DDS Devices Generate High-Purity Waveforms Simply, Efficiently, and Flexibly

By Brendan Cronin

Direct digital synthesis (DDS) technology is used to generate and modify high-quality waveforms in a broad range of applications.

Introduction

A key requirement across a multitude of industries is to accurately produce, easily manipulate, and quickly change waveforms of various frequencies and types. Whether a wideband transceiver requires an agile low-phase-noise frequency source with excellent spurious-free dynamic performance or an industrial measurement and control system needs a stable frequency stimulus, the ability to quickly, easily, and cost effectively generate an adjustable waveform while maintaining phase continuity is a critical design criterion that *direct digital frequency synthesis* can fulfill.

The Task of Frequency Synthesis

Increasing spectrum congestion, coupled with the insatiable need for lower power, higher quality measurement equipment, calls for the use of new frequency ranges and better exploitation of existing ones. As a result, better control of frequency generation is being sought—in most cases with the assistance of *frequency synthesizers*. These devices use a given frequency, f_c, to generate a waveform at a related desired frequency (and phase), f_{out}. The general relationship can be written simply as

$$f_{out} = \varepsilon_x \cdot f_c$$

Where the scaling factor, ε_x, is sometimes called the *normalized frequency*.

The equation is always implemented using algorithms for step-by-step approximations of real numbers. When the scaling factor is a rational number, a ratio of two relatively prime integers, the output frequency and the reference frequency will be harmonically related. In most cases, however, ε_x can belong to a much broader set of real numbers, and the approximation process is truncated as soon as it falls within an acceptable limit.

Direct Digital Frequency Synthesis

One such practical realization of a frequency synthesizer is *direct digital frequency synthesis* (DDFS), often shortened to *direct digital synthesis* (DDS). The technique uses digital data processing to generate a frequency- and phase-tunable output related to a fixed frequency reference, or clock source, f_c. In a DDS architecture, the reference or system clock frequency is divided down by the scaling factor, set by a programmable binary tuning word.

Stated simply, a direct digital frequency synthesizer translates a train of clock pulses into an analog waveform, typically a sine, triangular, or square wave. As Figure 1 shows, its essential parts are: *a phase accumulator,* which produces a number corresponding to a phase angle of the output waveform, a *phase-*
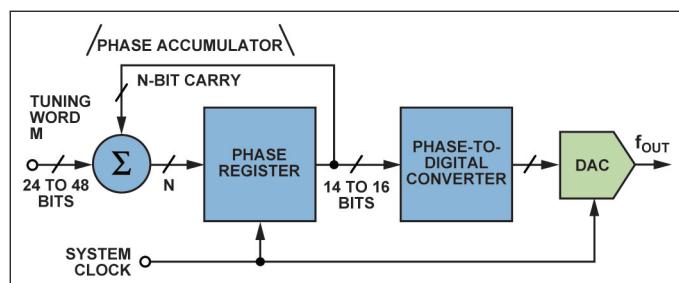

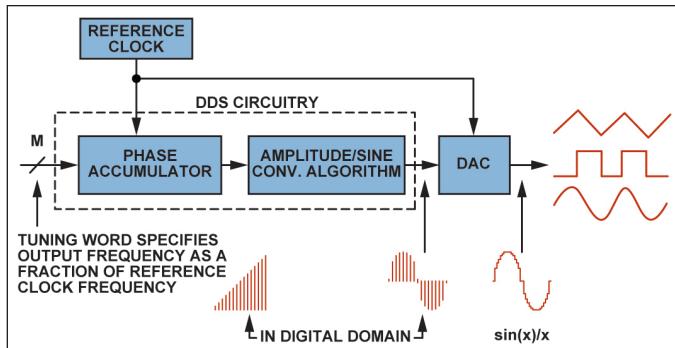**Figure 1 • Functional block diagram of a DDS system.**

Figure 2 • Typical DDS architecture and signal path with DAC.

to-digital converter, which generates the instantaneous digital fraction of the output amplitude occurring at a particular phase angle, and a *D/A converter* (DAC), which converts that digital value to a sampled analog data point.

