

Section 4. The Effect of DAC Resolution on Spurious Performance

By Ken Gentile, Systems Engineer, Analog Devices, Inc.

The resolution of a DAC is specified by the number of its input bits. For example, the resolution of a DAC with 10 input bits is referred to as having “10-bit resolution”. The impact of DAC resolution is most easily understood by visualizing the reconstruction of a sine wave.




Figure 4.1. Effect of DAC Resolution

Consider Figure 4.1 in which a 4-bit DAC (quantized black trace) is used to reconstruct a perfect sine wave (smooth red trace). The vertical lines are time markers and identify the instants in time at which the DAC output is updated to a new value. Thus, the horizontal distance between the vertical lines represents the sample period. Note the deviation between the DAC output signal and the perfect sine wave. The vertical distance between the two traces at the sampling instants is the error introduced by the DAC as a result of its finite resolution. This error is known as **quantization error** and gives rise to an effect known as **quantization distortion**.

To understand the nature of the quantization distortion, note the sharp edges in the DAC output signal. These sharp edges imply the presence of high frequency components superimposed on the fundamental. It is these high frequency components that constitute quantization distortion. In the frequency domain, quantization distortion errors are aliased within the Nyquist band and appear as discrete spurs in the DAC output spectrum.




Figure 4.2. 4-Bit vs. 8-Bit DAC Output Spectra

As the DAC resolution increases the quantization distortion decreases; i.e., the spurious content of the DAC output spectrum decreases. This makes sense because an increase in resolution results in a decrease in quantization error. This, in turn, results in less error in the reconstructed sine wave. Less error implies less distortion; i.e., less spurious content. This is graphically depicted in Figure 4.2. Note that the spurs associated with the 8-bit DAC are generally lower than those of the 4-bit DAC.

In fact, the relationship between DAC resolution and the amount of distortion is quantifiable. If the DAC is operated at its fullscale output level, then the ratio of signal power to quantization noise power (SQR) is given by:

$$\text{SQR} = 1.76 + 6.02B \quad (\text{dB})$$

Where B is the number of bits of DAC resolution.

For example, an 8-bit DAC exhibits an SQR of 49.92dB. It should be noted that the SQR equation only specifies the total noise power due to quantization errors. It does not provide any information as to the distribution of the spurs or the maximum spur level, only the combined power of all the spurs relative to the fundamental.

A second point to consider is that the SQR equation applies only if the DAC operates at fullscale. At output levels below full scale the power in the fundamental is reduced, but the quantization error remains constant. The net effect is a reduction in SQR; that is, the quantization noise becomes more significant relative to the fundamental. The effect of operating the DAC at less than fullscale is quantifiable and is given as:

$$A = 20\log(\text{FFS}) \quad (\text{dB})$$

where FFS is the fraction of fullscale at which the DAC operates. Thus, the SQR equation becomes:

$$\begin{aligned} \text{SQR} &= 1.76 + 6.02B + A \\ &= 1.76 + 6.02B + 20\log(\text{FFS}) \quad (\text{dB}) \end{aligned}$$

Continuing the previous example, if operate the DAC at 70% of fullscale ($A=0.7$) the resulting SQR is 46.82dB (a 3.1dB reduction from the original SQR performance).

The Effects of Oversampling on Spurious Performance

In oversampling, a sample rate is used that is higher than that required by the Nyquist criteria. Remember, Nyquist requires that the bandwidth of the sampled signal be constrained to $\frac{1}{2}$ of the sample rate. If the bandwidth of the sampled signal is intentionally constrained to a fraction of the Nyquist requirement, then the sample rate is in excess of the Nyquist requirement and oversampling is employed.

Figure 4.3 shows how oversampling improves SQR. The amount of quantization noise power is dependent on the resolution of the DAC. It is a fixed quantity and is proportional to the shaded area. In the oversampled case, the total amount of quantization noise power is the same as in the Nyquist sampled case. Since the noise power is the same in both cases (it's constant), and the area of the noise rectangle is proportional to the noise power, then the height of the noise rectangle in the oversampled case must be less than the Nyquist sampled case in order to maintain the same area. Note that in the band of interest the area of the noise rectangle is less for the oversampled case. Thus, for a given amount of signal power in the band of interest, the signal to noise ratio is greater when oversampling is employed.

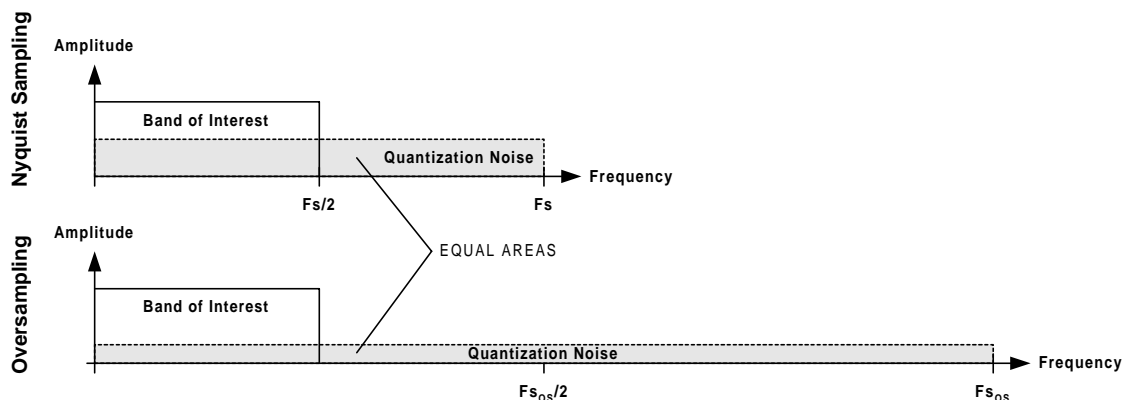


Figure 4.3. The Effect of Oversampling on SQR

The effect of oversampling is quantifiable and is given as:

$$C = 10\log(F_{sOs}/F_s) \text{ (dB)}$$

Where F_s is the Nyquist sampling rate and F_{sOs} is the oversampling rate. The modified SQR equation is:

$$\begin{aligned} \text{SQR} &= 1.76 + 6.02B + A + C \\ &= 1.76 + 6.02B + 20\log(FFS) + 10\log(F_{sOs}/F_s) \text{ (dB)} \end{aligned}$$

Returning to the previous example, if we operate the DAC at 70% of fullscale and oversample by a factor of 3, the SNR becomes 51.59dB. This constitutes an overall improvement of 1.67dB

over the original fullscale SQR performance. In this case, oversampling more than compensated for operating the DAC at only 70% of fullscale.

