

Application Note #5

Direct Digital Synthesis—Impact on Function Generator Design

Introduction

Function generators have been around for a long while. Over time, these instruments have accumulated a long list of features. Starting with just a few knobs for setting the amplitude and frequency of a sinusoidal output, function generators now provide wider frequency ranges, calibrated output levels, a variety of waveforms, modulation modes, computer interfaces, and, in some cases, arbitrary functions.

The many features added to function generators have complicated their design and increased their cost. There is an opportunity for a radical re-design of the familiar function generator using direct digital synthesis (DDS).

DDS provides remarkable frequency resolution and allows direct implementation of frequency, phase and amplitude modulation. These features which were 'tacked-on' to function generators are handled in a clean, fundamental way by DDS.

Direct Digital Synthesis

Many of the concepts of DDS are illustrated by the way in which a sine wave is generated. Figure 1 shows a block diagram of a simple DDS function generator. The

sine function is stored in a RAM table. The RAM's digital sine output is converted to an analog sine wave by a DAC. The steps seen at the DAC output are filtered by a lowpass filter to provide a clean sinewave output.

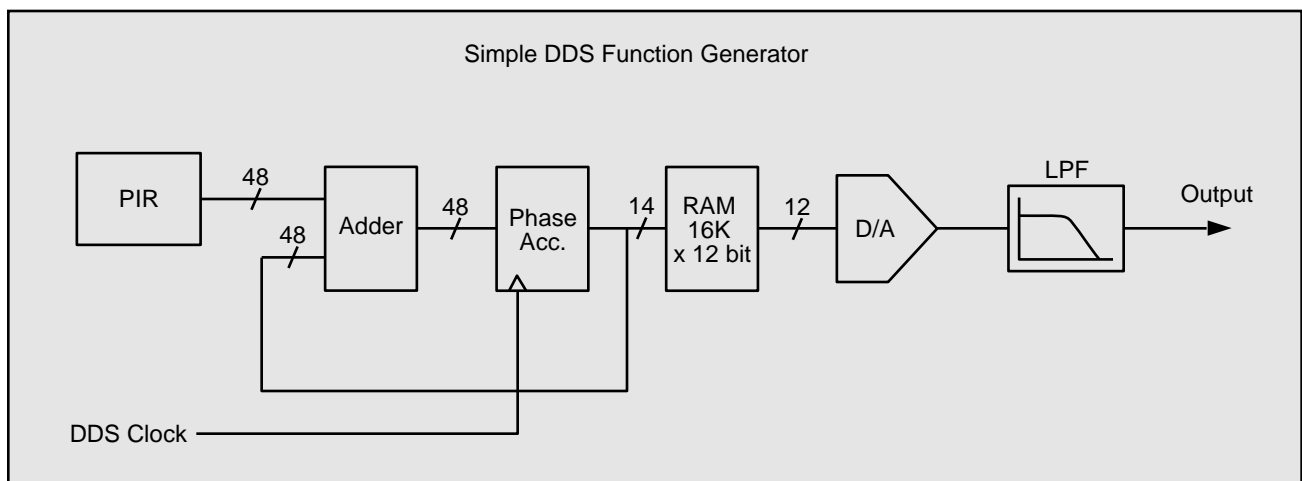
The frequency of the sine wave depends on the rate at which addresses to the RAM table are changed. Addresses are generated by adding a constant stored in the phase increment register (PIR) to the phase accumulator. Usually, the rate of additions is constant, and the frequency is changed by changing the number in the PIR.

The frequency resolution depends on the number of bits in the PIR. If the PIR, adder, and phase accumulator support 48-bit additions, then the fractional frequency resolution is one part in 247, or about one part in 1014. That means a 48-bit DDS function generator can provide better than 1 μ Hz resolution on a 10 MHz output.

Some Details

There are a few more details which need to be addressed in order to understand DDS in this application. Questions about sample rate, RAM size, DAC resolution, filter characteristics, and spectral purity of the output must be answered.

Figure 1. Simple DDS Block Diagram



Samples per Cycle

Our intuition might suggest that a large number of samples are required for each cycle of the sine wave to achieve good spectral purity of the output. A sketch of a sine which is approximated by a small number of samples hardly looks like a sine wave. Remarkably, only about 3 samples are required during each cycle. In fact, if we could make an arbitrarily sharp low-pass filter, we would need only two samples per cycle.

