Development and Characterization of NMOS Ring Oscillators Using a Four-Mask Process

Robin Teng, Deven Tseng, Tony Doan, Anish Gupta, and Asli Cebeci
Department of Electrical and Computer Engineering, University of California, Santa Barbara
(Dated: June 6, 2025)

This report details the fabrication steps and electrical characterization of 3-stage, 5-stage, and 7-stage NMOS ring oscillators with both resistor and NMOS loads. The fabrication sequence involved oxidation, photolithography, diffusion, etching, and metallization to create fully functional devices. We characterized the devices using electrical measurements, including I-V analysis, to extract key performance parameters. Our fabricated ring oscillators were functional, although measured frequencies deviated from simulation results which are likely due to errors introduced through imperfect fabrication.

I. INTRODUCTION

A ring oscillator as shown in Figure 1 is a fundamental integrated circuit structure composed of an odd number of inverter stages connected in a feedback loop. The odd count of inverters ensures that the signal continually toggles, as there is no stable logic state that satisfies the system. When the output of the final inverter is fed back to the input of the first, the logic level propagates through the chain and inverts at each stage, creating an unstable loop that results in a periodic oscillation. For example, a 3 stage oscillator with an input of 1 at the first inverter will propagate in the following form: $1 \rightarrow 0 \rightarrow 1 \rightarrow 0$ where this final 0 is connected directly back to 1 to create the aforementioned logic instability.

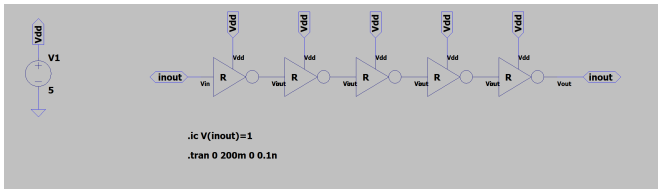

FIG. 1. 5-stage ring oscillator circuit diagram

Here we implemented a diverse set of ring oscillator configurations to evaluate the effects of stage count and delay element type on performance. The layout includes two 7-stage ring oscillators — one with discrete resistors (Figure 2) and one with pull-down NMOS transistors (Figure 3); four 5-stage ring oscillators—two resistor-based and two NMOS-based; and four 3-stage ring oscillators, again divided evenly between resistive loads and NMOS pull-downs. All RC components were kept identical across designs to ensure that differences in performance arise solely from the stage count and the type of pull-down element. This setup enables controlled comparison of how passive versus active delay mechanisms and inverter chain length affect oscillation frequency, power consumption, and waveform quality.

In this lab, the fabrication process incorporated thermal oxidation, dopant diffusion, photolithography, and aluminum metallization. These steps, along with the challenges encountered during each, are detailed in the methodology section of this report. A notable aspect of this process is its ability to produce multiple types of devices simultaneously on a single silicon wafer. Using the same four-mask process to fabricate n-type MOSFETs, we also produced MOS capacitors, Van Der Pauw structures and TLM (Transmission Line Method) test resistors. These additional structures are not only important components in their own right but also serve as reference devices for validation of both fabrication integrity and electrical performance. This report provides an in-depth study of the four-mask NMOS fabrication process used in the ECE120B laboratory. We begin by outlining the electrical testing methodology and expected results, followed by a detailed walkthrough of the fabrication steps, including the difficulties encountered.

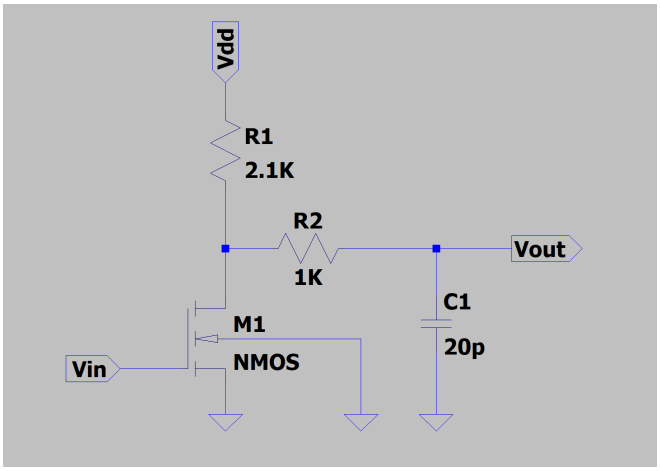

FIG. 2. Resistive load NMOS

II. METHODOLOGY

A. Sample Preparation

The fabrication process flow is similar to the one used for our individual transistor designs in ECE120A. The

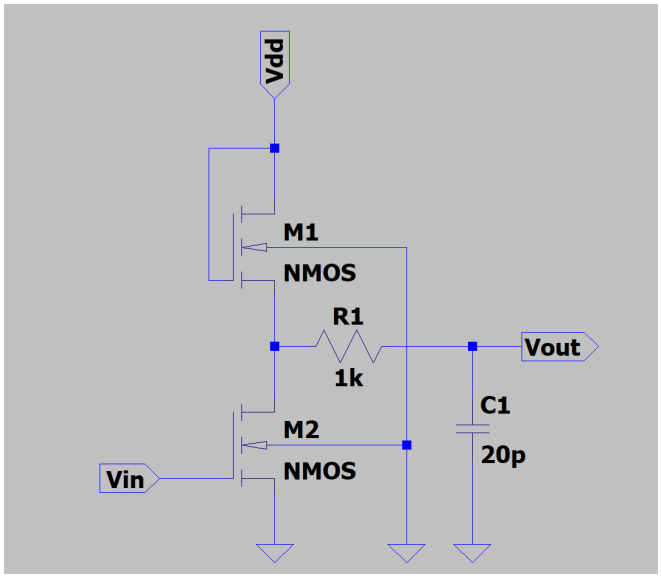

FIG. 3. Enhancement mode NMOS (pull-down NMOS transistor)

