From November 2008 *High Frequency Electronics*

A Wide Dynamic Range Playback System for Radar Signals

By David Friedman, PhD, and Paul Hiller
Symtx, Inc.

Here is a technique for reproducing a digitally recorded radar signal (or other high frequency analog signal) with 105 dB amplitude dynamic range and 16-bit resolution

Improvements in the speed and resolution of both digital to analog converters (DACs) and analog to digital converters (ADCs) have resulted in a continual push to move more functionality into the digital signal processing arena. RF and analog signal processing methodologies such as filtering and frequency translation are being handled increasingly in the digital domain, where near-ideal filters are achievable and analog errors are eliminated.

However, as discussed in the article entitled "A Wide Dynamic Range Radar Digitizer," [1] converting to the digital domain introduces errors which limit overall system performance. One of the most important limitations is dynamic range, which is the range of signal amplitudes that can be captured by an ADC. This is determined by the number of conversion bits as well as by the signal-to-noise ratio (SNR) of the analog components (amplifiers, mixers, etc.) which precede the ADC.

This article describes the use of a dual high-speed 16-bit DAC for reproducing a Doppler weather radar signal. The signal is played back from a digital recording produced using the digitizer described in the preceding article. A simplified block diagram for the system is shown in Figure 1. Note that the dual 16-bit DAC is used to effectively emulate a 20-bit DAC by the means described in this article.

This system is required to digitally record and reproduce an analog reflection-return radar signal down-converted to an IF of 30 MHz in a 1-MHz bandwidth. A key requirement is the ability to accommodate a dynamic range of at least 105 dB between the maximum-capable and minimum-detectable amplitudes that may occur in the course of a single radar trace.

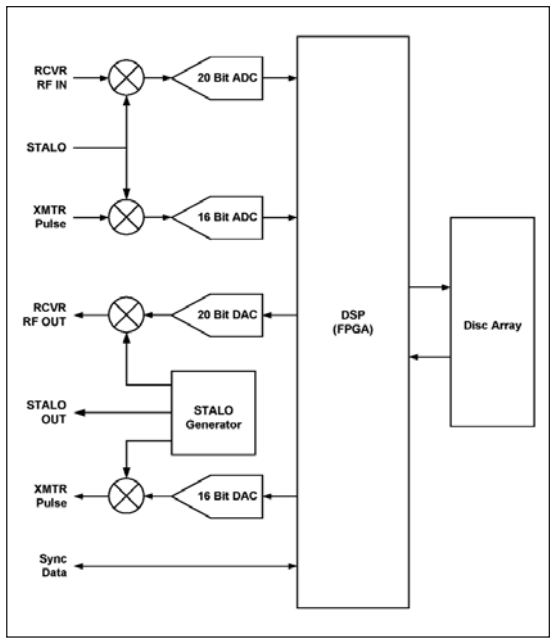

Figure 1 · Weather radar capture and playback system.

The actual system operates above 5 GHz and includes RF mixers, filters, amplifiers, tunable frequency sources, and other analog devices that are not shown in the figure. However, the dynamic range and SNR are set primarily by the IF devices in this diagram.

Weather Radar Signal Handling

The receiver in the radar itself as originally designed used analog AGC to compress the

signal amplitude range prior to digitization. However, this was found to cause distortion and other undesirable effects. The AGC was later eliminated by converting to an all-digital receiver using an arrangement of two 14-bit ADCs with a 24-dB gain offset. One or the other ADC output is used according to the instantaneous amplitude of the signal, and the resulting digital value is bit-shifted as needed to compensate for the gain offset, resulting in an effective 20-bit ADC. Note that this does not provide 20 bits of resolution, since only 14 bits are used for any given sample, but the ratio of maximum-to-minimum signal level is equivalent to that of a 20-bit ADC.

The record-and-playback system developed by Symtx uses a similar approach. Two 16-bit ADCs are used on the input side, offset by 24 dB to give a 20-bit data word, with essentially the reverse procedure at the output. Two 16-bit DACs with 24 dB gain offset generate the output signal, switching between the two according to the signal level.

