Analog Dialogue

Volume 53, Number 2, 2019

Your Engineering Resource for Innovative Design

5 Making Sense of Sounds, or How AI Can Boost Your Machines' Uptime

Visit *analogdialogue.com*

ANALOG DEVICES

AHEAD OF WHAT'S POSSIBLE™

In This Issue

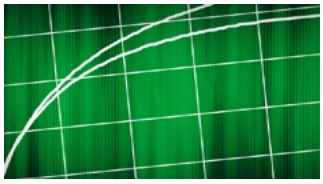

Bernhard Siegel, Editor

Bernhard became editor of *Analog Dialogue* in March 2017. He has been with Analog Devices for over 25 years, starting at the ADI Munich office in Germany. In his current role as the chief technical editor, he is responsible for the worldwide technical article program within Analog Devices.

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

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

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

AnalogDialogue

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

Making Sense of Sounds, or How Analog Devices' AI Can Boost Your Machine Uptime

By **Sebastien Christian**

Introduction

Anyone familiar with the necessity of maintaining a mechanical machine knows how important the sounds and vibrations it makes are. Proper machine health monitoring through sound and vibrations can cut maintenance costs in half and double the lifetime. Implementing live acoustic data and analysis is another important approach for condition-based monitoring (CbM) systems.

We can learn what the normal sound of a machine is. When the sound changes, we identify it as abnormal. Then we may learn what the problem is so that we can associate that sound with a specific issue. Identifying anomalies takes a few minutes of training, but connecting sounds, vibrations, and their causes to perform diagnostics can take a lifetime. There are experienced technicians and engineers with this knowledge, but they are a scarce resource. Instinctively recognizing a problem from sound alone can be difficult, even with recordings, descriptive frameworks, or in-person training with experts.

Because of this, our team at Analog Devices has spent the last 20 years on understanding how humans make sense of sounds and vibrations. Our objective was to build a system able to learn sounds and vibrations from a machine and decipher their meaning to detect abnormal behavior and to perform diagnostics. This article details the architecture of OtoSense, a machine health monitoring system that enables what we call computer hearing, which allows a computer to make sense of the leading indicators of a machine's behavior: sound and vibration.

This system applies to any machine and works in real time with no network connection needed. It has been adapted for industrial applications and it enables a scalable, efficient machine health monitoring system.

This article delves into the principles that guided OtoSense's development, and the role of human hearing in designing OtoSense. The article then discusses the way sound or vibration features were designed, how meaning is derived from them, and the continuous learning process that makes OtoSense evolve and improve over time to perform increasingly complex diagnostics with increasing accuracy.

Guiding Principles

To be robust, agnostic, and efficient, the OtoSense design philosophy followed some guiding principles:

- **Get inspiration from human neurology.** Humans can learn and make sense of any sound they can hear in a very energy efficient manner.
- **Be able to learn stationary sounds as well as transient sounds**. This requires adapted features and continuous monitoring.
- **Perform the recognition at the edge, close to the sensor**. There should not be any need of a network connection to a remote server to make a decision.
- **Interaction with experts and the necessity to learn from them must happen with minimal impact on their daily workload** and be as enjoyable as possible.

The Human Hearing System and Translation to OtoSense

Hearing is the sense of survival. It's the holistic sense of distant, unseen events, and it matures before birth.

The process by which we humans make sense of sounds can be described in four familiar steps: analog acquisition of the sound, digital conversion, feature extraction, and interpretation. In each step, we will compare the human ear with the OtoSense system.

- **Analog acquisition and digitization**. A membrane and levers in the middle ear capture sounds and adjust impedance to transmit vibrations to a liquid-filled canal where another membrane is selectively displaced depending on spectral components present in the signal. This in turn bends flexible cells that emit a digital output that reflects the amount and harshness of the bending. These individual signals then travel on parallel nerves arranged by frequency to the primary auditory cortex.
  - In OtoSense, this job is performed by sensors, amplifiers, and codecs. The digitization process uses a fixed sample rate adjustable between 250 Hz and 196 kHz, with the waveform being coded on 16 bits and stored on buffers that range from 128 samples to 4096 samples.
- **Feature extraction** happens in this primary cortex: frequency-domain features such as dominant frequencies, harmonicity, and spectral shape, as well as time-domain features such as impulsions, variations of intensity, and main frequency components over a time window spanned around 3 seconds.
  - OtoSense uses a time window that we call chunk, which moves with a fixed step size. The size and step of this chunk can range from 23 ms to 3 s, depending on the events that need to be recognized and the sample rate, with features being extracted at the edge. We'll provide more information on the features extracted by OtoSense in the next section.

- **Interpretation happens in the associative cortex**, which merges all perceptions and memories and attaches meaning to sounds, such as with language, which plays a central role in shaping our perceptions. The interpretation process organizes our description of events far beyond the simple capacity of naming them. Having a name for an item, a sound, or an occurrence allows us to grant it greater, more multilayered meaning. For experts, names and meaning allow them to better make sense of their environment.
  - This is why OtoSense interactions with people start from visual, unsupervised sound mapping based on human neurology. OtoSense shows a graphical representation of all the sounds or vibration heard, organized by similarity, but without trying to create rigid categories. This lets experts organize and name the groupings seen onscreen without trying to artificially create bounded categories. They can build a semantic map aligned with their knowledge, perceptions, and expectations regarding the final output of OtoSense. The same soundscape would be divided, organized, and labelled differently by auto mechanics, aerospace engineers, or cold forging press specialists—or even by people in the same field but at different companies. OtoSense uses the same bottom-up approach to meaning creation that shapes our use of language.

From Sound and Vibration to Features

A feature is assigned an individual number to describe a given attribute/quality of a sound or vibration over a period of time (the time window, or chunk, as we mentioned earlier). The OtoSense platform's principles for choosing a feature are as follows:

- **Features should describe the environment** as completely as possible and with as many details as possible, both in the frequency domain and time domain. They have to describe stationary hums as well as clicks, rattles, squeaks, and any kind of transient instability.
- **Features should constitute a set as orthogonally as possible**. If one feature is defined as the average amplitude on the chunk, there should not be another feature strongly correlated with it, as a feature such as total spectral energy on the chunk would be. Of course, orthogonality is never reached, but no feature should be expressed as a combination of the others—some singular information must be contained in each feature.
- **Features should minimize computation**. Our brain just knows addition, comparison, and resetting to 0. Most OtoSense features have been designed to be incremental so that each new sample modifies the feature with a simple operation, with no need for recomputing it on a full buffer or, worse, chunk. Minimizing computation also implies not caring about standard physical units. For example, there is no point in trying to represent intensities with a value in dBA. If there is a need to output a dBA value, it can be done at the time of output if necessary.

A portion of the OtoSense platform's two to 1024 features describe the time domain. They are extracted either right from the waveform or from the evolution of any other feature over the chunk. Some of these features include the average and maximal amplitude, complexity derived from the linear length of the waveform, amplitude variation, the existence and characterization of impulsions, stability as the resemblance between the first and last buffer, skinny autocorrelation avoiding convolution, or variations of the main spectral peaks.

The features used on the frequency domain are extracted from an FFT. The FFT is computed on each buffer and yields 128 to 2048 individual frequency contributions. The process then creates a vector with the desired number of dimensions—much smaller than the FFT size, of course, but that still extensively describe the environment. OtoSense initially starts with an agnostic method for creating equal-sized buckets on the log spectrum. Then, depending on the environment and the events to be identified, these buckets adapt to focus on areas of the spectrum where information density is high, either from an unsupervised perspective that maximizes entropy or from a semisupervised perspective that uses labelled events as a guide. This mimics the architecture of our inner ear cells, which is denser where the speech information is maximal.

Architecture: Power to the Edge and Data on Premises

Outlier detection and event recognition with OtoSense happen at the edge, without the participation of any remote asset. This architecture ensures that the system won't be impacted by a network failure and it avoids having to send all raw data chunks out for analysis. An edge device running OtoSense is a self-contained system describing the behavior of the machine it's listening to in real time.

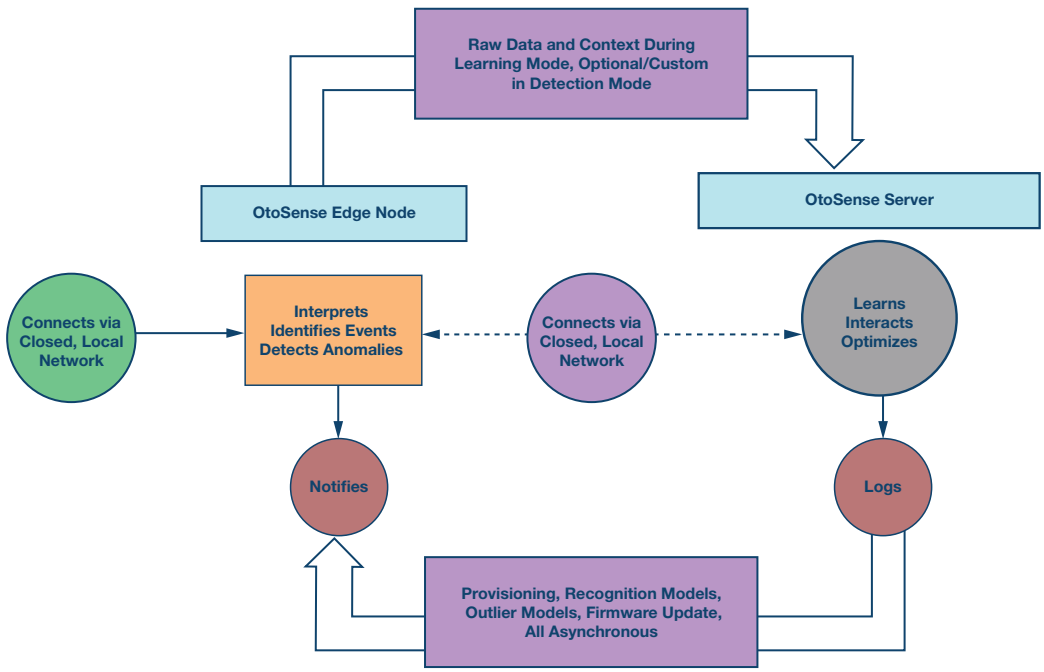

Figure 1. The OtoSense system.

The OtoSense server, running the AI and HMI, is typically hosted on premises. A cloud architecture makes sense for aggregating multiple meaningful data streams as the output of OtoSense devices. It makes less sense to use cloud hosting for an AI dedicated to processing large amounts of data and interacting with hundreds of devices on a single site.

From Features to Anomaly Detection

Normality/abnormality evaluation does not require much interaction with experts to be started. Experts only need to help establish a baseline for a machine's normal sounds and vibrations. This baseline is then translated into an outlier model on the OtoSense server before being pushed to the device.

We then use two different strategies to evaluate the normality of an incoming sound or vibration:

- The first strategy is what we call usualness, where any new incoming sound that lands in the feature space is checked for its surrounding, how far it is from baseline points and clusters, and how big those clusters are. The bigger the distance and the smaller the clusters, the more unusual the new sound is and the higher its outlier score is. When this outlier score is above a threshold as defined by experts, the corresponding chunk is labelled unusual and sent to the server to become available for experts.
- The second strategy is very simple: any incoming chunk with a feature value above or below the maximum or minimum of all the features defining the baseline is labelled as extreme and sent to the server as well.

The combination of unusual and extreme strategies offers good coverage of abnormal sounds or vibrations, and these strategies perform well for detecting progressive wear and unexpected, brutal events.

From Features to Event Recognition

Features belong to the physical realm, while meaning belongs to human cognition. To associate features with meaning, interaction between OtoSense AI and human experts is needed. A lot of time has been spent following our customers' feedback to develop a human-machine interface (HMI) that enable engineers to efficiently interact with OtoSense to design event recognition models. This HMI allows for exploring data, labelling it, creating outlier models and sound recognition models, and testing those models.

The OtoSense Sound Platter (also known as splatter) allows for the exploration and tagging of sounds with a complete overview of the data set. Splatter makes a selection of the most interesting and representative sounds in a complete data set and displays them as a 2D similarity map that mixes labelled and unlabelled sounds.

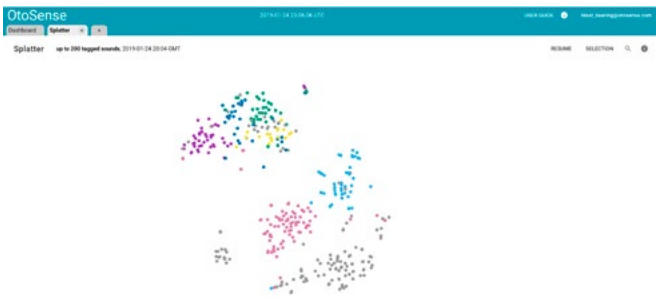

Figure 2. A 2D splatter map of sound in the OtoSense Sound Platter.

Any sound or vibration can be visualized, along with its context, in many different ways—for example, using sound widgets (also known as swidgets).

Figure 3. OtoSense sound widgets (swidgets).

At any moment, an outlier model or an event recognition model can be created. Event recognition models are presented as a round confusion matrix that allows OtoSense users to explore confusion events.

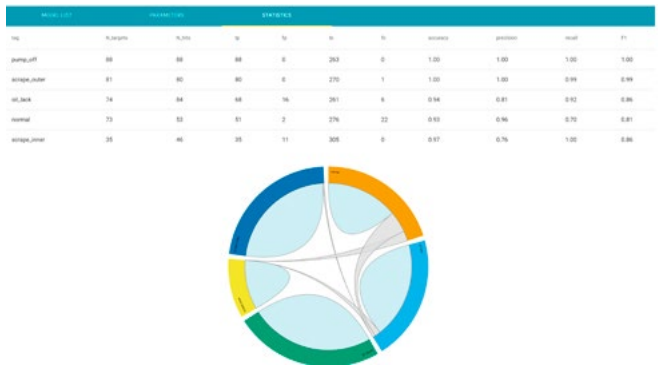

Figure 4. An event recognition model can be created based on the required events.

Outliers can be explored and labelled through an interface that shows all the unusual and extreme sounds over time.

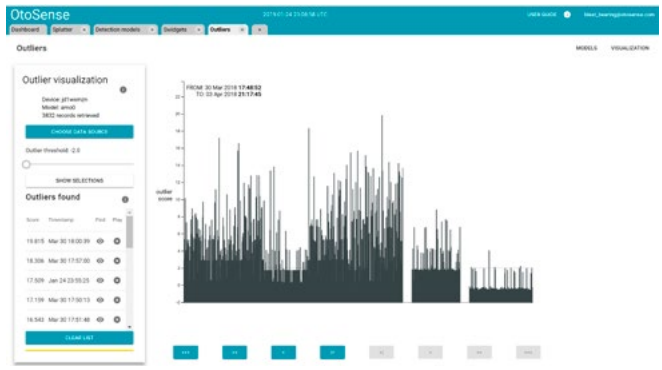

Figure 5. Sound analytics over time in the OtoSense Outlier visualization.

The Continuous Learning Process, from Anomaly Detection to Increasingly Complex Diagnostics

OtoSense has been designed to learn from multiple experts and allow for more and more complex diagnostics over time. The usual process is a recurring loop between OtoSense and experts:

- An outlier model and an event recognition model are running at the edge. These create output for the probability of potential events happening, along with their outlier scores.
- An unusual sound or vibration above the defined threshold triggers an outlier notification. Technicians and engineers using OtoSense can then check on the sound and its context.

- These experts then label this unusual event.
- A new recognition model and outlier model that includes this new information is computed and pushed to edge devices.

Conclusion

The objective of the OtoSense technology from Analog Devices is to make sound and vibration expertise available continuously, on any machine, with no need for a network connection to perform outlier detection and event recognition. This techonology's growing use for machine health monitoring in aerospace, automotive, and industrial monitoring applications has shown good performance in situations that once require human expertise and in situations involving embedded applications, especially on complex machines.

References

Christian, Sebastien. "How Words Create Worlds." TEDxCambridge, 2014.

Sebastien Chistian [sebastien.christian@analog.com] had an early passion for understanding how we humans build an inner, sharable model of the world, using our senses, and how we use this model to describe the world they live in.

Sebastian earned an M.S. in quantum physics, which he followed with an M.S. in neuroscience and a third degree in semantics. Sebastien's education combined research, development, and field experiments. He worked as a speech and language pathologist with psychotic and deaf children for 10 years, refining his understanding of sensor-based meaning creation and sharing, with an emphasis on hearing. Sebastien says that this practice, where he worked with the same young patients for years, is what brought all the scattered pieces of knowledge together into a single, coherent picture.

During the same period, Sebastien became an expert for the French Ministry of Health, where he advised on hearing loss policies, taught in medical school and at Paris Sorbonne University, and, in 2011, created the first independent private R&D laboratory dedicated to bringing AI driven innovations to people with sensing and cognitive disabilities.

In 2013, Sebastien completed a full prototype of his machine hearing project, which earned him laureate of the NETVA tech competition in Cambridge, MA. The massively positive feedback from his fellows at MIT and from the business community led him to found OtoSense in early 2014, and to develop what is the first AI focused on making sense of any sound. This machine hearing platform revealed itself to be well adapted to complex environments and complex machine monitoring.

After receiving multiple awards, which included a Best App of the Year award at GSMA Mobile World Congress in 2015, OtoSense focused on machine monitoring in the industrial and transportation verticals, with vast and growing range of potential applications ahead.

Sebastien is now leading OtoSense inside product development at Analog Devices.

Sebastien Christian

Sigma-Delta ADC Clocking—More Than Jitter

By **Pawel Czapor**

One of the key benefits of modern SAR and sigma-delta (Σ-Δ) analog-to-digital converters (ADCs) is that they are designed with ease of use in mind, where ease of use was an afterthought for previous generations. This simplifies the task for system designers and, in many instances, allows for a single reference design to be used and recycled for multiple generations and across a great variety of applications. In many instances, it allows you to build one reference design that can be used for different applications over a long period of time. The hardware of a precision measurement system stays the same while the software implementation adapts to the different system needs. That's the beauty of reusability, but nothing in life is entirely advantageous—there's always a penalty. The primary drawback to having a single design for multiple applications is that you forego the customizations and optimizations necessary to achieve the absolute highest possible performance for dc, seismic, audio, and higher bandwidth applications. In the rush to reuse and complete designs, precision performance is often sacrificed. One of the primary oversights and areas of neglect is in clocking. In this article, we will discuss the importance of the clock and offer guidance on proper designs for high performance converters.

ADC Fundamentals

Relation Between Jitter and Signal-to-Noise Ratio

When looking at the available literature, the dependency of ADC performance on the jitter specification is well described and, usually for good reason, such titles includes the words "high speed."[1] To examine the relationship between jitter and signal-to-noise ratio (SNR), the starting point is the relationship between the SNR figure and the rms jitter.

If jitter is the primary source of noise in the system, this relationship simplifies to:

$$SNR = -20 \times \log_{10}(2\pi f_{IN}\delta t_{RMS}) \quad (1)$$

If there are different sources of noise contributing, you need to use Equation 2 to calculate the combined SNR:

$$SNR = 10 \times \log_{10}\left[\left(\frac{A}{\sqrt{2}}\right)^2 / \left\{\left(\sqrt{2} \times \pi f_{IN} \times A \times \delta t_{RMS}\right)^2 + e_v^2\right\}\right] \quad (2)$$

Where:

A, Fin—parameters of input signal, amplitude, and input frequency

e_v is simplified voltage noise rms

δt_{RMS} is total rms jitter estimated as the rms sum of various contributions:

$$\delta t_{RMS} = \sqrt{\delta t_{INT}{}^2 + \delta t_{EXT}{}^2} \quad (3)$$

For an in-depth tutorial of Equation 3 usage, see: *analog.com/MT-008*.

The summing is valid on uncorrelated noise sources. With Equation 2, we show the SNR depending on thermal noise (e^2v) and jitter noise. The jitter contribution to SNR is dependent on the input frequency (f_{IN}). This means at higher frequencies, the SNR is mainly defined by the jitter. Figure 1 includes the curves from Equation 1 and Equation 2 for ideal and real-life ADCs affected by jitter. Plots such as Figure 1 are common for high speed ADC data sheets, but they generally start in the MHz range. For precision ADCs, we will show the same dependencies further down in the kHz ranges. We are pushing SNRs in excess of 108 dB (see Figure 1), which precision ADCs are capable of nowadays. This is where the AD7768-1 comes in handy.

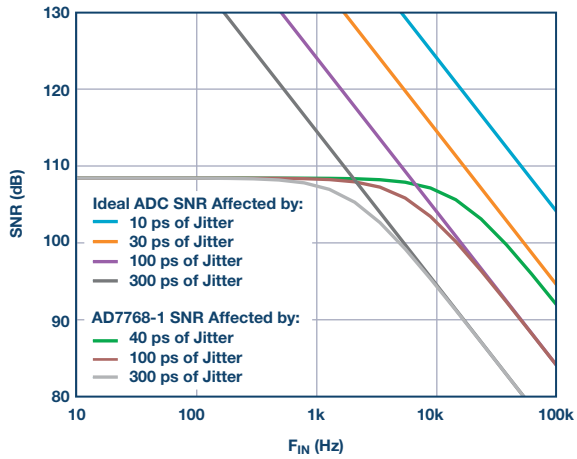

Figure 1. SNR vs. f_{IN} at different jitter levels.

Upon review of the plot in Figure 1, you can see that the AD7768-1 converting a 1 kHz signal (gray line) will be affected by clock jitter only if σt_{RMS} exceeds 300 ps. We can rearrange variables and show jitter requirement for specific ENOB and f_{IN}:

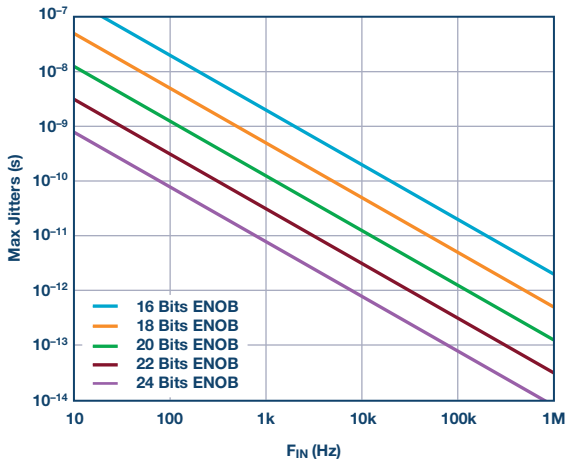

Figure 2. Maximum allowed jitter vs. f_{IN} at different ENOB of converter.

Target jitter for today's high precision converters will prevent the designer from using common relaxation oscillators (like the 555 timer-based oscillator) or many microcontroller or FPGA-based clock generators. This leaves us with crystal (XTAL) and phase-locked loop (PLL) oscillators. New technological advances in MEMS oscillators will be suitable as well.

Can Oversample Techniques Help Here?

An important observation in Equation 1 and Equation 2 is that there is no explicit dependency on the sampling frequency. This tells us that the contribution of jitter will be hard to mitigate with oversampling techniques (plain or noise shaped). Oversampling is very common in high precision systems, but offers almost nothing to fight in terms of jitter noise. Relation to the sampling frequency can be found in Equation 4:

$$\delta t_{RMS} = \frac{\sqrt{2\int_{f_{min}}^{f_{max}} L(f)}}{2\pi f_s} \quad (4)$$

Where:

L(f) is the phase noise spectral single-sideband (SSB) density function

f_{min} and f_{max} are the frequency span relevant to the particular measurement.

For an in-depth tutorial of Equation 4 usage, see: *analog.com/MT-008*.

Generally, jitter contribution should be considered only poorly improved by increasing f_S.[2] In theoretical discussions, the oversampling ratio of an ADC offers some reduction of broadband jitter contributions.[3] For quantization and thermal noise, noise shaping is a very efficient way of suppressing noise in the band of interest. Increasing oversampling ratios suppresses quantization noise a lot faster (Equation 5) than noise jitter suppression as shown in Equation 6. This makes jitter stand out even more in oversampling structures that utilize noise shaping. In Nyquist converters, this might not pronounce itself as severely. Figure 3 illustrates this phenomenon using the example of a second-order Σ-Δ ADC and a new, forth-order Σ-Δ ADC.

The relation between quantization noise shaped by an N[th]-order shaper with base error Δ at oversampling ratio *M*:

$$S_{QUANTIZATION} \approx \frac{\pi^{2N}}{(2N+1)} \times \frac{1}{M^{2N+1}} \times \frac{\Delta^2}{12} \quad M >> 1 \quad (5)$$

The relation between the oversampling ratio *M* and the amount of jitter:

$$S_{jitter} \approx \frac{(2 \times \pi f_{IN} \times \delta t_{RMS})^2}{M} \times \frac{\Delta^2}{8} \quad (6)$$

Equation 7 shows the second-order noise shaping (N = 2). Your attention should go to *M* as it changes now with the power of 5.

$$S_{QUANTIZATION(N=2)} \approx \frac{\pi^4}{5} \times \frac{1}{M^5} \quad \frac{\Delta^2}{12} \quad M >> 1 \quad (7)$$

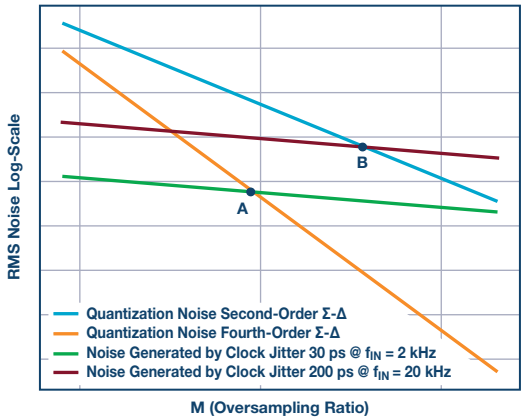

Figure 3. Oversampling reduces quantization noise below jitter. Point (A) shows a fourth-order Σ-Δ ADC requiring a clock with a jitter below 30 ps. Point B illustrates older technologies with second-order shapers immune to jitter up to 200 ps levels converting 20 kHz.

Universal relations will be seen on different generations of converters. A first-order noise shaper will hide jitter for the longest time, progressing with a cubic relation to ~1/M[3], whereas a fourth-order Σ-Δ will get us to the relation of ~1/M[9]. Jitter, at best, will be reduced by 1/M, and this generously assumes the presence of strong wideband frequency components as opposed to the relation of 1/(f[N]).

Will the Amplitude of the Signal Change Things?

Equation 2 shows that amplitude is in both the numerator and the denominator, preventing good trade-off between the amplitudes and the SNR figure. One can make the SNR worse with attenuated signal where, in addition to jitter, thermal noise starts to limit the dynamic range. So, we can see new precision ADCs will be exposed to jitter restrictions in almost all but dc/seismic applications if pushed to achieve low enough noise.

Clock Jitter Will Have a Spectrum, Too

In the introduction, we established the relationship between the signal, the overall voltage noise, and the rms of clock jitter. The SNR figure connects those three in a fairly straightforward Equation 2. The SNR figure is a good benchmark to compare circuits, but it doesn't necessarily determine usability in actual applications. In many applications, designing especially for SNR isn't good enough. For those interested in those specifications, spurious-free dynamic range (SFDR) becomes the design target. In new high precision systems, 140 dB or even 150 dB of SFDR is achievable.

The process where a signal is distorted by a clock source can be examined by looking at it as mixing both. For analysis in frequency domain, FM modulation theory is employed.[3] The resulting fast Fourier transform (FFT) spectrum is a product of mixing clock source spectrum with the input signal spectrum. To review how our ADC is affected by this, we introduce phase noise. Jitter and phase noise describe the same phenomena, but, depending on the application, one will be preferred. We have already shown

how translation between phase noises into jitter figures can be done in Equation 3. In the integration process, spectral nuances will be lost.

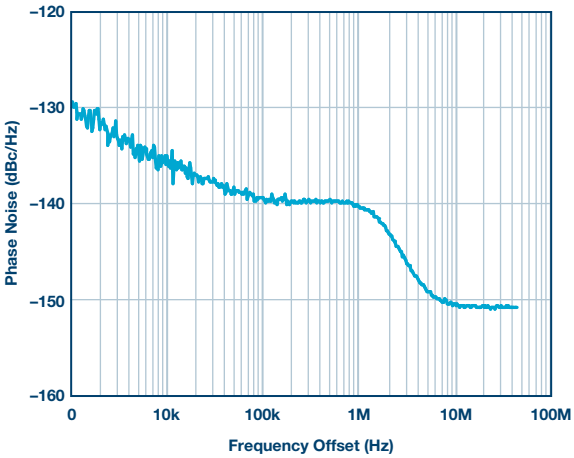

Figure 4. Phase noise density plot for the 100 MHz/33.33 MHz clock generator AD9573.

