Understanding AC Behaviors of High Speed ADCs

David Kress
Director of Technical Marketing, Analog Devices, Inc.

Abstract

Understanding common converter ac performance characteristics and concepts—quantization, sampling, signal-to-noise and distortion (SINAD), effective number of bits (ENOB), aperture jitter noise, distortion products, spurious-free dynamic range (SFDR)—empowers designers to optimize converter component choices for various design objectives, such as performance and cost.

A growing number of electronic products in consumer, medical, automotive, and even industrial sectors make use of high speed signal technologies for data and voice communication, audio, and imaging. Though these application categories process signals with different bandwidths and use correspondingly different converter architectures, they share certain characteristics when it comes to comparing analog-to-digital converters (ADCs) and evaluating the performance of specific implementations. In particular, designers working in these disparate application categories are concerned with a number of common converter ac performance characteristics that can set their systems' performance limits.

Quantization

All ADCs take input signals that are continuous in time and amplitude and output quantized discrete-time samples. The ADCs' dual functions—quantization and sampling—provide an efficient transformation from the analog to the digital signal domain, but each has consequences with regard to the converters' ac performance.

Because digitizers have only a finite number of codes into which to parse the continuous input signal, their outputs generate an error function in the shape of a sawtooth wave. The sawtooth edges correspond to the ADCs' code transitions.

To gauge the effect of the quantization error's best case noise contribution, postulate a full-scale sine wave input to a perfect digitizer:

$$e_i(t) = \frac{q2^N}{2} \, sin(2\pi ft) \tag{1}$$

where q is the size of an LSB, and N is the number of bits. The root mean square (rms) amplitude of this waveform is simply the amplitude divided by the square-root of two,

$$e_i|_{RMS} = \frac{q2^N}{2\sqrt{2}} \tag{2}$$

and the rms quantization noise is

$$e_n|_{RMS} = \frac{q}{\sqrt{12}} \tag{3}$$

The ratio of the rms full-scale signal to the rms quantization noise gives an ADC's ideal SNR, which we can express in decibels:

$$
\begin{aligned}
SNR(ideal) &= 20\log_{10}\left(\frac{e_i|_{rms}}{e_n|_{rms}}\right) \\
&= 20\log_{10}\left(\frac{2^N\sqrt{12}}{2\sqrt{2}}\right) \\
&= 20\log_{10}2^N + 20\log_{10}\sqrt[3]{\tfrac{3}{2}} \\
&= (6.02N + 1.76) \text{ dB}
\end{aligned}
\tag{4}
$$

Keep in mind that this expression gives the theoretical limit for an N-bit converter. Real quantizers do not attain this level of performance and real converters have additional noise sources, but you can use this number as a reference with which to judge candidate ADCs.
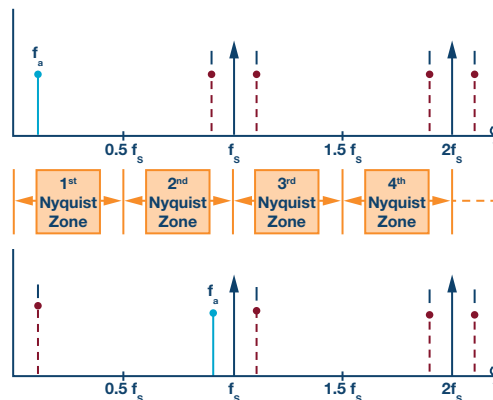

Figure 1. (Top) A sampler causes images (red) of the baseband signal f_a (blue) to appear offset from the sampling frequency f_S and its harmonics. (Bottom) The spectral offsets are equal to $\pm f_a$. Signals, noise, and interference spectra occurring near the sampling rate alias down to the baseband. Images will also appear in the upper Nyquist zones.

Sampling

The best known sampling feature is the aliasing signal energy that occurs at frequencies greater than half the sampling rate—$f_S/2$. This half-sampling-rate limit, named the Nyquist frequency, serves to divide the spectrum into equal sized segments called Nyquist zones. The first Nyquist zone extends from dc to $f_S/2$, the second Nyquist zone occupies the spectrum between $f_S/2$ and f_S, and so on.

