ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

An Almost Pure DDS Sine Wave Tone Generator

Patrick Butler, Field Applications Engineer

Introduction

The test and verification of ac performance of high precision fast analog-to-digital converters (ADCs) with resolution better than 16 bits require a near perfect sine wave generator capable of covering a 0 kHz to 20 kHz audio bandwidth at least. Usually, expensive laboratory instruments are used to perform these evaluations and characterizations such as the audio analyzer AP27xx or APx5xx series from Audio Precision. Most of the time, modern high speed SAR and wideband sigma-delta (Σ-Δ) ADCs exhibiting 24 bits or more feature single-supply and full-differential inputs, and therefore require the signal source used for the DUT to be dc and ac accurate, while providing full differential outputs (180° out of phase). Similarly, the noise and distortion level of this ac generator should be much better than the specifications of these ADCs, resulting in a noise floor level well below −140 dBc and distortion lower than −120 dBc with an input tone frequency of 1 kHz or 2 kHz and up to 20 kHz according to most supplier specifications. A typical configuration of a typical bench test setup suited for high resolution wideband ADCs is illustrated on the Figure 1. The most critical component is the sine wave generator (single or multitone) and here a software-based direct digital synthesizer (DDS) can provide full flexibility with extremely fine frequency resolution and clock synchronization with the data acquisition system to perform coherent sampling to avoid leakage and FFT window filtering.

At a fraction of the cost of an audio precision analyzer, it is possible to design a very accurate sine wave generator based on the direct digital frequency synthesis (DDFS) principle, but implemented in software onto a floating-point DSP processor such as the SHARC® processor. A reasonably fast floating-point DSP will meet real-time expectations and fulfill all the arithmetic and processing conditions to achieve the distortion and noise performance level set up by the most advanced SAR ADCs. Taking the benefit of the full-word data length of the SHARC core architecture either in 32-bit or 64-bit fixed-point format for the NCO phase accumulation and the proprietary, 40-bit, floating-point, extended precision to execute the sine approximation function and the digital filters used to shape the spectrum, the quantization effects (rounding and truncation noise) are drastically reduced to be considered negligible compared to the digital-to-analog converter (DAC) imperfections used for the signal reconstruction.

Direct Digital Frequency Synthesis

The digital signal generator synthesizer patent filed in April 1970 by Joseph A. Webb[1] described what could be considered as the basis of DDS mechanics to generate various types of analog waveforms, including sine waves, simply with the use of a few digital logic modules. Then, in early 1971, the frequently cited reference paper from Tierney et al.[2] on direct digital frequency generation by deepening the DDS operation for quadrature generation as well as its limitations (word truncations and frequency planning) regarding the sampled systems theory was published. Practical realizations began to show up, mostly relying on discrete standard logic ICs such as the TTL 74xx or ECL 10K families. Less than 10 years later, fully integrated solutions came on the market introduced by companies like Stanford Telecom, Qualcomm, Plessey, and Analog Devices with the AD9950 and the AD9955. Designed for the best speed, power, and cost trade-off, the logic ICs' architectures were based on a lookup table (LUT) to ensure the phase-to-sine amplitude conversion with limited phase, frequency, and amplitude resolutions. Today, Analog Devices remains the largest and, perhaps, most unique supplier of DDS standalone integrated circuits, while current numerically controlled oscillators (NCOs) tend to be integrated in numbers in RF DACs like the AD9164 or the AD9174. Despite their impressive noise and linearity performance over a multiple GHz bandwidth, none of these devices are appropriate for the test of moderate speed, high resolution ADCs such as the LTC2378-20, the AD4020, or the AD7768.
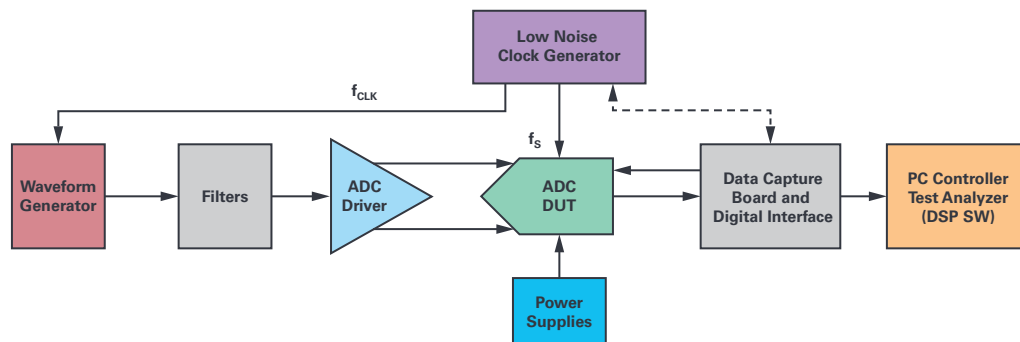

Figure 1. Processing chain of a typical ADC (ac) test setup based upon the IEEE 1241 standard. The DDFS makes the whole measurement system fully digital with a lot of benefits including full flexibility and coherent sampling acquisition.

Compared to traditional PLL-based synthesizers, NCOs and DDSs are mostly known for their very fine frequency resolution, fast agility, and ease of sine/cosine generation with perfect quadrature. They are also prized for their wide bandwidth coverage and dc accuracy. Their principle of operation is governed by digital signal processing and sampling systems theory, and their digital nature allows for fully digital and independent control of the phase, frequency, and amplitude of the output signals. The block diagram of Figure 2 depicts the architecture of a conventional DDS, which consists of three major functions:

▶ An N-bit phase accumulator;
▶ A phase-to-sine amplitude converter characterized by a W-bit truncated phase input word;
▶ A D-bit DAC and its associated reconstruction filter.