Phase noise density plots are commonly supplied with clock source equipment and PLL specifications. Plots such as Figure 4 become more scarce for lower frequency sources, which are used in current oversampling converters and total jitter (rms or peak) is reported instead.

Resistor and transistor elements can be forced to exhibit fairly flat noise behavior near dc with chopping schemes. There is no equivalent of chopping for clock circuits available.

When converting a high amplitude A_{IN} signal, the resulting FFT becomes an FM modulated spectrum, where A_{IN} acts as the carrier and the clock sidebands are equivalent to the signal. Note that phase noise will not be band-limited in your FFT and noise will simply deposit multiple alias contributions in slices (see Figure 6).

In precision ADCs, one can usually rely on the natural decaying nature of phase noise and not provide any clock antialiasing filter. There is some scope for jitter reduction by adding filtering to the clock source—for example, using a tuned transformer in the clock path to exhibit a desirable frequency response. Finding out upper integration bound for integral frequency (Equation 4) is not easy to pinpoint. Precision ADC data sheets do not provide much advice on this. In those circumstances, engineering assumptions are made about clock CMOS inputs.

A more common problem in precision ADCs happens very near the f_{IN} frequency where a $1/(f^N)$ shape of phase noise will deteriorate SFDR. A large A_{IN} signal will act as a blocker—a term more popular in radio receivers, which is also applicable here.

When aiming to record high precision spectrums with very long capture times, SFDR will suffer greatly due to the nature of clock phase noise spectral density. The SNR and a visual FFT plot can be improved by shorter capture times (wider frequency bins). For a given FFT capture, rms jitter should be counted as integrated phase noise from ½ of the bin frequency. This becomes obvious upon review of Figure 5.

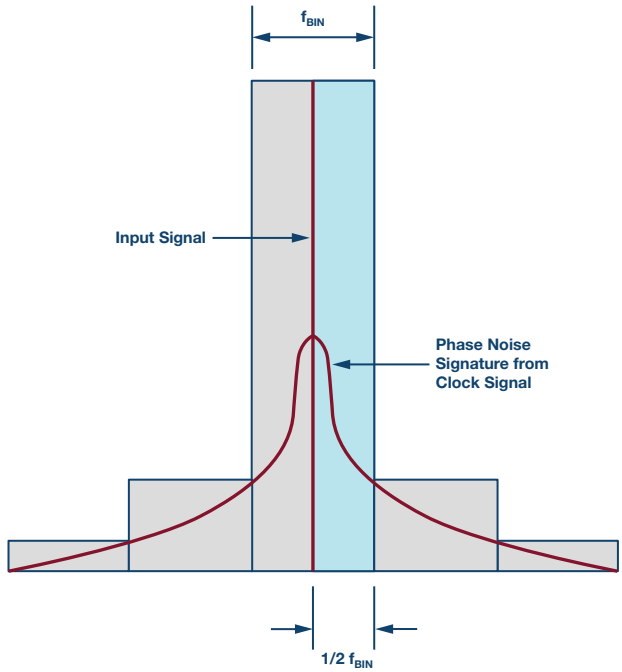

Figure 5. Close-in phase noise determines amplitude of the FFT bins around the primary bin.

While this trick may visually improve FFT plots and SNR figures, it will do nothing for the observation of signals near the blocker. An important generalization and simplification of FM modulation equations is that the heights of skirts are proportional to the ratio in Equation 8:

$$A \approx 10 \times \log_{10} \frac{f_{IN}}{f_S} \quad (8)$$

Elongating integration time for a single FFT hit is an uphill battle to collect further and more pronounced sections of phase noise. One will need to consider alternative ways of combining longer captures to improve this.

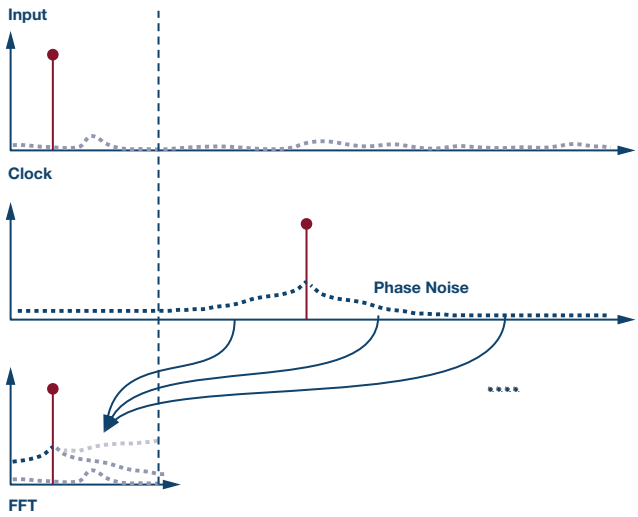

Figure 6. Phase noise aliasing down to baseband.

For practical purposes, SSB plots should be compared at a single point at $f_{BIN}/2$ offset frequency to pick a better source for clean, close-in spectrum and SFDR. If comparing sources to achieve better SNR, then integration in Equation 4 needs to be performed from $f_{BIN}/2$ to more than $3\times f_S$ (jitter aliases).

Sigma-Delta Modulators' Sensitivities to Clocks

The previous topics are universal for any ADC regardless of architecture and technology. The following topic will deal with challenges presented by specific technologies. One of the most prominent examples of jitter dependency is inside Σ-Δ ADCs. The distinction between discrete-time and continuous-time operation of modulators will have tremendous influence on jitter immunity.

Continuous-and discrete-time Σ-Δ ADCs suffer not only due to sampling-related jitter contribution, but also the fact that their feedback loops can be significantly disrupted by jitter. Linearity of DAC elements in both discrete-time and continuous-time modulators is key to achieving high performance. Intuitive understanding of the DAC's significance can be illustrated by drawing parallel to an operational amplifier (op amp). If one is tasked with designing a voltage amplifier with gain equal to 2, the first draft of anyone with a fundamental understanding of circuit design will be an op amp and two resistors. If external circumstances are not extreme, the circuit shown in Figure 7a will do its job. For the most part, the circuit designer doesn't have to understand op amps in order to achieve great performance. The designer has to pick resistors that are well matched and precise enough to achieve the right gain. For noise purposes, they have to be small. For thermal behavior, the thermal coefficients need to match. Note that none of these dependencies are dictated by the op amp. Op amp nonidealities are secondary for this circuit operation. Yes, the influence of input current or capacitive load can be devastating. The ability to slew needs to be reviewed as there might be noise contributions to consider if the bandwidth is not limited. But you only get to fix those problems if you haven't stunted your performance by choosing the wrong resistors. In Σ-Δ ADCs, feedback is more complicated than two resistors—in those circuits, we have DACs instead of resistors performing the corresponding function. Flaws in DAC operation are very detrimental while the remainder of the circuit will reap the benefits of loop gain in a manner similar to op amp circuits.

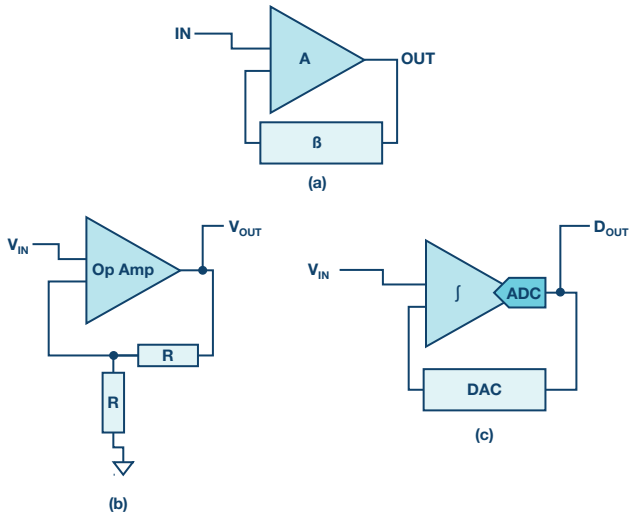

Figure 7. An op amp compared to a Σ-Δ ADC.

ADCs employ element shuffling, or calibrations, which provides a way to deal with mismatches of DAC elements. Those will move errors into high frequency, but will also use a lot more timed events with the potential to increase jitter-related deterioration. This leads to a situation where the noise floor will be polluted by jitter contribution, reducing the effectiveness of noise shaping. Since modulators can employ different DAC schemes—and their mixes—such as return to zero and half return to zero. It is beyond the scope of this article to drill down into analysis and numerical simulations of those schemes.

With regards to jitter in this article, we will limit ourselves to pictorial simplifications. Since jitter dependency problems are within ADC loops, some new designs will provide frequency multipliers on silicon that are designed with an appropriate amount of phase noise. While this takes away a chunk

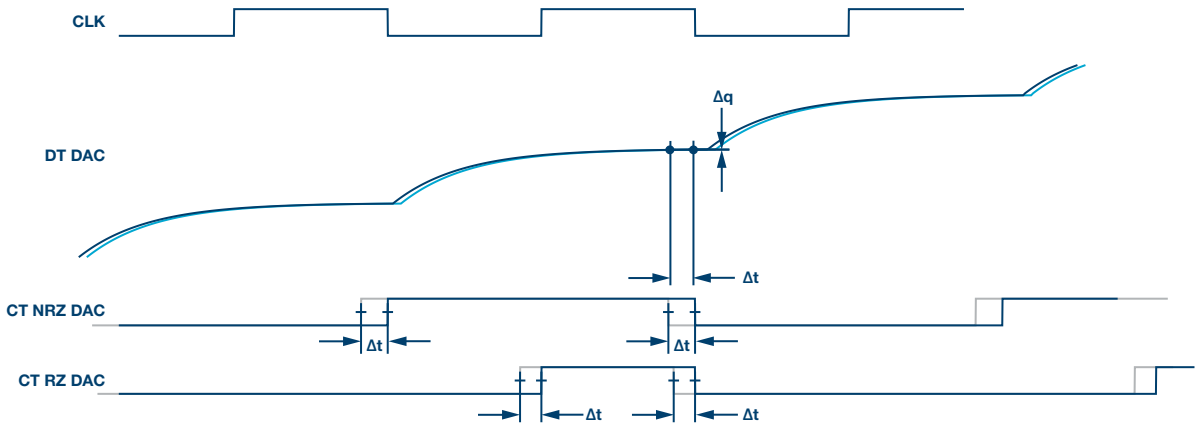

Figure 8. A discrete-time DAC is somewhat immune to jitter, whereas in a continuous-time DAC, narrowing pulses will create significant performance dependence on jitter.

of the work from system designers, please note that the frequency multipliers still rely on good external clocks and low noise power supplies. In those systems, one should consider reviewing PLL literature to see potential threats to observed phase noise. Figure 8 provides a visual illustration showing different DACs' immunity to jitter, showing exponentially smaller dependence when operating a discrete-time DAC.

Modern continuous-time Σ-Δ designs include on-board PLLs. Since timing is carefully tuned in those with agreement to passive elements, they do not offer a wide range of clock speeds. There is a somewhat artificial way of broadening the selection of ADC conversion rates that employs sample-rate conversions. While sample-rate conversions are not neutral on power dissipation with digital circuit advancements, those became affordable alternatives to highly tuned analog circuits. Analog Devices provides a number of ADCs providing sample rate conversion options.

Architecture Utilizing Switched Capacitor Filters

Another specific area where precise timing might influence your performance is switched capacitance filtering. When designing a precision ADC, one needs to make sure that all unwanted signals are excluded or sufficiently attenuated. The ADC might offer specific embedded analog and digital filtering. While an ADC's digital filtering will be immune to jitter, any form of clocked analog filtering will have jitter dependency.

This is particularly important when precision converters employ more advanced front-end switching. While the theory of switched capacitance filters could be beneficial, we will only reference the compendium for further research and analysis.[3]

One of the schemes common in converters is correlated double sampling (CDS). See Figure 9 to see how the performance of CDS rejection quality changes with a clock at three different quality levels. The plot shows signals near rejection band. A switched capacitor filter centered at 1 on the x-axis is shown. The center of the plot does not get suppressed by digital filtering and is dependent on the analog switch capacitance filter. A good quality clock is required to preserve a decent rejection level. Even for measuring dc signals, jitter can destroy noise performance by aliasing down unwanted signals that were supposed to be filtered by switched capacitance filters on silicon. The existence of on-board switched capacitance filters may not be explicitly mentioned in data sheets.

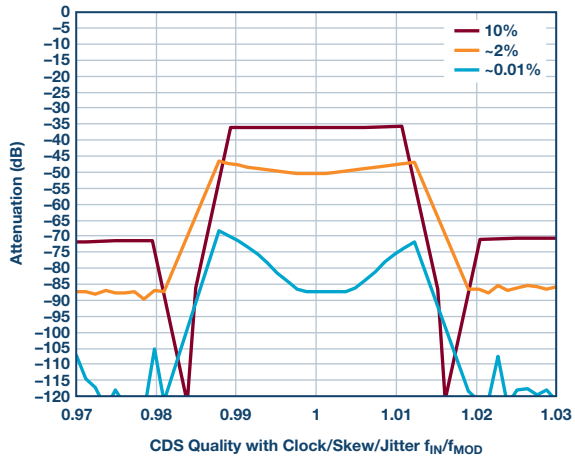

Figure 9. Switch capacitance filtering performance vs. quality of clock—mark-space ratio.

Practical Guide, Sources of the Problem, and the Usual Suspects

Now that we have shown a couple of ways that clocks will add to your troubles, it is time to look at techniques to help you build a system that minimizes the amount of jitter.

Clock Signal Reflections

A high quality clock source can have very sharp rise and fall times. This has the benefit of reducing jitter noise at transition time. Unfortunately, with the benefit of sharp edges comes quite stringent demand for proper routing and termination. If the clock line is not terminated properly, the line will suffer from the reflected waves added to the original clock signal. This process is very disruptive and associated jitter levels can easily account for hundreds of picoseconds. In extreme cases, the clock receiver is capable of seeing additional edges that can potentially lead to locked out circuits.

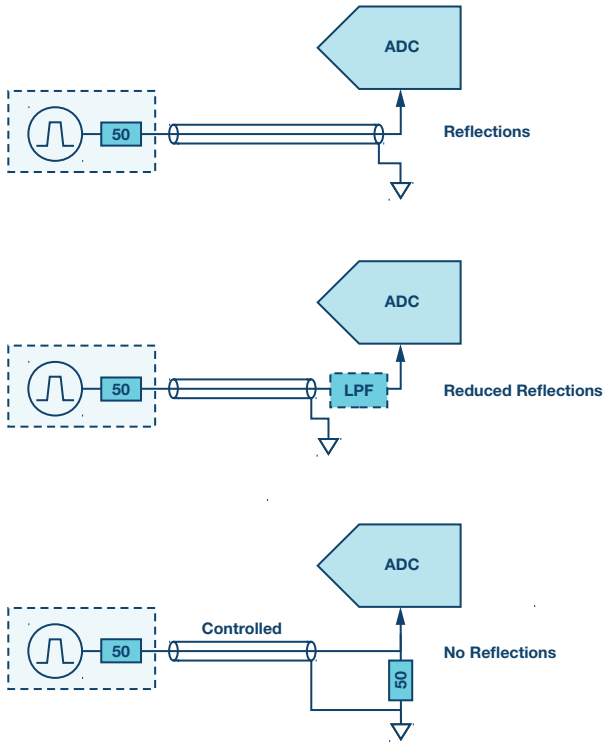

Figure 10. Bad, better, and best circuit designs (in descending order) on clock.

One of the methods that might be counter-intuitive is to slow down the edges with an RC filter, removing high frequency content. One can even use a sine wave as the clock source while waiting for the new PCB with 50 Ω track and termination. While the transition is relatively gradual, and the mark-space ratio can be skewed by hysteresis in digital input, this will reduce the reflection component of jitter.

Power Supply Noise

A digital clock might be routed inside the ADC through a variety of buffers and/or level shifters before the edge is delivered to the sampling switch. If the ADCs have analog supply pins, level shifters are employed and can be sources of jitter. Commonly, the analog side of a chip will have higher voltage devices with longer slewing times, thus jitter sensitivity rises. Some state-of-the-art devices split further analog power supplies between clocked and linear circuits on board.

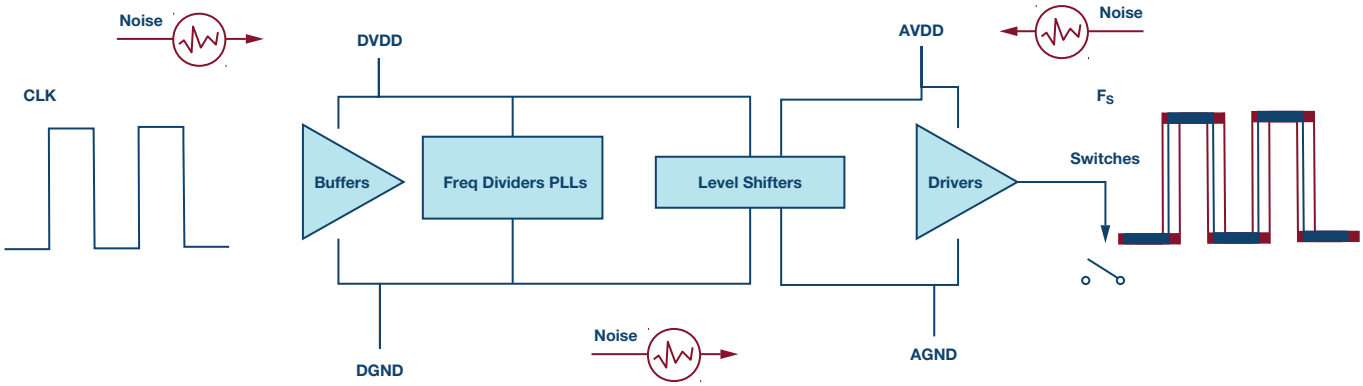

Figure 11. Sampling time disturbed by noise introduced by different power domains in DVDD, AVDD, and between AGND and DGND.

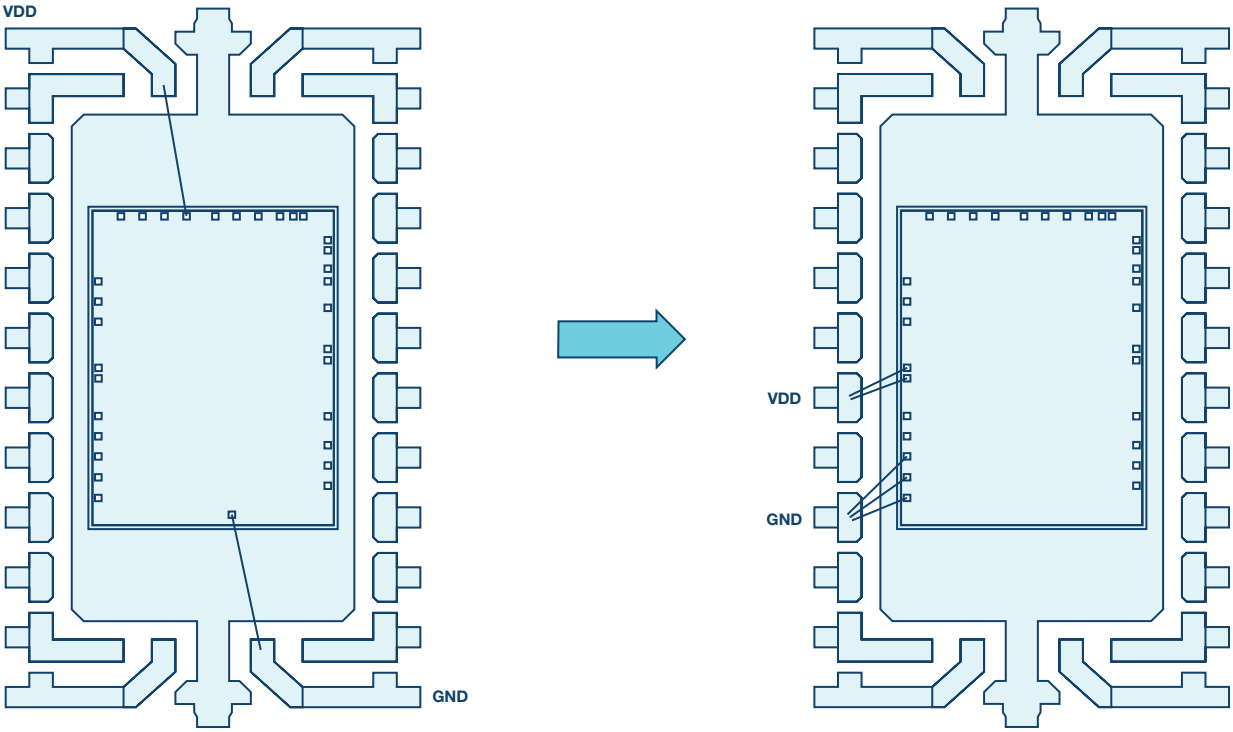

Figure 12. Power delivery schemes for linear circuit (left) and clocked circuit (right).

Decoupling Capacitor: Get the Right One

Jitter sourced by supply noise will be reduced or magnified by the quality of decoupling. Some of the Σ-Δ modulators will have heavy digital activity on the analog and digital sides. This could lead to noncharacteristic spurs with signal or digital data dependent interference. High frequency charge delivery should be limited to a short loop near the device. To accommodate the shortest bondwires, good designs use center pins along the elongated side of the chip. These restrictions are not common problems for amplifiers and low frequency chips, which can have V_{DD} and V_{SS} pins at the corners as in the left side of Figure 12. PCB design should take advantage of those features and keep good quality capacitors near pins.

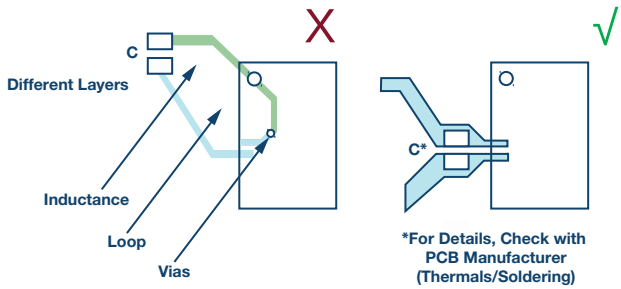

Figure 13. Incorrect (left) and correct (right) location of decoupling capacitor for lower jitter.

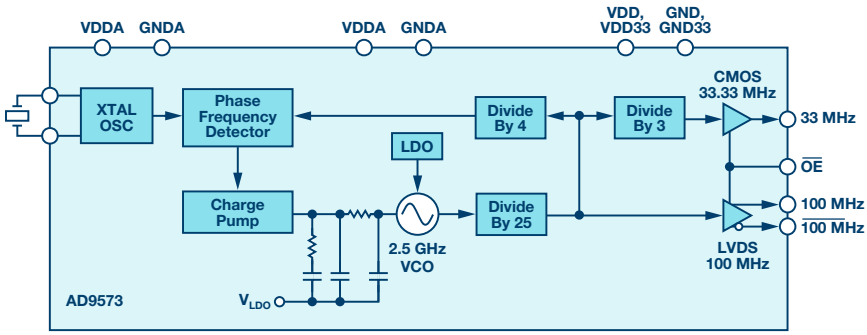

Figure 14. Detailed block diagram of the AD9573.

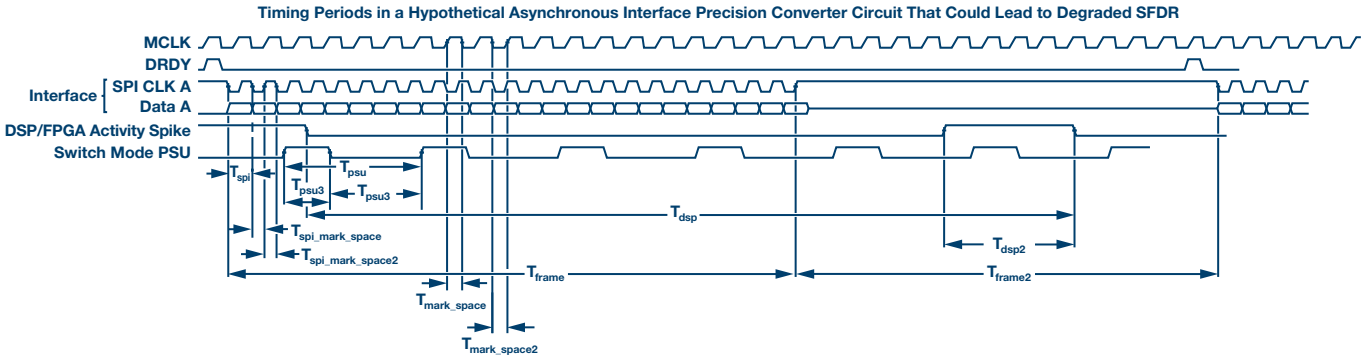

Figure 15. The presence of asynchronous communications and clocks is asking for trouble and investigative work for mix spurs.

Clock Dividers and Clock Signal Isolators

Faster clocks have less jitter, so if power constraints allow it, the use of dividers externally or internally to deliver a desired sampling clock can improve things. When designing a system with isolators, review their pulse widths. If there is a poor mark-space ratio, the skew can interfere with analog performance and, in extreme situations, can lock up the digital side of an IC. In precision ADCs, you might not need an optical fiber clock, but using higher frequencies can provide your final bit of performance. In Figure 14, AD9573 uses 2.5 GHz internally only to provide clean 33 MHz and 100 MHz for the same reasons. If there is no need for precise synchronisation between ADCs, the crystal circuit can be very robust with single-digit ps jitter. For precision ADCs, the crystal amplifier translates to better than 22 bits of performance at a 100 kHz input. This performance is hard to beat and explains why XTAL oscillators are here to stay for the foreseeable future.

Crosstalk from Other Signal Sources

Another source of jitter is related to clock disturbances originating in external lines. If the clock source is incorrectly routed near signals capable of coupling, it can have devastating effects on performance. If the interferer is unrelated to ADC operation—and random—it will add to your jitter budget rather gracefully. If the clock is polluted with ADC-related digital signals, one will observe spur. For slave ADCs, CLK lines and SPI lines can be independent clocks, but this can cause problems at frequencies defined in Equation 9 and aliased back to the first Nyqist zone.

$$N \times f_S \pm M \times f_{SPI} \tag{9}$$

It is advisable to use frequency-locked SPI and MCLK sources. Even with this precaution, SPI and MCLK sources can have spurs associated with the pulse duty cycle of a given clock. For example, if an ADC is decimating by 128 and an SPI reads only 24 bits, this introduces some risk of creating beat frequency relating to specific 1/(24t) and 1/(104t) measurements. Therefore, you should keep the MCLK away from locked SPI lines, as well as data lines.

Interface and Other Clocks

In Figure 15, a variety of timing periods are marked, which can easily disturb SFDR or contribute to jitter. When SPI communications are not frequency locked to the MCLK, spurs can occur. Mastery of layout techniques is your biggest asset in mitigating this problem. Frequencies present themselves as aliased down interferers, but also as beat frequencies and intermodulation products. For example, if the SPI is run at 16.01 MHz and the MCLK is at 16 MHz, one could expect spur at 10 kHz.

Outside of good layout, another way of reducing spurs is to move them outside the band of interest. If a MCLK and SPI can be frequency locked, a lot of disturbances can be avoided. Even then there is still the problem of idling periods in SPI contributing to how busy grounds are, which can still cause disturbance. You can use interface features to your advantage. Interface features in ADCs can offer status bytes or cyclic redundancy checks (CRCs). This may provide a great way to suppress spurs with the added benefit of those functions. Idle clocks—and even unused CRC bytes—can be beneficial to fill data frames evenly. You might choose to disregard the CRCs and still get the benefit of turning them on. Of course, this means extra power on digital lines (Figure 18).

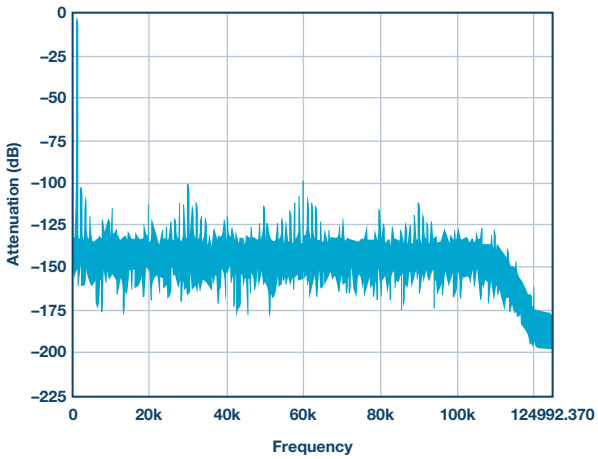

Figure 16. An MCLK routing running too close to a switch-mode PSU.

