A 100-MHz, Sub-100μW Temperature-Compensated RC Frequency Reference

Liam Crandall*, Nestor Gutierrez†, Prithvi Ulaganathan‡, Dhruv Dilbaghi§

*liamsc2@illinois.edu, †nguti4@illinois.edu, ‡pu9@illinois.edu, §dhruvd4@illinois.edu

***Abstract*—Clock generation is a necessity in practically every modern electronic device. Most modern devices utilize crystal oscillators, which are integrated off-chip. It is therefore necessary to explore alternative clock generation methods. Alternative methods such as LC time constant-based references have the issue of large power consumption and large inductor area. A promising approach is an RC-time-constant-based design, which occupies less area, consumes less power, and is much easier to integrate. The main issue of this approach, however, is the effect of temperature variations on its performance. Utilizing resistor temperature compensation techniques, we hope to minimize these drawbacks.**

I. Project Summary

The goal of this project was to complete the design, verification, layout, and tapeout of a 100 MHz temperature-compensated RC oscillator. This problem statement has been investigated and resulted in tape outs by Professor Hanumolu's group several times before, which allowed us to have unique access to a pool of previous work and experience which informed our own attempt. Previous work achieved 8.4 ppm/°C deviation of the clock signal, which we aim to meet and, ideally, surpass through additional circuit-level optimization and post-fabrication compensation. Post-layout extraction simulations of the individually extracted analog blocks resulted in a 6.4 ppm/°C deviation.

A. Target Specifications

- Frequency: $\sim$100 MHz
- Power: $<$ 1 μW/MHz
- TC: $<$ 10 ppm/°C
- Area: $<$ 1 mm^2

B. High-Level Architecture + Functionality

As stated previously, the function of this chip is to provide a stable frequency reference over a range of temperatures. This is achieved through a feedback loop that contains the following components:

- VCRO (Voltage-Controlled Ring Oscillator): Generates a frequency controlled by voltage VC
- FLL (Frequency-Locked Loop): Locks VCRO frequency to a reference by integrating the error between VREF and VF
- DCU (Digital Compensation Unit): Adjusts the effective resistance RREF using pulse-density modulation

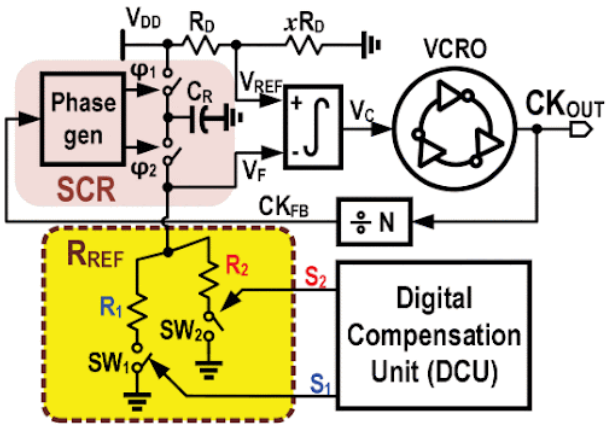

Fig. 1. High-Level Architecture Block Diagram

II. Block-Level Design Choices + Justifications

A. Full Chip Physical Design

In this design, 3 pins were used for power connections, 3 pins were used for ground connections, and 16 pins were used as data inputs which could feed either register of the DCU or the tuning switches of the VCRO (with each destination having its own enable pin). 1 pin was used for the 1μA current reference, 1 pin was used as a reset for the frequency dividers, and the last pin was used as the output frequency pin. Figure 2 will show the final layout, along with the location of each pin.

The following list describes the post-layout simulation performance. Note that area metric does not include the power supply smoothing capacitors.

- Frequency: $\sim$108 MHz
- Power: $\sim$2.5 μW/MHz
- TC: $\sim$6.4 ppm/°C
- Area: $\sim$.15 mm^2

B. Integrator

This block's design is primarily governed by achieving a large DC gain ($\sim$100 dB), supporting in achieving a loop gain sufficient to suppress VCRO temperature dependence. In addition, we found it important for its first pole to be very low ($<$ 100mHz) such that it can be modelled as a DC pole as in the case of an ideal integrator and allowing our closed loop system bandwidth ($\sim$1 kHz) to be dominantly tunable by

