ABSTRACT

Title of dissertation: MULTI-JUNCTION EFFECTS
IN DC SQUID PHASE QUBITS

Benjamin K. Cooper,
Doctor of Philosophy, 2013

Dissertation directed by: Professor Frederick C. Wellstood
Department of Physics

I discuss experimental and theoretical results on an LC filtered dc SQUID phase qubit. This qubit is an asymmetric aluminum dc SQUID, with junction critical currents 1.5 and 26.8 μA, on a sapphire substrate. The layout differs from earlier designs by incorporating a superconducting ground plane and weakly coupled coplanar waveguide microwave drive line to control microwave-qubit coupling.

I begin with a discussion of quantizing lumped element circuit models. I use nodal analysis to construct a 2d model for the dc SQUID phase qubit that goes beyond a single junction approximation. I then discuss an extension of this "normal modes" SQUID model to include the on-chip LC filter with design frequency ≈ 180 MHz. I show that the filter plus SQUID model yields an effective Jaynes-Cummings Hamiltonian for the filter-SQUID system with coupling $g/2\pi \approx 32$ MHz.

I present the qubit design, including a noise model predicting a lifetime $T_1 = 1.2$ μs for the qubit based on the design parameters. I characterized the qubit with measurements of the current-flux characteristic, spectroscopy, and Rabi oscillations.

I measured $T_1 = 230$ ns, close to the value 320 ns given by the noise model using the measured parameters. Rabi oscillations show a pure dephasing time $T_\phi = 1100$ ns. The spectroscopic and Rabi data suggest two-level qubit dynamics are inadequate for describing the system. I show that the effective Jaynes-Cummings model reproduces some of the unusual features.

MULTI-JUNCTION EFFECTS IN DC SQUID PHASE QUBITS

by

Benjamin K. Cooper

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2013

Advisory Committee:
Professor Frederick C. Wellstood, Chair/Advisor
Professor J. Robert Anderson
Professor Steven Anlage
Professor Chris J. Lobb
Professor Ichiro Takeuchi

For Kenzie and Iris

Acknowledgments

First and foremost, thanks to Fred Wellstood for serving as my thesis advisor. His even keel and extensive knowledge were both very helpful in facing the slings and arrows of superconducting qubit research, and his considerable patience was deeply appreciated. Chris Lobb raised interesting questions, offered good jokes, and was surprisingly eager to clamber around the dilution refrigerator when I was repairing it. Bob Anderson sent many articles and brought many visitors over the years to help keep the subbasement abreast of quantum information more generally. I would also like to thank Steve Anlage and Ichiro Takeuchi for serving as the remaining members of my committee.

Thanks to my co-namesakes of the BBC2 qubit Rangga Budoyo and Cody Ballard. Rangga, especially, was very tolerant of my coming on board with this qubit, which began as his design. Cody provided a great burst of new energy to the project and was very helpful during the period we were working on the refrigerator. During my time in the Wellstood lab, I overlapped with many other graduate students, including Sudeep Dutta, Hanhee Paik, Tauno Palomaki, Tony Przybysz, and Hyeokshin Kwon. While I spent much of my time with them pursuing my own project, they were an endless source of conversation and new ideas. Our former postdoc Rupert Lewis taught me a great deal about experimental work when I first started in the lab.

Thanks also to our collaborators at LPS. Ben Palmer and Kevin Osborn and their groups — including Zaeill Kim, Sergey Novikov, Bala Suri, Sergiy Gladchenko,

Micah Stoutimore and Moe Khalil — were another fertile source of discussions, ideas and techniques in superconducting qubits. I would especially like to thank Vitaley Zaretskey. We started performing e-beam lithography for BBC2 on an SEM in the Physics Building. After constantly competing for limited time on this machine, Vitaley suggested coming over to LPS to use the JEOL e-beam system there. He instructed me in the operation of the microscope, offered helpful suggestions about lithography more generally, and always had time to help diagnose problems. His support was instrumental in the building of BBC2.

The support staff in CNAM and the Physics Department have been extremely helpful over the years. Doug Bensen and Brian Straughn always had excellent advice on technical problems. Jesse Anderson was a great help in purchasing, and Al Godinez was very helpful in keeping the cryogens rolling into the lab. I owe Jane Hessing a great deal for her constant behind the scenes handling of various administrative issues.

Thanks to Andy York and Sanjay Varma for general physics chat and helping to keep the chin up over the years. Thanks to my wife MacKenzie. Waiting for me to finish this thesis has been trying for her — especially as we conclude our first year of parenting — and I thank her for her support, forbearance, and love. Lastly, thanks to my daughter Iris. Her constant joy is the greatest thing.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

The last twenty years have seen an immense amount of research in quantum computing. The promise of devices with previously unattainable computing power, along with the almost unlimited range of systems capable of being cast as two-level systems, has made the field broadly attractive. My thesis work was on understanding and improving the performance of a particular type of superconducting qubit. After a brief general discussion of quantum computing, I give an overview of the thesis.

1.1 Quantum computing

Classical computing concerns itself with bits, the 0s and 1s that represent the simplest possible unit of information. A bit can be set to either 0 or 1. Quantum computing, by contrast, asks what you can do if you exploit the superposition principle of quantum mechanics to get qubits that can be a mixture of $|0\rangle$ and $|1\rangle$. A talk by Feynman on the limits of simulating quantum mechanics with classical computers [1] is often considered the impetus for the idea, although there were earlier examinations of some of the ideas of quantum information, such as Wiesner's proposal for quantum money [2]. Feynman's point was that solving Schrödinger's equation required exponentially increasing resources as the system size increased. Trying to perform direct simulation of quantum systems — similar to building analog circuits with op-amps to solve particular differential equations — cannot be done classically, as Bell's inequality shows. By moving to quantum mechanical parts, a simulation could be performed.

In addition to this quantum simulation idea, people began to discover that quantum information could be exploited to perform other tasks better than classically possible. One early example was the Deutsch-Jozsa algorithm [3], demonstrating a quantum algorithm that outperformed any classical one, albeit on a very artificial problem. A more natural seeming application was superdense coding [4], where one spin 1/2 particle could be measured to extract two bits of classical information. By far the most famous quantum algorithm, though, is Shor's algorithm for factoring numbers [5]. The best approaches known for factoring numbers classically are exponential in the size of the number; Shor's algorithm is polynomial. Given that difficulty in factoring large numbers is the linchpin of the widely used public-key cryptosystem RSA, Shor's algorithm represents an amazing performance gain in a technologically important problem. Its existence catalyzed the broad interest in quantum computing.

A quantum computer would be of great practical interest, but they are still largely theoretical constructions. To reap the benefit of quantum algorithms, the qubits must behave quantum mechanically. Through coupling to an unmonitored environment, quantum mechanical pure states eventually decay to states analogous to classical mixtures, a process called decoherence [6]. Although it is possible to correct for errors introduced by decoherence [7], a minimum coherence threshold is still necessary for error correction to operate.

The coherence issue presents a dilemma. One may work with systems that are very weakly coupled to the environment. This gives excellent coherence, but leads to challenges in measurement and coupling multiple qubits. Alternatively, one can work with devices that are easy to measure and couple, but the stronger coupling that simplifies these tasks gives worse coherence. Many different physical systems have been proposed as qubits. Some examples include: donor spins in silicon [8]; polarization modes of photons, coupled using linear optics [9]; proton

spins of molecules, controlled using liquid-state NMR [10]; braids of quantum Hall quasiparticles [11]; electronic states of trapped ions, coupled to motional states [12]. Each provides a different set of benefits and challenges.

Superconducting circuits have been used to realize several different types of qubit. The precursor to superconducting qubits was the observation of energy level quantization in a Josephson junction [13]; this type of qubit, where the states correspond to resonances in tilted washboard potential of the Josephson junction, is called a *phase qubit* [14]. A different approach, and the first modern superconducting qubit, was the Cooper pair box of Nakamura et al [15]. In this *charge qubit*, the states correspond to the presence or absence of an excess Cooper pair on the box. Shortly after that experiment, a superposition of flux states of a SQUID was demonstrated [16]. This is the approach of *flux qubits*.

These different qubit designs correspond to the differing proportions of the charging energy E_C of the Josephson junctions to the Josephson energy E_J associated with the tunneling current through the junction. At one extreme is the charge qubit, where $E_C \gg E_J$. This results in a Coulomb blockade, where charge enters and leaves the box one Cooper pair at a time. At the other end is the phase qubit, where $E_J \gg E_C$ and instead the current through the junction is the better defined quantum number. Some modern designs like the transmon [17] operate in a hybrid regime between these two poles.

The focus of our lab has been phase qubits. To examine the strengths and weaknesses of the phase qubit, it is useful to consider what a qubit must do. To be practical, a qubit must satisfy the DiVincenzo criteria [18]. How well do phase qubits satisfy the criteria, quoted from that article?

- *A scalable physical system with well-characterized qubits*: One of the great appeals of solid-state approaches like superconducting qubits is that scalability seems very plausible. Lithographic and miniaturization techniques from the

semiconductor industry seem well suited to eventually placing many qubits on a chip. Phase qubits in particular seem well adapted to scaling up. A key feature of the phase qubit is the tunability of ω_{01} by changing the qubit bias. This would greatly improve yield for multi-qubit devices, as e.g. small local parameter fluctuations in the fabrication process would not break qubit coupling.

"Well-characterized" is used to indicate that the Hamiltonian of the qubits themselves — including coupling between qubits and to external control fields — is well understood. For simpler phase qubits, this holds, although as I will discuss later in this thesis some circuits can become quite complicated in their behavior.

