















PROPRIETARY AND CONFIDENTIAL THE INFORMATION CONTAINED IN THIS DRAWING IS THE SOLE PROPERTY OF -INSERT COMPANY NAME HERE-. ANY REPRODUCTION IN PART OR AS A WHOLE WITHOUT THE WRITTEN PERMISSION OF -INSERT COMPANY NAME HERE- IS PROHIBITED.		UNLESS OTHERWISE SPECIFIED:		NAME	DATE	TITLE: 40K STAGE STANDOFF
		DIMENSIONS ARE IN INCHES		DRAWN		
		TOLERANCES		CHECKED		
		FRACTIONAL ☐ DECIMAL ☒ ANGULAR ☐ HUND ☐		ENG APPR.		
		TWO PLACE DECIMAL ☐ THREE PLACE DECIMAL ☐		MFG APPR.		
		INTERPRET GEOMETRIC TOLERANCING PER		Q.A.		COMMENTS:
		MATERIAL: 304 STAINLESS STEEL				
		TEST:				
		NEXT ASSY	USED ON			SIZE B DWG. NO. FSP22008 REV A
		APPROVAL		DO NOT SCALE DRAWING		
						SCALE: 1:1 WEIGHT: SHEET 1 OF 1





FSP22_DRAWINGS_REV_A.pdf

RF Probe Magnet Calibration


July 11, 2022

Categories: General

Using a borrowed Lakeshore F71 teslameter and a vector sensor placed on the sample stage we applied a small current to the magnet at room temperature.

```
field_T = @(Amps) 0.01284*Amps;
```

There is a small intercept due to ambient field.




comsol diode TDGL v2

July 12, 2022


Categories: Comsol, General, Simulation, TDGL

Here are two simulations where I keep the applied current bias the same but vary the GL parameter. The first plot shows the final frame of a $t=20$ simulation. None of the wires switch. Increasing the time to $t = 100$, the $\kappa=100$ wire switches. How to resolve the time issue?.

$t = 20$



$t = 100$



36

1 Basics

From here we see that if $\delta\Psi = \Psi - \Psi_0 > 0$ then $\partial(\delta\Psi)/\partial t < 0$. Reciprocally, if $\delta\Psi < 0$, then $\partial(\delta\Psi)/\partial t > 0$. That means that small fluctuations will relax to zero, so that a weakly perturbed Ψ -function will relax to its steady-state value. In more formal language, (1.111) has a relaxational solution $\Psi(t) = \delta\Psi|_{t=0} \exp(-t/t_0)$, with a characteristic relaxation time $t_0 = \pi^3 T_c / [14\zeta(3)\Psi_0^2] \approx 0.5T_c/\Psi_0^2$.

Remark The relaxation time t_0 is temperature-dependent, since $\Psi_0 = \Psi_0(T)$. Ψ_0 is small near the transition temperature, and larger at $T \ll T_c$. Larger Ψ_0 means shorter relaxation time t_0 . This reflects the property of Bose-condensates (which are close relatives of Cooper-condensates): they have ability to successfully fight small fluctuations and heal themselves. That is why the superconducting state is much “quieter” than the normal metal state.


As always in numerical computations, it is convenient to represent these equations in a dimensionless form. Dividing (2.1) by $\eta = (T_c - T)/T_c$, we have

$$\frac{\pi}{8\eta T_c} \frac{\partial \Psi}{\partial t} - \frac{\pi D}{8\eta T_c} (\nabla - 2i\mathbf{A})^2 \Psi - \left(1 - \frac{|\Psi|^2}{\Psi_0^2}\right) \Psi = 0, \quad (2.3)$$

where $\Psi_0 = \sqrt{8\pi^2 T_c (T_c - T) / [7\zeta(3)]}$. From here it is obvious that time will become conveniently dimensionless if measured in units $t_0 = \pi / (8\eta T_c)$, and the same will happen with distance if measured in units $\sqrt{\pi D / (8\eta T_c)}$. The latter unit is called coherence length: $\xi \equiv \sqrt{\pi D / [8(T_c - T)]}$, since it characterizes spatial decay of the Ψ -function, which in turn describes coherent behavior of electrons in the Cooper condensate. Thus, we will pass to the dimensionless time $\tau = t/t_0$, but we will not use the coherence length ξ for the similar purpose with the spatial coordinates. Taking this into account, and denoting $\psi \equiv \Psi/\Psi_0$, we can write:

$$\frac{\partial \psi}{\partial \tau} - (\xi \nabla - 2i\xi \mathbf{A})^2 \psi - (1 - |\psi|^2) \psi = 0. \quad (2.4)$$



kappa = 10



Important links for BC problem

<https://www.comsol.com/forum/thread/41431/integral-boundary-condition-problem>

RAMP BIAS TEST




SNSPD 4-1 multiplexer

July 27, 2022

Categories: Electronics, SNSPD


Below is a spice simulation of a 4-1 SNSPD multiplexer. It is similar to the design that I have previously fabricated, except for the addition of two resistors on each of the not gates within the decoder. These resistors are needed to prevent sneak current that hinders proper operation if only inductive splitting is used.

The inputs A & B are current sources to the heaters of htrons X1 and X4. The current through the channel is split between resistors R_4+R_5 and R_3+R_6 respectively. With no input to the heater, the channel current is split between the resistors $I_{A'} = R_4/(R_4+R_5)$ and $I_{A'} = R_5/(R_4+R_5)$. With an input to the heater of X1, the current is split between the resistors and the channel hot spot ($\sim 1\text{k}\Omega$), ie $I_{A'} = R_4/(R_4+)$ and $I_A = /(R_4+)$. These currents I_A and $I_{A'}$ as well as I_B and $I_{B'}$ are routed to the heaters of an hTron tree that control the flow of current to the SNSPDs.



The red trace shows the output from the single readout port. The green and blue square waves are the input A and B. The bottom pane shows the current through each of the SNSPDs. The bias of the SNSPDs does not depend on the htron circuit.

Bonus 8-1




comsol diode TDGL v3


August 20, 2022

Categories: Comsol, General, Simulation



Finished formatting some simulation parameters. All units are dimensionless.



Electric field along the length of the wire. at $y=0$



The critical current vs field plot has similar characteristics as the **experimental data**, bottom half looks different.




Related Measurements:


Nanowire diode – SPG732 measurement

3×3 SNSPD array

September 2, 2022

Categories: General






10 μ m

EHT = 3.00 kV
WD = 4.0 mm

Signal A = InLens
Mag = 1.26 K X

Date :23 Aug 2022
Time :11:24:19





The plot above shows a single row input and readout of 3 snspsd columns. In the first column you can see the heater bias as well as some cross talk from other row biases. Pulses in column 1 that align with the heater pulse indicate a photon count from SNSPD Row1 Col1. The pulses in Column 1 that align with the ripple of other heater pulses indicate counts from SNSPD Row2/3 Col1. There is also a latching pulse that appears at high optical power and high current bias. This is likely from the SNSPD+loop switching.

Next step is to work on frame readout using additional pulses on the column bias.

Link post



| *SNSPD 1-Bit array*


NbN Anneal and RGA analysis

September 2, 2022

Categories: General

This deposition was the first deposition after replacing the target. Sample was heated to 800°C after deposition and RGA logging was started when the heat was turned on. The heat was turned off at 15:19, right when the water content decreases. There is an increase of nitrogen while the heat lamp is on but began decreasing before the heater was turned off.






snspsd multiplexer

September 3, 2022

Categories: General

This fabrication process did not yield any detectors. There appeared to be some problem with the smallest features adhering to the PECVD oxide. However all other features looked perfect and I was able to measure a few test structures.



20 μm



EHT = 3.00 kV

WD = 4.0 mm


Signal A = InLens

Mag = 1.72 K X

Date :23 Aug 2022

Time :11:38:45



2 μm 

EHT = 3.00 kV

WD = 4.0 mm


Signal A = InLens

Mag = 15.91 K X

Date :23 Aug 2022

Time :11:17:12



2 μm 

EHT = 3.00 kV

WD = 4.0 mm

Signal A = InLens


Mag = 15.27 K X

Date :23 Aug 2022

Time :11:17:39



NbN/PECVD SiO₂/NbN overlap test.

10 μm 

EHT = 3.00 kV

WD = 4.0 mm


Signal A = InLens

Mag = 2.03 K X

Date :23 Aug 2022

Time :11:29:18



2 μ m

EHT = 3.00 kV

WD = 4.0 mm


Signal A = InLens

Mag = 19.39 K X

Date :23 Aug 2022

Time :11:29:36



10 μm 

EHT = 3.00 kV

WD = 4.0 mm

Signal A = InLens


Mag = 2.95 K X

Date :23 Aug 2022



Time :11:33:09




Test current splitter.

20 μm
EHT = 3.00 kV
WD = 4.0 mmSignal A = InLens
Mag = 634 XDate :23 Aug 2022
Time :11:39:18

Red trace is port on the left, the with DC input. Blue and green traces are ports on the right side. Square wave inputs to the heater switched either the left or right branch. Depending on the heater select the pulse is either pos/neg.





Top branch resistance at 1.3K = 1.2 kohm


Bottom branch resistance at 1.3K = 0.6 kohm

snsdpd multiplexer v2

September 29, 2022

Categories: Electronics

Fab this time worked well (larger features) but heaters on the not gates did not survive. I calibrated the process for 6% HSQ but when going to expose the final layer only 4% was available... annoying. So the thinner resist meant the "tallest" features did not survive the etch. Because the heaters just need to be connected I went in and patched the broken circuit using TiAu liftoff.



20 μ m



EHT = 0.750 kV

WD = 4.6 mm



Signal A = InLens

Mag = 1.64 K X

Date :15 Sep 2022

Time :18:47:31



2 μm 

EHT = 0.750 kV

WD = 4.6 mm


Signal A = InLens

Mag = 9.49 K X

Date : 15 Sep 2022

Time : 18:49:38






200 nm

EHT = 3.00 kV
WD = 4.5 mm

Signal A = InLens
Mag = 60.41 K X

Date :15 Sep 2022
Time :18:51:27





200 nm



EHT = 3.00 kV

WD = 4.5 mm

Signal A = InLens



Mag = 49.07 K X

Date :15 Sep 2022

Time :18:57:24



example heater etched through

2 μ m

EHT = 3.00 kV

WD = 4.5 mm

Signal A = SE2

Mag = 11.45 K X

Date :15 Sep 2022

Time :19:04:30



Some alignment error. I believe there is a problem with the stage position after performing Mark Test Search. On the F125 after performing a test search the stage is placed at the center of the mark then the pattern is moved from the current mark position to the current stage position. Putting the GDS mark location at the physical mark location. I don't think that happens with the HS50


TESTING IN ICE OX

<https://docs.google.com/presentation/d/1j-G0y5nEgA2K3uxNg2TnQ7HrfdIo7EQzAfnhwdbG3io/edit?usp=sharing>

Daily Blog 1/24/2023

January 24, 2023

Categories: General, Superconductivity



This is the first of a daily blog that I hope will strengthen my writing and improve my note taking. I hope these blogs will mostly be literature review but I will also include some lab work. Will I actually post a new blog every day? I doubt it, but one day at a time. As a start lets see if it makes it to February. At a minimum I will at least try to include a photo in each blog– again either from a paper or from some project in the lab.

You may also be wondering, why the need for a daily blog? Well, Matteo recently pointed me to a paper “Vortex ratchet induced by controlled edge roughness” by Cerbu et al. published in 2013. Ten. years. ago. This paper solved the 2D TDGL equation in a superconducting strip with defects along one edge. Which I have been trying (with little success) to do in comsol for some time now. The fact that I have been unaware of this prior work is what prompted me to start this blog.

This mistake is really caused by two main factors: literature review and half-assery. I’ll address the latter first. I have been working on this comsol project for about a year–and has seem some improvement–but right now it is not living up to our research needs. Moving it to a main burner now means that it needs a serious treatment. I should have realized this a while ago (Matteo has had great diode measurements for months now) but now is a good of time as any, especially considering we found the exact work I envisioned from 2013. Which brings me to my second point. Literature review is vital to good research. In football terms it’s “special teams”. Gotta win all three phases. I hope this blog is a step in the right direction.


Ok so getting to that. Serious treatment means running my own numerical simulations and moving away from comsol. Initially, comsol was used for electrothermal simulations of nanowires. Adam and Zhao used it to model the nTron and yTron and Brenden and I expanded that model to work with the hTron & nMem. To continue to build our comsol simulation toolkit I tried to implement a 2D TDGL simulation in comsol. This was made easy when I found the “Shortcut to Superconductivity” by Armen Gulian, which implements TDGL in comsol using the PDE modules. I used this to create simulations of the diode and found similar behavior to the experimental work. The model applies a boundary on $\mathbf{A}$ on the edges of of the strip, a workaround specific to comsol which makes certain geometry difficult and reduces my confidence in the results.

What I plan to do, and will write about tomorrow (hopefully) is to follow the paper by Cerbu and build upon our own numerical simulator.

Daily Blog 1/25/2023

January 25, 2023

Categories: General



I spent most of earlier today trying to get all the probe station temperature sensors working. Turns out we had the wrong cable and now they are going to send us the correct one. No need to rush and make one, though we do have everything needed.

Returning to our discussion yesterday on TDGL simulations, Cerbu used the same model as Kato 1993 and Gropp 1995. To my surprise, it is the exact same model as the one we made in 6.336! Shocking that none of these papers came up when performing that initial research for class. Keyword selection is important, apparently, and should be considered carefully when searching and especially when publishing.

The above figure is from "Effects of the surface boundary on the magnetization process in type-II superconductors" by Kato et al. and bears a striking resemblance to the simulator we built. Hopefully that means our models are truly identical, although I haven't read either paper to ensure it truly is. I'll leave that investigation for another blog.

Assuming it is identical we are still left with a challenging problem. The same problem that caused so many headaches during 6.336. How do we define a current? When following the formulation by Winiecki 2002 we struggled to apply boundary conditions that applied the desired current and magnetic field. Part of our limitations might just have been the time constraint of the course and the mounting pressure to get some results worth presenting in time. Another part might just have been the explanation of the boundary conditions in Winiecki's work compounded by our limited understanding at the time. What I mean by that last point is our confusion with having the order parameter being zero at places where we applied a current. At the time we thought—or at least how I remember us thinking—that we should be able to flow a current with a boundary condition such that we are imagining an infinite strip of current flowing in the y-direction for example. Thus seeing the wave function go to zero where we apply the current felt like some miscalculation. However, when reading "Shortcut to Superconductivity", this is expected, and can be interpreted as the decay of superconductivity into normal

metal contacts. I saw this in every diode simulation I ran in comsol and to approximate a uniform current I simply made the wire longer so the effects from the normal contacts could be avoided.

This feels pretty optimistic. Next steps: find any minute differences between the models referenced and our model, and update accordingly; change the geometry to match the diode (might be more challenging than it sounds); change the boundary conditions so that we are able to apply fields and currents arbitrarily; improve simulator if necessary. Ok, feeling less optimistic now, but at least the simulator still runs without error as of today.

Daily Blog 1/26/2023

January 27, 2023

Categories: General




Figure 1: Evaluation points for ψ (•), A_x and $J_{s,x}$ (X), A_y and $J_{s,y}$ (Y), A_z and $J_{s,z}$ (Z), B_x (◻), B_y (◻), B_z (△)

Following up and answering the questions from yesterday's post. All models in question are derived from the same initial paper by Gropp but have subtle differences based on the area of interest. For example, general vs. not TDGL equation, current boundary conditions, approximations made, solution procedures, geometry. Although it has the same grid spacing, same use of the TDGL equations in either 2D or 3D, and same link variables.

Adjusting for those differences. The generalized TDGL has a more convenient dimensionless form but is strictly valid for $\gamma^2 < 5.79/(1-T/T_c)$, where 5.79 is a numerical factor from the generalization and γ is based on the material. The second difference is related to the current boundary conditions. They have a constraint only on the normal component of the current at the boundary, I don't think we were very careful when defining the current BC, and might have glossed over that $\mathbf{j} = \mathbf{j}_s + \mathbf{j}_n$. The other point of interest is that the entire superconducting volume is surrounded by a grid of normal metal. That was another point of confusion not well explained in Winiecki but is clearly stated in Gropp. With that new understanding, I am feeling ready to edit the code. But, in an attempt to correct past ambitions, I will continue reading because there is still a lot of math I don't understand and I will likely waste more time going down the wrong path.

Gropp has a nice figure showing the model variables on their coordinate grid. It shows where each link variable is located/evaluated on the half grid with respect to the order parameter. Between this difference and the use of the generalized TDGL equation it is worth remaking the model (not the simulator) from scratch. The simulator will be exactly the same. I am wondering now how to best restructure the model. For example, creating classdef for better creation and initialization of model variables, with functions for certain operations

Daily Blog 1/27/23

January 27, 2023

Categories: General



After a week-long delay, the diesaw is back up, and I plan on going to dice in an hour or so. After that it would be good to check some electrical contacts at room temp in the probe station as well as SEM the individual dies and determine the most optimal dose from the dose matrix. Additionally the alignment across the wafer can be determined from the test structures. I did not perform any wafer height mapping prior to the exposure and instead opted for continuous measurement using the height sensor (defaulting to the previously measured height if there is an error). It would be interesting to see if there is any significant change in resolution due to the change in height and therefore focus.

Dicing went well. Only errors came from sample vacuum, which I do not know of any solution. The sample is placed on an adhesive sheet of plastic and all air bubbles are scraped out on the back side. The plastic sheet overhangs the wafer by an inch on all sides and is placed in the center of the cutting plate. The vacuum is engaged and I then place tape around all sides of the plastic sheet, further adhering it to the plate in the event the vacuum is lost. Cuts were made (hairline: $\sim 150\text{ }\mu\text{m}$ wide, alignment: shown in figure) starting at the bottom and moving up. Vacuum errors only occur at the center of the wafer. I am not sure why. When the error occurs the blade feed stops mid cut and moves to local home, I simply acknowledge the error in the system and restart. On restart, the saw advances to the next scheduled cut, leaving the previous partial cut. I then go back and recut those lines. Alignment after is almost identical but I realigned just to be certain.

Back to the TDGL. Because I felt I was switching back and forth between papers too often I will collect a few notes below. That way they are all in one place. In the image with (6) and (7) we have the foundation for the model we created. Note $\$j_{\text{ext}}\$$ in (7). I should go check if that is actually included in our model, I can't remember if we have that term in our functions or if we tried to impose via boundary conditions. Also units.

We scale length in multiples of the coherence length, $\xi = \hbar/\sqrt{2m|a|}$, time in $\tau = \xi^2/D$, the wavefunction in $\psi_0 = \sqrt{|a|/b}$, the vector potential in $A_0 = \sqrt{2}\kappa H_c \xi$, where $H_c = \mu_0|a|^2/b$, the electric potential in $\Phi_0 = (\xi/\tau)A_0$ and resistivity in units of the normal resistivity $\sigma_0 = 1/\kappa^2 D\mu_0$. The so-called Ginzburg-Landau parameter is given by $\kappa^2 = 2m^2 b/e^2 \hbar^2 \mu_0$. The characteristic length scale for variations of the magnetic field is $\lambda = \kappa \xi$ and $\nabla \times \mathbf{A}$ measures the magnetic field in units of $\sqrt{2}\kappa H_c = H_{c2}$. In scaled units equations (2) and (3) become

$$(\partial_t + i\Phi) \psi = (\nabla - i\mathbf{A})^2 \psi + \psi - |\psi|^2 \psi \quad (6)$$

$$\kappa^2 \nabla \times \nabla \times \mathbf{A} = \underbrace{(\nabla S - \mathbf{A})|\psi|^2}_{\mathbf{j}_s} + \underbrace{(-\nabla \Phi - \partial_t \mathbf{A})}_{\mathbf{j}_n} + \underbrace{\kappa^2 \nabla \times \mathbf{H}}_{\mathbf{j}_{ext}}, \quad (7)$$

where S denotes the phase of ψ . The last term in equation (7) can be understood as an external current $\mathbf{j}_{ext}$ with $\nabla \cdot \mathbf{j}_{ext} = 0$. In the following, this term will be omitted. However, it can be easily included in the algorithm, for instance, to model magnetic impurities. In dimensionless units, the dynamics of the superconductor depends on the dimensionless Ginzburg-Landau parameter κ only. For values $\kappa < 1/\sqrt{2}$ one finds a behaviour characteristic of a type-I superconductor whereas for $\kappa > 1/\sqrt{2}$ a type-II superconductor is modelled.

Here they state the difference between their model and Groppe (which is ref 1 below). I forget the meaning of semi-implicit (probably that you only estimate the value at the next time step?) but I do remember having a discussion with Luca about it.

$$\partial_t \psi = (\nabla - i\mathbf{A})^2 \psi + \psi - |\psi|^2 \psi \quad (9)$$

$$\partial_t \mathbf{A} = (\nabla S - \mathbf{A})|\psi|^2 - \kappa^2 \nabla \times \nabla \times \mathbf{A}. \quad (10)$$

The most popular approach to the solution of the TDGL equations, (9) and (10), is a gauge-invariant discretization that is second order accurate in space and first order in time. In addition, a number of other finite difference and finite element methods have been developed. For large values of κ , the magnetic field is nearly homogeneous and equation (10) can be dropped. This case is often referred to as the London limit. The remaining equation has been solved by a semi-implicit Fourier spectral method which is second order accurate in time. An equation very similar to equation (9), the Gross-Pitaevskii equation, is used to model vortex dynamics in dilute Bose-Einstein condensates. Here, we modify the very robust and accurate semi-implicit Crank-Nicholson algorithm used in to include the equation for the vector potential.

Kappa is large for us. ~24 I believe for NbN.

Cerbu TDGL

$$\eta = 2\left(1 - \frac{T}{T_c}\right)$$

$$\frac{1}{\sqrt{1 + \gamma^2 |\psi|^2}} \left(i\phi \psi + \frac{\partial \psi}{\partial t} + \frac{\gamma^2}{2} \frac{\partial |\psi|^2}{\partial t} \right) = (\nabla - \mathbf{A})^2 \psi + \eta \psi - \eta |\psi|^2 \psi$$

$$\frac{\partial \psi}{\partial t} = (\nabla^2 - i\mathbf{A})\psi + \psi - |\psi|^2\psi$$

$$\frac{\partial \mathbf{A}}{\partial t} = (\nabla^2 - \kappa)(\mathbf{A} - \nabla \times \nabla \times \mathbf{A})$$

Daily Blog 1/30/23


January 30, 2023

Categories: General



After dicing on Friday, I plan on cleaning a few samples from the wafer and imaging them in the Gemini SEM. I should remember to determine which dose most had the highest resolution. Sadly It is booked until later tonight, so I will have to find something else to do in the mean time.



One good thing to do is mount a sample in the probe station. Before cooling there are a few things that I wanted to take care of. First, I wanted to add a second clip on the sample stage so that the chip is pressed down on both sides, hopefully improving thermal contact. While doing this I noticed that the camera is a bit unstable and sometimes needs to be held at a certain angle for the best view. A quick solution could be to fix a counter weight (the camera tilts a bit due to it being on one side of the arm) but I will leave it as is for now—because it's not too annoying. Returning to the thermal contact, the second sample holder could not be placed in the tapped hole symmetric to the first clip due to the probe hitting the hex head of the screw. A low profile screw would work but it was just as easy to set it up on a corner hole.



I started by probing a section of the gold ground plane. $\sim 80\Omega$ for a 3pt probe and $\sim 50\text{nm}$ Au. I then measured the resistance of the 20nm NbN nMem channel ranged between $100\text{k}\Omega$ - $1\text{M}\Omega$. Large spread could be do contact resistance and bad probing. I wasnt being too careful. Just happy it was neither open or short. For the second metal layer, the Au heaters, I was measuring a lot of open circuits. Basically the first four fields tested were open. I was thinking it was a probing problem, but in hindsight, I was able to probe the ground plan so it could not have been a prob issue. I carried on trying a few different probes and found the same result, open. Moved to the entire other side of the chip and tried again. This time I was considering the possibility of ESD on the probe tips and grounded the chip with one probe (terminating with 50Ω and probing the ground plane) before measuring the resistance with the other. The result was 3/4 of probed devices measuring resistances in the few 100Ω range. There are the heaters.

In the midst of all this, I got a little worried that the smallest features of the Au heater did not liftoff properly for this wafer run. I was confident due to a successful test run, but maybe something happened that couldn't be seen in the optical micrographs. Luckily the Sigma was free and I imaged another nMem test die. I scanned most of the die and only one (out

of maybe 20) heater was visibly damaged. Different from the percentage I would have predicted from my resistance measurements.



Thankfully, most of the devices looked like the second image.

Some other house keeping. I had to change the SQL query for the probe station logging file. Specifically, adding the `UNIX_TIMESTAMP(time)` function for the time column. Without it, the time was off about 8 hours and caused bugs when zooming in on grafana graphs. Also related to lab logging, the water level sensor just stoped working and began reading negative values. I tried a simple restart, nope. Checked all connections, fine. Actual water level looked ok. I removed the assembly from the chiller to debug on the bench. When running the python script on the Pi it still gave a negative value. I noticed the sensor was adhered to the plastic via the water. So maybe this is the problem. After drying things off a bit, I moved the sensor away from the plastic arm using a few round standoffs. I also wanted to have the Pi write the actual ADC value to the database. But I kept getting problems with the script, writing to the db, though I am 99% certain it should work.

Daily Blog 1/31/23

January 31, 2023

Categories: General

The probe station is cold and it would be good to use this cooldown to characterize a few devices, simple IV curves and T_c sweeps. Mainly to build the measurement setup and test the control scripts. Unfortunately, It appears that the NbN is not superconducting. Channel resistance is $\sim 80\text{--}200\text{k}\Omega$. Hopefully this is because the chip is not actually at the base temperature shown and not because the NbN is damaged. If there is poor thermal contact then I would expect probing the chip would not increase the temperature significantly. Alternatively, probing the sample stage would increase the temperature more significantly. However, neither option increased the temperature, meaning the probes are well thermalized. I will warm up.

In the meantime, I can perform a T_c measurement using the probe station. Hopefully, this will confirm that the NbN is fine, otherwise I'm S.O.L.

With the sample mounted to the PCB, I bonded to the channel of the memory and measured a resistance of $18\text{k}\Omega$. I then bonded a heater and measured 170Ω . At 4.3K the channel resistance was $21\text{k}\Omega$ and the gold resistance was $80\text{k}\Omega$. This is a bad result. I would expect a T_c of $\sim 10\text{K}$ for a 20nm film. Here the devices are still resistive at 4.3K .

