Analog Dialogue

Volume 53, Number 3, 2019

Your Engineering Resource for Innovative Design

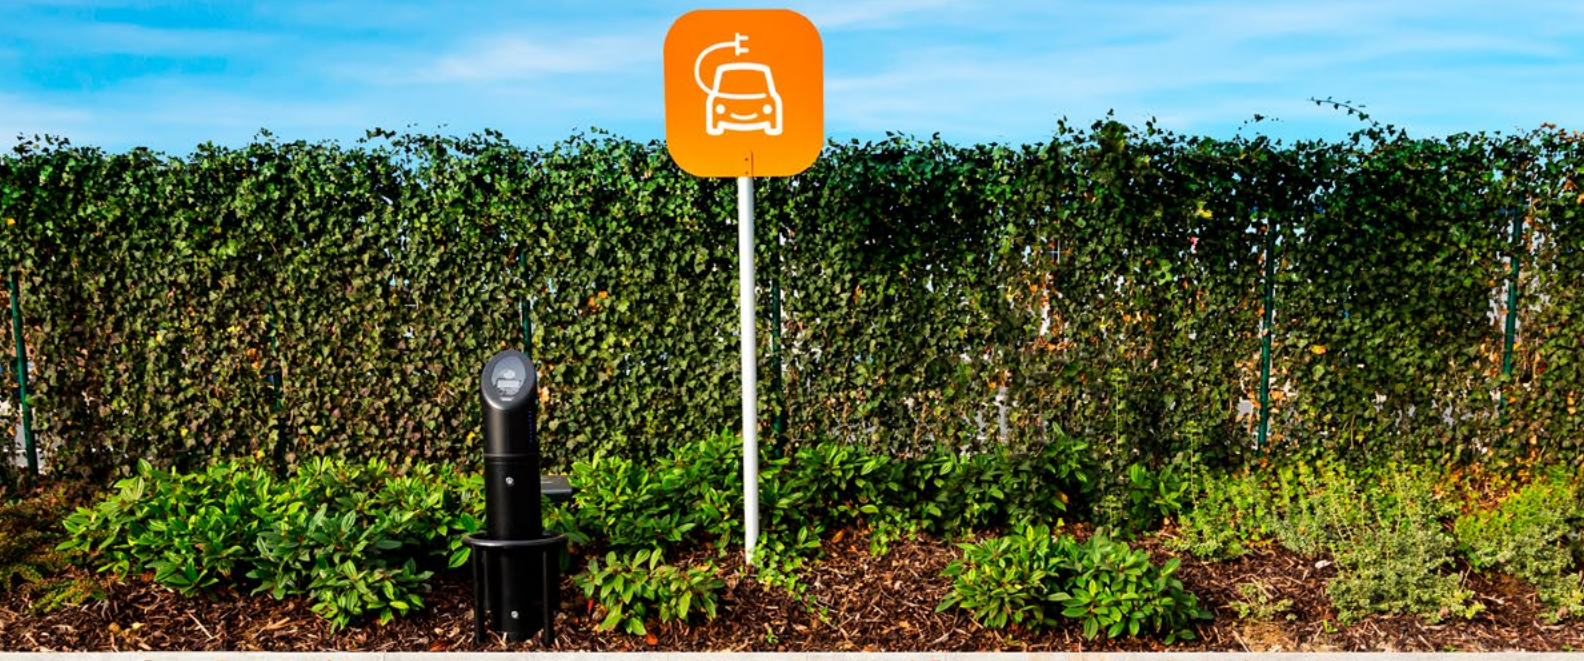

Visit *analogdialogue.com*

ANALOG DEVICES

AHEAD OF WHAT'S POSSIBLE™

In This Issue

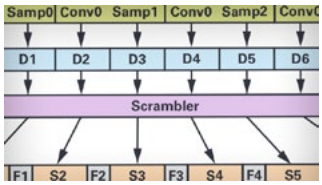

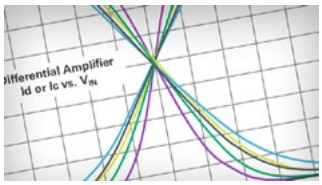

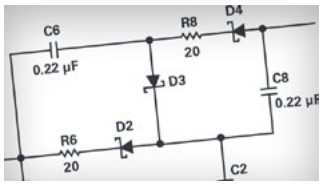

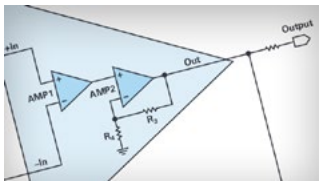

Bernhard Siegel, Editor

Bernhard became editor of *Analog Dialogue* in March 2017. He has been with Analog Devices for over 30 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*.

JESD204C Primer: What's New and in It for You—Part 2

Del Jones

In Part 1 of the JESD204C primer series, the new version of the JESD204 standard was justified by describing some of the problems it solves. The differences between the B and C version of the standard were summarized by describing new terminology and features, then providing a layer-by-layer overview of these differences. Since Part 1 laid a suitable foundation, let's take a closer look at a few of the more notable new features of the JESD204C standard.

The 64b/66b and 64b/80b Link Layers

For the 64b/66b link layer, the 66-bit block of data is two sync header bits followed by eight octets of sample data and is partially based on the block format defined in IEEE 802.3, clause 49. Unlike the IEEE standard, there is no encoding—the payload data is simply the converter sample data that has been packed into frames of data by the transport layer. Since there is no encoding to ensure a certain number of data transitions occur to provide dc balance, the sample data must be scrambled. These scrambled octets of framed data are placed directly into the link layer with the two sync header bits attached.

The 64b/66b block format is illustrated in Figure 1. The example shows the case where one lane of data consists of frames that contain one sample per frame from one converter. Block mapping rules are very similar to the frame mapping rules from the JESD204B standard. Mapping of octets into 64-bit blocks is done in order with D0 representing the first octet of a frame. For example, if F = 8, D0 represents the first octet of the JESD204C frame, and D7 represents the last octet of the JESD204C frame. The first octet of the frame is the octet whose MSB is the MSB of Sample0 of Converter0 (same as in JESD204B). For example, if F = 2, D0 and D1 represent the first frame, D2 and D3 represent the second frame, etc.

To remain consistent with the approach used in JESD204B, the octets within the multiblock are shifted into the scrambler/descrambler in the order of MSB to LSB.

For cases where E is 1, every multiblock starts on a frame boundary. If E > 1, the extended multiblock will (and must!) start on a frame boundary. This is covered more in the Multiblock (MB) and Extended Multiblock (EMB) section.

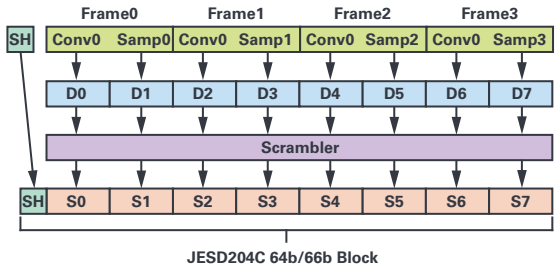

Figure 1. 64b/66b block format example for LMFS = 1.1.2.1, N = N' = 16.

The sync header is a 2-bit unscrambled value at the beginning of each block, the contents of which are interpreted to decode a single sync transition bit. These bits must be either a 0-1 sequence to indicate a logic 1 or a 1-0 sequence to indicate a logic 0. Table 1 enumerates the sync header sync transition bit values.

Table 1. Sync Header Values

Sync Header (0.1)	Sync Transition Bit
00	Invalid
01	1
10	0
11	Invalid

The 64b/80b block format is illustrated in Figure 2. In addition to the eight octets of sample data and the two sync header bits, there are two fill bits placed in-between each of the octets. The values of the fill bits are determined by a 17-bit PRBS sequence to reduce spurs and ensure an appropriate number of data transitions to maintain dc balance. The unscrambled fill bits are inserted into the block after the sample data has been scrambled.

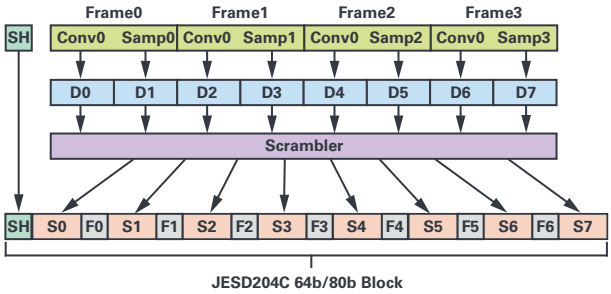

Figure 2. 64b/80b block format example for LMFS = 1.1.2.1, N = N′ = 16.

The 64b/80b option was provided to retain the same clock ratios as 8b/10b, which helps simplify the phase-locked loop (PLL) design while also minimizing spurs. This scheme would be preferred over 8b/10b in applications that want to use forward error correction or take advantage of the other functions provided by the sync word which will be discussed shortly.

Multiblocks (MB) and Extended Multiblocks (EMB)

There are 32 blocks in a JESD204C multiblock. The 32 sync transition bits in each multiblock make up a 32-bit sync word. These will be discussed in more detail later. An extended multiblock is a container of E multiblocks and must contain an integer number of frames. E > 1 is required when a multiframe does not contain an integer number of frames. The format of the multiblock and extended multiblock are illustrated in Figure 3.

A multiblock is either 2112 (32×66) or 2560 (32×80) bits depending on which 64-bit encoding scheme is used. For most implementations and configurations, an extended multiblock will be just one multiblock. The E parameter is introduced in JESD204C and determines the number of multiblocks in the extended multiblock. The default value for E is 1. As implied above, E > 1 is required for configurations where the number of octets in the frame, F, is not a power of two. The equation for E is: E = LCM(F, 256)/256. These configurations are typically preferred when transmitting 12-bit samples and N′ is set to 12 to maximize bandwidth efficiency in the link. This requirement ensures that the EMB boundary coincides with a frame boundary.

Figure 4 and Figure 5 illustrate an example of a JESD204C configuration where E > 1. The JESD204C configuration shown is for the case where LMFS = 2.8.6.1, N′ = 12, and E = 3. Figure 4 shows the transport layer mapping. In this configuration, there are four 12-bit samples per lane which translates to six octets. Since each block of a multiblock needs eight octets, the block is filled in with two octets (1.33 samples) from the subsequent frame.

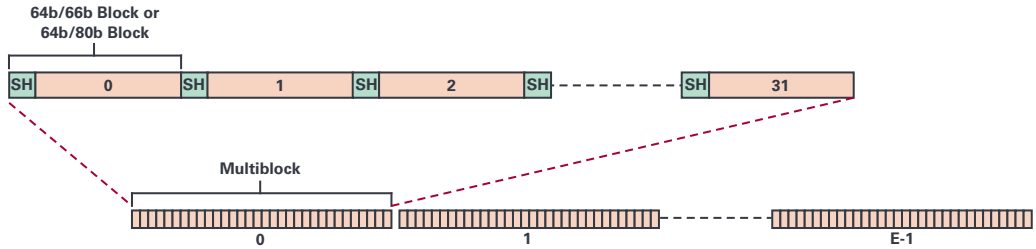

Figure 3. JESD204C multiblock and extended multiblock format.

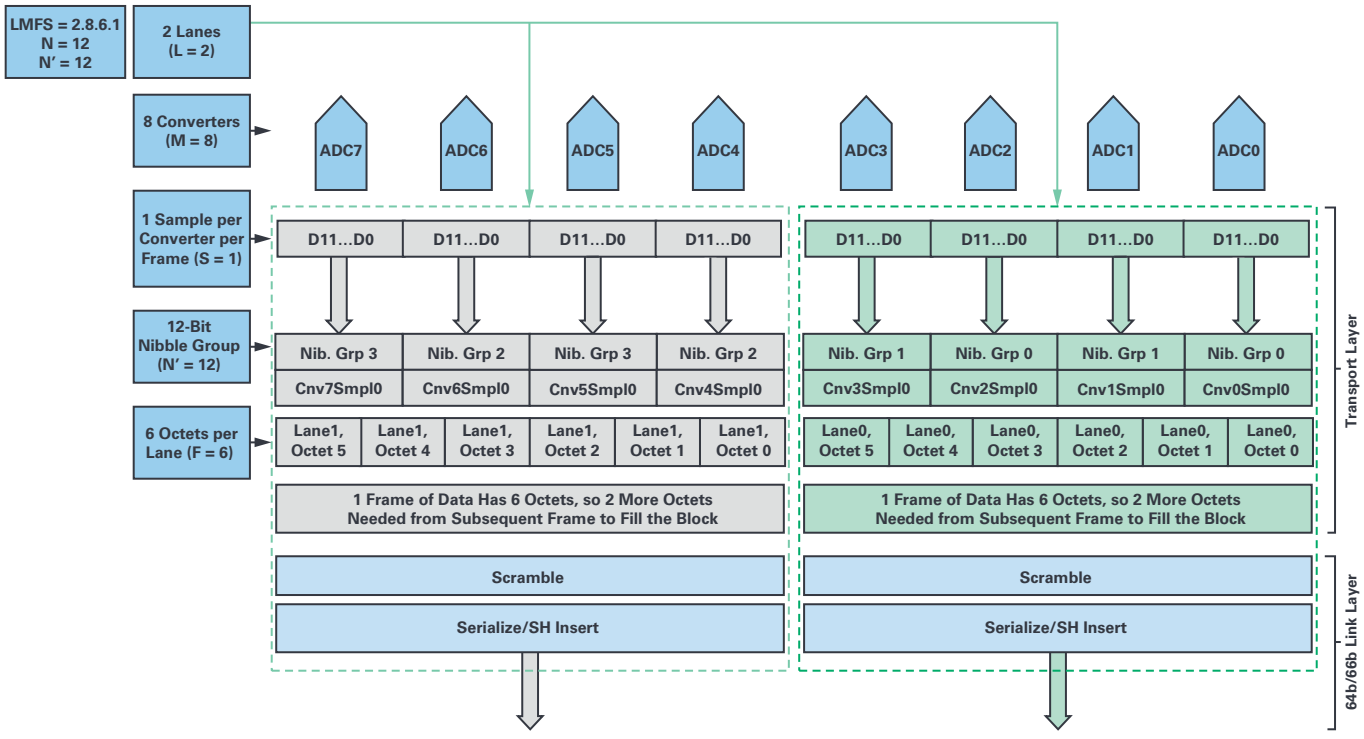

Figure 4. Transport layer mapping for LMFS = 2.8.6.1, N′ = 12, E = 3.

Figure 5 shows how the blocks and multiblocks are formed using the frames of data from the transport layer. As illustrated, you can see that frame boundaries align with block boundaries on every third block. Since multiblocks consist of 32 blocks, frame alignment to a multiblock is not achieved until after the third multiblock. Therefore, E = 3.

LEMC is the extended multiblock counter and is roughly equivalent to the LMFC in the 8b/10b link layer. SYSREF aligns all LEMCs in a system and the LEMC boundaries are used to determine synchronization and lane alignment.

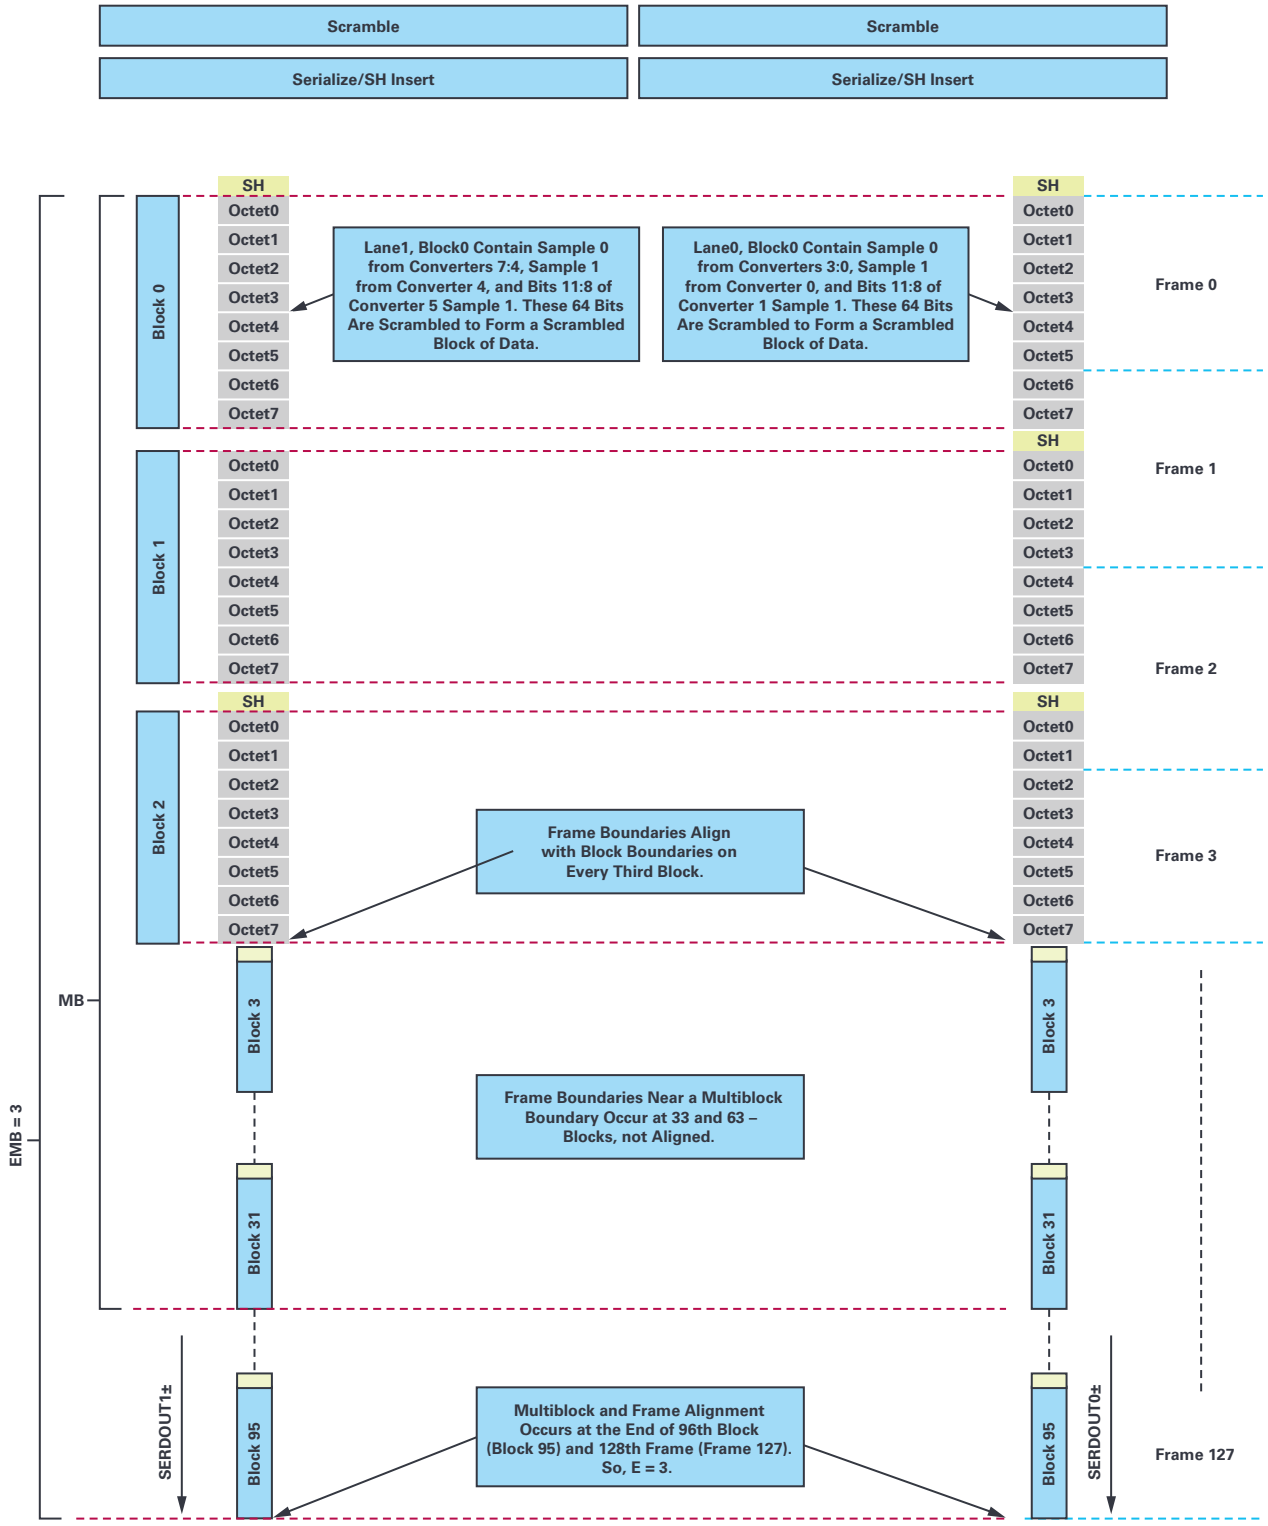

Figure 5. Serializer output multiblock/frame alignment for LMFS = 2.8.6.1, N′ = 12, E = 3.

Sync Word

The 32-bit sync word is made from each of the sample headers from the 32 blocks within the multiblock where bit 0 is transmitted first. The sync word is used to provide lane synchronization and enable deterministic latency. In addition, it can optionally provide for CRC error checking, forward error correction, or provide a command channel for the transmitter to communicate to the receiver.

There are three different formatting options for the 32-bit sync word. In each case, the end-of-multiblock sequence is required since it is used to acquire multiblock synchronization and lane alignment. Table 2 and Table 3 show the different bit fields available in the two most common use cases.

Table 2. CRC-12 Sync Word

Sync Word Bit	Field Name	Function
0	CRC11	Bits 11:9 of the 12-bit CRC value—applicable to the previous multiblock
1	CRC10	
2	CRC9	
3	1	Always 1
4	CRC8	Bits 8:6 of the 12-bit CRC value—applicable to the previous multiblock
5	CRC7	
6	CRC6	
7	1	Always 1
8	CRC5	Bits 5:3 of the 12-bit CRC value—applicable to the previous multiblock
9	CRC4	
10	CRC3	
11	1	Always 1
12	CRC2	Bits 2:0 of the 12-bit CRC value—applicable to the previous multiblock
13	CRC1	
14	CRC0	
15	1	Always 1
16	Cmd6	Bits 6:4 of the 7-bit command channel
17	Cmd5	
18	Cmd4	
19	1	Always 1
20	Cmd3	Bit 3 of the 7-bit command channel
21	1	Always 1
22	EoEMB	End-of-extended multiblock bit
23	1	Always 1
24	Cmd2	Bits 2:0 of the 7-bit command channel
25	Cmd1	
26	Cmd0	
27	0	End-of-multiblock pilot signal
28	0	
29	0	
30	0	
31	1	

Table 3. FEC Sync Word

Sync Word Bit	Field Name	Function
0	FEC[25]	Bits 25:4 of the 26-bit forward error correction word—applicable to the previous multiblock
1	FEC[24]	
2	FEC[23]	
3	FEC[22]	
4	FEC[21]	
5	FEC[20]	
6	FEC[19]	
7	FEC[18]	
8	FEC[17]	
9	FEC[16]	
10	FEC[15]	
11	FEC[14]	
12	FEC[13]	
13	FEC[12]	
14	FEC[11]	
15	FEC[10]	
16	FEC[9]	
17	FEC[8]	
18	FEC[7]	
19	FEC[6]	
20	FEC[5]	
21	FEC[4]	
22	EoEMB	End-of-extended multiblock bit
23	FEC[3]	Bits 3:0 of the 26-bit forward error correction word—applicable to the previous multiblock
24	FEC[2]	
25	FEC[1]	
26	FEC[0]	
27	0	End-of-multiblock pilot signal
28	0	
29	0	
30	0	
31	1	

64b/66b Link Operation

The link establishment process when using the 64b/66b link layer starts with sync header alignment, then progresses to extended multiblock synchronization, and finally to extended multiblock alignment.

Sync Header Alignment

The sync transition bit in the sync header ensures that there is a data transition at every block boundary (66 bits). A state machine in the JESD204C receiver detects a data transition and then looks for another transition 66 bits later. If the state machine detects bit transitions at 66-bit intervals for 64 consecutive blocks, sync header lock (SH_lock) is achieved. The machine is restarted if 64 consecutive transitions are not detected.

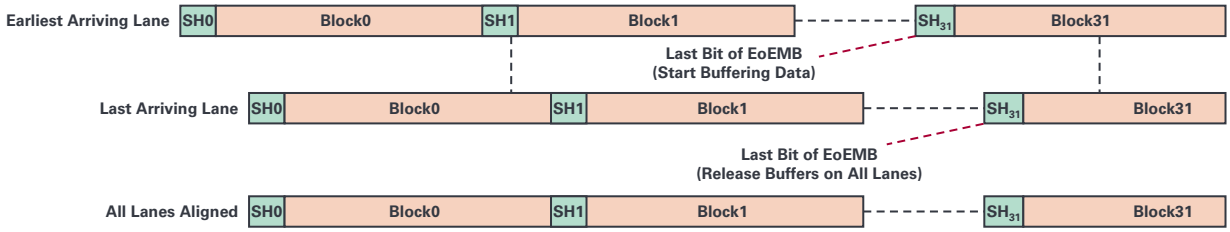

Figure 6. JESD204C extended multiblock (lane) alignment.

Extended Multiblock Sync

Once sync header alignment is achieved, the receiver looks for the end-of-extended-multiblock (EoEMB) sequence (100001) in the transition bits. The structure of the sync word ensures that this sequence can only happen at the appropriate time. Once an EoEMB is identified, the state machine examines every 32nd sync word to ensure the end-of-multiblock pilot signal (00001) is present. If E = 1, the EoEMB bit will also be present with the pilot signal. If E > 1, then every E × 32 transition bits, the pilot signal will have included the EoEMB bit. Once four consecutive valid sequences are detected, then end-of-extended-multiblock lock (EMB_LOCK) is achieved. Monitoring of every E × 32 transition bit continues and EMB_LOCK is lost if a valid sequence is not detected and the alignment process is reset.

Extended Multiblock (Lane) Alignment

Lane alignment when using the 64b/66b link layer is very similar to when using the 8b/10b link layer in that an elastic buffer is employed in the JESD204C receiver on each lane to store incoming data. This is called extended multiblock alignment and the buffers begin storing data at the EoEMB boundary (rather than the /K/ to /R/ boundary during ILAS when using the 8b/10b link layer). Figure 6 illustrates how lane alignment is achieved. Each lane's receive buffer starts buffering data once the last bit of the EoEMB is received except for the last arriving lane. When the last arriving lane's EoEMB is received, it triggers the release of all the lane's receive buffers so that all lanes are now aligned.

Error Monitoring and Forward Error Correction

JESD204C sync word options give the user the ability to either monitor or correct errors that may occur in the JESD204 data transmission. The trade-off associated with error correction is additional latency in the system. For most applications, error monitoring using the CRC-12 sync word is suitable since it provides for a bit-error rate (BER) of greater than 1×10^{-15}.

The CRC-12 encoder in the JESD204C transmitter takes in the scrambled data bits of each multiblock and computes 12 parity bits. These parity bits are transmitted to the receiver during the subsequent multiblock. The receiver will likewise compute 12 parity bits from each multiblock of data it receives, and these bits are compared to those that were received in the sync word. If all the parity bits do not match, there is at least one error in the received data and an error flag can be raised.

For error-sensitive applications that are not sensitive to added latency (like test and measurement equipment), using FEC can result in a BER of better than 10×10^{-24}. The FEC circuit in the JESD204C transmitter computes the FEC parity bits of the scrambled data bits in a multiblock and encodes these parity bits on the sync header stream of the next multiblock. The receiver calculates the syndrome of the received bits, which is the difference between the locally generated and the received parity. If the syndrome is zero, the received data bits are assumed to be correct. If the syndrome is non-zero, it can be used to determine the most likely error.

The FEC parity bits are calculated in a similar way to CRC. The FEC encoder takes in the 2048 scrambled data bits of the multiblock and adds 26 parity bits to construct a shortened binary cyclic code. The generator polynomial for this code is:

$$g(x) = (x^{17} + 1)(x^9 + x^4 + 1) = x^{26} + x^{21} + x^{17} + x^9 + x^4 + 1$$

This polynomial can correct up to a 9-bit burst error per multiblock.

Final Thoughts

To meet the faster data processing demands for data-intensive applications over the course of the next several years, JESD204C defines the multigigabit interface as a required communications channel between data converters and logic devices. The up-to 32 GSPS lane rates and 64b/66b encoding enables ultrahigh bandwidth applications with minimal overhead to improve system efficiency. 5G communications, radar, and electronic warfare applications all stand to benefit from these and other improvements in the standard. With the addition of error correction, cutting-edge instrumentation and other applications can count on operating without error for many years.

For more information on JESD204 and its implementation in Analog Devices products, please visit ADI's JESD204 serial interface page. For more information on ADI high speed converters, visit our RF converters page and 28 nm RF data converter page. For more information on ADI's transceiver products, visit our RadioVerse™ page.

About the Author

Del Jones is an applications engineer for the High Speed Converters Team in Greensboro, North Carolina. He has worked for ADI since 2000, supporting ADCs, DACs, and serial interfaces. Prior to ADI, he worked as a board and FPGA design engineer in the telecommunications industry. Del earned his bachelor's degree in electrical engineering from the University of Texas at Dallas. He can be reached at *del.jones@analog.com*.

Can You Really Get ppm Accuracies from Op Amps?

Barry Harvey

Industrial and medical design continually push to improve product accuracy and speed. The analog integrated circuit industry has generally kept up with speed requirements, but it is falling behind on accuracy demands. There is a march toward 1 ppm accurate systems, especially now that 1 ppm linear ADCs are becoming common. This article presents op amp accuracy limitations and how to choose the few op amps that have a chance of 1 ppm accuracy. We will also discuss a few application improvements to existing op amp limitations.

Accuracy is about numbers: how closely a system works to intended numerical value. Precision is about the depth of the numerical value in terms of digits. In this article, we will use accuracy as a term that includes all limitations to system measurements, such as noise, offset, gain error, and nonlinearity. Many op amps have some error terms at ppm levels, but none have all the errors at the ppm level. For instance, chopper amplifiers can provide ppm-level offset voltages, dc linearity, and low frequency noise, but they have problematic input bias currents and linearity at frequency. Bipolar amplifiers can provide low wideband noise and good linearity, but their input currents can still cause in-circuit errors (we will hence use the term application for in-circuit). MOS amplifiers have excellent bias currents but are generally deficient in the low frequency noise and linearity areas.

In this article we will use the rough equivalency of 1 ppm nonlinearity in the transfer function as –120 dBc distortion in harmonic distortion.

Non-ppm Amplifier Types

Let's discuss the types of amplifiers we reject as not highly linear. The least linearity is found in so-called video or line driver amplifiers. These are wideband amplifiers with terrible dc accuracies: offsets in the several millivolts and bias currents in the 1 µA to 50 µA range, and usually with poor 1/f noise. Expected accuracies are 0.3% to 0.1% at dc, although the ac distortion can be from –55 dBc to –90 dBc (2000 ppm to 30 ppm linearity).

The next category is older classic op amp designs, such as OP-07, that may have high gain, common-mode rejection ratio (CMRR), and PSRR, and good offsets and noise, but that cannot achieve better than –100 dBc distortion, especially into a 1 kΩ or heavier load.

Then there are the cheap amplifiers, new or old, that cannot best –100 dBc when loaded more heavily than 10 kΩ.

There is the audio amplifier class of op amps. They are fairly cheap, and their distortions can be very good. However, they are not designed for and do not offer good offsets nor good 1/f noise. They also cannot deliver distortion beyond perhaps 10 kHz.

There are op amps meant to support MHz signals linearly. These are usually bipolar throughout and have large input bias currents and 1/f noise. This application space sees more like –80 dBc to –100 dBc performance, and ppm performance is not practical with these op amps.

Current feedback amplifiers also cannot support deep linearity nor even modest accuracy, no matter how wideband nor huge their slew rates may be. Their input stage has a mess of error sources, and they do not have much gain nor input nor supply rejections. Current feedback amplifiers also have a thermal drift that extends fine settling times greatly.

Then we have the modern general-purpose amplifiers. They typically have a 1 mV offset and microvolts of 1/f noise. They support –100 dBc distortion, but usually not when heavily loaded.

Op Amp Error Sources

Figure 1 shows a simplified op amp block diagram with ac and dc error sources added. The topology is a single-pole amplifier with an input g_m that drives a gain node that is buffered as the output. While there are many op amp topologies, the error sources shown apply to them all.

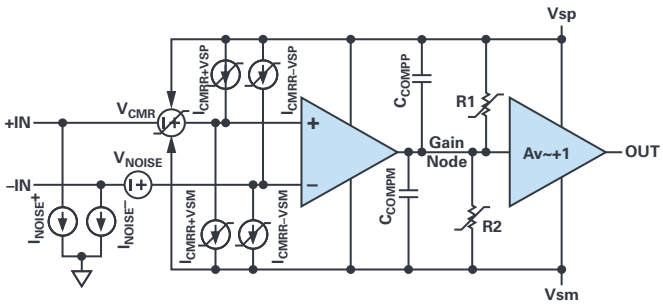

Figure 1. Simplified op amp and sources of error.

Input Noises

We have an input noise voltage V_{NOISE} with wideband and 1/f spectral content. You can't measure a signal accurately if the noise is of similar magnitude or more than a system LSB. For example, if we had a 6 nV/√Hz wideband noise and a 100 kHz system bandwidth, we would have 1.9 µV rms noise at the input. We could filter this noise down: for instance, dropping bandwidth to 1 kHz drops the noise to 0.19 µV rms, or about 1 µV p-p (peak-to-peak). Low-pass filtering in the frequency domain drops noise magnitude, as would averaging the output of an ADC over time.

However, 1/f noise cannot be practically filtered or averaged away because it is so slow. 1/f noise is usually characterized by peak-to-peak voltage noise generated in the 0.1 Hz to 10 Hz spectrum. Most op amps have between 1 µV p-p and 6 µV p-p low frequency noise and are thus not suitable for dc-accurate ppm levels, especially if providing gain.

Figure 2 shows the current and voltage noise of a good high accuracy amplifier, the LT1468.

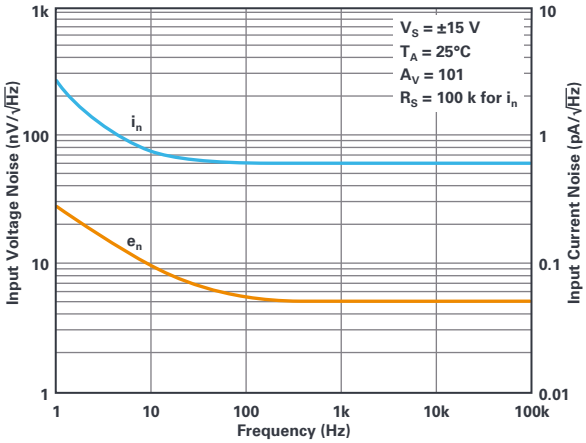

Figure 2. LT1468 input voltage and current noise.

At the inputs in Figure 1, we also have bias current noise sources I_{NOISE+} and I_{NOISE-}. They contain both wideband and 1/f spectral content. I_{NOISE} multiplies against application resistors to become more input voltage noise. Generally, the two current noises are uncorrelated and do not cancel with equal input resistors but add in rms fashion. Quite often I_{NOISE} times application resistors exceeds V_{NOISE} in the 1/f region.

Input Common-Mode Rejection and Offset Errors

The next error source is V_{CMRR}. This embodies the common-mode rejection ratio specification where an offset voltage changes in response to the input's level relative to both supply rails (the so-called common-mode voltage, V_{CM}). The symbol used indicates supply interaction at the arrows, and the segmented line through it suggests it's variable but might not be linear. The major effect of CMRR on signals is that the linear part is indistinguishable from a gain error. The nonlinear part will be a distortion. Figure 3 shows the CMRR of an LT6018. The added line intersects extreme points of the CMRR curve just before the curve diverges into overload. The slope of the line gives a CMRR = 133 dB. The CMRR curve diverges from a perfect line by only about 0.5 µV per 30 V span—a very successfully sub-ppm input. Other amplifiers can have much more curvature.

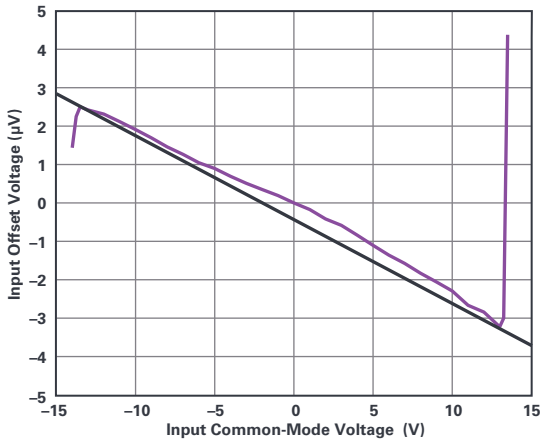

Figure 3. LT6018 input offset voltage vs. V_{CM}.

Offset voltage (VOS) will be lumped into CMRR here. Chopper amplifiers have sub-10 µV input offsets, and that's close to a single-ppm error, relative to typical signals of 2 V p-p to 10 V p-p. Even the best ADCs generally have as much as 100 µV offset. So, the onus of offset is not so much on the op amps; the system will have to auto-zero itself, anyway. Associated with the input signal's common-mode level is I_{CMRR}, which is the input bias current and its variation with supplies. The broken lines suggest that the bias currents are variable with voltage and also may not be linear. There are four I_{CMRR}s because both inputs can have independent bias currents and level dependencies, and because each input is varied by both supplies independently. The circuit effect of the I_{CMRR}s (which sum to form bias current) is to multiply against application circuit resistances to add to overall circuit offset. Figure 4 shows the bias currents of an LT1468 vs. V_{CM} (the I_{CMR} specification). The slope as shown by the added line is ~8 nA/V, which would be 8 µV/V with a 1 kΩ applications resistor, or a low ppm error. The deviation from straight line is about 15 nA, which in a 1 kΩ application environment creates 15 µV error over a 26 V span, or a 0.6 ppm nonlinearity.

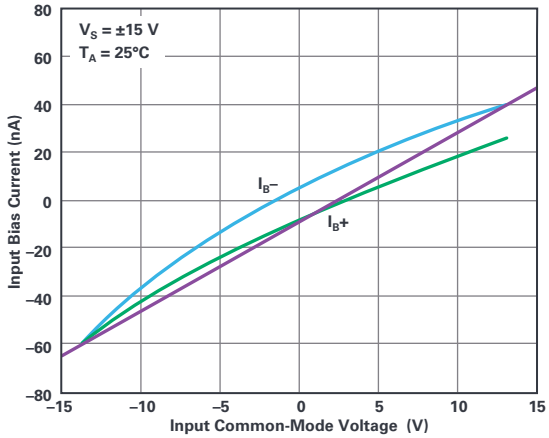

Figure 4. LT1468 input bias current vs. V_{CM}.

Input Stage Distortion

Figure 1 shows the input stage, which is generally a transconductor made from a differential pair of transistors. The top of Figure 5 shows the collector, or drain currents, of various differential amplifier types vs. differential input voltage. We simulate a simple bipolar pair, a translinear circuit that we will call *clever bipolar*, a subthreshold (that is, very large) MOS differential pair, a bipolar pair with emitter resistors (degenerated in Figure 5), and a smaller MOS pair operating out of the subthreshold region and into its square-law regime. All differential amplifiers are simulated with a 100 µA tail current.

Not a lot of information is obvious until we display transconductance vs. V_{IN}, as shown at the bottom of Figure 5. Transconductance (g_m) is the derivative of output current with respect to the input voltage, as generated using the LTspice® simulator. The syntax has d() to be mathematically equal to d()/d(VINP). The non-flatness of g_m is the basic distortion mechanism of op amps at frequency.

At dc, the open-loop voltage gain of the op amp is ~g_m(R1||R2), assuming the output buffer gain is about unity. R1 and R2 represent the output impedances of various transistors in the signal path, each connected to a supply rail or other. This is the basis of limited gain in an op amp. R1 and R2 are not guaranteed to be linear; they are a cause of unloaded distortion or nonlinearity. Aside from linearity, we need gains approaching or exceeding one million for ppm gain accuracies.

Observing the standard bipolar curve, we see it has the greatest transconductance of the group, but that transconductance fades quickly as the input moves from zero volts. This is concerning—a basic requirement for linearity is constant gain or g_m. On the other hand, who cares that the amplifier voltage gain is so high that the differential input would only move microvolts as the output moves volts? Time to introduce C_{COMP}.

C_{COMP} (the parallel of C_{COMPP} and C_{COMPM}) absorbs most of the g_m's output current over frequency. It sets the gain bandwidth product (GBW) of the

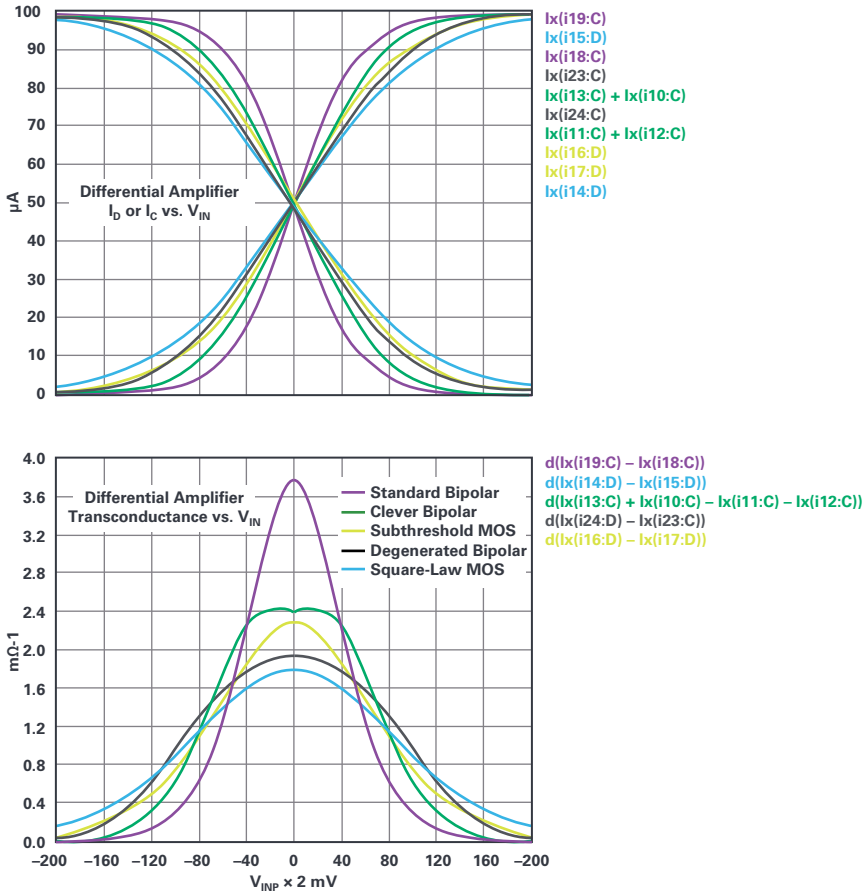

Figure 5. Various differential amplifiers' output currents and transconductances vs. input voltage.

amplifier. GBW establishes that, at a frequency f, the amplifier will have an open-loop gain of GBW/f. If the amplifier is outputting 1 V p-p at f = GBW/10 with a closed-loop gain of +1, then we have 100 mV p-p between the inputs. That's ±50 mV from balance. Note that the standard bipolar curve shown in Figure 5 has lost about half its gain at ±50 mV, guaranteeing massive distortion. However, the clever bipolar only lost 13% of its gain, the subthreshold MOS lost 26%, the degenerated bipolar lost 12%, and the square-law MOS lost 15%.

