

Ask The Application Engineer—33

All About Direct Digital Synthesis

By Eva Murphy [eva.murphy@analog.com]
Colm Slattery [colm.slattery@analog.com]

What is Direct Digital Synthesis?

Direct digital synthesis (DDS) is a method of producing an analog waveform—usually a sine wave—by generating a time-varying signal in digital form and then performing a digital-to-analog conversion. Because operations within a DDS device are primarily digital, it can offer fast switching between output frequencies, fine frequency resolution, and operation over a broad spectrum of frequencies. With advances in design and process technology, today's DDS devices are very compact and draw little power.

Why would one use a direct digital synthesizer (DDS)? Aren't there other methods for easily generating frequencies?

The ability to accurately produce and control waveforms of various frequencies and profiles has become a key requirement common to a number of industries. Whether providing agile sources of low-phase-noise variable-frequencies with good spurious performance for communications, or simply generating a frequency stimulus in industrial or biomedical test equipment applications, convenience, compactness, and low cost are important design considerations.

Many possibilities for frequency generation are open to a designer, ranging from *phase-locked-loop* (PLL)-based techniques for very high-frequency synthesis, to dynamic programming of *digital-to-analog converter* (DAC) outputs to generate arbitrary waveforms at lower frequencies. But the DDS technique is rapidly gaining acceptance for solving frequency- (or waveform) generation requirements in both communications and industrial applications because single-chip IC devices can generate programmable analog, output waveforms simply and with high resolution and accuracy.

Furthermore, the continual improvements in both process technology and design have resulted in cost and power consumption levels that were previously unthinkable low. For example, the AD9833, a DDS-based programmable waveform generator (Figure 1), operating at 5.5 V with a 25-MHz clock, consumes a maximum power of 30 milliwatts.




Figure 1. The AD9833—a one-chip waveform generator.

What are the main benefits of using a DDS?

DDS devices like the AD9833 are programmed through a high-speed *serial peripheral-interface* (SPI), and need only an external clock to generate simple sine waves. DDS devices are now available that can generate frequencies from less than 1 Hz up to 400 MHz (based on a 1-GHz clock). The benefits of their low power, low cost, and single small package, combined with their inherent excellent performance and the ability to digitally program (and re-program) the output waveform, make DDS devices an extremely attractive solution—preferable to less-flexible solutions comprising aggregations of discrete elements.

What kind of outputs can I generate with a typical DDS device?

DDS devices are not limited to purely sinusoidal outputs. Figure 2 shows the square-, triangular-, and sinusoidal outputs available from an AD9833.



Figure 2. Square-, triangular-, and sinusoidal outputs from a DDS.

How does a DDS device create a sine wave?

Here's a breakdown of the internal circuitry of a DDS device: its main components are a *phase accumulator*, a means of *phase-to-amplitude conversion* (often a sine look-up table), and a DAC. These blocks are represented in Figure 3.




Figure 3. Components of a direct digital synthesizer.

A DDS produces a sine wave at a given frequency. The frequency depends on two variables, the *reference-clock* frequency and the binary number programmed into the frequency register (*tuning word*).

The binary number in the frequency register provides the main input to the phase accumulator. If a sine look-up table is used, the phase accumulator computes a phase (angle) address for the look-up table, which outputs the digital value of amplitude—corresponding to the sine of that phase angle—to the DAC. The DAC, in turn, converts that number to a corresponding value of analog voltage or current. To generate a fixed-frequency sine wave, a constant value (the phase increment—which is determined by the binary number) is added to the phase accumulator with each clock cycle. If the phase increment is large, the phase accumulator will step quickly through the sine look-up table and thus generate a high-frequency sine wave. If the phase increment is small, the phase accumulator will take many more steps, accordingly generating a slower waveform.

What do you mean by a *complete* DDS?

The integration of a D/A converter and a DDS onto a single chip is commonly known as a *complete* DDS solution, a property common to all DDS devices from ADI.

Let's talk some more about the phase accumulator. How does it work?

Continuous-time sinusoidal signals have a repetitive angular phase range of 0 to 2π . The digital implementation is no different. The counter's carry function allows the phase accumulator to act as a *phase wheel* in the DDS implementation.

