Analog Dialogue

A forum for the exchange of circuits, systems, and software for real-world signal processing • Volume 50, Number 3, 2016

Analog Dialogue
Since 1967

analog.com/analogdialogue

ANALOG
DEVICES

AHEAD OF WHAT'S POSSIBLE™

Editor's Notes

IN THIS ISSUE

New RF DAC Broadens Software–Defined Radio Horizon

On the transmitter side of a communications system, a new class of high speed digital-to-analog converter, called an RF DAC, has the ability to directly synthesize RF signals without a conventional analog upconversion analog radio chain. In conjunction with the JESD204B high speed serial interface that enables practical management and transfer of output data, RF DACs enable conventional radio designs to be made more efficient than with previous generations of RF class or IF class DACs. This article explores in detail ADI's new RF DAC architecture, applications, and performance. (Page 3)

What's Up with Digital Downconverters—Part 1

Regardless of the end application, whether it is communications, aerospace/defense, or instrumentation, the frequencies of interest are pushing higher into the RF and microwave spectrum. Many current radio receiver architectures contain analog downconversion stages that translate an RF or microwave frequency band down to an inter-mediate frequency for baseband processing. This article explains how the new generation of RF ADCs with an integrated digital down-conversion function, eliminates the need for analog downconversion stages, and allows the spectrum in the RF frequency domain to be directly converted down to baseband for processing. (Page 7)

Crossing a New Frontier of Multiband Receivers with Gigasample ADCs—Part 1

Backhaul service providers who must respond to market demands for faster data rates and cheaper services are confronted with a dichotomous situation: faster data rates mean more bandwidth, which requires faster processing and increased system costs. The solution lies in RF sampling ADCs that integrate custom digital processing blocks to lower data rates, enabling the use of lower cost FPGAs as processors. This article discusses a use case for a multiband radio receiver design for TDD LTE bands 34 and 39, based on the AD9680 RF ADC. (Page 14)

Integrated Capacitive PGAs in ADCs: Redefining Performance

For many applications, capacitive programmable gain amplifiers (PGAs) offer better performance than traditional resistive PGAs. For instance, they provide higher common-mode voltage rejection of analog input signals, which is important in sensor applications. This article describes how a chopped capacitive amplifier operates and highlights the bene-fits of this architecture when a small signal from the sensor needs to be amplified close to the rails. Application examples analyzed include temperature measurement (RTDs or thermocouples) and Wheatstone bridges. (Page 17)

Where Zero–IF Wins: 50% Smaller PCB Footprint at ⅓ the Cost

Zero intermediate frequency (ZIF) transceiver architecture is a key enabler and optimal solution for high performance software-defined radio signal chain integration. But to achieve this, certain limitations in the ZIF approach had to be overcome in order to meet the performance, agility, and flexibility demanded by today's high end wideband wireless applications. Those limitations have been solved through innovation and the author walks through the architecture and performance data of a high performance ZIF transceiver implementation. (Page 21)

RF Transceivers Provide Breakthrough SWaP Solutions for Aerospace and Defense

Doing more with less is a mandate in military system design these days and reductions in SWaP (size, weight, and power) is a key system design requirement. The new generation of integrated RF agile wide-band transceivers enable SWaP reduction in many military applications. In this article's author's words, "The next-generation aerospace and defense platforms are demanding a new approach to RF design, one where several square inches of an existing platform are integrated into a single device. Where the boundary between software and hardware is blurred." This article explores an integrated programmable RF trans-ceiver signal chain solution that addresses SWaP reduction. (Page 28)

Jim Surber [jim.surber@analog.com]

Analog Dialogue

New RF DAC Broadens Software-Defined Radio Horizon

By Daniel E. Fague

Abstract

High speed data converters have been used in communications applications for many years and can be found in much of the equipment that forms the basis of our connected world, from cellular telephone base stations to cable head-end equipment to radar and specialized communication systems. Recent technological advances have enabled clock rates on high speed data converters to move to higher and higher frequencies. In conjunction with the JESD204B high speed serial interface that enables practical management and transfer of output data, these higher clock rate data converters form a new class of converter called an RF (for radio frequency) data converter. They have the ability to directly synthesize or capture RF signals without a conventional upconversion or downconversion with an analog radio chain.

This article will focus on a new RF digital-to-analog converter (RF DAC) family of products, the AD9162 and AD9164, and their ability to broaden the software-defined radio (SDR) definition. The AD9164 brings a new level of performance to the RF DAC class and enables conventional radio designs to be made more efficient than with previous generations of RF class or IF class DACs. The combination of world-best performance and a rich feature set make the AD9164 a natural choice for switching the context of the radio from one system to another, and moving one step closer to the reality of a truly software-defined radio.

Introduction

Conventional radio equipment has used high speed data converters in conjunction with quadrature modulators as some of the main building blocks for a wired or wireless communication link. The classic heterodyne, super heterodyne, and direct conversion architectures share the common need for data converters in the transmitter and the receiver to cross the boundary from digital processing to the real-world analog signals and back. Along with filter technology and power amplifier technology, data converter technology improvements set the pace for advances in radio design.

A classic radio transmitter implemented with a set of baseband high speed DACs is shown in Figure 1. The digital baseband data is sent through two synchronized high speed data converters, with in-phase data going through the I DAC and the quadrature data going through the Q DAC. The outputs of the DACs are sent to a quadrature modulator. Based on the type of modulator, its output can be a low intermediate frequency, such as 200 MHz to 400 MHz, a higher IF frequency such as 500 MHz to 1 GHz, or even an RF frequency in the 1 GHz to 5 GHz range. The diagram shows a subsequent upconversion to an eventual final frequency. The resulting signal is filtered with a band-pass filter and then sent through a power amplifier and another band-pass filter that may be part of a duplexer, for example.

The instantaneous bandwidth that is typically transmitted with such an architecture is tens to a few hundred MHz, limited mainly by the converter, power amplifier, and filter bandwidths. This is not enough for some systems, such as new E-band microwave backhaul radios that require 500 MHz, 1 GHz, or even 2 GHz radio channels. If a multiband radio is considered, such as might be implemented in a wireless infrastructure base station, an equally wide spacing of 500 MHz or 700 MHz or even 1 GHz may be required to cover some band combinations. A conventional radio would address that by implementing two radios, one for each band. It may be more desirable to combine radios into one radio chain, whether it be for cost or size or some other factor. In this case, a new approach is needed.
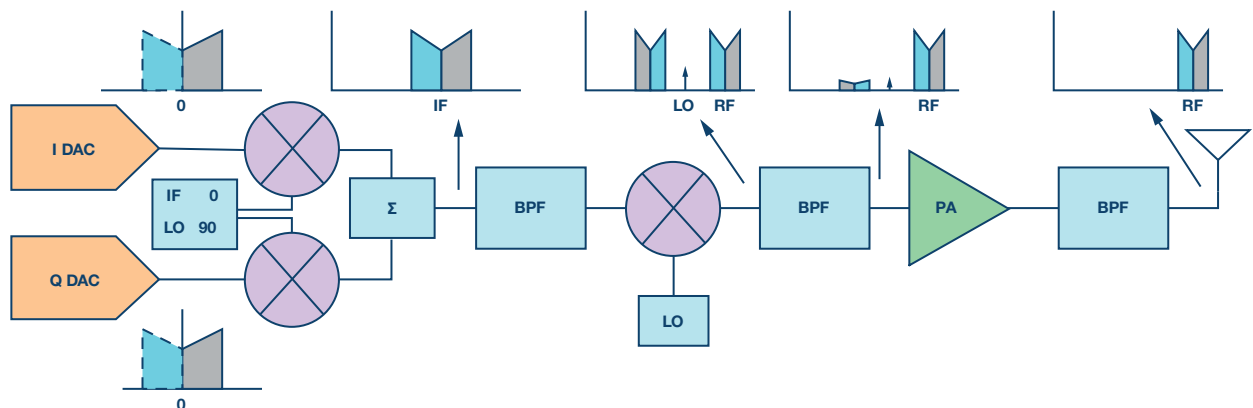

Figure 1. Illustration of a classic super heterodyne transmitter using high speed data converters.

Enabling Technology

The focus of technology development in high speed data converters has long been to push the rate of data conversion higher while maintaining a consistent performance figure of merit. The figure of merit includes such items as noise spectral density (NSD) and spurious-free dynamic range (SFDR). Intermodulation distortion (IMD) is also important, both with single-tone signals as well as modulated signals, such as those in popular wireless communications systems like GSM and 3G (W-CDMA) and 4G (OFDM), and cable applications where 256 QAM is used.

Higher rates of data conversion bring several advantages to the radio designer. First, the signal's image is pushed higher in frequency, making the design of the analog reconstruction filter simpler and more realizable. In addition, higher update rates create wider first Nyquist zones, which in turn enable the converter to directly synthesize higher output frequencies. When the directly synthesized signal is high enough, an entire stage of analog frequency translation, or upconversion, can be removed from the radio, simplifying frequency planning and reducing the power consumption and size of the radio. Higher update rates also increase the amount of bandwidth available to spread the quantization noise of the data converter, giving processing gain to the transmitter's noise spectral density.

As CMOS process technology has advanced, the addition of signal processing to data converters has become commonplace. The added feature sets of NCOs and interpolators in DACs relieves the FPGA or ASIC from the burden and power consumption of implementing those features, and enables the DACs to operate on lower data transfer rates than otherwise would be required. The lower data rates reduce overall power consumption in the system, and in some cases make it possible for the digital chip, on which the fabric speed may range up to 300 MHz to 400 MHz, to keep pace with the converter. Having NCOs on chip enable the first frequency translation in a radio to happen in the digital domain, and so it is common to find intermediate frequencies in the hundreds of MHz in today's radios, enabled by NCOs and interpolators on data converters.

Signal Processing RF DAC

What has changed with RF data converters is the ultimate update rate at which the RF converter is able to operate, and the addition of signal processing also capable of handling those speeds. This powerful combination of feature set and speed can enable dramatic changes to radio architecture design and opens new possibilities of reconfigurable and software-defined radios.

A good example of this is the AD9162 and AD9164 series of RF DACs. A block diagram of the AD9162 and AD9164 is shown in the Figure 2 block diagram of the AD9162 and AD9164 family of RF DACs. The AD9162 is a 16-bit, 6 GSPS RF DAC with several options of interpolation, from 1× bypass mode up to 24× interpolation. The interpolators operate on a classic 80% bandwidth or a wider 90% bandwidth for more instantaneous signal bandwidth at slightly higher power. The datapath also has a final half-band interpolator, FIR85, shown as the "HB 2×" block prior to the NCO in the Figure 2 block diagram of the AD9162 and AD9164 family of RF DACs, that effectively doubles the DAC update rate to up to 12 GSPS, moving images further away and easing filtering requirements. The optional FIR85 is followed by a 48-bit numerically controlled oscillator (NCO) that operates at either the 6 GSPS update rate or the 12 GSPS update rate when FIR85 is enabled. Following the NCO is an x/sinx compensation filter that corrects for the sinx/x roll-off of the DAC by pre-emphasizing the input to the DAC core.

The DAC core is designed with Analog Devices' patented Quad Switch architecture[1], delivering superior spurious-free dynamic range (SFDR) and noise spectral density (NSD), resulting in the industry's best dynamic range, while also providing the familiar DAC decoder options the Quad Switch enables: non-return-to-zero (NRZ) mode, return-to-zero (RZ) mode, and Mix-Mode™. The FIR85 adds a new feature to the DAC decoder called 2xNRZ mode, which will be described in more detail later.

The AD9164 has the base features of the AD9162, and it adds a direct digital synthesis (DDS) function in the form of a fast frequency hopping (FFH) NCO engine. The FFH NCO has several unique features that make it quite attractive for markets such as high speed test instrumentation, local oscillator replacement, secure radio communications, and radar exciters. The FFH NCO engine is implemented with thirty-two 32-bit NCOs, each with its own phase accumulator, and a selection block that enables the fast frequency hopping.

The AD9162 has two derivative products that are directed to specific markets. The AD9161 is an 11-bit, 6 GSPS RF DAC that has a minimum 2× interpolation. The SFDR and NSD of the AD9161 are suitable for cable head end and remote PHY applications and meet DOCSIS 3.0 specifications. The reduced signal bandwidth and dynamic range remove requirements for an export license for the AD9161. The AD9163 is a 16-bit,
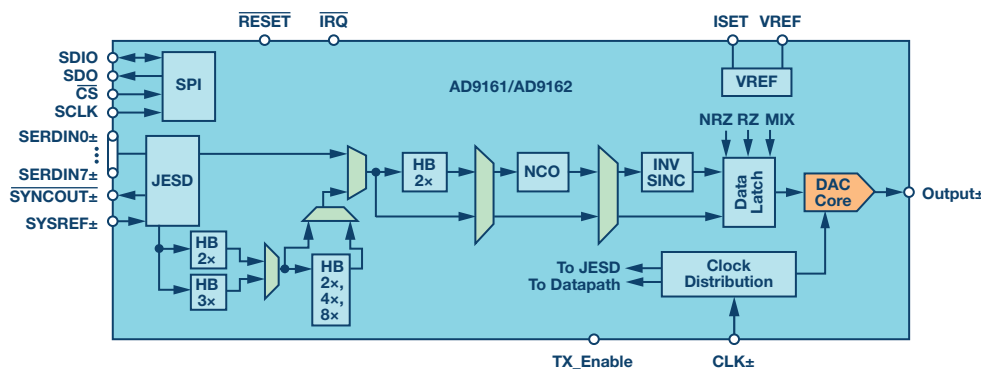

Figure 2. Block diagram of the AD9162 and AD9164 family of RF DACs.

6 GSPS RF DAC that has minimum 6× interpolation and retains the full dynamic range of the main product AD9162. The full dynamic range of the device and its wide 1 GHz instantaneous bandwidth plus full range NCO make the device suitable for single-band or dual-band wireless infrastructure base stations and also point-to-point microwave systems in the traditional bands, while also having the benefit of not requiring an export license. Table 1, Summary of AD9162 and AD9164 Family of 6 GSPS RF DAC Features and Target Markets summarizes the product family and key features.

Digital Datapath Highlights

Data is passed to the AD9162 and AD9164 by an 8-lane, 12.5 Gbps JESD204B interface. This high speed serial interface simplifies board layout complexity by reducing the number of wires needed to connect the digital baseband device to the DAC. A detailed guide to operation of the interface is given in the data sheet, and a comprehensive guide to the JESD204B interface is given on the Analog Devices website.

The first interpolator in the AD9162 and AD9164 datapath is either a 2× half-band or a 3× third-band filter. Either of these filters has a selectable 80% or 90% signal bandwidth. Both filters have 85 dB or greater of stop-band rejection. The 90% filters operate at higher power due to their sharper cutoff characteristic and therefore higher number of taps. The remaining 2× half-band filters all operate at 90% bandwidth to accommodate either of the first interpolators. The FIR85 also operates at 90% bandwidth. Because all of the subsequent filters are further down the line of interpolation, they can operate on the 90% bandwidth with an almost unnoticed increase in power.

The FIR85, which implements 2xNRZ mode when enabled, is implemented differently from the other interpolator filters. It takes advantage of the Quad Switch architecture of the DAC and uses the rising *and* falling edges of the DAC clock to sample data. This sampling method samples new data at each edge of the clock, and so it acts to double the sample rate of the DAC to up to 12 GSPS. This pushes the signal's image to $2xf_{DAC} - f_{OUT}$ from $f_{DAC} - f_{OUT}$, making it easier to filter the images with more realizable analog filters. This method of sampling and interpolating makes the DAC output more sensitive to clock balance, but there are adjustments to the DAC clock input that can be used to tune for better performance. These adjustments are made by programming registers through the serial peripheral interface (SPI). Details are given in the data sheet.

The 48-bit NCO is a full quadrature NCO that enables image-free frequency shifting of the input data signal or direct digital synthesis of a single tone. The NCO has two selectable modes of operation, either phase continuous or phase discontinuous frequency switching. In phase continuous switching, the

frequency tuning word (FTW) is updated but the phase accumulator is not reset, resulting in a continuous phase change in frequency. In phase discontinuous mode, the phase accumulator is reset when the FTW is updated. The serial peripheral interface (SPI) is guaranteed to 100 MHz to enable fast updating of the FTW.

The AD9164 adds an important feature to the NCO—the fast frequency hopping NCO (FFH NCO). The FFH NCO is implemented with an additional thirty-one 32-bit NCOs, each with their own phase accumulator. Each NCO has its own FTW, so that a total of 32 NCO FTWs can be programmed in the device. An FTW select register is provided so that a single SPI register byte write can accomplish a hop to a new frequency with accuracy to 32 bits. With the 100 MHz SPI, this means a new FTW can be chosen in 240 ns, with the single byte write.

The FFH NCO has an additional phase coherent frequency hopping mode that makes it attractive for instrumentation and military applications. Phase coherent frequency hopping is important for test applications and also radar applications that need to track the phase of an exciter signal for use later. Phase coherent frequency hopping enables changing from one frequency to another and back to the original one again, without losing track of the original frequency's phase accumulation. Stated another way, it enables changing from one frequency to another and back again, and *appearing as though the frequency was never changed.*

Applications and Measured Performance

The signal processing features and high sample rate of the AD9162 and AD9164 enable a simplification of the radio architecture in Figure 1. The updated drawing is shown in Figure 3. Because the RF data converter can directly synthesize signals at the desired output frequency, there is no longer a need for a quadrature modulator or an upconverting mixer. The signal is created in the digital processor and simply played out of the RF data converter. The amount of hardware needed to implement the transmitter is thus greatly reduced. In addition, the radio is simpler to implement, with no need to calibrate the LO and DAC inputs to a quadrature modulator to suppress LO leakage and unwanted image, because the modulator is implemented digitally inside the RF data converter.
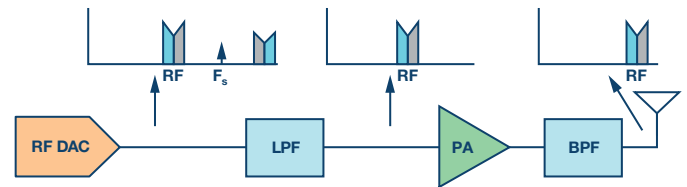
Figure 3. Radio transmitter architecture implemented with an RF data converter.

