Demystifying the Conversion Error Rate of High Speed ADCs

**by Ian Beavers, applications engineer,
Analog Devices, Inc.**

Just like many other semiconductor components, high speed analog-to-digital converters (ADCs) do not always perform perfectly, despite our best expectations. They have inherent limitations that permit them to make occasionally rare conversion errors that are outside their normal function. Many real-world sampling systems, such as test and measurement equipment, cannot tolerate a high rate of ADC conversion errors. Therefore, it is important to be able to quantify the frequency and magnitude of a high speed analog-to-digital conversion error rate (CER) so that engineers can design their systems with the proper expectation of performance.

The relatively infrequent occurrence of high speed or GSPS (gigasample per second) ADC conversion errors can make them not only challenging to detect, but it can also make the measurement process quite time consuming. This duration is typically not completed in micro or milliseconds, but rather hours, days, weeks, and even months. To help reduce this burden of lengthy testing, we can approximate the error rate with a confidence level of certainty and still maintain quality in the outcome.

BER vs. CER

Similar to its digital equivalent of bit error rate (BER) in a serial or parallel digital data transmission, the conversion error rate is a ratio of the number of conversion errors divided by the total number of samples. However, there are some distinct differences between BER and CER. BER testing in a digital data stream implements a long pseudorandom sequence that can be started within a transmitter using a common seed value at both ends of the transmission. The receiver will then have the expectation of an ideal transmission. The BER is precisely calculated by observing the difference in the received data compared to the ideal. Mismatches in the pseudorandom sequence data based on the seed value between both ends are counted as bit errors.

Contrasting with CER, the error determination is not quite as straightforward as a pure digital comparison. Since small nonlinearities are always present in the ADC conversion process, along with system noise and jitter, an exact difference between expected and actual data cannot always be determined. Instead, an error threshold needs to be established that determines the boundary between a conversion error and a sample with tolerable, but expected, noise. This contrasts with digital BER in that there is not an exact comparison of expected data transmitted and received. Instead, a sample's error magnitude must first be quantified before it is either considered a conversion error or within the expected nonidealities of the converter and system.

The ADC's back-end digital interface needs to have a lower error rate than the converter's core CER, and therefore cannot be ignored. If this is not the case, then the data output transmission error will swamp out the CER and become the dominant error contributor. System designers really don't care from what portion of the ADC the error came. However, for the purposes of this discussion, we will focus just on the ADC conversion error rate.

METASTABILITY

A typical cause of conversion error in a high speed ADC is a phenomenon called metastability. High speed ADCs often use many laddered comparators in different stages of the conversion process from the analog signal to a digital value. When a comparator is unable to make a decision on whether an analog input is above or below its reference point, a metastable outcome occurs that may cause an error code. This can happen when the difference between the two comparator inputs is very small or zero in magnitude, and a correct comparison cannot be made. As this incorrect value propagates down the pipeline, a significant conversion error can be output from the ADC as a result.

When the differential analog input is either relatively large positively or negatively, the comparator can quickly resolve the difference and make a clear decision. When the differential value is very small or zero, the time duration for the comparator decision is much longer. If the comparator output is latched before this decision point, a metastable outcome is created.

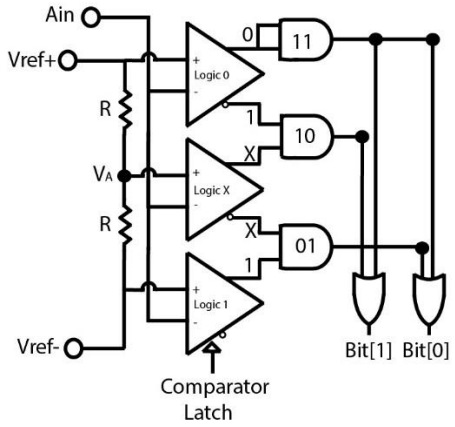

Figure 1. A Basic Laddered Comparator Design Showing a Probabilistic Point of Conversion Failure at Comparator Decision Points (Metastability)—In a Case Where A_{IN} = VA, the Middle Comparator May Be Unable to Discern a Stable Output Within a Finite Conversion Time Such That Bit[1] and Bit[0] Can Have Many Possible Error Combinations

Fortunately, there are some design remedies to mitigate this issue. The first, and most obvious, method is to design the comparator so that the indeterminate region is very small, forcing the comparator to make an accurate decision for the largest possible range of analog input conditions. This can cost circuit power and design area.

A second method is to delay the comparator sample time until the last possible moment, which gives the analog input the maximum amount of time to settle into a known comparator output value. However, there are limits to this method, as the delay can only be so long before time expires on the current sample and the comparator must move to resolve the next sample.