The details of the DAC output arrangement are the subject of this article. It should be clear that the basic approach is not limited to radar applications, but applies to any case where a wide-dynamic-range analog signal must be generated.

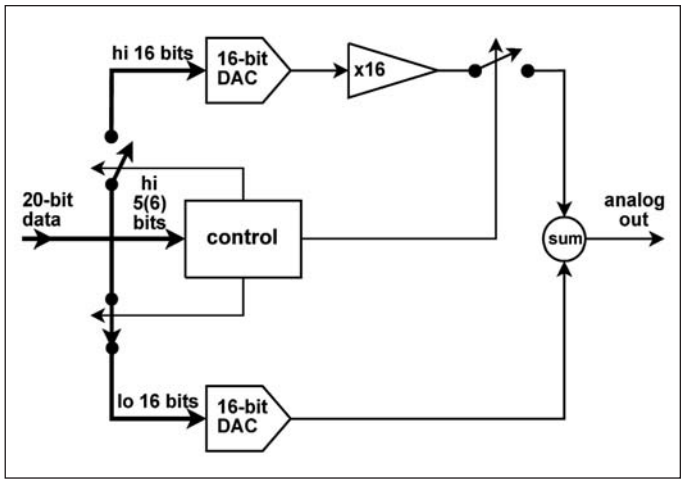

Figure 2 · Dual DAC implementation: two 16-bit DACs aer used to reproduce 20-bit input data.

Dynamic Range

For purposes of this discussion, consider the dynamic range to be the ratio of maximum-capable to minimum-detectable signal amplitude. In terms of the DAC alone, the minimum-detectable signal is determined by its quantization. For example, the dynamic range requirement of 105 dB corresponds to a ratio of approximately 217.5 or a shift of 17.5 bits. This can be shown to be accommodated by the 20-bit word as follows. Allowing for a sign bit leaves 19 bits for the peak magnitude of the largest signal. Shifting right by 17.5 bits leaves 1.5 bits for the peak of the smallest signal, or 1 bit for the RMS level, i.e., the RMS of the smallest signal is equal to the smallest value that can be represented (the magnitude difference corresponding to the low-order bit).

More generally, the dynamic range is determined by the SNR, defined as the ratio of the maximum signal amplitude to the noise floor when a small signal is present (so as to bring quantization noise into account). Assuming a signal must be above the noise floor by a certain amount (in dB) in order to be detectable, the dynamic range will be equal to the SNR less this amount. The noise present at the DAC output consists of quiescent (mostly thermal) circuit noise, which is fixed in absolute level, plus quantization noise and other noise generated in the DAC (as specified by the SNR given in the data sheet), both of which are relative to the DAC's full-scale output value. For a full-scale sinusoidal signal, the SNR defined by quantization noise alone is $6.02 \times N + 1.76$ dB, where N is the number of bits, giving approximately 98 dB for a 16-bit DAC, or 122 dB for 20 bits. This sets an upper limit on the data-sheet SNR, which takes all internal noise sources into account, including nonlinear intermodulation effects which generally will depend on the actual composition of the signal.

For the purposes of this article, we assume that the DAC output level is scaled so that the DAC-internal noise (including quantization noise) is above the quiescent (thermal) noise, so that the dynamic range is determined essentially by the DAC noise. If this is not the case, then the dynamic range is reduced by the amount by which the DAC noise falls below the quiescent noise.

Implementation Details

The scheme employed to obtain increased dynamic range is shown in Figure 2. It uses two identical 16-bit DACs in parallel, with an analog gain offset of 24 dB (a factor of 16) between the two. This results in an effective data width of 20 bits, compensated by a 4-bit shift between the two outputs. The "high-level" DAC (upper path in the figure) is followed by an analog amplifier with a gain of 16, while the "low-level" DAC (lower path in the figure) has no added amplification.

