

**Harnessing Unconventional and Multiple Parametric Couplings for
Superconducting Quantum Information**

by

Olivia Lanes

B.S., Dickinson College, 2014

M.S., University of Pittsburgh, 2016

Submitted to the Graduate Faculty of
the Dietrich School of Arts and Sciences in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2020

UNIVERSITY OF PITTSBURGH
KENNETH P. DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Olivia Lanes

It was defended on

June 12th 2020

and approved by

Michael Hatridge, Department of Physics and Astronomy

Benjamin Hunt, Department of Physics, Carnegie Mellon University

Jeremy Levy, Department of Physics and Astronomy

David Pekker, Department of Physics and Astronomy

Adam Leibovich, Department of Physics and Astronomy

Dissertation Director: Michael Hatridge, Department of Physics and Astronomy

Copyright © by Olivia Lanes
2020

Harnessing Unconventional and Multiple Parametric Couplings for Superconducting Quantum Information

Olivia Lanes, PhD

University of Pittsburgh, 2020

Josephson-junction based parametric amplifiers have become a ubiquitous component in superconducting quantum machines. Although parametric amplifiers regularly achieve near-quantum limited performance, they have many limitations, including low saturation power, lack of directionality, and narrow bandwidth. The primary aim of this thesis is to attack these problems by using multiparametric drives. We will discuss an assortment of multiparametric pumping schemes, and in particular, present an amplifying scheme that is created by applying imbalanced, simultaneous, strong gain and conversion pump tones in a Josephson Parametric Converter. This mode of amplification features broadband gain in transmission that is 10X the bandwidth of a typical, single-pumped JPC amplifier and is matched in reflection up to -15 dB, making it only one scattering parameter away from the ideal quantum-limited amplifier. Since this amplifier is matched in reflection, we also explore the possibility of being able to chain these amplifiers to enhance the saturation power and work towards a fully directional amplifier.

In addition, we discuss superconducting, thin-film van der Waals materials for use in superconducting circuits. We began by studying the nature of the contact made between 3D bulk aluminum and these 2D flakes, and discovered we could create truly superconducting contact at their interface. Next, we created a superconducting SQUID from two of these atypical junctions, and studied its quantum interference pattern. We were able to measure critical currents of 10-100's of μA , and also found that the effective area of both the junctions and the superconducting loop itself is much larger than their physical size due to the geometry of the flake. In addition, we studied the kinetic inductance of these flakes by placing them in NbTiN superconducting resonators, and found that the kinetic inductance we were measuring is around 2 orders of magnitude higher than what was predicted by theory.

The thesis paves the way for using nontraditional elements and pumping schemes in

superconducting circuits to create a new generation of devices for use in superconducting quantum information processes. In the future, hopefully, these fabrication techniques and experimental discoveries will bring the field closer to a fault tolerant, universal quantum computer.

Table of Contents

Preface	xv
1.0 Introduction	1
1.1 Motivation	1
1.2 Read-out and qubit control	2
1.3 Superconducting amplification	3
1.4 Enter the parametric amplifier	7
1.4.1 Measurement with the JPC	9
1.5 The ideal quantum amplifier	11
1.6 Van der Waals materials in quantum circuits	12
1.7 Dissertation organization	14
2.0 Quantization Overview	16
2.1 Transmon	16
2.2 The Josephson Parametric Converter and Josephson Ring Modulator	19
2.2.1 Quantizing the JPC	23
3.0 Parametric Couplings	25
3.1 Gain	25
3.1.1 Phase-preserving and phase-sensitive amplification	28
3.2 Conversion	30
3.3 GC amplification	32
3.4 GCI amplification	35
3.5 GG amplification	36
4.0 Hamiltonian Theory And Design	39
4.1 Software tools	39
4.1.1 Modeling amplifiers in Microwave Office	39
4.1.2 Ansys- HFSS	41
4.1.3 Mathematica	41

4.2	Microstrip JPC	42
4.3	4th-order nulling	43
4.4	Effects of stray inductance	48
4.5	Coupling terms and major/minor lobe regions	49
4.6	Shunted JRM	53
4.6.1	Meandering and lightning JRM	54
4.6.2	Wire JPC	57
4.7	Lumped JPC	58
4.8	Arrayed amplifier	61
4.8.1	Saturation power improvement	61
4.8.2	Array design	65
4.9	The SHARC amplifier	73
5.0	Incorporating van der Waals Superconductors Into Microwave Circuits	83
5.1	Partial edge contact with vdW materials	83
5.2	Kinetic inductance from resonator measurement	91
5.3	Fabrication	96
5.4	Experimental results at RF and DC	98
5.4.1	RF measurement	98
5.4.2	DC measurement	101
5.5	Future directions	106
6.0	Fabrication And Design Of All JPC Versions	112
6.1	Dolan bridge technique	112
6.2	Microstrip JPCs	116
6.2.1	Meandering JPC	117
6.2.2	Lightning JRM	118
6.2.3	Wire JPC	120
6.3	Lumped JPC	121
6.4	SNAIL/SHARC amplifier fabrication	123
6.4.1	First iteration	123
6.4.2	Iteration 2	125

7.0 Amplifier Data	127
7.1 WJRM data	127
7.1.1 X^2Z pumping	132
7.1.2 GC amplifier pumping	136
7.1.3 GCI amplifier pumping	143
7.1.4 GG amplifier pumping	145
7.2 Arrayed amplifiers	149
7.3 Superconducting Hybridized ARbitrary Cavity-based amplifier (SHARC)	152
8.0 Conclusions And Outlook	159
Appendix A. Recipes For Ebeam Lithography	161
A.1 Recipe for standard JPC Amplifier	161
A.2 Recipe for NbTiN shunts	162
Appendix B. Recipes for Optical Lithography	163
B.1 Recipe for NbTiN base layer	163
Bibliography	164

List of Tables

1	Calculated values for kinetic inductance	92
2	Example resonator shift modeling in sample 006-02	101
3	All resonator samples and parameters	102
4	All quality factors estimated from hanger fits shown in Fig. 58	106
5	MMA 13 spin recipe	161
6	PMMA 950K A4 spin recipe	161
7	PMMA 950K A4 spin recipe for NbTiN shunts	162
8	Shunt development recipe	162
9	LOR 5B spin recipe	163
10	S1805 spin recipe	163
11	Optical development recipe	163

List of Figures

1	Read-out in the dispersive regime	2
2	Qubit measurement with a coherent state with no amplification	4
3	Read-out of qubit with HEMT amplifier	5
4	Read-out of qubit with quantum parametric amplifier	6
5	Superconducting Qubit Read-out with a JPC	10
6	Circuit Schematic of Transmon Qubit	17
7	The eigenmodes of the JPC	19
8	Signal, idler, and common modes modes hooked up to the differential excitations of the JPC	22
9	Theoretical plots of typical gain in reflection and transmission	28
10	Theoretical plot of phase preserving vs. phase sensitive amplification in I	29
11	Theoretical plots of conversion in reflection and transmission.	31
12	Theoretical modes scattering diagram for GC amplification.	32
13	Theoretical plot of GC Amplification in reflection and transmission	34
14	Theoretical plots of GCI amplification in reflection, transmission, and compared to GC and single-G pumped amplifiers.	36
15	Frequency mode coupling between signal and idler modes from GG amplification	37
16	Theoretical plot of GG amplification compared to a single gain tone in reflection.	38
17	NLIND expansion in microwave office to create Josephson Junctions	40
18	Schematic of 8-junction, microstrip JPC	43
19	Schematic of the wire-shunted JPC	44
20	Theoretical strength of self and cross Kerr terms in 8-junction JRM vs. 4-junction wire JRM	47
21	Theoretical effects of stray inductance in the JRM	49
22	Example of flux sweep that displays both major and minor lobes	50
23	Current circulation in the JPC for major and minor lobes	51

24	Theoretical magnitude of the g_{abc} mode coupling coefficient responsible for phase-preserving gain	52
25	Theoretical magnitude of the g_{aac} mode coupling coefficient responsible for phase-sensitive gain	53
26	K-layout design for meandering JRM	55
27	K-layout design for the lightning JRM	56
28	K-layout design for the Wire-JRM	57
29	Lumped JPC circuit diagram	59
30	Lumped JPC HFSS design	60
31	Possible arrayed amplifier schematics	61
32	JPC array series schematic	62
33	First amplifier vs. second amplifier saturation power in series chain as a function of β	64
34	Maximum possible boost in saturation power as a function of amplifiers in the series array	65
35	2 JPCs arrayed in series in terms of system modes	67
36	Schematic of two WJPCs Arrayed in Series	68
37	HFSS design of 2 Lumped JPCs in series	69
38	Theoretical plots of 3-dB bandwidth in reflection and transmission for the two JPC arrayed amplifier	70
39	HFSS sweeps of avoided crossings	72
40	The circuit schematic of a SNAIL, shown in terms of individual junctions and its equivalent circuit symbol. The SNAIL consists of 3 (or more) large Josephson junctions in parallel with a smaller junction with an energy ratio of α	74
41	Schematic of triangle vs. array-type couplings	75

42	The first design of the SNAIL-based amplifier with linear array modes. The input signal should come from the left or right port, hop across the coupling capacitor "bridge", interact with the SNAIL at the center, and then leave via the same mechanism out the opposite side. The pump port for the SNAIL is created from a small coupling capacitor directly tied to the SNAIL's lumped capacitor pads at the top of the figure.	76
43	Theoretical scattering matrix of SHARC amplifier with GCI+C coupling	78
44	SHARC Schematic	80
45	SHARC Iteration 2	81
46	Final design of the SHARC amplifier	82
47	Schematic of TMD flake undergoing plasma etch and aluminum deposition	85
48	Zero-resistance contact between few-layer NbSe ₂ and aluminum.	86
49	TEM images of NbSe ₂ flake with aluminum contacts (Sample 009-4)	88
50	Zero-resistance contact between few-layer NbSe ₂ and aluminum.	89
51	Model of a 3D-3D Josephson junction compared to a 2D-3D Josephson junction	90
52	4 lumped resonators coupled to stripline to probe TMD kinetic inductance	93
53	Modified, final version of hanger resonator	95
54	Microscopic image of flake resonator	97
55	Resonator and control with and without NbSe ₂ flake (Sample 009-4)	100
56	dV/dI measurement of flake <i>NbSe₂flake</i> resonator (Sample 009-2)	102
57	Predicted vs. Experimental L_k Values	104
58	All flake resonators and controls	105
59	High power response of flake sample 013-03	108
60	New, low geometric inductance flake resonator design	109
61	New low inductance flake resonator images	110
62	E-beam exposure	113
63	Double angle deposition	115
64	Packed microstrip JPC and SEM of 8-junction JRM	116
65	Meandering JRM SEM image	118
66	Lightning JRM SEM image	119

67	SEM image of the wire-shunted JRM	120
68	Simulation and SEM image of the Lumped JPC	122
69	Microscopic image of the first iteration of the SHARC amplifier	124
70	Microscopic images of the final version of the SHARC amplifier	126
71	Visualization of how to measure cross Kerr	128
72	Flux sweep with self and cross Kerr duffing test (Sample WJPC21)	129
73	Example of self Kerr null point extraction (Sample WJPC30)	130
74	Precise Self Kerr duffing test of WJPC32	131
75	20 dB gain in reflection from X^2Z pumping scheme	132
76	X^2Z vs XYZ gain profiles	133
77	X^2Z Histograms measured at ideal and orthogonal pump phase alignment	134
78	Quantum jumps measured with X^2Z	135
79	GC data in transmission and reflection, and compared to a single-pumped G peak	137
80	Phase sensitivity of GC amplification, as shown by two idler peaks measured by the Signal Analyser	138
81	Phase sensitivity of GC amplification, as shown by two idler peaks measured by the Signal Analyser	139
82	Pulse sequence for quantifying measurement back-actionl	140
83	Weak measurement with orthogonal alignment	141
84	GC saturation power	142
85	GCI data measured in transmission and reflection, compared to GC and single G	144
86	GG vs. G data in reflection	146
87	Noise rise of GG vs. G pumping schemes	147
88	Idler forest produced by one input tone in GG pumping	148
89	Microscopic image of LJPCs in series	149
90	Flux sweeps in reflection from array amplifier	150
91	Flux sweep in transmission in array amplifier	151
92	Flux sweep and avoided crossing in SHARC amplifier	153
93	Example gain and conversion processes between different modes in the SHARC amplifier	154

94	Conversion in the SHARC amplifier, from mode a to c	156
95	Two pump conversion across SHARC	157
96	Preliminary GCI+C pumping scheme data in the SHARC amplifier	158

Preface

The first people I would like to acknowledge are those that without, this thesis would literally not exist– the lab family. First, I have to thank my advisor Michael Hatridge; First of His Name, Father of Hank and Sam, Founder of HatLab, and Knocker of Rings. Thank you for hiring me as Graduate Student 0. Thank you for your enthusiasm, willingness to explain things multiple times until they sank in, and kindness during the hard times, especially the occasional pandemic. Most importantly, thank you for building a lab culture that is healthy, supportive, and collaborative, that I have had the great privilege of being a part of.

Next I want to thank Tzu-Chiao Chien and Xi Cao. Thank goodness you guys were around at the very beginning when we were building the lab or I might still be sitting in the middle of an empty white lab surrounded by plumbing parts. Thank you for putting up with my goofiness, poor singing ability (especially to Despacito), and inability to plug in the helium spray nozzle. Tzu-Chiao especially, probably has spent more time with me over the past 6 years than any other human. He is the tamer of the EBL, finder of the 2-56 nuts, and somehow, miraculously, never gets tired of assisting or advising anyone in the lab, myself included. I will forever be grateful you transferred to Pitt, and taught me the importance of “using my eyeballs.”

I also want to thank Gangqiang Liu, who joined the lab shortly afterwards, as the first and only postdoc I’ve ever had. His patience and breadth of physics knowledge was always so inspiring and motivating. Michael Sinko of the Hunt lab has been an excellent colleague since the beginning and I am so thankful we were able to suffer through the resonator fab together while usually maintaining an excellent sense of humor. The other lab family members I have had the pleasure of working with are Pinlei Liu, Maria Mucci, Chao Zhou, Param Patel, Mingkang Xia, and Ryan Kaufman. Pinlei and Chao are far superior coders than myself, and I sincerely thank them for their help and advice. Pinlei has more enthusiasm for his work than anyone I have ever met, which makes working with him feel much less like work. Chao has an incomparable determination, and not to mention, he is champion of the one-liners.

Maria has truly impressed me with her unique ability to bring an experiment into existence with sheer willpower. I also shudder to think of what the lab would look like without her appropriately forcing us to have lab organization days. Param and Mingkang- we haven't been working together very long but you had the tedious task of editing this thesis for typos and I really do appreciate it. Lastly, I have never met anyone more willing to MacGyver something together (often with a 3D printer) than Ryan, which is a trait I really admire, even though it is tempting to chuckle at. I am very confident that after I leave, HatLab will be in excellent hands with these scientists. They are brilliant, creative, and most importantly, they are excellent people. Thank you for sciencing with me, but more importantly, thank you for your friendship.

The next group of people I want to acknowledge are those that helped me get to this stage in my life. I want to express immense gratitude towards my first year cohort: Nathan Herring, Cat Fielder, Brian Flores, and Justin Stickel. If I had to take the first year courses/quals without any of you, I am deeply suspicious I would even be here writing this dissertation. Maybe it's just the Stockholm syndrome talking, but doing horrible Fetter and Walecka/Jackson problems whilst sitting on my crappy grad student apartment floor, eating Thai food, will always be one of my fondest memories. I know people say after the 1st year of grad school you usually don't interact with the same group of people anymore, but I know we'll be talking about how the momentum operator is the generator of finite translations in space until we're all 80.

Others that have contributed to this dissertation are my thesis committee members: Jeremy Levy, Adam Leibovich, Benjamin Hunt, and David Pekker. Thank you all for your guidance, relevant and important questions, writing recommendation letters, and in general, being an awesome committee. Thank you to Leyla Hirschfeld for literally holding this department together. Thank you also, to my professors at Dickinson College for really igniting the curiosity and drive in me to succeed in physics. You are all exactly what an exemplary teacher looks like, and the reason students gain the necessary confidence to succeed in this challenging field.

Thank you to all the guys in the machine shop: Tom, Jeff, Jeff, Billy, and Sean. I know for a fact we have the coolest and nicest group of people working in any University machine

shop anywhere. I also can't tell you what a relief it was, being able to walk in there and talk about the Patriots or Steelers game for a while and not think about physics. Thank you to Jon and Linda Cohn, my next-door neighbors with whom I have developed a great friendship over the past year. Thank you for feeding me and keeping me company during the pandemic. While we are on the topic of food, thank you to Cate Hatridge for bringing us delicious lab snacks that I definitely made a dinner out of on more than one occasion, and advising us with fellowship/award/summer school application essays. Also thank you to Hank Hatridge who assisted by holding my hair that one time I had to fix the water chiller, and reminding us to "make good choices or the brown water comes out."

Next, my penultimate thanks goes to my partner and best friend Kevin Wilk, a fellow physicist who was the very first person I met from Pitt, before I even agreed to come here. Kevin has supported me through first year classes, qualifiers, broken lab equipment and just about every other seemingly terrible thing that happened over these past 6 years. Non coincidentally, he was also around for all of the happiest moments too.

Last, I want to thank my parents, Marsha and Stephan. You can rest assured that you have taught me the value of hard work, that nothing worth doing is easy, and somehow still provided me with constant encouragement, even if this did come in the form of sports metaphors involving Tom Brady. If you hadn't introduced me to *Cosmos* or let me be that weird kid who caught tadpoles and fish from the overflow of the pond, I don't think I would have discovered my love of science. Ya did good, and this dissertation is for you.

1.0 Introduction

1.1 Motivation

Quantum amplifiers are a necessary component in superconducting quantum machines due to their ability to produce high fidelity, quantum non-demolition (QND) qubit measurements [1, 2, 3]. Without these amplifiers, superconducting qubits would not have become a pioneering platform in the search for a universal, fault tolerant quantum computer [4, 5].

A universal, fault tolerant quantum computer is the ultimate goal for the field of quantum information. As fast and as impressive as modern day classical supercomputers are, some algorithms and simulations take these machines an unreasonable amount of time to complete, as they scale exponentially in terms of their time and hardware requirements. In the 1960s, Richard Feynman first theorized that simulating objects on the quantum scale could be done more efficiently on a new kind of computer that relied on quantum mechanical bits and gates [6]. After some time, others hypothesized that such a ‘quantum processor’ could not only better simulate quantum objects, but also provide a substantial speedup when performing algorithms [7, 8].

In the 90s, scientists showed for the first time concrete examples of quantum algorithms that could exchange exponential-scaling algorithms for a faster polynomial time. Famous examples are Shor’s and Grover’s algorithms, which are responsible for finding factors of large prime numbers and sorting through large data sets, respectively [9, 10]. The discovery of these algorithms reignited the field, and the race was on to find the best platform and best qubits that would be able to perform these calculations and make up a quantum computer.

There may be many physical building blocks that are capable of creating a quantum computer (neutral atoms, silicon dots, ect.), but none have made as much progress as quickly as the superconducting qubit [11]. Since the invention of the first superconducting qubit, decoherence times have gone from a few nanoseconds to hundreds of microseconds, scaling approximately by “Schoelkopf’s Law”: a factor of 10 every 3 years [12, 13, 14]

1.2 Read-out and qubit control

With the invention of circuit Quantum Electrodynamics (cQED) [15], it became possible to control and measure the qubit through 4th-order (self- and cross-Kerr) interactions with a coupled resonator. The qubit circuit can be coupled to a standard two- or three-dimensional resonator with a fundamental frequency in the microwave regime. When microwave tones are used to probe the resonator, the Kerr interactions will shift its frequency, depending on the excitation of the qubit. In other words, by operating the system in what we call the dispersive regime, meaning the frequencies of the cavity and qubit are far detuned from one another, we can deduce the state of the qubit by interacting only with the resonator [16], as shown in Fig 1.

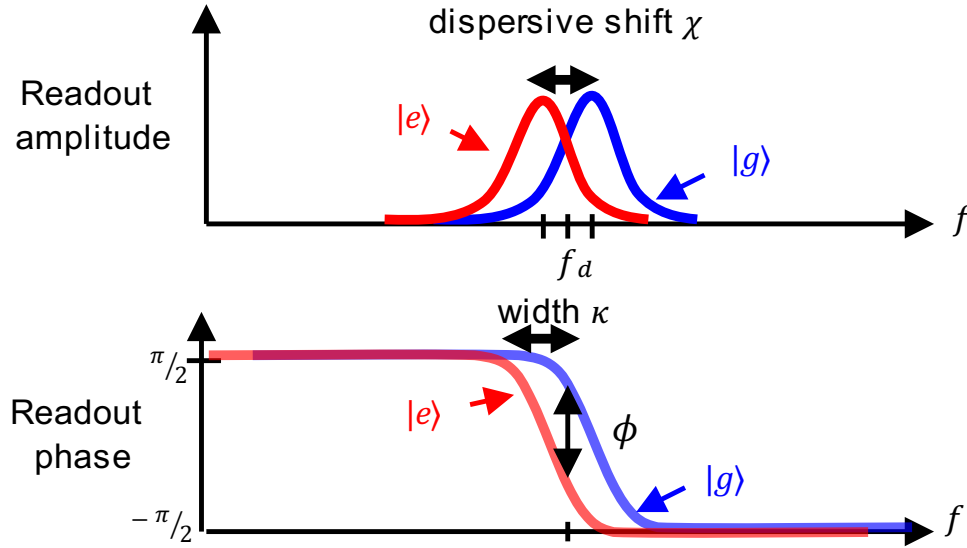


Figure 1: Reading out a superconducting transmon qubit in the dispersive regime. This pictorially shows how the frequency of the cavity shifts depending on whether the qubit state is in $|g\rangle$ or $|e\rangle$, and therefore, contains information that lets us know the state of the qubit. The dispersive shift, χ corresponds to how much the qubit state will effect the cavity, κ is the linewidth of the resonator mode and ϕ is the measurable phase difference between the two states. More details on this can be found in Chapter 2.

This theoretical plot shows how the center frequency of the cavity will shift, depending on if the qubit is in the ground or excited state. By monitoring the In-phase and Quadrature-phase (I-Q) components of the measurement signal in either phase and/or amplitude, the response of the cavity reveals the state of the qubit. Note that the measurement tone is a coherent microwave pulse, which we think of as a flying coherent state, and not a single photon (flying) Fock state.

At the time of writing this thesis, superconducting qubits and resonators have improved to the point where single-qubit and two-qubit gates perform well enough to run short, preliminary algorithms, albeit, with a moderate rate of error [17, 18, 19, 20, 21]. One way to decrease the error rate is to improve the individual components of the machine. An obvious first choice might be to improve the qubits themselves, and there is much work to be done in this regard, especially regarding coherence times [13, 22].

One of the Divincenzo criteria [23], which are the accepted, necessary components to build a quantum computer, require efficient qubit state read-out and preparation. More, building error-corrected quantum machines which can tolerate the inevitable errors due to finite qubit lifetime and gate precision, requires constant qubit measurement to detect and correct [7]. Since these measurements are used to act back on the system, they must not disrupt the quantum state more than is required, that is, the measurements must be quantum non-demolition (QND). Thus, a method for high-fidelity, QND readout is a vital pre-requisite for large-scale quantum machines. In cQED systems, where the measurement pulses are few-photon, coherent microwave pulses, this requires a quantum-limited amplification chain. Such a measurement capability is useful for not only error detection and final qubit readout, but can also help with qubit reset/ state preparation [24, 25] and, in some cases, can even produce needed quantum entanglement between qubits [26, 27, 28].

1.3 Superconducting amplification

To understand the reason quantum amplifiers are so crucial to the measurement of superconducting qubits, it is useful to first examine what would occur without one. Figure. 2

shows how the probe tone, a coherent microwave state, becomes conditionally altered as it passes through the cavity-qubit system, which separates the two gaussian distributions in I-Q space through a combination of phase and amplitude shifts. To distinguish between the two states of the qubit, the cavity output must yield distinguishable distributions when digitized at room temperature.

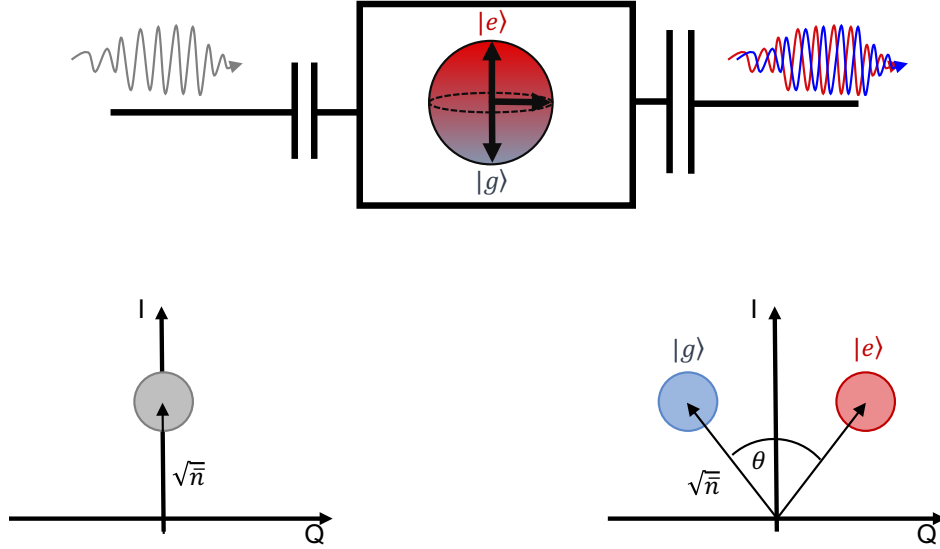


Figure 2: Read-out without amplification using Fresnell 'lollipop' diagrams. The coherent state measurement pulse is represented as a Gaussian disk in I-Q space with a displacement given by $\sqrt{\bar{n}}$, where \bar{n} is the average photon number in the coherent pulse (in grey). The signal enters the qubit/cavity system and interacts with the qubit so that it becomes phase-shifted depending on the state of the qubit (red for $|e\rangle$ and blue for $|g\rangle$). When the signal leaves the cavity, now phase-shifted by θ , this takes the coherent state and displaces it from its original position by $\frac{\theta}{2}$. The displacement from the origin and standard deviations are not affected.

These distributions are not, however, exactly observable as the quadrature operators for I and Q do not commute. However, they are also tiny amounts of microwave power (a few to a few tens of photons per pulse) and so the distribution would be essentially completely destroyed if we do not first substantially amplify the pulse. The best readily available

commercial amplifier we can use for these signals is the High Electron Mobility Transistor (HEMT) amplifier. Unfortunately, the HEMT amplifier still adds a large amount of noise to the outgoing phase-shifted signal, and thus, makes the gaussian distributions much larger in magnitude. However, it also makes them far less distinguishable.




Figure 3: Fresnell ‘lollipop’ diagrams of read-out with HEMT amplifier. When the phase-shifted signals leave the qubit/cavity, they are well distributed, but the signal is very small and weak. With the addition of the HEMT amplifier, the signal is greatly amplified, which stretches both the size and the displacement of the gaussian disks, but makes them overlap quite a bit as well due to the HEMT’s added thermal noise. Here, the gain of the amplifier has been divided out to avoid re-scaling I-Q and to highlight the added noise contribution.

As seen in the top part of Fig. 3, the read-out tones are essentially indistinguishable from one another in a single-shot read-out. To create the Fresnell diagrams seen here and have a good idea of the state of the qubit, this measurement experiment must be run hundreds or thousands of times and averaged. This is not ideal because averaging requires the experiment to be run multiple times, making single-shot read-out, and thus, error correction and state preparation, infeasible [29, 30].

Finally, we examine what the readout would look like with the addition of a quantum-limited (Q-L) amplifier inserted before the HEMT. The Q-L amplifier allows for single-shot measurement when used as the first amplifier in a read-out chain because it adds the minimal amount of noise allowed by quantum mechanics [31]. This can be seen between the 2nd and 3rd set of Fresnell diagrams; the standard deviations of the gaussian distributions have become larger, even though the gain is divided out, but the separation is still there. The HEMT amplifier is still used after the first amplifier to create even more gain (which can be seen in the top part of Fig. 4).




Figure 4: Fresnell ‘lollipop’ diagrams of read-out with quantum parametric amplifier. In this final readout scheme, our quantum-limited amplifier is used as the first stage of amplification. It adds a modest amount of gain to the readout signal (about 20 dB) and adds minimum noise. Only then is the signal amplified further by the HEMT. The Fresnell diagrams here have been normalized with respect to gain, so it is obvious that the noise added by the paramp, and then by the HEMT, does not destroy our quantum signal. The slight increase of the distributions after each stage is entirely due to added noise. The black dotted line represent the original coherent state and the white dotted lines represented the added quantum noise from the quantum-limited amplifier.

The HEMT amplifier still adds a small amount of additional noise, but it does not destroy the quantum signal any longer because the dominant noise contribution is now due to the quantum fluctuations. C. Caves tells us that these parametric amplifiers, or paramps, are capable of achieving the quantum-limit of minimum added noise, equal to $1/2$ a photon, on average ($T_Q = \frac{\hbar\omega}{2k_B}$) [31, 32] (this added noise also produces new output I and Q quadratures which now do commute). The Q-L amplifier is a crucial addition if you want to read-out quantum jumps, perform single-shot measurement, run algorithms, or have continuous qubit control [33].

However, not all gain is created equal. In order to read-out quantum signals efficiently, the noise of the pre-amplifier must dominate the total amount of noise in the read-out chain. It is not only necessary to have the amplifier operate at the quantum limit, but the gain must also be above ≈ 20 dB for significant SNR improvement, due to the HEMT’s noise temperature and typical system losses between the amplifier and HEMT [34].

1.4 Enter the parametric amplifier

To realize a quantum-limited amplifier, we leverage the concept of amplification via parametric modulation. One beautiful example, and one of the first observations, of a parametric oscillation is the “singing” wineglass [35]. Running a finger across the lip of a glass generates a familiar tone from edge waves through the liquid at a frequency exactly half that of the driving oscillation [36]. This asymmetrical correspondence of frequencies can be used to produce driven oscillations and amplification of signal as well. It was reported in 1948 by Aldert van der Ziel that parametric amplifiers could be used as extremely low-noise devices, since they produce gain through a variable reactance and not noisy resistors [37]. For a phase-preserving amplifier, the (classical) limit on the amplifier’s performance is given by the physical temperature of the amplifier and its idler mode.

However, diode and semiconductor amplifiers have moved on from parametric amplification because technological advancements in these instruments made it possible to achieve a noise temperature well below the physical temperature of the device. For example, the

Low Noise Factory HEMT amplifiers we use in our experiments operate at 4 K but have a noise temperature around 1.7 K. These state-of-the-art devices are impressive, but still add around 20 photons worth of added noise to the system, making them nowhere near the quantum limit.

Paramps on the other hand, match the temperature of their input modes because they operate from a very clean, simple Hamiltonian. This physical solution may not be desirable at higher temperatures, but at the base operating temperature of a dilution refrigerator (10 mK), it becomes advantageous because it makes the amplifier quantum-compatible with the primary requirements of the dilution refrigerator (minimal dissipated heat). By comparison, a cryogen HEMT requires operation power in the milliwatt range, which is enough to overwhelm the cooling power of the fridge (a few hundred microwatts at 100 mK).

The paramp itself is one or more superconducting resonators coupled to a Josephson junction-based element that assumes the role of the dynamic reactance. These modes can be modulated through a specific choice of pump frequency, which induces a nonlinear mixing capable of producing gain. There are many variables to work with when designing a paramp, such as basing the modes on cavities or transmission lines, and what type of Josephson junction mixing element is most suitable, etc. [38, 39, 40, 41, 42, 43].

As a result, there are multiple quantum parametric amplifiers available to implement in a quantum machine, each with their own pros and cons. At present, there are two families of parametric amplifiers: those based on discrete microwave modes and those based on nonlinear transmission lines, the so-called Traveling Wave Parametric Amplifiers (TWPAs). TWPAs have superior instantaneous bandwidth and saturation power, but their typical reported noise performance is a factor of a few higher than resonant mode-based amplifiers [44, 45, 46]. Additionally, although they are directional amplifiers, reflected pump and signal tones still require operation with external circulators or isolators. These external components are a significant problem in quantum information experiments because they are bulky and lossy, thus, lead to inefficiencies in our measurement.

In contrast, resonator-based amplifiers, such as the Josephson Parametric Converter (JPC), which is used extensively in this thesis, typically contain far fewer Josephson junctions, are easier to fabricate, and operate very near the quantum limit. Of course, they have

their own limitations, which will be discussed later in this chapter, but these amplifiers impressively operate within a factor of 2 of the quantum limit. They consist of three microwave modes, spaced out in frequency, which can be pumped at different frequencies to create different parametric interactions. They also contain rich physics and offer the possibility of improvements to qubit read-out, given new engineering and Hamiltonian design.

There are even more distinctions to be made in the category of mode-based parametric amplifiers. There is, for example, the single mode amplifier, or the SNAIL-based amplifier. This was invented at Yale and features a simple operation as well as Kerr-cancellation [47, 48, 49], which will be discussed more in further chapters. However, it also operates in reflection, isn't directional, and because it only contains one mode, offers little room for Hamiltonian engineering. Other mode-based amplifiers include the Field Programmable Parametric Amplifier [50] and the Josephson Array Mode Parametric Amplifier [51].

However, the hardware itself does not completely determine a paramp's function. This dissertation discusses substantial efforts to try and improve the fabrication and hardware of these systems, and in addition, explores how one can use the same physical platform, but different pump combinations, to enhance an amplifier's characteristics.

1.4.1 Measurement with the JPC

As briefly mentioned, the JPC is an amplifier that consists of three microwave modes. They couple to the Josephson Ring Modulator (JRM), the heart of the amplifier, consisting of a ring of Josephson junctions, often shunted with other inductors for better Hamiltonian control (see Chapter 2). The JPC is then biased with an external magnetic flux to create 3-wave nonlinear interactions in the JRM, which is described more thoroughly in Chapter 2. The two most-used interactions are gain and conversion, which occur when one of the modes is pumped at the sum or the difference of the other two modes, respectively.

