Effective Number of Bits for Arbitrary Waveform Generators

White Paper

Rev. 1.0

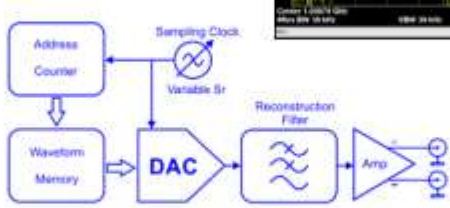

Table of Contents

1 Introduction

This white paper introduces the Effective Number of Bits (ENoB) concept and its impact on the real performance of AWGs in different applications. The ENoB specification is made of multiple parameters and the importance of each greatly depends on the signal being generated and the way it will be applied to the system under test.

The high-performance Tabor Proteus series of AWGs have been designed for high quality baseband and RF signal generation at frequencies up to 9 GHz. ENoB performance is an important component of this accomplishment. However, only the ENoB specification by itself cannot deliver the signal quality some applications require. This paper will deal with all aspects of linear and non-linear distortions and how they contribute to the overall performance, including ENoB, of an AWG such as the Proteus series from Tabor Electronics.

Figure 1.1 Tabor Electronics 9 GS/s Twelve Channel RF Arbitrary Waveform Generator P94812B

2 Effective Number of Bits (ENoB)

Arbitrary Waveform Generators (AWGs) convert a series of mathematically derived values (or samples) into a real analog signal. The signal samples are stored in the AWG's memory and are "played out" through the instrument's Digital-to-Analog Converter (DAC), at a rate defined by the instruments sample clock and a resolution equal to the number of bits the DAC has, see Figure 2.1, a), page 4. Already we can see that the waveforms will not be a perfect replica of their ideal analog counterparts, even if their mathematical definition is perfectly accurate. Each time-domain value is quantized, i.e., corresponds to a specific instance in time of a value defined by the sample clock or sample rate and DAC resolution. Therefore, not all continuous instantaneous levels of the signal are generated. The sampling rate defines the maximum frequency that can be generated by the instrument. This is defined by the Nyquist Sampling Theorem and it states the condition a given signal must meet to be reproduced without any loss of information. The theorem states that the highest component of the signal in the frequency domain must be strictly located below half of the sampling rate.

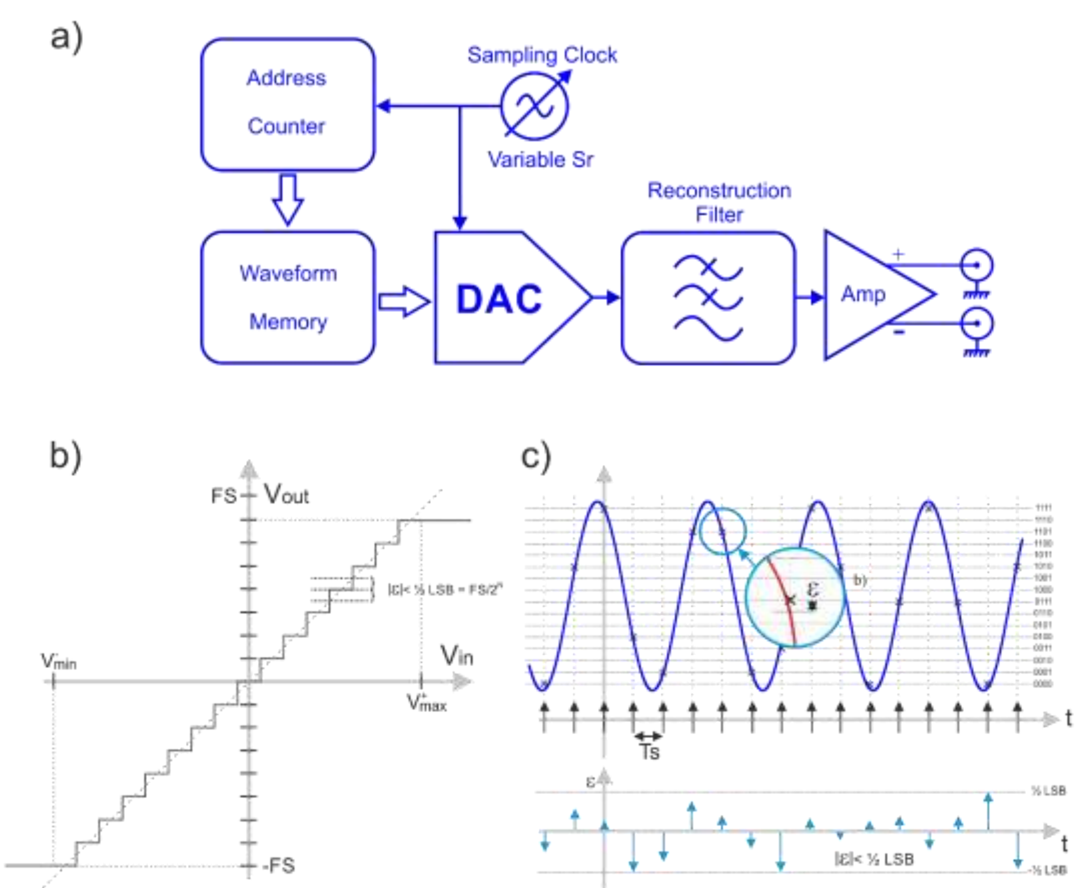

Figure 2.1 "True Arb" architecture AWGs (a) produce quantization noise because of the limited resolution of the DAC (Digital-to-Analog Converter) as seen in b). Quantization noise is random and wideband while the peak-to-peak amplitude is limited (c).

We also have quantization in the amplitude dimension, defined by the bit resolution of the DAC. Real-world DACs convert binary integers of a finite length (called resolution) to a given number of discrete integer levels, 2^N for a length of N bits, see Figure 2.1 b). Many modern DACs have excellent linearity across their range so the quantization level, or the distance between two consecutive levels, is the same throughout the full range of the part. Usually, the higher performance the DAC is, the better resolution is specified. However even high-performance DACs will still produce inaccurate waveforms as any level must be rounded to the nearest discrete integer level as discussed earlier. The error can be analyzed as a quantization noise signal added to the perfect signal, see Figure 2.1 c) . The peak-to-peak amplitude of this noise is equal to $Q = 2A/2^N$ (quantization level) and it is uniformly distributed over its full range. One way to characterize statistically the impact of finite sample resolution is by calculating the SNR (Signal-to-Noise ratio). For a simple signal such as a sinewave using the full DAC range:

$$S(t) = A * \sin(2*\pi*fc*t); \quad (1)$$

$$Q = 2A/2N \quad (2)$$

$$S'(nT) = Q \times \mathrm{round}((A + Q/2) * \sin(2* \pi *fc*nT) / Q) \quad (3)$$

For N>>1

$$S'(nT) = Q \times \mathrm{round}(A * \sin(2* \pi *fc*nT) / Q) \quad (4)$$

$$SNRQ(dB) = 6.02 *N + 1.76 \quad (5)$$

The above shows that the signal to quantization noise ratio improves by 6.02 dB for each additional resolution bit in the DAC. So, at least theoretically, the higher the resolution, the better signal fidelity. However, this assumes we have an ideal DAC.

Unfortunately, real world is not ideal at all. There are many linear and non-linear factors influencing the accuracy of any waveform being generated by an AWG. Even an ideal DAC in an AWG would incorporate thermal noise (unless it works at absolute zero temperature). A typical lab temperature is 25°C, therefore the thermal noise power density is -174 dBm/Hz. For the previously described sinewave being generated over a 50Ω load, the SNR for this thermal noise will be:

$$S(t) = A * \sin(2*p*fc*t); \quad (6)$$

$$PS(dBm) = 10 \log10(A2/(2*50)) + 30; \quad (7)$$

$$PN(dBm) = -174dBm + 10.0 \times \log10(BW (Hz)) = -174dBm + 10.0 \times \log10(SR / 2); \quad (8)$$

$$SNRT(dB) = PS(dBm) - PN(dBm); \quad (9)$$

An interesting exercise is applying the above expressions to a real case. For an SR = 9GSa/s AWG generating a 1Vpp (A = 0.5) sinewave, the signal to ideal thermal noise ratio would be:

PS(dBm) = 3.98dBm

PN(dBm) = -77.47dBm

SNR(dB) = 81.45dB

The above noise is Gaussian and not uniformly distributed as quantization noise. However, it is possible to define the resolution of an ideal DAC whose quantization noise is equal to the thermal noise for the Nyquist BW, in the above example:

6.02 x NE + 1.76 = SNR

NE = (SNR - 1.76) / 6.02 **(10)**

For the previous example, the above expression results in:

6.02 x NE + 1.76 = 81.45

NE = (81.45 - 1.76) / 6.02 = 13.24 bits

The N_E parameter is the equivalent number of bits and it basically expresses the maximum ideal resolution that can be implemented for a given sampling rate. In this case, beyond 14 bits of resolution the ideal quantization noise would be lower than thermal noise so there would not be any real advantage in increasing the resolution of the DAC, refer to Figure 2.2. In fact, given the Gaussian distribution of thermal noise, the peak-to-peak amplitude of quantization noise will be like that of thermal noise at a much lower resolution, or 11/12 bits. The equivalent number of bits decreases with sampling rate as thermal noise power is proportional to the bandwidth.

