Interleaving ADCs: Unraveling the Mysteries

By Gabriele Manganaro and Dave Robertson

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

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

As the n-bit outputs of the ADCs become sequentially available in the same order as just described for the sampling operation, these digital n-bit words are collected by the demultiplexer shown on the right hand side of the same figure. Here the recombined data out sequence $D_{OUT}(t_0 + L)$, $D_{OUT}(t_0 + L + T_s)$, $D_{OUT}(t_0 + L + 2T_s)$, ... is obtained. L stands for the fixed conversion time of each individual ADC and this recombined data sequence is an n-bit data series with sample rate f_s. So, while the individual ADCs, often referred to as the "channels," are n-bit ADCs sampling at f_s/M, the ensemble contained in the box is equivalent to a single n-bit ADC sampling at f_s, and we will refer to that as the time interleaved ADC (distinguishing it from the channels). Basically the input is sliced and separately processed by the ADCs in the array and then consistently reassembled at the output to form the high data rate representation D_{OUT} of the input V_{IN}.

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

This powerful technique is not free of practical challenges. The key issue manifests itself when the M data streams coming from the channels are digitally assembled together to reconstruct the original input signal V_{IN}. If we look at the spectrum of D_{OUT}, in addition to seeing the digital representation of V_{IN} and the distortion introduced by the analog-to-digital conversion, we will also see additional and substantial spurious content, termed "interleaving spurs" (or IL spurs, in short), IL spurs neither have the signature of polynomial type distortions like higher order signal harmonics (2nd, 3rd, and so on), nor the signature of quantization or DNL errors. IL artifacts can be seen as a form of time-domain fixed pattern noise and are introduced by analog impairments in the channels that, due to the interleaving process, modulate with the sliced converted signals and ultimately show up in the final digitized output D_{OUT}.

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

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

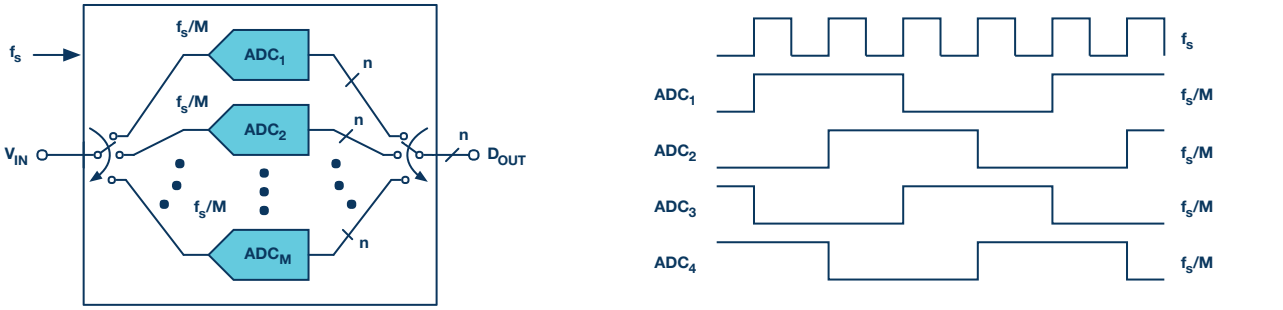

Figure 1. An array of M time interleaved n-bit ADCs. The sample rate of each one is f_s/M, the resulting sample rate of the time interleaved ADCs is f_s. An example of clocking scheme for the case of M = 4 is depicted on the lower part of this figure.

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

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

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

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

Ping Pong (Two-Way) Interleaving

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

If the input signal $V_{IN}(f)$ is completely bound between 0 and $f_s/4$, as in Figure 2(b), then the interleaving spurs are not frequency overlapping with the digitized input. In this case, the bad news is that we are only able to digitize in half of the Nyquist band, namely just like if we had a single channel clocked at $f_s/2$, though we are still consuming at least twice the power of such a single channel. The interleaving spurious image on the upper end of the Nyquist band can be suppressed by digital filtering after digitization and does not require correcting for analog impairments.