For sine-wave outputs, the phase-to-digital amplitude converter is usually a sine lookup table (Figure 2). The phase accumulator counts by *N*, to generate a frequency related to f_C according to the equation,

$$f_{out} = \frac{N}{2^M} f_c$$

where *M* is the resolution of the tuning word (24 bits to 48 bits) and *N* is the number of pulses of f_c corresponding to the smallest incremental phase change of the phase accumulator's output word.

Since changes to *N* result in immediate changes in the output phase and frequency, the system is inherently *phase-continuous,* a critical attribute in many applications. No loop settling time is required, in contrast to analog-type systems, such as *phase-locked-loops* (PLLs).

The DAC is usually a high-performance circuit specifically designed to work with the DDS core (phase accumulator and phase-to-amplitude converter). In most cases, the resulting device, often a single chip, is commonly referred to as a complete DDS or C-DDS.

Practical DDS devices often integrate multiple registers to allow various frequency- and phase-modulation schemes to be realized. When included, the phase register's contents are added after the phase accumulator. This enables the output sine wave to be phase-delayed in correspondence with a phase tuning word. This is extremely useful for phase-modulation applications in communication systems. The resolution of the adder circuit determines the number of bits in the phase tuning word and, therefore, the resolution of the delay.

Integrating a DDS engine and a DAC in a single device has advantages

and disadvantages, but whether integrated or not, a DAC is required to create a high quality analog signal of exceptional purity. The DAC converts the digital sine output into an analog sine wave and may be either single-ended or differential. A few of the key requirements are low phase noise, excellent wide-band (WB-) and narrow-band (NB-) *spurious-free dynamic range* (SFDR), and low power consumption. If it is an external component, the DAC needs to be fast enough to process the signal—so devices with a parallel port are common.

DDS versus Other Solutions

Other possibilities for frequency generation include analog *phase-locked loops* (PLLs), clock generators, and using an FPGA to dynamically program the output of a DAC. A simple comparison of the technologies can be made by examining spectral performance and power consumption, qualitatively demonstrated in Table 1.

A phase-locked loop is a feedback loop comprising: a *phase comparator,* a divider, and a *voltage-controlled oscillator* (VCO). The phase comparator compares a reference frequency with the output frequency (usually divided down by a factor, *N*). The error voltage generated by the phase comparator is applied to the VCO, which generates the output frequency. When the loop has settled, the output will bear an accurate relationship to the reference in frequency and/or phase. PLLs have long been recognized as superior devices for low phase noise and high *spurious-free dynamic range* (SFDR) applications requiring high fidelity and stable signals in a specific band of interest.

Their inability to accurately and quickly tune the frequency output and waveform and their slow response limits their suitability for applications such as agile frequency hopping and some frequency- and phase-shift keying applications.

Other approaches, including *field-programmable gate arrays* (FPGAs) with embedded DDS engines—in combination with off-the-shelf DACs to synthesize output sine waves—solve the frequency-hopping difficulties of PLLs, but have their own weaknesses. The main system disadvantages include higher operating and interface power requirements, higher cost, large size, and additional software, hardware, and memory overhead for the system developer. For example, up to 72 kbytes of memory are

	Power Consumption	Spectral Purity	Comment
DDS	Low	Medium	Easy to Tune
Discrete DAC + FPGA	Medium	Medium-High	Ability to Tune
Analog PLL	Medium	High	Difficult to Tune

Table 1 • DDS versus competing technologies—high level comparison.

required to generate a 10-MHz output signal with 60-dB dynamic range using the DDS engine option on modern FPGAs. In addition, the designer needs to be comfortable and familiar with subtle trade-offs and the architecture of the DDS core.

As a practical matter (Table 2), rapid advances in CMOS processing, together with modern digital design techniques and improved DAC topologies, have resulted in the DDS technology achieving power consumption, spectral performance, and cost levels that were previously unattainable for a wide range of applications. While complete DDS products will never match the highest performance and design flexibility achievable with custom combinations of high-end D/A converter technology and FPGAs, the size-, power- and cost benefits, coupled with the simplicity of DDS devices, may make them easily the first choice for many applications.

Also note that since a DDS device fundamentally embodies a digital method of generating an output waveform, it can simplify the architecture of some solutions or make it possible to digitally program the waveform. While a sine wave is normally used to explain the function and operation of a DDS, it is easily possible to generate triangular or square (clock) wave outputs from modern DDS ICs, avoiding the need for a lookup table in the former case, and for a DAC in the latter case, where the integration of a simple yet precise comparator will suffice.