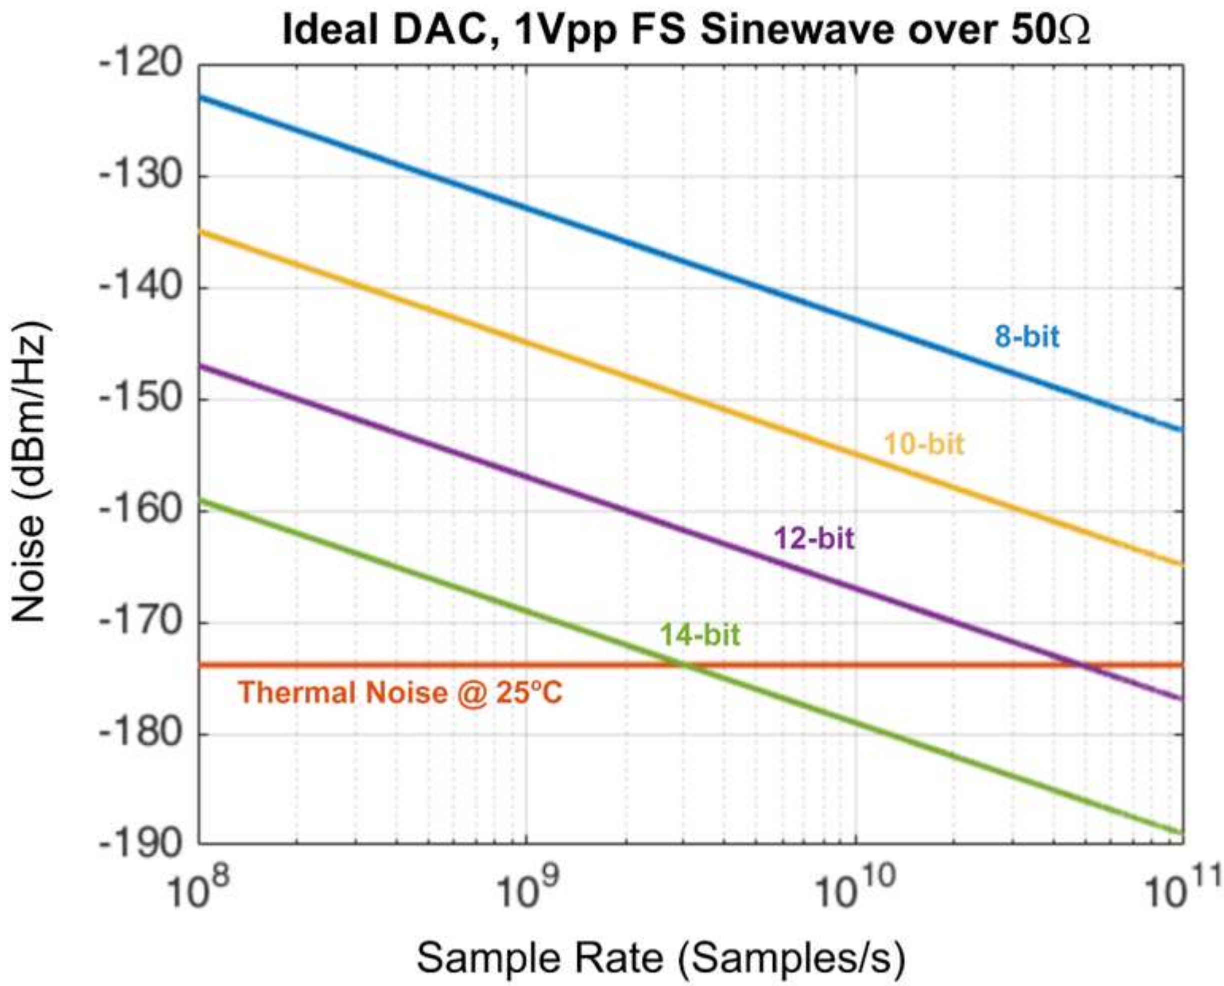

Figure 2.2: Quantization noise is not the only noise component at the output of an AWG. Thermal noise is an unavoidable noise source. While the total power of quantization noise remains constant with the sample rate, the thermal noise power is proportional to the usable bandwidth of the generator. In this graph, ideal quantization noise power density (NPD) for different resolutions of the DAC is shown. Thermal noise NPD remains constant, and for some specific sample rates, quantization noise is equal to it. For AWGs with higher sampling rates than this threshold, there is no real improvement in signal quality by using a higher resolution DAC.

SINAD (Signal-to-Noise-And-Distortion Ratio), see Figure 2.3: is another fundamental parameter that can be used to characterize any undesired signal being generated by a DAC. SINAD is defined as:

$$\mathrm{SINAD} = (\mathrm{PS} + \mathrm{PN} + \mathrm{PD}) / (\mathrm{PN} + \mathrm{PD}) \quad (11)$$

$$\mathrm{SINAD(dB)} = 10 \log_{10}(\mathrm{SINAD}) \quad (12)$$

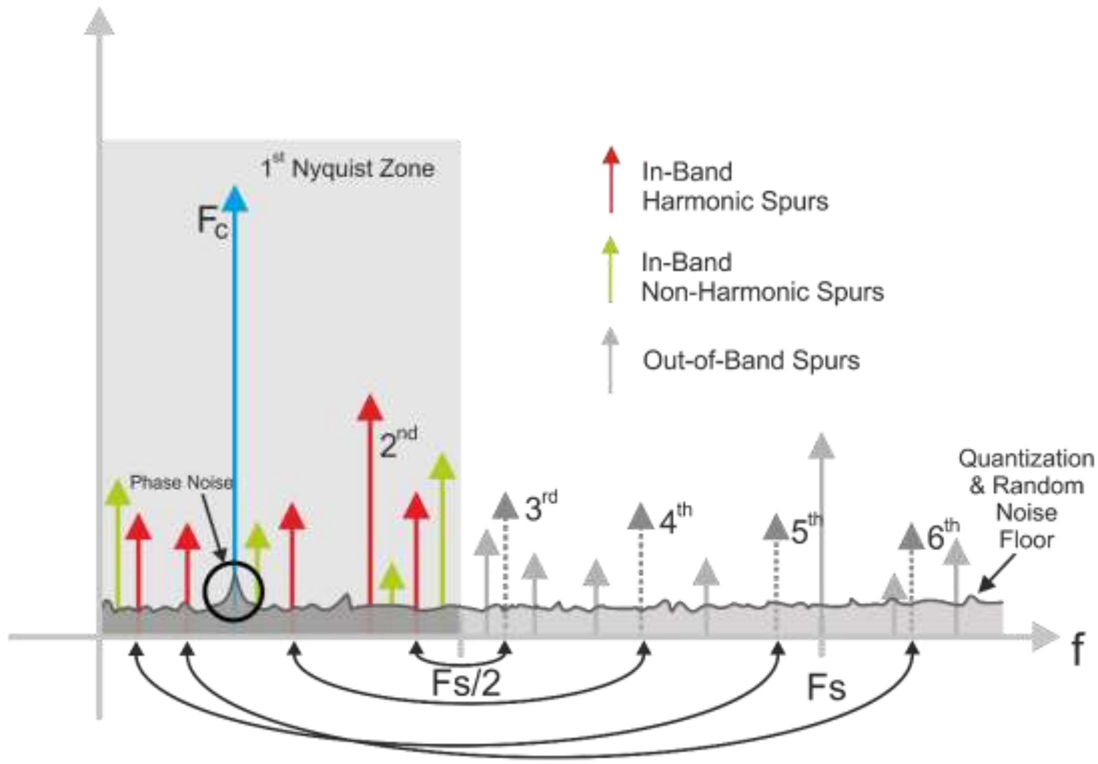

Figure 2.3: When generating a pure sinewave, an ideal, infinite resolution DAC would produce just the original sinewave plus multiple images around multiples of the sampling rate. Real-world AWGs produce many more components such as harmonics, random noise, phase noise and jitter, spurs, and quantization noise. The SINAD (Signal to Noise ratio and Distortion) quantifies the ratio between the desired signal and the undesired components within the Nyquist Bandwidth (DC-SR/2). Keep in mind that harmonics beyond the Nyquist frequency will be folded down to the first Nyquist Band as non-linear distortion will be also affected by sampling. SINAD is the basis for the ENoB (Effective Number of Bits) for AWGs.