Figure 6 shows the distortion vs. amplitude for the input stage. This will appear (times the noise gain) at the output of the application circuit. You may get more output distortion than this, but not less.

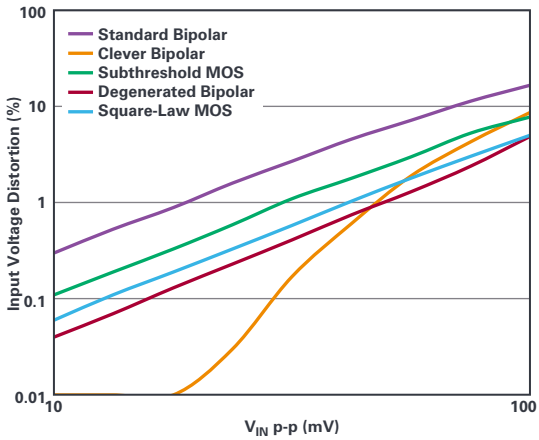

Figure 6. Total harmonic distortion of the input stage vs. differential input voltage.

Excluding the clever bipolar stage, the differential amps show that the distortion is proportional to the square of the input. In a unity-gain application, the output distortion contribution is equal to the input distortion. This is the dominant distortion source for most op amps.

Consider a unity-gain buffer with a bipolar input. For an output of V_{OUT} peak-to-peak volts, the input differential signal would be

$$f_{SIGNAL} \times \frac{V_{OUT}}{GBW} \quad (1)$$

We estimate that

$$\%Distortion_{INPUT} = 0.3\% \times \left(\frac{V_{IN}\ p\text{-}p}{10\ mV}\right)^2 \quad (2)$$

And

$$\%Distortion_{OUTPUT} = G_{NOISE} \times \%Distortion_{INPUT} = 0.3\% \times \left(\frac{V_{OUT,\ p\text{-}p} \times f_{SIGNAL}}{10\ mV \times GBW}\right)^2 \quad (3)$$

where G_{NOISE} is the noise gain of the application.

A 1 ppm nonlinearity is like –120 dBc harmonic distortion, and that's 0.0001%. Given an amplifier with a bipolar input stage, a 15 MHz GBW, and outputting 5 V p-p as a buffer, Equation 2 tells us that the maximum frequency for that linearity is just 548 Hz. This assumes the amplifier is at

least that linear at lower frequencies. Of course, when the amplifier provides gain, the noise gain increases and the –120 dBc frequency drops.

The subthreshold MOS input stage supports –120 dBc up to 866 Hz, square-law MOS up to 1342 Hz, and degenerated bipolar up to 1500 Hz. Clever bipolar does not follow the distortion prediction and one must get estimates from the data sheet.

We can use the simpler formula

$$\%\text{Distortion}_{\text{OUTPUT}} = K \times G_{\text{NOISE}} \times \left(\frac{V_{\text{OUT, p-p}} \times f_{\text{SIGNAL}}}{\text{GBW}}\right)^2 \quad (4)$$

where K is found from the distortion curves of an op amp's data sheet.

As a side-note, there are many op amps with rail-to-rail input stages. Most get this ability from two separate input stages that have a hand-off from one to the other over the input common-mode range. This hand-off generates changes in offset voltage, and potentially bias current, noise, and even bandwidth. It also essentially causes a switching transient at the output. These amplifiers cannot be used for low distortion if the signal ever traverses the crossover region. An inverting application may work, however.

We haven't discussed slew-enhanced amplifiers yet. These designs do not run out of current with large differential inputs. Unfortunately, small differential inputs still cause variations in g_m of similar magnitude to the inputs discussed, and low distortion still demands a large loop gain at frequency.

Since we are looking for ppm-level distortion, we will not operate the amplifier anywhere near its slew rate limit, so, oddly enough slew rate is not an important parameter for ppm linearity at frequency, just GBW.

We've discussed open-loop gain as modeled by a single-pole compensation design. Not all op amps are compensated that way. Generally, open-loop gain is taken from the data sheet curve, and GBW/($G_{\text{NOISE}} \times f_{\text{SIGNAL}}$) in the equation is that open-loop gain at frequency.

Gain Node Errors

The next items in Figure 1 to discuss are R1 and R2. These resistors, along with the input g_m, give the amplifier its open-loop dc gain of $g_m \times (R1||R2)$. These resistors have been drawn with the variable and nonlinear strikethrough in the schematic. Nonlinearities of these resistors embody the amplifier's unloaded distortions. Further, R1 injects influence from the positive supply such that the dc positive power supply rejection ratio (PSRR+) and is approximately equal to $g_m \times R1$. Similarly, R2 is responsible for PSRR–. Note how PSRR is almost equivalent to open-loop gain in magnitude. C_{COMPP} and C_{COMPM} have an analogous injection of supply signals to R1 and R2; they set PSRR+ and PSRR– over frequency.

It is possible that an amplifier with modest gain (<<10^6) can be quite linear, but that modest gain will limit gain accuracy.

The power supply terminals can be a source of distortion. When the output stage drives a heavy load, that load current flows from one of the supplies. At frequency, the distant source of the supply may have little remote regulation, such that the op amp's bypass capacitors are the real supply source. The supply current drops across the bypass capacitors. These drops are dependent on the ESR, ESL, and the reactance, and they cause a supply disturbance. Since the output is class AB, only half of the output current waveform modulates the supply, creating even harmonic distortion. PSRR over frequency attenuates the supply disturbance. As an example, if we observed 50 mV p-p supply disturbance and wish the PSRR-induced input disturbance to be less than 5 µV p-p, we need a PSRR of 80 dB at signal frequency. Estimating that PSRR(f)~Avol(f), an amplifier with a 15 MHz GBW would have enough PSRR at frequencies below 1500 Hz.

Output Stage Distortions

The last item in Figure 1 is the output stage, which is considered a buffer for this discussion. A typical output stage transfer function is shown in Figure 7.

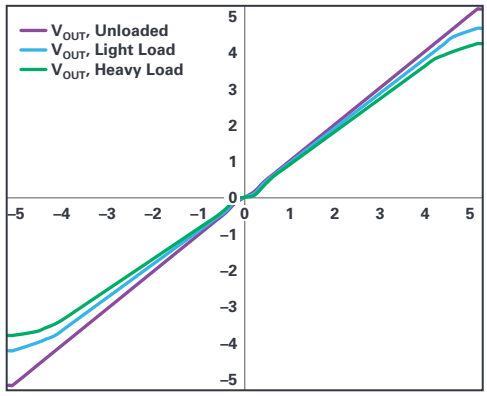

Figure 7. Transfer function of an output buffer with different loads.

For the different loads, we see four kinds of error. The first is clipping: although this hypothetical output stage has a gain of 1 nominally, it's not quite a rail-to-rail output stage. Even the unloaded output clips, in this case, 100 mV from each supply rail. The output clips at successively lower voltages as the load is increased (decreased load resistance). Obviously, clipping is a disaster for distortion and the output swing must be reduced to avoid it.

The next error is gain compression, which we see as curvature in the transfer function at the extremes of signal. The compression happens at earlier voltages as the load is increased, and like clipping no ppm-level distortion is generally possible in this regime. This compression is generally due to a small output stage that is struggling to output required current. A good rule of thumb is that the linear, uncompressed maximum output current available from the amplifier is only about 35% of output short-circuit current.

Another glaring source of distortion is the crossover region around $V_{IN} = 0$. Unloaded, the crossover kink may not be apparent, but with increasing loading we get something like the exaggerated kink of the green curve. Eliminating crossover distortion generally demands robust supply current.

The last distortion is harder to perceive. Because there are some bits of amplifier circuitry that output positive voltages and currents, and other bits for negative signals, there is no guarantee that they have the same gain, especially when loaded. Figure 7 shows a lesser gain for negative signals when loaded.

All these distortions are reduced by loop gain. If the output stage had 3% distortion, we would need a loop gain of 30,000 to get to a –120 dBc level. That of course happens below a frequency of GBW/(30,000 × G_{NOISE}), generally in the 1 kHz regime for a 15 MHz amplifier.

Some output stages' distortions are frequency-dependent, but many are not. The open-loop gain suppresses the output stage distortion, but that gain falls with frequency. If the output distortion is constant with frequency, then the gain loss creates an output distortion that increases linearly with frequency. Meanwhile, input distortion causes a total output distortion that increases with frequency. In this case, the input distortion will probably dominate the total closed-loop output distortion, masking the output stage distortion contribution.

On the other hand, if the output stage distortion does vary, say, linearly with frequency, the falling loop gain creates another output distortion that varies as frequency squared, additional to and indistinguishable from input distortion.

Low power op amps generally have starved output stages whose quiescent currents are low. These amplifiers' output stages may well dominate output distortion, rather than the input stage. It is somewhat true that it takes at least 2 mA supply current to make a low distortion op amp.

Required Specifications for ppm-Level Accuracies

In practical level-shift, attenuate/gain, and active filter circuits, we have some basic op amp requirements for an amplifier supporting ±5 V signal while working in a 1 kΩ environment and achieving 1 ppm linearity shown in Table 1.

Table 1. List of Op Amp Errors and Magnitude Needed for ppm Accuracy

Characteristic	Magnitude	Comment
V_{NOISE}	<6 nV/√Hz	Wideband input voltage noise
V_{NOISE} 0.1 Hz to 10 Hz	<1 ppm, p-p re. full-scale	Low frequency input voltage noise
I_{NOISE}	<6 pA/√Hz	Wideband input current noise
I_{NOISE} 0.1 Hz to 10 Hz	<10 nA, p-p re. full-scale	Low frequency input current noise
V_{OS}	<200 µV	Input offset voltage; usually digitally corrected
CMRR	>100 dB	Input common-mode rejection ratio; <10 ppm gain error
CMRR Linearity	>120 dB	The curvature of CMRR
I_{BIAS}	<200 nA	Input bias current; times a 1 kΩ resistor, <V_{OS}
I_{BIAS} vs. V_{CM} (I_{CMR})	<10 nA/V	Times a 1 kΩ resistor, less than 10 ppm gain error
I_{BIAS} vs. V_{CM} Linearity	<1 nA to 5 nA	Times a 1 kΩ resistor, less than 1 ppm re. full-scale
PSRR	>90 dB over some BW	Power supply rejection ratio; <1.6 µV V_{OS} shift per 50 mV supply variation
GBW	>1000× signal BW	Gain-bandwidth product; generally necessary for low distortion
Linear Output Current	>15 mA	Generally necessary for low distortion (= 35% of output short-circuit current)
DC V_{OS} vs. V_{OUT} Linearity	<1 ppm nonlinearity	Without dc linearity, you can't achieve ac linearity

So, we see the limitations of op amps in the ppm accuracy world—can we do anything to improve them?

Noise: Obviously the first step is to select an op amp with input noise voltage no higher than an application resistors' combined noise. One could reduce the overall impedance of applications circuits to reduce their noises. Of course, as the applications' impedance drops, the signal currents through them increase and can increase load-induced distortion. In any case, there's no reason to reduce the output noise of an op amp stage much below the input noise of the stage it drives.

Current noise will multiply against application impedances as more voltage noise. MOS inputs are attractive in having very low current noise, but they often have more 1/f voltage noise than bipolar inputs. Bipolar inputs have pA/√Hz levels of current noise that make non-trivial applications noise, but the 1/f current content can produce applications voltage noise greater than the 1/f voltage noise of the amplifier. A general rule of thumb is that the application impedance should be less than V_{NOISE}/I_{NOISE} of the amplifier to avoid I_{BIAS}-dominated application noise. The lower the V_{NOISE} of a bipolar amplifier, the higher the I_{NOISE} will be.

Helping Op Amps Achieve Best Performance

Reducing Input Errors

Aside from selecting an op amp with superior CMRR, designers can use op amps in inverting circuits rather than noninverting. In inverting circuits, the inputs hug ground or some reference and do not induce CMRR errors at all. Not all applications circuits can be made inverting, and often no negative power supply is available for the negative signal excursions. Figure 8 shows two-pole Sallen-Key filters in noninverting and inverting realizations.

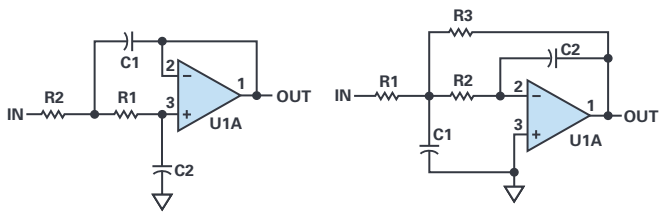

Figure 8. Noninverting (left) and inverting (right) Sallen-Key active filters.

I_{CMR} errors can be cancelled if both inputs have application resistors such that each input's bias current gets cancelled as an output error by the appropriate resistances. For instance, if an amplifier is set up at a gain of 10 with 900 Ω feedback and 100 Ω ground resistors, placing a series 90 Ω at the positive input will cancel perfectly equal bias currents at the output. The bias currents of most bipolar op amps are so well-matched that it is useful to choose 0.1% rather than common 1% resistors to achieve the best I_{CMR} rejection. In Figure 8, the compensating resistors would be placed in series with each –input. They should probably be bypassed across. The extra input resistor contributes more noise, unfortunately.

Inverting gain allows us to use an op amp that has rail-to-rail inputs without the signal traversing a switching point—assuming we bias the supplies and common-mode input level to avoid that switching voltage.

Supply Considerations

Output currents will modulate the local supply voltage. That supply signal will find its way to the input through PSRR. The induced input will produce an output signal that runs around its loop. At 1 kHz, a 1 µF local bypass capacitor has 159 Ω of impedance, much less than the impedance of the line between the supply source plus that of the source itself. Thus, the local bypass capacitor won't really be effective below 100 kHz. At 1 kHz, the remote supply controls regulation. At 1 kHz the amplifier may have, say, 90 dB supply rejection. Noting that most of the current from the op amp's supply terminals is made of even harmonics of the signal, we want the gain from the output to the offending supply to be lower than 30 dB to make our 120 dBc goal. The gain of 30 dB requires that the supply impedance should be <30× the load impedance. So, a 500 Ω load requires a supply with less than 17 Ω impedance. This is practical, but it does not allow a series isolating resistor or inductor between the supply and op amp. Things are tighter at 10 kHz; PSRR would drop from 90 dB to 70 dB, and the supply impedance would have to drop to 1.7 Ω. Doable, but tight. A large local bypass would help.

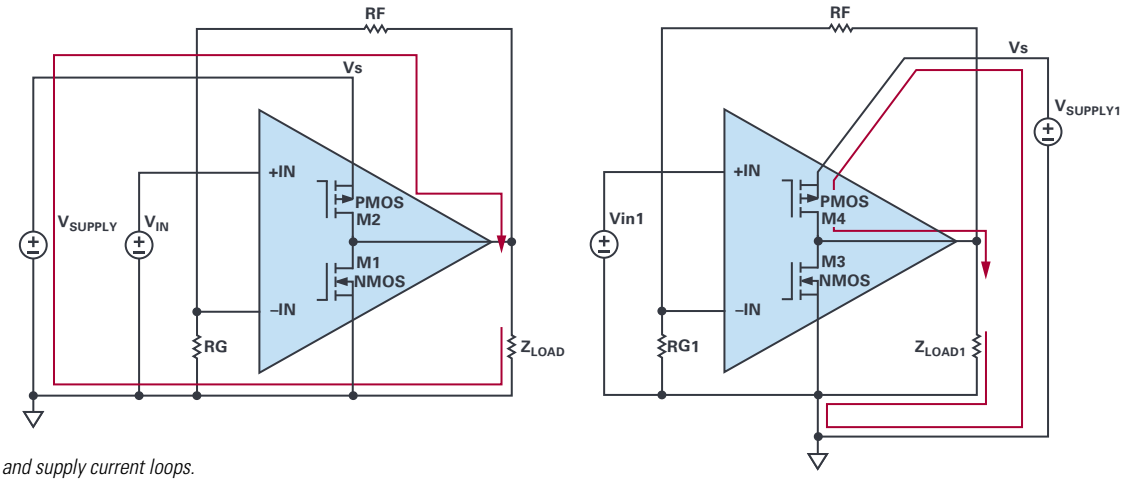

Figure 9. Load and supply current loops.

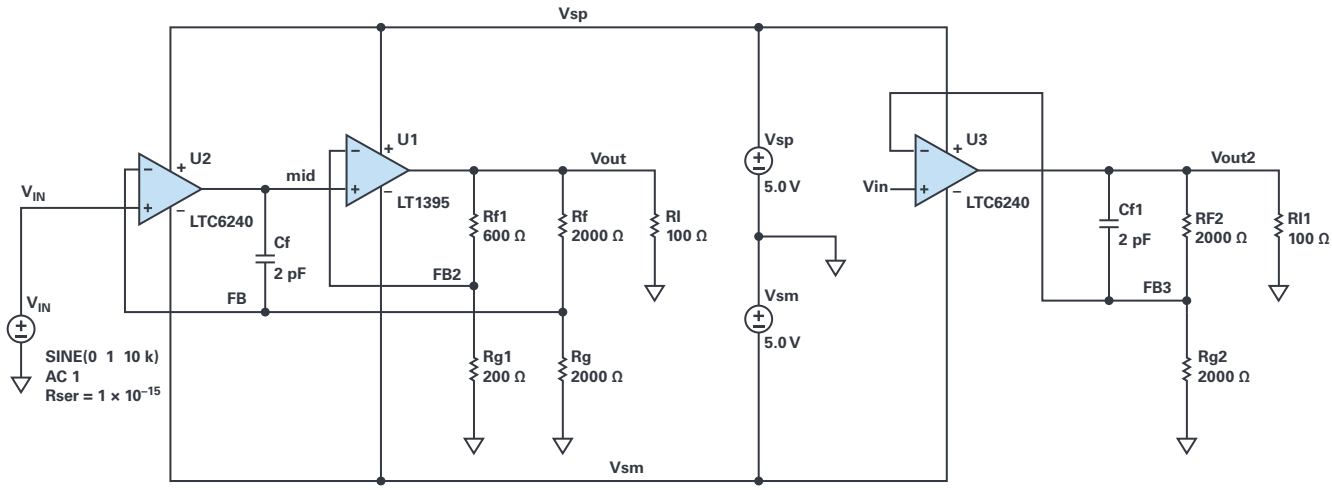

Figure 10. Composite amplifier vs. single amplifier distortion test.

From a layout point of view, it's important to see where the output current loops go, as seen in Figure 9.

The diagram on the left of Figure 9 shows positive supply current driven into a load, coming from a supply, then returning through ground back to the load. There can be voltage drops all along the ground path, such that even-harmonic supply current drops voltage from the signal source to the output, and from the feedback divider to output or input ground. This ground is not that ground. The right side of Figure 9 shows a better way to route the supply currents. The supply current is dressed away from input and feedback nodes.

At higher frequencies, above 100 kHz, the magnetic radiation of the supply lines can be a source of distortion. The even-harmonic currents of the supply can magnetically couple to the input or feedback network, dramatically increasing distortion with frequency. Careful layout is essential at these frequencies. Some amplifiers have nonstandard pinout; they keep the supply pins away from the inputs, and a few even offer an extra output terminal on the input side for to avoid magnetic interactions.

Reducing Load-Dominated Distortion

Many op amps' output stages become dominant distortion sources when loaded heavily. There are a couple of tricks to improve loaded distortion. One is the composite amplifier: one amplifier drives the output and another amplifier controls it, as seen in Figure 10.

This from an LTspice simulation. The LTC6240 and LT1395 have macromodels that include distortion playback. Most macromodels do not attempt to display distortion, and even if they do the simulated value can be inaccurate. This writer was able to look at the text of the macromodels, and yes, in these macromodels distortion was modeled fairly well.

On the right of Figure 10 is an LTC6240 providing a gain of 2 while driving 100 Ω—a difficult load for this amplifier. On the left of Figure 10 is a composite amplifier with another LTC6240 at the input, and a nice wideband current-feedback amplifier (CFA) to drive the same load as the standalone amplifier. The idea of the composite amplifier is that the output op amp already has a modestly low distortion, and that distortion can be reduced further by the input amplifier's loop gain over frequency. We have the same closed-loop gain of 2 for both standalone and composite amplifiers, but in the composite amplifier the LT1395 can be set up with its own gain (4 as set by Rf1 and Rg1) to reduce the output swing of the control amplifier. Since input-induced distortion increases as the square of output amplitude, there is a further reduction of distortion for the control op amp.

Figure 11 shows the spectrum of 10 kHz, 4 V p-p outputs.

Harmonic distortion would be measured as each harmonic level (dB) minus the level of the fundamental (at 10 kHz). As seen in the figure bottom the input signal has about –163 dBc distortion, good enough to trust the simulation. V(out2) is from the unassisted LTC6240 and has –78 dBc distortion. Not bad, but certainly not ppm-level.

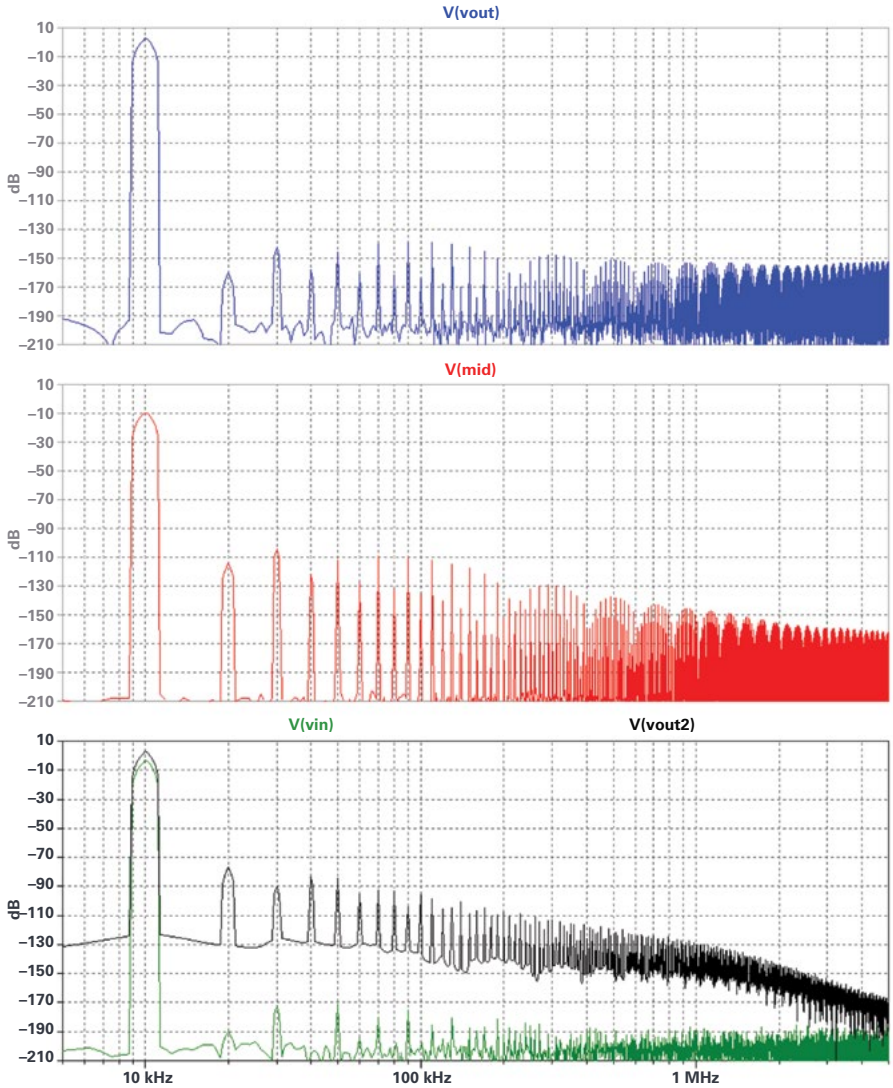

Figure 11. Composite and normal amplifier distortion spectra.

The top of Figure 11 shows the composite amplifier distortion at –135 dBc—pretty darn spectacular. Can we trust such a good result? For verification, the distortion at the schematic node mid is shown in the middle. If the output of a composite amplifier has near-zero distortion, but the output amplifier itself does have finite distortion, the feedback process will place the negative of the output amplifier's distortion at its input (mid). The distortion at mid is –92 dBc, and it actually matches the LT1395 data sheet curve! I would still wonder if the physical LTC6240 input CMRR or I_{CMR} curvature is expressed in the macromodel—they may yet increase real circuit distortion.

Unfortunately, few macromodels include distortion. You'd have to read the header in macromodel .cir files to see if it is supported. Some simulation is required to see if distortion matches the data sheet curves.

Compensation of a composite amplifier can be a little tricky, but in our example, we have a second amplifier whose bandwidth is over 10 times that of the input amplifier's, and just a little Cf compensates the circuit. In this compensation scheme, if the control amplifier has a bandwidth of BW at an overall gain, the output amplifier should have a bandwidth of >3 × BW, and the overall bandwidth will be conservatively set to ~BW/3.

To avoid losing bandwidth, we can use the boosted amplifier trick. It provides less distortion improvement than the composite approach but loses no bandwidth nor settling time. Figure 12 shows a test schematic.

The right side of Figure 12 shows U2, our standalone LTC6240, and the left side shows two LTC6240 amplifiers. U1 controls the output and has a gain of two, like the standalone; U2 has a gain of three. U2's output voltage at the boost node is larger than U1's, so U2 drives current into the output. R_{BOOST} and U2's gain are configured such that U2 drives 96% of the load current into RI, leaving a light load for U1 and thereby improving its distortion. One needs to assure that U2 has enough output headroom for its extra swing.

The LTC6240 has input-dominated distortion for loads in the kΩ range, but is output stage distortion dominated with our 100 Ω load.

Figure 13 shows the spectral results.

Again, we have –78 dBc distortion at 10 kHz for the standalone amplifier. The boosted amplifier delivers –106 dBc; not nearly as good as the composite amplifier, but almost 30 dBc better than the standalone. However, the boosted amplifier suffers only a modest bandwidth reduction.

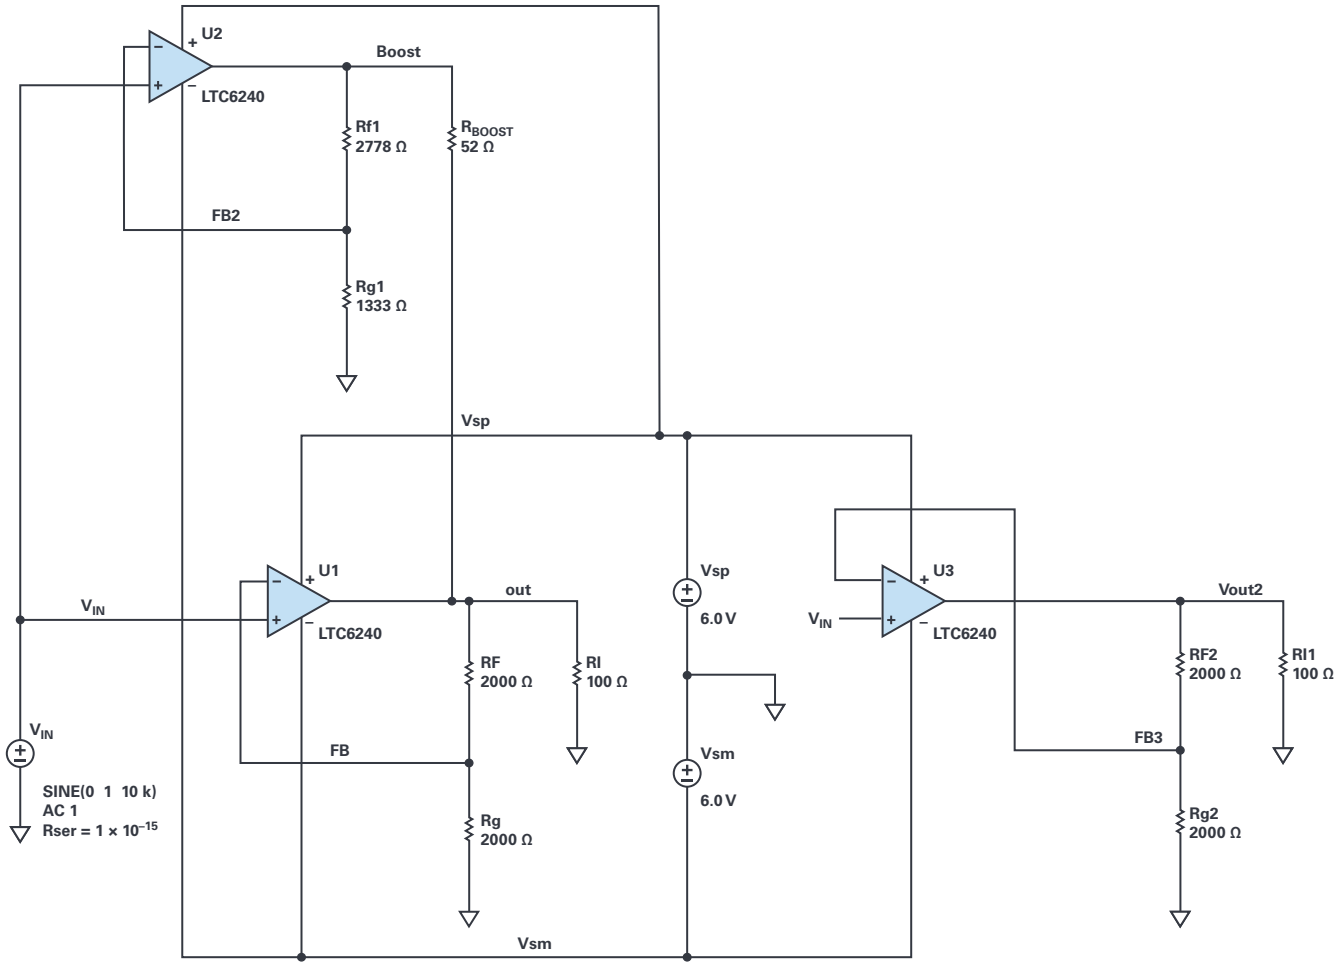

Figure 12. Boosted vs. standalone amplifier simulation setup.

Note that R_{BOOST} is tweaked; if we vary it as 52 ± 2 Ω, the boosted distortion degrades by 10 dBc, although little change thereafter happens for up to ±10 Ω. It would appear U1 likes to have some modest load of the expected polarity; ideal (no load) or excess boost current causes more distortion.

Ideally, U2 would have the same group delay as U1 so that the boost signal occurs at the same time as the output. U2 has 50% more gain than U1 and thus has less closed-loop bandwidth, suggesting the boost output lags the main output at frequency. U1's bandwidth could be reduced to be equal to U2's by installing a resistor across U1's inputs. This would increase the noise gain of U1 to be equal to U2 and achieve equality between the group delays. The simulator showed no improvement at 10 kHz; U1 gave best distortion with no delay balancing. Knowing whether or not this is true at higher frequencies will require trying it. If the amplifiers were of current-feedback type, Rf1 and Rg1 could be reduced to bring U2's bandwidth up to U1's.

Recommended ppm-Quality Amplifiers

Table 2 shows the salient specifications for some suggested amplifiers approaching ppm linearity.

Entries are in red to alert the reader that a parameter may violate ppm-level distortion. The easy to use winners of the group are the AD8597, ADA4807, ADA4898, LT1468, LT1678, and LT6018.

There are some amplifiers that have input problems that must be dealt with (noninverting applications might be an issue) but that can still provide good distortion: AD797, ADA4075, ADA4610, ADA4805, ADA4899, and LTC6228.

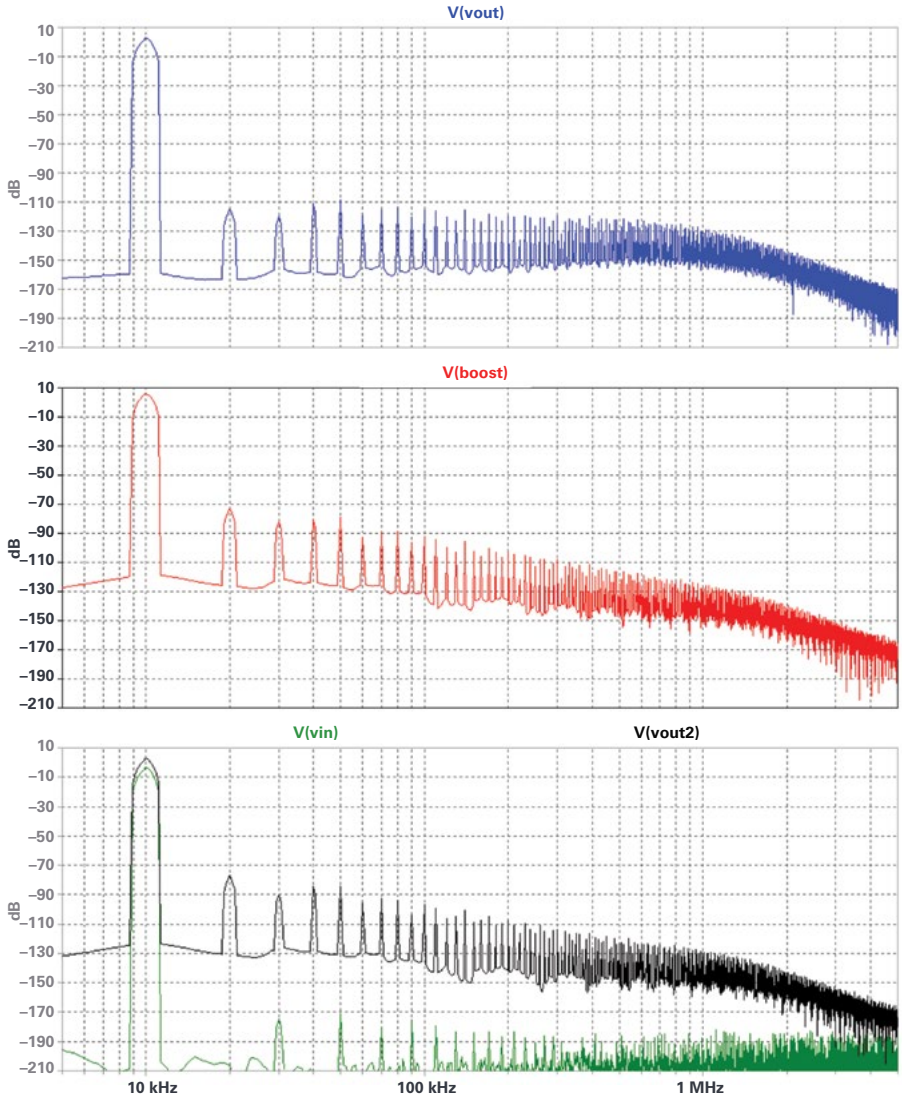

Figure 13. Boosted and normal amplifier distortion spectra.

Table 2. Salient Specifications for Some Suggested Amplifiers Approaching ppm Linearity

Device #	V_{NOISE} nV/√Hz	V_{NOISE} 0.1 Hz to 10 Hz µV p-p	I_{NOISE} pA/√Hz	I_{NOISE} 0.1 Hz to 10 Hz pA p-p	V_{OS} µV (max)	CMRR dB (min)	CMRR Nonlinearity µV/V	I_{BIAS} nA (max)
AD797	0.9	0.05	2	220	60 to 180	110 to 114		2000 or 3000
AD8597	1.1	0.08	2.4	190	120	105	0.1	200
ADA4075	2.8	0.06	1.2	60	1000	106	1.5	100 to 150
ADA4610	7.3	0.45	Very small		800 to 1800	96		0.025 to 1500 (hot)
ADA4805	5.2	0.1	0.7	140	125	103		800
ADA4807	3.1	1.6	0.7	370	125	103	0.2	800
ADA4898	0.9	0.05	2.4	130	125	103		400
ADA4899	1.0	0.4	5.2	4800	230	98		1000
LT1468	5	0.3	0.6	3	150 to 400	96		10, 40
LT1678	3.9	0.09	0.7	26	350	96		35
LT6018	1.2	0.03	0.75 or 3	110 or 750	75 to 95	120	0.02	150 to 900
LTC6228	0.9	0.94	6.3	9000 or 20,000	95 to 250	94	0.1	4000 or 44,000

Table 3. Op Amp Comparison Continued

Device #	I_{BIAS} vs. V_{CM} (nA/V)	I_{BIAS} vs. V_{CM} Nonlinearity (pA/V)	PSRR dB (min)	GBW (MHz)	Linear Output Current mA (min)	DC V_{OS} vs. V_{OUT} Nonlinearity ppm	Distortion: AV = 2, 2 V p-p Out 10 kHz dBc	Macromodel Shows Distortion?
AD797			110 or 114	110	±30		–120	Simulation model is too optimistic
AD8597	5	0.2	118	14	±20		~–120	Compare with data sheet
ADA4075	2	3000	100	6.5	~±15		~–130	Optimistic
ADA4610	Very small	0.1	100 or 103	12	~±30		~–114	None
ADA4805	2.2	4000	100	30	~±30	30	–125	None
ADA4807	0.7	~140	98	17	±50		~–130	None
ADA4898			98	120	±40		~–120	None
ADA4899			84	280			–117	None
LT1468	3.5	600	100	55	±15		~–120	Yes
LT1678			1000.7	10	~±10		~–120	Simulation model is too optimistic
LT6018			128	12	~±15	0.02	~–115	Yes
LTC6228	300	28 or 140	95	800	±20	0.2	–120	

Conclusions

Sadly, commercially available ppm-accurate amplifiers are difficult, if not impossible, to find. There are ppm-linear amplifiers, but attention must be paid to the amplifier's input currents that create distortion against application impedances. Those impedances can be lowered, but driving them in feedback runs the risk of creating distortion at the op amp input. By using an op amp with particularly low input currents and variations, the application impedance can be raised to obtain the best distortion from the op amp, but this will raise system noise. Careful op amp selection and applications circuitry optimization are required for ppm linearity and noise.

About the Author

Barry Harvey has worked as an analog IC designer, designing high speed op amps, voltage references, mixed-signal circuits, video circuits, DSL line drivers, DACs, sample-and-hold amplifiers, multipliers, and more. He has an M.S.E.E. from Stanford University. He holds more than 20 patents and has published about as many articles and papers. Barry's hobbies include repairing used test equipment, playing guitar, and working on Arduino-related projects. He can be reached at *barry.harvey@analog.com*.

High Voltage Boost and Inverting Converters for Communications

Jesus Rosales

The field of electronic communications is rapidly expanding into every aspect of ordinary life. Detection, transmission, and reception of data require a wide array of devices such as optical sensors, RF MEMS, PIN diodes, APDs, laser diodes, and high voltage DACs, to name just a few. In many cases, these devices require several hundred volts to operate, calling for dc-to-dc converters that meet stringent efficiency, space, and cost requirements.

Analog Devices' LT8365 is a versatile monolithic boost converter that integrates a 150 V, 1.5 A switch, making it ideal for high voltage applications found in the communications field, including portable devices. High voltage outputs are easily produced from inputs as low as 2.8 V and as high as 60 V. It features optional spread spectrum frequency modulation, which can help mitigate EMI, and many other popular features detailed in the data sheet.

The converters shown in Figure 1 and Figure 2 have been used to provide the positive and negative voltage rails for high voltage DACs, MEMS, RF switches, and high voltage op amps, from a 12 V input source. These converters operate in discontinuous conduction mode (DCM) and deliver as much as 10 mA, with +250 V and –250 V output voltages with a conversion efficiency of about 80%.

Step-Up Ratios > 1:40

One benefit of DCM operation in a boost converter is the ability to achieve a high step-up ratio independent of duty cycle. Additionally, the values and physical sizes of the inductor and output capacitor can be reduced, which leads to a smaller overall footprint solution on the PCB. The circuit in Figure 3 can easily fit in an area less than 1 cm^2.

There are situations when only a very low input source is available and a high output voltage is needed. The converter shown in Figure 3 could be used to drive a variety of avalanche photo diodes, PIN diodes, and other devices requiring high bias voltages. This boost converter produces 125 V from a 3 V input source with up to 3 mA of load current.

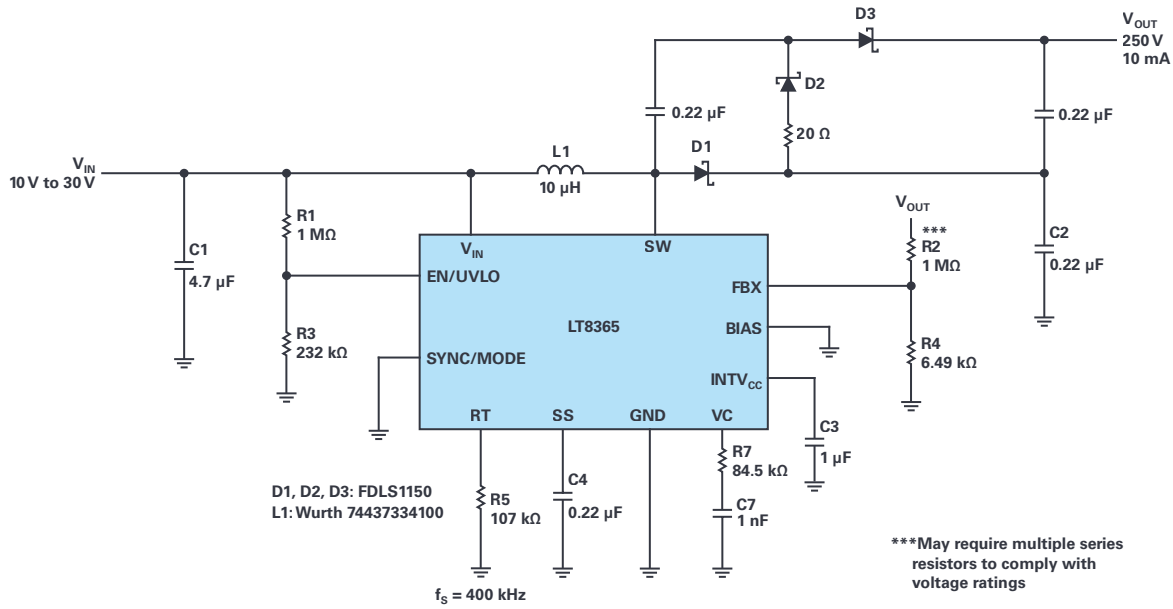

Figure 1. 12 V input to 250 V output 2-stage boost converter.