DDS Performance and Limitations

The actual output of the DAC is not a continuous sine wave but a train of pulses with a sinusoidal time envelope. The corresponding frequency spectrum is a set of images and aliases. The images lie along a sin x/x envelope (see |amplitude| plot in Figure 3). Filtering is necessary to suppress frequencies outside the band of interest, but it cannot suppress higher-order aliases (due to DAC nonlinearities, for example) appearing within the pass band.

The *Nyquist Criterion* dictates that a minimum of two samples per cycle are required to reconstruct a desired
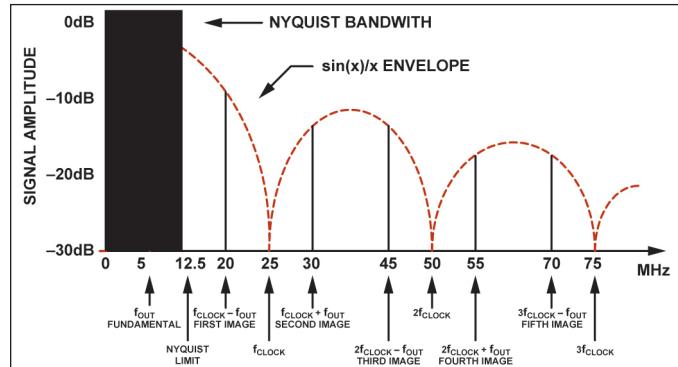

Figure 3 • Sin x/x rolloff in a DDS.

output waveform. Image responses are created in the sampled output spectrum at $K_{fCLOCK} \pm f_{OUT}$. In this example, where f_{CLOCK} = 25 MHz and f_{OUT} = 5 MHz, the first and second images occur (see Figure 4) at $f_{CLOCK} \pm f_{OUT}$, or 20 MHz and 30 MHz. The third and fourth images appear at 45 MHz and 55 MHz. Note that the sin x/x nulls appear at multiples of the sampling frequency. In the case where f_{OUT} is greater than the *Nyquist bandwidth* (1/2 f_{CLOCK}), the first image response will appear within the *Nyquist bandwidth* as an *aliased* image (a 15-MHz signal will alias down to 10 MHz, for example). The aliased image cannot be filtered from the output with a traditional Nyquist *anti-aliasing* filter.

In typical DDS applications, a low-pass filter is utilized to suppress the effects of the image responses in the output spectrum. To keep the cutoff requirements of the low-pass filter reasonable and the filter design simple, an accepted guideline is to limit the f_{OUT} bandwidth to approximately 40% of the f_{CLOCK} frequency using an economical low-pass output filter.

The amplitude of any given image in response to the fundamental can be calculated using the sin x/x formula. Because the function rolls off with frequency, the amplitude of the fundamental output will decrease inversely with its tuned frequency; in a DDS system, the decrease will be –3.92 dB over the dc to Nyquist bandwidth.

The amplitude of the first image is substantial—within 3 dB of the fundamental. To simplify filtering requirements for DDS applications, it is important to generate a frequency plan and analyze the spectral considerations of the image and sin x/x amplitude responses at the desired f_{OUT} and f_{CLOCK} frequencies. Online interactive design tools supporting the Analog Devices DDS product family allow for quick and easy simulation of where images lie

	Phase-Locked Loops	DAC + FPGA	DDS
Spectral Performance	High	Medium-High	Medium
System Power Requirements	High	High	Low
Digital Frequency Tuning	No	Yes	Yes
Tuning Response Time	High	Low	Low
Solution Size/Footprint	Medium	High	Low
Waveform flexibility	Low	Medium	High
Cost	Medium	High	Low
Design Re-use	Medium	Low	High
Implementation complexity	Medium	High	Low

Table 2 • Benchmark Analysis Summary—Frequency-Generation Technologies (<50 MHz).

and allow the user to choose frequencies where images are outside the band of interest. See the links section for additional useful information.

