AC 2011-2914: EVALUATING OSCILLOSCOPE SAMPLE RATES VS. SAMPLING FIDELITY

Johnnie Lynn Hancock, Agilent Technologies

About the Author

Johnnie Hancock is a Product Manager at Agilent Technologies Digital Test Division. He began his career with Hewlett-Packard in 1979 as an embedded hardware designer, and holds a patent for digital oscilloscope amplifier calibration. Johnnie is currently responsible for worldwide application support activities that promote Agilent's digitizing oscilloscopes and he regularly speaks at technical conferences worldwide. Johnnie graduated from the University of South Florida with a degree in electrical engineering. In his spare time, he enjoys spending time with his four grandchildren and restoring his century-old Victorian home located in Colorado Springs.

Contact Information:

Johnnie Hancock Agilent Technologies 1900 Garden of the Gods Rd Colorado Springs, CO 80907 USA +1 (719) 590-3183 johnnie_hancock@agilent.com

Evaluating Oscilloscope Sample Rates vs. Sampling Fidelity: How to Make the Most Accurate Digital Measurements

Introduction

Digital storage oscilloscopes (DSO) are the primary tools used today by digital designers to perform signal integrity measurements such as setup/hold times, rise/fall times, and eye margin tests. High performance oscilloscopes are also widely used in university research labs to accurately characterize high-speed digital devices and systems, as well as to perform high energy physics experiments such as pulsed laser testing. In addition, general-purpose oscilloscopes are used extensively by Electrical Engineering students in their various EE analog and digital circuits lab courses.

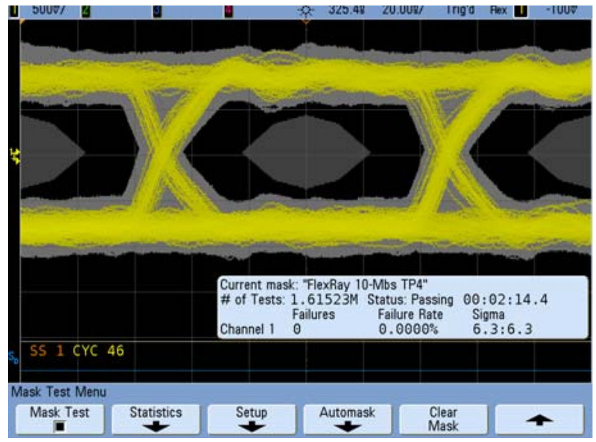

The two key banner specifications that affect an oscilloscope's signal integrity measurement accuracy are bandwidth and sample rate. Most engineers and EE professors have a good idea of how much bandwidth they need for their digital measurements. However, there is often a lot confusion about required sample rate — and engineers often assume that scopes with the highest sample rate produce the most accurate digital measurements. But is this true?

When you select an oscilloscope for accurate, high-speed digital measurements, sampling fidelity can often be more important than maximum sample rate. Using side-by-side measurements on oscilloscopes with various bandwidths and sample rates, this paper demonstrates a counterintuitive concept: scopes with higher sample rates *can* exhibit poorer signal fidelity because of poorly aligned interleaved analog-to-digital converters (ADCs). This paper also will show how to easily characterize and compare scope ADC sampling fidelity using both time-domain and frequency-domain analysis techniques.

In the field of academics, this paper can be first applied as a practical application and demonstration of theories presented in courses on digital signal processing. Secondly, when selecting high performance test equipment for electrical engineering and physics research labs, this paper will provide tips on how to select and evaluate digital storage oscilloscopes for accurate reproduction of captured high-speed signals.

Let's begin with a discussion of minimum required sample rate and a review of Nyquist's sampling theorem.

Nyquist's Sampling Theorem

How much sample rate do you need for your digital measurement applications? Some engineers have total trust in Nyquist and claim that just 2X sampling over the scope's bandwidth is sufficient. Other engineers don't trust digital filtering techniques based on Nyquist criteria and prefer that their scopes sample at rates that are 10X to 20X over the scope's bandwidth specification. The truth actually lies somewhere in between. To understand why, you must have an understanding of the Nyquist theorem and how it relates to a scope's frequency response. Dr. Harry Nyquist (Figure 1) postulated:

Figure 1: Dr. Harry Nyquist, 1889-1976, articulated his sampling theorem in 1928.

