AnalogDialogue

Volume 54, Number 2, 2020

Your Engineering Resource for Innovative Design

26 Phased Array Antenna Patterns—Series

Visit *analogdialogue.com*

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

In This Issue

Bernhard Siegel, Editor

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

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

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

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

AnalogDialogue

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

Leveraging the On-Chip FIR and IIR Hardware Accelerators on a Digital Signal Processor

Mitesh Moonat, Applications Engineer and
Sanket Nayak, Applications Engineer

Abstract

Finite impulse response (FIR) and infinite impulse response (IIR) filters are the most frequently used digital signal processing algorithms—especially for audio processing applications. Thus, a significant portion of a processor core's time is consumed for FIR and IIR filtering in a typical audio system. The on-chip FIR and IIR hardware accelerators, also referred to as FIRA and IIRA, respectively, on a digital signal processor can be used to offload the FIR and IIR processing tasks, thus freeing up the core for other processing. In this article, we will discuss how to make use of these accelerators in practice with the help of different usage models illustrated with tested real-time examples.

Introduction

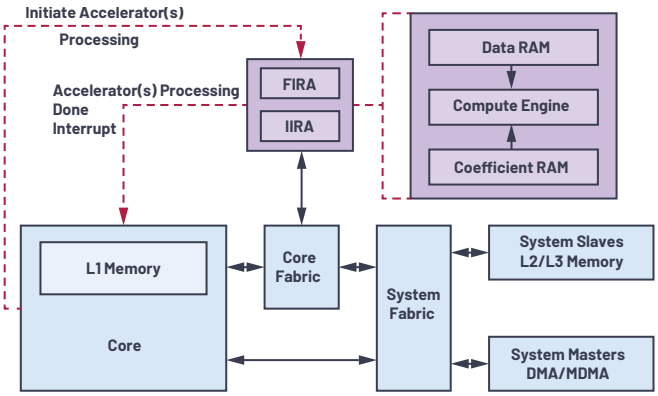

Figure 1. FIRA and IIRA system block diagram.

Figure 1 shows a simplified block diagram of FIRA and IIRA and how they interact with the rest of the processor system and resources.

- Both the FIRA and IIRA blocks mainly consist of a compute engine—multiply and accumulate (MAC) units—along with a small local data and coefficient RAM.
- To start the FIRA/IIRA processing, the core initializes a chain of DMA transfer control blocks (TCBs) in processor memory with channel-specific information. The core then writes the FIRA/IIRA chain pointer register with the start address of this TCB chain and then configures the FIRA/IIRA control register to start the accelerator processing. Once processing of all the channels is complete, an interrupt is sent to the core so that it can use the processed output for further operations.
- Theoretically, the best approach is to offload all the FIR and/or IIR tasks from the core to the accelerator(s) and allow the core to do something else in parallel. But in practice, this may not always be feasible, particularly when the core needs to use the output from the accelerator(s) for further processing and has no other independent tasks to finish in parallel. In such cases, we need to choose the appropriate accelerator usage model to achieve the best results.

In this article, we'll discuss various models to use these accelerators optimally for different application scenarios.

Using FIRA and IIRA in Real Time

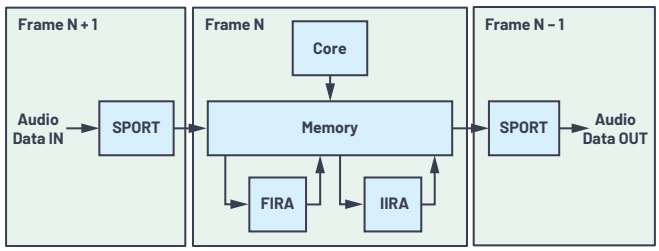

Figure 2. Typical real-time audio data flow.

Figure 2 shows a typical real-time PCM audio data flow diagram. One frame of digitized PCM audio data is received over a synchronous serial port (SPORT) and sent to memory via direct memory access (DMA). While reception of frame N+1 is going on, frame N is processed by the core and/or the accelerator(s), and the output of the previously processed frame (N-1) is sent out via SPORT to the DAC for digital-to-analog conversion.

Accelerator Usage Models

As discussed earlier, depending upon the application, the accelerators may need to be used in different ways to offload the maximum FIR and/or IIR processing tasks and to save the most possible core cycles for other operations. At a high level, the accelerator usage models can be divided into three categories: direct replacement, split task, and data pipelining.

Direct Replacement

- The core FIR and/or IIR processing is directly replaced by the accelerator(s) and the core simply waits for the accelerator(s) to finish the job.
- This model is effective only when the accelerator can process faster than the core; that is, using the FIRA block.

Split Task

- The FIR and/or IIR processing tasks are divided between the core and the accelerator(s).
- This model is especially useful when multiple channels are available to be processed in parallel.
- Based on a rough timing estimation, the total number of channels can be divided between the core and the accelerator(s) in such a way that both finish at approximately the same time.
- As shown in Figure 3, this usage model results in more core-cycle savings than the direct replacement model.

Data Pipelining

- The data flow between the core and accelerator(s) can be pipelined in such a way that both can work in parallel on different data frames.
- As shown in Figure 3, the core processes the N^{th} frame and then initiates the accelerator's processing of this frame. The core then continues in parallel to further process the $N\text{-}1^{th}$ frame output produced by the accelerator(s) in the previous iteration. This sequence allows the complete offloading of the FIR and/or IIR processing task to the accelerator(s) at the cost of additional output latency.
- The pipeline stages and, consequently, the output latency can increase depending upon the number of such FIR and/or IIR processing stages in the complete processing chain.

Figure 3 illustrates how the audio data frames flow between three stages—DMA IN, core/accelerator processing, and DMA OUT—for various accelerator usage models. It also shows how free core cycles increase compared to the core only model by fully or partially offloading the FIR/IIR processing to the accelerator across different accelerator usage models.

FIRA and IIRA on SHARC Processors

The following Analog Devices SHARC® processor families support on-chip FIRA and IIRA (oldest to the newest).

- ADSP-214xx (for example, ADSP-21489)
- ADSP-SC58x
- ADSP-SC57x/ADSP-2157x
- ADSP-2156x

Across processor families:

- Compute speed varies.
- The basic programming model remains the same except for the auto configuration mode (ACM) on ADSP-2156x processors.
- FIRA has four MAC units while IIRA has a single MAC unit.

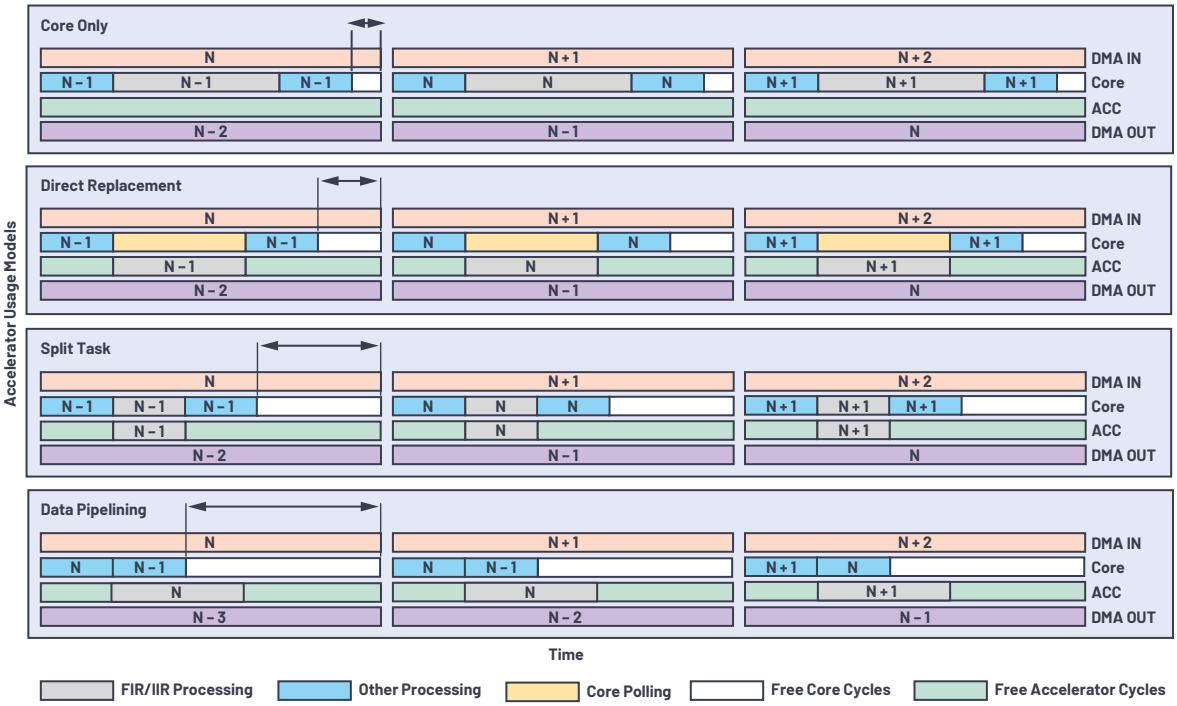

Figure 3. Accelerator usage models comparison.

FIRA/IIRA Improvements on ADSP-2156x Processors

ADSP-2156x is the latest addition to the SHARC processor family. It is the first 1 GHz single-core SHARC processor with FIRA and IIRA also capable of running at 1 GHz. FIRA and IIRA on ADSP-2156x processors offer various improvements over its predecessors, the ADSP-SC58x/ADSP-SC57x processors.

Performance Improvements

- Compute speed increased by eight times (SCLK-125 MHz to CCLK-1 GHz).
- Lesser data and MMR access latency between the core and the accelerators is possible due to closer integration of the core and the accelerators with the help of a dedicated core fabric.

Functional Improvements

Support for ACM was added to minimize core intervention required to handle accelerator processing. This mode comes with the following new, main features:

- Allows halting of the accelerator for dynamic task queuing.
- No channel number limitation.
- Trigger generation (master) and trigger wait (slave) support.
- Selective interrupt generation for each channel.

Experimental Results

In this section, we'll discuss the results of two real-time multichannel FIR/IIR use cases implemented with the help of different accelerator usage models on an ADSP-2156x evaluation board.

Use Case 1

Figure 4 shows the block diagram of use case 1. The sample rate is 48 kHz, the block size is 256 samples, and the ratio of core to accelerator channels used in the split task model is 5:7.

Table 1 shows the measured core and FIRA MIPS numbers along with the resultant core MIPS savings compared to the core only model. The table also shows additional output latency added by the corresponding usage model. As we can see, with accelerator usage, up to 335 core MIPS could be saved with a data pipelining usage model, at the cost of 1 block (5.33 ms) of output latency. Direct replacement and split task usage models also result in 98 MIPS and 189 MIPS savings, respectively, without any additional output latency.

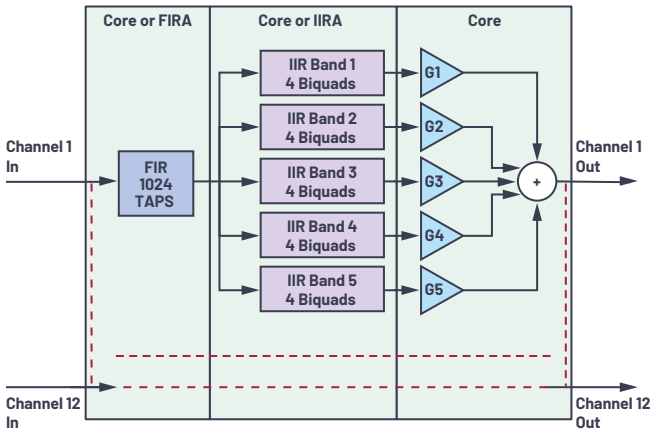

Figure 4. Use case 1 block diagram.

Table 1. Core and FIR/IIRA MIPS Summary for Use Case 1

Usage Model	Core MIPS	FIRA MIPS	IIRA MIPS	Core MIPS Saving	Usage Model Latency (ms)
Core Only	337	–	–	–	0
Direct Replacement	239	162	75	98	0
Split Task	148	96	44	189	0
Data Pipelining	2	161	75	335	5.33 (1 frame)

Use Case 2

Figure 5 shows the block diagram of use case 2. The sample rate is 48 kHz, the block size is 128 samples, and the ratio of core to accelerator channels used in the split task model is 1:1.

Like Table 1, Table 2 shows the results for this use case. As we can see, with accelerator usage, up to 490 core MIPS could be saved with a data pipelining usage model at the cost of 1 block (2.67 ms) of output latency. A split task usage model results in 234 core MIPS savings without any additional output latency. Note that unlike in use case 1, frequency-domain (fast convolution) processing is used for the core instead of time-domain processing. This is the reason why the core MIPS taken to process one channel is less than the FIRA MIPS taken, which results in a negative core MIPS savings for the direct replacement usage model.

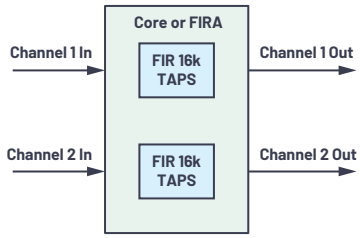

Figure 5. Use case 2 block diagram.

Table 2. Core and FIR/IIRA MIPS Summary for Use Case 2

Usage Model	Core MIPS	FIRA MIPS	Core MIPS Saving	Usage Model Latency (ms)
Core Only	493	–	–	0
Direct Replacement	515	511	–22	0
Split Task	259	257	234	0
Data Pipelining	3	511	490	2.67 (1 frame)

Conclusion

In this article, we have seen how significant core MIPS can be offloaded to FIRA and IIRA accelerators on ADSP-2156x processors, utilizing different accelerator usage models to achieve desired MIPS and processing profiles.

Further Reading

"Graphical Demonstration of ADSP-2156x FIR/IIR Accelerator Performance and Real Time Usage." Analog Devices, Inc.

Nayak, Sanket and Mitesh Moonat. "Engineer-to-Engineer Note EE-408: Using ADSP-2156x High Performance FIR/IIR Accelerators." Analog Devices, Inc., August 2019.

About the Author

Mitesh Moonat is currently working as an applications engineer in processor applications team at Bangalore (ADBL), India. He works on pre-/postsilicon validation, peripheral driver development, and support for SHARC® processors. He also worked on Blackfin® and ADSP-21xx processors in his career at ADI. His interests include processor architecture, digital signal processing algorithm optimization, module, and system-level debug of embedded systems. Mitesh joined Analog Devices in 2006. He graduated with a bachelor's degree in electronics and communication engineering from National Institute of Technology, Warangal, India. He can be reached at mitesh.moonat@analog.com.

About the Author

Sanket Nayak is a product applications engineer in the Processor Applications Team in Bangalore (ADBL), India. He joined ADI in 2016 and has worked on pre-/postsilicon validation of automotive DSPs, drivers/FuSa ROM design, development, and testing. He earned his bachelor's degree in electronics and communication engineering from PES Institute of Technology, Bangalore. He can be reached at sanket.nayak@analog.com.

Fast Reacting, Optical Encoder Feedback System for Miniature Motor Driven Applications

Jonathan Colao, Applications Engineer

Abstract

This article provides insight into common problems faced by designers of position-sensing interfaces for motor control in industrial automation—that is, being able to sense position on faster and smaller size types of applications. Capturing the information from encoders to accurately measure motor position is important to the successful operation of automation and robotic equipment. Fast and high resolution, dual, simultaneous sampling analog-to-digital converters (ADCs) are significant components in this system.

Introduction

Motor rotation information such as position, speed, and direction must be accurate in order to produce precise drivers and controllers across a wide variety of emerging applications—for example, in pick-and-place machines that mount microscopic components in the limited PCB area. Recently, motor controls have been miniaturized, enabling new applications in surgical robotics for healthcare and in drones for aerospace and defense. The smaller motor controllers also enable new applications in industrial and commercial installations. The challenge for designers is to meet the high accuracy requirement of the position feedback sensor in a high speed application, while at the same time infusing all components into the limited PCB space to be fit inside tiny enclosures, such as a robotic arm.

Motor Control

Motor control loops, as seen in Figure 1, are mainly made up of a motor, a controller, and a position-feedback interface. The motor turns a rotating shaft that causes the arms of a machine to move accordingly. The motor controller tells the motor when to apply force, stop, or continue rotating. The position interface in the loop provides rotational speed and position information to the controller. This data is central to the proper operation of a pick-and-place machine for the assembly of a tiny surface-mount PCB. All these applications require accurate position measurement information about the rotating object.

The position-sensor resolution must be very high—enough to accurately detect the motor shaft position, correctly pick up a tiny component, and place it accurately on a board. Also, higher motor rotational speeds lead to higher loop bandwidth and lower latency requirements.

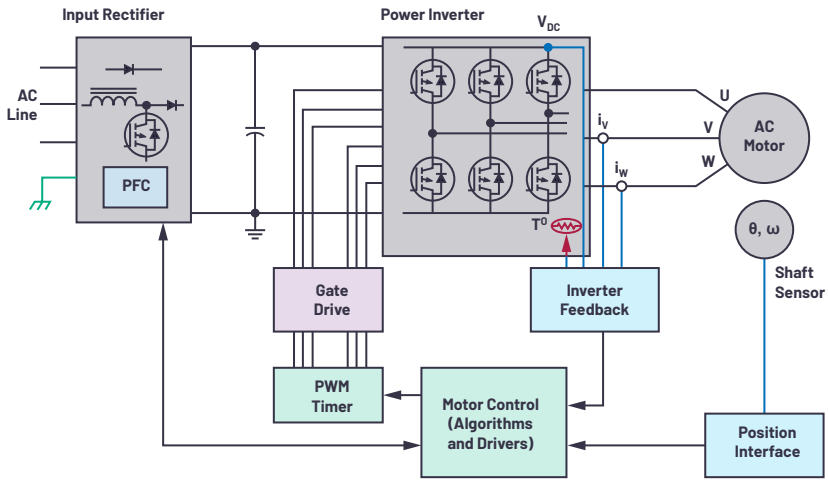

Figure 1. Closed-loop motor control feedback system.

Position-Feedback System

In a lower end application, an incremental sensor along with a comparator may suffice for position sensing, while a higher end application will require more complex signal chains. These feedback systems comprise the position sensor, followed by analog front-end signal conditioning, the ADC, and its driver before data gets into the digital domain. One of the most precise position sensors is the optical encoder. An optical encoder is composed of an LED light source, a marked disc attached to the motor shaft, and a photodetector. The disc features a masked pattern of opaque and transparent areas that obscure the light or allow it to pass through. The photodetectors sense the resulting light and the on/off light signals are converted to electric signals.

As the disc turns, the photodetectors—in conjunction with the patterns of the disc—produce small sine and cosine signals, in the mV or µV level. This system is typical in an absolute position optical encoder. These signals are fed to an analog signal conditioning circuitry, usually consisting of a discrete amplifier or an analog PGA to gain the signal up to the 1 V p-p range—commonly to fit an ADC input voltage range for maximum dynamic range. Each of the amplified sine and cosine signals are then acquired by a simultaneous sampling ADC's driver amplifier.

The ADC must feature simultaneous sampling on its channels such that the sine and cosine data points are taken at the exact same point in time, as that combination provides the shaft position information. The ADC conversion results are passed to an ASIC or microcontroller. The motor controller queries the encoder position every PWM cycle and uses this data to drive the motor based on the instructions it receives. In the past, system designers would have to trade ADC speed or channel count to fit restrictive board footprints.

Optimizing Position Feedback

The demands of evolving technology have resulted in innovation in motor control applications that require high accuracy position detection. The optical encoder resolution can be based on the number of slots inscribed from fine lithography in a disc, usually hundreds or thousands. Interpolating these sine and cosine signals to a high speed, high performance ADC will enable us to create higher resolution encoders without requiring system changes to the encoder disc. For example, when an encoder sine and cosine signal are sampled at a slower rate, fewer values of the signal are captured, as shown in Figure 3; this also limits the accuracy of the position cap. In Figure 3, when the ADC samples at a faster rate, more detailed values of the signal are captured, and a higher accuracy position is determined. A high speed sampling rate of the ADC allows oversampling, further improving the noise performance, removing some digital postprocessing needs. At the same time, it reduces the output data rate from the ADC; that is, allowing for slower serial frequency signals, hence simplifying the digital interface. The motor position feedback system is mounted in the motor assembly, which can be pretty small in certain applications. So size is vital to fit in the limited PCB area of the encoder module. The emerging of multiple channel components in a single, tiny package are best suited for space saving.

Optical Encoder Position Feedback Design Example

An example of an optimized solution for an optical encoder position feedback system is shown in Figure 4. The circuit can be easily interfaced to an absolute type of optical encoder where differential sine and cosine signals from the encoder can be easily captured by the circuit. The ADA4940-2 front-end amplifier is a dual-channel, low noise, fully differential amplifier that drives the AD7380, a dual-channel, 16-bit, fully differential, 4 MSPS, simultaneous sampling SAR ADC, housed in a small 3 mm × 3 mm LFCSP package. The on-chip 2.5 V reference would allow minimum component requirements for this circuit. The V_{CC} and V_{DRIVE} of the ADC and the supply rails of the amplifier driver can be powered by an LDO regulator, such as the LT3023 and LT3032. When these reference designs are interfaced—for example, with a 1024-slot optical encoder that produces 1024 cycles of sine and cosine in one revolution of the encoder disc—the 16-bit AD7380 samples each encoder slot at 2^{16} codes, overall increasing the encoder resolution up to 26 bits. The 4 MSPS throughput rate ensures that detailed sine and cosine cycles are captured and encoder positions are up to date. The high throughput rate enables oversampling on-chip, which reduces the time penalty of digital ASICs or microcontrollers feeding the precise encoder position to the motor. An extra benefit of the AD7380's on-chip oversampling is that it allows for an additional 2 bits of resolution, which can be easily used with an on-chip resolution boost feature. The resolution boost can further improve the accuracy up to 28 bits. Application note AN-2003 details this oversampling and resolution boost feature of the AD7380.

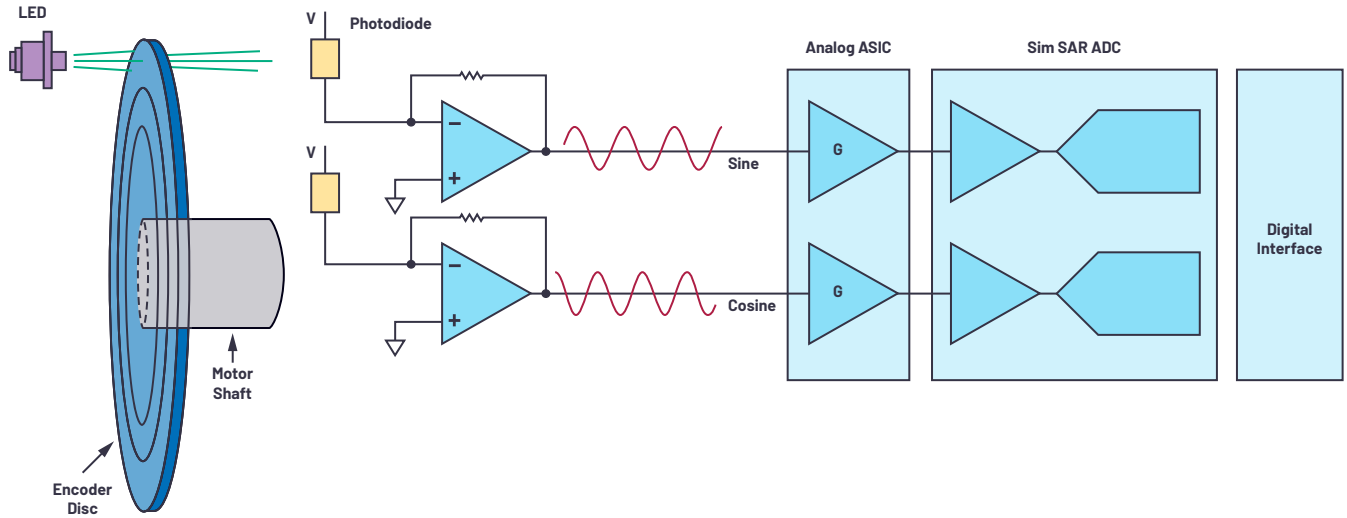

Figure 2. Position feedback system.

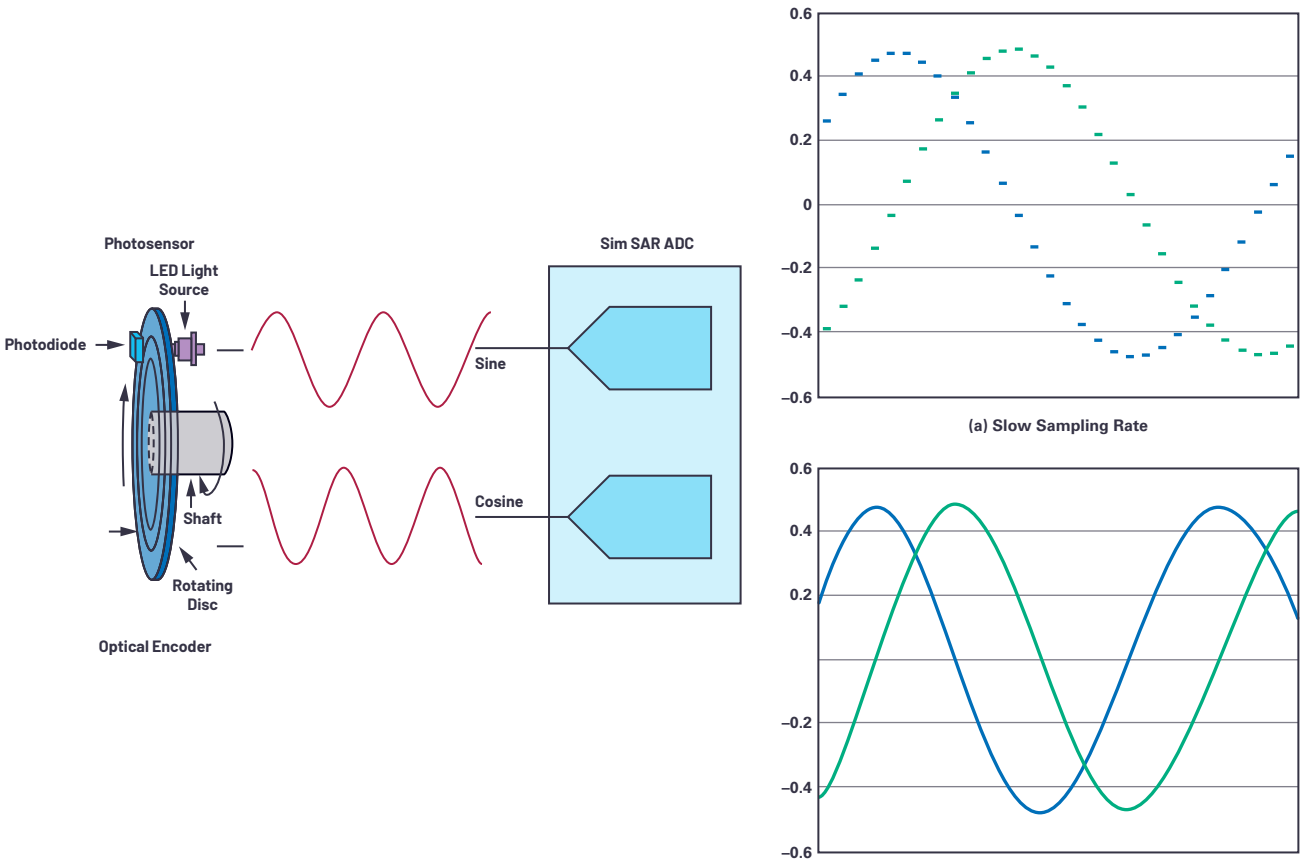

Figure 3. Sampling rate.

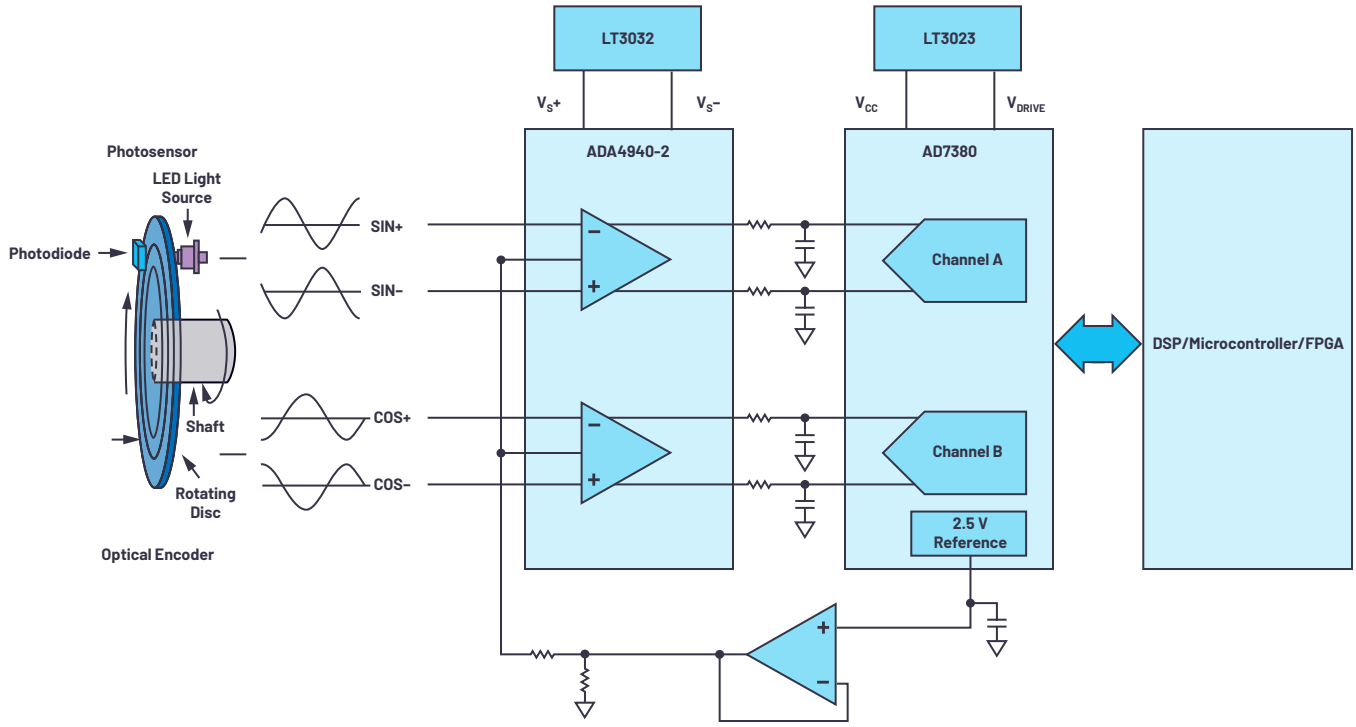

Figure 4. Optimized feedback system design.