- *The ability to initialize the system to a simple fiducial state, such as 000*: Any useful quantum computation will require the ability to prepare the qubits in a given state. To do this efficiently, it is necessary to start in a known state with high probability. The suggested state $|000\ldots\rangle$ is a natural one, as it is the ground state of the system. Most superconducting qubit experiments are performed in dilution refrigerators at temperatures k_BT 10-20 times lower than $\hbar\omega_{01}$. At these temperatures, barring excess noise specifically at ω_{01}, the Boltzmann factor leads to an overwhelming likelihood of being in the ground state, provided the system has had time to relax thermally[1].

- *A "universal" set of quantum gates*: The notion of universality comes from classical computing [19]. If we say that a particular "logical function" returns 0 or 1 given any choice of input bits, then a set of gates is universal if all possible logical functions can be built using those gates. The quantum analogue

[1] In a system with many qubits, although the ground state is still more probable than any given one qubit excited state, eventually a point will be reached where it is more probable that at least one qubit is excited than not. Since for the parameters discussed this requires somewhere between $\approx 10^5$ and 10^9 qubits, that would be a good problem to have.

involves being able to prepare any state in the Hilbert space generated by the qubits. A well-known result from quantum information theory [20] states that one-qubit gates (i.e. arbitrary unitary operations on a single qubit) and the two-qubit gate CNOT are universal. CNOT is short for "controlled NOT," where one qubit is the control and one is the target. A CNOT gate inverts the target qubit if the control qubit is in the 1 state:

Input	Output
$\|00\rangle$	$\|00\rangle$
$\|01\rangle$	$\|01\rangle$
$\|10\rangle$	$\|11\rangle$
$\|11\rangle$	$\|10\rangle$

There are other two-qubit gates that can be used in lieu of CNOT to yield universality. The key point is the notion of a controlled gate. This is a technique that often facilitates the creation of reversible gates. Reversibility is important since quantum mechanics is itself a reversible theory, and thus any prospective gate must be reversible. CNOT is logically reversible because the output is sufficient to recreate the input. By contrast, the classical AND gate is not reversible; knowing that AND on two inputs yields 0 is not sufficient to determine the inputs.

The Rabi oscillation control system described is sufficient (provided one uses both in-phase and quadrature fields) to perform arbitrary unitaries on a single qubit. Gate designs such as a controlled phase gate [21] are workable on two capacitively coupled phase qubits, allowing universal computation.

- *A qubit-specific measurement capability*: Quantum tunneling through the potential barrier provides a method for phase qubit readout. A general result from WKB theory is that the tunneling rate is exponentially sensitive to the

height of the barrier, as the decay of the wavefunction under the barrier is set by this energy. Thus, there is a great enhancement in the escape rate when the qubit is in the $|1\rangle$ state. This tunneling leads to a measurable voltage across the qubit that tunnels. If each qubit has its own set of leads, then each can be measured individually.

While there are some advantages to this approach, a major shortcoming is that tunneling based measurements are not projective. Tunneling may be a strong indicator that the system was in $|1\rangle$, but following the tunneling event the junction isn't even superconducting, let alone in the original logical basis. This type of measurement, then, can be quite useful for initial diagnostics, but poses a severe hurdle for key quantum information tasks such as error correction (which rely on projective measurement.)

- *Long relevant decoherence times, much longer than the gate operation time*: At the moment, this is the great challenge of Josephson qubits, and phase qubits in particular. The relative ease of coupling and measurement comes at the expense of relatively strong decoherence, which is nothing more than relatively easy coupling to other, unmeasured quantum systems. "Much longer" is quantified by the threshold theorem [22], which roughly says that quantum error correction can work well enough to allow computation provided the error rate is sufficiently low per gate operation. The particular threshold depends strongly on the choice of code and architecture. A coherence time 10^4 longer than the gate time was an oft-used number, but some codes [23] can approach 10^2. This latter number is feasible in some modern superconducting qubit setups.

Since decoherence is driven by coupling to an uncontrolled environment, building a better qubit is largely about controlling the qubit's environment as rigorously

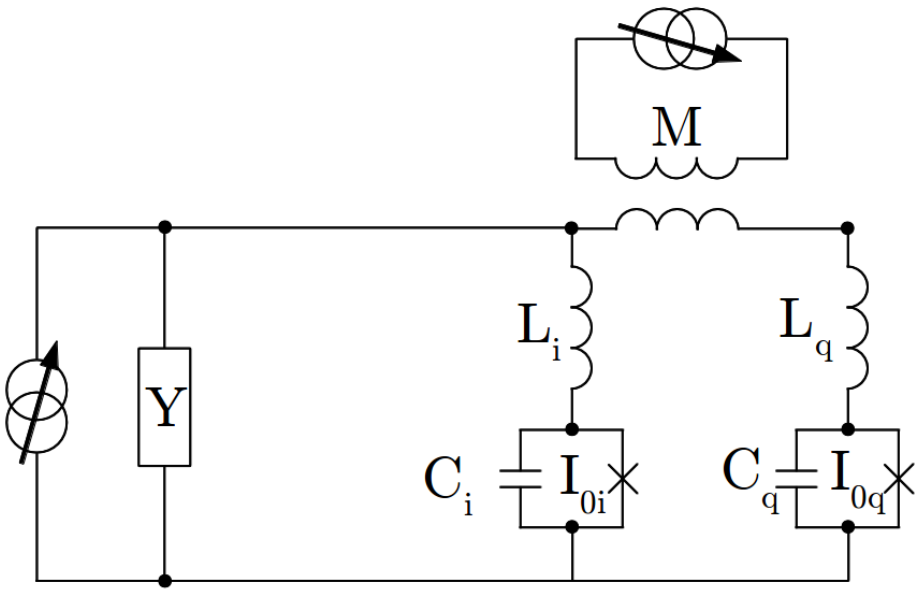

Figure 1.1: Schematic of the dc SQUID phase qubit, including bias line admittance Y. I discuss dc SQUID phase qubit design in Section 2.5.

as possible. This demands circuit designs more complicated than a single current biased Josephson junction. Most of the research in our lab has focused on using dc SQUIDs as phase qubits. A schematic is shown in Figure 1.1. Originally suggested by Martinis et al. [24], the idea is to provide an inductive filtering network that shields a qubit junction from noise on the bias line. To continue biasing the junction as before, a flux coil is added to modulate the SQUID.

Our lab has produced several theses [25–29] devoted to the study of the dc SQUID phase qubit. Over this time, the design has changed. We now employ smaller junctions, especially for the intended qubit junction. These changes have been for the better. Over the years, we have improved the coherence times by an order of magnitude. Yet a consistent puzzle is why the improvement has not been even better.

My work on dc SQUID phase qubits had two chief aspects. One aspect was experimental. With Rangga Budoyo, I fabricated our most recent dc SQUID phase qubit "BBC2." I then performed measurements on it to characterize its behavior. The lifetime $T_1 = 230$ ns of the qubit was comparable to the most recent qubits our lab has studied, and the phase coherence (which I characterize by the pure dephasing time $T_\phi = 1100$ ns) appears substantially better than any phase qubit we had previously measured.

The BBC2 data showed features that were unexpected. This led into the second aspect of my work, which was improving the theoretical models used to describe these qubits. In the original model, the dc SQUID phase qubit behaved as a single junction phase qubit with a filtering network attached to the front of it. For some of the older devices we studied this was a good approximation, but as I show, the move to smaller junctions makes the dynamics of the full circuit more important. I developed models that incorporated more of the circuit into the analysis — first the SQUID, and later the on-chip LC filter as well.

1.2 Overview of thesis

Chapter 2 examines much of the basic physics underlying the dc SQUID phase qubit. It begins with a derivation of the Josephson effect. Following this, I explain how network theory can be applied to give the appropriate quantum mechanical Hamiltonian for lumped element circuit designs. I apply this technique to several systems of increasing complexity and show that it yields sensible results. This network theory approach considers a lossless system, but to capture the all-important decoherence phenomena it is essential to include dissipation. I give a brief overview of loss in quantum systems and its application to our qubit.

The final section of Chapter 2 discusses the model I developed for the dc SQUID phase qubit. The modern design featuring small, heavily asymmetric junctions was arrived at after years of tinkering. Earlier models of the behavior of the SQUID treated it as a single junction with a classical filtering network in front of it. I explain the limitations of this approach and apply the nodal analysis formalism to derive a model for the full SQUID. For many applications, the normal modes of this model Hamiltonian — corresponding to a harmonic approximation to the qubit behavior — are sufficient. This model was found to successfully describe similar designs studied by previous students in our lab [28, 29].

Chapter 3 begins with an explanation of the Jaynes-Cummings model. Consisting of a linear coupling between a two-level system and a harmonic oscillator, it is a simple, exactly solvable model that has nevertheless proved highly successful in describing the interaction of two-level systems with light. In particular, it is a good descriptor of cavity QED experiments, where the field can be limited to a particular relevant mode. One of the most influential ideas in superconducting quantum computing has been the extension of this idea to Josephson junction circuits. In this context, the techniques have come to be known as circuit QED [30].

In circuit QED, a cavity in a coplanar waveguide is coupled to the qubit. The

interplay between cavity and qubit is then exploited to make dispersive microwave measurements of the qubit. The LC filtered dc SQUID phase qubit I studied does not have this waveguide cavity. However, the LC filter is itself a lumped element microwave resonator. Using the techniques of Chapter 2, I expand the SQUID model to include the LC filter. I then extract an effective Jaynes-Cummings model for this filter plus SQUID system.

