can be selected based on which rail needs higher accuracy. The example shown in Figure 3 is setup to regulate the positive voltage more than the negative voltage. One reason to regulate the positive voltage could be if the gate being driven has a tight tolerance on the gate voltage requirements, such as in the case of some GaN devices. Regulating the positive supply also has the added benefit of allowing the UVLO of the gate driver to act as expected, since any fluctuation in $V_3$ will be attenuated by the Zener diode until $V_3$ is too low to support the Zener voltage.

Using the Zener diode method to create two supplies out of a single supply also has the benefit of layout savings. Not only does a Zener diode and a resistor effectively replace an entire isolated voltage source, but by using a unipolar isolated gate driver, a six pin device can be used, such as ADI's ADuM4120 with *i*Coupler technology—saving even more space around the gate driver IC along the isolated creepage area.

A reference example of the Zener diode bipolar setup was created using ADI's ADuM4121 and GaN Systems' GS66508T to create a half bridge. The example was designed to have a +5 V and –4 V drive referenced to the device source. The example could easily be adapted to have +6 V and –3 V drive by using a different Zener diode, and the same 9 V isolated power supply. A large deadtime is used to separate the Miller bump from other turn-off transients visually, but in practice the ADuM4121 allows for much shorter deadtimes in the tens of ns range, which is an important metric for high efficiency GaN designs.

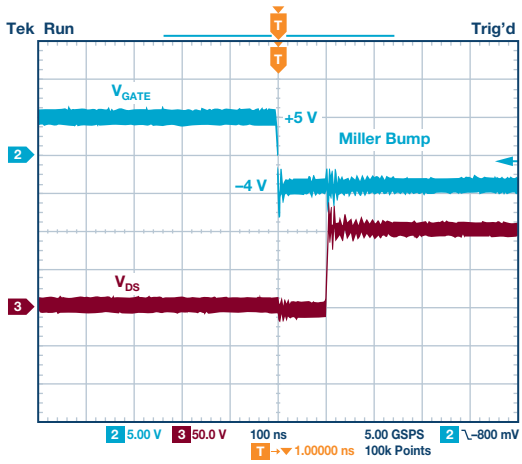


*Figure 4. ADuM4121 and GS66508T experimental results.*

Creating a negative gate volt drive that can mitigate Miller effect parasitic turn-on does not have to be complicated. Many existing gate drivers that are unipolar in operation can be operated to easily drive a gate negative with minimal external circuitry. There are some implications to consider, such as the effective UVLO voltage, but the benefit of such operation is great.


Ryan Schnell [ryan.schnell@analog.com] is an applications engineer at Analog Devices. His responsibilities include isolated gate drivers that use *i*Coupler technology to achieve isolation, as well as various power management products. He holds a B.S. and an M.S in electrical engineering, and a Ph.D. in power electronics from the University of Colorado.


**Ryan Schnell**

# Over-the-Air (OTA) Updates in Embedded Microcontroller Applications: Design Trade-Offs and Lessons Learned

By **Benjamin Bucklin Brown**

Share on


## Abstract

Many embedded systems are deployed in places that are difficult or impractical for a human operator to access. This is especially true for Internet of Things (IoT) applications, which are typically deployed in larger quantities and with limited battery life. Some examples would be embedded systems that monitor the health of a person or a machine. These challenges, coupled with the rapid software lifecycle, cause many systems to require support for over-the-air (OTA) updates. An OTA update replaces the software on the microcontroller or microprocessor of the embedded system with new software. While many people are very familiar with OTA updates on their mobile devices, the design and implementation on a resource constrained system leads to many different challenges. In this article, we will describe several different software designs for OTA updates and discuss their trade-offs. We will see how hardware features of two ultra low power microcontrollers can be leveraged in OTA update software.


## Building Blocks

### Server and Client

An OTA update replaces the current software on a device with new software, with the new software being downloaded wirelessly. In an embedded system, the device that runs this software is typically a microcontroller. A microcontroller is a small computing device with limited memory, speed, and power consumption. A microcontroller typically contains a microprocessor (core) as well as digital hardware blocks for specific operations (peripherals). Ultra low power microcontrollers that typically consume 30 µA/MHz to 40 µA/MHz in active mode are ideal for this type of application. Using specific hardware peripherals on these microcontrollers and placing them into low power modes is an important part of the OTA update software design. An example of an embedded system that might require OTA updates is shown in Figure 1. Here we see a microcontroller connected with a radio and sensor, which may be used in an IoT application that gathers data about the environment using the sensor and reports it periodically using the radio. This portion of the system is referred to as the edge node or client and is the target of the OTA update. The other portion of the system is referred to as the cloud or server and is the provider of the new software. The server and client communicate over a wireless connection using transceivers (radios).

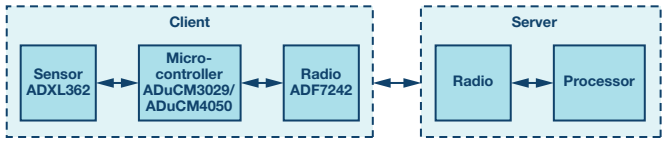


*Figure 1. Server/client architecture in an example embedded system.*

## What Makes a Software Application?

Much of the OTA update process is the act of transferring the new software from the server to the client. The software is transferred as a sequence of bytes, after it has been converted into a binary format from the source format. The conversion process compiles the source code files (for example, *c, cpp*), links them together into an executable file (for example, *exe, elf*), and then the executable is converted into a portable binary file format (for example, *bin, hex*). At a high level, these file formats contain a sequence of bytes that belong at a specific address of memory in the microcontroller. Typically, we conceptualize the information being sent over a wireless link as data, such as a command to change the system's state or sensor data collected by the system. In the case of the OTA update, the data is the new software in binary format. In many cases, the binary file will be too large to send in a single transfer from the server to the client, meaning that the binary file will need to be placed into separate packets, in a process called packetizing. To visualize this process better, Figure 2 demonstrates how different versions of the software will produce different binary files, and thus different packets to be sent during the OTA update. In this simple example, each packet contains 8 bytes of data, with the first 4 bytes representing the address in the client's memory to store the next 4 bytes.

## Major Challenges

Based on this high level description of the OTA update process, three major challenges arise that the OTA update solution must address. The first challenge relates to memory. The software solution must organize the new software application into volatile or nonvolatile memory of the client device so that it can be executed when the update process completes. The solution must ensure that a previous version of the software is kept as a fallback application in case the new software has problems. Also, we must retain the state of the client device between resets and power cycles, such as the version of the software we are currently running, and where

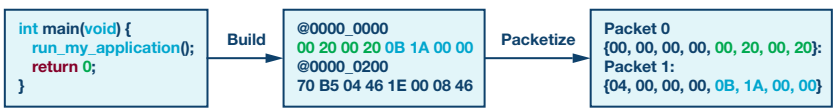


*Figure 2. Binary conversion and packetization process of a software application.*

it is in memory. The second major challenge is communication. The new software must be sent from the server to the client in discrete packets, each targeting a specific address in the client's memory. The scheme for packetizing, the packet structure, and the protocol used to transfer the data must all be accounted for in the software design. The final major challenge is security. With the new software being sent wirelessly from the server to the client, we must ensure that the server is a trusted party. This security challenge is known as authentication. We also must ensure that the new software is obfuscated to any observers, since it may contain sensitive information. This security challenge is known as confidentiality. The final element of security is integrity, ensuring that the new software is not corrupted when it is sent over the air.

## The Second-Stage Boot Loader (SSBL)

### Understanding the Boot Sequence

The primary boot loader is a software application that permanently resides on the microcontroller in read-only memory. The region of memory the primary boot loader resides in is known as info space and is sometimes not accessible to users. This application executes every time a reset occurs, generally performing some essential hardware initializations, and may load user software into memory. However, if the microcontroller contains on-chip nonvolatile memory, like flash memory, the boot loader does not need to do any loading and simply transfers control to the program in flash memory. If the primary boot loader does not have any support for OTA updates, it is necessary to have a second-stage boot loader. Like the primary boot loader, the SSBL will run every time a reset occurs, but will implement a portion of the OTA update process. This boot sequence is illustrated in Figure 3. In this section, we will describe why a second-stage boot loader is necessary and describe how specifying the role of this application is a key design trade-off.

### Lesson Learned: Always Have an SSBL

Conceptually, it may seem simpler to omit the SSBL and place all the OTA update functionality into the user application, as it would allow an existing software framework, operating system, and device drivers to be seamlessly leveraged for the OTA process. The memory map and boot sequence of a system that chose this approach is illustrated in Figure 4.

Application A is the original application that is deployed on the microcontroller in the field. This application contains the OTA update-related software, which is leveraged to download Application B when requested by the server. After this download is complete and Application B has been verified, Application A will transfer control to Application B by performing a branch instruction to the reset handler of Application B. The reset handler is a small piece of code that is the entry point of the software application and runs on reset. In this case, the reset is mimicked by performing a branch, which is equivalent to a function call. There are two major issues with this approach:

- Many embedded software applications employ a real-time operating system (RTOS), which allows the software to be split into concurrent tasks, each with different responsibilities in the system. For instance, the application presented in Figure 1 may have RTOS tasks for reading the sensor, running an algorithm on the sensor data, and interfacing with the radio. The RTOS itself is always active and is responsible for switching between these tasks based on asynchronous events or specific time-based delays. As a result, it is not safe to branch to a new program from an RTOS task, since other tasks will remain running in the background. The only safe way to terminate a program with a real-time operating system is through a reset.

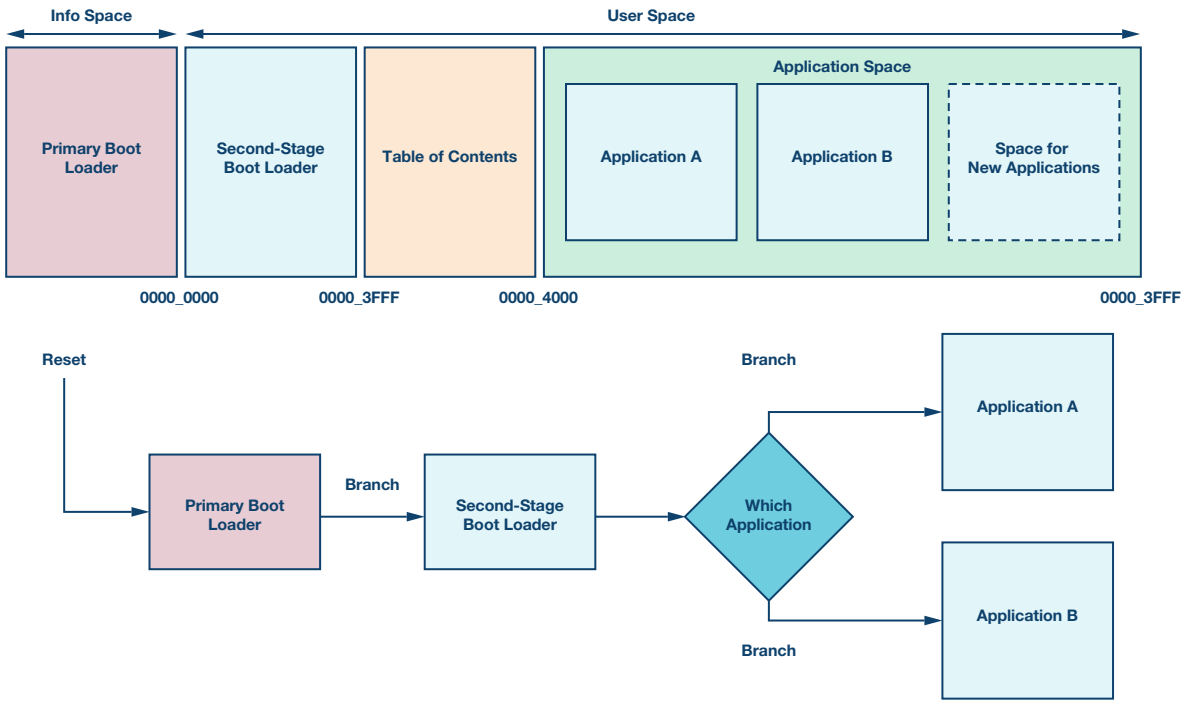


*Figure 3. An example of a memory map and boot flow with SSBL.*

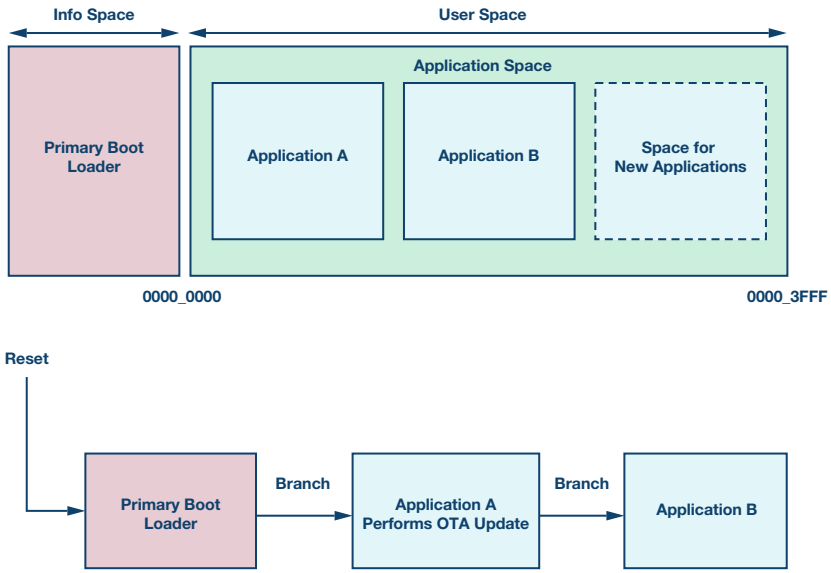


*Figure 4. Example memory map and boot flow without SSBL*

- Based on Figure 4, a solution to the previous issue would be to have the primary boot loader branch to Application B instead of Application A. However, on some microcontrollers, the primary boot loader always runs the program that has its interrupt vector table (IVT), a key portion of the application that describes interrupt handling functions, located at address 0. This means that some form of IVT relocation is necessary to have a reset map to Application B. If a power cycle occurs during this IVT relocation, it could leave the system in a permanently broken state.

These issues are mitigated by having an SSBL fixed at address 0, as illustrated in Figure 3. Since the SSBL is a non-RTOS program, it can safely branch to a new application. There is no concern of a power cycle placing the system in a catastrophic state since the IVT of the SSBL at address 0 is never relocated.

## Design Trade-Off: The Role of the SSBL

We've spent a lot of time discussing the SSBL and the relationship it has with the application software, but what does this SSBL program do? At the bare minimum, the program must determine what the current application is (where it begins) and then branch to that address. The location of the various applications in the microcontroller memory is generally kept in a table of contents (ToC) as shown in Figure 3. This is a shared region of persistent memory that both the SSBL and application software use to communicate with each other. When the OTA update process completes, the ToC is updated with the new application information. Portions of the OTA update functionality can also be pushed to the SSBL. Deciding what portions is an important design decision when developing OTA update software. The minimal SSBL described above will be extremely simple, easy to verify, and most likely will not require modifications during the life of the application. However, this means that each application must be responsible for downloading and verifying the next application. This can lead to code duplication in terms of the radio stack, device firmware, and OTA update software. On the other hand, we can choose to push the entire OTA update process to the SSBL. In this scenario, applications simply set a flag in the ToC to request an update and then perform a reset. The SSBL then performs the download sequence and verification process. This will minimize code duplication and simplify the application specific software. However, this introduces a new challenge of potentially having to update the SSBL itself (that is, updating the update code). In the end, deciding what functionality to place in the SSBL will depend on the memory constraints of the client device, the similarity between downloaded applications, and the portability of the OTA update software.

## Design Trade-Off: Caching and Compression

Another key design decision in the OTA update software will be how to organize the incoming application in memory during the OTA update process. The two types of memory that are typically found on a microcontroller are nonvolatile memory (for example, flash memory) and volatile memory (for example, SRAM). The flash memory will be used to store the program code and read-only data of an application, along with other system-level data such as the ToC and an event log. The SRAM will be used to store modifiable portions of the software application, such as nonconstant global variables and the stack. The software application binary illustrated in Figure 2 only contains the portion of the program that lives in nonvolatile memory. The application will initialize the portions that belong in volatile memory during a startup routine.

During the OTA update process, every time the client device receives a packet from the server containing a portion of the binary it will be stored in SRAM. This packet could be either compressed or uncompressed. The benefit of compressing the application binary is that it will be smaller in size, allowing for fewer packets to be sent and less space needed in SRAM to store them during the download procedure. The disadvantage of this approach is the extra processing time that the compression and decompression add to the update process, along with having to bundle compression related code in the OTA update software.

Since the new application software belongs in flash memory but arrives into SRAM during the update process, the OTA update software will need to perform a write to flash memory at some point during the update process. Temporarily storing the new application in SRAM is called caching. At a high level, there are three different approaches the OTA update software could take to caching.

- No caching: Every time a packet arrives containing a portion of the new application, write it to its destination in flash memory. This scheme is extremely simple and will minimize the amount of logic in the OTA update software, but it requires that the region of flash memory for the new application is fully erased. This method wears down the flash memory and adds overhead.
- Partial caching: Reserve a region of SRAM for caching, and when new packets arrive store them in that region. When the region fills up, empty it by writing the data to flash memory. This can get complex if packets arrive out of order or there are gaps in the new application binary, since a method of mapping SRAM addresses to flash addresses is required. One strategy is to have the cache act as a mirror of a portion of flash

memory. Flash memory is divided into small regions known as pages, which are the smallest division for erasing. Because of this natural division, a good approach is to cache one page of flash memory in SRAM and when it fills up or the next packet belongs in a different page, flush the cache by writing that page flash memory.

- Full caching: Store the entire new application in SRAM during the OTA update process and only write it to flash memory when it has been fully downloaded from the server. This approach overcomes the shortcomings of the previous approaches by minimizing the number of writes to flash memory and avoiding complex caching logic in the OTA update software. However, this will place a limit on the size of the new application being downloaded, since the amount of available SRAM on the system is typically much smaller than the amount of available flash memory.

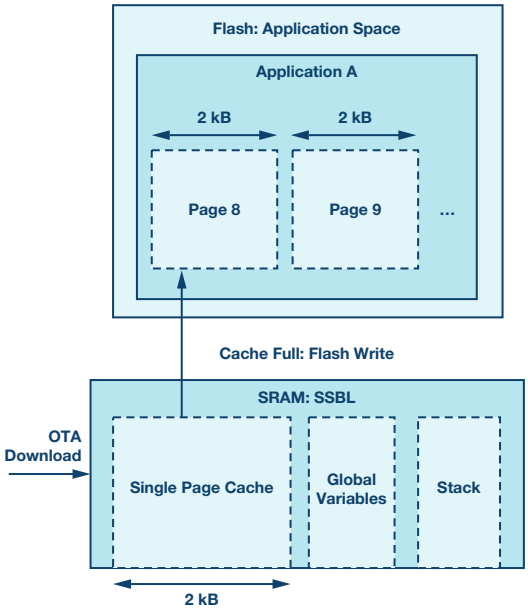


*Figure 5. Using SRAM to one page of cache flash memory.*

The second scheme of partial caching during an OTA update is illustrated in Figure 5, where the portion of flash memory for Application A from Figures 3 and 4 is magnified and a functional memory map of the SRAM for the SSBL is illustrated. An example flash page size of 2 kB is shown. Ultimately this design decision will be determined based on the size of the new application and the allowed complexity of the OTA update software.

## Security and Communication

### Design Trade-Off: Software vs. Protocol

The OTA update solution must also address security and communication. Many systems like the one shown in Figure 1 will have a communication protocol implemented in hardware and software for normal (non-OTA update related) system behavior like exchanging sensor data. This means that there is a method of (possibly secure) wireless communication already established between the server and the client. Communication protocols that an embedded system like Figure 1 might use would be, for example, Bluetooth® Low Energy (BLE) or 6LoWPAN. Sometimes these protocols have support for security and data exchange that the OTA update software may be able to leverage during the OTA update process.

The amount of communication functionality that must be built into the OTA update software will ultimately be determined by how much abstraction is provided by the existing communication protocol. The existing communication protocol has facilities for sending and receiving files between the server and client that the OTA update software can simply leverage for the download process. However, if the communication protocol is more primitive and only has facilities for sending raw data, the OTA update software may need to perform packetizing and provide metadata along with the new application binary. This also applies to the security challenges. The onus may be on the OTA update software to decrypt the bytes being sent over the air for confidentiality if the communication protocol does not support this.

In conclusion, building facilities like custom packet structure, server/client synchronization, encryption, and key exchange into the OTA update software will be determined based on what the system's communication protocol provides and what the requirements are for security and robustness. In the next section, we will propose a complete security solution that solves all the challenges introduced earlier and we will show how to leverage a microcontroller's cryptographic hardware peripheral in this solution.

### Solving Security Challenges

Our security solution needs to keep the new application sent over-the-air confidential, detect any corruption in the new application, and verify that the new application was sent from a trusted server as opposed to a malicious party. These challenges can be solved using cryptographic (crypto) operations. Specifically, two cryptographic operations known as encryption and hashing can be used in the security solution. Encryption will use a shared key (password) between the client and server to obfuscate the data being sent wirelessly. A specific type of encryption that the microcontroller's crypto hardware accelerator may support is called AES-128 or AES-256, depending on the key size. Along with the encrypted data, the server can send a digest to ensure that there is no corruption. The digest is generated by hashing the data packet—an irreversible mathematical function that generates a unique code. If any part of the message or digest is modified after the server creates them, like a bit being flipped during wireless communication, the client will notice this modification when it performs the same hash function on the data packet and compares the digests. A specific type of hashing that the microcontroller's crypto hardware accelerator may support is SHA-256. Figure 6 shows a block diagram of a crypto hardware peripheral in the microcontroller, with the OTA update software residing in the Cortex®-M4 application layer. This figure also shows the support for protected key storage in the peripheral, which can be leveraged in the OTA update software solution to safely store the client's keys.

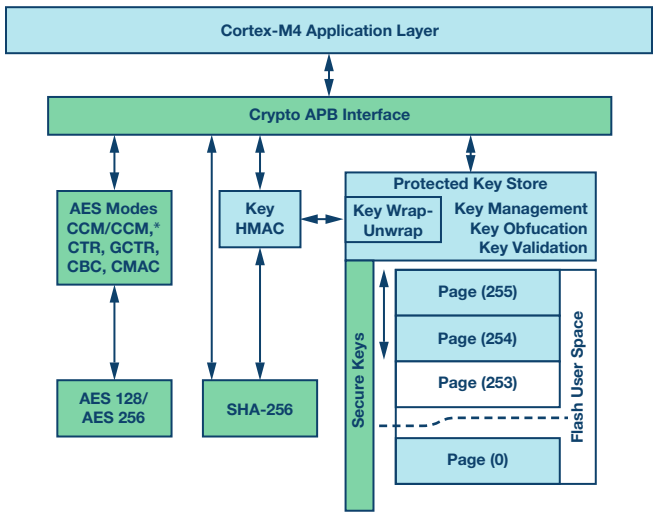


*Figure 6. Hardware block diagram of the crypto accelerator on the ADuCM4050.*

A common technique to solve the final challenge of authentication is to use asymmetric encryption. For this operation, the server generates a public-private key pair. The private key is known only by the server and the public key is known by the client. Using the private key, the server can generate a signature of a given block of data—like the digest of the packet that will be sent over the air. The signature is sent to the client, who can verify the signature using the public key. This enables the client to confirm the message was sent from the server and not a rogue third-party. This sequence is illustrated in Figure 7, with solid arrows representing function input/output and dashed arrows showing the information that is sent over the air.

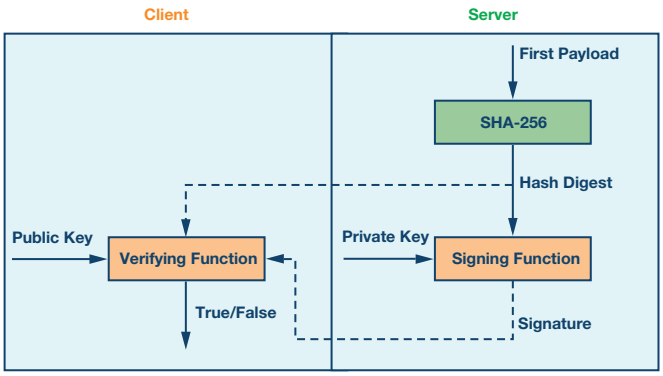


*Figure 7. Using asymmetric encryption to authenticate a message.*

Most microcontrollers do not have hardware accelerators for these asymmetric encryption operations, but they can be implemented using software libraries such as Micro-ECC, which specifically targets resource constrained devices. The library requires a user-defined random number generating function, which can be implemented using the true random number generator hardware peripheral on the microcontroller. While these asymmetric encryption operations solve the trust challenge during an OTA update, they are costly in terms of processing time and require a signature to be sent with the data, which increases packet sizes. We could perform this check once at the end of the download, using a digest of the final packet or the digest of the entire new software application, but that would allow third-parties to download untrusted software to the client, which is not ideal. Ideally, we want to verify every packet that we receive is from our trusted server without the overhead of a signature each time. This can be achieved using a hash chain.