Note that only one DAC is active at a time. Data bits are routed to the active DAC, while a string of zeroes is supplied to the other DAC. The output from the inactive DAC consists of quiescent noise alone.

When the signal amplitude is small, the lower 16 bits of the 20-bit input data stream are sent to the low-level DAC, with a zero word sent to the high-level DAC. When the amplitude is large, the upper 16 bits are sent to the high-level DAC, with the lower four bits of the 20-bit data effectively truncated, and zeroes sent to the low-level

DAC. This can be understood as a 4-bit right shift of the digital data, with a coarsening of the quantization, followed by an analog gain of 16 to restore the level. Once again, note that this approach does not provide 20-bit resolution, but it does provide a dynamic range equivalent to a 20-bit DAC.

The two analog outputs (from the low-level DAC, and from the high-level DAC via the amplifier) are then summed to provide the resulting output signal. However, a simple summation would cause the (amplified) quiescent noise of the high-level DAC to be summed with the output of the low-level DAC, leading to a degradation of up to 24 dB in SNR for low-level signals.

To prevent this, an analog switch is inserted in the high-level path between the amplifier and the summing point. When the high-level DAC is active, the switch passes its output via the amplifier to the summation point, and when the low-level DAC is active, the high-level DAC output is blocked by the switch. Thus, the effective noise level for a small-amplitude signal is that of the low-level DAC, while the maximum output signal level is determined by the saturation point of the high-level DAC. This provides a dynamic range that is 24 dB higher than that of a single DAC.

Note that there is a performance cost associated with adding the analog switch, as it introduces its own sources of error. These are discussed later in this article.

Control for Switching

The "control" block in Figure 2 controls switching between the two DAC paths according to signal amplitude. One indication of amplitude is given by the five high-order bits of the 20-bit data word. When all five bits are equal (either all zeroes or all ones, assuming two's-complement representation), this means the upper four bits effectively contribute nothing, and therefore the data value is within the capability of the low-level DAC alone. If the bits are not all equal, this means the data word is outside the range of the low-level DAC, and the high-level DAC path is required. Alternately, if the upper six bits are equal, this means the data value is at least 6 dB below the maximum of the low-level DAC, giving a 6-dB margin. In either case, this effectively defines a switching threshold on the signal peak-amplitude envelope.

The control level as defined above is generated by taking the AND of the upper five or six bits, ORed with the AND of the inversions of the same bits. If the resulting value equals one (the bits are all equal), the low-level DAC receives the lower 16 data bits, the high-level DAC receives zeroes, and the analog switch is set to block the high-level DAC output. Otherwise, the high-level DAC receives the upper 16 bits, the low-level DAC receives zeroes, and the analog switch is set to pass the high-level DAC output. This describes the switching operation in static terms.

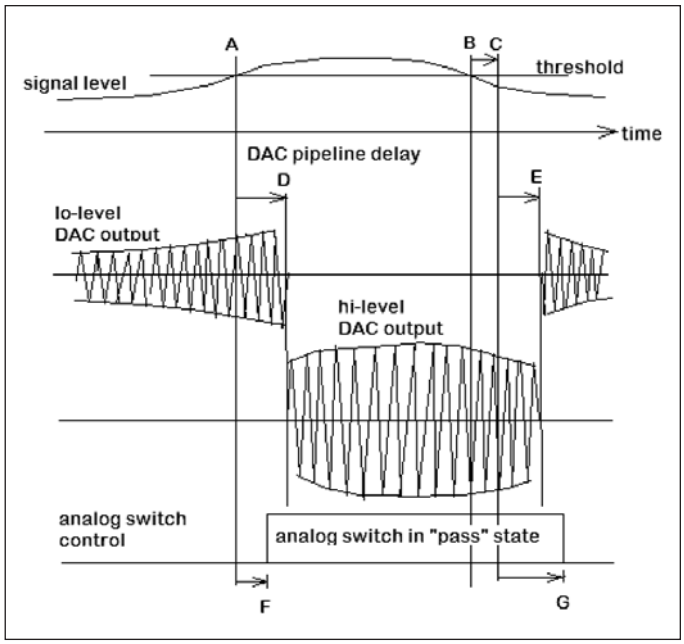

Figure 3 · Switch timing diagram, showing the sequence of events required for the transition between the two DACs.