Chapter 4 discusses the fabrication of BBC2. The initial design for BBC2 was done by Rangga Budoyo. I later made some modifications to the design, and Rangga and I jointly did the fabrication. For several years all of the qubits our lab has built have used aluminum deposited on sapphire substrates. Originally we only used photolithographic processes, but a move toward smaller junctions compelled us to begin using e-beam lithography as well. Sapphire presents a problem for e-beam lithography, as it accumulates charge if nothing is done to prevent this. I discuss how I modified an anti-charging e-beam process originally developed by Przybysz [28] for BBC2 and mention some of the difficulties encountered.

Chapter 5 discusses the experimental setup. Since our qubits are superconducting devices it is obviously essential to cool them to the point where they superconduct. Beyond that, we rely on thermal relaxation to prepare the system in the $|0\rangle$ state. Thus, it is beneficial to work at the lowest temperature possible. Accordingly, I mounted the qubit in a dilution refrigerator with base temperature 20 mK. Although this refrigerator has been used for several generations of experiments, repairs and wiring changes for BBC2 led to changes in the layout of the fridge. I begin by reviewing these changes.

Chapter 5 also examines some techniques I used in the measurements. One big issue in the initial BBC2 measurements was an apparent difficulty with grounding. This was likely introduced after the measurement electronics were hooked up to a set of recently installed electrical receptacles. I review the nature of the problem

and the techniques I used to fix it.

Chapter 5 concludes with a discussion of two different ways of measuring the state of the qubit. Both employ the dependence of the voltage switching rate Γ on the state of the qubit. The two methods differ in their ways of achieving a measurably high Γ. The first method (swept measurement) uses a linearly ramped bias to increase Γ to the point where escapes are measurable. The second method (pulsed measurement) uses fixed bias conditions and instead applies a short pulse to increase Γ.

My experimental results on BBC2 are given in Chapters 6 and 7. Chapter 6 begins by discussing the BBC2 design. Like the generations immediately preceding it, BBC2 is a dc SQUID with an on-chip LC filter. Where BBC2 differs is in the handling of the microwave drive for the qubit. Previous samples used a short microstrip running from a contact pad to a point near the qubit. The goal was to create a weak capacitive coupling to the qubit. BBC2 instead places the qubit in a gap in an aluminum ground plane. The microwave drive is now provided by a nearby coplanar waveguide. This change was motivated by experiments suggesting that the previous qubit had its lifetime limited by overcoupling to the microwave line [28]. Figure 1.2 is a photograph of a BBC2 patterned qubit in its ground plane, with the coplanar microwave line colored red. The ground plane and waveguide design were adapted from experiments measuring dielectric loss in lumped element microwave oscillators [31, 32].

Following the presentation of the design, I examine the expected lifetime of the qubit. I assume a particular set of channels for dissipation that seem likely to dominate the behavior of the qubit. I then consider how strong the loss to each of these channels is. I arrive at a value for the expected qubit lifetime, given the design parameters, of 1.2 μs. While this would still represent a short time compared to other recent superconducting qubit architectures, it would be among the best

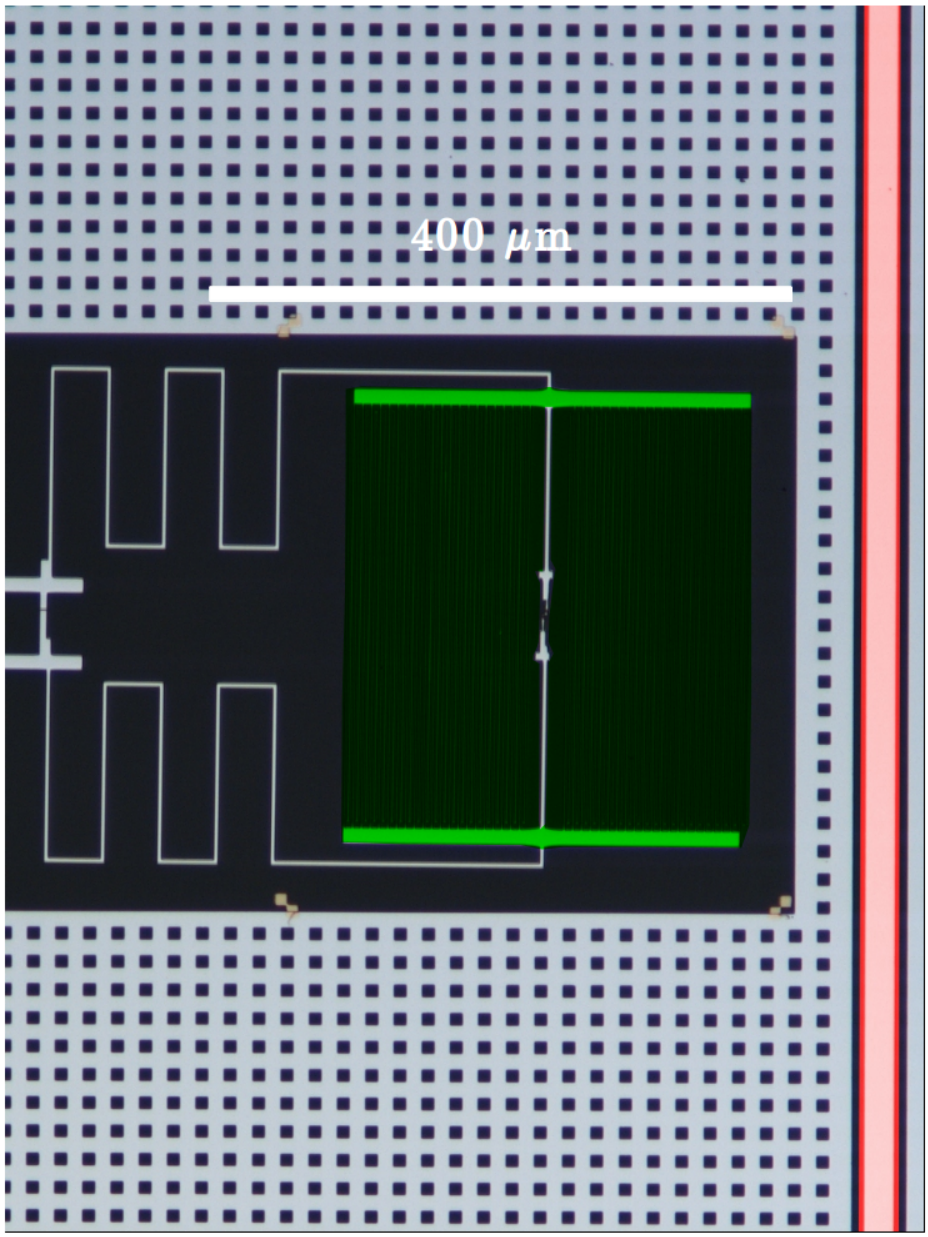

Figure 1.2: Photo of the BBC2 qubit in its ground plane. The fabrication of BBC2 is described in Chapter 4; the design of BBC2 is discussed in Chapter 6.

phase qubit lifetimes. Chapter 6 concludes with a discussion of some of the dc measurements I used to measure device parameters in situ.

In Chapter 7 I present some of the microwave measurements I performed on BBC2. These measurements represent coherent control of the state of the qubit. I begin by looking at spectroscopic studies performed using the swept measurement technique. While this technique allows for much more rapid data acquisition than the pulsed technique, it was of limited use beyond initial characterization. I discuss some of the shortcomings of this technique and how the pulsed measurement resolved them.

Using the pulsed measurement technique, I performed three principal types of measurements: spectroscopy, Rabi oscillations, and lifetime measurements. These measurements exhibited some unexpected properties. Spectroscopically, I anticipated the usual Lorentzian structure for a driven two-level system, with the possibility of other lines corresponding to higher level transitions like $|1\rangle \rightarrow |2\rangle$, etc. Instead, I measured a line with an asymmetric shape, steeper toward the red frequencies. Further, there were unexpected sidebands present in the spectrum. The asymmetry in the lineshape extended to the behavior of the Rabi oscillations, which I found to be much longer lived when detuned to the blue side of the resonance.

We thought the sidebands were suggestive of a characteristic feature of the Jaynes-Cummings model. That model predicts that as the photon occupancy of the harmonic oscillator coupled to the qubit increases, the qubit frequency is shifted a corresponding amount. The sidebands, then, might correspond to these photon peaks. I developed the model of Chapter 3 — where the on-chip LC filter is considered as the oscillator in an effective Jaynes-Cummings model — to understand this spectroscopic behavior. Because the LC filter is very low frequency, many photon states are likely to be relevant, thus yielding a large Hilbert space making full simulation difficult. I consider how a simplified approximation to the predicted photon

shifts reproduces the spectroscopic lineshape. The lineshape does appear generally consistent with the model, although further tests are necessary to determine whether this is the real source of the asymmetry. Many of the results of this chapter are published in [33].

Chapter 8 concludes the thesis by considering what some of those tests may entail. I lay out a theoretical and experimental program for addressing some of the remaining questions concerning the significance of the filter plus SQUID Jaynes-Cummings model. In addition, I also discuss one of the most mysterious aspects of our recent dc SQUID qubits — the anomalous behavior of the tunneling measurement. Since the thesis of Hyeokshin Kwon [34], all of the qubits our lab have studied have shown a strange response to the pulsed measurement when placed in a superposition state between 0 and 1. In BBC2, I also found that two voltage pulses, well-separated in time, were seen to occur following a single measurement pulse. I discuss approaches to a better model of pulsed measurement in dc SQUIDs.