A hash chain incorporates the cryptographic concepts we have discussed in this section into a series of packets to tie them together mathematically. As Figure 8 shows, the first packet (number 0) contains the digest of the next packet. Instead of the actual software application data, the payload of the first packet is the signature. The second packet (number 1) payload contains a portion of the binary, and the digest of the third packet (number 2). The client verifies the signature in the first packet and caches the digest, H0, for later use. When the second packet arrives, the client hashes the payload and compares it to H0. If they match, the client can be sure that this subsequent packet was from the trusted server without all the overhead of doing a signature check. The expensive task of generating this chain is left to the server, and the client must simply cache and hash as each packet arrives to ensure packets arrive uncorrupted, with integrity, and authenticated.

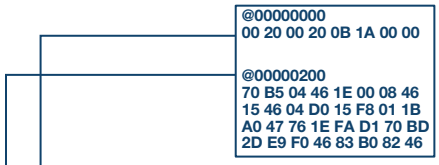


| Number | Payload | Hash |
|---|---|---|
| 0 | P0 = Signature(H0, Public Key) | H0 = H(P1 + H1) |
| 1 | P1 = 00 20 00 20 0B 1A 00 00 | H1 = H(P2 + H2) |
| 2 | P2 = 70 B5 04 46 1E 00 08 46 | H2 = H(P3 + H3) |
| n | Pn = … | Hn = H(Pn+1 + Hn+1) |

*Figure 8. Applying the hash chain to a packet sequence.*

## Experimental Setup

The ultra low power microcontrollers that solve the memory, communication, and security design challenges mentioned in this article are the ADuCM3029 and ADuCM4050. These microcontrollers contain the hardware peripherals discussed throughout the article for OTA updates such as flash memory, SRAM, crypto accelerators, and a true random number generator. The device family packs (DFPs) for these microcontrollers provide software support for building an OTA update solution on these devices. The DFP contains peripheral drivers that provide simple, flexible interfaces for using the hardware.

### Hardware Configuration

To verify and validate the concepts discussed here, an OTA update software reference design was created using the ADuCM4050. For the client, an ADuCM4050 EZ-KIT® is connected to an ADF7242 using the transceiver daughter board horseshoe connector. The client device is pictured on the left of Figure 9. For the server, a Python application was developed that runs on a Windows PC. The Python application communicates over the serial port to another ADuCM4050 EZ-KIT that also has an ADF7242 attached in the same arrangement as the client. However, the right EZ-KIT in Figure 9 performs no OTA update logic, and simply relays packets received from the ADF7242 to the Python application.

*Figure 9. Experimental hardware setup.*

### Software Components

The software reference design partitions the flash memory of the client device as shown in Figure 3. The main client application was designed to be very portable and configurable such that it could be leveraged in other arrangements or on other hardware platforms. Figure 10 shows the software architecture of the client device. Note that while we sometimes

refer to this entire application as the SSBL, in Figure 10 and from now on we logically separate the true SSBL portion (in blue) from the OTA update portion (in red), as the latter is not necessarily required to be implemented entirely in the same application as previously discussed. The hardware abstraction layer shown in Figure 10 keeps the OTA client software portable and independent of any underlying libraries (shown in orange).

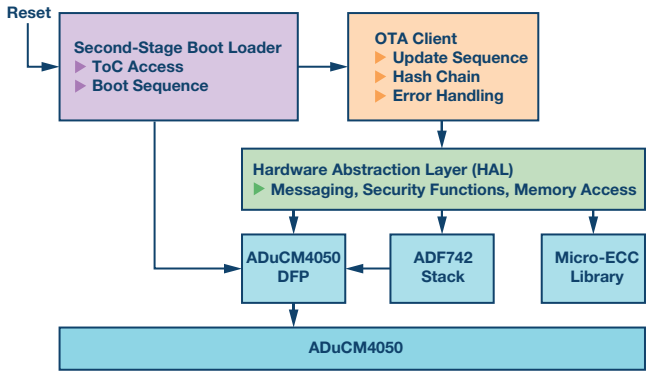


*Figure 10. Client software architecture.*

The software application implements the boot sequence in Figure 3, a simple communication protocol for downloading the new application from the server, and the hash chain. Each packet in the communication protocol has a 12-byte metadata header, 64-byte payload, and a 32-byte digest. In addition, it has the following features:

- Caching: Support for both no caching or caching one page of flash memory, depending on user configuration.
- Table of contents: The ToC is designed to only hold two applications, and the new application is always downloaded into the oldest spot, to keep a fallback application. This is called an A/B update scheme.
- Messaging: Support for either the ADF7242 or UART for messaging, depending on user configuration. Using the UART for messaging eliminates the left EZ-KIT in Figure 9, leaving the kit on the right for the client. This over-the-wire update scheme is useful for initial system bring-up and debugging.

## Results

Along with meeting the functional requirements and passing a variety of tests, the performance of the software is also critical to determining project success. Two metrics that are commonly used to measure the performance of embedded software are footprint and cycles. Footprint refers to how much space the software application takes up in volatile (SRAM) and nonvolatile (flash) memory. Cycles refers to the number of microprocessor clock cycles the software uses to perform a specific task. While being similar to software run-time, it accounts for the fact that the software may enter low power modes while performing the OTA update where the microprocessor is inactive, and no cycles are consumed. While the software reference design was not optimized for either of these metrics, they are useful for benchmarking the program and comparing design trade-offs.

Figure 11 and Figure 12 show the footprint of the OTA updates software reference design implemented on the ADuCM4050 with no caching. The figures are partitioned according to the components illustrated in Figure 10. As Figure 11 shows, the entire application uses around 15 kB of flash memory. This is quite small considering the ADuCM4050 contains 512 kB of flash memory. The true application software (the software developed for the OTA update process) only takes about 1.5 kB, with the rest being used for libraries such as the DFP, Micro-ECC, and ADF7242 stack. These results help to illustrate the design trade-off of what role the SSBL should have in the system. The majority of the 15 kB footprint is for the update process. The SSBL itself only takes around 500 bytes of footprint, with an additional 1 kB to 2 kB of DFP code for device access like the flash driver.

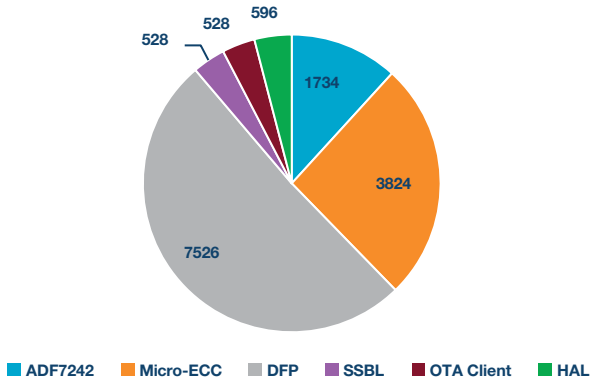


*Figure 11. Flash footprint (bytes).*

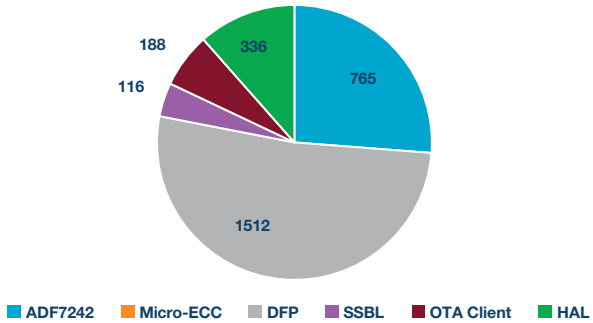


*Figure 12. SRAM footprint (bytes).*

To evaluate the overhead of the software, we perform cycle counting every time a data packet is received and then look at the average number of cycles consumed per packet. Each data packet requires AES-128 decryption, SHA-256 hashing, a write to flash memory, and some packet metadata validation. With a packet payload size of 64 bytes and no caching, the overhead is 7409 cycles to process a single data packet. Using a 26 MHz core clock, this is about 285 microseconds of processing time. The value was calculated using the cycle counting driver located in the ADuCM4050 DFP (unadjusted cycles) and is the average taken during a 100 kB binary download (about 1500 packets). The minimal overhead per packet can be attributed to the drivers in the DFP leveraging the direct memory

access (DMA) hardware peripheral on the ADuCM4050 when performing bus transactions and the drivers placing the processor into a low power sleep state during each transaction. If we disable the use of low power sleeping in the DFP and change the bus transactions to not use DMA, the overhead per data packet increases to 17,297 cycles. This illustrates the impact that efficient use of device drivers has on an embedded software application. While the overhead is also kept low by having a small number of data bytes per packet, doubling the data bytes per packet to 128 only yields a small increase in cycles—resulting in 8362 cycles for the same experiment.

Cycles and footprint also illustrate the trade-off discussed earlier of caching packet data instead of writing to flash memory each time. With one page of flash memory caching enabled, the overhead per data packet reduces from 7409 to 5904 cycles. This 20% reduction comes from the ability to skip the flash write for most packets and only perform a flash write when the cache is full. The reduction comes at a price of SRAM footprint. Without caching, the HAL only requires 336 bytes of SRAM, as shown in Figure 12. However, when caching is used we must reserve space equal to a full page of flash memory, which increases the SRAM utilization to 2388 bytes. The flash memory utilization of the HAL also increases by a small amount due to the extra code needed to determine when the cache must be flushed.

These results demonstrate the tangible impact the design decisions will have on the performance of the software. There is no one-size-fits-all solution—each system will have different requirements and constraints, and the OTA update software will need to be tuned to address them. Hopefully this article has shed some light on common problems and trade-offs that were faced when designing, implementing, and validating an OTA update software solution.

## References

Nilsson, Dennis Kengo and Ulf E. Larson. "Secure Firmware Updates over the Air in Intelligent Vehicles." ICC Workshops—2008 IEEE International Conference on Communications Workshops, May 2008.

Benjamin Bucklin Brown [benjamin-b.brown@analog.com] joined ADI in 2016 after graduating from McGill University with a Bachelor of Engineering in electrical engineering. Currently he works as an embedded software engineer in the Consumer Sensing and Processing Technology (CSPT) Group, developing firmware for application-specific integrated circuits. Previously he worked in the IoT Platform Technology Group, developing device drivers and software reference applications for the ADuCM3029 and ADuCM4050 microcontrollers.

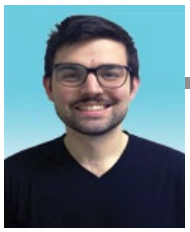

Benjamin Bucklin Brown

# Functional Safety in a Data Acquisition System

By **Chris Norris**

Share on

## Introduction

Functional safety is part of an overall safety strategy within many industries that attempts to reduce, to a tolerable level, the probability of harm coming to humans or operating equipment. The requirement for systems to be functionally safe has grown significantly in recent years. From nuclear power plants to medical devices, an errorless system has become an ideal for some and a necessity for others. For example, in the sensing world, acquiring incorrect or corrupted data can be devastating and potentially lethal depending on the system and the level of risk involved.

Traditionally, the onus was on the system developer to incorporate diagnostic and failure prevention mechanisms onto their products to ensure integrity of the data coming from the sensing IC. This came at the cost of PCB area, bill of materials, processing overhead, and, ultimately, expense. Since then, through extensive engagement with system design engineers, a solution has been developed to address this problem for them. To that effect, functional safety features have begun to be designed in at the IC level.

This article explores the functional safety potential of ADCs in terms of ensuring the overall integrity of a data acquisition system.

## A Legacy Functionally Safe Solution vs. a Better Way

In Figure 1, we see an example of a functionally safe system as it was in years gone by, and we compare it with a more modern solution. At the epicenter is the data acquisition ADC that converts the analog inputs and transmits the data to a microcontroller. To achieve this solution, however, requires many external components, repeated SPI transactions, and even a redundant ADC, which greatly increases the bill of materials, PCB area, processing overhead, and cost. It also places extra burden, such as development time and reliability, on system designers to develop this solution.

There is a single IC solution available, with minimal external components required, to operate the functional safety features.

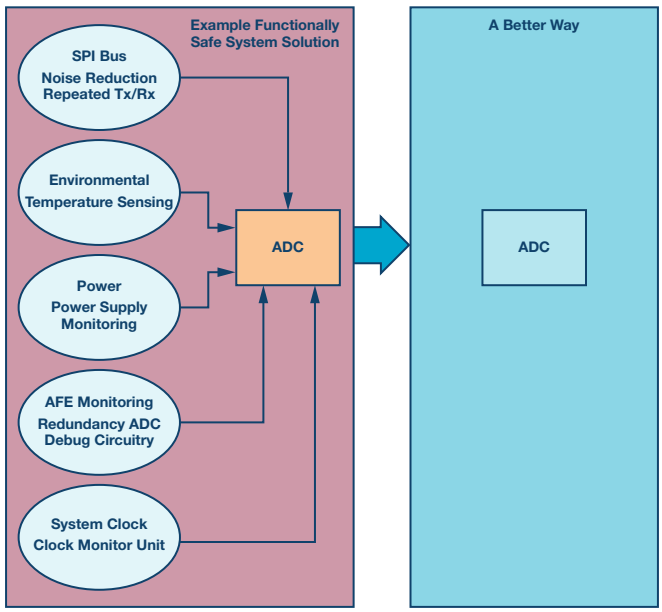


*Figure 1. Integration from multicomponent functional safety system to single-chip ADI solution.*

## An Example System with Functional Safety Requirements

In data acquisition systems that contain an ADC, many faults can occur that may increase the risk to human or machine health depending on the application. System designers must distinguish between acceptable and nonacceptable risk.

As an example, in a system that measures and regulates the pressure in a gas chamber, using a sensor that has a tolerance of 5% may be seen as

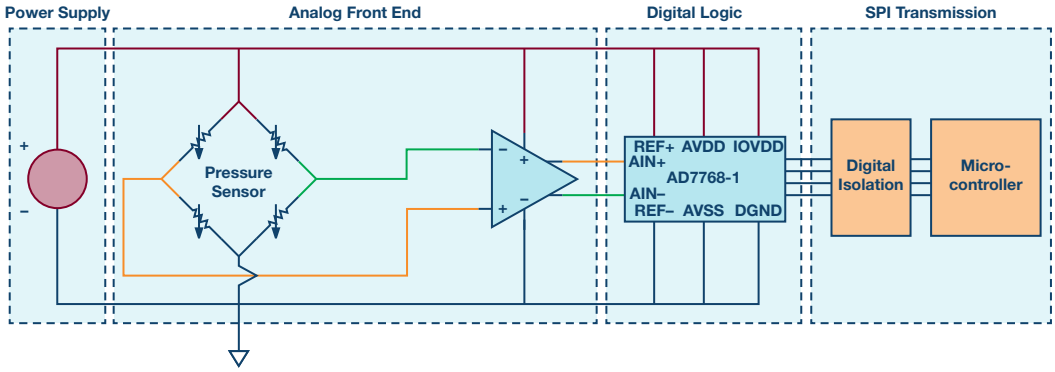


*Figure 2. Identifying the potential sources of failure in a pressure sensor system.*

an acceptable risk if the pressure inside the tank should not deviate greatly from the outside pressure. However, if the microcontroller receives incorrect ADC data, this could lead to a potentially fatal occurrence whereby the pressure in the chamber causes an implosion or explosion, both of which can injure or kill people nearby. This level of risk is unacceptable. Therefore, some functional safety measures should be put in place to ensure the integrity of the information being received by the controller.

Some sources of fault that could cause these type of errors are

- Power supply: low power supply voltage, low voltage output of the low dropout (LDO) regulators.
- Analog front end (AFE): damaged sensors or amplifier driving an incorrect voltage to the ADC.
- Digital logic: bit errors in the digital domain that can affect the converted result. For example, the factory gain or offset trim coefficient.
- SPI transmission: bit errors in the transmission of the converted data and receiving of commands due to a noisy environment of the transmission line.
- Environmental: out of the specified ambient temperature for the IC.

The AD7768-1, one of the Σ-Δ ADCs within ADI's functional safety portfolio, has a vast suite of diagnostic features intended to allow users the ability to detect and diagnose errors and more. Figure 2 highlights the sources of some of the possible faults in a typical pressure sensing system.

## Using the ADC to Diagnose the System Errors

Within the functional safety portfolio of ADI ADC products comes the ability to use an ADC to help diagnose and/or reduce the system errors. This ability to measure system errors is important in maintaining accurate measurements, and in a system that has functional safety requirements, this accuracy is even more crucial.

Positive and negative full-scale voltages, taken from the reference inputs, are used to measure the gain error of the system. A zero-scale internal short is used to measure the offset error. Users can then trim the offset and gain error performance of their system using the gain and offset trimming registers of the ADC.

A temperature sensor identifies changes in the temperature locally of the IC, including out of bounds temperatures. In a system sensitive to offset and gain error drift over temperature, this can be an attractive function. If a sizable temperature change has occurred, users may decide to trim out the gain and offset errors at this new temperature. Figure 3 illustrates how an analog diagnostic mux is connected to the ADC internally in the AD7768-1.

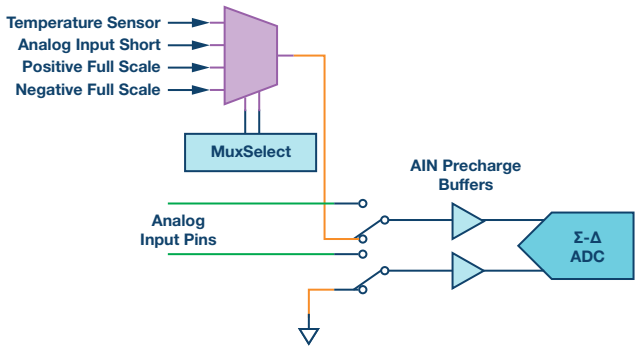


*Figure 3. Switching to convert the analog diagnostic mux.*

## Diagnostic Error Flags: The Register Map Diagnostic Status Indicators

Several diagnostic features may be enabled and their status can be flagged to the user, typically via the register map. If a fault occurs, an error flag is set in a register. Users can investigate further once they are alerted to the fault.

Let's speculate on some real life faults that can occur and that can be diagnosed using the portfolio of ADI functional safety ADCs. Let us first presume that our pressure sensor system is based in an industrial plant, with fluctuating temperatures, several shutdowns in power due to essential maintenance, and electromagnetic interference (EMI) from the surrounding industrial setting that can be conducted onto the system PCB.

### ADC Supply Error

Let's assume that, due to the high temperatures present in the environment and inrushes of current caused by the power cycling of the system, the LDO capacitors tasked with holding charge close to the ADC's LDO supply outputs have become worn and damaged. Maintaining these outputs at a known voltage requires an external capacitor and is essential for correct operation. If the capacitors are damaged due to this fault, users can notice that the converted ADC data or performance of other functions are unexpected. By enabling LDO monitors, once the voltage level drops below a certain trip point, an error flag will be set to alert users of issues at the LDO outputs.

### Analog Front-End Error

Let's presume this is a system where the inputs to the ADC should not exceed the full-scale range of the ADC. If users accidentally program an incorrect value to the gain register that increases the voltage seen by the

ADC to be greater than the full-scale range, then this will greatly affect the gain error performance of the system and should be seen as a serious risk. However, the **Filter Saturated** error checker monitors ADC output and will alert users to an out-of-range analog input.

### Digital Logic Random Bit Errors

Random bit errors occur occasionally within digital logic and memory blocks. In our example pressure system, let us say that a bit error has occurred loading the default factory offset setting during power-up. This is an intolerable fault as it disturbs the default offset error of the system, affecting the converted result. Within the ADI portfolio of functional safety ADCs there are functions available that run cyclic redundancy checks (CRCs) on the various memory blocks at regular intervals and flag faults to the user whenever a bit error occurs. A reset of the system will solve all of these faults.

### SPI Transmission Errors

Every system that transmits data along a medium will incur some bit errors along the way.

The rate at which this occurs can be estimated for every system, called a bit error rate (BER).

In our example pressure system, a BER of less than $10^{-7}$ can be assumed if transmitting to a microcontroller on the same PCB over a distance of 10 cm through digital isolation.

Let us presume that some electromagnetic interference is conducted onto the SPI lines and this results in a bit error in the transmission of converted ADC data from the AD7768-1 to the microcontroller. A bit error in the ADC data could be potentially devastating if it masks any building pressure in the gas chamber. By appending a CRC to the end of the transmitted data, users can identify if a bit error has occurred during transmission and can recheck the ADC conversion result.

### External Master Clock Error

If users are concerned about rejecting the frequencies of the main power supply (50 Hz/60 Hz) in a pressure sensor application, then an accurate low jitter external master clock source is important to align the notch of the digital filter to the correct frequency. If a source becomes disconnected, worn, or damaged, this is a huge concern as some frequency components of the main supply may become visible in the converted ADC data.

The external clock qualifier can flag an error to the users if the external clock source has not been connected successfully or if it has been removed. Users can then perform emergency conversions with the internal RC oscillator while essential maintenance is carried out on the external master clock source.

### POR Flag

Once a system is powered-up or successfully reset, the POR flag within the ADC will be set.

However, if an unexpected reset occurs, users may see unexpected results in the ADC data. They can identify this unexpected reset by checking the POR flag.

Figure 4 shows how many of these internal diagnostic features within the AD7768-1 hook up to the functions they will be monitoring.

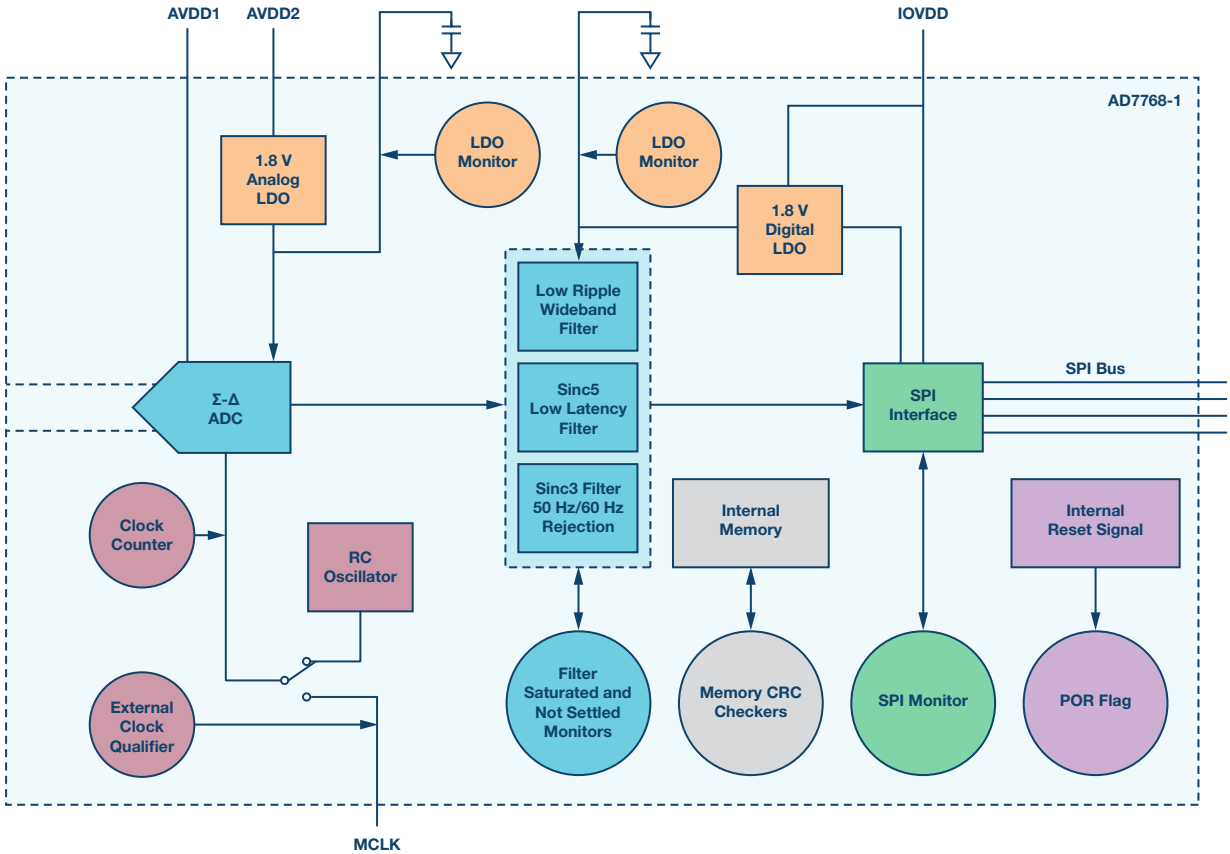


*Figure 4. Internal diagnostic monitors of the AD7768-1.*

## The Ultimate Functional Safety Solution Using the AD7768-1

Using the functional safety features provided by the AD7768-1, the following data acquisition system is possible. Users can power up the part and enable the following functional safety features:

- SPI integrity monitors
- LDO regulator output level monitoring
- A filter saturation monitor
- An external clock qualifier
- Internal logic and memory CRC monitors

System calibrations can be validated with the internal analog diagnostic mux. LDO regulator outputs can also be verified this way.