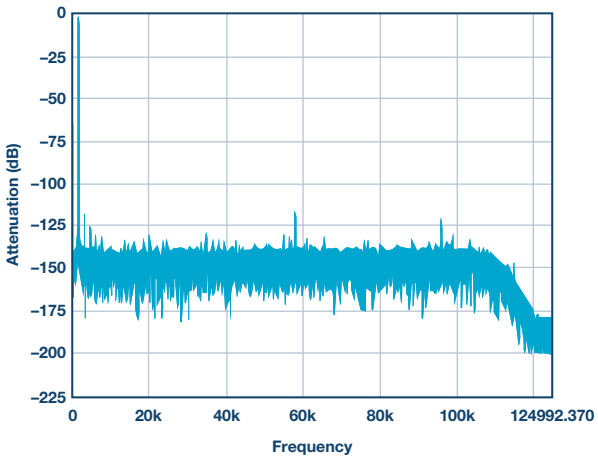

Figure 17. A locally sourced MCLK with an XTAL amplifier with SPI-related spur.

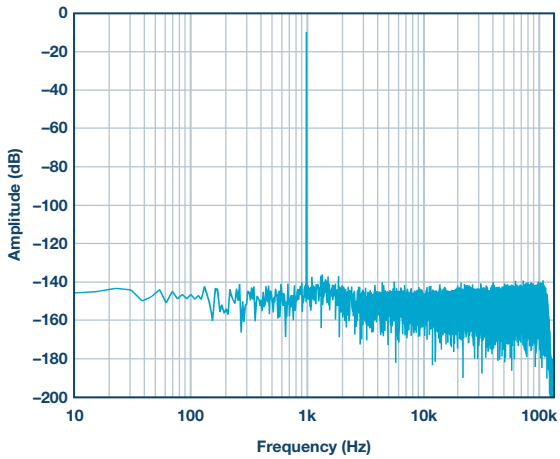

Figure 19. Spectrum of the AD7768-1 with properly designed PCB and clock circuit.

Conclusion

In 2018, ADI released the AD7768-1, a very high precision ADC with sub-100 µV of offset and flat frequency response all the way to 100 kHz. It has been successfully designed into systems capable of SFDR in excess of 140 dB where jitter has been proven negligible beyond audio bands with full scale input. It contains an on-board RC oscillator capable of providing reference points to debug disturbed clock sources. This internal RC, while not providing low jitter, can offer differentiation methods to uncover spur sources. The ADC implements internal switched capacitance filtering techniques, but also uses a clock divider to relieve pressure on the antialias filter. The internal clock divider ensures consistent performance that enables operation with skewed clocks commonly received from isolators. The supply positions are ideal for limiting external ESR/ESL effects with short internal bonds. Glitch rejection is implemented in clock input pads. Performance sweeps with applications boards show performance indicating jitter in 30 ps rms, which should satisfy the broad spectrum of applications. If you are tasked with measuring 140+ dB of SFDR, AD7768-1 might be your fastest way of getting the measurement at a fraction of the power previously needed with convenient power supply rails.

References

[1] Derek Redmayne, Eric Trelewicz, and Alison Smith. *Understanding the Effect of Clock Jitter on High Speed ADCs*. Analog Devices, Inc., 2006.

[2] Bernhard. E. Boser and Bruce A. Wooley. "The Design of Sigma–Delta Modulation Analog-to-Digital Converters." *IEEE Journal of Solid-State Circuits*, Vol. 23, No. 6, December 1988.

[3] Steven Harris. "The Effects of Sampling Clock Jitter on Nyquist Sampling Analog-to-Digital Converters, and on Oversampling Delta–Sigma ADCs." *Journal of the Audio Engineering Society*, Vol. 38, No. 718, pp. 537—542, July 1990.

Beex, A. A. and Monique P. Fargues. "Analysis of Clock Jitter in Switched-Capacitor Systems." IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications, Vol. 39, No. 7, July 1992.

Cherry, James A. and W. Martin Snelgrove. *Continuous-Time Delta-Sigma Modulators for High-Speed A/D Conversion: Theory, Practice and Fundamental Performance Limits.* Kluwer Academic Publisher, 2002.

Kester, Walt. "Converting Oscillator Phase Noise to Time Jitter." Analog Devices, Inc., 2009.

Markell, Richard. *Application Note: Take the Mystery Out of the Switched-Capacitor Filter: The System Designer's Filter Compendium.* Analog Devices, Inc., March 1990.

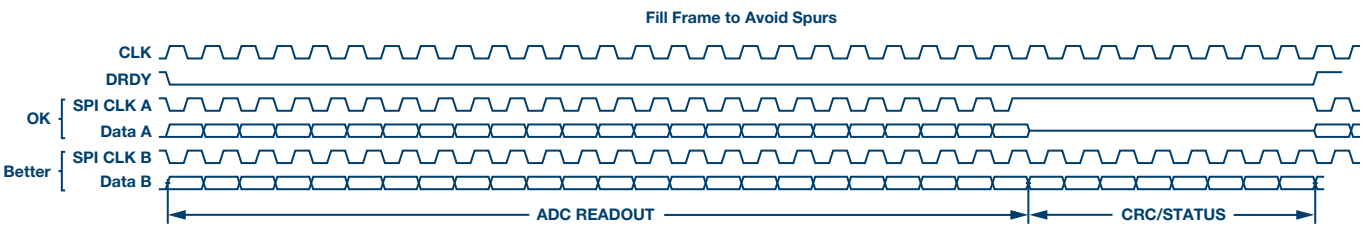

Figure 18. One can improve the frame with dummy CRC or status to remove spur.

Pawel Czapor [pawel.czapor@analog.com] joined Analog Devices in 2007 as a precision ADC test engineer. Currently, he works in the High Precision Converters Design Group in Limerick, Ireland. He graduated with an M.Eng. degree in electronic engineering from Wroclaw University of Science and Technology in Poland.

Pawel Czapor

Silent Switcher Devices Are Quiet and Simple

By **Tony Armstrong**

Introduction

It goes without saying that PC board layout determines the success or failure of every power supply design. It sets functional, electromagnetic interference (EMI), and thermal behavior. While switching power supply layout is not black magic, it is often overlooked until it is too late in the design process. Therefore, having a proven way to mitigate the potential threats from EMI generation from the onset can ensure a quiet and stable power supply. While many switch mode power supply designers are familiar with the design complexities and nuances of switch-mode supplies, there are simply not enough designers in many companies to get all the designs done for all of their project needs. They are retiring and leaving the industry! So, how does this problem get solved?

Well, for one thing, more and more digital designers are being asked to take on switch-mode power supply designs, if for no other reason than there are simply not enough analog power supply designers to get the job done! While it is safe to say that most digital designers know how to design with a simple linear regulator, not all of their power requirements are step-down (buck mode). In fact, many are step-up mode (boost) or even a buck-boost topology (buck and boost modes combined).

Clearly, an obvious question facing many electronic systems manufacturers is this: How will all the switch-mode power supply circuits needed in my systems get finished?

Solving a Design Resource Shortfall

During the course of this discourse, I will go over some of the basics of buck regulator operation including how high di/dt and parasitic inductance in the switcher hot loop cause electromagnetic noise and switch ringing. And then we will see what can be done to reduce the high frequency noise. I will also introduce ADI's Power by Linear™ Silent Switcher® technology, including how it is constructed, and demonstrate how it helps solve EMI problems without any compromises. This will also include how Silent Switcher devices work.

I will also give an overview of Silent Switcher packaging and layout and discuss how these can enhance the overall performance of the step-down converters. Moreover, I will show how to achieve a higher level of integration of our Silent Switcher device by demonstrating how this technique can be incorporated into our µModule® regulators. These provide simple and easy to use solutions for those users who are not sophisticated in switch-mode power supply design techniques.

Basic Buck Regulator Circuit

One of the most basic power supply topologies is the buck regulator, as shown in Figure 1. EMI starts off from the high di/dt loops. The supply wire, as well as the load wire should not have high ac current content. Accordingly, the input capacitor, C2 should source all the relevant ac to the output capacitor, C1, where any ac ends.

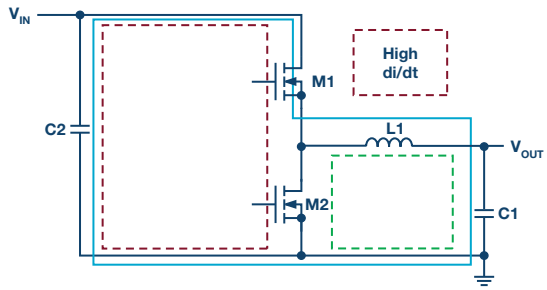

Figure 1. A synchronous buck regulator schematic.

Still referring to Figure 1, during the on cycle with M1 closed and M2 open, the ac follows in the solid blue loop. During the off cycle, with M1 open and M2 closed, the ac follows the green dotted loop. Most people have difficulty grasping that the loop producing the highest EMI is not the solid blue nor the dotted green. Only in the dotted red loop flows a fully switched ac, switched from the zero to I peak and back to zero. The dotted red loop is commonly referred to as a hot loop since it has the highest ac and EMI energy.

It is 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 dotted red loop as much as possible. If we could reduce the PC board area of the dotted red loop to zero and buy an ideal capacitor with zero impedance, the problem would be solved. However, in the real world, it is the design engineer who must find an optimal compromise!

So where does all this high frequency noise come from anyway? Well, in electronic circuits, the switching transitions coupled though parasitic resistors, inductors, and capacitors create high frequency harmonics. So, knowing where the noise is generated, what can be done to reduce the high frequency switching noise? The traditional way to reduce noise is to slow down the MOSFET switching edges. This can be accomplished by slowing the internal switch driver or by adding snubbers externally.

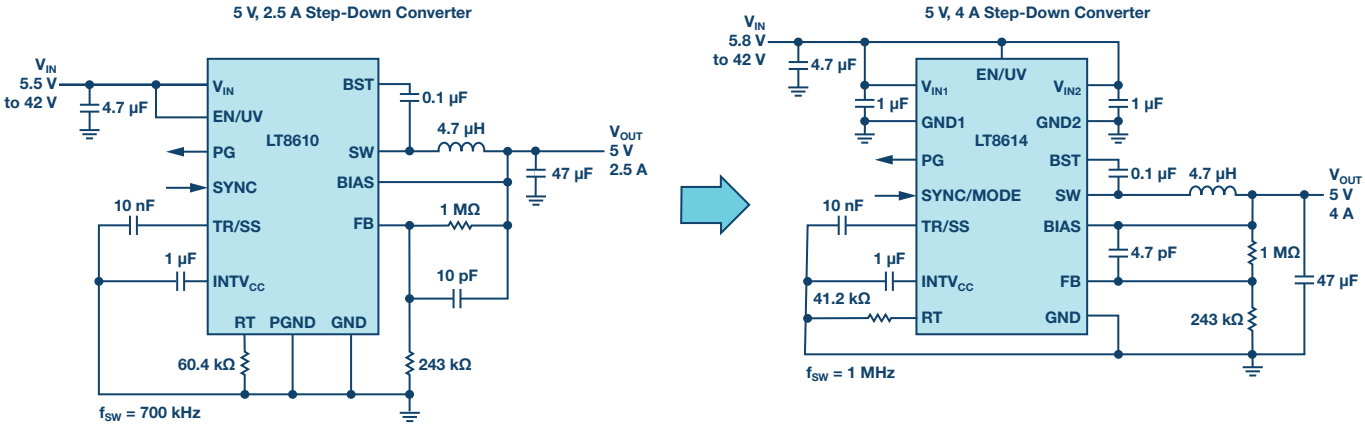

Figure 2. How to make an LT8610 into a Silent Switcher device—the LT8614.

However, this will reduce the efficiency of the converter due to increased switching loss—especially if the switcher is running at a high switching frequency of, say, 2 MHz. Speaking of which, why would we want to run at 2 MHz? Well, for several reasons actually:

- It enables the use of smaller (size) external components such as capacitors and inductors. For example, every doubling of switching frequency is a halving of inductance value and output capacitance value.
- In automotive applications, switching at 2 MHz keeps noise out of the AM radio band.

Filters and shielding can also be employed, but this costs more in terms of external components and circuit board area. Spread spectrum frequency modulation (SSFM) could also be implemented—this technique dithers the system clock within a known range. SSFM helps to pass the EMI standards. The EMI energy is distributed over the frequency domain. Although the switching frequency is most often chosen to be outside the AM band (530 kHz to 1.8 MHz), unmitigated switching harmonics can still violate stringent automotive EMI requirements within the AM band. Adding SSFM significantly reduces EMI within the AM band as well as in other regions.

Or, one could simply use ADI's Silent Switcher technology instead since it delivers on all the points outlined with no trade-offs:

- High efficiency
- High switching frequency
- Low EMI

Silent Switcher Technology

A Silent Switcher device breaks the trade-off between EMI and efficiency by not needing to slow down the switch edge rates. But how can this be accomplished? Consider the LT8610, as shown on the left side of Figure 2. It is a 42 V input capable, monolithic (FETs inside) synchronous buck converter that can deliver up to 2.5 A of output current. Notice that it has a single input pin (V_{IN}) at its top left corner.

However, when contrasting the LT8610 to the LT8614 (a 42 V input capable, monolithic synchronous buck converter that can deliver up to 4 A of output current), one can see that the LT8614 has two V_{IN} pins and two ground pins on the opposite side of the package. This is significant, since it is part of what makes it *silent* switching!

How to Make a Switcher Silent

So how can we do what we do? Placing two input capacitors on opposite sides of the chip between the V_{IN} and ground pins will cancel the magnetic fields. This is highlighted in the slide with the red arrows pointing to the capacitor placement, both on the schematic and the demo board, as shown in Figure 3.

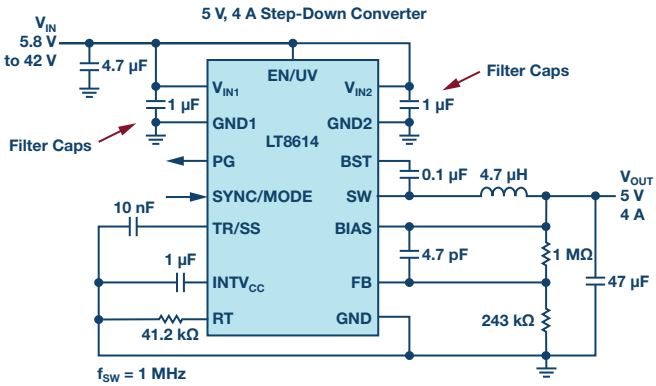

Figure 3. Diagram of the LT8614 showing the filter caps between V_{IN} and ground pins on opposing sides of the IC.

The LT8614 in More Detail

The LT8614 incorporates Silent Switcher capability. With it, we were able to reduce the parasitic inductance by using copper pillar flip-chip packaging. Furthermore, there are opposing V_{IN}, ground, and input caps to enable magnetic field cancellation (right-hand rule applies) to lower EMI emissions.

Reducing the package parasitic inductance is achieved by eliminating the long bond wires of a wire-bonded assembly technique, which induces parasitic resistance and inductance. The opposing magnetic fields from the hot loops cancel each other out and the electric loop sees no net magnetic field.

We compared the LT8614 Silent Switcher regulator against a current state-of-the-art switching regulator, the LT8610. Testing was performed in a GTEM cell using the same load, the same input voltage, and the same inductor on the standard demo boards for both parts. We found that a

20 dB improvement is made when using the LT8614 compared to the already very good EMI performance of the LT8610, especially in the more difficult to manage higher frequency areas. This enables simpler and more compact designs where the LT8614 switching power supply needs less filtering and distance compared to other sensitive systems in the overall design. Furthermore, in the time domain, the LT8614 exhibits very benign behavior on the switch node edges.

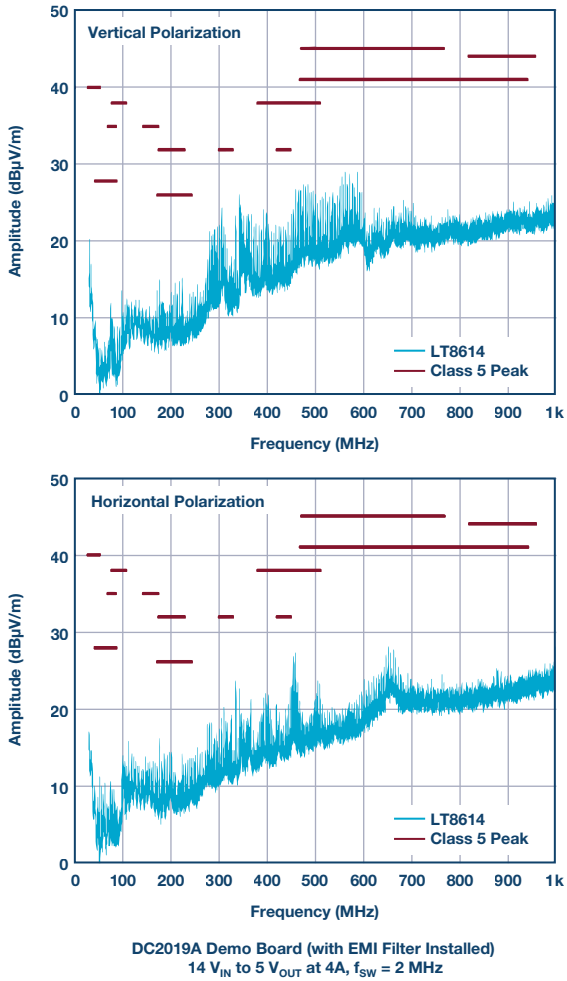

Figure 4. The LT8614 radiated EMI performance passes the most stringent CISPR 25 Class 5 limits.

Further Enhancements to Silent Switcher Devices

Although the LT8614 has impressive performance, we did not stop trying to improve upon its performance. Accordingly, the LT8640 step-down regulator features Silent Switcher architecture designed to minimize EMI/EMC emissions while delivering high efficiency at frequencies up to 3 MHz. Assembled in a 3 mm × 4 mm QFN, the monolithic construction with integrated power switches and inclusion of all necessary circuitry yields a solution with a minimal PCB footprint. Transient response remains excellent and output voltage ripple is below 10 mV p-p at any load, from zero to full current. The LT8640 allows high V_{IN} to low V_{OUT} conversion at high frequency with a fast minimum top switch on-time of 30 ns.

To improve EMI/EMC, the LT8640 can operate in spread spectrum mode. This feature varies the clock with a triangular frequency modulation of 20%. When the LT8640 is in spread spectrum frequency modulation mode, a triangular frequency modulation is used to vary the switching frequency between the value programmed by RT to approximately 20% higher than that value. The modulation frequency is approximately 3 kHz. For example, when the LT8640 is programmed to 2 MHz, the frequency will vary from 2 MHz to 2.4 MHz at a 3 kHz rate. When spread spectrum operation is selected, Burst Mode® operation is disabled, and the part will run in either pulse-skipping mode or forced continuous mode.

Nevertheless, despite all we stated in our Silent Switcher data sheets, such as showing the schematic and layout recommendations, as well as the input capacitors being placed as close as possible to the IC on both side—some of our customers still make mistakes. Furthermore, our in-house engineers were spending too much time fixing our customer's PCB layouts. So, our designers came up with a solution to this problem—the Silent Switcher 2 architecture.

Silent Switcher 2

With Silent Switcher 2 technology, we simply integrated the capacitors inside a new LQFN package: V_{IN} caps, IntV_{CC}, and boost caps—allowing for placement as close as possible to the pins. The benefits included all the hot loops and ground planes inside, with all of this resulting in lower EMI. Fewer external components meant a smaller solution footprint. Furthermore, we also eliminated PCB layout sensitivity.

As shown in Figure 5, you can see how the schematics for the LT8640 and LT8640S differ. And, the marketing breakthrough was to give the new, higher integrated version with internal caps an "S" suffix. Because it is more "silent" than the first generation!

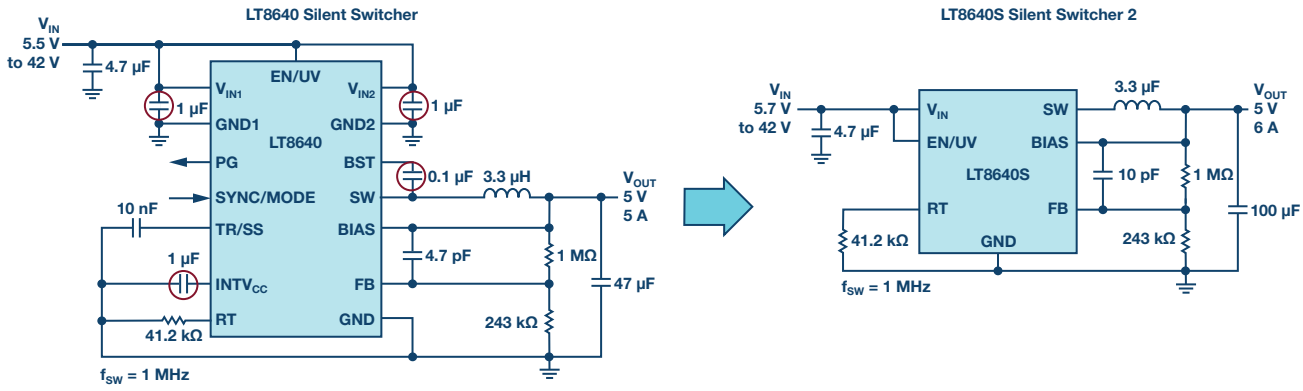

Figure 5. The LT8640S is a Silent Switcher 2 device with a higher level of integration of capacitors.

Silent Switcher 2 technology enables better thermal performance. The large multiple ground exposed pads on the LQFN flip-chip package facilitate the pulling of heat out of the package and into the PCB. We also get higher conversion efficiency since we have eliminated the high resistance bond wires. The EMI performance of the LT8640S easily passes the radiated EMI performance CISPR 25 Class 5 peak limits with a wide margin.

The Next Step: Everything Integrated with Silent Switcher 2 µModule Regulators

Silent Switcher technology is so compelling that we elected to incorporate it into our µModule regulator product line. In this form factor, everything is integrated inside a single package and provides the user with simplicity, reliability, performance, and high power density. The LTM8053 and LTM8073 are micromodule regulators where everything is virtually integrated with just a few capacitors and resistors external to them.

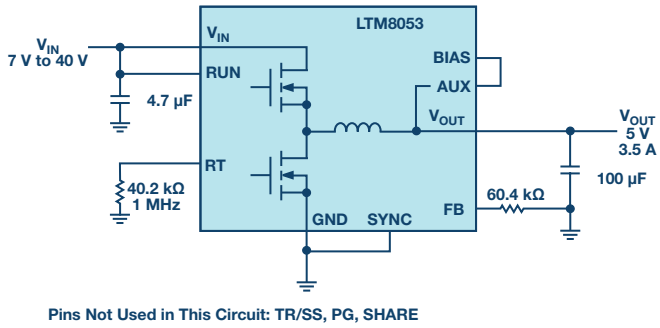

Figure 6. The LTM8053 Silent Switcher 2 µModule regulator.

Summary

In conclusion, Silent Switcher capabilities and benefits will make it easier for your switch mode power supply designs to pass the various noise immunity standards such as CISPR 32 and CISPR 25. They can do this easily and effectively due to their following attributes:

- High efficiency conversion at greater than 2 MHz switching frequency with minimum impact on conversion efficiency.
- Internal bypass capacitors reduce EMI radiation and provide for a more compact solution footprint.
- PCB layout sensitivity is essentially eliminated with Silent Switcher 2 technology.
- Optional spread spectrum modulation helps mitigate noise sensitivity.
- Using Silent Switcher devices saves on PCB area and can also reduce the number of layers needed.

Tony Armstrong [tony.armstrong@analog.com] is currently the product marketing director for Analog Devices' Power by Linear product group. He is responsible for all aspects of power conversion and management products from their introduction through obsolescence. Prior to joining ADI, Tony held various positions in marketing, sales, and operations at Linear Technology, Siliconix Inc., Semtech Corp., Fairchild Semiconductors, and Intel. He attained his B.S. (Honors) in applied mathematics from the University of Manchester, England.

Tony Armstrong

Rarely Asked Questions—Issue 164 Guidelines for Placing the Inductor on a Switch-Mode Power Supply Printed Circuit Board

By **Frederik Dostal**

Question:

Where should the coil go?

Answer:

Switching regulators for voltage conversion use inductors for temporarily storing energy. These inductors are often very large components and must be positioned in the printed circuit board (PCB) layout of the switching regulator. This task is not that hard because the current through an inductor can change, but not instantaneously. There can only be continuous, usually relatively slow changes.

Switching regulators switch the current flow back and forth between two different paths. This switching occurs very rapidly and the speed depends on the switching edge durations. The resulting traces, which conduct current in one switching state and do not conduct current in the other switching state, are called hot loops or ac current paths. They should be kept especially small and short in the PCB layout so that parasitic inductance is minimized in these traces. Parasitic trace inductances generate an undesired voltage offset and result in electromagnetic interference (EMI).

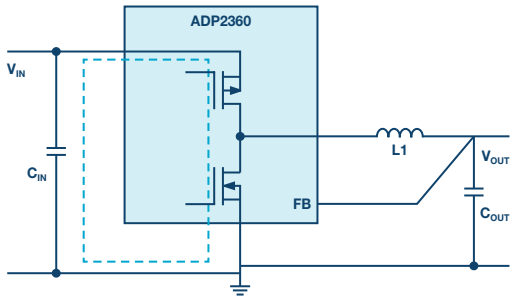

Figure 1. A switching regulator for step-down conversion with a critical hot loop shown as the dotted line.

Figure 1 shows a step-down regulator in which the critical hot loop is shown as a dotted line. It can be seen that coil L1 is not part of the hot loop. Thus, it can be assumed that placement of this inductor is not critical. It is correct to have the inductor lie outside the hot loop—therefore placement is, in the first instance, secondary. Still, a few rules should be followed.

No sensitive control traces should be routed under an inductor, neither right on the PCB surface nor below, in internal layers, or on the backside of the PCB. Due to the flow of current, the coil generates a magnetic field, which can affect weak signals in a signal path. In a switching regulator, one of the critical signal paths is the feedback trace, which connects the output voltage to the switching regulator IC or to a voltage divider.

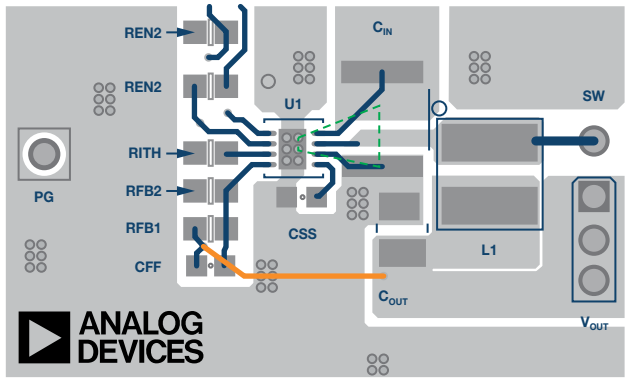

Figure 2. Example circuit with an ADP2360 step-down (buck) converter with coil placement.

It should also be noted that a real coil has a capacitive effect as well as an inductive one. The first coil windings are directly connected to the switch node of a step-down switching regulator, as shown in Figure 1. As a result, the voltage changes just as strongly and rapidly as the voltage at the switch node does. With very short switching times and high input voltages in the circuit, a considerable coupling effect is yielded on other paths on the PCB. Thus, for this reason as well, sensitive traces should be kept away from the location of the coil.

Figure 2 shows a sample layout with the ADP2360. Here, the important hot loop from Figure 1 is marked in green. The yellow feedback path can be seen at a distance to coil L1. It is on an internal layer of the PCB.

Some circuit designers even go so far as to not want any copper layers in the PCB under the coil. They would, for example, provide for a cutout under the inductor, even in a ground plane layer. The goal is to prevent eddy currents in the ground plane under the coil resulting from the magnetic field of the coil. This approach is not wrong, but there are arguments for a solid ground plane with no interruptions:

- A ground plane for shielding works best when it is not interrupted.
- The more copper a PCB has, the better the heat dissipation.
- Even if eddy currents are generated, these currents flow locally, only cause small losses, and hardly affect the function of a ground plane.

Thus, I am in favor of a solid ground plane layer, even under the coil.

In summary, we can conclude that the coil of a switching regulator is not part of the critical hot loop, but that it makes sense to not route control traces under or very close to the coil. Various planes on the PCB—for example, for ground, or also for V_{DD} (supply voltage)—may be created continuously, without cutouts.