The Effect of Truncating the Phase Accumulator on Spurious Performance

Phase truncation is an important aspect of DDS architectures. Consider a DDS with a 32-bit phase accumulator. To directly convert 32 bits of phase to a corresponding amplitude would require 2^{32} entries in a lookup table. That's 4,294,967,296 entries! If each entry is stored with 8-bit accuracy, then 4-gigabytes of lookup table memory would be required. Clearly, it would be impractical to implement such a design.

The solution is to use a fraction of the most significant bits of the accumulator output to provide phase information. For example, in a 32-bit DDS design, only the upper most 12 bits might be used for phase information. The lower 20 bits would be ignored (truncated) in this case.

To understand the implications of truncating the phase accumulator output it is helpful to use the concept of the "digital phase wheel". Consider a simple DDS architecture that uses an 8-bit accumulator of which only the upper 5 bits are used for resolving phase. The phase wheel depiction of this particular model is shown in Figure 4.4.

With an 8-bit accumulator, the phase resolution associated with the accumulator is $1/256^{\text{th}}$ of a full circle, or 1.41° ($360/2^8$). In Figure 4.4, the accumulator phase resolution is identified by the outer circle of tic marks. If only the most significant 5 bits of the accumulator are used to convey phase information, then the resolution becomes $1/32^{\text{nd}}$ of a full circle, or 11.25° ($360/2^5$). These are identified by the inner circle of tic marks.

Now let us assume that a tuning word value of 6 is used. That is, the accumulator is to count by increments of 6. The first four phase angles corresponding to 6-count steps of the accumulator are depicted in Figure 4.4. Note that the first phase step (6 counts on the outer circle) falls short of the first inner tic mark. Thus, a discrepancy arises between the phase of the accumulator (the outer circle) and the phase as determined by 5-bit resolution (the inner circle). This discrepancy results in a phase error of 8.46° ($6 \times 1.41^{\circ}$), as depicted by arc **E1** in the figure.

On the second phase step of the accumulator (6 more counts on the outer circle) the phase of the accumulator resides between the 1st and 2nd tic marks on the inner circle. Again, there is a discrepancy between the phase of the accumulator and the phase as determined by 5 bits of resolution. The result is an error of 5.64° ($4 \times 1.41^{\circ}$) as depicted by arc **E2** in the figure. Similarly, at the 3rd phase step of the accumulator an error of 2.82° ($2 \times 1.41^{\circ}$) results. On the 4th phase step, however, the accumulator phase and the 5-bit resolution phase coincide resulting in no phase error. This pattern continues as the accumulator increments by 6 counts on the outer circle each time.




Figure 4.4. Phase Truncation Error and the Phase Wheel

Obviously, the phase errors introduced by truncating the accumulator will result in errors in amplitude during the phase-to-amplitude conversion process inherent in the DDS. It turns out that these errors are periodic. They are periodic because, regardless of the tuning word chosen, after a sufficient number of revolutions of the phase wheel, the accumulator phase and truncated phase will coincide. Since these amplitude errors are periodic in the time domain, they appear as line spectra (spurs) in the frequency domain and are what is known as **phase truncation spurs**.

It turns out that the magnitude and distribution of phase truncation spurs is dependent on three factors (Ref. [3]):

1. Accumulator size (**A** bits)
2. Phase word size (**P** bits); i.e., the number of bits of phase after truncation
3. Tuning word (**T**)

Phase Truncation Spur Magnitude

Certain tuning words yield no phase truncation spurs at all while others yield spurs with the maximum possible level. If the quantity, A-P, is 4 or more (usually the case for any practical DDS design), then the maximum spur level turns out to be very closely approximated by **-6.02P** dBc (i.e., 6.02P decibels below the level of the tuning word frequency). So, a 32-bit DDS with a 12-bit phase word will yield phase truncation spurs of no more than -72dBc regardless of the tuning word chosen.

Tuning words that yield the maximum spur level are those that satisfy the following:

$$\text{GCD}(T, 2^{(A-P)}) = 2^{(A-P-1)}$$

Where $\text{GCD}(X, Y)$ is the Greatest Common Divisor of **both** X and Y. In order for this equation to be true, a tuning word bit pattern for the tuning word must be as shown in Figure 4.5 below.




Figure 4.5. Tuning Word Patterns That Yield Maximum Spur Level

An **A**-bit word is shown, which corresponds to a phase accumulator with **A** bits of resolution. The upper **P** bits constitute the **phase word** (the bits that are to be used for conversion from phase to amplitude). The lower **A-P** bits are truncated, that is, ignored as far as phase resolution is concerned. The tuning word, **T**, is made up of the **A-1** least significant bits (the most significant bit of the tuning word must be a 0 to avoid the problem of aliasing). As shown in the above figure, any tuning word with a 1 in bit position $2^{(A-P-1)}$ and 0's in all less significant bit positions will yield the worst case phase truncation spur level (-6.02P dBc).

At the other extreme are tuning words that yield no phase truncation spurs. Such tuning words must satisfy,

$$\text{GCD}(T, 2^{(A-P)}) = 2^{(A-P)}$$

In order for this equation to be true, the tuning word bit pattern must be as shown in Figure 4.6 below.




Figure 4.6. Tuning Word Patterns That Yield No Phase Truncation Spurs

Thus, tuning words that yield no phase truncation spurs are characterized by a 1 in bit position $2^{(A-P)}$ and 0's in all less significant bit positions. All other tuning word patterns that do not fit the two categories above will yield phase truncation spur levels between the two extremes.

