Analog Dialogue

Volume 54, Number 1, 2020

Your Engineering Resource for Innovative Design

Visit *analogdialogue.com*

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

In This Issue

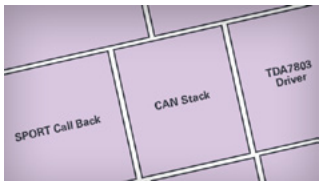

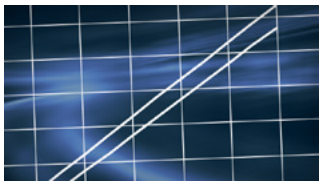

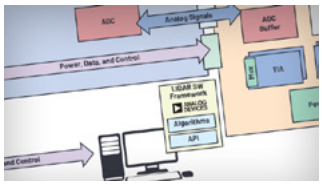

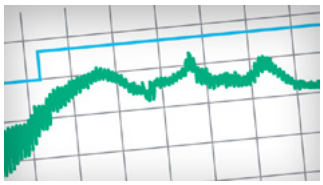

Bernhard Siegel, Editor

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

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

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

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

AnalogDialogue

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

Electric Vehicle Warning Sound System

Andreas Pellkofer, Applications Engineer,
Jagannath Rotti, Software Lead,
and **Danny Ko**, Applications Engineer

Introduction

Traditional combustion engine vehicles emit engine sound, even at low travel speeds. Typically, pedestrians and other traffic participants recognize an approaching or departing vehicle through sight and auditory identification of tire sounds and other emitted noise when the vehicle is out of sight.

Electric vehicles (EVs) do not emit engine sound. Hybrid electric vehicles (HEVs) or plug-in hybrid electric vehicles (PHEVs) move almost silently when traveling at low speeds and before the conventional internal combustion engine (ICE) kicks in. These vehicles are difficult to hear when travelling at speeds less than 19 mph. At greater speeds, tire sound becomes dominant.

Global governing bodies are exploring legislation that seeks to establish a minimum level of sound for PHEVs and HEVs when operating in electric mode so that visually impaired people, pedestrians, and cyclists can hear these vehicles approach and determine from which direction these vehicles are approaching. An example of this legislation can be found on the National Highway Traffic Safety Administration (NHTSA) website.

An electric vehicle warning sound system (EVWSS) produces a series of sounds designed to alert pedestrians to the presence of EVs, HEVs, and PHEVs. The driver can initiate warning sounds (similar to the sound from a car horn, but less urgent); however, the sounds must automatically be enabled at low speeds. These sounds vary from artificial tones to realistic sounds that mimic engine noise and tires moving over gravel.

Analog Devices offers two different solutions for advanced applications with an in-cabin engine sound for an EV as well as an external engine sound. Analog Devices developed a solution based on the ADSP-BF706. For entry-level systems, Analog Devices developed a solution based on the ADAU1450 SigmaDSP®. These solutions can synthesize sound and adjust frequency, sound volume, and other parameters depending on the traveling speed, and these solutions can send the audio to an audio power amplifier. Depending on the requirements of certain legislation, the warning sound can be simulated using combustion engine sounds or any other synthesized tones.

Blackfin-Based Solution

The ADSP-BF706 Blackfin+® processor provides a single-chip solution for audio processing and interfacing to the control area network (CAN) bus. Analog Devices developed a CAN software stack that runs on the ADSP-BF706, which enables users to build automotive grade demonstrations with minimal effort (a Vector® CAN stack can also be used). Additionally, Analog Devices provides full hardware and software reference design and SigmaStudio® compatibility for live tuning of parameters.

Figure 1 shows the different processing blocks inside the ADSP-BF706. External waveform audio files (WAVs) store signature engine sounds or audio tones. Up to 25 WAV files can be accessed simultaneously from the external serial peripheral interface (SPI). These files are frequency shifted and mixed internally in the digital signal processor (DSP) before adding the dynamic volume control.

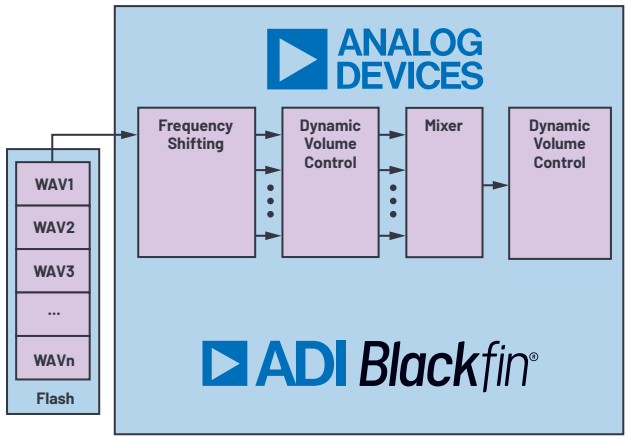

Figure 1. Processing blocks on Blackfin+ processor.

The ADSP-BF706 utilizes a memory mapped SPI interface that provides faster, simplified access to the external memory, which eliminates the need for an external double data rate (DDR) memory for this application. Up to 25 WAV files can be accessed simultaneously from the SPI flash memory. The large number of accessible WAV files helps to create more realistic engine sounds.

The ADSP-BF706 can also implement up to 16× pitch shifting, a recommendation from the United States NHTSA, which increases the frequency of the output sound as the vehicle speed increases. The ADSP-BF706 can dynamically control the volume as the vehicle speed from the CAN bus increases.

Figure 2 shows a detailed system block diagram. A Power by Linear LT8602 quad monolithic synchronous, step-down regulator provides all voltage rails required in the system, derived from the 12 V car battery supply. The 2 MHz switching frequency allows users to avoid critical, noise sensitive frequency bands—for example, the AM band. The 3 V to 42 V input voltage range of the LT8602 makes the device ideal for automotive applications, which must regulate through cold crank and start stop scenarios with minimum input voltages as low as 3 V and load dump transients more than 40 V.

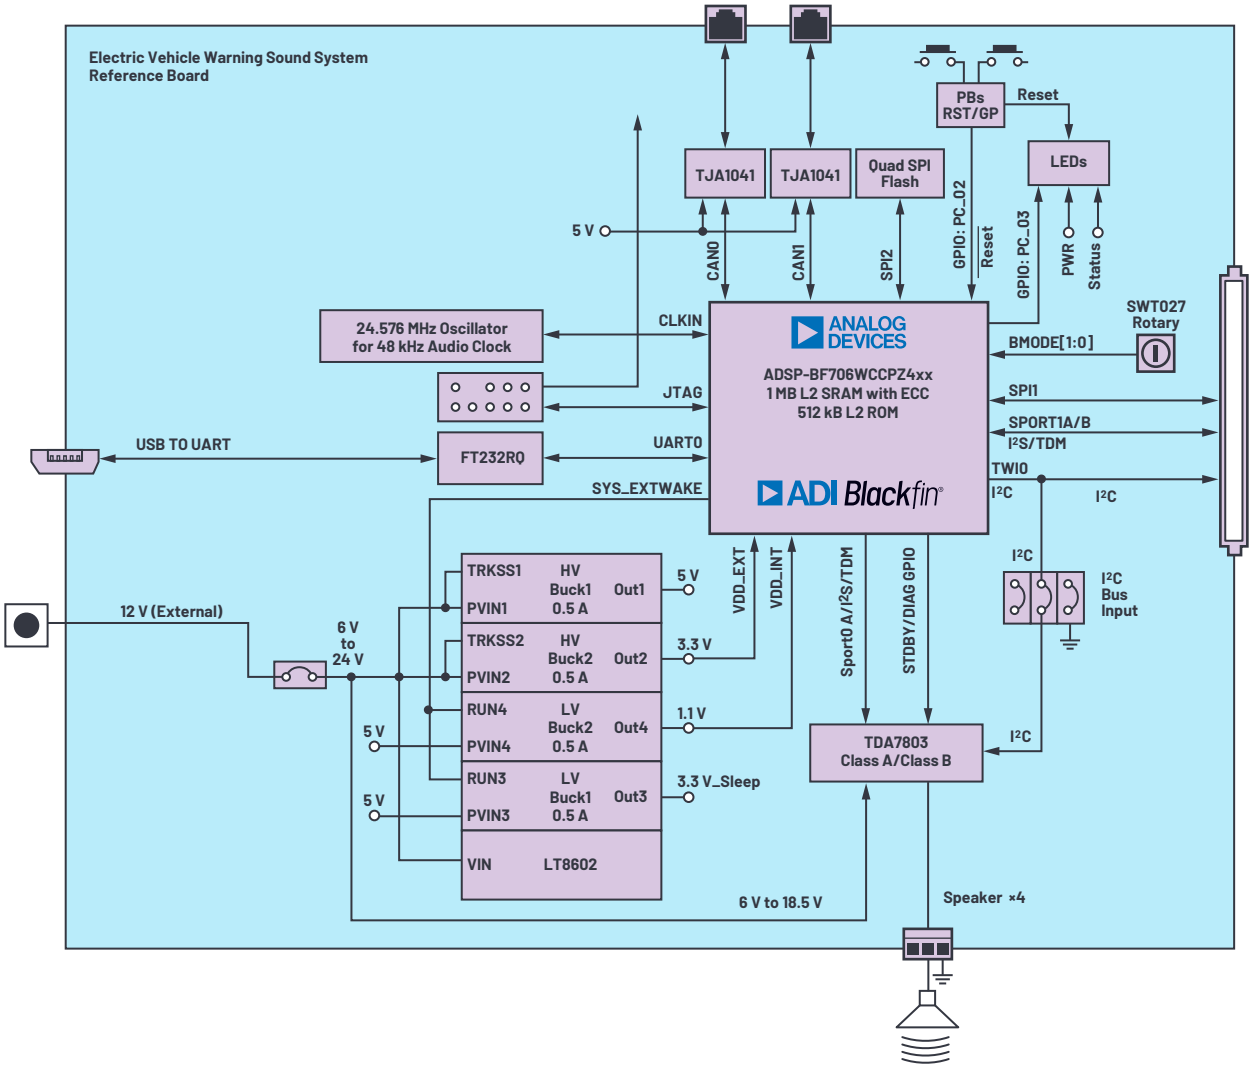

Figure 2. Detailed system block diagram with a Blackfin+ processor on a full featured board.

Figure 3 shows an alternative system block diagram with connectors, a reduced set of peripherals, and one automotive qualified connector carrying all relevant signals. This configuration allows the design of a smaller form factor board.

This system solution results in a reduced system bill of materials (BOM) cost because the ADSP-BF706 acts as a microcontroller and an audio processor.

For full details on the solution, see the EVWSS v1 Demo Manual and the EVWSS v2 Demo Manual, which are provided in the software download package. This software package (EVWSS-BF_SRC-Rel2.0.0) is available by request from the Software Request Form page on the Analog Devices website. For full details on the ADSP-BF706, see the ADSP-BF70x Blackfin+™ Processory Hardware Reference and the ADSP-BF7xx Blackfin+™ Processor Programming Reference.

EVWSS Software Architecture for the ADSP-BF706 Blackfin+ Processors

The EVWSS software architecture is based on the ADSP-BF706 hardware architecture. The processor dependence on the hardware architecture is due to the memory mapped SPI. The CAN interface reads directly from the flash memory using memory mapped SPI. This feature reduces the complexity of the EVWSS library and makes the memory access efficient for warning sound generation.

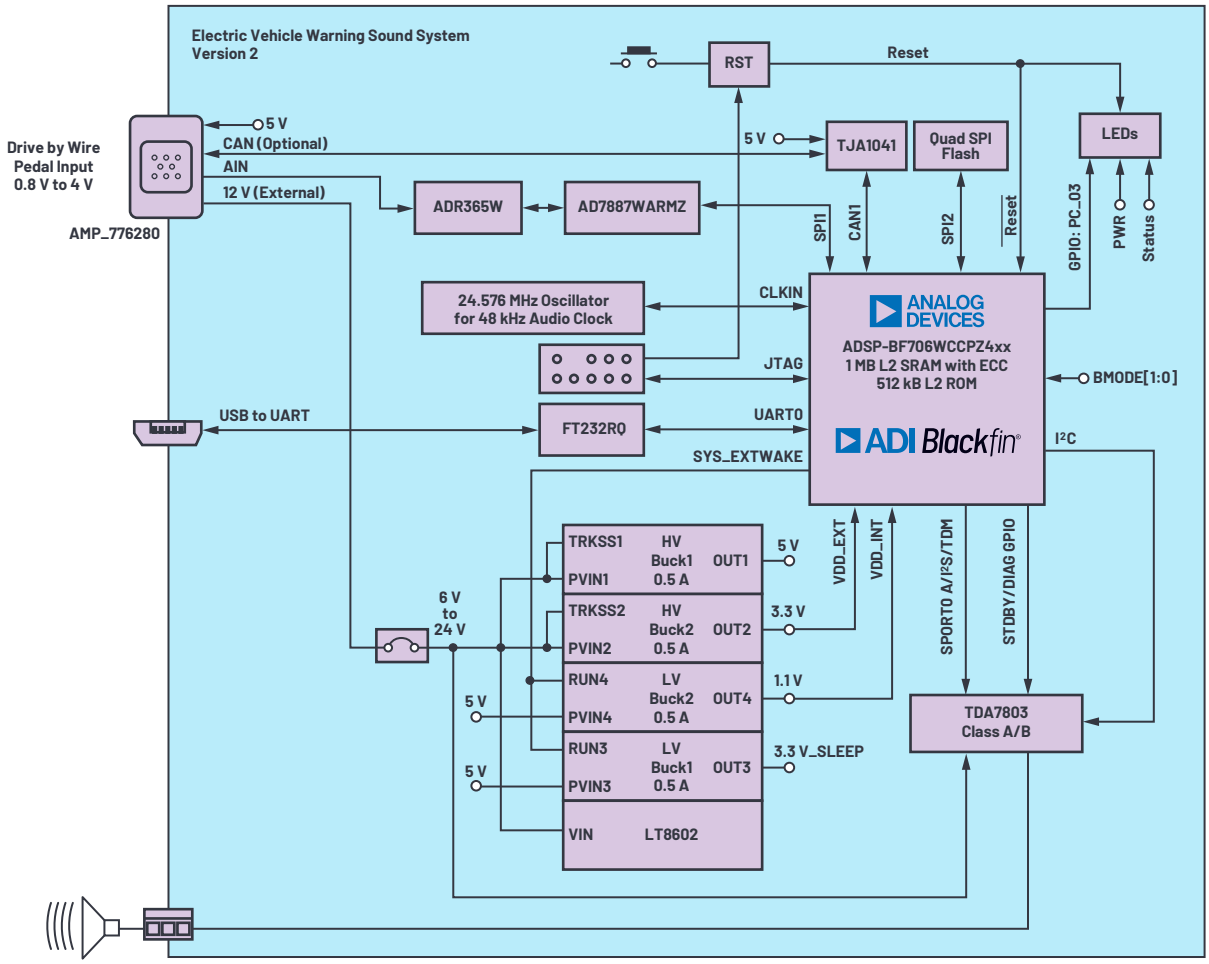

Figure 3. Detailed system block diagram with a Blackfin+ processor on a board with reduced components.

Software Components

The EVWSS software architecture consists of the components shown in Figure 4.

This section details the software components. The SPORT callback feature maps to the audio data sample rate and runs in the SPORT transceiver interrupt service routine (ISR) context, reading flash files (SPI memory mapped), performing audio manipulation using the EVWSS library, and sending out modified audio on the SPORT transceiver interface. The EVWSS library holds the different functions to synthesize the warning sound. The EVWSS library also receives vehicle speed input from the CAN stack (or the universal asynchronous receiver and transmitter (UART) interface for debugging). The TDA7803 driver controls the external power amplifier to generate the warning sound. The EVWSS application framework configures the system peripherals, CAN stack, and the TDA7803 driver.

EVWSS Library Functions

The following sections describe the functions of the EVWSS library. For full details, see the Electronic Vehicle Warning Sound System Release Notes, which can be found in the software download package.

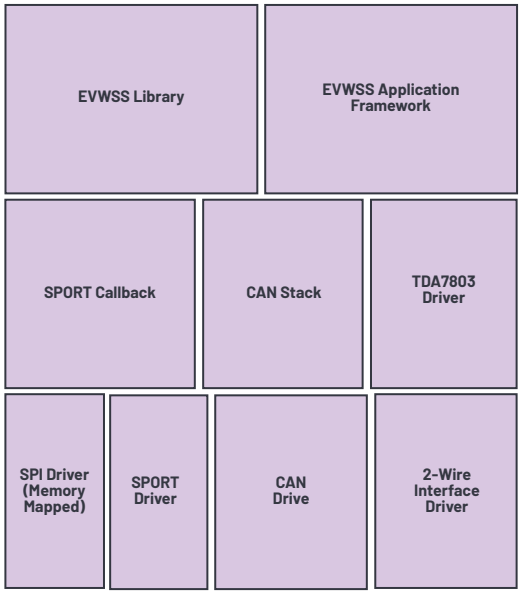

Figure 4. EVWSS software architecture block diagram.

Pitch Control

Pitch shifting is the concept of shifting the spectrum of the audio signal based on a control input. In the EVWSS application, the base pitch of the WAV file is shifted depending on the vehicle speed input.

Frequency Modulation and Amplitude Modulation

The engine sound is dependent on the engine strokes, which include intake, compression, power (expansion), and exhaust. These strokes create frequency modulated tones instead of just pure tones. Vary the pitch shift parameter across the samples to achieve frequency modulation.

Two kinds of modulation (sawtooth and triangular) are included in this application. In sawtooth modulation, the frequency ramps from lowest to highest and then back to lowest with a jump. In triangular modulation, the frequency ramps from lowest to highest and then ramps back down to lowest.

Audio Mixing in Slew

For audio mixing, configure the various gains with respect to vehicle speed.

Playback of WAV Files

Although the required WAV files are present in flash, the user can play or stop some of the WAV files, depending on the dynamic conditions.

SigmaDSP-Based Solution

For entry-level applications, an ADAU1450 SigmaDSP processor can be used as an alternative to the ADSP-BF706 processor. For evaluation purposes, the EVAL-ADAU1452 evaluation board can be used.

Figure 5 shows the different processing blocks inside the SigmaDSP processor.

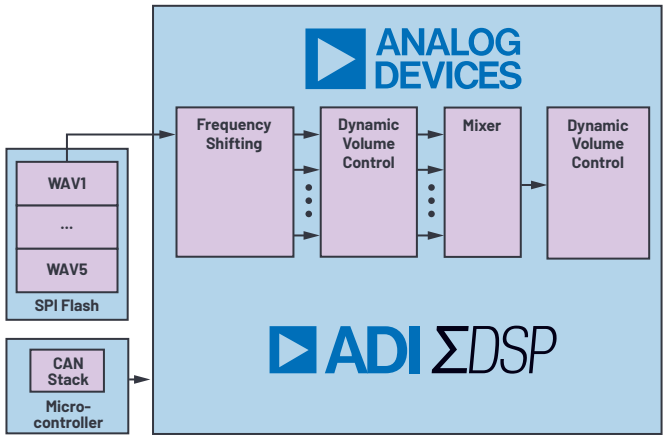

Figure 5. Processing blocks on SigmaDSP processor.

The following software requirements are supported in the ADAU1450 using the SigmaStudio programming environment:

- Multiple tone generation
- Dynamic volume control with up to 64 ranges
- Sound mixing
- Limiter
- Pitch shifting that increases audio pitch as speed increases
- Simultaneous playback of up to 5 WAV files from the SPI flash memory

Analog Devices provides an engine sound simulator module in SigmaStudio to simplify engine sound tuning and reduce the number of external, simultaneous WAV files required. The engine sound simulator can internally generate up to 32 harmonics. The order and amplitude of these harmonics can be programmed through the graphical user interface (GUI).

The harmonic generator module is available as part of the ESS Toolbox. It can be downloaded directly from the SigmaStudio download page (Rev 4.4. upward). Note that SigmaStudio cannot support a CAN software stack, and an external microprocessor is required.

SigmaStudio

SigmaStudio is a graphical programming environment originally designed for the SigmaDSP portfolio of processors. This software has a built in library of algorithms developed specifically for automotive applications. The GUI simplifies the tuning process and provides controls and filter coefficients that can be varied on the fly without the requirement for writing code. SigmaStudio can be downloaded from the SigmaStudio page on the Analog Devices website.

Conclusion

Analog Devices offers comprehensive solutions for an entry-level system and for an advanced engine sound system that supports in-cabin engine sounds and external engine sounds. This article aims to ease the decision making process for the user, as well as reduce user time to market. Analog Devices offers a complete system solution, including the necessary software components for rapid prototyping and product development.

About the Author

Andreas Pellkofer graduated from Technical University of Munich in electronic and information technology. He joined Analog Devices 2006 as an applications engineer working with the Blackfin processor family. Later he mainly dealt with automotive customers for DSPs. In 2013, he moved into a system engineer role in the Digital Video Products Group, focusing on video transport and camera systems for automotive. In 2018, he became part of the Emerging Systems and Technologies Group working on in-cabin vital sign monitoring and hands-on detection solutions. He can be reached at *andreas.pellkofer@analog.com*.

About the Author

Jagannath Rotti graduated from PES Institute of Technology, Bangalore with a degree in electronics and communication. He has 13 years of automotive software experience. Prior to joining Analog Devices, he worked at Robert Bosch and Autoliv in the powertrain and safety domains, respectively. At ADI, he is a software lead in the Automotive SW Team, mainly working in cabin electronics in general and automotive audio buses. His areas of interest include automotive networks, network security and cryptography, audio algorithms, autonomous driving, sensor fusion, and Sanskrit literature. He can be reached at *jagannath.rotti@analog.com*.

About the Author

Danny Ko is an automotive system applications engineer for audio and emerging technologies based in Seoul, Korea. Danny joined ADI in 2004 as a DSP FAE and supported Samsung, LG, and broad market for three years before changing his focus to automotive in 2007. In 2010, Danny transferred to the automotive segment as an automotive system application engineer and worked in the infotainment area, primarily in audio applications. Since 2018, his work has extended to emerging technology. He can be reached at *danny.ko@analog.com*.

ADI SDR Transceivers Enable Amateur Space Communication

Diego Koch, Applications Engineer

Radio amateurs recently received another way to provide uninterrupted worldwide radio coverage. By means of a new geostationary satellite, it is now possible to reliably cover one-third of the Earth in just one hop. In order to contact the satellite, it is necessary to use dedicated equipment because the access frequencies are different from the ones used to bounce radio signals from the ionosphere. The new software-defined radio (SDR) approach to radio transceivers offers multiple advantages, such as flexible reconfiguration and the capability to observe the whole band of interest at a glance, just to name a few.

This article begins with an overview of this satellite, the story behind it, the areas covered, and how it can be accessed. Then the realization of a practical radio station by using the ADALM-PLUTO SDR, based on one of the ADI SDR transceivers, will be presented.

The Satellite

Launched in 2018 from Cape Canaveral, the Es'hail-2 communication satellite of the Qatar satellite company Es'hailSat provides television, voice, internet, corporate, and government communication services across Europe, the Middle East, Africa, and beyond. It has been operational since February 2019 and has been positioned above central Africa in a geostationary orbit. From a height of 36,000 km, it covers an area spanning from Brazil to Malaysia, from the Faroe Islands to Antarctica, as shown in Figure 1.

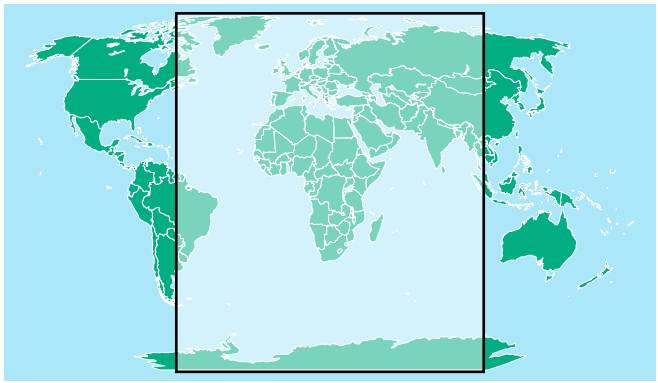

Figure 1. The Earth seen by Es'hail-2.

Es'hailSat was established in 2010. Based in Doha, Qatar, the company owns and operates satellites to serve broadcasters, businesses, and governments. In order to promote and foster space technology development in Qatar, Es'hailSat initiated the development of new technology for the Qatar Amateur Radio Society (QARS), a national nonprofit organization for amateur radio enthusiasts, together in partnership with Amateur Satellite Corporation (AMSAT), another global nonprofit organization. AMSAT designs, builds, arranges, launches, and operates satellites carrying amateur radio payloads. AMSAT affiliated national organizations exist in various countries, including AMSAT Germany (AMSAT-DL), which became involved on behalf of QARS in December 2012. This collaboration has made it possible to equip the Es'hail-2 satellite with two dedicated transponders, providing the first amateur radio geostationary communication capability that connects users across the visible globe in a single hop and in real time.

Many amateur satellites receive an OSCAR (orbiting satellite carrying amateur radio) designation. These satellites can be used free of charge by licensed amateur radio operators for voice and data communications. So far, they have been launched into low Earth orbits (LEOs) and into highly elliptical orbits (HEOs), and what all of them have in common is that it is necessary to track them with antennas when they appear above the horizon for just a few minutes. Once they disappear below the horizon, the communication is no longer possible. Satellites on a geostationary orbit have the advantage that, as observed from Earth, their positions do not shift in the sky. Although the antennas on Earth do not have to move to access them, the big distance of 36,000 km sets new challenges in terms of free space power loss, antenna pointing accuracy, and latency—about 250 ms for a trip from one ground-based transmitter to the satellite and back to another ground-based transmitter. The nickname given to Es'hail-2 is OSCAR100 because it is the one hundredth satellite to carry an amateur radio payload.

Access to Es'hail-2

Radio amateurs have worked with satellites for many years. Traditionally, this has been done using analog downconverters and upconverters that shift the received and transmitted signals to and from the amateur bands where transceivers operate. The uplink (from Earth to satellite) and downlink (from satellite to Earth) frequencies used by satellites are sometimes beyond the capabilities of available transceivers. Es'hail-2 has two transponders: one for narrow-band (NB) transmissions and one for wideband (WB) transmissions. In this section, we will talk about the narrow-band transponder. Since, on this transponder, the available bandwidth is only 250 kHz, to accommodate multiple channels it is necessary to use appropriate modulation techniques. The types of analog modulation that are most commonly used are telegraphy (Morse code, also called continuous-wave (CW)) or telephony (voice, also called single sideband (SSB)).

Uplink is at 2.4 GHz (13 cm band) in right hand circular polarization (RHCP) and downlink is at 10.45 GHz (3 cm band) in either horizontal (H) or vertical (V) polarization. Radio amateurs have the privilege to transmit in the 13 cm band (2300 MHz to 2310 MHz and 2390 MHz to 2450 MHz) as licensed radio operators for satellite communication with sufficient power and high gain antennas. This band overlaps with the civilian radio allocation 2400 MHz to 2500 MHz, which is part of the industrial, scientific, and medical (ISM) bands. One of the most popular unlicensed emissions on the ISM band is wireless LAN. The transponders are detailed in Figure 2.

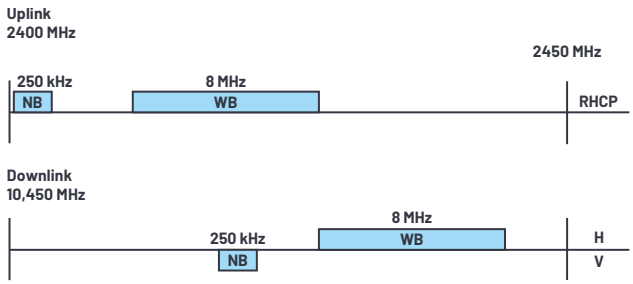

	Uplink (MHz)			Downlink (MHz)			BW (MHz)
	Pol	Begin	End	Pol	Begin	End	
NB	RHCP	2400.05	2400.3	V	10489.55	10489.80	0.25
WB	RHCP	2401.5	2409.5	H	10491	10499	8

Figure 2. Transponders of Es'hail-2.

The Innovative SDR Approach

The change that came with the introduction of SDR in its many variants also impacted the amateur radio world. Even though most transceivers still have the same controls of the older analog ones, many of them incorporate digital signal processor (DSP) technology after the mixer, at the intermediate frequency level. Some of them are also able to directly sample the whole shortwave portion of the spectrum (dc to 30 MHz). One advantage of the SDRs is that their performances do not degrade with time, since many critical analog components are partially replaced by digital algorithms. Another advantage is that the same performances that require expensive components in analog radios, like mixers or filters, can be obtained in a more cost-effective way by complementing them with different elements like analog-to-digital converters (ADCs) and DSPs. The integration of multiple blocks, such as image rejection mixers, oscillators, and ADCs, in the same silicon device has made feasible new receiver architectures that are very critical to implement with discrete technology. One example is devices like the AD9363/AD9364 RF agile transceivers that combine all RF front-end, mixed-signal, and digital blocks in a single device for both receiving and transmitting. When paired with an FPGA that manages the digital data flow into and out from the device, the elements remaining to build a complete station are the antennas, the power amplifier, and the software algorithms running on a computer.

ADI offers the ADALM-PLUTO SDR to demonstrate the capabilities of the AD9363, shown in Figure 3. This is a cost-effective hardware tool that can be used by engineers to develop applications where radio is involved based on the new SDR approach. The AD9363 has a receive and transmit bandwidth of 20 MHz and it can easily receive both the narrow and wide downlink transponders of the Es'hail-2, once they are downconverted externally to its frequency range of 235 MHz to 3.8 GHz. It can transmit on the uplink frequencies without any external upconverter. Another beneficial feature, when compared to devices of the same class and price, is that it has two connectors for receive and transmit, so it supports full-duplex operation. The normal amateur radio interaction is half-duplex (you either talk or listen), but the ability to receive your own transmission in real time allows you to understand whether you are modulating in a clear way, or whether you need to increase/decrease the transmitted power. It also helps to have the ability to point the transmit antenna to the sky once the receive antenna has been adjusted.

The ADALM-PLUTO is supported for both transmission and reception by some free software packages, often written by radio amateurs themselves. One example is the SDR Console by Simon Brown (amateur radio callsign G4ELI). This software manages the interaction between the user and the transceiver, and implements demodulation and modulation in software.

An SDR Satellite Station

Radio amateurs are well known for building their own hardware and repurposing existing equipment to fit their needs. With receive antennas and downconverters, the cheapest alternative is an ordinary satellite dish for commercial satellite television and a low noise block (LNB). The LNB contains the waveguide and the downconverter that translate the incoming downlink signal at 10.450 GHz to less than 1 GHz, which falls inside the receivable band of the SDR. Narrow-band modulation types such as CW (a few tens of Hz) or SSB (less than 3 kHz) mandate highly stable local oscillators to avoid continuous retuning, which is less critical in wideband modulation types such as the ones used by broadcast television (some MHz). In modern digital communications, compensation for frequency offset and long-term drift due to thermal issues is built into the standards and implemented by everyone. Unfortunately, this is not standardized, or not implemented, for many narrow-band modulation schemes implemented by amateur radio operators, and the expectation is that PLL or sample rate accuracy and drift either in the LNB or baseband signals is perfect. To ensure this assumption is correct, sometimes high precision/low drift reference clocks are used. Since many amateur radio operators are more comfortable swapping a reference clock than implementing complex digital signal processing techniques, many will recommend this easy fix.

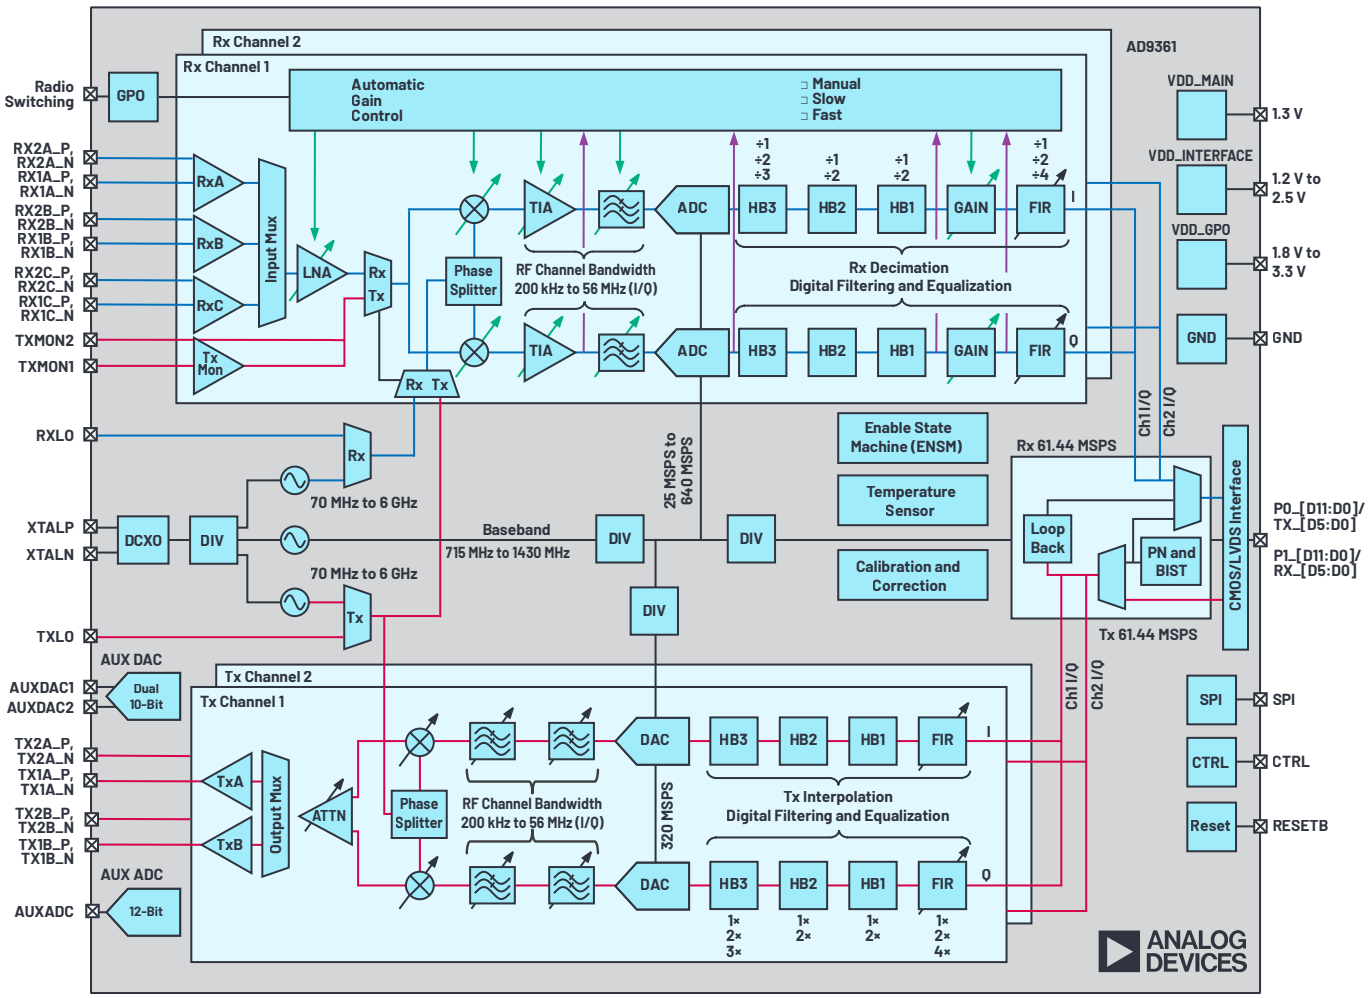

Figure 3. ADALM-PLUTO and its transceiver AD9363.

Since the uplink frequencies are within the WLAN 2.4 GHz band it is possible for licensed operators to repurpose existing WLAN equipment like power amplifiers and high gain antennae. The ADALM-PLUTO has about 5 dBm power output, which is insufficient to drive a power amplifier that has an output power of a few watts. The CN-0417 reference design, based on the ADL5606 20 dB power amplifier and powered by the LTM8045 SEPIC micromodule converter, yields enough power gain to overcome this limitation. Figure 4 shows how a communication station can be laid out. The station can also be rapidly deployed in the field to support emergency communication.

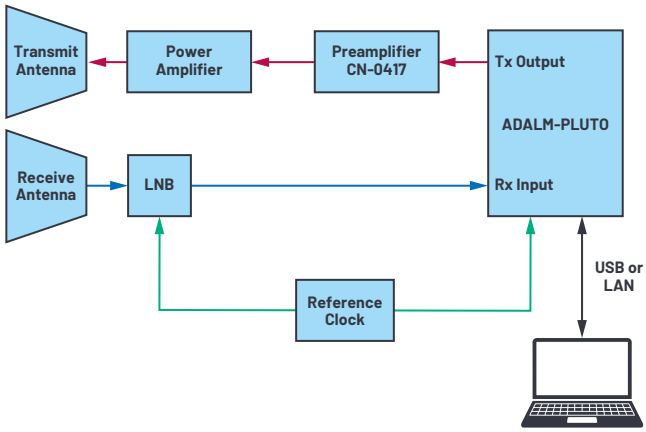

Figure 4. SDR satellite radio station.

Summary

In conclusion, we see a shift toward SDR technology in radio communication. This has been possible by integrating multiple analog and mixed-signal blocks in one device. Immediate advantages are cost-effectiveness, improved reliability, and reconfigurability.

Quoting the words of Drew Glasbrenner, KO4MA, AMSAT VP Operations, "May the 100th OSCAR satellite be the guide star to future amateur radio satellites and payloads to geostationary orbit and beyond."

References

"Es'hail-2/QO-100." AMSAT-UK, 2019.

"Space Communication." American Radio Relay League, 2019.

Taylor, Wyatt and David Brown. "RF Transceivers Provide Breakthrough SWaP Solutions for Aerospace and Defense." *Analog Dialogue*, September 2016.

About the Author

Diego Koch is an applications engineer with the European Centralized Applications Center based in Munich, Bavaria, Germany. He has been working in the semiconductor industry since 1998 and joined Analog Devices in 2017, providing design support for the Power by Linear portfolio for the European broad market. He holds a master's degree in electronics engineering from Politecnico di Milano, Lombardy, Italy. He is a radio amateur, licensed in two countries with callsigns IZ2MZL and DK2MZL. He can be reached at *diego.koch@analog.com*.

Paralleling Very Low Dropout Linear Regulators for Increased Output Current and Even Heat Distribution

Molly Zhu, Senior Applications Engineer
and **Fei Guo**, Field Applications Engineer

Introduction

Each successive generation of computing systems demands more total power and lower supply voltages than its predecessor, challenging power supply designers to keep pace with higher output current in a small area. At high power densities with low output voltages, the problem of heat dissipation rises to the top of the design priority list, especially for linear regulators in low noise applications. Paralleling LDO regulators can increase the supply current capability and mitigate heat dissipation, reducing the temperature rise of any particular component and the required size and number of cooling devices.

This article shows how to parallel the 3 A LT3033 very low dropout (VLDO) regulator for applications that require a current higher than 3 A, with the added benefit of spreading heat dissipation. The LT3033 can be easily paralleled—and current balanced—due to its built-in output current monitoring feature.

The LT3033 converts 1.14 V to 10 V input supplies to outputs as low as 0.2 V with load currents up to 3 A. The dropout voltage is only 95 mV at full load. Quiescent current is 1.8 mA during operation, dropping to 22 µA when shutdown. Programmable current limit and thermal protection give it the necessary robustness for high current, low voltage applications.

For Reference: 3 A, Single VLDO Application

Figure 1 shows the LT3033 delivering 0.9 V at 3 A from a 1.2 V input supply. A minimum of 10 µF very low ESR ceramic capacitor is required at the IN and OUT pins for stability. Adding a feedforward capacitor (C_{FF}) from V_{OUT} to the ADJ pin can improve the transient response and lower the output voltage noise. A 10 nF bypass capacitor from the REF/BYP pin to GND typically reduces output voltage noise to 60 µV rms in a 10 Hz to 100 kHz bandwidth and soft starts the reference. The minimum input voltage required for regulation equals the regulated output voltage V_{OUT} plus the dropout voltage or 1.14 V, whichever is greater. A demonstration board is shown in Figure 2.

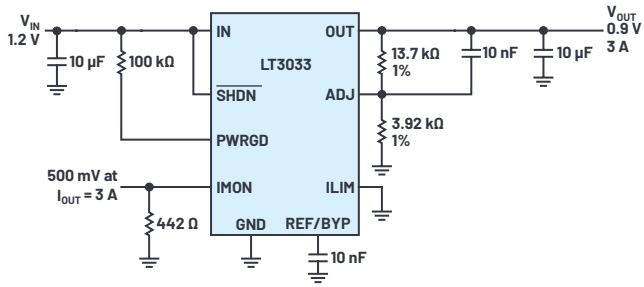

Figure 1. LT3033 typical application.

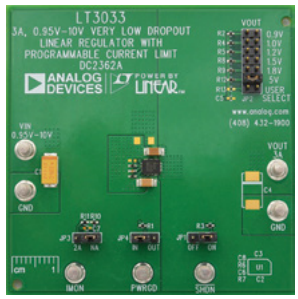

Figure 2. LT3033 demonstration board.

Current limit is programmable by connecting a single resistor from the ILIM pin to GND, accurate to ±12% over a wide temperature range. The external current limit may be overridden by internal current limit with foldback when the input and output differential voltage exceeds 5 V.

LT3033 provides an output current monitor by pulling the IMON pin to GND through a resistor. The IMON pin is the collector of a PNP, which mirrors the LT3033 output PNP at a ratio of 1:2650. The resistor voltage is proportional to the output current if it is not higher than V_{OUT} – 400 mV.

$$I_{OUT} = 2650 \times (V_{IMON}/R_{IMON}) \quad (1)$$

This output current monitor enables current sharing for multiple LT3033s.

Despite its tiny size, the LT3033 also includes a number of protection features, including internal current limit with foldback, thermal limit, reverse current, and reverse-battery protections.