Conclusion

Motor control system demands for higher accuracy, higher speed, and miniaturization are increasing. Optical encoders are used as motor position-sensing devices. To do this, the optical encoder signal chain must have a high level of accuracy when measuring the motor position. A high speed, high throughput ADC accurately captures information and feeds motor position data to the controller. The AD7380's speed, density, and performance answers the industry's demand while enabling higher levels of accuracy and optimization in the position-feedback system.

About the Author

Jonathan Colao is product applications engineer in the Precision Converter Technology Group at Analog Devices. He joined ADI in Cavite, Philippines, in 2006 as a product engineer. He graduated from Xavier University in Cagayan de Oro, Philippines, with a bachelor's degree in electronics engineering. He can be reached at jonathan.colao@analog.com.

Speed Up the Design of EMI Filters for Switch-Mode Power Supplies

Henry Zhang, Applications Director and
Sam Young, Applications Manager

Introduction

Switch-mode power supplies are used throughout modern electronics systems mainly because of their high efficiency power conversion. One side effect of the proliferation of switch-mode supplies is the noise they produce. This is commonly referred to as electromagnetic interference (EMI), EMI noise, or just noise. For example, the input side switch current of a typical buck converter is a pulsating current rich in harmonic content. Fast turn-on and turn-off of power transistors create sudden interruptions of current flow, resulting in high frequency voltage ringing and spikes.

The problem is that high frequency noise can couple to other devices in the system, degrading the performance of sensitive analog or digital signal circuits. Because of this, many standards have arisen to set acceptable limits on EMI. To meet these limits for a switch-mode supply, one must first quantify its EMI performance and, if necessary, add proper input EMI filtering to attenuate the EMI. Unfortunately, EMI analysis and filter design can be a difficult task, typically requiring a time-consuming iterative process of design, build, testing, and redesign—that is, assuming one has proper test equipment. To speed up the process of EMI filter design to meet EMI specifications, this article shows how conducted EMI noise analysis and filter design can be easily estimated and prebuilt using ADI's LTpowerCAD® program.

Different Types of EMI: Radiated and Conducted Noise, Common Mode, and Differential Mode

There are two major type of EMI: radiated and conducted. In a switch-mode supply, radiated EMI is usually generated by high dv/dt noise at the switching nodes. Industry standards for radiated emissions usually cover the frequency band from 30 MHz to 1 GHz. At these frequencies, radiated EMI from switching regulators is produced mainly by switching voltage ringing and spikes, and can depend heavily on PCB board layout. Other than what is inherently built into good layout practices, it is nearly impossible to precisely predict how much radiated EMI a switch-mode supply will transmit "on paper." One must simply build the board and measure its EMI in a sufficiently well-designed EMI lab to quantify its radiated noise level.

Conducted EMI results from the rapid changes in a switching regulator's conducted input current, including common-mode (CM) and differential-mode (DM) noise. Standard industry limits for conducted emissions usually cover a lower frequency range than radiated emissions, namely from 150 kHz to 30 MHz. Figure 1 shows the generic conduction paths of the common-mode and differential-mode noise of a dc-to-dc power supply (the DUT in an EMI lab).

To quantify conducted input EMI, a line impedance stabilization network (LISN) is placed at the regulator's input, providing a standard input source impedance. CM conducted noise is measured between each input line and earth ground. CM noise is generated at high dv/dt switching nodes, couples through the device's parasitic PCB capacitance, CP, to earth ground, then travels to the supply input LISN. Like radiated EMI, the high frequency switching node ringing and parasitic capacitance cannot be easily and accurately modeled in a paper design.

DM noise is measured differentially between two input lines. DM conducted noise arises from the high di/dt, pulsating input current of the switch-mode supply. Fortunately, unlike the other EMI types, the pulsating input current and the resulting relatively low frequency EMI generated at the input capacitors and LISN circuit can be predicted by software, such as LTpowerCAD, with acceptable accuracy.

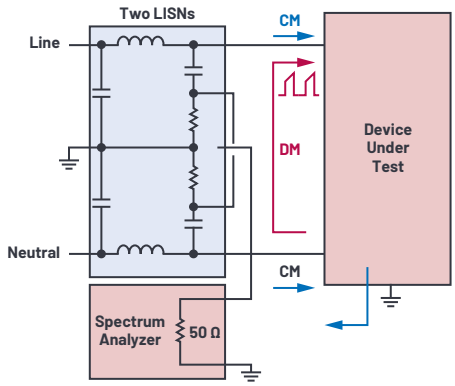

Figure 1. Conceptual overview of LISN-based measurement of differential-mode and common-mode conducted EMI of a switch-mode supply.

Figure 2 shows a typical EMI noise plot of a switch-mode, step-down buck supply *without* an input EMI filter. The most significant EMI spike occurs at the switching frequency of the supply, followed by additional spikes at its harmonic frequencies. Figure 2 shows an EMI plot where the peak values of these spikes exceed the CISPR 22 EMI limits. To meet the standard, an EMI filter is required to attenuate the differential-mode EMI.

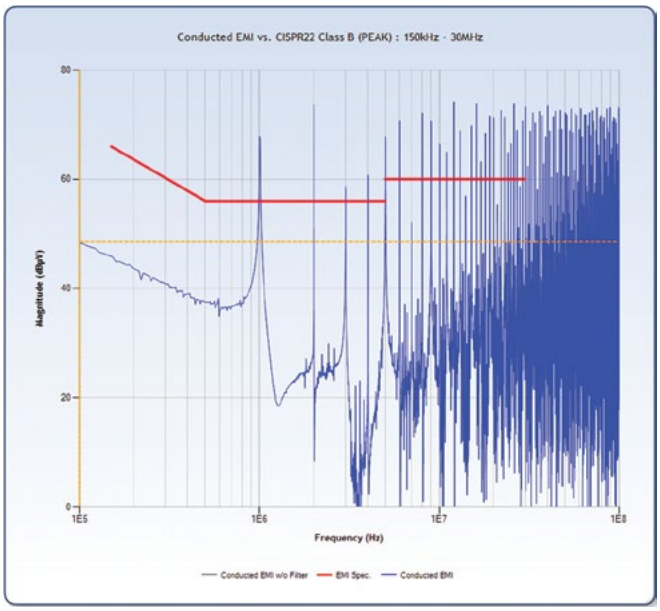

Figure 2. Typical EMI plot of a switch-mode buck supply without an input EMI filter.

Differential-Mode Conducted EMI Filter

Figure 3 shows a typical differential-mode conducted EMI noise filter on the input side of a switch-mode supply. In this case, we've added a simple first-order, low-pass L_fC_f network between the supply's local input capacitors C_{IN} (EMI noise source side) and the input source (LISN receiver side). This matches a standard EMI lab test setup, where the LISN network is inserted on the filter capacitor C_f side of the LC filter. The differential signal across the LISN resistor R2 is measured by the spectrum analyzer to quantify DM conducted EMI noises.

Figure 4 shows the LC filter attenuation gain plot. At very low frequency, the inductor is low impedance, essentially shorted, while the capacitor is high impedance, essentially open circuit. The resulting LC filter gain is 1 (0 dB), allowing dc to pass through without attenuation. As frequency rises, a gain spike appears at the resonant frequency of L_fC_f. As the frequency rises above the resonant frequency, the filter attenuates at a rate of –40 dB/decade. At relatively higher frequencies, the filter gain increasingly becomes a function of parasitic components: namely, the filter capacitor's ESR and ESL and filter inductor's parallel capacitance.

Because this filter's ability to attenuate quickly rises over frequency, the magnitude of first few low frequency noise harmonics overwhelmingly determine the size of the EMI filter—where the fundamental component of the supply's switching frequency (f_{SW}) is the most significant target. Therefore, we can focus on the EMI filter's lower frequency gain in efforts to meet industry standards.

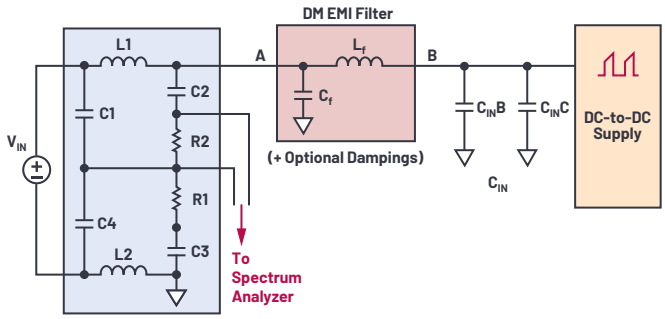

Figure 3. Differential-mode EMI noise filter (from Node B to Node A).

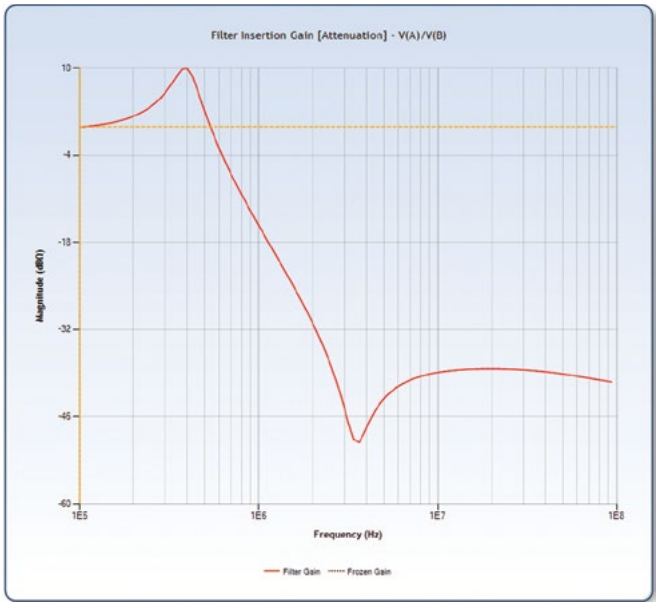

Figure 4. Typical single LC EMI filter insertion gain vs. frequency plot.

LTpowerCAD Can Predict Supply-Specific Filter Performance

LTpowerCAD is a power supply design assistance tool that can be downloaded at no charge at analog.com/LTpowerCAD. The program is designed to enable engineers to design and optimize complete power supply parameters in a few simple steps and in a few minutes.

LTpowerCAD leads a user through the entire power supply selection and design process, beginning with the user's power supply specifications. From there, LTpowerCAD narrows the range of suitable solutions and then helps in the selection of power stage components, as well as optimizes supply efficiency, design loop compensation, and load transient response.

The feature we are interested in here is LTpowerCAD's input EMI filter design tool, which enables an engineer to quickly estimate the differential-mode conducted EMI and determine what filter components may be required to meet EMI standards. LTpowerCAD's filter tool can significantly reduce design time and cost by producing realistic results—before a single circuit board is built and tested.

EMI Filter Design in LTpowerCAD

Overview

Let's look at a DM EMI filter design example. Figure 5 shows the LTpowerCAD schematic design page with the component selection for a supply using the LTC3833 buck converter, operating with 12 V input and 5 V/10 A output, running at 1 MHz switching frequency, f_{SW}. Before designing an EMI filter, design the buck converter by selecting the switching frequency, power stage inductor, capacitors, and FETs.

After power stage components are selected, click the EMI design icon to open the integrated DM EMI filter tool window, as shown in Figure 6. The EMI design window shows a detailed input filter network, L_fC_f, between the power supply input capacitors C_{INB}/C_{INC} and the source LISN. There are optional damping circuits, such as networks C_{dA}/R_{dA} on the LISN side, network C_{dB}/R_{dB} on the supply input capacitor side, and the optional damping resistor Rf_p across the filter inductor L_f. The estimated conducted EMI noise plot and the selected EMI standard limits appear on the right side of Figure 6.

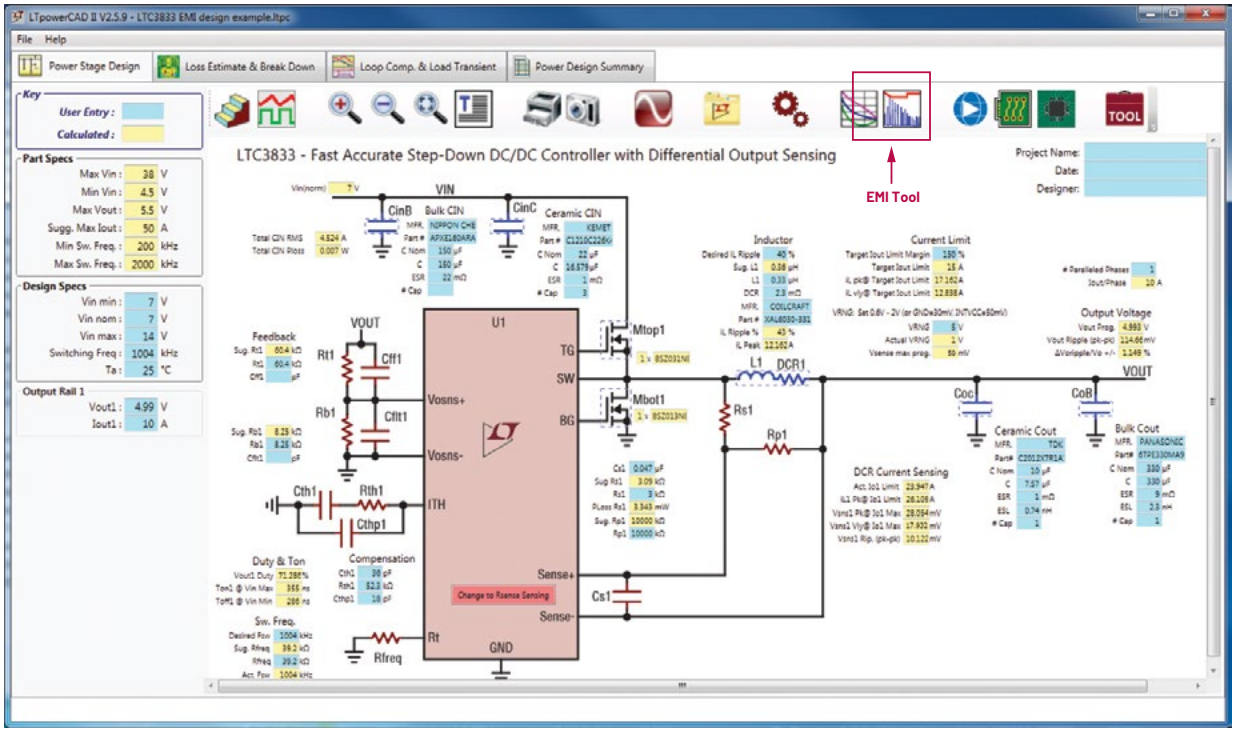

Figure 5. The LTpowerCAD schematic design page and the integrated EMI tool icon.

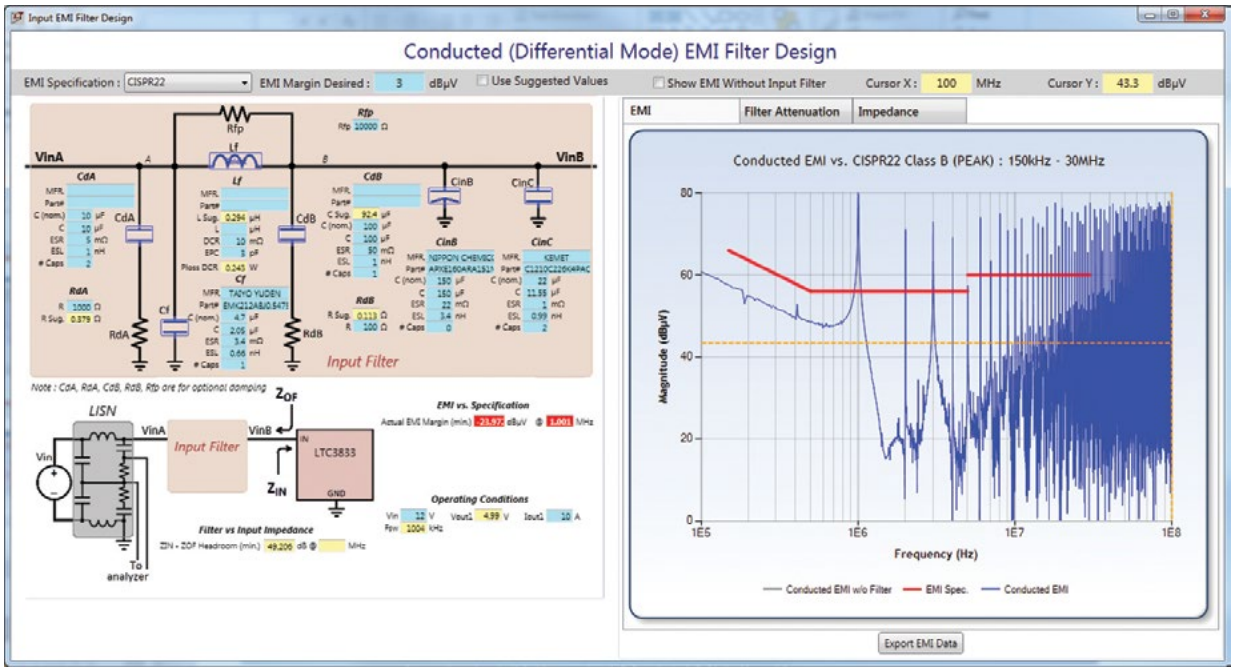

Figure 6. LTpowerCAD conducted DM EMI filter design window (L_f = 0, no filter).

Select an EMI standard

When designing an EMI filter, you'll want to see the design's targets—namely, the EMI standard itself. LTpowerCAD includes built-in plots for CISPR 22 (for IT equipment), CISPR 25 (for automotive devices), and MIL-STD-461G standards. Simply select your desired standard from the **EMI Specification** dropdown menu.

For example, in Figure 6, the filter inductor value is set at 0 to show the design's EMI results without an input filter. EMI spikes at the fundamental and harmonic frequencies, all exceeding the displayed CISPR 25 limits, resulting in a red warning on the **EMI vs. Specification** schematic display.

Set the EMI Filter Parameters

After selecting the desired EMI standard, enter the desired EMI margin—what distance do you want between the selected standard limits and the peak values of the fundamental. A 3 dB to 6 dB margin is generally a good starting point. From these choices, for a given filter capacitor, C_f, and supply operating condition, the program calculates the suggested filter inductor value, L sug., shown in a yellow cell in LTpowerCAD. Enter an inductor value in the L cell slightly greater than the suggested value to meet the EMI limit with the desired margin.

In this example, Figure 7 shows the design tool recommending a 0.669 µH filter inductance, along with the entered 0.72 µH inductance to meet the requirement. The benefits of the filter can be explored by comparing the results with and without the filter. Turn on the **Show EMI Without Input Filter** option to see the filtered results overlayed above a gray no-filter plot.

There is an important detail in choosing the filter capacitor C_f. If it is a multilayer ceramic capacitor (MLCC) with X5R, X7R, etc. type of dielectric material, its capacitance value can drop significantly with dc bias voltage. Because of this, in addition to the LTpowerCAD nominal capacitance, C(nom), the user should also enter its real capacitance under the applied dc bias voltage (V_{INA} or V_{INB}).

The derating curve can be found from capacitor vendors' data sheet. If the MLCC capacitor is chosen from the LTpowerCAD library, its derating with dc bias voltage is automatically estimated by the program.

Another component variation arises in the input filter inductor, which can have a nonlinear inductance due to its saturation with direct current. The inductance value may drop noticeably with increased load current, especially for a ferrite bead type of inductor. Users should enter the real inductance to produce accurate EMI predictions.

Check Filter Attenuation Gain

In the Figure 7 EMI plot with the input filter, there is a noise spike due to the LC input filter resonance at 245 kHz, a frequency *lower* than the supply switching frequency. Figure 8 shows the filter attenuation gain plot in place of the EMI results in the LTpowerCAD EMI window (click the **Filter Attenuation** tab), revealing the filter's resonant attenuation gain at 245 kHz.

In some cases, the LC resonance peak can result in a spike that exceeds the EMI standard. To attenuate this resonant peak, a pair of optional damping components C_{dA} and R_{dA} can be added in parallel with the filter capacitor C_f. In addition to showing the attenuation plot, LTpowerCAD simplifies the selection process for these components. In general, choose a damping capacitance, C_{dA}, that is around two to four times the real filter C_f value. LTpowerCAD will suggest a damping resistor R_{dA} value to push down the resonant peak.

Check Filter Impedance and Supply Input Impedance

When adding an input EMI filter in front of a switch-mode supply, the filter output impedance, Z_{OF}, can interact with the supply input impedance, Z_{IN}, causing undesirable oscillation. To avoid this unstable situation, the magnitude of the EMI filter output impedance, Z_{OF}, should be much lower than the magnitude of the supply input impedance, Z_{IN}, with enough margin. Figure 9 shows the concept of Z_{OF} and Z_{IN} and the stability margin between them.

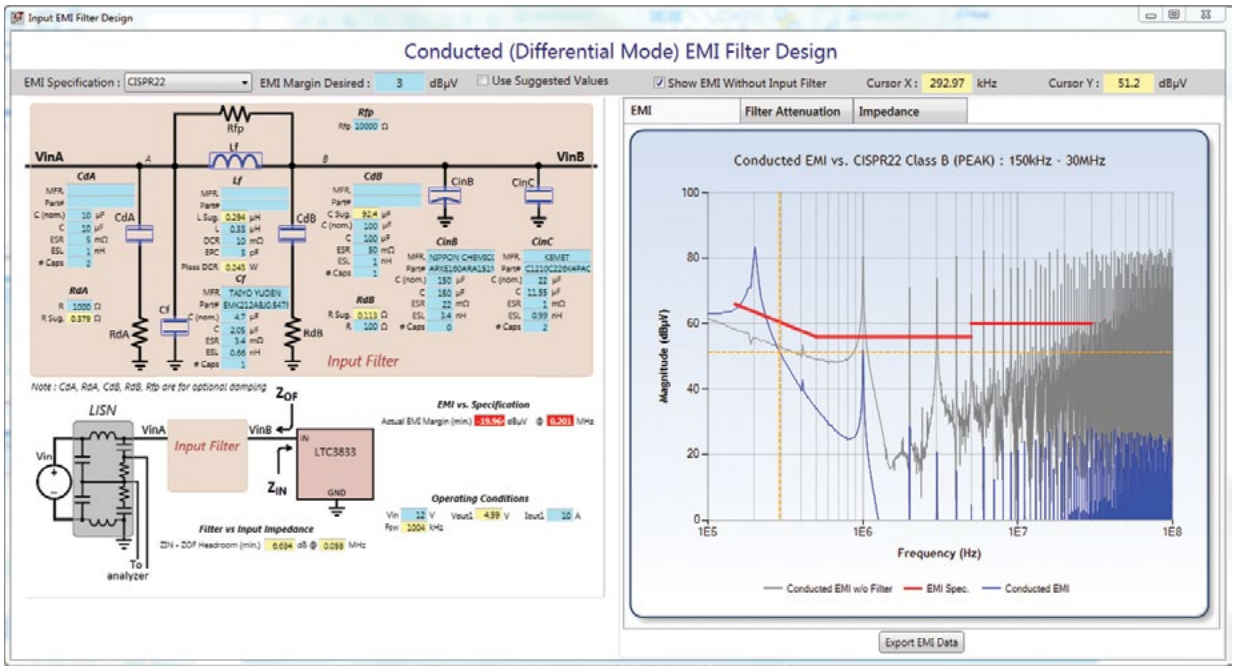

Figure 7. Select filter inductor value to meet EMI standard limit.

To simplify the problem, an ideal power supply with high feedback loop bandwidth can be treated as a constant power load; that is, input voltage V_{IN} times input current is constant. As input voltage increases, its input current decreases. Therefore, the ideal power supply has a negative input impedance $Z_{IN} = -(V_{IN}^2)/P_{IN}$.

To make it easy to design the input filter, LTpowerCAD displays the filter output impedance Z_{OF} and supply input impedance Z_{IN} on the impedance plot shown in Figure 10. Note the supply input impedance is a function of input voltage and input power. The worst case, which is the lowest level of impedance, occurs at the minimum V_{IN} and maximum P_{IN} condition.

As shown in Figure 10, the EMI filter output impedance has a peak point at the resonant frequency caused by filter inductor L_f and supply input capacitor C_{IN}. In a good design, the magnitude of this peak should be lower than worst-case Z_{IN}, with enough margin. In case it is necessary reduce this peak level, there is another pair of optional damping components, capacitor C_{dB} and resistor R_{dB}, in parallel with the supply input capacitor C_{IN}. This C_{IN} side damping network can effectively reduce the Z_{OUT} peak. The suggested C_{dB} and R_{dB} values are provided by LTpowerCAD EMI tool.

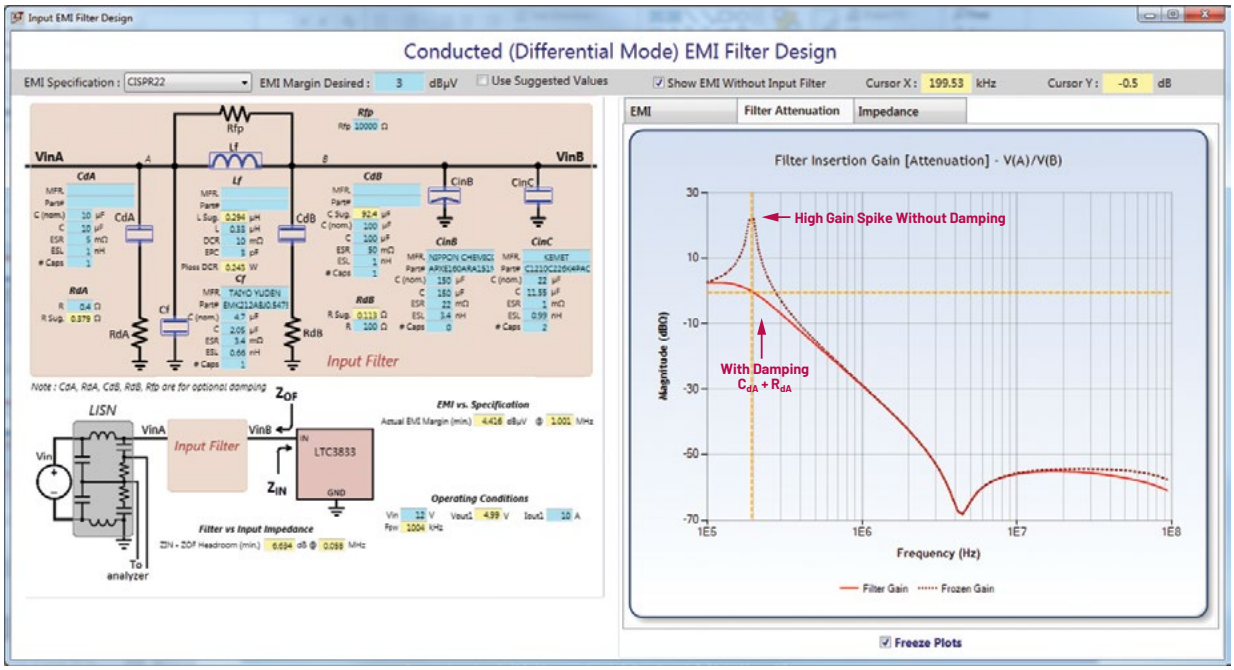

Figure 8. EMI filter attenuation gain (with and without damping on the LISN side).

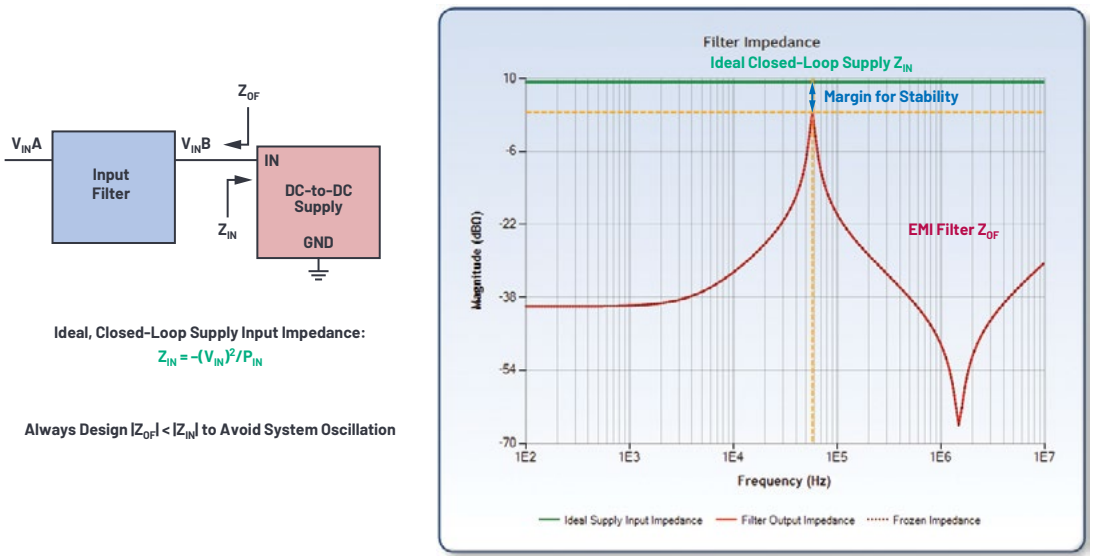

Figure 9. Check the EMI filter output impedance and supply input impedance for stability.