> **Nyquist Sampling Theorem**
>
> For a limited bandwidth signal with a maximum frequency f_{MAX}, the equally spaced sampling frequency f_S must be greater than twice of the maximum frequency f_{MAX}, in order to have the signal be uniquely reconstructed without aliasing.

Nyquist's sampling theorem can be summarized into two simple rules — but perhaps not-so-simple for DSO technology.

1. The highest frequency component sampled *must* be less than half the sampling frequency.
2. The second rule, which is often forgotten, is that samples *must* be equally spaced.

What Nyquist calls f_{MAX} is what we usually refer to as the Nyquist frequency (f_N), which is *not* the same as oscilloscope bandwidth (f_{BW}). If an oscilloscope's bandwidth is specified exactly at the Nyquist frequency (f_N), this implies that the oscilloscope has an ideal brick-wall response that falls off exactly at this same frequency, as shown in Figure 2. Frequency components below the Nyquist frequency are perfectly passed (gain =1), and frequency components above the Nyquist frequency are perfectly eliminated. Unfortunately, this type of frequency response filter is impossible to implement in hardware.

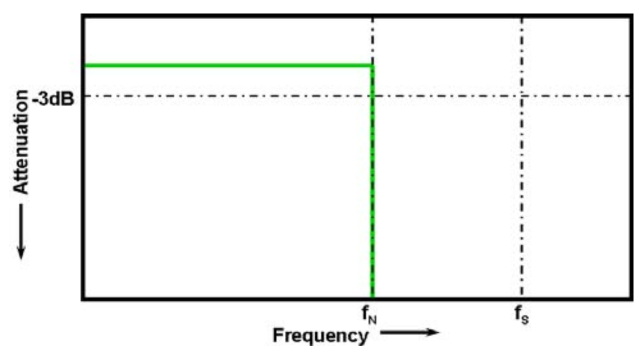

Figure 2: Theoretical brick-wall frequency response

Most oscilloscopes with bandwidth specifications of 1 GHz and below have what is known as a Gaussian frequency response. As signal input frequencies approach the scope's specified bandwidth, measured amplitudes slowly decrease. Signals can be attenuated by as much as 3 dB (~30%) at the bandwidth frequency. If a scope's bandwidth is specified exactly at the Nyquist frequency (f_N), as shown in Figure 3, input signal frequency components above this frequency – although attenuated by more than 3 dB — can be sampled (red hashed area) — especially when the input signal contains fast edges, as is often the case when you are measuring digital signals. This is a violation of Nyquist's first rule.

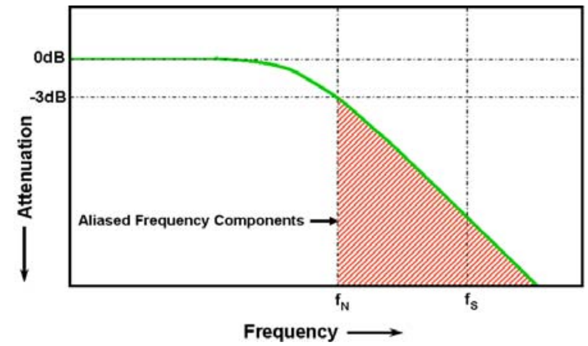

Figure 3: Typical oscilloscope Gaussian frequency response with bandwidth (f_{BW}) specified at the Nyquist frequency (f_N)

Most scope vendors don't specify their scope's bandwidth at the Nyquist frequency (f_N) – but some do. However, it is very common for vendors of waveform recorders/digitizers to specify the bandwidth of their instruments at the Nyquist frequency. Let's now see what can happen when a scope's bandwidth is the same as the Nyquist frequency (f_N).

Figure 4 shows an example of a 500-MHz bandwidth scope sampling at just 1 GSa/s while operating in a three- or four-channel mode. Although the fundamental frequency (clock rate) of the input signal is well within Nyquist's criteria, the signal's edges contain significant frequency components well beyond the Nyquist frequency (f_N). When you view them repetitively, the edges of this signal appear to "wobble" with varying degrees of pre-shoot, over-shoot, and various edge speeds. This is evidence of aliasing, and it clearly demonstrates that a sample rate-to-bandwidth ratio of just 2:1 is insufficient for reliable digital signal measurements.