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

Frederik Dostal

Also by this Author:

When Grounds are Separated

Volume 52, Number 4

Analysis of Input Current Noise with Even Harmonics Folding Effect in a Chopper Op Amp

By **Yoshinori Kusuda**

Abstract

This article presents a theoretical analysis and measurements of the input current noise of a chopper operational amplifier that has a 10 pF input capacitance, a 5.6 nV/√Hz voltage noise power spectral density (PSD), and a 4 MHz unity-gain bandwidth. With a higher closed-loop gain configuration, the input current noise is dominated by the thermal noise of the dynamic conductance that occurs at the input chopper. Additionally, this theoretical analysis identifies another source of the input current noise that is caused by the amplifier's voltage noise as sampled by the dynamic conductance at the input chopper. Moreover, upon the sampling, the broadband voltage noise spectral densities are folded back to the low frequency so that the resulting current noise spectral density actually increases with wider closed-loop bandwidth, hence with smaller closed-loop gain configuration. The measured current noise is 0.28 pA/√Hz with a closed-loop gain of 10, but it increases up to 0.77 pA/√Hz with the unity-gain configuration.

I. Introduction

The chopping technique periodically corrects an amplifier's offset voltage, and therefore can achieve microvolt-level offset voltage and very little 1/f noise with its corner frequency below sub-hertz.[1,2] Hence, many chopper operational amplifiers (op amps) and instrumentation amplifiers (in-amps) have mainly been intended for sensing small input voltage that has relatively low source impedance and low signal frequency. One of its important applications is to amplify millivolt-level sensor signals representing light, temperature, magnetic field, and force, whose signal frequencies are mostly lower than kilohertz.[2] However, the switching by the input chopper introduces input bias current and input current noise that are substantially higher than those of a conventional CMOS amplifier with no chopping.[3,4] When the amplifier's input is driven by a high source impedance, this input current noise will be converted to voltage noise, which may dominate the overall amplifier's noise.[3,4]

In the article "Measurement and Analysis of Input Current Noise in Chopper Amplifiers,"[4] various possible sources of input current noise were explained, and the shot noise associated with the charge injection of the input MOS switches was identified as the dominant noise source. However, in the article "Excess Current Noise in Amplifiers with Switched Input,"[5] the thermal noise of the dynamic conductance that occurred at the input chopper was identified as the dominant noise source. In all these previous measurements, the amplifier's output voltage noise was isolated from the input chopper by feedback attenuation from the amplifier's output to input.

Although chopper op amps have been traditionally used in high closed-loop gain configurations, their low offset voltage and low 1/f noise are also demanded in low closed-loop gain and/or high source impedance configurations.[2] Therefore, it is important to understand their current noise behavior in such configurations. This brief presents input current noise analysis and measurements of a chopper op amp with both high and low closed-loop gain configurations as presented in the article "A 5.6 nV/√Hz Chopper Operational Amplifier Achieving a 0.5 µV Maximum Offset Over Rail-to-Rail Input Range with Adaptive Clock Boosting Technique."[6] It identifies another source of the input current noise that is caused by the op amp's broadband voltage noise sampled by the dynamic conductance at the input chopper. Moreover, upon the sampling, the voltage noise PSDs from even harmonic frequencies of the chopping are folded back to the low frequency, which can increase the resulting current noise PSD. Therefore, this noise source can dominate the overall input current noise when the closed-loop gain is lower so that the output voltage noise of the op amp reaches the input chopper with less attenuation.

Section II reviews previously reported input current noise sources, and then Section III explains the mechanism of the input current noise source caused by the sampled broadband voltage noise and the associated noise spectrum folding effect. Section IV conducts some numerical calculations of various current noise sources of the op amp.[6] Section V then compares the calculated current noise with simulations and measurements to validate the analysis. Section VI provides some recommendations to reduce the input current noise, and the article ends with some conclusions in Section VII.

II. Previously Reported Input Current Noise Sources

The following three current noise sources were explained in the article "Measurement and Analysis of Input Current Noise in Chopper Amplifiers." First, channel charge injections of the input switches can be approximated as an average current I_{q_ave}, leading to shot noise:

$$i_{n\_SHOT} = \sqrt{2qI_{q\_ave}} = \sqrt{4qf_{CHOP} \times (WLC_{ox})_{SW} \times (V_{GS} - V_{TH})_{SW}} \quad (1)$$

where f_{CHOP} is the chopping frequency, while $(WLC_{ox})_{SW}$ and $(V_{GS} - V_{TH})_{SW}$ are the gate oxide capacitance and the overdrive voltage of the switches, respectively.

Second, the clock drivers produce kTC noise charges sampled onto the gate oxide capacitances of the switches, and then the noise charges flow into the amplifier's inputs at every chopping:

$$i_{n\_kTC} = \sqrt{8kT\,(WLC_{ox})_{SW} \times f_{CHOP}} \quad (2)$$

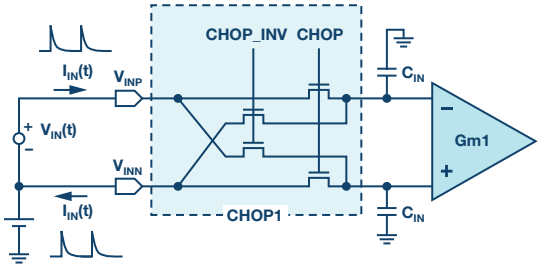

Figure 1. Dynamic input current due to chopping and input capacitances.

Third, as shown in Figure 1, a dynamic input current $I_{IN}(t)$ flows into the amplifier's input capacitors C_{IN} every time the input chopper, CHOP1, switches. When a dc voltage source $V_{IN}(t) = V_{IN\_DC}$ is applied, the averaged input currents over time I_{IN_ave} is given by:

$$I_{IN\_ave} = 2C_{IN}f_{CHOP} \times V_{IN\_DC} \quad (3)$$

The associated dynamic input conductance G_{IN_ave} and thermal noise i_{n_GIN} are then given by:

$$G_{IN\_ave} = \frac{I_{IN\_ave}}{V_{IN\_DC}} = 2C_{IN}f_{CHOP} \quad (4)$$

$$i_{n\_GIN} = \sqrt{4kTG_{IN\_ave}} = \sqrt{8kTC_{IN}f_{CHOP}} \quad (5)$$

Note that any one of the three noise equations of Equation 1, Equation 2, and Equation 5 consists of a unique set of the circuit and switch parameters, and therefore can dominate the overall noise depending on the values of the parameters. The shot noise shown in Equation 1 dominates the overall current noise in all three measured amplifiers[4]: an open-loop chopper in-amp and two chopper op amps with closed-loop gains of 100. This open-loop in-amp only had a 125 fF input capacitor, and thus the thermal noise of the dynamic conductance shown in Equation 5 was insignificant.

In the article "Excess Current Noise in Amplifiers with Switched Input," a chopper made of discrete FETs was measured, and the thermal noise shown in Equation 5 dominated the overall current noise when discrete capacitors ranging from 10 pF to 100 pF were added. Note that the current noise increased with the capacitor value.

III. Current Noise Caused by Sampled Voltage Noise and Noise Spectrum Folding Effect

The dynamic conductance itself generates the thermal current noise as suggested by Equation 5, but its sampling action also converts the voltage noise across the input chopper to current noise.

Dynamic Input Current Caused by Sampled AC Input Voltage

The dynamic input current with a dc input voltage is given by Equation 3. Let us now consider a case with an ac sinusoidal differential input voltage $V_{IN}(t)$) at the frequency of $2 \times f_{CHOP}$, as shown in Figure 2. It can be seen that $V_{IN}(t)$ reaches its peak value V_{IN_AC} when the chopping clocks CHOP and CHOP_INV switch. Consequently, this ac differential input voltage results in a dynamic input current $I_{IN}(t)$ in the same manner as a dc differential input voltage does, so that its time-averaged current I_{IN_ave} is given by:

$$I_{IN\_ave} = 2C_{IN}f_{CHOP} \times V_{IN\_AC} \quad (6)$$

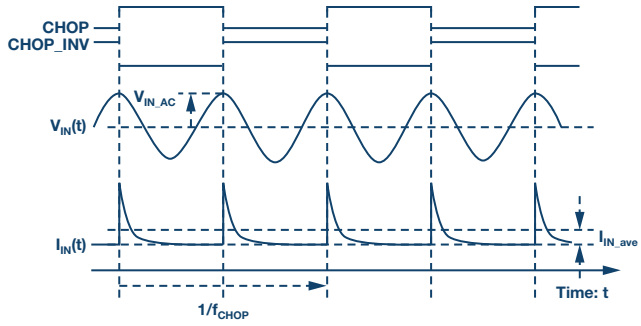

Figure 2. Dynamic input current waveform with ac differential input voltage.

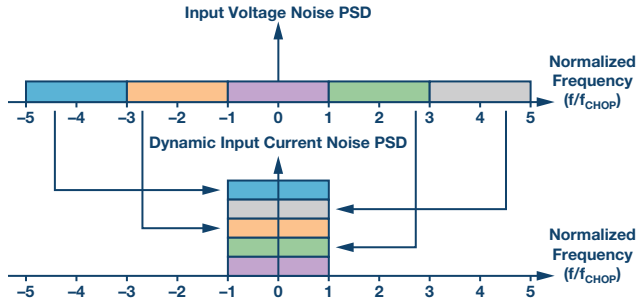

Figure 3. Noise spectrum folding effect while voltage noise PSD is sampled and converted to current noise PSD.

When the phase difference between the input voltage and the chopping clocks is random, the equation can be rewritten using the rms values of the input voltage V_{IN_RMS} and the resulting input current $I_{IN\_ave\_RMS}$:

$$I_{IN\_ave\_RMS} = 2C_{IN}f_{CHOP} \times V_{IN\_RMS} \quad (7)$$

The input current will also occur in the same manner, when an ac input differential voltage at a higher even harmonic frequency of the chopping (for example, $4 \times f_{CHOP}$ or $6 \times f_{CHOP}$) is applied.

Input Current Noise PSD Caused by Sampled Voltage Noise PSD and Noise Spectrum Folding Effect

When the input voltage has frequency spectrum including multiple even harmonic frequencies of the chopping, they are all folded back to low frequency, which is known as the noise spectrum folding effect.[1] The chopping is considered a modulation technique rather than a sampling technique. However, this dynamic input current occurs based on the sampled input voltages, rather than the continuous input voltage, so that noise spectrum folding occurs. In other words, the amount of the averaged dynamic current is only determined by the differential input voltages at the instance of the chopping, rather than the differential input voltage at any other time.

Figure 3 illustrates the noise spectrum folding effect with the consideration that an input voltage noise PSD is equal to e_n from dc to $5 \times f_{CHOP}$ but is zero above $5 \times f_{CHOP}$. This results in an input current noise PSD from dc to $\pm f_{CHOP}$, the Nyquist frequency. The input voltage noise PSD $e_n(f_{en})$ between $\pm f_{CHOP}$ will contribute to the input current noise PSD $i_{n\_en\_GIN\_0}(f)$ with no frequency shift:

$$i_{n\_en\_GIN\_0}(f_{in}) = 2C_{IN}f_{CHOP} \times e_n(f_{en}) \quad (8)$$

where f_{en} and f_{in} are the frequencies of the input voltage noise PSD and the resulting input current noise PSD, respectively. The input voltage noise PSD above f_{CHOP} and below $3 \times f_{CHOP}$ will contribute to the input current noise PSD with a frequency shift of $-2 \times f_{CHOP}$:

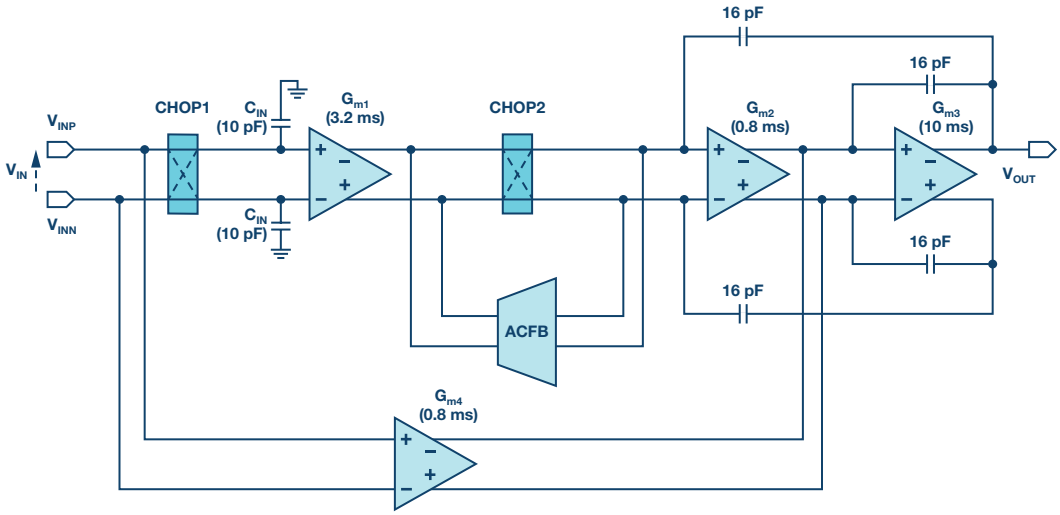

Figure 4. Chopper op amp diagram.

$$i_{n\_en\_GIN\_2fCHOP}(f_{in}) = 2C_{IN}f_{CHOP} \times e_n(f_{en} - 2f_{CHOP}) \quad (9)$$

The total input current noise PSD $i_{n\_en\_GIN\_RSS}(f)$ is obtained by summing the PSDs folded from all frequencies within the op amp's closed-loop bandwidth including those in Equation 8 and Equation 9, in the root sum square (RSS) manner:

$$i_{n\_en\_GIN\_RSS}(f_{in}) = 2C_{IN}f_{CHOP}\sqrt{\sum_{n=-\infty}^{\infty} e_n^2(f_{en} - 2nf_{CHOP})} \quad (10)$$

When the voltage noise PSD is flat at e_n and is band limited at a frequency of f_{en_BW}, the resulting low frequency current noise PSD is given by:

$$i_{n\_en\_GIN\_RSS} = 2C_{IN}f_{CHOP} \times e_n \times \sqrt{1 + \frac{f_{en\_BW}}{f_{CHOP}}} \quad (11)$$

When $f_{en\_BW}/f_{CHOP} >> 1$, the equation can be approximated to:

$$i_{n\_en\_GIN\_RSS} \cong 2C_{IN}\sqrt{f_{CHOP} \times f_{e\_BW}} \times e_n = 2C_{IN}\sqrt{f_{CHOP}} \times e_{n\_RMSINT} \quad (12)$$

where $e_n \times \sqrt{f_{en\_BW}}$ is replaced by the integrated rms voltage noise e_{n_RMSINT}. This input current noise source is approximately proportional to the rms voltage noise at the differential inputs, the input capacitor size, and the square root of the chopping frequency.

IV. Input Current Noise Estimation in a Chopper Op Amp

Chopper Op Amp Block Diagram

The chopper op amp presented in "A 5.6 nV/√Hz Chopper Operational Amplifier Achieving a 0.5 µV Maximum Offset Over Rail-to-Rail Input Range with Adaptive Clock Boosting Technique" is analyzed, simulated, and measured in this and later sections. This op amp is realized in a 0.35 µm CMOS process augmented by 5 V transistors, and it achieves a voltage noise PSD of 5.6 nV/√Hz and a unity-gain bandwidth of 4 MHz. Its block diagram is shown in Figure 4, and Table 1 summarizes the parameters of the input chopper (CHOP1). To realize rail-to-rail input common-mode range, the input transconductance amplifier stage G_{m1} consists of n-channel and p-channel differential pairs, both of which contribute to the input capacitances C_{IN}. Moreover, the larger size of the input MOS devices is needed to increase the transconductance of G_{m1} in a power efficient manner. Each of the four switches in the input chopper CHOP1 is realized by an NMOS, and its gate voltage is adaptively biased based on the input voltage, so that its overdrive voltage is constant at 0.5 V with the changes of the input voltage.

Table 1. Parameters of the Input Chopper (CHOP1)

Parameter	Explanation	Value	Unit
f_{CHOP}	Chopping frequency	200	kHz
C_{IN}	Input capacitance of G_{m1}	10	pF
R_{FB}	Gate oxide capacitance of a switch in CHOP1	30	fF
$(V_{GS}–V_{TH})_{SW}$	Gate overdrive voltage of a switch in CHOP1	0.5	V
k	Boltzmann constant	1.38×10^{-23}	J/K
T	Absolute temperature	300	K
q	Unit electron charge	1.60×10^{-19}	C

Voltage Noise Across Differential Input Terminals

To calculate the current noise PSD shown in Equation 12, the integrated rms voltage noise v_{in_RMSINT} needs to be known. The chopper op amp is simulated with closed-loop gains = 1, 2, 5, and 10. Figure 5 (a) and (b) show the voltage noise PSDs and their integrated rms noise, respectively, across the differential inputs of the op amp. All the simulations in this article are conducted by the SpectreRF periodic noise simulation (P_{NOISE}) to consider switching effects of the chopping.[7] The noise PSDs are flat below 100 kHz thanks to the chopping, but peak at the chopping frequency of 200 kHz.[6] Note that the figures present the noise at the op amp's differential inputs rather than its output, so that the noise PSDs below 100 kHz are constant with different closed-loop gains. The noise PSDs also increase above 1 MHz and are dominated by the thermal noise of G_{m2}, G_{m3}, and G_{m4} due to the gain drop of G_{m1}. Therefore, their integrated rms noise also increases above 1 MHz, especially with lower closed-loop gain, mainly due to the higher closed-loop bandwidth. The integrated rms voltage noise across the differential inputs is 11 µV rms with gain = 10, but is 68 µV rms with gain = 1.

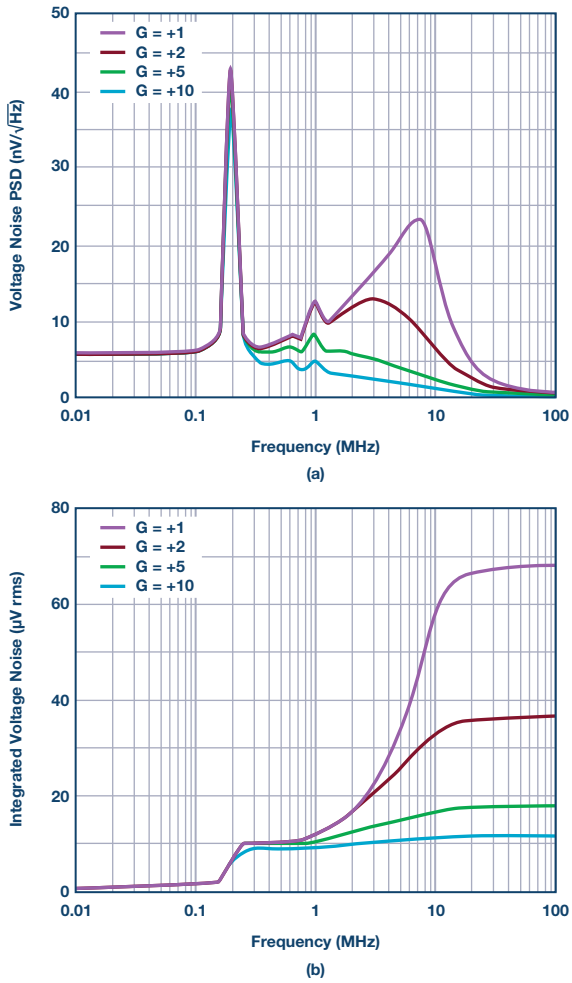

Figure 5. Simulated differential input voltage noise of the chopper op amp.

Estimation of Each Input Current Noise Source

This simulated integrated rms voltage noise is then applied to Equation 12 to calculate the current noise PSDs. Additionally, the current noise PSDs caused by the other noise sources[4] are calculated by applying the parameters in Table 1 to Equation 1, Equation 2, and Equation 5. Figure 6 shows the calculated current noise PSDs of the four noise sources with closed-loop gains from 1 to 10. The current noise PSD caused by the sampled broadband voltage noise PSD (Equation 12) dominates the total current noise PSD at closed-loop gains of 1 and 2. It decreases with higher closed-loop gains and only contributes to the total input current noise PSD by 7% at the closed-loop gain of 10. Instead, the total current noise PSD is dominated by the thermal noise of the dynamic conductance itself (Equation 5), and thus becomes nearly constant with the closed-loop gain above 5. Therefore, it is sufficient to evaluate the current noise with the closed-loop gain up to 10 for this op amp.[6]

V. Simulation and Measurement Results

To validate the analysis, the calculated total current noise PSDs shown in Figure 6 are compared with the simulation and measurement results. Both P_{NOISE} simulation and measurement are performed using a circuit setup, as shown in Figure 7. The voltage noise PSD e_{n_OUT} is measured by shorting R_S, and then the overall noise PSD $e_{n\_OUT\_RS}$ is measured with R_S = 100 kΩ. The current noise PSD i_{n_IN} is then given by:

$$i_{n\_IN} = \frac{\sqrt{\frac{(e_{n\_OUT\_RS}^2 - e_{n\_OUT}^2)}{G_{TOT}^2} - 4kTR_S}}{R_S} \quad (13)$$

$$G_{TOT} = \left(1 + \frac{R_F}{R_G}\right) \times G_{POST} \quad (14)$$

where $(1 + R_F/R_G)$ is the closed-loop gain around the op amp and G_{POST} = 100 is a post gain to ease the measurement by the dynamic signal analyzer HP 35670A. Note that in Equation 13 $e_{n\_OUT\_RS}$ and e_{n_OUT} are subtracted in RSS manner, because the current noise PSD is mostly caused by the folded noise from the higher frequencies and is thus uncorrelated with the voltage noise PSD.

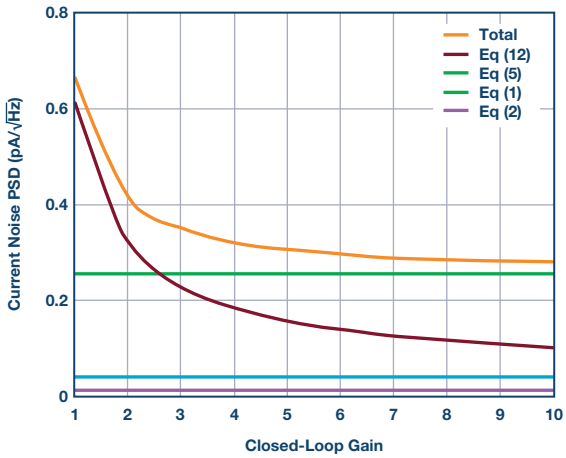

Figure 6. Calculated input current noise contribution from the different sources.

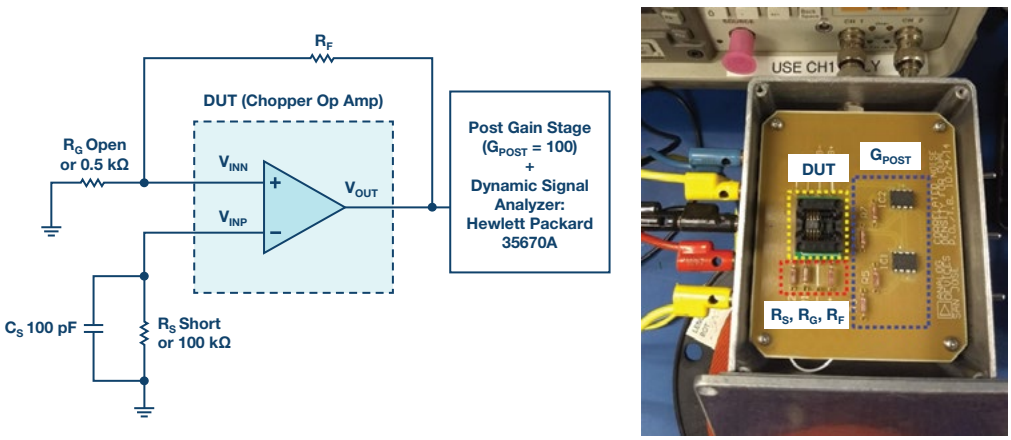

Figure 7. Circuit setup for input current noise simulations and measurements.

An external capacitor C_S = 100 pF limits the noise bandwidth of R_S at the cutoff frequency of 16 kHz. In this case, the thermal noise of R_S is sufficiently attenuated at the first even harmonic frequency of the chopping (400 kHz), and thus does not contribute to the current noise through the noise spectrum folding effect. On the other hand, the op amp's broadband output voltage noise reaches the negative input V_{INN}, sampled by the dynamic conductance at the input chopper, and can significantly contribute to the current noise. This resulting current noise PSD in the low frequency is then converted to the voltage noise again by R_S, which can be measured at the output of the post gain stage.

Figure 8 shows the simulated and measured input current noise PSDs over the frequency with gain = 1 configuration (R_G is open and R_F is short in Figure 7). At 0.01 kHz, the simulated and measured noise PSDs are 0.69 pA/√Hz and 0.78 pA/√Hz, respectively. The noise PSDs then start dropping at the 16 kHz cutoff frequency made by R_S and C_S. Figure 9 shows the input current noise PSDs at 0.01 kHz with different closed-loop gains to compare the calculated values in Figure 6 with the simulation and measurement results. Both the simulated and measured current noise PSDs increase with lower closed-loop gains, and present good correlation with the calculation. The measured input current noise PSD is 0.28 pA/√Hz with gain = 10, but increases up to 0.77 pA/√Hz with gain = 1.

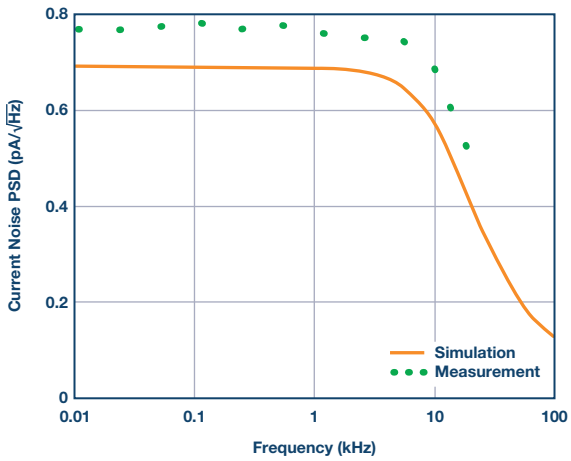

Figure 8. Input current noise PSD vs. frequency.

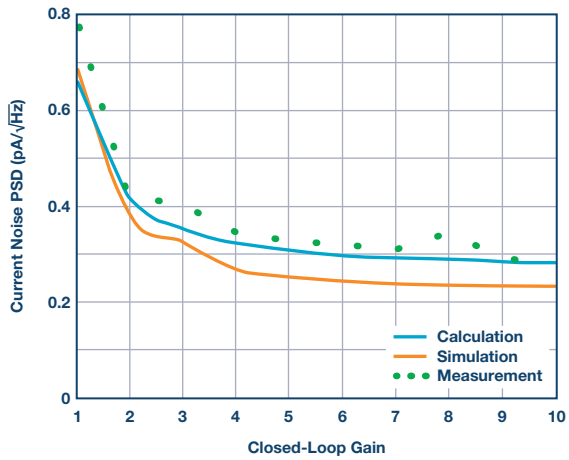

Figure 9. Input current noise PSD at 10 Hz vs. closed-loop gain.

VI. Recommendations to Reduce Input Current Noise

All the current noise sources given by Equation 1, Equation 2, Equation 5, and Equation 12 increase proportionally to the square root of the chopping frequency. Additionally, the current noise sources related to the dynamic conductance at the input chopper (Equation 5 and Equation 12) increase with an amplifier's input capacitance. This implies that chopper op amps designed for lower voltage noise PSD tend to have higher input current noise PSD, since the size of their input devices needs to be increased. This trade-off must be understood to achieve optimum voltage noise and current noise PSDs with a given source impedance. When possible, complementary input pairs or input transistors under a weak inversion region should be avoided to reduce the input capacitances.

Equation 12 identifies that the current noise PSD increases with the integrated rms voltage noise across the amplifier's differential inputs, and hence with noise bandwidth. Compared to open-loop chopper in-amps, chopper op amps are more vulnerable to this noise source, since their output noise can reach their input through the feedback network. When possible, a higher closed-loop gain can be used to decrease the noise bandwidth. Another way to decrease the noise bandwidth is to place capacitor(s) in parallel with R_G, R_S, and/or across the amplifier's differential inputs as shown in Figure 7.