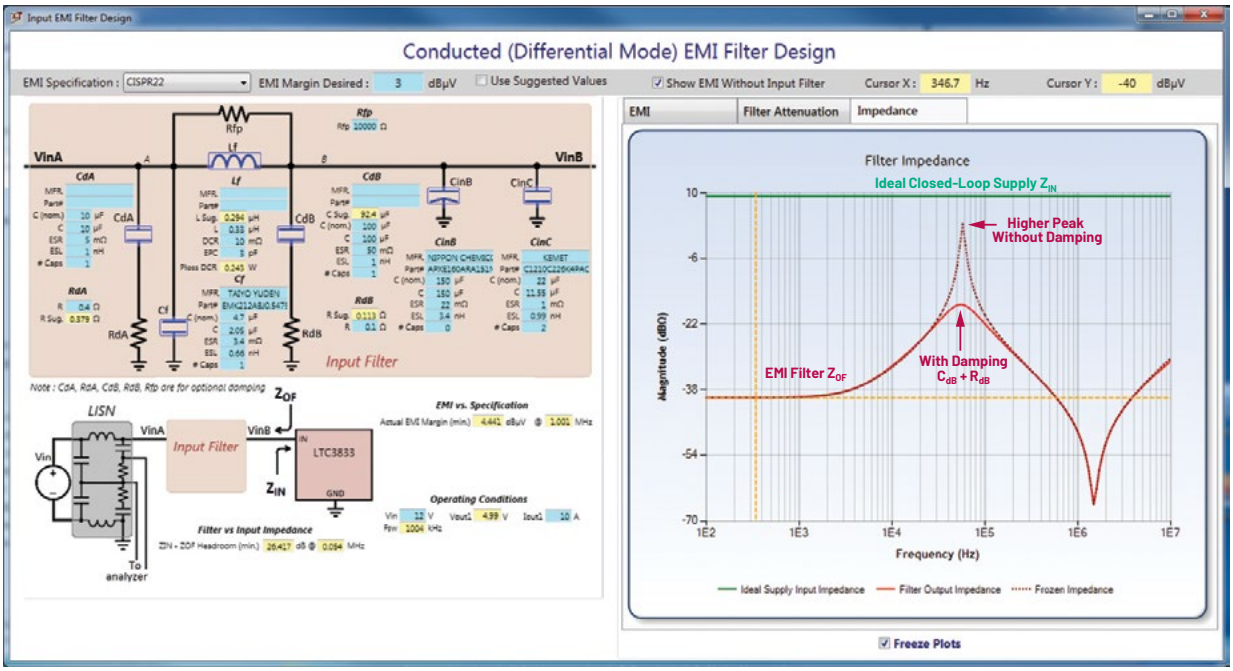

Figure 10. LTpowerCAD EMI filter impedance plot (with and without damping).

Accuracy of the LTpowerCAD EMI filter tool

The accuracy of the LTpowerCAD EMI filter tool can be seen by running an LTpowerCAD design against a real board, lab EMI test. Figure 11 shows the results of a comparison where the real-world test was performed using a modified LTC3851 buck supply demo board running at 750 kHz with a 12 V input voltage, 1.5 V output voltage, and 10 A load current. As shown in Figure 11, tested EMI data and LTpowerCAD modeled EMI data match well for lower frequency noise peaks, while real-world tested peaks are a few dB lower than modeled EMI peaks.

There are larger mismatches at higher frequency noise peaks, but these are of lower importance because the DM conducted EMI filter size is mainly determined by lower frequency noise spikes. Some of this discrepancy is due to the accuracy of inductor and capacitor parasitic models, including PCB layout parasitic values; accuracy that is beyond what is possible in a PC-based design tool, for now.

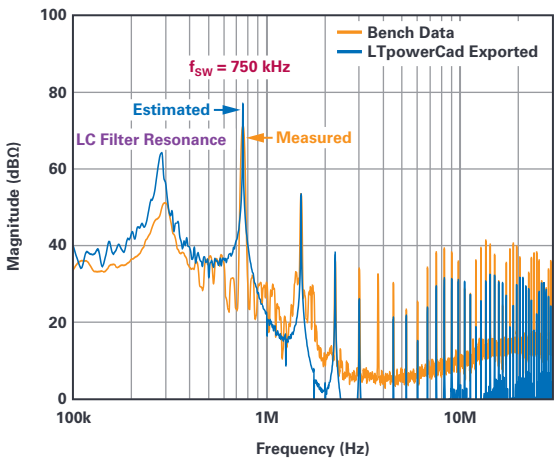

Figure 11. Real board lab measured vs. LTpowerCAD-estimated EMI (12 V_{IN} to 1.5 V_{OUT}/10 A buck example).

Please note the LTpowerCAD filter tool is an estimation tool, providing an initial design point for EMI filters. Nothing can replace a real lab test of a prototype supply board for truly accurate EMI data.

Summary

Many industries are using systems that require increasingly careful control of transmitted electromagnetic signals. To this end, there are a number of published clear standards on EMI. At the same time, the number of switch-mode power supplies is increasing, and they are being placed ever closer in proximity to sensitive circuits. Switch-mode power supplies are strong sources of EMI, so their noise output must be quantified and reduced in many cases. The problem is, EMI filter design and testing is a time-consuming and costly iterative process.

LTpowerCAD enables designers to save time and cost by precluding real-world design and test with predictive computer-based design tools. Its easy to use EMI filter tool predicts the results of differential-mode conducted EMI filters, including optional damping networks, to minimize EMI while retaining a stable supply. Lab test results verify the accuracy of predictive models.

About the Author

Henry Zhang is ADI's Power by Linear™ applications director. He received a B.S.E.E. degree from Zhejiang University, China in 1994 and his M.S. and Ph.D. degrees in electrical engineering from Virginia Polytechnic Institute at State University, Blacksburg, Virginia in 1998 and 2001, respectively. He has been with Linear Technology (now part of ADI) since 2001. He can be reached at henry.zhang@analog.com.

About the Author

Sam Young is an applications manager for Analog Devices, supporting µModule® regulator products and LTpowerCAD. He has over 10 years of experience in supporting regulator designs for for a broad range of applications and products. He can be reached at samuel.young@analog.com.

RAQ Issue 176: Ultralow Noise, 48 V, Phantom Microphone Power Supply Using a Tiny DC-to-DC Boost Converter

Thomas Mosteller, Field Applications Engineer and
Christopher Jarboe, Applications Engineer

Question:

Can I produce a compact, ultralow noise, phantom power supply (48 V) from a 5 V, 12 V, or 24 V input?

Answer:

You can, using a simple boost converter, a filter circuit to reduce EMI, and a little trickery to keep the size small.

Professional condenser microphones require a 48 V supply to charge the internal capacitive transducer and power the internal buffer for the high impedance transducer output. This supply is low current, typically only a few mA, but it must be very low noise because the microphone's output levels are quite low and the buffer doesn't have very good power supply ripple rejection. In addition, the phantom supply must not inject EMI into adjacent low level circuits, which is always a challenge in tightly packed products.

A very high performance supply can be built using the LT8362 boost converter, which features a 60 V, 2 A switch and is capable of running at frequencies up to 2 MHz, all in a package as small as 3 mm × 3 mm. The circuit presented here is based on the standard LT8362 demo board DC2628A, the schematic of which is shown in Figure 1.

The input EMI filter on the demo board does a good job of taking care of high frequency noise, aided by the switching inductor, which appears in series with the input. The situation is not as good on the output. The output EMI filter effectively suppresses noise in the MHz region, but has little effect on noise in the audio range. This noise mostly arises due to the 30× gain in the feedback loop amplifying the reference noise of the LT8362.

One approach to cure this noise is to add capacitance at the output. Given enough capacitance this would work, but with a 48 V output, the lowest practical capacitor working voltage of 63 V means that the required caps are both big and expensive. A second approach would be to increase the LT8362 output a volt or two and add an LDO regulator to the output. This requires a high voltage LDO regulator, which usually costs more than the low voltage counterpart. In addition, while these regulators can have low noise at lower output voltages, devices that use a voltage reference also suffer from the same reference noise multiplication issue as the LT8362.

A third approach is to take advantage of the fact that the sensitivity of the microphone output is not highly dependent on the supply voltage, so the phantom supply does not require perfect regulation. This means we can put some resistance in series with the output caps to increase their effectiveness; however, this only partly mitigates the size of the high voltage caps.

A better approach is to make the output capacitors seem bigger than they really are. This can be accomplished with an old school technique called capacitance multiplication. This simple circuit can be seen in the gray shaded area of Figure 2.

Here, the 100 µF cap controls the ripple on the base current, so its effect on the collector current is amplified by the beta of the NPN transistor. The effect is dramatic. Figure 3a shows the output of the LT8362 circuit at C4 (before the filter) with a 1 kΩ load (50 mA).

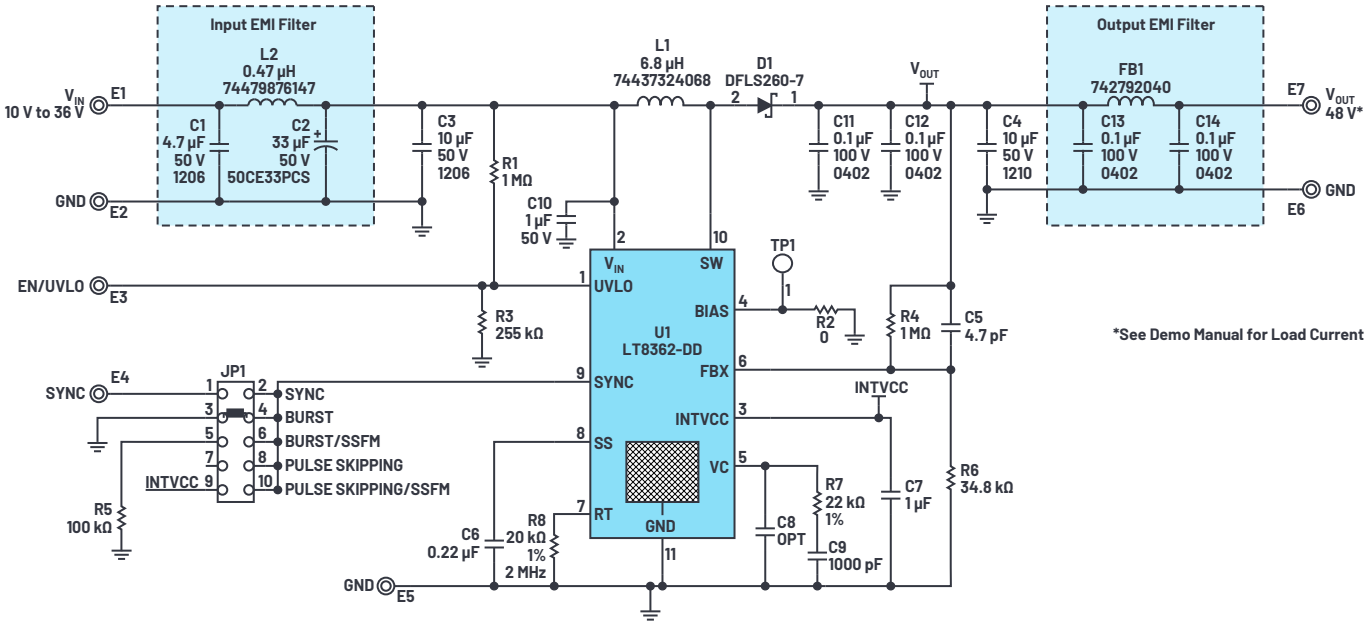

Figure 1. Schematic of demonstration circuit DC2628 used to build a phantom power supply.

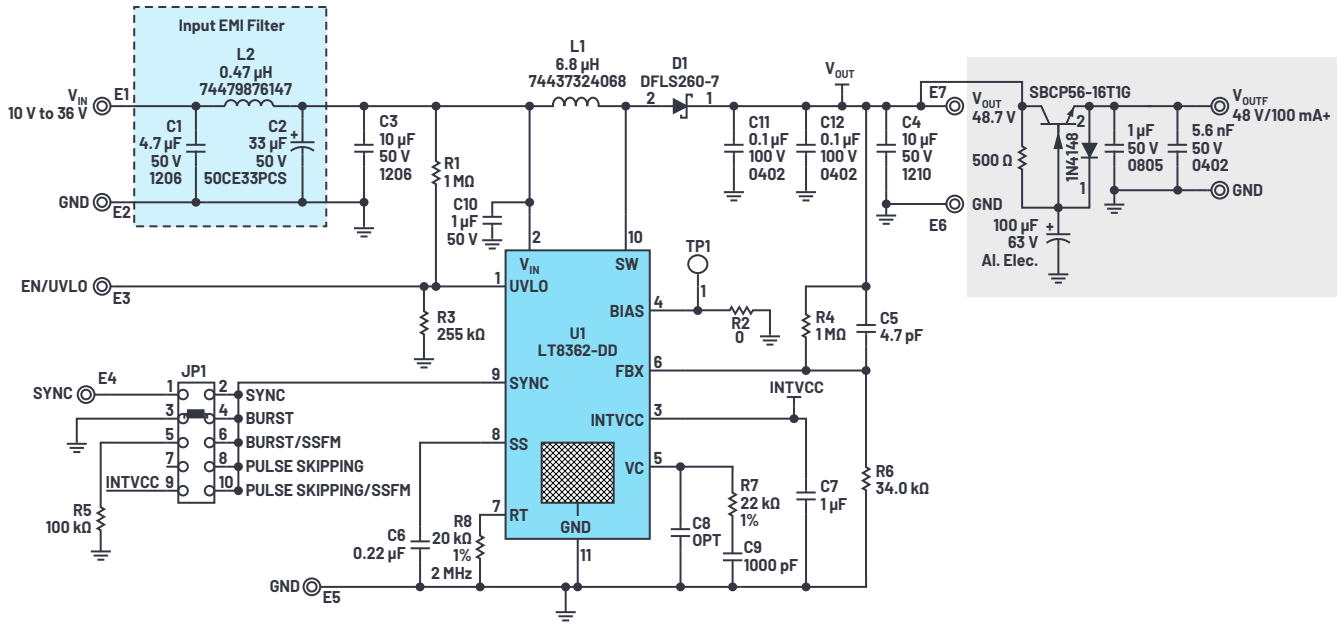

Figure 2. The same circuit as Figure 1, but with a capacitance multiplier (gray) at the output to suppress audio frequency noise produced by the switching regulator.

The noise is roughly 80 mV p-p, which represents about 0.2% noise content. While this may be sufficient for noncritical applications, the output after the filter is substantially better at approximately 1 mV p-p, as shown in Figure 3b. This represents about 0.002% or 20 ppm noise content—sufficient for even the most demanding applications. Figure 4 shows the benchtop setup.

The transistor SBCP56-16T1G was chosen for high V_{CBEO} (80 V) and high beta at low currents. The high beta gives the capacitance multiplier high apparent capacitance and relatively constant drop as output current varies. The output voltage drops from 47.8 V with a 2 kΩ load to 47.5 V at 500 Ω load, which is sufficient for microphone applications. Do not substitute another transistor without testing for noise and regulation.

The tests were run with 16 V input, but the performance will be similar with 12 V to 24 V_{IN}. Some applications may require boosting from 5 V, which can be accomplished by decreasing the LT8362's switching frequency from 2 MHz to 1 MHz to allow for the 75 ns minimum off time. This would also require increasing L1 to about 10 µH to 15 µH and doubling up on bulk output cap C4 to maintain equivalent performance.

References

Nelson, Carl. "Application Note 19: LT1070 Design Manual." Linear Technology, Inc., June 1986.

Williams, Jim. "Application Note 101: Minimizing Switching Regulator Residue in Linear Regulator Outputs." Linear Technology, Inc., July 2005.

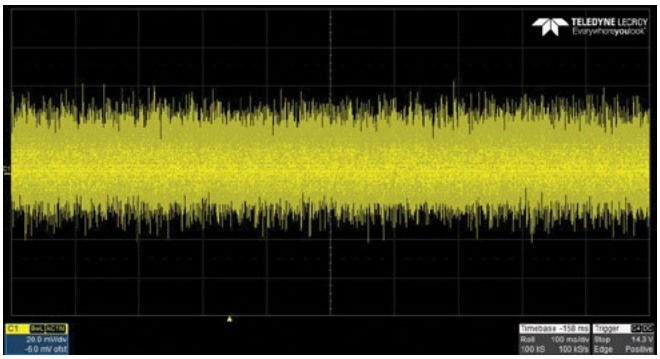

(a)

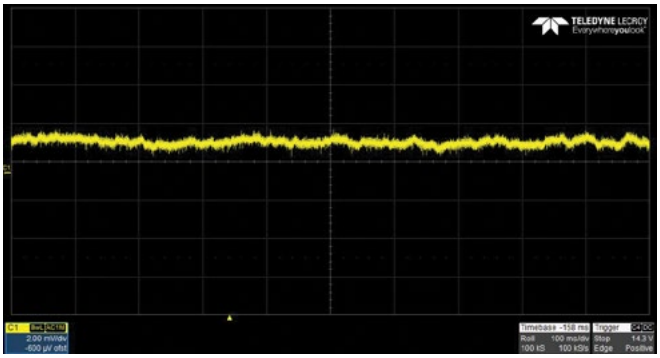

(b)

Figure 3. Before and after filter. (a) The output of the boost regulator presents about 0.2% noise content when measured at C4 (before filter). (b) The postfilter output contains a much improved 0.002% noise content.

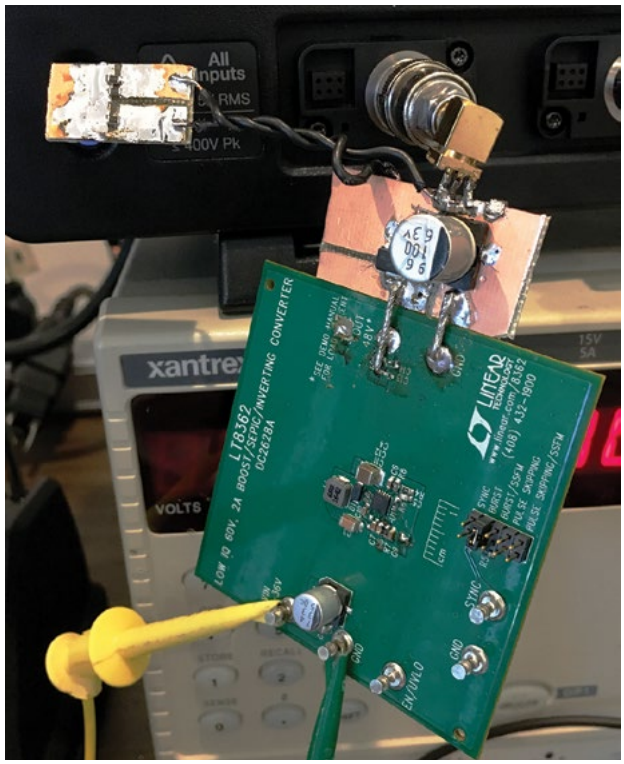

Figure 4. Benchtop setup of a clean phantom power supply using demonstration circuit DC2628.

About the Author

Thomas Mosteller was a field applications engineer for Analog Devices' Mid-Atlantic region starting in 1990. He assisted a wide range of military, commercial, and industrial customers with designs in many fields such as power distribution and supplies, analog signal conditioning, data conversion, and RF and communications applications. Prior to joining Linear Technology (now part of Analog Devices), Thomas designed a wide range of medical equipment for 10 years and holds a patent in the design of infusion pumps. He earned a B.S.E.E. degree from Drexel University in 1977.

About the Author

Christopher Jarboe received his B.S.E.E. from Western Kentucky University in 2005 and joined Linear Technology (now part of Analog Devices) in 2010. He is an applications engineer supporting customers in the Northeast and Mid-Atlantic United States. He can be reached at christopher.jarboe@analog.com.

Timing Challenges in Multiaxis Robotics and Machine Tool Applications

Dara O'Sullivan, System Applications Manager

Introduction

Industrial robotics and machine tool applications involve the precise, coordinated movement of several axes in space in order to accomplish the job at hand. Robots typically have six axes that need to be controlled in a coordinated manner, and sometimes seven if the robot is moving along a rail. In CNC machining, 5-axis coordination is common, although there are applications that utilize up to 12 axes in which tools and workpieces are both being moved with respect to each other in space. Each axis comprises a servo drive, a motor, and sometimes a gearbox between the motor and the axis joint, or end effector. The system is then interconnected over an Industrial Ethernet network, usually in a line topology, as shown in Figure 1. A machine controller converts the required spatial trajectory to individual position references for each servo axis, and these are communicated over the network on a cyclic basis.

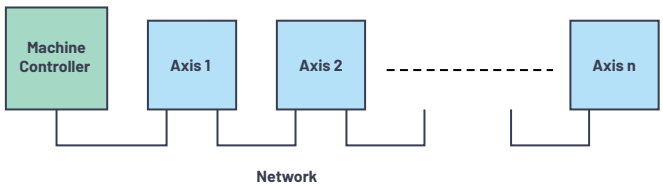

Figure 1. Network topology of a multiaxis machine.

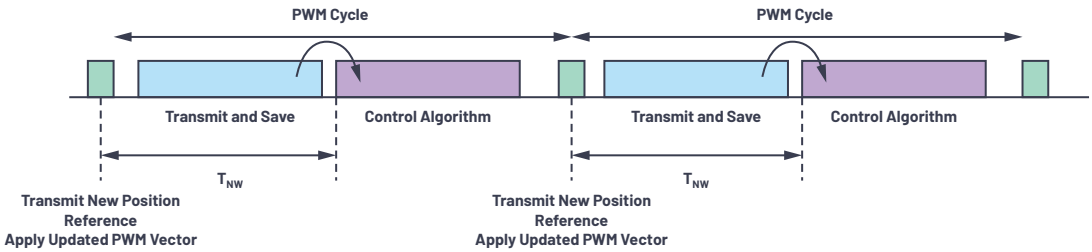

Figure 2. PWM cycle and network transmission time.

The Control Cycle

These applications run on a defined cycle time that is usually equal to, or a multiple of, the fundamental control/pulse-width modulation (PWM) switching cycle of the underlying servomotor drive. End-to-end network transmission latency is a key parameter in this context as illustrated in Figure 2. Within each cycle period, the new position reference and other relevant information must be transmitted from the machine controller to each node of Figure 1. Then there needs to be sufficient time remaining within the PWM cycle for each node to update the servo control algorithm calculation using the new position reference, as well as any new sensor data. Each node then applies the updated PWM vector in the servo drive at the same point in time via a distributed clock mechanism that is Industrial Ethernet protocol dependent. Depending on the control architecture, part of the control loop algorithm may be implemented in the PLC, and it requires sufficient time to be available, having received any relevant sensor information update across the network.

Data Transmission Delays

Assuming that the only traffic on the network is the cyclic data flowing between the machine controller and the servo nodes, the network latency (T_{NW}) is determined by the number of network hops to the furthest node, the network data rate, and the delays encountered in each node. In the context of robotics and machine tools, the propagation delay of the signal along the wire can be neglected, as the cable length is typically relatively short. The dominant delay is the bandwidth delay; that is, the time needed to get the data onto the wire. For a minimum size Ethernet frame (typical for machine tools and robotics control), the bandwidth delay is illustrated in Figure 3 for both 100 Mbps and 1 Gbps bit rates. This is simply the packet size divided by the data rate. A typical data payload for a multiaxis system from controller to servo would consist of a 4-byte speed/position reference update and a 1-byte control word update for each servo, which means a 30-byte payload for a 6-axis robot. Of course, some applications will carry more information in the update and/or will have more axes, in which case packets larger than the minimum size may be needed.

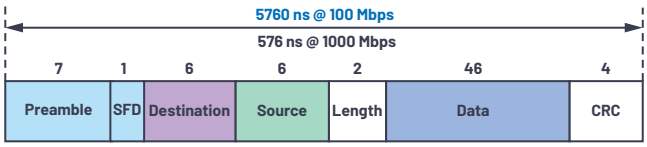

Figure 3. Bandwidth delay of a minimum length Ethernet frame.

Apart from the bandwidth delay, the other delay elements occur as a result of the Ethernet frame passing through the PHYs and 2-port switch at each servo network interface. These delays are depicted in Figure 4 and Figure 5, where the movement of the frame is shown through the PHY into the MAC (1-2), through the destination address analysis where only the preamble and destination parts of the frame must be clocked through. Path 2-3a represents extraction of payload data for the current node, whereas path 2-3b represents the onward journey of the frame to the destination node(s). Figure 4a shows only the payload being passed to the application in 2-3a, whereas Figure 4b shows the bulk of the frame being passed; this is indicative of small differences that can occur between Ethernet protocols. Path 3b-4 represents the outbound transmission of the frame through the transmit queue, through the PHY and back out onto the wire. This path does not exist on a line end node as shown. Cut-through packet switching is assumed here, rather than store-and-forward, which has much higher latency as the entire frame is clocked into the switch before it is forwarded on.

The delay elements of the frame are also shown in Figure 5 along a timeline, where the total frame transmission time through one axis node is illustrated. T_{BW} represents the bandwidth delay, while T_{L_1node} represents the latency of the frame through a single node. Apart from delays related to the physical transmission of the bits over the wire and the clocking in of address bits for destination address analysis, PHY and switch component latencies are the other elements that impact the transmission delays within the system. As the bit rates on the wire increase and the node count expands, these latencies become even more important in the overall end-to-end frame transmission delays.

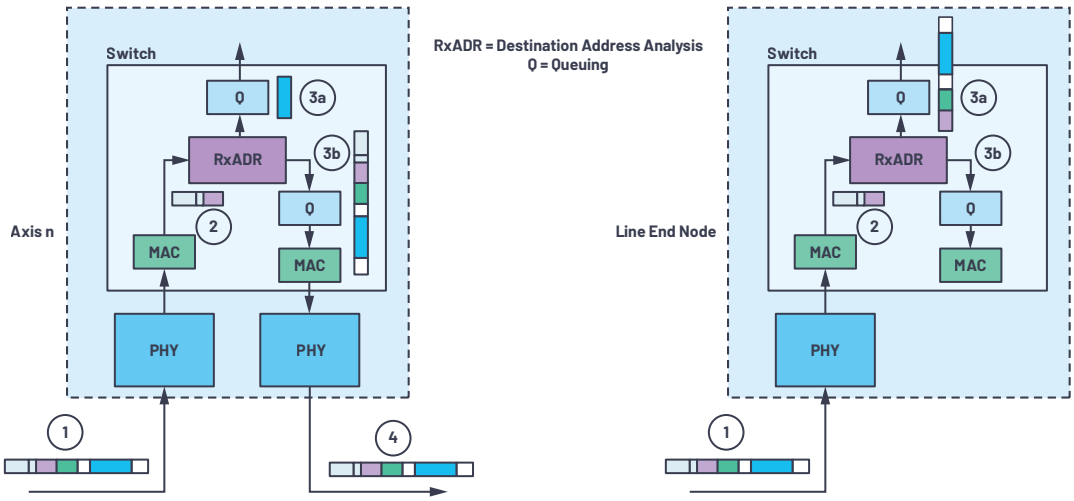

Figure 4. Frame latencies: (a) 2-port node frame latencies and (b) line end node.

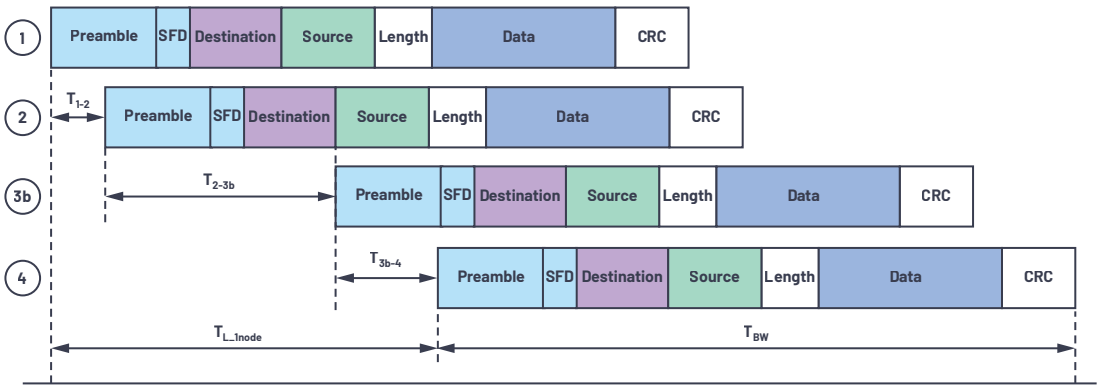

Figure 5. Frame transmission timeline.

Low Latency Solutions

Analog Devices has recently released two new Industrial Ethernet PHYs designed to operate reliably in harsh industrial conditions over extended ambient temperature ranges up to 105°C and with industry-leading power and latency specifications. The ADIN1300 and ADIN1200 were developed specifically to address the challenges outlined in this article and make ideal choices for industrial applications. With the fido5000 real-time Ethernet, multiprotocol, embedded 2-port switch, Analog Devices enables solutions for deterministic time-sensitive applications.

The latencies introduced by the PHY and switch are listed in Table 1, assuming that the receive buffer analysis is destination address based and assuming a 100 Mbps network.