Chapter 2

Quantum Mechanical Treatment of Circuits

Although quantum information can in principle work with systems with more than two states (qudits), most of the quantum circuit language used to describe algorithms focuses on qubits [20]. Consequently, experimental research has focused on two-level systems. In many of the architectures I mention in Section 1.1, the qubits are intrinsically two-level systems such as spins. By contrast, the superconducting circuits used as qubits admit many levels. The restriction to two levels requires the non-linearity of Josephson junctions.

I begin this chapter by discussing the physics of Josephson junctions. In particular, I discuss the "tilted washboard" potential that is the basis for the single junction phase qubit. A single current biased Josephson junction is too strongly coupled to its leads and makes a poor qubit. Improving the performance requires more complicated circuits. I discuss my approach, based on a discussion by Devoret [35], to deriving Hamiltonians for superconducting qubit designs. After discussing the technique, I apply it to relatively simple circuits to give examples of the method.

The Hamiltonian derivation assumes lossless circuits. In practice, loss is one of the most important aspects of a qubit. I discuss how loss is incorporated in qubit models through master equations and derive the Bloch equations describing the response of a damped two-level system to a microwave drive. The fluctuation-dissipation theorem [36] says that with any source of dissipation there are associated fluctuations. One of the major sources of dissipation in superconducting circuits is radiation to the control lines, modeled as a resistance. I discuss how the Johnson

noise associated with this resistance determines the lifetime T_1 of the qubit, and how filtering can reduce this Johnson noise and improve the lifetime.

Finally, with this background in place, I describe the design of the dc SQUID phase qubit studied in our lab. First proposed by Martinis et al. [24], this design uses an inductive current divider to protect a current biased Josephson junction from the bias line. Placing a second junction on the divider arm keeps the qubit junction from being shorted and allows flux bias of the resulting dc SQUID. A quantum model of this design was proposed by Mitra et al. [37] to describe an earlier dc SQUID phase qubit measured in our lab [25].

While Mitra's model succeeded at describing this earlier qubit, it makes several assumptions regarding parameters and bias conditions that are not accurate for the more recent dc SQUID phase qubits. The key problem is that the Josephson junctions we use now are much smaller than those we used then. As I discuss, this makes the inductive coupling between the junctions more important. With increased coupling, modeling the qubit as a single junction is no longer accurate. I present a model I developed to instead describe the qubit as a normal mode of the SQUID. In the large junction limit, this model reduces to the single junction approximation. Measurements on the more recent dc SQUID phase qubits [28, 29] show this model gives better agreement with the data.

Though I did give a public presentation of this model [38], I did not publish it. Independently, Lecocq et al. [39] developed a similar model to describe their "camelback" [40] dc SQUID phase qubits. These are operated in a different parameter regime, with more emphasis on interplay between the two modes [41]. Another similar examination of these ideas is "black-box quantization" [42], where the modes perturbed by the junction nonlinearities are generalized to consider numerical field analysis instead of the lumped element approximation. These other works reinforce the fruitfulness of this normal modes picture.

2.1 Josephson junctions

2.1.1 Josephson relations

A Josephson junction is a weak link between two superconducting regions. This is a broad definition. The various types of weak link correspond to different types of junction. In this thesis I exclusively discuss superconducting-insulating-superconducting (SIS) junctions. The most common junction insulator is an oxide on the surface of one of the superconductors. Early experiments on Josephson junctions often formed SIS junctions through point contacts between a superconducting needle and another surface [43]. In modern experiments, junctions are more typically formed through lithographic processes and the oxide junction barrier explicitly grown. (The fabrication techniques I used in building the BBC2 qubit are discussed in Chapter 4.)

The ideal SIS Josephson element involves only the tunneling of Cooper pairs through the insulating barrier. There are several derivations of the Josephson effect; the one I discuss draws heavily from Devoret's Les Houches lectures on quantum behavior in electrical circuits [35]. The earliest presentation of this derivation appears to be Ferrell and Prange [44].

Consider a zero-temperature limit where only Cooper pairs are tunneling through the junction. Define a number operator

$$\hat{N} = \sum_{N} N \left|N\right\rangle \left\langle N\right|, \tag{2.1}$$

where the eigenstates represent the total of all Cooper pairs that have tunneled from the first electrode to the second electrode. To make the Hamiltonian Hermitian, the tunneling processes can go in either direction, so a tunneling event can increase or

decrease $\hat{N}$. The tunneling Hamiltonian is

$$\hat{\mathcal{H}}_J = -\frac{E_J}{2} \sum_N \left(|N+1\rangle \langle N| + |N\rangle \langle N+1| \right), \tag{2.2}$$

where the tunneling energy scale E_J is called the Josephson energy.

In analogy to tight-binding models where hopping terms like this are common, I take the discrete Fourier transform of these number states. This leads to a phase variable γ,

$$|N\rangle = \frac{1}{2\pi} \int_0^{2\pi} d\gamma \; e^{-\imath N\gamma} |\gamma\rangle, \tag{2.3}$$

$$|\gamma\rangle = \sum_N e^{iN\gamma} |N\rangle. \tag{2.4}$$

In this phase basis, the Hamiltonian (2.2) takes on a simple form. Looking at one of the two outer products,

$$\begin{aligned} \sum_N |N+1\rangle \langle N| &= \frac{1}{2\pi} \int_0^{2\pi} d\gamma \sum_N e^{\imath N\gamma} |N+1\rangle \langle \gamma|, \\ &= \frac{1}{2\pi} \int_0^{2\pi} d\gamma \; e^{-\imath\gamma} \sum_N e^{\imath(N+1)\gamma} |N+1\rangle \langle \gamma|, \\ &= \frac{1}{2\pi} \int_0^{2\pi} d\gamma \; e^{-\imath\gamma} |\gamma\rangle \langle \gamma|. \end{aligned} \tag{2.5}$$

Combining this term with its Hermitian conjugate, I rewrite the Hamiltonian as

$$\hat{\mathcal{H}}_J = -\frac{E_J}{2\pi} \int_0^{2\pi} \cos\gamma \, |\gamma\rangle \langle \gamma|. \tag{2.6}$$

In an even more compact notation, remembering that functions of operators represent series expansions of the operator,

$$\hat{\mathcal{H}}_J = -E_J \cos\hat{\gamma}. \tag{2.7}$$

In this notation, Equation 2.5 shows that

$$e^{-\imath\hat{\gamma}} = \sum_N |N+1\rangle \langle N| \tag{2.8}$$

behaves very similarly to a raising operator. $\hat{N}$ and $\hat{\gamma}$ are conjugate operators, with the commutation relation $\left[\hat{\gamma}, \hat{N}\right] = \imath$.

Since $\hat{N}$ represents the net number of Cooper pairs that have tunneled through the junction, the current through the junction depends on the time derivative of this operator. Using the Heisenberg representation,

$$\begin{aligned}
\frac{d\hat{N}}{dt} &= \frac{1}{\imath\hbar}\left[\hat{N}, \hat{\mathcal{H}}_J\right], \\
&= \frac{1}{\imath\hbar}\left[\hat{N}, -\frac{E_J}{2}\sum_N \left(|N+1\rangle\langle N| + |N\rangle\langle N+1|\right)\right], \\
&= \frac{E_J}{\hbar}\frac{1}{2\imath}\sum_N |N\rangle\langle N+1| - |N+1\rangle\langle N|, \\
&= \frac{E_J}{\hbar}\frac{1}{2\imath}\int_0^{2\pi} d\gamma\ e^{\imath\gamma} - e^{-\imath\gamma}\,|\gamma\rangle\langle\gamma|, \\
&= \frac{E_J}{\hbar}\sin\hat{\gamma},
\end{aligned} \tag{2.9}$$

where I have used (2.5) between lines three and four. Multiplying by the $2e$ charge per Cooper pair yields the current

$$I \equiv 2e\frac{d\hat{N}}{dt} = \frac{2eE_J}{\hbar}\sin\hat{\gamma}. \tag{2.10}$$

This is the dc Josephson effect — we get a current across the junction depending on the phase operator $\hat{\gamma}$, limited by the critical current $I_0 \equiv 2eE_J/\hbar$.

Given these results, a less formal argument yields the voltage across the junction. The energy associated with the tunneling current for a given phase difference γ is given by $-E_J\cos\gamma$ as in (2.7). If the phase difference across the junction changes, then there is a corresponding change in the energy stored and thus work is done on

the junction. Since this model includes no dissipation the dynamics are conservative and thus path independent. Thus, we should be able to recover the change in energy by finding the time integral of the power applied to the system:

$$\begin{aligned} -E_J\left(\cos\gamma_2 - \cos\gamma_1\right) &= \int_{t(\gamma_1)}^{t(\gamma_2)} IV dt, \\ &= \int_{t(\gamma_1)}^{t(\gamma_2)} \frac{2eE_J}{\hbar}\sin\gamma V\ dt. \end{aligned} \tag{2.11}$$

This equation is clearly satisfied if $V\ dt$ is proportional to $d\gamma$, yielding the ac Josephson relation

$$V = \frac{\hbar}{2e}\frac{d\gamma}{dt} \equiv \frac{\Phi_0}{2\pi}\dot{\gamma}, \tag{2.12}$$

where $\Phi_0 \equiv h/2e$ is the magnetic flux quantum.

