

Application Note: Understanding Clock Jitter in High-Speed ADCs

Introduction

Clock jitter is often the least understood aspect of ADC performance for many end users. In this application note, we examine how jitter manifests in ADCs, describe its relationship to phase noise, and clarify how to interpret jitter specifications for both ADCs and their associated clock sources. We also outline the fundamental steps for deriving system-level jitter requirements. In addition, we provide practical guidance for diagnosing degraded ADC performance and determining whether clock jitter is the underlying cause. Finally, we offer recommendations for improving jitter performance within an application.

Jitter in ADCs

An ADC can be described as a device that has one signal input and two reference inputs: the reference voltage and the sampling clock. The reference voltage acts as a reference level against which the input signal is measured, while the sampling clock defines the time points of the ADC samples, i.e. determines when the measurement is performed. Generally, the dynamic inaccuracy of these time instants can be defined as jitter. The sampling clock is a logic-level, practically a square wave, signal when it reaches the sampler in the ADC and clock jitter is the deviation of the sampling clock edge from its ideal position.

Jitter can originate from the clock source or from the ADC device itself. In the latter case it is called additive jitter, which is usually specified in the ADC datasheet, where it is often referred to as *aperture uncertainty*. It is also possible that the clock signal gets contaminated when it is transferred from the clock source to the ADC. The ADC clock input has a very wide bandwidth to support clock signals with sharp transitions; therefore, it passes any noise and disturbance practically unfiltered to the sampling circuit.

Besides ADCs, other systems, such as wire line and SerDes links, also suffer from the effects of jitter, but the mechanisms and mitigation strategies are not directly transferrable to ADCs.

Characteristics of Jitter

Thermal and $1/f$ noise present in semiconductor devices and other circuit elements is typically the main source of jitter. Jitter can also be caused by coupling from digital signals in the system or outside of it. The coupling can be capacitive, inductive, or resistive, for instance, via a power supply or a shared ground. RF interference can also show up as clock jitter.

Jitter can be random or deterministic; the physical circuit element level jitter sources fall in the first category. In the frequency domain random jitter shows up as an elevated noise floor. A periodic signal coupling to the sampling clock produces a deterministic pattern of shifts in sampling points. In frequency domain a periodic signal shows up as spurious tones in the clock spectrum. The frequencies can be sub-harmonics of the sampling clock or completely at unrelated frequency.

Noise from digital logic can show up as random or deterministic jitter or a combination of the two.

If the ADC input signal couples to the sampling clock, it produces a replica of itself at two times the signal frequency in the ADC output spectrum, which can easily be confused with second order harmonic distortion.

The MSB of the ADC digital output strongly correlates with the analog input and can have a similar effect, with somewhat less pure frequency. Coupling in the opposite direction, the sampling clock coupling to the ADC input, is usually harmless.

Jitter vs. Phase Noise

Jitter and phase noise are two ways of looking at the same thing. Jitter is a time domain parameter while phase noise is usually used in the frequency domain. Jitter is expressed in units of second. It is an absolute measurement that does not change with clock frequency, for instance, when the clock is divided. In contrast, the phase noise, expressed in radians, is linked to the carrier frequency (the sampling clock) and must be scaled when the clock goes through a division. This is because the phase is the time-integral of the frequency and is intimately linked to it.

When jitter is specified, it is typically just a single number: the root-mean square (RMS) of all time points. The assumption is that jitter follows the normal distribution, and the RMS value is sufficient to describe it. This is typically true when dealing with random jitter but can be inaccurate when the origin of jitter is a digital source. Deterministic jitter is generally not normally distributed in amplitude. Many system models use additive white gaussian (AWG) noise to estimate the effects of noise and jitter to the application and can give misleading results when the assumption is not correct.

Peak-to-peak jitter is a specification used in wireline communications and SerDes but typically not in ADCs.

Jitter originating from digital noise coupling or RF interference can be time varying in nature, for instance, appearing in bursts when the interfering system is active.

Consider a sine wave. Noise that stretches the waveform amplitude up and down is called amplitude noise while noise that moves the waveform left and right (in time or in phase) is called phase noise. When this sinewave is subsequently used as a sampling clock, the phase noise around the waveform zero crossings represents the clock jitter. Whether the zero crossing displacements correlate with one another, or are uncorrelated, affects how the phase noise is spectrally distributed in the frequency domain.

When the phase noise is analyzed in frequency domain, for instance using a spectrum analyzer, a phase noise analyzer, or a simulation tool such as periodic noise analysis, the results are plotted using the frequency offset from the carrier as the X-axis. The Y-axis shows the phase noise in units of dBc/Hz. The use of single side band spectrum is common. When viewed this way, the phase noise plot presents a far more complete picture of the phenomenon than the RMS jitter, which condenses all the information into a single number.