The phase accumulator is built around a simple N-bit adder combined with a register whose content is updated at the rate of the sampling clock F_{CLK} with the input phase increment $\Delta\theta$, also commonly called the frequency tuning word (FTW). The accumulator can periodically overflow and operates like a fractional divider between the sampling or reference clock F_{CLK} and the DDS output frequency F_{OUT}, or like a gearbox with a divide ratio equal to:

$$\frac{FTW}{2^N}$$

The overflow rate gives the output frequency of the generated waveform such that:

$$F_{OUT} = \frac{FTW \times F_{CLK}}{2^N}$$

where $0 \leq FTW \leq 2^{N-1}$. Because of the divider effect, the contribution of the reference or sampling f_S clock phase noise at the NCO output will be reduced by

$$-20 \times log_{10}\left(\frac{F_{CLK}}{F_{OUT}}\right)$$

The output of the phase accumulator register represents the current phase of the generated waveform. Each discrete accumulator output phase value is then translated into an amplitude sine or cosine data or sample thanks to the phase-to-sine or phase-to-cosine mapper engine. This function is usually accomplished by means of trigonometric values stored in a LUT (ROM) and sometimes by the execution of a sine approximation algorithm or a combination of the two. The output of the phase-to-sine amplitude converter feeds a DAC, which produces a quantized and sampled sinusoid before being filtered to smooth the signal and avoid spectrum aliasing. This amplitude quantization imposed by the DAC finite resolution puts a theoretical limit on the noise floor and the resulting signal-to-noise ratio (SNR) of the

synthesizer. Moreover, as a mixed-signal device, the DAC exhibits a whole bunch of dc and ac nonlinearities due to its INL, DNL, slew rate, glitches, and settling time characteristics, which create spurious tones and reduce the overall dynamic range of the sine wave generator.

Practical sine waveform generator implementations based on the architecture of Figure 2 differ mostly by the phase-to-amplitude converter block, which is generally optimized for speed and power consumption rather than high precision because of the market orientation for digital radio applications. The simplest approach for the realization of the phase-to-sine amplitude converter is to use a ROM to store sine values with one-to-one mapping. Unfortunately, the length of the LUT grows exponentially (2^N), with the width N of the phase accumulator and linearly with the wavetable data word precision W. Unfortunately, trade-offs consisting in the reduction of the accumulator size or truncating its output result in the loss of frequency resolution and a severe degradation of the SFDR. It is shown that spurs caused by phase or amplitude quantization follow a –6 dB/bit relationship. Since a large N is normally desired to achieve a fine frequency tuning, several techniques have been promoted to limit the ROM size while maintaining adequate spur performance. Simple compression methods are commonly used by exploiting the quarter wave symmetry of the sine or cosine function to reduce the phase argument range by 4. For further range reduction, brutal truncation of the phase accumulator output is the de facto method, although it does introduce spurious harmonics. Despite that, this approach is always adopted because of the fine frequency resolution requirements, memory size, and cost compromise. Various angular decomposition methods have been suggested to lower the memory requirements with LUT-based methods. Combined with amplitude compression using various types of segmentation, linear, or polynomial interpolation, the idea is to accurately approximate the first quadrant of the sine function or over the [0, π/4] interval in the case of I/Q synthesis for which both sine and cosine functions are needed. Similarly, complex signal generation with no ROM LUT is efficiently supported by angle rotation-based algorithms just calling for shift and add operations in a successive approximations scheme. This method, represented by the popular CORDIC, is generally faster than other approaches when a hardware multiplier is not available or when the number of gates required to implement the functions should be minimized (in an FPGA or an ASIC) for speed or cost considerations. Conversely, when a hardware multiplier is available—as is always the case in a DSP microprocessor— table-lookup with interpolation methods and full polynomial calculations, such as Taylor-series expansion, Chebyshev polynomials are faster than CORDIC, especially when high accuracy is a must.
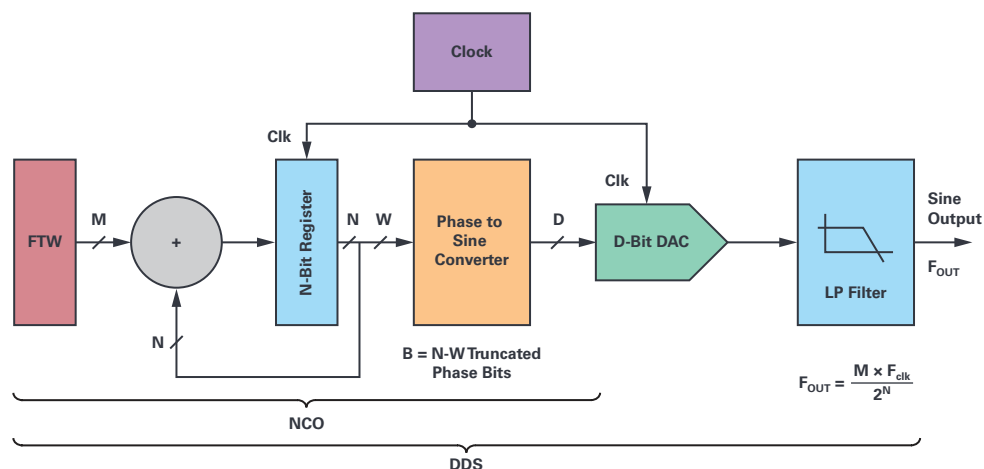

Figure 2. Main functional sections of an NCO and distinction with the complete direct digital synthesizer, which includes the reconstruction DAC and its associated AAF. The NCO section can be used to test or stimulate DACs.