VII. Conclusion

This article identified another input current noise source that is caused by the amplifier's broadband voltage noise sampled by the dynamic conductance at the input chopper. It also identified that, unlike the other noise sources previously reported, this current noise PSD increases with wider closed-loop bandwidth because of the noise spectrum folding effect associated with the input chopper. This analysis was confirmed by the measurements that showed 0.28 pA/√Hz current noise with gain = 10, and 0.77 pA/√Hz current noise with gain = 1 due to the increased closed-loop bandwidth. Some recommendations were provided for amplifier designers and users to reduce input current noise of chopper amplifiers. Table 2 compares the overall performance of the chopper op amp evaluated in this article[6] with other recent chopper op amps that have similar voltage noise PSD.[8, 9, 10]

Table 2. Specifications of the Chopper Op Amp

Parameter	This Work	LMP2021	MAX44250	OPA388
Supply Current (mA)	1.4	0.95	1.17	1.7
Chopping Frequency (kHz)	200	30	60	150
Gain Bandwidth Product (MHz)	4.0	5.0	10.0	10.0
Max Offset Voltage (µV)	0.5	5.0	8.5	5.0
Max Input Bias Current (pA)	400	100	1400	350
Voltage Noise PSD (nV/√Hz)	5.6	11.0	6.2	7.0
Current Noise PSD (pA/√Hz)	0.28	0.35	0.60	0.10

References

[1] Christian Enz and Gabor C. Temes. "Circuit Techniques for Reducing the Effect of Op Amp Imperfections: Auto-Zeroing, Correlated Double Sampling, and Chopper Stabilizationv." *Proceedings of the IEEE*, Vol. 84, No. 9, pp. 1320-1324, September 1996.

[2] Yoshinori Kusuda. "Reducing Switching Artifacts in Chopper Amplifiers," Ph.D. thesis. Delft University of Technology, the Netherlands, May 2018.

[3] Qinwen Fan, Johan Huijsing, and Kofi Makinwa. "Input Characteristics of a Chopped Multi-Path Current Feedback Instrumentation Amplifier." *2011 4th IEEE International Workshop on Advances in Sensors and Interfaces (IWASI)*, June 2011.

[4] Jiawei Xu, Qinwen Fan, Johan Huijsing, Chris Van Hoof, Refet Firat Yazicioglu, and Kofi Makinwa "Measurement and Analysis of Input Current Noise in Chopper Amplifiers." *IEEE Journal of Solid-Sate Circuits*, Vol. 48, No. 7, pp. 1575 to 1584, July 2013.

[5] Dietmar Drung and Christian Krause. "Excess Current Noise in Amplifiers with Switched Input." *IEEE Transactions on Instrumentation and Measurement*, Vol. 64, No. 6, pp. 1455 to 1459, June 2015.

[6] Yoshinori Kusuda. "A 5.6 nV/√Hz Chopper Operational Amplifier Achieving a 0.5 µV Maximum Offset Over Rail-to-Rail Input Range with Adaptive Clock Boosting Technique." *IEEE Journal of Solid-Sate Circuits*, Vol. 51, No. 9, pp. 2119 to 2128, September 2016.

[7] Ken Kundert. "Simulating Switched-Capacitor Filters with SpectreRF." Designer's Guide Consulting, Inc., July 2006.

[8] LMP2021 Data Sheet. Texas Instruments, September 2009.

[9] MAX44250 Data Sheet. Maxim Integrated, October 2011.

[10] OPA388 Data Sheet. Texas Instruments, December 2016.

Yoshinori Kusuda [yoshinori.kusuda@analog.com] joined Analog Devices Japan Design Center in 2004 after graduating from Tokyo Institute of Technology with an M.S. degree in physical electronics. He is currently based in San Jose, California, working for ADI's Linear Products and Solutions Group. He has been working with precision CMOS analog Designs, including standalone amplifiers and application specific mixed-signal products. This work has resulted in presentations and papers at IEEE conferences and journals, as well as 10 issued U.S. patents. From 2015 through 2018, he was registered as a guest in the Electronic Instrumentation Laboratory of TU Delft and obtained his Ph.D. on the subject of reducing switching artifacts in chopper amplifiers.

Yoshinori Kusuda

Maximize the Performance of Your Sigma-Delta ADC Driver

By **Stuart Servis** and **Miguel Usach Merino**

Introduction

Have you ever checked how many entries are in the web for "design buffer for an ADC"? It can be hard to find what you are looking for among more than 4 million references. Probably not a big surprise to most analog and mixed-signal data acquisition system design engineers as designing the external front end for an unbuffered analog-to-digital converter (ADC) requires a lot of patience and advice. It is often seen as an art form, the preserve of eccentric gurus who have mastered their craft over many years. To the uninitiated, it is a frustrating task of trial and error. Most of the time the frustration becomes an annoying companion due to the number of interrelated specifications, which forces many trade-offs (and evaluations) until the optimum results are achieved.

The Challenge

The design of the amplifier stage consists of two different stages related between them, so the problem becomes difficult to model mathematically, especially due to the nonlinearity associated with both stages. The first step is to select the amplifier that will buffer the sensor output and drive the ADC inputs. The second step is to design a low-pass filter to reduce the input bandwidth, which minimizes the out-of-band noise.

The ideal amplifier provides just enough bandwidth to correctly buffer the signal generated by the sensor or transducer, without adding extra noise, and providing zero power consumption, but the ideal is far away from the real amplifier. In most cases, the amplifier specifications are going to define the overall system performance, especially in terms of noise, distortion, and power. To get a better picture about the problem, the first step is to understand how a discrete time ADC operates.

A discrete-time ADC gets a sample of a continuous time analog signal that is later converted into a digital code. When the signal is sampled, depending on the type of analog converter, there are two different scenarios with the same inherent problem.

SAR ADCs integrate a sample-and-hold, also known as track-and-hold, which is fundamentally a switch and a capacitor that freeze the analog signal until the conversion is done, as shown in Figure 1.

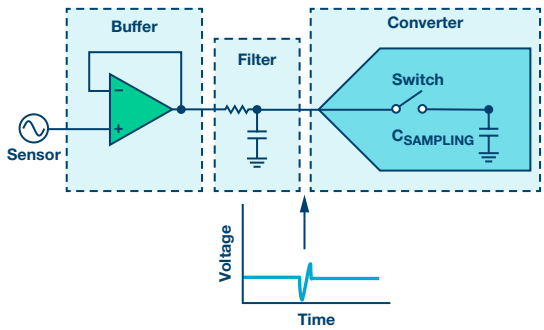

Figure 1. A diagram of sample-and-hold circuitry.

Discrete-time Σ-Δ ADCs, or oversampling converters, implement a similar input stage, that is, an input switch with some internal capacitance. In the case of Σ-Δ ADCs, the sampling mechanism is slightly different, but a similar input architecture of sampling occurs where switches and capacitors are used to hold a copy of the analog input signal.

In both cases, the switch is implemented in a CMOS process with a non-zero value of resistance when closed, typically a few ohms. The combination of this series resistance with the sampling capacitor, in the range of pF, means that the ADC input bandwidth is often very large, and it is in many cases much larger than the ADC sampling frequency.

The Bandwidth Problem

The input signal bandwidth is a problem for the converter. In sampling theory, we know that frequencies above the Nyquist frequency (half of the ADC sampling frequency) should be removed, otherwise those frequencies are going to generate images, or alias, into the band of interest. Noise, typically, has a spectrum where a significant amount of power can exist in the frequency band above the Nyquist frequency of the ADC. Unless we deal with this noise, it will alias down below the Nyquist frequency and add to the noise floor, as shown in Figure 2, effectively reducing the dynamic range of the system.

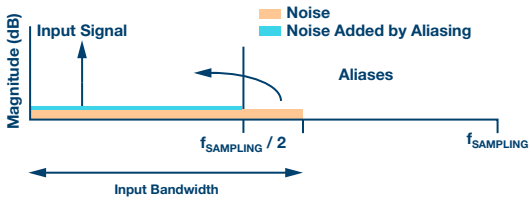

Figure 2. Nyquist folding images.

The ADC input signal bandwidth and, by extension, the buffer output bandwidth represent the first problem to be solved. To ensure that the noise does not get aliased down, the bandwidth of the ADC input signal must be limited. This is not a trivial problem.

Typically, the amplifier is chosen based on the specification for a large signal bandwidth—that is, slew rate—and gain bandwidth product to cover the worst case condition for our input signal, which defines the faster event our ADC can track.

However, the effective noise bandwidth of the amplifier is equivalent to the small signal bandwidth (typically considered for signals less than 10 mV p-p), and this is often at least four to five times higher than the large signal bandwidth.

In other words, if our large signal specifications are selected for 500 kHz, the small signal bandwidth could easily be 2 MHz or 3 MHz, potentially allowing a lot of noise sampled by the ADC. Consequently, the small signal bandwidth should be limited externally before feeding the analog signal into the ADC or the noise measured is going to be three to four times the ADC data sheet specifications, which is not good.

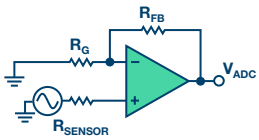

Figure 3. A noninverting amplifier configuration.

Table 1. An Amplifier Noise Referred to Output, RTO

Noise Source	Noise Referred to Output
R_{SENSOR}	$\left(1+\frac{R_{FB}}{R_G}\right)\times\sqrt{4\times k\times T\times R_{SENSOR}}\times\sqrt{BW}$
R_G	$\left(-\frac{R_{FB}}{R_G}\right)\times\sqrt{4\times k\times T\times R_{FB}}\times\sqrt{BW}$
R_{FB}	$\sqrt{4\times k\times T\times R_{FB}}\times\sqrt{BW}$
Amplifier Current Noise	$\sqrt{(I_{NOISE}\times R_{FB})^2+\left[\left(1+\frac{R_{FB}}{R_G}\right)\times I_{NOISE}\times R_{SENSOR}\right)\right]^2}\times\sqrt{BW}$

Remember that the thermal noise generated by the amplifier depends on the amplifier gain, and the total system bandwidth. An example for the circuit is shown in Figure 3 and the noise sources are summarized in Table 1, where:

T is the temperature in Kelvin,

k is Boltzmann's constant (1.38×10^{-23} J/K),

resistor values are expressed in Ω,

and BW refers to small signal bandwidth.

Previous equations make clear the importance of adding a low-pass filter with enough attenuation before the ADC input pin to minimize the noise sampled, as the noise is proportional to the square root of the bandwidth. Typically, a first-order low-pass filter implemented with a discrete resistor and capacitor, with a low enough cutoff frequency eliminates much of the wideband noise. A first-order low-pass filter has the added benefit of reducing the amplitude of any other larger signals outside of the band of interest before they are sampled, and potentially aliased, by the ADC.

However, this is not the whole story. The internal ADC switch resistance and the capacitor defines the analog input bandwidth, but also creates a time-domain charge-discharge cycle due to the varying input signals. Each time the switch (the external circuitry connected to the sampling ADC capacitor) is closed, the internal capacitor voltage may be different from the voltage previously stored on the sampling capacitor.

What Is the Kickback Problem?

The classic analog question: "If you have two capacitors in parallel connected with a switch, the switch is open and one capacitor stores some energy, then what happens to both capacitors when the switch is closed?"

The answer depends on the energy stored by the charged capacitor and the ratio between the capacitors. For example, if both capacitors are of the same value, the energy will be shared between them and the voltage measured between capacitor terminals will be halved, as shown in Figure 4.

Figure 4. Charged (left) and uncharged (right) capacitors.

This is the kickback problem.

Some ADCs perform internal calibrations to compensate for internal errors, known as auto-zero calibration. These procedures bring the sampling capacitor to a voltage close to the rails or another voltage, such as the reference voltage divided by two.

This means that the external signal buffered by the amplifier and the sampling capacitor—that must hold the analog value in order to acquire a fresh sample—are very often not at the same potential (voltage). Consequently, the sampling capacitor must be charged or discharged to bring it to the same potential as the buffer output. The energy required in this process will come from the external capacitor (the one from the low-pass RC filter) and the external buffer. This redistribution of charge, and settling of the voltages, will take a finite amount of time during which the voltage seen at the various points in the circuit will be disturbed, as shown in Figure 1. There can often be a significant amount of charge being redistributed, which is equivalent to current flowing to or from the amplifier and into the capacitors.

The consequence is that the amplifier should be able to charge/discharge the external capacitor of the low-pass filter and the sampling capacitor of the ADC in a very limited time, with the current limiter added by the low-pass filter resistor.

Being more specific, the amplifier should be able to charge/discharge the capacitors within a given error, from the sampling capacitor and external sources. The cutoff frequency of the external low-pass filter should be a little bit higher than the band of interest, which is defined by the time constant of the filter, the number of bits of the ADC, and the worst-case transition between samples—that is, the worst input step that we should able to measure accurately.

How Do We Solve the Kickback Problem?

The easier answer to solve the problem would be to choose an amplifier with enough slew rate, bandwidth gain product, open-loop gain, and CMRR, and place the highest capacitor you may find in the market at the output with a resistor small enough to satisfy low-pass filter bandwidth requirements.

As the capacitor is really big, the kickback problem will be negligible, and the bandwidth is limited by the LP filter, so problem solved, right?

Bad news. The previous solution is not going to work, but if you are curious and try the previous setup, then you will discover two things: the capacitor is going to be the size of a condensed milk container and the amplifier will not like imaginary impedances connected at the output.

The amplifier's performance is dependent on the imaginary load seen by the amplifier. In this case, the low-pass filter penalization is a degradation on THD and settling time. An increase in settling time would have the effect of leading the amplifier to become unable to charge the capacitors such that the voltage that the ADC sample is the correct final voltage. This would lead to further nonlinearity in the output of the ADC.

To illustrate the previous statement, Figure 5 shows the performance difference between different amplifier output currents or resistive loads. Figure 6 shows the small signal overshoot due to the capacitive load, which affects settling time and linearity.

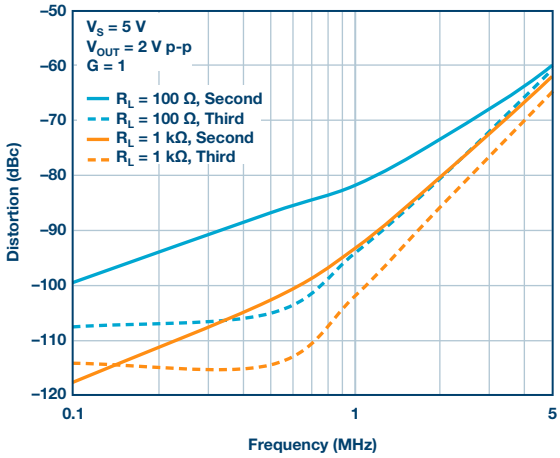

Figure 5. AD4896-2 THD performance vs. load.

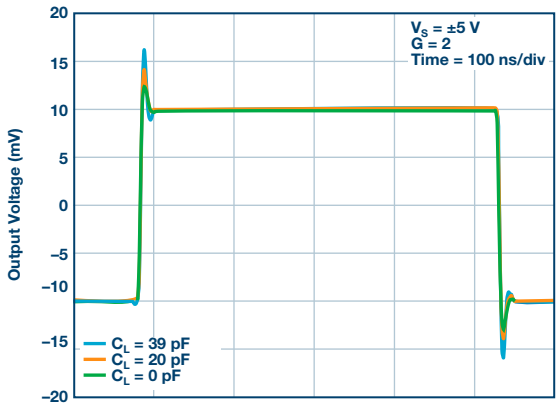

Figure 6. The ADA4896-2's small signal transfer response vs. load.

To minimize this problem, the amplifier output should be isolated from the external capacitor by the series resistor of the low-pass filter.

The resistor should be high enough to guarantee that the buffer is not going to see the imaginary impedance, but small enough to satisfy the required input system bandwidth and minimize the IR drop across the resistor due to the current flowing from the buffer, which can cause a voltage drop that may not be settled quickly enough by the amplifier. In parallel, the resistor should allow the external capacitor to reduce to a value small enough to minimize the kickback without affecting the settling.

You can find more information here.

Luckily for us, there are tools that allow us to predict the combined performance of the DAC, the amplifier, and the filter—like the Precision ADC Driver Tool.

The tool can simulate the kickback, as well as noise and distortion performance, as shown in Figure 7.

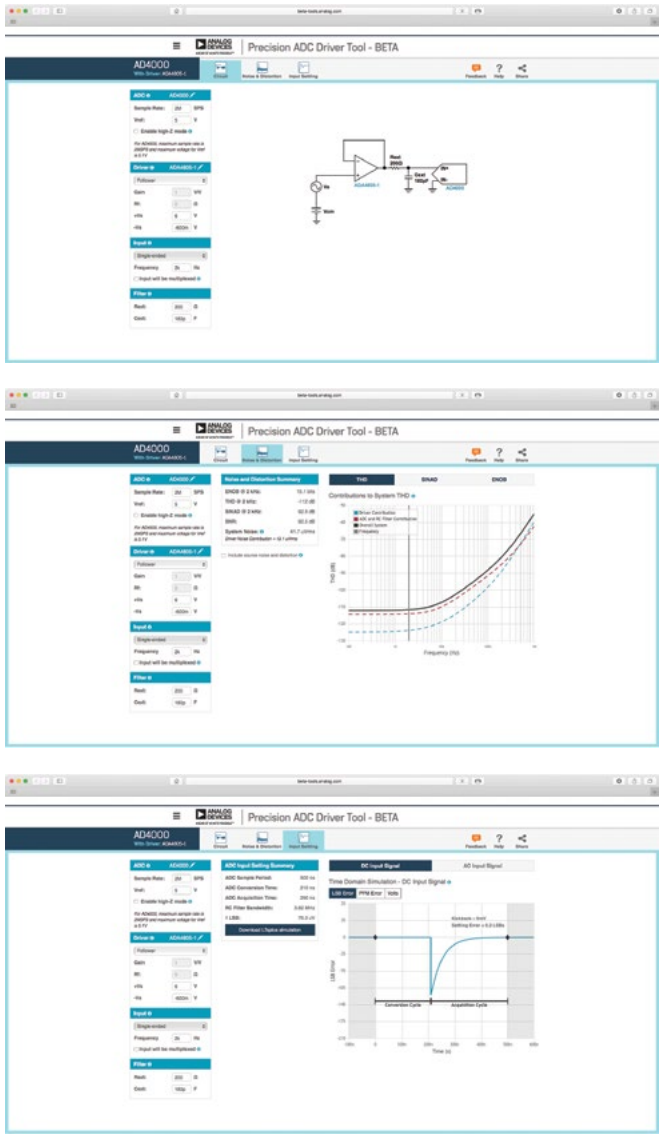

Figure 7. Precision ADC Driver Tool playground simulations.

Rule of Thumb for Low-Pass Filters

Typically, a first-order low-pass filter is seen in many recommendations, but why doesn't anyone use a higher filter order? Unless your device is going to be used in an application with specific requirements to remove larger out-of-band interferers or harmonics in the input signal, increasing the order of the filter will add an extra layer of complexity to your system. In general, the trade-off is leaving the small signal bandwidth a little bit higher than you need, which will impact the noise at the expense of being able to drive the ADC input stage easily—and reduce power and cost—due to the amplifiers.

Reducing the Burden

We previously mentioned that the amplifier does not like imaginary impedance and/or deliver high currents, and this is an element added by the capacitor that is needed to minimize the kickback problem.

The only way to improve the situation is by reducing the kickback itself. This solution has been adopted by the latest ADI converters, like the AD7768 and the AD4000.

The solutions adopted in each of the devices are different due to different converter architectures. The AD4000 SAR ADC can operate at supplies below the analog input range. The solution adopted is called high-Z mode and is only available for sampling frequencies below 100 kHz.

In the AD7768, the supplies equal or become higher than the analog input range. The solution adopted in the AD7768 is called precharge buffer and, as opposed to the high-Z mode, this operates up to the maximum ADC sampling frequency.

Both solutions are based on the same principle of operation; the main difficulty in driving the ADCs is the capacitive charge redistribution. In other words, the lower the voltage drop seen by the input buffer and the low-pass filter when the internal switch reconnects the sampling capacitor, the lower the voltage kick, which minimizes the ADC input current. Consequently, the easier it is to drive the ADC, the greater the settling time is reduced. The voltage drop across the filter resistance reduces, so the ac performance receives a boost.

The effect on the input current against precharge buffer and high-Z enable and disable can be seen in Figure 8.

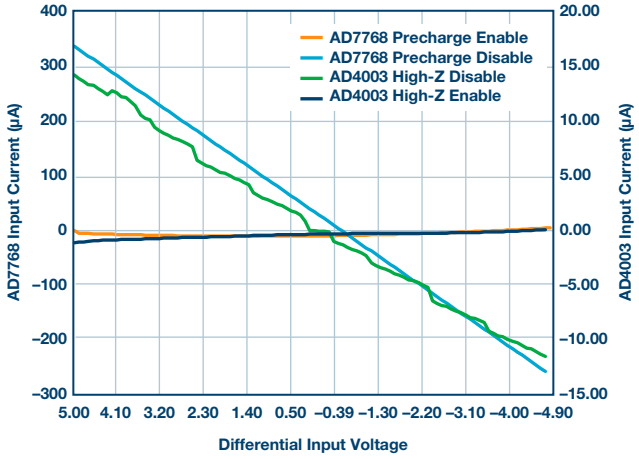

Figure 8. Input current.

The higher the input current, the higher (that is, faster) the amplifier bandwidth should be. Therefore, the higher the input low-pass filter bandwidth should be, and this impacts the noise as well.

For instance, using the SINAD accounts for harmonics as noise performance, for a 1 kHz input signal sampled at 1 MSPS. At different filter cutoff frequencies, we get something like shown in Figure 9.

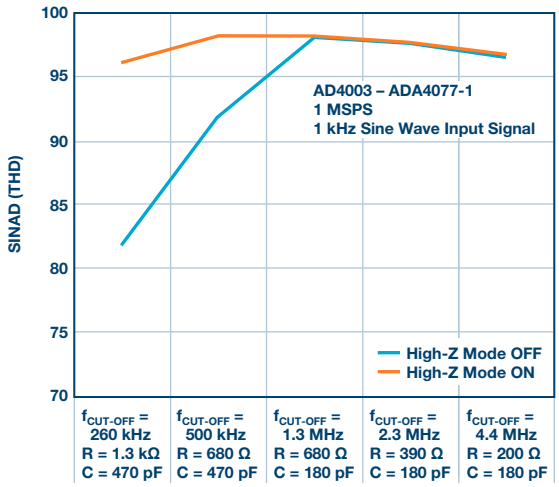

Figure 9. AD4003 SINAD vs. input BW with and without high-Z mode.

The previous figure shows that low input current (high-Z mode on) reduces the cutoff filter frequency requirements, as well as IR drop in the filter resistance value, boosting the ADC performance, compared with the exact same configuration but high-Z mode off.

In Figure 9, it can be observed that by increasing the input filter cutoff frequency, the external amplifier can charge/discharge the sampling capacitor faster at the expense of higher noise. For instance, with high-Z mode on, the noise sampled at 500 kHz is less than at 1.3 MHz. Consequently, the SINAD is better at 500 kHZ input bandwidth. In addition, the capacitance required by the low-pass filter gets reduced, improving amplifier driver performance.

Circuit Design Benefits

The addition of these easier-to-drive, or burden-reducing, features implemented in ADI's latest ADCs has some significant impacts on the overall signal chain. The key advantage that the ADC designer has in bringing some of the drive problem into the ADC silicon itself is that the solution can be designed to be as efficient as possible for the signal requirements of that ADC, thereby solving a few problems including input bandwidth and amplifier stability.

The reduction of the current into the ADC input, and therefore reduction of the kickback, means that the amplifier has a lower voltage step to deal with, but still has the same full sampling period as a standard switched-capacitor input.

Having a smaller step voltage to settle (ramp to final value) over a given period is the same as having a longer period to settle a larger step. The net effect is that the amplifier now does not need to have such a wide

bandwidth to sufficiently settle the input to the same final value. A reduced bandwidth usually means a lower power amplifier.

Another way to look at this is to imagine how an amplifier that would normally not be expected to have enough bandwidth to settle the input of a given ADC can now achieve enough settling when the precharge buffers are enabled.

The ADI application note AN-1384 shows the performance achievable with a range of amplifiers when paired with the AD7768 in each of its three power modes. One of the amplifiers shown in this document, the ADA4500-2, shows that it struggles to settle the input of the AD7768 in median power mode (THD is >−96 dB) when the precharge buffers are not used. However, when precharge buffers are enabled, the performance improves dramatically to better than −110 dB THD.

Since the ADA4500-2 is a 10 MHz bandwidth amplifier and the bandwidth required to settle the AD7768 in the given mode is approximately 12 MHz, we see that the use of this lower bandwidth amplifier has now been enabled by the easy-to-drive feature. In this way, these features not only enable easier design of the front-end buffer circuit, but also allow more freedom in selecting components to stay within system power or thermal ceilings.

A secondary advantage to having reduced current into the analog input pin of the ADC is that there is now less current flowing through the series resistor that is used as part of the input RC network.

For traditional ADC inputs, the relatively large current meant that only small value resistors could be used, otherwise large voltage drops would be developed across this resistor. A large voltage drop here can lead to gain error or linearity errors in the ADC conversion result.

However, using smaller resistor values also has its challenges. Achieving the same bandwidth of the RC using a smaller resistor means using a larger cap instead. However, this large cap and small resistor combination can lead to instability of the buffer amplifier.

The reduced current, encountered when using the easy-to-drive features, means that larger value resistors can be used without affecting the performance, and ensuring the system's stability.

Circuit Performance Benefits

Taking what we have already stated as the benefits to the circuit design, it becomes clear that there are also performance benefits, or opportunities to further improve performance, using these features.

The benefit already mentioned, being able to achieve better performance with lower bandwidth amplifiers, can also be employed to extend the performance of more optimal systems. For example, even with a well-settled input signal, there may still be some mismatch between the inputs as that final settling is happening. So, enabling the precharge buffers, for example, will mean that this final settling is much smaller, so supreme levels of THD can then be achieved where previously this would not have been possible.

The reduction in the current going through the series resistor of the RC network also benefits the performance. Also, not only is the input current significantly reduced, it has almost no dependence on the input voltage. Improved THD can be achieved since any mismatches in the resistors on the input pair will lead to smaller voltage differences seen at the ADC input, as well as the voltage drops not having a signal dependency.

The lower input current also has an effect on offset and gain accuracy. Because of the reduction in absolute current, as well as the reduction in signal dependent current changes, there are less chance that variation in component values across each channel or across each physical board will lead to large variation in offset and gain errors (for the same reason, lower current leads to smaller voltages across series resistances). With precharge buffers, better absolute offset and gain error specifications can be achieved, as well as consistent performance across boards or channels within a system.

There is another benefit to the lower current in systems where the ADC sampling rate changes to adapt to different signal acquisition needs, such as in data acquisition cards. Without precharge buffers, the voltage drops across the input passive components vary with the sample rate of the ADC, since the ADC input cap will charge and discharge more often at higher sampling rates. This applies to both the analog input path and the reference input path, and this variation in voltage is seen by the ADC as sample rate dependent offset and gain errors.

However, with precharge buffers enabled, the absolute current, and therefore the absolute voltage drop, is much smaller to begin with, so the change in voltage as the ADC sample rate changes will also be much lower. In the end system, this means that there is less need to recalibrate the system offset and gain errors as the sample rate is adjusted, and that offset and gain errors are less sensitive to changes in ADC sampling rates.

Cost Benefits

One of the main benefits of the easy to use features is related to the total cost of ownership. The different facets of the design and performance benefits lead to a potentially lower development cost and operational cost.

- Easier design means less design effort and faster time to a first prototype.
- Easier design means greater chance of success the first time in prototype design.
- The easier-to-drive feature may allow lower bandwidth and therefore lower cost amplifiers to be used.
- Offset and gain benefits may allow reduced calibration at factory.
- Performance improvement may allow reduced calibration in field or on-demand, thus leading to reduced downtime and/or improved throughput.

Real Examples Using AD7768-1

Table 2 shows some measured data from the AN-1384 application note, which helps designers choose the appropriate amplifier to drive the AD7768-1 ADC. The examples in the table show that there are significant possible improvements when the precharge feature is enabled on some amplifiers. The reason for the improvement in THD in particular is due to a combination of the previously mentioned effects of the reduced burden that the ADC places on the drive circuitry. For example, the configuration using the ADA4945-1 amplifier achieves a 4 dB improvement in THD when the precharge buffers are enabled. Similarly, the ADA4807-2 circuit can achieve an increase of 18 dB in THD. What these examples show is that amplifiers that are able to achieve reasonable performance on their own can achieve headline levels of performance when used in conjunction with the easy-to-drive features available in many of ADI's newest ADCs.

Table 2. AD7768-1 Performance with Various Amplifiers

Amplifier	Precharge Buffer	SNR (dB)	THD (dB)	SINAD (dB)
ADA4940-1	Disabled	105.4	–114.5	105.0
ADA4940-1	Enabled	105.2	–120.4	105.1
ADA4807-2	Disabled	105.1	–105.7	102.6
ADA4945-1	Disabled	105.9	–116.6	105.6
ADA4896-2	Disabled	106.7	–118.0	106.5
ADA4807-2	Enabled	104.9	–123.7	104.8
ADA4945-1	Enabled	106.0	–120.7	105.8
ADA4896-2	Enabled	105.5	–130	106.4

Conclusion

Designing a circuit to drive an unbuffered ADC is not a trivial thing and requires a proper methodology and trade-off considerations due to the kick-back of the converter and bandwidth requirements. Many times, the required circuitry is going to define the overall system performance in terms of THD, SNR, and power consumption.

