INTRODUCTION TO DIRECT DIGITAL SYNTHESIS

Technical Staff
March, 1990
Revised June, 1991

Basic Theory

Perhaps the easiest way to understand how DDS systems work is to explain the system from the output back toward the frequency control input. We will begin with the familiar analog RF signal and work backward through the digital system. The digital information provided to the DAC must represent the instantaneous amplitude value of the RF waveform. The digital section of a DDS system is concerned with generating this sinusoid oscillating digital number. The most efficient and flexible known method of generating this oscillating amplitude number is to first digitally define a frequency number. The frequency number is then digitally integrated in an accumulator to provide instantaneous digital phase information. Finally, the digital phase word is converted to the digital amplitude word in a Look-up Table (LUT). Figure 1 shows a basic DDS system.

The usual desired output from a synthesizer is a sinusoidal waveform. However, any waveform can be generated using DDS provided that the highest Fourier component is within the Nyquist limitation ($^1/_2$ the clock frequency). This discussion will concentrate on the sinusoidal waveform, however. For any repetitive waveform there is a one-for-one correspondence between values of phase and values of amplitude. This one-for-one conversion is easily realized using a ROM. Within the ROM is stored a "look-up table". The phase information is applied to the address bus and the amplitude information appears on the data bus. Each discrete phase point corresponds to a discrete amplitude value. The relationship is a trigonometric function, in this case a sine or cosine. The look-up table can be considered a digital phase-to-amplitude converter. Figure 2 shows the relationship between address and data bus numbers in a typical look-up table ROM. All the sinusoidal amplitude information necessary for a complete 360° cycle is contained in 90° of information. Therefore, only 90° of mapping is required in the look-up table. The phase accumulator simply clocks in both directions and then reverses sign of the amplitude for the other 180°. This allows more economy in the look-up table ROM. All Stanford Telecom NCO families use 90° look-up tables.

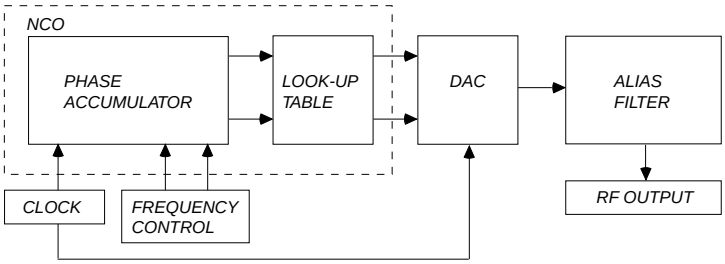

In a DDS System all the parameters of the waveform: Frequency, Phase, & Amplitude are defined by digital words

The first Analog Signal appears at the DAC output

Figure 1. Basic DDS System

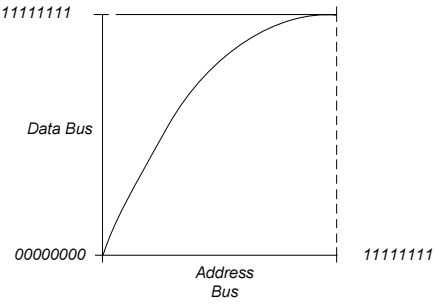

Look up table showing the 90° sinusoid relationship between the address and data buses

Figure 2. 90° Look-up Table

For a constant frequency ω, dϕ/dt will be constant. The phase change of a constant frequency signal is a linear progression. Digital accumulators are excellent generators of linear progressions of digital numbers. Indeed, the linearity of the time progression is simply dependent upon the time integrity of the clock. Different ω values require different slopes of phase (dϕ/dt) output. This is achieved by adding different values to the phase accumu-