Implementing a High Precision NCO in Software

Building a high precision ac tone generator with similar or better distortion performance than the best analog oscillators, as in the most famous Hewlett-Packard analyzers or as described in the application note AN-132[3] is not a trivial thing, even if dedicated to the audio frequency spectrum (dc to 20 kHz range). Nevertheless, as written previously, a full software implementation, performing the phase calculations (ωt) and sine function ($\sin(\omega t)$) approximations using the adequate arithmetic precision of an embedded processor can certainly help to minimize the quantization side effects, noise, and resulting spurs. This means that all the NCO functional blocks of Figure 2 are translated in lines of code (no VHDL!) to realize a software version that will meet real-time constraints to ensure the minimum sampling rate and the desired frequency bandwidth.

For the phase-to-sine amplitude conversion engine, the full LUT scheme or any variation demands too much memory or too many interpolation operations to achieve a perfect sine conformity. On the contrary, the polynomial method for sine approximation offers a very good complexity vs. accuracy trade-off by allowing the use of a very low cost, general-purpose DSP. Polynomial series expansion is also very attractive for its relative simplicity and ability to provide full flexibility in the choice of the type of power series, in tailoring the algorithm for a given precision. It does not require a large memory space, less than 100 lines of SHARC DSP assembly lines, and just a few RAM locations to store the polynomial coefficients and variables as sine values are only computed at sampling time instants.

At first, the obvious choice for a sine approximation function would be to use a straight Taylor/MacLaurin power series with the appropriate order to meet the targeted accuracy. However, since power series tend to lose effectiveness at endpoints, it is mandatory to reduce the argument input range to a smaller interval before performing any polynomial evaluation. Without argument range reduction, high precision over the function domain such as $[-\pi, +\pi]$ can only be supported with very high order polynomials. Thus, some transformations need to be applied to the elementary function to get the reduced argument such as $\sin(|x|) = \sin(f + k \times \pi/2)$ and $\sin(f) = \sin(x - k \times \pi/2)$ with $0 \leq f < \pi/2$. Consequently, extreme care should be taken with the trigonometric functions to avoid subtraction cancellations, which would lead to a serious loss of precision and produce catastrophic results, particularly with a poor arithmetic precision. In our case, this might occur when the phase input is large or close to an integer multiple of $\pi/2$.

Besides the periodicity and modulo-2π repetitions, the symmetric properties of the $\sin(x)$ function can be applied to further reduce the range of approximation. Given the fact that the sine function is antisymmetric about the point $x = \pi$ for the interval $[0, 2\pi]$, so it is possible to use the following relationship:

$$\sin(x + \pi) = -\sin(x)$$

to reduce the range to $[0, \pi]$. In the same manner, $\sin(x)$ shows a symmetry about the line defined by $x = \pi/2$ for the interval $[0, \pi]$, such that:

$$\sin(x + \pi/2) = \sin(\pi/2 - x)$$

for x in the interval $[0, \pi/2]$, which reduces the angle input approximation range even more. Further argument reductions to smaller intervals like $[0, \pi/4]$ to improve the accuracy is not efficient because it requires both the evaluation of the sine and cosine functions at the same time as dictated by the common trigonometric relationship: $\sin(a+b) = \sin(a) \times \cos(b) + \cos(a) \times \sin(b)$, worthwhile for the generation of quadrature tones.

Analog Devices' ADSP-21000 Family Application Handbook Volume 1 describes an almost ideal (for embedded systems) sine approximation function based on an optimized power series written for the first ADI DSP floating-point processor, namely the ADSP-21020, which is basically a SHARC core. This implementation of $\sin(x)$ relies on a minimax polynomial approximation that was published by Hart et al.[4] and refined by Cody and Waite[5] for floating-point arithmetic

to mitigate round-off errors and to avoid the occurrence of cancellations as previously mentioned. The minimax method relies on Chebyshev polynomials and the Remez exchange algorithm to determine the coefficients for a desired maximum relative error. As shown with MATLAB® in Figure 3, small changes in the set coefficients result in a dramatic increase in accuracy for minimax compared to Taylor for a seventh-order Taylor polynomial.[6] For the best accuracy vs. speed trade-off, the angle input range of this sine approximation function is shrunk to the $[-\pi/2$ to $+\pi/2]$ interval and the software routine includes an efficient range-reduction filter, which counts for about 30% of the total "sine" subroutine execution time.
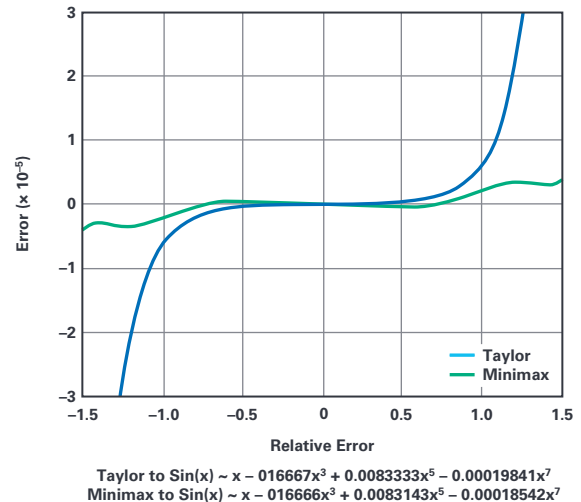

Taylor to Sin(x) ~ x – 016667x^3 + 0.0083333x^5 – 0.00019841x^7
Minimax to Sin(x) ~ x – 016666x^3 + 0.0083143x^5 – 0.00018542x^7

Figure 3. Unlike the Taylor-MacLaurin method defined around 0, the minimax sine approximation approach minimizes and equalizes the maximum relative error over the $[-\pi/2$ to $+\pi/2]$ interval.