Table 1. Summary of AD9162 and AD9164 Family of 6 GSPS RF DAC Features and Target Markets

Part #	# Bits	Min. Interp.	FFH?	Target Market	Notes
AD9161	11	2×	N	Cable	11-bit version for cable customers
AD9162	16	1×	N	Cable, WIFR, instrumentation	Full performance DAC for world cable, WIFR, instrumentation markets
AD9163	16	6×	N	WIFR	1 GHz BW version for WIFR customers or nonfull band cable customers (such as MDU)
AD9164	16	1×	Y	Instrumentation, military, cable, WIFR	Full performance DAC and DDS for instrumentation, military, cable, WIFR, markets; has phase coherent fast frequency hopping

This type of architecture, with only an analog low-pass filter to filter the data converter's image, opens up possibilities for reconfigurable or software-defined radios. The same digital part, RF data converter, and reconstruction low-pass filter could be used, with only a power amplifier and band-pass filter change, to implement a number of different radios. Figure 4 shows an example of a wireless base station dual-band transmitter output of five 5 MHz W-CDMA carriers at 1800 MHz and three 5 MHz W-CDMA carriers at 2100 MHz. Figure 5 shows an example of a compliant cable head-end transmitter output of 194 6 MHz wide 256 QAM carriers in the 50 MHz to 1.2 GHz spectrum of DOCSIS 3.1. Figure 6 shows an example of a fast frequency hopping dwell time of 260 ns, with 240 ns of register programming (single byte write) and 20 ns of hop time. Figure 7 shows the excellent phase noise performance of the AD9164, with better than –125 dBc/Hz at 10 kHz offset when running from a 4 GHz ovenized crystal oscillator and synthesizing a 3.9 GHz sine wave.

Figure 6. Fast frequency hopping performance of the AD9164—260 ns dwell time per hop.
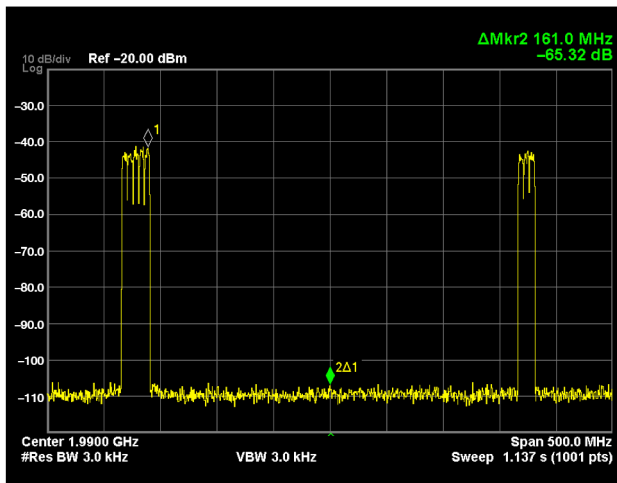
Figure 4. Dual–band W-CDMA signal at 1.8 GHz and 2.1 GHz bands.
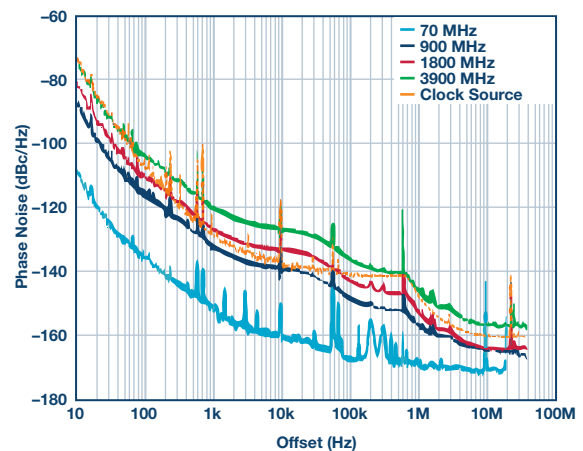
Figure 7. Aggregate phase noise performance of the AD9164. DAC clock signal source: 4 GHz ovenized crystal oscillator up to 600 kHz offset, then signal generator above 600 kHz offset.

Conclusion

RF data converters can simplify radio architecture designs and reduce their size by eliminating many components from the radio signal chain. The AD9162 and AD9164 combine an exciting set of features and superior RF performance into an RF data converter that is able to address a wide range of radio transmitter applications, demonstrating that the reality of a truly software-defined radio is closer than ever.

[1] US patent numbers #6,842,132 and #7,796,971

Figure 5. 194 6 MHz 256 QAM signals in the DOCSIS 3.1 frequency band (50 MHz to 1.2 GHz).

Daniel E. Fague

Daniel E. Fague [dan.fague@analog.com] is the high speed DAC applications engineering manager at Analog Devices. He received his B.S.E.E. from Gonzaga University in 1989 and his M.S.E.E. from the University of California at Davis in 1991. He joined Analog Devices' Wireless Handset Group in 1995, where he focused on handset radio architecture design, including direct conversion radios for GSM, EDGE, CDMA, and Bluetooth.® Prior to that, he worked for five years at National Semiconductor doing radio architecture design for DECT and PHS. Since joining the High Speed DAC Group in 2011, Dan has focused on RF DAC development. He holds seven patents and has published more than 30 articles and papers.

What's Up With Digital Downconverters—Part 1

By Jonathan Harris

Many current radio architectures contain downconversion stages that translate an RF or microwave frequency band down to an intermediate frequency for baseband processing. Regardless of the end application, whether it is communications, aerospace and defense, or instrumentation, the frequencies of interest are pushing higher into the RF and microwave spectrum. One possible solution to this scenario is to use an increasing number of downconversion stages, such as what is shown in Figure 1. However, another more efficient solution is to utilize an RF ADC with an integrated digital downconverter (DDC) as shown in Figure 2.

Integrating DDC functionality with an RF ADC eliminates the need for additional analog downconversion stages and allows the spectrum in the RF frequency domain to be directly converted down to baseband for processing. The capability of the RF ADC to process spectrum in the gigahertz frequency domain alleviates the need to perform potentially multiple downconversions in the analog domain. The ability of the DDC allows for tenability of the spectrum as well as filtering via the decimation filtering, which also provides the advantage of improving the dynamic range within the band (increases SNR). Additional discussion on this topic can be found here, "Not Your Grandfather's ADC," and here, "Gigasample ADCs Promise Direct RF Conversion." These articles provide some additional discussion on the AD9680 and the AD9625 and their DDC functionality.

The primary focus here will be on the DDC functionality that exists in the AD9680 (as well as the AD9690, AD9691, and AD9684). In order to understand DDC functionality and how to analyze the output spectrum when the DDC is employed with an ADC, we will take a look at an example with the AD9680-500. As an aid, the Frequency Folding Tool on the Analog Devices website will be utilized. This simple yet powerful tool can be used to aid in understanding the aliasing effects of an ADC, which is the first step in analyzing the output spectrum in an RF ADC with integrated DDCs such as the AD9680.

In this example, the AD9680-500 is operating with an input clock of 368.64 MHz and an analog input frequency of 270 MHz. First, it is important to understand the setup for the digital processing blocks in the AD9680. The AD9680 will be set to use the digital downconverter (DDC) where the input is real, the output is complex, the numerically controlled oscillator (NCO) tuning frequency is set to 98 MHz, half-band filter 1 (HB1) is enabled, and the 6 dB gain is enabled. Since the output is complex, the complex to real conversion block is disabled. The basic diagram for the DDC is shown below. In order to understand how the input tones are processed, it is important to understand that the signal first passes through the NCO, which shifts the input tones in frequency, then passes through the decimation, optionally through the gain block, and then optionally through the complex to real conversion.
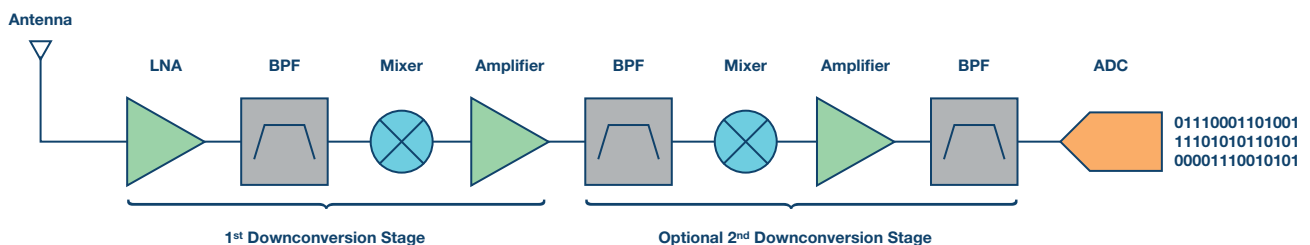
Figure 1. Typical receiver analog signal chain with downconversion stages.
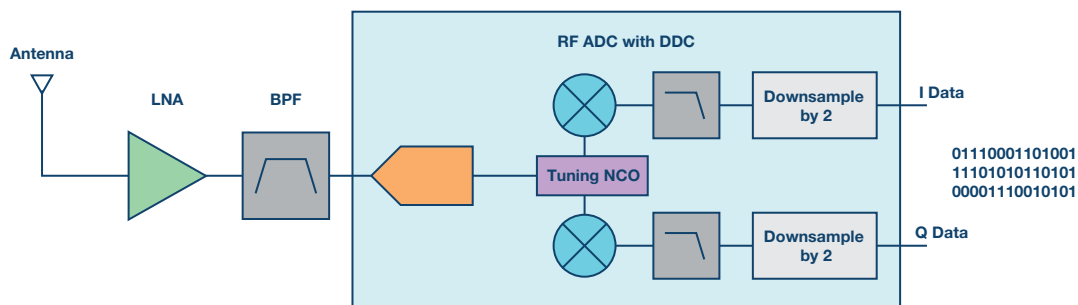
Figure 2. Receiver signal chain using an RF ADC with a DDC.

It is important to understand the macro view of the signal flow through the AD9680 as well. The signal enters through the analog inputs, passes through the ADC core, into the DDC, through the JESD204B serializer, and then out through the JESD204B serial output lanes. This is illustrated by the block diagram of the AD9680 shown in Figure 4.
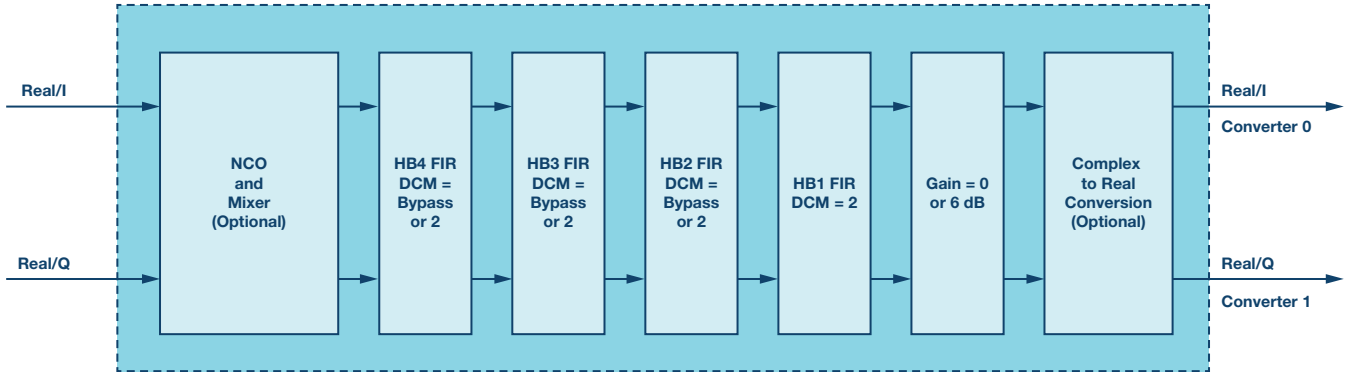
Figure 3. DDC signal processing blocks in the AD9680.
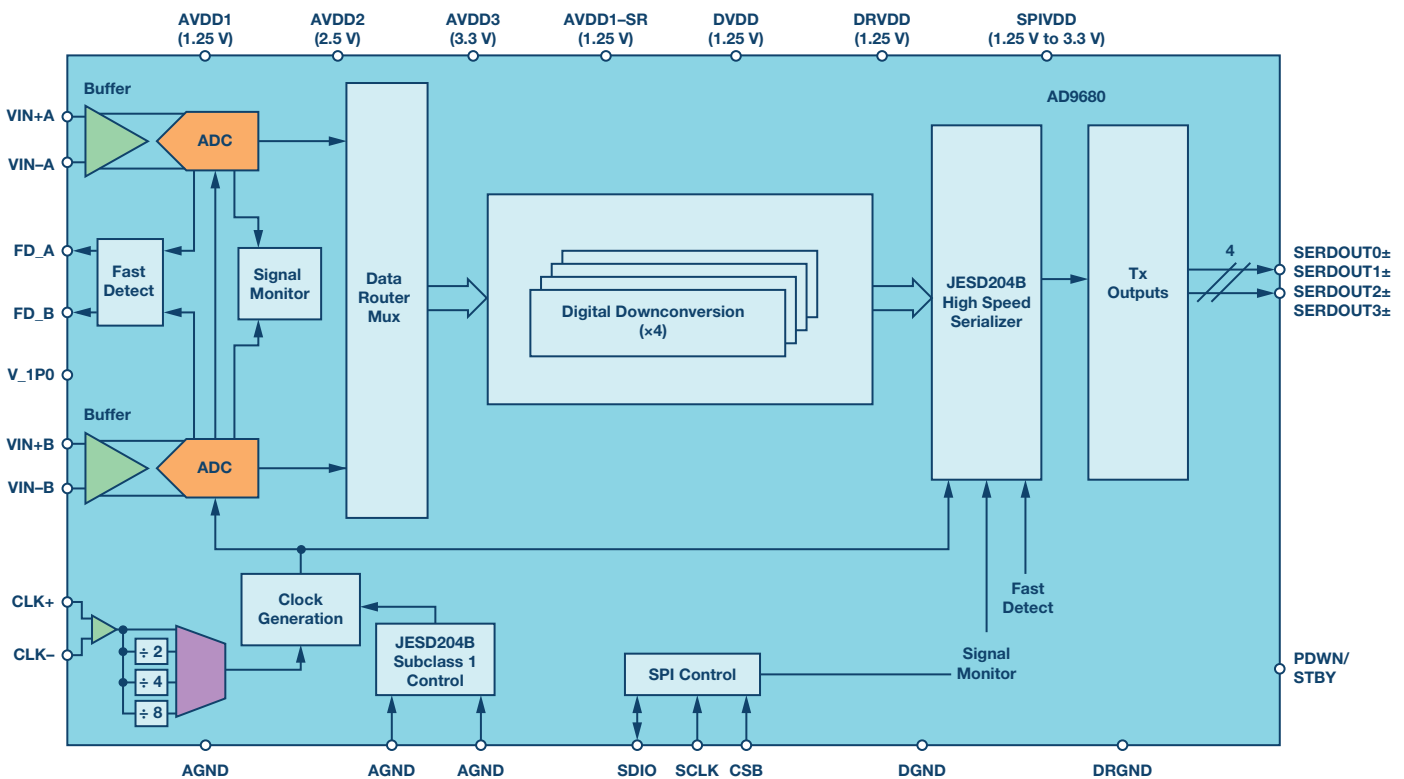
Figure 4. AD9680 block diagram.
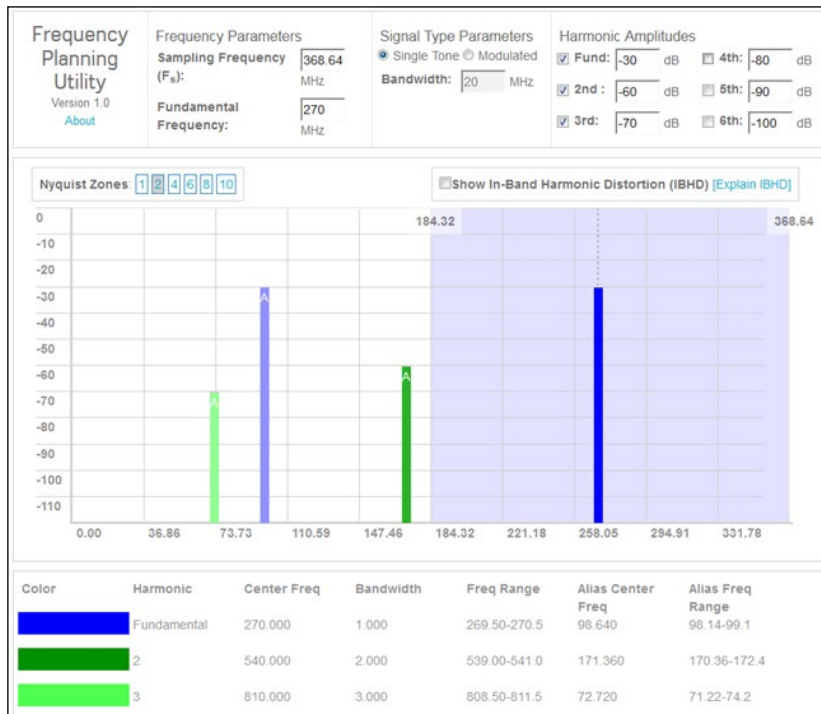
Figure 5. ADC output spectrum illustrated by the *Frequency Folding Tool*.

