DesignCon 2017

Understanding Vertical Resolution in Oscilloscopes

Peter J. Pupalaikis, Teledyne LeCroy
PeterP@LeCroy.com, +1 (845) 425-2000

Abstract

This paper discusses analog-to-digital converter (ADC) resolution (mostly with regard to oscilloscopes) and considerations for improving resolution using various ADC deployment methods.

Patent Disclosure

Portions of this document are the subject of patents applied for.

Biography

Peter J. Pupalaikis was born in Boston, Massachusetts in 1964 and received the B.S. degree in electrical engineering from Rutgers University, New Brunswick, New Jersey in 1988.

He joined LeCroy Corporation (now Teledyne LeCroy), a manufacturer of high-performance measurement equipment located in Chestnut Ridge, New York in 1995 where he is currently Vice President, Technology Development, managing digital signal processing development and intellectual property. His interests include digital signal processing, applied mathematics, signal integrity and RF/microwave systems. Prior to LeCroy he served in the United States Army and has worked as an independent consultant in embedded systems design.

Mr. Pupalaikis holds forty-two patents in the area of measurement instrument design and has contributed a chapter to one book on RF/microwave measurement techniques. In 2013 he became an IEEE fellow for contributions to high-speed waveform digitizing instruments.

He is a member of Tau Beta Pi, Eta Kappa Nu and the IEEE signal processing, instrumentation, and microwave societies and serves on the DesignCon TPC and the industry advisory boards for ECEDHA and the Rutgers ECE department.

Introduction

THIS PAPER ADDRESSES RESOLUTION in oscilloscope channels, specifically the metric of effective-number-of-bits or effective number of bits (ENOB). We will step through the math involved in ENOB calculations from a noise standpoint. We will see that noise can be improved by considering correlation between noise in various signal paths and by considering spectral content of noise. A complete discussion of enhanced resolution filtering is provided. Finally, ADC deployment strategies will be discussed for trading sample rate for resolution.

Sampling and Quantization

Old style oscilloscopes - the analog kind - operate by sweeping an electron beam across the face of a cathode-ray tube that is coated with phosphor. The sweep is initiated by a trigger event and the beam sweeps ideally linearly in time across the screen and moves vertically in response to the waveform input voltage. This type of oscilloscope produces an essentially continuous waveform, albeit one that could only be viewed. Today, these oscilloscopes have been replaced with digital storage oscilloscopes (DSOs) that produce digital waveforms. The digital waveform produced by a DSO is discrete in both time and amplitude. The discrete time effect is quite different from the discrete amplitude effect.

The discrete time effect is caused by sampling the input signal. Sampling ideally takes a snapshot in time of a waveform at times that are ideally separated by the sample period. This sampling does not in itself discretize the voltage of a signal and in fact samplers produce impulses whose height or area contain the analog voltage. Track-and-holds are related to samplers in that they produce fully analog, steppy waveforms where the step changes begin at these sample period boundaries and the final, settled value contains the analog voltage.

Track-and-hold elements are very important in digital oscilloscopes because these elements hold the signal so that it can be quantized. Usually, this quantization of the signal occurs by determining a digital voltage value held on a capacitor. There are many ways of quantizing the voltage value including flash, successive approximation, and pipeline conversion. All of these quantization methods end with a number called a code that gets stored in memory.

Thus, sampling followed by quantization allows a DSO to capture and store waveforms as a sequence of discrete numbers.

After bandwidth, sample rate is the primary feature of the oscilloscope. We all know that the sample rate needs to be more than twice the highest frequency content possible in the waveform so that no aliasing occurs. While the bandwidth point of the oscilloscope is usually close to, but not at the highest frequency content possible, we generally need the sample rate to be at least three times the bandwidth to meet the Nyquist criteria. That being said, in order to make good use of the waveform from a measurement perspective, we generally need waveforms with sample rates that are at least ten times the bandwidth. We'll talk more about the disconnect between the roughly $3\times$ and $10\times$ sample rate to bandwidth requirement later, but for now it's important to understand that while sampling causes the resulting waveform to be discrete time, there is ideally no error in this discrete time. Furthermore, if Nyquist's criteria is met from a sampling standpoint, a waveform so sampled contains ALL of the information contained in the continuous analog waveform as explained in much literature [1].

After bandwidth and sample rate, the most important features of the oscilloscope are its ability to accurately capture a signal. One of the key aspects of accuracy is the main subject of this paper - the

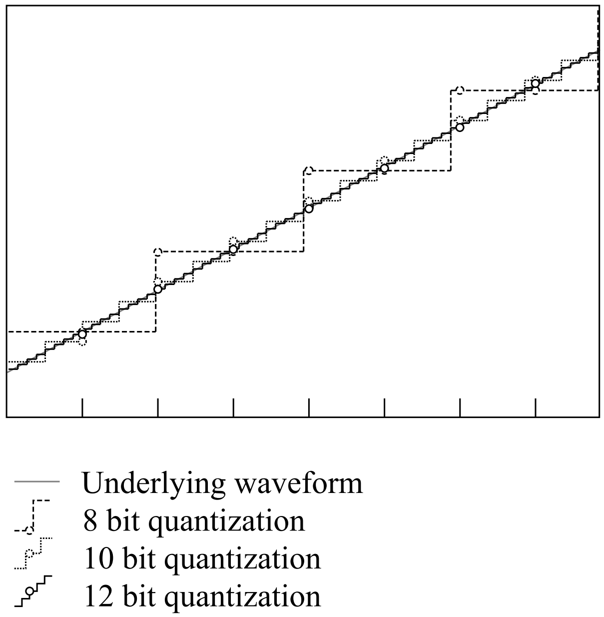

Figure 1: Quantization and Sampling

resolution of the oscilloscope. Resolution is the quantification of the vertical quantization of the waveform. Quantization comes about by converting the analog voltage sample into an integer number. This integer number is formed by comparing the analog voltage to discrete voltage levels. These voltage levels are fixed and there is a fixed number of these levels across the vertical range of the oscilloscope. Usually, the number of levels is a power of two, and the power of two is referred to as the number of bits in the converter. Each level produces a code, and thus the number of codes possible in a given vertical range of the oscilloscope is:

$$codes = 2^B$$

where B is the number of bits.

Since there is no restriction on the analog waveform, we know that a code at a sample point must be an approximation of the analog waveform sample because the sample code, under the most ideal conditions, could be in error from the analog waveform by $\pm\frac{1}{2}$ code.

Consider Figure 1 where we see an analog waveform being sampled by three different quantizers with eight, ten, and twelve bits. This figure is extremely zoomed in on a noise-free, ideal waveform. Here we can see the error created by sampling this waveform with different resolution. At this extreme zoom setting, the eight bit converter produces relatively large errors that diminish to a relatively small error with the twelve bit converter. It is important to note that this is an ideal, noise-free waveform and that it is zoomed way in to show the error. As we will discuss later, the quantization error produced is relatively small, even in an eight bit converter.

Thus, when talking about the characteristics of a digital oscilloscope, at the highest level, we talk about the bandwidth, the sample rate, and the number of bits of resolution. For a given bandwidth need, we want the hardware to sample at a *high enough* sample rate and want waveforms produced with sample rates that are about ten times the bandwidth and as many bits of resolution as possible.