The Effective Number of Bits (or ENoB) specification is just a convenient way to specify the same fundamental phenomena in an easy-to-understand single figure. ENoB simply applies expression (10) to SINAD

$$\mathbf{ENoB = (SINAD(dB) - 1.76) / 6.02} \qquad (13)$$

One way to interpret the ENoB specification is by assuming the number of bits in the DAC that are meaningful for a specific frequency. However, this specification is difficult to apply directly to a signal more complex than a sinewave. This becomes especially problematic when attempting to compare the overall noise and non-linear performance of different digital-to-analog conversion devices. Noise Power (P_N) in expression (11) includes several components such as quantization noise, switching noise, thermal noise, sampling clock phase noise and crosstalk. Distortion Power (P_D) is composed of harmonics and sub-harmonics. It is also important to remember that harmonics beyond the Nyquist frequency will show up as folded-down images in the first Nyquist band.

At first sight, ENoB should be always lower than the nominal DAC resolution as the overall noise will always include quantization noise and some additional noise and distortion. However, in many applications, just a fraction of the Nyquist band is used. This fraction can start at DC or be located around any frequency with a specified bandwidth. In this case noise and distortion out of the band of interest can be removed from the analysis. A good rule of thumb is to assume ENoB grows by one additional bit when the sample rate is multiplied by four. This process is called Oversampling which creates what we call Processing Gain. As a result, ENoB can be higher than the nominal resolution of the DAC when oversampling is applied.

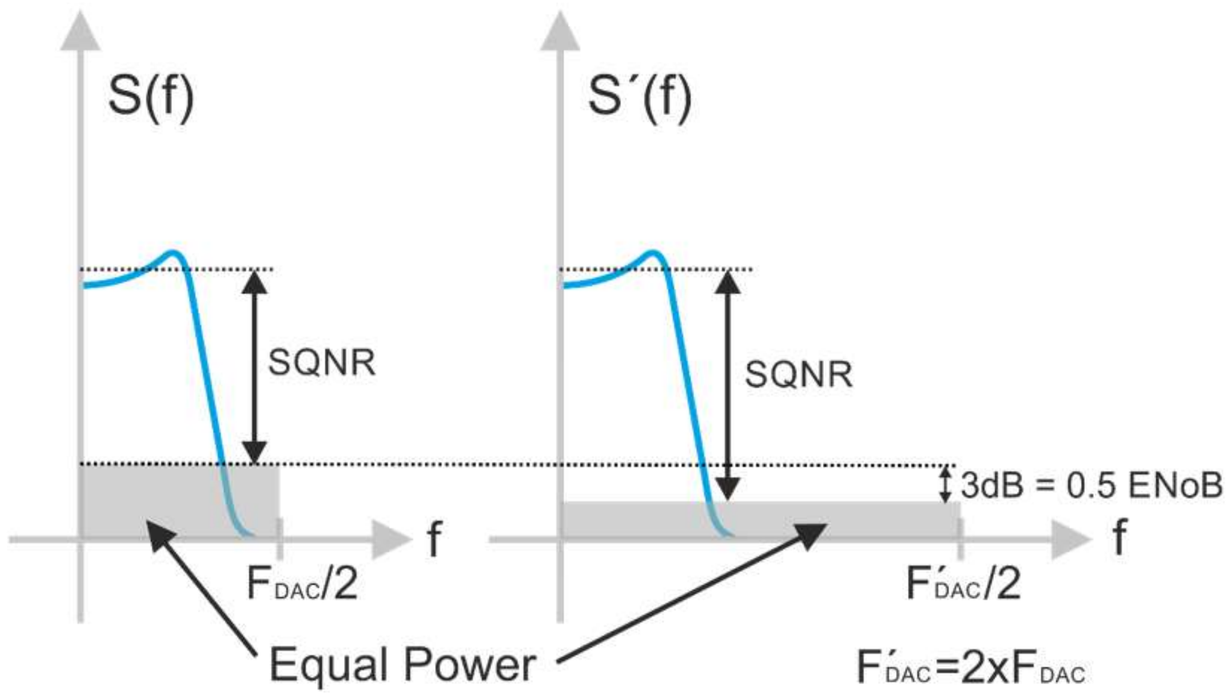

Figure 2.4: Quantization noise is spread all over the usable bandwidth of the AWG. This means that increasing the sample rate while keeping the same resolution of the DAC will result in a reduction of the NPD (Noise Power Density) as total power remains the same. The SQNR (Signal to Quantization Noise Signal) will improve by 3dB when sampling rate is doubled. This is equivalent to using a DAC with 0.5 more bits of effective resolution.

Finally, there are additional sources of error and noise:

- DAC and output stage non-linearities, both static and dynamic
- DAC switching glitches
- Increased thermal and random noise caused by electronic noise and crosstalk
- Non-flat frequency responses including zeroth-order hold DAC response
- Sampling uncertainty and sampling clock phase noise.

Noise, static, dynamic, linear, and non-linear performance of DACs (and AWGs) can be characterized through many different specs. Although many of them are straight-forward, others can be measured in different ways depending on the device manufacturer. To standardize the performances of DAC's the IEEE released the standard "IEEE 1658-2011 - IEEE Standard for Terminology and Test Methods of Digital-to-Analog Converter Devices". The standard defines noise and distortion measurements using sinewaves which can represent different sampling rate, test tone frequencies and amplitude.

3 Measuring ENoB in High-Speed AWGs

The IEEE 1658-2011 standard covers some test methodologies to obtain the SINAD parameter, so the ENoB specification can be directly calculated using expression (13). One of the methods, Figure 3.1: a), is based on acquiring the test signal (a sinewave) being generated by the DAC with a suitable ADC. Examining the signal in the frequency domain or breaking the signal into its constituent parts using an FFT (Fast Fourier Transform) allows us to identify and evaluate the signal. Now we can add together the power of all the spectral lines and extract the one corresponding to the test tone provides us with all the terms required to evaluate SINAD by using expression (11). A note of caution when making measurements - if the measurement device, say a Digital Storage Oscilloscope (DSO) - has an ENoB less than that of the DAC you would be simply measuring the ENoB of scope and not the device. To produce meaningful results, the ADC of the test instrument (DSO) used for the characterization must be better than the expected ENoB performance by at least 4 times (2 bits). Although, this may be feasible for low sample rate DACs, it is not for high-speed DAC's or AWG's. However, you can apply a tuned notch filter to the DAC output, so the fundamental test tone is removed, Figure 3.1: b). The signal at the output of the filter will contain all the unwanted signal components so its overall power combines all the noise and distortion components. This signal can be easily digitized and analyzed by the DSO utilizing the full dynamic range of the instrument, making the scope's own noise and distortion irrelevant. Frequency-domain processing (FFT or digital filtering) can be used to remove any remaining component of the test-tone further improving accuracy, also making the accuracy of the notch frequency not as critical. This method though does have drawbacks. The first one is related to the notch filter. The frequency response of the filter must be properly characterized so noise and distortion components are corrected according to their frequency. Additionally, the notch bandwidth in high-frequency filters may be wider than ideally required, so a fraction of the noise power may be lost resulting in a decrease in accuracy. The second drawback is an operational one, as we require one filter for each frequency to be characterized, which adds time and effort with respect to designing and characterizing multiple filters.

Improvement of the measurements can be achieved by utilizing a spectrum analyzer and some new measurement techniques developed by Tabor Electronics. Spectrum Analyzers differ from Oscilloscopes as they are essentially narrowband tunned receivers (defined by the resolution bandwidth of the analyzer), as opposed to an oscilloscope that is a broadband instrument. By utilizing the narrow bandwidth, the dynamic range and bandwidth required to characterize the high-speed DACs can now be realized. The spectrum analyzer can be tuned to measure the power of the test tone and all the signal components, Figure 3.1: c), thus accurately obtaining the SINAD parameter.

