Getting the Most Out of Your High Speed ADC
Part 1

By
Derek Redmayne
Linear Technology, Applications Engineer, Mixed Signal Products

Alison Steer
Linear Technology, Product Marketing Manager, Mixed Signal Products

Achieving full performance from an analog-to-digital converter (ADC) with 100dB SFDR, GHz input bandwidth, and SNR in the neighborhood of 80dB can be a challenge. Board level designers need to be equipped with an intimate understanding of the clock and sampling mechanism. This two part article looks at the most common ways in which the ADC performance can be compromised. There are a number of issues that we see recur, and which can render a design unusable:

- Believing the term “low jitter”
- Thinking the clock is a digital signal
- Thinking that differential signaling provides noise margin
- Copying the demo board
- Squaring up the clock signal
- Not thinking about GHz frequencies when designing for baseband
- Isolating analog and digital ground planes

Not all of these are necessarily always bad but each requires careful consideration. Please follow these words of caution to avoid the same traps of your predecessors……

Caution # 1: Believing the Term “Low Jitter”

Many clock sources, synthesizers, repeaters, zero delay buffers, are designed for high speed serial communications, not sampling. Although datasheets for these devices may mention the term low jitter, the term is relative. For de-serialization, low jitter may be 30-50 psec. For a DLL in an FPGA to maintain lock, 30-50 psec is on the order of what is required. However, for IF undersampling, with high dynamic range ADCs, low jitter means somewhere below 1 psec (see Fig. 1).

The encode clock of a high performance ADC is much like the local oscillator of a radio receiver. As in radio, noise, in the form of phase noise, and thermal noise in the clock will limit the sensitivity of the ADC if there is a strong signal present. Much like the mixer in a super-heterodyne radio, the ADC mixes the clock with the analog input. Indeed, this similarity in behavior extends to producing mixing products on the input port, again like the LO –RF feed-through of a true mixer. In the case of the highest speed ADCs, those with high input bandwidth, and the highest performance, like

Linear's 16-bit, 185Msps LTC2209 family, there are mixing products produced that extend out to GHz frequencies. This issue will be discussed in further detail in Part 2 of this article.

In fact, the sensitivity of the ADC to jitter, or phase noise is greater than that of a mixer as the bandwidth of the subsequent outputs is greater. In the case of the ADC, the output bandwidth is really not limited, even though it all folds down into a band from DC to half the sample frequency (fs/2). In the case of a mixer as used in radio design, the output of the mixer is typically band limited by a filter, generally selecting only the difference frequencies of interest; not the sum and difference products of the input, nor all the harmonics of the clock bandwidth extending out to approximately 1 GHz on the analog input, and 2-3 GHz on the clock. In addition, most radio design does not involve undersampling, except in the case of harmonic mixers, and undersampling, often used with ADCs, exaggerates the effect of jitter by $10\log_{10}(fin/fs)$. To relate this back to radio design, this is the reason that the super-heterodyne is generally chosen in receiver design, as the LO being higher than the RF mitigates the effect of jitter.