Using Two LT3033s in Parallel for a 6 A Application

For applications that require more than 3 A, multiple LT3033s can be paralleled by taking advantage of its current monitor feature. Figure 3 illustrates two LT3033s in parallel together with two 2N3904 NPN devices to generate a 1.5 V, 6 A output. The individual IN pins and OUT pins are tied together, respectively. One LT3033 acts as master, controlling the LT3033 slave device.

The IMON pins combined with an NPN current mirror create a simple amplifier. This amplifier injects current into the feedback divider of the slave LT3033 to force the IMON currents from each LT3033 to be equal. The 100 Ω resistors provide 113 mV emitter degeneration at full load to guarantee good current mirror matching. The output voltage of the slave LT3033 is set at 1.35 V, 10% lower than the circuit output to ensure the master LT3033 stays in control. The feedback resistors of the slave LT3033 are split into sections to ensure adequate headroom for the slave NPN. A 10 nF, 5.1 kΩ capacitor and resistor combination added to the IMON pin of the slave device frequency compensates the feedback loop.

Though this circuit can supply 6 A load current, current sharing accuracy is limited by the mismatch between two NPN devices—mismatches result in uneven heat distribution on the board. Higher current sharing accuracy can be achieved by using a matched monolithic transistor, such as the MAT14 from Analog Devices, to replace the two discrete NPN devices. The MAT14 is a quad monolithic NPN transistor that offers excellent parametric matching. Its maximum current gain matching is 4%.

Figure 4 compares the output current of each LDO regulator using discrete and matched NPN devices, respectively. Compared to the 2N3904, the MAT14 current mirror reduces the current mismatch from 5.3% to 1.6%.

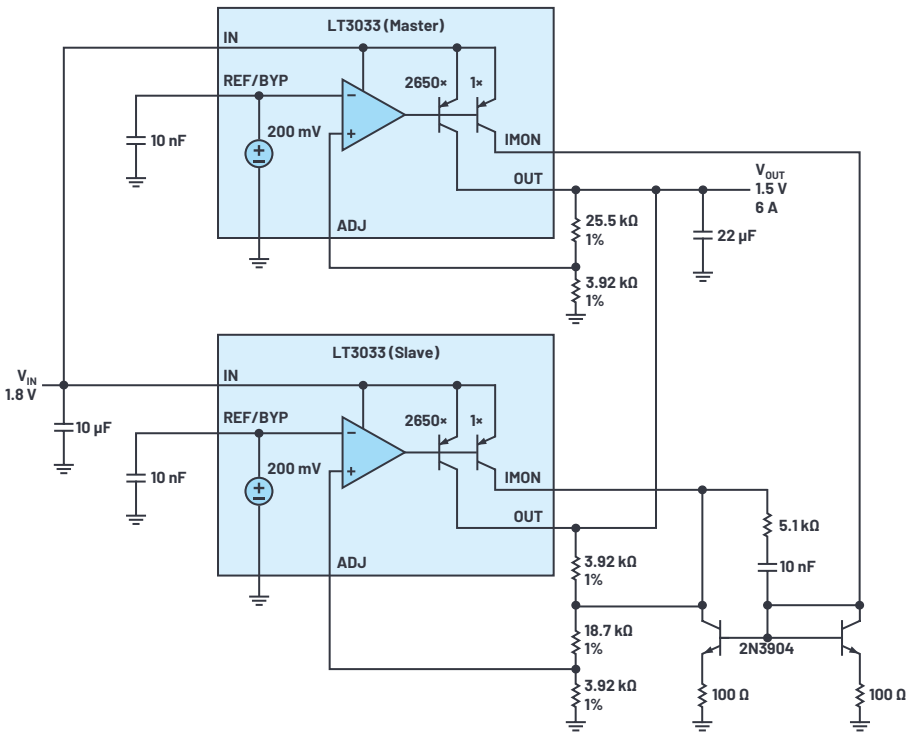

Figure 3. Paralleling two LT3033s.

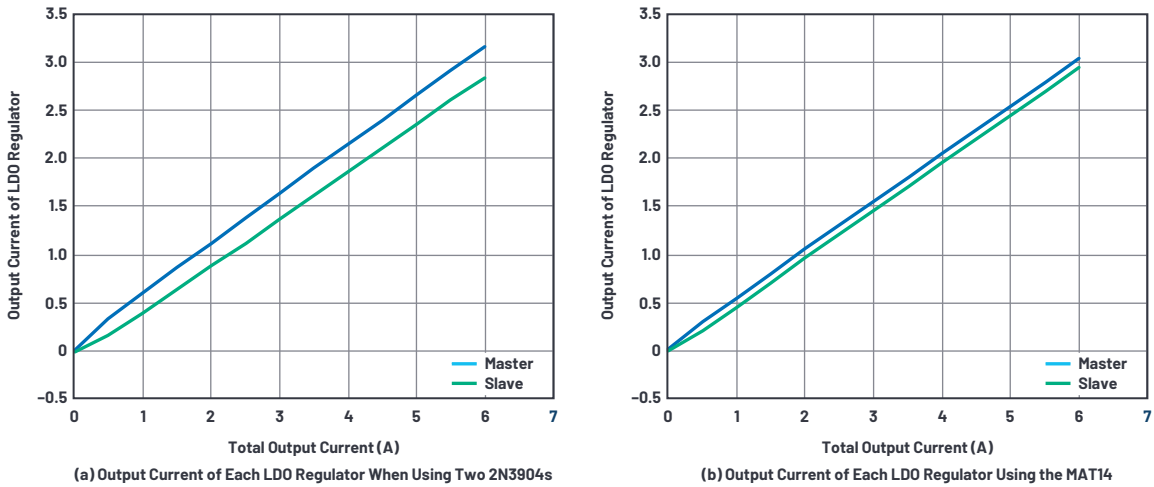

Figure 4. Current sharing mismatch is reduced by using the MAT14 matched monolithic quad transistor with parallel LDO regulators.

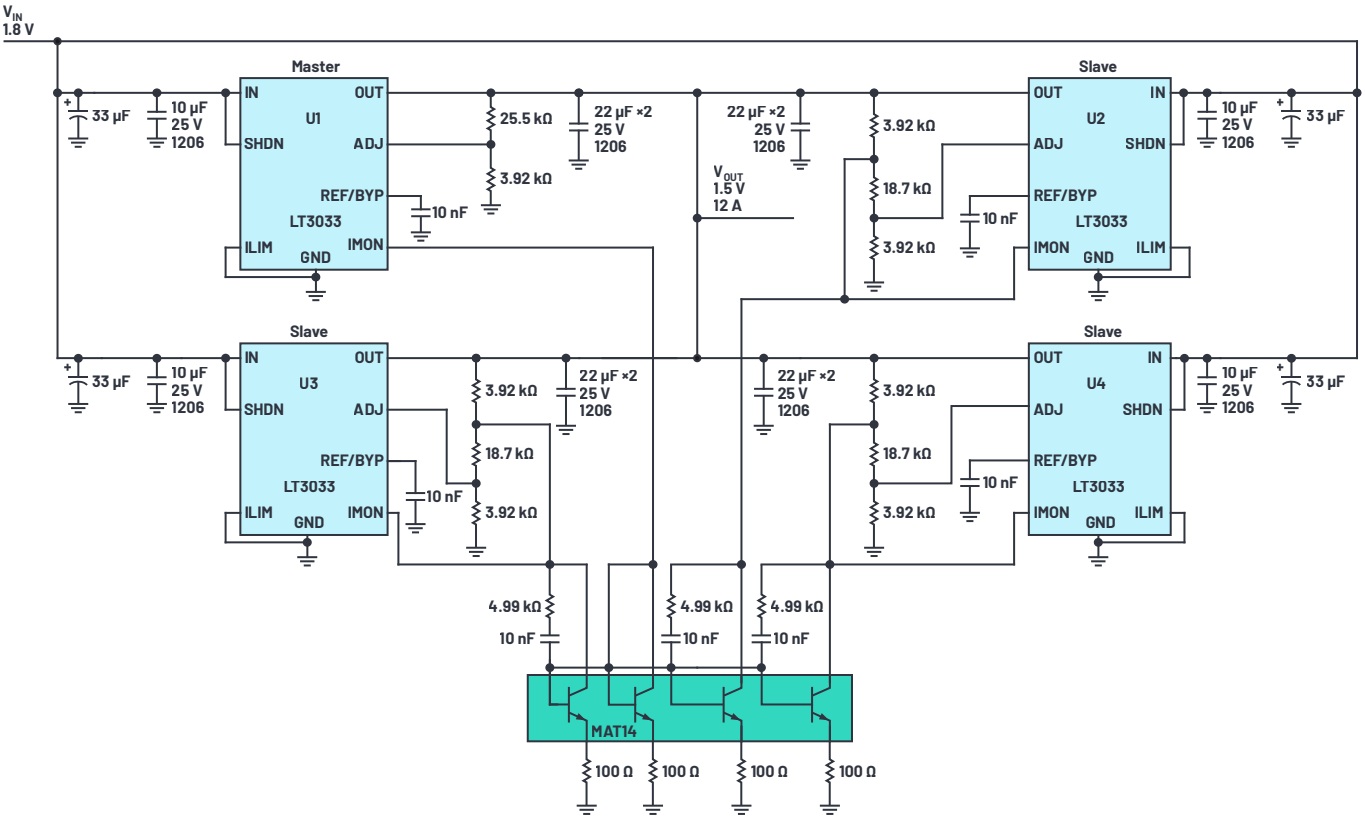

Figure 5. Four LT3033s in parallel using the MAT14.

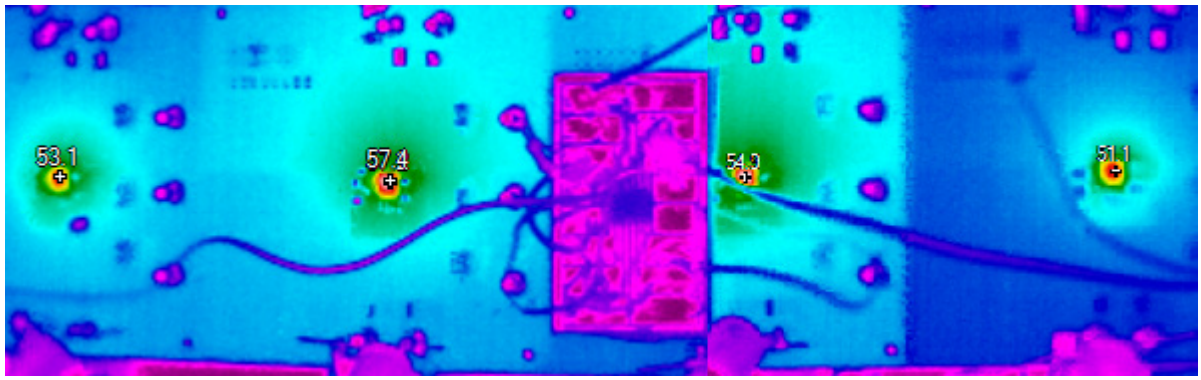

Figure 6. Thermal performance of four LT3033s in parallel.

Paralleling Four LT3033s Using Matched Components for Well Balanced Currents and Even Heat Spreading

The paralleling circuit architecture is scalable to as many LT3033s as needed by extending the current mirror and adding slave LT3033 devices. Figure 5 shows a solution with four LT3033s in parallel using the MAT14 for current sharing. The thermal performance is shown in Figure 6. The four LT3033s reach temperatures ranging from 51°C to 58°C. Considering the voltage drop along the input trace for each part, the heat spreads evenly on the board, indicating balanced current sharing for this solution. Figure 7 shows the transient response of a 1.5 V output, 12 A power supply operating from a 1.8 V input.

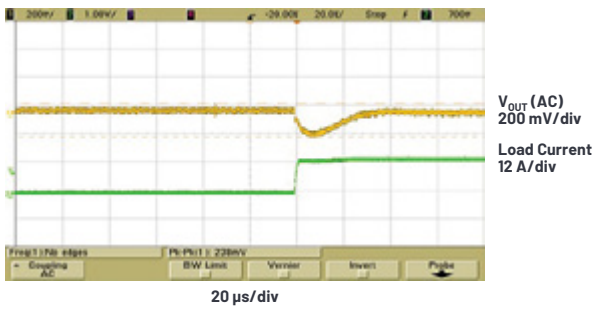

Figure 7. Load transient response of four LT3033s in parallel.

Conclusion

The LT3033 is a 3 A VLDO regulator in a 3 mm × 4 mm package. Multiple LT3033 VLDO regulators can be paralleled for high current applications because of their built-in output current monitor feature. With only 95 mV typical voltage drop at full load, LT3033 is ideal for low input voltage to low output voltage, high current applications, yielding comparable electrical efficiency to a switching regulator. Other features include programmable current limit, power good flag, and thermal limiting for reliable and robust solutions. Battery-powered systems benefit from low quiescent current and reverse-battery protection.

About the Author

Huiyu (Molly) Zhu is a senior applications engineer in Analog Devices' Power by Linear Group. She received her B.S.E.E. and M.S. degrees from Tsinghua University, China in 1998 and 2000, respectively, and a Ph.D. degree in electrical engineering from Virginia Polytechnic Institute and State University, Blacksburg, Virginia in 2005. He can be reached at *molly.zhu@analog.com*.

About the Author

Fei Guo is a field applications engineer for Analog Devices, supporting the San Francisco Bay area's broad market. His previous role was a senior applications engineer in Analog Devices, supporting the isolated power converter product line. Fei has an M.S. degree in electrical engineering from Santa Clara University. He has designed power converters for instrumentation, automotive, industrial, and communications products. Fei's hobbies include hi-fi audio, photography, and playing with his two young kids. He can be reached at *fei.guo@analog.com*.

RAQ Issue 173: On Building Physically Accurate Analog Switch Macromodels

Barry Harvey, *Analog IC Designer*

Question:

Can I improve the LTspice model of an analog switch in case my analog design contains switches and muxes?

Answer:

Sure; it is not difficult to generate your own models.

Introduction

I was testing a circuit and found many discrepancies from the paper design I used to create it. The dynamics of the circuit were a bit unexpected, and the noise level was much larger than required. I needed to bring the circuit to a simulator to fully understand it.

The circuit involved analog switches and op amps. There are good macromodels for the op amps employed, but the analog switch macromodel was not designed for generality. In the header of the switch macromodel file is the warning that modeled parameters were only valid for a specific supply and temperature. Well, wouldn't you know it: my circuit has different operating conditions from the modeled one. The thing about analog switches is that they are so general-purpose that one operating point is not enough. The existing industry-standard models provide a good start, but if you enter the analog performance arena, you might need a new macromodeling approach that brings your simulation to a higher level.

As I began to browse through various analog switch macromodels from Analog Devices and other IC companies, I noticed that all their headers tell of no supply nor temperature dependence being modeled. Thus, I would have to make my own macromodel.

My philosophy in this work is that full transistors in the analog switches using the simplest device models provide all the behaviors to be emulated, but the interface from control pin to MOS gates should be the simplest behavioral components.

All work here is done with the LTspice simulator; the code would work on other simulators with a translation of the LTspice behavioral devices to SPICE-like polynomial functions.

We will develop the simulated behaviors in a specific sequence.

Developing LTspice Model Parameters for On Resistance

We will use the simplest model to run real MOS devices. To model on resistance, we will employ:

- W/L, the width (W) divided by the length (L) of an MOS device. W/L is the size or relative strength of the device.
- V_{T0}, the threshold voltage; and gamma, which modifies V_{T0} with device back-bias. The back-bias is the voltage between the on device and its body voltage; the body is frequently connected to the positive supply for the PMOS and to the negative supply for the NMOS in the switch.
- K_P, in the model, also known as K' or K-prime. This parameter models the strength of the process and is multiplied by W/L to scale MOS currents. For a given process, the NMOS will have ~2.5× the K_P of the PMOS.
- RD, the parasitic resistance of the device's drain.

Different MOS processes have different intrinsic parameters. Table 1 is a collection of common CMOS processes, their characteristics, and estimated intrinsic parameters related to on resistance.

Table 1. Typical Semiconductor Process Parameters

Voltage Node (V)	Device Construction	Gate Oxide Thickness (m)	V_{TO}, n/p, V	Gamma, n/p, $V^{0.5}$	K_P, n/p, $\mu A/V^2$	L, µ	R_D, n/p, Ω
40	Drain drift region	10^{-7}	0.7/–0.9	0.4/–0.57	11/5	2	~80% of $R_{DS(ON)}$
15	Soft drain diffusion	4×10^{-8}	0.7/–0.9	0.4/–0.57	22/10	1.5	~20% of $R_{DS(ON)}$
5	Simple	1.4×10^{-8}	0.7/–0.9	0.4/–0.57	80/28	0.5	~0

Let us look at the ADG333A R_{ON} curves we wish to reproduce in Figure 1.

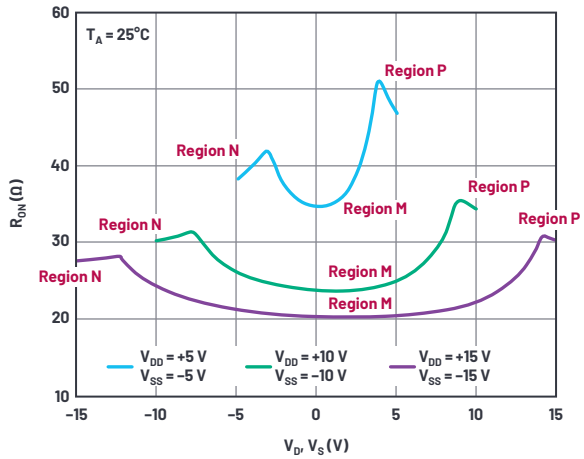

Figure 1. R_{ON} as a function of V_D (V_S), dual supply.

We see a general trend for this and any other analog switch that higher supply voltage reduces on resistance. As more voltage is applied to the switch MOS gates, the on resistance is reduced. We also see a distinct variation of on resistance with the analog level. In the N regions, the NMOS transistor in a switch is fully on, and as the analog voltage rises above the negative rail, the PMOS transistor turns on and helps reduce overall on resistance. The inflection point at region N is roughly a PMOS V_{TO} above the negative supply.

Similarly, in regions P, the PMOS device of the switch is fully on and the NMOS starts assisting the PMOS transistor roughly an NMOS V_{TO} below the positive supply.

Regions M are in the middle of the N and P regions with the NMOS and PMOS working in parallel, but each varying in on resistance depending on the analog signal level between the supplies.

To start the curve-fitting process, we first estimate the size of each transistor. The low voltage curve gives the best curve-fit for transistor $R_{DS(ON)}$. In region N, with the analog signal at the negative supply, the PMOS device is off and R_{ON} of the part is equal to the R_{ON} of the NMOS transistor. With

$$R_{DS(ON)} = \frac{1}{k_P\left(\frac{W}{L}\right)(vgs - V_{TO})} \quad (1)$$

using the 40 V NMOS typical process values, we set $R_{DS(ON)}$ = 38 Ω from the curve in Figure 1 and using the process quantities given find WNMOS = 2 µA/(38 Ω × (11 × 10^{-6} µA/V_2) × (10 V – 0.7 V)) = 514 µm. The PMOS switch would have an on resistance of 47 Ω from the above curve and thus a width of 936 µm.

I used the LTspice test circuit in Figure 2. Note that parameters R_{DN} and R_{DP}, the parasitic drain resistances, are of modest value. I started with a value of 1 µ, which caused simulator convergence slowdown. The R_{DN} value of 1 allows proper simulation speed. Adding $R_{CONVERGENCE}$ improved simulator noise and speed by giving the toggle node a convergeable conductance. I tested a floating current source for measuring on resistance.

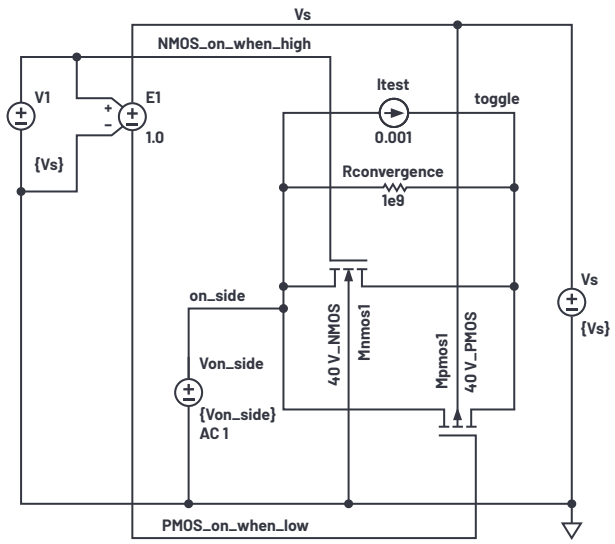

Figure 2. On resistance test circuit.

```
.model 40V_NMOS nmos (Vto=0.7 Kp=11e-6 Gamma=0.4 Rd={Rdn})

.model 40V_PMOS pmos (Vto=0.9 Kp=5e-6 Gamma=0.57 Rd={Rdp})

.dc Von_side 0 {Vs} 0.05

.step param Vs list 10 20 30
***.step param Rdn 0 20 2
***.step param Rdp0 20 2
***.step param Wn 300u 800u 50u
***.step param RDn 1 11 2

.param Von_side={Vs/2}
.param Vs=10
.param Wn=514u
.param Wp=936u
.param L=2u
.param Rdn=1
.param Rdp=1
.param k=0.4

.options plotwinsize=0
```

Figure 3 shows the simulated results for various supplies.

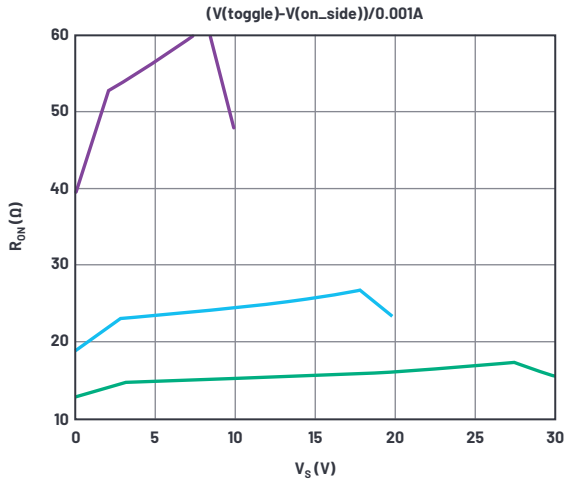

Figure 3. On resistance simulation results with initial model values.

This is a good start. The kink at the low voltage end for V_S = 30 V is at 3.6 V in the simulation and 2.7 V in the data sheet. This suggests we reduce the PMOS V_{TO}, but 0.9 V is already a realistic minimum. Better to adjust the gamma of the PMOS, which was only a guess anyway.

The kink near maximum supply is 2.5 V below the 30 V rail, where in the data sheet it should be ~1 V. Various values of gamma exaggerated the kink voltage from the rail; we will just set the NMOS V_{TO} to 1 V and its gamma to zero. A zero gamma is unexpected, but we're only trying to curve-fit. Figure 4 shows simulation results from these values with the gamma of the PMOS stepped for several supplies. We focus on the 30 V curves, which maximize the gamma effect compared to lower supplies.

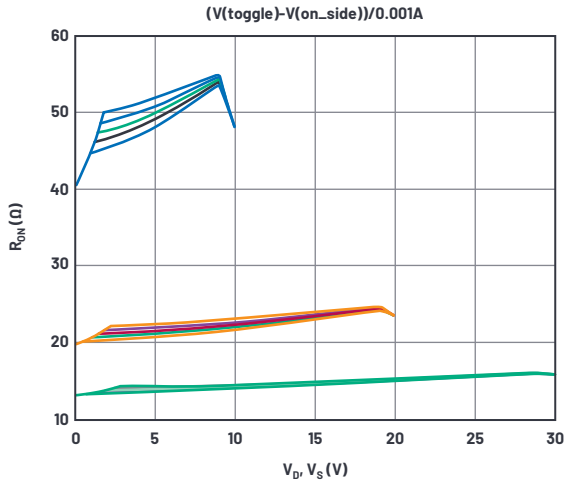

Figure 4. On resistance simulation results with gamma-p varied.

From the stepped curves, we'll choose a PMOS gamma = 0.4.

On to R_{ON}. Observe that the 10 V curves are representative of the data sheet curve at the supply extremes, but the simulation produces too low a R_{ON} for the 20 V and 30 V curves. The R_{ON}s are equal to $R_{DS(ON)}$(NMOS) + R_D(NMOS) at the negative supply extreme and $R_{DS(ON)}$(PMOS) + R_D(PMOS) at the positive supply extreme. For high supplies, the R_D parameter will be more significant than W/L, and for low supplies, W/L will dominate. We have two variables to juggle here; too laborious. We will posit that R_{ON} varies with supply due to the NMOS being variably enhanced, but the R_D value doesn't change with supply voltage (okay, it probably does in the case of drains with drift regions, but let's keep this simple). If we note the difference in data sheet R_{ON} between 10 V and 30 V supplies (11.4 Ω), we can compare that to the above curves where we step only W_N (width of the NMOS in the switch). After a bit of iterations of W_N in simulations it's clear that we need W_N = 1170 µm to get the required ΔR_{ON}, quite a lot more than the initial guess. Figure 5 shows our current results.

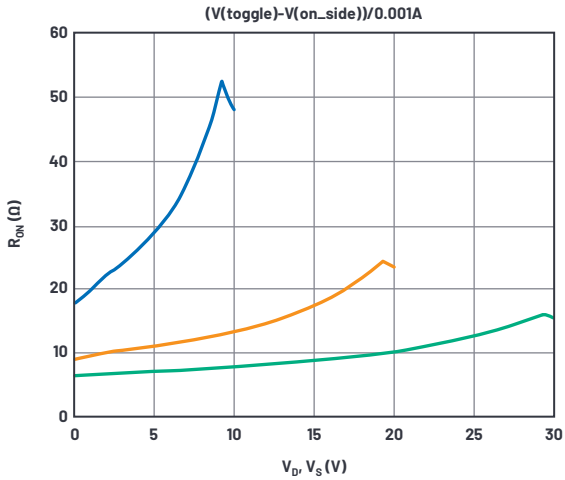

Figure 5. On resistance simulation results with W_N determined.

While the R_{ON} of the NMOS has the right supply sensitivity, the curves are too low a value at zero volts, and we must increase the fixed R_{DN}. After increasing and iterating R_{DN}, we get a best value of R_{DN} = 22 Ω, and the resulting curves are in Figure 6.

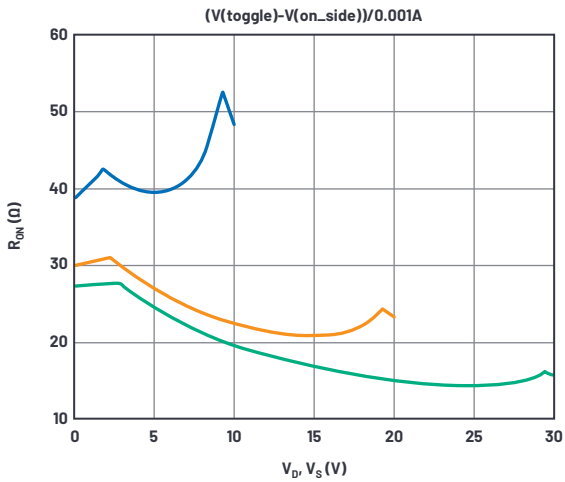

Figure 6. On resistance simulation results with R_{DN} determined.

We next determine W_P (width of the PMOS in the switch) to simulate the R_{ON} at maximum voltage, and get W_P = 1700 µm, again quite a lot more than initially guessed. With R_{DP} also set to 22 Ω, we get the final R_{ON} curve in Figure 7.

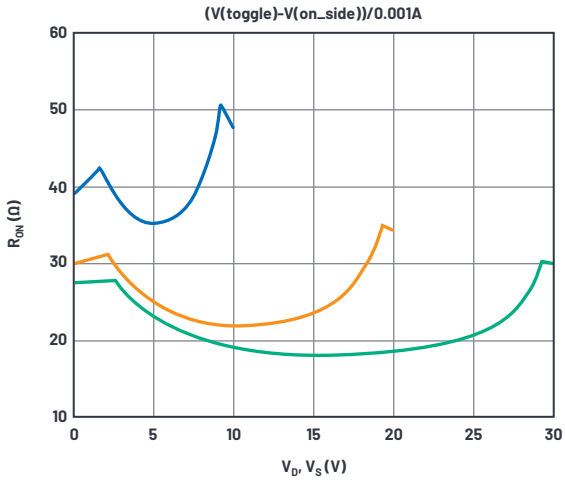

Figure 7. On resistance simulation results with W_P and R_{DP} determined.

Pretty good agreement here; there are only a few features different from the data sheet. One is that the inflection points are smooth in the data sheet curve but truly pointed in simulation. This is probably because the simple MOS model used does not support subthreshold conduction, and the simulated device turns truly off at V_{TO}. Real devices are not off at V_{TO}, but smoothly reduce current below that voltage.

Another error is most obvious in the 30 V curve. R_{ON} is 15% low at midsupply compared to data sheet. Perhaps this is due to JFET effects within the drain drift region, also not modeled.

As for temperature, there is fair but not strong compliance, seen in Figure 8.

The simulation has temperature dependence, but not as much as the data sheet curves. In the simulation model the R_D terms do not have tempco. R_Ds could be modeled by external resistors with correct tempco, but we will leave it as is for simplicity.

Obtaining LTspice Model Parameters for Charge Injection

When MOS transistors turn off, the charge in the channel must go somewhere, so it squirts out of the drain and source terminals. When an analog switch is turned off, charge also goes out and is called charge injection. A common way of measuring it is to place a fixed voltage on one end of an on switch and a large capacitor at the other end. When turned off, the charge is captured by the capacitor and a small voltage step occurs. We will now the add gate oxide thickness $T_{OX} = 1 \times 10^{-7}$ to the MOS models (gate capacitance is the largest source of charge injection). Our simulation setup is shown in Figure 9.

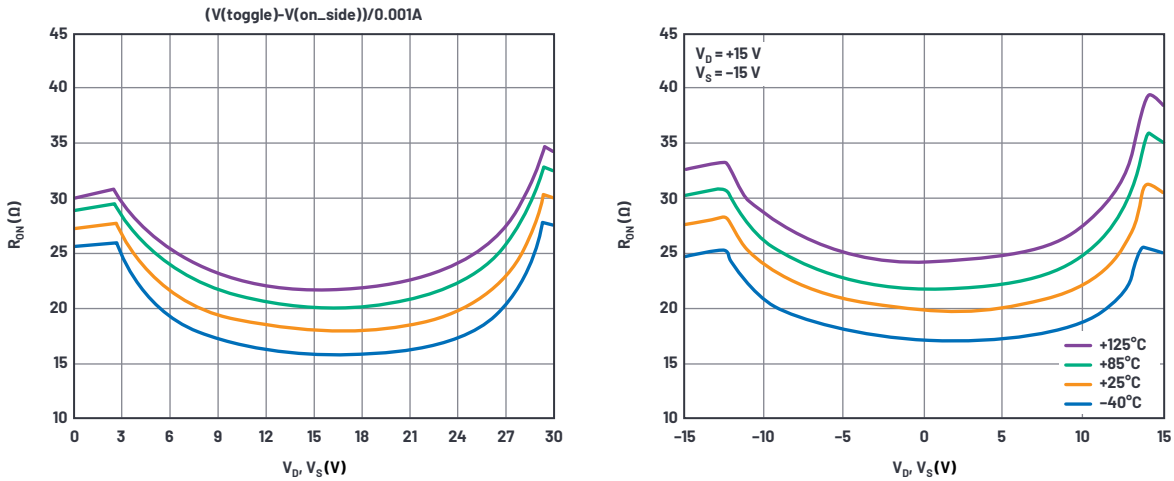

Figure 8. On resistance simulation and data sheet results over temperature.

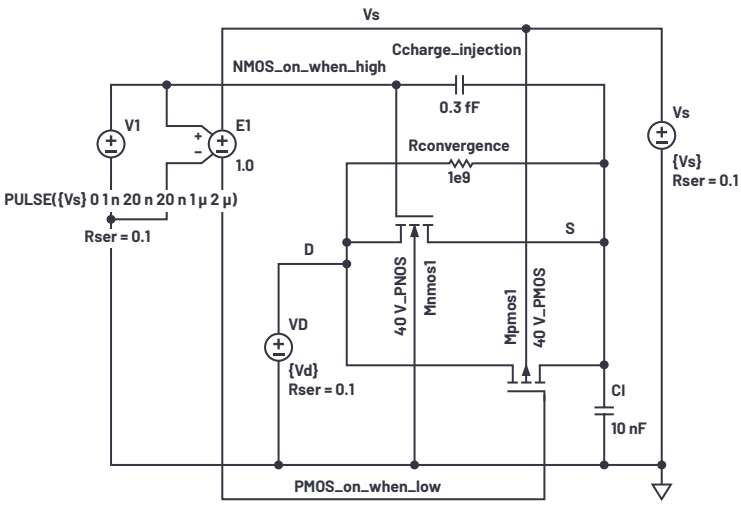


```
.model 40V_NMOS nmos (Vto=0.7 Kp=11e-6 Gamma=0 Rd=22 tox=0.1u)
.model 40V_PMOS pmos (Vto=0.9 Kp=5e-6 Gamma=0.4 Rd=22 tox=0.1u)

***.dc Vd 0 {Vs} 0.05
.tran 100n

***.step param Vs list 12 33
.step param Vd 0 30 0.2

***.step temp list –40 25 85 125

.param Vs=30
.param Vd={Vs/2}

.param Wn=1170u
.param Wp=1700u
.param L=2u

.options plotwinsize=0

.meas TRAN Charge_Injection find (V(S)–V(D))*10n at=99n
```

Figure 9. Charge injection simulation setup.

The data sheet charge injection test circuit places a voltage source at the D terminal of a switch, and the capacitor CI at the S terminal of the switch. When the switch transistors are turned off, CI is isolated and integrates charge pumped into it by the switches. The waveform of such an event with V_D held to 24 V with a 30 V supply is shown in Figure 10.

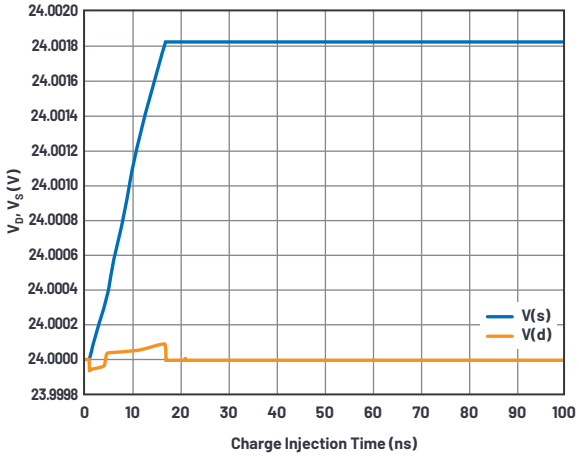

Figure 10. Charge injection simulation waveforms.

The charge injected is the voltage jump between V(S) and V(D) times the 10 nF hold capacitor. We can step the switch voltage V_D across the supply voltage and use the .meas statement to capture the value of charge injection at each voltage. Figure 11 shows the data sheet curve and simulated results.

Our simple MOS model does not mimic the shape of the data sheet curve very well, but the peak-to-peak charge injection is 32 pC in the data sheet curves and 31 pC in simulation. Surprisingly close, but if we had to, we could tweak T_{OX} to perfect the simulation results.

There is an offset between the curves that we can compensate for using $C_{CHARGE\_INJECTION}$. After fiddling with some values, we choose an optimal $C_{CHARGE\_INJECTION}$ = 0.28 pF. If an opposite polarity of shift were needed $C_{CHARGE\_INJECTION}$ would be reconnected to the PMOS_on_when_low node.

The tweak capacitor $C_{CHARGE\_INJECTION}$ was a convenient way to offset the charge injection vs. the analog voltage simulation curve. What if the peak-to-peak injection simulated were too small? Well, most of the charge injection is mostly the switches' gate voltage swings sending charge through the gate-channel capacitance of the switch transistors. If we simulate too little injection, we can simply increase one or both gate areas. To do this we would increase the parameters L and W of a switch device by the same factor, being careful to not modify the W/L ratio that sets on resistance. Rather than use $C_{CHARGE\_INJECTION}$ we could have increased the NMOS W and L.

Alternatively, we could adjust T_{OX} in each device to get better charge injection correlation. This would not be physically possible, but hey—it's just a simulation. With the simple models we are using, T_{OX} does not influence other behaviors.

Obtaining LTspice Model Parameters for Capacitances

Having set up parameters for good R_{ON} and charge injection simulation results, we now simulate S and D terminal capacitances.

One important point is that both the drain and source regions of high voltage MOS switches must have drift regions. As a switch, you can't tell the functional difference between sources and drains, and the body potential to both drains and sources will require the drift regions in each. This is also true of the medium-voltage soft diffusions, but non-existent in low voltage MOS. We have lumped the drift region resistance that would exist in both drain and source into R_D, and that works fine for switches, but not transistors in saturation.

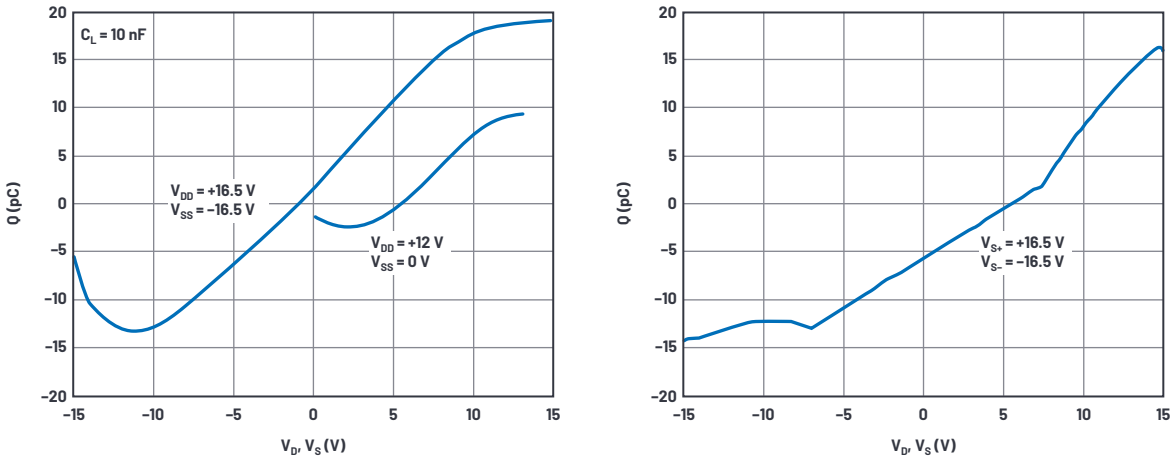

Figure 11. Charge injection data sheet and simulation waveforms.

Figure 12 shows our simulation setup.

In LTspice, you can run an .ac on only one frequency, using the list option in .ac, but offer only one frequency argument (1 MHz here). Then you run a .step V_{SOURCE} dc voltage across the supply range to get a capacitance vs. voltage sweep.

The D terminal of the off-switch device is held to midsupply. The S terminal, renamed source here to prevent confusion with V_S, is driven by a voltage source with dc value sweeping from 0 V to V_S and with an ac drive of 1 V. Capacitance is derived from $I(V_{SOURCE})/(2 \times \pi \times 1\ MHz \times 1\ V)$. The logic drive V1 is changed to 0 V to turn the transistors off.

Drain and source capacitances to bulk are C_{BD} and C_{BS} respectively in the model statement. There are built-in default concentrations, built-in voltage, and exponent in the model that make C_{BD} and C_{BS} voltage variable. Because they are symmetrical, drain and source capacitances would be made equal. Further, because the PMOS is a different width from the NMOS, the ratio of $C_{BD,NMOS}/C_{BD,PMOS} = C_{BS,NMOS}/C_{BS,PMOS} = W_N/W_P$, which we established in the on resistance modeling. Figure 13 shows the simulation results.

The displays are $I(V_{SOURCE})/(2 \times \pi \times 1\ MHz)$, which is capacitance. LTspice doesn't know this and displays pA instead of pF.

Unfortunately, we have no data sheet curves to compare to. We do know from the specification table in the data sheet that the capacitance—probably at midsupply, but not specified in the data sheet—is typically 7 pF at 30 V supply and 12 pF at 12 V supply. I adjusted the CBs to obtain the 7 pF curve at 30 V, but only simulated 10 pF at a 12 V supply. After fiddling with built-in potential and capacitance formula exponent, the model used allows no flexibility to improve the 12 V/30 V compliance.

Figure 14 shows the on-state capacitance simulation setup.

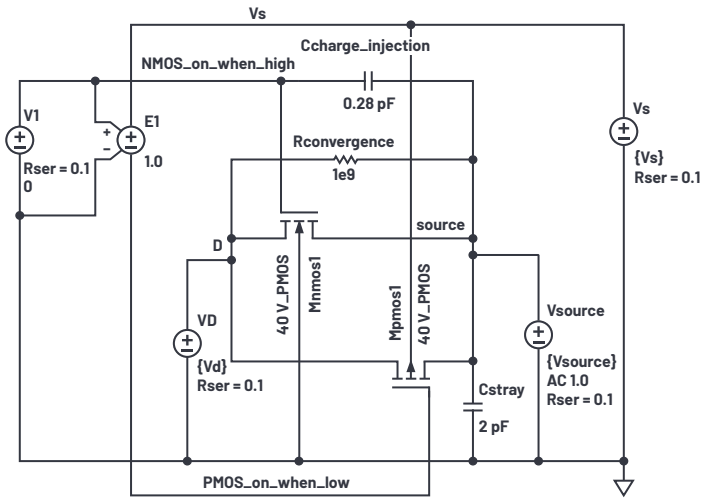


```
.model 40V_NMOS nmos (Vto=0.7 Kp=11e-6 Gamma=0 Rd=22
+ tox=0.1u cbd=9p cbs=9p)

.model 40V_PMOS pmos (Vto=-0.9 Kp=5e-6 Gamma=0.4 Rd=22
+ tox=0.1u cbd=13.5p cbs=13.5p)

***.dc Vsource 0 {Vs} 0.05
***.tran 100n
.ac list 1e6

***.step param Vs list 12 30
.step param Vsource 0 12 0.2

***.step temp list -40 25 85 125

.param Vs=12
.param Vd={Vs/2}
.param Vsource={Vs/2}

.param Wn=1170u
.param Wp=1700u
.param L=2u

.options plotwinsize=0

***.meas TRAN Charge_Injection find (V(S)-V(D))*10n at=99n
.meas AC Isource FIND mag(i(Vsource)) at=1e6
```

Figure 12. Off-capacitance test simulation setup.