AWGs generate a tone by looping an integer number of cycles stored in the waveform memory. The signal must be calculated in such a way that all the quantization levels are used, and the exact same sequence of samples is not generated more than once in a single loop. What this means is that any SINAD component linked to the waveform, such as quantization noise, will also repeat, Figure 3.2 . This means that quantization noise will show up as pure tones at multiples of the repetition frequency and

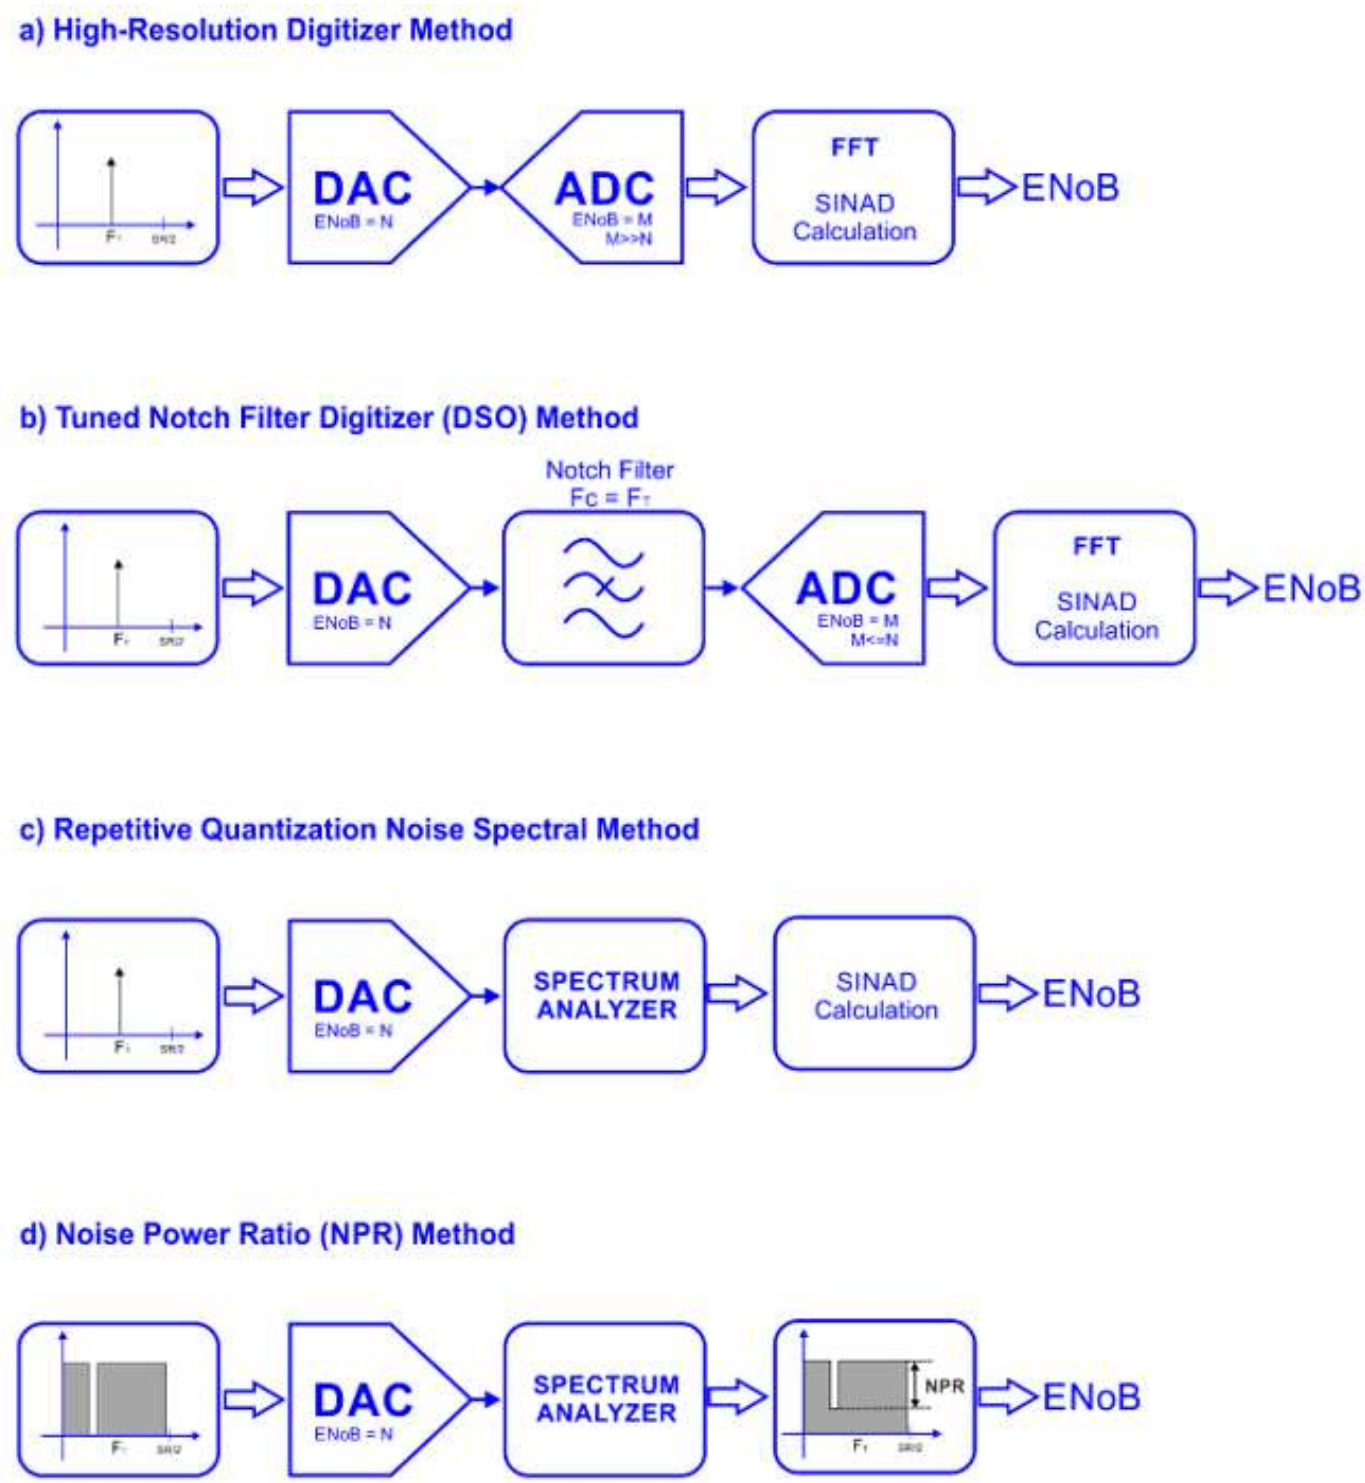

Figure 3.1: There are several methods to obtain the ENoB (Effective Number of Bits) parameter. Some of them (a, b. c) try to measure SINAD directly using multiple methodologies. Others, like the NPR (Noise Power Ratio) method (d) show the ENoB parameter indirectly. Test signals are also different. The SINAD based methodologies use single frequency sinewaves while the NPR method uses wideband random-like test signals.

they will be visible over the background, wideband noise. The background wideband noise must be corrected to extract the noise generated by the analyzer itself. Obtaining the overall narrowband components (quantization noise, clock feed-through, harmonics) and the corrected wideband noise results in the calculation of the SINAD parameter and ENoB. As the different components of SINAD (wideband noise, quantization noise, clock feed-through, harmonics, sub-harmonics, etc.) can be separated, this information can help to fully characterize the behavior of the AWG.

Another ENoB measurement technique using a Spectrum Analyzer is the NPR (Noise Power Ratio) technique, Figure 3.1: d). This methodology is very popular for ADC (and digitizer) testing however the same principle can be used for DACs as well. This technique uses a random wideband noise (ideally