Figure 4: 500-MHz bandwidth scope sampling at 1 GSa/s produces aliased edges

So, where should a scope's bandwidth (f_{BW}) be specified relative to the scope's sample rate (f_S) and the Nyquist frequency (f_N)? To minimize sampling significant frequency components above the Nyquist frequency (f_N), most scope vendors specify the bandwidth of their scopes that have a typical Gaussian frequency response at 1/4th to 1/5th, or lower, than the scope's real-time sample rate, as shown is Figure 5. Although sampling at even higher rates relative to the scope's bandwidth would further minimize the possibility of sampling frequency components beyond the Nyquist frequency (f_N), a sample rate-to-bandwidth ratio of 4:1 is sufficient to produce reliable digital measurements.

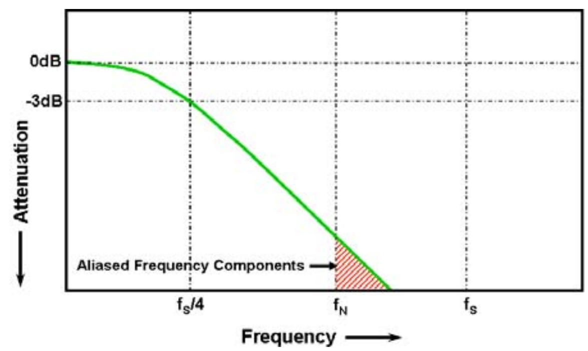

Figure 5: Limiting oscilloscope bandwidth (f_{BW}) to ¼ the sample rate ($f_{S/4}$) reduces frequency components above the Nyquist frequency (f_N)

Oscilloscopes with bandwidth specifications in the 2-GHz and higher range typically have a sharper frequency roll-off response/characteristic. We call this type of frequency response a "maximally-flat" response. Since a scope with a maximally-flat response approaches the ideal characteristics of a brick-wall filter, where frequency components beyond the Nyquist frequency are attenuated to a higher degree, not as many samples are required to produce a good representation of the input signal using digital filtering. Vendors can theoretically specify the bandwidth of scopes with this type of response (assuming the front-end analog hardware is capable) at $f_S/2.5$. However, most scope vendors have not pushed this specification beyond $f_S/3$.

Figure 6 shows a 500-MHz bandwidth scope capturing a 100-MHz clock signal with edge speeds in the range of 1 ns (10% to 90%). A bandwidth specification of 500 MHz would be the minimum recommended bandwidth to accurately capture this digital signal. This particular scope is able to sample at 4 GSa/s in a 2-channel mode of operation, or 2 GSa/s in a three- or four-channel mode of operation. Figure 6 shows the scope sampling at 2 GSa/s, which is twice the Nyquist frequency (f_N) and four times the bandwidth frequency (f_{BW}). This shows that a scope with a sample rate-to-bandwidth ratio of 4:1 produces a very stable and accurate representation of the input signal. And with Sin(x)/x waveform reconstruction/interpolation digital filtering, the scope provides waveform and measurement resolution in the 10s of picoseconds range. The difference in waveform stability and accuracy is significant compared to the example we showed earlier (Figure 4) with a scope of the same bandwidth sampling at just twice the bandwidth (f_N).

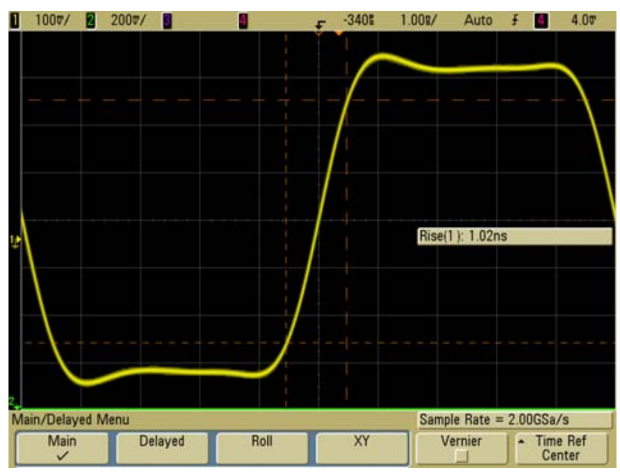

Figure 6: A 500-MHz bandwidth scope sampling at 2 GSa/s shows an accurate measurement of this 100-MHz clock with a 1-ns edge speed

