Plan for a 50 MHz Analog Output Channel

Lawrence R. Doolittle, LBNL

August 8, 2002

To digitize a narrow band 50 MHz analog input, it is well understood that one can cheat the normal Nyquist sampling rate rule. While sampling at 200 MS/s will pull in sequential I/Q/-I/-Q samples (suitable for clean vector manipulation) such samples are also available when sampling at $200/(2n+1)$ MS/s. The $n = 2$ case, 40 MS/s, is used repeatedly in the SNS project. Of course, as n increases, the analog input filter must be made narrower and narrower to avoid unwanted aliasing. This method works well because the analog bandwidth of modern ADC chips is much higher than 1/4 their maximum sample rate; the ADS808 used in the LBNL MEBT LLRF board has a maximum sample rate of 70 MS/s, but an analog bandwidth of 1 GHz.

The converse problem, that of generating a narrow band 50 MHz waveform, appears less well developed. Certainly a 200 MS/s DAC would work in theory, but as above, we would like to cheat and reduce the update rate. Such data rates are not plausible with mid-range FPGAs, the selection of high resolution DAC chips is sparse, the edge speeds required to drive the DAC are high and noisy, and clock synchronization margins are tight.

One technique, used by both the LBNL MEBT LLRF board and the LANL LLRF board, brings in a separate 50 MHz source, which is vector modulated with signals from twin, low data rate (20 MS/s) DACs. Beyond the usual questions of channel balance and phasing, the LBNL experience was that the particular vector modulator used was rather nonlinear, and the extra 50 MHz input wiring decreased reliability.

This note develops a design based on a direct synthesis technique, using a data stream considerably slower than 200 MS/s. As a lead-in to the final design, consider what happens if one feeds I/Q/-I/-Q samples at 40 MS/s to a single high speed DAC. The analog output will generate spectral lines not only at 10 MHz, but at $10\,\text{MHz} \cdot (2n+1)$. Low order lines include 30 MHz, 50 MHz, and 70 MHz. Power in these lines scales as f^{-2}, ignoring the intrinsic DAC output filter ($\tau \sim 2$ ns for a DAC902). Assuming +7 dBm total DAC output power, and accounting for the use of an existing 50 MHz RLC output filter, the power at the first few harmonics is:

f (MHz)	P_{DAC} (dBm)	A_{filt} (dB)	P_{out} (dBm)
10.0	6.09	-86.09	-80.00
30.0	-3.45	-56.45	-59.90
50.0	-7.89	-2.50	-10.39
70.0	-10.81	-52.29	-63.11
90.0	-13.00	-70.00	-82.99
110.0	-14.74	-80.55	-95.29
130.0	-16.19	-88.22	-104.41

Next, let's turn this concept into something more usable. The 10 MHz, 70 MHz, and 90 MHz components can be suppressed digitally by clocking the DAC at 80 MS/s instead of 40 MS/s. This rate is still within the "easy" realm for mid-grade FPGAs and DACs. The interpolation logic within the FPGA can be fairly simple: given a string of 40 MS/s data points, nominally representing $\pi/2$ phase advances in a 10 MHz waveform

$$a_i, \quad a_{i+1}, \quad a_{i+2}, \quad \ldots$$

interleave with them points with the same amplitude, but a phase that is interpolated to $-3\pi/4$.

$$a_i, \quad -(a_i + a_{i+1})/\sqrt{2}, \quad a_{i+1}, \quad -(a_{i+1} + a_{i+2})/\sqrt{2}, \quad a_{i+2}, \quad \ldots$$

The resulting waveform nominally represents a 30 MHz sine wave, but like the previous case includes a strong 50 MHz component. An effective low-dispersion analog notch filter can be built from a delay line and a splitter-recombiner pair, to suppress the 30 MHz component. In concept:

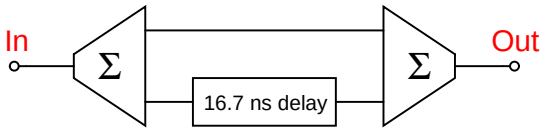

Combining the digital and analog filters, and assuming 3% error in notch frequency, the power flow now looks like this:

f (MHz)	P_{DAC} (dBm)	A_{notch} (dB)	A_{filt} (dB)	P_{out} (dBm)
30.0	4.89	-25.33	-56.45	-76.89
50.0	0.45	-0.84	-2.50	-2.89
110.0	-6.40	-0.47	-80.55	-87.42
130.0	-7.85	-2.78	-88.22	-98.84

Compared to the 40 MS/s concept, this shows 7.5 dB more signal at 50 MHz, and the largest sideband improves from -49.5 dBc to -74.0 dBc.

The above description is relatively abstract; perhaps a graph of the waveforms will make the concept more apparent. The following figure shows the nominal 30 MHz output of a perfect 80 MS/s DAC, the result of averaging that waveform with a 16.7 ns delayed copy of itself, and the 50 MHz component of that result. Note that the waveforms repeat every 8 DAC samples, or 100 ns.