lator each cycle of the clock. The lowest frequency that can be generated occurs when just 1 LSB (least significant bit) is added to the accumulator which then increments through all possible states, one at a time. This will take 2^N clock cycles for an N-bit accumulator, resulting in an output frequency of $f_{clk}/ 2^N$. It is easy to see that adding any number M to the accumulator each cycle of the clock results in the accumulator count incrementing M times as fast, so that the frequency of the sinewave generated will be $M^* f_{clk}/ 2^N$. Since the number M represents the phase increment in each cycle it is generally referred to as Delta-Phase ($\Delta\phi$). Thus the output frequency of an NCO is represented by the equation:

$$f_{out} = \frac{\Delta\phi * f_{clk}}{2^N}$$

A typical value for N in practice is 32, giving a resolution of 1 part in 2^{32}, or $4.29 * 10^9$. This translates to a resolution of 12 milli-Hz at a clock frequency of 50 MHz or 0.23 Hz at a clock frequency of 1 GHz.

All three parameters of the waveform; frequency, phase, and amplitude (ω, ϕ, and A) are defined by digital words. Frequently the resolution of these three parameters is confused. ω resolution is determined by the number of bits in the phase accumulator. ϕ resolution is determined by the number of bits in the ROM look-up table. Amplitude resolution is determined by the resolution of the DAC. A frequently asked question is; how can you achieve 48 bit frequency resolution using only an 8 bit DAC? The answer is that the DAC is simply approximating the ideal amplitude value for a given clock cycle. The resolution of the DAC will determine the amplitude accuracy of the desired waveform, not the frequency of the waveform. A one bit DAC can be represented by the MSB output of the phase accumulator. The output of the MSB will be a rectangular waveform. The spectrum of this digital waveform shows the most prevalent component to be the desired output frequency! As might be expected, as more bits are added, the output spectrum becomes progressively "cleaner". The frequency of this most prevalent spectral component is adjustable with 2^N resolution although its amplitude resolution is only one-bit.

Comparison with PLL Synthesizers

The technique of direct digital synthesis (DDS) has recently become a viable alternative or complement to phase-locked loop (PLL) synthesis for a wide range of applications. The DDS technique circumvents the traditional trade-offs associated with PLL architectures. Rather, a new set of trade-offs appear:

Trade-offs for PLL Synthesizers

1. Frequency resolution (step size)
2. Settling time
3. Tuning range (bandwidth)
4. Phase noise (spectral purity)
5. Cost, complexity and power consumption

Trade-offs for DSS Synthesizers

1. Clock speed (bandwidth)
2. Spurious responses (spectral purity)
3. Cost, complexity and power consumption

A comparison of these two trade-off lists implies a preference for one architecture or the other for a given application or specific design criteria. In DDS systems the frequency resolution is determined primarily by the word length of the phase accumulator. Furthermore, in DDS systems the settling time is usually determined by the bandwidth of the alias filter. Consequently, frequency resolution, settling time, and spectral purity are independent variables, unlike PLLs where they are interrelated.

It is important to distinguish between settling time and latency. Settling time is defined as the time a signal settles on a new frequency after being switched from an initial frequency. Latency is defined as the number of clock cycles from when a new value of $\Delta\phi$ is loaded into the NCO until the output signal switches to the corresponding frequency. Latency will be determined by the architecture used within the NCO device. The actual latency time can be calculated by multiplying the period of the clock by the latency clock cycles. A few more clock cycles must be added to account for delays in the DAC and other possible signal processing steps, such as multipliers used for AM.

In DDS systems the important trade-off is between bandwidth and spectral purity. This trade-off can manifest itself in several ways. If the clock speed is reduced, the Nyquist frequency is consequently reduced, and thus the bandwidth is also reduced. Lower clock frequencies allow for higher resolution DACs which can provide better spectral purity. Alternatively, higher clock speeds can be used in conjunction with frequency dividers. Dividing the DDS output frequency will reduce the bandwidth but also improve

spectral purity. Both DDS and PLL system performance can be improved by careful design, higher quality components, using more power to achieve higher speeds, and of course, spending more money.