To motivate this, consider the case where we have four samples per sine cycle. This situation is shown in Figure 2. The sampled sine is reduced to a pulse train (or a square wave, if we started sampling at 45 degrees instead of at 0 degrees).

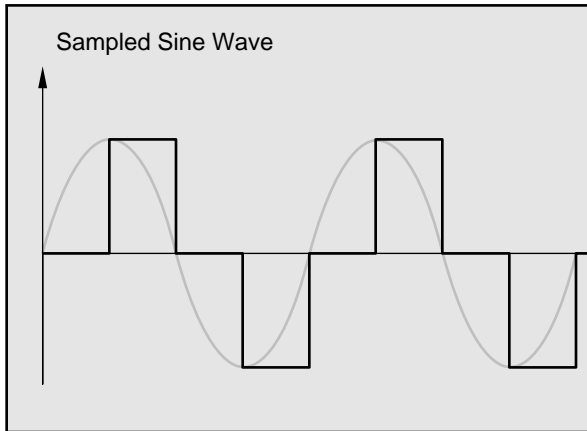


Figure 2. Sampling a Sine Four Times per Cycle.

The Fourier spectra of this pulse train has components at f , $2f$, $3f$, etc. If we can arrange the lowpass filter to eliminate the harmonic components of the pulse train, then we are left with the fundamental, a pure sine wave at frequency f .

In the more general case, generating an output at f by sampling at a rate of f_s , the lowest frequency Fourier component at a frequency of $f_s - f$. This simple result becomes the basis of the lowpass filter specification: the filter should pass f but stop $f_s - f$.

Filters

Figure 3 shows a lowpass filter transfer function. As we have seen, the filter must pass the highest frequency which we wish to generate (f_{max}), but must begin their stop-band at $f_s - f_{max}$. Steep rolloff filters with high stop-band attenuation are hard to build. A reasonable compromise in this trade-off occurs when $f_{max} = f_s / 3$. This allows the filter a one octave transition band.

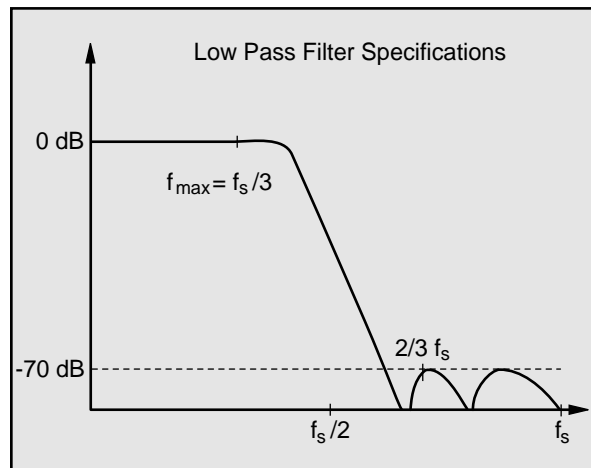


Figure 3. Lowpass Filter for DDS Outputs.

What stopband attenuation is needed? This depends on the spurious component specification of the output. A typical specification for a function generator application would be -70 dBc.

Cauer (elliptic) filters are a good choice for this application. They have fast transition bands, and may be designed with very low ripple in the pass-band. The specification for this example is met by a ninth degree Cauer filter.

Bessel Filters

While Cauer filters are the best choice for CW applications, they are unusable for arbitrary waveform generation. In the time domain, Cauer filters have a very nasty overshoot. A much better choice for arbitrary waveforms (or ramps and triangles) is the Bessel filter. The Bessel filter has a slower rolloff when compared to the Cauer filter, but it is nearly phase-linear. The lack of

dispersion in a phase-linear filter will preserve the pulse shape and prevent any ringing in the time domain. A seventh degree Bessel filter with a -3 dB cutoff of $f_c = f_s/4$ is a good choice for filtering arbitrary waveforms. This filter will exhibit an output risetime of $0.35/f_c$.

DAC and RAM Requirements

Big, fast RAMs and high speed, high resolution DACs have made DDS a viable technology for function generator applications. How big, how fast, and what resolutions are required?

As we have seen, a maximum practical output frequency is $f_s/3$. So the DDS phase accumulator, RAMs, and DACs must run at three times the maximum desired output frequency.

