


High-Speed DACs

Application Note







Table of Contents

Table of Contents	2	High-speed DACs for AWGs	10
DAC Basics	3	Binary Weighted Element DACs	11
The Sampling Theorem	4	Unary Element DACs	12
Zeroth-Order Hold Response	5	Segmented DACs	12
Reconstruction Filters	5	Switching glitches and resampling techniques	13
DAC Performance and Specifications	6	Extended Bandwidth DAC Architectures	14
Quantization noise	6	Interleaving DACs	14
Non-Linear Distortions	7	Return-to Zero (RZ) DACs	16
SFDR and Effective Bits	9	Mix-Mode (Doublet Mode) DACs	16
		Conclusions	17



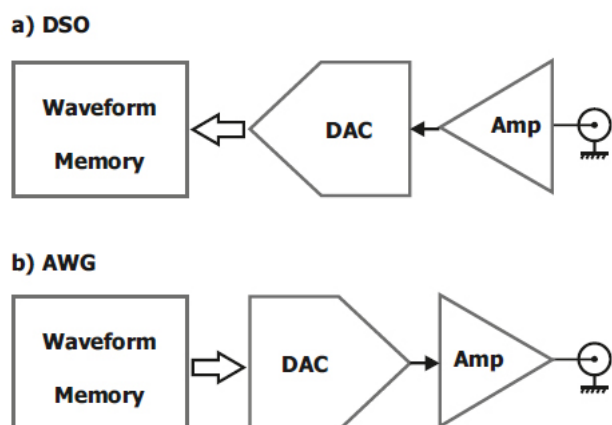


Figure 1. The block diagram of an Arbitrary Waveform Generators (AWG, b) is very similar to that of a real-time Digital Storage Oscilloscope (DSO, a). The only visible difference is the flow direction of the signal. The Digital-to-Analog Converter (DAC) is probably the most critical component in an AWG just as the Analog-to-Digital Converter (ADC) is for a DSO.


DAC Basics

The basic block diagram of an Arbitrary Waveform Generator (or AWG for short) looks very similar (Figure 1) to that of a Real-Time Digital Storage Oscilloscope (or DSO). In fact, the only visible difference is the direction taken by the waveform. In a DSO, a real-world signal becomes a digital mathematical entity known as “waveform” after being converted to a series of equally spaced digital samples by an Analog-to-Digital Converter (or ADC). These samples are transferred to a waveform memory for later processing, display and analysis. In an AWG, everything starts with the waveform stored in a memory from where the samples are read at a constant pace (or sampling rate) and converted to a real-world signal



Figure 2. The Tektronix AWG70000 series of Arbitrary Waveform Generators include 25GSa/s, 2-channels and 50GSa/s, 1-channel models. Both instruments use the TDAC25 DAC designed and manufactured by Tektronix Components Solutions.

through a Digital-to-Analog Converter or DAC. Most of the same mathematics applies to both cases, and both devices, DSOs and AWGs, handle sampled signals and share most limitations and performance issues. They also share many of their specifications and characteristics such as sampling rate, vertical resolution or record length, just to name a few. For both instruments classes, the most critical component is, perhaps, the data converter that must be the bridge between the ideal digital realm of waveforms and the real analog world of signals, the DAC for any AWG. The purpose of this paper is to show the characteristics of the DACs being used in state-of-the-art, high-speed Arbitrary Waveform Generators and how the required level of performance is reached. The term “high-speed”, as in any other area of technology, changes over time. Just 10 years ago, state-of-the-art AWGs reached around 4GSa/s. Nowadays, Tektronix has introduced the AWG70000 series (Figure 2) with a maximum sampling speed of 50GSa/s. Such an accomplishment is only possible through the availability of faster DACs and the clever implementation of architectures that overcome some of their intrinsic limitations.



Application Note

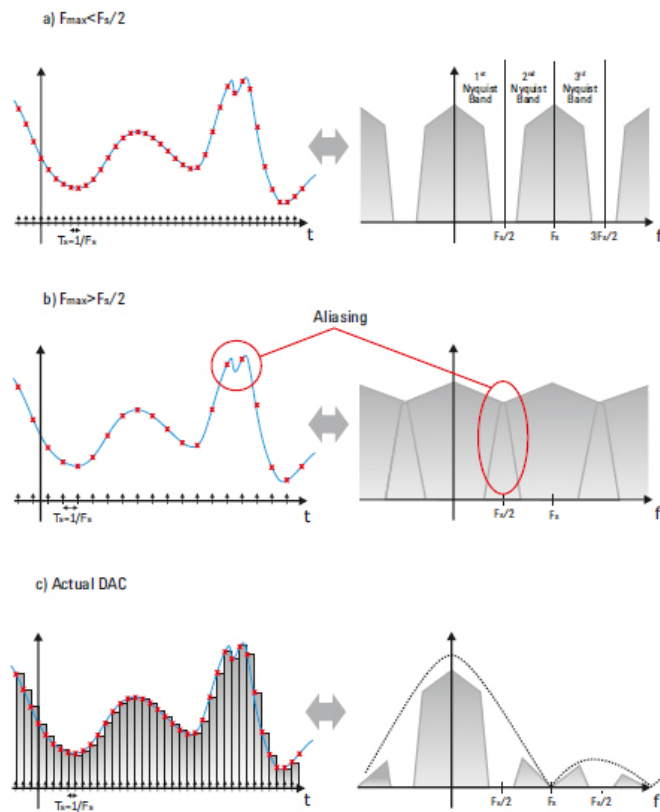


Figure 3. Waveforms are discrete versions of signals made by samples taken at a constant sampling rate (S_r). Sampled signals in the time domain when analyzed in the frequency domain result in infinite replicas of the original signal spectrum (called images) located around multiples of the sampling rate (a). If the signal bandwidth is larger than half the sampling rate, information about the original signal is lost and cannot be recovered. This phenomenon is known as aliasing (b). Actual DACs generate square samples so ideally the same level is held during the sampling period ($1/S_r$), this zeroth-order hold response in the time-domain corresponds to the well-known $\text{Sinc}(\pi f/S_r)$ response in the frequency domain (c).

The Sampling Theorem


Waveforms are a digital representation of a real-world signal. Waveforms used in AWGs are made of a series of equally spaced, limited resolution binary samples. It is important to understand that, unlike in scopes, the user must supply the waveforms used in AWGs. These waveforms may come from simulation or mathematical packages, user programs, or even from a waveform capture instrument such as a DSO or a Real-Time Spectrum Analyzer (RTSA). Waveform calculation may be very simple in some cases and extremely difficult in others.

Most waveform calculation operations use floating-point math as it provides the best accuracy but calculated samples must be converted to integer binary values of a limited resolution when they are transferred to the waveform memory attached to the DAC. Generally speaking, the more bits it uses for conversion the more accurate and higher quality analog level will result, so vertical resolution stated in bits is a primary specification of any DAC.