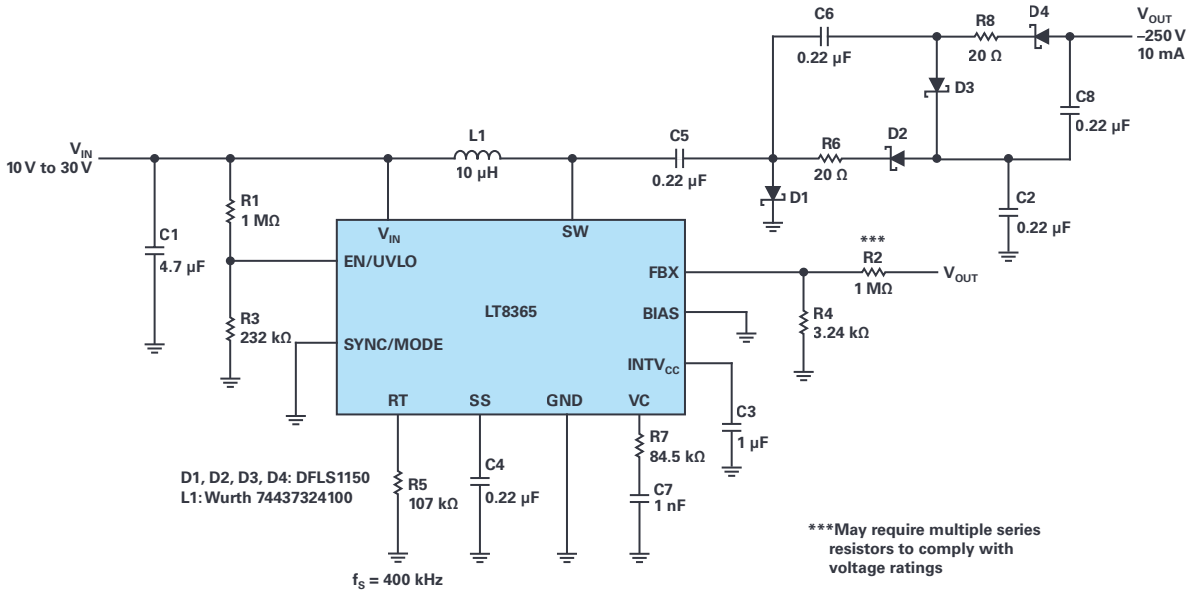

Figure 2. 12 V input to –250 V output 2-stage inverting converter.

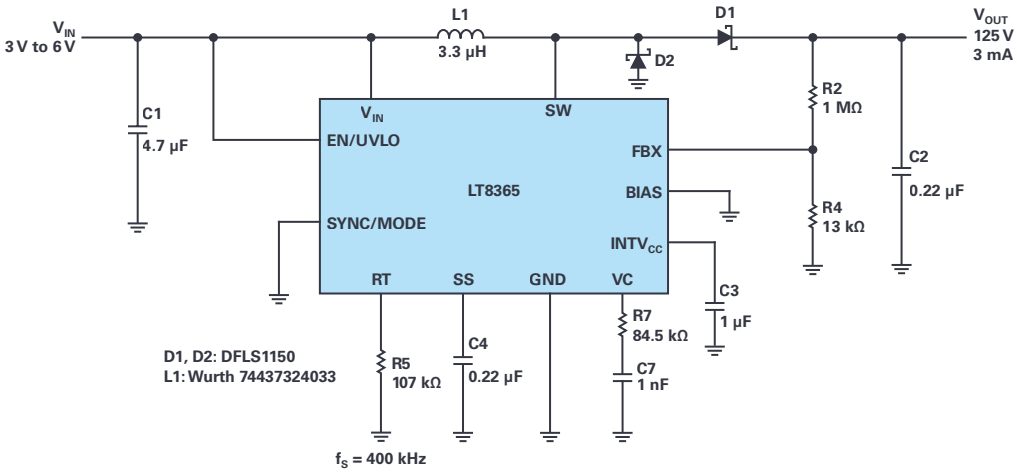

Figure 3. 3 V input to 125 V output boost converter.

The converter shown in Figure 4 extends the 125 V output to 250 V from a 3 V input source and supports about 1.5 mA. There are many devices in the communications field requiring such high bias voltages from low input voltage sources.

How High or Low Can You Go?

For situations where very high voltage is needed, whether positive or negative, a boost converter can use multiplier stages to boost the output 2×, 3×, or more. The converters in Figure 1 and Figure 2 show how to double the switch voltage in both directions, positive and negative. The 3-stage boost converter in Figure 5 delivers 375 V at 8 mA from a 12 V input source.

Note that the available output current must decrease as output voltage increases, since the switch capability does not change. As an example, a single-stage converter designed to deliver 20 mA will deliver about 10 mA when a second stage is added. As additional stages are added, always ensure the peak switch current stays within the guaranteed switch current limit.

Output Voltage Sensing Simplified

The LT8365 offers a single FBX pin to sense the output voltage. A simple resistor divider connected to the FBX pin senses the output voltage, independent of output polarity, as observed on all the schematics presented in this article.

Conclusion

The LT8365 enables applications that require compact, efficient, high output voltage boost conversion from input voltages as low as 2.8 V, which is common in the field of communications. It can also be used as an inverting converter and in popular topologies such as CUK and SEPIC converters. The LT8365 is available in a small, thermally enhanced, 16-lead MSOP package.

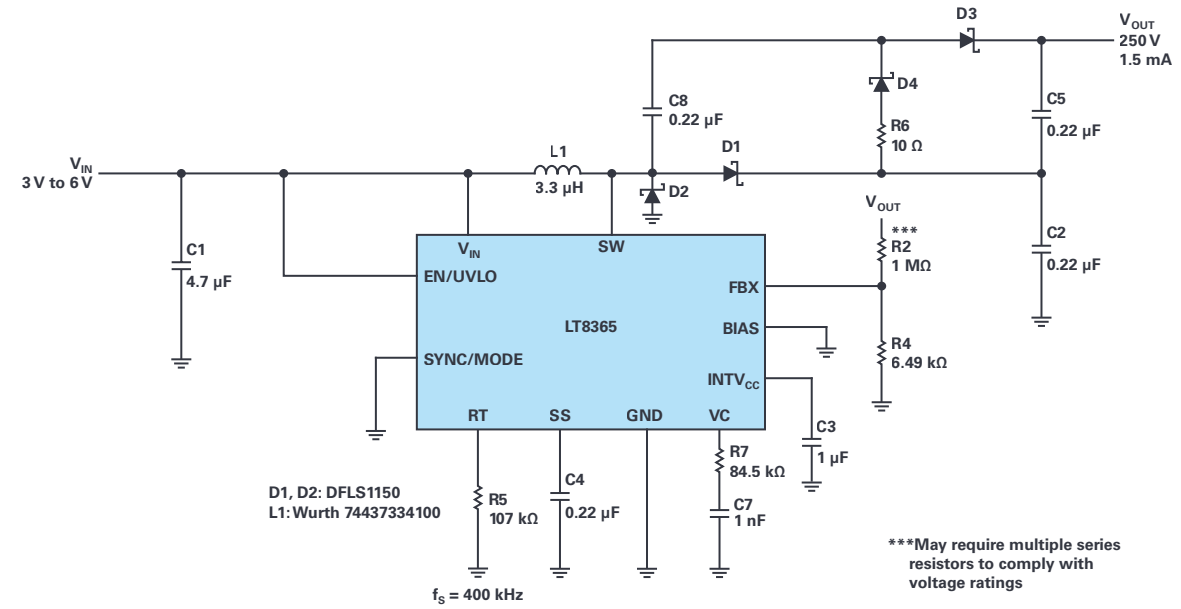

Figure 4. 3 V input to 250 V output 2-stage boost converter.

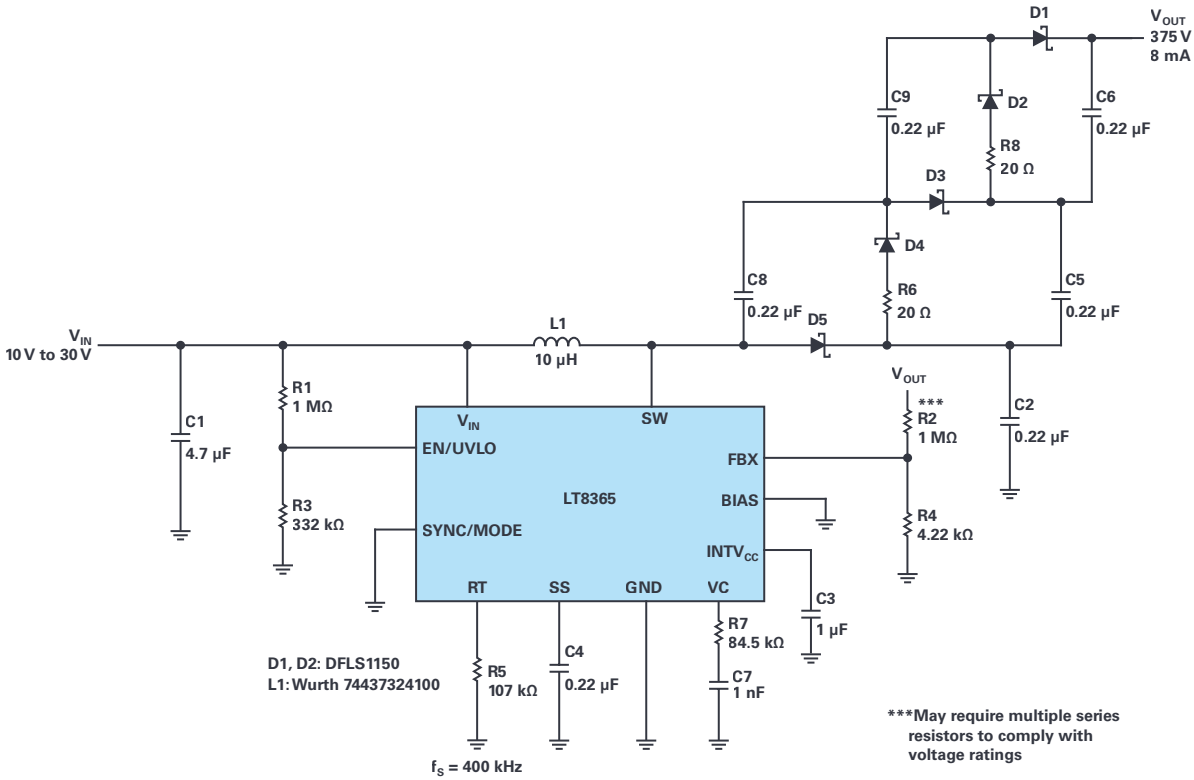

Figure 5. 12 V input to 375 V output 3-stage boost converter.

About the Author

Jesus Rosales is an applications engineer in Analog Devices' Applications Group in Milpitas, CA. He joined Linear Technology (now a part of ADI) in 1995 as an associate engineer and was promoted to applications engineer in 2001. He has since supported the boost/inverting/SEPIC family of monolithic converters and some offline isolated application controllers. He received an associate degree in electronics from Bay Valley Technical Institute in 1982. He can be reached at *jesus.rosales@analog.com*.

Rarely Asked Questions—Issue 167 When Smaller Is Better

Frederik Dostal

Question:

How does a µModule® regulator fit into such a small space?

Answer:

Many required components are already integrated.

Power modules have been on the market for many years now. A power module is a packaged, usually switched-mode power supply that can simply be soldered to a board and fulfills its task of converting an input voltage into a regulated output voltage. Compared with a switching regulator IC, where usually only the controller and the power switches are integrated into a chip, a power module also offers the integration of numerous passive components. Usually, the term *power module* is used when the inductor is integrated. Figure 2 shows the necessary components for a switched-mode step-down converter (buck topology). The dashed lines delineate the switching regulator IC and the power module. The development process for the voltage conversion circuit for these modules is assumed by the manufacturer of the power module, so the user does not have to be a power supply expert. There are other advantages besides this one. Through the high degree of integration in the module, the size of the switched-mode power supply can be especially small.

Quieter and Smaller DC-to-DC Regulation

Switching regulators naturally produce radiated EMI, as their operation requires high dI/dt events at relatively high frequency. EMI compliance is often mandatory and a critical design challenge for signal processing in medical equipment, RF transceivers, and test and measurement systems. For example, if a system fails EMI compliance or if the switching regulator impacts integrity of high speed digital or RF signals, debugging and redesign not only create long design cycles, they also elevate cost due to re-evaluation. Furthermore, the chance of noise is more pronounced in a denser PCB layout where the dc-to-dc switching regulators are in close proximity to noise sensitive components and signal routes.

Instead of relying on cumbersome EMI mitigation techniques—such as lowering the switching frequency, adding filter circuitry to PCB, or installing shielding—a better approach is to suppress the noise at the source: the dc-to-dc silicon itself. For a more compact dc-to-dc solution, all the components including MOSFETs, inductors, dc-to-dc ICs, and supporting components can be housed in a tiny overmolded package resembling a surface-mountable IC. See Figure 1.

Figure 1. The LTM8074 uses Silent Switcher® architecture for a complete low noise solution in a tiny package.

In addition to a quieter dc-to-dc conversion which meets most EMI compliance specifications such as EN 55022 Class B and small footprint, it's important to minimize the number of other components on the PCB such as output capacitors. With a fast transient response dc-to-dc regulator, dependency on output capacitance is reduced. What this means is that the design is simplified by optimized internal feedback loop compensation, which provides sufficient stability margins under a wide range of operating conditions with a broad range of output capacitors.

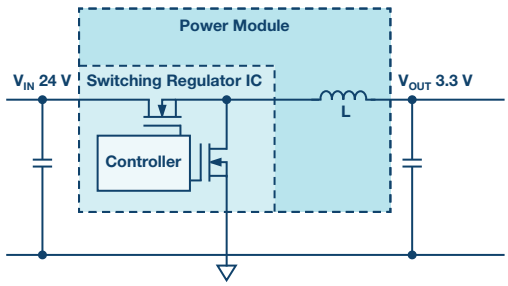

Figure 2. A step-down (buck) switching regulator highly integrated with the inductor in a power module.

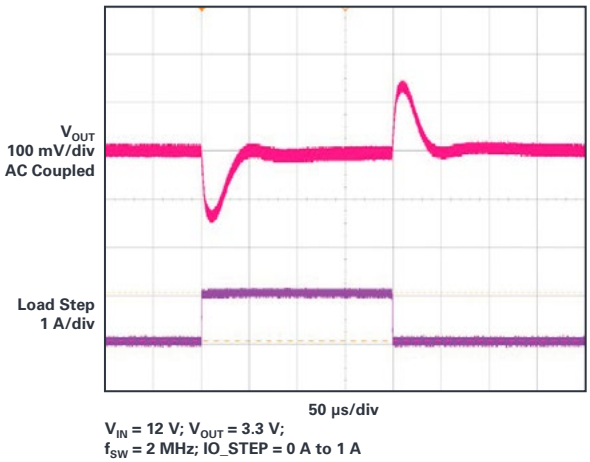

Figure 3. With minimal output capacitors (2 µF × 4.7 µF ceramics), the LTM8074 provides a quick transient response (12 V_{IN}, 3.3 V_{OUT}).

The LTM8074 is a 1.2 A, 40 V_{IN} µModule step-down regulator in a tiny 4 mm × 4 mm × 1.82 mm, 0.65 mm pitch BGA package. Its total solution size is 60 mm^2 for 3.2 V_{IN} to 40 V_{IN}, 3.3 V_{OUT} requiring only two 0805 capacitors and two 0603 resistors. The low profile and light-weight (0.08 g) package permits the device to be assembled on the backside of a PCB where the topside is often very densely populated. Its Silent Switcher architecture minimizes EMI emissions, enabling the LTM8074 to pass CISPR22 Class B, and reduce the possibility of EMC susceptibility to other sensitive circuits.

It is not usually possible to integrate all of the external components. There is a simple reason for this. If, for example, certain settings such as the switching frequency or the soft-start time should be adjustable, the circuit must be told what to do. This could be done in a digital manner. However, this would mean using a microcontroller and nonvolatile memory with the associated costs in the system. A common way of getting around this is to use external passive components for making these settings.

Input and output capacitors are often integrated into the power module but also sometimes externally required. Figure 2 shows a circuit with the new LTM8074 from Analog Devices.

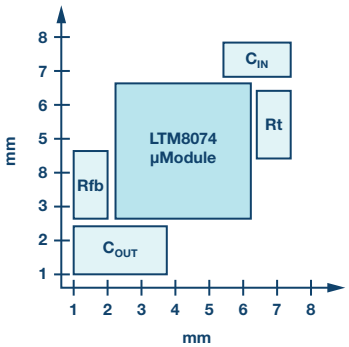

Figure 4. The LTM8074 with a V_{IN} of up to 40 V and an output current of 1.2 A in a space of just 4 mm × 4 mm.

Through the use of just one external resistor for setting the desired output voltage, the type variability is reduced and a certain amount of flexibility is provided for the application. If soft-start is not required, no capacitor has to be connected to the corresponding pin. All these capabilities make voltage conversion possible in an extremely small board area. With just the 4 mm × 4 mm edge length of the LTM8074 and minimal external wiring, the complete power supply unit can be operated on only approximately 8 mm × 8 mm of board area—with an input voltage of up to 40 V and a permissible output current of up to 1.2 A. Figure 3 shows an example layout with the minimal number of necessary external components.

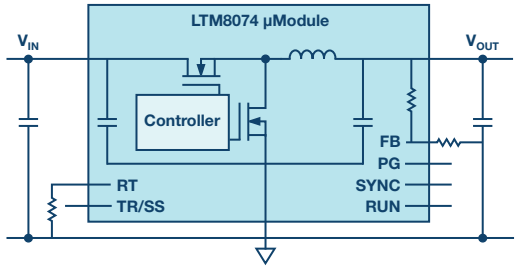

Figure 5. Example layout on a board area of approximately 8 mm × 8 mm.

For small power supplies, it is very important to offer an especially high conversion efficiency, otherwise there could potentially be problems with heat dissipation.

The new LTM8074, with an extremely compact size, is an ideal choice for this. Through its integrated Silent Switcher technology, it can even be used in circuits that are especially noise sensitive and are usually supplied by linear regulators.

Highly integrated power modules are not only suitable for simplifying the design of switched-mode power supplies, but are also useful for enabling efficient voltage conversion in an extremely small space.

The key performance characteristics of ADI µModule devices are:

- Lower noise (ultralow noise and Silent Switcher devices)
- Ultrathin packages
- 6-sided efficient cooling (CoP)
- Precision V_{OUT} regulation over line, load, and temperature
- Extreme reliability testing
- Minimum ground loops
- Multiple output on substrate
- Extreme temperature testing

About the Author

Frederik Dostal studied microelectronics at the University of Erlangen in 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 field applications engineer for power management at Analog Devices in München. He can be reached at *frederik.dostal@analog.com*.

Composite Amplifiers: High Output Drive Capability with Precision

Jino Loquinario

Introduction

It is normal, and almost expected, to be faced with applications for which a solution does not appear to exist. To meet their requirements would require us to think of a solution that is beyond the performance of current products that the market offers. For example, an application may require an amplifier that is high speed and high voltage with high output drive capability, but may also demand excellent dc precision, low noise, low distortion, etc.

Amplifiers that meet the speed and output voltage/current requirements and amplifiers with outstanding dc precision are readily available in the market—many of them, in fact. However, all the requirements may not exist in a single amplifier. When faced with this problem, some would think it is impossible for us to meet the demands of such applications, and that we must settle for a mediocre solution and go with either a precision amplifier or a high speed amplifier, perhaps sacrificing some of the requirements. Fortunately, this is not entirely true. There is a solution for this in the form of a composite amplifier, and this article will show how it is possible.

The Composite Amplifier

A composite amplifier is an arrangement of two individual amplifiers configured in such a way as to realize the benefits of each individual amplifier while diminishing the drawbacks of each amplifier.

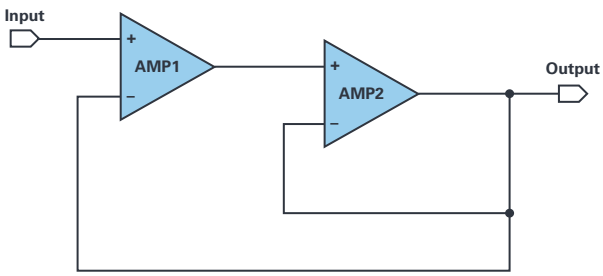

Figure 1. Simple composite amplifier configuration.

Referring to Figure 1, AMP1 should have excellent dc precision as well as the noise and distortion performance required by the application. AMP2 should provide the output drive requirements. In this arrangement, the amplifier (AMP2) with required output specifications is placed inside the feedback loop of an amplifier (AMP1) with the required input specifications. Some of the techniques and benefits of this arrangement will be discussed.

Setting the Gain

When initially encountering a composite amplifier, the first question that may arise is how to set the gain. To address this, it is helpful to view the composite amplifier as a single noninverting op amp contained within the large triangle, as in Figure 2. If we imagine the triangle is blacked out so that we couldn't see what's inside, then the gain of the noninverting op amp is 1 + R1/R2. Revealing the composite configuration inside the triangle doesn't change anything—the gain of the whole thing is still controlled by the ratio of R1 and R2.

In this configuration it is tempting to think that changing the gain of AMP2 by means of R3 and R4 will affect the output level of AMP2, indicating a change in composite gain, but this is not the case. Increasing the gain around AMP2 via R3 and R4 will simply decrease the effective gain, and output level, of AMP1 such that the output of the composite (AMP2 output) remains unchanged. Alternatively, decreasing the gain around AMP2 will serve to increase the effective gain of AMP1. So, in general, the gain of the composite amplifier is only dependent on R1 and R2.

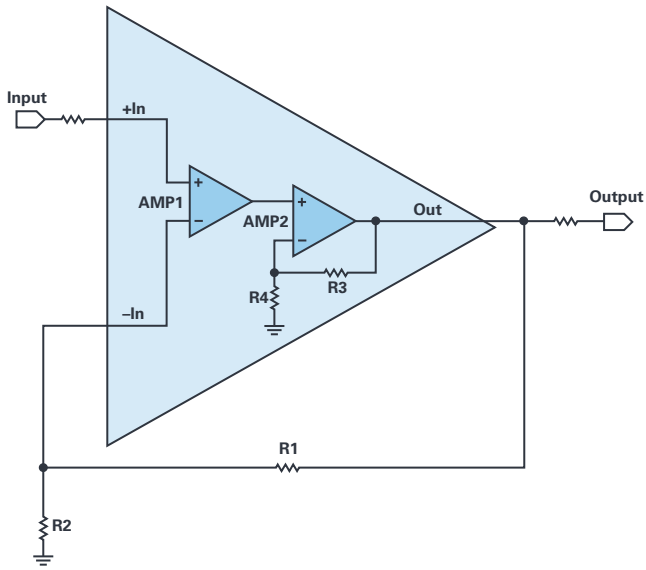

Figure 2. Composite amplifier seen as a single amplifier.

This article will discuss the major benefits and design considerations when implementing a composite amplifier configuration. The effects on bandwidth, dc precision, noise, and distortion will be highlighted.

Bandwidth Extension

One of the major benefits of implementing a composite amplifier is the extended bandwidth as compared to a single amplifier configured with the same gain.

Referring to Figure 3 and Figure 4, let's say we have two separate amplifiers each having a gain-bandwidth product (GBWP) of 100 MHz. Putting them together in a composite configuration will increase the effective GBWP

of the combination. At unity gain, the composite amplifier offers a ~27% higher –3 dB bandwidth, albeit with a small amount of peaking. However, at higher gains this benefit becomes much more noticeable.

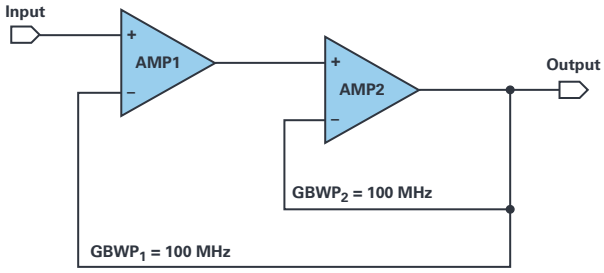

Figure 3. Composite amplifier at unity gain.

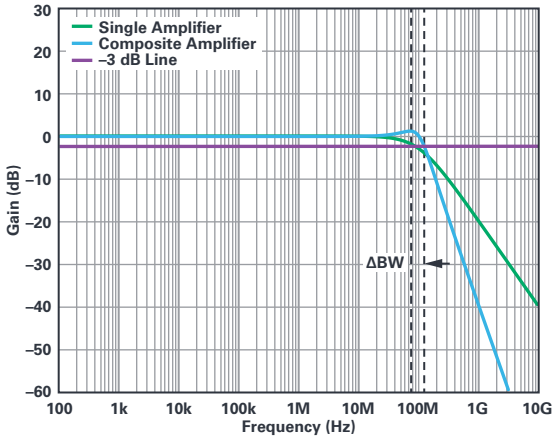

Figure 4. –3 dB BW improvement at unity gain.

Figure 5 shows the composite amplifier in a gain of 10. Note the composite gain is set to 10 via R1 and R2. The gain around AMP2 is set to approximately 3.16, forcing the effective gain of AMP1 to be the same. Splitting the gain equally between the two amplifiers yields the greatest possible bandwidth.

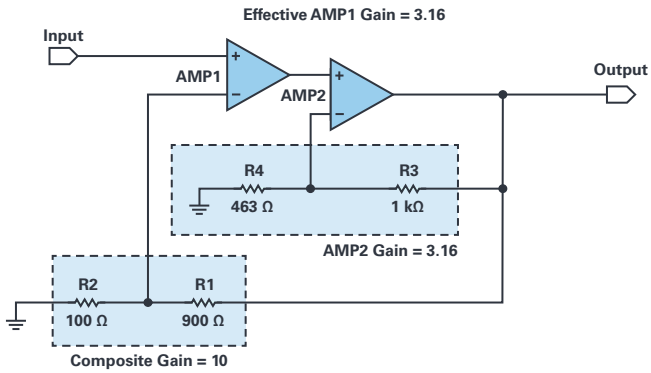

Figure 5. Composite amplifier configured for gain = 10.

Figure 6 shows the frequency response for a single amplifier at a gain of 10 compared to a composite amplifier configured with the same gain. In this case, the composite offers a ~300% increase in –3 dB bandwidth. How is this possible?

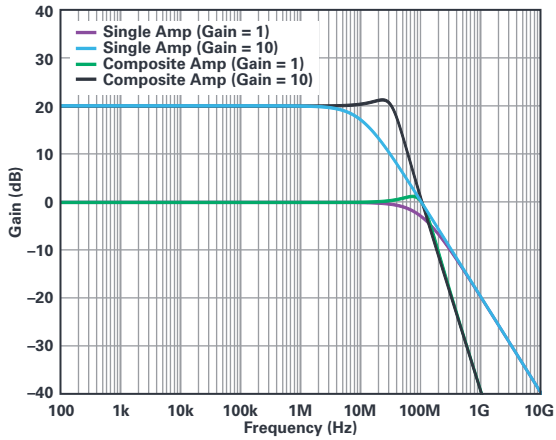

Figure 6. –3 dB BW improvement for gain = 10.

For a specific example, refer to Figure 7 and Figure 8. We require a system gain of 40 dB and will use two identical amplifiers, each with an open-loop gain of 80 dB and a GBWP of 100 MHz.

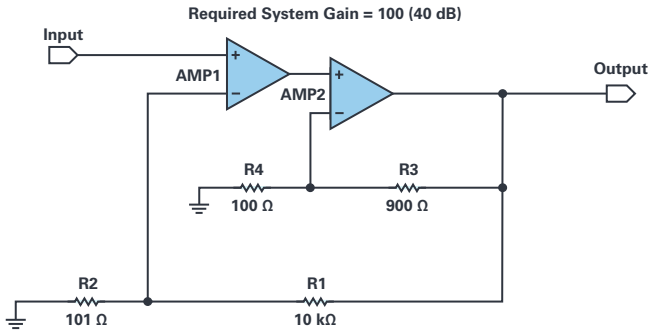

Figure 7. Gain splitting for maximum bandwidth.

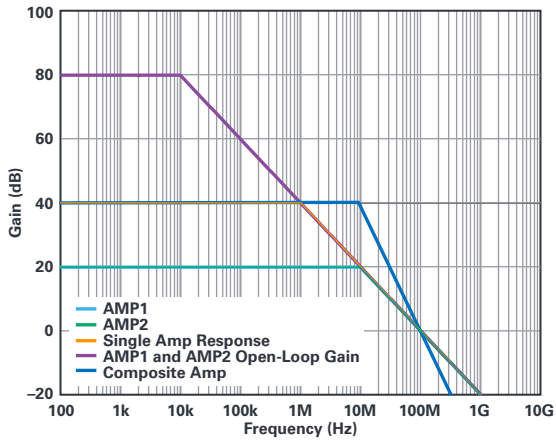

Figure 8. Expected response of a single amplifier.

To realize the highest possible bandwidth for the combination, we will split the required system gain equally between the two amplifiers, giving each of them a gain of 20 dB. So, setting the closed-loop gain of AMP2 to 20 dB forces the effective closed-loop gain of AMP1 to 20 dB as well. With this gain configuration, both amplifiers operate lower on the open-loop curve

than either of them would at a gain of 40 dB. As a result, the composite will have higher bandwidth at the gain of 40 dB as compared to the single amplifier solution of the same gain.

Although this may sound relatively simple and easy to implement, proper care should be taken in designing the composite amplifier to have the highest possible bandwidth without sacrificing the stability of the combination. In real-world applications where amplifiers are nonideal, and probably non-identical, a proper gain arrangement must be ensured to maintain stability. Also, note that the composite gain will roll off at –40 dB/decade, so one must be careful when distributing the gain between the two stages.

In some cases, splitting the gain equally may not be possible. To that point, equal distribution of the gain between the two amplifiers requires that the GBWP of AMP2 must always be greater than or equal to GBWP of AMP1, otherwise peaking—and possibly instability—will result. In a case where AMP1 GBWP must be greater than AMP2 GBWP, the instability can typically be corrected by redistribution of the gain between the two amplifiers. In this case, reducing the gain of AMP2 causes an increase in the effective gain of AMP1. The result is that AMP1 closed-loop BW is decreased as it operates higher on the open-loop curve and AMP2 closed-loop bandwidth is increased as it operates lower on the open-loop curve. If this slowing down of AMP1 and speeding up of AMP2 is adequately applied, the stability of the composite combination is restored.

For this article, the AD8397 was picked as the output stage (AMP2), interfaced with various precision amplifiers for AMP1 to demonstrate the benefits of a composite amplifier. The AD8397 is a high output current amplifier capable of delivering 310 mA.

Table 1. Bandwidth Extension on Different Amplifier Combinations for Gain of 10 and V_{OUT} = 10 V p-p

Amplifier	Single Amp BW (kHz)	Composite Amp BW (kHz)	% BW Extension
ADA4091	30	94	213
AD8676	165	517	213
AD8599	628	2674	325

Preserved DC Precision

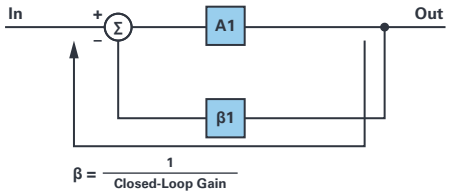

Figure 9. Operational amplifier feedback loop.

In a typical operational amplifier circuit, a portion of the output is fed back to the inverting input. Errors that are present on the output which were generated in the loop are multiplied by the feedback factor (β) and subtracted out. This helps maintain the fidelity of the output with respect to the input multiplied by the closed-loop gain (A).

Figure 10. Composite amplifier feedback loop.

For the composite amplifier, amplifier A2 has its own feedback loop, but A2 and its feedback loop are all inside the larger feedback loop of A1. The output now contains the larger errors due to A2 which are fed back to A1 and corrected. The larger correction signal results in the precision of A1 being preserved.

The effect of this composite feedback loop can be clearly seen in the circuit and results in Figure 11 and Figure 12. Figure 11 shows a composite amplifier comprised of two ideal op amps. The composite gain is 100 and the AMP2 gain is set to 5. $V_{OS}1$ represents a 50 µV offset voltage for AMP1 while $V_{OS}2$ represents a variable offset voltage for AMP2. Figure 12 shows that as $V_{OS}2$ is swept from 0 mV to 100 mV, the output offset is not affected by the magnitude of error (offset) contributed by AMP2. Instead, the output offset is proportional only to the error of AMP1 (50 µV multiplied by the composite gain of 100) and remains at 5 mV regardless of the value of $V_{OS}2$. Without the composite loop, we would expect the output error to increase as high as 500 mV.

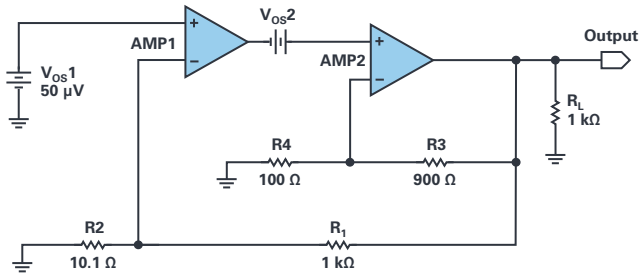

Figure 11. Offset error contribution.

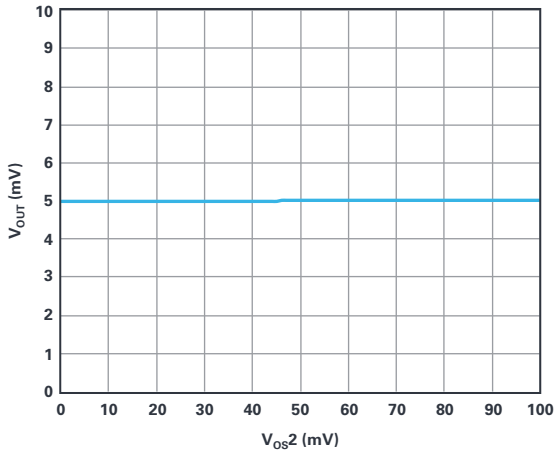

Figure 12. Composite output offset vs. $V_{OS}2$.

Table 2. Output Offset Voltage for Gain of 100

Amplifier	Effective V_{OS} (mV)	V_{OS} Reduction (Composite Configuration)
AD8397	100	
AD8397 + ADA4091	3.5	28.6×
AD8397 + AD8676	1.2	83.3×
AD8397 + AD8599	1	100×

Noise and Distortion

The output noise and harmonic distortion of the composite amplifier are corrected in a similar fashion as the dc errors, but, in the case of ac parameters, the bandwidth of the two stages also comes into play. We will look at an example using output noise to illustrate this with the understanding that distortion cancellation occurs in much the same manner.

Referring to the example circuit in Figure 13, for as long as the first stage (AMP1) has enough bandwidth, it will correct for the larger noise of the second stage (AMP2). As AMP1 begins to run out of bandwidth, the noise from AMP2 will begin to dominate. However, if AMP1 has too much bandwidth, and peaking is present in the frequency response, a noise peak will be induced at the same frequency.

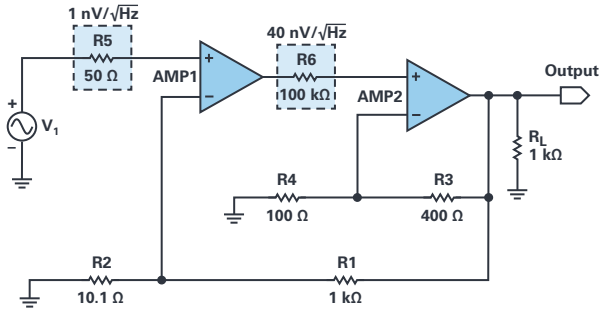

Figure 13. Noise sources of composite amplifier.

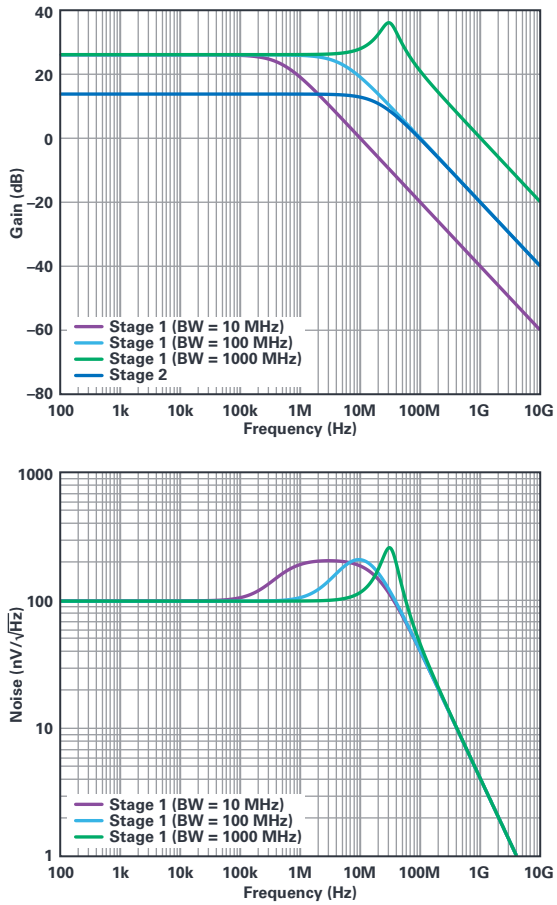

Figure 14. Noise performance vs. stage 1 bandwidth.

For this example, resistors R5 and R6 in Figure 13 represent the inherent noise sources for AMP1 and AMP2, respectively. The top plot of Figure 14 shows the frequency response for various AMP1 bandwidths as well as that of AMP2 for a single fixed bandwidth. Recall from the section on gain-splitting that a composite gain of 100 (40 dB) and AMP2 gain of 5 (14 dB) will force an effective AMP1 gain of 20 (26 dB) as can be seen here.

The bottom plot shows the wideband output noise density for each case. At low frequencies, the output noise density is dominated by AMP1 (1 nV/√Hz times the composite gain of 100 equals 100 nV/√Hz). This will continue for as long as AMP1 has enough bandwidth to compensate for AMP2.

For the cases where AMP1 has less bandwidth than AMP2, the noise density will begin to be dominated by AMP2 as AMP1 bandwidth begins to roll off. This can be seen in two of the traces of Figure 14 as the noise climbs to 200 nV/√Hz (40 nV/√Hz times the AMP2 gain of 5). Lastly, in the case where AMP1 has much greater bandwidth than AMP2, resulting in a peaking in the frequency response, the composite amplifier will exhibit a noise peak at the same frequency, also shown in Figure 14. Since the frequency response peaking results in excessive gain, the amplitude of the noise peak will also be higher.

Table 3 and Table 4 shows the effective noise reduction and the THD+n improvement when using various precision amplifiers as the first stage in a composite amplifier with the AD8397.

Table 3. Noise Reduction Using Different Front-End Amplifiers with Effective Gain = 100 at f = 1 kHz

Configuration	Noise, e_n (nV/√Hz)	Effective Noise Reduction (%)
AD8397 Only	450	
AD8397 + ADA4084	390	13.33
AD8397 + AD8676	280	37.78
AD8397 + AD8599	107	76.22

Table 4. THD+N Comparison Using Different Front-End Amplifiers with Effective Gain = 10 at f = 1 kHz and I_{LOAD} = 200 mA

Configuration	Effective THD+n (dB)	THD+n Improvement (dB)
AD8397 Only	–100.22	
AD8397 + ADA4084	–105.32	5.10
AD8397 + AD8676	–106.68	6.46
AD8397 + AD8599	–106.21	5.99

System-Level Application

In this example, the goal for a DAC output buffer application is to provide an output of 10 V p-p into a low impedance probe with a current of 500 mA p-p, low noise and distortion, excellent dc precision, and as high of a bandwidth as possible. The output of a 4 mA to 20 mA current-out DAC is to be converted to a voltage by the TIA, then to the input of the composite amplifier for more amplification. With AD8397s on the output, the output requirements are attainable. AD8397 is a rail-to-rail, high output current amplifier capable of delivering the needed output current.

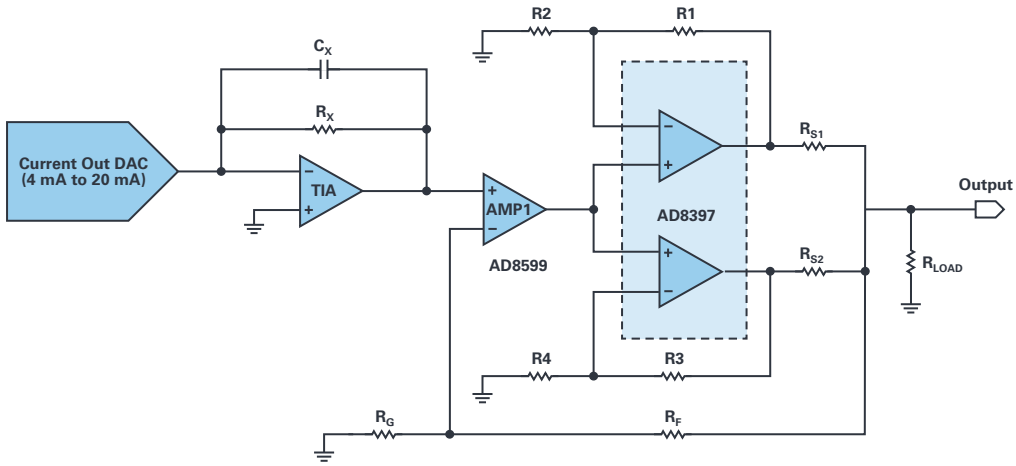

Figure 15. Application circuit for DAC output driver.

AMP1 could be any precision amplifier that has the desired dc precision needed for the configuration requirement. In this application, various front-end precision amplifiers could be used with AD8397 (and other high output current amplifiers) to attain both the excellent dc requirements and the high output capability drive needed for the application.

