7-bit 5MSPS Current Steering DAC

Kevin C. Fronczak

***Abstract*— This paper analyzes the design of a 7-bit, 5MSPS, 225 mW Current Steering DAC in AMI 0.5µm technology. Detailed circuit analysis is presented along with simulations and specification compliance. A full IC layout is also presented.**

I. Introduction

The goal of this project was to create a 7-bit DAC that operated at 5 MHz and utilized Thermometer decode logic to guarantee monotonicity at the output. The design process took place within roughly an 8-week period from an initial 'pencil-and-paper' design to a full integrated circuit layout.

In order to complete the design, multiple block has to be created, shown functionally in Figure 1. The Current Reference block consisted of an opamp-driven constant current source attached to an external 0.1% precision resistor. The current is then mirrored to the Current Manifold block which consists of eight current-mirrors which drive each row of the Current Switching Cell block. This block is divided into eight rows and sixteen columns. Each cell corresponds to one LSB of current and they are selected via signals from the decoder. Each cell has localized logic and re-timing circuitry as illustrated in Figure INSERT FIG NUM. The currents from each LSB cell are then summed at the output which is loaded by an external 37.5 Ω resistor.

II. Theory of Operation

A) Current Reference

The current reference circuit is, without a doubt, the most crucial cell within the entire DAC. Without an accurate reference current, the rest of the DAC would be useless since no static specification could possibly be met, bar power consumption. As such, a proper analysis of current-source transistor sizing and of the op-amp design is incredibly important. Figure 2 shows the architecture for the block.

The specification was for the output to be within 30 mA by ±0.5LSB which corresponds to a ΔI_{FS} of 118.1 µA.

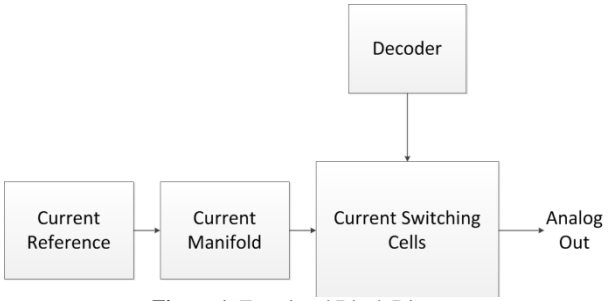

Figure 1. Functional Block Diagram

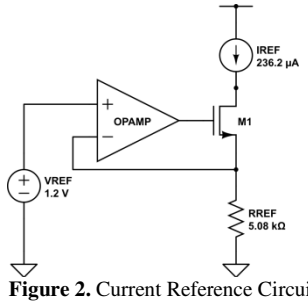

Figure 2. Current Reference Circuit

$$A_v = \frac{v_o}{(v_{in+} - v_{in-})} = \frac{v_o}{\Delta V_{FS}/127} \quad (1)$$

Because the DAC drives a load of 37.5 Ω, this ΔI_{FS} will correspond to a maximum tolerable full-scale voltage difference of ΔV_{FS} = 4.429 mV. From here, the minimum tolerable op-amp gain can be calculated, shown in (1).

Since the maximum tolerable output voltage difference is known, dividing this value by the number of current cells will provide a value for the maximum tolerable offset error across the external reference resistor. Since one node of this external resistor is being fed into v_{in-}, and v_{in+} sees a 1.2 V reference, this offset is also the maximum tolerable offset of the opamp which evaluates to 34.87 µV.

$$A_v = \frac{v_o}{\Delta v_{in}} = \frac{V_O - (V_T + V_{eff} + V_{REF})}{\Delta v_{in}} \quad (2)$$

The last remaining unknown is the output voltage, v_o. In order to determine this, (2) was used. V_T was assumed to be 1 V due to body-effect (which was verified by simulation) and V_{eff} was taken as 0.35 V. Because the opamp is operated off of a single rail and referenced to ground, it was assumed the output would sit at 2.5 V. This is a necessary consideration since the value for v_o will need to be referenced to this voltage. Plugging these values into (2) yields a $|v_o|$ of 0.05 V and a corresponding gain of 1433.9 V/V, or 63.13 dB. Obviously, this value has some flexibility due to the dependence on V_{eff} (and thus, the transistor size) and the voltage at the output of the amplifier with no signal applied to the inputs. Using values extracted from simulation, the output of the amp ended up sitting at 2.55 V with a V_{eff} of 0.312 V which results in a gain of 1089.76 V/V (60.75 dB) – a 24% decrease in required gain with just a 2% increase in amp output voltage and 11% decrease in the drive transistor's V_{eff}. That said, however, the required gain could easily **increase** after the layout due to various parasitics. Thus, to be safe, a minimum gain of 75 dB was used for the actual op-amp design.

The opamp architecture used is shown in Figure 3. It features a simple PMOS differential pair with an NMOS common-source amplifier as the second stage. A constant-gm network was used to generate bias voltages, but there are limitations in this application as will be explained later. Finally, a 2 pF compensation capacitor was added to improve phase response of the circuit.

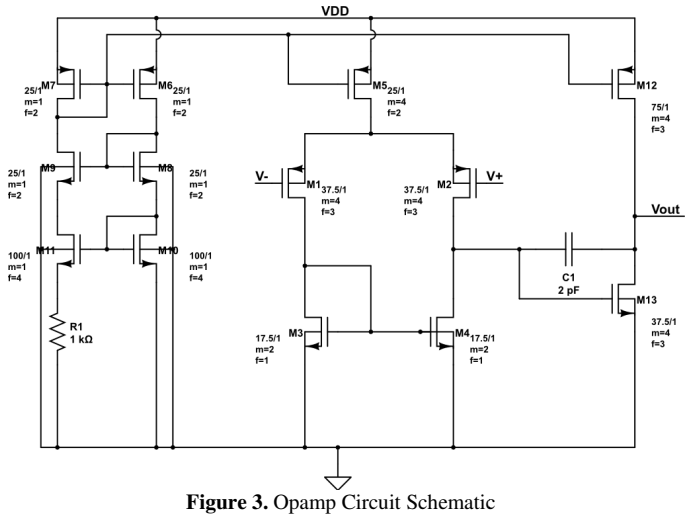

Figure 3. Opamp Circuit Schematic

The design for the constant-gm network was taken from [1]. From that, a few design choices were made: first that each differential transistor would receive 15 µA, and second that the common-source stage would receive 90 µA. These values were chosen rather arbitrarily based off of the performance of the bias network which runs at roughly 7.5 µA based on simulation.

Using these values for current, the minimum requires size for the differential pair, via (3), with $\mu_p C_{ox}$ = 40 µA/V^2 and V_{eff} = 0.1 V is 75 µm/µm. For the common-source stage, with $\mu_n C_{ox}$ = 120 µA/V^2 and V_{eff} = 0.1 V the minimum size requirement is 150 µm/µm. Because of the DAC's critical dependence on an accurate reference current (and, thus, a low offset opamp), the differential pair size was doubled to be the same size as the common-source NMOS. The pair was also split into four parallel transistors with three fingers each in order to allow for a common-centroid layout (as will be seen in Section V).

$$\left(\frac{W}{L}\right) = \frac{2I_D}{\mu C_{ox}\left(V_{eff}\right)^2} \tag{3}$$

The gains associated with each stage, then, are 140 V/V for the differential stage and 100 V/V which translated to a total gain of 14000 V/V, or 83 dB. This leaves roughly a 20 dB error margin before reaching the minimum allowable gain for a ±0.5LSB variance at the output. The thinking was that by allowing the opamp to have a large margin of gain, any offsets introduced by layout parasitics or any other non-idealities can be minimized and the offset error requirement will still be met.