Phase Truncation Spur Distribution

To precisely analyze the distribution of phase truncation spurs is quite complicated. A detailed analysis may be found in [3]. Rather than delve into the details of the analysis, a more intuitive presentation follows.

Remember, first of all, that the DDS core consists of an accumulator which recursively adds the tuning word value. Several iterations of this process are shown in Figure 4.7. Initially, the accumulator contains the value of the tuning word (in this case, an arbitrary binary number which has been assigned the variable, K). On each successive cycle of the DDS system clock, the tuning word is added to the previous contents of the accumulator. Remember, however, that the accumulator is modulo 2^A , so bits that would carry beyond the MSB are simply dropped. As the accumulator sequence proceeds the value of the accumulator will eventually return to the original tuning word value and the sequence will repeat. The number of steps (or clock cycles) required to accomplish this is known as the Grand Repetition Rate (**GRR**). The formula for determining the GRR is:

$$GRR = 2^A / \text{GCD}(T, 2^A)$$

For example, in case shown, A is 20 and T is 182,898 (base 10), which yields a GRR of 524,288. From this result it can be seen that over a half a million clock cycles are required before the accumulator begins to repeat its sequence. Although this may seem like a long repetition period, keep in mind that some DDS cores use 48-bit accumulators ($A=48$), which can yield enormous GRR values.

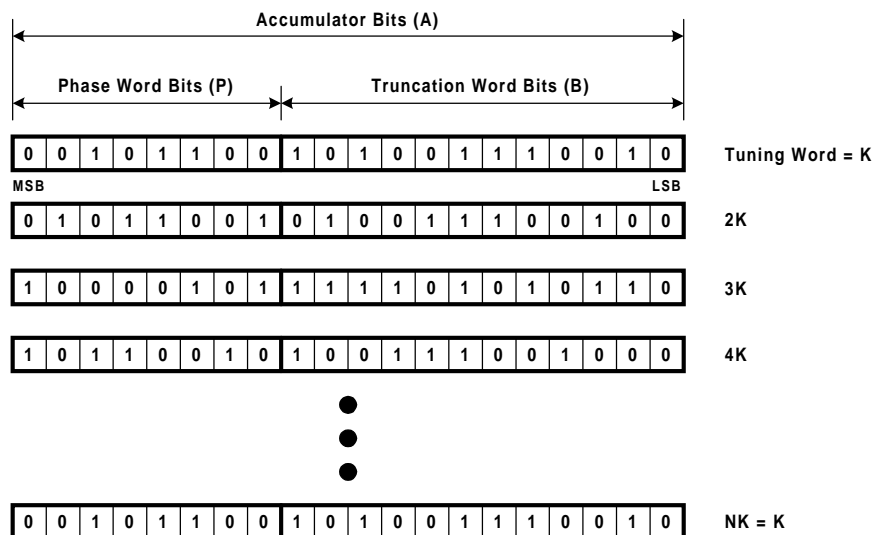


Figure 4.7. Accumulator Sequence

Refer, once again, to Figure 4.7. The P-bits of the phase word are passed along to the phase-to-amplitude conversion portion of the DDS, which is used to produce the output waveform. However, the B-bits of the truncation word are not passed along to the phase-to-amplitude converter. Therefore, if the full A bits of the accumulator represent the true phase, but only P-bits of the phase word are used for determining amplitude, then the output signal is essentially in error by the value of the truncation word. Thus, the output signal can be thought of as a composite of a full resolution signal (that which would be obtained with no phase truncation) and an error signal due to the B-bits of the truncation word.

The error signal, then, is a source of spurious noise. Since the error signal is defined by the truncation word, then analysis of the behavior of the truncation word should allow some insight into the nature of the error signal. Thus, we shall focus on only the truncation word and ignore the phase word.

If only the truncation bits are considered, it is possible to determine the period over which the truncation word repeats; i.e., the GRR of the truncation word. For example, for the conditions given in Figure 4.7, the value of A becomes 12 (the number of truncation bits). The truncation word behaves as a B-bit accumulator with an equivalent tuning word (**ETW**) given by,

$$\text{ETW} = T \text{ modulus } 2^B$$

Where T is the original tuning word. The result of this operation is nothing more than the value of the truncation word portion of the original tuning word. For the given example the ETW is 2,674 (base 10). So, with A=12 and T=2674, the GRR is 2,048. Thus, every 2,048 clock cycles, the truncation word will repeat the pattern of its sequence. So, at this point, we know we have an error signal that is periodic over a time interval of 2,048 clock cycles.

But what is the behavior of the truncation word within this period? That question can be answered by noting that the “**capacity**” of the truncation word is 2^B . Dividing the capacity by the ETW determines the number of clock cycles required to cause the accumulator to overflow. The capacity of the truncation word is easily calculated because in the example given B is 12. This yields a truncation word capacity of $2^{12} = 4096$.

Before we divide by the ETW, however, notice that the MSB of the ETW is a 1. This implies an overflow period of less than 2 clock cycles, which, in turn, implies that the frequency produced would be an alias. So, we must adjust the ETW by subtracting it from the capacity of the truncation word (4096). So, the adjusted ETW is 1422 (4096 – 2674). If the MSB of the ETW had been a 0, the alias adjustment procedure would not have been necessary.

Now that we know the capacity of the truncation word and the properly adjusted ETW, we can determine the overflow period of the truncation word as:

$$\text{Capacity/ETW} = 2^B/1422 = 4096/1422 = 2.88045$$

This value is the average number of clock cycles required for the truncation word to overflow. Since we know that the GRR of the truncation word is 2048 clocks and that it takes ~2.88 clocks for the truncation word to overflow, then the number of overflows that occurs over the period of the GRR is:

$$\text{Number of Overflows} = \text{GRR}/(\text{Capacity}/\text{ETW}) = 2048/(4096/1422) = 711$$

With this information it is possible to visualize the behavior of the truncation word as shown in Figure 4.8 below.