While all the computations could be executed with 32-bit fixed-point arithmetic, the most common and convenient format for mathematical calculations especially when dealing with long numbers has been for years the IEEE 754 floating-point standard. As a DSP VLSI chip manufacturer, Analog Devices pioneered the IEEE 754-1985 standard from the very beginning. At the time, there was no single-chip floating-point DSP processor at all, but only simple floating-point multiplier and ALU computation ICs such as the ADSP-3212 and the ADSP-3222, respectively. This format replaced most of the proprietary formats of the computer industry and became the native format for all the SHARC DSP processors, in single precision 32-bits, extended precision 40-bits, and recently, double precision 64-bits for the ADSP-SC589 and ADSP-SC573.

The SHARC 40-bit extended single precision floating-point format with its 32-bit mantissa provides enough precision (u 2^{-32}) for this sine wave generation application and to keep things equal, Cody and Waite show that a 15th order polynomial is appropriate for an overall accuracy of 32 bits with an evenly distributed error over the $[0$ to $+\pi/2]$ input domain. The final tweak to minimize the number of operations and maintain accuracy is to implement the Horner's rule for the polynomial calculation, a fast exponentiation method to evaluate a polynomial in one point, such that:

$$y(x) = \sin(x) \simeq ((((((((R7 \times x^2 + R6) \times x^2 + R5) \times x^2) + R4)$$
$$\times x^2 + R3) \times x^2 + R2) \times x^2 + R1) \times x^2 + 1) \times x,$$
$$\text{with } x \in [0, +\pi/2].$$

R1 to R7 are the Cody and Waite coefficients of the polynomial series and only eight multiplies and seven additions are necessary to evaluate the sine function for any input argument $\epsilon[0, \pi/2]$. The complete $\sin(x)$ approximation code written in the form of an assembly subroutine is executed in about 22 core cycles on a SHARC processor. The original assembly subroutine was modified to perform simultaneous double memory accesses when fetching the 40-bit polynomial floating-point coefficients to save six cycles.
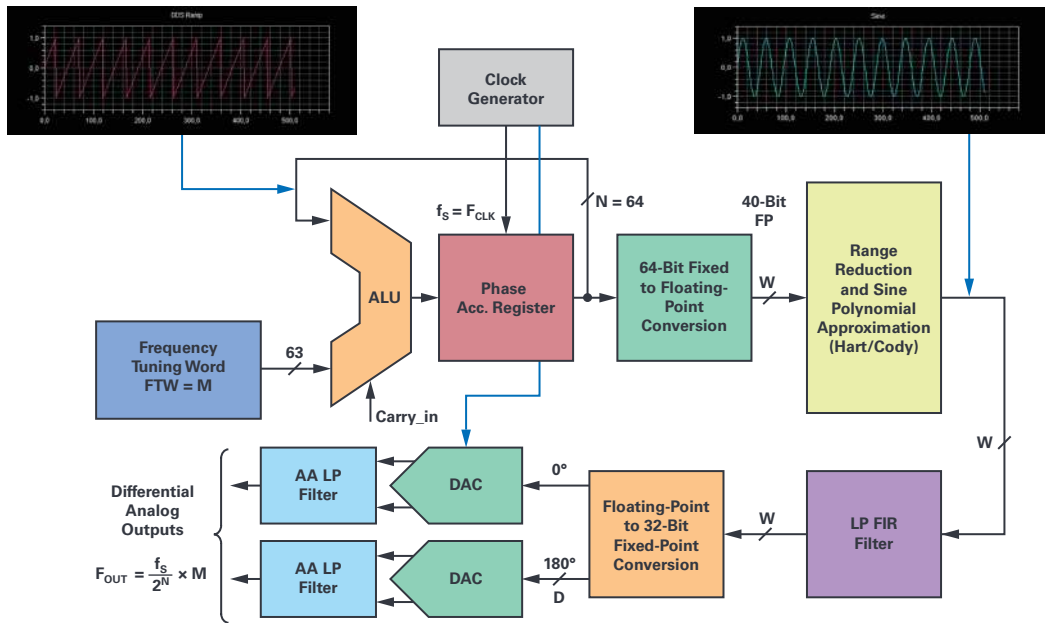
Figure 4. The software DDS simplified block diagram gives the data arithmetic formats and locations of the various quantization steps between the processing elements.

The NCO 64-bit phase accumulator itself is making use of the SHARC 32-bit ALU in double precision two's complement fractional format for its execution. A complete phase accumulator execution with memory update costs 11 core cycles, and as a result, every NCO output sample is generated in about 33 core cycles.

The diagram in Figure 4 shows the functional block implementation of the software DSP-based NCO with some reference to the arithmetic format precision at each stage. In addition, one or two DACs and their analog anti-aliasing filter circuitry are required for the signal analog reconstruction, and to realize the complete DDFS. The key elements of the processing chain are:

▸ the 64-bit phase accumulator (SHARC ALU double precision addition with overflow);

▸ the 64-bit fractional fixed-point to 40-bit FP conversion block;

▸ the range reduction block [0 to $+\pi/2$] and quadrant selection (Cody and Waite);

▸ the sine approximation algorithm (Hart) for the phase-to-amplitude conversion;

▸ the sin(x) reconstruction and normalization stage over the −1.0 to +1.0 range;

▸ the LP FIR filter and sin(x)/x compensation if necessary;

▸ and the 40-bit FP to D-bit fixed-point conversion and scaling function to fit with the DAC digital input.

