

**Parametric Cavity Quantum Electrodynamics With
Superconducting Circuits and Large Cavities**

by

Adam J. Sirois

B.A., Bowdoin College, 2002

M.S., University of Colorado-Boulder, 2009

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Physics

2016

This thesis entitled:
Parametric Cavity Quantum Electrodynamics With Superconducting Circuits and Large
Cavities
written by Adam J. Sirois
has been approved for the Department of Physics

Dr. Raymond Simmonds

Prof. John Price

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Sirois, Adam J. (Ph.D., Physics)

Parametric Cavity Quantum Electrodynamics With Superconducting Circuits and Large Cavities

Thesis directed by Dr. Raymond Simmonds

The research reported here details experimental progress toward future superconductor-based quantum computing technologies. Specifically, we present techniques for moving quantum information between various cavity resonators – spatially as well as between frequencies. Rather than using traditional *resonant* coupling though, we induce *non-resonant* coupling through parametric frequency conversion. This technology is mediated by chip-based, micro-fabricated, Josephson Junction circuits.

Parametric processes have the advantage over traditional coupling by allowing the constituents of the system – for example: cavity modes or qubits – to remain fixed in frequency, tuned to their optimal operation frequency, thereby avoiding unwanted resonant interactions. As the number of constituents grows – as quantum computing architectures expand to more and more bits – these techniques will become necessary to optimize performance.

Outside of the focus on quantum computation, these techniques have wide application for materials research, quantum optics, and extending traditionally optical-frequency experiments to microwave-frequencies. To that end, we demonstrate a hybrid quantum technology which expands the toolbox of superconducting quantum information to a new system, namely sapphire whispering gallery mode resonators. These resonators have been studied in the optical frequency domain, but here we demonstrate their usefulness at microwave frequencies.

All of the experiments in this dissertation should be considered proof-of-principle demonstrations of a future technology. Thus, we note when and where improvements for future devices will be necessary.

Dedication

To my wife, Gina.

Acknowledgements

My path to graduate school was serendipitous to say the least. After finishing undergrad I found myself wanting to work in a lab setting, but not wanting to pursue grad school (who would ever do that?). By chance I read an article about the Cryogenic Dark Matter Search in the popular press and naïvely decided to email one of the PIs asking for a job. *He* didn't hire me, but Dan Akerib at Case-Western *did* almost sight unseen. It was in that group that I really fell in love with physics and the day-to-day hands-on benchwork of physicists though still wasn't considering grad school. I especially thank Dan Akerib, Mike Dragowsky, Aaron Manalaysay, and Emily Dragowsky for their patience and support.

With my mind still not made up about graduate school, when my contract at CWRU ended I again cast about looking for employment. Martin Huber, a member of the CDMS collaboration with connections to NIST, passed along my resume to Ray Simmonds. Ray had recently taken over John Martinis' lab at NIST and was in need of warm bodies. The two years I spent as his lab tech were well-spent learning from Ray, Kat Cicak, Kevin Osborn and Josh Strong and cemented the idea that graduate school was probably a good idea.

Boy, was I wrong! Eight+ years later, I have out-lived two other graduate students: Jed Whittaker and Shane Allman; and many postdocs – chiefly Fabio Altomare, Dale Li, and Leonardo Ranzani. Fabio's unending energy and passion was a welcome pick-me-up, Dale is an extremely careful and well-rounded scientist (also always willing to chat about life and dole out good advice), and Leonardo's depth of knowledge in microwaves and mathematics was always helpful. A special thanks to Florent Lecocq is merited here as well – he has been

a good friend, a mentor, and level head to have around.

The big turning point in my grad school career came about in \sim 2012. This was essentially a *reset* of everything that had come before. We built up new lab space in a new building; our group merged with Joe Aumentado and John Teufel to form what is now the Advanced Microwave Photonics group; and Ray had two kids. At this point I was adopted by Joe and his two postdocs – Manuel Castellanos-Beltran and Mike DeFeo. I think without the combination of Joe, Manny, and Mike I would have struggled mightily. John Teufel’s consistency, passion, and deep understanding of *everything* microwaves served as the glue holding together the group often.

The current group – Ray, Joe, John, Florent, Kat – with the addition of Jeremy Clark, Shlomi Kotler, Xiaoyue Jin, Heli Vora, and Gabe Peterson is just now seeing the results of the new lab space and cleanroom. Shlomi and Jeremy have been a huge asset to the group – their background in atomic/optical physics has really helped us look at problems in a new light (har har).

I am also indebted to my collaborators at the University of Western Australia and to the American-Australian Association. The work with the sapphire whispering gallery mode resonator which features heavily in this dissertation again happened serendipitously. Mike Tobar, Eugene Ivanov, and Daniel Creedon gave a talk at NIST presenting some of their recent, unrelated, results and we wondered if their system could be used for quantum computing applications! They have been generous with their time and expertise and hosted me for a month in Perth with welcome arms. I hope I was able to do the same for Daniel when he visited us.

To close my (already very) long acknowledgements, I would really like to say a heartfelt thank you to everyone at NIST who has taken the time to help me over the years. In particular Jiansong Gao, Gene Hilton, Paul Dresselhaus, Maggie Crews, Dave Rudman, Jay Koch, and Norm Bergren have always been available and willing. NIST is full of world-class

physicists and world-class folks who help them do their jobs. Not once have I felt unwelcome or unwanted. NIST is a special place to do science.

Added post-defense: thank you to Murray Holland, John Price, Chuck Rogers, and Nils Halverson for making my dissertation defense enjoyable and for their insightful questions and comments!

Contents

Chapter	
1	Introduction 1
1.1	Historical motivation 3
1.1.1	Atomic physics 3
1.1.2	Josephson devices 5
1.2	Quantum information processing 6
1.3	Thesis overview 8
2	Background Physics 10
2.1	Microwave engineering 11
2.1.1	Transmission-lines 12
2.1.2	Lumped-element circuit analysis 15
2.2	Cavity Quantum Electrodynamics 19
2.2.1	Low-loss cavities 19
2.2.2	Quantization of a cavity field 31
2.2.3	Interactions between an atom and a cavity mode 34
2.3	Artificial atoms 41
2.4	Parametric interactions 46
2.4.1	Input-Output Introduction 47
2.4.2	Parametric Input-Output Model 51

2.4.3	Coupler circuit	56
3	Experimental Design	59
3.1	New lab space	59
3.2	Microwave measurement setup	61
3.3	Electromagnetic simulation and modeling	71
3.4	Micro-fabrication	75
3.4.1	Angle-evaporated junctions	76
3.5	Chapter conclusions	83
4	Motivation for Specific Experiments	84
4.1	Practical motivation	84
4.2	Parametric Hamiltonians	87
4.3	Quantum optics	87
5	Experiment I: Frequency conversion between square cavities	90
5.1	Device details and characterization	91
5.2	Frequency conversion between square cavities	96
5.3	Chapter conclusion	101
6	Experiment II: Freq. conversion between sapphire WGMs	103
6.1	Device details	104
6.2	Device characterization	109
6.3	Frequency conversion between modes	113
6.4	Chapter Conclusions	117
7	Experiment III: A qubit coupled to WGM	118
7.1	Device details	119
7.2	Qubit frequency inferred from splittings	119

7.3	Dispersive qubit-cavity measurements	122
7.4	Qubit characterization	127
7.5	Chapter conclusions	133
8	Experiment IV: Qubit and coupler and WGM	134
8.1	Device details	136
8.2	Sideband gate theory	137
8.3	Device characterization	141
8.4	Is high-power Jaynes-Cummings style readout possible?	146
8.5	Future improvements	147
9	Conclusions	149
	Bibliography	154
	Appendix	
A	List of relevant publications:	160

Tables

Table

2.1	Comparison of various cavity geometries.	34
6.1	List of high-Q sapphire modes	112
7.1	Vacuum Rabi anti-crossing and linewidth of the transmon and WGM	122
7.2	Extracted qubit parameters	133
8.1	Expected parametric coupling rates between transmon and WGM resonator	140

Figures

Figure

1.1	Sketch of <i>Von Neumann</i> classical and quantum computer	9
2.1	Equivalent ways to look at coax transmission lines	14
2.2	RLC circuit impedance and external loading	16
2.3	Cartoon of cavity quantum electrodynamics	20
2.4	Various cavity geometries and field distributions	23
2.5	Sketch of a single-mode field	32
2.6	Sketch of a single-mode field including an atom	35
2.7	Vacuum Rabi splitting modeled via the Jaynes-Cummings Hamiltonian . . .	40
2.8	Predicted transition frequency and anharmonicity for our transmon qubit . .	45
2.9	Cartoon of parametric frequency conversion	48
2.10	Lumped element model of coupled modes	52
2.11	Parametrically induced normal mode splitting theory	55
2.12	A comparison between qubit and coupler circuits	58
3.1	Schematic of the wiring inside a typical fridge	64
3.2	DC line filtering response	65
3.3	Instrument setup and schematic outside the fridge	67
3.4	Details of the IQ demodulation and digitization setup	70
3.5	Sample Microwave Office model	73

3.6	Sample HFSS model	74
3.7	Sketch of angle evaporation techniques	77
3.8	SEM of fabricated junction	79
3.9	Different resist styles	81
3.10	Summary of fabrication of coupler circuit	82
4.1	Summary of experimental goals	85
4.2	Two hierarchical views of the field of quantum computing	89
5.1	Double square cavity details.	92
5.2	Frequency vs flux characterization data	94
5.3	Parametric frequency conversion in the frequency- and time- domains	98
5.4	Store and retrieve time traces	100
5.5	Store and retrieve experiment with phase	102
6.1	Sketch of the orientation and setup for this experiment	104
6.2	Photo of the WGM crystal	108
6.3	Whispering gallery mode resonances vs coupler flux bias	110
6.4	Continuous-wave parametric frequency conversion between WGMs	113
6.5	Storage and retrieval between WGMs	115
6.6	Frequency conversion between high-Q WGM resonances	116
7.1	Vacuum Rabi splitting of the transmon with whispering gallery modes.	121
7.2	Strongly-driven Jaynes-Cummings behavior	125
7.3	Qubit 0-1 transition vs power	129
7.4	Qubit level spectroscopy	130
7.5	Driven Rabi and T_1 measurements of the qubit	132
8.1	Cartoon of storing a photon in the WGM resonances using frequency conversion.	135

8.2	Sapphire WGM, qubit, and coupler measurement sketch	136
8.3	Comparing the WGM frequency modulation between cooldowns	142
8.4	Vacuum Rabi anti-crossings compared between cooldowns	143
8.5	Vacuum Rabi splitting drift over time	145
8.6	Double-Kerr effect due to the qubit and coupler	148
9.1	Proposed improvements to the experiment	153

Chapter 1

Introduction

The fundamental ideas for quantum computing date to the early 1990s [1]. These ideas — using quantum mechanics to solve intractable problems for classical computers — have a close analogy to the development of classical computing in the 1940s. Just as with the early theories of Alan Turing, early quantum computing theorists have developed the *software* — algorithms, gates, etc — before even the *hardware* has been designed. With classical computing, it took decades of work for the hardware to be developed enough for complex computational programs to be run. The quantum computing world is in precisely the same situation today — a reliable, scale-able hardware is needed.

This is not to say that steady progress has not happened in the field of quantum computing. In a time span of approximately 30 years systems containing a few (up to about 10) quantum bits (qubits – the quantum version of a classical bit) have been built and tested [1]. Many of the proof-of-principle test algorithms have also been demonstrated [2, 3, 4]. But, the problem of developing a reliable quantum computer hardware of thousands to tens of thousands of qubits seems orders of magnitude harder than the equivalent problem using classical hardware.

It is also remarkable how many sub-fields of physics are involved in the quest for making a useful quantum computer. The potential technologies being explored to make qubits encompass almost every field of physics, excluding high-energy physics. For example, the leading contenders for quantum computing technologies cover condensed-matter physics:

superconducting [2] and semiconducting [5] devices; electromagnetism: photonic qubits [6]; atomic physics: trapped ions [7] or atoms [8]; and nuclear physics: nuclear magnetic resonance [9] or spins in solids [10].

Each technology has various pros and cons, but a useful metric one can use to differentiate them is comparing coherence times versus ease of control (or gate times). On one end of the spectrum are the ‘natural’ systems – trapped ions, atomic spins, nuclear spins, etc – and on the other are ‘man-made’ devices such as superconducting qubits or quantum dots. In this thesis we will often call the engineered, superconducting qubits ‘artificial atoms’.

On the natural side of things coherence times – the time a quantum state will remain coherent – are orders of magnitude higher than the man-made side. Seeing coherence times of seconds or more is commonplace compared to the best state of the art superconducting qubits which are approaching a few hundred microseconds [11, 12]. Alternatively, engineered systems allow the inclusion of control electrodes very close to the device and the device itself can be fabricated to couple to those fields strongly. The state of a superconducting qubit can be manipulated in a few tens of nanoseconds (compared to sometimes many microseconds for trapped ions)[13]. This trade-off between control and coherence will be described and emphasized continually in this dissertation.

The big lesson here is that it is likely that the eventual technology which a useful quantum computer will be made from will be some hybrid of man-made and natural quantum systems hopefully exploiting the strengths of both natural and engineered regimes. Clever schemes are continually being proposed to work on the problem from both ends — speeding up control times as well as increasing coherence — for all of the considered systems.

This thesis describes a proof-of-principle demonstration of some viable, scale-able techniques for coupling qubits (and quantum memories) as well as a novel hybrid system which exploits the benefits of natural systems as well as the engineered systems. Specifically, we describe using a fabricated (engineer-able) superconducting circuit to couple to a crystalline

sapphire resonator (natural). While this may not be the be-all-and-end-all technology for quantum computers, we feel it is an important blueprint of useful techniques and architectures.

1.1 Historical motivation

The possibility of combining quantum mechanics with man-made devices is a very modern idea – many of the foundational experiments were done in the 1980s – specifically, the development of Josephson parametric amplifiers and superconducting quantum interference devices (SQUIDs) [14, 15, 16]. Thus in order to motivate where this research fits in the broader physics domain, it is useful to describe a bit of historical background. Coincidentally, this also provides a good history of two research focuses of NIST-Boulder.

We begin with a brief review of atomic physics which serves as the inspiration for many of the artificial atom-ic experiments our group performs. We then show how the engineerability of superconducting circuits allows for exploration of parameter space unavailable to naturally occurring systems. Finally, the more practical side of these experiments – developing quantum information processing capabilities – is discussed.

1.1.1 Atomic physics

The development of quantum mechanics in the early 20th century has opened broad new avenues of research today. Famously, though, the early pioneers were dubious about the feasibility of such experiments. In 1952, Erwin Schrödinger wrote on the history of quantum theory, highlighting its contentiousness:

Planck was even more hesitant to adopt the view that radiation itself be divided up into portions or light-quanta or ‘photons,’ to use the present terminology. In all this his hesitance had good reasons. Yet only a few years later (1905) Einstein advanced the hypothesis of light-quanta, clinching it with irresistible arguments; and in 1913 Niels Bohr, by taking the discrete states of the atoms seriously and extending Planck’s assumptions in two

directions with great ingenuity, but irrefutable consistency, could explain quantitatively some of the atomic line spectra, which are all patently discrete, and which had in their entirety formed a great conundrum up to then: Bohr's theory turned them into the ultimate and irrevocable direct evidence, that the discrete states are a genuine and real fact. Bohr's theory held the ground for about a dozen of years, scoring a grand series of so marvellous and genuine successes, that we may well claim excuses for having shut our eyes to its one great deficiency: while describing minutely the so-called 'stationary' states which the atom had normally, i.e. in the comparatively uninteresting periods when nothing happens, the theory was silent about the periods of transition or 'quantum jumps' (as one then began to call them). [17]

It is remarkable, then, that beyond using the new theories for explaining physical phenomena such as the photoelectric effect or atomic spectra; physicists began thinking about and building experiments to manipulate simple quantum systems as early as the 1950s and 1960s [3]. Indeed, many of the early *gedankenexperiments* from the minds of the founders of quantum theory – say, being able to study isolated, single atoms – have become manifest in experimentalist's laboratories during the past few decades. To highlight this point, we contrast to what Schrödinger continued:

...we never experiment with just one electron or atom or (small) molecule. In thought-experiments we sometimes assume that we do; this invariably entails ridiculous consequences. In the first place it is fair to state that we are not experimenting with single particles, any more than we can raise Ichthyosauria in the zoo...We can never reproduce the same single-particle-event under planned varied conditions; and this is the typical procedure of the experimenter [18].

Beginning in the 1980s, the sub-field of quantum optics which theoretically and experimentally explores just such single-atom experiments has gained incredible momentum [3].

In general, quantum optics describes the study of the interactions between quantum (or coherent) light and quantum matter – e.g. using a single photon to promote an atomic energy level from its ground to excited state. The technologies involved for trapping and manipulating single atoms required many advanced technologies: ion traps, frequency combs and precise laser frequencies – many invented or perfected at NIST-Boulder [19].

It is from these experiments that much of this thesis is inspired. Can we reproduce atomic experiments with artificial atoms? Can we explore parameter space unavailable to those experiments and learn new things?

1.1.2 Josephson devices

The answer to both of those questions has been proven to be 'yes' with the development of quantum devices based on superconducting circuits. Again, here, NIST-Boulder has played a leading role in producing such novel devices.

The Josephson junction (details Chapter 2) is a key ingredient for three devices in the Quantum Devices Group at NIST (of which I am a member). First, arrays of thousands of junctions in series, each with a well defined superconducting gap, can be used as a voltage standard [20]. Or, one can use SQUID-based devices as sensitive magnetometers. These magnetometers are very useful for reading out 'TES' (Transition Edge Sensor) based devices – i.e. converting a change in resistance and current to measurable change in magnetic flux. The third device which was featured very strongly in our group (until very recently) was the phase qubit or RF SQUID [21].

The phase qubit is a device formed by a Josephson Junction (JJ) in parallel with a capacitance and inductance. The JJ is equivalent to a non-linear, tunable, inductor (see next chapter) thus the RF SQUID is then equivalent to a non-linear, tunable LC resonator circuit and can be operated as a qubit [22]. This was the birth of the quantum computing effort at NIST and was one of the early technologies which seemed to be a promising candidate for quantum computing with superconducting circuits.

Unfortunately, the large capacitance that is needed for the phase qubit to be operated at a reasonable frequency is a big source of loss through lossy materials, especially lossy dielectrics [23]. A big effort of my early research at NIST was exploring different materials and fabrication techniques to improve this loss mechanism. In some sense, this has been

fruitful in that it allowed our group to develop ‘vacuum-gap’ capacitors which then spawned a whole new, rich, branch of research – electro-mechanics [24, 25].

The phase qubit is now out-of-favor for use in superconducting quantum computing due to the invention of the ‘transmon’ qubit at Yale [26, 27]. The transmon’s coherence times are about an order of magnitude higher than other flavors of superconducting qubit, and are specifically designed to easily couple to electric fields of resonant structures. This is the technology that I chose to use in my research.

Additionally, a big draw for using engineered circuits is the possibility for extending the fabrication of these to the scale of industrial fabrication used in say, the semiconductor industry in making CPU (central processing unit) chips. Since the transmon circuit is simpler than other flavors of qubit to fabricate it is favored for this reason also.

1.2 Quantum information processing

There exist multiple proposals for how a quantum computer should even be built but if one considers the basic building blocks of a classical computer we can begin to see the requirements for one possible quantum computer and some of the technical hurdles that will need to be overcome. Figure 1.1 shows the cartoon version of the *Von Neumann* (classical) computing architecture. The CPU takes inputs from some device, performs operations (gates) using a control unit and logic unit and spits out the result to an output device. Intermediary to that, the CPU can have interactions with a memory unit. The CPU can store, modify, or retrieve information from the memory.

An example quantum computer will also need to have input states, a processing unit, a memory, and produce output states. The input and output states are quantum in nature though. The rest of the operations can be performed by purely classical fields but with the additional complication that quantum states are very sensitive to external noise. The focus of much of the work of this dissertation is to develop and demonstrate new ways (and

devices) which incorporate the controlled manipulation of quantum states from a memory element to a processing element which could also avoid sources of noise. In essence, how does one shuttle information around (in frequency and space) between inputs, memories, and outputs without losing the quantum properties we desire?

As briefly mentioned above, as opposed to a classical computer, the quantum states, qubits, memories, and control fields for a quantum computer are not at DC frequencies, but are all based on resonances at high frequencies (in our case $\sim 10^9$ Hz). To make the best resonance possible requires careful materials design, and clever geometries to minimize loss/noise sources. This is an active area of research and many groups are pursuing variants on the multi-mode memory ideas presented here. In particular, the Schoelkopf group at Yale has recently demonstrated storage times of their memory cavities to be around 1 millisecond [12]. It is believed that further improvements of our whispering gallery mode resonator device could rival that.

We think the devices and techniques discussed within this dissertation can be a viable platform for storing and manipulating a few tens of bits of quantum information. In that sense, it is not a fully scalable quantum computing solution, but could certainly be used as a test-bed or stepping-stone for further technology.

It is also worth mentioning the stringent rules needed for the quantum operations and quantum states that are input and output to/from the quantum computer. These were first laid out by David DiVincenzo and are known as the *DiVincenzo Criteria* [28]. These are very important factors which make the quantum computing problem many factors harder than the equivalent classical computing problem. This thesis does not focus on them, but they are in the background of most every quantum computer engineer's mind. The DiVincenzo criteria state that a scalable, useful quantum computer must have: 1.) well-defined qubits; 2.) reliable (quantum) state preparation; 3.) low decoherence; 4.) accurate quantum gate operations; and 5.) strong quantum measurements.

Again, just focusing on how to best fulfill all of the criteria is an active and fruitful area of research.

1.3 Thesis overview

This thesis follows the previous historical development loosely. In **Chapter 2**, we summarize the necessary background physics needed for this work – mostly all known fully since about 1930 and the days of Schrödinger and modified as the quantum optics field grew. In fact there is not much *new* physics described here. This dissertation is mostly concerned with novel devices to exploit the engineer-ability of superconducting-based quantum optics.

To that end we describe the practical aspects of experimental design in **Chapter 3** and motivations for the experiments we chose to perform in **Chapter 4**.

Chapters 5 - 8 detail the experimental data from those experiments; and **Chapter 9** offers some conclusions and outlooks for future devices and experiments.

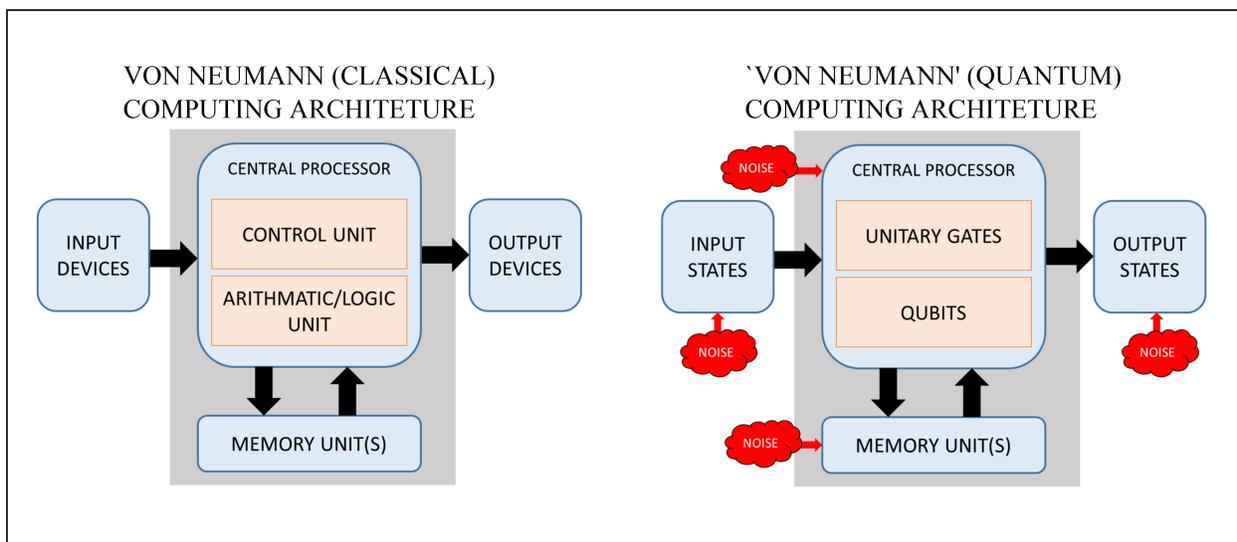


Figure 1.1: Quantum and classical computers will most likely share some of the same elements – a central processor, memories, control gates, etc – much of the challenge in porting the Von Neumann architecture from the classical to quantum regime involves careful avoidance of noise sources and determining which techniques are best for implementing transfer of information from the central processor to/from memories and inputs/outputs.

Chapter 2

Background Physics

As discussed briefly in the previous chapter, almost all of the measurements for quantum information technologies rely on some sort of resonance. This is true of the qubits themselves (non-linear resonators) and for cavity resonators (linear resonators) which are used to enhance measurements or for storage of quantum states. This chapter is meant to be an introduction to how and why we perform those measurements, so many of the topics are only briefly introduced here, but then expanded on in later chapters where specific devices are discussed.

In our field of superconducting quantum devices, we tend to idealize everything to an RLC (Resistor-Inductor-Capacitor) resonant circuit model which maps to the familiar simple harmonic oscillator. But, more realistically, we are usually more interested in less lumped element structures where each element has some capacitance *and* some inductance per unit volume (or length or area) – an example is the coaxial waveguide discussed below which has a lumped element approximation, but is really not lumped per se. More realistically, we model our devices as a frequency-dependent impedance. This impedance can sometimes be built up analytically by hand, but more often than not must be simulated with very good simulation software (details next chapter). This software is often extremely useful when trying to model the behavior of complex, coupled circuits.

Additionally, while impedance is a useful tool for imagining what the energy of a circuit is doing, we are almost always *measuring* incident, reflected, or transmitted microwaves which

probe our structures. In other words, in practice, we use microwave radiation to learn about our devices by propagating waves through a 50 Ohm environment (coax) and record the magnitude and phase of the waves which *bounce off* of our devices. These techniques are described in many microwave engineering textbooks by microwave network analysis theory [29]. We summarize the salient features of microwave circuit analysis, below, then expand on that to build up the necessary machinery for calculating more quantum systems. Then, we introduce the Josephson-junction based devices used in the experiments in later chapters and parametric interactions.

2.1 Microwave engineering

Again, all of our microwave-frequency measurements can be thought of as sending a stimulus wave to our device and recording the response of the reflected or transmitted wave as it ‘bounces off’ of our device. This is a mature science and commercial tools are available which record and calibrate the stimulus and response tones very well. Microwave engineers have termed this technique *network analysis*. We cover the basics in this section and highlight a few cases of interest to us.

The simple version is that for any microwave network of N ports we can imagine a probe wave incident and reflected off of some arbitrary impedance (described by an N -by- N impedance matrix). At each port we have an incident (‘+’) and reflected (‘-’) voltage and current. Together we can build up a N -by-1 column vector of voltages, $[V]$ and similarly for currents, $[I]$. Now, $[V]$ and $[I]$ are related by the impedance matrix $[Z]$ thusly

$$V_n = V_n^+ + V_n^- \quad (2.1)$$

$$I_n = I_n^+ + I_n^- \quad (2.2)$$

$$[V] = [Z][I]. \quad (2.3)$$

Defining currents and voltages is difficult though when describing complex networks where distributed circuit elements often have a bit of each. If we restrict ourselves to just looking at

incident and reflected voltages at each node (V_n^+ and V_n^-) we can define a better description of what we actually measure by using the scattering matrix

$$[V^-] = [S][V^+], \quad (2.4)$$

whose meaning becomes clear when looking at a specific element of the matrix. For example if we are interested in the response of port 1 due to *only* applying a drive at port 3 we need to look at $S_{13} = V_1^-/V_3^+$. Measurements of S_{ii} are termed ‘reflection’ measurements and terms like S_{ij} ($i \neq j$) are termed ‘transmission’ from port j to i .

The probe tones to/from our device propagate along coaxial transmission lines. It is pedagogical to describe this propagation and work our way through an imaginary measurement of a ‘device’ (load impedance) connected at the end.

2.1.1 Transmission-lines

The coaxial transmission line of Figure 2.1.1a consists of a center conductor on which we would like to send our probe voltages/waves and an outer conductor held at ground. The two conductors are separated by a dielectric. 2.1.1b imagines an idealized unit cell of the coax as a set of circuit elements – R (resistance of the conductors), L (self-inductance), G (conductance, accounting for dielectric loss), and C (capacitance between the two conductors) – *per unit length*. We can then write Kirchohoff’s voltage and current laws for this unit cell, using current (i) and voltage (v) as defined in the figure, as:

$$v(z, t) - R\Delta z i(z, t) - L\Delta z \frac{\partial i(z, t)}{\partial t} - v(z + \Delta z, t) = 0 \quad (2.5)$$

$$i(z, t) - G\Delta z v(z + \Delta z, t) - C\Delta z \frac{\partial v(z + \Delta z, t)}{\partial t} - i(z + \Delta z, t) = 0. \quad (2.6)$$

Then, after dividing by Δz , taking the limit as $\Delta z \rightarrow 0$, and assuming that $v(z, t)$ and $i(z, t)$ are sinusoidal propagating waves of frequency ω , we get the differential equations:

$$\frac{dV(z)}{dz} = -(R + i\omega L)I(z) \quad (2.7)$$

$$\frac{dI(z)}{dz} = -(G + i\omega C)V(z) \quad (2.8)$$

which can be solved simultaneously to find the propagating waves, $V(z) = V_o^+ e^{-\gamma z} + V_o^- e^{+\gamma z}$ and $I(z) = I_o^+ e^{-\gamma z} + I_o^- e^{+\gamma z}$. These are the signals which we use to probe and measure the transmitted and reflected signals from our device under test. The propagation constant is defined as $\gamma = \sqrt{(R + i\omega L)(G + i\omega C)}$. Often, the loss from the R and G factors is very small and we can idealize a bit further to get $\gamma = i\omega\sqrt{LC}$ (with $R \rightarrow 0$ and $G \rightarrow 0$). We also define the characteristic impedance $Z_o \equiv \sqrt{L/C}$. Z_o is set by the geometry — center conductor radius, dielectric constant thickness, and outer conductor radius — to 50 Ohms.

This 50 Ohm, idealized, transmission line will always be connected to source electronics and components which are designed carefully to be as close to 50 Ohms as possible. Thus once the source voltage travels down the transmission line, it will propagate along with minimal disturbance and no reflections backward. The devices we are interested in almost always have impedances which differ from 50 Ohm by design — we call this the load impedance. This mis-match between 50 Ohms and the load impedance will cause some of the propagating wave to be reflected off of the load and be returned to the measurement device — shown in Figure 2.1.1c. This is parameterized by a reflection coefficient

$$\Gamma = \frac{V_o^-}{V_o^+} = \frac{Z_L - Z_o}{Z_L + Z_o}, \quad (2.9)$$

which is precisely what was defined as the scattering parameter for a 1-port network (S_{11}), above. From this equation, we can see that if a 50 Ohm load is connected to the 50 Ohm transmission line, no reflection occurs. It is only when the load impedance is *different* from 50 Ohms when reflections occur. In this thesis, our load impedance will always be some kind of ‘cavity’ resonance which can be thought of as a load impedance which changes rapidly with frequency and is almost always designed to be a much higher impedance. Again, we will continue an imagined measurement below of a cavity resonance.

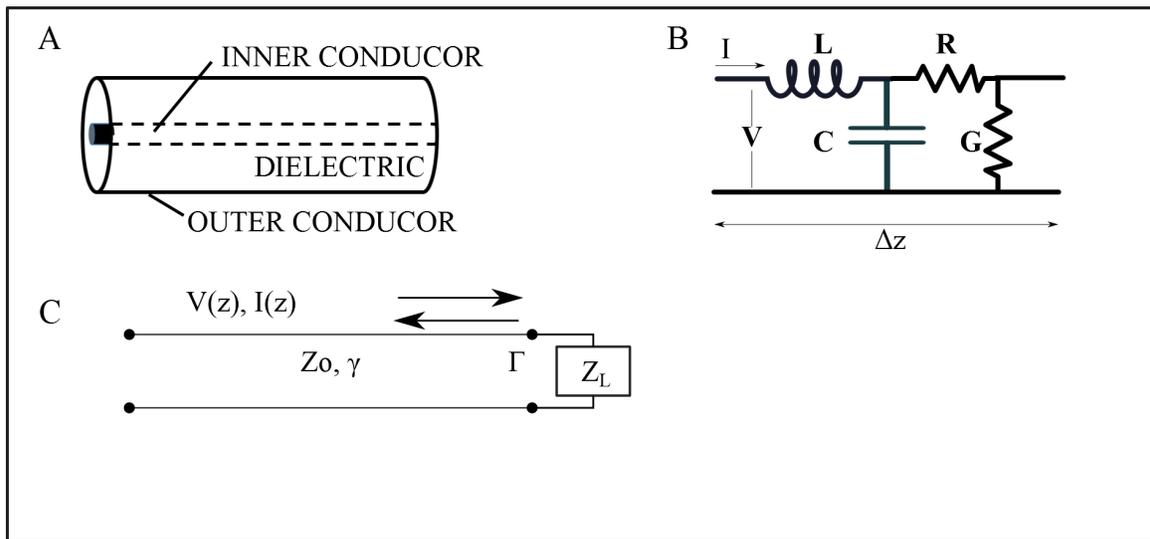


Figure 2.1: Three equivalent ways to picture a coaxial transmission line: a.) a sketch of a physical section of coaxial microwave transmission line labeling the two metallic conductors separated by a dielectric; b.) coaxial transmission line represented as ideal circuit elements with per unit length (Δz) inductance (L), capacitance (C), resistance (R), and conductance (G); and c.) propagating signals ‘bouncing off’ of a mis-matched load impedance. The signal propagates without loss until the load impedance causes reflection or transmission/absorption

2.1.2 Lumped-element circuit analysis

Before introducing cavities with larger spatial extent, it is again pedagogical to begin with a short review of lumped-element circuit models to gain intuition first. The term ‘lumped element’ refers to the fact that we model the circuit with only ideal circuit elements – in our case: L’s, C’s, and R’s (inductors, capacitors, and resistors). This is a useful approximation and applies to circuit elements that are less than one wavelength. This is often not always true with microwave circuits where the wavelength is typically 10^{-1} to 10^{-3} meters in free space (~ 3 GHz to 300 GHz). For now, we will use lumped-element circuits to introduce ideas and measurements which we will then apply later.

The familiar parallel RLC circuit of Figure 2.1.2a has an input impedance (Z_{in} , the impedance as if we connected it directly to an RF source) of

$$Z_{in} = \left(\frac{1}{R} + \frac{1}{i\omega L} + i\omega C \right)^{-1} \quad (2.10)$$

where R, L, and C are the parallel resistance, inductance, and capacitance of the circuit. In the large-R (low-loss) limit, and near resonance, equation 2.10 reduces to

$$Z_{in} = \frac{R}{1 + 2i(\omega - \omega_o)RC}. \quad (2.11)$$

The input impedance tells us almost everything we would like to know about how this circuit acts – its resonance frequency, and how it reacts below and above that frequency. The resonance frequency is $\omega_o = 1/\sqrt{LC}$ and below $\omega = \omega_o$ the resonator’s input impedance looks inductive (i.e. it has a positive imaginary component of Z_{in}), while above ω_o it looks capacitive (a negative imaginary component of Z_{in}). Precisely at resonance the input impedance is purely real $Z_{in} = R$. The width of the frequency vs impedance curve tells us the Q of the resonance; but there are subtleties in the Q-factor calculation.

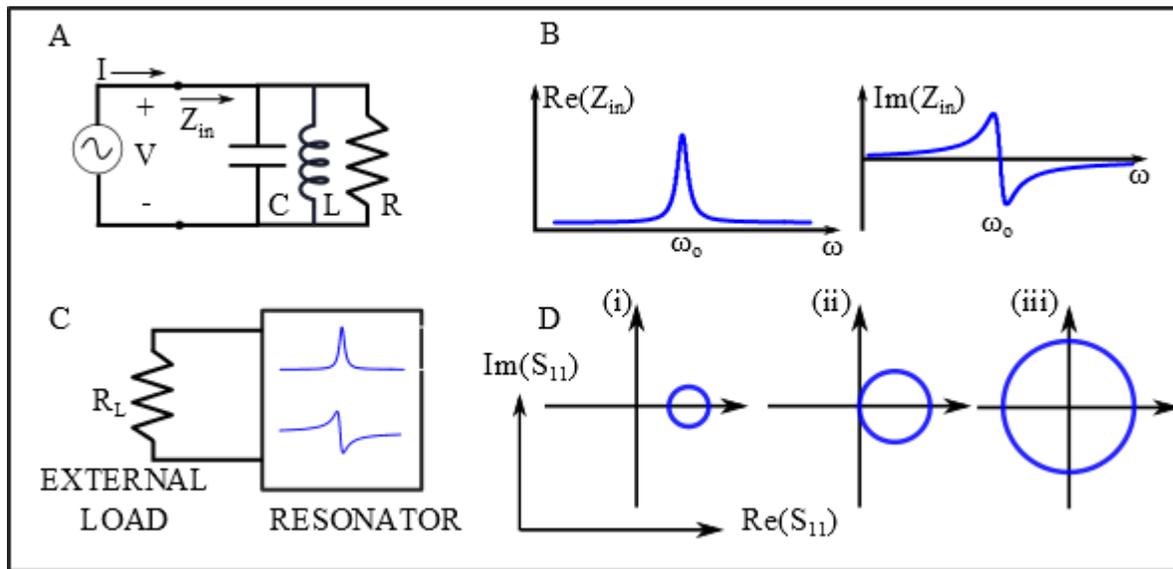


Figure 2.2: Resonator measurements: a.) a lumped-element model of a parallel RLC resonator; b.) the real and imaginary parts of the input impedance versus frequency of the circuit resonator from (a); c.) a more realistic model of a resonator connected to the outside world (modeled as a load resistor); and d.) real and imaginary parts of the scattering parameters off of an (i) under-coupled, (ii) critically-coupled, and (iii) over-coupled resonator.

More precisely Q is defined as

$$Q = \omega_o \frac{(\text{average energy stored})}{(\text{energy loss/second})} \quad (2.12)$$

$$= \omega_o \frac{W_{mag} + W_{elec}}{P_{loss}} \quad (2.13)$$

$$= \frac{R}{\omega_o L} = \omega_o RC \quad (2.14)$$

where $W_{mag} = \frac{1}{4}|V|^2 \frac{1}{\omega^2 L}$ and $W_{elec} = \frac{1}{4}|V|^2 C$ are the time-averaged magnetic energy stored in the inductor and the electric energy stored in the capacitor, and $P_{loss} = \frac{1}{2} \frac{|V|^2}{R}$ is the power dissipated by the resistor (V is the voltage as defined in Figure 2.1.2a). The last line follows from the fact that on resonance, the magnetic and electric energies are equal.

We shorten the equation for Q further by associating the ‘RC’ constant as a relaxation time and thus Q is simply written as $Q = \omega_o \tau$. τ is precisely the time constant one would measure (in the time-domain) for the cavity to ring-up (or ring-down) and determines how fast or slow the cavity will respond to stimulus.

The Q as defined in equation 2.14 is more specifically known as the *internal* quality factor – the quality factor of an ideal circuit completely isolated from the outside world. We will see, below, that we can separate out ‘de-Q-ing’ effects due to coupling to the 50 Ohm environment from the internal Q simply by modeling the external world as an external load resistor to our resonant circuit – see Figure 2.1.2c. This external load will add in parallel with the R in the RLC resonator such that the effective overall resistance is $RR_L/(R + R_L)$ which results in an overall Q less than Q_{int} . We call the external load as an ‘external Q’ or Q_{ext} . This external loading can be mediated by using a coupling capacitor or mutual inductance to isolate the ‘internal’ resonance by transforming the external load such that the resonator ‘sees’ a higher Q background [30].

We illustrate the use of the scattering parameter measurements with a simulated measurement of the RLC resonator above, now probed via a reflection measurement through say, a coupling capacitor which sets the external Q. The reflection scattering parameter (from

port 1 to port 1) can be calculated to be an inverted Lorentzian lineshape (cf: equation 2.11) with a linewidth κ :

$$S_{11} = \Gamma = 1 - 2\eta \frac{\kappa}{\kappa + 2i(\omega - \omega_o)} \quad (2.15)$$

where we have defined, $\eta = \kappa_{ext}/\kappa = Q/Q_{ext}$ – here Q is ‘as measured’ for the total resonance (i.e. what the experimentalist sees), including internal and external Q s.