Of course, by having so much gain in each stage, the op-amp is just waiting to oscillate once feedback is introduced. The given architecture, taking away the compensation capacitor, is quite unstable (a phase-margin of -23°, in fact). In order to stabilize the amp, Miller's Theorem was used to choose a capacitance that would decrease the bandwidth of the amp to the point where it is unconditionally stable. The equation used can be seen in (4) where the gain, A_v, is equal to 100 V/V (as calculated earlier) and $L_{ov}C_{ox} \cong 0.2$ fF/µm. This yields a compensation capacitor of 3 pF. This yielded a phase margin of roughly 70°, so to try and save on the total area of the opamp, the value was decreased until it was at 2 pF which gave a phase margin of 63°. This was the final value chosen due to the decrease in area while still maintaining a good value for phase-margin.

$$C_{eq} = A_v W L_{ov} C_{ox} \tag{4}$$

As was previously mentioned, there is a bit of a problem in using a constant-gm bias network in this architecture. Due to the nature of the DAC, there is no room for a secondary external precision resistor. This means that the resistor in the constant-gm network must be on chip which, in turn, implies a very low precision resistor. This is non-ideal since the transconductance of the amp is inversely proportional to the bias resistor. If there is a 30% variance in the resistor there will be a 30% variance in transconductance and since gain is proportional to transconductance there will be a 30% variance in gain of each stage. At worst case, this means that each stage could conceivably be 30% lower than what was designed. However, since gain was overdesigned in this circuit, the worst-cast variance results in a system gain of 76.8 dB (50% lower voltage gain than with a high precision resistor) which is still 13 dB higher than the minimum required gain for a low offset, so it is acceptable. Realistically, there is no reason to use a constant-gm bias network in this circuit because the fact that there is an on-chip resistor with high variance causes it to cease to be a constant-gm network.

Figure 4 shows the current source given in Figure 2. This is a simple wide-swing mirror and is connected to the NMOS shown in Figure 2. In order to select proper sizes, first M1 in Figure 2 needs to be considered. Since the opamp will be sitting at roughly 2.5 V by default, V_{G1} is 2.5 V. It is also known that V_{S1} will be 1.2 V during operation (of course, it is unlikely that V_{G1} will also be at 2.5 V in this case, but it is a good starting point for sizing considerations). Using values of V_{T0} = 0.7 V, γ = 0.5 V$^{1/2}$, and ϕ_F = 0.35 V and plugging into (5), the threshold due to body-effect was calculated to be 0.97 V.

$$V_T = V_{T0} + \gamma\left(\sqrt{V_{SB} + 2\phi_F} - \sqrt{2\phi_F}\right) \tag{5}$$

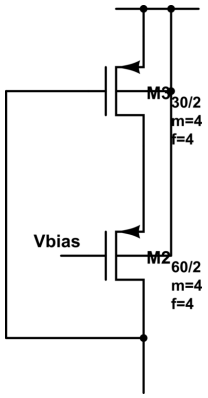

Figure 4. Wide Swing Mirror

Using these values, $V_{eff,1}$ was calculated to be 0.33 V which, when plugged into (3), yields a required W/L of 36.15 µm/µm which was rounded up to 40. This new value yields a $V_{eff,1}$ of 0.314 V (as opposed to the 0.33 V calculated).

From here, the value for the drain voltage of M1 can be found which also corresponds to the gate voltage of M3. IN order to be in saturation, V_{D1} must be at least equal to a threshold voltage above the gate, or 3.47 V (again, this is also V_{G3}). The next step is to calculate the required value for V_{bias}. Here, V_{T2} was estimated to be 1.1 V due to the body effect which results in a required V_{bias} of 2.37 V.

To find the necessary size requirements of M2 and M3 in Figure 4, $V_{eff,2}$ and $V_{eff,3}$ need to be calculated. Since V_{eff} can be given by (6) where $V_S = V_{DD} = 5$ V, $V_G = 3.47$ V, and $V_{T0} \cong 1$ V, $V_{eff,3}$ is equal to 0.53 V. This results in a required W_3/L_3 of 42 µm/µm. However, having a $V_{eff,3}$ of 0.53 is rather worrying (since it decreases the amount of swing available to the transistor pair) so the sizing ratio was increased to 60/1 to allow for a little more swing. This results in a $V_{eff,3}$ of 0.44 V which allows for roughly 100 mV more of swing.

$$V_{eff} = \sqrt{\frac{2I_D}{\mu C_{ox}\frac{W}{L}}} \quad (6)$$

Finally, to find the size of M2, a decision was made to simply set its V_{eff} equal to that of $V_{eff,1}$ (which was 0.314 V). This yields a size requirement of 119.8 µm/µm ≈120 µm/µm.

However, there's one more consideration: transistor matching. Because this pair will be providing the reference current to eight other mirrors (which then break out to sixteen LSB cells each), increasing the size is incredibly important for matching. Obviously the sizing ratios should be kept the same, so the option here is to just increase the total area while maintain the 60/1 and 120/1 ratios for transistors M3 and M2, respectively. The final sizes were chosen as 120/2 for M3 and 240/2 for M2 after various Monte Carlo matching simulations. One oversight, however, is that the sizing of M3 is more critical for matching that that of M2. This is because M3 is the transistor that acts, primarily, as the current source while M2 is more useful for increasing output impedance. As such it would've made more sense for M2 to be **smaller** than M3, rather than what is presented here.

B) Current Manifold

The design of the current manifold directly hinges on the sizes of the wide-swing source analyzed in the previous section. These sizes must be maintained to ensure proper reference current matching. Figure 5 shows how one manifolds are cascaded (there are a total of eight in the DAC). The transistors on the left are used to mirror current into the diode-connected NMOS transistor which mirror over to the NMOS on the right. The PMOS transistors on the right then drive sixteen LSB cells. $V_{bias1,2}$ are generated from the mirror in the current reference block and correspond to the gate voltages on M3 and M2 in Figure 4.

The sizing of the transistors was rather straight-forward; values for V_{eff} were chosen to be 0.3 V which yields a sizing ratio o f 43.7/1 (decreased to 42.5/1 after simulation).

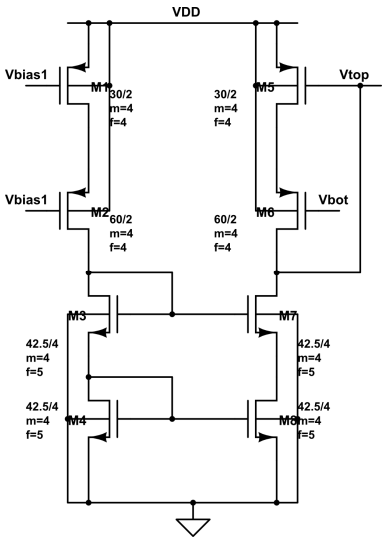

Figure 5. Current Manifold (2 stages)

C) Thermometer Decode Cells

In order to guarantee that the DAC is monotonic, thermometer logic was used. In order to properly address the 127 LSB current cells, the incoming 7-bit digital word needs to be decoded. To simplify the logic, the decoder was split to address rows and columns, as shown in Figure 6 (adapted from [1]). To achieve this, the row decode would need to be a 3:7 decoder and the column decode a 4:15 decoder. Because the decoder is split up to handle the three MSBs (row decode) and four LSBs (column decode), 15 cells need to have their row active when the MSBs are all 0 (binary 0 to 15) and seven cells need their column tied high when at least one MSB is high and all LSBs are low. This adds up to the needed 127 cells (the architecture of which is described in part D).