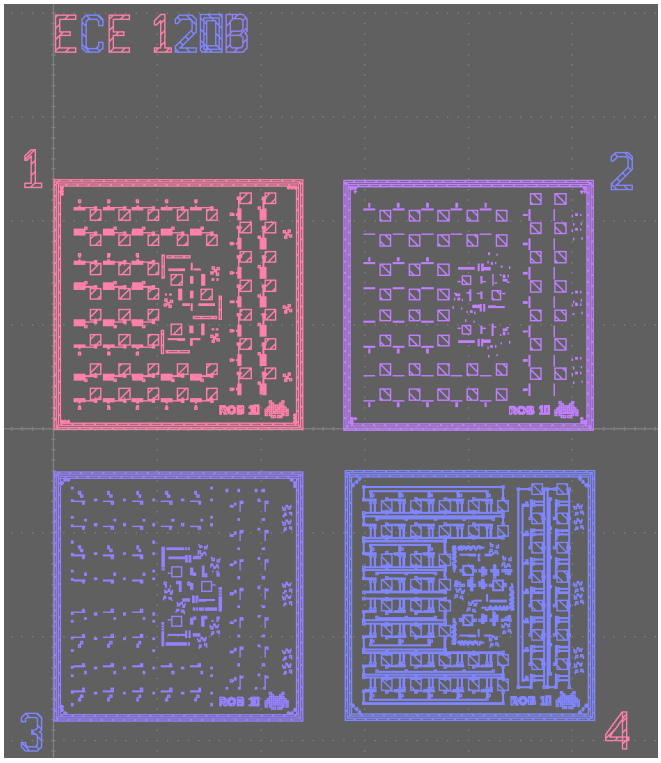

FIG. 4. Mask Design

key difference is the starting oxide thickness - we opted for 4000 Å of field oxide to ensure the oscillation frequency of our ring oscillators was observable within the limitations of electrical testing equipment. During the fabrication process, the first three steps—including wafer cleaving and the initial field oxide growth—were precompleted for us. Our task was to verify that these steps had been correctly performed. The first measurement we conducted was to confirm the quality and thickness of the initial oxide layer. We began with a rough estimate using the reflectance box, followed by a precise measurement using the Filmetrics SiO_2 thickness tool to verify that the oxide thickness was approximately 4000 angstrom. The reflectance box relies on the principle of thin-film interference, where the oxide layer's thickness determines the observed color. This visual estimate provides a reference thickness range. This estimate is then input into the Filmetrics system, which shines light onto the sample and analyzes the reflected signals. Some light reflects off the oxide surface while other portions reflect from underlying interfaces. The resulting interference pattern is captured and analyzed by a spectrometer, which calculates the oxide thickness based on the wavelength-dependent reflectance. This method allows for a non-destructive and accurate measurement of the oxide layer. Following this, we measured the thickness and sheet resistance of the original wafers, both with and without the silicon dioxide layer. These measurements were performed using a four-point probe system. The device applies current through the outer two probes and measures the voltage across the inner two probes. Using the known probe spacing and Ohm's Law, the system determines the sheet resistance of the sample.

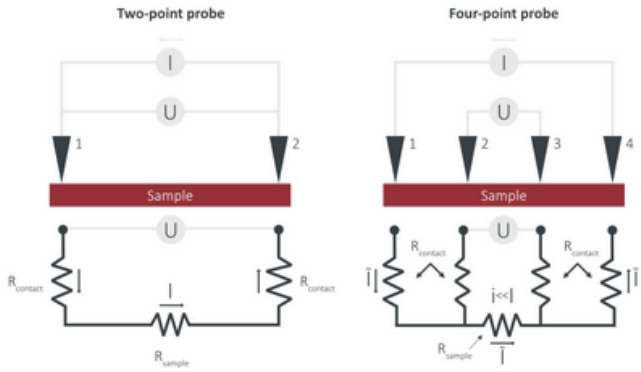

FIG. 5. Four point probe schematic

TABLE I. Initial Oxide Growth Verification

Sample #	Q1	Q2	Q3	Q4	Avg
Control					
1	4069.1	4078.7	4038.4	4079.8	4066.5
2	4088.9	4041.2	4097.5	4009.1	4059.1
3	4084.6	4037.2	4070.1	4084.2	4068.9
4	4063.2	4037.5	4062.4	4039.2	4050.5

Although the actual oxide layer ended up slightly thicker than intended, this deviation is expected to have minimal impact on the final performance of the fabricated FETs. Based on resistivity measurements, the high resistance observed in our sample confirms that the oxidation process was successful. Additionally, since we

were provided with p-type silicon, we can conclude that the majority carriers are holes. The acceptor concentration, N_a, is estimated to be approximately $10 \times 10^{15}\ \text{cm}^{-3}$

TABLE II. Oxide 4 Point Probe Measurements

Sample #	**Resistance (Ohms)**	**Resistivity (Ohms × Thickness)**
Control	88.4 Ω	0.0237 Ω·mm
1	≥ 9.99 kΩ	At least 4.90 Ω·mm
2	≥ 9.99 kΩ	At least 4.93 Ω·mm
3	≥ 9.99 kΩ	At least 4.91 Ω·mm
4	≥ 9.99 kΩ	At least 4.92 Ω·mm

B. Mask 1 Photolithography

The first photolithography step involves creating openings in the field oxide using Mask 1 to define regions for n-type dopant diffusion. The process begins with a cleaning sequence using acetone, 2-propanol, and deionized water, followed by a dehydration bake to eliminate moisture from the wafer surface. A thin layer of hexamethyldisilazane (HMDS) is vapor-deposited to promote photoresist adhesion, and AZ 4110 photoresist is then spin-coated onto the wafer. After a soft bake to partially cure the resist, the wafer is exposed to UV light through Mask 1 and developed using AZ 400K developer to reveal the intended pattern. The resulting photoresist pattern is inspected under a microscope to verify alignment and development quality, followed by a hard bake to improve resist adhesion and thermal stability. Additional microscope images are taken to confirm pattern integrity, and an oxygen plasma descum is performed to remove any residual photoresist from the exposed regions. This step is critical for accurately defining the diffusion windows for n-type dopants, which is essential for proper transistor function and structural integrity. Afterward, further microscope inspections were conducted, and a profilometer measurement was taken to determine the thickness of the photoresist. The collected data is shown below:

TABLE III. PR Height

Sample #	**Δ Height** (µm)
1	1.119
2	1.078
3	1.081
4	1.069

C. Oxide Etching for Diffusion

It's important to note that in the previous step, the photoresist was hard-baked and subjected to an oxygen plasma descum during the photolithography process for Mask 1. This preparation is essential for the subsequent etching step, as our etchant—hydrofluoric acid (HF)—selectively removes silicon oxide while having minimal effect on the photoresist. As a result, the photoresist serves as a protective mask, ensuring that only the exposed silicon oxide is etched away as shown in Figure 6.