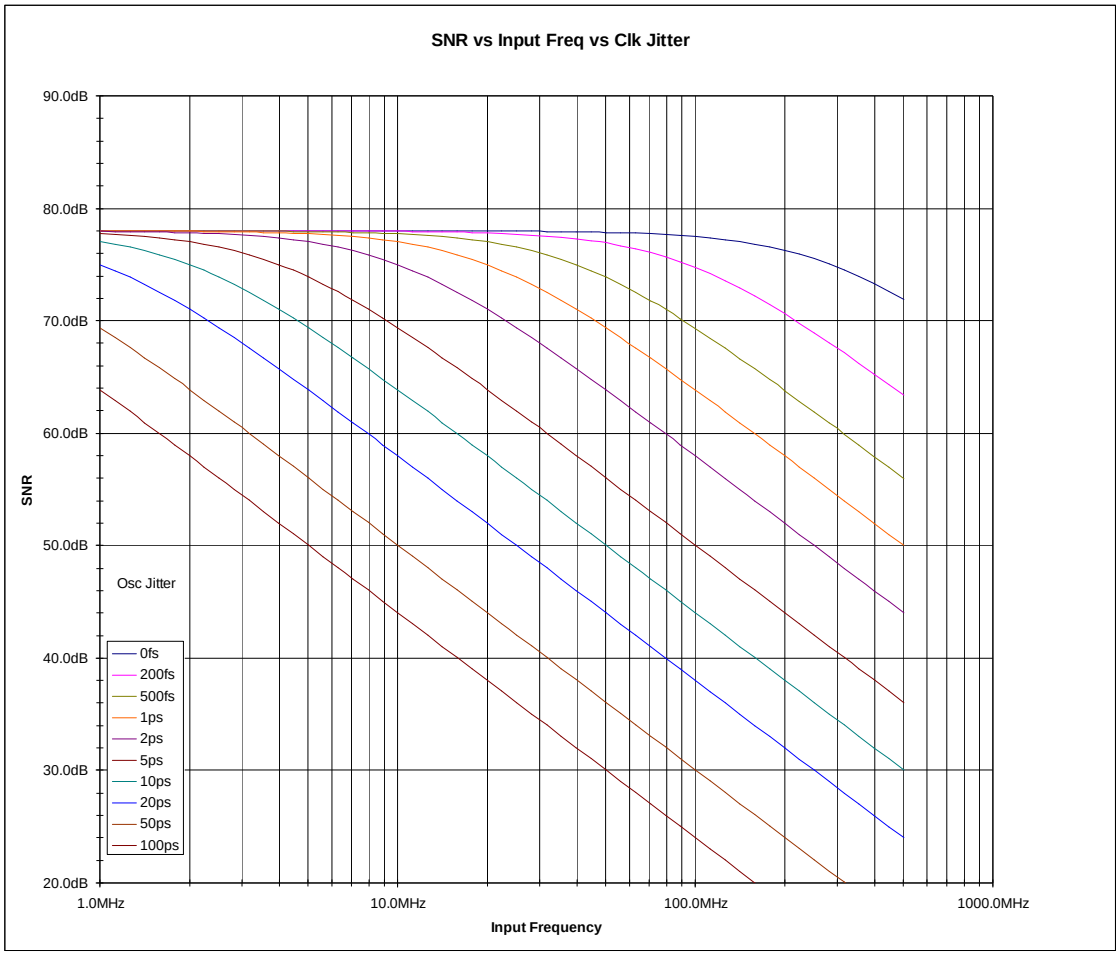

Figure 1: Jitter degradation of SNR as a function of input frequency

The term “low jitter” has become meaningless, and should not be taken at face value. There are a considerable number of clock generator devices on the market that range from rather poor monolithic devices with internal VCOs, loop filters, and poorly grounded packaging, to fairly expensive multi-loop modules with either external VCOs, or VCXOs. We do not want to endorse or condemn clock generator products from other manufacturers, but must warn that you must evaluate these early in your design process, and that this may be the greatest value in getting a demoboard for the ADC, as it, and a facsimile of your signaling, is the best way to evaluate the effect of potential clock sources. We see a great deal of effort spent, trying to confirm the datasheet performance of an ADC, when the usual outcome, after coming face to face with the limitations of signal generators, is that yes, the parts do perform as specified. All lab bench generators require filters to produce data sheet SNR, or SFDR. Some PCB mount clock generators may have specs that at first glance may seem attractive, but may have jitter integrated over a frequency range that opportunistically avoids spurs produced by the architecture, and which may be difficult to eliminate. The common 12 kHz to 20 MHz integration for jitter may be meaningless if there is a 25MHz spur just outside those limits. Note that single chip clock generators are likely to be poor unless they use a very high frequency internal VCO at more than just a few GHz.

Caution # 2: Thinking the Clock Is a Digital Signal

Now, assuming that a low phase noise oscillator or synthesizer of some type is adopted, and incorporated into the system, effort must then be expended to ensure that this is preserved until it reaches the ADC. There are many ways in which a clean clock source can subsequently be compromised. The most popular are:

- Passing the clock through an FPGA.
- Routing through an internal layer of the PCB alongside digital data lines.
- Routing across a back plane.
- Routing through clock fan-out devices that have features such as zero delay, programmable skew etc.