Next, users can enable the functions to append the 8-bit status byte to the end of the 24-bit data stream and the 8-bit SPI CRC word. The 8-bit CRC is calculated based on the 8-bit command word, the 24-bit data stream, and the 8-bit status word. If users are concerned about the amount of processing overhead, they can enable **Continuous Read-Back** mode, which removes the need to provide the 8-bit command. Instead, users may clock out the data register contents upon providing serial clocks to the part, as shown in Figure 6.

The result of this procedure is a data acquisition system whose gain and offset error have been verified and that is providing diagnostic information to the user every time they read back data the ADC data.

The LDO regulator outputs, analog front-end inputs, internal digital logic, and memory are continuously monitored. The users can be certain of the integrity of the SPI communications and that the temperature of the IC is known.

## Conclusions

As the requirements for functional safety within many industries grow, so too must the technology that supports these requirements. Analog Devices is continuing to develop the technology within our portfolio of products to support system designers in their quest for functionally safe operation.

The AD7768-1 takes much of the burden off the shoulders of the customer and has provided a solution that is more compact, less complex, and reducing processing overhead and the bill of materials required to produce the required solution. This single component approach can also ease the burden on system designers who wish to earn a Safety Integrity Level (SIL) certification for their designs.

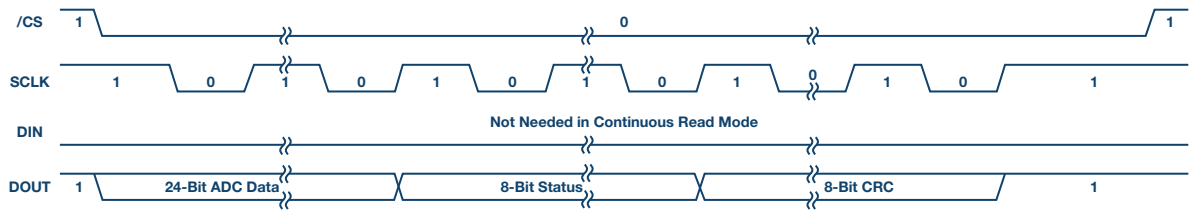


*Figure 5. Reading back the data register with appended status byte and CRC byte of the AD7768-1 in **Continuous Read-Back** mode.*


Chris Norris [christopher.norris@analog.com] is an ADC design evaluation engineer with Analog Devices in Limerick, Ireland. He received his Bachelor of Science in electronic engineering from Waterford Institute of Technology in 2012 and joined ADI as a graduate the same year.


Chris Norris

# Care and Feeding of FPGA Power Supplies: A How and Why Guide to Success

By **Nathan Enger**

Share on

## Introduction

Modern FPGAs are among the most complex integrated circuits ever created. They employ the most advanced transistor technology and cutting-edge architectural structures to achieve both incredible flexibility and the highest performance. Over time, as technology has advanced, this complexity has dictated certain compromises in the design and implementation of systems using the FPGAs. This is nowhere more apparent than in the power supplies, which must be ever more accurate, more agile, more controllable, smaller, more efficient, and more fault aware with each new FPGA generation.

In this article, we look specifically at some of the constraining specifications for the Altera® Arria 10 FPGA, and what they mean for a power supply design. Then we discuss the best power delivery solutions, and lay out the plan to successfully meet all of the specifications and make our FPGA perform at its optimal efficiency, speed, and power level using Analog Devices' complete set of power system management (PSM) ICs, including the LTC3887, LTC2977, and LTM4677.

## FPGA Power Requirements (Interpreting the Data Sheet)

Engineers should spend most of their time programming—they don't want to spend time and energy thinking about designing suitable power supplies. Indeed, the best approach to power delivery is to use a robust, flexible, proven design that meets the requirements and expands with the project. Here we take a closer look at some of the important power specifications and what they mean.

## Voltage Accuracy

Core power supply voltage is one of the most important keys to balancing FPGA power and performance. The specification documents give a range of acceptable voltages, but the total range is not the complete picture. As with all things, there are trade-offs and optimizations to be made.

Table 1 is an example of core voltage specifications from the popular Altera Arria 10 FPGA.[1] Though these numbers are specific to the Arria 10, they are representative of other FPGA core voltage requirements. The range amounts to a ±3.3% tolerance around the nominal voltage. The FPGA will operate just fine within this voltage window, but the complete picture is more complicated.

**Table 1. Altera Arria 10 Core Voltage Specification**

| Symbol | Description | Condition | Minimum | Typical | Maximum | Unit |
|---|---|---|---|---|---|---|
| $V_{CC}$ | Core voltage power supply | Standard and low power | 0.87 | 0.9 | 0.93 | V |
| | | | 0.92 | 0.95 | 0.98 | V |
| | | SmartVID | 0.82 | | 0.93 | V |

Notice the line labeled "SmartVID," with a range of 0.82 V to 0.93 V. This represents a broad range of voltages that are possible when the FPGA is requesting its own core voltage through the SmartVID[2] interface (more on this later). This SmartVID specification is an indication of an underlying truth about the FPGA: it can operate at different voltages, depending upon its particular manufacturing tolerance, and upon the particular logic design that it is implementing. The static voltage required by one FPGA might be different from another FPGA. The power supply must be able to respond and adapt.

The goal is to produce just the right performance level to operate the programmed functions without burning unnecessary power. We know from semiconductor physics, as well as published data from Altera, Xilinx® (Figure 1), and others, that both dynamic and static power will increase dramatically with increased core $V_{DD}$, so the goal is to give the FPGA just enough voltage to meet its timing requirements, but no more. Excess power dissipation does nothing to increase performance. In fact, it makes things worse because transistor leakage current increases with higher temperatures, dissipating even more unwanted power. For these reasons, the imperative is to optimize the voltage for the design and operating point.

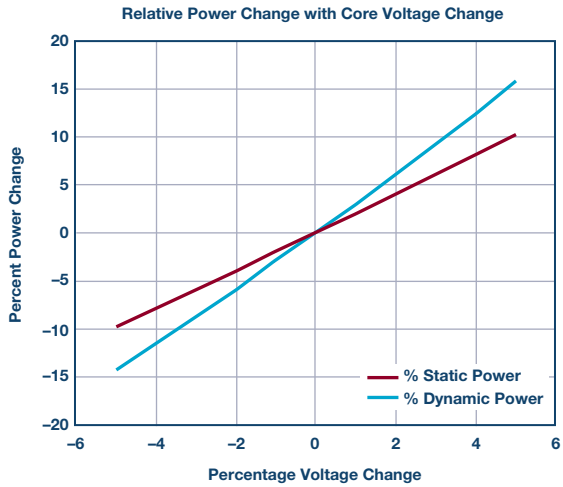


*Figure 1. Xilinx Virtex V power vs. core $V_{CC}$.*

This optimization process requires a very accurate power supply for success. Regulator inaccuracy must be baked into the error budget and subtracted from the available voltage range that can be used for optimization. If the core voltage drops below the requirement, the FPGA may fail due to timing errors. If core voltage drifts above the maximum specification, it may damage the FPGA, or it may create hold-time failures in the logic. All of these scenarios must be guarded against by taking into account the power supply tolerance range, and only commanding voltages that are guaranteed to remain within the specification limits.

The problem is that most power regulators are not accurate enough. The regulated voltage may be anywhere within the tolerance range around the commanded voltage, and it may drift with load conditions, temperature, and age. A power supply that guarantees a ±2% tolerance may regulate anywhere within a 4% voltage window. In order to compensate for the possibility that the voltage may be 2% too low, the commanded voltage must be raised by 2% above that required to meet timing. If the regulator then drifts to 2% above the commanded voltage, it will operate 4% above the minimum voltage required at that operating point. This still meets the specified voltage required by the FPGA, but it wastes a lot of power (Figure 2).

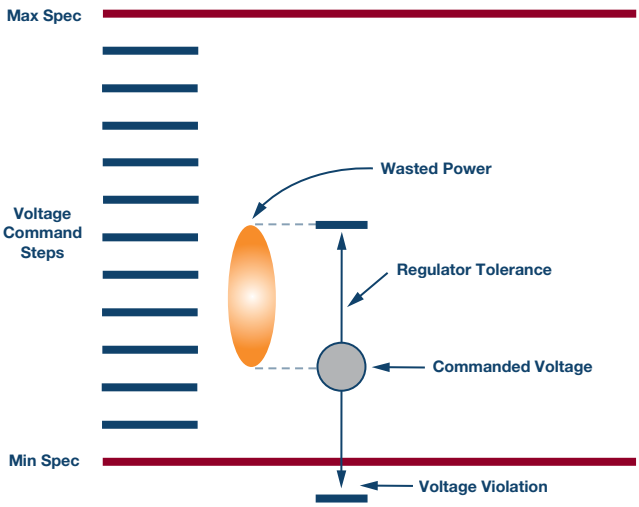


*Figure 2. Power supply regulator tolerance trade-offs.*

The solution is to choose a power supply regulator that can operate with a much tighter voltage tolerance. A regulator with a ±0.5% tolerance can be commanded to operate much closer to the minimum required specification at the desired operating frequency, and it is guaranteed to be less than 1% from that required voltage. The FPGA will be functional and it will be dissipating the minimum power possible at that operating condition.

The LTC388x family of power supply controllers guarantees regulated output voltage tolerance better than ±0.5% over a wide, configurable voltage range. The LTC297x family of power system managers guarantee a trimmed voltage regulator tolerance of better than ±0.25%. With these accuracies, there is a clear path to optimizing the power vs. performance trade-off for any FPGA.

## Thermal Management

A more subtle implication of power supply accuracy manifests itself in the thermal budget. Because static power dissipation is far from negligible, the FPGA heats up even when it is doing nothing. The increased temperature causes more static power dissipation, which further increases operating temperature (Figure 3). Adding unnecessary voltage to the power supply only makes this problem worse. An inaccurate power supply requires a guard band in operating voltage to ensure that there is enough voltage to do the job. The power supply voltage uncertainty that results from variability in tolerance, system assembly variability, and operating temperature can create voltages significantly higher than the minimum required. This additional voltage, when applied to the FPGA, can cause thermal complications, or even thermal run-away at high processing loads.

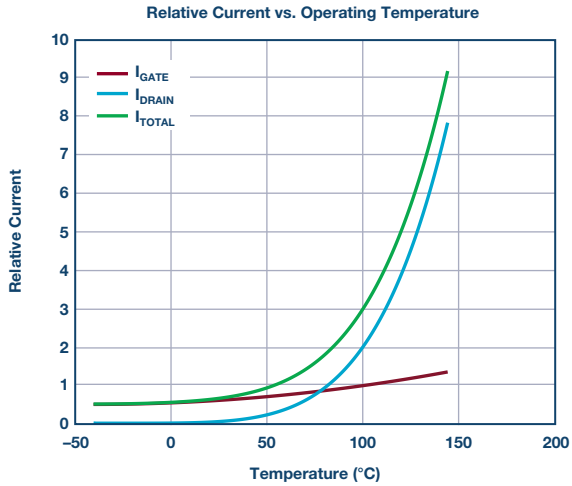


*Figure 3. Power supply current vs. operating temperature.*

The remedy is a very accurate power supply that creates just the right voltage, and no more than necessary, which is exactly what the ADI power system management (PSM) devices do well.

## SmartVID

SmartVID is the Altera name for the technique of operating each individual FPGA at its optimal voltage, as requested by the FPGA itself. There is a register inside of the FPGA that contains a device-specific voltage (programmed at the factory) at which the FPGA is guaranteed to operate efficiently. A piece of compiled IP inside of the FPGA can read this register and make a request over an external bus to the power supply to provide this exact voltage (Figure 4). Once the voltage is reached, it remains static during operation.

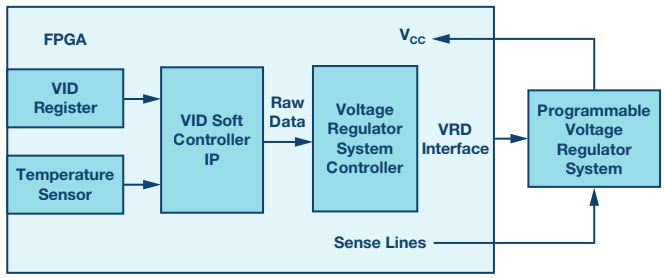


*Figure 4. Altera SmartVID structure.*

The demands of the SmartVID application on the power supply include the specific bus protocol, voltage accuracy, and speed. The bus protocol is one of several methods that the FPGA uses to communicate its required voltage to the power regulator. Of the available methods, PMBus is the most flexible because it addresses the widest variety of power management ICs. The SmartVID IP uses two PMBus commands: VOUT_MODE, and VOUT_COMMAND, with which it commands the PMBus-compliant power regulator to the correct voltage.

The voltage accuracy and speed requirements for the regulator include an autonomous boot voltage (before the PMBus is active), the ability to accept a new voltage command every 10 ms, the ability to take 10 mV steps every 10 ms during the voltage adjustment phase, and the ability to settle to within 30 mV (~3%) of the target within the 10 ms step time, ultimately ramping to the commanded voltage and remaining static during FPGA operation.

Though Altera uses SmartVID, there are other, similar techniques in use around the industry that accomplish much the same thing. One of the simplest is to test each board at the factory and program into the power supply's nonvolatile memory an exact voltage that optimizes performance for that particular board. This technique does not require any further intervention for the power supply to operate at the correct voltage. This is an advantage of a power supply manager or controller with EEPROM.

All of the requirements for Altera SmartVID can be met by the LTC388x family of power supply controllers. In addition, the LTM4675/LTM4676/LTM4677 µModule regulators easily meet the requirements and offer a complete solution in a single compact form.

## Timing Closure

The computing speed of any logic block depends on its supply voltage. Within limits, a higher voltage gives faster performance. We have already seen why we cannot simply operate at the highest voltage in order to guarantee the best speed. On the other hand, we must operate at a high enough voltage for the application, as shown in Figure 5.

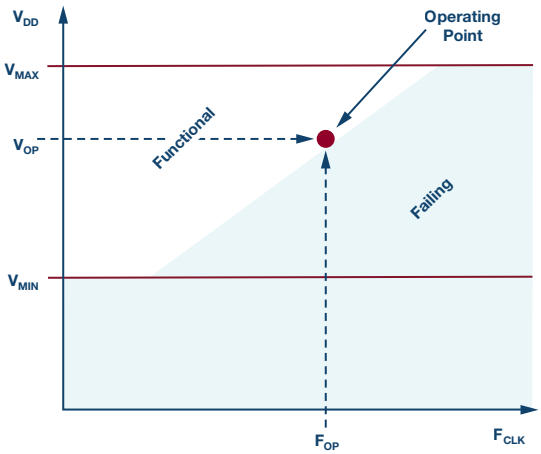


*Figure 5. FPGA operating frequency vs. $V_{DD}$ trade-off.*

An important implication of Figure 5 is what can be done when a particular design does not meet its logic timing requirements, and falls into the failing area. Often the edge between functional and failing is not well defined before a design is committed to hardware, and the particular voltage at which it will pass timing cannot be predetermined. The only options are to either commit in advance to a voltage that is well above the minimum, thus wasting power to guarantee functionality, or to design a flexible power supply that can adapt to the needs of the hardware at test time, or even, as in the case of SmartVID, at power-up time. The ability to adapt to unknown demands makes the accuracy of ADI PSM devices more valuable, as FPGA designers can trade off power for performance in the real design, and at any stage of development.

## Power Supply Sequencing 101

Moore's law drives the trend of shrinking transistors in modern FPGAs and forces the trade-offs involved in using these tiny transistors, which are very fast and small, but much more fragile. A chip containing hundreds of millions of transistors must be segmented into cores, blocks, and partitions that can be designed and managed independently. The practical result of these considerations is an FPGA with many power domains. Some recent FPGAs have upward of a dozen power supplies that need to be kept happy. This includes, in addition to voltage, current, ripple, and noise, the sequence order during start-up, shutdown, and fault conditions.

Recent FPGA specifications call out specific requirements for sequence order when starting-up and shutting down the power supplies. Both Xilinx and Altera recommend specific ordering and timing to ensure that the FPGA gets reset properly, maintains minimal current draw, and maintains its I/Os in the proper tristate configuration during the power transitions.

Given the number of power supplies for each FPGA, the complexity of the sequencing task is considerable.

The Altera Arria 10 prescription divides the power supplies into three sequence groups (1, 2, and 3), and requires that they sequence up in order 1, 2, and 3, and down in the reverse order: 3, 2, and 1.[3]

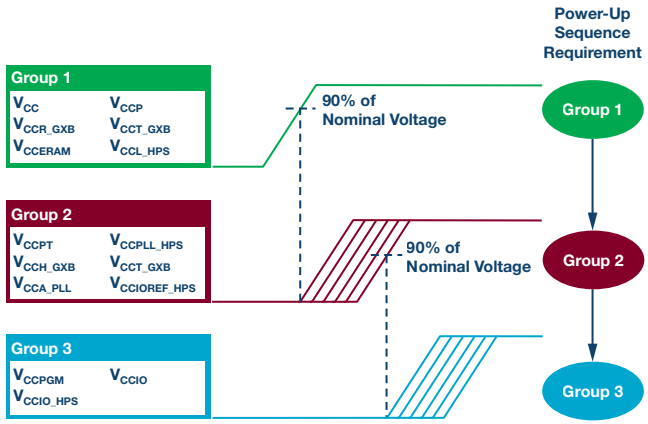


*Figure 6. Altera Arria 10 up-sequence group order.*

Similarly, the Xilinx recommendation for the Virtex UltraScale FPGA up-sequence is: $V_{CCINT}$/$V_{CCINT\_IO}$, $V_{CCBRAM}$, $V_{CCAUX}$/$V_{CCAUX\_IO}$, and $V_{CCO}$. Down-sequencing is the reverse of the up-sequence order.[4]

These are just two of the many FPGAs available. Nearly every modern FPGA system has multiple power supply rails, and one of the most obvious questions to ask is, in what order should they turn on and off? Even if there is no explicit sequencing requirement there are good reasons to enforce a deterministic sequence of events. Here are some of the available design options.

- No sequencing: Let the power supplies rise and fall on their own. What could go wrong?
- Hardware cascade sequencing: Each rising power supply is hardwired to enable the next one. This only works when the supplies are rising.
- CPLD-based sequencing: Use programmable logic to create a custom solution. This is flexible, but the entire challenge rests on the designer.
- Event-based sequencing: Event-based sequencing is similar to cascade sequencing, but more flexible because it can operate both up and down. A dedicated sequencer IC can be programmable and take care of many fault scenarios and corner cases.
- Time-based sequencing: Time-based sequencing triggers each event at a specified time. Coupled with comprehensive fault management, time-based sequencers can be flexible, deterministic, and safe.

The following sections explore these options in more detail.

## No Sequencing

It is possible to turn-on a power supply system with no management at all. When main power becomes available, or the ON switch is activated, the regulators start regulating. When power is removed, or the ON switch turns off, the regulators stop regulating. Of course, the problems with this approach are many. Some more obvious than others.

Lack of timing determinism can have various effects in a system. The first is simply that it stresses the sensitive FPGA. This might cause immediate catastrophic failure, or it might cause premature aging that slowly degrades performance. Neither is good. It may also cause unpredictable power-on-reset behavior or indeterminate logic states at power-up, which make system stability questionable and difficult to debug. Questions of fault detection and response, energy management, and debug support are left entirely unanswered in this scheme. In general, to avoid power supply sequencing is to invite disaster.

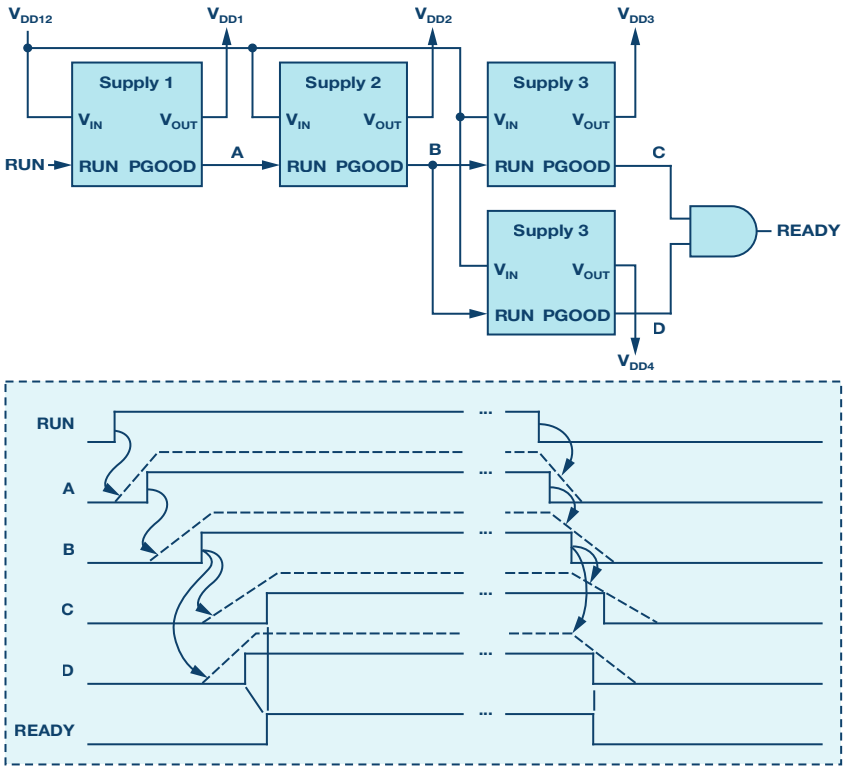


*Figure 7. PGOOD-to-RUN cascade sequencing.*

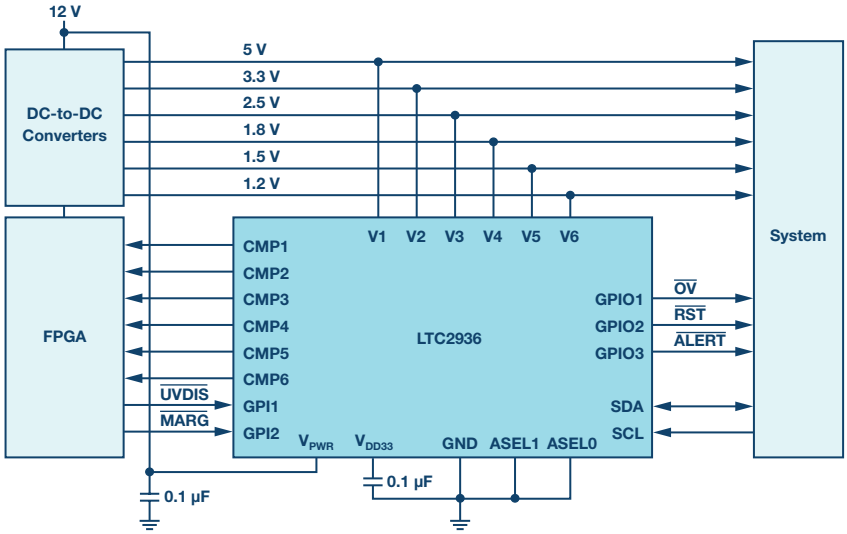


*Figure 8. LTC2936 programmable voltage supervisor.*

## Cascade Sequencing

A slightly more organized approach to sequencing is the classic PGOOD-to-RUN hardwired cascade shown in Figure 7. This is like dominoes falling: each one taps the next in the sequence, which guarantees progress in order. This technique has the benefit of simplicity. Unfortunately, it also has its shortcomings. While it often works adequately for sequencing up a power system, it cannot operate in reverse (or in any other order) for down-sequencing. There can be only one sequence order. Additionally, this scheme cannot gracefully handle faults or manage energy in uncertain operating conditions. It is not smart enough to make any decisions. If one stage of the sequence fails, what happens next? If one operating power supply browns-out, what happens? The answers are undefined, and debugging these problems is not easy.