DDS should not be regarded as a direct replacement for PLL circuits. Rather, both techniques have inherent advantages and disadvantages in specific systems. DDS systems usually require relatively high power consumption, with almost one watt being a state-of-the-art minimum. However, there are many specific synthesizer design objectives which render the DDS architecture quite attractive. Some of these include:

1. Minimum settling time (submicrosecond)
2. Multi-octave operation
3. Fine or very fine frequency resolution (<1μHz is possible)
4. Phase continuous frequency changes
5. Exceptionally linear analog or data modulation (16-bit typical)
6. Exceptionally linear sweep and chirp modulation (16-bit typical)
7. Quadrature generation over multiple octaves

Some applications where these synthesizer attributes are important include:

1. Frequency hopping and spread spectrum radios
2. EW and jammer systems
3. Doppler and chirp radars
4. High speed and/or high resolution PSK or FSK
5. Radio and television broadcast equipment (also HDTV)
6. Test equipment
7. MRI (magnetic resonance imaging)

Although DDS has traditionally been associated only with high cost military systems, today a wide range of commercial applications can be considered. The level of performance, both speed and purity, is being paralleled by a drop in prices. These simultaneous trends in performance and price are opening up a multitude of new applications for this fascinating technique.

Spectral Purity in DDS Systems

Phase noise, for practical purposes, is not a problem in DDS systems. The digital phase number is exceedingly linear because it is produced by a stable fixed clock. This is a major advantage over PLL systems. Although phase noise is typically not an issue when discussing DDS systems, the discrete spurious signals are significant. Therefore, spectral impurities in DDS systems are usually discrete narrow band spurs rather than broad band phase noise as in PLL systems.

Spurious responses arise from two levels; quantized errors and DAC errors. Since the DAC's actual output is a quantized approximation of an ideal amplitude value, quantization errors result. Therefore even if a "perfect" DAC existed, spurious signals would appear in the output. Additionally, because the DAC is never "perfect", DAC errors are also introduced. The DAC and quantization errors show themselves as harmonically related spurs. Predicting the location of the spurs is relatively simple. Predicting the amplitude of these spurs is more difficult. If a perfect DAC is assumed, a Fourier analysis can be performed to predict the quantized limitation of a given system. For a detailed discussion on alias and spurious responses in DDS systems, AN102 can be consulted.

Digital Modulation

Since ω, ϕ, and A are all defined by digital words at various locations within the system, we can effect digital modulation by changing these digital words in accordance to some digitized modulating waveform. FM can be realized by changing the frequency control word. PM can be realized by changing the instantaneous phase word. AM can be realized by placing a digital multiplier between the look-up table and the DAC. A four quadrant multiplier will yield a carrier suppressed double sideband signal. A single quadrant multiplier will yield a full carrier double sideband signal. Any combination of FM, PM, and/or AM can be realized by simultaneous modulation of the applicable words. Since FM, PM, and AM represent the only three known techniques of waveform modulation, any modulation scheme known can be realized using DDS techniques. These statements hold true for either digitized analog modulation signals or digital modulation schemes proper. A more detailed treatment of numeric modulation is presented in AN103.

Other DDS Features

DDS offers some other interesting possibilities. Since the phase of the synthesized signal is precisely defined by the look-up table, two look-up table/DAC combinations working from the same phase accumulator will have precise phase relationships. The most common need is quadrature generation. Sine and cosine look-up tables working from the same phase accumulator will provide very precise I and Q outputs, from DC to the Nyquist frequency. The most difficult practical problem for maintaining quadra-

ture lies in the alias filters. It is very difficult to build identical alias filters, even at HF. For example, to maintain 1° phase accuracy at 10MHz, the I and Q alias filters must maintain 270 pS accuracy! The Stanford Telecom STEL-1177 "MNCO" provides I and Q channel outputs while allowing independent 12-bit resolution PM on each channel and/or 16-bit FM to both channels simultaneously. The source of the PM and FM signal sources can be either digital or analog.