With an input sample clock of 368.64 MHz and an analog input frequency of 270 MHz, the input signal will alias into the first Nyquist zone at 98.64 MHz. The second harmonic of the input frequency will alias into the first Nyquist zone at 171.36 MHz while the third harmonic aliases to 72.72 MHz. This is illustrated by the plot of the *Frequency Folding Tool* in Figure 5.

The *Frequency Folding Tool* plot shown in Figure 5 gives the state of the signal at the output of the ADC core before it passes through the DDC in the AD9680. The first processing block that the signal passes through in the AD9680 is the NCO that will shift the spectrum to the left in the frequency domain by 98 MHz (recall our tuning frequency is 98 MHz). This will shift the analog input from 98.64 MHz down to 0.64 MHz, the second harmonic will shift down to 73.36 MHz, and the third harmonic will shift down to –25.28 MHz (recall we are looking at a complex output). This is shown in the FFT plot from Visual Analog in Figure 6 below.
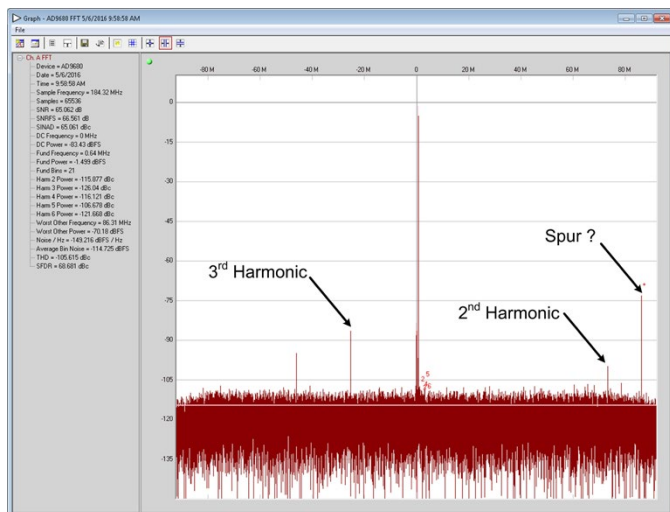
Figure 6. FFT complex output after a DDC with NCO = 98 MHz and decimate by 2.

From the FFT plot in Figure 6, we can see clearly how the NCO has shifted the frequencies that we observed in the *Frequency Folding Tool*. What is interesting is that we see an unexplained tone in the FFT. However, is this tone really unexplained? The NCO is not subjective and shifts all frequencies. In this case, it has shifted the alias of the fundamental input tone 98 MHz down to 0.64 MHz and shifted the second harmonic to 73.36 MHz and the third harmonic to –25.28 MHz. In addition, yet another tone has been shifted as well and appears at 86.32 MHz. Where did this tone actually come from? Did the signal processing of the DDC or the ADC somehow produce this tone? Well, the answer is no … and yes.

Let's look at this scenario a bit more closely. The *Frequency Folding Tool* does not include the dc offset of the ADC. This dc offset results in a tone present at dc (or 0 Hz). The *Frequency Folding Tool* is assuming an ideal ADC that would have no dc offset. In the actual output of the AD9680, the dc offset tone at 0 Hz is shifted down in frequency to –98 MHz. Due to the complex mixing and decimation, this dc offset tone folds back around into the first Nyquist zone in the real frequency domain. When looking at a complex input signal where a tone shifts into the second Nyquist zone in the negative frequency domain, it will wrap back around into the first Nyquist zone in the real frequency domain. Since we have decimation enabled with a decimation rate equal to two, our decimated Nyquist zone 92.16 MHz wide (recall: f_s = 368.64 MHz and the decimated sample rate is 184.32 MHz, which has a Nyquist zone of 92.16 MHz). The dc offset tone is shifted to –98 MHz, which is 5.84 MHz delta from the decimated Nyquist zone boundary at 92.16 MHz. When this tone folds back around into the first Nyquist zone it ends up at the same offset from the Nyquist zone boundary in the real frequency domain, which is 92.16 MHz – 5.84 MHz = 86.32 MHz. This is exactly where we see the tone in the FFT plot above! So technically, the ADC is producing the signal (since it is the dc offset) and the DDC is moving it around just a bit. This is where good frequency planning comes in. Proper frequency planning can help to avoid situations such as this one.

Now that we've looked at an example using the NCO and HB1 filter with a decimation rate equal to two, let's add a little more to the example. Now we will increase the decimation rate in the DDC to see the effects of frequency folding and translating when a higher decimation rate is employed along with frequency tuning with the NCO.

In this example we'll look at the AD9680-500 operating with an input clock of 491.52 MHz and an analog input frequency of 150.1 MHz. The AD9680 will be set to use the digital down-converter (DDC) with a real input, a complex output, an NCO tuning frequency of 155 MHz, half-band filter 1 (HB1) and half-band filter 2 (HB2) enabled (total decimation rate equals four), and 6 dB gain enabled. Since the output is complex, the complex to real conversion block is disabled. Recall from Figure 3 the basic diagram for the DDC, which gives the signal flow through the DDC. Once again the signal first passes through the NCO, which shifts the input tones in frequency, then passes through the decimation, through the gain block, and, in our case, bypasses the complex to real conversion.

Once again we will use the Frequency Folding Tool to help understand the aliasing effects of the ADC in order to evaluate where the analog input frequency and its harmonics will be located in the frequency domain. In this example we have a real signal, a sample rate of 491.52 MSPS, the decimation rate is set to four, and the output is complex. At the output of the ADC, the signal appears as illustrated below in Figure 7 with the Frequency Folding Tool.

With an input sample clock of 491.52 MHz and an analog input frequency of 150.1 MHz, the input signal will reside in the first Nyquist zone. The second harmonic of the input frequency at 300.2 MHz will alias into the first Nyquist zone at 191.32 MHz while the third harmonic at 450.3 MHz aliases into the first Nyquist zone at 41.22 MHz. This is the state of the signal at the output of the ADC before it passes through the DDC.
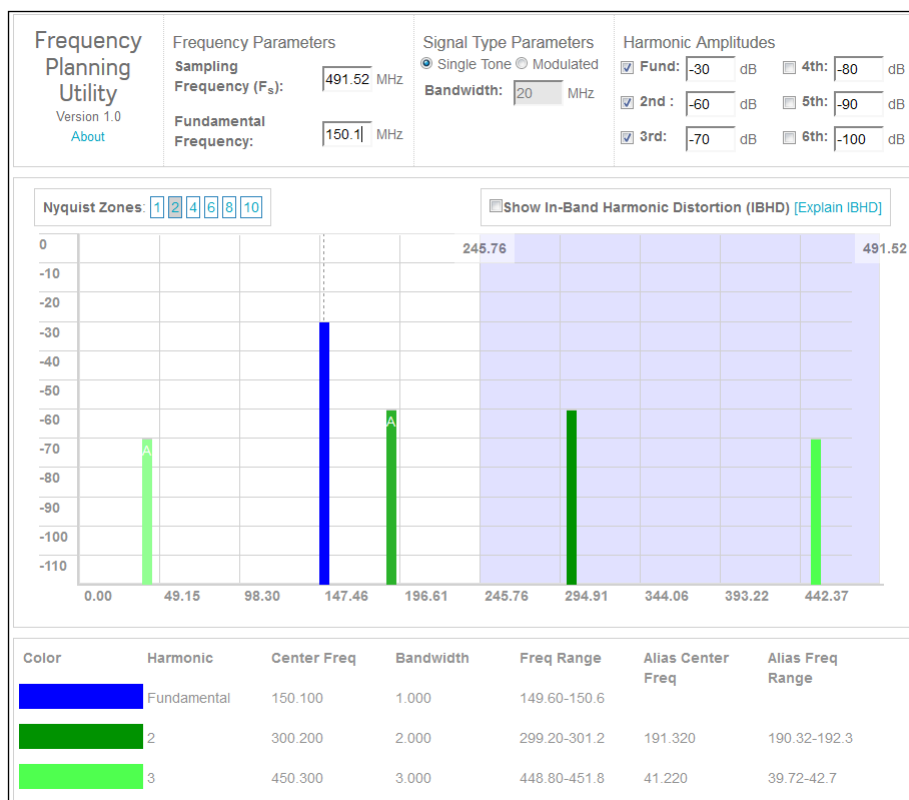
Figure 7. ADC output spectrum illustrated by the *Frequency Folding Tool.*

Now let's look at how the signal passes through the digital processing blocks inside the DDC. We will look at the signal as it goes through each stage and observe how the NCO shifts the signal and the decimation process subsequently folds the signal. We will maintain the plot in terms of the input sample rate, 491.52 MSPS and the f_s terms will be with respect to this sample rate. Let's observe the general process as shown in Figure 8. The NCO will shift the input signals to the left. Once the signal in the complex (negative frequency) domain shifts beyond $-f_s/2$, it will fold back around into the first Nyquist zone. Next the signal passes through the first decimation filter, HB1, which decimates by two. In the figure, I am showing the decimation process without showing the filter response even

though the operations occur together. This is for simplicity. After the first decimation by a factor of two, the spectrum from $f_s/4$ to $f_s/2$ translates into frequencies between $-f_s/4$ and dc. Similarly, the spectrum from $-f_s/2$ to $-f_s/4$ translates into the frequencies between dc and $f_s/4$. The signal now passes through the second decimation filter, HB2, which also decimates by two (the total decimation now is equal to four). The spectrum between $f_s/8$ and $f_s/4$ will now translate to the frequencies between $-f_s/8$ and dc. Similarly, the spectrum between $-f_s/4$ and $-f_s/8$ will translate to the frequencies between dc and $f_s/8$. Although decimation is indicated in the figure, the decimation filtering operation is not shown.
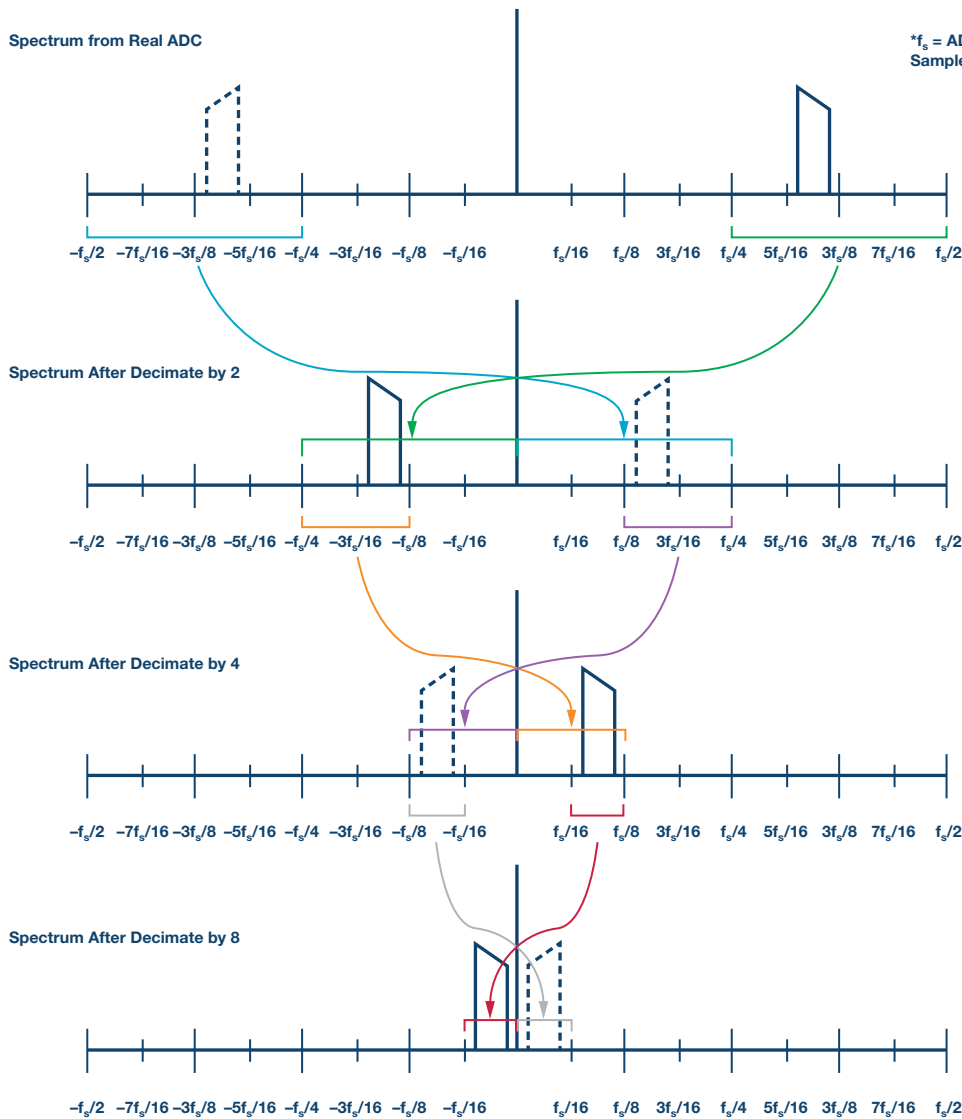
Figure 8. Effects of decimation filters on ADC output spectrum—generic example.

Recall the example previously discussed with an input sample rate of 491.52 MSPS and an input frequency of 150.1 MHz. The NCO frequency is 155 MHz and the decimation rate is equal to four (due to the NCO resolution, the actual NCO frequency is 154.94 MHz). This results in an output sample rate of 122.88 MSPS. Since the AD9680 is configured for complex mixing we will need to include the complex frequency domain in our analysis. Figure 9 shows the frequency translations are quite busy, but with careful study we can work our way through the signal flow.

Spectrum after the NCO shift:

1. The fundamental frequency shifts from +150.1 MHz down to –4.94 MHz.
2. The image of the fundamental shifts from –150.1 MHz and wraps around to +186.48 MHz.
3. The second harmonic shifts from 191.32 MHz down to 36.38 MHz.
4. The third harmonic shifts from +41.22 MHz down to –113.72 MHz.

Spectrum after decimate by 2:

1. The fundamental frequency stays at –4.94 MHz.
2. The image of the fundamental translates down to –59.28 MHz and is attenuated by the HB1 decimation filter.
3. The second harmonic stays at 36.38 MHz.
4. The third harmonic is attenuated significantly by the HB1 decimation filter.

Spectrum after decimate by 4:

1. The fundamental stays at –4.94 MHz.
2. The image of the fundamental stays at –59.28 MHz.
3. The second harmonic stays at –36.38 MHz.
4. The third harmonic is filtered and virtually eliminated by the HB2 decimation filter.

Figure 9. Effects of decimation filters on ADC output spectrum—actual example.