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

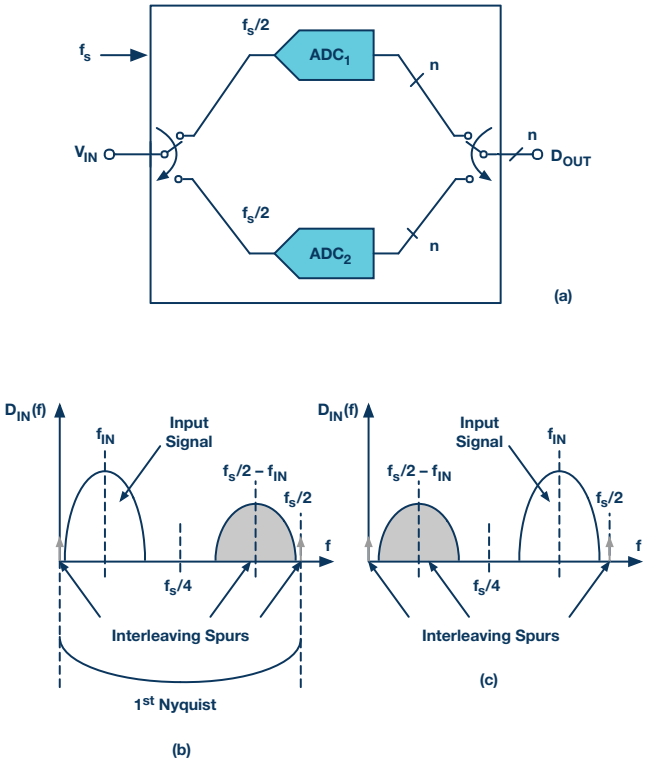

Figure 2. (a) A ping pong scheme, (b) the output spectrum when a narrow-band input signal lies below $f_s/4$, and (c) when the input signal lies between $f_s/4$ and the Nyquist frequency $f_s/2$.

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

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

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

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

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

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

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

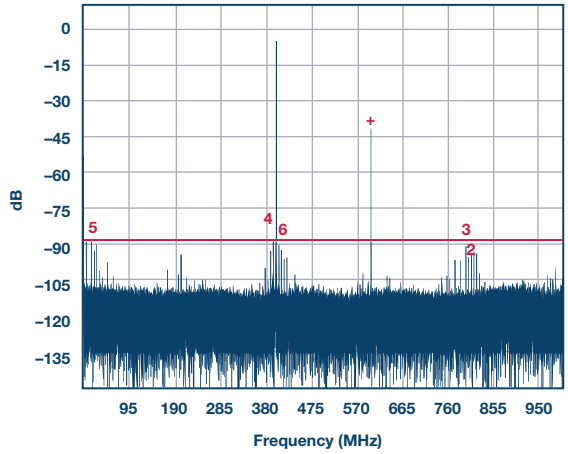

Figure 3. Spectrum of the combined 2GSPS output data of a ping pong scheme obtained by clocking the two ADCs of the AD9680 with a 1 GSPS clock but with 180° sampling phase shift.

Higher Order Interleaving

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

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

Compensation approaches are limited by the accuracy with which the mismatches can be measured and ultimately corrected. To further suppress the residual spurs beyond the level achieved via calibration, it is possible to intermittently and randomly shuffle the order with which the channels sample the input. In doing so, the previously discussed modulation effects of the converted input signal due to the uncalibrated mismatches turn from a fixed pattern noise to pseudorandom. As a result, IL tones and undesired periodic patterns turn into pseudorandom noise-like content that adds with the converter quantization noise floor and leads to the disappearance or, at least, to spreading of the undesired spurious images and tones. In this case, the power associated with the IL spurious content adds to the power of the noise floor. Hence, while distortion improves, SNR can degrade by the amount of IL spurious power added to the noise. SNDR (SINAD) is essentially unchanged as it combines both distortion and noise and randomization; it simply moves the IL contribution from a component (distortion) to the other (noise).