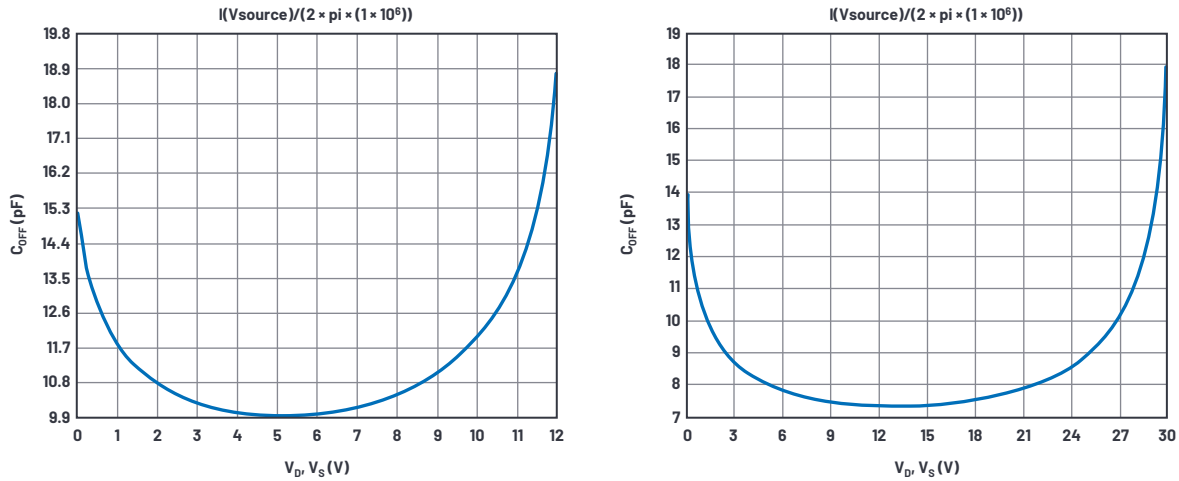

Figure 13. Off-capacitance vs. dc voltage at V_S = 12 V (left) and 30 V (right) results.

Here the right switch of a full spdt switch is on, and the left switch is off and connected to a $V_S/2$ source. The capacitances of the right half of the left switch and full capacitances of the right switch, plus inevitable parasitic capacitances at D and S terminals are all paralleled and driven by our 1 MHz test signal at the V_s source, whose dc level is stepped across ground to V_S. Figure 15 shows the results.

We simulate 29.5 pF and 21.4 pF where the data sheet gives 26 pF and 25 pF. Considering the variability in circuit-board layout capacitance, we'll call this close enough.

Leakage Currents

The data sheet curves show voltage-dependent pA-level leakage currents at 25°C, but the data sheet specification only guarantees hundreds of pA. I am swayed more by the curves' results at 25°C. The small leakage currents apparently were not considered important enough in this device to guarantee at test. To be fair, measuring single pA takes a lot of engineering development effort as well as long test times.

At 85°C, the guarantee is a few nA (which can be measured efficiently) with a typical result in the range of a few hundred pA. I'm going to accept these typical results as good.

Leakage current is a product shortcoming; it doesn't have tight statistics and varies wildly with temperature. It is not the kind of specification that we design to—rather, it's a quantity that disrupts the circuits it's connected to. For macromodel use, any leakage of proper magnitude will be simulated as a circuit defect and be a useful warning to the designer. I'll choose a target of 1 nA for an on switch at 85°C.

The model we have shows no leakage beyond $R_{CONVERGENCE}$ and G_{MIN} currents. G_{MIN} is a resistor the simulator places across junctions to assist convergence. It is normally 1×10^{-12} conductance, but in the presence of 30 V supplies we can get multiples of 30 pA currents, way too high for this work. G_{MIN} will be reduced to 1×10^{-15} in the .options line of the simulation and $R_{CONVERGENCE}$ raised to 1×10^{15}.

The physical origin of these leakages is probably mostly from electrostatic discharge (ESD) protection diodes connected to every pin. We will insert them into the simulation setup in Figure 16.

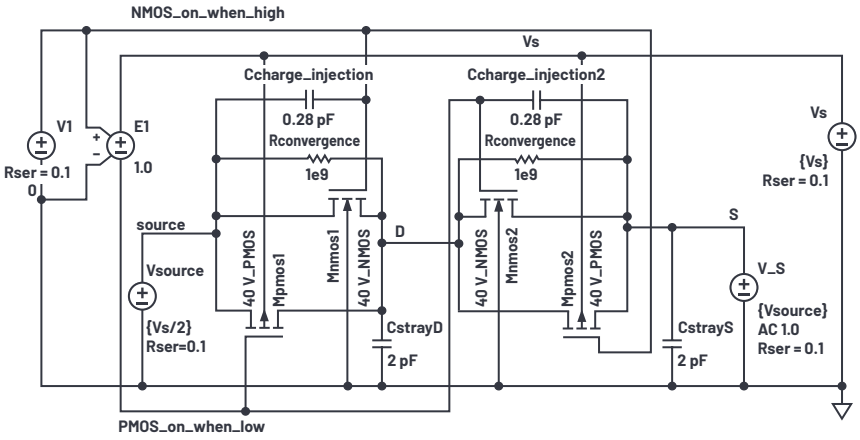


```
.model 40V_NMOS nmos (Vto=0.7 Kp=11e-6 Gamma=0 Rd=22
+ tox=0.1u cbd=9p cbs=9p))

.model 40V_PMOS pmos (Vto=-0.9 Kp=5e-6 Gamma=0.4 Rd=22
+ tox=0.1u cbd=13.5p cbs=13.5p)

***.dc Vsource 0 {Vs} 0.05
***.tran 100n
.ac list 1e6

***.step param Vs list 12 30
.step param Vsource 0 12 0.2

***.step temp list -40 25 85 125

.param Vs=12
.param Vd={Vs/2}
.param Vsource={Vs/2}

.param Wn=1170u
.param Wp=1700u
.param L=2u

.options plotwinsize=0

***.meas TRAN Charge_Injection find (V(S)-V(D))*10n at=99n
.meas AC Isource FIND mag(i(Vsource)) at=1e6
```

Figure 14. On-capacitance test simulation setup.

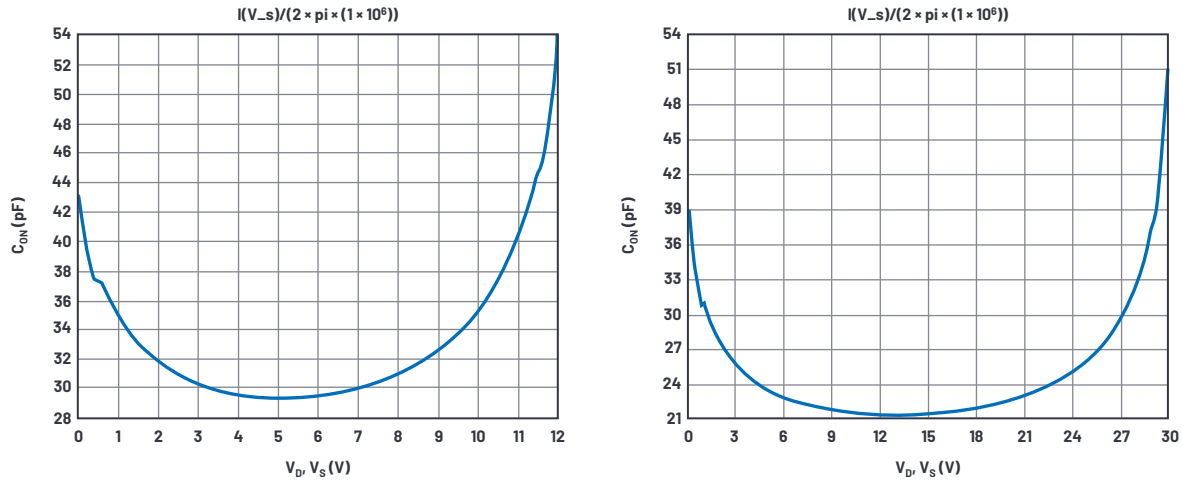

Figure 15. On-capacitance vs. dc voltage at V_S = 12 V (left) and 30 V (right) results.

After fiddling with IS in the diode model, we get leakage over temperature in Figure 17.

Logic Interface and Gate Drivers

A purely behavioral logic-to-gate drive circuit is shown in Figure 18.

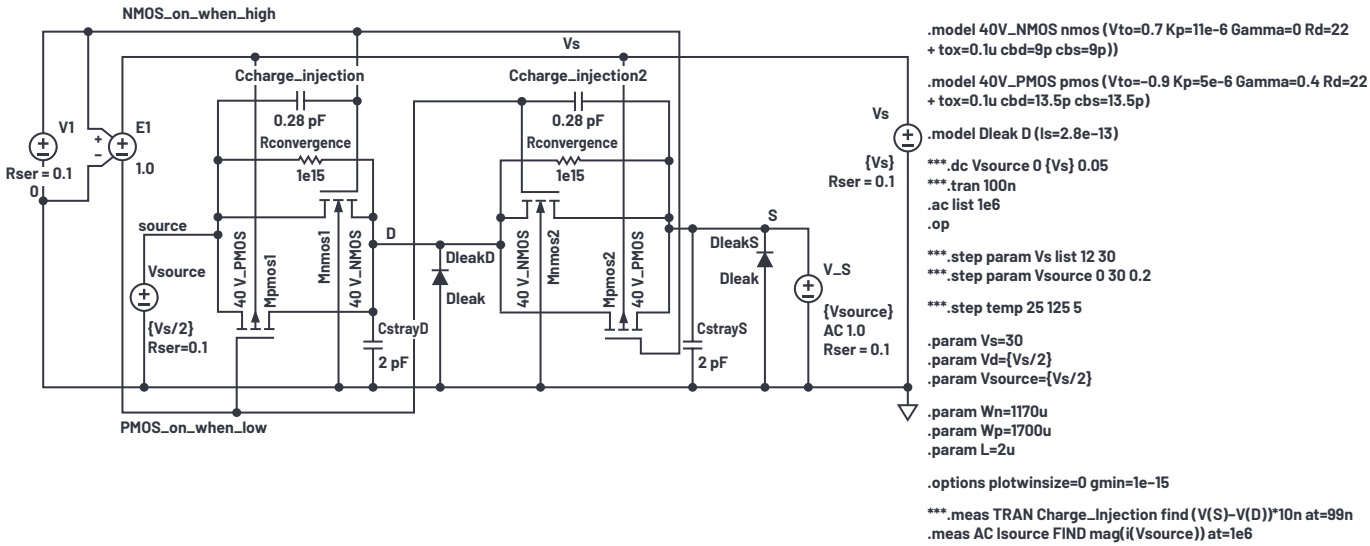

Figure 16. Leakage test simulation setup.

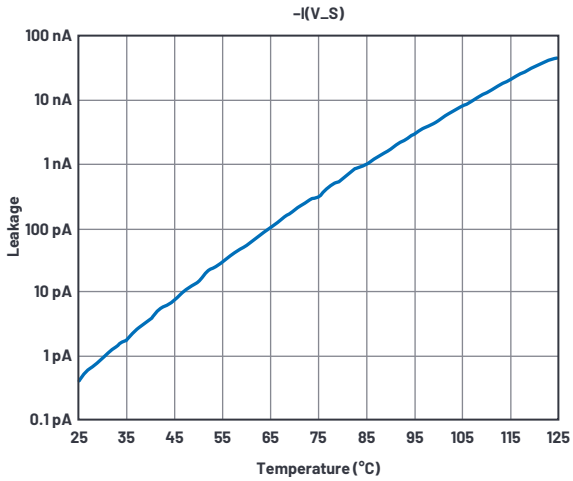

Figure 17. Leakage test over temperature simulation results.

The external logic input is at the In terminal at the left of Figure 18. It is the input of an ideal transconductance Glogic_thresholda, which has a piecewise-linear transfer function. For logic inputs below 1.37 V, the output at logica node is 0 V; for inputs above 1.43 V logica is at 1 V; and between 1.37 V and 1.43 V in logica moves linearly from 0 V to 1 V. Glogic_thresholda thus ignores supply variations to provide a 1.4 V input threshold.

Transiently, Cdelaya slows down the logica node so that we can pick off some time points from it. To make a comparator we again use a transconductance, here Gbreakbeforemakena whose output goes from 0 V to 1 V again but with the threshold skewed a bit above 0.5 V. As seen in Figure 19, the skewed pickoff voltages 0.52 V and 0.57 V rather than 0.5 V allow faster turn-off from exponentials falling from 1 V than the turn-on time for exponentials rising from 0 V.

Full gate drive voltage is produced by the B_non and B_pon behavioral current sources. B_nona sources a current of V_{DD}/1000 when node n_breakbeforemakena >0.5 V, driving the voltage at node nona to V_{DD}, as loaded by a 1000 Ω resistor. When node n_breakbeforemakena <0.5 V, the node nona is driven to V_{SS}. Thus, we have a nice rail-to-rail gate drive that complies with supply voltages and has a fixed 1.4 V input threshold.

One more characteristic needs explanation. Note that in Figure 20, higher supply voltages reduce the delay times. This is implemented by B_supplysensitivitya, which feeds back to Cdelaya a fraction of its own dynamic current that varies with V_{DD}. Rsupply_sensitivitya drops very little voltage due to Cdelaya current, leaving Cdelaya's behavior mostly a pure capacitor. Feeding a replica of Cdelaya's current back to Cdelaya essentially creates a controllable variable capacitor, and the math inside Bsupply_sensitivitya creates the delay vs. V_{DD} curve in Figure 20.

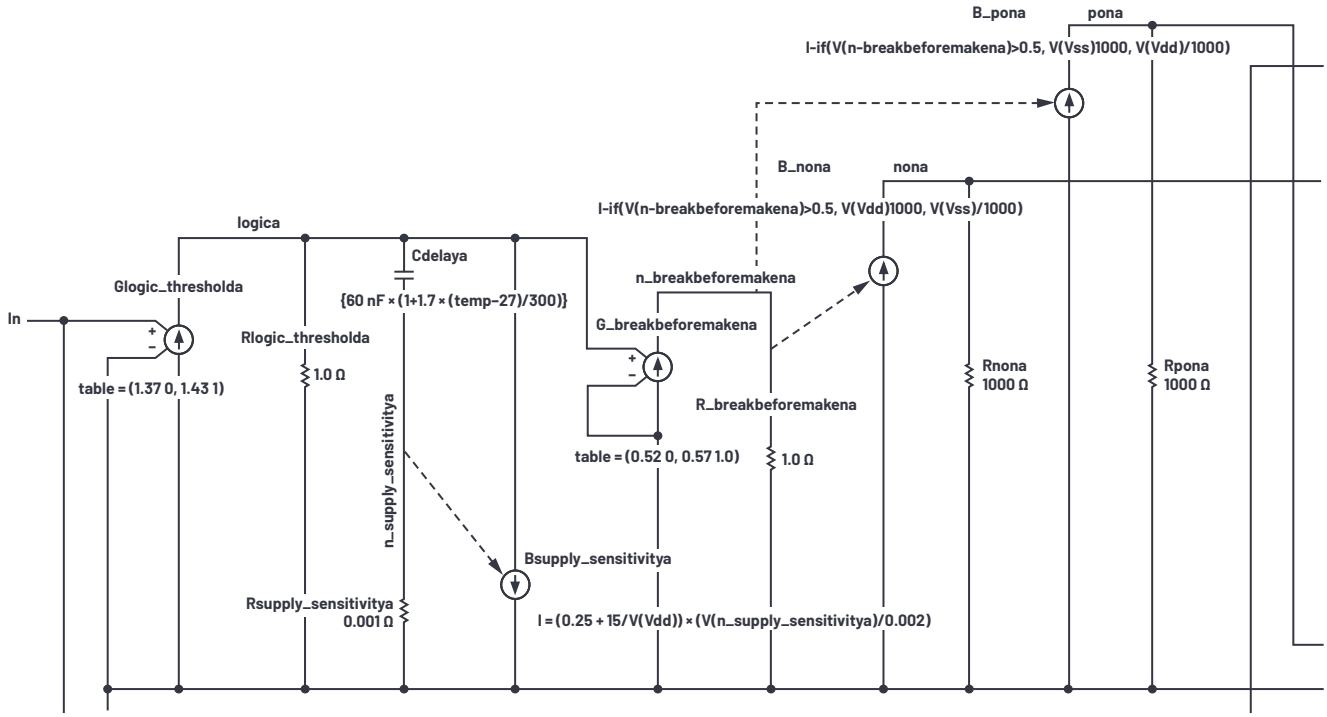

Figure 18. Behavioral logic-to-gate interface.

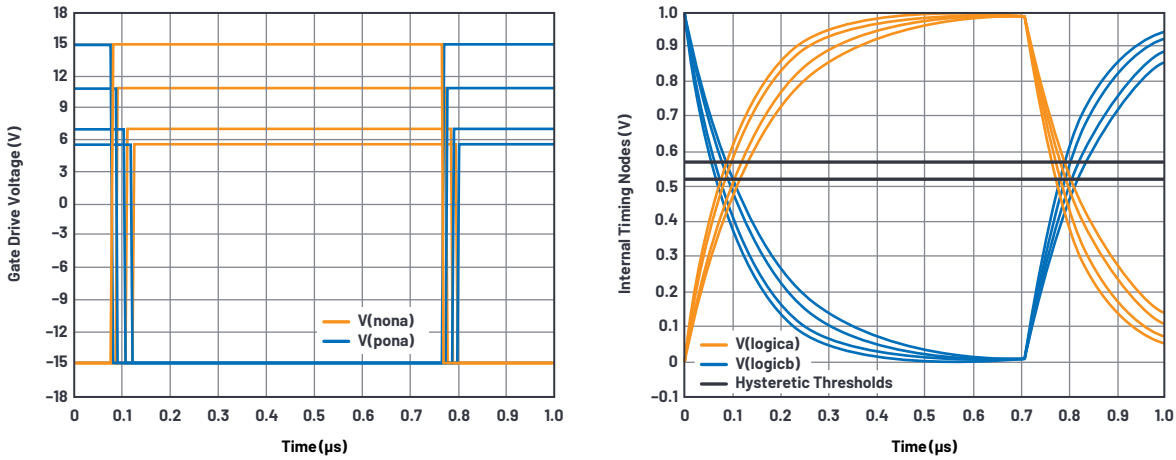

Figure 19. Break-before-make timing.

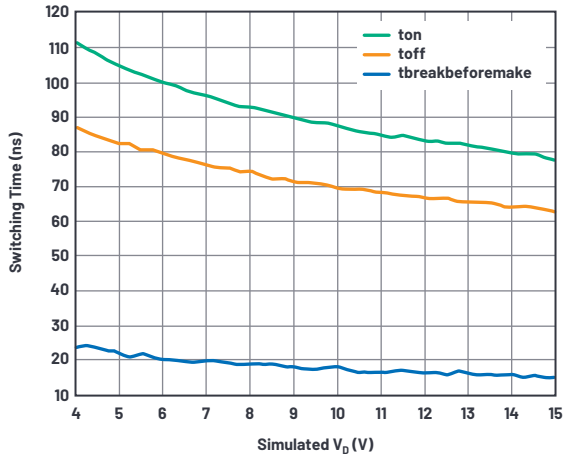

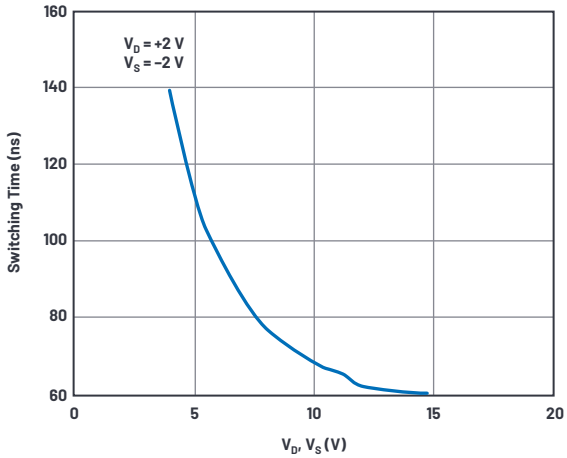

Figure 20. Break-before-make timing results from simulation and data sheet curve.

Well, our circuit emulates the T_{ON} delay as 111 ns for V_{DD} = 4 V while the data sheet curve says 140 ns; and for V_{DD} = 15 V simulated delay is 77 ns vs. data sheet delay of 60 ns. Not great correlation; I'll leave it to the reader to refine the Bsupply_sensitivity function to do better. At least the break-before-make varies nicely between 15 ns and 24 ns.

While we don't have much data sheet data on delay vs. temperature, I added a temperature term in Cdelaya to at least model slowdown when hot, seen in Figure 21.

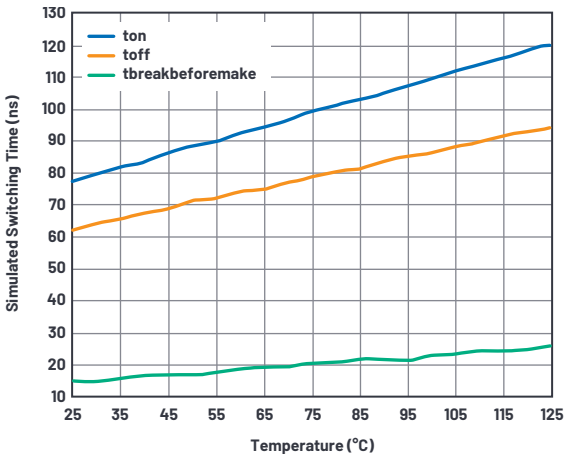

Figure 21. Timing delays vs. temperature.

Assembling the Macromodel

Figure 22 shows the assembled analog switch that will become a subcircuit. Hard L and W numbers were placed into the transistor symbols instead of parameters, and all excitation and I/O are removed in favor of pin connections SA, D, SB, In, V_{DD}, V_{SS}, and Gnd_pin.

A second logic interface is provided for the other switch of the spdt pair. ESD protection diodes are installed between analog terminals and V_{SS} and between the logic In and ground. Note that the "-a" suffix in names of the upper logic interface devices and nodes are replicated as "-b" suffix in the lower interface. Glogic_thresholdb interface has the opposite output from the table in Glogic_thresholda to allow one or the other switch pair to operate rather than be turned on simultaneously.

An alternative ESD protection scheme involves diodes from a protected pin to both V_{DD} and V_{SS}, and a clamp between V_{DD} and V_{SS}. The data sheet generally gives insight as to the protection scheme, and leakage currents are assigned to both supplies.

The spdt subcircuit is given a symbol and used four times in the master schematic ADG333A.asc of Figure 23.

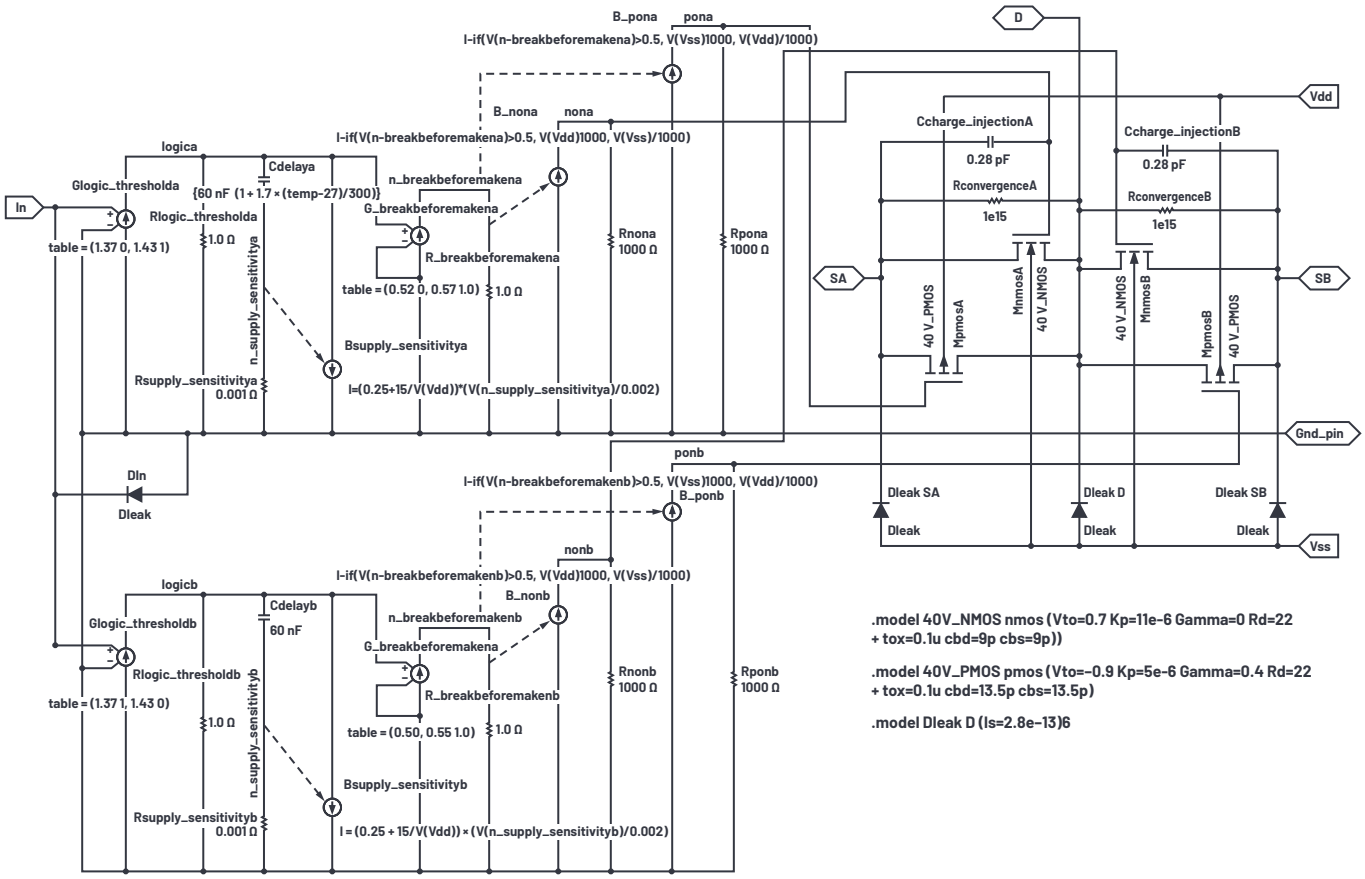

Figure 22. Assembled SPDT subcircuit spdt 40V.asc.

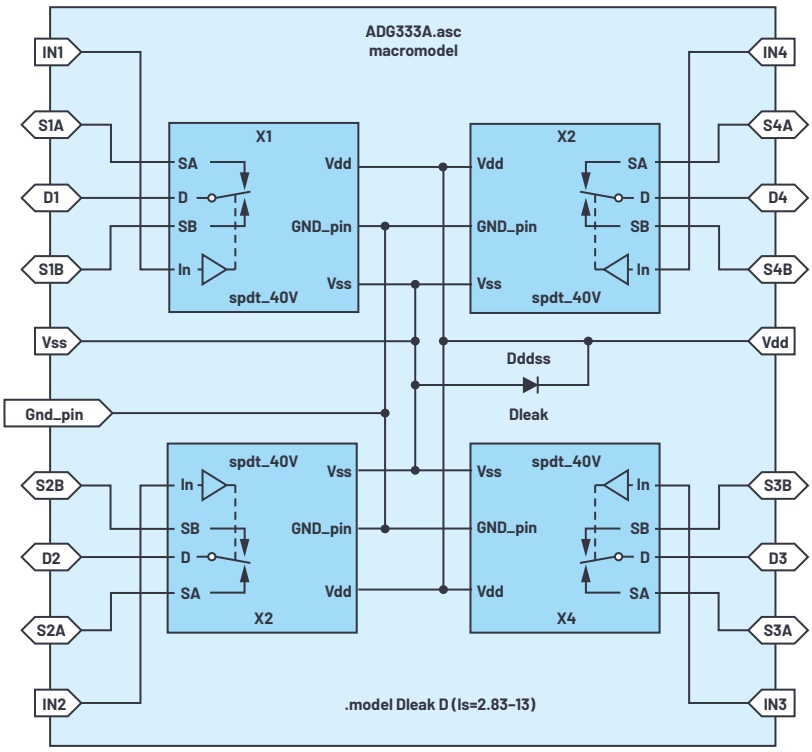

Figure 23. ADG333A macromodel circuit schematic.

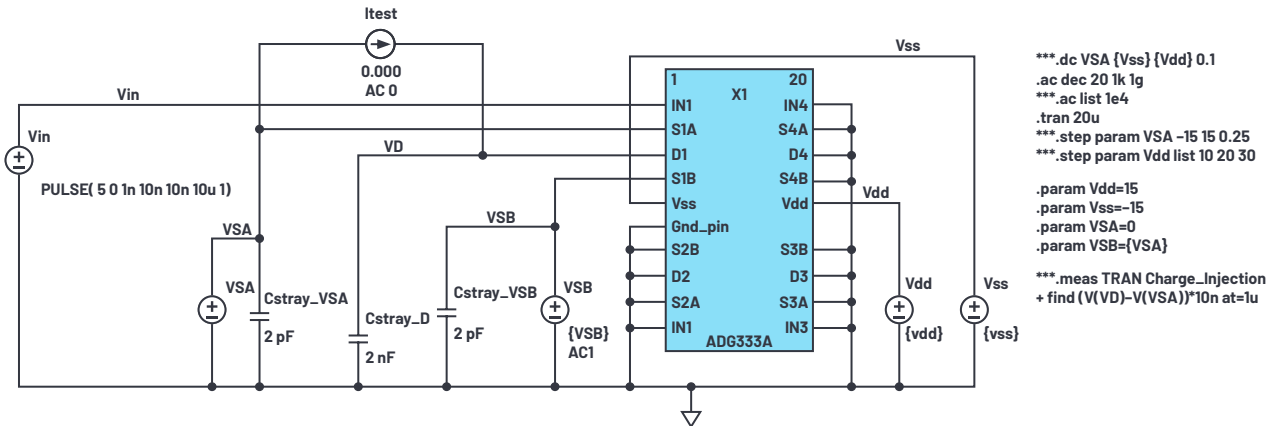

Figure 24. ADG333A macromodel test bench.

Figure 24 is the test bench schematic for verifying final macromodel results.

Summary

We've seen how to realize a decent macromodel for a specific analog switch and how to obtain parameters that support a few different semiconductor processes used to realize the physical device. The resulting macromodel displays defects such as on resistance and its variations, charge injection as a function of supply and signal level, parasitic capacitances and their variations over voltage, logic interface delays, and leakages. Hopefully, the macromodels will be helpful in simulating the real performance of analog switches.

Addendum

To download LTspice, please visit analog.com/ltspice.

Here is the LTspice text file of the macromodel symbol, to be filed under the name ADG333.asy. It contains subcircuit simulation details. Rather than copy the ADG333.asc schematic into every schematic that uses it, we use a symbol that refers to it as the .asy. Within the ADG333 symbol are individual switch symbols. This is the symbol simulation content to be filed as spdt_40V.asc. The actual symbol is to be filed as spdt_40V.asy.

About the Author

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

Open-Source LIDAR Prototyping Platform

István Csomortáni, FPGA Design Engineer,
Dragos Bogdan, Software Development Engineer Manager,
Cristian Orian, System Design Engineer, and
Andrei Cozma, Engineer Manager

Abstract

This article discusses Analog Devices' new broad market LIDAR prototyping platform and how it helps shorten the customers' product development time by providing a complete hardware and software solution that customers can use to prototype their algorithms and custom hardware solutions. It goes over the details of the modular hardware design, including the light receive and transmit signal chains, FPGA interface, and optics for long range sensing. The system partitioning decisions that were made are explained, helping highlight the importance of good system design, interface definition, and right level of modularity. The components of the open source LIDAR software stack and the platform specific API are described showing how customers can benefit from these during product development and integrate them in their final solutions.

Introduction

As autonomous vehicles and robots continue to move from science fiction to reality, automotive and industrial customers are seeking new environment perception solutions to enable these machines to navigate autonomously. LIDAR is one of the fastest growing technologies in this field and is seeing wider adoption as the technology becomes more mature and reliable, opening up a huge market opportunity. As many startups and well renowned sensor companies are working toward developing more precise, less power hungry, smaller form factor, and more cost-effective LIDAR sensors, they all face the same challenges when it comes to system hardware design and implementing the software infrastructure to talk to all the components in the system. These are the exact areas where ADI can bring value through hardware reference designs accompanied by open-source software stacks, enabling customers to easily integrate into their products ICs from the ADI LIDAR portfolio, as well as software modules and HDL IPs, shortening their time to market.

System Architecture

As customers develop their LIDAR sensors, there are few areas of differentiation in the system design: receive and transmit optics, number and orientation of lasers, laser firing patterns, laser beam steering, and number of light receive elements. But, regardless of these choices, there is a high degree of commonality in the receive signal chain and laser drive signal requirements. Based on these assumptions, Analog Devices designed a modular LIDAR prototyping platform, AD-FMCLIDAR1-EBZ, intended to allow customers to easily configure or replace parts of the design with their own hardware, designed according to specific applications requirements, but still be able to use the platform as a whole system. The system is partitioned into three different boards with standardized digital and analog interfaces:

- A data acquisition (DAQ) board containing a high speed JESD204B ADC and corresponding clocking and power. This board has an FMC compliant interface to connect to the users' preferred FPGA development board. It serves as the baseboard in the system by having the other two boards connected to it via digital connectors that route control and feedback signals between these boards and the FPGA and through coaxial cables for analog signals.
- An analog front-end (AFE) board containing the avalanche photodetector (APD) light sensor and the entire signal chain required for conditioning the APD output signal so that it can be fed into the ADC on the DAQ board.
- A laser board containing the lasers and drive circuit.

As always, in system design, modularity means flexibility, but it also comes with drawbacks such as increased complexity, performance degradation, and increased cost, which must be thoroughly evaluated when deciding on the system partitioning. In this case, the system was broken down into these three boards for the following reasons:

- The ADC and clocking are very likely to stay the same regardless of the implementation of the analog front end and chosen laser solution.
- The analog front-end hardware design and form factor is subject to change depending on the chosen APD, overall system receive sensitivity, and chosen optics.
- The laser board design and form factor are also subject to change depending on the chosen illumination solution and optics.
- The system must provide a lot of flexibility in positioning and orienting the receiver and the transmitter so that they are aligned with each other or other targets, which is why flex cables are being used for the digital signals and coaxial cables for the analog signals that go between the boards.

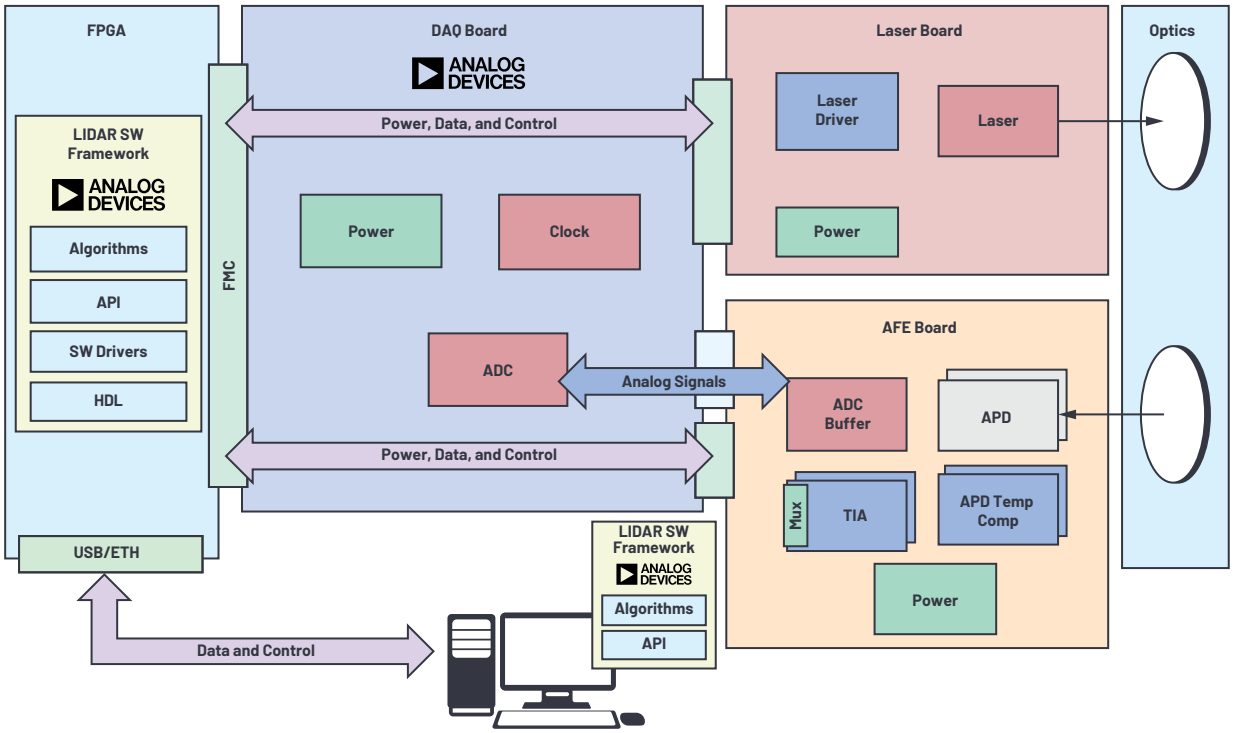

Figure 1. LIDAR platform system design.

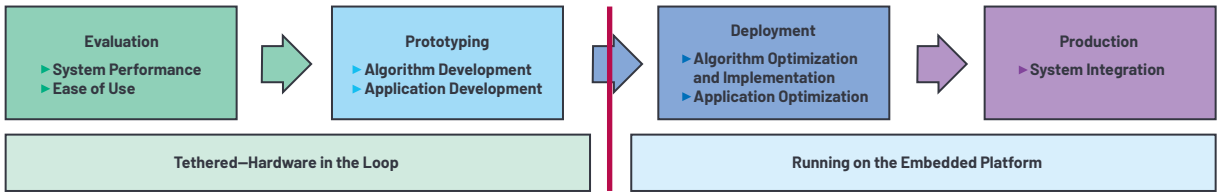

Figure 2. Product development cycle.

The software stack that accompanies the hardware design is based on a hierarchical approach with a few layers dividing it into OS specific drivers and interfaces, a system specific API, and an application layer. This allows the upper layers of the stack to remain unchanged regardless of whether the software is running on an embedded target or a PC talking to the system via the network or a USB connection. This is very valuable in the different product development stages, shown in Figure 2, since it means that the same application software that was developed during the prototyping stages, when the system is tethered to a PC for ease of development, can be easily deployed onto the embedded production system without ever touching the low level interfaces.

Hardware Design

A LIDAR sensor computes the distance to a target by measuring the time it takes for the light pulse to travel to the target and back. The time is measured in increments of the ADC sampling rate since that determines the resolution with which the system samples the received light pulse. Equation 1 shows how the distance is computed relative to the ADC sampling rate.

$$d = N \times \frac{L_S}{2 \times f_S} \quad (1)$$

Where:

L_S is the speed of light, 3×10^8 m/s

f_S is the ADC sampling rate

N is the number of ADC samples since the light pulse is generated until it is received back

Given the 1 GHz sampling rate of the AD9094 JESD204B quad ADC being used in the system, the results of each sample correspond to a distance of 15 cm. For this reason, it is critical that there aren't any sampling uncertainties in the system since even an uncertainty of a few samples can result in large distance measurement errors. Traditionally, LIDAR systems are based on parallel ADCs that inherently provide zero sampling uncertainty. As the number of receive channels keeps increasing and the power and PCB size requirements become more stringent, these types of ADCs do not scale well. Another option is to use ADCs with high speed serial outputs, such as JESD204B, which solve the issues the parallel ADCs have. This option comes with an increased complexity on the data interface and makes it harder to achieve zero sampling uncertainty.

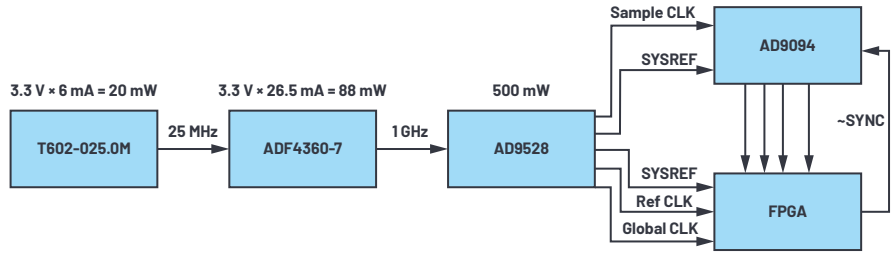

Figure 3. DAQ board clocking and data path.

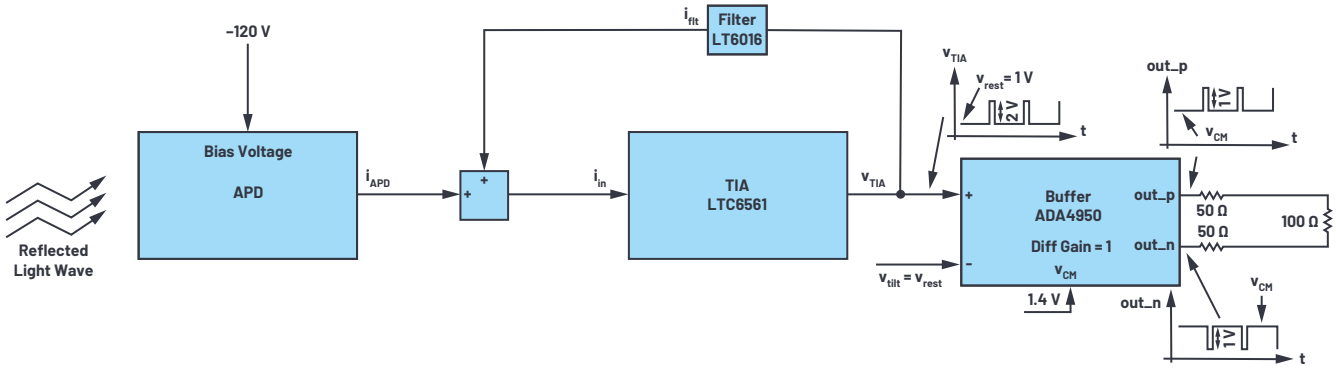

Figure 4. AFE board signal chain.