Now that we have determined that thermal contact is not an issue (since it is irrelevant in immersion measurements), I will do a more thorough scan of device resistance tomorrow. I could also try a chip towards the center of the wafer where the thickness is possibly greater. As to why the devices are not superconducting? My best hypothesis is that the selectivity of the maN was not sufficient when etching the NbN and resulted in a much thinner film. I would think that is unlikely.

Daily Blog 2/1/23

February 1, 2023

Categories: General

I started today by measuring the resistance of every device on two dies from the first wafer scale process. The results below are not too promising. The x and y axes indicate the device under test, as well as the physical location on the die. All units are in Ohms. I am assuming that lower resistance is better (the film needs to be thick enough to superconduct at 4K) and "working" devices have resistances the 16-20kOhm range and "broken" devices in the >50kOhm range. I should determine the approximate sheet resistance for each case. Below that are the results for Au heater measurements. Here I assume 170-300Ohms to be working, and I am using black to indicate an open circuit.


Channel							
	1	2	3	4	5	6	7
A	18.762E+3	46.60E+3	64.10E+3	38.30E+3	108.60E+3	160.50E+3	38.10E+3
B	55.10E+3	79.40E+3	43.50E+3	53.30E+3	73.50E+3	44.40E+3	59.30E+3
C	55.60E+3	17.20E+3	16.10E+3	38.10E+3	24.20E+3	26.90E+3	129.60E+3
D	51.60E+3	41.60E+3	19.10E+3	37.60E+3	56.30E+3	27.80E+3	38.60E+3
E	58.30E+3	27.00E+3	122.60E+3	154.40E+3	31.30E+3	55.70E+3	40.80E+3
F	20.00E+3	3.40E+6	39.00E+3	43.60E+3	39.30E+3	1.77E+6	29.10E+3
G	10.30E+6	56.60E+3	17.50E+3	100.00E+6	132.00E+3	1.20E+6	111.60E+3
Heater							
	1	2	3	4	5	6	7
A						30.00E+6	170.80E+0
B							187.80E+0
C						190.40E+0	197.30E+0
D			195.80E+0	220.00E+0		60.00E+6	198.90E+0
E				195.50E+0		323.00E+0	174.00E+0
F					176.00E+0		173.00E+0
G		232.00E+0	208.30E+0		212.00E+0	210.00E+0	214.70E+0


From this you might be thinking that there is a spacial dependence to our yield; the left side has more working devices. I partially agree with that hypothesis, however, I likely damaged devices in the top left via ESD, so I wouldn't draw any strong conclusions from this.

Below are the results from the second die. This die contained wires of various widths and identical lengths, so there should be a clear trend, i.e. column 1 has the highest resistance while column 7 has the lowest. This is clearly not the case.

Wire	1	2	3	4	5	6	7
A	10.80E+6	2.40E+6	10.50E+6	132.60E+3	6.40E+6	2.40E+6	1.96E+6
B	10.90E+6	25.00E+6	3.70E+6	330.00E+3	394.00E+3	250.00E+3	10.90E+6
C	80.00E+6	3.50E+6	4.50E+6	6.50E+6	4.30E+6	5.00E+6	1.40E+6
D	5.10E+6	4.40E+6	11.00E+6	25.00E+6	4.50E+6	3.00E+6	30.00E+6
E	5.10E+6	27.00E+6	70.00E+6	16.10E+6	100.00E+6	4.50E+6	20.00E+6
F	2.60E+6	75.00E+6	4.70E+6	35.00E+6	18.00E+6	21.70E+6	4.10E+6
G	35.00E+6	4.50E+6	6.30E+6	1.05E+6	500.00E+3	10.50E+6	6.50E+6

I went and imaged the nMem chip again, the same as the data shown above and attempted to see any visual issue with device. Really I would hope to see something wrong with the NbN, but it is covered with 100nm of SiO₂. For the gold, I am expecting to see lots of broken wires. Spoiler alert, I do.





I have some more SEMs looking at the interface between layers but there seems to be an issue with WordPress that is preventing me from uploading them right now. I will try and come back and put them here. Overall the conformality of the gold over the oxide looks good. There are some sections where the oxide on the tapered portion of the gold pad was ripped off. This is likely from aggressive sonication performed when striping the resist. However, this is only a cosmetic issue and will be avoided completely with a higher temperature (70deg C) strip.

An obvious reason for the wide range of NbN resistances was not discernible from the SEMs. I plan on testing another die tomorrow; one towards the center of the chip that might have a thicker film. It would be beneficial to outline the NbN process in detail and I will do so here.


- Deposit NbN SPG796, Thickness 20.053.
- Spin: ma-N 2401 100rpm for 10s, 3000krpm for 60s.
- Bake: 90°C for 60s
- Expose: $. * 350\mu\text{C}/\text{cm}^2$
- Develop: AZ300 30s (not the recommended developer but similar TMAH concentration <3%) + DI rinse
 - Etch 50W Bias, 0W ICP, at an etch rate of 6.5nm/min gives 185s, plus a 20% over etching factor is 220s.
 - Measure resistance: I measured 5KOhms between blank corner dies. Meaning there is still NbN.
 - Etch 50W Bias, 0W ICP, additional 60s (total 280s)
 - Measure resistance: Open. NbN film removed other than defined structures.

Total etch time is just under 5 minutes. According to Emily's paper that should leave about 50nm of maN. The new etch process has a slightly higher etch rate than the one she used and I had not considered this to be an issue because of the 50nm buffer. I am now reconsidering.

Daily Blog 2/3/23

February 12, 2023

Categories: General






Yesterday the blog was down and I was unable to make a post. I believe this is due to out of date plugins. Bill Adams is currently working on updating the blog so hopefully these problems will be resolved soon.

Quick recap of yesterday. I measured a test structure at the center of the wafer and found that the resistance values were nominally lower.

Wire	1	2	3	4	5	6	7
A		78.00E+6	48.00E+6	4.00E+6	1.20E+6	14.50E+6	9.80E+6
B	22.00E+6	1.70E+6	1.50E+6	460.00E+3	240.00E+3	1.00E+6	2.30E+6
C	21.50E+6	6.10E+6	470.00E+3	3.40E+6	732.00E+3	50.00E+6	1.40E+6
D	4.10E+6	860.00E+3	780.00E+3	429.00E+3	523.00E+3	1.00E+6	276.00E+3
E	2.30E+6	400.00E+3	1.30E+6	2.60E+6	163.00E+3	1.80E+6	33.00E+6
F	2.70E+6	1.05E+6	340.00E+3	250.00E+3	260.00E+3	225.00E+3	2.30E+6
G	40.00E+6	380.00E+3	2.40E+6	2.00E+6	1.00E+6	1.20E+6	6.00E+6


During the cooldown the probe moved due to the thermal contraction. This caused the resistance to either measure open (no contact) or low (contact with the gold ground plane). I could have removed this data but I decided to keep it in so we


have an idea of when the most probe movement is expected. (Ignore the long periods where the probe had lost contact.) I included this 300-3K sweep just to observed the RRR, which was very high, ~ 4 .




I am not sure what to make of this decrease (is it the beginning of a transition?) or how I should interpret the large RRR (increase from $\sim 300 k\Omega \rightarrow 1.2 M\Omega$). I would like to redo this measurement with samples with a higher T_c to better characterize the measurement setup (seeing as this was the first ever measurement with the new probe station) and I will try again once the new cables for the temperature sensors arrive.

I still find it hard to believe that a 20nm thick NbN film is not superconducting at 4K. I also find it hard to believe that the maN did not last through the etch and that there is a thinner NbN film. Below are a few SEMs of the NbN layer after etching but before deposition of oxide. I was very happy with these results.





10 μm

EHT = 3.00 kV

WD = 4.7 mm


Signal A = SE2

Mag = 1.41 K X

Date: 26 Dec 2022

Stage at T = 0.0 °

MIT.nano



500 nm



EHT = 3.00 kV

WD = 4.7 mm


Signal A = SE2

Mag = 20.45 K X

Date: 26 Dec 2022

Stage at T = 0.0 °

MIT.nano



200 nm



EHT = 3.00 kV

WD = 4.7 mm


Signal A = SE2

Mag = 50.53 K X

Date: 26 Dec 2022

Stage at T = 0.0 °

MIT.nano




The nMem array has some odd features that I hope are from residual resist. I would be interested to hear others' opinion.

Daily Blog 2/13/2023 – Super Blog


February 13, 2023

Categories: General, IoT

Phase 2 (NbN) - Thickness color map



Phase 1 (sio2) - Thickness color map




Well that was a long absence due to an error on the blog's host machine. Thankfully, Bill was able to get it back up and running on Friday and will continue working on updates. The blog will either be transferred to a new machine or updated here, depending on what RLE, or Karl, deem to be the best action.


I am starting to write this blog a day early and will try to collect my notes over the past few days in one big Super Blog (I was of course inspired by tonight's Super Bowl between the Philadelphia Eagles and the Kansas City Chiefs). The first part of the blog will contain some uniformity measurements on 4in wafers and $\$T_c$ measurements in the probe station. I performed these measurements to rule out the possibility of over etching or reduced thickness that may have caused my samples to have low $\$T_c$. The second part of the blog will contain some information about a water level sensor installed on the AJA chiller. Most of this work was copied either from online tutorials or from work by Mackenzie. Lastly, I'd like to compose some notes on papers I've come across over the past week.

The difference in room temperature resistance measurements of the test structures lead me to believe there might be a uniformity issue with the NbN deposition. I performed three maps of the wafer thickness using the ellipsometer. The first was just the SiO₂ thickness on a test 4in wafer. The second was a measurement of a 500s deposition of NbN. The third was a map of the NbN deposition after a 60s NBN_ETCH (2Pa, 0W ICP, 50W Bias).


The uniformity of the 300nm oxide is quite poor. Measured values range between 305 and 350nm with the lowest value observed in the center of the wafer. This gradient is visible by eye; the thinner areas correspond to the purplish regions in the photo.

Phase 1 (sio2) - Thickness color map






Phase 2 (NbN) - Thickness color map




Phase 1 (sio2) - Thickness color map




The measured thickness of the NbN deposition ranges between 22 and 25nm. Which is quite good, but might not explain the difference in measured resistance values. After etching the thickness range is 19 to 21.6nm.

Phase 2 (NbN) - Thickness color map



Phase 1 (sio2) - Thickness color map




To confirm that there is no issue with the measurement setup (i.e. there is proper thermal contact, low noise, low heat transfer from the probe, etc.), I performed a 500s deposition of NbN on two thermal oxide samples. The first sample was placed in the center of the chuck and the second towards the edge. I measured the oxide thickness prior to deposition for improved ellipsometry fitting. The NbN thickness and sheet resistance were measured to be and respectively. Thus the edge is thinner as expected.

The samples were then placed in the probe station and cooled to base temperature. I then performed T_c measurements by ramping the heater to 20K and recording the temperature and resistance every second. This measurement also confirmed that we are able to log to the database while also recording our measurement. See the two lines of code I run from the command line to start up logging with pythonw.


Start: `pythonw probe_station_logging.pyw` Just save whatever python script you want as a .pyw
 Stop: `taskkill /f /im pythonw.exe` Stop will also close out your Spyder window.

I first performed setpoint ramp tests with three different ramp rates. The dashed line represents the PID setpoint and the solid line represents the actual temperature. Ramp rates of 10, 20, and 40 K/min were used. I am mostly interested in


determining the fastest possible ramp rate before there is considerable lag, as well as determining optimal values for the PID.




I was able to observe a transition at $\sim 8.7\text{K}$ for SPG802 and $\sim 7.5\text{K}$ for SPG803. Showing that there is a 1K difference in T_c from depositions in the center of the chuck relative to the edge. Clearly there is some issue with the measurement setup. I am using a source meter with $10\mu\text{A}$ bias and simply probing the center of the die. I will have to investigate further.



Below is an image of the water level sensor. The **sensor** is resistive and heavily dependent on its position so it is fixed to a piece of acrylic and housed in a 1.5" PVC pipe. The ADC (MCP3008 8-Channel 10-Bit ADC) is on a breadboard mounted to the acrylic which contains the peripherals and external connectors to the Raspberry Pi (I followed this [guide](#) almost exactly). All are housed in a 3D printed inclosure (the black box in the photo).



After a few days of logging to the database I noticed the level reading change drastically without any perturbation to the sensor or to the real water level. I attributed this to water between the sensor and the acrylic; the data sheets recommends placing in a 1" PVC for support so there is nothing directly in contact with the sensor, but we only had 1.5" PVC in the lab. To resolve this I moved the sensor away from the acrylic using some round standoffs as shown below. The white bands of tape hold down the standoffs and the sensor.



After moving the sensor tape away from the acrylic the water level reading is stable but a little noisy. On Grafana I am tracking the most recent value along with a 30day moving average. It consistently shows that the water level is dropping and an alert for a minimum water level can now be created. The water level is simply a linear mapping of the ADC value, I did not convert it directly to inches.



Papers

This fabrication paper has fewer layers than the MITLL SFQ5ee process, making it a bit easier to understand. One take away is the use of much thicker wiring layers. We should further break up lithography layers such that any features that need to cross oxide layers are much thicker.

[Fabrication Process for Monolithic Integration of a Nitride Superconductor-Based Superconducting Qubit with a Single Flux Quantum Control Circuit](#)

The DARPA NaPSAC BAA kicked off a discussion of nanowires for multiply-accumulate operations (MAC). The closest work from the group that resembles the BAA is Murat's crossbar array. Also [Dirk's group](#).

[Design and characterization of superconducting nanowire-based processors for acceleration of deep neural network training](#)
[Si microring resonator crossbar arrays for deep learning accelerator](#)

Other papers.

[Solving Graph Problems Using Gaussian Boson Sampling](#)

[Superconducting fluctuation effect on epitaxially layered films of superconductor NbN and half-metallic Heusler alloy](#)
[Ultra Strong Magnon-Photon Coupling Achieved by Magnetic Films in Contact with Superconducting Resonators](#)

This is Benedikt's paper from the NIST group.

[Trap-Integrated Superconducting Nanowire Single-Photon Detectors with Improved RF Tolerance for Trapped-Ion Qubit State Readout](#)

[A 1-GS/s 6–8-b Cryo-CMOS SAR ADC for Quantum Computing](#)

This paper below is really cool. They fabricated an SNSPD directly on the end of a fiber.

[All-fiber device for single-photon detection](#)

Daily Blog 2/15/23

February 15, 2023

Categories: General

Ok, I didn't post yesterday. What's my excuse? Meetings.

On Monday, when celebrating Torque's time in the group, we revised an old argument between Di and me. What is more elite? Playing in the NBA, or being a Ph.D. at MIT? I changed the wording from *harder* to *more elite* to clarify that we are not talking about the job description but about achieving that position relative to others trying to do the same. We mostly agreed that how you define the question is more difficult than coming up with some answer. Here is how I am going to define it.

$$\frac{\text{NBA players}}{\text{College players}} > \frac{\text{MIT PhD EECS}}{\text{Other EECS PhD}}$$

From the first Google results I found.

$$\frac{450}{5,510} < \frac{814}{7,783}$$

That gives $0.081 < 0.104$. I chose the denominators to be just Division I men's athletes and just US EECS graduate students (no international). I'll stop there and get back to more important matters.

Yesterday I began the He3 condensation cycle in the 300mK cryostat. While waiting, I performed several updates to the PC. There was some issue with the Disk usage always being maxed at 100%. After updating it seems to be resolved. I was also able to get the logging script to the database working. I have also been expanding the instrument library for the Lakeshore temperature controller. Allowing us to control the various output loops from python. I hope that by adding the data to the database/Grafana we will be able to observe the temperature dynamics at each stage of the He3 condensation. This will undoubtedly improve my understanding of the cooling process but I also hope it will allow us to improve the process in the long term. Letting us cool down faster and remain cold for longer.

The cooldown was interrupted by discussions for the DARPA NaPSAC proposal. Because the abstract deadline is tomorrow and this program is not the best fit for our group, we will use this opportunity as an exercise. Build those writing muscles.

The goal of this BAA is to advance ultralow power in-memory architectures for scientific modeling, simulation, and analysis. One proof of concept that they intend this program to build upon is shown in the figure below. It is a Si microring resonator (MRR) crossbar array published in *ACS photonics*. Through a PIC they perform optical matrix-vector multiplication. The basis of our proposal is to argue that microwave matrix-vector multiplication using superconducting nanowires is more energy efficient (<10W/petaOPS) and more scalable (1×128 array).




Figure 1. Fully integrated MRR crossbar array. (a) Schematic of the MRR crossbar array including couplers, MZI switches, and PDs. (b) Plan-view image of the fabricated MRR crossbar array corresponding to (a). (c) Measured output spectra of the MRR crossbar array when optical signal is injected from Port 1.

Karl, Reed, and I met and discussed what a nanowire architecture would look like and performed some back-of-the-envelope power calculations. We had more success with the latter.

Daily Blog 2/16/23

February 16, 2023



Categories: General



The 300mK fridge is currently at its base temperature. I took a screenshot of the temperature readings during the condensation. Overall, it took much longer than I remembered to reach the base temperature. In Grafana, I added annotations for each step of the SOP. A few notes on this. When heating the adsorber I used a low heater range (either 2 or 3) and the increase was quite slow with minimal influence on other stages. In the future, it would be better (faster) to use range 4. Once the sorb reached 70K the SOP instructs you to turn off the heater and wait for the sorb to reach 40K; this should take 30 minutes according to the SOP. I waited more than 2 hours and the sorb never reached 40K. I proceeded with the temperature around 50K. Following the SOP from then on, a base temperature of 0.33K was reached. The plotted temperature data makes it very easy to monitor the temperature increase on the 3K stage when in contact with HS1, i.e. HS1 in position 4.



With my sample at the coldest temperature we can achieve in the lab, I was ready to see some superconducting behavior. If not, I would have no idea what went wrong. Thankfully, the IV characteristics showed almost exactly what I was expecting. Large initial contact resistance followed by a switch at $\sim 250\mu\text{A}$. I believe this is due to the NbN over the gold pad being too thin and not superconducting, inducing some corner resistance. I am not aware of the correct expression to describe this. But regarding the fabrication process, the deposition and liftoff steps should be swapped.



On the left, two of the three nMem channels had clear switching and the other (X1) had a resistance of $\sim 100\text{k}\Omega$. On the right, two of the three nMem heaters had resistances of $\sim 64\Omega$ while the third (X2) had a resistance of $\sim 100\text{k}\Omega$. Again, probably corner resistance. On a positive note, there seems to be some consistency between devices with good connections. Similar switching currents and both heaters measuring 64Ω imply the lithography is good. Actually, now that I am thinking about it more carefully, I would have no explanation for one of the heaters being $100\text{k}\Omega$ (X2 on the right figure). Seeing as there is little worry of conformality and gold-to-gold contact resistance should be minimal. What is more likely is that I forgot I bonded to an extra channel, thinking I cared more about what was going on with the NbN rather than the heaters. That thought is in my memory. Meaning the $100\text{k}\Omega$ X2 trace on the right plot really belongs on the left plot as X4. I needn't be too careful in labeling this cooldown. The agreement between the two pairs of superconducting and resistive IV curves also indicates that there is good uniformity across deposition and etch.

During an IV sweep up to 1mA and $R_{\text{switch}}=15\text{k}\Omega$ the He3 pot temperature increased to 0.5 , as seen at the 18:30 mark, and again later under different bias conditions. The characteristic valley occurs in the middle of the decreasing sweep; the time between the first retrapping and the switch during the negative bias. In other words, when the hotspot has cooled. At the second annotation, marked at 19:10, the He3 temperature increased sharply without any applied heat and did not recover. I assumed this indicated that the He3 condensate had been depleted. About 23 hours at 0.3K . I tried to speed through the regeneration process. Time to return to $<0.4\text{K}$ was just under three hours.



Daily Blog 2/17/23

February 20, 2023


Categories: General

(and a bit over into the long weekend)

Before warming up the 300mK there are a few more things I could measure. First would be the critical current as a function of heater bias. The second would be an R-T sweep as the coldest stage warms up. Outside of that, I need to install a temperature sensor in the probe station that sits directly on the sample stage.


I had some difficulty performing a nice R-T sweep up from 0.3K, up to say 10K, but I was able to observe a transition at 6K. This is about what I expected. However, it is quite low compared to our historical values for films of similar thickness. Might be worth performing a few depositions and measuring the T_c .

Thinking that my film has a much lower T_c than expected, I decided to query the database for all films with entries for resistivity and T_c . As it turns out my assumption was incorrect. The two test depositions I performed for testing in the probe station fell in the center of the historical data. Below, the color represents the date of the deposition and the two white dots in the center of the point cloud are the two test depositions.



This indicates that our film growth is ok. Any decrease in T_c for the nMem chip likely came from processing, i.e. thin corners over the gold pads. Besides swapping the deposition and liftoff it would also be beneficial to impose that any nonplanar structure is patterned with photolithography such that the material deposition can be much thicker; photoresist is much thicker than electron beam resist.


The probe station needs a temperature sensor installed on the sample stage and it will be wired through a feedthrough on one of the probe arms. I don't really want to disassemble the already installed arms and will use this as an opportunity to finally get around to machining a back plate for a third arm assembly. In the image below you see some aluminum stock. On it is, a probe arm back plate (I am making up that part name), an outline drawing of the backplane, and an XY positioner—but not one provided by Lakeshore.



For some reason, we do not have a third back plate anywhere in the lab, nor did we have a third vacuum bellow, but we did have a third probe arm feedthrough and Z positioner. My best guess as to *why* we are missing a few parts is that an arm broke, and was reworked in some capacity. I am mainly thinking of the optical fiber that was mounted on an arm with no Z positioner—which was found in the black toolbox next to the Free Space with the other optical components.

Anyway, I made the part but with mounting holes specific to the XY positioner in the photo. Below is a photo of the machined part and what the COTS probe arm assembly.






Daily Blog 2/21/23

February 21, 2023

Categories: General




Back from the long weekend. On the schedule for today, we have a review of Matteo's counter paper, a meeting about Val's master's thesis, the superconductivity meeting, and a proposal meeting. So forgive me if this blog is a bit short.

I installed the part I machined for the probe station and added photos to the blog post from 2/17. Thankfully, everything fit properly, and I will look to install the probe arm after rewiring the temperature sensor feedthrough.

Regarding the layout, there are a few changes I am considering. First, I am wondering if it is worthwhile just to reuse the same GDS from the first wafer run. This would be advantageous in terms of time; not needing any redesign and the same processing time. On the other hand, a new layout would be able to accommodate new devices and resolve mistakes in the old GDS. This would be slower; longer write times (assuming more devices on the wafer) and more time producing the layout. That being said redoing the layout is not overly laborious.

The primary concern for future wafer runs is the gold heater liftoffs. In the current process, I successfully lifted off 120nm of gold over a 100nm oxide step making contact with a previous gold liftoff pad. I am realizing now, that this process probably should not have worked. The reason I followed this process was because of the narrow heater widths, which constrains us to ebeam, which then constrains the heater thickness to some fraction of the resist thickness. In this case, 2.5krpm ZEP 530 was between 140nm-150nm. Almost a 1:1 ratio of resist thickness to metal thickness is nonideal (typically I hear 1:3 as a guideline) but I was able to achieve 100nm wide heaters on a bare test wafer. With a working liftoff subprocess, I performed the same process on the main process wafer. An SEM of that wafer is shown below.



400 nm


WD = 4.8 mm
Signal A = SE2

EHT = 3.00 kV
Stage at T = 20.0 °

Mag = 20.18 KX
I Probe = 100 pA

MIT.nano

In the above SEM, we are primarily focused on the heater traveling horizontally in the middle of the image where we can interrogate the sidewall profile. On the left, the heater is elevated on a 100nm thick oxide layer, moving further left of the screen we would run into the DUT. To the right, the heater makes contact with a craterous pad; also made of gold. The craters come from etching. In the middle, we have the substrate oxide. The middle region is not present in the GDS and the oxide spacer was designed to abut the edge of the pad, indicating about $1\mu\text{m}$ of alignment error with the MLA. I find that to be sufficient, but I am not sure if that meets others' expectations of the tool.

Why change the process if it was successful, i.e. "if it ain't broke, don't fix it."? Well, here I see a high probability of future broken circuits with a (potentially¹) easy solution. I will break up the liftoff into two steps: planar for nanostructures, and microstructures for nonplanar. I should have thought of this earlier (everyone does this), but I was probably fixed on the cost and time savings of one less lithography and evaporation step. For the next process run, I will target $>200\text{nm}$ liftoff metal thickness and $\sim 1\mu\text{m}$ thick photoresist.

There are a few other points learned from this image. The first is that the etch is well-calibrated. A longer etch would have made the valley deeper—potentially causing the heater to shear on the edges—while a shorter etch would have left oxide between the two gold contacts—creating an open circuit.

Actually, this blog turned out longer than expected.

¹I say potentially because I haven't run this *exact* process yet.

Daily Blog 2/23/23


February 23, 2023

Categories: General




Today (and yesterday because I did not finish the blog in time) will be focused on continuing the redesign of the wafer scale process layout. One problem I was running into was routing wires between nearby ports. Ideally, each port would be indexed and routed automatically in a single loop. But this is not always possible due to the non-numeric naming of ports. Some phidl devices name ports with strings like 'drain' or 'W1' (West 1) to make identification easier, but makes scripting more complicated.


I wrote a simple auto-routing script that first sorts the ports by distance from the origin and then routes each port starting from the furthest out.. An example is shown below; note the 180° rotation of the blue square.



It doesn't understand crossing; it just sorts the ports.




With the semiautomatic routing script, I am able to connect ports using a loop as opposed to indexing by hand. An example use case for an 8×8 nMem array is shown below.



Another issue found with the old layout is the clearance around the die for mounting in the probe station. Dies containing the arrays of SNSPDs and memories only had $500\mu\text{m}$ clearance. It was possible to mount without contacting the ground plane (maybe better if we do?), though $\geq 1\text{mm}$ clearance would be more comfortable.

The microscope works well. On the right, you can see the copper clips holding the sample to the stage. You can also see some resist that was redeposited during the strip.






Daily Blog 2/24/23


February 24, 2023

Categories: General




Finalizing the layout. I am having trouble staying organized and I am going to use this as a reference. If anyone sees this prior to fab I could try to include their design.

Most of the dies will be formatted like the photo below. First, alignment marks, then a grid of sub-circuits; each sub-circuits has an ID (e.g. A1) and field alignment marks¹. Depending on the die, the grid of sub-circuits may contain a single device repeated, or a sweep of some geometry. In the previous process run, I swept geometry along the x-axis and dose along the y-axis.



Below is a collage of the different circuits that I am including on this wafer.