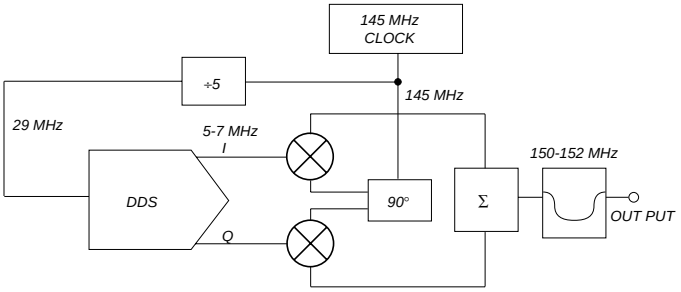

Figure 3. Single-sideband up-conversion scheme to achieve good spurious attenuation

Another common synthesizer requirement is to provide for decimal frequency step sizes from a decimal frequency clock. System clock standards are often 5, 10, or 100MHz. Decimal step sizes are possible when using binary phase accumulators, but they require a clock whose frequency is a binary multiple of the step size. A better solution is to employ a BCD phase accumulator. This will provide the desired 10, 1, or 0.1Hz resolutions from a 10, 50, or 80MHz clock, for example. The Stanford STEL-1176 uses a BCD counter that provides exact 0.1Hz step sizes over 0-30MHz from an 80MHz clock, which can easily be synthesized from a 10MHz reference.

DDS, PLL, or DDS+PLL?

Before the availability of economical and high performance DDS systems the synthesizer designer was limited to PLL techniques. Today, the synthesizer designer has three options; DDS, PLL, and DDS+PLL. The DDS architecture is a viable architecture for many applications previously given to PLL designs. However, there remain many applications better suited to more traditional PLL techniques. Also, by combining the advantages of both techniques, many of the advantages of both systems can produce specifications that were impossible to achieve just a few years ago. For example, a single package DDS system running at 1GHz can now be configured using the Stanford Telecom GaAs STEL-2373.

Good engineering practice mandates finding the optimum design solution within the limitations of technical, economic, time, and manufacturability constraints. Experienced PLL synthesizer designers can quickly determine optimum design architectures from a set of objective specifications. Alternatively, the designer might determine that the objective specifications are not realistic. With the emergence of the DDS system how can the designer know which architecture will be optimum for a set of objective specs? Furthermore, what are these new limitations on synthesizer performance? Can the specs be met? If so, which architecture is optimum, DDS, PLL, or DDS+PLL?

Design intuition can be gained through comparison of PLL and DDS trade-offs. However, several examples of objective specifications contrasted against DDS and PLL possibilities are provided.

Synthesizer 1:

Frequency range: 150-152MHz

Frequency resolution: <1Hz

Settling Time: 1μS

Worst-case discrete spurs: -65dBc

The combination of narrow bandwidth, fine resolution, and fast settling time clearly indicates a DDS architecture. Furthermore, the -65dB spur levels can be realized using lower frequency clocks, perhaps 30MHz. A 2MHz spectrum (perhaps 5 to 7MHz) is selected and then a mixer is used to upconvert to the 150-152 MHz band (Figure 3).

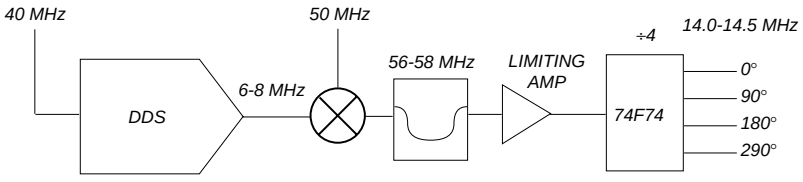

Figure 4. Low-spurious Quadrature Output Narrow-band DDS Synthesizer

Synthesizer 2:

Frequency range: 300-600MHz

Frequency resolution: 100KHz

Settling time: 100μS

Discrete Spurs: -80dBc

Phase noise: -80dBc/Hz @ 10KHz offset

The combination of wide frequency range, wide resolution, and slow settling time clearly indicates a PLL architecture. Furthermore, the discrete spur specification precludes an unprocessed DDS signal.