So what happens if we double the sample rate to 4 GSa/s in this same 500-MHz bandwidth scope (f_{BW} x 8)? You might intuitively believe that the scope would produce significantly better waveform and measurement results. But as you can see in Figure 7, there is some improvement, but it is minimal. If you look closely at these two waveform images (Figure 6 and Figure 7), you can see that when you sample at 4 GSa/s (f_{BW} x 8), there is slightly less pre-shoot and over-shoot in the displayed waveform. But the rise time measurement shows the same results (1.02 ns). The key to this slight improvement in waveform fidelity is that additional error sources were *not* introduced when the sample-rate-to-bandwidth ratio of this scope increased from 4:1 (2 GSa/s) to 8:1 (4 GSa/s). And this leads us into our next topic: What happens if Nyquist's rule 2 is violated? Nyquist says that samples *must* be evenly spaced. Users often overlook this important rule when they evaluate digital storage oscilloscopes.

Figure 7: A 500-MHz bandwidth scope sampling at 4 GSa/s produces minimal measurement improvement over sampling at 2 GSa/s

Interleaved Real-Time Sampling

When ADC technology has been stretched to its limit in terms of maximum sample rate, how do oscilloscope vendors create scopes with even higher sample rates? The drive for higher sample rates may be simply to satisfy scope users' perception that "more is better" — or higher sample rates may actually be required to produce higher-bandwidth real-time oscilloscope measurements. But producing higher sample rates in oscilloscopes is not as easy as simply selecting a higher sample rate off-the-shelf analog-to-digital converter.

A common technique adopted by all major scope vendors is to interleave multiple real-time ADCs. But don't confuse this sampling technique with interleaving samples from repetitive acquisitions, which we call "equivalent-time" sampling.

Figure 8 shows a block diagram of a real-time interleaved ADC system consisting of two ADCs with phase-delayed sampling. In this example, ADC 2 always samples ½ clock period after ADC 1 samples. After each real-time acquisition cycle is complete, the scope's CPU retrieves the data stored in each ADC acquisition memory and then interleaves the samples to produce the real-time digitized waveform with twice the sample density (2X sample rate).

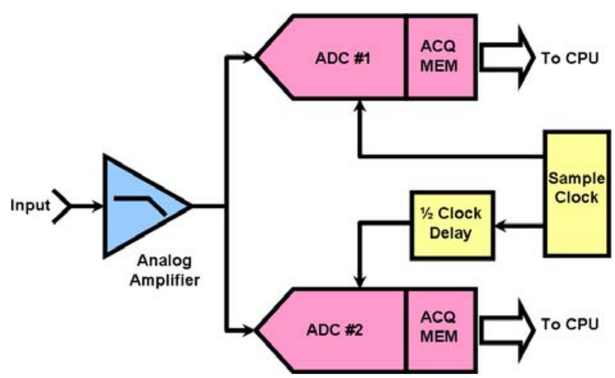

Figure 8: Real-time sampling system consisting of two interleaved ADCs

Scopes with real-time interleaved sampling must adhere to two requirements. For accurate distortion-free interleaving, each ADC's vertical gain, offset and frequency response must be closely matched. Secondly, the phase-delayed clocks must be aligned with high precision in order to satisfy Nyquist's rule 2 that dictates equally-spaced samples. In other words, the sample clock for ADC 2 must be delayed precisely 180 degrees after the clock that samples ADC 1. Both of these criteria are important for accurate interleaving. However, for a more intuitive understanding of the possible errors that can occur due to poor interleaving, the rest of this paper will focus on errors due to poor phase-delayed clocking.