## FPGA or CPLD Sequencing

Using an auxiliary CPLD or FPGA on the board to sequence the supplies is an option that many designers choose. In a system designed by and for digital designers, it has a certain appeal. There is a natural fit to designing a digital control block that can be programmed into an FPGA to control the power supply of another FPGA. Here the decision can be deceptive because a power supply system is not as simple as it may appear from a digital control perspective.

If designers wish to tackle the power supply sequencing, control, and management problem from top to bottom, they must first thoroughly understand its complexities. We have already discussed many of these, and there are more as well, such as detecting and responding to overvoltage and undervoltage conditions that can happen on the microsecond time scale, detecting hazardous currents and temperatures, logging telemetry and status, and providing bring-up and debug services to make life easier for the hardware guys. All of these considerations require dedicated analog hardware in addition to the digital algorithms.

For the intrepid designers who wish to take this path, Analog Devices provides several analog front-end ICs to help with the task. At the interface between the digital bits and the analog power supply, the LTC2936 provides six rugged, highly accurate programmable threshold analog comparators to

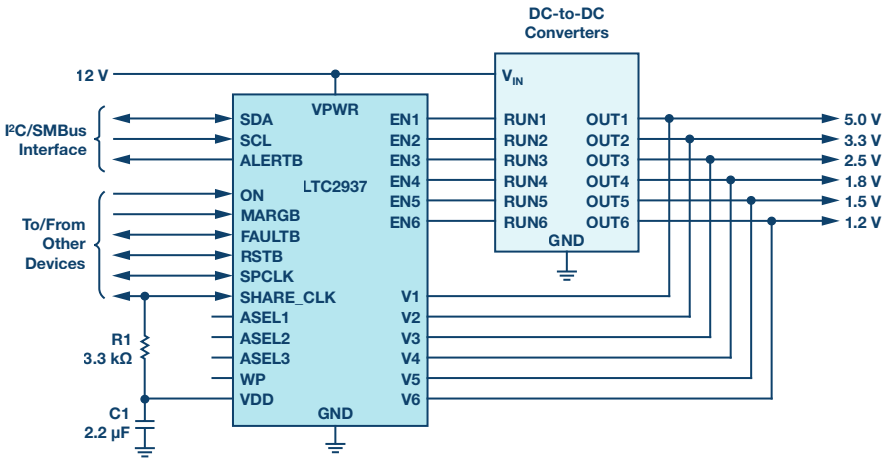


*Figure 9. LTC2937 power supply supervisor and sequencer.*

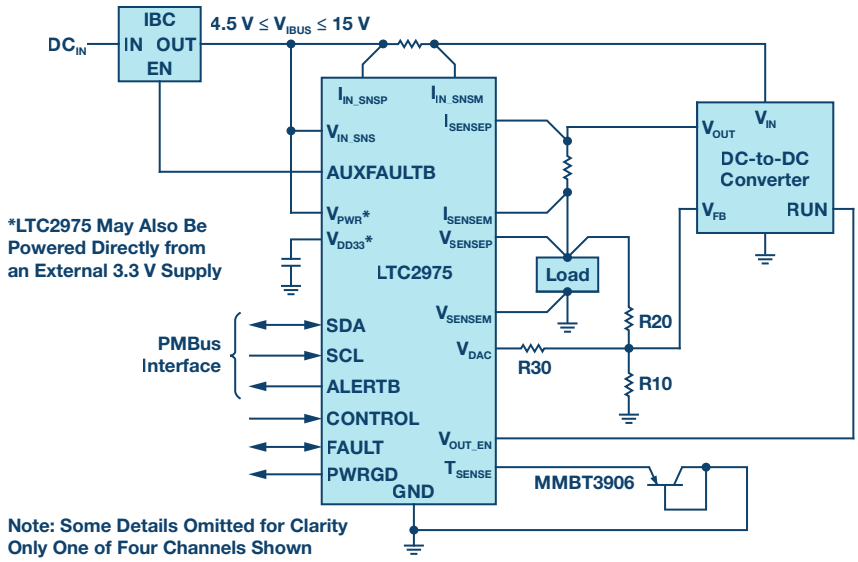


*Figure 10. LTC2975 4-channel power system manager.*

detect fast events and send digital status to the logic. It also has three programmable GPIO pins for additional functionality. This programmable IC has an EEPROM for nearly instant-on functionality at start-up, and has the ability to store fault telemetry for debugging through its I²C/SMBus interface. A convenient way to use LTC2936 is shown in Figure 8.

In addition to the fast comparator functions, there must be an analog-to-digital converter (ADC) to gather telemetry. A proven choice is the LTC2418, which can monitor up to 16 channels of analog signals with its fast-settling 24-bit Σ-Δ ADC and 4-wire SPI interface. The board controller can readily stream measurements and monitor many points of interest in the system.

In general, there are many, many options for using an FPGA or CPLD to control power sequencing. This approach works, but somebody must own the digital and analog designs, including all of the inevitable design bugs, opportunities for unimaginable corner cases and faults, and the unhappy question of support. There are certainly easier ways to build a power system.

## Simple Sequencer/Supervisors

Solving the puzzle of robust sequencing and fault handling is the domain of the simple sequencer/supervisors. These do the important job of sequencing the power rails and ensuring that they remain within their specified limits during operation (supervision). The LTC2928 is an easy to use pin-strap configurable sequencer with configurable sequence timing (down is the reverse of up), and configurable supervisor voltage thresholds. It has the potential to meet the requirements, but has no frills and offers no digital programmability or telemetry.

In the category of programmable sequencer and supervisor with EEPROM is the LTC2937. It features full digital programmability, features time-based and event-based sequencing, and can sequence and supervise any number of power supplies, handling faults, and logging fault status to the EEPROM black box. It is a worthy solution for cases where voltage management and telemetry are not required.

## Power System Management

To fully capture all of the benefits of complete PSM, use one of the Analog Devices PSM ICs. These introduce the ability to autonomously sequence up and down any number of power rails; accurately control rail voltages to better than 0.5% (or 0.25% in some cases); measure and report voltage, current, temperature, and status telemetry; cooperatively handle complex fault scenarios; and record detailed fault information to EEPROM.

Sequencing is done by a system of timing handshaking, with all ICs agreeing on time zero and the time base, and all sequence events happening at preprogrammed times (time-based sequencing). This allows any number of rails to sequence up and sequence down autonomously.

The family of PSM ICs includes controllers that have their own switch

drivers and analog loop control to handle the switching power supply in all aspects. Alternatively, power supply managers contain a servo loop that wraps around an external power supply, adding all of the features of power supply management, including sequencing, supervision, and monitoring, to any power rail, from switching power regulators to LDO regulators. An example of a power supply manager is the LTC2975, pictured in Figure 10.

## µModule Devices

The most tightly integrated solutions, providing the most functionality per square centimeter, in a BGA or LGA footprint, are PSM µModule devices. These are complete power supply systems in a single package, including controller ICs, inductors, switches, and capacitors. Some µModule regulators, such as the LTM4650, do not contain digital functions, so they can benefit from additional sequencing and management with the LTC2975. Some µModule regulators, such as the LTM4676A, contain their own PSM functions and can easily integrate with other PSM ICs in the system.

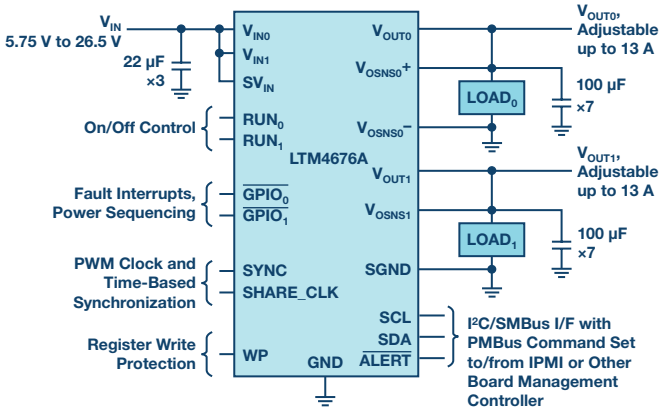


*Figure 11. LTM4676A PSM µModule dual 13 A regulator.*

## Shared Sequencing

The PSM micromodules, manager ICs, and controller ICs all cooperate together in sequencing up and sequencing down by sharing timing information through a simple one-wire bus called SHARE_CLK. Through this single wire, all of the PSM ICs share information about when sequencing should begin (time zero), when each tick of the clock occurs, and other status information that affects sequencing. It is enough to simply connect all of the SHARE_CLK pins together in the system to enable this coordination. Each IC has its own programming for sequence timing that can use the shared time base to accurately and reliably time events such as enabling and disabling, ramping, and timing-out in case of a fault.

The SHARE_CLK pin, at its most basic level, is an open-drain, 100 kHz clock pin. The open-drain nature means that an IC can either pull-down actively or let go and allow the bus to float. When all devices on the bus let go, the pull-up resistors pull the voltage to 3.3 V. This allows one device to stop the clock by pulling down until it is ready, and means that all devices must agree before the clock can start: an effective mechanism for communicating time zero, as well as indicating sequencing status by stopping the clock.

## Shared Fault Handling

Similar to the SHARE_CLK pin is the FAULT bus. Each of the PSM ICs in the system is connected to the shared FAULT wire, and can either pull it low using its open-drain output, or respond when another device pulls low. This provides a simple, quick way for the entire family of PSM devices to communicate and respond to faults. The behavior is fully configurable, and allows a coordinated response when something goes wrong, either during sequencing or during steady state. The system can be configured to remove power and attempt to re-sequence according to specified timing, while recording black-box information about the state of the system and causes of the fault when it occurred. This EEPROM black-box information is available for later processing over the I²C bus.

## Down-Sequencing and Managing Stored Energy

There is an additional consideration when sequencing down the supplies: energy management. Increasingly, it is important to provide deterministic timing for power supplies as they sequence down, and this requires carefully considering where the stored energy in the system is dissipated. A high power supply is likely to have dozens of large electrolytic capacitors as bulk charge storage elements, and these will be charged to the supply voltage, holding enough energy to blow-up an improperly protected device under unfortunate conditions. To avoid this, FPGA manufacturers specify a down-sequence that protects the device. For the Altera Arria 10 this sequence is shown in Figure 12.[5]

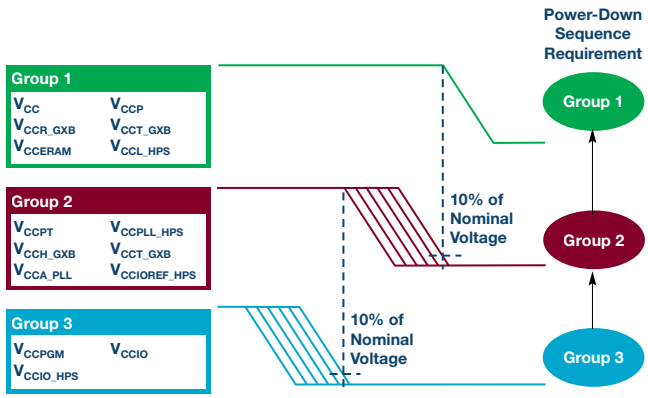


*Figure 12. Altera Arria 10 down-sequence group order.*

Implicit in this down-sequence is the requirement that all of that stored energy in the capacitors goes somewhere and is safely dissipated. There are several ways to do this. The simplest is a fixed resistor across the capacitor. This resistor always dissipates power while the supply is turned on, but its resistance can be made large enough that the comparative loss is minimal, and the RC discharge time constant is acceptably short. The time that it takes to sufficiently discharge the supply is a multiple (usually 5×) of the RC time constant, and should be optimized to make the static power dissipated in the resistor acceptable (<¼ W, for instance). For a 1 mF capacitance and a 1.0 V, supply a resistor value of R = 4 Ω will have a time constant of τ = 4 ms, and will discharge the supply to below 50 mV in just about 13 ms. This approach is sufficient as long as the resistor is rated for at least ¼ W, and the system works with a constant ¼ W loss and a 13 ms discharge time.

A more complex, but very safe, option is to switch the resistor across the capacitor only when it is time to discharge the supply. This approach pulls charge out of the bulk capacitors when it is time to do so, and safely dissipates it in the resistance of the switch FET and in a supplemental series resistor, but it avoids the continual power drain of a fixed resistor. The circuit is shown in Figure 13.

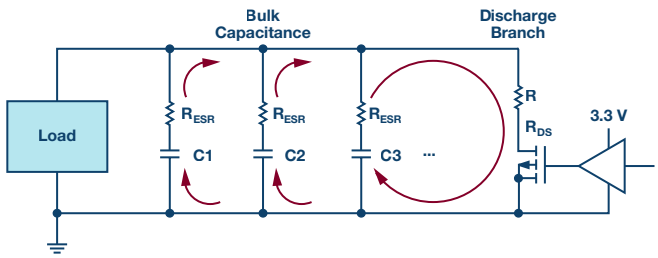


*Figure 13. Discharging power supply capacitance with a FET.*

There are several considerations with this approach: control, discharge time, and power dissipation. There must be an available signal to command the discharge switch to close at the appropriate time. The switch FET is NMOS, so the control signal must rise above the $V_{TH}$ of the FET suffi-

ciently to drive it into saturation. For common FETs, this gate drive voltage may be as much as 3 V to 5 V.

The typical electrolytic capacitor will have hundreds of milliohms of equivalent series resistance (ESR), which will dissipate some of the energy as the capacitor discharges, but there are many of these capacitors in parallel, so the total parallel capacitance may add up to tens of millifarads, and the equivalent resistance will be tens of milliohms or less. It is a safe assumption that the capacitor ESRs will dissipate a small fraction of the stored energy.

In order to discharge the capacitance in a reasonable amount of time, the discharge RC time constant must be less than 1/5th the desired discharge time (to allow the voltage to fall below a few millivolts). This is a simple calculation (Equation 1) using the sum of all of the capacitors and the sum of the FET and series R, as well as the parallel combination of the $R_{ESR}$ resistances, where N is the number of parallel capacitors.

$$\tau = \left[\frac{R_{ESR}}{N} + R + R_{FET}\right] \times [C_1 + C_2 + ... + C_N] \quad (1)$$

For a larger system with a 50 mF capacitor bank and the sum of $R_{DS}$ + R = 500 mΩ, the voltage will fall below 50 mV in about 125 ms. The peak current (and power) during this time is 1 V/500 mΩ = 2 A, or 2 W. Because a large majority of the stored energy is burned in the first two time constants, we can decide if a series resistor is necessary by looking at our FET's safe operating area chart, such as the example in Figure 14.[6] In this case, our FET will safely withstand a 2 W pulse in excess of 10 seconds, so there is no danger of damaging it. This FET, however, has an $R_{DS}$ of less than 20 mΩ, so the series R must be 480 mΩ. We must size the series resistor to handle the heat, since it will dissipate most of the power. In general, the pulse duration will be much shorter than the thermal time constant of the resistor. The resistor data sheet gives more information.

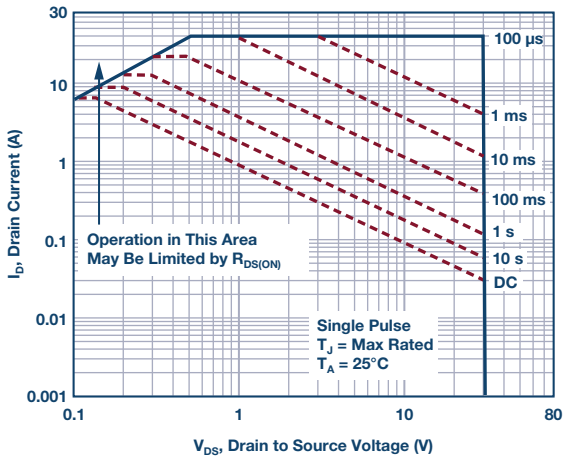


*Figure 14. NMOS FET safe operating area.*

The most robust discharge circuit can safely dissipate energy over a wide variety of conditions. The circuit in Figure 15 shows a tried and true method. It uses the ON semiconductor FDMC8878 discharge FET and a physically large SMD 1210 sized 0.5 Ω resistor.

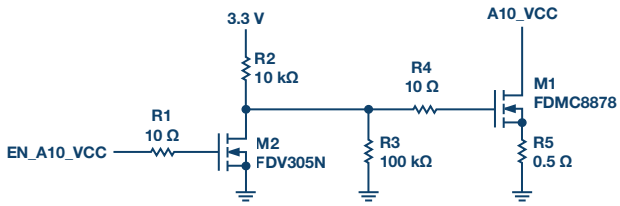


*Figure 15. Discharge FET circuit.*

## Meeting the Challenges with Power System Management

As we have seen, the best solution for managing all of the requirements in an FPGA power system is with Analog Devices PSM. The benefits of this portfolio include:

- Best-in-class voltage accuracy (better than ±0.5%)
- Full autonomy with EEPROM memory
- Integrated, fully programmable power supply sequencing, and independent up and down timing across the whole system
- Integrated, robust, system-wide fault management
- Comprehensive telemetry: voltage, current, temperature, and status
- Coordinated family of ICs addresses all areas of the power supply system

The Altera Arria 10 SoC development kit showcases the Analog Devices power system management solutions for the Altera Arria 10 SoC IC (Figure 16).

In this design (Figure 17), the core power supply operates at 0.95 V and 30 A. With these relatively relaxed power requirements, a single LTM4677 module easily supplies the necessary current (up to 36 A) as shown in Figure 18. For more demanding applications requiring more current, up to four LTM4677 modules can be run in parallel to provide up to 144 A, as in Figure 19.

*Figure 16. Altera Arria 10 SoC development kit.*

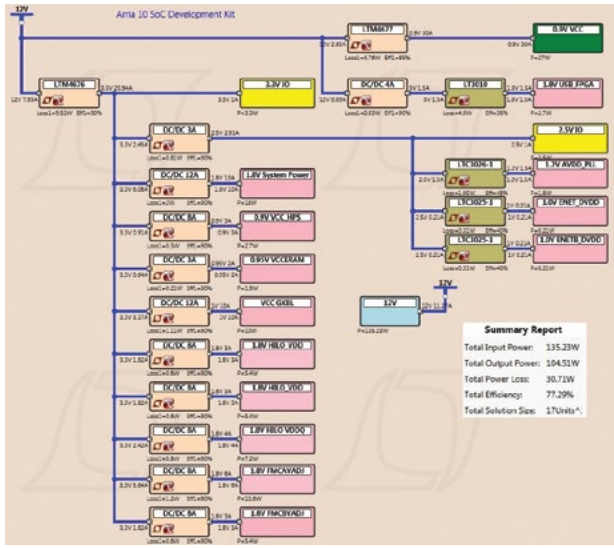


*Figure 17. Arria 10 SoC development kit power distribution.*

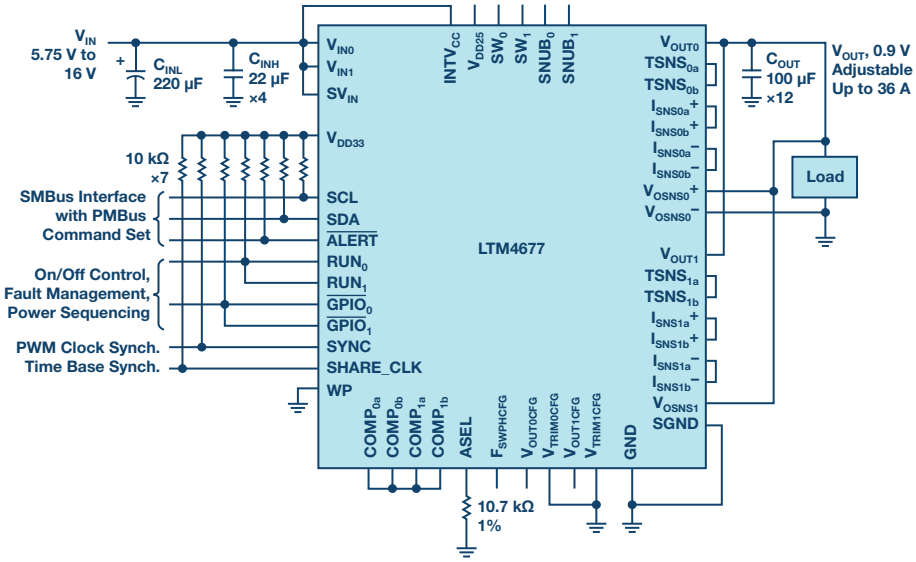


*Figure 18. A single LTM4677 providing up to 36 A.*

This solution provides the best board space utilization because the integrated µModule devices require very few external components, and the PMBus interface makes them configurable without hardware modification. Micromodules offer the lowest complexity solution because many of the complex analog considerations, such as power switches, inductors, current and voltage sensing elements, loop stability, and thermals, are included.

Because the LTM4677 module includes PSM, it guarantees that the core power supply will always operate within ±0.5% of the dc voltage target. It also allows for voltage adjustments through the PMBus interface, both from the SmartVID IP inside of the FPGA and from the LTpowerPlay® graphical user interface (GUI) that gives full user control over the power supplies.

To manage power supply regulators that do not include their own PSM capabilities, we simply include the LTC2977, which is an 8-channel PMBus-compatible power system manager. Each channel wraps around a power supply to servo the voltage to within 0.25% of the programmed target (Figure 19). It cooperates seamlessly with the LTM4677 µModule devices for both sequencing and fault response, making the entire power system coherent and easily programmable.

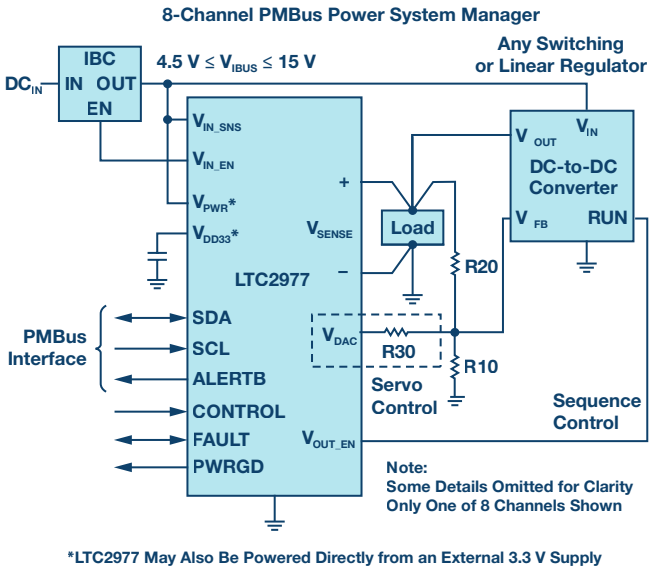


*Figure 20. LTC2977 managing any power supply voltage.*

System power sequencing is provided by the cooperative partnership of the LTM4677 core supply, the LTM4676A 3.3 V power supply, and the LTC2977 that manages all of the other power regulators on the board. These ICs have common PMBus timing commands (stored in EEPROM) that easily configure start-up and shutdown sequencing in any order, and with any timing. These guarantee the proper autonomous sequence of events specified for Group 1, Group 2, and Group 3 power supplies (Figure 6).

In addition to the voltage accuracy and sequencing control, the LTM4677, LTM4676A, and LTC2977 on this board provide complete fault handling as well. In the event of an overvoltage, undervoltage, brownout, overcurrent, or complete failure of one or more rails, the system can be configured to respond quickly and automatically shut down to protect the sensitive FPGA and restart if it is possible.

Most of the power supply rails in the system require modest currents (less than 13 A) and modest voltage tolerance. These can be supplied by non-PSM devices such as the LTM4620, and sequenced and managed by the LTC2977. This provides a very effective balance between board area, complexity, and cost.

There are also power supply rails, such as PLL and transceiver power, that require lower noise than a switching regulator can provide, and these need a linear regulator. The LTC3025-1and LTC3026-1 serve these functions well, removing the switching and load-induced noise from their outputs. These, too, can be managed by the LTC2977 to sequence, trim, and handle fault conditions.