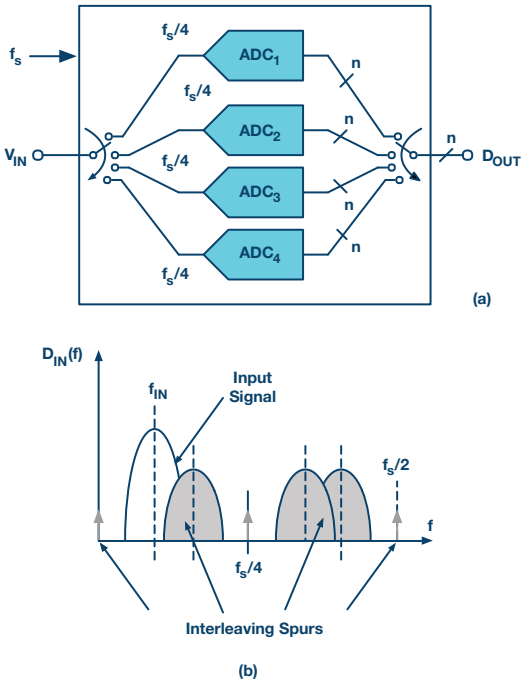

Figure 4. (a) A four-way interleaved ADC and (b) the corresponding 1st Nyquist output spectrum showing the interleaving spurs.

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

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

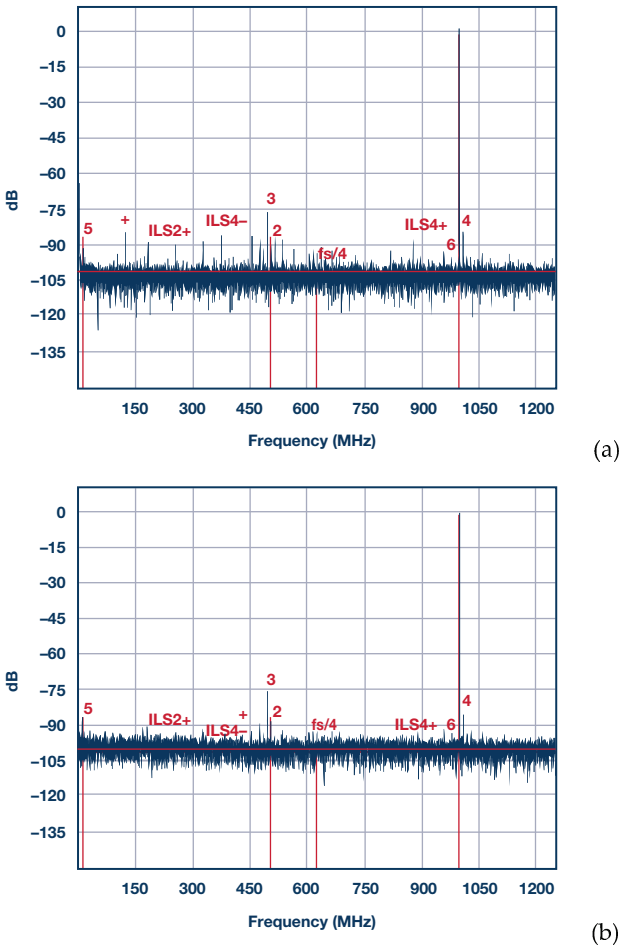

Figure 5. The output spectrum of the AD9625, clocked at 2.5 GSPS and with an input tone close to 1 GHz. (a) Sequential three-way interleaving; SNR = 60 dBFS, the SFDR = 72 dBc is limited by the third harmonic, near 500 MHz; however, a number of interleaving spurs are visible all across the spectrum.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

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

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

Lastly, the spectral purity in Figure 7 can be further improved by randomizing the channel order as shown in Figure 8. In this case the randomization uses a proprietary technique that while intermittently scrambling the order of the four channels doesn't require a spare (5th) channel to be added, hence saving its associated power. It can be seen in Figure 8 that, after randomization, only regular harmonic distortion is left on the resulting spectrum.