The maximum frequency at what samples can be converted is another important specification of any DAC. The speed challenge impacts many aspects of the DAC internal design and its interconnection with the external, typically massive, waveform memory. Conversion frequency limits the maximum frequency contents of a sampled signal that the converter can recover. The spectrum (Figure 3a) of an ideal sampled waveform (a series of Dirac's delta in the time-domain) consists of the original signal spectrum repeated infinite times around multiples of the sampling rate (S_r). A simple visual analysis of this spectrum allows extracting two conclusions:

- Any signal with frequency contents beyond $S_r/2$ will not be fully recoverable through low-pass filtering as the spectrum will be distorted by the unwanted superposition of multiple images of the original one, the frequency-domain consequence of the phenomena known as aliasing (Figure 3b).
- Any waveform designed for DAC conversion should be bandwidth-limited from DC to $S_r/2$ (also known as Nyquist Frequency). This fact is known as the Sampling Theorem (Figure 3a). For bandwidth-limited signals (such as a modulated carrier) any convenient image may be selected by applying a band-pass filter. Any of the valid frequency ranges (DC- $S_r/2$, $S_r/2$ - S_r , etc) is known as Nyquist Band (or Region) and numbered according to its position starting with #1 (DC- $S_r/2$).

Most AWGs are designed to work mainly in the first Nyquist band as this is the only band where both low-frequency (even DC) and high-frequency components can be generated at once. However, some applications such as RF signal generation may use signals in the second or even the third Nyquist band.



Zeroth-Order Hold Response


Real-world DACs do not generate anything close to a train of Dirac's deltas. Typically, they are designed to quickly change the output level at the sampling instant and keep this level for the whole sampling period T_s ($1/S_r$). This response is known as "zeroth-order hold" and it is equivalent to the time-domain convolution of the original, ideal train of Dirac's deltas and a square pulse of T_s duration. This operation results in a characteristic $\text{Sinc}(\pi f/S_r) = \text{Sin}(\pi f/S_r)/(\pi f/S_r)$ low-pass response (Figure 3c) with nulls at frequencies multiples of S_r . The first consequence of this response is that there is some linear distortion in the first Nyquist Band. The second one is that the amplitudes of images located in Bands # 2 and higher is reduced. On one hand this fact makes them easier to filter-out, on the other hand it reduces the usability of the images. Typically, only signals located in Nyquist Bands #2 and #3 are actually used in some applications. The previous discussion does not include the effects of the analog bandwidth of the DAC itself or any amplifier or filter attached to its output. Analog bandwidth of a DAC (or any related spec such as its rise-time) is, then, another factor of merit to consider.


Reconstruction Filters

Recovering the original signal from a sampled waveform requires a bandwidth-limited signal, a sufficient sampling speed, and a reconstruction filter. The functions of any reconstruction filter must be the following:

- Remove all the unwanted images from the output signal. Those frequency-domain images result in non-smooth transitions between samples (stair-step response). For Nyquist Band #1 this requirement results in a low-pass filter response.
- Equalize the response in the target band (usually Nyquist Band #1) to remove the linear distortion caused by the DAC' $\text{Sinc}(\pi f/S_r)$ response.

Cut-off frequencies for reconstruction filters depend on the DAC sampling rate and the target Nyquist Band. As real-world filters require some roll-off band, the need to avoid any energy from unwanted images at the output signal requires reducing the bandwidth of the waveform below the requirements of the sampling theorem. Most AWGs manufacturers specify $S_r/2.5$ (instead of $S_r/2$) as the maximum usable frequency for this specific reason. High-performance AWGs implement the "true-arb" architecture. In this architecture the sampling rate is variable and set by the user according to the waveform characteristics. This means that theoretically a different reconstruction filter should be required for each sampling rate. As this is unfeasible, manufacturers include just a few fixed filters or none at all. Users must, then, add their own filters according to their application and the target sampling rates required by them. AWG' embedded fixed filters are typically designed to obtain a good step response (with no ringing) and low group-delay distortion. Bessel-Thompson filters are much appreciated as they show both characteristics although their out-of-band attenuation performance is rather limited when compared to elliptic or Chebyshev filters of the same order. In many situations, the Sinc response of the DAC is not equalized by the reconstruction filter, if any. A solution to deal with fixed reconstruction filtering (or no filtering at all) is waveform compensation. In this scheme, the waveform is preprocessed using a correction digital filter. The response of that filter can be obtained through calibration of the response of the AWG. Typically, the correction filter boosts the higher frequency components to compensate for the DAC low-pass response. Boosting any frequency component will result in a higher amplitude waveform. As this higher amplitude waveform must fit in the available DAC dynamic range, this means that the amplitude of the corrected waveform must be reduced. This operation results in lower amplitude and in a reduction of the SNR (Signal-to-Noise Ratio).





Application Note

DAC Performance and Specifications

Even ideal DACs cannot perfectly recover the original signal from the data stored in the waveform memory due to fundamental limitations regarding bandwidth and vertical resolution. Real-world devices and the supporting circuitry attached to them (i.e. Sampling Clocks or amplifiers) are never ideal and they add distortions and imperfections of their own. In the previous section, bandwidth limitations and linear distortions were shown. The effects of noise sources and non-linear behavior will be shown in the current section. It is worth commenting that most of the DAC figures of merit and specifications as a device can be applied to AWGs as a system.

Quantization noise

Any DAC cannot generate any analog level as the available resolution is limited. For a unipolar DAC, the DAC's quantization level (or LSB, the difference between two consecutive levels) is $Q_{level} = Sr / (2^N - 1)$. Any analog level must be rounded to the available levels. Rounding will show up at the DAC's output as an error signal. This signal has some interesting properties:

- Generally speaking, the error signal will look as a random noise which amplitude is limited to $\pm 1/2$ LSB. Distribution will be constant over the full range so its rms (root-mean-square) value can be easily calculated as a function of the DAC resolution.
- Its power is also uniformly distributed in the frequency domain over the whole Nyquist band so the faster the DAC the lower noise power density will be obtained for a given vertical resolution.

Quantization noise (ϵ) power can be characterized by its rms value:

$$\epsilon_{RMS} = Q_{level} / 12^{1/2}$$

As quantization noise power is basically independent from the signal being generated, the Signal to Quantization Noise Ratio (SQNR) can be calculated for a given signal. For a sinusoidal signal $s(t)$ using the full range of a N-bit DAC, SQNR will be as follows:

$$SQNR(dB) = 20 \log_{10}(S_{RMS} / \epsilon_{RMS}) = 20 \log_{10}(2^N) + 20 \log_{10}((3/2)^{1/2})$$

$$SQNR(dB) = (6.02N + 1.76) dB$$

The above equation shows that adding a bit of resolution to the DAC will result in a 6dB improvement in the SQNR parameter.