The timing diagram shown in Figure 9 illustrates incorrect timing of interleaved samples if the phase-delayed clock system of two interleaved ADCs is not exactly ½ sample period delayed relative to each other. This diagram shows where real-time digitized points (red dots) are actually converted relative to the input signal. But due to the poor alignment of phase-delayed clocking (purple waveforms), these digitized points are not evenly spaced, thus a violation Nyquist's second rule.

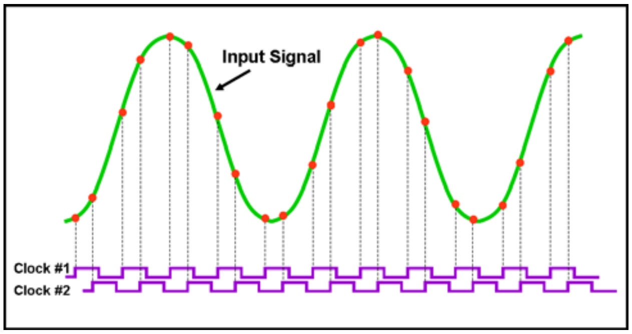

Figure 9: Timing diagram showing non-evenly spaced samples

When the scope's CPU retrieves the stored data from each ADC's acquisition memory, it assumes that samples from each memory device *are* equally spaced. In an attempt to reconstruct the shape on the original input signal, the scope's Sin(x)/x reconstruction filter produces a severely distorted representation of the signal, as shown in Figure 10.

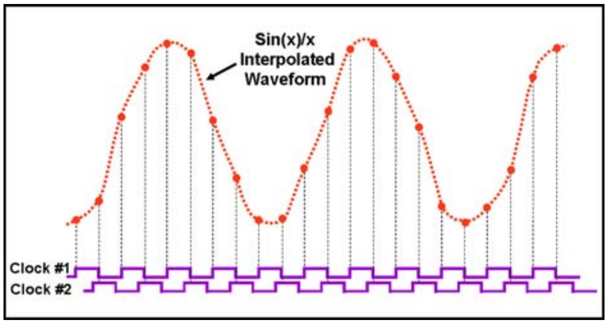

Figure 10: Timing diagram showing distorted reconstruction of waveform using Sin(x)/x filter due to poor phase-delayed clocking

Since the phase relationship between the input signal and the scope's sample clock is random, real-time sampling distortion, which is sometimes referred to as "sampling noise," may be interpreted mistakenly as random noise when you are viewing repetitive acquisitions. But it is not random at all. It is deterministic and directly related to harmonics of the scope's sample clock.

Testing for Interleave Distortion

Unfortunately, oscilloscope vendors do not provide their customers with a specification in their DSO data sheets that directly quantifies the quality of their scope's digitizing process. However, there are a variety of tests that you can easily perform to not only measure the effect of sampling distortion, but also identify and quantify sampling distortion. Here is a list of tests you can perform on scopes to detect and compare interleave distortion:

Interleave distortion tests

1. **Effective number of bits analysis using sine waves**
2. **Visual sine wave test**
3. **Spectrum analysis**
4. **Measurement stability**

Effective number of bits analysis

The closest specification that some scope vendors provide to quantify sampling fidelity is effective number of bits (ENOB). But ENOB is a composite specification consisting of several error components including input amplifier harmonic distortion and random noise. Although an effective number of bits test can provide a good benchmark comparison of overall accuracy between scopes, effective bits is not a very well understood concept, and it requires exporting digitized data to a PC for number crunching. Basically, an effective number of bits test first extracts a theoretical best-fit sinusoidal signal from the digitized sine wave. This sine wave curve-fit algorithm will eliminate any errors induced by oscilloscope amplifier gain and offset inaccuracies. The test then computes the RMS error of the digitized sine wave relative to the ideal/extracted sine wave over one period. This RMS error is then compared to the theoretical RMS error that an ideal ADC of "N" bits would produce. For example, if a scope's acquisition system has 5.3 effective bits of accuracy, then it generates the same amount of RMS error that a perfect 5.3-bit ADC system would generate.

A more intuitive and easier test to conduct to see if a scope produces ADC interleave distortion is to simply input a sine wave from a high-quality signal generator with a frequency that approaches the bandwidth of the scope. Then just make a visual judgment about the purity of the shape of the digitized and filtered waveform.