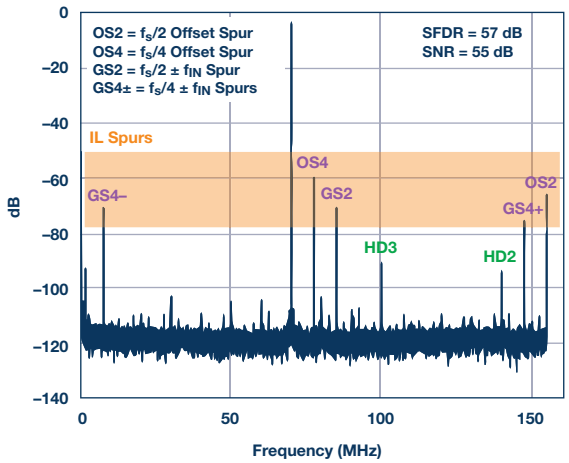

Figure 6. The output spectrum of the AD9652, clocked at f_S = 310 MHz and with a sinusoidal input at f_{IN} ~ 70 MHz. In this case, no channel calibration and randomization is applied. The 2nd (HD2) and aliased 3rd (HD3) harmonics are visible at ~140 MHz and ~100MHz respectively. Interleaving (IL) spurs are visible as well. These are the offset tones at dc, f_S/2 (OS2 in the graph) and f_S/4 (OS4 in the graph). Moreover the gain (/timing) spurs can be found at f_S/2 – f_{IN} (GS2 in the graph), f_S/4 + f_{IN} (GS4+ in the graph) and f_S/4 – f_{IN} (GS4– in the graph). The SNR quote in this graph is artificially poor due to the fact that some of the spurious content has been lumped with the noise power.

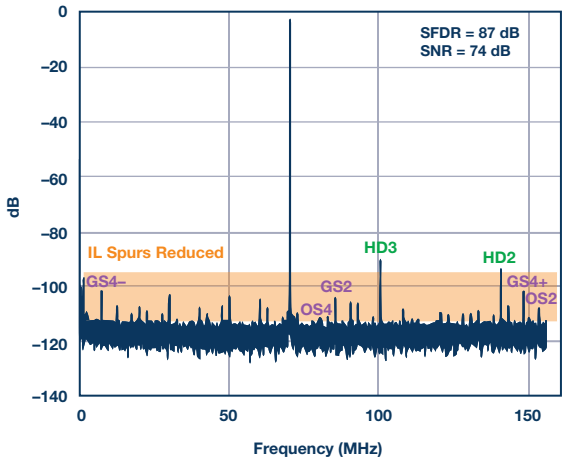

Figure 7. The output spectrum for the same AD9652, with the same input but after calibrating the four channels to reduce their mismatch. Comparing with Figure 6, while the 2nd and 3rd harmonics are unaffected, the interleaving spurs' power has substantially been reduced and the SFDR has improved by 30 dB going from 57 dBc to 87 dBc.

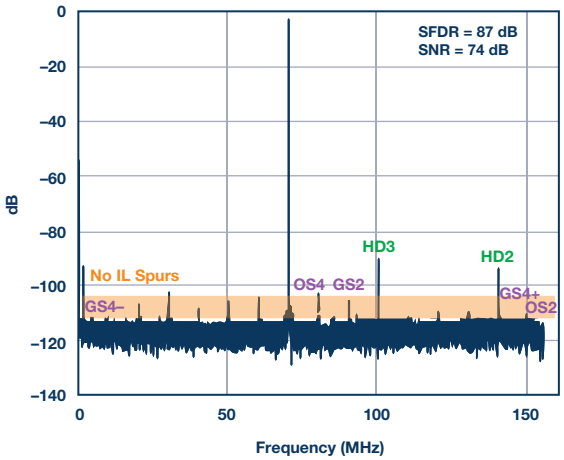

Figure 8. The output spectrum for the previous case once the randomization of the interleaving order is turned on. Randomizing the residual interleaving spurs distributes their power over the noise floor and the corresponding peaks disappear. Only regular harmonic distortion is left in sight. The SNR is nearly unaffected as the distributed spurious power from the interleaving tones is rather negligible after the mismatch calibration.

Conclusion

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

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

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

Acknowledgments

The authors would like to acknowledge Siddharth Devarajan, Prawal Shrestha, Antony DeSimone, Ahmed Ali, Umesh Jayamohan, and Scott Bardsley for capturing and providing some of the experimental results.

References

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

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

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

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

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

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

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

Gabriele Manganaro

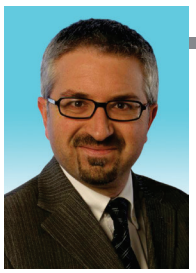

Gabriele Manganaro [gabriele.manganaro@analog.com] holds a Dr.Eng. and a Ph.D. in electronics from the University of Catania, Italy. Starting in 1994, he did research with ST Microelectronics and at Texas A&M University. He worked in data converter IC design at Texas Instruments, Engim Inc, and as Design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

David H. Robertson

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

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