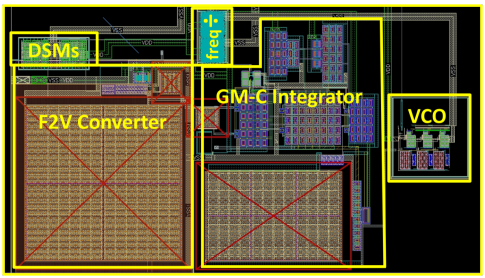

Fig. 2. Layout of the Active Area

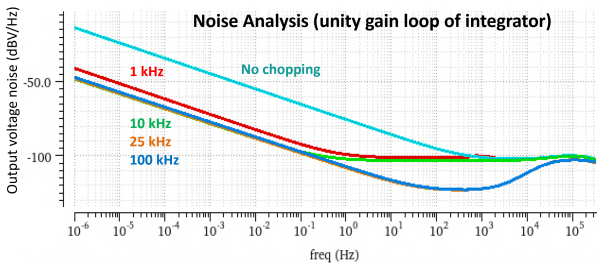

Fig. 3. Effect of Chopping

the block gains selected. Finally, the second pole of the block should be far outside of that system bandwidth so as to not degrade stability.

With these specs in mind, we opted to employ a GM-C integrating amplifier implemented through 2 stages, with the first being a folded telescopic cascode to allow for low input mode common mode voltage while delivering a high output impedance to meet the constraint of the first pole. The integrating capacitor also helps in realizing this pole location due to the Miller effect acting on the same node as the first stage's output. While this capacitor does introduce an RHP zero which could degrade stability, it can be effectively removed by introducing a series resistor of rough particular value in that same feedback path (the rough value makes its effect insensitive to expected process variation..

While achieving this performance with limited headroom (VDD = 1V) is nontrivial, we were able to meet specification by using thick oxide (2.5V) devices of long length (L = 5um) for lower overdrive and higher incremental resistance, respectively.

The resulting design achieves a DC gain of 116dB with a 3dB BW at 24 mHz in nominal conditions, both within the range required for loop performance. Even with Monte Carlo modeled mismatch, these metrics remain within spec at 3 σ, ensuring our design is resilient against process variation and device mismatch. We similarly tested the integrator over the process and temperature corners we might expect, observing worst case performance with slow-slow devices in the coldest expected environment. Even still, the 92 dB DC gain and 195 mHz 3dB BW achieved are still reasonable for operating in our application. This will ensure our loop is capable of meeting applicable performance even with undesirable process conditions.

After the core integrator design was completed, we moved onto implementing chopping switches for offset and noise reduction. Minimizing input-referred offset to the integrator is essential for temperature stability. Minimum sized devices were sufficient for the implementation of these switches, having almost no effect on the gain and dc operating points within the integrator. To generate the complimentary overlapping/non-overlapping clock signals to drive the switch gates, a basic SR latch is used, with additional capacitance added to create the necessary dead-time/delay. The effect of chopping at different frequencies on the output voltage noise of the integrator is shown in Figure 3, with 25 kHz being the final choice of chopping frequency in our design. The schematic of this completed block is included below.