The LIDAR DAQ board provides a solution to these challenges by showing how to design the power, clocking, and data interface for a JESD204B data acquisition system operating in Subclass 1 mode, ensuring deterministic latency so that zero sampling uncertainty can be achieved, while taking advantage of all the benefits that the JESD204B interface provides and having the lowest possible power for the clocking scheme. To operate in JESD204B Subclass 1 mode, there is a total of five clocks needed in the system:

- ADC sampling clock: drives the ADC signal sampling process.
- ADC and FPGA SYSREF: source synchronous, high slew rate timing resolution signals responsible for resetting device clock dividers to ensure deterministic latency.
- FPGA global clock (also referenced as core clock or device clock): drives the output of the JESD204B PHY layer and FPGA logic.
- FPGA reference clock: generates the PHY layer internal clocks needed by the JESD204B transceivers; needs to be equal to or an integer multiple of the device clock.

All the clocks are generated by one AD9528 JESD204B clock generator, thus ensuring they are all synchronized with each other. Figure 3 shows the clocking scheme and the data interface with the FPGA.

The AFE board receives the optical reflected signal, converts it to an electrical signal, and transfers it to the ADC on the DAQ board. This board is probably the most sensitive part of the entire design since it mixes signal condition circuits working with microampere current signals generated by the 16-channel APD array, converting the optical signal to an electrical signal, with high voltage power supplies in the range of –120 V to –300 V needed to power the same APD. The 16 current outputs are fed into four low noise, 4-channel, transimpedance amplifier (TIA) LTC6561s with an internal 4-to-1 mux to select the output channel that is afterward fed into one of the four ADC inputs. The input section of the TIAs needs a lot of attention to achieve the desired level of signal integrity and channel isolation so that there is no additional noise added to the very low current signal generated by the APD, thus maximizing the SNR and the object detection rate of the system. The design of the AFE board shows the best practices to achieve the maximum signal quality by keeping the length of the traces between the APD and the TIA as short as possible, adding vias in between the TIA inputs for maximum channel-to-channel isolation, and positioning the signal condition circuits so that they do not interfere with the other power circuits on the board. Another important feature is the ability to measure the temperature of the APD to be able to compensate for APD signal output variations due to temperature changes as the APD's temperature increases during normal operation. A few knobs are provided to control the offsets of the signal chain and the APD bias, which translates into APD sensitivity, to be able to maximize the ADC input range for maximum SNR. Figure 4 shows a block diagram of the AFE board signal chain.

The laser board generates the optical pulses with a wavelength of 905 nm. It uses four lasers that are driven simultaneously for an increase in beam strength, resulting in a longer measurement range. A PWM signal generated by the FPGA carrier board with programmable pulse width and frequency is used to control the lasers. The signal is generated on the FPGA as LVDS to make it less susceptible to noise as it travels from the FPGA to the laser board through the DAQ board and the ribbon cable connecting the DAQ and the laser boards. The drive signal can be fed back to one of the ADC channels for time of flight reference. An external power supply is used to power the lasers. The design complies with International Standards IEC 60825-1:2014 and IEC 60825-1:2007 for a Class 1 laser product.

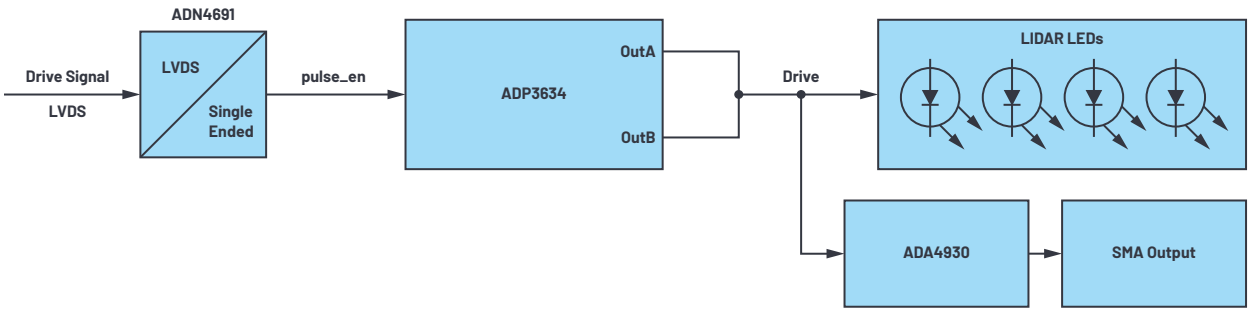

Figure 5. Laser board signal chain.

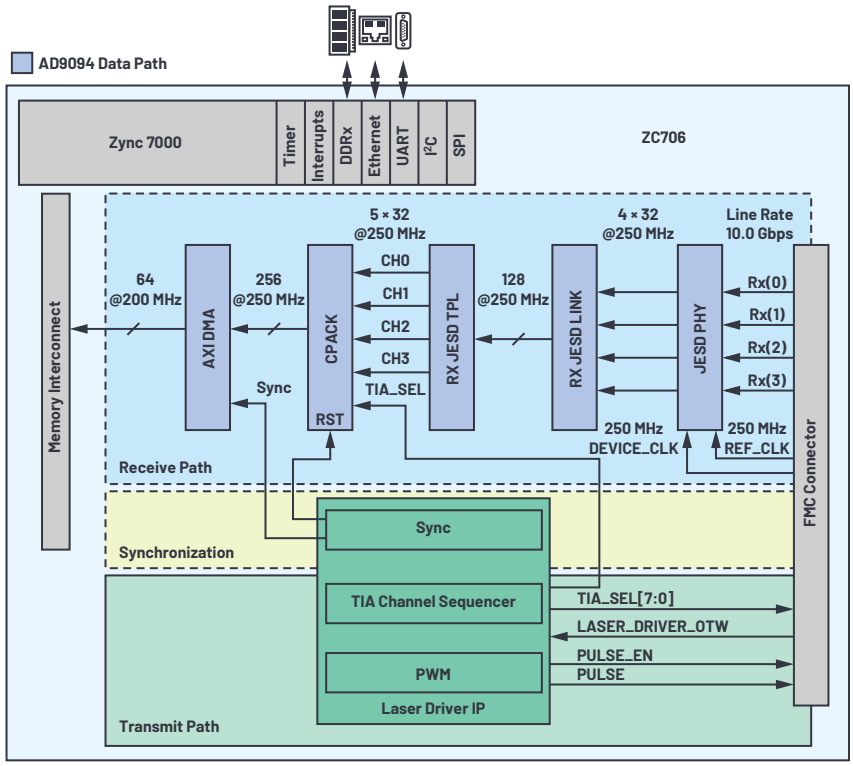

Figure 6. HDL design block diagram.

Both the AFE and laser boards require optics for long distance operation. The system was proven to operate at 60 m using fast axis collimators[1] for the laser diodes that narrow the vertical FoV to 1° while keeping the horizontal field of view unchanged, and an aspherical lens for the receive side.

HDL Reference Design

The HDL design constitutes the primary interface to the hardware and implements all the logic to transfer data from the JESD link to the system's memory, drive the lasers, synchronize the receiver and transmitter for accurate time of flight measurement, and implement the communication interfaces to all the components in the hardware design. Figure 6 shows a simplified block diagram of the HDL design. The generic architecture of the ADI's HDL reference designs makes the framework scalable and more accessible to port to another FPGA carrier. The design is using the Analog Devices JESD204B framework[2] along with several SPI and GPIO interfaces to receive data from the AD9094 ADC and to control all the devices on the prototyping platform.

The JESD204 link is configured to support four data converters (M) using four lanes (L) with a lane rate of 10 Gbps for a converter resolution of eight bits. The device clock, which will be the same as the reference clock for the high speed transceivers, is set to 250 MHz and is provided by the DAQ board. The link is running in Subclass 1 mode, which ensures deterministic latency between the high speed converter and FPGA.

One of the biggest challenges of a LIDAR system is how to synchronize various functions to the transmitted pulse and how to process just the necessary amount of data, received from the high speed ADC. To address this, the HDL design contains an IP that provides the required logic to generate a laser pulse, control the internal multiplexers of the TIAs and provides a back-pressure for the DMA. All these control functions are synchronized with the transmitted pulse, so the system doesn't have to save all the raw digitized high speed data stream. This way, the overall data rate of the system is highly reduced.

Software

The key points that define the software stack for the LIDAR platform are *free* and *open source*. The user gets through them the "freedom to run, copy, distribute, study, change, and improve the software."[3] Starting from the Linux® kernel and continuing with the user space tools, everything respects this.

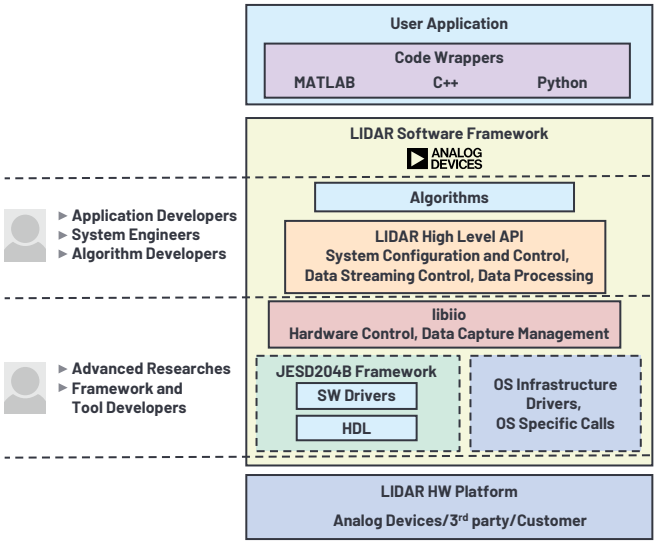

Figure 7. Software stack.

The software drivers implemented in the kernel initialize the hardware components and expose to the user all their useful functionalities. Most of these drivers are part of the Industrial I/O (IIO) Linux subsystem.[4] The drivers are all platform agnostic, so hardware changes, including carrier ones (for example, migrating from a Xilinx® FPGA to an Intel® one), do not require them to be changed.

To ease the development of software interfacing IIO devices, ADI has developed the libiio library.[5] It abstracts the low level details of the hardware and provides a simple yet complete programming interface that can be used for advanced projects. The various available libiio back ends (for example, local, network, USB, serial) make possible the usage of IIO devices locally and remotely from applications running on different operating systems (for example, Linux, Windows®, macOS®).

The IIO oscilloscope, developed by ADI, is an example of an application that uses libiio to interface with IIO devices and can be used during the system evaluation stage. The tool can capture and plot data in different modes (for example, time domain, frequency domain, constellation, cross-correlation), transmit data, and allow users to view and modify the settings of the detected devices.

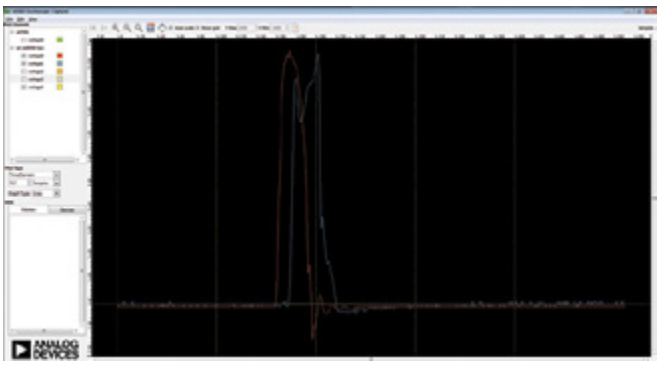

Figure 8. IIO oscilloscope capture window displaying LIDAR data.

While libiio provides a low level programming interface, in most cases, users expect a platform-specific API that abstracts the low level driver calls and exposes a set of functions to access and configure the various system parameters and stream data from the system. For this reason, the LIDAR prototyping platform comes with a specific API with bindings for popular frameworks and programming languages such as C/C++, MATLAB®, or Python®,[6] enabling users to interface with the system using their preferred programming language and focus on algorithms and applications, which is the value proposed to their customers.

Conclusion

In any system design, when the architecture is being established and design decisions are being made, there is a certain degree of ambiguity. This represents the risk that the system will not work or perform as expected after being built, resulting in multiple design cycles, increased development cost, and longer time to market for a product. Being built on pre-engineered systems intended to interoperate with each other, reference designs reduce risk and improve overall predictability and reliability as compared to using a custom, one-off design built from scratch. Using a reference design as a starting point in the planning process helps bring new designs to market faster and helps ensure there are fewer surprises and problems. System designers are always seeking reference platforms to prove out the design decisions, thus reducing the risk and improving reliability. Starting a project with clear and standardized design options facilitates the planning process. It does so by using a common language to help align goals, encourage cooperation and participation across multiple functions, and make it easier to evaluate the trade-offs between design goals. The LIDAR prototyping platform tries to meet these needs by providing open-source hardware and software designs that can be referenced in the initial system architecture phases. The hardware platform and the software stack can be used through all the phases of product development, from initial system evaluation, development, and integration into the final product. The contents of the reference design such as engineering drawings and a BOM provide a head start in getting to buildable, legal, and localized design system. This shortens the design cycle and likely saves money in the process. The modular hardware design allows various configuration options to meet the specific application requirements, while the open-source software stack, based on industry-standard frameworks and programming languages, and accompanied by application examples, allows customers to focus on developing the applications that bring value to their products without having to spend effort on the low level portions of the stack.

References

[1] Fast Axis Collimator Lenses. FISBA, 2019.

[2] JESD204 Interface Framework. Analog Devices, Inc., 2019.

[3] What Is Free Software? Free Software Foundation, Inc., 2019.

[4] The Linux Driver Implementer's API Guide. Analog Devices, Inc., 2019.

[5] About libiio. Analog Devices, Inc., 2019.

[6] "PyADI-IIO: Analog Devices Python Interfaces for Hardware with Industrial I/O Drivers." GitHub, Inc., 2020.

Michael Hennerich and Robin Getz. "Free and Open-Source Software—An Analog Devices Perspective." *Analog Dialogue*, Vol. 44, No. 3, Mar. 2010.

About the Author

István Csomortáni is an FPGA design engineer for ADI, supporting the design and development of FPGA-based reference designs. He holds a B.S. degree in industrial automation and informatics and an M.Sc. degree in integrated circuits design. He has worked for ADI since 2012, supporting various system-level reference design for high speed converters and RF transceivers. He can be reached at *istvan.csomortani@analog.com*.

About the Author

Dragos Bogdan currently leads a small embedded software team inside the SDG group that adds open-source Bare Metal and Linux support for various types of platforms and components. Dragos joined ADI in 2011 as a software engineer. From 2010 to 2011, he was with Pergamon RD, where he developed embedded hardware and software for printing equipment. Prior to that, he participated in internship programs at National Instruments and Continental Automotive. He holds a bachelor's degree in electronics and a master's degree in automation from Technical University of Cluj-Napoca. He can be reached at *dragos.bogdan@analog.com*.

About the Author

Cristian Orian is a system design engineer at ADI, working on the hardware development of evaluation platforms. He holds a Ph.D. in electronics. He has also worked on the design of power supplies. He can be reached at *cristian.orian@analog.com*.

About the Author

Andrei Cozma is an engineering manager for ADI, supporting the design and development of system-level reference designs. He holds a B.S. degree in industrial automation and informatics and a Ph.D. in electronics and telecommunications. He has been involved in the design and development of projects from different industry fields such as motor control, industrial automation, software-defined radio, and telecommunications. He can be reached at *andrei.cozma@analog.com*.

Wireless Water Quality Monitoring System

Piyu Dhaker, *Applications Engineer*

Introduction

Several industries such as beverage production, pharmaceutical plants, wastewater treatment plants, and so on, rely on water quality monitoring systems to measure and control important water quality indicators. Parameters defining physical, chemical, and biological characteristics of water can be used as water quality indicators. Examples include:

- Physical: temperature and turbidity
- Chemical: pH, oxidation reduction potential (ORP), conductivity, and dissolved oxygen
- Biological: algae and bacteria

The article focuses on chemical measurement parameters that have been historically indispensable and unreliable, posing an implementation burden. Electrochemistry is a branch of chemistry that characterizes the behavior of reduction-oxidation (redox) reactions by measuring the transfer of electrons from one reactant to another. Electrochemical techniques can be used directly or indirectly to detect and measure example water quality indicators previously listed. An electrochemical measurement system consists of two main blocks:

- Sensor: a device used to measure a water quality indicator and generate a corresponding electrical signal.
- Measurement and processing unit: circuitry that measures and processes the electrical signal.

Typically, individual wired sensors are deployed across the processing plants. The in-field sensors need to be cleaned and calibrated frequently and replaced often. Wireless networks could reduce some of this burden, yet they are not typically considered robust enough to deploy in the harsh environments of these applications.

In process, high reliability wireless sensing networks are now possible with new measurement and networking technologies. This article describes a demonstration platform that combines the universal sensor interface capability of the ADuCM355 and the wire-like reliability of Analog Devices' SmartMesh® IP technology to form a robust and low power, wireless water quality monitoring system focused on pH measurement as the example water quality indicator. The principle can readily be extended to other electrochemical parameters to form a suite of water quality measurements for each wireless sensing node.

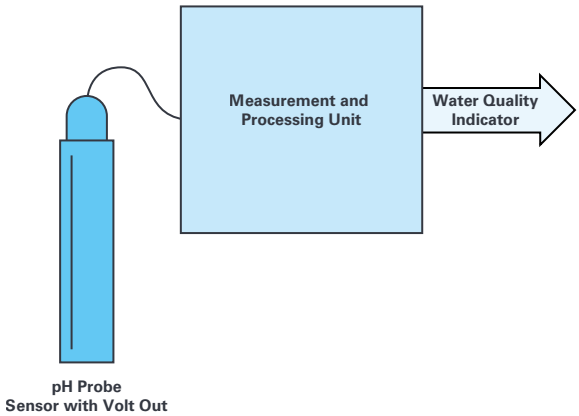

Figure 1. Typical electrochemical measurement system.

pH Measurement and pH Probe

The pH value is a measure of the relative amount of hydrogen and hydroxide ions in an aqueous solution. A neutral solution is one in which the hydrogen ion concentration exactly equals the hydroxide ion concentration. pH is another way of expressing the hydrogen ion concentration, measuring acidity or basicity of the solution, and is defined as:

$$pH = -\log_{10}(H+) \quad (1)$$

Where H+ is the hydrogen ion concentration in mol/liter.

The pH value of a solution ranges from 0 to 14, with a neutral solution having a pH of 7, an acidic solution having a pH less than 7, and an alkaline solution having a pH greater than 7.

The pH probe is an electrochemical sensor consisting of a glass electrode and a reference electrode.

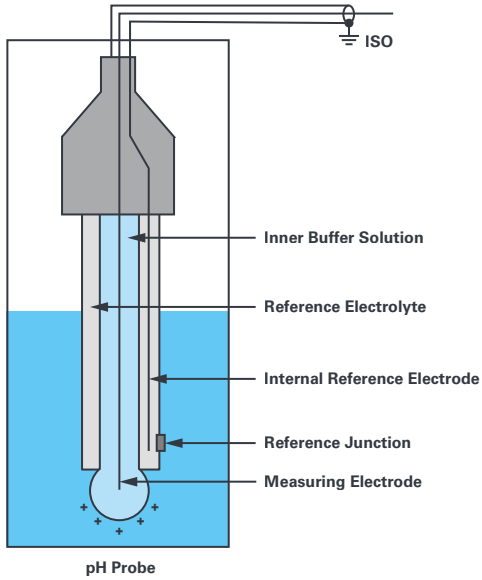

Figure 2. pH probe.

When the pH probe is inserted in the solution, the measuring electrode generates a voltage that depends on the hydrogen ion activity of the solution, which is then compared to the potential of the internal reference electrode. The difference between the measuring and reference electrode is the measured potential and is expressed in the Nernst equation as:

$$E = a - \frac{2.303\ R(T + 273.1)}{\text{nF}} \times (pH - pH_{ISO}) \quad (2)$$

Where:

E is the voltage of the electrode with unknown activity

a = ±30 mV, zero point tolerance

T is the ambient temperature in °C

n = 1 at 25°C, valence (number of charges on ion)

F = 96485 coulombs/mol, Faraday constant

R = 8.314 volt-coulombs/°K mol, ideal gas constant

pH = the hydrogen ion concentration of an unknown solution

pH_{ISO} = hydrogen ion concentration of the reference electrolyte; consult the probe documentation; typical pH_{ISO} = 7

The equation indicates generated voltage that varies with the pH in a known manner. It also shows that the generated voltage is directly proportional to the temperature of the solution. As the solution temperature increases, the potential difference between two electrodes increases and vice-versa. An ideal pH probe produces ±59.154 mV/pH units at 25°C.

Change in temperature can also change the sensitivity of the measuring electrode, which in turn causes measurement error. This error is predictable and can be accounted for by calibrating the probe across temperature, and then correcting for temperature during subsequent measurements. Typically, a temperature sensor is integrated into the pH probe. The temperature sensor could be a negative temperature coefficient (NTC) thermistor or RTD such as PT100 or PT1000. A pH probe with a temperature sensor is shown in Figure 3.

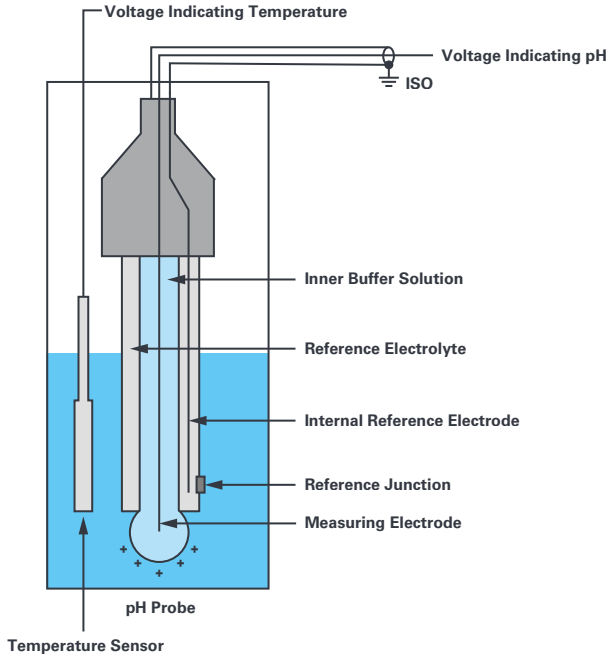

Figure 3. pH probe with temperature sensor.

If the temperature sensor measures change in temperature, the correction factor is applied to the final pH reading, and the meter then shows a corrected and more accurate reading. This mechanism works well to compensate for any error in pH that can arise due to temperature variations.

pH Measurement Unit with ADuCM355

ADuCM355 is the industry's most advanced and highly integrated chemical sensor measurement front end provides a platform solution for pH measurement with all the necessary measurement functions integrated with a low power microprocessor. The ADuCM355 enables a very low power measurement platform in a small form factor, making it small enough to be implemented inside the sensor housing while providing the capability and performance of benchtop instruments. Figure 4 shows ADuCM355's pH measurement board with BNC and RCA connectors for pH probe and temperature sensor connections. This board is from the CN-0428 reference design, and more details about the board can be seen in Figure 5.

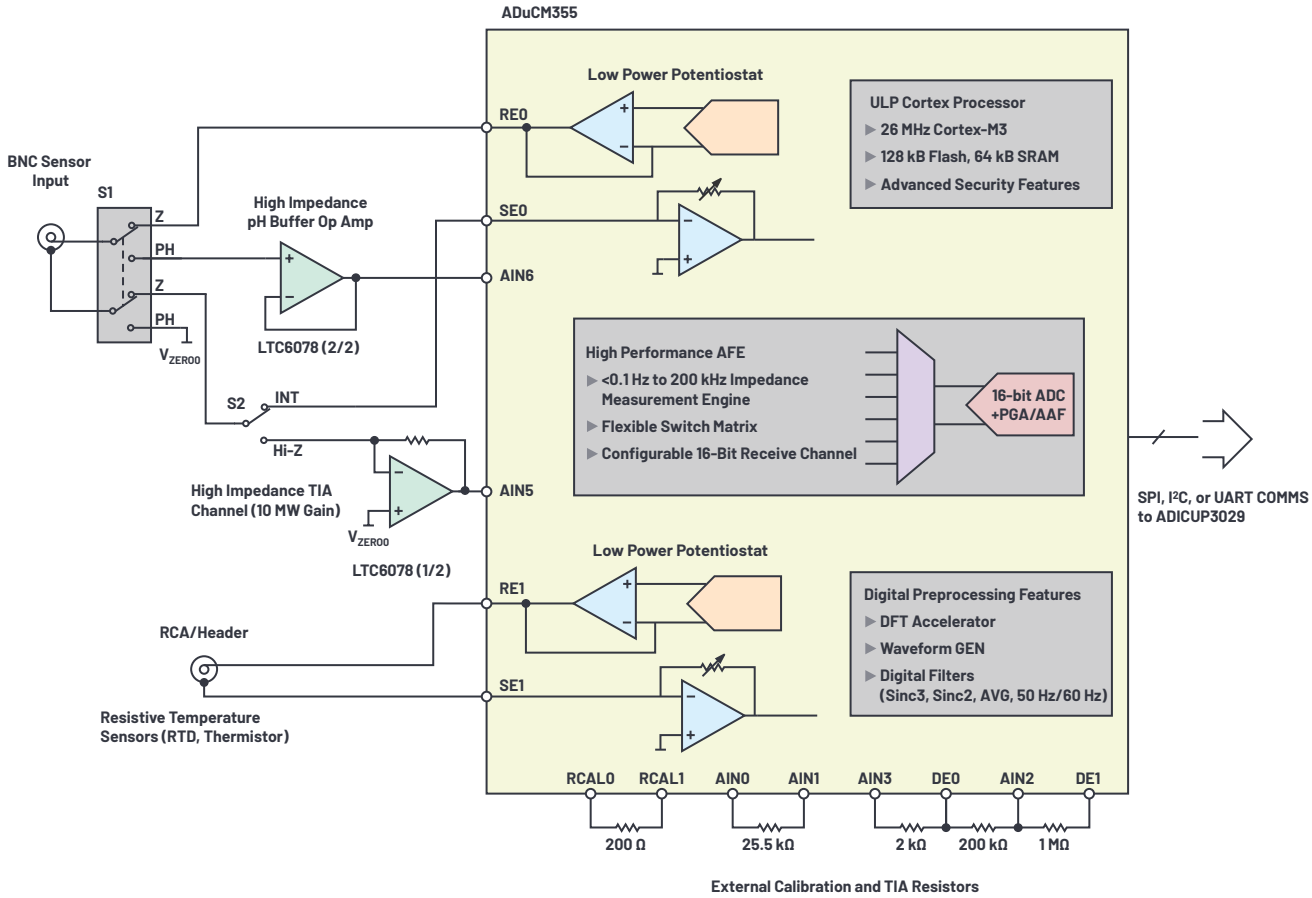

Figure 4. pH probe with integrated temperature sensor connected to ADuCM355.

Figure 5. ADuCM355 pH measurement PCB with BNC and RCA connectors.

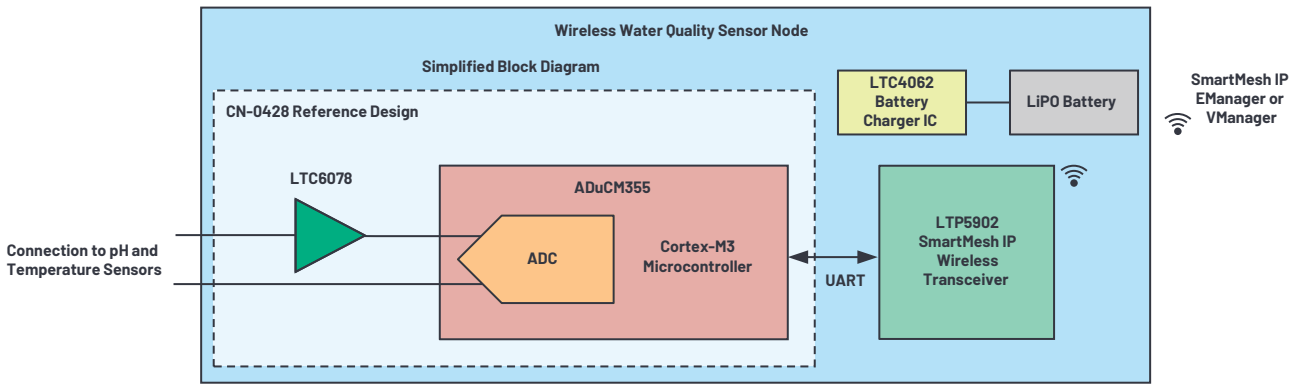

Figure 6. pH sensor connected to the ADuCM355 and the SmartMesh wireless sensor node.

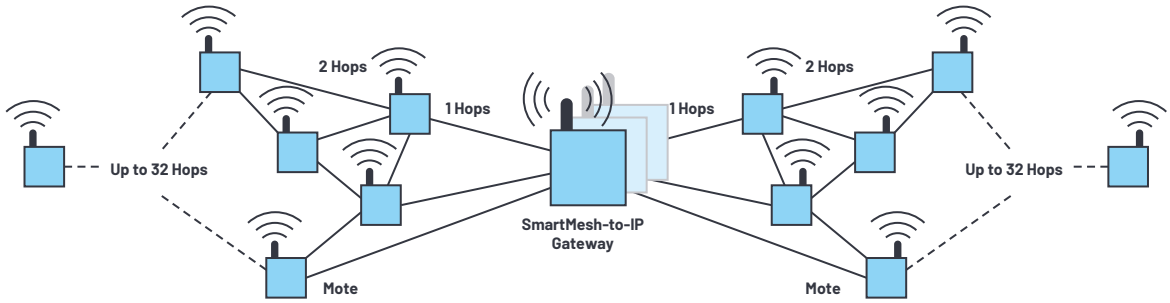

Figure 7. SmartMesh network with sensor nodes and network IP manager within the gateway.

Connecting pH Measurement Sensor Node to SmartMesh

By combining the ADuCM355 with an Analog Devices SmartMesh transceiver, we created a small, low power pH measurement sensor node. The ADuCM355 provides measured pH data as a digital output. This digital data is then connected via UART to Analog Devices' LTP5902 SmartMesh IP wireless transceiver. The LTP5902 transmits digital data over SmartMesh network to the SmartMesh IP manager.

SmartMesh is Analog Devices' proprietary 2.4 GHz multihop wireless mesh networking solution based on IEEE 802.15.4e standards. It includes AES 128 encryption and authentication, providing robust end-to-end security. It is ultra low power and energy efficient, enabling each sensor node to run on batteries.

SmartMesh networks communicate using a time slotted channel hopping (TSCH) link layer, which provides triple-play redundancy. The SmartMesh network manager (part of the gateway) coordinates the schedule, manages security, performs over-the-air programming (OTAP), and automatically optimizes connectivity 24 hours a day, 7 days a week. The network manager also provides detailed network health reports via an API. For small networks, one embedded manager can support up to 100 sensor nodes (also called motes). Truly enormous installations of up to 50,000 nodes are supported with VManager.

Rigorous network stress testing ensures >99.999% data reliability, making SmartMesh the perfect solution for industrial wireless sensor networks that must maintain high network availability and not lose packets.

Wireless Water Quality Monitoring System:

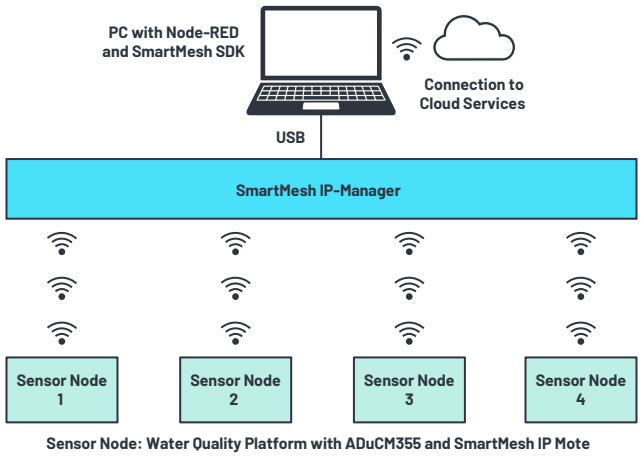

Figure 8. Wireless water quality monitoring system using ADuCM355 and SmartMesh.

The wireless water quality network demonstration as shown in Figure 8 consists of:

- Four sensor nodes:
  - Each sensor node consists of an off-the-shelf glass electrode pH probe with an integrated temperature sensor connected to ADuCM355 and SmartMesh IP mote as shown in Figure 6.
  - The pH probe senses the pH, ADuCM355 performs the measurement and calculations, and provides measured pH in digital output, which is then transferred over the wireless SmartMesh network to the SmartMesh IP manager.
  - SmartMesh IP manager connected via USB to the PC.
- A gateway in this system is performed by a PC. This PC has Node-RED and SmartMesh SDK installed. The SmartMesh SDK is used to create a JavaScript object notation (JSON) server for the data, which is connected to Node-RED. Node-RED is used to display measured pH data from each sensor node and allows connection to cloud services like IBM Watson, Amazon AWS, etc.

Hardware Setup

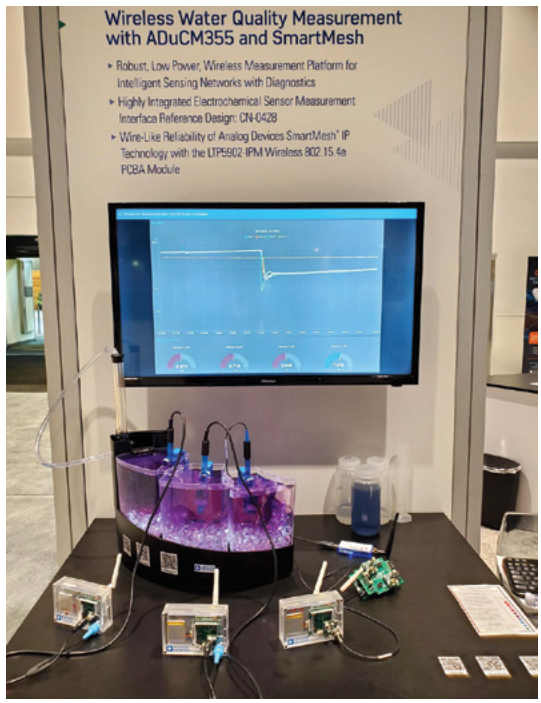

Figure 9. Wireless water quality network.

For demonstration purpose, we are using a staggered 3-chamber fish tank with water flowing from the top to subsequent chambers. A pH sensor probe is immersed in each chamber. We have a fourth sensor placed in a reference solution at a distance (not shown in Figure 9) to illustrate SmartMesh wireless communication over distance. As we change the pH of the solution in the top chamber, the data on Node-RED is updated, indicating a new pH. As this new pH solution flows from the top to subsequent chambers, the other two pH sensors update their measurements and the data is displayed on the screen. Since the fourth sensor is immersed in a reference solution with no pH change, the readings from this sensor are constant. More information on Node-RED and measured data is provided in following sections. You can also view the demo recording here.

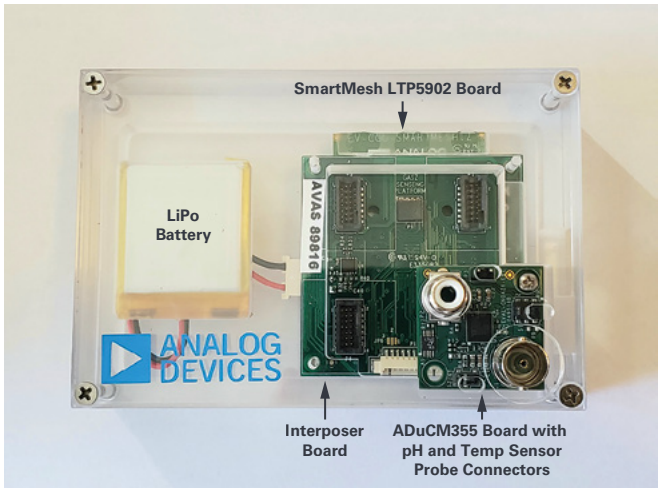

Figure 10. Sensor node.

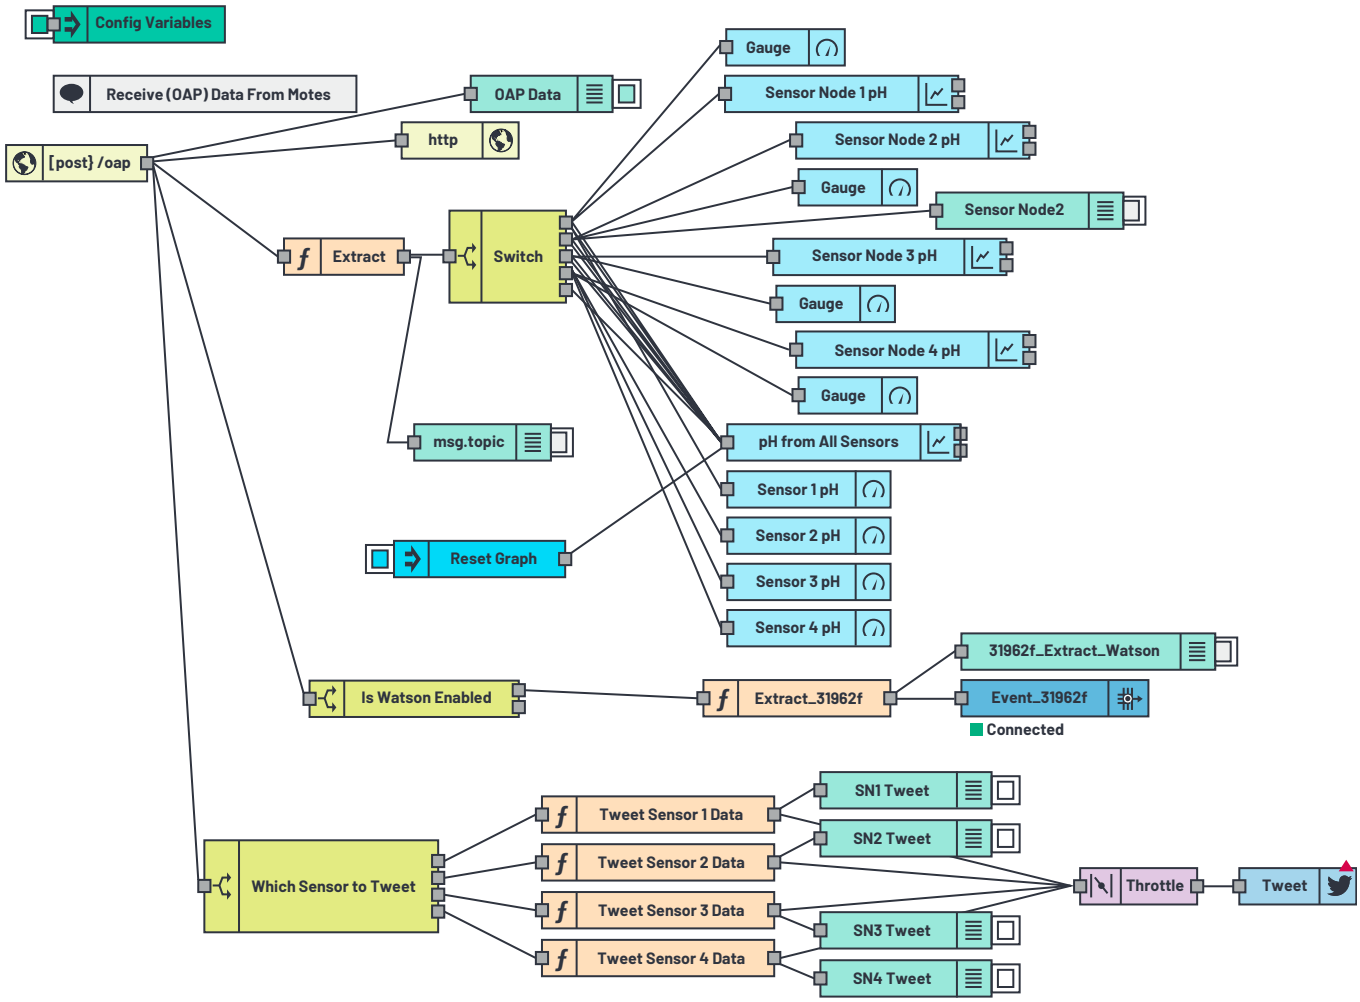

Figure 11. JSON flow.

Measurement Data:

pH measurement from four sensor nodes is displayed on the PC using Node-RED.

Node-RED is a programming tool with a web-based browser that allows connection of hardware devices, APIs, and other online services together. The JSON flow for the demo is displayed in Figure 11.

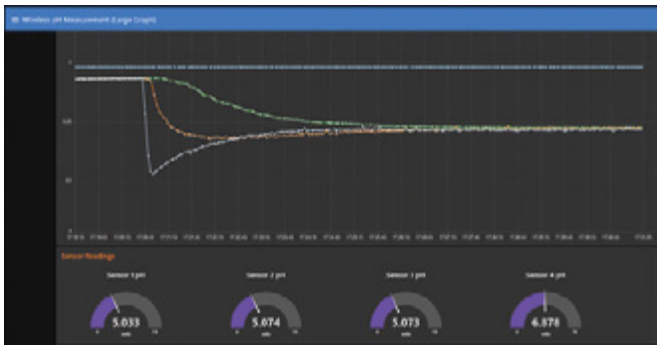

Figure 12. Wireless water quality demo dashboard.