In practice, the high-level-to-low-level switching transition will need to be delayed in time for at least the duration of one signal cycle, to prevent the switching state from changing sample-to-sample within a cycle.

Timing Considerations

The actual timing of the switching events is shown in Figure 3 for a typical excursion of signal level past the switching threshold and back down. The threshold is exceeded at time A, and digital input is immediately switched from the low-level DAC to the high-level DAC. The outputs of the DACs typically show the change only after a pipeline delay (the processing time of the DAC) has elapsed, from A to time D.

At time B, the signal level returns below threshold and remains there through C. An inertial delay (B to C) is applied to the falling transition to avoid fluctuation when the signal level is rapidly varying (alternately, hysteresis can be used). From time C, digital input is again routed to the low-level DAC, with a pipeline delay to time E, until the DAC outputs are affected.

The analog switch "pass" time interval, from F to G, brackets the time during which the high-level DAC output is carrying the signal (D to E), so that the switching takes place when no signal is present. Therefore, the block-to-pass transition is delayed by an amount (A to F) that is less than the DAC pipeline delay, and the pass-to-block transition by an amount (C to G) that is greater than the pipeline delay.

Possible Sources of Error

With the digital data switching between high- and low-level paths and the addition of the analog switch, the following factors may give rise to additional noise and/or distortion components in the analog output:

- High amplitude compression and video feedthrough in the analog switch
- Crosstalk to the low-level DAC output from the digital data inputs
- Gain, direct current (DC) offset, and/or phase mismatch between the paths

Switch Distortion

The analog switch is required to accommodate the maximum level of the amplified high-level DAC output without significant compression or nonlinearity, and without excessive insertion loss. These are minimized through proper selection of the analog switch.

In addition, video feedthrough (capacitive coupling of the switch control transitions into the signal path) will occur to some degree, depending on the characteristics of the switch and the speed of the transitions. Use of a switch with double-rail control inputs may reduce this coupling through partial cancellation of the opposite-polarity coupling components. The magnitude of these effects is difficult to analyze, since they are dependent on the characteristics of the switch.

The effect of video feedthrough can be minimized by slowing down the edges of the switching signals, and/or the frequency of switching by increasing the B-to-C inertial delay. For the weather radar, for example, the selection between the two DACs can be made so as to accommodate the maximum signal within a range gate, rather than on a cycle-by-cycle basis, to reduce the frequency of switching transitions.

Digital Crosstalk

Crosstalk between the digital inputs to the DAC and the resulting analog output can cause a significant level of in-band noise. This is particularly true for small signals. The two's complement extension of the sign bit results in many bits changing simultaneously at the signal frequency as the sampled data values pass through zero.

An interesting characteristic of this crosstalk is that, since it bypasses the DAC internal circuitry, it is not subject to the internal DAC pipeline delay, and appears to precede the analog signal being generated by that amount. The crosstalk can be reduced by proper board design, but a further improvement can be gained simply by offsetting the sample values from zero by adding a small constant value to the digital data. The value should have zeroes and ones alternating, to reduce the length of carry propagation for small signal values.

A better approach (but one requiring a change on the part of the DAC IC vendor) would be to use a sign-and-magnitude representation for the digital input data, based on two's-complement but with the non-sign bits inverted when the sign bit is one. The result would be that only the sign bit changes at the signal frequency. The low-order bits would change at even multiples of the signal frequency, since they represent the sample magnitude, while the high-order bits remain at zero.