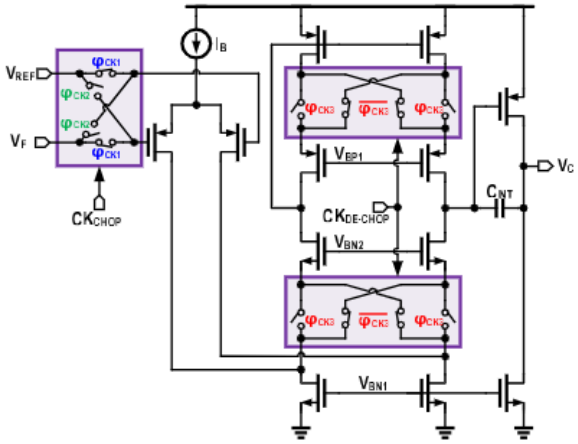

Fig. 4. Integrator Schematic

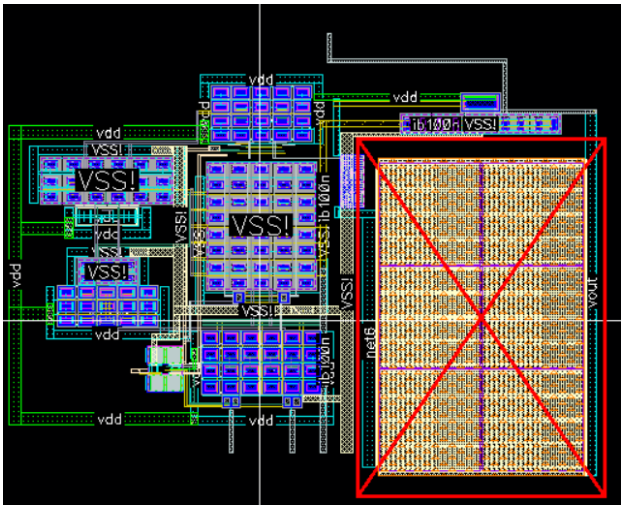

Fig. 5. Completed Integrator Layout

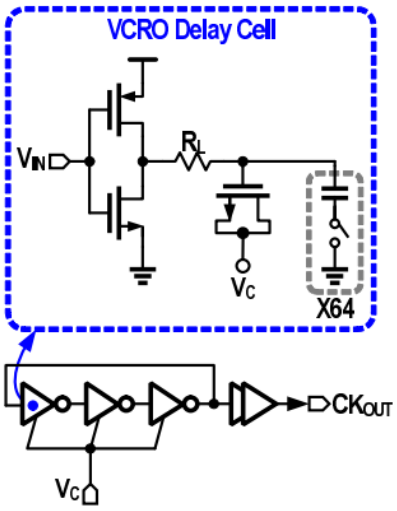

Fig. 6. VCO and Buffer Schematic

C. Voltage-Controlled Ring Oscillator

This block's design was primarily focused in achieving a $K_{VCO} = 86\frac{MHZ}{V}$ around a center point of $V_c = .5V$, where V_c represents the input to the VCRO from the integrator. This design choice was done to keep the unity gain frequency of the loop at the desired point and allow for sufficient fine tuning of the frequency through the varactor.

From Figure 6, it can be noted that the VCRO design consists of an inverter stage, a resistor R_L, and a capacitor bank in parallel with a varactor. To account for process variations post-tapeout, the capacitor bank was designed to be tuneable using 5-bit off-chip inputs. Worst-case corner simulations were done to ensure that there would be a binary code for the capacitor bank such that the oscillator can operate at 100 MHz at $V_c = .5V$.

R_L was implemented using two resistors in series: one P+ poly resistor and one N+ diffusion resistor. This was done to lower the temperature dependence of the VCRO and reduce dependence on the loop feedback as much as possible. The size of $R_L = 25k\Omega$ was chosen such that the on resistance R_{ON} of either transistor in the inverter cell is $< 1\%$ of R_L. $R_{ON,PMOS}$ was found to be approximately 220Ω.

The varactor was implemented using a low-threshold voltage N-type transistor in order to ensure a linear region in the V_c - Frequency curve is present near $V_c = .5V$. The proportion of stage capacitance coming from the varactor and the bank was tuned such that the "dynamic capacitance" of the varactor will result in the desired K_{VCO} of $86\frac{MHZ}{V}$. Figure 7 will demonstrate the Voltage-Frequency transfer curve of the VCO under nominal conditions at room temperature. An issue in preliminary design was that of the varactor being too large (i.e. high capacitance and low resistance per delay cell). A large varactor caused oscillation at the output of the integrator, which caused settling issues.

To ensure proper rail-to-rail oscillation for the digital components of the feedback loop, a two-inverter chain at the output of the VCRO was implemented. The first inverter was made the same size as a VCRO inverter stage and the second was made twice as wide. The extra gain from the second stage pushed it to rail. Figure 8 demonstrates the effect of the buffer on the VCRO's output.

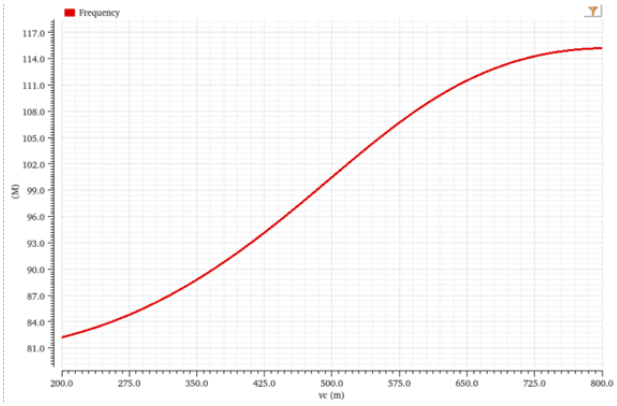

Fig. 7. V_C vs. Frequency Curve

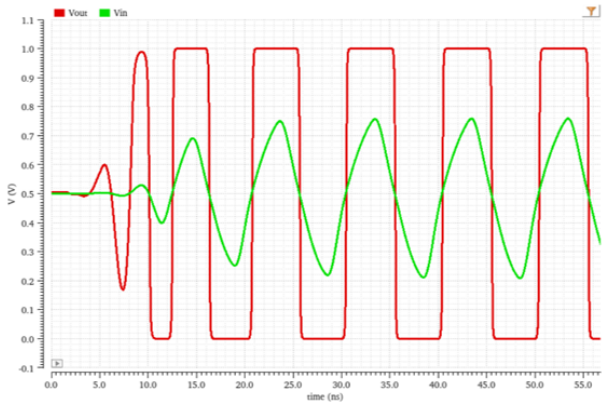

Fig. 8. VCRO output pre-buffer (green) and post-buffer (red)