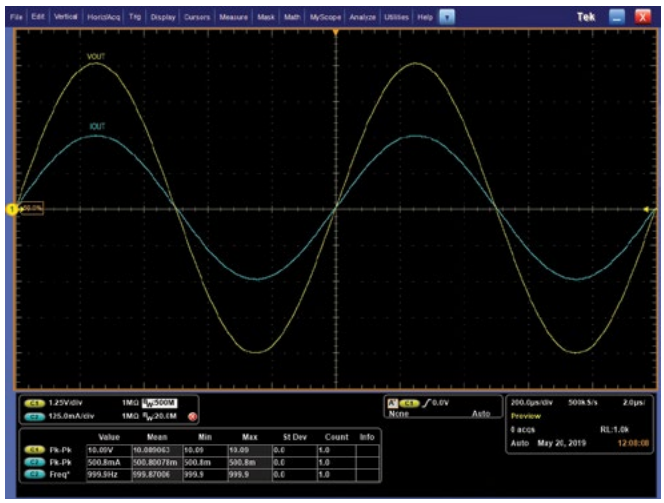

Figure 16. V_{OUT} and I_{OUT} for AD8599 and AD8397 composite amplifier.

Table 5. AD8599 + AD8397 Composite Amplifier Specifications

Parameter	Value
Gain	10 V/V
–3 dB Bandwidth	1.27 MHz
Output Voltage	10 V p-p
Output Current	500 mA p-p
Output Offset Voltage	102.5 µV
Voltage Noise (f = 1 kHz)	20.95 nV/√Hz
THD+n (f = 1 kHz)	–106.14 dB

This configuration is not limited to AD8397 and AD8599, but is possible with other combinations of amplifiers to cater this output drive specification that requires excellent dc precision. The amplifiers in Table 6 and Table 7 are also suited for this application.

Table 6. Amplifiers with High Output Current Drive

High Output Current Amplifiers	Current Drive (A)	Slew Rate	V_S Span, Max (V)
ADA4870	1	2.5 kV/µs	40
LT6301	1.2	600 V/µs	27
LT1210	2	900 V/µs	36

Table 7. Precision Front-End Amplifiers

Precision Amplifiers	V_{OS} (µV)	V_{NOISE}, e_n (nV/√Hz)	THD+n, 1 kHz (dB)
LT6018	50	1.2	–115
ADA4625	80	3.3	–110
ADA4084	100	3.9	–90

Conclusion

With the composite amplifier, the marriage of two amplifiers realizes the best specifications that each one offers while compensating for their limitations. Amplifiers with high output drive capability combined with precision front-end amplifiers could provide solutions to applications with challenging requirements. When designing, always consider stability, noise peaking, bandwidth, and slew rate for optimum performance. There are plenty of possible options to cater a wide range of applications. With the proper implementation and combination, striking the right balance for the application is highly achievable.

Acknowledgements

I would like to thank Zoltan Frasch and Bruce Petipas for their technical contributions to this article.

About the Author

Jino Loquinario joined ADI in 2014 and works as a product applications engineer for the Linear Products and Solutions Group. He graduated from Technological University of the Philippines–Visayas with a bachelor's degree in electronics engineering. He can be reached at *jino.loquinario@analog.com*.

High Accuracy Resolver Simulator System with Fault Injection Function

Nandin Xu

Introduction

Considering that resolvers have outstanding reliability and high precision performance in harsh and severe environments for a very long time, they are widely used in EV, HEV, EPS, inverter, servo, railway, high speed train, aerospace, and other applications where position and velocity information is needed.

Many resolver-to-digital converters (RDCs), such as Analog Devices' AD2S1210 and AD2S1205, used in previous systems to decode the resolver's signal to get digital position and velocity data. Interferences and fault issues tend to happen in customers' systems, and most of the time they want to evaluate the accuracy performance of angle and velocity under the interference conditions, find and validate the root cause, then fix and optimize the system. A high accuracy resolver simulator (which simulates a resolver attached to the real motor at constant speed or position) with fault injection can solve the interference and fault pain-points without setting up a complicated motor control system.

This article will analyze the error contribution in resolver simulator systems and give some error calculation examples to help understand why high precision is so important in resolver simulators. It will then show the fault case under interference conditions in field applications. Following that will be a description of how to build a high accuracy resolver simulator with fault simulation and injection functions using the latest high precision products. Finally, it will show some of the resolver simulator's capabilities.

Error Contributions in Resolver Simulator Systems

First, this section will introduce the ideal resolver structure. Then, five commonly nonideal characteristics and error analysis methods will be given to help you understand why high precision is necessary in resolver simulator systems.

The resolver simulator will simulate a resolver, as shown in Figure 1, attached to a real motor at constant speed or position. For classic or variable reluctance resolvers, a rotor and a stator are included. A resolver can be thought of as a special transformer. On the primary side, as expressed in Equation 1, EXC is the excitation sinusoidal input signal. On the secondary side, as expressed in Equation 2 and Equation 3, SIN and COS are the modulated sinusoidal signal at the two outputs.

$$EXC = A_0 sin(\omega t) \quad (1)$$

$$SIN = A_0 T sin(\theta) sin(\omega t) \quad (2)$$

$$COS = A_0 T cos(\theta) sin(\omega t) \quad (3)$$

where:

θ is the shaft angle, ω is the excitation signal frequency, A_0 is the excitation signal amplitude, and T is the resolver transformation ratio.

The modulated SIN/COS signals are shown in Figure 2. For constant angular θ in different quadrants, the SIN/COS signal will have in-phase and antiphase cases. For constant velocity, the frequency of the SIN/COS envelope is constant, indicating the velocity information.

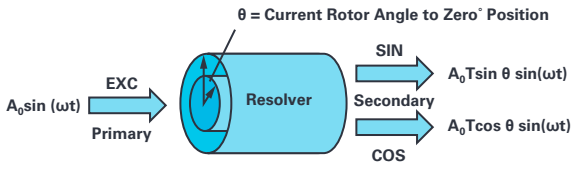

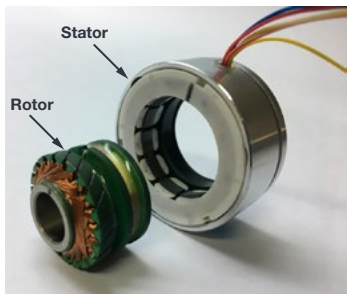

Figure 1. Resolver structure.

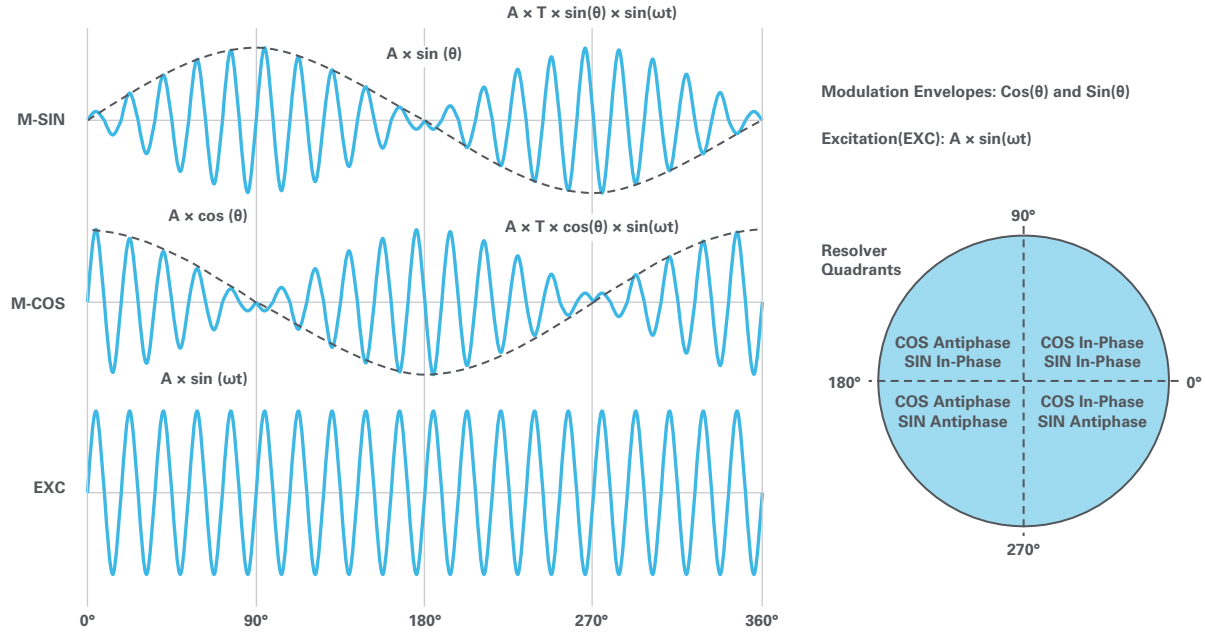

Figure 2. Resolver electrical signal.

For all of ADI's RDC products, the demodulation signal is expressed in Equation 4. The Type II tracking loop will be accomplished when φ (output digital angle) equals θ (position of the rotor), the resolver angle. In a real resolver system with amplitude mismatch, phase shift, imperfect quadrature, excitation harmonic, and inductive harmonic, any of these five nonideal conditions may happen and contribute error.

$$A_0T(sin(\theta)cos\varphi - cos(\theta)sin(\varphi)) \quad (4)$$

Amplitude Mismatch

Amplitude mismatch is the difference in the peak-to-peak amplitudes of the SIN and COS signals when they are at their peak amplitudes, with 0° and 180° for COS, and 90° and 270° for SIN. Mismatch can be introduced by variations of the resolver windings or by unbalanced gain control of the SIN/COS input. To determine the position error created by amplitude mismatch, Equation 3 can be rewritten as Equation 5.

$$COS = A_0T(1 + a)cos(\theta)sin(\omega t) \quad (5)$$

Where a represents the amount of mismatch between SIN and COS signals, the remaining envelope signal after demodulation can readily be shown, as in Equation 6. When driving the envelope signal to zero in a Type II tracking loop, by setting Equation 6 to equal zero, it is possible to find the position error ε = θ – φ. Then we can receive error information, as shown in Equation 7.

$$A_0T(sin(\theta)cos(\varphi) - cos(\theta)sin(\varphi) - acos(\theta)sin(\varphi)) \quad (6)$$

$$\varepsilon = sin^{-1}\left(\frac{a}{a+2}sin(\theta + \varphi)\right) \quad (7)$$

For the realistic case when a is small, the position error is also small, which implies that sin(ε) ≈ ε and θ + φ ≈ 2θ. So, Equation 7 becomes Equation 8, and the error term is expressed in radians.

$$\varepsilon \approx \frac{a}{2}sin(2\theta) \quad (8)$$

As shown in Equation 8, the error term oscillates at twice the rate of rotation while a maximum error of a/2 is reached at odd integer multiples of 45°. Assume the amplitude mismatch is 0.3%, substitute the variables in Equation 8, and, using an odd integer multiple of 45°, the maximum error will be expressed in Equation 9, where m is an odd integer.

$$\varepsilon_{MAX} = \frac{0.003}{2}sin(2 \times 45° \times m) = 0.0015 \quad (9)$$

The error, calculated in radians, can be converted to LSBs via Equation 10 when the RDC mode is 12 bits, or about 1 LSB.

$$\varepsilon_{12BIT\_LSB} = \left\{\frac{180}{\pi} \times (0.0015) \times \left(\frac{360}{2^{12}}\right)^{-1}\right\} = 0.9778 \text{ LSB} \quad (10)$$

Phase Shift

Phase shift refers to both differential phase shift and common phase shift. Differential phase shift is the phase shift between the resolver's SIN and COS signals. Common phase shift is the phase shift between the excitation reference signal and the SIN and COS signals. To determine the position error created by differential phase shift, Equation 3 can be rewritten as Equation 11.

$$COS = A_0Tcos(\theta)\ sin(\omega t + a) \quad (11)$$

Where a represents the differential phase shift, the envelop signal remaining after demodulation can be expressed as Equation 12 when the quadrature term cos(wt)(sin(a)sin(θ)cos(φ)) is ignored. For the realistic case when a is small, cos(a) ≈ 1 – a²/2. When driving this signal to zero in a Type II tracking loop, with Equation 10 set to zero, it is possible to find the position error ε = θ – φ that results. Then we can get error information as shown in Equation 13.

$$A_0T(sin(\theta)cos(\varphi) - cos(\theta)sin(\varphi)cos(a)) \quad (12)$$

$$\varepsilon = \frac{a^2}{2}sin(\theta)cos(\varphi) \quad (13)$$

For θ ≈ φ, sin(θ)cos(φ), has a maximum of 0.5 at θ ≈ 45°. So, Equation 13 becomes Equation 14 with the error term expressed in radians.

$$\varepsilon = \frac{a^2}{2} \times 0.5 \quad (14)$$

Assume the differential phase shift is 4.44°, the error, which can be converted to LSBs using Equation 15 when the RDC mode is 12 bits, is about 1 LSB.

$$\varepsilon_{12BIT\_LSB} = \left\{\frac{180}{\pi} \times \left(0.25\left(4.44 \times \frac{\pi}{180}\right)^2\right) \times \left(\frac{360}{2^{12}}\right)^{-1}\right\} = 0.9778 \text{ LSB} \quad (15)$$

When the common phase shift is β, Equation 2 and Equation 3 can be rewritten as Equation 16 and Equation 17, respectively.

$$SIN = A_0 T sin(\theta) sin(\omega t + \beta) \quad (16)$$

$$COS = A_0 T cos(\theta) sin(\omega t + a + \beta) \quad (17)$$

Similarly, the error term can be expressed in Equation 18.

$$\varepsilon = 0.53 \times a \times \beta \quad (18)$$

Under static operating conditions, common phase shift will not affect the converter's accuracy, but resolvers at speed will generate speed voltages due to the reactive components of the rotor impedance and the signals of interest. Speed voltages, which only occur at speed and not at static angles, are in quadrature to the signal of interest. When the common phase shift is β, the tracking error can be approximated as Equation 19, where ω_M is the motor speed and ω_E is the excitation speed.

$$\varepsilon = \beta \times \frac{\omega_M}{\omega_E} \quad (19)$$

As shown in Equation 19, the error is proportional to the resolver speed and phase shift. Thus, in general, it is beneficial to use a high resolver excitation frequency.

Imperfect Quadrature

Imperfect quadrature indicates the two resolver signals that SIN/COS refer to in this situation are not exactly 90° quadrature. This occurs when the two resolver phases are not machined or assembled in perfect spatial quadrature. When β represents the amount of imperfect quadrature, Equation 2 and Equation 3 can be rewritten as Equation 20 and Equation 21.

$$SIN = A_0 T sin(\theta) sin(\omega t) \quad (20)$$

$$COS = A_0 T cos(\theta + \beta) sin(\omega t) \quad (21)$$

As before, the envelop signal remaining after demodulation can readily be shown as Equation 22. When you set Equation 22 to zero, assume β is small, cos(β) ≈ 1 and sin(β) ≈ β, it is possible to find the position error ε = θ − φ that results. Then we can receive error information, as shown in Equation 23.

$$A_0 T(sin(\theta)cos(\varphi) - cos(\theta + \beta)sin(\varphi)) \quad (22)$$

$$\varepsilon \approx -\beta sin^2(\theta) = -\beta(1 - cos(2\theta))/2 \quad (23)$$

As shown in Equation 23, the error term oscillates at twice the rate of rotation, while a maximum error of β/2 is reached at odd integer multiples of 45°. Compared to the error due to amplitude mismatch, in this case, the mean error is nonzero and the peak error is equal to the quadrature error. From the amplitude mismatch example, when β = 0.0003 radian = 0.172°, this can cause about 1 LBS error in 12-bit mode.

Excitation Harmonic

In all the preceding analysis, it was assumed that the excitation signal was an ideal sinusoid and contained no additional harmonics. In a real-world system, the excitation signal does contain harmonics. So, Equation 2 and Equation 3 can be rewritten as Equation 24 and Equation 25.

$$SIN = Tsin(\theta)\sum_{(n=0)}^{\infty} A_{sn}\, sin(n+1)\omega t \quad (24)$$

$$COS = Tcos(\theta)\sum_{(n=0)}^{\infty} A_{cn}\, sin(n+1)\omega t \quad (25)$$

The envelope signal remaining after demodulation can readily be shown, as in Equation 26. Driving this signal to zero in Type II tracking loop.

$$sin(\theta)cos(\varphi)\sum_{(n=0)}^{\infty} A_{sn} - cos(\theta)sin(\varphi)\sum_{(n=0)}^{\infty} A_{cn} \quad (26)$$

Set Equation 26 to zero and it is possible to find the position error ε = θ − φ that results. Then we can get error information as shown in Equation 27.

$$\varepsilon \approx \frac{\sum_{(n=0)}^{\infty}(A_{cn} - A_{sn})}{\sum_{(n=0)}^{\infty}(A_{cn} + A_{sn})}\, sin(2\theta) \quad (27)$$

If the resolver excitation has identical harmonics, the numerator of Equation 27 is zero and no position error is incurred. That means the common excitation harmonic has negligible effect on the RDC, even at very large values. However, if the harmonic content is different in SIN or COS, the position error incurred has the same functional shape as the amplitude mismatch shown in Equation 8. This will greatly affect the accuracy of the position.

Inductive Harmonic

In the real world, it is impossible to construct a resolver with inductance profiles that are perfect sinusoidal and cosinusoidal functions of position. Normally, the inductances will contain harmonics, and VR resolvers will contain dc components. Thus, Equation 2 and Equation 3 can be rewritten as Equation 28 and Equation 29, respectively, where K_0 indicates the dc component.

$$SIN = Tsin(wt)(K_0 + \sum_{(n=0)}^{\infty} K_n sin(n\theta)) \quad (28)$$

$$COS = Tsin(wt)(K_0 + \sum_{(n=0)}^{\infty} K_n cos(n\theta)) \quad (29)$$

$$\sqrt{2}K_0 cos\left(\varphi + \frac{\pi}{4}\right) + \sum_{(n=0)}^{\infty} K_n sin(n\theta - \varphi) \quad (30)$$

The remaining envelope signal after demodulation can be shown, as in Equation 30.

Driving this signal to zero in a Type II tracking loop, when the harmonic amplitudes are small, Kn << 1 for n > 1, the error information ε = θ − φ can be derived from Equation 31.

$$\varepsilon \approx \sqrt{2}K_0 cos\left(\varphi + \frac{\pi}{4}\right) - \sum_{(n=1)}^{\infty} K_n sin(n-1)\theta \quad (31)$$

According to the expression, the error is more sensitive to the dc term than the harmonic effect, it is proportional to the inductive harmonic amplitude. In the meantime, the nth inductance harmonic determines the amplitude of the $(n-1)^{th}$ harmonic of the position error.

Summary of Error Contribution in a Resolver Simulator System

Except for the error source mentioned, the interferent coupled to the SIN and COS lines, amplifier offset error, bias error, etc., can also contribute to the system error. The error source and contribution summary in the resolver simulator system, with the worst example of 1 LSB of 12-bit mode included, are shown in Table 1. Another RDC resolution mode can be calculated by referring to the table.

Table 1. Summary of Error Source and Contribution in a Resolver Simulator System

Error Source	Error Expression	Description	1 LSB Example
Amplitude Mismatch	$\frac{a}{2} \times \sin(2\theta)$	a = amplitude mismatch	0.003 amplitude mismatch contributes 1 LSB error
Phase Shift	$\frac{a^2}{2} \times 0.5$	a = differential phase shift	4.44° differential phase shift contributes 1 LSB
	$\beta \times \frac{\omega_M}{\omega_E}$	β = common phase shift ω_M = motor speed, ω_E = excitation speed	
Imperfect Quadrature	$-\beta(1-cos2\theta)/2$	β = angular deviation from perfect quadrature	0.172° imperfect quadrature contributes 1 LSB
Excitation Harmonic	$\frac{\sum_{(n=0)}^{\infty}(A_{cn}-A_{sn})}{\sum_{(n=0)}^{\infty}(A_{cn}+A_{sn})} sin(2\theta)$	A_{cn}, A_{sn} = harmonics amplitude	
Inductive Harmonic	$-\sqrt{2}K_0 cos\left(\varphi+\frac{\pi}{4}\right)-\sum_{(n=1)}^{\infty}K_n sin(n-1)\theta$	K_0 = dc component, K_n = harmonics amplitude	

Fault Types in an RDC System

In real RDC systems, a lot of fault cases can appear. The following sections will show different fault types and some fault signals from field tests, and how the fault type can be simulated when using the resolver simulator solution described in the third section. Except for the fault type mentioned, there may be random interference that leads to another fault, or some faults may happen at the same time.

Misconnection Fault

Misconnection means connected resolver excitation and SIN/COS pairs to RDC SIN/COS input and excitation output pins via incorrect connections. When a misconnection happens, RDC can also decode the angular and velocity information, but the angular output data will show some a jump, like an offset error in DAC output. The misconnection case and result data are shown in Figure 3. Where the first column shows the EXC/SIN/COS pins and output angular, the remaining columns show the misconnection cases.

AD2S1210 Signal	North Atlantic Signals, Input Angle 0°, 45°, 90°, 135°, 180°, Rotation Direction 0° to 360°									
EXC	EXC	EXC	EXC	EXC	EXC	EXC	EXC	EXC	/EXC	/EXC
/EXC	/EXC	/EXC	/EXC	/EXC	/EXC	/EXC	/EXC	/EXC	EXC	EXC
Sin	S3 (SinHi)	S1 (SinLo)	S1 (SinLo)	S3 (SinHi)	S2 (CosHi)	S4 (CosLo)	S4 (CosLo)	S2 (CosHi)	S3 (SinHi)	S2 (CosHi)
SinLo	S1 (SinLo)	S3 (SinHi)	S3 (SinHi)	S1 (SinLo)	S4 (CosLo)	S2 (CosHi)	S2 (CosHi)	S4 (CosLo)	S1 (SinLo)	S4 (CosLo)
Cos	S2 (CosHi)	S2 (CosHi)	S4 (CosLo)	S4 (CosLo)	S3 (SinHi)	S3 (SinHi)	S1 (SinLo)	S1 (SinLo)	S2 (CosHi)	S1 (SinLo)
CosLo	S4 (CosLo)	S4 (CosLo)	S2 (CosHi)	S2 (CosHi)	S1 (SinLo)	S1 (SinLo)	S3 (SinHi)	S3 (SinHi)	S4 (CosLo)	S3 (SinHi)
Output Angle (°)	0	0	180	180	90	270	270	90	180	270
	45	315	225	135	45	315	225	135	225	315
	90	270	270	90	0	0	180	180	270	0
	135	225	315	45	315	45	135	225	315	45
	180	180	0	0	270	90	90	270	0	90
Output Direction (°)	0–360	360–0	0–360	360–0	360–0	0–360	360–0	0–360	0–360	0–360

Figure 3. Resolver misconnection and angular output.

Phase Shift Fault

From the error contribution section, we know phase shift contains differential phase shift and common phase shift. Considering the differential phase can be thought of as the difference of common phase shift, in this section, the phase shift fault means the fault caused by common phase shift.

The common phase shift error contribution is shown in Figure 4. Phase 1 is excitation filter delay. Phase 2 is resolver phase shift. Phase 3 is line delay. Phase 4 is SIN/COS filter delay. In a field RDC system, when phase shift error happens, it means the total value of phases 1, 2, 3, and 4 is bigger than 44°. Normally, the resolver phase shift error is 10°. In bad cases, the total phase value can reach 30°. For MP consideration, enough phase margin needs to be left.

When the phase shift for SIN/COS are different, it can cause phase shift mismatch fault. If this happens, angular and velocity accuracy will be affected.

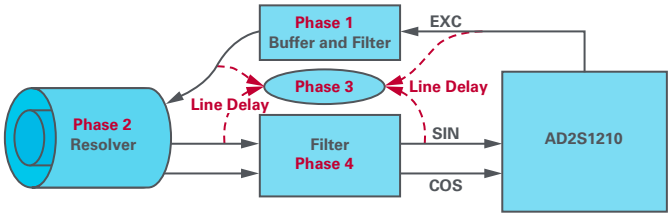

Figure 4. Phase shift error contribution.

Disconnection Fault

Disconnection fault happens when any line of the resolver is disconnected from the RDC platform interface. With the product safety upgrade, line disconnection detection is always mentioned by customers. This fault can be simulated to set SIN/COS to zero voltage. When disconnection happens, LOS/DOS/LOT fault can be triggered in AD2S1210.

Amplitude Mismatch/Exceed Fault

Amplitude mismatch happens when the circuit gain control or resolver ratio of SIN/COS are different, which also means the amplitude value of the SIN/COS envelope is different. When the amplitude is close to AVDD, it will trigger an amplitude exceed fault. For AD2S1210, this is called a clipping fault. A good SIN/COS signal example is shown in Figure 5.

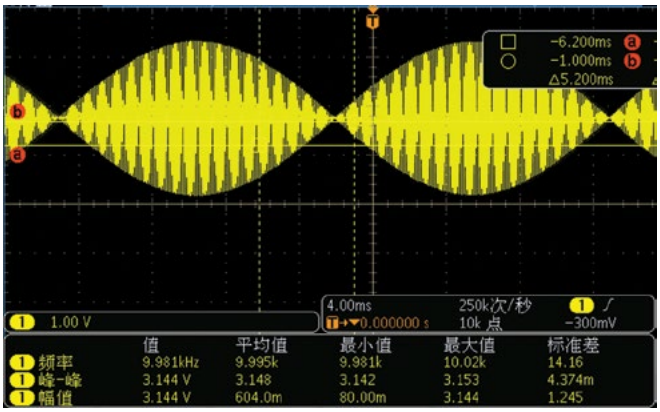

Figure 5. An ideal SIN/COS signal.

IGBT Disturbance Fault

Figure 6. A SIN/COS coupled IGBT disturbance.

An IGBT disturbance means that the interference signal coupled with the IGBT switch on/off effect. When the signal is coupled with the SIN/COS line, position and velocity performance can be affected, the angular value will have a jump, and the direction of the velocity may change. An example from the field is shown in Figure 6, where Channel 1 is the SIN signal, Channel 2 is the COS signal, and the spur indicates interference coupled with IGBT turn on/off.

Velocity Exceed Fault

Velocity exceed fault happens when the electrical velocity is higher than the resolver decode system. For example, in 12-bit mode, the max velocity that AD2S1210 can support is 1250 SPS, and when the resolver electrical velocity is 1300 SPS, velocity exceed fault will be triggered.

Resolver Simulator System Architecture and Description

From the first section, we know the amplitude and phase error directly determine the performance of decode angular and velocity performance. Luckily, ADI has a vast portfolio of precision products from which to choose and build your resolver simulator system. The following description will show how to build a high accuracy resolver simulator and discuss which parts to choose.

The simulator block diagram as shown in Figure 7 has seven modules to focus on:

1. The process control platform for data analysis and control.
2. The sync clock generation module generates sync clock for the subsystem.
3. The fault signal generation module generates different fault signals.
4. The SIN/COS generation module generates modulated SIN/COS signals as resolver outputs.
5. The signal capture module acts as the excitation and feedback signal capture module.
6. The SIN/COS output module handles SIN/COS output with buffer, gain, and filter included.
7. The excitation signal input module comes with a built-in buffer and filter circuit.
8. The power module delivers the supply power for ADC, DAC, switch, amp, etc. components.

The resolver simulator system works by having the signal capture module sample the excitation signal from the input module, where the processor will analyze the frequency and amplitude. The processor will calculate the SIN/COS DAC output data code by using the CORDIC algorithm and, through the SIN/COS module, generate the same frequency sinusoidal signal as the excitation input. Then the system will recapture the excitation and SIN/COS signals at the same time, calculate and adjust the SIN/COS phase/amplitude, compensate for the phase error between excitation and SIN/COS so that it equals zero, and calibrate the SIN/COS amplitude to the same level. Finally, the system will generate the modulated SIN/COS signal and fault signal to simulate the angular performance, velocity, and fault case.

The specific signal chain in Figure 8 shows a dual 16-bit sim SAR ADC AD7380 used to capture the excitation and feedback signal when OSR is enabled and a 98 dB SNR can be reached. It is very suitable for simultaneous high precision data acquisition for phase and amplitude calibration. An ultra low power, low distortion ADA4940-2 is used as an ADC driver. While a high precision, low noise 20-bit DAC AD5791 is used to generate the SIN/COS and

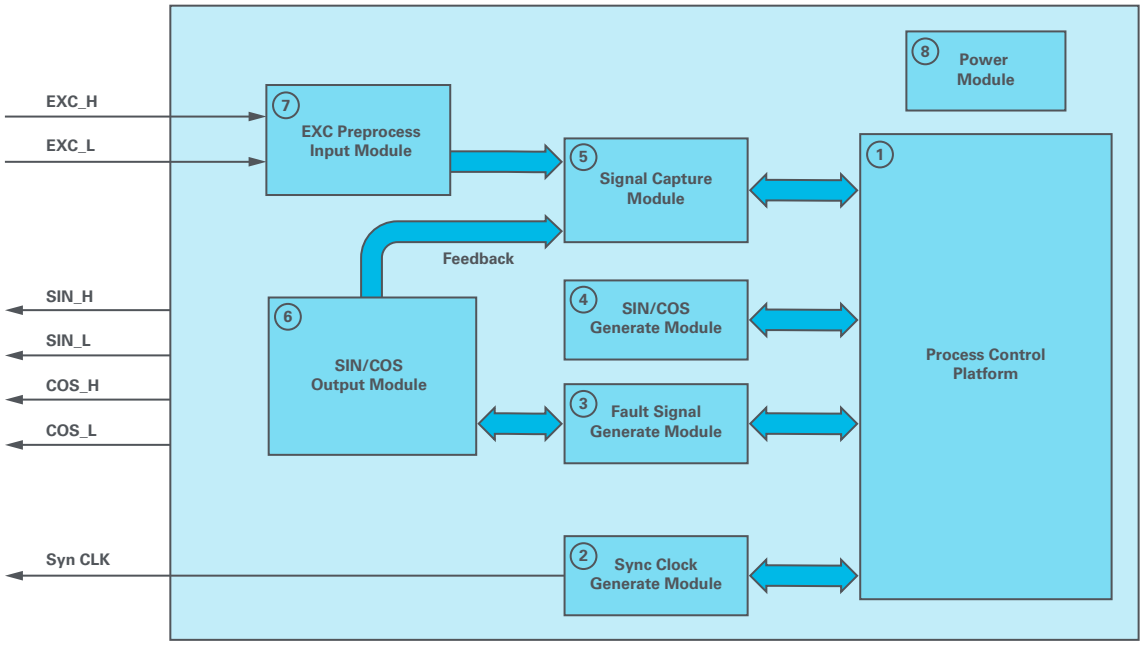

Figure 7. Resolver simulator block diagram.

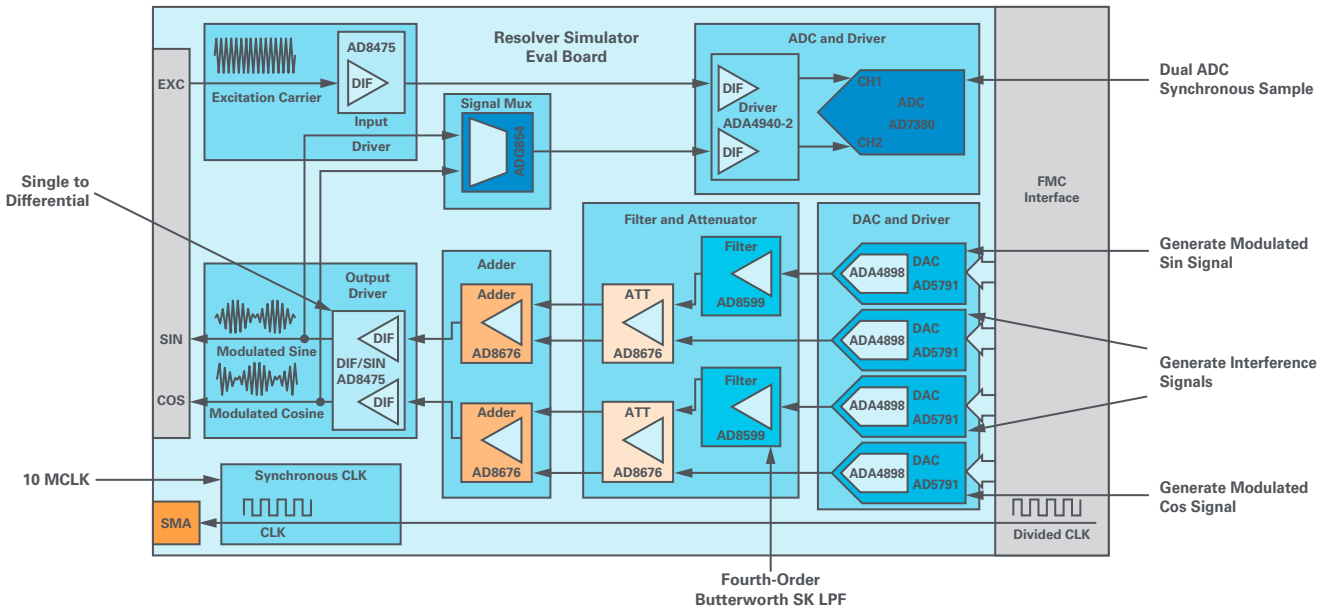

Figure 8. The resolver simulator signal chain.

fault signal for lower resolution and lower cost consideration, the AD5541A or AD5781 is recommended to replace the AD5791. A precision, selectable gain differential amp, the AD8475, is used as the input/output buffer. Precision, rail-to-rail operational amplifiers with ultralow offset drift and voltage noise amps, the AD8676 and AD8599, are used to build an active filter and adder circuit. A single-supply, rail-to-rail, 0.8 Ω max, dual SPDT, the ADG854, is used to switch and choose the SIN/COS signal, which is then sent to the data capture module.

The whole resolver simulator system is powered through an external 12 V adapter with different voltage levels generated by using dc-to-dc converters and LDO regulators. A detailed power supply signal chain is shown in Figure 9. Positive and negative 16 V voltages are generated by an ADP5071, but clearer and more stable positive and negative 15 V voltages can be generated by using the ADP7118 and ADP7182. These power sources are mainly used to supply power for the DAC-related circuit. Similarly, clear and stable +3.3 V, +5 V, –5 V, and –2 V powers are generated by using ADP2300, ADP7118, ADM660, and AD7182. These powers are mainly used to supply power for ADC-related circuits and detailed design requirements.

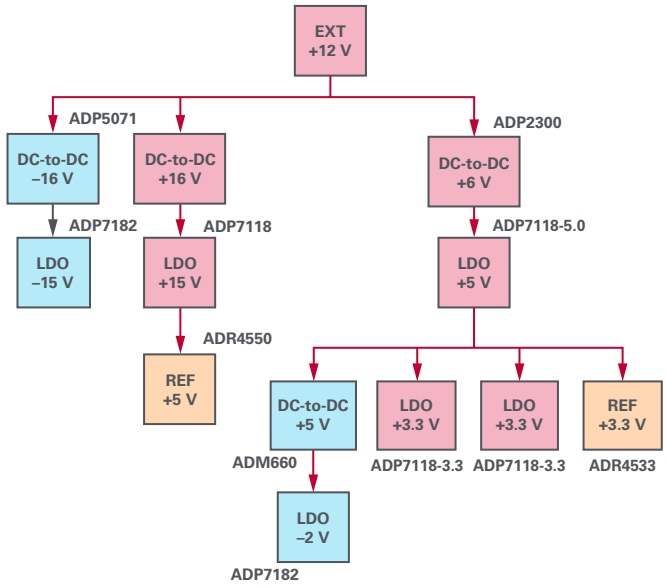

Figure 9. The power supply signal chain.

Resolver Simulator Bench Test and Result

The whole system bench test is shown in Figure 10. It contains a resolver simulator board, an AD2S1210 eval board, and a GUI. The GUI and bench test picture is shown in Figure 11. An AD2S1210 GUI is used to directly evaluate the performance of the resolver simulator, especially velocity and angular performance. Through resolver simulator GUI, velocity, angular performance, and fault signal can be configured.

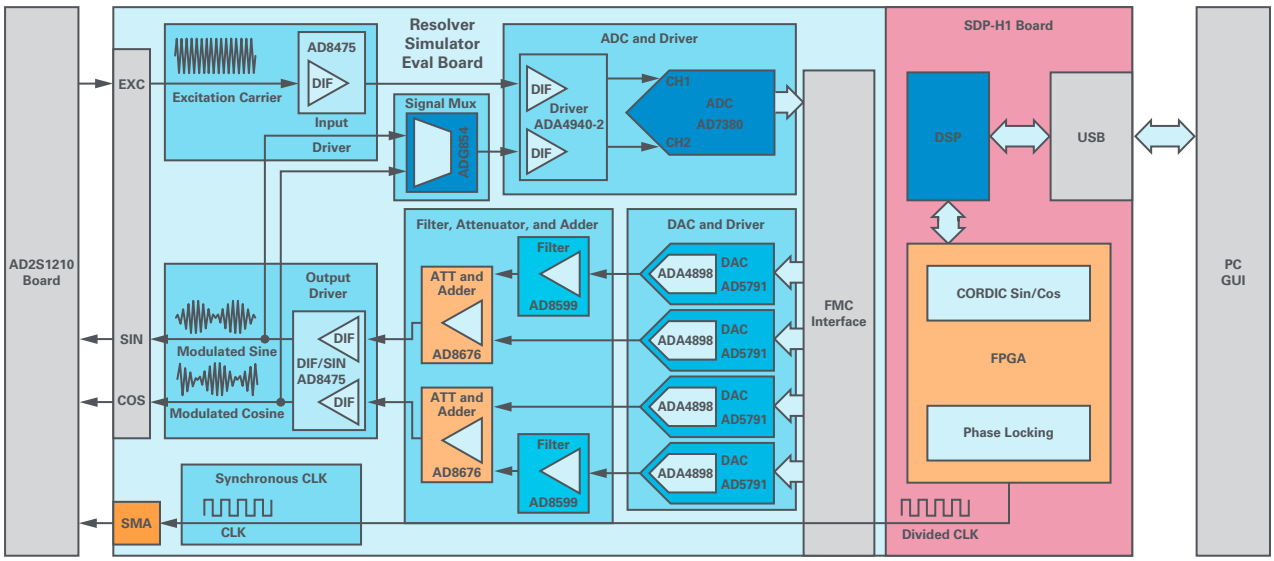

Figure 10. Bench test block diagram.

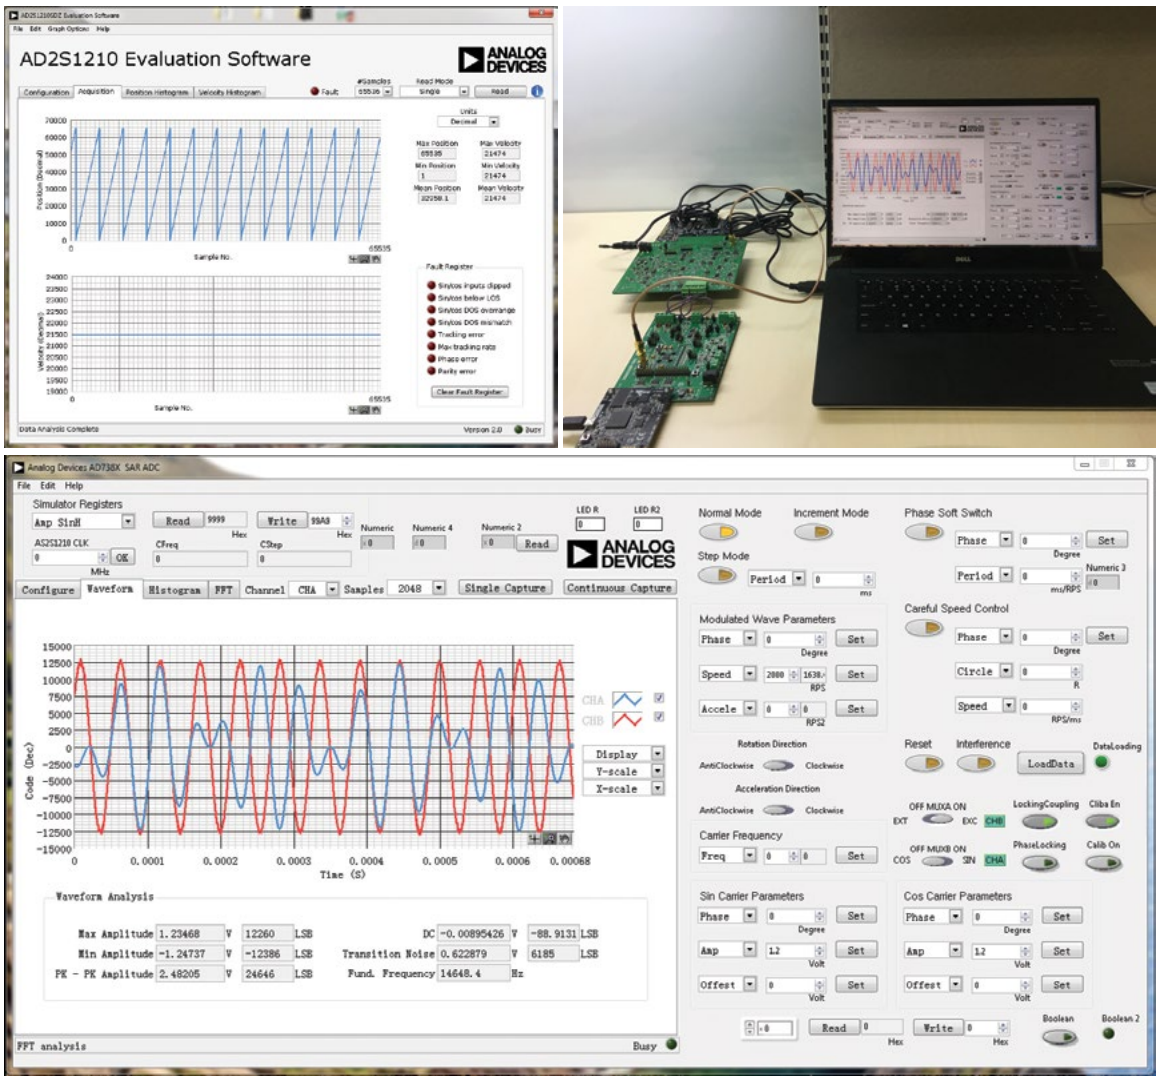

Figure 11. Bench test and GUI.