A clock source, such as a PLL, shows a phase noise profile that is not flat but has higher noise level close to the carrier (*close-in phase noise*) eventually levelling off to a flat noise floor beyond certain offset (*far-out phase noise*) from the carrier. This type of colored noise profile with prominent close-in phase noise can be understood in time domain as correlation of the edge displacements: adjacent clock edges move together, the lower the phase noise frequency, the larger the number of edges that are correlated i.e. the longer the disturbance persists in time.

The specifications of clock sources and signal generators tell the phase noise density at one or more frequency offsets from the carrier expressed in radians per Hz or dBc/Hz. The integrated double-sided phase noise, with specified lower and upper integration limits, is often also given. If not, the user can do the calculation using the given noise profile. The phase noise, which is in radians, can be converted to jitter using the fact that one period

of the carrier is $1/f$ seconds long in time and equivalent to 2π radians in phase. It is important to note that the conversion must be done at the carrier frequency the phase noise is specified at, not at the final sampling clock frequency in situations where clock dividers are used.

When comparing the phase noise performance of clock sources that are specified at different frequencies, one must convert the numbers to the same frequency by applying the conversion factor $f1/f2$ or $20*\log_{10}(f1/f2)$, if operating in dB units. When integrated phase noise or jitter numbers are specified, attention must be paid to the integration limits.

How Jitter Affects Sampling

Clock jitter alters the sampling instant of an ADC. The faster the signal moves, the greater the effect. Sine wave is the most used test signal in ADC characterization. Its rate of change is fastest around the zero crossings. Furthermore, the higher the frequency and amplitude of the signal, the steeper the slope, making the signal more susceptible to the effects of jitter. The relationship between the jitter sensitivity and both the signal frequency and the amplitude is linear.

When increasing the signal amplitude, the jitter-induced noise floor follows the signal level because they both grow at the same rate, keeping the jitter-associated signal-to-noise ratio (SNR) constant. Increasing the signal frequency makes the effect of jitter worse, increasing the jitter-generated noise by 6dB every doubling of the signal frequency.

Most real-world signals are less sensitive to jitter than the sinewave, which has a relatively low peak-to-average ratio (PAR) of 3dB. Jitter calculations made with sinewaves can be adjusted by the difference of the PAR of the two signals. When dealing with wide band signals, the average frequency should be used.

When studied in the frequency domain, the sampling operation performed by the ADC is mathematically a convolution. If the test signal is a sinewave, this produces a spectral copy of the phase noise profile of the clock around the signal tone. The spectral density of the phase noise is scaled by the ratio of the frequencies of the input sine wave and the sampling clock. The same scaling applies to tonal sidebands manifesting as periodic jitter.

Sampled systems exhibit a phenomenon called aliasing. In the case of ADCs this folds the entire infinitely long continuous time frequency axis to the range between 0 and $F_s/2$, where F_s is the ADC sampling frequency. This means that the clock phase noise at frequencies above $F_s/2$ cannot be ignored. All the phase noise up to the noise bandwidth of the ADC clock input, which can be several GHz wide, must be considered.

In many applications, the frequency profile of the clock phase noise matters. Usually the close-in phase noise up to a certain frequency offset can be ignored. This is very application dependent. For instance, in wireless communication systems the phase noise spreads the signal power of strong channels into the adjacent channels. In this case, the phase noise closer than some fraction of the channel spacing/bandwidth can be ignored. Demodulators can usually track and compensate for slow phase variations. The update rate of this loop defines the limit below which the phase noise can be ignored.

Identifying Jitter Effects with ADC Measurements

Sometimes one comes across a situation where ADC yields degraded performance without a clear root cause, and jitter becomes the first thing to investigate. This section provides some practical steps that can be used to

gain a better understanding. The outlined tests assume a sinusoidal test signal. Using high signal frequency, which is more sensitive to jitter, makes the effects easier to observe.

One should be aware that even a good RF signal source almost always shows harmonic distortion far exceeding the levels present in a high-performance ADC and often has noise floor exceeding that of the ADC as well. The use of a narrow band-pass filter can fix these issues. Some of the tests require changing the signal frequency, which may necessitate the use of multiple bandpass filters, complicating the analysis by introducing the filter bandwidth and amplitude into the picture.