The loop sweeps of the nMem and SNSPD single cells vary the size of the loop, increasing the inductance. For the nMem, I simulated the inductance ratio and heater suppression as a function of the loop size. Previously these two parameters were slightly coupled, leading to errors induced by the uneven heating of each branch. The double crossing of the heater on the

right branch ensures that an equivalent number of squares are heated between the left and right branches. This was the main result of the nMem_v2 experiment and the simulations confirmed this. For the SNSPD loop sweep, I did not take as much consideration in the design. The primary goal of this circuit (or the array) is to see if we can yield a large number of detectors coupled with another nanowire device. Larger loop inductance should lead to a larger number of counts before the detection efficiency goes to zero. This is assuming an operation similar to that of Onen's work, but that is not the intent of this design.

The nTron subcircuit has additional ports added to the drain, gate, and load. This is mainly to add local voltage probes and calibrate out cable resistances. For the gate, the additional port will be used to supply a DC offset to the choke and investigate our minimum gate current. The sweep over the die shifts the gate from centered with the channel to slightly below. This is to investigate the nTron gain. The hTron is configured in a 4-point measurement for both the channel and the heater. The sweep of the hTron die varies the width of the channel and heater.

Each die is then placed on the wafer with an ID and dicing alignment marks. Thus, a device would be completely identified as follows: SPX999 A1B2.




Not using them. Field alignment is not accessible via the software at this time.

Daily Blog 2/27/23


February 27, 2023

Categories: General



It is the beginning of a short week for me. I will be out Wednesday to Friday for some skiing in Colorado. I am excited about the trip but there are a lot of things I'd like to do before then.




First, I was finally able to get a hold of our contacts at NRL and I am hopeful that we can get our irradiated samples back for testing. This would be a great year 1 result for the BSF program. Second, the layout for the second round of the wafer scale process is complete. Thought I am reconsidering the number of dies included. Fewer dies would decrease writing times and speed up the fabrication process overall. I could split the process into two different wafers such that I can process it in parallel and test it intermediately.




Before moving on, I need to think about other fabrication questions that can be answered with the previous wafer run. For example, the line width as a function of the dose was determined. I should measure a few more wires to add to this plot. In fact, there is a lot more data I should take.

Following that idea, I went down and measured the 8×8 SNSPD arrays at 3.3K. I already know this device is unlikely to work—having already measured a few devices. The measurements below tell us that the heater yield was 100%. The channel (superconducting layer) yield was much worse. The first column of each device had a low resistance of $\sim 200 \Omega$ while the remaining columns had much higher resistance. Again, I believe this is from the NbN being too thin over the gold contact pad edge due to the wide range of measured values. $I_{\text{src}} = 10 \mu\text{A}$ for all measurements. As to why the first column specifically has low resistance, I am not sure.

The heaters for central rows are shorter than the rows toward the edge. In other words, rows 1 and 8 have higher resistance than 4 and 5 as expected. The longest heater has 765 squares giving an estimate of $\sim 0.35 \Omega/\text{sq}$ for 20nm Ti+100nm Au.



Despite the redeposited resist covering many of the contacts, the connections were fine.



Daily Blog 3/1/23

March 1, 2023

Categories: General

I had half of a blog written yesterday before my Italian colleagues pointed out a paper from Zhao, published just one day ago. The title, "[Splitter trees of superconducting nanowire cryotrons for large fan-out](#)" appears to be the first publication of the new APL Advances in Superconducting Logic special issue; editor's pick. Without having to even look at the paper, the title itself makes clear how closely related this work is to ours. I will discuss it towards the end.

Looking back on my past blogs there is still a heavy bias toward experimental work and very little discussion of recent literature. Recall, the impetus for starting a daily blog was to motivate more academic endeavors. Let's correct that.

Last week, I came across this [short paper](#) about designing a superconducting qubit using the [Qiskit-Metal](#) toolkit. What I liked about this paper—aside from its brevity—was its simplicity. I have read more prominent and detailed qubit design papers (such as "[A quantum engineer's guide to superconducting qubits](#)") and always felt a little overwhelmed afterward.

This week there were Google Scholar notifications for two papers I found interesting. The first is from the NIST and German groups demonstrating an [all-optical operation of SNSPDs](#). They bias the detector using a photodiode and readout counts using an electro-optically modulated Michelson interferometer. This greatly reduces the thermal load of large arrays of SNSPDs; optical fibers have less heat load than RF cables. However, this does not completely solve the scaling problem. There is still a need for 3N fibers for N detectors as well as the added footprint of the photodiode and interferometer.

The second is a paper from Yabuno and the rest of the Japan group on [micrometer-wide detectors patterned with photolithography](#). The detector was not saturated at 1550nm but was nearly saturated at 850nm. This work is something we have been attempting to do here at MIT for the past few years. Attempting is perhaps a generous term. I performed a single experiment during my master's thesis which failed due to design, fabrication, or materials... or all three. More specifically, the detector latched and only exhibited a few counts under UV illumination. If we are to continue this work I believe our collaboration with Lincon Lab is the best direction to take. Demonstrating SNSPDs can be fabricated using photolithography is more beneficial in the long term if there is a foundry that runs our process for us.

In the splitter tree paper, one of the main findings is an R-L-R interface circuit used between stages of nTrons. This was actually presented in an earlier paper of his—the encoder paper I think—to separate the two-stage amplifier. The design and methods presented here are valuable for our future designs.

Aside from the lithography, the fabrication process is almost identical to ours. One improvement I would say is the use of photolithography for wiring between stages. I currently use ebeam for because all of the patterns are contained within a single field and write times are pretty short. It might be beneficial for future throughput to write all but the smallest of features with photolithography. They do all lift-off processes with photolithography.

They tested a 1-to-4 fanout circuit. Why not do a larger circuit? I would have to assume it is limited by their measurement setup, an AWG, a four-channel scope, and a counter. This is a limitation we have been working to surpass with the custom analog FPGA board.

Daily Blog 3/6/23

March 6, 2023

Categories: General

Back from my trip out west. I had a great time skiing but I am happy to be back at work. And with no breaks on the horizon for the next few months I am excited to get to work. Although, the first day back isn't always the most productive.

Today's main goal was to prepare two wafers for MLA writing tomorrow night. While at the AJA, I spent some time trying to set up a link between the AJA laptop and the network. Creating this link will allow for better logging of deposition parameters. I think this is valuable because our material growth is vital to high-performance devices. NbN growth has historically been a competitive advantage for the group, but recent literature suggests we might be falling behind. The laptop is running Windows XP, making it less straightforward than I had hoped. Additionally, there are some complications because the pressure controller is connected via ethernet. I tried, and failed, to map a network drive; it would not accept the password. I was spending too much time debugging and moved on to deposition.

I performed a test deposition on a 10x10mm die to calibrate the deposition rate and later use it for T_c debugging in the probe station. The deposition parameters for the two wafers were kept constant, except for the temperature which was RT and 800C for the two wafers. The resistivity of each wafer was 235 and 184 $\mu\Omega\text{-cm}$. Both were 24nm thick.

Some other cleanup notes. While performing the depositions I remapped the water level sensor ADC and updated some Grafana settings. We are planning on opening the main chamber to the AJA on Friday to remove the pure Nb target and install an aluminum target for Matt.



Daily Blog 3/7/23


March 7, 2023

Categories: General




This morning I spent my time installing the third probe arm before heading to the cleanroom.

The last time I installed a temperature sensor I was unaware that the bellow could be removed from the probe arm. After reading the manual—and learning the bellow should be removed—installation and thermalization to the probe arm were more manageable. I modified one of the feedthrough blanks to accommodate a hermetically sealed SMA feedthrough.



The temperature sensor was wound to a bobbin (yellow zip tie) that was mounted to the radiation shield. It might be better to move this to the sample stage in the future, but I used the radiation shield out of convenience. The sensor was mounted directly to the sample stage (left of the sample) with indium foil for improved thermal contact.



The probe station was closed and the pump turned on. I am a bit worried about vacuum leaks due to the additional o-rings. After 15 minutes the pressure is 2.5×10^{-4} Torr. Good so far.

In the clean room, I patterned two 4-inch wafers using the MLA 150 with optical autofocus. The two wafers had 24nm of NbN deposited on them, SPG805 & SPG806. There are two different pad layouts to be patterned. The first, WSP-A, was

patterned on SPG806 using nLof2035. While the second, WSP-B, was patterned on SPG805. WSP-B contained pads for single-layer device testing, mainly nTrons and SNSPDs.

Across the entire wafer, there was some change in the quality of the gold liftoff. Some dies looked better than others. For the test structure dies, the calibration lines were completely removed (no gold) and others had lines down to approximately 2 μ m. For a negative tone process, I believe this is due to the exposed resist not being completely removed during development. Other structures—like alignment marks—were completely removed, likely for the same reason.

Another possibility could be the exposure. Dennis and Scott were troubleshooting the MLA right before my reservation. The development looked good in the optical microscope. During the liftoff in NMP I observed 90% of the gold removed and floating above the wafer. There were a few blank dies on the wafer that had gold remaining but were close to coming off. Because these dies were empty, and it was after midnight, I decided to take the sample out of the heat bath and leave it to cool in NMP overnight. The next day (3/8), the gold was still stuck to the blank dies. Some, but not all, of the gold lifted off with solvent spray. I mechanically removed the remaining gold with a swab and cleaned it with solvents. Not ideal, but overall, it is sufficient to proceed.

Daily Blog 3/9/23

March 9, 2023


Categories: General




Now that the pads of both wafers successfully lifted off, I spent today preparing the files for ebeam exposure. For a wafer, the process of setting up a write file is a bit more complicated than a single die. The complications arise from how I construct the GDS, how Beamer compiles the CAR file, and how ELMS constructs the schedule file. Let's break that down.

On the wafers I am running now, there are a few unique dies that I tile across the wafer. Tile implies grid, but actually, I want to place the dies at specific locations. For test structure dies, containing calibration and alignment structures, I want them placed around the edges and center of the wafer so uniformity data can be determined. For device dies, I want them arranged in the center of the wafer where I expect uniformity to be the best. Also, I try to avoid groups of the same die so that local errors in fabrication don't completely kill the yield. These geographical constraints make dose matrices difficult. Dose matrices are easily set up in ELMS but are limited to fixed rectangular spacing. Furthermore, I would like the dose to vary within a die, meaning each subcircuit would have to have its own car file. Meaning I will have to set up a dose matrix in Beamer. As I am typing this I am realizing how convoluted this all is (there must be a better way) but this is the best I have currently.

To change the dose in Beamer you need to use the Field Dose Assignment (FDA) block. It is probably easier to explain the whole process in person so feel free to ask anytime. The process flow from Beamer is shown below. There is a top group of 4 blocks that import and extract the layer to be written on the wafer. The split then selects the specific die, three different kinds on this wafer. The loop selects the die subcircuit and sets the dose incrementally, merging the output together. An output of one of the three branches (test structures) is shown below colored by the relative dose. With this method, you have greater freedom to assign the dose.



Why am I breaking things from a full GDS, into smaller pieces, and back into a full layout for the Elionix? The answer is how alignment is set in the ELMS software. Since I want to perform local alignment at each 10x10mm die I need the CAR files to contain a single die. That way the relative local alignment settings are easily defined. Again, if this sounds confusing, it is because it is. Please let me know if you have come up with a better solution. The full exposure schedule file is shown below. Note the global and local alignment marks as well as the colored dose gradients.




The center die had a local alignment error and I continued exposure with no alignment. The total write time for 13 (out of a possible 45 that can fit on a 4" wafer) dies was just under 3 hours.

Daily Blog 3/10/23

March 10, 2023

Categories: General



APS March meeting was this week. The topic gaining the most attention is that of near ambient superconductivity, but not exactly for the right reasons. The driving force behind all this interest is, of course, controversy. More specifically, the authors **report** on reaching near ambient superconductivity in **N-doped lutetium hydride** with a T_c of 294K at just 10kbar—an amazing result—but have faced compelling accusations of data manipulation in past works. Any of the linked articles (**NYT**, **Science**, **Quanta**) will be a far better explanation than anything I could regurgitate here, so I'll leave that reading up to you. But, perhaps even better than that, is this virtual talk by **James Hamlin**, a professor of physics at the University of Florida. He was a former colleague of one of the authors and was even a co-author of one of the retracted papers. He states in his talk that being an author on a retracted paper motivated him to look closely at the difference between the published and raw data of past works by the authors. What Hamlin presented amazed me. You should watch for yourself to see how he unravels each layer of data deception. For example, when asking for raw data, he received 100 pages of scanned numerical tables. Forcing him to write a script that performed optical character recognition on each page before showing the exact manipulations used to produce the published data. Again, I'll let you watch for yourself if you are actually interested.


The slide below is from Hamlin's presentation. The plots show the steps taken to convert the raw data into the published data for an older paper by the group in question. These steps were determined by his own investigation and not part of any disclosure by the authors.

Re-analysis of Eu data

I obtain the original data including:

- 65 raw data files
- An OriginLab file used to create the susceptibility plot
- Lab notebook scans

Problems very quickly become apparent.



It is pretty clear that the data was fabricated but it is unclear what will happen to future papers by the same group. It brings up an old question about data availability and best practices within our own group. *What is our process for data availability?* I bet the answer is in the QNN student guide (I read it I promise) but Hamlin makes a few suggestions that we might consider adopting. The complete list is one of the last slides in the talk but the two I found relevant to the group were: a fully documented script showing all modifications made to the raw data for publication (e.g. Matlab), and having a reviewer whose only responsibility is to confirm that the data analysis is well documented. To reiterate, they are confirming that the analysis process is documented—not that the analysis is correct—just that it is documented.

Daily Blog 3/13/23

March 13, 2023

Categories: General



Over the weekend I completed the conditioning of the AJA main chamber after target replacement. To perform the target replacement we first set the pressure controller to local and closed the automatic gate, isolating the cryopump. Then, the main chamber and load lock were vented. Target replacement was performed (check [wiki](#) for details) and the chamber was pumped down via the load lock. Once the main chamber was at $\sim 1\text{e-}6$ Torr (ideally it is the same as the pressure on the other side of the gate, which is $\sim 1\text{e-}8$) the automatic shutter was opened and the main chamber was pumped with both pumps. Heat straps were turned on (low) for ~ 24 hours, then turned off. The chamber continued to pump for a day and the base pressure was $8.6\text{e-}9$ Torr. Pretty good.

Titanium was sputtered for 10 minutes. During this deposition, I was reading the manual of the QCM thickness monitor (which can be found at S:\SC\Documentation\Equipment\aja\qcm) to try and get it working. I didn't get very far..




The shutter attached to the QCM was, apparently (I don't remember fully) installed by Ilya. The shutter is actually the same as those on the target guns and looks to be an actual AJA mounting system and not something Ilya fashioned. Perhaps that's needed for the QCM—giving a larger area to deflect the arm—but I don't know much about how this works. I should read more. Regardless, the shutter does not fully cover the glass backing behind the chuck and a larger shutter with identical mounting holes should be cut. Assuming that the size of the shutter does not change QCM operation. The standard NbN recipe was sputtered for 10 minutes. The base pressure this morning was $7.6\text{e-}9$ Torr.

Daily Blog 3/14/23

March 14, 2023

Categories: General




I have had an equal number of triumphs and failures working in the clean room the past two days—c'est la vie. But like all failures, there is plenty to learn from.

The biggest failure was over-etching the device layer on wafer B and likely thinning the devices past an operable point. I have not tested them, but all of the nTron gates are disconnected and the SNSPDs have a rough appearance. The etch time was determined by first performing a deposition, measuring the thickness, etching, and measuring again to calculate the rate. Using this value ($\sim 3\text{nm/min}$) I determined I needed approximately 8 minutes of etching to eat through a 25nm film. That number sounded high, but I had no reason to doubt it given the thickness **measurements**. Feeling a little hesitant, I decided to etch for half the calculated time, then check the resistance, and perform the second half. At the halfway mark I measured open resistance and began to worry. That bad feeling stuck with me while trying to figure out what went wrong. The easy thing would be to blame the etcher but I have not observed significant variations in the etch rate, assuming the plasma sparked and the process was the same, i.e. no drift over time. That leaves the ellipsometer measurement/model. I have not investigated the model accuracy since but I did check using our Filmsense and found a faster rate that would explain the over etch. Lesson learned.

The second failure was due to poor documentation and bad intuition, but thankfully with less negative impact. I performed a dose matrix on the test wafer and varied the dose in the y-direction with the highest dose at the top. Unfortunately, I flipped the order in the Beamer for loop, leaving me with a trend that suggested a higher dose led to smaller features. I should have picked up on this mistake right away, as a larger dose should return larger features for a negative resist. I went back and corrected the **plot**.

Now the good news. This run did not have any issues of redeposition during either the liftoff or strip steps and contamination was near zero. Additionally, I performed the device layer exposure on wafer A, developed and etched. After half of the etch time, I measured a non-zero resistance, performed the second half of the etch and finally measured open. Perfect.

All of the images below were taken on our old NSL SEM which is now in the basement of Nano. Also, the resist has not been stripped yet.

10 μm 

EHT = 3.00 kV


WD = 4.8 mm

Signal A = InLens

Mag = 2.83 K X


Date :14 Mar 2023

Stage at T = 0.0 °




103.9nm

EHT = 3.00 kV
WD = 4.8 mmSignal A = InLens
Mag = 124.50 K XDate :14 Mar 2023
Stage at T = 0.0 °



100 nm



EHT = 3.00 kV

WD = 4.8 mm


Signal A = InLens

Mag = 298.16 K X

Date :14 Mar 2023

Stage at T = 0.0 °





Daily Blog 3/15/23

March 15, 2023

Categories: General

The group is currently in proposal writing mode. As a younger member of the group I never liked these meetings/months because they usually took time away from lab work. As an older group member, I am starting to enjoy the process, having more time to think about big challenges and future directions. Below are some notes and a rough draft of the introduction.

- Defining industry needs, requirements, restrictions, and other factors
- Identifying new capabilities needed at DOE user facilities
- Training a diverse next generation workforce
- Communicating discoveries to industry, including new and existing companies
- Catalyzing spin-off companies and supporting entrepreneurs
- Leveraging and building regional expertise and resources
- Providing impact to underserved communities
- Supporting principles of environmental justice
- Establishing specific metrics and assessments to ensure success

Proposal requirements

1. Identify key gaps that need to be overcome at the early stages of the research to facilitate transitioning to later stage activities.
2. Incorporate two or more of the four key technology areas for accelerating innovations (in table above).
3. Integrate additional elements to accelerate innovation (such as the examples listed in the bulleted list above).

Large-scale superconducting nanowire detectors and circuits would benefit a number of applications of interest to the U.S. Department of Energy. Such applications range from single photon and charged particle detection over a wide range of energies, from 1~eV (infra-red photon energy) to 100~GeV (ultra-relativistic proton energy), to low-power computation and data storage. However, the development of these technologies, to the point of industrial commercialization, has been limited by a lack of well-defined fabrication and performance metrics. This proposal will outline the basic R&D needed to create a large-scale fabrication process.

The development of superconducting nanowire technology has led to advances in single photon counting and low-level logic circuits. However, only one of these areas has overcome the "valley of death", transitioning from basic R&D to commercial deployment. Today, many companies produce high-efficiency single-photon detectors based on nanowires (although only on the scale of 10s of detectors) for applications such as LIDAR, quantum information science, and astronomy. This commercialization is made possible by the large number of small-scale applications, where only a few detectors are needed. The other area, that of nanowire-based (nano-cryotron) logic circuits, has not matured to commercial development at any scale.

Limits on the commercialization of nano-cryotrons arise from the low utility of small-scale circuits and the challenges related to developing high-utility large-scale circuits. However, the demand for large-scale nano-cryotron circuits has drastically increased in the past five years. One area would be the realization of large arrays of superconducting nanowire detectors which are currently limited to approximately 1000 detectors due to the thermal load of individually addressing each pixel. Larger arrays could be made possible by the heterogeneous integration of nano-cryotron logic circuits to reduce the number of cryogenic cables. These arrays would find immediate applications in particle accelerators where the high radiation environment limits the use of other detector technologies. Another area would be low-power hardware for classical computation or AI/ML applications.


The size of nano-cryotron circuits has steadily grown but has been limited to circuits containing just 10s of devices. Larger circuits have been limited by underdeveloped fabrication processes and poorly-defined measurement protocols. Unlike single flux quantum circuits, which have been scaled to XXX in a highly controlled CMOS foundry process and benefited from a number of design and measurement tools, nano-cryotron fabrication has relied on university cleanrooms and low throughput cryogenic testing. Fabrication of nano-cryotron circuits requires electron beam lithography for achieving critical dimensions less than 50nm. Past attempts at patterning these fine features at the wafer scale have failed due to electron beam stability or process uniformity. Thankfully, advances in commercially available electron beam lithography and other state-of-the-art fabrication tools, available at MIT.nano, have made high-volume production of nano-cryotron circuits

possible. However, the production at the wafer scale is only useful if it is possible to characterize the process at a similar scale. Current cryogenic testing is limited by the number of cables installed on dry or wet measurement systems. Where each device is wire bonded to an RF port and requires the entire system to be warmed up and opened before bonding to another device on the chip. This thermal cycling and breaking of vacuum are time-consuming, and potentially influential on device performance, making this process impractical for high-throughput testing. Fortunately, we plan to use a close-cycle cryogenic probe station (purchased through FIXME DOE Microelectronics) capable of testing hundreds (if not thousands) of devices within a single cooldown. The combination of high-volume fabrication and measurement makes us ideally positioned to accelerate the development of large-scale superconducting nanowire detectors and circuits.

Daily Blog 3/17/23

March 17, 2023


Categories: General



One of the remaining issues with the fabrication process is the removal of nLOF resist after etching into the oxide layer. We typically call this issue "resist burning" and we can reduce this effect by shortening etch times. However, the resist isn't actually burning, especially when considering that the system has helium cooling on the back side of the wafer, but actually surface polymerization. Still, I'll refer to this problem as burning. I did some further research on the formation of [fluorocarbon thin films on the surface of the resist](#).

An example of the burnt resist is shown below. The large chunks are actually areas of the resist that have lifted off and then redeposited on the wafer. I am confident in saying this is redeposition because the level of contamination is highest towards the top of the wafer, away from where I grabbed it with tweezers. Meaning burnt chunks floating in the resist would stick to the top as I was lifting the wafer out of NMP. I also observed that the resist was more likely to stick to gold rather than the oxide covering the device layer. A silver lining to the redeposition problem as the gold area is less critical. Sonication did not help with this problem but ashing potentially would. The edges of the patterned resist were more likely to have withstood the stripping process, not a problem of redeposition but actually related to the polymerization of the resist surface.

I first observed this problem on the test wafer. I thought a higher temperature and longer duration strip would solve the problem. I adjusted the process from 60°C for 1 hour to 75°C for more than 4 hours. As discussed, this did not solve the problem. Dennis recommended using a positive tone process for etching. His explanation is that this resist will still undergo some fluorocarbon formation but because this only occurs on the surface; the bottom portion of the resist should still be removed easily. I am doubtful this will work, the same reasoning should be valid for the current process but we still see the contamination. Other than the difference in tone, the positive tone resist is about twice as thick.



I talked to Matt about these problems and he had some good suggestions. In order of things to try: ashing, putting the wafer in developer again, or performing heated acetone strip.

Daily Blog 3/22/23

March 22, 2023
Categories: General



I haven't posted a blog in a few days. I have mainly been working in the clean room and taking notes on my process tracking doc. One thing I have been thinking about is the number of processes still needing development. These are all processes I have run in the past with various degrees of success. A basic stack-up of each process is shown below.

Microwave	
20 nm	NbN
50 nm	Gold
10-20 nm	HfO2
20 nm	NbN
100 nm	SiO2
50 nm	Gold

Memory	
20 nm	NbN
50 nm	Gold
100 nm	SiO2
50 nm	Gold


Multiplexer	
20 nm	NbN
50 nm	Gold
100 nm	SiO2
20 nm	NbN

Resistor	
20 nm	NbN
50 nm	Gold
5nm	MoSi
20 nm	NbN
100 nm	SiO2
50 nm	Gold

Currently, my focus has been on the memory process—because it has been fully simulated and demonstrated previously. The other processes have steps that have not been developed. For example, an integrated resistor process using a lower TiC film like MoSi would allow us to integrate bias resistors on chip. To do this, we need to ensure that our ebeam resist is completely stripped from the NbN layer so that the following MoSi layer has good contact. Second, the etch needs to be well-calibrated to avoid over-etching into the NbN film. It might make more sense to reorder this stack-up. First pattern the MoSi using the MLA—with alignment—then deposit the NbN and pattern, then add the gold contacts. Again, this will need more thought and development.

Other processes require less thought but need more design and simulation work. For the microwave process, I'd like to use ALD HfO2 for the dielectric spacer which should remove an etch step, as there is no need to open sections of this oxide. I imagine it would be easy to probe through this oxide and contact the ground plane.


There were a total of 34 alignments for this wafer write. The outliers ($>20\text{nm}$) underwent a second round of searching.



Daily Blog 3/23/23



March 23, 2023

Categories: General



Prior to performing the final exposure, I wanted to prepare other wafers for gold evaporation. Doing so requires me to produce a layout, but I have not had time to prepare enough designs to fill a wafer. My compromise was to prepare standard pad layouts and determine the nanowire circuit later. I am betting on the fact that there are enough single-layer designs that I (or others) would want to be fabricated. I am producing two wafers with 20nm of NbN and 5nm Ti +50nm of gold on top. I will share the gds and python scripts here so others can quickly create layouts.

The scripts should create a grid of subcircuits. Each subcircuit has ports surrounding an area equal to the field size. These ports can be referenced within a device, `D.references.ports`, and the center of the device reference is the center of that given field, `D.references.center`. Each function has an example of how to place your circuit and connect to that field's port, i.e. routing. It also shows how you would create a sweep of a geometric parameter. An example of the output is below.



If you have a design the easiest thing is to send me the same as the gds below but with each subcircuit filled in. Any way of doing that is fine, doesn't have to be like the example at all.

I've Been behind on the blogs lately.

[3_port_standard.gds](#)


[9_port_standard.gds](#)

[Code](#)

Daily Blog 3/27/23


March 27, 2023

Categories: General




The first chip from the “working” wafer is in the probe station and at base temperature. Hopefully, this will help generate daily blogs over the next few weeks with lots of data to present. Before we can really crank up the measurement throughput, we need to write a new driver for the scope we have on loan from the [Rogers lab](#). I don’t think there is a website yet but there is a Slack workspace where Steven Nagle posts regularly. He is super helpful and thinking a lot about how to grow the space, I highly recommend talking to him (also to get added to the Slack channel) if you have a chance.