In reality, samplers alias signals across all Nyquist zones. For example, images of a baseband signal at frequency f_a will appear at $f_S \pm f_a$, $2f_S \pm f_a$, and so on (Figure 1, top). Similarly, signals occurring near the sampling frequency will alias down to the first Nyquist zone. Images of that signal will appear in the third and fourth Nyquist zones as well (Figure 1, bottom). Therefore, a sampler with input signal energy in a Nyquist zone other than the one you're interested in will generate an image of that signal in the Nyquist zone that you care about simply by aliasing.

The out of band signal energy, shown as f_a (Figure 1, bottom), doesn't have to come from an intended signal source. Instead, that energy could derive from noise sources, an out of band interferer, or distortion products created by circuit elements operating on the intended input signal. This is an important consideration when determining the requisite distortion performance for your application.

You can reduce the amount of out of band signal energy that is available to the sampler by including a baseband antialiasing filter in the signal chain ahead of the sampler's inputs. Though, theoretically, you could sample at only twice the highest frequency you're interested in digitizing, so-called brickwall filters—filters with zero transition band—do not exist in the analog domain. Oversampling—sampling at a frequency greater than $2f_S$—provides some spectral space for the antialiasing filter's transition band.

If the ADC's quantization noise is uncorrelated with the ac input signal, the noise distributes across the first Nyquist zone. In such cases, oversampling also reduces the effective quantization noise by broadening the Nyquist zone, increasing the SNR (signal-to-noise ratio) by 3 dB for each doubling of the sampling rate. This assumes an antialiasing filter with a fixed pass band. With sufficient oversampling, the antialiasing filter can attenuate out of band signal content so that their aliased images remain below the noise floor.

Note that if the input signal locks to an integer submultiple of the sampling frequency, the quantization noise will no longer appear as a uniform energy distribution across the Nyquist zone. In this case, the quantization noise will appear to cluster about the signal's harmonics. For this reason, you should carefully consider your application signal's spectral characteristics when choosing a sampling rate.

SINAD and ENOB

If distortion products and the aliases of out of band spectral content do not remain below the noise floor, they contribute to SINAD. A converter's data sheet will express SINAD in dB under specified conditions of input signal. The converter's ENOB, perhaps the most commonly quoted ac specification for ADCs, is simply the SINAD expressed in bits instead of dB:

$$ENOB = \frac{(SINAD - 1.76 \text{ dB})}{6.02 \text{ dB/bit}} \tag{5}$$

For cases in which distortion products and aliased signal energy remain below the noise floor, SINAD = SNR. In this case, Equation 5 is simply a rearrangement of Equation 4, solving for N. The more usual case is one in which SINAD < SNR. Because the converter's SINAD is dependent on operating and signal conditions, the SINAD (and correspondingly, the ENOB) your application realizes depends on how it drives the ADC.

Though often quoted, the ENOB is insufficient to describe a high speed converter's performance. High speed converters are famously multiparametric, and no single number can hope to capture what takes an entire specification table to describe. ENOB does make a reasonable starting point for comparing candidate converters, so long as you do not depend excessively on the number's significance.
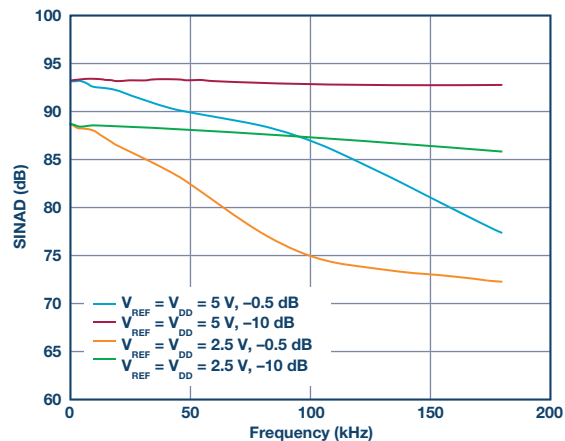

Figure 2. Though ENOB gives a useful (if coarse) comparison between candidate high speed ADCs, characteristic curves depicting SINAD vs. frequency provide substantially more insight into a converter's performance.