Figure 4.8. Behavior of the Truncation Word

Note that the truncation word accumulates up to a maximum value of 2^B . It has the shape of a sawtooth waveform with a period of $4096/1422$ clock cycles. It should be apparent that the sawtooth shape results from the overflow characteristic of the accumulator. Also note that the complete sequence of truncation word values repeats after a period of 2048 clock cycles. Since the behavior of the truncation word is periodic in the time domain, then its fourier transform is periodic in the frequency domain. Also, the truncation word sequence is a real sequence, so the fourier transform may be represented by half as many frequency points as there are periodic time domain points (because the fourier transform of a real time domain sequence is symmetric about the origin in the frequency domain). Hence, there will be 1024 discrete frequencies associated with the behavior of the truncation word, and these frequencies constitute the truncation spurs.

Furthermore, the spectrum of the truncation word sequence will be related to that of a sawtooth waveform. The fundamental frequency of the sawtooth is $F_s \times (\text{ETW}/\text{Capacity})$ or $0.3472 F_s$ for the example given. The spectrum of a sawtooth waveform is comprised of harmonics of its fundamental. Since we know that there are 1024 discrete frequencies associated with the truncation word sequence, then the spectrum consists of triangle waveform with 1024 frequencies spaced at intervals of $0.3472 F_s$. This spans a frequency range of $355.5 F_s$. This, of course, results in aliasing of the higher order harmonics into the Nyquist bandwidth, $F_s/2$. Figure 4.9 below illustrates this phenomenon.




Figure 4.9: Spectrum of Truncation Word Sequence

The upper trace of Figure 4.9 shows the partial spectrum of the sawtooth waveform. The middle trace shows the remapping of the spectral lines due to aliasing. Note that aliasing causes spurs in frequency bands that are odd integer multiples of $F_s/2$ to map directly into the region of $F_s/2$. While spurs that occur in frequency bands that are even multiples of $F_s/2$ map as mirror images into the region of $F_s/2$. Such is the nature of the aliasing phenomenon. The bottom trace of the figure shows only the region $F_s/2$ (the Nyquist band) with the remapped spectral lines. This is the actual truncation spur spectrum produced by the DDS. Keep in mind however, that Figure 4.9 only displays the frequency range of 0 to $3F_s$. The full spectrum of the sawtooth waveform actually spans $355.5 F_s$. Thus, there are many more truncation spurs present than are actually shown in the Figure 4.9 (the intent of Figure 4.9 is to demonstrate the concept rather than to be exhaustively accurate).

Phase Truncation Summary

In summary, truncation of the phase accumulator results in an error in the DDS output signal. This error signal is characterized by the behavior of the truncation word (the truncation word being the portion of the phase accumulator which contains the truncated bits). Furthermore, the truncation error signal causes discrete frequency spurs to appear in the DDS output and these spurs are referred to as phase truncation spurs.

The magnitude of the phase truncation spurs has an upper bound that is determined by the number of bits in the phase word (P). The value of that upper bound is $-6.02P$ dBc and this upper bound occurs for a specific class of tuning words. Namely, those tuning words for which the truncated bits are all 0's except for the most significant truncated bit. However, a second class of tuning words results in no phase truncation spurs. These are characterized by all 0's in

the truncation word and a 1 in *at least* the LSB position of the phase word. All other classes of tuning words produce phase truncations spurs with a maximum magnitude *less than* $-6.02P$ dBc.

The distribution of the truncation words is not as easily characterized as the maximum magnitude. However, it has been explained that the truncation word portion of the accumulator can be thought of as the source of an error phase signal. This error signal is of the form of a sawtooth waveform with a frequency of:

$$F_s(\text{ETW}/2^B),$$

where F_s is the DDS system clock frequency, ETW is the equivalent tuning word represented by the truncated bits (after alias correction), and B is the number of truncation bits. The number of harmonics of this frequency which must be considered for the analysis for phase truncation spurs is given by:

$$2^{B-1} / \text{GCD}(\text{ETW}, 2^B),$$

where $\text{GCD}(x,y)$ is the Greatest Common Divisor of x and y. The result is a spectrum which spans many multiples of F_s . Therefore, a remapping of the harmonics of the sawtooth spectrum must be performed due to aliasing. The result of the remapping places all of the spurs of the sawtooth spectrum within the Nyquist band ($F_s/2$). This constitutes the distribution of the phase truncation spurs as produced by the DDS.

Additional DDS Spur Sources

The previous two sections addressed two of the sources of DDS spurs; DAC resolution and phase truncation. Additional sources of DDS spurs include:

1. DAC nonlinearity
2. Switching transients associated with the DAC
3. Clock feedthrough

DAC nonlinearity is a consequence of the inability to design a perfect DAC. There will always be an error associated with the expected DAC output level for a given input code and the actual output level. DAC manufacturers express this error as DNL (differential nonlinearity) and INL (integral nonlinearity). The net result of DNL and INL is that the relationship between the DAC's expected output and its actual output is not perfectly linear. This means that an input signal will be transformed through some nonlinear process before appearing at the output. If a perfect digital sine wave is fed into the DAC, the nonlinear process causes the output to contain the desired sine wave plus harmonics. Thus, a distorted sine wave is produced at the DAC output. This form of error is known as *harmonic distortion*. The result is harmonically related spurs in the output spectrum. The amplitude of the spurs is not readily predictable as it is a function of the DAC linearity. However, the location of such spurs is predictable, since they are harmonically related to the tuning word frequency of the DDS. For example, if the DDS is tuned to 100kHz, then the 2nd harmonic is at 200kHz, the 3rd at 300kHz, and so on. Generally, for a DDS output frequency of f_o , the nth harmonic is at nf_o . Remember, however, that a DDS is a sampled system operating at some system sample rate, F_s . So, the Nyquist criteria are applicable. Thus, any harmonics greater than $1/2F_s$ will appear as aliases in the frequency range between 0 and $1/2F_s$ (also known as the first Nyquist zone). The 2nd Nyquist zone covers the

range from $\frac{1}{2}F_s$ to F_s . The 3rd Nyquist zone is from F_s to $1.5F_s$, and so on. Frequencies in the ODD Nyquist zones map directly onto the 1st Nyquist zone, while frequencies in the EVEN Nyquist zones map in mirrored fashion onto the 1st Nyquist zone. This is shown pictorially in Figure 4.10.