There are three regimes that η is allowed to have: if $1/2 < \eta \leq 1$, the resonance is ‘over-coupled’; if $\eta = 1/2$ the resonance is ‘critically’ coupled; and if $0 \leq \eta < 1/2$, the resonance is ‘under-coupled’. We illustrate all three examples below. Figure 2.1.2d shows the results of plotting the real part of S_{11} versus the imaginary part of S_{11} (a frequency sweep from low to high frequency would move clockwise around the circle), for three values of an external Q and the same value for Q_{int} . This is the data which a vector network analyzer would give you in the lab. We can see for the three values of the external coupling the ‘resonance circle’ sweeps out more or less of the $\text{Re}(S_{11})$ vs $\text{Im}(S_{11})$ plane. If the coupling is small (undercoupled, figure 2.1.2d(i)) we see that the resonance circle does not include the origin. In essence we are probing our circuit very weakly and most of the power (voltage) returns unaffected. We can calculate the Q -factor of this resonance nonetheless, and we see that we get very close to the result, above, for the bare internal Q – i.e. $Q \rightarrow Q_{int}$.

If we measure the Q of the resonance where we are strongly over-coupled Figure 2.1.2 d(iii), the circle sweeps out an area with radius $\simeq 1$. In this case if we calculate the Q -factor we see that we have de- Q 'ed our resonance by coupling it too strongly to the drive port. In this manner, it is hard to separate the internal and external Q s.

In the intermediate case (Figure 2.1.2d(ii)), we are neither overcoupled nor undercoupled to our probe and we can separate the internal and external Q s via

$$Q_{total} = \left(\frac{1}{Q_{int}} + \frac{1}{Q_{ext}} \right)^{-1} \quad (2.16)$$

There are other, important, aspects about fitting and extracting circuit parameters accurately and correcting data to account for experimental artifacts (e.g. [31, 32, 33]), but

the main idea is that in order for a high-Q resonator to store energy for long times, it must be well-isolated from the 50 Ohm environment. This inherently limits how fast energy can be stored or retrieved. In quantum computing terms, this limits how fast one can read or write quantum information to/from the resonant cavities.

After that brief review of *how* we measure resonating devices, we will now move on to *why* we are interested in these devices in terms of usefulness for quantum mechanical experiments.

2.2 Cavity Quantum Electrodynamics

The general theme of the experiments described in this dissertation falls under the broad umbrella of *quantum optics* or more specifically *cavity quantum electrodynamics* (cavity QED). We will unpack the physics in turn: first describing various types of ‘cavities’, then the ‘quantum’ part (atoms and artificial atoms), and finally putting it all together by describing the interactions between the two. As sketched out in Figure 2.3 the light fields inside the cavities are described by Maxwell’s equations, though in our case the light is at microwave frequencies and the electromagnetics are more easily described by circuit theory. Quantum mechanics is needed to describe the behavior of the atoms and interactions.

2.2.1 Low-loss cavities

As the name suggests, cavity QED requires *some* kind of cavity. The defining characteristics of any cavity are the two impedance mis-matches at each end (or boundary conditions), the resonance frequency, and the quality factor (or ‘Q’). The end boundary conditions can be thought of as two mirrors with high reflectivity which ‘bounce’ light back and forth between them. They can take the form of actual mirrors (in the Fabry-Perot case), or coupling capacitors (in the microwave resonator case). The resonance frequency is characterized by the wavelength of the standing wave of electromagnetic field that exists due to the two

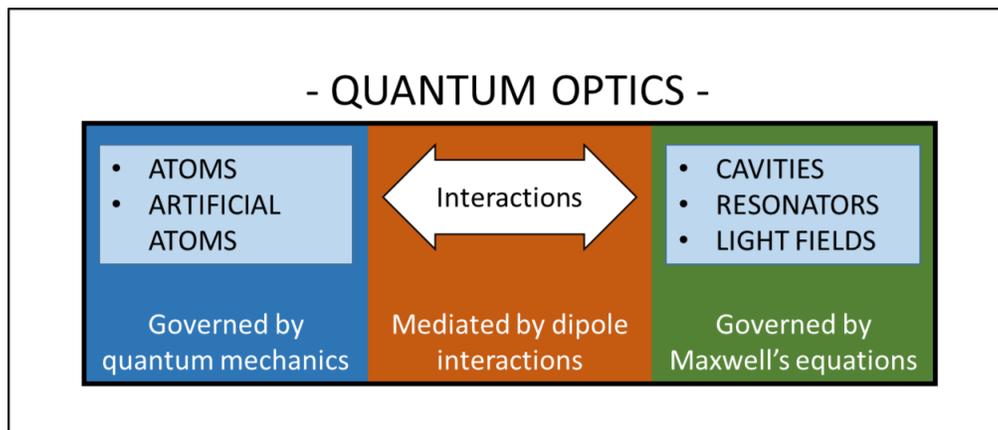


Figure 2.3: The overall theme of this thesis describes experiments in *cavity quantum electrodynamics* – essentially described by a quantum object (atom or artificial atom) interacting via dipole interactions with a cavity containing a light field

boundary conditions; and the Q tells us how long in time the standing wave will exist in the cavity without decaying. In general, microwave resonators deal with low-order resonances – typically the first or second harmonic of the standing wave. Optical cavities support much higher harmonics due to the very short wavelengths of laser light compared to microwave light.

In general, for optics and for circuits, producing cavities with the very highest Q s necessitates a low-loss, superconducting cavity resonator [34]. As we saw in the lumped element case above, on resonance the internal Q is set by the resistance of the resonator. This is true in the three-dimensional cavity resonator as well, but we are now concerned with the surface losses (i.e. surface resistance) of the material making up the resonant cavity. Luckily, below a certain critical temperature, T_C , certain superconducting materials completely lose any resistance – i.e. loss — at DC frequencies. We will be focusing on aluminum cavities which have a critical temperature $T_C = 1.2$ K. Even at frequencies above DC, the resistance of superconducting aluminum is much lower than even the best normal conductor. This finite, but small, loss is manifest through the fact that above $\omega = 0$ superconducting materials have a surface resistance approximated by:

$$R_S \approx A \frac{\omega^{1.7-2}}{T} e^{-\frac{\Delta(0)}{k_B T}} + R_0 \quad (2.17)$$

where A is a materials constant, $2\Delta(0)$ is the superconductor's energy gap at $T=0$, k_B is Boltzmann's constant, and R_0 is termed the residual resistance [34]. At temperatures low compared to T_C and frequencies low compared to $2\Delta(0)/\hbar$ — always the case in this work — the exponential is vanishingly small and only the R_0 term is left. This can be reduced very low by careful choice of materials and surface preparation to $R_0 \sim 10^{-9}$ or lower resulting in Q -factors as high as 10^{11} [35, 36]. This is not the full picture though. The structure of the mode fields has a strong impact on how much of the standing waves can induce lossy currents in the walls of the cavity.

At finite frequency, the finite fields inside the cavity can penetrate the superconductor

to a depth called the ‘skin depth’ — typically, a few tens of nm at GHz frequencies — and is denoted $\delta_{sk} = R_S/\omega\mu$. The fields can induce surface currents in this lossy layer.

Indeed if there is any surface current flowing in the walls of the cavity we can apply Joule’s Law to find the power loss in the conducting walls due to surface resistance:

$$P_c = \frac{1}{2\sigma} \int_{Surf} \int_0^{\delta_{sk}} \frac{|J_S|^2}{\delta_{sk}} dz ds \quad (2.18)$$

$$= \frac{R_S}{2} \int_{Surf} |J_S|^2 ds \quad (2.19)$$

$$= \frac{R_S}{2} \int_{Surf} |H_t|^2 ds \quad (2.20)$$

where $J_S = \hat{n} \times H_t$ is the surface current due to the tangential magnetic field of the cavity mode. This shows that the field distribution inside the cavity plays an important role in determining power lost through surface resistance. Cavities with large spatial extent, namely three-dimensional volumes where the field is mostly distributed in the lossless volume, will yield the highest Qs.

2.2.1.1 Three-dimensional cavities

Recalling equation 2.11, we can re-write the formula for the Q of a ‘3D-cavity’ resonator as

$$Q = \omega \frac{(\text{average energy stored})}{(\text{energy loss/second})} \quad (2.21)$$

$$= \frac{1}{\delta_{sk}} \frac{\int_{Vol} (\text{Stored field in (lossless) vaccum})}{\int_{Surf} (\text{Field in lossy surface resistance})} \quad (2.22)$$

$$= \frac{\mu_o\omega}{R_S} \frac{\int_{Vol} |E|^2}{\int_{Surf} |H_t|^2} \quad (2.23)$$

which suggests that we can look at various geometries with different volume-to-surface area ratios as a means to increasing our Q value in addition to various materials and surface preparations.

RECTANGULAR CAVITIES

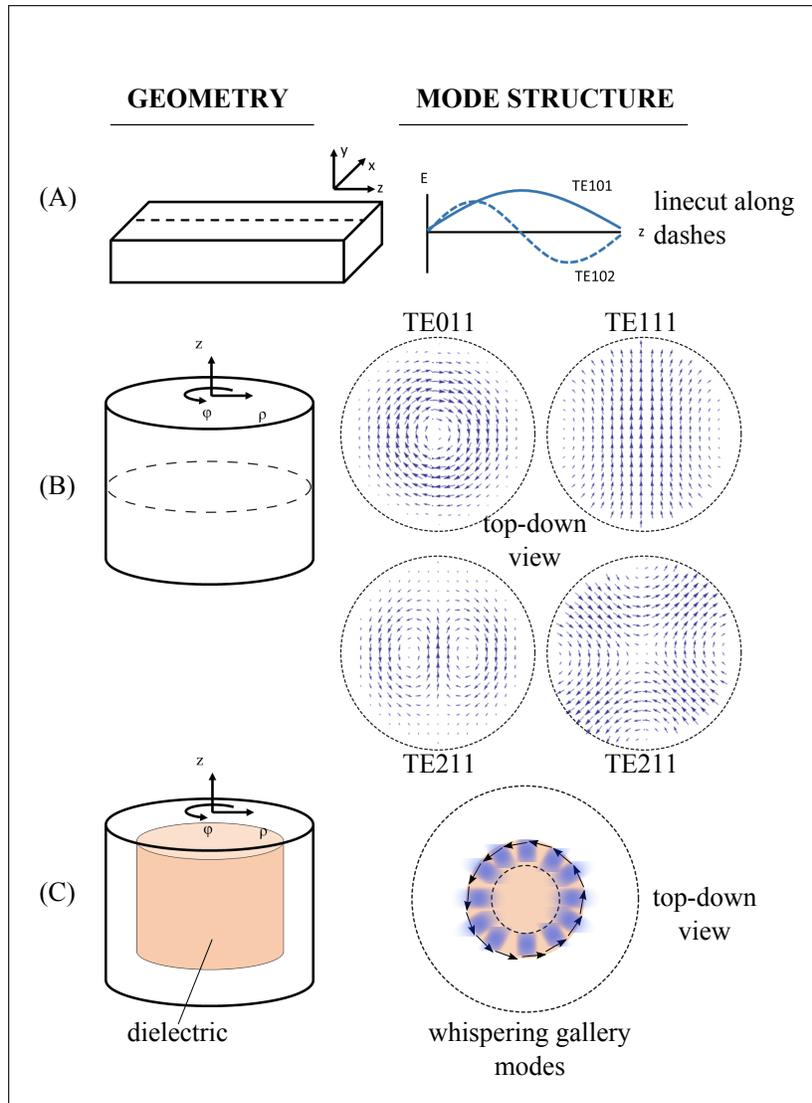


Figure 2.4: Qualitative comparison of mode structure of the electric field eigenmodes for: a.) rectangular cavities – the linecut plots the field along the dashed line; b.) cylindrical cavities which support more complex modes; and c.) whispering gallery modes in a dielectric cylinder – most of field is confined at the outer edge (outside the dashed line in the top-view). Since our qubit couples like a dipole to the electric field of a cavity, we would like a cavity with easily accessible fields *and* a good volume to surface area ratio which maximizes the geometric contribution to the Q-factor.

The first geometry we consider is the aluminum rectangular cavity pictured in Figure 2.4a, this has become ubiquitous in the superconducting qubit community after dramatic increases in coherence times were seen by placing transmon qubits inside these cavities [27]. This type of cavity is featured in experiments described in Chapter 5.

A rectangular cavity with dimensions $a \times b \times d$ supports transverse electric (TE) standing waves with frequencies

$$f_{mnl} = \frac{1}{2\pi\sqrt{\mu_r\epsilon_r}} \sqrt{\left(\frac{m\pi}{a}\right)^2 + \left(\frac{n\pi}{b}\right)^2 + \left(\frac{l\pi}{d}\right)^2} \quad (2.24)$$

where m , n , and l take on integer values, $= 0, 1, 2, \dots$, and μ_r/ϵ_r are the relative permeability and permittivity of any material filling the volume. In general, our cavities will be mostly filled with vacuum so it is a good approximation to use the vacuum permittivity, but a very thin layer of oxidized aluminum dielectric metal could exist on the surfaces; and any device we insert into the cavity will be lithographically patterned on a dielectric substrate. The rectangular cavities we use in our work have $b < a < d$ and thus the first few modes are the TE101, TE102, and TE103 modes.

The relevant fields of these cavity modes can be calculated to be:

$$E_y = -iE_o \frac{k_x}{k_c} \sin\left(\frac{\pi x}{a}\right) \sin\left(\frac{l\pi z}{d}\right) \quad (2.25)$$

$$H_x = -H_o \frac{k_x k_z}{k k_c} \sin\left(\frac{\pi x}{a}\right) \cos\left(\frac{l\pi z}{d}\right) \quad (2.26)$$

$$H_z = H_o \frac{k_c}{k} \omega \sqrt{\mu_r \epsilon_r} \eta a \cos\left(\frac{\pi x}{a}\right) \sin\left(\frac{l\pi z}{d}\right) \quad (2.27)$$

$$\text{with } k_x = m\pi/a, \quad k_y = n\pi/b, \text{ and } k_z = l\pi/d \quad (2.28)$$

and are plotted (via linecuts) in Figure 2.4a. $k_c = \sqrt{k_x^2 + k_y^2}$ refers to a cutoff wavenumber in which limits propagation below a cutoff frequency.

On resonance, just like in the RLC resonator case the magnetic and electric energies are equal so we can just calculate the stored electric energy and multiply by two. Thus, the

numerator in equation 2.23 is

$$2 \times \int_{Vol} |E|^2 = \frac{\epsilon}{2} \int_{Vol} E_y E_y^* dv = \frac{\epsilon abd}{8} E_o^2. \quad (2.29)$$

The denominator is more complicated since we must integrate over each interior face of the cavity but gives:

$$\begin{aligned} R_S \int_{Surf} \frac{1}{2} |H_t|^2 ds &= \frac{R_S}{2} \left\{ \int_{y=0}^b \int_{x=0}^a 2 |H_x(z=0)|^2 dx dy \right. \\ &+ \int_{z=0}^d \int_{y=0}^b 2 |H_z(x=0)|^2 dy dz \\ &+ \left. \int_{z=0}^d \int_{x=0}^a 2 (|H_x(y=0)|^2 + |H_z(y=0)|^2) dx dz \right\} \end{aligned} \quad (2.30)$$

$$= R_S \frac{E_o^2 \lambda^2}{8 \eta^2} \left(\frac{l^2 ab}{d^2} + \frac{bd}{a^2} + \frac{l^2 a}{2d} + \frac{d}{2a} \right), \quad (2.31)$$

with $\lambda = 2\pi / \sqrt{(\pi/a)^2 + (l\pi/d)^2}$ and $\eta = 377\Omega$, the impedance of free space.

Finally, we calculate the Q due to finite surface resistance for rectangular cavities to be:

$$Q_{rect} = \left(\frac{1}{R_S} \right) \frac{(2\pi ad)^3 b \eta}{2\pi^2 (2l^2 a^3 b + 2bd^3 + l^2 a^3 d + ad^3)}. \quad (2.32)$$

Which we note is a factor that only depends on the surface loss multiplied by a purely geometric factor. The geometric factor is what determines the resonant frequency and mode structure, but also determines how much current is induced to flow in the lossy skin depth layer.

If we limit ourselves to the TE101 and TE102 modes that correspond to resonant frequencies in the 5 to 10 GHz range of our measurement band the geometrical factor is limited. For a typical rectangular cavity of dimensions $\sim 6.35\text{mm} \times 19\text{mm} \times 30\text{-}50\text{mm}$ we get:

$$Q_{rect} \approx \frac{1}{R_S} (3 \times 10^{-3}) \text{ to } \frac{1}{R_S} (6 \times 10^{-3}). \quad (2.33)$$

We compare different geometries for reasonable frequencies at the end of the chapter in table 2.1.

The biggest draw of this geometry is its simplicity. The TE modes are easily excited by inserting a coaxial cable through an aperture in the side of the cavity volume. The placement of the aperture and depth of the coaxial drive determines the external Q. The coupling is strongest for the fundamental mode (TE101) with higher order modes being weaker.

Another advantage of this type of cavity is that it is relatively compact – to get the same frequencies cylindrical cavities are much larger. These volumes are also compared in table 2.1.

CYLINDRICAL CAVITIES

Cylindrical cavities benefit from the fact that for some modes there is very little field near the surface of the cavity to induce surface currents in the lossy skin depth layer. This is not true for all of the modes though, some induce strong surface currents and therefore have lower Qs.

The electric and magnetic fields for the eigenmodes of a cylindrical geometry are precisely the cylindrical harmonic solutions to Maxwell's equations. Before writing the formulae for the fields we can make a few simplifications to make the calculations easier. First, we will restrict ourselves to right, square, cylinders – i.e. cylinders whose diameter equals the height (square in cross section) – of height, d (radius $d/2$). This will maximize the volume to surface area ratio in equation 2.23. Second, we will group all of the constants describing the mode impedance and amplitude into one constant out front, A_o , which allows us to write the equations for the electric field more simply as:

$$E_\rho = 2A_o \frac{\omega n}{p'_{nm}\rho} J_n\left(\frac{2p'_{nm}\rho}{d}\right) \sin(n\phi) \sin\left(\frac{l\pi z}{d}\right) \quad (2.34)$$

$$E_\phi = 2A_o \frac{\omega n}{p'_{nm}\rho} J'_n\left(\frac{2p'_{nm}\rho}{d}\right) \cos(n\phi) \sin\left(\frac{l\pi z}{d}\right) \quad (2.35)$$

$$E_z = 0. \quad (2.36)$$

Here n , m , and l denote the number of nodes in the azimuthal (ϕ), radial (ρ), and z directions

and take on integer values. The J_n and J'_n functions are Bessel functions and derivatives of Bessel functions with roots p'_{nm} . A few representative solutions of these electric fields are plotted in Figure 2.4b. Just like in the case of rectangular cavities, the field distribution of the modes will influence how much current is induced to flow on the surfaces, and in turn how much dissipation is induced.

With the more complex mode structure though, an additional difficulty is that it can be hard to couple to the desired mode without unwanted coupling to other modes. This is especially true for modes which are contained mostly in the volume (as opposed to near the edges). In other words, the simple coaxial drive line technique for coupling to rectangular cavity modes will excite the modes nearest the cavity walls easily, but exciting the high-Q ‘internal’ modes is more difficult. Looking ahead to experimental details, it remains a challenge as to how one inserts a chip into the volume without negatively affecting the Q of that mode. A recent paper from the Yale group has come up with a clever solution which modifies the cylindrical geometry and enables them to do just that [12].

As with the rectangular geometry we can write down a formula for the Q:

$$Q_{cyl} = \frac{1}{R_s} \frac{\omega [1 - (\frac{n}{p'_{nm}})^2] [p'^2_{nm} + (\frac{l\pi}{2})^{3/2}]}{2\pi c [p'^2_{nm} + \frac{d}{2l} (\frac{2nl\pi}{d^2 p'^2_{nm}})^2]} \quad (2.37)$$

$$\approx \frac{1}{R_s} (5 \times 10^{-3}) \text{ to } \frac{1}{R_s} (2 \times 10^{-2}). \quad (2.38)$$

Again, we have written the Q as a geometric factor multiplied by the factor describing the surface resistance, and in the second line have estimated the geometric factor with realistic frequencies and dimensions. We note that these modes often have a factor of 10 or more better Q than the rectangular shape just based on geometry.

As discussed above, this results in larger volumes overall, but many more modes in a usable frequency range. The next logical step to maximize the volume-to-surface area ratio would be investigate spherical cavities. But due to the already complex mode structure

and difficulties with including devices inside the cavities we went in search of other cavity geometries and materials.

DIELECTRIC, WHISPERING GALLERY MODE, CAVITIES

While we mentioned the *dielectric loss* contributions to the Q factor above, it did not play into our estimations of Q_{int} since most of our cavity was filled with vacuum. There exist *dielectric* resonators though which are fully made of very low-loss dielectric materials. These dielectric resonator's Q are set by this factor rather than the loss due to surface resistance. Chief among these are sapphire Whispering Gallery Mode (WGM) resonators. The world of *optical* quantum optics has used small spherical or toroidal WGMs as resonators to interact with atoms for many years [37]. In this thesis, we will again follow their lead and use an artificial atom to interact with the modes of a microwave-frequency WGM. Those experiments will be detailed in Chapter 8, but here we introduce the WGM geometry and mechanism.

Our WGM resonator is a right, square, cylindrical piece of crystalline sapphire (Al_2O_3) about 3 cm in diameter and 3 cm in height. Protruding from each of the ends of the cylinder are smaller posts used to suspend the WGM resonator away from the surrounding metallic cavity which houses it. The overall shape is then a spindle rather than perfect cylinder. The posts are also used to clamp and thermalize the crystal. It is on loan from the frequency standards and metrology group at the University of Western Australia led by Prof. Michael Tobar [38]. The cylindrical geometry supports many of the TE or TM modes discussed above, but most of those are moved to lower frequency due to the slower speed of light in the dielectric material than in vacuum. To wit, WGMs would also exist in a vacuum filled resonator, but would require much larger scales impractical for our intended uses. The use of dielectric as a propagation medium forces those WGMs down into the range which we are interested.

The whispering gallery mode itself occurs when an electric (or magnetic) field enters the

sapphire and gets continuously totally internally reflected around the circumference of the dielectric/vacuum interface – see Figure 2.4c – inside the crystal. The resonance frequency is determined by how many wavelengths ‘fit’ around the circumference of the sapphire. In other words an integral number of wavelength must equal the circumference $n\lambda/\sqrt{\epsilon\mu} = \pi d$, where n is an integer, λ is the wavelength, ϵ the permittivity, and μ the permeability of the sapphire. Just as with the cylindrical eigenmodes, the modes here are labeled with three numbers. Each denotes the number of nodes in the standing wave in the azimuthal (m), radial (n), and z (p) directions. The highest Q modes have $n = p = 0$ and we will restrict ourselves to studying these.

The detailed calculation of these modes requires the volume to be split into three regions of solution [39]. The fields are *excluded* from the very inner volume (small radii) since the modes require internal reflections at the circumference it is safe to ignore this volume. Again, the spindle post is set to be smaller than this region in order to minimize clamping losses. Solutions are *allowed* in a torus-shaped volume around the circumference where almost all of the field is concentrated. Outside the sapphire, the fields exponentially decay with distance from the sapphire/vacuum surface. In the region of interest, shown by the dashed line in Figure 2.4, the azimuthal field of an example WGM is:

$$E_\phi = A_o J_m(k\rho) \cos(\beta z) \begin{cases} \cos(m\phi) \\ \sin(m\phi) \end{cases} \quad (2.39)$$

where the sine and cosine terms after the brace denote two degenerate solutions. In general, the two solutions are not perfectly frequency-degenerate due to imperfections in the crystal. The propagation constants, k and β , describe the field propagation in the z - and radial-directions.

The Q-factor calculation for a dielectric-filled resonator follows equation 2.23. If we

imagine a cavity where the fields are fully within the dielectric, the Q is:

$$Q_d = \omega_o \frac{(\text{average energy stored})}{(\text{energy loss/second})} \quad (2.40)$$

$$= \omega_o \frac{W_M + W_E}{P_d} \quad (2.41)$$

$$= \omega_o \frac{2W_E}{P_d} \quad (2.42)$$

where W_M and W_E again denote the stored energy of the electric and magnetic fields and P_d describes the power lost in the *dielectric*. Just as with the RLC resonator case, on resonance, the electric and magnetic fields are equal so we only need to know two factors. Each in turn yields:

$$W_E = \frac{\epsilon'}{4} \int_V |E|^2 dV \quad (2.43)$$

$$P_d = \frac{1}{2} \int_V J \cdot E^* dV = \frac{\omega \epsilon''}{2} \int_V |E|^2 dV. \quad (2.44)$$

The last line uses the fact that $J = \sigma \cdot E$ and the effective conductivity for dielectrics is defined as $\sigma = \omega \epsilon''$. The ratio of the real and imaginary parts of the permittivity ($\epsilon = \epsilon' - i\epsilon''$) is defined as the dielectric loss tangent: $\tan \delta = \epsilon''/\epsilon'$. Additionally, we note that the same volume integral appears in the formula for the electric energy as well as the dielectric power loss. Thus, the formula for Q due solely to dielectric loss (ignoring conductor loss) is simply:

$$Q_d = \frac{1}{\tan \delta}. \quad (2.45)$$

Essentially, the loss tangent represents how much power is lost *in a material* versus how much is stored. If we consider conductor (wall) losses as well as dielectric losses, the total Q goes as

$$Q = \left(\frac{1}{Q_c} + \frac{1}{Q_d} \right)^{-1}. \quad (2.46)$$

The loss tangent for sapphire, at 20mK has been measured to be $\sim 10^{-8}$ and thus the Q of a perfect sapphire resonator could be very high [40]. The loss tangent depends on materials properties such as defects and impurities.

Another attractive aspect of this geometry is that the electric fields which are circulating around the sapphire are accessible just outside the crystal in the exponentially decaying fields at the sapphire/vacuum interface. In other words, there is evanescent field outside the sapphire in which we could use a microfabricated antenna to couple to (see Chapter 6). This is slightly complicated by the fact that the highest-Q modes are *more* contained in the sapphire, while the lower-Q modes ‘leak out’ to the environment more (and are easier to couple to). This negates the complications of the vacuum-filled cylindrical cavity modes in that a chip can be easily inserted into the system, the whispering gallery modes can be excited via the traditional coaxial probe, and *many* modes will be available in a compact geometry (slightly smaller volume-wise than the empty cylinder).

The details of the experiments performed with these cavities is saved until later chapters, next we discuss how our artificial atomic circuits can affect such standing waves and why these types of cavities are interesting.

2.2.2 Quantization of a cavity field

The previous section described various ways of creating standing waves of an electromagnetic field inside cavities that differed in geometry, material, and size. The physics involved was fully classical, and in fact fell more into the engineering camp than the physics camp. In this section we ease our way into the quantum physics world by considering an ideal cavity — a single mode, one-dimensional, electromagnetic wave trapped between two perfectly reflecting mirrors (see Figure 2.5). This section will introduce second quantization notation for linear harmonic oscillators, as well as serve as a counterpoint to the nonlinear harmonic oscillators in later sections.

We begin by writing the electric field of a standing wave inside the cavity as

$$\vec{E} = \hat{\varepsilon}q(t)\sqrt{\frac{2\Omega^2}{\epsilon_0 V}} \sin(kz) \quad (2.47)$$

where ε denotes the polarization (in the \hat{x} -direction), $q(t)$ is the (dimensionless) amplitude

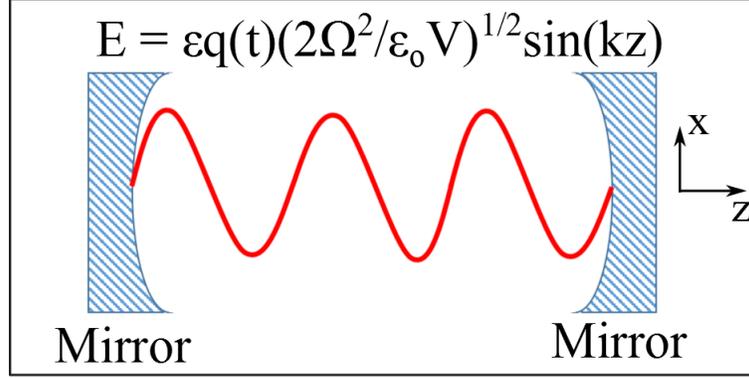


Figure 2.5: An idealized cavity which supports a standing wave of an electric field. This field can be quantized and written in terms of creation and annihilation operators.

of the field, Ω is the cavity resonance frequency ($\Omega = kc$, with c being the speed of light), and V is the mode volume. The field must obey Maxwell's equations, in particular

$$\hat{\nabla} \times \hat{B} = \frac{1}{c^2} \frac{\partial \hat{E}}{\partial t} \quad (2.48)$$

which implies that the magnetic field, and energy density are

$$\hat{B} = \frac{\hat{y}}{c^2 k} \dot{q}(t) \sqrt{\frac{2\Omega^2}{\epsilon_0 V}} \cos(kz) \quad (2.49)$$

$$U = \frac{1}{2} (\epsilon_0 |E|^2 + \frac{1}{\mu_0} |B|^2). \quad (2.50)$$

To find the total energy, we just take the integral over the mode volume (in this case just one-dimensional) and arrive at the Hamiltonian:

$$H = \int U dV \quad (2.51)$$

$$= \frac{1}{2} \int_V dV \epsilon_0 q^2 \left(\frac{2\Omega^2}{\epsilon_0 V} \right) \sin^2(kz) + \frac{1}{2} \int_V dV \frac{1}{\mu_0} \frac{\dot{q}^2}{(c^2 k)^2} \left(\frac{2\Omega^2}{\epsilon_0 V} \right) \sin^2(kz) \quad (2.52)$$

$$= \frac{1}{2} \Omega^2 q^2 + \frac{1}{2} \dot{q}^2 \quad (2.53)$$

$$= \frac{1}{2} (\Omega^2 q^2 + p^2) \quad (2.54)$$

where we have simplified using $\frac{1}{c^2} \equiv \mu_0 \epsilon_0$, and recognized this as the harmonic oscillator Hamiltonian with mass = 1 and canonically conjugate position and momentum: q and p ($= \dot{q}$).

From here we can introduce normal mode amplitudes (i.e. creation and annihilation operators if they are wearing hats):

$$\hat{a} = \frac{1}{\sqrt{2\hbar\Omega}}(\Omega\hat{q} + i\hat{p}) \quad (2.55)$$

$$\hat{a}^\dagger = \frac{1}{\sqrt{2\hbar\Omega}}(\Omega\hat{q} - i\hat{p}) \quad (2.56)$$

by defining $\hat{q} = \sqrt{\frac{\hbar}{2\Omega}}(\hat{a} + \hat{a}^\dagger)$ and $\hat{p} = -i\sqrt{\frac{\hbar}{2\Omega}}(\hat{a} - \hat{a}^\dagger)$. To finally arrive at the Hamiltonian, $\hat{H} = \hbar\Omega(\hat{a}^\dagger\hat{a} + 1/2)$, for a single mode quantum mechanical oscillator with eigenfrequency Ω and energy eigenstates $|n\rangle$. These states have the following properties (ignoring the hats):

- $a^\dagger a |n\rangle = n |n\rangle$ for $n = 1, 2, 3, \dots$;
- eigenenergies are $\hbar\Omega(n + 1/2)$;
- n describes the number of photons in the cavity ($a^\dagger a$ is the photon number operator)
- a removes a photon, a^\dagger creates a photon
- coherent states are built up from $|\alpha\rangle = e^{-|\alpha|^2/2} \sum_n \frac{\alpha^n}{\sqrt{n!}} |n\rangle$
- the electric field can now be written $E = \hat{q}\sqrt{\frac{2\Omega^2}{\epsilon_0 V}} \sin(kz) = \frac{1}{2}\epsilon_\Omega(a + a^\dagger)\sin(kz)$.

The final item, above, defined a new term, $\epsilon_\Omega = \sqrt{\hbar\Omega/\epsilon_0 V}$ which has dimensions of ‘electric field per photon’. Below, this will help determine the strength of interaction between our atom and the cavity modes. In essence it describes, with the field distribution, the field strength at each position in the cavity. This is most used for point-like atoms, but in the case of artificial atoms this field must be integrated over the spatial extent of the artificial atom if it is a significant fraction of the wavelength.

Here we have just described the fully harmonic case (equally spaced energy levels) of the quantum harmonic oscillator. Under any coherent drive, the cavity will keep absorbing photons and ‘climb the ladder’ of energy levels. There is no way to selectively drive transitions

from one rung of the ladder to another. We need to include a nonlinear element to have that ability.

It is useful to pause and compare this term for various forms of cavity QED. In table 2.1, we compare the approximate volume, typical resonance frequency, field per photon, and maximum Q for the three cavity geometries discussed above *and* a typical optical cavity. This should serve as motivation for the later design of our artificial atoms. If we are to mimic the experiments from atomic physics, we are handicapped by much lower fields in which to interact. Before discussing these designs, though, we must describe the interactions.

System	Volume	Freq., Ω	$\sim \epsilon_\Omega$	Max Q
Optical cavity	$1mm^3$	10^4 GHz	1V/m	10^9
Rect cavity	$4,000 mm^3$	4-8 GHz	5×10^{-3} V/m	10^6
Cyl cavity	$15,000 mm^3$	4-10 GHz	10^{-4} V/m	10^8
Whispering gallery	$15,000 mm^3$	7-15 GHz	10^{-4} V/m	10^{8+}

Table 2.1: Table comparing Q, etc, for cavity geometries. By spreading the field out over a larger volume though we have decreased the pre-factor which controls the coupling between an atom and cavity, ϵ_Ω . We will see, below, that artificial atoms can be lithographically defined to have much larger dipole moments and thus overcome this complication with larger cavities.

2.2.3 Interactions between an atom and a cavity mode

The second necessary ingredient in cavity quantum electrodynamics is the purely ‘quantum’ part. This role will be played by artificial atoms in the experimental sections to follow, but the necessary machinery for predicting the behavior of an atom coupled to a cavity was developed for real atoms. Indeed, by choosing to work at microwave frequencies we have hurt ourselves with smaller electric fields per photon to couple to – see table 2.1. This section will try to make the connection between atomic and artificial atomic coupling clear, and highlight the power of engineering the coupling (dipole strength) of our artificial atoms.

We begin in a similar way to the previous section by considering the same standing wave between to cavity mirrors. But now have inserted an atom with two internal energy

states into the cavity – see Figure 2.6 – which is allowed to exchange energy with the cavity field. The idealized, atomic system which we have in mind can be a single electron in a Rydberg orbit with two allowed states, or say, a magnetic spin 1/2 magnetic dipole as part of a nuclear magnetic resonance. This is clearly an approximation — atoms and artificial atoms have many energy levels, but the basic physics is best described by only looking at the ground and first excited state of our atomic system. Later we will show that this two-level system behavior can be reproduced by a non-linear circuit.

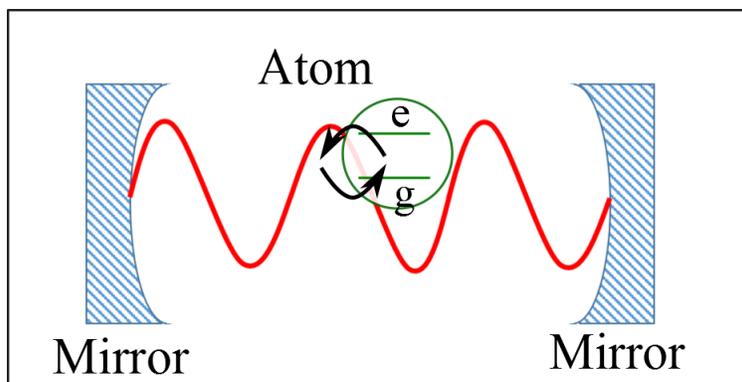


Figure 2.6: A sketch of a single-mode cavity including a simple two-level atom. The atom is allowed to swap energy (photons) with the cavity and this interaction is described by the Jaynes-Cummings Hamiltonian.

The two energy levels of interest can be labeled as $|e\rangle$ and $|g\rangle$, and the quantum state of the atom is described by $|\Psi\rangle = C_e |e\rangle + C_g |g\rangle$. The coefficients $|C_g|^2$ and $|C_e|^2$ denote the probability that the atom will be measured in either the $|g\rangle$ or $|e\rangle$ energy state; constrained by $|C_e|^2 + |C_g|^2 = 1$. The bare atomic Hamiltonian, without any interactions, is given by:

$$H_o = \frac{1}{2} \hbar \omega_a \sigma^z \quad (2.57)$$

where $\hbar \omega_a$ is the energy difference between the two energy levels, and σ^z is the 2-by-2 Pauli matrix with +1 and -1 along the diagonal and 0s on the off-diagonal.

To drive transitions between the two energy states we can supply some energy (e.g. a laser or microwave drive) whose energy is close to the energy spacing between the two levels.

We will assume that this drive is provided by the electric field standing wave inside of our cavity from the previous section. This assumes that the cavity resonance frequency is near the atom's.

In the case of Rydberg atoms and artificial atoms, the drive field interacts with the atoms through dipole interactions. The big difference, though, is that natural atoms have a fixed dipole moment; an artificial atom can be engineered to have a much much larger dipole moment.

The dipole interaction is described as $V = -\vec{E} \cdot (e\vec{r})$ where e is the electric charge, \vec{r} a position vector, and \vec{E} describes the electric field. The electric field can excite the atom from the ground state to the excited state via a σ^+ type interaction, or de-excite the atom from the excited state to the ground state via a σ^- interaction. Thus, we write the interaction energy as:

$$V = -\vec{E} \cdot (e\vec{r}) \quad (2.58)$$

$$= -E(\wp\sigma^+ + \wp\sigma^-) \quad (2.59)$$

where \wp describes the matrix element of the dipole moment along the electric field — i.e. $\wp = \langle e | e\hat{x} \cdot \vec{r} | g \rangle$. A simple atomic model with a single electron orbiting a nucleus one Bohr radius (a_o) away has $\wp = ea_o \sim 10^{-29}$ C-m ~ 10 Debye and a large Rydberg atom has $\wp \sim 40,000$ Debye. In comparison, an artificial atom can have dipole moments in excess of 10^7 Debye quite easily. To achieve this, micro-fabricated antennae can be used to extend the dipole moment of the electrical circuit over millimeters or more. This is okay since the wavelength of microwave standing waves are *also* much longer than in the optical case. We can increase the wavelength and dipole moment proportionally. Making even longer dipole antennae though allows artificial atoms to have couplings even larger than set by the limits of purely atomic physics – another example of the powerful engineer-able approach to quantum optics.

To continue, we can insert the single-mode quantized field from above ($E = 1/2\epsilon_{\Omega_c}(a +$

$a^\dagger \sin(kz)$) into 2.59 to get:

$$V = \frac{-\varepsilon_{\Omega_c} \wp}{2\hbar} \sin(kz)(a + a^\dagger)(\sigma^+ + \sigma^-) \quad (2.60)$$

$$= \hbar(a + a^\dagger)(g\sigma^+ + g^*\sigma^-). \quad (2.61)$$

If we define a (real) coupling rate as $g = -(\wp\varepsilon_{\Omega_c}/2\hbar) \sin(kz)$ we can write the Hamiltonian for the atom-cavity system as:

$$H = \frac{1}{2}\hbar\omega_a\sigma^z + \hbar\Omega_c a^\dagger a + \hbar(a + a^\dagger)(g\sigma^+ + g\sigma^-). \quad (2.62)$$

The final term can be expanded and simplified further using the rotating wave approximation (RWA) if we note that if $\omega_a + \Omega_c \gg |\omega_a - \Omega_c|$ and that $\omega_a + \Omega_c \gg g$ the far off resonant terms listed below (the second two) can be ignored/dropped:

- $a\sigma^+$: denotes absorption of a photon from the field by the atom; the atom transitions from the ground to excited state — allowed by RWA
- $a^\dagger\sigma^-$: denotes emission of a photon from the atom to the field; the atom transitions from the excited to the ground state — allowed by RWA
- $a\sigma^-$: denotes absorption of a photon from the field and transitioning from the excited to the ground state — *not* allowed by RWA
- $a^\dagger\sigma^+$: denotes emission of a photon and transitioning from ground to excited state — *not* allowed by RWA.