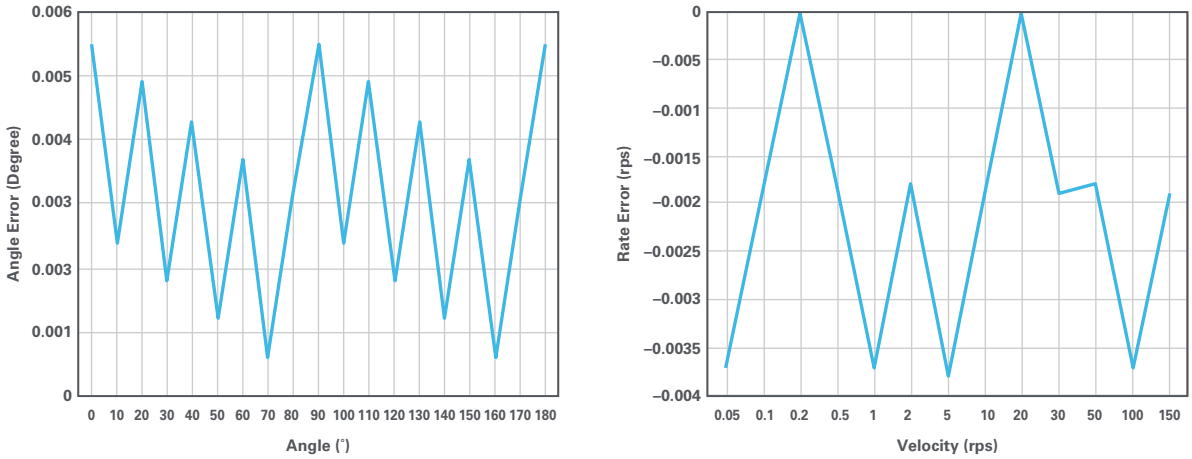

Figure 12. Angular/velocity INL.

The INL of angular and velocity performance of a 16-bit AD2S1210 with hysteresis mode disabled is shown in Figure 12.

Compared with standard resolver simulator devices, the performance of this solution can be seen in Table 2. A 0.006° angular accuracy can be reached in a real bench test—with 0.0004° in theory when using AD5791—as well as 3000 rps maximal velocity output, 0.004 rps velocity accuracy, and can easily meet the 10-bit to approximately 16-bit mode of AD2S1210.

The fault mode supported in this simulator is shown in Table 3. For phase-related fault, a 0° to approximately 360° range can support a SIN/COS signal. For amplitude-related fault, a 0 V to approximately 5 V range can support a SIN/COS signal. Exceed velocity, IGBT, disconnection, and other faults can also be simulated by using this solution.

Table 2. Performance Comparison

Product/Parameter	North Atlantic 5330A	North Atlantic 5300A	This Solution	AD2S1210 Required
Excitation Frequency	47 Hz to 10 kHz	360 Hz to 20 kHz	2 kHz to 20 kHz	2 kHz to 20 kHz
Angle Accuracy	0.003° to ~0.015°	0.00055556° to ~0.0167°	Better than 0.006° @ 12.2070 kHz Carrier	0.0417° and 1 LSB
Rate Range	Up to 277 rps	Up to 278 rps	Up to 3000 rps	Up to 2500 rps @ 8.19 MHz clock
Rate Accuracy	±1%		0.004 rps (<150 rps)	±0.0305 rps (<125 rps)

Table 3. Fault Mode and Supported Range

Fault Mode	Phase Shift	Phase Shift Mismatch	Amplitude Mismatch	Amplitude Exceed	IGBT Disturbance	Random Disturbance	Exceed Velocity	Disconnection
Range	0° to ~360°	0° to ~360°	0 V to ~5 V	0 V to ~5 V	✓	✓	0 rps to ~3000 rps	✓

Figure 13. IGBT interference example.

For IGBT fault, one test example is shown in Figure 13. Configure the simulator output to 45°, then add a periodic interference signal to the SIN/COS output. As the angular and velocity performance of the AD2S1210 evaluation board GUI shows, the angular performance has fluctuation around 45°, while at the same time the velocity will fluctuate around 0 rps.

Conclusion

While interference exists in most RDC-related applications, many types of fault can be triggered under serious conditions. When you build your own resolver simulator, follow this solution as it can help you to not only evaluate your system performance under interference, but also calibrate and verify your products as a standard simulator does. Detailed error analysis is very helpful in understanding why precision analog SIN/COS signals are necessary, and all the fault types discussed in the article can be simulated to help with some functional safety verification.

References

Boyes, Geoffrey. "Synchro and Resolver Conversion." Analog Devices, Inc., 1980.

Hanselman, Duane C. "Resolver Signal Requirements for High Accuracy Resolver-to-Digital Conversion." IEEE Trans.Ind. Electron., Vol. 37, No. 6, Dec. 1990.

Lynch, Michael. "High Precision Voltage Source." Analog Devices, Inc., Oct. 2017.

O'Meara, Shane. AD7380 Evaluation Kits. Analog Devices, Inc., 2019.

Symczak, Jakub, Shane O'Meara, Johnny Gealon, and Christopher Nelson De La Rama "Precision Resolver-to-Digital Converter Measures Angular Position and Velocity." Analog Devices, Inc., Mar. 2014.

Acknowledgements

Many thanks to ADI intern Edward Luo and applications engineers Shane O'Meara, Steven Xie, Karl Wei, and Michael Lynch for their advice and support in the design and test efforts required for this article.

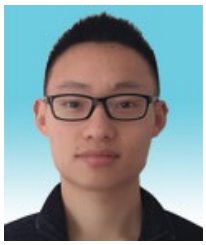

About the Author

Nandin Xu is an applications engineer at Analog Devices in Shanghai, China. He is responsible for technical support for RDC, isolated modulators, and precision DAC products across China. He joined ADI in 2013 after graduating from Huazhong University of Science and Technology in Wuhan, China with a master's degree in control science and control technology. In his spare time, he is a super fan of basketball and soccer. He can be reached at *nandin.xu@analog.com*.

Comprehensive Power Supply System Designs for Harsh Automotive Environments Consume Minimal Space, Preserve Battery Charge, Feature Low EMI

by Bin Wu and Zhongming Ye

Advances in automotive technology have significantly increased the electronic content of modern automobiles to enhance safety, improve the driving experience, enrich entertainment functions, and diversify the power and energy sources. We continue to commit engineering resources to improving power management solutions for the automotive market. Many of the technologies from that effort have resulted in significant advances in power supply efficiency, compactness, robustness, and EMI performance.

Power supplies for automotive applications must perform without failure in the face of harsh conditions—the designer must consider all exigencies, including load dump, cold crank, battery reverse polarity, double battery jump, spikes, and other transients defined in LV 124, ISO 7637-2, ISO 17650-2, and TL82066, as well as mechanical vibration, noise, extremely wide temperature ranges, etc. This article focuses on the critical requirements in automotive power supply specifications and solutions to meeting automotive specifications, including:

- Automotive input transients
- Input voltage range
- Output voltage/current
- Low quiescent current (I_Q)
- Electromagnetic interference (EMI)

Several example solutions are shown to illustrate how combinations of high performance devices can easily solve what would otherwise be difficult automotive power supply problems.

Harsh Automotive Environments

Figure 1 illustrates a complete power solution that meets the demanding requirements of automotive applications. At the front end, the LT8672 acts as an ideal diode, protecting the circuit from brutal conditions under the hood and destructive faults, such as reverse polarity. Following the ideal diode is a family of low quiescent current (I_Q) buck regulators that feature wide input ranges—working down to 3 V and up to 42 V—to deliver regulated voltages for the cores, I/O, DDR, and other rails required by peripheral devices.

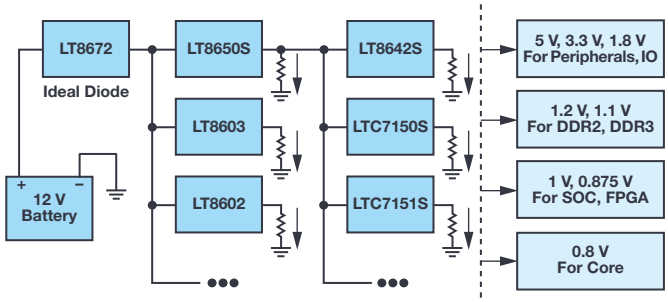

Figure 1. Overview of ADI's Power by Linear solutions for automotive electronics that meet transient immunity requirements.

These regulators feature ultralow quiescent current, extending battery run time for always-on systems. Low noise power conversion technology minimizes the need for costly EMI mitigation, as well as design and test cycles to meet stringent automotive EMI standards. For many critical functions that must ride through cold crank events, the LT8603 multichannel low I_Q buck regulators with a built-in preregulation boost controller delivers a compact solution with at least three regulated voltage rails. The LT8602 can deliver four regulated voltage rails required for many advanced drive assistance system (ADAS) applications, such as collision warning, mitigation, and blind spot monitoring.

Figure 2 shows a traditional automotive electrical system where the engine drives an alternator. The alternator is essentially a 3-phase generator, with its ac output rectified by a full diode bridge. The output of this rectifier is used to recharge a lead-acid battery and power 12 V circuits and devices. Typical loads include the ECU, fuel pump, brakes, fan, air conditioner, sound systems, and lighting. Increasing numbers of ADAS are added to the 12 V bus, including peripherals, I/Os, DDRs, processors, and their power supplies.

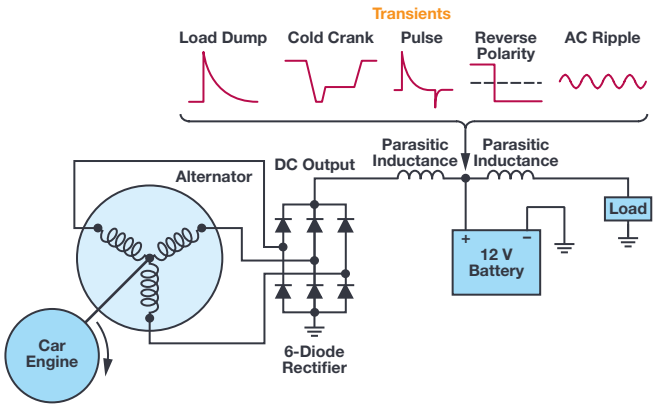

Figure 2. A typical electrical system in a car.

Electric cars change the picture somewhat. The engine is replaced with an electric motor, where a dc-to-dc converter converts a 400 V high voltage lithium-ion (Li-Ion) battery stack to 12 V, instead of an alternator. Nevertheless, traditional 12 V alternator devices are here to stay, along with their transient pulses—including fast pulses.

An engine runs at its peak efficiency in a narrow range of rpms, so the steady state output of the alternator and the battery voltage are relatively stable, say ~13.8 V, under most conditions (more about that below). Every circuit powered directly from the car battery must run reliably over the range of 9 V to 16 V, but robust automobile electronic designs must also operate during outlier conditions that will inevitably occur at the most inconvenient time.

Although output of the alternator is nominally stable, it is not stable enough to avoid the need for conditioning before it powers the vehicle's other systems. Unwanted voltage spikes or transients are harmful to downstream electronic systems and, if not properly addressed, can cause these systems to malfunction or cause permanent damage. In the past few decades, many automotive standards such as ISO 7637-2, ISO 16750-2, LV 124, TL82066 have been produced to define the spikes and voltage transients that automotive power supplies will face, and set design expectations.

One of the most critical and challenging high voltage transients is load dump. In automotive electronics, load dump refers to the disconnection of the vehicle battery from the alternator while the battery is being charged. During a load dump transient, the excitation field of the alternator remains high given its large time constant—the alternator still outputs high power even without the load. A battery is a big capacitor and will normally absorb the extra energy, but when it is disconnected due to a loose terminal or other issues, it can no longer provide this service. As a result, all the other electronics see the voltage surge and must be able to survive load dump events. An unsuppressed load dump could generate voltages upward of 100 V. Thankfully, modern car alternators use avalanche-rated rectifier diodes, limiting the load dump voltage to 35 V—still a significant diversion from the norm. A load dump event can last up to 400 ms.

Another high voltage event is jump-start. Some tow trucks use two batteries in series to assure effective jump starts to revive a dead car battery, so an automobile's circuits must survive the doubled nominal battery voltage of 28 V for a couple of minutes. Many Power by Linear high voltage step-down regulators, such as the Silent Switcher and Silent Switcher 2 families, including the LT8650S and LT8640S (Table 1) operate up to 42 V, exceeding this requirement. In contrast, lower voltage rated options would require a clamp circuit, adding cost and lowering efficiency. Some Power by Linear regulators, such as the LT8645S and LT8646S, are rated for 65 V to accommodate truck and airplane applications, where a 24 V system is the norm.

Table 1. Silent Switcher and Silent Switcher 2 Monolithic Buck Regulators for Automotive Applications

Device	Number of Outputs	V_{IN} Range (V)	Output Current	Peak Efficiency f_{SW} = 2 MHz V_{IN} = 12 V V_{OUT} = 5 V	I_Q at 12 V Input (TYP) (µA)	EMI Feature	Packages
LT8650S	2	3 to 42	4 A on both channels 6 A on either channel	94.60%	6.2	Silent Switcher 2	6 mm × 4 mm × 0.94 mm LQFN
LT8645S	1	3.4 to 65	8 A	94%	2.5	Silent Switcher 2	6 mm × 4 mm × 0.94 mm LQFN
LT8643S	1	3.4 to 42	6 A continuous 7 A peak	95%	2.5	Silent Switcher 2 external compensation	4 mm × 4 mm × 0.94 mm LQFN
LT8640S	1	3.4 to 42	6 A continuous 7 A peak	95%	2.5	Silent Switcher 2	4 mm × 4 mm × 0.94 mm LQFN
LT8609S	1	3 to 42	2 A continuous, 3 A peak	93%	2.5	Silent Switcher 2	3 mm × 3 mm × 0.94 mm LQFN
LT8641	1	3 to 65	3.5 A continuous, 5 A peak	94%	2.5	Silent Switcher	3 mm × 4 mm, 18-Lead QFN
LT8640 LT8640–1	1	3.4 to 42	5 A continuous 7 A peak	95%	2.5	Silent SwitcherLT8640: pulse skipping LT8640–1: forced continuous	3 mm × 4 mm, 18-Lead QFN
LT8614	1	3.4 to 42	4 A	94%	2.5	Silent Switcher low ripple Burst Mode operation	3 mm × 4 mm, 18-Lead QFN
LT8642S	1	2.8 to 18	10 A	95%	240	Silent Switcher 2	4 mm × 4 mm × 0.94 mm LQFN
LT8646S	1	3.4 to 65	8 A	94%	2.5	Silent Switcher 2	6 mm × 4 mm × 0.94 mm LQFN

Another voltage transient occurs when a driver starts an automobile and the starter draws hundreds of amperes of current from the battery. This pulls down the battery voltage for a short period of time. In a traditional automobile, this happens only when the car is started by the driver—for instance, when one starts a car to drive to the supermarket and starts it again to drive back home. In modern automobiles with start-stop features to save fuel, start-stop events can occur a number of times on that supermarket trip—at every stop sign and every red light. The additional start-stop events put significantly more strain on the battery and starter than in a traditional automobile.

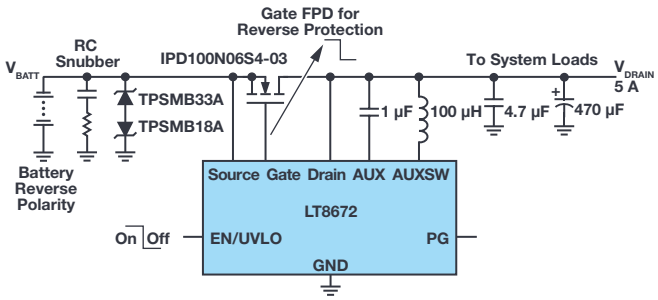

Figure 3. LT8672 response to battery reverse polarity.

Furthermore, if a start event happens on a cold morning, the starter draws more current than at higher ambient temperatures, pulling the battery down to 3.2 V or lower for around 20 ms—this is called cold crank. There are functions that must remain active even in cold crank conditions. The good thing is that, by design, such critical functions typically do not require significant power. Integrated solutions, such as the LT8603 multiple channel converter, can maintain regulation even if their inputs drop below 3 V.

ISO 7637-2 and TL82066 define many other pulses. Some have higher positive or negative voltages but also higher source impedances. Those pulses have relatively low energy compared to the events described above, and can be filtered or clamped with proper selection of input TVS.

An Ideal Diode Satisfies Automotive Immunity Norms

The active rectifier controller LT8672, featuring high input voltage rating (+42 V, –40 V), low quiescent current, ultrafast transient response speed, and ultralow external FET voltage drop control, provides protection in 12 V automotive systems with extremely low power dissipation.

Battery Reverse Polarity

Whenever the battery terminals are disconnected, there is a chance the car battery polarity is reversed by mistake and the electronic systems can be damaged from the negative battery voltage. Blocking diodes are commonly placed in series with supply inputs to protect against supply reversal, but blocking diodes feature a voltage drop, resulting in an inefficient system and reducing the input voltage, especially during a cold crank.

The LT8672 is an ideal diode replacement to the passive diode to protect the downstream systems from the negative voltages, as shown in Figure 3.

Under normal conditions, the LT8672 controls an external N-channel MOSFET to form an ideal diode. The GATE amplifier senses across DRAIN and SOURCE and drives the gate of the MOSFET to regulate the forward voltage to 20 mV. D1 protects SOURCE in the positive direction during load steps and overvoltage conditions. When a negative voltage appears in the input side, GATE is pulled to SOURCE when SOURCE goes negative, turning off the MOSFET and isolating DRAIN from the negative input. With its fast pull-down (FPD) capability, LT8672 can quickly turn off the external MOSFET.

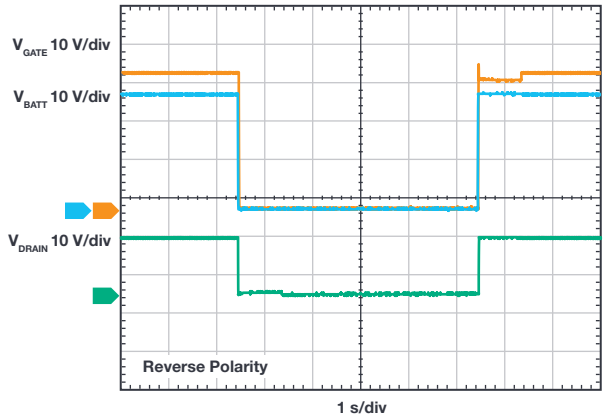

Figure 4. Waveform of LT8672 response to reverse polarity.

Superimposed Alternating Voltage

A common disturbance on the battery rail is a superimposed ac voltage. This ac component can be an artifact of the rectified alternator output or a result of frequent switching of high current loads, such as motors, bulbs, or PWM controlled loads. According to automotive specifications ISO 16750 and LV 124, an ECU may be subjected to an ac ripple superimposed on its supply, with frequencies up to 30 kHz and amplitudes of up to 6 V p–p. In Figure 5, a high frequency ac ripple is superimposed on the battery line voltage. Typical ideal diode controllers are too slow to react, but the LT8672 generates high frequency gate pulses up to 100 kHz to control external FETs as needed to reject these ac ripples.

The unique ability of the LT8672 to reject common ac components on a power rail are a function of its fast pull-up (FPU) and FPD control strategy and its strong gate driving capability, where the gate driver is powered by an integrated boost regulator. Compared with a charge pump gate power solution, this boost regulator enables the LT8672 to maintain a regulated 11 V voltage to keep the external FET on, while providing strong gate souring current to reduce switching loss for high frequency ac ripple rectification. Its 50 mA source current capability enables super-fast turn-on of the FET, minimizing power dissipation; its 300 mA sinking current capacity realizes fast turn-off, minimizing the reverse current

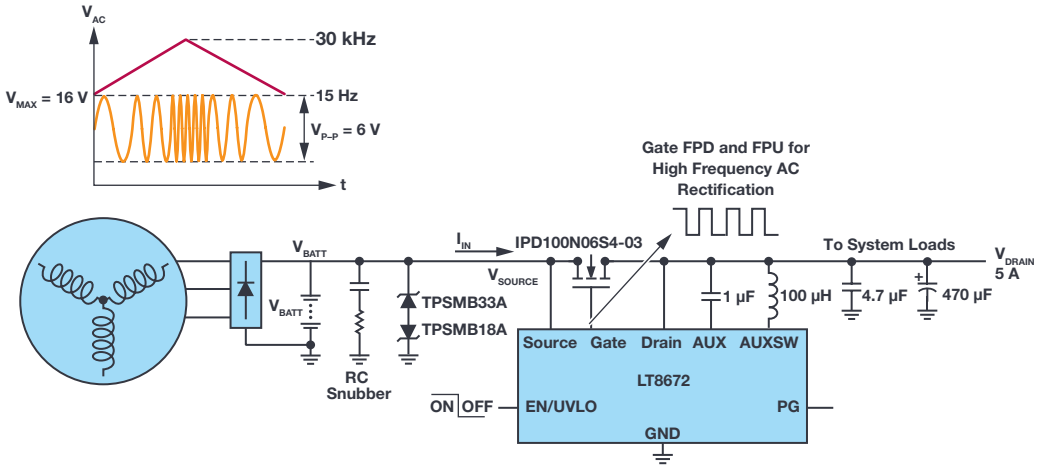

Figure 5. Waveform of LT8672 response to superimposed alternating voltage.

conduction. In addition, this significantly reduces the ripple current in the output capacitor. Typical rectification waveforms for a superimposed alternating voltage are shown in Figure 6.

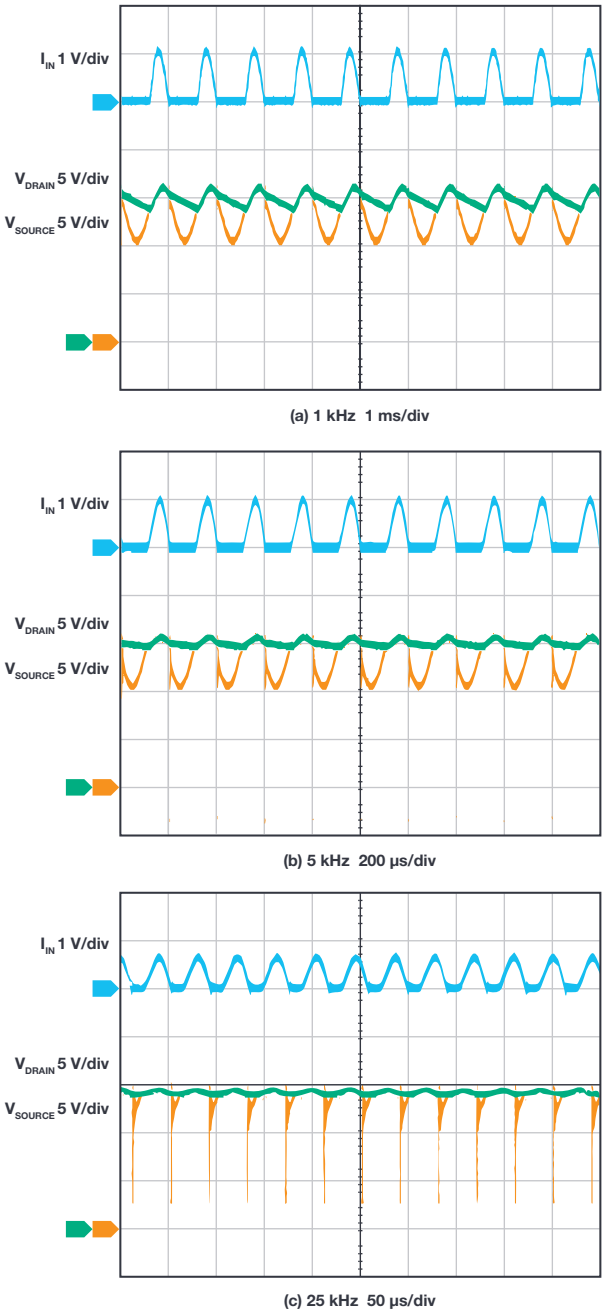

Figure 6. Waveform of LT8672 response to superimposed alternating voltage.

In addition, the LT8672 effectively reduces conduction losses when compared with a traditional Schottky diode solution under the same load conditions. As seen in the thermal images of Figure 7, the solution using the LT8672 is almost 60°C cooler than a traditional diode-based solution. It not only improves the efficiency, but also eliminates the need for a bulky heat sink.

High peak, narrow pulses that appear on input of automotive electronic systems usually come from two sources:

- The disconnection of input power supply when there is inductive load in series or parallel.
- The switching processes of a load influencing the distributed capacitance and inductance of a wire harness.

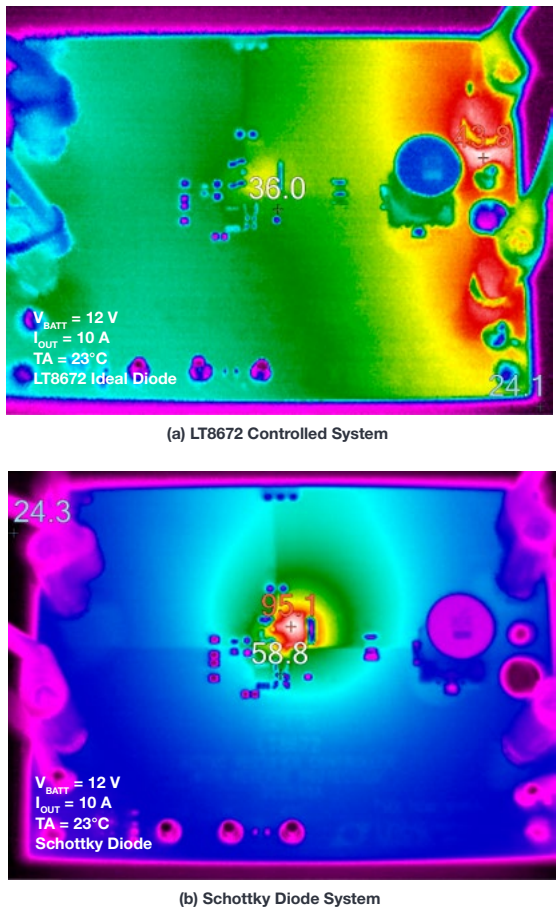

Figure 7. Thermal performance comparison.

Some of these pulses could have high voltage peaks. For example, pulse 3a defined in ISO 7632-2 is a negative spike whose peak voltage exceeds –220 V, while pulse 3b defines a pulse with maximum peak voltage of 150 V, on top of the battery's initial voltage. Although they feature a large internal impedance and very narrow duration time, downstream electronics could be easily damaged if they see these pulses.

Two properly sized TVSs are installed in the front end to suppress such spikes. In fact, some of the low energy pulses could be absorbed directly by filter effect of input capacitor and parasitic wire inductance.

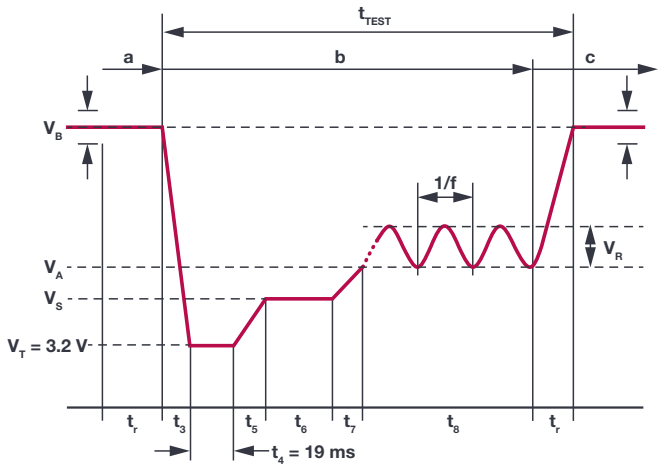

Figure 8. Severe cold crank for the 12 V system defined in LV 124.

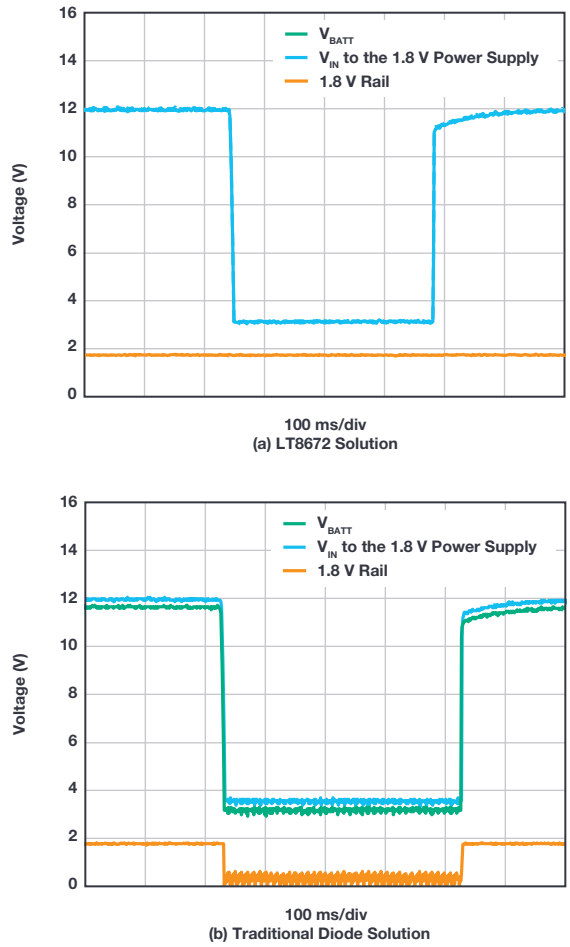

Figure 9. Cold crank event.

Multiple Rail Regulator Rides Through Cold Crank Events

The LT8602 provides compact solutions for up to four regulated rails (for example, 5 V, 3.3 V, 1.8 V, 1.2 V) with an input voltage range from 5 V to 42 V, suitable for functions that do not necessarily need to be on during a cold crank. Otherwise, for functions that must operate even during cold crank— such as the spark plug controller or alarm—solutions like the LT8603 work down to 3 V (or lower) inputs.

LV 124 has defined the worst case of cold crank, shown in Figure 8. It indicates that the lowest battery voltage can go down to 3.2 V and last for 19 ms at car startup. This specification challenges applications to keep running as low as 2.5 V when faced with the extra diode voltage drop from battery reverse protection in a traditional (nonideal diode) solution. In a passive diode protection scheme, buck-boost regulators may be required instead of less complex and more efficient buck regulators to provide a stable 3 V supply often required by many microcontrollers.

The LT8672 controller features a minimum input operating voltage of 3 V V_{BATT}, enabling the active rectifier to operate through the cold crank pulse with a minimum drop (20 mV) between input and output. Downstream power supplies during a cold crank event see an input voltage no lower than 3 V. This allows use of a buck regulator with a minimum operating voltage of 3 V and low dropout characteristics, such as the LT8650S, to generate a 3 V supply.

Like the LT8650S, many ADI Power by Linear automotive ICs feature minimum input voltage rating of 3 V.

Figure 9 shows the comparison of 1.8 V power supply with the LT8672 and with a traditional diode. The step-down regulator works down to 3 V. As shown, with a traditional diode, V_{IN} to the buck regulator drops to near 2.7 V when the battery voltage V_{BATT} drops to 3.2 V, due to high voltage drop of the diode, triggering the UVLO shutdown of the downstream switching regulator, and its 1.8 V output collapses. In contrast, voltage remains nearly constant at the LT8672 output during a cold crank event, and the downstream step-down regulator is able to maintain a 1.8 V output.

Numerous critical functions require regulated 5 V and 3.3 V rails, plus sub-2 V rails to power content, processor I/O, and core in analog and digital ICs. If V_{BATT} drops below its outputs or V_{IN} (MIN), a pure buck regulator would lose regulation if directly powered from V_{BATT}. However, such critical functions typically do not require much power, so a highly integrated compact solution can be used, such as the 6 mm × 6 mm LT8603 quad output, triple monolithic buck converter plus boost controller.

The LT8603's integrated boost controller works down to below 2 V, making it an ideal preregulator to its three buck regulators. Figure 10 shows a Power by Linear state-of-the-art solution for these applications that can ride through a cold crank event. The two high voltage buck regulators are powered by the preboost converter. When V_{BATT} drops below 8.5 V, the boost controller starts switching and the output (OUT4) is regulated to 8 V. It can keep the output regulated with the input voltage down to 3 V once it is started. Therefore, the two high voltage bucks can ride through the cold crank condition, while providing constant 5 V and 3.3 V outputs, as shown in Figure 11. Once V_{BATT} recovers to above 8.5 V from cold crank, the boost controller simply works as a diode pass through. The high voltage bucks can handle V_{BATT} up to 42 V. The low voltage buck is powered from OUT2, providing 1.2 V through the cold crank event.

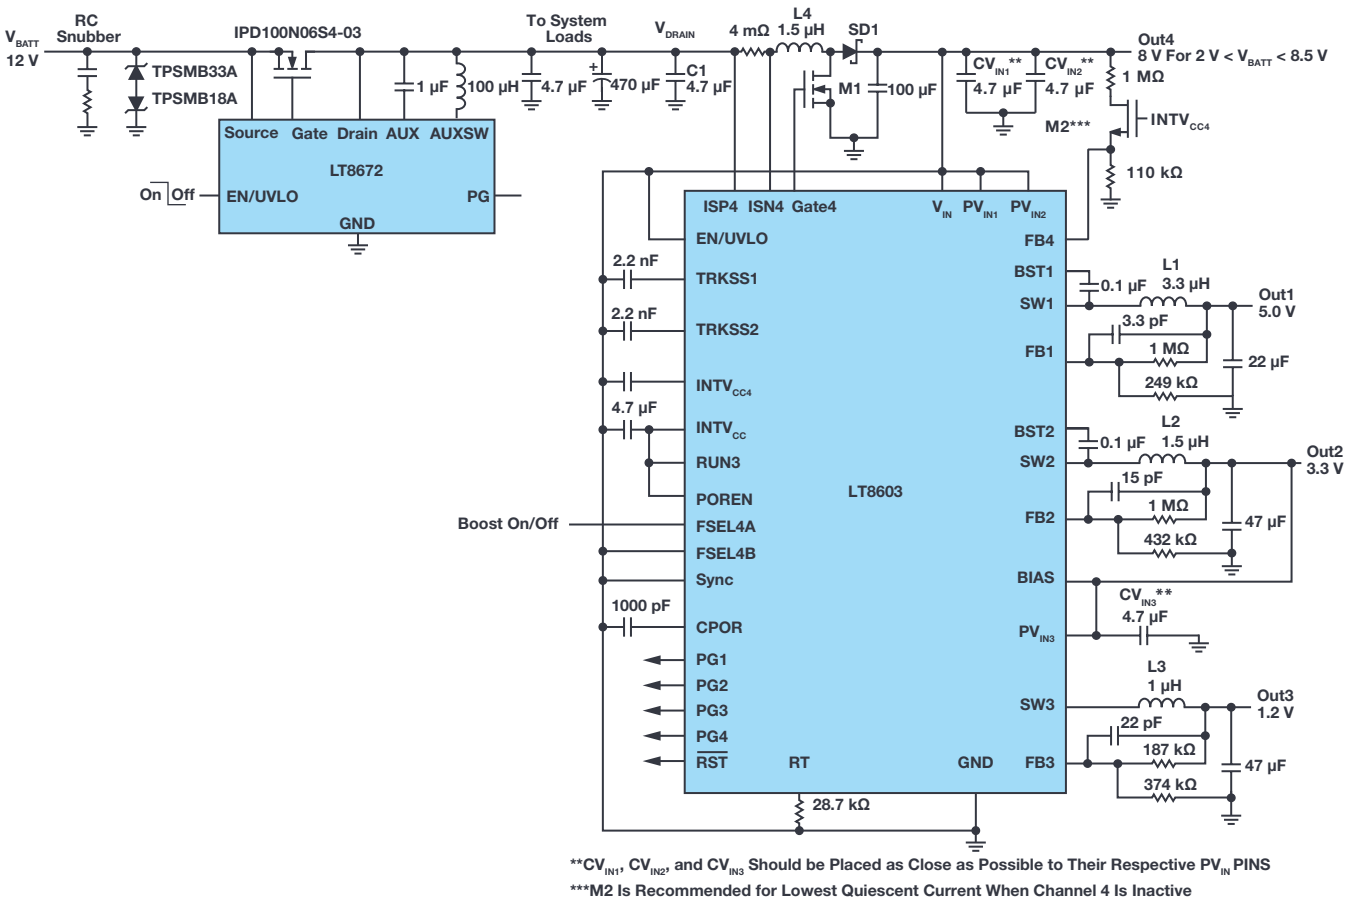

Figure 10. LT8672 and LT8603 solution tolerates cold crank events.

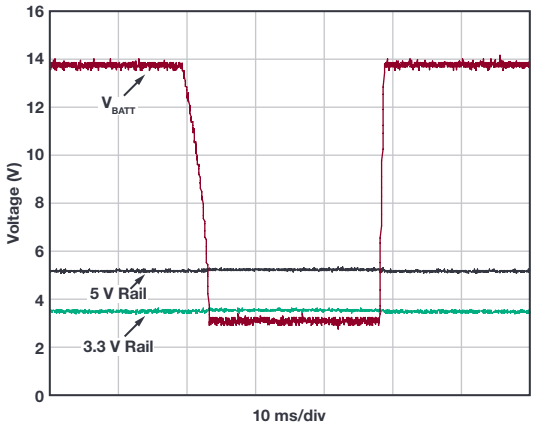

Figure 11. The LT8672 and LT8603 combination produces 5 V and 3.3 V outputs that ride through cold crank events.

Ultralow I_Q Extends Battery Run Time for Always-On Systems

For always-on systems connected to V_{BATT} for weeks or months without a battery recharge, light load and no-load efficiency are, in some cases, more important than full load efficiency. The Power by Linear family of ultralow quiescent current (I_Q) devices preserve battery charge while withstanding challenging transient conditions and wide input voltage ranges, from 3 V to 42 V, and wide temperature ranges. To optimize efficiency and maintain regulation at light loads and no load, the regulator features Burst Mode® operation. Between bursts, all circuitry associated with controlling the output switch is shut down, reducing the input supply current to a few microamps. In contrast, a typical buck regulator might draw hundreds of hundreds of microamps from V_{BATT} when regulating with no load, draining the battery orders of magnitude faster.

The Burst Mode efficiency at a given light load is mainly affected by the switching loss, which is a function of switching frequency and gate voltage. Because a fixed amount of energy is required to switch the MOSFET on and off, and keep the internal logic alive, a lower switching frequency reduces gate charge losses and increases efficiency. The switching frequency is primarily determined by the Burst Mode current limit, the inductor value, and the output capacitor. For a given load current, increasing the burst current limit allows more energy to be delivered during each switching cycle, and the corresponding switching frequency is lower. For a given burst current limit, a larger value inductor stores more energy than a smaller one, and the switching frequency is lower as well. For the same reason, a bigger output capacitor stores more energy and takes longer to discharge.

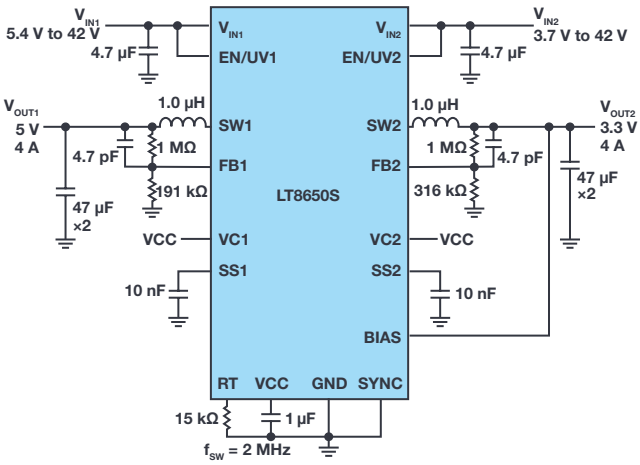

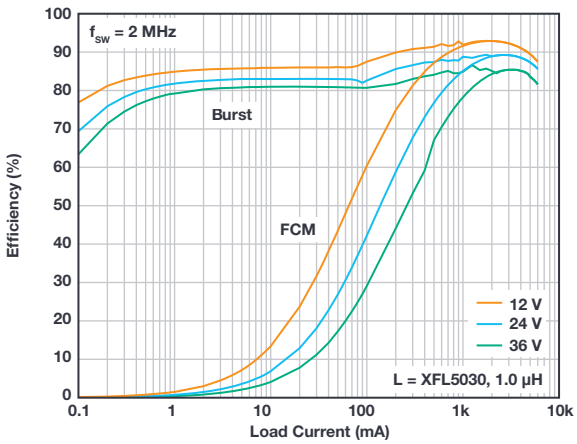

Figure 12. Low I_Q LT8650S maintains very high light load efficiency to support always-on applications without significantly draining the battery.

Figure 12 shows the ultralow I_Q synchronous buck regulator LT8650S in a solution that features high efficiency over wide input voltage and load current ranges. With integrated MOSFETs, this device can deliver up to 8 A total output current at fixed output voltages of 3.3 V or 5 V. Despite the simple overall design and layout, this converter includes options that can be used to optimize the performance of specific applications in battery-powered systems.

Table 1 lists low I_Q monolithic regulators that are well-suited to the automotive market, with inputs up to 42 V or 65 V. Typical quiescent current for these devices is only 2.5 µA, thanks to the low I_Q technologies developed by Analog Devices. With minimum turn-on time of 35 ns, these regulators deliver 3.3 V output voltage from input 42 V at switching frequency of 2 MHz, which is common in the automotive industry.

Silent Switcher Portfolio Takes Complexity out of EMI Design

Automotive applications demand systems that do not produce electromagnetic noise that could interfere with the normal operations of other automotive systems. For instance, switching power supplies are efficient power converters, but by nature generate potentially unwelcome high frequency signals that could affect other systems. Switching regulator noise occurs at the switching frequency and its harmonics.

Ripple is a noise component that appears at the output and input capacitors. Ripple can be reduced with the low ESR and ESL capacitors, and low-pass LC filters. A higher frequency noise component, which is much more difficult to tackle, results from the fast switching on and off of the power MOSFETs. With designs focused on compact solution size and high efficiency, operating switching frequencies are now pushed to 2 MHz to reduce the passive component size and avoid the audible band. Furthermore, switching transition times have been reduced to the nanosecond realm to improve efficiency—by reducing switching losses and duty ratio losses.

Parasitic capacitance and inductance from both the package and PCB layout play important roles in distributing noise, so if the noise is present, it can be difficult to eliminate. EMI prevention is complicated by the fact that switching noise covers the domain from tens of MHz to beyond GHz. Sensors and other instruments subjected to such noise could malfunction, resulting in audible noise or serious system failure. Therefore, stringent standards have been set up to regulate EMI. The most commonly adopted one is the CISPR 25 Class 5, which details acceptable limits at frequencies from 150 kHz to 1 GHz.