D. Superbuffer

The output of the VCRO must be able to drive an off-chip capacitive load. This is because PCB traces on the test board may have capacitances on the order of multiple picofarads, (up to 10 pF for a long PCB trace). To do this, an additional buffering circuit called a superbuffer is used. A superbuffer is implemented as a cascade of N inverters, with each inverter in the chain being f times larger than the previous one. It is known that for a given value of N, choosing $f = (C_{out}/C_{g1})^{1/N}$ minimizes the propagation delay between the input and output of the superbuffer (C_{g1} is the gate capacitance of the first inverter and C_{out} is the load capacitance to be driven). Through schematic and layout design iterations, it was found that choosing $N = 6$ allows for a reasonable form factor in terms of the inverter sizing and the chain length. Assuming C_{out} to be 10 pF and plugging in values for the gate capacitance and the number of stages,

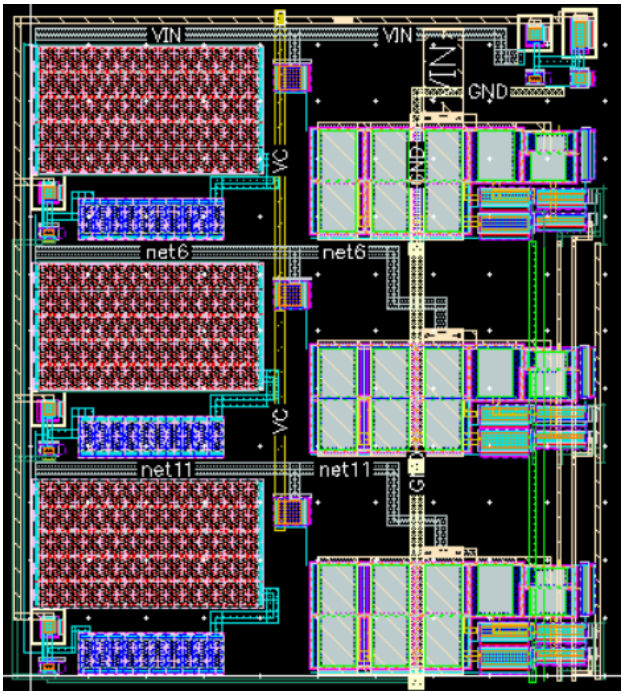

Fig. 9. Completed VCRO Layout

the nearest integer value of f turns out to be 4. Figure 10 shows the post-extraction simulation results of the superbuffer properly driving a 10 pF load with the VCRO oscillation given as the input.

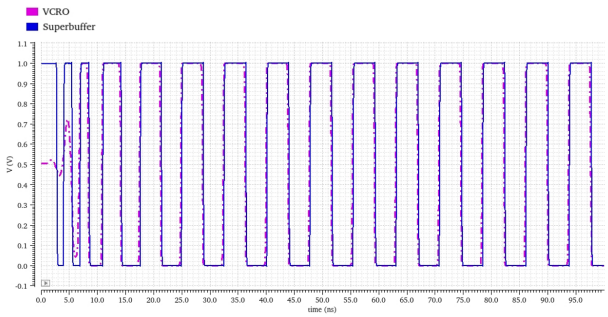

Fig. 10. VCRO input (pink) and superbuffer output (blue)

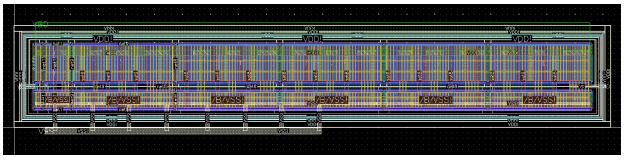

Fig. 11. Superbuffer layout

E. Frequency-to-Voltage Divider

Frequency-to-voltage (F2V) conversion in the main loop of the FLL is achieved through a voltage divider consisting of an upper switched capacitor resistor (SCR) and a lower switched resistor (R_{ref}). The resulting voltage is subtracted from the voltage formed by a reference divider (R_D and xR_D) through the integrator. In addition to the ratio of the reference divider and the division factor of the frequency divider in the feedback path of the FLL, the transfer function of the F2V also critically sets the operating output frequency and has the greatest variation across temperature among these three factors.