Below is a classic IV sweep of an SNSPD coupled to an hTron loop. This leads to two retrapping currents, the first being the loop and the second being the detector. This kind of measurement is how we have characterized single devices in the group but is too slow for the large datasets we plan on building. This was taken simply to verify that the fabrication and thermal coupling is sufficient for further testing. So far so good.



Another quick check of the setup is this R-T curve which indicates a $T_c \approx 8K$.



I need to spend a lot of time writing code.

Daily Blog 3/28/23

March 28, 2023

Categories: General

Another day in the lab. Progress is slow. In part due to the setup and partially due to the device under test.

When recording an IV sweep using the classic SRS+meter setup I did not see any dependence between the switching current and heater bias, but when using the scope I saw a clear and controllable variation of the switching current using the heater. I saw three clear transitions: the nanowire, the loop, and the routing trace (presumably). By applying the heater bias, I would expect the largest effect would be on the second transition, i.e. the loop. Contrary to my expectation, the first transition is the only one affected by the heater bias.

The last time this type of device was tested the counts out of the SNSPD were visible when the heater was enabled. Thus, the bias current was entirely passed through the detector and when fired it was shunted by the loop resistance. This was not the intended use—where we would store the number of counts as flux in the loop to be read out later— but was a good debugging point. I tried to replicate this measurement by biasing the detector and heater and looking for counts. There were pulses resembling photon counts but much longer falling edges ($>1\mu\text{s}$). I believe this difficulty is due to the primitive setup and not due to the device (maybe operating temperature).

Regardless of the DUT, there are a number of important fabrication yield parameters to be extracted from this chip. I'll list them below.

- Heater resistance (I measured a few all within a few ohms of each other)
- Critical current device variation (due to fabrication)
- Critical current spread (due to electrical or thermal noise)
- Critical current vs heater bias vs temperature (a plot showing the effective temperature of the applied heat)

All of these need to be determined in a high-throughput manner, as there are 49 devices on this die.

Daily Blog 3/29/23

March 29, 2023

Categories: General



Measure twice, cut once. A phrase we have all probably heard emphasizing the importance of planning before action, and it sums up the last two days at the probe station. I'd like to just dive in and start getting as much data as possible but I have made that mistake in the past—generating incomplete or inconsistent datasets.

Thinking about fast statistical measurements, the scope I borrowed from Roger's lab seems like an ideal piece of equipment, thanks to the built-in AWGs. I spent a few hours writing and debugging a python driver for the scope and was getting annoyed by how long it was taking me. So I asked ChatGPT to do it for me. It did a surprisingly (or unsurprisingly now that everything it does is surprising) good job. I did not bother to test it but the SCPI commands all looked correct and the syntax was much better than the garbage I was writing. The free version would not generate a full driver, only an example, it may be worth buying in the future.

After writing all the basic functions I thought I would need for the measurement I started to look at the scope's built-in measurement functions. Ideally, we would perform some time delay measurements and pull the statistical data from the scope. This is much faster than pulling each waveform and performing the analysis on a PC. Unfortunately, the scope has a limited number of measurement functions compared to the ones available on the Lecroy scopes. For example, the delay measurement does not have options for setting where it measures the delay; the only options were slope and edge detection. With any of the available settings, I noticed the measurement would trigger on the noise and not the edge of interest. So after spending all that time writing a driver for that scope I decided to go back to using the Lecroy.


Below is a typical voltage trace of the nanowire under a ramp bias. There is a skew measurement tracking the time between the clock signal of the ramp input and the first voltage edge of the nanowire output. In the second panel, the timing information is plotted as a trend line and a histogram. The trend shows a clear periodic behavior which is synced exactly to the thermal cycle of the cold head. This certainly is going to make analysis more difficult but it should be possible to filter out.



Daily Blog 3/30/23


March 30, 2023

Categories: General




Some quick literature housekeeping. Nature electronics recently published a review of **cryogenic memories** that will certainly be useful when writing the nMem paper (they don't cite it). I also came across a theoretical paper on a **Non-local superconducting single-photon detector** using a double loop with SNS junctions. I didn't read it yet but it seemed interesting.


Today I got to a good point with the measurement setup and started recording some of the device behavior. First, different bias currents are applied to the device to capture the three switching events: the detector, the loop, and the routing line.



Second, I plot the recorded switching currents for each of the 1000 traces captured and plot them sequentially. The periodic behavior is assumed to be thermal—it aligns with the pulsing noise from the cold head. I forgot to record the time stamp of each sample so the actual frequency is not known. Just under 2 μA bounce in I_c .





I thought that by heating the stage the thermal noise amplitude would be decreased. However, it appeared to increase until near T_c where I believe the rising edge was significantly rounded. The slight downward trend from the first sample to the thousandth sample is due to the continued heating of the stage during the capture period. I could have used longer hold times to allow the temperature to stabilize but I prioritized speed over accuracy.



The hold time was 30s and the steps are clearly shown in the logged data.







Daily Blog 4/3/23

April 3, 2023


Categories: General



I captured the plot below late last Friday and went back over the weekend to collect more data. When I got settled in I noticed I could no longer reproduce the plot, and it appeared that the heater no longer worked. This raised some concerns, did I over bias the heater? did I discharge some built-up static potential? I went and checked some other devices and found they too were not functioning. I tried changing the room temperature circuit because this was the biggest change between measurements, but the problem persisted. I eventually determined that the low noise voltage source (SRS batteries) had failed– the batteries need to be serviced. Once I changed the voltage source to the AWG I was able to reproduce the expected behavior. phew...




The high current traces do not show the thermal bounce. That is due to the trigger on the slower-rising edge, the thermal bounce is still there. (Sorry for the change in colors, this is the same measurement)




Constant biasing of the detector is difficult due to the thermal bounce. Biasing below the 2 μA ripple might be too far from the critical current to observe a significant dark count rate, while too high might lead to latching. Ideally, this will be resolved with some thermal damping to the sample, but for this cooldown and early investigation, I am using a very slow sinusoidal bias. This allows the detector to reset.

In the plot below, I am capturing what I believe is the detector firing for different bias currents, i.e. peak amplitude of the aforementioned sinusoidal signal. Unlike the switching current characterization (the first two plots), the scope is triggered on the RF signal out of the device using the standard amplifier chain readout circuit with a constant voltage threshold ($\sim 30\text{mV}$). At low-bias currents, there are no pulses. Between 176 and 197 μA there is a train of evenly spaced pulses with decreasing amplitude. The decreasing amplitude might be explained by the storage of circulating current in the loop which then negates a fraction of the bias current. The spacing would ideally represent the dead time of the detector, however, it is too consistent and is more likely some relaxation oscillation between the detector and the hTron. The current bias of the hTron heater was kept at a constant 150 μA .




The plot below is for a bias current of 180 μA and the heater bias is swept from 0 to 300 μA . The plot clearly shows the region (150-190) where an appropriate amount of current is directed from the loop to the detector and pulses are observed.



Daily Blog 4/4/23


April 4, 2023

Categories: General



This morning I came across a new [paper](#) from the MQO group at Paderborn where they demonstrated SPDC in PPLN at cryo. I found it to be quite relevant to the QUIC TAQs project and I learned a lot about how they characterize quantum light sources. This also led me to find a [paper](#) that was scooped (we sort of stopped working on it) from us on optically written SNSPDs. Their results are not that impressive but it proves the point.


The measurements performed over the past few days hope to pinpoint the optimal bias conditions for photon counting. While the device is too thick to count low-energy photons, I have been keen on observing some sensitivity to light, for device characterization but also for characterizing the probe station. The probe station currently does not have an optical fiber installed and the only light that can be controlled is from the microscope. I was able to see some change in the spiking behavior as a function of this light power. This is likely just the chip warming up but I figured it would be good to capture.



Daily Blog 4/5/23


April 6, 2023

Categories: General




The thermal bounce is looking to be the main problem with measuring the dark count rate—or anything with a constant bias. Before warming up I thought it would be good to capture some of the pulsing behavior with some averaging. In each of the plots below the trace is averaged over 100 sweeps. There are three different sweeps of the bias current through the device which is indicated by the color and legend. The heater current is constant for each of the plots and is indicated in the title: $0\mu\text{A}$, $100\mu\text{A}$, $150\mu\text{A}$.


For zero heater bias current the lowest detector bias current (red) does not have any pulsing behavior; the red spike at time zero is noise and could be removed with a higher threshold, but then no traces would be captured. As the bias current increases, more current passes through the detector, and pulses start to appear at $\sim 190\mu\text{A}$. There appear to be two peaks in the first pulse. The amplitude of the trailing pulses decreases until no longer visible. Interestingly, the highest bias current produces the lowest amplitude train of pulses (blue). One explanation is that for higher bias currents the amount of current being added to the loop is greater. Thus, the circulating current which counteracts the bias current effectively reduces the detector bias current resulting in a lower amplitude voltage pulse.



With $100\mu\text{A}$ heater current—which is barely sufficient to switch the loop—the pulses look similar but appear 1-2 increments earlier, say $188.75\mu\text{A}$.



At 150 μ A bias of the heater current, we see pulses at the onset of the bias sweep and the dual peak is even more evident (red). As the bias current increases we no longer see pulses, meaning the entire device has switched and we are just triggering on the noise (blue).



Daily Blog 4/10/23

April 10, 2023

Categories: General





Fig. 1. Schematic diagram of the detection of secondary electrons using a TES array. Secondary electrons are generated when the ion beam is irradiated on samples which simulate living tissues. Those electrons will eventually irradiate on the TES array for the detection.


The group has been thinking about an experiment that uses SNSPDs for electron detection. The idea, which would engage both sides of the group, has floated around for so long that it has all but become a QNN trope. One reason there has not been much progress in this area is the complexity of the experimental apparatus. Do we install a cold stage in an SEM? Or do we install an electron source in a cryostat? Either option will require significant mechanical design (not impossible, just time-consuming) and neither option has a clear advantage; more design work is needed. Another reason for stagnation is a clear motivation—not to mention funding. Some motivations discussed were probing SNSPD physics (i.e. vortex motion & formation), low-energy STEM detectors with high timing resolution, or EELS with high spatial resolution (ideally sub 100nm pixels). I do not recall the broader applications of these experimental demonstrations (maybe talk to Karl, Donnie, John, or Stewart who were also in this meeting) but a big-picture application is likely needed for a large proposal.

A superconducting electron detector with a big-picture application was recently published in [IEEE TAS](#). In this work, a superconducting detector (TES) is used to detect secondary electrons generated by carbon ion beams in tissue. Carbon ion beams are being used to locally irradiate cancer cells, killing them and curing the patient. However, the ion beam generates secondary particles such as electrons that can damage DNA and kill other cells in the tissue. They specify an energy resolution of 1eV or better is needed, which is possible with SNSPDs. Personally, I don't think this is a good paper, with poor writing and Excel/Sheets figures, but might provide some direction (and references) for proposal writing in the future.


On Friday, I warmed up the probe station, applied some Apiezon thermal grease, and cooled it back to the base temperature. I was hoping that this will further couple the sample to the 3K stage and reduce the variation in I_c . Val suggested that this might be due to vibration and I partially agree with him. I don't believe that this is directly due to the vibration of the sample or probe but some thermal swing that is caused by the probing strength. It is quite easy to see that stronger probing (i.e. really pushing the probe into the pad) reduces the critical current by warming up the chip. Lightly probing the sample makes electrical contact with minimal heat transfer. The plot below shows the critical current of the same device with and without thermal grease; the orange trace is with the thermal grease. The right subplot shows the critical current increased from 120 μ A to just under 400 μ A.




I was hoping that the thermal bounce would be further damped, but relative to the switching current it appears slightly worse.





Change in critical current due to probing strength.






Under certain bias conditions, it produces a fast sinusoidal signal. I am not sure if this is the device or some other part of my setup. Adding here for reference only.



Daily Blog 4/13/23

April 13, 2023


Categories: General



No blogs in the last two days. I might argue that I was busy measuring and wanted to wait until I had a complete data set to present here, but I should have squeezed out 500 words on some other topic. Always room to improve.


The chip currently in the probe station (SPG806-C5) is an array of SNSPDs and loop hTrons, with varying loop sizes along the x-axis and doses along the y-axis. I did not intend these to be good detectors. They were fabricated on a 20nm thick NbN film that was chosen for the nMem and these devices were included just on the wafer because they fit within the same fabrication stack up. Still, I had hoped to see some dark counts but I did not. I believe this is due to the setup. I would like to try again with an optical fiber and better probe thermalization.

The other purpose these devices serve is a characterization of the fabrication yield. I measured the critical current of each device as a function of the heater current. The full dataset is below.



The error bar plot shows that all but one of the heaters are functioning. The device with the broken heater (B7) was over-biased, i.e. I broke it. That gives a yield of 98% (1mk if you get this inside joke). The large error on some of the measurements is likely just noise from the skew measurement trigger threshold. I could go through each trace and exclude points that are clearly outliers but I rather move on to something else.

The colormap below shows the average critical current for each device positioned according to the location on the die. The color bar limits are set to the minimum and maximum of the array, which is more than a 60μA spread. I believe this is not truly indicative of the device or fabrication but is solely due to the measurement. If the measurement were completely accurate, we would expect to see some spacial dependence on the I_c due to the varying dose. However, from previous measurements outlined in the past few blog posts, we know that the strength of probing, the thermal cycle, and the measurement setup can have large effects on the I_c . I would like to revisit this once some of the measurement problems have been solved. On the bright side, all of the devices work. All the data can be found on the NAS (<https://qnn-nas-rle.mit.edu:5001/>) at Data\SC\Measurements\SPG806\C5\{device name}. Time to warm up the fridge.



Now is the time to decide what exactly to do next. I will certainly load a nMem chip into the probe station and produce a similar plot with $I_c(1)/I_c(0)$, but to do that I need to write some code creating the waveforms. I previously used Brenden's MATLAB code; I vaguely recall already rewriting it in python but I have to check. Other thing is fixing the probe arm upgrades.

Daily Blog 4/14/23

April 14, 2023

Categories: General



Earlier this year I flew to Colorado to snowboard with some friends. On the flight there I decided to rewatch, "The Ballad of Buster Scruggs" a Coen brothers movie with tales of western settlers. I have always been interested in that time period, especially the gold rush. It made me think of an analogy between mining for gold and working in the lab; there is hidden treasure everywhere, it just takes a keen eye and lots of work to dig up.

Lab cleanups can be thought of as prospecting for gold. Digging through old bins, drawers, and cabinets uncovers old parts and setups that could be used for new experiments. Unfortunately, these parts usually require some work to fit our current needs and it is up to us to find parts with the lowest opportunity cost.


That brief intro describes the work I have been doing on a third probe arm for the probe station. After prospecting the lab I determined that we had enough parts lying around to construct this third arm without buying any new components (a \$10k savings). The photo below shows two probe arms with the bellows removed. The top arm has a vacuum manifold with an aluminum standoff that is press fit into a G10 arm for thermal isolation. The bottom arm has a similar standoff that I made to replace a crude version shown below that. The makeshift standoff was constructed using three bolts, some nylon locking nuts, and lots of torr seal. The torr seal might have been an attempt at further reducing the thermal coupling between the vacuum manifold (room temperature) and the G10 arm (40K). I do not know how much of an improvement this made but when I installed the arm with the makeshift standoff the probe arm was angled such that the probe would not be able to land on the chip. A new standoff needed to be made.

To make a new standoff coupling the vacuum manifold to the G10 rod I found two standoffs (placed in the drawer labeled standoffs {nice}) that would closely match the stock version. Unfortunately, these standoffs were 8-32 threaded and the vacuum manifold was M6. To join the two standoffs I found a 8-32 set screw which I inserted with Loctite into one of them. On the other end, I inserted an 8-32 to M6 external thread adapter, surprisingly we had a few of these. The mating end to the G10 rod required an internal M6 thread, which would be a strange adapter. It was simple enough to go down to the machine shop, bore out a 5mmØ hole, and tap an M6 thread. The whole process took about an hour but it was much faster and cheaper than ordering new parts. The finished part looks very similar to the stock version and holds the arm normal to the vacuum manifold.



The copper sleeve surrounding the G10 arm thermalizes the cabling to the 40K stage. For the arm that I was reworking, this copper sleeve was loose and was previously taped to the rod. I secured it in place using a two-part epoxy. With both of the arms mechanically sound, I began wiring up the temperature sensors. To do this I found the wiring schematic in the probe station manual and soldered one end of the four-wire cryo cable to the MIL spec connector. The other end of the cable was soldered to a 2x2 solder cup socket header. These headers were the kind that comes in a long row that you then snap off to the desired number of contacts. One problem with these is that the break isn't always perfect and the terminals can become loose, especially when heating to apply solder. Another problem is that the force required to disconnect these connectors sometimes causes breaks in the hair-width wires. To solve both problems I used the same two-part epoxy on the sides of the headers (not the terminals but the plastic bit) and some heat shrink. After shrinking and curing the epoxy the heat shrink gives improved strain relief to the wires and gives the user something sturdy to hold onto when connecting or disconnecting the cables. It also ensures that the heat shrink won't slide off of the connector—they make glue shrink for this reason (also for ingress protection) but we don't have any in the lab. I used this same method on the DC lines on the ICE.

The photos below show the blue heat shrink on both mating ends. The powered side should always have sockets and the unpowered side has pins. The harness was then installed and thermalized to the anchor point on the copper sleeve using GE varnish. I did not bother to wind the cable around the arm—which is usually done to improve thermalization—because the extra cable length will be wound on a bobbin and anchored radiation shield. I made sure to mark all the connectors with the correct mating orientation and verify that the sensors work before the final installation. (I prospected for more temperature sensors in the dip probe graveyard). The black electrical tape was used just to hold the wire in place while the GE varnish cured and was replaced with "cryo tape" courtesy of Mark Mondol. I don't know what makes this tape cryo-compatible, maybe a better adhesive isn't affected by thermal cycling?



Yesterday there was an AI event, "Imagination in Action" held at the MIT Samberg Center. The event had 15+ speakers all day but I was most interested in talks given by Lex Fridman, Stephen Wolfram, and Sam Altman. Turns out none of these were individual talks but more interview-style conversations that explored the connections between large language models (like GPT) and powerful computational languages (like Mathematica). Lex interviewed Stephen Wolfram for about 30 minutes and he gave interesting answers about truth and the future impact of LLMs. I'll note here just two interesting points that I remember.

The first point worth mentioning was regarding truth in large language models (LLMs) like GPT-4. As many of us know, these models just produce the next most probable token (or word) based on a given prompt. The result is a human-like response that can be extremely accurate but perhaps not objectively accurate depending on the question. We often think of mathematics and computation as something that is objective and true, but Wolfram argued that that is only because we agree on a set of axioms beforehand. If we agree that a point is an infinitesimally small region in space then we can deduce from that, and other axioms, all of the geometric truths, e.g. a triangle has three angles whose values add to 180°. His point is that truth in mathematics is based on a set of prior assumptions (axioms) that we must agree on. The same must be true for more human questions.

The second point about ChatGPT that Wolfram discussed is how LLMs are unable to really compute anything. Of course, this is where he began advertising the importance of his computation product Mathematica. Read his [blog post](#) to try out the Mathematica plugin for ChatGPT. One final thing I'll add is his discussion of how ChatGPT will change the future of

education. For example, he described a scenario where the LLM can better understand the exact point that a student is lost on. Or a scenario where a student has learned something and since forgotten and the LLM is able to quickly bring up the language needed to relink the student's memory. He envisions that most of the industrial aspects of education will be automated using LLMs and educators will have more time to focus on teaching students to think and learn. Feels like we are living in a sci-fi world.



Daily Blog 4/20/23

April 20, 2023

Categories: General

Arb Download Times GPIB / RS-232 (115 Kbps)³

Arb Length	Binary	ASCII Integer	ASCII Real
64K points	23 sec	92 sec	154 sec
16K points	6 sec	23 sec	39 sec
8K points	3 sec	12 sec	20 sec
4K points	1.5 sec	6 sec	10 sec
2K points	0.75 sec	3 sec	5 sec

You can't work unless you're comfortable

This blog was originally scheduled for 4/17 but I was delayed by some hair-pulling Python troubleshooting. This story begins with a need to program arbitrary waveforms to correctly store circulating currents in the memory loop. This was previously accomplished using a MATLAB driver written by Brenden and was integrated with his optimization code, but this driver is limited compared to qnnp and a Python version of the script would be beneficial. Adam had written a Python function that allows a user to define an array and send it to the AWG. However, this function stored the data in volatile memory which limits the user to only one channel. If the waveform could be saved to the internal memory, then both channels could be used.

Writing to the internal memory of the AWG is one of the slowest processes in the measurement and it is important to transfer the data as compactly as possible. Below is a table showing the download times for different encoding formats and waveform lengths. This formatting is the primary source of my confusion and stress.

Arb Download Times GPIB / RS-232 (115 Kbps)³

Arb Length	Binary	ASCII Integer	ASCII Real
64K points	23 sec	92 sec	154 sec
16K points	6 sec	23 sec	39 sec
8K points	3 sec	12 sec	20 sec
4K points	1.5 sec	6 sec	10 sec
2K points	0.75 sec	3 sec	5 sec

Why did this take me more than 3 days? Well, my first excuse is that I was interrupted by other responsibilities and I really needed some chunks of time to learn a new topic. The second excuse is that the documentation for PyVISA and Keysight is limited. Long story short, I simply had to chip away at each line of code and really understand what I was sending to the instrument. Some snippets of documentation are added below.

[Link to this programming guide](#)

TOC

- Keysight Trueform Series Ope
- Preliminary Information
- Operating Information
- SCPI Programming Referer
- Introduction to SCPI Lan
- Internal Function Waveform
- Alphabetical List of SCP
- ABORT
- AM Subsystem
- APPLy Subsystem
- BPSK Subsystem
- BURSt Subsystem
- CALibration Subsystem
- COMBine FEED
- DATA Subsystem
- DATA:ARB2:FORMat
- DATA:ARBitrary[1|2]
- DATA:ARBitrary[1|2]:C
- DATA:ATTRibute:AVE
- DATA:ATTRibute:CFA
- DATA:ATTRibute:POI
- DATA:ATTRibute:PTP
- DATA:SEQUence


DATA Subsystem

The DATA subsystem manages user-defined arbitrary waveforms:

- [SOURCE[1|2]:]DATA:ARBitrary2:FORMat {AABB|ABAB} - specifies the order for bytes in a dual arbitrary waveform file (requires optional IQ player).
- [SOURCE[1|2]:]DATA:ARBitrary[1|2] <arb_name>, {<binary_block>|<value>, <value>, ...} - downloads arbitrary waveform normalized values to waveform memory
- [SOURCE[1|2]:]DATA:ARBitrary[1|2]:DAC <arb_name>, {<binary_block>|<value>, <value>, ...} - downloads arbitrary waveform DAC codes to waveform memory
- [SOURCE[1|2]:]DATA:ATTRibute:AVERage? [<arb_name>] - returns arithmetic average of all data points for an arbitrary waveform or sequence
- [SOURCE[1|2]:]DATA:ATTRibute:CFActor? [<arb_name>]ATA:ATTRibute:CFActor? - returns crest factor of all data points in an arbitrary waveform or sequence
- [SOURCE[1|2]:]DATA:ATTRibute:POINts? [<arb_name>] - returns number of data points for arbitrary waveform or sequence
- [SOURCE[1|2]:]DATA:ATTRibute:PTPeak? [<arb_name>] - returns peak-to-peak value of all data points in an arbitrary waveform or sequence
- [SOURCE[1|2]:]DATA:SEQUence <block_descriptor> - combines previously loaded arbitrary waveforms into a sequence
- [SOURCE[1|2]:]DATA:VOLatile:CATalog? - returns the contents of volatile waveform memory, including arbitrary waveforms and sequences
- [SOURCE[1|2]:]DATA:VOLatile:CLEar - clears volatile waveform memory
- [SOURCE[1|2]:]DATA:VOLatile:FREE? - returns number of points available (free) in volatile memory

Using the IEEE-488.2 Binary Block Format

In the binary block format, a *block header* precedes the waveform data. The block header has the following format:



The function generator represents binary data as 16-bit integers, which are sent as two bytes. Therefore, **the total number of bytes is always twice the number of data points in the waveform** (and must always be an **even number**). For example, 32,768 bytes are required to download a waveform with 16,384 points.


Use the FORM:BORD command to select the byte order for binary transfers in block mode. If you specify FORM:BORD NORM (default), the most-significant byte (MSB) of each data point is assumed first. If you specify FORM:BORD SWAP, the least-significant byte (LSB) of each data point is assumed first. Most computers use the "swapped" byte order.

Some GitHub links that helped [A](#), [B](#)

Daily Blog 4/21/23

April 21, 2023

Categories: General



Now that the AWG debacle is over the real challenge begins; reliably addressing the memory. Before I dive into that, a number of papers published in the last week piqued my interest and deserve some mention at least. I think if I really read and understood these papers I would gain many new research insights. However, I feel oversaturated with my current research goals and will leave a deep dive to a later date.


A nice [paper](#) from Nobuyuki Yoshikawa titled, "Reconfiguration of nondestructively readable superconductor memory by direct injection of magnetic flux to storage loop", has reduced the footprint and number of JJ's compared to a previous memory cell architecture. It also outlines future applications for such a memory to be used to construct LUTs, a basic component of FPGAs.

A fifth force? I haven't worked on the dark matter project much but I enjoyed reading the introduction of this [paper](#), "Novel constraints on fifth forces and ultralight dark sector with asteroidal data".

Zhao published another [paper](#) that uses planar hTrons for addressing SNSPDs. I should read and update here.

A large [study](#) of the impact of ion (Ar, Ga) irradiation of superconducting films was published last month. I shared this with Marco Z and Francesca. I think it will be a good guide for future He ion experiments. (I have an idea for this).


A [paper](#) with wafer scale SNSPD fab data? What is the yield? Oh... 39%? I am only joking. I need to come back and update after I've read it.



Daily Blog 4/24/23


April 24, 2023

Categories: General




Reading and writing to the memory was initially so tricky that Brenden wrote an optimizer to find the optimal bias levels. I don't plan on doing that in Python so careful documentation will be needed. Let's start by defining the waveforms.

A ramp-read bias was used to find bias levels that gave the largest separation in switching current, i.e. I_{T_0} and I_{T_1} in the figure below. This allows us to see the degree of separation between the two memory states. Once the distribution of each state is known the read pulse height is set between the two distributions. A flat read pulse could be used to find the optimal level, however, it would not reveal how separated the two states are.




The waveforms can be defined with numpy arrays in Python. The amplitude is normalized and will be set to an absolute value in the AWG. The number of points used to define the waveform is kept to a minimum for simplicity. The actual number of points sent to the AWG can be defined using the updated `awg.set_arb_wf()` function which interpolates to the desired number of samples. The number of samples sets the maximum operating frequency.