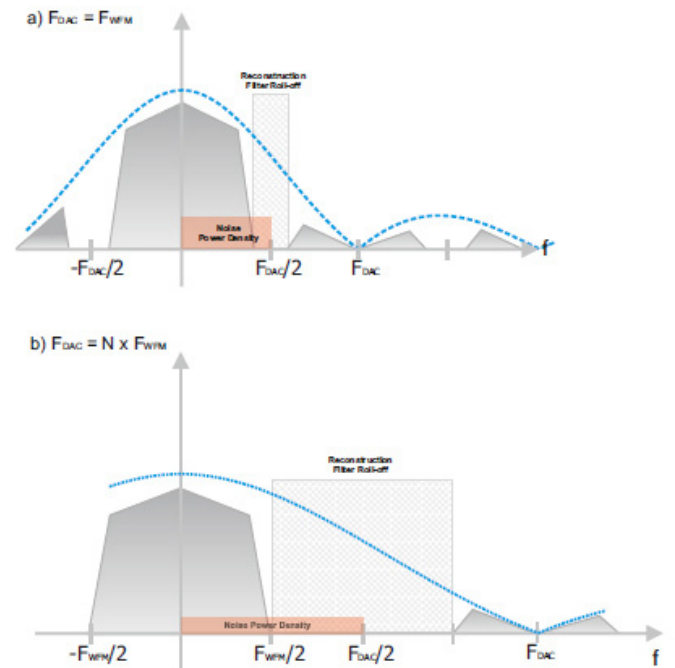




Figure 4. Oversampling DACs can generate extra samples by interpolation at a multiple of the original waveform sampling rate. This capability can improve signal quality by reducing the in-band noise while signal images are located farther away. In this way, less demanding, lower cost reconstruction filters may be used (b) while without having to sacrifice valuable bandwidth in the first Nyquist band (a).

An interesting development is that increasing the sample rate beyond the Nyquist requirements results in a reduction of the SQNR when it is measured within the band occupied by the signal instead of the full Nyquist band. As total power of quantization noise remains constant, the noise power density goes down by 3dB when sample rate doubles. Increasing sampling rate by a factor of 4 results in a 6dB improvement and, after applying the above formula, the improvement is equivalent to adding an extra bit of resolution to the DAC running at the original speed. This is the trick played by oversampling DACs, where digital-to-analog conversion takes place at a multiple of the waveform sampling rate. Additional intermediate samples are obtained within the DAC through some interpolation process. Oversampling improves the usability of the signal in another way: as effective sampling rate increases, unwanted images are located farther away and, as a happy outcome, they are easy to remove with a simpler filter. Additionally, as there is no need to provide for filter roll-off band within the waveform's Nyquist band, it can be fully occupied by the signal without suffering of aliasing (Figure 4).





Non-Linear Distortions

Although quantization noise is an inevitable consequence of the finite resolution of any DAC, its importance depends on two factors, the full-scale amplitude supported by the DAC and its operating sampling rate. As any real-world device will produce thermal noise at room temperature, once quantization noise power density is lower than that of thermal noise, the SNR improvement obtained by increasing the vertical resolution is not significant. Historically, quantization noise for commercial DACs running at tens or hundreds of MSa/s was much higher than thermal noise at the typical lab temperatures. Nowadays, nominal quantization noise for the Tektronix TDAC-25 10 bit DAC running at 25 GSa/s is comparable to thermal noise. Increasing vertical resolution in those devices is pointless as there is no performance to gain. Does this mean that a signal quality wall has been hit? The answer is no. There is an even more important factor in the DAC behavior affecting signal quality: linearity. Non-linear behavior produces harmonics and intermodulation products that show-up inside and outside the band occupied by the signal when analyzed in the frequency domain, thus degrading the signal to noise ratio.

Non-linearity can be categorized in two groups: static and dynamic. Static non-linearity can be analyzed by comparing the actual input code / output level transfer curve obtained in static conditions (constant input code) against a linear reference. As the slope and offset for the static response of any DAC can be easily corrected through linear networks, the linear reference must be obtained by statistical analysis of the actual transfer curve, typically by obtaining an optimal linear fit through linear regression analysis (least-squares method). The most important parameters (Figure 5) that can be extracted from the comparison between the actual curve and the reference are the following:






Figure 5. Real-world DACs show some static non-linear behavior. Here, the most important categories are shown.

- **Differential Non-Linearity (DNL):** It basically expresses the difference between any quantization level in the actual response and the nominal value for it. It provides a good idea on how the transfer function behaves for small variations of the signal amplitude. It is measured in LSB (a nominal quantization level) and it can be positive (when it is bigger than nominal) or negative (when it is smaller than nominal). When it takes the -1 value or lower for a given input code it means that the output decreases when the code increases and then there is a monotonicity error for that particular code. Given that a different curve would be obtained for each DAC, the common practice is to specify the worst-case DNL.
- **Integral Non-Linearity (INL):** It can be seen as the overall deviation of the actual transfer curve from the reference one. It can be expressed in either LSB units or as a percentage of full scale. Again, usually the worst-case deviation is specified for a given DAC part.





Application Note





Figure 6. Switching glitches largely contribute to dynamic non-linear behavior. Glitches are produced by non-perfect timing alignment (skew) between the different elements in the DAC (a). As skew grows the impact in the output signal also grows due to the combined effect of a longer duration and the distribution of the glitch power in the frequency domain (b). Low-noise, low-distortion signal generation requires excellent glitch performance.

Dynamic non-linearity can be observed at the DAC while running at its operating sample rate. There are many contributors but the following are the most significant ones:

- **Switching glitches:** These are spikes produced when the sample code changes due to timing skew between the different switching devices within the DAC (Figure 6). A more detailed explanation will be provided in the following section of this paper.
- **Settling Time / Slew Rate:** Time to transition for one level to another may depend on the initial and final levels. Typically, it will grow as the difference between the initial and final level grows. Rise and fall times associated to bandwidth limited signals do not depend on the voltage excursion at the output as they are caused by a linear distortion. Transition times will be a combination of both linear and non-linear distortions.

A very convenient way to observe the effects and the level of non-linearity in a DAC is generating a pure sinewave or a set of them (multi-tone). As this is the ultimate bandwidth limited signal, any non-linear component in the DAC response will show-up as unwanted components located at different frequencies. There may be other tones in the DAC output coming from the sampling clock (or its harmonics and sub-harmonics) or other external sources. It is relatively easy to tell as tones coming from the test sinewave and tones coming from the clock show-up at predictable frequencies. However, things may be more complex as harmonics beyond the Nyquist frequency ($Sr/2$) will be folded back to the first Nyquist band but, again, their expected frequencies can be easily computed.