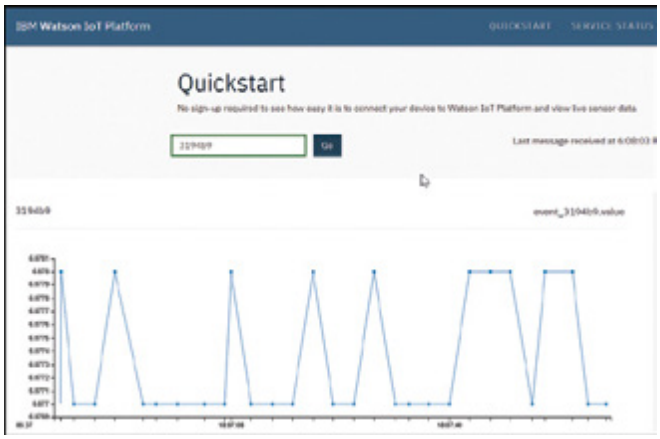

Figure 13. pH measurement data as displayed on IBM Watson.

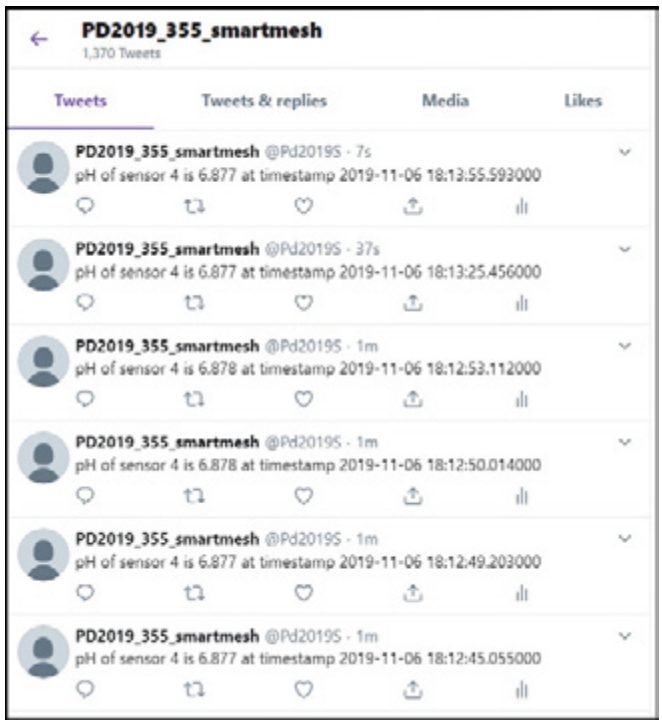

Figure 14. pH measurement data as a tweet.

Conclusion

This article describes the wireless water quality monitoring system using Analog Devices' ADuCM355 and SmartMesh IP technology. The small form factor and low power features of these products allow the sensor node to be powered by a battery. Robust SmartMesh technology reliably transfers data, even in harsh external environments. The demonstration illustrates a highly reliable wireless monitoring system and connectivity to cloud. This has tremendous potential depending on the end application, as this technology allows for monitoring of water quality at a location that is hard to reach, enables users to create alarms and warnings for different water quality thresholds, and leverages data to get more robust and better information about the continuous water quality.

References

CN-0398: Soil Moisture and pH Measurement System with Temperature Compensation. Analog Devices, Inc., January 2020.

CN-0409: Low to High Level Water Turbidity Measurement System. Analog Devices, Inc., January 2020.

IBM Watson IoT Platform. IBM Corporation, 2019.

Kämmerer, Christoph. Liquid Measurements—From Water to Blood. Analog Devices, Inc., October 2019.

Node-RED: Low Code Programming for Event-Driven Applications. Node-RED, January 2020.

Tzscheetzsch, Thomas. "Isolated pH Monitor with Temperature Compensation." Analog Devices, Inc., July 2019.

Water Quality. Fondriest Environmental, Inc., January 2020.

@Pd2019S. "pH of sensor 4 is 6.877 at timestamp 2019-11-06 18:13:55.593000." Twitter, November, 2019, 6:13 PM.

Acknowledgements

The author would like to thank Scott Hunt and Bill Lindsay for their assistance during the demo development and also for taking time to review the article. I would also like to thank Dan Braunworth and Dan Burton for reviewing and proofreading the article, and sharing their insights on the article.

About the Author

Piyu Dhaker is an applications engineer in the North America Central Applications Group of Analog Devices. She graduated from San Jose State University in 2007 with a master's degree in electrical engineering. Piyu joined the North America Central Applications Group in June 2017. She also previously worked in the Automotive Power Train Group and Power Management Group within ADI. She can be reached at *piyu.dhaker@analog.com*.

Protecting and Powering Automotive Electronics Systems with No Switching Noise and 99.9% Efficiency

David Megaw, Senior Design Engineer

Introduction

Powering automotive electronics systems can be challenging due to requirements of high reliability while contending with a relatively unstable battery voltage. The variety of electrical and mechanical systems that interface with a vehicle's battery can cause wild voltage excursions in the nominal 12 V supply. In reality, 12 V can vary from –14 V to +35 V for extended periods of time and experience voltage spikes with extremes ranging from +150 V to –220 V. Some of these surges and transients arise from everyday use, others from fault conditions or human error. Regardless of cause, the damage they can produce in a vehicle's electronics system can be difficult to diagnose and expensive to fix.

The experience of automakers over the last century has led to a catalog of electrical conditions and transients that are known to disrupt operation and cause damage. The International Organization for Standardization (ISO) has compiled this industry knowledge into the ISO 16750-2 and ISO 7637-2 specifications for road vehicles. At a minimum, the power supply for an automotive electronic control unit (ECU) should survive these conditions without damage. For critical systems, functionality and tolerances must be maintained. This requires that the power supply regulates the output voltage through the transient to preserve ECU operation. Ideally, a complete power solution avoids the need for fuses, minimizes power dissipation, and features low quiescent current to support always-on systems without draining the battery.

ISO 16750-2 Conditions for Automotive Electronics Systems

Analog Devices has several publications covering the ISO 7637-2 and ISO 16750-2 specifications in detail, along with how to simulate them using LTspice.[1,2,3,4]

In its latest iteration, the ISO 7637-2 electromagnetic compatibility specification focuses on high amplitude (>100 V), short duration (150 ns to 2 ms) transients from relatively high impedance sources (2 Ω to 50 Ω). These voltage spikes can often be mitigated with passive components. Figure 1 shows ISO 7637-2 pulse 1 as defined and with an added 330 µF bypass capacitor. The capacitance reduces the spike's amplitude from –150 V to –16 V, well within the range of reverse battery protection circuitry. ISO 7637-2 pulses 2a, 3a, and 3b have significantly lower energy than pulse 1 and require less capacitance to suppress.

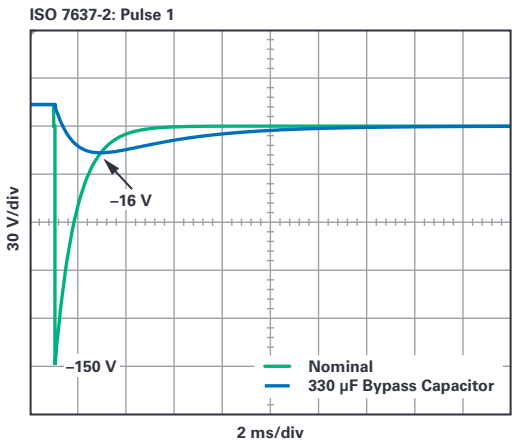

Figure 1. ISO 7637-2: pulse 1 with and without a 330 µF bypass capacitor.

ISO 16750-2 focuses on longer duration pulses from low impedance sources. These transients cannot easily be filtered and frequently require active, regulator-based solutions. Some of the more challenging tests include the load dump (test 4.6.4), the reverse battery condition (test 4.7), the superimposed alternating voltage test (test 4.4), and the engine starting profile (test 4.6.3). Figure 2 gives a visual overview of these test pulses. The variety of conditions presented in ISO 16750-2, along with the voltage and current requirements of the ECU, frequently require a combination of approaches in order to satisfy them all.

Load Dump

A load dump (ISO 16750-2: test 4.6.4) is a severe overvoltage transient that models a battery disconnect while the alternator is sourcing substantial current. The peak voltage during a load dump is classified as either suppressed or unsuppressed, depending on whether avalanche diodes are used on the outputs of the 3-phase alternator. A suppressed load dump pulse is limited to 35 V, whereas an unsuppressed pulse peak ranges from 79 V to 101 V. In either case, it can take up to 400 ms to recover due to the large amount of magnetic energy stored in the alternator's stator windings. While most automakers utilize avalanche diodes, increasing reliability demands are driving some manufacturers to require that ECUs satisfy peak load dump voltages approaching that of the unsuppressed case.

One solution for facing load dumps is adding a transient voltage suppressor (TVS) diode to locally clamp the ECU supply. A more compact and tighter tolerance approach is to use an active surge stopper, such as the LTC4364, which linearly controls a series N-channel MOSFET to clamp the maximum output voltage to a user-programmed level (for example, 27 V). Surge stoppers add the ability to disconnect the output, allowing programmable current limit and undervoltage lockout, and often provide reverse battery protection when back-to-back NFETs are used.

The concern with any linearly regulated power device such as a surge stopper is the potential for significant power dissipation in the N-channel MOSFET(s) when limiting the output voltage during a load dump, or when limiting current with a shorted output. The safe operating area (SOA) constraints of the power MOSFET ultimately limit the maximum current possible with a surge stopper. It also puts a time limit (typically set with a programmable timer pin) on how long regulation can be maintained before the N-channel MOSFET must be shut off to avoid damage. These SOA imposed limitations become more acute at higher operating voltages, making surge stopper use trickier for 24 V and 48 V systems.

A more scalable approach is to use a buck regulator capable of operating with a 42 V input, such as the LT8640S. A switching regulator doesn't have the MOSFET SOA limitations of a linear regulator, but it is certainly more complex. The efficiency of a buck regulator allows for very high current operation, and its top switch permits output disconnect and current limiting. Concerns about buck regulator quiescent current have been put to rest with the latest generation of parts that draw only a few microamps while in regulation under no load conditions. Switching noise has also improved substantially with Silent Switcher® technology and spread spectrum frequency modulation techniques.

Additionally, some buck regulators are able to operate at 100% duty cycle such that the top switch is on continuously, passing the input voltage to the output through the inductor. Switching operation is triggered during overvoltage or overcurrent conditions to limit output voltage or current, respectively. These buck regulators, such as the LTC7862, act as switching surge stoppers, achieving low noise, low loss operation while still retaining the robustness of a switch-mode power supply.

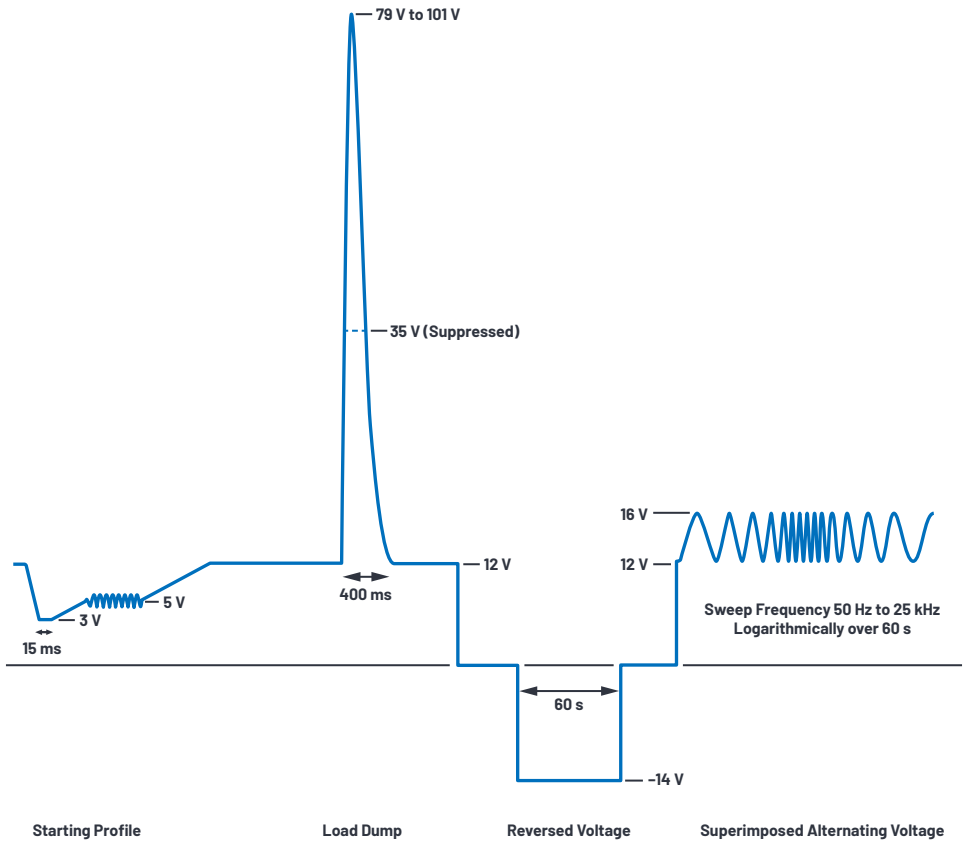

Figure 2. Overview of some of the tougher ISO 16750-2 tests.

Reverse Voltage

A reverse voltage condition (also known as a reverse battery condition) occurs when battery terminals or jumper cables are connected backward due to operator error. The relevant ISO 16750-2 pulse (test 4.7) applies –14 V to the DUT for 60 second durations, repeatedly. Some manufacturers add their own dynamic versions of this test where the part is initially powered (for example, V_{IN} = 10.8 V) before a reverse bias (–4 V) is abruptly applied.

A quick survey of data sheets shows that few ICs are designed to tolerate negative biases, with IC absolute minimum pin voltages typically limited to –0.3 V. Voltages more than a diode below ground can cause excessive current flow through internal junctions such as ESD protection devices as well as body diodes of power MOSFETs. Polarized bypass capacitors such as aluminum electrolytics may also be damaged during the reverse battery condition.

A Schottky diode can prevent reverse currents, but this approach leads to significant power loss at higher forward currents in normal operation. A simple protection scheme based on series P-channel MOSFET shown in Figure 3 reduces this loss, but may not work well at low input voltages (for example, engine start) due to the device threshold voltage. A more efficient approach is to use an ideal diode controller, such as the LTC4376, which drives a series N-channel MOSFET that cuts off input voltages below ground. In normal operation, an ideal diode controller regulates the source to drain the voltage of the N-channel MOSFET to 30 mV or less, reducing forward voltage drop and power dissipation by more than an order of magnitude compared to a Schottky diode.

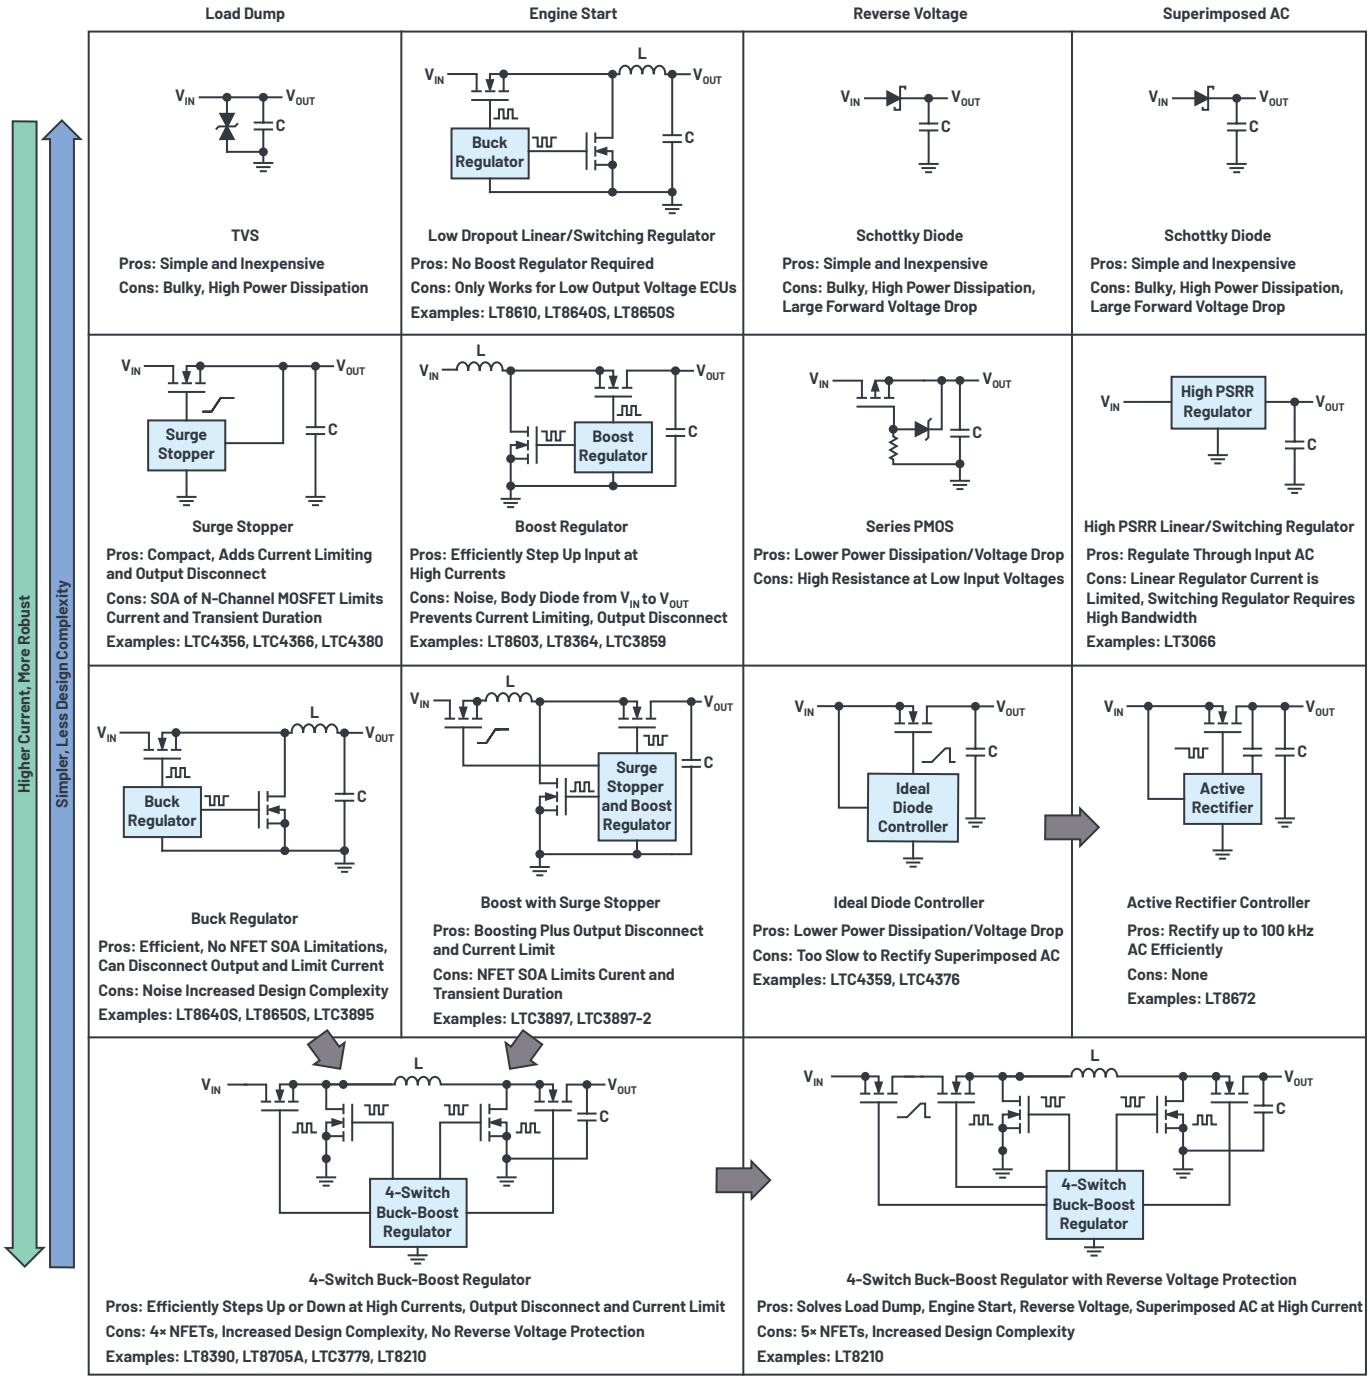

Figure 3. Different approaches to solving difficult ISO 16750-2 tests.

Superimposed Alternating Voltage

The superimposed alternating voltage test (ISO 16750-2: test 4.4) models the impact of the ac output of the vehicle's alternator. As the name implies, a sinusoidal signal is superimposed on the battery rail with peak-to-peak amplitude of 1 V, 2 V, or 4 V, depending on severity level classification. For all severity levels, the maximum input voltage is 16 V. The frequency of the sinusoid is swept logarithmically from 50 Hz to 25 kHz and then back to 50 Hz over 120 seconds and repeated a total of five times.

This test causes large amplitude current and voltage swings in any connected filter network with a resonance below 25 kHz. It can also cause problems for switching regulators where loop bandwidth limitations make it a struggle to regulate through high frequency input signals. One solution is an intermediate rectifying element such as a power Schottky diode, but as with reverse voltage protection, this is a poor way to solve the problem.

An ideal diode controller will not work here like it did for reverse voltage protection because it cannot switch the N-channel MOSFET sufficiently fast to keep up with the input. The limiting factor is the gate pull-up strength, which is typically limited to 20 µA or so by an internal charge pump. While the ideal diode controller can quickly turn the MOSFET off, turn on is painfully slow, unsuitable for rectification of anything beyond very low frequencies.

A more elegant approach is to use the LT8672 active rectifier controller, which can toggle an N-channel MOSFET fast enough to rectify the input voltage at frequencies up to 100 kHz. An active rectifier controller is an ideal diode controller with two important additions: a large reservoir of charge boosted up from the input voltage and a powerful gate driver to toggle the N-channel MOSFET both on and off quickly. This method can reduce power loss by more than 90% compared to a Schottky. The LT8672 also protects downstream circuitry from a reverse battery condition just as an ideal diode controller would.

Starting Profile

The engine starting profile (ISO 16750-2: test 4.6.3) is an extreme undervoltage transient, sometimes referred to as the *cold crank* pulse because the worst-case battery droop occurs at lower temperatures. Specifically, when the starter turns over, the 12 V battery voltage can momentarily drop to 8 V, 6 V, 4.5 V, or 3 V, depending on severity level classification (I, IV, II, III, respectively).

In some systems, a low dropout (LDO) linear regulator or switching buck regulator is sufficient to allow power rails to ride through this transient, provided that ECU voltages are less than the lowest input voltage. For example, if the highest ECU output voltage is 5 V, and it must satisfy severity level IV (minimum input voltage of 6 V), then a regulator with a dropout voltage less than 1 V is enough. The lowest voltage segment of an engine starting profile lasts for only 15 ms to 20 ms, so a rectifying element (Schottky diode, ideal diode controller, active rectifier controller) followed by a large bypass capacitor may be able to ride through this portion of the pulse if the voltage headroom briefly dips below the regulator dropout.

If, however, the ECU must support voltages higher than the lowest input voltage, then a boost regulator is required. Boost regulators can efficiently maintain a 12 V output voltage from inputs less than 3 V at high current levels. There is one problem with a boost regulator, though: the diode path from input to output prevents disconnect, so current is not naturally limited at startup or due to a short. To prevent current runaway, specialized boost regulators, such as the LTC3897 controller, incorporate a surge-stopper front end to allow output disconnect and current limiting, as well as provide reverse voltage protection when back-to-back N-channel MOSFETs are used. This solution can address the load dump, engine start, and reverse battery conditions with a single integrated circuit, but available current is limited by the SOA of the surge-stopper MOSFET.

A 4-switch buck-boost regulator lifts this constraint by combining a synchronous buck regulator and a synchronous boost regulator through a shared inductor. This approach can satisfy both the load dump and engine starting profile tests without MOSFET SOA limitations on current level or pulse duration, while retaining the ability to disconnect the output and limit current.

The switching operation of a buck-boost regulator depends on the relationship between the input and output voltages. If the input is significantly higher than the output, the boost top switch turns on continuously while the buck power stage steps down the input. Similarly, when the input is significantly lower than the output, the buck top switch turns on continuously while the boost power stage steps up the output. When the input and output are roughly equal (within 10% to 25%), the buck and boost power stages switch simultaneously in an interleaved fashion. In this way, efficiency is maximized in the various switching regions (buck, buck-boost, boost) by limiting switching to only the MOSFETs required for regulation for input voltages above, roughly equal to, or below the output, respectively.

ISO 16750-2 Solution Summary

Figure 3 summarizes the various solutions to addressing the load dump, reverse input voltage, superimposed alternating voltage, and engine starting profile tests along with the pros and cons of each approach. Several key takeaways begin to emerge:

- A series N-channel MOSFET with its drain facing the input is extremely desirable because it can be used to limit current and disconnect the output whether it's used as a switch (for example, in a buck power stage) or linearly controlled (for example, in a surge stopper).
- In the case of reverse input protection and superimposed alternating voltage, using an N-channel MOSFET as the rectifying element (source facing the input) significantly reduces power loss and voltage drop compared to a Schottky diode.
- A switch-mode power supply is preferable to a linear regulator because it alleviates the reliability concerns and output current restrictions that come with operating within the SOA of the power device. It can regulate at input voltage extremes indefinitely, whereas linear regulator and passive solutions have inherent time limitations that complicate design.
- A boost regulator may or may not be necessary depending on the severity classification of the starting profile and the details of the ECU (what's the highest voltage it must provide).

If boost regulation is called for, then a 4-switch buck-boost regulator combines the above desirable traits into a single part. It can efficiently regulate through severe undervoltage and overvoltage transients at high current levels for extended periods of time. This makes it the most robust and straightforward approach, from an applications standpoint, despite the increased design complexity. Nevertheless, a typical 4-switch buck-boost regulator has some drawbacks. For one, reverse battery protection is not naturally provided and must be addressed with additional circuitry.

The primary shortcoming of the 4-switch buck-boost regulator is that it spends much of its operational life in the lower efficiency, noisier buck-boost switching region. When the input voltage is nearly equal to the output (V_{IN} ~ V_{OUT}) all four N-channel MOSFETs are actively switching to maintain regulation. Efficiency drops as a result of increased switching losses and application of maximum gate drive current. Radiated and conducted EMI performance suffer in this region as both the buck and boost power-stage hot loops are active and the regulator input and output currents are discontinuous.

A 4-switch buck-boost regulator can regulate through occasional large amplitude undervoltage and overvoltage transients, but at the price of operating with high quiescent current, lower efficiency, and increased noise in the far more common, nominal conversion region.

PassThru Operating Mode Brings High Efficiency and EMI Performance Buck-Boost Region

The LT8210 is a 4-switch buck-boost dc-to-dc controller that can operate conventionally with a fixed output voltage and also features a new PassThru™ operating mode (Figure 4), which eliminates the switching losses and EMI through a programmable input voltage window. It operates from 2.8 V to 100 V, allowing it to regulate through the most severe battery drop during engine start to the peak amplitude of an unsuppressed load dump. It comes with built-in reverse battery protection to –40 V, which is implemented with the addition of a single N-channel MOSFET (DG in Figure 5).

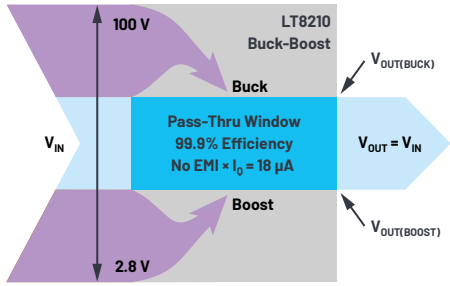

Figure 4. Buck-boost controller with PassThru mode solves many of the problems imposed by automotive standards tests.

In PassThru mode, the output voltage is regulated to the edges of a voltage window when the input voltage is outside the window. The top of the window and the bottom are programmed via the FB2 and FB1 resistor dividers. When the input voltage is within this window, the top switches (A and D) turn on continuously, passing the input voltage directly to the output. In this nonswitching state, the total LT8210 quiescent current drops to tens of microamps. The absence of switching means no EMI and no switching losses, enabling efficiencies of greater than 99.9%.

For those who want the best of both worlds, the LT8210 allows transitioning between its different operating modes on-the-fly by toggling the MODE1 and MODE2 pins. In other words, the LT8210 can operate as a conventional buck-boost regulator with fixed output voltage (in CCM, DCM, or Burst Mode™) in some situations and then change to PassThru mode as the conditions in the application change. This can be a useful feature for always-on systems and in start-stop applications.

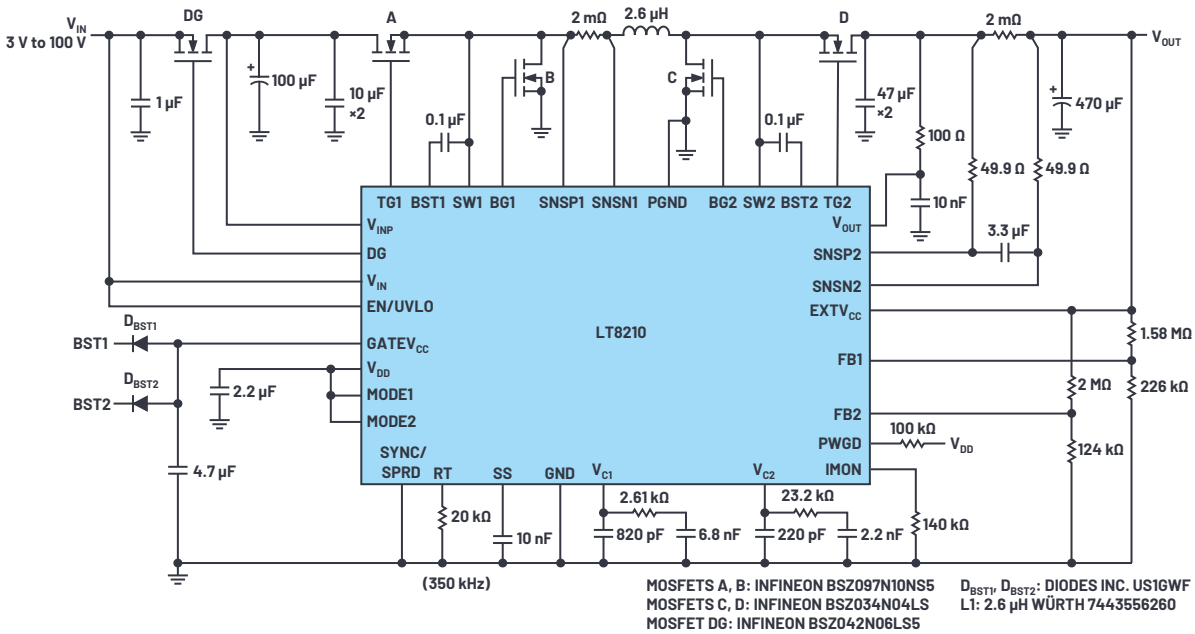

Figure 5. This 3 V to 100 V input buck-boost controller operates with 8 V to 17 V PassThru outputs.

PassThru Performance

The PassThru solution in Figure 5 passes the input to output in the window between 8 V and 17 V. When the input voltage is above the PassThru window, the LT8210 bucks it down to a regulated 17 V output. If the input drops below 8 V, the LT8210 boosts the output to 8 V. As a protection feature, switching operation is triggered within the PassThru window to control current if either the inductor current limit or programmed average current limit (via the IMON pin) are exceeded.

Figure 6, Figure 7, and Figure 8 show the response of the LT8210 circuit to the load dump, reverse voltage, and starting profile tests, respectively. Figure 9 and Figure 10 demonstrate the efficiency improvement and low current operation possible while in the PassThru window (it's amazing how efficient doing nothing can be). Figure 11 shows the dynamic transition between PassThru mode and CCM operation. An LTspice simulation of this circuit along with sped up versions of the toughest ISO 16750-2 test pulses is available at analog.com/media/en/simulation-models/LTspice-demo-circuits/LT8210_AutomotivePassThru.asc.

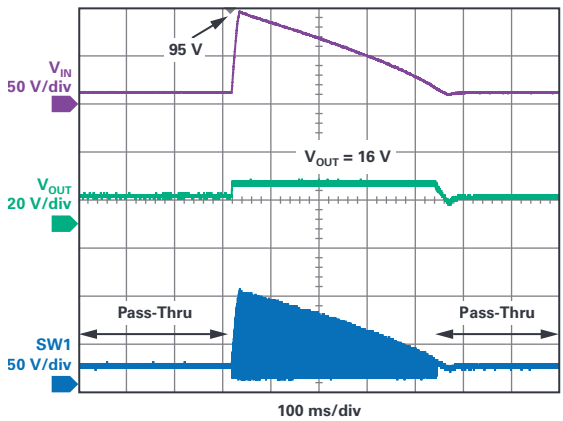

Figure 6. PassThru response to unsuppressed load dump.

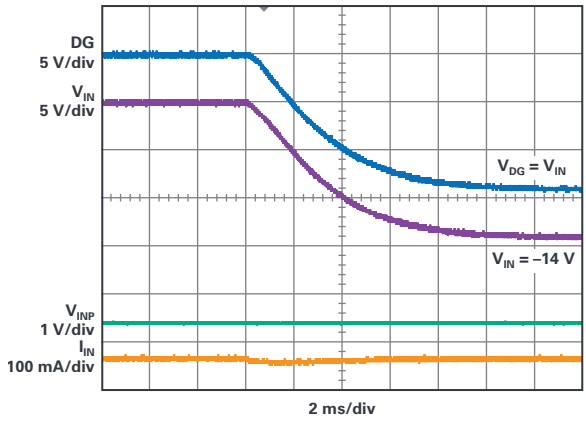

Figure 7. LT8210 response to reverse battery condition.

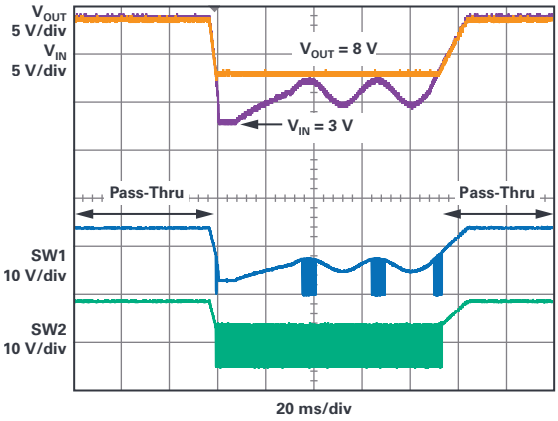

Figure 8. PassThru response to engine cold crank.

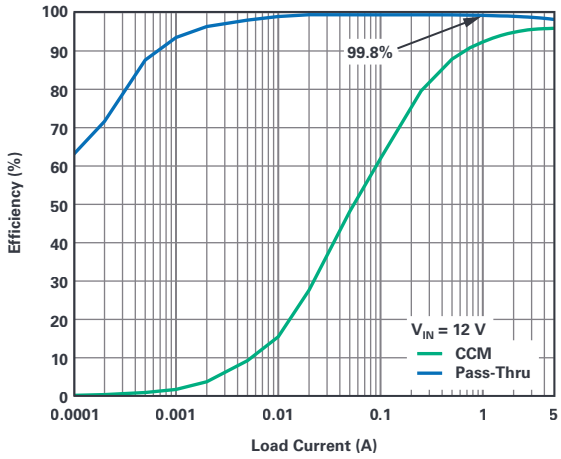

Figure 9. Efficiency for CCM and PassThru operation.

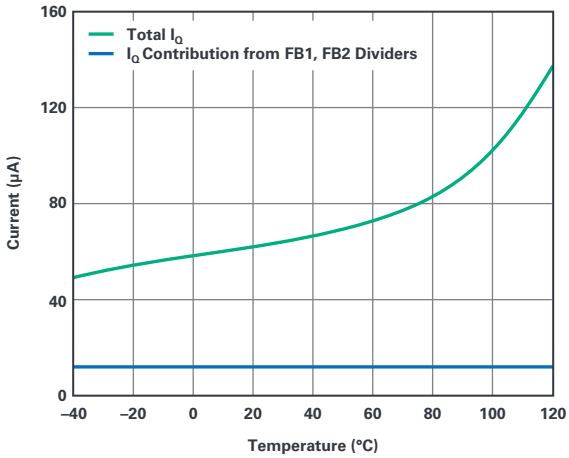

Figure 10. No load input current in PassThru mode (V_{IN} = 12 V).

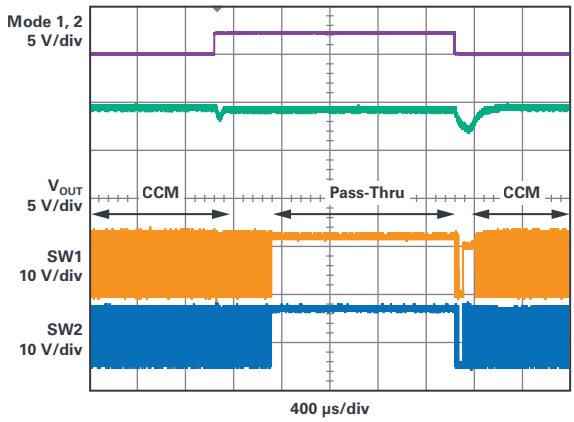

Figure 11. Dynamic transition between PassThru and CCM operation.

Conclusion

When designing a power supply for an automotive electronic system, the LT8210 4-switch buck-boost dc-to-dc controller offers a superior solution with its 2.8 V to 100 V input operating range, built-in reverse battery protection, and its new PassThru operating mode. PassThru mode improves upon buck-boost operation, achieving no switching noise, no switching losses, and ultralow quiescent current while regulating the output to a user-programmed window rather than a fixed voltage. The output voltage minimum and maximum values are bounded during large amplitude transients such as load dump and cold crank without MOSFET SOA concerns, or the current or timing limitations that come from linear solutions.

The novel LT8210 control scheme results in clean and fast transitions between the different switching regions (boost, buck-boost, buck, and nonswitching) allowing it to regulate through large signal, high frequency ac voltage on the input. For more details, please refer to the video at *analog.com/en/education/education-library/videos/6136638907001.html*. The LT8210 can be toggled between PassThru operation and conventional, fixed output voltage, buck-boost operating modes (CCM, DCM, or Burst Mode) while running, and the fixed output can be set to any voltage within the PassThru window (for example, V_{OUT} = 12 V for an 8 V-to-16 V window). This flexibility allows the user to alternate between PassThru and normal buck-boost operation to trade the low noise, low I_Q, high efficiency operation of PassThru mode for the tighter regulation and improved transient response in CCM, DCM, or Burst Mode.

References

[1] Dan Eddleman. "Low Quiescent Current Surge Stopper: Robust Automotive Supply Protection for ISO 7637-2 and ISO 16750-2 Compliance." *LT Journal of Analog Innovation*, January 2017.

[2] Christian Kueck. "Distilled Automotive Electronics Design." Analog Devices, Inc., April 2013.

[3] Bin Wu and Zhongming Yi. "Comprehensive Power Supply System Designs for Harsh Automotive Environments Consume Minimal Space, Preserve Battery Charge, Feature Low EMI." *Analog Dialogue*, Vol. 53, August 2019.

[4] Dan Eddleman. "LTspice: Models of ISO 7637-2 and ISO 16750-2 Transients." Analog Devices, Inc., 2019.

ISO 7637-2:2011. International Organization of Standardization, March 2011.

ISO 16750-2:2012. International Organization of Standardization, November 2012.

About the Author

David Megaw is a senior design engineer with ADI. He can be reached at *david.megaw@analog.com*.

RAQ Issue 174: Current Noise in FET Input Amps

Kaung Win, *Senior Field Applications Engineer*

Question:

Why is my design noisier at higher frequencies?

Answer:

The phenomenon of current noise increasing with frequency is well known to IC design engineers and circuit designers, but it was elusive to many engineers as a result of either too few articles in the field or incomplete information from manufacturers.

Many semiconductor manufacturer data sheets, including ADI's, specify the current noise of an amplifier in the specification tables, typically at a frequency of 1 kHz. It isn't always clear where the current noise specifications come from. Is it measured or is it theoretical? Some manufacturers are transparent in how they come up with this number by providing an equation of

$$i_n = \sqrt{2qi_b} \quad (1)$$

known as the shot noise equation. Historically, ADI had provided most current noise numbers this way. Does this calculated number hold up to 1 kHz for every amplifier?

Over the past few years, there has been a growing interest regarding current noise over frequency in amplifiers. Some customers—as well as manufacturers—assume that current noise for FET input amps follows a similar shape as bipolar input amplifiers—for example, 1/f or flicker noise component and flat wideband component, as shown in Figure 1. This is not the case in FET input amps; rather, in Figure 2, it looks like a bizarre noise shape that is not well known and is ignored in many simulation models.

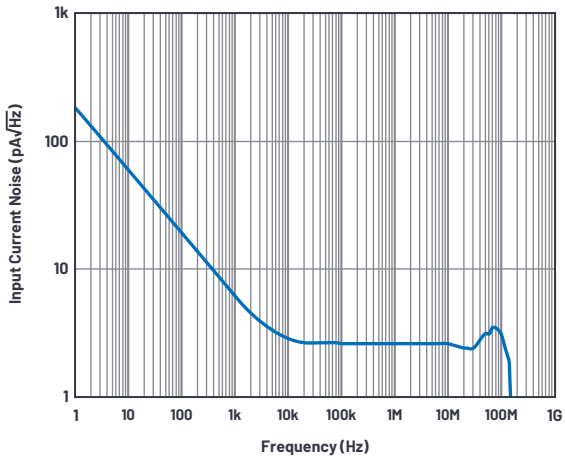

Figure 1. Current noise of AD8099, a bipolar input amplifier.

