Analog Dialogue

A forum for the exchange of circuits, systems, and software for real-world signal processing • Volume 49, Number 3, 2015

Analog Dialogue
Since 1967

analog.com/analogdialogue

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

Editor's Notes

IN THIS ISSUE

Four Quick Steps to Production: Using Model–Based Design for Software–Defined Radio

Part 1—the Analog Devices/Xilinx SDR Rapid Prototyping Platform: Its Capabilities, Benefits, and Tools

This article series examines the advances in platforms and tools that allow developers to quickly simulate and prototype wireless radio systems while establishing and maintaining a deployable path to production. As a real-world walkthrough of the process, the authors will design and prototype a wireless SDR platform that receives and decodes ADS-B aircraft signals. In this first part, the article discusses the Analog Devices/Xilinx® SDR prototyping system, its capabilities and benefits, and a brief description of the tool flow. (Page 3)

New Advances in Energy Harvesting Power Conversion

Today many power management integrated circuits are available that were specifically designed for use in energy harvesting applications. They enable systems to run with smaller harvesters or make energy harvesting solutions possible that could not have been designed a few years ago. This article looks at several energy harvesting applications and describes a high efficiency dc-to-dc conversion and voltage regula-tion solution that solves energy harvesting challenges. (Page 10)

A Low Power Data Acquisition Solution for High Temperature Electronics Applications

A growing number of applications require data acquisition systems that must operate reliably at very high ambient environments, such as downhole oil and gas drilling, avionics, and automotive. This article presents a new reference design for high temperature data acquisition, characterized from room temperature to 175°C. (Page 13)

Analyzing, Optimizing, and Eliminating Integer Boundary Spurs in Phase–Locked Loops with VCOs at up to 13.6 GHz

A phase-locked loop (PLL) and voltage controlled oscillator (VCO) outputs an RF signal at a certain frequency, and ideally this signal would be the only signal present at the output. In reality, there are unwanted spurious signals and phase noise at the output. This article discusses the simulation and elimination of one of the more trouble-some spurious signals—integer boundary spurs. (Page 19)

Interleaving ADCs: Unraveling the Mysteries

Time interleaving is a technique that allows the use of multiple identical analog-to-digital converters to process regular sample data series at a faster rate than the operating sampling rate of each individual data converter. This technique is frequently utilized in military and electronic instrumentation applications where there is a need to continually push the state-of-the-art in data conversion speeds, resolutions, and perfor-mance. This article explains the data converter interleaving technique in technical detail as well as focusing on some of the practical challenges associated with implementing this technique. (Page 22)

Zero–Drift Amplifiers: Now Easy to Use in High Precision Circuits

A zero-drift amplifier, as the name suggests, is an amplifier with offset voltage drift very close to zero. It uses auto-zero or chopping technology, or a combination of both, to continuously self-correct for dc errors over time and temperature. This enables the amplifier to achieve microvolt level offsets and extremely low offset drifts, making it uniquely suited for signal conditioning circuits that require high gain and precision performance. This article explores the architecture of zero-drift ampli-fiers and provides insight into considerations for designing with these precision devices in drift-critical applications. (Page 27)

Jim Surber [jim.surber@analog.com]

Product Introductions: Volume 49, Number 3

Data sheets for all ADI products can be found by entering the part number in the search box at analog.com.

July

Analog Dialogue

Four Quick Steps to Production: Using Model–Based Design for Software–Defined Radio

Part 1—the Analog Devices/Xilinx SDR Rapid Prototyping Platform: Its Capabilities, Benefits, and Tools

By Di Pu, Andrei Cozma, and Tom Hill

Abstract

There is a significant gap between the concept of a wireless system and the realization of that working design. Bridging this gap typically involves teams of engineers with a variety of different skill sets (such as RF, SW, DSP, HDL, and embedded Linux®), and in many cases projects get derailed early in the development stage because of the difficulty in coordinating the efforts of these varied design entities.

In this four part article, we will examine the advances in platforms and tools that allow developers to quickly simulate and prototype wireless systems while establishing and maintaining a deployable path to production. As a real-world example of the process, we will prototype a wireless SDR platform that receives and decodes automatic dependent surveillance broadcast (ADS-B) signals to allow us to detect and report the position, altitude, and velocity of the commercial aircraft flying in our vicinity. The resources required in this case are MATLAB® and Simulink and the skills to integrate and embed hardware/software. The hardware platform will be the Analog Devices/Xilinx software-defined radio (SDR) prototyping system. Using MATLAB and Simulink® the following tasks will be performed:

- Design of signal processing algorithms used to decode ADS-B messages
- Simulation of the RF transceiver receiving ADS-B signals
- Generation of C and HDL code
- Verification of the HDL code with recorded and live data on the target transceiver and FPGA

The final result will be a working RF SDR design running on production-worthy hardware, which we will take to a local airport and verify its performance and functionality.

The first part of this four part article will discuss the Analog Devices/Xilinx SDR prototyping system, its capabilities and benefits, and a brief description of the tool flow. The second part will review the automatic dependent surveillance broadcast signals and explain how to decode their information in MATLAB and Simulink in simulation. The third part will describe and showcase how to use hardware in the loop (HIL) and capturing signals with the target transceiver, but still doing the signal processing on the host in Simulink for verification. The fourth part will show how to take the algorithm developed in Part 2, verified in Part 3, and use HDL Coder and Embedded Coder from MathWorks to generate code and deploy it in the production hardware, and finally we'll operate the platform with real-world ADS-B signals at an airport.

Introduction

With the exponential growth in the ways and means by which people need to communicate, modifying radio devices easily and cost effectively has become business critical. Based on this requirement, software-defined radio technology has been widely employed recently since it brings the flexibility, cost efficiency, and power to drive communications forward.[1] The purpose of an SDR system is to implement as much as possible of the modulation/demodulation and data processing algorithms in software and reprogrammable logic so that the communication system can be easily reconfigured just by updating the software and the reprogrammable logic and not making any changes to the hardware platform.

With the advent of system on chip (SoC) devices like the Xilinx Zynq® All Programmable SoC that combine the versatility of a CPU and the processing power of an FPGA, designers have the means to consolidate the data processing functions of an SDR system into a single device while integrating additional processing tasks. Processing intensive tasks like the data modulation/demodulation algorithms are offloaded to the programmable logic of the device while tasks like data decoding and rendering, system monitoring and diagnosis, and user interface are deferred to the processing unit.

At the same time, prototyping wireless systems has been a discussion topic for decades but has only in recent years evolved into a complete design flow for FPGAs—from model creation to complete implementation—due to the evolution of the modeling and simulation tools like MATLAB and Simulink from MathWorks. Prototyping wireless systems is transforming the way engineers and scientists work by moving design tasks from the lab and field to the desktop.[2] Now the entire wireless system, such as an SDR system, can be modeled, allowing the engineer to observe the system's behavior and to tune it before it is actually implemented in the field. This has several benefits, such as accelerating system integration and reducing the dependency on equipment availability. Moreover, once the Simulink model for the SDR system is complete, C and HDL code can be generated automatically for implementation on Zynq SoCs, saving time and avoiding the introduction of manually coded errors. The risk is further reduced by linking the system model to a rapid prototyping environment that allows the SDR system to be exercised under real-world conditions.

This first part of the four part article series will discuss the Analog Devices/Xilinx SDR rapid prototyping system, its capabilities and benefits, and a brief description of the tool flow. The article showcases how Analog Devices RF IC technology and reference design hardware and software require a reduced design skill subset, thus enabling customers to mitigate risk and shorten their time to market.
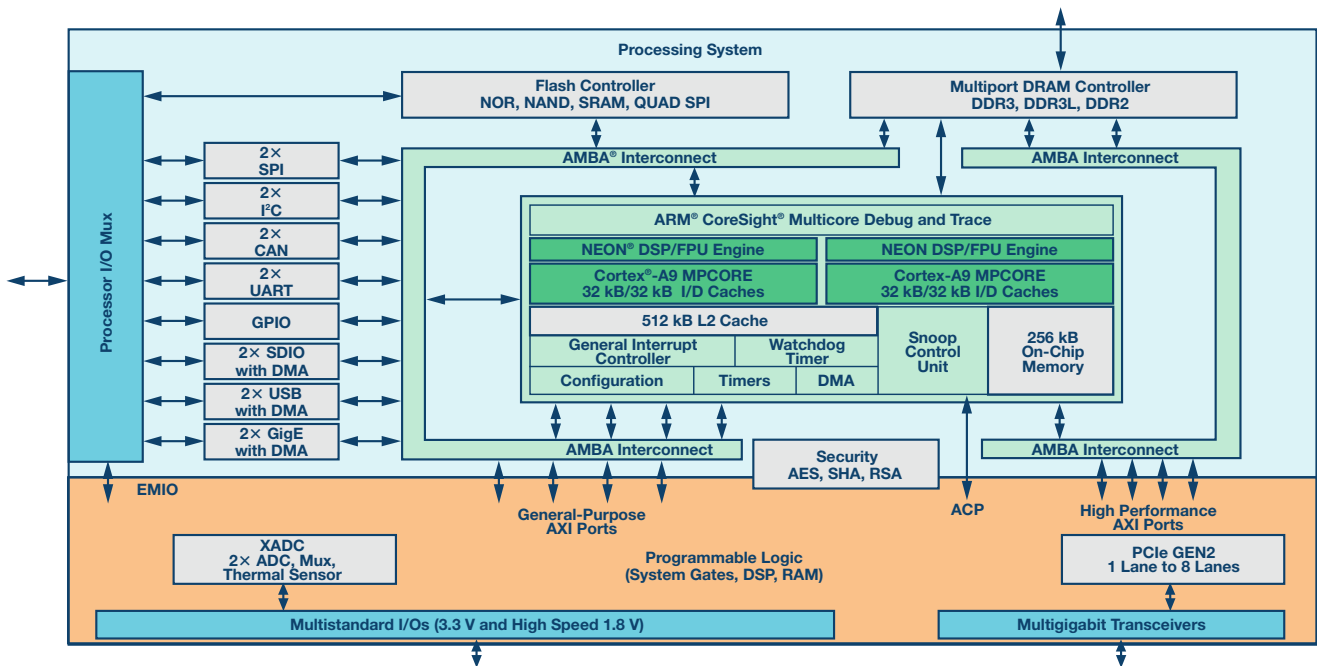
Figure 1. Xilinx Zynq SoC block diagram.

Zynq for SDR

Advanced SDR systems are required to execute a combination of data processing, communication, and user interface tasks that have different processing bandwidth requirements and real-time constraints. The hardware platform chosen to implement such a system must be robust and scalable at the same time allowing for future system improvements and expansion. Xilinx Zynq-7000 All Programmable SoCs fulfill these requirements by supplying a high performance processing system combined with programmable logic as shown in Figure 1.[3] The combination of programmable logic and processing system delivers superior parallel processing power, real-time performance, fast computational speeds, and connectivity versatility.

The processing system side of the Zynq SoC consists of a dual-core ARM® Cortex®-A9 processor combined with a NEON coprocessor and floating-point extensions to accelerate software execution. Embedded Linux or real-time operating systems can be deployed on the dual-core ARM processor to fully benefit from the system's capabilities. The processor is self-contained and can be used without the need to configure the programmable logic, which is a critical element for software developers who will want to start developing code in parallel to hardware developers who will design the FPGA fabric.

On the programmable logic side, the device has up to 444,000 logic cells and 2200 DSP slices that supply massive processing bandwidth, allowing the Zynq device to tackle a variety of challenging signal processing applications. Five high through-put AMBA®-4 AXI high speed interconnects tightly couple the programmable logic to the processing system with the equivalent of more than 3000 pins of effective bandwidth.[4]

AD9361 Agile Wideband RF Transceiver IC for SDR

In recent years, Analog Devices has brought to market revolutionary SDR products to support increasingly evolving SDR requirements and system architectures. Some of the most important Analog Devices products in this field are the AD9361/AD9364 integrated RF agile transceivers. The AD9361 (2 × 2)[5] and AD9364 (1 × 1)[6] are high performance, highly integrated RF transceiver ICs intended for use in SDR architectures in applications such as wireless communications infrastructure, defense electronics systems, RF test equipment and instrumentation, and general software-defined radio platforms. The devices combine an RF front end with a flexible, mixed-signal baseband section and integrated frequency synthesizers, simplifying design-in by providing a configurable digital interface to a processor or FPGA. The chips operate in the 70 MHz to 6 GHz range, covering most licensed and unlicensed bands, and support channel bandwidths from less than 200 kHz to 56 MHz by changing the sample rate, digital filters, and decimation, all programmable within the AD9361 and AD9364 devices.[7] Figure 2 shows the block diagram of a AD9361 device.
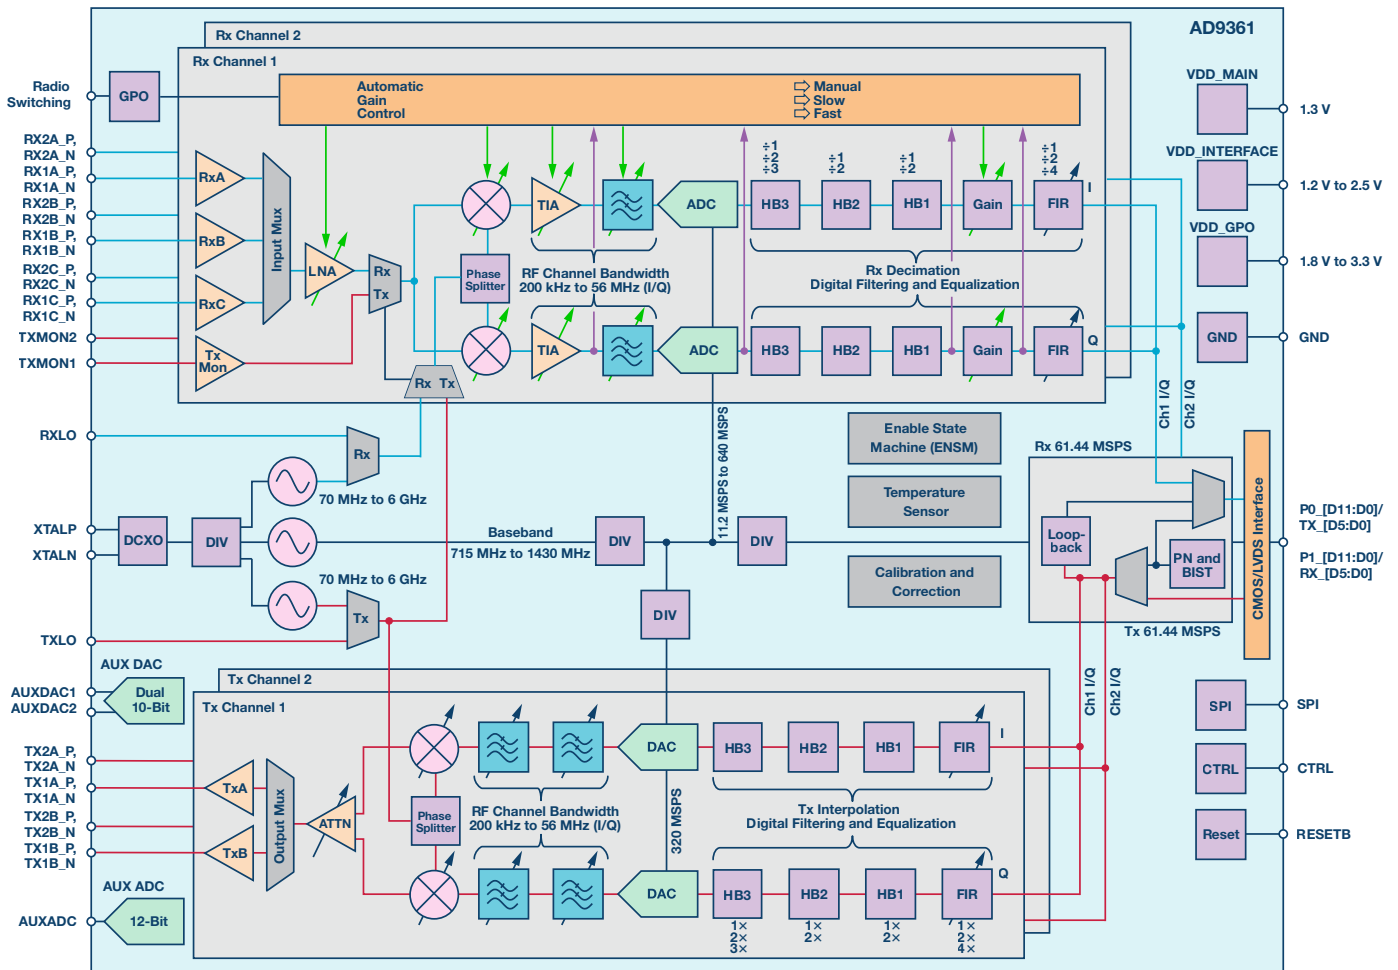
Figure 2. AD9361 block diagram.

In order to help customers shorten time to market and overall development effort, Analog Devices has gone a step further by providing SDR solutions within a complete ecosystem of seamless FPGA connectivity, enabling a rapid prototyping and development environment for complete radio system design. The AD-FMCOMMSx-EBZ rapid development and prototyping boards are a family of high speed analog FMC modules, incorporating AD9361 or AD9364 agile RF transceiver ICs or a discrete signal chain that seamlessly connects to the Xilinx FPGA development platform ecosystem. These boards are fully customizable by software without any hardware changes and come with downloadable Linux drivers and bare metal software drivers, schematics, board layout, and design aid reference materials, all contained on their respective Analog Devices wiki sites. Table 1 summarizes the features of the different FMCOMMSx platforms.