These could all be lumped under the banner of considering the encode clock to be a digital signal.

Even relatively good clock management devices intended for ADC use can be made marginal by compromises in layout in their immediate surroundings: Ignoring the effects of lead inductance, trace inductance, via inductance in bypass, poor or absent isolation barriers between neighboring traces in the layout. If you route a clock output from a device equipped with programmable dividers alongside an output at a different frequency than that intended for the ADC, you will see some manifestation of the other clock frequency. For example, routing lines producing ½ fs right next to the encode clock will produce image frequencies mirrored around ½ Nyquist. This is due to phase modulation of the clock that produces aliasing of products as if they were sampled at ½

fs. The unwanted images may be 70 dB down from the originator, but that can be a show stopper in many applications. What works for OFDM or WCDMA, both fairly tolerant of images and jitter, will not necessarily work for general purpose software defined radio.

Many clock management devices will have some limits in terms of the isolation that they can maintain between outputs. Make sure that you test with all outputs enabled, at different frequencies if that is an eventuality, and make sure that any digital signals, such as other optional ADCs are active. Testing a multiple ADC design, without all the ADCs operational is a big mistake. The performance of clock management devices can be made worse by introducing heavy loading, creating asymmetry in the loading, and reflections from loads, and of course routing lines alongside each other. Not following recommendations in placement of bypass, placement of ground vias can have a very negative effect. Routing clock lines, even at the same frequency, alongside each other is not recommended. We have seen cases where serpentine routing has placed these in close proximity at different points in time, and this can cause waveform distortion.

One last note: digital signals are often daisy chained: Do not do this with encode clocks; the earlier devices will see a compromised edge, often producing a dwell in the threshold region.

Caution # 3: Thinking That Differential Signaling Provides Noise Margin

On the subject of routing clock lines: Even differential signaling such as PECL or CML will couple into neighbors if they are too close. If you have had, or are having problems along these lines, there are experiments that can be performed on your board to determine how effectively neighboring lines are isolated from each other. Take a bare board, and a network analyzer, and introduce the stimulus into one line that is in close proximity to your clock line, and examine S12, or the power transferred into your clock line, sensed at the ADC. You may need to use a high frequency Guanella type balun such as the M/A-COM ETC1-1-13 to translate the single ended stimulus to differential, and back again for the analyzer, emulating the differential input of the ADC. You will need to add source and end termination to your board on those traces involved, or partially populate and power select devices to get realistic measurements. If you have intervening traces, or neighboring traces, they may need to be terminated in realistic fashion to get an accurate reading of how much crosstalk you will see.

One of the pitfalls we have observed is the use of signal integrity software that gives yes and no answers, and is intended to be used to predict crosstalk in digital signaling, being used to validate the clock layout.

The encode clock of an ADC has no noise margin. The acceptable amount of crosstalk into the clock from digital lines carrying 3V logic is in the neighborhood of -100dB or more.

If you make the mistake of thinking of the encode clock as a digital signal and run it alongside the data bus, you will allow digital feedback to phase modulate the clock. One line running at minimal separation from another, with power passing in opposite directions is very much how a directional coupler is constructed. The fact that one of these lines is part of a differential pair, has little impact on the signal that is induced in the nearest member of the pair. If sandwiched between closely spaced ground planes, the rejection of near channel interference may be less than 30 dB. The nearer member will have a stronger signal induced, so rejection is minimal. If the clock source were located close to the FPGA (and this is not recommended) the clock line should be isolated from any digital lines, even those with low repetition rates by a layer of solid copper, or at least an uninterrupted line of vias between grounded copper, above and below the clock. This should be regarded as a coaxial cable buried in the board. Similar issues will be found on the surface.

One feature that we often see in clock distribution is some attention to signals in the same layer, but an apparent lack of diligence in ensuring that there is nothing offensive in those layers above and below the clock lines. This seems like stating the obvious, but it occurs very frequently that it must be a kind of blind spot in common CAD practices.