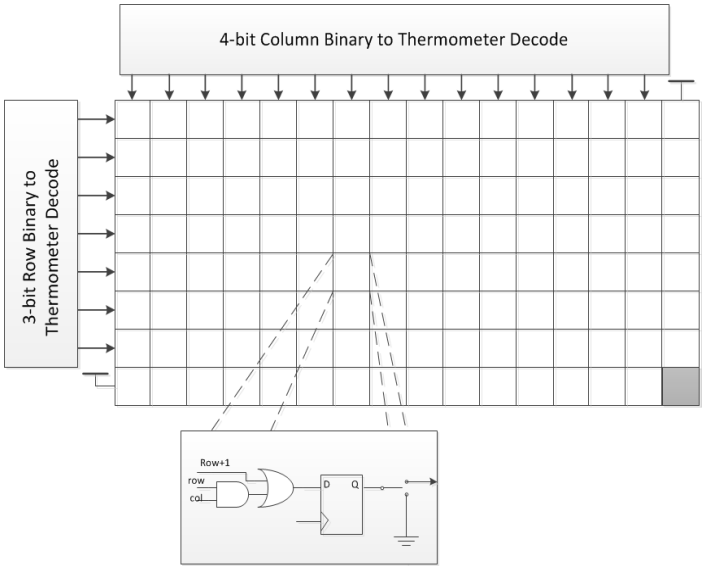

Figure 6. Thermometer Decode and LSB Current Cell Block Diagram

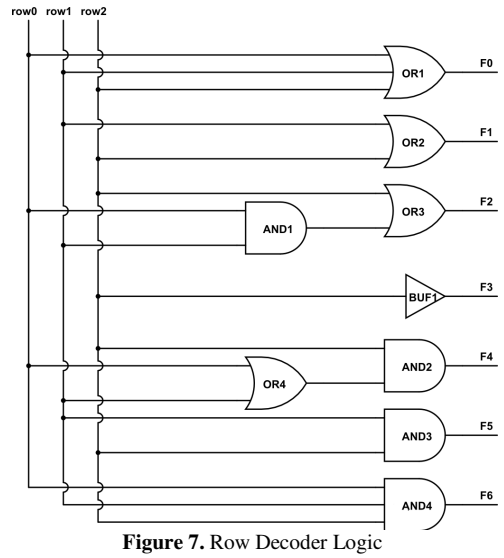

Figure 7. Row Decoder Logic

Figure 7 shows the logic used for the row decoder. The devices were sized by equating the PMOS and NMOS devices to unit resistors where a single NMOS has unit resistance R and a single PMOS has unit resistance $2R$ (essentially, the resistance of a single transistor is given by $R = \frac{L}{W}\frac{1}{\mu C_{ox}}$ and because μC_{ox} is a constant, only the transistor size can be changed to modify its resistance). For an example calculation, a 2-input NAND-gate will, in the worst-case, have one PMOS on and two NMOS on This means that the output of the NAND-gate will see $R_P||R_N = 2R||2R = R$. This means that each transistor should be sized as being equivalent to one unit resistance. This size can be determined in two ways: arbitrarily, or using a concept known as fan-out. Using fan-out allows a designer to size transistors to minimize the amount of propagation delay throughout a logic circuit by taking into account both load and input capacitances (among other things). A design using fan-out is crucial if one requires a high-speed DAC. However, seeing as how 5 MHz isn't really high-speed by any means, the device sizes were chosen rather arbitrarily. Given that this is a design with a very fast turn-around time, forgoing some delay-time optimization for a quick design and simulation schedule seemed a logical choice. Given this, a single value for R was set to correspond to a W/L ratio of 8/1.

The process to design the decoder shown in Figure 7 was fairly simple: create a truth table. From there, the relationships between the input and output codes became clear and trivial to implement. The same method was used for the column-decode which isn't shown here. If propagation delay were in issue, aside from using fan-out, another technique would be to utilize only NAND or NOR gates since there is only one delay associated with either gate as opposed to two for an AND or OR (since those are just NANDs/NORs with inverters on the output).

D) LSB Current Switching Cell

The current switching cell shown below in Figure 8, and adapted from [2], required a bit of care to design properly. First, the current source needed to be analyzed to guarantee that the output resistance was significantly larger than the load resistor of 37.5 Ω in order to ensure that the switching process didn't exhibit any non-ideal loading behaviors. Since the sizes of the cascoded pair were determined in part A, calculating R_{out} is fairly straight forward. The approximate output resistance for a cascoded source is given by (7) where g_m was calculated to be 1.5 mS and r_{ds} was calculated to be 425 kΩ with a λ equal to 0.01 V^{-1}. This results in an R_{out} equal to 271 MΩ and, thus, a full-scale R_{out} of 2.13 MΩ which, clearly, is significantly larger than 37.5 Ω.

$$R_{out} \cong g_m {r_{ds}}^2 = \sqrt{2I_D \mu C_{ox} \frac{W}{L}} \cdot \left(\frac{1}{\lambda I_D}\right)^2 \qquad (7)$$

The next important task was to size the actual current switches (M3 and M4 in Figure 8) as well as determine the necessary bias voltage for M4. The reason a bias voltage was used for M4 was so that it didn't experience any switching transients caused by sharing a bias voltage with multiple switches. It is always on; it's just more weakly conducting than the on-state for M3 so it ends up sinking close to zero current when it should be off (which is ideal).

To determine the required switch bias voltage, V_{sw}, the maximum swing on the drain of M4 needs to be determined. Since the full-scale output current is known to be 30 mA, multiplying this by the 37.5 Ω load results in a $V_{D4,max}$ = 1.125 V. Since the gate of M4 must sit a threshold above the drain, Vsw must be around 2 V (ignoring body-effect for the moment). Thus the minimum voltage on the source of M4 will be at least a threshold above the gate of M2 which equates to 3.47 V (using previously calculated values for V_{G2} and V_{T2} of 2.37 V and 1.1 V respectively). From here, the V_{eff} of M4 can be found by taking the above $V_{S,4}$ voltage and subtractive off the gate voltage on M4 and the threshold voltage of M4 (which is assumed to be roughly 1 V). This results in a $V_{eff,4}$ = 0.45 V which results in a size of approximately 60/1. In order to increase the response time of the circuit, decreasing the capacitance at the output node is critical which implies a decrease in switch area. Thus the length of M4 was set to the minimum of 0.6 μm which results in a width of 40 μm. This yields a new V_{eff} of 0.54 V and after simulation it was determined that $V_{T,4}$ is close to 1.2 V which results in a V_{sw} of 1.73 V.

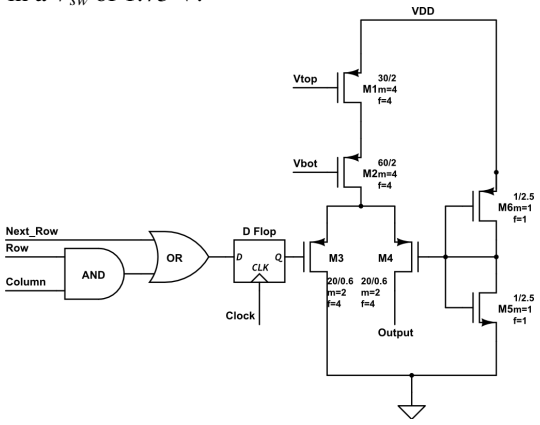

Figure 8. LSB Current Cell with Local Decode Logic

In order to create this 1.73 V bias voltage, a simple CMOS voltage divider was used. The reasoning, here, is that the device can be made very small, designed to consume very little power, and will decouple the gate voltage of each switch from the rest of the network (as opposed to driving each switch from the same node). The design was rather straight-forward: set the currents of the NMOS and PMOS equal to each other with their gat voltages set to the desired voltage (1.73 V in this case) and then solve for the W/L ratio. This resulted in $\left(\frac{W}{L}\right)_P = 0.618\left(\frac{W}{L}\right)_N$. Setting the NMOS size to 1.2/1 yielded a PMOS size of 1/1.35 though this needed to be iterated on in simulation to achieve the correct bias voltage (the final PMOS size was 1/2.5). It would've made significantly more sense to choose an NMOS size based off of target power consumption rather than choosing it arbitrarily as, in simulation, this divider consumes 20 µA of current which results in 12.7 mW of power (for, essentially, no good reason). This could, and should, be decreased substantially so that the divider still maintains a relatively stable voltage with a small amount of power consumed.