Figure 4.10. Nyquist Zones and Aliased Frequency Mapping

The procedure, then, for determining the aliased frequency of the Nth harmonic is as follows:

1. Let R be the remainder of the quotient $(Nf_o)/F_s$, where N is an integer.
2. Let $SPUR_N$ be the aliased frequency of the Nth harmonic spur.
3. Then $SPUR_N = R$ if $(R \leq \frac{1}{2}F_s)$, otherwise $SPUR_N = F_s - R$.

The above algorithm provides a means of predicting the location of harmonic spurs that result from nonlinearities associated with a practical DAC. As mentioned earlier, the magnitude of the spurs is not predictable because it is directly related to the amount of non-linearity exhibited by a particular DAC (i.e., non-linearity is DAC dependent).

Another source of spurs are **switching transients** that arise within the internal physical architecture of the DAC. Non-symmetrical rising and falling switching characteristics such as unequal rise and fall time will also contribute to harmonic distortion. The amount of distortion is determined by the effective ac or dynamic transfer function. Transients can cause ringing on the rising and/or falling edges of the DAC output waveform. Ringing tends to occur at the natural resonant frequency of the circuit involved and may show up as spurs in the output spectrum.

Clock feedthrough is another source of DDS spurs. Many mixed signal designs include one or more high frequency clock circuits on chip. It is not uncommon for these clock signals to appear at the DAC output by means of capacitive or inductive coupling. Obviously, any coupling of a clock signal into the DAC output will result in a spectral line at the frequency of the interfering clock signal. Another possibility is that the clock signal is coupled to the DAC's sample clock. This causes the DAC output signal to be modulated by the clock signal. The result is spurs that are symmetric about the frequency of the output signal.

Proper layout and fabrication techniques are the only insurance against these forms of spurious contamination. The spectral location of clock feedthrough spurs is predictable since a device's internal clock frequencies are usually known. Therefore, clock feedthrough spurs are likely to be

found in the output spectrum coincident with their associated frequencies (or their aliases) or at an associated offset from the output frequency in the case of modulation.

Wideband Spur Performance

Wideband spurious performance is a measure of the spurious content of the DDS output spectrum over the entire Nyquist band. The worst-case wideband spurs are generally due to the DAC generated harmonics. Wideband spurious performance of a DDS system depends on the quality of both the DAC and the architecture of the DDS core. As discussed earlier, the DDS core is the source of phase truncation spurs. The spur level is bounded by the number of non-truncated phase bits and the spur distribution is a function of the tuning word. Generally speaking, phase truncation spurs will be arbitrarily distributed across the output spectrum and must be considered as part of the wideband spurious performance of the DDS system.

Narrowband Spur Performance

Narrowband spurious performance is a measure of the DDS output spectrum over a very narrow band (typically less than 1% of the system clock frequency) centered on the DDS output frequency. Narrowband spurious performance depends mostly on the purity of the DDS system clock. To a lesser degree, it depends on the distribution of spurs associated with phase truncation. The latter is only a factor, however, when phase truncation spurs happen to fall very near the DDS output frequency.

If the DDS system clock suffers from jitter, then the DDS will be clocked at non-uniform intervals. The result is a spreading of the spectral line at the DDS output frequency. The degree of spreading is proportional to the amount of jitter present. Narrowband performance is further affected when the DDS system clock is driven by a PLL (phase locked loop). The nature of a PLL is to continuously adjust the frequency and phase of the output clock signal to track a reference signal. This continuous adjustment exhibits itself as phase noise in the DDS output spectrum. The result is a further spreading of the spectral line associated with the output frequency of the DDS.

Predicting and Exploiting Spur "Sweet Spots" in a DDS' Tuning Range

In many DDS applications the output frequency need not be constrained to a single specific frequency. Rather, the designer is given the liberty to choose any frequency within a specified band that satisfies the design requirements of the system. Oftentimes, these applications specify fairly stringent spurious noise requirements, but only in a fairly narrow passband surrounding the fundamental output frequency of the DDS. In these applications, the output signal is usually bandpass filtered so that only a particular band around the fundamental output frequency is of critical importance. In these instances, the designer can select a DDS output frequency that lies within the desired bandwidth but yields minimal spurious noise within the passband.

As mentioned earlier, harmonic spurs (such as those due to DAC nonlinearity) fall at predictable locations in the output spectrum. Knowledge of the location of these spurs (and their aliases) can aid the designer in the choice of an optimal output frequency. Simply choose a fundamental frequency that yields harmonic spurs outside of the desired passband. This topic was detailed in the section entitled, "Additional DDS Spur Sources".

Also, knowledge of the location of phase truncation spurs can prove helpful. Choosing the appropriate tuning word can result in minimal spurs in the passband of interest, with the larger phase truncation spurs appearing out of band. This topic was detailed in the section titled, “The Effect of Truncating the Phase Accumulator on Spurious Performance”.

Using the above techniques, the designer can select the output frequency which results in minimum spurious noise within the desired passband. This may have the affect of increased out of band noise, but in many applications bandpass filters are employed to suppress the out of band signals. The net result is a successful implementation of a DDS system. Many times designers pass over a DDS solution because of the lack-luster spurious performance often associated with a DDS system. By employing the above techniques coupled with the improvements that have been made in DDS technology, the designer can now use a DDS in applications where only analog solutions would have been considered.

Jitter and Phase Noise Considerations in a DDS System

The maximum achievable spectral purity of a synthesized sinewave is ultimately related to the purity of the system clock used to drive the DDS. This is due to the fact that in a sampled system the time interval between samples is expected to be constant. Practical limitations, however, make perfectly uniform sampling intervals an impossibility. There is always some variability in the time between samples leading to deviations from the desired sampling interval. These deviations are referred to as **timing jitter**. There are two primary mechanisms that cause jitter the system clock. The first is thermal noise and the second is coupling noise.