There must be no lines paralleling the clock lines in layers above or below the clock lines; this also means power planes. Making a sandwich of power planes, ground planes and clock lines can be an expensive mistake. There is some conventional wisdom that causes some grief with clock distribution. A large collection of lines that cross perpendicular in an adjacent layer, common in digital signaling, can still electro-statically couple into the clock line. If the clock lines are differential, this would be a common mode signal, but there are limitations to common mode rejection in a clock receiver. It is unrealistic to expect more than about 30-40 dB common mode rejection at high frequency, and that can be further compromised by asymmetry in the clock path to the ADC.

If a clock originates at some distance from the ADC, the odds of picking up digital noise en route to the ADC increases dramatically. If a single ADC, or a small group of ADCs, is involved in a design, the clock source should be located as close as possible to the ADC. Of course this should be done without compromising the layout of either the ADC, or the clock synthesizer, or compromising thermal design. If however, there is a master clock at the encode rate, that is centrally located, and must be distributed to multiple boards you must consider either a filter where that clock is received from the back plane, or via a coaxial connector, or you must use a jitter cleaner PLL. A jitter cleaner is effectively a very narrow band filter. It cannot improve the jitter within the loop BW, but will reject wideband noise, and distant spurs. The bandwidth of the PLL as a filter is essentially 2x the bandwidth of the loop filter. If the passage of a clock through a backplane only acquires far-out-of-band interferers, it may be practical to filter it with an LC filter before it is presented to the ADC or to any kind of clock fan-out device. The use of a low-pass filter to filter out noise acquired when passing through an area of digital circuitry may be unwise, as low repetition rates, or low frequency content in a digital data bus may get into the clock trace. A band-pass filter is better in this case. A

frequent justification for routing the clock among digital signals is that they are low repetition rate or static. If these originate in an FPGA, a DSP, or a microcontroller, there will be problematic noise present. Any low rate signal that passes through RF or clock generation circuitry should be band limited, and low pass filtered.

Part 1: Conclusion

It is difficult to do any of these subtopics justice within the constraints of a single article. Linear can answer any questions you have on these topics or offer design support including layout review for faster time to production. Further information on Linear's extensive High Speed ADC portfolio can be found at http://www.linear.com/ad/highspeedADC.jsp, which includes our new 1.8V lowest power, LTC2261 ADC family that consumes just 127mW at 125Msps. Linear offers complete ADC evaluation systems comprisedof clock source, signal source and ADC demo boards for use with our free Quick-Eval software tool. These systems are available through your local Linear sales representative.

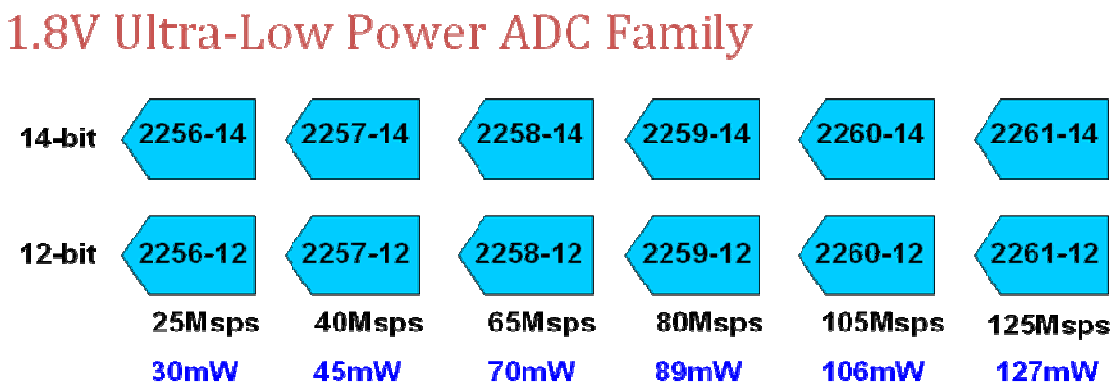

Getting the Most out of Your High Speed ADC
Part 2

By
Derek Redmayne
Linear Technology, Applications Engineer, Mixed Signal Products