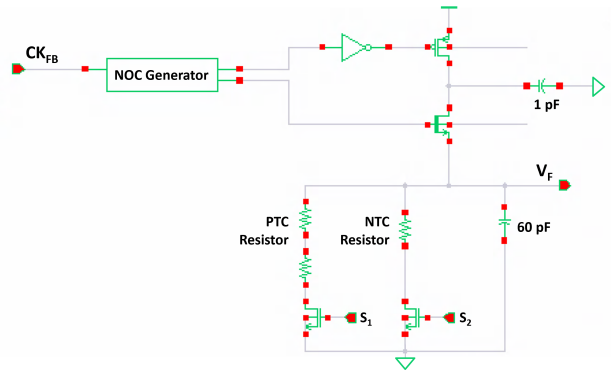

Fig. 12. F2V Schematic

Temperature compensation is achieved through digital trimming of the switched lower resistor. Two nominally 100 kΩ resistors with opposite first-order temperature coefficients (a negative-TC p-poly resistor and a composite positive-TC n-poly resistor) are switched in parallel using pulse density modulated bit streams that variedly multiply both resistances to flatten the combined resistance curve around room temperature. A two-point trimming process close to room temperature allows the optimal modulation indices to be determined and programmed into the chip as 16-bit inputs to delta-sigma modulators, compensating not only for the resistors' temperature variations but also for mismatches in other circuit elements.

The NMOS switches used to modulate each resistance branch are implemented as high threshold voltage devices with on resistances below 1% of a single branch, reducing their contributions to temperature variation without inducing significant off-state leakage. A challenge with these switches was selecting the right device type to reduce overshoot and undershoot during switching, which contribute significantly to error when switching around 100 MHz (which is needed to minimize delta-sigma modulator-induced VCO jitter). The high threshold voltage devices were found to be optimal when driven with fast rise/fall times to minimize the over/undershoot time.

The SCR is implemented using a high threshold voltage PMOS and a thick oxide NMOS around a 1 pF sampling capacitor. An additional 60 pF filter capacitor in parallel with the switched resistor is added with a considerably larger capacitance than the sampling capacitor to minimize charge-sharing losses and reduce ripple at the output to approximately 10 mV. The switches are sized wide enough to guarantee settling of charge transfer within one nominal clock cycle at the worst process/temperature corner without being too wide

as to induce significant temperature-dependent leakage. Selecting the optimal device type was a significant optimization challenge due to the differing sizing limitations of each one; ultimately, the PMOS did not require as thick of an oxide to reduce leakage due to its already lower mobility, whereas the NMOS needed to have the thickest oxide to prevent significant leakage.

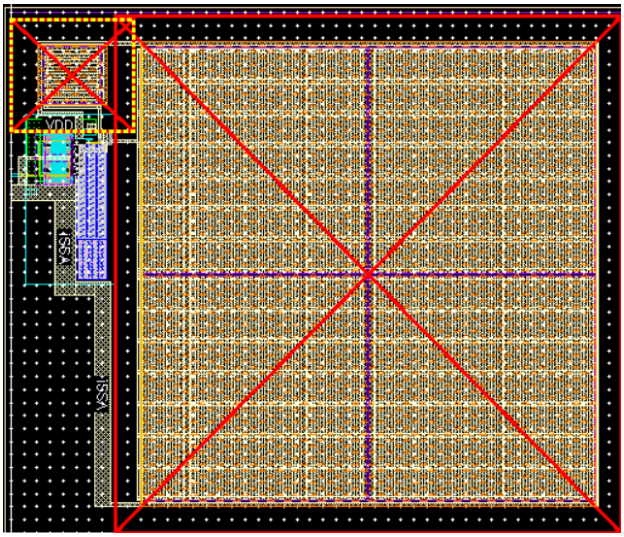

Fig. 13. Completed F2V Layout

F. Digital Components

Certain parts of this design were implemented using digital synthesis: frequency dividers for the feedback loop, delta-sigma modulators, and the data multiplexer/register. This subsection will explore the purpose, implementation, layout, and design justifications for each of these digital components.