Gain and Offset Mismatch Error

To reduce gain mismatch between the high-level and low-level channels, the insertion loss of the switch must be balanced out as closely as possible, either by increasing the gain of the amplifier or by adding corresponding attenuation in the low-level path. If possible, a dual DAC (such as the TI DAC5687) should be used which provides built-in vernier adjustment of the relative gains (this dual DAC is intended for I-and-Q communication channel use, with the two channels intended to be gain-matched).

If the DAC outputs are DC-coupled, any baseline offset in the high-level DAC output will be modulated on and off by the analog switch. A trim adjustment may be needed to prevent a change in the DC offset of the summed output when the switch changes state. The possibility of phase mismatch should also be taken into account, depending on the frequency range. The high- and low-level analog paths should be matched as closely as possible to reduce phase differences.

The magnitudes of the possible mismatch error components are calculated below. In this analysis, we assume as the worst possible case a narrowband signal with mean frequency f_c and nearly constant amplitude fluctuating around the switching threshold. Switching is assumed to take place constantly at a mean frequency $f_{sw} << f_c$ and with a duty cycle of approximately 50 percent. The amount of DC offset mismatch can be represented as a percentage of the full-scale output value. Likewise, the amount of gain mismatch can be represented as a percentage difference between one channel and the other. Phase mismatch for a narrowband signal is simply the phase difference expressed in degrees or radians.

DC Offset

Let b be the size of the step change in DC offset at the output, as a fraction of full scale (–FS to +FS), when the switch changes state. Then the effect of the continual switching is to introduce a square wave signal with peak-to-peak amplitude b into the output. With a baseline shift of $b/2$, this is equivalent to a symmetrical square wave of amplitude $b/2$. The wideband level of this square wave is $-20 \log_{10} b - 3$ dB relative to a full-scale sinusoid, giving a wideband SNR of $20 \log_{10} b + 3$ dB (the 3 dB term, rather

than 6 dB, accounts for the RMS-to-peak of 1.0 for a square wave vs. 0.7071 for a sinusoid).

The frequency spectrum of a square wave of amplitude $b/2$ and frequency f_{sw} consists of terms at the odd multiples, nf_{sw} (n = 1, 3, 5, …), with peak amplitudes $(b/2)(\pi/2n)$. The amount of square wave noise power falling within a band of width Δf centered at f_c would therefore be approximately equal to $(1/2)\ (b/2)^2(\pi/2n_c)^2$ with $n_c = f_c / f_{sw}$, the power of a harmonic falling in mid-band, times the approximate number of harmonics falling within the band, $(1/2)(\Delta f / f_{sw})$. This reduces to:

$$(1/2)(b/2)^2\ (\pi/2)^2\ (f_{sw} / f_c)^2 \cdot (1/2)(\Delta f / f_{sw})$$

$$= (1/2)^2\ (b/2)^2\ (\pi/2)^2\ (\Delta f / f_c)\ (f_{sw} / f_c)$$

In other words, the amount of square-wave noise power falling in the band is proportional to the DC offset, times the bandwidth as a fraction of the center frequency, times the frequency of the square wave as a fraction of the center frequency. Expressed as SNR relative to the maximum signal amplitude, this gives:

$$\text{SNR} = -0.943 - 20\log_{10} b + 10\log_{10}(f_c / \Delta f) + 10\log_{10}(f_c / f_{sw})$$

deducting 9 dB for the RMS/peak-to-peak of the narrowband signal, and with the frequency ratios inverted to give quantities greater than 1. The actual SNR will likely be greater since the above assumes an ideal square wave with zero switching time. A finite switching time will cause the square-wave harmonics to decrease in amplitude at a rate faster than $1/n$, and accelerating as frequency increases.

For $\Delta f > f_{sw}$, as is likely, the in-band noise will appear in the time domain as a series of transient bursts centered at f_c, arising from the individual step changes of the DC offset convolved with the time response of the Δf band-pass. The peak amplitude of the bursts will be independent of f_{sw} and depends only on the amplitude b of the step changes, the switching time, and the shape of the band-pass step response.