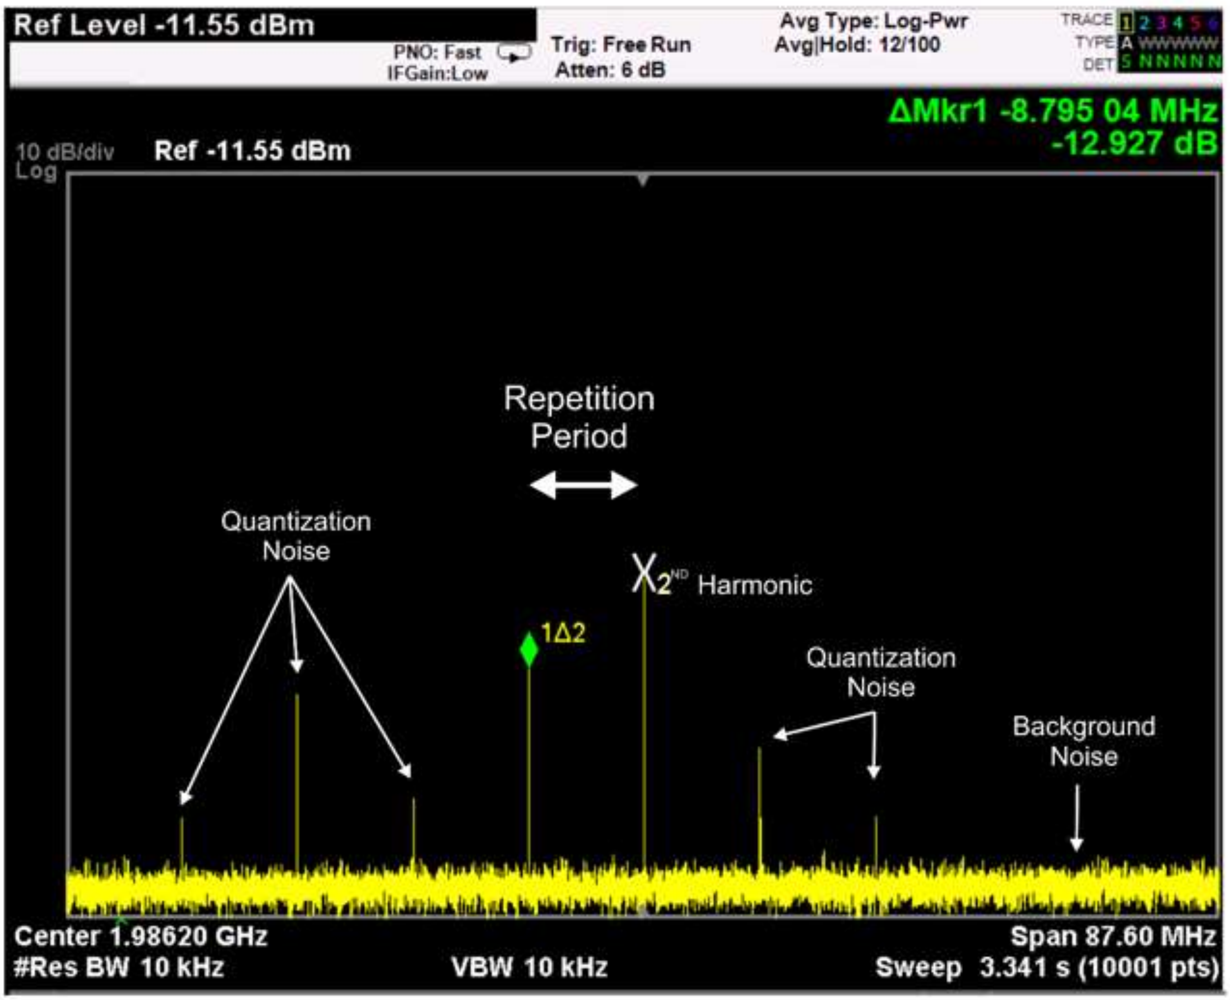

Figure 3.2 The Repetitive Quantization Noise Spectral Method can be implemented using Spectrum Analyzers as the measurement device. As the test signal repeats with some known period, quantization noise will also repeat, it shows up in the spectrum as discrete spectral lines at known frequencies. In this example, some of these spectral lines are shown for a 1GHz test tone. The line in the middle corresponds to the second harmonic while the other lines are caused by quantization noise. The background noise must be also measured and requires a separate procedure.

Gaussian) with a "notch" in it. The width of the notch must be a fraction of the AWG Nyquist band and it may be located at different central frequencies. For a perfect notch filter and DAC, the ratio between the noise power in the notch and that of a similar BW out of the notch as seen by a Spectrum Analyzer depends on the DAC resolution (NoB) and the PAPR (Peak-to-Average-Power-Ratio) is:

$$\mathbf{NPRdB = (NoB - 1) * 6.02 - PAPRdB + 10.79} \qquad \textbf{(14)}$$

In a real DAC, the actual NPR result can be used to obtain the ENoB by using the below expression:

$$\mathbf{ENoB = 1 + (NPRdB + PAPRdB - 10.79) / 6.02} \qquad \textbf{(15)}$$

This method can be used at different frequencies so ENoB can be characterized as a function of frequency, Figure 3.3: . However, the test conditions are very different from the test-tone methods described previously. In this case, a wideband signal is used, so the ENoB parameter will be more realistic for wideband signals. This is key when we want to understand the performance of the DAC with modulated signals, not just discrete sine waves.

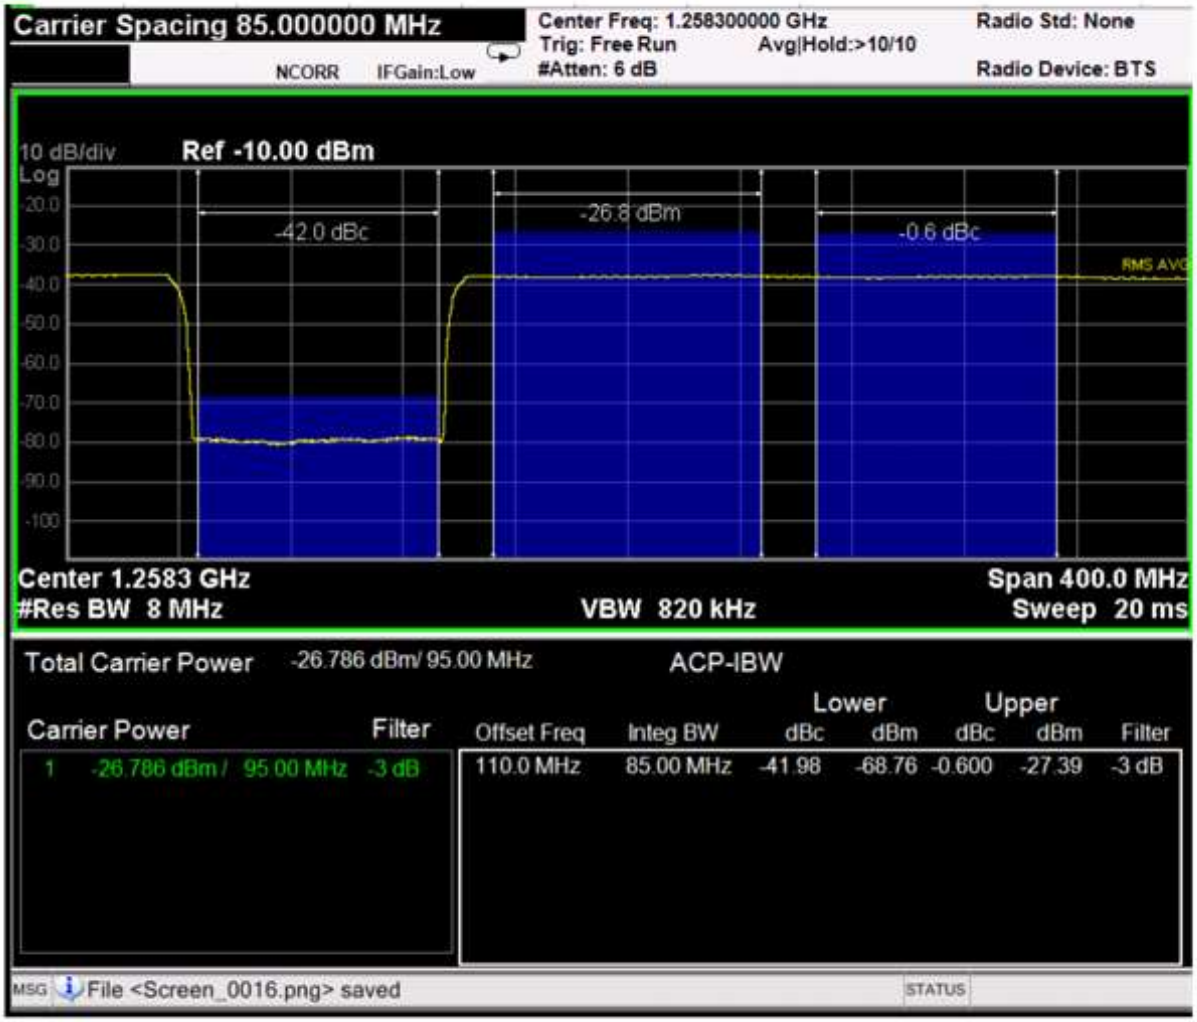

Figure 3.3: The NPR method derives the ENoB parameter from the power ratio between a notch in the spectrum and a close section of the same bandwidth of the non-filtered signal. In this example, an 8-bit, 9GSa/s test signal (multi-tone waveform with a notch around 1.1GHz) is generated. After computing the NPR (42.0 dB) and PAPR (10.3 dB) values, ENoB for this DAC is 7.85 bits.

4 Impact of ENoB in Actual AWG Performance

We have learned so far that the ENOB summarizes the impact of multiple linear and non-linear distortions and noise sources with respect to the quality of the output signal of an AWG. However, it does not provide all the information required to assess the performance of an AWG for different types of signal conditions. For example, if you are generating a signal and the most important factor is timing, i.e. the position of its rising or falling edges, such as those used in high-speed digital or radar applications; the vertical resolution is not as important as the position of the edges on the horizontal axis. Using the case of high-speed digital, an AWG can generate a multi-level PAM (Pulse-Amplitude Modulation) signal by defining the level corresponding to each successive symbol. If the symbol rate is a multiple of the sampling rate, then edges can be positioned in the right places without any added jitter (other than the sampling clock jitter itself). If this is not feasible, or the application requires a continuous edge positioning (i.e. to generate a controlled amount of jitter), the ideal transition instant would be located anywhere between two consecutive samples, so the resulting signal will incorporate an additional jitter component with a peak-to-peak amplitude equal to the sampling period of the AWG.