ADC distortion due to misalignment can also be measured in the frequency domain using a scope's FFT math function. With a pure sine wave input, the ideal/non-distorted spectrum should consist of a single frequency component at the input frequency. Any other spurs in the frequency spectrum are components of distortion. You also can use this technique on digital clock signals, but the spectrum is a bit more complex, so you have to know what to look for.

Another easy test you can perform is to compare parametric measurement stability, such as the standard deviation of rise times, fall times, or Vp-p, between scopes of similar bandwidth. If interleave distortion exists, it will produce unstable measurements — just like random noise.

Visual sine wave comparison tests

Figure 11 shows the simplest and most intuitive comparative test – the visual sine wave test. The waveform shown in Figure 11a is a single-shot capture of a 1-GHz sine wave using an Agilent 1-GHz bandwidth scope sampling at 4 GSa/s. This scope has a sample-rate-to-bandwidth ratio of 4:1 using non-interleaved ADC technology. The waveform shown in Figure 11b is a single-shot capture of the same 1-GHz sine wave using another vendor's 1-GHz bandwidth scope sampling at 20 GSa/s. This scope has a maximum sample-rate-to-bandwidth ratio of 20:1 using interleaved technology.

Figure 11a: 1-GHz sine wave captured on an Agilent 1-GHz bandwidth oscilloscope sampling at 4 GSa/s

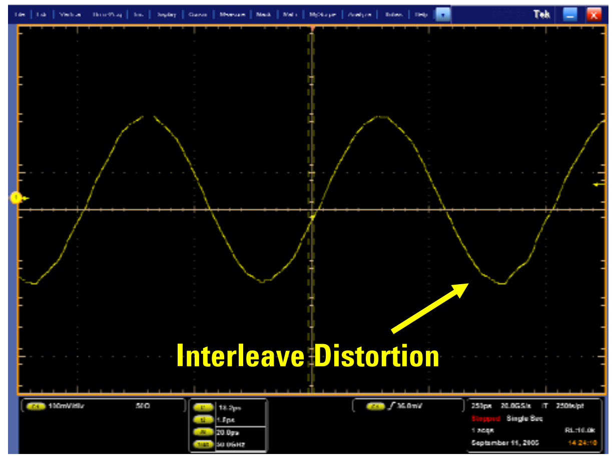

Figure 11b: 1-GHz sine wave captured on another vendor's 1-GHz bandwidth oscilloscope sampling at 20 GSa/s

Although we would intuitively believe that a higher-sample-rate scope of the same bandwidth should produce more accurate measurement results, we can see in this measurement comparison that the lower sample rate scope actually produces a much more accurate representation of the 1-GHz input sine wave. This is *not* because lower sample rates are better, but because poorly aligned interleaved real-time ADCs negate the benefit of higher sample rates.

Precision alignment of interleaved ADC technology becomes even more critical in higher-bandwidth and higher-sample-rate scopes. Although a fixed amount of phase-delayed clock error may be insignificant at lower sample rates, this same fixed amount of timing error becomes significant at higher sample rates (lower sample periods).

Spectrum analysis comparison tests

The visual sine wave test doesn't really prove where the distortion is coming from. It merely shows the effect of various error/components of distortion. However, a spectrum/FFT analysis will positively identify components of distortion including harmonic distortion, random noise, and interleaved sampling distortion. Using a sine wave generated from a high-quality signal generator, there should be only one frequency component in the input signal. Any frequency components other than the fundamental frequency detected in an FFT analysis on the digitized waveform are distortion.

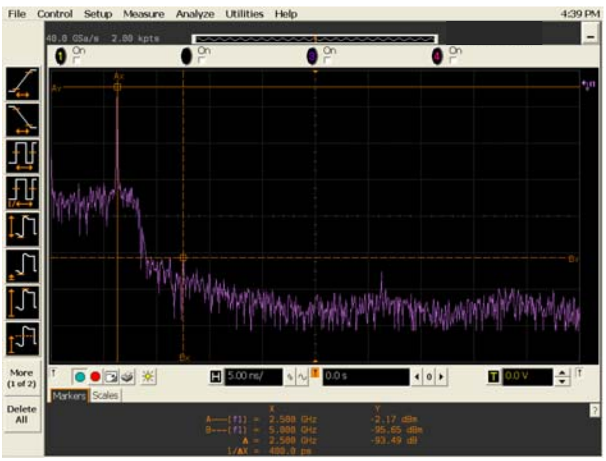

Figure 12a: FFT analysis of 2.5-GHz sine wave captured on an Agilent 2.5 GHz bandwidth scope sampling at 40 GSa/s