This derivation is less intuitive than some; among other things, this approach does not make evident that the phase variable γ is associated with the phase difference between the superconducting condensate in the two electrodes. Feynman's coupled Schrödinger equation approach [45] is clearer in this regard. I have chosen Devoret's method because the number operator form of the tunneling Hamiltonian is fairly widely used with superconducting qubits, especially with charge qubits. Thus some familiarity with this approach is useful in reading the literature.

2.1.2 Junction dynamics

The discussion of the previous section considered the ideal Josephson element. A physical junction also includes other processes. The physics of SIS junctions are well captured by the resistively and capacitively shunted junction (RCSJ) model [46, 47]. The resistive shunt represents dissipative processes the junction is exposed to. These include effects like the presence of thermal quasiparticles in the superconducting terminals, or coupling to bias leads. The capacitive shunt represents

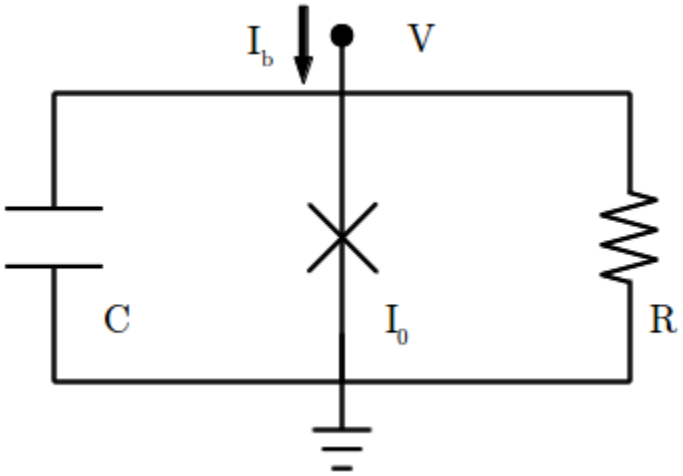

Figure 2.1: Schematic drawing of a Josephson junction in the RCSJ model. The X in the center represents the ideal Josephson element.

displacement currents associated with the changing fields between the terminals as the phase difference changes. In a typical tunnel junction the geometry is essentially that of a parallel plate capacitor, so the source of the capacitance is physically obvious. Figure 2.1 shows a schematic of a physical Josephson junction in this model.

The RCSJ model presents the junction as a classical lumped element circuit. The arguments of the previous section involving energy storage in the ideal junction suggest that in a small-signal limit we could consider it as having some sort of reactance. Given the flowing current, a logical guess is that the Josephson element is inductive. Using equations (2.10) and (2.12), I get the Josephson inductance[1]

$$\begin{aligned} L_J \frac{dI}{dt} &= V, \\ L_J(I_0 \dot{\gamma} \cos\gamma) &= \frac{\Phi_0}{2\pi}\dot{\gamma}, \\ L_J &= \frac{\Phi_0}{2\pi I_0 \cos\gamma}. \end{aligned} \tag{2.13}$$

The nonlinearity in the Josephson relations results in a tunable inductance; as the

[1] A more rigorous derivation of this result is given in Section 2.5.1.

current bias of the junction changes, so does L_J. As $\gamma \to \pi/2$ and the junction is biased near the critical current, the inductance diverges. From the simple argument above, we see that this is because the current is insensitive to small changes in γ near that point.

Applying Kirchhoff's laws to the RCSJ model yields equations of motion for a Josephson junction. Since this is a prelude to a quantum treatment, I ignore the resistance R for now; I discuss how resistance is incorporated into the circuit models in Section 2.4. The bias current I_b is split between the Josephson element and the capacitor:

$$\begin{aligned} I_b &= I_0 \sin\gamma + C\frac{dV}{dt}, \\ &= I_0 \sin\gamma + \frac{C\Phi_0}{2\pi}\ddot{\gamma}. \end{aligned} \tag{2.14}$$

In analogy with mechanical systems, the capacitive term corresponds to a kinetic energy. This suggests describing the rest of the equation with a potential. The potential

$$U = -\frac{\Phi_0}{2\pi}\left(I_b\gamma + I_0\cos\gamma\right) = -E_J\left(\frac{I_b}{I_0}\gamma + \cos\gamma\right) \tag{2.15}$$

yields the correct equation of motion; the flux quantum factor $\Phi_0/2\pi$ gives the right units to the potential.

Figure 2.2 shows plots of the tilted washboard potential. The x-axis is the phase difference $\gamma/2\pi$ across the junction. The y-axis is U/E_J, the potential (2.15) normalized by the Josephson energy. Three different bias currents are shown. At low bias current, the corrugations from the cos term are more prominent, yielding periodically spaced wells. The critical points of the potential are located at γ that satisfy (2.10). The minima correspond to locations with a positive second derivative, and the maxima to the locations with negative second derivative. Defining the bias factor $b \equiv I_b/I_0$, the two solutions to $\sin\gamma = b$ are symmetrically located about $\pi/2$.

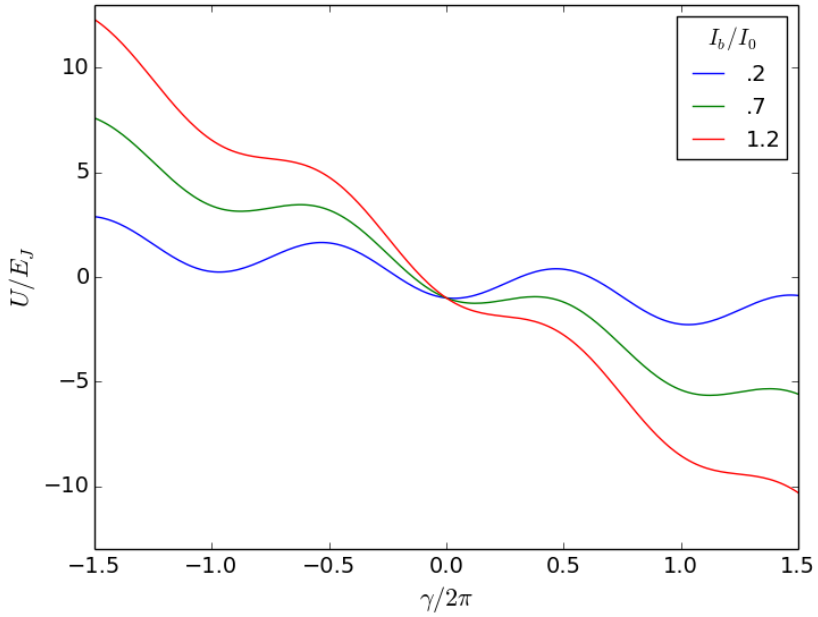

Figure 2.2: Tilted washboard potential for three bias currents I_b.

The height of the potential well is given by

$$\begin{aligned}\frac{\Delta U}{E_J} &= \frac{1}{E_J}(U_{\max} - U_{\min}),\\ &= -b\left(\pi - \sin^{-1} b\right) + I_0 \cos\sin^{-1} b + b\sin^{-1} b + I_0 \cos\sin^{-1} b,\\ &= 2\sqrt{1-b^2} + b(2\sin^{-1} b - \pi).\end{aligned} \tag{2.16}$$

This function is plotted in Figure 2.3.

As the bias current increases, we see two significant effects. The second derivative at the potential minimum is proportional to $I_0 \cos\gamma$. As the bias increases, $\gamma \to \pi/2$ and this curvature vanishes. This corresponds to the increasing Josephson inductance of Equation 2.13. Further, as seen in Figure 2.3, as b approaches 1 the height of the well decreases to 0. When $I_b > I_0$, there is no longer any stable solution to the Josephson relations — this is why I_0 is called the critical current. In our mechanical analogy, there is always a positive force on the phase particle at all points, leading to nonstop acceleration and a continuously changing γ.

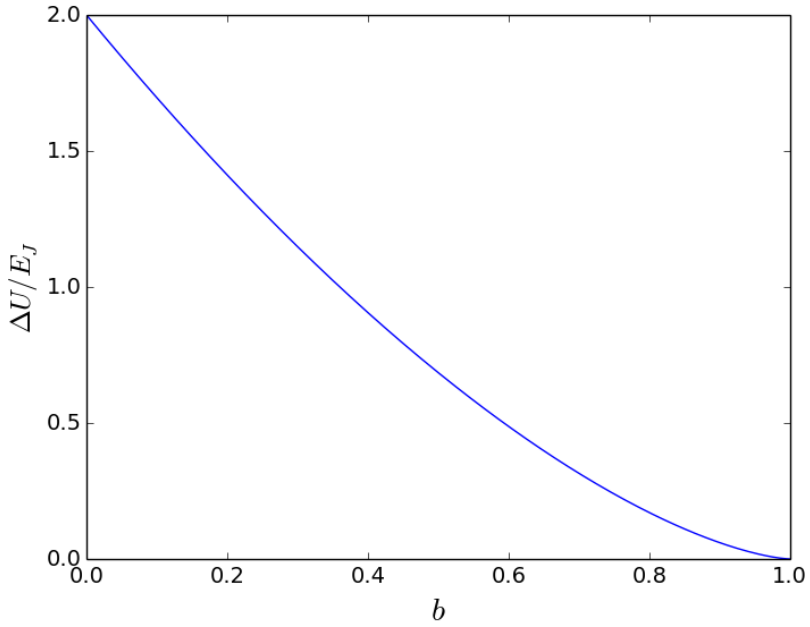

Figure 2.3: Plot of the well height $\Delta U/E_J$ as a function of bias factor b.