Writing short waveforms ($<2^{10}$) to the AWG does not take much time. With the current script, I am rewriting the waveforms to memory on every measurement run. This is perhaps inefficient if the time to write to memory is long and the only change between measurements is an amplitude or offset setting, i.e. not requiring the internal memory to be updated. I believe it is possible — though I need to reread the manual — to write a list of waveforms to the internal memory and then output a defined sequence of waveforms from memory. In other words, we could define a PRBS using already stored internal segments.


With the waveforms stored in internal memory, it makes sense to give a brief description of the room temperature setup. Unlike the past memory work — which took place in the dewar and had SMD components at 4.2K — this setup cannot be matched to 50Ω impedance. In an attempt to avoid using the box resistors, I tried using a -3dB attenuator as a bias resistor. At DC, the attenuator has a series resistance of $\sim 15\Omega$ and a shunt impedance of $\sim 150\Omega$. The signals were then read out at the two points (indicated by colored circles below) using the scope channels set to $1M\Omega$ impedance. The colors correspond to channels C2 and C4 in the scope screenshot below.






I did not determine the series impedance of the attenuator at high frequencies and the actual bias current is currently unknown. I am ok with not knowing the actual current because this level of characterization is focused on the fabrication yield and relative separation of the memory states. The future array test will take place in the ICE with on-chip bias resistors and the exact currents will be known. The following data will be presented in substitutional units like voltage or time in replace of current.


The plots below show the output of the channel and enable ports under different bias conditions. The top two plots show the output for two different channel bias levels with zero enable bias. At 30mV, the write pulse does not switch both sides of the memory and there is no output pulse. At 40mV, the write pulse switches both sides of the memory loop causing a positive or negative (1 or 0) output pulse to be visible.




The two plots below fix the channel bias and increase the enable bias. At 20mV, the enable pulse does not trigger an output during the write operation. At 25mV, the enable bias causes the memory to switch during the write operation.





The previous four plots show some limits on the channel and enable bias levels. The rest of this blog will adjust those bias levels and capture timing information from the ramp-read. The time between the start of the ramp-read and the switching event is visible in the scope screenshot; skew measurements P1 and P2. The first input the to measurement is the enable signal C4 at the start of the read pulse and the second input is the channel signal C2. The same is true for P2, except shifted in time to the zero read. The plots below show the timing distributions for two consecutive read/write operations. The written values are , , and . Please forgive the lack of units. The top x-axis is time in microseconds and the bottom is time in seconds. The bottom y-axis is counts.





The same measurement (writing 1 then writing 0) was performed for different enable voltage levels. Histograms contain 1000 measurements.



The above data is plotted below with error bars indicating 1 standard deviation from the mean.




For an enable voltage of 40mV, a sweep of the channel bias was performed. The time for W1 and W0 cross at ~30mV, 42mV, 52mV. Above 40mV we are switching the channel during the writing step, a non-ideal operation. The correct operating region is between 25 and 30mV while the inverted region is between 32 and 40mV.



Daily Blog 4/27/23

April 27, 2023

Categories: General




Another light week for the “daily” blog. I need to keep pushing myself to write.


Yesterday was both productive and educational — perhaps a valid excuse for not publishing a blog. The day’s productive portion was centered on measuring and analyzing another memory cell. The educational portion was centered around integrated photonics; one talk was a thesis defense by Saumil Bandyopadhyay from Dirk’s group and the second was the Haus lecture by John Bowers. I have never worked on photonic circuits, but I see that field as a blueprint for the future of superconducting nanowire circuits and it is worth paying attention to the problems and challenges they face. Both talks heavily relied on the progress made in commercial foundries.

I measured the distribution of the time delay between the read enable pulse and the device switch during the read ramp for both memory states. The zero state is colored in blue and the one state is colored in orange. Each distribution contains 1000 measurement samples — other than the low bias values where the device might not have switched at all. The x-axis on the animations below shows the channel voltage and each frame displays a different enable voltage. Yesterday’s post had a device with a moderate loop size. The plots below show the two extremes ends of the loop sizes on the chip.

The largest loop memory has very little separation between states.




The smallest loop has better separation. There is a very interesting folding/wrapping/inverting of the two states as the enable voltage increases. It can be thought of as a wavefront moving from right to left as the enable voltage is increased.




It takes a keen eye to pick out the correct operating region in the animation above. To make the operating region more visible, I plotted the difference between the tails of the two histograms as a color and coordinated each point by the channel voltage and the enable voltage. A positive value indicates a working memory — including inverted operation — and a negative value indicates an inoperable memory.

The plot below more clearly shows the state inverting through the curved blue regions. The dark red square at (17, 40) is a measurement error where the measured delay time was essentially at the limit (150ns).



The large loop device almost never works. Except at high bias values where the channel voltage is high enough to switch both branches during the write operation; which is about 40mV on the channel.



It is possible that the larger loop device (C7) has an operating point outside of the bias levels shown. I should go back and investigate more.


Overall, these results are quite promising. The next measurement to perform would be an 8-bit PRBS BERT test at the most operable point. The most operable point could be first determined for each loop size, then perform the BERT, or, find the most operable point for the smallest loop and keep that point fixed as the loop size changes. Having to determine each operating point will be time-consuming, while not performing this optimization might be misleading. I'll have to think about it.

One problem with all of this data processing is the amount of time it takes to both collect and process the data — about 2 hours to perform the dual sweep of the channel and enable voltages and on the order of minutes to import all the data. I will say, it is much faster and smoother than any dataset I have previously collected.

Daily Blog 4/28/23

April 29, 2023


Categories: General



One thing I could do to improve blog writing frequency is to populate future blog drafts with paragraph topics. For example, I thought of this yesterday and I put a note in a draft for writing a paragraph about writing paragraphs (so meta). It might be obvious but writing it here will help me remember. Overall, after four months, the daily blog has been a huge personal success. There are still things to improve — like frequency, quality, or content — but the daily blog has kept me on task, improved my literature search, and made presentations much easier.

The first measurements of the nMem have been exciting but I seem to have hit a wall; I am collecting a lot of data with little gain in progress or understanding. The measurement I have been trying to perform is to measure the state separation as a function of the loop size. I started by performing the same bias voltage sweeps on two different devices. The results (which are in a previous blog post) showed that one device had an operable region and one did not. The next logical step would be to measure the devices with intermediate loop sizes and find where the operating point became non-existent. Two problems arose. The first problem was that one of the devices in the same row had a broken heater (either due to ESD or over biasing). The second problem was that I was not certain that the operable point of all the DUTs would be located within the bias sweep. Extending the range of the bias sweep would be one possible solution, but measurement time would increase. Instead of increasing the scope of the bias sweep it makes more sense to focus on each operating point. The plot to be produced would then be a contour plot showing operating rings for the various loop sizes. Before doing that. It makes sense to try and reduce the size of the files, thus decreasing the time to process the data.

Another roadblock was figuring out how to describe the “falsely” separated states when the write bias is too high. Hopefully, I’ve illustrated it well in the plot below. In the plot below the axis correspond to histograms, while the waveforms captured from the oscilloscope do not have axes and are included to explain the observed changes in the distribution. The enable voltage is higher than in previous plots due to an additional 20dB attenuator. As the enable voltage increases the two states begin to separate, visible in both the waveform and the distribution. At 179mV, a small secondary peak for the “1” state appears and at 180mV it is bigger. This second peak is due to the write pulse switching both sides of the memory. The write pulse is visible in the waveforms with $V_{en} > 179\text{mV}$. The same thing happens for the “0” state at 189mV.



Inverting readout




Figure 2-4: Read process for a non-inverting cell within an array. The basic read operation consists of the application of a read bias to the column, and selecting the row to be read by applying an enable signal to that row. If a selected cell is in the “1” state, then no voltage is seen at the top of that column. On the other hand, if the cell is in the “0” state, then the hTrons switch and a voltage is seen at the top of the column.





Figure 2-5: Read process for a non-inverting cell within an array. The process and evolution of cell state for the non-inverting memory cell is similar to that for an inverting cell (see figure 2-4). Again, to read the cell, a read bias is applied to the column, and the row to be read is selected by asserting the associated row-enable signal. If the selected cell is in the “1” state, then a voltage is seen on that column. On the other hand, if the selected cell is in the “0” state, then the cell remains superconducting and no voltage is seen at the top of the column.

Daily Blog 5/1/23

May 1, 2023


Categories: General



Back to programming the AWG. I keep getting errors when trying to define a sequence of waveforms; looking forward to the day AI does this automatically. I will update at the bottom of the blog if I make any progress. While I work on that here are some papers that caught my eye last week.

The NIST group [published](#) their work on SNSPD-integrated ion traps. The detector performance was previously limited by the RF tones used to address the trapped ions. The energy from the RF signal would couple to the SNSPDs and increase the local temperature (I think). This work shows how a RF mirror improved the detector performance. I got to meet Benedikt at WOLTE and ASC, where we talked about this work. Cool to see their results.

Here is — yet another — [paper](#) on microwave photon transduction. Though it was interesting that they used aluminum wire bond stubs as a flip-chip interconnect.




The errors from the AWG were figured out. The first problem was the byte length of the block descriptor was different from the length used in the `SOURce:DATA:ARbitrary` function. The difference is that the *name of the waveform* was not included in the block length while the *name of the sequence* was included. The second problem was that the waveform I had saved in the internal memory needed to be loaded into the volatile memory `'MMEM:LOAD:DATA "INT:\WRITE0.arb"'`.

Daily Blog 5/2/23


May 2, 2023

Categories: General




Yesterday I figured out how to output a sequence of waveforms saved to the internal memory of the AWG. This ability allowed for the creation of just three waveforms: write zero, write one, and enable. Then, these three waveforms could be stored in the internal memory and sequentially combined to create the desired bit string. This process eliminates the need to upload long waveforms to the AWG, which is a slow process. There is still the potential to improve this process. Currently, each sequence must be created and saved to the AWG for each bit string. This is still pretty fast but one further improvement would be to precompile all of the sequence strings and save them on the AWG.


Here are the three waveforms saved to the internal memory. They are similar to the ones in previous blog posts with the main difference being the flat read pulse (opposed to a ramp).




Below is the output when writing and reading the bit string 00001000. Note the waveform captures a leading and trailing bit.





Below is a plot of the bit strings for 0-7 offset by an arbitrary voltage for clarity. There is one error on the #4 trace. I am confident that this can be removed with more optimal bias conditions.



I repeated this experiment with the enable signal off. As expected there is no output. This is a good sign for array operation.




I repeated this for 8-bit strings from 0 to 64; a total of 512 write/read operations. I calculated the bit error rate in MATLAB (slow) and found there were a total of 4 errors: W0R1 3 W1R0 1. Giving a first BER of 0.78%. Actually, after some investigation, the errors were a result of the script. After adjusting some thresholds and search locations the actual error rate was 0%. Below is a figure showing the locations (red trace when high) where I am taking an average of the measured voltage. If the mean voltage is greater than 6.5mV then the state is 1, otherwise 0. This is not an ideal way to discriminate the logic levels as it filters out some noise. I will have to program a different threshold in the scope.





Daily Blog 5/3/23

May 3, 2023


Categories: General



The estimation of the bit error rate (BER) is shown in the two plots below. The distributions of the two states are fit with a Gaussian and the intercept of the two center tails is calculated. The intercept is then taken to be an estimate of the BER and is highlighted in cyan on the right plot. This process was repeated for other memory cells.





Previous searches for the most optimal bias point were difficult to process due to the large file sizes. An obvious solution would be to reduce the size of each file. To implement this, the code was rewritten to only save the bias voltages used and the distances between the center of each distribution. These three values were used to create the plot below.




The sweep above seemed to cut off some operating points in the northwest and southeast corners. Let's do a bigger sweep.

In the plot below I changed a few things. First, I formatted the data into a single .mat file. Second, I excluded — by setting to zero — any inverted or overlapping measurements. The positive values are then either: the correct operation, switching during the write, or measurement errors. The center diagonal is the correct operation and the northeast and southeast corners are switching during the write.




Just for fun, I collected 100k samples. Only took 10mins.



Daily Blog 5/5/23

May 5, 2023


Categories: General




Happy Cinco de Mayo.

The chip in the probe station has been cold for a few weeks and it is about time to warm up and move on to the next measurement. Before doing that, all the data needed to produce the plot this chip was designed to yield needs to be collected. The two plots would be the BER as a function of loop size and a PRBS BERT of the best-performing memory.


To speed up the process of locating the optimal bias point for each loop size, a stripped-down measurement script was created. The measurement script only returns the bias points and the distance between the states. The distance is only returned if the states do not overlap and are in the non-inverting mode. The plot below shows the results from 6 different memory cells. The central diagonal region is the correct operating region. Other reddish portions outside of that diagonal are either measurement errors or in the write-switching regime. There is a general downward trend in the size of the operating region with the size of the loop. Note that the axes may shift to capture the region of interest, however, the scale remains constant. The enable voltage required increases with loop size due to the resistance of the heater increasing. Thus, the current — and therefore heat — is mostly constant.





The BER as a function of loop size is not as interpretable as I had hoped but there is a general upward trend in the error rate as the loop size increases. The fact that the operating region is not consistent made the measurement more difficult. T

The left plot shows that the channel voltage needed for correct operation decreases with loop size. More accurately, you could say it decreases with increasing enable voltage, which is expected. The clearest example of the operating modes can be observed in the red line for loop size 1. Starting at low channel voltage, the memory states are slightly inverted but mostly overlapped. As the bias increases the states separate (the "1" switching current becomes lower) and the error rate decreases exponentially until it flattens. Then there is a small dip which peaks to an ultimate estimated BER of $\sim 1e-10$. The peak is the write-switching state and not an operable memory. After the peak, the memory is mostly overlapped and works in the second inverted regime.




The 10k PRBS took 3.3 hours.

Daily Blog 5/9/23


May 9, 2023

Categories: General




Colang's defense today. An end of an era. Marco really set the tone when it came to working and research. He was always down to fab. So even though there was lots of celebration today, we still need to write a blog, even if short. Let's do it.

I noticed while running a PRBS that the memory would get into modes where it was always producing errors. More specifically, there would be no output pulses, even with multiple 1's being written.




Do you also see some periodic behavior? A period of correct operation, followed by an increase in errors, and repeat. A total of three peaks. It is interesting that there are ~long periods of correct operation. For example, the time between 1000 measurements is about 20 minutes. It is also interesting to see if there are multiple errors within a word. Below shows that the majority of errors are single errors and have a maximum of 3 errors in a word.



Daily Blog 5/12/23


May 12, 2023

Categories: General



Tuesday night I reserved time on the HS50 for a long overnight write. Unfortunately, the exposure did not go well. Though, I will say, the results were exactly as I expected, which is arguably better. Still, the failed exposure weighed on me, and I struggled to see the positive outcomes.

Yesterday, I went back into the cleanroom to etch the wafer with the new process developed by Dip Joti. This was also an opportunity to monitor the etch using the probe station. This showed two things. First, that a wafer I previously thought was over-etched had devices that were still connected and could be tested. Second, the “failed” exposure was properly etched and also had devices that could be tested. Both wafers will be diced and will be ready for testing by next week.




The images above are screenshots from the ELMS exposure software. The “!” symbol represents the numerous height sensor errors (HSE) during the exposure. Normally, these height sensors are infrequent and do not have a large impact on the yield; when an error occurs the previously measured height is used, which is typically fine. The image on the right shows the height sensor output for a failed measurement. The tool will move up and down attempting to maximize the signal above the threshold. The small peak on the right — a reflection off of a non-planar surface — caused the tool to continually search for a new height. This search lasts for a few seconds before the tool triggers a HSE and uses the previously measured height. This added search time doubled the total write time.


The non-planar geometry was apparent when spinning (streaking) as well as in the SEMs below. This error came from underdeveloping the liftoff resist. My previously “optimized” development time was for a less populated wafer. The density

of this pattern requires a longer development time. Because the feature size is not critical in this step I will plan to purposefully overdevelop in the future.

The left SEM shows a gold pad that was folded over during a sonication step. The edges of this pad were evaporated onto undeveloped portions of the resist, thus not making contact with the substrate. The right SEM shows a pad that had proper contact.



Additionally, the resist was still burnt.



<https://assets.researchsquare.com/files/rs-2823538/v1/c10d59ed-b445-407c-80c0-7f8f84a10b6a.pdf?c=1683648866>

A more descriptive title – Daily Blog 5/31/23

May 31, 2023

Categories: General

We're back! The past two weeks were pretty tough. The ANL proposal is mostly done, the workshop at ASU has concluded, and the VTCMR is over.

I made up that last acronym for the Vermont City Marathon and Relay. I have been training for the marathon for about 6 months and finished in 3:39:22. Given it was the first time I've ever run a marathon, I am pretty happy with that time and it is about what I expected based on my training. The first half of the race was great and I was projected to finish under 3:30, but the hills in the second half really killed me. I had an average heart rate of 175bpm and peaked at 195bpm. Tough. I am still recovering.

I am in Arizona right now for the Breakthrough Starshot Workshop. It was a day and a half of presentations and discussions covering every aspect of the communications system. My presentation was on the progress of scaling nTron circuits. It was not the best presentation as I did not have enough time to work on the slides or rehearse. But it was an exercise in speedrunning a talk — something I'll likely have to do again in the future.

While in the airport, I've gone back and looked at some of my daily blogs and I've noticed that they are difficult to reference — especially when describing them to you all. This is for two reasons, the title and the categories. The title is just the date, which is redundant because all the blogs are dated anyway. I wanted to simply put the date in order to reduce the barrier to me writing a blog (don't think just write) and to keep me on schedule by always typing "Daily". However, it would be better to have a descriptive title so it is easier to find (i.e. more keywords). I'll use the format in this blog's title in the future and revert back if needed. The second reason, categories, is just laziness on my part. I think it is useful and beneficial to new people joining the group. I can remember clicking on each section when I was a tech, it was useful. So two easy fixes.

I'll have another blog tomorrow, hopefully with something more interesting.


Another TDGL simulation attempt – Daily Blog 6/8/23

June 8, 2023


Categories: Electronics, Simulation, Superconductivity, TDGL

I started writing this blog back in the Pheonix airport and I hoped to have some results before publishing it. However, I ran into some roadblocks and now it would be better to document these problems here.

Below is the device we are interested in simulating. A $1\mu\text{m}$ wide wire, $5\mu\text{m}$ long with a notch on one side that extends one-third of the way into the wire width, i.e. the diode.




To create the mesh, the geometry is resampled from the number of points used to define the geometry (e.g. 4 points for a square) to a large value that sets the mean edge length to be less than the maximum edge size of the mesh. Here the maximum edge length is set to $1/2$ of the order parameter (ξ), which gives a resample value of ~ 250 points. In the mesh image, the consequence of resampling is visible in the rounded corners of the device. More points will make the corners sharper but will also increase the number of mesh elements. The selection of the resample points depends on the number of malformed Voronoi cells.




Meshing with different resampling



To run the simulation, the TDGL equations are solved for a given number of timesteps for thermalization (equilibrate) before simulating the output solution. Below is an output of the sheet current density with and without streamlines. The current is measured at the cut line to verify that it is equal to the applied current $-1\mu\text{A}$ set and $1.023\mu\text{A}$ measured. For the simulations below I used a slightly different geometry (due to mesh size) and different material parameters: $T = T_c/2 \rightarrow \xi = 6.2\text{nm}$, $\lambda = 200\text{nm}$ and the phonon parameter $\gamma = 23.8$, referenced from Jason Allmaras' thesis.



Order parameter and phase for $1\mu\text{A}$, $15\mu\text{A}$, and $60\mu\text{A}$ The magnitude of the order parameter is 1 for most of the device; except for the edges where the normal-metal source and drain terminals make contact. The phase gradient gradually increases with the current.



Ok so the I_c is way higher than I thought, here's $200\mu\text{A}$, $1000\mu\text{A}$. (I should just calculate it)





Alright here we are, 3000 μA , vortex crossing.



https://qnn.mit.edu/wordpress/wp-content/uploads/2023/06/download.mp4?_=1

From the video, we can take snapshots at different timesteps. Here is the supercurrent and normal current at $\tau = 0$.







By the number of results above you might have gathered that simulation times aren't too bad. Ill try to keep them coming.

Appendix

These malformed Voronoi cells were annoying. Here are some examples




More TDGL and some Autoprobe updates – Daily Blog 6/12/23

June 12, 2023


Categories: Devices, General, Simulation, TDGL

A bit of multitasking today, here is a recap.


I performed some extended simulations of the diode and measured the voltage as a function of time. I still need to look at the units, but I am more confident it will be straightforward with this software package than any previous attempt (e.g. Comsol). This is because the rest of the material parameters are already in proper units and I just need to look up the normalization ($\$V_0\$$). Here are some snapshots for a 2000uA bias at different time steps.



The rest of my day was spent checking on the autoprobe and processing the data as it was fed out. There are a lot of measurements, upwards of 3000 I believe.



nTron source and nTron gate.




TDGL IV of diode with 30mT applied field – Daily Blog


6/16/23


June 16, 2023

Categories: General, Simulation, TDGL






Voltage measured between the probe points shown above. The plot below shows the voltage as a function of time with an applied magnetic field of 30mT. I ran this long simulation to gauge how long is needed to reach steady state. At $600\mu\text{A}$ we see the voltage dynamics corresponding to phase slips entering the wire.





There is some code to make animations directly in python, but I found it easier to create a figure for each simulation and then merge them in photoshop using this [guide](#).






TDGL IV continued – Daily Blog 6/20/23


June 20, 2023



Categories: General, Simulation, TDGL

I repeated the simulation of the diode with a finer current step and a longer simulation time.



Here is the same data with fewer traces shown, focusing on the initial phase slip.






The voltage at $I=900\ \mu\text{A}$ is 1.5mV.

nMem array testing – Daily Blog 07/12/23


July 12, 2023

Categories: General



The past two months have seen quite a drop off in the number of blogs I have published; this is visible in the chart below. The chart also shows, that even at my best, I only publish 50% of the days. Maybe I should lose the “daily” modifier, but 0.500 still gets you into the **Hall of Fame**.

number of blogs vs. Month



May and June were spent mostly writing proposals, creating posters and presentations, and doing some other experimental work with collaborators. For my own research, I attempted to measure an 8×8 array of memory cells. My first attempt was unsuccessful and left me head scratching. In the process of exchanging samples from the probe station, I decided to clean my sample of some Apiezon grease. Now that we have the sonicator from the NSL I thought to use it to thoroughly remove the grease. I wasn’t thinking that sonication will destroy the metal traces on my sample (D’oh!), but it might have. I measured open resistances at room temperature but in the SEM I did not observe any damage, or at least nothing catastrophic as I had expected to see. So that gives me some hope, and I have a second sample with 8×8 arrays in the worst case.


Semi-during that mishap, I loaded a 4×4 array from the same wafer into the probe station. Measuring the 4×4 array went better than I hoped, and I realized past measurement problems are likely just due to setup. On to the data.

Below is a 4×4 grid of plots. Each grid represents a different cell in the array. Rows 2 & 4 could not be measured, the heater measured open. To measure the memory, probes were landed at the respective coordinate and supplied a two-parameter sweep. The read ramp measurement as shown [here](#) (and [here](#)) was supplied, write a “1” then read; write a “0” then read. The time of each read operation was tagged and assembled in histograms for each bit. Gaussian fits were made to give an estimate of the ultimate BER (the measured BER is 1e-2). The estimated BER is plotted for each channel voltage and enable voltage applied. Two important **notes**: the color bars are not consistent between cells, instead min and max values; the x and y axis are consistent, meaning operation was not optimized for each cell.

On the first note, the black pixels indicate when the cell is in an inverted operation. In some cases, although indicated at unity BER, the inverted states are well separated.

On the second note, because the x and y axes are fixed, the diagonal shift of the operating margins is visibly moving up and down. This needs more investigation.

This plot does measure the fabrication yield—minus the broken rows, but that could be due to ESD. The difference in BER due to fabrication, like switching current variations, is overshadowed by thermal or vibrational effects. With the greatest contribution due to changes in probing strength.




This is 1 of 16 arrays on the chip, lots more to measure...


nMem array testing 2 – Daily Blog 7/13/23

July 13, 2023


Categories: General



On the same chip as yesterday I performed two measurements to investigate the environmental influence on the BER. Quickly, I did not notice any change in error rate due to stray photons.




In the most isolated environment (window closed, lights off, door closed, overnight measurement) I collected 100k samples to see if there was any time correlation between errors. My thinking is that there may be vibrational effects that cause multiple errors in a brief period.



I started to look around for other working arrays but before doing that, I should take some time to rewrite the measurement code. Currently, the code does not return the images that I showed in yesterday's blog. Those images were produced after processing waveforms pulled from the scope—this is time-consuming. I need a fast way to see the operating margin for each cell, while balancing the time and resolution.


The plots below show the distance between the two states in seconds. A positive distance represents normal operation and a negative distance is inverted operation. I swapped the axes so the graph resembles typical hTron behavior—a plateau at low heater bias and then a downward slope. At the knee, there is evidence of some instability, visible by a line of ± 20 ns points.



nMem array testing 3 – Daily Blog 7/14/23


July 14, 2023

Categories: General



Below shows the operating state of each cell in a 4x4 memory array. The axes are constant for all plots. Red regions represent correct operation. To save time, each pixel is the result of a single measurement, not an average.


The ideal operating point moves around quite a bit for each cell. I am working on understanding if this is just due to probing and potential ways to mitigate it.



nMem array testing 4 – Daily Blog 7/24/23


July 24, 2023

Categories: General



No blogs last week. I hit a roadblock with measuring the memory and have been struggling to find a solution. The main issue is the delay between writing and reading, i.e., preserving the memory state.


This problem arose from difficulties measuring a PRBS, which was eventually resolved by adjusting the time between a write and a read. Seeing that longer ($>500\text{ns}$) delay times no longer had measurable states, the persistent current degradation became a focus.



nMem array testing 5 – Daily Blog 7/27/23

July 27, 2023

Categories: General, nMem



Big news this week from researchers in Korea—who have [published](#) evidence of a room temperature (and ambient pressure) superconductor. Of course, this has come with much skepticism, following several [retracted papers](#) claiming the same thing. For a better scoop, see this [article](#) that Donnie shared. Overall, it's still exciting, but superconductivity upto 127°C, not exactly [room temperature](#). Oh, remember these [guys](#)?