Gain is created by driving mode c , for example, at the sum of the mode a and b frequencies. This produces an amplified tone in a and b at the corresponding mode frequency in both transmission and reflection. Similarly, noiseless frequency conversion occurs between a and b when mode c is driven at the difference of their frequencies.




Figure 5: Superconducting qubit read-out with a JPC. As before, when coherent signal tones enter a qubit-cavity system, they become phase shifted, depending on the state of the qubit. Once the signal leaves the the cavity, it is directed forward by two commercial circulators before being amplified by the JPC. The circulators are crucial for directing amplified signal away from the delicate qubit, but are a large source of loss and inefficiency in the read-out chain.

The JPC traditionally works in reflection, meaning amplified signal comes back along the same direction the original signal was input. If this large amplified signal were to go back and interact with the qubit, it could easily disrupt or scramble the state we are trying to read out.

To combat this issue, we install cryogenic circulators in the read-out chain, as shown in Fig. 5. These cryogenic circulators are commercial components that are bought, not fabricated (although quantum-limited microwave circulators have been created [52, 53]). Any time a physical component is added after the qubit, additional loss is generated due to leakages in the cable connections or the components themselves. Any source of loss can drastically hurt the quantum efficiency.

Though not obvious from the diagram, to accurately read-out the qubit-cavity system,

the amplifier's frequency has to be aligned with the cavity. Generally, the bandwidth of the amplifier is only on the order of a few MHz at 20 dB of gain, because the JPC suffers from a gain-bandwidth product [54], meaning the gain and bandwidth of the amplifier are inversely proportional. The JPC must be tuned so that its rather narrow bandwidth aligns with the frequency of the read-out cavity, which is not possible in all scenarios. However, if one desires to create a practically useful quantum computer, many qubits and cavities are necessary, each requiring amplification. If the amplifier's bandwidth is narrow, then each cavity will require its' own amplifier. This is not ideal because of the time and resources required to tune up and bias each amplifier. In addition, each amplifier occupies room in the finite space of a dilution refrigerator.

Another aspect that holds us back from multiplexing the read-out chain with multiple qubits through one amplifier is the saturation power of the amplifier. The saturation power is the amount of power that can be input to the amplifier before the gain begins to fall. We typically measure this in terms of the $P_{\pm 1dB}$ point, which refers to the point at which the gain diverges by 1 dB [55]. There are some subtleties with this, since we have discovered that there are certain bias conditions at which the JPC begins to saturate upwards [56]. However, in general, JPC's have a typical low saturation point of ≈ -125 dBm [56].

1.5 The ideal quantum amplifier

We have now described two broad categories of imperfect parametric amplifiers; the TWPA and mode-based paramps. Both are practical devices, but neither is ideal for qubit measurement. At this point, let's step back from reality and instead imagine the quintessential, ideal, quantum parametric amplifier.

The ideal quantum amplifier needs to have four main criteria. First, it needs to have wide bandwidth in order to simultaneously measure multiple qubits at different frequencies [39, 57]. Previously, this was thought to be unreachable due to the fixed gain-bandwidth product of singly pumped parametric amplifiers. Second, the amplifier has to have a rather large input saturation power, which again, would make it possible to measure multiple qubits

at the same time [58].

We also need our amplifier to be as close to the quantum limit as possible. The amplifier should only add, on average, $1/2$ photon of noise to the system if operated in a phase-preserving mode. The closer the amplifier gets to this fundamental limit, the better the quantum efficiency. On this front, we might say we have achieved our goal. We regularly operate the JPC at about twice the standard quantum limit.

Lastly, we want our amplifier to be directional, meaning it amplifies in transmission, not reflection, and breaks reciprocity [54, 59, 60]. It is not enough to have an amplifier that only works in transmission, as it may amplify in transmission in both directions (i.e. from a to b and b to a) simultaneously. The amplifier has to break this symmetry in order to eliminate any signal that could travel back towards the qubit and disrupt its fragile state. This would allow us to get rid of the cryogenic circulators and greatly boost our quantum efficiency.

We are working towards creating the ideal amplifier through a combination of better hardware and pump interactions. We have explored a variety of pumping schemes and ideal bias points, such as phase-sensitive amplification only possible at specific ranges of magnetic flux bias, and also specific biases which tune away higher order, unwanted Hamiltonian terms. We also use combinations of gain and conversion pumps to create specific types of parametric interference, which can eliminate the gain-bandwidth product and widen the amplifier's bandwidth without use of low Q modes. Lastly, we discuss how to build an amplifier based on a one-dimensional chain of modes (very reminiscent of a short TWPA) with the superior Hamiltonian control of a JPC in a few-mode array-type amplifier.

1.6 Van der Waals materials in quantum circuits

Outside of the realm of amplifiers, I have also worked during my PhD to explore novel materials for use in superconducting circuits for quantum information. In collaboration with Michael Sinko in Benjamin Hunt's lab at Carnegie Mellon University, we studied van der Waals materials, specifically superconducting transition metal dichalcogenides (TMDs). These materials offer some advantages compared to traditional materials used in supercon-

ducting quantum information experiments. Like graphene, TMDs can be exfoliated only a few atomic layers thick, but unlike graphene, some TMDs, specifically NbSe₂ and TaS₂ are intrinsically superconducting, even in single layer films [61]. Since these materials tend to degrade in contact with oxygen, they are traditionally encapsulated with the insulator hexagonal Boron Nitride (hBN).

Traditional methods of contact to vdW films, such as a graphene segment in contact with the flake being studied [62], have the disadvantage that they have significant ($k\Omega$ or greater) normal resistance in series with the flake. This precludes us from forming devices such as SQUIDs [63] or superconducting microwave circuits from these TMDs, unless the entire circuit is contained in a single flake, which is an overwhelming fabrication challenge. Thus, our first task has been to develop a TMD-bulk superconductor interface with zero resistance.

We have been able to create these joints [64], which are a form of novel Josephson junction. The junctions themselves are unusual, and we have discovered that their thin geometry allows much more of the magnetic flux to penetrate the flake, creating larger effective areas for the junction. We have even been able to fabricate a Superconducting Quantum Interference Device (SQUID) from the combination of NbSe₂ and aluminum, which displays the characteristic interference pattern paramount to these devices. We have also explored how to measure and extract other important properties of these 2D materials, such as the kinetic inductance and strength of nonlinear terms, by embedding vdW flakes/junctions in microwave resonators. More exotic devices could be developed from these materials once there is a procedure for understanding and channeling these properties.

We are very optimistic that junctions made from these materials can be used in amplifiers or even qubits, as they feature some highly desirable qualities such as magnetic field tolerance, precise layer stacking control, and, possibly, non-sinusoidal current-phase relations. This project is still in its early stages, but now with a working fundamental fabrication plan and promising preliminary results, we expect new devices with complex physics to be built from this foundation.

1.7 Dissertation organization

This thesis is organized as follows. First, Chapter 2 will discuss the quantization of superconducting circuits, including the qubits and amplifiers used throughout this dissertation, and the circuit QED framework we use for qubit control and readout. Chapter 3 will discuss several multiparametric processes which are available in any device that has a Hamiltonian with a 3-body mixing term (such as the JPC). Further, we look at how interference from multiple, simultaneous parametric drives in a few-mode device can combine to give rise to novel modes of amplification that have increased bandwidth, no gain-bandwidth product, and even give our amplifiers matched inputs (no reflections). In Chapter 4, I will review how to physically realize a desired Hamiltonian using the multiple simulation tools at our disposal, including Microwave Office, HFSS, and Mathematica. In addition, I will present some subtleties in the designs including stray inductance and higher order effects. I will walk through how we arrive at the final design for the amplifiers used in this thesis.

Chapter 5 is devoted to the theory, fabrication, and results of our progress made with 2D vdW materials. It will show how we supplement dc characterization tools with rf measurements of these circuits based on the techniques we use to characterize parametric amplifiers. We show how to make true superconducting contact between 2D TMD flakes and bulk aluminum circuits, while examining the resultant junctions' strange and unique properties due to substantial flux penetration through the few-layer thick flakes. We were able to make both individual Josephson junctions and SQUIDS, and measured very high critical currents, in the tens of microamp range, from these devices. At the end, we present a study of the kinetic inductance of few-layer flakes of NbSe₂, which we measure to be orders of magnitude higher than simple expressions for kinetic inductance predict. If this can be confirmed and explained, this could be highly useful for components in superconducting circuits such as super-inductances.

Chapter 6 presents all of the fabrications recipes and techniques used throughout the dissertation, and Chapter 7 contains the experimental data collected from our parametric amplifiers. Specifically, we show that multiparametric pumping schemes, such as a combination of gain and conversion (which we term GC for convenience), do not hinder quantum

efficiency, and in fact, boost the bandwidth and improve scattering parameters of the device. The initial multiparametric pumping data is presented all from the same amplifier to compare their features, and finishes with the imbalanced GC pumping scheme (termed GCI), which achieves an order of magnitude larger bandwidth than a traditional JPC operated at 20 dB gain. More, imbalanced GC amplification results in matched devices, and thus, unlike all of our previous amplifier designs, can be arrayed. Next, the new preliminary data from our array amplifier and our new Superconducting Hybridized ARbitrary Cavity-based (SHARC) platform, is discussed. Lastly, Chapter 8 presents my conclusions and future outlook.

The field of superconducting quantum computing is still some time away from not only an ideal amplifier, but a firm understanding of how to best control and use nonlinear elements in a circuit. However, my work shows how a combination of multiparametric drives in the JPC can lead to significant improvements of its characteristics and additionally, pave the way for a device that may be fully directional. Furthermore, it lays the groundwork for how to incorporate 2D superconducting TMD materials into these circuits to study their contact properties and nonlinearity.

2.0 Quantization Overview

This chapter will serve as the first theoretical look into cQED, which governs the superconducting transmon qubit and the Josephson Parametric Converter. It will begin with a brief description of the transmon, followed by a quick derivation of the quantized Hamiltonian and a look at mathematical description of the coupled qubit-cavity system. Next, we apply these same physical principles to the JPC with linear shunts, and work through the eigenmodes of this new geometry. Finally, we show how to quantize the JPC in terms of the three modes' raising and lowering operators, up to 3rd order. We extend this to the 4th order (and higher) terms, which limit amplifier performance but are not fundamental to its operation, in the next chapter.

2.1 Transmon

All qubits that are discussed and measured in this thesis are superconducting transmon qubits, a variety that was invented at Yale [65, 66, 67] and soon adopted as a qubit of choice around the world, both in academic and industrial labs. The transmon is the simplest superconducting qubit to both fabricate and measure and consists of a Josephson junction in parallel with a large capacitance, as shown in Fig. 6.




Figure 6: The circuit schematic of a transmon qubit, consisting of a single Josephson junction of energy E_J , shunted by a capacitor with energy $E_C = \frac{Q^2}{2C_J}$.

The Hamiltonian that describes this circuit is given by:

$$H = \frac{Q^2}{2C_J} - E_J \cos(\phi). \quad (2.1)$$

Where the Q is the charge across the capacitor, ϕ is the phase difference of the superconducting order parameter across the junction, and C_J is the capacitance. We can expand the second, nonlinear term in ϕ , where $\phi = \frac{2\pi\Phi}{\Phi_0}$ as in [15, 68],

$$-E_J \cos(\phi) \approx -E_J \left(1 - \frac{1}{2} \left(\frac{2\pi\Phi}{\Phi_0} \right)^2 + \frac{1}{4!} \left(\frac{2\pi\Phi}{\Phi_0} \right)^4 + \dots \right), \quad (2.2)$$

and rewrite the energy in terms of the inductance, $E_J = \frac{\Phi_0 I_0}{2\pi} = \frac{\phi_0^2}{L_J}$. The transmon regime occurs when $E_J \gg E_C$ [65]. From here, we can quantize the Hamiltonian by seeing that the flux in the nonlinear inductor, the junction, Φ , and the charge across the capacitor, Q act as quantum operators, thus following a commutation relation given by

$$[\hat{\Phi}, \hat{Q}] = i\hbar. \quad (2.3)$$

With this, our quantum operators, $\hat{\Phi}$ and \hat{Q} can be mapped onto bosonic raising and lowering operators by

$$\hat{\Phi} = \sqrt{\frac{\hbar Z}{2}}(\hat{a} + \hat{a}^\dagger) \quad (2.4)$$

$$\hat{Q} = i\sqrt{\frac{\hbar}{2Z}}(\hat{a} - \hat{a}^\dagger), \quad (2.5)$$

where Z is the characteristic impedance, given by

$$Z = \sqrt{\frac{L}{C}}. \quad (2.6)$$

Finally, we rewrite the transmon Hamiltonian in Eqn. 2.1 by expressing it in terms of the raising and lowering operators:

$$H \approx \hbar\omega a^\dagger a - \frac{\alpha}{2}\hbar a^\dagger a a^\dagger a + \dots \quad (2.7)$$

where ω and α are the frequency and anharmonicity of the qubit, respectively. Now, it becomes clear that we are left with the typical Hamiltonian for a quantum harmonic oscillator, but with the addition of a higher order perturbation. This 4th-order perturbation is crucial, because it is responsible for introducing nonlinearity to the system, and effectively stretches the harmonic potential. This now anharmonic potential features energy levels that are unevenly spaced, making it possible to isolate lowest energy levels of the system, $|g\rangle$ and $|e\rangle$, with a specific drive frequency.

In order to perform QND measurements with the transmon, the qubit itself is coupled to a superconducting resonator, which can be either 2 or 3-dimensional [69, 70, 71]. By coupling the qubit to the resonator, the resulting Hamiltonian that describes the system is famously [72]

$$\frac{H}{\hbar} = \tilde{\omega}_c c^\dagger c + \tilde{\omega}_a a^\dagger a - \frac{\alpha}{2}(a^\dagger a)^2 + \chi(a^\dagger a c^\dagger c), \quad (2.8)$$

where the c/a and c^\dagger/a^\dagger operators represent the raising and lowering operators for the cavity and qubit respectively. The χ is the dispersive shift, symbolizing how changing the photon number of the qubit shifts the frequency of the cavity (and vice versa) and the $\tilde{\omega}_c$ and $\tilde{\omega}_a$ are the coupled frequencies of the cavity and qubit.

2.2 The Josephson Parametric Converter and Josephson Ring Modulator

The JPC was also invented at Yale [73, 74] and further developed at ENS [40]. The core of this amplifier is the Josephson Ring Modulator (JRM); a ring of four nominally identical Josephson junctions. The ring is threaded with an external magnetic flux which changes the frequency of the intrinsic resonant modes.

The original device had no shunting inductors, linear or nonlinear, but this led to instability in the lowest order solution of the ring when magnetic flux was applied. To extend the stability of the JRM there needs to be an added inductance in the center of the JRM itself [75]. There are many papers that work through the derivation of the JRM's Hamiltonian with Josephson junction shunts, but here we will focus on the linear shunted JRM [34, 73].

We refer to Fig. 7 (a), where all the elements as well as the node phases of the device have been labeled.




Figure 7: (a) The circuit diagram for the shunted JPC with phases across all elements labeled. (b) The eigenmodes of the JPC displayed each with their differential voltage drives. The X, Y, and Z are the usable eigenmodes of the system, but the W mode does not couple to the outside world and is a trivial mode, described by 2.13.

If we write down the Hamiltonian of the JRM with these elements we get

$$H_{\text{shuntedJRM}} = -E_J [\cos(\varphi_a) + \cos(\varphi_b) + \cos(\varphi_c) + \cos(\varphi_d)] + \frac{E_L}{2}(\varphi_\alpha^2 + \varphi_\beta^2 + \varphi_\gamma^2 + \varphi_\delta^2), \quad (2.9)$$

where $E_L = \frac{\Phi^2}{L}$. We can also directly substitute $\varphi_\alpha, \varphi_\beta, \varphi_\gamma, \varphi_\delta$ with their corresponding closest flux node minus the center flux [76],

$$\begin{aligned} \varphi_\alpha &= \varphi_1 - \varphi_E \\ \varphi_\beta &= \varphi_2 - \varphi_E \\ \varphi_\gamma &= \varphi_3 - \varphi_E \\ \varphi_\delta &= \varphi_4 - \varphi_E. \end{aligned} \quad (2.10)$$

In addition, φ_E can be set to 0 as it has no constraint.

We also express the flux across the Josephson junctions as the difference in the node fluxes plus the added external flux which is equally distributed:

$$\begin{aligned} \varphi_a &= \varphi_2 - \varphi_1 + \varphi_{\text{ext}}/4 \\ \varphi_b &= \varphi_3 - \varphi_2 + \varphi_{\text{ext}}/4 \\ \varphi_c &= \varphi_4 - \varphi_3 + \varphi_{\text{ext}}/4 \\ \varphi_d &= \varphi_1 - \varphi_4 + \varphi_{\text{ext}}/4. \end{aligned} \quad (2.11)$$

For the sake of clarity, the Hamiltonian with all these substitutions, can now be described by the slightly less obvious expression given by

$$\begin{aligned} H_{\text{shuntedJRM}} &= -E_J [\cos(\varphi_2 - \varphi_1 + \varphi_{\text{ext}}/4) + \cos(\varphi_3 - \varphi_2 + \varphi_{\text{ext}}/4) \\ &\quad + \cos(\varphi_4 - \varphi_3 + \varphi_{\text{ext}}/4) + \cos(\varphi_1 - \varphi_4 + \varphi_{\text{ext}}/4)] \\ &\quad + \frac{E_L}{2}(\varphi_1^2 + \varphi_2^2 + \varphi_3^2 + \varphi_4^2). \end{aligned} \quad (2.12)$$

Lastly, the normal modes, or eigenmodes, of the JRM can be expressed as [73, 76]:

$$\begin{aligned}
\phi_X &= \phi_1 - \phi_3 \\
\phi_Y &= \phi_2 - \phi_4 \\
\phi_Z &= \phi_1 + \phi_3 - \phi_2 - \phi_4 \\
\phi_W &= \phi_1 + \phi_2 + \phi_3 + \phi_4.
\end{aligned} \tag{2.13}$$

These substitutions can easily be seen by examining Fig. 7 (b) in terms of the differential drives governing each eigenmode. In order to perform the substitution, it is more useful to solve for the inverse transformation, which gives,

$$\begin{aligned}
\phi_1 &= \frac{-\varphi_X}{2} - \frac{\varphi_Z}{2} + \frac{\varphi_W}{2} \\
\phi_2 &= \frac{-\varphi_Y}{2} + \frac{\varphi_Z}{2} + \frac{\varphi_W}{2} \\
\phi_3 &= \frac{\varphi_X}{2} - \frac{\varphi_Z}{2} + \frac{\varphi_W}{2} \\
\phi_4 &= \frac{\varphi_Y}{2} + \frac{\varphi_Z}{2} + \frac{\varphi_W}{2}.
\end{aligned} \tag{2.14}$$

When substituted into Eqn. 2.12, we recover our Hamiltonian explicitly written in terms of these eigenmodes as in [34],

$$\begin{aligned}
H_{shuntedJRM} &= -4E_J \left[\cos\left(\frac{\varphi_X}{2}\right) \cos\left(\frac{\varphi_Y}{2}\right) \cos(\varphi_Z) \cos\left(\frac{\varphi_{ext}}{4}\right) \right. \\
&\quad \left. + \sin\left(\frac{\varphi_X}{2}\right) \sin\left(\frac{\varphi_Y}{2}\right) \sin(\varphi_Z) \sin\left(\frac{\varphi_{ext}}{4}\right) \right] \\
&\quad + \frac{E_L}{4} (\varphi_X^2 + \varphi_Y^2 + 2\varphi_Z),
\end{aligned} \tag{2.15}$$

where it now becomes clear that the 4th eigenmode, W, is a trivial mode that does not couple to the outside world. We can see from this equation that the four Josephson junctions on the outer arms of the JRM provide nonlinear couplings between the eigenmodes of the circuit [41]. Assuming that the ground state of the circuit is $\varphi_X = \varphi_Y = \varphi_Z = 0$ and it is stable as

we tune the external magnetic flux bias, we can expand the nonlinear coupling terms around this ground state [76],

$$\begin{aligned}
H_{shuntedJRM} \approx & \left[\frac{E_L}{4} + \frac{E_J}{2} \cos\left(\frac{\varphi_{ext}}{4}\right) \right] (\varphi_X^2 + \varphi_Y^2) + \left[\frac{E_L}{2} + 2E_J \cos\left(\frac{\varphi_{ext}}{4}\right) \right] \varphi_Z^2 \\
& - \underline{E_J \sin\left(\frac{\varphi_{ext}}{4}\right) \varphi_X \varphi_Y \varphi_Z} - \frac{1}{96} E_J \cos\left(\frac{\varphi_{ext}}{4}\right) (\varphi_X^4 + \varphi_Y^4 + 16\varphi_Z^4) \quad (2.16) \\
& + \frac{1}{16} \cos\left(\frac{\varphi_{ext}}{4}\right) (\varphi_X^2 \varphi_Y^2 + 4\varphi_X^2 \varphi_Z^2 + 4\varphi_Y^2 \varphi_Z^2) + \dots,
\end{aligned}$$

where the underlined term represents the desired 3rd order mixing term that is responsible for gain and conversion in the JRM [77]. The higher order terms that follow are significantly smaller but are all still unwanted and will hinder the performance of the amplifier, as is discussed in further chapters [78].

If we picture the JRM now as a three mode coupler, we can imagine attaching three resonators to the intrinsic, nontrivial modes of the system, X, Y , and Z , through their differential drives. When these resonators are connected to these normal modes, they are traditionally referred to as the a, b , and c modes, or the signal, idler, and common, as shown below in Fig. 8. This makes up the full Josephson Parametric Amplifier

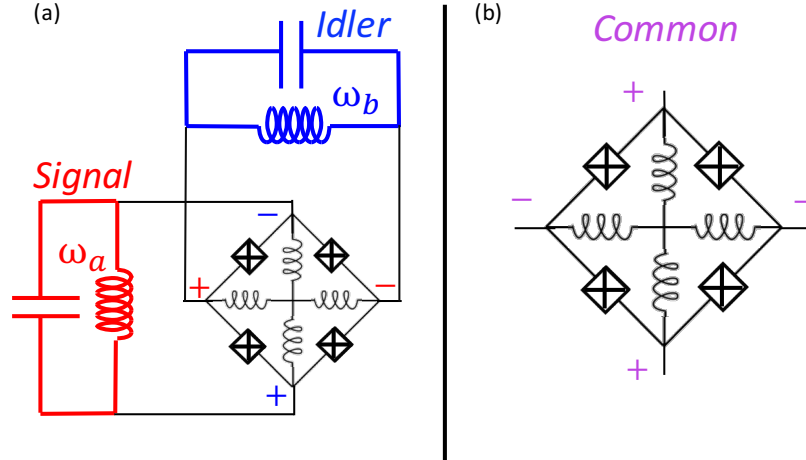


Figure 8: (a) Signal and Idler modes hooked up to the JRM through the differential excitations of the system. (b) The common mode is not a physical mode with its own port, but as is clear from the differential excitations, it is capable of leaking into either of the other two ports. These three modes together make up the JPC.

2.2.1 Quantizing the JPC

When we move on from the JPC device itself to actual pumping and operating the device, we need to rewrite our equation for the JRM in terms of the bosonic raising and lowering operators, much like we did for the transmon [79]. Adding the resonators to our JRM is equivalent to placing 3 harmonic oscillators in our system:

$$H_{resonators} = \frac{\Phi_X^2}{2L_a} + \frac{\Phi_Y^2}{2L_b} + \frac{\Phi_Z^2}{2L_c} + \frac{Q_X^2}{2C_a} + \frac{Q_Y^2}{2C_b} + \frac{Q_Z^2}{2C_c} \quad (2.17)$$

These 3 oscillators add capacitance and additional inductance to our system. Combining H_{JRM} (Eqn. 2.16) with $H_{resonators}$ gives us the total Hamiltonian of the JPC. However, for clarity, let's ignore 4^{th} order terms for now. We will revisit them in Chapter 3. Eqn. 2.16 can be rewritten in terms of the Josephson inductance, and the node fluxes can be rewritten in terms of the flux quantum, $\Phi_0 = h/2e$, such that:

$$\begin{aligned} \varphi_X &\rightarrow \frac{\Phi_X}{\Phi_0} \\ \varphi_Y &\rightarrow \frac{\Phi_Y}{\Phi_0} \\ \varphi_Z &\rightarrow \frac{\Phi_Z}{\Phi_0} \\ \varphi_{ext} &\rightarrow \frac{\Phi_{ext}}{\Phi_0}. \end{aligned} \quad (2.18)$$

Now, the 2^{nd} order terms in Eqn. 2.16 combine with Eqn. 2.17 to give quantized harmonic oscillators that can be written in terms of the bosonic modes, as in [68]:

$$\begin{aligned}
\Phi_X &= \sqrt{\frac{\hbar Z_a}{2}}(\hat{a} + \hat{a}^\dagger) \\
Q_X &= i\sqrt{\frac{\hbar}{2Z_a}}(\hat{a} - \hat{a}^\dagger) \\
\Phi_Y &= \sqrt{\frac{\hbar Z_b}{2}}(\hat{b} + \hat{b}^\dagger) \\
Q_Y &= i\sqrt{\frac{\hbar}{2Z_b}}(\hat{b} - \hat{b}^\dagger) \\
\Phi_Z &= \sqrt{\frac{\hbar Z_c}{2}}(\hat{c} + \hat{c}^\dagger) \\
Q_Z &= i\sqrt{\frac{\hbar}{2Z_c}}(\hat{c} - \hat{c}^\dagger)
\end{aligned} \tag{2.19}$$

Making all of these substitutions and grouping terms gives the familiar Hamiltonian:

$$\frac{H_{JPC}}{\hbar} = \omega_a a^\dagger a + \omega_b b^\dagger b + \omega_c c^\dagger c + g_3(a + a^\dagger)(b + b^\dagger)(c + c^\dagger), \tag{2.20}$$

where the quantum hats have now been left off in the name of brevity. All operators should still be considered quantized. g_3 represents the strength of the 3-body coupling,

$$g_3 = \sqrt{\hbar \frac{p_a p_b p_c \omega_a \omega_b \omega_c}{E_J^{eff}}}, \tag{2.21}$$

where $p_{a,b,c} \approx L_{Jo}/L_{a,b,c}$ are the participation ratios of the inductances that arise from the junctions themselves, as well as the linear, geometric inductance from each mode. E_J^{eff} is the effective Josephson energy, which is flux dependent.

This quantized JPC Hamiltonian shows 3 quantum harmonic oscillators that are coupled together, and it is this last term that is responsible for every interesting interaction in the JPC. Furthermore, the strength of these couplings depends on the external magnetic flux, and is periodic.

3.0 Parametric Couplings

This chapter will be a more in-depth look at many (but not all) of the possible parametric couplings available in the JPC, or in any system with an all-to-all 3-body coupling. It will begin with singly-pumped parametric processes; gain and conversion, and move on to multiparametric interactions including GC, GCI, and GG (which are each named for their pump configuration). This chapter is entirely theoretical, with mathematical descriptions of the scattering parameters that arise from these pumping schemes. To best compare all the multiparametric pumping schemes, the bandwidths of all the modes are held constant throughout ($\kappa = 20$ MHz) and all gain profiles achieve 20 dB. This chapter aims to show that, at least in theory, turning specific couplings on and off through a choice of pump frequencies and powers can greatly enhance the performance of an amplifier.

3.1 Gain

The Hamiltonian of the JPC up to 3rd order is given by:

$$H_{JPC} = \hbar\omega_a a^\dagger a + \hbar\omega_b b^\dagger b + \hbar\omega_c c^\dagger c + \hbar g(a + a^\dagger)(b + b^\dagger)(c + c^\dagger). \quad (3.1)$$

This way of grouping terms in the Hamiltonian allows us to see there are a great deal of possible 3-body couplings intrinsically available to the user. One can change the type of parametric process with a specific choice of pump frequency, as the terms that are not at the same frequency rotate away in the Rotating Wave Approximation (RWA). However, there are two typical modes of operation in the JPC: gain and conversion. Gain is the process by which pump photons are turned into signal and idler photons and increase the magnitude of the signal output from the device. The standard mode of amplification in the JPC is phase-preserving gain, created when the common mode, c , is pumped at the sum of the other two modes resonant frequencies: $\omega_p = \omega_a + \omega_b$

Under the rotating wave approximation (RWA) this leaves us with:

$$H_{RWA}^{gain} = \hbar\omega_a a^\dagger a + \hbar\omega_b b^\dagger b + \hbar\omega_c c^\dagger c + \hbar g(a^\dagger b^\dagger c + abc^\dagger). \quad (3.2)$$

In order to account for dissipation, loss, and incoming photons, we turn to the quantum Langevin equation,

$$\frac{d\hat{a}}{dt} = \frac{i}{\hbar}[\hat{H}, \hat{a}] - \frac{\kappa}{2} + \sqrt{\kappa}\hat{a}^{in}(t), \quad (3.3)$$

which allows us to calculate the full dynamics of the system, as κ is the linewidth of the mode. The first term on the right-hand side of the equation is the typical Heisenberg equation of motion. The second and third term represent the damping caused by the coupling to the transmission line and the incoming drive term, respectively. \hat{a}^{in} is the traveling photon amplitude wave for incoming signals. The bosonic operator a is related to the incoming and outgoing fields by [80],

$$\hat{a}^{in} + \hat{a}^{out} = \sqrt{\kappa}\hat{a}. \quad (3.4)$$

Generalizing these equations for all modes allows us to calculate the equations of motion for the system, as in [34]:

$$\dot{a} = -i\omega_a a - \kappa_a/2a - ig_3 b^\dagger c + \sqrt{\kappa_a}a^{IN} \quad (3.5)$$

$$\dot{b} = -i\omega_b b - \kappa_b/2a - ig_3 a^\dagger c + \sqrt{\kappa_b}b^{IN} \quad (3.6)$$

We want to solve for the steady-state solution, meaning the left hand side of the equations will go to zero. In order to read off the scattering parameters from these equations of motion, we must first apply the stiff pump approximation. When the third mode, c is driven significantly far off resonance from the other two modes, this quantum operator can be exchanged for the average number of pump photons in the resonator: $\langle c(t) \rangle = \sqrt{n}e^{-i(\omega_a + \omega_b +)t}$. Combining this

and the input/output equations allows us to derive the scattering parameters for the gain mode of operation on resonance:

$$\begin{bmatrix} a^{out}(\omega_a) \\ b^{out\dagger}(\omega_b) \end{bmatrix} = \begin{bmatrix} \sqrt{G_0} & -ie^{-i\phi_P}\sqrt{G_0-1} \\ ie^{i\phi_P}\sqrt{G_0-1} & \sqrt{G_0} \end{bmatrix} \begin{bmatrix} a^{in}(\omega_a) \\ b^{in\dagger}(\omega_b) \end{bmatrix}, \quad (3.7)$$

where ϕ_P is the pump phase and the $a^{in/out}$ and $b^{in/out}$ parameters represent the incoming and outgoing signals in each mode. Note, this form necessarily implies that gain is created in both reflection and transmission in this mode of operation [81]. The gain, in voltage, can be written as:

$$\sqrt{G_0} = \frac{1 + |\rho_0|}{1 - |\rho_0|}, \quad (3.8)$$

where ρ_0 is a dimensionless pump parameter given by:

$$\rho_0 = \frac{2g_3\sqrt{n_c}e^{-i\phi_P}}{\sqrt{\kappa_a\kappa_b}}. \quad (3.9)$$

As one can tell by inspection, the gain diverges as ρ_0 approaches unity, making the gain process an example of positive feedback. The gain profile itself is a Lorentzian centered about the idler and signal resonant frequency,

$$G_0(\Delta\omega) = G_0 \frac{1}{1 + \left(\frac{\Delta\omega}{B(G_0)/2} \right)^2}, \quad (3.10)$$

where $\Delta\omega$ is the difference between the frequency and the resonant frequency. This shows that the bandwidth of the gain profile is given by:

$$B(G_0) = \frac{2\kappa_a\kappa_b}{\kappa_a + \kappa_b} G_0^{-1/2} \quad (3.11)$$

κ_a and κ_b are the bandwidths of the resonator modes. Equation 3.11 shows the gain-bandwidth product of the Lorentzian line shape, meaning that the bandwidth is inversely proportional to the strength of the gain and vice versa. The constant in the equation is essentially an average of the two bandwidths involved in the process. The theoretical reflection and transmission profiles are plotted below in Fig. 9 .