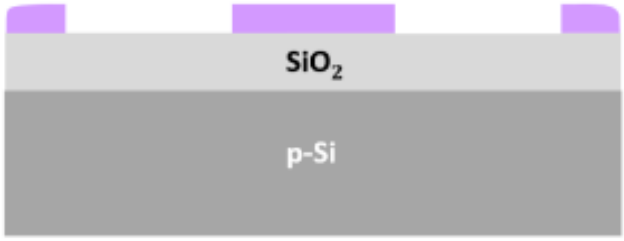

FIG. 6. Expose PR with source/drain diffusion mask (Mask 1) and develop

The next step involves etching away the silicon oxide by immersing the sample in hydrofluoric acid (HF). To accurately determine the etch rate, a control sample with a known oxide thickness is periodically dipped into the HF solution for short intervals. After each dip, the remaining oxide thickness is measured using the Filmetrics reflectometer. This process allows us to calculate the etch rate, which was determined to be approximately 100 $\frac{\text{nm}}{\text{min}}$. Based on our measurements, the slope of the data indicates an actual etch rate of 118.67 $\frac{\text{nm}}{\text{min}}$. To ensure complete oxide removal—especially important for creating electrically conductive contact regions—we intentionally over-etched the sample slightly. Etching could leave behind residual oxide, preventing proper electrical contact. Therefore, we added a 20% buffer to our calculated etch time of 4 minutes and 27 seconds, resulting in a final etch time of 5 minutes and 20 seconds. Profilometer data of this step is shown in the following table.

TABLE IV. PR and Oxide Height

Sample #	**Δ Height** (µm)
1	1.703
2	1.641
3	1.575
4	1.604

The photoresist was then removed by dissolving it in acetone, followed by a standard cleaning sequence using isopropyl alcohol, deionized water, and a dehydration bake. This process prepared the chips to match the appearance shown in Figure 7.

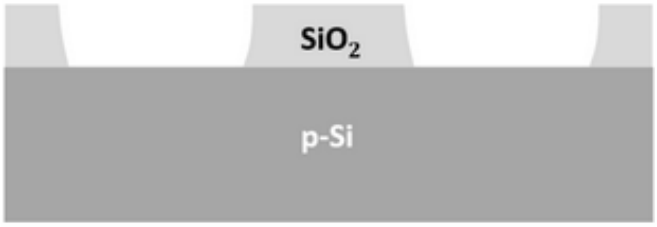

FIG. 7. PR Removal

Next, we captured microscope images and collected profilometer data, as shown below. It is important to compare these results with the initial oxide thickness measurements to confirm that the oxide was fully removed before proceeding to the next fabrication step.

TABLE V. Oxide Height

Sample #	**Δ Height** (µm)
1	0.531
2	0.552
3	0.479
4	0.482

D. Diffusion (n-type dopant)

With a precisely defined window opened into the silicon, the sample is now ready for doping. Before entering the diffusion furnace, the samples must be thoroughly cleaned—first in a RCA solution for 20 minutes in total to remove organic contaminants, followed by a 10-second dip in hydrofluoric acid (HF) to eliminate any remaining residues. The samples are then subjected to a high-temperature diffusion process at 950 °C for 20 minutes, allowing phosphorus atoms to diffuse into the exposed silicon regions.

Due to the photolithographically defined openings, only the exposed silicon areas are doped, resulting in well-defined, n-type doped wells, as illustrated in Figure 8. It's also important to note that this diffusion process leaves behind a layer of phosphorus-rich glass on the wafer surface, which must be removed with an additional 10-second HF dip.

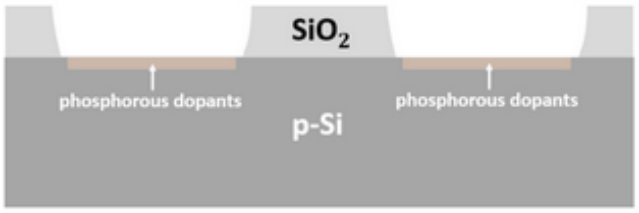

FIG. 8. HF etch to remove phosphorous glass

By measuring a test sample processed alongside the rest of the batch, we determined the inner and outer resistances, resistivity, and carrier concentration.

$$\rho = R \times \text{thickness} \tag{1}$$

$$N_d = \frac{1}{q\mu\rho} \tag{2}$$

These values were calculated assuming an oxide thickness of 525 µm and an electron mobility of 600 $\frac{\text{cm}^2}{\text{Vs}}$.