Another source of noise is sample jitter. Sample jitter can be generated internally within the DAC or it may be caused by jitter in the sampling clock. It may also be observed as phase noise in CW or modulated carriers generated by the DAC.





SFDR and Effective Bits

It is difficult to fully characterize non-linear behavior and noise with a single specification. Typically, specifications are obtained by generating a single tone at a given frequency and measuring the power of any signal component out of the tone itself (Figure 7). There are several figures of merit that can be used to evaluate and compare DAC's non-linear behavior:

- SFDR: Spurious-Free Dynamic Range is, perhaps, the most popular single specification to assess non-linear behavior in a DAC (and, as a natural extension, in an AWG). This number basically reflects the ratio in dB between the levels of a single tone being generated by the DAC and the highest unwanted component within the Nyquist band. Typically, the test conditions used to obtain this figure must be included in the spec. These include information about tone frequency, sample rate, and the percentage of the DAC range being used. The spec may be also referred to the Full-Scale level even when the test tone is not covering the complete DAC range. In this case, unit for the spec is dBFS.
- THD: Total Harmonic Distortion measures the ratio between the addition of all the rms vales of harmonics and the rms value of the fundamental. It is typically specified as a percentage (%). As DACs are sampled systems, harmonics over the Nyquist frequency fold down to the Nyquist band at predictable frequencies so they can be easily differentiated from tones generated by non-harmonic sources such as sampling clock harmonics and sub-harmonics.

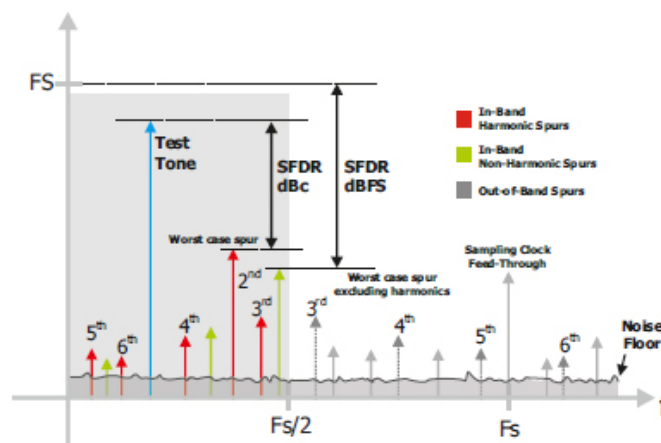




Figure 7. Dynamic non-linearity performance can be better established in the frequency domain. The basic technique consists in generating a single test tone and look to any signal component (spur) at a different frequency of that of the tone. Not all spurs are the same. Relevant spurs are located in the target Nyquist band (typically the first) and they can be generated by the signal itself in the form of harmonics or by other sources. Probably the most popular parameter is the Spurious-Free Dynamic Range (SFDR). It may be referred to the test tone amplitude (in dB units) or the DAC's full-scale (in dBFS units). Notice that harmonics beyond $Sr/2$ show-up as folded-down spurs in the first Nyquist band.

- SNR/SINAD: The Signal-to-Noise Ratio (SNR) shows the relative power level (in dB) between the test tone and all the other components within the Nyquist band excluded harmonics. The Signal-to-Noise-and-Distortion Ratio (SINAD) includes also the power of the harmonics.
- ENoB: The Effective Number of Bits is a different way to specify SINAD as it consists in computing the number of bits of an ideal DAC resulting in an equivalent level of quantization noise. Taking the $SQNR(dB) = (6.02N + 1.76)dB$ formula, ENoB can be easily computed:

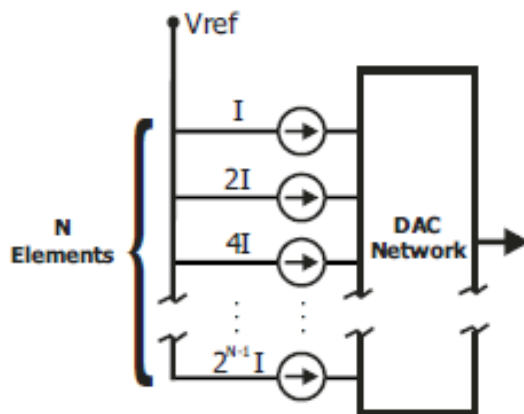
$$ENoB = (SINAD_{dB} - 1.76) / 6.02$$



a) DAC Basic Element



b) N-Bit Binary-Weighted Element DAC



c) N-Bit Unit Element DAC

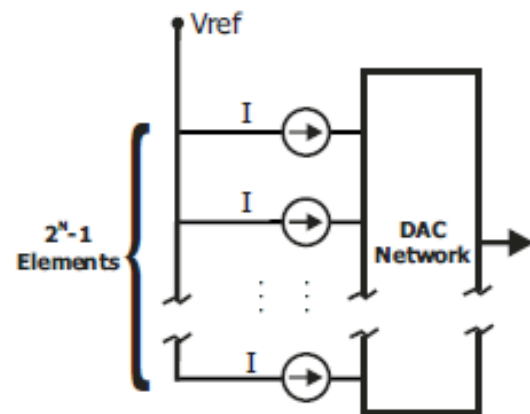


Figure 8. DACs may be built in many different ways. Typically, they are designed around basic building blocks called “elements”. Elements are made of some sort of current source and a data-controlled solid state switch (a). There are two basic ways to arrange elements to create an effective DAC. In binary-weighted elements N-bit DACs, N elements assigned to one bit each in the sample word are required. In unary-element N-bit DACs the sample word controls the number of identical elements being switched so $2N-1$ elements are required. Each type shows advantages and drawbacks.

High-speed DACs for AWGs

There are many ways to produce an analog level from a digital word representing that level. DACs are used in multiple applications and operating environments. Some may hold constant or near-constant values where extremely high accuracy and resolution are required. Others may be supplying a modulated carrier at multiple GHz frequencies where high sampling rate, low jitter and linearity are much more important than raw resolution. This paper concentrates in high-speed Digital-to-Analog Conversion so the most popular architectures implemented in such devices will be shown.

High-speed DACs are typically based on a series of current sources switched (each current source/switch set is known as “element”), directly or through some decoding process, by the input digital word representing the sample level (Figure 8a). The current is converted into a voltage value through some network. DACs can be divided primarily according to the way these current sources and switches are arranged in relation with the input digital word resulting Binary-Weighted (Figure 8b) element and Unary element DAC (Figure 8c) categories.