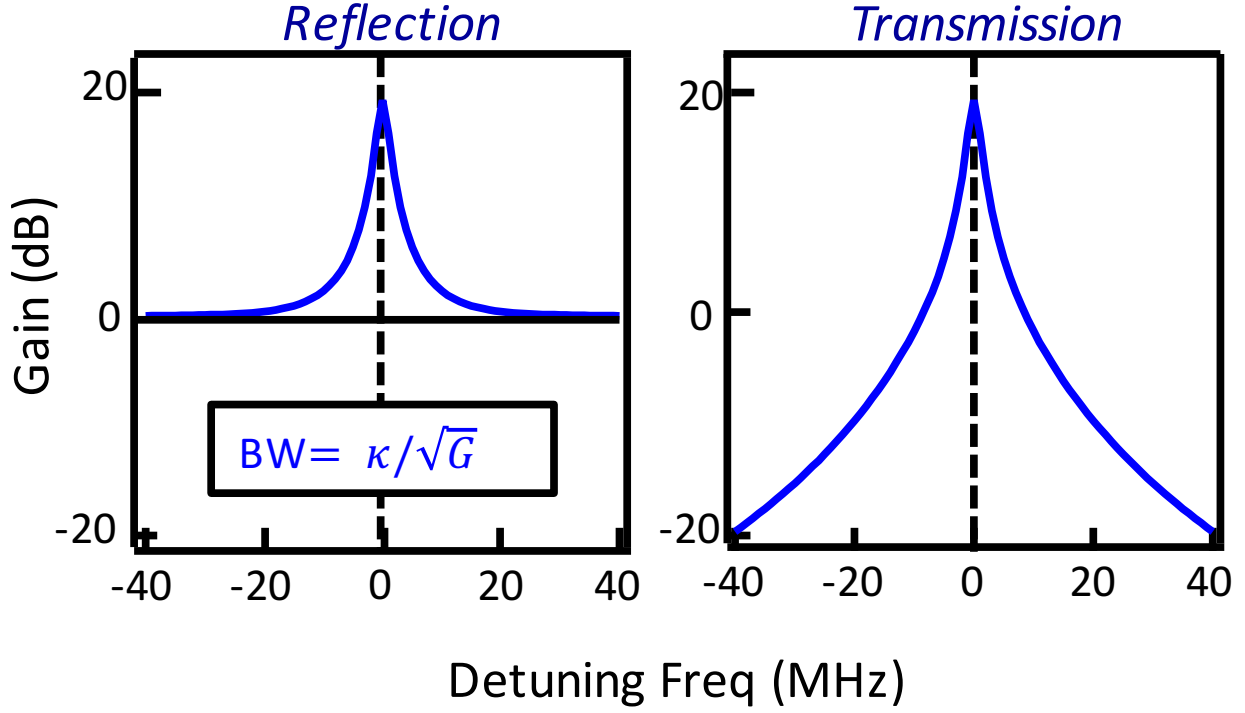


Figure 9: Theoretical plots of a one-pump, phase-preserving gain process in reflection and transmission as a function of the input frequency. The bandwidth (BW) is inversely dependent on the square root of the gain, G , leading to the gain-bandwidth product. It is also dependent on the effective average of both modes involved, a and b , where $\kappa = 2\kappa_a\kappa_b/(\kappa_a + \kappa_b)$.

3.1.1 Phase-preserving and phase-sensitive amplification

Gain can be broken down further into two broad categories of either phase-preserving or phase-sensitive amplification. Although the JPC is typically operated in the phase-preserving regime, the phase-sensitive mode works equally as an amplifier and may even be preferable in some instances.

Quantum mechanics does not allow us to measure both the In-phase (I) and Quadrature-phase (Q) of the signal perfectly at once, due to the Heisenberg uncertainty relation. In the

phase-preserving mode, quantum noise is equally distributed amongst both quadratures (see Fig. 10 (a) [82]. In the ideal phase-sensitive case, however, one quadrature is amplified by \sqrt{G} , while the other is deamplified by $1/\sqrt{G}$ [83, 84]. This creates a one mode squeezing effect, which is visualized as an ellipse signal shape in the I-Q plane, shown in Fig. 10 (b) [39].




Figure 10: Theoretical plot in I-Q space of phase-preserving (a) vs. phase-sensitive (b) amplification. When amplified by gain G , the phase-preserving case adds noise equally to both quadratures, equal to $1/2$ a photon, on average. When amplified by the same gain, the phase-sensitive case noiselessly amplifies one quadrature (in this case, I) at the expense of squeezing the other (here, Q).

In (a), a coherent state with amplitude α is amplified with a gain of $\sqrt{G_0}$. The phase of the original coherent state remains the same after the amplification. The quantum noise ($\sigma^2 = \sigma_I^2 + \sigma_Q^2 = 1/2$) is added to the radius of state. This means that the SNR deteriorates

by a factor of 2 in this process. In contrast, (b) shows the same coherent state α amplified through a phase-sensitive process by the same amplitude factor. This time, one quadrature (in this case, I) will be amplified by $\sqrt{G_0}$ while the other quadrature (Q) will be deamplified by $1/(\sqrt{G_0})$ [85]. The elliptic shape of the signal highlights how the overall SNR in each quadrature stays constant.

This phase-sensitive type of amplification can be achieved in the JPC when the common mode is now pumped at twice that of each of the other two modes: $\omega_P = 2\omega_a$ or $\omega_P = 2\omega_b$. This creates a phase-sensitive regime, as the two signal or two idler photons can now interfere with one another in the same physical mode. This scheme might be preferable if all the information from a signal is contained in one quadrature, if one wishes to keep track of the phase in an experiment, or if phase is a variable parameter throughout.

3.2 Conversion

The other one-pump mode commonly used in the JPC is coherent frequency conversion from one mode to another [86, 74]. To convert photons between two modes, the pump frequency now must be the difference of the other two resonant frequencies: $\omega_P = \omega_a - \omega_b$. The Hamiltonian that describes this is given by,

$$H_{RWA}^{Conv} = \hbar\omega_a a^\dagger a + \hbar\omega_b b^\dagger b + \hbar\omega_c c^\dagger c + \hbar g(a^\dagger bc + ab^\dagger c^\dagger), \quad (3.12)$$

which has already taken into account the RWA, so only one 3-wave coupling term remains. Once the stiff pump approximation is applied, the scattering parameters can be found in a similar fashion to the gain,

$$\begin{bmatrix} a^{out}(\omega_a) \\ b^{out}(\omega_b) \end{bmatrix} = \begin{bmatrix} \sqrt{C} & ie^{-i\phi_P}\sqrt{1-C} \\ -ie^{i\phi_P}\sqrt{1-C} & \sqrt{C} \end{bmatrix} \begin{bmatrix} a^{in}(\omega_a) \\ b^{in}(\omega_b) \end{bmatrix}, \quad (3.13)$$

where C is related to the dimensionless pump parameter as follows:

$$\sqrt{C} = \frac{2|\rho_0|}{1 + |\rho_0|}. \quad (3.14)$$

As depicted theoretically in Fig. 11, conversion theoretically produces a deep dip in reflection, meaning that the port is matched and essentially has no light exiting that port. In transmission, perfect conversion looks like a broad peak that reaches zero in dB, or unity in linear units. This corresponds to all the signal going from the input port to the output without adding additional noise. This 0 dB point is especially handy because it allows us to normalize signal in transmission, since the VNA shows real losses in the system and not actually 0 dB. We call this point "full conversion".

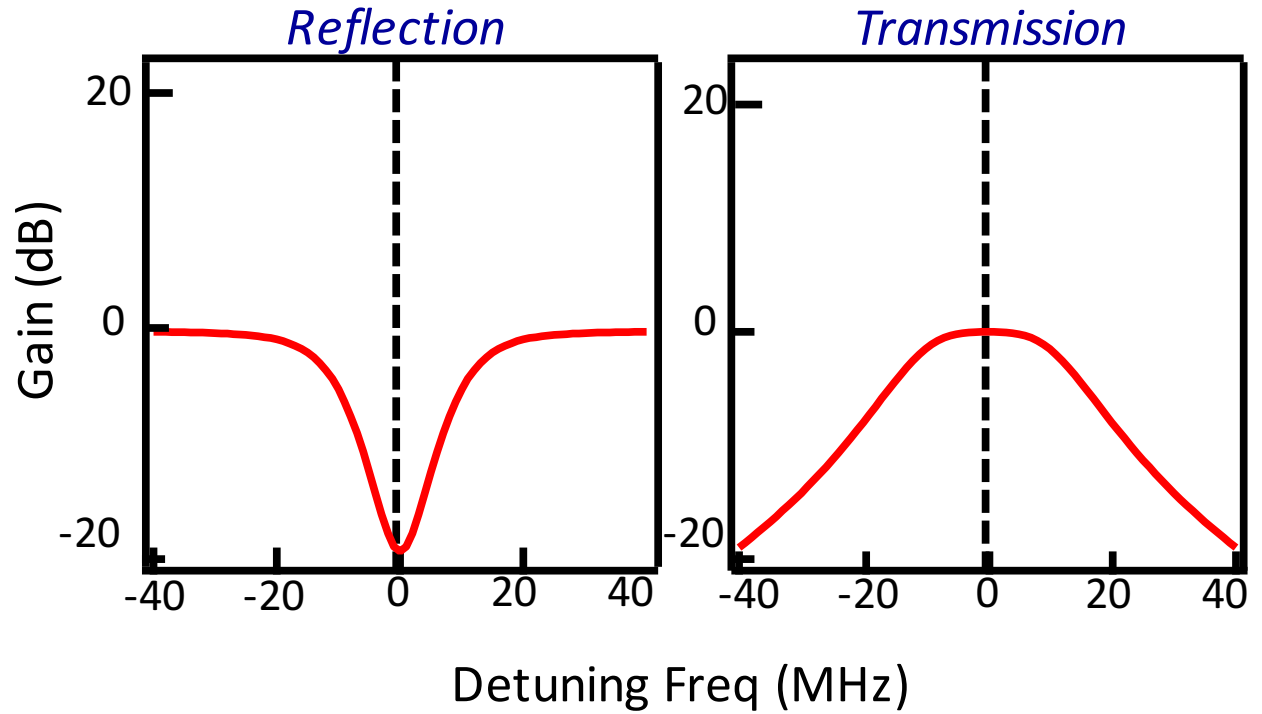


Figure 11: Theoretical plots of "full conversion" in both reflection and transmission as a function of input frequency. In reflection, a deep dip is produced for strong conversion, which corresponds to a broad peak that goes to 0 dB in transmission.

3.3 GC amplification

GC amplification is a special mode of amplification that was first proposed by A.A. Metelmann [87, 59, 88]. As the name implies, it requires one gain pump and one conversion pump between the same modes to be applied to be applied to the third spacial mode simultaneously, as seen in Fig. 12.

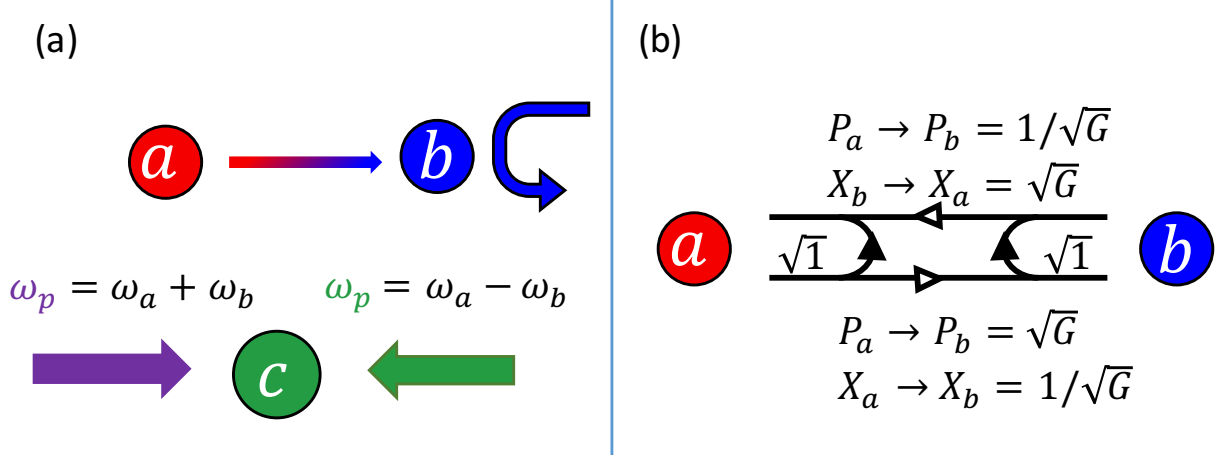


Figure 12: (a) Theoretical mode diagram from GC amplification, showing that if mode c is pumped with a gain pump and a conversion pump simultaneously, gain is created in transmission and unity in reflection. (b) Scattering matrix parameters in terms of the quantum variables X and P , showing that GC is a phase-sensitive amplification process, meaning that at the phase-sensitive point, one quadrature is amplified at the expense of the other.

This combined gain and photon hopping interaction is described by the coupling Hamiltonian:

$$H_{GC} = \hbar g_G (a^\dagger b^\dagger + ab) + \hbar g_C (a^\dagger b + ab^\dagger), \quad (3.15)$$

where g_G represents the strength of the gain pump and g_C represents the strength of conversion. This leads to the scattering matrix (where we have switched the basis to $a, b, a^\dagger, b^\dagger$

to highlight asymmetry):

$$S(\omega) = \begin{bmatrix} R(\omega) & 0 & T(\omega) & r_G T(\omega) \\ 0 & R(\omega) & -r_G T(\omega) & -T(\omega) \\ T(\omega) & r_G T(\omega) & R(\omega) & 0 \\ -r_G T(\omega) & -T(\omega) & 0 & R(\omega) \end{bmatrix} \quad (3.16)$$

where r_G is the ratio of the gain and conversion strengths: g_G/g_C . We use a slightly different convention for the scattering matrix here, where the 4 ports now represent $a, b, a^\dagger, b^\dagger$.

The reflection and transmission amplitudes are given by:

$$R(\omega) = \frac{C(1 - r_G^2) - \left(1 + i\frac{2\omega}{\kappa}\right) \left(1 - i\frac{2\omega}{\kappa} + \frac{\delta\kappa}{\kappa}\right)}{C(1 - r_G^2) + \left[1 - i\frac{2\omega}{\kappa} \left(1 - i\frac{2\omega}{\kappa} + \frac{\delta\kappa}{\kappa}\right)\right]}. \quad (3.17)$$

$$T(\omega) = \frac{i2\sqrt{C}\sqrt{1 + \frac{\delta\kappa}{\kappa}}}{C(1 - r_G^2) + \left(1 - i\frac{2\omega}{\kappa}\right) \left(1 - i\frac{2\omega}{\kappa} + \frac{\delta\kappa}{\kappa}\right)}, \quad (3.18)$$

where the so-called coopertivity is given by: $C = \frac{4g_G^2}{\kappa^2}$, and we set $\kappa_1 = \kappa$ and $\kappa_2 = \kappa_1 + \delta\kappa = \kappa + \delta\kappa$. GC amplification also requires equal pumping strengths between gain and conversion, meaning we can set $r_G = 1$.

Calculating the square of the transmission and reflection amplitudes gives us the gain as a function of frequency:

$$G(\omega) = T(\omega)^2 = \frac{4C \left(1 + \frac{\delta\kappa}{\kappa}\right)}{\left[\left(1 + \frac{\delta\kappa}{\kappa}\right)^2 + \frac{4\omega^2}{\kappa^2}\right] \left(1 + \frac{4\omega^2}{\kappa^2}\right)} \quad (3.19)$$

$$R(\omega)^2 = 1 \quad (3.20)$$

From 3.19 and 3.20 we can see how this pumping scheme realizes an amplifier that produces gain only in transmission and has unity in reflection. Furthermore, although not initially obvious, the interference of the two pumps destroys the gain-bandwidth profile. The

bandwidth, $\Delta\omega$ is now independent of the gain strength, and only a function of the line width of the resonator.

We can also better understand this amplifier by examining the scattering terms graphically. In reflection, it is easy to see how a peak added to a dip would add to zero, as is shown in Fig. 13. However, in transmission, this multiparametric scheme produces an even broader bandwidth than a single gain profile alone. The conversion pump converts the frequency of photons between the two modes in such a way that it spreads out the range of where gain occurs in frequency space. If both decay rates, κ_a and κ_b are equal, then, $\Delta\omega = \sqrt{\sqrt{2} - 1}\kappa$.




Figure 13: Theoretical plot of GC amplification in transmission (purple) and reflection (grey) as a function of input frequency. Note that gain is only created in transmission, and has a bandwidth that is proportional to the κ of the resonator alone, while unity gain is reflected back. In this theoretical plot $\kappa = 20$ MHz to best compare to experimental data.

This is a superior amplifier compared to a single gain pump, as it destroys the gain-bandwidth product and no longer amplifies in reflection. Furthermore, this is a phase-sensitive mode of amplification, which allows the user to keep track of phase in the experiment. It also remains quantum-limited, and adds no extra noise than a single gain pump alone. The measurement efficiency and discussion of this amplifier is addressed further in

3.4 GCI amplification

The GCI amplifier is very similar to the GC amplifier and is governed by the same Hamiltonian (see Eqn. 3.15). The only difference is now $g_C > g_G$. To account for the new variance in r_G , it is easiest to first set the decay rates equal to one another, i.e. $\delta\kappa = 0$. Note that for stability reasons, we confine $r_G < \sqrt{1 + \frac{1}{C}}$. Reflection and transmission amplitudes can now be written as:

$$R(\omega)^2 = \frac{\left(\frac{4\omega^2}{\kappa^2} + 2 - \epsilon\right)^2}{\frac{16\omega^4}{\kappa^4} + \frac{8\omega^2}{\kappa^2}(2 - \epsilon) + \epsilon^2} \quad (3.21)$$

$$T(\omega)^2 = \frac{4C}{\frac{16\omega^4}{\kappa^4} + \frac{8\omega^2}{\kappa^2}(2 - \epsilon) + \epsilon^2}. \quad (3.22)$$

where we parameterize $r_G^2 = 1 + \frac{1 - \epsilon}{C}$. $\epsilon = 1$ gives us the GC amplifying regime while $\epsilon = 2$ gives us the GC Imbalance regime that is the focus of this section. On resonance $R(0)^2 = 0$ for this regime, meaning that reflections are negated. In dB units, this would look like an infinite dip corresponding to a fully matched amplifier. The bandwidth of the gain in transmission becomes: $\Delta\omega = \sqrt{2}\kappa$. This makes GCI's bandwidth 2.2x that of GC's.

Even though GCI is quite similar to GC, it has some very important differences which make it an all around better amplifier. Both pumping schemes do away with the gain-bandwidth product, but GCI has a broader, usable bandwidth. Further, GC does not amplify in reflection, but it still sends unity signal back. GCI is theoretically matched, and shouldn't produce reflections at all.

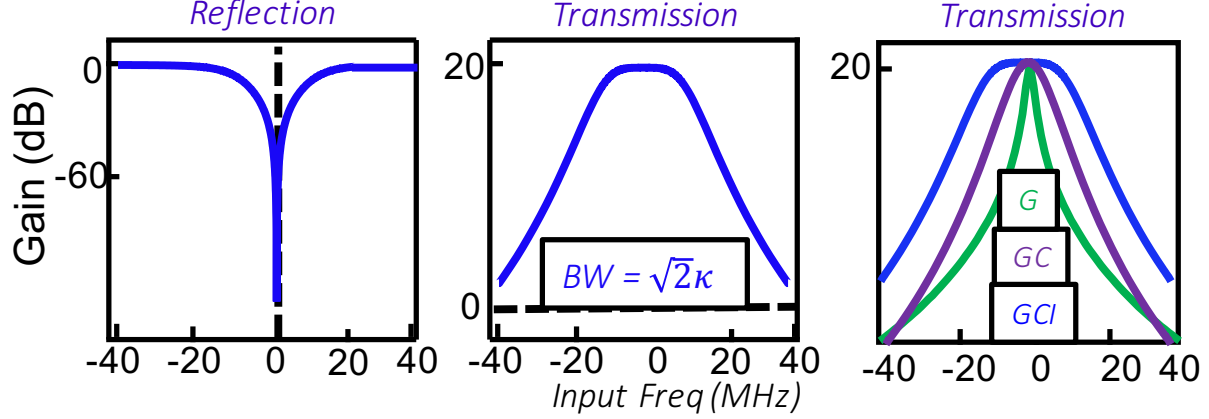


Figure 14: Theoretical plot of GCI amplification in reflection, transmission, and in comparison to the GC amplifier mode and the singly-pumped G amplifier mode. In reflection the signal becomes a deep dip, showing that it is a matched port. In transmission, we see a bandwidth that is a broad, flat peak, and only dependent on κ , much like the GC amplifier. However, in the last pane, it is clear that GCI has a significantly broader bandwidth than any other mode discussed in this section. The linewidth of the modes are both given by $\kappa = 20$ MHz.

3.5 GG amplification

GG amplification is created when two gain frequency pumps are detuned from the center frequency by $\pm\Delta$, making them $\omega_{P1} = \omega_a + \omega_b + \Delta$ and $\omega_{P2} = \omega_a + \omega_b - \Delta$, where Δ is a small frequency shift, on the order of a few MHz.

The two offset gain pumps add together to create a broadband amplifier centered about the original gain frequency. It is governed by the coupling Hamiltonian:

$$H_{GG} = \hbar g_{G1}(a^\dagger b^\dagger c + ab)e^{i(\phi_{P1} + \Delta t)} + \hbar g_{G2}(a^\dagger b^\dagger + ab)e^{i(\phi_{P2} - \Delta t)}. \quad (3.23)$$

Like the single gain tone process, this amplifier creates gain both in reflection and trans-

mission. However, due to the detuning in the pump frequencies, each tone will create multiple idler tones in both physical modes (see Fig. 15). This adds additional noise channels to the amplifier, meaning, it will add more noise to an output signal than a single gain tone alone. This is a significant hindrance of this pumping scheme; however, compared to the other modes of operation discussed in this chapter, it is the easiest by far to implement.

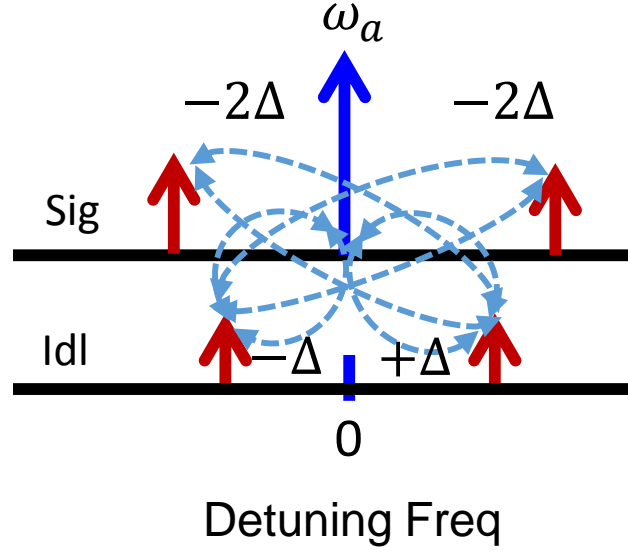


Figure 15: When GG pumping is applied, instead of having the signal couple to the idler at 0 detuning, like typical phase-preserving gain processes, the input tone creates an array of idler tones which couple to the idler mode, then back to the signal mode, and so on, at frequency intervals of Δ .

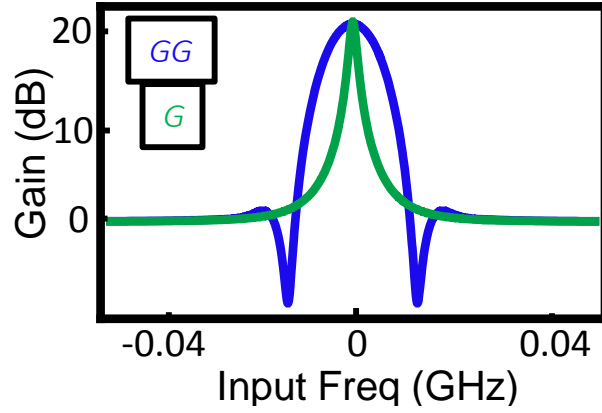


Figure 16: Comparison of the GG (blue) vs. singly pumped G amplifier (green) in reflection, showing the differences in bandwidth as a function of input frequency.

From Fig. 16 we see that the biggest benefit of this amplifying scheme is the enhanced bandwidth. It is not as broad as the GC or GCI amplifier, and more theoretical work could be done to derive an expression for this bandwidth in terms of Δ and κ . However, Mathematica predicts that the bandwidth seems to increase proportionally to Δ , until Δ reaches a significant fraction of the mode's bandwidth κ , at which point the gain diverges quickly.

4.0 Hamiltonian Theory And Design

Carefully designing a new amplifier or device is a multi step process even before any actual fabrication of the sample begins. In this rather bulky chapter I will breakdown the design stages a sample goes through before it is fabricated. In addition, I will briefly summarize the role each simulation software, such as HFSS, Microwave Office, and Mathematica, has in the design stages. These subsections will then be followed by theoretical descriptions of specific amplifier iterations we invented, such as the Meandering JPC, Wire-JPC, and Lumped JPC, to combat unwanted 4th order Hamiltonian terms. Next, I will outline all the theoretical work that has been done, so far, to incorporate the Lumped JPC into a series array, and finally, show the evolution of the amplifier series to a brand new platform we call the SHARC amplifier.

4.1 Software tools

4.1.1 Modeling amplifiers in Microwave Office

Microwave officeTM (MO) is a powerful suite of software that is capable of modeling nonlinear processes in the JPC, as well as standard microwave devices such as filters. In MO, the user can select lump elements for circuit design and test new designs and couplings by designing the circuit schematically. The two most common required elements are the lumped capacitor and inductor, whose values can be modified and changed, but usually remain in the pH and μF regime for microwave circuits.

To introduce nonlinear elements into the circuit, such as a Josephson Junction, MO offers a series of nonlinear elements. Since a Josephson Junction can be thought of as a nonlinear inductor in parallel with a capacitor, we represent it with a special type of nonlinear, lumped inductor. This element, called NLIND, can be uniquely altered to include higher order terms of the inductance. At least one NLIND element must be manually changed, and then copies

Element Options: NLIND - (Obsolete) Nonlinear Inductor Properties

Parameters | Statistics | Display | User Attributes | Symbol | Layout | Model Options | Vector

Name	Value	Unit	Tune	Opt	Limit	Lower	Upper	Step	Description
ID	L5								Device ID
SC	1		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>				Scale factor
ITH	-1000000000000000	mA	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>				Threshold current
L0	3.29106e-7/I0	nH							Polynomial coefficient
L1	0	nH	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>				Polynomial coefficient
L2	(3.29106e-7)/(2*(I0)^3)	nH							Polynomial coefficient
L3	0	nH	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>				Polynomial coefficient
L4	(3*3.29106e-7)/(8*(I0)^5)	nH							Polynomial coefficient
L5	0	nH	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>				Polynomial coefficient
L6	(5*3.29106e-7)/(16*(I0)^7)	nH							Polynomial coefficient
L7	0	nH	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>				Polynomial coefficient
L8	(35*3.29106e-7)/(128*(I0)^9)	nH							Polynomial coefficient
L9	0	nH	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>				Polynomial coefficient
L10	(63*3.29106e-7)/(256*(I0)^11)	nH							Polynomial coefficient
L11	0	nH	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>				Polynomial coefficient
L12	(231*3.29106e-7)/(1024*(I0)^13)	nH							Polynomial coefficient
L13	0	nH	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>				Polynomial coefficient
L14	(429*3.29106e-7)/(2048*(I0)^15)	nH							Polynomial coefficient
L15	0	nH	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>				Polynomial coefficient

Device ID

☒ Enable ☐ Freeze Part Number

With this addition of the modified NLIND variable, MO can be used to study nonlinear processes like gain and conversion. One of the most important reasons to use MO when designing circuits is to check for mode leakages. Photons can accidentally leak from one mode to another if frequencies are designed too close to one another, or if coupling capacitors are too large. Substantial leakages (generally ≥ -20 dB) will cause a device to behave incorrectly.

4.1.2 Ansys- HFSS

HFSSTM is Ansys’ High Frequency Structure Simulator. It was created commercially to aid in the simulation of radio antenna and complex RF circuit design, including transmission lines and filters. From a quantum information perspective, it is vital to testing new circuit designs, because although it cannot deal with nonclassical elements, it can tell us a great deal about signal reflections and leakages.

HFSS has two primary modes of operation: eigenmode and driven modal. Eigenmode offers a rougher estimate for designed frequencies and Quality factors. It is not as accurate as driven modal, but can be performed much faster, which makes it perfectly suitable as a first step in simulation. Eigenmode cannot directly tell you which circuit elements are responsible for which resonant frequency response, however, by plotting the E-field either in magnitude or vector form, and then changing which resonant frequency excitations are visible, the user can see where the density of these excitations are most concentrated.

In driven modal, wave ports to the “outside world” must now be defined, as well as the direction of input signal. For JPCs, the port connections we attach to physical microwave lines can be defined on the vacuum box. A vacuum box “lid” must typically be applied to the top of the wafer a design sits on. Any of the four sides of this vacuum box can be individually highlighted and defined as ports. Once the ports are correctly defined, driven modal can report the magnitude and phase of the signal flow in reflection or transmission through any two ports. Plotting modes in reflection is the easiest way to extract Quality factors, and plotting in transmission is the most accurate way to read leakages to other ports. This is the code that has been used to generate all the theoretical gain profiles in the previous section.

4.1.3 Mathematica

The last software tool readily at our disposal to study parametric processes is Mathematica. For specific details on the lab code used to create and model parametric processes, see [89]. The core functionality of the code though, is to input the number of defined modes in your circuit and the number of parametric processes or intrinsic couplings. The code

will then generate the S-matrix and γ -Matrix, as defined in [81]. Once these are calculated, the reflections and transmissions from each port can be plotted as a function of frequency detuning, δ , to get a good sense of the amplifier's dynamics. Important parameters to adjust are the coupling strengths (for either gain or conversion), the linewidths of each resonator mode, and the phase, if working with a phase-sensitive amplifier.

4.2 Microstrip JPC

The hybrid-free microstrip JPC is the first version of the JPC as well as the original design used in the HatLab [73, 34]. It consists of two $\lambda/2$ resonators which we refer to as the signal and idler modes, capacitively coupled to external ports, as seen in Fig. 18. The size of the microstrip resonators controls the frequency of the modes, although the common mode cannot be controlled independently, and lies somewhere between the signal and idler modes. The coupling capacitors control the linewidth of the modes, and thus their Q-factor as well. The strength of the pump port is controlled in the same way, although this capacitor is much smaller than the others, to account for the strong microwave pumps that will be applied on this port.

Zooming in on the intersection of the resonators shows the JRM, which in this older design, consists of four small Josephson junctions shunted by four larger junctions to lift the degeneracy of the ground state and get rid of hysteresis [34]. The microstrip JPC was a reliable and well-loved staple of superconducting qubit experiments [1, 40, 90]. However it was not without its problems. For one, the transmission line design creates higher order modes in the frequency spectrum of the device, which had to be identified, avoided, and designed around. In addition, we discovered theoretically, and confirmed experimentally, that higher order terms, such as the 4th-order Kerr terms do not identically cross zero in the same place in this design, which will be discussed further in the following subsection.




Figure 18: Schematic of 8-junction microstrip JPC. The length of the $\lambda/2$ transmission line resonators define the frequencies of the signal (blue) and idler (Red) modes, while small coupling capacitors fix the bandwidth of the modes. A much larger capacitor allows for the input of the pump, represented by the purple port.

4.3 4th-order nulling

One of the most important design choices for a JPC is the ratio of inductance on the inner part of the JPC ring vs. the outer ring itself. This is because for a specific range of ratios, the 4th-order Hamiltonian terms of the JPC can be cancelled. The equations shown in Chapter 2 for the various 3-wave mixing Hamiltonians are true, but do not tell the whole story. 4th-order terms have been shown to have multiple negative consequences such as saturation power suppression and resonant frequency shifts [56].

In order to cancel these unwanted terms, we no longer shunt the JRM with Josephson Junctions, but instead, use linear inductances. The effective circuit diagram of our inductively shunted JRM is depicted in Fig. 19.




Figure 19: Schematic circuit diagram of the Wire-shunt JRM (WJRM). The linear inductive shunts are labeled as L_{shunt} , and the geometric inductance in the outer arms as L_{out} .

The key ingredients of the model are (1) the Josephson junctions, (2) the shunt inductors labeled L_{shunt} , and (3) the stray inductance on the outer arms of the JRM, labeled L_{out} . The shunt inductors are typically used to control the degeneracy of the ground state of the JPC at finite magnetic flux bias. They lift the energy of states in which current flows through the shunt inductors, thus preventing the device from switching to these undesired configurations. The stray inductors are an inherent property of the aluminum traces which we use to make the outer ring of our JRM.

The circuit is described by the Hamiltonian $H_{JRM} = H_{shunt} + H_{out}$, composed of the shunt Hamiltonian

$$H_{shunt} = \sum_{i=1}^4 \frac{\phi_0^2}{2L_{shunt}} (\varphi_i - \varphi_E)^2, \quad (4.1)$$

and the outer ring Hamiltonian

$$H_{\text{out}} = \sum_{i=1}^4 H_{\text{seg}}(\varphi_{i+1} - \varphi_i - \phi_{\text{ext}}/4). \quad (4.2)$$

Here, φ_i is the superconducting phase at the i -th vertex of the JRM (see Fig. 19) and we use the notation $\varphi_5 = \varphi_1$ ¹, φ_E is the phase at the middle point which is constrained by the outer node phase by $\varphi_E = \sum_{j=1}^4 \varphi_j/4$. The phase gain due to the externally applied flux bias in each of the four outer arms of the JRM is $\phi_{\text{ext}}/4 = (2e/\hbar) \Phi_{\text{ext}}/4$. The energy of an outer ring segment phase biased to φ is

$$H_{\text{seg}}(\varphi) = \min_{\chi} \frac{\phi_0^2}{2L_{\text{out}}} (\varphi - \chi)^2 - E_J \cos(\chi), \quad (4.3)$$

and the minimization with respect to χ results in a transcendental equation that ensures that the current in the outer inductor is identical to the current in the Josephson junction.