TABLE VI. Diffusion Data

Region	**R** (Ω)	ρ ($\frac{\Omega}{\text{cm}}$)	N_d (cm^{-3})
Inner	20.4	1.071	9.7×10^{15}
Outer	22.4	1.176	8.7×10^{15}

E. Drive-In Oxidation Growth

Drive-in diffusion is essential for achieving sufficient junction depth, which directly influences the electrical characteristics of the FET. This is accomplished by growing a new oxide layer over the doped wells, promoting deeper dopant penetration into the silicon substrate. To prepare the samples for furnace processing, a thorough cleaning cycle is performed—starting with acetone, isopropyl alcohol, and deionized water, followed by a 10-minute piranha acid dip to remove organic contaminants. A final 10-second hydrofluoric acid (HF) dip ensures the removal of any residual materials. Once cleaned, the samples are placed in the furnace for a dry-wet-dry oxidation process aimed at producing high-quality oxide layers. The target was to grow an oxide layer between 300 nm–350 nm. Using the BYU Oxidation Growth Calculator, we determined the required oxidation durations : 10 minutes of dry oxidation, 41 minutes of wet oxidation, and an additional 10 minutes of dry oxidation. The resulting oxide thickness measurements and a cross-sectional diagram of the fabricated structure are shown in Figure 9 and the table below:

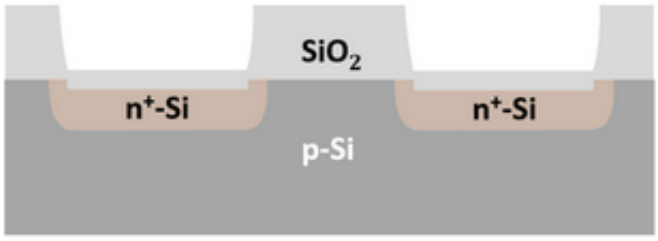

FIG. 9. Drive-in diffusion

TABLE VII. Oxide Growth

Sample #	**Oxide Growth** (Å)
Control	4000
1	6000
2	6000
3	6000
4	6000

F. Mask 2 Photolithography

The second photolithography step involves defining windows in the field oxide using Mask 2 to specify regions for thin gate oxide growth. This step follows the same pre-photolithography cleaning and preparation procedures as the first, up to the point of mask exposure. However, since the sample already contains patterned features from the previous step, precise alignment of Mask 2 is required. This is done using alignment markers such as crosshairs, corner squares, and wafer edges to ensure accurate overlay of the new pattern onto the existing features as seen in Figure 10.

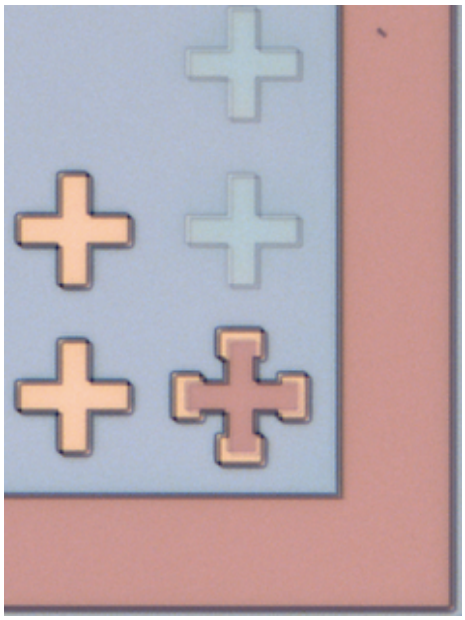

FIG. 10.

Although this alignment process takes a few additional minutes, once proper alignment is achieved, the standard photolithography procedure resumes. The sample is developed in AZ 400K developer, followed by a hard bake and an oxygen plasma descum to remove residual photoresist. Profilometry data was then collected to verify that the photoresist thickness was within the expected range. These measurements were taken over the border wells and are intended solely to confirm the depth of the patterned features. The collected data is presented below:

TABLE VIII. PR Height

Sample #	**Δ Height** (µm)
1	1.31
2	1.31
3	1.24
4	1.27

G. Oxide Etch for Gate Oxide Growth

As noted in the previous step, the photoresist was hard-baked and descummed during the photolithography of Mask 2. This preparation was necessary for the subsequent etching of the silicon oxide region left between the wells, which must be removed to allow for gate formation. Following the same procedure used earlier, we determined the etch rate for the day to be 107 $\frac{\text{nm}}{\text{min}}$. To etch a target depth of 500 nm with an added 20% buffer to ensure complete removal, we etched the sample for 5.5 min.

TABLE IX. PR and Oxide Height

Sample #	**Δ Height** (µm)
1	1.65
2	1.65
3	1.56
4	1.59

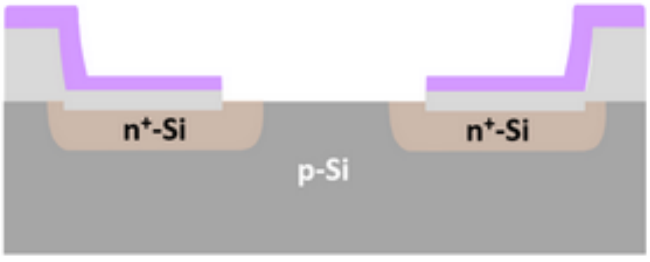

FIG. 11. HF etch to remove oxide from gate region

After etching, the photoresist was removed using an acetone cleaning cycle to prepare the sample for measurement. We then captured microscope images and collected profilometer data, as shown below. These measurements were crucial to confirm that the silicon surface had been fully exposed and no residual oxide remained, ensuring accurate control over the subsequent gate oxide thickness.

TABLE X. Etch Height

Sample #	**Δ Height** (µm)
1	0.694
2	0.699
3	0.687
4	0.690

H. Oxidation Growth of Gate Oxide

This step began with a standard solvent cleaning sequence followed by a dehydration bake, and then a 10-minute RCA clean to thoroughly remove any organic contaminants. Since the piranha solution leaves behind a thin oxide layer on the surface, which could interfere with controlled oxide growth in the furnace, we performed a brief 10-second buffered HF (BHF) dip to remove it. Following this, a new oxide layer was grown to electrically isolate the gate from the substrate. This step mirrored previous oxidation procedures but targeted a significantly thinner oxide thickness—specifically, 40 nm to 50 nm of dry oxide. This reduced thickness is critical for the gate region, as it enables more precise control during electrical characterization. A schematic representation of this oxide growth is shown in Figure 12.