Other anomalies in the output spectrum, such as integral and differential linearity errors of the DAC, glitch energy associated with the DAC, and clock feedthrough noise, will not follow the $\sin x/x$ roll-off response. These anomalies will appear as harmonics and spurious energy in many places in the output spectrum—but will generally be much lower in amplitude than the image responses. The general noise floor of a DDS device is determined by the cumulative combination of substrate noise, thermal noise effects, ground coupling, and other sources of signal coupling. The noise floor, performance spurs, and jitter of a DDS device are greatly influenced by circuit board layout, the quality of the power supplies, and—most importantly—the quality of the input reference clock.

Jitter

A perfect clock source would have edges occurring at precise intervals in time that would never vary. This, of course, is impossible; even the best oscillators are constructed from non-ideal components and have noise and other imperfections. A high-quality, low-phase-noise crystal oscillator will have jitter on the order of picoseconds, accumulated over many millions of clock edges. Jitter is caused by thermal noise, instabilities in the oscillator's electronic circuitry, and external interference through the power, ground, and output connections—all contributing to disturbances in the oscillator's timing. In addition, oscillators are influenced by external magnetic or electric fields, and RF interference from nearby transmitters. A simple amplifier, inverter, or buffer in the oscillator circuitry will also add jitter to a signal.

So choosing a stable reference clock oscillator with low jitter and sharp edges is critical. Higher frequency reference clocks allow greater oversampling, and jitter can be somewhat ameliorated by *frequency division,* since dividing the frequency of the signal yields the same amount of jitter across a longer period, and so reduces the percentage of jitter on the signal.

Noise—Including Phase Noise

Noise in a sampled system depends on many factors, starting with reference clock jitter, which shows up as phase noise on the fundamental signal. In a DDS system,

truncation of the phase register output may introduce code-dependent system errors. Binary-coded words don't cause truncation errors. For non-binary coded words, however, the phase-noise truncation error produces spurs in the spectrum. The frequency/magnitude of the spurs is determined by the code word. The DAC's quantization and linearity errors will also add harmonic noise in the system. Time-domain errors, such as undershoot/overshoot and code glitches, all contribute distortion to the output signal.

Applications

DDS applications can be segmented into two primary categories:

- Communication and radar systems that require agile frequency sources for data encoding and modulation applications
- Measurement, industrial, and optical applications that require a generic frequency synthesis function with programmable tuning, sweeping, and excitation

In both cases, an increasing trend towards higher spectral purity (lower phase noise and higher spurious free dynamic range) is coupled with low operating power and size requirements for remote or battery operated equipment.

DDS in Modulation/Data Encoding and Synchronization

From its exclusive origins in radar and military applications, some of the advances in DDS product characteristics (performance improvements, cost, and size) have combined to make DDS technology very popular in modulation and data encoding applications. This section will discuss two data encoding schemes and their proposed implementation with a DDS system.

Binary frequency shift keying (BFSK, or simply *FSK*) is one of the simplest forms of data encoding. The data is transmitted by shifting the frequency of a continuous carrier between one (binary 1, or *mark*) and the other (binary 0, or *space*) of two discrete frequencies. Figure 4 shows the relationship between the data and the transmitted signal.

Binary 1s and 0s are represented as two different frequencies, f_0 and f_1, respectively. This encoding scheme is
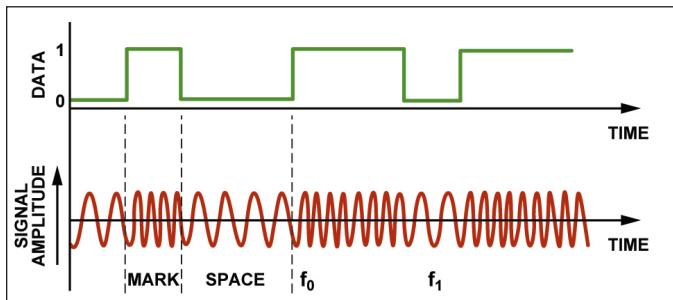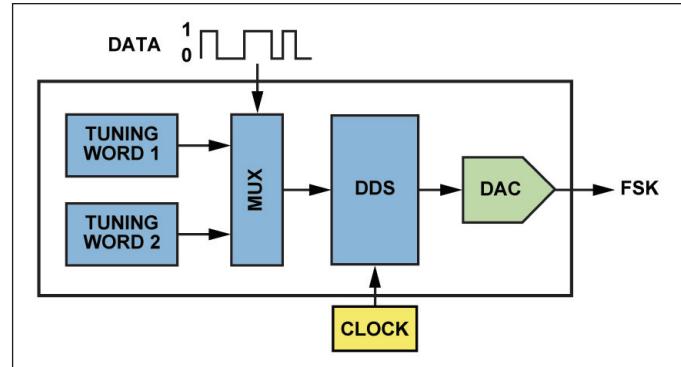

Figure 4 • Binary FSK modulation.