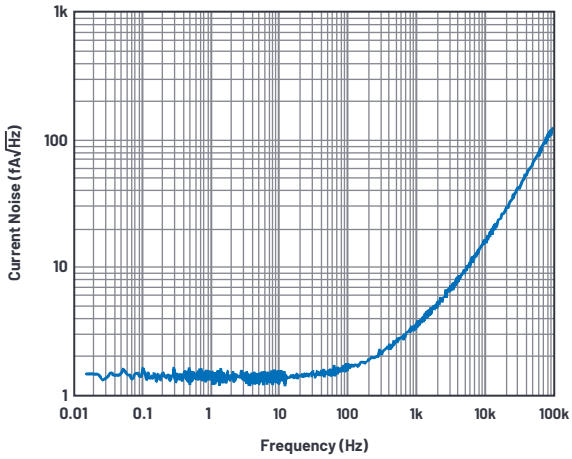

Figure 2. Current noise of AD8065, a FET input amplifier.

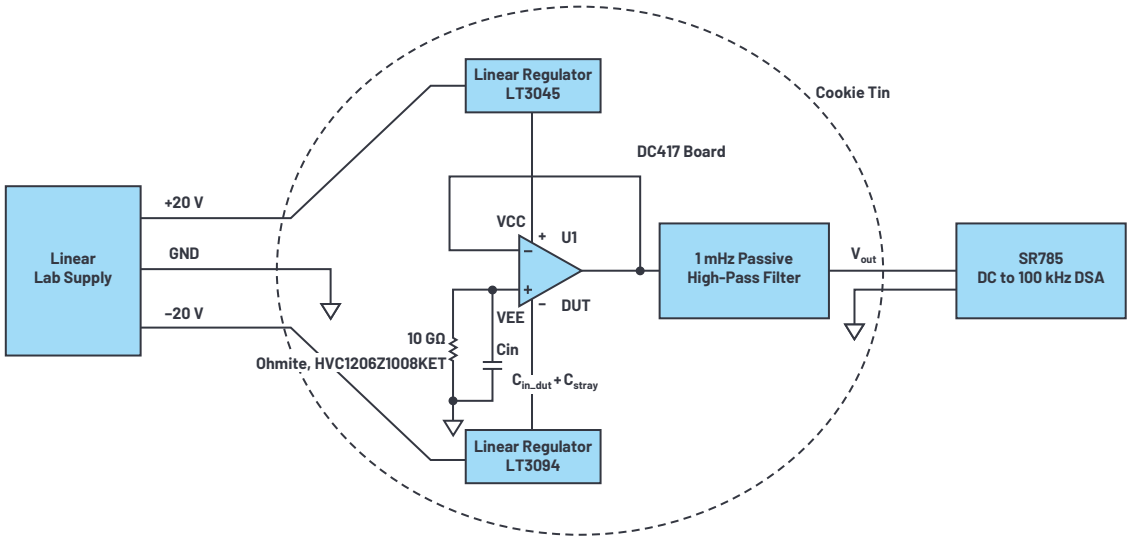

Figure 3. Measurement setup.

Measurement Setup Is the Key

Before we get into why this is the case, let's quickly look at the measurement setup. Getting an easy-to-reproduce, reliable measurement method is required so that the measurement can be repeated over many different parts.

The DC417B single amplifier evaluation board may be used. The power supplies to the device under test (DUT) must be low noise and low drift. Linear supplies are preferred over the switching supplies so that any supply variation, such as switching artifacts, does not add to the measurement. The LT3045 and LT3094, positive and negative ultrahigh PSRR, ultralow noise linear regulators, may be used to further reduce the ripple from the linear supply. Using the LT3045 and LT3094, a single resistor can be used to configure any output voltage necessary up to +15 V and down to –15 V. These two parts are ideal bench top supplies for low noise measurements.

A 10 GΩ SMT resistor from Ohmite (HVC1206Z1008KET) was used to convert current noise to voltage noise at the noninverting pin of the DUT. Typical bias current of FET input amps is about 1 pA, which equals 0.57 fA/$\sqrt{\text{Hz}}$

if the equation

$$i_{n\_dut} = \sqrt{2qi_b} \tag{2}$$

is correct. 10 GΩ source impedance thermal noise is

$$e_{n\_R} = \sqrt{4kTR} = 12.8\frac{\mu\text{V}}{\sqrt{\text{Hz}}} \tag{3}$$

This gives us the measurement current noise floor of

$$i_{n\_R} = \sqrt{\frac{4kT}{R}} = 1.28\frac{\text{fA}}{\sqrt{\text{Hz}}} \tag{4}$$

and it can be subtracted out in postprocessing. However, it becomes impossible to measure accurately if the resistor current noise dominates the current noise of the DUT. So, we would need a resistor value of at least 10 GΩ to see some of the noise. 100 MΩ source impedance thermal noise is about 1.28 μV/$\sqrt{\text{Hz}}$ (= 12.8 fA/$\sqrt{\text{Hz}}$) and it will not be enough to distinguish between DUT and resistor noise. The noise, if uncorrelated, adds in root sum squared (RSS) fashion. Figure 4 and Table 1 show the RSS impact on the ratio of two numbers. n:n adds about 41%, n:n/2 adds about 12%, n:n/3 adds about 5.5%, and n:n/5 about 2%. With enough averaging, we might be able to extract about 10% (0.57 fA/$\sqrt{\text{Hz}}$ and 1.28 fA/$\sqrt{\text{Hz}}$ RSS).

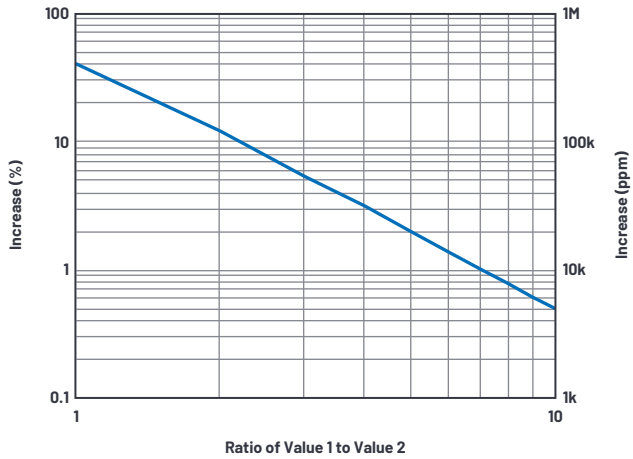

Figure 4. RSS addition based on the ratio of two numbers.

Table 1. RSS Addition Based on the Ratio of Two Numbers

Value 1	Value 2	RSS Sum	% Increase
n	n	1.414 n	41.42 %
n	n/2	1.118 n	11.80%
n	n/3	1.054 n	5.41%
n	n/4	1.031 n	3.08%
n	n/5	1.020 n	2.00%
n	n/6	1.014 n	1.38%
n	n/7	1.010 n	1.02%
n	n/8	1.008 n	0.78%
n	n/9	1.006 n	0.62%
n	n/10	1.005 n	0.50%

Why Are the Results So Strange?

Figure 5 shows the voltage noise density of the setup with the AD8065, a 145 MHz FET input op amp with a common-mode input impedance of 2.1 pF. The 10 GΩ resistor thermal noise is 12.8 μV/√Hz until the input capacitance along with the board and socket stray capacitance roll off the voltage noise. Ideally, this should keep rolling off at −20 dB/dec, but the curve starts to change shape around 100 Hz and flattens around 100 kHz. What's going on here? Our intuition tells us that the only way to stop the −20 dB/dec roll-off and cause a flatness is to provide a +20 dB/dec slope. The culprit is the current noise, increasing at higher frequencies with +20 dB/dec slope.

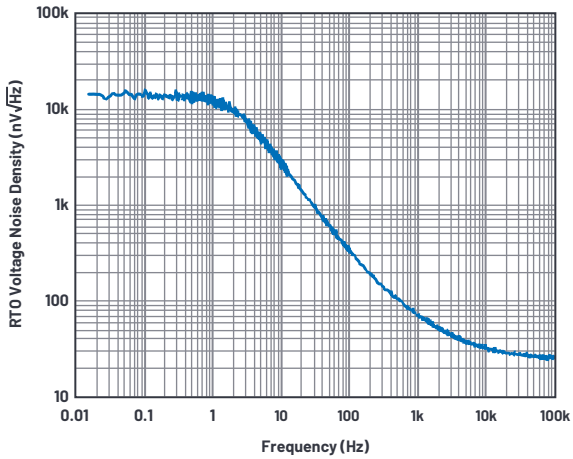

Figure 5. Output referred voltage noise density.

The SR785 dynamic signal analyzer or an FFT instrument can be used to measure the output voltage noise; however, a noise floor of less than 7 nV/√Hz is preferred. When the output voltage noise of the DUT roll-off is close to 20 nV/√Hz to 30 nV/√Hz, we want the analyzer noise floor to add as little noise as possible. A ratio of 3 times only adds about 5.5%. We can live with a 5% error in the noise domain (see Figure 4).

The Art Is in Back-Calculation

Measuring this way, the two main parameters necessary to plot current noise were obtained in just one measurement. First, we got the total input capacitance—that is, stray capacitance and input capacitance—which was necessary to back-calculate the roll-off. Even if there is stray capacitance, the information was captured. The input capacitance dominates over the 10 GΩ resistance. This total impedance converts the current noise into voltage noise. Therefore, knowing this total input capacitance is important. Second, it shows where the current noise starts to dominate—that is, where it starts deviating from the −20 dB/dec slope.

Let's look at an example with this data in Figure 5. The 3 dB roll-off point is read at 2.1 Hz, which corresponds to

$$C = \frac{1}{2\pi R_f} = 7.6 \text{ pF} \quad (5)$$

capacitance at the input. The data sheet mentions that the common-mode input capacitance is only about 2.1 pF, which means that there is about 5.5 pF of stray capacitance. Differential-mode input capacitance is bootstrapped by negative feedback, so it doesn't really come into play at low frequencies. With 7.6 pF capacitance, the impedance that the current noise sees is shown in Figure 6.

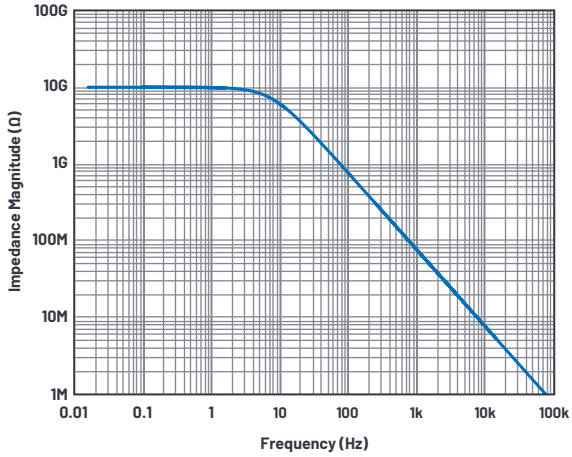

Figure 6. Total impedance magnitude of 10 GΩ resistor and 7.6 pF input capacitance in parallel.

Taking the referred to output (RTO) voltage noise measured on the AD8065 (Figure 5) and dividing by the impedance vs. frequency (Figure 6) gives us the equivalent current noise of AD8065 and the 10 GΩ resistor combined in RSS (Figure 7).

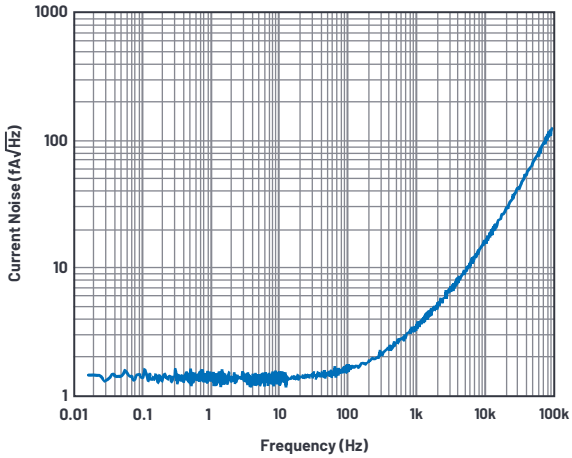

Figure 7. RTI current noise of the AD8065 and a 10 GΩ resistor.

After removing the current noise of 10 GΩ, the input referred noise of the AD8065 looks as shown in Figure 8. Below 10 Hz, it was very fuzzy because we were trying to fish out the 0.5 fA/√Hz to 0.6 fA/√Hz out of 1.28 fA/√Hz (10% on RSS scale) and only 100 averages were done. Between 15 mHz to 1.56 Hz, there are 400 lines with 4 mHz bandwidth. That's 256 seconds per average! 100 averages of 256 is 25,600 seconds, slightly more than 7 hours. Why is measurement down to 15 mHz required, and why spend that much time? Input capacitance of 10 pF with 10 GΩ creates a low-pass filter of 1.6 Hz. Low noise FET amplifiers have large input capacitances that can be up to 20 pF, which puts the 3 dB point at 0.8 Hz. To measure the 3 dB point correctly, we would need to see a decade before—that is, down to 0.08 Hz (or 80 mHz).

If we eyeball the fuzzy lines below 10 Hz, 0.6 fA/√Hz through

$$i_{n\_dut} = \sqrt{2qi_b} \quad (6)$$

can be verified. This equation is not entirely false for current noise. In the first-order approximation, it still shows the low frequency current noise behavior of the part because this current noise density value was obtained through dc input bias current. At high frequencies, however, current noise does not follow this equation.

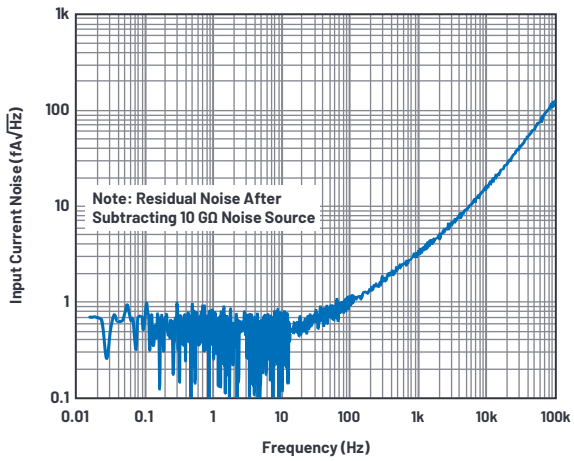

Figure 8. RTI current noise of AD8605.

At higher frequencies, the DUT current noise dominates the resistor current noise significantly, and the resistor noise can be ignored. Figure 9 shows the input referred current noise of various FET input amplifiers at 10 GΩ noise, measured with the setup shown in Figure 3. It seems that 100 fA/√Hz at 100 kHz is the typical performance that can be expected from most precision amplifiers.

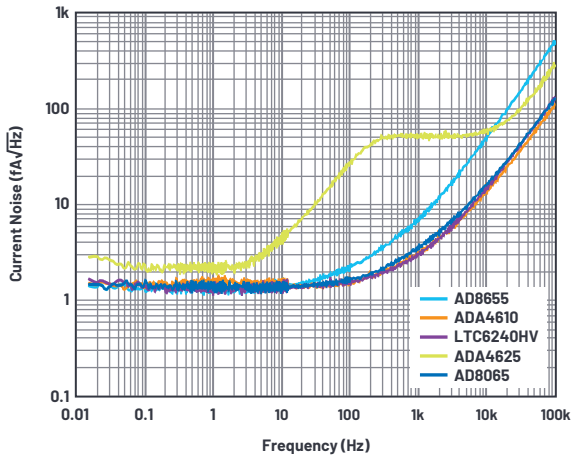

Figure 9. RTI current noise of selected ADI amplifiers.

There are exceptions: LTC6268/LTC6269 current noise is 5.6 fA/√Hz at 100 kHz. These parts are great for high speed TIA applications where high bandwidth, low input capacitance, femtoampere-level bias current are required.

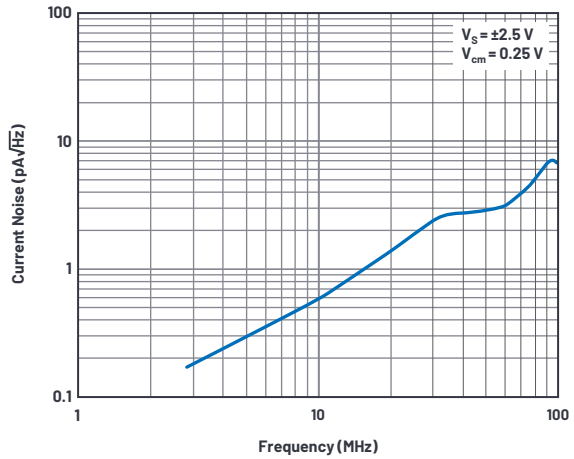

Figure 10. Input referred current noise of LTC6268.

Is This All There Is to Current Noise in FET Input Amps?

There are four major current noise sources that contribute to total input current noise in high source impedance applications and, so far, we have covered two. A simplified TIA amplifier with major noise sources is shown below in Figure 11. MT-050 is a good reference for the op amp noise sources.

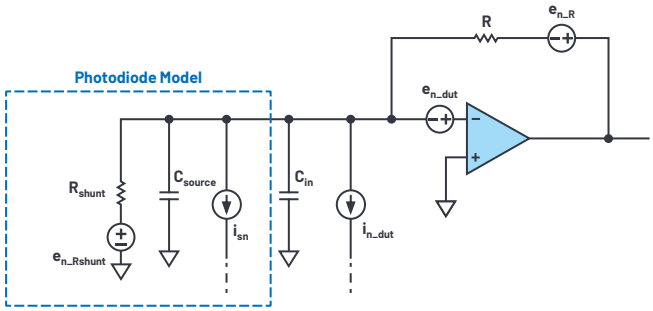

Figure 11. Simplified TIA amplifier with major noise sources.

Current Noise from the FET Input Amplifier (i_{n_dut})

The shape of the current noise depends on the amplifier input stage topology. Generally, the noise is flat in low frequencies, but gets larger as frequency gets higher. See Figure 8. Eventually, the noise will roll off at –20 dB/dec as the amplifier runs out of gain at higher frequencies.

Current Noise from the Resistor (i_{n_R})

This can be calculated from the thermal voltage noise of the resistor e_{n_R} divided by the impedance of the resistor, R. 1 MΩ contributes roughly 128 fA/√Hz and 10 GΩ contributes 1.28 fA/√Hz.

$$i_{n\_R} = \frac{e_{n\_R}}{R} = \frac{\sqrt{4kTR}}{R} = \sqrt{\frac{4kT}{R}} \quad (7)$$

The thermal voltage noise of the resistor is ideally flat over frequency, until it sees a capacitor and rolls off at –20 dB/dec. Figure 5 shows this behavior between 10 mHz to 1 Hz range.

Current Noise from the Sensor (i_{n_source})

The sensor itself contributes current noise, and we have to live with it. It can have any shape over frequency. For example: a photodiode exhibits shot noise, I_{sn}, from photocurrent, I_P, and dark current, I_D, as well as Johnson noise, I_{jn}, from the shunt resistance.[1]

$$i_{n\_source} = I_{sn} + I_{jn} = \sqrt{2q(I_P + I_D)} + \sqrt{\frac{4kT}{R_{shunt}}} \quad (8)$$

Current Noise from the Amplifier Voltage Noise Itself

The current noise from the amplifier voltage noise is coined as enC noise and is explained very well in *The Art of Electronics* by Horowitz and Hill.[2] Similar to resistor voltage noise being converted into current noise by the resistance, the amplifier voltage noise e_{n_dut} is converted into current noise by the total input capacitance, which includes the sensor capacitance, board stray capacitance, and the amplifier input capacitance

$$C_{in\_total} = C_{source} + C_{stray} + C_{in\_dut} \quad (9)$$

On first order, we get

$$i_{n\_enC} = \frac{e_{n\_dut}}{Z_{cin\_total}} = \frac{e_{n\_dut}}{1/\omega C} = \omega e_{n\_dut} C = 2\pi f\, e_{n\_dut} C \quad (10)$$

This equation tells us three things. First, the current noise gets larger with increasing frequency—yet another current noise component that gets larger with frequency. Second, the larger the input voltage noise of the amplifier, the larger the current noise. Third, the larger the total input capacitance, the larger the current noise. This results in the figure of merit e_nC where both the voltage noise of the amplifier and the total input capacitance should also be considered for a given application.

The shape of current noise for TIA applications, ignoring the DUT current noise, is shown in Figure 12. The flat portion is mainly the resistor noise

$$i_{n\_R} = \sqrt{\frac{4kT}{R}} \tag{11}$$

and the capacitor induced current noise is

$$i_{n\_enC} = 2\pi f e_{n\_dut} C \tag{12}$$

increasing with 20 dB/dec. From the two equations, the crossover point can be calculated as

$$f_x = \frac{\sqrt{4kT}}{2\pi} \times \frac{1}{e_{n\_dut} C \sqrt{R}} \tag{13}$$

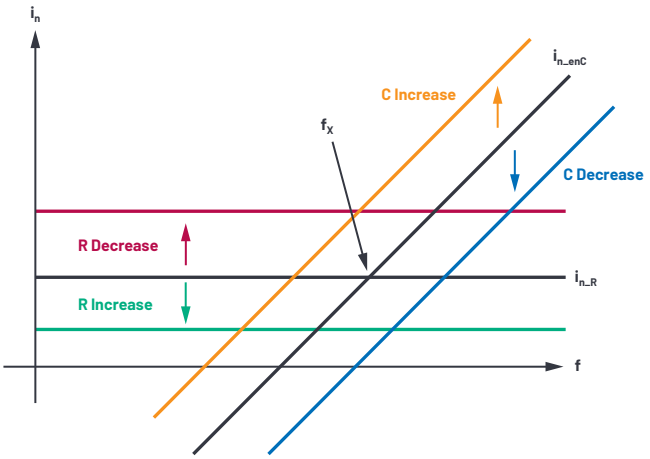

Figure 12. e_nC noise over frequency.

Depending on the C_{in}, e_nC noise can be larger or smaller than the DUT current noise. For inverting configuration such as TIA applications, C_{dm} is not boot-strapped; that is,

$$C_{in\_dut} = C_{cm} + C_{dm} \tag{14}$$

For instance, at 100 kHz, the LTC6244 with C_{cm} = 2.1 pF, C_{dm} = 3.5 pF, and e_n = 8 nV/√Hz will have e_nC current noise of

$$i_{n\_enC} = 2\pi \times 100\ \text{kHz} \times 8\frac{\text{nV}}{\sqrt{\text{Hz}}} \times (2.1 + 3.5)\ \text{pf} = 28\frac{\text{fA}}{\sqrt{\text{Hz}}} \tag{15}$$

This is much less than the DUT current noise of 80 fA/√Hz

However, when a photodiode is connected, an extra C_{source} or C_{pd} is added to the equation and the current noise can be recalculated. It takes only 16 pF of extra capacitance from C_{pd} to be equal to the DUT current noise. Low speed, large area photodiodes tend to be in the order of 100 pF to 1 nF, while high speed, small area photodiodes can be 1 pF to 10 pF.

Summary

The phenomenon of current noise increasing with frequency, in both CMOS and JFET input amplifiers, is well known to IC design engineers and seasoned circuit designers, but it was elusive to many engineers as a result of either too few articles in the field or incomplete information from manufacturers. The goal of this article is to bridge the understanding of the current noise behavior toward a higher frequency domain and to show a technique to reproduce the measurement on the op amp of choice.

Further Reading

Choosing op amps to get the best performance is not a simple task. Based on the applications, trade-offs are performed between noise, bandwidth, gain, and accuracy. References 1, 2, 3, 4, 5, 6, and 7, along with many amplifier data sheets, detail how these trade-offs can be made.

References

[1] Photodiode Characteristics and Applications. OSI Optoelectronics, August 2007.

[2] Paul Horowitz and Winfield Hill. *The Art of Electronics*, 3rd edition. Cambridge University Press, April 2015.

[3] ADA4530-1 Data Sheet. Analog Devices, Inc., November 2019.

[4] CN-0407. Analog Devices, Inc., February 2019.

[5] "ADA4530-1R-EBZ User Guide: UG-865." Analog Devices, Inc., October 2015.

[6] "MT-050: Op Amp Total Output Noise Calculations for Second-Order System." Analog Devices, Inc., February 2009.

[7] *Low Level Measurements Handbook: Precision DC Current, Voltage, and Resistance Measurements*. Tektronix, Inc., February 2016.

Brisebois, Glen. "Signal Conditioning for High Impedance Sensors." Analog Devices, Inc..

Brisebois, Glen. "Transimpedance Amplifier Noise Considerations." Analog Devices, Inc..

Acknowledgements

Kaung would like to thank Glen Brisebois and Aaron Schultz for their support, and Henry Surtihadi, Scott Hunt, Barry Harvey, Harry Holt, Philip Karantzalis, and Jordyn Ansari for their input.

Appendix

Measuring noise in a high impedance environment, 10 GΩ impedance with FET input, doesn't come without fighting with the environment and its subtleties.

In a typical single amplifier pin layout, Pin3 (Vin+) is next to Pin4 (V–). Layout of the board matters significantly when there is no guard ring in place. There was significant dc shift at the output as the supplies were swept. The 10 GΩ SMD was originally soldered in parallel with the V– (R10 in Figure 13) and the leakage from solder paste was unbearable. As a result, the 10 GΩ SMD was moved to another location (R8) and the leakage disappeared. The data sheet of ADA4530-1 (electrometer-grade amplifier with 20 fA at 85°C) shows all the precautions necessary regarding solder paste selection, contamination, humidity effects, and other juicy details regarding high impedance measurements. The data sheet and user guide UG-865, as well as circuit note CN-0407, are worth studying.

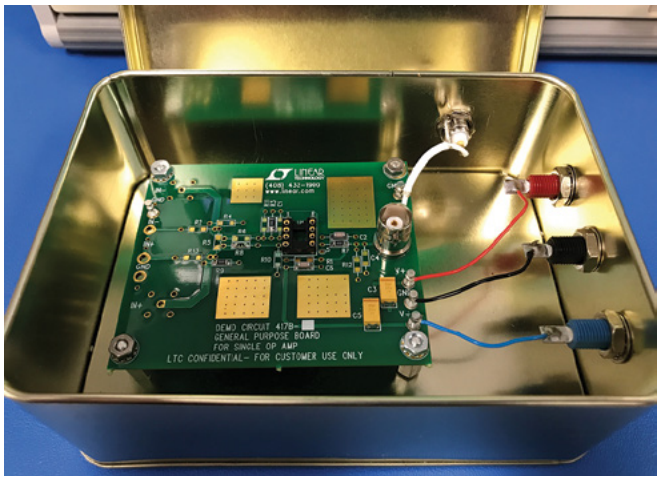

Figure 13. Measurement setup.

Devices that are high impedance and not soundproof are prone to triboelectric, piezoelectric, or microphonic effects. One day, I accidentally dropped my keys and saw the noise spiking up at audible frequencies (1 kHz and beyond in particular). I did not think measurement with 10 GΩ at the high impedance FET input amp would be this sensitive to sound. I whistled just to double check. And there it was, a spike between 1 kHz and 2 kHz. Even with a significant amount of averages, one sharp whistle would bring up a noise spike on the CRT screen of the SR785. The hermetically sealed glass resistors mentioned in CN-0407 would be a better choice for piezoelectric/triboelectric effects.

To confirm, I measured the lab environment sound with a laptop microphone, processed the data with MATLAB®, and found that the noise correlates well with the measurement. A significant noise spike was seen at 768 Hz and other frequencies as seen in Figure 14. The culprit was the large ac duct running a few meters away from my bench. To make sure I was not picking up the noise of the laptop itself, I went into one of the phone booths, the acoustically quietest place, and took data. There was no 768 Hz measurement. Noise spikes at other frequencies were at least 100 times lower.

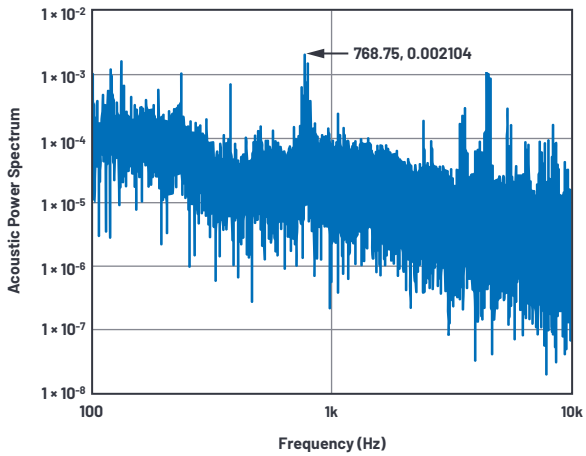

Figure 14. Lab acoustic noise.

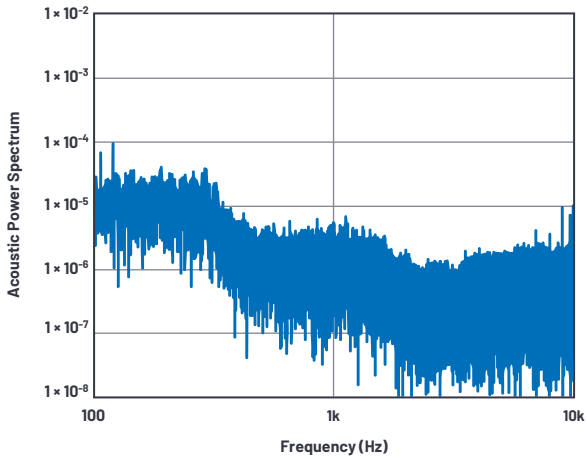

Figure 15. Phone booth acoustic noise.

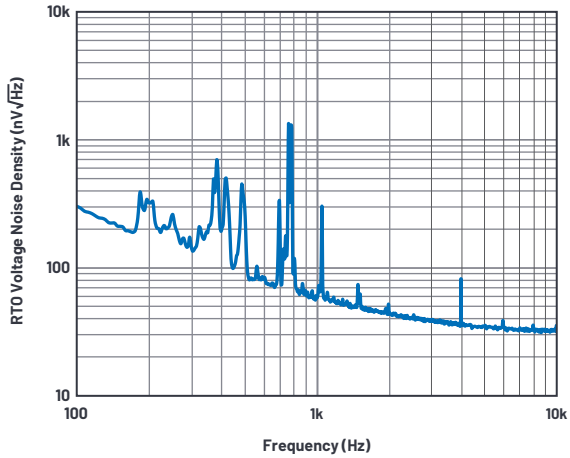

Figure 16. Output referred voltage noise density without acoustic shielding.

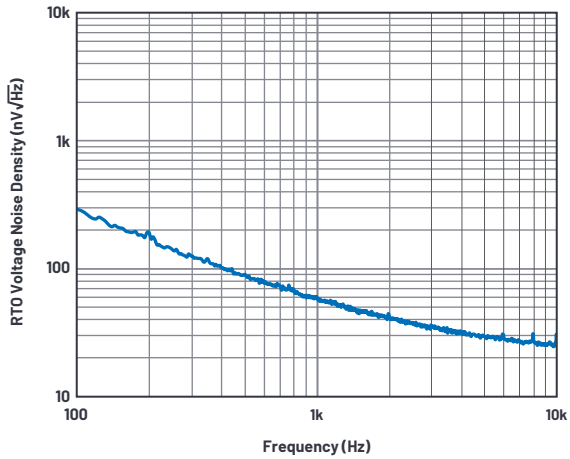

Figure 17. Output referred voltage noise density with acoustic shielding.

To attenuate the audible noise, a Temptronix box was used. The box seems to be thermally isolated, meaning no significant air flow. All I needed was for it to shield the acoustics enough so that the microphonic effects would not show up in the measurements. And it did the job. See Figure 16 and Figure 17.

Instrument Specific Issue:

FET input amps have input bias currents in the order of pA. 10 pA going into 10 GΩ still reads about 100 mV of offset at the output of the amplifier. The SR785 has an ac coupling feature that works well to remove this offset and measure the output noise with the best range of –50 dB V peak (3.2 mV peak). However, the ac coupling features cuts into the frequency of interest less than 1 Hz, which makes it hard to determine the flat 12.8 μV/$\sqrt{\text{Hz}}$ and read 3 dB off of the point. DC coupling must be used, but now the most sensitive range of the instrument cannot be used. A 1 mHz passive filter, made with two 270 μF polarized caps in series (135 μF cap) and 1 MΩ resistor, was put in between the output of the DUT and the SR785. Due to the long leads in capacitor—that is, more loops—it tends to pick up the magnetic field produced by the SR785 CRT screen at 20 kHz and its harmonics. Because magnetic fields are three-dimensional in nature, angling and rotating the passive filter box solved the issue. Notice the angled blue box in Figure 18. E&M black magic!

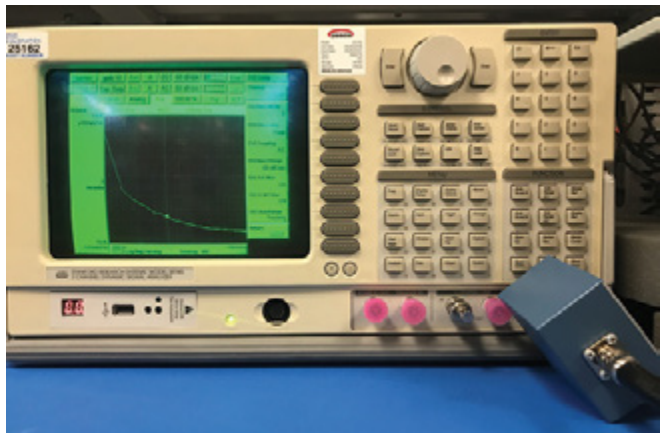

Figure 18. Filter box rotated to be less sensitive to the magnetic fields.

About the Author

Kaung Win joined ADI in 2013 as a product evaluation engineer for the Linear Products and Solutions Group and transferred to an application engineer position in 2019. He has a bachelor's degree in electrical and computer engineering from Worcester Polytechnic Institute and a master's degree in electrical engineering from Santa Clara University. Kaung specializes in amplifier signal chain solutions. He can be reached at kaung.win@analog.com.

Why Does Voltage Reference Noise Matter?

Anshul Shah, Applications Engineer

There is increasing demand for ultrahigh precision measurements that can achieve greater than 24-bit resolution in industries from aerospace and defense and gas exploration to pharmaceutical and medical device manufacturers. For example, the pharmaceutical industry uses high precision lab balances that offer 0.0001 mg resolution over a 2.1 g full-scale range that would require an analog-to-digital converter (ADC) with greater than 24-bit resolution. Calibration and testing of these high precision system challenges the instrumentation industry to offer test equipment that can achieve greater than 25-bit resolution with at least 7.5 digit measurement precision.

To achieve this high resolution, a signal chain with exceptionally low noise is required. Figure 1 shows the relationship of noise vs. effective number of bits (ENOB) and signal-to-noise ratio (SNR). Note, noise is calculated based on voltage reference (V_{REF}) equal to 5 V and ADC input set to full-scale range. To provide perspective, to achieve 25-bit resolution, or 152 dB dynamic range, the maximum allowable system noise is 0.2437 µV rms.

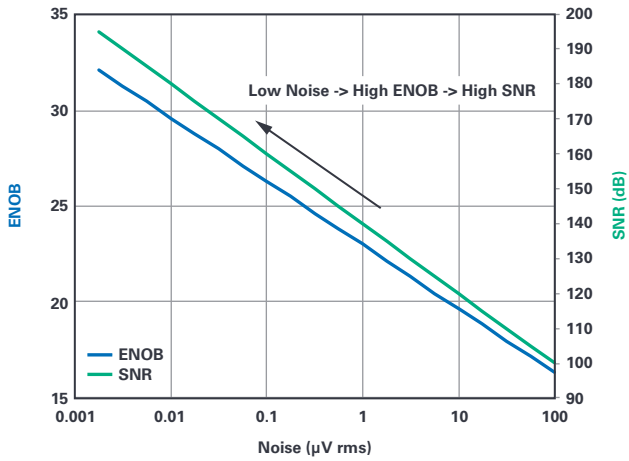

Figure 1. Noise vs. ENOB and SNR.

The voltage reference sets the limit to the input analog signal that the ADC can resolve. Equation 1 is the ideal transfer function of an ADC where the output code—in decimal form—is computed by the analog input signal V_{IN}, voltage reference V_{REF}, and number of ADC bits N.

$$ADC\ Code = V_{IN} \times \frac{2^N}{V_{REF}} \quad (1)$$

Typically, the resolution stated in the ADC data sheet is based on an input shorted technique where the ADC input is connected to the GND or the ADC differential inputs are connected to a common source. The ADC input shorted technique helps to characterize the absolute limit of the ADC resolution by omitting the ADC input source noise and eliminating the effect of V_{REF} noise. This is true because V_{IN} is set to 0 V, resulting in the ratio V_{IN}/V_{REF} being equal to 0 V.

To investigate the effects of voltage reference noise on overall system noise, Figure 2 shows the relationship between total system noise (rms) with ADC input dc source voltage. For this test, we used the AD7177-2 32-bit ADC with the V_{REF} input connected to the LTC6655-5 (5 V) and the ADC input connected to a low noise dc source. The ADC output data rate was set to 10 kSPS. Note, throughout the ADC input voltage range, the ADC noise remains constant (35 nV/$\sqrt{Hz}$) while the ADC dc input source noise rises (≤6 nV/$\sqrt{Hz}$) but remains low in comparison to the voltage reference noise (96 nV/$\sqrt{Hz}$). As shown in Figure 2, the total noise is proportional to the ADC dc input voltage. This is because as V_{IN} increases, the ratio V_{IN}/V_{REF} increases and so the V_{REF} noise dominates the overall system noise when the ADC is at full-scale input. The individual noise of each component in the signal chain adds together in root sum square (RSS) fashion and gives rise to the shape of the curve in Figure 2.

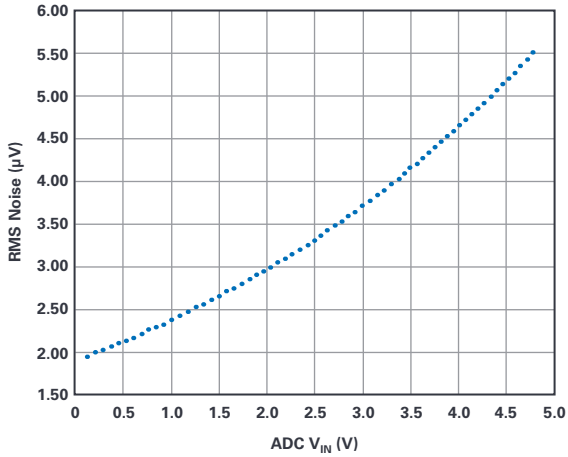

Figure 2. Relationship of ADC V_{IN} with rms system noise. V_{REF} set to LTC6655-5.

To achieve a high measurement resolution of 25 bits or beyond, even the best standalone voltage reference available in the market with a low noise specification needs some help to attenuate its noise. Adding external circuitry such as a filter can help attenuate noise to achieve the desired ADC dynamic range.

The remainder of this article explains various types of low-pass filters and how they can be applied to attenuate voltage reference noise. Filter design techniques and filter trade-offs will be discussed. Two types of low-pass filters that will be discussed in the context of attenuated voltage reference noise are simple passive RC low-pass filters (LPFs) and active-based signal flow graph (SFG) low-pass filters. System evaluation results using a sigma-delta (Σ-Δ) ADC will be presented in the circuit performance section.

Noise Reduction Using a Passive Low-Pass Filter

Figure 3 shows the voltage reference driving an ADC via a low-pass filter implemented with an external reservoir capacitor, C1, the equivalent series resistance (ESR) of the reservoir capacitor and the output impedance of voltage reference operational amplifier (op amp). The passive RC LPF cutoff frequency is determined by

$$f_C = \frac{1}{2\pi RC} \quad (2)$$

which states that bandwidth is inversely proportional to resistance R and capacitance C.

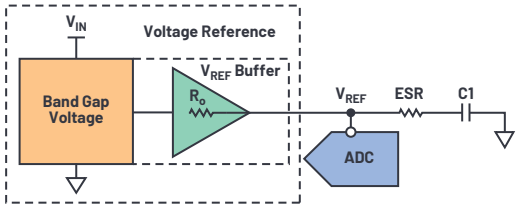

Figure 3. Low-pass filter between series voltage reference and ADC.

Reservoir capacitor C1 also works as local energy storage to compensate for voltage spikes caused when ADC voltage reference circuitry demands sudden change in load current. Figure 4 shows the sigma-delta AD7177-2 and SAR AD7980 ADC dynamic reference current response.

The user can choose the value of the C1 capacitor to meet the LPF cutoff frequency requirement, but some SAR ADCs require 10 µF minimum capacitor on the reference input in order to operate correctly. This minimum 10 µF C1 capacitor reduces the phase margin of the reference buffer. As the phase margin reduces, the buffer feedback is no longer negative.[1] The signals near the unity-gain crossover frequency are fed back in-phase with the incoming signals.[1] This causes the closed-loop response to introduce noise peak near the crossover frequency.[1] Since the bandwidth from the cutoff frequency (–3 dB point) reaches up to 16 MHz, the total integrated noise (rms) is dominated by the noise peak. Even though voltage reference reservoir capacitor C1 operates as a noise filter and compensates for voltage spikes, the caveat is the noise peak. Figure 5 shows the noise peak of the LTC6655 voltage reference introduced by reservoir capacitor C1. The noise peak magnitude is determined by the value of the reservoir capacitor and its ESR rating.

Most voltage references are designed with a complex output stage to drive a large load capacitance suitable for ADC reference circuitry. For example, the LTC6655 output stage is designed to be critically damped with a reservoir capacitance set to 10 µF. When the LTC6655's reservoir capacitance is set to a minimum of 2.7 µF and a maximum of 100 µF, noise peaking is introduced.

The equivalent series resistance of the V_{REF} output reservoir capacitance does mitigate the primary noise peak but introduces a secondary noise peak at 100 kHz and above. This can be explained by the fact that the ESR of the cap introduces a zero, which leads to improving phase margin and reducing primary noise peak. However, this zero combines with the inherent zero of the LTC6655 and creates secondary noise peaking. Note, the noise response in Figure 5 is only valid for the LTC6655 voltage reference.