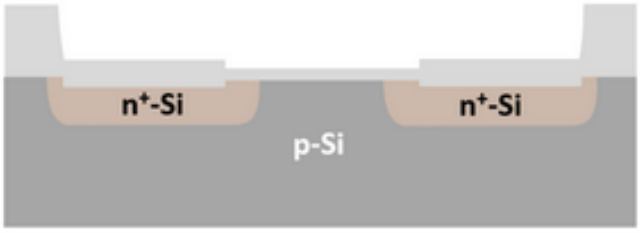

FIG. 12. Remove PR and perform gate oxidation

Using the BYU Oxide Growth Calculator, we determined that approximately 42 minutes of dry oxidation was required to achieve our target oxide thickness. After preheating the furnace and coordinating with the lab manager, we carefully placed our samples into the oxidation furnace and initiated the timed growth process. Upon completion, we used a control sample and the Filmetrics system to measure the resulting oxide thickness. The data indicated that approximately 400.2 angstrom (or 40.02 nm) of oxide had been successfully grown. The corresponding measurement data is presented below:

TABLE XI. Oxide Growth

Sample #	**Oxide Growth** (Å)
1	412
2	415

I. Mask 3 Photolithography

The third photolithography step involves patterning the source and drain vias using Mask 3. This process follows the same procedure outlined previously for Mask 2, including surface preparation, photoresist application, exposure, and development. The primary objective of this step is to define precise openings in the oxide layer to guide the upcoming etching process, which will remove the oxide in the source and drain regions. The data collected for this step is shown below:

TABLE XII. PR Height

Sample #	**Δ Height** (µm)
1	1.631
2	1.656

J. Oxide Etch of Source and Drain Vias

The oxide etching step is critical for opening contact windows in the oxide layer, enabling metal connections to reach the underlying silicon at the source, drain, and body-bias regions. As with previous etching steps, the process begins by determining the etch rate using buffered hydrofluoric acid (BHF) on a control sample. Once the rate is established, all samples are etched for a slightly extended duration beyond the calculated minimum to ensure complete removal of the oxide. Particular care is taken with the body-bias contacts, where multiple layers of oxide from earlier processing steps must be fully removed to guarantee proper electrical connection to the substrate. Based on a previously grown oxide thickness of approximately 320 nm during drive-in diffusion and an etch rate measured at 115 $\frac{\text{nm}}{\text{min}}$, the required etch time was calculated to be approximately 3 minutes and 30 seconds.

It should be noted that during mask design, the via windows were made too large, leading to potential undercutting of other features during etching. Fortunately, this did not seem to impact the functionality of the devices.

Step heights were measured using the profilometer, and microscope images were captured to ensure the device was free of contaminants and that the features were properly developed.

TABLE XIII. PR and Etch Height

Sample #	**Δ Height** (µm)
1	1.688
2	1.642

Finally, the photoresist is stripped, and follow-up measurements are conducted to verify that the etch depth aligns with the intended specifications:

TABLE XIV. Etch Height

Sample #	Δ Height (µm)
1	0.656
2	0.621

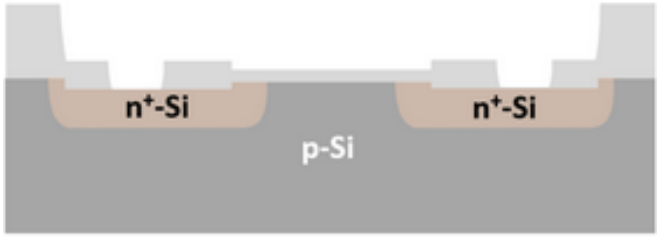

FIG. 13. PR Removal

K. Mask 4 Photolithography

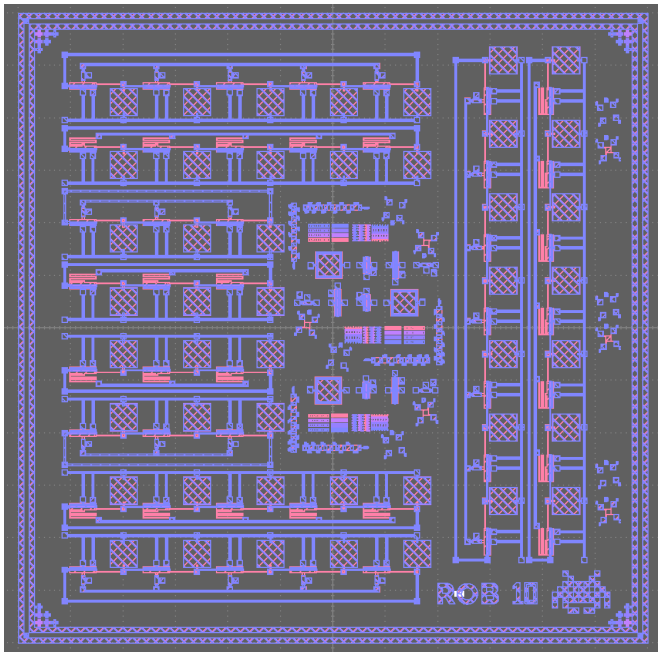

FIG. 14. Mask 4

The fourth photolithography step utilizes Mask 4 to define the pattern for metal deposition. While the overall process is similar to previous photolithography steps, there are a few key differences specific to preparing for metal lift-off. The sample undergoes the standard preparation sequence: acetone cleaning, HMDS vapor deposition, spin-coating with photoresist, and UV exposure. However, after exposure, the wafer is soaked in toluene to soften the photoresist—a critical step for ensuring clean lift-off during metallization. Following the toluene soak, the pattern is intentionally overdeveloped in AZ 400K developer for an additional 30–60 seconds to ensure clear definition of the metal deposition regions. A water rinse is avoided to preserve the effect of the toluene soak. The developed pattern is then inspected under a microscope to verify accuracy and completeness. As with earlier steps, an oxygen plasma descum is performed to remove any residual photoresist. Additional profilometer data shown below:

TABLE XV. PR Height

Sample #	Δ Height (µm)
1	1.721
2	1.756

L. Metal Evaporation and Lift-Off

This step began with a brief 10-second HF dip to remove any residual oxide from the sample surface. Following this, a 300 nm layer of aluminum was deposited using e-beam physical vapor deposition (PVD). This deposition was performed by the lab instructor, ensuring consistent results without complications. After metallization, the samples were submerged in acetone overnight to perform the metal lift-off process. The acetone targets the underlying photoresist, dissolving it and, in turn, lifting away the unwanted aluminum. Thanks to the design of Mask 4, the remaining metal adheres only to the intended source, gate, and drain regions, as shown in the figures below. This marks the final step of the metallization process, successfully removing excess metal and revealing the patterned features on the main die. To verify the results, we conducted a quick metrology pass, capturing both microscope images and profilometer data:

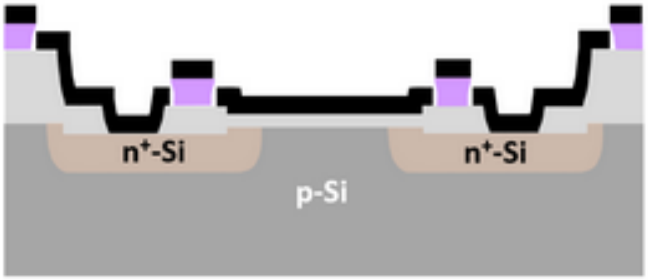

FIG. 15. Aluminum evaporation to form source, drain, and gate electrodes

At this stage, we evaluated the Vernier patterns on the chip to assess the alignment accuracy of our photomasks. These patterns function like rulers, placed on each mask layer and overlaid on the main die, allowing for precise measurement of any misalignment. In an ideally fabricated chip, these markers should be perfectly aligned. Due to inherent human error in mask alignment, small deviations are expected. Our Vernier analysis showed that Masks were misaligned by approximately

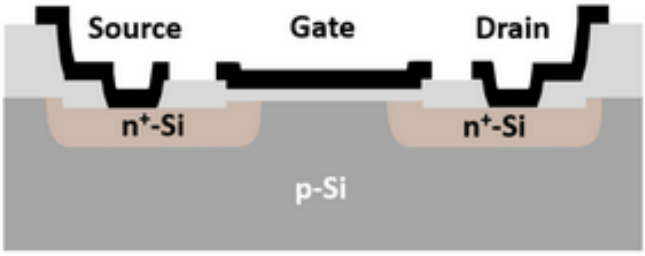

FIG. 16. Metal liftoff to form final MOSFET device

1 µm, while Masks 1 and 2 exhibited a negligible offset ($\ll 1$ µm). This level of accuracy is excellent, especially considering the smallest gate width in our design is 5 µm. The Vernier patterns confirming these results are shown below. From this point, we proceeded to electrical testing:

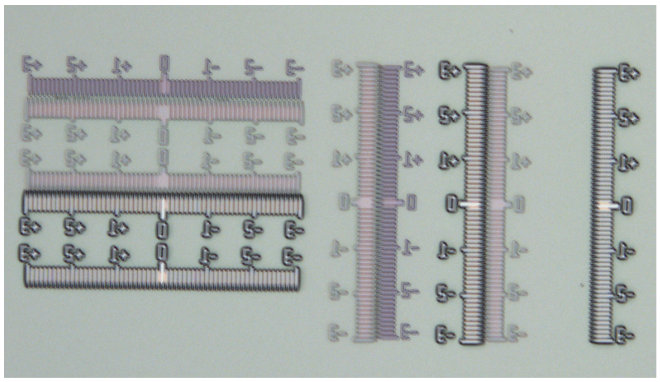

FIG. 17. Vernier patterns for checking alignment

III. RESULTS & ANALYSIS

A. Initial Electrical Characterization and Sintering

While the fabrication process ensures that devices are physically aligned and appear clean under microscopic inspection, the true validation of successful processing comes from electrical characterization. Electrical testing was conducted using a Signatone manual probe station in conjunction with a Hewlett-Packard 4145 Semiconductor Parameter Analyzer for I-V measurements. Each sample contained four 3-stage, four 5-stage, and two 7-stage ring oscillators utilizing NMOS transistors, along with a Van der Pauw structure for resistivity measurements. In addition, the samples included test transistors specifically designed for characterization, featuring gate widths of 100 µm, 200 µm, and 400 µm, and gate lengths of 5 µm, 10 µm, 15 µm, and 20 µm. Transmission Line Method (TLM) structures were also included for extracting contact resistance.

Sample 4 was chosen for electrical characterization due to its superior alignment and well-defined features compared to the other samples. Photolithography for this sample showed minimal mask misalignment, ensuring proper gate overlap with the source and drain regions. Furthermore, the metal lift-off process was notably cleaner, reducing the risk of unintended shorts or incomplete contacts. These factors made Sample 4 the most reliable candidate for accurate electrical measurements.

B. Van der Pauw Testing

The Van der Pauw (VDP) method is a widely used technique for accurately measuring the sheet resistance of thin, uniformly thick conductive films, independent of their exact geometry. In nanofabrication, it provides a non-destructive, high-precision way to characterize material properties critical to device performance. The method relies on a four-point probe configuration placed at the periphery of a planar sample, allowing for current injection and voltage sensing along orthogonal paths. By applying a known current and measuring the resulting voltage drops across opposing contacts, the sheet resistance R_s can be extracted using the Van der Pauw equation, which assumes isotropic conductivity and negligible edge effects. In our lab, we integrated VDP test structures alongside ring oscillators to assess the uniformity and doping quality of the fabricated layers, providing essential feedback for both process calibration and device modeling.

This graph shows the current-voltage (I-V) characteristics measured during a Van der Pauw (VDP) sheet resistance test using a 4-point probe setup. The yellow traces represent two sets of I-V data, likely taken across two orthogonal configurations of the probe contacts—commonly labeled as configuration A and B, where current is sourced through one pair of adjacent contacts and voltage is measured across the opposite pair.

As shown in Figure 18 the X-axis and Left Y-axis (IF1P) is the sourced current in mA swept symmetrically from approximately -10 mA to 10 mA. The right Y-axis (VF2P) measures the voltage in Volts (up to ± 2 V), likely across a voltage probe pair orthogonal to the current path. The linearity of the I-V curves indicates ohmic behavior of the contacts and the film, which is a necessary condition for a valid Van der Pauw measurement. There are two distinct lines because the test was likely repeated in two perpendicular configurations to account for anisotropy or geometric asymmetries.