Thermal noise is produced from the random motion of electrons in electric circuits. Any device possessing electrical resistance serves as a source for thermal noise. Since thermal noise is random, its frequency spectrum is infinite. In fact, in any given bandwidth, the amount of thermal noise power produced by a given resistance is constant. This fact leads to an expression for the noise voltage, V_{noise} , produced by a resistance, R , in a bandwidth, B . It is given by the equation:

$$V_{\text{noise}} = \sqrt{(4kTRB)}$$

Where V_{noise} is the RMS (root-mean-square) voltage, k is Boltzmann’s constant (1.38×10^{-23} Joules/°K), T is absolute temperature in degrees Kelvin (°K), R is the resistance in ohms, and B is the bandwidth in hertz. So, in a 3000 Hz bandwidth at room temperature (300°K) a 50Ω resistor produces a noise voltage of 49.8nVrms. The important thing to note is that it makes no difference where the center frequency of the 3kHz bandwidth is located. The noise voltage of the room temperature 50Ω resistor is 49.8nVrms whether measured at 10kHz or 10MHz (as long as the bandwidth of the measurement is 3kHz).

The implication here is that whatever circuit is used to generate the system clock it will always exhibit some finite amount of timing jitter due to thermal noise. Thus, thermal noise is the limiting factor when it come to minimizing timing jitter.

The second source of timing jitter is coupled noise. Coupled noise can be in the form of locally coupled noise caused by crosstalk and/or ground loops within or adjacent to the immediate area of the circuit. It can also be introduced from sources far removed from the circuit. Interference

that is coupled into the circuit from the surrounding environment is known as EMI (electromagnetic interference). Sources of EMI may include nearby power lines, radio and TV transmitters, and electric motors, just to name a few.

The existence of jitter leads to the question: “How does timing jitter on the system clock of a DDS effect the spectrum of a synthesized sine wave?” This is best explained via Figure 4.11, which is a Mathcad simulation of a jittered sinusoid.

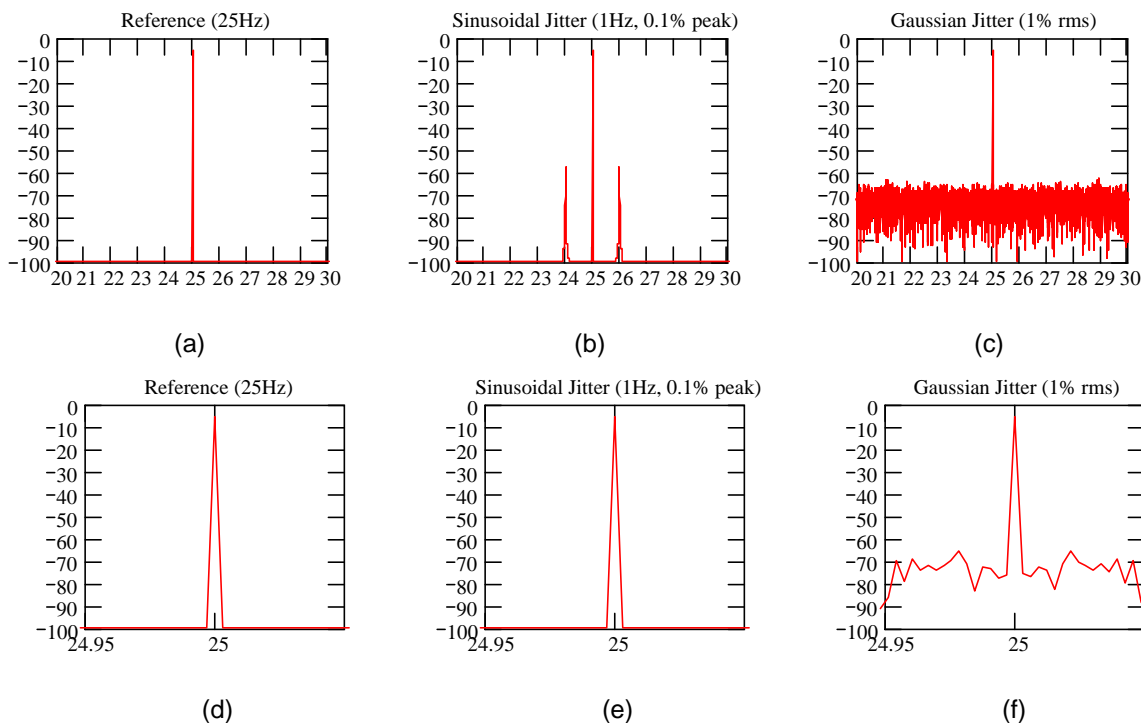


Figure 4.11. Effect of System Clock Jitter

Figure 4.11(a)-(c) span a 10Hz range centered on the 25Hz fundamental frequency, while Figure 4.11(d)-(e) is a “zoom in” around the fundamental frequency to show spectral detail near the fundamental.

Figure 4.11(a) and (d) show the spectrum of a pure sinusoid at a frequency of 25Hz. Note the single spectral line at 25Hz. This is the spectral signature of a pure sinusoid. The widening of the spectral line in Figure 4.11 (d) is a result of the finite resolution of the FFT used in the simulation.

Figure 4.11(b) and (e) show the same sinusoid but with sinusoidal timing jitter added. The jitter varies at a frequency of 1Hz and has a magnitude that is 0.1% of the period of the 25Hz fundamental. Since the period of the fundamental is 40ms, the magnitude of the jitter is 40μs peak. Thus, the sampling of the fundamental occurs at intervals that are not uniformly separated in time. Instead, the sampling instants have a timing error which causes the sampling points to occur around the ideal sampling points with a sinusoidal timing error. Thus, for the example given, the timing error oscillates around the ideal sampling points at a 1Hz rate with a peak deviation of 40μs. Note that sinusoidal jitter in the sampling clock causes modulation sidebands

to appear in the spectrum. Also, the spectral line is unchanged as can be seen by comparing Figure 4.11(d) and (e).