1) Frequency Dividers: As shown in Figure 1, the NOC portion of the SCR is fed the output frequency. With 100 MHz being far too large for the SCR, the frequency dividers must be placed on this feedback path to ensure the voltage division ratio of the SCR-R_{ref} is kept as close to that of the V_{ref} path. Another use for the frequency dividers is for the chopping switches, which as seen in Figure 3 requires a frequency much smaller than 100 MHz. Although the division ratio for the SCR (32) and chopping blocks (128*32) are different, the implementation for these two dividers were largely the same.

These frequency dividers were implemented using a register-based approach, with counters counting up to the division ratio (either 32 or 128) and outputting a logic "high" signal upon reaching half the ratio to give a 50% duty cycle.

2) Delta-Sigma Modulators: In Figure 1, the 16-bit delta-sigma modulators produce the pulse density signals S_1 and S_2, which allow for the R_{ref} network to be largely temperature independent after conducting a two-point trim. The input being 16 bits allows for the pulse density to be tuned very precisely.

3) Data Multiplexer/Register: In order to minimize pin usage, a data multiplexer was implemented so that each of the three destinations (2 DSMs + VCRO Capacitor Bank) share the same 16 bits. Since the capacitor bank only uses 5 bits, the top 11 bits are not routed to its respective register. Each destination has its own clock signal, which acts as the "select".

A register was also added for each destination to eliminate the need to continuously provide the digital signal to these three destinations.

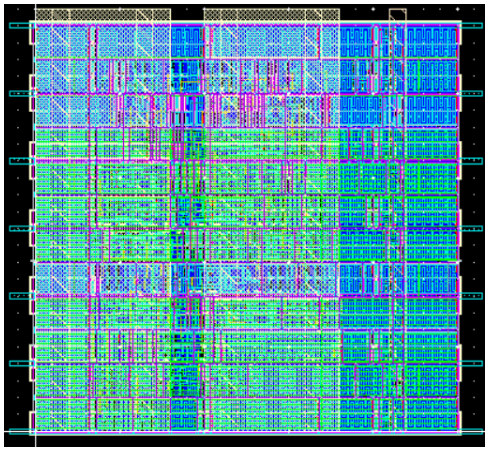

Fig. 14. Layout of the 128 divider, the 32 divider is visually identical

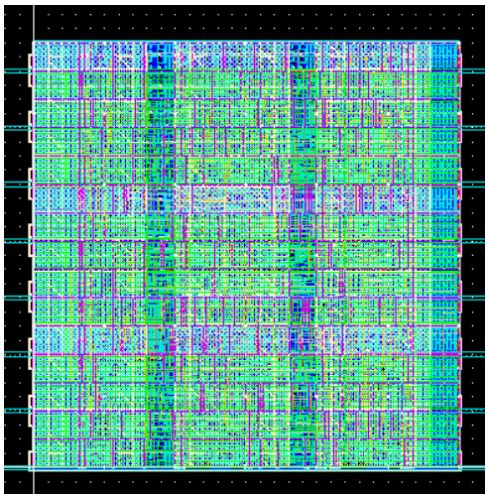

Fig. 15. Layout of the DSM

G. Closed-loop Simulations + Post-Fabrication Testing

Prior to beginning layout work, all blocks were individually tested and then validated in a closed-loop simulation. A 2-point trim at 10°C and 40°C was used to determine optimal modulation indices for the F2V switched resistor. The resulting steady-state output frequencies are plotted across temperature in Figure 18. Figure 19 demonstrates the change in settling time and steady-state value of the VCO control voltage as temperature is swept. Figure 12 demonstrates the change in performance of the whole system post-extraction. For our purposes, each block was extracted separately and then tested together.

These results demonstrate an 861 ppm variation in output frequency across the entire temperature range from -40°C to