After applying the RWA we are left with the Jaynes-Cummings Hamiltonian:

$$H_{JC} = H_o + V = \frac{\hbar\omega_a}{2}\sigma^z + \hbar\Omega_c(a^\dagger a + \frac{1}{2}) + \hbar g(a^\dagger\sigma^- + a\sigma^+). \quad (2.63)$$

This Hamiltonian is ubiquitous in the cavity QED literature and describes the resonant (or near resonant) swapping of photons back and forth between the cavity mode and atom [41].

The result of this structure is that the resulting mathematics only needs to diagonalize, individual, 2-by-2 Hamiltonians which is always possible to do. This problem is the classic ‘Rabi’ problem where, near resonance ($\delta = 0$), the atom will periodically ‘flop’ an excitation (photon) back and forth with the cavity field at the Rabi oscillation rate of $\Omega_R = \sqrt{\delta^2 + 4g^2(n)}$. If this rate is faster than the inverse linewidth (κ^{-1}) of the cavity, this time-domain behavior results in the cavity frequency being ‘split’ by this interaction – i.e. if one tunes the atom through the cavity resonance the initial resonance will have an avoided crossing behavior such that the splitting size is $2g\sqrt{n}$ as shown in figure 2.7. Specifically, the energy levels of the system are now

$$E_{g,n} = \hbar\left(n - \frac{1}{2}\right)\omega_c - \frac{1}{2}\hbar\Omega_R \quad (2.70)$$

$$E_{e,n} = \hbar\left(n - \frac{1}{2}\right)\omega_c + \frac{1}{2}\hbar\Omega_R \quad (2.71)$$

$$|E_{e,n} - E_{g,n}|/\hbar = 2g\sqrt{n} \text{ for } \delta = 0. \quad (2.72)$$

This is clearly an idealization since no decay or dissipation mechanism has been introduced yet. We have hinted at this mechanism previously in that we spent a good deal of text describing how to build high Q-factor cavities, but we will delay the introduction of decay mechanisms until later experimental chapters and focus on the basic physics of artificial atoms now.

The second thing to note about the Jaynes-Cummings model is that there is not a way to turn the interactions ‘off’. If the atom and cavity are on resonance with one another (or nearly so) the Rabi swapping will happen continuously until moved off-resonant. As the number of atoms or cavity modes increases (as it will with larger systems) this can lead to unwanted interactions. For example, if by moving off-resonant with one mode the atom becomes on-resonance with another mode the atom will swap energy with that mode rather than undergo free evolution.

Later, we will use the fact that our qubit design is frequency-tunable to sweep it

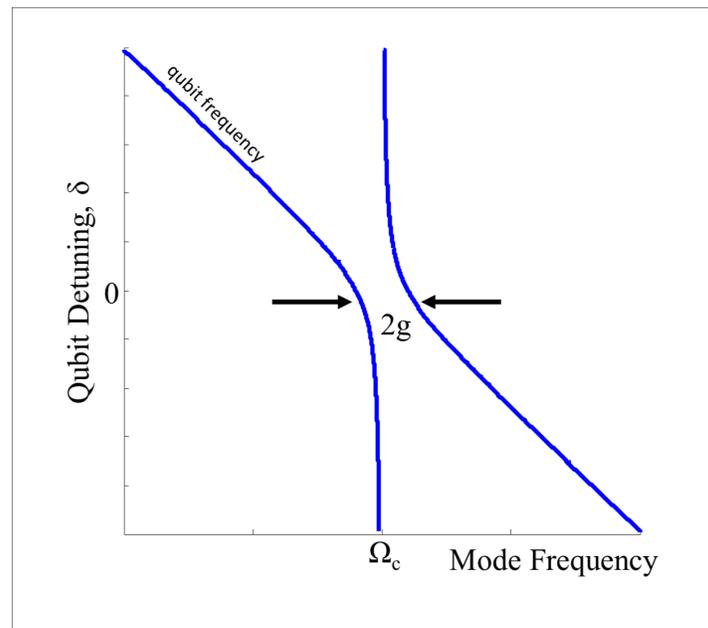


Figure 2.7: The Jaynes-Cummings model predicts that as the qubit, originally in its ground state, approaches the cavity frequency, an anti-crossing between the two resonance frequencies appears. The two frequencies are split by $2g$, where g is the rate at which they are exchanging photons in the time-domain.

through various modes of interest. By inducing anti-crossings in each mode individually we get a measure of the interaction rate, g , for each mode.

To get around the mode-crowding problem we will use parametric interactions. In that case, the mode frequency and atom frequency are all tuned to an optimum (e.g. highest Q) and the energy swapping is induced by introducing an external pump. We will show that these parametric interactions can induce Rabi-like swapping, but can be turned on and off with an external parameter/drive.

First, though, we need to introduce our artificial atom (‘transmon’) and then move on to parametric coupling physics.

2.3 Artificial atoms

In this section we combine the ‘microwave engineering’ of the first section and the ‘cavity QED’ of the previous section to build a model of our artificial atoms. In our group we chose to use the transmon style ‘atom’ (interchangeably ‘qubit’) whose use dominates the field at the moment. The development of the various flavors of qubit have been discussed *ad nauseam* in references [2, 13, 22], and so we just introduce the design parameters and outline the physics needed to understand the rest of the work in this thesis.

When we introduced superconducting cavities earlier we only mentioned that below the critical temperature (T_C) of certain materials, currents can flow with extremely low resistance. The well-known, microscopic, ‘BCS’ theory of superconductivity describes how below T_C electrons in the metal begin to pair up into Cooper pairs. This behavior is due to a small, attractive interaction mediated by the lattice of metal nuclei [42]. Since electrons are spin-1/2 particles, the combined two-electron ‘particle’ can have either spin-0 or spin-1 — thus acting like a Boson rather than a Fermion. This is an important distinction since Bosons can all relax to a single quantum ground state, or condensate, separated from unpaired electrons (‘quasiparticles’) by an energy gap of $2\Delta \simeq 3.52k_B T_c$ (at $T = 0$). This condensate

of Cooper pairs forms a single, macroscopic quantum state which we can describe by a single wavefunction. This wavefunction can maintain coherence throughout macroscopic, bulk materials. In our case both the metals making up our cavities and the metals which will eventually form our artificial atoms are based on this property.

If instead of a continuous piece of metal we now imagine two spatially distinct superconducting wires separated by a thin, insulating barrier — for example, superconducting aluminum metal + insulating, oxidized aluminum + superconducting aluminum metal — the wavefunctions describing the Cooper pairs must be discontinuous across the forbidden region of the dielectric (in the forbidden region, the wavefunction is exponentially suppressed). Yet, if this insulating layer is thin enough compared to a ‘coherence length’ the two wavefunctions can become coherent and form one wavefunction again. In order to maintain coherence between the two metals a phase difference between the two wavefunctions is accumulated.

The dynamics of this *Josephson Junction* are described by the two Josephson relations:

$$I = I_o \sin(\delta) \tag{2.73}$$

$$\frac{d\delta}{dt} = \frac{2e}{\hbar} V = \frac{\Phi_o}{2\pi} V \tag{2.74}$$

which describes the current, I , that can flow by applying a voltage between the two metal layers [43]. δ is the phase difference (discontinuity) between the two superconducting wavefunctions on either side of the barrier, I_o is the maximum possible *supercurrent* which can be supported through the barrier, and Φ_o is the magnetic flux quantum. The maximum supercurrent, I_o is determined by junction fabrication parameters – i.e. the thickness of the insulating barrier. More details can be found in the next chapter regarding the specific geometry and fabrication of these junctions.

One can lithographically define various *linear* circuit elements — inductors, capacitors, and transmission lines — by patterning thin films of aluminum easily. But, by including Josephson Junctions one can also fabricate *non-linear* circuits. This is the key idea that allows us to make artificial atoms. This can be seen by noting that if the current through

the junction is changing in time, there will be a voltage across the junction. This is very much the behavior one would expect from an inductor so we can define a ‘Josephson inductance’,

$L_J = \frac{\Phi_o}{2\pi} \frac{1}{I_o \cos(\delta)}$ and write down the analog of Faraday’s Law of induction:

$$V = L_J(\delta) \frac{dI}{dt}. \quad (2.75)$$

The Josephson junction can be now used in configurations with capacitors and other elements, below, to form the two circuits which are used in the experimental chapters to come.

2.3.0.1 Transmon Qubit

The transmon qubit is simply a Josephson Junction in parallel with a capacitor. The typical energy scales of the two circuit elements are: $E_J = \frac{I_o \Phi_o}{2\pi} = \frac{\Phi_o^2}{L_J}$, the energy of an electron tunneling across the junction; and $E_C = \frac{e^2}{2C_\Sigma}$ the charging energy of an electron stored between the capacitor plates ($C_\Sigma = C_g + C_J$ — the sum of the capacitance to ground and the junction capacitance, respectively).

By choosing the ratio of E_J/E_C — or equivalently the values of C and L_J — one can produce various shapes of the potential energy of the circuit. For large ratios (above about 80), the potential is almost harmonic with a non-linear correction. The Hamiltonian accordingly describes a harmonic oscillator (with frequency $\omega = \sqrt{8E_J E_C} \sim \frac{1}{\sqrt{L_J C}}$) summed with nonlinear corrections thusly (cf: equation 2.54):

$$H = H_o + V_{nl} \quad (2.76)$$

$$= \frac{1}{2} \frac{Q^2}{C_\Sigma} + \frac{1}{2} \frac{\Phi^2}{L_J} + V_{nl} \quad (2.77)$$

$$= \frac{(2e)^2 n^2}{2C_\Sigma} + \frac{E_J \cos(\delta)}{2} \quad (2.78)$$

$$= 4E_C n^2 + \frac{E_J \delta^2}{2} + E_J (-O(\delta^4) + O(\delta^6) + \dots) \quad (2.79)$$

where Q and Φ are conjugate charge and flux variables and V_{nl} are nonlinear perturbations to the harmonic potential. The third line uses the fact that n Cooper pairs of charge $2e$

can tunnel. The nonlinearity is expressly folded into the $\cos(\delta)$ term also. The fourth line expands the $\cos \delta$ term and should be compared to the second line. At this point we have now switched our basis from Q and Φ to n and δ . Inherent in doing the cosine expansion is the fact that in the transmon limit of large E_J/E_C it is small variations in δ that are relevant – in other words: δ is the ‘good quantum number’ to quantize, not n (or charge).

The quantization of δ follows as before. Specifically, we can write $\delta = \frac{2E_C}{E_J}(a + a^\dagger)$. The Hamiltonian (keeping $O(4)$ terms) becomes

$$H = \hbar\omega\left(a^\dagger a + \frac{1}{2}\right) - \frac{1}{12}E_C(a + a^\dagger)^4 \quad (2.80)$$

$$\approx \hbar\omega\left(a^\dagger a + \frac{1}{2}\right) - \frac{E_C}{2}(a^\dagger a^\dagger aa) - E_C(a^\dagger a) \quad (2.81)$$

$$= \hbar\omega'\left(a^\dagger a + \frac{1}{2}\right) - \frac{E_C}{2}(a^\dagger a^\dagger aa) \quad (2.82)$$

where we have again applied a rotating wave approximation to drop rapidly rotating terms from the expansion of the $(a + a^\dagger)^4$ – only leaving $(a^\dagger a^\dagger aa)$ and $(a^\dagger a)$ terms. As promised, equation 2.82 has nonlinear energy levels – the first two energy levels are spaced by $\omega' = \omega_{01} = \sqrt{8E_J E_C} - E_C$ and $\omega_{12} = \sqrt{8E_J E_C} - 2E_C$.

If we wire two Josephson junctions in parallel rather than just a single junction we form a SQUID (Superconducting Quantum Interference Device) loop. By changing an external DC magnetic flux, Φ , threading the loop formed by the two junctions we can tune the effective E_J of the circuit [16]. If we choose to make the two junctions symmetrically, each with half the original E_J . The total E_J (or inductance) remains the same since inductors add in parallel. Thus the resonant frequency of the 0-1 transition of the qubit can be modulated via $\omega_{01} \simeq \sqrt{8E_C \frac{E_J}{2} |\cos(\frac{\pi\Phi}{\Phi_0})|} - E_C$. A plot of predicted ω_{01} and ω_{12} frequencies versus applied SQUID flux is shown in Figure 2.8 for the device featured in Chapter 7.

Practically speaking, our transmon is formed by two, thin-film, metal pads lithographically patterned on a sapphire substrate (see next chapter). The capacitance between the two plates is typically about 70fF and the Josephson inductance at zero external flux is approxi-

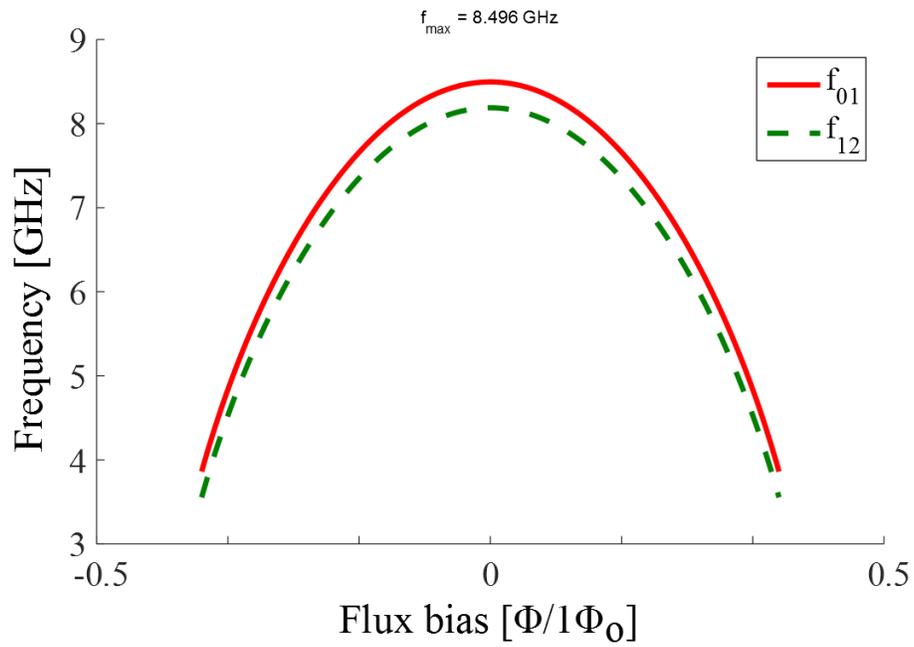


Figure 2.8: The predicted transmon f_{01} and f_{12} frequencies (energies) for the transmon circuit featured in Chapter 7. An external flux bias can tune the effective critical current of the SQUID which forms the non-linear inductance of the circuit.

mately 5-10nH. Using parameters in this range results in transmon resonance frequencies in the 5-10 GHz range.

To reiterate, even though our nonlinear circuit is essentially just a harmonic oscillator with unequally spaced internal energy states, if we restrict ourselves to the bottom two energy levels and apply the mechanics of cavity QED we can derive similar vacuum Rabi splittings between our artificial atom (with a large, engineer-able dipole moment) and cavity modes. The higher energy levels *do* play an important role in corrections to calculations and models, but we note that much of the Jaynes-Cummings machinery is still valid. The non-linear and linear circuit models are compared in Figure 2.12, below.

2.4 Parametric interactions

As discussed after introducing the Jaynes-Cummings Hamiltonian, and above, if a system has more than one mode, more than one qubit, and only static/resonant interactions to work with, the system develops unwanted coupling very quickly. But, if one is able to turn on and off interactions without having to pass through other resonances the system is less constrained. These non-resonant interactions can be implemented via parametric physics.

The use of non-resonant interactions implies the use of only static biasing fields and all-microwave frequency control pulses to switch the interactions on and off. As will be shown in later chapters, it also allows for additional control over the coupling phase, as well as the ability to sequence coupling between any arbitrary pair of modes. This general approach possesses the potential to reduce wire count in quantum computing architectures since coupling in multi-mode systems can be frequency-multiplexed. Finally, this work demonstrates a methodology which may be useful in novel approaches to quantum computing that use linear modes for state storage [44].

The concept of parametric interactions has a simple intuitive picture, shown in figure 2.9. If two masses on springs (harmonic oscillators) of different frequency are *weakly* coupled

by a third spring they will not exchange energy. This notion of weak coupling implies that any static coupling constant be less than the decay rate of either of the two oscillators.

Now, if we modulate the coupling-spring at a frequency that equals the difference between the two oscillator's frequencies, we can *pump* energy from one oscillator to the other. The bottom panel of Figure 2.9 shows an oscillation starting in mode A (red), and being swapped back and forth between the two oscillators while changing the spring constant of the coupling spring in time with $k(t) = k_o \cos(\omega_P t + \phi_P)$. Here $\omega_P = |\omega_A - \omega_B|$, and ϕ_P is the phase of the pump.

Of course, we'll be using electrical (or cavity) harmonic oscillators and the role of the modulate-able spring will be played by a circuit element containing Josephson Junctions. We will describe why we chose to use this particular element for our modulate-able 'spring' in the next section, for now though consider it a tunable impedance. This should be a familiar concept from equation 2.75.

Rather than deriving this result from considering voltages and currents of various circuit elements, we would like to use the more general approach for calculating fields by using 'input-output' formalism. This formalism has the added advantage that it produces results for the input field, output field, and intracavity field all of which will be necessary to consider when building up more complex systems later. We begin by sketching the theory for a single cavity connected via a single port to the outside world; then consider parametrically coupled cavities afterward.

2.4.1 Input-Output Introduction

There are many ways to calculate and model the behavior of our system, but in this section we introduce *input-output* formalism to describe the dynamics [45]. We very well could have used Master equations or rigorous field theory approaches. Input-output theory has the advantage that it is very similar to the ideas of network analysis from microwave

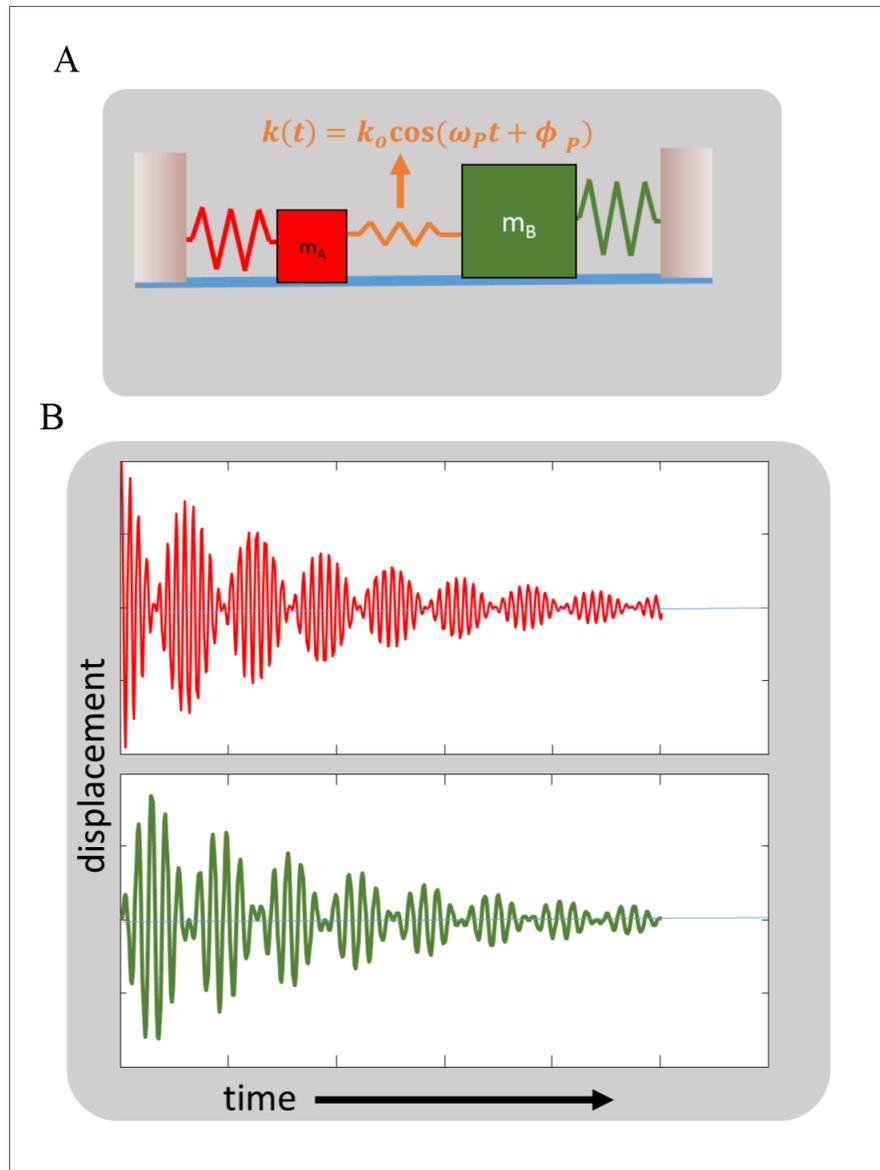


Figure 2.9: If two simple harmonic oscillators (mass A and B) are weakly coupled, energy can be swapped back and forth by modulating the coupling spring constant as in (a). If the first mass is initially set in motion as shown in the top panel of (b), the oscillations of mass A can be periodically transferred to mass B and back.

engineering – the dynamics of the cavity depend on the input and output, or equivalently the incident and reflected waves sent to probe the cavity dynamics. For a full derivation of input-output theory Chapter 3 of reference [46] is recommended.

For all of the calculations in this section we will be in a rotating frame of reference, centered on some cavity mode frequency Ω . The main implication of this is that when performing integrals over frequency, the limits of integration are shifted from $\int_0^{+\infty}$ to $\int_{-\Omega}^{+\infty}$. In fact, extending the limits to $\int_{-\infty}^{+\infty}$ should not introduce significant error to our calculation if $\Omega \gg 0$.

To make the connection to the previous sections clear we begin with a similar Hamiltonian as equation 2.63 – the sum of three terms describing a cavity mode amplitude, the external environment, and interactions:

$$H = H_s + H_r + V \quad (2.83)$$

$$= H_s + \hbar \int_{-\infty}^{+\infty} \omega (b^\dagger(\omega)b(\omega)) d\omega + i\hbar \int_{-\infty}^{+\infty} [b^\dagger(\omega)g(\omega) - g^\dagger(\omega)b(\omega)] d\omega. \quad (2.84)$$

H_s describes the system Hamiltonian – a function only of internal mode operators a and a^\dagger . H_r describes a continuum of heat baths, described by $b(\omega)$ and $b^\dagger(\omega)$, of the external environment which are linearly coupled to our system. The interaction term V describes this coupling between our system and bath with strength $g(\omega)$.

We can now write down the Heisenberg equations of motion, in the interaction picture, for $a(t)$ and $b(\omega, t)$ as

$$\dot{a}(t) = \frac{i}{\hbar} [H_s, a] - i \int d\omega g(\omega) b(\omega, t) \quad (2.85)$$

$$\dot{b}(\omega, t) = -i\omega b(\omega, t) - ig^*(\omega) a(t). \quad (2.86)$$

The solution to equation 2.86 (for the fields outside the cavity, the reservoir) is found by integration. This integration is split into two regimes though. The first solution takes the integration limit from t_o to t will denote input fields. The second solution uses the limit

t to t_1 and denotes output fields. An alternative interpretation of these integration limits is for right-going (incident) and left-going (reflected) propagation.

In frequency-space, the two solutions are:

$$b(\omega, t) = b_o(\omega)e^{-i\omega(t-t_o)} - ig^*(\omega) \int_{t_o}^t dt' a(t')e^{-i\omega(t-t')} \quad (2.87a)$$

$$b(\omega, t) = b_1(\omega)e^{-i\omega(t-t_1)} + ig^*(\omega) \int_t^{t_1} dt' a(t')e^{-i\omega(t-t')}. \quad (2.87b)$$

Here we make our second approximation: we assume that $g(\omega)$ is a slowly varying function of frequency. We set

$$|g(\omega)|^2 = \frac{\kappa}{2\pi}. \quad (2.88)$$

This statement is termed the Markov approximation – which states that our bath modes have no ‘memory’. Intuitively, having a broad slowly varying bath coupling in frequency-space means very short time-correlations in the time-domain. These time-correlations are what we refer to as a bath memory – all memory effects are destroyed very quickly. Mathematically, this says that $\int_{-\infty}^{+\infty} e^{-i\omega(t-t')} d\omega = \delta(t-t')$.

How does the input bath mode affect the cavity fields? We now insert the input solution (equation 2.87a), into equation 2.86 and apply the Markov approximation simplifications.

The *intracavity* equation of motion becomes

$$\begin{aligned} \dot{a}(t) &= \frac{i}{\hbar}[H_s, a] \\ &\quad - i\sqrt{\frac{\kappa}{2\pi}} \int b_o(\omega)e^{-i\omega(t-t_o)} - i\frac{\kappa}{2\pi} \int_{t_o}^t dt' a(t')e^{-i\omega(t-t')} \end{aligned} \quad (2.89)$$

$$\dot{a}(t) = \frac{i}{\hbar}[H_s, a] - i\sqrt{\kappa}a_{in} - \frac{\kappa}{2}a(t). \quad (2.90)$$

Where we have defined $a_{in}(t) \equiv \frac{i}{\sqrt{2\pi}} \int d\omega b_o e^{-i\omega(t-t_o)}$ – which describes how the external reservoir affects the intracavity field at an initial time, t_o . The third term is interpreted as the decay of the intracavity field to the bath at a rate $\kappa/2$.

By doing a similar simplification to equation 2.87b and defining

$$a_{out}(t) \equiv -\frac{i}{\sqrt{2\pi}} \int d\omega g(\omega)b_1 e^{-i\omega(t-t_1)} \quad (2.91)$$

we can write equation 2.86 as either

$$\dot{a}(t) = \frac{i}{\hbar}[H_s, a(t)] - \frac{\kappa}{2}a(t) + \sqrt{\kappa}a_{in}(t) \quad (2.92a)$$

$$\dot{a}(t) = \frac{i}{\hbar}[H_s, a(t)] + \frac{\kappa}{2}a(t) - \sqrt{\kappa}a_{out}(t). \quad (2.92b)$$

Now, by subtracting 2.92b from 2.92a we have

$$a_{out}(t) + a_{in}(t) = \sqrt{\kappa}a(t). \quad (2.93)$$

This equation relates the input (incident) field, the output (reflected) field, and the intra-cavity field just as we set out to do.

To check ourselves, if our system Hamiltonian is simply that of a harmonic oscillator: $H_s = \hbar\omega_o a^\dagger a$; equation 2.92a is just

$$\dot{a}(t) = -i\omega_o a - \frac{\kappa}{2}a(t) + \sqrt{\kappa}a_{in}(t). \quad (2.94)$$

which has the solution, in the frequency-domain:

$$\tilde{a}(\omega) = \frac{2\sqrt{\kappa}}{\kappa - 2i(\omega - \omega_o)}. \quad (2.95)$$

Which is precisely the same Lorentzian lineshape we calculated for e.g. the RLC harmonic oscillator, above.

2.4.2 Parametric Input-Output Model

We would now like to extend the input-output formalism to a system of two cavities, a parametric coupler (tunable impedance) between them and one output port. We will further make the connection between the input-output formalism and circuit theory more concrete.

To reiterate the concept, even if two *classical* RLC circuits are connected as in Figure 2.10 with $Z(t) = \text{constant}$ the two resonators would *not* exchange energy if and only if they have different resonance frequencies. In that case, resonator A sees a high impedance from

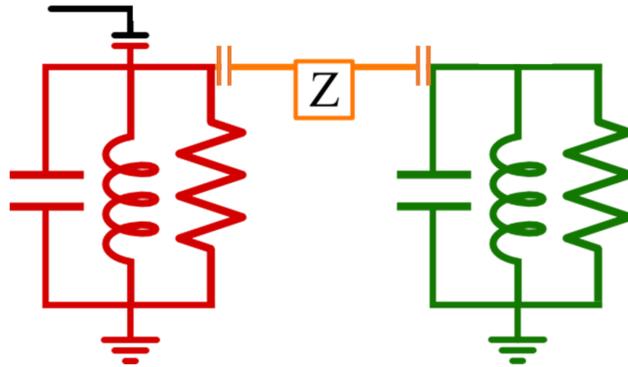


Figure 2.10: Lumped element model of coupled RLC circuits. Now the coupling element has been generalized to a modulate-able impedance, rather than a modulate-able spring constant.

resonator B at $\omega_A = 1/\sqrt{L_A C_A}$ and reflects all of the power back to A; and resonator B see a high impedance from resonator A at $\omega_B = 1/\sqrt{L_B C_B}$ (recall e.g. equation 2.15).

The fact that the two resonators each ‘see’ each other as a high impedance when they are not co-resonant means that we can insert our coupling element between the two circuits and use that to change the boundary condition that either one sees. To be more concrete, if the $Z(t)$ in Figure 2.10 is replaced by a flux-tunable inductor (a SQUID), by changing the inductance of the SQUID at a ‘pump’ frequency ω_P , one can change the boundary condition of the circuits and therefore induce parametric coupling. This must happen at a frequency nearly equal to the difference frequency between the two resonators – i.e. $\omega_P \sim |\omega_A - \omega_B|$.

By this rationale, we can write the equations of motion for each cavity as

$$\dot{A} = -i\omega_A A - igB - \frac{\kappa_A}{2}A + \sqrt{\kappa_A}A_{in} \quad (2.96a)$$

$$\dot{B} = +i\omega_B B - igA - \frac{\kappa_B}{2}B + \sqrt{\kappa_B}B_{in}. \quad (2.96b)$$

These are very similar to 2.94, but with an added term. To avoid confusing notation, the two modes are labeled with capital A and B – B does not refer to the b-operator (bath operator) above. Equation 2.96a has four terms. In order they are: the free harmonic oscillator term, a coupling term to cavity B, a term describing the field leaving A to the external bath, and the effect of an incoming field. The terms in equation 2.96b have similar interpretations.

Here g , though, denotes a parametricly-induced coupling between the two cavities. This factor depends on the particulars of the coupling circuit so to keep this general we will hold off on describing the specific form for g until Chapters 4 and 5. To write g as a static term we have converted into a doubly-rotating frame of reference where $H_s = \hbar\frac{\Delta\omega_P}{2}(a^\dagger a - b^\dagger b) + \hbar g(a^\dagger b + ab^\dagger)$ [47]. In this frame of rotation, the pump detuning from the optimal ($= \omega_A - \omega_B$) is denoted $\Delta\omega_P$.

In addition to equations 2.96, we also need the boundary conditions $A_{out}(t) + A_{in}(t) = \sqrt{\kappa_A}A(t)$ and $B_{out}(t) + B_{in}(t) = \sqrt{\kappa_B}B(t)$.

Proceeding to the frequency domain via Fourier Transforms of 2.96 yields the equations

of motion:

$$-i\omega A = -i\frac{\Delta\omega_p}{2}A(\omega) - igB(\omega) - \frac{\kappa_A}{2} + i\sqrt{\kappa_A}A_{in}(\omega) \quad (2.97a)$$

$$-i\omega B = +i\frac{\Delta\omega_p}{2}B(\omega) - igA(\omega) - \frac{\kappa_B}{2} + i\sqrt{\kappa_B}B_{in}(\omega) \quad (2.97b)$$

Looking ahead, our final experimental setup will only have access to the fields of mode A, so we can set $B_{in} = 0$ – i.e. there will be no driving inputs from the environment to the B-cavity. Additionally, just as with the case of network analysis, above, we are interested in reflections off of the A-cavity. We define a similar reflection coefficient as before to be $\Gamma = \frac{A_{out}}{A_{in}}$ which can be calculated to be:

$$\Gamma = \frac{4g^2 + (i\kappa_A + \Delta\omega_P - 2\omega)(i\kappa_B + \Delta\omega_P + 2\omega)}{4g^2 + (-i\kappa_A + \Delta\omega_P - 2\omega)(i\kappa_B + \Delta\omega_P + 2\omega)} \quad (2.98)$$

In order to move back to the non-rotating frame – i.e. back to the lab frame, we make the substitution $\omega \rightarrow \omega - \omega_A + \Delta\omega_P/2$. In this reference frame, the reflection coefficient, as a function of frequency driving near ω_A is:

$$\Gamma = \frac{g^2 + (i\kappa_A/2 + \omega_A - \omega)(i\kappa_B/2 + \Delta\omega_P + \omega - \omega_A)}{g^2 + (-i\kappa_A/2 + \omega_A - \omega)(i\kappa_B/2 + \Delta\omega_P + \omega - \omega_A)}. \quad (2.99)$$

This result is plotted in Figure 2.11. The qualitative form of this behavior is reminiscent of the avoided crossing as the result of the Jaynes-Cummings model. Indeed the overall effect is similar – energy (photons) are induced to swap between two systems and results in normal-mode splittings. Yet, the two systems are never in resonance. Thus this technique is called parametric *frequency conversion*.

There are two important differences here though. First, if we do not pump the parametric interaction, the coupling is nearly zero and the two systems undergo free evolution. Second, the two systems connected here are at a static frequency – there is no need to move the frequencies quickly or through other modes (and thereby risking unwanted couplings via the Jaynes-Cummings model).

Below, we describe the design parameters and device desiderata for the coupling element which we will use in the experiments to follow.

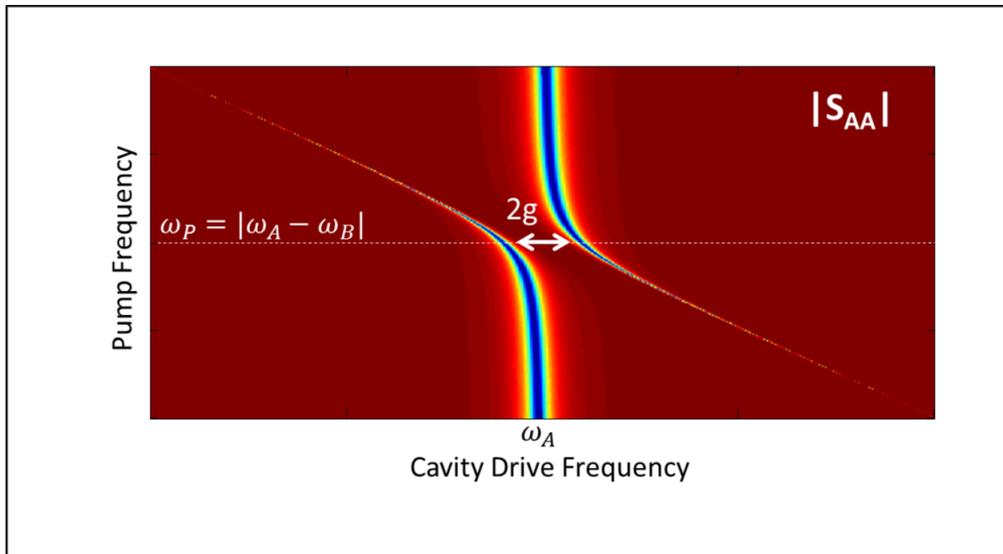


Figure 2.11: Predicted parametrically-induced mode splitting by pumping a tunable impedance connecting two RLC circuits.

2.4.3 Coupler circuit

Just as with the artificial atom or transmon circuit described above, our coupler device consists of a long antenna which dipole-couples to cavity electric fields. The electric fields induce current to flow along it, but the antenna wire is interrupted by Josephson Junctions. Again we will be using SQUIDs rather than junctions since we anticipate needing to modulate the effective inductance of this element via an external pump.

Multiple SQUIDs in series are used rather than a single SQUID as in the case of the transmon. The use of multiple junctions is necessitated by two factors – the need for large pump amplitudes, and to build up significant ‘participation ratios’.

First, as we show theoretically and experimentally in Chapter 5, the rate of parametric frequency conversion – the induced splitting size – is proportional to the amplitude of the pump. Yet, if the pump amplitude induces currents in excess of the critical current of the junctions, the junctions will cease to operate as Josephson Junctions and become ‘normal’ – i.e. dissipative. This would result in strong losses and destruction of any delicate quantum states. To fabricate a high critical current SQUID is contradictory to the goal of a device with high Josephson inductance – see equation 2.75.

The second factor that necessitates multiple junctions is due to a participation ratio argument. The participation ratio for an example mode is $\partial\omega_n/\partial\Phi$ – meaning if we change our tuneable parameter, how much does that affect the mode’s frequency? The participation ratio can be understood by looking at the two limits and thinking about where energy is stored. If the ratio is very small, the Josephson inductance (tuneable boundary condition) will have very little effect on the mode – almost all of the energy in the system will be stored in the mode inductor, not in the coupler. If the ratio is very large – i.e. comparable to the mode inductor – the coupler will hybridize with the cavity’s inductor and energy will be equally stored in both. Thus, we want a large, but not too large L_J – typically about 1-10%. This is contradictory to the goal of having larger critical currents since $L_J \propto 1/I_o$. But, many

high critical current junctions connected in series will achieve both goals simultaneously.

Separately, in order to use a flux pump, we chose to split these junctions into two, and thus form SQUID loops. This choice allows for a DC bias to tune to the optimal participation ratio by choosing the point in the $\partial\omega_n/\partial\Phi$ curve with the largest slope. We can then use our pump to apply modulation around that point. For example, see data in Figure 5.2.

All of these design constraints result in a much more harmonic structure of the internal energy states of the coupler as compared to the very anharmonic states of the transmon. Yet, we can still exploit the tunability of the Josephson junction/SQUID configuration as a circuit element. In Figure 2.12 we summarize the salient features of the two design goals before moving on to Chapter 3 where we detail the practical experimental design aspects of this work. We are less concerned with the self-resonance frequency of the coupler circuit since it can just as easily be inductive (below its resonant frequency) or capacitive (above its self-resonance) and still achieve the desired effect.

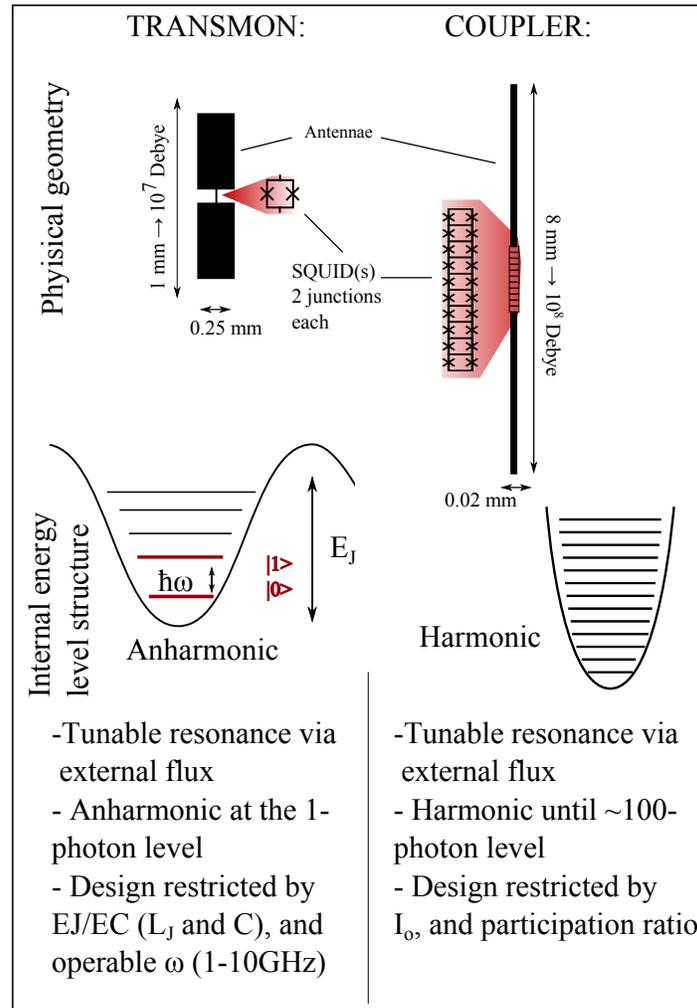


Figure 2.12: A summary of the difference between the ideal qubit and coupler circuit highlighting some of the length scales and energy scales mentioned in this chapter. In particular, we note that transmons require much smaller junction areas (low critical currents), whereas the coupler circuit is designed to use high critical current junctions.