Table 1. FMCOMMSx Platforms

| Platform	Features
AD-FMCOMMS5-EBZ	Integrating two AD9361 2 × 2 agile transceiver ICs, this SDR rapid prototyping board provides full synchronization capability for four receiver channels and four transmitter channels, enabling any subset of a 4 × 4 MIMO system to be created. Wideband 70 MHz to 6 GHz and 2.4 GHz tuned ports are accommodated. AD-FMCOMMS5-EBZ resource wiki page: http://wiki.analog.com/resources/eval/user-guides/ad-fmcomms5-ebz
AD-FMCOMMS4-EBZ	Integrating the AD9364 agile RF transceiver IC, this 1 × 1 SDR rapid prototyping board can be software configured for highest RF performance in the 2400 MHz to 2500 MHz region, or can be software configured to operate over the AD9364's complete RF tuning range of 70 MHz to 6 GHz for system prototyping and development purposes. AD-FMCOMMS4-EBZ resource wiki page: http://wiki.analog.com/resources/eval/user-guides/ad-fmcomms4-ebz
AD-FMCOMMS3-EBZ	Integrating the AD9361 agile RF transceiver IC, this 2 × 2 version of SDR rapid prototyping board supports the AD9361's full RF tuning range of 70 MHz to 6 GHz. This kit is ideal for the wireless communications SDR system architect seeking a unified development platform with wide tuning capabilities. AD-FMCOMMS3-EBZ resource wiki page: http://wiki.analog.com/resources/eval/user-guides/ad-fmcomms3-ebz
AD-FMCOMMS2-EBZ	Integrating the AD9361 agile RF transceiver IC, this 2 × 2 SDR rapid prototyping board is tuned for highest RF performance in the 2400 MHz to 2500 MHz region. This kit is ideal for the RF engineer seeking optimized system performance meeting AD9361 data sheet specifications within this defined range of RF spectrum. AD-FMCOMMS2-EBZ resource wiki page: http://wiki.analog.com/resources/eval/user-guides/ad-fmcomms2-ebz
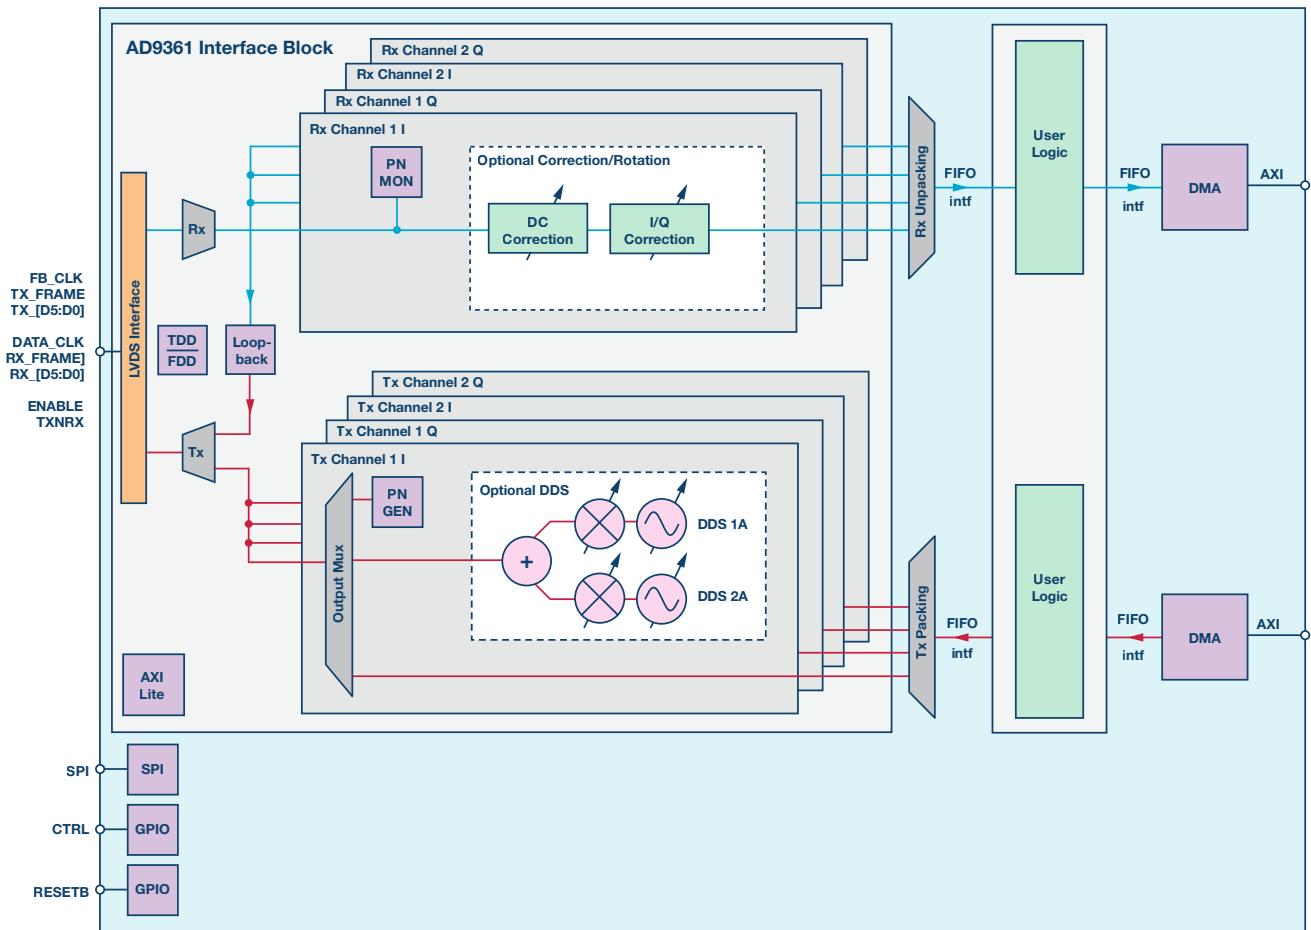
Figure 3. ADI HDL and software infrastructure.

Zynq SDR Rapid Prototyping Platform

Reference Design

Together with the FMCOMMSx platforms, Analog Devices provides a complete Vivado framework, with a Linux and bare metal software infrastructure that can be used both for prototyping purposes as well as a part of the final production system. Figure 3 shows the Analog Devices Zynq Infrastructure to support the FMCOMMSx boards.

This high level diagram shows how the ADI reference design is partitioned on a Xilinx Zynq SoC. An HDMI output is used to display the Linux interface on a monitor while a keyboard and mouse can be connected to the system on a USB 2.0 port. The ARM Cortex-A9 processing system runs Ubuntu Linux provided by Analog Devices. This includes the Linux IIO drivers needed to interface with the Analog Devices FMCOMMS hardware, the IIO Oscilloscope (Scope)[8] user space application for monitoring and control, a libiio server[9] that allows real-time data acquisition and system control over TCP together with clients running on a remote computer, and optional user applications that incorporate C code generated by the Embedded Coder for the controller's Simulink model.

Software Infrastructure

All ADI Linux drivers are based on the Linux Industrial I/O (IIO) subsystem, which is now included in all mainline Linux kernels. The IIO Scope is an open-source Linux application developed by Analog Devices that runs on the dual ARM Cortex-A[9] cores inside the Xilinx Zynq and has the ability to display real-time data acquired from any Analog Devices FMC card connected to the Xilinx Zynq platform. The data can be displayed either as a time domain, frequency domain, or constellation plot. Different popular file formats like comma separated values or .mat MATLAB data files are supported to save the captured data for further analysis. The IIO Scope provides a graphical user interface for changing or reading back the configuration of the Analog Devices FMC cards. The libiio server allows real-time data acquisition and system control over transmission control protocol (TCP) together with clients running on a remote computer.[10] The server runs on an embedded target under Linux and manages real-time data exchange over TCP between the target and a remote client. This library abstracts the low level details of the hardware, and provides a simple yet complete programming interface that can be used for advanced projects. Its modular architecture, well designed API, and built-in network capabilities allow the users to create applications that will run on the system not only where the IIO devices are connected, but also remotely through the network. At first targeted at Linux, it can now be used under Windows as well by using the remote back end of the library. Written in C and licensed under the LGPL, it features bindings for C#, Python, and MATLAB. A MathWorks IIO client[11] is available as a system object to be integrated in native MATLAB and Simulink applications. It is designed to exchange data over Ethernet with an ADI hardware system connected to a FPGA/SoC platform running the ADI Linux distribution, which enables a MATLAB or Simulink model to perform the following functions:

- Stream data to and from a target
- Control the settings of a target
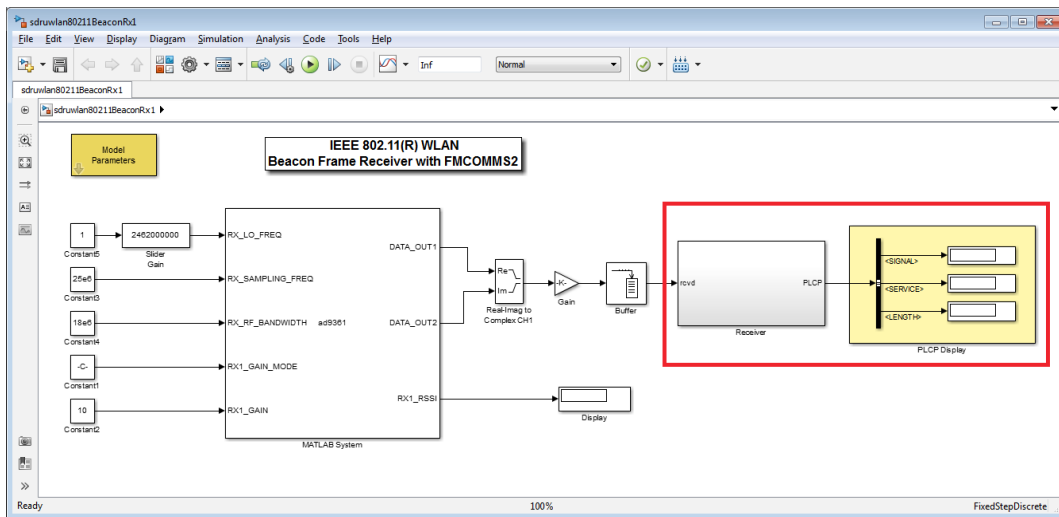- Monitor different target parameters

Figure 4. Screen capture of the beacon frame receiver example.

The IIO System Object™ is available in both MATLAB and Simulink, depending on whether the user calls it from a MATLAB script or incorporates it into a MATLAB System Block. The Linux software and HDL infrastructure provided by ADI for the FMCOMMS platforms is a great environment for prototyping SDR applications together with the tools provided by MathWorks and Xilinx, and it also contains production ready components that can be integrated into the SDR system—helping to reduce the time and cost needed to move from concept to production.

In order to help customers ramp up quickly and easily with the IIO System Object, we provide several MATLAB and Simulink examples based on this interface, such as a beacon frame receiver,[12] QPSK transmitter and receiver,[13] as well as a LTE transmitter and receiver.[14] In these examples, FMCOMMSx platforms are configured by IIO System Object, and are used as RF front ends, which transmit or receive the analog signals over the air. These signals are streamed to or from the target via the IIO System Object. All the other signal processing happens in MATLAB or Simulink. Figure 4 is a screen capture of the beacon frame receiver example, which shows a typical connection between the IIO System Object and the other Simulink blocks.

MathWorks Support for Zynq

MathWorks support for Zynq-based SDR comes from the following four aspects:

1. AD9361 Simulink Model

Since the AD9361 is an integrated RF transceiver chip, signal probing and internal operation monitoring is not really possible. For this reason, MathWorks and Analog Devices have codeveloped a SimRF™ model of the AD9361 that allows a simulation of the chip's operation so that customers can see exactly what's going on under the hood and how the chip performs under different test conditions that are hard to replicate in real life. SimRF provides a component library and simulation engine for designing RF systems using equivalent baseband or circuit envelope blocks, such as amplifiers, mixers, and S-parameter blocks. It is a useful and appropriate tool to model the AD9361 RF transceiver. The system-level AD9361 Agile RF Transceiver model, shown in Figure 5, replicates exactly the functionality of the AD9361 and is available to the users as a MathWorks hardware support package.[15]
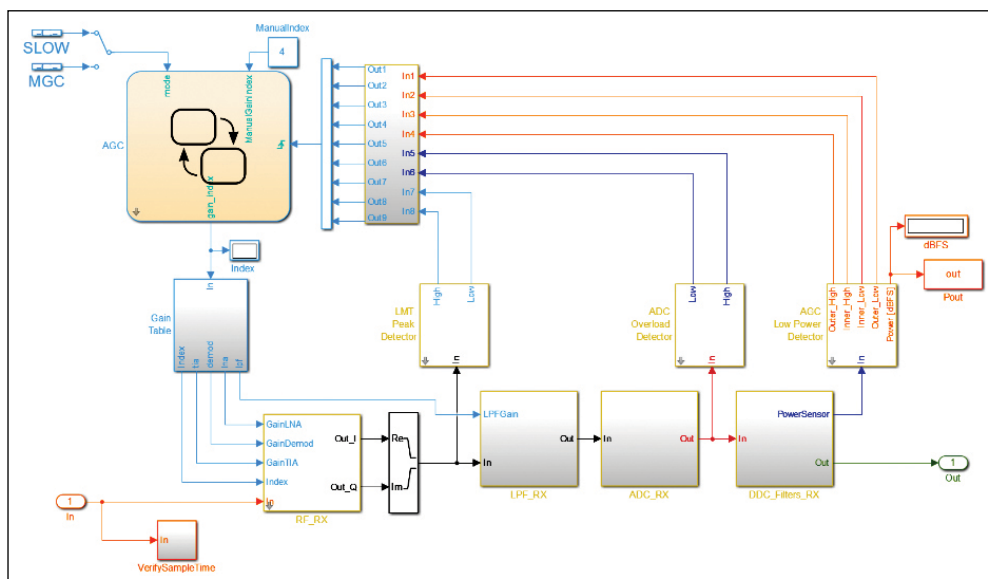

Figure 5. MathWorks SimRF model of AD9361 Agile RF receiver.

The SimRF models have been validated in a lab with power spectral measurements. The characterization of the transceiver's noise and nonlinearity at different frequencies and power levels are identified. The models are then designed to generate the same characterizations, which validates them across the range of design.

With the AD9361 transceiver SimRF models, the users can do the following:

- Predict the impact of the RF imperfections on the test signals
- Use reference tones and LTE signals
- Generate or import test vectors and evaluate the effects of nonlinearity, noise, gain, and phase imbalance, spectral leakage, and other imperfections introduced by the RF transmitter and receiver
- Add interfering signals and evaluate the results in the time or frequency domains

2. Communications and DSP System Toolbox Functions

MathWorks products such as the *Communications System Toolbox*,[TM16] Signal Processing Toolbox,[TM17] DSP System Toolbox,[TM18] and SimRF[19] provide industry-standard algorithms and apps for systematically analyzing, designing, and tuning SDR systems. All of these tools provide the means to create high fidelity SDR models that can be used to verify the behavior and performance of the *communications* system before moving to the actual physical implementation.

3. Simulink Workflow for Zynq

MATLAB and Simulink from MathWorks are environments for multidomain simulation and model-based design that are well suited to simulating SDR systems with communication algorithms. Communication algorithms adjust gain, frequency offset, timing offset, and other performance variables, often for better synchronization between transmitter and receiver systems. Evaluating communication algorithms using simulation is an effective way to determine the suitability of SDR designs and reduce the time and cost of algorithm development before committing to expensive hardware testing. Figure 6 depicts an efficient workflow for designing a communication algorithm by following these steps:

- Build accurate SDR models using the libraries provided by the model-based design environment.
- Simulate system behavior to verify that the system is performing as expected.

- Generate C code and HDL for real-time testing and implementation.
- Test communication algorithms using prototyping hardware.

Once the performance of the SDR system is proven to be satisfactory through simulation and testing on the prototyping hardware, it is safe to take the system implementation and deploy it onto the final production system.
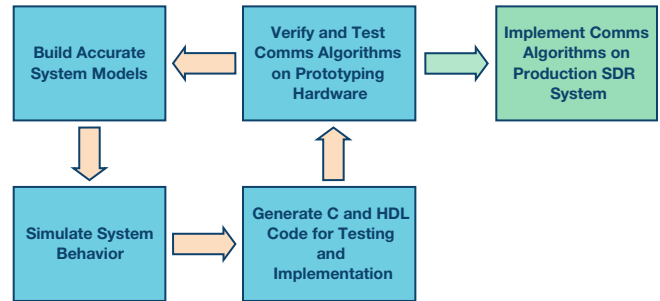
Figure 6. Workflow for communication algorithm design.

4. Simulink Platform Integration to Zynq SDR Kit

Once the SDR system is fully verified in the simulation environment using tools like the Embedded Coder®[20] and the HDL Coder[TM21] from MathWorks, the user can generate C code with Embedded Coder and VHDL or Verilog using HDL Coder, and then deploy the code to prototyping hardware for testing, and afterward, onto the final production system. At this point, software and hardware implementation requirements are specified, such as fixed-point and timing behavior. Automatic code generation helps to reduce the time needed to move from concept to actual system implementation and avoids the introduction of manual coding errors, ensuring that the actual SDR implementation matches the model. Figure 7 depicts a real-life process of the steps needed to model a SDR system in Simulink and transfer it onto the final production system based on a Xilinx Zynq SoC.

The first step is to model and simulate the SDR system in Simulink. At this stage, the communication algorithm is partitioned into blocks that will be implemented in software and blocks that will be implemented into the programmable logic. Once the partitioning and the simulation are complete the SDR model is converted into C code and HDL using Embedded Coder and HDL Coder. A Zynq-based prototyping system is used to verify the performance of the communication
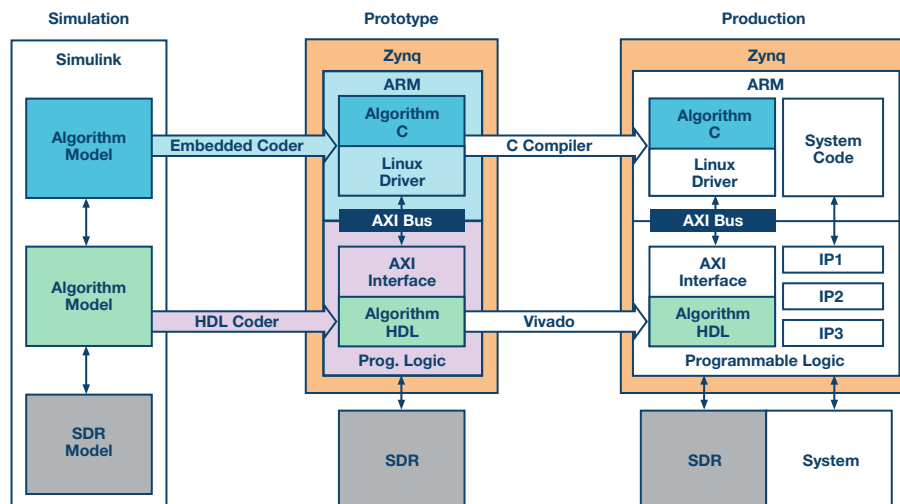
Figure 7. Path from simulation to production.

algorithm and to help further tune the SDR model before moving to the actual production stage. In the production stage, the automatically generated C code and HDL are integrated into the complex production system framework. This workflow ensures that once the communication algorithm reaches the production stage it is fully verified and tested and provides a lot of confidence in the system's robustness. Zynq Hardware Support Packages for Embedded Coder and HDL Coder make it easier to program the Zynq platform by providing a framework for integrated hardware/software design, simulation, and verification that integrate model-based design into the workflow, enabling rapid design iteration cycles and helping to detect and correct design and specification errors early.[22]

Conclusions

This article illustrated the requirements and trends of modern SDR systems and the tools and systems that MathWorks, Xilinx, and Analog Devices bring to the market in order to meet these requirements and help drive toward more performant SDR solutions. By combining the model-based design and automatic code generation tools from MathWorks with the powerful Xilinx Zynq SoCs and Analog Devices integrated RF transceivers, SDR systems design, verification, testing, and implementation can be more effective than ever, leading to higher performance radio systems and reducing the time to market. Analog Devices FMCOMMS platforms paired with the Avnet Zynq-7000 AP SoC provide a great prototyping environment for the SDR algorithms designed using MATLAB and Simulink from MathWorks. The FMCOMMS platforms are accompanied by a set of open source reference designs intended to give a starting point for anyone who wants to evaluate the system and help kick-start any new SDR project.

In the next article in this series, we will advance down the SDR design process as we review the characteristics of automatic dependent surveillance broadcast signals and explain how to decode their information in MATLAB/Simulink in simulation.

For more information about the topics presented in this article, documentation, videos, and reference designs, check out the References section.

References

[1] "What Is Software-Defined Radio?" Wireless Innovation Forum.

[2] Model-Based Design. MathWorks.

[3] Zynq-7000 All Programmable SoC. Xilinx.

[4] Hill, Tom. "Motor Drives Migrate to Zynq SoC with Help from MATLAB." *Xcell Journal*, Issue 87, Second Quarter, 2014.

[5] AD9361.

[6] AD9364.

[7] "Software-Defined Radio Solutions from Analog Devices." Analog Devices.

[8] IIO Oscilloscope. Analog Devices Wiki.

[9] Simulink Libiio. Analog Devices Wiki.

[10] What Is Libiio? Analog Devices Wiki.

[11] IIO System Object. Analog Devices Wiki.

[12] Beacon Frame Receiver Example. Analog Devices Wiki.

[13] QPSK Transmitter and Receiver Example. Analog Devices Wiki.

[14] LTE Transmitter and Receiver Example. Analog Devices.

[15] AD9361.

[16] "Communications System Toolbox." MathWorks.

[17] "Signal Processing Toolbox." MathWorks.

[18] "DSP System Toolbox." MathWorks.

[19] "SimRF." MathWorks.

[20] "HDL Coder." MathWorks.

[21] "Embedded Coder." MathWorks.

[22] "Xilinx Zynq Support from Simulink." MathWorks.