1. FFT shows spurious tones (spurs), which can be an indication of deterministic jitter. The following steps help to identify some commonly occurring cases.
 - a. Turn off the signal source. If the tone persists, it is not caused by jitter but rather direct coupling to the ADC input signal path.
 - b. If the tone frequency is a harmonic (2x, 3x, 4x, ...) of the signal frequency it is unlikely to be caused by jitter. The second harmonic is an exception that needs to be investigated further. More on that later.
 - c. Because of the aliasing behavior, it can sometimes be tricky to determine if a certain spur is a harmonic tone. One good way to test this is to move the signal frequency by a small amount, say 100kHz, and watch what happens to the spur frequency. Say it moves 500kHz in either direction. That indicates it is the fifth harmonic, not caused by jitter.
 - d. As mentioned earlier, the input signal coupling to the clock produces a tone at two times the signal frequency. When the signal amplitude is changed, say by 3dB, and the spur level follows, the evidence indicates that coupling is indeed happening. The classic behavior of the second harmonic predicts a 6dB HD2 change for a 3dB change in signal amplitude. The ADC nonlinearity effects can be more complex than that simple model, producing a level change that is not exactly 6dB.
 - e. When doing the frequency shift test, the spurious tone moves the same amount as the input frequency, i.e. the frequency offset stays the same. Both side bands ($F_{in} + \Delta f$ and $F_{in} - \Delta f$) are present with similar amplitude levels. Sometimes the aliasing effect can cause confusion, folding one of the side bands, and making it move in the opposite direction. The above is a strong indication that either the amplitude or the phase of the signal is being modulated. In ADCs, the phase modulation happens via the clock.
 - i. Distinguishing between amplitude and phase modulation can be done by observing how the amplitude or the spur behaves when the signal frequency is changed. To make the effect observable, the frequency change needs to be substantial, a large fraction of the signal frequency. Increasing the signal frequency, say, by a factor of 1.5, predicts a 3dB increase in the spur level in the case of phase modulation and no change in the case of amplitude modulation. In general, the tone level produced by phase modulation is expected to show a linear response to the signal frequency.
 - f. A special case of the previous point is the situation where the modulating frequency is a sub harmonic of the sampling clock. For instance, modulation with $F_s/4$ produces replicas of the signal tone on both sides of $F_s/4$.
2. An increased flat noise floor can be an indication of random wide-band jitter. Unfortunately, because of its random nature, tracing down the root cause can be trickier than in the case of spurs. The ADC and the presented input signal have their own noise floors. Separating those from the one caused by jitter can be difficult.
 - a. Turn off the signal source. If the noise floor doesn't change, the cause is not jitter.
 - b. Change the signal amplitude and observe how the noise floor behaves. Start from a low signal amplitude and start increasing it towards the ADC input full scale. When approaching the ADC full scale, the noise floor starts to rise, eventually following the signal level a dB by dB. This strongly

indicates the possibility of a jitter effect. While amplitude modulation via the ADC voltage reference is also possible, this path is typically not a wide band one and unlikely to produce a noise floor increase over a wide frequency range.




Figure 1. Noise floor follows the signal amplitude when it is dominated by jitter.

- c. Further evidence for the jitter effect can be collected by repeating the previous amplitude test using different signal frequencies. Doubling the signal frequency should result in a 6dB higher jitter limited noise floor. Carrying out the test at a very low signal frequency gives a good idea how the ADC noise floor behaves in the absence of jitter. Some noise floor increase with signal amplitude is expected as the increased circuit activity in the ADC shows its effects.
3. Signal tone shows a skirt around it.
 - a. The skirt can be just spectral leakage from the FFT, which can be significantly reduced by windowing. Use coherent sampling, if possible. Lock the signal generator and the clock source to the same time base and select a signal frequency that falls exactly in an FFT bin: $F_{in} = \text{round}(f/F_s * \text{FFT_len}) * F_s / \text{FFT_len}$, where F_{in} is f rounded to the closest bin frequency. Even when locking the time bases is not practical, selecting the signal frequency this way is often helpful.
 - b. When a bandpass filter is used to clean up the signal generator output, the observed skirt may simply be the signal generator's noise floor appearing within the filter passband. Verify that the skirt width is consistent with the bandwidth of the filter.
 - c. Change the signal frequency. If the noise spectral density in the skirt region increases and decreases with the signal frequency, clock phase noise is the likely cause. It is possible, though, that what is seen is the phase noise of the signal generator. Check its specifications to rule it out. If the noise density that doesn't change points to something else, possibly amplitude modulation via the ADC reference voltage.
 4. Analyzing jitter in time domain gives another perspective that can be helpful. A sinewave fitting algorithm can be used to fit a captured ADC waveform to an ideal sine wave. By subtracting the ideal signal from the measured one isolates the error signal, which can then be plotted against the signal level or observed as a function of time.
 - a. When the error magnitude plotted against the signal level shows a pattern that is wide in the middle, around sine wave zero crossings where the signal rate of change is the fastest, there is a strong indication of the presence of jitter.