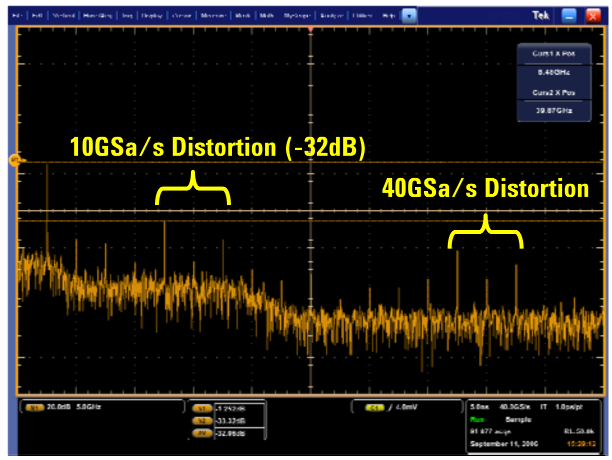

Figure 14b: FFT analysis of 2.5-GHz sine wave captured on another vendor's scope sampling at 40 GSa/s

Figure 12a shows an FFT analysis of a single-shot capture of a 2.5-GHz sine wave using an Agilent 2.5 GHz bandiwth oscilloscope sampling at 40 GSa/s. The worst-case distortion spur measures approximately 90 dB below the fundamental. This component of distortion is actually second harmonic distortion, most likely produced by the signal generator. And its level is extremely insignificant and is even lower than the scope's in-band noise floor.

Figure 12b shows an FFT analysis of a single-shot capture of the same 2.5-GHz sine wave using another vendor's scope — also sampling at 40 GSa/s. The worst-case distortion spur in this FFT analysis measures approximately 32 dB below the fundamental. This is a significant level of distortion and explains why the sine wave test (Figure 13b) produced a distorted waveform. The frequency of this distortion occurs at 7.5 GHz. This is exactly 10 GHz below the input signal frequency (2.5 GHz), but folded back into the positive domain. The next highest component of distortion occurs at 12.5 GHz. This is exactly 10 GHz above the input signal frequency (2.5 GHz). Both of these components of distortion are directly related to the 40-GSa/s sampling clock and its interleaved clock rates (10 GHz). These components of distortion are *not* caused by random or harmonic distortion. They are caused by real-time interleaved ADC distortion.

Summary

As you've read in this paper, there's more to oscilloscope signal fidelity than just sample rate. In some cases a lower-sample-rate scope may produce more accurate measurement results.

To satisfy Nyquist criteria, you need a scope that samples at least 3 to 5 times higher than the scope's bandwidth specification, depending on the scope's frequency roll-off characteristics. Achieving higher sample rates often requires that scope vendors interleave multiple real-time ADCs. But if real-time interleaving is employed, it is critical that the interleaved ADCs be vertically matched and the timing of phase-delayed clocking must be precise. It should be noted that the problem is *not* the number of interleaved ADCs; the issue is the level of precision of interleaving. Otherwise, Nyquist's second rule (equally-spaced samples) can be violated, thereby producing distortion and often negating the expected benefit of higher sample rates.

References

H. Nyquist, "Certain Topics in Telegraph Transmission Theory", Trans. AIEE, vol. 47, pp. 617-644, Apr. 1928

R.J. Marks II, "Introduction to Shannon Sampling and Interpolation Theory", Spinger-Verlag, 1991

Howard Johnson, Martin Graham, "High-Speed Digital Design: A Handbook of Black Magic", Prentice Hall, 1993

Robert A Witte, "Electronic Test Instruments: Theory and Application", Prentice Hall, 1992

Johnnie Hancock, "Evaluating Oscilloscope Bandwidths for your Applications", Agilent Technologies Application Note publication number 5989-5733EN, 2011