The frequency of the jitter can easily be determined by the separation of the sidebands from the fundamental (1Hz in this case). The magnitude of the jitter can be determined by the relative amplitude of the sidebands. The following formula can be used to convert from dBc to the peak jitter magnitude:

$$\text{Peak Jitter Magnitude} = [10^{(\text{dBc}/20)}]/\pi$$

For the above case, where the jitter sidebands are at -50dBc , the peak jitter magnitude is found to be:

$$[10^{(-50/20)}]/\pi = 0.001 \text{ (or 0.1\%)}$$

This value is relative to the period of the fundamental. Thus, the absolute jitter magnitude is found by multiplying this result by the period of the fundamental (40ms). Thus, the peak jitter magnitude is $40\mu\text{s}$ (0.1% of 40ms).

Figure 4.11(c) and (f) show a pure sinusoid but with random timing jitter added. This implies that the actual sampling instants fluctuate around the ideal sampling time points in a random manner. The jitter in the example follows a Gaussian (or normal) distribution. The mean (μ) and standard deviation (σ) are 0 and 0.0004, respectively. The standard deviation of 0.0004 represents 1% of the fundamental period (or 0.4ms). The timing jitter is defined as Gaussian with a σ value of 0.0004. Thus, statistically, there is a 68% probability that the timing error of any given sampling instant is in error by no more than 0.4ms. Notice in Figure 4.11(c) that random jitter on the sampling clock results in an increase in the level of the noise floor. Furthermore, comparing Figure 4.11(f) to Figure 4.11(d), note that there is a broadening of the fundamental. The broadening of the fundamental is known by the term, **phase noise**.

Output Filtering Considerations

Fundamentally, a DDS is a sampled system. As such, the output spectrum of a DDS system is infinite. Although the device is “tuned” to a specific frequency, it is inferred that the tuned frequency lies within the Nyquist band ($0 \leq f_o \leq \frac{1}{2}F_s$). In actuality, the output spectrum consists of f_o and its alias frequencies as shown below in Figure 4.12.




Figure 4.12. DDS Output Spectrum

The sinc (or $\sin[x]/x$) envelope is a result of the zero-order-hold associated with the output circuit of the DDS (typically a DAC). The images of f_o continue indefinitely, but with ever

decreasing magnitude as a result of the sinc response. In the Figure 4.12, only the result of generating the fundamental frequency by means of the sampling process have been considered. Spurious noise due to harmonic distortion, phase truncation, and all other sources have been ignored for the sake for clarity.

In most applications, the aliases of the fundamental are not desired. Hence, the output section of the DDS is usually followed by a lowpass “antialiasing” filter. The frequency response of an ideal antialias filter would be unity over the Nyquist band ($0 \leq f \leq \frac{1}{2}F_s$) and 0 elsewhere (see Figure 4.13). However, such a filter is not physically realizable. The best one can hope for is a reasonably flat response over some percentage of the Nyquist band (say 90%) with rapidly increasing attenuation up to a frequency of $\frac{1}{2}F_s$, and sufficient attenuation for frequencies beyond $\frac{1}{2}F_s$. This, unfortunately, results in the sacrifice of some portion of the available output bandwidth in order to allow for the non-ideal response of the antialias filter.




Figure 4.13. Antialias Filter

The antialias filter is a critical element in the design of a DDS system. The requirements which must be imposed on the filter design are very much dependent on the details of the DDS system. Before discussing the various types of DDS systems, it is beneficial to review some of the basic filter types in terms of their time domain and frequency domain characteristics.

First of all, it is important to clarify the relationship between the time and frequency domains as applied to filters. In the time domain, we are concerned with the behavior of the filter over time. For example, we can analyze a filter in the time domain by driving it with a pulse and observing the output on an oscilloscope. The oscilloscope displays the response of the filter to the input pulse in the time domain (see Figure 4.14).




Figure 4.14. Time Domain Response

When dealing with filters (or any linear system, for that matter) there is a special case of time domain response that is fundamental in characterizing filter performance. This special case is known as **impulse response**. Impulse response is conceptually identical to the time domain figure above. The only difference is that the rectangular pulse is replaced by an ideal impulse (i.e., an infinitely large voltage spike of zero time duration). Obviously, the concept of an ideal impulse is theoretical in nature, but the response of a filter to such an input would constitute that filter's impulse response. The impulse response of a hypothetical filter is depicted below in Figure 4.15.




Figure 4.15. Impulse Response

Usually, when describing the behavior of a filter, a frequency domain point of view is chosen instead of a time domain point of view. In this case, the earlier oscilloscope analogy can not be used to observe the behavior of the filter. Instead, a spectrum analyzer must be employed, because it is capable of measuring magnitude vs. frequency (whereas an oscilloscope measures amplitude vs. time). A filter's frequency response is a measure of how much signal the filter will pass at a given frequency. A hypothetical lowpass filter response is shown in Figure 4.16. Typical filter parameters of interest are the cutoff frequency (f_c), the stopband frequency (f_s), the maximum passband attenuation (A_{max}) and the minimum stopband attenuation (A_{min}).




Figure 4.16. Frequency Response

Mathematically, there is a direct link between impulse response and frequency response; namely, the Fourier transform. If a filter's impulse response is known (that is, its time domain behavior), then the Fourier transform of the impulse response yields the filter's frequency response (its frequency domain behavior). Likewise, the Inverse Fourier transform of a filter's frequency response yields its impulse response. Thus, the Fourier transform (and its inverse) is the platform by which we can translate our viewpoint between the time and frequency domains.

There is an important reason for exploring the relationship between the time and frequency domains in regard to filters. Specifically, the choice of a particular filter type depends on whether an application requires a filter with certain time domain characteristics or a filter with certain frequency domain characteristics. One must realize that there exists a trade off between the desirable characteristics the two domains. Namely, a smooth time domain response and a sharp frequency domain response. Unfortunately, a filter that exhibits a sharp, well defined passband will necessarily have ringing and overshoot in its impulse response. Likewise, a filter with a smooth time domain characteristic will not yield a sharp transition between its passband and stopband.