The slope $\frac{V}{I}$ of each line gives the resistance measured in each configuration. These resistances are plugged into the Van der Pauw equation:

$$e^{-\pi \frac{R_A}{R_s}} + e^{-\pi \frac{R_B}{R_s}} = 1 \tag{3}$$

where R_A and R_B are resistances from each configuration and R_s is the sheet resistance. Using this we found the sheet resistance to be around 17.8 $\frac{\Omega}{\square}$.

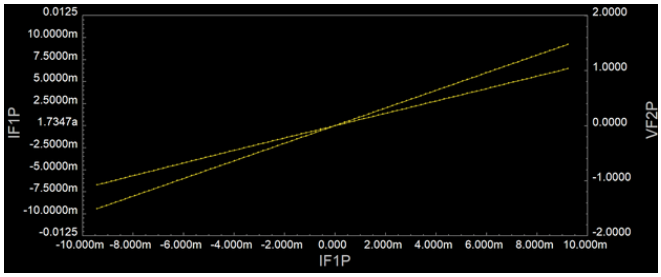

FIG. 18. Van der Pauw plot

C. TLM Measurements

I-V curves were taken for each of the TLM pad separation distances. Although the sample may have already exhibited Ohmic characteristics before annealing, we still annealed our samples because it generally improves device performance as seen from previous labs in ECE120A. An IV curve of our TLM data is shown in Figure 19. The linear Ohmic trend indicates that our devices should be functional.

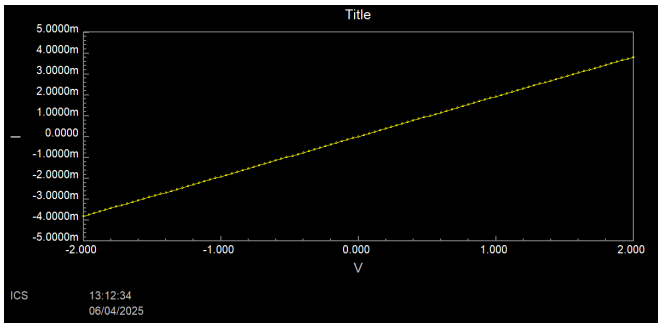

FIG. 19. IV for TLM post-anneal for 80 µm pad separation

D. Ring Oscillator Measurements

To control the oscillation frequency, capacitive and resistive elements are introduced between inverter stages, forming RC delay elements. These RC networks impose a time constant $\tau = RC$, which delays the signal transition at each stage. The total delay through one full loop determines the oscillation period. For an N-stage ring oscillator, the fundamental frequency is approximately:

$$f_{\text{osc}} = \frac{1}{2Nt_{\text{delay}}} \tag{4}$$

where $t_{\text{delay}} \propto \tau$.

To characterize the fabricated ring oscillator devices, we constructed a test setup consisting of a shared ground rail and a 5 V power rail to supply the NMOS transistors. A waveform generator was used to inject a 5 V square pulse into the input of each ring oscillator, serving as a trigger to initiate oscillation. The signal was applied using a probe referenced to the common ground. The output of the oscillator was monitored using a high-impedance oscilloscope probe connected to a digital oscilloscope. This allowed us to capture and analyze the resulting waveform, verifying the presence and frequency of oscillations across different configurations.

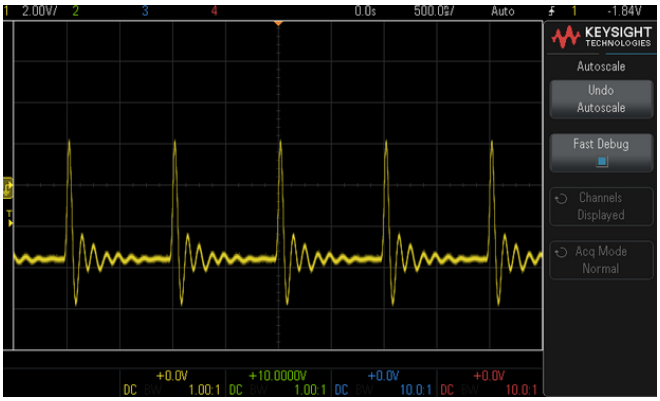

FIG. 20. 5-stage ring oscillator waveform.

The ring oscillators were characterized using the oscilloscope to extract their output waveforms and determine their oscillation periods. The measured periodicities were then compared to theoretical expectations derived from the design parameters. Whereas simulations show a perfectly oscillating waveform, in reality the tested waveform exhibits significant damping effects which may be caused by a variety of imperfect fabrication issues. Nevertheless, the oscillations still demonstrate that our ring oscillators work as expected.

For the 3-stage ring oscillators, the expected oscillation period was approximately 180 ns, based on RC delay calculations using an assumed oxide thickness of 3000 Å and a sheet resistance of 21.93 $\frac{\Omega}{\square}$, values obtained from the previous quarter's fabrication run. However, the measured period was 190 ns, indicating a relatively close match with a slight deviation.

For the 5-stage ring oscillators, the discrepancy was more pronounced. The expected period was 300 ns, but the measured period was only 240 ns. This indicates a faster-than-anticipated propagation delay per stage.

The expected RC delays were based on ideal, uniform process parameters that informed the physical dimensions and layer thicknesses for both the resistive and capacitive elements in the layout. However, process variation in the actual fabrication introduced significant deviations. Post-fabrication ellipsometry and four-point probe measurements revealed that the oxide thickness was around 4000 Å, substantially thicker than expected. Since capacitance in a parallel-plate model is inversely proportional to dielectric thickness, this increase in oxide thickness reduced the capacitance, thereby reducing the RC time constant.