Quantization Error

The resolution limitation of an oscilloscope waveform leads to small errors. There is an error on every sample taken. To quantify this error, consider the somewhat confusing equation (1) that describes the sampling process:

Mathematically speaking, if we have an analog waveform $v(t)$ and acquire K samples at a sample rate Fs (implying a sample period $Ts = 1/Fs$) with a B bit quantizer at a given voltage per division (VDIV) setting, we have a vector **x** of waveform samples, for $k \in 0 \ldots K-1$:

$$x[k] = \left\lfloor (v(k \cdot Ts) + 4 \cdot VDIV + Offset) \cdot \frac{2^B}{8 \cdot VDIV} \right\rfloor \cdot \frac{8 \cdot VDIV}{2^B} - 4 \cdot VDIV - Offset + \frac{1}{2^B} \quad (1)$$

Let's step through (1) to understand it:

- The sample rate is Fs and the sample period is $Ts = 1/Fs$.
- $k \cdot Ts$ is the time that the voltage waveform is sampled relative to the trigger time.
- VDIV is the number of Volts per division vertically on the oscilloscope screen and is how the gain setting of the oscilloscope is referred to by oscilloscope users. There is an implicit assumption here that there are eight vertical divisions across the screen (this may be different on different scopes).
- 2^B is the number of codes across the screen, where B is the resolution of the scope.
- $8 \cdot VDIV$ is the number of volts from the top to the bottom of the screen.
- $4 \cdot VDIV$ is added to the sampled analog voltage because, with no offset applied, the middle of the screen vertically is zero Volts.
- $Offset$ is the offset applied to the analog waveform, in Volts.
- $8 \cdot VDIV/2^B$ is the number volts per one code of the quantizer ($VoltsPerCode$).
- $\lfloor x \rfloor$ means the floor of x which means the next lower integer that is less than or equal to x.
- We are assuming in (1) that the waveform fits on the screen. When the waveform goes beyond the screen boundaries of $\pm 4 \cdot VDIV$, the value inside the floor function must be clipped to the minimum code (0) or maximum code ($2^B - 1$).
- The final addition of $1/2^B$ is not intuitive and is added so that when a full-scale sinusoid is applied (i.e. one with an amplitude of $4 \cdot VDIV$ or a peak-peak amplitude of $8 \cdot VDIV$), that the mean quantization error is zero.
- The result $x[k]$ is in quantized Volts (the floor function produced the actual integer codes).

Ideally, if the waveform were only sampled, and not quantized, we would have $x[k] = v(k \cdot Ts)$, but because it is quantized, we have instead:

$$x[k] = v(k \cdot Ts) + \varepsilon[k]$$

or said differently, a quantization error of:

$$\varepsilon[k] = x[k] - v(k \cdot Ts)$$

As such, we have an error vector that when subtracted from the waveform produces the exact voltage. Said differently, we can think of the acquired waveform as exact samples of the analog waveform with a noise waveform added to the analog waveform. This separation of noise and actual waveform is important in the determination of the quality of the acquisition channel and will be covered in the next section.

ENOB

When describing the signal-to-noise-ratio (SNR) of a system, we compare the root-mean-square (rms) value of both the signal and noise, with units of Volts (rms). The rms voltage of a waveform is often referred to as the *effective* voltage. The rms value of a vector is expressed as:

$$V_{rms}^2 = \frac{1}{K} \sum_k x[k]^2$$

This is the same as the standard deviation of the signal with the mean value of the signal (the zero frequency (DC) component removed).

The mean or DC component of a signal is:

$$\mu = \frac{1}{K} \cdot \sum_{k=0}^{K-1} x[k]$$

and the standard deviation is:

$$\sigma^2 = \frac{1}{K} \sum_{k=0}^{K-1} (x[k] - \mu)^2$$

We can see that if $\mu = 0$, the standard deviation equals the rms value. It is customary, when measuring SNR to compare signals not including the DC component of either the signal or the noise and to attribute these errors elsewhere.

Returning to the quantization error $\varepsilon[k]$, we know that, statistically speaking, we have an error that is uniformly distributed within $\pm 1/2$ code.

Therefore, the quantization noise can be expressed as[1]:

$$\sigma^2 = \int_{-\frac{1}{2}}^{+\frac{1}{2}} x^2 \cdot dx \tag{2}$$

and therefore:

[1] The rms value is calculated by integrating the PDF multiplied by x^2 over the entire range of possibilities:

$$\sigma^2 = \int_{-\infty}^{\infty} \text{PDF}(x) \cdot x^2 \cdot dx$$

For example, for a normal, zero mean distribution, we have:

$$\sigma = \frac{1}{2 \cdot \sqrt{3}} \cdot VoltsPerCode \tag{3}$$

Thus, the rms noise in an oscilloscope due to quantization according to (1) can be expressed as:

$$\sigma = \frac{1}{2 \cdot \sqrt{3}} \cdot \frac{8 \cdot VDIV}{2^B} \tag{4}$$

which we usually express in decibels (dB) relative to 50 mW into 50 Ohms (dBm) as:

$$\begin{aligned}\sigma_{dBm} &= 20 \cdot \log(\sigma) + 13.010 = \ldots \\ \ldots &= 20 \cdot \log\left(\frac{1}{2 \cdot \sqrt{3}} \cdot \frac{8 \cdot VDIV}{2^B}\right) + 13.010 = \ldots \\ \ldots &= 20 \cdot \log(VDIV) + 10 \cdot \log\left(\frac{16}{3}\right) - 20 \cdot B \cdot \log(2) + 13.010\end{aligned}$$

The rms value of a full-scale sinewave is:

$$V_s = \frac{4 \cdot VDIV}{\sqrt{2}}$$

This can also be expressed in dBm as:

$$\begin{aligned}V_{dBm} &= 20 \cdot \log(V_s) + 13.010 = \ldots \\ \ldots &= 20 \cdot \log\left(\frac{4 \cdot VDIV}{\sqrt{2}}\right) + 13.010 = \ldots \\ \ldots &= 20 \cdot \log(VDIV) + 10 \cdot \log(8) + 13.010\end{aligned}$$

The SNR is calculated by subtracting the noise from the signal as:

$$\begin{aligned}\sigma^2 &= \int_{-\infty}^{\infty} \frac{1}{\sqrt{2 \cdot \sigma^2 \cdot \pi}} \cdot e^{-\frac{x^2}{2 \cdot \sigma^2}} \cdot x^2 \cdot dx = \ldots \\ \ldots &= \frac{\sqrt{2} \cdot \sigma \cdot \mathrm{csgn}(\sigma)}{4 \cdot \sqrt{\pi}} \cdot \left[\sigma \cdot \sqrt{2 \cdot \pi} \cdot \mathrm{erf}\left(\frac{1}{2} \cdot \frac{\sqrt{2}}{\sigma} \cdot x\right) - 2 \cdot x \cdot e^{-\frac{x^2}{2 \cdot \sigma^2}}\right]\Bigg|_{-\infty}^{\infty} = \ldots \\ \ldots &= \frac{\sqrt{2} \cdot \sigma}{4 \cdot \sqrt{\pi}} \cdot \left[2 \cdot \sigma \cdot \sqrt{2 \cdot \pi}\right] = \ldots \\ \ldots &= \sigma^2\end{aligned}$$

$$SNR = V_{dBm} - \sigma_{dBm} = \dots$$

$$\dots = 20 \cdot \log(VDIV) + 10 \cdot \log(8) + 13.010 - \dots$$

$$\dots - \left[20 \cdot \log(VDIV) + 10 \cdot \log\left(\frac{16}{3}\right) - 20 \cdot B \cdot \log(2) + 13.010\right] = \dots$$

$$\dots = 10 \cdot \log\left(\frac{3}{2}\right) + 20 \cdot B \cdot \log(2) \quad (5)$$

Note that (5) is a function only of the number of bits in this scenario. In other words, given an ADC resolution B, we obtain an SNR as given in (5). In practice, the SNR is a function of many other things, and for a given resolution, we can say that (5) is the absolute best case SNR possible; it is the SNR measured in a system with only quantization.

Thus, for an eight, ten, and twelve bit scope, the absolute best SNR possible is 49.92, 61.96, and 74 dB, respectively, but as we will see, these absolute limits apply to digitizers whose useful frequency content spans the entire Nyquist band.

This is why we said previously that the discretization of time has ideally no impact on the measurement of the continuous, analog signal, but quantization adds an imperfection in the form of noise added to the measurement.

The equation in (5), when solved for B provides a measurement of the effective resolution called ENOB [2]:

$$ENOB = \frac{SNR - 10 \cdot \log\left(\frac{3}{2}\right)}{20 \cdot \log(2)} \approx \frac{SNR - 1.76}{6.02} \quad (6)$$

Thus, given a measurement of SNR, we can convert this to a system figure of merit measured entirely as effective resolution.