Di Pu

Di Pu [di.pu@analog.com] is a system modeling applications engineer for ADI, supporting the design and development of software-defined radio platforms and systems. She has been working closely with MathWorks to solve mutual end customer challenges. Prior to joining ADI, she received her B.S. degree from Najing University of Science and Technology (NJUST), Nanjing, China, in 2007 and her M.S. and Ph.D. degrees from Worcester Polytechnic Institute (WPI), Worcester, MA, U.S.A., in 2009 and 2013—all in electrical engineering. She is a winner of the 2013 Sigma Xi Research Award for Doctoral Dissertation at WPI.

Andrei Cozma

Andrei Cozma [andrei.cozma@analog.com] 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.

Also by this Author:

FPGA-Based Systems Increase Motor-Control Performance

Volume 49, Number 1

Tom Hill

Tom Hill [tom.hill@xilinx.com], system generator product manager, Xilinx, Inc. Tom Hill has over 18 years experience in the EDA industry. Hill oversees all products, strategic, and corporate marketing activities related to Xilinx's Target Design Platforms for DSP. Hill was most recently at AccelChip, Inc., where he was technical marketing manager responsible for product direction and application of high level design methodologies and tools to DSP applications. Prior to AccelChip, Hill held positions as product manager, technical marketing manager, technical marketing engineer, and field applications engineer for various FPGA and ASIC synthesis tools. Hill began his career as a hardware and ASIC design engineer at Allen-Bradley and Lockheed. Hill holds a B.S. in electrical engineering from Cleveland State University.

New Advances in Energy Harvesting Power Conversion

By Frederik Dostal

Energy harvesting has been around for a very long time. I still remember my 1980s pocket calculator with a solar cell powering the computing unit as well as the LCD display. But even much before that, usable energy was harvested in the early days of the electrical revolution by putting generators on river mills that were powered by running water to generate electricity. Today, when we talk about energy harvesting, we usually use the term for a power source to replace battery cells in electric equipment. So the example of the pocket calculator from the 1980s fits quite well into what we try to achieve today with "energy harvesting."

Setup of an Energy Harvesting System

The most important item in energy harvesting systems is obviously a harvester, and the most common one is a solar cell. The electricity generated by the harvester needs to be converted into a useful voltage or current to power the system or charge intermediate energy storage devices such as super capacitors and batteries. When the system is powered, the correct voltage for the electronics needs to be generated. Figure 1 shows the power management unit fulfilling many different tasks. Matching the input impedance to allow for maximum harvested energy, charging an intermediate energy storage, routing power from a conventional primary cell battery, generating the correct output voltage for the system, and monitoring current flows and voltages to generate a reliable system. All these tasks need to be fulfilled at extremely little supply power so that the system can work with a small harvester or sensor. High integration of these functions in the dc-to-dc converter can help to reduce the power needed for such tasks.

The system in Figure 1 shows a typical energy harvesting system for a wireless environmental sensor. These sensors are typically used to sense temperature, humidity, or different gases such as CO_2. There are many other applications for energy harvesting. Industrial applications can be found in security and surveillance in wireless occupancy sensors or in industrial monitoring such as asset tracking and machine monitoring.

Energy harvesting is also used in consumer electronics such as in portable and in wearable devices. In home healthcare applications, wireless patient monitoring has a need to run without, or extend, battery life.

Today energy harvesting is a very popular topic. Many engineers must evaluate whether or not an energy harvesting solution can replace or complement an existing power solution. The reason why such systems are so popular today is that we are finally reaching a point of equilibrium, where power being harvested from relatively low cost and small size harvesters is enough to power very low energy consumption microcontrollers and RF circuitry. Advances have been made with both electricity generation and in energy consumption in the last few years, so that today many applications that were unrealistic five to 10 years ago are now possible and economically feasible.
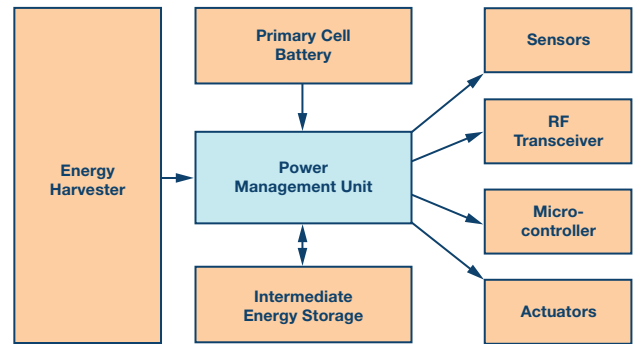

Figure 1. Energy harvesting system setup.

Different Sources of Energy

There are different sources of energy and the most common ones are photovoltaic (PV), thermoelectric (TEG), electromagnetic, piezoelectric, and RF. Photovoltaic and thermoelectric harvesters generate dc voltages, while electromagnetic, piezoelectric, and RF harvesters generate changing or ac voltages. This makes the requirement for a power conversion technology slightly different.

Figure 2 shows different harvesting types and the amount of energy that can roughly be generated with a harvester size of 10 square centimeters. It shows energy generation on the left side and energy consumption for different tasks on the right side. Notice that the power scale in the center is logarithmic. This graph is very important to get a realistic idea of the feasibility of an idea. Often times, designers put work and effort into evaluating an energy harvesting solution only to find out that the harvested energy is not enough to power a given system.
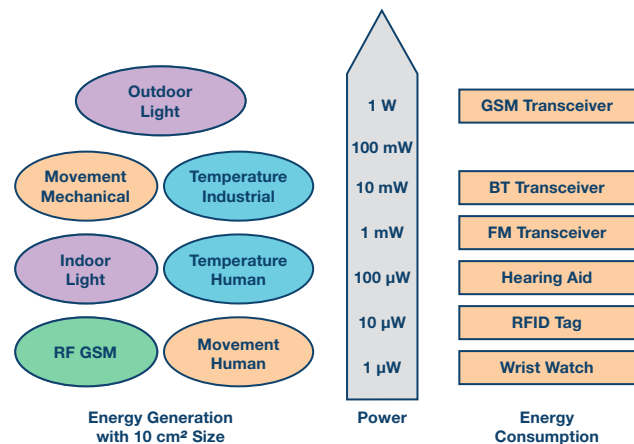

Figure 2. Different sources of energy and the required energy requirements of different applications.

The Importance of the DC–to–DC Converter Unit

Power conversion and management is generally the core of modern energy harvesting systems. While some applications do not use sophisticated power devices, many do. Examples of systems without intelligent power management are a stack of daisy-chained solar cells generating a relatively high dc voltage to either directly power a system or with a simple linear regulator in between. Such systems usually do not have the optimum energy efficiency or do not have a supply voltage that is well regulated. While some loads might work with a widely varying supply voltage, others will not. Future, more advanced systems are more likely to require some form of voltage converter and management block.
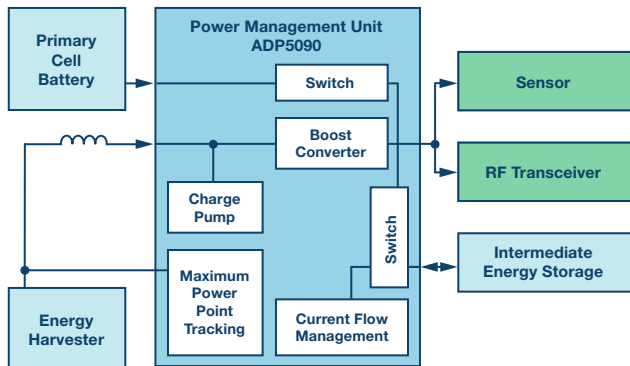
Figure 3. Block diagram of a power management device for energy harvesting applications.

Figure 3 shows the block diagram of a modern power management device for energy harvesting applications. It contains start-up circuitry with a charge pump to allow a start-up voltage of 380 mV on the input. Once the system is running, the internal circuitry of the ADP5090 is powered from the output voltage of the ADP5090. This is the node that powers the load of the energy harvesting system. Once this node is above 1.9 V, the input voltage may go down as low as 80 mV while still being able to harvest energy. This is very helpful in systems that spend a lot of time in suboptimal situations. For example, a solar cell powered indoor sensor. In the morning hours and evening hours, there might be very little light on the solar cell, which generates very little electrical power. Using this time to harvest some energy can help the total power budget in a given time period. One other aspect that helps in such situations is the low quiescent current of the ADP5090. While alert, the current consumption is only 260 nA. Figure 4 shows a typical real-world application. This plot shows different places in a residential building and the typical times a sensor with a solar cell spends in darkness. Certainly this is just a typical case. The amount of light a sensor will see depends on the architecture of the house, including the amount of windows, the amount of electric light that is being used, and the exact location of the sensors. Also the time of year and the location of the house will influence such a chart. The point is, in such varying light conditions, the low power consumption of the ADP5090 helps the total power budget immensely, especially in places that spend most of their time in darkness.
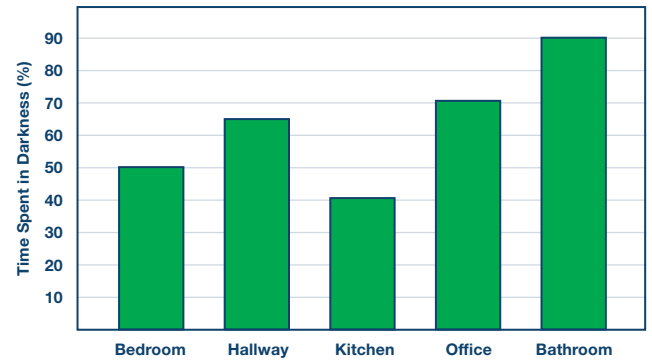
Figure 4. Typical time a sensor spends in darkness in different residential building locations.

The dc-to-dc converter stage in the ADP5090 is quite interesting. It has a regulation loop just like most dc-to-dc converters. However, it does not regulate the output voltage, nor the output current. The regulation loop is primarily set up in a way to regulate the input impedance.

Solar cells have a current and voltage behavior as shown in Figure 5. In an open-loop condition, with no current flow, the provided voltage is at its maximum. Then as current starts to flow, the voltage goes down. At very high currents, the voltage collapses sharply. In the middle of the curve, there is a knee that is the point of peak power. This is the point where the voltage is still relatively high but also where quite a lot of current is drawn. To operate close to the maximum peak power point, we need to track this point. Just setting a fixed current value we draw will not work because the curve of a given solar cell in Figure 5 will shift depending on different light conditions. To track the MPP (maximum peak power point), the ADP5090 stops conducting current on the input, checks the solar cell voltage without it being loaded, and then sets the MPP for the next 16 seconds. After this time period, an open-loop check is performed again. Sixteen seconds turns out to be a good compromise between drifting away from the MPP and interrupting the harvesting action too often.
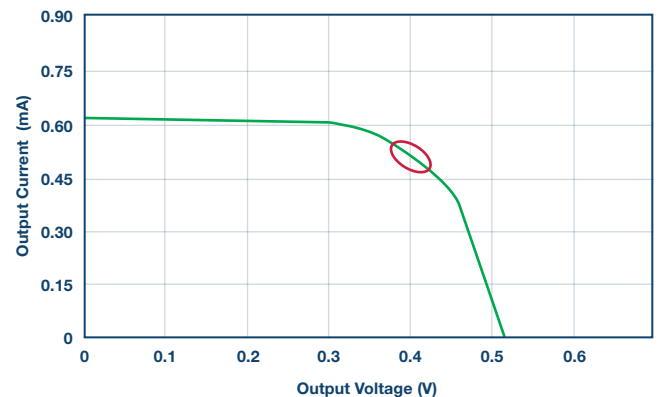
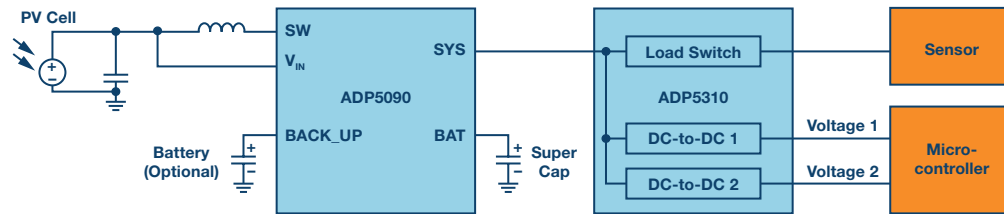Figure 5. Voltage and current plot of a typical photovoltaic cell.

Figure 6. Example of a power management stage for energy harvesting applications.

MPP tracking ensures that the most energy is harvested from a power source, such as a photovoltaic cell or a thermoelectric generator. Still the power management unit has additional tasks. For instance, it needs to control the output voltage in a certain voltage window. The ADP5090 acts like a current source to charge a super cap or a battery. This element is important to decouple energy harvesting with energy consumption. This enables many systems that do not have a constant source of available energy to harvest and perform certain system tasks in set intervals. For example, a sensor in a wireless sensor network that needs to send temperature values every five minutes. If the sensor is powered by a solar cell, it can still operate during dark times due to the intermediate energy storage.

A quite popular architecture today involves attaching energy harvesting to systems that are powered by a primary cell battery. Successful products using a nonrechargeable battery could extend the lifetime of the system by having energy harvesting added. This brings extended operating time without compromising system reliability. For such hybrid systems, the ADP5090 offers the capability to control a primary cell battery. Once there is not enough harvested energy available, the power path from the primary cell battery is routed to power the load directly.

Figure 6 shows a complete energy harvesting power stage with not only the main ADP5090 MPPT energy harvesting IC, but also one second IC, the ADP5310. It is a dc-to-dc converter very efficiently generating two output voltages. The efficiency is close to 90% at 100 µA output current. Additionally, the ADP5310 also has one load switch integrated. This load switch can be used to turn off loads that otherwise would constantly consume power, even when they are not in use.

The ADP5310 step-down converter can accept input voltages all the way up to 15 V. This enables the device to be used directly in ac voltage generators, such as piezoelectric or electromagnetic types. All that is needed is a bridge rectifier and the output voltage can be fed directly into the ADP5310.

Today many power management integrated circuits are available that were specifically designed for use in energy harvesting applications. They enable systems to run with smaller harvesters or make energy harvesting solutions possible that could not have been designed a few years ago. System designers have great ideas that are being implemented right now and which we will be able to see and marvel at in the very near future.

Frederik Dostal

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

A Low Power Data Acquisition Solution for High Temperature Electronics Applications

By Jeff Watson and Maithil Pachchigar

Introduction

A growing number of applications require data acquisition systems that must operate reliably at very high ambient environments, such as downhole oil and gas drilling, avionics, and automotive. While the end uses in these industries are quite different, there are several common signal conditioning needs. The majority of these systems require precision data acquisition from multiple sensors or require a high sample rate. Furthermore, many of these applications have stringent power budgets because they are running from batteries or cannot tolerate additional temperature rise from self-heating of the electronics. Therefore, a low power analog-to-digital converter (ADC) signal chain that maintains high precision over temperature and can be easily used in a wide variety of scenarios is required. Such a signal chain is shown in Figure 1, which depicts a downhole drilling instrument.

While the number of commercially available ICs rated for 175°C is still small, they are increasing in number in recent years, especially for core functions such as signal conditioning and data conversion. This has enabled electronics engineers to rapidly and reliably design for high temperature applications and achieve performance that was not possible in the past. While many of these ICs are well characterized over temperature, this tends to be limited to the function of that device only. There is clearly a lack of circuit level information for these components that demonstrates best practices to achieve high performance in real-world systems.

In this article, we present a new reference design for high temperature data acquisition, characterized from room temperature to 175°C. This circuit is intended to be a complete data acquisition circuit building block that will take an analog sensor input, condition it, and digitize it to an SPI serial data stream. It is versatile enough to be used as a single channel, or it can be scaled for multiple channel simultaneous sampling applications. Recognizing the importance of low power consumption, the power consumption of the ADC scales linearly with the sample rate. The ADC can also be directly powered from the voltage reference, eliminating the need for an additional power rail and the associated power conversion inefficiencies. This reference design is available off the shelf to facilitate testing by designers and includes all schematics, bill of material, PCB artwork, and test software.

Circuit Overview

The circuit shown in Figure 1 is a 16-bit, 600 kSPS successive approximation analog-to-digital converter system using devices rated, characterized, and guaranteed at 175°C. Because many harsh environment applications are battery-powered, the signal chain has been designed for low power consumption while still maintaining high performance.
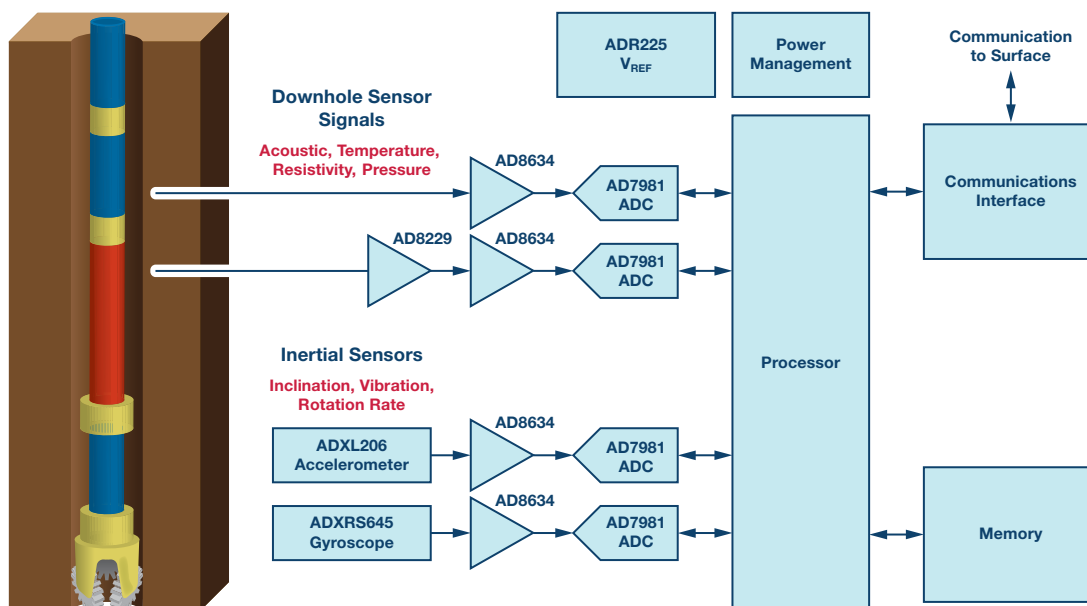

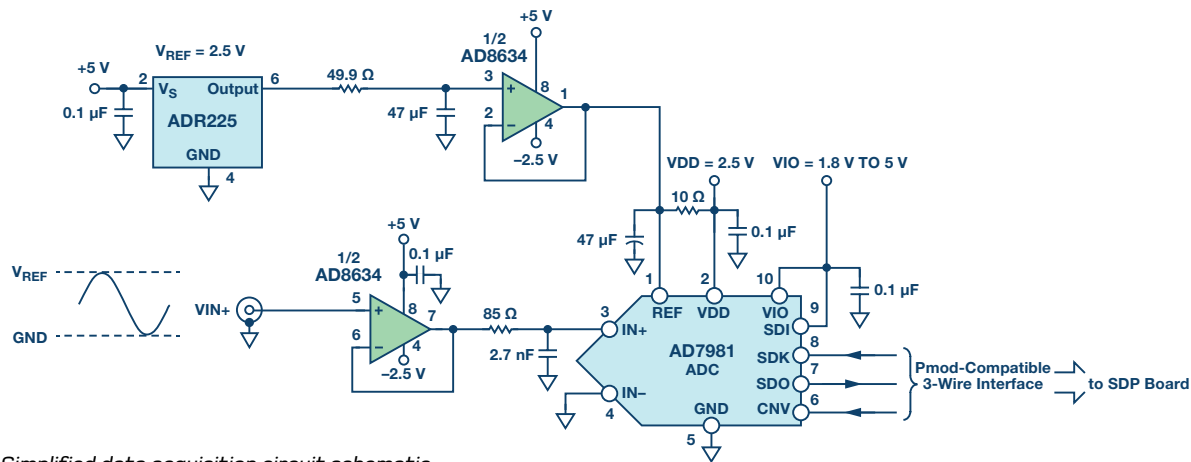Figure 1. Downhole instrument data acquisition signal chain.