The latest ADI precision converters for both SAR and Σ-Δ technologies integrate a set of features to minimize the converter input current. This minimizes the kickback, greatly reducing and simplifying the external circuitry, achieving specifications numbers that were not previously possible. This makes the SAR and Σ-Δ technologies easier to use, reduces the engineering time, and improves system specifications.

Stuart Servis [stuart.servis@analog.com] is a product applications engineer at Analog Devices, where he works in the precision signal chains team within the Instrumentation and Precision Technology Group. His area of expertise is precision data acquisition signal chains, based on Σ-Δ and SAR ADCs. He received his B.Sc. in applied physics and electronics from the National University of Ireland in Galway.

Stuart Servis

Miguel Usach Merino [miguel.usach@analog.com] received his degree in electronic engineering from Universitat de Valencia. Miguel joined ADI in 2008 and works as an applications engineer in the Linear and Precision Technology Group in Valencia, Spain

Miguel Usach Merino

The Evolutionary Path to the 100 A µModule Regulator

By **Tony Armstrong**

Introduction

A µModule device resembles a surface-mount IC, but they include all the necessary support components that would normally be used to construct a power conversion circuit. This includes a dc-to-dc controller, MOSFET dice, magnetics, capacitors, resistors, and so on, all mounted on a thermally efficient laminate substrate. They are then encapsulated using a plastic mold cap. The result is a complete power supply that can be simply adhered to a printed circuit board (PCB).

Built to the industry's highest quality standards, this product family dramatically reduces risk, time, and effort to successfully design high performance, high power density solutions. It is as if we have taken all our power supply expertise and know-how within Analog Devices and put it into an IC-like form factor. For those of you who have had time constraints for designing power conversion circuits where mass production is only weeks away—and you have had to spend many late nights debugging your supply late into the early morning—those times will be a thing of the past if you use a µModule regulator instead of a "do-it-yourself" discrete solution.

Taking a closer look at the internal construction of a typical µModule product, you will notice that the package options are a land grid array (LGA) or ball grid array (BGA). The internal components used to form the internal switch-mode power conversion circuit can be in die form, while others are finished packaged products. Nevertheless, these components are all mounted onto a bismaleimide triazine (commonly known as BT) laminate substrate, which has excellent electrical and thermal properties. Moreover, µModule products are not just about integration, as they provide other properties and performance characteristics over competing alternatives.

Power design expertise is declining on a global basis and there are simply not enough power supply design resources to develop every single power supply at most customers' sites. It is reported by the trade press that the average age of a degreed engineer is 57 years—and this is a global statistic, with China having the youngest average.

The top three concerns of power design engineers are:

- Insufficient people to get the job done.
- Finding the optimal components for their design.
- Time-to-market pressures.

Because of these trends, we wanted to deliver a complete power supply that is ready to use off-the-shelf and with all the performance criteria required for the end application. Furthermore, at the same time, PCB area is at a premium since everyone is trying to pack even more functionality and capability into a smaller space. If this is not bad enough already, thermal design constraints are becoming more complex as designers try to pack more capability into an ever-smaller space while delivering more power in an environment that has limited airflow for cooling purposes.

Finally, time-to-market pressures are intense as the power supply is one of the last items to be designed in a system and mass production starts can be just weeks away!

PCB area is a critical priority in most designs. For example, any given datacom or telecommunications board is bound to be laden with many digital processors, ASICs, and memory. All these need to be powered on the board while voltage levels vary from just north of 5 V to as low as 0.6 V, after an intermediate system bus voltage which varies from 12 V to 48 V. At the same time, system designers are continually being asked to pack more functionality into these ever-shrinking form factors—which are probably mutually exclusive!

Design Problems That Needed to Be Solved

Thermal design constraints are becoming more severe. As more and more functionality is packed onto the PCB, the overall power levels needed to power them on the board are increasing. Meanwhile, cooling is at a premium due to heat sinking space constraints and limited air flow volume. This is a headache for the designer, since there is a maximum internal ambient temperature constraint on the system that cannot be violated without compromising performance and long-term reliability.

Time-to-market pressures have increased dramatically in recent years due to both competitive pressures and the need for faster revenue streams. So, the power supply designer is under the gun for getting his power conversion circuits designed and functional in weeks, if not days!

Simply put, µModule products provide a "simple and done" proven power conversion solution. So, using them means no more late nights of debugging power supplies in the lab!

Of course, these products need to have rigorous quality and long-term reliability to ensure long operational life once they have been deployed in an end system. Accordingly, ADI has engaged in rigorous quality and reliability testing to ensure long-term deployments in harsh environments.

The following is a summary of the testing and data we have accumulated since the introduction of our first µModule product, the LTM4600, back in October of 2005. This includes:

- Over 22 million power cycles.
- Over 5 million device hours of hot temperature operational life.
- Over 2 million hours of mounted temperature cycles to ensure that these modules can operate 24 hours per day, 7 days per week, 365 days per year for a decade without having any intermittent contact issues from the package leads to the PCB.
- Over 25 million temp cycles from –65°C to +150°C.
- Over 16 million thermal shock cycles from –65°C to +150°C. And remember, this is liquid-to-liquid on a finished power supply!

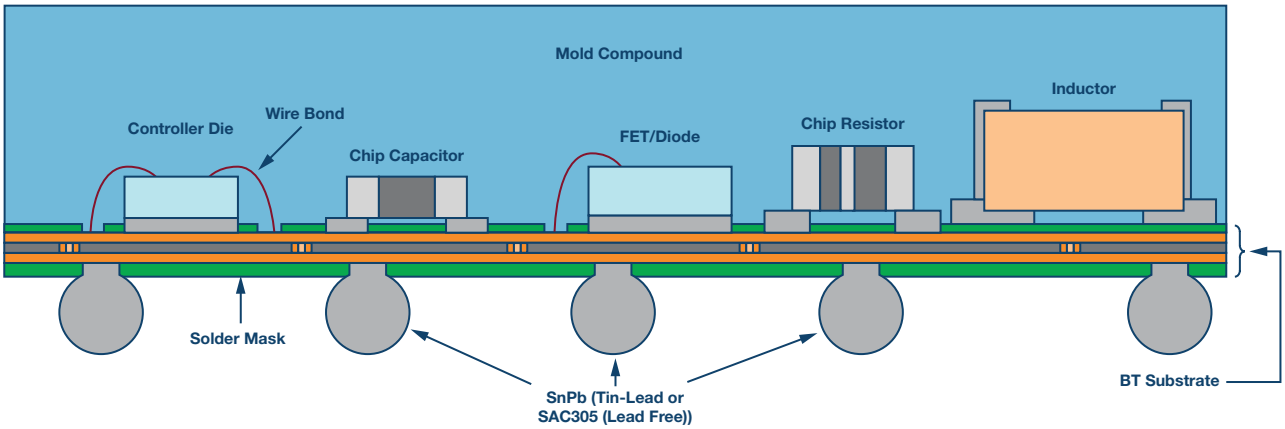

Figure 1. A µModule regulator BGA package cross-section.

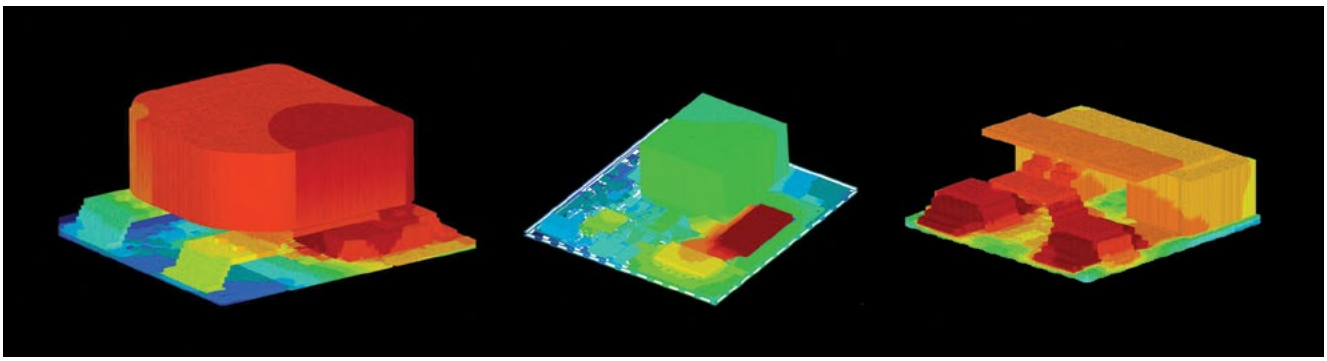

Figure 2. Thermal imaging of µModule regulator construction and heat dissipation.

The end outcome is an FIT rate of <0.4. To put this in perspective, this equates to 0.4 device failures in every billion device hours of operation. And this is on a complete power supply. To put this in context, many of our competitors' integrated circuits (single silicon in a package) have higher FIT rates!

Packaging Evolution

Let's take a closer look at the µModule packaging options. When we first introduced the LTM4600 way back in 2005, we used an LGA package option. The thinking at the time was that since a lot of the VLSI digital ICs had similar LGA form factors, it would be easy for the user to use our µModule products. While this was true some of the time, it was not true all of the time.

Accordingly, it was decided that having a BGA package option would be an innovative idea, too. This turned out to be fortuitous for two reasons. First, it was easier for users who were not used to high volume production of the LGAs. And, secondly, it was easier to put solder balls on the round pin pads. Moreover, it allows for both leaded and lead-free solder balls. And since many µModule device users are in the aerospace and military market segments, they were very happy about this.

Accordingly, the first part we introduced in this product category was named the LTM4600: a 4.5 V to 20 V input/0.6 V to 5.5 V output and up to 10 A of continuous output current. It was in a 15 mm × 15 mm × 2.82 mm surface-mount LGA package. Its application was 12 V_{IN} to 3.3 V_{OUT} at 10 A with 90% efficiency. Remember, this was October 2005, so this level of performance was ground breaking stuff.

Nevertheless, one of our key metrics was to improve the µModule regulator's thermal performance so that we could increase its output current density while remaining in same the 15 mm × 15 mm footprint. Since there was clearly a significant thermal issue, we needed to address getting the heat out of the package. To facilitate this goal, our designers had decided to use a BT laminate substrate because it had excellent thermal properties and facilitated taking heat through the bottom of the µModule package and into the PCB where it could be dissipated. While this was acceptable in the mid-2000s, by the time another 5 years went by, our customer-base informed us that they could no longer dissipate most of the heat through their PCB. Instead, it needed to be able to pull heat out of the top of the package and dissipate it into the air! Therefore, we designed a special heat sink that was encapsulated inside the package and connected to the top of the internal MOSFETs and inductors. This heat sink was exposed on the top side of the µModule regulator. Now the user could add their own heat sink on top of the µModule device to improve pulling heat out of it. If they had 200 LFM or airflow, they could also facilitate better thermal performance. Truly, a win-win scenario.

Regardless of this enhancement, we continued to evolve and developed µModule regulators with inductors on top since these acted as heat sinks to further improve thermal dissipation qualities.

Finally, it should be mentioned why we introduced our ultrathin µModule devices. We realized that, in many cases, our customers would only utilize the underside of their PCBs with discrete components due to space limitations. It turns out that for many rack-mounted systems, there was a 2.2 mm height restriction for mounting components on the underside of a PCB. Therefore, we developed µModule regulators with 1.8 mm and 1.9 mm maximum heights so that they would easily fit while also helping with space and density issues.

Now, with this background, it is easy to comprehend what is going on with a µModule devices' thermal performance. This can be thought of as an evolutionary pathway that has allowed for a continuous improvement of our µModule thermal performance from inception through to our current offerings—a decade plus journey.

Figure 2 shows three thermal imaging photographs, representing several types of µModule regulators with varying types of construction with the goal of increasing the device's ability to pull heat out of the top of the µModule into free air where it can facilitate additional cooling from air flow within the system, or could also have an additional heat sink that is shared with the VLSI digital ICs that are usually present. The color blue shows low temperature (minimal power dissipation) and colors orange through red show elevated temperatures (high power dissipation). Of course, this is what we want to occur when we want the heat generated by the power conversion process to be pulled out into free air and not into the PCB.

While we have been improving the thermal performance characteristics of this product offering, we have simultaneously continued to increase µModule regulators' power density by putting them into ever-shrinking form factors. Figure 3 shows the LTM4627, a 20 V input device that can deliver a 15 A output current as low as 0.6 V with an efficiency in the nominal 90% range depending on the specific V_{IN} and V_{OUT} conditions. Next to this is the LTM4638, which is also a 20 V input device and can deliver a 15 A output as low as 0.6 V with nominal 86% efficiency—pretty close! However, volumetrically, the LTM4638 is 5.6 times less than the LTM4627. See Figure 3 for a size comparison.

The point being that there is only a small delta in conversion efficiency between these two parts under the same operating conditions, but the footprint and space required for its implementation are orders of magnitude less. All this improvement has all been attained in less than 4 years.

Figure 3. LTM4627 (15 mm × 15 mm × 4.92 mm) vs. its new, smaller equivalent, the LTM4638 (6.25 mm × 6.25 mm × 5.02 mm).

Sojourn to a Single 100 A µModule Device

For a long time, our existing users of high power µModule packages kept asking us for smaller, more efficient, and higher current density devices—even though this feature set might be considered mutually exclusive. Nevertheless, our design team took this request to heart and began to figure out a way to get us there.

From a historic perspective, back in the 2013 to 2016 timeframe, we had µModule regulators in the 15 mm × 15 mm footprint that are capable of delivering output currents in the 26 A to 50 A per device range. It should also be noted that a key matrix measurement for our high power µModule devices is that they should be able to deliver full rated output current from a 12 V input to a 1 V output with 90% conversion efficiency. The reasoning is that dealing with 10% power lost as heat is usually thermally acceptable within most applications. By late 2016, our 40+ A µModule regulators had efficiencies in the 88% to 89% range—which is very close to this goal.

The progression to get to a 100 A single µModule regulator required us to use multiple devices, as shown below:

Thus in 2010, having 12 LTM4601s in a Polyphase® parallel configuration allowed us to deliver a 100 A output from a 12 V to 1 V output.

In 2012, only four LTM4620s in a Polyphase parallel configuration allowed us to deliver a 100 A output from a 12 V to 1 V output.

In 2014, only three LTM4630s in a Polyphase parallel configuration allowed us to deliver a 100 A output from a 12 V to 1 V output.

In 2016, only two LTM4650s in a Polyphase parallel configuration allowed us to deliver a 100 A output from a 12 V to 1 V output. And we have ±1% total dc error over line, load, and temperature.

Finally, in November 2018, we introduced the LTM4700—a dual 50 A or single 100 A output µModule regulator. See Figure 4 for an image of the actual device.

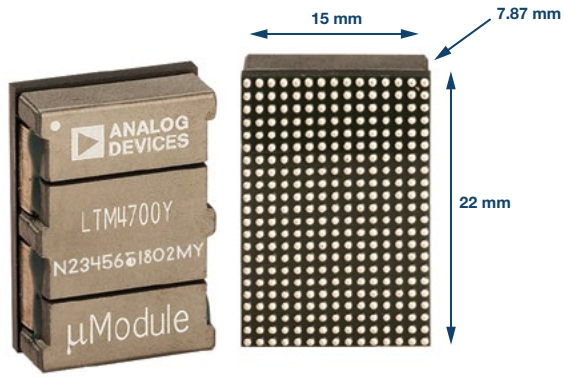

Figure 4. The LTM4700 is capable of delivering up to 100 A output current in a single package.

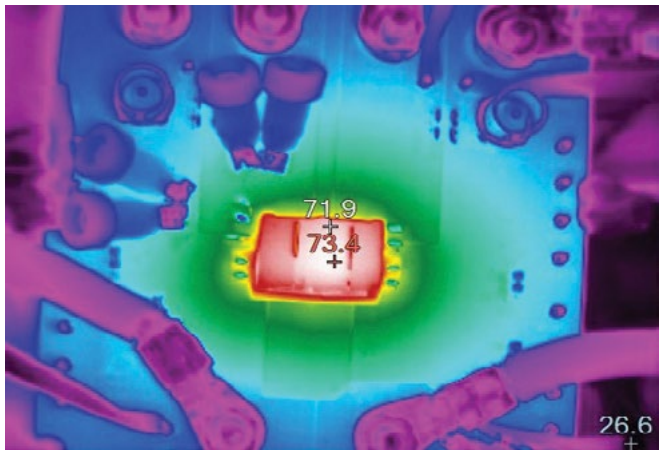

Figure 5. The LTM4700 100 A µModule (89.6% efficiency).

Figure 5 shows a thermal image of the LTM4700 during normal operation. The operation conditions are 12 V to 1 V delivering 100 A of current with high conversion efficiency and only 200 LFM of airflow. As a result, its best-in-class energy efficient performance makes it an excellent choice to reduce data center infrastructure cooling requirements.

Taking a closer look at some of the LTM4700's key specifications:

- It is a single 100 A output capable µModule device. It can also be used as two 50 A outputs.
- It is very close to 90% conversion efficiency when stepping down from 12 V to 1 V at 100 A with only 200 LFM air flow. And it has ±0.5% maximum dc error over temperature.
- Its x, y, z footprint is 15 mm × 22 mm × 7.82 mm.

In addition to having a dual 50 A, or single 100 A output, the LTM4700 also incorporates a PMBus I²C interface or power system management (PSM).

This enables many different capabilities, including:

- Configure voltages, define complex on/off sequencing arrangements, define fault conditions such as OV and UV limits, and set important power supply parameters such as switching frequency, current limit, etc., over a digital communication bus.
- Over the same communication bus, you can readback important operating parameters such as input and output voltage, input and output current, input and output power, internal and external temperature, and, in some of our products, measure energy consumed.
- Users can implement very precise closed-loop margin testing of their designs, as well as trim power supply voltages to very precise levels.
- PSM devices enable higher reliability and quality.
- Our built-in servo loops will maintain higher power supply accuracy over the life of the product, improving reliability.
- The readback features of our PSM devices can be used to improve test coverage at in-circuit test and screen out possible defective devices before they get into the field.
- During the life of the customer's product, our PSM devices continue to monitor important parameters. Trends in voltage, current, and temperature can be used to profile the power system. Once a good system signature can be found, a flawed system, or one that is about to fail, can be identified.

Conclusion

We introduced our first µModule regulator back in 2005, the LTM4600. It came in a 15 mm × 15 mm × 2.8 mm LGA package and could deliver 10 A of output current from a 12 V input to a 1.2 V output with 89% efficiency. Fast forward 13 years, and the LTM4700 can deliver 100 A from 12 V to 1 V with 89.6% efficiency (and 200 LFM air flow). But that's not all: our designers are already working on other modules that can enable even more performance and capabilities.

Tony Armstrong [tony.armstrong@analog.com] is currently the product marketing director for Analog Devices' Power by Linear™ product group. He is responsible for all aspects of power conversion and management products from their introduction through obsolescence. Prior to joining ADI, Tony held various positions in marketing, sales, and operations at Linear Technology, Siliconix Inc., Semtech Corp., Fairchild Semiconductors, and Intel. He attained his B.S. (Honors) in applied mathematics from the University of Manchester, England.

Tony Armstrong

Rarely Asked Questions—Issue 165 Discrete Difference Amplifier vs. an Integrated Solution

By **Jordyn Ansari** and **Chau Tran**

Question:

Why pay more for less?

Answer:

The classic discrete difference amplifier design is quite simple. What can be complicated about an op amp and a four resistor network?

However, the performance of this circuit may not be as great as the designers would like. Based on actual production designs, this article shows some of the disadvantages encountered with discrete resistors, including gain accuracy, gain drift, ac common-mode rejection (CMR), and offset drift.

The classic four resistor difference amplifier, shown in Figure 1.

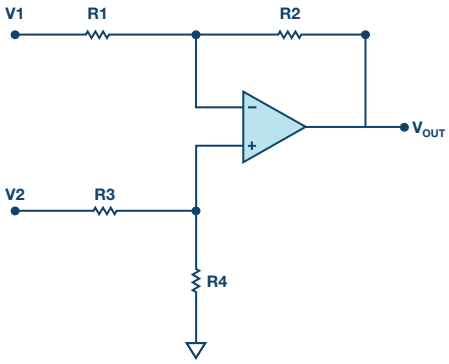

Figure 1. The classic discrete difference amplifier.

The transfer function of this amplifier is:

$$V_{OUT} = \left(\frac{R4}{R3+R4}\right)\times\left(\frac{R1+R2}{R1}\right)\times V2 - \left(\frac{R2}{R1}\right)\times V1$$

With R1 = R3 and R2 = R4, Equation 1 simplifies to

$$V_{OUT} = \left(\frac{R2}{R1}\right)\times(V2 - V1)$$

This simplification can be a quick way to approximate the expected signal, but those resistors are never exactly equal. In addition, the resistors usually have low accuracy and a high temperature coefficient, contributing significant errors to the circuit.

For example, using a good op amp and standard 1%, 100 ppm/°C gain-setting resistors, the initial gain error can be up to 2% and may vary by up to 200 ppm/°C. One solution to this issue would be to use monolithic resistor networks for precise gain setting, but those are bulky and expensive. In addition to the low accuracy and significant drift over temperature, most discrete differential op amp circuits have poor CMR and an input voltage range smaller than the power supply voltage. Also, monolithic instrumentation amplifiers (in-amps) will have a gain drift because the internal resistor network of the pre-amps does not match with the external gain setting resistor going to the RG pin.

The best solution to all of these issues is to use a difference amplifier with internal gain setting resistors, such as the AD8271. Typically, these products consist of a high precision, low distortion op amp and several trimmed resistors. These resistors can be connected to create a wide variety of amplifier configurations, including difference, noninverting, and inverting configurations. The resistors on the chip can be connected in parallel for a wider range of options. Using the on-chip resistors provides the designer with several advantages over a discrete design.

Much of the dc performance of op amp circuits depends on the accuracy of the surrounding resistors. Those internal resistors are laid out to be tightly matched, and are laser trimmed and tested for their matching accuracy. Because of this, there is a guarantee of high accuracy for specifications such as gain drift, common-mode rejection, and gain error. The circuit in Figure 1, when integrated, can provide 0.1% gain accuracy with less than 10 ppm/°C gain drift, as seen in Figure 2.

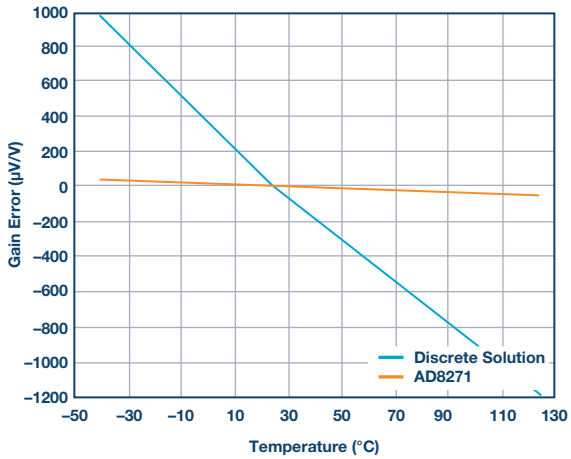

Figure 2. Gain error vs. temperature comparison—AD8271 vs. discrete solution.

AC Performance

The circuit size is much smaller in an integrated circuit than on a printed circuit board (PCB), so the corresponding parasitic factors are also smaller and therefore help the ac performance. For example, the positive and negative input terminals of the AD8271 op amp are not pinned out intentionally. By not connecting these nodes to the traces on the PCB, the capacitance remains low, resulting in both improved loop stability and common-mode rejection over frequency. See Figure 3 for the performance comparison.

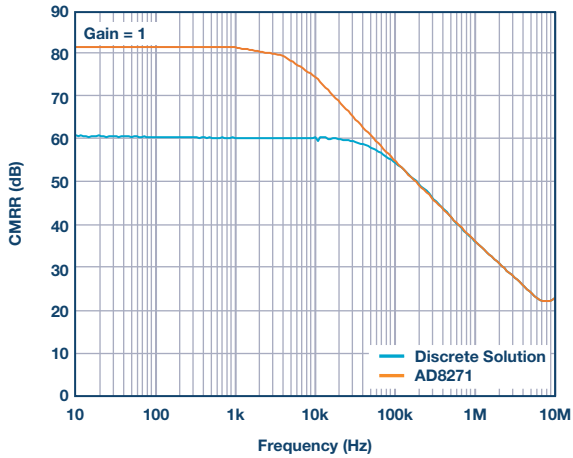

Figure 3. CMRR vs. frequency comparison—AD8271 vs. discrete solution CMRR.

An important function of the difference amplifier is to reject signals that are common to both inputs. Referring to Figure 1, if the resistors R1 through R4 are not perfectly matched (or if R1, R2 and R3, R4 are not ratio matched when the gain is greater than 1), part of the common-mode voltage will be amplified by the difference amplifier and appear at V_{OUT} as a valid difference between V1 and V2 that cannot be distinguished from a real signal. If the resistors are not perfect, part of the common-mode voltage will be amplified by the difference amplifier and appear at V_{OUT} as a valid difference between V1 and V2 that cannot be distinguished from a real signal. The ability of the difference amplifier to reject this is called common-mode rejection. This can be expressed as common-mode rejection ratio (CMRR) or converted to decibels (dB). With the discrete solution, the resistors are not as well matched as the laser trimmed ones within the integrated solution, as seen by the output voltage vs. CMV plot in Figure 4.

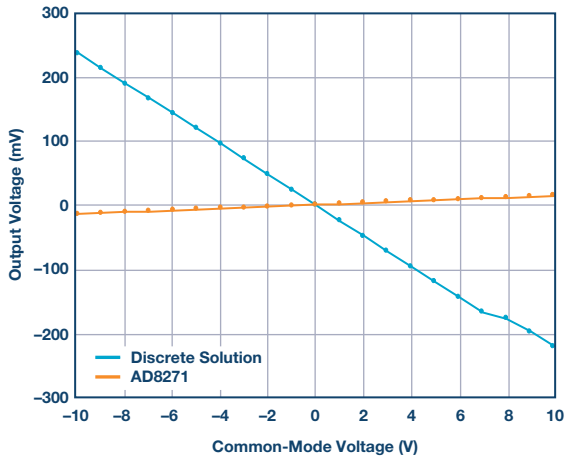

Figure 4. Output voltage vs. common-mode voltage—AD8271 vs. discrete solution.

The CMRR, assuming a perfect op amp, is:

$$CMRR \cong \frac{A_d + 1}{4t}$$

where A_d is the gain of the difference amplifier and t is the resistor tolerance. Thus, with unity gain and 1% resistors, the CMRR is 50 V/V, or about 34 dB; with 0.1% resistors, the CMRR increases to 54 dB. Even given a perfect op amp with infinite common-mode rejection, the overall CMRR is limited by resistor matching. Some low cost op amps have a minimum CMRR in the 60 dB to 70 dB range, making the errors worse.

Low Tolerance Resistors

While amplifiers usually perform well inside of their specified operating temperature range, the temperature coefficient of the external discrete resistors must be taken into account. In the case of an amplifier with integrated resistors, the resistors can be drift trimmed and matched. The layout usually has the resistors close together, so that they all drift together, thus reducing their offset temperature coefficient. In the discrete case, the resistors are spread out further around the PCB and are not as well matched as the integrated case, producing a worse offset temperature coefficient, as shown in Figure 5.

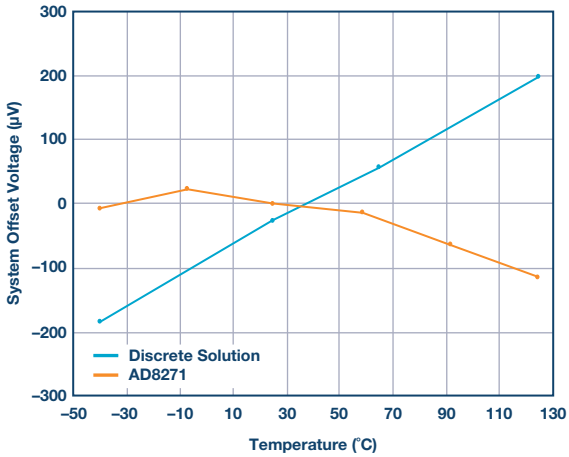

Figure 5. System offset vs. temperature—AD8271 vs. discrete solution.

The four resistor difference amplifier—whether discrete or monolithic—is widely used. With only one part, rather than several discrete components, placed on the PCB, the board can be built more quickly, efficiently, and with significant area savings.