## LTpowerPlay

The entire family of PSM devices is supported by the comprehensive LTpowerPlay GUI (Figure 21). Because much of the functionality of PSM is accessed through the rich set of configuration registers in the EEPROM of the ICs, one tool can consolidate the entire collection of PSM ICs on the bus into one easy to use view. The LTpowerPlay tool provides a deep set of features to accelerate all phases of design and development. It can operate offline to present a view of the ICs before they are programmed, or in real-time, communicating over the I²C bus with a complete system containing from one to hundreds of power supply rails controlled by many PSM devices. LTpowerPlay simplifies and streamlines complex configurations by providing detailed information about registers and functions. It graphically represents all of the configuration, status, and telemetry information available in the system, making it clear and understandable while the system is operating. It simplifies programming and maintenance of the complete

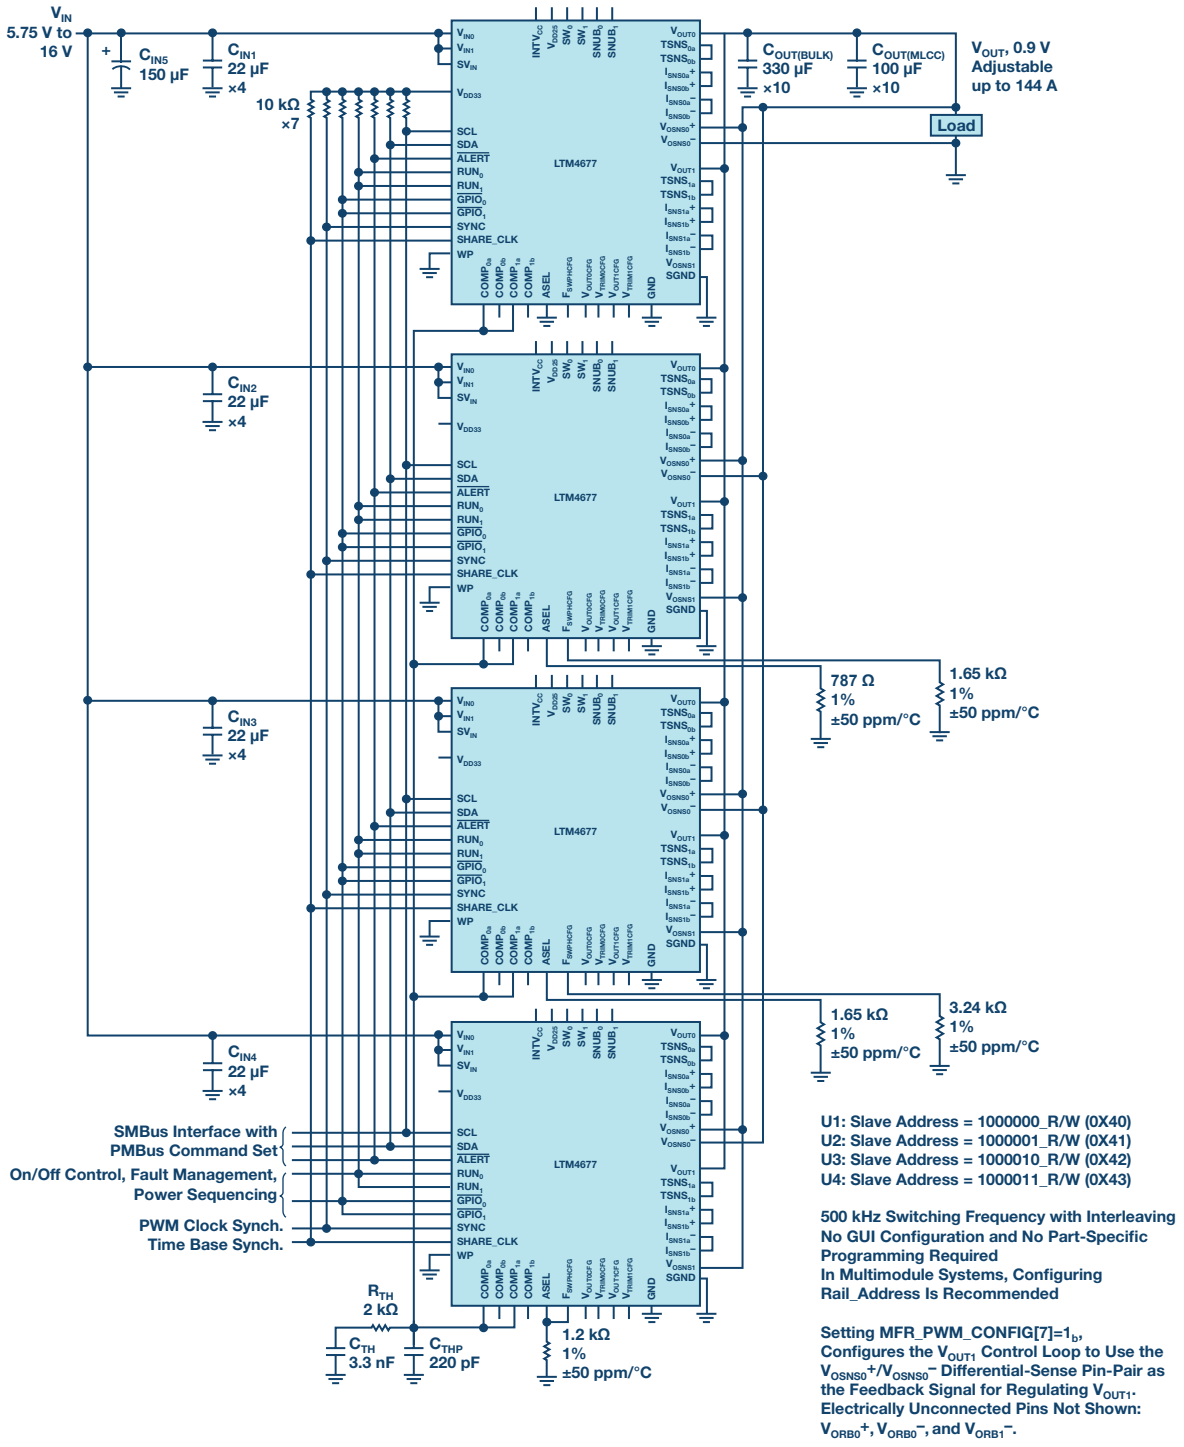


*Figure 19. Four LTM4677s providing up to 144 A at 0.9 V.*

register sets, providing a simple way to create and save configurations on a Microsoft® Windows® PC. When power supply faults occur, LTpowerPlay makes it easy to see where in the system the fault happened and what the status, telemetry, and black-box information indicate about what happened. It also provides detailed debugging help for common fault scenarios. In the event that someone needs a helping hand, LTpowerPlay also has the ability to call for help, enlisting a live support person who can view the GUI running in real time and see what you see.

Download the free LTpowerPlay tool here.

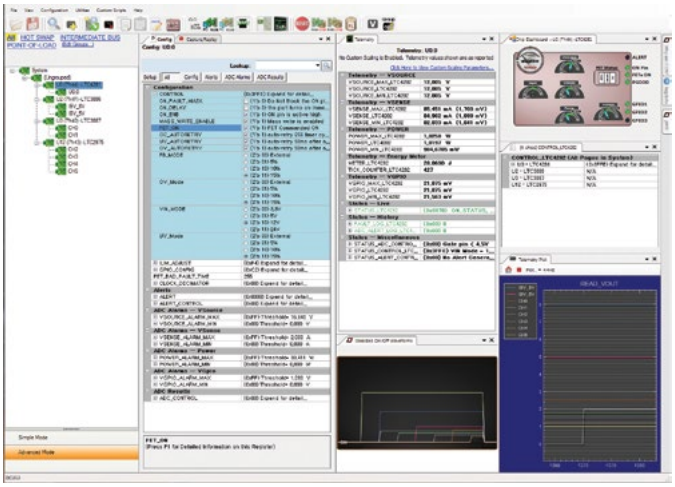

*Figure 21. LTpowerPlay graphical user interface.*

Analog Devices provides a comprehensive set of demonstration platforms for Altera, Xilinx, and NXP FPGAs. These fully functioning boards are working examples of how PSM provides the cleanest, most flexible, and robust power solutions for FPGA systems. In addition, your local Analog Devices applications engineers can provide detailed help in selecting and using the complete portfolio of PSM ICs. Read more, download reference material, and order the FPGA boards here.

## The Journey to FPGAs

Now that we understand how best to power an FPGA system, we can take a whimsical aside and look at why things are this way. In order to understand why things are the way that they are today, we need a brief lesson in history.

## Moore's Law

In 1965 Gordon Moore published his famous article in *Electronics Magazine*,[8] stating his observation that the number of transistors on a single chip had been doubling every year, and his prediction that it should continue to do so, at least through the year 1975. Later enhancements and additional observations of the larger electronics market caused him to revise his model, but the basic principle of a sustained exponential growth rate in the number of transistors on a chip has become an axiom in the electronics industry. This is a curious self-fulfilling prophecy that exists in no other industry, and at no other time in history. In fact, it has become a primary motivator for engineers across the globe, creating innovations, and forcing trade-offs that were as-yet unimagined when Gordon Moore first published his simple observation.

As a result of this technological race against ourselves, the decision-making process has always favored technologies that squeeze more devices into a smaller area at the expense of cost, power, usability, and even durability. In the technology race, size is everything. Some of the implications of this trajectory are that advanced chips use more power, become leakier, are more fragile and more sensitive, and they are much more difficult to manage and to protect.

## Transistor Engineering

As transistors have shrunk to feature sizes on the nanometer scale, important side effects have become increasingly dominant. The most obvious is voltage headroom. Whereas 5 V would have been a fine power supply for transistors a few decades ago, such a voltage would break down all of the junctions and oxides in a more recent FET transistor. As transistor features shrink, the internal electric fields become much stronger, and the tolerable operating voltage shrinks to prevent damage. Recent transistor generations can only tolerate about 1.0 V as a maximum power supply voltage. In addition, the absolute voltage tolerance also shrinks proportionally: 2% of 1.0 V is a much smaller range than 2% of 5 V, making accuracy an ever more pressing concern.

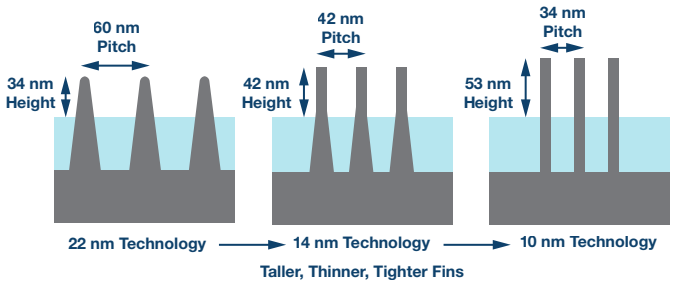


*Figure 22. Shrinking transistor sizes.*

Along with shrinking voltages comes increasing transistor current drive ($I_{DSAT}$). Increased drive strength accomplishes at least two purposes. First, it allows a transistor with a smaller gate voltage to drive a significant current—making it strong enough to switch at useful frequencies. Second, it allows a physically smaller transistor. The smaller transistor can be faster. Unfortunately, the increased transistor drive strength comes with its own penalty: leakage current.

There are two kinds of power dissipated by the transistors on a chip. Dynamic power is the familiar cost of switching between Logic 1 and Logic 0 at some frequency, and dynamic power is caused by charging and discharging the tiny parasitic capacitors associated with the transistors themselves and the wires on the chip connecting devices together. Dynamic power is proportional to the frequency of logic transitions and to the square of the power supply voltage.

Less obvious is the power spent by leaking transistors. This power leaks away any time the circuit is powered, regardless of whether it is active or idle, clocking or not. Increased transistor drive strength causes more leakage current because junctions and structures made to conduct more current are harder to turn off. A stronger transistor tends to leak more than a weaker one. With every transistor generation, the effects of leakage have grown. It is only the combination of heroic transistor engineering (chemical, metallurgical, lithographical, and physical) and accurate, flexible power supply management that keeps leakage power under control.

A decade ago, Gordon Moore observed these facts and noted two important points. First, if dynamic power continued to rise at the same pace, then the junction temperature on an operating chip would approximate that of the surface of the sun. Second, if nothing else was done, leakage power would overtake dynamic power as the primary energy dissipation mode, further exacerbating the power dissipation problem (Figure 23). To address these effects, the IC industry adopted several new techniques at that time. One of these was clock management—slowing or stopping the clocks to curb dynamic power—and another was the use of multiple processing cores on a single chip to leverage the growing transistor count.

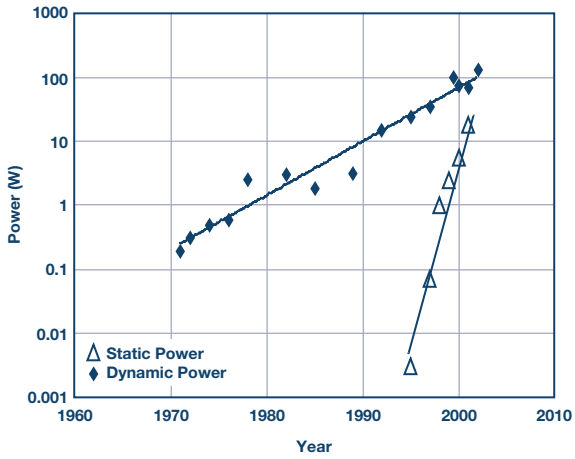


*Figure 23. Static and dynamic power growth.*[9]

Even with all of this advanced architecture, the problem of leakage power remains troublesome. Transistor engineering is a powerful way to bend the curve downward, but it is not enough. As each smaller transistor generation demands a lowered supply voltage, the problem of dynamic power remains tame, but the resulting increased transistor strength and leakage, coupled with the ever-growing number of devices on the chip, produces a demand for voltage management. Power supply voltage must be tightly controlled, as well as actively adjustable, to meet the needs of each particular device.

## Advanced Architectures

Architectural developments up until the turn of the millennium had mainly focused on optimizing a single computing core to perform as many computations possible as quickly as possible. This involved the free technique of raising the clock rate to just under the speed at which the circuit failed: its maximum operating frequency. It also involved architectural optimizations, but these were mainly intended to squeeze more performance out of every clock cycle.

Following the startling realization that power mattered, engineers began to redirect resources away from raw speed and into more subtle kinds of optimization. This new trend showed up in computing architecture first as a plateau in the ever-increasing clock speeds, and a leveling-out of the rate of performance increase per transistor in each generation (Figure 24). This was the most obvious way to tame the beast of dynamic power: stop slopping charge from $V_{DD}$ to $V_{SS}$ so fast.

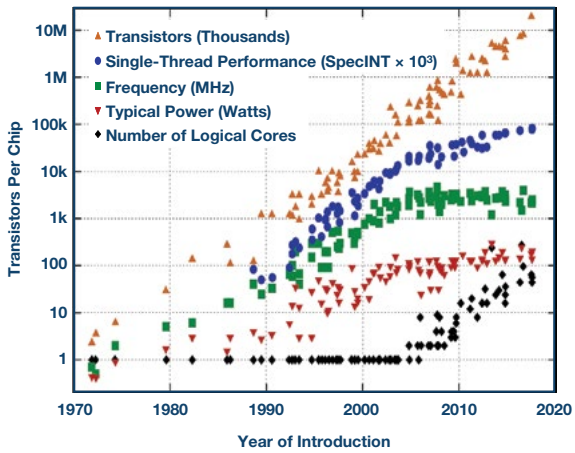


*Figure 24. Growth trends in computing hardware metrics.*

But the number of transistors on a single chip continued to climb at the inexorable rate predicted (demanded?) by Gordon Moore. Something had to be done with all of these transistors. This necessitated the second great innovation: multicore architectures. At about the same time that the clock speeds stopped growing, the number of cores on a single chip started growing. The advantages of multiple cores include simpler chip design through reuse, simpler software design with familiar building blocks, and the ability to individually throttle each core to meet the demands of the computational load. The multicore revolution began with fixed computing platforms, but one might say that this event was the moment when FPGAs came into their own: when the world realized that maximizing the number of cores is best. In a sense, nothing has more cores than an FPGA with its sea of identical, programmable logic blocks!

## Anatomy of an FPGA

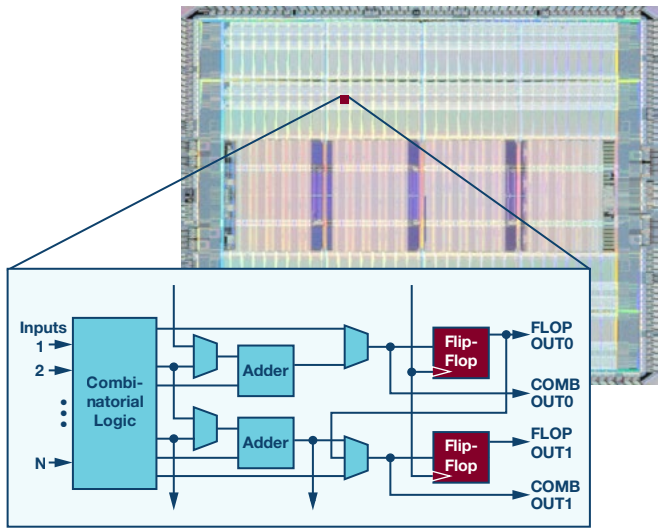


*Figure 25. FPGA configurable fabric.*

An FPGA, at its most basic level, is a collection of primitive configurable logic cells tied together through a configurable mesh of interconnections. Together with a compiler they form a highly flexible computing fabric that transforms into nearly any imaginable general-purpose digital function, including combinatorial and sequential logic blocks. At the top level, this fabric is surrounded by additional features to support and augment functionality. Some blocks, such as bias circuits, RAM, and PLLs, support functions internal to the chip. Various configurable GPIO cells, high speed communication hard macros (LVDS, DDR, HDMI, SMBus, etc.), and high speed transceivers allow the logic inside of the chip to communicate with the outside world with various voltages, speeds, and protocols. Other blocks, such as integrated CPU and DSP cores, support optimized functions that are commonly needed, and optimized for power, speed, and compactness.

The FPGA core fabric is composed of thousands or millions of primitive cells called configurable logic blocks (CLBs). Each CLB is a collection of combinatorial and sequential logic elements that together can produce an elementary computation and hold the value in one or more flip-flops. The combinatorial logic usually takes the form of a programmable look-up table (LUT) than can translate a few input bits to a few arbitrary output bits. Each LUT performs one basic logic function, as programmed, and passes the result into the configurable interconnect for subsequent processing (Figure 26). The specific CLB and LUT design is one of the secrets that make one FPGA family different from another. Inexpensive FPGAs use simpler CLBs with fewer inputs, outputs, and interconnections, and fewer flip-flops. The highest end FPGAs use much more complex CLBs, each capable of more inputs, more logic combinations, and higher speed. This optimization allows more computation per CLB, and more optimized performance in the compiled design. Naturally, the added inputs and outputs in more complex FPGAs have different dynamic power trade-offs than simpler, less interconnected devices.

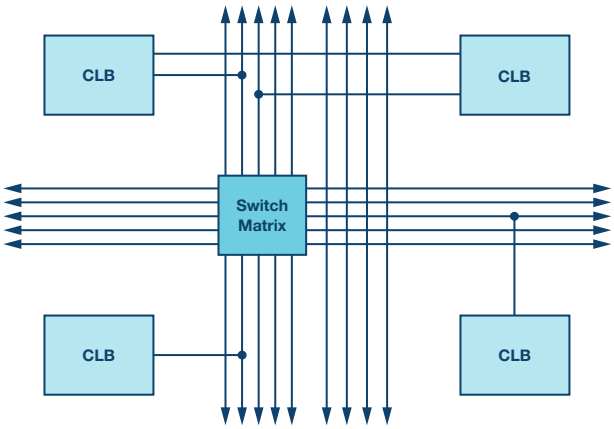


*Figure 26. Configurable switch matrix between CLBs.*

The fundamental concept of configurable logic functions continues outside of the core fabric itself into the I/O cells, which are also highly configurable to meet a broad array of voltages, drive strengths, and logic styles (push-pull, tristate, open-drain, etc.). Like the configurable LUTs and interconnect matrix, programmable I/Os receive their configuration at boot-up time from configuration memory, which has implications for power supply sequence order.

There are also functional blocks that cannot, or should not, be implemented using general-purpose CLBs and GPIOs. These are the so-called hard macros. They are functions that benefit from optimization, or that simply cannot be made fast enough or small enough, and need-dedicated circuits. These include gigabit transceivers, arithmetic logic and DSP elements, specialized controllers, memories, and dedicated processor cores. These are hard macros in contrast to soft blocks that can be compiled like software and loaded into the configurable fabric. Hard macros usually have their own power supplies, specific voltages, and timing requirements.

All of these various functional blocks have diverse power needs that the power supply system must accommodate. The core fabric usually requires both the lowest voltage and the highest power on the chip. In a modern FPGA, the fabric, when fully utilized, may demand more than 100 A from a power supply operating at 0.85 V. Similar voltages are found in the CPU core, but at different currents, and with different sequencing requirements. Other on-chip analog functions may be powered by 1.8 V or 3.3 V and must be energized before anything else. At the same time, the GPIO banks may operate at 3.3 V or 1.8 V, and must not be energized until the power-on reset for the core fabric is complete. Each of these power supply sequence requirements must be enforced by the system.

The final piece of FPGA architecture is the tool chain (Figure 27). In order to transform the blank slate of configurable logic fabric into a high performance circuit, a comprehensive set of tools exists to translate a set of Verilog or VHDL code into logical blocks, assign clock, reset, and testability resources; optimize the functions for speed, power, or size constraints; then load the result into the FPGA's configuration EEPROM. Without these tools the FPGA would never achieve its full potential. In fact, the tools and programming languages are so important that they often overshadow the fundamental circuit design that enables the FPGA to function. Engineers spend most of their time programming and don't want to spend time and energy thinking about delivering suitable power supplies. Often overlooked, however, are the power supply requirements implied by the tools. Because so much effort goes into the digital design, it is only late in the game when the compiled design comes together that power requirements are known, and problems with power supplies may be discerned. Here in the digital design and software tools, just as much as in the hardware design, flexible power supply architecture is crucial to success.

History, economics, and human factors continue to drive the trends in transistors and architectures that create FPGAs. At every level, and every design stage, the power supply plays a critical, unseen role in the success of the FPGA. The best choice of power supply is one that is accurate, robust, flexible, compact, and easy to use. In all of these qualities, Analog Devices' family of PSM products set the standard for the industry.

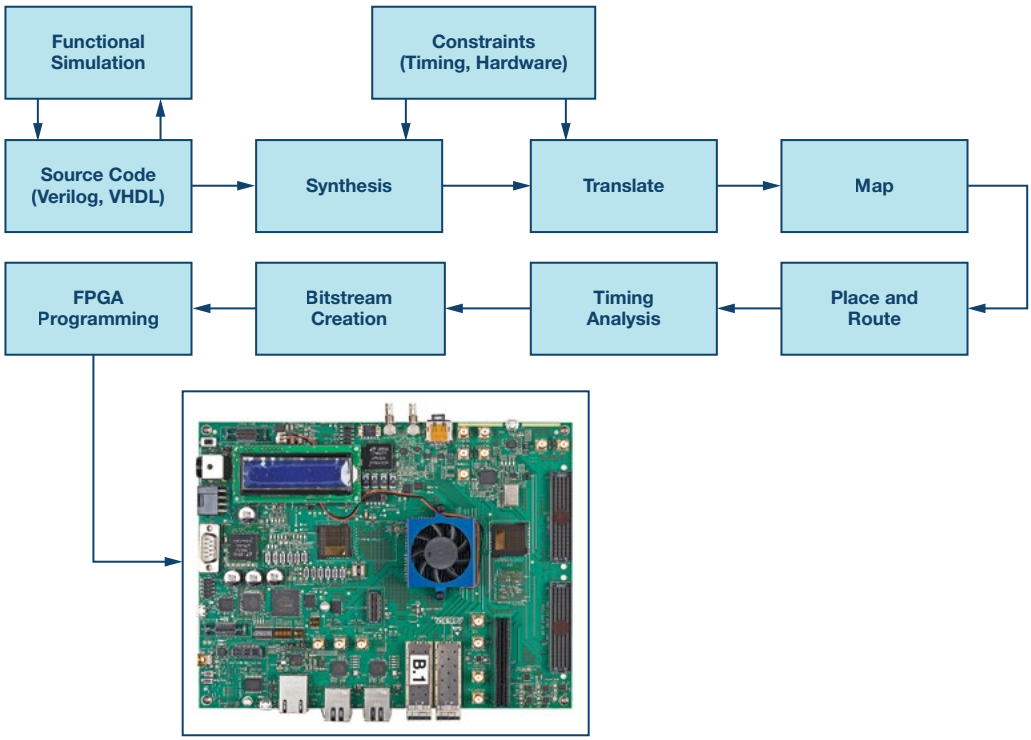


*Figure 27. FPGA design flow (hardware comes last).*

## Acknowledgements

We would like to thank M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, I. Hammond, C. Batten, and K. Rupp[11] for their contributions in collecting and plotting the data in Figure 24.