To understand this basic function, visualize the sine-wave oscillation as a vector rotating around a phase circle (see Figure 4). Each designated point on the phase wheel corresponds to the equivalent point on a cycle of a sine wave. As the vector rotates around the wheel, visualize that the sine of the angle generates a corresponding output sine wave. One revolution of the vector around the phase wheel, at a constant speed, results in one complete cycle of the output sine wave. The phase accumulator provides the equally spaced angular values accompanying the vector's linear rotation around the phase wheel. The contents of the phase accumulator correspond to the points on the cycle of the output sine wave.




Figure 4. Digital phase wheel.

The phase accumulator is actually a modulo- M counter that increments its stored number each time it receives a clock pulse. The magnitude of the increment is determined by the binary-coded input word (M). This word forms the phase step size between reference-clock updates; it effectively sets how many points to skip around the phase wheel. The larger the jump size, the faster the phase accumulator overflows and completes its equivalent of a sine-wave cycle. The number of discrete phase points contained in the *wheel* is determined by the resolution of the phase accumulator (n), which determines the tuning resolution of the DDS. For an $n = 28$ -bit phase accumulator, an M value of 0000...0001 would result in the phase accumulator overflowing after 2^{28} reference-clock cycles (increments). If the M value is changed to 0111...1111, the phase accumulator will overflow after only 2 reference-clock cycles (the minimum required by Nyquist). This relationship is found in the basic tuning equation for DDS architecture:

$$f_{OUT} = \frac{M \times f_C}{2^n}$$

where:

f_{OUT} = output frequency of the DDS

M = binary tuning word

f_C = internal reference clock frequency (system clock)

n = length of the phase accumulator, in bits

Changes to the value of M result in immediate and *phase-continuous* changes in the output frequency. No loop settling time is incurred as in the case of a phase-locked loop.

As the output frequency is increased, the number of samples per cycle decreases. Since sampling theory dictates that at least two samples per cycle are required to reconstruct the output waveform, the maximum fundamental output frequency of a DDS is $f_C/2$. However, for practical applications, the output frequency is limited to somewhat less than that, improving the quality of the reconstructed waveform and permitting filtering on the output.

When generating a constant frequency, the output of the phase accumulator increases linearly, so the analog waveform it generates is inherently a ramp.

Then how is that linear output translated into a sine wave?

A phase-to-amplitude lookup table is used to convert the phase-accumulator's instantaneous output value (28 bits for AD9833)—with unneeded less-significant bits eliminated by truncation—into the sine-wave amplitude information that is presented to the (10-bit) D/A converter. The DDS architecture exploits the symmetrical nature of a sine wave and utilizes mapping logic to synthesize a complete sine wave from one-quarter-cycle of data from the phase accumulator. The phase-to-amplitude lookup table generates the remaining data by reading forward then back through the lookup table. This is shown pictorially in Figure 5.

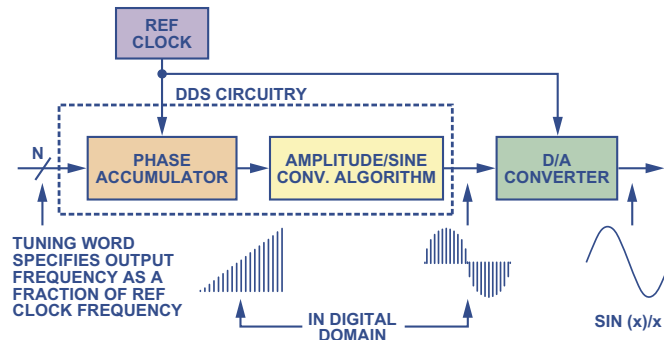


Figure 5. Signal flow through the DDS architecture.

What are popular uses for DDS?

Applications currently using DDS-based waveform generation fall into two principal categories: Designers of communications systems requiring agile (i.e., immediately responding) frequency sources with excellent phase noise and low spurious performance often choose DDS for its combination of spectral performance and frequency-tuning resolution. Such applications include using a DDS for modulation, as a reference for a PLL to enhance overall frequency tunability, as a local oscillator (LO), or even for direct RF transmission.

Alternatively, many industrial and biomedical applications use a DDS as a programmable waveform generator. Because a DDS is digitally programmable, the phase and frequency of a waveform can be easily adjusted without the need to change the external components that would normally need to be changed when using traditional analog-programmed waveform generators. DDS permits simple adjustments of frequency in real time to locate resonant frequencies or compensate for temperature drift. Such

applications include using a DDS in adjustable frequency sources to measure impedance (for example in an impedance-based sensor), to generate pulse-wave modulated signals for micro-actuation, or to examine attenuation in LANs or telephone cables.

What do you consider to be the key advantages of DDS to designers of real-world equipment and systems?