- b. Using a beat frequency signal, for instance, one with its frequency close to F_s , produces an output waveform that appears to have a low frequency. Looking at the time domain error waveform alongside the ideal waveform can reveal the same effect as the observation versus the signal level: increased noise around the areas where the waveform is the steepest. Alternatively, by observing only every other sample, a similar visual effect can be obtained with a beat frequency around $F_s/2$.

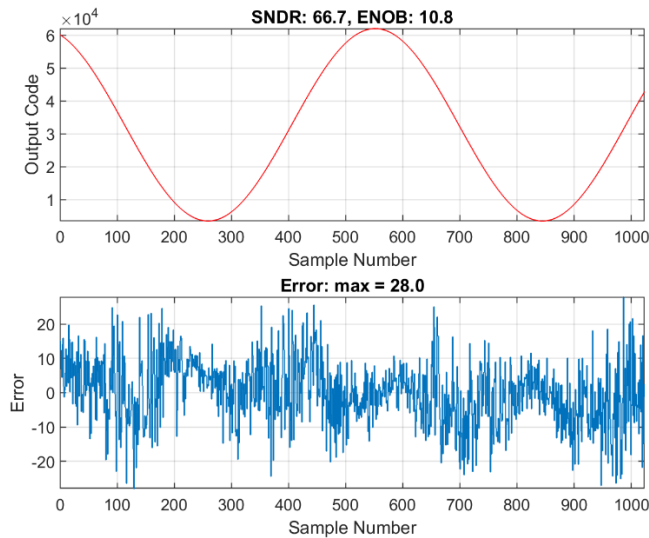


Figure 2. Beat frequency signal sampled with a jittery clock and fitted into a sine wave shows increased noise around the steepest parts of the waveform.

- c. Time domain investigation of a case that has already been determined to be jitter related can reveal something new. A bursty error pattern can indicate noise coupling or RF interference from a source that has time varying activity level. A PLL clock source that has difficulties staying in lock can show somewhat similar effect.

Estimating Jitter from ADC Measurements

The jitter limited SNR can be calculated as $SNR_j = -20 * \log_{10}(2\pi * F_{in} * t_j)$, where t_j is the RMS jitter. Measuring the ADC SNR and solving t_j from the equation can give an upper estimate for the jitter. Creating a situation where the jitter effect is strong, by using a high signal frequency, can improve this estimate.

A more effective way to improve the estimate is to separate the jitter effect from the other noise sources. This can be done knowing that the measured SNR can be expressed as $signal/(noise_{ADC} + noise_{JITTER})$, where the noise sources are added in power. The ADC noise can be obtained with a second measurement in conditions where the jitter effect is made small, for instance by reducing the signal amplitude or the signal frequency.

Assuming that $noise_{ADC}$ and t_j are independent from the signal amplitude and frequency and using $noise_{JITTER} = t_j * F_{in} * 2\pi * A_{in}$, the two unknowns can be solved with two measurements taken at different amplitude levels and/or with different signal frequencies:

$$t_j = \sqrt{\left(\left(\frac{A_1}{snr_1}\right)^2 - \left(\frac{A_2}{snr_2}\right)^2\right) / \left((2\pi)^2 \cdot (F_{in1}^2 A_1^2 - F_{in2}^2 A_2^2)\right)}, \quad (1)$$

where snr is in linear units i.e., $snr = 10^{(SNR/20)}$. The equation assumes that $A_1 \geq A_2$ and $F_{in1} \geq F_{in2}$.

There are algorithms that use sine wave fitting to estimate jitter. First, the nonlinear effects are estimated and subtracted out, leaving just the signal with residual noise. This waveform is fitted to a sine wave. Next the error signal is analyzed point by point and the amplitude noise and the phase noise are separated with the aid of the instantaneous phase of the signal. This works because at the peaks of the sinewave the rate of change is zero, eliminating the effect of phase noise there. Everywhere else the measured noise is a sum of the amplitude noise and the phase noise (or jitter induced noise) modulated by the time derivative of the signal, which for the sinewave is the corresponding cosine wave. The accuracy of the method depends heavily on how well the nonlinear effects have been subtracted out.

Ensuring Good Jitter Performance

Know What is Needed

The first step in achieving a good jitter performance is to understand what is needed. Ideally, one has a system level model that includes jitter and phase noise effects and yields separate requirements for close in phase noise, wideband jitter / phase noise, and ADC SNR.