However, an accurate edge position can be achieved by controlling the amplitude of the samples in the transition between consecutive symbols. Maximum control on edge position is accomplished when the transition time for any edge is at least twice the sampling period, so two samples are used to implement the edge, see Figure 4.1: . The amplitude of those two samples will define the precise location of the corresponding threshold crossing event. In this case timing resolution is influenced by the vertical accuracy (and resolution) of the samples. In practice, it may be small compared to the intrinsic jitter of the generator caused by the sampling clock jitter and additive noise. Consider a 12-bit DAC @ 4GSa/s, this can generate a 4Gbps serial binary signal, however, a $250ps_{pp}$ jitter is added in the process. Alternatively, an 8-bit, 8GSa/s DAC can produce the same signal, but the jitter is now reduced to less than 1ps, refer to Figure 4.2: . The difference between the sampled signal and the ideal signal is an error signal, which can be considered as a reduction in the effective number of bits. In this example, the accuracy of the 8GS/s, 8-bit DAC will be much better than the one corresponding to the 4GHz, 12-bit DAC. In other words, for many signals, sampling rate and/or analog bandwidth may be more critical than the ENoB specification for a particular DAC.

Another source of confusion is the impact of vertical resolution and ENoB in the quality and accuracy of narrowband signals, such as those used in wireless, mobile telephony, and broadcast networks. The direct translation of equation (5) to the actual SFDR or SNR of a narrowband signal generated by an AWG at first glance may make the choice of an AWG unsuitable. First, it is important to remember that quantization noise is spread all over the Nyquist band, so the most important number is noise power density (NPD). NPD depends roughly on the ENoB specification and the sampling rate. The higher the sampling rate, the lower NPD level for the same ENoB specifications. SNR for a narrowband signal is defined by the noise within the channel as the noise outside the channel can either be filtered out at the output of the generator or at the input of the DUT. However, ENoB does not provide a complete description of the different contributors to the degradation of the vertical accuracy of a DAC. While quantization noise is spread rather uniformly over the full Nyquist band, the nonlinear behavior of the

DAC or the output stage results in harmonics and Third Order IMD (Inter-Modulation Distortion). Harmonics and Intermodulation can appear in or out of the band of interest. A very-linear DAC with a lower ENoB specification may be much better than a not-so-linear DAC with a higher ENoB number as the latter will concentrate the power of the error signal in the proximity of the signal being generated, see Figure 4.3 . If a signal is generated in the center region of the Nyquist band, all the harmonics created by the non-linear behavior of the generator will show up (folded down) in the same band as the desired signal, thus leading to interference and ACPR degradation. Again, a lower ENoB specification may be much less important than the overall linearity of the DAC and the output stage.

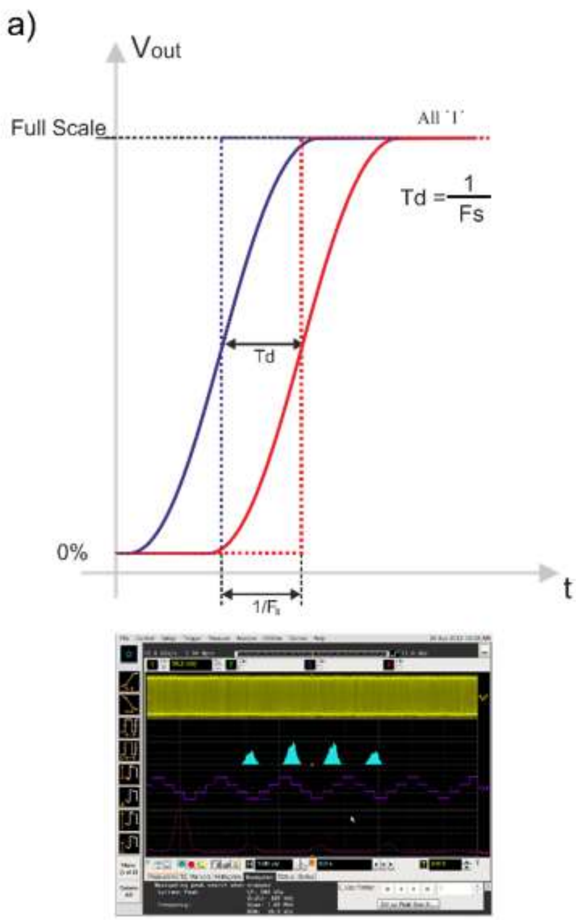

Figure 4.1: In some areas, timing information may be more important than the instantaneous amplitude of the waveform. The capability of an AWG to position an edge depends on the number of samples used to implement that edge. In a) the edge is produced by simply switching the amplitude from minimum to maximum value in one sampling time (top). This will produce a very fast edge, but the granularity of the edge position will be 1/Fs. If the application requires the generation of a sinusoidal jitter profile, an analysis of the jitter behavior will show the quantization of the instantaneous jitter to multiples of 1/Fs (bottom). If two or more samples are used in the transitions, as shown in b), edge position granularity will be greatly improved, even for low resolution DACs. In the bottom jitter analysis, the same amount of sinusoidal jitter is produced without any noticeable quantization.

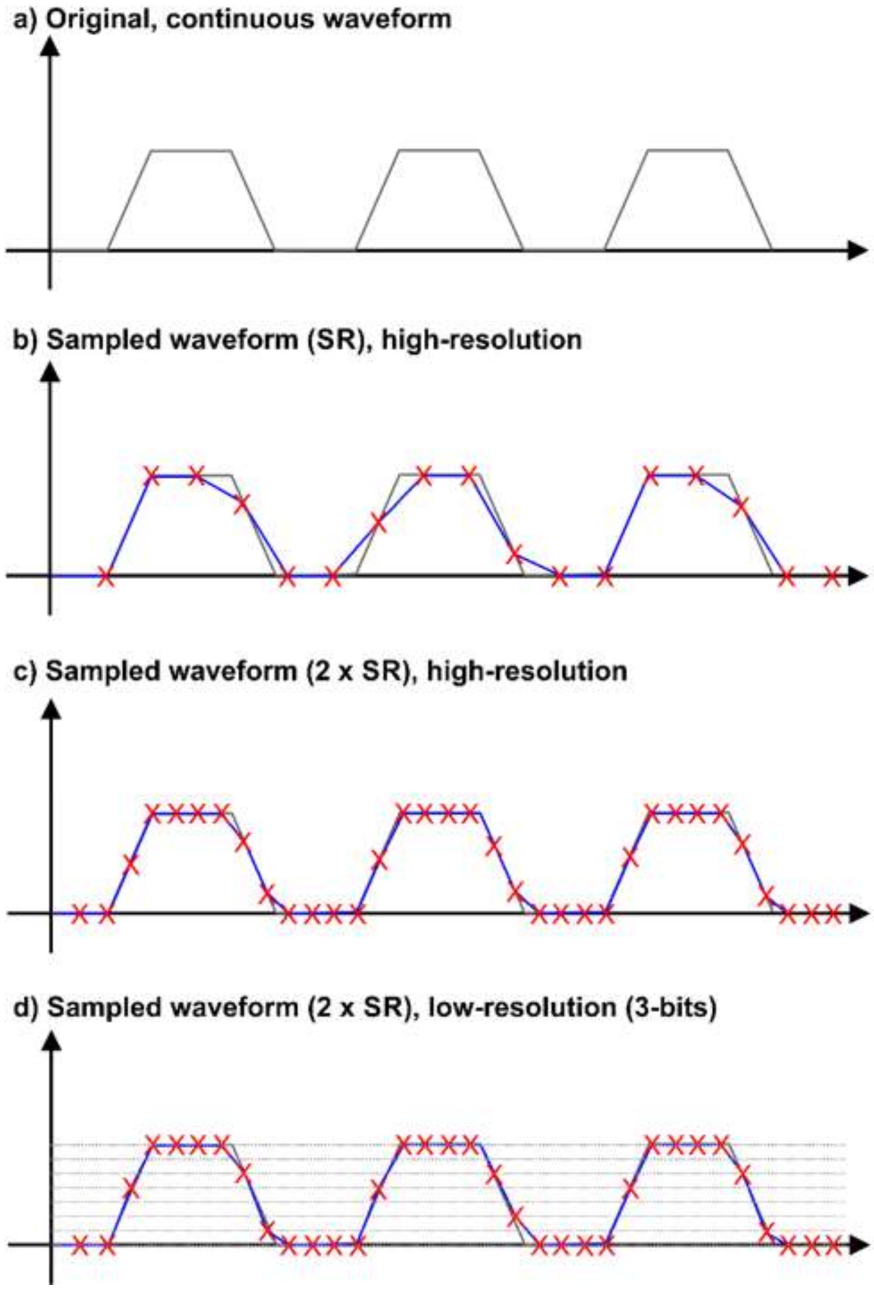

Figure 4.2: Increasing the ENoB specification is not always useful to improve signal generation accuracy. Often, the sample rate and bandwidth may be more critical to the application. In this example, a clock signal (not a submultiple of the sampling rate) is generated (a). The sampled version of this signal is shown in b). The differences with the original waveform are very visible and affect to the timing location of some of the edges. In c) the same signal is generated at twice the sampling rate. As the sampling rate is higher and the 2-sample time per edge criteria is met, the resulting signal is much close to the original one while the 50% crossings occur exactly at the same time as the original signal. Finally, in d) a low-resolution (3-bit) quantized version of the previous waveform is shown. Even with such a low resolution (just 8 quantization levels), signal fidelity and timing consistency is much better than the ones corresponding to the high-resolution, half sample rate version of it.