Table 1. PHY and Switch Latencies

Delay Element	Component	Time
PHY Receive	ADIN1200	248 ns
PHY Transmit	ADIN1200	52 ns
Switch Preamble and Destination	fido5000	1120 ns (14 bytes at 100 Mbps)
Switch MAC, Queuing, and Receiver	fido5000	330 ns

As an example, aggregating these delays up to a 7-axis line network, and including the clocking of the full payload into the final node (3a in Figure 4), the total transmission delay becomes

$$6 \times T_{L\_1node} + TBW + Tnode7 = 6 \times (248\ \text{ns} + 330\ \text{ns} + 1120\ \text{ns} + 52\ \text{ns}) + 5760\ \text{ns} + (248\ \text{ns} + 1120\ \text{ns} + 58 \times 80\ \text{ns}) = 22.3\ \mu\text{s} \quad (1)$$

where the 58 × 80 ns represents the remaining 58-byte payload after the preamble and destination address bytes have been read.

This calculation assumes that there is no other traffic on the network or that the network is managed to enable priority access for time sensitive traffic. It is also somewhat protocol dependent, with slight variations in the calculation being introduced depending on the exact Industrial Ethernet protocol used. Referring back to Figure 2, in a machine system with cycle times down to 50 µs to 100 µs, the frame transmission to the furthest node can take up to almost 50% of the cycle, reducing the time available to update the motor control and motion control algorithm calculations for the next cycle. Minimizing this transmission time is important for performance optimization, as it allows longer and more complex control calculations. Given that delays associated with data on the wire are fixed and related to the bit rate, utilizing low latency components, such as the ADIN1200 PHY and the fido5000 embedded switch, is key to performance optimization, especially as node count increases (for example, 12-axis CNC machine) and cycle times reduce. Moving to gigabit Ethernet dramatically reduces the impact of bandwidth delay, but increases the proportion of overall latency introduced by the switch and PHY components. For example, a 12-axis CNC machine on a gigabit network will have a network transmission delay of approximately 7.5 µs. The bandwidth element of this is negligible, and it makes little difference whether minimum or maximum Ethernet frame sizes are used. The network delay is split approximately equally between the PHYs and the switches, underlining the value in minimizing the latency in these elements as industrial systems move toward gigabit speeds, control cycle times reduce (EtherCAT® has demonstrated 12.5 µs cycle times), and node count expands with the addition of Ethernet connected sensors in the control network and the flattening of network topologies.

Conclusion

In high performance, multiaxis, synchronized motion applications, control timing requirements are precise, deterministic, and time critical, with a requirement to minimize end-to-end latency, especially as control cycle times get shorter and control algorithm complexity increases. Low latency PHYs and embedded cut-through switches are important elements in optimizing these systems. To address the challenges outlined in this article, Analog Devices has recently released two new robust Industrial Ethernet PHYs, the ADIN1300 (10 Mb/100 Mb/1 Gb) and ADIN1200 (10 Mb/100 Mb). For more information on ADI's physical layer technology, visit analog.com/ADIN1300 and analog.com/ADIN1200. To learn about ADI's Chronous™ portfolio of Industrial Ethernet solutions and how they are accelerating real-world Industrial Ethernet networks, please visit analog.com/chronous.

About the Author

Dara O'Sullivan 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. from University College Cork, Ireland, and he has worked in industrial and renewable energy applications in a range of research, consultancy, and industry positions since 2001. He can be reached at dara.osullivan@analog.com.

Phased Array Antenna Patterns—Part 1: Linear Array Beam Characteristics and Array Factor

Peter Delos, Technical Lead, **Bob Broughton,** Director of Engineering, and **Jon Kraft,** Senior Staff Field Applications Engineer

Introduction

With the proliferation of digital phased arrays in commercial and aerospace and defense applications, there are many engineers working on various aspects of the design who have limited phased array antenna familiarity. Phased array antenna design is not new, as the theory has been well developed over decades; however, most of the literature is intended for antenna engineers well versed in the electromagnetic mathematics. As phased arrays begin to include more mixed-signal and digital content, there are many engineers who could benefit from a much more intuitive explanation of phased array antenna patterns. As it turns out, there are many analogies between the behavior of phased array antennas and the discrete-time sampled systems that the mixed-signal and digital engineers work with every day.

These articles are not intended to create antenna design engineers, but rather to help the engineer working on a subsystem or component used in a phased array to visualize how their effort may impact a phased array antenna pattern.

Beam Direction

First, let's look at an intuitive example of steering a phased array beam. Figure 1 provides a simple illustration of a wavefront striking four antenna elements from two different directions. A time delay is applied in the receive path after each antenna element, and then all four signals are summed together. In Figure 1a, that time delay matches the time difference of the wavefront striking each element. And in this case, that applied delay causes the four signals to arrive in phase at the point of combination. This coherent combining results in a larger signal at the output of the combiner. In Figure 1b, that same delay is applied; however, in this case, the wavefront is perpendicular to the antenna elements. That applied delay now misaligns the phase of the four signals, and the output of the combiner is significantly reduced.

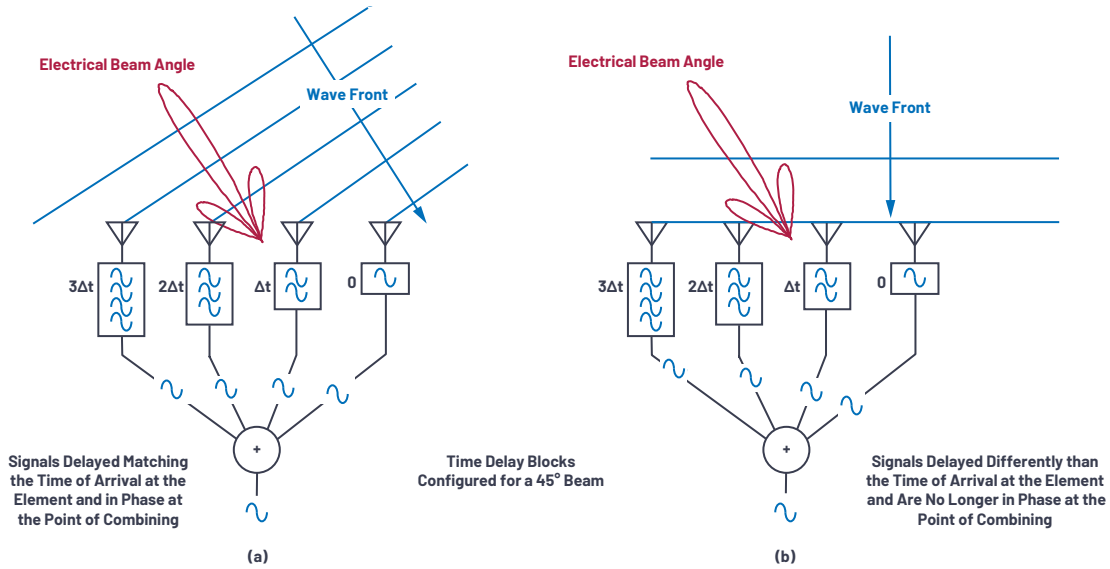

Figure 1. Understanding steering angle.

In a phased array, time delay is the quantifiable delta needed for beam steering. But time delay can also be emulated with a phase shift, which is common and practical in many implementations. We will discuss the impact of time delay vs. phase shift in the section on beam squint, but for now let's look at a phase shift implementation, and then derive the calculation for beam steering with that phase shift.

Figure 2 shows this phased array arrangement using phase shifters rather than time delay. Note that we define the boresight direction ($\theta = 0°$) as perpendicular to the face of the antenna. A positive angle θ is defined to the right of boresight, and a negative angle is defined to the left of boresight.

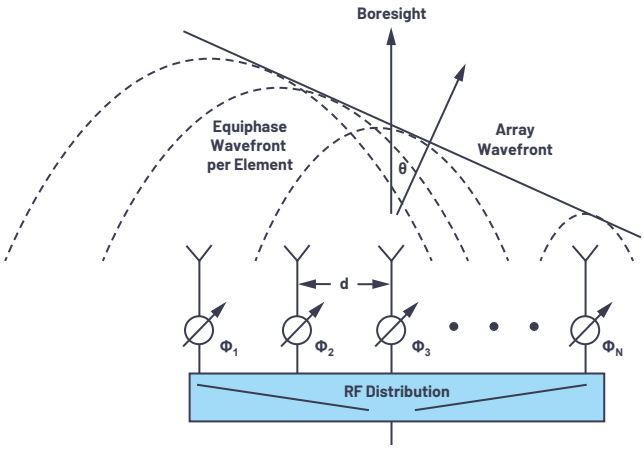

Figure 2. Phased array concept using RF phase shifters.

To visualize the phase shift needed for beam steering, a set of right triangles can be drawn between adjacent elements, as shown in Figure 3, where $\Delta\Phi$ is the phase shift between those adjacent elements.

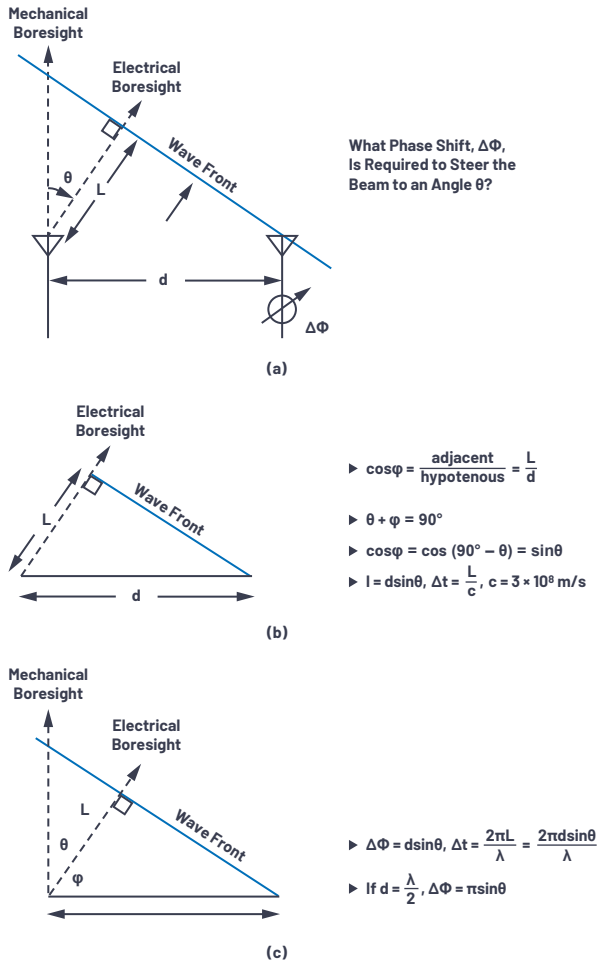

Figure 3. Derivation of phase shift ΔΦ vs. beam steering angle.

Figure 3a defines the trigonometry between those elements, with each element separated by a distance (d). The beam is pointed in a direction off boresight, θ, which is an angle, φ, from the horizon. In Figure 3b, we see that the sum of $\theta + \varphi = 90°$. This allows us to compute L, the delta distance of wave propagation, as $L = d\sin(\theta)$. The time delay to steer our beam is equal to the time it will take for the wavefront to traverse that distance, L. If we think of L as a fraction of the wavelength, then a phase delay could be substituted in for that time delay. The equations for $\Delta\Phi$ can then be defined relative to θ, as shown in Figure 3c and repeated in Equation 1.

$$\Delta\Phi = \frac{2\pi d \sin\theta}{\lambda} \quad (1)$$

If the spacing between elements is exactly one half of the signal wavelength, then this can further be simplified to:

$$\Delta\Phi = \pi \sin\theta, \text{ for } d = \frac{\lambda}{2} \quad (2)$$

Let's work out an example with these equations. Consider two antenna elements spaced 15 mm apart. If a 10.6 GHz wavefront is arriving at 30° from mechanical boresight, then what is the optimal phase shift between the two elements?

- θ = 30° = 0.52 rad
- λ = c/f = (3 × 108 m/s)/10.6 GHz = 0.0283 m
- ΔΦ = (2π × d × sinθ)/λ = 2π × 0.015 × sin(0.52)/0.0283 m = 1.67 rad = 95°

So, if our wavefront is arriving at θ = 30°, then if we shift the phase of the neighboring element by 95°, we will cause the individual signals of both elements to add coherently. This will maximize the antenna gain in that direction.

For a better appreciation of how the phase shift varies with the beam direction (θ), these equations are plotted for a variety of conditions in Figure 4. Some interesting observations can be made from these graphs. For the case of d = λ/2, there is an approximate 3 to 1 slope near boresight, which is the π multiplier in Equation 2. This case also shows a full 180° shift between elements provides a theoretical 90° shift in beam direction. In practice, with real element patterns, this is not realizable, yet the equations do show the theoretical ideal. Note that for d > λ/2, no amount of phase shift provides a full beam shift. Later, we will see this case can lead to grating lobes in the antenna pattern, and this graph provides a first indicator that something is different with the d > λ/2 case.

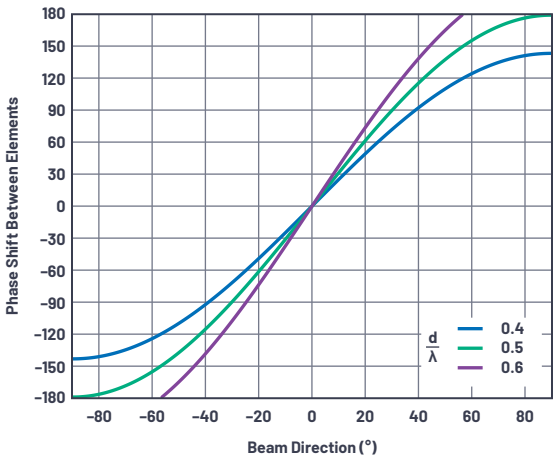

Figure 4. Phase shift ΔΦ between elements vs. beam direction (θ) for three cases of d/λ.

A Uniformly Spaced Linear Array

The equations developed above have applied to just two elements. Yet a real phased array can be thousands of elements spaced across two dimensions. But for our purposes here, let's just consider one dimension: a linear array.

A linear array is a single element wide with N number of elements across. The spacing may vary, but often it is uniform. Therefore, in this paper, we will set the spacing between each element to a uniform distance, d, (Figure 5). Although simplified, this uniformly spaced linear array model provides the foundation for insight into how the antenna pattern is formed vs. a variety of conditions. We can further apply the principles of the linear array to understand two-dimensional arrays.

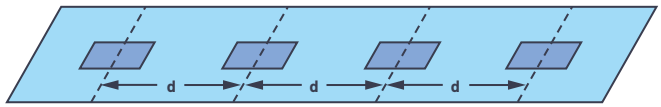

Figure 5. A uniformly spaced linear array (N = 4).

Near Field vs. Far Field

So how can we take the equations previously developed for an N = 2 linear array and apply them to an N = 10,000 linear array? Right now, it seems that each antenna element has a slightly different angle pointing to the spherical wavefront, shown in Figure 6.

With the RF source near, the incident angle varies for each element. This situation is called the near field. We can work out all these angles, and sometimes we need to do this for antenna testing and calibration as our test setup can only be so large. But, if we instead just assumed that the RF source was far away, then we'd have the case in Figure 7.

With the RF source far away, the large radius of the spherical wavefront results in wave propagation paths that are approximately parallel. Therefore, all our beam angles are equal, and each adjacent element has a path length that is L = d × sin*θ)* longer than its neighbor. This simplifies the math and means that the two element equations we derived can be applied to thousands of elements, provided they have uniform spacing.

But when can we make the far field assumption? How far is far? It's a little subjective, but in general, far field is considered anything greater than:

$$\text{Far Field} > 2\,D^2/\lambda \quad (3)$$

where D is the diameter of the antenna ((N-1) × d for our uniform linear array)

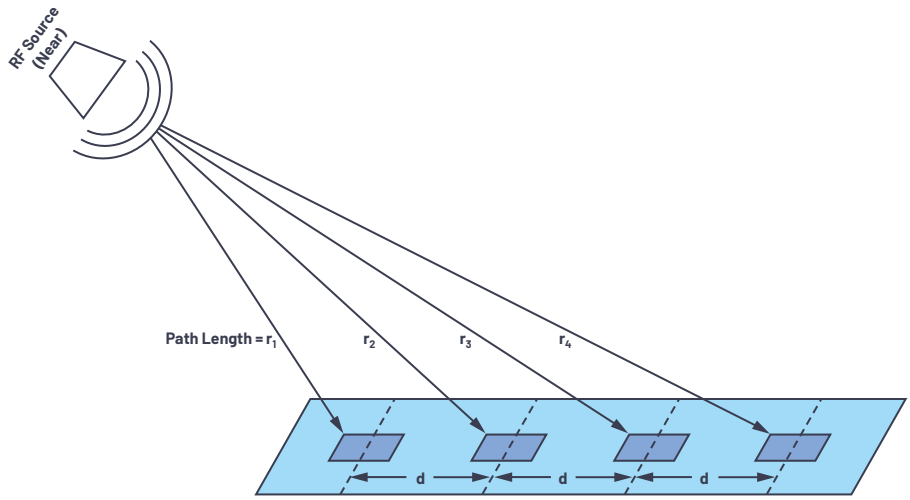

Figure 6. An RF source near the linear array.

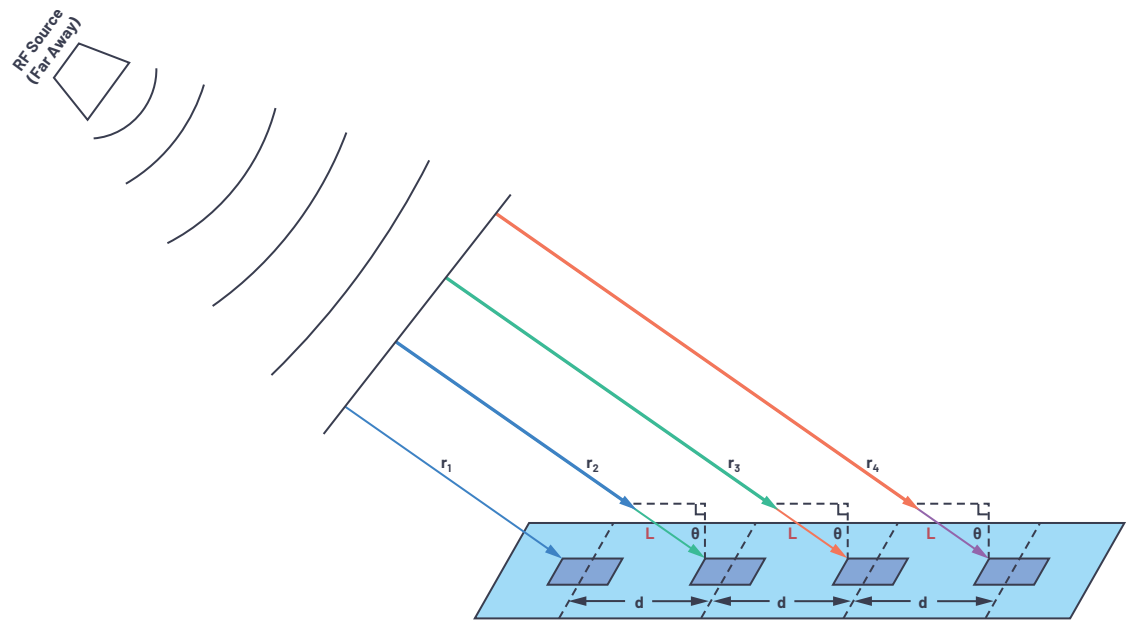

Figure 7. An RF source far from the linear array.

For a small array (small D) or a low frequency (large λ), the far field distance is small. But for a large array (or high frequency), the far field distance could be many kilometers! That makes it hard to test and calibrate the array. For those conditions, a more detailed near model may be used, and then bridge this back to the far field, real-world use of the array.

Antenna Gain, Directivity, and Aperture

Before we go too far, it is helpful to define antenna gain, directivity, and aperture. Let's start with a clarification on gain vs. directivity, as the two are often interchanged. Antenna gain and directivity can be compared to an isotropic antenna—which is an ideal antenna that radiates evenly in all directions. Directivity is a comparison of the maximum power measured, $P_{max,}$ in a particular direction to the average power radiated across all directions, P_{av}. When no direction is defined, directivity is determined by Equation 4.

$$D = \frac{P_{max}}{P_{av}} \tag{4}$$

Directivity is a useful metric when comparing antennas as it defines the ability to focus radiated energy. Gain has the same pattern of directivity, but gain includes the antenna losses.

$$\text{Gain} = G = kD, \text{ where } k = \frac{P_{rad}}{P_{in}} \tag{5}$$

P_{rad} is the total power radiated, P_{in} is the input power to the antenna, and k accounts for losses in the antenna radiation process.

Next, let's consider an antenna pattern as a function of a three-dimensional direction and directivity as a function of beamwidth.

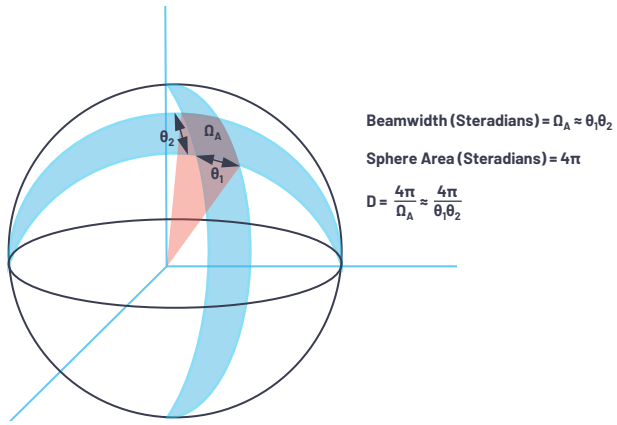

Figure 8. A three-dimensional view of an area projected onto a sphere.

The total surface area of a sphere is $4\pi^2$, and an area on a sphere is defined in units of steradians with 4π steradians in a sphere. Therefore, the power density from an isotropic radiator is

$$\frac{P_{rad}}{4\pi r^2} \tag{6}$$

and has units of (W/m^2).

There are two angular directions for an area of a sphere. In radar systems, these are commonly referred to as azimuth and elevation. Beamwidths can be described as a function of each angular direction as θ_1 and θ_2: the combination creates an area on the sphere of Ω_A.

Ω_A is the beamwidth in steradians and can be approximated as $\Omega_A \approx \theta_1 \times \theta_2$.

Recognizing Ω_A as an area on the sphere, directivity can then be expressed as

$$D = \frac{4\pi}{\Omega_A} \approx \frac{4\pi}{\theta_1\theta_2} \tag{7}$$

The third antenna term we'll consider is aperture. Antenna aperture represents an effective area for receiving electromagnetic waves and includes a function relative to wavelength. The aperture of an isotropic antenna is

$$A\_isotropic = \frac{\lambda^2}{4\pi} \tag{8}$$

Gain is relative to isotropic radiation, making the effective aperture of an antenna

$$A_e = \frac{G\lambda^2}{4\pi} \tag{9}$$

Pulling these three terms together, we can see that gain can be considered a function of angle that defines a radiation pattern and accounts for efficiency (or losses) in the antenna.

Array Factor for a Linear Array

At this point, we're able to predict the optimal time (or phase) delta between elements to achieve maximum antenna directivity. But we'd really like to understand and manipulate the complete antenna gain pattern. There are two main parts to this. First, there is the gain of each individual element of our array (perhaps one patch), called the element factor (G_E). Secondly, there is the impact that we can exert through beamforming the array, called the array factor (G_A). The full array antenna gain pattern is the combination of the two factors, as shown in Equation 10.

$$G(\theta) = G_E(\theta) + G_A(\theta), \text{ in dB} \tag{10}$$

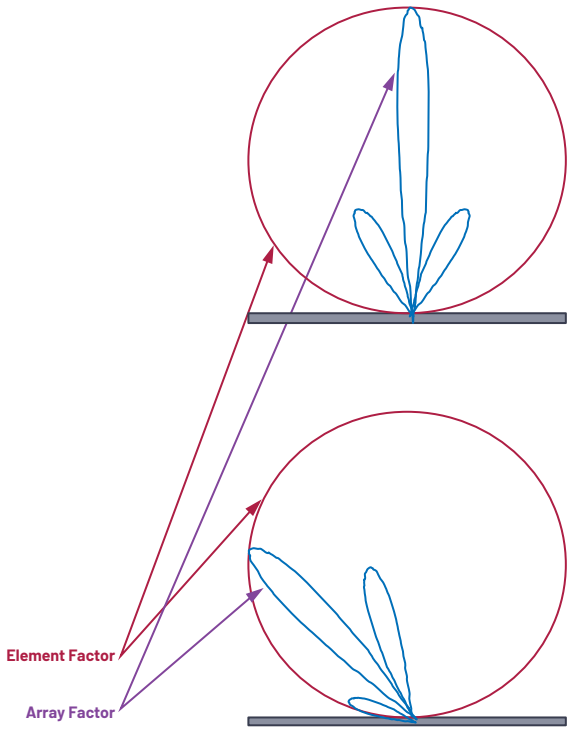

Figure 9. Element and array factor.

The element factor, G_E, is the radiating pattern of a single element in the array. This is defined by the geometry and construction of the antenna, and not something that is varied in operation. It is important to know, as it will limit the gain of our total array—particularly near the horizon. But since we can't control it electrically, we'll leave it as a fixed influencer to our total phased array gain equation. For this article, we'll assume that all the individual elements have the same element factor.

The focus then is on the array factor, G_A. The array factor is calculated based on array geometry (d for our uniform linear array) and beam weights (amplitude and phase). Deriving the array factor for a uniform linear array is straightforward, but the details are best covered in the references cited at the end of this article.

There are some variations in equations used across literature depending on how parameters were defined in the linear array. We use the equations from this article, which results in consistency with our definitions in Figure 2 and Figure 3. Since our primary concern is how the gain changes, it is often more instructive to plot the normalized array factor relative to unity gain. That normalized array factor can be written as Equation 11.

$$AF[\theta] = \frac{sin\left(\frac{N\pi d}{\lambda}\left[sin(\theta) - sin\left(\theta_0\right)\right]\right)}{Nsin\left(\frac{\pi d}{\lambda}\left[sin(\theta) - sin\left(\theta_0\right)\right]\right)} \tag{11}$$

$$\theta_0 = beam\ angle$$

We have already defined beam angle θ_0 as a function of phase shift between elements $\Delta\Phi$; therefore, we can also write the normalized antenna factor as Equation 12.

$$AF[\theta, \Delta\Phi] = \frac{sin\left(N\left[\frac{\pi d}{\lambda}\,sin(\theta) - \frac{\Delta\Phi}{2}\right]\right)}{Nsin\left(\frac{\pi d}{\lambda}\,sin(\theta) - \frac{\Delta\Phi}{2}\right)} \tag{12}$$

The conditions assumed in the array factor equation include:

- The elements are equally spaced.
- There is an equal phase shift between elements.
- The elements are all at equal amplitude.

Next, using these equations, we plot the array factor for several array sizes.

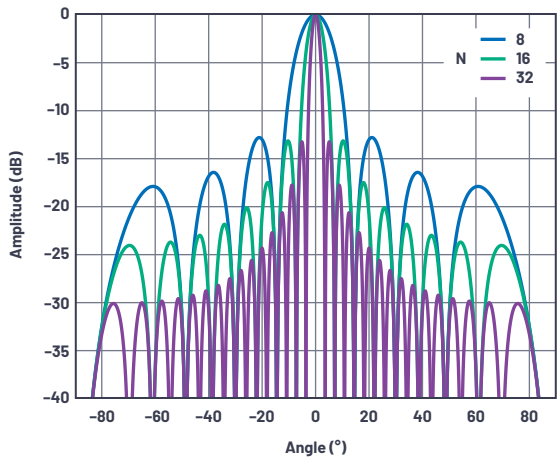

Figure 10. Normalized array factor at boresight of a linear array with an element spacing of d = λ/2 and an element count of 8, 16, and 32.

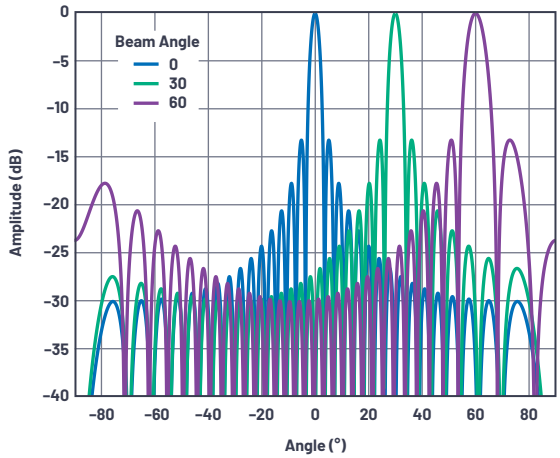

Figure 11. Normalized array factor of a 32-element linear array at several beam angles with an element spacing of d = λ/2.

Some observations from these figures include:

- The first sidelobe is at –13 dBc regardless of the element count. This is due to the sinc function in the array factor equation. The sidelobes can be improved with tapering the gain across elements and will be the subject of an upcoming section in this series.
- The beamwidth reduces with the number of elements.
- The beamwidth widens as the beam is scanned away from boresight.
- The number of nulls increases as the number of elements increases.

Beamwidth

Beamwidth provides a metric of angular resolution for antennas. Most commonly, beamwidth is defined by either the half-power beamwidth (HPBW) or the null-to-null spacing of the main lobe (FNBW). To find the HPBW, we move 3 dB down from the peak and measure the angular distance, as shown in Figure 12.

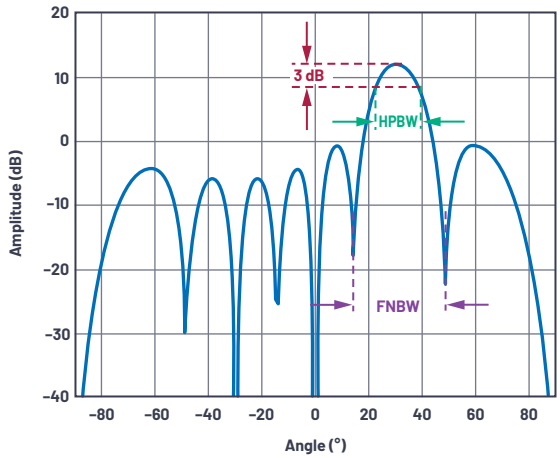

Figure 12. Definition of antenna beamwidth (linear array of N = 8, d = λ/2, θ = 30° shown).