Alison Steer
Linear Technology, Product Marketing Manager, Mixed Signal Products

Achieving full performance from an analog-to-digital converter (ADC) with 100dB SFDR, GHz input bandwidth, and SNR in the neighborhood of 80dB can be a challenge. Board level designers need to be equipped with an intimate understanding of the clock and sampling mechanism. This two part article looks at the most common ways in which the ADC performance can be compromised. There are a number of issues that we see recur, and which can render a design unusable:

- Believing the term “low jitter”
- Thinking the clock is a digital signal
- Thinking that differential signaling provides noise margin
- Squaring up the clock signal
- Copying the demo board
- Not thinking about GHz frequencies when designing for baseband
- Isolating analog and digital ground planes

Not all of these are necessarily always bad but each requires careful consideration. Please follow these words of caution to avoid the same traps of your predecessors……

Caution # 4: Squaring up the Clock

If a noisy clock is amplified by a limiting amplifier, then subsequently band limited, the damage is already done. Limiting produces commutation of the various frequency components involved, hence the generation of intermodulation components. Limiting also translates “out of band” components to in-band components, which cannot be filtered out.

The transmission of a filtered sinusoidal encode clock on a back plane, that sees no band limiting upon reception will pass all the out of band noise acquired en route to the ADC.
By noise, we are talking about interferers, digital noise rather than thermal noise. If the clock is amplified after transmission over some distance, the amplifier must not be driven non-linear, or these out-of-band components will fold into a lower frequency, perhaps very close to the fundamental, and become impossible to

remove. This situation is often overlooked in the case where PECL repeaters, CMOS to PECL translation, gates, hysteresis, comparators or any other manner of limiting device is used to “square-up” the signal. Squaring up a polluted clock means mixing products, and filtering after this point is like closing the barn door after the horses have bolted. Amplification of the signal to increase dV/dt, in order to satisfy the requirements of the clock receiver, must occasionally be done, and you can use RF amplifiers to do this. You can also use transformers with 1:4 turns ratio or more, or impedance transformation in a band-pass filter. The signal must be clean before being “squared up”.

Passing a good clock through an FPGA, to give some flexibility in divide ratios, then passing it to a clock repeater, which is essentially a limiting amplifier, will lock-in the loss by translating the spurious content picked up in the FPGA to lower frequencies. A limiting amplifier behaves in this respect like a mixer. Any spurs at some offset frequency from harmonics of the fundamental will appear at the same offset frequency from tones in your input spectrum.

If the crosstalk mechanisms, ground bounce, power supply noise that occur in an FPGA were primarily high frequency, following the FPGA with a low pass filter prior to using a clock repeater may work, but it is risky. If the FPGA was later altered to produce a different output spectrum, on unrelated I/O, you may get a surprise. However, if an FPGA is used to produce a programmable clock divider, or a means of gating the clock, if that clock is re-timed using a D type flip flop clocked by the original clock source, you may get good results if you are careful.

This leads into a discussion of some of the more subtle ways of degrading the clock.

If clock management devices are used to drive multiple loads, some of these offensive loads may feed back into the clock path to the ADC. If one of the outputs of a clock fan-out device is used to drive an FPGA, an FPGA producing unrelated or sub-harmonic outputs that result in ground bounce on its substrate, some of this can reflect back into the clock fan-out device. If this is a question of a differential clock, LVPECL for example, the reflected offenders would result in common mode to a large extent; however, asymmetry in the layout of the clock along the path from the offender through to the ADC may compromise the common mode rejection of either the ADC, or the repeater.

If an FPGA is to be used to implement the digital portions of a PLL, you must great take care. This is not impossible if an external VCO, loop filter, and a re-timing stage are used. You should be aware that the output of the VCO driving the FPGA could reflect back into the oscillator. The use of a splitter to provide isolation between the FPGA, and the re-timing stage for example may be a mistake. Splitters, either Wilkinson splitters using transmission lines, or lumped element splitters, 90 degree, 180 degree, or 0 degree, all have limited isolation out of band. Resistive splitters also do not provide much isolation, although it is

broad band. The out of band noise reflected from the FPGA can phase modulate the clock as it originates in the VCO, and will not be recoverable. The solution is to use an asymmetrical split, sending most of the VCO output power to the re-timing stage, and buffering the VCO output prior to the FPGA. The reverse isolation of the buffer will then compound the isolation from the other elements in the path. It may be advisable in this case to introduce a filter between the VCO and the buffer, such that the reverse isolation is much improved at higher frequencies where it is likely that the bulk of the ground bounce in the FPGA will appear. The use of, for example, 20dB gain to drive the divider input of the PLL in the FPGA would also improve reverse isolation as higher gain amplifiers tend to have higher reverse isolation due to reduced feedback, and if you attenuate by, for example 20dB into the amplifier, you will gain another 20dB of reverse isolation.