An optional, digital low-pass filter can be placed at the output of the NCO to remove any spur and noise that could fold in the band of interest. Optionally, this filter can provide interpolation and/or inverse sin(x)/x frequency response compensation depending upon the DAC selected for the analog reconstruction. Such a low-pass FIR filter could be designed with the MATLAB Filter Designer tool. As an example, assuming a 48 kSPS sampling frequency and a dc to 20 kHz bandwidth with a 0.0001 dB in-band ripple and a −150 dB out of band attenuation, a high quality equiripple filter could be implemented with 40-bit floating-point coefficients. With only 99 filter coefficients, its total execution time will consume about 120 SHARC core cycles in single instruction, single data (SISD) single-computation unit mode. After digital filtering, the pairs of calculated samples are sent by DMA to the DACs using one of the DSP synchronous serial ports. For a better speed performance, chaining

DMA operation is also possible with large ping-pong memory buffers to support processing by block operation. For example, the block data size could be equal to the length of the FIR data delay line.

Final Tweaks at the NCO for an Optimal SFDR

As mentioned earlier, the NCO suffers from spurs mainly due to the truncation of the phase accumulator output and, to a lesser extent, from the amplitude quantization done on the sinusoidal values obtained by calculation or by tabulation. The error due to phase truncation generates spurs around the carrier frequency by phase modulation (sawtooth), while sine amplitude quantization causes harmonically related spurs although were considered as random errors and noise for a long time. Today, the operation of the phase accumulator is mathematically perfectly known as described in a technical paper[7] from Henry T. Nicholas and H. Samueli. After a thorough analysis, a model is presented such that the phase accumulator is considered a discrete phase sample permutation generator from which the frequency spurs can be predicted. Whatever the phase accumulator parameters (M, N, W), the length of the phase sequences equal to

$$\frac{2^L}{GCD(M,2^L)}$$

(where GCD is the greatest common divisor) is determined by the rightmost bit position, L, of the frequency tuning word, M, as shown in Figure 4. Hence, the value of L defines sequence classes, each sharing their own set of phase components, but permutated according to the

$$\frac{2^L}{GCD(M,2^L)}$$

ratio. These sequences of truncated phase samples generated in the time domain are used to determine, by DFT, the respective location and magnitude of each spurious line in the frequency domain. These sequences also demonstrate that odd values of M (FTW) exhibit the lowest frequency spur's amplitudes and suggests a simple modification of the phase accumulator to satisfy these minimum conditions by simply adding 1 LSB to the FTW. This way, the phase accumulator output sequences are forced to always have the same 2^N phase elements, whatever the M value and the initial content of the phase accumulator. The level of the worst spurious tone magnitude is then reduced by 3.922 dB and equal to SFDR_min (dBc) = 6.02 × W. The Nicholas modified phase accumulator confers several benefits to the NCO, as first it eliminates the cases where the rightmost bit of the FTW is too close to
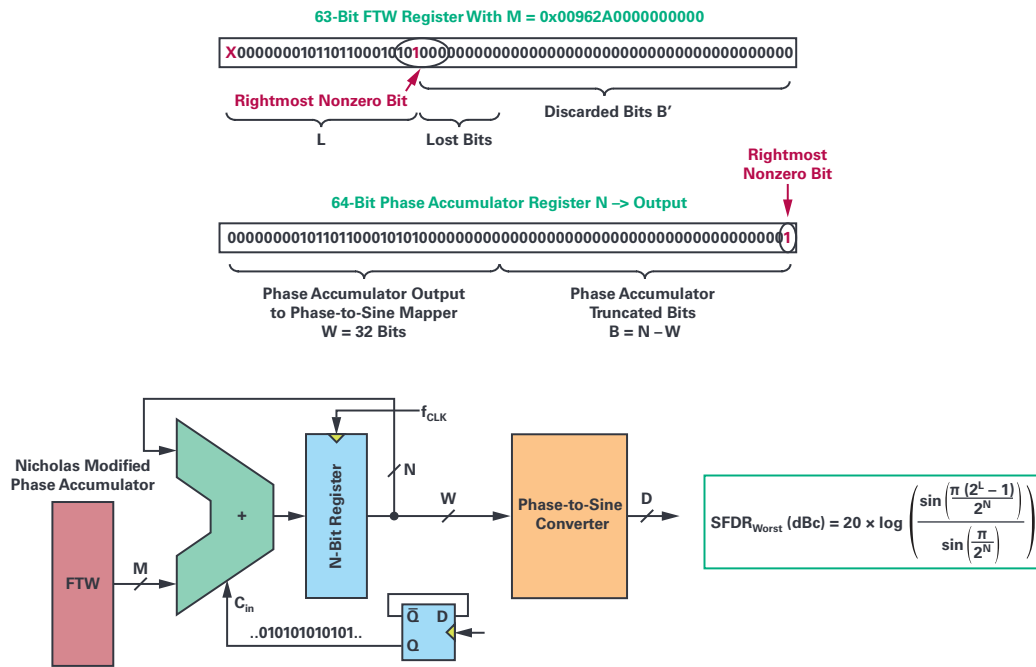
Figure 5. The position of the rightmost, nonzero bit of the FTW sets the theoretical SFDR worst-case level. The Nicholas modified phase accumulator solves the issue for any value of N and maximizes the SFDR of the NCO.

its MSB (frequency sweep in FMCW applications), and, secondly, it makes the spur's amplitude independent of the frequency tuning word, M. This modification is easily implemented in software by toggling the ALU LSB at the sampling rate f_S, the same behavior of the phase accumulator could be simulated as if the FTW LSB was set to logic 1. With a phase accumulator size N = 64 bits, a ½ LSB offset can be considered as a negligible error regarding the accuracy of the desired frequency F_{OUT}.