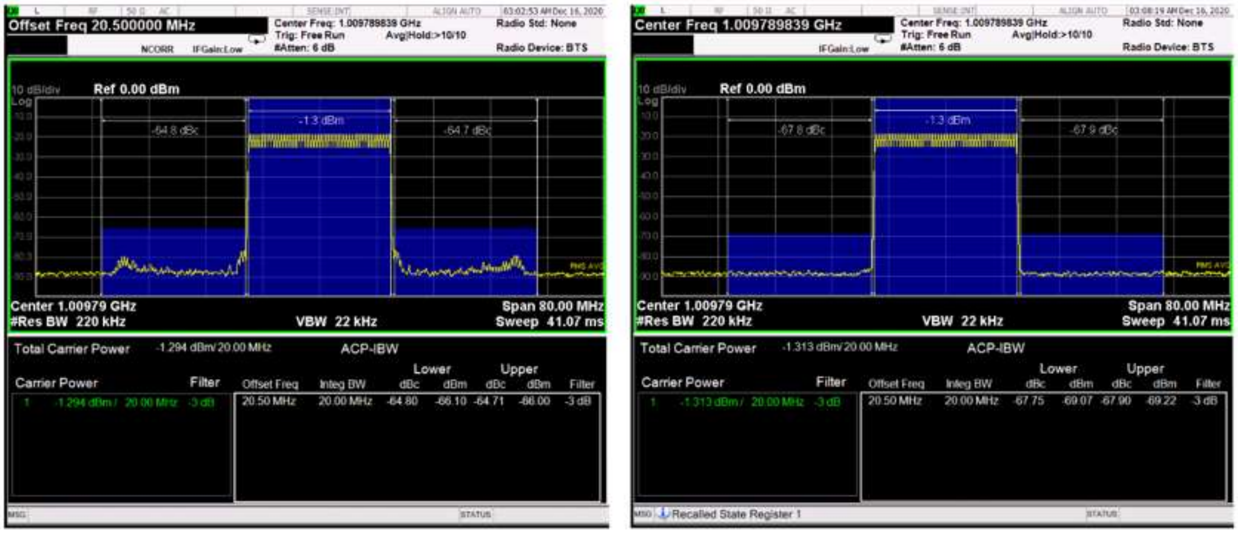

Figure 4.3 ENoB groups multiple sources of error. Narrowband signals may be more influenced by some components. On the right a 20MHz bandwidth signal @ 1GHz is measured using two AWGs with similar ENoB specifications. However, the signal is generated using two different output stages. In the left, where an amplified output stage is used, intermodulation products and skirts are visible in the adjacent channels. On the righthand side, using a direct DAC connection, these unwanted components are buried in the background noise. As a result, the ACPR (Adjacent Power Ratio) improves by more than 3dB.

Avoiding the central region of the Nyquist band is a good practice but it is not always possible, especially for signals with very high modulation bandwidths. Finally, in some applications it is possible to improve the SNR and SFDR of the signal by applying some advanced correction techniques. These are based on the “noise shaping” principle. (A technique used in Hi-Fi audio.) We know that quantization noise will always be generated by a digital-to-analog converter, but the signal could be created in such a way, so the noise is not uniformly distributed all over the Nyquist band and removed from some band of interest. But remember the noise out of this band will increase accordingly. Sigma-Delta modulation is a technique that can be applied after a careful characterization of the DAC. However, characterizing the DACs to the level of accuracy required by the signal-delta technique may be very difficult, especially with high speed AWGs. Nulling-Tones, refer to Figure 4.4: , is a different technique where IMD products can be removed within the band (i.e. for a notch in a multi-tone signal) or in the adjacent channels (i.e. for a more accurate ACPR measurement of an RF amplifier). The nulling-tones technique is applied by calculating tones that interfere destructively with the undesired IMD products. An iterative calibration process using an external Spectrum Analyzer can result in an excellent improvement of the SFDR that would be impossible to accomplish even with the highest performance DACs.

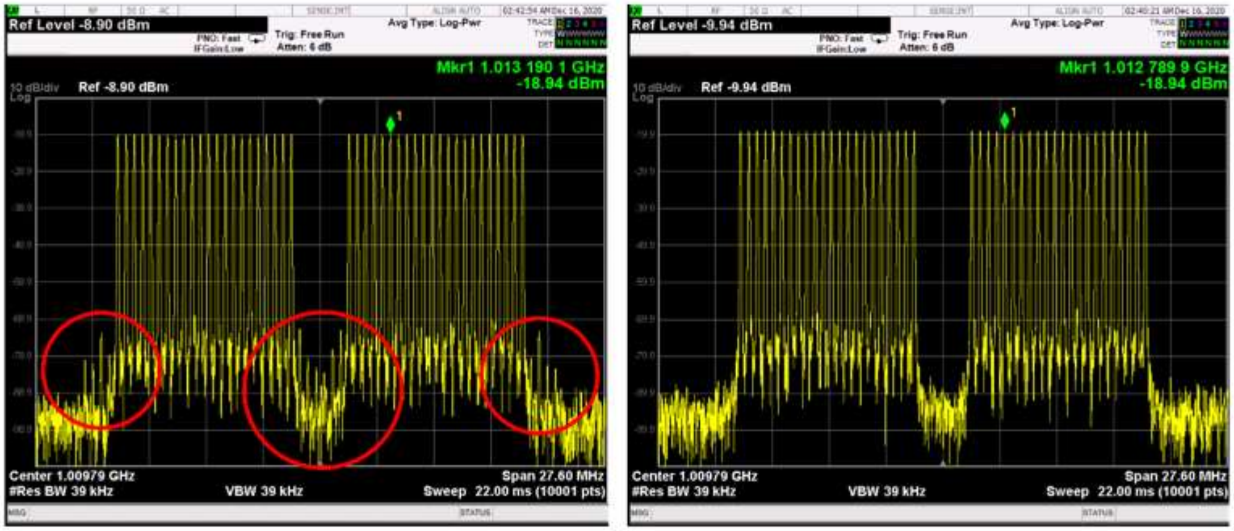

Figure 4.4: Multiple correction techniques can be applied to improve the SNR of the generated signals and is equivalent to use a higher ENoB DAC. Noise shaping techniques, where noise is removed from some bands of the spectrum, is one of these techniques. In this example, the "nulling-tones" technique is used. On the left, an uncorrected bandwidth-limited multi-tone signal with a notch is shown. IMD products are clearly visible in the notch and the adjacent channels (circled in red). On the right, after applying the "nulling-tone" technique, the SFDR has been improved by over 20dB, and the IMD products are buried in the noise - improving the dynamic range of any NPR test.

The frequency response of an AWG can also impact the ENoB performance as the amplitude frequency response for most AWGs goes down as frequency increases. The reduced amplitude at higher frequencies results in the reduction of the ENoB. However, we live in a world of wideband signals such as WLAN, 5G and UWB, therefor to maintain excellent flatness linear correction of the frequency response is required. Corrections in theory boost the higher frequencies components while keeping the amplitude of the lower frequency components. However, you cannot boost the signal amplitude more than the DAC range, so in reality you reduce the overall amplitude of the lower frequencies (renormalization). This yields a new corrected signal with a lower amplitude than the original one. Yet SINAD stays the same effectively reducing the ENoB for the lower frequency components (one bit for every 6dB of loss). This must be understood as trade-off, as in some cases better frequency response may be much more valuable in terms of signal generation accuracy than extra effective bits.

5 Sample Resolution vs. DAC Resolution

Some AWGs and DACs use samples with a higher resolution than the DAC itself. The Tabor Proteus uses 16-bit samples while the DAC core is 14-bit. The advantage of this method is that SNR will be minimally affected by any calculations for any intermediate signal processing. For example, a common processing stage is an x/sin(x) digital filter that is usually applied before data conversion to flatten the frequency response. Using higher resolution samples means "calculation noise" will not affect the validity of any of the actual bits being applied to the DAC.

Other real-time processing would also result in the loss of effective bits with respect to the input signals. Interpolation, for example, see Figure 5.1: . Oversampling (interpolating) DACs are common as they can increase the Nyquist BW, moving the nearest out of the band of interest, or they can spread the quantization noise over a larger BW, thus reducing the noise density.

Oversampling is like increasing the number of bits, you get one additional bit every time the sampling rate is multiplied by four. However, to achieve a processing gain like this the input samples must have a resolution equal or better than the improved resolution. Let us look at how an interpolator works; Interpolators take the input sequence of samples, insert zeros to increase sampling rate (zero-padding), and then apply a sharp roll-off LPF to get rid of the unwanted images. As input signals are finite resolution integers, these are already carrying the quantization noise corresponding to the integer length being used over the non-interpolated sample rate Nyquist band. Zero-padding will result in the replication of the spectrum of the signal over the new Nyquist band, including the original quantization noise. The final Low-Pass filter will add some additional calculation noise. Yielding little or no improvement to the SNR.