Generally speaking, distortion components are also considered in measurements of ENOB, so strictly speaking, (6) can be thought of as the ENOB due to noise effects only. If there are no other sources of noise in the system, ENOB will evaluate to the number of bits in the quantizer or converter.

In the design of high resolution oscilloscopes, one goal is to have the highest ENOB, which will always be limited at some point by the number of bits of resolution. In subsequent sections, we will talk about other sources of noise in the system that lower the SNR and therefore the ENOB.

Overstatement of ENOB

On a slight diversion, strictly speaking, (2) (and therefore (3)) is only truly valid for converting the resolution of the quantizer to the rms noise expected when the distribution of noise is completely uniform. Therefore, (6) is true only when the error is absolutely uniformly distributed as implied by (3) and this would be true if a ramp spanning all of the codes were applied to the system. If pure DC was applied to the system and there are no other sources of noise, this would clearly not be true - the error would be constant.

When measuring effective bits, we typically measure with respect to frequency. This is because distortion components are typically a function of frequency. But SNR can also be a function of frequency when jitter is involved. Jitter tends to raise the noise floor at higher frequencies. In any case, SNR is usually measured with a full-scale, or near full-scale sinusoid applied [3]. The equation in (3) assumes a uniformly

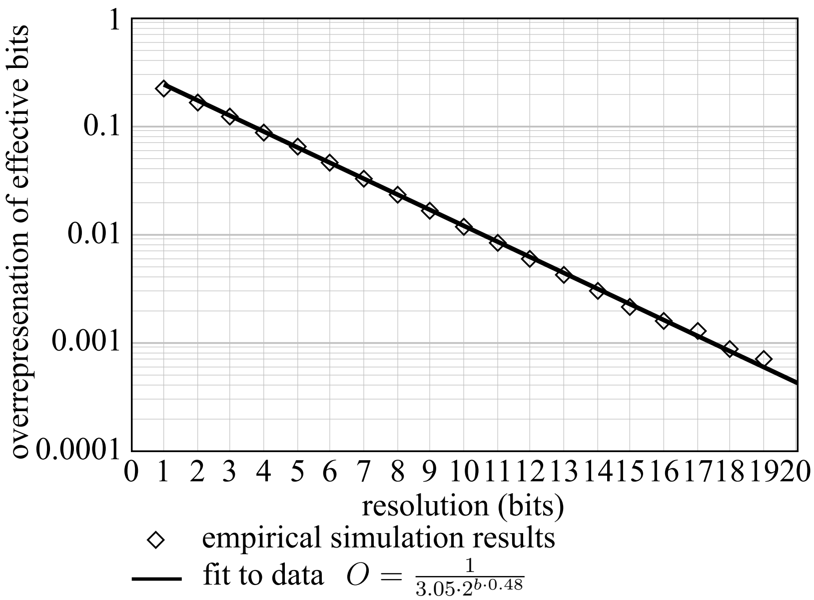

Figure 2: Overstatement of Effective Bits with Applied Sinusoids

distributed error. It can be shown through simulation that a sinusoid mostly contains uniformly distributed quantization noise. It can be further shown that because of the slight non-uniformity, (6) tends to slightly overstate the effective bits for applied sinusoids, but this overstatement diminishes significantly with number of bits. In Figure 2 we see this overstatement. Here we show through empirical measurements that a function can be fitted that is log-linear that expresses this overstatement as:

$$O = \frac{1}{3.05 \cdot 2^{b \cdot 0.48}}$$

For quantizers in excess of five bits, this error is less than 0.05 bits and for quantizers in excess of seven bits, the error is below 0.025 bits.

Sources of Noise in Oscilloscopes

The effective resolution shown in (6) is limited by the resolution of the ADC but quantization error is not the only source of noise in the oscilloscope, and usually is not even the dominant noise source. The dominant source of noise in the oscilloscope tends to be the front-end amplifier.

Figure 3 shows an example of noise sources in an oscilloscope channel. Here we have a user input signal V_{in}, itself with noise added to it entering the oscilloscope. The front-end amplifier adds noise to the signal. In this particular design, two front-end outputs drive two ADC inputs which themselves add their own noise. Each ADC input drives four, internal ADCs which add noise, mostly in the form of quantization noise due to limited resolution.

It is important to realize that the often neglected noise source: the noise on the user's signal, if present, cannot be removed. This is because the oscilloscope does not know that the noise is even noise. As far as it is concerned, the noise is signal and its job is to faithfully reproduce the signal. This is an important point

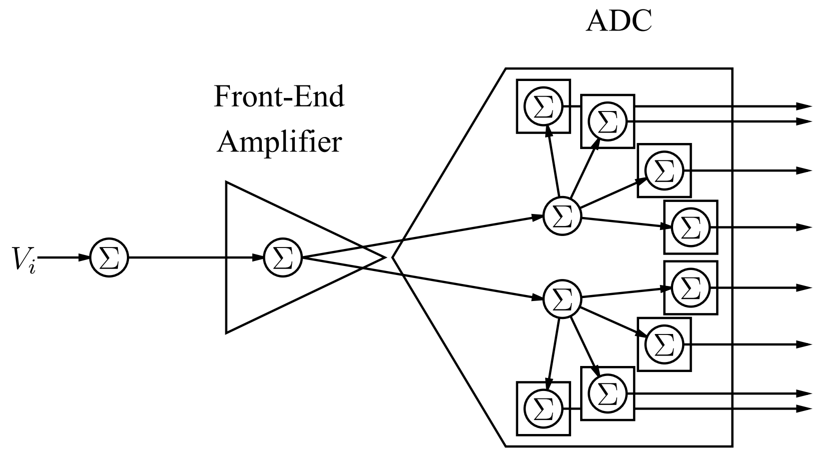

Figure 3: Example Noise Sources in an Oscilloscope Channel

to remember when measuring ENOB and noise and we must remember to use a high fidelity low-noise source and sometimes filters, otherwise we might be measuring the noise in the source and not the channel.

The front-end amplifier adds the most noise. Usually, it consists of multiple, selectable gain stages, so often, there is an implementation strategy that can provide the least noise relative to signal size. It's important to understand that the noise added to the signal in the front-end amplifier causes a problem in that it is indistinguishable also from the noise on the user's input signal and is common to all downstream paths as shown.

The job of the ADC is, as with all elements, to faithfully digitize the waveform presented to it, but the fact that the two ADC inputs add different noise is an opportunity which we will discuss. Finally, each internal ADC adds its quantization noise and perhaps other types of noise.

The important parameters of each noise source in the system are:

- Obviously, the magnitude of the noise.
- The correlation between a source and other sources.
- The location of the noise source in the signal path and how common the noise source is to the other paths through the system.
- The spectral characteristics of the noise.

We will talk subsequently about how knowledge of the noise sources can be exploited to improve effective resolution.

Correlation of Noise

Correlation or noise sources is a term in statistics that defines how *related* one noise source is to another. Correlation of noise sources can be an advantage or disadvantage for noise removal through various processing techniques. A possible advantage is presented if, in a given stream of waveform data, a given sample is unrelated or uncorrelated with other samples.