With an output phase word, W, of 32 bits, the maximum spur's amplitude due to phase truncation is therefore limited to a value of −192 dBc! Finite quantization of the sine sample values also leads to another set of frequency spurs, and it is commonly considered as noise and estimated by the well-known relationship $SNR_q(dB) = 6.02 \times D + 1.76$. This must be added to the parasitic elements due to the approximation errors of the phase-to-sine amplitude-conversion algorithm stage which, however, are considered negligible, given the extreme care in the choice of the phase-to-sine approximation algorithm and the calculation's precision.

These results indicate that both the linearity and the noise of our software sinusoidal NCO are at theoretical levels well beyond the required thresholds to test most of the high precision ADCs available on the market. It remains to find the last, but most critical elements of the signal chain: the reconstruction DAC and its complementary analog antialiasing filter and associated driver circuitry susceptible to meet the expected level of performance.

The Reconstruction DAC: The Achilles' Heel of the Thing!

The first temptation would be to select a high precision DAC with the best specification in terms of nonlinearity error (INL and DNL), like the superb AD5791, a 20-bit accurate DAC. But its resolution is only 20 bits and its R-2R architecture does not favor the reconstruction of signals, and especially the production of very pure sinusoids, because of its large glitches during input code transitions. Conventional DAC architectures built around binary weighted current generators or resistor networks are sensitive to digital feedthrough and digital switching impairment such as external or internal timing skew and other switching asymmetries of the digital input bits, particularly during major transitions for which the energy variation is consequent. This induces code-dependent transients, resulting in harmonic spurs of high amplitude.

At 20+-bit resolution, the use of an external ultralinear and fast sample and hold amplifier to deglitch the output of a DAC does not help much as it generates its own transients in tens of LSBs and introduces group-delay nonlinearity due to the resampling. For signal reconstruction, primarily in communication applications, the glitch issue is solved with the use of segmented architectures mixing fully decoded sections for the MSBs and binary weighted elements for the lowest significant bits. Unfortunately, no such commercial DAC currently exists beyond 16-bit precision. Instead of the fully predictable behavior of the NCO, the DAC errors are difficult to estimate and simulate accurately, especially when the manufacturers' dynamic specifications are rather weak or nonexistent, except for the DACs or ADCs dedicated to audio applications. The interpolating oversampled and multibit sigma-delta DAC then seems the only solution to be good enough for the job. With a resolution up to 32 bits, ultralow distortion, and high SNR, these state-of-the art converters are the best candidates for signal reconstruction over low to medium bandwidths. Trying to get the best noise and distortion performance within the audio spectrum or a slightly wider band (20 kHz or 40 kHz bandwidth), the best sigma-delta DAC within the Analog Devices portfolio is the AD1955 audio stereo DAC, still one of the best audio DACs available on the market, despite its resolution being limited to 24 bits.

Introduced in 2004, this audio DAC is based on a multibit sigma-delta modulator and oversampling techniques aided with various tricks to mitigate distortion and other plagues inherent to this principle of conversion.[8]

The AD1955 has got one of the best interpolation LP FIR filters of its kind, even today. It has a very high stop-band attenuation ($\approx$−120 dB) and a very low in-band ripple ($\approx$±0.0001 dB). Its two (left and right channels) DACs can operate up to 200 kSPS, but the best ac performance is achieved at 48 kSPS and 96 kSPS with a typical EIAJ standard, A-weighted, 120 dB figure for both its dynamic range and SNR in stereo mode. In mono mode, for which the two channels are simultaneously combined out of phase, a performance improvement of 3 dB can be expected. However, for wideband applications, these specifications are somewhat unrealistic since they are synthetic and restricted to the 20 Hz to 20 kHz bandwidth. Out-of-band noise and spurs are not considered beyond 20 kHz, partly because of the EIAJ standard, A-weighted filter, and audio industry specification definitions. This band-pass filter specific to audio measurements mimics the human ear frequency response and yields 3 dB better results over unfiltered measurements.

DDFS Hardware Demonstration Platform: Sine Wave Reconstruction with the AD1955

The complete DDFS has been implemented using two evaluation boards, one supporting the DSP processor and one for the analog signal reconstruction with the AD1955 DAC. The second-generation SHARC ADSP-21161N evaluation board was chosen for availability reasons as well as its ease of use and lean configuration for any audio applications. Still in production, the ADSP-21161N was designed a while ago, to support industrial, high-end consumer and professional audio applications, providing up to 110 Mips and 660 MFlops or 220 MMACS/s capabilities. Compared to the most recent generations of SHARC processors, the ADSP-21161N differs mostly by its short, 3-stage instruction pipeline, an on-chip, 1 Mb, only triple port RAM, and a reduced set of peripherals. The final and most critical stage of the precision tone generator is based upon the AD1955 evaluation board, which must faithfully reconstruct the analog signals from the samples delivered by the software NCO. This evaluation board carries an antialiasing filter (AAF) optimized for the audio bandwidth to meet the Nyquist criterion and has a couple of serial audio interfaces to support PCM/I²S and DSD digital streams besides the usual S/PDIF or AES-EBU receiver. The PCM/I²S serial link connector is used to connect the AD1955 DAC board to serial ports 1 and 3 connector (J) of the ADSP-21161N EVB. Both boards can be configured for I²S PCM or DSP modes of operation at 48 kSPS, 96 kSPS, or 192 kSPS sampling rates. The DSP serial port 1 generates the left and right channel data, the word select or L/R frame sync and the SCK bit clock signals needed by the digital input interface of the dual-channel DAC. The serial port 3 is just used to generate the DAC master clock, MCLK, required for the operation of the DAC interpolation filters and the sigma-delta modulators running 256 times (by default) faster than the input sampling frequency (48 kSPS). As all the DAC clocking signals are generated by the DSP, the board original, low cost Epson clock oscillator has been changed for an ultralow noise oscillator CCHD-957 from Crystek. Its phase noise specification could be as low as −148 dB/Hz at 1 kHz for a 24.576 MHz output frequency.