Next, we analyze the non-linear couplings between the JPC modes. Here, we focus on the case in which the JPC has a non-degenerate ground state centered on $\varphi_1 = \varphi_2 = \dots = 0$. In this case, the phases at the JRM vertices can be expressed in terms of the canonical variables φ_a , φ_b , and φ_c that correspond to the a , b , and c eigenmodes of the JPC:

$$\begin{aligned} & \{\varphi_1, \varphi_2, \varphi_3, \varphi_4\} \\ &= \{\varphi_a + \lambda_1 \varphi_c, \varphi_b - \lambda_2 \varphi_c, -\varphi_a + \lambda_1 \varphi_c, -\varphi_b - \lambda_2 \varphi_c\}. \end{aligned} \quad (4.4)$$

The coefficients $\lambda_1 = C_3/\sqrt{C_1^2 + C_3^2}$ and $\lambda_2 = C_1/\sqrt{C_1^2 + C_3^2}$ come from the JPC eigenmode analysis and their values depend on the capacitors added to the JRM to make the JPC. We obtain the non-linear couplings between the normal modes of the JPC by taking derivatives of H_{JRM} with respect to the canonical variables. For the self- and cross-Kerr terms we find:

$$K_{aa} = \frac{1}{4!} \frac{\partial^4 H_{\text{JRM}}}{\partial \varphi_a^4} = \frac{1}{6} H_{\text{seg}}^{(4)}(\varphi_{\text{ext}}/4) = K_{bb} = K_{cc}, \quad (4.5)$$

$$K_{ab} = \frac{1}{4} \frac{\partial^4 H_{\text{JRM}}}{\partial \varphi_a^2 \partial \varphi_b^2} = H_{\text{seg}}^{(4)}(\varphi_{\text{ext}}/4) = K_{ac} = K_{bc}. \quad (4.6)$$

Observe that all of the Kerr terms are proportional to the fourth derivative of $H_{\text{seg}}(\varphi)$ with respect to φ evaluated at $\varphi = \varphi_{\text{ext}}/4$. The cosine Hamiltonian of the junction means that

¹We have assumed that the phase on the middle node is zero, details can be found in [91]

the Kerr naturally passes through zero, if the desired flux configuration holds to the required flux. More, all six Kerr terms vanish identically at a single null point. In the absence of stray inductance (i.e. $L_{\text{out}} \rightarrow 0$) the null occurs at $\phi_{\text{ext}} = 2\pi$; the null point persists and shifts towards $\phi_{\text{ext}} = 4\pi$ as L_{out} increases.

We now introduce two dimensionless parameters that we will use to characterize our circuits. The first parameter, $\beta = L_J/L_{\text{shunt}} = \phi_0^2/L_{\text{shunt}}E_J$, measures the strength of the shunt inductors relative to the Josephson junction energy. The second parameter, $\alpha = L_{\text{out}}/L_J = L_{\text{out}}E_J/\phi_0^2$, measures the strength of the Josephson junction with respect to the stray inductance. Preferably $\alpha \ll 1$ to ensure that the Josephson energy dominates the stray inductance. At the same time, β should be sufficiently large to ensure that $\varphi_1 = \varphi_2 = \dots = 0$ is the global minimum of H_{JRM} and hence the JPC has a single, well defined ground state at the null point. For the case $\alpha = 0$, the non-degeneracy of the ground state for all values of φ_{ext} is ensured by setting $\beta > 4$. As α increases, so does the minimum required value of β .

A similar analysis of the JRM with Josephson-junction based shunts (as in [34, 56]), shows that the cross-Kerr terms null together at a similar point to our linearly shunted JRM, but the self-Kerr null point is shifted to a larger flux, which can be seen in Fig. 20.

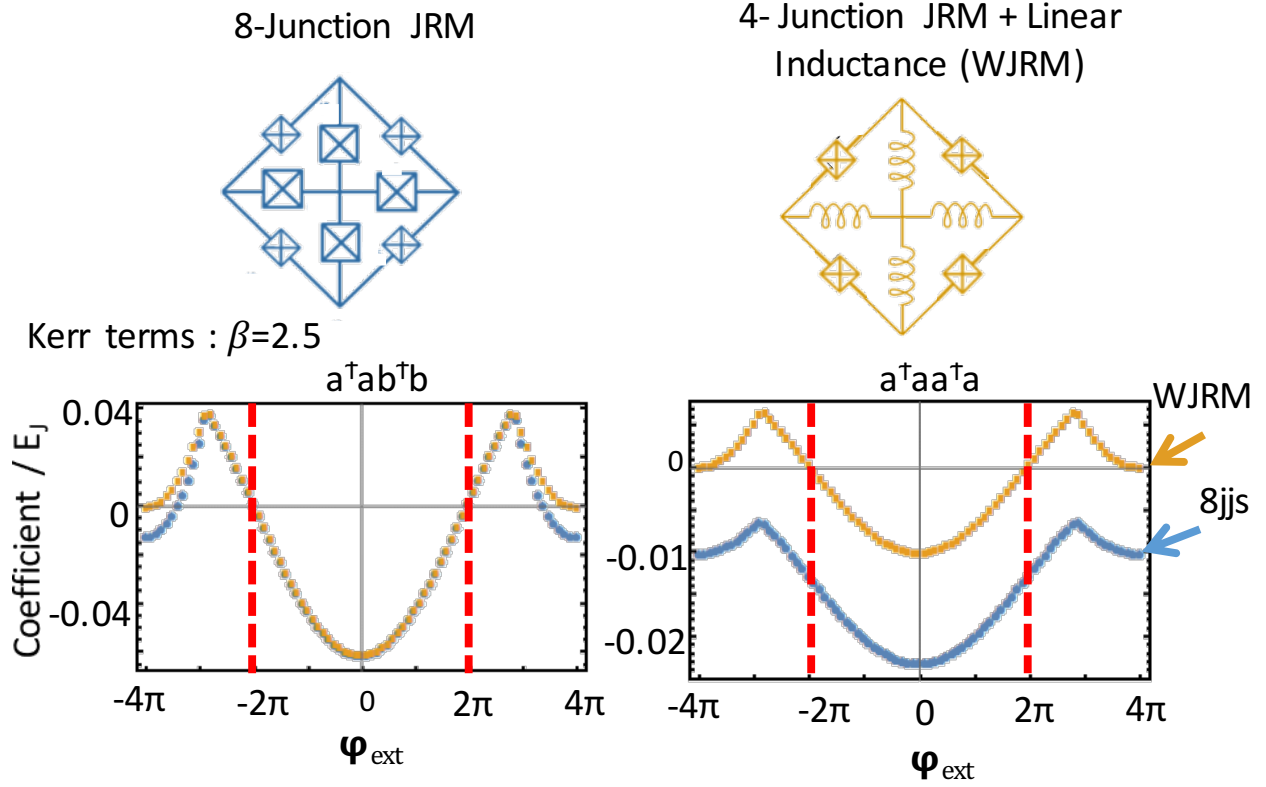


Figure 20: Theoretical plots of the strength of 4th-order cross and self Kerr terms for the 8-junction JRM and linear-shunted JRM for a β ratio of 2.5. The blue data points represent the 8-junction device terms and the yellow, the linear-shunted terms. The red dotted line shows the 0 point crossing for both cross and self Kerr terms. Note that both 4-th order terms cross 0 at the same point for the linear-shunted device but not for the 8-junction device.

As the behaviors of the various Kerr terms are very similar in their effect on device performance, the junction-shunted JRM ring will not realize the benefits of a linearly shunted JRM but will instead perform essentially no different from the un-shunted version. We have also found, in both theory and experiment, that asymmetries in junction critical current and applied bias flux to the JRMs four loops must be minimized, as they can cause substantial difference in the flux response, and hence Kerr nulling, of the JRM.

In terms of device design, this means our overall goal is to achieve $\beta > 4$ and $\alpha \approx 0$. In real life, achieving $\alpha = 0$ is of course, infeasible. However, if we can make our L_{shunt} large enough than we have mitigated this problem.

Chapter 6 will describe experimental attempts to test this 4th-order nulling theory with a Stark shift-type test.

4.4 Effects of stray inductance

The stray inductance in the outer ring of the JPC is an inherent, geometric property of the aluminum traces themselves. They can be an important factor in JPC design depending on how nonlinear the circuit is [76]. For instance, in the Wire-JPC design introduced in the previous section, the stray inductance changes the degeneracy of the ground state of the system. We then find the minima that correspond to smallest values of the Hamiltonian. As seen in Fig. 21 we plot the degeneracy of the JRM ground state as we sweep the external magnetic flux and the shunting parameter, β .

The blue region in the plots corresponds to conditions under which the JRM has a single global minimum (the JPC has a non-degenerate ground state), while the orange and pink correspond to multiply-degenerate regions. Additionally, all even order terms (beyond 2nd order) become zero when the external magnetic flux bias is set to 2π . We find that at this flux bias the ground state is non-degenerate for $\beta > 1.0$ (minimum requirement for operation of the nulled JPC). Additionally, the ground state is never degenerate as long as $\beta > 4$.

If we now include stray inductance, and keep β constant at 4.5 (close to our real life samples and greater than the cutoff of 4), we can study α as a function of external flux. As α increases, the degeneracy reappears. Further, we observe that as α increases, the Kerr nulling point moves from $\phi_{ext} = 2\pi$ towards 3.5π before it hits the degenerate area.

Other effects of stray inductance include flux modulation and resonance frequencies themselves. Obviously with a greater total inductance in the JPC, the resonant frequencies of the modes are often lower than expected, since the JPC is typically modeled as just a lumped inductor in HFSS. When modeling frequency modulation with flux, without the addition of




Figure 21: (a) Theoretical plot of the degeneracy of the ground state of the JRM without stray inductance. The blue region corresponds to a single solution, orange to a two-fold degenerate solution and pink to four-fold degeneracy. (b) The JRM's ground state degeneracy with stray inductance for $\beta = 4.5$ as a function of α and external flux. The thin green line shows the nulling point location, and we observe it shift to higher external flux values as α increases before the unstable region is encountered.

stray inductance, the modulation curves also do not reach exactly zero.

4.5 Coupling terms and major/minor lobe regions

Other theoretical parameters that are useful to examine before fabrication are the coupling terms relative strengths. It has always been important to flux bias the JPC amplifier at a place where the 3-body coupling coefficient between the a , b , and c modes (g_3) is significantly strong [92]. For example, biasing the amplifier at intervals of 4π won't work because

the term $\sin(\Phi_{ext})$ in the JRM Hamiltonian goes to zero.

Through this theoretical exploration, one interesting feature that was discovered was the presence of other 3-body couplings strengths in different parts of the flux modulation. Before the invention of the WJPC, flux sweeps of JPC's included regions we refer to as the minor lobe and major lobe, as shown in Fig. 22.




Figure 22: Example flux sweep from the signal mode of a JPC. The major and minor lobe regions are clearly visible throughout the flux modulation. One full oscillation is shown in the blue box, where the minor lobe region is highlighted by the purple dotted lines.

The major lobe, or the taller oscillation, is created when the current is flowing through the JRM in the direction one would expect, namely, all four quadrants have current flowing in the same direction together. The minor, smaller lobes, on the other hand, occur when this is not the case, and some current in a quadrant is flowing in a different direction [34], see Fig. 23.




Figure 23: Current circulation direction in the JPC in major and minor lobe region. Major lobes arise when the eigenmode solutions to the Hamiltonian create current that flows through all individual loops in the JPC in the same direction. Minor lobe oscillations arise when the current switches direction in two of the sub-loops.

Zooming in on the region outlined in blue in Fig. 22, Fig. 24 shows the theoretical, normalized strength of the 3-wave mixing coefficient g_{abc} which corresponds to achieving phase-preserving gain. It is obvious that the minor lobe region, outlined in purple dashed lines, is where this coefficient is smallest, and almost reaches zero (and reaches zero exactly for some choices of β). Thus, choosing to bias the JPC amplifier for phase-preserving gain in this region would be a poor choice.




Figure 24: The theoretical magnitude of the g_{abc} coefficient which is responsible for phase-preserving gain in the JPC for $\beta = 2.5$. The purple dotted lines correspond to the same region in Fig 22, which outlines the minor lobe area. The strength of this coefficient is very small, and often zero for certain parameters in this region.

On the other hand, in the same minor lobe region, a different 3-wave mixing term has significant strength. The g_{aac} term (which is equal to g_{bbc}) in contrast, goes to zero in the major lobe region and is strongest in the minor lobe external flux regions, as shown theoretically in Fig. 25.

However, one thing to note is that these theoretical calculations only hold true for a perfectly symmetric JRM. If the device is fabricated with any such asymmetry, deliberately or accidental, then the strength of the g_{aac} coefficient, which represents the ability to achieve phase-sensitive gain, becomes larger in the major lobe region, meaning, that it is actually possible to see this mode of amplification everywhere in a flux sweep. This is in fact, what we observe experimentally.

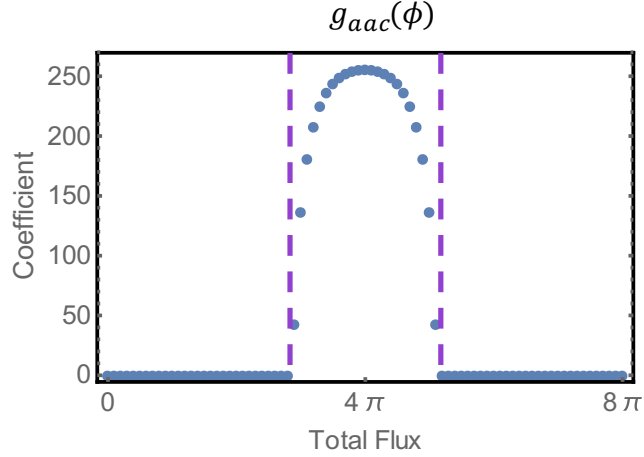


Figure 25: The theoretical magnitude of the g_{aac} coefficient which is responsible for phase-sensitive gain in the JPC for when $\beta = 2.5$. The purple dotted lines correspond to the same region in Fig. 22, which outlines the minor lobe area. The strength of this coefficient is only significant in the minor lobe region and goes to zero in the major lobe region.

4.6 Shunted JRM

As previously stated, 4^{th} order terms are known to shift the resonant frequencies of the device's modes, which can cause significant problems when trying to accurately pump the device at the sum or difference of the mode's resonant frequencies. This becomes an even bigger issue in the case of multiparametric operation, as now there are at least twice as many pumps that need to be accurate.

In order to make Fig. 19 a reality, and create a linearly shunted JRM, we tried multiple approaches. First, we designed a JRM that had long, curved, aluminum traces through the center of the devices, known as the "Meandering" JRM. This evolved into a larger JRM which instead had long, jagged shunts through the center. Last, we also designed the Wire JRM, which instead uses NbTiN shunts that feature a boost from kinetic inductance. All these designs will be discussed in this section, and their experimental results are reported in Chapter 6.

4.6.1 Meandering and lightning JRM

We began to explore how we could get rid of nonlinear shunts in the center of the JRM while still maintaining favorable ratios of α and β . The first attempt to do this resulted in the "meandering" JRM, where the large shunting junctions were replaced with zig-zagging linear shunts. The geometric inductance here was estimated from the approximate relationship of $300 \text{ pH}/\mu\text{m}$. We determined that to keep the JRM outer ring the same size, at $38\mu\text{m} \times 26\mu\text{m}$, the inner linear shunts would each have to be about $68 \mu\text{m}$ across, which can be seen in the K-layout design in Fig. 26. Since this is obviously longer than the width of the outer ring, the meandering shape became necessary. Although we had some success with this design, we were not reliably creating the β ratios that we intended. We realized that the mutual inductance created from the shunts in close proximity to one another was not being accurately simulated.




Figure 26: The K-layout design for the meandering JRM. The outer ring is $26 \times 38 \mu\text{m}^2$ and the inner shunts are each about $105 \mu\text{m}$ in length. The Josephson junctions are created from the four structures on the outer ring when they overlap in the double-angle deposition.

To erase the effects of mutual inductance, we decided to try stretching the size of the outer JRM ring itself, to about $105 \mu\text{m} \times 105 \mu\text{m}$. To minimize α in this scenario, it is necessary to increase the nonlinear inductance by decreasing the physical size of the Josephson junctions. From making this trade off, we estimated we could achieve an α of 0.35.




Figure 27: The K-layout design for the lightning JRM. The outer ring's size has been increased to be able to fit the longer, less winding shunts. In order to balance the ratio α , we also had to decrease the size of the Josephson junctions.

The lightning JRM is a perfect example of how no matter how precisely you simulate and calculate devices, it is still crucial to test them experimentally. This is because the biggest issue with the Lightning JRM is something that cannot be detected from simulation. We found that the size of the JRM became so large that the magnetic field from the bias magnet was not flowing through the sub-loops evenly. Thus, we cannot assume the ground state of the phases $\varphi_1 = \varphi_2 = \varphi_3 = \varphi_4 = 0$. This distorts the Hamiltonian, making our solutions somewhat invalid. We decided stretching the JRM size wasn't the best design, and we would need to reduce it to the original size.

4.6.2 Wire JPC

In order to minimize $\alpha = \frac{L_{out}}{L_J}$, $\beta = \frac{L_J}{L_{shunt}}$, and keep the size of the JRM ring small, we finally decided to use linear shunts made out of NbTiN instead of Aluminum. NbTiN has a relatively high kinetic inductance that was found to add a boost of ≈ 5 to the geometric inductance [91]. This added boost made the meandering or jagged shunt design no longer necessary, which eliminated the mutual inductance problems. In addition NbTiN does not grow a native oxide, so Aluminum can be deposited straight on top of the shunts and still make good contact. Fig. 28 shows the NbTiN shunts in blue, and highlights the importance of having good overlap area between the shunts and the aluminum resonators fabricated on top to ensure good contact.




Figure 28: The K-layout design for the Wire-JRM. The aluminum shunts in the center have been replaced with NbTiN (in blue) to take advantage of the inductance boost provided by the kinetic inductance. In addition, the ring size has been decreased to create a more uniform magnetic field in each sub-loop. The junction sizes remain small to have an even smaller α .

The ring size has been decreased to $20\text{ }\mu\text{m} \times 20\text{ }\mu\text{m}$ to ensure uniform magnetic flux in each of the sub-loops. The junction size was also decreased to $0.4\text{ }\mu\text{m} \times 0.5\text{ }\mu\text{m}$ to create a smaller α . The final, fabricated design can be seen in Fig. 67. Results from this amplifier can be seen in Chapter 6.

4.7 Lumped JPC

The Wire-shunted JRM is the most promising ring design at the time of this dissertation. However, even though this ring is an improvement over past designs, we still hoped to improve the JPC design as a whole, which is the focus of this subsection. First and foremost, we wanted to get rid of the transmission line based design and replace it with lumped resonators. The reasons for this are numerous: the lumped design is smaller, more flexible, and will have no higher order mode frequencies [93, 94]. A circuit schematic for this new JPC design is shown in Fig. 29.




Figure 29: The Lumped JPC (LJPC) circuit diagram, featuring the WJRM at the center. The yellow inductors represent the stray, geometric inductance, the grey, the shunting inductance. The green inductors represent the lumped resonator inductance, along with the lumped capacitors in red and blue, which form the signal, idler, and common modes of the JPC. The pump port is also coupled to the device by a coupling capacitor.

The JRM of this JPC will not need to be altered, and will remain the same WJRM with NbTiN shunts as discussed in the previous subsection. Only the resonators need to be altered. JPC's require a differential drive for correct functionality. For the microstrip design, that was easily accomplished using $\lambda/2$ resonators. For the lumped design however, the differential drive had to be created by having each resonator be tied to the same ground plane, which is not necessarily obvious in Fig. 29.

The lumped nature of this circuit also means it can be made much smaller, as seen in Fig 30. The long, microstrip resonators used to be fabricated onto chips that were 8.5 mm x 11 mm in size, but the entire LJPC fits onto a chip that is 4 mm x 4 mm. Due

to the smaller chip size, different PCB boards for the JPC box itself were required. Fig. 30 shows the design for the new resonators. The large triangle pads control the frequency of the idler (vertical) and signal (horizontal) modes. The yellow ground plane is also made out of NbTiN, albeit, a thicker layer. There exists a dielectric AlO_x layer between the ground plane and the resonator pads to create capacitors. Details of the LJPC fabrication can be found in Chapter 5.




Figure 30: The Lumped JPC (LJPC) design in HFSS with all dimensions labeled in units of μm . The blue triangles serve as the top layer of the capacitor pads and the bottom layer is the yellow ground plane. The green insulating layer sits between them. The triangle pads can be coupled to the transmission lines (orange) through a coupling capacitor or through direct ohmic contact.

4.8 Arrayed amplifier

4.8.1 Saturation power improvement

With the discovery of the GCI parametric pumping scheme, outlined in Chapter 2, it is now possible to realize an arrayed chain of amplifiers. Since GCI pumping features a matched port in reflection, it should be possible to bridge amplifiers in series or in parallel to create enhanced, bidirectional amplification.

Arrayed amplifiers are an appealing concept for several reasons. Arraying modes back to back resembles a TWPA, and it is possible such a design could be used to create a one-way directional amplifier. In addition, distributing high gain photons across multiple JRMs could improve the saturation power of the device.

There are also two options for how one might go about arraying a JPC: either in series or in parallel, as shown in Fig. 31.




Figure 31: (a) Schematic for an amplifier array in parallel. (b) Schematic for an amplifier array in series.

Arraying JPCs in parallel requires the use of both an on-chip frequency splitter and

a combiner. In addition, the parallel circuit only boosts the saturation power (in dB) by n , where n is the number of amplifiers in the array. In contrast, the series option doesn't require additional components, and also offers a much more interesting saturation power boost factor.

The saturation power of these parametric amplifiers is related to the gain by equation 4.7.

$$P_{-1dB} \sim \frac{1}{G^{3/2}} \quad (4.7)$$

Where the $3/2$ factor is an inherent property of these amplifiers that has been derived theoretically [86] and verified experimentally ourselves numerous times.

In order to deduce how to get the largest boost of saturation power from this series array, we theoretically examined what would happen if all amplifiers in the series could produce different amounts of gain, as shown in Fig. 32




Figure 32: Schematic of n JPCs in series where every individual amplifier can have a different amount of gain. All exponential factors must simply add up to one.

To demonstrate this idea, we can show the theoretical derivation with just two amplifiers in series. The first amplifier will be tuned to a gain amplitude of G_0^β and the second then must be $G_0^{1-\beta}$. We know that in the intermediate stage, before the last amplifier, the P_{-1dB}

saturation point will be equal to:

$$P_{-1dB}^{int} = \frac{A}{G_0^{(1-\beta)\alpha}} \quad (4.8)$$

where we have replaced $3/2$ with α for convenience in calculations and A is an unknown constant. To find the power that will saturate the entire chain, we recognize that:

$$P_{-1dB}^{total} G_0^\beta = P_{-1dB}^{int}. \quad (4.9)$$

Rewriting gives,

$$P_{-1dB}^{total} = \frac{A}{G_0^{\alpha(1-\beta)+\beta}}. \quad (4.10)$$

Lastly, to find how much this enhances our saturation power, we take the ratio of this new saturation point with the original one given in equation 4.7:

$$P_{-1dB}^{boost} = \frac{A}{G_0^{\alpha(1-\beta)+\beta}} \div \frac{A}{G_0^\alpha} = G_0^{\alpha-\alpha(1-\beta)+\beta} \quad (4.11)$$

If we desire the standard 20 dB of gain, then we set $G_0 = 100$ and solve for the β which maximizes the boost. To find the ideal β we plot the saturation point for each individual amplifier in the chain as a function of β and find their intersection point, as seen in Fig 33.

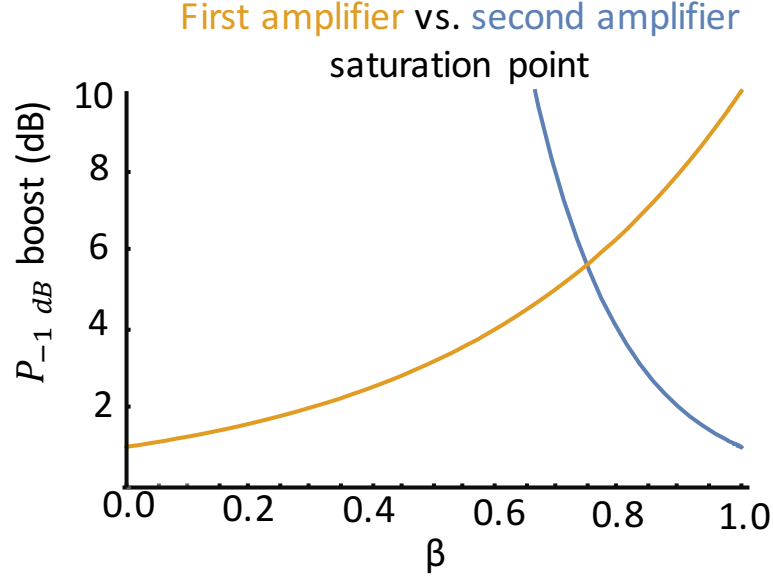


Figure 33: The yellow curve is first amplifier's saturation power (Equation 4.7) and the blue curve is the second amplifier's saturation power (Eq.4.8), both as functions of β . Their intersection point gives us the ideal value for β to achieve the maximum boost of saturation power.

For only two amplifiers in series, the ideal value of $\beta = 3/4$, meaning to achieve 20 dB gain from the entire amplifier chain, then ≈ 15 dB of gain should be achieved on the first amplifier, and ≈ 5 dB on the second.

This theoretical model can easily be extended for n amplifiers in series. Equation 4.12 gives the max saturation boost for n amplifiers with varying gain strengths:

$$P_{-1dB}^{boost} = G^{\alpha - (1 - \sum_{\beta}^n \beta)\alpha + \sum_{\beta}^n \beta)} \quad (4.12)$$

Calculating the ideal gain factors for now, say 8 amplifiers, allows us to see how effective this method can possibly be. From Figure 34, one can see that the boost continues to increase until we have about 4 amplifiers in our series, at which point the addition of more amplifiers in the chain stops helping.




Figure 34: Plotted here is the maximum possible boost of saturation power in dB as a function of the number of amplifiers in a series array.

It is clear that adding an infinite number of amplifiers in the chain has diminishing returns. However, acquiring a boost of about 5.5 dB for two amplifiers, 8.4 for three, or 9.5 for four, is both feasible and worthwhile. We will of course, first attempt this idea with the simplest case, chaining just two JPCs in series.

4.8.2 Array design

In order to create an arrayed chain of amplifiers, there must exist an innate way for signal to flow from one JPC to another. The easiest way to do this is to create an inherent conversion process between two of the modes in our system. The simplest way to do this is to have both idler modes of the two JPCs be on resonance with one another, and coupled together with a capacitor. The Hamiltonian that describes this interaction along with bi-directional GCI coupling in both JPCs, after the stiff-pump approximation, is given by Equation 4.13,

$$\begin{aligned}
\frac{H_{array}}{\hbar} = & \omega_a a^\dagger a + \omega_{b_1} b_1^\dagger b_1 + \omega_{b_2} b_2^\dagger b_2 + \omega_c c^\dagger c \\
& + g_{G1}(a^\dagger b_1^\dagger + ab_1) + g_{G2}(c^\dagger b_2^\dagger + cb_2) \\
& + g_{C1}(a^\dagger b_1 + ab_1^\dagger) + g_{C2}(c^\dagger b_2 + ab_2^\dagger) \\
& + g_2(b_1^\dagger b_2 + b_1 b_2^\dagger),
\end{aligned} \tag{4.13}$$

where the $\omega_{a,c}$ give the frequency of the signal modes and ω_{b_1,b_2} give the frequencies of the idler modes, which should ideally be on resonance with one another. It is not possible for $b_1 = b_2$ exactly because since these modes are directly interacting with one another, there will be an avoided crossing. Controlling this conversion or g_2 term that leads to this avoided crossing will be discussed further on in the section.

To further understand how these modes will interact with each other, given this Hamiltonian, please refer to the mode diagram in Fig 35.

From this picture, it is easy to see how signal should be input into the first signal mode, or mode a , pumped with GCI through to mode b_1 , converted to b_2 , and then output with GCI coupling again from mode c . In order to actualize this design, we will turn the Lumped JPC.

The Lumped JPC design lends itself perfectly to the goal of arraying amplifiers. The design is much smaller than the Microstrip JPC, and thus, multiple amplifiers can fit on the same size chip without having to redesign sample boxes or PCB boards.

As seen in Fig. 35, the goal of the chain is to have signal be input on one mode of the first amplifier and pass through both JRMs, to exit through one port of the second amplifier. This proposed design is shown schematically in Fig. 36.

The figure shows that the signal modes are both physically and spectrally distinct, however, although both JPCs have their own idler mode, they must be coupled together, and thus on resonance with one another. It is also important to note that although it is left off in the diagram, both JPCs will retain their own pump port to apply GCI coupling. Before proceeding further in HFSS simulation, it is important to keep in mind that the signal frequencies should be spaced such that both gain pumps and conversion pumps (four in total) are at least 1 GHz away from one another to avoid frequency crowding.




Figure 35: Mode schematic of 2 JPCs arrayed in series. Signal should be input on one of the signal modes (a), couple to the first idler mode ($b1$) through GCI coupling, convert to the second idler ($b2$) when tuned on resonance and then couple back to the second signal mode (c) for the output.

The two JPCs were designed on the same $4 \times 4 \text{ mm}^2$ chip, with the same fabrication and design structure as the Lumped JPC, as seen in Fig. 37.

The key elements visible in this design are the size of the capacitor pads, which control the resonant frequencies of all modes, the coupling capacitors to the ports, which control κ and Q , and the coupling capacitors in the center, which control the magnitude of the g_2 coupling in Equation 4.13 and hence, the inherent photon conversion, or hopping.

To get a better sense of how big each of these parameters should be, Mathematica was used to sweep through a large range of coupling strengths that corresponded to physically feasible designs, while holding the conditions necessary for 20 dB of GCI gain constant. Specifically, the assumptions we held were:




Figure 36: Schematic of two WJPCs arrayed in series. The modes signified by the red and purple ports and capacitors represent the signal and the input/output modes. The blue capacitors represent the idler ports with a coupling capacitor in green linking them. The pump port through the idler modes is not explicitly shown.

$$g_{G1} = \frac{1}{2} \sqrt{4g_{C1}^2 - 2\kappa_a g_2} \quad (4.14)$$

$$g_{G2} = \frac{1}{2} \sqrt{4g_{C2}^2 - 2\kappa_c g_2} \quad (4.15)$$

$$g_{C1} = g_{C2} = \frac{\sqrt{101/2} \sqrt{\kappa_a^2 \kappa_c - 20ig_2 \kappa_a^{3/2} \kappa_c^{3/2} + g_2 \kappa_a \kappa_c^2}}{\sqrt{\kappa_a^2 + 402\kappa_a \kappa_c + \kappa_c^2}} \quad (4.16)$$

4.14 and 4.15 are the conditions for GCI pumping which relates the conversion pump strength to the gain pump strength. 4.16 is the condition that was solved for when GCI pumping gives 20 dB gain in transmission. We also assumed both the gain and conversion




Figure 37: 2 Lumped JPCs in series, where the red mode indicates signal 1, the purple, signal 2, and the blue the 2 idler modes. The ports on the top and bottom of the figure are the input/output ports of the device, and the ports on the left and right, controlled by a minuscule coupling capacitor, indicate the pump ports.

pumps were equal ($g_{G1} = g_{G2}$, $g_{C1} = g_{C2}$), and in addition, that all phases are 0 as well. Figure 38 shows the resultant theoretical max bandwidth achieved in reflection and transmission for different values of $\kappa_a = \kappa_c$ and the g_2 coupling.




Figure 38: Theoretical plots of the 3-dB bandwidth in reflection and transmission of the two JPC arrayed amplifier. The linewidth of $\kappa_a = \kappa_c$ is swept vs. the g_2 coupling between the idler modes ($\kappa_{b1,b2}$). The star represents the chosen conditions, which was selected for having a good balance of fabrication feasibility and desirable properties.

From examining these plots, we selected a bandwidth for $\kappa_{a,c}$ of 100 MHz, which we knew was feasible from past LJPC designs. To get the ideal bandwidth of the amplifier in reflection, this corresponded to a g_2 of 30 MHz, producing about 100 MHz of bandwidth in reflection and transmission. This seemed like an attractive starting point so we proceeded with these parameters as the goal.

Adjusting the κ s of the modes a and c is not challenging, as their coupling capacitors act as an individual knob for this parameter. Their corresponding phase responses in reflection from HFSS can be fit and κ can be extracted. The more challenging parameter to extract is g_2 .

The magnitude of g_2 is dictated by the size of the coupling capacitors between the two idler modes on the coupling capacitor "island." The island consists of the base layer, which is separated from the rest of the chip, the insulating layer, and the two capacitor pads on

top. Since, the insulating barrier is so thin (100 nm) the top capacitor pads themselves don't exchange energy through their fringe fields directly. Instead, capacitance is created from both of the top pads coupling to the same ground pad below.

To simulate how strong this coupling, g_2 , is between the idler modes we had to simulate how an applied external field could be used to tune the modes' resonance frequencies. To approximate this effect, one of the idler modes capacitor pads was swept from a size that began smaller and ended larger than its twin's ideal size. This mimicked the external magnetic field's effect of frequency tunability.

As one can see from Fig. 38, the data is a little noisy due to the uncertainty in the HFSS convergence, but an avoided crossing is clearly visible in all three capacitor sizes tested. These avoided crossing can be fit in python, using the derived equation for this behavior, given by equation 4.17 and g_2 is extracted.

$$y = \frac{1}{2}(-x \pm \sqrt{x^2 + 4g_2^2} + 2\omega) \quad (4.17)$$

ω is the center frequency of the avoided crossing, x represents the the arbitrary scaling unit of the capacitor size, and y is the frequency. We use this equation to produce the green and blue line in the fits below:




Figure 39: HFSS sweeps of avoided crossings with fit and extracted g_2 . The size of the coupling capacitor is noted for each, and is held constant throughout the sweep. One idler mode's size is also held constant, and the other is swept through a range of sizes until the avoided crossing is evident.

Our theoretical fits match our HFSS data fairly well, and in addition, one can see that the magnitude of the individual response dips was lower (darker color) towards the center where the avoided crossing takes place. This is an indicator that our simulation is working correctly, since in real life, the response should get larger the closer the frequency is the crossing region as well.

Fortunately, one of these sweeps, namely the last one, displayed the exact avoided crossing response we hoped to see. The g_2 term that we extracted was 30 MHz, which is the value we selected from the Mathematica plots. We now have completed our first design of the the array amplifier enough to proceed to fabrication.

4.9 The SHARC amplifier

Another element that can be used to mitigate higher-order Hamiltonian terms is the Superconducting Nonlinear Asymmetric Inductive eLement (SNAIL) [48, 47]. The SNAIL consists of a superconducting loop of n large Josephson junctions, with energy E_J and one single smaller junction with energy αE_J where $\alpha < 1$. The circuit diagram that depicts this geometry is shown in Fig 40.