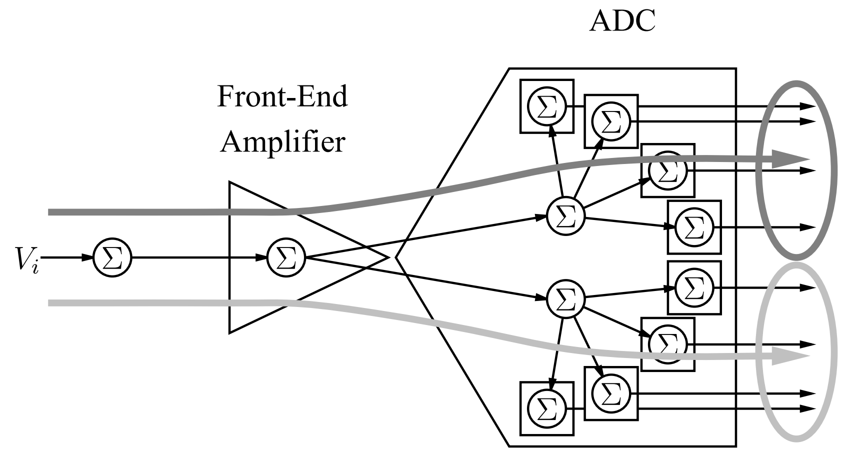

Figure 4: Multiple Paths in an Oscilloscope Channel

A more common advantage is gained through lack of correlation of noise in multiple paths through a system. For example, in Figure 3, regarding the multiple front-end amplifier connections to the digitizer, we can separate the signal into at least two system paths as shown in Figure 4. Here, we have one path shown in gray and another path shown in silver. In an ordinary arrangement, the internal ADCs to the digitizer are time interleaved, meaning that every other sample comes from the gray path and the silver path. Since the two paths have in common the noise in the user's signal and the noise from the front-end, this noise must be correlated in each of the paths - it is the same noise. But the noise in the two separate front-end outputs/ADC inputs may not be correlated, as well as the noise from the individual digitizers. To the extent that portion of the noise in the separate portions of the paths are uncorrelated, an advantage is presented that could be exploited.

Assuming, for example, that the internal ADCs are sampling at 5 GS/s, we have two separate 20 GS/s streams in the gray and silver paths. If the input signal content is above 10 GHz, the two streams could be arranged to sample the signal simultaneously and their average value would have less noise than an individual 20 GS/s stream to the extent that the noise in the two paths is uncorrelated. In this case, however, the resulting averaged 20 GS/s stream would be undersampling the signal which is generally undesirable. If the input signal content is below 10 GHz, then simultaneous sampling and averaging of the two streams would result in a sufficiently sampled 20 GS/s stream. Not only could this resulting stream be resampled back to 40 GS/s, providing the oversampling benefit provided in a time-interleaved system, each of the 20 GS/s streams could have been time interleaved originally and provide the same benefit. This is because one of the two streams could be resampled to the same sample phase as the other, averaged, and upsampled.

The benefit of this averaging would depend on two things:

- The magnitude of the noise in the separate portions of the path relative to the common portions of the path.
- The degree of the lack of correlation between the noise in the separate portions of the path.

The quantification of the benefit can be described, in a limiting sense, as follows:

1. If the all of the noise comes in the common portion of the path, the benefit is zero.

2. If the noise in the separate portions of the path are totally correlated, the benefit is zero.

3. If all of the noise comes in the separate portions of the path and is totally uncorrelated, the benefit of averaging the two streams is the equivalent of half a bit of resolution.

Thus, the noise benefit due to averaging two streams can be bounded and is between zero and one-half bit of resolution. The zero bound is clear, averaging two streams that are the same results in no change. The half-bit resolution improvement is described statistically:

Given two random variables with the same mean (signal content) and normally distributed, uncorrelated noise, we can write these as:

$$X \sim N\left(\mu, \sigma_X^2\right)$$

$$Y \sim N\left(\mu, \sigma_Y^2\right)$$

Where μ is the mean and σ is the standard deviation. The average of these two streams can therefore be written as:

$$Z = \frac{1}{2} \cdot (X + Y) \sim N\left(\mu, \frac{1}{2} \cdot \left(\sigma_X^2 + \sigma_Y^2\right)\right)$$

If the magnitude of σ_X and σ_Y is the same, then the resulting standard deviation (and therefore rms value) is $\sigma/\sqrt{2}$. Looking back at (5), we see that this means that the SNR improves by $20 \cdot \log\left(\sqrt{2}\right) \approx 3\,dB$ and therefore the ENOB (due to noise only) according to (6) increases by 0.5 bits.

Spectral Content of Noise

Spectral content of noise is important to understand because this is often an area for improvement. If we have a spectral density expressed as $R\left(f\right)$ in units of $V\ \left(rms\right)/\sqrt{Hz}$, the total noise in the system is:

$$\sigma^2 = \int_0^{BW} R\left(f\right)^2 \cdot df$$

If the noise is white (i.e. evenly distributed throughout the bandwidth BW), such that $R\left(f\right) = \overline{R}$, then we have:

$$\sigma^2 = \int_0^{BW} \overline{R}^2 \cdot df = \overline{R}^2 \cdot BW$$

and therefore:

$$\overline{R} = \frac{\sigma}{\sqrt{BW}}$$

If we were to halve the bandwidth in such a system, by employing a hard cutoff at $BW/2$, we have

$$\sigma_{bwl}^2 = \int_0^{\frac{BW}{2}} \overline{R}^2 \cdot df = \overline{R}^2 \cdot \frac{BW}{2} = \frac{\sigma^2}{BW} \cdot \frac{BW}{2} = \frac{\sigma^2}{2}$$

and therefore:

$$\sigma_{bwl} = \frac{\sigma}{\sqrt{2}}$$

Looking back at (5), we see that this means that the SNR improves by $20 \cdot \log\left(\sqrt{2}\right) \approx 3\,dB$ and therefore the ENOB (due to noise only) according to (6) increases by 0.5 bits.

This is the precept behind enhanced resolution (ERES) filtering, although the typical ERES filter applied tends not to have the brick-wall shape in this example. The principle, however, is exactly the same. A filter is crafted that causes the bits to increase by 0.5 bits for every halving of bandwidth.

Remember that we understood previously that the absolute high limit on ENOB for a B bit digitizer is only true if the usable signal content spans the entire Nyquist band. If the sample rate of the system is higher then two-three times the bandwidth, then resolution can be improved by simply filtering the waveform digitally after acquisition. This means that simply sampling at a higher rate and filtering can be used to improve resolution. This might seem counterintuitive, but think of it this way - if we sample at a high rate and the noise is only due to quantization, then we could average two adjacent samples that are presumed to be uncorrelated to obtain a half bit improvement. This would of course come at the expense of lowering the bandwidth, but only lowering it half the Nyquist rate. But if our useful signal content was only contained in the first half of the Nyquist band, the improvement comes essentially for free. We will talk more about this in the next section.

ERES Filtering

A digital filter response is expressed as $H\left(z\right)$ where . Given a noise spectral density $R\left(z\right)$, the effect of the filter on the noise spectrum is defined as:

$$\sigma_{df}^2 = \frac{1}{2 \cdot j} \oint_C \frac{\left[\left|H\left(z\right)\right| \cdot R(z)\right]^2}{z} \cdot dz$$

where the contour of integration is along the rim of the unit circle.

This can be simplified by substituting for z where[2]:

$$z = e^{j \cdot \omega}$$

$$dz = j \cdot z \cdot d\omega$$

$$\sigma_{df}^2 = \frac{1}{2} \cdot \int_0^{2 \cdot \pi} \left[\left|H\left(z = e^{j \cdot \omega}\right)\right| \cdot R\left(\omega\right)\right]^2 \cdot d\omega$$

Assuming evenly spread noise from DC to the Nyquist rate, we have originally[3]:

[2] We have chosen the sample rate arbitrarily as unity.

[3] The reader might be expecting $\sigma/\sqrt{2 \cdot \pi}$ here, but the duplicated noise in the negative spectrum has already been factored in and it's appropriate to state the density of the noise in the positive spectrum only.