Chapter 3

Experimental Design

With chapters 1 and 2 behind us we now turn to practical aspects of the experimental design of this work. I focus mostly on device modeling/simulation, device fabrication, and microwave measurement techniques. One benefit of doing my research at NIST is that I was able to be involved in *all* of these, some of which could be whole chapters on their own. Part of the reason for this is that my graduate career was interrupted during mid-2011 to early-2013 by our lab being shut down, packed-up, and moved to new space in a new building on the NIST campus. This new laboratory space expanded our capabilities greatly and building-up a lab ‘from scratch’ was a good (though often painful) learning experience.

The new space updated our 1950s-era NIST labs to the modern era. The Precision Measurement Lab (PML) features excellent temperature and humidity control, a stable power supply, modern lighting and increased collaborative space. Additionally, we received new cryogen-free dilution refrigerators and the use of a substantially larger, all-around better cleanroom. The devices featured in this work are some of the first Josephson-junction based devices born of the new cleanroom.

3.1 New lab space

The new Precision Measurement Lab at NIST was a big improvement to the measurement capabilities of our group. The biggest improvement has been the installation of three new ‘dry’ dilution fridges. While these specific models have been difficult to get up and

running (leaks, poor design, more leaks, etc) they now seem to be stable and reliable. The old ‘wet’ dilution fridges were then freed up for use with the elcro-mechanics experiments of John Teufel et al, the fear being that the pulse tube cooling system used by the wet fridges would be noticeable in the delicate mechanical oscillators they fabricate. Through careful isolation and mounting this has been shown to not affect their measurements recently.

The overall design philosophy has been to wire the fridges to be capable of testing multiple devices. To that end, the fridge where most of these measurements were taken had twelve microwave coaxial cables, and six DC capable lines. Two ‘output’ lines were installed with low-noise amplifiers. With this wiring setup, and using two Radiall microwave switches (50 Ohm, six pole single-throw 6PST) up to twelve single-port devices could be measured in one cooldown. In general, though, most of our devices require a few ports or pumps so the fridge was never operated fully-loaded. This setup was very useful early-on to test various cavity geometries and couplings (external Qs).

A very important, and often overlooked aspect of a new laboratory setup is to have proper grounding of the dewar to avoid ground loops. A ground loop can be thought of as an uncontrolled antenna in the lab space which can pick up stray fields from equipment or act as unwanted resonances. A fair amount of effort was needed to achieve this and also allow DC voltages to be applied to bias lines, etc. To avoid this, all microwave signals were sent through DC blocks – a device which breaks the inner and outer conductors independently – before entering the fridge; DC signals were produced using floating sources; and measurement equipment was carefully isolated and grounded separately.

With those steps taken, our microwave-frequency devices still need careful measurement and amplification to avoid destroying our delicate quantum states. This design is as follows.

3.2 Microwave measurement setup

The energy of a room temperature photon corresponds to a frequency of $k_B \times 300K / \hbar \sim 2\pi \times 6$ THz. But, as discussed in the previous chapter, our circuits operate in the 1-15 GHz range. The energy of those photons is equivalent to $(\hbar\omega)/k_B = (1.05 \times 10^{-34})(2\pi \times 10^9)/(1.38 \times 10^{-23}) \approx 0.5K$. Thus we need to operate at temperatures $T \ll 0.5K$. Additionally, since our devices are usually fabricated from aluminum, we need to be far below the superconducting transition temperature of $T_C = 1.2K$.

In order to probe our GHz circuits using room temperature equipment we must therefore reduce the energy of the probe photons dramatically. If we did not reduce the energy of incoming photons, they would have enough energy to permanently excite our delicate quantum states out of their ground state and into higher levels.

For six of the microwave coax lines, to lower the energy (temperature) of these incoming probe photons we can simply attenuate them at different cold stages of the dilution refrigerator. In these ranges of temperature and frequency the voltage Johnson noise spectrum, S_V , is approximately linear

$$S_V(\omega) \approx 4k_B T R \quad (3.1)$$

thus a cold attenuator will reduce the effective noise spectrum by a factor of $A = T_{hot}/T_{cold}$. We use 20dB of attenuation at 4K, 10 dB at 1K and 20dB at the 20mK mixing chamber (plus 10dB attenuation from the slightly resistive Copper-Nickel cabling) for a total of about 60dB. Which naively reduces the noise spectrum by six orders of magnitude.

The situation is slightly more complicated than that though. Each stage of attenuation thermalizes the incoming photons from the previous stage and re-radiates photons at the current stage's temperature. The number of added photons at each stage must be combined with the number subtracted by attenuation. Each subtraction is determined by the attenuation factor, A , and the number of photons added at stage i can be calculated from the

Bose-Einstein distribution $n_i = 1/(e^{h\omega/kT} - 1)$. For example:

$$n = n_{20mK} + A_{20mK} \times (n_{Still} + A_{Still} \times (n_{4K} + A_{4K} \times (n_{300K}))). \quad (3.2)$$

The 60dB of attenuation which we chose should produce a photon noise background at our range of frequencies < 0.1 – acceptable for our measurements.

After spending all of that effort attenuating the microwave signals on the way into the fridge, directly measuring the now very small signals coming out of our device is impractical, thus we need to now ‘un-do’ some of the attenuation via an amplification chain. It is important to efficiently amplify the signal as it propagates *out* of the fridge such that we collect as much of the signal as possible. This is complicated by the fact that the amplifier we use is a semiconductor-based High Electron-Mobility Transistor (HEMT) based amplifier. The amplifiers were purchased from CalTech/JPL and have been used for low noise observational astronomy applications for many years. Their bandwidth, $\sim 1 - 14GHz$, sets the bandwidth of the measurement setup.

These amplifiers dissipate a significant amount of power since they are resistive by nature yet still have a low noise temperature of 5-10K. This dissipated power (estimated to be 20mW) would overwhelm the cooling power of the dilution system at the mixing chamber and heat our device up though (the cooling power at 30mK is $\sim 10\mu W$!). Thus, we operate these amplifiers at the 4K stage of the fridge where the cooling power is capable of handling an extra load without detrimental base temperature changes. To get the signal from our device(s) to the HEMT amp properly we must use very low-loss (superconducting niobium) coax to carry our signal.

This is a two-way street though – blackbody photons (at about 5K) from the amp could then propagate with low loss in the *opposite* direction (from amp to devices). To avoid this, a dual-stage, wide-band, cryogenic-compatible, microwave isolator (manufactured by Pamtech/Quinstar) is used at the base temperature stage of the fridge. This non-reciprocal device allows signals to propagate in one direction, but the reverse propagation is suppressed

– typically by about 70dB. This blocks the blackbody radiation from the amplifiers at higher temperatures from shining back onto our devices, but comes at a slight loss – the isolators have some unavoidable insertion loss. It is possible to minimize this again, with the inclusion of another amplification stage (with a lower noise temperature) directly at the mixing chamber, but these amplifiers are quantum, parametric, devices on their own – typically a Josephson Parametric Amplifier or Josephson Traveling Wave Parametric Amplifier which come with complications in and of themselves. For the sake of simplicity we avoided using a JPA, but a future device or quantum computer will most likely need to incorporate just such a device.

Another complication arrives because superconductors (Niobium) have very poor thermal conductivity [20]: it can be difficult to make sure the inner conductor of the Nb coax is properly thermalized to the 4K stage (the isolator thermalizes the other end well). To accomplish this, a home-made circuit is used to electrically (and thus thermally) connect the inner conductor at 4K to ground through a large inductor. This inductor looks like a high impedance to our signals propagating to the 50 Ohm amp, so the signal is not distorted yet allows electrons to thermalize.

The main measurement chain – heavily attenuated coax plus careful amplification – carries our main signal, but we often need to provide strong pump tones too. The circuits which are being pumped are very often more linear and need to be pumped with higher powers than say a qubit. Thus, the remaining six (of twelve) coax line on the input side are only attenuated by 10dB at 4K, 10dB at 1K, and 10-20dB at 20mK. This results in a higher noise background, but none of that extra noise reaches our delicate quantum states nor has enough power to ‘turn on’ parametric processes which we will later turn on via a pump tone. We often combine the AC pump tones with DC bias fields using cryogenic-compatible bias tees (mfg: Anritsu), though in later devices it was found to be easier to provide an external DC field using superconducting magnet coils.

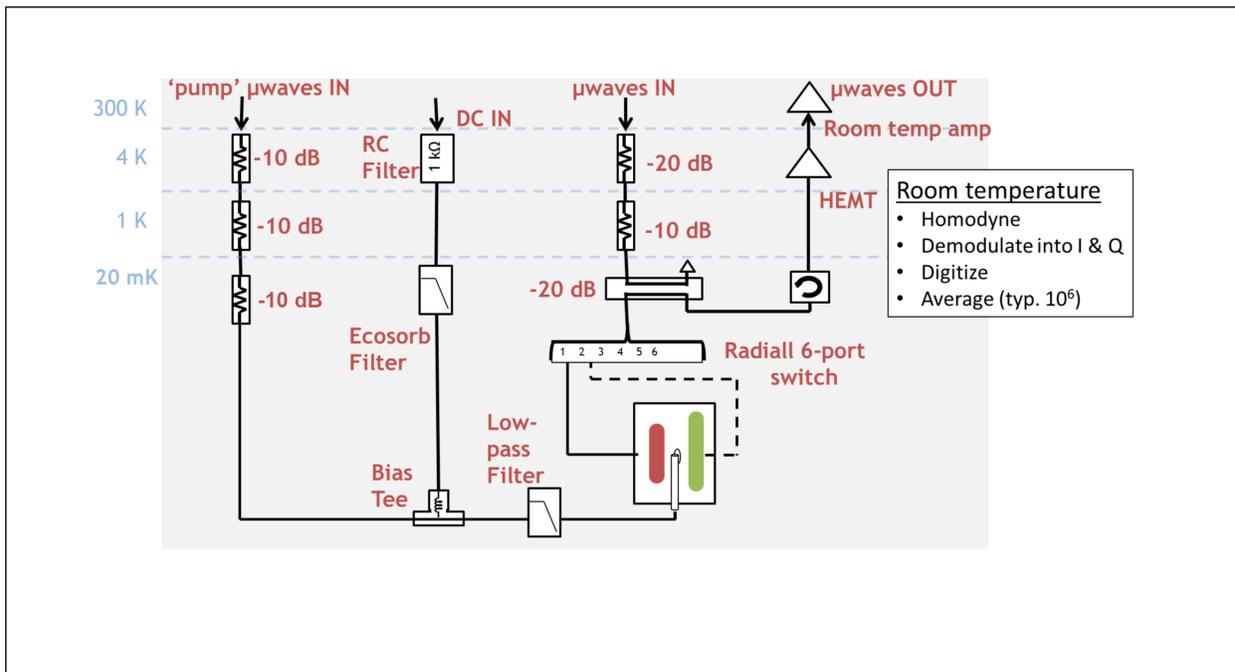


Figure 3.1: Typical measurement setup that allows for multiple devices to be measured each cooldown. Not shown are the 5 additional pump lines, 5 additional DC lines and 5 additional input lines.

DC biases

The DC bias currents can reach several mA , so the bias lines cannot be attenuated like the microwave signals since Joule heating would produce too much heat load to the fridge. Thus, we employed a strategy of using strong low-pass filters. These filters allow the DC currents to pass with little dissipation, but absorb (and dissipate at higher temperatures) the unwanted noise at higher frequencies. This strategy has two parts.

The first part – anchored to the 4K stage of the fridge is a homemade RC filter. Surface mount components (made from Nickel-Chrome which is stable with temperature) are soldered to a custom printed circuit board (PCB) such that they form a low-pass filter with cutoff 10 MHz. With *only* the RC components, the filter loses efficacy at higher frequencies (near 1 GHz) as shown in Figure 3.2a. Thus, an epoxy (Emerson-Cummings brand Eccosorb CR-124) is then injected into the filter body to suppress any stray signals over the desired pass band. This epoxy acts as another low-pass filter, but begins absorbing around 2GHz. The ‘R’ of the RC filter was chosen to be 1 kOhm. This is reasonable for the currents we expect to need and provides a stiff current bias.

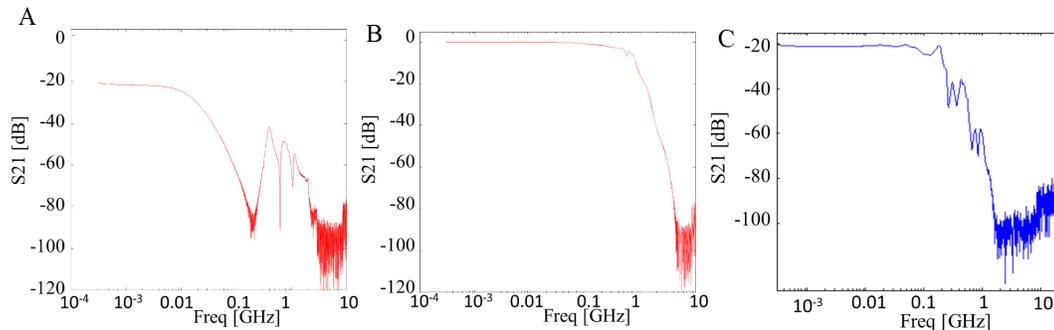


Figure 3.2: Frequency response for the low-pass DC line filters at room temperature A.) shows the response for the (lossy) RC filters located at the 4K stage of the fridge. B.) shows the response for the low-pass LC chip filters. Both are measured with the CR-124 epoxy already in place. The combined response results in a 10MHz low-pass filter that is flat at high frequencies, shown in (c), measured at low temperature (note the change in plotted x and y axes).

At the mixing chamber we don't want to have resistive components which would produce Joule heating, thus another PCB was designed to have chip based LC filters soldered to it (Mini-circuits LFCN style filters). These filters, in their non-chip form factor, have been used at low temperatures for many years in our labs so they are empirically cryogenics-compatible. Again, though their low-pass properties only extend to about 1 GHz so they were also injected with CR-124. These filters and pass bands are shown in Figure 3.2b. In Figure 3.2c, we show a full measurement of a typical DC line including both low-pass filters installed in situ.

3.2.0.1 Outside the fridge

Microwaves

We perform three types of measurement in this thesis: continuous-wave (CW) stimulus-response measurements; time domain long-pulse (quasi-CW) measurements; and short-pulse time-domain measurements. The CW measurements are all performed using a vector network analyzer (Agilent VNA 5071C). This instrument scans a probe tone in frequency to the device under test and measures the amplitude and phase of the returning signal. At each stimulus frequency we can also drive other ports of our device using a separate CW tone from say another RF generator. The drawback of using a VNA is that it only gives continuous-wave, frequency-domain, data. The benefit of using a VNA is its automation and ability to calibrate frequencies and powers very well.

We can replicate the measurements done by a VNA by using an IQ mixer and digitization hardware in the setup shown in Figure 3.3 with the added benefit that we can retrieve non-CW, time-domain responses. By using a software-controlled room-temperature switch, shown in the figure we can quickly switch between VNA and IQ measurement mode. This acts as a coarse way to calibrate and compare between the two.

An IQ mixer compares its 'RF' port to a CW local oscillator (LO). In our measurement

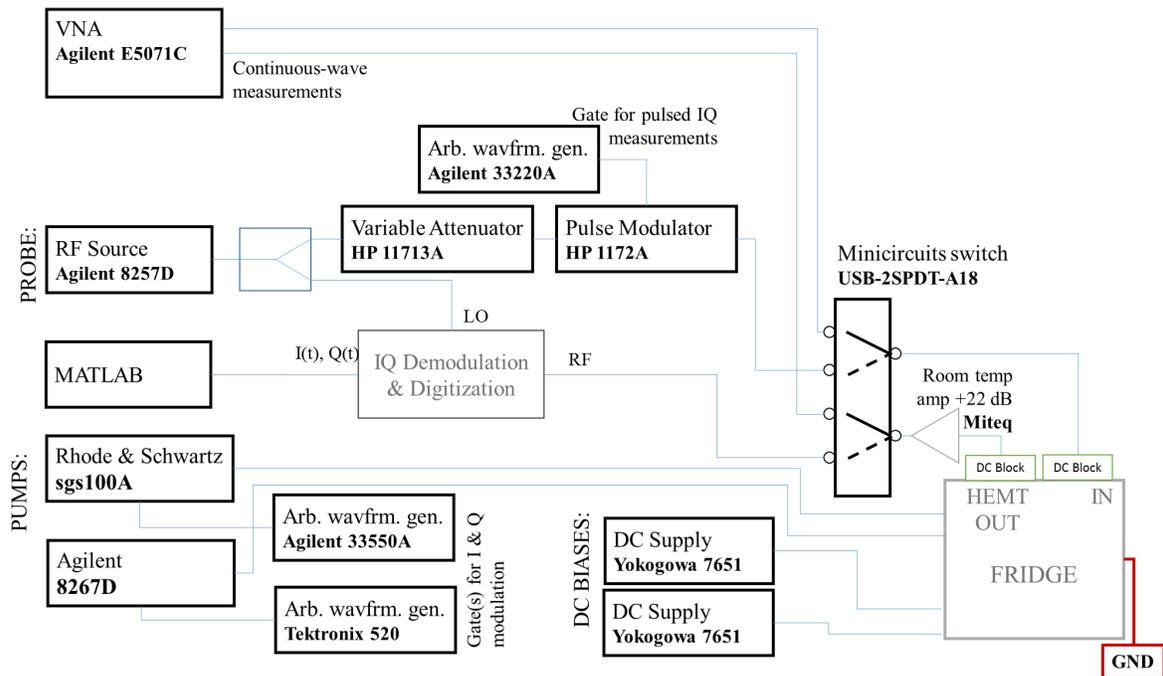


Figure 3.3: The measurement schematic for the instruments outside the fridge. All of the instruments are on their own, separate ground from the ‘fridge ground’ to avoid ground loops. The IQ modulation and digitization is described and detailed below.

setup, the LO provides both signals to the IQ mixer which can be helpful to negate drift between sources or other DC offsets. The LO signal from the source is split such that most of its power drives the IQ mixer, but ~ 20 dB less power is sent toward the fridge. This split signal, which we now call the RF tone, can be gated in time using a pulse modulator diode (fast (few ns) switching with ~ 80 dB on/off ratio), and its amplitude further attenuated if required with a step attenuator. This pulsed RF signal probes our device through the heavily attenuated lines inside the fridge, then is amplified by the HEMT and another room temperature amplifier (Miteq) on its way out.

The returning signal coming out of the fridge is compared to the LO signal. Since they originated from the same source, they are both at the same frequency and thus the mixer mixes the signal down to DC frequency. A DC voltage is sampled which tells us information about the two quadratures of that signal. If the pulses are long enough, the system will equilibrate to a steady-state response, which agrees with the CW case of the VNA above.

In contrast to the VNA though, once the pulsed RF tone is turned off, we can continue recording the time-domain responses. In this manner we can record transient or decay signals. Making the RF pulse very short results in controlled dynamics where, for example, only one swap (or two, or three swaps) are allowed to happen at very precise times. Since this is at the heart of most of the measurements described in this thesis we expand on this measurement technique below.

Homodyne with IQ mixers

The general problem we are faced with is that electronics do not exist to sample and record our signals at GHz frequencies. Thus we require some mechanism to down-convert those signals to a lower (heterodyne) intermediate frequency, ω_{IF} , or DC frequencies (homodyne). The device which does this is a mixer.

A mixer takes two inputs – the ‘RF’ tone and the ‘LO’ tone – and multiplies them together. If our RF signal is $A(t) \sin(\omega_{RF}t + \phi(t))$ and the LO is $\sin((\omega_{IF} + \omega_{LO})t)$, how do

we determine the amplitude $A(t)$ and phase $\phi(t)$ of the RF signal?

Internal to the IQ mixer the LO is split into two branches – with a phase shift of $\pi/2$ added to the second branch. The mixer compares the two LO signals to the RF and outputs the resulting two quadratures (we get from multiplying the two sines together):

$$I(t) = A(t) \cos(\omega_{IF}t + \phi(t)) \rightarrow A(t) \cos(\phi(t)) \quad (3.3)$$

$$Q(t) = A(t) \sin(\omega_{IF}t + \phi(t)) \rightarrow A(t) \sin(\phi(t)) \quad (3.4)$$

Since we want our resulting signal to be at DC frequency we choose $\omega_{IF} = 0$ (or $\omega_{RF} = \omega_{LO}$ (and filter out the 2ω mixing products). Each of those signals are output to two ports labeled I and Q .

In this manner the I signal (‘In-phase’) records the real part – $I(t) = B(A(t) \cos(\phi)) = \text{Re}(\exp(i\omega_{RF}t + \phi))$ – of the signal relative to the LO; and the Q signal (‘Quadrature’) records the imaginary part – $Q(t) = B(A(t) \sin(\phi)) = \text{Im}(\exp(i\omega_{RF}t + \phi))$. Here B is an overall constant which describes the mixer gain.

From these two numbers a phasor is described by $V(t) = I(t) + iQ(t)$. This is precisely the same signal recorded by a network analyzer and is easily converted into magnitude and phase – i.e. $|V| = \sqrt{I^2 + Q^2}$ and $Phase = \angle(V)$.

The I and Q signals can now be digitized using a two-channel, variable sampling rate, data acquisition card (GaGe CG4300). At each sample step the DC voltage is mapped to 12 bits of DAQ range. To assure Nyquist sampling, and maximal signal to noise the signal coming into the card is filtered and amplified as shown in Figure 3.4. Typically, we average these measurements 10^4 to 10^6 times.

The inverse of this process – using a DC signal to modulate an RF signal happens internal to *vector signal generators* and is used to produce fast pulses of RF tones. Aside from careful construction of the pulse shape and being watchful to maximize the pulse on/off ratio, this mixing is transparent.

IQ Demodulation and digitization details:

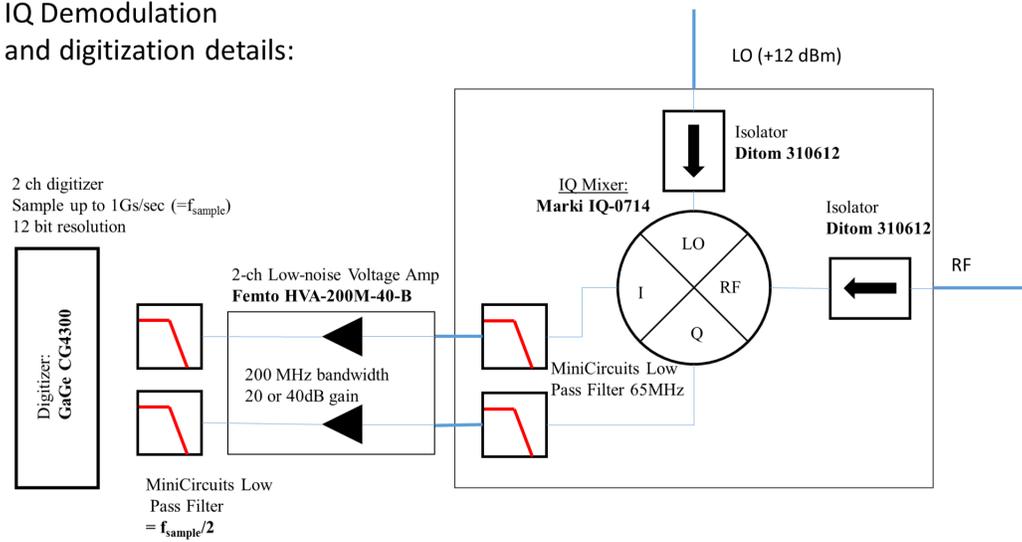


Figure 3.4: The IQ demodulation scheme used in this experiment for recording time-domain responses.

3.3 Electromagnetic simulation and modeling

A powerful tool of our field which has reduced the computational and theory overhead has been the use of advanced electromagnetic simulation software to design and model the behavior of our circuits. As discussed in Chapter 2, aside from a few very simple lumped-element circuits, it becomes very hard to calculate precisely the behavior of more complex circuits. Indeed, this software gives us the ability to quickly and efficiently simulate the *environment* in which we place our circuits. As we showed in Chapter 2, the field structure of the cavities can determine the internal Q of a resonance and the frequency of those resonances is determined almost solely by geometrical considerations.

I used *FastHenry/FastCap*, *Microwave Office*, and *HFSS* software extensively throughout this work to explore various parameter spaces of design ideas. In general, my work flow begins with a simple Microwave Office simulation with lumped elements. Then on to FastHenry/FastCap which are specifically tailored to calculating superconducting circuit element properties. Then, back to Microwave Office to add more accurate models for capacitances and inductances. Finally, a full scale HFSS simulation is performed to verify or modify the simpler simulations.

All of these simulators work by breaking up a drawn design into a ‘mesh’. The software then applies the appropriate boundary conditions for each mesh element and solves Maxwell’s Equations – i.e. finite element analysis. Often, this procedure is iterated a handful of times by the software. With each iteration the previous solution is compared to the current one to determine in which areas of the geometries the solutions changed the most. In these regions, the software re-meshes the design (usually increases the number of mesh ‘slices’ or decreases the mesh size) and recalculates the electromagnetic fields.

FastHenry/FastCap is the simplest of the three softwares, and only does this mesh/solve process once. Additionally, the solutions are in only two dimensions, Microwave Office includes some interlayer connections and is thus termed ‘2.5’ dimensions, and HFSS is fully

3D. The added 0.5 dimension in each software adds complexity and therefore solution time, but gives very accurate answers when compared to what is finally fabricated.

FastHenry and FastCap

FastHenry and FastCap are programs that are built into our CAD software, *Xic*. These two programs use simple mesh shapes (squares and rectangles) to tile the surface of the drawn circuits. To accommodate this mesh shape we often modify a design to only have straight lines and ninety degree angles in the drawing to get an approximate solution for capacitance and inductance. If the number of mesh tiles gets to be too big the software is unable to compute a solution – so again this is usually a starting point for more complex designs.

When compared to other methods of simulation, FastCap seems to underestimate the capacitance by about 5%, most likely due to the fact that the ground reference of the box/chip/cavity is not accurately simulated in only two-dimensions.

FastHenry was developed for simulating superconducting circuits, it includes corrections to the electromagnetic simulation to account for the kinetic inductance of the materials used. For example, our coupler chip described below has long antennae which form the circuit’s inductance. These wires have an additional kinetic inductance of:

$$L_k \approx \frac{\lambda\mu}{2 \tanh(\frac{d}{2\lambda})} \left[\frac{L}{W} \right], \quad (3.5)$$

where d ($\sim 100nm$) is the thickness of the thin-film aluminum and L and W are the length and width of the antenna. The London penetration depth, λ of aluminum is about $15nm$. This correction can be made small – for our antenna it is $L_k \sim 0.001nH/mm$ compared to a typical inductance for a wire $L \sim 1nH/mm$.

Microwave Office

Microwave office has a suite of built in lumped-element and transmission line models which can be easily ‘wired’ together to yield complex designs very quickly. The solution to these circuits can often be calculated in a few seconds – compared to a few hours with complex HFSS models.

This software also includes a CAD-like drawing program in order to simulate the actual lithographically-patterned, less-lumped, circuits. Again, though the simulation is performed in 2.5 dimensions but, since HFSS performs fully three-dimensional solutions this CAD-feature was only used sparingly. We tend to use HFSS to feed back results (stray capacitances, resonant frequencies, etc) into the lumped-element model.

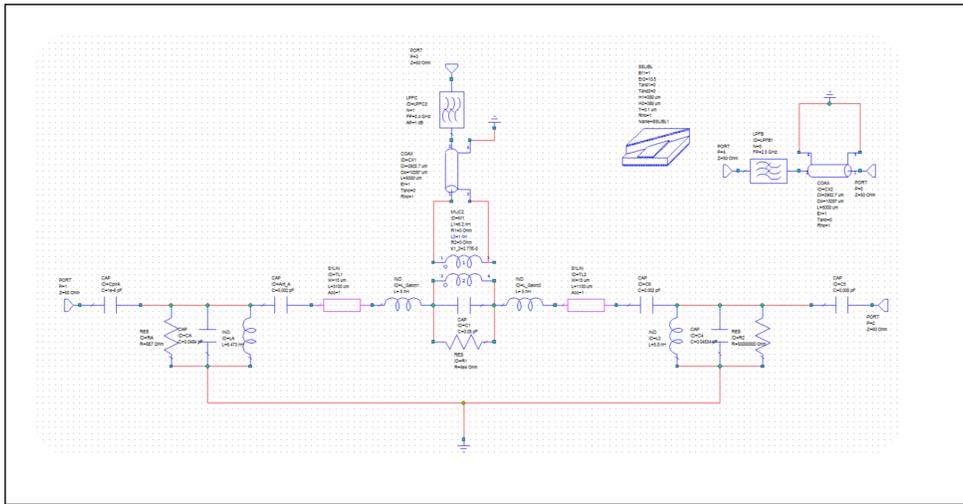


Figure 3.5: A typical Microwave Office model of the experiment described in Chapter 5 – the simple coupled RLC circuits of Figure 2.10 has been modified after noting different frequencies/stray capacitances/external couplings from the HFSS model. The simulated performance of this model is calculated in a few minutes.

In Figure 3.5, we show a Microwave Office model of the experiment in Chapter 5, which is nominally two RLC resonators connected by a coupler circuit. The simple model was modified with coupling capacitors, stray capacitances, antenna models, and drive ports to reflect the actual, fabricated device.

HFSS

A typical HFSS simulation is shown in Figure 3.6 – again for the device described experimentally in Chapter 5. Rather than modeling the two cavities as RLC equivalent circuits, the cavities are drawn fully to scale. Additionally, the drive ports (grey and yellow in the Figure) can be modeled true to life. Other factors, like the coupling circuit, though

much smaller in scale can also be drawn as patterned thin films on a chip.

HFSS has two styles of solution. The first calculates eigenmodes for a given geometry and will plot the resulting field strengths as vectors. Since our circuits are designed to couple to electric-field vectors via a dipole coupling, this can be useful for determining where to place the chip to maximize or minimize couplings to certain modes (and not others). By including estimates for the surface resistance of the metal used, this style of solution will also estimate the internal Q possible.

The second style of solution will propagate a wave to a designated port in the drawing. This wave then propagates, and resonates, in the drawn structure. In this way, external Q s and drive couplings can be simulated. This method of solution solves for the electric (and magnetic) fields for a user defined list of frequencies, and therefore can plot those fields not only on resonance, but off resonance also.

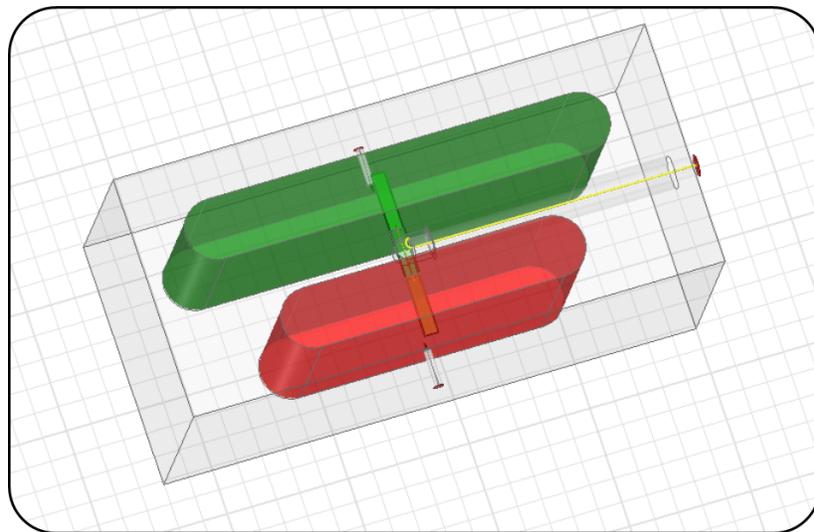


Figure 3.6: The HFSS software will simulate Maxwell's equations at all points in 3D space. This is a screenshot of the software simulation for the experiment described in Chapte 5. The red and green volumes represent 'rectangular' cavities milled out of a block of superconductor (clear outline). The probe signals propagate along the three yellow coaxial transmission lines.

3.4 Micro-fabrication

The key technology that allows our group to explore the physics of quantum optics with artificial atoms is cleanroom microfabrication. We are lucky at NIST to have a production-ready, but research-focused cleanroom (again, brand new as of 2012) focused on superconducting devices. This section of the thesis would be much longer if I were to list all of the fabrication projects I have worked on, as this has been a focus of my time and effort for much of my years at NIST (before and during graduate school). I will thus only describe the fabrication relevant to the devices which are used in the experiments described in the following chapters.

A feature of the NIST cleanroom is its very stable and robust wafer-level processing capabilities. Wafer-level fabrication and lithography generally uses *optical* wavelengths to expose patterns on a wafer. This is in contrast to many other superconducting quantum computing groups around the world who almost always produce devices at the single chip level (a full wafer can contain hundreds of single chips). The lithography in these groups use electron beams (e-beams) to write patterns rather than optical light. There are tradeoffs between the two technologies, and our cleanroom can do both. The biggest benefit of doing optical lithography is for producing mass quantities of devices with feature sizes greater than $0.1 \mu m$. E-beam lithography is better at producing a few devices with small feature sizes (nm resolution).

For readers unfamiliar with the basics of microfabrication it proceeds as follows. The device designs are written onto a photoresist covered quartz plate called a mask. The mask-making machine, a pattern generator, writes the design to the quartz at five times the intended final dimensions. This allows for very accurate writing – it’s easier to write big things than little. After developing and cleaning the mask (similar to steps involved with film processing), the mask is inserted into a stepper aligner machine. This machine shines light through the mask, then through a 5x reducing lens, onto the surface of a wafer coated with

photoresist. Since the designs are much much smaller than the full 3 inch diameter wafer's area, the machine accurately *steps* the area being exposed in a grid pattern to maximally cover the wafer. Many different designs/masks can be swapped in and out of the machine while preserving alignment between various 'layers'. Indeed, many *wafers* can be produced back-to-back with extremely high reproduce-ability.

The benefit of 'e-beam' lithography is the ability to reliably fabricate 50 nm^2 Josephson Junctions (using the angle-evaporation techniques, below). E-beam lithography is limited though in that it is much slower than wafer-level optical lithography – to cover a single wafer e-beam exposures might take several hours vs minutes for a typical optical lithography exposure.

3.4.1 Angle-evaporated junctions

The fabricated devices in this thesis are very simple in comparison to most cleanroom fabricated devices, so should be ideally suited for for mass production using optical lithography with high yield. But, the design parameters (EJ/EC) dictate a small junction area which is less than the resolution of optical lithography. Using techniques used mainly by e-beam lithographers, we were able to push the limits of the optical lithography to make nm scale junctions with some success though. Later, as quantum computing devices become more complex and require more and more junctions, these techniques and processes could be developed further, and integrated smoothly.

A cartoon of an ideal fabricated device is shown in figure 3.7. The process begins by spreading a two-layer photoresist (LOR5A and SPR660) on a blank, 3-inch, sapphire wafer. After being exposed to the designed pattern (optically, through a mask) the two resists are developed and removed in the areas exposed to light. The two resists develop at different rates though.

The bottom resist, LOR5A, was chosen such that it develops at a faster rate than the

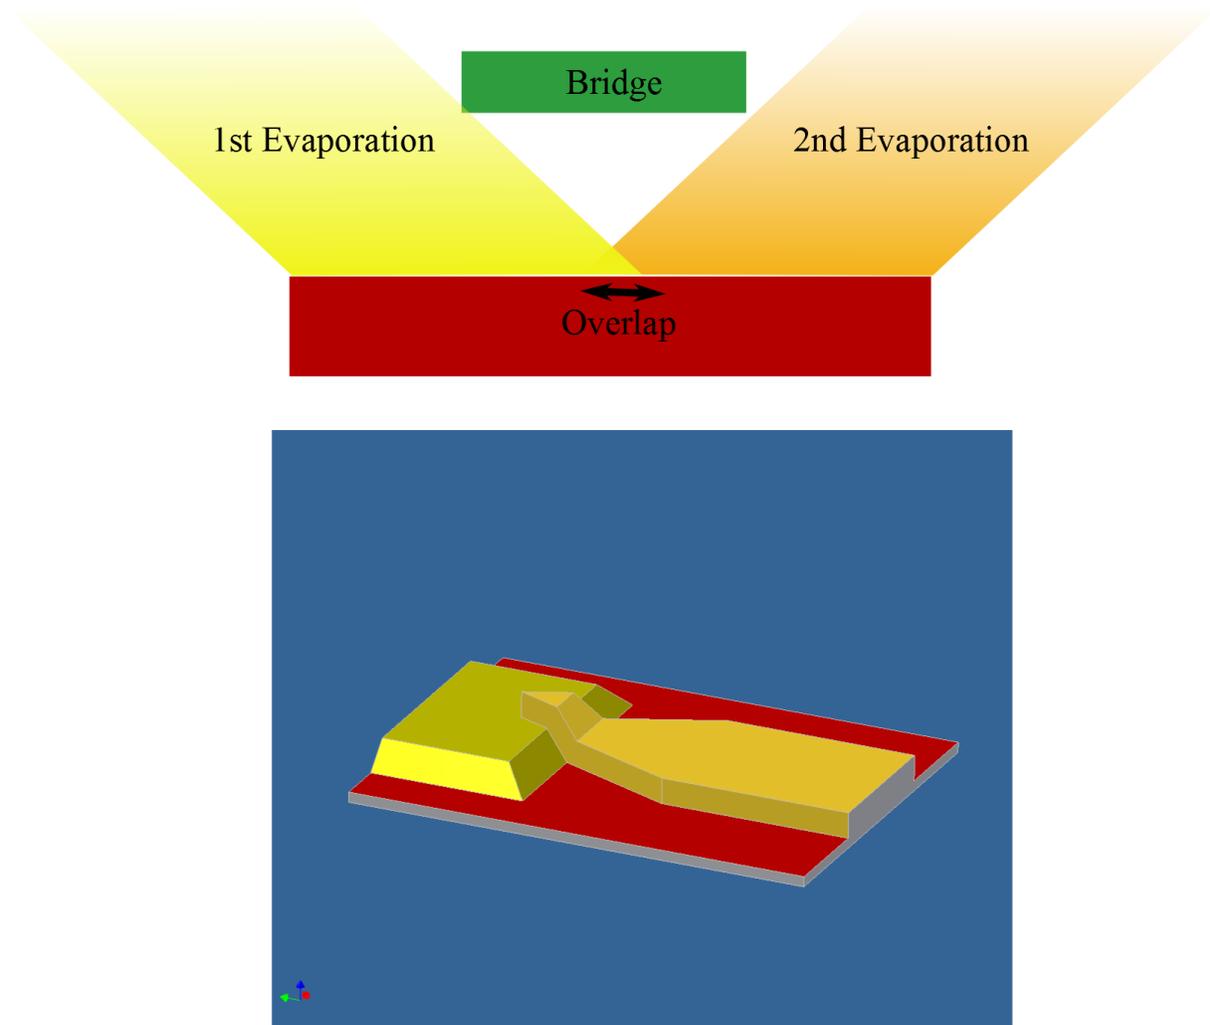


Figure 3.7: A side-view of the angle evaporation process – the suspended bridge and accurate angle control allows for careful control of the resulting amount of overlap. Below, the idealized result of the process. Nonidealities in the resist thickness, sharpness of the resist features (rounding), and bridge nonuniformity need to be considered.

top resist. This results in the same pattern being sharply defined in the SPR660, but much larger and rounded underneath in the LOR5A as if it were ‘overexposed’. This is a standard ‘lift-off’ technique [48]. We modify this lift-off technique by creating features that are smaller than the amount of ‘undercut’ which the LOR5A overshoots in the lower pattern. If this pattern is small enough (less than $1\ \mu m$), this results in a suspended layer of SPR660, with no LOR5A left underneath – a bridge-type structure.

Once this bridge is fabricated, two layers of metal can be evaporated at two different, opposite angles from the normal of the wafer such that the two evaporations overlap under the bridge as shown in Figure 3.7a. Between the two evaporation steps oxygen is introduced to the evaporation chamber to oxidize the first evaporated aluminum without breaking vacuum. The resulting structure is pictured in figure 3.7b. As shown schematically, the bottom layer (the first evaporation) is designed to be rectangular and the second layer is triangular. The resulting overlap is exaggerated in the cartoon, but the junction area is defined by this overlap.

If we were drawing this pattern using e-beam lithography techniques, the final junction would appear very similar to the cartoon in figure 3.7. But, optical lithography does not have the necessary built-in resolution necessary to make fine details like the sharp corner of the triangular second evaporation layer. A typical junction, as fabricated looks more like figure 3.8. The sharp triangle is now rounded. This is due to portions of the photoresist being overexposed and therefore rounder.