Figure 40: The circuit schematic of a SNAIL, shown in terms of individual junctions and its equivalent circuit symbol. The SNAIL consists of 3 (or more) large Josephson junctions in parallel with a smaller junction with an energy ratio of α .

The Hamiltonian that describes this circuit up to 4th order is given by equation 4.18,

$$H/\hbar = \hbar\omega_a aa^\dagger + g_3(a + a^\dagger)^3 + g_4(a + a^\dagger)^4, \quad (4.18)$$

where ω_a is the resonant frequency of the SNAIL, and g_3 and g_4 are the strengths of the nonlinear 3rd and 4th order terms respectively [47]. These terms can be tuned in-situ with an external magnetic field, much like the JPC. A full description of the SNAIL is outside the breadth of this thesis, this section will instead focus only on its design for use as a parametric amplifier.

The JPC and the SNAIL fundamentally differ in the number of inherent eigenmodes in the system. Unlike the JRM, the SNAIL only has one eigenfrequency, so any parametric processes or mixing in the SNAIL occur through only one mode. In order to achieve 3-wave

mixing, the device should be pumped at $\omega_p = 2\omega_a$. This can be an attractive feature, since the JPC does not allow one to independently control all 3 modes of the system. Furthermore, the SNAIL can be coupled arbitrarily to any number of linear modes and "share" its nonlinearity, and thus 3-wave coupling. There is no straightforward analogue for this in a JRM.

The idea of using the SNAIL as our nonlinear element in place of the JRM stemmed from the desire to create a fully directional amplifier. The GCI parametric coupling, as outlined in Chapter 2 is fully realizable in a JPC, and can even be used to array JPCs in series, as shown in the previous section. However this is cumbersome and proved difficult to engineer for several reasons. Instead, the SHARC took its place to circumvent some of these problems.




Figure 41: Schematic of triangle vs. array-type couplings as possible ways to achieve directional amplification. Triangle mode-type coupling relies on pump interference in a loop, which can also be used in the array-type coupling, but this geometry more closely resembles that of a small TWPA.

There are two paths one can walk down when thinking of how to design a directional parametric amplifier. In one school of thought, we have the "triangle" amplifier, which aims to create directionality and break reciprocity through an interference of mode couplings and phases in a loop, as shown in Fig 41. On the other hand, one can also think about creating an array in series, as was described in the previous section. This is not unlike the TWPA

amplifier, which relies on thousands of junctions hooked up back to back [95, 96]. For an in depth look into the triangle amplifier theory, refer to [89]. This work instead focuses on the array coupling idea.

I designed an arrayed SNAIL-based amplifier in HFSS, as seen in Fig. 42.




Figure 42: The first design of the SNAIL-based amplifier with linear array modes. The input signal should come from the left or right port, hop across the coupling capacitor "bridge", interact with the SNAIL at the center, and then leave via the same mechanism out the opposite side. The pump port for the SNAIL is created from a small coupling capacitor directly tied to the SNAIL's lumped capacitor pads at the top of the figure.

The first way to test the idea of an array amplifier with this platform is to create one input, one output, and one interior mode which does not couple to the outside world. In this design, the goal was for signal to be input on the port on either the left or right, and exit through the opposite side with no reflections. The triangle-shaped capacitor pads on either side of the image control the frequency of what I call the signal 1 and signal 2 modes, while the larger triangle pads in the middle control the inherent frequency of the SNAIL (which is not visible in this image, but is marked with a blue circle). In order to achieve bi-directional coupling, GCI amplification would have to be created across the entirety of the amplifier,

from parametrically coupling signal 1 (mode a) and the SNAIL and signal 2 (mode b) with the SNAIL. Much like the array built up from multiple JPC's, as described in the previous section, we assumed to achieve GCI arrayed-amplification, we would need two conversion and two gain pumps. However, exploring this theory using the parametric coupling Mathematica notebook showed that would we need both conversion pumps ($\omega_{p1} = \omega_s - \omega_a$, $\omega_{p2} = \omega_s - \omega_b$) but only one gain pump, between the SNAIL and the output mode. This reduction in necessary pumps made this platform all the more desirable.

In this design, the SNAIL would "share" its nonlinearity to the linear a and b modes through a g_2 "hopping" term created from coupling capacitors. The coupling capacitors are again created from islands, or metal isolated from the rest of the ground plane with an insulating layer beneath it. The size of the pads from both the SNAIL and the linear modes control the strength of the g_2 coupling, and to some extent, the g_3 coupling as well. The only pump port in this design is the port at the top of the image, which feeds directly to the SNAIL through another coupling capacitor. To achieve directionality, it is crucial for this port to be "hidden," meaning, it must have a very small bandwidth and high Q so light does not leave the array in the middle of the circuit. The theoretical scattering matrix for this design, when pumped with GCI+C pumps is shown in Fig. 43.

From the scattering matrix we can see that when these three pumps are applied to the SHARC amplifier we retain the bi-directionality from the GCI coupling, if modes a and c act as our input-output ports. Mode b acts as the high Q mode, and ideally no light should escape through this middle mode. In reality, we must balance the ratio of the mode bandwidths to create linewidths that are actually experimentally achievable.

The challenge with using a combination of linear and nonlinear elements was deciding how big to make the 3-wave coupling and 2-body coupling coefficients in order to achieve reliable photon hopping between the modes and also create parametric amplification. The relationship between the pure 3-wave nonlinearity of the SNAIL and the linear modes is given by equation 4.19:

$$g_{xss} = g_{sss} \frac{g_2}{\Delta} \quad (4.19)$$

Where g_{xss} represents the 3-body coupling to a linear mode, g_2 is the 2-body hopping




Figure 43: Theoretical scattering matrix of SHARC amplifier with GCI+C coupling. From the scattering matrix, we can see that just like with 2 mode GCI coupling, this is a bi-directional amplifier that is matched in reflection. Mode b should ideally be a hidden mode in reality, but we include it here to show that no light should escape from this mode if $\kappa_b \ll \kappa_{a,c}$.

coefficient, and Δ is the frequency separation between the SNAIL and specific linear mode (a, b, c) .

One feature that is left off in this simulation is the "healing capacitor." In real life, there must exist a slit in the ground plane that extends from the SNAIL to the corner of the entire sample so that flux can still be used to bias the SNAIL. Without it, the superconducting Meissner effect would prevent magnetic flux from being able to reach the SNAIL. We create a capacitor across this slit to hold the two sides of the ground plane to the same potential,

which is important in achieving the correct microwave design. We did not simulate the large slit or the healing capacitor in every HFSS simulation because such a large object required substantial computational time, and we found it did not effect the underlying results.

To order to change the SNAIL's frequency modulation response, one can adjust the energy ratio α . Much like the JPC, the real g_3 coefficients can be extracted by tuning up a 20-dB gain point and using the measured values for ω_a , κ , and a calibrated pump power. 4th-order, terms, on the other hand, can be found from a Stark shift measurement, which is described in Chapter 6. Much like the WJPC, the SNAIL can then be biased at the point where the frequency detunings brought on by Stark shifts disappear, and hence, the 4th-order nonlinearity becomes negligible.

For this first design, we designated $\alpha = 0.28$ to achieve a g_3 coefficient between the SNAIL (g_{sss}) of about 4.5 MHz. This gives a 3-body coupling to the linear modes of about 0.55 MHz and 0.18 MHz, which we deemed sufficient to achieve 20 dB gain based on past JPC designs. The frequencies for the linear modes were designed as: $\omega_a = 5.5$ GHz and $\omega_b = 7.5$ GHz. This ensured good separation of at least 1 GHz between all gain and conversion combinations of pump frequencies. The SNAIL was designed to modulate from about 5 GHz to 2.5 GHz. All g_3 and g_2 values were extracted using the Black box procedure.

However, we wished to create a platform that could be viable for both types of directional amplification theories; triangle and array. We designed a SNAIL-based amplifier, that we call the Superconducting Hybridized ARbitrary Cavity-based amplifier (SHARC), which couples via an intrinsic conversion term to 3 linear modes, described by the Hamiltonian:

$$\frac{H}{\hbar} = \omega_a a^\dagger a + \omega_b b^\dagger b + \omega_c c^\dagger c + \omega_s s^\dagger s + g_2^{sa}(s a^\dagger + s^\dagger a) + g_2^{sb}(s b^\dagger + s^\dagger b) + g_2^{sc}(s c^\dagger + s^\dagger c) + \dots \quad (4.20)$$

where $\omega_{a,b,c}$ are the linear mode frequencies and ω_s is the frequency of the SNAIL itself. The Hamiltonian is truncated here after the 2nd-order hopping terms for clarity, since virtually all 3-body couplings should be possible, and is in fact what this idea relies upon. The mode schematic for this design is shown in Fig. 44.




Figure 44: SHARC amplifier mode diagram, where the g_{xs} terms correspond to the conversion terms between the SNAIL and the linear modes and the Δ terms correspond to the frequency difference between the modes. Here, the coupling strength created from the coupling capacitors determines the magnitude of the g_2 terms.

The SNAIL is then coupled to the 3 linear modes by 3 separate coupling capacitors in the same way as the SNAIL-array amplifier. From here, the 3 linear modes, as well as the SNAIL's pump port, also couple back to the outside world through ports created with another coupling capacitor. This can be seen in the HFSS schematic in Fig. 45.




Figure 45: HFSS schematic of the SHARC amplifier, showing the nonlinear SNAIL mode in the middle (orange), which is coupled to 3 linear modes (seen in red, maroon, and blue) through coupling capacitors.

Since this platform is designed to be used to create a directional triangle amplifier and a directional array-type amplifier, the mode bandwidths will have to shift from the previous design. The triangle mode coupling regime requires a "dump" mode, or a very lossy mode for signal to leave the amplifier. However, the array amplifier, as previously mentioned, requires the opposite and instead needs a very high Q mode in the center, between the input and output mode, so light cannot escape. We determined that it would be best for the high Q mode to no longer be the SNAIL, since that made for the most convenient dump mode as it is also the only mode being directly pumped. The new high Q mode was designed to be the b mode, at around 7.5 GHz, and the new mode, c , designed at about 9 GHz.

For this new design, we decided to increase the strength of the g_2 couplings between the nonlinear and linear modes, as we were not seeing much gain experimentally. We had

some success pumping couplings that were designed to be around 50 MHz, although at this strength he was also seeing a lot of mode leakage. Going forward, after a few iterations, we aimed to design $g_{sss} = 50$ MHz, and $g_{a,b,c} = 5$ MHz. α was also decreased to a ratio of 0.2.

This amplifier design had many promising properties, but suffered from fairly lossy resonators, either due to the insulating barrier or the NbTiN ground plane. To circumvent this, we came up with a new, final design for the SHARC which is shown in Fig. 46.




Figure 46: Final design of the SHARC amplifier, where the resonators have been redesigned for RIE fabrication. The SNAIL has its own resonant mode, which is controlled only by the size of the finger capacitors in the center, since it already has its own inductance. The other three modes resonant frequencies are determined by the ratio of the lumped, geometric inductors and the finger capacitors.

To get rid of the NbTiN ground plane, this amplifier is to be fabricated out of pure aluminum. The fabrication of which, will be discussed in Chapter 5. Theoretically, this amplifier it is the same as the previous design, the lumped resonators are just constructed in a different way to prevent loss. The properties of the SNAIL and the resonators themselves, including κ' s and frequencies, were held constant between design iterations.

5.0 Incorporating van der Waals Superconductors Into Microwave Circuits

The following chapter will diverge from the main topic of this thesis somewhat; however, much of the theory and guiding principles behind this work are applicable in the study of parametric amplifiers and superconducting measurement in general. This section will focus on the idea of using novel 2D van der Waals (vdW) materials in combination with more traditional materials to design new superconducting circuits for use in quantum information platforms. The first section will describe making contact between 2D and 3D superconductors, and studying the physics of this interface with dc transport. The next subsection will focus on measuring the kinetic inductance of these 2D materials, namely, NbSe₂. The following chapters on fabrication and measurement incorporate many of the ideas in this chapter.

5.1 Partial edge contact with vdW materials

Two-dimensional vdW superconductors, such as $2H_a - NbSe_2$ and $2H_a - TaS_2$ [97, 98] (hereafter NbSe₂ and TaS₂), possess several appealing qualities, due to their atomic thinness. For example, kinetic inductance L_K , which results from the inertia of the charge carriers, is inversely proportional to cross-sectional area. Thus, it can be quite large in a vdW superconductor due to the fact that they can be exfoliated with only a handful of atomic layers at a time. A large kinetic inductance is often desirable in superconducting circuits like fluxonium qubits [99], or even for use as inductive shunts [100]. Additionally, these vdW materials can be easily incorporated with heterostructures such as hexagonal boron nitride (hBN) which could be used as an insulating barrier in Josephson junctions [12, 101]. Finally, their atomic thinness suppresses orbital pair breaking because they are so thin that vortices cannot form, leading to exceptionally high in-plane upper critical fields in few-layer devices [97, 98, 102].

In order to take advantage of the extraordinary properties of 2D van der Waals superconductors, the most promising pathway is to integrate the 2D materials into already well

understood circuits made from more traditional materials. This strategy relies upon making true superconducting contact between the 2D vdW and 3D material, which minimizes dissipation and allows the use of standard microwave drive and readout techniques. In this section, we present such a method demonstrating zero resistance contacts between 2D NbSe₂ and 3D aluminum. In addition, we study the magnetic-flux response of a 2-terminal device and a SQUID-type geometry [64].

We begin by describing the fabrication and how we make contact to the vdW flakes, since every other experiment we are interested in relies on this ability. All of the experiments in this section are done with the material NbSe₂. The flake begins its journey by being exfoliated in a glove box and then encapsulated between two layers of hBN. The encapsulation from the hBN protects it from the atmosphere after being removed from the glove box since NbSe₂ tends to degrade in oxygen. The flake stack is then transferred to a silicon chip and we perform a version of the “edge” contact method [103] to aluminum.

We expose the NbSe₂ to a reactive ion etch to reveal a cross-section of the hBN-NbSe₂-hBN stack. We then transfer the stack immediately into an evaporation chamber for an argon ion mill process to clean the exposed cross section of the stack, after which we tilt the sample *in situ* to the appropriate orientation and evaporate Al onto the exposed NbSe₂ edge (Fig. 47). We have successfully made devices using both single- and double-angle evaporation.

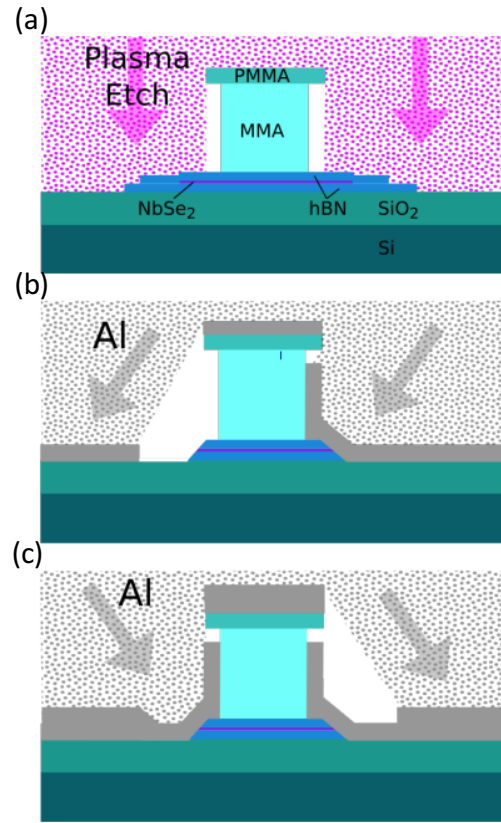


Figure 47: Schematic of TMD flake undergoing plasma etch and aluminum deposition. (a) After the lithographic mask is patterned via EBL, a reactive ion etch is used to expose a cross-section of the hBN/NbSe₂/hBN heterostructure. The stack is then loaded into a Plassys Evaporator with ION Gun where a 3 minute Ar ion mill step is used to expose a fresh facet in situ before evaporation. (b) Aluminum is evaporated in the Plassys at an angle 30 degrees from vertical. (C) Optionally, additional double-angle deposition can be used as well. (Figure created by M.Sinko)




Figure 48: Zero-resistance contact between few-layer NbSe₂ and aluminum. (a) Semilogarithmic plot of $R(T)$ showing a zero-resistance state, limited by the noise floor of roughly 135 m Ω . Inset (bottom-right): linear plot of $R(T)$ showing the Al and two NbSe₂ transitions. Inset (top-left) 4-pt measurement setup for isolating and measuring the contact resistances of the Al/NbSe₂ edge contacts (Al:gray NbSe₂:purple). NbSe₂ is approximately 1.5 μm X 3.5 μm and 10 nm thick. (b) Differential resistance (dV/dI) vs. DC current I_{DC} of same device, showing four DC critical currents 1 through 4 with $I_c \approx 56.5, 78.5, 125, 280$ μA . The temperature dependence of $I_c^{(3)}$ and $I_c^{(4)}$ identifies them as bulk aluminum and NbSe₂, respectively; $I_c^{(1)}$ and $I_c^{(2)}$ are oscillatory in magnetic field and are therefore associated with the Al/NbSe₂ contacts, which behave as Josephson junctions (see [64]).

In order to first ensure that we can indeed make zero resistance contact from 3D to 2D, we use a pseudo-4 terminal measurement configuration. This measurement configuration allows us to eliminate resistances from the voltage and current leads, but it includes the interface resistance between the 3D and 2D superconductors which is shown in Fig. 48.

Figure 48 shows a resistance measurement as a function of temperature between 2.5 K and 175 mK in a dilution refrigerator. The transition at 7 K is that of the NbSe₂ flake, since the T_c matches the bulk value of NbSe₂, which means that the flake in this case is probably thicker than 6 atomic layers [102]. Another transition occurs at 1.2 K, which we attribute to the deposited aluminum leads and the Al/NbSe₂ contacts, since this is very close to the known T_c of aluminum. After this transition, the total resistance of the device falls to zero within the noise floor of our measurement ($R < 10^{-1}\Omega$). This confirms that this sample, at least, has real superconducting contact between the two materials. We also took TEM images of a similar sample that shows these contacts, shown in Fig. 49.




Figure 49: TEM images of NbSe₂ flake with aluminum contacts (Sample 009-4). (a) Cross-sectional image of the flake with both aluminum contacts labeled on either end. (b) Cross-sectional image of the center of the TMD flake stack, showing the substrate, the hBN encapsulation, and the NbSe₂. (c) & (d) Black and white and false-color EDS map of the NbSe₂-Al contact. Interestingly, there is a small gap that was created in the stack between the NbSe₂ and the top hBN layer, see [64].

If we now expand the breadth of this investigation to study the nature of the contacts themselves, we can create a device that mimics a Superconducting Quantum Interference Device (SQUID) (Fig. 50) by creating a loop with two Al-NbSe₂ contacts.

A typical SQUID in 3 dimensions has a critical current that oscillates in a magnetic field where the periodicity is given by $\Delta B = \frac{\Phi_0}{S}$, S being the area enclosed by the superconducting loop. This interference forms a Fraunhofer pattern, and can be modeled by the equation:

$$\frac{I_c(\Phi)}{I_{c,max}} = \left| \cos \left(\pi \frac{\Phi}{\Phi_0} + \Lambda \int_C J_s \cdot dl \right) \right| \quad (5.1)$$

where $\Lambda = m/n_s e$. For a traditional SQUID, $J_s = 0$ in the interior of the 3D superconductor outlined by the contour C and so the oscillation frequency is simply determined by the loop area. However, this is not what we see experimentally.

We sweep through a range of magnetic field strengths vs. applied current and measure the differential resistance response from the SQUID. Fig. 50 shows three oscillations, with frequencies of $1/\Delta B = 2.7, 6.8$, and 9.5 mT^{-1} , which correspond, respectively, to areas $S = \Phi_0/\Delta B = 5.4, 13.6$, and $19.1 \mu\text{m}^2$. None of these areas match the physical area of the enclosed loop, shown in Fig. 50. The “beat” pattern is also inconsistent with the ratio of the Josephson junction size to the loop size [63].




Figure 50: Zero-resistance contact between few-layer NbSe₂ and aluminum. (a) Tilted SEM image of the SQUID-flake sample formed by two Al edge contacts to the NbSe₂ flake encapsulated by hBN on an silicon substrate. (b) Quantum interference pattern observed in a device C at 30 mK. A minimum resistance of $dV/dI = 81 \Omega$ has been subtracted from the data. Inset: Fast Fourier transform (FFT) amplitude of a linecut along $I_{DC} = 0$. These data are adapted from [64].

To make sense of this mismatch, we first note that in a 2D superconductor of thickness $d \ll \lambda$, where λ is the bulk penetration depth, the magnetic field can actually penetrate the entirety of the flake. This means that in comparison to a 3D Josephson junction, the effective area penetrated by the flux must be much larger in the 2D-3D junction, and the contour around which the superconducting phase winds by 2π is much larger than the physical area.

We formulated a theory based on this idea that the supercurrent actually flows through more of the flake than it would in a 3D bulk superconductor making a unique contour shape, which is highly dependent on the flake's geometry. In 3D superconductors, J_s is constrained to λ in the surface, but in 2D, it flows through the bulk, leading to almost uniform B-field penetration. The circulating supercurrent through the body/bulk of the 2D flake results in fluxoid quantization condition with a contour through the body of the 2D flake, see Fig. 51.




Figure 51: (a) Model of a 3D-3D Josephson junction. The phase difference across the junction changes by 2π when one flux quantum, Φ_0 , is threaded through the insulating barrier, shown in purple. The junction's area is $w(\lambda_1 + \lambda_2 + l)$. (b) Model of a 2D-3D JJ. The 2D flake is depicted in a circular geometry, for simplicity. In response to a field B , circulating supercurrent J_s (grey arrows) flows in the 2D bulk. As with the 3D-3D JJ, the phase difference across the junction winds by 2π from point P to point Q when one Φ_0 is threaded through the effective area of the 2D-3D junction. This requires the effective area to be bounded by a contour along which there is no additional change of the phase in the interior of the 2D flake, i.e. the contour must be perpendicular to J_s everywhere. This is satisfied by the wedge-shaped contour \overline{POQ} , and thus the effective area of the 2D-3D JJ is enclosed by the area shown in magenta (Fig. credit M. Sinko).

The theory and the simulation that backs this hypothesis are simply numerical solutions to the Ginzberg-Landau equations, and details can be found here: [64]. We find that the effective SQUID area at contact 3 is $S_{eff} = 19.7 \mu\text{m}^2$, which is very close to the experimentally calculated area of $19.1 \mu\text{m}^2$ we found from the Fourier transform of the current oscillations. This corresponds to the largest frequency, $1/\Delta B = 9.5 \text{ mT}^{-1}$ in the Fourier transform.

Contact 4 (labeled in Fig. 50) should be a 2D-3D Josephson junction, and our simulation predicts again, a larger effective area of $6.3 \mu\text{m}^2$ compared to our experimentally determined value of $5.4 \mu\text{m}^2$. The last frequency is still somewhat of a mystery, but we theorize it has to do with the phase slips through both of the junctions, giving rise to slightly different frequencies when that occurs.

Lastly, it should be noted that the residual resistance in the SQUID sample in Fig. 50 is much larger than the previous sample. This method for creating contact is not perfect, and we still experience some samples that fail outright since the source of the residual resistance is not clear. However, with improved fabrication, especially more precise ion milling, more than 60% of the devices created to date have resistances of 0.2Ω or less as temperature approaches zero.

5.2 Kinetic inductance from resonator measurement

Kinetic inductance is a direct result of the inertial response of charge carriers in a superconductor. That means that atomically thin superconductors, like vdW materials, would possess inductance not just from their planar geometry but from this kinetic response, as well. This added boost could be beneficial to many circuits that require a fixed geometry, for example, our WJPC's, which is further explored in the next chapter. However, this would require an approximate measurement of the kinetic inductance boost. To do this, we would need to be able to measure how much a vdW flake would shift an already known resonant frequency, and then subtract the geometric inductance to leave L_k as the only unknown variable.

We have estimated the kinetic inductance of these flakes to be on the order of hundreds

of fH based on the equation for kinetic inductance,

$$L_k = \frac{m}{n_c(2e)^2} \left(\frac{l}{AN} \right), \quad (5.2)$$

where n_c is the carrier density, m and e are the mass and charge of the carriers, in this case, the cooper-pairs, N is the number of atomic layers and l and A are the length and cross sectional area of the channel. For the carrier density, we used $n_c = 1.9 \times 10^{15} \text{ cm}^{-2}$, the value reported by [102]. The predicted values of all of the samples that survived fabrication are shown below:

Table 1: Calculated values for kinetic inductance

Sample	Calculated L_K (pH)
006-2	0.555
007-5	0.823
009-2	0.145
009-4	0.322
013-03	0.359

We began this project by designing a stripline the same length as the chip used in the microstrip JPC, so we could re-purpose sample boxes. We then designed multiple lumped resonators that would capacitively couple to the stripline resonator, so we could test numerous TMD flakes with one measurement setup. This geometry is called “hanger resonators” and the first design can be seen in Fig. 52.




Figure 52: 5 lumped hanger resonators of varying frequencies are coupled to a stripline so their responses can be read out in transmission. The zoomed in image of the last resonator shows the ideal placement of a TMD flake in parallel with the rest of the resonator (TMD flake not shown to scale).

It is important to space out the hanger resonators in frequency, so it is obvious which response came from which sample. For every resonator we make that would have an encapsulated vdW flake transferred into it, there also needs to be a control resonator with the same geometry (especially reproducing the shapes needed to contact to the flake which is unique to each sample) so the only unknown element generating the frequency shift is the flake inductance itself. Note that the aluminum film which replaces the film in the control sample is rather thick (100 nm) and so is assumed to have negligible kinetic inductance.

However, not all of this added inductance is due to kinetic inductance. Connecting the aluminum leads to the cross section of the TMD flake creates a Josephson junction between the Al-NbSe₂ as we explored in section 1, which comes with its own Josephson

inductance that scales inversely with the critical current. The total added inductance is given by Equation 5.3, where we have hopefully, already accounted for L_{geo} by mimicking the geometry in HFSS (and geometrically in the control sample):

$$\Delta L = L_k + L_{J1} + L_{J2}. \quad (5.3)$$

In addition to the RF hanger measurements, we also need to perform DC measurements to extract the critical currents of the Josephson junctions to calculate the Josephson inductances (Eqn. 5.4). After we subtract the inductances created from both junctions, we will be left with the value of L_k .

$$L_J = \frac{\Phi_0}{2\pi I_0} \quad (5.4)$$

Wire bonding to the sample (after rf measurement) for dc transport measurement necessitated adding large 'ears' to the device. In addition, transferring the flakes close to the resonator itself, as shown in Fig. 52, proved to be very challenging. We decided to modify the hanger geometry slightly to create DC bonding pads and more room for the transfer site away from the resonator itself. The design was modified to look like Fig. 53:




Figure 53: The modified hanger resonator design. The hanger resonator now features 2 large pads for DC wire bonding on either side. In addition, the flake outline is drawn more realistically in size, and is moved further away from the rest of the resonator to ease in transfer. The original design is left in magenta and the new details are shown in blue.

The bulk of the lumped resonator was not altered, but the transfer area, or where the flake stack is to be deposited, is more clearly outlined by alignment markers. The wirebond sites are used for a separate cooldown with dc leads attached after we measure the RF response.

Contact is made from the flake to the two aluminum leads by using the exact same method described in the previous section. In addition, we measure NbSe₂ flakes of varying thicknesses, and hope to develop a theory for how atomic thickness relates to kinetic inductance.

5.3 Fabrication

Fabrication of the resonators was initially done with a single-angle aluminum deposition E-beam process on an $8.5 \times 11 \text{ mm}^2$ silicon chip. This one step fabrication is exactly the same as the microstrip JPC fabrication and recipes can be found in Appendix A.

However, this design suffered from fabrication issues that the microstrip JPC did not. When we did lift-off in hot acetone, the aluminum between the finger capacitors would often get stuck and not lift-off. The reliability of the lift-off was so poor that we would often leave the chips in acetone for hours, which helped somewhat, but still did not alleviate all of the problems.

Since the aluminum lift-off was so inconsistent, we decided to change the entire fabrication process. We began making all of the resonators out of NbTiN using an optical process. This not only proved to be more reliable, but also allowed us to make large wafers filled with the resonator pattern all at once, instead of doing it chip by chip.

The fabrication steps that made up this process are as laid out in order below, but detailed explanations follow. Specific recipes can be found in Appendix B.

- Spin optical resist and use maskless aligner (MLA) to pattern resonators.
- Develop in 450/351 developer
- Deposit NbTiN in the AJA (creates pattern that is fuschia in Fig 53)
- Liftoff in PG remover
- Transfer the encapsulated vdW flake to to the site (area outlined with black alignment markers in Fig. 53). This step is done (with thanks!) by collaborators at CMU
- Write connecting segments (blue in Fig. 53) with E-beam
- Use side-contact procedure from previous section to expose side of flake stack
- Deposit aluminum connectors in Plassys
- Liftoff aluminum in acetone

Once the stack has been successfully exfoliated into the correct site area, the E-beam resist bilayer (MMA + PMMA) is spun on top of the stack. After the EBL is used to

write the connecting leads, the chip undergoes the reactive ion etch used to expose a cross-section of the stack. After the plasma etch, the chip is loaded into a Plassys evaporator and a short, 3 minute Ar ion mill is used to remove any contaminants. Lastly, in the same Plassys evaporator, aluminum is deposited at a 30 degree angle to make contact with the cross section of the stack. This completes the entire fabrication of the resonator with the flake. For a typical chip, there are 5 possible resonator sites on each stripline, so 5 flakes can be exfoliated onto one sample. As previously mentioned, because the geometry is so important, for every resonator that contains a flake, a control resonator that mimics the exact geometry of the leads and flake, but doesn't contain any vdW material, needs to be fabricated as well. This is done on a completely separate chip, since each resonator on the stripline is designed to be at a different frequency. The fabrication for the control is exactly the same for comparison. The finished flake resonator is shown in Fig. 54.




Figure 54: NbTiN flake resonator (orange) with aluminum leads (white). The NbSe₂ flake is in between the 2 aluminum leads and appears black in the image. The hBN protecting layer surrounds the flake in a light grey color.

5.4 Experimental results at RF and DC

Before any of the samples were cooled down in the dilution refrigerator, they were probed at DC for contact resistance. We learned that DC probing can tell us if the resonator is dead (if the contact method failed somehow in fabrication) but not necessarily if it is alive. Any flakes that were probed across the DC pads that had resistances between 20 – 60 k Ω were found to be dead. Any samples that had less resistance had the potential to have a visible signal at RF (what we call alive). Furthermore, we measured resistances that were roughly inversely proportional to flake thickness.

5.4.1 RF measurement

Once cooled down to base temperature in the dilution refrigerator, the resonators' magnitude was measured in transmission across the stripline. The hanger geometry of the flakes creates dips in the transmission response because the signal becomes phase shifted on resonance, and then reflected back along the stripline.

The quantities we are particularly interested in from this design are the internal and external, or coupling, quality factors Q . The internal Q factor is dependent on the ratio of the geometric inductance and capacitance of each resonator, and is also greatly effected by the material quality. The external Q depends on the strength of the coupling between the lumped resonators and the stripline. To get a correct estimate for both quantities, they must be approximately within an order of magnitude of one another or one term will dominate the total Q :

$$\frac{1}{Q_{tot}} = \frac{1}{Q_{int}} + \frac{1}{Q_{ext}}. \quad (5.5)$$

When we first fabricated these resonators, they were made entirely out of aluminum, and found the internal quality factors to be $\approx 3,000$. In our first design, the resonators were also 200 μm away from the stripline, which made their Q_{ext} much smaller than the Q_{int} ; thus, it was too difficult to get a reliable value. In the next version of the device we moved the resonators to 20 μm from the stripline, which increased the Q_{ext} by an order of magnitude,

making it on par with 3,000. We were then able to get a much more accurate fit and reliable parameters from the measurements.

The first resonators that were fabricated were made entirely of aluminum, as described in the previous section. The designed frequencies in HFSS were approximately 4.5, 5.0, 6.1, 6.7, and 8.0 GHz \pm 100 MHz. The first aluminum resonators were fairly close to this, and on average, their frequencies were 4.73, 5.23, 6.25, 7.25, and 8.11 GHz (without flakes). However, as mentioned before, the extracted Q-fit parameters were often in the millions, which is simply not accurate. Even though the fits looked accurate, the discrepancy between the Q_{ext} and Q_{int} was too great to extract accurate values. Furthermore, due to the fabrication difficulties, no full aluminum resonators were measured with a flake in them.

When we switched from aluminum to NbTiN resonators, there was a noticeable shift in the resonance frequencies. The solid NbTiN (control) resonators, including DC pads and leads, were: 2.55, 3.29, 3.88, 4.49, and 4.97 GHz. That means, on average, the resonators decreased in frequency by about $37\% \pm 3.6\%$

It was not totally clear why the frequencies had shifted so drastically, but we hypothesized that most of the shift was due to the added inductance of the large DC pads, which were left off in the original simulation, along with the contact leads (as in Fig 52).

However, re-simulating the resonators in HFSS with the DC pads and contact leads (as in Fig. 53) showed that the frequencies only decreased by about $20\% \pm 5.4\%$ on average. The extra discrepancy must be from the new material itself, although it remains a bit mysterious.

Nevertheless, the exact frequency of the resonators does not matter, as it is only the response shift relative to the control that is crucial. Furthermore, once we moved the resonators much close to the stripline, the extracted Q values became much more sensible. Figure 55 shows a typical flake resonator with the new NbTiN fabrication compared to its control.




Figure 55: Example response for a resonator with the flake and its corresponding control with Q-factors noted. Note the downwards shift in resonant frequency with the additional L_k .