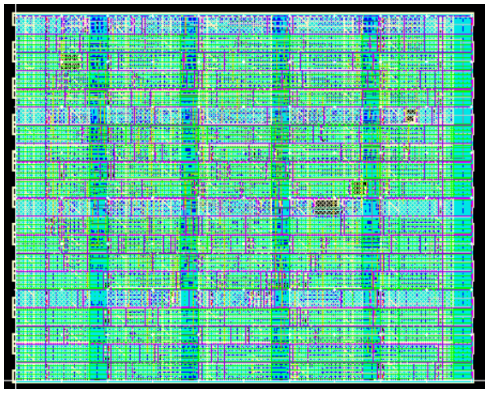

Fig. 16. Layout of the Data Multiplexer/Register

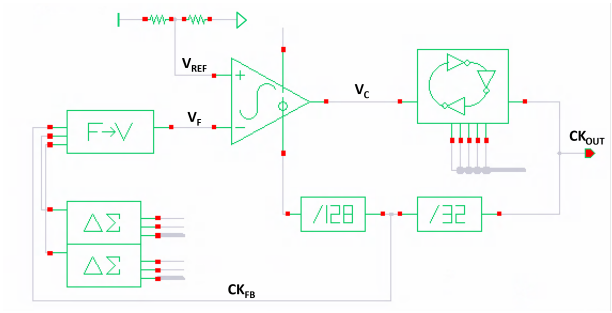

Fig. 17. Closed Loop Schematic

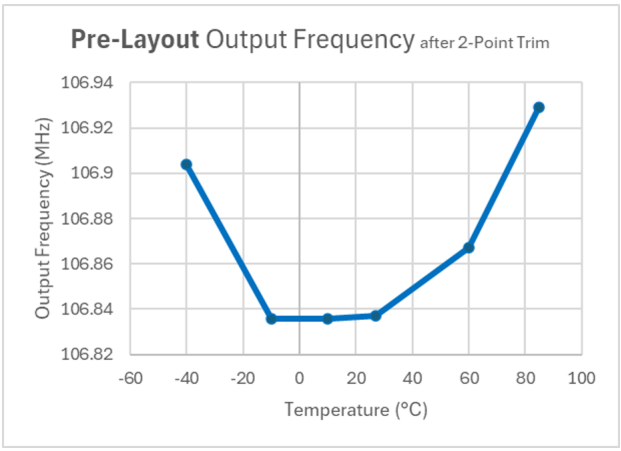

Fig. 18. Frequency Deviation across Temperature (Pre-Layout)

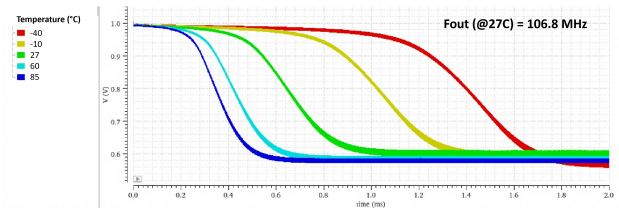

Fig. 19. Settling of VCO Control Voltage across Temperature

85°C, which is 6.9 ppm/°C with a center frequency of 106.8 MHz. Output frequency jitter was measured to be below 5 ps across the entire temperature range.

For post-fabrication testing, we intend on varying the input to each DSM in the TCU (i.e. a two-point trim) to determine the optimal pulse density for each part of the parallel switch resistor network. This process will be done with a temperature chamber to ensure the resulting optimal pulse density will largely cancel the temperature variation of the switch resistor network.

H. Variation from Initial Proposal

Ultimately, our final design was quite similar to the one described in the proposal. The first difference was the input reference current, which increased to 1μA instead of 100nA to ensure a typical bench power supply can reliably supply the needed reference. The second and final major difference was power consumption, which due to the increased current reference and other design choices described previously, made the original target of < 1 μW/MHz infeasible.

III. Post-Silicon Validation Plan

For testing the fabricated chips, we plan on using the general structure as shown in Figure 21, with a bench power supply providing the current reference, an off-chip LDO supplying the 1V V_{DD}, and the FPGA providing the data input to the DCU + capacitor bank and the corresponding select bits.

The initial steps will be conducting the two-point trim utilizing a temperature chamber to determine the optimal digital control word for the DCU followed by the tuning of the capacitor bank of the VCRO.

Currently, it is planned that Liam, Nestor, and Prithvi will work on the post-silicon validation during the Spring 2026 semester.

IV. Major Challenges + Lessons Learned

A. Challenges + Future Directions

Throughout the duration of this project, the major challenges we encountered came from the stability and coupling issues that arose when integrating the different blocks of the circuit together. Although many of these issues could not be completely eliminated, they were improved by employing layout techniques such as guard rings and isolation of high-frequency blocks.

We were able to largely follow our initial timeline from the proposal, with the digital and analog circuit implementation taking much less time than anticipated.

Although the overall design performs well, some noteworthy improvements could be the inclusion of support blocks such as on-chip LDOs and reference voltage/current references.

B. Lessons Learned

The major takeaways from this project were exposure to the complete analog design flow, which can be useful for developing the intuition needed to become competent designers. This project being done as a group setting has also helped us communicate more effectively in a team environment.