Finally, the last important part of the current switching cell is the local decode logic. Since the thermometer decoder was split into row and column decoders, each cell needs to be able to interpret whether or not it should turn on. This is done simply by saying "if row and column are on OR the next row is on, divert current to output". The 'or' is necessary because as the row MSBs begin to increase, all previous rows must stay on to properly indicate that code. The flip-flop is there in order to guarantee all bits arrive at the same time within each LSB cell, since the bits coming out of the decoder are un-latched. A flip-flop was chosen over a latch in order to eliminate any potential glitching at the input from appearing at the output.

III. Specification Compliance

Table 1, below, shows the specification compliance table. The corners are designated as follows: **NP** = worst-case power, **np** = worst-case speed, **T** = high temperature (100℃), **t** = low temperature (0℃), **v** = low voltage (4.5 V), and **V** = high voltage (5.5 V). Any specifications that were not met are highlighted in red.

Figure 9 shows the full-scale settling time over corners, along with the full-scale current. As shown in the specification table, the worst case settling time happened to be 161.7 ns which translates to a worst-case conversion time of 6.18 MSPS. The full-scale current is also shown here where it differs by 0.34LSBs from the ideal value of 30 mA and occurs at high temperature and high voltage (which, incidentally, corresponds to the worst corner for nearly all of the specifications). As is obvious from the waveforms in Figure 9, the DAC experiences some overshoot before settling to its final value. This can be directly controlled by proper sizing of the current switch. For example, decreasing the switch width by a factor of 2 eliminates all overshoot but the DAC then is subject to some loading effects which decreases the rise-time and pushes the settling time further out. Two very quick ways to decrease both the settling time and the overshoot would be to drop the control logic down to, say, 2.5 V while simultaneously decreasing the current switch width; decreasing the width, as previously discussed, will have a direct impact on decreasing overshoot while dropping the control logic supply will decrease the amount of transition time between the "on" and "off" states which will directly impact the settling time. The issue here is that care needs to be taken to ensure that when the switch should be directing current to the output that it is significantly more conductive than the switch attached to the flip-flop, otherwise some of the reference current will be steered to ground which will result in incorrect output values. Since this design met spec, the logic supply was kept at 5 V.

Performance Metric	Specification	Actual Performance	Worst Corner
Resolution	7 bits	8.64 bits	NPTV
Settling Time	< 200 ns	161.7 ns	npTv
Power Dissipation	< 280 mW	226.2 mW	NPtV
Current Mirror Matching	σ < 1.746 µA	σ = 1.82 µA	t
Full Scale Current	30 mA	30.08 mA	NPTV
INL	±1.0 LSB	+0.35LSB	NPTV
DNL	±0.5 LSB	-0.026 LSB	NPTV

Table 1. Specification Compliance Table

Figure 10 shows the Monte Carlo simulations for both the reference current (top) and current manifold (bottom) with temperature as 0℃, 27℃, and 100℃ left to right. The reference should have a very tight distribution with a mean, ideally, at 236.2 µA. The worst-case mean occurred at high temperature with a mean of 236.253 µA (which translates to 0.00014 LSBs) and the worst-case deviation occurred at low temperature with σ = 577.523 nA, or a 3σ value of 0.007 LSBs. This incredibly tight distribution is absolutely critical for a cell that is supposed to provide a current reference for the rest of the DAC. The current manifold was measured in reference to the reference current so the mean should sit, ideally, at 0 A. The worst-case mean occurs at high temperature with a value of 138.581 nA. This translates to 0.00059 LSBs, or 0.075 LSBs at full-scale. The worst-case standard deviation (and the only one that fails the matching specification) occurs at low temperature with a value of 1.82 µA. This is a 3σ value of 5.46 µA which, using (8), is an

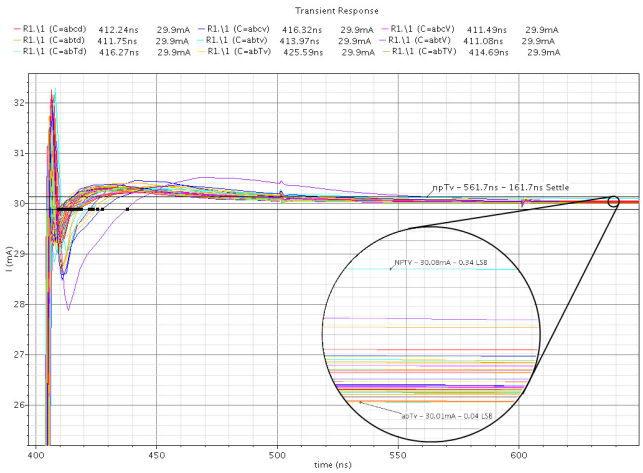

Figure 9. Settling Time and Full-Scale Current over Corners

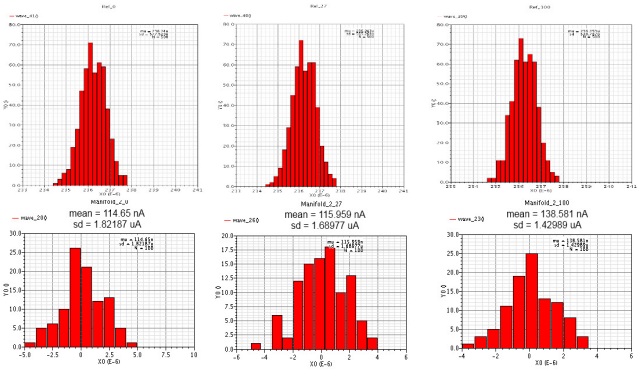

Figure 10. Monte Carlo Simulation for Mirror Matching and Reference Current

variance of 0.26 LSBs, which differs by 4.2% from the specification. Noted earlier, the size of these current mirrors directly influences the deviation in matching. The mirrors also exhibit better matching when the source transistor (M3 in Figure 4) size is increased. The spec could have easily been met if said transistor was increased in size, however there was concern that the DAC would not fit within the 1.3mm x 1.3mm frame given so the sizes were left unmodified.

$$\sigma = \sqrt{\sum_i \sigma_i^2} \therefore \sigma_i = \frac{\sigma}{\sqrt{127}} \therefore \frac{3\sigma_i\sqrt{127}}{LSB} = \pm LSB_\sigma \qquad (8)$$

Figure 11 shows the static power consumption for a code value of 1000000 (which is half of full-scale). The worst corner was at NPTV with a value of 226.2 mW. Over all corners, the power consumption varied by roughly 60 mW with 170 mW being the best-case power consumption. As mentioned previously, the localized current switch voltage biases consume 20 µA and, at 127 total cells, this attributes to 12 mW that could easily be minimized. Overall, the power-consumption exhibited by this design fell well within the required spec so optimizing the consumption wasn't really of much concern.

