

Direct Digital Synthesis (DDS) with a Programmable Modulus

by Ken Gentile




Figure 1. Typical Accumulator-Based DDS Architecture

OVERVIEW

The programmable modulus is a modification of the typical accumulator-based DDS architecture. It extends the use of DDS to applications that require exact rational frequency synthesis (any-rate applications, for example). The [AD9913](#) is the first accumulator-based DDS product by Analog Devices, Inc., to offer the programmable modulus architecture.

TYPICAL ACCUMULATOR-BASED DDS

A typical accumulator-based DDS relies on an accumulator to recursively sum the digital input tuning word at the rate of the sample clock (see Figure 1). This produces a time series of digital words at the output of the accumulator that increases linearly until the accumulator rolls over at its maximum value of 2^C . Hence, the accumulator output has a fixed modulus of 2^C .

Usually the accumulator output is truncated to P bits (using only the MSBs) to reduce the size and complexity of the angle-to-amplitude conversion block that immediately follows the accumulator. This causes the time series of digital words produced by the accumulator to appear at the input to the angle-to-amplitude converter as P -bit word(s) ranging in value from zero to $2^P - 1$.

The angle-to-amplitude converter maps the P -bit word(s) to one revolution on the unit circle; that is, it linearly maps binary values from 0 to 2^P to radian angles from 0 to 2π . This mapping arrangement allows the angle-to-amplitude converter to translate the P -bit word(s) to D -bit amplitude value(s) (A) in a very efficient manner.

The translation process relies on the trigonometric relationship

$$x = \sin(2\pi k/2^P)$$

where:

P is the number of bits taken from the accumulator.

k is the binary value of those bits at any given instant.

The value of A for each k is the value of x scaled, offset, and rounded such that A takes on integer values between zero and $2^D - 1$.

It is possible to readily extend the angle-to-amplitude conversion function to perform both sine and cosine conversions concurrently. Therefore, DDSs that offer both conversion schemes are commonly found.

Following the angle-to-amplitude converter is a D -bit DAC. It converts the D -bit digital amplitude values produced by the angle-to-amplitude converter to analog levels. The result is a sinusoidal waveform at the output of the DAC with a frequency determined by the average rollover rate of the accumulator.

The following equation expresses the frequency of the sinusoid that appears at the DAC output for a typical accumulator-based DDS:

$$f_o = \frac{M}{2^C} f_s \quad (1)$$

where:

f_o is the synthesized frequency.

f_s is the sampling frequency.

$M/2^C$ is the fractional scale factor (M and C are positive integers).

Normally, M is constrained to be less than 2^{C-1} . Otherwise, the DDS synthesizes a Nyquist image frequency. The quantity, 2^C , is the modulus of the accumulator, where C is the accumulator width in bits. The integer, M , is often referred to as the frequency tuning word. Because M , by definition, is an integer, f_o is constrained to the following set of frequencies:

$$f_o \in \left\{ 0, \frac{f_s}{2^C}, \frac{2f_s}{2^C}, \frac{3f_s}{2^C}, \dots, \frac{(2^{C-1}-1)f_s}{2^C} \right\} \quad (2)$$

Thus, the set of frequencies from dc to almost $\frac{1}{2}f_s$ comprises equally spaced intervals of $f_s/2^C$. This spacing, $f_s/2^C$, is referred to as the frequency resolution of the DDS. Inspection of the f_o frequency set indicates that the modulus of the DDS accumulator, namely 2^C , determines both the frequency resolution of the DDS and the number of possible output frequencies. For example, the output frequency set of a DDS with a 32-bit accumulator has 2,147,483,648 (that is, 2^{31}) possible output frequencies for a given sample rate (f_s).

Even with such a large set of available frequencies, the typical accumulator-based DDS is not capable of generating some useful frequencies (like precisely $f_s/10$) because the accumulator modulus is fixed and the tuning word (M) must be an integer. Rearranging Equation 1 reveals that the ratio of the output frequency to the sampling frequency must satisfy

$$\frac{f_o}{f_s} = \frac{M}{2^C} \quad (3)$$

Now for output frequencies that are an integer submultiple of the sample rate (for example, $f_s/10$), f_o can be expressed as $f_o = f_s/Q$ (Q is an integer). Substituting f_s/Q for f_o in Equation 3 leads to

$$\frac{1}{Q} = \frac{M}{2^C} \quad (4)$$

Solving for M yields $M = 2^C/Q$. Because M and Q must both be integers, the only values of Q that satisfy Equation 4 are those that take the form 2^K , where K is an integer. That is, $M = 2^C/2^K = 2^{C-K}$ (both C and K are integers). To demonstrate, let $f_o = f_s/8$ and assume that C is 32. Then $Q = 8 = 2^3$ and $M = 2^{32-3} = 2^{29} = 536,870,912$ (an integer). On the other hand, if $f_o = f_s/10$, Q is 10 and $M = 2^{32}/10 = 429,496,729.6$ (not an integer), which means that the typical DDS is not capable of generating the exact frequency $f_s/10$. The typical DDS can synthesize an output frequency extremely close to $f_s/10$ (based on the frequency resolution) but not exactly $f_s/10$. In most cases, extremely close is acceptable, but in certain applications (network clocks, for example), the exact frequency ratio is a necessity.