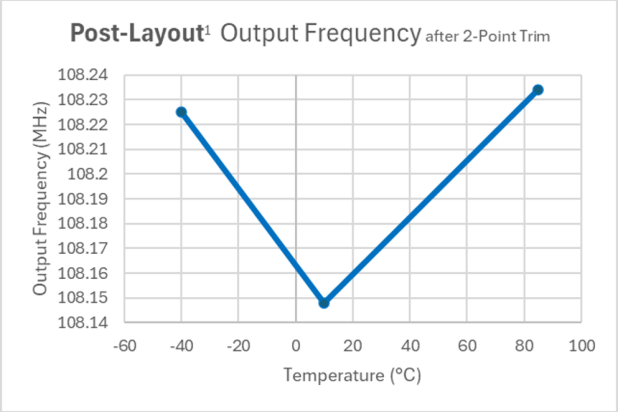

Fig. 20. Frequency Deviation across Temperature (Post-Layout)

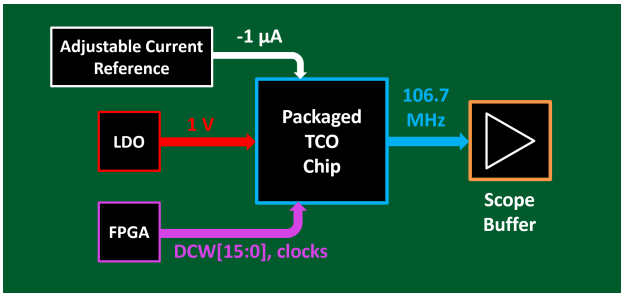

Fig. 21. Illustration of Completed PCBA

V. Group Ownership

Liam

- Primary Ownership: Systematic Model, Integrator Schematic Design
- Partial Ownership: Integrator Layout (w/ Dhruv)

Dhruv

- Primary Ownership: Chopping Circuits, Biasing, Analog Layout, Superbuffer Design, General Design Consulting

Nestor

- Primary Ownership: Voltage-Controlled Oscillator Schematic Design + Layout, Frequency Divider Design + Layout

Prithvi

- Primary Ownership: Frequency-to-voltage divider, Non-overlapping clock generator, Loop simulation

VI. Testing Instructions

Design files are stored under the directory `/groups/ece427-group6/rc-oscillator/post`, where Cadence can be launched by running the command `source launch_virtuoso`. Design files are in the `TCO_core_blocks` and `TCO_support_blocks` libraries, while testbench files are under the `TCO_core_TB` library.

To run the closed loop simulation, open the maestro file for `TCO_tb` under the `TCO_core_TB` library. A transient simulation should already be set up for 3 ms of runtime, which guarantees that full settling of the loop will occur for any temperature from -40°C to 85°C. Run this simulation and wait until the control voltage (`Vc`) has completely settled as depicted in Figure 19; output frequency can then be measured using the expression `frequency(clip(VT("/Fout") ta tb))`, with ta and tb replaced by the desired start and end times to evaluate the frequency.

TABLE I
Timeline for the Fall 2025 Semester

Week of	Planned Tasks
Sep 1	Design Review I with Professor Hanumolu
Sep 8	Finalize initial block designs
Sep 15	Loop simulations and block revision
Sep 22	Design Review II with Professor Hanumolu
Sep 29	Initial block layouts and extraction
Oct 6	
Oct 13	Final block layout, initial floorplanning, and chip top integration
Oct 20	
Oct 27	Finalize chip top and prepare for GDS submissions
Nov 3	
Nov 10	Trial GDS submission on 11/19
Nov 17	Final GDS submission on 11/26

VII. References

- A. Khashaba, J. Zhu, N. Pal, M. G. Ahmed and P. K. Hanumolu, "A 32-MHz 34-uW temperature-compensated RC oscillator using pulse density modulated resistors", IEEE J. Solid-State Circuits, vol. 57, no. 5, pp. 1470-1479, May 2022.
- K.-S. Park et al., "A 1-uW/MHz RC Oscillator With Three-Point Trimmed 2.1-ppm/°C and Single-Point Trimmed 8.7-ppm/°C Stability From 40 °C to 95 °C," in IEEE Journal of Solid-State Circuits, vol. 58, no. 7, pp. 2064-2074, July 2023.
- K.-S. Park et al., "A Temperature- and Aging-Compensated RC Oscillator With ±1030-ppm Inaccuracy From 40 °C to 85 °C After Accelerated Aging for 500 h at 125 °C," in IEEE Journal of Solid-State Circuits, vol. 58, no. 12, pp. 3459-3469, Dec. 2023.