## Appendix A: Care and Feeding?

There may be some readers who question the title of this article, especially those who are not familiar with English language vernacular. At first it may seem entirely out of place to refer to the care and feeding of an FPGA. The answer to this objection, however, is quite simple: English is a funny language. While no one agrees on the precise moment in history that the term "care and feeding" came into popular usage, it is well understood that the term originates in the agrarian roots of a simpler time, and has come into popular use (abuse) to refer to just about anything that may be fragile or temperamental. In this case we have hit the nail on the head. While it is arguable whether one must "feed" an FPGA, one certainly must "care" for it!

While the term "care and feeding" is broadly applied in our modern internet era to such things as infants, children, husbands, bosses, expatriates, scientific data, and even digital pulse-shaping filters, one of the earliest, and perhaps the oddest, references can be found in this classic text, which is interesting not only for its title, but for the fact that it is instantly available on the internet.


CARE AND FEEDING
—OF—
Indian Runner Ducks
BY MRS. GEO. R. SIMPSON
FIRST EDITION
OWENSVILLE, - - INDIANA
PRICE, ONE DOLLAR
E. D. HULFISH, PRINTER
1911

*Figure 28. Care and Feeding of Indian Runner Ducks by Mrs. Geo. R. Simpson.*

## References

[1] Intel Arria 10 device data sheet. Intel, June 2018.

[2] AN 711: Power Reduction Features in Intel Arria 10 Devices. Intel, July 2018.

[3] Intel Arria 10 Core Fabric and General Purpose I/Os Handbook. Intel, August 2018.

[4] Virtex UltraScale FPGAs data sheet: DC and AC Switching Characteristics. Xilinx, January 2018.

[5] AN 692: Power Sequencing Considerations for Arria 10 and Stratix 10 Devices. Intel, April 2018.

[6] DMN1032UCB4 data sheet. Diodes Incorporated, January 2015.

[7] Joachim Von der Ohe. Application Note: Pulse Load on SMD Resistors: at the Limit. Vishay Beyschlag, August 2015.

[8] Gordon Moore. "Cramming More Components onto Integrated Circuits." *Electronics Magazine.* Vol. 38, No. 8. April 19, 1965.

[9] Gordon Moore. "No Exponential Is Forever: But 'Forever' Can Be Delayed!" International Solid-State Circuits Conference, 2003.

[10] Samuel Fullers and Lynette Millett. *The Future of Computing Performance: Game Over or Next Level?* National Academy of Science, Washington D.C.

[11] Karl Rupp. "40 Years of Microprocessor Trend Data." February 2018.

Nathan Enger [nathan.enger@analog.com] is a mixed-signal applications engineer at Analog Devices in Colorado Springs, Colorado. Nathan has a B.S.E.E. and an M.S.E.E. from the University of Colorado and Colorado Technical University, respectively. He has 23 years of mixed-signal IC design experience, including high speed PLL and transceiver development at Ford, Intel, and Marvell Corporations.

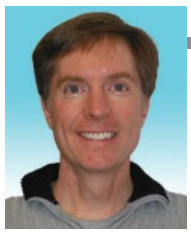

Nathan Enger

# Rarely Asked Questions—Issue 159 When Grounds Are Separated

By **Frederik Dostal**

Share on

## Question:

Where do I connect the grounds of switching regulators?

## Answer:

How should you proceed with a switching regulator with an analog ground (AGND) and a power ground (PGND)? This is a question asked by many developers designing a switching power supply. Some developers are accustomed to dealing with a digital GND and an analog GND; however, their experience frequently fails them when it comes to the power GND. Designers then often copy the board layout for a selected switching regulator and stop thinking about the problem.

PGND is the ground connection over which higher pulsed currents flow. Depending on the switching regulator topology, this means the currents through a power transistor or the pulsed currents of a power driver stage. This is especially relevant in the case of switching controllers, for example, with external power switches.

AGND, sometimes called SGND (signal ground), is the ground connection that the other, usually very calm, signals use as a reference. This includes the internal voltage reference needed for the regulation of the output voltage. Soft start and enable voltages are also referenced to the AGND connection.

There are two different technical philosophies, and thus different opinions among experts regarding the handling of these two ground connections.

According to one philosophy, the AGND and PGND connections on a switching regulator IC should be joined to each other right next to the respective pins. This keeps the voltage offset between the two pins relatively low. Thus, the switching regulator IC can be protected from disturbances and even destruction. All of the circuit's ground connections and a possible ground plane would be linked to this common point in a star topology. Figure 1 shows an example implementing this philosophy. The board layout for an LTM4600 is shown here. It is a 10 A step-down micromodule. The separate ground connections on the board are joined right next to each other (see the blue oval in Figure 1). Due to the parasitic inductance of the respective bonding wires between silicon and the housing, as well as the inductances of the respective pins, there is already a certain amount of decoupling of PGND and AGND, resulting in a low amount of mutual interference between the circuits on the silicon.

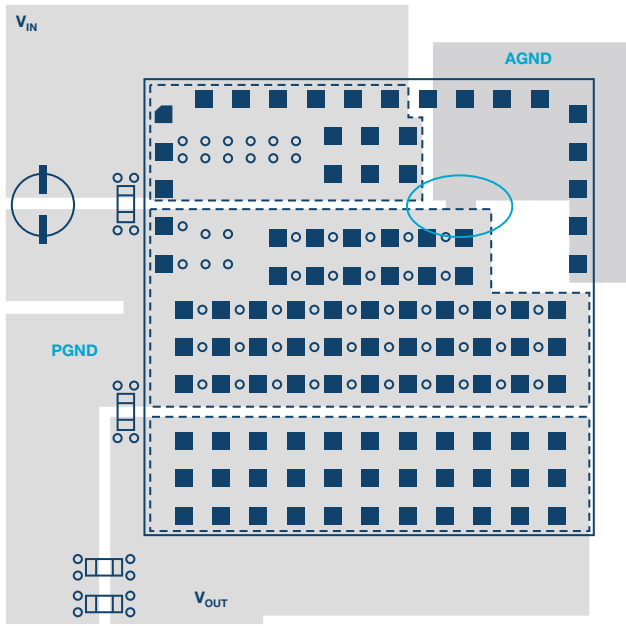


*Figure 1. Local connection of PGND and AGND right at the solder contacts.*

The other philosophy involves additional separation of AGND and PGND on the board into two separate ground planes connected to each other at one point. Through this connection, interfering signals (voltage offset) remain largely in the PGND region, while the voltage in the AGND region remains

very calm and decoupled very well from PGND. However, the disadvantage is that, depending on the transients in the pulsed currents and the current intensities, there may be a significant voltage offset between PGND and AGND at the respective pins. This can lead to improper functioning of, or even damage to, a switching regulator IC. Figure 2 shows an implementation of this philosophy. This comes from an ADP2386, a 6 A step-down switching regulator.

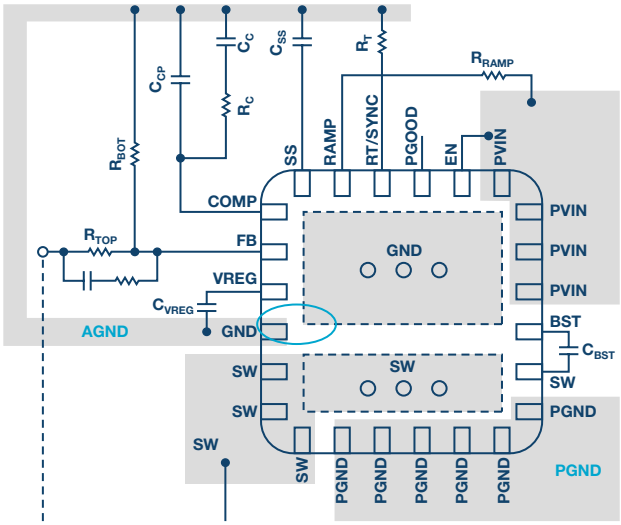


*Figure 2. Separated AGND and PGND connected under the GND tab by vias.*

The grounding question comes down to a trade-off between strong separation with the advantage of separating noise and disturbances, and running the risk of generating voltage offsets between the two grounds and thus causing harm to silicon and compromising functionality. The right decision to make in regard to this trade-off is heavily based on the IC design, including switching transition speeds, power levels, parasitic inductances on bonding wires and IC package, and the latch-up risk of each IC design involving the individual semiconductor process.

## Conclusion

The answer to the question of how to deal with the grounds AGND and PGND is not that simple. That's why this discussion continues. At the beginning, I mentioned that many users of switching regulators adopt the board layout and the ground connection type from the example circuit supplied by the IC manufacturer. This procedure is useful because you can usually assume that the manufacturer also tested the respective IC in this configuration. It can also be seen in the examples given in Figure 1 and Figure 2 that the respective IC pinout is suitable for local ground connection close to PGND and AGND, or for separate grounding.

Of course, an IC manufacturer may make mistakes when designing example circuits. That's why it's good to have some further information about the underlying philosophies.


Frederik Dostal [frederik.dostal@analog.com] studied microelectronics at the University of Erlangen-Nuremberg, Germany. Starting work in the power management business in 2001, he has been active in various applications positions, including four years in Phoenix, Arizona, working on switch-mode power supplies. He joined Analog Devices in 2009 and works as a power management technical expert for Europe.



**Frederik Dostal**


Also by this Author:

Optimized Power Supply Measurement Setup

Volume 52, Number 3

# Programmable Gain Instrumentation Amplifiers: Finding One that Works for You

By **Kristina Fortunado**

Share on

## Introduction

Data acquisition (DAQ) systems find their use in many industries, for a wide range of applications, such as research, analysis, design verification, manufacturing, and testing. By nature, these systems interface to various sensors, which poses a challenge to the front end. Different sensor sensitivities must be considered—for example, a system may need to interface to a load sensor that has a maximum output of 10 mV and submicrovolt sensitivity, while also interfacing to a sensor that is preconditioned for a 10 V output. With a single gain, the system needs to have a very high resolution to sense both inputs. Even then, signal-to-noise ratio (SNR) is sacrificed at the lowest inputs.

In these applications, programmable gain instrumentation amplifiers (PGIAs) are a good solution for the front end to accommodate the sensitivities of the various sensor interfaces, while optimizing SNR. Integrated PGIAs are available to achieve good dc and ac specifications. This article discusses the various integrated PGIAs and the advantages in using them. Limitations will also be discussed, along with guidelines for building a discrete PGIA when trying to meet a specific requirement.

## Integrated PGIAs

ADI offers a number of integrated PGIAs in its portfolio. Integrated PGIAs provide the benefit of shorter design time and a smaller footprint. The digitally adjustable gain is achieved with internal precision resistor arrays. On-chip trimming of these resistor arrays can be done to optimize gain, CMRR, and offsets, resulting in good overall dc performance. Design techniques can also be utilized for compact IC layout to minimize parasitics, as well as provide excellent matching, resulting in good ac performance. Because of these advantages, it is always recommended to choose an integrated PGIA, if there is one that meets the design requirement. Table 1 lists available integrated PGIAs along with some key specifications.

The choice of PGIA is dependent on the application. The AD825x is very useful in multiplexed systems due to its fast settling time and high slew rate. The AD8231 and LTC6915, which have zero-drift architectures, are useful for systems where precision performance is required over a wide range of temperatures.

**Table 1. Programmable Gain Instrumentation Amplifiers Specifications**

| | AD825x | AD8231 | LTC6915 |
|---|---|---|---|
| Gain Settings | 1, 2, 5, 10 (AD8250)<br>1, 2, 4, 8 (AD8251)<br>1, 10, 100, 1000 (AD8253) | 1 to 128 in 6 dB steps | 1 to 4096 in 6 dB steps |
| CMRR (G = 1) | 80 dB | 80 dB | 125 dB |
| Gain Drift | 10 ppm/°C | 10 ppm/°C | |
| Quiescent Current | 4.5 mA | 4 mA | 2 mA |
| Bandwidth | 10 MHz | 2.7 MHz | 200 kHz |
| Settling time | 0.78 µs | 4 µs | |
| Offset Voltage (G = 1) | 1.05 mV | 45 µV | 10 µV |
| Offset Voltage Drift | 6.2 µV/°C | 50 nV/°C | 50 nV/°C |
| Input Bias Current | 50 nA | 500 pA | 10 nA |
| Noise (G = 1) | 45 nV/$\sqrt{Hz}$ | 66 nV/$\sqrt{Hz}$ | 2.5 µV p-p (0.1 Hz to 10 Hz) |
| Gain Nonlinearity | 6 ppm | 3 ppm | 15 ppm |
| Rail-to-Rail In | No ($-V_s$ + 1) to ($+V_s$ – 1.5) | Yes (0.2 V beyond the rails) | Yes |

**Table 2. DAQ System Specifications**

| | ADAS3022 | ADAS3023 | AD7124-8 |
|---|---|---|---|
| Description | 16-bit, 1 MSPS, 8-channel DAQ system | 16-bit, 8-channel simultaneous sampling DAQ system | 8-channel, low noise, low power, 24-bit Σ-Δ ADC with PGA and reference |
| Gain Settings | 0.16, 0.2, 0.4, 0.8, 1.6, 3.2, 6.4 | 0.2, 0.4, 0.8, 1.6 | 1 to 128 in 6 dB steps |
| CMRR (G=1) | 90 dB | 95 dB | 85 dB |
| Gain Drift | 0.1 ppm/°C | 1 ppm/°C | 2 ppm/°C |
| Power (Max Gain)<br>Conversion Rate (Max Channels) | 12 mA<br>125 kSPS | 10.5 mA<br>125 kSPS | 1.2 mA<br>19.2 kSPS (full power) |

There are also a number of component solutions that integrate the multiplexer, PGIA, and ADC to form a complete DAQ solution. Examples of these are the ADAS3022, ADAS3023, and the AD7124-8.

The choice of these solutions will mainly depend on the specifications of the input signal sources. The AD7124-8 is designed for slower applications that require very high precision, such as temperature and pressure measurements. The ADAS3022 and ADAS3023 are useful in relatively higher bandwidth applications such as in process control or power line monitoring. They do, however, consume more power in comparison to the AD7124-8.

## Implementing a Discrete PGIA

Some systems may need to meet one or two specifications that are not achievable with the aforementioned integrated devices. Typically, the following requirements would necessitate the user to instead build their own PGIA using discrete components:

- Very high scan rates in multiplexed systems that require higher bandwidth
- Ultra low power
- Customized gains or attenuation in the system
- Low input bias current for high impedance sensors
- Very low noise

One of the approaches commonly used in designing discrete PGIAs is to use an instrumentation amplifier with the desired input characteristics (for example, the low noise of the AD8421) accompanied with a multiplexer to switch in the gain resistor to vary the gains.

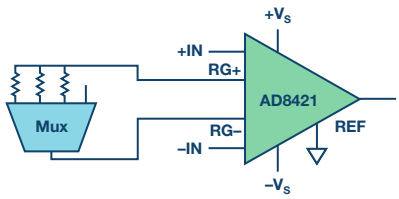


*Figure 1. AD8421 with mux for switched-in gain.*

In this configuration, the on-resistance of the mux is effectively in series with the gain resistor. A problem arises as this on-resistance changes with respect to the voltage present at the drain. Figure 2, taken from the ADG1208 data sheet, demonstrates this relationship.

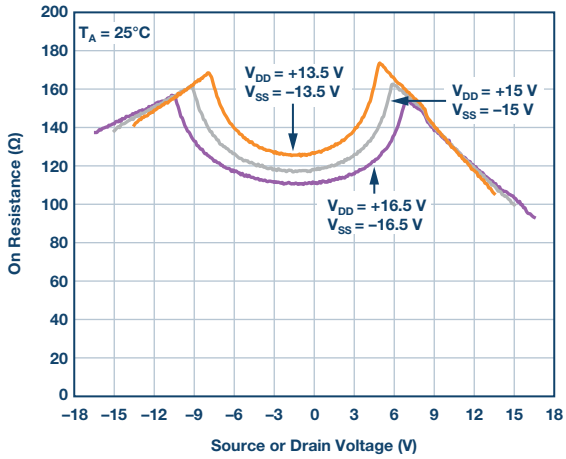


*Figure 2. On-resistance vs. drain voltage for ADG1208.*

The series combination of the on-resistance and the gain resistor results in a nonlinearity in the gain. This implies that the gain will change along with the common-mode voltage, which is not desirable. For example, the AD8421 requires a gain resistor of 1.1 kΩ for a gain of 10. For the ADG1208, the on-resistance changes by as much as 40 Ω with the source or drain voltage varying from ±15 V. This results in a gain nonlinearity of about 3%. For larger gains, this error will become even more apparent, and the on-resistance may even start to become comparable to the gain resistor.

Alternatively, one could use multiplexers with low on-resistance to minimize this effect, but this comes at the price of higher input capacitance. Table 3 demonstrates this with a comparison of ADG1208 and ADG1408.

### Table 3. On-Resistance vs. Capacitance Trade-Off in Multiplexers

| | ADG1208 | ADG1408 |
|---|---|---|
| On-Resistance (typ) | 120 Ω | 4 Ω |
| Drain and Source Capacitance (typ) | 7 pF | 135 pF |

The input capacitance of the switch leads to another problem in the configuration in Figure 1, because the $R_G$ pins on any given three pin op amp in-amp are very sensitive to capacitances. The capacitances of the switch could cause peaking or instability in this circuit. A bigger problem is that the capacitive imbalance on the $R_G$ pins leads to a reduction of ac common-mode rejection ratio (CMRR), which is a key specification of instrumentation amplifiers. The simulation plot in Figure 3 shows how CMRR degrades with different multiplexers across the gain pins of the AD8421. The plot clearly shows that, with increasing capacitance, there is more degradation in CMRR.

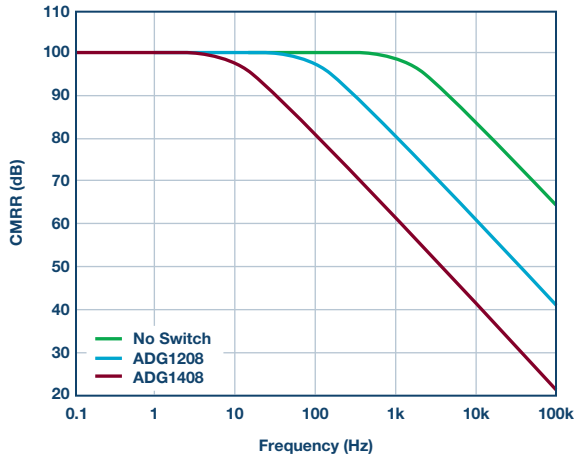


| | No Switch | ADG1208 | ADG1408 |
|---|---|---|---|
| CMRR at 10 kHz | 84 dB | 61 dB | 41 dB |

*Figure 3. Simulated CMRR with different switches.*

In order to alleviate the ac CMRR degradation, the best solution is to ensure that the $R_G$ pins see the same impedance. This can be done by balancing the resistors and placing the switching component in between the two resistors, as shown in Figure 4. A multiplexer does not work in this case due to the inherent capacitive imbalance on the two ends of the switches. Also, because the drains of the multiplexer are shorted together, only a single resistor can be used on one side of the $R_G$ pins, which still leads to imbalances.

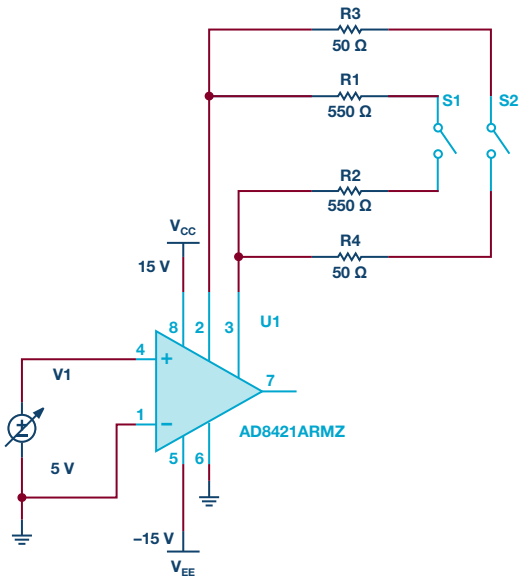


*Figure 4. Discrete PGIA using a balanced configuration.*

A quad SPST switch such as the ADG5412F is recommended in this case. Apart from the fact that the switch gives the flexibility to use balanced resistors, the capacitances are also balanced for the drain and source, which reduces the degradation on CMRR. Figure 5 shows the comparison of the ac CMRR when a multiplexer is used across the gain pins of the AD8421 vs. when a quad SPST switch is used.

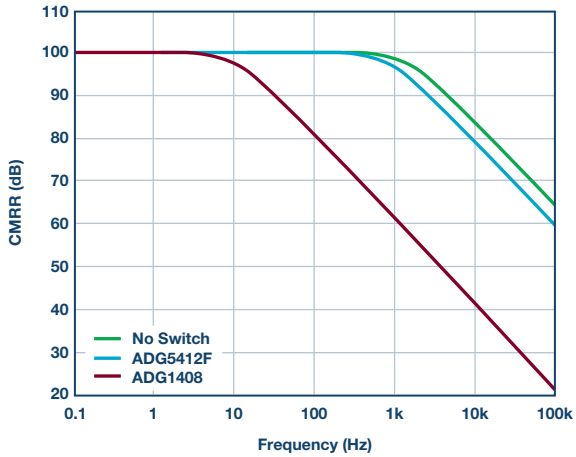


| | No Switch | AD5412F | ADG1408 |
|---|---|---|---|
| CMRR at 10 kHz | 84 dB | 79 dB | 41 dB |

*Figure 5. CMRR simulation with an SPST switch vs. mux configurations.*

The ADG5412F also has low on-resistance, which is very flat across drain or source voltage, as shown in Figure 6. It is specified to change to a maximum of 1.1 Ω over the drain or source voltage. Going back to the initial example, with the AD8421 in a gain of 10 and a gain resistor of 1.1 kΩ, the switch would introduce only 0.1% of gain nonlinearity. There is still a component of drift though, which will be more pronounced at higher gains.

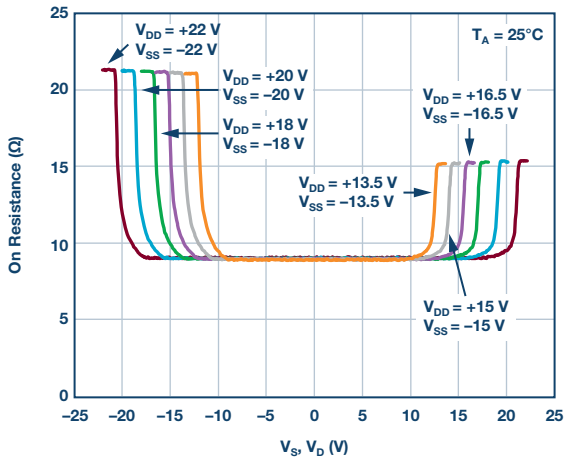


*Figure 6. On-resistance of ADG5412F over common-mode voltage.*

In order to eliminate the parasitic resistance effects of the switch, an in-amp with a different architecture can be used to implement arbitrary gains. The AD8420 and AD8237 have an indirect current feedback (ICF) architecture, and are good choices for applications that require low power and low bandwidth. In this configuration, the switch is placed in a high impedance sense path, so the gain is not affected by the changing on-resistance of the switch.

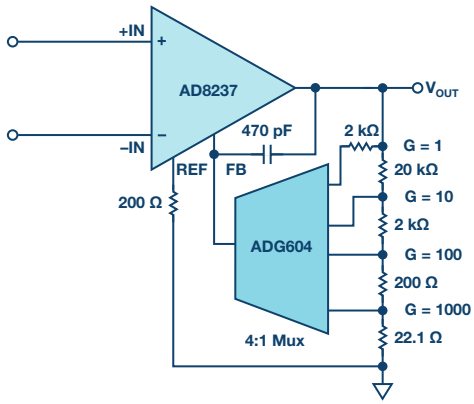


*Figure 7. Discrete PGIA using in-amps with indirect current feedback.*

For these amplifiers, the gain is set with the ratio of the external resistors in the same way as a noninverting amplifier. This gives the user more flexibility as the gain-setting resistors can be chosen depending on the design requirement. Standard thin film or metal film resistors can have temperature coefficients as low as 15 ppm/°C. This gives better gain drift than standard instrumentation amplifiers that set the gain with a single external resistor, where mismatch between the on-chip and external resistors typically limit the gain drift to around 50 ppm/°C. For the best gain error and drift performance, a resistor network can be used for its tolerance and temperature coefficient tracking. This does come at the expense of cost, though, so unless it is needed, discrete resistors are preferred.