More valuable are the SINAD vs. frequency characteristic curves, which many high speed converters present in their data sheets (Figure 2). These allow you to identify at least typical performance at the frequencies of interest to your application instead of at the spot frequencies the converter manufacturer chose for the data sheet's specification table.

Aperture Jitter Noise

The quantization noise discussion that led to Equation 4 assumed an ideal digitizer. Built into that assumption were noise-free signal and clock sources. In real circuits, signals arrive at the ADC's inputs with noise and distortion products already donated by previous signal processing stages. The noise content is usually uncorrelated with the quantization noise and so adds in a root square sum:

$$e_{n(total)} = \sqrt{\sum_{i=1}^{m} e_{n(i)}^2} \tag{6}$$

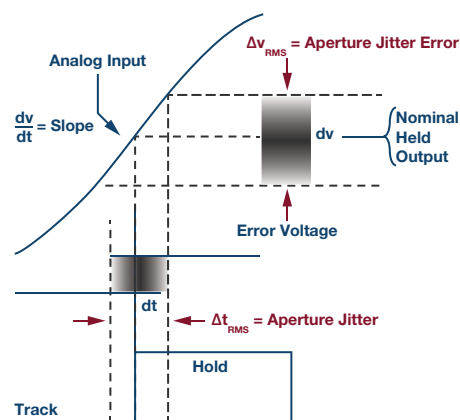where $e_n(i)$ is the noise from a contributing source in a system of m uncorrelated sources.

Figure 3. Aperture jitter—uncertainty in the sampling time—gives rise to a noise amplitude that is dependent upon signal frequency due to the signal slewing during the jitter time.

One of the contributing noise sources derives from uncertainty in the sampling clock edge timing, resulting in aperture jitter noise. This noise draws from the fact that the sampler is capturing an ac signal—aiming at a moving target, so to speak. Variations in the sampling edge timing result in a statistical distribution of amplitudes that the sampler captures—noise (Figure 3). The higher the signal frequency, the greater the signal slope or slew rate and, thus, the greater the amplitude error resulting

from a given change in edge timing. Consequently, the effect of a given amount of aperture jitter is dependent on signal frequency.

The SNR due to aperture jitter is

$$SNR|_{t_j} = 20\log\left(\frac{1}{2\pi f t_j}\right) \tag{7}$$

where f is the signal frequency, and t_j is the rms aperture jitter. Often, when selecting among candidate ADCs, the question is what maximum aperture jitter your application can tolerate given its SNR requirements with signals of a given frequency. Rearranging Equation 7 yields

$$t_j = \frac{1}{(2\pi f 10^{\frac{SNR}{20}})} \tag{8}$$

Note that, in addition to sources of jitter within the converter, there are jitter sources that derive from your application circuit. The net performance your circuit achieves, therefore, is a function of both your converter selection and the quality of other aspects of your design, notably your clock generation circuitry and circuit board layout.

To get a sense of how jitter affects the maximum signal frequency for a given ENOB, consider two systems in which 1 ps and 2 ps of jitter noise, respectively, dominate other performance limiting parameters. Rearranging Equation 8 allows us to calculate, for a given jitter, the maximum signal frequency that results in a specified ENOB (or SNR).

Table 1. Comparison of Systems with Jitter Times that Differ by a Factor of 2

ENOB (bits)	SNR (dB)	f_{MAX} (t_j = 1 ps)	f_{MAX} (t_j = 2 ps)
20	122	124 kHz	62 kHz
18	110	496 kHz	248 kHz
16	98	1.98 kHz	9.93 kHz
14	86	7.94 kHz	3.97 kHz
12	74	31.7 kHz	15.8 kHz
10	62	127 kHz	63.5 kHz

Distortion Products

Nonlinearities within the signal chain give rise to a number of distortion products, notably HD2 (second-harmonic distortion), HD3 (third-harmonic distortion), IMD2 (second-order intermodulation distortion), and IMD3 (third-order intermodulation distortion). Distortion in linear circuits tends to increase gradually as the signal approaches the extremes of the active elements' linear operating range. This is not the case in ADCs where code spaces end abruptly.

It is important, therefore, to allow sufficient range in the input span to accommodate the expected input amplitudes for which you anticipate low distortion quantization, particularly when processing complex broadband signals. Ultimately, your choice of nominal input amplitude results from balancing signal span headroom to avoid clipping against the desire   to optimize SNR.