Gain Mismatch

Let us assume that a particular data bit pattern within the range of the low-level DAC yields an output of V volts through one DAC channel, and $(1 + d)V$ volts through the other, i.e., the path gains are in the ratio $1:(1 + d)$ where d is a small value, which may be taken to be positive for the present. Then, the effect of the continual switching as in our worst-case example is to introduce an intermittent term in the output, equal to the narrowband signal multiplied by d or by zero according to the switching state.

For $d << 1$, this is approximately the same as a term equal to the narrowband signal multiplied by $+d/2$ or $-d/2$ according to the switching state. This is essentially a frequency-shifted version of the square wave in the preceding section, now centered in the signal band itself with amplitude of $-20\log_{10} d - 6$ dB relative to the maximum signal amplitude, giving SNR of $20\log_{10} d + 6$ dB within the signal band. The spectrum of the signal will appear to have noise sidebands spaced by multiples of the frequency f_{sw} of the square wave representing the switching, and falling off inversely with distance from the signal center frequency.

Phase Mismatch

The analysis for gain mismatch also can be applied to phase. Let d be the difference in phase expressed in radians (i.e., the value in degrees times $\pi/180$). Then for $d << 1$, this gives a term equal to the narrowband signal multiplied by $+d/2$ or $-d/2$ as above, but rotated in phase by 90 degrees, so that it adds in quadrature to the original narrowband signal. The SNR is again $20\log_{10} d + 6$ dB, and the added sidebands will behave similarly.

Mismatch Correction

The mismatch errors analyzed above can be corrected by analog and/or digital means.

Offset error involves only the high-level DAC channel, and can be corrected by injecting a small DC current at the switch input using a resistor and trimmer potentiometer. If the DAC is DC-coupled, a small constant value can be added to the digital input, but this creates the possibility of overflow if the signal value is close to full scale. If the DAC output is AC-coupled (i.e., through a transformer), digital correction will have no effect. The correction is then adjusted so as to cancel any output step when the analog switch changes state.

Gain mismatch can be corrected either by introducing a small amount of analog attenuation in one DAC channel, or by multiplying the digital input to the low-level DAC by a value $(1 + x)$ where x will likely be a small positive or negative correction value. The multiplication is done on the full length data word to allow for growth of the sample value when x is positive, and the amplitude threshold test is now applied to the result of the multiplication (if the high-level DAC is selected, it is driven with the high bits of the original unmultiplied data word). Again, the correction is applied to the low-level channel to avoid the possibility of overflow in the data to the high-level DAC. The value of x is chosen by adjusting it so as to minimize modulation of the output by the switching frequency while applying a constant-level sinusoidal digital signal a few dB below the saturation level of the low-level DAC, and with the switching forced at a fixed rate by applying a square wave from an external source.

The DAC itself may include provision for gain adjustment (this is the case with the TI DAC5687, which

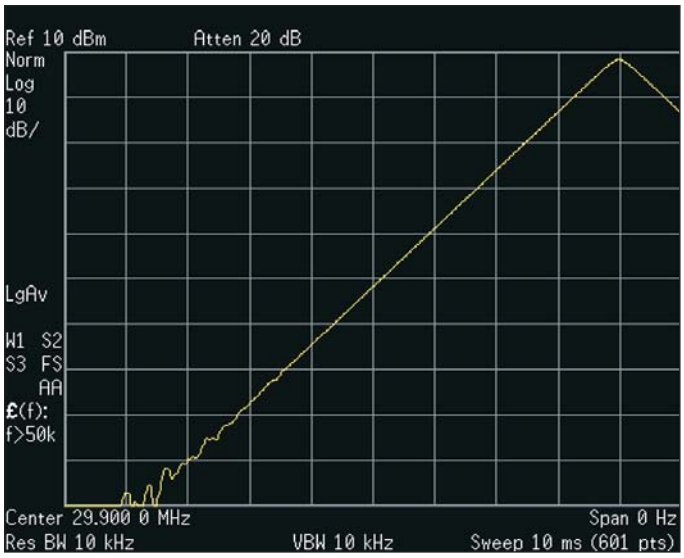

Figure 4 · DAC output (spectrum analyzer, zero span mode)—Positive ramp.