To achieve a solid, production-worthy design, carefully consider noise gain, input voltage range, and a CMR of 80 dB or better. These resistors are also manufactured from the same low drift, thin film material, so their ratio match over temperature is excellent.

Conclusion

It is easy to see the difference between amplifiers with internal gain setting resistors and discrete difference amplifiers.

Jordyn Ansari

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

Chau Tran

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

JESD204C Primer: What's New and in It for You—Part 1

By **Del Jones**

Data-intensive applications across many industries continue to push the boundaries for delivering payload data fast and efficiently. 5G communications networks employ systems demanding more bandwidth in the infrastructure and its connecting components. In aerospace and defense industries, this translates into processing more information in a shorter amount of time in radar applications and complex data analysis instruments. Relatedly, testing and analysis of this rapid expansion in bandwidth translates into the need for higher speed and capacity in electronic test equipment.

This ever-increasing demand for data has led to the need for the JEDEC Solid State Technology Association to introduce the latest evolution in the JESD204 standard for high speed serial links between data converters and logic devices. The B revision of the standard, released in 2011, pushed the serial link data rates to 12.5 Gbps and ensured deterministic latency from one power cycle to the next while enabling the higher bandwidth requirements of the converter-based applications at the time. The newest revision of the standard, JESD204C, was released late in 2017 to continue to support the upward trend in performance requirements for this and next generation's multigigabit data processing systems. The JESD204C subcommittee established four high level goals for this new revision of the standard: increase the lane rates to support even higher bandwidth applications' needs, improve the efficiency of payload delivery, and provide for an improved robustness of the link. In addition, they wanted to write a spec that had more clarity than JESD204B while also fixing some of the errors that were in that version of the standard. It was also desired that a backward-compatible option to JESD204B be made available. The complete JESD204C specification is available through JEDEC.

This two-part primer serves as an introduction to the JESD204C standard by highlighting the differences from JESD204B and detailing the key new features intended to meet the previously stated goals and make for a more user-friendly interface while delivering the bandwidth capability needed for a variety of industries. Part one of this series provides a high level view of these differences and the new features. Part two will dive a little deeper into the most important new features.

Summarizing the Changes for JESD204C

The JESD204C specification has been organized for improved readability and clarity, and it includes five major sections. The "Introduction and Common Requirements" section covers requirements that apply to all layers of the implementation. The sections for the physical, transport, and each of the data link layers (8b/10b, 64b/66b, and 64b/80b) cover requirements that apply specifically to those layers of the implementation. Several new terms are introduced throughout the standard, mostly associated with the new 64b/66b and 64b/80b link layers as well as the new synchronization process for these link layers. While the transport layer remains intact from JESD204B, the physical layer has undergone quite a bit of change. The aforementioned changes, along with small changes to clocking and synchronization and the addition of forward error correction (FEC), are all summarized in the following sections.

New Terminology

There are several new terms and configuration parameters introduced in JESD204C that are primarily used to describe the functions associated with the 64b/66b and 64b/80b link layers. Table 1 lists the most relevant terms and parameters along with a brief description of each. These will be described further in the following sections.

Table 1. New Terms and Parameters

Term	Definition
Block	A structure starting with a 2-bit sync header containing either 66 or 80 (BkW) bits total
BkW	Block width; the number of bits in a block
cmd	Command, as related to the command channel
Command Channel	Data stream using extra bandwidth afforded from sync headers
E	The number of multiblocks in an extended multiblock
EMB_LOCK	A state that asserts that extended multiblock alignment has been achieved
EoEMB	End of extended multiblock identifier bit
EoMB	End-of-multiblock sequence (00001); also known as the pilot signal
Extended Multiblock	A set of data containing one or more multiblocks
FEC	Forward error correction
Fill Bit	A bit used to artificially extend the block size in 64b/80b encoding mode
LEMC	Local extended multiblock clock
Multiblock	A set of data containing 32 blocks
PCS	Physical coding sublayer
SH	Sync header
SH_LOCK	A state that asserts that sync header alignment has been achieved
Sync Header	Two bits, which guarantee a transition preceding every block

Transport Layer

For JESD204C, the transport layer remains intact from JESD204B. The frames of data assembled in the transport layer are sent across the link in 8-octet blocks. Changes to the organization, text, and figures have been made to this section of the standard to provide improved clarity.

Due to the nature of the 64-bit encoding schemes, there are some configurations where frame boundaries will not align with the block boundaries (frames may not include exactly eight octets). The details and implications of this will be covered in part two of this series.

Data Link Layer

As previously implied, there are two major sections of the standard that cover the different data link layer schemes. The 8b/10b encoding scheme from previous versions of the JESD204 standard, including the use of the SYNC~ pin and use of K.28 characters for synchronization, lane alignment, and error monitoring, remains intact as a backward-compatible option. However, most applications, in the long-term, are likely to use one of the new 64-bit encoding schemes that have been added in JESD204C. The 64b/66b scheme will provide the highest efficiency and is based on IEEE 802.3. While it is referred to as encoding, there is not actually any encoding (similar to 8b/10b) going on. The scheme just adds two header bits to 64 bits of payload data. Since this is the case, scrambling is made mandatory so that dc balance is maintained and to ensure enough transition density so that the clock and data recovery (CDR) circuitry in the JESD204C receiver can reliably recover the clock. This will be covered in more detail in part two of this series. A 64b/80b option has also been added that keeps the same clock ratios as the 8b/10b scheme while allowing for the use of new features like forward error correction. Neither of the 64-bit encoding schemes is compatible with the 8b/10b encoding used in JESD204B.

Physical Layer

JESD204C has increased the upper limit on lane rates to 32 Gbps while maintaining the lower limit of 312.5 Mbps established in earlier revisions. The upper limit in JESD204B is 12.5 Gbps. While not strictly forbidden, 8b/10b encoding is not recommended for lane rates above 16 Gbps and neither of the 64b schemes are recommended for lane rates below 6 Gbps.

JESD204C introduces two categories of classes to define the characteristics of the physical interface. Table 2 lists the lane rate associated with each class. Table 3 lists the channel types within Class C and the associated emphasis and equalization characteristics.

Table 2. Lane Data Rates for Data Interface Classes

Data Interface Class	Minimum Data Rate (Gbps)	Maximum Data Rate (Gbps)
B-3	0.3125	3.125
B-6	0.3125	6.375
B-12	6.375	12.5
C	6.375	32

Table 3. JESD204C 32 Gbps Interface Device Class Features

Class	Relative Power	Transmitter FFE (Minimum)	Receiver CTLE (Minimum)	Receiver DFE Taps (Minimum)
C-S	Low	9.5 dB	6 dB	0
C-M	Medium	9.5 dB	9 dB	3
C-R	High	9.5 dB	12 dB	14

JESD204C also introduces the concept of the JESD204 channel operating margin (JCOM), which is used to confirm compliance to the Class C PHY layer standard. This calculation of the operating margin supplements the eye masks that apply the Class B PHY layer implementations that are described in this and in previous revisions of the standard.

Clocking and Synchronization

JESD204C will retain the use of SYSREF and device clock as defined in JESD204B. However, when using either of the 64-bit encoding schemes, instead of aligning the LMFC, the SYSREF is used to align the local extended multiblock counter (LEMC) to provide a mechanism for deterministic latency and multichip synchronization.

The synchronization process for the 64-bit encoding schemes is completely different than the one used in JESD204B. The SYNC signal has been eliminated and sync initialization and error reporting will now be handled in the application layer software. Therefore, there is no code-group sync (CGS) or initial lane alignment sequence (ILAS). Sync header sync, extended multiblock sync, and extended multiblock alignment are new sync-related terms used to describe the synchronization process. Each of these synchronization phases are achieved using a 32-bit sync word. This is discussed in detail in part two of this series.

Note that for 8b/10b encoding, both the SYNC pins and the ILAS are retained.

Deterministic Latency and Multichip Synchronization

As implied above, the mechanism for achieving deterministic latency and multichip synchronization remains mostly intact from JESD204B. When using one of the 64-bit encoding schemes, there is no Subclass 2 option. Instead, only Subclass 1 operation is supported and the SYSREF signal is used to align the LEMC across all devices in the JESD204 subsystem.

Forward Error Correction

To meet the goal of providing a more robust link at higher lane rates, an FEC option has been included in JESD204C. This algorithm is based on fire codes and may be particularly useful for instrumentation applications. This is an optional feature that is only available when using one of the 64-bit encoding schemes.

Fire codes are cyclic codes that correct single-burst errors. The advantage of cyclic codes is that their codewords can be represented as polynomials—as opposed to vectors—over a finite field. Fire codes use a syndrome that can be split into two components for faster decoding.

More Information

Coming soon, in part two of the JESD204C primer series, we will dive a little deeper into the key elements of the JESD204C standard that enable the problem-solving technology we described in the opening paragraphs. Specifically, the bandwidth efficiency improvements enabled by the 64b/66b encoding scheme is given a closer look as is the bandwidth-increasing 32 Gbps physical layer specification. More depth is also provided on the new synchronization process as well as the optional forward error correction aspect of the standard that improves link robustness.

For more information on JESD204 and its implementation in Analog Devices products, please visit ADI's JESD204 serial interface page.

Del Jones [del.jones@analog.com] 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.

Del Jones

How Sensor Performance Enables Condition-Based Monitoring Solutions

By **Pete Sopcik** and **Dara O'Sullivan**

Advancements in semiconductor technologies and capabilities are enabling new opportunities to sense, measure, interpret, and analyze data in industrial applications and, in particular, condition-based monitoring (CbM) solutions. Next-generation sensors based on MEMS technology, combined with advanced algorithms for diagnostic and prognostic applications, expand opportunities to measure a variety of machines and improve the ability to effectively monitor equipment, improve uptime, enhance process quality, and increase throughput.

To enable these new capabilities and capture the benefits of condition-based monitoring, new solutions must be accurate, reliable, and robust so that real-time monitoring can expand beyond basic detection of potential equipment faults to deliver insightful and actionable information. Performance of next-generation technologies combined with system-level insights enable a deeper understanding into the application and requirements necessary to solve these challenges.

Vibration, one of the key components of machine diagnostics, has been reliably used to monitor the most critical equipment across a wide range of industrial applications. A significant amount of literature exists to support the various diagnostic and predictive capabilities required to enable advanced vibration monitoring solutions. Less well covered is the relationship between vibration sensor performance parameters, such as bandwidth and noise density, and end application fault diagnostic capabilities. This article addresses the major machine fault types in industrial automation applications and identifies the key vibration sensor performance parameters that are relevant to the specific faults.

Several common fault types and their characteristics are highlighted below to provide insights into some of the key system requirements that must be considered when developing a condition-based monitoring solution. These include—but are not limited to—imbalance, misalignment, gear faults, and rolling bearing defects.

Imbalance

What is imbalance and what causes it?

Imbalance is an unequal distribution of mass that causes the load to shift the center of mass away from the center of rotation. System imbalances can be attributed to improper installations such as coupling eccentricity, system design errors, component faults, and even accumulation of debris or other contaminates. As an example, the cooling fans built into most induction motors can become unbalanced due to an uneven accumulation of dust and grease, or due to broken fan blades.

Why is an unbalanced system a concern?

Unbalanced systems create excess vibrations that mechanically couple to other components within the system such as bearings, couplings, and loads—potentially accelerating the deterioration of components that are in good operating condition.

How to detect and diagnose imbalance

Increases in overall system vibration can point to a potential fault created by an unbalanced system, but diagnosis of the root cause of the increased vibration is performed through analysis in the frequency domain. Unbalanced systems produce a signal at the rotational rate of the system—typically referred to as 1×—with a magnitude that is proportional to the square of the rotational rate, $F = m \times w^2$. The 1× component is typically always present in the frequency domain, so identification of an unbalanced system is done by measuring the magnitude of the 1× and the harmonics. If the magnitude of the 1× is higher than the baseline measurement and the harmonics are much less than the 1×, then an unbalanced system is likely. Both horizontally and vertically phase-shifted vibration components are also likely in an unbalanced system.[1]

What system specifications must be considered when diagnosing an unbalanced system?

Low noise is required to reduce the sensor influence and enable detection of small signals created by an unbalanced system. This is important for the sensor, signal conditioning, and acquisition platform.

Sufficient resolution of the acquisition system to extract the signal (especially the baseline signal) is required to detect these small imbalances.

Bandwidth is necessary to capture sufficient information beyond just the rotational rates to improve the accuracy and confidence of a diagnosis. The 1× harmonic can be influenced by other system faults, such as misalignment or mechanical looseness, so analysis of the harmonics of the rotation rate, or 1× frequency, can help differentiate from system noise and other potential faults.[1] For slower rotating machines, fundamental rotation rates can be well below 10 rpm, meaning the low frequency response of the sensor is critical for capturing the fundamental rotation rates. Analog Devices' MEMS sensor technology enables detection of signals down to dc and provides the ability to measure slower rotation equipment, while also enabling measurement of wide bandwidths for higher frequency content typically associated with bearing and gearbox defects.

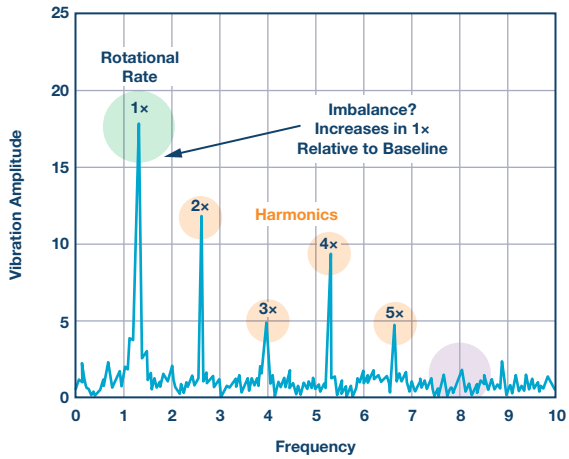

Figure 1. Potential for an unbalanced system exists based on increased amplitude at the rotational rate, or 1× frequency.

Misalignment

What is misalignment and what causes it?

System misalignments, as the name suggests, occur when two rotating shafts are not aligned. Figure 2 shows an ideal system where alignment is achieved starting with the motor, then the shaft, the coupling, and all the way to the load (which, in this case, is a pump).

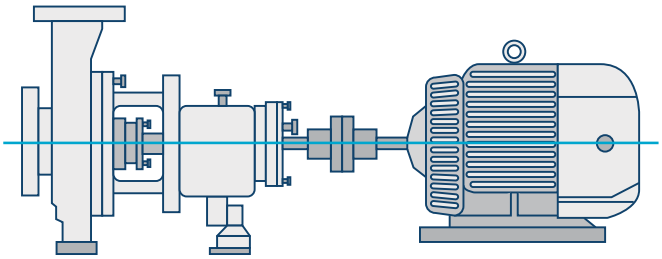

Figure 2. An ideally aligned system.

Misalignments can occur in the parallel direction as well as in the angular direction and can also be a combination of both (see Figure 3). Parallel misalignment occurs when the two shafts are displaced in the horizontal or vertical directions. Angular misalignment occurs when one of the shafts is at an angle relative to the other.[2]

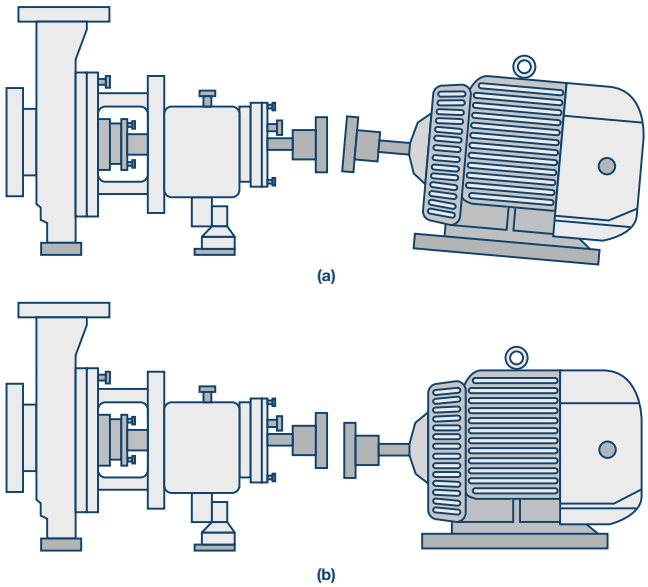

Figure 3. Examples of different misalignments include (a) angular, (b) parallel, or a combination of both.

Why is misalignment a concern?

Misalignment errors can impact the greater system by forcing components to operate under higher stresses, or loads, than what the components were originally designed to handle and can ultimately cause premature failures.

How to detect and diagnose misalignments

Misalignment errors typically show up as the second harmonic of the rotational rate of the system, referred to as 2×. The 2× component is not always present in the frequency response, but when it is, the relationship of the magnitude to the 1× can be used to determine whether a misalignment is present. Increased misalignments can excite harmonics out to 10× depending on the type of misalignment, the location at which it is measured, and the directional information.[1] Figure 4 highlights the signatures associated with potential misalignment failures.

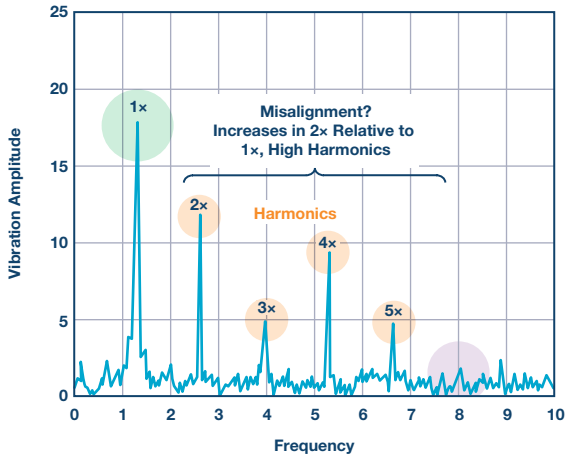

Figure 4. An increasing 2× harmonic, combined with increasing higher order harmonics, indicates a potential misalignment.

What system specifications must be considered when diagnosing a misaligned system?

Low noise and sufficient resolution are required to detect small misalignments. Machine types, system and process requirements, and rotational rates dictate the allowable misalignment tolerances.

Bandwidth is necessary to capture sufficient frequency range and improve the accuracy and confidence of a diagnosis. The 1× harmonic can be influenced by other system faults, such as misalignment, so analysis of the harmonics of the 1× frequency can help differentiate from other system faults. This is especially true for higher rotational speed machines. As an example, machines operating above 10,000 rpm, such as machine tools, will typically require quality information beyond 2 kHz in order to accurately detect imbalance with high confidence.

Multidirectional information also improves the accuracy of the diagnosis and provides insight into the type of misalignment error and the direction of the misalignment.

The phase of the system, combined with directional vibration information, further improves the diagnostics of a misalignment error. Measuring the vibration at different points on the machine and determining the difference in the phase measurements or across the system provides insights into whether the misalignment is either an angular, parallel, or combination of the two misalignment types.[1]

Rolling Element Bearing Defects

What are rolling element bearing defects and what causes them?

Rolling element bearing defects are typically artifacts of mechanically induced stresses or lubrication issues that create small cracks or defects

within the mechanical components of the bearing, leading to increased vibration. Figure 5 provides some examples of rolling element bearings and depicts a couple of the defects that can occur.

Figure 5. Examples of (top) rolling element bearings and (bottom) lubrication and discharge current defects.

Why are rolling element bearing failures a concern?

Rolling element bearings are found in almost all types of rotating machinery, ranging from large turbines to slower rotating motors all the way from relatively simple pumps and fans to high speed CNC spindles. Bearing defects can be a sign of contaminated lubrication (Figure 5), improper installations, high frequency discharge currents (Figure 5), or increased loading from the system. Failures can cause catastrophic system damage and have significant impacts on other system components.

How are rolling element bearing faults detected and diagnosed?

There are a number of techniques used to diagnose bearing faults and because of the physics behind bearing design, each bearing's defect frequencies can be computed based on the bearing geometries, the speed of rotation, and the defect type, which aids in diagnosing faults. Bearing defect frequencies are listed in Figure 6.

Analysis of the vibration data from a particular machine or system often relies on a combination of both time and frequency domain analysis. Time domain analysis is useful for detecting trends in the overall increase of system vibration levels. However, very little diagnostic information is contained in this analysis. Frequency domain analysis improves diagnostic insights, but identifying the fault frequencies can be complex due to influences from other system vibrations.

For early diagnosis of bearing defects, harmonics of the defect frequencies are used to identify the early stage, or incipient, faults so that they can be monitored and maintained before a catastrophic failure. In order to detect, diagnose, and understand the system implications of a bearing fault, techniques such as envelope detection, shown in Figure 7, combined with spectral analysis in the frequency domain typically provide more insightful information.

What system specifications must be considered when diagnosing a rolling element bearing fault?

Low noise and sufficient resolution are critical to the detection of early stage bearing defects. Typically, these defect signatures are low in amplitude during the onset of a defect. Mechanical slip, inherent to bearings due to design tolerances, further reduces the magnitude of the vibrations by spreading amplitude information across multiple bins in the frequency response of a bearing, thus requiring low noise to detect the signals earlier.[2]

Bandwidth is critical for early detection of bearing defects. An impulse containing high frequency content is created each time the defect is struck during a revolution (see Figure 7). Harmonics of the bearing defect frequencies, not the rotational rate, are monitored for these early stage faults. Because of the relationship of the bearing defect frequencies to rotation rates, these early signatures can occur in the several kilohertz range and extend well beyond the 10 kHz to 20 kHz range.[2] Even for lower speed equipment, the inherent nature of bearing defects requires wider bandwidths for early detection to avoid influences from system resonances and system noise that influence the lower frequency bands.[3]

Dynamic range is also important for bearing defect monitoring as system loads and defects can impact the vibrations experienced by the system. Increased loads lead to increased forces acting on the bearing and the defect. Bearing defects also create impulses that excite structural resonances, amplifying the vibrations experienced by the system and the sensor.[2] As machines ramp up and down in speeds during stop/start conditions or normal operation, the changing speeds create potential opportunities for system resonances to become excited, leading to higher amplitude vibrations.[4] Saturation of the sensor can result in missing information, misdiagnosis, and—in the case of certain technologies—damage to the sensor elements.

Bearing Fundamental Defect Frequencies

FTF (Fundamental Train Frequency)	Frequency of the cage	$FTF = \frac{F}{2} \times \left(1 - \frac{B}{P} \times cos\Theta\right)$
BSF (Ball Spin Frequency)	Circular frequency of each rolling element as it spins	$BSF = \frac{P}{2B} \times F \times \left[1 - \left(\frac{B}{P} \times cos\Theta\right)^2\right]$
BPFO (Ball Pass Frequency Outer Race)	Frequency generated when rolling elements roll across a defect in the outer ring	$BPFO = \frac{N}{2} \times F \times \left(1 - \frac{B}{P} \times cos\Theta\right)$
BPFI (Ball Pass Frequency Inner Race)	Frequency generated when rolling elements roll across a defect in the inner ring	$BPFI = \frac{N}{2} \times F \times \left(1 + \frac{B}{P} \times cos\Theta\right)$

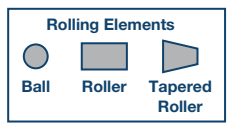

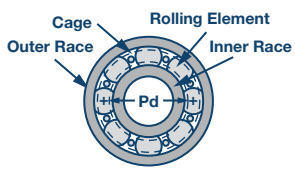

Figure 6. Bearing defect frequencies are dependent on the bearing types, geometries, and rotation rates.

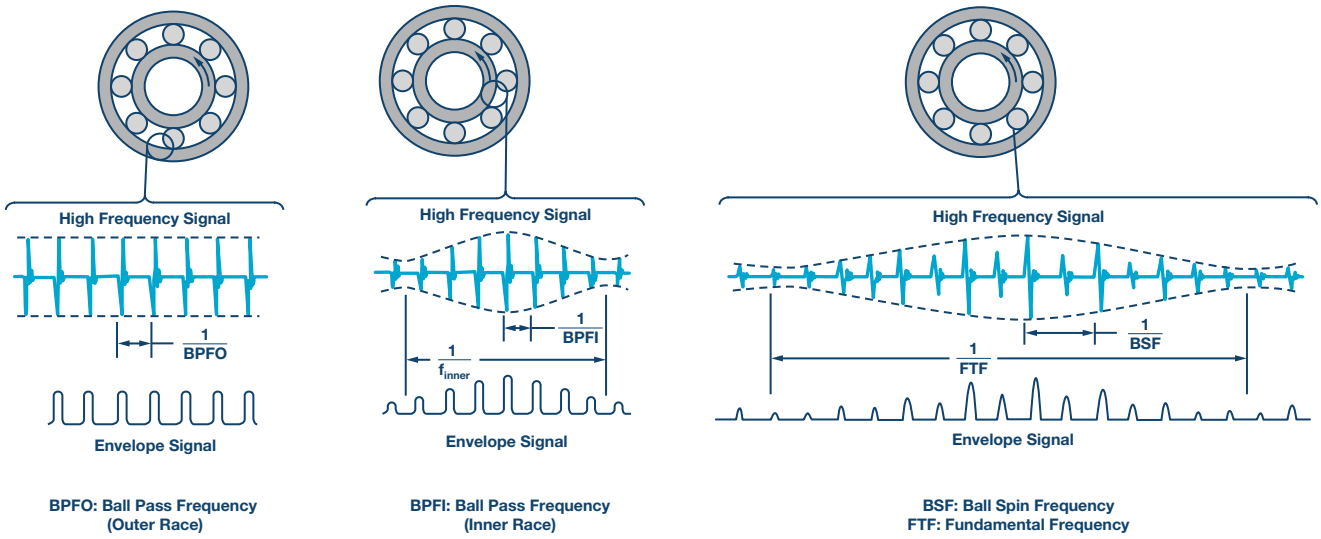

Figure 7. Techniques such as envelope detection can extract early bearing defect signatures from wide bandwidth vibration data.

Gear Defects

What are gear defects and what causes them?

Gear faults typically occur in the teeth of a gear mechanism due to fatigue, spalling, or pitting. These can be manifested as cracks in the gear root or removal of metal from the tooth surface. They can be caused by wear, excessive loads, poor lubrication, backlash, and occasionally improper installation or manufacturing defects.[5]

Why are gear faults a concern?

Gears are the main elements of power transmission in many industrial applications and are subjected to significant stresses and loading. Their health is critical to the proper operation of the entire mechanical system. A well-known example of this in the renewables field is the fact that the greatest contributor to wind turbine downtime (and consequent revenue erosion) is the failure of the multistage gearbox in the main powertrain.[5] Similar considerations apply in industrial applications.

How are gear faults detected and diagnosed?

Gear faults are tricky to detect due to the difficulty in installation of vibration sensors close to the fault and the presence of significant background noise due to multiple mechanical excitations within the system. This is especially true in more complex gearbox systems, in which there can be multiple rotational frequencies, gear ratios, and meshing frequencies.[6] Consequently, multiple and complementary approaches can be taken in the detection of gear faults, including acoustic emissions analysis, current signature analysis, and oil debris analysis.

In terms of vibration analysis, the gearbox casing is the typical mounting location for an accelerometer, with the dominant vibration mode being in the axial direction.[7] Healthy gears produce a vibration signature at a frequency known as the gear mesh frequency. This is equal to the product of the shaft frequency and the number of gear teeth. There typically also exist some modulation sidebands related to manufacturing and assembly tolerances. This is illustrated for a healthy gear in Figure 8. When a localized fault such as a tooth crack occurs, the vibration signal in each revolution will include the mechanical response of the system to a short duration impact at a relatively low energy level. This is typically a low amplitude, broadband signal that is generally considered to be non-periodic and non-stationary.[7,8]

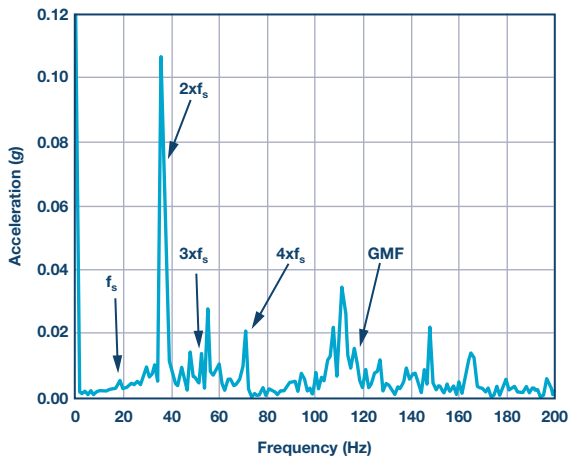

Figure 8. Frequency spectrum of a healthy gear with crank shaft speed at ~1000 rpm, gear speed at ~290 rpm, and gear teeth = 24.