Figure 5 • FSK encoding using the tuning-word selector of an AD9834 or AD9838 DDS.

easily implemented with a DDS device. The DDS frequency tuning word representing the output frequencies is changed so that f_0 and f_1 are generated from 1s and 0s to be transmitted. In at least two members of Analog Devices complete-DDS product families (the AD9834 and the AD9838—see also the Appendix), the user can simply program the two current FSK frequency tuning words into the IC's embedded frequency registers. To shift output frequency, a dedicated pin, FSELECT, selects the register containing the appropriate tuning word (Figure 5).

Phase-shift keying (PSK) is another simple form of data encoding. In PSK, the frequency of the carrier remains constant, and the phase of the transmitted signal is varied to convey the information. Several schemes can be used to accomplish PSK. The simplest method, commonly known as binary PSK (or BPSK), uses only two signal phases: 0° (logic 1) and 180° (logic 0). The state of each bit is determined according to the state of the preceding bit. If the phase of the wave does not change, the signal state stays the same (low or high). If the phase of the wave changes by 180°, i.e., if the phase reverses—the signal state changes (low to high, or high to low). PSK encoding is easily implemented with a DDS product as most of the devices have a separate input register (a *phase register*) that can be loaded with a phase value. This value is directly added to the phase of the carrier without changing its frequency. Changing the contents of this register modulates the phase of the carrier, generating a PSK output. For applications that require high-speed modulation, the AD9834 and AD9838, which have pairs of phase registers, allow signals on a PSELECT pin to alternate between the preloaded phase registers to modulate the carrier as required.

More complex forms of PSK employ four or eight wave phases. This allows binary data to be transmitted at a faster rate per phase change than is possible with BPSK modulation. In four-phase modulation (*quadrature* PSK), the possible phase angles are 0°, +90°, –90°, and +180°; each phase shift can represent two signal elements. The AD9830, AD9831, AD9832, and AD9835 provide four
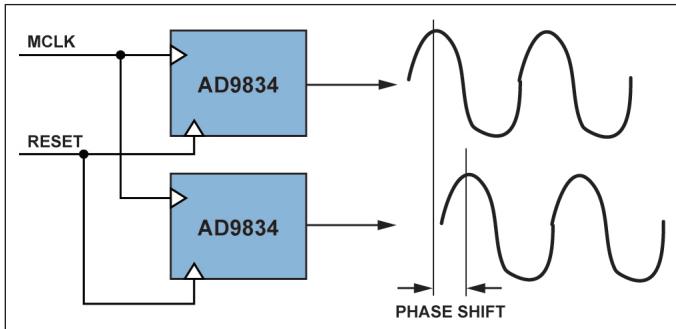
Figure 6 • Synchronizing two DDS components.

Figure 7 • Typical network analysis architecture using frequency stimulus.

phase registers to allow complex phase modulation schemes to be implemented by continuously updating different phase offsets to the registers.

I/Q Capability Using Multiple DDS Components in Synchronous Mode

Many applications require the generation of two or more sinusoidal or square wave signals having a known phase relationship. A popular example is *in-phase and quadrature modulation* (I/Q), a technique wherein signal information is derived from a carrier frequency at its 0° and 90° phase angles. Two single DDS components can be run from the same source clock to output signals whose phase relationship can be directly controlled and manipulated. In Figure 6, the AD9834 devices are programmed using one reference clock; the same RESET pin is used to update both devices. In this way, simple I/Q modulation can be achieved.

A reset must be initiated after power-up and before transferring any data to the DDS. This establishes the DDS output in a known phase, which becomes the common reference angle that allows synchronization of multiple DDS devices. When new data is sent simultaneously to multiple DDS devices, a coherent phase relationship can be maintained—or the relative phase offset between multiple DDS devices can be predictably shifted by means of the phase offset register. The AD983x series of DDS products have 12 bits of phase resolution, providing an effective resolution of 0.1°.