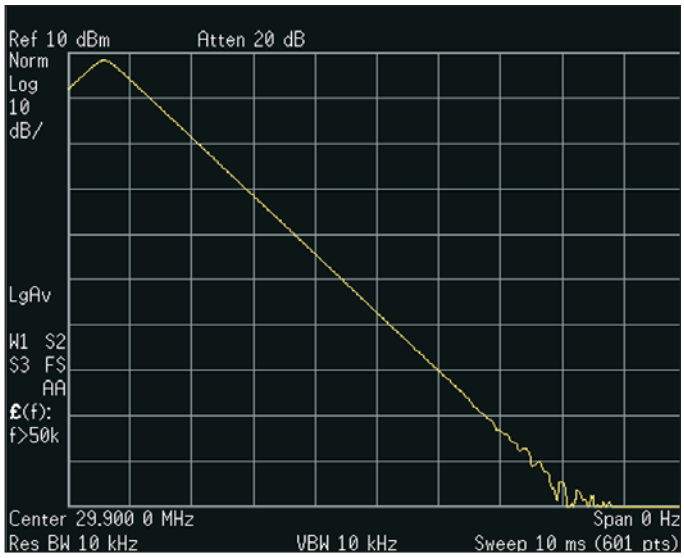

Figure 5 · DAC output (spectrum analyzer, zero span mode)—Negative ramp.

includes both coarse and fine adjustments, intended to match the two channel paths in compensation for mismatch in the external circuitry). If so, then the effect is equivalent to adjusting the analog gain or attenuation in each channel, and digital overflow need no longer be taken into account.

Phase mismatch is more difficult to compensate, but one digital approach that is valid for small error and narrow-band signal is augmenting the $(1 + x)$ factor to include the sum and difference with the previous sample, giving the expression $[1 + x(1 + z^{-1}) + y(1 - z^{-1})]$ for the multiplier, where z^{-1} represents a delay by one sample. Applying the sum and difference $(1 + z^{-1})$ and $(1 - z^{-1})$ to the narrow-band signal gives two components approximately 90 degrees apart in phase, so that adjustments to the coefficients x and y are essentially orthogonal. Re-expressing the above, if at any time s_0 represents the current (full-word) input data sample and s_{-1} the previous sample, then

s_0 is replaced by: $s_0 + (x + y)\, s_0 + (x - y)\, s_{-1}$

as the input to the low-level DAC (taking the lower bits). The values x and y then can be adjusted in sequence to minimize switching-frequency modulation sidebands at the output, thus cancelling the gain and phase mismatch errors simultaneously.

Measured Results

Figures 4 and 5 show the envelope of a digitally generated 29.9 MHz signal modulated by a 105 dB exponential ramp. The output is produced using a dual 16-bit DAC with 24 dB offset according to the method described, and displayed on a decibel scale using an Agilent E4440A spectrum analyzer in zero-span mode. The digital sample rate is 120 MHz, with the DAC output interpolated to 240 MHz.

Since the scale range of the analyzer is only 100 dB, the lower end of the ramp is not visible. Note that the ramp is linear and that there is no discernible discontinuity at the switching level (30 dB below the maximum).

References

1. Ed Crean, Paul Hiller, "A Wide Dynamic Range Radar Digitizer," *High Frequency Electronics*, Sept. 2008.

2. TI DAC5687 data sheet, Texas Instruments, September 2006.

Author Information

Dr. David Friedman serves as a senior principal engineer for Symtx. He can be reached at: dfriedman@symtx.com. Paul Hiller is chief technology officer for Symtx. Located in Austin, Texas, Symtx Inc. is a manufacturer of functional test systems for electronic systems including engineering, depot, and production line applications; commercial and military satellite systems; and commercial and military avionics systems.