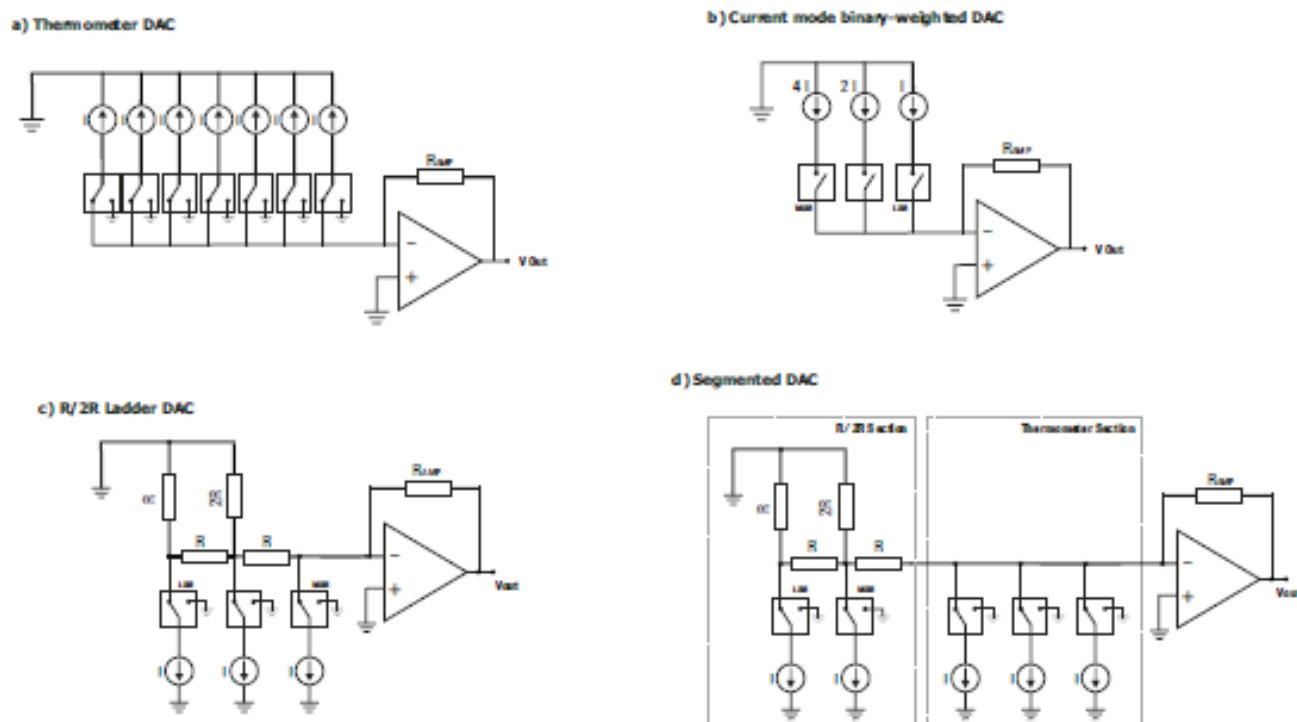




Figure 9. High speed DACs basic architectures include the thermometer DAC (a) and the current-mode binary weighted DAC (b). A very popular implementation of the current-mode binary weighted type is the R/2R ladder DAC (c) as it is based in identical current sources and similar resistor values, improving both accuracy and simplicity. High-speed DACs required in state-of-the-art AWGs typically use an hybrid architecture known as segmented DACs (d) where a binary weighted segment is associated to the LSBs and a thermometer section is connected to the MSBs. In this way a good linearity can be obtained while keeping the design simple. The Tektronix TDAC25 10-bit DAC use a segmented architecture where the 3 MSBs (7 elements) are attached to the thermometer section and the 7 LSBs are connected to the binary-weighted section (7 elements again). A pure thermometer DAC with the same resolution would require 1023 ($2^{10}-1$) elements.

Binary Weighted Element DACs

Probably the simplest way to implement a N-bit DAC converter is by using a series of N elements where each element contributes with a power of 2 corresponding to the associated bit position in the input digital word ($2^0 \dots 2^{N-1}$). Figure 9b shows a possible implementation of such arrangement. Although binary-weighted elements DACs minimize the number of current sources and switches, they also show important shortcomings. First of all, DAC linearity is very sensitive to inaccuracies in the current elements. The DAC response may even be non-monotonic (so sometimes a higher data word results in a lower output level). Another important practical issue is the difficulty to keep a good

accuracy in all the elements given the big disparity between them. A good solution to the latest problem is the popular R/2R ladder architecture (Figure 9c) often used in high-speed DACs. In this scheme, all the current sources have the same value while the associated network is made of two values of resistors: R and 2R. As a 2R resistor can be obtained as two R resistors in series (or one R resistor as two 2R resistors in parallel), the complete network can be made of resistors of the same value. Given that all current sources and resistors can be made with exactly the same characteristics on the same silicon die, higher accuracy and repeatability can be obtained. It is very difficult, though, to obtain good-enough linearity beyond 8 bits of resolution for high-speed DACs.





Application Note

Unary Element DACs

The alternative structure consists in a series of 2^N-1 identical elements where the number of elements activated depends on the input word. This type is also known as “Thermometer” DAC (Figure 9a) for obvious reasons. Linearity for this structure is much less sensitive to the individual element accuracy and monotonicity is assured as adding an additional element will always result in an additional output voltage. The fact that all the elements are identical simplifies obtaining a good accuracy and repeatability. The big disadvantage of such an arrangement is the unacceptable amount of elements required to implement DACs with sufficient vertical resolution for some applications. As an example, a 10-bit DAC would require 1023 elements (level 0 is obtained by switching off all the elements). An additional feature of unary elements DACs is the capability of improving linearity by randomizing the specific elements used to generate a given level. Given that all the elements are identical, what matters to reach a given level is the number of them and not the specific elements. Such operation results in different elements being activated in successive appearances of the same codeword. This results in a statistical distribution of levels which average is always the exact expected level. In fact, it is like using a perfectly linear DAC with some random noise added. The noise power is spread all over the spectrum while non-linearities result in distorted waveforms. For bandwidth limited signals (i.e. a CW carrier) non-linearity shows up as harmonics, IMD or shoulders in a modulated carrier. This strategy increases SFDR by converting non-linearity into random noise.


Segmented DACs

Both basic element arrangements have advantages and disadvantages. A good way to stress the advantages of each type while reducing the effects of their disadvantages is to combining both types in a single device. This hybrid structure is known as segmented DAC (Figure 9d) and it has become the most popular solution used in most high-speed DACs found in high-performance AWGs. Additionally, when properly designed, both sections (the binary weighted and the thermometer ones) can be implemented using the

same current sources, switches, and resistors improving the overall quality and performance. This is why the R/2R ladder architecture is so popular for the binary-weighted section. The thermometer section can still use the randomization procedure to further improve linearity.