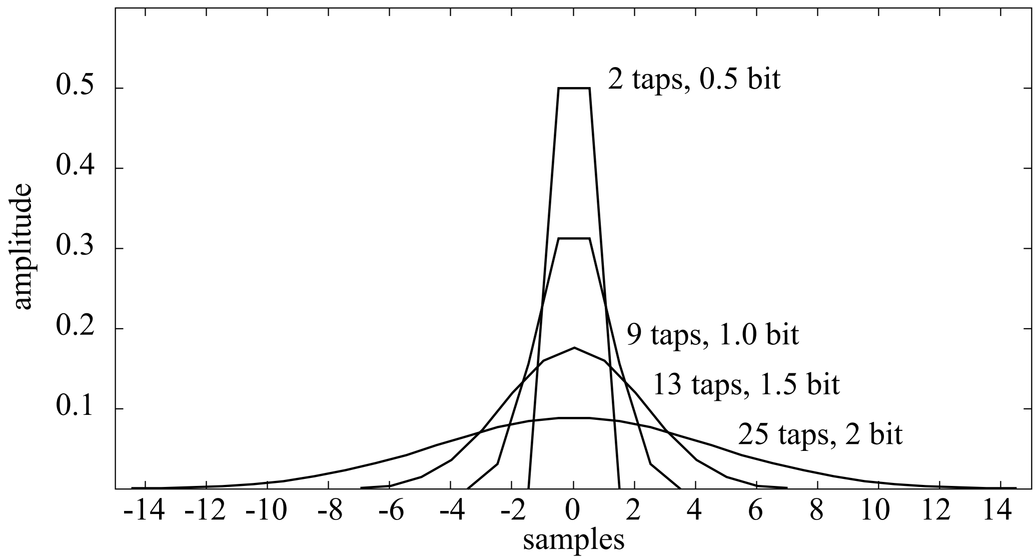

Figure 5: ERES Impulse Responses

$$R(\omega) = \overline{R} = \frac{\sigma}{\sqrt{\pi}}$$

and we have:

$$\sigma_{df}^2 = \sigma^2 \cdot \frac{1}{2 \cdot \pi} \cdot \int_0^{2 \cdot \pi} \left| H\left(z = e^{j \cdot \omega}\right) \right|^2 \cdot d\omega$$

ERES filtering is typically performed with a Gaussian filter [4]. A Gaussian filter is a filter whose impulse response is a Gaussian which provides an ideal pulse response in that there is no overshoot. The central limit theorem states that convolving many rectangular, or boxcar filters produces a response that tends towards a Gaussian shape and that is how ERES filters are created. With this in mind, the simplest ERES filter is the two-tap boxcar filter with taps $\left[\frac{1}{2}, \frac{1}{2}\right]$. As such, this filter simply averages two points. The response of this filter is:

$$H(z) = \frac{1}{2} + \frac{1}{2} \cdot z^{-1}$$

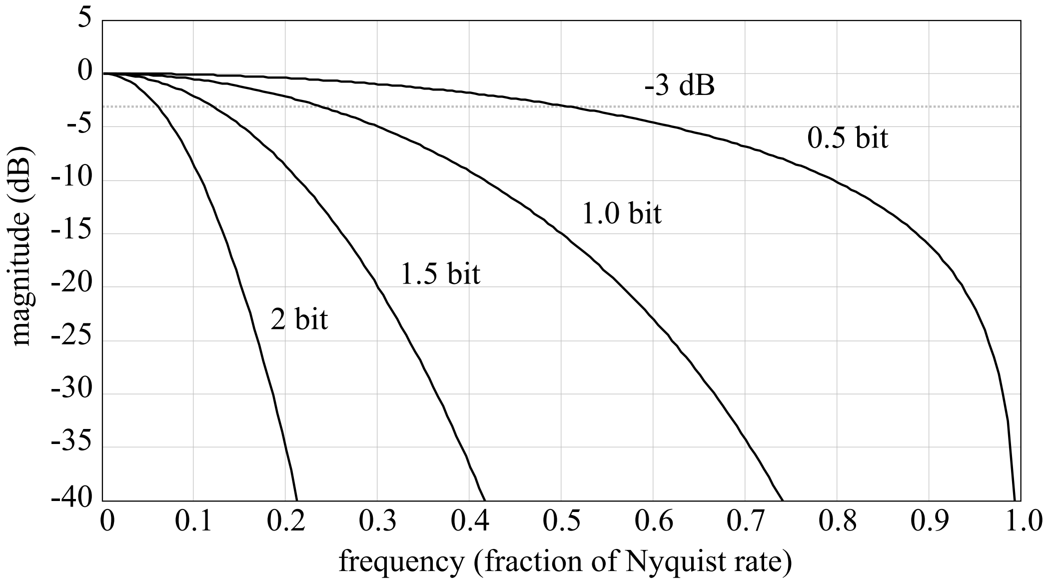

Figure 6: ERES Frequency Responses

$$\sigma_{eres}^2 = \sigma^2 \cdot \frac{1}{2 \cdot \pi} \cdot \int_0^{2 \cdot \pi} \left| \frac{1}{2} + \frac{1}{2} \cdot e^{-j \cdot \omega} \right|^2 \cdot d\omega = \ldots$$

$$\ldots = \sigma^2 \cdot \frac{1}{2 \cdot \pi} \cdot \int_0^{2 \cdot \pi} \cos^2 \left(\frac{\omega}{2} \right) \cdot d\omega = \ldots$$

$$\ldots = \sigma^2 \cdot \left[\frac{\omega}{4 \cdot \pi} + \frac{1}{2 \cdot \pi} \cdot \sin \left(\frac{\omega}{2} \right) \cdot \cos \left(\frac{\omega}{2} \right) \right] \Bigg|_0^{2 \cdot \pi} = \ldots$$

$$\ldots = \sigma^2 \cdot \frac{1}{2}$$

and thus:

$$\sigma_{eres} = \frac{\sigma}{\sqrt{2}}$$

This is the so-called 0.5 bit ERES filter. This result should not be surprising because we are effectively averaging two adjacent points and since the noise in the two adjacent sample points are assumed to be not correlated, we get the half bit improvement foretold in the averaging discussion[4].

To improve resolution further, we can cascade many stages S of these two-tap averaging filters, and we find that the effective bits improvement is:

[4]this half bit improvement is only possible if the noise bandwidth extends to the Nyquist rate as we've stipulated. If the noise bandwidth did not extend to the Nyquist rate, the spectral density of the noise could not be uniform and we could not get the full half bit improvement.

$$B=-\frac{1}{2}\cdot\log_2\left(\frac{1}{\pi}\cdot\int\limits_0^{\pi}\cos^{2\cdot S}\left(\frac{\omega}{2}\right)\cdot d\omega\right)$$

To solve this, we make use of the identity[5]:
$\int\cos^n\left(a\cdot x\right)\cdot dx$

$$\int\cos^n\left(a\cdot x\right)\cdot dx=\frac{1}{n\cdot a}\cdot\cos^{n-1}\left(a\cdot x\right)\cdot\sin\left(a\cdot x\right)+\frac{n-1}{n}\int\cos^{n-2}\left(a\cdot x\right)\cdot dx$$

and thus:

$$\frac{1}{2\cdot\pi}\cdot\int\limits_0^{2\cdot\pi}\cos^{2\cdot S}\left(\frac{\omega}{2}\right)\cdot d\omega=\ldots$$

$$\ldots=\frac{2\cdot S-1}{2\cdot S}\cdot\frac{1}{2\cdot\pi}\cdot\int\limits_0^{2\cdot\pi}\cos^{2\cdot(S-1)}\left(\frac{\omega}{2}\right)\cdot d\omega=\ldots$$

$$\ldots=\prod_{s=1}^{S}\frac{2\cdot s-1}{2\cdot s}$$

So, for a given number of stages, the improvement in bits can be written as:

$$B=-\frac{1}{2}\cdot\log_2\left(\prod_{s=1}^{S}\frac{2\cdot s-1}{2\cdot s}\right)$$

The 3 dB bandwidth as a fraction of the Nyquist rate of such a number of cascaded sections is found by solving:

$$\left|\left(\frac{1}{2}+\frac{1}{2}\cdot e^{-j\cdot 2\cdot\pi\cdot f}\right)^S\right|=\frac{1}{\sqrt{2}}$$

to arrive at[6]:

$$f_{3\,dB}=\frac{1}{\pi}\cdot\tan^{-1}\left(\frac{\sqrt{2^{\frac{2\cdot S-1}{S}}\cdot\left(1-2^{-\frac{1}{S}}\right)}}{2^{\frac{S-1}{S}}-1}\right)$$

The bit improvement for some numbers of stages are shown in Table 1.