Similar issues occur in a clock architecture that involves a clock fan-out device that must drive a device with ground bounce present on its substrate. A poorly designed buffer, with only one ground pin, and one power pin, and multiple outputs is unlikely to be able to reject reflections from the other loads. Devices in QFN, or other exposed pad packages have a better chance of providing isolation. Current Mode Logic (CML) clock repeaters may be the best choice as the CML output stage is not as susceptible to induced ground bounce as is PECL or LVDS, and certainly CMOS single ended clock architectures are considerably more sensitive to reflected noise from loads.

Caution # 5: Copying the Demo Board

It is useful to step back and review a basic characteristic of transmission lines at this point.
If a transmission line, perhaps a coaxial cable, is linking two subsystems where there is potential difference between the two grounds, the rejection of this potential difference will only be effective if the transmission line is terminated at its characteristic impedance at both ends, over the entire bandwidth that may be observed at the receiving end.
We have seen many cases where this appears to have been forgotten. If there is a reflective filter at the driven end of the cable, and essentially wideband termination at the receiving end, any potential difference outside the pass-band, where the transmission line is misterminated at the filter, will manifest itself at the receiving end, the ADC. This is true at either the analog input, or the clock input. Most filters are reflective; they produce their desired characteristics by reflecting the unwanted stop-band power back to the source. Absorptive filters are far better in these cases as the transmission line is properly terminated out of band.

The typical ADC demo board is a compromise as it must be usable over a very wide frequency range, on either the analog input, or the clock input. In the lab, we place the filters as close as possible to the demo board, as this reduces the

potential difference that is developed between the filter, and the termination on the demo board. But this is often not done in cases where the circuitry of the demo board is replicated in a design. Ideally, if a signal is to be transmitted between two boards or subsystems, the signal should be filtered on reception, not on transmission. This is not like the case of a string of two terminal devices, nor is it reciprocal like transmission and reception. If the signal is received by a filter that has good return loss in-band, and is transmitted by a source that also has good return loss in-band, or broadband, any potential difference between the grounds out-of-band will be rejected by filter, and in-band, by the transmission line. There are limitations to the rejection that occurs, and this is a question of how well matched the transmission may be. It is common practice to consider that 20dB return loss is acceptable, but if there is a potential difference at some frequency of around -80dBm, this may be apparent as a visible spur in the spectrum of the ADC.

Caution # 6: Not Thinking about GHz Frequencies when Designing for Baseband

Direct sampling ADCs produce mixing products extending out to a few GHz.

This is a not uncommon complaint about high speed ADCs, as it is often unanticipated by inexperienced designers, and is often found rather late in the design process to be a cause of interference in the front end of software radios, often located in the same cavity in an enclosure. If the mixer is located in the same cavity as the ADC, and this cavity has reflective walls, those mixing products at GHz frequencies will be picked up by the LNA, or the RF port of the mixer, and translated to the IF band, and digitized, appearing at a different frequency. Poor placement of the local oscillator or routing of the local oscillator drive to the mixer may also result in unwanted spurs picked up from that ADC.

Those frequencies that could be received by exposed antenna area prior to the mixer must be suppressed as close as possible to the input port of the ADC, or the ADC should be in a separate compartment. I often suggest that the ADC should be in the same compartment as the digital circuitry rather than in the shielded RF section. If the ADC and RF must be in the same enclosure, the antenna area in both cases, at the RF port of the mixer, and the input network of the ADC, must be minimized.