A first design decision is how to split the available resolution bits between the two sections. The designers must reach a trade-off between complexity and performance. The more bits in the thermometer section, the better the linearity. The more bits in the binary-weighted section, the fewer elements will be necessary and the lower overall complexity. As an example, the TDAC-25 25GSa/s, 10 bits Digital-to-Analog Converter manufactured by Tektronix Component Solutions and incorporated into the AWG70000 series AWG, features a segmented architecture with 7 segments (the 3 MSBs) in the thermometer section and another 7 segments (the remaining 7 LSBs) in the binary weighted section.

An important component in every DAC element is the switch. Switches are designed so that instead of interrupting the current flow from the source, they transfer the current from the branch connected to the DAC network to another branch connected to ground. In this way, current always flows and there are not transients and loss of bandwidth caused by the sudden current change. Alternatively, the second branch can be connected to an identical conversion network. For any given code, one of the networks will produce the level corresponding to that code while the other network will generate the level corresponding to the complementary code. This complementary output may be used to obtain a pseudo-differential output. As both branches share the same current sources and switches, amplitude and timing will be very well matched. An interesting outcome of this arrangement is that linearity of the differential output will be improved as most power of the even harmonics (i.e. the 2nd harmonic) caused by distortion is removed. For some single-ended applications requiring improved SFDR performance it is better to use a balun to convert the differential output into single-ended if the signal to generate does not include DC or low-frequency components.



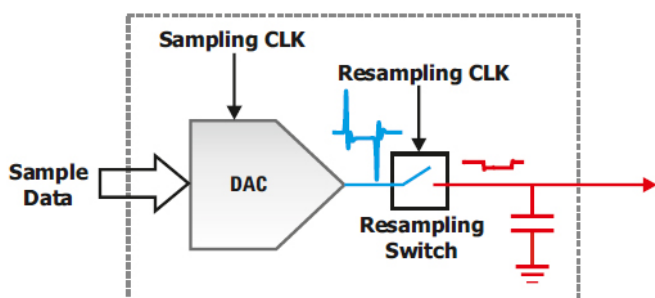



Figure 10. Resampling is a technique that can be applied to remove switching glitches in the output signal by disconnecting the DAC during the switching period. The problems with this method are that keeping the output voltage during the time the resampling switch is off is difficult or even impossible and the resampling switch linearity. Unlike switches in DAC elements which only handle two levels, resampling glitches must handle the complete DAC range so linearity is a big issue.

Switching glitches and resampling techniques


No matter the architecture of the DAC, the output must transition from the previous level to the next at some specific moments in time. Ideal DACs would activate all the switches in perfect synchronicity immediately after the new sample code reaches the device. Feeding a high-speed DAC with sample words at the required speeds is a challenging endeavor by itself and it deserves a separate discussion. Sample words can be stored in latches (even a FIFO) within the DAC so any timing or skew problems in the interconnections between the waveform memory and the DAC will not generate any problem in the DAC output. Internal skew between the interconnections in the parallel bus connecting the input latches (or decoding block) and the switches is not negligible in high-speed DACs as even tiny timing errors may take an important fraction of the sampling period (40 ps for a 25GSa/s DAC). This problem is even worse for segmented DACs as the thermometer and the binary weighted sections require very different pre-processing operations and show different bandwidth/settling time performances. While the binary-weighted section can directly use the MSBs, the thermometer section bits must go through a decoding process that may include randomization. Any timing difference in the operation of the switches will show up at the output as a switching glitch (Figure 6). Switching glitches may look like ringing or inter-symbol interference (ISI) but this is not the case. While the ringing and ISI are

linear distortions and their effects are proportional to the level change, switching glitch is an extremely non-linear phenomenon which is related to the switches being switched and not to the level change. The switching glitch may be very small for a high change in level if the involved switches are very well aligned while a tiny change in level can cause a high-energy glitch if the involved switches are not well aligned. Switching glitches may be one of the most important sources of noise and non-linearity in DACs and, as a consequence, in AWGs.

A traditional solution to switching glitches is the addition of a resampling switch at the DAC's output (Figure 10). The resampling switch disconnects the output from the DAC during the period the transient glitches show-up. Although this technique greatly reduces the effects of the switching glitches, the resampling glitch also introduces problems of its own:

- The DAC is disconnected from the output for a fraction of the sampling period. This leads to a RZ (Return-to-Zero) operation so amplitude is reduced. More complex architectures use dual core DACs (two synchronized devices) to generate the output level during half the sampling period but complexity increases.
- The resampling switch itself is a non-linear device forced to handle the full DAC dynamic range, unlike the switches in the elements within the DAC which handle only two levels (ON/OFF). This may be an important source of non-linear behavior and make any DAC unusable when very high dynamic range signal generation is necessary in applications such as RF/Wireless test. More complex resampling schemes such as distributed resampling may partially solve this problem although, again, it introduces new problems such as timing issues between the distributed resampling switches and bandwidth reduction.

Tektronix Components Solutions has followed a completely different path to solve the switching glitch issue in the TDAC-25. Instead of hiding or making-up the effects of the switching glitches through re-sampling, the development team has been very careful in the design to obtain an excellent time alignment within the DAC chip. In this way, the energy of the glitches is so tiny and spread over such a large frequency range, than their effects are negligible (Figure 6a and Figure 6b).