Synthesizer 3:

Frequency range: 14.0-14.5MHz

Frequency resolution <10Hz

Settling time: 10μS

Discrete spurs: -80dBc

Phase noise: -100dBc/Hz @ 500Hz offset

I & Q outputs required within 1° accuracy

The combination of range resolution, settling time, phase noise and I & Q certainly point to DDS. However, -80dBc spur performance is probably not achievable from DDS systems clocked at 40MHz or so. -70dBc spurs are possible over 2MHz bandwidths below 25% of the clock frequency, perhaps 6-8MHz. Then the 6-8MHz bandwidth is up-converted to 56-58Mhz. The 56-58MHz signal is then divided by 4. With 74AC74 dual flip-flops, all four quadrants can be derived from a divide-by-4 circuit. Simultaneously, the spurs are attenuated by 12dB through the division process, thus meeting the spurious specification

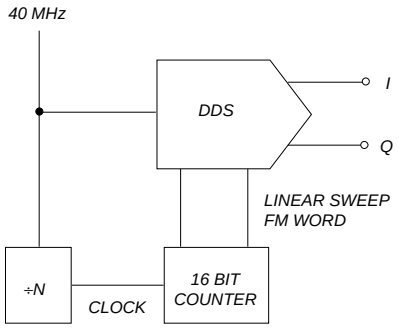

Figure 5. Chirp Generator for Radar System with I&Q Outputs for SSB Up-Conversion

(Figure 4).

Synthesizer 4:

Baseband chirp generator for Radar system

Frequency range: 4-8MHz

Frequency resolution: 16-bit chirp linearity

Settling time: maintain 16-bit chirp linearity

Discrete spurs: -60dBc

I & Q outputs required within 3° accuracy

Clearly this is a DDS design. It is very difficult, if not impossible to achieve 16-bit chirp linearity using analog techniques. 3° phase accuracy is achievable using careful active alias filter design techniques. The balance of the specifications should also be within state-of-the-art DDS capabilities (Figure 5).

Synthesizer 5:

Frequency range: 100-1000MHz

Frequency resolution: <1Hz

Settling time: 10μS

Discrete spurs: -70dBc

Phase noise: -100dBc/Hz @ 1KHz offset

The combination of resolution and settling time requirements eliminate PLLs, while the bandwidth and spur requirements eliminate DDS. However, perhaps the bandwidth of a PLL can be coupled with the resolution and speed of DDS in a DDS/PLL hybrid synthesizer (Figure 6).

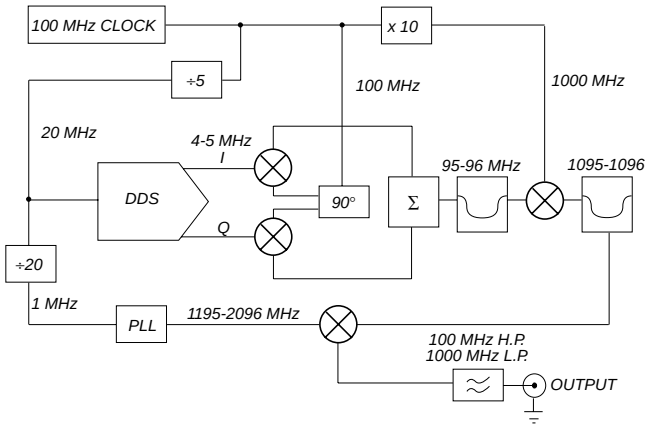

Figure 6. 100-1000 MHz DDS+PLL hybrid using a triple conversion architecture

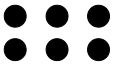

For Further Information Call or Write

INTEL CORPORATION
Cable Network Operation
350 E. Plumeria Drive, San Jose, CA 95134
Customer Service Telephone: (408) 545-9700
Technical Support Telephone: (408) 545-9799
FAX: (408) 545-9888