Passing automotive EMI regulation at high current usually means a complicated design and test procedure, including numerous trade-offs in solution footprint, total efficiency, reliability, and complexity. Traditional approaches to controlling EMI by slowing down switching edges or lowering switching frequency come with trade-offs such as reduced efficiency, increased minimum on-and off-times, and larger solution size. Alternative mitigation, including a complicated bulky EMI filter, snubber, or metal shielding, adds significant costs in board space, components, and assembly, while complicating thermal management and testing.

Our Silent Switcher technology addresses the EMI issue in an innovative way, enabling impressive EMI performance in high frequency, high power supplies. Second-generation, Silent Switcher 2 devices simplify board design and manufacture by incorporating the hot loop capacitors into the packaging. For a buck regulator such as the 42 V/4 A LT8650S, the hot loop consists of an input capacitor and the top and bottom switches. Other noisy loops include the gate drive circuit and boost capacitor charge circuit. In Silent Switcher 2 devices, the hot loop and warm loop capacitors are integrated into the packaging and laid out to minimize EMI. This reduces the effect of final board layout on the EMI equation, simplifying design and manufacturing. Further peak EMI reduction can be achieved by using the optional spread spectrum frequency modulation feature incorporated into these parts, making it even easier to pass stringent EMI standards.

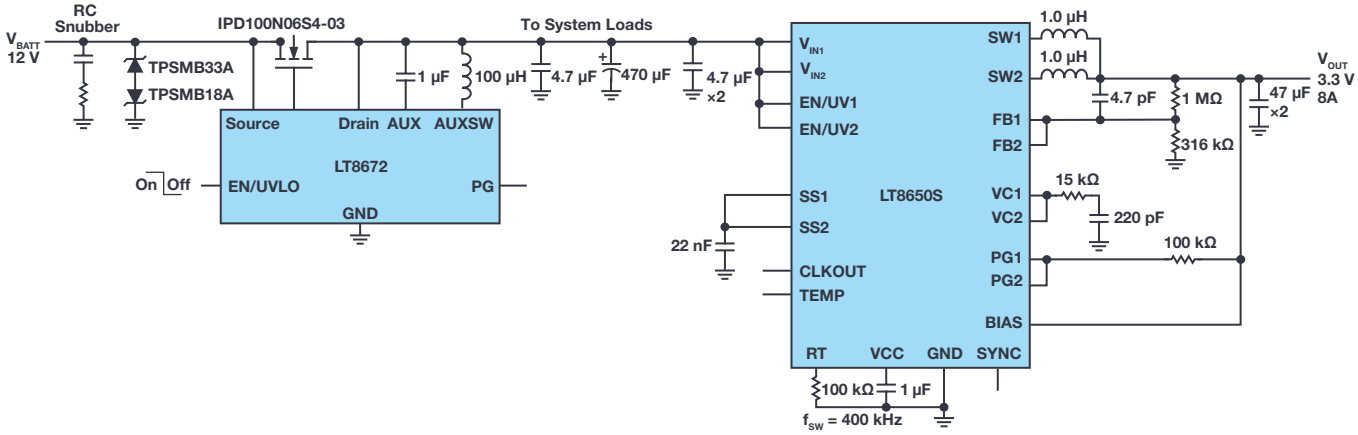

Figure 13. LT8672 and LT8650S configuration for high output current.

Figure 13 exhibits a low I_Q, low noise solution for a high current application for automotive I/Os and peripherals. The LT8672 at the front end protects the circuit from reverse battery faults and high frequency ac ripple with only tens of mV of forward voltage drop. The LT8650S switches at 400 kHz with input ranging from 3 V to 40 V, and an output capability of 8 A by operating two channels in parallel. Two decoupling capacitors are placed close to the input pins of the LT8650S. With Silent Switcher 2 technology, the high frequency EMI performance is excellent even without an EMI filter installed. The system passes the CISPR 25 Class 5 peak and average limit with significant margins. Figure 14 shows the radiated EMI average test results over the range of 30 MHz to 1 GHz, with vertical polarization. A complete solution features a simple schematic, minimal overall component count, compact footprint, and EMI performance that is immune to changes in board layout (Figure 15).

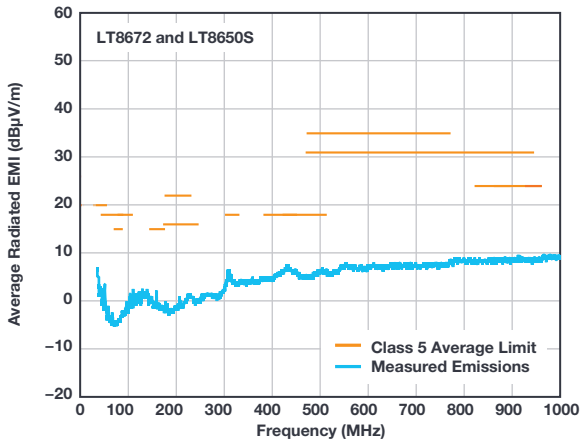

Figure 14. LT8672 and LT8650S EMI performance: 30 MHz to 1 GHz.

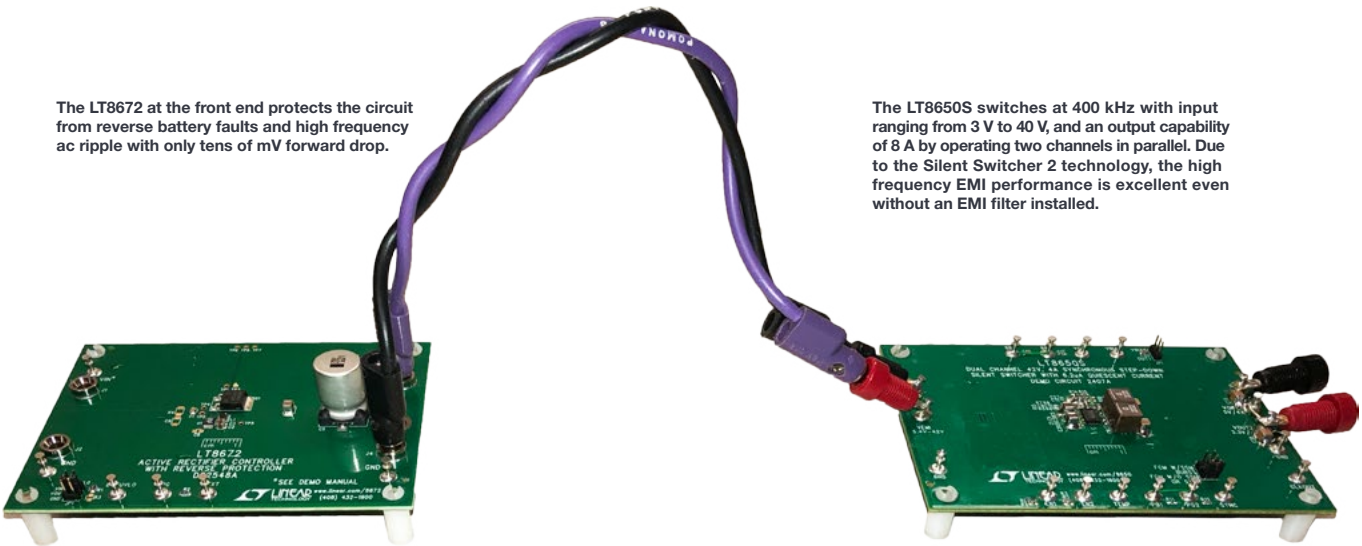

Figure 15. A complete power supply solution for 3.3 V and 5 V outputs from an automotive battery.

Conclusion

Automotive applications call for low cost, high performance, reliable power solutions. The cruel under-the-hood environment challenges power supply designers to produce robust solutions, taking into account a wide variety of potentially destructive electrical and thermal events. Electronic boards connected to the 12 V battery must be carefully designed for high reliability, compact solution size, and high performance. The Power by Linear device catalog includes innovative solutions specifically addressing automotive requirements: ultralow quiescent current, ultralow noise, low EMI, high efficiency, wide operating ranges in compact dimensions, and wide temperature range. By eliminating complexity while improving performance, Power by Linear solutions reduce power supply design time, lower solution costs, and improve time to market.

About the Author

Bin Wu (S'14) was born in Zhejiang, China, 1985. He received his Ph.D. degree in electrical engineering from University of California, Irvine, California in April 2016. From April 2016 to July 2017, he was a post-doctoral research associate in University of Maryland, College Park. After that, he worked at Maxim Integrated, Inc. Since November 2017, he has been an applications engineer with Analog Devices, San Jose. His interests include electrical vehicle power architecture, high power density step-up/step-down dc-to-dc converters, switched capacitor converters, modeling, and renewable energy integration systems. He can be reached at *bin.wu@analog.com*.

About the Author

Zhongming Ye is a senior applications engineer for power products at Analog Devices in Milpitas, California. He has been working at Linear Technology (now part of Analog Devices) since 2009 to provide application support to various products including buck, boost, flyback, and forward converters. His interests in power management include high performance power converters and regulators of high efficiency, high power density, and low EMI for automotive, medical, and industrial applications. Prior to joining Linear Technology, he worked at Intersil for three years on PWM controllers for isolated power products. He obtained a Ph.D. degree in electrical engineering from Queen's University, Kingston, Canada. Zhongming was a senior member of IEEE Power Electronics Society. He can be reached at *zhongming.ye@analog.com*.

Rarely Asked Questions—Issue 168 Bootstrapping a Low Voltage Op Amp to Operate with High Voltage Signals and Supplies

Barry Harvey

Question:

Can I bootstrap a low voltage amplifier to get a high voltage buffer?

Answer:

You can take an op amp with rare input characteristics and elevate it to achieve higher voltage range, better gain accuracy, higher slew rate, and less distortion than the original op amp.

I was designing the input of a precision voltmeter and needed a sub-picoampere input unity-gain amplifier/buffer with less than 1 µV p-p low frequency noise, a low offset voltage of approximately 100 µV, and a nonlinearity of <1 ppm. It also needed to have very low ac distortion over audio frequencies and 60 Hz to make use of ever-deepening ADC resolution. That's ambitious enough, but it must buffer ±40 V signals using ±50 V supplies. The buffer input would be connected to either a high impedance divider or directly to external signals. Thus, it must also tolerate electrostatic discharges and inputs beyond the supplies.

There aren't many sub-picoampere bias current op amps available. Those that are available are often called electrometer grade and offer low tens of femtoampere bias current. Those electrometer amplifiers, unfortunately, have a low frequency noise (0.1 Hz to 10 Hz) of several microvolts peak-to-peak. They also generally have input offset voltage and offset tempco that don't meet requirements. Their common-mode rejection ratio (CMRR) and open-loop gain are not good enough to support 1 ppm linearity. Finally, none of the electrometers can tolerate high supply voltages.

The LTC6240 family offers 0.25 pA typical bias current and 0.55 µV p-p low frequency noise. That's good enough for the input buffer except that the part only works on supplies up to 12 V maximum. We will have to add circuitry around the amplifier to adapt it to higher voltages.

Design Approach

Figure 1 shows a simplified schematic of a bootstrapped amplifier.

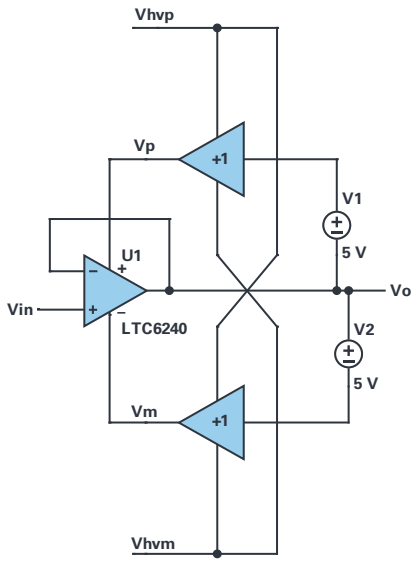

Figure 1. Basic bootstrapped supply circuit topology.

The LTC6240 is powered by Vp, which follows the output plus 5 V through a gain of +1 buffer amplifier, and by Vm, which follows 5 V below the output driven by another buffer.

Because the supplies always follow the input signal, as buffered by the output of the LTC6240, there is no common-mode input error at all, ideally.

Even a mediocre CMRR is bootstrapped up by at least 30 dB. That 30 dB value is due to the finite gain accuracy of the Vp and Vm buffers.

The open-loop gain of the LTC6240 is similarly boosted. Gain limitations arise in amplifier circuits when transistor output impedances exist between an internal gain node(s) and a power supply rail. Since the supplies are bootstrapped to the output, little signal current flows through said impedances, and open-loop gain is raised by amounts like the CMRR benefit. However, loading of the output can still limit open-loop gain.

Less obviously perhaps, the overall circuit slew rate is also raised by bootstrapping. Normally, it is limited by internal LTC6240 quiescent currents and compensation capacitors referenced to supplies. When the supplies follow the input and output, little dynamic current flows into these capacitors and the amplifier does not enter limited slew rate. The buffer amplifiers will ultimately limit overall slew rate.

The high voltage supplies Vhvp and Vhvm may have disturbances, but the buffer outputs will largely reject them and the power supply rejection ratio (PSRR) of the LTC6240 will be greatly enhanced.

So, this is great; the buffer is improved in several ways by bootstrapping the supplies. What could go wrong? Well, the circuit shown in Figure 1 will almost certainly oscillate. The best way to think of the supply terminals' behavior is as part of a feedback loop: the output terminal voltage times the buffer amplifier frequency response, then times 1/PSRR is added to the input, finally multiplied by the open-loop gain to become the output, and 'round the loop evermore. Figure 2a shows the PSRR over frequency.

We don't get phase data in our PSRR plot, but let's say it has a +90° phase. Yes, that's +90° like a differentiator. The open-loop gain, seen in Figure 2b, has a –90° phase from low frequencies to 100 kHz, after which it becomes increasingly negative. The buffers will have finite frequency response and they will exhibit phase lag as well. Adding up all the phase lags in the loop guarantees a few frequencies wherein the feedback phase is 0° or multiples of 360°. If the supply loop gain is >1 at such phases, we have an oscillator. The PSRR magnitude drops to a low of 4 dB (that's attenuation = –4 dB→ gain = 0.63 in non-dB) so it appears that the loop might never have enough gain to oscillate. That's probably wrong, since the PSRR applies to both Vp and Vs, and their PSRR gains may well add up to a magnitude more than one. Further, the buffers could have some peaking before their gain rolls off at high frequency, pushing the overall feedback magnitude of >1. We will also see that the buffers must drive moderately large capacitors and will have more phase lag. In any event, simulating the circuit in LTspice® showed large signal oscillations (the frequency response and nonlinearities of the LTC6240 are embodied in the macromodel).

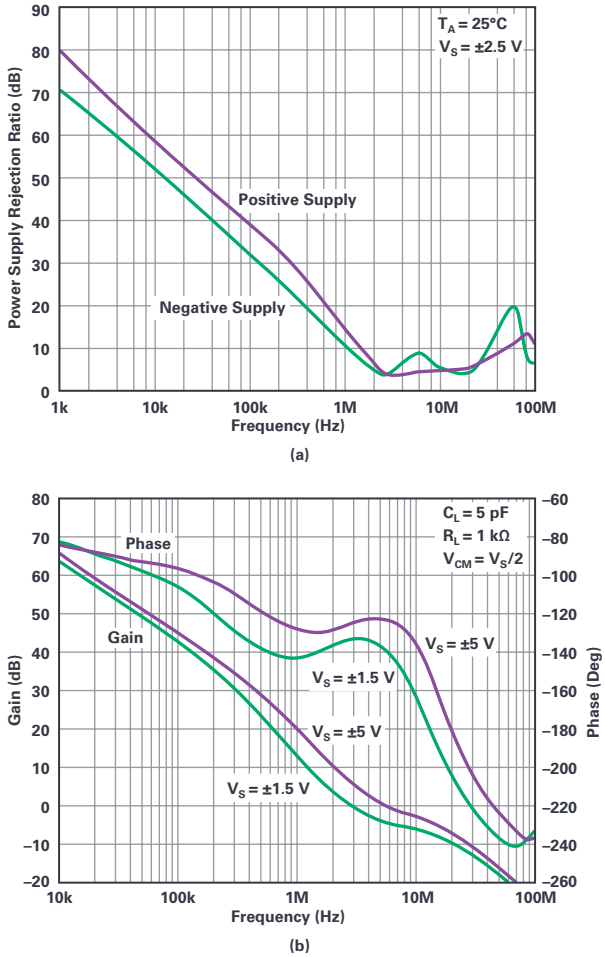

Figure 2. (a) PSRR of the LTC6240, (b) open-loop gain of the LTC6240.

Actual Implementation

Figure 3 shows the full circuit.

Note the 1000 pF bypass capacitors must be closely connected to the LTC6240 supply terminals. Op amps have dozens of internal transistors that, in this amplifier, have F_ts on the order of GHz. They are often connected in feedback to each other and, unless bypass capacitors are installed, can oscillate against a high ac impedance supply. 1000 pF is sufficient to quash those oscillations. We also want the supply bypass capacitor to be much greater than any output load capacitor, since at high frequencies voltage transitions across a load capacitor cause currents that flow to a supply rail and can modulate the supply voltage, feeding back through PSRR to cause oscillation. Our bypasses thus reduce supply modulation at frequency, equivalent to reducing feedback gain from output to supply.

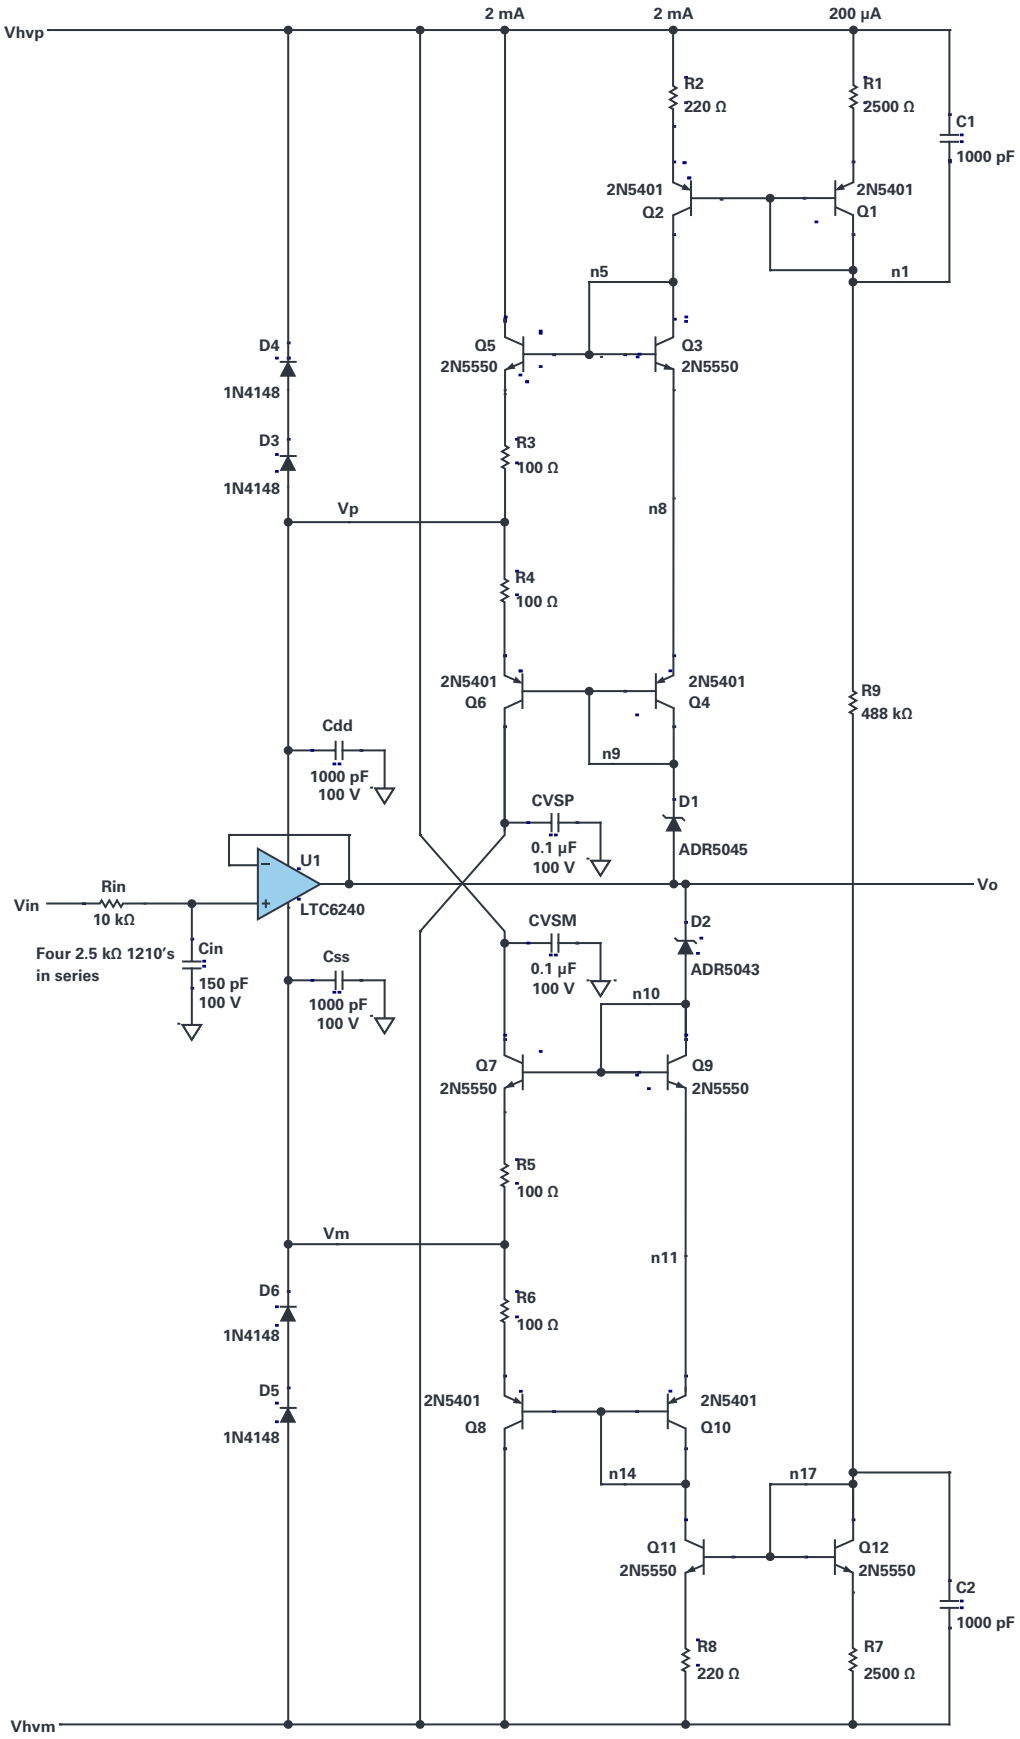

Figure 3. Full circuit.

Slewing those bypass capacitors will take serious current and must be bidirectional. Q5 and Q6 are emitter followers that can drive the bypass' slew currents. Q3 and Q4 are biasing diodes to set Q5 and Q6 quiescent currents. Q2 provides the bias current for those diodes and for Zener D1 (really a shunt reference IC), which sets a positive supply voltage relative to the output. Q2's collector is the output of a current mirror biased by R9 between the high voltage rails. R9 might be replaced with two current sources if the supply voltages are not constant.

Q7 through Q12 form the Vm minus supply driver equivalent to the previous description. Note the intentional mismatch in the Zener voltages: 5 V above the input/output for Vp and 3 V below input/output for Vm. The mismatch centers the input voltage within the LTC6240's supply-limited input range to optimize slew waveforms.

Normally, the supply current of the LTC6240 pulls against Q5's emitter and substantially turns off Q6, so that the Vp buffer output impedance is mostly R3. The bandwidth of the supply feedback Vp path is thus ~1/(2π × 100 Ω × 0.001 µF) = 1.6 MHz. This guarantees that the Vp loop gain is substantially less than one around 10 MHz and above, where LTC6240 open-loop phase is moving toward oscillation. The 100 Ω resistor also allows follower Q5 to not have to drive the 1000 pF directly. Emitter followers display output inductance that can resonate with capacitive loads, causing ringing or even oscillations.

Having designed the bootstrapping to fail at frequencies above 1.6 MHz, we will see that perfect behavior of the overall circuit will degrade beyond ~100 kHz. If the output cannot exactly follow the input, the benefits of bootstrapping will be degraded. Rin with Cin limits bandwidth to 100 kHz, part of a system antialias filter for an ADC to follow the buffer, and it also attenuates radio interference and unsupportable slew rates.

The circuit must tolerate any unlimited-slew input signal or ESD, so Rin also serves to limit input fault current. The resistor has four series segments to split up input overdrive and tolerate 1 kV temporarily. Depending on the signal source and anticipated overloads, the input resistor can be reduced.

There are protection diodes within the LTC6240 that guide input overvoltage currents to either Vp or Vm. The maximum fault current allowed into the LTC6240's input is 10 mA, but if there is surrounding circuitry that can quickly disconnect the input fault, that current can be increased for a short time. In the intended application of this circuit, there is an SPDT relay that, when unpowered, connects the input of the buffer to a ÷10 network. When powered, the relay connects the input directly. Thus, when unpowered the buffer connects to much more than 10 kΩ source impedance and the fault voltage and current are reduced commensurate with that 10 mA continuous rating. The input range of my application is continuously ±400 V with a fault tolerance of ±1000 V. This can only be safely done if there are two comparators that sense input overvoltage and quickly release the relay. This can be done in 1 ms to 2 ms, allowing a transient 100 mA input current that will not melt the protection diodes of the LTC6240. Note the inclusion of D3 through D6 to guide the input overload current which had been directed to Vp or Vm through the LTC6240 to the Vhvp or Vhvm supplies. These supplies probably cannot absorb the overload current since that current is backward to normal supply operation; we would depend on large enough bypass capacitance to hold the supply voltage safely while waiting for the relay switch relief. We would need 100 µF to hold the supply to within a 2 V change in 2 ms from a 100 mA overload.

A High Voltage Signal Source

When it came time to test the lab prototype, I realized that I had no signal generator with enough output voltage swing of any waveform to exercise the circuit. I do have generators that produce various waveforms to ±10 V p-p. It is time to come up with an amplifier that can cleanly reproduce waveforms at large amplitudes. Figure 4 shows a high voltage discrete realization of a current-feedback amplifier (CFA).

CFAs have fabulously high slew rate and, usually, wide bandwidth.[1] Because we are using high voltage transistors, though, the bandwidth is modest. High voltage transistors have higher parasitic capacitances and lower F_ts than lower voltage types.

Some warnings here. There is no current or dissipation limitation built into the circuit, so heavy sustained load currents more than 10 mA will burn out the output stage and maybe more stages. Further, it's best to not add bypass capacitors >0.1 µF to the high voltage supplies. A short-circuit can cause welding if a big capacitor is used. Having said that, I had to add 100 µF bypass capacitors to the high voltage supplies to suppress second-harmonic distortion. I crank the lab supplies up and down by hand to avoid hard turn-ons and turn-offs. Please note that even 50 V can cause enough current through a human to cause heart arrest. It's best to turn the current limit of the high voltage supply down to 60 mA as well. 50 V is high enough to be respected.

In Figure 4, the ADA4898 op amp controls the CFA and keeps its accuracies and distortions controlled. CFAs generally have high dc errors and poor settling to high accuracies; the op amp fixes those.

The positive input of the CFA is node n25 and its negative input is n5 (yes, that's an input). By itself, Rff and Rgg set the gain of the internal CFA to about 27. This high gain allows the controlling op amp's output to swing only ±2 V. The CFA could have been set to higher gain to unburden the control amplifier further, but then the CFA would lose bandwidth and increase its distortion. Overall gain is set by Rf and Rg and is 20. Ctweak and Ctweak2 work with Rf to remove the phase lag of the CFA from the overall op amp feedback above 215 kHz, enhancing op amp stability.

The n13 is the CFA gain node and is driven by current mirrors involving Q1/Q2/Q20 and Q11/Q12/Q19.

Q7/Q8/Q10/Q13 form the output buffer as a compound complementary emitter follower. There is no current limit circuitry—don't short the output to anything!

The CFA section of the high voltage amplifier has a 35 MHz, –3 dB bandwidth and does not peak, on its own. The overall circuit has a 33 MHz, –3 dB bandwidth but with 8 dB of peaking. Normally, the second amplifier of a composite amplifier design has at least 3× the bandwidth of the input control amplifier to avoid the peaking; but we could not get so favorable of a ratio. At least the 8 dB peak does not have a high Q and ringing damps reasonably fast. The intended 100 kHz signals are reproduced just fine below the peaking frequency. The distortion at an output of 80 V p-p at 100 kHz measured –82 dBc, dropping to –100 dBc for outputs of 32 V p-p and less at 100 kHz. A square wave response has a ~60% overshoot for fast edges and little to no overshoot occurs with output slew rates less than 250 V/µs. The maximum slew rate is about 1900 V/µs.

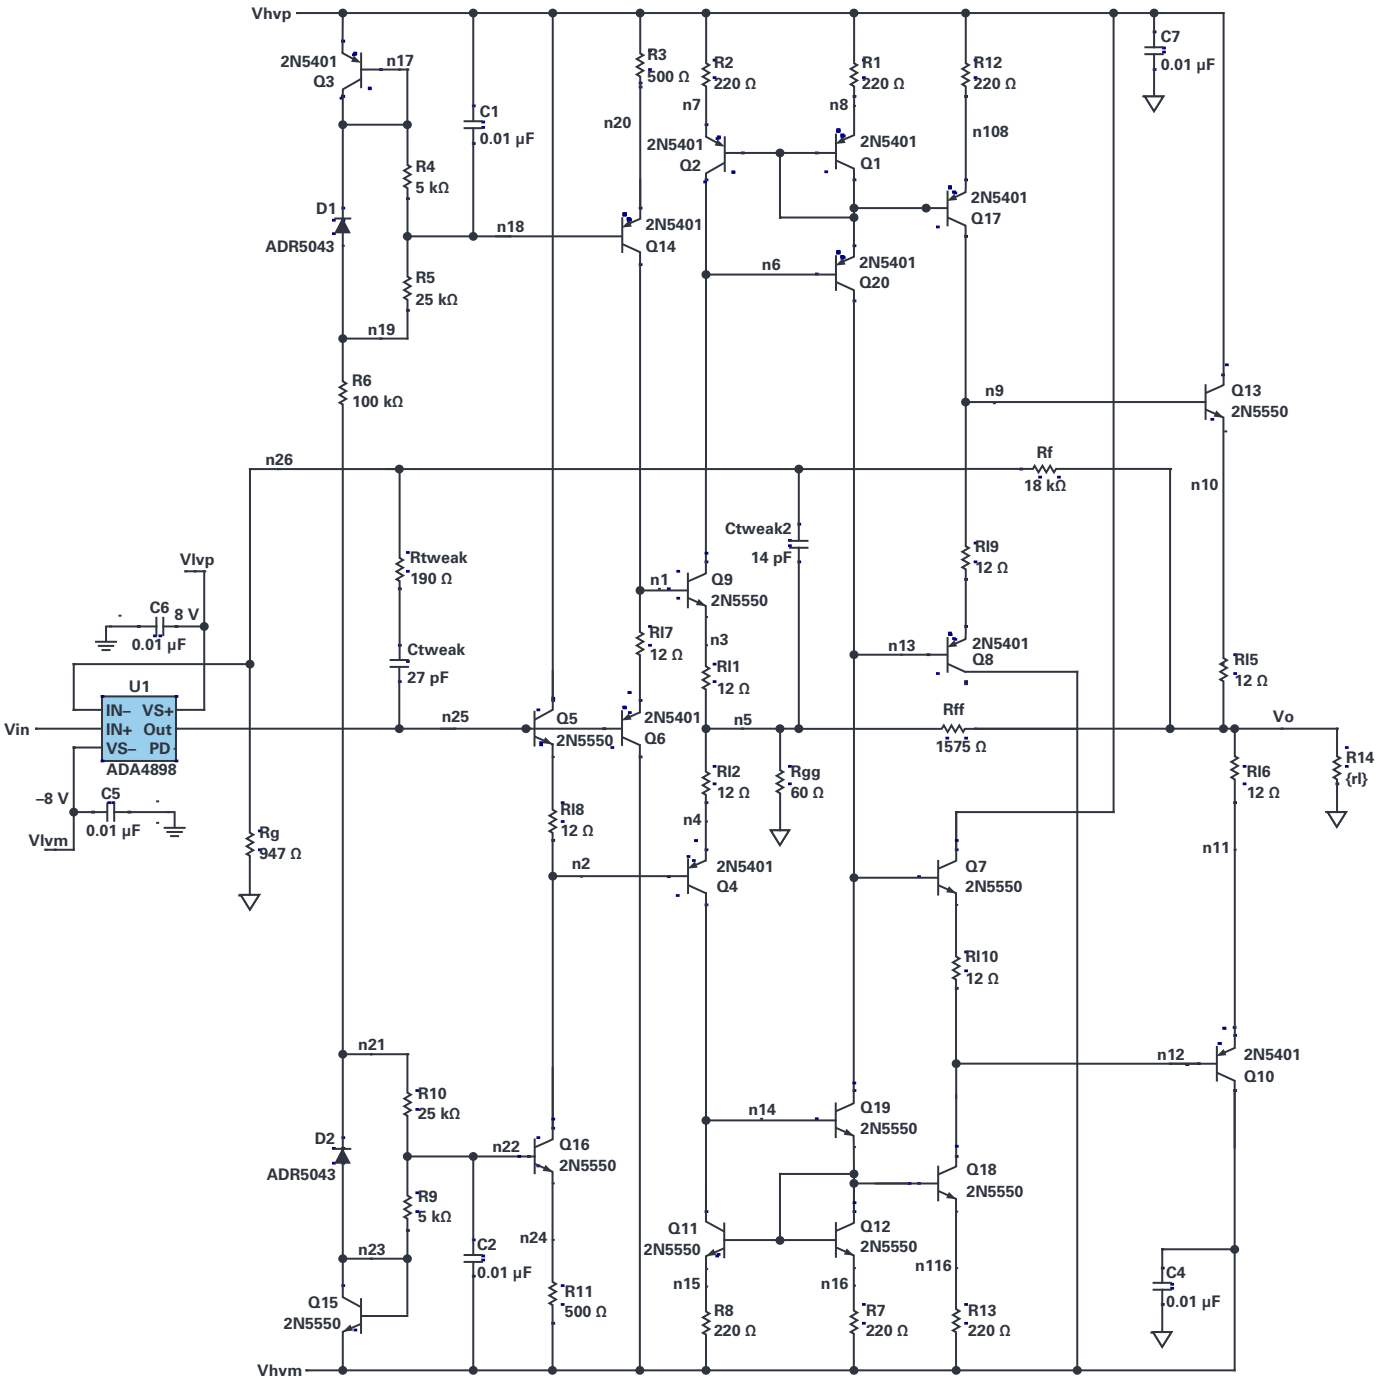

Figure 4. High voltage amplifier.

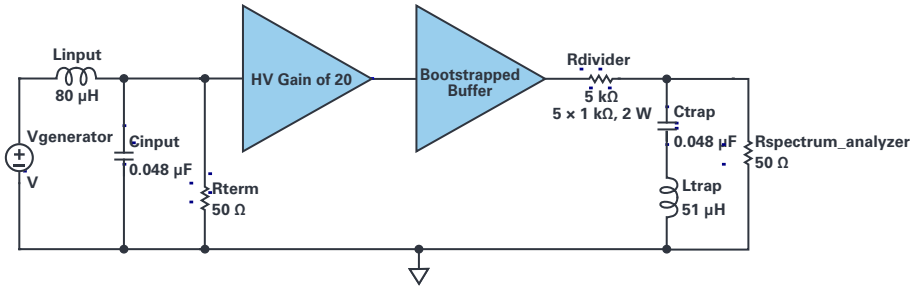

Figure 5. Distortion test setup.

Measurement Setup

Now that we have big signals, how do we use ordinary lab gear to measure the ±40 V outputs? Neither the high voltage amplifier nor the high voltage buffer should output more than 10 mA, nor can they work into more than ~40 pF load stably. At 27 pF/ft, coaxial cables are too capacitive. An oscilloscope ÷10 probe will have only ~15 pF||10 MΩ loading, so that will be fine for coupling to an oscilloscope.

For measuring distortion, none of the audio analyzers in our lab can beat –80 dBc at 100 kHz, so we must turn to spectrum analyzers. These unfortunately have only 50 Ω inputs—far too low for our circuits to drive. My solution was to raise the impedance to 5050 Ω (see Figure 5); that is, place a 5 kΩ divider resistor between the signal and the 50 Ω analyzer input, making close to a ÷100 divider. It is important that the 5 kΩ resistor not exhibit thermal shifts during low frequency signals, since these shifts are $V_{OUT}2$ related and cause even harmonics. I chose to put five 1 kΩ, 2 W resistors in series to make Rdivider. A 2 W resistor will have about 37°C/W thermal resistance, and the five 1 kΩ resistors have 7.5°C/W thermal resistance. With a ±40 V sinewave across it, there is a 160 mW dissipation, and that will cause 7.5 × 0.16 = 1.2°C resistors heating in the resistors. They have around 100 ppm/°C resistors shift, so at dc there would be a 120 ppm shift, or around 0.01% nonlinearity and –80 dBc generated distortion. How can this ever be accurate enough for our measurements? The good news is that the divider resistors have fairly long thermal time constants, and we expect little actual resistor shift in the middle of 100 kHz cycles. We would ironically see worse distortion at lower frequencies, probably 1 kHz and below.

The 80 V p-p signal had to be attenuated anyway because of limited analyzer input range, but it is still too large to get the best spectrum analyzer performance. Our analyzer can only offer –80 dBc distortion unaided, as a trade-off between its noise swamping the harmonics and large inputs causing additional distortion. A solution is to place a 100 kHz trap at the analyzer input to kill the fundamental amplitude. With less than a few millivolts of signal (harmonics only) we can approach –120 dBc measurement range. Figure 5 shows the test setup.

The generator drives Rterm through a low-pass filter Linput and Cinput that attenuates our generator's harmonics of 100 kHz. This improves its distortion to about –113 dBc, lower than the circuits to be measured. The cleaned-up signal is boosted by the high voltage amplifier and passed by the buffer, which drives the divider.

The inductors are constructed of magnet wire wound on large bobbins intended for power E-I cores. Core materials of any kind cannot be used due to added distortions; air-wound is mandatory. You just wind and measure repeatedly.

Ltrap was found to be magnetically radiating harmonics to adjacent, sloppy unshielded wiring, my usual approach, so I put the trap components in a cookie tin with a grounded BNC jack connection. We use cookie tins in our lab; I like roasting pots, but any shielding steel box would do.

For calibration, I replaced the two amplifiers with a through wire and logged the gain from the voltage at Rterm to the spectrum analyzer inputs at second through fourth harmonic frequencies. When measuring a harmonic in the distortion test, I use the stored gain at that frequency to infer the harmonic content at the output of the buffer. I have an oscilloscope monitoring the amplitude of the buffer fundamental frequency output, rms the normalized harmonics, and divide by fundamental amplitude to get overall distortion.

Results

Using the setup of Figure 5, the spectrum analyzer showed a distortion of –81 dBc at 70 V p-p and 80 V p-p output, –82 dBc at 50 V p-p and 60 V p-p out, and –86.5 dBc for 16 V p-p and 32 V p-p out, all at 100 kHz.

DC linearity, gain accuracy, and input range were then measured. Figure 6 shows the input offset of the buffer as we sweep the input dc signal.

Any amplifier with useful input characteristics can be bootstrapped as discussed to operate with high voltage signals. Very low input noise or exquisitely low offset amplifiers could run at hundreds of volts.

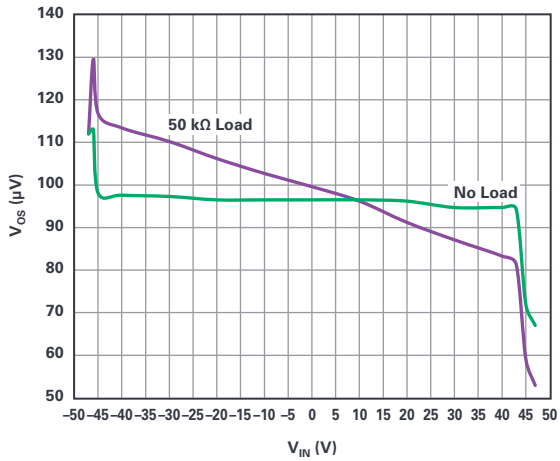

Figure 6. V_{OS} vs. V_{IN} of the buffer. Rl = 50 kΩ and ∞.

It is difficult for multimeters to resolve sub-microvolt variations against a background of ±40 V signals, but since this is a buffer we can simply connect a voltmeter from input to output to find offset shifts and use a sensitive range. The common-mode rejection of my multimeter was less than 1 µV for ±40 V inputs (inputs shorted for that test).

The perturbations in the curve are caused by low frequency noise and especially thermal perturbations. Just having a human in proximity or air conditioning can cause drafts and thermal variations that cause Seebeck and thermocouple voltage errors in a circuit at the microvolt level. I did not have a good shield or screen room, but I did cover my circuits with some clothing to prevent drafts. Even still, there is 0.6 µV rms wander in the results.

Amidst the noise, the unloaded (green) curve suggests a gain error of ~0.03 ppm. Not bad. The un-bootstrapped LTC6240 would have a nominal gain error of 5.6 ppm, and worst-case 100 ppm due to CMRR error. When loaded with 50 kΩ (purple), we see a gain error of –0.38 ppm. The loaded gain error is equivalent to an output impedance of 0.02 Ω. It's hard to know what the source of that 0.02 Ω comes from—it could be load currents modulating Vp or Vm and acting through common-mode rejection or gain limitation processes within the LTC6240, or it could simply be wire and circuit board resistances. In any event, to keep gain precise we could connect the feedback of the LTC6240 remotely to the final load to affect a Kelvined connection.

Figure 7 shows the small-signal pulse response.

Figure 7. Small-signal pulse response.

All apologies for the ringing of the green channel, which is the output of the high voltage amplifier. It doesn't ring on its own, I just had a mediocre oscilloscope probe and board-to-board grounding. The yellow channel is the buffer output, and it's a simple exponential dominated by the Cin + Rin.

Figure 8 shows the large signal pulse response with an input slew rate of ±32 V/µs—a nice, smooth response.