As a result of these particular characteristics, standard frequency domain techniques on their own are not regarded as suitable for accurate identification of gear faults. Spectral analysis may be unable to detect early stage gear failures as the impact energy is contained in sideband modulation, which can also contain energy from other gear pairs and mechanical components. Time domain techniques such as time-synchronous averaging or mixed-domain approaches such as wavelet analysis and envelope demodulation are generally more appropriate.[9]

What system specifications must be considered when diagnosing a gear fault?

Wide bandwidth is generally very critical in gear fault detection, since the number of gear teeth acts as a multiplier in the frequency domain. Even for relatively low speed systems, the required detection frequency range is quickly pushed up in to the multiple kHz region. Moreover, localized faults further extend the bandwidth requirement.

Resolution and low noise are extremely critical for several reasons. The difficulty of mounting vibration sensors in close proximity to specific fault zones means that there is potentially higher attenuation of the vibration signal by the mechanical system, making it vital to be able to detect low energy signals. Furthermore, since the signals are not static periodic signals,

standard FFT techniques to extract low amplitude signals from a high noise floor cannot be depended on—the noise floor of the sensor itself must be low. This is particularly true in a gearbox environment in which there is a mixing of multiple vibration signatures from different elements of the gearbox. Added to these considerations is the importance of early detection not just for asset protection reasons, but for signal conditioning reasons. It has been shown that vibration severity can be higher in the case of a one-tooth breakage fault, as opposed to a fault with two-or-more-tooth breakage, implying that detection may be relatively easier at the early stages.

Summary

While common, imbalance, misalignments, rolling element bearing defects, and gear tooth faults are just a few of the many fault types that can be detected and diagnosed with high performance vibration sensors. Higher sensor performance, combined with the appropriate system-level considerations, enable next-generation condition-based monitoring solutions that will deliver deeper levels of insight into the mechanical operation of a wide range of industrial equipment and applications. These solutions will transform how maintenance is performed and how machines operate, ultimately reducing downtimes, improving efficiencies, and delivering new capabilities to next-generation equipment.

Table 1. Requirements on Each Sensor Parameter

Fault Type	Bandwidth	Noise Density	Dynamic Range	Resolution
Imbalance	Low	Medium	High	Medium
Misalignment	Medium	Low/medium	High	Medium
Bearing	High/very high	Low	Medium	High
Gears	Very high	Low	Low	High

For Table 1, a low bandwidth is considered <1 kHz, a medium bandwidth is between 1 kHz to 5 kHz, and a high bandwidth is considered >5 kHz. A low noise density is considered >1 m$g/\sqrt{Hz}$, a medium noise density is between 100 µ$g/\sqrt{Hz}$ to 1 m$g/\sqrt{Hz}$, and a high noise density is considered <100 µ$g/\sqrt{Hz}$. A low dynamic range is considered <5 *g*, a medium dynamic range is between 5 *g* to 20 *g*, and a high dynamic range is considered >20 *g*.

References

[1] Jason Mais. "Spectrum Analysis: The Key Features of Analyzing Spectra." SKF USA, Inc., 2002.

[2] Robert Bond Randall. *Vibration-based Condition Monitoring: Industrial, Aerospace, and Automotive Applications*. John Wiley & Sons, Ltd., December 2010.

[3] Scott Morris. "SKF Pulp and Paper Practices." *SKF Global Pulp & Paper Segment*, No. 19, 2016.

[4] Chris D. Powell, Erik Swanson, and Sorin Weissman. "A Practical Review of Rotating Machinery Critical Speeds and Modes." *Sound and Vibration*, May 2005.

[5] Shahin Hedayati Kia, Humberto Henao, and Gérard-André Capolino. "Trends in Gear Fault Detection Using Electrical Signature Analysis in Induction Machine-Based Systems." 2015 IEEE Workshop on Electrical Machines Design, Control and Diagnosis (WEMDCD), Torino, pp. 297-303, 2015.

[6] Alexander Bliznyuk, Ido Dadon, Renata Klein, and Jacob Bortman. "Gear Diagnostics—Fault Type Characteristics." PHM Society, 2014.

[7] Giorgio Dalpiaz, Alessandro Rivola, and Riccardo Rubini. "Effectiveness and Sensitivity of Vibration Processing Techniques for Local Fault Detection in Gears." *Mechanical Systems and Signal Processing*, Vol. 14, No. 3, 2000.

[8] Wenyi Wang. "Early Detection of Gear Tooth Cracking Using the Resonance Demodulation Technique." *Mechanical Systems and Signal Processing*, Vol. 15, No. 5, 2001.

[9] Kiran Vernekar, Hemantha Kumar, and K. V. Gangadharan. "Gear Fault Detection Using Vibration Analysis and Continuous Wavelet Transform." *Procedia Materials Science*, Vol. 5, 2014.

Pete Sopcik [pete.sopcik@analog.com] is the marketing manager for condition-based monitoring at Analog Devices. He is responsible for the strategy supporting the sensors, signal chains, and solutions required to enable condition monitoring in industrial applications. Prior to this role, Pete spent 11 years working in several different areas within Analog Devices, ranging from the design and development of integrated high speed converters and sensor modules to the project management of system level solutions, including MEMS-based inertial measurement units. He has a B.S.E.E. from the University of Florida and an M.B.A. from Wake Forest University.

Pete Sopcik

Dara O'Sullivan [dara.osullivan@analog.com] is a system applications manager with the connected motion and robotics team within the automation and energy business unit at Analog Devices. His area of expertise is power conversion, control, and monitoring in industrial motion control applications. He received his B.E, M.Eng.Sc., and Ph.D. degrees from University College Cork, Ireland, and has worked in industrial and renewable energy applications in a range of research, consultancy, and industry positions since 2001.

Dara O'Sullivan

Bipolar, Single Output, and Adjustable Power Supplies Based on the Common Buck Converter

By **Victor Khasiev**

Introduction

A bench-top power supply (PS) tends to have an even number of terminals (ignoring the chassis port)—with one positive terminal and one negative terminal. Using a bench-top supply to produce a positive polarity output is easy: set the minus output to GND and the positive output voltage at the plus output. It is just as easy to produce a negative supply by reversing the setup. But what about producing a bipolar supply, where positive and negative voltages are both available to the load? This is relatively easy, too—just connect the positive terminal of one lab channel to the negative of another channel and call that GND. The other two terminals, minus and plus, are the positive and negative supplies, respectively. The result is a three-terminal bipolar power supply with available GND, positive, and negative voltage levels. Because three terminals are used, there must be some switch between positive and negative supplies downstream of the power supply.

What if an application calls for the same power supply terminal to be positive or negative—a setup where only two terminals are provided to the load? This is not a purely academic question. There are applications in automotive and industrial environments that require bipolar, adjustable two terminal power supplies. For instance, two terminal bipolar power supplies are used in applications ranging from exotic window tinting to test and measurement equipment.

As noted earlier, a traditional bipolar PS produces two outputs using three output terminals: positive, negative, and GND. In contrast, a single output power supply should be equipped with only two output terminals: one GND and another that can be positive or negative. In such applications, the output voltage can be regulated relative to the GND by a single control signal, in the full range from the minimum negative to maximum positive.

There are controllers that are specifically designed to implement the bipolar supply function, such as the LT8714, a bipolar output synchronous controller. Nevertheless, for many automotive and industrial manufacturers, testing and qualifying a specialized IC requires some investment in time and money. By contrast, many manufacturers already have prequalified step-down (buck) converters and controllers, as they are used in countless automotive and industrial applications. This article shows how to use a buck converter to produce a bipolar PS when a dedicated bipolar supply IC is not an option.

Circuit Description and Functionality

Figure 1 shows a buck converter-based solution for a bipolar (two-quadrant) adjustable power supply. The input voltage range is 12 V to 15 V; the output is any voltage in the ±10 V range, adjusted by the control block, that supports loads up to 6 A. The dual output step-down controller IC is the central component to this design. One output, connected per buck-boost topology, generates a stable –12 V (that is, the –12 V negative rail in Figure 1, with its power train comprising L2, Q2, Q3, and output filter C_{02}).

The –12 V rail serves as ground for the second channel with the controller's ground pins connected to the –12 V rail as well. Overall, this is a step-down buck converter, where the input voltage is the difference between –12 V and V_{IN}. The output is adjustable and can be either positive or negative relative to GND. Note the output is always positive relative to the –12 V rail and includes a power train comprising L1, Q1, Q4, and C_{01}. The feedback resistor divider R_B–R_A sets the maximum output voltage. The value of this divider is adjusted by the output voltage control circuit, which can regulate the output down to the minimum output voltage (negative output) by injecting current into R_A. The application start-up characteristics are set by the termination of the RUN and TRACK/SS pins.

Both outputs function in forced continuous conduction mode. In the output control circuit, the 0 µA to 200 µA current source, I_{CTRL}, is connected to the negative rail as tested in the lab, but it can be referenced to the GND as well. The low-pass filter R_{F1}–C_F reduces fast output transients. To reduce the cost and size of the converter, output filters are formed using relatively inexpensive polarized capacitors. The optional diodes D1 and D2 prevent developing the reverse voltage across these capacitors, especially at start-up. There is no need for the diodes if only ceramic capacitors are used.

Converter Testing and Evaluation

This solution was tested and evaluated based on the LTC3892 and evaluation kits DC1998A and DC2493A. The converter performed well in a number of tests, including line and load regulation, transient response, and output short. Figure 2 shows startup to a 6 A load, with a +10 V output. The linearity of the function between the control current and output voltage is shown in Figure 3. As control current increases from 0 µA to 200 µA, the output voltage decreases from +10 V to –10 V. Figure 4 shows the efficiency curves.

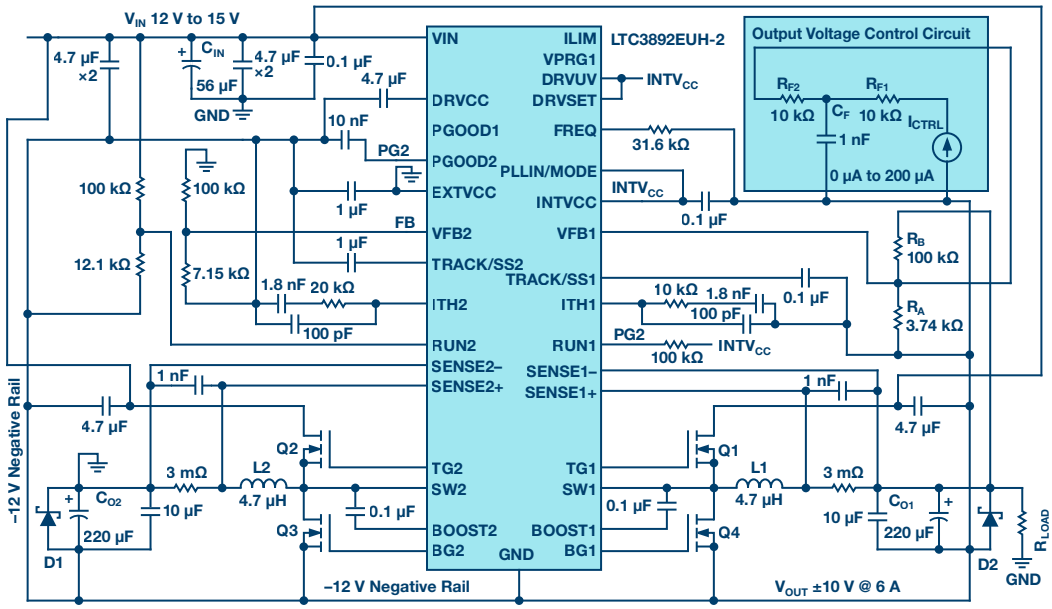

Figure 1. Electrical schematic of the two terminal, bipolar, adjustable power supply.

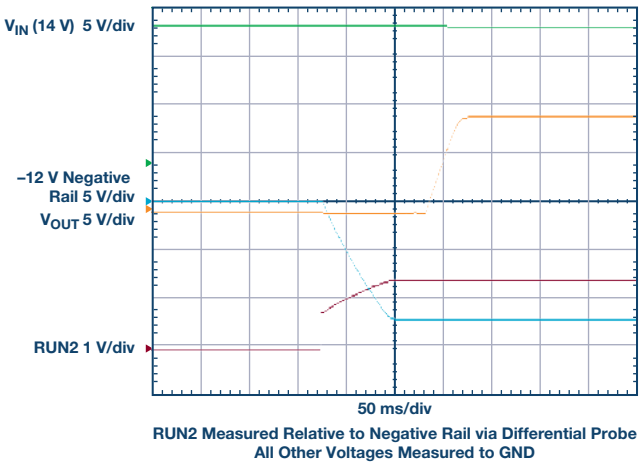

Figure 2. Start-up waveforms into resistive load.

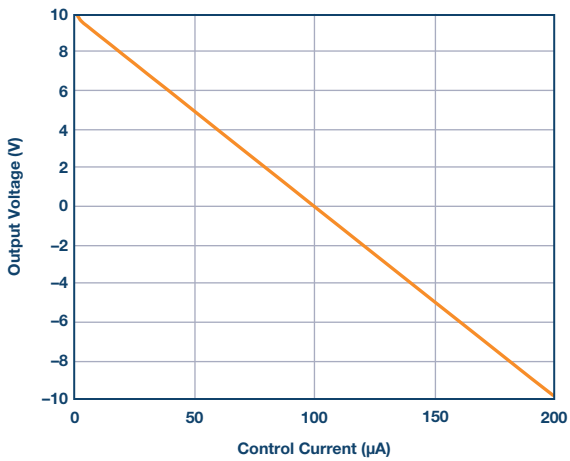

Figure 3. V_{OUT} as a function of control current I_{CTRL}. As I_{CTRL} increases from 0 A to 200 µA, the output voltage drops from +10 V to –10 V.

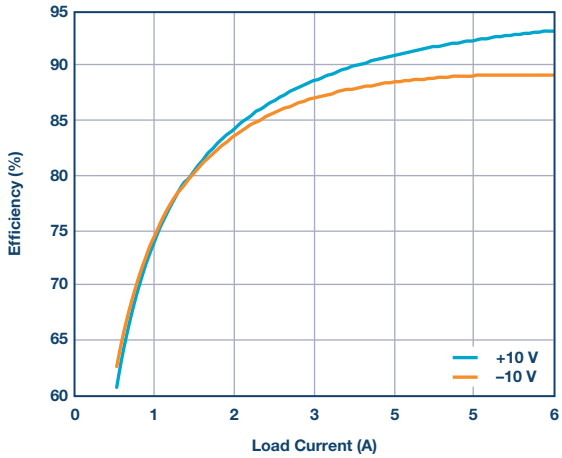

Figure 4. Efficiency curves for positive and negative output.

An LTspice® model of the bipolar, two terminal power supply was developed to simplify adoption of this approach, allowing designers to analyze and simulate the circuit described above, introduce changes, view waveforms, and study component stress.

Essential Formulas and Expressions Describing this Topology

This approach is based on the negative rail, V_{NEG}, generated by the buck-boost section of the design.

$$V_{NEG} = V_{OUT} + V_{OUT} \times K_m \quad (1)$$

Where V_{OUT} is the absolute value of maximum output voltage and K_m is a coefficient ranging from 0.1 to 0.3. K_m limits the minimum duty cycle of the step-down converter. V_{NEG} also sets the minimum value of V_{IN}:

$$V_{IN} \geq |V_{NEG}|$$
$$V_{BUCK} = |V_{NEG}| + V_{IN} \quad (2)$$

Where V_{BUCK} is the input voltage for the step-down section and thus presents the maximum voltage stress on the converter's semiconductors:

$$V_{BUCK(MAX)} = |V_{NEG}| + V_{OUT}$$
$$V_{BUCK(MIN)} = |V_{NEG}| - V_{OUT} \quad (3)$$

$V_{BUCK(MAX)}$ and $V_{BUCK(MIN)}$ are the maximum and minimum voltages of the step-down section of this topology, respectively. The maximum and minimum duty cycles and inductor current of the step-down section can be described by the following expressions, where I_{OUT} is output current:

$$D_{BUCK(MAX)} = V_{BUCK(MAX)}/V_{BUCK}$$
$$D_{BUCK(MIN)} = V_{BUCK(MIN)}/V_{BUCK} \quad (4)$$
$$I_{L(BUCK)} = I_{OUT} + \Delta I_1$$

The duty cycle of the buck-boost section of the PS:

$$D_{BB} = |V_{NEG}|/(V_{IN} + |V_{NEG}|) \quad (5)$$

The input power of the step-down section and, correspondingly, output power of the buck-boost:

$$P_{OUT(BB)} = (V_{OUT} \times I_{OUT})/\eta \quad (6)$$

Output current of the buck-boost section and its inductor current

$$I_{OUT(BB)} = P_{OUT(BB)}/|V_{NEG}|$$
$$I_{L(BB)} = I_{OUT(BB)}/(1 - D_{BB}) + \Delta I_2 \quad (7)$$

The converter power and input current.

$$P_{BB} = P_{OUT(BB)}/\eta$$
$$I_{BB} = P_{BB}/V_{IN} \quad (8)$$

The output voltage changes are executed by injecting current into the feedback resistor divider of the step-down section. Setting up the output voltage control is illustrated in the output voltage control circuit section of Figure 1.

If R_B is given, then

$$R_A = V_{FB} \times RB/(V_{BUCK(MAX)} - V_{FB}) \quad (9)$$

where V_{FB} is the feedback pin voltage.

When the current source I_{CTRL} injects zero current into R_A, the output voltage of the buck converter is the maximum positive value ($V_{BUCK(MAX)}$) relative to the negative rail and maximum output voltage (+ V_{OUT}) relative to GND. To produce a negative output voltage to the load (relative to GND), the output voltage is reduced to its minimum value, $V_{BUCK(MIN)}$, relative to the negative output voltage (–V_{OUT}), by injecting ΔI into resistor R_A of the buck's voltage divider.

$$\Delta I = I_{FB} - I_{RAL}$$
$$I_{FB} = V_{FB}/R_A \quad (10)$$
$$I_{RAL} = (V_{BUCK(MIN)} - V_{FB})/R_B$$

Numerical Example

By using the previous equations, we can calculate voltage stress, current through the power train components, and the parameters of the control circuit for the bipolar power supply. For instance, the following calculations are for a supply generating ±10 V at 6 A from a 14 V input voltage.

If K_m is 0.2, then V_{NEG} = –12 V. Verifying conditions of minimum input voltage $V_{IN} \geq |V_{NEG}|$. The voltage stress on the semiconductor's V_{BUCK} is 26 V.

The maximum voltage of the step-down section is $V_{BUCK(MAX)}$ = 22 V, relative to negative rail, setting the output voltage +10 V relative to GND. The minimum voltage, $V_{BUCK(MIN)}$ = 2 V, corresponds to the output voltage of –10 V relative to GND. These maximum and minimum voltages correspond to the maximum and minimum duty cycles, $D_{BUCK(MAX)}$ = 0.846, $D_{BUCK(MIN)}$ = 0.077, and D_{BB} = 0.462.

Power can be calculated by assuming an efficiency of 90%, producing $P_{OUT(BB)}$ = 66.67 W, $I_{OUT(BB)}$ = 5.56 A, $I_{L(BB)}$ = 10.37 A, and P_{BB} = 74.074 W.

For an output voltage of +10 V (as per Figure 1), the control circuit current, ΔI, is 0 µA, whereas for an output voltage of –10 V, ΔI = 200 µA.

Conclusion

This article presents a design for bipolar, two terminal power supplies. The approach discussed here is based on step-down converter topology, which is a staple of modern power electronics, and thus available in a variety of forms, from simple controllers with external components to complete modules. Employment of step-down topology gives the designer flexibility and an option to use prequalified parts, which saves time and cost.

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

Victor Khasiev

Rarely Asked Questions—Issue 166 How to Convert Light Intensity Into an Electrical Quantity

By **Thomas Brand**

Question:

How could I measure the light intensity of different light sources?

Answer:

An LED that is red, green, and blue.

Determination of the light intensity can be crucial, for example, if you want to design the lighting of a room, or in preparing for a photo shoot. In the era of the Internet of Things (IoT), however, light intensity also plays an important role in so-called smart agriculture. Here, one key task is to monitor and control important plant parameters that contribute to maximizing plant growth and accelerating photosynthesis. Light is thereby one of the most important factors in smart agriculture. Most plants usually absorb light in the wavelengths of the red, orange, blue, and violet regions of the visible spectrum. As a rule, light in the wavelengths of the green and yellow regions of the spectrum is reflected and contributes only slightly to growth. By controlling parts of the spectrum and the intensity of light exposure throughout various life stages, growth can be maximized, and the yield ultimately increased.

A corresponding circuit design for measuring the light intensity over the visible spectrum, in which plants are photosynthetically active, is shown in Figure 1. Here, three differently colored photodiodes (green, red, and blue) are used, which respond to different wavelengths. The light intensity measured via the photodiodes can now be used to control the light source according to the requirements of the respective plants.

The circuit shown here is made up of three precise current-to-voltage converter stages (transimpedance amplifiers), one for each of the colors green, red, and blue. They are connected to the differential inputs of a Σ-Δ analog-to-digital converter (ADC), which, for example, provides the measured values as digital data to a microcontroller for further processing.

Conversion of Light Intensity into Current

Depending on the light intensity, more or less current flows through the photodiodes. The relationship between current and light intensity is approximately linear, which is illustrated in Figure 2. It shows the characteristic curves of the output current as a function of light intensity for a red (CLS15-22C/L213R/TR8), a green (CLS15-22C/L213G/TR8), and a blue photodiode (CLS15-22C/L213B/TR8).

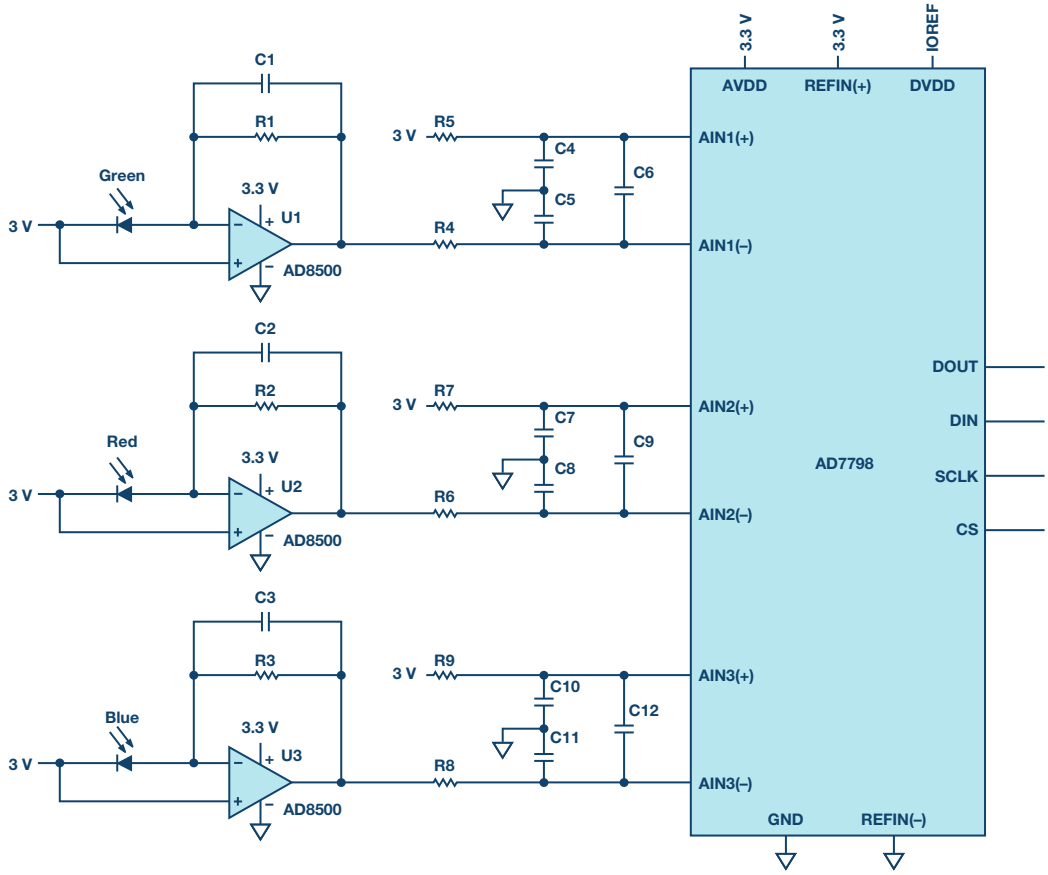

Figure 1. Circuit design for measurement of light intensity.

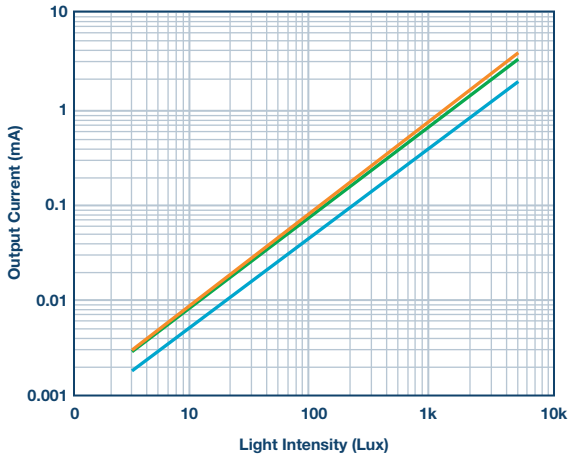

Figure 2. Characteristic curves of current to light intensity for red, green, and blue photodiodes.

However, the relative sensitivities of the red, green, and blue diodes are different, so the gain of each stage must be determined separately through the feedback resistance R_{FB}. For this, the short-circuit current (I_{SC}) of each diode has to be taken from the data sheet and subsequently the sensitivity, S (pA/lux), at the operating points determined from it. R_{FB} is then calculated as follows:

$$R_{FB} = \frac{V_{FS,P\text{-}P}}{S \times INT_{MAX}}$$

$V_{FS,P\text{-}P}$ represents the desired full output voltage range (full-scale, peak-to-peak) and INT_{MAX} the maximum light intensity, which is 120,000 lux for direct sunlight.

Current-to-Voltage Conversion

For a high quality current-to-voltage conversion, the minimal bias current of the operational amplifier is desirable because the output current of the photodiode is in the picoampere range and thus can cause considerable errors. A low offset voltage should also be present. With a bias current of typically 1 pA and a maximum offset voltage of 1 mV, the AD8500 from Analog Devices is a good choice for these applications.

Analog-to-Digital Conversion

For further processing of the measured values, the photodiode current that was first converted into a voltage has to be provided to the microcontroller as a digital value. For this purpose, ADCs with multiple differential inputs can be used, such as the 16-bit ADC AD7798. Thus, the output code for the measured voltage is as follows:

$$Code = \frac{(2^N \times A_{IN} \times GAIN)}{V_{REF}}$$

where

A_{IN} = input voltage,

N = number of bits,

GAIN = gain factor of the internal amplifier,

V_{REF} = external reference voltage.

For further noise reduction, a common-mode and a differential filter are used on each of the differential inputs of the ADC.

All of the depicted components are extremely power-saving, making the circuit ideal for battery-operated portable field applications

Conclusion

Error sources such as bias currents and offset voltages of the components must be considered. Also, unfavorable converter stage amplification factors can impact the quality and thus the result of the circuit. With the circuit shown in Figure 1, light intensity can be converted into an electrical value for further data processing in a relatively simple way.

Thomas Brand

Thomas Brand [thomas.brand@analog.com] began his career at Analog Devices in Munich in 2015 as part of his master's thesis. After graduating, he was part of a trainee program at Analog Devices. In 2017, he became a field applications engineer. Thomas supports large industrial customers in central Europe and specializes in the field of Industrial Ethernet. He studied electrical engineering at the University of Cooperative Education in Mosbach before completing his postgraduate studies in international sales with a master's degree at the University of Applied Sciences in Constance.

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*

AnalogDialogue

EngineerZone® Support Community

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

EngineerZone®
SUPPORT COMMUNITY

Facebook
facebook.com/AnalogDevicesInc

@ADI_News
twitter.com/ADI_News

LinkedIn
linkedin.com/company/analog-devices

Sina Weibo (China)
weibo.com/analogdevices

WeChat (China)

Analog Devices, Inc.
Worldwide Headquarters

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

Analog Devices, Inc.
Europe Headquarters

Analog Devices GmbH
Otl-Aicher-Str. 60-64
80807 München
Germany
Tel: 49.89.76903.0
Fax: 49.89.76903.157

Analog Devices, Inc.
Japan Headquarters

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

Analog Devices, Inc.
Asia Pacific Headquarters

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

M02000-10/19

analog.com

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™