This rounding of the triangle results in a larger junction area than the designed value. To compensate for this, the original pattern was actually designed to be *even smaller* than wanted. Additionally, the angle which forms the triangle shape was optimized (as shown in Figure 3.8 and Figure 3.9a by the tapering shape) in order to avoid overexposure and therefore roundings. A detailed study of which mask patterns produce the *smallest* junction areas was a necessary process development step.

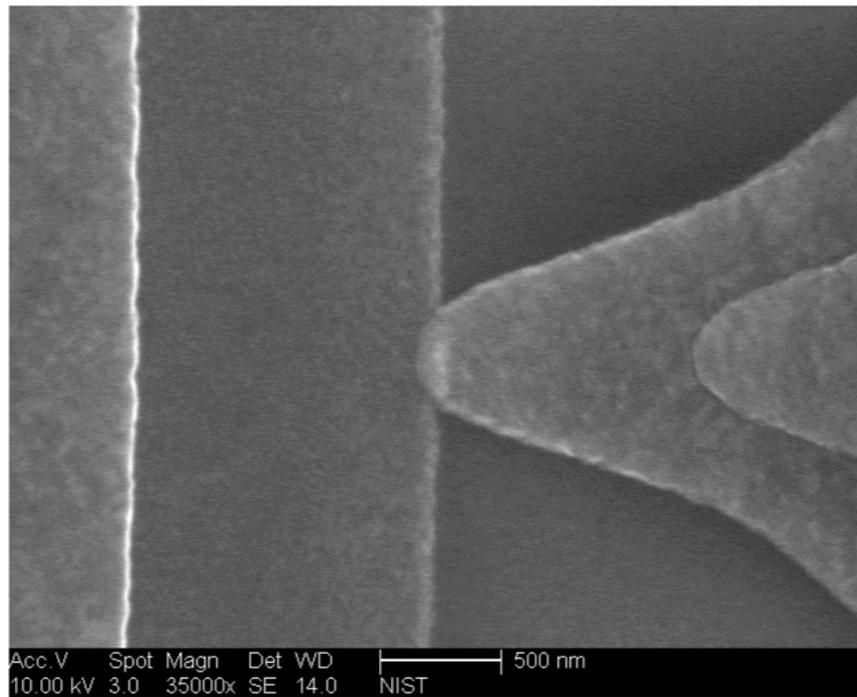


Figure 3.8: Scanning electron micrograph of a fabricated transmon Josephson junction. The two metal layers, evaporated at two angles are both visible. The small area where the triangle overlaps the square shape forms the junction. Compare to the idealized cartoon in Figure 3.7.

Qubit fabrication parameters

All of this effort is required to create $\sim 50nm^2$ area junctions. This number results from the design parameters of the transmon device. A typical E_J needed for a transmon device (converted to GHz) is $\frac{1}{2\pi} \frac{E_J}{\hbar} \approx 40GHz$. This then determines the critical current which produces that Josephson energy via $E_J = \frac{I_o \Phi_o}{2\pi}$, all other factors being constants.

The critical current *density* achievable using aluminum, aluminum oxide, aluminum Josephson junctions is set by the Ambegokar-Baratoff formula which says that the critical current density, J_o (critical current per area) is:

$$J_o = \frac{1}{R_n A} \frac{\pi \Delta}{2e} \tanh\left[\frac{\Delta(T)}{2k_B T}\right] \quad (3.6)$$

where Δ is the superconducting energy gap and R_n is the resistance of the junction in the normal state (room temperature), and A is the area. At low temperatures, the Ambegokar-Baratoff formula simplifies to $J_o = \frac{\pi \Delta}{2e R_n A} = I_o/A$.

This resistance-area product in equation 3.6 is determined by the thickness of the oxide barrier between the two aluminum layers. Controlling this thickness is done by changing the time and pressure of oxygen which the first evaporation is exposed to (before the second oxidation). Our specific aluminum evaporation chamber (JoeTek industries) controls these parameters very accurately, and we chose to use a recipe which yields $J_o = 2.9\mu A/\mu m^2$. From there we can calculate the normal resistance to aim for during fabrication to yield the proper transmon behavior. Thusly, we have connected the desired E_J to a critical current value, and then to an area which will yield that. By also choosing the oxidation time and pressure the normal resistance can be calculated, and measured at room temperature. For our transmon designs a junction resistance of 3-10 kOhms corresponds to a junction critical current below 100 nA, and qubit resonance frequency between 4 and 10 GHz.

Figure 3.9b shows histograms of a typical wafer fabrication run with three different junction designs. Again, the wafer-level fabrication allows us to make many test junctions spread across a wafer in order to produce these statistics (~ 65 junctions of each design

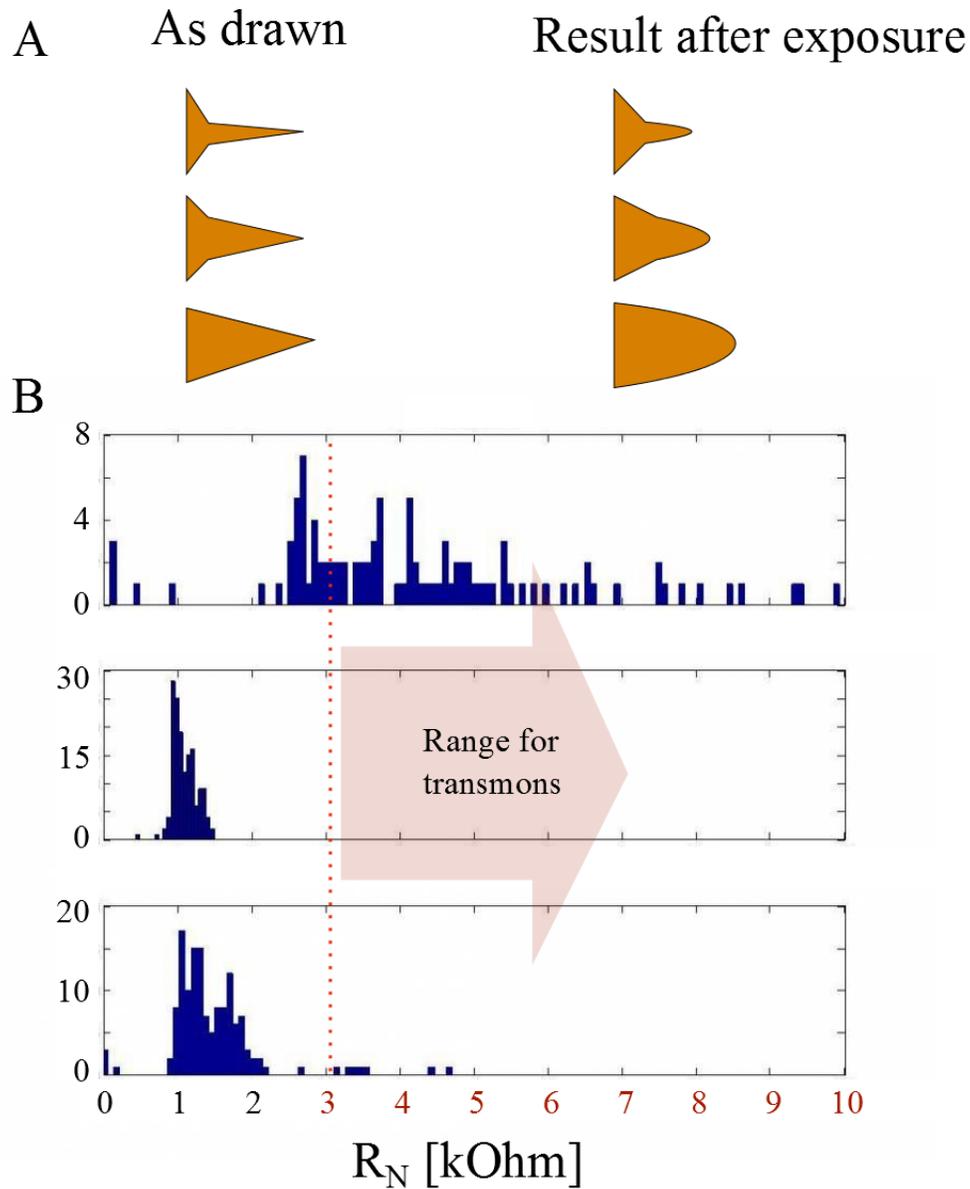


Figure 3.9: a.) A cartoon of the resist rounding problem with optical JJ lithography. To achieve sharp angles the resist must be overexposed, which results in wider 'triangles' than intended. But if we modify the design to compensate for this, sharp angled 'triangles' can be fabricated which result in small junction areas. b.) Histograms for the three junction designs from part (a): by over compensating for the rounding, small area (high R_n) junctions can be fabricated with low yield. The other two styles result in high-yield devices, but R_n 's that are too low.

per fab run). The top panel shows a design which tried to push the limits of the optical lithography to the very edge of working. This yielded *some* junctions in the good range, but the spread is large. The bottom two panels show junctions which were less at the limits of lithography. These junction designs were more reliable, but were too big area-wise (too low R_n) to be useful. Future work could certainly improve from both ends to achieve reliable, small (high R_n) junction fabrication. It is also possible to change the oxidation parameters such that the normal resistance decreases or increases.

As an experimental note, producing very pure films of aluminum, reliable oxide barriers, and very high-vacuum conditions is critical and worthy of pages of text. I benefited from the careful work of JoeTek industries greatly.

Coupler fabrication parameters

Since the junction critical currents for the ten SQUIDs for the coupler device necessitates larger junctions, the coupler junctions push the limits of optical fabrication much less than the transmon junctions but we require twenty of them (2 for each of 10 SQUIDs). These junction areas are about $0.5\mu m^2$, which yields $I_c = 1.5\mu A$.

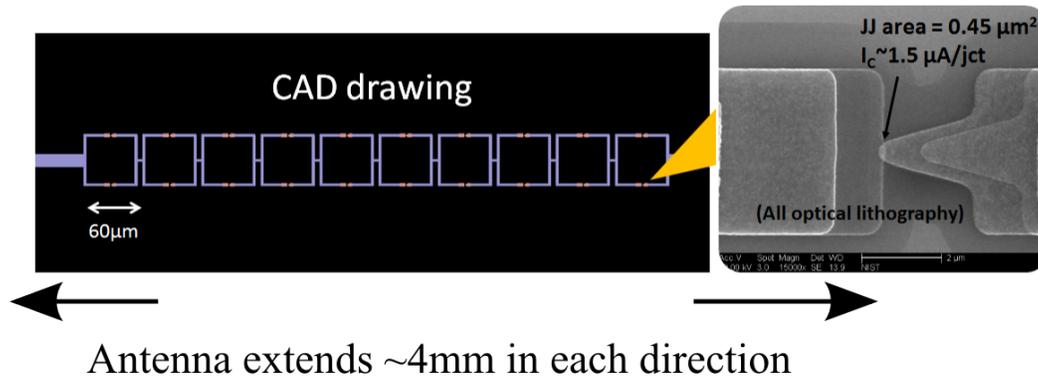


Figure 3.10: The coupler device is a $15\mu m$ wide wire (antenna) of 80 nm thick aluminum. It extends 8mm in length total. At the midpoint of the wire a series array of 10 SQUIDs interrupts the wire. Each SQUID junction is about $0.45\mu m^2$, and has a critical current of $\sim 1.5\mu A$. The inset is an SEM photo of a sample junction, note the difference in scale from the transmon junctions.

It's harder to conclusively say how reliably we made the twenty coupler junctions – i.e. what does the histogram of the coupler devices look like – because the individual junctions cannot be probed to determine their normal state resistance, only the series resistance. We have cooled down and tested five different coupler chips, though, and all have yielded very similar behavior. Thus, we suspect our fabrication yield is very good for these devices.

The SQUID layout, from the *Xic* software, and another SEM micrograph of the larger coupler junctions is shown in Figure 3.10 – note the larger scale bar in the SEM. The larger junction areas also place a less-stringent constraint on the photo-lithography and do not require the optimization of the mask designs, simple triangles produce reliable results.

3.5 Chapter conclusions

Chapters 2 and 3 described almost all of the necessary theory and experimental knowledge needed for the rest of the work described in this thesis. The next chapter reminds readers of some of the goals and approaches to keep in mind for the experimental sections which follow.

Chapter 4

Motivation for Specific Experiments

Before continuing with the second, data-focused, half of this dissertation it is worthwhile to pause and re-consider the motivations for the experiments to follow. First, we stress the practical aspects of using these systems and techniques for building quantum computing technologies, and second we discuss a more abstract benefit of the parametric interactions – namely the ability to design and modify Hamiltonians in situ. Then discuss why this provides additional motivation for coupling quantum systems in interesting ways.

4.1 Practical motivation

If we return to figure 1.1 for a moment we can begin to see how the ingredients from chapters 2 and 3 could be combined to be useful for future quantum computing technologies. Broadly, the necessary technologies are: 1.) at least one qubit to produce input states, 2.) the ability to manipulate that quantum state after it has been input to the ‘computer’ – this includes performing gates as well as moving the state between different constituents of the computer, 3.) a memory element – preferably a high-Q resonator, and 4.) an efficient and lossless readout mechanism.

The remaining chapters in this dissertation describe proof-of-principle experiments which demonstrate each of these individual aspects of the Von Neumann computing architecture. First, parametric frequency conversion is used to shuttle information around between modes efficiently and with minimum loss as highlighted in the blue dashed box in

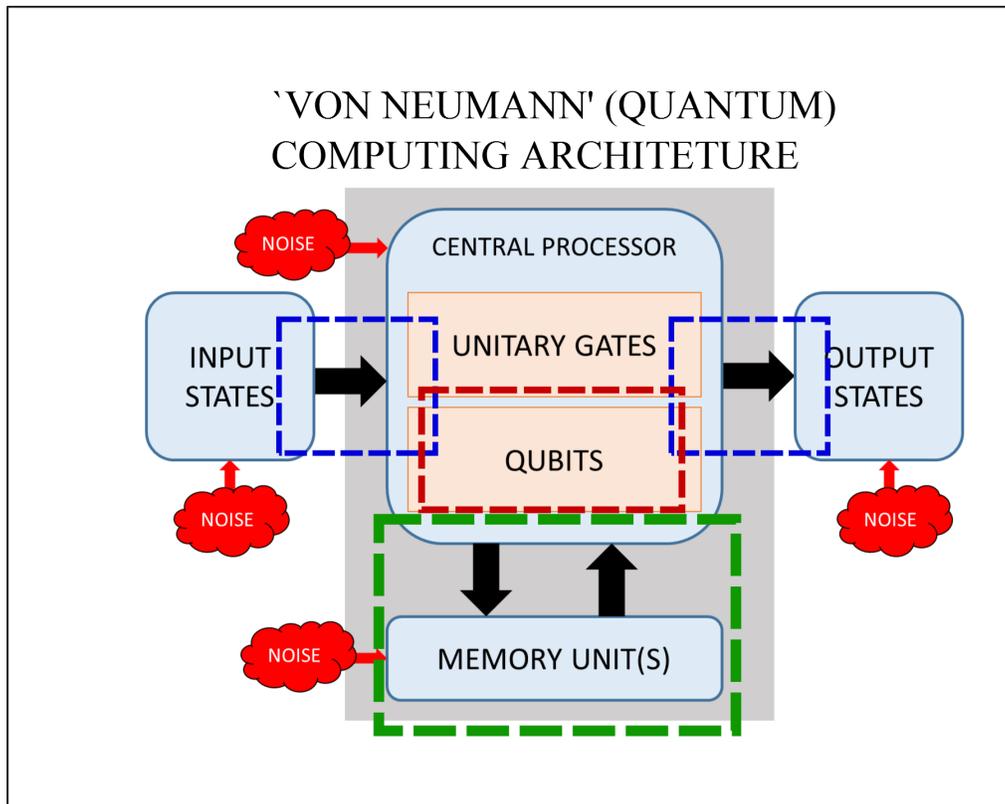


Figure 4.1: Returning to the ‘Von Neumann quantum computing architecture’ from chapter 1, we now highlight the goals of the experimental chapters which follow. The dark blue dashes show the goals of chapter 5; green chapter 6; red chapter 7. Chapter 8 attempts to connect all of these into one. See details in text.

figure 4.1. This technique will be used in chapter 5 to move coherent states from one cavity to another – converting between two frequencies *and* two spatially isolated cavities. If one of the cavities is high-Q and the other low-Q we can implement a basic read/write memory that fulfills requirement (3) above.

Designating a full cavity volume to read/write takes up a lot of valuable experimental real estate, so in chapter 6 we implement a very similar protocol between frequencies of a whispering gallery mode resonator. This experimental geometry has the advantage of *many* modes in a compact volume, but also the additional advantage that those modes can be some of the highest-Q resonances measured ever at microwave frequencies [38]. This implements the scheme showed in the green dashed box of Figure 4.1. We submit that this is, in fact, a very valuable asset for future quantum memory applications.

Those two experiments are purely classical in nature – no quantum bit is involved – only mediated by quantum circuits that contain Josephson junctions. Thus in chapter 7 (red dashes), we use the WGM crystal to perform simple qubit manipulations and measurements. The goal of this chapter is to prove whether we need to combine the ideas of chapter 5 and chapter 6 – namely to separate the quantum part of the experiment spatially from the memory part of the experiment – or if the qubit can just simply be combined with the WGM and coupling element in one very compact volume.

The answer to that question seemed to be yes after performing those experiments, but in reality the combined system resulted in a ‘no-go’ proof-of-principle for the system – described in chapter 8. Further research needs to be performed and a second generation of devices will probably need to be fabricated to implement a full Von Neumann-like experimental device. The concluding chapter comments on what would be needed and wanted for that device to work fully.

4.2 Parametric Hamiltonians

One way to think about the physics describing the parametric frequency conversion process described in previous sections is to write the interaction as a ‘beam-splitter’ Hamiltonian:

$$H_{INT}(t) = \hbar(g_P(t)a^\dagger b + g_P(t)^* ab^\dagger), \quad (4.1)$$

which describes an interaction Hamiltonian governing the exchange of energy from mode a to mode b (or vice versa) at a rate set by g_P , the parametric coupling rate. We expressly write this parametric coupling term as a function of time. If the parametric coupling is turned on for a period of time such that it converts the incident signal with half its amplitude in mode a and half in b (or e.g. from cavity a to b), this is the equivalent to a 50:50 beam-splitter from optics. In optics, this is implemented via a physical piece of half-silvered mirror, but here we only need to turn on a microwave source for a period of time! In fact, our beam-splitter can be programmed to be a 50:50 beam-splitter or any other opacity *in situ*.

To be concrete, we can finally introduce the form of g_P and again explicitly note where the time dependence enters. The parametric coupling rate is:

$$g_P(t) \times e^{i\phi_P(t)} = \frac{\delta\Phi(t)}{4} \sqrt{\frac{\partial\omega_A}{\partial\Phi} \frac{\partial\omega_B}{\partial\Phi}} \times e^{i\phi_P(t)}. \quad (4.2)$$

$\delta\Phi(t)$ is the pump amplitude (which we can, again, turn on and off in time) which modulates our *parameter*, Φ . Our mode frequencies, ω_A and ω_B , must also have a dependence on the same parameter Φ . Another aspect of the parametric coupling rate is the phase term at the end – again, something not available to traditional optics experiments. This phase can also have a time dependence.

4.3 Quantum optics

Our group has recently demonstrated the usefulness and applicability of these parametric Hamiltonians in other systems with applications to amplification, squeezing, metrology,

and optomechanics [49, 50, 51, 52]. The analogy between optics and circuits can be furthered by noting that by combining all of these beam-splitters, amplifiers, etc we are building up an ‘optical table’ inside a fridge. These components can be sequenced with room temperature electronics and can yield on-demand Hamiltonians at base temperature. The experiments which follow, we hope, will demonstrate the power of this.

Building more interesting non-classical states into the cited works, above, will almost certainly necessitate incorporating a transmon qubit. This dissertation also is a stepping-stone toward that.

The ideas presented in this chapter represent a different perspective than was presented in Chapter 1. To conclude the chapter we highlight the two ways of thinking – each equally valid. Figure 4.2a describes the general ideas of that chapter. The tradeoff between coherence and control divides the field into two general approaches. Most research groups in the past could be generally categorized into either one. Some were focused on materials science and design to build the highest-coherence time qubits, whereas other groups focused on developing algorithms and control techniques.

If one considered approaching the problem as in Figure 4.2b, though, we can see how these two are related more closely. If we consider the general problem as ‘how do we couple quantum systems?’ – the same sub-fields arise more naturally. The highest coherence systems are well-isolated from the environment, and how one engineers the coupling between the constituents determines what algorithm is ‘computed’.

With the second point of view comes a vast area of research. How does one couple two quantum systems efficiently? What sorts of systems can be coupled? This approach should remind the reader again of motivations raised in Chapter 1 which highlighted the engineer-ability of the artificial atoms!

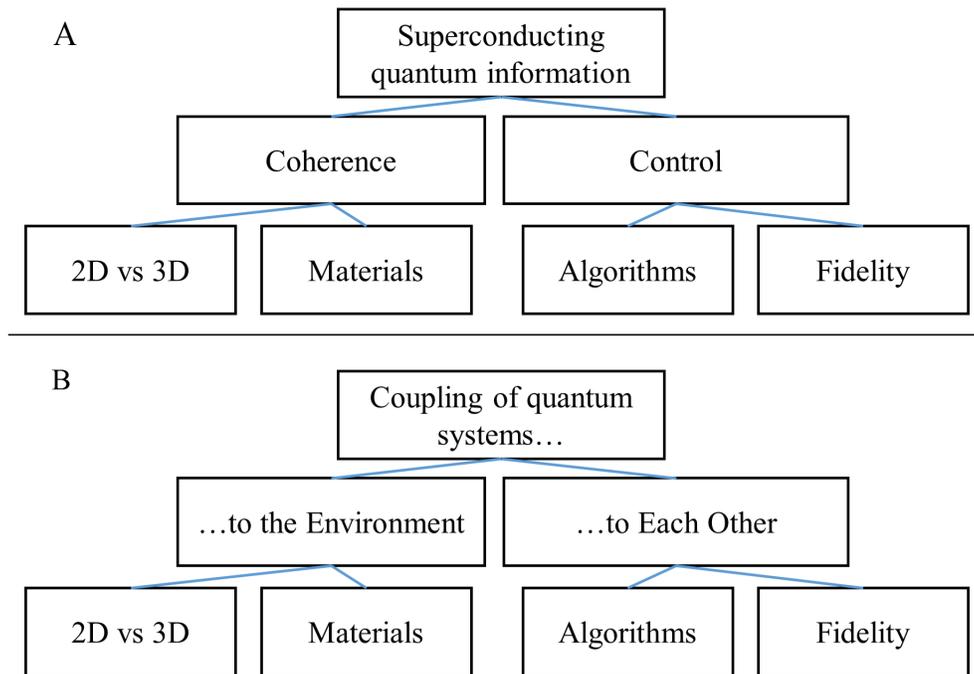


Figure 4.2: Two views of the ideas presented in this thesis as hierarchies. The first places quantum computing at the apex; the other considers the focus to be coupling between quantum systems. Each should be kept in mind reading the following experimental chapters.

Chapter 5

Experiment I: Frequency conversion between square cavities

This chapter is the first of four chapters that details experiments and shows data demonstrating the ideas discussed in Chapter 1 through Chapter 4. Indeed this is an extended version of the text as published in reference [53]. The first data comes from a device used to demonstrate frequency conversion of coherent states both in frequency and in space. This experiment is nominally inspired by recent developments in the field of superconducting quantum computing where metallic, 3D cavities have been shown to be a viable platform for housing qubits [2].

These rectangular cavities, as discussed earlier, can possess large internal quality factors, $Q_{int} > 10^{6+}$ and have been demonstrated to be useful for long-lived quantum memories which far exceed the coherence time of microfabricated circuits [27]. The increase of coherence time can, in part, be attributed to the restricted density of states that the qubits ‘sees’. In other words, by placing a qubit in a cavity with a few, high-Q modes surrounding it, the possible decay pathways of the qubit are limited. The high-Q modes themselves are useful and interesting also though. For example, a quantum state may be stored in these modes.

However, to be useful as a storage element, these cavities require a fast ‘read/write’ mechanism. In other words, with these high Qs comes long ring-up and ring-down times which, in turn sets how fast we can learn about the state stored in the cavity. Our proposed solution implements a tunable coupling between two cavities with very different Qs. This allows a fast read/write cavity (with low-Q) to be parametrically coupled — on-demand —

to the high-Q storage/memory cavity.

States (information) can be swapped back and forth between the two cavities by driving the coupler at the difference frequency of the two cavities. As a proof-of-principle demonstration of this technique we swap (classical) coherent states between the two cavities in this dual cavity system. These measurements proceed over the course of two cooldowns. The first aimed at characterizing the performance of the cavities and coupler chip, the second cooldown to demonstrate this ‘store and retrieve’ protocol mediated by the coupler chip.

These techniques have been used to convert single microwave photons between different frequencies of a coplanar waveguide resonator[47], and between a qubit and a linear mode [54, 55]. They have also been used to exchange coherent states between propagating modes at different frequencies [56], and between a cavity mode and the internal modes of a Josephson Ring Modulator [57]. Here we extend these results by demonstrating parametric frequency conversion of coherent states between two spatially-distinct, three-dimensional cavity resonators.

5.1 Device details and characterization

The device consists of two spatially-separated, approximately rectangular, cavities milled out of a split block of 99.995% pure aluminum. The purity of the aluminum determines the surface resistance of the metal, and therefore predicts a higher Q than less-pure aluminum. After machining, the cavities were immersed in heated (20 C) aluminum etchant (Transene Type A) for 30 min to smooth out machining scratches, and dissolve impurities in the aluminum surface. This results in a metal surface that is speckled with the large (mm-scale), smooth, grains of the bulk aluminum, indicating that we have indeed etched away impurities. We use this as an indication of when the etching process is ‘done’ [36].

The cavities are rounded at each end due to milling/machining constraints, but still support TE₁₀₁ modes. A photograph of half of the full device is shown in Figure 5.1a. The

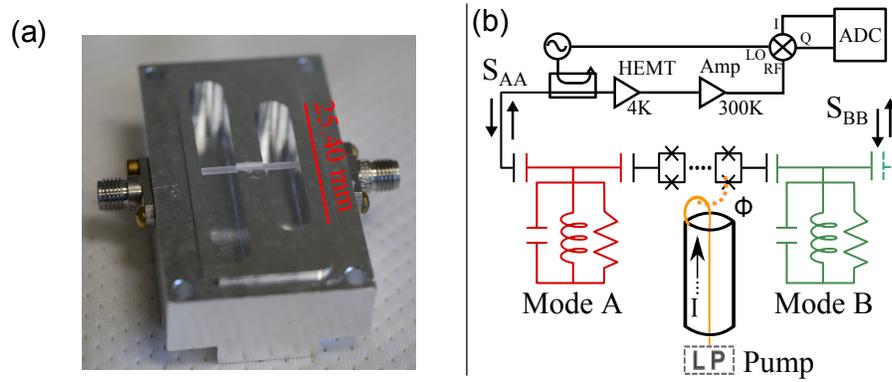


Figure 5.1: a.) A photograph of half of the split-block, double-cavity experiment. The coupling chip can be seen extending into both cavities. b.) A sketch of the measurement setup for initial characterization. The six-port microwave switch allows for probing through the A or B ports. Transmission can be measured through the coupler drive coax; and below a DC flux is swept to the coupler to characterize the participation ratio. Later, the B port is blocked off, and a low-pass filter added to the coupler pump line.

two halves of the split block are screwed together with an indium gasket to also help any stray fields from leaking out. Recent results from other groups also using this cavity geometry report that the seams are a significant source of loss [12]. The surface currents (and therefore surface loss mechanism) strongly ‘participate’ in that area of the cavity. This is another reason we decided to move away from the square cavity geometry for other experiments later.

The readout cavity (A-mode) is $6.35 \text{ mm} \times 19.05 \text{ mm} \times 38.10 \text{ mm}$. The fundamental mode (TE101) resonates at $\omega_A = 2\pi 8.07 \text{ GHz}$. The second cavity volume (B-mode) is $6.35 \text{ mm} \times 19.05 \text{ mm} \times 25.40 \text{ mm}$ and resonates at $\omega_B = 2\pi 9.33 \text{ GHz}$ (also TE101). A shallow slot is milled between the two cavities just deep enough to support the coupler chip and allow a flux pump coax – a connectorized, Niobium semi-rigid coax with a loop welded near the end – to allow DC and AC flux to thread the SQUIDs on the chip (with an external bias tee), see Figure 5.1b and Figure 3.6. This allows us to drive both AC and DC flux to the SQUIDs of the coupler. The mutual inductance between the flux pump line and each SQUID is $\simeq 0.05 \text{ pH}$. The slot is milled such that when placed in it, the coupler chip’s ‘antennae’ extend into the maximum of the electric field of the TE101 standing wave by about 1.5mm each. In this way, the coupler chip provides a modulate-able reactance between the two electromagnetic fields of the two volumes.

In order to characterize the participation ratio of the coupler with the cavity volumes, an initial cooldown was performed where this device was probed and excited through three ports — one port for each cavity volume, and one for the coupler drive. The two cavity ports were connected to the six-port switch – allowing them to be measured individually.

The coupling to the outside world for the cavities is through a simple probe wire (high-purity copper) pushed into an SMA connector. This simple connector allows for quick turn-around testing of various lengths of probe wire length and since we’d eventually like to have one cavity isolated, easy removal when not wanted. With both cavity probes in place

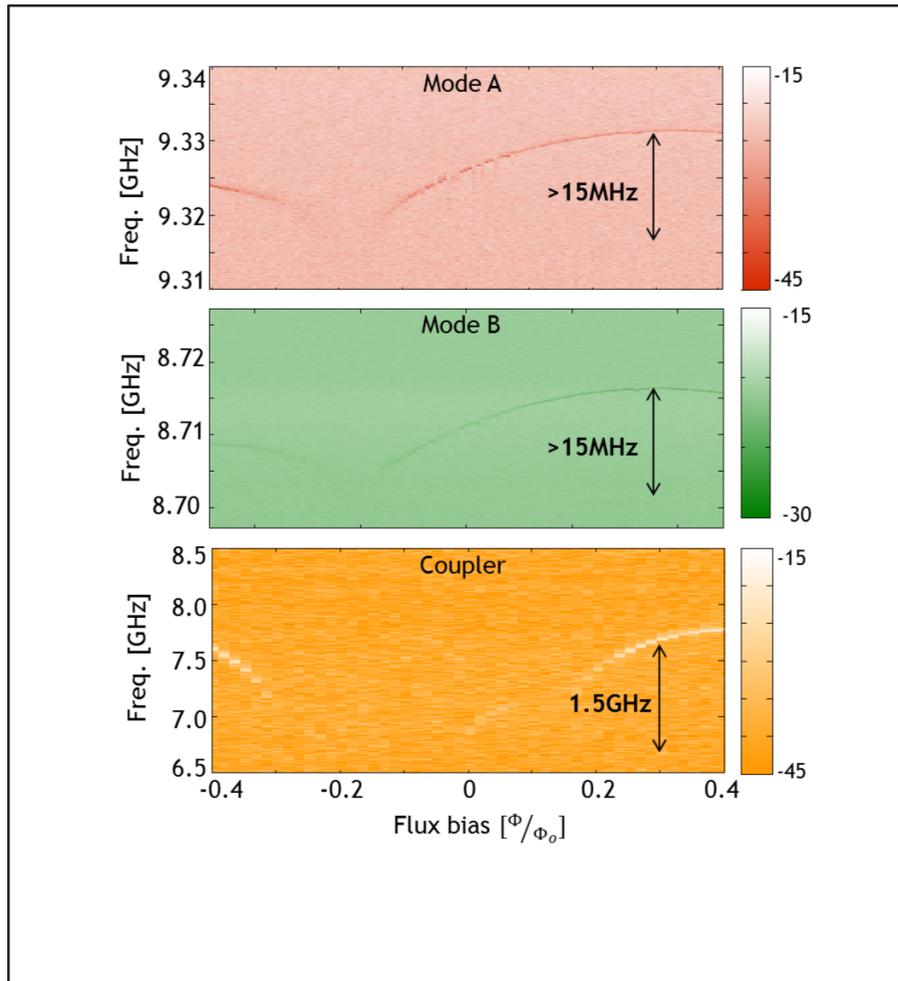


Figure 5.2: By sweeping the DC bias to the pump coax, the SQUID-coupler participates in each cavity mode, and therefore loads the cavity frequency with its Josephson inductance – lowering its resonant frequency. a and b.) show reflection measurements off of cavity A and B (selectable via switch). c.) shows transmission from the pump port through cavity A, and demonstrates that the coupler’s resonance can also be read out and characterized.

we can fully characterize the transmission and reflection from all **three** ports. Then, when it comes time to have this cavity be high-Q we remove one connector to isolate one cavity as much as possible from the external world.

The characterization data is shown in figure 5.2 – we measure reflection off of each cavity volume while changing the DC bias to the coupler chip (a and b), and transmission from the coupler port through the cavity volume (c) via a pump line, then out the cavity. As the DC bias to the coupler port is swept, the resonance frequency of the cavity mode modulates downward. This is due to the effective reactance of the coupler chip participating more and more strongly in the cavity resonator-coupler system combined impedance.

This is confirmed by measuring the transmission from the coupler port to a cavity port which allows us to ring up the coupler also. Figure 5.2c shows that we are also able to read out the self-resonance of the coupler and note that its modulation matches that of the mode volume’s resonance modulation. The self-resonance of the coupler chip is due to the flux-tunable SQUID inductance, shunt capacitance to the aluminum box, and distributed capacitance and inductance from the microstrip antenna – all modeled and simulated via Microwave Office and HFSS. In the end, the coupler’s resonance has a maximum natural frequency of $\omega_C \simeq 2\pi 7.7$ GHz. We should also note that the coupler’s self-resonance is below that of both of the cavity’s, this means that the reactance ‘seen’ by each cavity is capacitive.

From this data, we can estimate the parametric coupling rate amplitude. As mentioned in Chapter 4, the parametrically induced coupling rate depends on the slope of the frequency vs flux curve, as well as the amplitude of the pump:

$$g_P = \frac{\delta\Phi}{4} \sqrt{\frac{\partial\omega_A}{\partial\Phi} \frac{\partial\omega_B}{\partial\Phi}}, \quad (5.1)$$

where $\delta\Phi$ is the amplitude of the pump flux modulation and the partial derivatives are the first order terms in a Taylor expansion of the flux modulation curves shown in Figure 5.2a-b expanded around a given DC flux bias point.

With these initial spectroscopic characterization measurements finished we then warmed up the device, changed the external couplings to the 50 Ohm environment, and cooled it back down to perform the intended experiment described above with one low-Q and one high-Q cavity. To that end, the SMA connector on cavity B was removed (and the hole blocked with superconducting indium foil) to isolate that cavity as fully as possible. Additionally, the measured transmission from the flux pump line to the cavity (A or B) shown in Figure 5.2c is an undesirable source of dissipation — i.e. energy can leak out from the cavity through the pump line. This was mitigated by adding a commercial low-pass filter in-line with the pump coax. The low-pass cutoff of the filter was chosen to be, $f_{LPF} \simeq 1$ GHz. This frequency band allows the intended pump frequencies ($|\omega_A - \omega_B| \approx 630$ MHz) to pass with low impedance, while frequencies above that are reflected back to the device (including frequencies at ω_A and ω_B).

5.2 Frequency conversion between square cavities

With the SMA connector removed from cavity B and the pump line filtering in place we measured the internal quality factor as $Q_A^{int} \simeq 900,000$, while the external quality factor is $Q_A^{ext} \simeq 50,000$, set by the length of the microwave connector pin extending into the cavity volume. The corresponding total dissipation for this mode yields $\kappa \simeq 2\pi 170$ kHz and $T_1^A \simeq 0.94$ μ s. This provides us with a reasonable time to do our read/write operations and parametric frequency conversion pulses.

As in Chapter 2, equation 2.96, the dynamics of the coupled mode system are encapsulated by the equations of motion

$$\frac{da(t)}{dt} = -i(\omega_A - i\frac{\kappa_A}{2})a(t) - ig_P(t)e^{i(\omega_P t + \phi_P)}b(t) \quad (5.2)$$

$$\frac{db(t)}{dt} = -i(\omega_B - i\frac{\kappa_B}{2})b(t) - ig_P(t)e^{-i(\omega_P t + \phi_P)}a(t) \quad (5.3)$$

where $a(t)$ and $b(t)$ are the complex mode amplitudes of the electromagnetic field in the readout (“A”) and storage (“B”) cavities (normalized to the square root of the photon

energies, $\sqrt{\hbar\omega_{A/B}}$, $\omega_{A/B}$ are the natural frequencies of the readout and storage cavities, and $\kappa_{A/B}$ are the corresponding total dissipation rates. In the highly off-resonant regime where the mode separation $|\omega_A - \omega_B|$ is much greater than the static coupling and dissipation rates, energy cannot be exchanged between modes residing within these cavities. However, if the coupling term is modulated at the difference frequency, $\omega_P = |\omega_A - \omega_B|$, such modes can resonantly couple in a mean rotating frame, transferring energy in both directions. This process, parametric frequency conversion, can be turned on and off by modulating the coupling envelope amplitude $g_P(t)$. In addition, it provides an extra degree of control over the swap process through the pump phase ϕ_P . We show each of these data in turn below.

The first measurements shown, in Figure 5.3a, are in the continuous-wave domain and were thus performed using a VNA to measure the reflected signal from cavity A while driving the pump line with signals near the difference frequency all with a static DC bias to the coupler. This sets the static, DC, slope (the $\frac{\partial\omega_{A/B}}{\partial\Phi}$ terms in equation 5.1) of the modulation curves which we can then use to ‘turn on’ parametric coupling by increasing the drive strength (the $\delta\Phi$ term in equation 5.1). The observed splitting in the reflected S_{11} measurements indicates energy exchange with a mode at $\omega_A + \omega_P = 2\pi 9.33$ GHz, as expected. With pump modulation of $\Delta\Phi \simeq 0.2\Phi_o$ (pump powers of $P_P = -52$ dBm at the device) the splitting size is $2g_P = 2\pi 2.4$ MHz, consistent with equation 5.1. By pumping with a higher amplitude (higher pump power), the splitting size is able to be increased to a coupling rate of $2g_P \simeq 2\pi 3.5$ MHz before the system begins to heat from excess pump power dissipation at the mixing chamber cooling stage.

These continuous wave measurements show that strong coupling (where $g_P \geq \kappa$) is possible between our two cavity volumes. To demonstrate this in the time domain we inject an on-resonance microwave pulse which populates a coherent state into mode A (such that the photon number population is $\bar{n}_A = 10$). After the population pulse ends, we then pulse the pump drive for a fixed length of time, τ_{fc} , swapping some fraction of the energy between

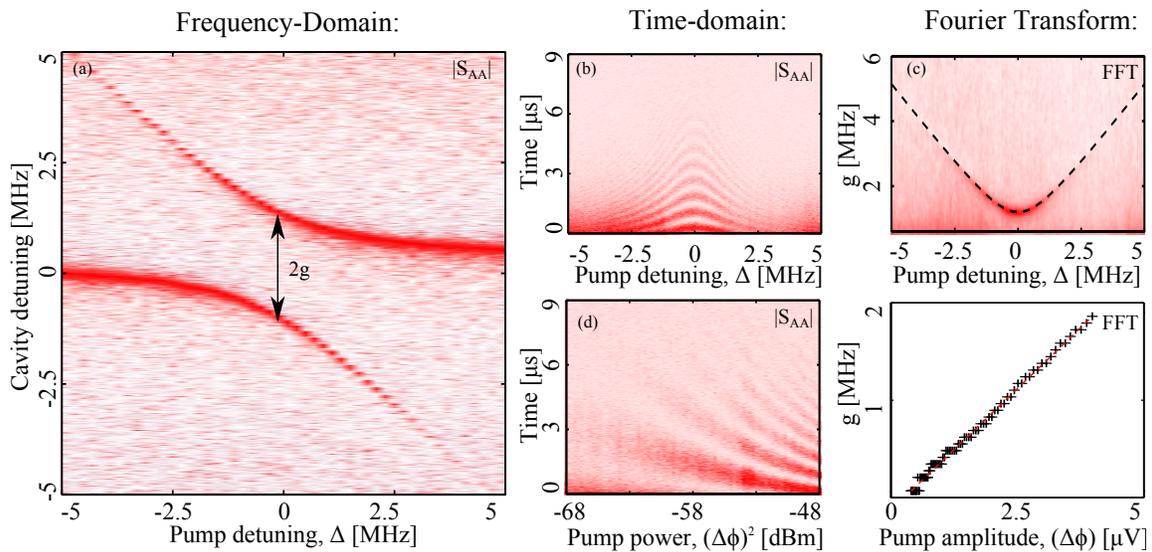


Figure 5.3: a.) The continuous-wave signature avoided-crossing behavior of coupled modes. b.) In the time-domain this periodic swapping is seen clearly as the pump detuning is swept through the optimal $|\omega_A - \omega_B|$ difference frequency. c-d.) By changing the pump power, we can confirm the expected $\sqrt{\text{power}} \propto \Delta\Phi$ dependence from equation 5.1. The Fourier transforms of (b) and (d), shown in (c) and (e) show the range over which the frequency conversion operates.