Figure 2. Simplified data acquisition circuit schematic.

This circuit uses the AD7981, a low power (4.65 mW @ 600 kSPS), high temperature, PulSAR® ADC, driven directly from the AD8634 high temperature, low power op amp. The AD7981 ADC requires an external voltage reference between 2.4 V and 5.1 V, and in this application, the voltage reference chosen is the micropower ADR225 precision 2.5 V reference, which is also qualified for high temperature operation and has a very low quiescent current of 60 µA maximum at 210°C. All of the ICs in this design have packaging specially designed for high temperature environments, including monometallic wire bonds.

Analog–to–Digital Converter

The heart of this circuit is the AD7981, a 16-bit, low power, single-supply ADC that uses a successive approximation architecture (SAR), capable of sampling up to 600 kSPS. As shown in the diagram in Figure 1, the AD7981 uses two power supply pins: a core supply, VDD, and a digital input/output interface supply, VIO. The VIO pin allows a direct interface with any logic between 1.8 V and 5.0 V. The VDD and VIO pins can also be tied together to save on the number of supplies needed in the system, and they are independent of power supply sequencing. A simplified connection diagram is shown in Figure 3.

The AD7981 typically consumes only 4.65 mW at 600 kSPS and powers down automatically between conversions in order to save power. Therefore, the power consumption scales linearly with the sampling rate, making the ADC well suited for both high and low sampling rates—even as low as a few Hz—and enables very low power consumption for battery-powered systems. Additionally, oversampling techniques can be used to increase the effective resolution for low speed signals.
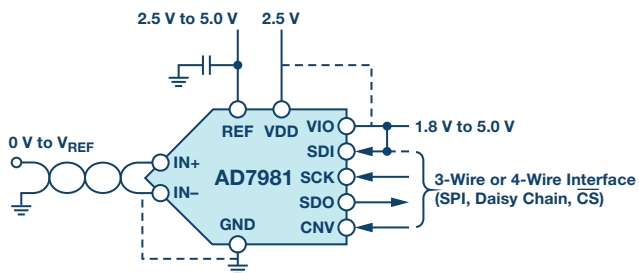

Figure 3. AD7981 application diagram.

The AD7981 has a pseudo differential analog input structure that samples the true differential signal between the IN+ and IN− inputs and rejects the signals common to both inputs. The IN+ input can accept the unipolar, single-ended input signal from 0 V to V_{REF}, and the IN− input has a restricted range of GND to 100 mV. The pseudo differential input of AD7981 simplifies the ADC driver requirement and lowers power

dissipation. The AD7981 is available in a 10-lead MSOP rated for 175°C.

ADC Driver

The input of AD7981 can be driven directly from low impedance sources, but high source impedances significantly degrade the ac performance, especially total harmonic distortion (THD). Therefore it is recommended to use an ADC driver or op amp, such as AD8634, to drive the input of the AD7981 as shown in Figure 4. At the start of the acquisition time the switch closes, and the capacitive DAC injects a voltage glitch (kickback) on the ADC input. The ADC driver helps to settle this kickback as well as isolate it from the signal source.

The low power (1 mA/amplifier) AD8634 dual precision op amp is suited for this task because its excellent dc and ac specifications are a good fit for sensor signal conditioning and elsewhere in the signal chain. While the AD8634 has rail-to-rail outputs, the input requires 300 mV headroom from the positive and negative rails. This necessitates the negative supply, which was chosen to be –2.5 V. The AD8634 is available in an 8-lead SOIC rated for 175°C and an 8-lead flatpack rated for 210°C.
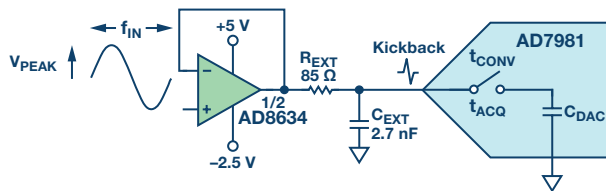

Figure 4. ADC front–end amplifier circuit.

The RC filter between the ADC driver and AD7981 is used to attenuate the kickback injected at the input of the AD7981 and band limits the noise coming to its input. However, too much band limiting can increase settling time and distortion. Therefore, it is very important to find the optimal RC values for this filter. The calculation is primarily based on the input frequency and throughput rate.

From the AD7981 data sheet, internal sampling cap C_{IN} = 30 pF and t_{CONV} = 900 ns, so as described, for a 10 kHz input signal, assuming the ADC is running at 600 kSPS and C_{EXT} = 2.7 nF, the voltage step for a 2.5 V reference would be:

$$V_{STEP} = \frac{2\pi f_{IN} V_{PEAK} t_{CONV} C_{IN}}{C_{EXT} + C_{IN}}$$

$$V_{STEP} = 7.768e - 4 \text{ V}$$

Therefore, the number of time constants required to settle to ½ LSB at 16 bits is:

$$N_{TC} = \ln\left(\frac{V_{STEP}}{\frac{V_{REF}}{2^{N+1}}}\right) = \ln\left(\frac{7.768e-4}{\frac{2.5\,V}{2^{16+1}}}\right) = 3.707$$

The acquisition time of AD7981 is

$$t_{ACQ} = \left(\frac{1}{f_S}\right) - t_{CONV} = \left(\frac{1}{600\,kSPS}\right) - 900\,ns = 7.67e-7$$

We can then calculate the bandwidth of the RC filter using the following equation:

$$\tau = \left(\frac{t_{ACQ}}{N_{TC}}\right) = \left(\frac{7.67e-7}{3.707}\right) = 2.068e-7$$

$$f_{-3dB} = \left(\frac{1}{2\pi\tau}\right) = 769.5\,kHz \rightarrow R_{EXT} = 76.6\,\Omega$$

This is a theoretical value with first-order approximation that should be verified in the lab. We determined through testing that the optimum values were $R_{EXT} = 85\,\Omega$ and $C_{EXT} = 2.7\,nF$ ($f_{-3dB} = 693.48\,kHz$), which gave excellent performance over the extended temperature range to 175°C.

In the reference design, the ADC driver is in unity-gain buffer configuration. Adding gain to the ADC driver will reduce the bandwidth of the driver and lengthen the settling time. In this case, the throughput of the ADC may need to be reduced or an additional buffer as a driver should be used after the gain stage.

Voltage Reference

The ADR225 2.5 V voltage reference uses only 60 µA maximum of quiescent current at 210°C and has a very low drift of 40 ppm/°C typical, making it an ideal part for this low power data acquisition circuit. It has an initial accuracy of ±0.4% and can operate over a wide supply range of 3.3 V to 16 V.

The voltage reference input of the AD7981, like other SAR ADCs, has a dynamic input impedance and should therefore be driven by a low impedance source with efficient decoupling between the REF pin and GND as shown in Figure 5. The AD8634 is well suited as a reference buffer in addition to its ADC driver application.

Another advantage to using a reference buffer is that the noise on the voltage reference output can be further reduced by adding a low-pass RC filter, as shown in Figure 5. In this circuit, a 49.9 Ω resistor and 47 µF capacitor gives a cutoff frequency of approximately 67 Hz.
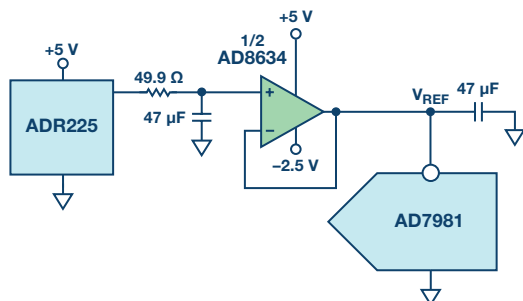
Figure 5. SAR ADC reference buffer and RC filter.

During conversions, current spikes as high as 2.5 mA can occur on the AD7981 reference input. A high value reservoir capacitor is placed as close as possible to the reference input to supply that current and keep the reference input noise low. Typically, a low ESR—10 µF or more—ceramic capacitor is used, but for high temperature applications this is problematic due to the lack of availability of high value, high temperature ceramic capacitors. For this reason, a low ESR 47 µF tantalum capacitor was chosen that has minimal impact to the performance of the circuit.

Digital Interface

The AD7981 offers a flexible serial digital interface compatible with SPI, QSPI, and other digital hosts. The interface can be configured for a simple 3-wire mode for the lowest I/O count, or 4-wire mode that allows options for the daisy-chained readback and busy indication. The 4-wire mode also allows independent readback timing from the CNV (convert input), which enables simultaneous sampling with multiple converters.

The Pmod-compatible interface utilized on this reference design implements the simple 3-wire mode with SDI tied high to VIO. The VIO voltage is supplied externally from the SDP-PMOD interposer board. The interposer board connects the reference design board to the Analog Devices System Development Platform (SDP) board and allows connection to a PC through USB in order to run software to evaluate performance.

Power Supplies

This reference design requires external low noise power supplies for the +5 V and –2.5 V rails. Because the AD7981 is low power, it can be supplied directly from the reference buffer. This eliminates the need for an additional power supply rail—saving power and board space. The proper configuration to power the ADC from the reference buffer is shown in Figure 6. VIO can also be supplied if logic levels are compatible. For the reference design board VIO is supplied externally through the Pmod-compatible interface for maximum flexibility.
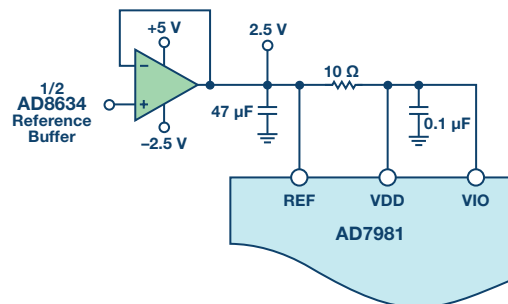
Figure 6. Supplying ADC from reference buffer.

The typical total power consumption for an entire data acquisition solution at 175°C can be calculated as follows:

ADR225: 30 µA × 5 V = 0.15 mW

AD8634: (1 mA × 2 amplifiers) × 7.5 V = 15 mW

AD7981: 4.65 mW @ 600 kSPS

Total power consumption = 19.8 mW

IC Packaging and Reliability

Devices in the Analog Devices high temperature portfolio go through a special process flow that includes design, characterization, reliability qualification, and a production test. Part of this process includes special packaging designed specifically for extreme temperatures. A special material set is used for the 175°C plastic packages in this circuit.

One of the major failure mechanisms in high temperature packaging is the bond wire-to-bond pad interface, particularly when gold (Au) and aluminum (Al) metals are mixed, as is typical in plastic packages. Elevated temperature accelerates the growth of AuAl intermetallic compounds. It is these intermetallics that are associated with bond failures, such as brittle bonds and voiding, which can occur in a few hundred hours as shown in Figure 7.
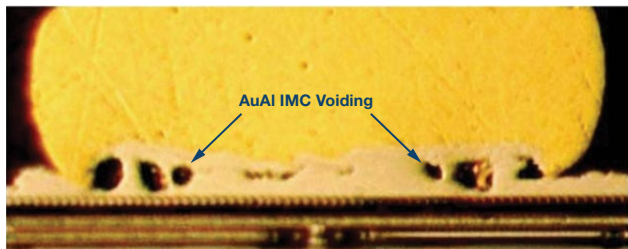

Figure 7. Au ball bond on Al pad, post 500 hours at 195°C.

In order to avoid these failures, Analog Devices uses an over pad metallization (OPM) process to create a gold bond pad surface for the gold bond wire to attach. This monometallic system will not form intermetallics and has been proven reliable in qualification testing with over 6000 hours soak at 195°C, as shown in Figure 8.
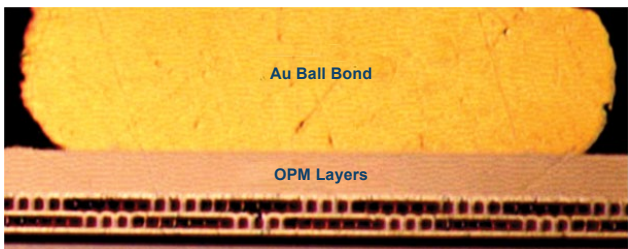

Figure 8. Au ball bond on OPM pad, post 6000 hours at 195°C.

Although Analog Devices has shown reliable bonding at 195°C, the plastic package is rated for operation only to 175°C due to the glass transition temperature of the molding compound. In addition to the 175°C rated products used on this circuit, 210°C rated models are also available in a ceramic flatpack package. Known good die (KGD) are also available for systems that require custom packaging.

Analog Devices has a comprehensive reliability qualification program for high temperature (HT) products that includes high temperature operating life (HTOL), with the parts biased at the maximum operating temperature. HT products are data sheet specified for a minimum of 1000 hours at the maximum rated temperature. Full production testing is the last step required to guarantee performance for each device that is manufactured. Each device in Analog Devices' high temperature portfolio is production tested at elevated temperature to ensure performance is met.

Passive Components

Passive components chosen should be rated for high temperatures. For this design, >175°C thin film, low TCR resistors were used. COG/NPO capacitors were used for low value filter and decoupling applications and have a very flat coefficient over temperature. High temperature rated tantalum capacitors are available in larger values than ceramic and are commonly used for power supply filtering. The SMA connector used on this board is rated for 165°C, so it should be removed for long duration testing at elevated temperatures. Similarly, the insulation material on the 0.1" header connectors (J2 and P3) is only rated for short durations at high temperature but should also be removed for prolonged high temperature testing. For

production assemblies, there are a number of options for HT rated connectors from multiple vendors, such as Micro-D style connectors.

PCB Layout and Assembly

The PCB for this circuit is designed so that the analog signals and digital interface are on separate sides of the ADC, with no switching signals running under the ADC IC or near analog signal paths. This design minimizes the amount of noise that is coupled into the ADC die and supporting analog signal chain. The pin out of the AD7981, with all its analog signals on the left side and all its digital signals on the right side, eases this task. The voltage reference input, REF, has a dynamic input impedance and should be decoupled with minimal parasitic inductances, which is achieved by placing the reference decoupling capacitor as close as possible to the REF and GND pin and making the connection to the pin with a wide, low impedance trace. The layout of this board was purposely designed with components only on the top side of the board in order to facilitate testing over temperature where heat would be applied from the bottom of the board. A photo of the complete assembly is shown in Figure 9. For further layout recommendations, see the AD7981 data sheet.

Figure 9. Reference design circuit assembly.

For high temperature circuits, special circuit materials and assembly techniques should be used to ensure reliability. FR4 is a common material used for PCB laminates, but commercial grade FR4 has a typical glass transition temperature around 140°C. Above 140°C, the PCB will begin to break down, delaminate, and cause stress on components. A widely used alternative for high temperature assemblies is polyimide, which typically has a glass transition temperature of greater than 240°C. A four layer polyimide PCB was used in this design.

The PCB surface is also a concern, especially when used with solders containing tin because of the tendency to form bronze intermetallics with copper traces. A nickel-gold surface finish is commonly used, where the nickel provides a barrier, and the gold provides a good surface for the solder joint bonding. High melting point solder should also be used with a good margin between the melting point and maximum operating temperature of the system. SAC305 lead free solder was chosen for this assembly. With a melting point of 217°C, there is a margin of 42°C from the highest operational temperature of 175°C.

Performance Expectations

The AD7981 is specified for typical SNR of 91 dB with a 1 kHz input tone and a 5 V reference. However, when using low reference voltages such as 2.5 V, as is common in low power/low voltage systems, some degradation in SNR is expected. We can calculate the theoretical SNR based on the specifications of the components used in the circuit. From the AD8634 amplifier data sheet, its input voltage noise density is 4.2 nV/√Hz and current noise density is 0.6 pA/√Hz. Since the noise gain of AD8634 in buffer configuration is 1, and assuming negligible series input resistance for the current noise calculation, the equivalent output noise contribution from the AD8634 would be:

$$\sqrt{(4.2e-9)^2 + 0 \times (0.6e-12)^2} = 4.2 \text{ nV}/\sqrt{\text{Hz}}$$
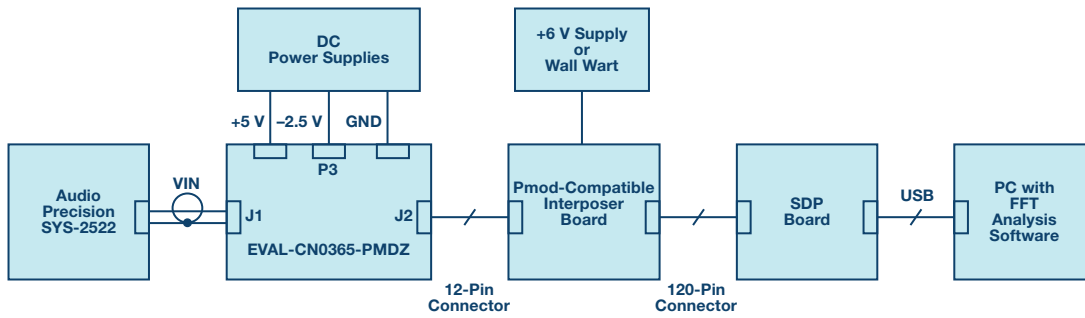
Figure 10. Characterization test setup.

The total integrated noise at the ADC input

(after RC filter $\left(\dfrac{1}{2\pi(85)(2.7e-9)}\right)$) would be:

$$4.2 \text{ nV}/\sqrt{\text{Hz}} \times \sqrt{\left(693.48e3 \times \frac{\pi}{2}\right)} = 4.38 \text{ μV rms}$$

The rms noise of AD7981 can be calculated from its data sheet typical SNR of 86 dB for a 2.5 V reference.

$$e_{AD7981} = 10^{\left(-\frac{\text{SNR}}{20}\right)} \times V_{signal\text{-}rms} =$$

$$10^{\left(-\frac{86}{20}\right)} \times 0.884 \text{ V} = 44.3 \text{ μV rms}$$

The total rms noise of the complete data acquisition system can be calculated by using root-sum-square (RSS) of AD8634 and AD7981 noise sources:

$$V_{noise\text{-}rms} = \sqrt{(4.38e-6)^2 + (44.3e-6)^2}$$

$$= 44.51 \text{ μV rms}$$

So, the theoretical SNR of the data acquisition system at room temp (25°C) can be estimated as shown below:

$$\text{SNR} = 20 \times \log\left(\frac{V_{signal-rms}}{V_{noise-rms}}\right)$$

$$= 20 \times \log\left(\frac{0.884 \text{ V}}{44.51 \text{ μV rms}}\right) = \sim 86 \text{ dB}$$

Test Results

The ac performance of the circuit was evaluated over temperature from 25°C to 185°C. It is critical to use a low distortion signal generator to characterize performance. For this test, the Audio Precision SYS-2522 was used. In order to facilitate testing in an oven, extension harnesses were assembled so that only the reference design circuit was exposed to an elevated temperature. The block diagram of the test setup is shown below in Figure 10.

From our calculations in the previous setup, we expect to achieve approximately 86 dB SNR at room temperature. This compares well to our measured value of 86.2 dB SNR at room temperature as shown in the FFT summary in Figure 11.
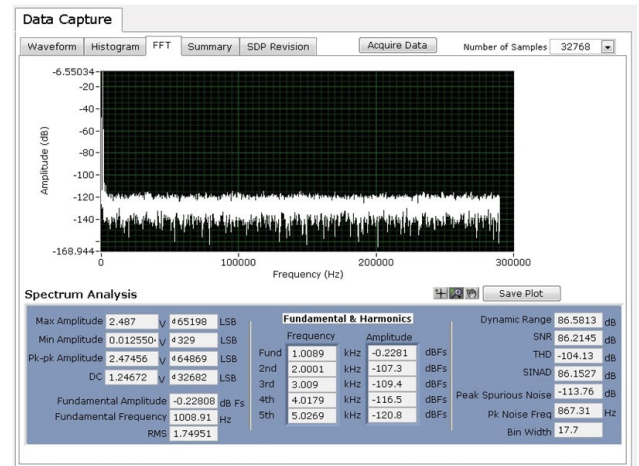

Figure 11. AC performance with 1 kHz input tone, 580 kSPS, 25°C.