Using our normalized array factor equation, we can solve for this HPBW by setting Equation 3 equal to the half-power level (3 dB or $1/\sqrt{2}$). We'll assume mechanical boresight (θ = 0°), N = 8, and d = λ/2.

$$1/\sqrt{2} = \frac{sin\left(8\left[\frac{\pi\lambda}{2\lambda}\,sin(\theta) - \frac{\Delta\Phi}{2}\right]\right)}{8\,sin\left(\frac{\pi\lambda}{2\lambda}\,sin(\theta) - \frac{\Delta\Phi}{2}\right)} \tag{13}$$

Then solving for ΔΦ gives 0.35 rad. Use Equation 1, and solve for θ:

$$0.35 = \frac{2\pi\lambda\,sin\theta}{2\lambda} \rightarrow \theta = 0.11 \text{ rad} = 6.4° \tag{14}$$

That θ is the peak to 3 dB point, which is half of our HPBW. Therefore, we simply double it to arrive at the angular distance between the 3 dB points. This results in an HPBW of 12.8°.

We could repeat this for an array factor equal to 0 and obtain the first null-to-null spacing angle of FNBW = 28.5°, for the previously mentioned conditions.

For uniform linear arrays, an approximation for HPBW [1,2] is given in Equation 15.

$$\theta_B \sim \frac{0.886\lambda}{Ndcos\theta} \tag{15}$$

Figure 13 plots beamwidth vs. beam angle for several element counts in the condition of a λ/2 element spacing.

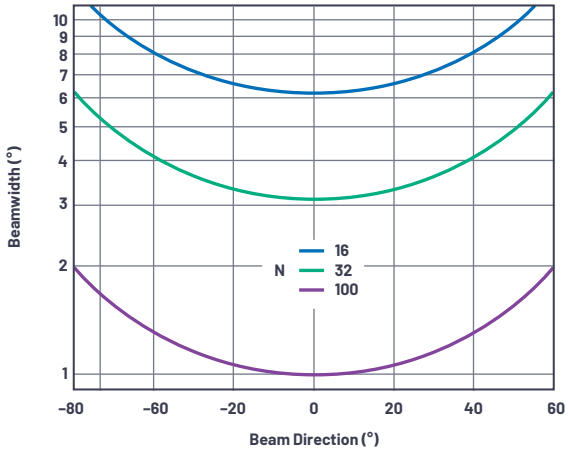

Figure 13. Beamwidth vs. beam angle at an element spacing of λ/2 for an element count of 16, 32, and 100 elements.

From this graph, it is worth noting some observations relative to array sizes under development in the industry.

- A 1° beam accuracy requires 100 elements. If this is desired in both azimuth and elevation, this results in a 10,000 element array. The 1° accuracy is only at boresight under near ideal conditions. Maintaining 1° accuracy in a fielded array across a variety of scan angles will increase the element count further. This observation then sets a practical limit for beamwidth with very large arrays.
- A 1000-element array is common in the industry. 32 elements in each direction provides an element count of 1024, and can yield a beam accuracy less than 4° near boresight.
- A 256-element array, which can be mass produced at low cost, can still have a beam pointing accuracy less than 10°. This may be perfectly acceptable for many applications.
- Also note that for any of these cases, the beamwidth doubles at 60° offsets. This is from the cosθ in the denominator and is due to the foreshortening of the array; that is, the array appears to be a smaller cross section when viewed from an angle.

Combining Element and Array Factors

The previous section only considered the array factor. But to find the total antenna gain, we also require the element factor. Figure 14 illustrates an example. In this example, we use a simple cosine shape as the element factor, or normalized element gain, $G_E(\theta)$. The cosine roll-off is common in phased array analysis and can be visualized if considering a flat surface. At broadside, there is a maximum area. As the angle moves away from broadside, the area visible reduces following a cosine function.

The array factor, $G_A(\theta)$, was used for a 16-element linear array, with a $\lambda/2$ spacing, and a uniform radiation pattern. The total pattern is a linear multiplication of the element factor and array factor, so in a dB scale, they can be added together.

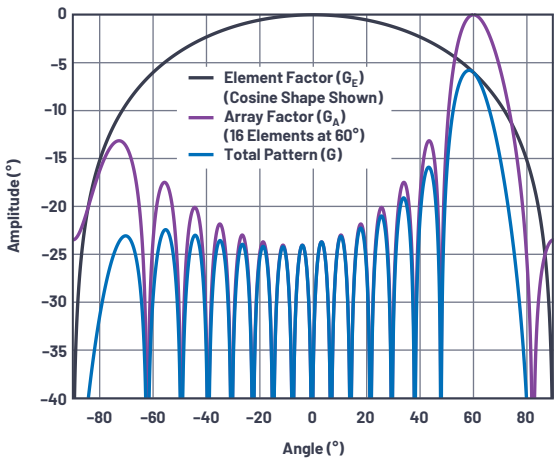

Figure 14. Element factor and array factor combine to form the total antenna pattern.

A few observations as the beam moves off boresight:

- The main beam loses amplitude at the rate of the element factor.
- The sidelobes on boresight have no amplitude loss.
- The result is the sidelobe performance of the overall array degraded off boresight.

Antenna Plots: Cartesian vs. Polar

The antenna pattern plots used so far have been in Cartesian coordinates. But it is common to plot antenna patterns in polar coordinates as they are more representative of energy radiating spatially outwards from the antenna. Figure 15 is a redrawn version of Figure 12, but using polar coordinates. Note that this is the exact same data, point for point—it's just redrawn with a polar coordinate system. It is worthwhile to be able to visualize the antenna pattern in either representation as both are used in literature. For most of this text, we will use cartesian coordinates as in this representation it can be easier to compare beamwidth and sidelobe performance.

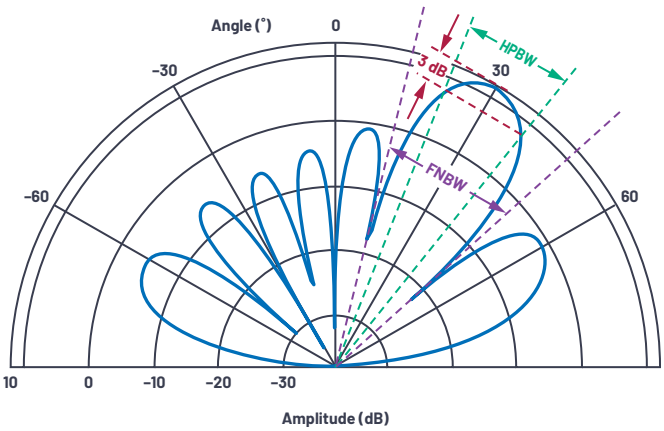

Figure 15. Polar coordinate antenna directivity plot for N = 8, d = λ/2, θ = 30°.

Array Reciprocity

Up until this point, all the diagrams and text have described a signal that the array is receiving. But how would this change for a transmit array? Fortunately, most antenna arrays are reciprocal. Therefore all of the diagrams, equations, and terminology are the same for transmit as they are for receive. Sometimes it is easier to think of the beam as being received by the array. And sometimes, perhaps in the case of grating lobes, you may find it more intuitive to think of the array as transmitting a beam. In this article, we generally describe the array as receiving a signal. But if this is harder to visualize for you, then you can equally think of the same concepts on the transmit side.

Summary

This concludes Part 1 of the series. The concept of beam steering with a phased array was introduced. The equations to calculate phase shift across the array for beam steering were derived and shown graphically. Then array factor and element factor were defined with observations of how the number of elements, the spacing between elements, and the beam angle impacts the antenna response. Finally, a comparison of antenna patterns in cartesian vs. polar coordinates was shown.

Upcoming articles in this series will further explore phased array antenna patterns and impairments. We'll study how antenna tapering reduces sidelobes, how grating lobes are formed, and the impact of phase shift vs. time delay in wideband systems. The series will finish with an analysis of the finite resolution of the delay block and how it can create quantization sidelobes and degrade beam resolution.

References

Balanis, Constantine A. *Antenna Theory: Analysis and Design*, third edition. Wiley-Interscience, 2005.

Mailloux, Robert J. *Phased Array Antenna Handbook*, second edition, Artech House, 2005.

O'Donnell, Robert M. "Radar Systems Engineering: Introduction." IEEE, June 2012.

Skolnik, Merrill. *Radar Handbook*, third edition, McGraw-Hill, 2008.

About the Author

Peter Delos is a technical lead in the Aerospace and Defense Group at Analog Devices in Greensboro, NC. He received his B.S.E.E. from Virginia Tech in 1990 and M.S.E.E. from NJIT in 2004. Peter has over 25 years of industry experience. Most of his career has been spent designing advanced RF/analog systems at the architecture level, PWB level, and IC level. He is currently focused on miniaturizing high performance receiver, waveform generator, and synthesizer designs for phased array applications. He can be reached at peter.delos@analog.com.

About the Author

Bob Broughton started at Analog Devices in 1993 and has held positions as a product engineer and an IC design engineer, and is currently the director of engineering in the Aerospace and Defense Business Unit. Prior to ADI, Bob worked at Raytheon as an RF design engineer and at Peregrine Semiconductor as an RFIC designer. Bob graduated with a B.S.E.E. from West Virginia University in 1984. He can be reached at bob.broughton@analog.com.

About the Author

Jon Kraft is a senior staff FAE in Colorado and has been with ADI for 13 years. His focus is software-defined radio and aerospace phased array radar. He received his B.S.E.E. from Rose-Hulman and his M.S.E.E. from Arizona State University. He has nine patents issued, six with ADI, and one currently pending. He can be reached at jon.kraft@analog.com.

Add an Efficient Positive Rail to a Unipolar Negative Supply

Victor Khasiev, Senior Applications Engineer

Introduction

Sometimes you need a positive power supply, and the most available rail (or only available rail) is negative. In fact, negative-to-positive voltage conversion is used in automotive electronics, and the biasing circuitry for a variety of audio amplifiers, and industrial and test equipment. Even though power is distributed in many of these systems via a negative—relative to ground—rail, the logic boards, ADCs, DACs, sensors, and similar devices found in them still require one or more positive rails. This article presents a simple, low component count, efficient circuit for the generation of a positive voltage from a negative rail.

Circuit Description and Power Train Functionality

Figure 1 shows a complete solution for efficient conversion of a negative voltage to a positive voltage. This particular solution uses a boost topology. The power train includes switching MOSFETs, bottom Q1, top Q2, inductor L1, and input/output filters. The synchronous high efficiency boost controller IC regulates output voltage by changing the state of the switching MOSFETs in the power train. For the purposes of describing this circuit, system ground (SYS_GND) serves as the reference regarding polarity, with a negative—relative to SYS_GND—input rail ($-V_{IN}$), and positive—relative to SYS_GND—output rail ($+V_{OUT}$).

The converter works as follows. If transistor Q1 is on, then current flows from the SYS_GND to the negative rail. The transistor Q2 is off and inductor L1 stores energy in its magnetic field. To complete the switching period, Q1 turns off, and Q2 turns on. Current starts to flow from the SYS_GND to the $+V_{OUT}$ rail, discharging L1 energy to the load.

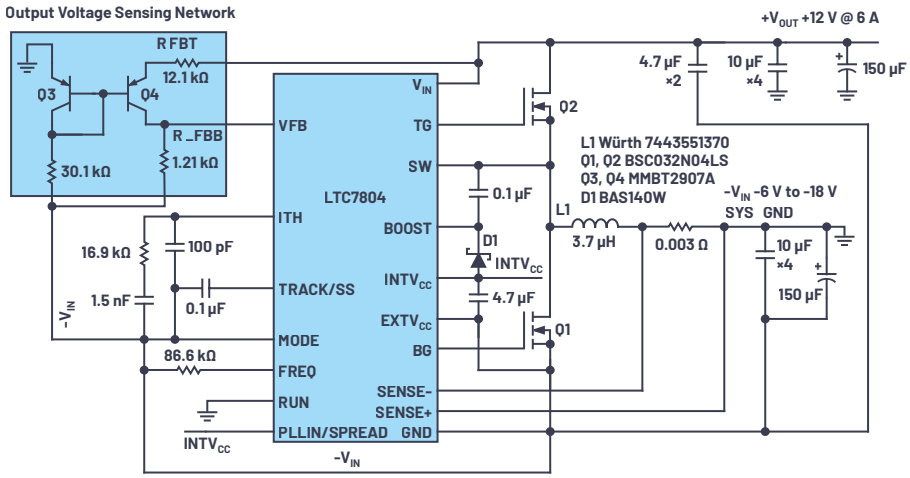

Figure 1. A negative-to-positive converter electrical schematic, with V_{IN} –6 V to –18 V (–24 V peak) and with V_{OUT} +12 V at 6 A.

Basic Expressions for the Power Train Components Selection

The topological diagrams in Figure 2 of the switching behavior illustrate the negative-to-positive converter behavior. For the first interval of a switching cycle, over a length of time defined by the duty cycle, the bottom switch, B_{SW}, is shorted and the top switch, T_{SW}, is open. The voltage across the inductor, L, is equal to $-V_{IN}$. Throughout this interval, current in inductor L increases, generating a voltage polarity matching $-V_{IN}$ across the inductor. At the same time, the output filter capacitor discharges, supplying current to the system load.

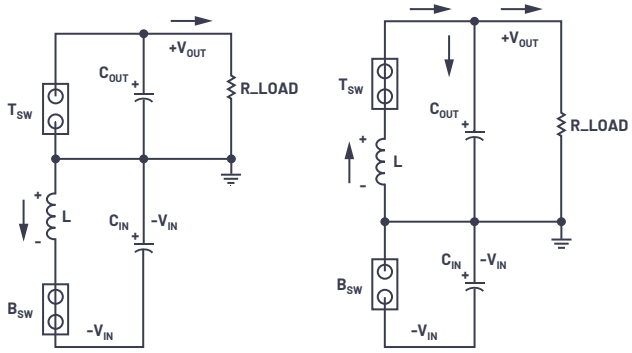

Figure 2. Negative-to-positive converter topological diagrams.

The second interval of the cycle flips both switches—B_{SW} is open and T_{SW} is shorted. The polarity across inductor L changes, and the inductor starts sourcing current (stored in the first interval of the cycle) to both the load and C_{OUT}, the output filter capacitor. The inductor sees a corresponding decrease in current over this part of the cycle. The volt-second balance of the inductor defines the duty cycle, D, of the converter in continuous conduction mode.

Calculating Timing and Component Stresses

Here are the formulas describing the timing and stresses of the power train components.

The duty cycle determines the on/off time of the switches

$$D = \frac{V_{OUT}}{(|V_{IN}| + V_{OUT})}$$

The average value of the input current, I_{OUT}, is input current

$$I_{IN} = \frac{I_{OUT}}{(1 - D)}$$

The peak value of the inductor current

$$I_L = I_{IN} + \frac{(dI/dt)}{2}$$

The voltage stress on the switching MOSFET

$$V_{DS} = |V_{IN}| + V_{OUT}$$

The average current through the bottom MOSFET

$$I_{Q1} = I_{IN} \times D$$

The average current through the top MOSFET

$$I_{Q2} = I_{OUT}$$

These expressions are useful for a general understanding of the functionality of the topology and for preliminary selection of the power train components. For final selection and detailed design, please use LTspice® modeling and simulation.[1]

Converter Control Description and Functionality

Sensing of the output voltage and level shifting of the control voltage are managed by the current mirror based on the PNP transistors Q3 and Q4. The feedback current I_{FB} (1 mA in this circuit) determines the value of the resistors in the feedback loop.

$$R_{FB(B)} = \frac{V_C}{I_{FB}}$$

where V_C is the reference voltage of error amplifier.

$$R_{FB(T)} = \frac{V_{OUT}}{I_{FB}}$$

where $R_{FB(T)}$ is the output voltage sensing resistor.

The feedback circuit presented in Figure 1 is an inexpensive solution, but tolerance of discrete transistors can be affected by the differences in base emitter voltage and temperature variations. To improve accuracy, a matched pair transistor can be used.

Control of the converter power train is left to the LTC7804 boost controller. This chip was selected due to its high efficiency by means of synchronous rectification, easy implementation, high switching frequency operation (if a small inductor size is desired), and low quiescent current.

Test Results and Topology Limitations

This solution was meticulously tested and verified. Figure 3 shows that efficiency remains high over a wide range of load currents—reaching 96%. Note that as the absolute value of the input voltage decreases, the input and inductor current increases. At a certain point, the inductor current can exceed the maximum, or saturation current, on the inductor. The derating curve showing this effect is shown in Figure 4. The maximum load current is 6 A in the range from –9 V to –18 V, falling below that for input voltages with absolute values below –9 V. Thermal performance is shown in Figure 5 for the solution board in Figure 6.

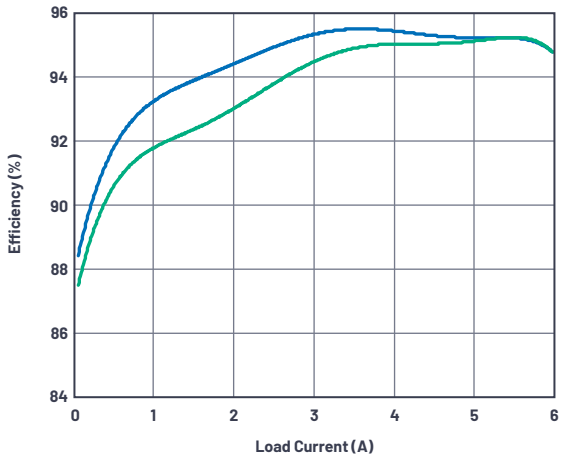

Figure 3. Efficiency curve for V_{IN} –12 V and –18 V with natural convection cooling.

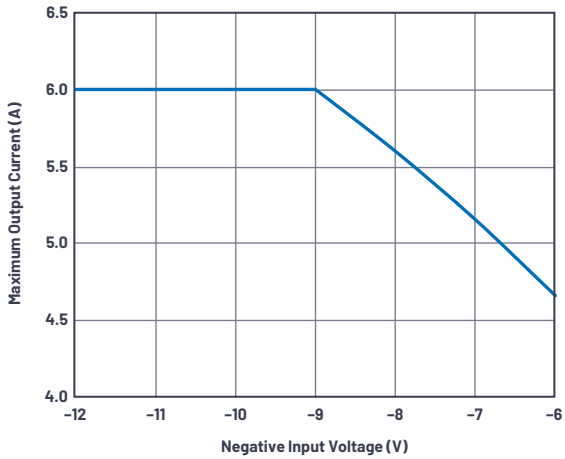

Figure 4. Output current derating curve for absolute value input voltages below –9 V.

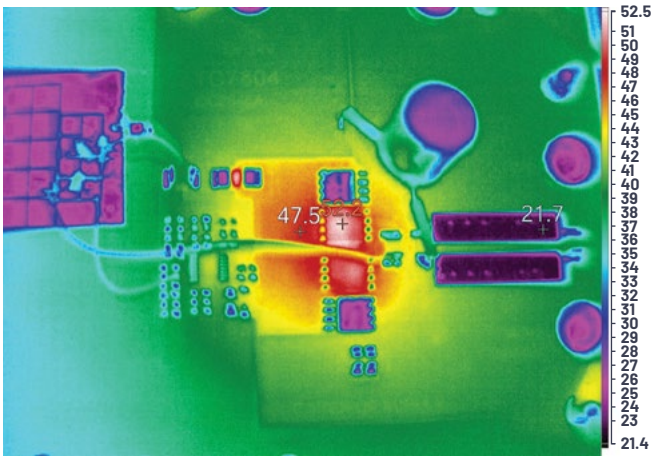

Figure 5. Thermal image of the converter with V_{IN} –12 V and V_{OUT} +12 V at 6 A, using natural convection cooling with no air flow.

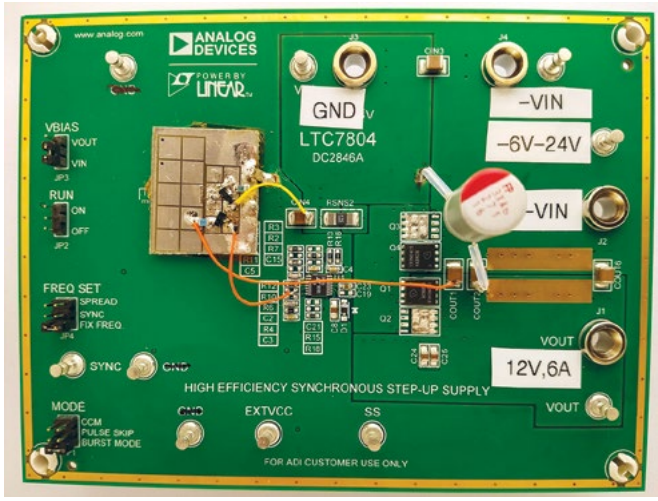

Figure 6. Converter photo.

Conclusion

This article presents a complete solution for a very efficient and relatively simple design for adding a positive rail to a unipolar negative power supply using a boost controller. It also provides electrical schematics and calculations for timing, power conversion components, and electrical stress. Test data confirms high efficiency and good thermal performance. Furthermore, the boost topology used in this solution gives the designer the option to use a prequalified boost controller, saving development time and cost. Conversely, qualifying a boost controller for a negative-to-positive converter can prequalify it for future boost applications.

About the Author

Victor Khasiev is a senior applications engineer at ADI with extensive experience in power electronics both in ac-to-dc and dc-to-dc conversion. He holds two patents and wrote multiple articles. These article 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 customer by answering questions about ADI products, designing and verifying power supplies schematics, laying out printed circuit boards, and troubleshooting and participating in testing final systems. He can be reached at victor.khasiev@analog.com.

RAQ Issue 177: RMS Power vs. Average Power

Doug Ito, Applications Engineer

Question:

Should I use units of root mean square (rms) power to specify or describe the ac power associated with my signal, system, or device?

Answer:

It depends on how you define rms power.

You do not want to calculate the rms value of the ac power waveform. This produces a result that is not physically meaningful.

You do use the rms values of voltage and/or current to calculate average power, which does produce meaningful results.

Discussion

How much power is dissipated when a 1 V rms sinusoidal voltage is placed across a 1 Ω resistor?

$$P = \frac{V^2}{R} = \frac{1^2}{1} = 1 \text{ W} \quad (1)$$

This is well understood[1] and there is no controversy here.

Now, let's see how this compares with the value from an rms power calculation.

Figure 1 shows a graph of a 1 V rms sinusoid. The peak-to-peak value is 1 V rms × 2 $\sqrt{2}$ = 2.828 V, swinging from +1.414 V to –1.414 V.[2]

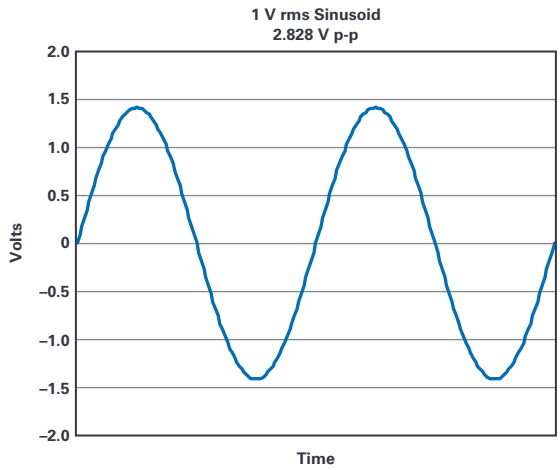

Figure 1. Graph of a 1 V rms sinusoid.

Figure 2 is a graph of the power dissipated by this 1 V rms sinusoid across a 1 Ω resistor ($P = V^2/R$) that shows:

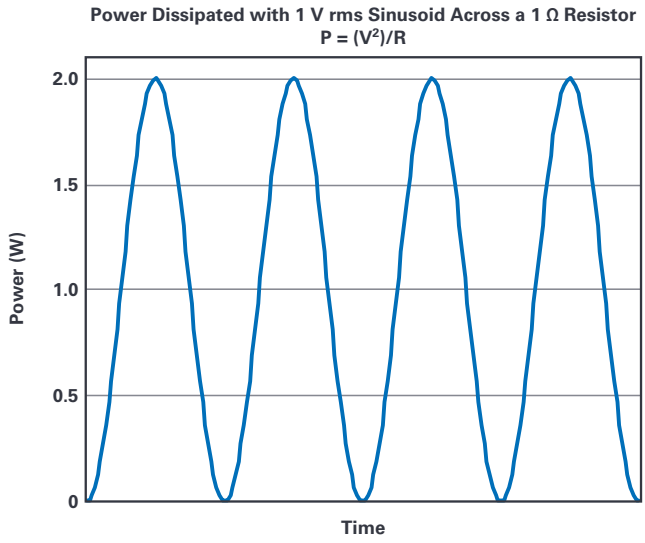

Figure 2. Graph of the power dissipated by a 1 V rms sinusoid across a 1 Ω resistor.

- The instantaneous power curve has an offset of 1 W and swings from 0 W to 2 W.
- The rms value of this power waveform is 1.225 W.
  - One method of calculating this number is Equation 2[3]:

$$Root\ Mean\ Square\ Value = \sqrt{Offset^2 + (AC\_rms)^2} = \sqrt{1\ \mathrm{W}^2 + (2\ \mathrm{W}/2\sqrt{2})^2} = \sqrt{1 + 0.5} \approx 1.225\ \mathrm{W} \quad (2)$$

  - This can be verified by using a more detailed formula[4] in MATLAB® or Excel.
- The average value of this power waveform is 1 W. This is obvious from inspection; the waveform swings symmetrically above and below 1 V. The same value results from calculating the numerical average of the waveform datapoints.
- The average power value matches the power calculated using rms voltage.

The power dissipated by a sinusoidal 1 V rms across a 1 Ω resistor is 1 W, not 1.225 W. Thus, it is the average power that produces the correct value, and thus it is average power that has physical significance. The rms power (as defined here) has no obvious useful meaning (no obvious physical/electrical significance), other than being a quantity that can be calculated as an exercise.

It is a trivial exercise to perform the same analysis using a 1 A rms sinusoidal current through a 1 Ω resistor. The result is the same.

Power supplies for integrated circuits (ICs) are generally dc, so rms power is not an issue for IC power. For dc, average and rms are the same value as dc. The importance of using average power, as opposed to rms power as defined in this document, applies to power associated with time-varying voltage and current—that is, noise, RF signals, and oscillators.

Use rms voltage and/or rms current to calculate average power, resulting in meaningful power values.

[1] Power dissipated from voltage across a resistor is a fundamental relation that is easily derived from Ohm's law (V = IR) and the fundamental definitions of voltage (energy/unit of charge) and current (unit of charge/time). Voltage × current = energy/time = power

[2] The peak-to-peak amplitude of a sinusoid is the rms value multiplied by $2\sqrt{2}$. For a sinusoidal voltage, V p-p = V rms × $2\sqrt{2}$, where V p-p is the peak-to-peak voltage and V rms is the rms voltage. This is a well-known relation that is documented in countless text books, as well as here: en.wikipedia.org/wiki/Root_mean_square.

[3] This is adapted from the rms value calculated from a constant dc offset value plus a separate rms ac value and in the application note "Make Better AC RMS Measurements with Your Digital Multimeter" by Keysight.

[4] The standard textbook definition is one example of a more detailed formula.

$$x_{RMS} = \sqrt{\frac{1}{n}\left(x_1^2 + x_1^2 + \ldots x_n^2\right)} \quad (3)$$

About the Author

Doug Ito is an applications engineer for the High Speed ADC team at Analog Devices, Inc., San Diego, California. He earned a bachelor's degree in electrical engineering from San Diego State University. Doug is a member of ADI's EngineerZone® High Speed ADC Support Community. He can be reached at douglas.ito@analog.com.

Phased Array Antenna Patterns—Part 2: Grating Lobes and Beam Squint

Peter Delos, Technical Lead, **Bob Broughton,** Director of Engineering, and **Jon Kraft,** Senior Staff Field Applications Engineer

Introduction

This is the second article of our three-part series on phased array antenna patterns. In Part 1, we introduced the phased array steering concept and looked at the influencers on array gain. In Part 2, we'll discuss grating lobes and beam squint. Grating lobes can be hard to visualize, so we'll draw on their similarity with signal aliasing in digital converters, then use that to think of a grating lobe as a spatial alias. Next, we explore the issue of beam squint. Beam squint is an unfocusing of the antenna across frequency when we use phase shift, instead of a true time delay, to steer the beam. We'll also discuss the trade-off between these two steering methods and understand the impact of beam squint on typical systems.

An Introduction to Grating Lobes

So far, we have only seen the case where the element spacing is d = λ/2. Figure 1 begins to illustrate why an element spacing of λ/2 is such a common metric in phased arrays. Two cases are shown. First, in blue, the same 30° plot from Figure 11 in Part 1 is repeated. Next, the d/λ spacing is increased to 0.7 to show how the antenna pattern changes. With this increase in spacing, note how the beamwidth reduces, which is a positive result. The decreased spacing of nulls brings them closer together, which is also an acceptable result. But now there is a second angle, in this case at -70°, where there is full array gain. This is a most unfortunate result. This replica of antenna gain is defined as a grating lobe and can be considered spatial aliasing.

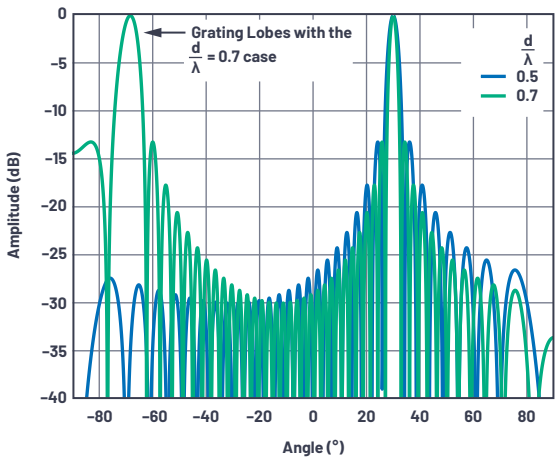

Figure 1. Normalized array factor of a 32-element linear array at two different d/λ spacings.