Figure 12 shows the output staircase for a ramping input code, while Figures 13 and 14 show the worst-case INL and DNL, respectively (both occurred at the same corner of NPTV and were generated using MATLAB – no offset or gain correction was used).

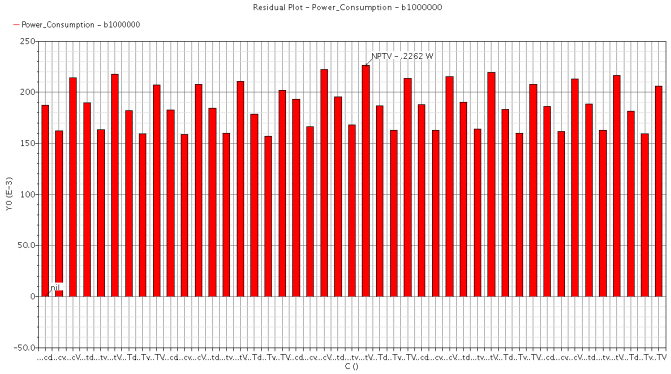

Figure 11. Static Power Consumption for input code 1000000

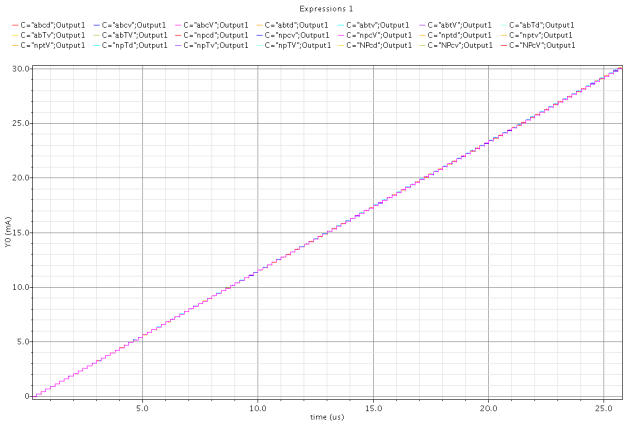

Figure 12. Output for Ramping Input Code over Corners

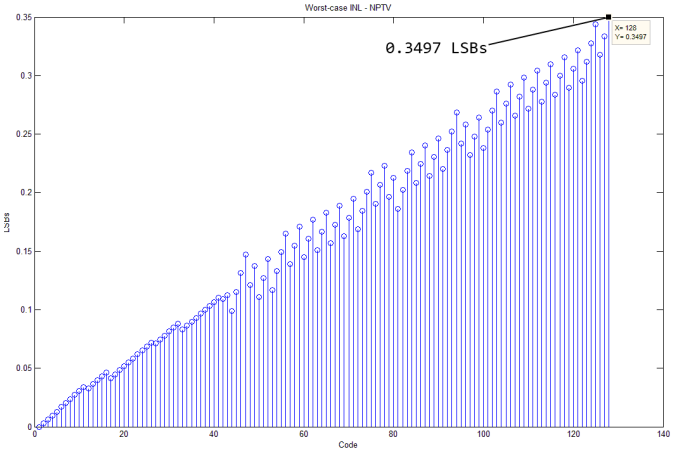

Figure 13. Worst-case INL at corner NPTV

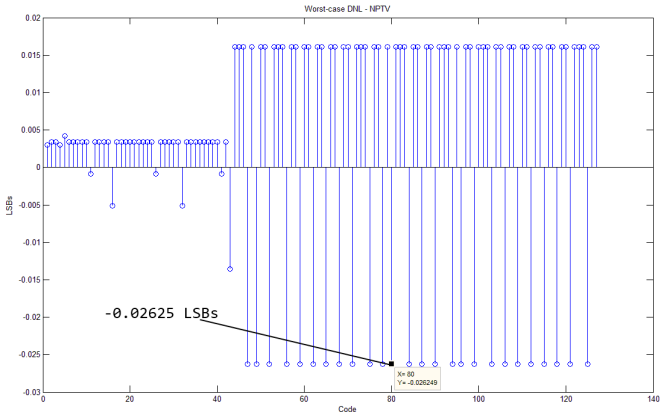

Figure 14. Worst-case DNL at corner NPTV

An interesting characteristic of the INL plot is that is in increasing at an almost constant rate. This indicates that each cell sinks slightly more current than the reference of 236.2 µA and could be fixed via gain correction or by tightening the current mirror matching distribution. A further interesting characteristic is the oscillation between positive and negative DNL after about a code value of 0101000 (decimal equivalent of 40). The reason for this oscillation is slightly confusing but it is perhaps due to variances in the mirror bias voltages due to switching cascading back into the manifold and then coupling back into the LSB cells again.

IV. Simulation And Test Circuits

A) Opamp Test Bench

Figure 15 shows the test bench used to measure the gain and phase margin (Figure 16 and Figure 17, respectively) of the opamp. The lowest gain occurred at corner NPTv with a value of 74.65 dB (which, when coupled with a potential 50% voltage gain decrease due to the gm-bias network resistor variance, results in a minimum possible gain of 68.6 dB; still 5 dB higher than required). The worst-case phase margin was found to be 60° at corner npTd.

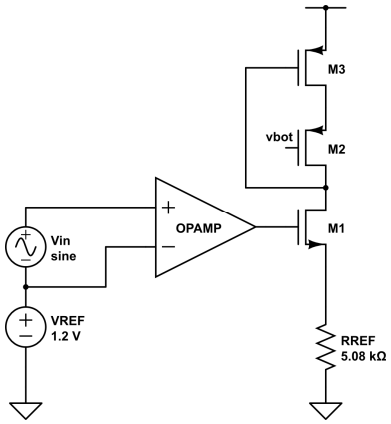

Figure 15. Opamp AC-Analysis Test Bench

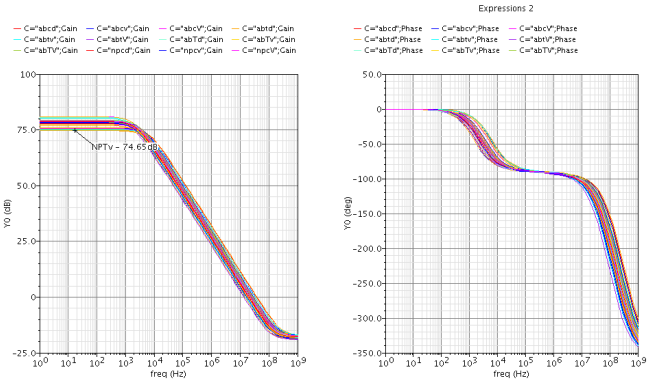

Figure 16. Gain and Phase for Opamp with Worst-Case Gain of 74.65 dB at Corner NPTv

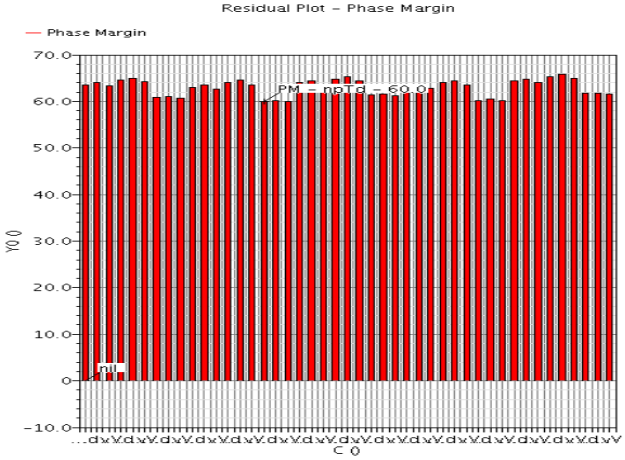

Figure 17. Phase Margin Plot for Opamp with Worst-Case PM of 60° at Corner npTd

As the simulations show, the gain is high enough to guarantee current reference accuracy and stable enough to theoretically not oscillate (since 60° is plenty of phase margin to guarantee amp stability). However, AC simulations can be misleading as they linearize the circuit before running. If the circuit is not, in fact, linear then the simulation is simply incorrect. Thus, to absolutely guarantee opamp stability, it is wise to run a transient simulation to verify this. In order to run this transient simulation, the architecture of Figure 2 was used in conjunction with a stepped V_{REF} with a 1 ns rise time. The current through R_{REF} was then measured over all corners which helped to guarantee both reference accuracy as well as stability. The outcome can be seen in Figure 18 which depicts an incredibly stable reference step-response with a very miniscule overshoot before immediately settling to the desired reference value. The current only varied between 236.2 µA and 236.3 µA (or a difference of 0.0004 LSBs) and exhibited no oscillation which verified that the opamp is, in fact, stable.