Now let's look at the actual measurement on the AD9680-500. We can see the fundamental resides at –4.94 MHz. The image of the fundamental resides at –59.28 MHz with an amplitude of –67.112 dBFS, which means that the image has been attenuated by approximately 66 dB. The second harmonic resides at 36.38 MHz. Notice that VisualAnalog does not properly find the harmonic frequencies since it does not interpret the NCO frequency and decimation rates.
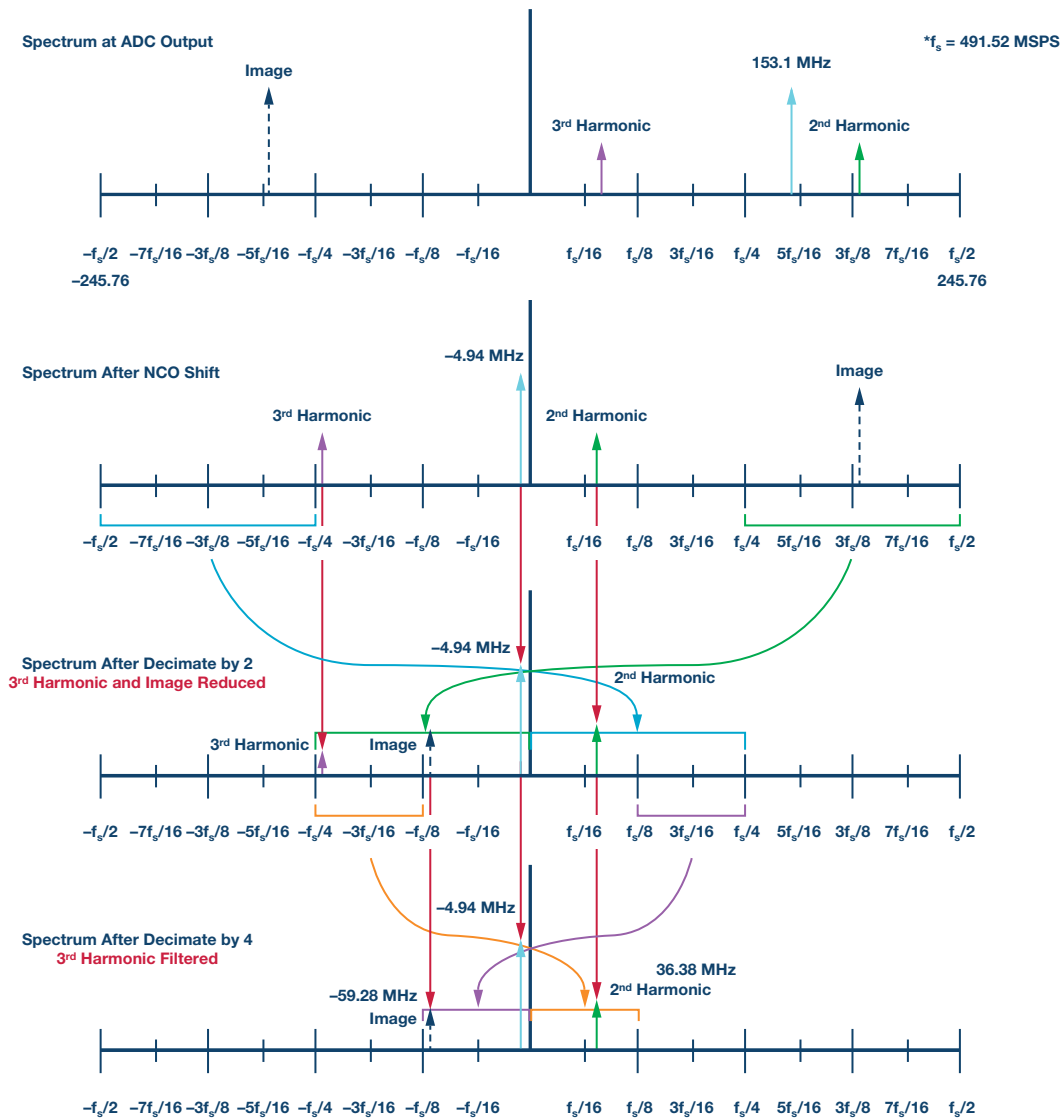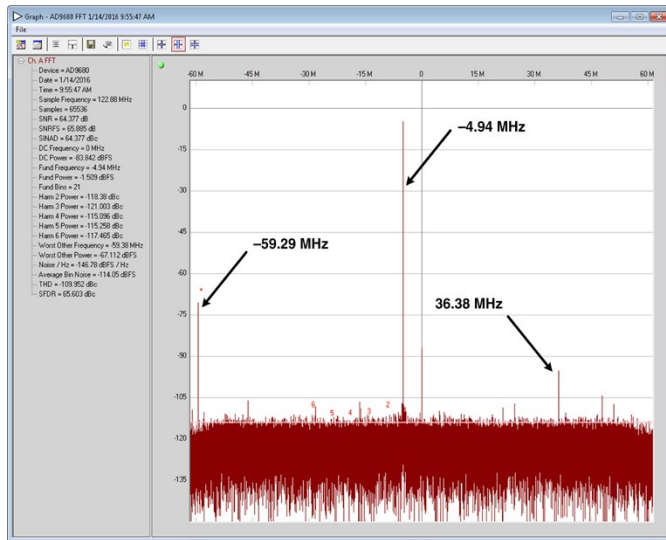

Figure 10. FFT complex output plot of signal after DDC with NCO = 155 MHz and decimate by 4.

From the FFT we can see the output spectrum of the AD9680-500 with the DDC set up for a real input and complex output with an NCO frequency of 155 MHz (actual 154.94 MHz), and a decimation rate equal to four. I encourage you to walk through the signal flow diagram to understand how the spectrum is shifted and translated. I also would encourage you to walk through the examples provided within this article carefully to understand the effects of the DDC on the ADC output spectrum. I recommend printing out Figure 8 and keeping it handy for reference when analyzing the output spectrum of the AD9680, AD9690, AD9691, and AD9684. While supporting these products, I have had many questions related to frequencies that are in the output spectrum of the ADCs that are considered unexplainable. However, once the analysis is done and the signal flow is analyzed through the NCO and the decimation filters, it becomes evident that what were at first considered unexplained spurs in the spectrum are actually just signals residing exactly where they should be. It is my hope that after reading and studying this article you are better equipped to handle questions the next time you are working with an ADC that has integrated DDCs. Stay tuned for part two, where we will continue looking at additional aspects of the DDC operation and also how we can simulate its behavior. We will look at the decimation filter responses due to ADC aliasing, more examples will be provided, and Virtual Eval will be used to observe operation of the DDC in the AD9680 and its effects on the ADC output spectrum.

Jonathan Harris

Jonathan Harris [jonathan.harris@analog.com] is a product applications engineer in the Space Products Group at Analog Devices in Greensboro, NC. He has over 10 years of experience as an applications engineer supporting products in the RF industry. Jonathan received his M.S.E.E. from Auburn University and his B.S.E.E. from UNC Charlotte. In his spare time he enjoys motorcycle riding, college football, mobile audio, and spending time with his family.

Crossing a New Frontier of Multiband Receivers with Gigasample ADCs—Part 1

By Umesh Jayamohan

Introduction

The analog-to-digital converter (ADC) has been a staple of communications receiver design for quite some time now. As communications technology continues to evolve, consumers are demanding faster data rates and cheaper services. Backhaul service providers who enable this technology are confronted with a dichotomous situation. Faster data rates mean more bandwidth, which translates to faster data converters that convert the analog air waves to be digitally processed. However, faster data converters (GSPS or gigasample per second converters)—widely known as RF sampling ADCs—also produce huge amounts of data that have to be processed at a much higher speed in these DSP chips. This invariably increases the cost of operating a radio receiver.

The solution lies in clever design of the silicon that forms the RF sampling ADC. Taking advantage of the silicon processing advancements (thank you, Moore's law), RF sampling ADCs mix in custom digital processing blocks that are more power and area efficient compared to existing FPGAs. The use of these digital signal processing blocks also results in lower data rates, which enable the use of lower cost FPGAs. This is a win-win situation for operators since they can sample at high frequencies using these GSPS ADCS, use the internal digital downconverters (DDCs) to process the data at speed, and send it out at a manageable (low) data rate to a cheaper FPGA (or existing generation ASIC) for further baseband processing.

The other advantage of using RF sampling ADCs with the DDCs is that this enables a more flexible, compact, cost-effective way to implement a dual-band radio system. Dual-band radio systems have been around for years now. Base station systems designers have traditionally implemented the dual-band radio systems by making use of two separate radio paths, one for each band. This article discusses a method to utilize a multiband radio receiver using an RF sampling ADC like the AD9680 to digitize and process two separate widely used bands. Part 1 of the article explains a block diagram level

implementation and discusses advantages of using a GSPS ADC for a dual-band radio system. Part 2 of the article will discuss an implementation and data analysis of TDD LTE bands 34 and 39 (also known as band A and band F, respectively) with data analysis to show converter performance.

Traditional Dual–Band Radio Receiver

In order to cater to customer demands for dual-band radios and meet overall system-level performance, base station designers resorted to what they knew best: replicate a radio's design twice and tune one for each band. This meant that the designer had to have two separate radio hardware designs tuned to the two bands that the customer had chosen.

For example, if there was a need to build a radio receiver that could support TDD LTE band 34 (band A: 2010 MHz to 2025 MHz) and band 39 (band F: 1880 MHz to 1920 MHz)[1], the designer would pack two radio receiver designs. The frequency plan for the TDD LTE bands is shown in Figure 1.
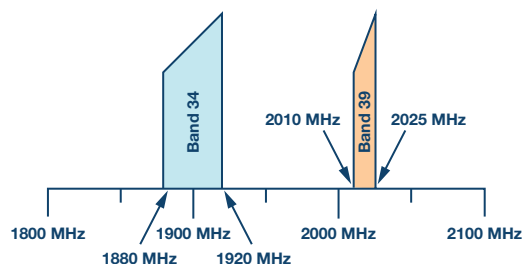

Figure 1. Frequency plan showing TDD LTE bands 34 and 39.

The traditional approach to designing a dual-band radio receiver to accommodate these bands would be to implement two separate receiver chains, one for each band. A block diagram representation of the dual-band radio receiver is shown in Figure 2 below.[2]
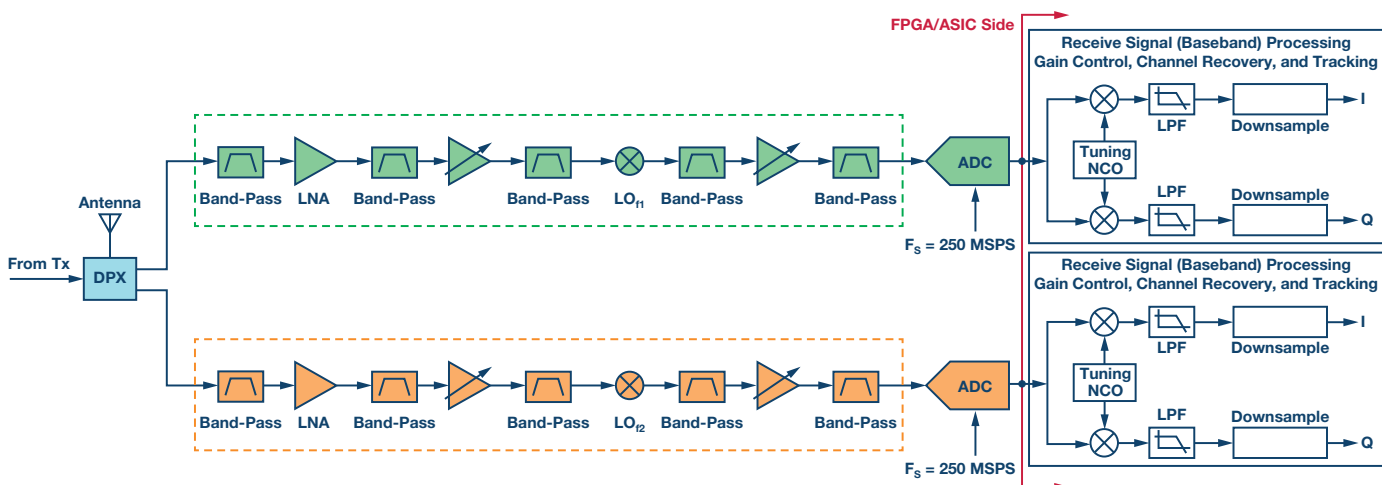

Figure 2. Traditional approach to a dual–band radio receiver design.

Figure 2 shows a traditional implementation of a dual-band radio. This implementation is fairly expensive to implement because it practically is two radio receivers in one system. Every processing element is duplicated in order to accommodate the respective bands. This also applies to FPGA resources. Every processing element is duplicated in order to accommodate the respective bands resulting in duplication of FPGA resources, increased system cost and complexity, and additional power. In terms of the interface to the FPGA, the FPGA resources will have to be doubled to accommodate the two ADC data streams. Figure 3 shows a block diagram representation of the FPGA I/O resource requirements or a dual-band radio receiver system design. It shows both an LVDS and a JESD204B ADC interface. The LVDS data rates are lower, but the FPGA will need a higher I/O count. The JESD204B interface requires a lesser number of I/O resources from the FPGA, but the lane rates can be higher, which could warrant a more expensive FPGA.
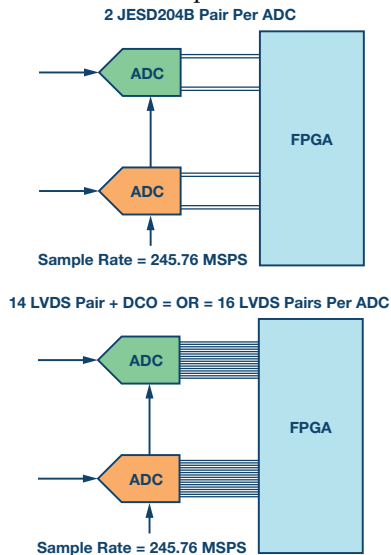

Figure 3. FPGA interface requirements for a traditional approach to a dual-band radio receiver.

Dual–Band Radio Receiver Using RF Sampling (GSPS) ADCs

The RF sampling or GSPS ADC can offer system design flexibility. By taking advantage of the deep submicron process technology, the GSPS ADCs can pack digital processing blocks that can manipulate the data at speed with much less power consumption compared to an FPGA. At the heart of the RF sampling ADC is a high bandwidth analog sampling core that samples at GHz speeds. Following the analog core is a plethora of digital signal processing elements. These digital downconverters can be used to extract the respective bands. A block diagram of the internals of an RF sampling ADC setup for a dual-band receiver is shown in Figure 4. The DDCs, in addition to processing the signals, also reduce the lane rate of the data on the JESD204B lanes.

With the added digital signal processing blocks, the GSPS ADC can now single-handedly accommodate two bands for processing. This is a win-win situation for operators since they can sample at high frequencies using these RF sampling ADCs, use the internal digital downconverters (DDCs) to process the data at speed, and send it out at a manageable (low) data rate to a cheaper FPGA (or existing generation ASIC) for further baseband processing. The high bandwidth front end offered by these ADCs enables the system designer to capture a wide swath of frequencies (two radio bands, for example) and digitize it for signal processing. Figure 5 below shows a dual-band receiver system using an RF sampling ADC and internal DDCs to extract the bands. As is evident by comparing with the implementation in Figure 2, the dual-band receiver using the RF sampling ADC is significantly simpler in implementation. In this implementation the RF is mixed down to a high IF that is many hundreds of MHz wide, as opposed to many tens of MHz wide in the traditional dual-band approach. The BPF and VGA stage is optional and is dependent on what level of system performance is expected.
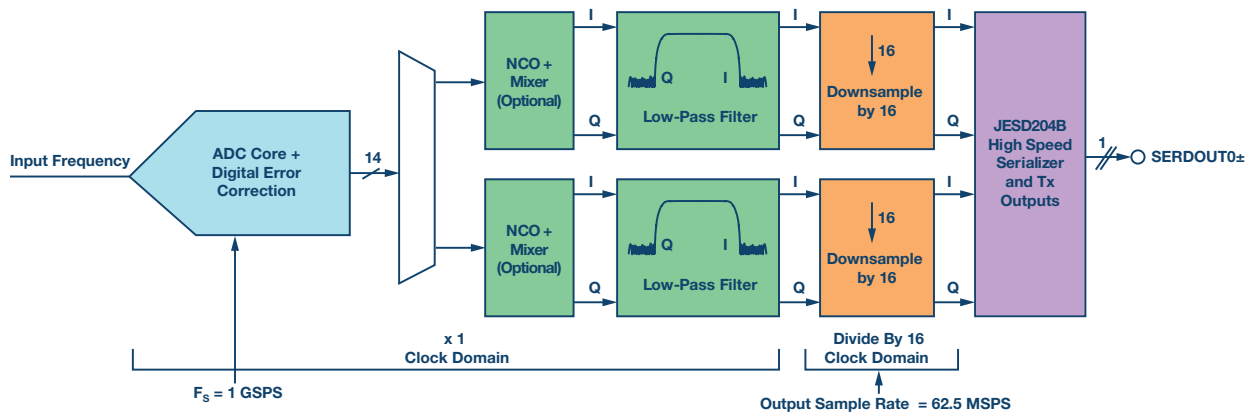

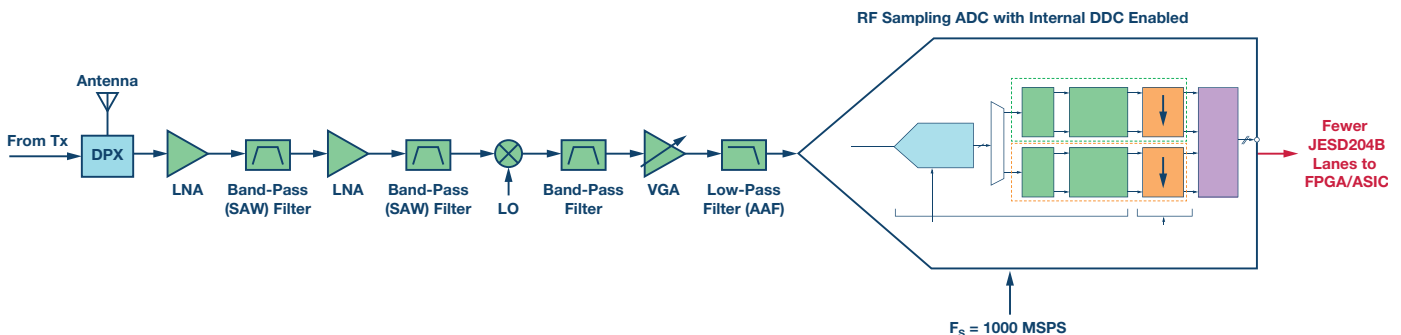Figure 4. Block diagram representation of the RF sampling ADC showing the internal DDCs.

Figure 5. A dual–band radio receiver using the RF sampling ADC and internal DDCs to extract the bands.

Some of the advantages of using an RF sampling ADC for a dual-band radio system are explained below:

Simpler Front–End Design

A dual-band radio system design using an RF sampling ADC greatly simplifies the front-end network. For starters, there is only a need for a single front-end design instead of two (one for each band). This heavily reduces the bill of material for the system board. Then there is the AAF (antialiasing filter) requirement, which is a band-pass filter (BPF) for the two IF converter case, compared with a low-pass filter (LPF) for the GSPS ADC case. This is because the GSPS ADC oversamples the input signal.[3, 4] Now that the data is oversampled, the digital downconverters can do their job of decimating and filtering. If the frequency plan is such that the second and third harmonics fall out of band, this eases the requirements on the AAF.

Lower System Power, Smaller Form Factor

Instead of two LNAs, two mixers, and two IF ADCs as shown in Figure 2, there is only a need for one front end in the RF sampling case (Figure 5). This results in major power savings from a system-level power consideration. The lower system power clubbed with the need for a simpler front-end design allows systems to be made in smaller form factor.