Analogy to Sampled Systems

An analogy to visualize grating lobes is to think of aliasing in a sampled system. In an analog-to-digital converter (ADC), undersampling is often used when frequency planning a receiver architecture. Undersampling involves purposefully reducing the sample rate (f_S) such that the sampling process translates frequencies above $f_S/2$ (the higher Nyquist zones) to appear as aliases in the first Nyquist zone. This causes those higher frequencies to appear as if they were at a lower frequency at the output of the ADC.

A similar analogy can be considered in phased arrays, where the elements spatially sample the wavefront. The Nyquist theorem can be extended to the spatial domain if we suggest that two samples—that is, elements—per wavelength are required to avoid aliasing. Therefore, if the element spacing is greater than λ/2, we can consider this spatial aliasing.

Calculating Where Grating Lobes Appear

But where will these spatial aliases (grating lobes) appear? In Part 1, we showed the phase shift applied to the elements across the array as a function of beam angle.

$$\Delta\Phi = \frac{2\pi d \sin\theta}{\lambda} \quad (1)$$

Inversely, we can compute the beam angle as a function of phase shift.

$$\theta = arcsin\left(\frac{\Delta\Phi}{2\pi} \times \frac{\lambda}{d}\right) \quad (2)$$

The arcsin function only produces real solutions for arguments between –1 and +1. Outside of these bounds, the solution is not real—the familiar "#NUM!" in spreadsheet software. Also note that the phase in Equation 2 is periodic and repeats every 2π. So, we could replace ΔΦ with (m × 2π + ΔΦ) in the beam steering equation to give Equation 3.

$$\theta = arcsin\left(\frac{m \times 2\pi + \Delta\Phi}{2\pi} \times \frac{\lambda}{d}\right) \quad (3)$$

where m = 0, ±1, ±2, ...

To avoid grating lobes, our goal is to obtain a single real solution. Mathematically, this is accomplished by keeping

$$\left|\frac{m \times 2\pi + \Delta\Phi}{2\pi} \times \frac{\lambda}{d}\right| > 1 \text{ for all } m \geq 1 \quad (4)$$

If we do so, then all the spatial images (that is, m = ±1, ±2, etc.) will produce non-real arcsin results, and we can ignore them. But if we can't do this, and therefore some values of m > 0 produce real arcsin results, then we end up with multiple solutions: grating lobes.

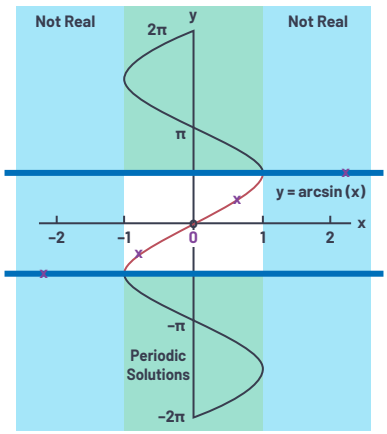

Figure 2. The arcsin function application to grating lobes.

Grating Lobes for d > λ and λ = 0°

Let's try some examples to better illustrate this. First, consider the case at mechanical boresight where θ = 0, and therefore ΔΦ = 0. Then Equation 3 simplifies to Equation 5.

$$\theta = arcsin\left(m \times \frac{\lambda}{d}\right), for\ \Delta\Phi = 0 \quad (5)$$

From this simplification, it is evident that if λ/d is > 1, then only m = 0 could give an argument that is bounded between –1 and +1. And that argument will just be 0, and the arcsin(0) = 0°, the mechanical boresight angle. So, this is all as we would expect. Furthermore, for any m ≥ 1, the arcsin argument will be too large (>1) and the resulting answer will not be real. We will see no grating lobes for θ = 0 and d < λ!

However, if d > λ (therefore λ/d is <1), then multiple solutions, grating lobes, could exist. For example, if λ/d = 0.66 (that is, d = 1.5λ), then real arcsin solutions would exist for m = 0 and for m = ±1. That m = ±1 is the second solution, which is the spatial aliasing of the desired signal. Therefore, we can expect to see three main lobes, each with approximately equal amplitude, located at arcsin(0 × 0.66), arcsin(1 × 0.66), and arcsin(–1 × 0.66). In degrees, these angles are 0° and ±41.3°. In fact, this is what our array factor plot shows in Figure 3.

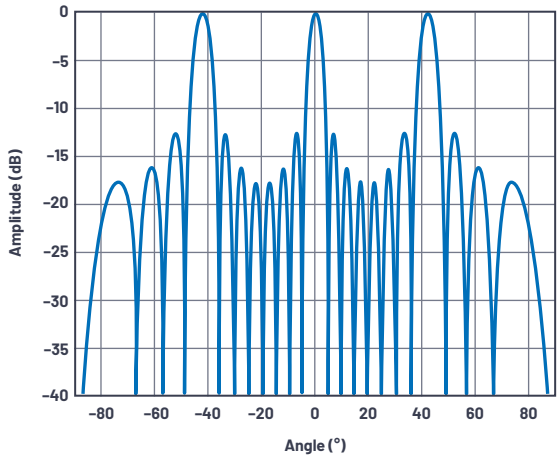

Figure 3. Array factor at boresight for d/λ = 1.5, N = 8.

Grating Lobes for λ/2 < d < λ

In simplifying the grating lobe equation (Equation 5), we chose to only look at mechanical boresight (ΔΦ = 0). And we saw that, at mechanical boresight, grating lobes would not appear for d < λ. But from our analogy of sampling theory, we know that we should also expect to see some kind of grating lobe for any spacing greater than λ/2. So where are the grating lobes for λ/2 < d < λ?

First, recall how the phase changed with steering angle in Figure 4 from Part 1. We saw ΔΦ range from 0 to ±π as the main lobe deviated from mechanical boresight. Therefore,

$$\frac{m \times 2\pi + \Delta\Phi}{2\pi} \times \frac{\lambda}{d} \tag{6}$$

will range

$$\text{from } 0 \text{ to} \left(\pm 0.5 \times \frac{\lambda}{d} \right) \; for \; m = 0 \tag{7}$$

And for |m| ≥ 1, it will always be something beyond

$$\pm 0.5 \times \frac{\lambda}{d} \tag{8}$$

This restricts the minimum permissible λ/d if we want to keep the entire arcsin argument > 1 for all |m| ≥ 1. Consider two cases:

- If λ/d ≥ 2 (that is, d ≤ λ/2), then you could never have multiple solutions, regardless of the value of m. All m > 0 solutions will result in an arcsin argument > 1. This is the only way to avoid grating lobes to the horizon.
- But if we purposefully restricted ΔΦ to something less than ±π, then we could tolerate a smaller λ/d and still not see grating lobes. Reducing the range of ΔΦ means reducing the maximum steering angle of our array. It is an interesting trade-off that will be explored in the next section.

Element Spacing Considerations

Should the element spacing always be less than λ/2? Not necessarily! This becomes a trade-off for the antenna designer to consider. If the beam is steered completely to the horizon, then θ = ±90°, and an element spacing of λ/2 is required (if no grating lobes are allowed in the visible hemisphere). But in practice, the maximum achievable steering angle is always less than 90°. This is due to the element factor and other degradations at large steering angles.

From the arcsin figure, Figure 2, we can see that if the y axis, θ, is restricted to a reduced limit, then grating lobes only occur at scan angles that are not used anyway. What would this reduced limit (θ_{max}) be for a given element spacing (d_{max})? We had said previously that our goal is to keep

$$\left| \frac{m \times 2\pi + \Delta\Phi}{2\pi} \times \frac{\lambda}{d} \right| > 1 \text{ for all } |m| \geq 1 \tag{9}$$

We can use this to calculate where our first grating lobe (m = ±1) would appear. Making this change, and using Equation 1 from Part 1 for ΔΦ, gives:

$$\frac{\pm 1 \times 2\pi + \frac{2\pi d_{max} sin\theta_{max}}{\lambda}}{2\pi} \times \frac{\lambda}{d_{max}} = 1 \tag{10}$$

Which simplifies to

$$\pm\lambda + d_{max} sin\theta_{max} = d_{max} \tag{11}$$

Then solving for d_{max}

$$d_{max} = \frac{\lambda}{1 + |sin\theta_{max}|} \; for \; \theta_{max} from \; 0 \; to \; \pm\pi/2 \tag{12}$$

This d_{max} is the condition for no grating lobes in the reduced scan angle (θ_{max}), where θ_{max} is less than π/2 (90°). For example, if the signal frequency is 10 GHz and we need to steer ±50° without grating lobes, then the maximum element spacing is:

$$d_{max} = \frac{(3 \times 10^8)/(10 \times 10^9)}{1 + \sin(50°)} = 17 \text{ mm} \tag{13}$$

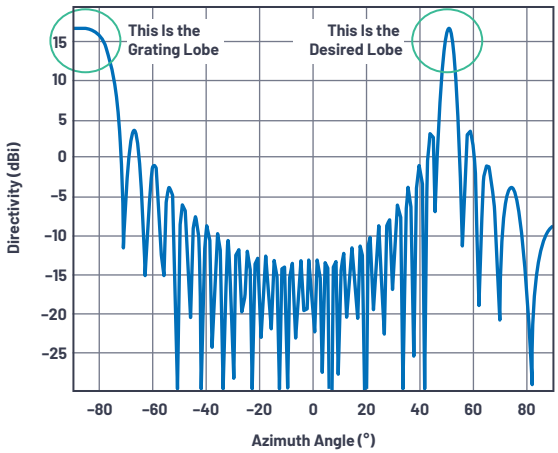

Figure 4. Grating lobes starting to appear at the horizon for θ = 50°, N = 32, d = 17 mm, and Φ = 10 GHz.

Restricting the maximum scan angle, then, brings a freedom to extend the element spacing to increase the physical size per channel and to extend the aperture for a given number of elements. An example application that could exploit this phenomenon is for an antenna assigned to a fairly narrow predefined direction. The element gain can be increased for directivity in the predefined direction and the element spacing can also be increased for a larger aperture. Both result in larger overall antenna gain within the narrowed beam angle.

Note that Equation 3 indicates a maximum spacing of one wavelength, even for zero steering angle. This is the case if grating lobes cannot be tolerated in the visible hemisphere. In the case of a GEO satellite, for example, the entire Earth is covered with a steering angle of 9° from mechanical boresight. It may be the case that grating lobes can be tolerated, as long as they don't land on the Earth's surface. In such a case, the element spacing can be several wavelengths, resulting in even more narrow beamwidths.

There are also antenna architectures worth noting that attempt to overcome the grating lobe problem by producing a nonuniform element spacing. These are categorized as aperiodic arrays, with spiral arrays as an example. For mechanical antenna construction reasons, it may be desirable to have a common building block that can be scaled to a larger array, but this would produce a uniform array that is subject to the grating lobe conditions described.

Beam Squint

We opened Part 1 by describing how, when a wavefront approaches an array of elements, there is a time delay between elements based on the wavefront angle θ relative to boresight. For a single frequency, the beam steering can be accomplished by replacing the time delay with a phase shift. This works for narrow-band waveforms, but for wideband waveforms where the beam steering is produced by a phase shift, the beam can shift direction as a function of frequency. This can be intuitively explained if we remember that a time delay is a linear phase shift vs. frequency. Thus, for a given beam direction, the phase shift required changes as a function of frequency. Or conversely for a given phase shift, the beam direction changes as a function of frequency. The concept of the beam angle changing as a function of frequency is called beam squint.

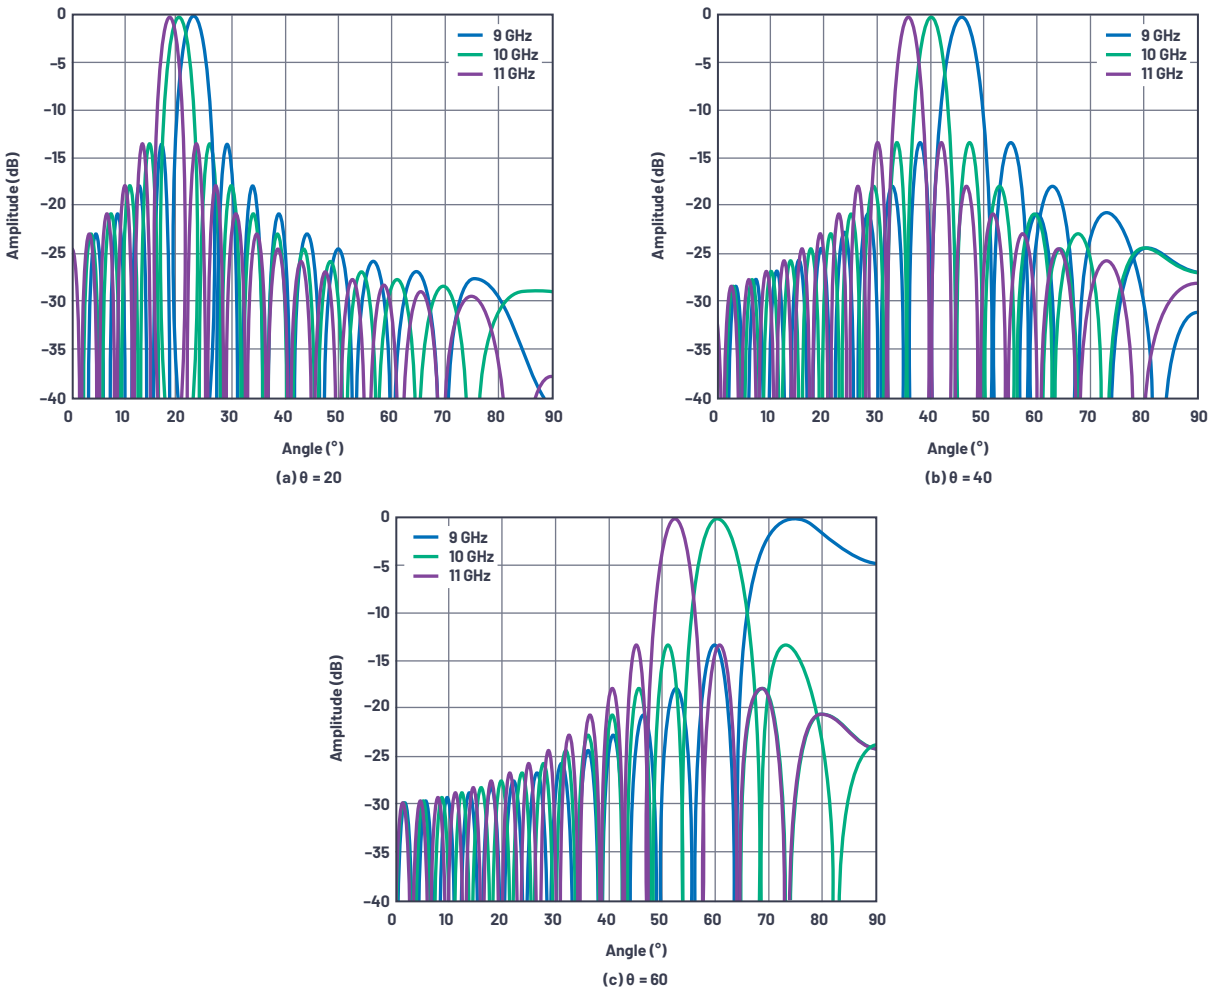

Figure 5. Beam squint example at X-band for a 32-element linear array with a λ/2 element spacing.

Also consider that at boresight, θ = 0, there is no phase shift across the elements and therefore no means to produce any beam squint. Therefore, the amount of beam squint must be a function of angle, θ, as well as the frequency variation. Figure 5 shows an X-band example. In this example, the center frequency is 10 GHz, the modulation bandwidth is 2 GHz, and it is apparent that the beam changes direction as a function of both frequency and the initial beam angle.

Beam squint can be calculated directly. Using Equation 1 and Equation 2, the beam direction deviation, beam squint, can be calculated as

$$\Delta\theta = arcsin\left(\frac{f_0}{f}\ sin\theta_0\right) - \theta_0 \qquad (14)$$

This equation is shown graphically in Figure 6. In Figure 6, the f/f_0 ratio shown is intentional. The reciprocal of (f_0/f) from the previous equation provides an easier way to visualize the change relative to a center frequency.

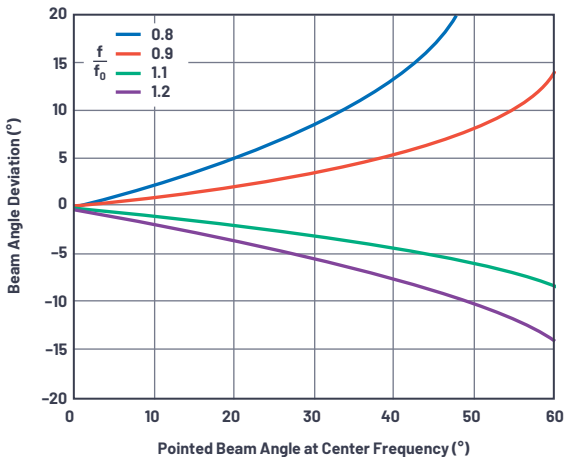

Figure 6. Beam squint vs. beam angle for several frequency deviations.

A few observations on beam squint:

- The deviation in beam angle vs. frequency increases as beam angle away from boresight increases.
- A frequency below the center frequency causes a larger deviation than a frequency above the center frequency.
- A frequency below the center frequency moves the beam further away from boresight.

Beam Squint Considerations

The beam squint (the deviation in steering angle vs. frequency) is caused by approximating a time delay with a phase shift. Implementing beam steering with true time delay units does not have this problem.

With the beam squint problem so clearly visible, why would anyone use a phase shifter over a time delay unit? Typically, this comes down to design simplicity and IC availability of phase shifters vs. time delays. Time delays are implemented in some form of transmission line and the total delay needed is a function of the aperture size. To date, most available analog beamforming ICs are phase shift based, but there are families of true time delay ICs emerging and these may become much more common for phased array implementations.

In digital beamforming, true time delay can be implemented in the DSP logic and digital beamforming algorithms. Therefore, a phased array architecture where every element is digitized would lend itself naturally to overcome the beam squint problem, while also providing the most programmable flexibility. However, the power, size, and cost of such a solution can be problematic.

In hybrid beamforming, there is a combination of analog beamforming for subarrays followed by digital beamforming for the full array. This can offer some natural beam squint mitigation worth considering. Beam squint is only subject to the subarray, which is a much wider beamwidth, so it is more tolerant to a beam angle deviation. Thus, as long as the subarray beam squint is tolerable, then the hybrid beamforming architecture can be implemented with phase shifters in the subarrays followed by true time delay in the digital beamforming.

Summary

This concludes Part 2 of a three-part series on phased array antenna patterns. In Part 1, we introduced beam pointing and the array factor. In Part 2, we introduced imperfections of grating lobes and beam squint. In Part 3, we will discuss tapering as a method to reduce sidelobes, and also provide insight into the impact of phase shifter quantization errors.

References

Balanis, Constantine A. *Antenna Theory: Analysis and Design*, third edition. Wiley-Interscience, 2005.

Longbrake, Matthew. "True Time Delay Beamsteering for Radar." *2012 IEEE National Aerospace and Electronics Conference (NAECON),* IEEE, 2012.

Mailloux, Robert J. *Phased Array Antenna Handbook*, second edition. Artech House, 2005.

O'Donnell, Robert M. "Radar Systems Engineering: Introduction." IEEE, June 2012.

Skolnik, Merrill. *Radar Handbook*, third edition. McGraw-Hill, 2008.

About the Author

Peter Delos is a technical lead in the Aerospace and Defense Group at Analog Devices in Greensboro, NC. He received his B.S.E.E. from Virginia Tech in 1990 and M.S.E.E. from NJIT in 2004. Peter has over 25 years of industry experience. Most of his career has been spent designing advanced RF/analog systems at the architecture level, PWB level, and IC level. He is currently focused on miniaturizing high performance receiver, waveform generator, and synthesizer designs for phased array applications. He can be reached at peter.delos@analog.com.

About the Author

Bob Broughton started at Analog Devices in 1993 and has held positions as a product engineer and an IC design engineer, and is currently the director of engineering in the Aerospace and Defense Business Unit. Prior to ADI, Bob worked at Raytheon as an RF design engineer and at Peregrine Semiconductor as an RFIC designer. Bob graduated with a B.S.E.E. from West Virginia University in 1984. He can be reached at bob.broughton@analog.com.

About the Author

Jon Kraft is a senior staff FAE in Colorado and has been with ADI for 13 years. His focus is software-defined radio and aerospace phased array radar. He received his B.S.E.E. from Rose-Hulman and his M.S.E.E. from Arizona State University. He has nine patents issued, six with ADI, and one currently pending. He can be reached at jon.kraft@analog.com.

Using LTspice to Analyze Vibration Data in Condition-Based Monitoring Systems

Simon Bramble, Staff Field Applications Engineer

Abstract

This article describes how to use LTspice to analyze the frequency content of vibration data in condition-based monitoring systems in order to give an early warning of motor failure in industrial machinery. The article explains how to take X, Y, and Z plane data from a Microsoft Excel® spreadsheet and convert it into a format to which LTspice can apply a Fourier transform to produce a plot of the harmonic content of the vibration data.

Introduction

The advancement of digital technology shows no sign of slowing down, pervading every area of our lives. Giving machines intelligence is far from an Orwellian dystopia; it instead yields efficiency improvements in factory automation, as automated feedback loops can reduce direct maintenance time.

Industry 4.0 describes the concept of bringing the advantages of big data to the factory floor. Machines fitted with sensors can monitor their own performance and communicate with each other, enabling them to share the overall workload while providing important diagnostic information to the back office, be it in the same building or on another continent.

A quick survey of Analog Devices' product offerings shows that ADI is heavily committed to providing solutions for the Industrial Internet of Things (IIOT), namely by offering robust, high performance signal chain components from the sensor to the cloud.

One such area of industrial automation is that of condition-based monitoring (CbM), whereby the nominal operating characteristics of a machine are carefully calibrated, then the machinery itself is closely monitored with local sensors. Conditions that deviate from the nominal signal show that the machine needs maintenance. Thus, machines equipped with condition-based monitoring systems can be serviced when they actually need it instead of as part of a relatively arbitrary servicing schedule.

The preeminent way to determine the state of health of a motor is to examine its vibration signature. Analog Devices' MEMS technology enables the vibration signature of a motor to be continuously monitored, revealing the health of the motor when its signature is compared to a known no-fault motor. Indeed, each motor fault has its own unique harmonic signature. By looking at the harmonic content of the vibration pattern, faults can be detected in the bearings, inner and outer races, and even in the gearbox teeth.

Analyzing Vibration Data in LTspice

To produce data for Fourier analysis in LTspice, three ADXL1002 accelerometers were connected to a motor, as shown in Figure 1, to measure vibration in the side-to-side, vertical, and fore-to-aft directions—X, Y, and Z, respectively.

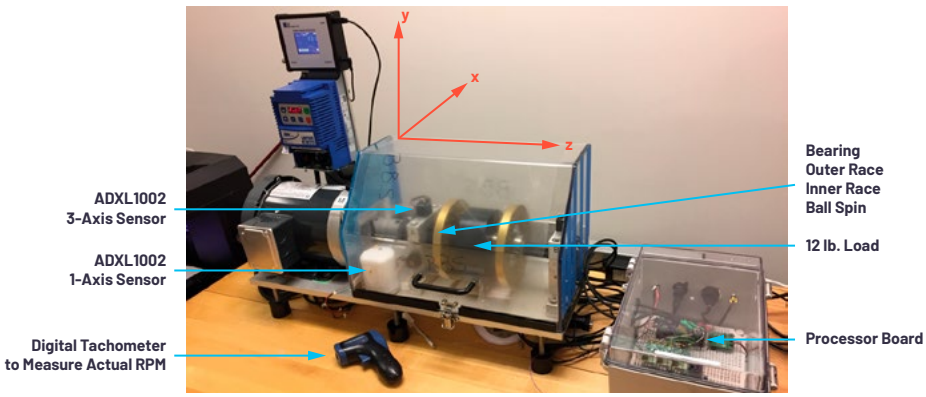

Figure 1. Channels X, Y, and Z measured vibration in the side-to-side, vertical, and fore-to-aft directions, respectively.

The vibration data was downloaded and saved in a Microsoft Excel spreadsheet. The data was sampled at 500 kSPS, so one second of vibration data resulted in three columns of Microsoft Excel data, each 500,000 lines long. A sample of the X, Y, and Z data is shown in Figure 2.

	A	B	C
1	35403	34899	35171
2	35411	34900	35180
3	35403	34910	35184
4	35404	34912	35181
5	35412	34921	35185
6	35404	34913	35174
7	35401	34915	35177
8	35388	34917	35181
9	35399	34927	35181
10	35399	34924	35178

Figure 2. An extract of the X, Y, and Z data.

The harmonic content of this data can now be examined to determine the state of health of the motor. Fourier analysis is the mathematical process of extracting the component frequency content from a waveform. The spectral content of a pure sine wave consists of only one frequency, called the fundamental. If the sine wave is distorted, other frequencies aside from the fundamental appear. By analyzing the frequency content of the vibration pattern of the motor, an accurate diagnosis of its state of health can be determined.

Hardware and software capable of conducting Fourier analysis can be expensive, so here we show a method of performing Fourier analysis on MEMS data that is essentially no-cost.

LTspice is a powerful, and free, circuit simulator that has the ability to plot the frequency content of any waveform using Fourier analysis, including the vibration data obtained from the MEMS sensor in a condition-based monitoring system.

LTspice can produce a Fourier analysis plot when the data is in the format shown in Figure 3, where each vibration data point is paired with its corresponding timestamp.

	A	B
1	time1	value1
2	time2	value2
3	time3	value3
4	.	.
5	.	.
6	.	.
7	.	.
8		

Figure 3. Format of time and voltage instances.

It is relatively easy to massage the data into this format using Microsoft Excel. Here is the process for doing so.

First, separate the columns of data in Figure 2 into three worksheets within the Excel file, named X, Y, and Z, as shown in Figure 4.

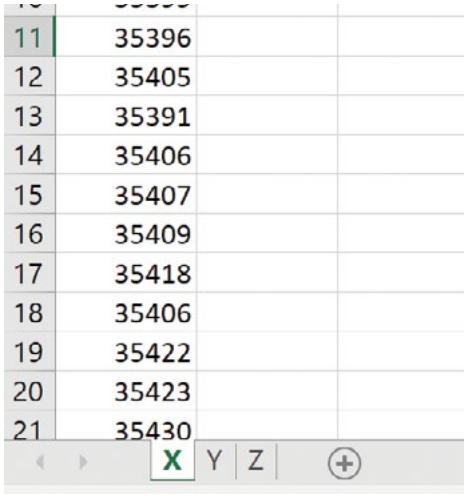

Figure 4. Three sheets were created and the X, Y, and Z data were copied to their respective sheets.

Insert a column to the left of the data—this column is for the timestamp of each data value.

Since there were 500,000 samples of data taken over a one second timespan, each data point was taken 2 µs apart. Therefore, in the first cell of the new column, enter

2E-6

representing the first timestamp at 2 µs.

The easiest way to fill in the rest of the timestamp column is to use the **Series** command. In the **Search** box in Microsoft Excel, type "Series" to bring up the menu options shown in Figure 5.

From here, choose **Fill Series** or **Pattern**, then **Series**... from the dropdown menu.

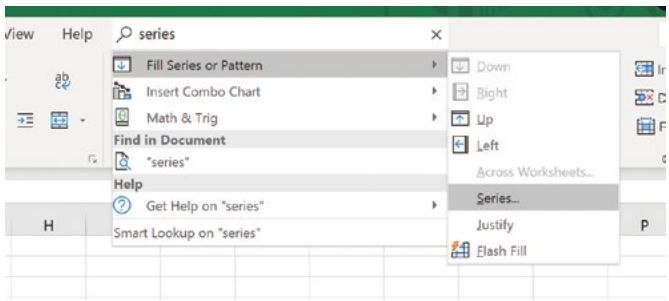

Figure 5. How to fill many cells in Microsoft Excel.

The dialog box shown in Figure 6 appears, with the **Columns** and **Linear** radio buttons selected. Enter a **Step value** of 2E-6 and a **Stop value** of 1.

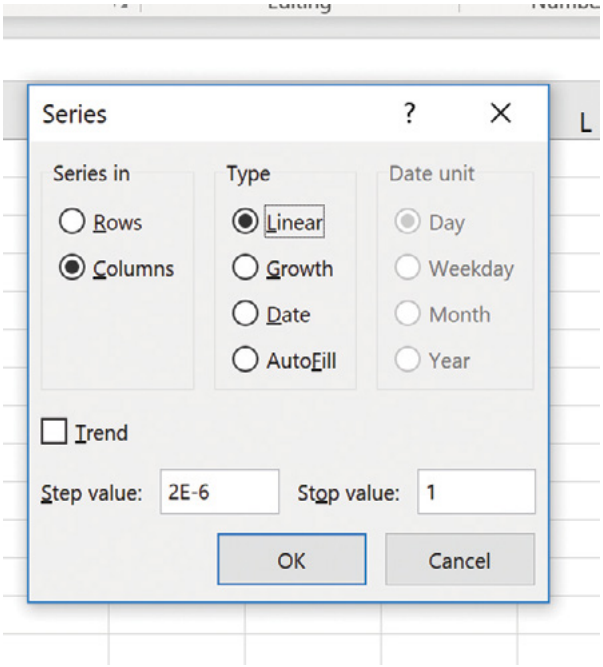

Figure 6. Fill the cells with a linearly expanding dataset.

Click **OK** to fill the left column data timestamps, incrementing from 2 µs to 1 second. The same can be achieved by filling in the first few values, then dragging the cursor to the bottom cell at the end of the data range—but with 500,000 lines of data, this would be a long drag.