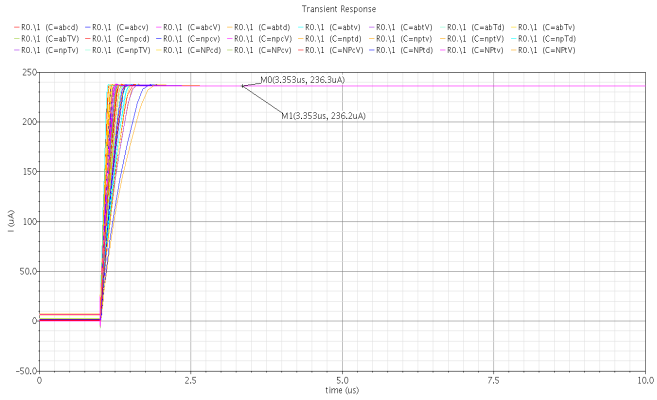

Figure 18. I_{REF} Step-Response Over Corners

C) Current Manifold and Current Cell Test Benches

In order to test the current manifold matching, the test bench in Figure 19 was constructed. The current at the I_{out} node was measured and plotted in reference to I_{REF}. The results have already been shown in Figure 10 and will not be replicated again here.

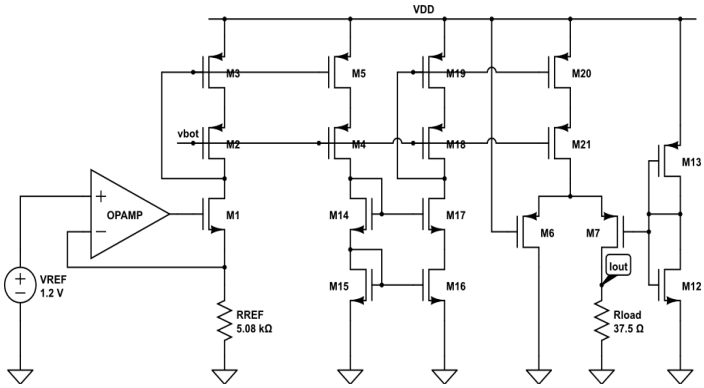

Figure 19. Current Manifold Matching Test Bench

The current cell output resistance bench is very similar to that of Figure 19 and is replicated below in Figure 20. Here the load resistor is replaced with a DC voltage source which is swept between 0 V and 1.125 V (which corresponds to the full-scale current of 30 mA). The current at that node was then saved and plotted against the swept DC voltage to obtain an I-V curve. From here, the derivative of said curve was then calculated within the Virtuoso ADE using Spectre, and the average of the inverse of that derivative was plotted over corners to result in Figure 20. As can be seen, there is quite a large variance in the output resistance but the absolute lowest value occurs at corner npTv with a value of 180 MΩ. Dividing this value by 127 yields the equivalent full-scale output resistance of 1.5 MΩ which is significantly larger than the load resistance.

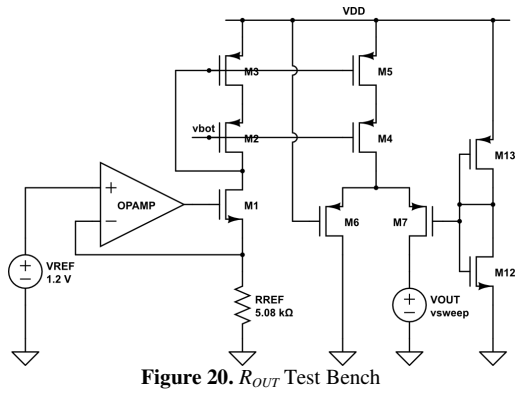

Figure 20. R_{OUT} Test Bench

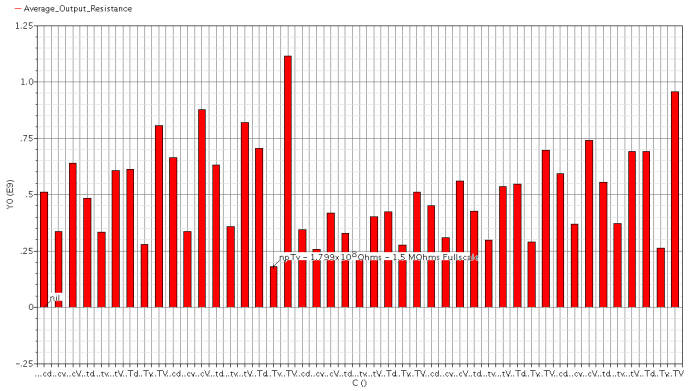

Figure 21. Average R_{OUT} over Corners

C) Full DAC Benches

In order to test the various DAC parameters, the test bench shown in Figure 22 was used. To generate the plots for static power consumption, the current out of the source V_{DD} was measured and multiplied by the actual value of V_{DD} (shown in Figure 11). To generate the plots for settling time, each bit value was turned on at the same time (to represent a code transition from 0000000 to 1111111) which resulted in the previously reported Figure 9. Finally, to generate the staircase shown in Figure 12, the input code was changed one bit at a time over a given time period.

Once this staircase waveform was generated, it was sampled every 100 ns and saved to a CSV file. This file was then imported into MATLAB where a custom script extracted the

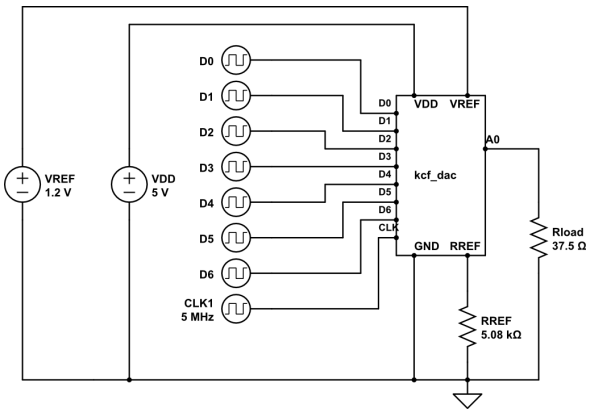

Figure 22. DAC Test Bench

relevant data. This was done by using equations (9), (10), and (11) for INL, DNL and N_{ABS}, respectively. Once this data was calculated, it was plotted for each corner. The corner plots are shown for INL (Figure 23), DNL (Figure 24), and N_{ABS} (Figure 25). The worst-case values for each have already been presented.

$$INL_i = \left(\frac{I_{out,i}}{I_{LSB}} - i\right)\Big|_{i=0}^{i=127} \tag{9}$$

$$DNL_i = (INL_i - INL_{i-1})\Big|_{i=1}^{i=127} \tag{10}$$

$$N_{ABS} = \frac{ln\left(\frac{V_{REF}}{V_{offset,max}}\right)}{ln(2)} \tag{11}$$