More Efficient FPGA Utilization

When using an RF sampling ADC to implement a dual-band radio system, the DDCs are employed to extract the individual bands. Since DDCs decimate the data, the output sample rate is reduced. This results in many flexible configurations for the JESD204B interface. For example, if a dual ADC is sampling at 1 GSPS and is in full bandwidth mode, the line rate is calculated at 10 Gbps/lane over four lanes. The line rate for JESD204B converters from Analog Devices can be calculated as follows:

$$Lane\ Line\ Rate = \frac{M \times N' \times \left(\frac{10}{8}\right) \times F_{OUT}}{L}\ \text{where,}$$

M = number of converters (in this example, 2)

N' = number of converter bits per sample (in this example, 16)

10/8 = 8B10B overhead

F_{OUT} = output sample rate (Fsample/Decimation_Ratio; in this example, Decimation_Ratio = 1 for full bandwidth)

L = number of JESD204B lanes (in this example, 4)

If the same dual ADC is utilizing a total of four DDCs in a decimate-by-8 configuration, for example, then there are many configurations that the ADC will support, based on the number of lanes. The output sample rate becomes 125 MSPS (1 GSPS ÷ 8). The different configurations are listed in Table 1:

Table 1

DDC Configuration	M	L	Line Rate (Gbps/Lane)
Real	4	1	10
Real	4	2	5
Complex	8	2	10
Complex	8	4	5

These flexible configurations provide the systems designer the freedom to either use an expensive FPGA with higher line rate but better I/O lane density usage, or use an existing FPGA/ASIC that has line rate limitations.

Conclusion

The advent of the GSPS ADC in deep submicron silicon processes has ushered in a new era of radio architecture discussion and design. The GSPS ADC with its high bandwidth sampling core and the digital downconverter options offer a flexible pathway to rethink and redefine the radio architecture that will cater to the growing demands of the consumer. The reduction in power and space offered by these GSPS ADCs will lower the cost of ownership for these radio boxes. The flexible output options provided by the current generation ADCs with the JESD204B interface do not lock the systems designer into using an expensive high line rate FPGA or digital logic.

Part 2 of this article will discuss a use case involving TDD LTE bands 34 and 39 and its analyses in a multiband radio receiver using the AD9680.[5]

References

[1] E-UTRA Bands.

[2] Walt Kester. *The Data Conversion Handbook*. Analog Devices, Inc., 2005.

[3] Umesh Jayamohan. "Not Your Grandfather's ADC: RF Sampling ADCs Offer Advantages in Systems Design." Analog Devices, Inc., 2015.

[4] Oversampling.

[5] AD9680. Analog Devices, Inc.

Umesh Jayamohan

Umesh Jayamohan [umesh.jayamohan@analog.com] is an applications engineer with Analog Devices in the High Speed Converter Group (Greensboro, NC). He has been a part of Analog Devices since 2010. Umesh received his B.S.E.E. from the University of Kerala, India, in 1998 and his M.S.E.E. from Arizona State University in 2002.

Also by this Author:

Rarely Asked Questions— Issue 129, May 2016 Who Ate My dBs?

Integrated Capacitive PGAs in ADCs: Redefining Performance

By Miguel Usach Merino and Gerard Mora–Puchalt

Abstract

Patented ADI capacitive programmable gain amplifiers (PGAs) offer better performance than traditional resistive PGAs, including higher common-mode voltage rejection of analog input signals.

This article describes how a chopped capacitive amplifier operates, highlighting the benefits of this architecture when a small signal from the sensor needs to be amplified close to the rails—like in temperature measurements (RTDs or thermocouples) and Wheatstone bridges.

Σ-Δ analog-to-digital converters (ADCs) are widely used in applications with sensors that have small responsivity and reduced bandwidth, such as strain gage or thermistors, due to the high dynamic range offered by this architecture. The reason behind the high dynamic range is the low noise performance compared with other ADC architectures.

The Σ-Δ converters base their operation on two principles: oversampling and noise shaping. When an ADC samples the input signal, the quantization noise, which is independent from the sampling frequency, spreads across the entire frequency spectrum up to half of the sampling frequency. Therefore, if the input signal is sampled at a much higher frequency than the minimum dictated by the Nyquist theorem, the quantization noise in the band of interest reduces.

Figure 1 shows an example of the quantization noise density for different sampling frequencies.
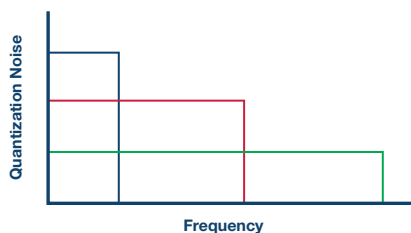

Figure 1. Quantization noise density over frequency for different sampling frequencies.

In general and for a given band of interest, the dynamic range improves by 3 dB every oversampling factor of 2 (assuming a white noise spectrum). The second benefit in a Σ-Δ converter is the noise transfer function. It shapes the noise to higher frequencies, as shown in Figure 2, which reduces even more of the quantization noise at the band of interest.
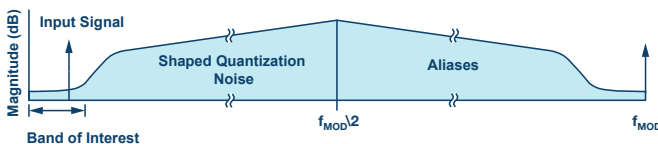

Figure 2. Σ–Δ noise shaping.

In addition, the Σ-Δ may incorporate a digital filter to remove the quantization noise outside the band of interest, equating to an excellent dynamic range performance, as shown in Figure 3.
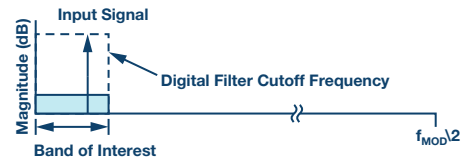

Figure 3. Quantization noise after the LPF.

Input Buffer

One of the disadvantages of oversampling architectures is that the requirements for an input buffer to drive the Σ-Δ modulator may become more stringent, compared with other architectures that operate at lower sampling frequencies. The acquisition time becomes shorter and, therefore, the buffer requires a higher bandwidth. Modern Σ-Δ converters integrate the input buffer on chip to maximize the ease of use.

Moreover, in sensing systems, presenting a very high input impedance with high precision to the sensing element is critical for the accuracy of the measurement. This makes the requirement for input buffers even more critical.

Integrating an input buffer generates other challenges. The Σ-Δ modulator offers very low noise at low frequencies, but any additional components like the input buffer will add thermal noise and, more importantly, flicker noise at low frequencies, as shown in Figure 4.
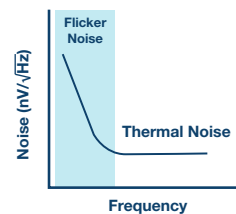
Figure 4. Flicker noise.

In addition, the offset of the buffer may contribute to the overall system error. The offset can be compensated by a system calibration but if the offset drift is relatively high, this approach may become impractical, as it would require that the system be recalibrated every time the temperature of operation changes to compensate the buffer offset contribution.

As an example, when having an offset drift of 500 nV/°C, a 10°C temperature increment will equate into 5 µV offset change, which in a ±2.5 VREF 24-bit ADC equates to 16.8 LSBs, which is around 4 bits.

The typical way to solve these two problems is by chopping the inputs and outputs of the buffer, as shown in Figure 5.

Figure 5. Chop amplifier.

By chopping the inputs, the input frequency is modulated to higher frequencies. The buffer offset and flicker noise remain at their original low frequencies, as they are not affected by the input chopping.

The output dechopper mechanism demodulates the input frequency back to baseband, while it modulates the offset and flicker noise added by the buffer up to higher frequencies that will be removed by the ADC low-pass filter.
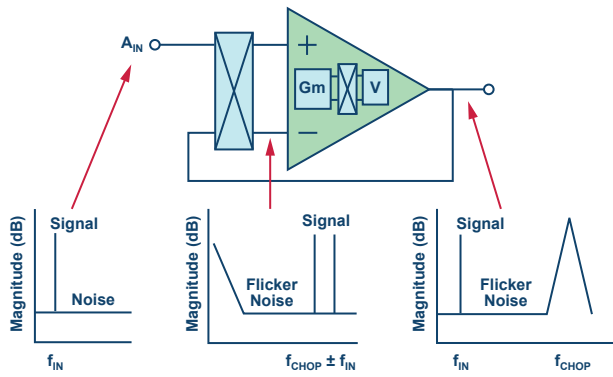
In some cases, the input buffers are replaced by a resistor-based instrumentation amplifier (resistive PGA) to accommodate a small sensor signal to the full modulator input range, maximizing the dynamic range. Just note that a resistor-based instrumentation amplifier is preferred over a differential resistive amplifier due to the higher input impedance required in discrete sensors. The resistive PGA implements a similar chopping scheme, as shown in Figure 6.
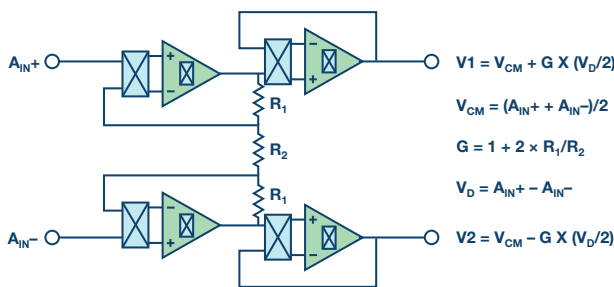

Figure 6. Resistive PGA.

The resistive PGA may require a second set of buffers connected in cascade, as the amplifier may not provide enough bandwidth to drive the modulator directly. Concurrently, the current consumption should be kept low, which dictates the value of the resistors and, consequently, the bandwidth of the amplifier.

The main restriction in using this amplifier topology is the limitation on the common-mode voltage, especially with gains other than one, as the resistive PGA has a floating common mode that depends on the input signal, as shown in Figure 6.

Additionally, the resistive network mismatch and its drift is also a concern in the overall error budget, as it may have an impact on most of the precision specifications.

To avoid these limitations, recent ADI Σ-Δ converters have employed a capacitive PGA.

The capacitive PGA amplification principle is similar to the resistive PGA; the gain depends on the capacitor ratios, as shown in Figure 7.

In order to amplify dc signals, the capacitive PGA introduces a chopping mechanism at the PGA inputs, the dc input signal is modulated to the chop frequency, and then it is amplified by the capacitive amplifier. Finally, the signal is demodulated back to dc by the output dechopper. Additionally, the amplifier offset and flicker noise is modulated to the chop frequency and low-pass filtered at a later stage.
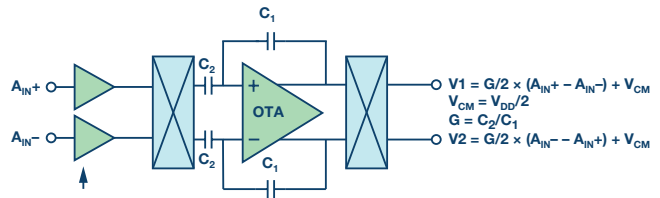
Precharge Buffers, Only Connected for a Short Period of Time to Minimize the Input Currents Due to the Chopping Mechanism

Figure 7. Capacitive PGA (some blocks have been removed for clarity).

There are some benefits associated with this capacitive architecture, as compared with the resistive one:

It offers a better noise vs. power trade-off, as it contains fewer noise sources. Fewer amplifiers are needed and capacitors don't contribute noise as opposed to the resistors.

Capacitors offer a wide range of advantages over resistors. Apart from being noiseless, they don't suffer from self-heating and normally offer a better matching and temperature drift. This has a positive impact on offset, gain error, and drift specifications.

The capacitors decouple the input common mode from the rest of the signal chain common mode. This offers an advantage in terms of CMRR, PSRR, and THD.

One of the most powerful advantages is the fact that the capacitive PGA input common-mode range may be rail-to-rail and beyond. That gives the possibility to bias the sensor common-mode voltage pretty much anywhere from the positive rail down to the negative rail.

This capacitive architecture combines the benefits of an instrumentation amplifier, which is a really high input impedance, as the input impedance is a capacitor, with the benefits of capacitors over resistors as the gain element, increasing the dynamic range of the amplifier, not only in terms of signal swing but also noise efficiency.

A common solution to overcome the resistive PGA common-mode limitation is to increase or shift the power rails or, alternatively, to recenter the sensor signal common mode. This comes at the expense of higher power consumption, supply design complexity, additional external components, and cost.

Practical Examples

In a Wheatstone bridge, the common-mode voltage is defined by the impedance connected in each of the legs and is proportional to the applied power supply. Weigh scale applications implement this sensing topology due to the benefit of linear sensing in strain gages. Figure 8 shows a half-bridge type II.
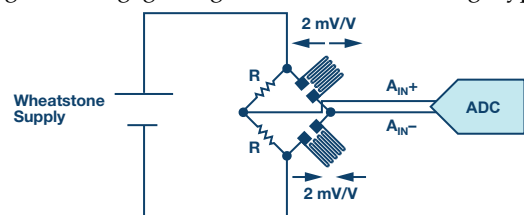

Figure 8. Weigh scale with strain gage in Wheatstone topology.

Table 1. Comparative of Resistive and Capacitive PGA in a Wheatstone Bridge Assuming Standard Supplies and Gains

PGA	Resistive PGA	Capacitive PGA
Maximum Wheatstone Supply	3.3 V	6 V
Strain Gage Differential Sensitivity	3.3 mV	6 mV
Dynamic Range Improvement (dB)		5.2 dB

The sensitivity of a strain gage is typically 2 mV/V. The higher the Wheatstone supply, the higher the sensitivity obtained. To increase the dynamic range of the strain gage and maximize the SNR, the bridge may be powered at higher supplies than the ADC.

In a resistive PGA, due to its common-mode limitation, the bridge should be powered at the same supply voltage as the ADC supply in order to maximize the dynamic range, while in a capacitive PGA the bridge can be powered at almost twice the ADC supply voltage as there is no input common-mode limitation.

For example, assuming standard supply levels and powering the ADC at 3.3 V, the improvement of a capacitive PGA over a resistive PGA for the same selected gain can be summarized in Table 1.

Another probable issue is a potential difference between grounds when the bridge is connected at some distance from the ADC. This could shift the common-mode voltage, unbalance the ADC input common mode with respect to the bridge, and reduce the maximum allowable gain in the resistive PGA.

A possible way to match the capacitive PGA performance with the resistive PGA is by powering the bridge at a higher supply voltage. For instance, powering the bridge with a bipolar supply, ±3.3 V, to increase the sensitivity to the strain gage at the expenses of increased system complexity and power dissipation.

Another case example that could benefit from a capacitive PGA is a temperature measurement using resistance temperature detectors (RTD) or thermocouples.

A popular RTD resistor, such as the PT100, may be used to sense the temperature directly or indirectly sensing the cold junction of a thermocouple, as shown in Figure 9.
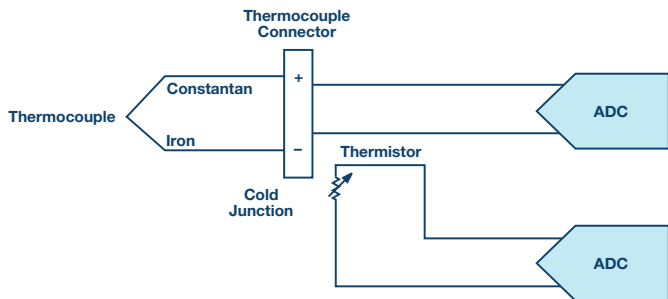

Figure 9. Typical thermocouple setup.

The PT100 is offered with different wires per element, being the most popular and cost-effective 3-wire configuration.

A conventional way to measure the temperature while canceling the lead error is proposed in Figure 10. In this example the internal current sources of the AD7124-8 Σ-Δ ADC with PGA, drive two wires of the RTD with the same current, generating an offset error equal in both leads and proportional to the lead resistance.

Due to the small value of the lead resistance and the currents provided by the AD7174-8 to minimize the self-heating effect, the offset voltage generated in RL3 is close to the negative rail, significantly reducing the maximum allowable gain in a resistive PGA as its input common mode would also be very close to the rails as opposed to a capacitive PGA that will internally set the common-mode voltage to half of the supply rails, allowing for a higher gain configuration and, therefore, increasing the overall dynamic range.

The proposed solution significantly reduces the complexity of the system and the hardware connections, as the third cable should not be returned to the ADC PCB and can be connected to ground near the RTD location.
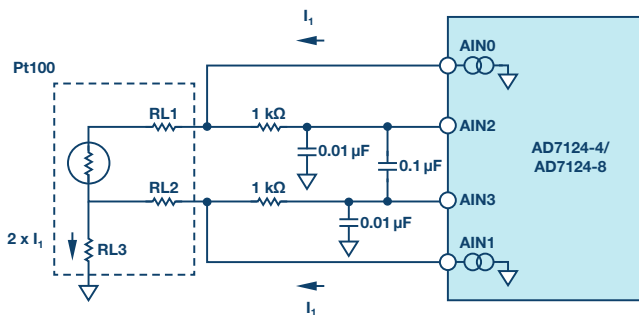
Figure 10. 3-wire RTD measurement.

To increase the precision on the temperature measurement, 4-wire measurements are preferred. In this case, only one current reference is used. To avoid the imprecision on the current source, a ratiometric measurement with a precision resistance used as the ADC reference voltage generator can be used, as shown in Figure 11.

Figure 11. Ratiometric 4-wire RTD measurement.

The value of the external precision resistor is chosen so that the maximum voltage generated across the RTD equals the reference voltage divided by the PGA gain.