When this circuit is evaluated over temperature, SNR performance only degrades to approximately 84 dB at 175°C as shown in Figure 12. THD remains better than –100 dB, as shown in Figure 13. The FFT summary for the circuit at 175°C is shown in Figure 14.
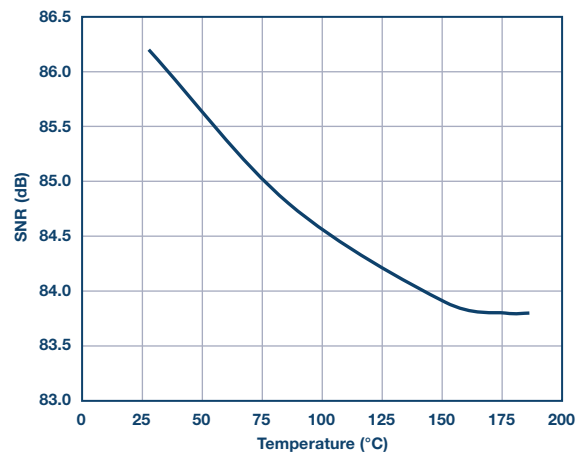

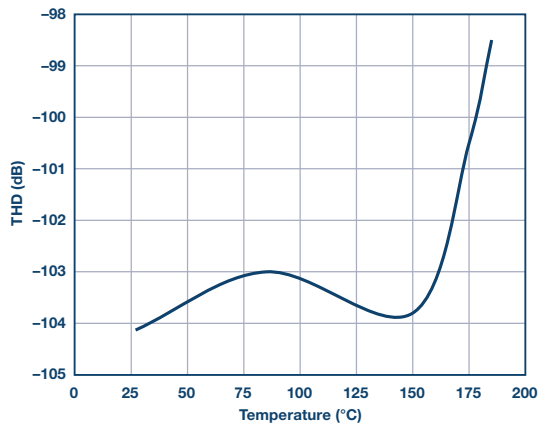Figure 12. SNR over temperature, 1 kHz input tone, 580 kSPS.

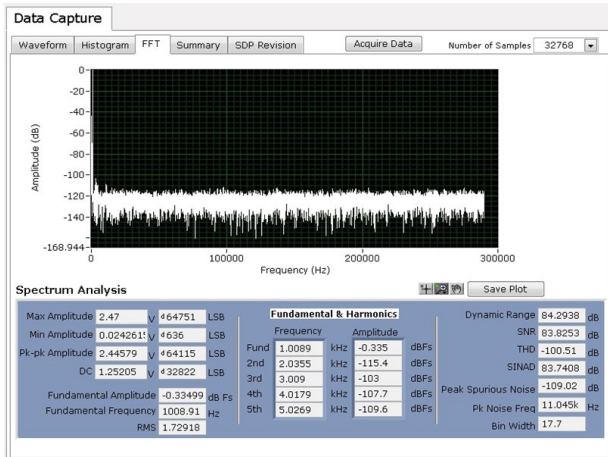*Figure 13. THD over temperature, 1 kHz input tone, 580 kSPS.*

Figure 14. AC performance with 1 kHz input tone, 580 kSPS, 175°C.

Summary

In this article we presented a new reference design for high temperature data acquisition, characterized from room temperature to 175°C. This circuit is a complete low power (<20 mW) data acquisition circuit building block that will take an analog sensor input, condition it, and digitize it to an SPI serial data stream. This reference design is available off the shelf to facilitate testing by designers and includes all schematics, bill of materials, PCB artwork, test software, and documentation.

For more information on this reference design, please visit analog.com/CN0365. For more information on ADI's high temperature portfolio, please visit analog.com/hightemp.

References

Arkin, Michael, Jeff Watson, Michael Siu, and Michael Cusack. "Precision Analog Signal Conditioning Semicondcutors for Operation in Very High Temperature Environments." Proceeding from the High Temperature Electronics Network, 2013.

AD7981.

Digilent Pmod Interface Specification.

Harman, George. *Wire Bonding in Microelectronics*. McGraw Hill, Feb 2010.

Phillips, Reggie, et al. "High Temperature Ceramic Capacitors for Deep Well Applications." CARTS International 2013 Proceedings. March 2013.

Siewert, Thomas, Juan Carlos Madeni, and Stephen Liu. "Formation and Growth of Intermetallics at the Interface Between Lead-Free Solders and Copper Substrates." Proceedings of the APEX Conference on Electronics Manufacturing. Anaheim, California. April 2003.

Walsh, Alan. "Front-End Amplifier and RC Filter Design for a Precision SAR Analog-to-Digital Converter." *Analog Dialogue*, Volume 46, Number 4, 2012.

Walsh, Alan. "Voltage Reference Design for Precision Successive-Approximation ADCs." *Analog Dialogue*, Volume 47, Number 2, 2013.

Watson, Jeff and Gustavo Castro. "High Temperature Electronics Pose Design and Reliability Challenges." *Analog Dialogue*, Volume 46, Number 2, 2012.

Zedníček, Tomas, Zdeněk Sita, and Slavomir Pala. "Tantalum Capacitor Technology for Extended Operating Temperature Range."

Jeff Watson [jeffrey.watson@analog.com] is a systems applications engineer in the Instrumentation, Aerospace and Defense business unit at Analog Devices, focusing on high temperature applications. Prior to joining ADI, he was a design engineer in the downhole oil and gas instrumentation industry and off-highway automotive instrumentation/controls industry. Jeff received his bachelor's and master's degrees in electrical engineering from Penn State University.

Jeff Watson

Also by this Author:

High Temperature Electronics Pose Design and Reliability Challenges

Volume 46, Number 2

Maithil Pachchigar [maithil.pachchigar@analog.com] is an applications engineer in in the Instrumentation, Aerospace and Defense business unit at Analog Devices in Wilmington, MA. He joined ADI in 2010 and focuses on the precision ADC product portfolio and customers in the instrumentation, industrial, healthcare, and energy segments. Having worked in the semiconductor industry since 2005, he has published numerous technical articles. He received an M.S.E.E. degree from San Jose State University in 2006 and an M.B.A. degree from Silicon Valley University in 2010.

Maithil Pachchigar

Also by this Author:

RF-to-Bits Solution Offers Precise Phase and Magnitude Data for Material Analysis

Volume 48, Number 4

Analyzing, Optimizing, and Eliminating Integer Boundary Spurs in Phase–Locked Loops with VCOs at up to 13.6 GHz

By Robert Brennan

A phase-locked loop (PLL) and voltage controlled oscillator (VCO) outputs an RF signal at a certain frequency, and ideally this signal would be the only signal present at the output. In reality, there are unwanted spurious signals and phase noise at the output. This article discusses the simulation and elimination of one of the more troublesome spurious signals—integer boundary spurs.

PLL and VCO combinations (PLL/VCOs) that are only capable of operating at integer multiples of the phase frequency detector reference frequency are known as integer-N PLLs. PLL/VCOs capable of much finer frequency steps are known as fractional-N PLLs. Fractional-N PLL/VCOs offer much more flexibility and are more widely used. Fractional-N PLLs achieve this feat by modulating the feedback path in the PLL at the reference rate. While capable of much finer frequency steps than the phase detector reference frequency, fractional-N PLL/VCOs have spurious outputs called integer boundary spurs (IBS). Integer boundary spurs occur at integer (1, 2, 3 … 20, 21 …) multiples of the PLL's phase frequency detector's reference (or comparison) frequency (f_{PFD}). For example, if f_{PFD} = 100 MHz, there will be integer boundary spurs at 100 MHz, 200 MHz, 300 MHz … 2000 MHz, 2100 MHz. In a system where the desired VCO output signal is 2001 MHz, then there will be an IBS at 2000 MHz—this will appear at a 1 MHz offset from the desired signal. Due to effective sampling in the PLL system, this 1 MHz offset IBS is aliased to both sides of the desired signal. Therefore, when the desired output is 2001 MHz, spurious signals will be present at 2000 MHz and 2002 MHz.

Integer boundary spurs are undesirable for two main reasons:

- If they are at low frequency offsets from the carrier (the desired signal), then the IBS power contributes to integrated phase noise.
- If they are at large frequency offsets from the carrier, then the IBS will modulate/demodulate adjacent channels to the desired channel and result in distortion in system.

In some systems, high integer boundary spurs render some output channels unusable. If a system has 1000 channels in a certain spectrum bandwidth, and 10% of the channels have spurious signals above a certain power level, those 100 channels may be unusable. In protocols where spectrum bandwidth costs a lot of money, it is wasteful if 10% of the available channels cannot be used.

Integer boundary spurs are strongest when the integer boundary falls within the PLL bandwidth from the desired output frequency. That is, if the output frequency is 2000.01 MHz and the loop bandwidth is 50 kHz, the IBS will be strongest. As the output frequency moves away from the integer boundary, the power of the IBS reduces in a calculable and repeatable manner. Analog Devices' new, free simulator—ADIsimFrequencyPlanner™—uses this predictable behavior to accurately simulate integer boundary spur power (and much more).

Figure 1 shows the worst-case integer boundary spur power at each output frequency from 1900 MHz to 2150 MHz (1 MHz steps). It can be seen that, at 2001 MHz, the worst-case IBS power is –70 dBc (70 dB below the carrier power). At 2000 MHz, there is no IBS because the output frequency falls on an integer boundary. The IBS power reduces as the carrier moves away from the integer boundary until the carrier starts getting close to the next integer boundary.

The spurious signals seen halfway between the integer boundaries (2049 MHz and 2051 MHz in Figure 1) are second-order integer boundary spurs. Second-order integer boundary spurs occur halfway between integer boundaries. Typically, second-order IBS are 10 dB to 20 dB lower than first-order IBS. ADIsimFrequencyPlanner simulates first-, second-, third-, fourth-, and fifth-order integer boundary spurs.
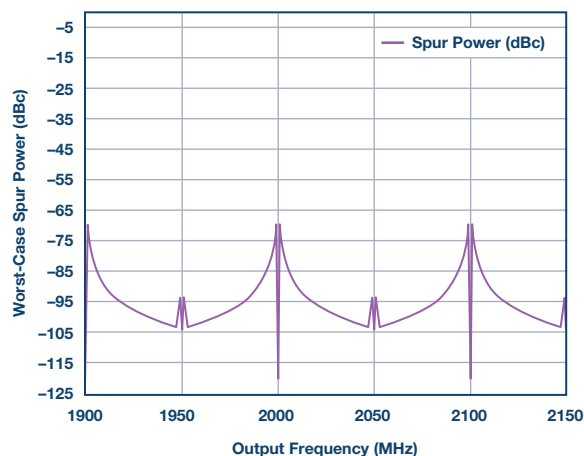

Figure 1. Worst–case integer boundary spur power at each output frequency from 1900 MHz to 2150 MHz (1 MHz steps; 100 kHz loop bandwidth; HMC830).

Suppose a certain modulation scheme states that channels with integer boundary spur power above –80 dBc are unusable; then about 10% of the channels in Figure 1 are no longer available. To overcome this problem, ADIsimFrequencyPlanner can optimize the PLL/VCO configuration to reduce and, in most cases, eliminate integer boundary spurs. Recall that the integer boundary spurs occur at integer multiples of the PFD frequency, and that they are strongest when near the carrier frequency. If the PFD frequency can be changed so that the integer multiple of the PFD frequency falls at a large enough offset from the carrier frequency, then the IBS power will be reduced to a nonproblematic level. This is what the algorithm in ADIsimFrequencyPlanner does—while taking into account the relative powers of the first- through fifth-order integer boundary spurs, ADIsimFrequencyPlanner finds the optimum solution that results in the lowest possible integer boundary spurs at the VCO output.

How can the PFD frequency be changed? Traditionally, in a PLL/VCO system, the PFD frequency is kept fixed. However,

by making the most of a programmable clock distribution source, the PLL reference input divider, and the PLL fractional-N modulator architecture, it is now easy to change the PFD frequency for each output channel.

In the recommended solution, the new HMC7044 clock generation and distribution chip is used. The HMC7044 has 14 ultralow noise outputs; each of the 14 outputs has a programmable divider. By connecting one of these outputs to the PLL reference input, and then programming the output divider as needed, an array of reference frequencies becomes available to the PLL.

The HMC7044 is a clock distribution system applicable to applications that use numerous synchronized clocks for ADCs, DACs, and other system components. Simpler applications that don't require as many outputs can use a simpler alternative, such as the HMC832 or ADF4351—both are integrated PLL and VCO chips.

Then, at the PLL reference input, the reference input divider (R divider) can be programmed as needed to divide the array of available reference frequencies to a larger array of PFD frequencies (the PFD frequency is the frequency at the output of the R divider). Thanks to the high order fractional-N modulator in the PLL, a change to the PFD frequency does not cause a problem in achieving the desired output frequency. Also, the programmable charge pump current of the PLL can be used to compensate for any change in the PFD frequency and therefore maintain a constant-loop bandwidth.
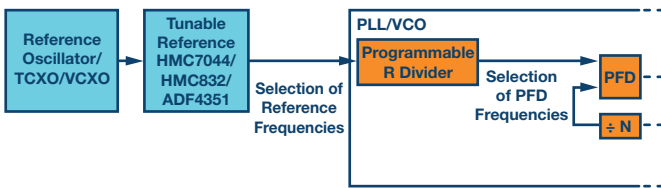

Figure 2. Block diagram showing a selection of PFD frequencies.

Example:

f_{PFD} (MHz)	N	I_{CP}	$f_{PFD} \times N = RF_{OUT}$ (MHz)	Comments
100	20.01	2.08	2001	IBS at ±1 MHz
75	26.68	1.88	2001	IBS at ±24 MHz

where:
I_{CP} = programmable charge pump current
f_{PFD} = PLL PFD frequency;
N = PLL fractional-N value;
RF_{OUT} = VCO output frequency/carrier frequency/ desired signal

The programmable charge pump current changes inversely with the PFD frequency—as PFD frequency increases, the charge pump current must decrease. This serves to keep the loop filter dynamic constant.

When using ADIsimFrequencyPlanner, the user inputs the required output frequency range, step size, PFD frequency and reference frequency constraints, and loop filter parameters. The user also selects the available clock generator output dividers and PLL reference input dividers. ADIsimFrequencyPlanner then steps through each desired frequency step and calculates the optimum PFD frequency from the array of available PFD frequencies. ADIsimFrequencyPlanner then returns the required divider settings and charge pump current to the user. The data can be easily exported to a lookup table that

the end application's firmware can read and then program the HMC7044 and the PLL/VCO accordingly. ADIsimFrequencyPlanner also generates a series of plots to show the user what is happening.

In Figure 3, the user has used the same configuration as Figure 1, except this time, the PFD frequency is optimized by changing the HMC7044 output divider and the PLL reference input divider. The unoptimized simulation is also shown in gray for comparison.
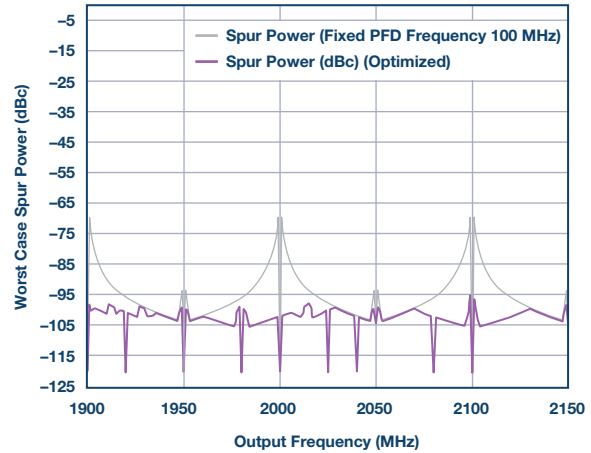

Figure 3. Same output configuration as Figure 1, but now the PFD frequency is optimized.

It can be seen in Figure 3 that across the output range (1900 MHz to 2150 MHz in 1 MHz steps), all integer boundary spurs are now < –95 dBc. This represents a dramatic improvement and makes a very high percentage of the desired outputs all the same excellent quality.

Applying ADIsimFrequencyPlanner to a Wideband VCO

In an experiment to measure the accuracy and effectiveness of ADIsimFrequencyPlanner, several of Analog Devices high performance parts were put together and evaluated in the laboratory. In the experiment, the following parts were used:

- HMC7044 clock generation and distribution:
  - Up to 3.2 GHz output.
  - JESD204B compatible.
  - Ultralow noise (<50 fs jitter, 12 kHz to 20 MHz).
  - –142 dBc/Hz at 800 kHz offset from 983.04 MHz output.
  - 16 programmable outputs.

- ADF5355 integrated PLL and VCO:
  - 55 MHz to 13.6 GHz output.
  - 5 mm × 5 mm LFCSP package.
  - –138 dBc/Hz at 1 MHz offset from a 3.4 GHz output.

- HMC704 ultralow noise PLL:
  - RF input up to 8 GHz.
  - 100 MHz maximum PFD frequency.
  - –233 dBc/Hz normalized phase noise floor.

Although the ADF5355 has an internal PLL, the HMC704 was used to externally lock the ADF5355 VCO. There are two main benefits to this technique:

1. The overall phase noise benefits from the industry-leading VCO phase noise of the ADF5355 and from the industry-leading PLL phase noise of the HMC704.

2. Isolating the VCO and PLL results in less unwanted signal coupling and therefore reduces the power of spurious signals.

ADIsimFrequencyPlanner was used to optimize an output range from 4800 MHz to 6300 MHz in 250 kHz steps (6000 steps). At each step, the optimum divider settings (therefore optimum PFD frequency) and charge pump current was programmed to the HMC7044, ADF5355, and HMC704. Once the parts were programmed to an output step, a spectrum analyzer measured the carrier power and power of the first-order and second-order integer boundary spurs. The spectrum analyzer used a very narrow frequency span and resolution bandwidth—even so, at most channels only noise was measured because the integer boundary spur power was lower than the instrument's noise floor.

The following measurement was taken with the PFD frequency constrained between 60 MHz and 100 MHz. The loop bandwidth and phase margin were 17 kHz and 49.6,° respectively.

Figure 4 shows the measured and simulated results for the HMC7044, ADF5355, and HMC704 solution.

- 6000 output channels were simulated and measured.
- Most integer boundary spurs are simulated around –120 dBc. This is below the noise floor of the spectrum analyzer so only noise was measured.
- Most frequencies have spurs below –100 dBc! A typical requirement is –70 dBc to –80 dBc.
- The only region where the optimization doesn't improve the IBS is less than 2 MHz wide and occurs at 2 × HMC7044 master clock—at this frequency, no combination of dividers can improve the IBS performance. Alternative solutions are offered below.
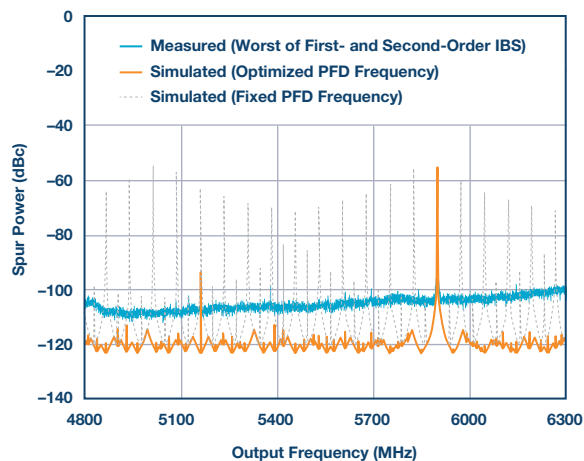

Figure 4. Measured and simulated results for the HMC7044, ADF5355, and HMC704. Note the narrow frequency range where optimization was not possible is correctly simulated by ADIsimFrequencyPlanner. At most other frequencies, the measurement was limited by the noise floor of the spectrum analyzer.