Today's cost-competitive, high-performance, functionally integrated DDS ICs are becoming common in both communication systems and sensor applications. The advantages that make them attractive to design engineers include:

- digitally controlled micro-hertz frequency-tuning and sub-degree phase-tuning capability,
- extremely fast *hopping speed* in tuning output frequency (or phase); phase-continuous frequency hops with no overshoot/undershoot or analog-related loop settling-time anomalies,
- the digital architecture of DDS eliminates the need for the manual tuning and tweaking related to component aging and temperature drift in analog synthesizer solutions, and
- the digital control interface of the DDS architecture facilitates an environment where systems can be remotely controlled and optimized with high resolution under processor control.

How would I use a DDS device for FSK encoding?

Binary *frequency-shift keying* (usually referred to simply as FSK) is one of the simplest forms of data encoding. The data is transmitted by shifting the frequency of a continuous carrier to one of two discrete frequencies (hence *binary*). One frequency, f_1 , (perhaps the higher) is designated as the mark frequency (binary one) and the other, f_0 , as the *space* frequency (binary zero). Figure 6 shows an example of the relationship between the mark-space data and the transmitted signal.




Figure 6. FSK modulation.

This encoding scheme is easily implemented using a DDS. The DDS frequency tuning word, representing the output frequencies, is set to the appropriate values to generate f_0 and f_1 as they occur in the pattern of 0s and 1s to be transmitted. The user programs the two required tuning words into the device before transmission. In the case of the AD9834, two frequency registers are available to facilitate convenient FSK encoding. A dedicated pin on the device (FSELECT) accepts the modulating signal and selects the appropriate tuning word (or frequency register). The block diagram in Figure 7 demonstrates a simple implementation of FSK encoding.




Figure 7. A DDS-based FSK encoder.

And how about PSK coding?

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.

Of the schemes to accomplish PSK, the simplest-known as binary PSK (BPSK)—uses just two signal phases, 0 degrees and 180 degrees. BPSK encodes 0° phase shift for a logic 1 input and 180° phase shift for a logic 0 input. 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 reverses (changes by 180 degrees), then the signal state changes (from low to high, or from high to low).

PSK encoding is easily implemented with DDS ICs. 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, thus generating a PSK output signal. For applications that require high-speed modulation, the AD9834 allows the preloaded phase registers to be selected using a dedicated toggling input pin (PSELECT), which alternates between the registers and modulates the carrier as required.

More sophisticated 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 or QPSK), the possible phase angles are 0, +90, -90, and 180 degrees; each phase shift can represent two signal elements. The AD9830, AD9831, AD9832, and AD9835 provide four phase registers to allow complex phase modulation schemes to be implemented by continuously updating different phase offsets to the registers.

Can multiple DDS devices be synchronized for, say, I-Q capability?

It is possible to use two single DDS devices that operate on the same master clock to output two signals whose phase relationship can then be directly controlled. In Figure 8, two AD9834s are programmed using one reference clock, with the same reset pin being used to update both parts. Using this setup, it is possible to do I-Q modulation.




Figure 8. Multiple DDS ICs in synchronous mode.

A *reset* must be asserted after power-up and prior to transferring any data to the DDS. This sets the DDS output to a known phase, which serves as the common reference point that allows synchronization of multiple DDS devices. When new data is sent simultaneously to multiple DDS units, a coherent phase relationship can be maintained, and their relative phase offset can be predictably shifted by means of the phase-offset register. The AD9833 and AD9834 have 12 bits of phase resolution, with an effective resolution of 0.1 degree. [For further details on synchronizing multiple DDS units, please see Application Note AN-605.]

What are the key performance specs of a DDS based system?

Phase noise, jitter, and spurious-free dynamic range (SFDR).

Phase noise is a measure (dBc/Hz) of the short-term frequency instability of the oscillator. It is measured as the single-sideband noise resulting from changes in frequency (in decibels below the amplitude at the operating frequency of the oscillator using a 1-Hz bandwidth) at two or more frequency displacements from the operating frequency of the oscillator. This measurement has particular application to performance in the analog communications industry.

Do DDS devices have good phase noise?

Noise in a sampled system depends on many factors. Reference-clock jitter can be seen as phase noise on the fundamental signal in a DDS system; *and phase truncation* may introduce an error level into the system, depending on the code word chosen. For a ratio that can be exactly expressed by a truncated binary-coded word, there is no truncation error. For ratios requiring more bits than are available, the resulting phase noise truncation error results in spurs in a spectral plot. Their magnitudes and distribution depends on the code word chosen. The DAC also contributes to noise in the system. DAC quantization or linearity errors will result in both noise and harmonics. Figure 9 shows a phase noise plot for a typical DDS device—in this case an AD9834.