A and B , and then integrate the power that leaked from the A port into our amplification chain. Figure 5.3b shows the (uncalibrated) leaked energy from A as a function of both pump drive detuning and length of the pump envelope. These data show Rabi-like oscillations as a function of swap time and pump frequency detuning as expected for coupled classical linear oscillators [58, 59]. In Figure 5.3c we show the Fourier transform of these oscillations demonstrating a $\sqrt{\Delta^2 + g_P^2}$ detuning dependence [47]. In a similar manner we fix the pump frequency ($\Delta = 0$) and measure the pump power dependence of these swap oscillations (figure 5.3d). The resulting swapping rate, g_P , follows the expected linear dependence on pump amplitude up to a few megahertz (Figure 5.3e).

The swap data shown in Figure 5.3b provides the necessary pulse-length calibration for a full state swap, analogous to the π -pulse calibration one would do when driving to the excited state in a qubit. Fixing the pump frequency, power, and envelope time accordingly ($\omega_P = 2\pi 632.5$ MHz; $P_P = -52$ dBm, $t_{swap} = 0.6$ μ s), we can perform a storage and retrieval measurement sequence. In this sequence, we inject a coherent state into readout mode A , frequency convert/swap it into storage mode B , wait a delay time t_{delay} , then swap it back to A where it leaks out rapidly into the amplification chain. Typical time domain data for the leaked power (square of the sum of in-phase (I) and quadrature (Q) components of the field) are shown in Figure 5.3a. Again, we integrate these traces to provide a measure of the total energy retrieved. Plotting this as a function of delay time yields a B -mode decay time $T_1^B = 14.9$ μ s (Figure 5.3b). This time far exceeds the measured readout mode decay time, T_1^A noted earlier.

The integrated energy retrieved (after the shortest storage time) can be compared to the energy stored in cavity A with the swap sequence disabled to derive an efficiency measure, η , of the swap process – this yields $\eta = 74.7\%$. We note that this efficiency is dominated by the decay of cavity A during the swap pulses. If we account for the energy loss from the decay of cavity A during the swap, the swapping efficiency increases to $\eta' = 99.6\%$,

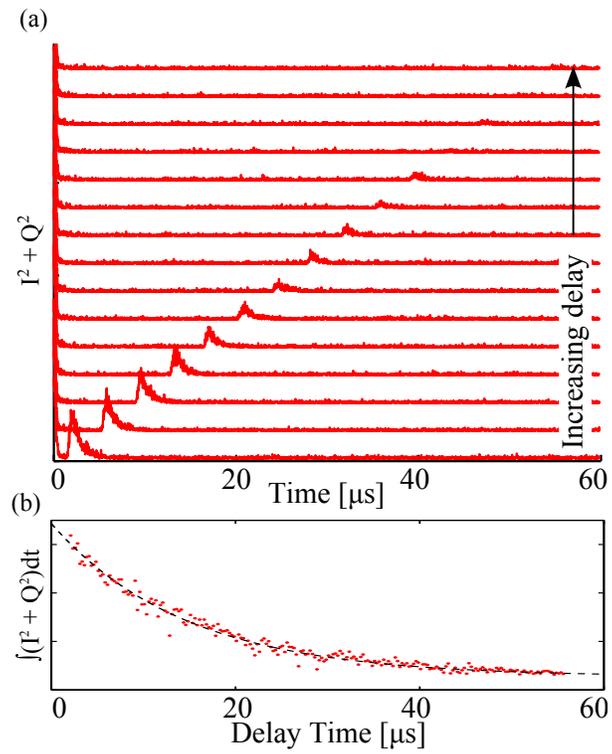


Figure 5.4: a.) The state remaining in cavity A after loading a coherent state in A, swapping to B, storing the state, and swapping back to A for various storage times (offset vertically). b.) By integrating the pulse area under each returned pulse, a measure of the decay time of cavity B can be plotted.

indicating that the swap process itself does not add significant loss.

We complete our characterization of the parametric swap process by demonstrating control over the coupling phase. As before, we perform a storage and retrieval sequence (see Figure 5.5(a-e)), but fix the pump power, duration, and delay time, changing only the pump phase on the retrieval pulse. We can clearly see that the pump phase is imprinted on the resulting retrieved state that leaked from the A-mode into the amplification chain, Figure 5.5f. In Figure 5.5(g) we show the phase angle of the retrieved signals as a function of relative pump phase. This ability to easily control the phase in the swap process using vector modulation sets parametric frequency conversion apart from resonant coupling and opens other possibilities for manipulating the state of dual-mode systems.

To expand on this point further: in, for example, the Jaynes-Cummings model to acquire a phase the qubit is moved off-resonance with the cavity, a phase is acquired by waiting (in time), then the qubit is moved back on-resonance. These three steps can introduce errors and complications not present in the parametric case.

5.3 Chapter conclusion

This was the first demonstration in our lab of the usefulness and reconfigure-ability of 3D cavities other than just measuring different mode resonances and comparing geometries. Having the ability to change couplings or add filtering easily without having to do any micro-fabrication was a big change from the all-microfabrication modality we had been operating under before. In my mind, this opens up a broad area of experiments where we can take just about *any* cavity or system with an accessible electric field and use simple coupler or qubit chips to couple to them. In fact, this sort of ‘chip stacking’ idea could provide an interesting avenue for further study.

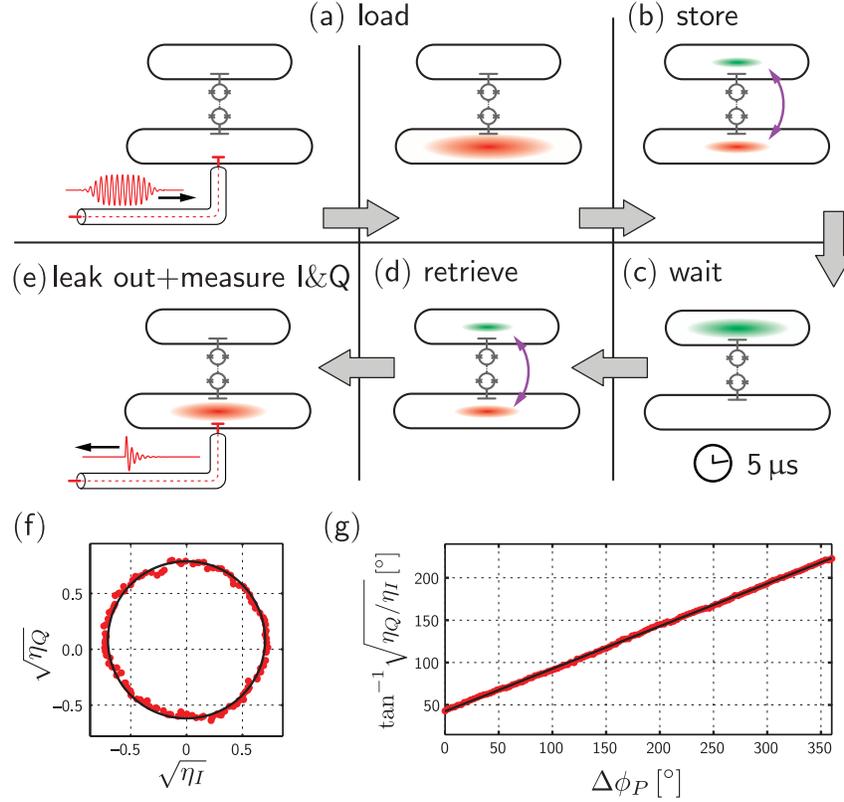


Figure 5.5: State initialization and measurement sequence. (a) a coherent state is loaded into cavity A with a $20 \mu\text{s}$ pulse; (b) the coherent state is immediately swapped to cavity B with a $0.6\mu\text{s}$ pulse applied to the pump at the difference frequency, ω_p ; (c) the coherent state is stored in mode B for $5 \mu\text{s}$; (d) the same pulse as in (b) is applied again, but with a phase shift $\Delta\phi_P$ relative to the first swap-pulse; (e) the final state in mode A is then read out and demodulated. (f) Efficiency of the I and Q integrated signals over a 2π range of relative phase shifts. The corresponding data (red dots) and its fit (solid black) are scaled such that they encompass an area $= \pi\eta$, where $\eta = 0.49$, the state retrieval efficiency at $5 \mu\text{s}$ delay. (g) Phase of the retrieved integrated signal **vs.** relative phase shift between the storage and retrieval frequency conversion drives. The fit slope (solid black) is 1.

Chapter 6

Experiment II: Freq. conversion between sapphire WGMs

The device in Chapter 5 was designed from the beginning to limit our working density of states to just two modes – the TE₁₀₁ modes supported by each of the two cavity volumes. Other modes certainly existed, but either did not couple to the coupler’s antennae, or coupled very weakly. For example, the TE₁₀₂ mode (next higher resonance frequency) has a node of the electric field precisely where the coupler antenna sits. The next highest mode, TE₁₀₃, has an anti-node, but the field strength is roughly (1/3) the strength as the TE₁₀₁ mode at that location and resonates at higher frequency.

If we take the same coupler chip from the experiment in chapter 5 and place it on top of a sapphire whispering gallery mode resonator we can reproduce the same experiment as detailed in that chapter. In this case, though, we have now expanded our density of states in which we are allowed to frequency convert between. In the frequency range of 7 and 12 GHz, there are an average of about 20 modes per GHz. Though, not all of them are whispering gallery modes, and not all are high-Q (though almost all have linewidths below a few MHz).

We do this with a caveat that many of the assumptions used to derive the parametric frequency conversion (PFC) model, as well as the rotating wave approximation *may* not fully apply in all cases any more. First, the model for PFC assumed that there was a small, static coupling between modes. Since the mode density per bandwidth is much higher in this case static couplings may play a role. Additionally, the rotating wave approximation assumes that only terms proportional to $\omega_A - \omega_B$ are involved in determining the dynamics of the

couplings. In our case, since there are so many modes over such a wide range in frequencies, terms say, with $\omega_A + \omega_B$ (and other multiples) could also play a role. With that said, with so many modes available to us, it is likely that we can just ignore frequencies which do not suit our needs.

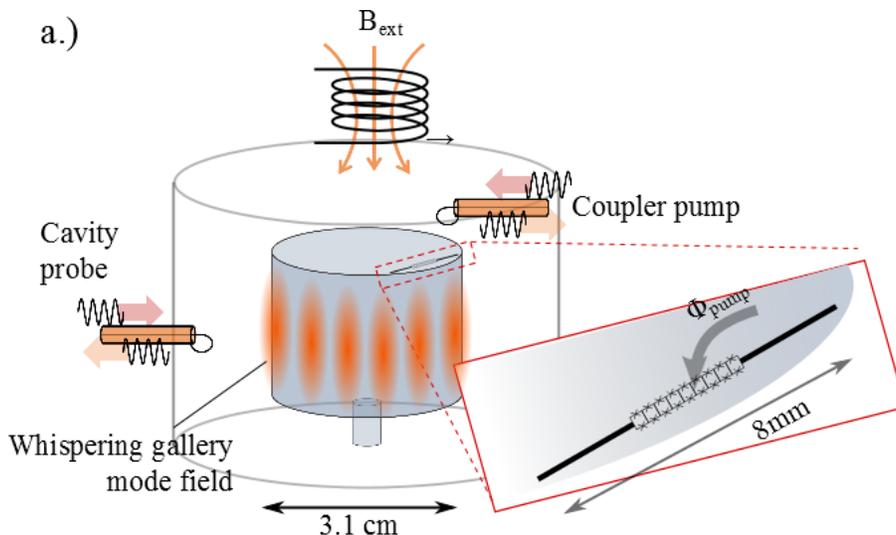


Figure 6.1: A cartoon of the measurement setup for this experiment and the whispering gallery mode standing wave fields.

In this chapter we again perform mode characterization experiments to show how the coupler chip modulates the mode frequencies, then use the same parametric frequency conversion protocol to perform a ‘store and retrieve’ experiment with coherent states containing approximately ten photons.

6.1 Device details

By using the fields of the sapphire whispering gallery mode resonator, there exist *many* modes in a volume similar to the previous device which effectively only had two modes to work with. The whispering gallery mode resonator’s fields are concentrated around the circumference of the crystal (recall Chapter 2). Thus the coupler chip is placed very nearly tangent to the perimeter. This allows the coupler circuit’s dipole moment to align most

strongly with the electric fields in the crystal pointing in the E_ϕ direction. The angle of this placement was tested experimentally by changing the coupler angle from radial, to a 45 degree angle between radial and azimuthal, to azimuthal. This confirmed that the strongest effect was seen in the azimuthal-orientation. Additionally, to maximize the dipole coupling of the chip to the sapphire fields, the chip is placed ‘face-down’, with its metalized surface contacting the surface of the WGM crystal.

Maximizing the coupling to the sapphire electric fields is necessary because as the mode number increases – i.e. as frequency increases or wavelength decreases – the fields are more and more confined in the sapphire volume. This is due to the shallower angle at which the standing wave electric field impinges on the sapphire/vacuum interface and the field is more efficiently internally reflected back into the sapphire. Higher order modes leak out into the vacuum to a lesser extent than lower order modes. Thus the small amount of electric field, which the coupler’s dipole couples to become increasingly smaller. Having the mode *more* confined in the sapphire volume, though, results in higher Qs. This is analogous to having an optical cavity with mirrors whose reflection becomes better with frequency.

So, there are two effects – the electric field that the coupler couples to is weaker, and the Q of the resonator becomes higher – the highest Q modes are generally tenth to fourteenth order modes with frequencies $\sim 12.5+$ GHz. The reason for this is a current topic of research at the University of Western Australia and is most likely due to iron-group impurities in the sapphire at the part-per-billion level which resonate near 11 – 12GHz [60]. These impurities act as a low-Q bath which washes out any resonances near them. A possible way to mitigate these effects is to use larger WGM crystal diameters. This allows the 10th or 11th order WGMs with the highest Qs to move to lower frequencies, and thereby avoid the detrimental resonances. This would be attractive for future devices since that would move the highest-Q frequencies closer to the optimal bandwidth of our amplifier/electronics chain. The longer wavelength standing wave might also allow for less stringent tolerances for placing of chips

on the sapphire surface.

The sapphire is housed in a silver-plated copper cylindrical enclosure to allow for well-controlled isolation from the outside world and a low-loss environment in the fridge. The housing for the WGM was designed and tested at the University of Western Australia to make sure to thermalize the sapphire crystal via the post at the bottom of the spindle-shaped crystal (shown as the copper colored piece in Figure 6.2). This should also minimize clamping losses through the post as no fields exist in that region.

The resonance frequencies are probed and read out through the loop probe indicated in Figure 6.2. The sapphire resonator came from the Perth with the external drive port's coax coupling loops set to already determined values. The external Q to the sapphire modes was set to be critically coupled or less for almost all of the whispering gallery modes such that their internal Q can be measured accurately and not de- Q them with an external load. Some of the lower- Q 'box' modes – modes that are formed by the silver cylinder's boundary conditions – are overcoupled though. This gives us a wide range of fast and slow modes which we can experimentally determine if they couple parametrically to each other. An annotated photograph is shown in figure 6.2. The diameter of the sapphire is 3.158 cm, which is the same dimension as the height.

We predict that the WGMs will all couple to one another via the coupler chip since the wavelength of the WGM resonances is about the same length as the 8mm dipole antenna length of the chip. It is much harder to predict where the box mode field will overlap with the WGM fields or with the coupler's dipole. These box mode's orientation in the box/housing can differ significantly due to slight imperfections or asymmetries. Rather than fully calculating/simulating these complications by modeling them with HFSS, it was deemed easier to experimentally explore them.

Though not shown in Figure 6.2, a custom, high-purity copper lid completes the enclosure and allows for an additional probe to extend into the cavity from above to provide

DC and AC pump tones near the coupler circuit placement – sketched in Figure 6.1. This coaxial loop probe is centered above the coupler chip’s SQUID array and couples to the SQUIDs with a mutual inductance of about $M \simeq 0.15pH$. In a similar manner to the device in chapter 5, this probe provides an AC pump tone to the coupler. A low pass filter was not needed to inhibit transmission from the coupler bias out the cavity probe due to the structure of the WGMs involved. Specifically, by exciting WGM resonances, the fields above the sapphire crystal are exponentially suppressed in the z-direction – again the maximum of the electric field is azimuthal. The transmission from the coupler’s coax probe out the WGM’s coax probe was isolated by at least $40dB$, and in general it required at least $20dB$ more power to excite a whispering gallery resonance through the coupler coax than through the WGM coax.

Aligning the pump coax was a bit of an art form rather than an easily predicted/modeled coupling. A future design could easily overcome this aspect of the experiment by making the coupler chip integral to a larger chip that covers the top side of the WGM crystal. Additionally, the newest designs from UWA show that larger WGM crystals (that, again, also have higher Q at lower frequencies) do not need to have the top spindle post like our’s — more ‘mushroom’ shaped than spindle [38]. This geometry would allow a large chip containing bias circuitry, a coupler, and other circuits to be placed on the top surface of the crystal. In that case, all of the couplings between the coupler and pump/bias could be determined solely by reliable lithography.

To provide the static DC bias field needed to choose the slope of the modulation curves, a 2,000 turn, 1.2 mm radius solenoid coil was screwed to the copper lid outside the resonance volume. This coil, wound with superconducting niobium-titanium wire, provides a very stable DC field over a large area of the crystal which relaxes the constraint of aligning the coupler SQUIDs directly under the coax loop.

After a careful survey of resonances over the full range of measurement bandwidth –

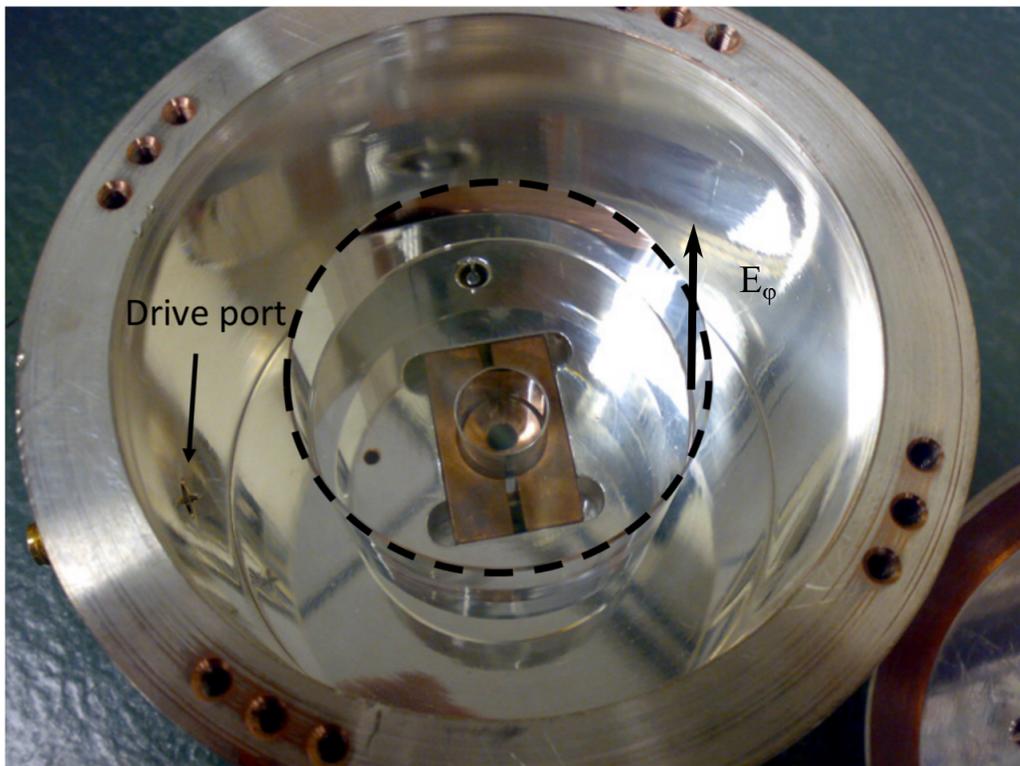


Figure 6.2: A photograph of a whispering gallery mode sapphire crystal. The dashed outline is for clarity. The coaxial, loop drive port is noted on the left, and the alignment of the electric field (along the circumference) in the ϕ -direction determines where the coupler chip will be placed.

$\sim 1 - 14$ GHz, a list of all of the modes which changed with DC flux bias to the coupler was collected. There were well over 40 modes, but in this experiment we focus on the top ten highest Q modes as our storage modes, then chose a moderately high Q mode as our read/write mode. An additional complication occurs in that the two modes can now have different participation ratios and therefore different $\partial\omega/\partial\Phi$ terms in equation 5.1. We now need to write that equation as:

$$g_P = \frac{\delta\Phi}{4} \sqrt{\frac{\partial\omega_{R/W}}{\partial\Phi_{R/W}} \frac{\partial\omega_n}{\partial\Phi_n}}, \quad (6.1)$$

where the subscript on the $\partial\Phi_n$ terms denotes the fact that the two mode frequencies do not modulate the same amount, and ‘R/W’ indicates the read/write mode. In the data below, the read/write mode was chosen to be a low-Q mode at 10.104 GHz. So all of the difference frequencies are relative to that – i.e. $\omega_P = |\omega_{R/W} - \omega_n|$. Similar to the experiment in Chapter 5, this low-Q mode again has a decay time of $T_{R/W} \simeq 1\mu s$. Looking ahead, $\kappa = 1/1\mu s$ is narrow compared to the expected linewidths of the sapphire modes, and $1\mu s$ is close to the predicted coherence time of qubits used by our group (typically $1 - 5\mu s$). This allows us enough time to perform parametric frequency conversion processes which are typically completed in $200ns = (5MHz)^{-1}$.

6.2 Device characterization

The measurements in this chapter differ very little from the previous measurements. The first measurements are continuous-wave measurements of the reflected microwave signal off of various modes while changing the DC flux bias to the coupler. In this manner we can determine how strongly the coupler’s tunable reactance (via the tunable SQUID reactance) participates in the various mode resonances. From these measurements we can also determine which modes are acceptable for use as storage modes and which for read/write modes.

Some example data is shown in Figure 6.3. The horizontal axis shows the change in coupler flux bias through a full flux quantum, and the vertical axis represents VNA

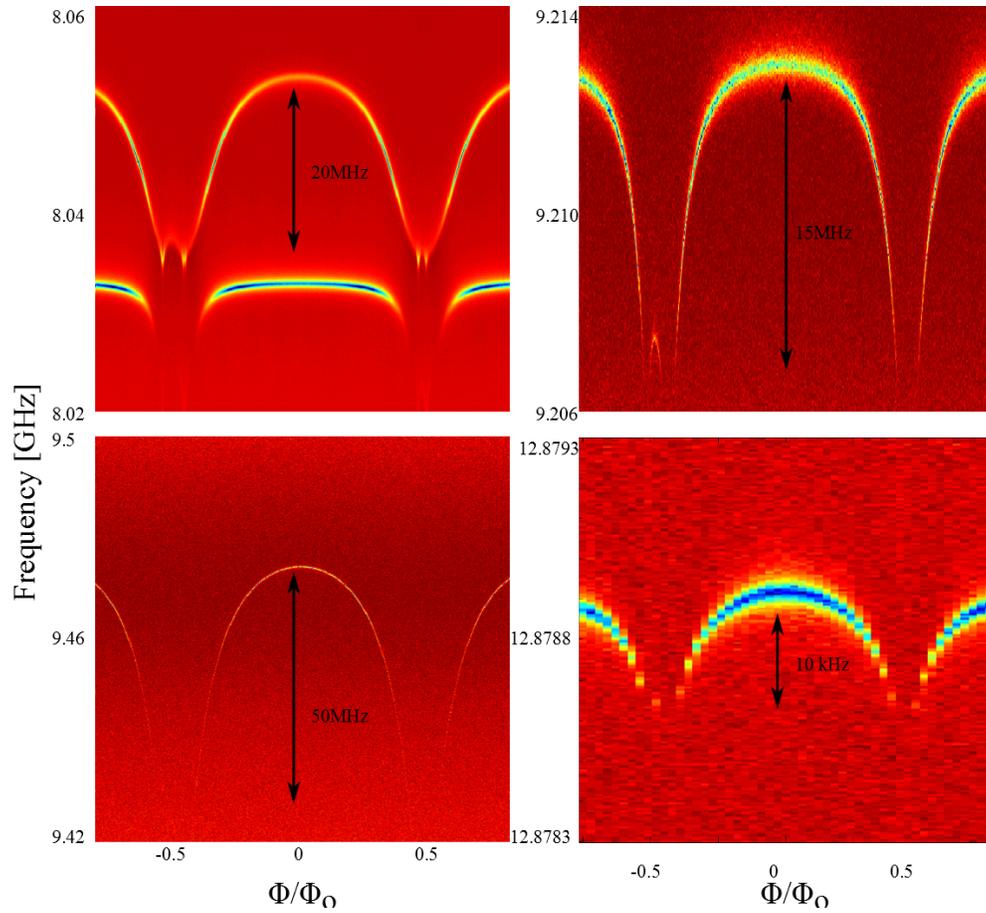


Figure 6.3: Four sample WGM crystal resonances which change with coupler flux bias measured with a VNA while changing the external flux bias field. The bottom right panel shows that even very high Q modes move with flux bias more than a linewidth.

scans (different scales on each y-axis). We note that often some of the sapphire modes will even modulate enough to cross and split with modes that do not modulate with flux. This ‘vacuum’ splitting may be interesting to explore for later experiments, and gives us an accurate measure of the static couplings between modes. The static coupling between modes is typically 1 – 5 MHz. But, by choosing to operate at DC fluxes which avoid these crossings or using frequencies with no crossings this can be avoided.

As shown in the example modes of Figure 6.3, many resonance frequencies are affected by the coupler chip, but by different amounts. This variation in modulation of the mode frequency is due to a few factors. First, the electric field which extends outside the sapphire crystal, which in turn induces current to flow along the coupler chip, is different for each mode. Second, the actual mode inductance differs from mode to mode depending on how much of the fields are confined in the vacuum or in the sapphire. A third factor that determines how much a mode modulates is the alignment of the whispering gallery mode. The amount of current induced along the coupler antenna is proportional to $\vec{E} \cdot \vec{d}$ – thus misalignment between the electric field and dipole moment, or misalignment between the maximum of the electric field strength affects the participation. If the standing wave is not aligned along the coupler’s antenna, or the coupler is not placed at a maximum, or if the coupler’s antenna is longer than a wavelength (which might be true for higher frequency modes) – the coupler will not affect the mode as much as desired.

Clearly, we have *many* modes that modulate with coupler bias now. But we must keep in mind the goal of our experiment. Modes that have a low Q (linewidths ~ 1 MHz) and modulate by many MHz are good for read/write modes. High-Q modes, on the other hand, don’t need to modulate too much to be useful for use as storage modes. Thus, it is just a matter of making a list of the highest Q modes, and checking that they modulate even a small amount – see Table 6.1.

Frequency (GHz)	Linewidth (kHz)	Lifetime (μs)
10.103 (Read/Write)	150	0.985
7.9040	8.4	18
10.0925	8	20
10.9835	8	20
12.1571	10	14
12.3225	4	39
12.3234	8	19
12.5050	10	14
12.8650	1.5	105
12.8785	0.326	490
13.2515	4	40

Table 6.1: List of high-Q WGM frequencies with their corresponding linewidths and decay time constants. We also include a low-Q ‘read/write’ mode for reference.

6.3 Frequency conversion between modes

Once we have the list of high-Q storage modes, the next measurements are continuous-wave measurements of the reflected microwave signal as we sweep a second tone through various frequencies which correspond to the difference frequencies between those modes and the designated read/write mode: these mode-splittings are shown in Figure 6.4.

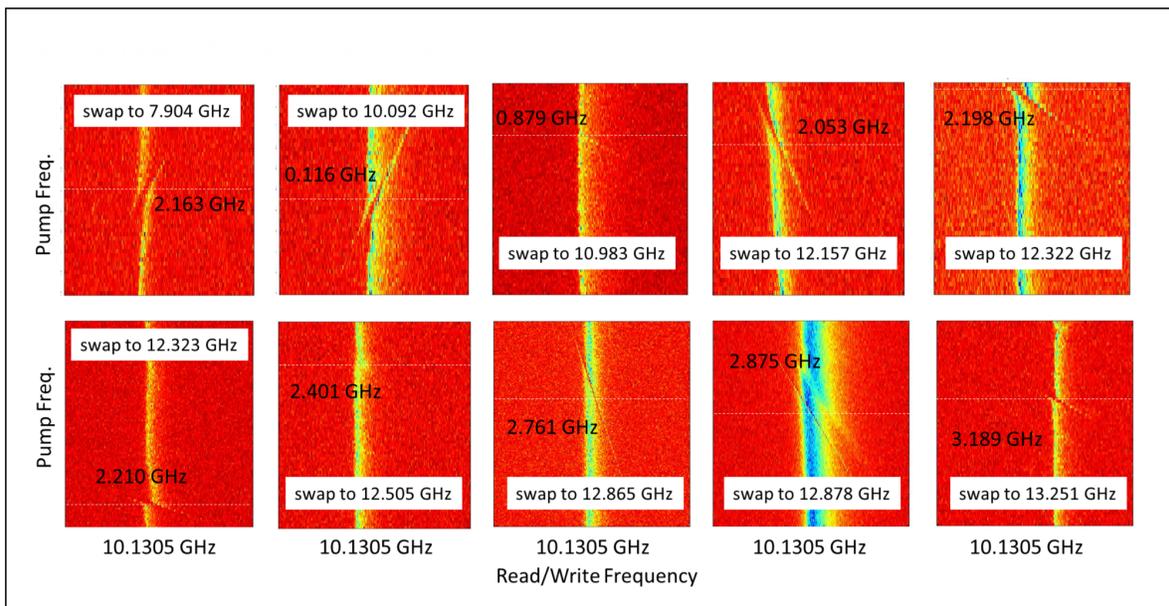


Figure 6.4: Parametric frequency conversion between the a single read/write mode and the 10 highest-Q modes that also modulate with coupler flux bias. Compare to Figure 5.3a.

The avoided crossings in each panel of Figure 6.4 show that indeed the coupler chip is acting in a very similar way as it did when placed between the two square cavities. Unfortunately, the splitting sizes are often less than $\sim 1\text{MHz}$ ($= 1/1\mu\text{s}$) and therefore the read/write mode is just barely fast enough to efficiently load, swap, then retrieve. These scans were performed at a moderate level of pump power though, so this rate can almost certainly be increased.

We can again implement the time-domain ‘store and retrieve’ experimental protocol: loading a coherent state into the read/write mode, then pulsing on the coupler tone, holding

the state in a high-Q mode, then later turning the coupler back on to retrieve the state. In this case we focus on three of the modes which showed the largest CW-splittings – i.e. the modes with the fastest induced parametric splitting rates. For this reason the data shown in Figure 6.5 are for swapping from the 10.104 GHz mode to the 7.904, 10.092, and 12.322 GHz modes.

In Figure 6.5, the returning digitized time traces after this protocol are integrated (cf: Figure 5.4) and plotted for various delays between the store and retrieve pulses. Again, we have shown that we can store coherent states for much longer than the limits of our read/write mode’s inverse-linewidth efficiently. In order to perform these measurements for other modes a stronger pump power or coupler with higher critical current would be desired.

In Figure 6.6, we also show parametrically induced anti-crossings between two, neighboring, high-Q modes. Though not a focus of this dissertation, the connections between high-Q modes will also be an important tool for quantum computing – i.e. not only will a quantum computer need to store information in memory modes, it may also need to manipulate the stored states or shift them to other frequencies. In this case the two modes are only 0.65 MHz apart, but by pumping the difference frequency, we are able to induce parametric coupling.

The two resonances featured in Figure 6.6 can tell us a lot about how the coupler chip interacts with the electric field of the sapphire. These two modes form a doublet – the sine and cosine solutions from equation 2.39, and are a pair of whispering gallery modes with a high mode number (and therefore a short wavelength compared to the coupler length). The two modes split by a similar amount which suggests that they both couple to the coupler chip in a similar way. For these modes near 13 GHz, that correspond to the 13th harmonic, the wavelength of the standing wave is ~ 4 mm. Our coupler dipole is 8 mm long though and extends more than a full wavelength. Indeed even though the sine and cosine electric fields are orthogonal, the coupler interacts and couples to *both*. One theory as to why this

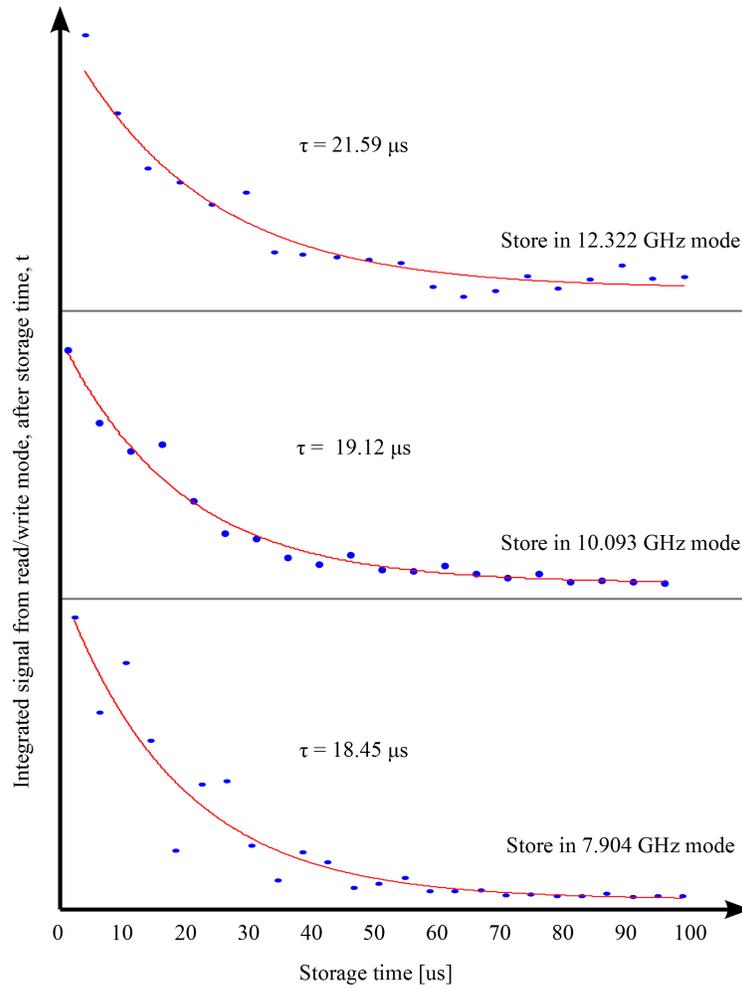


Figure 6.5: We implement the storage and retrieval protocol between the ‘read/write’ mode and three high-Q storage modes.

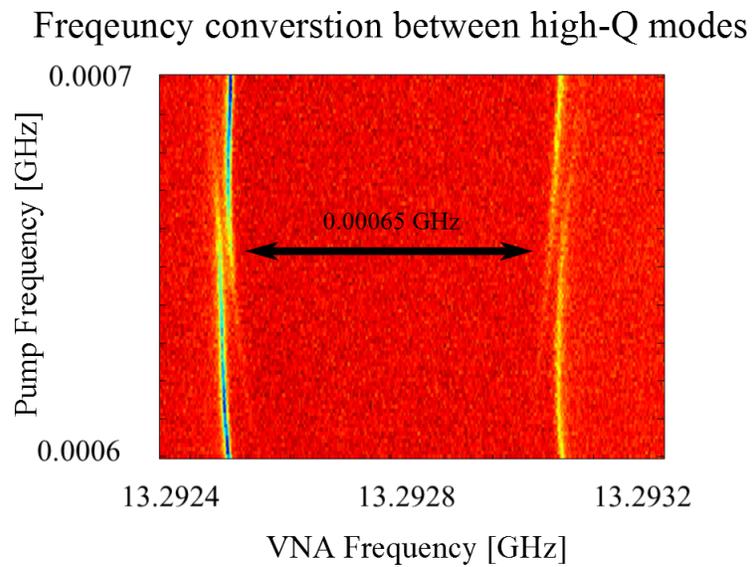


Figure 6.6: Parametric-induced anti-crossing between two nearby, high-Q WGM resonances. The ability to swap between high-Q modes allows for parametric interactions between the memory modes, not just between the memory modes and the read/write mode.

is happening is that the WGMs self-align to defects in the crystal structure normally. But, since the coupler chip is essentially a large defect, the fields align to *it* instead. In this case, this allows us to couple by the same amount to both fields of the doublet, in the next chapter this may explain some difficulties we have seen.

6.4 Chapter Conclusions

This chapter extended the results of Chapter 5 by introducing the ability to parametrically frequency convert between whispering gallery modes of the sapphire dielectric resonator. There exist whispering gallery modes with very narrow linewidth (which corresponds to long memory times) alongside useful ‘read/write’ modes. As a demonstration of this, we implemented the store and retrieve protocol from Chapter 5. Since there are many modes available to us, in a practical volume, we think that optimizing the performance of this device by modifying the crystal geometry or coupler antenna is a worthwhile area of further study. But, since we are interested in incorporating quantum devices into this experiment, we pursue that next.

Chapter 7

Experiment III: A qubit coupled to WGM

In this chapter we focus on demonstrating qubit readout using the sapphire WGM cavity and performing some classic atomic/optical quantum optics experiments with our engineered system. By doing vacuum Rabi (Jaynes-Cummings) measurements, we learn how the transmon couples to each of the WGM resonances. And by sweeping the flux bias to the transmon's SQUID we can confirm all of the design/fabrication parameters of our qubit. This information is needed before chapter 8 to prove if the final experiment putting it all together is feasible.

On its own, this chapter represents a new 'hybrid' quantum system which will be interesting in its own right for further study. The fact that a qubit easily couples to a large crystal introduces the possibility for using single photons to study the materials properties of interesting bulk materials. This coupling is fully engineer-able via lithography of the qubit.

In some sense this is the opposite approach used to couple to bulk samples previously [61, 62]. In the two cited works, a bulk crystal was set on top of a chip-scale circuit. The circuit was used to probe the behavior of the sample under study. These samples often contain an ensemble of nuclear spins which the chip-resonator probes. In our case, we are using chip-scale circuits (smaller in scale) to modify the behavior of a large sample which we can probe directly. The qubit acts as a single spin whose behavior we wish to learn about.

In this chapter we also introduce some physics not covered in Chapter 2 to describe qubit readout and dissipation.

7.1 Device details

The set up for this experiment is very similar to the previous set up just with the coupler chip replaced with the transmon chip (see Figure 6.1). The only difference is that a different magnet coil is mounted above the transmon. The transmon SQUID loop area is much much smaller than the coupler's ($3500 \mu m^2$ for the coupler versus $\sim 3 \mu m^2$ for the transmon) and thus requires a larger amount of external flux to bias it, so to control the DC flux to the qubit we mounted a $10,000$ turn solenoid above the qubit. This solenoid provides the static DC flux to the qubit in all subsequent experiments.

We chose such a small SQUID area for the qubit to minimize the amount of flux noise that might decohere our qubit. But, as shown later in Chapter 9: if this design constraint is relaxed, it allows for more flexibility in the parametric interactions available to us. A later qubit design should include a larger SQUID area.

No other modifications were needed for this cooldown.