In Practice, we don't like to have fractional bits improvements, so instead of cascading averaging stages, we find the fractional value of S that produces the improvement in bits that we want and simply find the

[5]If you want to solve this yourself, approach $\int\cos^n\left(a\cdot x\right)\cdot dx$ using $u\cdot dv$ substitution with $u=\cos^{n-1}\left(a\cdot x\right)$ and $dv=\cos\left(a\cdot x\right)\cdot dx$.

[6]when $S=1$, you need to take the limit as $S\rightarrow 1$ and you arrive at $f_{3\,dB}=1/2$

stages	$\frac{1}{2\cdot\pi}\cdot\int_0^{2\cdot\pi}\cos^{2\cdot S}\left(\frac{\omega}{2}\right)\cdot d\omega$	bits improvement	bandwidth (fraction of Nyquist rate)
1	$\frac{1}{2}$	0.5	0.5
2	$\frac{3}{8}$	0.708	0.364
3	$\frac{5}{16}$	0.839	0.3
4	$\frac{35}{128}$	0.935	0.261
5	$\frac{63}{256}$	1.011	0.234
6	$\frac{231}{1024}$	1.074	0.214
7	$\frac{429}{2048}$	1.128	0.199

Table 1: ERES bits improvement as a function of stages

stages	$\frac{1}{2\cdot\pi}\cdot\int_0^{2\cdot\pi}\cos^{2\cdot S}\left(\frac{\omega}{2}\right)\cdot d\omega$	bits improvement	bandwidth (fraction of Nyquist rate)
1	$\frac{1}{2}$	0.5	0.5
4.84	$\frac{1}{4}$	1	0.238
20.1	$\frac{1}{8}$	1.5	0.118
81.2	$\frac{1}{16}$	2	0.059
326	$\frac{1}{32}$	2.5	0.029
1305	$\frac{1}{64}$	3	0.015
5220	$\frac{1}{128}$	3.5	0.007

Table 2: ERES bits improvements in 0.5 bit increments

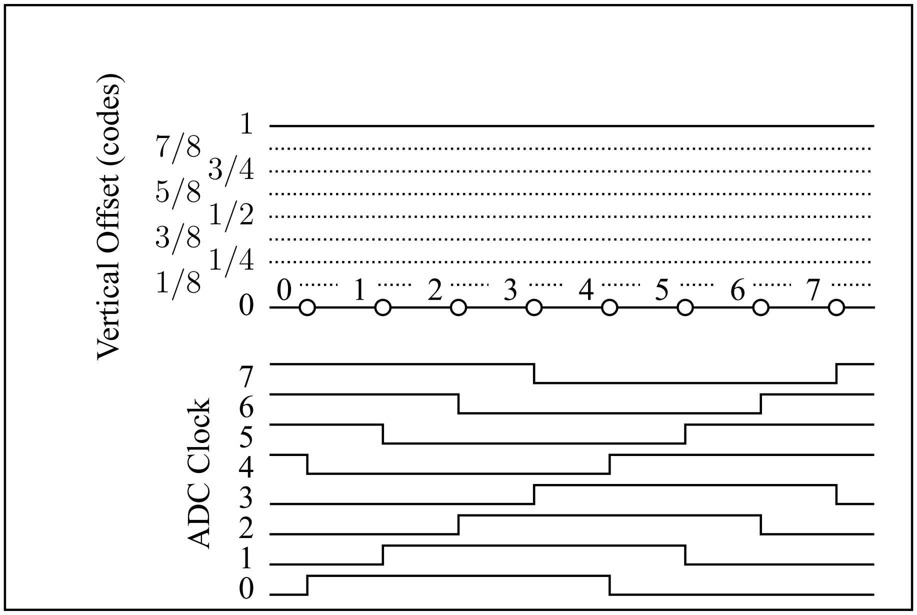

Figure 7: Standard 40 GS/s, Eight Bit Time Interleave

filter that fits the response. Thus, the bit improvements and the number of stages for these improvements is tabulated in Table 2. When the ERES filters are calculated, we find that we need very large numbers of stages and therefore very long filters. In practice most of the filter coefficients are insignificant as the large number of cascades better approximates a true Gaussian shape and we know that the tails of a Gaussian tend to be very small values. These filter responses are shown in the time domain in Figure 5 and in the frequency domain in Figure 6.

Time and Vertical Interleaving

This section will now discuss certain ADC deployment strategies for improving resolution. Most often, ADCs are deployed in a time-interleaved arrangement to improve sample rate and to address practical limits to speed of individual ADCs. In Figure 7 we see a typical situation of combination of ADCs for time interleaving. Here we have eight 5 GS/s ADCs arranged such that their 5 GHz sample clocks are delayed by a fraction of the 5 GHz sample phase. Each ADC is offset by the same amount. ADC 0 samples with a given sample clock, ADC 1 samples at one eighth of the period later, ADC 2 samples at one quarter the sample period later, and so on. After the data has been acquired, the data from each ADCs data is merged together in the proper time order forming an effective 40 GS/s acquisition.

In preparation for the subsequent discussion, in Figure 8 we show a conceptual arrangement of N ADCs in both time and offset. Here we have N^2 possible deployment locations for the N ADCs. Any of the N ADCs can be placed in any of the locations (even overlapping) offset by multiples of $1/N$ of an eight bit code vertically or the individual ADC sample period horizontally. The traditional time interleaved arrangement exemplified in Figure 7 is shown using this conceptual arrangement in Figure 9.

While Eight ADCs all arranged in a time interleaved arrangement is one extreme that provides for the maximum 40 GS/s sample rate, but restricted to the eight bit resolution of each ADC, Figure 10 shows an