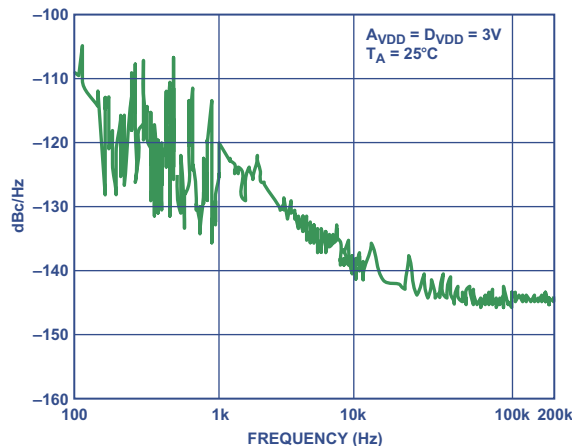


Figure 9. Typical output phase noise plot for the AD9834. Output frequency is 2 MHz and M clock is 50 MHz.

What about jitter?

Jitter is the dynamic displacement of digital signal edges from their long-term average positions, measured in degrees rms. A perfect oscillator would have rising and falling edges occurring at precisely regular moments in time and would never vary. This, of course, is impossible, as even the best oscillators are constructed from real components with sources of noise and other imperfections. A high-quality, low-phase-noise crystal oscillator will have jitter of less than 35 picoseconds (ps) of period jitter, accumulated over many millions of clock edges

Jitter in oscillators is caused by thermal noise, instabilities in the oscillator electronics, external interference through the power rails, ground, and even the output connections. Other influences include external magnetic or electric fields, such as RF interference from nearby transmitters, which can contribute jitter affecting the oscillator's output. Even a simple amplifier, inverter, or buffer will contribute jitter to a signal.

Thus, the output of a DDS device will add a certain amount of jitter. Since every clock will already have an intrinsic level of jitter, choosing an oscillator with low jitter is critical to begin with. Dividing down the frequency of a high-frequency clock is one way to reduce jitter. With frequency division, the same amount of jitter occurs within a longer period, reducing its percentage of system time.

In general, to reduce essential sources of jitter and avoid introducing additional sources, one should use a stable reference clock, avoid using signals and circuits that slew slowly, and use the highest feasible reference frequency to allow increased oversampling.

Spurious-Free Dynamic Range (SFDR) refers to the ratio (measured in decibels) between the highest level of the fundamental signal and the highest level of any spurious signal—including aliases and harmonically related frequency components—in the spectrum. For the very best SFDR, it is essential to begin with a high-quality oscillator.

SFDR is an important specification in an application where the frequency spectrum is being shared with other communication channels and applications. If a transmitter's output sends spurious signals into other frequency bands, they can corrupt or interrupt neighboring signals.

Typical output plots taken from an AD9834 (10-bit DDS) with a 50-MHz master clock are shown in Figure 10. In (a), the output frequency is exactly 1/3 of the master clock frequency (MCLK). Because of the judicious choice of frequencies, there are no harmonic frequencies in the 25-MHz window, aliases are minimized, and the spurious behavior appears excellent, with all spurs at least 80 dB below the signal (SFDR = 80 dB). The lower frequency setting in (b) has more points to shape the waveform (but not enough for a really clean waveform), and gives a more realistic picture; the largest spur, at the second-harmonic frequency, is about 50 dB below the signal (SFDR = 50 dB).