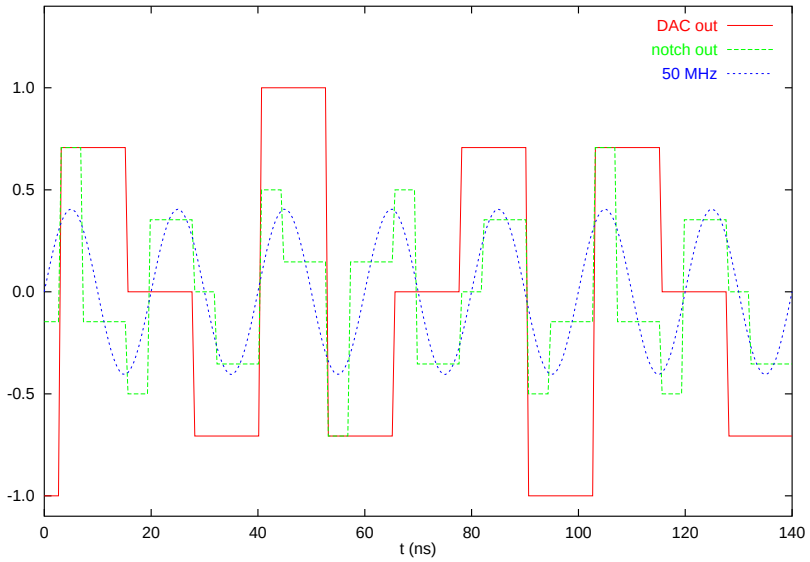

The following circuit abandons the input splitter in favor of directly driving a 25 Ω node with the high impedance output of the DAC, where a 2:1 turns ratio transformer increases the effective impedance seen by the DAC (differential mode) to 100 Ω, a good match for the output compliance of a DAC902. Each output pin sources 0 to 20 mA of current, so the signal strength is ±10 mA, and each pin delivers $\frac{1}{2} \cdot 10\,\text{mA} \cdot 0.5\,\text{V} = 2.5\,\text{mW}$ average power, totalling +7 dBm for the two pins. The 30 MHz component of the DAC output is absorbed by the combiner. The conventional 50 MHz RLC post filter is not shown.

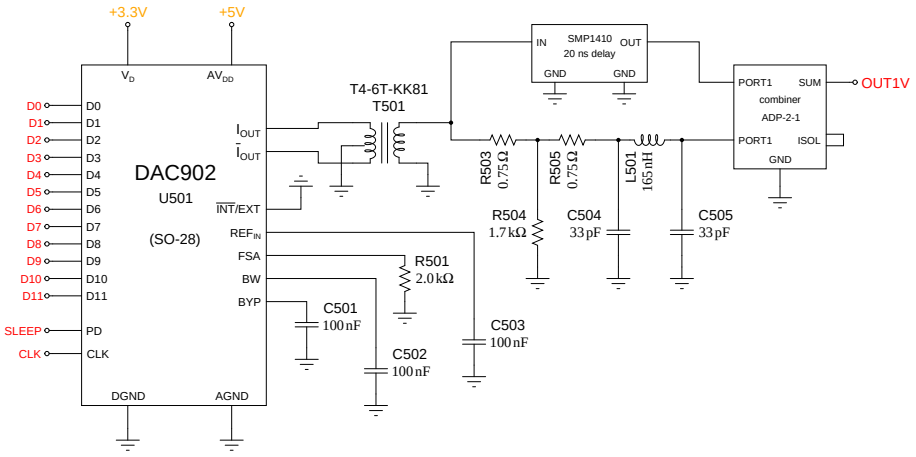

Since a 16.7 ns delay line is not available off-the-shelf, the circuit employs a 20 ns part, and adds 3.3 ns of compensating delay to the other leg. An imbalanced 165 nH inductor and a 66 pF capacitor add such a delay in a 50 Ω system;

the π configuration shown uses capacitors of half that value. R503 through R505 are set up to balance the attenuation through the delay line, nominally 3%. Depending on the amount of 30 MHz rejection desired, one could hand select the components, install the nominal values, or just line up three zero-ohm jumpers for R503, R505, and L501.

To start up a cold DTL tank, the controller needs to generate output in the 50.0 to 50.4 MHz band. This is readily accomplished by a gradual, constant rate phase shift in the 10 MHz signal sent to the interpolator and DAC. Another description of the implementation is a 400 kHz single sideband modulation of the carrier. Successive 40 MS/s samples, that nominally show an I/Q/-I/-Q pattern with $\pi/2$ phase advance between them, have their phase shift adjusted to $(1+\delta)\pi/2$, where to reach 10.4 MHz, δ can be as high as 0.04. The interpolation step described earlier multiplies the waveform by a 40 MHz carrier, which turns the original 10.4 MHz into 29.6 and 50.4 MHz outputs.

The digital processing necessary to introduce controlled phase modulation is well understood. A high resolution phase accumulator is followed by a CORDIC rotator. Both are easily implemented in an FPGA; a 24-bit accumulator plus a 12-bit accurate high-throughput rotator take up about 340 cells. The current LBNL controller uses 40% of the 3456 available in a Xilinx XC2S150. The rotator introduces about 12 cycles (40 MHz) of latency, which is not an issue because the phase velocity will be constant during a macropulse, and this block is outside the primary feedback loop.

Sideband modulation imbalances both the 30 MHz notch filter, and the digital filter that suppresses the 10 MHz and 70 MHz lines. These latter two lines will reappear proportionally to δ – calculations at 400 kHz modulation show a 10.4 MHz spur at -29 dBc at the output of the notch filter, probably unmeasurable after the analog filter. The 70 MHz component is even smaller.

In theory, one could change a handful of components (VCXO and delay), and operate this circuit at 66/133 MHz instead of 40/80 MHz. The notch filter moves from 30 MHz to 83.3 MHz (16.7 ns to 6.0 ns delay), and the divisors programmed into the PLL chip (which locks the acquisition clock to the 50 MHz phase reference) change. The data sheets say that this is still within spec for the ADC, DAC, and PLL, but the margins for these parts and the FPGA get mighty thin. While the output power, output spectrum, and system latency would all improve, the technical risk is considered too high to pursue this option in the short term.