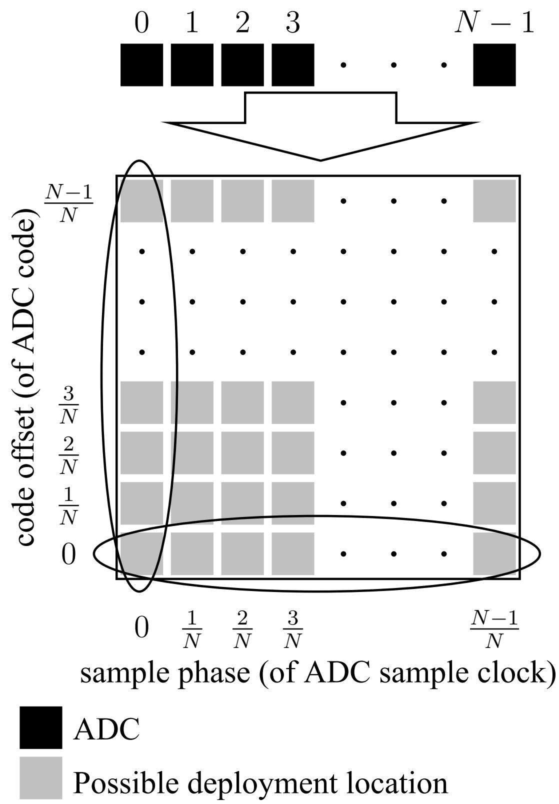

Figure 8: Conceptual ADC Deployment Locations

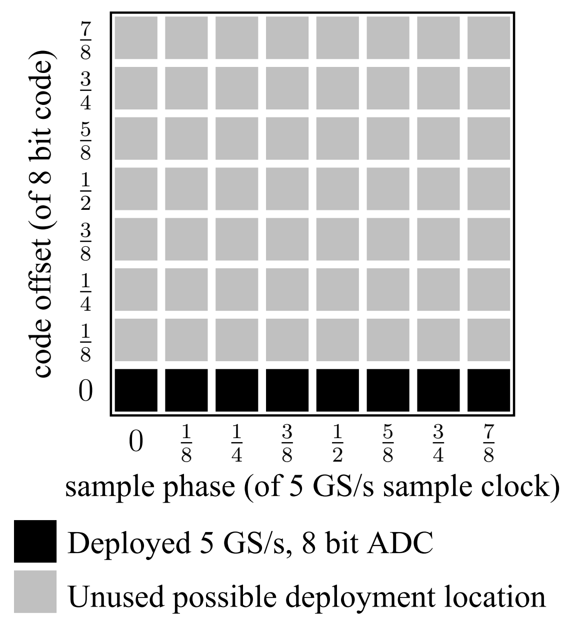

Figure 9: Conceptual 40 GS/s, Eight Bit Time Interleave

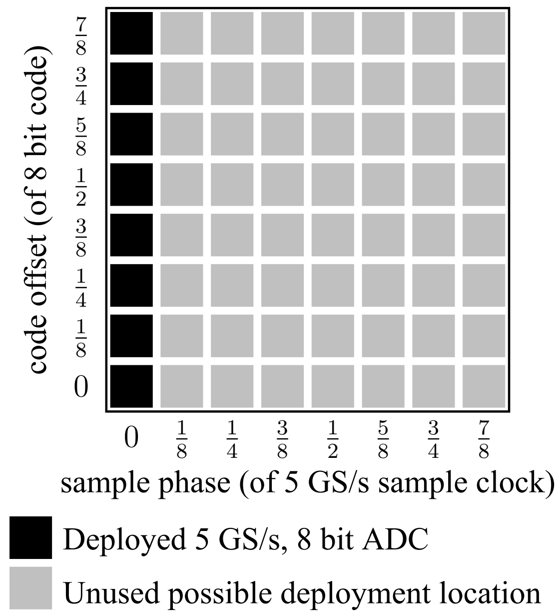

Figure 10: Conceptual 5 GS/s, Eleven Bit Vertical Interleave

opposite extreme. In this arrangement, all eight of the ADCs samples simultaneously, but each is offset by one eighth of a code. This produces a resulting 5 GS/s acquisition that can provide for eight levels in between each of the eight bit codes, thereby expanding the resolution to eleven bits albeit with a severe reduction of sample rate.

We've shown here that there is a trade-off that can be made in how ADC resources are deployed between sample rate and resolution. It is important to understand how this trade-off might be considered.

High resolution and high sample rate are both desired. What we really want is high effective resolution based on the previous ENOB discussion. The high sample rate is an absolute requirement, based on the bandwidth, but the absolute requirement is really on *hardware* sample rate. The hardware sample rate must satisfy the Nyquist rate limitation. Once the Nyquist rate limit has been satisfied, interpolation methods can be used to provide the $10\times$ oversampling desired by the user. Users want waveforms with $10\times$ over sampling because it can be shown that, in this situation, lines can be drawn between the samples and the actual sample points. Not only are lines drawn between points (i.e. linear interpolation) a very good approximation of the underlying analog waveform when the sample rate is at least ten-times the bandwidth, the eye interpolates linearly when viewing a waveform. Therefore, an ADC deployment strategy can be to deploy the ADCs in a manner to sufficiently sample the signal meeting the Nyquist criteria, interpolate to reach the $10\times$ oversampling requirement, and deploy the remaining ADCs in a manner to improve the ENOB. Deploying the ADCs only in the traditional time-interleaved arrangement to maximize sample rate could be considered as a wasteful extravagance when effective resolution could be improved. That being said, the best operating modes of the ADCs will depend on the characteristics of the channel.

An example of this determination is shown in Figure 11 for a 4 GHz analog bandwidth. Since the Nyquist requirement is for a sample rate of $2 \cdot 4GHz = 8GS/s$, we can probably sample at 10 GS/s[7].

[7] We say probably because the bandwidth of the channel is the -3 dB point, not the limit to signal content. But most oscilloscope responses tend to drop off rapidly after the bandwidth has been reached. While the bandwidth is a good indicator for frequency content possible, it is not the limit and the limit would need to be confirmed through measurements.

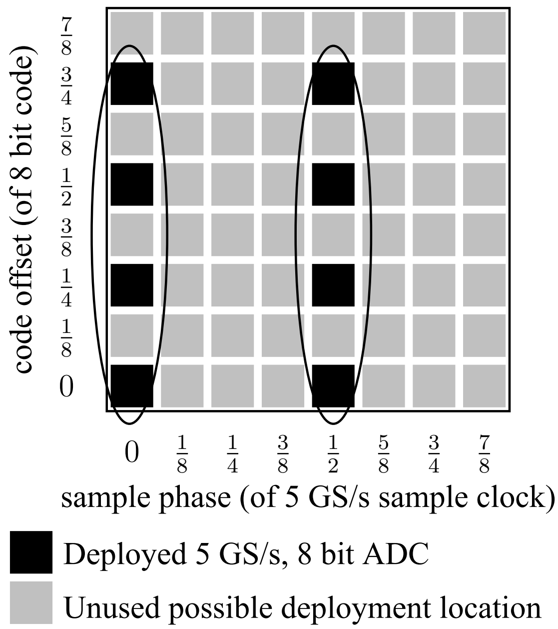

Figure 11: Conceptual 10 GS/s, Ten Bit Time and Vertical Interleave

Thus, having eight ADCs, we can deploy them in a mixture of time and vertical interleave arrangements as shown. Here, groups of four ADCs sample the waveform simultaneously with one group at phase zero and another group at one-half a sample period later. Within each group of four, each ADC is offset by a multiple of one quarter of a code producing four levels in between each eight bit code for ten bits of resolution. In this manner, the ADCs are deployed to produce a 10 GS/s waveform at ten bits of resolution. Since $10\times$ oversampling is desired, three points are interpolated in between each of the sample points acquired by the hardware producing the desired 40 GS/s, ten bit waveform.

Considering that in Figure 11, we have at phases 0 and $1/2$, four ADCs 0, 1, 2, and 3 forming a vertically interleaved group. The code corresponding to a given ADC and a given voltage follows the formula:

$$voltage = voltsPerCode \cdot code + offset\left[adc\right]$$

and therefore the code provided by a given ADC due to a given voltage and offset is:

$$code = \left\lfloor \frac{voltage - offset\left[adc\right]}{voltsPerCode} \right\rfloor = \ldots$$
$$\ldots = \left\lfloor \frac{voltage}{voltsPerCode} - \frac{offset\left[adc\right]}{voltsPerCode} \right\rfloor = \ldots$$
$$\ldots = \left\lfloor \frac{voltage}{voltsPerCode} + codeFraction\left[adc\right] \right\rfloor = \ldots$$
$$\ldots = \left\lfloor \frac{voltage}{voltsPerCode} + \frac{adc}{vilv} \right\rfloor$$

where here $vilv$ is the number of vertically interleaved ADCs and $adc \in 0 \ldots vilv - 1$ is an adc number. Here, $vilv = 4$.

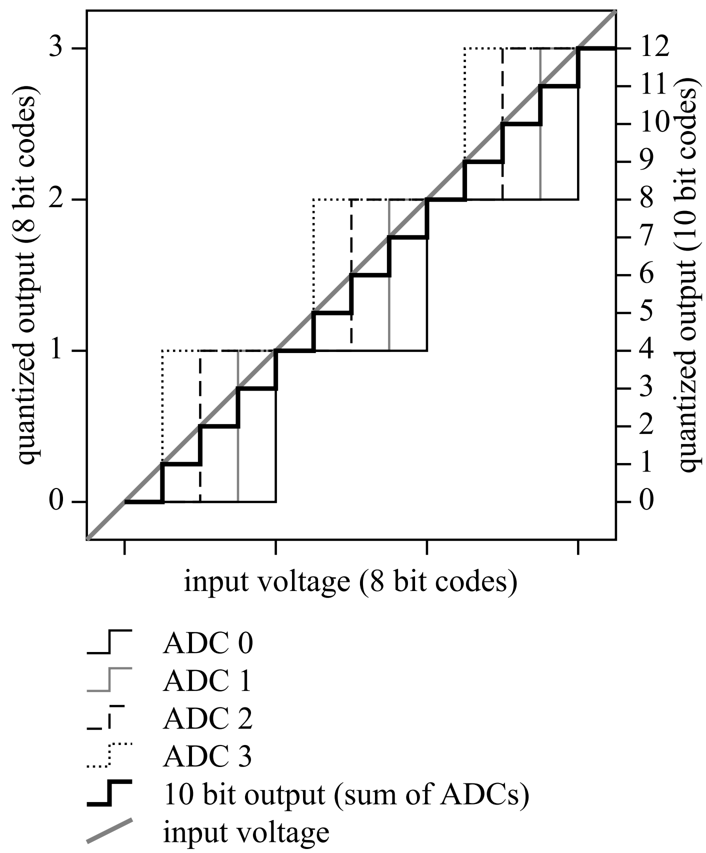

Figure 12: Transfer Characteristics for Four Way Vertical Interleave

Let's assume for simplicity that $voltsPerCode = 1$ (i.e. we can talk about voltages applied in terms of eight bit codes). for ADC 0, with an offset of zero, this means that it will output a code of zero for any voltage greater than or equal to zero and less than one. Any voltage equal to or above one, and it will output a one, and so on. ADC 1, because of the one-quarter code offset applied, outputs a zero for any voltage less than three quarters and greater than negative one quarter. ADC 2 has a half code offset applied and ADC 3 has a three quarter code offset.