On the analog output side, active I/V converters must be used to hold the AD1955 current differential outputs at a constant common-mode voltage, typically 2.8 V, to minimize the distortion. Ultralow distortion and ultralow noise high precision operational amplifiers like the AD797 are used for this purpose and to handle analog signal reconstruction as well. As the two differential outputs are processed separately by the DSP, the stereo output configuration with its AAF topology has been selected instead of the mono mode. This AAF was simulated with LTspice® XVII with results given in Figure 6. As the last section of the filter is passive, an active differential buffer stage should be added like the recently introduced ADA4945. This low noise, ultralow distortion, fast settling time, fully differential amplifier is the almost perfect DAC companion to drive any high resolution SAR and sigma-delta ADCs. With a relatively large common-mode output voltage range and superb dc characteristics, the ADA4945 provides exceptional output balance and contributes to the suppression of even-order harmonic distortion products.

The EVB third-order filter has a −3 dB cutoff frequency of 76 kHz with an attenuation of only −31 db at 500 kHz. The in-band flatness is very good, but the out-of-band attenuation of this LP filter must be seriously improved, even if restricted to a pure reconstruction audio application. This is mandatory to reject the DAC shaped noise as well as the modulator clock frequency MCLK. Depending upon the use of the software DDS either for a single tone generator or an arbitrary waveform generator (AWG for complex waveforms), the AAF will be optimized for out-of-band attenuation or group delay distortion. As a practical example and comparison, the old but renowned SRS DS360 ultralow distortion function generator has been designed with a seventh-order Cauer AAF for a similar sampling rate. The signal reconstruction lies on the AD1862, a serial input 20-bit segmented R-2R DAC aimed

at digital audio applications. The AD1862 was able to sustain 20-bit word sampling rates up to 768 kHz (×16 f_S) and exhibits exceptional noise and linearity specifications. Its single-ended current output leaves the choice to use the best amplifier for the external I-to-V conversion stage.
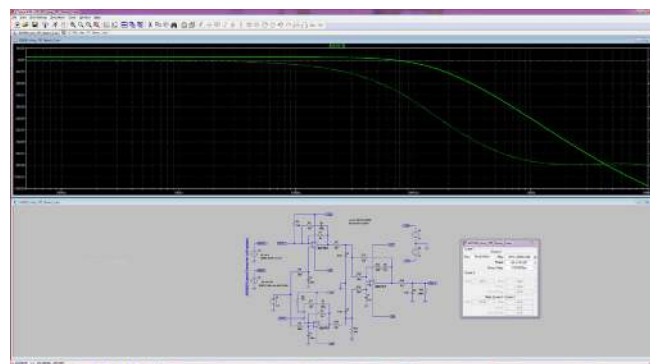

Figure 6. The LTspice simulated frequency response of the AD1955 EVB third-order antialiasing filter (stereo configuration).

The AD1955 and SHARC DSP combination was tested against several high resolution SAR ADCs such as the AD4020 with no external selective passive filters in between. By default, the basic AD4020 evaluation board offers no other choice than the on-board ADA4807 drivers. The simple circuitry to bias the ADC inputs at the V_REF/2 common-mode voltage imposes a rather low input impedance of 300 Ω and requires either signal isolation, ac coupling, or the use of an external differential amplifier module such as the EVAL-ADA4945-1. The AD4020 reference design board described in the circuit note CN-0513 is a better choice. It includes a discrete programmable gain instrumentation amplifier (PGIA), which provides a high input impedance and accepts ±5 V differential input signals (G = 1). Although these AD4020 boards and their SDP-H1 controller lack the capability to support coherent sampling acquisition, they allow decent waveform capture lengths for samples, ranging up to 1M. Thus, long FFTs with selective windowing are possible, providing both fine frequency resolution and a low noise floor. For example, with the seven-term Blackman-Harris window, the 1 Mpts FFT plot shown in Figure 7 illustrates the level of distortion of the AD1955 for a 990.059 Hz generated sine wave. The second harmonic is the largest distortion component and the largest spur at −111.8 dBc over a 350 kHz bandwidth. However, when considering the whole ADC Nyquist bandwidth of 806 kHz, the SFDR is limited by the DAC sigma-delta modulator and interpolating filter frequency and its second harmonic (384 kHz and 768 kHz).
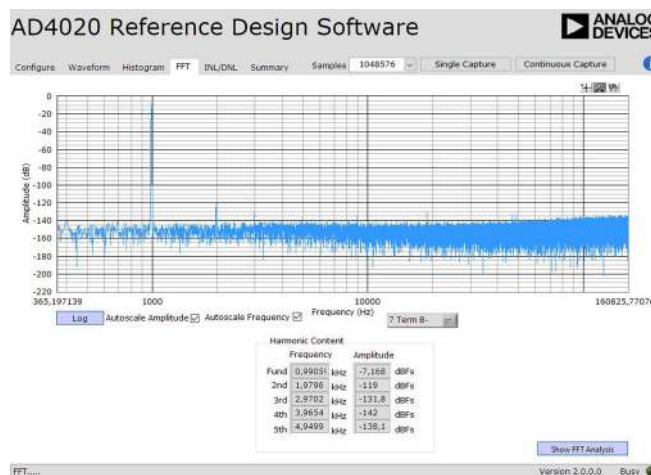

Figure 7. 1 M points FFT analysis shows pretty good distortion with H2 lower than −111 dBc, with the largest spur in the 10 kHz to 200 kHz band for a 1 kHz input frequency. The noise floor sits at about −146 dBFS.