From (2.12), continuous change in γ creates a voltage across the junction. The acceleration means that the voltage grows larger and larger with time. In practice, what actually occurs is that when the voltage grows sufficiently large there is enough energy to break some of the Cooper pairs tunneling through the junction. When this happens, the newly formed quasiparticles increase the dissipation of the junction. This loss behaves like friction and stops the acceleration of the phase, leading to a steady dc component to the voltage. I refer to this process of going from a superconducting state to a state with finite voltage as switching, and it is essential for the operation of our qubits.

Depending on the damping already shunting the junction (i.e. the R in the RCSJ model), there may be a unique voltage associated with a particular current or the junction may be hysteretic [48]. Hysteresis occurs in the case of low damping. With low damping, once the phase starts rolling downhill there is plenty of energy to get over successive crests even when we turn the current back down and raise the barrier height (2.16). Since low damping is very desirable for quantum computing,

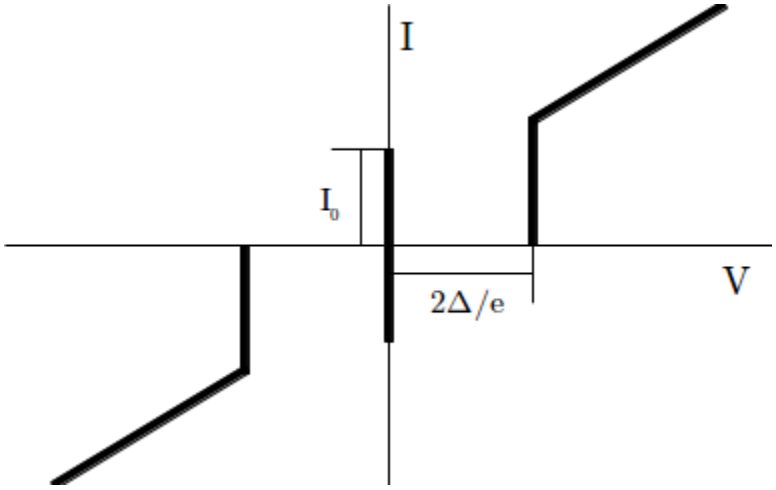

Figure 2.4: Schematic I-V curve for a strongly hysteretic current biased Josephson junction.

superconducting qubits operate in this hysteretic regime.

Figure 2.4 shows a schematic representation of a typical I-V curve for a hysteretic junction. The x-axis is the dc voltage measured across the junction in a four-wire measurement. The y-axis is the current running through the junction. The vertical line at zero voltage is the supercurrent branch of the I-V curve, corresponding to the zero voltage current (2.10). The supercurrent branch can support current in either direction up to I_0. Above I_0, a nonzero voltage develops across the junction. Because of the runaway phenomenon mentioned in the preceding paragraph, for hysteretic junctions voltages below $2\Delta/e$ are not observed. For larger currents, the junction behaves ohmically, with an ideal normal resistance given by the Ambegaokar-Baratoff relation [49]:

$$I_0 R_N = \frac{\pi\Delta}{2e} \tanh\left(\frac{\Delta}{2k_B T}\right). \tag{2.17}$$

At low temperatures, this yields the height of the I-V curve at $V = 2\Delta/e$ as $I = 4I_0/\pi$.

2.1.3 Quantum mechanics of junctions: the phase qubit

Although the Josephson effect was derived quantum mechanically, my discussion of its circuit dynamics was classical. The Josephson relations were used to define a circuit element with specified I-V relations characterized by a definite phase γ and classical I and V. In 1964, Anderson [50] suggested that the Hamiltonian associated with (2.15) could itself be quantized, imposing the usual commutation relation between γ and its conjugate momentum. Although this is a confusing maneuver[2], it has been an excellent tool for understanding quantum effects in circuits based on Josephson junctions.

A major motivation for the study of quantum effects in Josephson circuits was Leggett's proposal to use Josephson junctions to test quantum mechanics in macroscopic systems [52] by looking for macroscopic quantum coherence. Other researchers suggested macroscopic quantum tunneling (MQT) might be easier to observe. Since (2.15) has a finite barrier defining the well, one can use WKB theory to calculate a tunneling rate through the barrier [53]:

$$\Gamma = \omega_0 \sqrt{60 \frac{36\Delta U}{10\pi\hbar\omega_0}} \exp(-\frac{36\Delta U}{5\hbar\omega_0}), \tag{2.18}$$

where ΔU is the height of the barrier (2.16) and $\omega_0 \equiv \sqrt{(1/L_j C)}$ is the plasma frequency of the junction. Classically, thermal activation over the barrier is also possible [54]. However, by operating at sufficiently low temperatures the classical contribution is minimal and quantum tunneling becomes dominant. Voss and Webb [55] first unambiguously observed MQT in Josephson junctions.

Another major quantum signature is the presence of quantized energy levels in Josephson junctions, first observed in 1985 [13]. Because of the tunneling, the washboard potential wells do not admit true energy eigenstates. However, suffi-

[2]For a detailed discussion of many of the issues this raises, see Fred Strauch's thesis [51].

ciently deep wells can sustain long-lived resonances which function very similarly. A simple approximation to the number of states the well can support at a given bias is $\Delta U/\hbar\omega_0$. However, although the assumption of equal spacing is used in this estimate, an important aspect of the washboard is that anharmonicity yields unequal spacing between successive resonances.

This is the underlying idea of the phase qubit, a two-level system using the lowest lying states of the washboard [14]. Qubit control is performed by applying microwaves at the transition frequency between the two lowest states in a washboard well ($|0\rangle$ and $|1\rangle$), using Rabi oscillations [56] to coherently change the qubit state. Provided the control pulses are long enough to have minimal spectral weight at ω_{21} [57], the qubit will remain in the $|0\rangle$ and $|1\rangle$ subspace.

Measuring the phase qubit requires a measurable property of the qubit that has a state dependence. The most common approach uses voltage switching. The escape rate (2.18) is orders of magnitude higher for $|1\rangle$ than $|0\rangle$. When microwaves are applied on resonance, the coherent oscillations between $|0\rangle$ and $|1\rangle$ provide enhancement in the measured escape rate. This enhancement allows reconstruction of the state of the qubit. This measurement technique is described in greater detail in Chapter 7.

2.2 Hamiltonians of lossless circuits: nodal analysis

To study the quantum mechanics of electrical circuits, the first step is to generate Hamiltonians corresponding to a particular circuit design. The simplest thing to begin with is a lossless circuit, where the only elements I consider are inductances, capacitors, and ideal Josephson elements. The approach I use is adapted from Michel Devoret's lectures on quantum fluctuations in electrical circuits [35], which is itself a generalization of an article on quantum network theory by Yurke and Denker [58].

Each element of the circuit — capacitor, inductor, or Josephson junction — is a two-terminal *branch*. Branches terminate at *nodes*. Each branch has associated with it a current running through it and a voltage difference across it. I choose a sign convention so that the voltage difference runs opposite the current. Given a branch current I_b and a branch voltage V_b, the branch charge and branch flux are the time integrals of these quantities:

$$Q_b(t) = \int_{-\infty}^{t} I_b(t')dt', \tag{2.19}$$

$$\Phi_b(t) = \int_{-\infty}^{t} V_b(t')dt', \tag{2.20}$$

To quote Devoret, "branch fluxes and charges do not directly constitute the degrees of freedom of the circuit because they are not independent variables. They must follow constraints imposed by the topology of the circuit" [35] These constraints are Kirchhoff's laws. The current law states all current entering a node must leave the node if the node is remaining at fixed charge; the voltage law states all branch voltages in a loop must sum to zero if the flux through the loop is constant. Using these constraints, there are two complementary approaches to solving the circuit. One can use mesh analysis, where the branch currents are the variables and the voltage law generates the equations. Alternatively, one can use nodal analysis, where the fluxes at each node are the variables and the current law generates the equations. In both approaches, the choice of variable enforces one of Kirchhoff's constraints automatically and uses the other to generate equations of motion. I use nodal analysis for my models.

Nodal analysis uses the flux at each node, but thus far I have only defined branch fluxes. For each circuit, I can choose a particular node to represent ground, a node always at zero voltage. In analogy with branch flux, which is the time integral of branch voltage, this suggests that the nodal flux of ground is zero. This gives

a convenient way of defining the nodal fluxes of the remaining nodes. Given any other node, I define the flux of that node as the sum of the branch fluxes over a path leading from ground to that node. However, a problem with this definition is that it is potentially ambiguous. There might be many paths between ground and a given node.

To assign node fluxes unambiguously, I consider a reduced version of the circuit such that there is exactly one path from each node to ground. An equivalent way of specifying such a network is that it is connected and contains no loops. To see this, suppose there is a node in the network with two paths leading to ground. By going from ground to the node along one path, and returning from the node to ground on the other path, I have produced a loop. In graph theory, a connected network without loops is called a tree, and since the network includes every node, the reduced version of the circuit is called the *spanning tree*.

To construct the spanning tree, begin by choosing a node as ground, then perform a depth-first search [59] of the nodes in the circuit. From whatever node we are at now, we look for a neighboring node (i.e. a branch connects the node we are at and the new node) that is not yet in the spanning tree. If we find one, we incorporate the branch between the present node and the new node into the tree, then go to the new node and repeat the process. When we run out of new nodes connected to our present node, we work back to the previous node and continue looking for new nodes. By only adding branches attached to nodes not yet in the tree, we guarantee that we never introduce a loop. If the network is connected, then every node will eventually be reached and the result is a spanning tree. While a formulaic approach is useful for automatically generating spanning trees of a given network, in practice it is pretty easy to construct a spanning tree by eye.