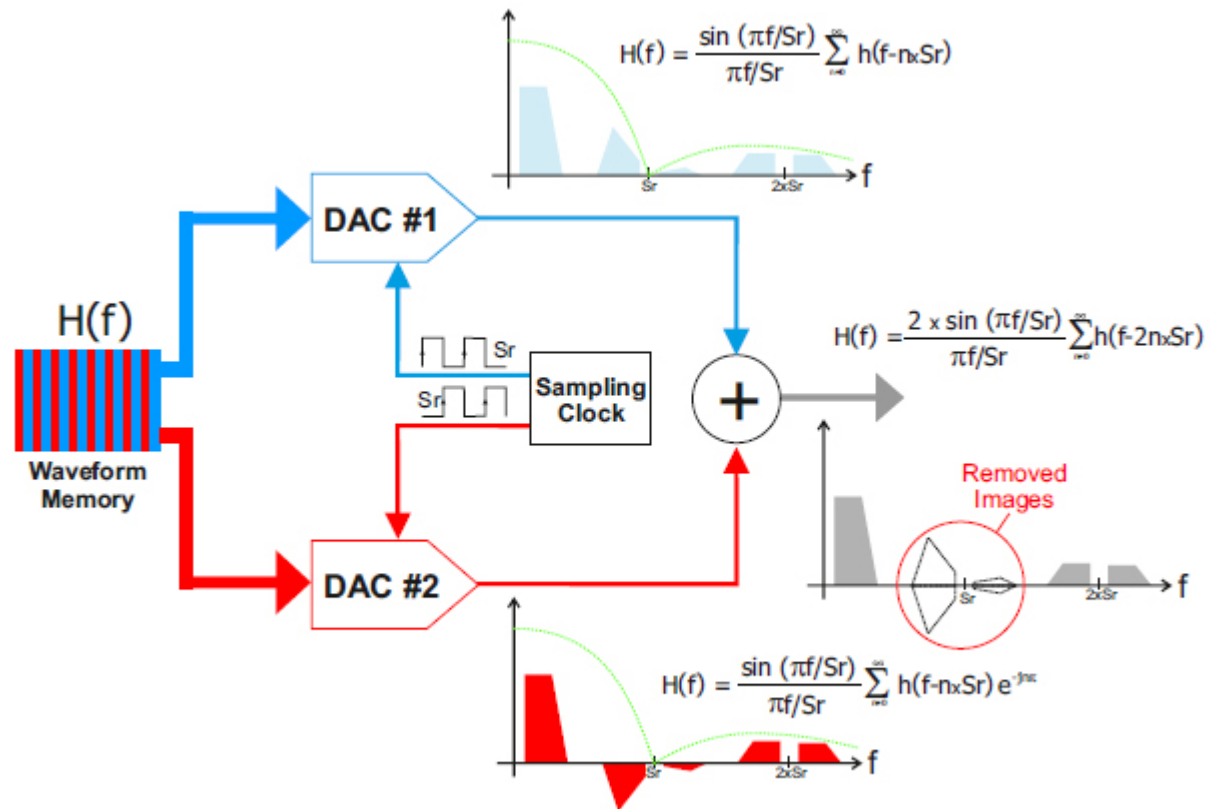




Figure 11. Tektronix implements the interleaving-DAC architecture in the AWG70001 so an effective 50GSa/s sampling rate can be produced out of two 25GSa/s DACs. Ideal interleaving DACs should use a switch to interleave both DACs. However, as it is not possible to produce switches at these speeds with sufficient quality to keep the signal clean, an alternative method based in the addition of signals from both DACs (after applying a differential delay of half the sampling period) is actually implemented in the Tektronix AWGs.


Extended Bandwidth DAC Architectures

Interleaving DACs

Real-Time digital storage oscilloscopes have been using for years a multiple ADC architecture where the sampling process for all of them is interleaved to increase the overall sampling rate. The trick is that each ADC sees the instantaneous level of a signal through a Sample&Hold device timed in time sequence that makes sure equally spaced samples are taken at S_r speed while each ADC works at S_r/N , where N is the number of ADC. This process is equivalent to the usage of a 1 input/ N output switch connecting the signal to each ADC

at $1/S_r$ intervals. Although all the sampling limitations apply to each ADC individually, the addition of a non-linear device (the 1: N switch) results in a sampling rate and bandwidth gain equivalent to a faster ADC converter. The same idea could be applied to DACs in order to obtain a higher equivalent conversion rate out of N lower-speed DACs combined with a $N:1$ switch. This is the interleaving-DAC architecture. The real problem for DACs is that there is no replacement for an actual $N:1$ switch as it happens in ADCs with properly timed S&H circuits connected in parallel to the input signal. At the sampling rates required by high-speed AWGs, building such a switch, if feasible, would end up in an unacceptable level of noise and switching glitches.





Tektronix took a different path to handle the interleaving-DAC switch issue. In the Tektronix implementation (currently implemented in the AWG70000 series), two high-speed DACs take the odd and even samples of a signal and their outputs are added after skewing one of them by half the sampling period (Figure 11). The output of both DACs consists in the same signal sampled at different moments in time. If those signals are analyzed in the frequency domain the following expressions are obtained:

$$H1(f) = \frac{\sin\left(\frac{\pi f}{Sr}\right)}{\pi f/Sr} \times \sum_{n=0}^{\infty} h(f - n \times Sr)$$

$$H2(f) = \frac{\sin\left(\frac{\pi f}{Sr}\right)}{\pi f/Sr} \times \sum_{n=0}^{\infty} h(f - n \times Sr) \times e^{-jn\pi}$$

When both signals are added together, this is the combined output:

$$H(f) = \frac{\sin\left(\frac{\pi f}{Sr}\right)}{\pi f/Sr} \times \sum_{n=0}^{\infty} h(f - 2 \times n \times Sr)$$


Although this expression may look quite complex, the actual facts around it are quite easy to visualize:


- The $\text{Sinc}(\pi f/Sr)$ is still there so there is no bandwidth gain in this case as the addition operation is purely linear.
- Odd images (including the one in the first Nyquist band) are the same and they have the same phase.
- Even images for each DAC are the same but with opposite phase. As a result they cancel each other in the combined signal.
- The two previous facts result in an extension of the Nyquist bandwidth to Sr , equivalent to a DAC running at twice the speed ($2 \times Sr$).

Signals now can be synthesized from DC up to frequencies well over $Sr/2$ (the Nyquist band for each individual DAC) although the Sr theoretical limit cannot be reached due to the zeroth-order hold response of the individual DACs. Typically, usable bandwidth may extend to 80% of the Sr frequency, instead of the usual 80% of $Sr/2$ for an isolated DAC.

The interleaving-DAC architecture, though, is not free of challenges. Any imbalance between the two participating DACs in terms of frequency response (both amplitude and phase) as well as any inaccuracy in the relative delay ($1/2Sr$) will result in an incomplete odd image cancellation. The problem, in this case, is that partially cancelled images are located in the new Nyquist bands so they show up as an in-band interference that cannot be fully eliminated by filtering. The AWG70000 series incorporates differential gain and fine skew controls so almost perfect DAC-to-DAC balance can be accomplished over a wide frequency range.

Another problem of the interleaving DAC architecture is the relative importance of the zeroth-order hold response. The additional attenuation at high frequencies (especially beyond $Sr/2$) can be compensated through a combination of mechanisms. One of them is to use an adequate reconstruction filter that boosts the high frequencies. The problem with this method is that it also boosts noise (including quantization noise) at those frequencies. Another method is to design the DAC to directly generate pulses with the right frequency contents, the same de-emphasis methodology applies to high-speed serial links. Tektronix applies both methods in the AWG70000 series AWGs so both good flatness and SNR characteristics can be obtained simultaneously.





Application Note

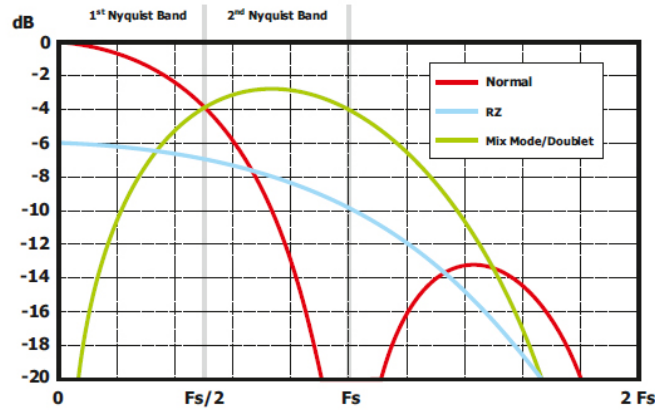


Figure 12. Mix-Mode (doublet mode) and RZ-DACs can be used to improve signal quality for generation in the second (or even the third) Nyquist Band. Here the relative frequency response of normal, RZ, and Mix-Mode DACs running at the same sample rate can be seen. Notice that Mix-Mode DACs cannot generate signals in the first Nyquist band.