As the name suggests, harmonic distortion generates signal artifacts at multiples of the signal frequency. By contrast, intermodulation distortion derives from signal processing nonlinearities with signals comprising two or more frequencies—virtually any complex waveform—generating sums and differences of the input frequencies.

In narrow-band applications, a tightly tuned antialiasing filter can attenuate some harmonic distortion products and even the additive components of the IMD2 (Figure 4). On the other hand, the subtractive components of IMD3, which appear at $2f_2 - f_1$ and $2f_1 - f_2$, are pernicious because they can appear within the signal spectrum.
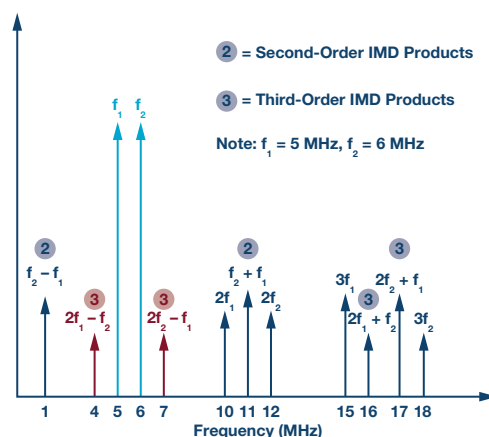

Figure 4. A 5 MHz and 6 MHz two-tone input signal demonstrates HD2 (at 10 MHz and 12 MHz), HD3 (at 15 MHz and 18 MHz), IMD (at 1 MHz and 11 MHz), and IMD3 (at 4 MHz and 7 MHz). Of these, the IMD3 products are the hardest to attenuate with an antialiasing filter due to their close proximity to the source signals.

SFDR

SFDR (spurious-free dynamic range) is simply a measure of the worst-case spectral artifact compared to either the converter's full-scale range (dBFS) or to the input signal level (dBc). Be sure to identify both the reference levels and the operating and signal conditions when comparing ADCs. Direct comparisons between data sheet specifications require matching references and signals (Figure 5).
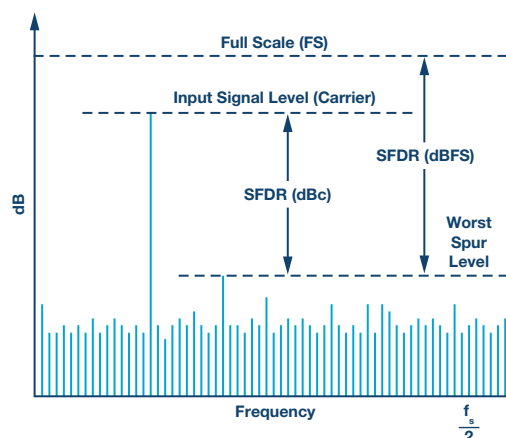

Figure 5. Converter manufacturers can specify SFDR performance with respect to the converter's full scale (dBFS) or with respect to a specific input signal amplitude (dBc). Be sure that candidate converters are specified in a like manner before making numeric comparisons.

Though SFDR appears as a numeric datum in converter specification tables, the measure is itself parametric in sampling rate, signal amplitude, signal frequency, and common-mode operating point. Examine candidate converter characteristic curves to gain insight into a converter's performance under operating and signal conditions similar to those that your application will impose.

Resources

Explore ADC evaluation tools and behavioral models at *analog.com/ADIsimADC*.

About the Author

David Kress earned both his B.S.C.E. and M.S.C.E. degrees from the Massachusetts Institute of Technology. David was a medical electronic equipment designer before joining Analog Devices in 1972 as an IC designer. Over the years, he has been involved in analog chip design, wafer fabrication engineering, and new product development and management. For many years, David led ADI's technical support team and was the director of technical marketing before he retired in 2018.

Online Support Community

EngineerZone™
SUPPORT COMMUNITY

Engage with the Analog Devices technology experts in our online support community. Ask your tough design questions, browse FAQs, or join a conversation.

Visit ez.analog.com

ANALOG
DEVICES

AHEAD OF WHAT'S POSSIBLE™