So far, two significant aspects of filters have been presented; the time domain and the frequency domain response. Another important filter parameter is **group delay** (which is related to the time domain response). Group delay is a measure of the rate at which signals of different frequencies propagate through the filter. Generally, the group delay at one frequency is not the same as that at another frequency; that is, group delay is typically frequency dependent. This can cause a problem when a filter must carry a group of frequencies simultaneously in its passband. Since the different frequencies propagate at different rates the signals tend to spread out from one another in time. Which becomes a problem in wideband data communication applications where it is important that multi-frequency signals that are sent through a filter arrive at the output of the filter at the same time.

There are many classes of filters that exist in technical literature. However, for most applications the field can be narrowed to three basic filter families. Each is optimized for a particular characteristic in either the time or frequency domain. The three filter types are the Chebyshev, Gaussian, and Legendre families of responses. Filter applications that require fairly sharp frequency response characteristics are best served by the Chebyshev family of responses. However, it is assumed that ringing and overshoot in the time domain do not present a problem in such applications. Conversely, filter applications that require smooth time domain

characteristics (minimal overshoot and ringing and constant group delay) are best served by the Gaussian family of filter responses. In these applications it is assumed that sharp frequency response transitions are not required. For those applications that lie in between these two extremes, the Legendre filter family is a good choice. A brief description of the three filter families follows.

The Chebyshev Family of Responses

The Chebyshev family generally offers sharp frequency domain characteristics. As such, the time domain response is rather poor with significant overshoot and ringing and nonlinear group delay. This makes the Chebyshev family suitable for applications in which the frequency domain characteristics are the dominant area of concern, while the time domain characteristics are of little importance.

The Chebyshev family can be subdivided into four types of responses, each with its own special characteristics. The four types are the *Butterworth* response, the *Chebyshev* response, the *Inverse Chebyshev* response, and the *Cauer-Chebyshev* (also known as *elliptical*) response. Figure 4.17 shows the generic lowpass response of each of the Chebyshev filter types.




Figure 4.17. The Chebyshev Family of Responses

The Butterworth response is completely monotonic. The attenuation increases continuously as frequency increases; i.e., there are no ripples in the attenuation curve. Of the Chebyshev family of filters, the passband of the Butterworth response is the most flat. Its cutoff frequency is identified by the 3dB attenuation point. Attenuation continues to increase with frequency, but the rate of attenuation after cutoff is rather slow.

The Chebyshev response is characterized by attenuation ripples in the passband followed by monotonically increasing attenuation in the stopband. It has a much sharper passband to stopband transition than the Butterworth response. However, the cost for the faster stopband rolloff is ripples in the passband. The steepness of the stopband rolloff is directly proportional to the magnitude of the passband ripples; the larger the ripples, the steeper the rolloff.

The Inverse Chebyshev response is characterized by monotonically increasing attenuation in the passband with ripples in the stopband. Similar to the Chebyshev response, larger stopband ripples yields a steeper passband to stopband transition.

The Elliptical response offers the steepest passband to stopband transition of any of the filter types. The penalty, of course, is attenuation ripples. In this case, both in the passband and

stopband. For applications involving antialiasing filters, the elliptical is usually the filter of choice because of its steep transition region.

The Gaussian Family of Responses

The Gaussian family of responses are well suited to applications in which time domain characteristics are of primary concern. They offer smooth time domain characteristics with little to no ringing or overshoot. Furthermore, group delay is fairly constant. Since the time domain characteristics are so well behaved, it follows that the frequency domain response will not exhibit very sharp transitions. In fact, the frequency response is completely monotonic. The attenuation curve always maintains a negative slope with no peaking of the magnitude in either the passband or stopband.

The Gaussian family can be subdivided into three types of responses, each with its own special characteristics. They are the *Gaussian Magnitude* response, the *Bessel* response, and the *Equiripple Group Delay* response. Figure 4.18 shows the generic lowpass response of each of the Gaussian filter types. Although the magnitude responses of all three types seem to exhibit the same basic shape, each has its own special time domain characteristic for which it has been optimized, as explained below.




Figure 4.18. The Gaussian Family of Responses

The Gaussian Magnitude response is optimized to yield a response curve that most closely resembles a Gaussian distribution. The time domain characteristic offers nearly linear phase response with minimal overshoot or ringing. Group delay is not quite constant, but is dramatically better than that of the Chebyshev family.

The Bessel response is fully optimized for group delay. It offers maximally flat group delay in the passband. The Bessel response is to the time domain as the Butterworth response is to the frequency domain. This makes the Bessel the filter of choice where group delay is of primary concern. It offers nearly linear phase response with minimal overshoot or ringing.

The Equiripple Group Delay response is optimized to yield ripples in the group delay response that do not exceed a prescribed maximum in the passband (much like the magnitude response of the Chebyshev filter). Because the entire passband offers a certain maximum group delay, this filter well suited to wideband applications in which group delay must be controlled over the entire band of interest. Like the other Gaussian filters, the phase response is mostly linear with minimal overshoot or ringing.

The Legendre Family of Responses




Figure 4.19. The Legendre Response

The Legendre filter family consists of a single type. Its passband response has slight ripples and is similar to a Chebyshev response with 0.1dB ripple. The stopband response monotonically decreases. The attenuation rate after the cutoff frequency is steeper than that of the Butterworth type, but not as steep as that of the Chebyshev type. Group delay is virtually constant over the first 25% of the passband, but shows ever increasing deviations as the cutoff frequency is approached.

References:

- [1] Bellamy, J., 1991, Digital Telephony, John Wiley & Sons
- [2] Gentile, K., 1998, “Signal Synthesis and Mixed Signal Technology”, RF Design Magazine, Aug.
- [3] Nicholas III, H. and Samuelli, H., 1987, “An Analysis of the Output Spectrum of Direct Digital Frequency Synthesizers in the Presence of Phase-Accumulator Truncation”, 41st Annual Frequency Control Symposium
- [4] Zverev, A., 1967, Handbook of Filter Synthesis, John Wiley & Sons
- [5] High Speed Design Seminar, 1990 Analog Devices, Inc.