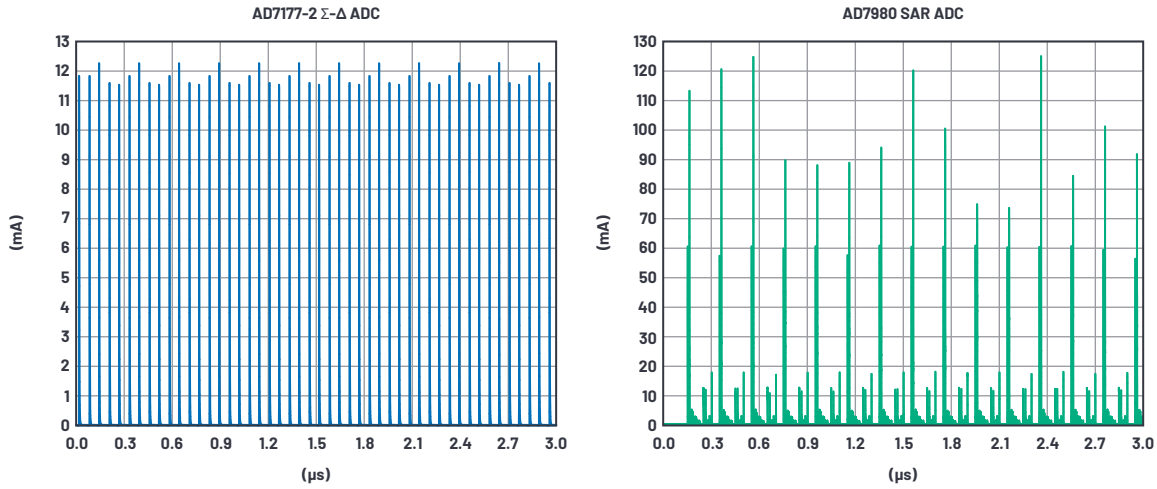

Figure 4. AD7177-2 and AD7980 simulated dynamic reference current response.

One of the other solutions to filter voltage reference noise, remove the noise peak, and properly drive the ADC is to add a passive RC LPF followed by a buffer. By adding a buffer, we separate the design constraints of the LPF and the ADC reference input capacitor. See Figure 6.

Setting the passive RC LPF cutoff frequency well below the unity-gain crossover frequency will not only reduce broadband and low frequency noise but also avoid noise peaking. For example, Figure 7 shows the LTC6655 noise response with C1 = 100 µF (ESR = 0 Ω), followed by a passive LPF where R = 10 kΩ and C2 = 10 µF (ESR = 0 Ω), creating a pole at 1.59 Hz.

Increasing the low-pass filter resistor R can help achieve a low cutoff frequency, but can also result in dc accuracy degradation of precision voltage reference. When adding a passive RC LPF, the user must also consider the impact on the load regulation and the impact on the V_{REF} buffer response (τ = RC), which affects its transient performance when driving an ADC.

To achieve the required transient performance, it is suggested to use a buffer as shown in Figure 6. Critical specifications to consider in terms of selecting buffer includes ultralow noise, capability to support high load capacitance, low distortion, excellent slew rate, and wide gain bandwidth. Recommendations for reference buffers are the ADA4805-1 and ADA4807-1.

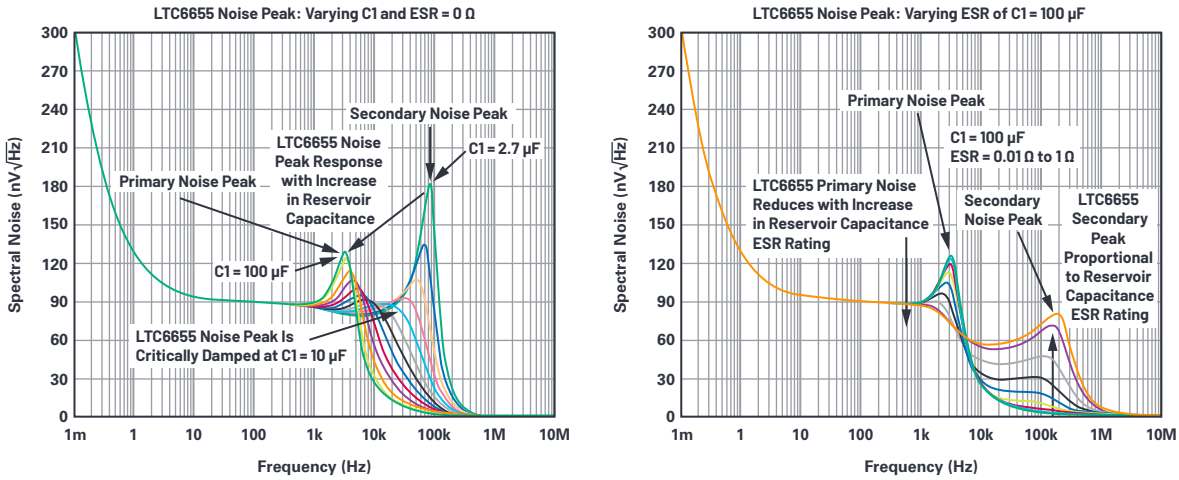

Figure 5. LTC6655 voltage reference noise peaking density.

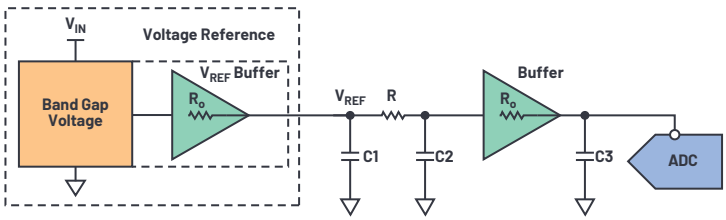

Figure 6. Passive RC LPF followed by a buffer.

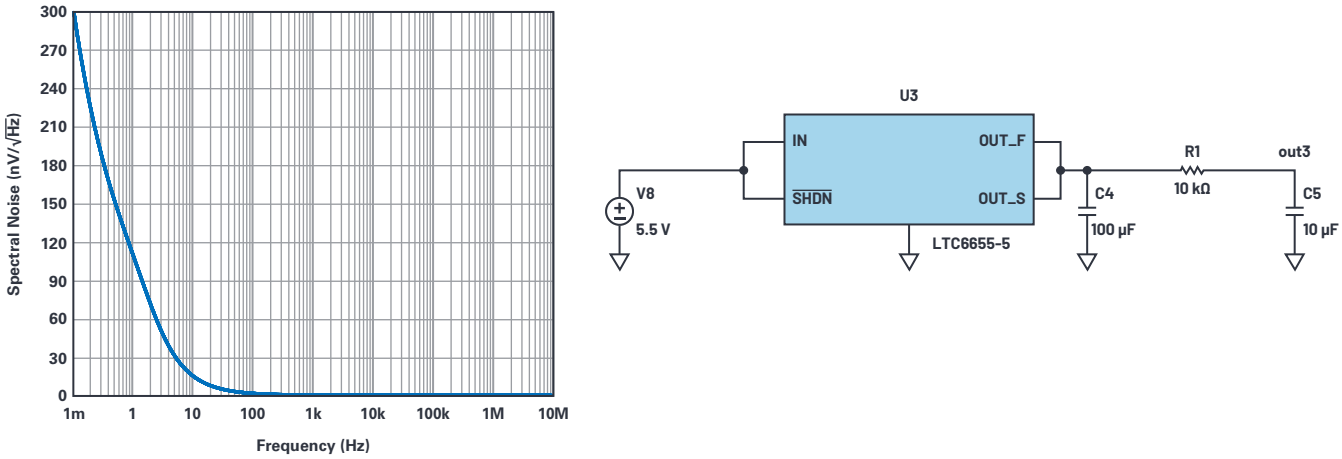

Figure 7. LTC6655-5 followed by passive RC LPF noise response.

Noise Reduction Using an Active LPF

Table 1 states the required dynamic range and maximum allowable system noise that must be met in order to achieve the desired ENOB ADC resolution. Depending on the ADC bandwidth, a single-pole, low-pass filter attenuating at 20 dB/decade may not achieve the desired wideband noise reduction. Cascading passive low-pass filters creates a ladder structure that can generate a higher order filter, but each section's input impedance will be a load on the previous section. This can degrade the dc accuracy of the precision voltage reference. However, designing a higher order LPF based on active components will provide excellent isolation between input to output, minimizing voltage reference dc accuracy degradation, and provide low output impedance to drive the reference circuitry of the ADC.

$$SNR = 6.02N + 1.76 \text{ dB} \quad (3)$$

$$LSB = \frac{V_{REF}}{2^N} \quad (4)$$

Table 1. Condition: V_{REF} = 5 V and ADC Input Set to Full-Scale Range

ENOB	SNR (dB)	Noise (µV rms)
20	122.16	7.798301
21	128.18	3.89942
22	134.2	1.949845
23	140.22	0.97499
24	146.24	0.487528
25	152.26	0.243781
26	158.28	0.121899
27	164.3	0.060954
28	170.32	0.030479
29	176.34	0.015241
30	182.36	0.007621
31	188.38	0.003811
32	194.4	0.001905

There are several different types of active low-pass filters—for example, Bessel, Butterworth, Chebyshev, and elliptic—as shown in Figure 8. Having a band-pass that is flat or does not exhibit ripple will keep the precision voltage reference's dc accuracy degradation to a minimum. Out of all filter types, designing LPF based on the Butterworth topology can achieve flat band-pass and steep attenuation.

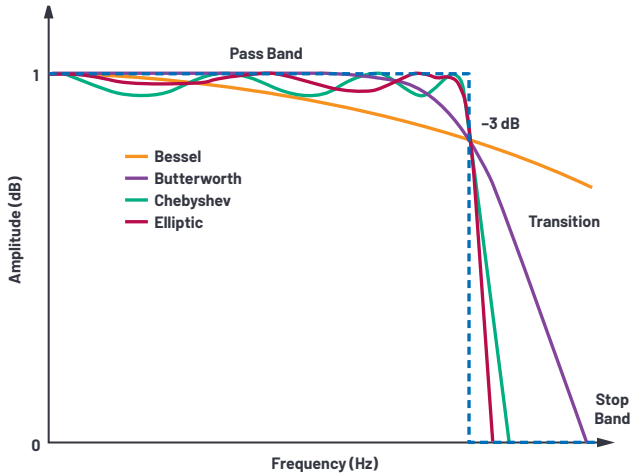

Figure 8. Filter amplitude response examples.

Active Low-Pass Filter Design Technique

A signal flow graph is a graphical representation of a system derived from a set of linear equations.[2] An SFG provides a bridge from a transfer function to a corresponding circuit topology of a system.[2] This theory can be applied to designing analog filters based on active circuitry. The key advantage of an SFG filter design approach is that the damping factor, Q, and cutoff frequency can be individually controlled. An SFG LPF can help to attenuate noise and improve SNR, but comes at the cost of additional bill of material (BOM) expenses, PCB area, and power. Furthermore, an SFG LPF can affect the reference output voltage with temperature leading to a small PPM error and hence dc accuracy degradation. Figure 9 shows an example of second-order low-pass filter transitioning from transfer function to circuit blocks via the SFG method. The scaling resistor (R) and capacitor (C) configures for the cutoff frequency (please see Equation 5).

For more details about signal flow graph theory, please refer to *Feedback Control of Dynamic Systems*, published by Addison-Wesley.[2]

$$R = Rs \times Rn \qquad C = \frac{Cn}{Ws \times R} \quad (5)$$

where

Rs is scaling factor

Cn is scaling factor

Ws is cutoff frequency (Rad/s)

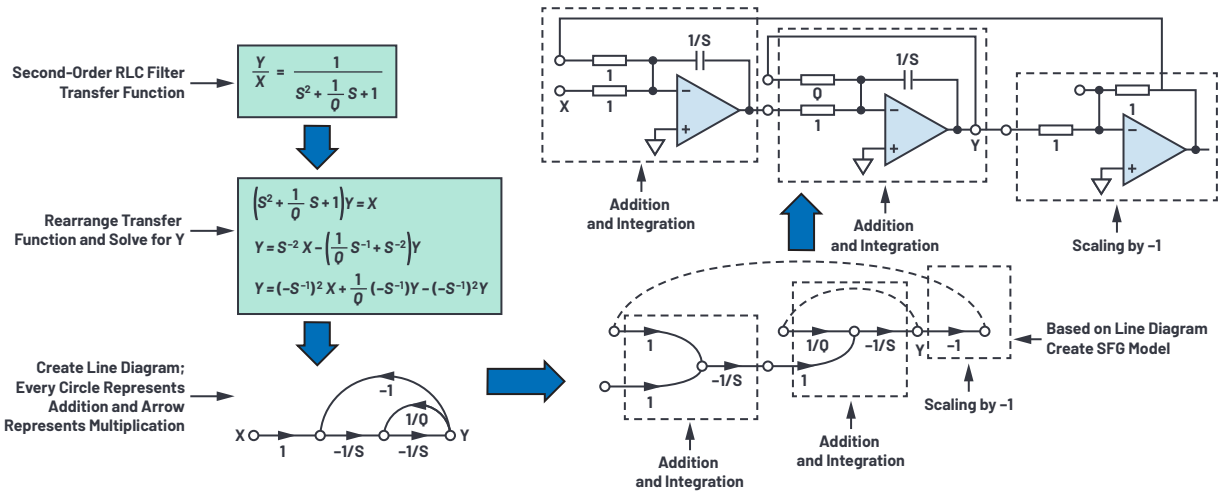

Figure 9. Active RC low-pass filter implementation based on SFG method.

The following is a calculation example for designing a second-order, 0.5 Hz cutoff frequency SFG low-pass Butterworth filter:

- For simplicity of this example, select Rs = 1 Ω, Cn = 1 F.
- Select Fs = 0.5 Hz to maximize rejection of broadband noise. Ws = 2 × π × 0.5 = 3.141 rad.
- Set the damping factor Q = 0.71. Select this value to achieve a flat band-pass and steep attenuation to reflect Butterworth topology.
- R, C, and Rq values were chosen based on an iterative process to achieve low thermal noise and the availability of component values for surface mount.

$$R = 7.32\ \text{k}\Omega$$

$$C = \frac{1}{2 \times \pi \times 0.5\ \text{Hz} \times 7.32\ \text{k}\Omega} = 44\ \mu\text{F} \quad (6)$$

$$Rq = R \times Q = 7.32\ \text{k}\Omega \times 0.71 = 5.2\ \text{k}\Omega$$

Introducing LTC6655LN

Considering the RC LPF and SFG LPF trade-offs, a better solution is to have a low-pass filter placed before the integrated low noise buffer of the voltage reference as shown in Figure 10. This implementation will not only reduce the PCB area but also not hinder the voltage reference buffer response. Using a voltage reference buffer with fast settling, high input impedance, and the capability to sink and source current will help overcome poor load regulation, maintain dc accuracy, and improve transient performance. The LTC6655LN takes advantage of this architecture. It comes with a noise reduction pin that enables reduction of wideband noise and an integrated output stage buffer. LTC6655LN is internally equipped with R3 resistor (see Figure 10) and allows users to connect external capacitor at the noise reduction (NR) pin to create a low-pass filter. With LTC6655LN architecture, users can configure the low-pass cutoff frequency based on their system requirements.

Table 2. The 3 dB Cutoff Frequencies for Different Values of the Capacitor Connected to the NR Pin

CNR	2.500	4.096	5.000	V
0.1 µF	5305	4233	3969	Hz
1 µF	531	423	397	Hz
10 µF	53	42.3	39.7	Hz
100 µF	5.3	4.2	4.0	Hz

The LTC6655LN RC LPF is connected to the noninverting node of the buffer, which is the most sensitive pin on this device. Precaution must be taken when selecting a low leakage type for the external capacitor to prevent leakage current flow through the R3 resistor, which can degrade dc accuracy. Furthermore, the variation of R and C do not track each other and therefore the RC time constant and LPF cutoff frequency can change due to process, voltage, and temperature (PVT) variation.

Table 3. Resistance Value of R3 for the Three Voltage Options

Voltage Option	2.500 V	4.096 V	5.000 V
R3 ± 15%	300 Ω	376 Ω	401 Ω

A voltage reference such as the LTC6655LN with an internally built-in LPF provides the best solution in simplifying noise filter design and eliminating the need for external buffer to drive ADC voltage reference circuitry.

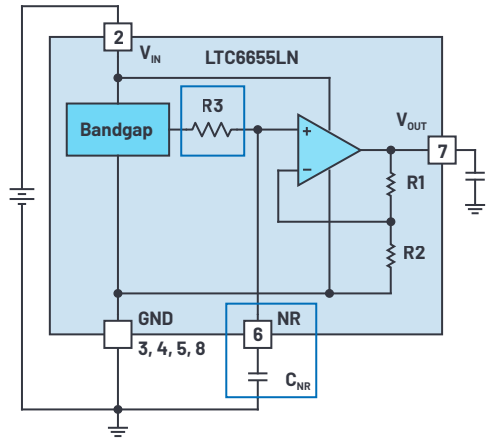

Figure 10. LTC6655LN block diagram.

Test Circuit Description

The AD7177-2 precision ADC was used to benchmark the performance of LTC6655/LTC6655LN with a 10 µF NR capacitor and LTC6655 followed by an active SFG filter. The AD7177-2 is a high resolution, 32-bit, low noise, fast settling, 2-channel/4-channel, sigma-delta, analog-to-digital converter for low bandwidth inputs. AD7177-2 is integrated with a programmable digital low-pass filter that allows users to control the output data rate (ODR) from 5 SPS to 10 kSPS.

The components used in designing SFG LPF (Figure 11) were two ADA4522-1 op amps, an AD797 op amp, 25 ppm surface-mount resistors, multilayer surface-mount ceramic capacitors, and a 10 µF WIMA film capacitor. ADA4522 is a rail-to-rail output op amp with a broadband noise density of 5.8 nV/√Hz and 177 nV p-p flicker noise. AD797 is a low noise op amp with 0.9 nV/√Hz broadband noise, 50 nV p-p flicker noise, excellent slew rate of 20 V/µs, and gain bandwidth of 100 MHz, which makes it suitable for driving an ADC.

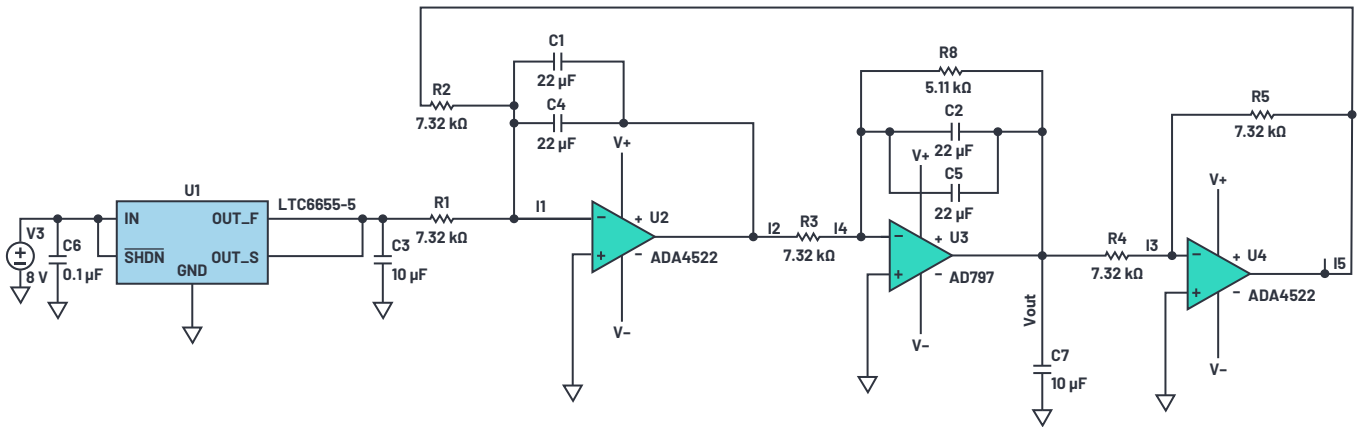

Figure 11. SFG LPF.

In order to correctly evaluate the performance when using an LTC6655 and an LTC6655LN with an AD7177-2, a dc source with overall noise lower than the ADC voltage reference and the ADC noise is required. Therefore, an ideal source was used, namely a 9 V battery supply as can be seen in Figure 12.

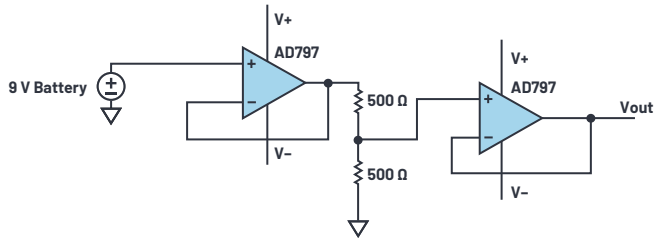

Figure 12. Low noise dc source.

Circuit Performance

Figure 13 displays spectral noise density and Figure 14 displays output data rate (ODR) vs. ENOB, depicting the performance of AD7177-2 with its V_{REF} input connected to a LTC6655/LTC6655LN with 10 µF NR capacitor or a filtered LTC6655 (SFG). To provide perspective of spectral noise density comparison at 1 kHz, see Table 4. Both Figure 13 and Figure 14 have two important regions.

Table 4. Spectral Noise Density Comparison at 1 kHz

	LTC6655	LTC6655LN with 10 µF NR Capacitance	LTC6655 SFG Filter	ADC Input DC Source
Spectral Noise Density at 1 kHz (nV/√Hz)	96	32	2.4	6.7

Region A:

Spectral noise density plot Figure 13 shows that at ODR of 500 SPS and higher, both the filtered LTC6655 (SFG) and ADC dc input source noise are significantly lower noise than the ADC. This results in the least amount of deviation from the maximum performance achievable by the ADC as shown in region A in Figure 14. The key takeaway based on ODR vs. ENOB and spectral noise density plot is that, within region A, the rise of the total integrated noise (rms) prevents the signal chain from achieving 25-bit measurement resolution.

Region B:

In this region, the spectral noise density plot (Figure 13) shows that the flicker noise of the three voltage reference options and the dc source increase and overall system noise are dominated by the dc source noise. This increase in flicker noise within region B explains the rise in deviation of ENOB between the measured performance and maximum achievable by the ADC (Figure 14).

According to ODR vs. ENOB plot, filtered LTC6655 (SFG) achieves 25-bit resolution at 20 SPS and lower while the LTC6655LN-5 with 10 µF NR cap and the LTC6655 cannot achieve better than 24.6-bit resolution.

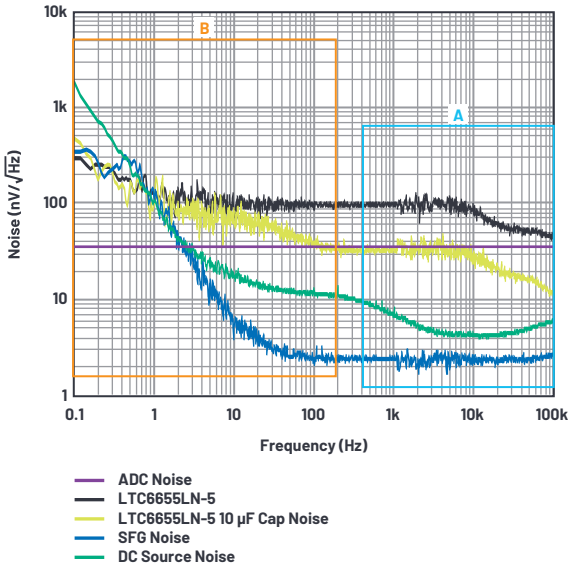

Figure 13. Spectral noise density.

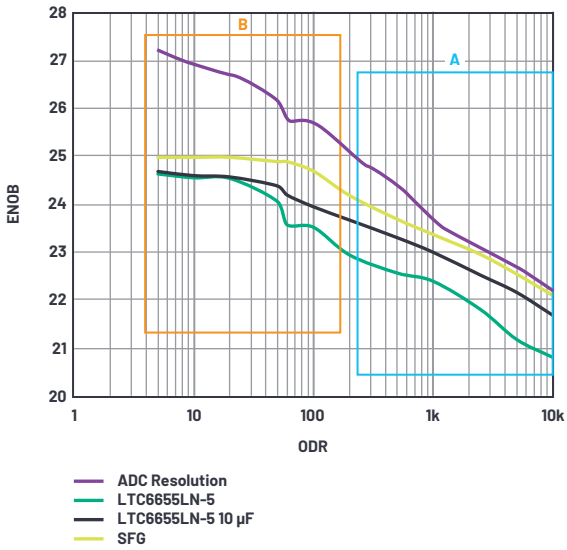

Figure 14. ODR vs. ENOB.

Table 5 below is a summary of the AD7177-2 ADC performance with the V_{REF} input either connected to a LTC6655/LTC6655LN with 10 µF NR capacitance or filtered LTC6655 (SFG). With ADC inputs tied together and the V_{REF} input connected to LTC6655, the zero-scale column establishes the best dynamic range AD7177-2 can achieve. With the ADC inputs nearly set to full-scale range, LTC6655LN-5 with 10 µF NR cap increases on average 4 dB dynamic range for up to 59.96 SPS compared to LTC6655. On the other hand, a filtered LTC6655 (SFG) achieves on average a 7 dB increment in dynamic range compared to LTC6655 for up to 59.96 SPS. The dynamic range delta does not vary much below 59.96 SPS and the variance is mainly due to the dominated low frequency flicker noise induced by the ADC input dc source.

Compared to LTC6655/LTC6655LN with 10 µF connected to NR pin reduces the broadband noise at 1 kHz by 62% and filtered LTC6655 (SFG) reduces broadband noise by 97%.

Conclusion

A precision system that is attempting to achieve a 25-bit resolution or higher must account for the significance of voltage reference noise. As shown in Figure 2, the contribution of V_{REF} noise to system noise is proportional to the utilization of the ADC's full-scale range. This article shows that adding a filter to a precision voltage reference attenuates V_{REF} noise, which leads to reducing overall system noise. An LTC6655 voltage reference followed by an SFG filter can reduce broadband noise by 97% of the LTC6655 with no filter. This comes at a cost of additional BOM, more PCB area, more power consumption, a few PPM of dc accuracy degradation, and can vary precision reference output with temperature. Considering SFG LPF trade-offs, LTC6655LN has leverage in terms of simple design, low power, only requires a single capacitor to reduce broadband noise, and eliminates the need for an external buffer to drive an ADC. LTC6655LN with 10 µF NR capacitor does reduce broadband noise by 62% of the LTC6655 with no filter. Hence, users can now take advantage of the built-in LTC6655LN low-pass filter to enable precision systems to achieve their desired resolution.

Addendum

To download LTspice, please visit *analog.com/ltspice*.

Click here to download the LTspice simulation for the Figure 7 circuit, the SFG LPF circuit in Figure 11, and the low noise dc source circuit shown in Figure 12.

References

[1] Mark Reisiger. "Reduce Amplifier Noise Peaking to Improve SNR." *ElectronicDesign*, October 2012.

[2] Gene F. Franklin, J. David Powell, and Abbas Emami-Naeini. Feedback Control of Dynamics Systems. Addison-Wesley Longman Publishing Co., Inc., November 1993.

Acknowledgements

I would like to thank author Robert Kiely for his previous work on sigma-delta ADCs, precision amplifiers, and voltage references.

Table 5. Dynamic Range Comparison

ODR	ADC Dynamic Range Zero Scale (dB)	LTC6655 Dynamic Range (dB)	LTC6655LN 10 µF Dynamic Range (dB)	LTC6655 (SFG) Dynamic Range (dB)	Dynamic Range Delta (LTC6655LN 10 µF—LTC6655) (dB)	Dynamic Range Delta (LTC6655 (SFG)—LTC6655) (dB)
10000	135.40	126.88	132.22	134.65	5.33	7.77
5000	138.41	129.14	135.08	137.37	5.94	8.23
2500	140.82	132.91	137.23	139.86	4.32	6.95
1000	144.43	136.50	140.11	142.42	3.61	5.92
500	148.65	137.55	141.95	144.37	4.40	6.83
200	152.86	139.83	144.15	147.40	4.32	7.57
100	156.47	143.32	145.82	150.49	2.49	7.17
59.96	157.08	143.66	147.31	151.71	3.65	8.05
49.96	159.48	146.58	148.43	151.72	1.85	5.14
20	162.49	149.51	149.56	152.26	0.06	2.76
10	163.70	149.58	149.72	152.26	0.14	2.68
5	165.50	150.07	150.25	152.26	0.18	2.19

About the Author

Anshul Shah graduated from Arizona State University with an M.S. degree in electrical engineering. He is currently an applications engineer in ADI's Instrumentation and Precision Technology Group in Santa Clara, California, with a focus on precision voltage reference. Anshul joined Analog Devices in 2018. Prior to joining Analog Devices, Anshul held various positions in product test and validation at NXP Semiconductor. He can be reached at *anshul.shah@analog.com*.

Contactless Fluid-Level Measurement Using a Reflectometer Chip

Bruce Hemp, Senior Applications Engineer

Fluid-level measurements can be accurately measured through the wall of a non-metallic tank by placing an air-dielectric transmission line up against the side of the tank and sensing the RF impedance. This article provides an empirical design example that illustrates how a reflectometer device such as the Analog Devices ADL5920 can simplify the design.

Compared to traditional methods of fluid-level sensing that might involve mechanical floats, the reflectometer-based solution offers several benefits, including:

- Fast, real-time fluid-level measurements
- Extensive electronic postprocessing becomes possible
- Contactless design (no contamination of the liquid)
- No moving parts
- Minimal radiated RF field (far field cancels)
- No holes in the tank for an internal sensor (reduce possibility of leaks)
- Intrinsic safety, due to no electrical wires or parts in the tank

Fluid-Level Measurement Overview

Figure 1 shows a block diagram of the overall system, consisting of an RF signal source driving a balanced and terminated air-dielectric transmission line with a reflectometer located inline.

Principle of Operation

Transmission lines suspended in air can be designed for precise characteristic impedance and low RF loss as a result of low loss conductors and the lack of solid dielectric material. Classical plots of E and H vectors show that the electric and magnetic fields are concentrated around the conductors, and their magnitude decays quite rapidly with distance, where distance is measured relative to the size and spacing of the transmission line structure itself. Any nearby dielectric material such as a fluid tank wall and the fluid within will alter the transmission line's electrical characteristics,[1] which can be summarily measured with a reflectometer such as the ADL5920.

Detailed Description

Consider the case of an air-dielectric, low loss transmission line designed for a specific characteristic impedance Z_0 in air. Any added dielectric substance such as a fluid in the near field of the transmission line will:

- Lower the characteristic impedance of the transmission line,
- Reduce the velocity of propagation, thus increasing the effective electrical length of the line, and
- Increase attenuation of the line.

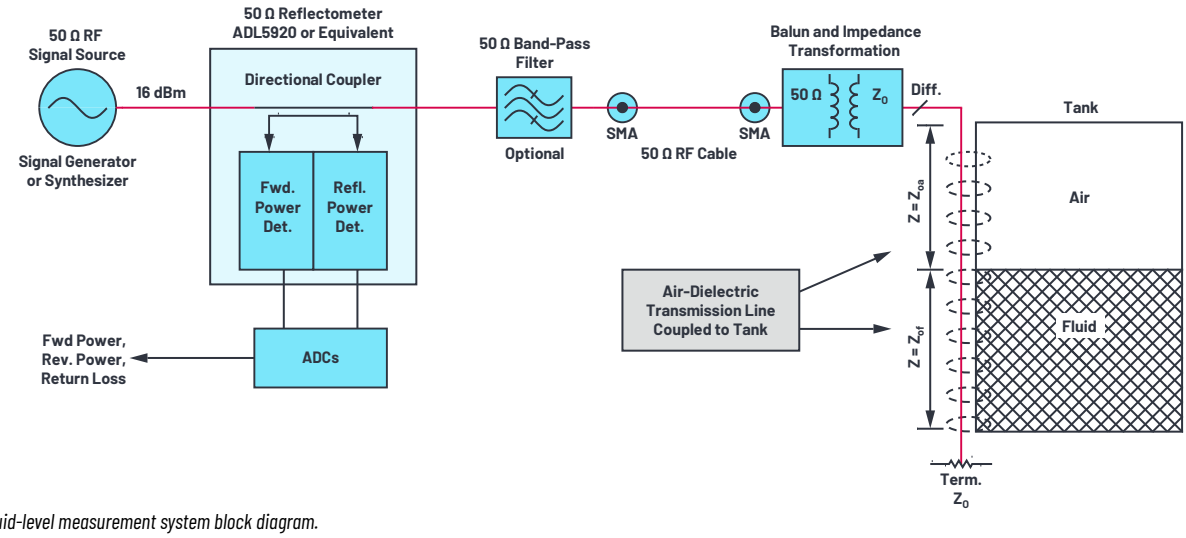

Figure 1. Fluid-level measurement system block diagram.

All three of these effects can combine to create a reduction in return loss, which is directly measurable with a reflectometer device or instrument. With careful design and calibration, return loss can be correlated to fluid level.

To simplify the analysis, consider the air-dielectric transmission line of Figure 1 with impedance set equal to Z_0 before attaching the line to the tank. Because the line is terminated with Z_0, theoretically, there is no reflected energy, and return loss is infinite.

After the transmission line is affixed to the side of a tank, what was one transmission line now behaves as two separate transmission lines, cascaded in a series configuration:

- Above the fluid level, the transmission line is air dielectric, except for the tank wall material. Transmission line impedance Z_{0A} is changed little from its air dielectric value, Z_0. The same is true for transmission line velocity of propagation.
- Below the fluid level, the transmission line impedance Z_{0F} becomes lower compared to Z_{0A}. Electrical length effectively increases, as does attenuation, all because of the extra dielectric material present in the near field of the transmission line.

The impedance of the termination Z_0 at the far end of the transmission line will be transformed when measured by the reflectometer at the source end of the transmission line. The transformation is depicted graphically, approximately as shown in Figure 2. Because Z_{0F} is lower than Z_0, a clockwise Smith chart rotation is created, as shown by the arrows.

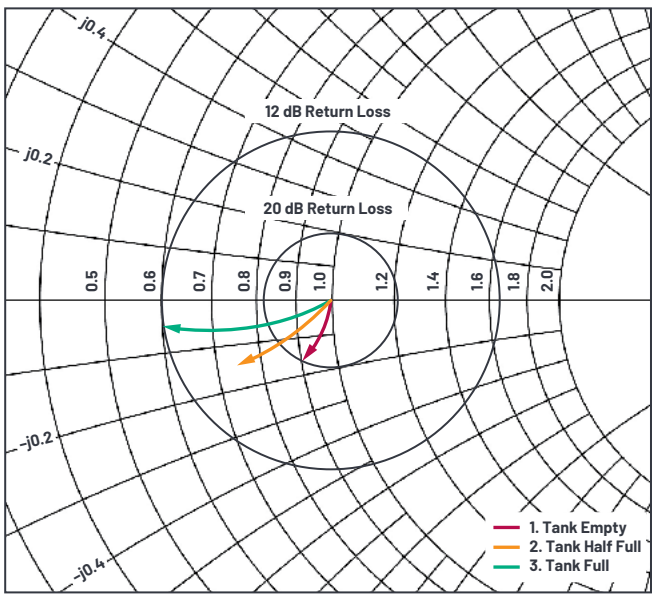

Figure 2. Expanded, normalized Smith chart representation of transmission line input impedance. Trace endpoints depict how fluid level translates to a return loss measurement.

When the transmission line impedance is precisely matched to the resistive termination at the end of the line, there shall be no impedance transformation due to the transmission line. This condition corresponds to the center of the Smith chart, Figure 2, which shows a normalized impedance of 1 + j0 Ω. Return loss should be at least 26 dB *before* the transmission line is attached to the tank.

After attaching the transmission line to an empty tank, the wall material of the tank will contribute some extra dielectric material to the transmission line, thus lowering the impedance of the line to Z_{0A}, and slightly increasing effective electrical length of the transmission line, Trace 1, as exemplified in Figure 2. Return loss should still measure quite well at approximately 20 dB.

As the fluid level rises in the tank, transmission line impedance becomes reduced due to fluid displacing a portion of the air as the dielectric transmission. Transmission line impedance that was Z_{0A} now becomes Z_{0F}. Hence, the center of rotation on the Smith chart moves lower. Simultaneously, the amount of Smith chart rotation increases, because the effective electrical length of the transmission line is increasing. This is depicted by Trace 2 and Trace 3 in Figure 2. Consequently, the reflectometer measures reduced return loss at the generator end of the line.

Because the ADL5920 measures reflection magnitude, not phase, the impedance transformation should be constrained to the bottom half of the Smith chart where the reactive component is negative. Otherwise, impedance is being transformed back toward the center of the Smith chart, causing a magnitude measurement ambiguity. This means the electrical length of the transmission line attached to a full tank should be 90° or less. If electrical length exceeds 90°, the measured return loss will appear to foldback.

A bidirectional RF detector such as the ADL5920 can measure both incident and reflected power in units of dBm, along an RF transmission line of characteristic impedance Z_0 = 50 Ω. The ADL5920 is also able to subtract these two readings, directly measuring return loss in dB.

What Is Return Loss?

Simply stated, when an RF source is connected to a load, some of the power will be transferred to the load, and the remainder will be reflected back toward the source. The difference between these two power levels is the return loss. It's essentially a measure of how well-matched the load is to the source.

Purpose of the Balun

The balun serves to drive each conductor with equal but opposite polarity ac voltage, and thus serves two primary purposes:

- Reducing stray RF coupling to and from the transmission line. This is important for regulatory emissions and susceptibility compliance. Far-field EMI in either direction is reduced by cancellation.
- Transforming impedance. Higher impedance means wider spacing of the transmission line elements, which means deeper electric field penetration into the container. The result is more change in return loss vs. fluid level, which means a more sensitive fluid-level measurement.

The balun should be designed to provide good common-mode rejection ratio (CMRR) over the entire pass band of the band-pass filter.

Is a Band-Pass Filter Necessary?

The optional band-pass filter of Figure 1 is recommended whenever stray RF could couple into the transmission line. A band-pass filter will be very helpful for reducing or eliminating interference from Wi-Fi, cellular, and PCS services, land mobile radio, and all other outside signals that are not in the same frequency band as the desired source.

For best results, it is recommended that the band-pass filter design features low insertion loss, with return loss commensurate with that of the return loss measurement; that is, approximately 30 dB or better if possible.

Basic Design Procedure

The design procedure outline is approximately as follows:

- Choose an operating frequency based on the length of the transmission line. Normally, the transmission line length will be about the same as the tank height or slightly longer. Operating frequency should be chosen such that transmission line length is typically one-tenth to one-fourth of the RF wavelength in the air. Figure 3 illustrates this approximate frequency range. A lower frequency will give the best linearity of return loss vs. fluid level, while a higher frequency will give a larger range of return loss signals, but linearity may not be as good, and measurement foldback may occur (Figure 2). If radiated emissions compliance is required, the frequency may be chosen from the list of applicable ISM frequencies.[2]
- Design or choose a balun for the chosen frequency, or frequency band. The balun can be a lumped-element LC or transformer based. The balun should exhibit excellent return loss when terminated at the balanced end.
- Calculate the conductor width and spacing dimensions of the transmission line. A transmission line impedance calculator such as an arbitrary transmission line calculator (ATLC) is useful for this purpose.[3]

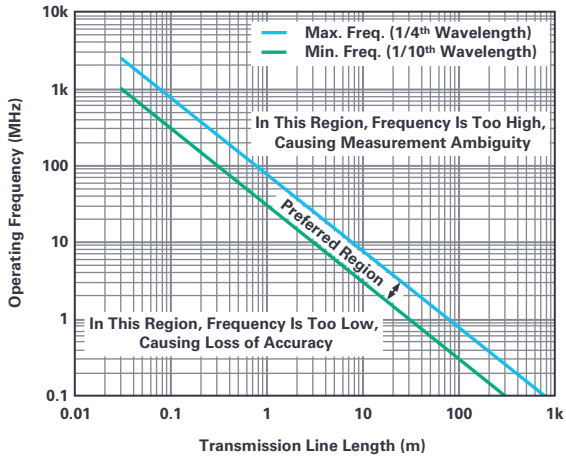

Figure 3. Recommended operating frequency vs. transmission line length.

A Simple Design Example

For demonstration purpose, a fluid-level monitor for an automotive windshield washer tank was devised. The test setup moves water between two identical tanks, one of which is to have a transmission line attached, for fluid-level measurement.

In accordance with the previous outline:

- Because tank height is approximately 6" (0.15 m), a target RF excitation of about 300 MHz is appropriate (see Figure 3).
- Next, an LC balun is designed and constructed for this frequency range. A slight step-up impedance transformation to Z0 is desired to increase sensitivity to the fluid-level variation[4] (see Figure 4). A network analyzer or reflectometer is used to verify approximately 30 dB or better return loss on the single-ended port, with the fixed resistive termination connected directly to the balun, before connecting the transmission line.
- A parallel transmission line is designed and fabricated with Z_0 equal to the resistor value previously used. The transmission line is connected in-circuit, and the resistor termination moves to the end of the line. See Figure 4 and Figure 5. The network analyzer or reflectometer is again used to verify that return loss remains good—approximately 25 dB or better.

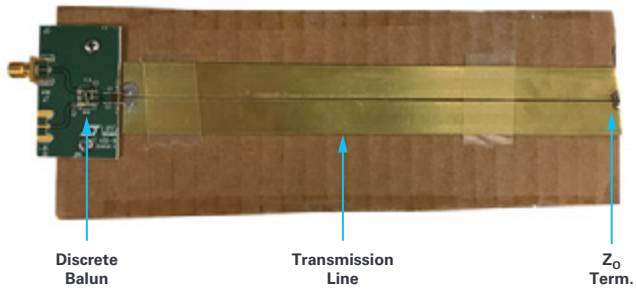

Figure 5. Discrete balun and terminated transmission line, before affixing to the tank.