Using this particular resonator as an example, we observe that the resonator shifted down in frequency about 62 MHz.¹ To analyze this data, I transfer this shift in frequency to a percentage compared to the original control resonator. Here this shift is about 1.4 %. To figure out what added inductance would create this percentage shift, I went back to HFSS and modeled the corresponding resonator. I drew a polygon that matched the geometry of the individual flake I was trying to simulate and assigned it to be a lumped RLC port. I started with adding 0 extra inductance and labeled this as my starting frequency. From there, I added more lumped inductance to the lumped port until the shift in the resonator frequency matched the percentage shift, not the magnitude, that I calculated experimentally. Again, it is important to use the percentage and not the magnitude because the HFSS design

¹We note that the internal quality factor has decreased moderately for both the control and flake resonator, which is not so surprising because this particular sample of NbTiN has been producing lower Q resonators in other labs for some time. If very high Q resonators are required later, the NbTiN sample should be replaced.

frequencies did not perfectly match what we were measuring in reality.

Table 2: Example resonator shift modeling in sample 006-02

Simulation Freq (GHz)	Lumped L (pH)	percentage shift
4.5651	0	original
4.5536	10	0.25
4.5425	20	0.495
4.5323	30	0.719
4.5214	40	0.957
4.5004	60	1.417

Using this specific resonator again as an example, Table 2 shows that the added lumped inductance due to the flake is about 53 pH. When the correct percentage shift was found, we used the corresponding lumped inductance assigned to the RLC port as the total ΔL for the flake.

5.4.2 DC measurement

Using this same sample as an example (Lk-NbSe2-006-2), we measured a frequency shift the corresponded to a ΔL of ≈ 53 pH. We then warmed up the sample, unloaded it, and cooled it back down in a different dilution refrigerator that was set up with DC filters. The critical current for both of the leads in this sample were found to be about $71.7 \mu\text{A}$ and the retrapping current was 52 mA, as shown below in Fig. 5.4.2:




Figure 56: dV/dI curves measured at DC. We scanned through I_0 multiple times to ensure that the critical current did not vary depending on the measurement current.

For this sample, both of the L_{JS} came out to be approximately 4.59 pH. Subtracting these from $\Delta L = 53$ lets you solve for L_k , which here is equal to ≈ 43.5 pH. Below, Table 3 shows the data and measured parameters for 5 flake samples that we measured.

Table 3: All resonator samples and parameters

Sample (pH)	Frequency (GHz)	$\Delta f(MHz)$	ΔL (pH)	$I_c(\mu A)$	L_k (pH)	l/w	thickness (nm)
006-2	3.27	45.75	53	70	43.6	9	12*
007-5	4.917	235.5	125	—	—	4	3.6*
009-2	3.24	52.7	51	56.5, 78.5	41	2.33	11.9
009-4	4.287	62.6	31	108, 128	25.4	3.8	8.8
013-3	3.6966	66.8	42	—	—	1.94	4*

*These are estimates and were not yet measured precisely with the TEM. Values are missing for the parameters of Sample 007-5 because the resonator died after RF measurement was performed but before DC measurements could be taken.

As noted in Table 11, Sample 007-5 somehow died after the RF measurement was performed. One possible explanation for this is that it sat at room temperature for too long before being transferred to the DC measurement fridge. TMDs are known to degrade in oxygen, and, for some reason, this sample proved to be particularly vulnerable. However, this was very unfortunate because this was one of our thinnest flake samples, and the fabrication on the thin flakes proved to be much more difficult than the thicker stacks. In addition, we saw a much larger Δf and ΔL in this sample, which we attribute to its highly unusual shape. This was not only the thinnest flake but the most narrow, at only $0.5 \mu\text{m}$. HFSS simulations show us that the added inductance is highly dependent on the geometry of the flake as well, which agrees with what we predicted in section 1.

In addition, sample 013-03 has not yet been measured at DC either, due to the Covid-19 pandemic. At the time of writing, the sample is still in the RF measurement fridge and not yet accessible. Hopefully, data will be collected from this flake soon, but we will proceed with the remaining 3 flakes. If we subtract the $L_{J1,J2}$ from the total ΔL we are left with our experimental estimate for the L_k of each flake. We can plot this versus the predicted estimate, as seen below in Fig. 57.




Figure 57: Predicted vs. experimental L_k values. Experimental values are ≈ 2 orders of magnitude higher than the estimated values.

It is hard to draw an all-encompassing conclusion from just these three data points; however, it is clear that the experimental results are significantly higher than what we predicted. There is some substantial shift here, shown by a crude trend line, that is worthy of future investigation.

Lastly, all data from the flakes and their controls are shown below. Their quality factors are listed as well. On average, there is not a discernible difference between the quality factors of the control versus the flake resonators. This is excellent, because it implies we are not degrading the Q of the resonators when we embed a flake. Although these quality factors are not yet high enough to be used as a qubit-type element, this could probably be improved by replacing the NbTiN sample. What is even more promising, is that both the Q_{int} and the size of the critical currents are already suitable for use in paramps.

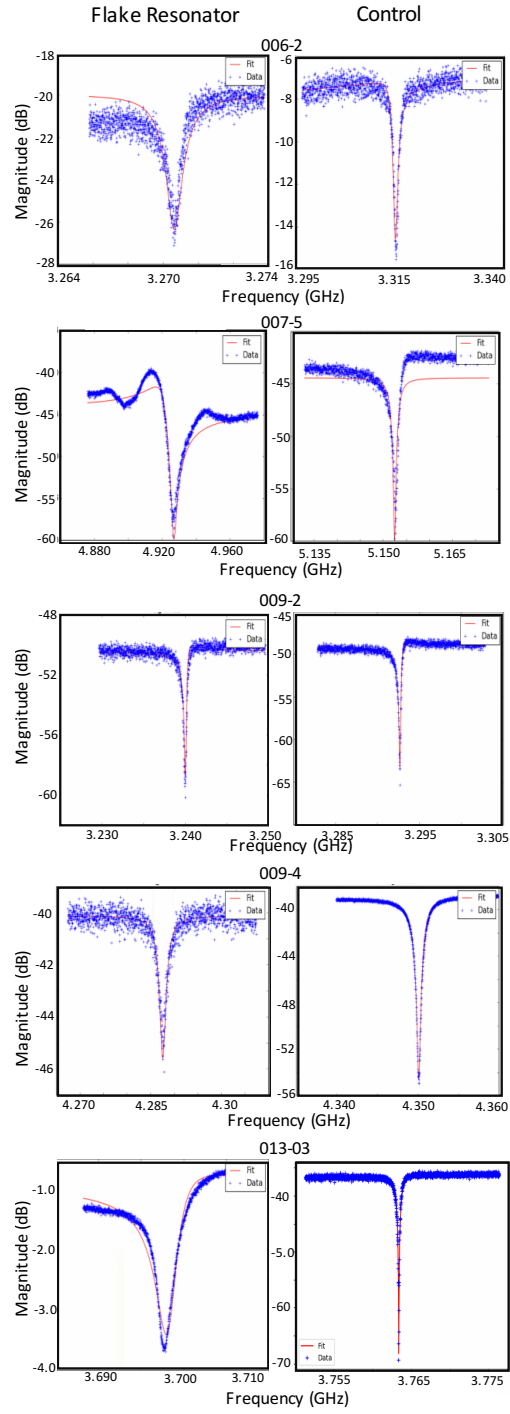


Figure 58: All flake resonators and controls from Table 3. The data is plotted in blue, while the red represents the theoretical fit from which the quality factors were extracted.

Table 4: All quality factors estimated from hanger fits shown in Fig. 58

Sample	Q_{int}	Q_{ext}	Q_{tot}
006-02-flake	4E+03	3.7E+03	1922
006-02-control	4.72E+03	5.72E+03	2587
007-05 -flake	2E+03	6E+03	462
007-05-control	1.5E+04	3E+03	2500
009-02-flake	1.47E+04	1.0E+04	5951
009-02-control	2.6E+04	7.0E+03	5515
009-04-flake	3.5E+03	4.0E+03	1867
009-04-control	1.5E+04	3.0E+03	2500
013-03-flake	1.4E+03	4.52E+03	1068
013-03-control	7.79 E+03	8.2E+03	3995

5.5 Future directions

In the near future, we hope to make other junction-type devices with the TMD flakes, as well. We have shown that it is both possible to make single junction and two junction devices and measure their critical currents. Once we have a good grasp of the strength of the nonlinearity of the materials, it should be possible to create junctions for qubits or even amplifiers. To do that, the next step would be to measure the nonlinear response of the devices.

We expect that the flake itself should have a good deal of inherent nonlinearity. It should be, in fact, possible to back out the full Hamiltonian of the system by breaking up the terms in even and odd components. The even terms can be found in a very similar fashion to the 4th order terms, in the procedure outlined in Chapter 5. By applying a strong tone off resonance, we can monitor how the resonator shifts in frequency vs. pump power. Next, this response can be fit to a polynomial curve, and the strengths of the individual terms can be

backed out.

The odd order terms, on the other hand, cannot be found from one pump tone alone, as this rotates away in the RWA. Instead, a method very similar to amplifier mixing can be used. A tone at 2ω can be applied, and by seeing how the resonator responds to power, or how much gain is created, the odd order terms can be extracted. Thus, the full Hamiltonian dynamics should be possible to examine.

We had been roughly looking for nonlinearity in the flake samples with a high power response, but until recently, had not measured a noticeable shift. However, upon removing some attenuators in the fridge at base, the first sample with a strong nonlinear response was found, and the rough response as a function of VNA power is shown in Fig. 59. This type of response is very promising. We should be able to start probing the flake with the pumping schemes described in this chapter and be able to derive the full system Hamiltonian.

In addition, we also developed a new way of making low geometric inductance resonators *a la* the lumped JPC. Since the kinetic inductance we calculated from the flake resonators is significantly higher than what was predicted from theory, we wanted to create a circuit where the flake itself was a larger fraction of the total inductance, to remove any error. These new lumped resonators have a much larger capacitance to inductance ratio, and thus, should be well suited to this task (see Fig. 60). Furthermore, they no longer require large bonding ‘ears’ which add more geometric inductance.

0013-3 Flake High Power Response




Figure 59: The high power response of the flake sample 013-03. At low powers, the characteristic hanger response is clearly visible, but at higher powers it becomes distorted and then disappears. This nonlinear response is indicative of higher order Hamiltonian terms we would like to extract in the future.




Figure 60: New, low geometric inductance flake resonator design. In this new design, four flakes can be placed, one in each of the four lumped resonators. The center of each has a slit to the outer part of the chip to prevent flux trapping, and the DC bonding sites are the large resonators on each side of the lumped inductor pads (the triangles).

The measurement scheme will be slightly different because we will be doing reflectometry instead of hanger transmission measurements. In addition, only 4 resonators can fit on these chips instead of 5 since the whole chip itself is the same size as the LJPC chip (4 mm x 4 mm).




Figure 61: (a) Optical image of one of the new flake resonators with an encapsulated NbSe₂ flake in the center. (b) One side of the new resonators, showing the DC bonding site where two bonds have been removed for imaging.

The flake stack will now be exfoliated in the middle of the two resonator pads, as labeled in Fig. 60 and shown in the optical image in Fig. 61. The DC bonding pads will be attached on either side of the resonator pads themselves, on the rectangular site, shown in Fig. 61 (b). The bonding has to occur on part of the chip that does not have dielectric underneath it, as the force of the wire bonder proved to be strong enough to push right through the upper metal layer. The new design will be created from NbTiN once again, with much smaller aluminum leads connecting the resonator to the flake. This design was recently fabricated, and soon, we expect to get a much more accurate kinetic inductance reading of NbSe₂.

Lastly, we still are interested in making junctions completely out of exfoliated, thin-film materials, such as by stacking NbSe₂-hBN-NbSe₂. Since we have found that the critical currents between the contacts themselves are so large, the I_0 of these exfoliated junctions should be significantly smaller and thus, they will dominate the nonlinear response. Hence,

we should be able to make even smaller junctions for use as qubits, and even scale up to ≈ 1 mA junctions for use in paramps, without having to worry about the edge contacts. There are so many possible circuits that can be created from this platform that it has become an exciting possibility for how to overcome problems in more traditional superconducting circuits.

6.0 Fabrication And Design Of All JPC Versions

This chapter will start off with a broad overview on how all amplifier samples were fabricated. Specific recipes for spin coating and development can be found in the appendices. My hope for this chapter is that it will serve as a manual of sorts for younger graduate students learning the nuances of lithography and nanofabrication. Many of the devices I discuss in this chapter could be classified as “failures”, as they were deemed unusable for a variety of reasons. However, they are documented here for posterity so future students are not doomed to repeat them, and with the hope that they can learn from my first drafts.

6.1 Dolan bridge technique

All Josephson junctions created in the HatLab have been created with the Dolan Bridge technique. The Dolan Bridge relies on having two layers of resist, which we have chosen to be MMA and PMMA. First, the two resists must be spun onto a clean substrate, which can be either silicon or sapphire, depending on the purpose of the device. The specific spinning recipes can be found in Appendix A and B.

After letting the resists bake for the appropriate time, E-beam lithography can be performed. The E-beam is able to be directed at certain parts of the wafer that we control with a mask. Wherever the wafer is exposed to the E-beam becomes damaged, and the molecules of the resist begin to break, depending on how strong the exposure dose was. A strong dose will destroy both the PMMA and MMA molecules, but a weaker dose will leave the PMMA intact, which creates the bridge that gives the Dolan Bridge its name (see Fig. 62).




Figure 62: E-beam exposure showing the layers of PMMA and MMA that are damaged from the proximity effect and then the full undercut. Both effects must be taken into account when assigning dosages for the EBL.

The stronger dose is generally associated with larger, less delicate features in a device, while the weak doses are used to create the junctions themselves. Specific doses need to be decided with a dose test. Careful attention should be paid to the junction doses as well as all small features. Generally, for our specific lithography machine, layers are broken down into small features and big features. Small features include the entire JRM itself (and anything the order of 10's of μm), and big features, include the resonator and coupling capacitors (generally anything on the order of 100's of μm). The mark of a successful dose test are

sharp corners, lines, and junctions that do not display any "peeling" effects. Furthermore, a wide range of doses should be tested and the ideal dose should be selected from the middle of acceptable doses, since a wide range of variables can cause the dose to change slightly from run to run.

In addition, the proximity effect must be taken into account. The proximity effect arises from the back-scattering of electrons bouncing off the substrate. Hence, not all parts of a pattern will receive the same dose, even if it is assigned that way. Generally speaking, corners and edges receive a weaker dose for this reason, which is why dose tests are so crucial.

After a suitable dose for each layer is selected, lithography can begin on the small features, followed directly by the big device features. Care, of course, should be given towards aligning both fields of view or alignment will be sub-optimal.

After the lithography step is completed, the device must be developed (recipes in appendix) and then aligned for double-angle deposition in the Plassys evaporator. It is also advantageous to make sure that one's device is not placed upside down in the evaporation chamber.

After loading the device in the evaporator, the chamber should be pumped down to $5.0e^{-7}$ to ensure a proper vacuum has been achieved. The junctions themselves are created from a double-angle deposition, which is shown in Fig. 63.




Figure 63: Double-angle deposition to form Josephson junctions. The first layer (blue) is created from depositing aluminum at an angle of 40 degrees, then the insulating barrier is created by flowing oxygen through the evaporation chamber, before the second layer (red) is created from depositing aluminum, now at an angle of -40 degrees. The junction length, $x = h_2 \tan \theta - g$ and the starting position $d = (2h_1 + h_2) \tan \theta$.

The substrate is rotated to one angle and metal is deposited, falling wherever is not blocked by the PMMA bridge. Next, oxide is allowed to grow naturally on the aluminum, before the substrate is rotated to the opposite angle while metal is deposited again, this time, about 4 times thicker. The overlap region, x created from the two depositions form the junctions.

Once the Plassys recipe has finished, the device needs to be left in hot acetone (at 65 Celsius) for about 20 minutes for lift-off. After about 20 minutes, a pipette should be used to spray the hot acetone over the sample while still in the beaker to help remove the metal film. After the sample looks clear of metal, use a tweezer to remove it from the beaker and spray it down first with IPA, then DI water to remove any residue.

6.2 Microstrip JPCs

The single-layer fabrication described here is what was used for all microstrip iterations of the JPC. The more complicated, multi-step fabrication that was used for the lumped JPC and SNAIL-based amplifier are discussed further in their respective subsections. The finished original 8-junction microstrip JPC and JRM are shown in Fig 64.




Figure 64: (a) The packaged microstrip JPC, showing the $\lambda/2$ resonators whose length dictates the frequency of the signal, idler, and common modes. Coupling to the outside world is controlled by a coupling capacitor on each mode, including a very weak coupling for the pump port. (b) Zooming in on the intersection of the two microstrips shows the JRM, which consists of 4 smaller Josephson junctions, which in this sample are shunted by 4 larger junctions to get rid of hysteresis.

The entire sample is glued down in the gold-coated JPC boxes with silver paste to ensure good thermal contact. Generally, we let the glue dry for a number of hours before wirebonding the sample to the 4 copper striplines.

This concludes the entire fabrication procedure for the one-layer JPCs. In the next subsections, I will specifically focus on the multiple iterations of the microstrip JPC and discuss their individual difficulties that lead to the Wire JPC.

6.2.1 Meandering JPC

We tried to improve upon the original microstrip design to get rid of 4th-order Hamiltonian effects, as described in Chapter 2. In order to do that, we had to replace the inner shunting junctions with linear inductors. The inductance has to be fairly significant and we wish to keep $\alpha = \frac{L_{out}}{L_J}$ as small as possible in order to have the desired effect.

Since we want the inductance to be large, fabricating straight aluminum lines across the center of the junction didn't add enough. We instead came up with what we call the "meandering" JPC, which features winding, connecting lines, as seen in Fig. 65.

In this design, the JRM ring was designed to be about $24 \mu\text{m} \times 24 \mu\text{m}$ while the inductors are $\approx 60 \mu\text{m}$ in length. It is also important to note that the size of all the sub-loops in the JRM are the same size so magnetic field is equally distributed amongst them.

However, both the fabrication and simulation of this device became an issue. Calculating the α and β parameters of this device yielded approximately 0.37 and 3.4, respectively. However, a β less than 4 should no longer feature minor lobes in the device's flux sweep. That is not what we observe in reality. In all versions of this meandering JRM minor lobes were not visible, meaning that the experimental β must be larger than we calculated. We hypothesized that the meandering shape gave rise to a high mutual inductance which was not accurately simulated in HFSS or analytic methods. Iterating between subtle fabrication changes in this device became too difficult when we attempted to account for this mutual inductance. In addition, the lithography doses for the inner linear shunts did not seem to be stable, and led to asymmetric shapes, which can be seen in Fig. 65. We decided to try to get rid of the meandering inductors all together.

“Meandering” JRM




Figure 65: SEM image of the meandering JRM showing the 4 Josephson Junctions in the outer ring and the meandering linear, shunting inductors in the center. The center shunts frequently became distorted because it was hard to find a correct, stable dose for this part of the design.

6.2.2 Lightning JRM

In order to reduce the effects of mutual inductance from the meandering JPC, we decided to use a much larger outer JRM ring ($100\text{ }\mu\text{m} \times 100\text{ }\mu\text{m}$), and shunt it with jagged lines, creating the aptly named “lightning” JPC, seen in Fig. 66.

“Lightning” JRM




Figure 66: SEM image of the lightning JRM where the meandering lines have been replaced by the jagged linear shunts to reduce mutual inductance. It is also important to note the size of the ring itself is much larger.

Remember, we want to keep α as close to zero as possible. Since there is now more linear inductance from the ring, we must increase the inductance from the junctions. We do this by decreasing the size of the junctions to $1 \mu\text{m}^2$, which made out $\alpha \approx 0.35$.

With this new design, we were able to reliably create β to be less than 4. However, the issue with this design is that the sub-loops of the JRM are not much larger, and although they are all the same area, it is much harder to ensure magnetic flux uniformity through them all. This altered the device's flux sweeps, and made them rather asymmetrical and difficult to figure out stable regions to bias the amplifier. Luckily, we had one more idea for how to shunt the JRMs.

6.2.3 Wire JPC

We decided we had to shrink the JRM outer ring back down to its original size, and we wanted to keep the junctions the same size as well. However, there didn't seem to be a simple solution if the entire fabrication was to still be done entirely of aluminum. In order to create shunts with the desired large inductance, we decided to explore the idea of kinetic inductance, which we had been doing for the vdW flake resonators and is discussed in the previous chapter.




Figure 67: The Wire-shunted JRM (WJRM) which features NbTiN shunts, which are $70\ \mu\text{m}$ long, shown here in light grey. These linear inductors add a boost of about ≈ 5 to the geometric inductance. They are deposited first, before the Al outer ring and Josephson junctions. This is due to the fact that NbTiN does not grow a native oxide, so good contact between metals can still be achieved.

NbTiN has a kinetic inductance boost of about 5, which is ideal for our design, since we wanted to shrink the size of our JRM ring significantly. This boost should help us compensate for the small size of the ring, and let us decrease the size by a factor of 5 as well. In addition NbTiN was an ideal material candidate because we had already become familiar

with its fabrication and properties from the resonators. Furthermore, NbTiN does not grow a native oxide, so making contact between the inner shunts and the outer ring should not be a problem. We held the junction sizes to $1.0 \times 0.8 \mu\text{m}^2$ to reach our desired β ratio. The NbTiN shunts were created to be $2 \times 20 \mu\text{m}^2$. The final device can be seen in Fig. 67 and proved to be the best way to create added linear inductance in the JRM ring. These WJRMJs have excellent properties and their results are shown in the next chapter.

6.3 Lumped JPC

The microstrip JPC is a reliable and well-loved device, but even with the addition of the WJRM, they are not without flaws. The microstrips have inherent higher harmonic frequencies that must be accounted for and avoided. In addition, using transmission line-based resonators does not allow for full Hamiltonian control. If we were instead to build this device entirely out of lumped elements, however, we would have very precise control over the Hamiltonian and be able to get rid of higher harmonics.

In exchange for all these nice features, the fabrication, of course, becomes more involved, as the capacitors, seen in Fig. 29 will have to be created in different layers, below and above an insulating barrier. The optical and E-beam recipes can again be found in Appendices A and B; however the multi-step fab is outlined as followed:

1. Spin optical 2-layer resist in Maskless Aligner (MLA).
2. Write base layer with MLA and develop in 351 and 400k.
3. Put down base layer of NbTiN with AJA ($t = 100 \text{ nm}$), then liftoff with PG remover.
4. Spin optical 2-layer resist.
5. Do MLA write with the automatic alignment procedure for the sapphire layer.
6. Develop in 351 and 400k.
7. Deposit sapphire in Plassys ($t = 100 \text{ nm}$), liftoff with PG remover.
8. Spin E-beam (EBL) resist: single layer only.
9. Write JRM shunts in EBL using manual 3-point alignment procedure.
10. Develop in MIBK/IPA.

11. Deposit NbTiN shunts using AJA ($t = 20$ nm).
12. Spin double-layer EBL resist, write final JRM using manual 3-point alignment.
13. Deposit Aluminum in Plassys and liftoff in acetone.

As one can see, the basic layers are the ground plane (NbTiN), the sapphire insulating layer, the thin NbTiN shunts, and the aluminum top layer of the capacitors, whose size determines the mode frequencies. These layers are most clearly shown in Fig. 68 (a) which depicts the HFSS design. Fig. 68 (b) is the corresponding SEM image.

Lumped JPC

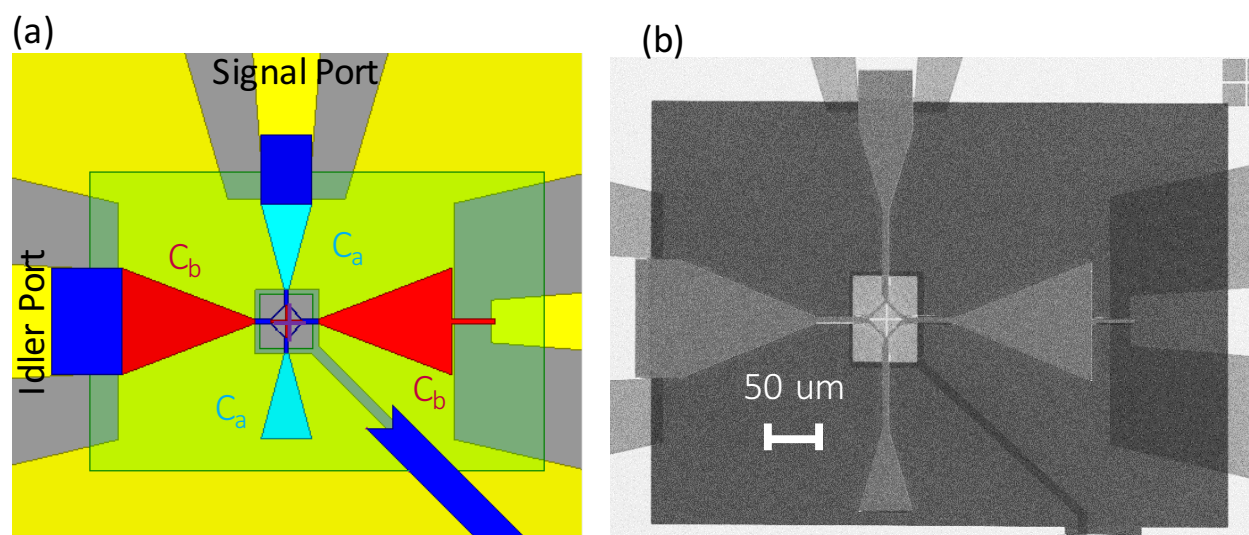


Figure 68: (a) The Lumped JPC (LJPC) simulated in HFSS. The yellow represents the base layer of NbTiN (100 nm thick), the green represents the sapphire (AlO_x) insulating layer (80 nm), the purple are the NbTi shunts (10 nm) and the red/blue pads are the top aluminum layers of the capacitors. (b) SEM image of corresponding device.

These JPC's can also have considerably large linewidths, on the order of 100's of MHz, for the resonator modes, due to the ability to create direct ohmic contact between the capacitor pads to the ports. Alternatively, by extending the AlO_x layer (shown in light green in Fig. 68 (a)) one can also create coupling capacitors like the pump port on the far right.

6.4 SNAIL/SHARC amplifier fabrication

The last type of amplifier that will be discussed in this chapter is the SNAIL-based amplifier. The need for this innovation is discussed in previous chapters, for now we will only examine the fabrication.

6.4.1 First iteration

The first SNAIL-based amplifier was designed to create an array of linear modes with the aim of studying directional amplification. It was designed in HFSS, as shown in Fig. 42, where the signal would ideally be input on one side of the amplifier, flow through the linear modes created from the same lumped fabrication and then to the SNAIL mode by means of a coupling capacitor created from a sort of “bridge” design between the lumped pads. The SNAIL and all 3-body couplings will be driven by a pump port created from a small coupling capacitor shown in the top of the figure.

The linear modes, which I call signal 1 and signal 2, are easy to design and fabricate, as their frequency is simply a function of the size of the capacitor pads. In addition, the signal modes will no longer modulate with flux at all, as they are designed to be significantly detuned from the SNAIL mode and have no nonlinearity of their own. They are created the same way as the resonators in the Lumped JPC.

The only complexity in this design’s fabrication comes from the SNAIL itself. The SNAIL is a nominally “pure” 3-wave coupling element that is designed to be free of 4th-order effects at a specific external bias flux, much like the WJPC [48, 47, 51]. The SNAIL element is similar to a DC SQUID, with asymmetry introduced from the difference in junction sizes, on either branch of the parallel arms, as seen in the circuit schematic in Fig 40. For this first attempt at the SNAIL-based design, specific attention was mostly given to α , which for the SNAIL element only, refers to the ratio of the size of the junctions.

We had some difficulties fabricating a SNAIL with identical large junctions sizes, probably due to the proximity effect causing the middle junction to receive a larger dose than the other two. In order to compensate for this, we adjusted our layout design to shift the

junction sizes by their offset discrepancy. All following fabricated SNAILs no longer had this problem.

In the first SNAIL-based amplifiers that were tested, we also saw some unusual jumps in the flux sweeps. We believed this was due to flux slips from the SNAIL, so in later devices, small holes were added around the SNAIL in order to trap this flux. The rest of this fabrication follows the exact same steps as the Lumped JPC. Our first fabricated version of the SHARC design is shown in Fig. 69.




Figure 69: Microscopic image of the first version of the SHARC amplifier, showing the lumped resonators that mimic the LJPC fabrication. The SNAIL is at the center, highlighted by the light blue box, and is responsible for all the nonlinearity in the system. The SNAIL shares its nonlinearity to the other three linear modes through the coupling capacitor islands. The flux slit and healing capacitor is visible in the top right-hand corner of the image. In addition, flux-trapping holes were added in the center of the device to prevent flux slips.

6.4.2 Iteration 2

There were some concerns that the previous design was not suitable for high Q resonators, which is necessary for the amplifier’s ideal purpose. We hypothesized that it was either the NbTiN or the sapphire, or some combination of the two, that was to blame. The NbTiN sample that we had been using was old, and had consistently been creating lower Q resonators than it should. In addition, the sapphire we were using began to look less uniform over time, as more oxide developed. Either one could have been hindering our ability to get high Q resonators by adding additional loss.

We took the ideas used to guide the design for the vdW flake resonators and incorporated it into our amplifier fabrication. The goal with this design was to eliminate the NbTiN and the sapphire to only leave aluminum. However, aluminum’s native oxide grows so quickly that making direct contact between different depositions of the metal is nearly impossible. We switched from an optical resist fabrication to a chlorine gas etch with an RIE. With aluminum deposited down first over the entire sample, the metal was then etched, leaving the three linear resonator modes, the coupling capacitors, and 1/2 of all the SNAIL mode finger capacitors. The other half of the finger capacitors and the SNAIL itself were deposited in the final layer, still done with E-beam lithography. By doing this, neither aluminum layer ever was required to make direct contact, and all of the mode coupling was done through the finger capacitors. Note that in between the aluminum layers, the healing capacitor was still created from sapphire, since it doesn’t matter if there is loss here. This is not modeled in every HFSS simulation, as shown in Fig. 46, in the interest of time. It is of course, visible in the microscopic images of the SHARC, as seen in Fig. 70.

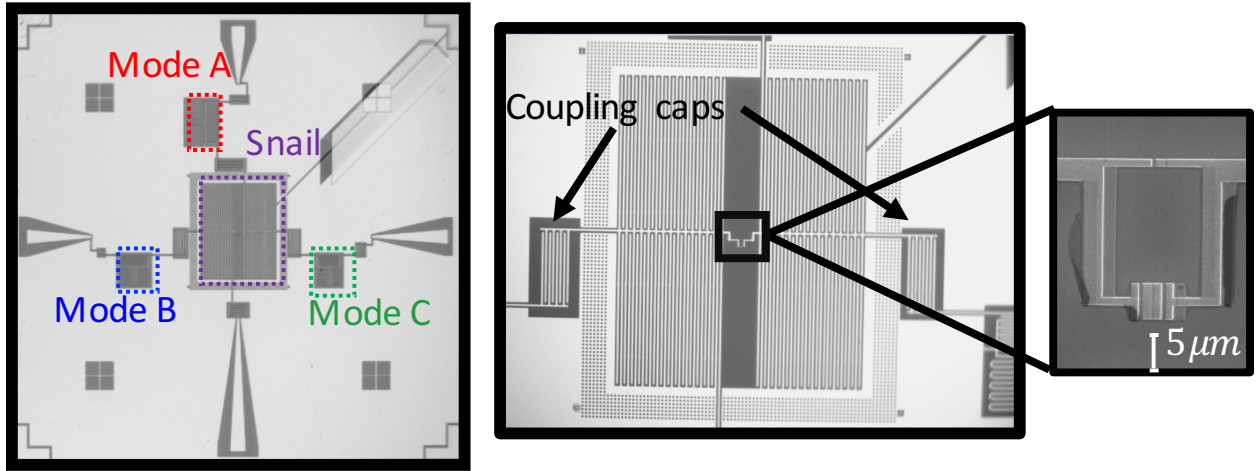


Figure 70: Microscopic image of the entire chip of the SHARC amplifier, and a zoomed in image of the SNAIL. The three linear modes, as well as the SNAIL mode are outlined in red, blue, green, and purple dotted lines. Coupling capacitors, visible in the second image, control the coupling between the nonlinear SNAIL mode to the three linear modes. The geometry of the SNAIL itself, shown in the last image, controls in inherent nonlinearity of the entire system.

7.0 Amplifier Data

This is the section where I will be reporting the data collected from each of the pumping schemes described in Chapter 3 and examine the effectiveness of design choices in Chapter 4. Specific attention is given to the effectiveness of 4th-order cancellation, phase-sensitivity, and qubit measurement. I will first present parametric coupling data, which is all performed with the same physical amplifier (WJPC32) for best comparison. Next I'll show and discuss the preliminary results of the array amplifier and SHARC amplifier.

7.1 WJRM data

The first JPC property we want to examine is the effectiveness of the 4th order cancellation, which was theoretically discussed in Ch. 2, since without this ability, multiparametric interactions becomes much more challenging. The two aspects we want to examine are: 1.) The ability to cancel the terms themselves, 2.) the ability to find the bias point where the 4th order terms are zero.

In order to measure the strength of higher-order even terms in the JPC Hamiltonian, we apply a Duffing test, which is visualized in Fig. 71. We know that the 4th order terms are the largest even order terms, and we assume that the ratio of these terms vs. even higher order terms is quite large.




Figure 71: In order to measure cross Kerr, We prepare a pump tone that is 5 times the linewidth away from one mode frequency, and measure the frequency shift of the other mode as a function of that pump power.

A duffing tests consists of a strong pump tone that is applied far off resonance from one of the modes, while either the same mode, for self Kerr, or the opposite mode, for cross Kerr, is monitored for frequency shifts as a function of the strong pump power compared to a weak tone.

In order to visualize the data, we plot the magnitude of the frequency shift as a function of the pump power over a range of external applied current. It is useful to plot these duffing responses with the flux modulation to see where on the flux sweep the responses are weakest, as shown in Fig. 72.