A third method is to have an intelligent error detection and correction algorithm that digitally compensates for the indecision by the comparator in later stages of the high speed ADC conversion process. When no comparator decision is made during the maximum allowed time period, the absence can be detected with logic. This information can then be appended to the sample in question for future internal adjustments. When this alert is identified, a postprocessing step can be used to correct for the error before the sample is output from the converter. This can be seen in Figure 2 for the AD9625, a 12-bit, 2.5 GSPS ADC from Analog Devices.

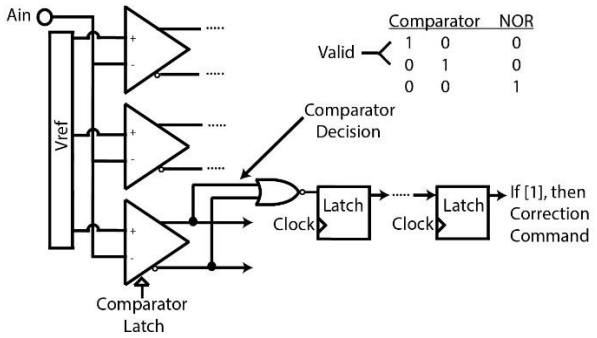

Figure 2. A Comparator's Indecision Can Be Identified Within the Analog-to-Digital Conversion Process of the AD9625; a Correction Command Can Be Executed in Future Steps to Correct the Sample with the Error Before It Is Output from the Converter

CONFIDENCE LEVEL

A CER confidence level (CL) is an extrapolated expectation of an error in the future, despite not measuring to a certain failure rate. This allows a reduction in the total number of samples taken for a given CER, at the expense of having less than 100% certainty. Measuring to an absolute 100% certainty would mathematically require an infinite duration of samples. Therefore, an industry rule of thumb is that a 95% confidence level is relatively close enough to a known value with a balance between some uncertainty and actual measurement time. If testing were to be repeated one hundred times, we would be able to accurately identify the error rate 95% of those times.

It is sometimes mistakenly thought that once an error is detected during testing, the process is over and the final conversion error rate has been found. This is neither accurate nor complete. A conversion error rate with an associated confidence level can be tested with or without errors during the process. However, if errors are detected for a given confidence level, the quantity of measured samples must be increased, compared to a sample count with no errors. This impact can be seen in Figure 3.

The natural logarithmic relationship for confidence level, error rate, and sample count can be represented mathematically with the following equation:

$$N = \frac{1}{CER} \times [-\ln(1-CL) + \ln(\sum_{k=0}^{E} \frac{(N \times CER)^k}{k!})]$$

N = measured sample count

CER = conversion error rate

CL = confidence level

E = error count detected

When there are no errors detected, the equation becomes much simpler, as the right term is equal to zero and the result only depends upon the left term. For a 95% confidence level with no measured errors, we must take only about three times the number of samples as the inverse expected CER. Measuring to a 100% confidence level, where CL = 1.0 for any CER value, mathematically takes an infinite amount of samples (N) as –ln(0) → infinity.

$$N \times CER = -\ln(1 - 0.95) = -\ln(0.05) = 2.996$$

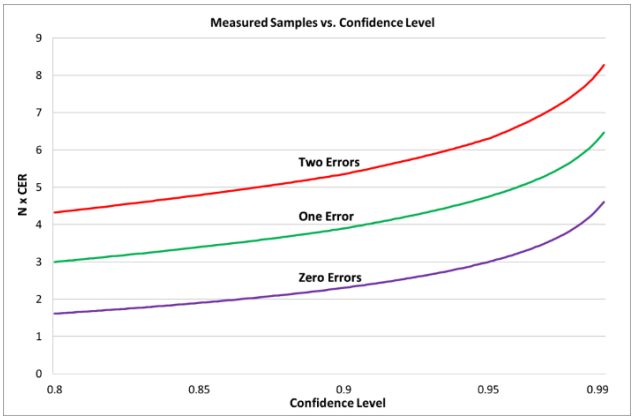

Figure 3. N × CER vs. Confidence Level and Error Detection Count—Notice that CER Testing Can Continue After a Detected Error, but Only at the Expense of an Increased Number of Measured Samples to Achieve the Same Confidence Level

ERROR THRESHOLD

Not all conversion errors within high speed ADCs are created equal. Error magnitude is critical, as some errors are definitely more important than others. For example, a one or two least significant bit (LSB) error may be within the expected noise floor of the system and may not even impact instantaneous performance. However, a most significant bit (MSB) error, or even a full-scale error, could potentially cause a system failure event. Therefore, the CER testing needs to have a mechanism or threshold to grade the severity of the error in the conversion.