Table 2. Comparative of Resistive and Capacitive PGA in a 4-Wire RTD Ratiometric Measurement

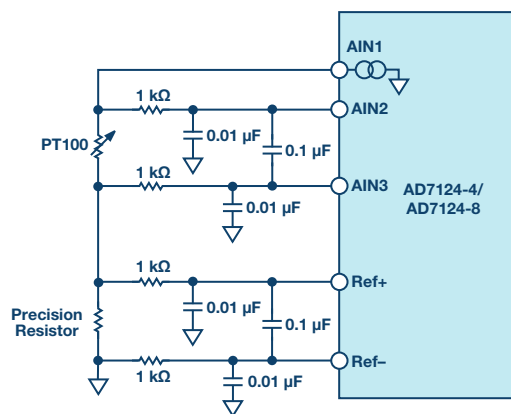| PGA	Resistive	Capacitive
Pt100 Output Voltage	500 μA × 313.7 = 156.85 mV	500 μA × 313.7 = 156.85 mV
VREF	1.65 V	2.5 V
Maximum PGA Gain	1.65 V/156.85 mV ≈ 11	2.5 V/156.85 mV ≈ 16
Improvement (dB)		3.6

For a 3.3 V supply, in a resistive PGA, the voltage generated on the precision resistance should be around 1.65 V, otherwise the PGA common-mode voltage will limit the maximum gain. The consequence is that the maximum gained signal should be equal to 1.65 V. In a capacitive PGA, there is no input common-mode limitation, therefore, the RTD common-mode signal can sit close to the top rail, which maximizes the ADC reference voltage generated by the precision resistor, and hence the highest selectable gain and the dynamic range can be maximized.

Table 2 summarizes the maximum gain of a resistive PGA against a capacitive PGA, with a maximum current source of 500 μA to limit the Pt100 self heating, assuming a Class B RTD, at a maximum temperature of 600°C, a maximum VREF of 2.5 V.

Conclusion

The capacitive PGA offers an important number of advantages compared with the resistive PGA. Critical specifications such as noise, common-mode rejection, offset, gain error, and tempera-ture drift are improved due to the inherent temperature stability and matching properties of the capacitors as gain elements.

Another key feature is the decoupling of the input common-mode voltage from the amplifier internal common-mode voltage. This is critical when the input signal to be amplified sits on a common-mode voltage close to the rails. The resistive PGA selected gain would be severely restricted by its common-mode limitation, or it would require higher supply rails or external components to rebias the input signal to half of the rails. On the contrary, the capacitive PGA could handle this sensing scenario easily.

Some of the latest Σ-Δ ADC products that include a capacitive PGA are the AD7190, AD7124-4, AD7124-8, and AD7779.

Miguel Usach Merino

Also by this Author:
Powering ICs On and Off
Volume 49, Number 2

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

Gerard Mora–Puchalt

Gerard Mora-Puchalt [gerard.mora–puchalt@analog.com] received his M.S. degree in tele–communications engineering from University Polytechnic of Valencia in 2005. He completed his Final Year Project at Analog Devices with the DAC Applications Group in 2005 and joined ADI after graduation in 2006 in the Temperature Sensor Design Evaluation Team. He moved to the Precision Converters Group in 2007, where he has been working as an analog design engineer ever since. He is based in Valencia, Spain.

Where Zero–IF Wins: 50% Smaller PCB Footprint at ⅓ the Cost

By Brad Brannon

Introduction

Zero-IF (ZIF) architecture has been around since the early days of radio. Today the ZIF architecture can be found in nearly all consumer radios, whether television, cell phones, or Bluetooth® technology. The key reason for this wide adoption is that it has proven time and again to offer the lowest cost, lowest power, and the smallest footprint solution in any radio technology. Historically, this architecture has been withheld from applications that demand high performance. However, with the demand for wireless growing around us and the rapidly crowding spectrums, a change is required in order to continue economically deploying radios in the infrastructure that supports our wireless needs. Contemporary zero-IF architectures can satisfy these needs as many of the impairments normally associated with these architectures have been resolved through a combination of process, design, partitioning, and algorithms. New advances in ZIF technology challenge the current high performance radio architectures and introduce new products with breakthrough performance to enable new applications previously beyond the reach of ZIF. This article will explore the many benefits of ZIF architectures and introduce new levels of performance that they bring to radio designs.

Challenges of the Radio Engineer[1]

Today's transceiver architect is challenged by a growing list of demands driven by our ever increasing requirements for wireless devices and applications. This leads to the continual need to access more bandwidth.

The designer has moved over the years from a single carrier radio to multicarrier. As spectrum becomes fully occupied in one band, new bands are allocated; now there are more than 40 wireless bands that must be served. Because operators have spectrum in multiple bands and these resources must be coordinated, the trend is toward carrier aggregation, and carrier aggregation leads to multiband radios. This all leads to more radios, with higher performance, requiring better out-of-band rejection, improved emissions, and less power dissipation.

While the demand for wireless is rapidly increasing, power and space budgets are not. In fact, with an ever increasing need to economize both in power and space, reducing both the carbon footprint and the physical footprint are very important. To achieve these goals, a new perspective on radio architectures and partitioning is required.

Integration

In order to increase the number of radios in a particular design, the footprints must be made smaller for each radio. The traditional way to do this is to progressively integrate more and more of the design onto a single piece of silicon. While this may make sense from a digital perspective, integration of analog functionality for the sake of integration doesn't always make sense. One reason is that many analog functions in a radio cannot effectively be integrated. For example, a traditional IF sampling receiver is shown in Figure 1. There are four basic stages to an IF sampling architecture: low noise gain and RF selectivity, frequency translation, IF gain and selectivity, and detection. For selectivity, SAW filters are typically used. These devices cannot be integrated and therefore must be off chip. While RF selectivity is provided by piezoelectric or mechanical devices, occasionally LC filters are used for the IF filter. While LC filters may occasionally be integrated on monolithic structures, the compromise in both filter performance (Q and insertion loss) and the required increase in sample rate of the digitizer (detector) increase the overall dissipation.
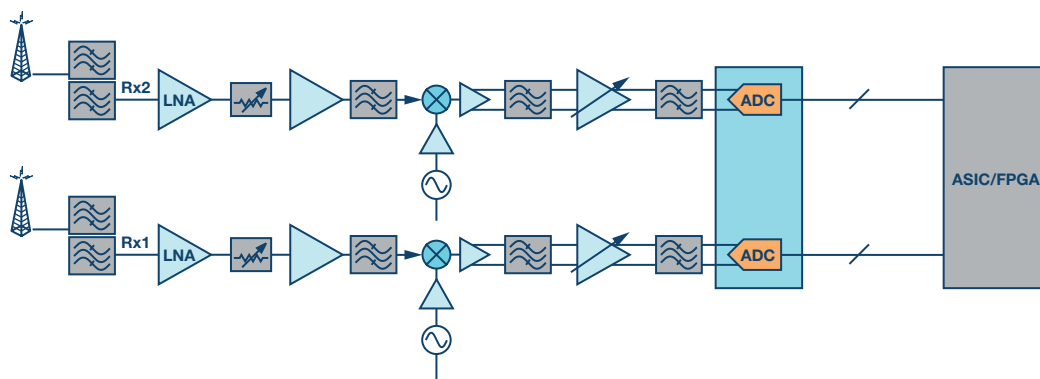

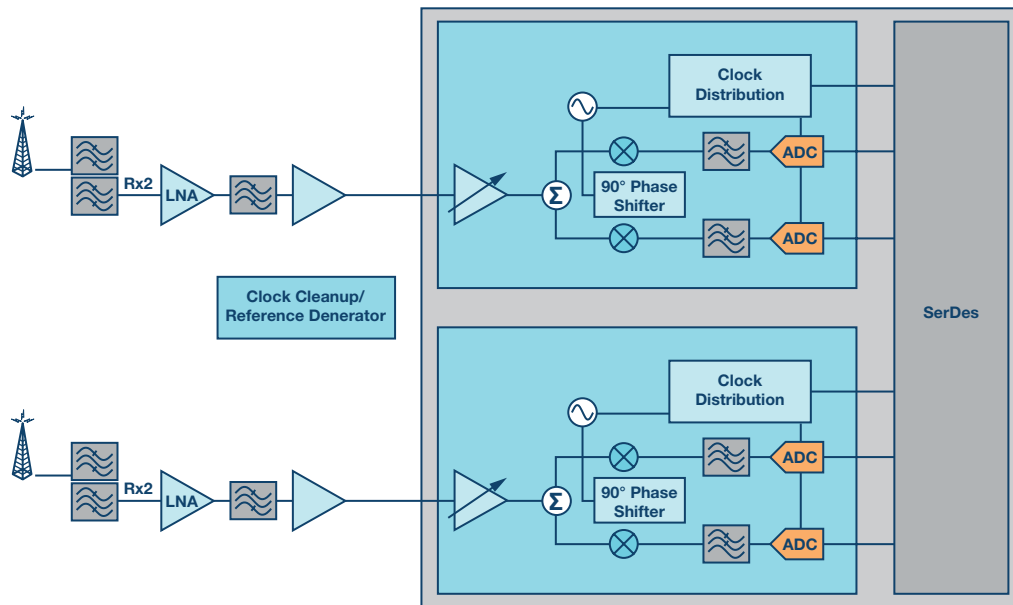*Figure 1. Traditional IF sampling receiver.*

Figure 2. Typical zero–IF sampling receiver.

Digitizers (analog-to-digital converters) must be done on low cost CMOS processes to keep the cost and power reasonable. While they certainly can be fabricated on bipolar processes, this results in both larger and more power hungry devices, which runs counter to optimization for size. Thus standard CMOS is the desired process for this function. This becomes a challenge for integration of high performance amplifiers, particularly the IF stage. While amplifiers can be integrated on CMOS processes, it is difficult to get the performance required from processes that are optimized for low power and low voltage. Furthermore, integrating the mixer and IF amplifier on chip requires that the interstage signals be routed off chip to access the IF and antialias filters prior to being digitized, foregoing much of the benefit of integration. Doing so is counterproductive to integration as it increases the pin count and package size. Additionally, each time critical analog signals pass through a package pin, a compromise in performance is made.

The optimal way to integrate is to repartition the system to eliminate the items that cannot be integrated. Since SAW and LC filters cannot be effectively integrated, the best option is to determine how to get rid of them by re-architecting. Figure 2 shows a typical zero-IF signal chain that achieves these goals by translating the RF signal directly to a complex baseband, completely eliminating the need for an IF filter and IF amplifiers. Selectivity is achieved by introducing a pair of low-pass filters into the I/Q baseband signal chain that can be integrated as active low-pass filters instead of off chip lossy fixed IF devices. Traditional IF SAW filters or LC filters are by nature fixed while these active filters can be electronically tuned often from the hundreds of kHz range through hundreds of megahertz. Changing the bandwidth of the baseband allows the same device to cover a broad range of bandwidths without having to change a bill of material or switching between different fixed IF filters.

Although not intuitive from the figure, zero-IF receivers can also cover a very broad range of RF frequencies simply by changing the local oscillator. Zero-IF transceivers provide a truly broadband experience with typical coverage continuously from several hundred MHz up to around 6 GHz. Without fixed filters, truly flexible radios are possible, greatly reducing and possibly eliminating the effort required to

develop band variations of the radio design. Because of the flexible digitizers and programmable baseband filters, zero-IF designs not only deliver high performance, but also significant flexibility in adopting to a wide range of frequency and bandwidths while maintaining nearly flat performance without the need to optimize analog circuits (such as filters) for each configuration—true software-defined radio (SDR) technology. This too adds greatly to the reduction of footprint by elimination of banks of filters for applications that must cover multiple bands. In some cases, the RF filter may be completely eliminated, introducing a completely wideband radio that requires virtually no effort to change bands. By elimination of some devices and integration of others, the required PCB footprint for a zero-IF design is greatly reduced, not only simplifying the rebanding process, but also reducing the effort to change the form factor when required.

Smallest Footprint

A direct comparison of the PCB area for each of these architectures (Figure 3 and Figure 4) shows that for a dual receive path, the respective PCB area for a reasonable implementation gives 2880 mm^2 (18 mm × 160 mm) for IF sampling and 1434 mm^2 (18 mm × 80 mm) for zero-IF sampling. Not counting the potential elimination of RF filters and other simplifications,[2] the zero-IF architecture offers the possibility of reducing the radio footprint by up to 50% as compared to current IF sampling technology. Future generation designs can potentially redouble these savings with additional integration.
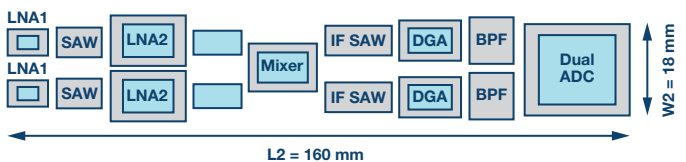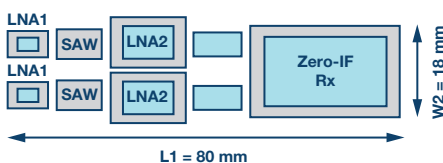

Figure 3. Typical IF sampling layout.

Figure 4. Typical zero–IF sampling layout.

Lowest Cost

From a direct bill of material point of view, the savings when moving from an IF sampling system to a zero-IF architecture are 33%. Cost analysis is always difficult. However, a thorough examination of Figure 1 and Figure 2 shows that many of the discrete items are eliminated, including the IF and antialias filtering, and that the mixer and baseband amplifiers are integrated. What is not obvious is that because zero-IF receivers inherently offer out-of-band rejection not offered in traditional IF sampling architectures, the overall external filtering requirements are greatly reduced. There are two contributors within the zero-IF architecture that drive this. The first is the active baseband filter that provides both in-band gain and out-of-band rejection. The second is the high sample rate low-pass Σ-Δ converter used to digitize the I/Q signals. The active filter reduces the out-of-band component while the high sample rate of the ADC moves the alias point out to a sufficiently high frequency that external antialiasing filtering is not required (because the active filter has sufficiently rejected the signals).
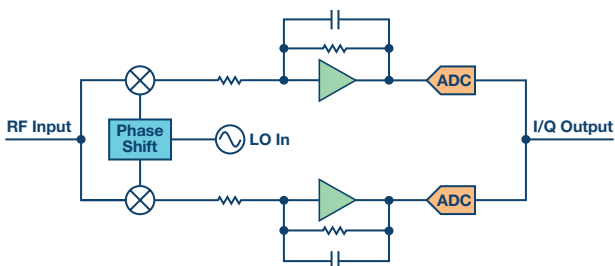

Figure 5. Active baseband filter and ADC.

By applying the baseband signals to an active filter, as in Figure 5, high frequency content is rolled off. The ADC then digitizes and ultimately filters any residual output from the low-pass filter. The cascaded results are shown in Figure 6. This figure shows what a typical receiver performance might look like with the compound effect of an active filter and Σ-Δ ADC. Shown here is a typical 3 dB desense of both in-band and out-of-band power. Note the improvement in out-of-band performance without any external filtering.

For similar levels of performance, IF sampling receivers rely on discrete IF filtering such as SAW technology for selectivity and protection from out-of-band signals and to prevent aliasing of wideband signals and noise alike from aliasing back in band. IF sampling architectures must also be protected from other unwanted mixer terms including the half-IF term, which drives additional RF and IF filtering requirements as well as restricts sample rates and IF planning. The zero-IF architecture has no such frequency planning restrictions.
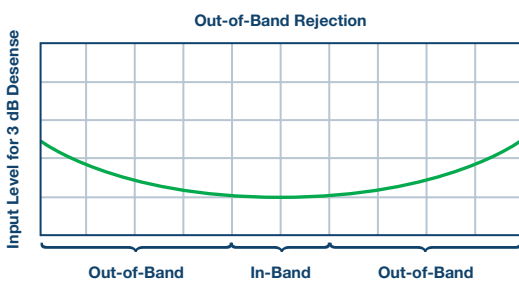

Figure 6. Typical zero–IF out–of–band rejection.

Depending on the design and application, this native rejection reduces or eliminates external RF filtering requirements. This results in a direct savings by their omission as external RF filters can be relatively expensive depending on the type. Secondarily, removal of these lossy devices may allow the elimination of RF gain stages, saving not only cost but reducing power and improving linearity. All of these add to the savings delivered by repartitioning and smart integration.

As noted, it is difficult to assess cost as this depends greatly on volume and vendor agreements. However, a detailed analysis shows that zero-IF architectures typically reduce the full system cost by up to ⅓ through the impact of integration, elimination, and reduction in requirements. It is important to remember that this is system cost and not device cost. Because more functions are being placed in fewer devices, some device costs may increase while overall system costs are reduced.

Beyond bill of material costs, the integrated zero-IF receiver addresses a few other areas. Because integrated systems reduce the number of devices in the system, assembly costs are lower and factory yields are higher. Because there are fewer discrete devices, alignment time is shorter. These items together reduce factory costs.

Because the zero-IF receiver is truly wideband, engineering costs are reduced to reband. IF frequencies must be carefully chosen in IF sampling systems, but with zero-IF systems, there is no careful planning required. New bands may be added largely by changing the local oscillator. Additionally, because many applications do not require an external RF filter when zero-IF is used, further simplifications may result. Overall cost savings can be substantial when considering a zero-IF solution when the direct cost is considered alongside the manufacturing and engineering costs outlined above.

Lowest Power

Simply taking an architecture like that shown in Figure 1 and directly integrating it into a system on chip will not result in a power or cost savings. Power savings come through selecting an efficient architecture that can be optimized for the process on which it is targeted. Architectures like the IF sampling receiver shown involve a lot of high and midrange frequencies that are difficult to scale on low cost processes and therefore require significant amounts of power be dissipated to support the frequencies required. However, the zero-IF architecture as shown in Figure 2 works to immediately reduce the frequencies of interest to dc (baseband), allowing implementation of the lowest frequency circuits possible.