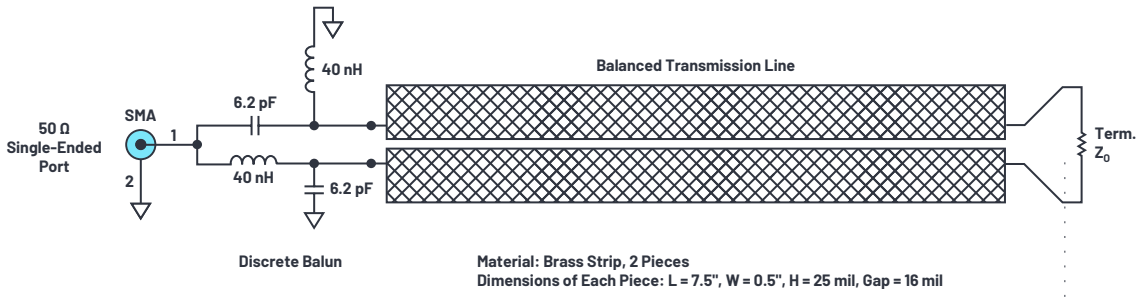

Figure 4. Balun and transmission line used for fluid level sensing example.

Now the transmission line may be attached to the side of the tank, as shown in Figure 6. It's normal to observe return loss drop slightly when affixed to an empty tank due to the detuning effect of the tank wall material as an additional dielectric layer on the transmission line.

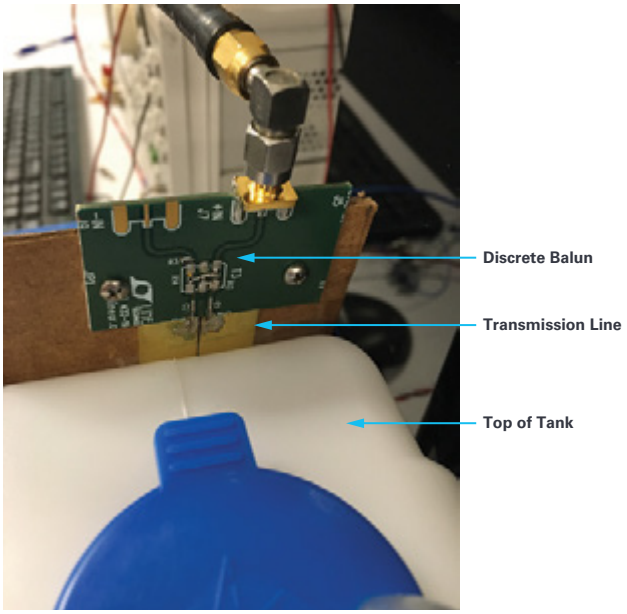

Figure 6. Example design showing the transmission line affixed to the side of the tank.

Example Test Results

Figure 7 shows a complete test setup. The transmission line is affixed to the side of a tank, and the tank has provision for filling and draining in a controlled manner.

Analog Devices' evaluation kit DC2847A is used to easily read ADL5920 reflectometer measurement results. This evaluation kit includes a mixed-signal MCU to read the forward and reflected detector analog voltages. PC software will automatically load and display results in graphical format vs. time. Return loss is easily calculated as the difference between forward and reflected power measurements. Figure 7 shows the complete test setup for the design example.

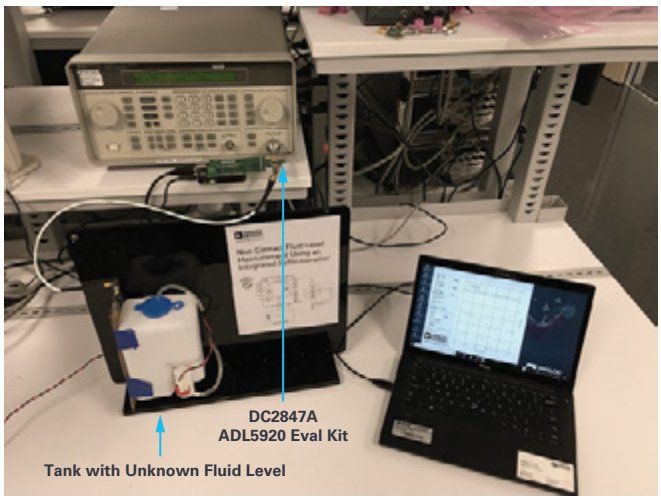

Figure 7. Complete test setup for the design example.

In this design example, fluid-level conditions are established by activating a pump on one of the two tanks. Mass flow rate is relatively constant when a pump is running, so ideally the fluid level in the tank ramps linearly with respect to time. In practice, the tank cross-section is not fully consistent from top to bottom.

Figure 8 shows the test results as fluid level goes from full to empty. As fluid is pumped out of the tank, forward power holds constant, while reflected power falls relatively linearly.

At t = 33 seconds, a visible change in slope occurs. This is believed to be due to the design of the tank. The cross-sectional area of the tank is reduced at the lower end of the tank, as seen in Figure 7, to create space for the pump motor. This introduces a measurement nonlinearity that could be easily corrected in the system firmware if necessary.

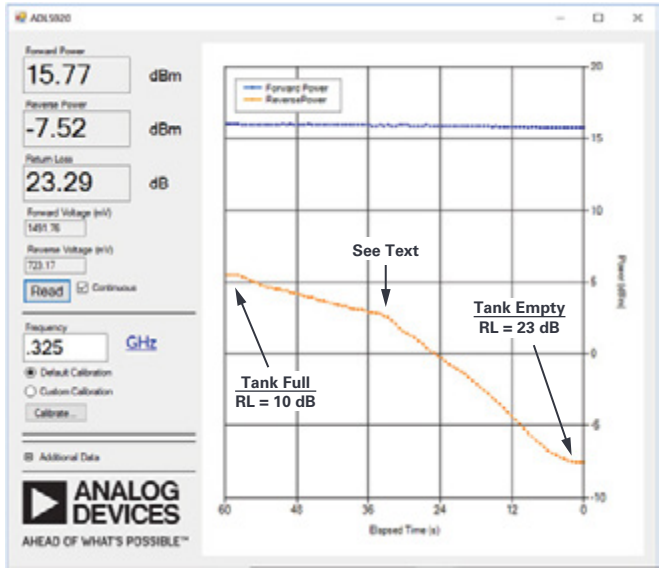

Figure 8. Example test results vs. fluid level. Fluid-level measurement is linear and monotonic, with an exception due to tank design as noted in the text.

Calibration

For best accuracy, reflectometer calibration is required. Calibration will correct for the manufacturing variation of the RF detectors within the reflectometer—namely slope and intercept. The DC2847A evaluation kit supports individual calibration, as seen in Figure 8.

At a higher level, the fluid level vs. return loss also needs calibration. This can be due to the following sources of uncertainty:

- Manufacturing variation of the distance between transmission line and tank wall.
- Variation in tank wall thickness.
- Fluid and/or tank wall dielectric properties could vary vs. temperature.

Systematic nonlinearities may exist, for example, the change in slope observed in Figure 8. If linear interpolation is used, a three-or-more point calibration becomes necessary in this case.

All calibration coefficients will typically be stored in the system's nonvolatile memory, which could be unused code space in an embedded processor application, or a dedicated nonvolatile memory device.

Fluid-Level Measurement Limitations

Directivity of any reflectometer is a key specification. Neglecting balun losses, when the transmission line is precisely terminated with its own Z_0, reflected power goes to zero, and the reflectometer measures its own directivity specification. The higher the directivity specification, the better the ability of the reflectometer to accurately separate the magnitudes of incident and reflected waves.

For the ADL5920, directivity is specified as 20 dB typical at 1 GHz, increasing to approximately 43 dB typical at 100 MHz or lower. This makes ADL5920 well suited for fluid-level measurements where tank height is about 30 mm or higher (see Figure 3).

Application Extensions

For some applications, the basic contactless fluid-level measurement principle can be extended in several ways. For example:

- The measurement may be performed at low duty cycle to conserve power.
- If the fluid level is held constant, return loss measurement may correlate to another fluid property of interest; for example, viscosity or pH.
- Each application is unique. For example, there are some techniques that might offer better accuracy up at the top end of the scale, compared to the bottom end, or vice versa, depending on the application.
- If the tank is metallic, the transmission line will need to go inside the tank. Depending on the application, the transmission line may be submersed.
- Measurements at more than one RF power level can help identify if external RF interference is a contributing error. Many single-chip PLL devices support this feature, which becomes a confidence test for the system, or a self-test.
- Transmission line sensors on two or four sides of the tank can compensate for container tilt along one axis or two axes, respectively.
- If fluid-level threshold measurement is the goal, one or more shorter transmission lines operated at higher frequency can be a good solution.

Conclusion

The development of a single-chip reflectometer device such as the ADL5920 brings with it new types of applications, such as fluid-level instrumentation. Eliminating moving parts, such as a mechanical float that has been used for years, will result in a huge reliability increase. Oil- and fuel-level monitoring may also be possible, opening up many new industrial and automotive applications.

Footnotes

[1] The presence of fluid affects transmission line impedance, loss, and velocity of propagation.

[2] Industrial, scientific, and medical frequencies. Visit *en.wikipedia.org/wiki/ISM_band*.

[3] ATLC: arbitrary transmission line calculator (for transmission lines and directional couplers). Visit *atlc.sourceforge.net*.

[4] Too large of an impedance step-up will make the transmission line difficult to design and transmission line losses may become excessive.

Acknowledgements

The author wishes to thank Michiel Kouwenhoven, James Wong, Bruce Nguyen, and John Chung. Without their guidance and help, this article would not be possible.

About the Author

Bruce Hemp graduated in 1980 from California State University, Fullerton with his B.S. degree in engineering. He has held various systems, board-level, and applications engineering positions. Since 2012, Hemp has been a senior applications engineer and section leader with Analog Devices. He can be reached at *bruce.hemp@analog.com*.

Reduce Power Supply Requirements for Ceramic Capacitors with a High Efficiency, High Frequency, Low EMI DC-to-DC Converter

Zhongming Ye, Senior Applications Engineer

The price of the multilayer ceramic capacitors (MLCCs) has risen sharply over the past several years, tracking the expansion in the number of power supplies used in automotive, data center, and telecommunications industries. Ceramic capacitors are used in power supplies on the output to lower output ripple, and to control output voltage overshot and undershot due to high slew rate load transients. The input side requires ceramic capacitors for decoupling and to filter EMI due to their low ESR and low ESL in high frequency.

The quest to increase the performance of industrial and automotive systems calls for increases in data processing speed of several orders of magnitude, with an increasing number of power-hungry devices squeezed into microprocessors, CPUs, system on chips (SoCs), ASICs, and FPGAs. Each of these complex device types requires a number of regulated voltage rails: typically, 0.8 V for cores, 1.2 V and 1.1 V for DDR3 and LPDDR4, respectively, and 5 V, 3.3 V, and 1.8 V for peripheral and auxiliary components. Buck (step-down) converters are widely used to produce the regulated power supplies from a battery or a dc bus.

For instance, the proliferation of the advanced driver assistance systems (ADAS) in automobiles has dramatically increased ceramic capacitor usage rates. With the rise of 5G technology in telecommunications, where high performance power supplies are required, ceramic capacitor usages will also significantly increase. Core supply currents have increased from several amps to tens of amps, with very tight control of supply ripple, load transient overshot/undershot, and electromagnetic interference (EMI)—features that require additional capacitance.

In many cases, traditional power supply approaches can't keep up with the pace of change. Overall solution size is too big, efficiency is too low, circuit design is too complicated, and the bill of materials (BOM) is too costly. For example, to meet the tight voltage regulation specifications for a fast load transient, a large number of ceramic capacitors are required at the output to store and source significant currents arising from load transients. The total cost of the output ceramic capacitors can reach several times that of the power IC.

Higher power supply operating (switching) frequencies can reduce the effect of transients on output voltage and reduce the capacitance requirements and overall solution size, but higher switching frequencies usually result in increased switching losses, reducing overall efficiency. Is it possible to avoid this trade-off and meet transient requirements at the very high current levels demanded by advanced microprocessors, CPUs, SoCs, ASICs, and FPGAs?

Analog Devices' Power by Linear™ monolithic Silent Switcher® 2 buck regulator family enables compact solution size, high current capability, high efficiency, and, more importantly, superior EMI performance. The LTC7151S monolithic buck regulator uses a Silent Switcher 2 architecture to simplify EMI filter design. Valley current mode reduces the output capacitance requirement. Let's look at a 20 V input to 1 V at 15 A output solution for an SoC.

15 A Solution from 20 V Input for an SoC

Figure 1 shows a 1 MHz, 1.0 V, 15 A solution for SoC and CPU power applications where the input is typically 12 V or 5 V and can vary from 3.1 V to 20 V. Only input and output capacitors, an inductor, and several small resistors and capacitors are necessary to complete a power supply. This circuit can be easily modified to produce other output voltages, such as 1.8 V, 1.1 V, and 0.85 V, down to 0.6 V. The negative return (to the V– pin) of the output rail enables remote feedback sensing of the output voltage close to the load, minimizing feedback errors caused by voltage drops across board traces.

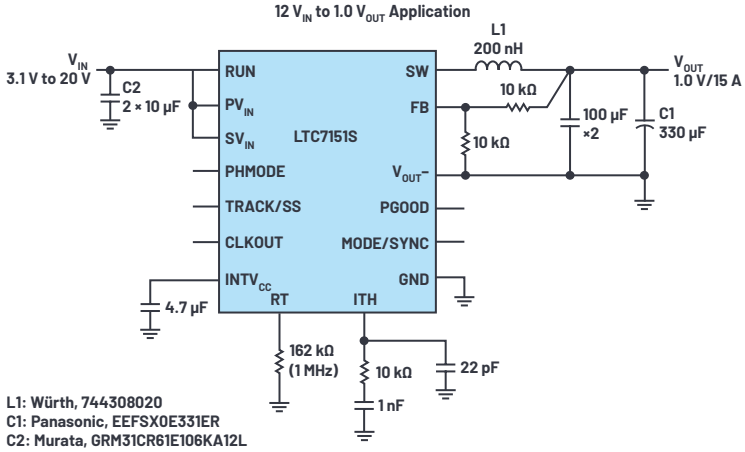

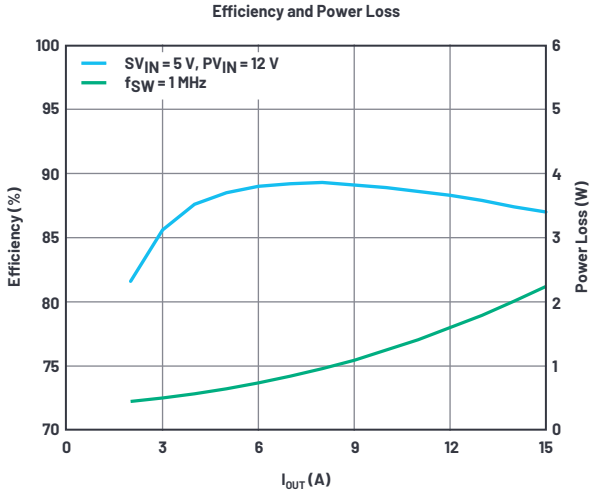

Figure 1. Schematic and efficiency of a 1 MHz, 15 A buck regulator for SoCs and CPUs.

The solution in Figure 1 uses an LTC7151S Silent Switcher 2 regulator, which features high performance integrated MOSFETs in a 28-lead, thermally enhanced 4 mm × 5 mm × 0.74 mm LQFN package. Control is via valley current mode. Protection functions are built-in to minimize the number of external protection components.

The minimum on time of the top switch is only 20 ns (typical), enabling direct step-down to core voltages at a very high frequency. Thermal management features enable reliable and continuous delivery current up to 15 A from input voltage up to 20 V without heatsink or airflow, making it a popular choice for SOCs, FPGAs, DSPs, GPUs, and microprocessors in telecom, industry, transportation, and automotive applications.

The wide input range of the LTC7151S enables use as a first-stage intermediate converter, supporting up to 15 A at 5 V or 3.3 V to multiple downstream point-of-load or LDO regulators.

Meet Tight Transient Specifications with a Minimum Output Capacitor

Typically, an output capacitor is scaled to meet requirements for loop stability and load transient response. These specifications are particularly tight for power supplies that serve processor core voltages, where load transient overshoot and undershoot must be well controlled. For instance, during a load step, the output capacitor must step in, instantaneously supplying current to support the load, until the feedback loop brings up the switch current enough to take over. Typically, overshoot and undershoot are suppressed by installing a significant number of multilayer ceramic capacitors at the output side, fulfilling charge storage requirements during fast load transients.

Additionally, or alternatively, pushing the switching frequency higher can improve fast loop response, but at the price of increased switching losses.

There is a third option: regulators with valley *current-mode control* can dynamically change the regulator's switch T_{ON} and T_{OFF} times to almost instantaneously meet the demands of load transients. This allows for significantly reduced output capacitance to meet fast response times. Figure 2 shows the results of the LTC7151S Silent Switcher regulator immediately responding to a load step from 4 A to 12 A with 8 A/µs slew rate. The controlled on-time (COT) valley current-mode architecture of the LTC7151S allows the switch node pulses to compress during the 4 A to 12 A load step transition. About 1 µs after the start of the rising edge, the output voltage starts its recovery, with overshoot and undershoot limited to 46 mV peak-to-peak. The three 100 µF ceramic capacitors shown in Figure 2a are sufficient to meet typical transient specifications, as shown in Figure 2b. Figure 2c shows typical switch waveforms during a load step.

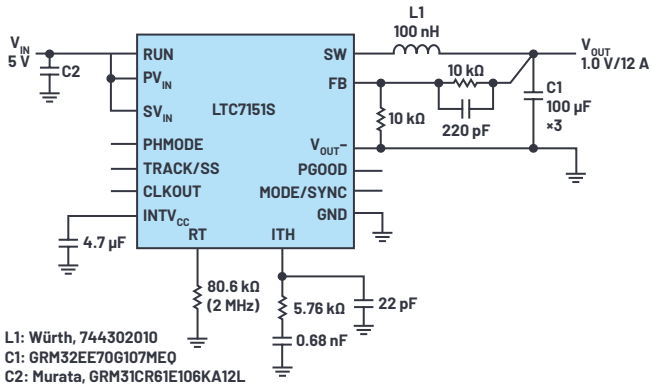

(a)

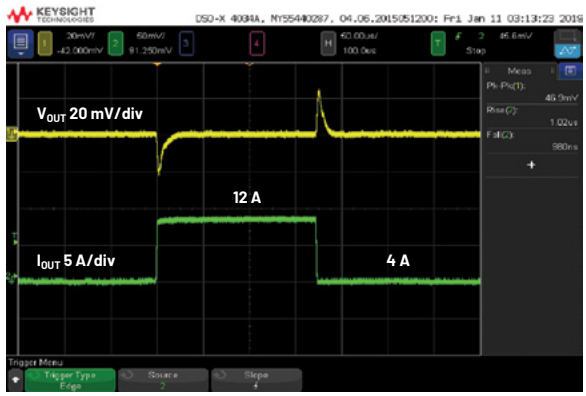

(b)

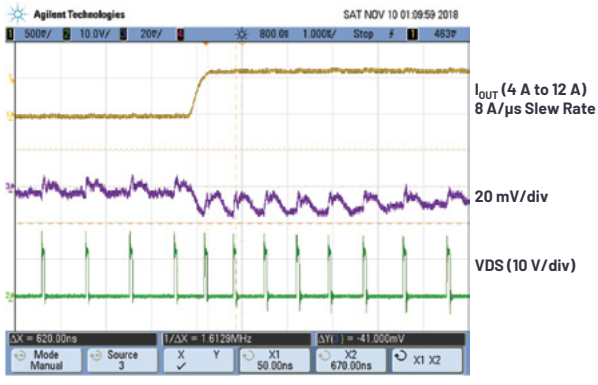

(c)

Figure 2. (a) This 5 V input to 1 V output application runs at 2 MHz, requiring minimal capacitance at the output to quickly and cleanly react to (b) load steps, along with (c) switching waveforms during the load step.

High Efficiency Step-Down at 3 MHz Fits Tight Spaces

In the 4 mm × 5 mm × 0.74 mm package of the LTC7151S are integrated MOSFETs, drivers, and hot-loop capacitors. By keeping these components close, parasitic effects are reduced, allowing for fast turn on/off of the switches with very narrow deadtime. The conduction loss of the antiparallel diode of the switches is greatly reduced. The integrated hot-loop decoupling capacitor and built-in compensation circuit also eliminate design complexity, minimizing the total solution size.

As previously mentioned, the 20 ns (typical) minimum on the top switch allows for very low duty ratio conversion at high frequency, enabling the designer to take advantage of very high frequency operation (such as 3 MHz) to reduce the size and value of the inductor, input capacitor, and output capacitor. Extremely compact solutions are possible for limited space applications, such as portable devices or instruments in automotive and medical applications. Bulky thermal mitigation components such as fans and heatsinks are not necessary when using the LTC7151S, thanks to its high performance power conversion, even at very high frequencies.

Figure 3 shows a 5 V to 1 V solution operating at a 3 MHz switching frequency. A small size 100 nH inductor from Eaton together with three 100 µF/1210 ceramic capacitors delivers a very low profile compact solution for FPGA and microprocessor applications. The efficiency curve is shown in Figure 3b. The temperature rise at room temperature is about 15°C at full load.

Silent Switcher 2 Technology Results in Excellent EMI Performance

Meeting published EMI specifications—such as CISPR 22/CISPR 32 conducted and radiated EMI peak limits—with a 15 A application could mean a number of iterative board spins, involving numerous trade-offs of solution size, total efficiency, reliability, and complexity. Traditional approaches control the EMI by slowing down switching edges and/or lowering the switching frequency. Both have undesired effects, such as reduced efficiency, increased minimum on and off times, and a larger solution size. Brute force EMI mitigation—such as complicated and bulky EMI filters or metal shielding—add significant costs in required board space, components, and assembly, while complicating thermal management and testing.

Analog Devices' proprietary Silent Switcher 2 architecture uses a number of EMI reduction technologies, including integrated hot-loop capacitors, to minimize noisy antenna size. The LTC7151S keeps EMI low by incorporating high performance MOSFETs and drivers, which allows the IC designer to produce a device with built-in minimized switch-node ringing. The result is that the associated energy stored in the hot loop is highly controlled, even when switching edges have high slew rates, enabling exceptional EMI performance while minimizing the ac switching losses at high operating frequencies.

The LTC7151S has been tested in an EMI testing chamber and passed the CISPR 22/CISPR 32 conducted and radiated EMI peak limit with a simple EMI filter in front. Figure 4 shows the schematic of a 1 MHz, 1.2 V/15 A circuit, and Figure 5 shows the radiated EMI CISPR 22 test result in gigahertz transverse electromagnetic (GTEM) cell.

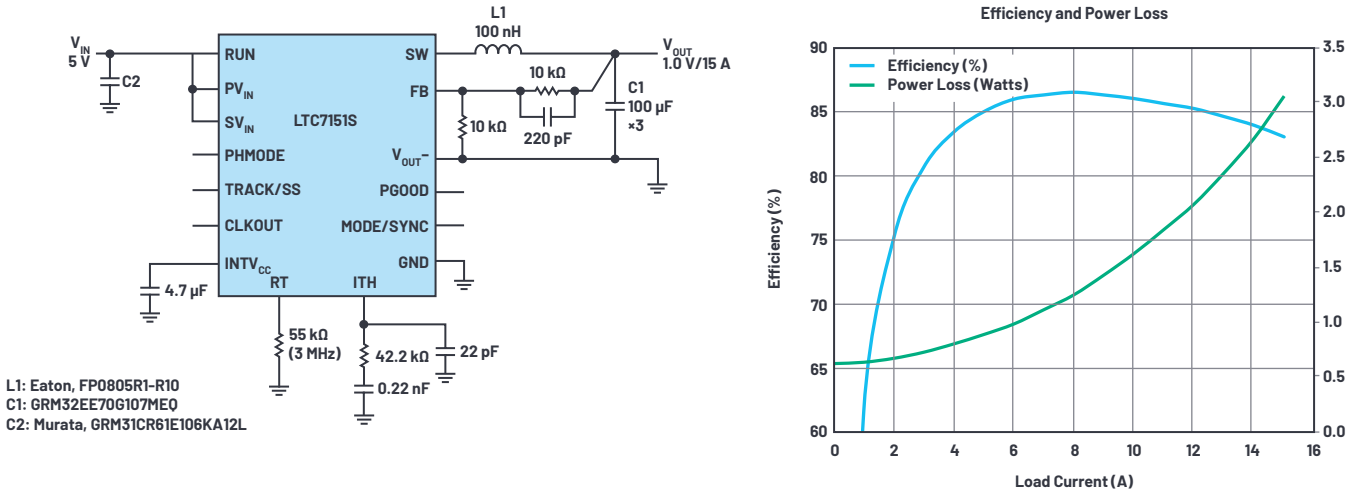

Figure 3. Schematic and efficiency for 5 V input to 1 V/15 A with f_{SW} = 3 MHz.

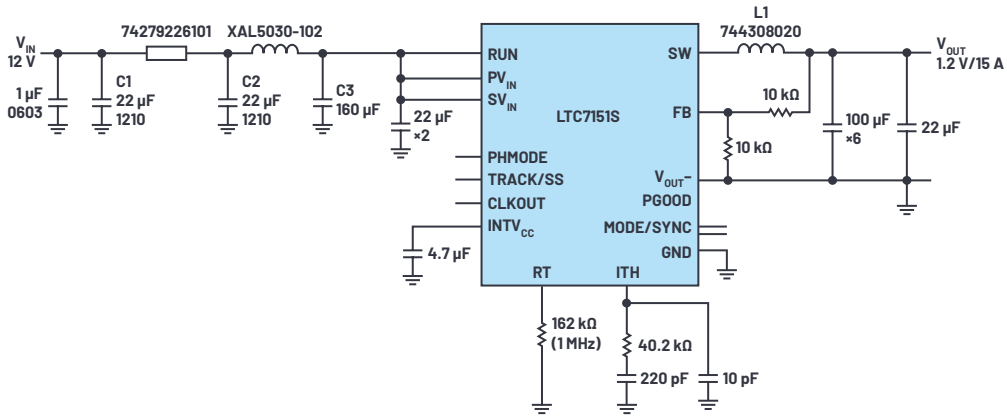

Figure 4. Schematic of a 1.2 V regulator with 1 MHz switching frequency.

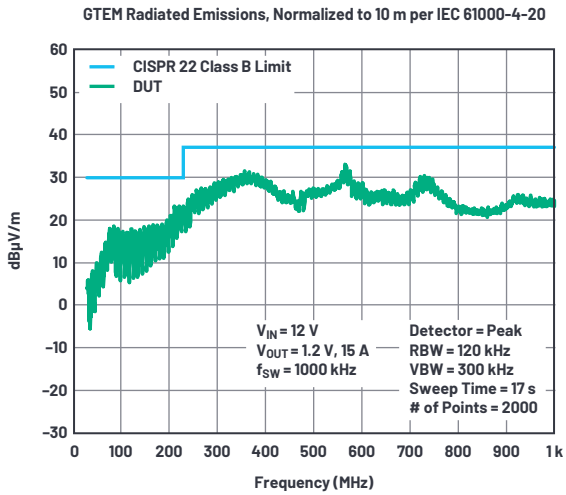

Figure 5. Radiated EMI in GTEM passes the CISPR 22 Class B limit.

Conclusion

The proliferation of intelligent electronics, automation, and sensors in industry and automotive environments has driven up the required number and performance requirements of power supplies. Low EMI, in particular, has increased in priority as a key power supply parametric consideration, along with the usual demands for small solution size, high efficiency, thermal proficiency, robustness, and ease of use.

The LTC7151S meets stringent EMI demands in a very compact size using Silent Switcher 2 techniques from the Power by Linear division of Analog Devices. Thanks to the valley current-mode control and high frequency operation, the LTC7151S dynamically changes the T_{ON} and T_{OFF} times to actively support load transients nearly instantaneously, allowing for much smaller output capacitance and fast response. Its integrated MOSFETs and thermal management enable robust and reliable delivery of current up to 15 A continuously from input ranges up to 20 V.

About the Author

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

RAQ Issue 175: LTspice Audio WAV Files: Using Stereo and Encrypting Voice Messages

Simon Bramble, Senior Field Applications Engineer

Question:

Can you utilize stereo data and encrypt voice messages with LTspice audio WAV files?

Answer:

If music be the food of love, simulate on.

This RAQ explains how to use LTspice audio WAV files for the syntax of stereo (and higher channel count).

LTspice can be used to generate WAV files as an output of a circuit simulation as well as import WAV files to excite a circuit simulation. It is well documented that mono WAV files can be used as an input in LTspice, and LTspice can be used to produce a WAV output. This article details how to use LTspice audio WAV files for the less well-known syntax of stereo (and higher channel count).

LTspice has many superpowers, but its handling of audio files is one of its more impressive talents. While it is fascinating to see a circuit come to life on the computer screen, creating a sound file that can be played outside of LTspice enables engineers to use another sense to evaluate simulations. Using mono LTspice audio WAV files is well documented. This article expands the discussion to stereo (or more channels), and shows how to export stereo data from, and import it to, LTspice via WAV files. It also illustrates a few tips and tricks with WAV files that will enable the reader to further utilize WAV files.

Generating a Stereo WAV File

Let's start by producing a stereo wave file from a mono signal. Figure 1 shows a circuit that generates a 1 V, 1 kHz sine wave and splits it into two channels, alternating the signal between them—the 1 kHz tone is switched in 2-second intervals between CH1 and CH2.

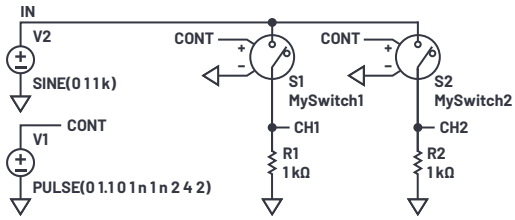

Figure 1. In this simulation, a 1 kHz sine wave is switched in two-second intervals between CH1 and CH2. The resulting two-channel signal is exported to an audio WAV file.

The command .wave **"C:\export.wav" 16 44.1k V(CH1) V(CH2)** digitizes each channel with 16-bit resolution, sampled at 44.1 kSPS, and stores the resultant audio in C:\export.wav. In the command above, each signal listed after the sample rate becomes its own channel in the WAV file. LTspice can store as many as 65,535 channels in a single LTspice audio WAV file—just append signals to the above command as desired.

By default, LTspice's .wave command saves the first listed channel as the *left* audio channel and the second as the *right* audio channel. In this case, when export.wav is played back through a media player, CH1 will be read as the left channel and CH2 will read as the right, regardless of the circuit node naming convention. Note that, by default, CH1 and CH2 are stored as chan 0 and chan 1, respectively, in the .wav file, which is important for reading the file discussed below.

This exported stereo audio file can be used to stimulate another circuit, shown in Figure 2, which uses the two channels from export.wav as signal inputs.

Figure 2. The two stereo channels from export.wav are used to stimulate two independent circuits.

The voltage sources, V1 and V2, are placed as usual, then the voltage signals from export.wav are assigned by holding down the CTRL key and right clicking over each voltage source, revealing the **Component Attribute Editor** as shown in Figure 3.

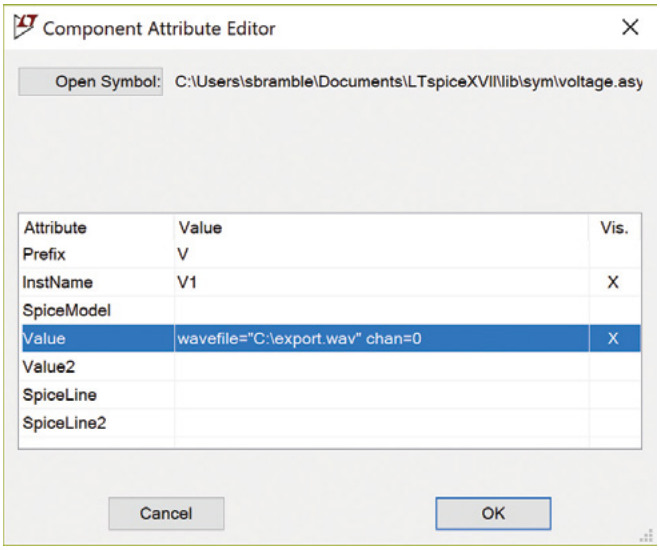

Figure 3. Stereo signals from export.wav are used as inputs to the circuit of Figure 2. Here is the assignment for V1, with the value set to pull Channel 0 from export.wav.

As mentioned above, when the LTspice audio WAV file is first generated, as many as 65,535 channels can be digitized into one WAV file—simply append as many channels as desired onto the end of the .wave command. Remember, by default, LTspice names the first channel Channel 0, the next one is named Channel 1, and so on. In this case, export.wav (generated by the simulation in Figure 1), stores the voltage V(CH1) as Channel 0 and V(CH2) as Channel 1. To play back these channels using a voltage source, specify the .wav file and channel in the value line of the voltage source. In this case:

- To instruct V1 to replay Figure 1's V(CH1): **wavefile="C:\export.wav" chan=0**
- To instruct V2 to replay Figure 1's V(CH2): **wavefile="C:\export.wav" chan=1**

Audio Separation

Theoretically, playing export.wav through a media player should switch between playing the 1 kHz tone entirely through the left speaker (or headphone) for two seconds then through the right speaker for two seconds. Nevertheless, complete stereo separation is not guaranteed and this depends on the quality of the media player used during playback.

Playing export.wav through a laptop showed about 30% of the left channel appearing on the right channel when measured on an oscilloscope as shown in Figure 4.

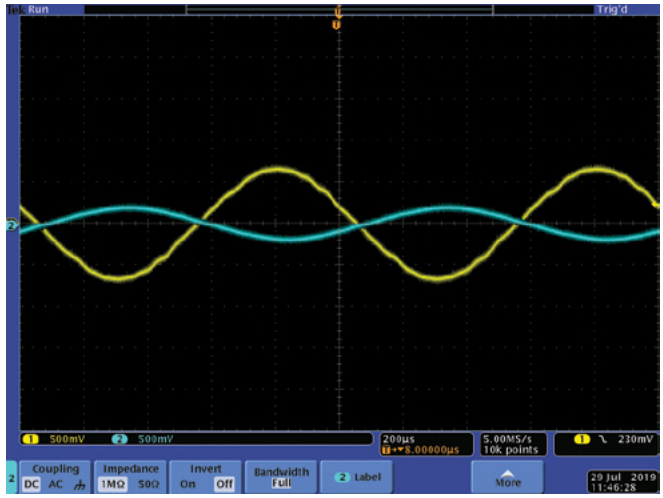

Figure 4. The left (yellow) channel shows about 30% feedthrough into the right (blue) channel when played on a laptop.

Playing the same file on a mobile phone handset (circa 2000) gave a more separated result, showing no perceivable crosstalk, but a slight amount of distortion at full volume as shown in Figure 5.

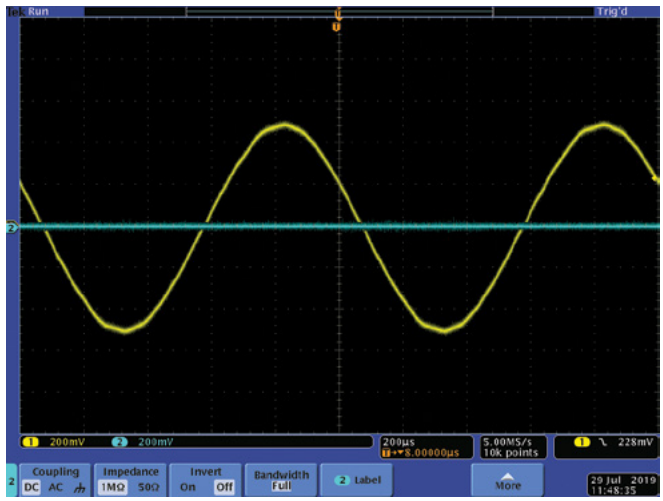

Figure 5. The mobile phone handset from the year 2000 showed no crosstalk, but distortion at full volume.

Repeating the experiment on a later generation 2018 handset showed no perceivable crosstalk, a full 1 V peak signal, and very little distortion as shown in Figure 6. Note the oscilloscope plots were taken with a 500 mV/div sensitivity.

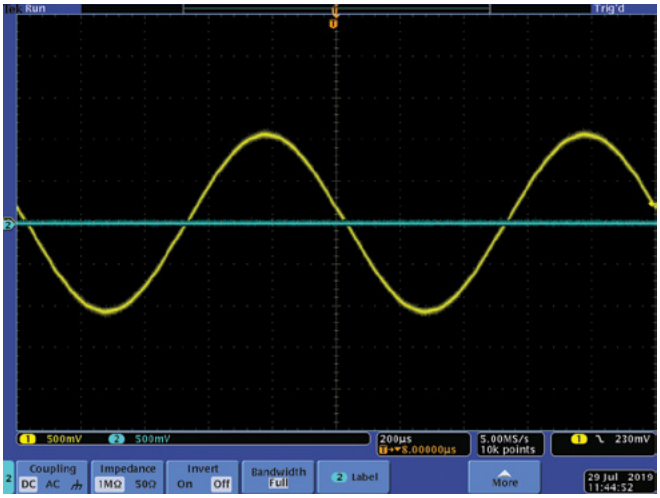

Figure 6. A later generation handset showed significantly better performance for crosstalk, distortion, and amplitude.

The same file was used across all three platforms, showing that LTspice is producing the WAV file with full separation, but resulting playback depends greatly on the quality of the player's audio stage.

Voice Encryption

The circuit in Figure 7 shows a basic method of voice encryption whereby an audio signal is encrypted using a random number sequence, then decrypted.

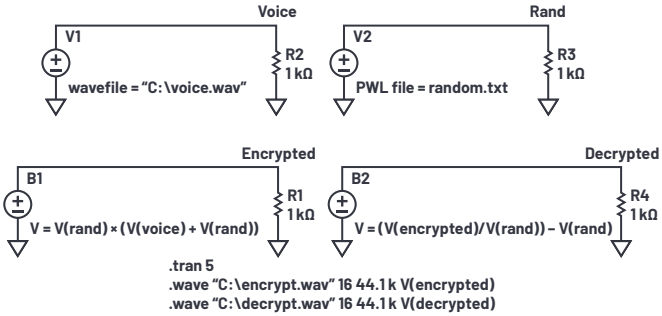

Figure 7. Using a random voltage source to encrypt/decrypt an audio file.

The file voice.wav contains the original audio. An Excel spreadsheet is used to generate a random number sequence with a 100 µs change period. The results are copied into a text file called random.txt. An extract of random.txt is shown in Figure 8.

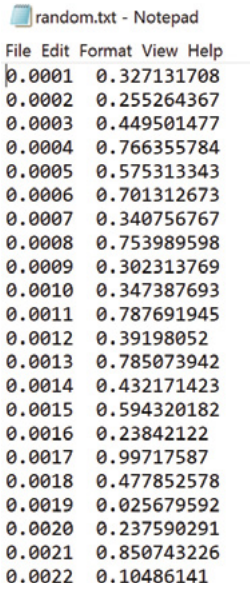

Figure 8. Random voltages generated using Excel and saved in a text file.

This file is used to generate a randomly changing voltage, V(RAND), using the piecewise linear (PWL) voltage source in LTspice.

V(RAND) is added to the voice signal using the behavioral voltage source B1. The output is then multiplied by V(RAND) and the result is sent to the encrypt.wav file. Listening to encrypt.wav showed that the original audio was barely perceivable.

Figure 9 shows the original voice, the encrypted voice, and the decrypted voice signals in the LTspice plot window.

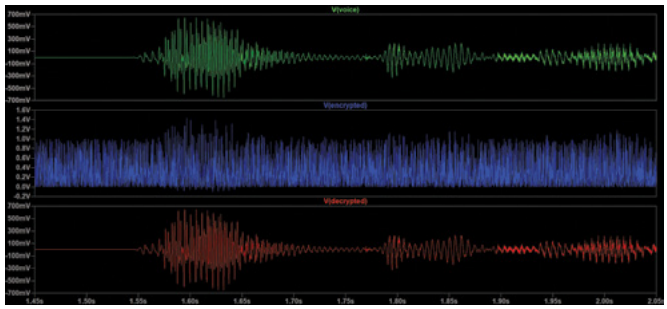

Figure 9. Output of original, encrypted, and decrypted voice signals.

A second behavioral voltage source was then used to decrypt the original audio signal and the result was sent to the file decrypt.wav.

Producing a WAV File from a Differential Voltage Source

The syntax of the .wave command does not allow for the digitization of differential voltages. However, this can easily be overcome by using a behavioral voltage source as shown in Figure 10.

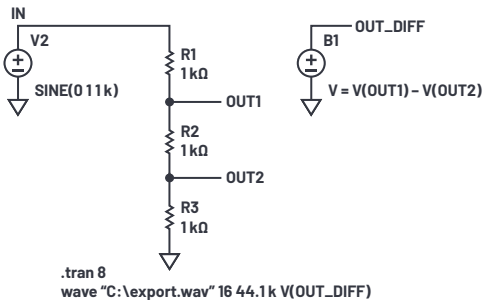

Figure 10. Creating a WAV file from a differential voltage.

The behavioral voltage source outputs a voltage equal to V(OUT1) – V(OUT2) and this can be used in the .wave command in the usual way as shown.

Indeed, the variables inside the behavioral voltage source's function can include any voltage or current in the circuit and these can be manipulated using any of the mathematic functions of LTspice. The net result can then be exported to an LTspice audio WAV file in the normal way.

LTspice is a powerful simulator, but its results do not have to be contained inside LTspice. Using the .wave command, LTspice can import, manipulate, and export audio files for playing on a media player.

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.

Notes

Notes

Notes

Stay Connected with *Analog Dialogue*

Your Engineering Resource for Innovative Design

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

EngineerZone® Support Community

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

Facebook
facebook.com/AnalogDevicesInc

@ADI_News
twitter.com/ADI_News

LinkedIn
linkedin.com/company/analog-devices

Sina Weibo (China)
weibo.com/analogdevices

WeChat (China)

Analog Devices, Inc.
Worldwide Headquarters

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

Analog Devices, Inc.
Europe Headquarters

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

Analog Devices, Inc.
Japan Headquarters

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

Analog Devices, Inc.
Asia Pacific Headquarters

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

M02000-5/20

analog.com

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™