The DAC resolution depends on the spurious component specification for the output (or the desired arbitrary waveform resolution). The DAC's quantization error and non-linearities lead to spurious outputs. To get a rough idea of the magnitude of the spurious frequency component, realize that the difference between the actual output of the DAC and the desired sine value is the source of these spurious output components. So a 12-bit DAC which is linear and monotonic to 2 LSBs will have output errors on order one part in 2048, or about -66 dB.

A short RAM table is another way to get the wrong value out of the DAC. To avoid 'phase quantization noise', there should be two more bits of address to the RAM than bits in the DAC.

Extending Frequency Range

The frequency range of the DDS output may be extended by a variety of techniques. Depending on which technique is used, some of the advantages of DDS may be lost. Just as with more conventional frequencies synthesizers, the DDS output may be doubled, mixed with other fixed sources, or used as a reference inside of a phase locked loop.

Modulation Techniques

The power and elegance of DDS are most apparent when a modulated source is required. The frequency of the output may be changed instantly to any frequency

from dc to f_{max} by simply changing the number in the phase increment register. Figure 4 shows the block diagram of a DDS phase accumulator with programmable modulation capabilities.

This phase accumulator, which has been optimized for function generator applications, has two phase increment registers, PIRA and PIRB. A 48-bit wide multiplexer can switch between the PIRs in a single clock. The modulation processor can modify the PIRs at a rate of up to 10 million bytes per second, filling one PIR while the other is used as an input to the adder.

Complex modulation programs may be stored in the modulation RAM. This RAM contains op-codes and data for the modulation processor. Frequency scans illustrate the operation of this processor. When programmed for a log frequency sweep, a list of up to 4000 discrete frequencies are stored in the modulation RAM by the host system. The modulation processor modifies PIRA while the adder is using PIRB and vice-versa.

More complex modulation programs may be stored, such as frequency modulation by any arbitrary function, linear or log sweeps, frequency hopping, etc. Phase modulation is easily done by programming PIRA with the nominal frequency, and using PIRB, which contains the nominal phase increment plus any desired phase shift, for a single clock cycle.

Wide frequency or phase deviations are no problem. Any phase or frequency hop may be programmed and executed in a single clock. And, since the PIRs may be modified very quickly, modulation frequencies up to several hundred kHz are possible.

In fact, arbitrary modulation programs may be stored. This feature allows the function generator to be used for modem testing, frequency agile communications, bit error rate determination, etc.

Amplitude Modulation

There are two approaches for amplitude modulation of the output waveform. Either the digital outputs from the RAM or the analog output from the DAC may be multiplied by the desired amplitude. The later approach is better for function generators so that either an internal or external source may be used for amplitude modulation.