This news led me to [Retraction Watch](#), a site that tracks retractions and other wrongdoings in academia. One thing that hasn't been widely discussed here at MIT is the recent resignation of Stanford president Marc Tessier-Lavigne—worth reading about if you haven't. The full investigation can be found [here](#), which presents evidence of data manipulation.

On Monday and Tuesday, I performed measurements of the memory array in the dewar. I keep switching measurement setups for two main reasons: having cold biasing elements (resistors, capacitors) and multiple connections (wire bonds).

The first reason, having biasing resistors and capacitors on the PCB, lets us directly replicate the biasing network used for the previous memory work. I tried this exact network in the past with little to no success. I also tried biasing with higher series resistance (5kΩ) and was again unable to have well-defined states in the memory.

The second reason was for multi-cell testing. Having the full array connected via wire bonds allows for operations where multiple cells are randomly accessed. This too was unsuccessful, but led me to think more about [state preservation](#).

So what do I do now? I have some results showing good BER for a 4×4 array. However, I have not yet shown that the memory state can be preserved for longer than half a microsecond, which makes the array inoperable. To move forward, let's go back to basics.





Above is the classic hTronn plot for device SPG806-D6D1A4. (I might have to come up with a new naming convention.) Error bars are min and max of 100 samples. I will make a plot which these curves for each of the cells.

It is time to finally address the thermal/vibrational oscillations in the probe station and put some quantitative values behind my qualitative assessments. Below shows two plots, the first a trend of each I_C measurement where each measurement was performed using the segments mode of the oscilloscope (I am assuming that this captures consecutive events), and then the time of each measurement is shifted by the period of the captured waveform. 10ms in this case.


Then I simply fit a sinusoidal to the data and extract the frequency. The frequencies fit for each heater current are shown in

the second plot. I performed some basic goodness of fit filtering, which did not plot the $\pm 500 \mu\text{A}$ values because they were less sinusoidal. This plot shows fluctuations have a frequency of approximately 2 Hz.



Lastly, here is a plot of the IV characteristics. There appear to be three steps in the IV, potentially due to the three other cells in the column switching. It would be interesting to do a measurement where two heaters are active, thus changing


how these steps appear if they indeed are the other cells.




nMem array testing 5.1 – Daily Blog 7/28/23

July 28, 2023


Categories: nMem




The plot below is similar to the one from yesterday's post but shows curves for each row in the array. This shows that cell A4 is more constricted than the other 3 cells, due to it having the smallest plateau at low heater current. For the other cells, the critical current is limited by A4, up to 250uA in the case of A3.



Here is the whole array. Pretty ugly, Ill have to make some changes.



Below shows each column as a separate subplot. Each row has the same color. The second column, B, appears to be the least constricted, as there is no plateau.



note

RC (50ohm, 82pf) filter 40MHz -3dB cutoff. If i use a 5kohm resistor I need a 0.82pf capacitor.

nMem room temperature measurement setup – Daily Blog 8/7/23

August 7, 2023



Categories: General

I have not had much success measuring over the past few weeks and I attribute that to poor planning. What I mean exactly is that I should be performing more calculations now that I know the fabrication and measurement infrastructure is suitable. Early on, calculating the exact bias voltages, delays, or sample rates is less helpful as time is better spent resolving larger issues like thermal stability, vacuum, or electrical wiring. With those larger issues now resolved I should break out the pen and paper.

The focus on the room temperature setup was motivated by the dependence on the [delay](#). I spent a few days performing some “guess and check”; changing the components and hoping to see some improvement. That was not successful. I will review the room temperature setups.

hTron characterization – 100Hz triangle bias.


This is the same setup used by most of the users in the group to characterize the switching current of a nanowire. The scope is set to $1\text{M}\Omega$ input impedance so the majority of the supplied current is directed to the nanowire. A large input resistance of $10\text{k}\Omega$ is used to approximate an ideal current source. A $1\text{k}\Omega$ resistor is used to bias the heater. Now that I am writing this, I probably should use a larger resistor as the heater resistance is on the order of 500Ω ; or go measure the resistance of each heater to accurately calculate the current. The schematic below is for reference only: ignore the bias voltage of 10mV , the SNSPD represents the channel, the spice schematic is not simulated, and the plot on the right is measured signals, not a simulation output.



This setup works well for the accurate determination of the supplied current, thanks to the large bias resistors. However, the bias resistors are not suited for high-frequency measurements, making any read/write operations impossible with this setup. This setup produced the plots [here](#).

State Separation – 1MHz ramp-readout

To measure the switching current of the “1” state and “0” state, and thus produce the waveforms and plots on the right, we substitute current biasing for voltage biasing. This allows us to keep everything mostly impedance matched. The scope is still kept in $1\text{M}\Omega$ input impedance, which causes reflections, however, 50Ω also works just with smaller voltages on the readout. The input signals are attenuated (pi configuration according to the [datasheet](#) and values estimated). The channel has a -3dB attenuator functioning as a bias resistor as well as a shunt. The heater has a -20dB attenuator placed after the T connection to the scope to reduce the current flowing to the heater (ideally preventing the destruction of the heater) while keeping the signal on the scope reasonably large.



However, I suspect that the reflections might be contributing to the state separation, hence the delay dependence.

Standard wafer process dies – Daily Blog 8/10/23


August 10, 2023

Categories: Electronics, Fabrication, nTron, SNSPD


This will be a short post documenting the **standard wafer run** I previously wrote about. All of the GDS files are in this folder **SPG850.zip** Each die has 64 devices with some geometrical sweep.

If you'd like a die, let me know. I have some experiments I'd like to perform on each but happy to share for any reason.


Diode (first layer)




Wire




Via




SQUID




CLOVER



SNSPD






nMem further analysis – Daily 8/11/23


August 11, 2023

Categories: General




Here is the plan. Let's pick a given setup and biasing and do a deep dive. The goal is to understand this perfectly, not to change the setup to something "better" 10 mins from now. In other words, I want to be able to say why it isn't working.

The waveforms below show the channel signals for "0" and "1", and the enable signal. Each state is written with a square pulse on the channel and then read out with a ramp pulse twice the amplitude. The enable line is delayed from the writing pulse and is on for most of the read. The length of this read pulse might change the switching current, via heat, but I haven't thought about it too much, just noting it here.




On the AWG the channel voltage is set to a peak-to-peak amplitude of 68mV. This puts the peak of the read pulse at 34mV and the write pulse at 16mV. Simulating this circuit in SPICE gives a channel current of 900 μ A. This is more than double the switching current at zero heater current...




The schematic above shows a few odd components that require some explanation. First, the pi attenuator is a **Mini-Circuits VAT-3A+**. The datasheet does not give the actual resistance values but we can use analytic values. Second, the 1 μ S resistors are simply to add current probe points in spice. Third, the bias-tee is a **Mini-Circuits ZFBT-6GW-FT+**. The actual circuit has more components than just the inductor and capacitor shown in the circuit diagram (to maintain wideband performance), and the values here are calculated based on a 3GHz signal (center band) and an impedance ratio of 1:100.

For the enable signal to the heater, we are using a power splitter (Minicircuits ZFRSC-42+) which is made of three $33.66\ \Omega$ resistors in a delta configuration. The values were determined directly from the device. The power delivered to the heater will be imbalanced due to the impedance mismatch. However, this should not pose any difficulty.




The resistance of the heater should be around $800\ \Omega$, based on measurements of similar heaters. With a bias voltage of 100mVp-p applied to the power splitter, the current through the heater is around $50\ \mu\text{A}$. This is far too low. More heat needs to be applied to ensure that only the selected row is affected by the read/write signals.

The waveforms below show a scenario of particular interest. Primarily, it shows the asymmetry in the voltage output for positive and negative ("1" or "0") pulses. The green trace shows that both branches of the memory loop are normal during the write "1" operation, while the blue trace shows that at least one branch remains superconducting during the write "0" operation.



The applied heater current should be closer to $200\ \mu\text{A}$.




nMem further analysis 2 – Daily Blog 8/14/23

August 14, 2023


Categories: General, nMem

A lot of changes were made during yesterday's analysis. To briefly summarize, the circuit used for initial BER measurements kept the scope at $1M\Omega$ input impedance, ensuring latching but also generating reflections that may affect the state. Now in 50Ω , we are free of reflections but much lower bias voltages are needed, thus amplification of the readout is needed. Additionally, the control waveforms were adjusted.


An applied voltage of just 5mV will generate a $\sim 300\mu A$ current. With a bias voltage $< 5mV$, the output signal will be small and amplification is needed. Additionally, a DC-80MHz low-pass filter is on the input port of the scope. The schematic is shown below for reference only.



The main changes to the control waveforms are as follows. First, the width of the write signal was increased and the write enable signal was decreased. It is possible that during previous tests the write signal was causing relaxation oscillations between the left and right hTrons, therefore storing some unknown amount of current in the loop. With a shorter heater pulse and longer channel pulse, it is easier to see if the state latches for a given bias point. Of course, we do not want to latch but to find the point just below it. Second, the read enable signal was decreased. This reduces the amount of heat delivered to the device during the read, allowing us to compensate for the short write-enable pulse. Think grey zone width:height relation.




Below is a waveform captured from the device. The designed asymmetry is clearly visible as the "0" (blue) trace shows no switching during the write procedure, while the "1" (green) trace latches.



Notes:

The left hTron has a width of 100nm, while the right hTron has a width of 300nm.



odd scope behavior – Daily Blog 8/17/23

August 17, 2023

Categories: General



The Lecroy scope has multiple sample modes, the most common being Real-Time and Sequence. In Real-Time, the scope acquires and displays each triggered event as they occur. In Sequence, a number (segments) of waveforms can be collected and displayed in chunks. The manual description is copied below for reference, but generally, we use Sequence when acquiring large datasets (>100 segments) where displaying each slows down the measurement.

From this understanding, we would expect that the sample mode selection would not affect the device operation. The screenshots below show the opposite...

The screenshot shows the typical response of the memory in the Real-Time sample mode. The NW quadrant shows the voltage output of the memory for a W0R0 and W1R1 operation. The SW quadrant shows the enable/heater signal and the subtracted trace ($F4 = C2 - F3$) for the timing measurements (P1 and P2). The result of the measurements P1 and P2—which is essentially the switching current measurement for each state—is recorded in the SE quadrant and is the primary area of focus for this blog.

In this first screenshot, we see the switching current measurements (trends in the SE quadrant) are quite noisy. The W0R0 measurement has a higher switching current than the W1R1, i.e. pink is above yellow. The W1R1 measurement (F1, yellow) is less stable than the pink trace F2, and appears to switch between a high value and a low value. Additionally, there are two downward spikes where the measurement triggered on some noise in the waveform, these can be ignored.



The noisy behavior of the yellow trace in the screenshot above is *exactly what we would expect* for a memory biased slightly off of its operating point. If the correct amount of current is stored in the loop the yellow trace is at its low value, and if the incorrect amount is stored it reads near the high value. (Another way of viewing this is if the write pulse switches the left branch only (correct) or both branches (incorrect)).

Keeping everything else constant, the sample mode was changed from Real-Time to Sequence. This is where our troubling observations begin. Let's first point out that each sequence contains 100 measurements (or segments). This is why there are discrete chunks equal to a horizontal division in the SE quadrant. The full trend contains 1000 samples so there are 100 samples per division.

What we observe is that for a given sequence, the state measurement is far less noisy than in real-time. In other words, the memory behaves consistently within the 100 segments (1 horizontal division). We observe that some segments show the memory poorly defined—like in the 2, 4, 5, 6, and 7th divisions starting from the left. We also observe the memory states being well defined/separated like in the 1, 3, 8, and 10th divisions. The 9th division shows the inverted operation (i.e. the yellow trace is above the pink trace).



The number of segments does not change the behavior. I have tried with 2 segments (the minimum) as well as 5000 segments (the maximum). Shockingly, for 5000 measurements the behavior is the same, meaning in some instances we see large state separation over long periods with no errors. This issue needs to be solved before more sophisticated measurements can be performed.

Real Time Sampling Mode

Real Time sampling mode is a series of digitized voltage values sampled on the input signal at a uniform rate. These samples are displayed as a series of measured data values associated with a single trigger event. By default (with no Delay), the waveform is positioned so that the trigger event is time 0 on the grid. The relationship between sample rate, memory, and time can be expressed as:

$$\begin{aligned}\text{Capture Interval} &= 1/\text{Sample Rate} \times \text{Memory} \\ \text{Capture Interval}/10 &= \text{Time Per Division}\end{aligned}$$


Usually, on fast timebase settings, the maximum sample rate is used when in Real Time mode. For slower timebase settings, the sample rate is decreased so that the maximum number of data samples is maintained over time.

Sequence Sampling Mode

In Sequence Mode sampling, the completed waveform consists of a number of fixed-size segments. The instrument uses the Timebase Sequence settings to determine the capture duration of each segment. The desired number of segments, maximum segment length, and total available memory are used to determine the actual number of samples or segments, and time or points.


Sequence Mode is ideal when capturing many fast pulses in quick succession or when capturing few events separated by long time periods. The instrument can capture complicated sequences of events over large time intervals in fine detail, while ignoring the uninteresting periods between the events. Measurements can be made on selected segments using the full precision of the timebase.

Vpp:427.5mV, Read Current:520uA, Write Current:205.0, enable_current:120.0




Daily Blog 8/30/23


August 30, 2023
Categories: General





Related to the previous blog: [Noise measurements](#). Here, using nMem and scope.



Vpp:353.8mV, Read Current:400uA, Write Current:200.0, enable_current:240.0
Channel_message: 1R0R0R1R, Channel_enable: WNNRWNNR














Sequence (The scope screenshot says real-time because it was set back after capturing the data)





Daily Blog 10/9/23

October 9, 2023


Categories: General

Ok, here we go, blog time.


The past month (or two) was spent trying to understand how to get the memory working with a low-impedance room-temperature setup. The first measurements, which showed BER estimates as low as $10e^{-8}$, were likely influenced by reflections. This large state separation was no longer visible after changing the setup to a completely 50Ω circuit (minus the DUT). I was a bit worried that the memory was not functioning, though I held onto hope as the hTrons were performing as expected. Additionally, the thermal fluctuations of the probe station and the issues with the scope sampling added to my confusion.


Little by little these issues were resolved. First, I loaded a sample into the ICE to minimize the thermal fluctuations (having all of the readout circuits at RT made this easy). Second, I built up the measurement code base so I had greater flexibility when investigating the nMem behavior—in other words, I could write and plot any type of sweep at will.

When measuring in the ICE I decided to start by finding the operating point at a lower frequency, a fixed read & enable current (determined from DC hTron curves), and without emulated read/write operations. I performed a 2D sweep of the read current (y; in A) and read width (x; in number of points). With some basic BER calculations, we get the plot below. It is quite low resolution but we see that there is little dependence on read width and some dependence on read current.




Some additional context, a BER of zero means that every "1" read had a switching event and every "0" read did not. A BER of one means the opposite. A BER of 0.5 means that they both switched, not necessarily a bad thing, this is typical if we are using the time difference for a BER estimate.






Vpp:1916.3mV, Read Current:2300uA, Write Current:950.0, enable_current:230.0
Channel_message: 0NNRN1NNRN, Channel_enable: WNNENWNNEN



Vpp:1975.2mV, Read Current:2300uA, Write Current:1050.0, enable_current:230.0
Channel_message: 0NNRN1NNRN, Channel_enable: WNNENWNNEN




some nMem plots


November 14, 2023

Categories: General


V_{cpp}:774.3mV, V_{pp}:90.8mV, Write Current:413.2, Read Current:900uA,
enable_write_current:300.0, enable_read_current:200.0
Channel_message: N0NNR1NNRN, Channel_enable: NWNNEWNNEN




Vc_{pp}:707.5mV, V_{pp}:90.8mV, Write Current:300.0, Read Current:900uA,
enable_write_current:300.0, enable_read_current:215.0
Channel_message: NONNR1NNRN, Channel_enable: NWNNEWNNEN



Vc_{pp}:825.5mV, V_{pp}:99.9mV, Write Current:400.0, Read Current:1000uA,
enable_write_current:330.0, enable_read_current:213.0
Channel_message: N0NNR1NNRN, Channel_enable: NWNNEWNNEN



Vc_{pp}:825.5mV, Ve_{pp}:90.8mV, Write Current:400.0, Read Current:1000uA,
enable_write_current:300.0, enable_read_current:165.0
Channel_message: NONNR1NNRN, Channel_enable: NWNNEWNNEN





posting for meeting


November 21, 2023

Categories: General


Vpp:860.8mV, Read Current:1040uA, Write Current:420.0, enable_current:190.0
Channel_message: NONNR1NNRN, Channel_enable: NWNNEWNNEN








Vpp:864.0mV, Read Current:1080uA, Write Current:385.3, enable_current:190.0
 Channel_message: NONNR1NNRN, Channel_enable: NWNNEWNNEN




S:\SC\Measurements\SPG806\D6\D1\nMem_ICE_writeEnable_sweep7\A1


G:\My Drive__Projects_electronics\wafer_scale_process\TESTING\SPG806_D6_nMemArray4


Vc_{pp}:828.4mV, Ve_{pp}:75.7mV, Write Current:405.0, Read Current:1000uA,
enable_write_current:250.0, enable_read_current:160.0
Channel_message:NONNR1NNRN, Channel_enable:NWNNEWNNEN






Vcpp:766.5mV, Vepp:95.4mV, Write Current:300.0, Read Current:1000uA,
enable_write_current:315.0, enable_read_current:160.0
Channel_message: NONNR1NNRN, Channel_enable: NWNNEWNNEN










S:\SC\Measurements\SPG806\D6\D1\nMem_ICE_writeEnable_sweep9\A1


Vc_{pp}:855.0mV, Ve_{pp}:95.4mV, Write Current:450.0, Read Current:1000uA,
enable_write_current:315.0, enable_read_current:120.0
Channel_message: N0NNR1NNRN, Channel_enable: NWNNEWNNEN







Vc_{pp}:855.0mV, V_{pp}:93.8mV, Write Current:350.0, Read Current:1100uA,
enable_write_current:310.0, enable_read_current:130.0
Channel_message: N0NNR1NNRN, Channel_enable: NWNNEWNNEN







Vc_{pp}:855.0mV, V_{pp}:93.8mV, Write Current:350.0, Read Current:1100uA,
enable_write_current:310.0, enable_read_current:130.0
Channel_message: NONNR1NNRN, Channel_enable: NWNNEWNNEN





Vc_{pp}:855.0mV, Ve_{pp}:93.8mV, Write Current:350.0, Read Current:1100uA,
enable_write_current:310.0, enable_read_current:220.0
Channel_message: NONNR1NNRN, Channel_enable: NWNNEWNNEN






```
def range_len_pattern():  
    a = [1, 2, 3]  
    for i, v in enumerate(a):  
        ...
```

```
def range_len_pattern():  
    a = [1, 2, 3]  
    b = [4, 5, 6]  
    for i, (av, bv) in enumerate(zip(a, b)):  
        ...
```

```
import time
```

```
def timing_with_time():  
    start = time.perf_counter()  
    time.sleep(1)  
    end = time.perf_counter()  
    print(end - start)
```

G:\My Drive__Projects_physics\py-tgd\InMem\test\IV_seed\2023-11-28-14-51-44\6current_1200uA



Return of the blog – Daily Blog 1/2/24


January 3, 2024

Categories: General



Last year part of my New Year's resolution was writing a daily blog. It wasn't perfect, nor daily, but it was a good attempt and it made me a better researcher. I got off track during the fall period. I will try again this year.


In the plot below let's first consider the subplot in the upper left corner. This plot shows the current flowing through each branch of the memory (i_L , i_R) as a function of the write current. The inductive ratio determines the slope and the y-offset is equal to the persistent current. For a write bias (I_W), we see that the left branch current (i_L) is greater than the left branch critical current, thus switching that side and pushing the current over to the right. Then looking at the subplot in the top right we can define the current flowing through the right branch (i_R) given that the left side is switched. Here we see for (I_W) the current flowing through the right branch is less than the right branch critical current. This demonstrates the correct write procedure: switching only the left side. We need to also verify that writing will be successful if the memory was previously in the 0 state with a negative persistent current. The bottom two plots show the same thing but for the "0" state.



In this picture, we can construct some bounds on valid write currents. In the "1" state, the minimum write current is $216 \mu A$ ($i_L > I_{c,H_L}$) and the maximum allowable write current is $350 \mu A$ ($i_R < I_{c,H_R}$). Giving an


operating window of $134\ \mu\text{A}$. In the "0" state we get $72\ \mu\text{A}$ and $350\ \mu\text{A}$ (the persistent current doesn't matter once the left branch is switched), giving a margin of $278\ \mu\text{A}$. Of course, we take the lesser of the two.

We can now analyze these margins as a function of the persistent current.








Having determined the operating margins for the write procedure we can turn our attention to the read. In the top left figure we see the branch currents plotted as a function of the read bias for the "1" state. It is worth mentioning here that the values used for I_W and I_R were determined from measurement while the I_{c,H_L} , I_{c,H_R} values were arbitrarily assigned (to be updated later).

In the top left the read current I_R produces a left branch current (i_L) greater than the left branch switching current (I_{c,H_L}) and a i_R less than the right branch critical current I_{c,H_R} . Thus we expect to see the redistribution shown in the top right where we see the right branch will switch producing a voltage. However, a similar analysis can be followed for the "0" state in the bottom two plots. Unfortunately, both branches again switch and a voltage appears. We can continue by performing a margin analysis like we did for the write. From this analysis we need to find a persistent current that sums with the read current to switch the right branch (I_{H_R}) and subtracts to not switch the left branch (I_{H_L}).



By calculating the branch currents less the critical currents we can see what persistent current will give the correct output for the chosen read current $I_R = 570 \mu\text{A}$. In other words, $i_R - I_{c,H_R} > 0$ and $i_L - I_{c,H_L} < 0$. Based on the plot below, we need a persistent current of at least $300 \mu\text{A}$ to prevent the left branch from switching. This does not make much sense as we expect the persistent current to, well, be persistent. That means it needs to be less than the critical current with no heater enabled (approx $2 * I_{c,H_L} \approx 200 \mu\text{A}$). Ignoring this problem for now (it could be due to my assumed critical currents, or more likely an arithmetic error) we can look at the redistribution. In the "1" state we have the right branch switching, then current is redistributed to the left branch (top right) and switches the left side producing a voltage when reading a "1", as we had hoped. Unfortunately, we also get a voltage in the "0" state. I will have to come back to this...




nMem operating region contour – Daily Blog 1/4/23

January 4, 2024

Categories: General

I spent most of today double-checking my calculations and producing the figure below. Using the governing equations we need to fit to the experimental data using I_P , I_{c,H_L} , I_{c,H_R} and $\alpha = L_R/(L_L+L_R)$ as fitting parameters.




more nMem maths – Daily Blog 1/5/24


January 5, 2024

Categories: Devices, General, nMem

I went back and measured the hTron curve, $I_c(I_h)$ of cell A4. There is a very broad plateau for this device, indicating that it has a higher switching current than the other cells in this column, i.e. the critical constriction is far from this row.



This relation is needed to convert from the write-enable plots to write-switch plots. That probably doesn't make much sense; here's what I mean. Most of the measurements have consisted of bit error rates as a function of two variables, most commonly the write current (I_W) and the write enable current ($I_{EN, W}$). To calculate other parameters like I_P or I_{c,H_L} we need to convert this enable current (i.e. I_h) into a critical current (I_c) via our hTron relation. To then subdivide the total channel critical (I_{ch}) current to branch critical currents (I_{c,H_L} and I_{c,H_R}), I will be assuming constant critical current density and taking the designed width ratio of 3.




I performed a simple edge detection and fit two curves. (Not showing the fit in the plot above just the line/points of the edge)

*I really should come up with a better naming convention.

Daily Blog 1/8/24

January 8, 2024
Categories: General

Publishing now but I will fill in with more information...



Rendering GDS/DXF in Blender


June 5, 2024

Categories: Devices, General, Memo, nMem, Simulation

This is a short post to share some code for generating objects in Blender. If you're not familiar with blender, check out [CG figures](#) on youtube.

In the code you'll see a dxf is imported on line 5. From there each layer is extracted, extruded, and elevated. I then manually changed some lighting and camera settings to generate the rendered image of the nMem array.

Try it out with a DXF of your own. There may be some issue with `phidl.Device.write_dxf()`, if you encounter a problem try generating the DXF using Klayout.



```
import bpy
```

```
bpy.ops.import_scene.dxf(filepath = "G:\\My
Drive\\__Projects\\_electronics\\nMem\\simulation2022\\DXF\\nMem22_array.DXF")
```

```
objectToSelect = bpy.data.objects
objectToSelect.select_set(True)
bpy.context.view_layer.objects.active = objectToSelect
bpy.ops.object.convert(target="MESH")
bpy.ops.object.mode_set(mode="EDIT")
bpy.ops.mesh.select_all(action='SELECT')
bpy.ops.mesh.fill()
bpy.ops.mesh.extrude_region_move(TRANSFORM_OT_translate={"value":(0, 0, 0.1)})
```

```
bpy.ops.object.mode_set(mode="OBJECT")
bpy.ops.object.origin_set(type='ORIGIN_CENTER_OF_VOLUME', center='MEDIAN')
```

```
bpy.ops.object.select_camera()
bpy.ops.object.constraint_add(type='TRACK_TO')
bpy.data.objects.constraints.target = bpy.data.objects
```

```
bpy.data.objects.location = bpy.data.objects.location
bpy.data.objects.location = bpy.data.objects.location
bpy.data.objects.location = 40
bpy.data.cameras.lens = 30
bpy.ops.object.select_all(action='DESELECT')

objectToSelect = bpy.data.objects
bpy.context.view_layer.objects.active = objectToSelect
ob = bpy.context.active_object

bpy.ops.material.new()
mat = bpy.data.materials.get("Material")
bpy.data.materials.metallic = 0.50
bpy.data.materials.node_tree.nodes.inputs.default_value = 0.5
bpy.data.materials.diffuse_color = (0.103283, 0.369067, 0.8, 1)
bpy.data.materials.node_tree.nodes.inputs.default_value = (0.103283, 0.369067, 0.8, 1)

ob.data.materials.append(mat)

bpy.ops.object.select_all(action='DESELECT')

objectToSelect = bpy.data.objects
objectToSelect.select_set(True)
bpy.context.view_layer.objects.active = objectToSelect
bpy.ops.object.convert(target="MESH")
bpy.ops.object.mode_set(mode="EDIT")
bpy.ops.mesh.select_all(action='SELECT')
bpy.ops.mesh.fill()

bpy.ops.mesh.extrude_region_move(TRANSFORM_OT_translate={"value":(0, 0, 0.3)})
bpy.ops.object.editmode_toggle()
bpy.ops.transform.translate(value=(0, 0, 0.4))

ob = bpy.context.active_object
mat = bpy.data.materials.get("Material.001")
bpy.data.materials.metallic = 0.50
bpy.data.materials.node_tree.nodes.inputs.default_value = 0.5
bpy.data.materials.diffuse_color = (0.8, 0.578599, 0.173389, 1)
bpy.data.materials.node_tree.nodes.inputs.default_value = (0.8, 0.578599, 0.173389, 1)

ob.data.materials.append(mat)

bpy.data.lights.use_shadow = False

#bpy.ops.mesh.primitive_plane_add(size=2, enter_editmode=False, align='WORLD', location=(0, 0, 0), scale=(1, 1, 1))
#bpy.context.object.active_material.diffuse_color = (0.103283, 0.369067, 0.8, 1)
```

Grafana Dashboard

June 18, 2024

Categories: Cryostat, General, Lab, Memo, Operations

Out of the many possible ways to collect and display environmental data, I've found the most reliable, capable, and free option is Grafana. Before learning about Grafana from Felix, we tried python+GoogleAPI, Matlab Thingspeak, and other less memorable configurations. What I like best is that Grafana interacts with a database—opposed to interacting with a python script or a raspberry pi directly.