Given a choice of spanning tree, there is a unique path from ground to each node in the circuit. A final detail is to impose a preferred direction on each branch.

This allows me to define the node flux as the sum of branch fluxes over the connecting path:

$$\phi_n = \sum_b S_{nb}\Phi_b, \tag{2.21}$$

where the sum is over all branches in the tree, S_{nb} is 0 if the branch is not in the path connecting n to ground, 1 if the branch is in the path and is traversed in the correct orientation, and -1 if it is in the path and traversed in the wrong direction. Given the definition of branch fluxes, this amounts to just the time integral of the voltage at node n. The arbitrariness of the spanning tree in this definition seems disconcerting. Devoret comments that it is analogous to a choice of gauge in an electromagnetic problem [35]. In Section 2.3 I show an example of different spanning trees generating the same Hamiltonian.

After choosing the spanning tree, there remain branches in the circuit that are not included in the spanning tree. These branches are called *closure branches*; I denote them with dashed lines in circuit schematics. Every closure branch corresponds to a loop in the circuit. To see this, imagine a closure branch connects nodes A and B. This corresponds to the loop

- ground $\rightarrow$ A along the spanning tree;
- A $\rightarrow$ B through the closure branch;
- B $\rightarrow$ ground along the spanning tree.

Since every closure branch defines a loop, this provides an excellent way of performing bookkeeping on external flux applied to a loop. These external fluxes are simply added to the branch flux of the closure branch defining the loop. With these definitions, the flux version of the Kirchhoff voltage rule

$$\sum_{b\in\text{loop}} \Phi_b = \tilde{\Phi}_l, \tag{2.22}$$

where Φ_l is the externally applied flux through the loop being summed over, is automatically satisfied.

I now use Kirchhoff's current law to give the equations of motion. For each non-ground node, the current in equals the current out. One neat organization of this that makes moving to the Lagrangian easier is to equate the currents moving in on inductive branches (i.e. currents are positive when entering a node on an inductive branch and negative when leaving it) and the currents moving out on capacitive branches (currents are negative when entering a node on a capacitive branch and positive when leaving it). I can write down the currents for inductors, capacitors and Josephson elements in terms of branch fluxes, remembering that the fluxes are just time integrals of voltage:

$$I_{\mathrm{C}} = C\ddot{\Phi}_b; \tag{2.23}$$

$$I_{\mathrm{L}} = \frac{\Phi_b}{L}; \tag{2.24}$$

$$I_{\mathrm{JJ}} = I_0 \sin \frac{2\pi\Phi_b}{\Phi_0}. \tag{2.25}$$

For branches in the spanning tree,

$$\Phi_b = \phi_{\mathrm{start}} - \phi_{\mathrm{end}}, \tag{2.26}$$

while closure branches must also include any applied external flux. The Josephson phase looks unfamiliar expressed this way, but a quick check that it is correct is that it produces the ac Josephson effect when you differentiate the node fluxes.

This discussion has been formal, abstract, and concerned largely with graph theoretic considerations. However, this is the justification for a largely automatic procedure. Establishing an algorithmic procedure for producing the equations of motion ensures that many issues (e.g. sign issues in definitions, where to assign loop

fluxes) are handled consistently. In practice, using this method is actually fairly simple:

- Write down a schematic for a circuit you'd like to quantize.
- Identify the nodes — places where circuit elements are connected together.
- Pick one node to be ground and give the other nodes labels.
- Decide on a spanning tree. The practical importance of this is that the choice of closure branches dictates where external flux shows up in the final model.
- For each node, write down an equation where the left hand side is the sum of ingoing currents on inductive arms (junctions are inductive, outgoing currents are subtracted) and the right hand side is the sum of outgoing currents on capacitive arms. These are the Euler-Lagrange equations for each node.
- Junctions go to cos terms in the Lagrangian. The rest of the Lagrangian, involving inductors and capacitors, will be a quadratic form in the ϕ and ϕ'. Differentiating each Euler-Lagrange equation with respect to these variables yields the second derivatives of the Lagrangian. This extracts the component of $\phi_i\phi_j$ or $\dot{\phi}_i\dot{\phi}_j$ in the Lagrangian (with the proviso that factors of 2 must be included if $i = j$.)

2.3 Hamiltonians of lossless circuits: examples

2.3.1 LC oscillator

The simplest interesting circuit to consider is the classic LC oscillator. There are only two nodes in the circuit, and one of them is ground. I denote the flux of the other node ϕ. With a network of only two nodes and two branches, one branch will be the spanning tree and the other branch will be a closure branch. It is usually

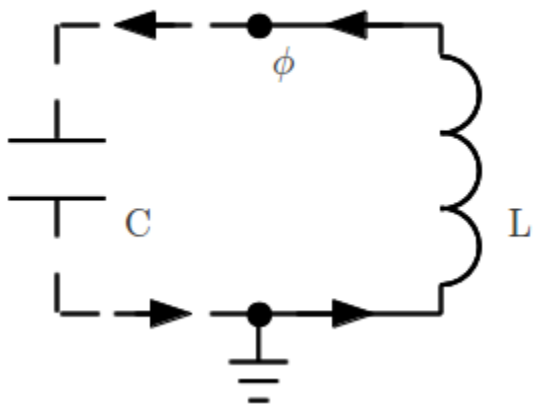

Figure 2.5: Schematic of the LC oscillator. The spanning tree is just one branch, the inductance L.

convenient when possible to choose capacitive branches to be closure branches. This is because capacitive branch currents involve second derivatives of the branch flux; thus, any static external flux disappears from the problem. Figure 2.5 shows the schematic given this choice of spanning tree.

Since there is only one relevant node, the current law yields only one equation of motion:

$$\begin{aligned} I_{\text{Lin}} &= I_{\text{Cout}}; \\ -\frac{\phi}{L} &= C\ddot{\phi}, \end{aligned} \tag{2.27}$$

where the simple form occurs because the branch flux associated with L $\Phi_b = \phi - 0 = \phi$. Integrating the right side with respect to time and then with respect to ϕ' yields $\frac{C\dot{\phi}^2}{2}$; integrating the left side with respect to ϕ gives $-\frac{\phi^2}{L}$. Combining these yields the Lagrangian

$$\mathcal{L} = \frac{C\dot{\phi}^2}{2} - \frac{\phi^2}{L} \tag{2.28}$$

for our LC oscillator.

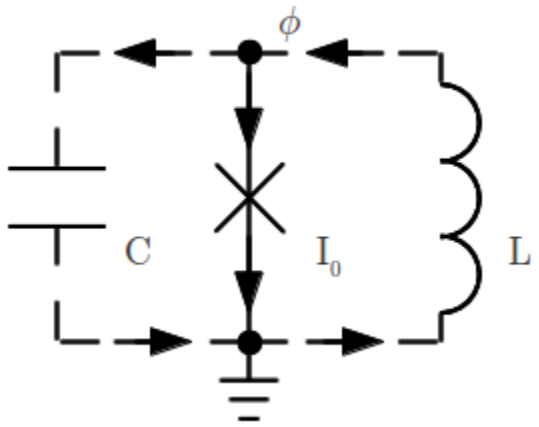

Figure 2.6: Schematic of an rf SQUID. In this setup, the SQUID loop and the capacitor are chosen to be closure branches.

2.3.2 rf SQUID

For a slightly more complicated example, I consider a Josephson junction embedded in a superconducting loop. This type of circuit is known as an rf superconducting quantum interference device, or rf SQUID. These devices exploit the interaction between the Josephson relations and fluxoid quantization, first observed in 1961 [60, 61]. Fluxoid quantization arises from the requirement that superconducting phase be single-valued around the loop. Since the Josephson current depends on the phase difference across the junction, this leads to oscillatory behavior in the current across the junction with changing flux. This quantum interference was first observed in the more complicated dc SQUID [62], a two-junction circuit that I discuss in Section 2.3.3. The single junction SQUID came later [63].

Figure 2.6 shows a schematic of the rf SQUID. Using the RCSJ model to capture the junction (but neglecting the effect of resistance), the inductive loop then corresponds to an inductor shunting the junction. Equivalently, one can consider the circuit as an LC oscillator shunted by a pure Josephson element. The addition of the Josephson junction means that it is now possible to trace a closed superconducting loop around the circuit. This is what makes the dc flux effects discussed

in the preceding paragraphs possible. Since there are three branches and only one active node, two of the branches will be closure branches. Choosing the capacitive branch as the spanning tree seems to cause ambiguity — which of the remaining two branches do we assign loop closure to? This suggests that it is usually simpler to choose inductive branches for the spanning tree.

I choose the Josephson branch to be the one in the spanning tree and let the loop inductance be a closure branch. I have chosen branch directions such that the loop inductor branch points into the node and the Josephson branch points out. Forming the equation of motion for our one active node:

$$\begin{aligned} I_{\text{Lin}} &= I_{\text{Cout}}; \\ -\frac{\phi + \Phi_e}{L} - I_0 \sin\frac{2\pi\phi}{\Phi_0} &= C\ddot{\phi}. \end{aligned} \tag{2.29}$$

This is a good point to consider what would have happened had I chosen the junction as the closure branch instead. The flux contained in the loop is the same either way, but since I have chosen the branches to go in two different directions, the sign of the external flux Φ_e must flip if I choose to include it with the junction instead. This gives the alternative equation of motion

$$-\frac{\phi_*}{L} - I_0 \sin\frac{2\pi(\phi_* - \Phi_e)}{\Phi_0} = C\ddot{\phi}_*, \tag{2.30}$$

which is clearly equivalent to the original if I make the substitution $\phi_* = \phi + \Phi_e$. This offers reassurance that the choice of spanning tree is not physically important.