There is only one very narrow range of frequencies where optimizing the PFD frequency does not improve the IBS performance. This frequency range is twice the system master clock (in this case, 2949.12 MHz × 2 = 5898.24 MHz). At this frequency, if the application is capable, it is recommended to shift the carrier frequency to a nearby, cleaner frequency and then shift the baseband frequency in digital (NCO) to compensate. For example, offset the carrier frequency 2 MHz and offset the digital baseband frequency 2 MHz to compensate. Alternatively, if possible in the system, change the master clock frequency to create a clean output frequency.

If the simpler solution mentioned above (using the HMC832 or ADF4351 instead of the HMC7044), then no problem frequencies exist!

From Figure 4, it can be seen that ADIsimFrequencyPlanner:

- Simulates integer boundary spurs accurately.
- Successfully optimizes the reference source and PLL/VCO system for excellent integer boundary spur performance.
- This makes more channels in a range usable and therefore increases value for money in expensive frequency spectrums.
- Simulates wide frequency range systems very quickly. Manually the process can take days or even weeks. The above 6000 step simulation takes less than one minute in ADIsimFrequencyPlanner.

Resources

ADF5355.

ADIsimFrequency Planner.

HMC704.

HMC7044.

Acknowledgments

Robert Brennan

Robert Brennan [robert.brennan@analog.com] joined Analog Devices in 2010 after graduating with a bachelor's in electronic engineering from University of Limerick, Ireland. He worked at the Analog Devices Limerick office for several years before relocating to the U.S. He is currently a senior applications engineer in the RF and Microwave Group based in Massachusetts. His primary focus is on PLLs, VCOs, and integrated PLL/VCOs. Follow Robert on Analog Devices EngineerZone community at https://ez.analog.com/people/rbrennan.

Interleaving ADCs: Unraveling the Mysteries

By Gabriele Manganaro and Dave Robertson

Time interleaving is a technique that allows the use of multiple identical analog-to-digital converters[1] (ADCs) to process regular sample data series at a faster rate than the operating sample rate of each individual data converter. In very simple terms, time interleaving (IL) consists of time multiplexing a parallel array of M identical ADCs, as shown in Figure 1, to achieve a higher net sample rate f_s (with sampling period $T_s = 1/f_s$) even though each ADC in the array is actually sampling (and converting) at the lower rate of f_s/M. So, for example, by interleaving four 10-bit/100 MSPS ADCs one could in principle realize a 10-bit/400 MSPS ADC.

To better understand the principle of IL, in Figure 1 an analog input $V_{IN}(t)$ is sampled by the M ADCs and results in a combined digital output data series D_{OUT}. ADC_1 will sample $V_{IN}(t_0)$ first and begin converting it into an n-bit digital representation. T_s seconds later, ADC_2 will sample $V_{IN}(t_0 + T_s)$ and begin converting it into an n-bit digital representation. Then, T_s seconds later, ADC_3 will sample $V_{IN}(t_0 + 2T_s)$ and so on. After ADC_M has sampled $V_{IN}(t_0 + (M-1) \times T_s)$, the next sampling cycle starts with ADC_1 sampling $V_{IN}(t_0 + M \times T_s)$ and this carousel carries on.

As the n-bit outputs of the ADCs become sequentially available in the same order as just described for the sampling operation, these digital n-bit words are collected by the demultiplexer shown on the right hand side of the same figure. Here the recombined data out sequence $D_{OUT}(t_0 + L)$, $D_{OUT}(t_0 + L + T_s)$, $D_{OUT}(t_0 + L + 2T_s)$, ... is obtained. L stands for the fixed conversion time of each individual ADC and this recombined data sequence is an n-bit data series with sample rate f_s. So, while the individual ADCs, often referred to as the "channels," are n-bit ADCs sampling at f_s/M, the ensemble contained in the box is equivalent to a single n-bit ADC sampling at f_s, and we will refer to that as the time interleaved ADC (distinguishing it from the channels). Basically the input is sliced and separately processed by the ADCs in the array and then consistently reassembled at the output to form the high data rate representation D_{OUT} of the input V_{IN}.
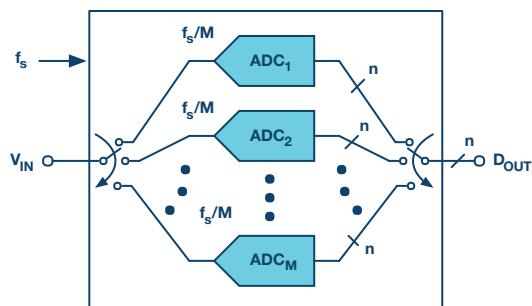
This powerful technique is not free of practical challenges. The key issue manifests itself when the M data streams coming from the channels are digitally assembled together to reconstruct the original input signal V_{IN}. If we look at the spectrum of D_{OUT}, in addition to seeing the digital representation of V_{IN} and the distortion introduced by the analog-to-digital conversion, we will also see additional and substantial spurious content, termed "interleaving spurs" (or IL spurs, in short), IL spurs neither have the signature of polynomial type distortions like higher order signal harmonics (2nd, 3rd, and so on), nor the signature of quantization or DNL errors. IL artifacts can be seen as a form of time-domain fixed pattern noise and are introduced by analog impairments in the channels that, due to the interleaving process, modulate with the sliced converted signals and ultimately show up in the final digitized output D_{OUT}.

Let's begin understanding what might be happening by analyzing a simple example. Consider the case of a two-way interleaved ADC with a sinusoidal input V_{IN} at frequency f_{IN}. Assume that ADC_1 has a gain, G_1 and that ADC_2 has a different gain, G_2. In such a two-way IL ADC, the ADC_1 and ADC_2 will alternate in sampling V_{IN}. So if ADC_1 converts the even samples and ADC_2 converts the odd samples, then all the even data of D_{OUT} has an amplitude set by G_1, while all the odd data of D_{OUT} has an amplitude set by G_2. Then D_{OUT} doesn't only contain V_{IN} along with some polynomial distortion, but it has been subject to the alternate magnification of G_1 and G_2 just as if we were instead amplitude modulating V_{IN} with a square wave at frequency $f_s/2$. That is what will introduce additional spurious content. Specifically, D_{OUT} will include a "gain spur" at frequency $f_s/2 - f_{IN}$ and, unfortunately, this spur's frequency tracks the input f_{IN} and it is located within the first Nyquist band of the interleaved ADC (that is, within $f_s/2$) and there are also aliases of it on all other Nyquist bands. The power/magnitude of this interleaving spur depends on the net *difference* between the two gains G_1 and G_2. In other words, it depends on the gain error mismatch.[2] And, finally, it depends on the magnitude of the input V_{IN} itself.

---

[1] While analog-to-digital converters are discussed here, all the same principles are applicable to the time interleaving of digital-to-analog converters.

[2] Note that it is the gain error mismatch that matters, not its absolute value. Because if both channels have the same gain (error), then G1 = G2. In that case, the two channels are equally scaled up, so the two data streams are recombined into a single D_{OUT} data stream without alternating amplitude (or modulation) and no gain spur is introduced.

Figure 1. An array of M time interleaved n–bit ADCs. The sample rate of each one is f_s/M, the resulting sample rate of the time interleaved ADCs is f_s. An example of clocking scheme for the case of M = 4 is depicted on the lower part of this figure.

If the input isn't a simple sine wave but, as in a real application case, it is a whole band limited signal, then the "gain spur" isn't simply an undesired tone, it is instead a complete scaled image of the band limited input signal itself that shows up within the Nyquist band. This to some extent negates the benefits of the increased bandwidth provided by interleaving.

While in the above example we have considered only the gain error mismatch between the channels, other impairments introduce interleaving spurs too. Offset mismatch (difference between the channels' offsets) introduces tones ("offset spurs") at fixed frequency and with power proportional to the offset mismatch.[3] Sampling time skew occurs when some of the channels sample a bit earlier or later than they should in the intended order. That introduces "timing spurs" that lie at the very same frequency (and add up to the same amplitude) as the gain spurs[4] but with power that is increasingly stronger as f_{IN} grows and as the input amplitude grows. Bandwidth mismatch between the individual channels introduces yet more spurious content at frequencies that depend on f_{IN} and, just like the timing spurs, the spurious power gets progressively stronger with f_{IN} itself, not just with the input amplitude. Again, in all cases, the severity of the spectral degradation of the output isn't dependent on the absolute value of the channels' impairments (offset, gain, timing, band), but on the relative mismatches/differences between them.

While the general technique of time interleaving has been around for several decades, the degree with which the IL spurs could be kept minimal has limited its past applicability to low resolution converters. However, recent advances in the calibration of channel mismatch and in the suppression of the residual IL spurious content is allowing today the realization of fully integrated very high speed 12-, 14-, and 16-bit IL ADCs.

At this point, we need to distinguish between some classes of interleaving. We generally refer to "ping pong" operation in the case of two interleaved channels. We can then distinguish between "lightly interleaved" and "highly interleaved" as we refer to the cases of a reduced number of channels—for example, three channels to four channels—or the case of a large number of channels, say more than four, and often eight of more, respectively.

Ping Pong (Two–Way) Interleaving

When we interleave only two channels to double the net sample rate as shown in the block diagram of Figure 2(a), we term that "ping pong." This is an especially simple case that has some interesting and useful features. In this case, within the 1st Nyquist band of the interleaved ADC, the interleaving spurs are located at dc, at $f_s/2$ and at $f_s/2 - f_{IN}$. So, if the input signal V_{IN} is a narrow-band signal centered at f_{IN}, as depicted in the first Nyquist output spectra of Figure 2(b), the interleaving spurs will consist of an offset spur at dc, another offset mismatch spur at $f_s/2$, and a gain and timing spurious image centered at $f_s/2 - f_{IN}$ that looks like a scaled copy of the input itself.

If the input signal $V_{IN}(f)$ is completely bound between 0 and $f_s/4$, as in Figure 2(b), then the interleaving spurs are not frequency overlapping with the digitized input. In this case, the bad news is that we are only able to digitize in half of the Nyquist band, namely just like if we had a single channel

[3] In general, for M channel interleaving, the offset spurs occur at $f_{OS} = (k/M) f_s$, for k = 0,1,2,... (Manganaro, 2011).

[4] In general, for M channel interleaving, the gain and the timing skew images occur at $f_{GS} = \pm f_{IN} + (k/M) f_s$ with k = 1,2,... (Manganaro, 2011).

clocked at $f_s/2$, though we are still consuming at least twice the power of such a single channel. The interleaving spurious image on the upper end of the Nyquist band can be suppressed by digital filtering after digitization and does not require correcting for analog impairments.

The good news, however, is that since the ping pong ADC is clocked at f_s, the digitized output benefits from a 3 dB processing gain in dynamic range. Moreover, compared to using a single ADC clocked at $f_s/2$, the antialiasing filter design has been relaxed for the ping pong ADC.
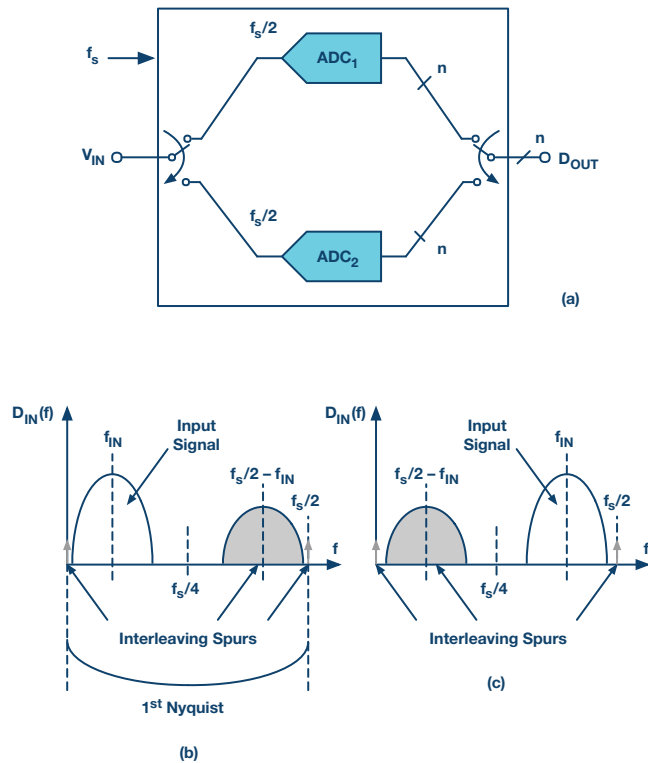

Figure 2. (a) A ping pong scheme, (b) the output spectrum when a narrow–band input signal lies below $f_s/4$, and (c) when the input signal lies between $f_s/4$ and the Nyquist frequency $f_s/2$.

All the same considerations can be repeated if the narrow-band signal is located on the upper half of the first Nyquist band, as shown in Figure 2(c), since the interleaving image spur is moved to the lower half of the Nyquist band. Once again, the gain and timing spur can be digitally suppressed after digitization by filtering.

Finally, the input signal and the interleaving spurs will frequency overlap and the input spectrum gets corrupted by the interleaving image, as soon as the input signal frequency location crosses the $f_s/4$ line. In this case, recovering the desired input signal is not possible and the ping pong scheme is not usable. Unless, of course, the channel-to-channel matching is sufficiently close to make the interleaving spurious content acceptably low for the application or if calibration is employed to reduce the causes leading to IL images.

In summary, frequency planning and some digital filtering allow recovering the narrow-band digitized input in a ping pong scheme even in the presence of channel mismatch. While the converter power consumption roughly doubles compared to the case of using a single ADC clocked at $f_s/2$, the ping pong scheme provides a 3 dB processing gain and relaxes antialiasing requirements.

An example of a ping pong without any correction for channel mismatch and its resulting interleaving spurs is shown in

Figure 3. In this case, the two ADCs of the dual 14-bit/1 GSPS ADC AD9680 sample at alternate times a single sine wave, hence returning a single combined output data stream at 2 GSPS. When we look at the 1st Nyquist band of the output spectrum of this ping pong scheme—that is between dc and 1 GHz—we can see the input tone, which is the strong tone on the left at f_{IN} = 400 MHz, we can also see the strong gain/timing mismatch spur at $f_s/2 - f_{IN}$ = 2G/2 – 400 M = 600 MHz. We also see a number of other tones due to the two channels' own distortion as well as other impairments, but these are all below the –90 dB line.
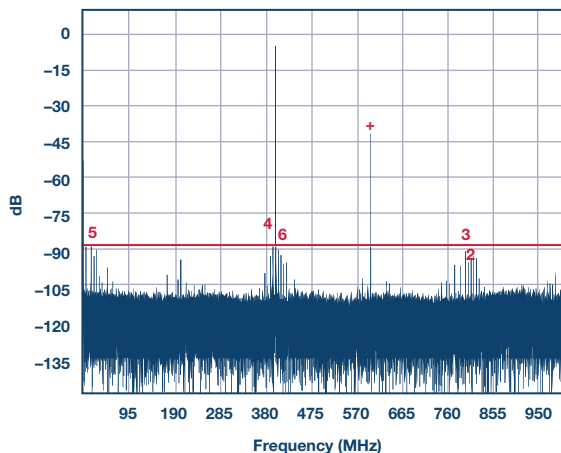
Figure 3. Spectrum of the combined 2GSPS output data of a ping pong scheme obtained by clocking the two ADCs of the AD9680 with a 1 GSPS clock but with 180° sampling phase shift.

Higher Order Interleaving

When we have more than two channels, frequency planning as described above is not very practical or attractive. The location of the interleaving spurs cannot be confined to a fraction of the Nyquist band. For example, consider the case of a four-way interleaved ADC as shown in Figure 4(a). In this case, the offset mismatches give rise to tones at dc, $f_s/4$ and $f_s/2$. While the gain and timing interleaving images are located at $f_s/4 - f_{IN}$, $f_s/4 + f_{IN}$ and $f_s/2 - f_{IN}$. An example of the spectrum of the interleaved ADC's output is shown in Figure 4(b). It can be clearly seen that, unless the input is within a bandwidth of less than $f_s/8$, no matter where we place f_{IN}, the input will overlap with some of the interleaving spurs and, if the input is a very narrow-band signal, we shouldn't try to digitize it with a wideband interleaved ADC.

In a case like this, we need to minimize the IL spurious power to obtain full Nyquist and a cleaner spectrum. In order to do that, calibration techniques are used to compensate for the mismatch between the channels. As the effect of the mismatches is corrected, the power of the resulting IL spurs decreases. Both the SFDR and the SNR benefit from the reduction of this spurious power.

Compensation approaches are limited by the accuracy with which the mismatches can be measured and ultimately corrected. To further suppress the residual spurs beyond the level achieved via calibration, it is possible to intermittently and randomly shuffle the order with which the channels sample the input. In doing so, the previously discussed modulation effects of the converted input signal due to the uncalibrated mismatches turn from a fixed pattern noise to pseudorandom. As a result, IL tones and undesired periodic patterns turn into pseudorandom noise-like content that adds with the converter quantization noise floor and leads to the disappearance

or, at least, to spreading of the undesired spurious images and tones. In this case, the power associated with the IL spurious content adds to the power of the noise floor. Hence, while distortion improves, SNR can degrade by the amount of IL spurious power added to the noise. SNDR (SINAD) is essentially unchanged as it combines both distortion and noise and randomization; it simply moves the IL contribution from a component (distortion) to the other (noise).
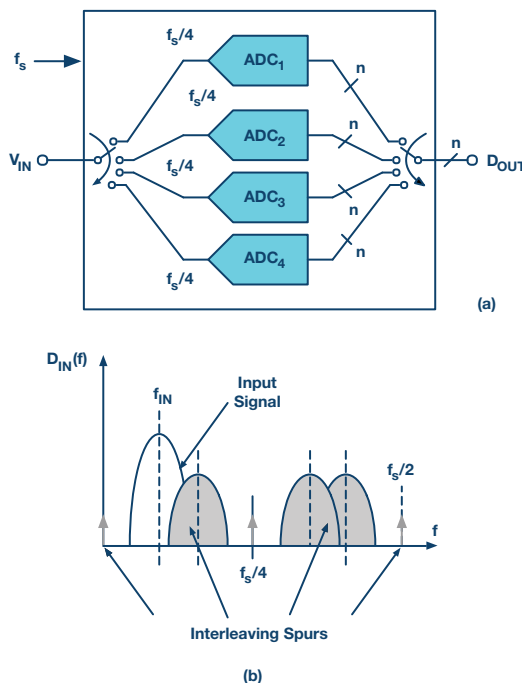

Figure 4. (a) A four-way interleaved ADC and (b) the corresponding 1st Nyquist output spectrum showing the interleaving spurs.

Let us consider some examples of interleaved ADCs. The AD9625 is a 12-bit/2.5 GSPS three-way interleaved ADC. The mismatches between the three channels are calibrated in order to minimize the interleaving spurs. An example of its output spectrum with an input close to 1 GHz is depicted in Figure 5(a). In this spectrum, besides the ~1 GHz input tone, it is possible to see the channels' 2nd and 3rd harmonic distortion near 500 MHz and the 4th harmonic distortion near the fundamental. The interleaving mismatch calibration substantially minimizes the power of the interleaving spurs and a large set of additional residual small spurious tones is visible across the entire spectrum.