The data is now in a format ready to be processed by LTspice, as shown in Figure 7.

	A	B	C
1	2.00E-06	35403	
2	4.00E-06	35411	
3	6.00E-06	35403	
4	8.00E-06	35404	
5	1.00E-05	35412	
6	1.20E-05	35404	
7	1.40E-05	35401	
8	1.60E-05	35388	
9	1.80E-05	35399	
10	2.00E-05	35399	
11	2.20E-05	35396	
12	2.40E-05	35405	
13	2.60E-05	35391	
14	2.80E-05	35406	
15	3.00E-05	35407	
16	3.20E-05	35409	
17	3.40E-05	35418	
18	3.60E-05	35406	
19	3.80E-05	35422	
20	4.00E-05	35423	

Figure 7. Columns showing the timestamp and corresponding data sample.

If the dataset is large and the sample interval is low, it is possible that Microsoft Excel might round the timestamps to an inappropriate number of decimal places. If this is the case, highlight the first column, then select **Format > Format Cells**, as shown in Figure 8.

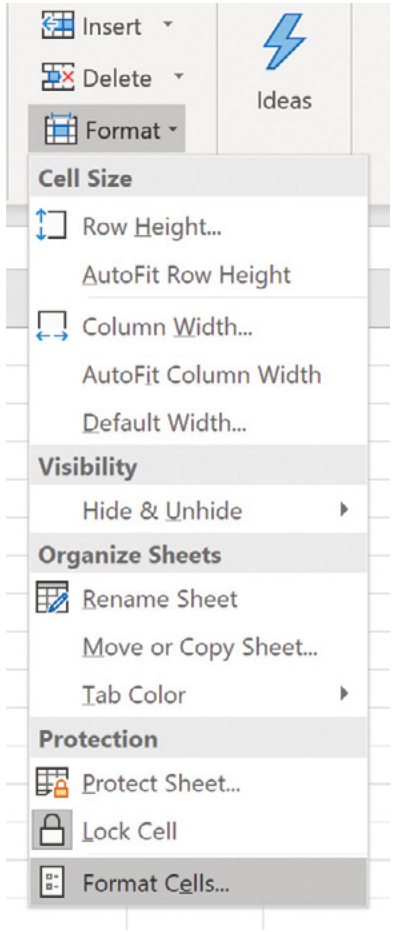

Figure 8. Reformat the cells to remove any rounding errors.

Select the appropriate number of decimal places, as shown in Figure 9.

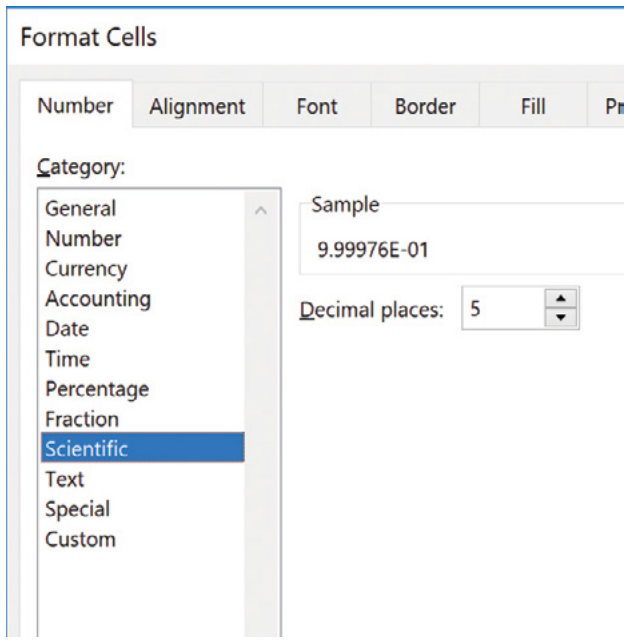

Figure 9. Increasing the timestamp resolution to five decimal places.

Once the timestamp column is populated and the significant digits expanded, copy both columns of each sheet to its own Notepad or another text editor file, as shown in Figure 10.

Vibration x fault_10Hz.txt - Notepad

File Edit Format View Help

2.00000E-06	35403
4.00000E-06	35411
6.00000E-06	35403
8.00000E-06	35404
1.00000E-05	35412
1.20000E-05	35404
1.40000E-05	35401
1.60000E-05	35388
1.80000E-05	35399
2.00000E-05	35399
2.20000E-05	35396
2.40000E-05	35405
2.60000E-05	35391
2.80000E-05	35406
3.00000E-05	35407
3.20000E-05	35409
3.40000E-05	35418
3.60000E-05	35406
3.80000E-05	35422
4.00000E-05	35423
4.20000E-05	35430
4.40000E-05	35439
4.60000E-05	35436
4.80000E-05	35438

Figure 10. Text file containing time and vibration data.

You should have three text files containing the vibration data for axes X, Y, and Z in the condition-based monitoring system.

This data can now be read directly into LTspice.

Construct a schematic in LTspice, as shown in Figure 11. In this design, there are six voltage sources corresponding to fault and no fault data for axes X, Y, and Z. This enables a Fourier analysis to be performed on the vibration data from a new motor so it can be compared with the Fourier analysis of the data from a suspected faulty motor. A big advantage of this method is that the frequency plot of a new (not faulty) motor can be overlaid on that of a suspected faulty motor, so the difference in performance can be seen.

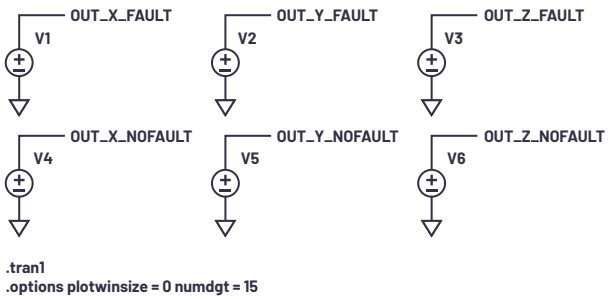

Figure 11. LTspice schematic showing voltage outputs for faulty and nonfaulty vibration data.

The LTspice command

.options plotwinsize=0 numdgt=15

removes the default compression in LTspice and sometimes produces clearer results. The simulation will run faster if this line is omitted but may produce less accurate results.

Once the schematic is complete, right-click each voltage source, select the **Advanced** button, select the **PWL File** radio button, then enter the file name of the appropriate text file containing the vibration data, as shown in Figure 12. This creates a piecewise linear voltage source consisting of a series of voltages and their corresponding time instances. Your life will be easier if these text files are stored in the same directory as the LTspice file.

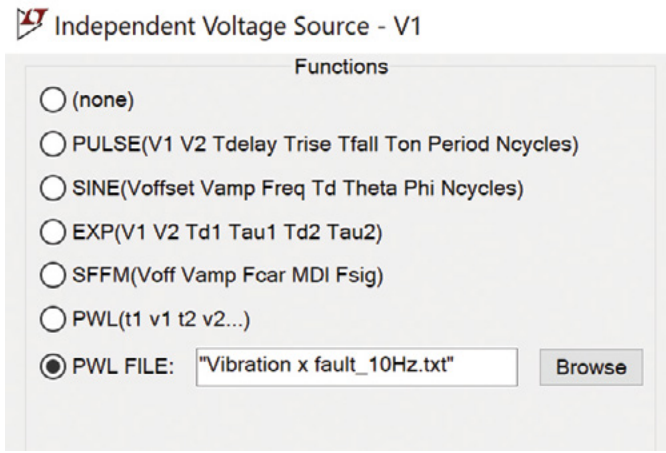

Figure 12. Creating a piecewise linear voltage source from the vibration data.

The transient analysis should then be configured to run for the duration of the original vibration test using the command

.tran 1

Run the simulation. The simulation may take some time to complete depending on the data points and the length of the transient analysis.

The simulation results of the faulty and nonfaulty motors are shown in Figure 13. The experiment was conducted on a motor rotating at 587.3 rpm with a bearing fault with the outer race misaligned and with a 12 pound load. The plots also show the vibration pattern of a motor with no fault rotating at the same speed. It is clear that the faulty motor has a significantly higher amplitude in the vibration signature compared with the nonfaulty motor.

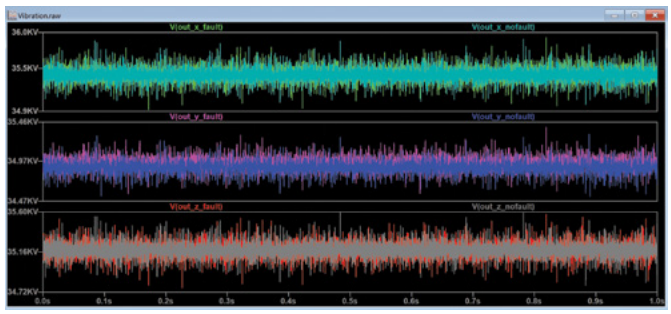

Figure 13. Time domain results of the vibration data for a faulty and nonfaulty motor.

With the **Waveform** window highlighted, select **View > FFT** from the menu bar. This will compute the FFT based on the transient data.

Looking at the data in Figure 2, the numbers show a small variation about a large offset of approximately 35,000. When simulated in LTspice, this translates to a dc offset voltage of 35,000 V with an ac waveform on top of this offset.

In the Fourier plot, this offset voltage manifests itself as a large spur at dc, so when LTspice autoscales the Y axis, the harmonics of interest are scaled too small. Right-click the X axis to specify a frequency range above dc, so the dc offset voltage is ignored—a range of 5 Hz to 1 kHz should suffice.

Right-click the Y axis and select the **Linear** radio button to view the harmonics, as shown in Figure 14.

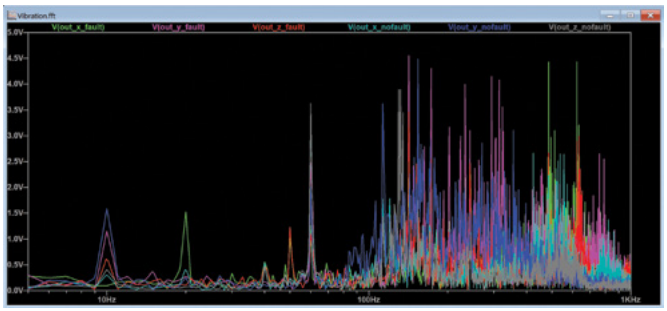

Figure 14. The Fourier plot with the dc spur removed and shown on a linear scale.

Right-clicking inside the plot area enables extra plot panes to be added, which enables the spectral content of the vibration to be separated into X, Y, and Z plots, as shown in Figure 15.

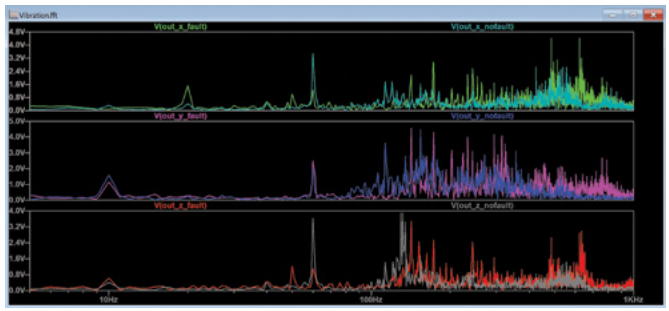

Figure 15. The X, Y, and Z vibration plots are separated out.

The 10 Hz rotation frequency of the motor can clearly be seen, as well as significant harmonics at 60 Hz, 142 Hz, and 172 Hz. It is beyond the scope of this article to analyze what component inside the motor has caused these harmonics, but there is no doubt that the vibration pattern has changed due to motor wear.

Conclusion

Analog Devices' range of MEMS accelerometers provide critical data to enable the early detection of motor failure, but that is only half the solution. The data must be carefully studied using Fourier analysis. Unfortunately, equipment or software capable of performing Fourier analysis is typically expensive. LTspice provides a free-of-charge route to accurately analyze CbM data, enabling the early detection and diagnosis of machine failure.

About the Author

Simon Bramble graduated from Brunel University in London in 1991 with a degree in electrical engineering and electronics, specializing in analog electronics and power. He has spent his career in analog electronics and worked at Linear Technology (now part of Analog Devices). He can be reached at simon.bramble@analog.com.

DC-to-DC Conversion Directly from Automotive Battery Input: 5 A, 3.3 V, and 5 V Supplies Meet Stringent EMI Emission Standards

Zhongming Ye, System Applications Manager

Introduction

Noise-sensitive applications in harsh automotive and industrial environments require low noise, high efficiency buck regulators that can fit into tight spaces. Monolithic buck regulators, which include the MOSFET power switches in the package, are often chosen because of their small overall solution size relative to a traditional controller IC and external MOSFETs. Monolithic regulators that can operate at high frequency—in the 2 MHz territory well above the AM band—also help to reduce the size of external components. Furthermore, if a regulator offers a low minimum on time (T_{ON}), the regulator can directly operate from higher voltage rails without intermediate regulation, saving space and complexity. Low minimum on times require fast switching edges and minimum deadtime control to effectively reduce the switching loss and allow high switching frequency operation.

Another way to save space is to reduce the number of components required to meet electromagnetic interference (EMI) standards and thermal requirements. Unfortunately, in many cases, simply shrinking the converter makes meeting these requirements more difficult. This article presents state-of-the-art solutions that save space while also achieving low EMI and excellent thermal performance.

Switch-mode power converters are chosen for their efficiency, especially at high step-down ratios, but one trade-off is the switching-action induced EMI. In a buck converter, EMI arises from the fast current changes (high di/dt) in the switches and from switch ringing due to the parasitic inductance in the hot loop.

EMI is just one of the parameters that system design engineers must struggle with when trying to design a compact, high performance power supply. A number of critical design constraints are often at odds, requiring critical compromises within the design limits and time to market.

Improving EMI Performance

To reduce EMI in a buck converter, one must reduce the radiating effect of the hot loop as much as possible and minimize signals from the source. There are a number of ways to reduce radiated EMI, but many also reduce the performance of the regulator.

For instance, in a typical discrete FET buck regulator, the switching edge is slowed down, with an external gate resistor, BOOST resistor, or snubber, as a last resort method to meet the stringent radiated emissions standards in the automotive industry. Such a quick fix for the EMI comes at the price of performance; namely lower efficiency, higher component count, and larger solution size. Slow switching edges increase switching losses, as well as duty ratio loss. The converter must operate at a lower frequency—for instance, 400 kHz—to achieve satisfactory efficiency and pass mandatory radiated EMI emission tests. Figure 1 shows typical switching node voltage waveforms with a fast switching edge and a slow switching edge, respectively. As shown, the switching edge is significantly slower, resulting in increased switching losses, and a significant increase in the minimum duty cycle, or step-down ratio, not to mention other negative effects on performance.

Slowing the switching frequency also increases the physical size of the converter inductor, output cap, and input cap. Meanwhile, a bulky π filter is necessary to pass the conducted emissions tests. The inductance, L, and capacitance, C, in the filter get bigger as the switching frequency goes down. The inductor current rating should be larger than the maximum input current at low line full load. Therefore, a bulky inductor and multiple capacitors are required on the front end to help pass the stringent EMI standards.

For example, at a 400 kHz switching frequency (as opposed to 2 MHz), in addition to increasing the size of the inductor and capacitor, the inductors and capacitors in the EMI filter must also be relatively large to pass the conducted EMI standard required in automotive applications. One reason is that they must not only attenuate the switching fundamental frequency at 400 kHz, but all of its harmonics up to 1.8 MHz. A regulator operating at 2 MHz does not have this problem. Figure 2 shows the size of a 2 MHz solution against a 400 kHz solution.

Shielding might be a last ditch remedy to reduce radiated emissions, but shielding takes space that might not be available in the application and would require additional mechanical design and test iterations.

To avoid the AM frequency bandwidth and maintain a small solution size, a 2 MHz or higher switching frequency is preferred in automotive applications. With the AM band avoided, it is just a matter of ensuring that higher frequency noise—also known as harmonics—and switch ringing are also minimized. Unfortunately, high frequency switching usually results in increased radiated emissions from 30 MHz to 1 GHz.

There are switching regulators that feature fast and clean switching edges, which reduce EMI, such as the Silent Switcher® devices in ADI's Power by Linear™ line. First, though, let's look at some other features that can help.

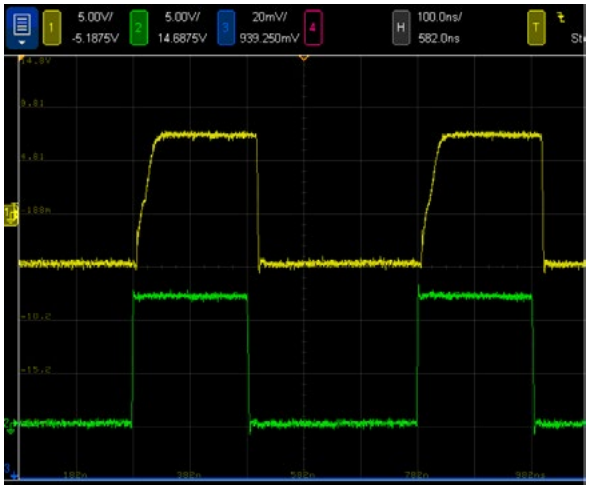

Figure 1. A slow switching edge means significant switching loss in addition to duty ratio loss.

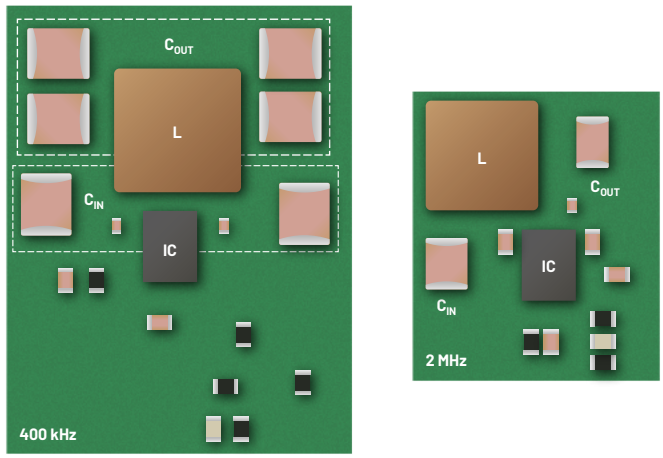

Figure 2. The size of a 2 MHz solution compared to that of a 400 kHz solution.

Spread spectrum frequency modulation (SSFM) is a technique that dithers the system clock within a known range, thus distributing the EMI energy over the frequency domain. Although the switching frequency is 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 and other regions.

Figure 3 shows an ultralow EMI and high efficiency 12 V to 5 V/5 A converter that operates at a switching frequency at 2 MHz using the LT8636 Silent Switcher monolithic buck regulator. Figure 4 shows conducted and radiated EMI performance for a tested demonstration circuit at a 14 V input and an output of 5 A at 5 V. At the front end, a small inductor and ceramic cap helps to filter out the conducted noise, while the ferrite bead and the ceramic capacitor help to reduce the radiated noise. Two small ceramic caps are placed to the input and ground pins to minimize the area of the hot loop, while also splitting the hot loop, which helps to cancel out high frequency noise.

To improve the EMI performance, the circuit is set to operate in spread spectrum mode: SYNC/MODE = $INTV_{CC}$. A triangular frequency modulation is used to vary the switching frequency between the value programmed by R_T to approximately 20% higher than that value—that is, when the LT8636 is programmed to 2 MHz, the frequency will vary from 2 MHz to 2.4 MHz at a 3 kHz rate.

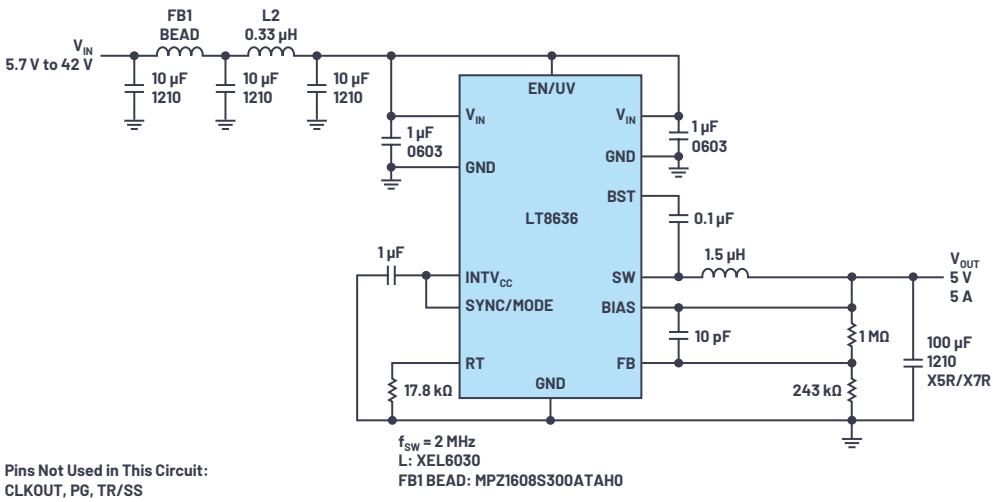

Figure 3. Ultralow EMI LT8636 5 V/5 A step-down converter in spread spectrum mode with 7 A peak works over 5.7 V to 42 V.

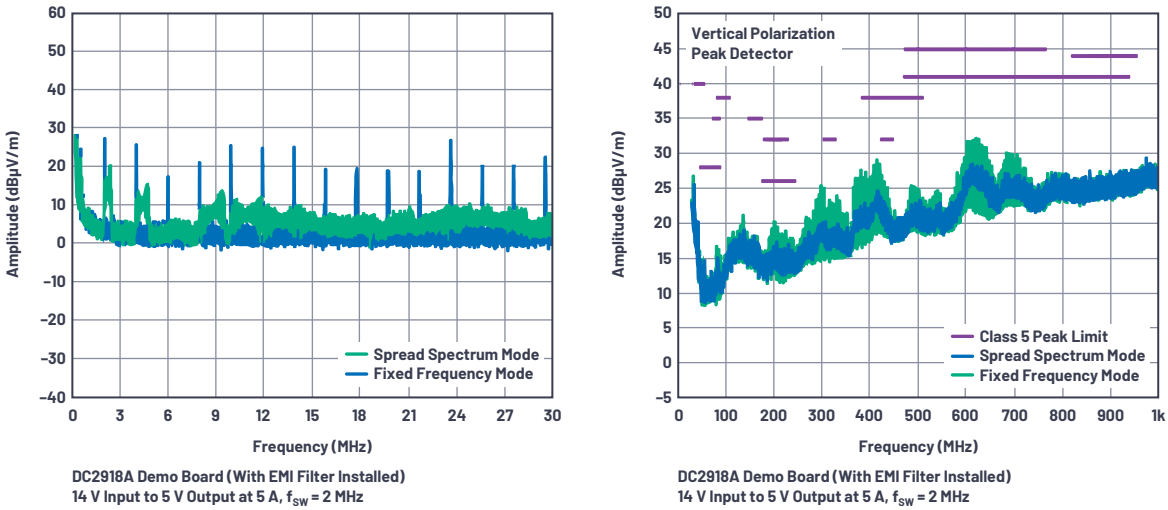

Figure 4. CISPR 25 radiated EMI emission with and without spread spectrum mode.

From the conducted EMI spectrum, it is obvious the peak harmonic energy is spread out, reducing the peak magnitude at any particular frequency—noise is reduced due to the spread spectrum function by at least 20 dBµV/m. From the radiated EMI spectrum, it is also obvious that spread spectrum mode reduces the radiated EMI as well. This particular circuit meets the stringent automotive CISPR 25 class 5 radiated EMI specification with a simple EMI filter at the input side.

High Efficiency over the Entire Load Range

The number of electronics devices in automotive applications is only increasing, with most devices demanding more supply current with each design iteration. With active load currents so high, heavy load efficiency and proper thermal management are top priorities—robust operation depends on thermal management, with unmitigated heat production possibly resulting in costly design problems.

System designers are also concerned with light load efficiency, which is arguably just as important as heavy load efficiency, since battery life is mostly determined by the quiescent current at light load or no load. Trades-off in the silicon, as well as system-level design, have to be made among full load efficiency, no load quiescent current, and light load efficiency.

It might seem straightforward that in order to achieve high efficiency at full load, the $R_{DS(ON)}$ of the FET, especially the bottom FET, should be minimized. However, a transistor with low $R_{DS(ON)}$ usually has a relatively high capacitance, with an associated increase in switching and gate drive losses, plus a larger die size and cost. In contrast, the LT8636 monolithic regulator has very low MOSFET conduction resistances, enabling exceptional efficiency in full load conditions. The maximum output current for the LT8636 is 5 A continuous and 7 A peak in still air without any additional heat sink, simplifying robust design.

To enhance light load efficiency, regulators that operate in low ripple Burst Mode® keep the output capacitor charged to the desired output voltage while minimizing the input quiescent current while minimizing output voltage ripple. In Burst Mode operation, current is delivered in short pulses to the output capacitor, followed by relatively long sleep periods, where most of the control (logic) circuits are shutdown.

In order to achieve higher light load efficiency, a larger value inductor is preferred since more energy can be delivered to the output during the short pulses and the buck regulator can remain longer in sleep mode between each pulse. By maximizing the time between pulses, and minimizing the switching loss of each short pulse, monolithic buck converter quiescent current can approach 2.5 µA in a monolithic regulator, such as the LT8636. This number is compared with tens of µA or hundreds of µA of the typical parts on the market.

Figure 5 shows a high efficiency solution for 3.8 V/5 A output from 12 V input for automotive applications using the LT8636. The circuit runs at 400 kHz for very high efficiency, and an XAL7050-103 10 µH inductor is used. It maintains efficiency above 90% with loads as light as 4 mA and as high as 5 A. The peak efficiency is 96% at 1 A.

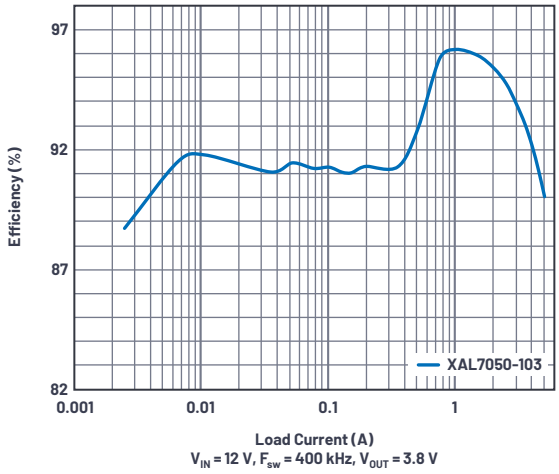

Figure 5. The efficiency of a 12 V to 3.8 V/5 A solution with an XAL7050-103 inductor (f_{sw} = 400 kHz).

Figure 6 shows the efficiency from µA to 5 A for this solution. The internal regulator is supplied from the 5 V output through the BIAS pin to minimize power dissipation. Peak efficiency reaches 95%; full load efficiency is 92% for a 5 V output from 13.5 V input. Light load efficiency remains at or above 89% for loads down to 30 mA for the 5 V application. The converter runs at 2 MHz and the inductor used for the test is an XEL6060-222 to optimize the efficiency in both heavy and light loads, in a relatively compact solution. Light load efficiency can be further improved—to above 90%—by using a larger inductor. The current in the feedback resistor divider is minimized as it appears to the output as load current.

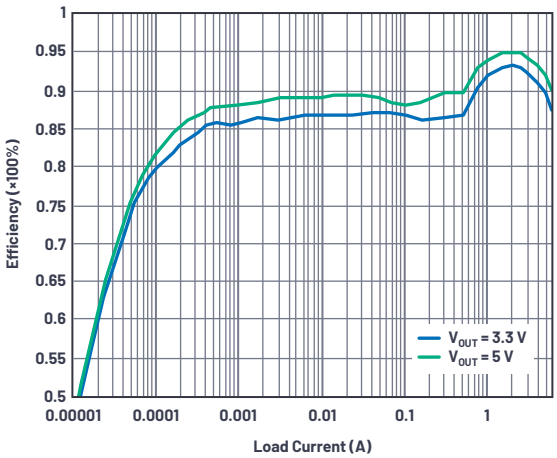

Figure 6. Efficiency of the LT8636 in a solution for 13.5 V to 5 V and 3.3 V using an XEL6060-222 inductor (f_{SW} = 2 MHz).

Figure 7 shows thermal performance for this solution under a 4 A constant load plus a 4 A pulsed load (8 A total at pulse) with a duty cycle of 10% (of 2.5 ms)—from a 13.5 V input and still air at ambient room temperature. Even at a 40 W pulsed power and 2 MHz switching frequency, the LT8636 case temperature remains below 40°C, enabling the circuit to run safely up to 8 A in short periods with no fans or heat sinks. This is possible with a 3 mm × 4 mm LQFN package because of enhanced thermal packaging technology and the LT8636's high efficiency at high frequency.

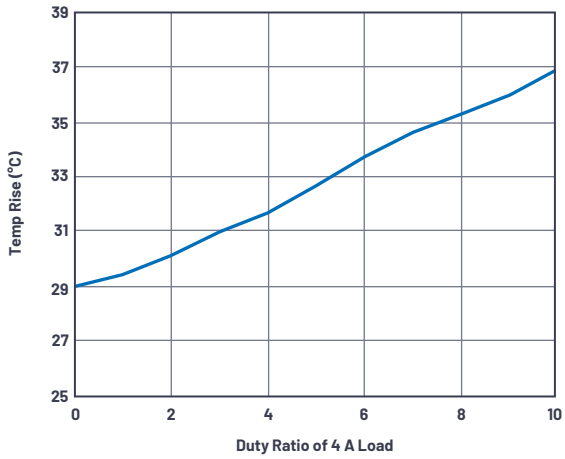

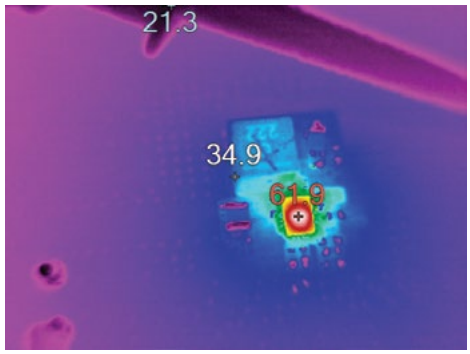

Figure 7. The 3 mm × 4 mm LT8636 in a 13.5 V to 5 V/4 A constant load plus a 4 A pulsed load (10% duty cycle) thermal picture showing temperature rise.