Because of these specific offsets, we find that for voltages between zero and one quarter, all ADCs are outputting a zero. For voltages between one quarter and one half, all ADCs are outputting a zero except for ADC 3 which is outputting a one. For voltages between one half and three quarters, ADCs 0 and 1 are outputting a zero and the other two are outputting a one. Finally, for voltages between three quarters and one, only ADC 0 is still outputting a zero, while all other ADCs are outputting a one.

Thus, if we simply sum the codes coming from each of the ADCs, we obtain a ten bit code. For the single ADC 0, we got a zero code for all fractional voltages between zero and one code, but now, after summing the ADCs, we obtain four codes 0, 1, 2, and 3 for fractional voltages and therefore have a ten bit quantizer. The transfer characteristics of each ADC in the discussion is shown in Figure 12 where we see that the eight bit steps of the ramp applied now become ten bit steps.

The ADC deployment options are large and configurable. They range from eight 8-bit ADCs vertically interleaved to provide an 11 bit, 5 GS/s acquisition, to the eight ADCs time interleaved to provide an 8 bit, 40 GS/s acquisition. Time and vertical interleaving can even be combined as shown in Figure 13[5].

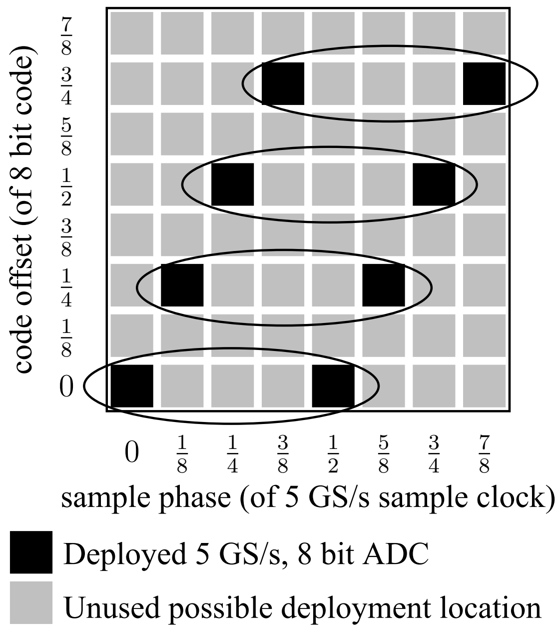

Figure 13: Conceptual 10 GS/s Hybrid Time and Vertical Interleave

Furthermore, if two ADC chips on two channels are combined, as is possible in many scopes, one could achieve between 80 GS/s eight bit acquisitions to 5 GS/s twelve bit acquisitions.

As a final, and not insignificant note, we've discussed here trade offs between sample rate and vertical resolution, but as we've seen, resolution might not be the main noise limitation. Sometimes, time interleaving combined with filtering or averaging ADCs values sampling simultaneously produces the largest noise reductions and therefore the highest ENOB.

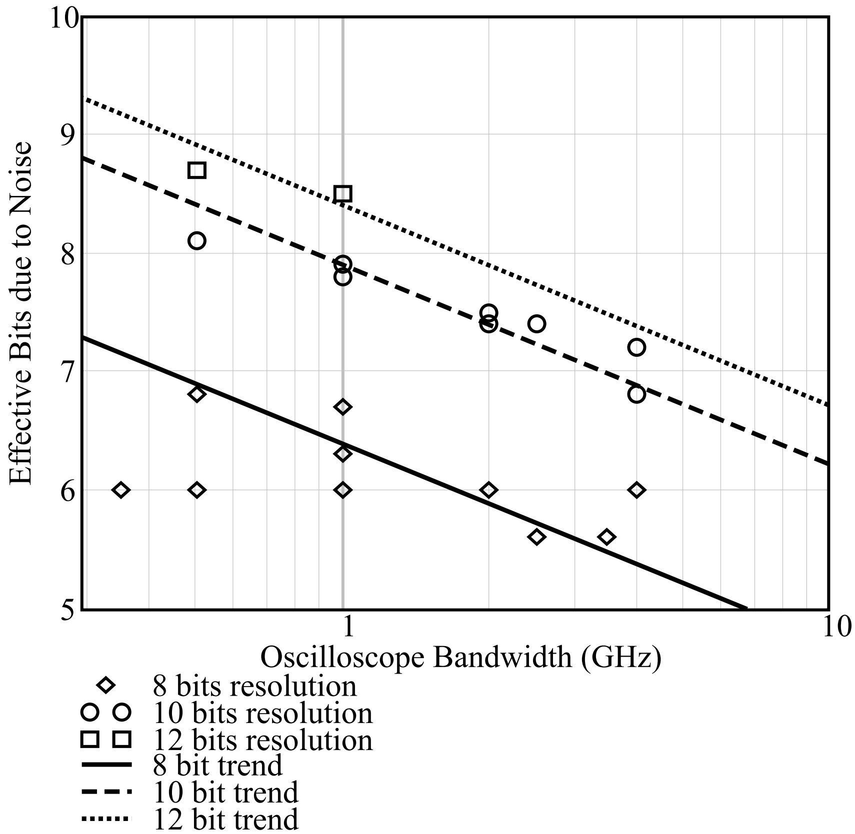

Figure 14: ENOB vs. Resolution for Various Bandwidths

Conclusion

Lately, ENOB and scope resolution have become very important topics. The chart in Figure 14 provides a comparison of oscilloscopes on the market from many manufacturers with eight, ten, and twelve bit resolution along with the ENOB provided. A trend line has been drawn for each of these resolutions that follows the expected half-bit improvement for every halving of bandwidth as previously discussed. To be clear, these trends weren't fitted to the data, but instead drawn according to the expectation. One conclusion you should draw is that for any scope resolution, ENOB decreases mostly as expected according to bandwidth. Another important conclusion is that although the higher resolution oscilloscopes tend to produce higher ENOB, the ENOB approaches but does not reach the resolution of the oscilloscope. Furthermore, one can observe that as resolution is improved, the main limit to ENOB seems to be other sources of noise in the oscilloscope channel.

References

[1] P. Pupalaikis, “The relationship between discrete-frequency s-parameters and continuous-frequency responses,” in *DesignCon*, IEC, February 2012.

[2] Teledyne LeCroy, “Computation of effective number of bits, signal to noise ratio, & signal to noise & distortion ratio using FFT.” February 2011.

[3] —, “Standard for Terminology and Test Methods for Analog-to-Digital Converters,” tech. rep., IEEE. Standard 1241-2010.

[4] B. Orwiler, *Oscilloscope Vertical Amplifers - Circuit Concepts*, pp. 21–38. Tektronix, 1 ed., 1969.

[5] M. McTigue and P. Byrne, “An 8-gigasample-per-second, 8-bit data acquisition system for a sampling digital oscilloscope,” *Hewlett-Packard Journal*, pp. 11–23, October 1993.