Similarly throwing bandwidth at the problem is also inefficient. Architectures like direct RF sampling provide wide bandwidths with a lot of flexibility. However, adding bandwidth to a system always adds extra power to the problem as documented by both Walden[3] and Murmann.[4]

Unless the raw bandwidth is required, addressing the problem with bandwidth alone doesn't provide an economical solution for most receiver applications. Data from these long-term studies show two regions of converter development. The technology front documents advances in technology that provide for meaningful increases in core ac performance in the form of dynamic range and bandwidth. The architecture front

documents advances in overall core architecture efficiencies. Typically the curve moves first to the right and then upward as designs are optimized. For communications applications, operation tends to be along the technology front where the slope of the line is about 10 dB per decade reduction in converter efficiency as shown in Figure 7. At this slope, doubling the bandwidth results in dissipating about three times the power. However, by the time these cores are integrated into functional devices, the efficiency has improved and typically carries a power penalty closer to two as it moves closer toward the architectural front.
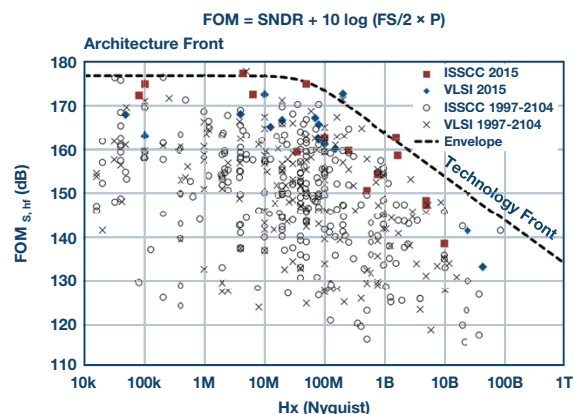

Figure 7. Figure of merit for core ADC technology.[4]

The conclusion for applications that are concerned about power is that the lowest power solution is one where bandwidth and sample rate are optimized for the application. Zero-IF sampling with Σ-Δ converters are optimized for such applications. Depending on the specific implementation, power savings implementing a zero-IF receiver may be 50% or more reduced compared to an IF sampling architecture and as much as 120% compared to direct RF sampling.

Power is also directly related to cost. Not only does higher power drive more expensive packaging and supply generation, but for each watt a circuit dissipates, at 12 cents per kW/hr, the cost of operation is more than $1 per year per watt. Given the low cost of many electronic devices, the power to run them for just a year can easily be more than their direct cost. Thus, as options for integrated radio solutions become available, applications that are sensitive to cost and power must choose the trade-offs carefully. Selecting architectures that unnecessarily increase dissipation may not only increase the power, but may also impact long-term operating costs of the solution.

Performance Enhancements

For a radio design, there are a number of key metrics that are considered important. These include specifications including noise figure (NF), linearity (IP3, IM3), desensitization, and selectivity to name a few. Beyond the normal radio specifications, there are additional specifications that are important but are often hidden from most users. These include specification distribution and drift as a function of time, supply, temperature, and process. Zero-IF architectures meet these and other key requirements for radio design.

Tracking by Temperature, Supply, and Process

One of the benefits of a fully integrated transceiver architecture is that device matching can be much better for a properly designed radio, not just initially, but devices can track effectively over process, temperature, supply, and frequency when properly designed. Any residual mismatch is readily removed with signal processing techniques that are typically embedded in these integrated solutions. While this is very typical of IC design, what is different about integration of the radio is that because all frequency dependent items are on chip with a zero-IF design, they too can be made to track. A typical radio as shown in Figure 1 includes an IF filter off chip. The characteristics of the IF filter will change as a function of time, temperature, or device-to-device, which will be uncorrelated to anything on chip and cannot be followed. However, one of the major advantages of integration of the filter is that because it is constructed with on-chip devices, devices can be scaled or made to ratiometrically track one another to keep performance stable. Those items that cannot be stabilized by design can easily be calibrated. The end result is that when budgeting device variations, much less margin is required than for a discrete design where all devices are uncorrelated.

For example, it is not uncommon to allocate NF variation of 1 dB for a mixer, IF filter, IF amplifier, and ADC each. When budgeting performance, these variations must be cascaded. However, in an integrated design where all critical specifications either track one another or are calibrated out, the result is a single device variation of 1 dB greatly simplifying signal chain variation. This can have a significant impact in a design as compared to a design with uncorrelated terms which would otherwise require extra system gain to offset the potential increase in noise—impacting cost, power, and linearity for the end product. In an integrated design such as that in Figure 2, the total variation in performance is considerably smaller than an uncorrelated design and therefore smaller system gain is required.

Advanced Correction Techniques

Zero-IF receivers typically have two areas that have caused concern in the past. Because complex data is generated and represented with a pair of real cascaded networks representing the real and imaginary components, errors are generated that represent gain, phase, and offset of the individual signal chains as represented in Figure 8.
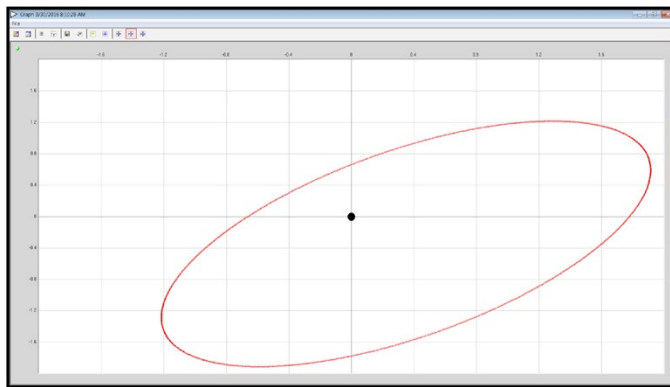

Figure 8. Quadrature errors showing gain, phase, and offset terms.

These errors manifest as images in the spectrum and are what has typically prevented these architectures from being more widely adopted. However, as an integrated solution, these artifacts can be easily controlled by both analog optimization and digital correction. Figure 9 shows a typical uncorrected representation of the complex data. Here both the LO leakage (and dc offset) and image rejection (quadrature error) can be seen.
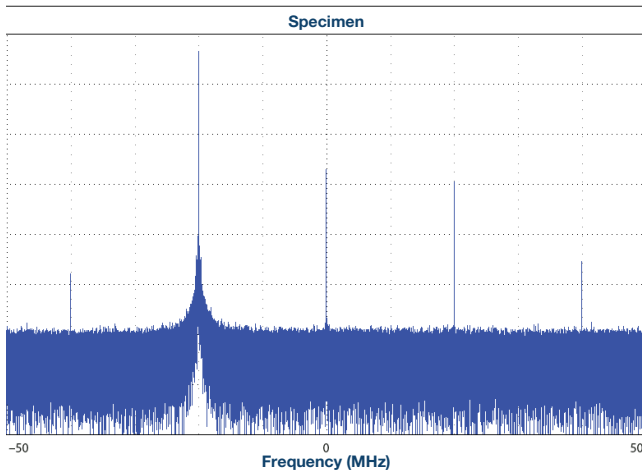

Figure 9. Typical uncorrected LO leakage and image rejection.

LO Leakage Control

LO leakage shows up as increased dc offset in the I or Q signal paths. This occurs as a result of LO coupling directly into the RF signal path and being coherently downconverted to the output. The result is a mixer product that appears as a dc offset that adds to any residual dc offset in the signal chain. A good zero-IF architecture will automatically track and correct for these errors both initially and as they shift over time, temperature, supply, and process resulting in performance better than –90 dBFS as shown in Figure 10.
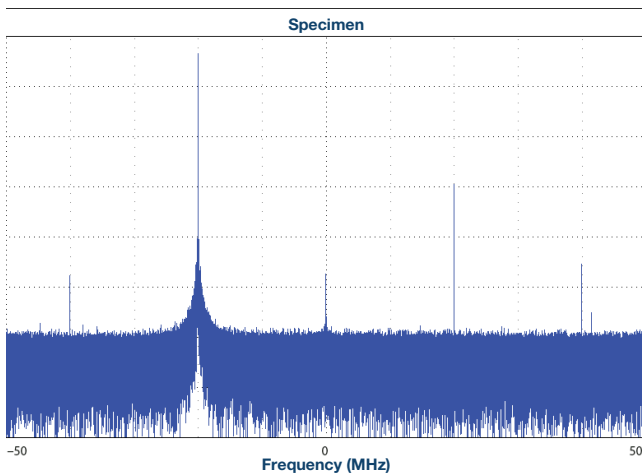

Figure 10. Typical LO leakage control.

QEC

To prevent images from disrupting performance, quadrature error correction (QEC) will typically be implemented. Figure 11 shows the impact such a function can make. In this example, the image improves to better than –105 dBc, which is more than adequate for most wireless applications. For both LO leakage and QEC, tracking is employed to ensure that as performance shifts over time, the corrections stay current ensuring that optimal performance is always achieved.
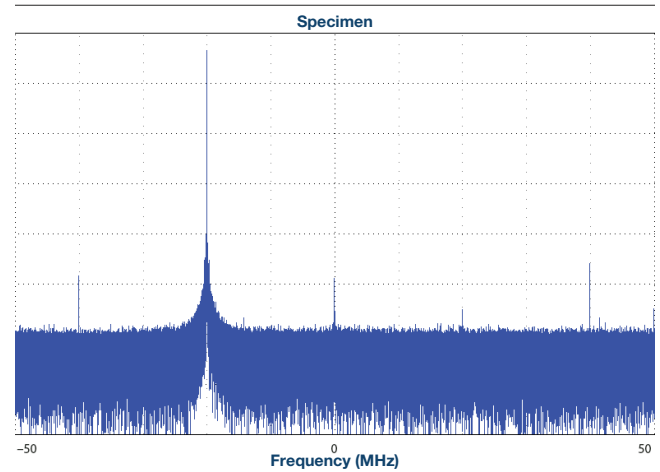

Figure 11. Typical quadrature correction with LO leakage control.

Quadrature error and LO feedthrough matters in a radio system. If the error is large enough, images of large blockers could mask out smaller desired signals. In Figure 12 the image of a large blocker falls at 15 MHz while a desired signal is centered at 20 MHz. If the image had fallen partly or wholly onto the desired signal, it would degrade the SNR of the desired signal resulting potentially in errors in the demodulation. Typically systems like LTE and W-CDMA have a reasonable tolerance for these kinds of images but are not totally immune. Typically these systems require image rejection of 75 dBc or better, which as shown in Figure 11 is easily met and maintained with a zero-IF architecture.
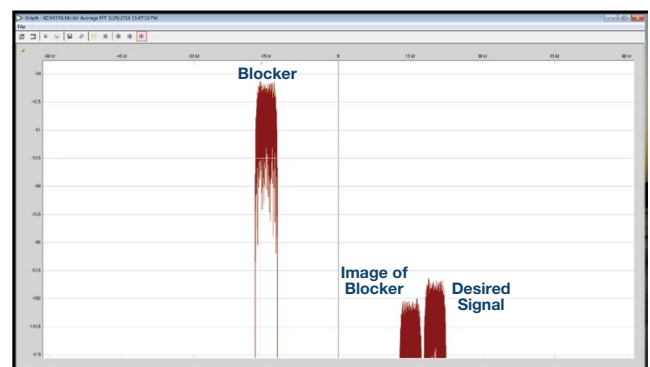

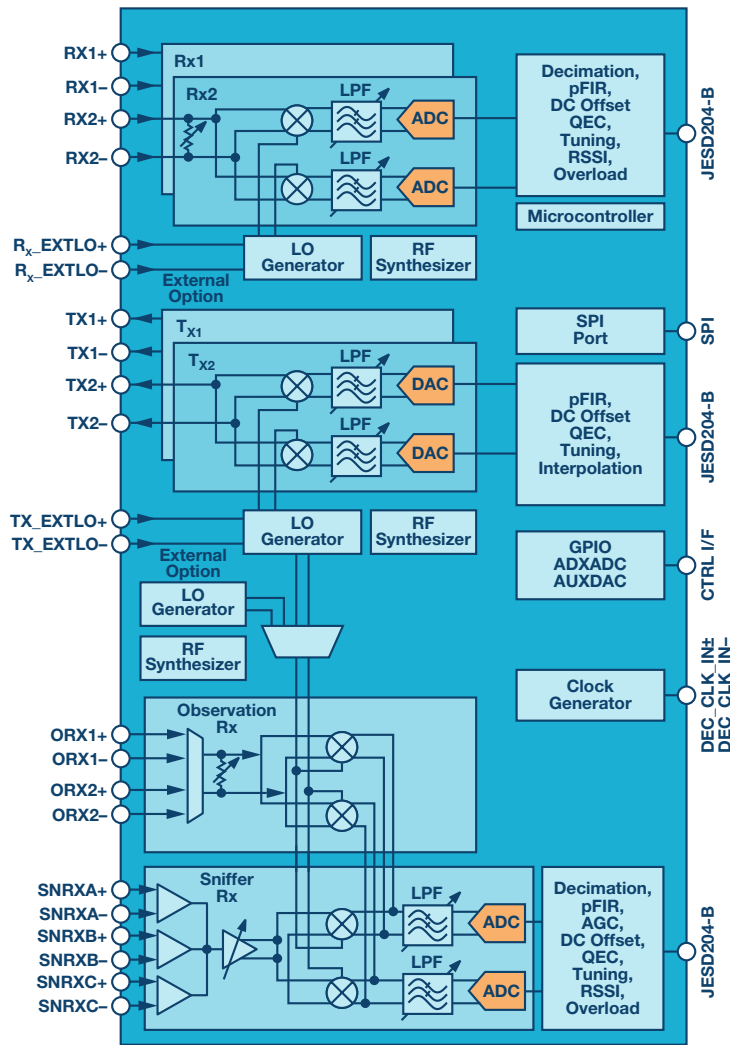Figure 12. Example of an image blocking a desired signal.

Figure 13. AD9371 integrated zero–IF transceiver.

AD9371

A typical example of zero-IF transmit and receive is the AD9371. As shown in Figure 13, the AD9371 provides a very high level of integrated functionality including dual transmit, dual receive along with additional functionality including an observation and sniffer receiver as well as integrated AGC, dc offset correction (LO leakage control), and QEC. The product offers wide RF coverage from 300 MHz to 6 GHz. Each transmitter can cover between 20 MHz and 100 MHz of synthesis bandwidth while each receiver is capable of between 5 MHz and 100 MHz. While this device is targeted at 3G and 4G applications, it is an ideal solution for many other general-purpose radios and software-defined applications up to 6 GHz.

The AD9371 offers a complete system integration including all of the frequency dependent devices discussed earlier as well as all of the calibration and alignment functionality in a 12 mm × 12 mm BGA package. Adding to the receive func-

tion from Figure 4, Figure 14 includes the required transmit functionality to the footprint to yield a very compact dual transceiver design. Power depends on the exact configuration including bandwidth and features enabled, but typical dissipation of the AD9371 is only 4.86 W, including the digital functionality to maintain LO leakage and image rejection.
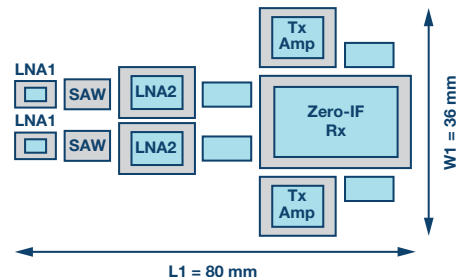

Figure 14. Typical zero–IF transceiver layout.

Key AD9371 Performance

Noise Figure

Figure 15 and Figure 16 show the typical NF characteristics of the AD9371. The first figure shows a broad sweep of RF frequency and that the NF is relatively flat across this spectrum. The input structure for this device is in the form of an attenuator so that the NF increases dB for dB. Assuming a worst-case NF of 16 dB with zero attenuation, allowing about 4 dB attenuation for external gain variations, a total NF of 20 dB could be assumed. An external LNA (0.8 dB) providing at least 24 dB of gain would provide a system NF of 2 dB.
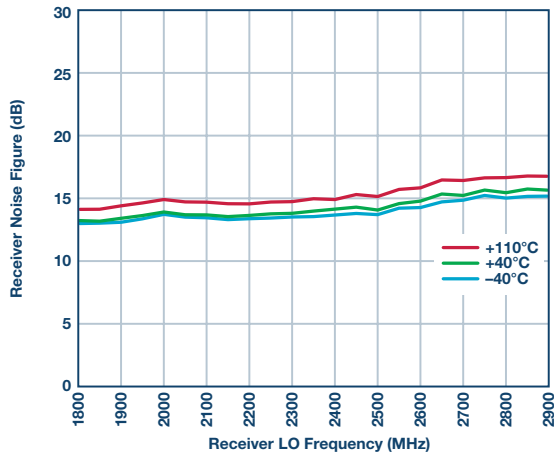

Figure 15. AD9371 NF with 0 dB attenuation and 40 MHz BW.

Figure 16 shows NF as a function of out-of-band blockers relative to the input of the AD9371. Assuming 24 dB of external gain, 0 dBm relative to the input of this device would occur at –24 dBm relative to the antenna connector. Considering only the impact of the AD9371, the overall NF degradation would be about 1 dB for a 3 dB degradation to the integrated receiver.
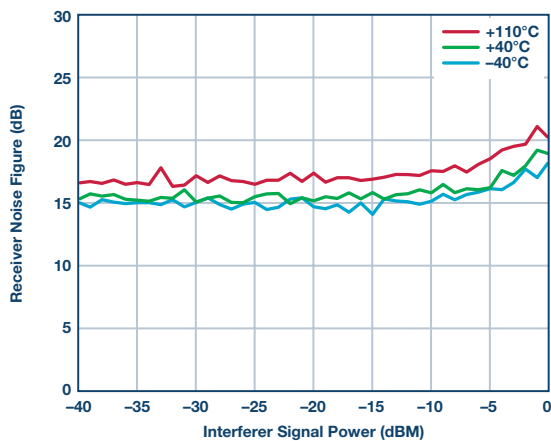

Figure 16. AD9371 NF vs. out–of–band signal power.