This means that the driver, or filter prior to the ADC must be as close are possible to the ADC. If, as is generally the case, a lowpass (LP) or bandpass (BP) filter is placed between the drive amplifier and the ADC, this filter must be designed to pass the band of interest, and the stop bands in neighboring Nyquist zones, but the layout must be designed to also suppress GHz frequencies, and avoid radiating at these frequencies. The use of differential signaling may mitigate radiation, or pick-up of unwanted signals to some extent, but any

mismatch in components, physical layout, including asymmetry in cavities limits the benefit. Widely spaced differential filters in asymmetrical cavities will behave more like single ended radiators.

The inputs of the high speed direct sampling ADCs contain mixing products in differential form, due to commutation of the input signal by the sample switches, but they also produce common mode artifacts that can radiate more effectively than the differential components, as this network is then much like a single ended transmission line. These common mode components can also produce disturbances in the substrate of a driver that may be expected to have some noise immunity. A poorly grounded or bypassed drive amplifier may do little to suppress either the differential components, or the common mode components. Occasionally we have seen the use of an RF power detector monitoring the path to the ADC. Quite frequently, in these cases, there is no isolation whatsoever. In these cases, the RF power detector will always show indication of some nominal power level, as long as the ADC is sampling. In some cases the RF power detector and the ADC are fed with a splitter, but again, a splitter may provide minimal isolation out of band, and especially so if the return loss seen by the input port of the splitter is poor. Occasionally, an amplifier with rather poor reverse isolation is found to cause the same disappointment.

In these cases, the power received by the RF power detector must be taken prior to the final filters and drivers. This should probably be done in any case as the power delivered to the ADC is evident in it's output, but the power delivered prior to the band limiting is not apparent, and generally should be known in order to manage AGC.

Caution # 7: Isolating Analog & Digital Ground Planes

There is a long history of isolating analog and digital ground planes. This, however, more often than not, leads to problems. The most serious blunder that we see on a regular basis is that of crossing a gap between two planes with CMOS digital outputs from an ADC.
Sometimes these planes are linked with an inductor, a ferrite bead, a resistor, or even just a trace. Sometimes they are returned to the power supply through long leads. Unfortunately, the propagation of a signal requires that ground current flow in opposition to the charges propagating on the transmission line. In differential signaling, this flows in the other signal, with little significant current in the ground plane. In the case of CMOS data crossing a gap, the ineffective path for the ground current results in a potential difference being developed between the two planes. This will be a high frequency potential difference if the detour in the ground path is short, extending to lower frequencies as ground path gets longer. If you must do this, you should use LVDS or CML outputs from the ADC. Linear Technology offers a 16-bit 105Msps serial ADC, LTC2274 (see fig 1), with CML 8B/10B encoded outputs streaming data at 2.1Gbps over the differential

pair. This makes it easier to isolate analog and digital circuitry, while utilizing the SerDes ports commonly found on FPGAs.