Figure 8. Large signal response to moderate input slew rate (±32 V/µs.)

Figure 9 shows the buffer response to an overloading slew rate. An 80 V p-p output at 100 kHz demands a peak slew rate of ±25 V/µs, within the ±32 V/µs capability shown.

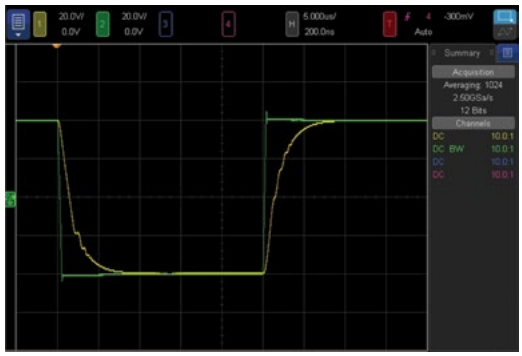

Figure 9. Large signal response to overloading input slew rate (±130 V/µs).

Note that the input filter limits the overloading slew rate to an amount the buffer can follow. The ripples are artifacts of the inability of the bootstrap circuitry to follow the output slew, which causes input headroom overloads repeatedly during the slew. Reducing Cin forces more input slew rate and the bootstrap circuitry will not follow, causing much uglier rippling.

Summary

A method to effectively bootstrap a low voltage op amp buffer to become a high voltage buffer has been shown. We have taken an op amp with rare input characteristics and elevated it to have higher voltage range, better gain accuracy, higher slew rate, and less distortion than the original op amp.

References

[1] Barry Harvey. "Application Note AN1106: Practical Current Feedback Amplifier Design Considerations." Renesas, March 24, 1998.

About the Author

Barry Harvey has worked as an analog IC designer, designing high speed op amps, voltage references, mixed-signal circuits, video circuits, DSL line drivers, DACs, sample-and-hold amplifiers, multipliers, and more. He has an M.S.E.E. from Stanford University. He holds more than 20 patents and has published about as many articles and papers. Barry's hobbies include repairing used test equipment, playing guitar, and working on Arduino-related projects. He can be reached at *barry.harvey@analog.com*.

Newest Sigma-Delta ADC Architecture Averts Disrupted Data Flow When Synchronizing Critical Distributed Systems

Lluis Beltran Gil

Abstract

This article introduces the way distributed data acquisition systems have been traditionally synchronized for both SAR ADC-based systems and sigma-delta (Σ-Δ) ADC-based systems, and explores the differences among the two architectures. Additionally, we will discuss the typical inconveniences when trying to synchronize multiple sigma-delta ADCs. Finally, a novel approach based on AD7770's sample rate converter (SRC) is presented, which shows how synchronization is achieved on sigma-delta ADC-based systems without interrupting the data flow.

Introduction

Have you ever imagined yourself in a supersonic aircraft breaking the sound barrier? Since Concorde was retired, this seems like an impossible dream, unless you are a military pilot or an astronaut.

As an electronic engineer, I'm fascinated with how everything works such as, for example, in a cuckoo clock, and how every single system works harmoniously in perfect synchronization with the others.

This applies to every single aspect of our lives. As we live in an interconnected world, everything is synchronized—from bank servers to the alarms of our smartphones. The only difference is the magnitude or complexity of the problem to solve in each particular case, the synchronization of different systems vs. the accuracy required—or tolerable error—or the size of the system to synchronize.

Distributed Systems

In a standalone design, the synchronization is inherent to the local clock or oscillator used. But when the standalone design should be integrated into a wider system—let's call it a distributed system—the perspective of the problem changes and the standalone design should be designed accordingly for the use case.

In a system, to calculate the instantaneous power consumption of an electric appliance, both current and voltage must be measured simultaneously.

By performing a quick analysis, you could solve the problem in three different ways:

- Use two synchronized single-channel ADCs to measure current and voltage.
- Use a multichannel simultaneous sampling ADC, either it has one ADC or one sample-and-hold circuit per channel.
- Use one multiplexed ADC and interpolate the measurements, in order to compensate for the time shift between the voltage and the current measurement.

At this point, you may have a solid solution to solve the problem, but let's expand the system requirements from the original single electric appliance to an application where the power must be measured in every single ac power socket in a factory. Now, your original instantaneous power consumption design should be distributed around a factory and somehow designed in such a way that the power in each ac power socket is measured and calculated at the same time.

You are now dealing with a distributed system that consists of a set of subsystems located apart from each other, but closely interrelated. Each subsystem needs to provide data sampled at the exact same point in time, such that the total instantaneous power consumption in the factory can be calculated.

Finally, if we keep expanding the hypothetical application example, imagine that your original design is going to be integrated into your country's power grid. Now, you are sensing millions of power watts, and any failure in a link could have terrible results, like damage to power lines due to stressful conditions that, in turn, may lead to a power outage, with dramatic consequences such as a wildfire or a hospital running without an energy supply.

So, everything needs to be precisely synchronized—that is, the data captured in the power grid shall be captured at the exact same point in time, independent of their geographical situation, as shown in Figure 1.

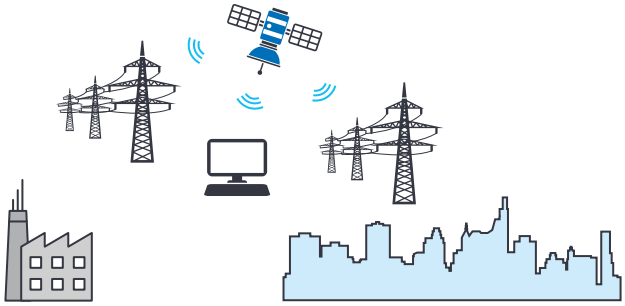

Figure 1. Power grid synchronization.

Under these circumstances, you may consider this as a critical distributed system, and you must get a continuous, fully synchronized stream of data from every single sense node.

Similar to the power grid example, these requirements apply to many other examples of critical distributed systems that may be found within the aerospace or the industry market, among others.

Nyquist and Oversampling ADCs

Before moving into explaining how to synchronize the sampling instant of many ADCs, it is good to understand how each ADC topology determines when to sample the analog input signal, and the benefits and drawbacks of each architecture.

- Nyquist or SAR ADC: This converter's maximum input frequency is dictated by Nyquist, or half-sampling frequency.
- Oversampling or sigma-delta ADC: The maximum input frequency is typically a fraction of the maximum sampling frequency, typically around 0.3×.

On one hand, the SAR ADC's input signal sampling instant is controlled through an external pulse applied to the conversion start pin. If a common conversion start signal is applied to every SAR ADC in the system being synchronized, as shown in Figure 2, they will all trigger the sampling simultaneously on the conversion start signal's edge. By just making sure there are no significant delays between the signals—that is, the conversion start pulses reach every SAR ADC at the same moment in time—the system synchronization can easily be achieved. Note the propagation delay between the pulse reaching the conversion start pin and the actual sampling instant should not vary from device to device, and it is not significant in precision ADCs where the sampling speed is relatively low.

Sometime after the conversion start pulse is applied, which is called conversion time, the conversion result will be available through the digital interface in all ADCs.

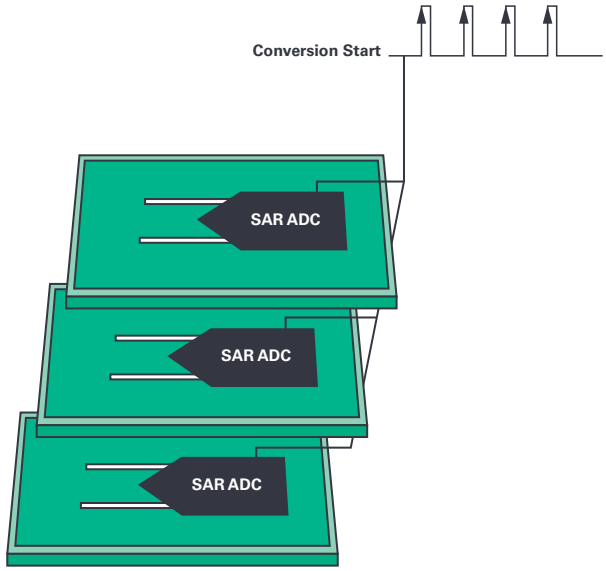

Figure 2. Synchronizing a SAR ADC-based distributed system.

On the other hand, sigma-delta ADC operation is slightly different due to its architecture. In this type of converter, the internal core—the modulator—samples the input signal at a much higher frequency (modulator frequency, f_{MOD}) than the minimum frequency dictated by Nyquist, which is the reason why it is called an oversampling ADC.

By sampling at a much higher frequency than strictly needed, an extra number of samples are gathered. All the ADC data is then postprocessed through an average filter for two reasons:

- The noise decreases 1 bit for every 4 averaged samples.
- An average filter transfer function is a low-pass filter. As the sigma-delta architecture pushes its quantization noise toward high frequencies, this shall be removed, as shown in Figure 3. So, this filtering is accomplished by this averaging filter.

The number of samples averaged, known as the decimation ratio (N), dictates the output data rate (ODR) that is the rate at which the ADC provides the conversion results, in samples per second, as indicated in Equation 1. The decimation ratio is typically an integer number with a set of predefined values discretely programmable (that is N = 32, 64, 128, etc.) on the digital filter. So, by keeping f_{MOD} constant, the ODR will be configured depending on the value of N, within the set of predefined values.

$$ODR = \frac{f_{MOD}}{N} \quad (1)$$

The averaging process is typically implemented internally by a sinc filter, and the analogous conversion start pulse for the modulator is generated internally as well, so there is no external control on triggering the conversion process. This type of converter is indeed continuously sampling,

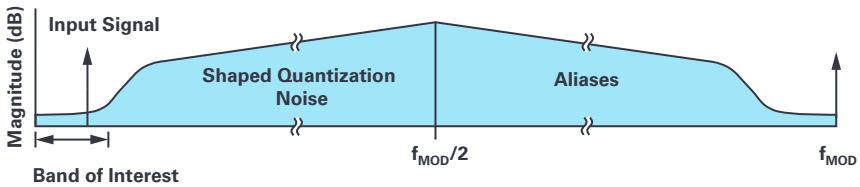

Figure 3. Sigma-delta noise shaping.

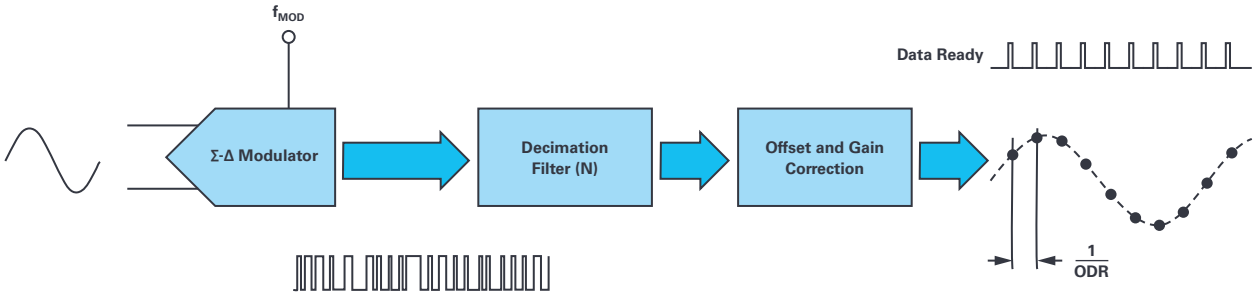

Figure 4. Sigma-delta ADC flow.

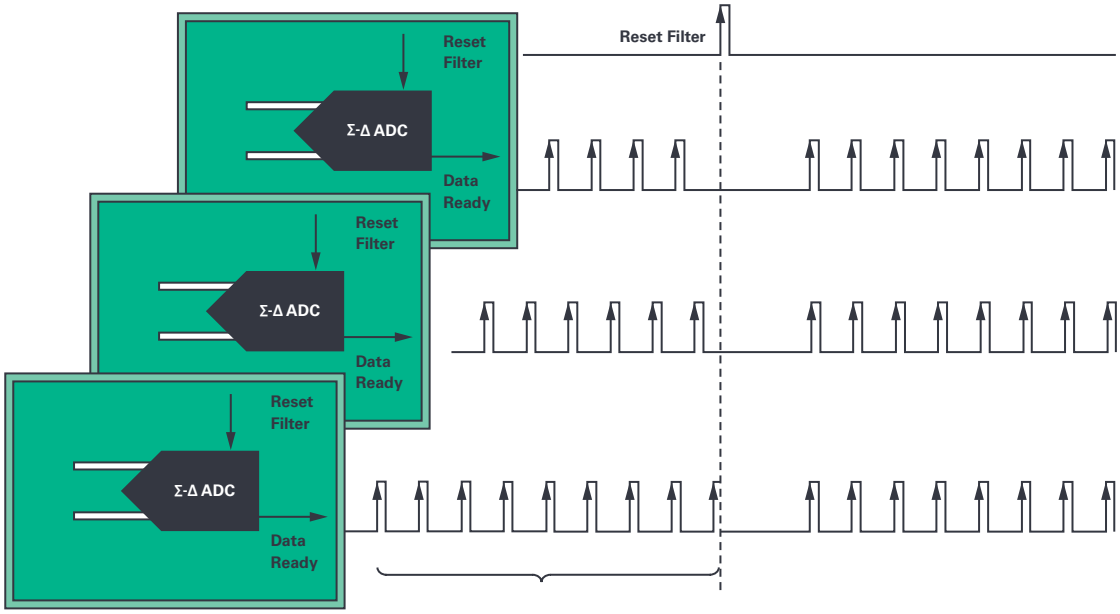

Figure 5. Sigma-delta system reset synchronization.

tracking the input signal, and processing the data acquired. Once the process (sampling and averaging) is completed, the converter generates a data ready signal to indicate to the controller that the data can be read back through the digital interface.

As shown in Figure 4, the flow for a sigma-delta can be summarized in four major steps,

- The modulator samples the signal at f_{MOD} frequency.
- The samples are averaged through the sinc digital filter.
- The result from the sinc filter is offset and gain corrected.
- The data ready pin toggles, indicating that the conversion result is ready to be readback by the controller.

As there is no external control on when the internal sampling triggers, in order to synchronize multiple sigma-delta ADCs within the distributed system, all the digital filters must be reset simultaneously as it is the digital filter that controls the start of the averaged conversion.

Figure 5 shows the effect on the synchronization assuming same ODR, and f_{MOD} in all the sigma-delta ADCs.

Similar to the SAR ADC-based system, it must be ensured here that the reset filter pulse reaches all subsystems at the same time.

However, note that every time the digital filter is reset, the data flow is interrupted as the filter must settle again. The duration of the data disruption in this case depends on the digital filter order, the f_{MOD}, and the decimation rate. An example is shown in Figure 6 where the LPF nature of the filter delays the time until a valid output is generated.

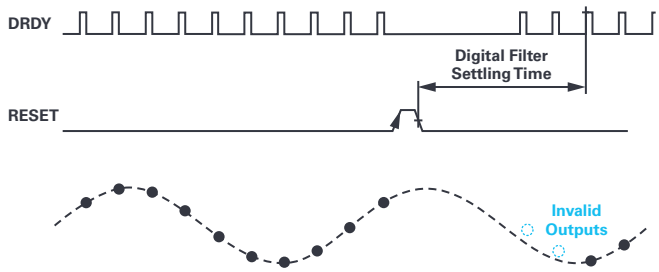

Figure 6. Data disruption due to digital filter settling time.

Implication on Synchronizing the Sampling in Distributed Systems

In a distributed system, the global synchronization signal (let's call it Global_SYNC) is shared across all the modules/subsystems. This synchronization signal could be generated by the master or by a third-party system, like the GPS 1 pps, as shown in Figure 1.

Once the Global_SYNC signal is received, each module must resynchronize the instantaneous sampling of each converter, and most probably its local clock, to guarantee simultaneity.

In a SAR ADC-based distributed system, the resynchronization is intrinsically easy as described in the previous section: the local clock (which manages the conversion start signal) realigns with the Global_SYNC signal, getting these signals in synchrony from then on.

This has an implication in terms of generating frequency spurs because, during the synchronization, there is one sample that has been gathered at different time and distance, as highlighted in blue in Figure 7. In distributed applications these spurs may be acceptable, while disrupting the data flow would have been indeed critical in applications like the power line monitoring mentioned earlier, for instance.

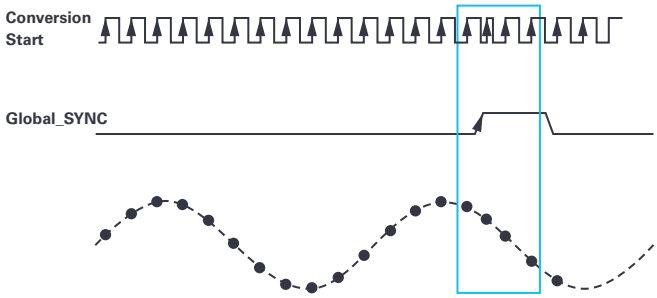

Figure 7. Aligning the SAR ADC conversion process to a global synchronization signal.

In sigma-delta-based distributed systems, the resynchronization with regards to a Global_SYNC signal is a little bit more complicated because the modulator is continuously sampling the analog input signal, and the conversion process is not externally controlled as it is in an SAR ADC.

The easiest way to synchronize multiple sigma-delta-based distributed systems is by resetting the digital filter: all modulator samples gathered and stored to be used on the average filter are dumped, and the digital filter is emptied. That means it will take some time, depending on the digital filter order, to settle its output again, as shown in Figure 5 and Figure 6.

Once the digital filter is settled, it will provide valid conversion data again, but resetting the digital filter on sigma-delta ADC implied data disruption may not be acceptable considering the amount of time the settling takes. The more often the distributed system needs to be resynchronized, the more data flow interruptions there will be, which can make sigma-delta ADCs impractical for critical distributed systems due to the continuous data flow disruption.

Traditionally, a method to minimize the data disruption has been the use of a tunable clock, like a PLL that decreases the error between the global sync and the f_{MOD} frequencies.

Once the Global_SYNC pulse is received, the uncertainty between the sigma-delta ADC conversion start and the Global_SYNC pulse can be calculated following a process similar to:

- The controller calculates the time difference between the sampling instant (calculated backwards from the data ready signal by knowing the group delay, as shown in Figure 8) and the Global_SYNC pulse. The group delay is a data sheet specification that accounts for the time between when the input is sampled until the data ready pin toggles, indicating that the sample is ready to be read.
- If there is a discrepancy between the sampling instant and Global_SYNC, the local controller quantifies the time difference (t_{ahead} or $t_{delayed}$), as shown in Figure 9.

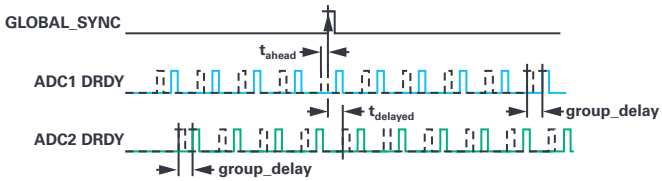

Figure 9. Quantifying the discrepancy from each ADC's sampling instant (provided the group delay is known) to the global synchronization signal.

- If there is a discrepancy, the sigma-delta filter could be reset or the f_{MOD} could be modified in order to adjust the sigma-delta sampling during a few samples. In both cases, a few samples would be missed. Note that by changing the local clock frequency (f_{MOD}), the sigma-delta ADC is modifying its output data rate (ODR = f_{MOD}/N), such that the ADC samples its analog input either slower or faster, with the intention for this ADC to catch up the rest of the ADCs on the system and with the Global_SYNC.
- If the f_{MOD} is updated, once in synchrony, the master clock frequency is reverted back to original frequency to return to previous ODR, while the subsystem becomes synchronous from then on.

This process of changing f_{MOD} during a certain period of time is shown in Figure 10.

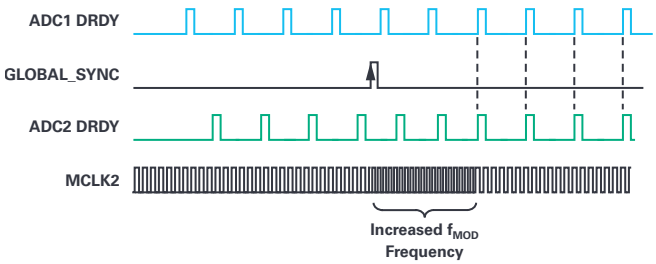

Figure 10. Synchronization method using PLL to tune the modulator frequency.

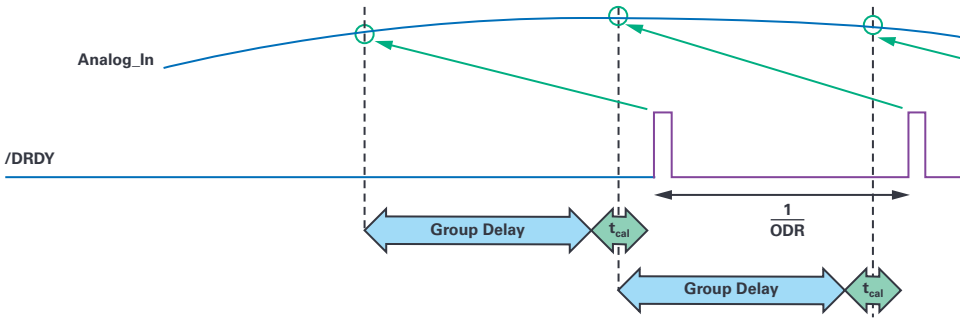

Figure 8. Time delay between the analog input being sampled and the data ready toggling.

This method may not be appropriated in some cases as there are a couple of details to consider:

- Changing the modulator frequency to non-multiple values may be unpractical.
- If fine frequency tuning is possible, the frequency steps while changing must be small, otherwise the digital filter may go out of bounds, so the lead time for the synchronization routine becomes longer.
- If the ODR change needed is big enough, it could be solved by changing the decimation rate (N) instead of the modulator frequency (f_{MOD}), but this would also imply missing samples.
- Using a PLL implies an extra amount of power being consumed plus its own settling time until reaching the desired modulator frequency.

In general, the overall system complexity and cost increase, scaling up with the system size, especially compared to SAR ADCs where this problem was solved more easily by just aligning the conversion start to the Global_SYNC signal. In addition, in many cases the use of a sigma-delta ADC must be not possible as per the system restrictions and limitations explained above.

Easily Resynchronize Sigma-Delta ADCs with No Data Disruption

The AD7770 family (which includes AD7770, AD7771, and AD7779) has a built-in SRC. With the introduction of this novel architecture, the restriction of the decimation factor (N) being a fixed value is no longer valid.

The SRC allows you to program decimal numbers, not only integers, as the decimation rate (N), which allows you to program any desired output data rate. On the previous synchronization method, as N was fixed, the external clock needed to change to adjust f_{MOD} in order to perform the synchronization routine.

Using AD7770 family of products, as the N is flexibly programmable and reprogrammable on-the-fly, any ODR can be programmed without having to change f_{MOD} and without data disruption.

This novel method for resynchronizing sigma-delta-based subsystems, making use of the SRC simplifies the resynchronization, minimizing the complexities described in the previous sections.

The new method is as follows:

- When the Global_SYNC signal is received, each subsystem checks if it is sampling synchronously or not, taking the data ready signal as a reference and using the group delay to find the actual sampling instant.
- If there is a discrepancy between the sampling instant and when the Global_SYNC signal has been received, the local controller quantifies the time difference (t_{ahead} or $t_{delayed}$) as shown in Figure 9.
- A new ODR is programmed to temporarily generate a faster or slower ODR through modifying the decimation factor (N) by means of the SRC. The whole operation of resynchronization will always take 4 samples (or 6 if the sinc5 filter is enabled on AD7771), but without interrupting the data flow due to these samples which are still valid and 100% settled.
- Once the required amount of DRDY has been received, the decimation factor is reprogrammed again to return to the desired ODR, which guarantees that the sigma-delta converter is synchronized with the rest of subsystems, as shown in Figure 11, with no data flow disruption.

Conclusions

Critical distributed systems require synchronous conversion in all the subsystems and a continuous data flow.

SAR converters provide an intuitive way to resynchronize the sampling by readjusting and aligning the conversion start signal with the Global_SYNC pulse.

In applications that require high dynamic range (DR) or signal-to-noise ratio (SNR) specifications, SAR is not an option, but traditional sigma-delta converters become a challenge to use due to their inflexibility to readjust without disrupting the data flow.

As seen in the example, the SRC provides a seamless synchronization routine with minimum latency and much lower cost and complexity than other solutions.

There are plenty of applications where the SRC can also be advantageous. As with the power line monitoring example, any line frequency change can be compensated for by immediately changing the decimation rate on the fly. That way, the power line is always sampled at a coherent sampling frequency. Here, in critical distributed systems, it's demonstrated that the SRC can also be very useful for resynchronizing the system without having to interrupt the data flow and with no need for extra components like PLLs. The AD7770 solves the traditional problem of synchronizing sigma-delta ADC-based distributed systems, without missing samples and avoiding the added cost and complexity of the PLL-based method.

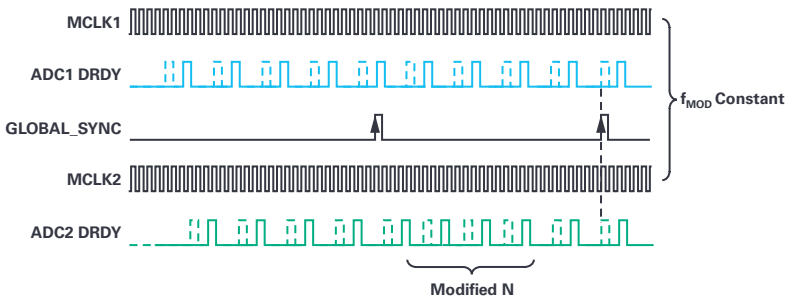

Figure 11. The sample rate converter adjusts the ODR on-the-fly in order to resynchronize the sampling on all devices.

About the Author

Lluis Beltran Gil received his B.S. degree in electronics engineering in 2009 and in industrial engineering in 2012, both from the Universitat Politècnica de València, UPV (Technical University of Valencia). After graduation, Lluis joined Analog Devices in 2013 as an applications engineer in the Precision Converter Group in Limerick, supporting temperature sensors. Currently, Lluis is working on the SAR ADC Applications Team within the Precision Converters Group, and he is based in Valencia, Spain. He can be reached at *lluis.beltrangil@analog.com*.

Does My Voltage Reference Design Hold Water? Methods of Managing Humidity and Performance in Precision Analog Systems

Paul Perrault and Robert Kiely

Introduction

Voltage references serve a crucial role in precision analog systems, often setting the noise/resolution floor within an analog-to-digital converter (ADC) for precision measurement systems in instrumentation, test and measurement, and energy metering applications. For the design engineer, a supplier's portfolio may include a bewildering array of silicon options. But among various voltage reference specifications (voltage noise, accuracy, drift, quiescent current, series vs. shunt, etc.) and their options for packaging (hermetic ceramic, plastic, die packaging), it's worth assessing whether the end electronics product is actually performing as optimally as desired. There are many design pitfalls that could easily undermine the µV or nV accuracy noise target you were looking to achieve. Through a holistic view of the PCB manufacturing process, this article explores ways that a design engineer or PCB assembly engineer can protect against damaging environmental effects while preserving analog performance.

Background

While every electronics design has different trade-offs in terms of performance, a general analog signal chain will have some manner of analog input signal conditioning, an ADC, and a voltage reference. For the purposes of this article, we'll use a medium speed 100 kSPS, 16-bit analog sensor input design, as shown in Figure 1. For more information on some of the trade-offs and design choices made in that signal chain, see the CN-0255 circuit note.

The 2.5 V voltage reference used in this application is the ADR4525 from the ADR45xx series of plastic-packaged voltage references, and it provides high precision, low power, low noise, and features ±0.01% (±100 ppm) initial accuracy, excellent temperature stability, and low output noise. System performance is improved by the ADR4525's low, thermally induced output voltage hysteresis and low long-term output voltage drift. A maximum operating current of 950 µA and a low dropout voltage of 500 mV maximum make the device optimum for use in portable equipment.

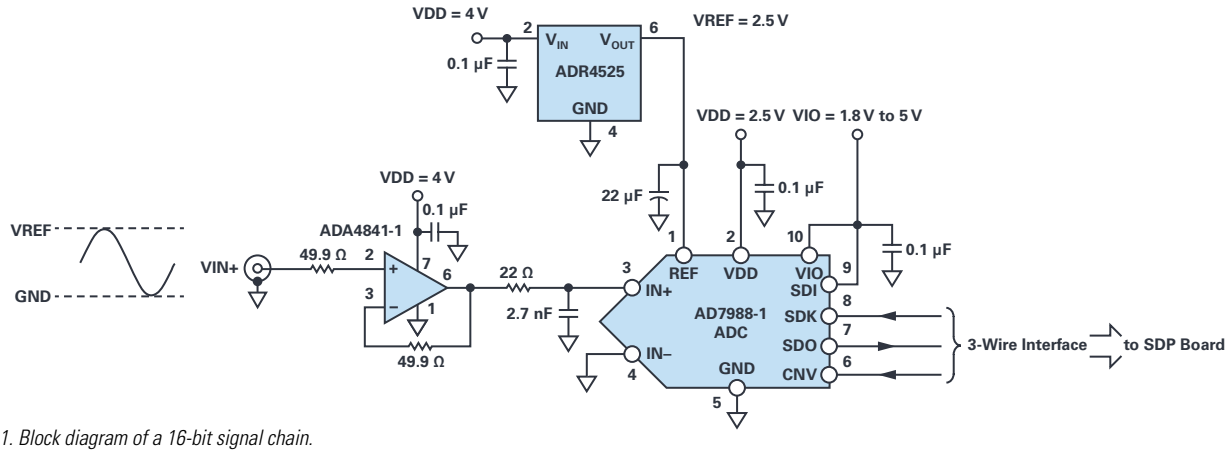

Figure 1. Block diagram of a 16-bit signal chain.

Once you've selected the components for your precision analog signal chain, it then falls to the PCB assembly team to manufacture a repeatable system using a printed circuit board as the substrate for your electronic design. As anyone who has worked with precision electronics knows, mechanical stresses at the board level can manifest themselves as dc offsets in a precision circuit design or similarly in a MEMS-based sensor design. This can often be proven by simply pressing on the plastic package of the reference and seeing the output voltage or sensor output change. Environmental factors like moisture and temperature can therefore affect electronic performance due to the differential stresses caused by moisture/humidity/temperature. Temperature causes mechanical stress in the package and board due to the different thermal expansion coefficients of the materials that make up the package and the board. Moisture causes mechanical stress in the package and board since both the plastic and the board absorb moisture, causing them to expand. The results of environmentally induced mechanical stress often take the form of increased drift over temperature/time in the case of plastic-packaged voltage references, or the form of increased offsets in plastic-packaged MEMS accelerometers. For plastic packages, the mechanical stress created by humidity is significant, and one way to manage this humidity effect is to package integrated circuits in ceramic or hermetically sealed packages. While this resolves a large amount of the challenge with humidity, this solution comes with the additional cost of a premium package, and often a larger sized component.

Conformal Coating Options

Another suggested method to isolate these stresses from the voltage reference is to use a conformal coating in the PCB manufacturing process so that any mechanical stress at the board level translates into less stress at the voltage reference. In this case, a thin layer of a coating compound across the voltage reference and corresponding PCB ensures that any stresses induced through moisture or temperature applied to the PCB don't translate to a differential stress to the voltage reference chip package and induce offsets. This also ensures that humidity from condensation due to cold temperatures has less effect on humidity ingress into the package.

HumiSeal®, a specialist coatings manufacturer, offers many conformal coatings that include acrylics, urethanes, silicones, epoxies, and water-based coatings for the protection of sensitive devices in PCB manufacturing. One of the parameters that allows the selection of the appropriate coating is called moisture vapor permeability (MVP), which is the rate at which moisture vapor passes through a coating. This is pertinent to our application since we are trying to make the PCB impervious to humidity.

The test method for MVP is to take a dry cup with the respective coating applied, put it into a temperature chamber at various humidities, and then periodically weigh the cup to assess how much moisture is traveling through the material and into the dry cup. After a week of this type of testing, it is apparent how effectively the coating slows down the progress of water.

Table 1 shows a selection of conformal coatings and their respective MVP normalized to a per-mil thickness value.

Table 1. Selection of HumiSeal Coatings and Their MVP

Material	Vapor Permeation ((g/m²)/day)	Normalized Vapor Permeation ((g/m²)/day/mil)	Thickness (mil)
HumiSeal 1A33	9.18	0.315	29.13
HumiSeal 2A64	13.54	0.249	54.33
HumiSeal 1A20	21.89	0.492	44.49
HumiSeal UV40	0.83	0.024	35
HumiSeal UV40	No data after one week of testing due to impermeability	No data after one week of testing due to impermeability	61.41
HumiSeal UV40-250	9.1	0.156	58.26
HumiSeal 1B73	25.1	1.2	20.86
HumiSeal 1C49LV	60.14	2.22	27
HumiSeal 1B51	0.78	0.026	35

Examining the table yields an important insight—in all cases (with the exception of a very thick UV-cured coating material called UV40), all of these coatings allow some amount of vapor through the coating over time. This is measured in the weight of water that permeates the coating through a given surface area in a given period of time; in these measurements the time period was seven days. Choosing the popular 1A33 coating (a polyurethane coating that is simple to apply, which also means it's cost-effective) shows that this coating is more than 10 times less effective at slowing down the rate of water vapor absorption than the rubber-based 1B51 coating at a similar thickness. The key message here, though, is that if they're left for a long enough period of time at high humidity, none of these coatings provide complete blockage against humidity.

This isn't to discount the usage of conformal coatings. Instead, it's useful to understand the environment in which the electronics will be deployed. Will the exposed electronics only experience short periods of high water vapor? Will the packaging/container of the electronics block water vapor, meaning that the conformal coating is as useful as wearing both a belt and suspenders? Does the environment of the electronics change so frequently that the purpose of the conformal coating is simply to block fast changes in the electronics? All of these questions are important for the product owner to consider before embarking on a conformal coating path.

One thought to consider here before moving ahead to actual data is that the use of conformal coatings in some cases can increase mechanical stress problems. This is because the coating, if improperly applied, can add stress to the package. For instance, if, in the PCB manufacture stage, the surface of the voltage reference package has moisture on it prior to coating, this almost certainly ensures that this moisture will migrate into the hydrophilic plastic package. From the data sheet for the 1A33 product: "Cleanliness of the substrate is of extreme importance for the successful application of a conformal coating. Surfaces must be free of moisture, dirt, wax, grease, flux residues, and all other contaminants. Contamination under the coating could cause problems that may lead to assembly failures." This is good guidance to anyone considering a conformal coating.

Data and Discussion: Does It Hold Water?

To assess the effect of conformal coatings, Analog Devices has produced a set of test boards. Each board has 27 of the same high performance voltage references soldered to the PCB using the recommended J-STD-020 reflow profile. The boards are placed in a humidity chamber and measured using a Keysight 3458A 8.5 digit digital multimeter (002 model) that achieves a 4 ppm/year drift by using the LTZ1000. The chamber is maintained at a constant temperature and humidity while the boards are allowed to settle. The boards sit in the chamber for up to a week before the humidity step is applied while keeping the temperature constant. Two different conformal coating processes were used on the plastic-packaged voltage references to assess the effect of humidity in the presence of the coatings.

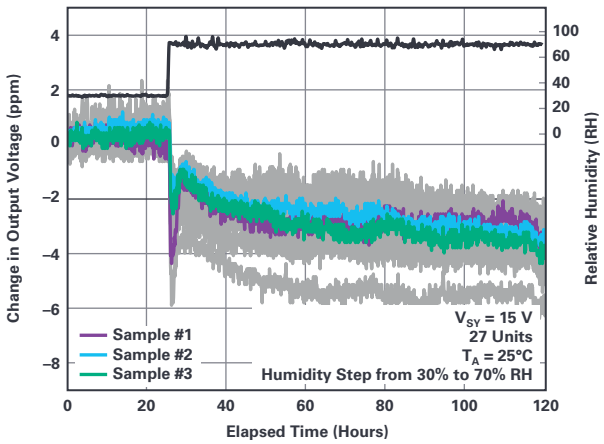

Figure 2. ADR4525 voltage reference in ceramic package.

Using the ceramic packaged ADR4525 as a baseline (Figure 2) establishes that 100 hours of being subjected to 70% humidity shows a change in the voltage output of ~3 ppm, or 0.075 ppm/% RH, which is excellent stability. The initial peak in the data is due to a jump in temperature that is caused by the sudden shift in humidity. The humidity chamber slowly recovers the temperature back to 25°C, as can be seen in the data. In contrast, the same voltage reference die when placed in plastic packages in the same environment and test conditions shows a voltage output change of ~150 ppm, as shown in Figure 3. Normalizing the data from Figure 3 with a 60% RH shift shows that the output drifts at ~2.5 ppm/% RH with no conformal coating applied. It also looks apparent that the drift hasn't completely settled out after 168 hours of soaking the boards in the high humidity environment.

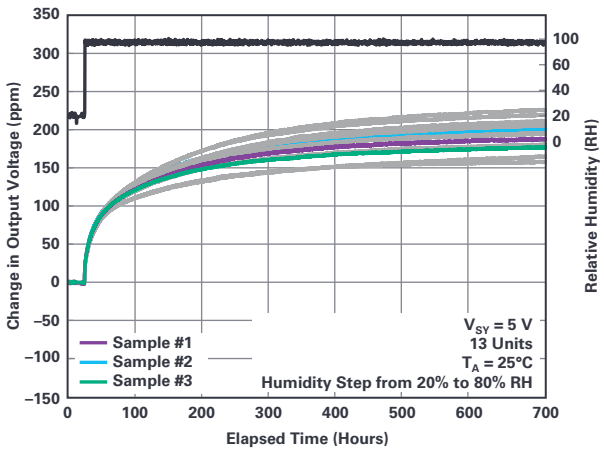

Figure 3. ADR4525 voltage reference in a plastic package subjected to a humidity step of 20% to 80%.

The HumiSeal 1B73 acrylic coating was tested next and the data is shown in Figure 4. The application procedure consisted of first washing and baking the board (submerge the boards quickly a few times into 75% isopropyl alcohol and 25% deionized water, lightly hand brush, then bake at 150°F for 2 hours), and then spraying the 1B73 coating to the specified thickness. The entire board was coated with the exception of the edge connector, which is required to be clean to enable measurement of the output voltage.

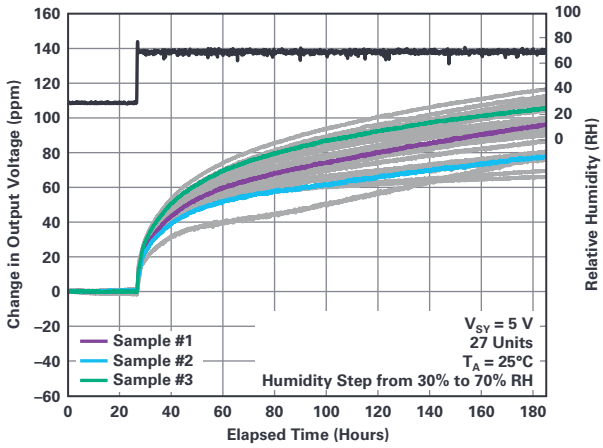

Figure 4. ADR45xx voltage reference coated with HumiSeal 1B73 acrylic coating using a spray application.

While the oven used in this test limited the humidity stress to 70% RH, the normalized drift looks like ~100 ppm/40% RH or 2.5 ppm/% RH, which isn't that different from the drift with no coating applied. In consultation with HumiSeal, it's possible that the coating didn't fully adhere to the underside of the voltage reference package along with the edges of the parts as well. It's also useful to note here that the ~168 hours of testing under the high humidity may still not be long enough as the voltage reference looks like it hasn't fully stabilized yet, similar to the uncoated parts. However, it is useful to see that the humidity effect does appear to have been slowed down in rate of change, at least at the initial time step, which lends credence to the concept of the moisture vapor permeability rate, where the coating isn't stopping the moisture, but is instead slowing it down.

The next test tried the same conformal coating (HumiSeal 1B73) but with a three-step application process that used a dip coating process to better ensure complete coverage of the board. This data is shown in Figure 5.

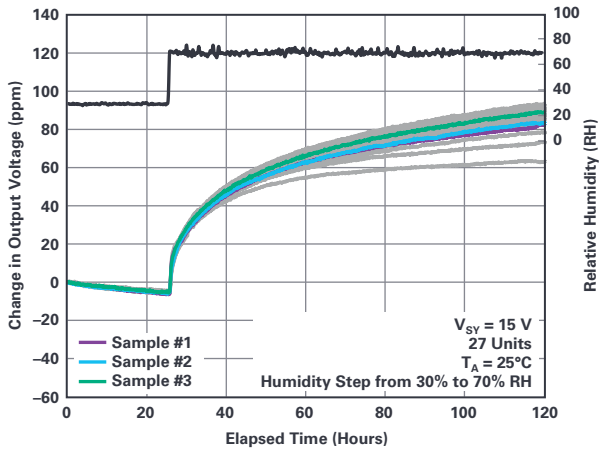

Figure 5. ADR45xx voltage reference coated with HumiSeal 1B73 acrylic coating using a 3-step dip coating application.

Issues with the oven prevented more than 96 hours of testing for this test. Normalizing the data set for a 30% RH to 70% RH step shows ~90 ppm or 2.3 ppm/% RH, which is not the massive improvement that was hoped for with this application process, but a slight improvement over the spray coating—although it's fair to say that if left for a longer test, this slight improvement may disappear. The three tests are summarized in Table 2.

Table 2. Summary of Humidity Testing with Conformal Coatings

	ADR4525 Plastic, No Coating	ADR45xx Plastic, 1B73 Spray Coated	ADR45xx Plastic, 1B73 Dip Coated	ADR4525 Ceramic
Hours of Testing	168	168	96	168
RH Test Regime	20% RH to 80% RH	30% RH to 70% RH	30% RH to 70% RH	30% RH to 70% RH
Output Drift Result	2.5 ppm/% RH	2.5 ppm/% RH	2.3 ppm/% RH	0.075 ppm/% RH

Future testing could include other types of conformal coatings (silicone, rubber, etc.) along with many other variations in the application process. In addition, doing cross-section analysis after coating would also confirm whether the application thickness matches the manufacturer's specifications, along with whether or not some of the edges were adequately coated. In short, the data from these experiments show that the ceramic hermetically sealed package is the single best defense against humidity ingress.