As mentioned earlier the Tabor Proteus architecture uses 16 bits samples, and interpolation capability up to 16x (2 bits), and utilizes a 14 bit DAC, so it can take advantage of processing gain with improved SNR. The same considerations apply to other signal processing blocks such as DUC (Digital-Up-Converters) and I/Q modulators where interpolation and real time arithmetic are also applied. When higher resolution samples are not used in interpolation it is the equivalent of losing effective bits from the DAC. If 14-bit input samples would be used in Proteus, then the output of the DAC would be equivalent to using a 12-bit DAC when interpolating by 16x.

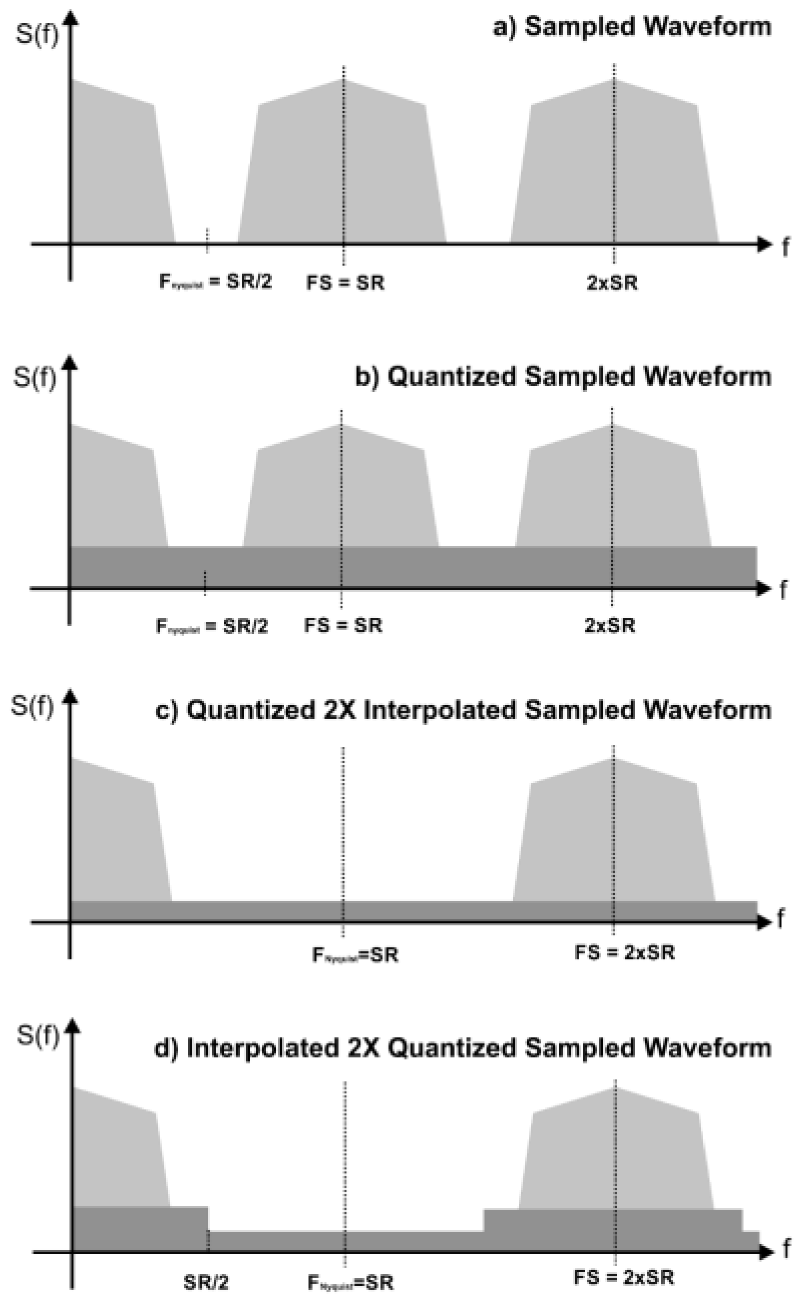

Figure 5.1: Sample resolution may become important when using DSP techniques to improve the behavior of an AWG. Interpolation is one of these cases. In the above example, a signal is calculated and stored with a high accuracy (a). When this signal is generated by an AWG, quantization noise will be added to it(b). If the signal is ideally interpolated (x2 in this example), the same quantization noise will be spread over a higher bandwidth so the SQNR will improve (c). However, if the signal is stored with the resolution of the DAC before being interpolated, quantization noise will be also interpolated so the SQNR will be kept and no processing-gain will be obtained (d). The Tabor Proteus uses 16-bit samples although the DAC is 14-bit for this specific reason.

Document Revision History

Table Document Revision History

Revision	Date	Description	Author
1.0	13-Jan-21	• Original release.	Joan Mercade

Acronyms & Abbreviations

Table Acronyms & Abbreviations

Acronym	Description
µs or us	Microseconds
ADC	Analog to Digital Converter
AM	Amplitude Modulation
ASIC	Application-Specific Integrated Circuit
ATE	Automatic Test Equipment
AWG	Arbitrary Waveform Generators
AWT	Arbitrary Waveform Transceiver
BNC	Bayonet Neill–Concelm (coax connector)
BW	Bandwidth
CW	Carrier Wave
DAC	Digital to Analog Converter
dBc	dB/carrier. The power ratio of a signal to a carrier signal, expressed in decibels
dBm	Decibel-Milliwatts. E.g., 0 dBm equals 1.0 mW.
DDC	Digital Down-Converter
DHCP	Dynamic Host Configuration Protocol
DSO	Digital Storage Oscilloscope
DUC	Digital Up-Converter
DUT	Device Under Test
ENoB	Effective Number of Bits

Acronym	Description
ESD	Electrostatic Discharge
EVM	Error Vector Magnitude
FPGA	Field-Programmable Gate Arrays
FW	Firmware
GHz	Gigahertz
GPIB	General Purpose Interface Bus
GS/s	Giga Samples per Second
GUI	Graphical User Interface
HP	Horizontal Pitch (PXIe module horizontal width, 1 HP = 5.08mm)
Hz	Hertz
IF	Intermediate Frequency
I/O	Input / Output
IP	Internet Protocol
IQ	In-phase Quadrature
IVI	Interchangeable Virtual Instrument
JSON	JavaScript Object Notation
kHz	Kilohertz
LCD	Liquid Crystal Display
LO	Local Oscillator
MAC	Media Access Control (address)
MDR	Mini D Ribbon (connector)
MHz	Megahertz
ms	Milliseconds
NCO	Numerically Controlled Oscillator
ns	Nanoseconds
PAM	Pulse-amplitude Modulation
PC	Personal Computer
PCAP	Projected Capacitive Touch Panel

Acronym	Description
PCB	Printed Circuit Board
PCI	Peripheral Component Interconnect
PXI	PCI eXtension for Instrumentation
PXIe	PCI Express eXtension for Instrumentation
QC	Quantum Computing
Qubits	Quantum bits
R&D	Research & Development
RF	Radio Frequency
RT-DSO	Real-Time Digital Oscilloscope
s	Seconds
SA	Spectrum Analyzer
SCPI	Standard Commands for Programmable Instruments
SFDR	Spurious Free Dynamic Range
SFP	Software Front Panel
SINAD	Signal-to-Noise-And-Distortion Ratio
SMA	Subminiature version A connector
SMP	Subminiature Push-on connector
SPI	Serial Peripheral Interface
SQNR	Signal to Quantization Noise Signal
SRAM	Static Random-Access Memory
TFT	Thin Film Transistor
T&M	Test and Measurement
TPS	Test Program Sets
UART	Universal Asynchronous Receiver-Transmitter
USB	Universal Serial Bus
VCP	Virtual COM Port
Vdc	Volts, Direct Current
V p-p	Volts, Peak-to-Peak

Acronym	Description
VSA	Vector Signal Analyzer
VSG	Vector Signal Generator
WDS	Wave Design Studio

Resources & Contact

For more information on Microwave signal generation challenges and solutions, review the following resources:

- White Paper: Multi-Nyquist Zones Operation-Solution Note
- Solution Brief: Quantum bits – Solution Note
- Data Sheet: Proteus – Arbitrary Waveform Generator
- Online Webinar: Advanced Microwave Topics for Quantum Physicists Topics

Stay Up To Date

- www.taborelec.com
- LinkedIn page
- YouTube channel

Corporate Headquarters

Address: 9 Hata'asia St., 3688809 Nesher, Israel
Phone: (972) 4 821 3393
Fax: (972) 4 821 3388

For Information

Email: info@tabor.co.il

For Service & Support

Email: support@tabor.co.il

US Sales & Support (Astronics)

Address: 4 Goodyear Irvine, CA 92618
Phone: (800) 722 2528
Fax: (949) 859 7139

For Information

Email: info@taborelec.com

For Service & Support

Email: support@taborelec.com