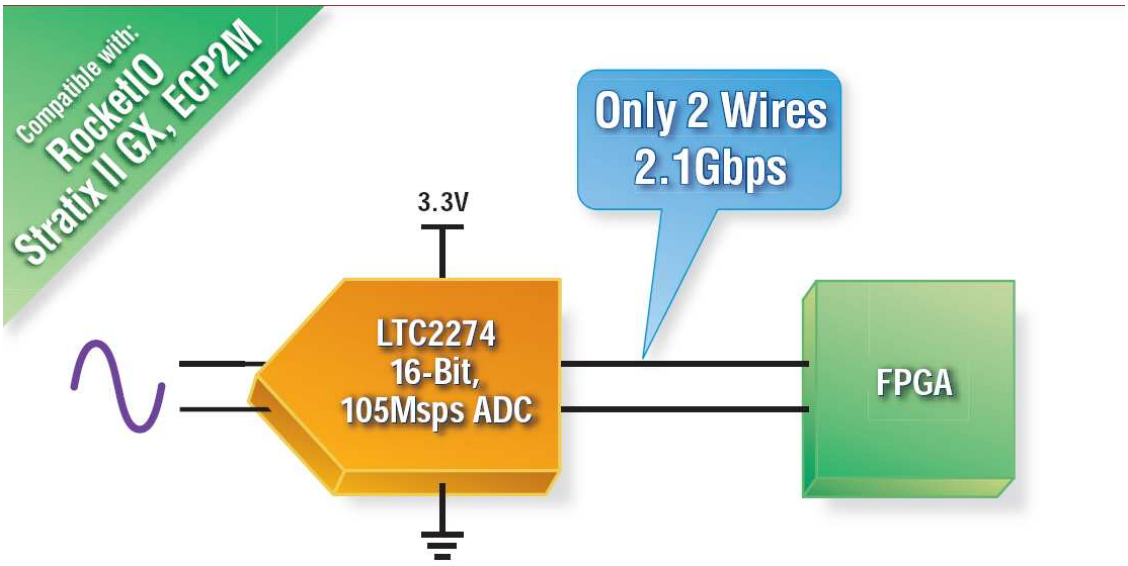

Fig 1: New 16-bit 105Msps ADC with JESD204 2-Wire Serial Interface

Otherwise, it is best to put the ADC at the edge of the digital plane, and cross the gap with a 1:1 transmission line transformer similar to the M/A-COM ETC1-1-13. Also known as the Guanella 1:1 balun, this device acts like a common mode choke, while passing up to 3GHz differential. If the transformer is followed by, or preceded with a low-pass filter, you can use a flux coupled transformer. As any direct sampling ADC produces transients with content out into the GHz region, this network should be devised to suppress these components before it extends any length, and before crossing those boundaries.

Any application that entails transmitting differential signaling between two ground planes needs to be protected against excessive voltages developed between the ADC and the load by events such as ESD strikes, line surges etc.

Conclusion

It is difficult to do any of these subtopics justice within the constraints of a single article. Linear Technology can answer any questions you have on these topics or offer design support including layout review for faster time to production. Further information on Linear's extensive High Speed ADC portfolio can be found at http://www.linear.com/ad/highspeedADC.jsp which includes our new 1.8V lowest power, LTC2261 ADC family consuming just 127mW at 125Msps. Linear offers complete ADC evaluation systems comprising of clock source and clock divider boards, signal source and ADC demo boards for use with our free Quick-

Eval software tool. These systems are available through your local Linear sales representative.

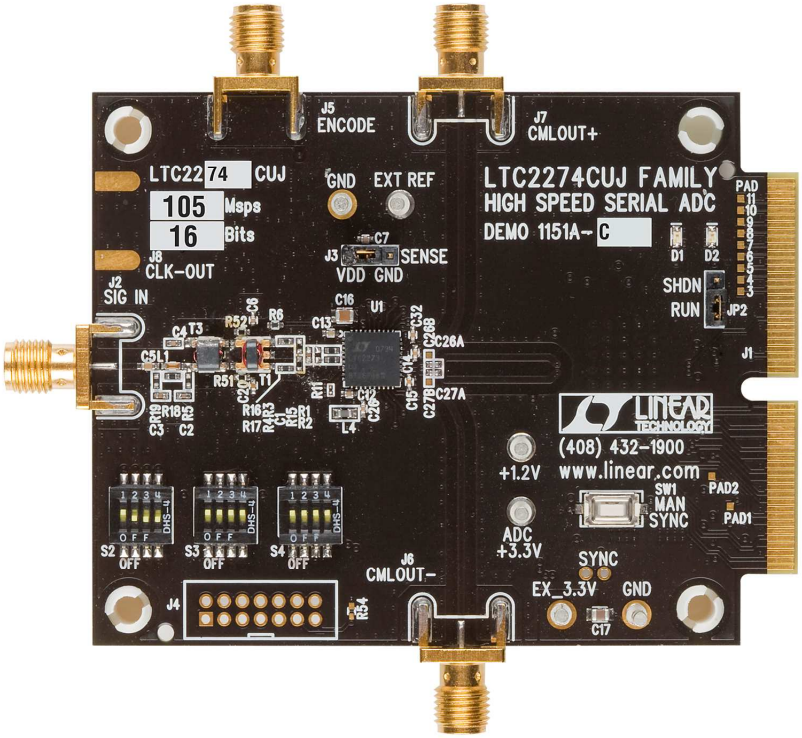

Fig 2: LTC2274 High Speed ADC Demo Board with CML Outputs