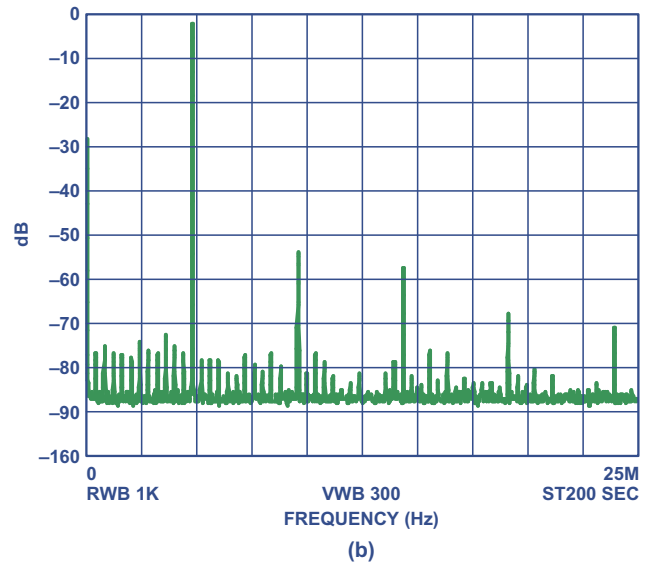



Figure 10. Output of an AD9834 with a 50-MHz master clock and (a) $f_{OUT} = 16.667$ MHz (i.e., $MCLK/3$); (b) $f_{OUT} = 4.8$ MHz.

Do you have tools that make it easier to program and predict the performance of the DDS?

The on-line *interactive design tool* is an assistant for selecting tuning words, given a reference clock and desired output frequencies and/or phases. The required frequency is chosen, and idealized output harmonics are shown after an external reconstruction filter has been applied. An example is shown in Figure 11. Tabular data is also provided for the major images and harmonics.

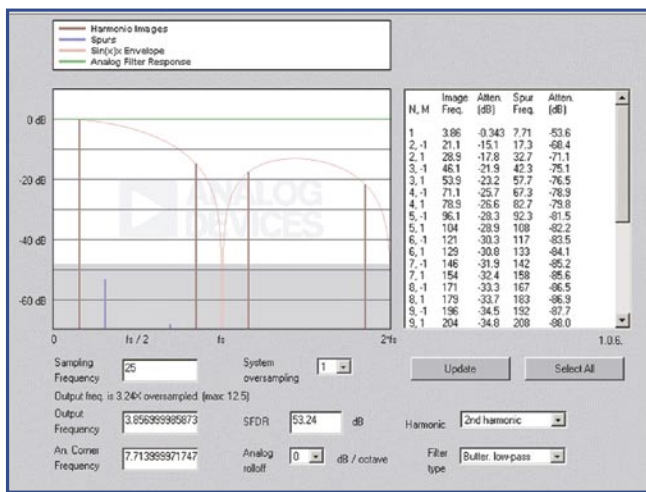


Figure 11. Screen presentation provided by an interactive design tool. A $\sin(x)/x$ presentation of a typical device output.

How will these tools help me program the DDS?

All that's needed is the required frequency output and the system's reference clock frequency. The design tool will output the full programming sequence required to program the part. In the example in Figure 12, the MCLK is set to 25 MHz and the desired output frequency is set to 10 MHz. Once the update button is pressed, the full programming sequence to program the part is contained in the Init Sequence register.

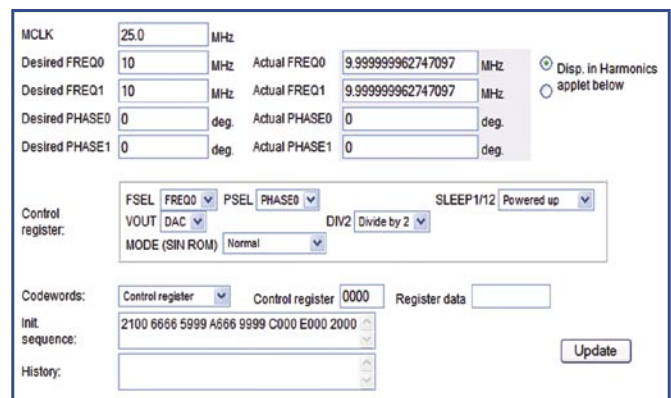


Figure 12. Typical display of programming sequence.

How can I evaluate your DDS devices?

All DDS devices have an evaluation board available for purchase. They come with dedicated software, allowing the user to test/evaluate the part easily within minutes of receiving the board. A technical note accompanying each evaluation board contains schematic information and shows best recommended board-design and layout practice.

Where can I find more information on DDS devices?

The main DDS homepage is located at www.analog.com/dds

Links to design tools are provided at http://www.analog.com/Analog_Root/static/techSupport/interactiveTools/#dds

An in-depth tutorial on DDS technology can be found at http://www.analog.com/UploadedFiles/Tutorials/450968421DDS_Tutorial_rev12-2-99.pdf

AN-605 can be found at http://www.analog.com/UploadedFiles/Application_Notes/3710928535190444148168447035AN605_0.pdf

The latest DDS selection guide can be found at http://www.analog.com/IST/SelectionTable/?selection_table_id=27