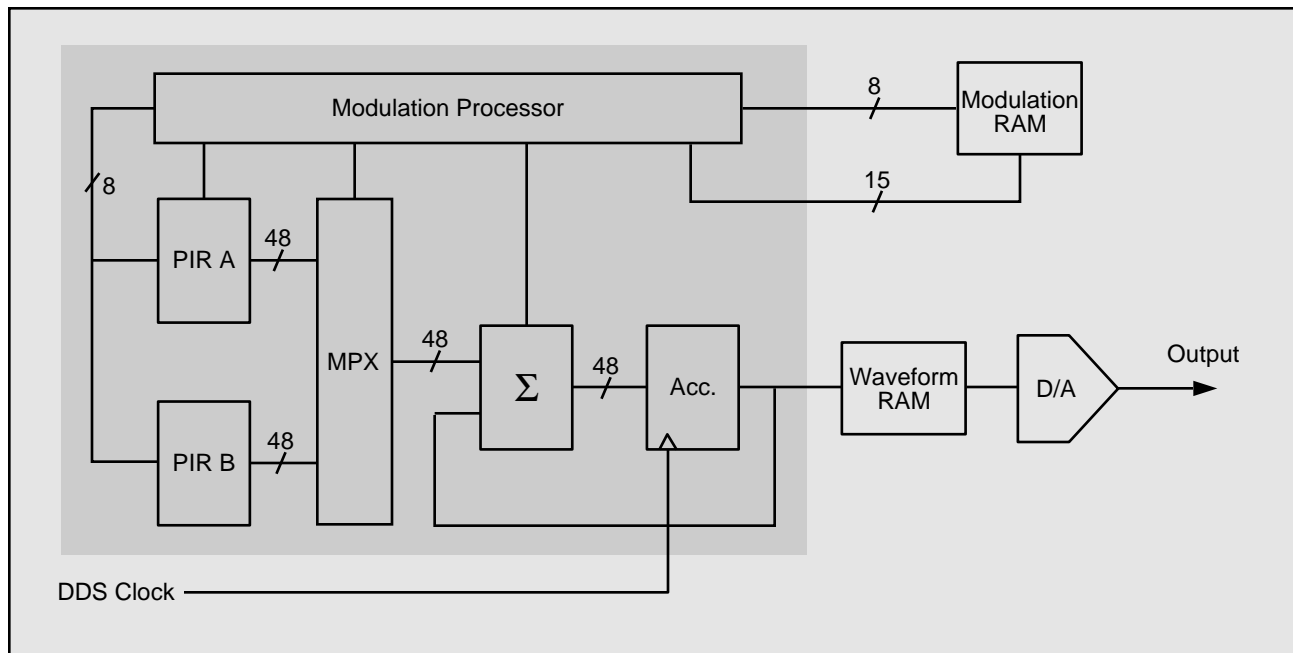


Figure 4. DDS Phase Accumulator with Modulation Processor.

Arbitrary Functions

One of the immediate benefits of the DDS architecture is that arbitrary waveform generation comes along for free. Instead of storing a sine table in the waveform RAM, a list of arbitrary values is saved. The phase accumulator is programmed to step through the stored values, one at a time, to play back the desired waveform through the output DAC.

The DDS's arbitrary waveform capability simplifies the task of generating the other 'standard' waveforms found in function generators. Ramp, sawtooth and even Gaussian white noise may be generated by changing the list of values in the waveform RAM.

The phase accumulator must be designed to support certain modes required for arbitrary waveforms. The rate at which RAM values are retrieved may be changed by simply using a different PIR value. However, variable record lengths, triggering functions, and wrap-around addressing are unique to arbitrary function generation.

As previously mentioned, a Bessel filter is required for arbitrary waveform generation. The Bessel filter will smooth the steps at the DAC output. With a -3 dB cutoff frequency, f_c , of $f_s/4$, the output will show a controlled risetime of $0.35/f_c$ without overshoot.

Square Waves

Square waves are a special case for the DDS. One might think that a square wave could be generated by loading +1 and -1 into the waveform RAM. Indeed they can, but with the unfortunate restriction that the square-wave edges must be synchronous with the DDS sample clock. This restriction would greatly limit the resolution of available frequencies, especially at high frequencies.

A much better approach for generating squarewaves is to generate a clean sinewave, then discriminate the sine into a square wave. In this way, square waves will have the same frequency range and resolution as sine waves.

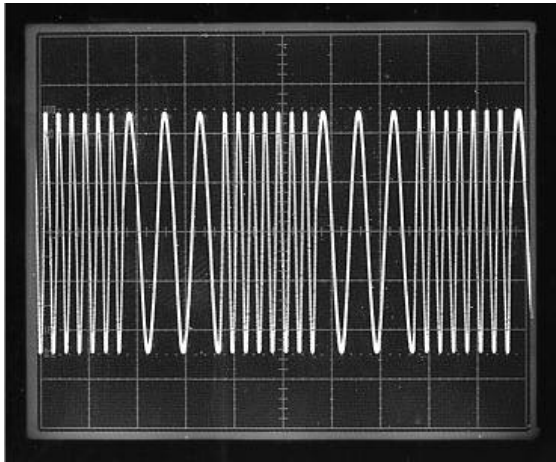


Figure 5. Frequency Shift Keying of Sinewave.

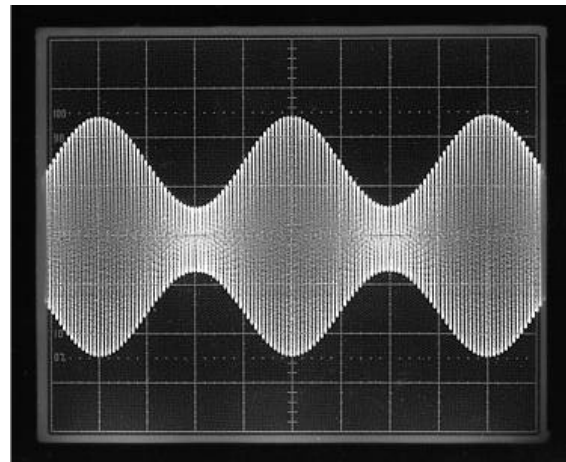


Figure 6. Amplitude modulation of Sine by Sine.