Shrink Solution Size with High Frequency Operation

Space is an increasing premium in automotive applications, necessitating that power supplies shrink to fit costly board footprints. Increasing a power supply's switching frequency enables the use of smaller external components such as capacitors and inductors. Plus, as previously mentioned, in automotive applications, switching frequencies above 2 MHz (or below 400 kHz) keep the fundamental out of the AM radio band. Let's compare a commonly used 400 kHz design to a 2 MHz design. In this case, quintupling the switching frequency to 2 MHz reduces the required inductance and output capacitance to one-fifth of the 400 kHz design. Seems easy. Nevertheless, even ICs that are high frequency capable may not be usable in many applications because of some of the trade-offs inherent in using a high frequency solution.

For instance, high frequency operation in high step-down ratio applications requires a low minimum on time. According to the equation $V_{OUT} = T_{ON} \times f_{SW} \times V_{IN}$, at a 2 MHz operating frequency, an on time for the top switch (T_{ON}) of about 50 ns is required to produce 3.3 V from 24 V input. If the power IC cannot achieve this low on time, pulses must be skipped to maintain the low regulated output—essentially defeating the purpose of the high switching frequency. That is, the equivalent switching frequency (due to pulse skipping) is likely in the AM band. With minimum top switch on time of 30 ns, the LT8636 allows direct high V_{IN} to low V_{OUT} conversion at 2 MHz. In contrast, many devices are limited to a >75 ns minimum, requiring they be operated at low frequency, 400 kHz, for higher step-down ratios to avoid skipping pulses.

Another common issue with high switching frequency is that switching losses tend to increase. The switching related losses include the switch turn on loss, turn off loss, and gate drive loss—all roughly linearly dependent on the switching frequency. Nevertheless, these losses can be improved with faster switch turn-on and turn-off times. The LT8636 switching turn-on and turn-off times are very short, less than 5 V/ns, resulting in minimum deadtime and minimum diode time, reducing the switching losses at high frequency.

The LT8636 used in the solutions here is assembled in a 3 mm × 4 mm LQFN, using a monolithic construction with integrated power switches and inclusion of all necessary circuitry yields a solution with a minimal PCB footprint. The large area exposed ground pad under the IC guides heat to the PCB through a very low thermal resistance (26°C/W) path, reducing the need for additional thermal management. The package is designed for FMEA compatibility. Silent Switcher technology reduces the PCB area of the hot loop, so radiated EMI with such high switching frequency can be easily addressed with simple filters, as shown in Figure 3.

Conclusion

With careful IC selection, it is possible to produce compact high performance power supplies for automotive applications without the usual trade-offs. That is, high efficiency, high switching frequency, and low EMI can all be achieved. To exemplify the types of compact designs that can be achieved, the solutions shown in this article use the LT8636, a 42 V, 5 A continuous/7 A peak monolithic step-down Silent Switcher regulator in a 3 mm × 4 mm LQFN package. In this IC, the V_{IN} pins are split and placed symmetrically on the IC, splitting the high frequency hot loop, mutually cancelling the magnetic fields to suppress the EMI radiated emission. Also, a synchronous design and fast switching edges improve efficiency at heavy load, while the light load efficiency benefits from the low ripple Burst Mode operation.

The LT8636 also fits automotive applications with a 3.4 V to 42 V input range and low dropout, enabling it to operate in automotive crank or load dump scenarios. In automotive applications, system designers are accustomed to facing a number of trade-offs when trying to shrink power supply solution size, but with the designs shown here, designers can achieve all their performance goals without trade-offs.

About the Author

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

RAQ Issue 178: Wide Voltage Range Automotive Circuit Protector

Albert Hinckley, Product Evaluation Engineer

Question:

Are there overvoltage and undervoltage protection devices available, especially for automotive applications?

Answer:

There are specific power path controllers available to protect your system.

Introduction

Ignition cranking during startup and load dumps during shutdown are common sources of voltage transients on an automotive supply line. These undervoltage (UV) and overvoltage (OV) transients can have significant magnitudes and will damage circuits that are not designed to operate during these extremes. Specialized UV and OV protection devices have been developed to disconnect sensitive electronics from power supply transients.

The LTC4368 is an example of a specialized UV and OV protection device. It utilizes a window comparator to monitor and validate the input supply. The supply voltage is monitored by a resistive divider network connected to the UV and OV monitor pins. The window comparator output drives the gates of two N-channel MOSFETs that make or break the connection between the supply and the load.

The LTC4368's window comparator is designed with 25 mV of hysteresis on its monitor pins to improve noise immunity. Hysteresis can prevent false MOSFET on/off switching due to ripple or other high frequency oscillations on the supply line. The 25 mV of hysteresis in the LTC4368 is equivalent to 5% of the monitor pin thresholds and is common for UV and OV protection devices.

For their own protection or to reduce ignition loading, some automotive accessory circuits must be disconnected from the supply line during startup or shutdown. Due to the large transients involved, these circuits may require more hysteresis than the LTC4368 can provide alone. For such applications, the increased hysteresis requirement can be satisfied by matching the LTC4368 with a supply monitor that has adjustable hysteresis, such as the LTC2966. Figure 1 is an example of a wide voltage range automotive circuit protector. In this circuit, the LTC2966 assumes the role of the window comparator and the LTC4368 is responsible for connecting the load to the supply.

Automotive UV/OV and Overcurrent Monitor with Circuit Protection

The solution shown in Figure 1 protects electronics that are sensitive to undervoltage, overvoltage, and overcurrent transients present on an automotive supply.

The LTC2966 monitors reverse voltage, undervoltage, and overvoltage conditions. Monitoring thresholds and hysteresis levels are configured by the resistor networks on the INH and INL pins and the voltages on the RS1 and RS2 pins. OUTA is the UV window comparator output and OUTB is the OV window comparator output. The polarity of these outputs can be selected to be inverting or noninverting with respect to the inputs via the PSA and PSB pins. In Figure 1, they are configured to be noninverting. The OUTA and OUTB outputs from the LTC2966 are pulled up to the REF pin of the LTC2966 and are fed directly to the UV and OV pins of the LTC4368.

The LTC4368 provides reverse current and overcurrent protection. The size of the current sense resistor, R11, determines the reverse current and overcurrent levels. The LTC4368 decides if the load should be connected to the supply based on its overcurrent comparators as well as the monitoring information from the LTC2966. The UV, OV, and SENSE (overcurrent) pins all participate in the decision-making process. If conditions are satisfied for all three pins, then the GATE pin pulls above V_{OUT} and the load will connect to the supply through the dual N-channel MOSFET power path. If any of the three pins become dissatisfied, the GATE pin pulls below V_{OUT} and the load is disconnected from the supply.

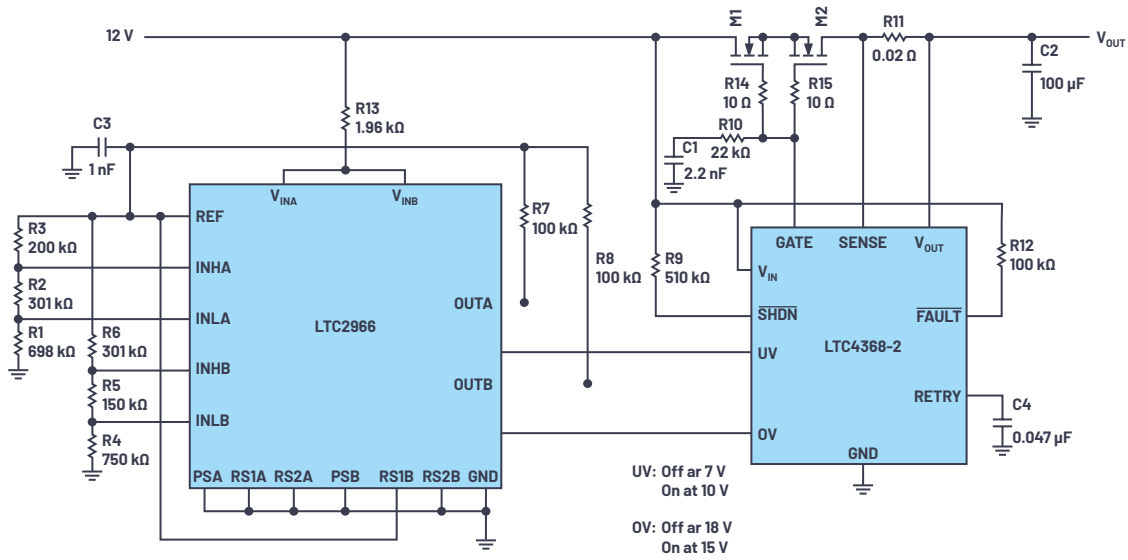

Figure 1. Power path control with wide voltage monitor hysteresis.

Automotive applications that are powered directly from the battery are subjected to large voltage swings during engine start and stop. In this protection solution, voltage monitoring thresholds are based on nominal operating voltages and those expected during automotive cranking or load dump situations, while ensuring that downstream electronics are protected.

Cranking transients are generated when the automobile ignition is energized to start the vehicle. In this application, Channel A of the LTC2966 is configured to detect a cranking transient.

Load dump transients are generated when the automobile is shut off. High amplitude voltage spikes occur on the battery bus when current flow is suddenly stopped in the automobile wiring harness. In this application, Channel B of LTC2966 is configured to detect a load dump transient.

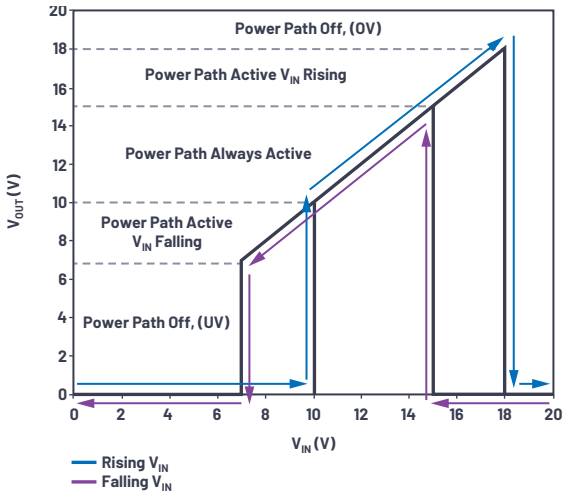

Figure 2. V_{OUT} vs. V_{IN}.

Figure 2 shows the input voltages where the power path is active. The cranking monitor, Channel A, is configured to have a falling voltage threshold of 7 V and a rising threshold of 10 V. The load dump monitor, Channel B, is configured to have a rising threshold of 18 V and a falling threshold of 15 V. These threshold values were obtained by looking at different cranking and load dump waveform specifications. If needed, different thresholds are easily configured by adjusting the resistive divider string for the INL and INH inputs of the LTC2966.

Configuration

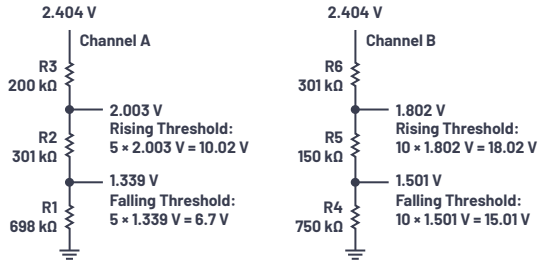

Figure 3. Resistive divider determination for voltage monitor thresholds.

Figure 3 shows how the resistive divider values were calculated for this application. The REF pin of the LTC2966 supplies the 2.404 V.

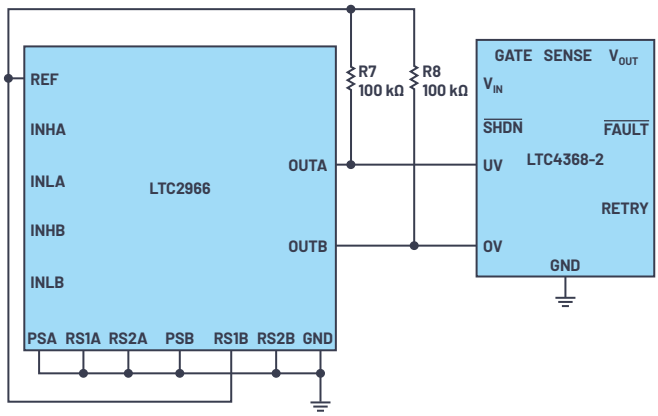

Figure 4. Range and comparator output polarity selection.

Figure 4 shows the range and output polarity configurations for the circuit. The range selections for each channel are based on the voltage range of the particular channel that it is to monitor. Range is configured by the RS1A/B and RS2A/B pins. The polarity of the LTC2966 output pins, whether they pull high or low, is determined by setting the PSA and PSB pins. In this application, the input pins of the LTC4368 determined the polarity of the LTC2966 output pins. For the load to be connected to the supply, the UV pin has to be greater than 0.5 V and the OV pin has to be less than 0.5 V.

Reverse Voltage Protection

In the solution shown in Figure 1, both the LTC2966 and the LTC4368 are protected from reverse voltages: the LTC4368 features built-in reverse supply protection to –40 V, while the LTC2966 requires a component choice.

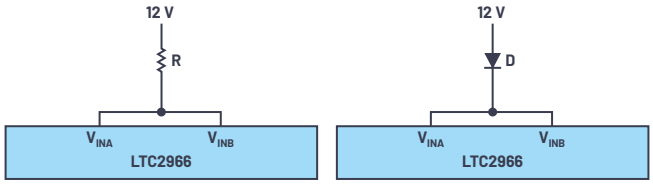

Figure 5. Possible reverse voltage protection methods for LTC2966.

Figure 5 shows two possible reverse voltage protection schemes for the LTC2966—a resistor solution and a diode solution—where choosing between them depends on the application.

In the diode solution, the diode is active only during normal circuit (that is, positive voltage) operation. The supply current for the LTC2966 is in the tens of microamps, meaning a low power diode is sufficient and provides a small footprint solution. During a reverse voltage event, the diode blocks current from coming out of the LTC2966 supply pins. Diode selection is driven by the reverse breakdown voltage of the diode. To match the LTC4368, a 40 V diode should be selected. Consequences of the diode solution are that the forward voltage drop may negatively affect the undervoltage lockout threshold and voltage monitoring threshold accuracies.

In the resistor solution, the resistor is chosen to be large enough to safely limit the current pulled from the LTC2966 supply lines during a reverse voltage event. However, mindful resistor sizing ensures minimal impact on the undervoltage lockout and voltage monitoring threshold accuracies. Correct package size selection ensures safe resistor power dissipation.

For this application, the monitored voltages are low enough that a diode forward voltage in series with the input significantly impairs the accuracies of the voltage monitoring thresholds. Using the resistor solution, a 1.96 kΩ current limit resistor is selected to protect the LTC2966 from a reverse voltage event. The resistor is sized to limit the current out of the input pins to 20 mA if the input voltage is pulled down to –40 V. The low value resistor results in only a several millivolt drop, so the impact on threshold accuracy from the resistor is negligible.

Overcurrent and Inrush Current Protection

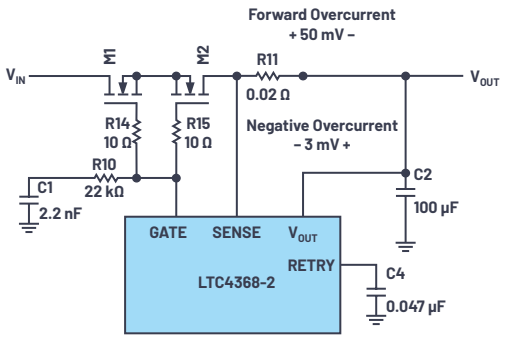

Figure 6. Application overcurrent and inrush current protection.

The LTC4368 is responsible for overcurrent and inrush current protection for the application. Figure 6 shows the responsible components. Comparators inside the LTC4368 monitor the voltage drop across current sense resistor R11.

In the forward direction, V_{IN} to V_{OUT}, the overcurrent comparator will trip when the SENSE to VOUT voltage exceeds 50 mV. In the negative direction, V_{OUT} to V_{IN}, the overcurrent comparator will trip when the SENSE to VOUT voltage exceeds –3 mV. This application utilizes a 20 mΩ sense resistor that sets the current limits to +2.5 A and –150 mA.

Inrush current limiting allows the application to power up without asserting the forward overcurrent protection. R10 and C1 are the inrush current limiting components.

For this application, inrush current is limited to 1 A, well below the forward current limit of 2.5 A. C1 selection is based on the desired inrush current limit and the size of C2. R10 prevents C1 from slowing down the reverse polarity protection, stabilizes the fast pull-down circuits, and prevents chatter during a fault condition.

C4 is the capacitor that sets the retry delay after a positive overcurrent event. The retry delay is the amount of time the MOSFET gate is held low after an overcurrent event is detected. In this application, the retry delay is 250 ms. The 10 Ω resistors R14 and R15 are added to the MOSFET gates to prevent circuit oscillations caused by PCB layout parasitics.

Functionality Demonstration

Cranking Events

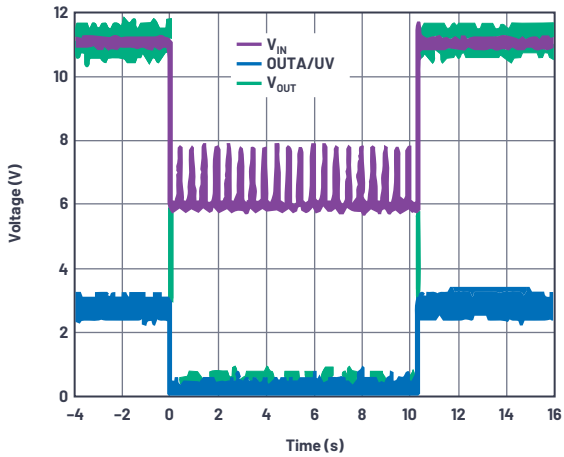

Figure 7. Complete cranking waveform.

Bench characterization of the prototype was conducted and the results are shown in Figure 7. Prior to ignition activation, V_{IN} is greater than the 10 V rising monitor threshold configured for Channel A. The LTC4368-2 UV pin is pulled above its 500 mV threshold by the OUTA pin of the LTC2966, allowing the power path to become active and $V_{OUT} = V_{IN}$.

During a cranking event, the 12 V bus is pulled down to 6 V. The 7 V falling voltage monitor threshold is crossed and OUTA immediately pulls down the UV pin of the LTC4368-2. The LTC4368-2 responds to this by pulling its GATE pin low, which de-energizes the switching element and V_{OUT} falls to 0 V. The 3 V hysteresis programmed by the voltage monitor resistive divider allows the LTC2966 to ignore the ripple on the bus during cranking. As a result, the switching element remains off until the cranking cycle completes. When the cranking cycle completes, the battery voltage returns to its nominal value, which is greater than the 10 V threshold. The OUTA pin pulls the LTC4368-2 UV pin high and the switching element is re-energized.

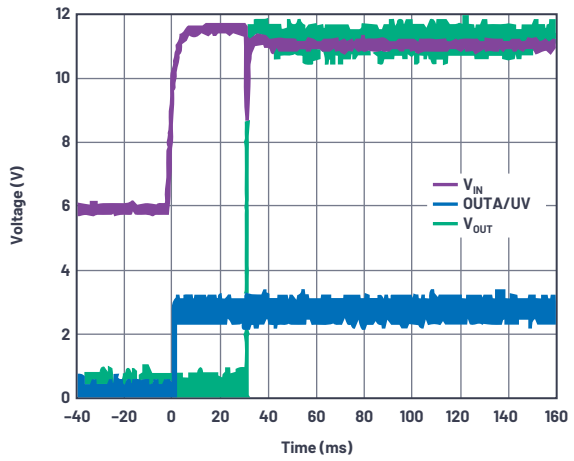

Figure 8. Expanded cranking recovery.

Figure 8 shows the cranking recovery behavior. It shows the LTC4368-2's internal recovery timer (36 ms typical) that is satisfied before the switching element is re-energized. Also observe that once the switching element is re-energized, V_{IN} is momentarily pulled low. This is due to charging the circuit's load capacitance and series input inductance. This demonstrates the need for wide voltage monitoring threshold hysteresis. This load capacitor charging transient is ignored by the LTC2966.

Load Dump Events

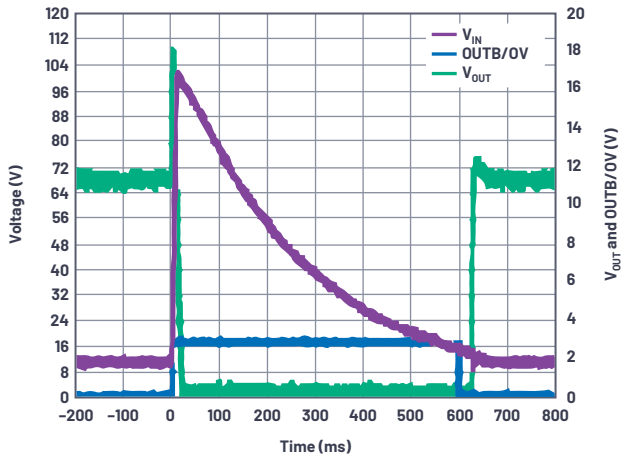

Figure 9. Complete load dump waveform.

Figure 9 shows the circuit's load dump behavior. Prior to ignition deactivation, V_{IN} is at its nominal value. The power path is active and $V_{OUT} = V_{IN}$. During the load dump event, the battery voltage pulls up to 100 V. The 18 V rising voltage monitor threshold is crossed and OUTB immediately pulls up the 0 V pin of the LTC4368-2. The LTC4368-2 responds to this by pulling its GATE pin low, which opens the power path and V_{OUT} falls to 0 V. The switching element remains open until the load dump discharges down to 15 V. Once the 15 V falling threshold is crossed, OUTB of the LTC2966 pulls down the 0 V pin of the LTC4368-2 and after the LTC4368-2 internal recovery timer expires, the LTC4368-2 energizes the switching element again.

Reverse Voltage Protection

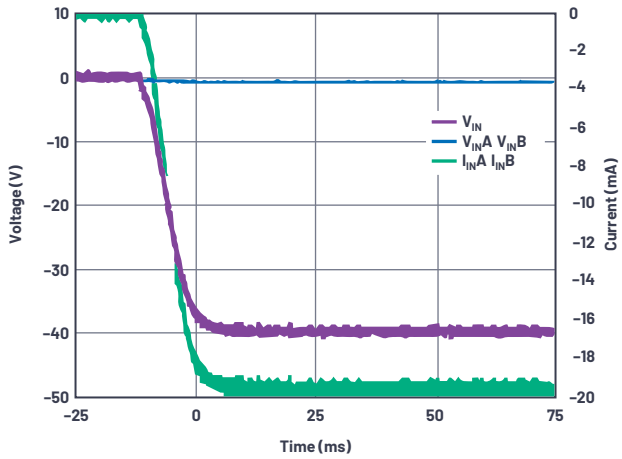

Figure 10. Reverse voltage protection measurement.

Figure 10 shows the 1.96 kΩ resistor limiting the current out of the LTC2966 supply pins during a reverse voltage event. The application's input voltage was ramped from 0 V to –40 V. The current out of the V_{INA} and V_{INB} pins is limited to 20 mA and the voltage of the V_{INA} and V_{INB} pins is held to several hundred millivolts below ground. The LTC2966 safely withstands the reverse voltage event.

Forward Overcurrent Protections

Figure 11 shows the inrush current limiting determined by R10 and C1. As expected, inrush current is limited to 1 A and V_{OUT} pulls up to 12 V cleanly without asserting the overcurrent limit.

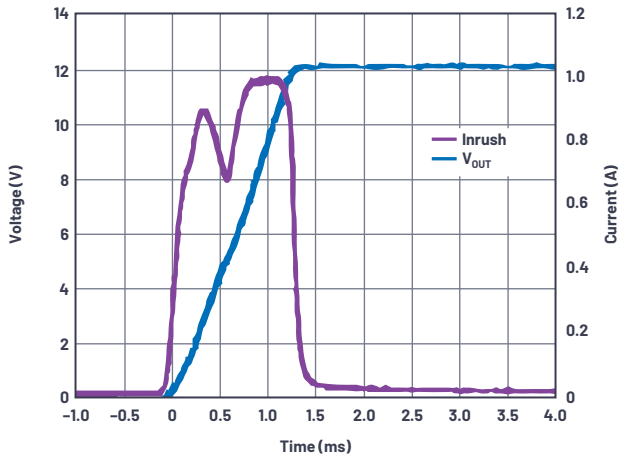

Figure 11. Inrush current limiting.

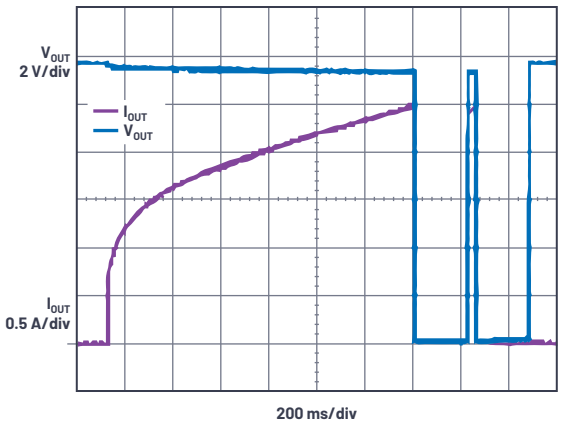

Figure 12. Assertion of forward overcurrent protection and retry delay.

Figure 12 shows the LTC4368 response to a positive overcurrent event. The positive overcurrent comparator in the LTC4368 trips when the voltage between the SENSE and V_{OUT} pins exceeds 50 mV. Current sense resistor R11 is 20 mΩ, which sets the current limit in the application to 2.5 A.

In this demonstration, the current is ramped up until the overcurrent protection asserts. As expected, the overcurrent protection activates at 2.5 A. The LTC4368 removes the load from the supply VOUT and load current drops to 0 V. After the LTC4368 retry timer is satisfied, the LTC4368 reconnects the supply to the load. If the overcurrent condition is gone, then the load remains connected to the supply. Otherwise the LTC4368 removes the load from the supply. The amount of retry delay can be increased by adding capacitance to the RETRY pin. If desired, V_{OUT} can be latched off by grounding the RETRY pin. In this circuit, the retry timer is set for 250 ms. The retry timer configuration is explained in the LTC4368 data sheet.

Reverse Overcurrent Protection

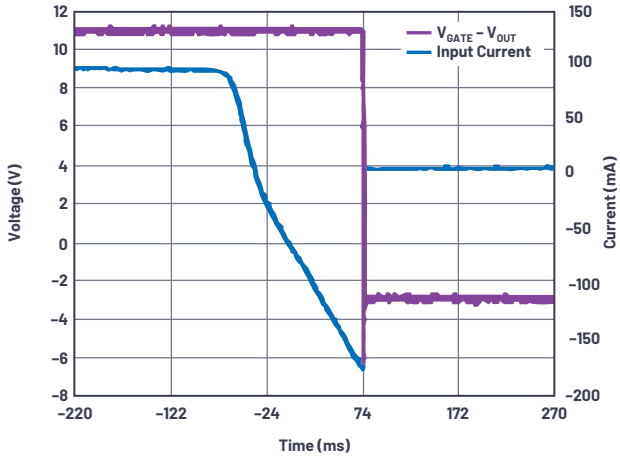

Figure 13. Assertion of reverse overcurrent protection.

Figure 13 shows the LTC4368 response to a reverse overcurrent transient. The reverse overcurrent comparator senses the voltage between the V_{OUT} and SENSE pins. The voltage threshold for reverse overcurrent assertion is version dependent. The LTC4368-1 will assert at 50 mV and the LTC4368-2 will assert at 3 mV. This application is designed with the LTC4368-2 version. Current sense resistor R11 is 20 mΩ. This sets the reverse overcurrent limit to 150 mA.

In this example, while the supply provides 100 mA to the load, a voltage step is introduced to V_{OUT}, so that V_{OUT} is larger than V_{IN}. As V_{OUT} increases, I_{LOAD} decreases. The voltage step is large enough to force current to flow from the load to the supply. This continues until the reverse current reaches 150 mA and the reverse overcurrent comparator trips. When the reverse overcurrent comparator trips, the GATE pin is pulled low. This removes the load from the supply and prevents the load from farther back driving the supply. The LTC4368 will hold the gate low until it senses V_{OUT} drop 100 mV below V_{IN}.

Conclusion

The automotive application developed in this article demonstrates that the use of specialized protection devices can simplify the implementation of automotive protection circuitry. With minimal additional circuitry, the LTC2966 and LTC4368-2 were combined to provide accurate, robust, and comprehensive transient protection. The flexibility of the devices allows them to be configured for use in numerous applications.

About the Author

Al Hinckley received his B.S.E.E. from Merrimack College, and later a Graduate Certificate in VLSI and microelectronics from UMass Lowell. He joined Linear Technology (now part of Analog Devices) in April 2005. He can be reached at albert.hinckley@analog.com.

Notes

Stay Connected with *Analog Dialogue*

Your Engineering Resource for Innovative Design

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

EngineerZone® Support Community

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

Facebook
facebook.com/AnalogDevicesInc

@ADI_News
twitter.com/ADI_News

LinkedIn
linkedin.com/company/analog-devices

Sina Weibo (China)
weibo.com/analogdevices

WeChat (China)

Analog Devices, Inc.
Worldwide Headquarters

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

Analog Devices, Inc.
Europe Headquarters

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

Analog Devices, Inc.
Japan Headquarters

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

Analog Devices, Inc.
Asia Pacific Headquarters

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

M02000-8/20

analog.com

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™