For more information about synchronizing multiple DDS devices, see Application Note AN-605, "Synchronizing Multiple AD9852 DDS-Based Synthesizers".

Network Analysis

Many applications in the electronic world involve the gathering and decoding of data from networks such as analog measurement and optical communications systems. Normally, the system analysis requirement is to stimulate a circuit or system with a frequency of known amplitude and phase, and analyze the signal characteristics of the response signal through the system.

The information gathered on the response signal is used to determine key system information. The range of
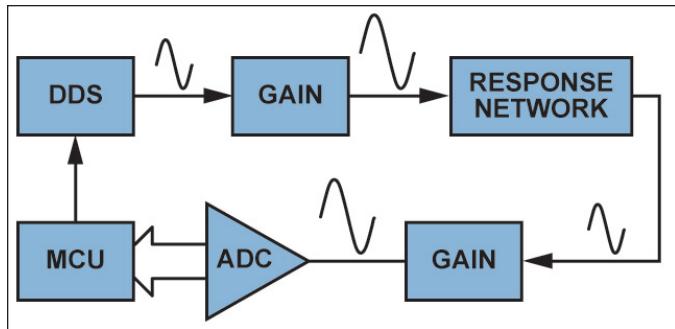
networks being tested (see Figure 7) can be quite wide, including cable integrity testing, biomedical sensing, and flow-rate measurement systems. Wherever the basic requirement is to generate frequency-based signals and compare phase and amplitude of the response signal(s) to the original signal, or if a range of frequencies needs to be excited through the system, or if test signals with different phase relationships (as in systems with I/Q capability) are required, direct digital synthesis ICs can be highly useful for digitally controlling stimulus frequency and phase through software with simplicity and elegance.

Cable Integrity/Loss Measurement

Cable integrity measurement is a noninvasive method of analyzing cables in applications such as airplane wiring, local area networks (LANs), and telephone lines. One way to determine performance is to see how much signal is lost through the cable. By injecting a signal of known frequency and amplitude, the user can calculate cable attenuation by measuring the amplitude and phase at remote portions of the cable. Parameters such as dc resistance and characteristic impedance will affect a particular cable's attenuation. The result is usually expressed in decibels below the signal source (0 dB) over the frequency range of the test. The frequencies of interest depend on the cable type. DDS devices, with their ability to generate a wide range of frequencies, can be used as a stimulus with the necessary frequency resolution.

Flow Meter

A related application area is in water, other liquids, and gas flow analysis in pipelines. An example is ultrasonic flow measurement, which operates on a phase-shift principle, as shown in Figure 8. Basically, a signal is transmitted from one side of the channel where the liquid is flowing and a transducer sensor is positioned on the opposite side to measure the phase response—which depends on the flow rate. There are many variations on this technique. Test frequencies depend on the substance being measured; in general, the output signal is often transmitted over a range of frequencies. DDS provides the flexibility to set and change the frequency seamlessly.
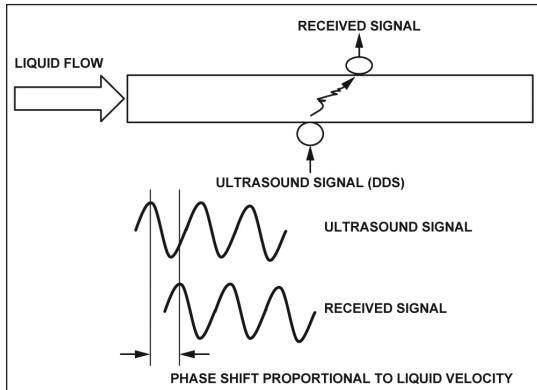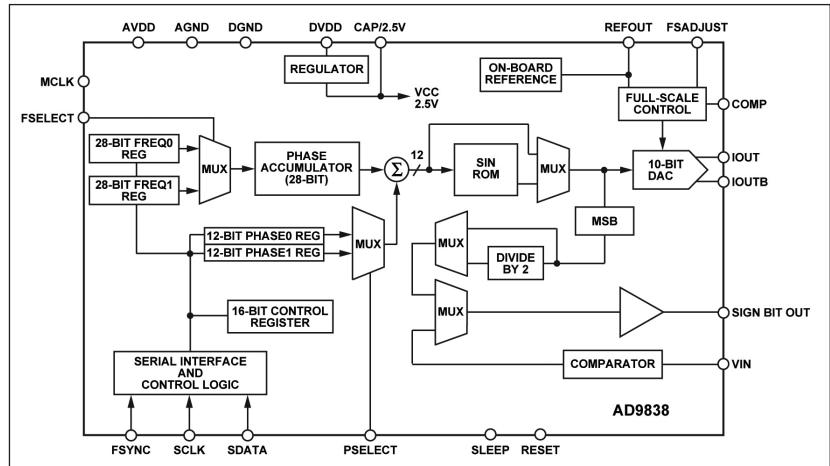
Figure 8 • Ultrasonic flowmeter.