7.2 Qubit frequency inferred from splittings

In contrast to the previous two chapters where we induced normal-mode splittings by parametrically pumping a coupling circuit, in this chapter we return to resonant interactions momentarily. We use the flux tunability of the qubit to move the resonance frequency of the qubit onto resonance with various WGMs. This is precisely the physics described by the Jaynes-Cummings interaction of Chapter 2.

To review, in the on-resonance case, energy is swapped back and forth between the qubit and mode at a rate determined by the coupling between the electric field of the mode and the dipole moment of the qubit. In fact, we must use the vacuum Rabi splittings to map out where the qubit induces anti-crossings with the WGM frequencies and therefore infer the qubit frequency as a function of bias flux. This is because the transmon is nonlinear after only one photon's worth of excitation and therefore it cannot be read out directly.

The anti-crossing width is $2g$ and is proportional to the dipole moment of the qubit and the strength of the electric field at the qubit's location. The dipole moment of the qubit is smaller than the coupler and is the result of a design tradeoff between fast vacuum Rabi rates, and the Purcell effect – discussed below.

Later, this ‘vacuum Rabi’ swap rate will be necessary for calculating off-resonant, dispersive dynamics of the system, which we would like to exploit for readout of the qubit in the next section. So the first characterization measurements done with the qubit/WGM system are to sweep the qubit frequency through a several GHz range – 5 GHz to 8.5 GHz – while monitoring individual sapphire modes. This is preparation for moving the qubit off-resonant from the modes to do dispersive measurements later.

Here it is also important to review why we would like to move the qubit off-resonant from the cavity modes. First, if multiple qubits are involved in the system it would be detrimental to all parts of the system if they were all co-resonant. In essence the coherence time of the system would be solely determined by the *least* coherent constituent. Second, if the qubit is non-resonant it can undergo free evolution (or *no* evolution) in isolation. If the qubit is too close to a mode its coherence can be limited by the Purcell effect. The Purcell effect [63, 64] can be modeled as energy leaking out unwanted modes by spontaneous emission. A simple formula

$$\gamma_{Purcell} = \left(\frac{g^2}{\Delta}\right)\kappa \quad (7.1)$$

relates the decay rate from the Purcell effect, $\gamma_{Purcell}$, to the linewidth of a nearby mode (κ), detuned from the qubit by Δ . We can see from this formula, that we can enhance our qubit *decay* by detuning from a large linewidth cavity, or *decrease* this decay rate by detuning from a narrow linewidth cavity.

This simple model for Purcell loss is complicated by the fact that we have *many* resonances now, each with different detunings and linewidths. In the limit of multiple modes near the qubit, a more accurate Purcell decay rate needs to account for the impedance ‘seen’ by

the qubit over a broad range of frequencies. To avoid some of this loss mechanism, though, we designed our qubit with a small dipole moment – and therefore a small g . Since the Purcell loss is proportional to g^2 we can quadratically avoid this loss mechanism. While this may limit our vacuum Rabi rate, it should help with regard to decay from Purcell losses.

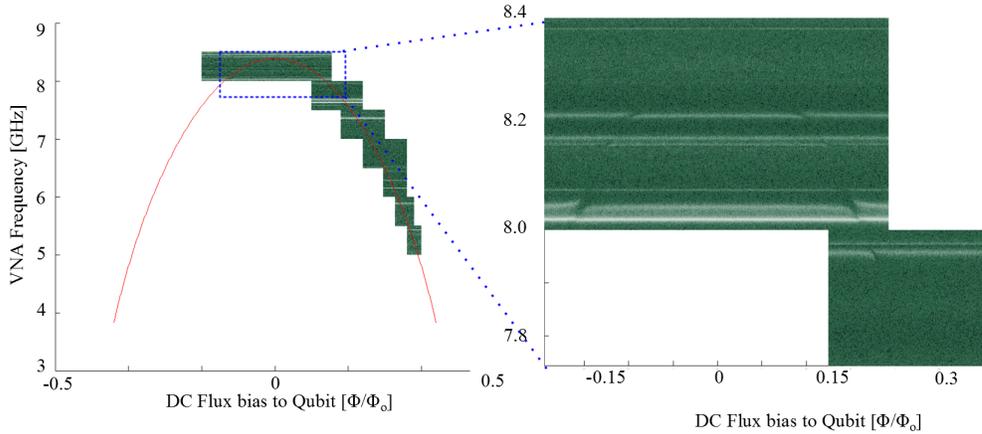


Figure 7.1: In (a), as the qubit’s frequency is swept (by changing the DC flux bias through the qubit’s SQUID loop) through various WGM frequencies, the qubit splits with some of the modes due to overlap of the qubit dipole with the electric field of the WGMs. The red line is the predicted qubit frequency (from fabrication parameters), and the horizontal white lines are individual modes. In (b), we zoom in on a few of the higher frequency splittings which show the characteristic splittings.

Thus, the first measurement we would like to do is sweep the static qubit DC bias through a full flux quantum which in turn changes the qubit frequency. While doing this we will monitor the WGM modes with the VNA at low probe powers. The low powers are necessary due to nonlinear, Kerr, effects that will be shown later. The power reaching the qubit is typically about -110 dBm, which corresponds to a few photons probing our qubit.

As the qubit crosses each WGM frequency a vacuum Rabi splitting should be observed if the Rabi rate is larger than the linewidth of the individual modes – in other words: if $g_n > \kappa_n$ for each of n modes. We summarize the g 's and κ 's in Table 7.1, below.

Additionally, we can also find areas of the WGM spectrum where the qubit does *not*

FREQ. (GHz)	$2g/2\pi$ (MHz)	$\kappa/2\pi$ (MHz)
8.215	15	4.2
8.16	9.8	3.1
8.04	34	7
7.96	22	5
7.69	16	9
6.52	7.7	2
6.44	24	8.5
6.28	15	2
5.43	15.5	1.5

Table 7.1: A list of the mode frequencies, their vacuum Rabi coupling rate to the transmon ($2g$), and linewidth (κ) as the qubit frequency is swept via a DC flux bias through various WGM resonances.

interact strongly. We can then set the qubit frequency away from other modes – i.e. we can chose operating fluxes where we will be least limited by the Purcell effect. In Figure 7.1 two example frequencies are near 8.3 GHz and 7.9 GHz – as far from other modes as possible. These frequencies are also convenient since when we add the coupler chip in the next chapter they are also in the range of frequencies where the coupler chip participates strongly enough to modulate the mode frequencies significantly. We would also like to have not only the ω_{01} frequency of the qubit isolated from WGM resonance, but also e.g. the ω_{12} frequency. Since we expect the anharmonicity of the qubit to be about 300 MHz, these two frequency ranges (separated by about 300 MHz) will be ideal for later measurements.

7.3 Dispersive qubit-cavity measurements

Returning to the Jaynes-Cummings Hamiltonian physics of Chapter 2, we again write our Hamiltonian (same as equation 2.63) as:

$$H = \hbar\Omega_c(a^\dagger a + \frac{1}{2}) + \frac{\hbar\omega_a}{2}\sigma^z + \hbar g(a^\dagger\sigma^- + a\sigma^+). \quad (7.2)$$

If we now consider the case where the qubit is at a different frequency than the mode that it is coupled to, we can expand the Hamiltonian in powers of g/Δ to second order where Δ

is the frequency difference between the qubit 0-1 transition frequency and the cavity mode: $\Delta = \omega_{01} - \omega_c$. This expansion requires that $g \ll \Delta$ and is derived fully in references [65] and [66] and gives:

$$H = \hbar(\Omega'_c + \chi\sigma_z)a^\dagger a + \frac{\hbar}{2}\omega'_q\sigma_z \quad (7.3)$$

where $\chi \approx \frac{g^2}{\Delta} \frac{E_C}{\hbar\Delta - E_C}$ denotes the effect of coupling two renormalized frequencies ω'_c and ω'_q . The old cavity and qubit frequencies are shifted, or ‘dressed’, by the interaction. This shift is known as the Lamb shift and shifts the original frequencies thusly: $\omega'_q = \omega_{01} + \chi_{01}$ and $\omega'_c = \omega_c - \chi_{12}/2$ with $\chi_{ij} = g_{ij}^2/(\omega_{ij} - \omega_c)$.

We have grouped the σ_z term with the cavity frequency to highlight the fact that this Hamiltonian now describes a system where the cavity frequency is dependent on the *qubit’s* state. This state-dependent frequency shift is the essence of many qubit readout schemes. If the qubit is in the ground state one reads out a particular magnitude or phase of the cavity mode, and if the qubit is in the first excited state one reads out a different magnitude or phase. If the dispersive shift is larger than the linewidth of the cavity mode of interest, this results in high-fidelity, reliable qubit measurements.

The problem with this type of measurement is that it requires low photon numbers (proportional to $a^\dagger a$ in the Hamiltonian, above) as will be discussed next. This constraint produces a signal to noise of the measurement that is very small. A good deal of study in our group, as well as others, has been the development of a very low-noise amplifier called a Josephson Parametric Amplifier (JPA), which can increase this signal-to-noise to acceptable levels. A future experiment and future quantum computing technologies will almost certainly include a JPA. An alternative method for qubit state readout uses much higher powers, and has a larger signal-to-noise ratio. This high-power dispersive readout is described below.

7.3.0.1 High-power dispersive qubit measurement

Dispersive measurement techniques exploit the qubit's effect on the cavity's frequency or phase. At higher and higher photon numbers in the cavity though (higher drive powers), this interpretation breaks down – the cavity occupation can instead affect the hybridized (dressed) states of the combined system. This can be quantified by adding a ‘Kerr’ term – proportional to $(a^\dagger a)^2$ (rather than the previous $a^\dagger a$) – to the Hamiltonian. Indeed with even higher cavity photon numbers the cavity-qubit systems undergoes dynamics described by a Duffing oscillator rather than a Harmonic oscillator and must be described semi-classically due to the size of the Hilbert space [67, 68].

This behavior is shown in Figure 7.2 for a sample WGM resonance near 10.0925 GHz. This plot shows a color plot of the cavity frequency (x-axis) as a function of drive power (y-axis); and sample linecuts on the right. The qubit is detuned from this frequency by 1.39GHz below ($\omega_q = 2\pi 8.28\text{GHz}$). At the lowest drive powers, we see that the cavity's natural frequency has been dressed to a higher frequency – this is the dispersive limit of the Jaynes-Cummings model. As power is increased, the Kerr terms become more relevant and the cavity frequency shifts lower (χ is negative), and eventually may even show bifurcation behavior.

Interestingly, at very high powers (large photon numbers) a curious phenomena was recently observed [69]. The cavity becomes much more linear, symmetric, and stable in frequency. This behavior persists for all higher powers too and denotes the return to ‘classical’ behavior. The surprising aspect of this transition is that the point, in power, for which this behavior switches to the classical response depends strongly on which state the qubit was in when starting to populate the cavity to begin with. In this manner, a new type of qubit readout was born which has better signal-to-noise (because one can measure at high power), and does not require a JPA. Two related things to note about this measurement are that it backacts on the qubit and therefore is not a Quantum Nondemolition Measurement, and

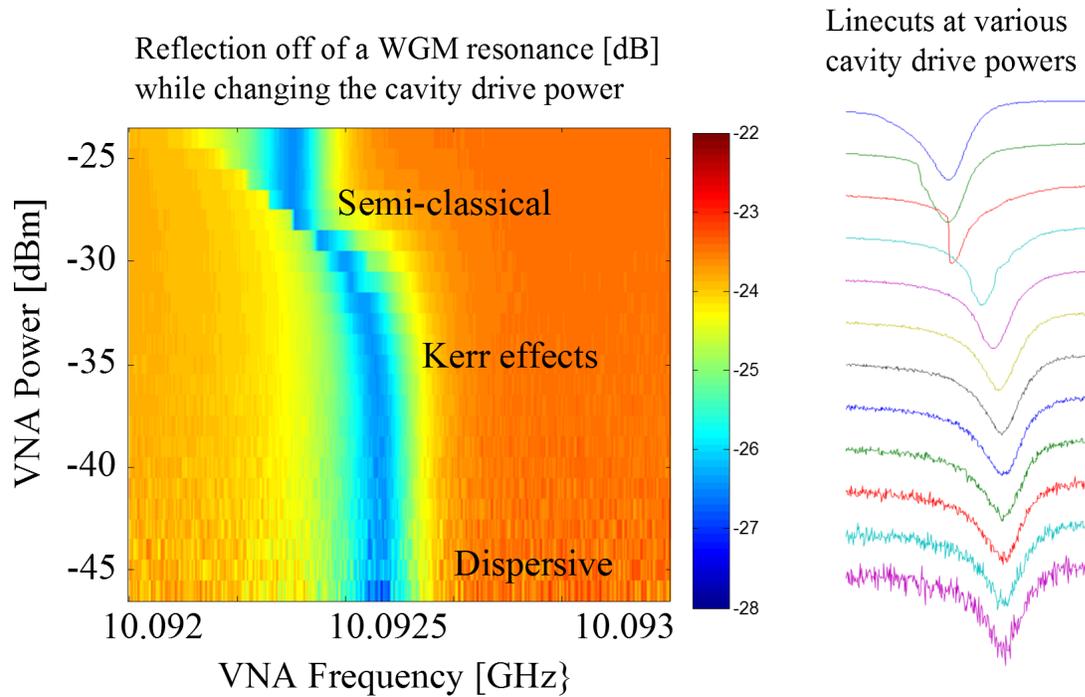


Figure 7.2: VNA reflection measurements [dB] of a cavity mode resonance as the drive power is changed (y-axis). At the lowest drive powers the cavity only exhibits dispersive shifts. At higher powers Kerr terms begin to play a role and move the resonance downward in frequency. At the highest powers, the original, classical behavior returns fully. The linecuts on the right side show the Lorentzian, then distorted Duffing-oscillator behavior, and return to Lorentzian lineshapes.

it is a ‘latching’ measurement. The qubit imprints its state on the cavity and undergoes no further controllable, coherent, time-evolution.

What is required for this measurement to be used? The strength of the effect is governed by χ as before, namely $\chi \approx g^2/\Delta$. Qualitatively, one can picture driving the cavity off-resonant from the dressed cavity resonance, but at the *classical* cavity frequency. If the qubit is in the ground state, the cavity takes some amount of power to latch to the classical behavior. But, if the qubit is in the first excited state, the cavity latches to classicality earlier (at a lower power). We note this is true if the qubit frequency is below the cavity frequency, which will be the case in this work. In the opposite case – with the qubit above the cavity – this measurement technique still works, but is more difficult. So, we prefer to have the readout cavity frequency above the qubit’s and we would like the g^2/Δ factor to be large compared to the linewidth of our cavity.

To maximize χ for the case of the WGM resonator, this is just a matter of sweeping frequency and power with the VNA for various cavity frequencies. We note that for some of these frequencies, we already know what g is based on the vacuum Rabi splitting data from above. So this is a good method to confirm that measurement. For WGM frequencies above the tunable range of the qubit, we set the qubit frequency at ~ 8.28 GHz, again, this frequency is as far away as possible from many of the WGM frequencies also to avoid Purcell effects.

After a search of many of the WGM frequencies, we decided that a mode near 10.0925 GHz provided the best balance of g , Δ , and linewidth – again, see Figure 7.2. To perform qubit readout, we probe the cavity with a power low enough to *not* reach the classical behavior if the qubit is in the $|0\rangle$ state. This results in very little cavity response. If the qubit is in the $|1\rangle$ state, though, this power will result in a large, classical resonance response.

In contrast to the low-power dispersive technique, this must be performed as a pulsed measurement because once the high power is provided to the cavity, there is significant

back-action to the qubit (this readout is again, not QND). After performing the measurement pulse, all pulses are turned off for typically $500\mu s$ to make sure all the systems relax, thermally, to their ground state. Even with this fairly long delay, this method results in faster readout because the signal to noise is so much higher we need not average the result as much compared to traditional dispersive measurements.

The readout pulse itself is typically set to $10\mu s$, though we only integrate the first microsecond to get the measurement signal desired [69]. The measurement frequency, integration time, and readout pulse power are carefully calibrated to maximize the signal to noise of the readout while performing the following qubit characterizations.

7.4 Qubit characterization

In order to characterize the qubit behavior we will perform two measurements. These are both classic quantum optics-type experiments and are fully described in e.g. reference [3]. First we use long pulse-length spectroscopy to find the transition frequencies for the first few transmon levels – i.e. ω_{01} and ω_{12} . This information provides us with confirmation that we have designed and fabricated the device such that it will be useful for later experiments – recall that the anharmonicity is equal to E_C , which in turn depends on the capacitance. Second, we shorten the driving pulses such that the two-level system dynamics of our artificial atom can be shown conclusively. This will also calibrate us to how long we will need to drive the qubit to produce a single photon on-demand – the π -pulse length.

To probe the ω_{01} transition frequency, a long ($\sim 50\mu s$) pulse near the expected ω_{01} frequency (inferred from the vacuum Rabi splitting curve above) is applied through the pump coax line. This long pulse drives the qubit to a superposition state, and may also drive higher transitions (i.e. ω_{12} and above), depending on the drive power. This technique is often called ‘saturation spectroscopy’ for this reason – we can saturate lower levels and see higher level dynamics start to appear. At the end of that pulse, the cavity (at 10.092 GHz)

is queried. Again, this query is a long-pulse ($\sim 10\mu\text{s}$) drive tone at a power below the power required for the classical behavior to emerge. The resulting cavity query level will depend on if we are on-resonance with the qubit or not.

Once that qubit frequency is determined we can change the qubit's flux and repeat the scan. This requires another calibration of the optimal readout power and frequency so was only performed for a few qubit fluxes. Performing the measurement with and without the qubit in the measurement range allows us to confirm that this indeed *is* the qubit which can be difficult in the forest of WGM resonances!

The saturation spectroscopy data is shown in Figure 7.3. First, we show this spectroscopy data with the qubit frequency fixed at ~ 8.283 GHz as we change the qubit drive power. The pulse schematic is sketched in Figure 7.3a. The linewidth of the qubit's ω_{01} transition frequency broadens with the amplitude of the drive. This again confirms the expected 'power broadening' behavior that lets us know that the resonance that we are probing is indeed the qubit, not a spurious resonance [45]. At the very lowest powers, the linewidth of the qubit resonance is about 400 kHz. This linewidth, κ_{01} accounts for both decoherence, and dephasing qubit behavior:

$$\kappa_{01} = \frac{1}{2T_1} + \frac{1}{T_2}. \quad (7.4)$$

Where T_1 is the energy relaxation time, and T_2 is the dephasing time. Thus, this linewidth gives us an upper bound on the qubit's coherence time (T_1) of approximately $2.5\mu\text{s}$ (if $T_1 = 2T_2$).

By performing a measurement including *two* drive tones we can directly measure the 1-2 transition – i.e. one tone, at a fixed frequency saturates the 0-1 only, the other tone scans in frequency to drive the 1-2 at the same time (and also the 0-2/2). This is a direct measure of the anharmonicity of the transmon as fabricated. From this data we can say that the 1-2 occurs ~ 306 MHz below that of the 0-1 transition which corresponds to a capacitance of 63 fF. This is very similar to the designed value of 70 fF. The red trace in Figure 7.4 shows the

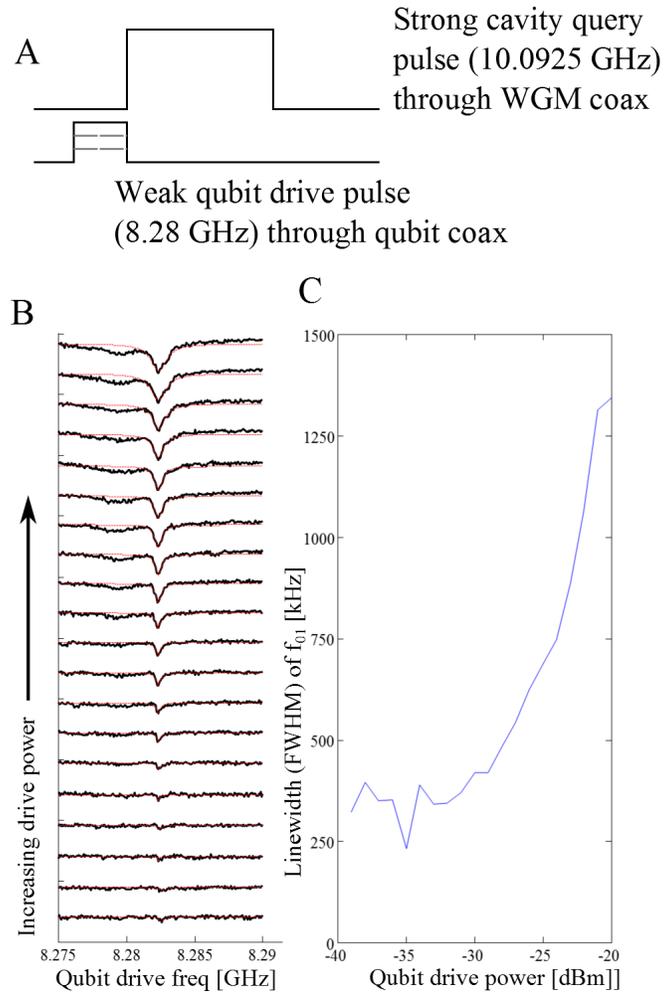


Figure 7.3: a.) A pulse diagram noting the two pulses in this measurement – the first saturates the qubit resonance and is driven through the qubit coax. The second pulse queries the cavity mode and infers the state of the qubit. b.) To determine the exact qubit frequency we drive the qubit with various powers near where we suspect ω_{01} exists then query the readout cavity at 10.0925 GHz. c.) With increasing qubit drive power the 0-1 transition become power broadened which we can see by fitting the linewidths of each of the qubit spectroscopies.

low-power driving of the qubit with only the 0-1 transition visible. The blue trace (offset for clarity) shows the two-pump case with all three qubit transitions visible.

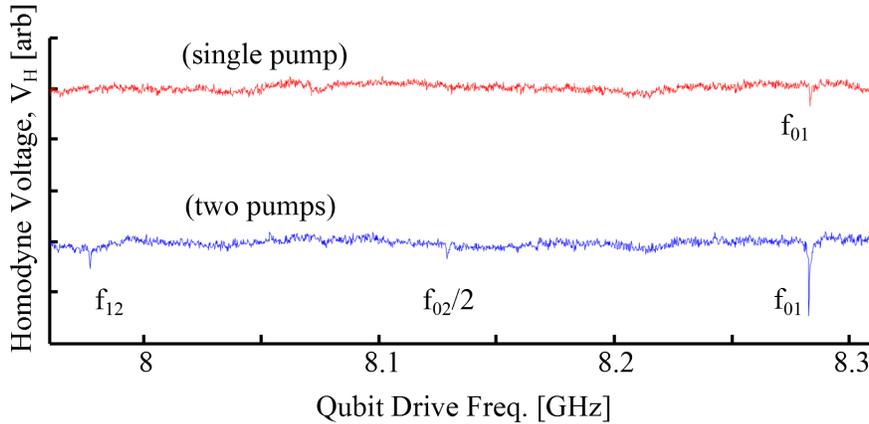


Figure 7.4: Red trace: by pumping the qubit drive at low power and scanning frequency the 0-1 transition can be isolated. Using two pumps (blue trace) the higher order qubit transitions can also be excited. The spacing between the 0-1, 1-2, and 0-2/2 transitions give a direct measurement of the fabricated qubit parameters.

Comparing Figures 7.4 and 7.1 we also note that by driving through the port directly above the qubit chip, we do *not* excite the whispering gallery modes near the qubit frequency directly. This is due to the fact that to excite those modes efficiently, a coaxial loop probe needs to be positioned where it is indicated in Figure 6.2, not near the qubit location. At very high qubit drive powers ($\sim 20 - 40$ dB more power than required to probe the qubit) the nearby WGM resonances would appear in Figure 7.4. Since qubit measurements are inherently low-power, though, this is acceptable. In Chapter 8, this will be a complication.

We can directly measure the qubit dynamics in time by performing short-pulse excitations of the qubit. This is similar to the spectroscopy above, and to the frequency conversion experiment performed in Chapter 5. We drive the qubit tone for varying lengths of time. The shortest pulse is 20 ns long and increase from there. For each pulse length we drive the qubit from the $|0\rangle$ to the $|1\rangle$ state or vice versa. The pulse scheme is shown schematically in Figure 7.5a. We note that we are now using a Gaussian pulse shape rather than a square

pulse. The square pulses have fast rise/fall times, and therefore have high-frequency Fourier components which could excite the higher qubit levels. Gaussian pulses suffer from this less so. Other, more complex pulse shapes could be used to null this effect [70]. At the end of the qubit driving pulse, we again query the readout cavity. This data is plotted in Figure 7.5c. As the pulse length is increased we can explicitly demonstrate the two-level system behavior of the transmon, and place a limit on the drive power we can apply to only drive transitions between the $|0\rangle$ and $|1\rangle$ energy levels. This *driven* Rabi experimental data is fit to a decaying sine function with a decay constant of $1.51 \mu\text{s}$. At higher drive powers the decaying sine does not saturate to 0, indicating that higher levels are being populated.

Now the driven Rabi experiment finished, we know the precise pulse length, pulse power, and frequency to perform a T_1 decay measurement – this is the ‘ π ’ pulse from quantum optics language. A π pulse is such that it maximally excites the $|1\rangle$ state – with our selected power and frequency this was ~ 90 ns. In this case we fix all of those parameters, but change the delay time between the end of the excitation pulse and the measurement pulse – see pulse sequence in Figure 7.5b. The corresponding data, in panel d of Figure 7.5, shows a decaying exponential with a relaxation time constant of $1.59 \mu\text{s}$. This acts to confirm the Rabi oscillation and spectroscopy data. The difference in fit parameters can be explained by an ambiguity of when the Gaussian pulse ‘ends’ – this was known only to ± 10 ns.

With these simple qubit demonstrations finished, another needed proof-of-principle aspect of our future experiments was checked off the list. Namely, the ability to create a non-classical state of (microwave light) – e.g. a single photon. The next step is to demonstrate not just the ability to read out the state of the qubit dispersively, but to actually use parametric processes to transfer that state into the mode structure of the WGM crystal. Chapter 8 will detail those efforts.

First, though, it is important to note that we surreptitiously added some physics to our discussion of qubit-cavity dynamics by introducing energy decay and dissipation. Recall that

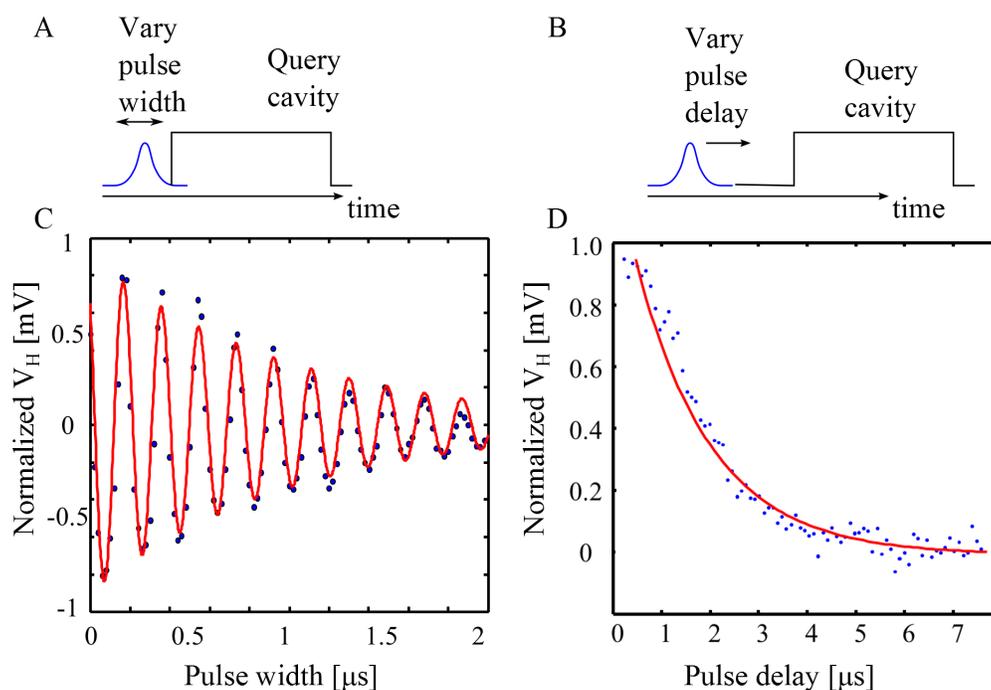


Figure 7.5: a.) & b.) The pulse sequences used to produce the driven Rabi oscillations shown in (c) and the pulse sequence used for the T1 decay data shown in (d). c.) By varying the pulse width driving the qubit 0-1 transition the qubit can be made to oscillate between the $|0\rangle$ and $|1\rangle$ energy levels. d.) An initial π pulse directly excites the qubit to the $|1\rangle$ state, then after a varying delay time we query the cavity to determine the leftover population.

f_{01}	f_{12}	E_J/h	E_C/h	L_J	I_o	C
8.283 GHz	7.977 GHz	31.4 GHz	0.308 GHz	5.23 nH	63 nA	63 fF

Table 7.2: Extracted qubit parameters for the transmon, as fabricated.

the Jaynes-Cummings model only describes coherent evolution. In the strong coupling regime (with $g > \text{decay rates}$) this is modeled simply by multiplying the oscillations by a decaying exponential. For a full, Master Equation approach to the problem, see e.g. reference [41].

7.5 Chapter conclusions

In this chapter we finally introduced the quantum part of the system via a transmon qubit (artificial atom). We measured two interactions of the qubit with the whispering gallery modes. First, vacuum Rabi anti-crossings were observed in the frequency-domain. We will use this data next chapter to determine the parametric swap rates we can induce. Then we moved to the time-domain and proved that our qubit is indeed acting as designed in this new hybrid setup and that the cavity modes are indeed usable for quantum measurements. This was also necessary to show before moving on to the next chapter's experiment.

To simplify the system dynamics, we avoided including the coupler chip in this experiment. In the next chapter we simply add that chip back in.

Chapter 8

Experiment IV: Qubit and coupler and WGM

The goal for this set of measurements is to combine the sapphire WGM resonator, the qubit, and the coupler into one system capable of creating, storing, and reading out a single photon. While we have used the Von Neumann quantum computing architecture as a motivation for doing this, there are other, interesting experimental uses for a system able to create and store single photons on demand. An example from our group used resonant and parametric interactions to show a Hong-Ou-Mandel interference between photons recently [47]. Other, more complex photon states could be created and interacted by using the long-lived whispering gallery modes. These modes, if long-lived compared to the individual operation times would allow for multiple operations.

In the referenced experiment, the qubit was able to be tuned on fast timescale into and out of resonance with the cavity modes. But, there exists experimental [54] and theoretical [71] evidence that we can use the same parametric interactions that we have been discussing in previous chapters with the qubit-cavity system, not just the cavity-cavity system of e.g. Chapter 5.

The experimental design is to excite the qubit to its excited state, use parametric frequency conversion between the qubit and a fast read/write mode to then load that photon into the sapphire mode structure. At that point, we would like to use another frequency conversion process to swap the single photon from the read/write mode to a higher-Q, long-lived storage mode. After a storage time, we would then like to swap the photon *back* to the

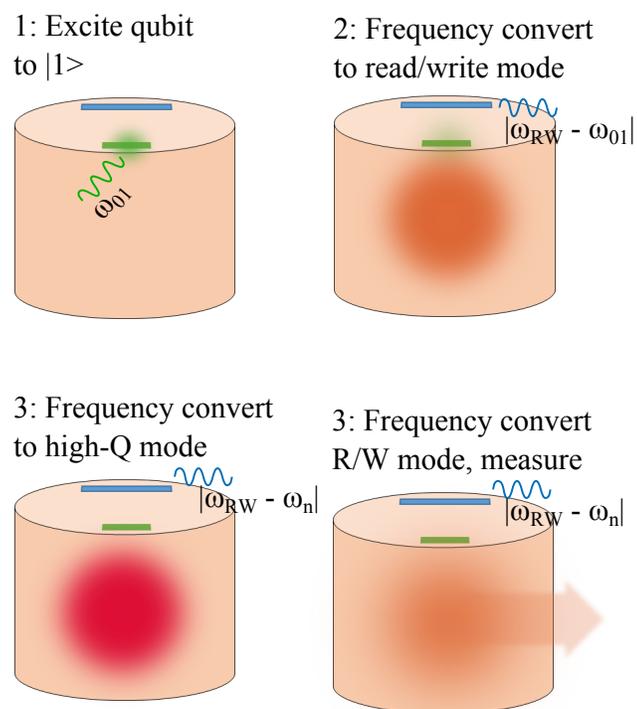


Figure 8.1: A cartoon of the experiment proposed for this chapter. A single photon is created using the qubit. It is then transferred via frequency conversion to a mode (the ‘read/write’ mode) of the WGM. Another frequency protocol transfers the photon to a high-Q mode where it can be stored for times much longer than the qubit’s decoherence time. The photon is then frequency converted back to the read/write mode and read out by the measurement chain.

read/write mode where we can finally measure the photon.

From previous chapters we have demonstrated almost all of these steps of the protocol individually. This setup is complicated by the fact that the frequency overlap of the qubit (maximum frequency of about 8.5GHz) is lower than the optimum WGM frequencies that have high Qs (typically above 10 GHz). From the data shown in Figure 6.3, though, there exist modes that modulate with coupler bias in that range.

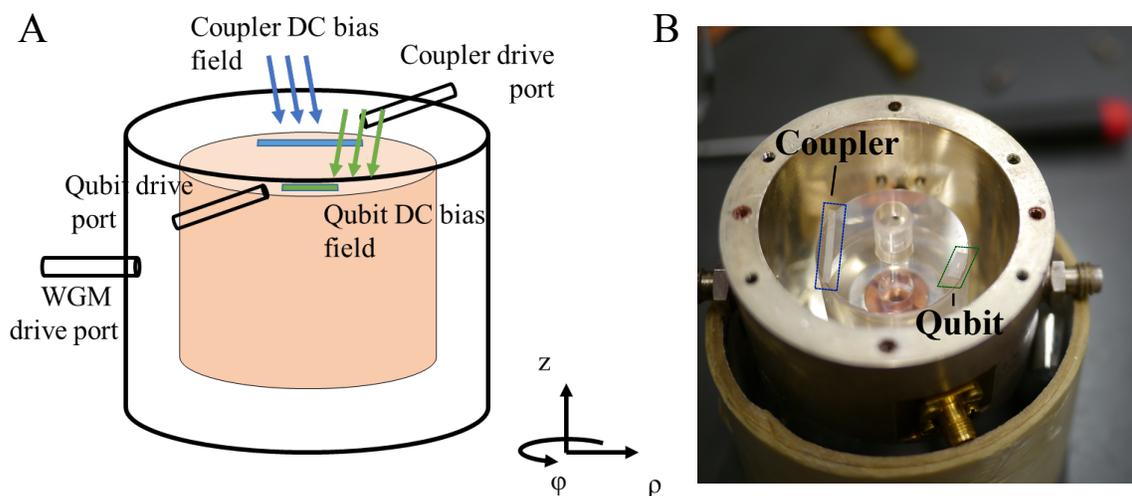


Figure 8.2: a.) Shows a cartoon of the three microwave ports used to drive the coupler, qubit, and WGM resonator, as well as the two DC bias fields directed over the coupler and qubit chips. b.) Shows a photograph of the system before placing the lid containing the pump coaxes and bias coils on top.

8.1 Device details

This experimental setup includes elements from the previous two chapters. Now we have mounted the coupler chip on one side of the WGM crystal and the qubit radially across the top surface on the other side of the crystal. Each chip has its own flux pump line and DC bias coil. This means there are two pump coils, two DC bias coils, and one cavity input/output probe.

With the added complexity of this experiment though we present data from multiple

cooldowns wherein we attempt to formulate a coherent picture of the device properties.

8.2 Sideband gate theory

The theory for parametrically swapping photons from a qubit to a resonator mode is outlined theoretically in a recent paper by Blais et al [71] and this protocol was experimentally demonstrated by Plourde’s group soon thereafter [54]. This experiment at Syracuse used a frequency-tunable qubit as the modulate-able frequency and an external pump to swap single photons into a lumped element resonator. We seek to extend these results by swapping the qubit state into our multi-mode whispering gallery mode resonator via the read/write mode by using our frequency tunable *cavities* rather than the qubit. Both of these protocols have also been demonstrated in single-cavity mode lumped-element devices measured in parallel to the experiments described here in a separate fridge in our lab [72]. These ‘sideband’ gates can be thought of as the limit of cavity-cavity parametric frequency conversion to cavity-two-level-system frequency conversion. There are a few subtleties and differences in the theory when comparing to what we introduced earlier that are worth mentioning.

First we outline the theory for the standard ‘sideband gate’ theory, where the qubit is modulate-able via a pump, then extend those results to a modulate-able cavity. We start with the Jaynes-Cummings Hamiltonian, again but assume that the qubit frequency can be changed via a pump in time ($\omega_{01}(t) = \omega_{01} + \frac{\epsilon}{2} \sin(\omega_{FC}t)$):

$$H(t) = \hbar\Omega_c a^\dagger a + \hbar \frac{\omega_{01}(t)}{2} \sigma^z + \hbar g (a^\dagger \sigma^- + a \sigma^+). \quad (8.1)$$

Just like in the frequency conversion case we assume that the qubit frequency is modulated at the difference frequency between the cavity and qubit resonances: $\Delta = |\omega_{01} - \Omega_c|$. We also require that this difference frequency is greater than g .

The theory paper continues by moving to a rotating frame defined by

$$U(t) = \exp(-i(\omega_{01}t - \frac{\epsilon}{2\omega_{FC}} \cos(\omega_{FC}t)\sigma^z - \Omega_c a^\dagger at)). \quad (8.2)$$

In this frame the Hamiltonian can be written as:

$$\begin{aligned}
 H'(t) = & \hbar \frac{\epsilon}{2} (\sin(\omega_{FC}t) + \cos \omega_{FC}t) \sigma^z \\
 & + \hbar g a^\dagger \sigma^- \sum_{m=1}^{\infty} (-i)^m J_m\left(\frac{\epsilon}{\omega_{FC}}\right) e^{i(m\omega_{FC}-\Delta)t} + H.C.
 \end{aligned} \tag{8.3}$$

where the J_m terms are Bessel functions of the first kind, and m denotes the m -photon transition. In our case, we can simplify greatly by assume we will only look at the single-photon transitions and assume that $\Delta = \omega_{FC}$. In that case, the second line of equation 8.3 simplifies to an effective coupling rate of:

$$g_{eff} \approx g_o \frac{\delta\omega_{01}}{2\Delta}. \tag{8.4}$$

Here g_o is the vacuum Rabi coupling (resonant coupling) between the qubit and cavity, and $\delta\omega_{01}$ is the amplitude at which we are pumping the qubit frequency. This effective coupling rate has an important caveat. It limits the amplitude of pumping ($\delta\omega_{01}$) to be less than the detuning (Δ) – i.e. $2\delta\omega_{01} \approx \Delta$ – which implies that the maximum effective coupling possible is limited by $\sim g_o$.

This approach has been very successful in our lab and the Plourde group. The vacuum Rabi rate for lumped element cavities coupled to transmons can be designed to be > 100 MHz and the detuning can be set to $\sim 20g$. For moderate pump amplitudes that modulate the qubit by say 200 MHz, this results in effective parametric rates of $g_{eff} = g \times \frac{200\text{MHz}}{20g} = 10\text{MHz}$.

In the specific case described in this work though, the mutual inductance between our qubit's SQUID loop and the qubit pump coax is too small to be able to pump the large amplitudes needed to modulate the qubit frequency by those amounts. In other words, the SQUID area is too small. This is another aspect of the design that will need to be changed in later designs, discussed in the conclusion chapter (Chapter 9).

There is good news though: if rather than modulate the qubit frequency with the cavity static; we modulate the cavity frequency (via the coupler), with the qubit static, the

effective parametric coupling rate goes as:

$$g_{eff} \approx g_o \frac{2\delta\Omega(C)}{\Delta}. \quad (8.5)$$

To maximize *this* rate we would like:

- Δ to be small: we will tune the qubit frequency to be 10-100 MHz away from the cavity mode to be coupled to.
- g_o to be as big as possible: see Table 7.1 and below, typical vacuum Rabi rates are 1-10 MHz.
- $\delta\Omega_c$ to be as big as possible: typically, the WGM modes modulate by 10-30 MHz.