In order to further reduce such residual spurious content, channel randomization is introduced. A fourth calibrated channel is added and the four channels are then three-way interleaved in randomly changing order by intermittently swapping one of the interleaved channels with the fourth one. One can liken that to a juggler playing with three Skittles up in the air while a fourth one is swapped in every so often. By doing so, the residual interleaving spurious power is randomized and spread out over the noise floor. As shown in Figure 5(b), after channel randomization, the interleaving spurs have nearly disappeared, while the power of the noise has marginally increased, hence degrading the SNR by 2 dB. Note, of course, that while the second spectrum shown in Figure 5(b) is considerably cleaner of distortion tones, the shuffling cannot affect the 2nd, 3rd, and 4th harmonic since these aren't interleaving spurs.
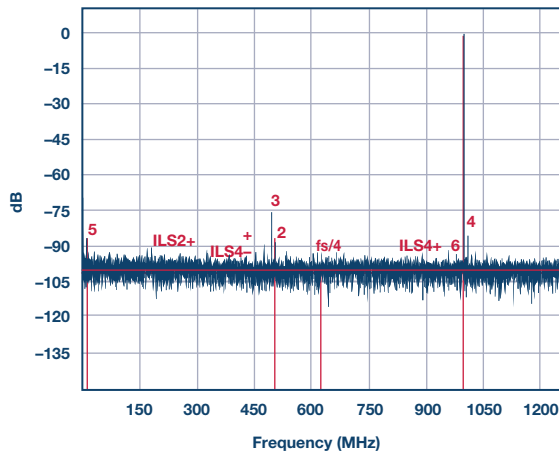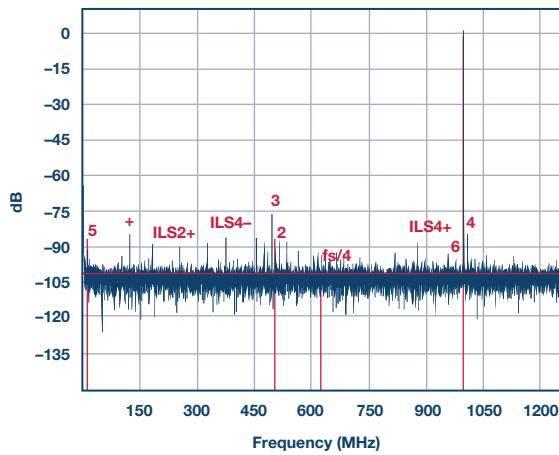
(a)

(b)

Figure 5. The output spectrum of the AD9625, clocked at 2.5 GSPS and with an input tone close to 1 GHz. (a) Sequential three–way interleaving; SNR = 60 dBFS, the SFDR = 72 dBc is limited by the third harmonic, near 500 MHz; however, a number of interleaving spurs are visible all across the spectrum. (b) Three–way inter–leaving with random channel shuffling; SNR = 58 dBFS, while the SFDR = 72 dBc is still set by the third harmonic, all the interleav–ing spurs have been eliminated by spreading their power over the noise floor.

Another example of an interleaved ADC using channel ran-domization is the one shown in the spectra of Figure 6. This is the case of the four-way interleaved 16-bit/310 MSPS ADC AD9652. In the case shown in Figure 6 the four channels are sequentially interleaved in a fixed sequence while no effort is made to calibrate them to reduce channel mismatch. The spectrum shows clearly the interleaving spurs at the predicted frequency locations and their large power is far greater than the 2nd and 3rd harmonics and limits the spurious-free dynamic range to only 57 dBc.

However, if the same ADC is foreground calibrated to reduce the channel mismatch, the power of the interleaving spurs is substantially reduced as shown in Figure 7. Similar to the case of the previous example, the channel harmonic distor-tion isn't affected, however the interleaving spurs are greatly reduced in power through channel mismatch calibration.

Lastly, the spectral purity in Figure 7 can be further improved by randomizing the channel order as shown in Figure 8. In this case the randomization uses a proprietary technique that while intermittently scrambling the order of the four channels doesn't require a spare (5th) channel to be added, hence saving its associated power. It can be seen in Figure 8 that, after ran-domization, only regular harmonic distortion is left on the resulting spectrum.
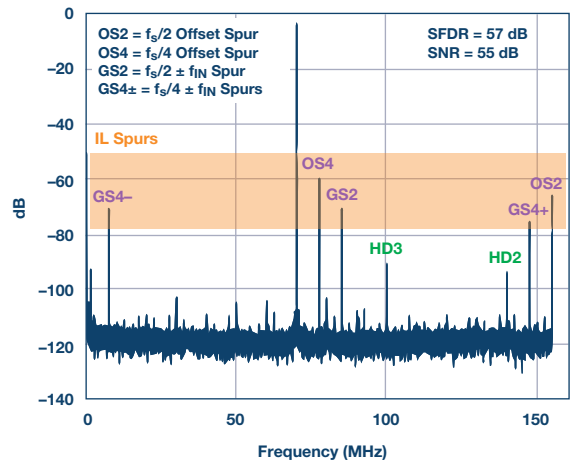
Figure 6. The output spectrum of the AD9652, clocked at $f_s = 310$ MHz and with a sinusoidal input at $f_{IN} \sim 70$ MHz. In this case, no channel calibration and randomization is applied. The 2nd (HD2), and aliased 3rd (HD3) harmonics are visible at ~140 MHz and ~100 MHz, respectively. Interleaving (IL) spurs are visible as well. These are the offset tones at dc, $f_s/2$ (OS2 in the graph) and $f_s/4$ (OS4 in the graph). Moreover the gain (/timing) spurs can be found at $f_s/2 - f_{IN}$ (GS2 in the graph), $f_s/4 + f_{IN}$ (GS4+ in the graph), and $f_s/4 - f_{IN}$ (GS4– in the graph). The SNR quote in this graph is artificially poor due to the fact that some of the spurious content has been lumped with the noise power.
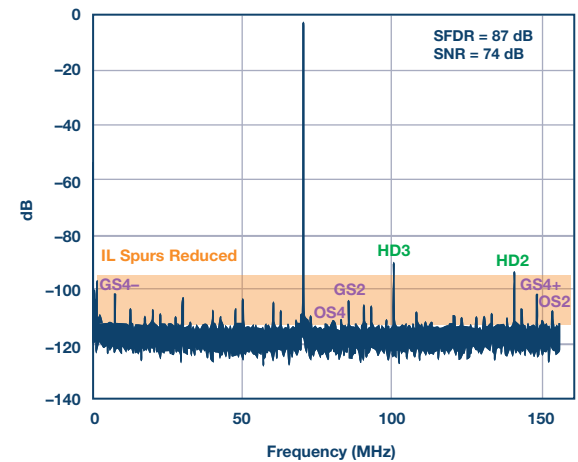

Figure 7. The output spectrum for the same AD9652, with the same input but after calibrating the four channels to reduce their mismatch. Comparing with Figure 6, while the 2nd and 3rd harmonics are unaffected, the interleaving spurs' power has substantially been reduced and the SFDR has improved by 30 dB going from 57 dBc to 87 dBc.
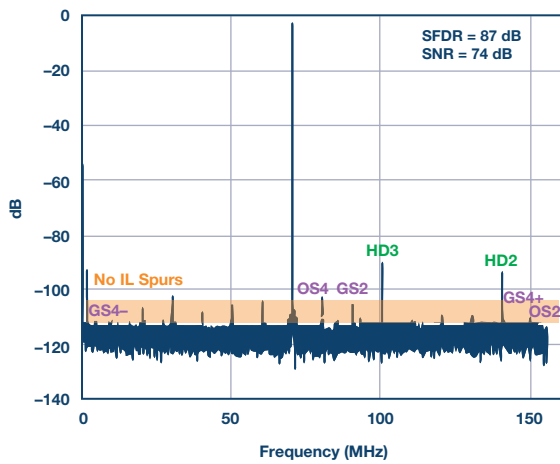
Figure 8. The output spectrum for the previous case once the randomization of the interleaving order is turned on. Randomizing the residual interleaving spurs distributes their power over the noise floor and the corresponding peaks disappear. Only regular harmonic distortion is left in sight. The SNR is nearly unaffected as the distributed spurious power from the interleaving tones is rather negligible after the mismatch calibration.

Conclusion

Time interleaving is a powerful technique to increase the bandwidth of data converters. Recent advances in mismatch compensation as well as cancellation of residual spurious content via randomization techniques have allowed the fully integrated realization of very high speed 12-, 14-, and 16-bit interleaved ADCs.

In the case in which the input signal is band limited, such as, for example, in a number of communication applications, a ping pong (two-way) interleaving approach allows allocating the undesired interleaving spurs away from the input band of interest via frequency planning. The spurious content can then be digitally filtered. While this approach consumes roughly twice the power compared to a noninterleaved ADC at half the IL sample rate required to capture the same spurious-free input bandwidth, on the other hand it both increases the dynamic range by 3 dB via processing gain and it also relaxes the roll-off of the antialiasing and roofing filters that precede the ADC thanks to the higher IL sample rate.

When the full input band of the IL converter is required to capture a wideband input signal, a higher order interleaving converter is appropriate. In this case, calibration and random shuffling allow interleaving distortion and spurious content compensation and cancellation.

Acknowledgments

References

Beavers, Ian. "Gigasample ADCs Run Fast to Solve New Challenges." Analog Devices, 2014.

Black, William and David Hodges. "Time Interleaved Converter Arrays." *IEEE Journal of Solid-State Circuit*, Vol. SC-15, No. 6, 1980.

Bosworth, Duncan. "GSPS Data Converters to the Rescue for Electronics Surveillance and Warfare Systems!" Analog Devices, 2014.

Elbornsson, Jonas, Fredrik Gustafsson, and Jan-Erik Eklund. "Analysis of Mismatch Effects in a Randomly Interleaved A/D Converter System." *IEEE Transactions on Circuits and Systems*, Vol. 52, No. 3, 2005.

Harris, Jonathan. "Further into the Alphabet with Interleaved ADCs." *EDN Network*, 2013.

Harris, Jonathan. "The ABCs of Interleaved ADCs." *EDN Network*, 2013.

Manganaro, Gabriele. *Advanced Data Converters.* Cambridge, UK: Cambridge University Press, 2011.

Gabriele Manganaro

Gabriele Manganaro [gabriele.manganaro@analog.com] holds a Dr.Eng. and a Ph.D. in electronics from the University of Catania, Italy. Starting in 1994, he did research with ST Microelectronics and at Texas A&M University. He worked in data converter IC design at Texas Instruments, Engim Inc, and as Design Director at National Semiconductor. Since 2010, he has been an engineering director for High Speed Converters at Analog Devices. He served in the technical subcommittee for Data Converters of the ISSCC for seven consecutive years. He was Associate Editor, then Deputy Editor in Chief and Finally Editor in Chief for IEEE Transactions on Circuits and Systems—Part I. He authored/coauthored 60 papers, three books, and holds 13 US patents with more pending. He is a Senior Member IEEE (since 2003) and a Fellow of the IET (since 2009).
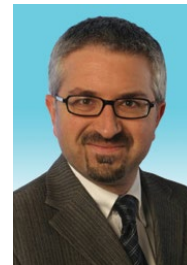
David H. Robertson

David H. Robertson [david.robertson@analog.com] has been with the Data Converter group of Analog Devices since 1985. He has worked on a wide variety of high speed digital-to-analog and analog-to-digital converters on complementary bipolar, BiCMOS, and CMOS processes. He has held positions as a product engineer, design engineer, and product line director, working with product development teams in the US, Ireland, Korea, Japan, and China. Dave is presently the Product and Technology Director for ADI's High Speed Converter group.

Dave holds 15 patents on converter and mixed-signal circuits, has participated in two "best panel" International Solid-State Circuits Conference evening panel sessions, and was coauthor of the paper that received the IEEE Journal of Solid-State Circuits 1997 Best Paper Award. He served on the ISSCC technical program committee from 2000 through 2008, chairing the Analog and Data Converter subcommittees from 2002 to 2008.

Zero–Drift Amplifiers: Now Easy to Use in High Precision Circuits

By Vicky Wong and Yoshinori Kusuda

A zero-drift amplifier, as the name suggests, is an amplifier with offset voltage drift very close to zero. It uses auto-zero or chopping technology, or a combination of both, to continuously self-correct for dc errors over time and temperature. This enables the amplifier to achieve microvolt-level offsets and extremely low offset drifts. Therefore, it is uniquely suited to be used in signal conditioning circuits with high gain and precision performance. For example, a sensor (such as a temperature, pressure, or load cell sensor) typically produces a low level output voltage and hence requires an amplifier to amplify its output without introducing additional errors. Zero-drift amplifiers, designed for ultralow offset voltage and drift, high common-mode rejection, high power supply rejection, and reduced $1/f$ noise, are an ideal choice to achieve a high level of resolution in a demanding system application, such as sensing, with a long product life cycle.

Basic Architecture of a Zero–Drift Amplifier

Figure 1 shows the circuit diagram of a basic chopper amplifier in unity gain configuration. The dc gain path consists of an input chopping switch network ($CHOP_{IN}$), a first transconductance amplifier (Gm1), an output chopping switch network ($CHOP_{OUT}$), a second transconductance amplifier (Gm2) and frequency compensation capacitors (C1 and C2). CHOP and CHOP' are controlled by a clock generator and function to correct unwanted amplifier dc offset voltage (V_{OS}).

Figure 2 shows the associated timing diagram and expected output voltage (V_{OUT}). When the CHOP clock signal is high (A phase), amplifier Gm1's differential input and output are connected to the signal path with no inversion. This results in a positive output voltage, V_{OUT}, due to the presence of V_{OS}. When the CHOP' clock signal is high (B phase), Gm1's input and output are connected to the signal path with inversion, resulting in a negative output voltage due to V_{OS}. The positive and negative output voltages from Gm1 result in an output voltage equal to $\pm V_{OS}$. This chopping concept in the time domain is similar to modulation in the frequency domain. In other words, the offset voltage of Gm1 is up-modulated by $CHOP_{OUT}$ to the chopping frequency. On the other hand, the input signal is chopped twice by $CHOP_{IN}$ and $CHOP_{OUT}$. This is equivalent to the input signal being up-modulated and then down-modulated to its original frequency. Hence, the input signal gets through to the output with no inversion.

The positive and negative output voltages ($\pm V_{OS}$) from Gm1 appear as voltage ripples at V_{OUT} (Figure 2). In addition, the CHOP and CHOP' clocks are coupled to the differential input pins through parasitic capacitances associated with the switches. When the clocks change state, charges are injected into the differential input pins. These charge injections are translated into output voltage glitches via the finite input source impedances. The magnitude and shape of the glitches depend on the amount and matching of the input source impedances and the charge injections at the differential input pins. These output ripples and glitches introduce switching artifacts that appear as increases in noise spectrum at the chopping frequency and its multiple integer frequencies. Also, the magnitude and frequencies of the switching artifacts differ for each zero-drift amplifier and from unit to unit. In this article, the term chopping and switching frequency are used interchangeably.
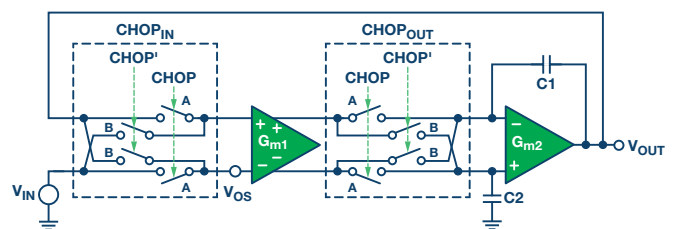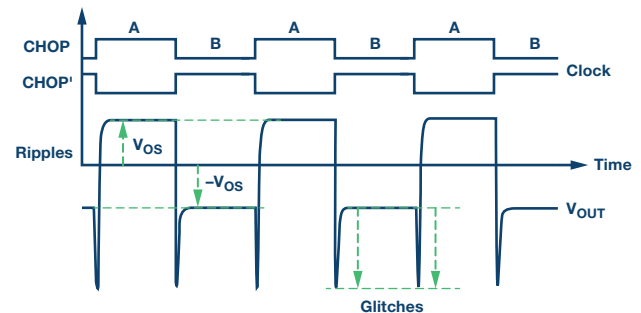

Figure 1. Chopping architecture.

Figure 2. Chopping timing diagram.

Switching Artifacts as Shown on a Data Sheet

Traditionally, zero-drift amplifiers have fairly large broadband noise and low switching frequencies, ranging from a few kilohertz to a few tens of kilohertz. This limits their usage to dc and sub-100 Hz applications so that the switching frequency remains out of the signal bandwidth of interest. For applications requiring high precision and low drift at a higher bandwidth, it is important to use a zero-drift amplifier with higher switching frequency. As a matter of fact, the switching frequency is sometimes viewed as the figure of merit for zero-drift amplifiers. With advanced design architectures, newer zero-drift amplifiers are designed to have smaller switching artifacts at much higher frequencies. For example, in addition to chopping the offset voltage at 4.8 MHz, the ADA4522-2, a high voltage, dual, zero-drift amplifier, uses a patented offset and ripple correction loop circuity to minimize switching artifacts. The correction loop operates at 800 kHz and functions to null out the offset voltage, $\pm V_{OS}$ (as shown in Figure 2). Reducing $\pm V_{OS}$ to 1% of its original value provides a 40 dB improvement in the switching artifact. This reduces the system designer's effort to achieve targeted system level precision.

The easiest way to detect the switching artifact is by observing the amplifier's voltage noise density spectrum. Figure 3 shows the input referred voltage noise density graph of the ADA4522-2. Note that Channel B exhibits an increase in noise spectrum at its switching frequency of 800 kHz. This increase in noise spectrum, as described in the earlier part of this article, is the byproduct of the charge injection mismatch. Since the mismatch is part-to-part and channel-to-channel dependent, the magnitude of noise spikes are different and not all units exhibit the noise spike. As an example, Channel A of the same unit does not exhibit any noise spikes at the switching frequency of 800 kHz. The switching frequencies could also differ up to a factor of 10% to 20% from unit to unit due to the on-chip clock oscillator frequency variation.
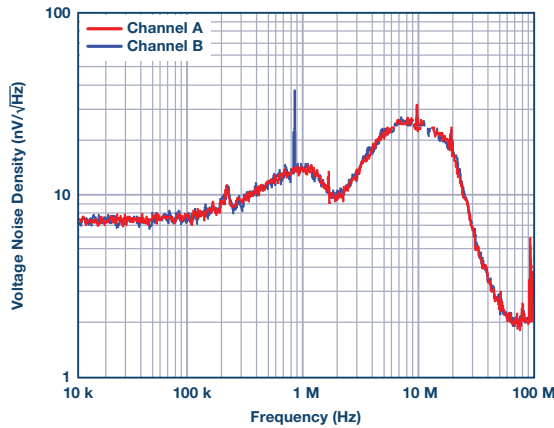

Figure 3. ADA4522–2 voltage noise density.

Noise Comparison Between Different Zero–Drift Amplifiers

Figure 4 shows the input referred voltage noise density of three different leading edge high voltage zero-drift amplifiers. Note that all three zero-drift amplifiers tested exhibit some sort of switching artifacts. Some of the switching artifacts also repeat at its multiple integer frequencies. These switching artifacts could be significant and could introduce errors in a circuit design. Hence, it is important to understand their impact on a circuit and find ways to mitigate the effect. If the amplifier has a closed-loop frequency that is higher than the switching frequency, this increase in noise spectrum will be integrated across the entire bandwidth and be reflected at the output. Not only that, this input-referred voltage noise will be gained up by the amplifier noise gain. For instance, assume that the amplifier is configured in a gain of 100, the effective output referred voltage noise density would also increase by a factor of 100.
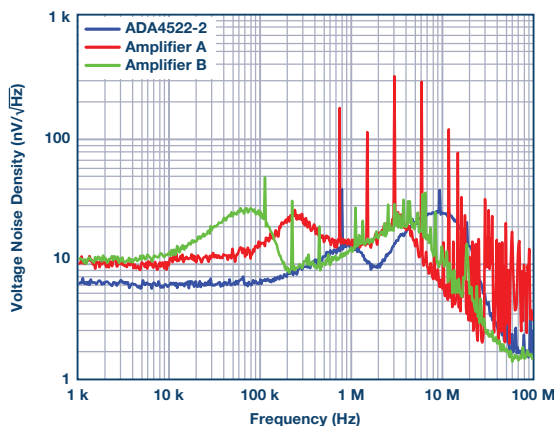

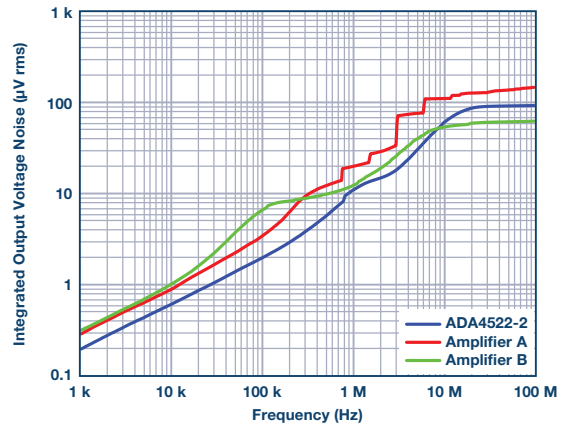Figure 4. Voltage noise density of different zero–drift amplifiers.