Another solution, and the one that provides the most flexibility, is the three op amp in-amp architecture with discrete components, as shown in Figure 8, with multiplexers to switch in gain resistors. There is a much larger variety of operational amplifiers to choose from compared to instrumentation amplifiers, so designers have more choices, which allows them to design around a specific design requirement. Special functions such as filtering can also be built into this first stage. A difference amplifier in the second stage completes this architecture.

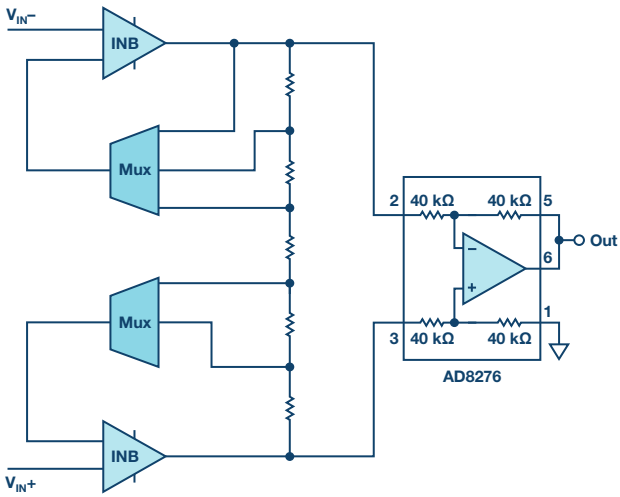


*Figure 8. Discrete PGIA.*

The choice of the input amplifiers is directly dependent on the DAQ requirement. For example, a low power design will require the use of amplifiers with low quiescent current, while systems that expect to see high impedance sensors at the input might make use of amplifiers with very low bias current to minimize error. Dual amplifiers should be used for better tracking over temperature.

It can be noted that, when using the configuration in Figure 8, the on-resistance of the switch is also in series with the high impedance input of the amplifier, and therefore it does not affect the gain. Looking back to the trade-off between on-resistance and switch input capacitance, since the constraint on the on-resistance is removed, low input capacitance switches, such as ADG1209, can be chosen for the design. In this way, instability and ac CMRR degradation is avoided.

As in the previous design, the resistors will dictate the gain accuracy and drift. Discrete resistors can be chosen with the right tolerance and drift that correspond to the design requirements of the application. Again, better accuracy can be achieved with resistor networks for better tolerance and temperature tracking at the expense of cost.

The second stage of the three op amp in-amps takes care of the rejection of common-mode voltages. Difference amplifiers, which have integrated resistor networks, are recommended for this stage to ensure best CMRR. For a single-ended output and relatively low bandwidth applications, the AD8276 is a good choice. The AD8476 can be used if a differential output and higher bandwidth are required. Another option for the second stage is to use the LT5400 as gain-setting resistors around a standard amplifier. While this may take up more board space, this again gives you more flexibility with the choice of your amplifier, allowing more capability to design around a specific design requirement.

It should be noted that care needs to be practiced in the layout for discrete PGIAs. Any imbalances in the layout of the board will cause the CMRR to degrade over frequency.

The table below gives a summary of the strengths and weaknesses of each method:

## Discrete PGIA Design Example

An example of a discrete PGIA built for a particular design specification is given in Figure 9. In this design, the PGIA should be built for very low power. For the input buffers, the LTC2063 was chosen for its low supply current of 2 µA max. For the switching component, the ADG659 was chosen for its low supply current of 1 µA max and low input capacitance.

Care also needs to be taken in choosing the passive components in the circuit—these also need to be chosen such that they meet the low power requirement. Not choosing passives appropriately will lead to larger current

**Table 4. Comparison of Different Programmable Gain Instrumentation Amplifier Implementations**

| Method | Strengths | Weaknesses |
|---|---|---|
| Integrated PGIA | ▶ Least design effort<br>▶ Optimized for good ac and dc performance<br>▶ Guaranteed CMRR performance<br>▶ Requires smaller board space<br>▶ Single component solution | ▶ Limited choices available |
| Balanced Configuration (In-Amp with Switch) | ▶ Minimal design effort compared to discrete three op amp in-amp | ▶ Prone to instability if not designed properly<br>▶ Less choices for in-amp compared to discrete three op amp in-amp<br>▶ On-resistance of switch contributes to gain error and drift |
| Indirect Current Feedback with Mux | ▶ Minimal design effort compared to discrete three op amp in-amp<br>▶ Easily configurable<br>▶ On-resistance of switch does not contribute to gain error and drift | ▶ Limited input range<br>▶ Less flexible, small number of ICF in-amps available |
| ADA4077 and Clamping OVP | ▶ Most flexible because of more op amp choices compared to in-amps<br>▶ Highly configurable—can incorporate filtering, differential output, etc.<br>▶ On-resistance of switch does not contribute to gain error and drift | ▶ More components required<br>▶ More board space<br>▶ Much higher design effort required to achieve key specifications |

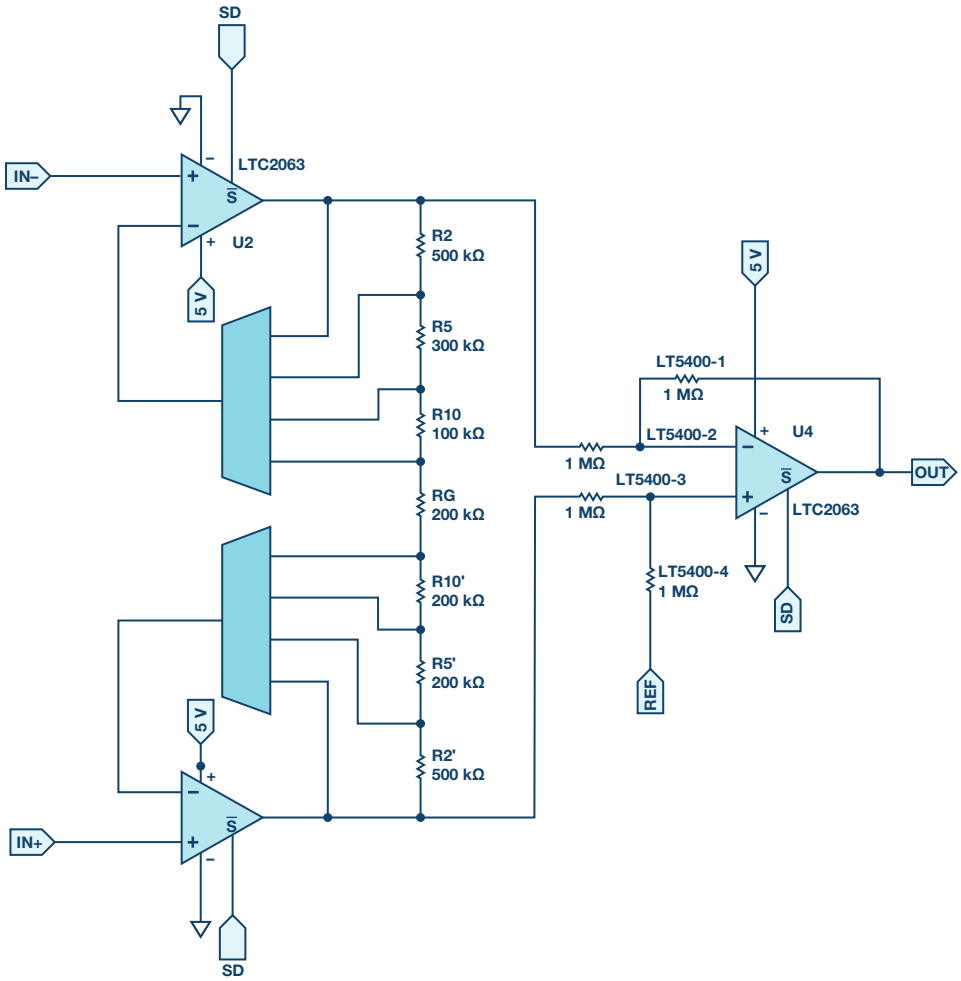


*Figure 9. Low power PGIA design.*

draw, which would negate the effect of using low power components. The gain resistors, in this case, need to be large enough so as not to draw too much current. The chosen resistor values, which were set to provide gains of 1, 2, 5, and 10, are shown in Figure 9.

For the second stage difference amplifier, the LTC2063 was used with the LT5400 quad matched resistor network, 1 MΩ option. This ensures that minimal current is drawn, and that CMRR is preserved due to the precise matching of the resistors.

The circuit runs on a 5 V supply and was evaluated with different common-mode voltages, differential input voltages, and gains. In the best condition where the reference and the inputs are held at mid-supply, the circuit draws only 4.8 µA of current.

Some increase in current is expected at varying differential inputs due to the current that flows through the gain resistors, given by $|V_{OUT} - V_{REF}|/(2\ M\Omega||1\ M\Omega)$. Figure 10 below shows the current drawn at different gains. Data is taken with respect to the output to account for the gains.

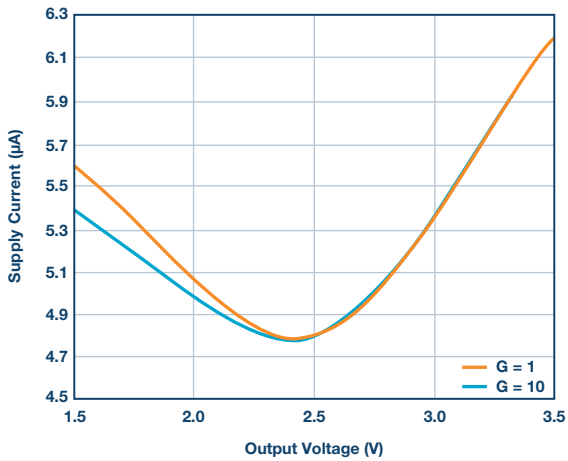


*Figure 10. Supply current over output voltage.*

An additional increase in current will also be expected when different common-mode voltages are applied to the inputs. The voltage applied will cause a current to flow through the resistors in the second stage leading to additional current draw. This is given by $|V_{CM} - V_{REF}|/1\ M\Omega$. The choice of 1 MΩ resistors for the LT5400 was made specifically to minimize this. Figure 11 below shows the effect of the common-mode voltage on the current drawn at different gains:

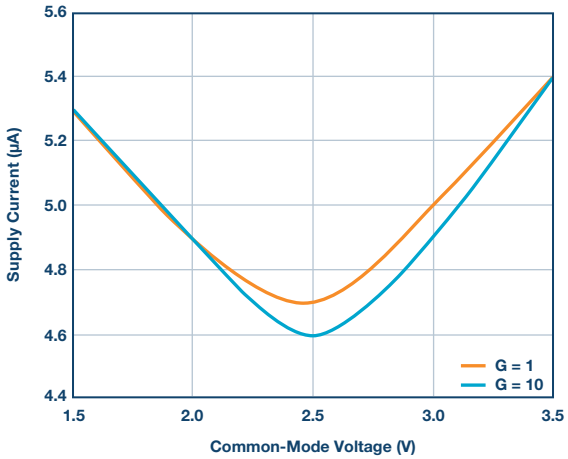


*Figure 11. Supply current over common-mode voltage.*

The static current of the circuit was also measured in shutdown mode. When all components are shut down, the circuit draws only 180 nA of current. This does not vary, even if variables such as common-mode voltage, reference, and differential inputs are changed, for as long as they are all kept within the supply. All components have the option for power down in case additional power needs to be saved and the user would like to do power-cycling. In portable, battery-powered applications, this circuit would be extremely useful and the key specification would not otherwise have been achieved with an integrated PGIA.

## Conclusion

Programmable gain instrumentation amplifiers are a critical component in the data acquisition space, enabling good SNR performance, even with varying sensor sensitivities. The use of integrated PGIAs allows for shorter design time and better overall dc and ac performance for the front end. Integrated PGIAs should generally be preferred in the design, if there is one that meets the requirements. However, when system requirements dictate specifications that are not attainable with available integrated options, it is possible to design a discrete PGIA. By following the right design recommendations, the optimal design can be achieved even with a discrete approach, and the various implementations can be assessed to determine the best configuration in a specific application.

The author would like to thank Scott Hunt and Paul Blanchard for their technical contributions to this article.

Kristina Fortunado [kristina.fortunado@analog.com] joined ADI in 2009 and works as a product applications engineer for the Linear Products and Solutions Group. She graduated from De La Salle University with a bachelor's degree in electronics and communications engineering.

Kristina Fortunado

# Remote Sensing Using a High Precision Instrumentation Amplifier

By **Hooman Hashemi**

Share on


## Abstract

The instrumentation amplifier (IA) is the workhorse of sensing applications. In this article, I will explore some ways to take advantage of these amplifiers' balance and excellent dc/low frequency common-mode rejection (CMR) for use with resistive transducers (for example, strain gage) when the sensor is physically separated from the amplifier. I will present methods to increase the noise immunity of such gain stages while making them less sensitive to supply variation and component drift. Measured performance values and results will also be presented to show the accuracy range in order to allow a quick evaluation for end-user applications.


## Detail

When it comes to sensors, there is little competition for what a Wheatstone bridge (Figure 1) can do. The bridge can produce a differential voltage that predictably changes in response to changes in a physical parameter—with the side benefit of providing temperature and time drift immunity. The differential voltage rides atop a large common-mode (CM) voltage. To amplify the small signal from a bridge, an instrumentation amplifier is used. The beauty of an IA is that, with little or no loading of the bridge elements, it can sense the differential voltage and reject the CM to a degree that is next to impossible to achieve by a traditional op amp, due to the high degree of external resistor matching required.

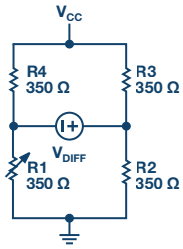


*Figure 1. Wheatstone bridge.*

The electronics involved in physical measurements are often located far from the physical parameter under measurement. For instance, a strain gage measurement, such as that buried under the tarmac at a truck weigh station or within the structure of a bridge, is unlikely to be located next to the electronics used to read the measurement. For example, when dealing with a two-wire quarter-bridge strain gage such as OMEGA's SGT-1/350-TY43, placing the sensor remotely from the sensing amplifier, as shown in Figure 2, yields unsatisfactory results, even if shielded twisted pair sensor leads are used.

The problem is that shielded twisted pairs are not immune to all interference over long cable runs. In such a case, the well-balanced input(s) of the instrumentation cannot be relied upon to eliminate the CM pickup. The positive and negative amplifier inputs are not equally affected by the interference picked up by the long cable, and the inputs contain uncorrelated signals that CMR cannot eliminate. Therefore, it's not a surprise to find significant noise at the output of the circuit, as shown in Figure 3, due to this unbalanced response to what seems to be CM noise.

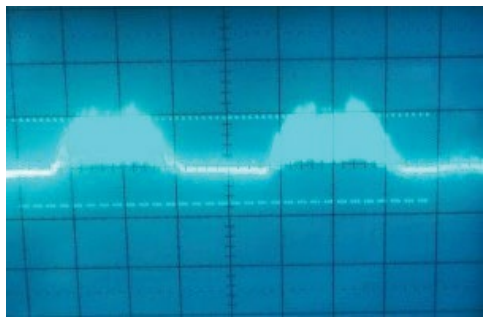

*Figure 3. Troublesome 120 Hz pickup at the amplifier output (0.1 V/div, 2 ms/div).*

One solution for the successful extraction of the small bridge differential voltage from the CM (dc and interference) is to use two pairs of shielded or unshielded twisted pair (UTP). Done this way, both IA inputs are balanced and subjected to the same CM noise pickup. This is shown in Figure 4. A device such as the LT6370, with excellent low frequency CMR (120 dB) can reliably go to work on rejecting what afflicts both IA inputs. The result is a clean output waveform at a large distance, even in noisy environments.

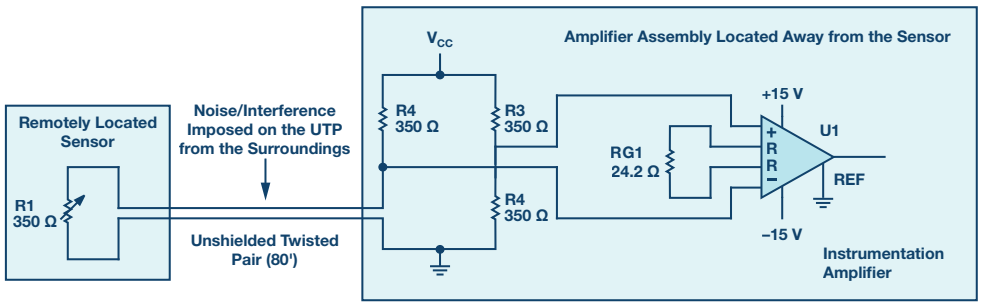


*Figure 2. Remotely located sensor setup suffers from environmental noise pickup.*

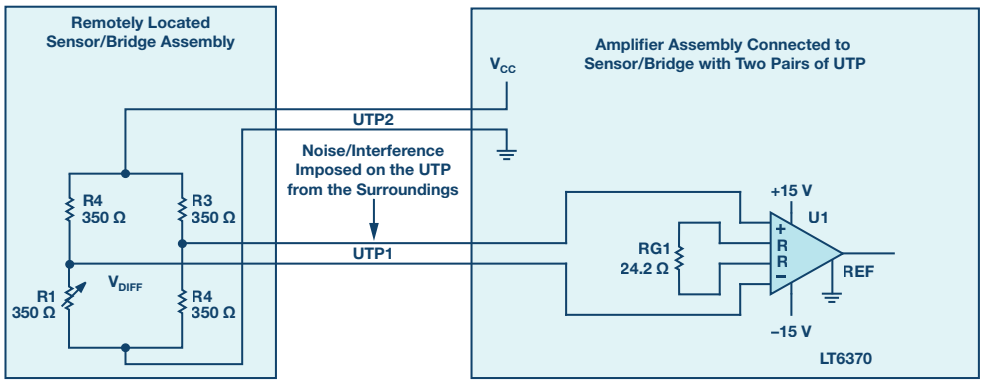


*Figure 4. Remote sensing using two unshielded twisted pairs of wire.*

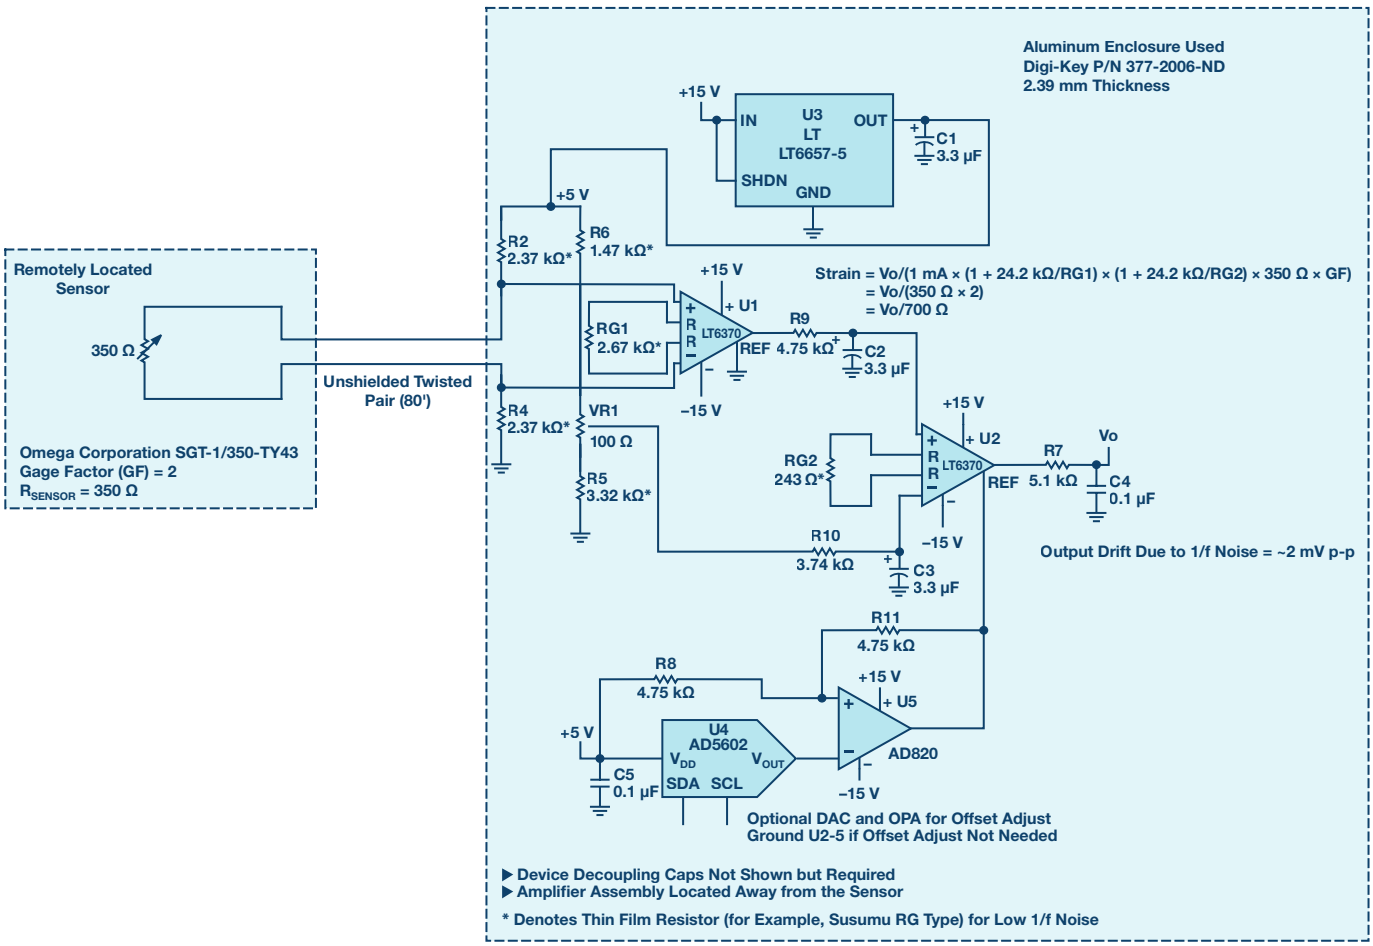


*Figure 5. A single UTP for remote sensing.*

Having all the CMR horsepower of LT6370 at hand, one could take this idea one step further and streamline the configuration by eliminating one pair of wiring, leaving a single UTP. This concept is illustrated in Figure 5, where U2's inputs are kept balanced for good CMR. Note how the UTP leads look identical to U2 and have identical impedance to ground (R2, R4).

For the component values shown in Figure 5, there will be about 1 mA flowing through the sensor $R_{SENSOR}$. With U1's RG1 value, that stage runs at G = 10 V/V and provides a 10× replica of the voltage across $R_{SENSOR}$ at its output voltage, about 3.5 V. U1's main task is to eliminate the interference present on the UTP long length of wire and responding only to the sensor voltage which is the sensor resistance times the ~1 mA current flowing through it. LT6370's excellent low offset voltage and drift along with its exceptional CMR make it the obvious choice.

The other half of the Wheatstone bridge is comprised of R5, R6, and VR1 with near identical current flow as the sensor half of the bridge. Both the sensor voltage at U1 output and the reference voltage at VR1 wiper reach the differential inputs of U2 after some low-pass filtering to eliminate unwanted noise. U2 is set for high gain (G = 1 + 24.2 kΩ/RG2 = 100 V/V) to magnify the very small sensor voltage on its positive input compared against the fixed, low noise reference voltage, derived from the LT6657-5 voltage reference, on its negative input. U1 output accurately represents the measured strain applied to the sensor, attached to the element or material of interest, to drive an ADC or other similar signal processing.

The optional DAC and OPA (U4, U5) tied to U2's REF pin (which can be grounded if no offset adjust is needed) can be used to provide output offset adjustment and zeroing. By using the DAC, it is possible to shift the U2 output voltage to a desired pedestal or CM level suitable for the selected ADC. For example, an ADC with a reference voltage of 5 V can be driven directly from U2 with its zero output set to 2.5 V using the DAC driving

U2 REF input. Done this way, 0 V to 2.5 V ADC analog input represents compression and 2.5 V to 5 V signal represents tension strain. It is important to note that the device driving U2 REF pin, AD820 in this case, should maintain a low impedance to eliminate any possible gain errors.