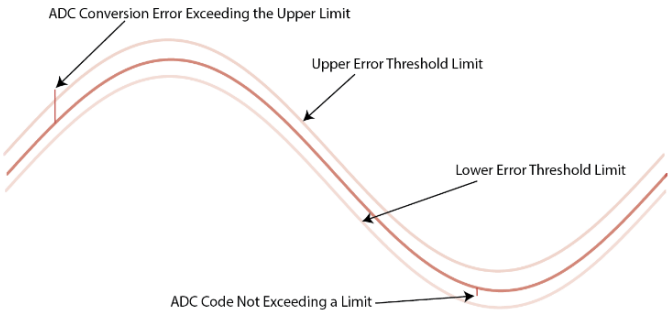

Figure 4. A Reconstructed Sine Wave from ADC Samples Can Be Seen with a Bounded Upper and Lower Threshold Limit—When a Code Exceeds the Limit, It Is Determined to Be a Conversion Error; a Benign Nonlinear Outlier Sample that Is Still Within the Thresholds Is Not Considered a Conversion Error

The error threshold for conversion should include the known linearity imperfections of the ADC, along with the clock jitter and other system noise that are outside the capabilities of the converter. This typically sums cumulatively to four or five least significant bits or 16 to 32 codes of a 14-bit ADC for any given sample. It may be slightly more or less depending upon the ADC resolution, system performance, and the application's error rate requirements. When this error band is used to compare against the ideal value, a sample that exceeds this limit will be counted as a conversion error. In legacy video ADCs, this error was called a "sparkle code," as it produced a bright white pixel flash on the video screen.

The acceptable converter error rate will largely depend upon the requirements of the signal processing system and system tolerance for errors. For example, a user of a backyard mobile Bluetooth® speaker system may be able to tolerate several errors per hour and not even notice. A mission critical sensor aboard a space satellite may need minimal converter ambiguity, or else satellites might start falling from the sky. Well, maybe not to that severity, but very bad things may start to happen, like poor television reception.

Historical measured GSPS ADC conversion error rates have typically been no better than 1e-14. For an error rate of 1e-12, this means that the converter should not produce a conversion error in 1e-12 (one trillion) samples. An error rate of 1e-15 means that the converter should not produce a conversion error within a span of 1e-15 (one quadrillion) samples. With the high sample rates of today's state-of-the-art converter technology, this may seem large, but still manageable to test for CER. However, for a 125 MSPS converter with an 8 ns sample rate, one trillion samples will take 800 seconds (1e-12 × 8 ns) or about 13 minutes. One quadrillion samples would take 800,000 seconds (1e-15 × 8 ns) or 9.24 days. For a 95% confidence level in these error rates, we would need to multiply each of those sample durations by 2.996 respectively.

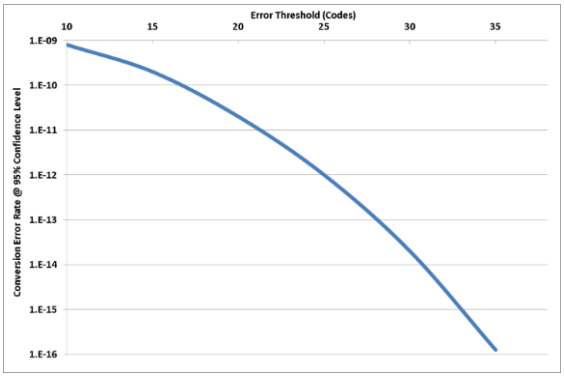

Figure 5. CER vs. Error Magnitude Threshold—the Error Threshold Limit Placed on the Testing (in ADC Codes) Will Have an Impact on the CER at a Given Confidence Level

TESTING FOR CER

The simplistic block diagram below depicts how an internal ADC core can be tested for its CER. A relatively slow frequency sinusoid can be used as the analog input, while sampling at or near the ADC maximum encode rate. The analog input signal is planned so that the expected absolute difference from one sample to the next is ideally no more than one LSB code, neglecting system noise. Ideally, the analog input signal is slightly larger than full-scale, so that all codes of the ADC are exercised. The analog input and the encode sample rate should be computed such that a long cycle of coherency is established and the ADC is not consistently sampling at the same code levels.

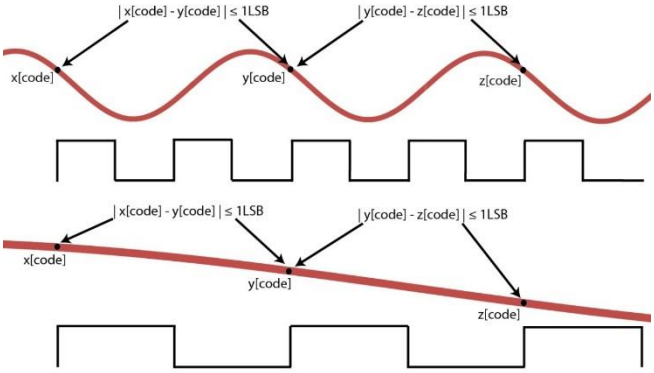

Figure 6. Two Sampling Cases for CER Testing—the Top Case Samples an Analog Signal Just Slightly Faster than Fs/2, Where Only Every Other Sample Is Compared; Two Successive Samples Are Ideally Different by No More Than One LSB Code; The Lower Case Oversamples a Relatively Slow Analog Input Such that Two Adjacent Samples Are also Different by No More Than One LSB Code