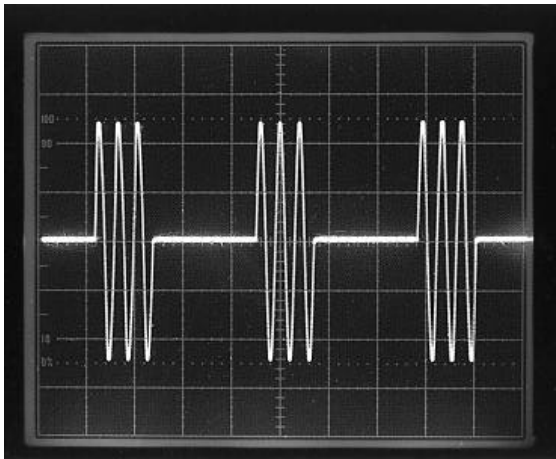


Figure 7. Three Cycle Burst of Sines

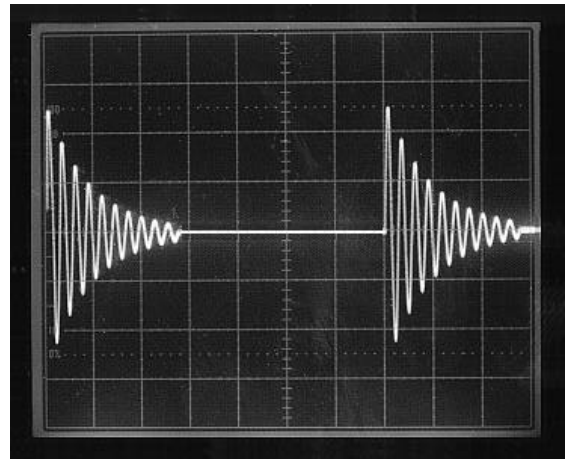


Figure 8. Arbitrary Waveforms.

Output Amplifiers

The output amplifier used in a DDS function generator must meet some stringent requirements. In order to preserve waveforms generated in the arbitrary mode, the amplifier must have a wide and flat passband, and exhibit a phase linear response well past the cutoff frequency of the Bessel filter.

The amplifier's bandwidth also determines the risetime of the squarewave output. Here again, a well behaved (phase linear) rolloff is required to prevent overshoot on the squarewave output.

Finally, the output amplifier must be able to drive 10 V_{pp} into a 50Ω load, meet distortion and settling specifications, and be protected against short circuits or con-

nection to external power supplies. The output amplifier should exhibit a 50Ω output impedance regardless of output level setting.

To generate low signal levels, most function generators have output attenuators. The attenuators allow the output amplifier to work within a limited range of output levels, so that distortion and signal-to-noise ratios remain constant as the output levels are changed.

Floating Generator

Many applications require that the function generators be able to provide a signal to a load which is not ground referenced. Even if the load is nominally ground referenced, a floating generator output will provide a much cleaner signal because system ground loops are

eliminated. It is important that the generator output shield is floating under all circumstances, even when the function generator is connected to a GPIB controller, or if an external frequency reference is connected to the instrument.

ASICs

DDS provides a new, clean, design approach for function generators. Much of the analog 'baggage' required for function generators is handled by digital logic circuits. Unfortunately, these logic circuits are big, complicated, and have to run fast. For example, a 15 MHz DDS requires a 48-bit adder operating at 40 MHz, with lots of glue logic. Fortunately, application specific integrated circuits (ASICs) provide a low cost solution to the problem.

A TTL prototype of the phase accumulator diagrammed in Figure 4 required about 150 ICs. The prototype was just able to work with a clock of 10 MHz. A CMOS gate array of the same design was fabricated in a 68-pin PLCC plastic package. The gate array operates at 40 MHz (worst case), uses about 1/4 watt of power, and has a recurring cost of about \$10.

Conclusions

DDS based function generators are just beginning to appear in the market. These function generators offer substantial performance improvements, at reduced costs, over conventional analog function generators. As the cost of ASICs, RAMs and DACs decline, while their speed and resolution increase, expect to see DDS based function generators soon replace their analog counterparts.