Figure 10 • Block diagram of the AD9838 DDS.

Interactive Design Tool

What is it? An online interactive design tool for DDS is an assistant for selecting tuning words, given a reference clock and desired output frequencies and/or phases. The tool shows the tuning word and other configuring bits encoded as a sequence of codes for use in programming the part via its serial interface. Idealized output harmonics can be shown for the selected reference clock and output frequency after an external reconstruction filter has been applied. Links to ADI's design tools can be found on the Interactive Design Tools homepage. One example is the AD9834 design tool.

Evaluation Kit

The AD983x series of products come with a fully functional evaluation kit with schematics and layout. The software provided in the evaluation kit allows the user to easily program, configure, and test the device (Figure 9).

Other useful DDS information can be found on the DDS website.

Murphy, Eva and Colm Slattery, "All About Direct Digital Synthesis" Ask The Applications Engineer—33, *Analog Dialogue,* Vol. 38 No. 3, (2004) pp. 8-12.

A Technical Tutorial on Digital Signal Synthesis (1999 122 pages).

The AD9838 in Brief: A block diagram of the AD9838 DDS appears in Figure 10. Built on a fine-line CMOS process, the device is an ultra-low power, (11-mW) complete DDS. The 28-bit frequency registers permit 0.06-Hz frequency resolution with a 16-MHz clock and 0.02-Hz frequency resolution with a 5-MHz clock. Phase- and frequency modulation are configured via on-chip registers using software or pin selection. The device features 68dBc wide-band and –97-dBc narrow-band SFDR and operates over the extended temperature range of –40°C to 125°C. The device is housed in a small 4-mm × 4-mm 20-lead LFCSP (lead-frame chip-scale package).

About the Author

Brendan Cronin is a Product Marketing Engineer working with the Core Products and Technologies (CPT) Group within Analog Devices. Brendan joined ADI in 1998 and spent six years working as a Mixed-Signal Design Engineer in the Industrial and Automotive products group. Brendan's focus today is on linear and related technologies.
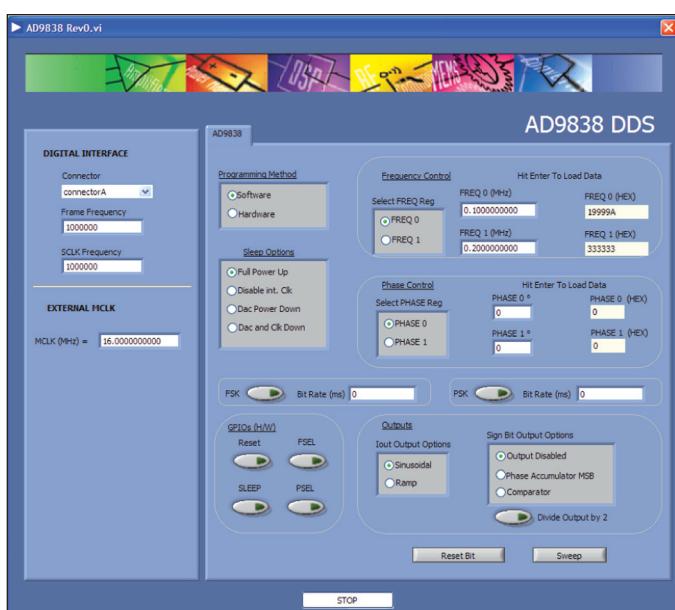

Figure 9 • AD9838 evaluation software interface.