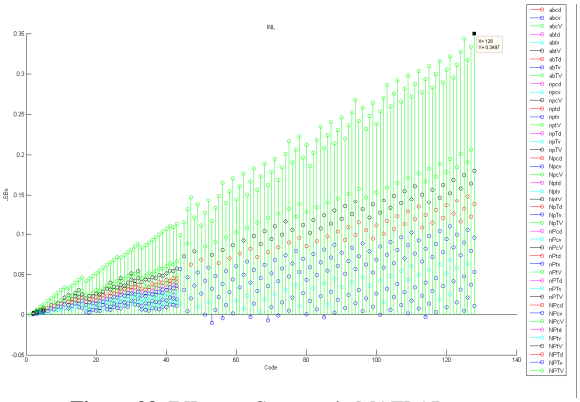

Figure 23. INL over Corners via MATLAB

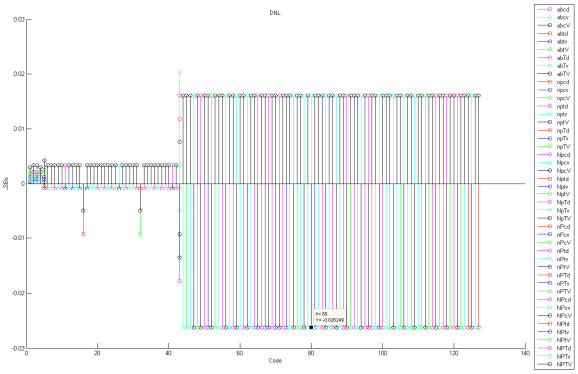

Figure 24. DNL over Corners via MATLAB

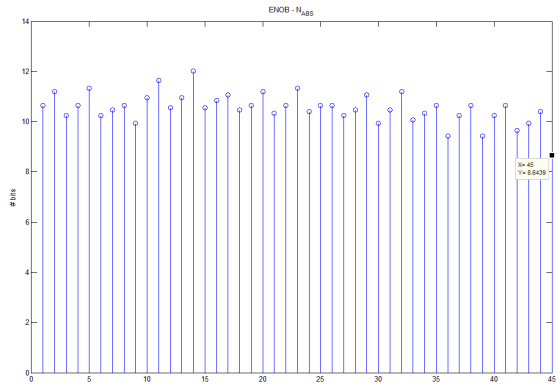

Figure 25. ENOB (Absolute) over Corners via MATLAB

V. Physical Layout

This section will serve as an overview for the layout of all the major block including problems encountered and areas for improvement.

A) Opamp

Figure 26 shows the layout for the opamp. Care was taken to layout the differential pair in a common-centroid fashion to help eliminate offset and mismatch errors after production (this pair can be see towards the bottom center of the Figure). The capacitor, shown as a series of 20 cells on the right, was split to try and minimize the effect of oxide-thickness gradients. This particular layout didn't pose any unforeseen challenges and the final product turned out to be quite compact and should operate well within the DAC.

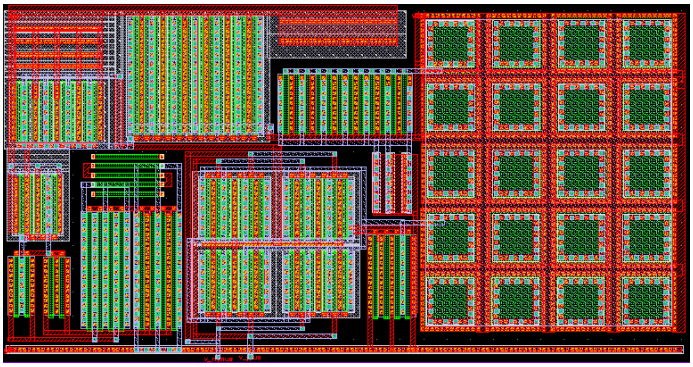

Figure 26. Opamp Layout

B) Decode Logic

Unfortunately, the layout of the decode circuitry did not progress as smoothly as the opamp. Figure 27 shows the layout for the row-decoder which Figure 28 shows the column decoder.

The strategy for the row decode was to have all inputs arrive on the left and be broken out through each row of logic cells. The output could then be taken and fed out to the top for easy access when laying out the DAC. While there is still some unused space, overall the layout is decently compact and the strategy worked well to help streamline the DAC layout itself.

However, moving to the column decode, it's clear that this same strategy doesn't work well with the larger number of devices. There is an unacceptable amount of unused space and some of the metal lines are unnecessarily long. This particular layout desperately needs to be optimized, but due to time constraints was left as-is. It consumes roughly twice the area of the row decoder on the full-chip and, as mentioned, should be optimized in the future to take up less area.

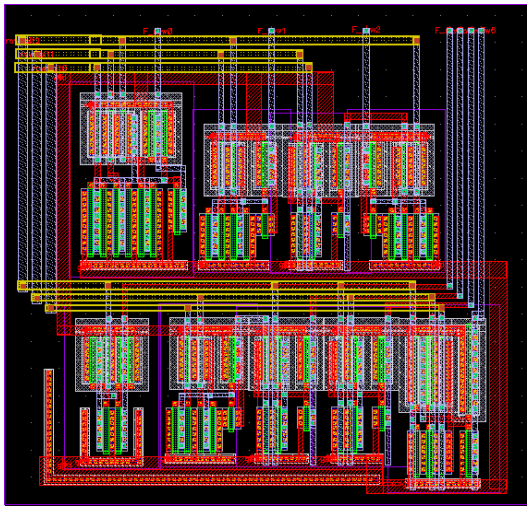

Figure 27. Row Decoder Layout

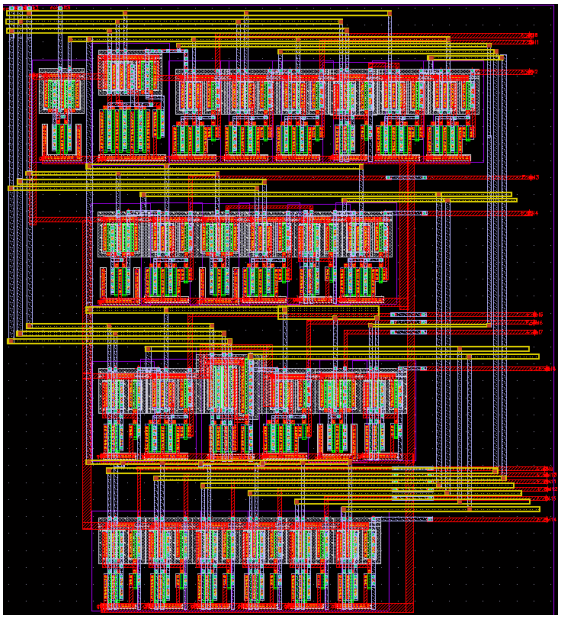

Figure 28. Column Decoder Layout

C) Current Manifold and LSB Current Cell

Figure 29 shows the layout of the current manifold. The strategy was to keep the cell as square as possible and to make sure all input and output were easily accessible to help in streamlining the full DAC layout. The design has some unused space, but is still fairly compact. The idea in keeping it square was to help minimize mismatch due to the distance of one cell to another. However, since the voltages being

generated are still being sent nearly the full way across the chip, minimizing the spacing in the manifold just skirts the actual mismatch issue (that is, cascading out to cells on the otherside of the chip). Retrospectively, the manifold should have been laid out in a more rectangular fashion so each mirror could line up with the row it needs to drive rather than needing to loop metal layers around to reach the correct row.