Return-to Zero (RZ) DACs

Some DACs can generate return-to-zero samples where the output is forced to zero during half the sampling period. The narrower isolated pulse results in a $\text{Sinc}(2\pi f/Sr)$ response so the first null is located at $2 \times Sr$. Although the amplitude of the resulting signal will be halved (6dB lower at DC), the response in the first Nyquist band is much flatter than a regular NRZ (Non-Return-to-Zero) DAC and their response in the second Nyquist band is significantly higher. In fact, this method is useful only to generate signals in the second Nyquist band.

Mix-Mode (Doublet Mode) DACs

Dual-core DACs use two different RZ DACs generating the same nominal level during half the sampling period each in order to eliminate switching glitches. Some dual-core DACs allow for the generation of two different levels during one single sampling period based in the same sample value.

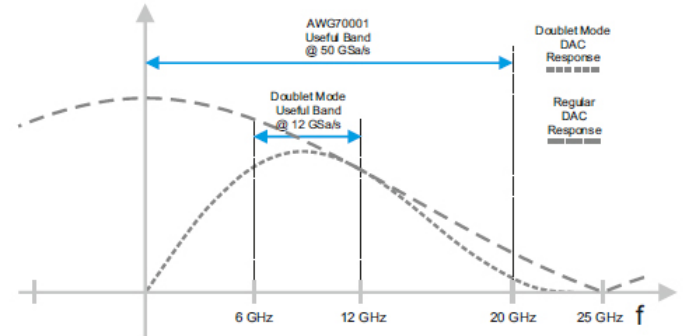




Figure 13. The frequency responses of a 50GSa/s interleaved DAC and a 12GSa/s Mix-Mode DAC can be seen in this graph. Notice the big difference in terms of signal bandwidth and frequency coverage. Tektronix has tailored the pulse shape in the AWG70000 series to boost high frequency components and improve flatness.

A useful generation method consists in the application of the sample code during the first half of the sampling period and the application of the complementary code during the second half. This strategy is called “Mix-Mode” or “Doublet-Mode” depending on the manufacturer. The advantage of this technique is that images located in the second Nyquist band are boosted while the signal located in the first Nyquist band is attenuated. This mode is ideal to extend the frequency range of the DAC but only signals with frequency components between $Sr/2$ and Sr can be generated. Figure 12 shows the relative frequency response of regular (NRZ), Return-to-Zero (RZ) and “Mix-Mode” DACs for devices with the same Full-Scale amplitude specification. Figure 13 shows the difference between the response of a mix-mode DAC running at 12GSa/s and an Interleaving DAC running at $2 \times 25\text{GSa/s}$. All these responses do not include the influence of the analog frequency response of the DAC.







Conclusions

With the growing applications associated with Digital to Analog Conversion, it is critical that users of such devices understand the characteristics and their uses. In this note we tried to explore the most important topics and hopefully after reading this you are more familiar with the DACs used in AWGs today. Here are some of the conclusions that you should have taken away after reading this document.

- DACs are basically the opposite of ADCs
- Applications generally only use only the 1st Nyquist zone
- Flat frequency response in 1st Nyquist zone is preferable
- Frequency response in upper Nyquist zones is not as important
- DAC linearity is as important as or more important than resolution
- Spurious Free Dynamic Range is one of the more important specifications regarding DAC performance
- Interleaving mode allows for higher sample rates
- Increasing the sample rate beyond the Nyquist requirements results in a reduction of the SQNR when it is measured within the band occupied by the signal instead of the full Nyquist band
- DAC systems usually need filtering to generate improved waveforms
- Amplifiers to compensate for the DAC $\text{Sin}(x)/x$ response is difficult
 - Compensation of both frequency response and requirement for constant group delay across the bandwidth is difficult with analog circuits
 - Bandwidth requirements for distortion minimization are significant
- $\text{Sin } x/x$ compensation for "true-arb" AWGs is unpractical because a unique compensation filter would be required for each sampling frequency.





Contact Tektronix:

- ASEAN / Australia (65) 6356 3900
- Austria* 00800 2255 4835
- Balkans, Israel, South Africa and other ISE Countries +41 52 675 3777
- Belgium* 00800 2255 4835
- Brazil +55 (11) 3759 7627
- Canada 1 (800) 833-9200
- Central East Europe and the Baltics +41 52 675 3777
- Central Europe & Greece +41 52 675 3777
- Denmark +45 80 88 1401
- Finland +41 52 675 3777
- France* 00800 2255 4835
- Germany* 00800 2255 4835
- Hong Kong 400-820-5835
- Ireland* 00800 2255 4835
- India +91-80-30792600
- Italy* 00800 2255 4835
- Japan 0120-441-046
- Luxembourg +41 52 675 3777
- Macau 400-820-5835
- Mongolia 400-820-5835
- Mexico, Central/South America & Caribbean 52 (55) 56 04 50 90
- Middle East, Asia and North Africa +41 52 675 3777
- The Netherlands* 00800 2255 4835
- Norway 800 16098
- People's Republic of China 400-820-5835
- Poland +41 52 675 3777
- Portugal 80 08 12370
- Puerto Rico 1 (800) 833-9200
- Republic of Korea +822-6917-5000
- Russia +7 495 664 75 64
- Singapore +65 6356-3900
- South Africa +27 11 206 8360
- Spain* 00800 2255 4835
- Sweden* 00800 2255 4835
- Switzerland* 00800 2255 4835
- Taiwan 886-2-2656-6688
- United Kingdom* 00800 2255 4835
- USA 1 (800) 833-9200

* If the European phone number above is not accessible,
please call +41 52 675 3777

Contact List Updated June 2013

For Further Information

Tektronix maintains a comprehensive, constantly expanding collection of application notes, technical briefs and other resources to help engineers working on the cutting edge of technology. Please visit www.tektronix.com



Copyright © 2014, Tektronix. All rights reserved. Tektronix products are covered by U.S. and foreign patents, issued and pending. Information in this publication supersedes that in all previously published material. Specification and price change privileges reserved. TEKTRONIX and TEK are registered trademarks of Tektronix, Inc. All other trade names referenced are the service marks, trademarks or registered trademarks of their respective companies.