Figure 5. Integrated output voltage noise.

Total rms noise that is integrated at the output of an amplifier depends on the amplifier's bandwidth. Output voltage noise rolls off with the available bandwidth; therefore, the higher the gain or the higher the bandwidth, the higher the amplitude of the output amplifier noise. Figure 5 shows a graph of integrated output voltage noise vs. frequency. This is a helpful graph for understanding the total integrated noise with respect to a frequency. As an example, if the amplifier's bandwidth is limited to 100 kHz by means of filtering, the total output noise due to the inherent amplifier voltage noise can be read off the graph and will be as follows:

Table 1. Output Integrated Noise

Amplifier	Output Noise (µV rms)	Peak–to–Peak Output Noise (µV p–p)
ADA4522-2	1.91	12.61
Amplifier A	3.33	21.98
Amplifier B	6.40	42.24

Using a common multiplier (called crest factor) to convert rms voltage to peak-to-peak voltage, a peak-to-peak noise estimation is shown in the third column of Table 1. In a 5 V system, the ADA4522-2 would provide 18.6 bits of peak-to-peak resolution, whereas Amplifier B provides 16.8 bits of peak-to-peak resolution. Having a lower total integrated output noise is always desirable as it increases the signal-to-noise ratio and enables a higher resolution to the entire system.

Another interesting thing to note about Figure 5 is that the integrated noise increases with a step-like function at the noise spike frequencies. The noise spikes (with increased noise energies), albeit narrow, add significantly to the total output integrated noise.

Switching Artifacts in the Time Domain

Often times, the switching artifacts can be clearly seen in the voltage noise density spectrum in the frequency domain. To understand the time-based behavior of the switching artifact, one can configure the amplifier in a buffer configuration with the noninverting pin grounded and directly monitor

the output with an oscilloscope. Figure 6 shows the typical output of two zero-drift amplifiers. Note that Amplifier A exhibits output voltage spikes in various amplitudes. The spikes repeat themselves every 0.66 µs This matches the noise spikes that are seen at 1.51 MHz in Figure 4. On the other hand, the ADA4522-2 does not exhibit any switching artifact in the time domain (blue graph). In other words, the noise spikes that exist are below the noise floor of the measurement system and cannot be detected. This allows designers to use the ADA4522-2 in applications such as driving an ADC with confidence that noise spikes will not be an issue.
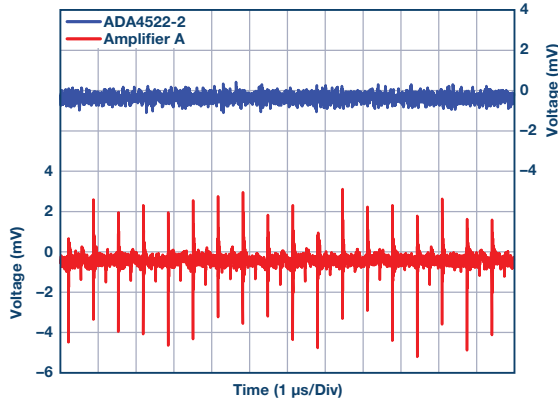

Figure 6. Output voltage noise in the time domain.

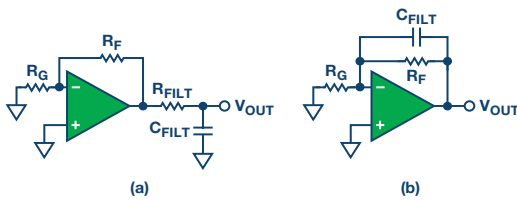## Filters to Mitigate Switching Artifact

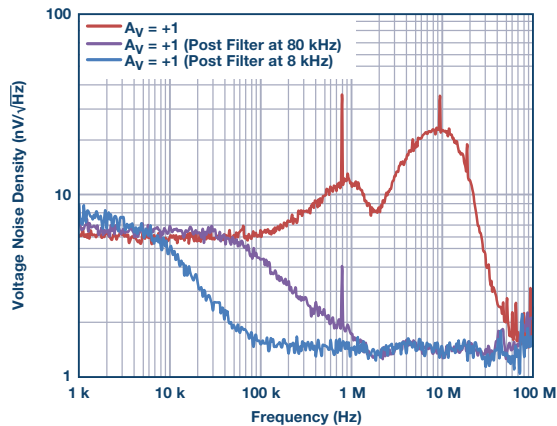*Figure 7. Zero–drift amplifier with filter setup.*

Figure 8. Voltage noise density of a unity–gain zero–drift amplifier with post filter.

To reduce the impact of the switching artifacts, there are a couple of methods that can be implemented. These methods ultimately lead to limiting the amplifier bandwidth such that it is less than the switching frequency. Using a filter is an effective way to suppress the noise spikes. The easiest design is to place a resistor-capacitor network at the amplifier output to create a low pass filter (Figure 7A). Figure 8 shows the voltage noise density of a zero-drift amplifier with a post filter designed at one or two decades below the switching frequency. The noise spike at 800 kHz reduces from 36 nV/$\sqrt{\text{Hz}}$ (no post filter) to 4.1 nV/$\sqrt{\text{Hz}}$ (post filter at 80 kHz),

which is below the amplifier's low frequency broadband noise level. With a post filter positioned two decades below the switching frequency (post filter at 8 kHz), the noise spike is no longer visible and the ADA4522-2 looks like any other traditional amplifier.

Some applications might not tolerate having an RC network at the output of the amplifier. Amplifier output current flowing through the filter resistor creates a voltage offset that introduces output error. In this case, one can opt to filter the noise spikes by placing a feedback capacitor across the feedback loop (Figure 7(b)). Figure 9 shows the output voltage noise density of an amplifier configured in gain of 10 with no filtering vs. having a post filter or a feedback filter positioned a decade below the switching frequency. The post filter configuration is more effective as a low-pass filter than the feedback capacitor.
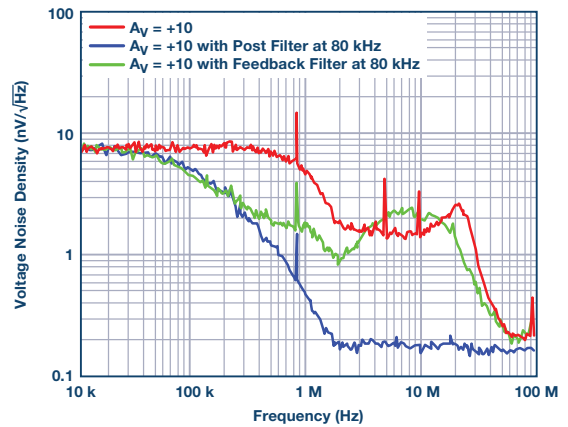

Figure 9. Switching artifact reduces with filters.

Using Zero–Drift Amplifiers in High Gain Configura–tion Helps

A lot of designers have used zero-drift amplifiers, but have not observed any switching artifacts in their system. One reason could be due to the configuration of the amplifier. Zero-drift amplifiers have low drift and offset, and are most often used to signal condition a low level amplitude sensor signal in a high gain configuration of, for example, a gain of 100 to 1000. Using the amplifier in a high gain configuration has the same effect as placing a low-pass filter on the amplifier. As gain increases, bandwidth decreases. Figure 10 illustrates how having a high gain configuration mitigates the switching effect. With a closed-loop gain of 100, the switching artifact can hardly be seen on the noise plots.
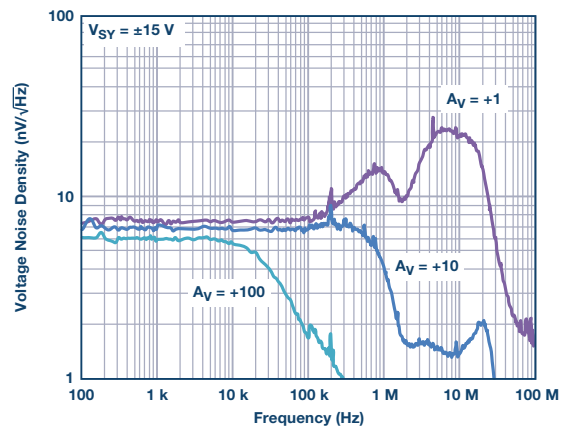

Figure 10. Amplifier bandwidth roll–off with gain.

The Benefits of ADA4522–2 as a Zero–Drift Amplifier

Analog Devices' newest zero-drift operational amplifier, the ADA4522-2, employs a patented and innovative circuit topology to achieve a high switching frequency and to minimize the switching artifacts in comparison to its predecessors. With a unity-gain bandwidth at 3 MHz and a switching frequency at 800 kHz and 4.8 MHz, a gain configuration of 40 is sufficient to filter the switching artifacts, and eliminates the need for external low-pass filtering. Its low offset voltage drift of 22 nV/°C maximum, low noise at 5.8 nV/$\sqrt{\text{Hz}}$ (gain of 100 configuration), low input bias current at 150 pA maximum, high common-mode rejection, and power supply rejection make it an ideal choice for precision applications such as weigh scale, current sensing, temperature sensor front ends, load cell and bridge transducers, and many more drift-critical applications.

Conclusions

Zero-drift amplifiers feature very low offset voltage and drift and are an ideal choice for applications requiring precision amplification of low-level signals. Here are a couple of insights when using one.

All zero-drift amplifiers exhibit some sort of switching artifacts and this can most commonly be detected in the voltage noise density plots.

The magnitude of the switching artifact differs from unit to unit.

The switching frequency could differ from unit to unit up to a factor of 20%.

Switching artifacts can be detected in the frequency and the time domain. Depending on the application, they could present errors.

Zero-drift amplifiers are often used in a high gain configuration, where bandwidth is reduced and thus many times, switching artifacts do not pose an issue.

It is important to mitigate the switching artifacts to reduce the amount of output error. Apply a low-pass filter (RC post filter or feedback capacitor) to roll off the amplifier's bandwidth before the switching frequency to suppress the artifacts.

A high switching frequency simplifies filter requirements for a wide, useful, and artifact-free bandwidth.

Acknowledgments

Vicky Wong

Vicky Wong [vicky.wong@analog.com] is an applications engineer at Analog Devices. She has been with Analog Devices since 2008 and is responsible for precision amplifiers and voltage references. She holds a B.S.E.E. and an M.S.E.E. from the University of Illinois at Urbana–Champaign.
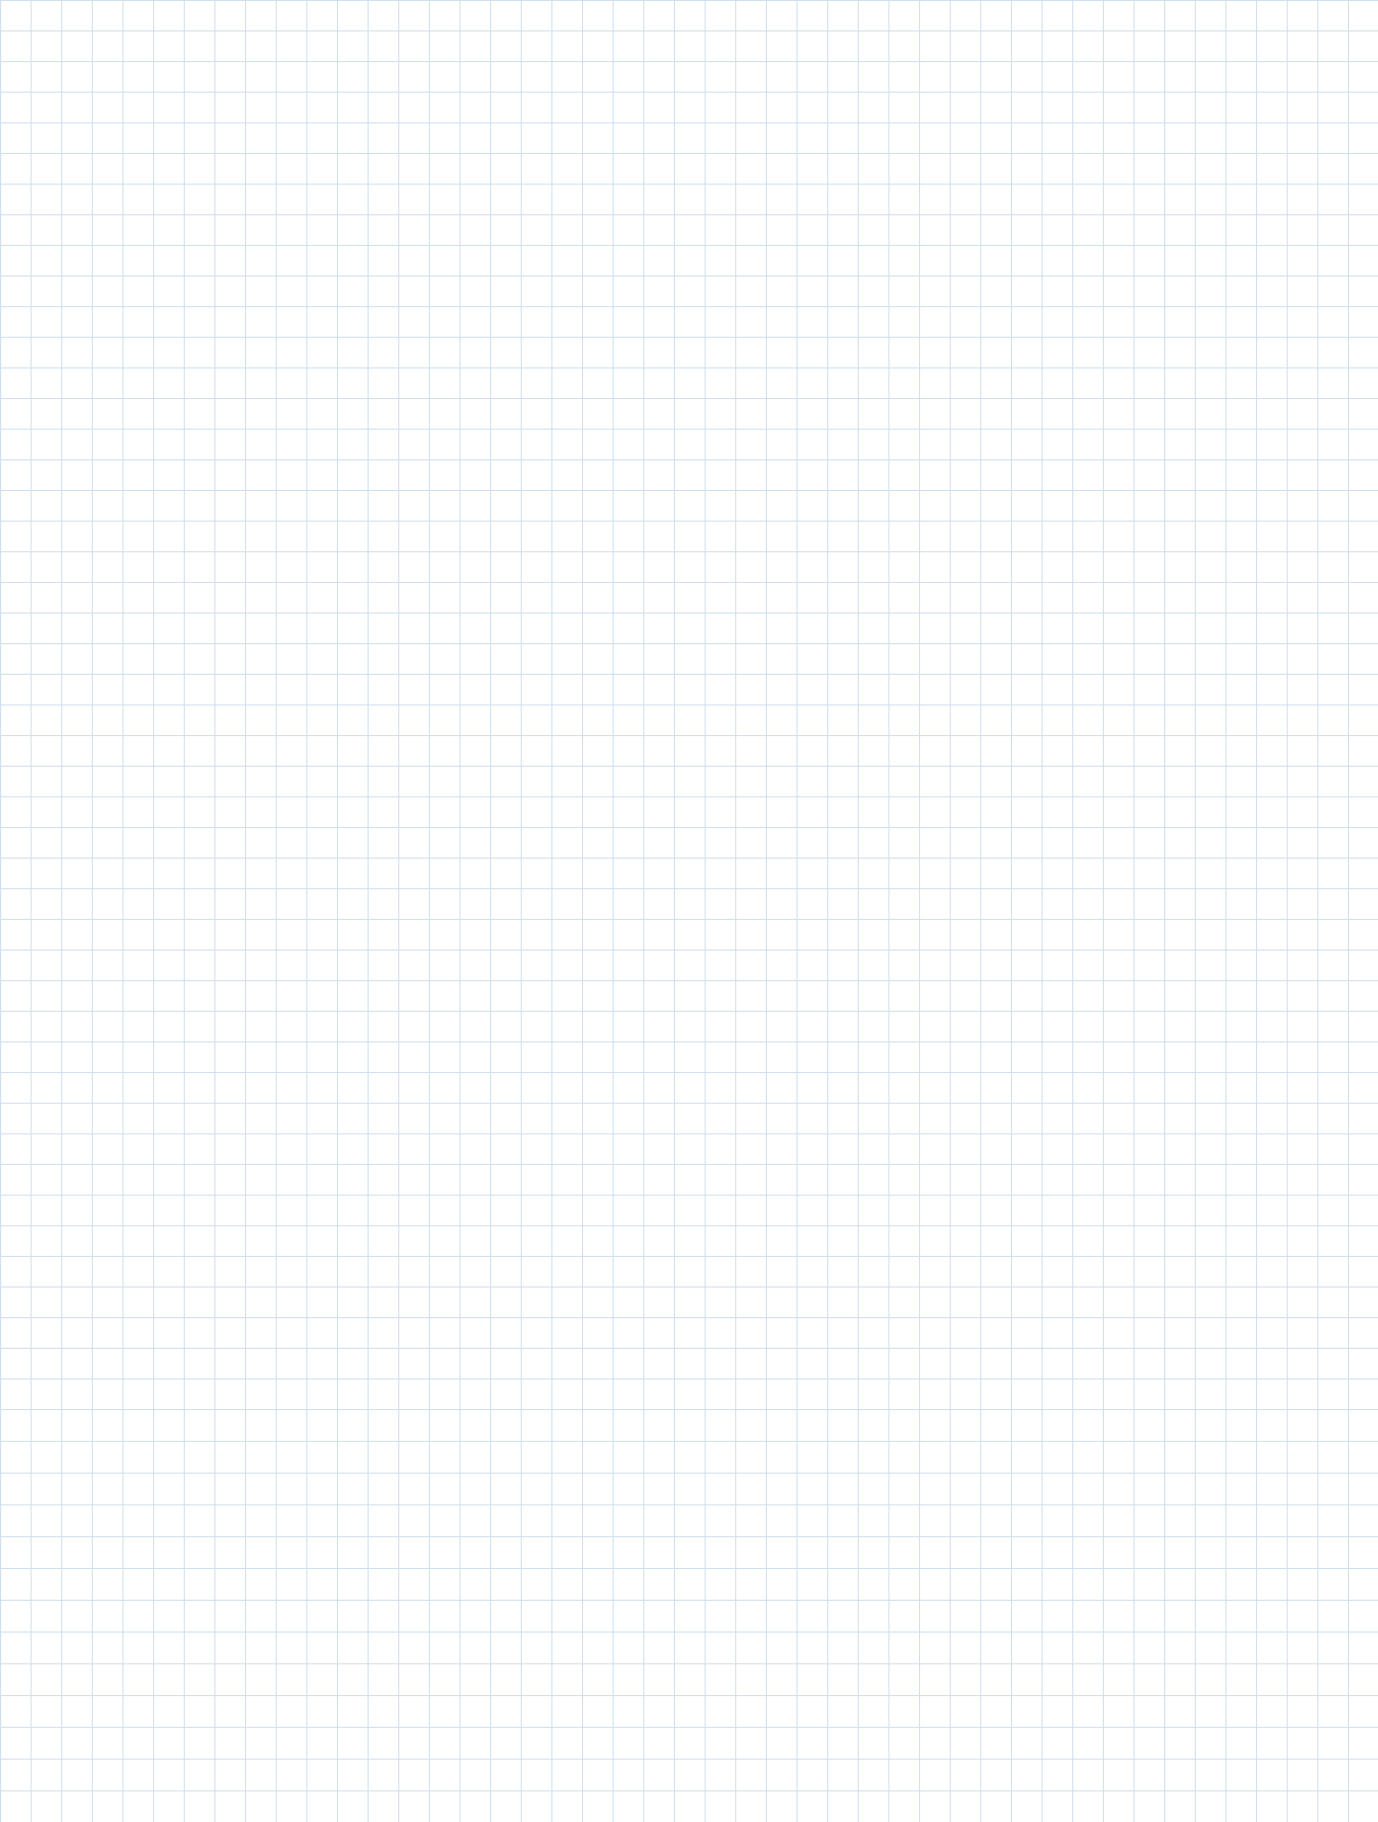
Yoshinori Kusuda

Yoshinori Kusuda [yoshinori.kusuda@analog.com] is an IC design engineer working for the Linear and Precision Technology Group in San Jose, CA. He has mainly worked for precision CMOS amplifiers and switched capacitor designs. He acquired B.S.E.E. and M.S.E.E. degrees from Tokyo Institute of Technology in 2002 and 2004, respectively.

Notes

analog.com/analogdialogue

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, Inc.
Wilhelm-Wagenfeld-Str. 6
80807 Munich
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

ANALOG DEVICES

AHEAD OF WHAT'S POSSIBLE™