If that is not the case, at the bare minimum, one must know what the SNR required by the application is. Then, using a graph like the one below, which shows the combined effect of jitter and ADC noise for an ideal 12-bit ADC, one can come up with the jitter requirement. The graph is calculated for a full-scale sinewave, which is a very pessimistic scenario compared to most real-world signals, except for systems limited by QPSK modulated or CW blockers.




Figure 3. Clock jitter effect on ADC SNR.

A more realistic requirement can be obtained by considering how jitter sensitive the input signal is. This is affected by the PAR of the signal and how it compares the 3dB PAR of a sinewave. Wideband signals should be represented by the average signal frequency instead of the highest frequency. For instance, a 30MHz wide signal that spans from 110MHz to 140MHz and has PAR of 10dB has jitter sensitivity equivalent to a 125MHz sine wave with -7dBFS amplitude. To be still able to use the graph provided, the equivalent full-scale sine wave frequency can be calculated as $125\text{MHz} * 10^{(-7/20)} = 56\text{MHz}$. Note the big difference to the commonly applied, overly pessimistic, estimate using a full-scale sine wave at the highest frequency of the band.

Make sure that ADC aperture uncertainty or additive jitter specification meets the requirements. The ADC additive jitter is typically dominated by wideband jitter with flat frequency profile.

Select a Clock Source Meeting the Requirements

The next step is to select the clock source. In some cases, the system may dictate the type of clock source, for instance a PLL vs. a crystal clock module. A digital clock is not a suitable sampling clock for a high-speed, high-performance ADC.

If PLL is used as a clock source, one must make sure that the phase noise profile is suitable for the application. This information can be found from the datasheet.

A crystal clock module combines a crystal oscillator and a buffer amplifier in the same package. Models with CMOS and differential outputs are available. One should be aware that the specified jitter is usually integrated up to a very low frequency offset such as 1MHz or 10MHz while the ADC clock input can have noise bandwidth of 1GHz or more. On the other hand, the integration often includes very close in phase noise which might be irrelevant to the application. A carefully selected XTAL clock module can deliver a very good jitter performance. The selection process often involves evaluating the performance of candidate parts in a lab setup together with the ADC.

When evaluating ADCs, one should be aware that a RF signal generator with its sinusoidal output is not a low jitter clock source at low frequencies due to the slow transitions of the sine wave. Using high amplitude and clipping the waveform with back-to-back diodes (which are often included in ADC evaluation kits) can improve the situation. Even at high frequencies it may be necessary to bandpass filter the clock to achieve sufficient purity. If the ADC includes a clock divider, it is recommended to input a higher frequency clock and divide it down in the ADC.

Proper Delivery of the Clock to the ADC

Once the clock source has been selected the clock signal has to be delivered to the ADC without introducing jitter. The clock must have sufficient amplitude and sharp transitions and be differential, if possible.

When the sampling clock enters the ADC, it goes through a clock receiver circuit that is typically a differential amplifier. While this amplifier has the requisite gain to amplify even a very low amplitude signal to a sufficient level to be used as a sampling clock in the ADC core, it adds its own noise to it. The edge rate of the received clock signal around the zero crossing determines conversion factor from noise to jitter. The ADC additive jitter specification assumes an input clock with sharp transitions.

The clock routing should be shielded from switching signal lines and a solid ground plane should be placed underneath the traces.

If the ADC has a clock divider, sending in a higher frequency clock and dividing it down in the ADC can sometimes yield better jitter.

Self-interference from ADC digital outputs can find its way to the clock input. LVDS interface should be used, if possible. Driving long lines or large capacitance with CMOS outputs should be avoided, which can be accomplished by adding a buffer amplifier with its power supply separated from the ADC. The output lines should have solid ground underneath to serve as current return path. The ADC power supply must be properly

decoupled and preferably provided with a dedicated linear regulator that is not shared with other blocks in the system.

System level frequency planning can help with spurious tones. The ADC is immune to coupling at the sampling frequency and its integer multiples, while divided versions of the ADC clock and unrelated frequencies can result in phase modulation via the clock input. This should be considered in selecting system level clock frequencies.

Revision History

Version	Date	Comment
1.0	Mar. 12, 2026	Initial Release.
2.0	Mar. 19, 2026	Added formula for estimating jitter from two measurements.

Silanna Semiconductor Proprietary and Confidential

Information furnished by Silanna Semiconductor is believed to be accurate and reliable. However, no responsibility is assumed for its use. Silanna Semiconductor makes no representation that the interconnection of its circuits as described herein will not infringe on existing patents rights. Specifications subject to change without notice.