Here is the expression for the output voltage as a function of the sensor resistance and relationship between the output voltage and the strain (ε) being measured:

$$\Delta V_O = 1\ \text{mA} \times \Delta R_{SENSOR} \times G1 \times G2$$

where $\Delta R_{SENSOR}$ is the change in sensor resistance due to strain

$$G1 = 1 + 24.2\ \text{k}\Omega/RG1 = 10\,\frac{\text{V}}{\text{V}}$$

$$G2 = 1 + 24.2\ \text{k}\Omega/RG2 = 100\,\frac{\text{V}}{\text{V}}$$

$$\frac{\Delta V_O}{\Delta R_{SENSOR}}\left(\frac{\text{V}}{\Omega}\right) = 1\ \text{mA} \times 10\frac{\text{V}}{\text{V}} \times 100\,\frac{\text{V}}{\text{V}} \approx 1\,\frac{V}{\Omega}$$

$$\textit{Gage Factor (GF)} = \frac{\frac{\Delta R_{SENSOR}}{R_{SENSOR}}}{\frac{\Delta L}{L}} = \frac{\frac{\Delta R_{SENSOR}}{R_{SENSOR}}}{\varepsilon}$$

Where:

L refers to the sensor length

ε refers to the amount of strain being measured

For the sensor chosen:

$R_{sensor}$ = 350 Ω

GF= 2

Resulting in the strain (ε):

$$\varepsilon = \frac{\Delta V_O}{1\ \text{mA} \times 1000\,\frac{\text{V}}{\text{V}} \times R_{SENSOR} \times GF} = \frac{\Delta V_O}{700}$$

The LT6370's exceedingly low gain error (<0.084% at G = 10 V/V) and low input offset voltage (<50 µV max specified over temperature) guarantee that U2 is presented with a true replica of the sensor voltage, minus the interference picked up by the UTP, to compare against the reference voltage developed at U2's inverting input. The LT6657-5 creates a stable, low noise, and low drift voltage reference, immunizing the entire circuit from supply voltage variation. Of particular importance is the LT6657-5's low 1/f noise, which can have a significant contribution due to the large gain in the circuit.

With the simple RC low-pass filters (R9, C2 and R10, C3) set to roll-off at about 10 Hz on each input of U2, the output noise can be reduced by limiting the bandwidth. The low (<10 Hz) LT6370 1/f noise corner frequency, as shown in Figure 6, provides an advantage by reducing the impact of the 1/f noise. Furthermore, the current noise density plot shows that it is much better to keep both input impedances balanced for the lowest current noise impact by taking advantage of the correlated component of noise at the input(s). Hence the value for R10 is reduced to 3.74 kΩ in order to match the R9 impedance of 4.75 kΩ due to the equivalent impedance looking into the wiper of VR1.

## Summary

Placing a bridge sensor at a distance from a signal processing amplifier requires an instrumentation amplifier that can cleanly extract the measured differential voltage. The attributes of the LT6370 instrumentation amplifier enable it to process successfully signals from distant sensors via long cable runs. The LT6370 manufacturing process, which invokes on-chip heaters during production testing to guarantee over temperature drift values, further enhances the LT6370 suitability to remote monitoring applications, and improves longevity and product life in hard-to-service installations.

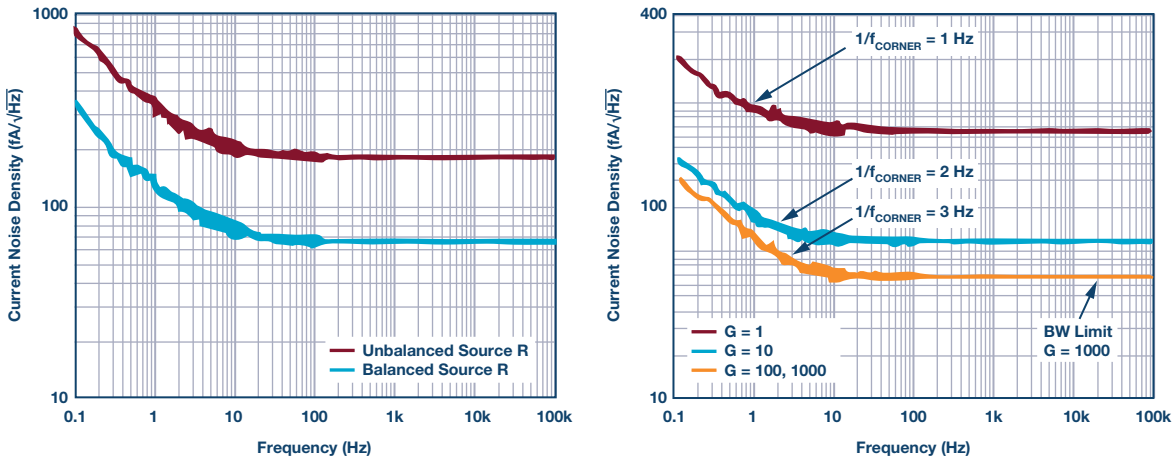


*Figure 6. LT6370 input referred current/voltage noise density.*


Hooman Hashemi [hooman.hashemi@analog.com] joined Analog Devices in March 2018, where he works on characterizing new products and developing applications that showcase the products features and uses. Hooman previously worked for Texas Instruments for 22 years as an applications engineer, concentrating on the high speed portfolio. He graduated from University of Santa Clara with an M.S.E.E. in August, 1989, and San Jose State University with a B.S.E.E. in December, 1983. He is married with one 14-year old daughter.



Hooman Hashemi

# Overcoming Constraints: Design a Precision Bipolar Power Supply on a Simple Buck Controller

By **Victor Khasiev**

Share on

## Introduction

Industrial, automotive, IT, and networking companies are major purchasers and consumers of power electronics, semiconductors, devices, and systems. These companies use the full array of available topologies for dc-to-dc converters that employ buck, boost, and SEPIC in different variations. In an ideal world, these companies or firms would use a specialized controller for each new project. However, adopting new chips requires significant investment due to the lengthy and costly process of testing new devices for compliance with automotive standards, verification functionality in the specific applications, conditions, and equipment. The obvious solution for reducing development and design cost is employing already approved and verified controllers in different applications.

The most used topology for generating a power supply is for step-down converters. However, employment of this topology is limited to generating positive outputs from the input voltages that are greater than the output. It cannot be used in a straightforward way for generating negative voltages or providing stable outputs when the input voltage drops below the output. Both aspects to generating output are important in automotive electronics when negative voltage is needed for supplying amplifiers or when a complete system must continuously work properly in case of cold cranking when the input voltage rails drop significantly. This article details a method for using a simple buck controller in SEPIC, Cuk, and boost converters.

## Generating Negative and Positive Voltage from a Common Input Rail

Figure 1 illustrates the design of a bipolar power supply based on a single buck controller with two outputs.

For maximum utilization of this chip, one output must be employed to generate a positive voltage and a second to generate negative voltage. The input voltage range of this circuit is 6 V to 40 V. The $V_{OUT1}$ generates positive 3.3 V at 10 A and $V_{OUT2}$ negative voltage –12 V at 3 A. Both outputs are controlled by U1. The first output $V_{OUT1}$ is the straightforward buck converter. The second output has a more complex structure. Because $V_{OUT2}$ is negative relative to GND, the differential amplifier U2 is employed to sense negative voltage and scale it to the 0.8 V reference. In this approach, both U1 and U2 are referenced to the system GND, which significantly simplifies the power supply's control and functionality. The following expressions help to calculate the resistor values for RF2 and RF3 in case a different output voltage is required.

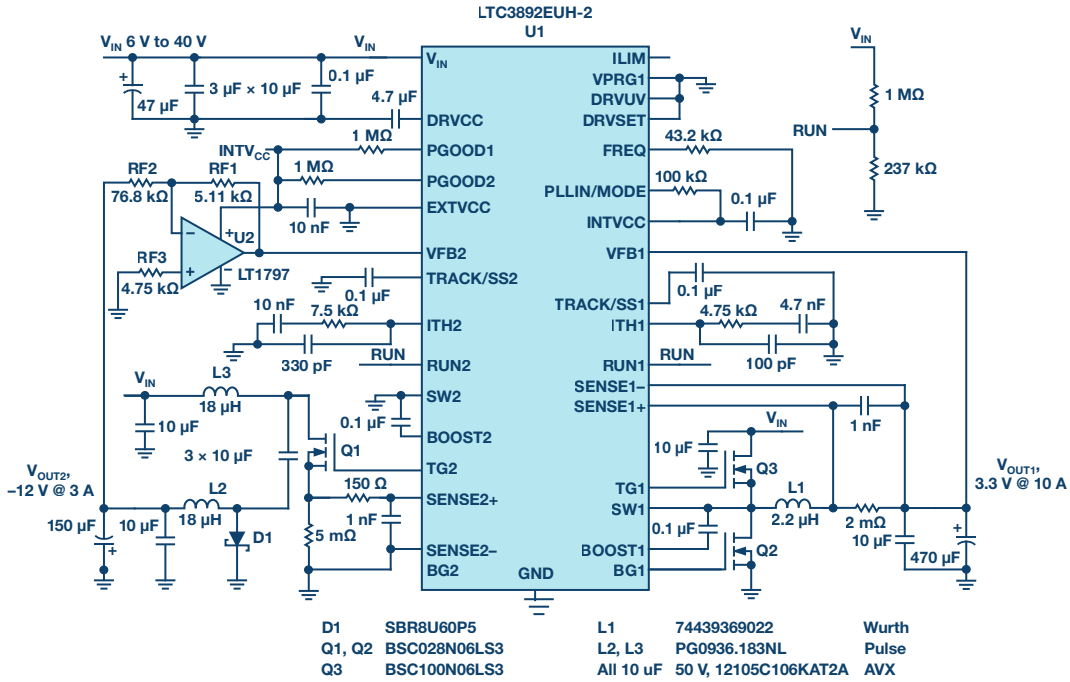


*Figure 1. An electrical schematic of LTC3892 that is generating positive and negative voltages. $V_{OUT1}$ is 3.3 V at 10 A and $V_{OUT2}$ is –12 V at 3 A.*

$$KR = \frac{0.8\ \text{V}}{|V_O|}$$

$$RF1 = 5.11\ \text{k}\Omega$$

$$RF2 = \frac{RF1}{KR}$$

$$RF3 = \frac{RF1 \times RF2}{RF1 + RF2}$$

The $V_{OUT2}$ power train employs a Cuk topology, which is widely covered in the relevant technical literature. The following basic equations are required to understand the voltage stress on the power train components.

$$D = \frac{|V_O|}{|V_O| + V_{IN}}$$

$$V_C = \frac{V_{IN}}{1 - D}$$

$$V_{DS} = V_D = V_C$$

The $V_{OUT2}$ efficiency curve is presented in Figure 2. The LTspice® simulation model of this approach is available here. In this example, the LTC3892 converter's input is 10 V to 20 V. The output voltages are +5 V at 10 A and –5 V at 5 A.

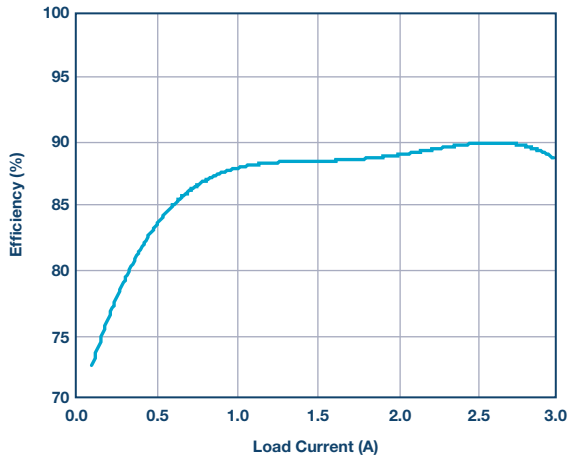


*Figure 2. Efficiency curve of the negative output at 14 V input voltage.*

## Generating Stable Voltages from a Fluctuating Input Rail

The electrical schematic of the converter shown in Figure 3 supports two outputs: $V_{OUT1}$ with 3.3 V at 10 A and $V_{OUT2}$ with 12 V at 3 A. The input voltage range is 6 V to 40 V. $V_{OUT1}$ is created in a similar fashion, as shown in Figure 1. The second output is a SEPIC converter. This SEPIC converter, as with Cuk above, is based on noncoupled, dual discrete inductor solutions. Use of the discrete chocks significantly expands the range of the available magnetics, which is very important for cost-sensitive devices.

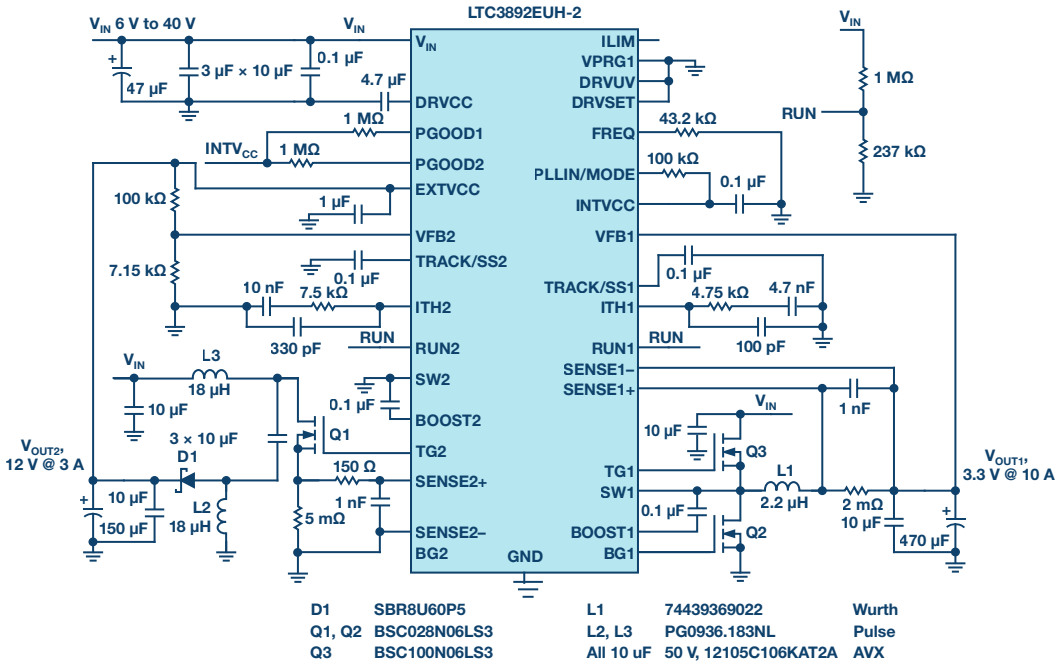


*Figure 3. Electrical schematic of LTC3892 in a SEPIC and in buck applications.*

$$D = \frac{V_O}{V_O + V_{IN}} \qquad V_C = V_{IN}$$

$$V_{DS} = V_D = V_{IN} + V_O$$

Figure 4 and Figure 5 illustrate the functionality of this converter at voltage drops and spikes; for example, at cold cranking or load dumps. The rail voltage $V_{IN}$ drops or rises at a relatively nominal 12 V. However, both $V_{OUT1}$ and $V_{OUT2}$ stay in regulation and provide a stable power supply to the critical loads. The two-inductor SEPIC converter can be easily rewired to a single inductor boost converter.

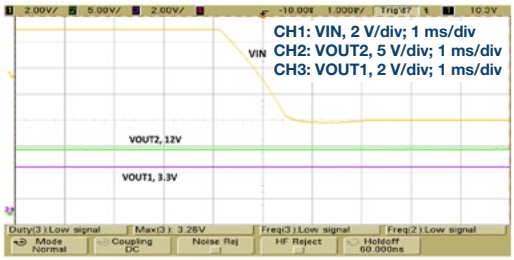


*Figure 4. If the rail voltage drops from 14 V to 7 V, both $V_{OUT1}$ and $V_{OUT2}$ stay in regulation.*

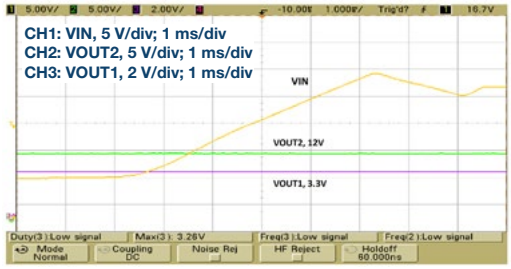


*Figure 5. The rail voltage rises from 14 V to 24 V. However, both $V_{OUT1}$ and $V_{OUT2}$ stay in regulation.*

The relevant LTspice simulation model can be found here. It shows the LTC3892 converter's input is 10 V to 20 V. The output voltages are +5 V at 10 A and –5 V at 5 A.

## Conclusion

This article explained the methods of building bipolar and dual-output power supplies based on the step-down controller. This approach allows for the use of the same controller in buck, boost, SEPIC, and Cuk topologies. This is very important for vendors of automotive and industrial electronics, as they can design power supplies with a variety of output voltages based on the same controller, once it is approved.


Victor Khasiev [victor.khasiev@analog.com] is a senior applications engineer at ADI. Victor has extensive experience in power electronics both in ac-to-dc and dc-to-dc conversion. He holds two patents and has written multiple articles. These articles relate to the use of ADI semiconductors in automotive and industrial applications. They cover step-up, step-down, SEPIC, positive-to-negative, negative-to-negative, flyback, forward converters, and bidirectional backup supplies. His patents are about efficient power factor correction solutions and advanced gate drivers. Victor enjoys supporting ADI customers, answering questions about ADI products, designing and verifying power supply schematics, laying out print circuit boards, troubleshooting, and participating in testing final systems.



Victor Khasiev

# Rarely Asked Questions—Issue 160 Simple Circuit Measures Relative Intensity of Two Light Sources

By **Chau Tran**

Share on

## Question:

Can I measure the difference between two light sources with an instrumentation amplifier?

## Answer:

Yes, by replacing the main setting resistor of an instrumentation amplifier with two photoresistors.

In many lighting applications, measuring the relative intensity of two light sources is more important than measuring their individual intensity. This ensures that two light sources shine with the same intensity. For example, it is helpful to compare the brightness inside a control room (Room 1) to a second room (Room 2) in the same building so the adjustment can be made any time of day as well as night, or, if you have a production system, you might want to ensure that the bright light conditions do not change.

One way to determine the relative intensity is to measure the different outputs of two additional light detectors. Their difference will be converted to a single-ended voltage signal with ground reference.

The circuit in Figure 1 shows a simple but effective way to solve this problem, by using an instrumentation amplifier with resistor gain control, such as the AD623.

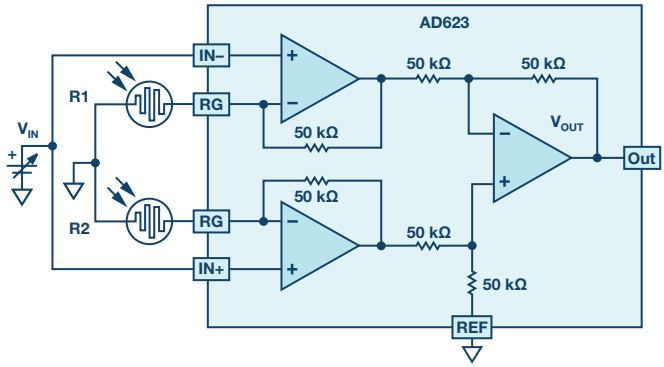


*Figure 1. A simple circuit that measures relative light densities.*

Note that in this circuit, there are two special resistors, R1 and R2. R1 (LDR1) measures the brightness of the two light sources.

What is an LDR? The term stands for light dependent resistor or photoresistor. It is a passive electronic component with a resistor that has variable resistance depending on light intensity. Light dependent resistors come in different shapes and colors and they are useful in many electronic circuits, especially in alarms, switching devices, clocks, and street lights.

In general, LDRs' resistance is very high in darkness—almost high as 1 MΩ—but when light falls on the LDR, the resistance falls to a few kΩ (10 kΩ to 20 kΩ at 10 lux, 2 kΩ to 4 kΩ at 100 lux), depending on the model.

The LDR used for this schematic is from RadioShack (part #276-1657).

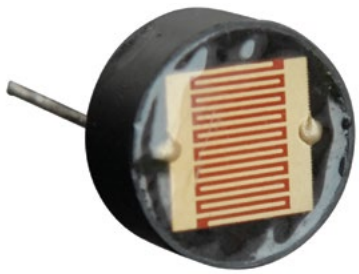

*Figure 2. RadioShack part #276-1657.*

The schematic in Figure 1 employs an AD623 and two LDRs. As the main sensor, photoresistor R1 is a reference point light source. It is used as a baseline for light intensity and it is located in the control room. If you are comparing more than two light sources, you should use this light source as a reference during every comparison. This comparison can happen at night or during the day. Keep in mind that it takes 8 ms to 12 ms for a change in resistance to take place. It also take seconds for the resistance to change back to its initial value.

The design is very simple. The power supplies of the system are ±5 V and the input voltage at both inputs is $V_{IN}$. Therefore, the same voltage is on one end of each photoresistor and ground on their other end. If the same amount of light fell on both photoresistors, the current difference between them would be zero because their resistances are equal. The result is that the output voltage is at zero volts. When the two rooms are not equally illuminated, there is a difference between the intensities of the two light sources and this creates a voltage at the output of the system. The polarity of this voltage indicates which room is brighter. If the output voltage was positive, this would mean that more light has fallen on LDR2 and vice versa.

Figure 3 shows a scope plot of the output waveform. With the input voltage is 1 V p-p square wave at a frequency of 1 kHz, the output (about 2 V) indicates that the light source in Room 2 is brighter.

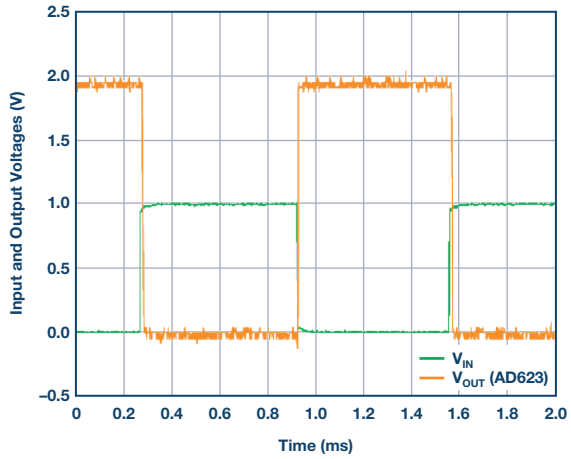


*Figure 3. Output voltage indicates the relative light intensity.*

In the circuit in this article, there are two LEDs at the output of the AD623. The red LED with the positive side connected to the output would turn on when the output is positive (light source 2 is brighter) and the yellow LED with the positive side to ground would be turned on when light source 1 is brighter. Note that the brightness of the LED shows the amplitude level of the relative intensity of the room.

When both rooms appear equally bright, their illuminance is equal and the output is 0 V while both LEDs are off.

The voltage at the output of the circuit is:

$$V_{OUT} = V_{IN} \times 50\ \text{k}\Omega \times \left(\frac{1}{LDR1} - \frac{1}{LDR2}\right)$$

The rms value of the output is the intensity level of two light sources.

After calibrating the value of LDR1 to find the precise resistance value at a certain brightness, LDR1 can be replaced by a pure resistor and then the system will compare the value of LDR2 to a certain brightness all the time. The fixed resistor then works as a reference for a known light reference. This kind of circuit can be a solar seeker—namely, a simple device that tracks a light source. Such a device can keep solar panels aligned with the Sun, or be used in search and rescue robots that try to guide trapped people toward light. In order to implement the solar seeker, a servo motor that rotates the photoresistors can be used.

Using AD623, one can find out the efficiency of two light bulbs by putting them in the different rooms together with LDR1 and LDR2, respectively. This circuit has low power consumption and can be powered just by two AA batteries, which is useful in power sensitive applications.


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



**Chau Tran**

Also by this Author:

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

Volume 51, Number 3

# Notes

# Notes

# Stay Connected with *Analog Dialogue*

## Your Engineering Resource for Innovative Design

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

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

AnalogDialogue

## EngineerZone® Support Community

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

EngineerZone®
SUPPORT COMMUNITY

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

**Analog Devices, Inc.**
**Europe Headquarters**

Analog Devices GmbH
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


M02000-0-3/19

*analog.com*

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™