Additionally, our sheet resistance measurement returned a value of 17.8 $\frac{\Omega}{\square}$, which is considerably lower than the 21.93 $\frac{\Omega}{\square}$ used in the original design calculations. The reduced resistance, coupled with decreased capacitance, results in a significantly smaller RC product than expected, meaning the delay introduced by each stage was much smaller than designed.

Together, these deviations explain the shortened oscillation periods in both cases, with a particularly large impact on the 5-stage design. The RC time delay effectively became negligible, and the dominant delay likely arose from the intrinsic inverter switching characteristics rather than the added passive components. This highlights the sensitivity of analog timing structures like ring oscillators to process variation and underscores the need for on-chip process monitoring or post-fabrication calibration when precise timing is required.

IV. CONCLUSION

Three-stage, five-stage, and seven-stage ring oscillator devices were successfully fabricated, tested, and electrically characterized. While the fabrication process yielded fully functional devices, they also exhibited frequency characteristics that deviate from SPICE simulations.

Key discrepancies were observed in the expected versus measured oscillation periods of the ring oscillators. These deviations were traced back to process-induced differences in sheet resistance and oxide thickness, as revealed by Van der Pauw and ellipsometry measurements. The measured sheet resistance was significantly lower than design expectations, and the oxide thickness was considerably greater, both of which led to a reduction in the RC time constants. Consequently, oscillation frequencies were higher than simulated, especially in the five-stage ring oscillators.

Despite these variations, the devices demonstrated functional oscillatory behavior, validating the overall design and fabrication approach. The results underscore the critical role of precise process control and the importance of on-chip characterization structures such as TLM pads and Van der Pauw configurations.

V. AUTHOR CONTRIBUTIONS AND ACKNOWLEDGMENT

We would like to show our greatest appreciation to the lab manager, Dr. Prashant Srinivasan, and the TAs. It would be impossible for us to successfully build the device without constant support and helpful instructions during every lab section. The following is the table of author contributions.

Section	**Contributor(s)**
Abstract	Deven
Introduction	Robin Deven
Methodology	Tony Deven
Results & Analysis	Robin Deven
Conclusion	Deven
Fabrication	Robin Deven Tony Anish Asli
Electrical Testing	Robin Tony Anish

REFERENCES

[1] BrightSpot Automation. *Contact Resistance Measurement*. Available at: https://brightspotautomation.com/wp-content/uploads/2020/05/Contact-resistance-paper-BrightSpot-PVSC2015.pdf. Accessed Mar. 19, 2025. 2025.

[2] Chinese Physics B. *Modeling Capacitance–Voltage Characteristic of TIW/P-INP Schottky Barrier Diode*. Available at: https://cpb.iphy.ac.cn/article/2018/1953/cpb$_2$7$_{90}$97203.*html*. Accessed Mar. 19, 2025. 2018.

[3] Bharadwaj. *Zener Diode Breakdown Characteristics*. Available at: https://engineeringtutorial.com/zener-diode-breakdown-characteristics/. Accessed Mar. 19, 2025. 2025.

[4] All About Circuits. *The Physical Operation of Light-Sensitive PN Junctions - Technical Articles*. Available at: https://www.allaboutcircuits.com/technical-articles/physical-operation-of-light-sensitive-pn-junctions/. Accessed Mar. 19, 2025. 2025.

[5] SURAGUS GmbH. *Four Point Probe Measurement Method*. Available at: https://www.suragus.com/en/technology/four-point-probe/. Accessed Mar. 19, 2025. 2025.

[6] Filmetrics Inc. *Thickness Measurement of Indium Tin Oxide (ITO) and Other Transparent Conductive Oxides (TCOs)*. Available at: https://www.filmetrics.com/applications/ito. Accessed Mar. 19, 2025. 2025.

[7] MKS Instruments. *Mosfet Physics*. Available at: https://www.mks.com/n/mosfet-physics. Accessed Mar. 19, 2025. 2025.

[8] Libretexts. *3: Ideal Diode Equation*. Accessed Mar. 19, 2025. 2025.

[9] MathWorks. *P-Channel MOSFET*. Accessed Mar. 19, 2025. 2025.

[10] R. K. Ratnesh et al. "Advancement and Challenges in MOSFET Scaling". In: *Materials Science in Semiconductor Processing* (2021). Accessed Mar. 19, 2025. URL: https://www.sciencedirect.com/science/article/pii/S1369800121003498.

[11] ResearchGate. *Measured Curves and Dual-Frequency Corrected C-V Curves for MOS Capacitors*. Available at: https://www.researchgate.net/figure/Measured-curves-and-dual-frequency-corrected-C-V-curves-for-MOS-capacitors-with-different$_fig4_2$25580100. Accessed Mar. 19, 2025. 2025.

[12] ResearchGate. *TLM of the Sidewall Ohmic Contacts*. Available at: https://www.researchgate.net/figure/TLM-of-the-sidewall-ohmic-contacts$_fig8_2$60319997. Accessed Mar. 19, 2025. 2025.

[13] ScienceDirect. *Ideality Factor - an overview, ScienceDirect Topics*. Available at: https://www.sciencedirect.com/topics/engineering/ideality-factor. Accessed Mar. 19, 2025. 2025.

[14] SparkFun. *Diodes - SparkFun Learn*. Available at: https://learn.sparkfun.com/tutorials/diodes/real-diode-characteristics. Accessed Mar. 19, 2025. 2025.

[15] TSMC. *3NM Technology - Taiwan Semiconductor Manufacturing Company Limited*. Accessed Mar. 19, 2025. 2025.

[16] Brigham Young University. *Oxide Growth Calculator*. Available at: https://cleanroom.byu.edu/oxidetimecalc. Accessed Mar. 19, 2025. 2025.

[17] Brigham Young University. *Oxide Thickness Calculator*. Available at: https://cleanroom.byu.edu/oxidethickcalc. Accessed Mar. 19, 2025. 2025.

[18] WaferPro. *What is a Silicon Diode? How Does It Work?* Available at: https://waferpro.com/what-is-a-silicon-diode-how-does-it-work/. Accessed Mar. 19, 2025. 2025.