This blog will be about creating Slack notifications from a Grafana time series; we'll need a future blog on creating logging scripts.

Accessing Grafana

The Grafana site is hosted on the Synology NAS here <http://18.25.16.44:3010/>. You can log in with lab, and the group password (viewer access). But if you'd like to make edits or create alerts as shown below, you'll need to have an account with edit privileges.

Dashboards, Panels, Alerts

Here is a dashboard for the 300mK fridge. It isn't much, but it's far better than the Lakeshore software.




Still, let's make some improvements. The Dashboard contains three panels: two Time series and one Stat panel. When editing a panel you can change the type using the drop-down in the top right.




The data in the panel is queried via SQL. An example is shown below.

```
SELECT
  UNIX_TIMESTAMP(time) AS "time",
  stage_3K,
  stage_He3Pot,
  stage_SORB,
  stage_40K
FROM janis_300mK_logging
WHERE
  $__timeFilter(time)
ORDER BY time
```





Internal to Grafana, there are still more ways to query and display data, feel free to try other methods.

To create an alert from this panel, click the Alert tab (in the image above) and create a new alert. You can also create an alert from the alerts page but I prefer this method because it will automatically panel information (title, UID). Below I edited the SQL prompt to only include the 3K stage temp. The alert condition on this value is specified in the expression B. After this screenshot, I changed B to alert when A is below 0.4K; thus reaching the base temperature.



Your panel will now have a colored heart next to the title and annotations when the alert is pending/firing/resolved will be added automatically.



The contact point takes a Slack channel Webhook URL, these can be created in the Slack app portal. While updating the contact point you can also add {{ template "slack_title" . }} in the Title text box and {{ template "slack_message" . }} in the Text Body box.

Next, the link between label and contact point is defined in Notification policies. Firing alerts with matching labels are routed to the specified contact point. I kept all of these things named eqpt-XX. Perhaps that is confusing too, but it isn't important.

Alerting
Alert rules and notifications

Alert rules | Contact points | **Notification policies** | Silences | Alert groups | Admin

Choose Alertmanager
Grafana

Root policy - default for all alerts Edit
All alerts will go to the default contact point, unless you set additional matchers in the specific routing area.

Contact point: grafana-default-email
Group by: -
Timings: Group wait: - | Group interval: - | Repeat interval: -

Specific routing
Send specific alerts to chosen contact points, based on matching criteria

Matching labels	Group by	Contact point	Mute timings	Actions						
Matches all alert instances Matching labels <table border="1"> <thead> <tr> <th>Label</th> <th>Operator</th> <th>Value</th> </tr> </thead> <tbody> <tr> <td>label</td> <td>=</td> <td>eqpt-300mk</td> </tr> </tbody> </table> + Add matcher Contact point: eqpt-300mk	Label	Operator	Value	label	=	eqpt-300mk	-	-	-	Edit ⌵
Label	Operator	Value								
label	=	eqpt-300mk								

Each Slack channel gets its own notification policy.

+ New policy

Matching labels	Group by	Contact point	Mute timings	Actions
> label=eqpt-300mk	-	eqpt-300mk	-	Edit ⌵
> label=eqpt-dewar	-	eqpt-dewar	-	Edit ⌵
> label=eqpt-ice	-	eqpt-ice	-	Edit ⌵

Lastly, the actual text in the Slack message is configured in Contact points under Message templates. There are templates named `slack_message` and `slack_title` as we referenced earlier when editing the Alert rule.

Some examples:

```
{{ define "slack_message" }}
{{ if gt (len .Alerts.Firing) 0 }}
  {{ range .Alerts.Firing }} {{ template "slack_alert_firing" . }} {{ end }} {{ end }}
{{ if gt (len .Alerts.Resolved) 0 }}
  {{ range .Alerts.Resolved }} {{ template "slack_alert_resolved" . }} {{ end }} {{ end }}
{{ end }}
```

which references:

```
{{ define "slack_alert_resolved" }}
*{{ .Labels.alertname }}*
{{ .Annotations.summary }}
{{ .Values }}
{{ .ValueString }}


{{ .StartsAt }}
{{ .EndsAt }}
```

{{ end }}


This is a rough start but hopefully improves over time.

nMem update

August 13, 2024
Categories: Devices, General, nMem







SPG806-D6-A4-D3 -- 20240813_nMem_ICE_ber
2024-08-13 08-34-45
Vcpp:451.7mV, Vepp:90.8mV, Write Current:121.0uA, Read Current:645uA,
enable_write_current:300.0, enable_read_current:200.0
Channel_message: N0RNR1RNRN, Channel_enable: NWNWEWNWEE




SPG806-D6-D1-B4 -- nMem_ICE_ber_read
2023-12-14 20-20-14


590μA







SPG806-D6-A4-D2 -- 20240813_nMem_ICE_ber
2024-08-13 10:42:53
Vcpp:457.1mV, Vepp:78.7mV, Write Current:250.0uA, Read Current:525uA,
enable_write_current:260.0, enable_read_current:250.0
Channel_message: N0RNR1RNRN, Channel_enable: NWNWEWNWEE




Errors persistent in the memory array

August 16, 2024

Categories: Devices, nMem

Also see [Daily Blog 8/30/23](#)

SPG806-D6-A4-C3 -- 20240816_nMem_ICE_ber
Vc_{pp}:287.8mV, V_{pp}:107.3mV, Write Current:80.1, Read Current:408uA,
enable_write_current:354.5, enable_read_current:286.0
Channel_message: NONNR1NNRN, Channel_enable: NWNNEWNNEW







enable_read_current: 2.86e-04, write_current: 8.01e-05

	Value
Write Current	80.100
Write Critical Current	176.681
Write Bias Fraction	0.453
Write Enable Fraction	0.898
Read Current	408.000
Read Critical Current	476.095
Read Bias Fraction	0.857
Read Enable Fraction	0.724
Max Heater Current	394.921
Write / Read Ratio	0.196
Enable Write / Read Ratio	1.240
Enable Write Power	31.832
Enable Read Power	20.719
Write 0 Read 1	13617.000
Write 1 Read 0	2967.000
Bit Error Rate	0.083

File Saved:

S:\SC\Measurements\SPG806\D6\A4\20240816_nMem_ICE_ber\C3\SPG806_20240816_nMem_ICE_ber_D6_A4_C3_2024-08-16 08-44-17
run time 11.03 minutes





Sample time in the code below is 10us. Each measurement is recorded after a 10us period from the previous.

```
import numpy as np
import scipy.io as sio
import matplotlib.pyplot as plt

datafile = r"SPG806_20240816_nMem_ICE_ber_D6_A4_C3_2024-08-16 08-44-17.mat"

measurement_dict = sio.loadmat(datafile)

write_1_read_0_error = np.where(
    measurement_dict > measurement_dict, 1, 0
).flatten()
write_0_read_1_error = np.where(
    measurement_dict < measurement_dict, 1, 0
).flatten()

# plt.plot(write_1_read_0_error, ".")
# plt.show()

fft_write_1_read_0_error = np.fft.fft(write_1_read_0_error)
fft_write_0_read_1_error = np.fft.fft(write_0_read_1_error)

plt.plot(np.abs(fft_write_0_read_1_error), label="fft of write_0_read_1 errors")
plt.plot(np.abs(fft_write_1_read_0_error), label="fft of write_1_read_0 errors")
plt.ylim()
plt.legend()
plt.ylabel("Amplitude")
plt.xlabel("Frequency")
plt.show()


fftfreq_write_1_read_0_error = np.fft.fftfreq(
    len(write_1_read_0_error), d=measurement_dict.flatten()
)
fftfreq_write_0_read_1_error = np.fft.fftfreq(
    len(write_0_read_1_error), d=measurement_dict.flatten()
)

plt.plot(
    fftfreq_write_0_read_1_error,
    np.abs(fft_write_0_read_1_error),
    label="fft of write_0_read_1 errors",
)
plt.plot(
    fftfreq_write_1_read_0_error,
    np.abs(fft_write_1_read_0_error),
    label="fft of write_1_read_0 errors",
)
plt.ylim()
plt.xlim()
plt.ylabel("Amplitude")
plt.xlabel("Frequency ")
plt.legend()
plt.show()
```



nMem optimization using scikit-optimize


August 18, 2024


Categories: Devices, General, nMem





SPG806-D6-A4-C3 -- 20240818_nMem_ICE_ber
Vcpp:481.8mV, Vepp:107.5mV, Write Current:71.2, Read Current:417uA,
enable_write_current:331.5, enable_read_current:268.8
Channel_message: N0NNR1NNRN, Channel_enable: NWNNEWNNEW
























Nominal and Inverting nMem optimization

August 21, 2024

Categories: Devices, General, nMem



Below is the result of two optimization passes on cell A3 of nMem array A4 of the die D6 from the wafer SPG806. The two passes are identical except for the objective function, whose output is to be minimized. That result is the `res` variable in the code blocks below. For the Nominal Operation optimization, the result is 1 if the cell is in the inverting region and 0 in the ideal or nominal region. For the Inverting + Nominal optimization, the result of the objective function is 0 if the cell is in either the nominal or inverting region.

The optimization passes returned two parameter sets that measured bit error rates of 0.015 for the Nominal region and 0.007 for the Inverting region. Some of the parameters have hit the limits of the defined space. They will be adjusted later.

Nominal Operation

Input


```
space =
x0=
```

Output

```
errors = data_dict + data_dict
res = errors / (NUM_MEAS * 2)
```

Result

```
optimal parameters:
optimal function value: 0.015
run time 16.77 minutes
```



Inverting + Nominal Operation

Input



```
space =
x0=
```

Output

```
errors = data_dict + data_dict
res = errors / (NUM_MEAS * 2)
if res > 0.5:
    res = 1 - res
```

Result

optimal parameters:
 optimal function value: 0.007
 run time 16.52 minutes




S:\SC\Measurements\SPG806\D6\A4\20240821_nMem_ICE_ber\A3

nMem Calculations

September 10, 2024


Categories: Devices, General, nMem



To better understand if the reading and writing operations are optimal—meaning the current pulse does not impede operation through noise or other parasitics—the cell's state should be defined analytically. This analytical definition will rely on a few assumptions, primarily that small currents to shunt components are negligible. The model takes the cell geometry and hTron response as inputs and returns the persistent currents, read currents, read margins, critical current ratio, and kinetic inductance ratio.

The model assumes that the total write current is stored as a persistent current when the write current is within the normal operating region. The normal operating region is defined as write currents large enough to switch the left branch, but smaller than the right critical current. The inverting operating region is defined as write currents greater than the right critical current but smaller than the left critical current plus the right retrapping current. The second-order operating region is where the write current is large enough to switch both branches, leading to resistive splitting, and a voltage pulse on the output. Because of this voltage pulse, this region is undesirable and the exact persistent current does not need to be calculated. The persistent current would depend on the size of each respective hotspot and thus the pulse length. We assume that the persistent current is equal to the maximum current able to be stored in the loop with that enable current, or equivalently $I_{C,L}$.

The model also assumes that the retrapping current is the equivalent fraction of each branch's critical current. They may need to be independent/different in the future. The same can be said about read and write retrapping currents.




The array of maximum persistent currents shown above is used to calculate the state critical currents. The zero state current is defined as the maximum current that can be supported without producing a voltage while a zero, or counter-clockwise, current is stored in the loop. The same is true for the one state.

To determine the zero state current we can calculate the current needed to switch the left branch. We can be certain that in the case of a zero, the left branch will switch first because the state current is flowing in the same direction as the read current, thus summing together. The current through the left branch during the read will be $I_L = I_R \cdot \alpha + I_P$. This current needs to be greater than the left critical current $I_L > I_{C,L}$. The read current needed to switch the left side is then $I_R = (I_{C,L} - I_P) / \alpha$. Read currents greater than this value will switch the left branch.


If this read current is also greater than the current needed to switch the right side then it will cause a voltage. The current needed to switch the right side is $I_R > I_{C,R}$ and after the left is switched, is equal to $I_R = I_R - I_{r,L}$. This can be rewritten as the current needed to switch the right branch as $I_R = I_{C,R} + I_{r,L}$. Because the zero state current adds on the left side, the left will always switch first and it is always the case that the current needed to produce a voltage is the current that switches the right branch.

To determine the one-state current we can perform a symmetric analysis of the zero-state calculation. For a one, the persistent current sums with the read current on the right branch. The right branch will switch if the current $I_R = (I_{C,L} - I_P) / \alpha$ is applied. If the rediverted current $I_L = I_R - I_{r,R}$ is also greater than the left critical current, we will get a voltage.


NOTE: the retrapping current is the retrapping current of the Enable-OFF critical current. A value of $0.2 I_{C,L}$ was obtained from IV curve.



The read current is then calculated as the midpoint between the state currents. For a set read current of 200 μ A the expected region of operation is outlined. In other words, a read current of 200 μ A is greater than the zero current and less than the one current within the contour.



The predicted operating regions resemble the observed behavior, though not exactly. More measurements to come.



nMem Calculations v2

September 12, 2024

Categories: Devices, General, nMem

Depending on the input parameters, the calculated operating region can change drastically. That, like everything, can be either good or bad. The good thing is that the highly variable operating region is like the highly variable operating region we see in measurement. The bad thing is that the operating region is highly variable.

Because I couldn't make sense of why small changes to both the retrapping current and the inductive ratio could drastically change the resulting operating region, I tried to split up the two retrapping current variables. First, IRETRAP is the percentage of the maximum (enable-off) critical current used to calculate the state currents. Second, IRETRAP_ENABLE is the percentage of the instantaneous (enable-on) critical current.


Why have two different values of the retrapping current ratio? If the wire is suppressed (enable-on) we expect the critical current to be almost equal to the retrapping current. Otherwise, we expect the retrapping current to be that at $T=T_{\text{sub}}$. Perhaps this is not the right thing to do, but it might help break the problem down.

Model inputs


```
current_cell = "C1"
HTRON_SLOPE = CELLS
HTRON_INTERCEPT = CELLS
WIDTH_LEFT = 0.1
WIDTH_RIGHT = 0.3
ALPHA = 1 - 0.3

MAX_CRITICAL_CURRENT = CELLS
IRETRAP = 0.2
IRETRAP_ENABLE = 0.9
```

IRETRAP = 0.2, IRETRAP_ENABLE = 0.9



IRETRAP=0.2, IRETRAP_ENABLE=0.2




It might be more correct to call IRETRAP_ENABLE, IRETRAP_WRITE, and likewise for read. Then, one would be able to determine if the enable is on for too long, heating the cell, by observing the operating region. If there is no inverting operation then the device can cool; there is no inverting region with both retrapping currents set to 0.2.

nMem state critical current measurements

September 12, 2024


Categories: Devices, General, nMem




The read margin of the nMem is proportional to the amount of current in the loop and is calculated as the difference between the state critical currents. When performing a read sweep we see that the bit error rate goes to zero when the read is greater than the zero state current and goes back up when the read current is greater than the one state current.


By changing the amount of current we apply during the write we expect a proportional change in the read margin. Here, we performed two read sweeps for $I_W = 20 \mu A$ and $I_W = 30 \mu A$. When more current is stored in the loop, the read margin increases.

SPG806-D6-A4-C1 -- 20240912_nMem_parameter_sweep
2024-09-12 14:32:14
Vc_{pp}:1322.2mV, V_{pp}:95.7mV, Write Current:20.2uA, Read Current:577uA,
enable_write_current:289.5, enable_read_current:212.8
Channel_message: NONNRN1NNR, Channel_enable: NWNNEWNNE




SPG806-D6-A4-C1 -- 20240912_nMem_parameter_sweep
2024-09-12 14:27:55
Vcpp:1344.4mV, Vepp:95.7mV, Write Current:30.2uA, Read Current:577uA,
enable_write_current:289.5, enable_read_current:212.8
Channel_message: NONNRN1NNR, Channel_enable: NWNENWNNE






100% remote, 100% programmed, highly recommended



nMem delay measurement

September 13, 2024


Categories: Devices, nMem







Presenting at ASC gave me a week to think about what measurements should be performed between this talk and this paper submission. Ideally minimizing the number of measurements and maximizing the knowledge gained. If accomplished, we will have a good paper in a short time.

Those thoughts have culminated in a new point of interest. Instead of looking at the ideal operating point, which is the midpoint between state currents, we should focus on the state currents themselves. These state currents can be determined using the read sweep (or read margin, or whatever) where at some arbitrary threshold on the downward slope we define the zero state current, and on the positive slope, the one state. If you're reading all of these posts, please forgive the redundancy.

Having decided that read sweeps are the preferred characterization method, we can look at this response for different amounts of time between writing and reading. Below is the result of 4 read sweeps at different time intervals. Thankfully, the response does not change with delay. Longer delays will need to be measured in the future.



The increasing time interval can more clearly be observed in the enable waveform below.



Code to generate the bit error rate plot

```
import matplotlib.pyplot as plt
import numpy as np
import scipy.io as sio

plt.rcParams =
```

```
plt.rcParams = 12
```

```
def plot_read_delay(data_dict: dict):
    cmap = plt.get_cmap("Reds")
    colors = cmap(np.linspace(0.2, 1, len(data_dict)))
    for key, data in data_dict.items():
        read_currents = data.flatten() * 1e6
        ber = data.flatten()
        ax = plt.plot(
            read_currents,
            ber,
            label=f"{key}$\mu s",
            color=colors,
            marker=".",
            markeredgcolor="k",
        )
        plt.xlim(read_currents, read_currents)
        plt.xlabel("Read Current ($\mu A)")
        plt.ylabel("Bit Error Rate")
        plt.grid(True)
    plt.legend(frameon=False, bbox_to_anchor=(1, 1))

if __name__ == "__main__":
    data0 = sio.loadmat(
        "SPG806_20240913_nMem_parameter_sweep_D6_A4_C1_2024-09-13 13-30-06.mat"
    )

    data1 = sio.loadmat(
        "SPG806_20240913_nMem_parameter_sweep_D6_A4_C1_2024-09-13 13-33-45.mat"
    )

    data2 = sio.loadmat(
        "SPG806_20240913_nMem_parameter_sweep_D6_A4_C1_2024-09-13 13-26-26.mat"
    )

    data3 = sio.loadmat(
        "SPG806_20240913_nMem_parameter_sweep_D6_A4_C1_2024-09-13 13-37-24.mat"
    )


    data_dict = {
        0: data0,
        1: data1,
        2: data2,
        3: data3,
    }

    fig, ax = plt.subplots()
    plot_read_delay(data_dict)
```

nMem enable write sweep

September 16, 2024

Categories: Devices, nMem




V1 of this blog had a very colorful and overly busy plot that made sense to me and only me. V2 of this post aims to correct that.

Below are 4 copies of the same read sweep measurement (see SPG806_20240916_nMem_parameter_sweep and the script parameter_sweep.py), each performed at a fixed enable write current. The enable write current defines the critical current of our cell during the write operation.

During the write operation, a $20\mu\text{A}$ write pulse is applied and its full amplitude is ideally stored in the loop. With a non-zero current stored in the loop, the critical current of the cell should be slightly different than a cell with zero current (or equally, zero flux). In the best-case scenario, a zero state current will be $I_C - I_P$ and the one state current will be $I_C + I_P$; the variables are the critical current at zero flux and the persistent current respectively.


The state currents can be determined from the shape of the read sweep measurement. For nominal operation, a negative edge marks the zero-state current (diamond), and a positive edge marks the one-state current (cross), the same is true for the inverting operation; however the rising edge occurs before the falling edge (i.e. this is our definition of inversion). The "edge" is made of only a few data points, each an average of 1000 measurements. The state current is then the read current at 50% of the peak height. Arguments for different percentages are valid and may be used in the future.



The precision of this measurement is quite low, and we should only take these numbers as an approximation. What is of interest is the relative difference between state currents for different write critical currents. If that can be determined, it may be possible to calculate the I_C ratio between branches and other parameters of interest.

Looking at the difference between state currents we can get an idea of how much current is being stored in the loop. I am hesitant to call this the persistent current: a) because it is larger than I_W and b) during the inverting operation we expect some dissipation which I am neglecting here. There is likely some error in the calculation in addition to the coarse precision of the measurement.

The difference between operating peaks is $\sim 20\text{-}25\mu\text{A}$.



I normally like to import files via path but I've opted instead to define each file name so it is more future-proof.

```

import matplotlib.pyplot as plt
import numpy as np
import scipy.io as sio
from matplotlib.ticker import MultipleLocator
plt.rcParams =
plt.rcParams = 14

def plot_enable_write_sweep_single(data_dict: dict, index: int, ax=None):
    if ax is None:
        fig, ax = plt.subplots()
        plt.sca(ax)
        cmap = plt.get_cmap("Spectral")
        colors = cmap(np.linspace(0, 1, len(data_dict)))

        data_dict = {index: data_dict}

    for key, data in data_dict.items():
        read_currents = data.flatten() * 1e6
        ber = data.flatten()
        plt.plot(
            read_currents,
            ber,
            label=f"$I_{\{EW\}}$ = {data*1e6:.1f} $\mu$A",
            color=colors,
            marker=".",
            markeredgcolor="k",
        )

        state0_current, state1_current = find_state_currents(data)
        plt.plot(
            read_currents,
            ber,
            color=colors,
            marker="D",
            markerfacecolor=colors,
            markeredgcolor="k",
            linewidth=1.5,
        )
        plt.plot(
            read_currents,
            ber,
            color=colors,
            marker="P",
            markerfacecolor=colors,
            markeredgcolor="k",
            markersize=10,
            linewidth=1.5,
        )

    ax = plt.gca()

    plt.ylim(0, 1)
    # plt.yscale("log")
    # plt.xlabel("Read Current ($\mu$A)")
    # plt.ylabel("Bit Error Rate")
    plt.grid(True)

```



```

ax.xaxis.tick_top()
ax.xaxis.set_label_position("top")
plt.legend(frameon=True, loc=2)
# plt.title("Enable Write Sweep")
return ax

def plot_enable_write_sweep(data_dict: dict, index: int = None, ax=None):
    if ax is None:
        fig, ax = plt.subplots()
    plt.sca(ax)
    cmap = plt.get_cmap("Spectral")
    colors = cmap(np.linspace(0, 1, len(data_dict)))

    if index is not None:
        data_dict = {index: data_dict}

    for key, data in data_dict.items():
        read_currents = data.flatten() * 1e6
        ber = data.flatten()
        ax = plt.plot(
            read_currents,
            ber,
            label=f"$I_{\{EW\}}$ = {data*1e6:.1f} $\mu$A",
            color=colors,
            marker=".",
            markeredgcolor="k",
        )

        state0_current, state1_current = find_state_currents(data)
        plt.plot(
            read_currents,
            ber,
            color=colors,
            marker="D",
            markerfacecolor=colors,
            markeredgcolor=colors,
        )
        plt.plot(
            read_currents,
            ber,
            color=colors,
            marker="P",
            markerfacecolor=colors,
            markeredgcolor=colors,
            markersize=10,
        )
    plt.xlabel("Read Current ($\mu$A)")
    plt.ylabel("Bit Error Rate")
    plt.grid(True)
    plt.legend(frameon=False, bbox_to_anchor=(1, 1.25))
    plt.title("Enable Write Sweep")
    return ax

def plot_enable_write_sweep_grid(data_dict: dict):
    fig, ax = plt.subplot_mosaic(

```

```

    ],
    figsize=(16, 9),
    tight_layout=True,
    sharex=False,
    sharey=False,
)
for i, j in zip(, ):
    ax = plot_enable_write_sweep_single(data_dict, j, ax=ax)
    if i == "A":
        print(ax)
        ax.set_ylabel("Bit Error Rate")
        plt.legend(loc=2)
    if i == "C":
        plt.legend(loc=3)
    if i == "D":
        plt.legend(loc=3)

    ax.set_xlabel("Read Current ( $\mu A$ )")

ax = plot_state_currents(data_dict, ax=ax)
plt.tight_layout()
plt.show()

def plot_persistent_currents(data_dict: dict, ax=None):
    persistent_currents = []
    for key, data in data_dict.items():
        state0_current, state1_current = find_state_currents(data)
        persistent_currents.append(state1_current - state0_current)

    if ax is None:
        fig, ax = plt.subplots(figsize=(6, 4))
        plt.sca(ax)

    plt.bar(
        * 1e6 for data in data_dict.values()],
        persistent_currents,
        width=2.5,
    )
    ax.xaxis.set_major_locator(MultipleLocator(20))
    ax.xaxis.set_minor_locator(MultipleLocator(5))

    plt.grid(True, axis="both", which="both")
    plt.xlabel("Enable Write Current ( $\mu A$ )")
    plt.ylabel("Diff. Between State Currents ( $\mu A$ )")

def find_state_currents(data_dict: dict):
    read_currents = data_dict.flatten() * 1e6
    ber = data_dict.flatten()

    if np.max(ber) > 0.6:
        max_ber = np.max(ber)
        ber_threshold = 0.5 + (max_ber - 0.5) / 2
        state1_current = read_currents
        state0_current = read_currents

    else:

```

```
min_ber = np.min(ber)
ber_threshold = 0.5 - (0.5 - min_ber) / 2
state0_current = read_currents
state1_current = read_currents

return state0_current, state1_current

def plot_state_currents(data_dict: dict, ax=None):
    if ax is None:
        fig, ax = plt.subplots()
    plt.sca(ax)
    cmap = plt.get_cmap("Spectral")
    colors = cmap(np.linspace(0, 1, len(data_dict)))
    enable_write_currents = []
    state0_currents = []
    state1_currents = []
    for key, data in data_dict.items():
        state0_current, state1_current = find_state_currents(data)
        # print(f"State 0 Current: {state0_current:.2f} μA")
        # print(f"State 1 Current: {state1_current:.2f} μA")

        enable_write_currents.append(data * 1e6)
        state0_currents.append(state0_current)
        state1_currents.append(state1_current)

    plt.plot(
        data * 1e6,
        state0_current,
        "D",
        color=colors,
        markeredgewidth=1,
    )
    plt.plot(
        data * 1e6,
        state1_current,
        "P",
        color=colors,
        markeredgewidth=1,
        markersize=10,
    )

    plt.plot(
        enable_write_currents,
        state0_currents,
        label="State 0",
        marker="D",
        color="grey",
        markeredgewidth=1,
        zorder=0,
    )
    plt.plot(
        enable_write_currents,
        state1_currents,
        label="State 1",
        marker="P",
        color="grey",
```

```
        markeredgecolor="k",
        markersize=10,
        zorder=0,
    )
    plt.xlabel("Enable Write Current ( $\mu\text{A}$ )")
    plt.ylabel("State Current ( $\mu\text{A}$ )")
    plt.grid(True)
    plt.legend(frameon=False)
    # plt.title("Enable Write Sweep")

    return ax

if __name__ == "__main__":
    data0 = sio.loadmat(
        "SPG806_20240916_nMem_parameter_sweep_D6_A4_C1_2024-09-16 14-51-36.mat"
    )
    data1 = sio.loadmat(
        "SPG806_20240916_nMem_parameter_sweep_D6_A4_C1_2024-09-16 14-48-50.mat"
    )
    data2 = sio.loadmat(
        "SPG806_20240916_nMem_parameter_sweep_D6_A4_C1_2024-09-16 14-45-47.mat"
    )
    data3 = sio.loadmat(
        "SPG806_20240916_nMem_parameter_sweep_D6_A4_C1_2024-09-16 14-43-16.mat"
    )
    data4 = sio.loadmat(
        "SPG806_20240916_nMem_parameter_sweep_D6_A4_C1_2024-09-16 14-40-40.mat"
    )
    data5 = sio.loadmat(
        "SPG806_20240916_nMem_parameter_sweep_D6_A4_C1_2024-09-16 14-37-15.mat"
    )
    data6 = sio.loadmat(
        "SPG806_20240916_nMem_parameter_sweep_D6_A4_C1_2024-09-16 14-54-34.mat"
    )
    data7 = sio.loadmat(
        "SPG806_20240916_nMem_parameter_sweep_D6_A4_C1_2024-09-16 14-25-14.mat"
    )
    data8 = sio.loadmat(
        "SPG806_20240916_nMem_parameter_sweep_D6_A4_C1_2024-09-16 14-28-20.mat"
    )
    data9 = sio.loadmat(
        "SPG806_20240916_nMem_parameter_sweep_D6_A4_C1_2024-09-16 14-31-44.mat"
    )
    data10 = sio.loadmat(
        "SPG806_20240916_nMem_parameter_sweep_D6_A4_C1_2024-09-16 14-34-33.mat"
    )
    data11 = sio.loadmat(
        "SPG806_20240916_nMem_parameter_sweep_D6_A4_C1_2024-09-16 15-04-25.mat"
    )
    data12 = sio.loadmat(
        "SPG806_20240916_nMem_parameter_sweep_D6_A4_C1_2024-09-16 15-19-21.mat"
    )
    data13 = sio.loadmat(
        "SPG806_20240916_nMem_parameter_sweep_D6_A4_C1_2024-09-16 15-22-26.mat"
    )
    data14 = sio.loadmat(
        "SPG806_20240916_nMem_parameter_sweep_D6_A4_C1_2024-09-16 15-26-08.mat"
```


```
)
```

```
data_dict = {  
    0: data0,  
    1: data1,  
    2: data2,  
    3: data3,  
    4: data4,  
    5: data5,  
    6: data6,  
    7: data7,  
    8: data8,  
    9: data9,  
    10: data10,  
    11: data11,  
    12: data12,  
    13: data13,  
    14: data14,  
}  
plot_enable_write_sweep_grid(data_dict)  
plot_persistent_currents(data_dict)
```

nMem state currents and pulse widths

September 18, 2024

Categories: Devices, nMem




This post will document some efforts to optimize the state currents via tuning waveform parameters. Its purpose is also to describe a new default input sequence.

Below is a typical figure produced during a read sweep measurement (see many earlier blogs). Each channel message contains a sequence of characters that represent different pulse waveforms stored in the AWG. The figure's preamble contains the channel message and the enable message sent to the column and row respectively.

Unlike previous read sweeps, the array emulation pulses on the enable and channel lines used here are slightly different. These additional signals were chosen to stress test the memory. First, enable write pulses were added before every data operation to ensure the enable write pulse does not overheat the cell. Second, read pulses were added after data write operations to ensure that the read pulse amplitude does not switch the cell when it is cooling post-write.

As we will see in later plots, overheating from high enable currents is not necessarily a deal breaker. Any residual heat eventually relaxes to some nominal value, reducing the I_C by a small amount, shifting our state currents lower.





Below are three sets of read sweep measurement data. In each set, there are two figures, the same except only a subset of curves are plotted on the right. As indicated in the figure title, a single waveform parameter is varied; the write width, the write enable width, and the read width. These parameters have been swept in the past, as variables in the scikit optimization routine. However, that process did not yield the information that was needed. What is needed, is to understand how the grey zone of the hTron can impact the state currents.

In general, changing the pulse width changes our location on the grey zone plot, longer pulses can switch the nanowire at lower amplitudes. By increasing the write width we expect to see some point where the longer pulse width has no improvement in BER, ensuring our write operation is near unity efficiency. By increasing the enable write width we are adding more heat to the system and can expect a lower bound of heat needed to switch the wire and an upper bound where the nominal temperature of the substrate is increased. By increasing the read width we are increasing the probability of switching (same as with the write width) and expect to see a lower bound where we are near unity efficiency when

reading a zero and an upper bound where we depart from near unity of reading a one. A reminder that reading a one is the device not switching, so longer pulse widths would be worse.


The pulse width is in units of points out of a maximum of 128 points. Each waveform has a rising edge of 10 points on either side and an amplitude-dependent slope. Because of the changing slope, the exact pulse width in seconds is not directly proportional to the number of points. 100 points is approximately 400ns.

A write width of zero shows that when no state is written we have no operation. For write widths greater than 40 points we see little change in the state currents.






For low write widths, the cell is not enabled enough for a state to be written. For widths of 20 and 30 points, we see an upper bound on the state 1 current. This rising edge is pushed to the left as the enable width increases (40, 60, 80). The falling edge on the left side of the graph, the state 0 current, moves in the same direction indicating the cell is warming. Widths greater than 30 points seem sufficient. Widths greater than 40 may warm up the

cell.



Similar observations about the state currents can be for the read width plot. There are three distinct groups of rising edges, perhaps more visible on the right plot with curves 60, 80, 120. The cause of this is uncertain. Read widths greater than 40 are sufficient to operate the cell.





nMem – write sweeps with a bit more precision

September 19, 2024


Categories: General




These aren't the results you're looking for

I was expecting to see the left side, negative edge, or zero state current to move to the left in clear increments of 5uA. You could say it moves left but not clear if in steady increments



5k measurements per point




10k measurements per point






Heating from the write pulses seems to have the worst increase in BER.





SPG806-D6-A4-C1 -- 20240918_nMem_parameter_sweep
2024-09-18 15:46:05
Vc_{pp}:1361.9mV, V_{pp}:95.9mV, Write Current:15.0uA, Read Current:600uA,
enable_write_current:290.0, enable_read_current:215.0
Channel_message: NONNRN1NNR, Channel_enable: EWEEEEEEEE






nMem calculations v3

September 20, 2024

Categories: Devices, nMem



If you don't know what to write, write the date.*

*This is a mantra I've been repeating since Donnie's presentation at the group meeting a few weeks ago. Inertia. It makes more sense if you are writing on paper.

To calculate the persistent current in the memory cell, we need to know the branch critical currents, the inductive ratio, and the write current.

Branch Critical Currents


The branch critical currents can be determined using the hTron enable-current relation. In the calculation below the enable write current is swept from 380 μ A to 440 μ A, which was converted to the left and right branch currents using the measured+fit values from the enable-current relation sweep ($I_C(I_H)$). The measured critical current (i.e. the channel critical current) is the sum of the left branch retrapping current and the right branch critical current, $I_{\{r,H_L\}}(I_{\{r,H_R\}}) + I_{\{C,H_R\}}(I_{\{r,H_R\}})$, we then assume critical current ratio proportional to the designed widths 300nm:100nm.

```
# Calculate the channel critical current
channel_current_enabled = htron_critical_current(
    enable_write_currents,
    HTRON_SLOPE,
    HTRON_INTERCEPT,
)

# Define the critical currents for the left and right branches
right_critical_currents = channel_current_enabled / (
    1 + IRETRAP * (1 / width_ratio)
)
left_critical_currents = right_critical_currents / width_ratio
```


Inductive Ratio

The inductive ratio is also based on the designed geometry. In past blog posts, I treated this parameter (α) as a fitting variable and maybe should still be for fabrication tolerances. Here, we can fix this value based on the number of squares in each branch. Calculated with FEM, we have $L_L=8.7\mu$ m and $L_R=13.7\mu$ m resulting in an inductive ratio of $\alpha = L_R/(L_L+L_R)=0.612$.



Write Current


In this calculation, the write current is an input array of positive values. This applied current is divided, proportional to our inductive ratio α , between the left and right branches. If the write current exceeds the left branch critical current the left branch will switch and the write current will be diverted to the right branch. Because the hotspot formation is so rapid and the impedance increases by orders of magnitude, we assume that all of the write current is redirected to the right branch. If the write current is less than the right branch critical current then the write current will be flowing through only the right branch and becomes the persistent current when the write current is turned off. Areas where this is the case are plotted in the figure below.



Note where the values are non-zero, the maximum persistent current is equal to the write current. The lower bound marks write currents that are large enough to switch the left branch. For example, a write current of $40\mu\text{A}$ supplies a left branch current of $24.5\mu\text{A}$. We then expect to see persistent currents when the left branch critical current is less than $24.5\mu\text{A}$. At that intercept, the maximum persistent current changes from zero (dark blue) to $40\mu\text{A}$ (greenish blue). The upper bound marks when the write current is enough to also switch the right branch. In the plot above, the region beyond this bound is zero but can be non-zero if the left branch is able to retrap.

Inverting Write


The region where the write current is large enough to switch the left branch (redirecting the full write current to the right branch) and is then greater than the right branch current was added to the calculation in the figure below. This is the inverting write region. In this region, the amplitude of the persistent current is equal to the write current less the right branch retrapping current, i.e. $I_W - I_{r,H_R}$. Because we assume the retrapping current is proportional to the critical current we see a maximum persistent current that depends on both axes. The sign of the persistent current will be inverted, though the memory state has been ignored thus far.




The upper bound of the inverting region marks where the input write current exceeds the channel critical current $I_{CH} = I_{C,H_R} + I_{C,H_L}$. Similar to before, inputs exceeding the channel critical current can still result in a non-zero persistent current, which is not calculated in the figure above. Beyond the upper bound of the inverting region, a voltage is formed during the write operation.

Voltage During Write

The persistent current stored in this region is difficult to calculate. It depends on the time evolution of the hot spot in either branch and how the two interact. This is something that could be analyzed further. A reasonable assumption for the maximum persistent in the case where a write current results in a voltage is the left branch critical current with the read enabled. That way, if an enable read pulse is high, then the persistent current (less an incremental amount) would not switch the cell. In the plot below where a voltage is formed during the write, we assign a maximum persistent current equal to the left branch critical current I_{C,H_L} .



Altogether we have these regions




In past blogs, I made the mistake of feeding this colorful plot of persistent currents into the calculation of the read current. However, for our read current measurements, a constant write current is used and therefore a constant persistent current should be used in the read current calculation.


Read Current

After the write operation, a state is stored in the loop and we can calculate the critical current for each state direction. We calculate the one-state current with a positive I_P , clockwise circulating current, and the inverse for the zero state. Using these two state currents the ideal read current is calculated as the midpoint between state currents and the margin the distance between state currents. (This might not be exactly so in practice, especially due to bias-dependent noise contributions.). If our input read current is within the margin (or between state currents) then our cell is operable.

Having a new input variable, the read current, I switched to a new colormap to distinguish these plots from the previous set. The state currents are calculated using the branch critical currents, the retrapping current, the inductive ratio, and the persistent current. The state currents are positive values less than the maximum (enable off) critical current. Because of this upper bound, we see a plateau in the one state current below. It should also be clear that the state current does not depend on the y axis, the set read current. Still, we want to plot this parameter as it is one of our measurable quantities.



The ideal read current margin is calculated by subtracting the zero-state current from the one-state current. The plot below shows the difference between state currents changes linearly with the channel critical current, peaking around a left branch critical current of 110 μ A. The geometric parameters used here result in a readout that is always non-inverting, and the state current is either $I_{\mathbf{0}} = I_{r,H_L} + I_{C,H_R}$ or $I_{\mathbf{1}} = I_{r,H_R} + I_{C,H_L}$. The maximum critical current with the enable off is marked with a dashed red line.



To better visualize our operating margins, we can trim out regions where the set read current (our y-axis input) is outside our calculated state currents. In python that looks like this:


```


read_margins = np.where(
    (read_current_mesh < read_current_dict)
    * (read_current_mesh < read_current_dict),
    0,
    read_margins,
)

read_margins = np.where(
    (read_current_mesh > read_current_dict)
    * (read_current_mesh > read_current_dict),
    0,
    read_margins,
)

```





Of course, a statement checking that the read current is greater than the zero state and less than the one state might make more sense, but I opted for the above expression to allow for inverting cases.

And updating the plot:



Now the color gradient makes a bit more sense. The color denotes the read margin, or the difference between the calculated state currents. The brightest point is where the vertical difference between the top and bottom bound is greatest. As the read margin decreases, the vertical distance between bounds equally decreases, they are the same value.


For a persistent current of 20 μA , we have a boundary that can be measured by sweeping the read current or the enable read current. That region is outlined in white below. In that region, a read current of 630 μA falls between the state currents.



nMem measurement observation. Periodic errors during read

September 25, 2024

Categories: Devices, General, nMem




This blog has been at the bottom of the last three posts and I keep trimming it out and moving it to the latest draft. Instead of doing that a third time, I'll just keep this post short.

Sweeping the enable write current with zero write current (i.e. no persistent current in the loop) shows some odd behavior. The bit error rate goes from 0.5 with zero switches to some random state with many read pulses. This observed behavior is for an unrealistic input and is less relevant to understanding the device operation. Still, it may shed light on future error mitigation.

My interpretation is that the enable write pulse, though not paired with a write data pulse, is large enough to heat the cell and push the read current into the nanowire's grey zone. Reasonable enough, but interestingly, the switching events at higher enable write currents to exhibit some periodic bunching. See the two plots of enable write current equal to 299 μ A and 300 μ A (green and blue read voltage points).


The bunching period seems to be related to the enable write current, higher currents and thus higher temperatures, leading to longer bunches.

A more general way of saying this is the temperature during the read operating can have some interesting (and unintended) consequences.



Great post from way back.


SPE829-G memory



nMem operating region dynamics

October 4, 2024

Categories: Devices, Electronics, nMem




Things are starting to align.

Below are two plots comparing the calculated nominal (blue) and inverting (red) regions with the measured nominal and inverting regions. In the calculation plot, the color value equals the y-coordinate value, besides the negative sign for inverting operations. The measured state currents are extracted from the BER measurement sweeps shown below. The precision of each value is low, but the accuracy of edge detection is reasonable for the span and noise of this data.





$$I_{\min} = (I_{C,H_L} - I_P) \frac{L_L + L_R}{L_R}$$


$$I_{\min} = I_{C,H_L} + I_{r,H_R}$$

$$\text{max}(I_{C,H_R}) = I_{C,H_R} = I_{r,H_L}$$




The following parameters were used to produce the calculated operating regions: $W_L:W_R=1:1.92$ (designed: 1:3), $L_L:L_R=6.5:10$ (designed: 6.1:10), $I_r=0.625I_c$. This is currently a ballpark estimation, though pretty close to the designed values.






SPG806-D6-A4-C1 -- 20240930 nMem_parameter_sweep
2024-09-30 09:23:55
Vc_{pp}:1252.7mV, Ve_{pp}:95.9mV, Write Current:30.0uA, Read Current:536uA,
enable_write_current:290.0, enable_read_current:215.0
Channel_message: NONNRN1NNR, Channel_enable: NWNENWNNE



SPG806-D6-A4-C1 -- 20240930 nMem_parameter_sweep
2024-09-30 09:31:23
Vc_{pp}:1252.7mV, Ve_{pp}:102.5mV, Write Current:30.0uA, Read Current:536uA,
enable_write_current:310.0, enable_read_current:215.0
Channel_message: NONNRN1NNR, Channel_enable: NWNENWNNE



nMem cell comparison

October 18, 2024

Categories: Devices, nMem



This blog will consist of just two figures. The first will show the measurement of a cell's operating margins over some parameter space, and the calculated operating margins predicted using a few geometric fitting parameters. The second will compare two different memory cells measured over a similar parameter space.




Fig1. **a** shows the waveforms applied and readout from the memory cell. The **a** subpanels show: (i) the column input, (ii) the row input, (iii) the column voltage correctly reading the two states discriminated by a threshold shown in red, (iv) the column voltage with two errors (i.e. inverted), (v) a single write 1 read 0 error, (vi) a single write 0 read 1 error. **b** The bit error rate measured as a function of the read bias current (I_R), the enable read current (I_{ER}) and a fixed write current $I_W = 30.0 \mu A$ at an enable write current (I_{EW}) of (i) $290 \mu A$, (ii) $300 \mu A$, (iii) $310 \mu A$. The bit error rate was calculated from $4e3$ measurements at each point (31 total) along the shaded data slice. Markers locate operating bounds for both nominal (shaded green) and inverting (shaded purple) regions. **c** The same operating bounds are plotted as a function of the channel critical current I_{CH} (which is a linear transformation of the enable read current $I_{CH} = I_{H_{cell}}(I_{ER})$) and the read current. The color bar represents the read current at the predicted operating regions, with inverting regions assigned a negative value for clarity. The model assumes an inductive ratio of $\alpha = 0.563$, $I_r = 0.573$ and a width ratio of 2.13:1. A persistent current of (i) $-I_W$, (ii) $0 \mu A$, (iii) $+I_W$ was also included in the calculation.

SPG806-D6-A4-C1

SPG806-D6-A4-C4





Fig2. The left figure shows Fig1**b**(iii) and the right figure shows the same measurement (mostly) performed on another cell in the same column. There is a small feature that can be noticed in both plots; that being a small peak at $I_{ER} = 160 \mu A$ on the left plot and a similar looking feature at $I_{ER} = 210 \mu A$ on the right plot.

Currently, I am not sure what to make of that particular feature. Also, the reason for the difference between cells is not well understood. It is likely due to the small differences in geometry and can hopefully be accounted for.

Circuitikz Nanowire hTron nMem

January 3, 2025

Categories: Devices, Electronics, nMem





Here is a Circuitikz script for a nanowire bipole element, hTron, and nMem array.

If we create a gitrepo to standardize these elements, and any future circuit elements, our work will look more professional.

[Circuitikz manual](#)







You should be able to drop this in overleaf and generate the images above.

```
\documentclass{article}
\usepackage{graphicx} % Required for inserting images
\usepackage{circuitikz}
\title{circuit_debug}
\author{Owen Medeiros}
\date{December 2024}
```

```
\makeatletter
```



```

% used to process styles for to-path
\def\TikzBipolePath#1#2{\pgf@circ@bipole@path{#1}{#2}}
\makeatother
\newlength{\ResUp}
\newlength{\ResDown}
\newlength{\ResLeft}
\newlength{\ResRight}

% Nanowire
\tikzset{bipoles/nanowire/height/.initial=.30}
\tikzset{bipoles/nanowire/width/.initial=.80}
\pgfcircdeclarebipole{
{\ctikzvalof{bipoles/nanowire/height}}
{nanowire}
{\ctikzvalof{bipoles/nanowire/height}}
{\ctikzvalof{bipoles/nanowire/width}}
{
\pgfsetlinewidth{\pgfkeysvalueof{/tikz/circuitikz/bipoles/thickness}}\pgfstartlinewidth}
\pgfextractx{\ResRight}{\northeast}
\pgfextractx{\ResLeft}{\southwest}
\pgfextracty{\ResDown}{\southwest}
\pgfpathmoveto{\pgfpoint{\ResLeft}{0}}
\pgfpathlineto{\pgfpoint{0.5\ResLeft}{0}}
\pgfpathlineto{\pgfpoint{1.1\ResLeft}{1.1\ResDown}}
\pgfpathlineto{\pgfpoint{1.1\ResRight}{1.1\ResDown}}
\pgfpathlineto{\pgfpoint{0.5\ResRight}{0}}
\pgfpathlineto{\pgfpoint{\ResRight}{0}}
\pgfusepath{draw} %draw nanowire
}
\def\nanowirepath#1{\TikzBipolePath{nanowire}{#1}}
\tikzset{nanowire/.style = {\circuitikzbasekey, /tikz/to path=\nanowirepath, l=#1}}

\ctikzsubcircuitdef{htron}{in 1, out 1, in 2, out 2, center}{%
% reference anchor is -center
coordinate(#1-center)
(#1-center) ++(-1.2,0.5) coordinate (#1-in 1)
(#1-center) ++(-1.2,-0.5) coordinate (#1-in 2)
(#1-center) ++(1.2,0.5) coordinate (#1-out 1)
(#1-center) ++(1.2,-0.5) coordinate (#1-out 2)
(#1-center) ++(0,1) coordinate (#1-up)
(#1-in 1) to (#1-out 1)
(#1-in 2) to (#1-out 2)
}

% \ctikzsubcircuitdef{nmem}{in 1, out 1, in 2, out 2, center}{%
% % reference anchor is -center
% coordinate(#1-center)
% %
% (#1-center) ++(0,1.5) coordinate (#1-in 1)
% (#1-center) ++(0,-1.5) coordinate (#1-out 1)
% %
% (#1-center) ++(-2, 0) coordinate (#1-in 2)
% (#1-center) ++(2,0) coordinate (#1-out 2)
% %
% (#1-center) ++(0,1) coordinate (#1-up)
% (#1-center) ++(0,-1) coordinate (#1-down)
% (#1-center) ++(-1, 0) coordinate (#1-left)

```

```

% (#1-center) ++(1, 0) coordinate (#1-right)
% %
% (#1-center) ++(-0.5, -0.5) coordinate (#1-sw)
% (#1-center) ++(0.5, 0.5) coordinate (#1-ne)
% (#1-in 1) to (#1-up) --++(-0.5, 0) coordinate (#1-nw)
% (#1-nw) to ++(0,-2) coordinate (#1-sw)
% (#1-up) to ++ (0.5, 0) coordinate (#1-ne)
% (#1-ne) to ++(0,-2) coordinate (#1-se)
% (#1-sw) -- (#1-se)
% (#1-out 1) to (#1-down)
% %
% (#1-nw) ++(-0.8, 0) to ++(0,-2) coordinate (r1out)
% (#1-ne) ++(0.8, 0) to ++(0,-2) coordinate (r2out)
% (r1out) --++(0, -.25) to ++(2.6,0) -- (r2out)
% %
% (#1-in 2) to ++ (0.2, 0) --++(0, 1) --++(0.5,0)
% (#1-out 2) to ++ (-0.2, 0) --++(0, 1) --++(-0.5,0)
% }

\ctikzsubcircuitdef{nmem}{in 1, out 1, in 2, out 2, center}{%
% reference anchor is -center
coordinate(#1-center)
%
(#1-center) ++(0,3) coordinate (#1-in 1)
(#1-center) ++(0,-3) coordinate (#1-out 1)
%
(#1-center) ++(-3, 0) coordinate (#1-in 2)
(#1-center) ++(3,0) coordinate (#1-out 2)
%
(#1-in 1) to ++(0, -1)
(#1-out 1) to ++(0, 1)
(#1-in 2) to ++(1.2, 0)
(#1-out 2) to ++(-0.5, 0)
%
(#1-center) ++(1, -2) coordinate (#1-se)
(#1-se) to ++(0, 2) to ++(0,2) coordinate (#1-ne)
%
(#1-center) ++(-1, 2) coordinate (#1-nw)
(#1-nw) to ++(0, -2) to ++(0, -2) coordinate (#1-sw)
(#1-nw)--(#1-ne)
(#1-sw)--(#1-se)
% (#1-nw) to ++(0,-2) coordinate (#1-sw)
% (#1-up) to ++ (0.5, 0) coordinate (#1-ne)
% (#1-sw) -- (#1-se)
% (#1-out 1) to (#1-down)
%
(#1-nw) ++(-0.8, -2) to ++(0,-2) --++(0, -0.5) coordinate (r1out)
(#1-ne) ++(0.8, 0) to ++(0,-2) to ++(0,-2) --++(0, -0.5) coordinate (r2out)
(r1out)--(r2out)
%
(#1-ne) ++(0.8,0)--++(0.7, 0)--++(0, -2)
% (#1-in 2) to ++ (0.2, 0) --++(0, 1) --++(0.5,0)
% (#1-out 2) to ++ (-0.2, 0) --++(0, 1) --++(-0.5,0)
%
}

\begin{document}

```

```

\section{Introduction}
\begin{circuitikz}
  \draw(0,0) to ++ (2,0) to ++ (2,0) to ++ (2,0);
\end{circuitikz}

\begin{tikzpicture}
\ctikzsubcircuitactivate{htrn}

\draw (0,0) \htrn{one}{};

\end{tikzpicture}
\vspace{2cm}

\ctikzsubcircuitactivate{nmem}

\begin{tikzpicture}
  \draw (0,0) \nmem{one}{};
\end{tikzpicture}

\begin{tikzpicture}
% Number of rows and columns
\def\nrows{4}
\def\ncols{4}
% Spacing between devices
\def\xspace{6} % Horizontal spacing
\def\yspace{6} % Vertical spacing
% Draw the array
\foreach \row in {1,...,\nrows} {
  \foreach \col in {1,...,\ncols} {
    % Position of the nmem
    \coordinate (pos) at (\col*\xspace, -\row*\yspace);
    % Draw the nmem subcircuit with a unique name
    \draw (pos) \nmem{one}{};
  }
}


\end{tikzpicture}
\end{document}


\end{document}



```


nMem array enable response analysis

January 10, 2025
Categories: General















Minimum Bit Error Rate







C3