A counter is used to track the instances where the magnitude difference between two adjacent samples in time exceeds a threshold limit, counting that instance as a conversion error. The counter must keep a cumulative total of errors throughout the duration of the testing. To ensure that the system is working as expected, the magnitude of the error vs. ideal should also be logged. The time needed for the test will be based upon the sample rate, the desired tested conversion error rate, and the confidence level desired.

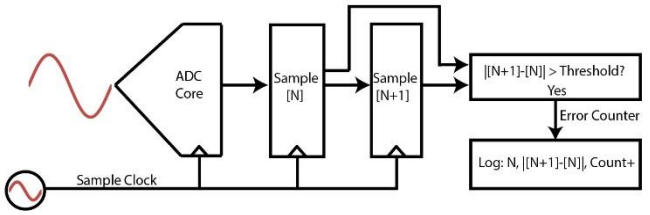

Figure 7. CER Testing Compares Two Successive ADC Samples to a Predetermined Error Threshold; a Counter Logs Error Occurrence, Magnitude, and Sample Location Identifier

MEASURED vs. SIMULATED

In selecting an ADC with a low CER, a system engineer should be able to differentiate between listed specifications that can actually be tested, compared to those that are only based on design simulation cases. For example, a claim for a CER of 1e-18 with a 95% confidence level and no errors using a 1 GSPS ADC must be based either only on circuit simulation or a nearly century long continuous measurement. A simple calculation can help prove this. To measure a CER of 1e-18 to a 95% CL for even a relatively fast 1 GSPS ADC with a 1 ns sample rate, it will consume 2.996 billion seconds (2.996 × 1e-18 × 1 ns), or about 95 years. Do you want your system's ADC conversion error rate specified solely on an extrapolated estimate from simulation, or actual real-world measured results from the lab?

SUMMARY

While differing in concept from digital bit error testing, even GSPS ADC conversion error rate testing can be lengthy to accurately measure. A confidence level for CER testing of less than 100% is needed as it is not feasible to measure indefinitely. ADC samples must be compared against a threshold before determining its significance as a true conversion error. A real-time test system compares adjacent samples for significant excursions that exceed a threshold.

While typical converter architectures may achieve a measured conversion error rate that is acceptable for some systems, new designs and error detection algorithms are pushing the limits to perform even better. A 12-bit, 2.5 GSPS ADC, AD9625, subranging pipeline core from Analog Devices uses a proprietary technique to detect ADC conversion errors in early stages of pipeline processing. It can then go on to resolve and correct for the errors in later stages. This allows for an industry-leading measured CER on 12-bit GSPS ADCs of better than 1e-15 at a CL of 95%.

ABOUT THE AUTHOR

Ian Beavers is an applications engineer for the High Speed ADC team at Analog Devices in Greensboro, NC. He has worked for ADI since 1999 and has over 18 years of experience in the semiconductor industry. Ian earned a bachelor's degree in electrical engineering from North Carolina State University and an M.B.A. from the University of North Carolina at Greensboro. Ian is a member of EngineerZone's High Speed ADC Support Community. Feel free to send your questions to IanB on Analog Devices EngineerZone® Online Technical Support Community.

REFERENCES

BER Calculator, Jitter Time Consulting LLC

Murden, Franklin M. and Michael R. Elliott. "Metastability Error Reduction in Signal Converter Systems." U.S. Patent No. 7,623,051. 2009.

Kester, Walt. MT-011 Tutorial. *Find Those Elusive Sparkle Codes and Metastable States*. Analog Devices, Inc., 2006.

Kester, Walt. MT-020 Tutorial. *ADC Architectures I: The Flash Converter*. Analog Devices, Inc., 2006.

Redd, Justin. "Explaining Those BER Testing Mysteries." *Lightwave Online*, 2004

Redd, Justin. "Calculating Statistical Confidence Levels for Error-Probability Estimates." *Lightwave Online*, 2000.

Wolaver, Dan H. "Measure Error Rates Quickly and Accurately." *Electronic Design*, pp. 89-98, May 30, 1995.

RESOURCES

Share this article on 

One Technology Way • P.O. Box 9106 • Norwood, MA 02062-9106, U.S.A.
Tel: 781.329.4700 • Fax: 781.461.3113 • www.analog.com

TA12902-0-12/14