The first item, above, limits the range of WGM frequencies in which we can operate. Since the qubit has a maximum frequency of about 8.5GHz, we will restrict ourselves to cavity modes from $\sim 5 - 8.5$ GHz. The third item also restricts the possible rates achievable – the coupler chip participates less and less at lower frequencies. The modes that modulate with coupler flux the most are typically 7+ GHz. These limits predict, though that we will be able to induce the largest parametric coupling rates in a range of frequencies between ~ 7 and ~ 8.5 GHz.

What parametric coupling rates are possible? Combining all of the typical rates from the list above yields $g_{eff} = 1 - 10$ MHz. For example, in Table 8.1 we take the data from Chapters 6 and 7 to predict g_{eff} . To be conservative, we take the half of the maximum cavity modulation to be $\delta\Omega_c$ and Δ to be about $2g$. Additionally, recall that our qubit has a lifetime of about $1.5 \mu s$, so we would like to induce swapping at a rate faster than ~ 1 MHz.

With this data in hand, we cooled the full experimental apparatus down and began to characterize the full device operation.

f_{cavity} [GHz]	$g_o/2\pi$ [MHz]	$\delta\Omega_c/2\pi$ [MHz]	$\Delta/2\pi$ [MHz]	$\sim g_{eff}/2\pi$ [MHz]
8.215	7.5	4	10	6
8.16	4.9	4	10	4
8.04	17	12.5	25	17
7.96	11	10	20	11
7.69	8	3	15	3
6.52	3.6	5	7	5
6.44	12	2	25	1.5
6.28	7.5	2	15	2
5.43	7.52	2	15	2

Table 8.1: By combining the data from Chapters 6 and 7, we can predict what effective parametric coupling rates are possible with the coupler + qubit + WGM system.

8.3 Device characterization

The data presented below are from four separate cooldowns of the full device: WGM crystal with coupler and qubit. The number of cooldowns is due to the fact that during the following characterization steps, unexpected behavior was seen. This behavior contradicted the data from Chapters 6 and 7. We note the discrepancies and compare the the data from those chapters, below.

WGM mode frequency modulation with coupler flux

We begin as usual by sweeping the DC flux bias to the Josephson circuits. There are now two DC biases, one centered on the coupler chip, the other centered on the qubit. The physical size of the DC bias coils ensures that there is significant field overlap between the two, so careful calibration to offset the effect of either coil is needed. In other words, while increasing one bias current we need to compensate the other bias coil such that for example, the qubit stays fixed in frequency while the coupler is allowed to sweep through a flux quantum.

First we show data from sweeping the coupler flux bias and monitoring the WGM resonance frequencies. In Figure 8.3 we compare this data from four different cooldowns. The leftmost panel shows data from the experiment described in Chapter 6, and the three panels to the right are from three cooldowns with the coupler placed at different locations on the crystal. The coupler is moved by about 10 degrees along the circumference from the original placement (panel 2) each subsequent cooldown.

There are two important pieces of information we can infer from the data in Figure 8.3. First, the placement of the coupler chip now has a very large effect on how much the WGM frequencies modulate. The experiment in Chapter 6 did not take special care to align the coupler in a particular place along the circumference. In fact, in that chapter we presented data showing that the WGM fields may actually be *aligning* to the chip's placement. We theorize that by including the second chip, the WGM fields align in a more complex way.

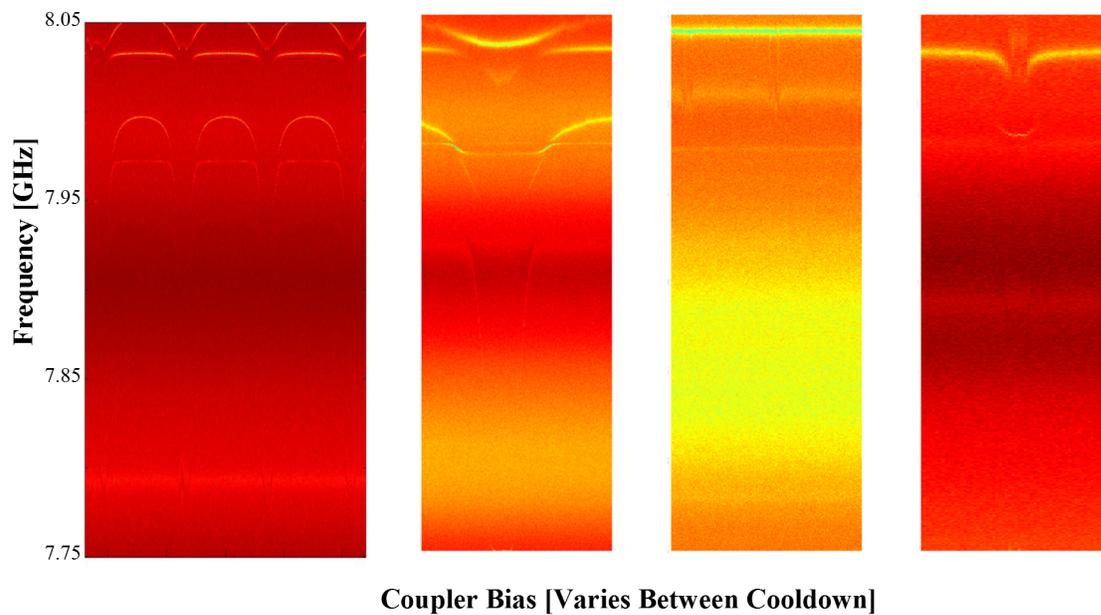


Figure 8.3: Comparing The original measurement with just the coupler chip (a), to three subsequent cooldowns which included the coupler chip *and* the qubit. b-d: the modes modulate by very different amounts and vary strongly with qubit chip placement. In fact, the WGM frequencies are also changing.

The second important thing to note from this data is that the mode frequencies change from cooldown to cooldown. Again, this behavior was *not* seen in data from Chapter 6 nor Chapter 7 which sometimes required multiple cooldowns and chip placements.

The amount that the modes modulate ($\delta\Omega_c$), in the frequency range which overlaps the qubit's range (about 5.5 to 8.5 GHz) also seems to vary strongly with chip placements *and* is qualitatively usually less than before (for a few mode frequencies it can be greater).

Qubit vacuum Rabi splittings

The qubit also exhibits different behavior than observed in Chapter 7. Figure 8.4 compares an example Vacuum Rabi splitting from the data in chapter 6 (left panel) and two of the cooldowns with both chips on the crystal. Again we notice that the WGM frequencies change between cooldowns, and also that the vacuum Rabi splitting changes. In almost all cases the vacuum Rabi anti-crossing is *smaller* than observed with just the qubit chip.

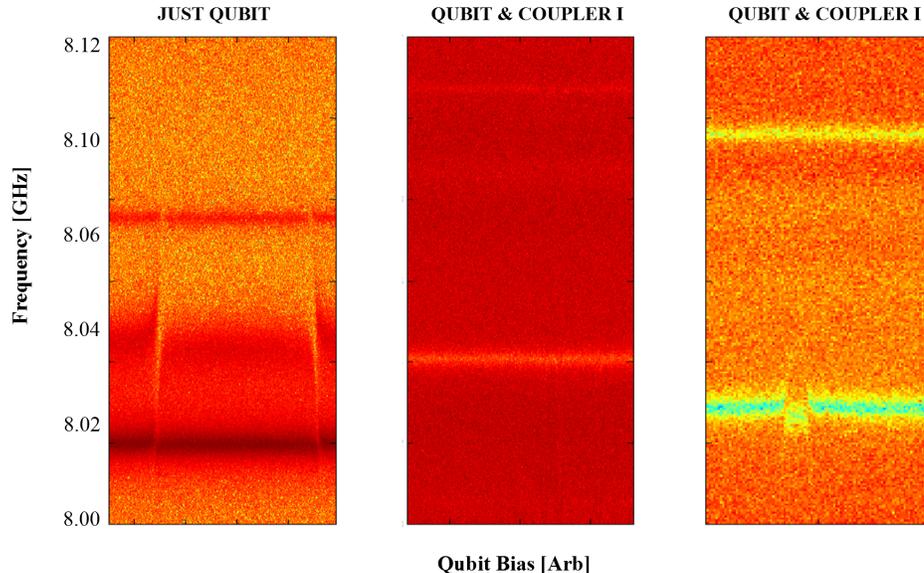


Figure 8.4: Rather than changing the coupler's bias field, we can measure the qubit's vacuum Rabi splitting with the WGM modes without (a) and with (b-c) the coupler chip also in place. Again the qubit splittings are almost always smaller, and sometimes nonexistent.

Again, we theorize that the electric fields inside the sapphire volume are now 'shared'

between the coupler and qubit circuits more than in the previous experiments. This serves to decrease the electric field which the qubit/coupler antennae coupler do via their dipole moments. This smaller dipole coupling translates to smaller modulations of the WGM frequencies *and* smaller vacuum Rabi rates. Again, by moving the coupler and qubit chips to different placements along the circumference of the sapphire crystal a large, heretofore, unobserved correlation between chip placement and mode coupling was observed.

This theory has support from the fact that over the two cooldowns shown in Figure 8.4 the number of anti-crossings induced by the qubit were fewer than the number from Chapter 7. Table 7.1 lists nine anti-crossings in the frequency range of about 5.5 to 8.5 GHz with an average vacuum Rabi rate of $\bar{g} = 2\pi \times 8.8\text{MHz}$. For the first cooldown described here, there were only four in that range with $\bar{g} = 2\pi \times 2.6\text{MHz}$. The second cooldown, after moving the chips, showed five with $\bar{g} = 2\pi \times 2.9\text{MHz}$. The anti-crossings observed, on average are smaller.

An additional complication incurred with this experimental setup is shown in Figure 8.5. This Figure shows repeated measurements at a precise mode resonance while changing the bias current to the qubit's flux bias (x-axis) such that the qubit is induced to cross that resonance. Where the qubit crossed the mode is indicated by the red color in the plot. This same measurement is repeated many times (y-axis) over the course of about 140 minutes. For each measurement in time, the point in qubit flux where the qubit crosses the mode frequency changes. This indicates that the qubit frequency is drifting in time now. Fortunately, the mode frequency is stable in time (e.g. data in other figures).

The explanation for this instability of the qubit (also not seen in Chapter 7) is unclear. This behavior could be explained by trapped magnetic vortices in the superconducting thin films forming the qubit antenna pads, but all magnetic fields are 'off' (or minimized) during cooldown to base temperature to avoid this as a standard practice in all four chapters. Why this would only manifest itself during these cooldowns is unknown.

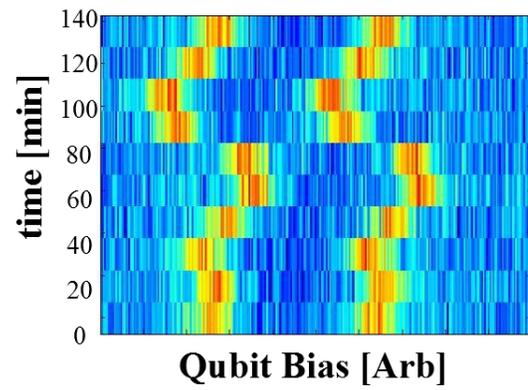


Figure 8.5: If we repeatedly scan over the vacuum Rabi splitting over the course of a few hours, the splittings move over time. This indicates that the qubit frequency is not stable either.

Consequences for qubit-cavity frequency conversion

From the discussion, above, we need to re-evaluate the feasibility of implementing the parametric qubit-cavity frequency conversion protocol introduced in this chapter. First, our cavity mode modulation, $\delta\Omega_c$ has been generally reduced. This could be remedied by pumping the interaction harder, but only to the point where power is dissipated too strongly at the mixing chamber stage of the fridge. Second, our vacuum Rabi rate, g_o is at least a factor of two on average less than previously. Third, we have fewer modes in which we can try this experiment with.

If we re-consider the effective parametric coupling rates shown in Table 8.1 in light of the results from this chapter almost all of the rates are reduced below $g_{eff} = 2\pi \times 1\text{MHz}$, the necessary rate to swap a photon before it decays from the qubit. A few modes are viable still, and the protocol shown in step 2 of Figure 8.2 was tried without success.

8.4 Is high-power Jaynes-Cummings style readout possible?

The experiments described above were not all bad news. We were also able to confirm that performing the high-power Jaynes-Cummings/Kerr readout technique described in Chapter 7 is still possible. The equivalent curve as Figure 7.2 is shown below in Figure 8.6. In this case, the qubit is detuned from the cavity *above the cavity frequency*. The low-power dispersive shift then moves the cavity frequency lower. As power is increased, the frequency again tends toward a semi-classical behavior and latches to a certain frequency. In this case, though, the frequency that we latch to is not the bare, classical, cavity still. The coupler's non-linear (Kerr) terms then start to affect the measured frequency at higher powers still. The coupler's self-resonance frequency is also above the bare cavity frequency so the dressed frequency of the cavity is lower. The two levels of 'dressing' of the cavity frequency gives us information about the power levels where the qubit and coupler circuits become non-linear. As expected the qubit incurs non-linearity at much lower power levels. We were required to

perform these measurement quickly compared the qubit frequency drift, though so the data quality is sub-par. The difference in frequency in the lowest three powers shows some of that drift imprinted on the dispersive shift of the cavity frequency.

8.5 Future improvements

By moving the qubit spatially off of the WGM we think we could implement the full experimental idea much more readily than in the current setup. This would involve combining the ideas from Chapter 5 with the previous three chapters. By isolating the qubit, we could also gain from the fact that the Purcell effect would have a beneficial increase in T_1 times for the qubit. We discuss these improvement in the next, concluding, chapter. To implement the proposed improvement would require significant time and effort, though.

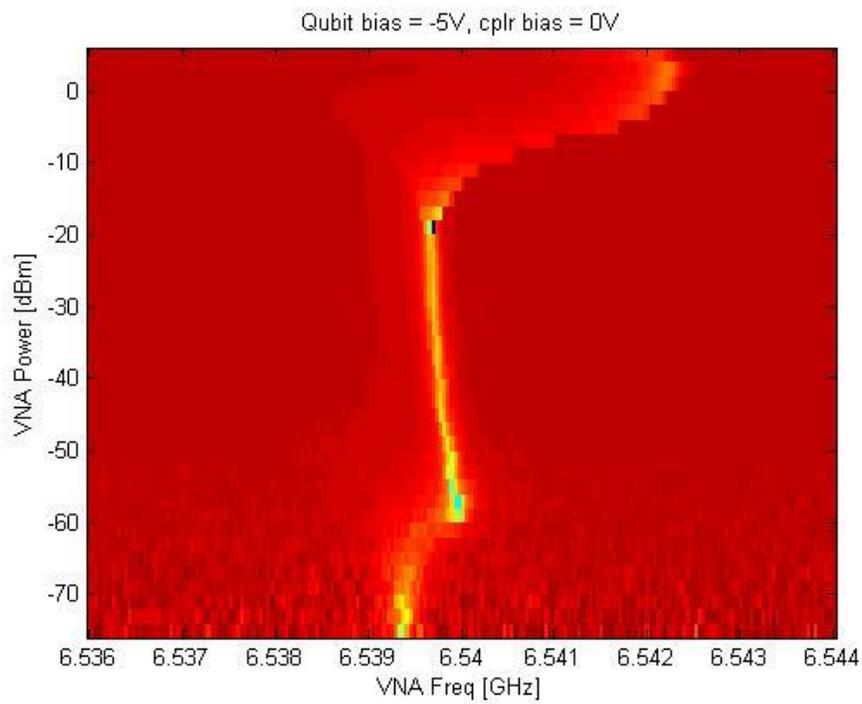


Figure 8.6: With both the qubit and coupler resonances above the 5.5 GHz WGM frequency, a 'double-Kerr' frequency vs power effect is shown. Driving at higher powers 'turns on' first the qubit's non-linear terms, then the coupler's.

Chapter 9

Conclusions

In conclusion, we would like to highlight the experimental results of this dissertation and offer some perspective about how they could be used in future quantum technologies. This thesis has described some of the practical aspects of quantum computing design which will become more and more important as the systems become more complex and more coherent. In particular we made use of parametric frequency conversion via Josephson-based circuits to connect photons between two frequencies and cavities.

Parametrically connecting a high-Q storage mode to a low-Q read/write mode allowed us to benefit from cavity designs isolated from their environment with long memory times, but also read out the information contained within efficiently. To demonstrate this we also demonstrated techniques useful for efficient *inter-cavity* coupling.

The sapphire whispering gallery mode resonator crystal also provides an extremely high-Q memory element in a compact space and allowed us to demonstrate the usefulness of *intra-cavity* coupling. Indeed, it provides *many* memory modes in which to store quantum information. The fact that we can parametrically couple them to each other makes this a useful technology for storing more complex quantum states. The WGM technology is a mature science, and these resonators are commercially viable for use in frequency standards and other precision measurements. In recent years larger, 5 cm diameter WGM crystals have shown even higher Qs than the 3 cm design in this work. The larger diameter crystals also push some of the useful frequencies down toward where our measurement setup is most

optimized. Thus these larger crystals could be a big improvement for storage times as well as ease of measurement.

In addition to using the sapphire resonators in new ways for storing coherent states, we demonstrated that the physics of cavity QED still apply and can be used to incorporate quantum devices and quantum states into this device. This will also be a necessary technique for later quantum technology. It is most likely that frequency-crowding will limit chip sizes unless the generated states can be moved from location to location on-chip or between cavities. In fact, this would allow experimentalists to exploit the competing design goals of creating isolated, highly-coherent, qubits with the need to interact the created states in other regions of the experiment.

An overlooked aspect of this experiment that I find attractive is the fact that we have mass-produced quantum systems and produced interactions by simply placing one atop the other. The promise of engineered quantum systems has often compared these systems to having ‘*LEGO*’ style building blocks which can be assembled in interesting and novel ways. This experiment certainly shows that this assembly is possible and can yield results easily.

Unfortunately, the final experimental assembly did not function as planned, though all of the individual pieces performed well. Along the way we mentioned improvements we would like to implement for later devices. The list of improvements is below, followed by a cartoon of a possible experimental setup which will avoid the problems faced in Chapter 8.

Use 5 cm WGM crystals:

- High-Q modes move lower in frequency.
- Electric fields spread out over larger circumference, allows for larger antennae for coupler chip.
- New designs not spindle-shaped, only one post at bottom – top surface is free for placing large chips.

- Larger area reduces some guess-work about where the fields are.

Coupler modifications:

- Lithographically define AC and DC bias coils to more accurately control pump and bias fields to coupler.
- Due to already reliable microfabrication parameters, could mass-produce various shapes of antennae to maximize dipole coupling and/or participation.
- Reliable fabrication means increasing the number of SQUIDs is possible, or changing critical currents to allow for stronger pump powers.

Qubit modifications:

- Larger SQUID area allows for direct sideband gating using the (large) qubit $\partial\omega/\partial\Phi$, rather than the smaller mode $\partial\omega/\partial\Phi$.
- More reliable optical or e-beam fabrication.
- Move qubit to separate, isolated, cavity: provides an ideal environment for the qubit, no Purcell effect – frequency convert, spatially, to the high-Q WGM crystal from there.

Explore other parametric interactions:

- Possible to pump at sum frequencies rather than difference frequencies to implement *in situ* amplification.
- Possible to pump two or more processes at once to implement other important interactions.

Other desiderata

- Incorporate Josephson Parametric Amplifier to output of the read/write mode to implement fully dispersive QND measurements, also allows for efficient state tomography.

Below, Figure 9.1 shows a possible architecture which could be used implementing, more reliably the ideas presented in this dissertation. Other groups have begun to think about this problem also [73, 74], but our approach is more modular and ‘LEGO-like’. We suggest that, in the short term, this will be a viable alternative to full-scale microfabrication solutions. Later, when industrial-level engineering and optimization, other ideas will certainly be necessary.

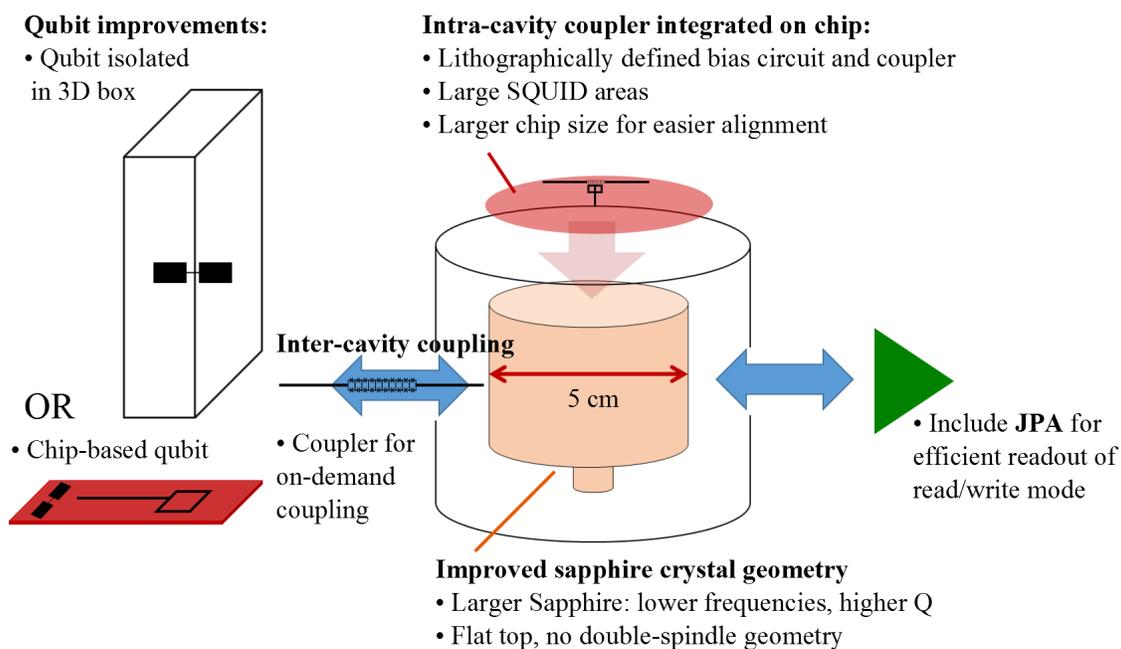


Figure 9.1: A sketch of the proposed improvements to the experiments shown in this thesis. Including a larger coupler + bias chip or wafer on top of a larger WGM crystal allows for coupling between even higher-Q modes in a more predictable manner. Moving the qubit to a different cavity (or chip) lets the qubit live in a well-isolated environment. And adding a JPA for readout would negate the need to do high-power Jaynes-Cummings readout.

Bibliography

- [1] M. A. Nielsen and I. L. Chuang. Quantum Computation and Quantum Information: 10th Anniversary Edition. Cambridge Univ. Press, 2011.
- [2] M. H. Devoret and R. J. Schoelkopf. Superconducting circuits for quantum information: An outlook. Science, 2013.
- [3] S. Haroche and J-M. Raimond. Exploring the Quantum: Atoms, Cavities, and Photons. Oxford University Press, 2006.
- [4] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O’Brien. Quantum computers. Nature, 464(08812), Mar 2010.
- [5] C. Kloeffel and D. Loss. Prospects for spin-based quantum computing in quantum dots. Annual Review of Condensed Matter Physics, 4:41–81, Apr 2013.
- [6] P. Kok, W. J. Munro, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn. Linear optical quantum computing with photonic qubits. Reviews of Modern Physics, 79(135), Jan 2007.
- [7] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. Quantum dynamics of single trapped ions. Reviews of Modern Physics, 75(281), Mar 2003.
- [8] P. S. Jessen, I. H. Deutsch, and R. Stock. Quantum information processing with trapped neutral atoms. Quantum Information Processing, 3:91–103, Oct 2004.
- [9] L. M. K. Vandersypen, M. Steffen, G. Breyta, C. S. Yannoni, M. H. Sherwood, and I. L. Chuang. Experimental realization of shor’s quantum factoring algorithm using nuclear magnetic resonance. Nature, 414:883–887, Dec 2001.
- [10] L. Childress and R. Hanson. Diamond nv centers for quantum computing and quantum networks. MRS Bulletin, 38(02):134–138, Feb 2013.
- [11] O. Dial, D. T. McClure, S. Poletto, J. M. Gambetta, D. W. Abraham, J. M. Chow, and M. Steffen. Bulk and surface loss in superconducting transmon qubits. Superconductor Science and Technology, 29(4), 2016.

- [12] M. Reagor, W. Pfaff, C. Axline, R. W. Heeres, N. Ofek, K. Sliwa, E. Holland, C. Wang, J. Blumoff, K. Chou, M. J. Hatridge, L. Frunzio, M. H. Devoret, L. Jiang, and R. J. Schoelkopf. A quantum memory with near-millisecond coherence in circuit qed. [arXiv:1508.05882 \[quant-ph\]](https://arxiv.org/abs/1508.05882), 2015.
- [13] I. Buluta, S. Ashhab, and F. Nori. Natural and artificial atoms for quantum computation. Reports on Progress in Physics, 2011.
- [14] B. Yurke, L. W. Rupp, and P. G. Kaminsky. Use of the josephson junction in fundamental quantum-electronics experiments. IEEE Transactions on Magnetics, 23(2):458–464, 1987.
- [15] B. Yurke, L. R. Corruccini, P. G. Kaminsky, L. W. Rupp, A. D. Smith, A. H. Silver, R. W. Simon, and E. A. Whittaker. Observation of parametric amplification and deamplification in a josephson parametric amplifier. Physical Review A, 39(5):2519–2533, 1989.
- [16] J. Clark and A. I. Braginski, editors. The SQUID Handbook Vol 1: Fundamentals and Technology of SQUIDs and SQUID systems. Wiley-VCH, 2004.
- [17] E. Schrodinger. Are there quantum jumps? part 1. The British Journal for the Philosophy of Science, 3(10):109–123, 1952.
- [18] E. Schrodinger. Are there quantum jumps? part 2. The British Journal for the Philosophy of Science, 3(11):233–242, 1952.
- [19] D. B. Sullivan. Time and frequency measurement at nist: the first 100 years. Online, 2001. <http://tf.nist.gov/timefreq/general/pdf/1485.pdf>.
- [20] T. Van Duzer and C. W. Turner. Principles of Superconductive Devices and Circuits. Prentice Hall, 2 edition, 1998.
- [21] J. M. Martinis, S. Nam, J. Aumentado, and C. Urbina. Rabi oscillations in a large josephson-junction qubit. Physical Review Letters, 89(117901), 2002.
- [22] M. H. Devoret and J. M. Martinis. Implementing qubits with superconducting integrated circuits. Quantum Information Processing, 3(1):163–203, 2004.
- [23] J. M. Martinis, K. B. Cooper, R. McDermott, M. Steffen, M. Ansmann, K. Osborn, K. Cicak, S. Oh, D. P. Pappas, R. W. Simmonds, and C. C. Yu. Decoherence in josephson qubits from dielectric loss. Physical Review Letters, 95(210503), 2005.
- [24] J. D. Teufel, D. Li, M. S. Allman, and J. D. Whittaker K. Cicak, A. J. Sirois, and R. W. Simmonds. Circuit cavity electromechanics in the strong-coupling regime. Nature, 471(7337):204–208, 2011.
- [25] J. D. Teufel, T. Donner, D. Li, J. W. Harlow, M. S. Allman, K. Cicak, A. J. Sirois, J. D. Whittaker, K. W. Lehnert, and R. W. Simmonds. Sideband cooling of micromechanical motion to the quantum ground state. Nature, 475(7356):359–363, 2011.

- [26] J. Koch, T. M. Yu, J. M. 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 from optimizing the cooper-pair box. Physical Review A, 76(042319), 2007.
- [27] 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. Physical Review Letters, 107(240501), 2011.
- [28] D. Loss and D. P. DiVincenzo. Quantum computation with quantum dots. Physical Review A, 57(120), 1998.
- [29] D. M. Pozar. Microwave Engineering. Wiley, 4th edition, 2011.
- [30] M. Neeley, M. Ansmann, R. C. Bialczak, M. Hofheinz, N. Katz, E. Lucero, A. OConnell, H. Wang, A. N. Cleland, and J. M. Martinis. Transformed dissipation in superconducting quantum circuits. Physical Review B, 77(180508(R)), 2008.
- [31] K. Geerlings, S. Shankar, E. Edwards, L. Frunzio, R. J. Schoelkopf, and M. H. Devoret. Improving the quality factor of microwave compact resonators by optimizing their geometrical parameters. Applied Physics Letters, 100(192601), 2012.
- [32] H. Paik and K. D. Osborn. Reducing quantum-regime dielectric loss of silicon nitride for superconducting quantum circuits. Applied Physics Letters, 96(072505), 2010.
- [33] P. J. Petersan and S. M. Anlage. Measurement of resonant frequency and quality factor of microwave resonators: Comparison of methods. Journal of Applied Physics, 84(3392), 1998.
- [34] V. B. Braginsky, V. P. Mitrofanov, and V. I. Panov. Systems with Small Dissipation. The University of Chicago Press, 1985.
- [35] J. P. Turneaure and I. Weissman. Microwave surface resistance of superconducting niobium. Journal of Applied Physics, 39(4417), 1968.
- [36] Bruno Sanguinetti. Building a Modern micromaser: Atoms and Cavities. PhD thesis, The University of Leeds School of Physics and Astronomy, 2009.
- [37] K. J. Vahala. Optical microcavities. Nature, 424:839–846, 2003.
- [38] J. G. Hartnett, C. R. Locke, E. N. Ivanov, M. E. Tobar, and P. L. Stanwix. Cryogenic sapphire oscillator with exceptionally high long-term frequency stability. Applied Physics Letters, 89(203513), 2006.
- [39] Daniel L. Creedon. Effects of Electron Sping Resonance in a Cryogenic Sapphire Whispering Gallery Mode Maser. PhD thesis, Univ. of Western Australia, 2011.

- [40] D. L. Creedon, Y. Reshitnyk, W. Farr, J. M. Martinis, T. L. Duty, and M. E. Tobar. High q-factor sapphire whispering gallery mode microwave resonator at single photon energies and millikelvin temperatures. Applied Physics Letters, 98(222903), 2011.
- [41] P. Meystre and M. Sargent III. Elements of Quantum Optics. Springer Verlag, 2007.
- [42] J. Bardeen, L. N. Cooper, and J. R. Schrieffer. Microscopic theory of superconductivity. Physical Review, 106(162), April 1957.
- [43] B. D. Josephson. The discovery of tunnelling supercurrents. Reviews Modern Physics, 46(251), April 1974.
- [44] M. Mirrahimi, Z. Leghtas, V. V. Albert, S. Touzard, R. J. Shoelkopf, L. Jiang, and M. H. Devoret. Dynamically protected cat-qubits: a new paradigm for universal quantum computing. New Journal of Physics, 16(045014), 2014.
- [45] D. F. Walls and G. J. Milburn. Quantum Optics. Springer-Verlag, 2nd edition edition, 2008.
- [46] B. Yurke. Quantum Squeezing. Springer, 2004.
- [47] E. Zakka-Bajjani, F. Nguyen, M. Lee, L. R. Vale, R. W. Simmonds, and J. Aumentado. Quantum superposition of a single microwave photon in two different colour states. Nature Physics, 7:599603, 2011.
- [48] S. A. Campbell. The Science and Engineering of Microelectronic Fabrication. Oxford, 2001.
- [49] J. D. Teufel, F. Lecocq, and R. W. Simmonds. Overwhelming thermomechanical motion with microwave radiation pressure shot noise. Physical Review Letters, 116(013602), 2016.
- [50] J. B. Clark, F. Lecocq, R. W. Simmonds, J. Aumentado, and J. D. Teufel. Observation of strong radiation pressure forces from squeezed light on a mechanical oscillator. Nature Physics, (online: 1601.02689), 2016.
- [51] F. Lecocq, J. B. Clark, R. W. Simmonds, J. Aumentado, and J. D. Teufel. Mechanically mediated microwave frequency conversion in the quantum regime. Physical Review Letters, 116(043601), 2016.
- [52] L. Ranzani and J. Aumentado. Graph-based analysis of nonreciprocity in coupled-mode systems. New Journal of Physics, (023024), 2015.
- [53] A. J. Sirois, M. A. Castellanos-Beltran, M. P. DeFeo, L. Ranzani, F. Q. Lecocq, R. W. Simmonds, J. D. Teufel, and J. Aumentado. Coherent-state storage and retrieval between superconducting cavities using parametric frequency conversion. Applied Physics Letters, 106(17), 2015.

- [54] JD Strand, Matthew Ware, Félix Beaudoin, TA Ohki, BR Johnson, Alexandre Blais, and BLT Plourde. First-order sideband transitions with flux-driven asymmetric transmon qubits. Physical Review B, 87(22):220505, 2013.
- [55] Michael Shane Allman, Jed D Whittaker, Manuel Castellanos-Beltran, Katarina Cicak, Fabio da Silva, MP DeFeo, Florent Lecocq, Adam Sirois, JD Teufel, Jose Aumentado, and RW Simmonds. Tunable resonant and nonresonant interactions between a phase qubit and l c resonator. Physical Review Letters, 112(12):123601, 2014.
- [56] B. Abdo, K. Silwa, F. Schackert, N. Bergeal, M. Hatridge, L. Frunzio, A. D. Stone, and M. Devoret. Full coherent frequency conversion between two propagating microwave modes. Physical Review Letters, 110(173902), 2013.
- [57] E. Flurin, N. Roch, J.D. Pillet, F. Mallet, and B. Huard. Superconducting quantum node for entanglement and storage of microwave radiation. Physical Review Letters, 114(090503), 2015.
- [58] W. H. Louisell. Coupled mode and parametric electronics. Wiley, 1960.
- [59] Martin Frimmer and Lukas Novotny. The classical bloch equations. American Journal of Physics, 82(10):947–954, 2014.
- [60] Michael Tobar. Private Communication.
- [61] D. I. Schuster, A. P. Sears, E. Ginossar, L. DiCarlo, L. Frunzio, J. J. L. Morton, H. Wu, G. A. D. Briggs, and R. J. Schoelkopf. High cooperativity coupling of electron-spin ensembles to superconducting cavities. Physical Review Letters, 105(140501), 2010.
- [62] Y. Kubo C. Grezes and, B. Julsgaard, T. Umeda, J. Isoya, H. Sumiya, H. Abe, S. Onoda, T. Ohshima, K. Nakamura, I. Diniz, A. Auffeves, V. Jacques, J.-F. Roch, D. Vion, D. Esteve, K. Moelmer, and P. Bertet. Towards a spin-ensemble quantum memory for superconducting qubits. ArXiv, 2015.
- [63] E. M. Purcell. Resonance absorption by nuclear magnetic moments in a solid. Physical Review, 69(37), 1946.
- [64] D. Kleppner. Inhibited spontaneous emission. Physical Review Letters, 47(4), 1981.
- [65] L. S. Bishop. Circuit Cavity Electrodynamics. PhD thesis, Yale University, 2010.
- [66] D. Shuster. Circuit Cavity Quantum Electrodynamics. PhD thesis, Yale University, 2007.
- [67] L. S. Bishop, E. Ginossar, and S. M. Girvin. Response of the strongly driven jaynes-cummings oscillator. Physical Review Letters, 105(100505), 2010.
- [68] Maxime Boissonneault, J. M. Gambetta, and Alexandre Blais. Improved superconducting qubit readout by qubit-induced nonlinearities. Physical Review Letters, 105(100504), 2010.

- [69] M. D. Reed, L. DiCarlo, B. R. Johnson, L. Sun, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf. High-fidelity readout in circuit quantum electrodynamics using the jaynes-cummings nonlinearity. Phys. Rev. Lett., (173601), 2010.
- [70] R. Schutjens, F. Abu Dagga, D. J. Egger, and F. K. Wilhelm. Single qubit gates in frequency-crowded transmon systems. Physical Review A, 88(052330), 2013.
- [71] F. Beaudoin, M. P. da Silva, Z. Dutton, and A. Blais. First-order sidebands in circuit qed using qubit frequency modulation. Physical Review A, 86(022305), 2012.
- [72] Manuel Castellanos-Beltran et al. Private Communication.
- [73] T. Brecht, W. Pfaff, C. Wang, Y. Chu, L. Frunzio, M. H. Devoret, and R. J. Schoelkopf. Multilayer microwave integrated quantum circuits for scalable quantum computing. ArXiv:1509.01127, 2015.
- [74] T. Brecht, M. Reagor, Y. Chu, W. Pfaff, C. Wang, L. Frunzio, M. H. Devoret, and R. J. Schoelkopf. Demonstration of superconducting micromachined cavities. ArXiv:1509.01119, 2015.

Appendix A

List of relevant publications:

- Coherent-state storage and retrieval between superconducting cavities using parametric frequency conversion, AJ Sirois, MA Castellanos-Beltran, MP DeFeo, L Ranzani, F Lecocq, RW Simmonds, JD Teufel, J Aumentado, Applied Physics Letters 106 (17), 172603 1 2015
- Realization of a single-Cooper-pair Josephson laser, Fei Chen, Juliang Li, AD Armour, E Brahim, Joel Stettenheim, AJ Sirois, RW Simmonds, MP Blencowe, AJ Rimberg, Physical Review B 90 (2), 020506 17 2014
- Tunable-cavity QED with phase qubits, Jed D Whittaker, FCS da Silva, Michael Shane Allman, Florent Lecocq, Katarina Cicak, AJ Sirois, JD Teufel, Jose Aumentado, Raymond W Simmonds, Physical Review B 90 (2), 024513 7 2014
- Tunable Resonant and Nonresonant Interactions between a Phase Qubit and L C Resonator, Jed D Whittaker, FCS da Silva, Michael Shane Allman, Florent Lecocq, Katarina Cicak, AJ Sirois, JD Teufel, Jose Aumentado, Raymond W Simmonds, Physical review letters 112 (12), 123601 8 2014
- Sideband cooling of micromechanical motion to the quantum ground state, JD Teufel, Tobias Donner, Dale Li, JW Harlow, MS Allman, Katarina Cicak, AJ Sirois, Jed D Whittaker, KW Lehnert, Raymond W Simmonds, Nature 475 (7356), 359-363 807 2011

- Introduction of a dc bias into a high-Q superconducting microwave cavity, F Chen, AJ Sirois, RW Simmonds, AJ Rimberg, Applied Physics Letters 98 (13), 132509 20 2011
- Circuit cavity electromechanics in the strong-coupling regime, JD Teufel, Dale Li, MS Allman, K Cicak, AJ Sirois, JD Whittaker, RW Simmonds, Nature 471 (7337), 204-208 405 2011
- Tripartite interactions between two phase qubits and a resonant cavity, F Altomare, JI Park, K Cicak, MA Sillanp, MS Allman, D Li, A Sirois, JA Strong, JD Whittaker, RW Simmonds, Nature Physics 6 (10), 777-781 32 2010
- Measurement crosstalk between two phase qubits coupled by a coplanar waveguide, Fabio Altomare, Katarina Cicak, Mika A Sillanp, Michael S Allman, Adam J Sirois, Dale Li, Jae I Park, Joshua A Strong, John D Teufel, Jed D Whittaker, Raymond W Simmonds, Physical Review B 82 (9), 094510 6 2010
- rf-SQUID-mediated coherent tunable coupling between a superconducting phase qubit and a lumped-element resonator, Michael S Allman, Fabio Altomare, JD Whittaker, K Cicak, D Li, A Sirois, J Strong, JD Teufel, RW Simmonds, Physical review letters 104 (17), 177004 45 2010
- Time-Domain Detection of Weakly Coupled TLS Fluctuators in Phase Qubits, MS Allman, F Altomare, JD Whittaker, K Cicak, D Li, A Sirois, J Strong, JD Teufel, RW Simmonds, arXiv preprint arXiv:1004.2738 2 2010
- Low-loss superconducting resonant circuits using vacuum-gap-based microwave components, Katarina Cicak, Dale Li, Joshua A Strong, Michael S Allman, Fabio Altomare, Adam J Sirois, Jed D Whittaker, John D Teufel, Raymond W Simmonds, Applied Physics Letters 96 (9), 093502 25 2010

- Coherent interactions between phase qubits, cavities, and TLS defects, RW Simmonds, MS Allman, F Altomare, K Cicak, KD Osborn, JA Park, M Sillanp, A Sirois, JA Strong, JD Whittaker, Quantum Information Processing 8 (2-3), 117-131 16 2009
- Frequency-tunable Josephson junction resonator for quantum computing KD Osborn, JA Strong, AJ Sirois, RW Simmonds, Applied Superconductivity, IEEE Transactions on 17 (2), 166-168 31 2007