Figure 72: (a) The flux sweep of WJPC21 showing an asymmetrical minor lobe. (b) and (c) Show the self and cross Kerr respectively from the flux sweep in (a). It is interesting to note how similar the self and cross kerr look in this sample. The ideal 4^{th} -order nulling points are the white areas of the duffing test

The strongly colored blue and red region correspond to where the frequency of the mode shifts most strongly. The dark blue corresponds to a strong positive Kerr, and the dark red to a strong negative Kerr. Obviously, we are looking for points where there is no Kerr, and thus, no 4^{th} order. This occurs in the white regions, between the blue and red. These regions can be obvious or difficult to see, depending on the sample. A more challenging example of this is WJPC30, who's duffing test is shown in Fig. 73 (a). Much of this region is fairly white, so one thing that can be done is to extract the exact magnitude of the Kerr terms

themselves, seen in Fig. 73 (b) and (c).

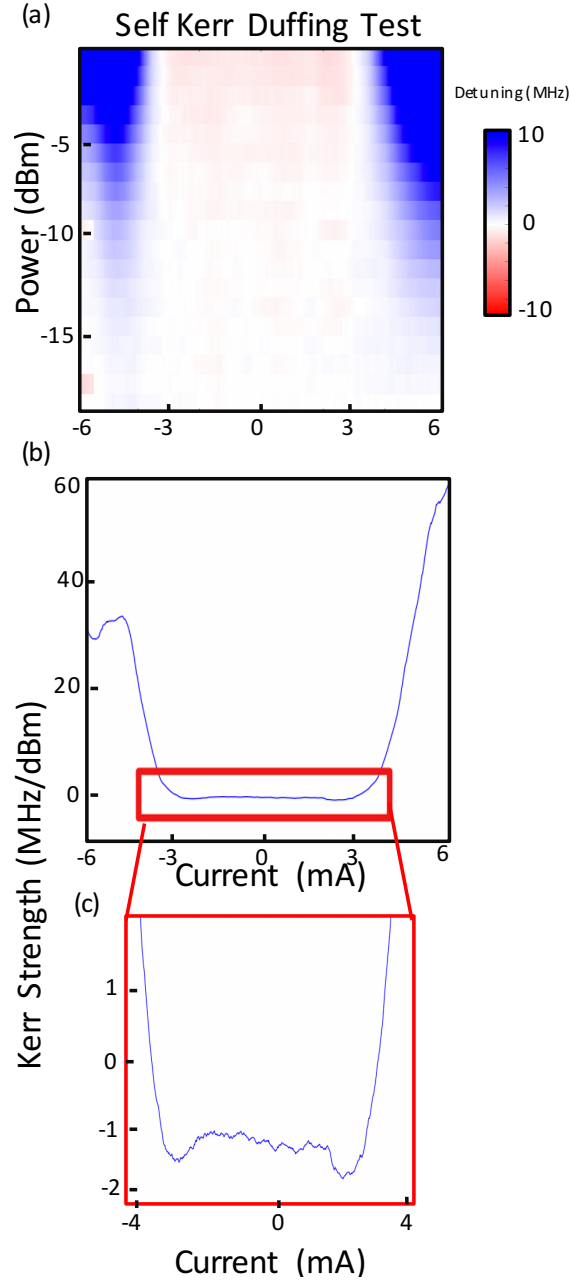


Figure 73: (a) Self Kerr Duffing test of WJPC30. (b) Self Kerr magnitude extracted from (a). (c) A close up of the bottom of (b), more precisely showing the nulling points, or where the magnitude crosses 0 (here at 3.130 and -3.203 mA).

Once the Kerr terms themselves are extracted from the Duffing test we can easily find

the bias points where they are zero. In Fig. 73, we can see that the magnitude is zero only at 3.130 and -3.203 mA. These are what we call our nulling points, and represent the best bias points available in the sample in this current range. Choosing one over the other is usually just a matter of trial and error.

In the WJPC we use for the parametric pumping comparison, the duffing test pointed to two clear nulling points, shown by the dotted black lines in Fig. 74. As a quick aside, the range of current that the duffing test should be performed over is related to the strength of the magnet. WJPC32 was operated with a weaker magnet than WJPC30, so the bias current range is much larger.




Figure 74: Self-Kerr Duffing test of WJPC32. The frequency shifts as functions of pump power are plotted over a range of currents/external flux. The nulling points are noted by the dotted black lines. The range of this Duffing test is much larger than the others shown because WJPC32 is used with a larger, lid-cap magnet.

This method of identifying and using 4th order nulling points has been highly successful. Multiparametric pumping in the JPC has been attempted in the past in other labs and by

ourselves, without accounting for higher-order Hamiltonian terms, but frequency shifts made it impossible to achieve high gain. One trick to keep in mind going forward is that if the JPC is biased to a nulling point, the cavity must also be tuned to that frequency as well in order to make it suitable for qubit read-out.

7.1.1 X^2Z pumping

The first mode of amplification that is presented in this section is the only singly- pumped X^2Z mode amplifier. The X^2Z mode, as previous described, is a phase-sensitive mode of amplification created by pumping a JPC at twice the idler or twice the signal frequency. This creates gain at $1/2$ the pump frequency in the corresponding mode.




Figure 75: 20 dB gain created in reflection in the a mode by pumping the JPC at $2\omega_a$, or, the X^2Z scheme. The profile shown here is actually of phase-preserving gain, since the phase-sensitive point is not visible with the VNA and stands 6 dB higher.

Probably nothing unusual stands out about this gain profile, because it looks identical to all phase- preserving gain profiles that are typically shown. In fact, the Vector Network Analyser can only show the phase-preserving gain profiles, and does not pick up the phase-

sensitive point after averaging. The phase-sensitive point of this amplifier, the point in the very center of the peak, is actually 6 dB higher than what appears here, because two photons interfering on resonance with one another adds twice the voltage, which is four times the power, or 6 dB. To compare the phase-sensitive vs. phase-preserving amplifier, we plot both side by side, shown in Fig 76.




Figure 76: X^2Z and XYZ gain profiles for multiple gain strengths. X^2Z is plotted in blue and XYZ is plotted in green. The profiles are almost identical, despite being achieved in different ways, until the last pane, when the phase-sensitive interference point becomes visible in the X^2Z mode. Also note that the pump power strengths for X^2Z are moderately higher than the XYZ pumps throughout, probably due to the difference in bandwidths between the two modes.

X^2Z and XYZ pumping schemes produce nearly identical gain profiles, just depending on how closely related κ_a is to κ_b . However, even though these phase-preserving profiles are identical, the pump power required to achieve them was not. On average, this mode of phase-sensitive amplification necessitates slightly high pump powers, probably due to the fact that the max gain is in fact 6 dB higher than is visible here.

When working with a phase-sensitive amplifier, the phase will obviously play a prominent role. As described in more detail in chapter 2, the phase effects how much squeezing is applied to the read out [104]. When the phase of the pump isn't optimal the I-Q component of the $|g\rangle$ and $|e\rangle$ state overlap with one another, making them indistinguishable.

Practically, what we do in an experiment to find the optimal phase for the pump is to scroll through the phase of the generator from 0 to 2π and select the angle which produces the histogram of $|g\rangle$ and $|e\rangle$ with the largest separation between the two Gaussian distributions. This can be seen below in the figure on the left. Rotating this phase by 180 degrees brings the mode to what we call the orthogonal alignment and the phase which creates the most overlap between the states.




Figure 77: X^2Z histograms measured at ideal and orthogonal alignment of the pump phase. In the first histogram, the optimal alignment case, the $|g\rangle$ and $|e\rangle$ state are well separated from one another. When this ideal phase is rotated by 90 degrees to the orthogonal alignment, the two states overlap with one another and become indistinguishable.

In order to get a better sense for how important this phase is to qubit read-out, we can attempt to measure quantum jumps. Reading out quantum jumps with the qubit is a pretty eye-opening experiment in and of itself, as it allows the observer to see the evolution of the qubit in real time, with no averaging. Figure 78 shows the quantum jumps of the qubit between the $|g\rangle$ and $|e\rangle$ state.

One subtle point to be made here is that qubit read-out is only possible while the phase of the pump remains at close to the optimal phase position. However, over time, generator phases drift. If one desires to take data with a phase-sensitive mode of amplification for a long, continuous period of time, an interferometric set-up might be needed. However, for one day or one night, the phase doesn't seem to drift enough to have a significant effect, which will become clear in the next subsection.




Figure 78: Quantum jumps measured by pumping the JPC at X^2Z and aligning the amplifier's pump phase to the optimal alignment point. Without controlling the phase, no quantum jumps are visible on these phase-sensitive modes of amplification.

7.1.2 GC amplifier pumping

Experimental data from the GC amplifier pumping scheme will be discussed in this section. First, we biased our amplifier (WJPC32) to the left nulling point, as seen in Fig. 74. At this point we identified the signal and idler frequencies, which are 7.4668 GHz and 4.8715 GHz respectively, and pumped at their sum, tuning the applied pump power to achieve 20 dB gain in reflection (G). We repeated the process with the pump at the difference frequency (C), tuning the applied pump power to find a 20 dB dip in reflection. These bias powers are each very close to the critical values, and tell us the room temperature ratio of applied microwave powers required to balance the G and C processes. We next fine tuned the applied pump frequencies until they linked to identical idler frequencies in transmission through the device.

Turning on both drives simultaneously achieves GC amplification. We control GC gain by increasing/decreasing both pumps simultaneously while maintaining the ratio established above. Measured GC amplification gain strengths between 8 and 15 dB, along with reflection performance, are shown in Fig. 79.

All transmission gain curves are measured with a VNA. An external mixer at the difference frequency converts device outputs from the idler mode back to the input, signal mode, frequency. The VNA is sensitive only to a single frequency, thus all gains shown are phase-preserving, and so for large gains, the phase-preserving gain peak sits ≈ 6 dB below the phase-sensitive gain of the device, so that the maximum phase-sensitive gain for the 15.5 dB curve (which we will use for qubit readout) is 21.5 dB. Above 10 dB gain, the frequency of the pump tones had to be adjusted slightly to account for the imperfectly nulled higher order terms, which shifted the device's modes, and hence, the frequency of peak gain, to lower frequencies. At the same time, the reflection curves at this higher power start to show small disturbances away from unity (Fig. 79 (a)). For clarity in these experiments, the amplifier is flux shifted slightly below the cavity frequency, which is evident as a notch to the right of the gain peaks. The other crucial feature of these gain curves is that they demonstrate the same bandwidth regardless of gain (Fig. 79 (b)). We compare these bandwidths with the standard 20 dB for one single pump amplification process represented by the black curve. The GC

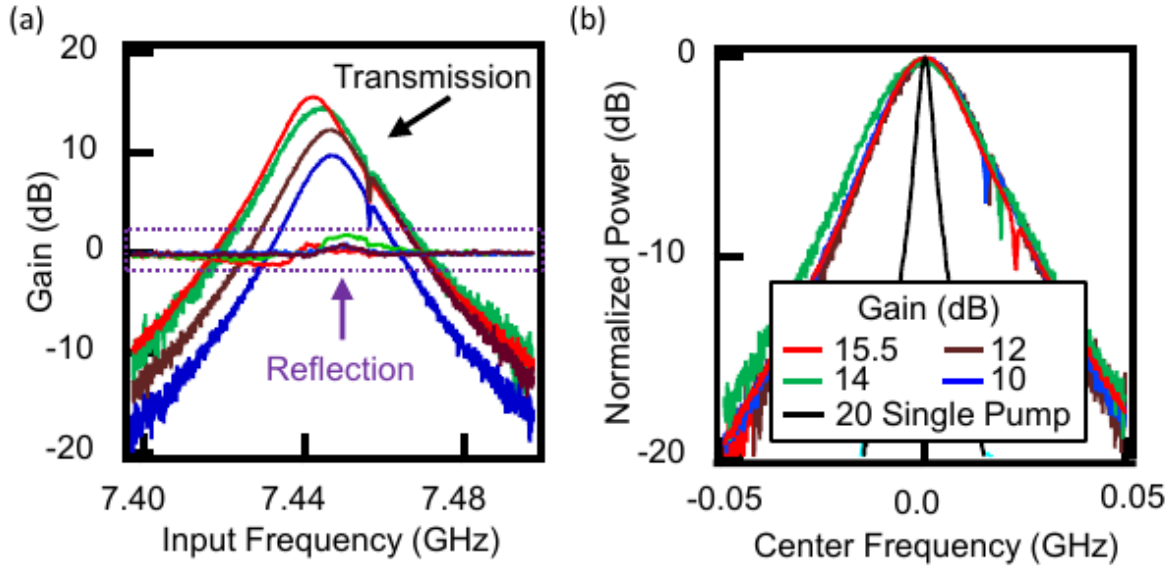


Figure 79: (a) GC data in transmission (marked with black arrow) and reflection (purple arrow), showing large, broad peaks in transmission and nearly 0 dB with minor fluctuations in reflection. (b) All gain curves for different pump powers shown in (a) normalized to one another's center frequency to highlight the constant bandwidth. Compared to a single G-pumped gain profile, seen in black, the bandwidth is clearly much larger.

amplifier shows 14 MHz of phase-preserving bandwidth. In contrast, a singly-pumped gain response of 20 dB only has a bandwidth of 2.33 MHz in the same device. This bandwidth is approximately 6 times larger, in good agreement with theory. The bandwidth is not quite as large as theory predicts because the linewidths of the signal and idler mode are not quite identical, as the theory assumes.

We can also take data from a spectrum analyzer (received power) when the GC amplifier is driven with a fixed tone 5 MHz detuned from its center frequency, shown in Fig. 80. The presence of a symmetrically detuned tone below the center frequency is a clear demonstration of the phase-sensitive nature of this devices' gain.




Figure 80: Signal analyser data of GC amplification when one signal tone is applied slightly off resonance from the peak. A corresponding tone appears on the other side, showing the interference of two photons in the same resonator.

Collectively, Fig. 79 and 80 demonstrate all key predicted features of GC amplification; fixed bandwidth, 0 dB gain in reflection, and phase-sensitive amplification in transmission.

However, we also want to demonstrate that our Kerr-nulled, GC pumped device is a practical, quantum-limited amplifier. To do so, we performed phase-sensitive strong/projective measurements on a superconducting transmon qubit. The qubit is first prepared in the superposition state $|\Psi\rangle = (|g\rangle + |e\rangle)/\sqrt{2}$.

Next, we determined the optimal alignment of the device's amplified quadrature by finding the largest separation between ground and excited states when projective measurement is performed, just as we did for the X²Z amplifier, as shown in Fig. 81 a. Rotating the relative pump phase by 180 degrees from this optimal point moves the signal to the squeezed quadrature, so that the $|g\rangle$ and $|e\rangle$ states overlap (Fig. 81(b)). Both histograms contain 80,000 measurements.

Finally, we calculate the quantum efficiency of our amplifier via the back-action of deliberately weak measurements on the qubit's state, as in [105, 106, 107]. The pulse sequence is detailed below. For phase-sensitive amplification, only one microwave quadrature can be received at room temperature. In a qubit measurement, this confines the qubit back-action to a single plane on the Bloch sphere, determined by which quadrature is amplified. To calibrate our quantum efficiency (η), we use the amplifier aligned as in Fig. 81(b) resulting in a pure back-action on the qubit's phase. We found that our quantum efficiency was $\eta = 55$, in good agreement with the efficiency of the device when operated as a typical phase-preserving amplifier at the same bias point. We also performed weak measurements with the amplifier aligned optimally (as in Fig. 81(a)) for which the back-action is only on the qubit's z-coordinate.




Figure 81: Signal analyser data of GC amplification when one signal tone is applied slightly off resonance from the peak. A corresponding tone appears on the other side, showing the interference of two photons in the same resonator.

Weak measurement is a type of measurement used to study the back-action on the qubit

and extract the quantum efficiency. The standard way to perform back-action measurement is shown in Fig. 82. We first strongly read out the qubit and record the outcome, which will be used to prepare the qubit in the ground state with state selection. Then the qubit is rotated about the y -axis and measured with a variable measurement strength, and the outcome is recorded again. For the final tomography, phase measures the x , y , or z component of the qubit with a strong measurement pulse. If we take the qubit to originally be oriented along the y -axis and the I quadrature to be perfectly squeezed, the back-action corresponds to stochastic, trackable motion of the qubit state in the $x - y$ plane, with the extent of the motion varying with the strength of the measurement. For weak strength, the back-action looks like a stochastic rotation in the $x - y$ plane with the degree of rotation encoded in Q .




Figure 82: We first strongly read out the qubit and record the outcome, which will be used to prepare the qubit in the ground state by postselection. Then, the qubit is rotated about the y -axis and measured with a variable measurement strength, and the outcome is recorded. The final tomographs, phase measures the x , y , or z component of the qubit Bloch vector with a strong measurement pulse.

This pulse sequence produces the results seen in Fig. 83. By fitting this data, we found the quantum efficiency of our read-out to be about 55 %, which is identical to what we have measured with a traditionally, single-pumped JPC. This implies that this amplifier is still just as quantum limited, and adds no additional noise.




Figure 83: (a) Histogram of weak measurement protocol with orthogonal alignment (Q-quadrature amplified) to determine quantum efficiency of the GC amplifier. Each plot contains 80,000 measurements. (b)-(d) Conditional expectation value of X/Y/Z after weak measurement plotted versus measurement outcome. Sinusoidal oscillation in both X and Y are a non-classical stochastic Ramsey process. Together with the nearly constant outcomes for Z, these plots show how an orthogonally-aligned measurement provides a ‘kick’ around the equator of the Bloch sphere.

Lastly, we look at the saturation power performance of this amplifier. Previously, it was thought and accepted that the poor saturation power performance of parametric amplifiers was due to pump depletion [86, 108]. However, if this was the case, then these multiparametrically pumped modes would have higher saturation power, since the pumps that drive

these operations are much stronger. This is not what we observe in reality. If we look to Fig. 84 we see it shows that the performance is no better.




Figure 84: The saturation power of our GC device at around 14 dB is measured. $P_{\pm 1 \text{ dB}}$ is around -130 dBm . We have concluded that still higher-order nonlinearities can contribute strongly to gain saturation and must be controlled to realize amplifiers with superior saturation power.

The saturation power for the singly-pumped gain process is very similar to that of GC. It varies, depending on the bias condition somewhat, but is within $\approx 2 \text{ dB}$ of what is shown here in Fig. 84. Unfortunately, even though we are working with an amplifier at a bias point that ideally has no 4^{th} order terms, we have concluded that even higher order nonlinearities greatly effect and hinder the saturation power [76].

7.1.3 GCI amplifier pumping

The next amplifying scheme that will be presented is the GC Imbalanced pumping or GCI. As the name suggests, GCI is very similar to GC amplification, with the slight modification that the strength of the conversion pump tone is greatly increased.

To achieve GCI gain experimentally, we start with a balanced GC process, which is described in the previous section. From here, the conversion tone is tuned up so that a large dip in reflection becomes visible. Then, we switch the read-out to transmission and use a mixer to tune our amplifier's outgoing frequency back to that of the VNA. To increase the gain, G and C pumps are both now individually increased slowly. As long as the dip in reflection is still present, the correct ratio of the pumps is approximately maintained.

Fig. 85 shows that when the conversion pump is tuned so that there is about a -15 dB dip in reflection, the corresponding amplitude in transmission is just short of 14 dB gain (in blue). Again, this is a phase-sensitive amplifier, so the center frequency which corresponds to the phase-sensitive point is 6 dB higher than what the VNA captures. Thus the phase-sensitive point here actually sits very close to 20 dB.

If we compare the response in reflection to GC (purple) and to both GC and single G in transmission, we can see that both scattering parameters are superior. GC is attractive because it only has unity gain in reflection, but GCI is matched in reflection. That means that it theoretically has no reflections, which is obviously much closer to our ideal amplifier. The amplifier is still bi-directional, meaning that you can amplify not only from mode a to b , but also in reverse. A fully directional amplifier is still necessary to achieve ideal performance, because stray photons can still end up traveling backwards down the amplification chain to the qubit [109, 110]. Thus, measurements should still be performed with commercial circulators.




Figure 85: GCI data measured in transmission and reflection, showing that GCI (blue) is broader than both GC (purple) and single G-pumped (green) amplifying schemes. In reflection, we compare GC, which shows unity gain, with GCI, which is matched up to almost -15 dB. With these scattering parameters, GCI proves to be the best multiparametric pumping scheme we have tried.

Furthermore, if we compare the bandwidths of these different pumping schemes in reflection we can see that GCI has the largest bandwidth of all, by quite a bit. At this bias point the single gain pump achieves only about 1.5 MHz bandwidth, which is typical for a standard JPC. The GC pumped mode here has a moderately different bandwidth than the data reported in the previous section, which can be explained since this is a different bias point (still a nulling point), and the device had aged somewhat between the cooldowns when this data was taken. Nevertheless, we can compare GCI and GC side by side at this new bias point, and we find that the bandwidth goes from 8.3 MHz for GC, to 13.3 MHz for GCI. Again, the theory predicts a broader bandwidth for both of these modes of amplification, but that rests upon an assumption that $\kappa_a = \kappa_b$, but this is not exactly the case here ($\kappa_a = 25$ MHz, $\kappa_b = 20$ MHz). The most important aspect to note is the ratio if these

bandwidths, because even though this bias point doesn't produce as wide a bandwidth as in the previous subsection, GCI will always be broader than GC as the same external flux conditions.

The GCI pumping scheme is also phase-sensitive, as in the GC mode of operation. This can be proved in a very similar way to what we did in the previous section, by using a spectrum analyzer and inputting a signal slightly off-resonance from the center frequency, and observing a corresponding idler tone appear on the opposite side. GCI amplification also squeezes the modes the same way GC does. We did not explicitly measure a qubit with this amplifier because it would look identical to the previous section.

7.1.4 GG amplifier pumping

The last parametric pumping scheme that we will discuss in this chapter is GG amplification. GG amplification is inherently different than the other two schemes already discussed because it amplifies in reflection as well as transmission. This is obviously a downside of the amplifier, however, the GG pumping scheme is much simpler to set up.

Unlike GC and GCI, GG pumping doesn't seem to be negatively effected by higher order Hamiltonian terms. This is due to the fact that GG doesn't require the intensely strong pump tones that the others do, which increases the effect of higher-order terms.

To experimentally acquire GG gain, we first find gain separately with two generators at the same frequency, where $\omega_p = \omega_a + \omega_b$. From there, one can start detuning both generators an equal amount from this center frequency, usually in 1 MHz increments. The data presented in Fig 86 was taken using this procedure, until the generators were detuned from each other by ± 5 MHz.




Figure 86: GG (blue) vs. G (green) gain profiles measured in reflection. GG amplification has a broader, flatter gain profile (5.6 MHz) than a single G pump alone (1.7 MHz). The dips below zero are characteristic of this amplifiers' pumps interfering with one another.

This data was taken in reflection, compared to a single gain pump to highlight the difference in bandwidths. The single gain tone has a bandwidth of 1.7 MHz and the GG profile has a bandwidth of 5.6 MHz. This is in good agreement with the theory presented in Chapter 2.

If the noise temperature of the standard HEMT is given by T_{HEMT} , and the noise temperature of the quantum fluctuations are given by T_Q , then we can examine a parameter called the Noise Visibility Ratio, defined by

$$NVR = \frac{(T_{preamp} + T_Q)G_{JPC} + T_{HEMT}}{T_{HEMT}}. \quad (7.1)$$

The Noise Visibility Ratio (NVR) of a mode of amplification can be observed using the Spectrum Analyzer. We normalize the signal by normalizing the received power with the

amplifier on by the the received power when just the HEMT amplifier is on. The NVR of the other two multiparametric modes of amplification are not significantly different than a single gain pump, which is typically $8 \text{ dB} \pm 1 \text{ dB}$.

In contrast, the NVM of the GG pumping scheme, shown in Fig. 87 is moderately higher than the singly pumped gain amplification when both are tuned up to produce gain at 20 dB. The singly pumped gain profile (green) has a noise rise of $\approx 8 \text{ dBm}$, but the GG amplification peaks at $\approx 12 \text{ dBm}$ (blue).




Figure 87: The noise rise of GG vs. G amplification at 20 dB. This shows that the noise produced by GG is moderately higher than G. When NVR of amplifiers creating the same amount of gain, are compared, the larger noise rise will adds extra noise to the system and degrades the quantum efficiency.

Again, since both types of amplification are producing gain of the same magnitude, it is fair to compare them side by side. For the same gain, the amplifier that is producing the lowest NVR is the one that is closer to approaching the quantum limit of noise. This implies that the GG amplifier adds more noise to the read-out and thus, is not as quantum-limited as a single gain tone amplifier. This is another downside of the GG amplifier, but for the read-out of a single qubit, a trade off of greater noise for a larger bandwidth might not be terrible.

However, if we go back to the spectrum analyzer and examine what happens when we input an additional tone on resonance at the signal frequency, we see that multiple corresponding “idler” tones appear in the mode, each separated by a frequency shift of 2Δ where Δ is the frequency detuning of the two pumps. Looking at the idler mode, we see something similar, which shows that the theory of idlers coupling one another in Fig. 15 appears to be true. One input signal in one mode couples to two idlers in the other mode, each at $\pm\Delta$, and then these idlers create two tones back in the original mode, and this pattern continues. It seems that this “idler forest” continues until it reaches the linewidth of the resonators. In a typical amplifier, the modes couple together on resonance and add noise to the system. More idler tones implies more coupling, and thus, creates more noise in the system, as evidenced by Fig. 87.




Figure 88: When one input tone is applied in the GG pumping scheme, it creates multiple idler tones in the signal and idler modes, each one 2Δ apart. All of these tones couple together and create additional noise channels in this method of amplification.

7.2 Arrayed amplifiers

We want to take advantage of our most successful multiparametric pumping scheme, GCI, as it finally allows for a device that is matched in reflection. This feature opens up the possibility of arraying amplifiers together, as we described in Chapter 3. We began by fabricating two Lumped JPC's in series on the same chip, as shown in Fig. 89.




Figure 89: Microscopic image of LJPCs in series. The two JPCs are at the center of their respective signal and idler resonator pads, and communicate through the coupling capacitor in the center. The input and output ports are connected to the signal modes, while the pump ports are coupled through a very weak coupling capacitor to the idler modes.

The two JPCs are coupled together by a coupling capacitor in the center. The input/output ports are connected to the two signal ports and the pump ports are weakly coupled to the idler modes with a coupling capacitor.

This device was actually cooled down with two magnets attached to the JPC box containing the array. One was placed in the usual location, on the back of the device, but the other was taped to the top lid of the box. The first magnet was placed approximately equally between both of the JPCs, but the second was placed closer to the second. We did this intentionally because we wanted to have individual control over one JPC, without

effecting the other substantially in order to bring the two JPCs' idler modes into resonance with one another.

Early samples of this design suffered from some very odd flux sweeps which had frequency jumps and led us to believe there were phase-slips occurring in the JRM. To combat this, we introduced small flux holes in the ground plane near the JRMs, which aren't included in the earlier version of the device we imaged here. Since we have incorporated the flux holes we have not observed significant jumps in the flux sweeps.




Figure 90: Flux sweeps in reflection from the pump to the signal ports of the array amplifier. This essentially measures each JPC, which lets us see the individual signal and common modes. Of course, if frequencies from any of these modes overlap then avoided crossings are created in the flux sweep, which we see in the white dotted box. We believe the avoided crossing highlighted here is from the idler mode crossing the common mode. The idler mode is visible here due to hybridization, it is usually not easily seen in reflection from the signal port.

Our best array sample was the first one we created with the addition of the flux holes. The flux sweeps of this device in reflection are shown in Fig. 90. These flux sweeps are taken in reflection from both signal ports. This lets us see that the signal and common mode from each JPC are separated from one another. The signal frequencies of the JPCs 1 and 2 were designed to be 11 and 9.5 GHz, respectively, but both are moderately lower than that. They are, however, are well separated in frequency which is the most crucial element. Another feature worth noting in flux sweep (a) is the avoided crossing in the white dotted box.

The avoided crossings signify that two modes come close enough together at this bias point to interact with one another. We would expect to be able to create an avoided crossing like this between the two idlers by tuning the bias of one of the JRMs. However, both magnets moved both JPCs at the same time, and it was very difficult to hold one at a constant flux bias. Even though we tried to place the magnets at optimal points for individual control, magnet 1 created a flux period of 0.27 mÅ for JPC 1 and 0.261 for JPC 2. Magnet 2 created a period of 4.8 mÅ for JPC1 and 5.15 mÅ for JPC 2. Even though these are not identical, which is good, they are not different enough to allow for individual control within the 10 mÅ range we are limited to without the risk of heating up the fridge.




Figure 91: Flux sweep in transmission of the array amplifier. This is taken by inputting signal from the pump port on JPC 2 and reading out from the signal 2 port. One can identify the signal and common mode of this JPC, but also both idler modes from each JPC.

The transmission flux sweep shown in Fig. 91 is produced from inputting signal into the first pump port and reading out from the signal port. Not only is another avoided crossing visible in this sweep, but we can also see both idler modes from each JPC. This is actually a good thing, because we want the idlers to hybridize with one another when tuned into resonance with one another, so the fact that they live in both JRMs is not so surprising. They are both close to their designed frequency of 5.5 GHz and are only separated by about 300 MHz. It is hard to say for certain which modes are hybridizing in the white dotted line region, although we believe it is one of the idlers and the other common mode, which isn't visible in this sweep. This area would have to be avoided when pumping the device, since we only want the idlers to couple to one another. Gain was achieved separately in both JPCs of this device. However, we were not able to hybridize the idlers in order to achieve gain across the array in its entirety, as described in Chapter 3.

In practice, this didn't work because the JRMs are only about 500 μm from one another, and the magnet is almost 2 cm in diameter. We tried multiple placements for the second magnet and calculated exactly where to place it to have the most independent control. However, none worked to the level that we required. We had a smaller, thinner magnet created in order to solve this problem. The small magnet was intended to be placed in the box itself, and be screwed to the lid just above the second JRM. However, before we had the chance to test this magnet idea, we developed an entirely new platform that should get rid of the magnet bias problem altogether— the SHARC.

7.3 Superconducting Hybridized ARbitrary Cavity-based amplifier (SHARC)

Since biasing two individual idler modes together became more trouble than we anticipated, we decided to pursue an entirely new design that could act as an array of coupled modes, with a single nonlinear element controlling all modes. We decided that the easiest way to do this was to remove the JRM entirely (as its 3-mode structure is complicated to connect to anything other than three microwave modes), and replace it with a SNAIL element which could act as its own, individual mode which also added the necessary nonlinear

coupling to the rest of the system.

After going through many iterations of the SHARC amplifier, we settled upon the design in Fig. 46. This design features 3 lumped resonators (which we refer to as a , b , and c) at frequencies 5.406, 7.442, 8.496 GHz, respectively. The SNAIL mode, s modulates from ≈ 6.6 –4.6 GHz. The linewidths of these modes are: $\kappa_s = 12$ MHz, $\kappa_a = 10$ MHz, $\kappa_b = 10$ MHz, and $\kappa_c = 60$ MHz. The modulation of the SNAIL crosses the a mode, and by fitting the flux sweeps and hybridization with our theoretical prediction, we are able to extract the strength of the g_2 coefficient, which is about 150 MHz, close to our theoretical prediction of 130 MHz. Operating the device when the two modes are strongly hybridized is not possible, as photons would constantly be leaked between modes. We can use this flux sweep to also show where these “forbidden” bias points are, which we define to be within 10% of the max coupling value, marked by red bars in Fig. 92.




Figure 92: Flux sweep of the SHARC amplifier showing the avoided crossing between the SNAIL mode and mode a . The flux sweeps can be fit with our theoretical model (red dotted lines) and allows us to extract the g_2 coupling strength, 150 MHz. The red solid bars indicate where the hybridization between the modes is too strong to have the device be properly operated here.

Since this is a brand new type of amplifier, our first order of business was to ensure that gain and conversion were possible between each pair of modes, as shown in Fig. 93.

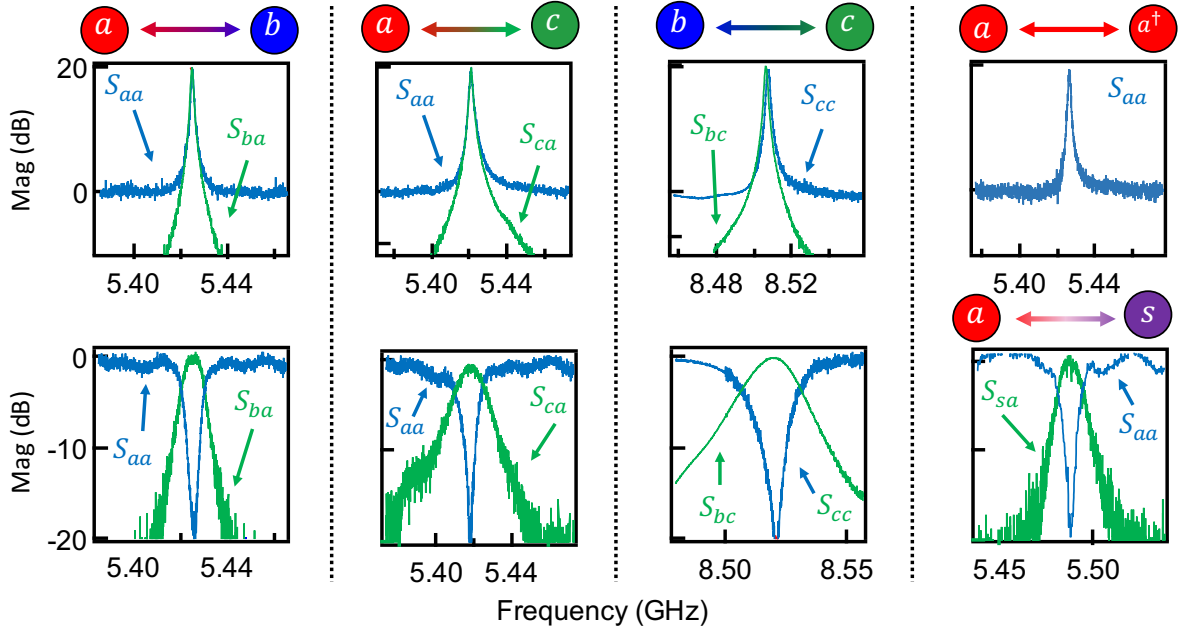


Figure 93: Example gain and conversion processes between different modes in the SHARC amplifier. Gain (top) and conversion (bottom) has been achieved between all combinations of linear modes. We can also create gain easily using the X²Z pumping scheme in every mode, shown here in mode *a*, for example. We can also find gain and conversion between combinations of linear modes and the SNAIL mode, shown in bottom right corner.