In the same conditions, test trials were conducted on the vintage AD1862, which exhibited a slightly different spectral behavior. Placed in differential configuration, the two 20-bit DACs clocked at about 500 kSPS reported a noise floor of −151 dBFS, a THD of −104.5 dB for a sine output level of 12 V p-p at 1.130566 kHz. The SFDR over the AD4020 Nyquist bandwidth (806 kHz) is close to 106 dB limited by the third-harmonic. The DAC reconstruction filter based around two AD743 low noise FET amplifiers is a third-order similar to the one of the AD1955 evaluation board, but with a cutoff frequency of 35 kHz at −3 dB.

To become effective, the DDS-based generator requires a decent filter capable of an attenuation greater than 100 dB at about 250 kHz for a generated dc to 25 kHz CW signal frequency range. This can be achieved with a sixth-order Chebyshev and even a sixth-order Butterworth LP filter for a perfect in-band flatness. The order of the filter will be minimized to limit the number of analog stages and their nonidealities such as noise and distortion.

Conclusion

Preliminary and out of the box tests performed on standard evaluation boards demonstrate that the processor-based DDS techniques for conventional sine wave CW generation with top performance are within reach. The −120 dBc harmonic distortion figure could be met with a careful design of the reconstruction filter and the analog output buffer stage. The DSP-based NCO/DDS is not only restricted to the generation of single tone sine waves. By using an optimized AAF (Bessel or Butterworth) with an appropriate cutoff frequency and no other hardware change, the same DSP and DAC combination can be disguised into a high performance AWG to produce any type of waveform, for example, to synthesize fully parametrizable multitone sine waves with a full control of the phase and amplitude of each component for IMD testing.

Since floating-point arithmetic is crucial for applications requiring high accuracy and/or high dynamic range, today the SHARC+ DSP processors such as the low cost ADSP-21571 or the SoC ADSP-SC571 (ARM® and SHARC) are the de facto standard for real-time processing up to an aggregated sampling rate of 10 MSPS. Clocked at 500 MHz, the dual SHARC cores and their hardware accelerators can provide more than 5 Gflops computation performance and offer tons of internal specialized SRAM, the basic ingredients demanded by the tasks for the generation of any kind of waveforms as well as for complex analysis processing. This type of application shows that the systematic use of hardware programmable solutions is not mandatory for handling precision digital signal processing. Floating-point processors and their complete development environments allow easy and fast code portability from simulators such as MATLAB, as well as rapid debugging thanks to Analog Devices' CCES and VDSP++ C and C++ compilers and their full suite of simulators and real-time debuggers.

References

[1] Joseph A. Webb. U.S. patent US3654450 April 1970.

[2] Joseph Tierney, Charles M. Rader, and Bernard Gold. "A Digital Frequency Synthesizer." *IEEE Transactions on Audio and Electroacoustics*, Vol. 19, Issue 1, March 1971.

[3] Jim Williams and Guy Hoover. AN-132: Fidelity Testing for A→D Converters Proving Purity. Analog Devices, Inc., February 2011.

[4] John F. Hart. *Computer Approximations*. Krieger Publishing Company, 1978.

[5] William J. Cody and William Waite. *Software Manual for the Elementary Functions*. Prentice-Hall, Inc., 1980.

[6] Robin Green. "Faster Math Functions, Part 2 Presentation." Sony Computer Entertainment America, May 2016.

[7] Henry T. Nicholas and Henry Samueli. "An Analysis of the Output Spectrum of Direct Digital Frequency Synthesizers in the Presence of Phase-Accumulator Truncation." IEEE, May 1987.

[8] Robert Adams, Khiem Nguyen, and Karl Sweetland. "A 113 dB SNR Oversampling DAC with Segmented Noise-Shaped Scrambling." IEEE, February 1998.

ADSP-21000 Family Application Handbook Volume 1. Analog Devices, Inc., May 1994.

A Technical Tutorial on Digital Signal Synthesis. Analog Devices, Inc., March 2001.

Butler, Oscar. "Internship Report Summer 2017: High Precision Oversampled 20-Bit Ultra Low Power Acquisition System." Analog Devices, Inc., 2017.

Crawford, James A. *Advanced Phase-Lock Applications: Frequency Synthesis*. AMI, LLC, May 2011.

Evaluation Board User Guide UG-048. Analog Devices, Inc., February 2010.

EV-4020-REF-DGNZ Reference Design Board User Guide UG-1280. Analog Devices, Inc., May 2019.

Goldberg, Bar-Giora. *Digital Techniques in Frequency Synthesis*. McGraw-Hill, August 1995.

Model DS360 Ultra Low Distortion Function Generator. Stanford Research Systems, 1999.

Symons, Pete. Digital Waveform Generation. Cambridge University Press, November 2013.

AD1862 data sheet. Analog Devices, Inc., July 2011.

1241-2010 - IEEE Standard for Terminology and Test Methods for Analog-to-Digital Converters. IEEE, January 2011.

About the Author

Patrick Butler is a field applications engineer with Analog Devices' south Europe sales organization, supporting the French global market and some ADEF customers. He has been with ADI since 1984, supporting the DSP building blocks ICs, as well as high speed converters. Previously, he worked as a design engineer in the ATE division of Schlumberger in Saint-Étienne, France for five years, and then occupied several application engineer and FAE positions at Matra-MHS in Nantes, AMD and Harris SC-Intersil. Today, his main hobby is collecting vintage sound components to build active, high efficiency horn loudspeaker-based systems with the help of his two sons.

Online Support Community

ADI *EngineerZone*™
SUPPORT COMMUNITY

Engage with the Analog Devices technology experts in our online support community. Ask your tough design questions, browse FAQs, or join a conversation.

Visit *ez.analog.com*