Image Rejection

Similar to LO leakage, receive image rejection may be estimated by the information in Figure 17. With a typical input level at the antenna of –40 dBm, the image can be estimated to be better than 80 dB lower or –120 dBm relative to the antenna port.
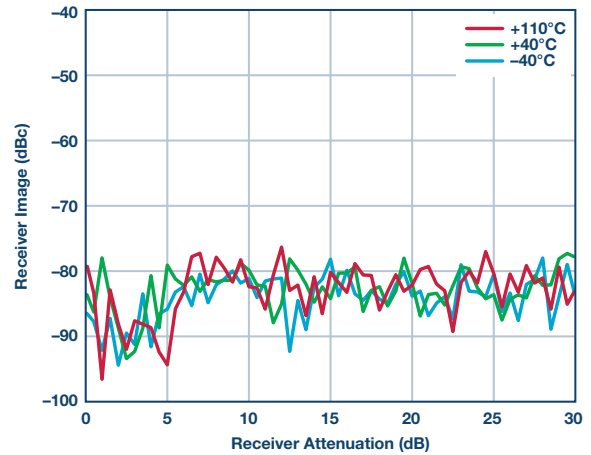

Figure 17. Receiver image rejection.

Conclusion

While historically zero-IF architectures have been confined to low performance applications, new products like the AD9371 offer game changing performance. Not only do these devices offer performance in line with IF sampling receivers, they go one step further by repartitioning the radio such that a more robust architecture is created that not only reduces manufacturing cost, but reduces the cost of operation once deployed. No longer does radio performance have to be compromised for a low solution cost design allowing users to focus time and resources on developing the application and not the radio implementation.

References

[1] While this discussion primarily focuses on receivers, this discussion applies to transmitters as well. For transmitters, zero-IF has been the accepted architecture of high performance for more than a decade.

[2] As examined here, the typical zero-IF receiver also includes a full transmit path (AD9371) within the same package.

[3] R. H. Walden. "Analog-to-Digital Converter Survey and Analysis." *IEEE Journal on Selected Areas in Communications*, April 1999.

[4] Boris Murmann. "ADC Performance Survey 1997-2015." Stanford University, 2015.

Brad Brannon

Brad Brannon [brad.brannon@analog.com] has worked at Analog Devices for 32 years following his graduation from North Carolina State University. At ADI he has held positions in design, test, applications, and in system engineering. Brad has authored a number of articles and application notes on topics that span clocking data converters, designing radios, and testing ADCs. Currently Brad is responsible for system engineering for 4G and 5G receive architectures.

RF Transceivers Provide Breakthrough SWaP Solutions for Aerospace and Defense

By Wyatt Taylor and David Brown

Next-generation aerospace and defense platforms introduce new challenges that require solutions beyond what can be achieved through individual device optimization. Integrating more software control and cognitive abilities to the radio demands a more frequency and bandwidth flexible RF design. To achieve this goal, static filters need to be removed and replaced with tunable filters. Similarly, the concept of a common platform would allow for shorter development times, enable reduced manufacturing costs, and provide greater interoperability between systems. The common platform demands that the RF system be capable of providing full performance for applications that traditionally had very different architectures. Finally, future platforms are pushing size and power demands to a new extreme.

Handheld soldier radios are becoming more capable and complex, but simultaneously requiring improved battery efficiency. Small UAVs lack the power generation of large aircraft and every milliwatt that the RF system consumes directly translates to payload battery weight and, thus, reduced flight time. To overcome these challenges and create the next generation of aerospace and defense solutions, a new radio architecture is needed.

Superheterodyne Architecture and Diminishing Returns

Since its inception, the superheterodyne architecture has been the backbone of radio design for aerospace and defense systems. Whether it is a soldier radio, unmanned aerial vehicle (UAV) data link, or a signal intelligence (SIGINT) receiver,

the single or dual mixing stage superheterodyne architecture is the common choice. The benefits of this design are clear: proper frequency planning can allow for very low spurious emissions, the channel bandwidth and selectivity can be set by the intermediate frequency (IF) filters, and the gain ditribution across the stages allows for a trade off between optimizing the noise figure and linearity.

For over 100 years of use, there have been significant gains in performance for the superheterodyne across the entire signal chain. Microwave and RF devices have improved their performance while decreasing power consumption. ADCs and DACs have increased the sample rate, linearity, and effective number of bits (ENOB). Processing capability in FPGAs and DSPs has followed Moore's law and increased with time, allowing for more efficient algorithms, digital correction, and further integration. Package technology has shrunk device pin density while simultaneously improving thermal handling.

However, these device specific improvements are beginning to reach the point of diminishing returns. While the RF components have followed a reduced size, weight, and power (SWaP) trend—high performance filters remain physically large and are often custom designs, adding to overall system cost. Additionally, the IF filters set the analog channel bandwidth of the platform, making it difficult to create a common platform design that can be reused across a wide range of systems. For package technology, most manufacturing lines will not go below a 0.65 mm or 0.8 mm ball pitch, meaning there is a limit on how physically small a complex device with many I/O requirements can become.
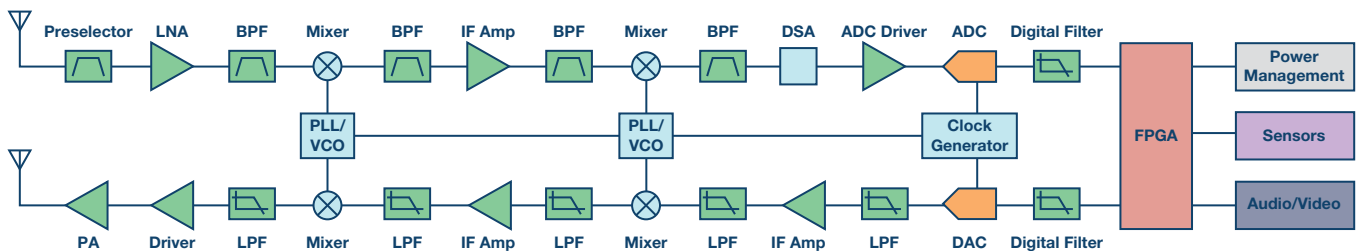

Figure 1. Basic superheterodyne architecture.

Zero–IF Architecture

An alternative to the superheterodyne architecture, which has reemerged as a potential solution in recent years, is the zero-IF (ZIF) architecture. A ZIF receiver utilizes a single frequency mixing stage with the local oscillator (LO) set directly to the frequency band of interest, translating the received signal down to baseband in phase (I) and quadrature (Q) signals. This architecture alleviates the stringent filtering requirements of the superheterodyne since all analog filtering takes place at baseband, where filters are much easier to design and less expensive than custom RF/IF filters. The ADC and DAC are now operating on I/Q data at baseband, so the sample rate relative to the converted bandwidth can be reduced, saving significant power. For many design aspects, ZIF transceivers provide significant SWaP reduction as a result of reduced analog front-end complexity and component count.
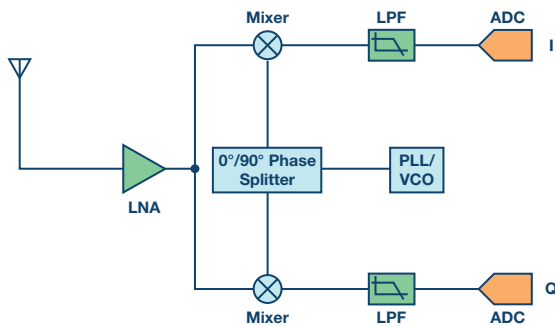

Figure 2. Zero–IF architecture.

However, there are drawbacks to this system architecture that need to be addressed. This direct frequency conversion to baseband introduces a carrier leakage and an image frequency component. Mathematically, the imaginary components of I and Q signals cancel out due to their orthogonality (Figure 3). Due to real-world factors, such as process variation and temperature deltas in the signal chain, it is impossible to maintain a perfect 90° phase offset between the I and Q signals, resulting in degraded image rejection. Additionally, imperfect LO isolation in the mixing stage introduces carrier leakage components. When left uncorrected, the image and carrier leakage can degrade a receiver's sensitivity and create undesirable transmit spectral emissions.

Historically, the I/Q imbalance has limited the range of applications that were appropriate for the ZIF architecture. This was due to two reasons: first, a discrete implementation of the ZIF architecture will suffer from mismatches both in the monolithic devices and also the printed circuit board (PCB). In addition to this, the monolithic devices could pull from different fabrication lots, making exact matching very difficult due to native process variation. A discrete implementation will also have the processor physically separated from the RF components, making a quadrature correction algorithm very difficult to implement across frequency, temperature, and bandwidth.
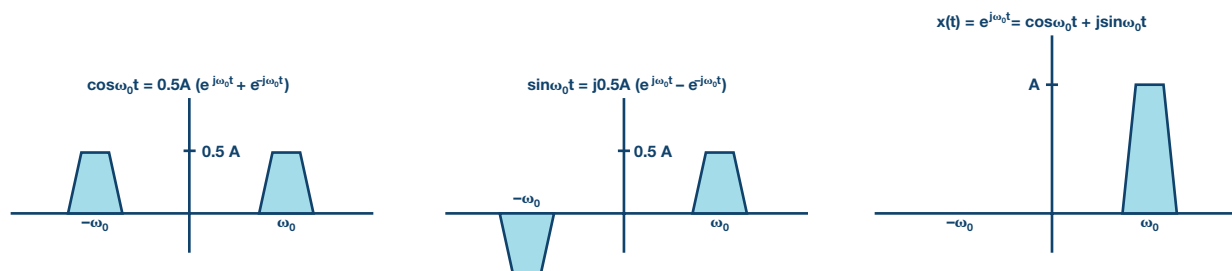
Integrated Transceivers Provide SWaP Solution

Integrating the ZIF architecture into a monolithic transceiver device provides the path forward for next-generation systems. By having the analog and RF signal chain on a single piece of silicon, process variation will be kept to a minimum. In addition, DSP blocks can be incorporated into the transceiver, removing the boundary between the quadrature calibration algorithm and the signal chain. This approach provides both unparalleled improvements in SWaP and can also match the superheterodyne architecture for performance specifications.

ADI now offers two transceivers to meet the demands of the aerospace and defense market, the AD9361 and AD9371. These devices integrate the full RF, analog, and digital signal chain onto a single CMOS device, and include digital processing to run quadrature and carrier leakage correction in real time across all process, frequency, and temperature variations. The AD9361 focuses on medium performance specifications and very low power, such as UAV data links, handheld and manpack communication systems, and small form factor SIGINT. The AD9371 is optimized for very high performance specifications and medium power. Additionally, this device has an integrated ARM® microprocessor for refined calibration control, as well as an observation receiver for power amplifier (PA) linearization and a sniffer receiver for white space detection. This opens up new design potential for a different suite of applications. Communication platforms using wideband waveforms or occupying noncontiguous spectrum can now be implemented in a much smaller form factor. The high dynamic range and wide bandwidth allows for SIGINT, EW, and phased array radar operation in locations with highly congested RF spectrum.

Next Generation Is Now

100 years of device optimization allowed the superheterodyne to achieve greater and greater performance in continually smaller and lower power platforms. Those improvements are beginning to slow down as physical limitations take hold. Next-generation aerospace and defense platforms will demand a new approach to the RF design, one where several square inches of an existing platform are integrated into a single device, where the boundary between software and hardware is blurred allowing for optimization and integration currently unavailable, and where decreased SWaP no longer means decreased performance.

The combination of the AD9361 and AD9371 provides aerospace and defense designers with the ability to now create systems that, just a few years ago, would have been impossible. The devices share many similarities—tunable filter corners, wideband LO generation, diversity capability, and calibration algorithms. There are key differences, though, that drive each part to be optimized for different applications. The AD9361 is focused on single carrier platforms, where SWaP is a primary driving force. The AD9371 is focused on wideband, discontinuous platforms, where performance specifications are even more difficult to achieve. These two transceivers will be key enablers for next-generation aerospace and defense signal chains.
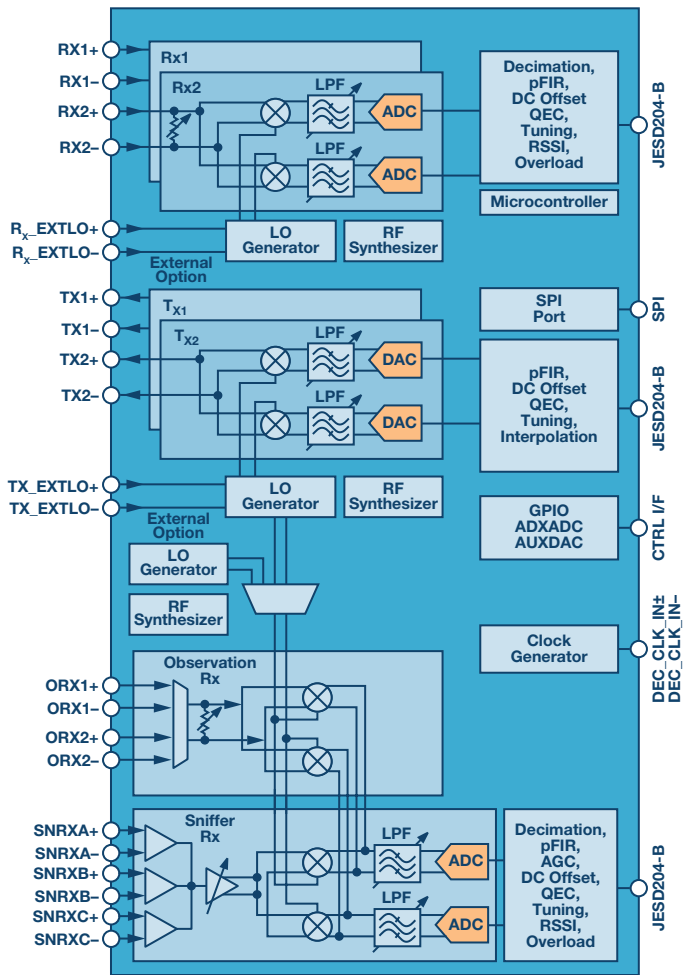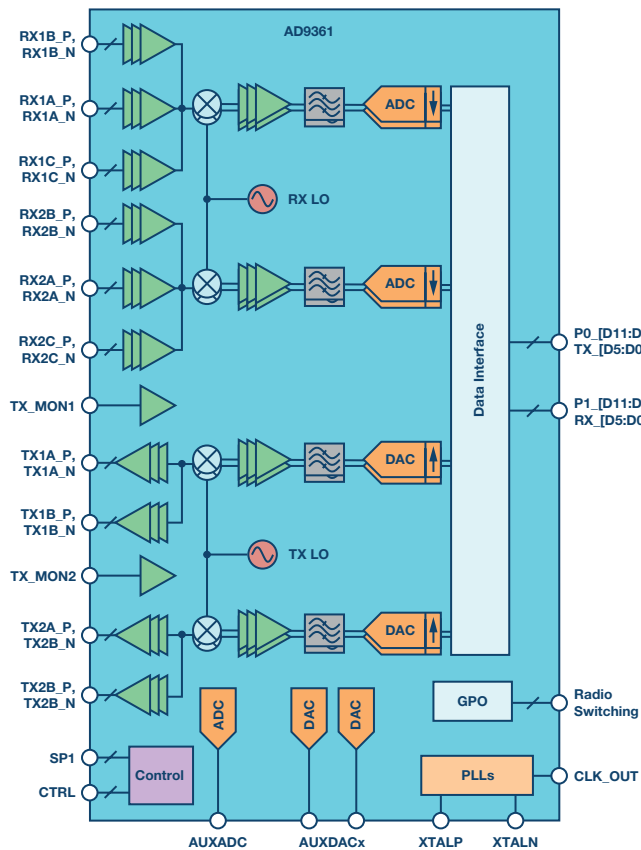

Figure 3. Zero–IF image cancellation.

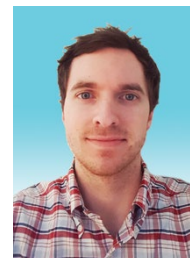*Figure 4. AD9361 and AD9371 block diagrams.*

Wyatt Taylor

Wyatt Taylor [wyatt.taylor@analog.com] is a senior RF systems engineer with Analog Devices, located in Greensboro, North Carolina. He is focused on aerospace and defense radio applications, with a particular emphasis on integrated RF transceivers, small form factor microwave design, and software-defined radio (SDR). Formerly, Wyatt was an RF design engineer at Thales Communications, Inc., and Digital Receiver Technology, Inc., in the Maryland area. Wyatt received his M.S.E.E. and B.S.E.E. from Virginia Tech in Blacksburg, VA., in 2006 and 2005, respectively.

David Brown

David Brown [david-w.brown@analog.com] is an RF system applications engineer with Analog Devices, in Greensboro, North Carolina. He joined ADI in 2015 and focuses primarily on aerospace and defense applications. David graduated from North Carolina State University in 2014 with a bachelor's degree in electrical engineering.

Notes

analog.com/analogdialogue

ANALOG
DEVICES

AHEAD OF WHAT'S POSSIBLE™