Figure 30 shows the layout of the LSB current cell. Initially, the cell was laid out to be square to help minimize total area consumption of all 127 cells. However, because the core is 8 cells by 16, it is inherently rectangular which resulted in all 127 cells not being able to fit. As such, the cell was laid out more rectangular, as shown, so that when they are placed in the DAC it can fit in both dimensions.

A big concern for the current cell was that the analog out lines should be noise free which means it should be isolated from the noisy digital logic. This was managed by simply defining clear digital and analog regions on the layout and then surrounding the analog section in a strip of grounded metal to act as shielding. The inputs, however, were not laid out in the most optimal fashion resulting in the need for digital lines to cross over portions of the analog out line when the full core was laid out. This, obviously, is quite non-ideal; however, again, due to time-constraints, not much could be done.

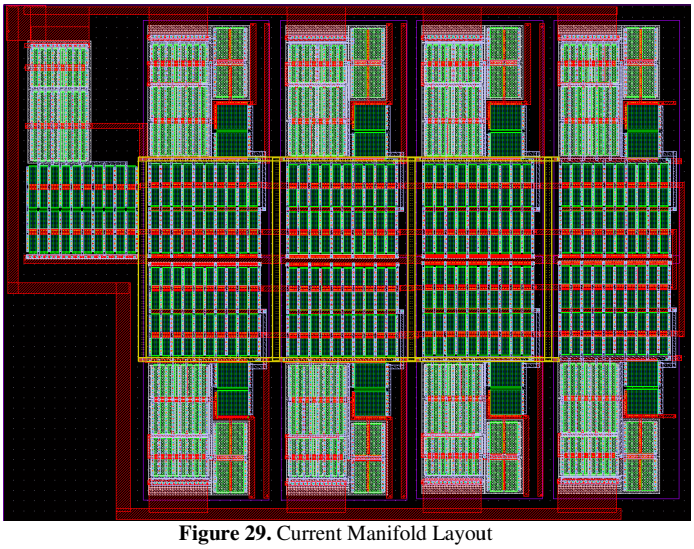

Figure 29. Current Manifold Layout

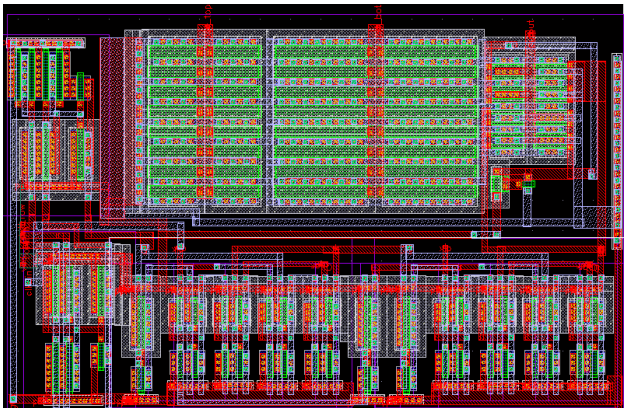

Figure 30. LSB Current Cell Layout

D) Full DAC

Figure 31 shows the layout for the full DAC. As mentioned part C, because of the oversight on input and output port placements on the LSB cells, digital lines needed to be physically routed over analog components, which is a very bad idea. Due to the fact that not much could be done to alleviate that problem, care was taken to try and route ground-lines as close to any digital-analog crossover points in order to try and shield the sensitive analog lines as much as possible. The actual analog output lines, shown on the top-left of Figure 31, are completely surrounded in ground lines on all metal layers in order to achieve isolation (functionally, it is very similar to a coaxial cable). Both V_{DD} and *GND* were routed along the edge of the chip to serve both as shielding and to allow easy rail taps from any point within the layout.

On the right hand side of the figure are the V_{REF} and R_{REF} input lines. Both the Current Manifold and Opamp were placed as close as possible to these lines in order to minimize any potential noise injection due to a long line of metal. Likewise, towards the bottom-right, the decoding logic was placed as close to the digital inputs as possible to help eliminate any skew due to parasitic resistance and capacitance within the metal lines. Any area that did not have any components or metal lines was filled with grounded metal planes in order to both meet the 30% metal utilization design rule and to help introduce as much ground shielding as possible.

On the following page, Figure 32 shows the pin-out of the 40-pin DIP that this layout will be placed in. Any pin that has no label receives no connection (there are only 16 pins on this DAC design and they are selected on an area of the package that minimizes bonding wire length and angle).

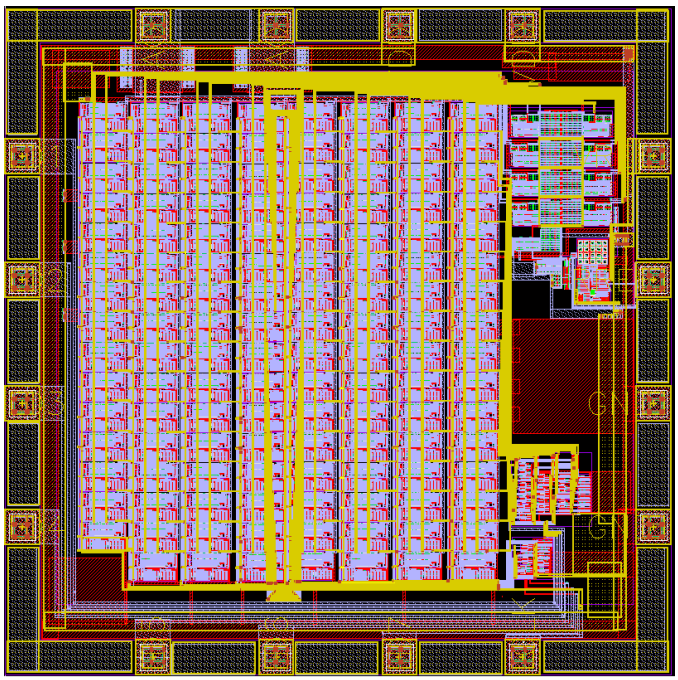

Figure 31. Full DAC Layout with Pad Ring

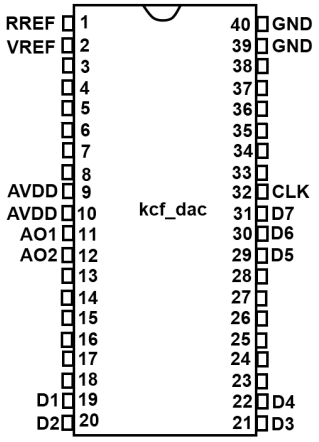

Figure 32. DAC Package Pin-out

VI. Conclusion

This paper presented a design for a 7-bit 5MSPS Current Steering DAC including full physical layout. Various performance metrics were analyzed and compared against specification. Potential problems resulting from the layout were presented including potential issues due to digital-analog cross-talk as a result of poor layout planning on a critical cell. In terms of simulation, the device performed very well against specifications but actual circuit performance will have to wait until after fabrication when more non-idealities will be introduced into the circuit.

VII. References

[1] G.I. Radulov, P.J. Quinn, P.C.W. van Beek, J.A. Hegt, A.H.M. van Roermund, "A Binary-To-Thermometer Decoder with built-in redundancy for improved DAC yield", *IEEE International Symposium on Circuits and Systems*, 2006.

[2] Takahiro Miki, Yasuyuki Nakamura, Masao Nakaya, Sotoju Asai, Yoichi Akasaka, Yasutaka Horiba, "An 80-MHz 8-bit CMOS D/A Converter", *IEEE Journal of Solid-State Circuits*, vol. sc-21, no. 6, December 1986.