One interesting thing to note about the SHARC amplifier is that the Hamiltonian inherently contains degenerate 3-body mixing terms, such as g_{aas} , at significant magnitude throughout the range of the flux sweep. This is due to the fact that the SNAIL is already a one mode device, and upon hybridization, can share this attribute with the linear modes. This implies that X²Z -type gain can be found anywhere in the flux modulation, without having to worry about finding a sweet spot, evidenced by the right-hand plot in the above figure.

Of course, the purpose of developing the SHARC amplifier was to have a platform for studying multiparametric processes. A few different examples of these, including a circulator and GC are discussed in [89]. Additionally, the SHARC was meant to replace the arrayed JPC amplifier, which was too hard to physically bias.

We want to see if it is possible to create the same type of couplings between modes in the SHARC to make an array. However, to do so, requires a very high Q mode to be in the center of the input/output mode, that this amplifier doesn't have. We designed the b -mode to have a linewidth on the order of 100 KHz, but apparently a fabrication or wirebonding error made this not the case.

However, we can still use the amplifier we have for a proof-of-concept scheme. Ideally, we should use the two lowest Q modes as the input/output mode, and the highest Q in the center. The only issue is that three of the modes have approximately the same κ and only one is significantly larger. We first attempted to not use the larger, c mode at all, and have the b mode be in the center, coupled to a and s , but we were having a difficult time getting stable conversion between b and a for an unknown reason.

Instead, we decided to put the SNAIL mode in the center, as it couples strongly to every mode, and have a be the input mode, and b be the output mode. We decided to ignore c altogether for this attempt, as we hypothesized something was wrong with the mode, leading incorrect bandwidth.

The array should feature GCI coupling between the first two modes in the series, in this case, a and s , followed by conversion between the last two, s and b . To do this, we first look for full conversion from mode a to s , and s to b . In order to have a fair point which we call our "zero" to compare gain to in transmission, we first examine conversion between a and b , as shown below in Fig. 94.




Figure 94: Conversion in the SHARC amplifier from mode a to c . The peak reaches -30.5 dB, which will be our zero point in transmission.

Looking in transmission, we find that the peak, symbolizing full conversion, reaches -30.5 dB. This will become our normalized zero point.

Next we break up this conversion process into two pumps: a to s and s to b . When the strengths of these pumps are equal, we observe an interesting pattern shown in Fig. 95.




Figure 95: Two pump tone conversion between modes a and s and between s and c . The profile shapes shown here are very typical for what we should see in transmission and reflection when frequencies are aligned between the two pumps.

We also note that in transmission, the peak lies ≈ 1 -2 dB below the zero point, implying that even conversion with two pumps still loses some photons.

Once we have a good normalized zero point, we turn this conversion pump off and look for GC amplification between a and s . After GC amplification has been achieved, we turn the conversion pump between a and s up to even higher powers to create GCI, and look for its characteristic dip in reflection. When this is found, we then turn on the other conversion pump again and look in transmission at the whole array.

At first, no gain is observed, but by slowly increasing the strength of the pumps individually, we can increase the peak in transmission. As a very simple, proof-of-concept, we were able to find ≈ 1.8 dB of gain stretching across the entire array, shown in Fig. 96.




Figure 96: Preliminary GCI+C pumping scheme data from the SHARC amplifier. When all three pumps are turned on at once, we were able to achieve a small amount (≈ 2 dB) gain in transmission across the entire array of modes. In reflection, our response remained matched up to -10 dB. Higher gain was difficult to achieve in this sample because the middle mode did not have a lower Q than the other two modes.

This is, of course, not a finished product, as this device's ability to produce gain across the array was severely hampered by the lack of a high- Q mode. However, every other component necessary is there, and fixing the bandwidth of one mode should be trivial once fabrication facilities reopen. We remain very optimistic about the SHARC as a new platform for quantum-limited amplification and multiparametric study.

8.0 Conclusions And Outlook

This thesis has focused on two different aspects of superconducting quantum information research: quantum-limited amplifiers and the incorporation of novel vdW materials. In our pursuit of the ideal, quantum-limited amplifier, we explored multiparametric pumping schemes (Chapter 3) and found that specific combinations of pump frequencies can enhance characteristics of the JPC amplifier.

Although we have not yet created the perfect amplifier with the necessary scattering parameters, we have shown GCI coupling is only one scattering parameter short of this goal. In order to build a JPC that could produce high gain from multiple pumps, we first had to build a new type of amplifier that minimized higher order Hamiltonian effects. We redesigned the JPC (Chapter 4), taking into account stray inductances and participation ratios, and found in order to cancel unwanted 4th order terms, we had to shunt the JPC with large, linear inductances.

The evolution of the JPC took a turn at this point, and we went from shunting the JPC with Josephson junctions to strips of aluminum to take advantage of the inherent geometric inductance. However, these designs all had various problems, but we ended up discovering we could use NbTiN for our linear shunts instead, which is an amazing material for this purpose because it grows no native oxide and has a boosted kinetic inductance (Chapter 6). In addition, we did away with the microstrip version of the JPC and replaced it with a lumped design, which not only made the entire amplifier more compact, but features a more precise Hamiltonian.

A very clean Hamiltonian allows us, as quantum engineers, to be more creative and accurate when we design and fabricate these amplifiers. We were able to test multiple multiparametric schemes in one of these new, shunted JPCs, and verified that GCI pumping is indeed possible. Using it not only broadens the bandwidth of the JPC, but it erases the gain-bandwidth product and reflections from the device (Chapter 7).

We believe that GCI pumping can be used in the future to create a fully directional amplifier, and explored the possibility of doing this in an array-type amplifier. The physical

array we created from multiple JPCs in series was very hard to bias, and quickly evolved into the new platform we are calling the SHARC amplifier. Although the SHARC data presented is very new, we are optimistic about the future of this platform, as it allows for individual control over every mode in the system, can be biased with one magnet only, is inherently arrayable [111], and is designed for multiparametric pumping. With some sadness, in the future, we expect to move on from the JPC entirely and replace it with the SHARC amplifier.

However, there is still reason to believe we might need to re-engineer our amplifiers further. The junctions made from aluminum in all of these devices are still very nonlinear, which can actually be a disadvantage in amplifiers, since higher order terms become stronger and are difficult to control. That is one reason we began to look to other materials, especially superconducting TMD flakes, especially NbSe₂, to see if their unique properties could be used in conjunction with traditional materials (Chapter 5). We can now make junctions and SQUIDS from these flakes and are currently studying the inherent kinetic inductance of the flakes themselves, which we have already seen to be highly useful in superconducting circuits.

Future work will hopefully combine the study of multiparametric pumping and these novel materials to build an ideal, quantum-limited amplifier. In addition, there is still a good deal of work to be done in alleviating confusions surrounding the saturation power of these devices. We have a new working theory for saturation power of singly-pumped JPCs [76], but there is much more room for theorists and instrumentalists alike to explore if this holds for multiple pumps, SNAILS, arrays of SNAILS, etc. Still, this work addresses many of the limitations of state-of-the-art quantum amplifiers, and by studying their nonlinear behavior in depth, hopefully the work presented in this dissertation can be built upon to significantly progress the field of superconducting quantum computing.

Appendix A Recipes For Ebeam Lithography

A.1 Recipe for standard JPC Amplifier

Table 5: MMA 13 spin recipe

Time	RPM (pH)
10	500
5	2000
90	2000
10	0

Bake at 175 degrees Fahrenheit for 90 seconds.

Table 6: PMMA 950K A4 spin recipe

Time	RPM
10	500
5	3000
90	3000
10	0

Bake at 175 degrees Fahrenheit for 15 minutes.

Liftoff: Acetone at 65 degrees for 20 minutes.

A.2 Recipe for NbTiN shunts

Table 7: PMMA 950K A4 spin recipe for NbTiN shunts

Time	RPM
10	500
8	6000
60	6000
10	0

Bake at 175 degrees Fahrenheit for 15 minutes.

Table 8: Shunt development recipe

Solvent	Time
MIBK	1 min
IPA	1 min

Liftoff: Acetone for at least 1 min.

Appendix B Recipes for Optical Lithography

B.1 Recipe for NbTiN base layer

Table 9: LOR 5B spin recipe

Time	RPM
5	500
45	3000
10	0

Bake at 195 degrees Fahrenheit for 9 minutes.

Table 10: S1805 spin recipe

Time	RPM
5	500
60	4000
10	0

Bake at 115 degrees Fahrenheit for 3 minutes.

Table 11: Optical development recipe

Solvent	Time
351/DI water 4:1	75 secs
DI water	10 secs
400K	45 secs
DI water	10 secs

Bibliography

- [1] M. H. Devoret and R. J. Schoelkopf. Superconducting circuits for quantum information: An outlook. *Science*, 339:1169, 2013.
- [2] U. Vool and M. H. Devoret. Introduction to quantum electromagnetic circuits. *International Journal of Circuit Theory and Applications*, 45:897–934, 2017.
- [3] L. Sun, A. Petrenko, Z. Leghtas, Z. Leghtas, G. Kirchmair, K. M. Sliwa, M. Hatridge A. Narla, S. Shankar, J. Blumoff, L. Frunzio, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf. Tracking photon jumps with repeated quantum non-demolition parity measurement. *Nature*, 511:444–448, 2014.
- [4] J. M. Martinis, M. H. Devoret, and J. Clarke. Quantum josephson junction circuits and the dawn of artificial atoms. *Nature Physics*, 16:234–237, 2020.
- [5] F. Yan T. P. Orlando S. Gustavsson P. Krantz, M. Kjaergaard and W. D. Oliver). A quantum engineer’s guys to superconducting qubits. *Appl. Phys. Rev.*, 6:21318, 2019.
- [6] Richard P. Feynman. ”there’s plenty of room at the bottom”. *Engineering and Science*, 23:pp. 22–36.
- [7] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2004.
- [8] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O’Brien. Quantum computers. *Nature*, 464:45–53, 2010.
- [9] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Review*, 41(2):303–332, 1999.
- [10] L.K. Grover. A fast quantum mechanical algorithm for database search. *arXiv*.
- [11] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, and J. I.-J. Wang. Superconducting qubits: current state of play. *Annu. Rev. Condens. Matter Phys.*, 11, 2020.

- [12] J. I-Jan Wang, L. Bretheau D. Rodan-Legrain, D. L. Campbell, B. Kannan, D. Kim, M. Kjaergaard, P. Krantz, G. O. Samach, F. Yan, J. L. Yoder, K. Watanabe, T. Taniguchi, T. P. Orlando, S. Gustavsson, P. Jarillo-Herrero, and W. D. Oliver. Coherent control of a hybrid superconducting circuit made with graphene-based van der waals heterostructures. *Nature Nanotech*, 14:120–125, 2019.
- [13] A. M. Place, L. V. H. Rodgers, P. Mundada, B. M. Smitham, M. Fitzpatrick, Z. Leng, A. Premkumar, J. Bryon, S. Sussman, G. Cheng, T. Madhavan, H. K. Babla, B. Jaeck, A. Gyenis, N. Yao, R. J. Cava, N. P. de Leon, and A. A. Houck. New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds. *arXiv:2003.00024*, 2020.
- [14] D. Koch, B. Martin, S. Patel, L. Wessing, and P. M. Alsing. Benchmarking qubit quality and critical subroutines on ibm’s 20 qubit device. *arXiv:2003.01009*, 2020.
- [15] D. Schuster. *Circuit Quantum Electrodynamics*. PhD thesis, Yale University, 2007.
- [16] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf. Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation. *Phys. Rev. A*, 69:062320, Jun 2004.
- [17] L. DiCarlo, J. M. Chow, J. M. Gambetta, Lev S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf. Demonstration of two-qubit algorithms with a superconducting quantum processor. *Nature*, 460(7252):240–244, July 2009.
- [18] J. M. Chow, L. DiCarlo, J. M. Gambetta, F. Motzoi, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf. Optimized driving of superconducting artificial atoms for improved single-qubit gates. *Phys. Rev. A.*, 82:040305, 2010.
- [19] T. Yamamoto, Y. A. Pashkin, O. Astafiev, Y. Nakamura, and J. S. Tsai. Demonstration of conditional gate operation using superconducting charge qubits. *Nature*, 425:941–944, 2003.
- [20] J. R. Wootton and D. Loss. Repetition code of 15 qubits. *Phys. Rev. A*, 97, 2018.
- [21] F. Arute, K. Arya, R. Babbush, and et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574, 2019.

- [22] J. A. Schreier, A. A. Houck, Jens Koch, D. I. Schuster, B. R. Johnson, J. M. Chow, J. M. Gambetta, J. Majer, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Suppressing charge noise decoherence in superconducting charge qubits. *Phys. Rev. B*, 77:180502, May 2008.
- [23] D. P. DiVincenzo. The physical implementation of quantum computation. *Fortschr. Phys.*, 48:771, 2000.
- [24] J. E. Johnson, C. Macklin, D. H. Slichter, R. Vijay, E. B. Weingarten, John Clarke, and I. Siddiqi. Heralded state preparation in a superconducting qubit. *Phys. Rev. Lett.*, 109:050506, 2012.
- [25] M. Takita, A. W. Cross, A.D. Córcoles, J. M. Chow, , and J. M. Gambetta. Experimental demonstration of fault-tolerant state preparation with superconducting qubits. *Phys. Rev. Lett.*, 119:180501, 2017.
- [26] D. Riste, M. Dukalski, C. A. Watson, G. de Lange, M. J. Tiggelman, Ya. M. Blanter, K. W. Lehnert, R. N. Schouten, and L. DiCarlo. Deterministic entanglement of superconducting qubits by parity measurement and feedback. *Nature*, 502:350–354, 2013.
- [27] S. Shankar, M. Hatridge, Z. Leghtas, K.M. Sliwa, A. Narla, U. Vool, S.M. Girvin, L. Frunzio, M. Mirrahimi, and M. H. Devoret. Autonomously stabilized entanglement between two superconducting quantum bits. *Nature*, 504:419–422, 2013.
- [28] M. Silveri, E. Zolys-Geller, M. Hatridge, Z. Leghtas, M. H. Devoret, and S. M. Girvin. Theory of remote entanglement via quantum-limited phase-preserving amplification. *Phys. Rev. A*, 93, 2016.
- [29] T. Walter, P. Kurpiers, S. Gasparinetti, P. Magnard, A. Potocnik, Y. Salathe, M. Pecha, M. Mondal, M. Oppliger, C. Eichler, and A. Wallraff. Rapid high-fidelity single-shot dispersive readout of superconducting qubits. *Phys Lett A*, 2017.
- [30] Gottesman D. *Stabilizer Codes and Quantum Error Correction*. PhD thesis, California Institute of Technology, 1997.
- [31] C. M. Caves. Quantum limits on noise in linear amplifiers. *Phys. Rev. D*, 26:1817–1839, Oct 1982.

- [32] C. M. Caves, J. Combes, Z. Jiang, and S. Pandey. Quantum limits on phase-preserving linear amplifiers. *Phys. Rev. A*, 86:063802, Dec 2012.
- [33] P. Campagne-Ibarcq, E. Flurin, N. Roch, D. Darson, P. Morfin, M. Mirrahimi, M. H. Devoret, F. Mallet, and B. Huard. Persistent control of a superconducting qubit by stroboscopic measurement feedback. *Phys. Rev. X*, 3:021008, May 2013.
- [34] Flavius Dietrich Octavian Schackert. *A Practical Quantum-Limited Parametric Amplifier Based on the Josephson Ring Modulator*. PhD thesis, Yale University, 2013.
- [35] M. M. Farraday. On a peculiar class of acoustical figures; and on certain forms assumed by a group of particles upon vibrating elastic surfaces. *Phil. Trans. Roy. Soc.*, 258, 1831.
- [36] J. A. Adam. *Mathematics in Nature: Modeling Patterns in the Natural World*. 2011.
- [37] A. Van der Ziel. *Noise*. 1954.
- [38] R. Vijay. *Josephson Bifurcation Amplifier: Amplifying quantum signals using a dynamical bifurcation*. PhD thesis, Yale University, 2008.
- [39] M. A. Castellanos-Beltran, K. D. Irwin, G. C. Hilton, L. R. Vale, , and K. W. Lehnert. Amplification and squeezing of quantum noise with a tunable josephson metamaterial. *Nat. Phys.*, 4:929–931, 2008.
- [40] N. Roch, E. Flurin, F. Nguyen, P. Morfin, P. Campagne-Ibarcq, M. H. Devoret, and B. Huard. Widely tunable, nondegenerate three-wave mixing microwave device operating near the quantum limit. *Phys. Rev. Lett.*, 108:147701, Apr 2012.
- [41] T. Yamamoto, K. Inomata, M. Watanabe, K. Matsuba, T. Miyazaki, W. D. Oliver, Y. Nakamura, and J. S. Tsai. Flux-driven josephson parametric amplifier. *Applied Physics Letters*, 93(4):042510, 2008.
- [42] M. Hatridge, R. Vijay, D. H. Slichter, John Clarke, and I. Siddiqi. Dispersive magnetometry with a quantum limited squid parametric amplifier. *Phys. Rev. B*, 83:134501, Apr 2011.
- [43] R. Vijay, M. H. Devoret, and I. Siddiqi. Invited review article: The josephson bifurcation amplifier. *Review of Scientific Instruments*, 80(11):–, 2009.

- [44] C. Macklin, K. O'Brien, D. Hover, M. E. Schwartz, V. Bolkhovskiy, X. Zhang, W. D. Oliver, and I. Siddiqi. A near-quantum-limited josephson traveling-wave parametric amplifier. *Science*, 350(6258):307–310, 2015.
- [45] T. C. White, J. Y. Mutus, I.-C. Hoi, R. Barends, B. Campbell, Yu Chen, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, J. Kelly, A. Megrant, C. Neill, P. J. J. O'Malley, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, S. Chaudhuri, J. Gao, and John M. Martinis. Traveling wave parametric amplifier with josephson junctions using minimal resonator phase matching. *Appl. Phys. Lett.*, 106(24):242601, 2015.
- [46] A. B. Zorin. Josephson traveling-wave parametric amplifier with three-wave mixing. *Phys. Rev. Applied*, 16:034006, 2016.
- [47] N. E. Frattini, U. Vool, S. Shankar, A. Narla, K. M. Sliwa, and M. H. Devoret. 3-wave mixing josephson dipole element. *Appl. Phys. Lett.*, 110:222603, 2017.
- [48] N. E. Frattini, V. V. Sivak, A. Lingenfelter, S. Shankar, and M. H. Devoret. Optimizing the nonlinearity and dissipation of a snail parametric amplifier for dynamic range. *Phys Rev Applied*, 10:054020, Nov 2018.
- [49] T. Roy, S. Kundu, M. Chand, A. M. Vadiraj, A. Ranadive, N. Nehra, M. Patankar, J. Aumentado, A. A. Clerk, and R. Vijay. Broadband parametric amplification with impedance engineering: beyond the gain-bandwidth product. *Applied Physics Letters*, 2015.
- [50] F. Lecocq, L. Ranzani, G. A. Peterson, K. Cicak, R. W. Simmonds, J. D. Teufel, and J. Aumentado. Nonreciprocal microwave signal processing with a field-programmable josephson amplifie. *Phys. Rev. Appl*, 7:024028, 2017.
- [51] V.V. Sivak, N.E. Frattini, V.R. Joshi, A. Lingenfelter, S. Shankar, and M.H. Devoret. Kerr-free three-wave mixing in superconducting quantum circuits. *Phys. Rev. Applied*, 11:054060, 2019.
- [52] K. M. Sliwa, M. Hatridge, A. Narla, S. Shankar, L. Frunzio, R. J. Schoelkopf, and M. H. Devoret. Reconfigurable Josephson Circulator/Directional Amplifier. *Phys Rev X*, 5:041020, 2015.
- [53] L. Ranzani and J. Aumentado. Circulators at the quantum limit: Recent realizations of quantum-limited superconducting circulators and related approaches. *IEEE Microwave Magazine*, 20(4):112–122, 2019.

- [54] T. Roy, S. Kundu, M. chand, A. M. Vadiraj, A. Ranadive, N. Nehra, M. P. Patankar, J. Aumentado, A. A. Clerk, and R. Vijay. Broadband parametric amplification with impedance engineering: Beyond the gain-bandwidth product. *Appl. Phys. Lett.*, 107:262601, 2017.
- [55] D. M. Pozar. *Microwave Engineering*. 1998.
- [56] G. Liu, T.-C. Chien, X. Cao, O. Lanes, E. Alpern, D. Pekker, and M. Hatridge. Josephson parametric converter saturation and higher order effects. *Appl. Phys. Lett.*, 111:202603, 2017.
- [57] J. Y. Mutus, T. C. White, R. Barends, Yu Chen, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, J. Kelly, A. Megrant, C. Neill, P. J. J. O’Malley, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, K. M. Sundqvist, A. N. Cleland, and John M. Martinis. Strong environmental coupling in a josephson parametric amplifier. *Appl. Phys. Lett.*, 104:263513, 2014.
- [58] C. Eichler and A. Wallraff. Controlling the dynamic range of a josephson parametric amplifier. *EPJ Quantum Technology*, 1(1):2, 2014.
- [59] A. Metelmann and A. A. Clerk. Nonreciprocal photon transmission and amplification via reservoir engineering. *Phys. Rev. X*, 5:021025, Jun 2015.
- [60] G. A. Peterson, F. Lecocq, K. Cicak, R. W. Simmonds, J. Aumentado, and J. D. Teufel. Demonstration of efficient nonreciprocity in a microwave optomechanical circuit. *Phys. Rev. X*, 7:031001, Jul 2017.
- [61] A. W. Tsen, B. Hunt, Y. D. Kim, Z. J. Yuan, S. Jia, R. J. Cava, J. Hone, P. Kim, C. R. Dean, and A. N. Pasupathy. Nature of the quantum metal in a two-dimensional crystalline superconductor. *Nature Physics*, 12(3):208–212, March 2016.
- [62] L. Wang, I. Meric, P. Y. Huang, Q. Gao, Y. Gao, H. Tran, T. Taniguchi, K. Watanabe, L. M. Campos, D. A. Muller, J. Guo, P. Kim, J. Hone, K. L. Shepard, and C. R. Dean. One-Dimensional Electrical Contact to a Two-Dimensional Material. *Science*, 342(6158):614–617, November 2013.
- [63] J. Clarke and A.I. Braginski. *The SQUID Handbook*. 2006.
- [64] M. R. Sinko, S. C. de la Barrera, O. Lanes, K. Watanabe, T. Taniguchi, D. Pekker, M. Hatridge, and B. M. Hunt. Superconducting edge contact and quantum inter-

ference between two-dimensional van der waals and three-dimensional conventional superconductors. *arXiv:1911.09711v1*.

- [65] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Charge-insensitive qubit design derived from the cooper pair box. *Phys. Rev. A*, 76:42319, 2013.
- [66] M.H Devoret and J. M Martinis. Implementing qubits with superconducting integrated circuits. *Quant. Infor. Proc.*, 3:1–5, 2004.
- [67] D. Vion P. Joyez D. Esteve Bouchiat, V. and M. H. Devoret). Quantum coherence with a single cooper pair. *Physica Scripta T76*, 165, 1998.
- [68] S. M. Girvin. Les houches lecture notes. *Phys. Rev. A*, 2012.
- [69] H. Paik, D. I. Schuster, L. S. Bishop, G. Kirchmair, G. Catelani, A. P. Sears, B. R. Johnson, M. J. Reagor, L. Frunzio, L. I. Glazman, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf. Observation of high coherence in josephson junction qubits measured in a three-dimensional circuit qed architecture. *Phys. Rev. Lett.*, 107:240501, Dec 2011.
- [70] A. Blais L. Frunzio R.-S. Huang J. Majer S. Kumar S. M. Girvin A. Wallraff, D. I. Schuster and R. J. Schoelkopf. Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics. *Nature*, 431:162–167., 2004.
- [71] S. Haroche, M. Brune, and J. M. Raimond. From cavity to circuit quantum electrodynamics. *Nature Physics*, 16:343–346, Jun 2020.
- [72] M. Reed. *Entanglement and Quantum Error Correction with Superconducting Qubits*. PhD thesis, Yale University, 2013.
- [73] N. Bergeal, F. Schackert, M. Metcalfe, R. Vijay, V. E. Manucharyan, L. Frunzio, D. E. Prober, R. J. Schoelkopf, S. M. Girvin, and M. H. Devoret. Phase-preserving amplification near the quantum limit with a josephson ring modulator. *Nature*, 465(7294):64–68, May 2010.
- [74] N. Bergeal, R. Vijay, V. E. Manucharyan, I. Siddiqi, R. J. Schoelkopf, S. M. Girvin, and M. H. Devoret. Analog information processing at the quantum limit with a josephson ring modulator. *Nat. Phys.*, 6(4):296–302, April 2010.

- [75] E. Flurin, N. Roch, F. Mallet, M. H. Devoret, and B. Huard. Generating entangled microwave radiation over two transmission lines. *Phys. Rev. Lett.*, 109:183901, Oct 2012.
- [76] C. Liu, T.-C. Chien, M. Hatridge, and D. Pekker. Optimizing josephson-ring-modulator-based josephson parametric amplifiers via full hamiltonian control. *Phys. Rev. A*, 101, 2020.
- [77] F. Schackert, A. Roy, M. Hatridge, M. H. Devoret, and D. A. Stone. Three-wave mixing with three incoming waves: signal-idler coherent attenuation and gain enhancement in a parametric amplifier. *Phys. Rev. Lett.*, 111:073903, 2013.
- [78] Bogdan A. Kochetov and Arkady Fedorov. Higher-order nonlinear effects in a josephson parametric amplifier. *Phys. Rev. B*, 92:224304, Dec 2015.
- [79] B. Yurke, L. Corruccini, P. Kaminsky, L. Rupp, A. Smith, A. Silver, R. Simon, and E. Whittaker. Observation of parametric amplification and deamplification in a josephson parametric amplifier. *Phys. Rev. A*, 39:2519, 1989.
- [80] A. Kamal, J. Clarke, and M. H. Devoret. Gain, directionality, and noise in microwave squid amplifiers: Input-output approach. *Phys. Rev. B*, 86:144510, 2012.
- [81] L. Ranzani and J. Aumentado. Graph-based analysis of nonreciprocity in coupled-mode systems. *New J of Phys*, 17:023024, 2015.
- [82] A. A. Clerk, M. H. Devoret, S. M. Girvin, Florian Marquardt, and R. J. Schoelkopf. Introduction to quantum noise, measurement, and amplification. *Rev. Mod. Phys.*, 82:1155–1208, Apr 2010.
- [83] S. K. Choi, R. D. Li, C. Kim, and P. Kumar. Traveling-wave optical parametric amplifier: investigation of its phase-sensitive and phase-insensitive gain response. *Opt. Soc. Am. B*, 14:1564–1575, 1997.
- [84] P. A. Andrekson Z. Tong M. Karlsson P. Petropoulos F. Parmigiani J. Kakande, C. Lundstrom and D. J. Richardson. Detailed characterization of a fiber-optic parametric amplifier in phase-sensitive and phase-insensitive operation. *Opt. Express*, 18:4130–4137, 2010.
- [85] A. A. Nadeem. Theory of a two-mode phase-sensitive amplifier. *Phys. Rev. A*, 46:1560–156, 1992.

- [86] B. Abdo, A. Kamal, and M. H. Devoret. Nondegenerate three-wave mixing with the josephson ring modulator. *Phys. Rev. B*, 87:014508, Jan 2013.
- [87] A. Metelmann and A. A. Clerk. Quantum-limited amplification via reservoir engineering. *Phys. Rev. Lett.*, 112:133904, Apr 2014.
- [88] A.A. Metelmann. *Private Communications*.
- [89] T.-C. Chien. *Creating directional quantum-limited amplification using multiple parametric drives*. PhD thesis, University of Pittsburgh, 2020.
- [90] R. Vijay, D. H. Slichter, and I. Siddiqi. Observation of quantum jumps in a superconducting artificial atom. *Phys. Rev. Lett.*, 106:110502, Mar 2011.
- [91] T.-C. Chien, O. Lanes, C. Liu, X. Cao, P. Lu, S. Motz, G. Liu, D. Pekker, and M. Hatridge. Multiparametric amplification and qubit measurement with a kerr-free josephson ring modulator. *arXiv:1903.02102v2*.
- [92] T. Yamamoto, K. Inomata, M. Watanabe, K. Matsuba, T. Miyazaki, W. D. Oliver, Y. Nakamura, and J. S. Tsai. Flux-driven josephson parametric amplifier. *Applied Physics Letters*, 93:042510, 2008.
- [93] J. D. Pillet, E. Flurin, F. Mallet, and B. Huard. A compact design for the josephson mixer: The lumped element circuit. *Applied Physics Letters*, 105, 2015.
- [94] M. R. Perelshtein I. Lilja E. V. Korostylev T. Elo, T. S. Abhilash and P. J. Hakonen). Broadband lumped-element josephson parametric amplifier with single-step lithography. *Appl. Phys. Lett.*, 114, 2019.
- [95] R. Movshovich B. Yurke, M. L. Roukes and A. N. Pargellis. A low-noise series-array josephson junction parametric amplifier. *Applied Physics Letters*, 69:3078–3080, 1996.
- [96] H. G. LeDuc H.E. Byeong, P. K. Day and J. Zmuidzinas. A wideband, low-noise superconducting amplifier with high dynamic range. *Nature Physics*, 8:623–627, 2012.
- [97] S. C. de la Barrera, M. R. Sinko, D. P. Gopalan, N. Sivadas, K. L. Seyler, K. Watanabe, T. Taniguchi, A. W. Tsen, X. Xu, D. Xiao, and B. M. Hunt. Tuning ising superconductivity with layer and spin–orbit coupling in two-dimensional transition-metal dichalcogenides. *Nat Commun*, 9:1427, 2018.

- [98] E. Navarro-Moratalla, J. O. Island, S. Mañas-Valero, E. Pinilla-Cienfuegos, A. Castellanos-Gomez, J. Quereda, G. Rubio-Bollinger, L. Chirolli, J. A. Silva-Guillén, N. Agraït, G. A. Steele, F. Guinea, Herre S., J. van der Zant, and E. Coronado. Enhanced superconductivity in atomically thin NbSe_2 . *Nat Commun*, 7:11043, 2016.
- [99] V. E. Manucharyan, J. Koch, L. Glazman, and M. H. Devoret. Fluxonium: Single cooper-pair circuit free of charge offsets. *Phys. Rev. A*, 326, 2009.
- [100] M. R. Vissers, J. Hubmayr, M. Sandberg, S. Chaudhuri, C. Bockstiegel, and J. Gao. Frequency-tunable superconducting resonators via nonlinear kinetic inductance. *Appl. Phys. Lett.*, 107:062601, 2015.
- [101] K.-H. Lee, S. Chakram, S. E. Kim, F. Mujid, A. Ray, H. Gao, C. Park, Y. Zhong, D. A. Muller, D. I. Schuster, and J. Park. Two-dimensional material tunnel barrier for josephson junctions and superconducting qubits. *Nano Lett.*, 19:8287–8293, 2019.
- [102] X. Xi, Z. Wang, W. Zhao, J.-H. Park, K. T. Law, H. Berger, L. Forró, J. Shan, and K. F. Mak. Ising pairing in superconducting NbSe_2 atomic layers. *Nature Phys*, 12:139–143, 2016.
- [103] L. Wang, I. Meric, P. Y. Huang, Q. Gao, Y. Gao, H. Tran, T. Taniguchi, K. Watanabe, L. M. Campos, D. A. Muller, J. Guo, P. Kim, J. Hone, K. L. Shepard, and C. R. Dean. One-dimensional electrical contact to a two-dimensional material. *Science*, 342:614–617, 2013.
- [104] Leila R. Vale Gene C. Hilton K. W. Lehnert M. Malnou, D. A. Palken. Optimal operation of a josephson parametric amplifier for vacuum squeezing. *Phys. Rev. Appl.*, 9, 2018.
- [105] M. Hatridge et. al. Quantum back-action of an individual variable-strength measurement. *Science*, 339:178, 2013.
- [106] K. Sliwa. *Improving the Quality of Heisenberg Back-Action of Qubit Measurements made with Parametric Amplifiers*. PhD thesis, Yale University, 2016.
- [107] K.W. Murch, S. J. Weber, C Macklin, and I Siddiqi. Observing single quantum trajectories of a superconducting quantum bit. *Nature*, 502:211–214, 2013.

- [108] C. Eichler, D. Bozyigit, and A. Wallraff. Characterizing quantum microwave radiation and its entanglement with superconducting qubits using linear detectors. *Phys. Rev. A*, 86:032106, Sep 2012.
- [109] F. Lecocq, L. Ranzana, G. A. Peterson, K. Cicak, R. W. Simmonds, J. D. Teufel, and J. Aumentado. Nonreciprocal microwave signal processing with a field-programmable josephson amplifier. *Phys. Rev. Appl.*, 7:024028, 2017.
- [110] W.-A. Li, G.-Y. Huang, and Y. Chen. Directional amplifiers in a hybrid optomechanical system. *J. Opt. Soc. Am. B*, 36(2):306–311, Feb 2019.
- [111] L. Planat, R. Dassonneville, J. P. Martínez, F. Foroughi, O. Buisson, W. Hasch-Guichard, C. Naud, R. Vijay, K. Murch, and N. Roch. Understanding the saturation power of josephson parametric amplifiers made from squid arrays. *Phys. Rev. Applied*, 11, 2019.