Conclusion

In a design that is only targeting 10 bits of accuracy (1 in 1000 type accuracy or ±5 mV in a 5 V reference), there is a lot of room for hiding inaccuracies from a variety of error sources. However, if your precision instrumentation system is targeting 16 bits and even 24 bits of accuracy, it's imperative to look at a whole system design, including PCB manufacture, to ensure complete accuracy over the life cycle of your design. This article shows that the best way to ensure humidity performance is to use hermetic packages like ceramic, and that conformal coatings can slow down the humidity effects within precision analog electronics. As the design engineer moves their design through to manufacture, it becomes necessary to leverage skills outside of electronics, consulting with coatings companies to get the absolute best performance in challenging environments. The expression "this argument holds water" typically means that your argument has merit and is true. In this case, following best practices will ensure that your voltage reference itself won't hold water, but instead will keep water out and preserve the performance you need in your precision design. This design approach holds water while your voltage reference doesn't!

References

ASTM E398-03, Standard Test Method for Water Vapor Transmission Rate of Sheet Materials Using Dynamic Relative Humidity Measurement. ASTM International, 2003.

Bryant, James. "Ask the Applications Engineer—11: How Good Must a Voltage Reference Be?" *Analog Dialogue*, Jan. 1992.

HumiSeal 1A33 Urethane Conformal Coating Technical Data Sheet. HumiSeal, 2019.

"IPC-HDBK-830: Guidelines for Design, Selection and Application of Conformal Coatings". IPC, Oct. 2002.

"MT-087 Tutorial: Voltage References." Analog Devices, Inc., 2009.

About the Author

Paul Perrault is a senior staff field applications engineer based in Calgary, Canada. His experience over the last 17 years at Analog Devices varies from designing 100+ amp power supplies for CPUs to designing nA-level sensor nodes and all current levels in between. He holds a B.Sc. degree from the University of Saskatchewan and an M.Sc. degree from Portland State University, both in electrical engineering. In his spare time, he enjoys back-country skiing in hip-deep powder, rock climbing on Rockies' limestone, scrambling and mountaineering in local hills, and spending time outdoors with his young family. He can be reached at *paul.perrault@analog.com*.

About the Author

Rob Kiely is a staff applications and marketing engineer in the Instrumentation and Precision Technology Group in Santa Clara, California. He joined Analog Devices in 2010. Rob's focus is in precision signal chains and products including Σ-Δ ADCs, precision amplifiers, and voltage references. Rob received his B.S.E.E. and M.E. in VLSI systems from the University of Limerick, Ireland. He can be reached at *rob.kiely@analog.com*.

Is It Possible to Fit Low EMI Power Supplies onto Crowded Boards?

Bhakti Waghmare and Diarmuid Carey

Limited and shrinking board real estate, tight design cycles, and stringent electromagnetic interference (EMI) specifications, such as CISPR 32 and CISPR 25, are limitations that make it difficult to produce power supplies that feature high efficiency and good thermal performance. Matters are further complicated by design cycles that often push power supply design to near the end of the design process—a recipe for frustration, as designers try to squeeze complex power supplies into tighter spots. Performance is compromised to finish designs on time, kicking the can down the road into test and validation. Simplicity, performance, and solution volume are traditionally at odds: prioritize one or two desired features and live with not having the third—especially when design deadlines loom. Sacrifices are accepted as normal; they should not be.

This article begins with an overview of a significant issue posed by power supplies in complex electronic systems: EMI, often simply called noise. Power supplies produce it, and it must be addressed, but what are the sources and what are typical mitigation strategies? This article addresses EMI reduction strategies, presenting a solution to reduce EMI, maintain efficiency, and fit power supplies into limited solution volumes.

What Is EMI?

Electromagnetic interference is an electromagnetic signal that disrupts the performance of the system. This disturbance affects circuitry by electromagnetic induction, electrostatic coupling, or conduction. It is a critical design challenge for automotive, medical, and test and measurement equipment manufacturers. Many of the limitations mentioned above and the increasing performance demands on power supplies—increasing power density, higher switching frequencies, and higher current—only serve to expand the effects of EMI, calling for solutions to reduce it. In many industries, EMI standards must be met, significantly impacting time to market if not considered early in the design cycle.

Types of EMI Coupling

EMI is a problem in electronics systems when the source of the interference couples with a receiver—namely some component in an electronic system. EMI is classified by its coupling medium: conducted or radiated.

Conducted EMI (Low Frequency, 450 kHz to 30 MHz)

Conducted EMI couples to components via conduction through parasitic impedances and power and ground connections. The noise is transmitted by conduction to another device or circuit. Conducted EMI can be further classified as either common-mode or differential-mode noise.

Common-mode noise is conducted via parasitic capacitance and a high dV/dt (C × dV/dt). It follows a path from any signal (positive or negative) to GND via parasitic capacitance, as shown in Figure 1.

Differential-mode noise is conducted via parasitic inductance (magnetic coupling) and a high di/dt (L × di/dt).

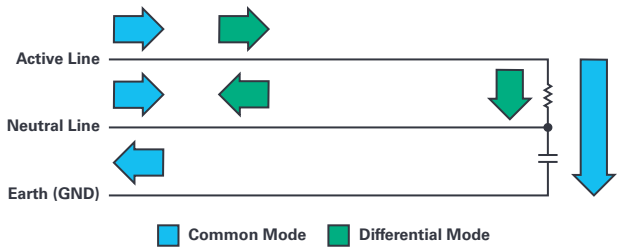

Figure 1. Differential-mode and common-mode noise.

Radiated EMI (High Frequency, 30 MHz to 1 GHz)

Radiated EMI is noise that is transmitted wirelessly via magnetic energy to the device under test. In a switching power supply, the noise is the result of high di/dt coupled with parasitic inductance. This radiated noise can affect nearby devices.

EMI Control Techniques

What is the typical approach to resolving EMI-related issues in a power supply? First, establish that EMI is an issue. It seems obvious, but acquiring this knowledge can be time consuming as it requires access to an EMI chamber (not available on every corner) to quantify how much electromagnetic energy is produced by a power supply and if it falls sufficiently within the standards posed by the system.

Assuming, after test, a power supply poses an EMI problem, one is faced with the process of reducing it via a number of traditional correction strategies, including:

- Layout optimization: Careful power supply layout is as important as choosing the right components for a supply. A successful layout depends greatly on the experience level of the power supply designer. Layout optimization is inherently iterative and an experienced power supply designer can help minimize the number of iterations, avoiding time delays and additional design costs. The problem: such experience is not often available in-house.

- Snubbers: Some designers plan ahead and provide footprints for simple snubber circuits (simple RC filter from the switch node to GND). This can dampen switch node ringing—an EMI contributor—but this technique results in increased losses, negatively impacting efficiency.
- Reduced edge rates: Reducing switch node ringing can also be achieved by reducing the slew rate of the gate turn on. Unfortunately, like a snubber, this negatively affects overall system efficiency.
- Spread spectrum frequency modulation (SSFM): This feature, implemented as an option in many Analog Devices Power by Linear switching regulators, helps designs pass stringent EMI test standards. In SSFM, the clock used to drive the switching frequency is modulated over a known range (for example, ±10% variation around the programmed f_{SW}). This helps to distribute the peak noise energy over a wider frequency range.
- Filters and shielding: Filters and shielding are invariably costly in money and space. They also complicate production.
- All of the above contingencies can reduce noise, but they all come with drawbacks. Minimizing noise at the power supply design is often the cleanest path, but difficult to achieve. ADI Silent Switcher and Silent Switcher 2 regulators achieve low noise at the regulator, avoiding the need for additional filtering, shielding, or significant layout iterations. Avoiding costly countermeasures speeds up product time to market and can save significant cost.

Minimizing Current Loops

To reduce EMI, one must determine the hot loop (high di/dt loop) in the power supply circuit and reduce its impact. The hot loop is shown in Figure 2. In one cycle of a standard buck converter, ac flows through the blue loop with M1 closed and M2 open. During the off cycle with M1 open and M2 closed, current follows through the green loop. It is not entirely intuitive that the loop producing the highest EMI is neither the blue nor green loops—only the purple loop conducts a fully switched ac, switched from zero to I_{PEAK} and back to zero. This loop is called the hot loop because it has the highest ac and EMI energy.

The high di/dt and parasitic inductance in the switcher hot loop that causes electromagnetic noise and switch ringing. To reduce EMI and improve functionality, one needs to reduce the radiating effect of the purple loop as much as possible. Radiated emission of the hot loop rises with its area, so reducing the PC area of the hot loop to zero and using an ideal capacitor with zero impedance could solve the problem, if that were possible.

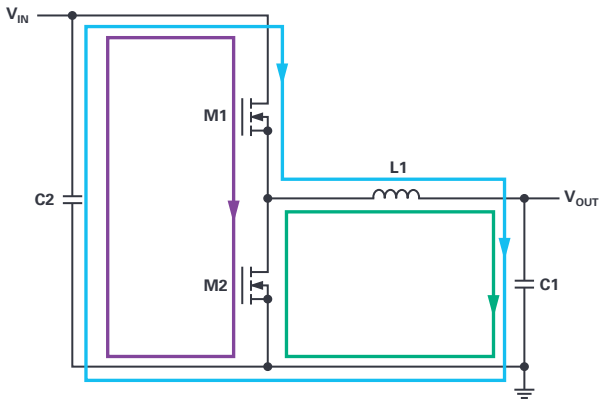

Figure 2. Buck converter hot loops.

Achieve Low Noise with Silent Switcher Regulators

Magnetic Cancellation

It is impossible to reduce hot loop area to zero, but we can split the hot loop into two loops with opposite polarities. This effectively contains the magnetic field locally, with the fields effectively mutually cancelling each other at any distance from the IC. This is the concept behind Silent Switcher regulators.

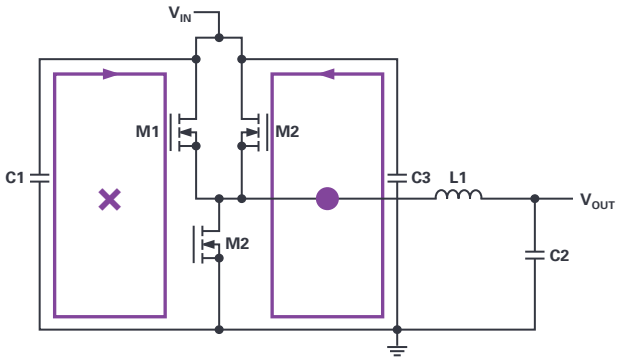

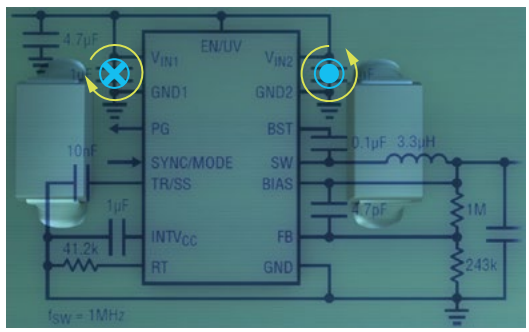

Figure 3. Magnetic cancellation in a Silent Switcher regulator.

Flip Chip Replaces Wirebond

Another way to improve EMI is to shorten the wires in the hot loop. This can be done by removing the traditional wirebond method of connecting the die to the package pins. In the package, the silicon is flipped and copper pillars are added. This further minimizes the area of hot loops by shortening the distance from the internal FET to the package pin and the input capacitors.

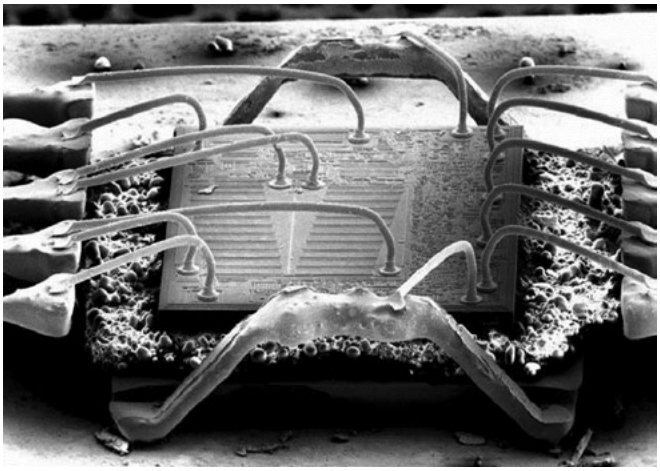

Figure 4. Wirebond shown in a disassembled LT8610.

Silent Switcher vs. Silent Switcher 2

Figure 6 shows a typical application using a Silent Switcher regulator, recognizable by the symmetrical input capacitors at the two input voltage pins. Layout is important in this scheme, as Silent Switcher technology requires these input capacitors are laid out as symmetrically as possible to provide the mutual field cancelling benefit. If not, the benefits of the Silent Switcher technology are lost. The problem, of course, is how does one ensure proper layout in design and all the way through production? The answer is Silent Switcher 2 regulators.

Silent Switcher 2

Silent Switcher 2 regulators take EMI reduction one step further. The EMI performance sensitivity to PCB layout is eliminated by incorporating the capacitors into the LQFN package—V_{IN} caps, $INTV_{CC}$, and boost caps—allowing for placement as close as possible to the pins. All of the hot loops and ground planes are internal, resulting in minimized EMI, and an overall smaller solution footprint.

Silent Switcher 2 technology also results in improved thermal performance. The large, multiple-ground exposed pads on the LQFN flip-chip package facilitate the extraction of heat from the package into the PCB. Higher conversion efficiency also results from the elimination of high resistance bond wires. When tested for EMI performance, the LT8640S passes CISPR 25 Class 5 peak limits with a wide margin.

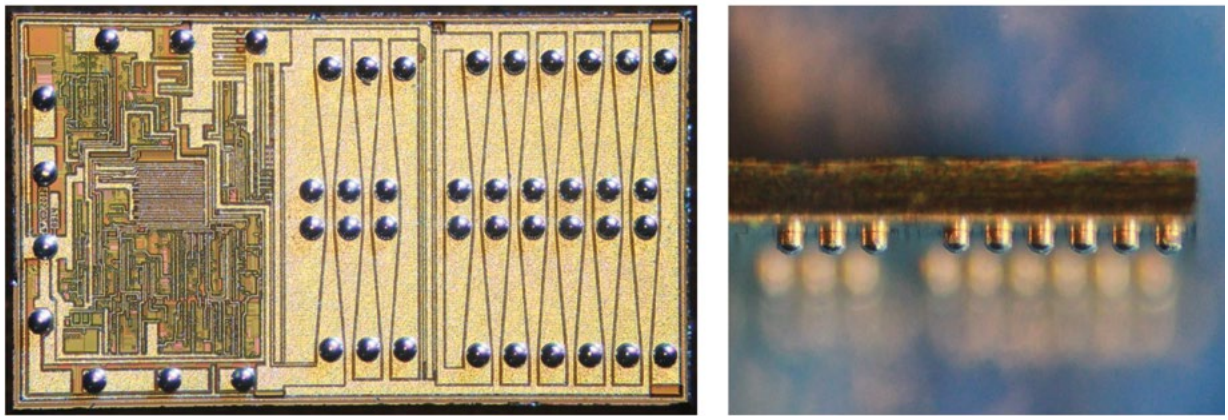

Figure 5. Flip chip with copper pillars.

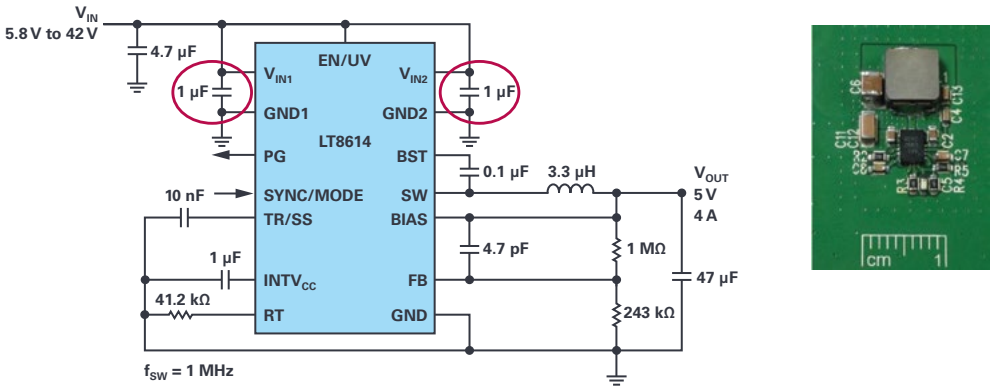

Figure 6. Typical Silent Switcher application schematic and how it looks on the PCB.

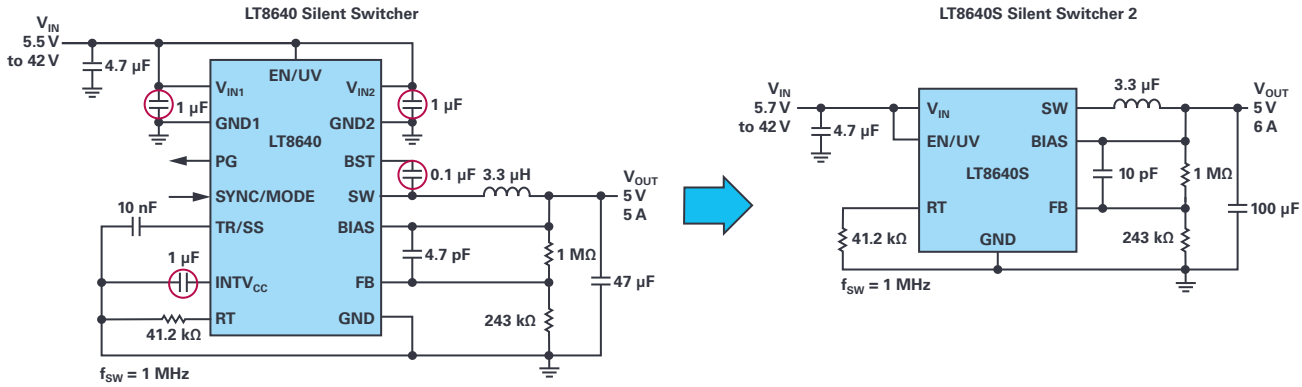

Figure 7. Silent Switcher application vs. Silent Switcher 2 application diagrams.

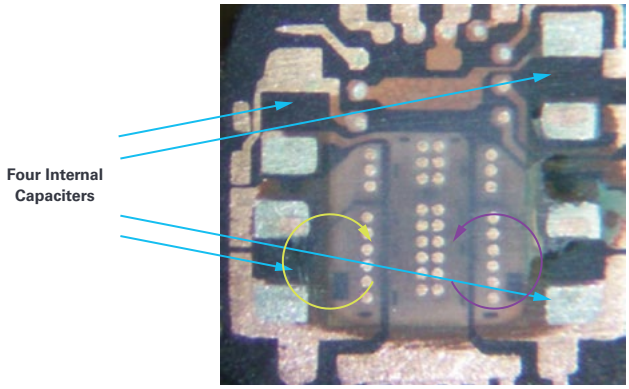

Figure 8. Decapped LT8640S Silent Switcher 2 regulator.

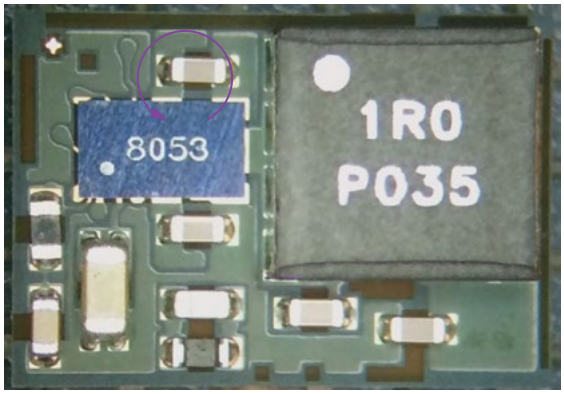

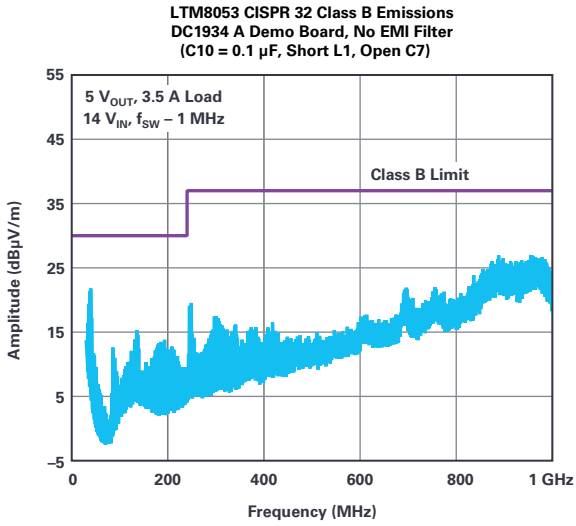

Figure 9. LTM8053 Silent Switcher exposed die and EMI results.

µModule Silent Switcher Regulator

Taking the knowledge and experience gained while developing the Silent Switcher portfolio and coupling it with an already vast µModule portfolio allows us to deliver a power product that is easy to design while meeting some of the most important metrics of a power supply—thermal, reliability, accuracy, efficiency, and great EMI performance.

Figure 9 shows the LTM8053 incorporating the two input caps, allowing for magnetic field cancellation, as well as a number of other passive components required for this power supply to operate. All of this is achieved in a 6.25 mm × 9 mm × 3.32 mm BGA package, which allows customers to focus their efforts on other areas of their board design.

Removing the Need for an LDO Regulator—Power Supply Case Study

A typical high speed ADC requires a number of voltage rails, and some of those rails must be very quiet to achieve the ADC's highest data sheet-listed performance. The generally accepted solution to achieving a balance of high efficiency, small board area, and low noise is to combine switching power supplies with LDO postregulators as shown in Figure 10. Switching regulators are able to achieve relatively high step-down ratios with high efficiency, but are relatively noisy. A low noise LDO postregulator is relatively inefficient, but it can reject much of the conducted noise produced by the switching regulator. Efficiency is helped by minimizing the step-down ratio of the LDO postregulator. This combination produces clean supplies, resulting in the ADC operating at top performance levels. The problem is a complex layout of numerous regulators, and LDO postregulators can have thermal issues at higher loads.

In the design shown in Figure 10, several trade-offs are evident. In this case, low noise is a priority, so efficiency and board space must suffer. Or, maybe not. The latest generation of Silent Switcher µModule devices combines low noise capability switching regulator design with µModule packaging—achieving a heretofore unachievable combination of easy design, high efficiency, compact size, and low noise. These regulators minimize board area, but also enable scalability—several voltage rails can be powered from one µModule regulator, providing further area and time savings. Figure 11 shows an alternative power tree using the LTM8065 Silent Switcher µModule regulator to power the ADC.

These designs have been tested against each other. ADC performance using the power supply designs in Figure 10 and Figure 11 have been tested and compared in a recent article published by ADI.[1] Three configurations were tested:

- A standard configuration using a switching regulator and an LDO regulator to power the ADC.
- Using the LTM8065 to directly power the ADC with no further filtering.
- Using the LTM8065 with the addition of an output LC filter to further clean the output.

Measured SFDR and SNRFS results showed that the LTM8065 could be used to power the ADC directly without compromising the performance of the ADC.

The core benefit of this implementation is a significant reduction in the number of components resulting in higher efficiency, far easier production, and reduced board area.

Summary

In conclusion, as we see a shift toward more system-level designs with ever more stringent specifications, it is important to utilize modular power supply designs where possible, especially where there is little power supply design expertise. With many market segments requiring that the system design pass the latest EMI specifications, the use of Silent Switcher technology is integrated into the small form factor, and ease of use of a µModule regulator can drastically improve your time to market while saving board area.

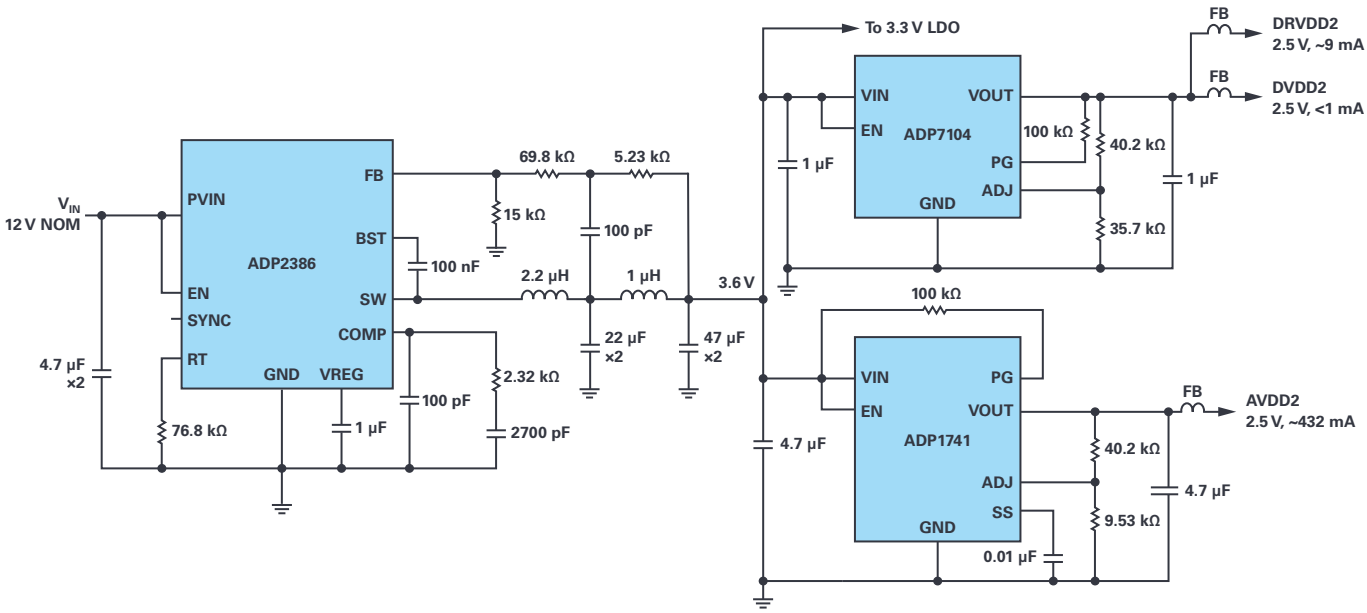

Figure 10. Typical power supply design for powering an AD9625 ADC.

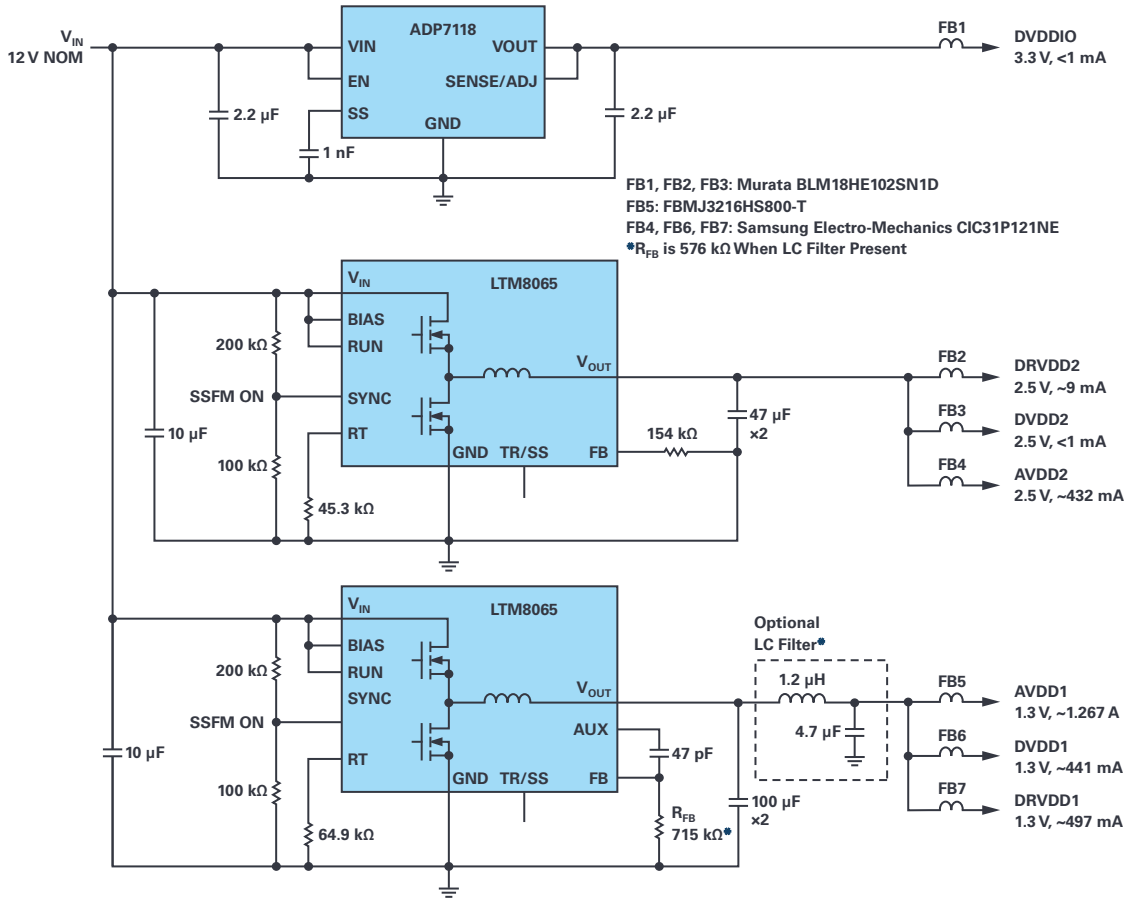

Figure 11. Space-saving solution to powering the AD9625 using a Silent Switcher µModule regulator.

Advantages of Silent Switcher µModule Regulators

- Saves PCB layout design time (no need to respin the board to rectify noise issues).
- No need for additional EMI filters (saving cost on components and board area).
- Reduces the need for an in-house power supply expert to debug the power supply noise.
- High efficiency at wide operating frequency range.
- Can remove the need for an LDO postregulator when powering noise-sensitive devices.
- Shortened design cycle.
- High efficiency in a minimal board area.
- Good thermal performance.

References

[1] Aldrick Limjoco, Patrick Pasaquian, and Jefferson Eco. "Silent Switcher µModule Regulators Quietly Power GSPS Sampling ADCs in Half the Space." Analog Devices, Inc., October 2018.

About the Author

Bhakti Waghmare is currently a product marketing engineer for the µModule Regulator Power by Linear Product Group based in Santa Clara, California. She supports marketing for µModule regulator power products. Bhakti joined Analog Devices in 2018. She attained a B.S. in mechanical engineering and an M.S. in industrial engineering from Wayne State University, Detroit, MI. She can be reached at *bhakti.waghmare@analog.com*.

About the Author

Diarmuid Carey is an applications engineer with the European Centralized Applications Center based in Limerick, Ireland. He has worked as an applications engineer since 2008 and joined Analog Devices in 2017, providing design support for the Power by Linear portfolio for European broad market customers. He holds a bachelor's of engineering in computer engineering from University of Limerick. He can be reached at *diarmuid.carey@analog.com*.

Rarely Asked Questions—Issue 169 A Simple Way to Measure Temperature Using One GPIO Digital Interface

Chau Tran and Naveed Naeem

Question:

How can I make an analog measurement if I only have a single GPIO left on the FPGA/microprocessor for my system?

Answer:

A voltage-to-frequency converter can be used instead of an analog-to-digital converter.

Abstract

As the need for sensing capabilities becomes more prevalent in modern applications focused on machine health and other Internet of Things (IoT) solutions, so does the need for simpler interfaces with fewer I/Os and smaller device footprints. The density of devices connected to a single microprocessor or FPGA is continuing to increase, while application space—and as a result, the number of I/O pins—can become constrained. In an ideal world, all applications would have an ASIC providing a small integrated solution. However, ASIC development is time consuming, costly, and doesn't provide the flexibility to be repurposed for other uses. As a result, more and more applications are using microprocessors or small form factor FPGAs to get product development done in a timely and cost-effective manner. In this article, we will explore a temperature-to-frequency converter that can provide accurate temperature results while only using a single GPIO pin. It will also demonstrate how using a voltage-to-frequency converter can be adapted to a variety of sensing applications.

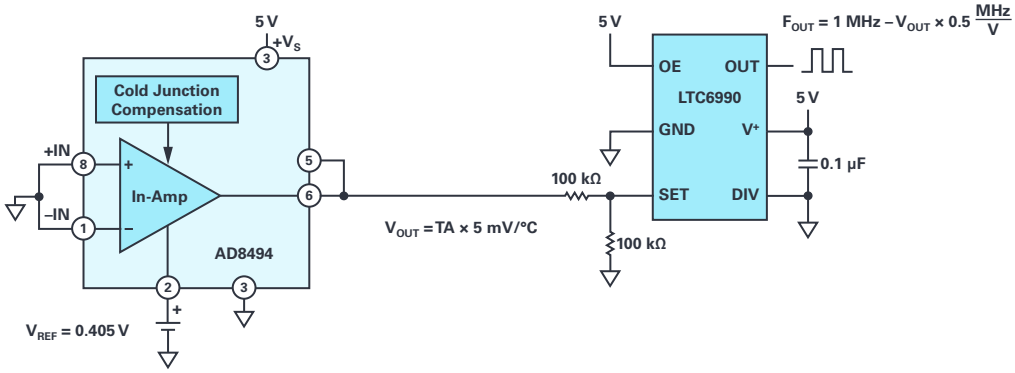

Figure 1. The simple temperature-to-frequency converter.

Motivation

Some sensor measurements, such as temperature, humidity, and barometric pressure, are inherently dc and don't change at rates fast enough—nor do they require resolutions accurate enough—to warrant an ADC and the design considerations that go along with it. Most ADCs require fast, accurate clock generation and timing, a stable voltage reference, a reference buffer with very low output impedance, and analog front-end circuitry to properly signal condition a transducer output before it can be digitally quantized and used by the system monitoring it. In the case of sensing ambient temperature, a discrete application might use a thermistor in a Wheatstone bridge whose output would then be gained up by an instrumentation amplifier and then fed into an ADC. This design is over-engineered and requires more space, power, and computation cycles necessary for an application where the measurement might only need to be made once every 15 seconds.

What alternative measurement solution can be used to reduce the number of components and complexity associated with an ADC signal chain but still measure an analog voltage? The solution is a voltage-to-frequency converter, such as the LTC6990, configured in voltage-controlled oscillator (VCO) mode can be used to measure an analog voltage without an ADC. In this example, the AD8494, a precision thermocouple amplifier, is configured as an ambient temperature sensor whose output voltage serves as the input to the LTC6990, resulting in a temperature-to-frequency converter signal chain.

How to Convert a Temperature Input to a Frequency Output

Today, many modern electronic appliances require an on-board temperature monitoring system. The method of converting analog signals to pulse-width modulated signals or digital signals is well documented. However, if the measurement solution would require an ADC, there are drawbacks associated with cost, accuracy, and speed. Typically, the more accurate the measurement, the more expensive the solution. This circuit delivers a low cost, versatile, and easily interfaceable solution whose accuracy can vary as determined by the needs of the temperature measurement system.

The AD8494 is a precision thermocouple amplifier, but it also can be used as an ambient temperature sensor by shorting its inputs to ground. The output is then defined as:

$$V_{OUT} = T_{ambient} \times \frac{5\ \text{mV}}{^{\circ}\text{C}} \tag{1}$$

In this circuit where a unipolar supply is used, $-V_S$ = ground (0 V) and an offset voltage must be applied to the REF pin of the AD8494 in order to bias the output voltage above ground, even if ambient temperatures are negative.

The output voltage of the temperature sensor, V_{OUT}, is defined as:

$$V_{OUT} = T_{ambient} \times \frac{5\ \text{mV}}{^{\circ}\text{C}} + V_{offset} \tag{2}$$

In VCO mode, the LTC6990 frequency output is defined as:

$$F_{OUT} = \frac{\left(1\ \text{MHz} - \left(\frac{R_{set}}{R_{VCO} + R_{set}} \times V_{CTRL}\right) \times \frac{\text{MHz}}{\text{V}}\right)}{N_{div}} \tag{3}$$

Since the output voltage of the AD8494 is V_{CTRL} to the LTC6990, Equation 1 can be substituted for V_{CTRL} in Equation 2, and setting $R_{SET} = R_{VCO}$ yields the following:

$$F_{OUT} = \frac{\left(1\ \text{MHz} - \left(\frac{1}{2} \times \left(T_{ambient} \times \frac{5\ \text{mV}}{^{\circ}\text{C}} + V_{offset}\right)\right) \times \frac{\text{MHz}}{\text{V}}\right)}{N_{div}} \tag{4}$$

From here, $T_{ambient}$ can now be solved for. The voltage units cancel out, which yields Equation 5:

$$T_{ambient} = \frac{-2((F_{OUT} \times N_{div}) - 1\ \text{MHz}) - (Offset \times 1\ \text{MHz})}{0.005\ \text{MHz}}{}^{\circ}\text{C} \tag{5}$$

Okay, I Have a Frequency Output. How Is This Useful?

The beauty of a frequency output is that you can use a single GPIO pin to get a sensor measurement. If the synchronous counter circuit in Figure 2 is used, then the rising edge of a clock will always be observed at its CLK_IN input. If the F_{OUT} of the LTC6990 is used as the input clock, then the counter will increment every time a rising edge on F_{OUT} is detected, creating a period counter. If the time interval between each measurement is constant, then the number of periods within a given time interval can be counted and frequency can be determined either by using floating point math or a lookup table. By dividing the acquisition time $T_{Acquisitions}$ by the number of periods counted, we get the period of F_{OUT}. Taking the inverse of this relationship results in Equation 6.

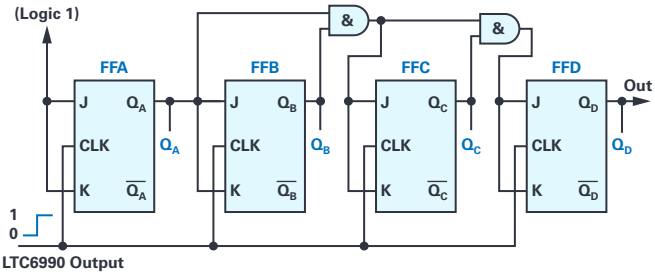

Figure 2. A 4-bit synchronous counter with LTC6990 output as its clock input.

$$F_{OUT} = \frac{Total\ Periods}{T_{Acquisitions}},\ F_{OUT} \geq 2 \times \left(\frac{1}{T_{Acquisitions}}\right) \tag{6}$$

The example Verilog code shows a function that can be used to count the number of periods by using a single GPIO input on an FPGA. The larger the acquisition period, the more accurate the measurement will be. In the case of the following code, a 16-bit counter is used to provide additional resolution. This also assumes the logic that controls the measurement acquisition time is performed at a higher level in the architecture.

```
module tempToFreqCounter (Clk_master, Clk_In, rst, PeriodCount)

input Clk_master, Clk_In,rst;
output PeriodCount;
reg[15:0] PeriodCount;


always @(posedge Clk_master)
begin

      always @(posedge Clk_In)
      begin
      PeriodCount = PeriodCount +1;
      end

      always @ (posedge rst)
      begin
      PeriodCount = 16'b0;
      end

end

endmodule
```

Figure 3. Example Verilog code.

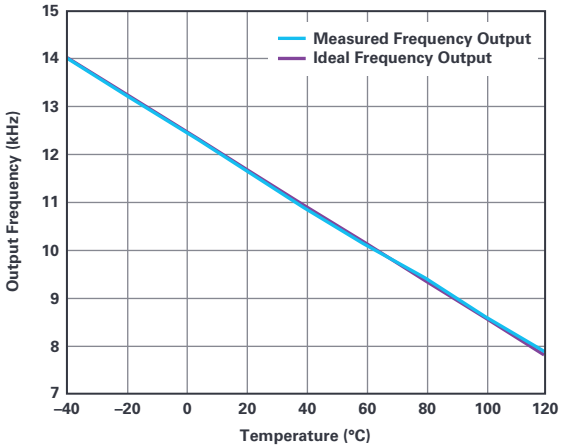

Figure 4. Transfer function of the temperature-to-frequency converter.

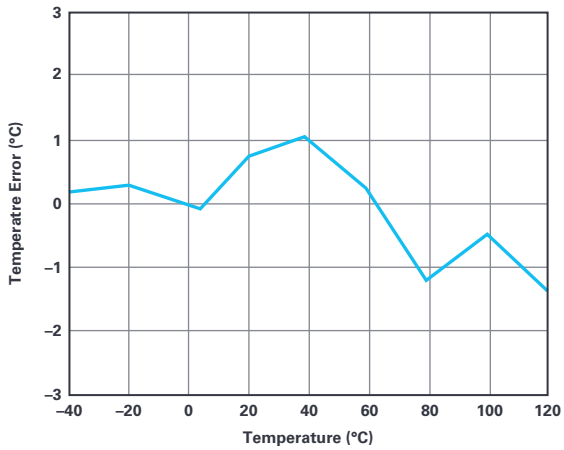

Figure 5. Temperature error.

Conclusion

In this application, a new type of temperature-to-frequency converter is discussed. It provides an accurate, low cost method to measure temperatures. If the temperatures exceed the industrial range of –40°C to +125°C, a thermocouple can be installed at the inputs of the sensor. As a conclusion, the following plot shows the error of the measurement system. It demonstrates the linear relationship between the ambient temperature and the output frequency as well as the accuracy of the system. While this solution may not provide a very fine temperature resolution result, for applications where roughly ±2°C error is acceptable, this provides a cheap and simple interface for measuring temperature. Additionally, the concept of using a voltage-to-frequency converter can be adapted to measure other types of transducer outputs without the need for an ADC.

References

"Fundamentals of the Electronic Counters." Hewlett Packard, 1997.

Johansson, Staffan. "New Frequency Counting Principle Improves Resolution." Proceedings of the 2005 IEEE International Frequency Control Symposium and Exposition, 2005.

Kamal, Ibrahim. "5 Hz to 500 kHz Frequency Meter." IKA Logic, April 2008.

Synchronous Counter. ElectronicsTutorials, 2019.

About the Author

Chau Tran 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. Chau holds more than 10 patents and has authored more than 10 technical articles. He can be reached at *chau.tran@analog.com*.

About the Author

Naveed Naeem joined Analog Devices in 2016 as a product test development engineer. Prior to joining ADI, he worked as an electrical engineer for an automotive sensors company for 3 years. His interests include sensor applications, analog-to-digital converters, audio applications, and software development. In 2013, he graduated with a bachelor's degree in electrical and computer engineering from Worcester Polytechnic Institute. He can be reached at *naveed.naeem@analog.com*.

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

EngineerZone® Support Community

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

Facebook
facebook.com/AnalogDevicesInc

@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

M02000-11/19

analog.com

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™