Inspection of Equation 4 indicates that the ratio on the right-hand side is established by the modulus of the accumulator (2^C), which is a fixed quantity governed by the accumulator size. The fixed modulus of the accumulator limits the set of available output frequencies. However, it seems reasonable to assume that an adjustable accumulator modulus would make it possible to accommodate values of Q other than just those that take the form 2^K . Such is the foundation of the programmable modulus DDS architecture.

PROGRAMMABLE MODULUS DDS

A programmable modulus is a modification of the typical accumulator-based DDS architecture (see Figure 2). Although the underlying function of the programmable modulus technique is to alter the accumulator modulus, the actual implementation is more complicated. The complication is twofold.

The first complication is that the angle-to-amplitude converter maps the entire P -bit input range (0 to 2^P) to 0 to 2π radians. It is this power-of-two mapping arrangement that allows the angle-to-amplitude converter to operate efficiently. Arbitrarily changing the accumulator modulus to something other than a power-of-two necessarily violates the mapping arrangement required by the angle-to-amplitude converter. To overcome this difficulty, the programmable modulus architecture uses a secondary accumulator to make the primary accumulator appear to have an altered modulus while still maintaining the power-of-two mapping required for the angle-to-amplitude converter.

The second complication involves spurious performance. If not done properly, modification of the accumulator modulus can introduce spectral artifacts. The programmable modulus implementation requires great care to minimize any negative impact on spurious performance. The architecture of the [AD9913](#) embodies such care as demonstrated by its excellent spurious performance.




Figure 2. Programmable Modulus DDS Architecture

07234-002

The programmable modulus DDS architecture results in a variation of Equation 3 as follows:

$$\frac{f_o}{f_s} = \frac{M}{N} \tag{5}$$

where M and N are integers and M < N. In practice, M should be constrained to less than N/2 in order to comply with the Nyquist sampling requirement.

Notice that the frequency ratio for the programmable modulus DDS (Equation 5) is very similar to that of the typical accumulator-based DDS (Equation 3). The only difference is that N is not required to be a power-of-two for the programmable modulus, but can be an arbitrary integer. In practice, hardware constraints place limits on the range of values for N.

Implementation of the programmable modulus function within the AD9913 is such that the fraction, M/N, is expressible per Equation 6. Note that the form of Equation 6 implies a compound frequency tuning word with X representing the integer part and A/B representing the fractional part. X, A, and B are inputs to the modulus control logic in Figure 2.

$$\frac{M}{N} = \frac{X + \frac{A}{B}}{2^{32}} \tag{6}$$

where X, A and B are 32-bit integers programmable with any integer from 0 to 2³² - 1, inclusive.

Although programmable over this range, the architecture imposes the following constraints on A and B:

$$B > 0$$

$$A < B$$

The fact that X, A, and B are 32-bit integers places the constraints on N. Specifically, for N ≤ 2³², N can be any integer from 0 to 2³², but for N > 2³² the choices for N are constrained per Table 1.

Table 1. Constraints on the Value of N

N	Constraints	Remarks
N ≤ 2 ³²	3 ≤ N ≤ 2 ³²	Note that N = 0 is meaningless and N = 1 or N = 2 is trivial.
N > 2 ³²	N = K × 2 ³²	K is an integer: 2 ≤ K ≤ 2 ³² - 1. This constrains N to an integer multiple of 2 ³² with an upper bound of (2 ³² - 1) × 2 ³² .

CALCULATING THE MODULUS PARAMETERS: A, B, AND X

Given the known values of f_o and f_s, use f_o/f_s to produce the fraction, M/N, but with both M and N expressed as integers. Next, reduce the fraction to its lowest terms and ensure that N satisfies the two constraints in Table 1. If N fails to satisfy the constraints of Table 1, then synthesis of the exact output frequency (f_o) from the given sample rate (f_s) is not possible.

With M and N expressed as integers, use long division to find X as summarized in Equation 7. That is, given M and N, choose X as the largest integer that yields a positive (or zero) value for Y.

$$\frac{f_o}{f_s} = \frac{M}{N} \longrightarrow N \overline{)M2^{32}} \tag{7}$$

$$Y \longrightarrow Y = M2^{32} - XN$$

where:

X is the integer quotient of the division process.

Y is the integer remainder.

Equation 7 constitutes a long division operation with N as the divisor and M as part of the dividend. Note that M and N are integers with the fraction M/N reduced to lowest terms and equivalent to the fraction f_o/f_s. The first step in the long division process is to choose X as the largest integer that satisfies XN ≤ M2³². Next, place the product of X and N below the dividend and subtract it from the dividend, which leaves the remainder, Y. Thus, Y is the value obtained by subtracting XN from M2³².

Equation 7 yields Y as an integer such that the fraction Y/N has the same numeric value as A/B in Equation 6. Reducing the fraction Y/N to its lowest terms yields the integers A and B.

Interestingly, M establishes the periodicity of the accumulator output sequence. That is, the accumulator is required to cycle through M overflows to complete one unique output sequence and return to its original starting point. This unique output sequence repeats indefinitely.

It is worth noting that if Y is 0, then A is 0 and the programmable modulus feature is unnecessary. That is, f_o can be synthesized directly from f_s using a typical accumulator-based DDS. In fact, when Y = 0 it is best to disable the programmable modulus feature of the AD9913, thereby reducing power consumption.

PROGRAMMABLE MODULUS EXAMPLE

Consider the case in which f_s = 250 MHz and the desired value of f_o is 25 MHz. This scenario synthesizes an output frequency that is not a power-of-two submultiple of the sample rate, namely f_o = (1/10)f_s, which is not possible with a typical accumulator-based DDS. The frequency ratio, f_o/f_s, leads directly to M and N, which are determined by reducing the fraction (25,000,000/250,000,000) to its lowest terms. That is:

$$M/N = 25,000,000/250,000,000 = 1/10$$

Therefore, M = 1 and N = 10 and using Equation 7 allows X and Y to be calculated as X = 429,496,729 and Y = 6. The fraction, Y/N, is 6/10, which reduces to 3/5. Thus, A = 3 and B = 5.

Programming these values into the AD9913 causes it to produce an output frequency of exactly 25 MHz given a 250 MHz sampling clock.

A typical accumulator-based DDS, regardless of the capacity of its accumulator, is not capable of synthesizing exactly 25 MHz given that $f_s = 250$ MHz. For example, the closest that a 32-bit accumulator-based DDS can get to 25 MHz with $f_s = 250$ MHz is 25.000000023283064365386962890625 MHz. This constitutes an error of 0.931 parts per billion. In a network clocking application, for example, this close value may not be acceptable because such applications are very sensitive to any fractional frequency error.

CONSIDERATIONS

Most computational programs (for example, Excel®, Mathcad®, and MATLAB®) can introduce numerical errors when calculating X, Y, A, and B. These errors are due to internal rounding or truncation, which shows up when multiplying or dividing with large integers. However, the calculator included as part of the Windows® XP operating system provides sufficient numeric precision to perform these computations without the aforementioned errors (at least in the context of the 32-bit numbers associated with programming the [AD9913](#)). In fact, having prior knowledge of the numeric precision issue led developers of the [AD9913](#) evaluation board software to mitigate any rounding or truncation errors.

To demonstrate the effects of rounding and truncation, consider again the case in which the desired value of f_o is 25 MHz, but f_s is 249,999,999.5 MHz (instead of 250 MHz). The frequency ratio, f_o/f_s , leads directly to M and N by reducing the fraction (25,000,000/249,999,999.5) to its lowest terms. This first requires multiplying M and N by 2 to eliminate the decimal point in N, as follows:

$$M/N = 25,000,000/249,999,999.5 = 50,000,000/499,999,999$$

The fraction, 50,000,000/499,999,999 is irreducible, therefore, use these values for M and N. This leads to $X = 429,496,730$ and $Y = 229,496,730$. Thus, $Y/N = 229,496,730/499,999,999$, which is irreducible leading to $A = 229,496,730$ and $B = 499,999,999$. These values yield $f_o = 25$ MHz, exactly.

Using Excel yields the same value for X, but Y calculates as 229,496,736. Notice that the last digit of Y is 6 instead of 0 (the correct value). Thus, $Y/N = 229,496,736/499,999,999$, which is irreducible leading to $A = 229,496,736$ and $B = 499,999,999$. These values yield $f_o = 25.0000000000000000698491930961609$ MHz. Applications requiring a precise f_o/f_s ratio might not be able to tolerate this small error.

More dramatically, consider the case where N is the maximum allowable value of $(2^{32} - 1) \times 2^{32}$. To express this integer requires 64 bits. Most computational programs are unable to represent integers in excess of 48 bits.

Be sure that the computing engine used to calculate X, Y, A, and B has sufficient numerical precision to fully represent very large integers. Otherwise, there is the potential to introduce rounding or truncation errors into the calculations, which yields incorrect results.