As before, the right hand side gives a kinetic energy analogue and the left side a potential. Performing the relevant integrations yields the Lagrangian:

$$\mathcal{L} = \frac{C\dot{\phi}^2}{2} + \frac{I_0\Phi_0}{2\pi}\cos\frac{2\pi\phi}{\Phi_0} - \frac{(\phi + \Phi_e)^2}{2L}. \tag{2.31}$$

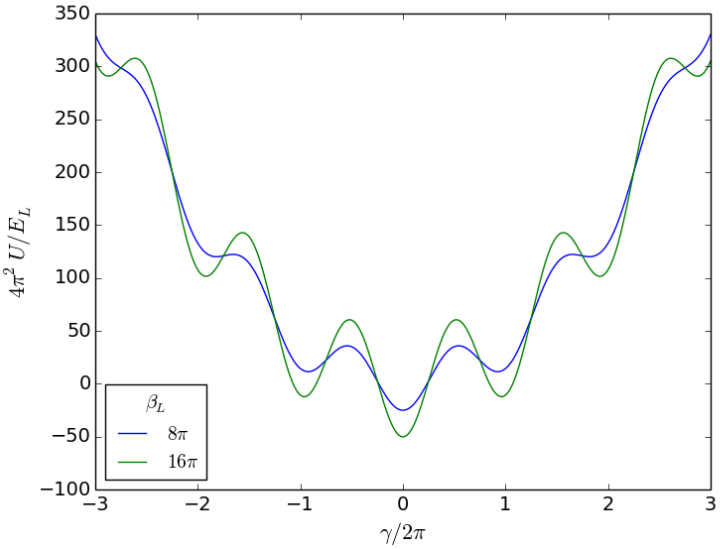

Figure 2.7: Potential $4\pi^2 U/E_L$ associated with an rf SQUID.

This corresponds to a potential

$$U = \frac{(\phi + \Phi_e)^2}{2L} - \frac{I_0 \Phi_0}{2\pi} \cos \frac{2\pi\phi}{\Phi_0}. \tag{2.32}$$

Converting this to dimensionless units by multiplying by the loop inductance and two factors of $2\pi/\Phi_0$, and making the substitution for the Josephson phase $\gamma \equiv 2\pi\phi/\Phi_0$ we get

$$\frac{4\pi^2 U}{E_L} = (\gamma + \gamma_e)^2 - \beta \cos\gamma, \tag{2.33}$$

where $E_L = \Phi_0^2/2L$ and $\beta_L = 2\pi L I_0/\Phi_0$ is an important SQUID parameter. Intuitively, β_L is 2π times the number of flux quanta that can enter the SQUID without exceeding the critical current of the junction.

Figure 2.7 shows a plot of this potential for two different choices of β. The x-axis is the Josephson phase $\gamma/2\pi$; the y-axis is the dimensionless potential $4\pi^2 U/E_L$. β_L controls the depth of the wells; this also determines the number of stable wells. One can get deeper wells by increasing the critical current and making the absolute scale of the cosine oscillations larger, or by increasing the inductance and making

the magnetic contribution to the potential less significant. The wells are very similar to those seen in the current biased Josephson junction potential (2.15), and indeed rf SQUIDs represent one way of realizing phase qubits. One benefit of the rf SQUID is that the parabolic contribution prevents the potential from running to negative infinity on one side, making numerical calculations much more tractable. Physically, this corresponds to the inability of the closed loop to tunnel to the voltage state.

2.3.3 dc SQUID

The last example I examine is the dc SQUID. First and foremost is the question of how to handle the bias current in this picture. I do this by adding an extra circuit element that functions as a current source. Since I want to model a steady state bias current a capacitor is inappropriate — the current sourced by a capacitor depends on the second derivative of fluxes on either side — so I use an inductor. An ideal current source provides current I_b independent of the dynamics of the circuit. To accomplish this, I consider an infinite inductor charged with infinite flux, with the limit taken using a fixed ratio $\Phi/L = I_b$.

The dc SQUID has 4 nodes, 3 of them active. I denote these ϕ_x, ϕ_i, and ϕ_q. I orient all branches so that they go clockwise if you close a loop with the bias inductor, including the capacitors. The arms containing the bias inductor and the arm containing L_q are chosen as closure branches — L_i and the two junctions are the actual branches of the spanning tree. Since there are 3 active nodes, there are

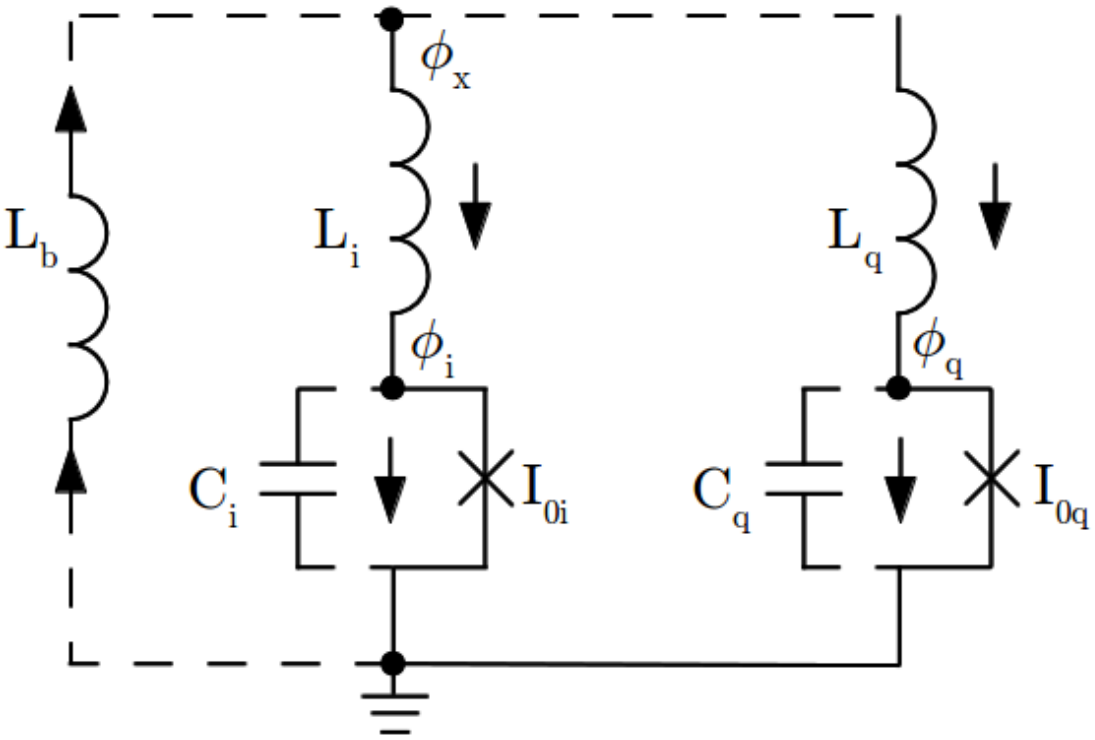

Figure 2.8: Schematic of the dc SQUID.

3 equations of motion for this system:

$$-\frac{\phi_x + \Phi_b}{L_b} + \frac{\phi_i - \phi_x}{L_i} + \frac{\phi_q - \phi_x + \Phi_e}{L_q} = 0; \tag{2.34}$$

$$-\frac{\phi_i - \phi_x}{L_i} - I_{0i} \sin \frac{2\pi\phi_i}{\Phi_0} = C_i \ddot{\phi}_i; \tag{2.35}$$

$$-\frac{\phi_q - \phi_x + \Phi_e}{L_q} - I_{0q} \sin \frac{2\pi\phi_q}{\Phi_0} = C_q \ddot{\phi}_q. \tag{2.36}$$

$$\tag{2.37}$$

I take the $\Phi_b \to \infty, L_b \to \infty, \Phi_b/L_b \to \infty$ limit right away, sending the first term in the first equation to I_b. This first equation is the other new feature of the method demonstrated by the dc SQUID. The presence of a node that is only connected to inductive branches gives a coordinate that doesn't have a kinetic energy term in the Lagrangian. I solve for ϕ_x in this equation and eliminate it elsewhere, thus using

the first equation of motion as a constraint on the other variables. Doing this, I get:

$$\begin{aligned} \phi_x(\frac{1}{L_i}+\frac{1}{L_q}) &= I_b + \frac{\phi_i}{L_i} + \frac{\phi_q+\Phi_e}{L_q}; \\ \phi_x &= \frac{L_q\phi_i + L_i(\phi_q+\Phi_e) + L_iL_qI_b}{L}, \end{aligned} \tag{2.38}$$

where $L \equiv L_i + L_q$ is the inductance of the full SQUID loop. Substituting for ϕ_x gives two equations of motion (which I have simplified algebraically):

$$-\frac{\phi_i-(\phi_q+\Phi_e)}{L} + \frac{L_q}{L}I_b - I_{0i}\sin\frac{2\pi\phi_i}{\Phi_0} = C_i\ddot{\phi}_i; \tag{2.39}$$

$$\frac{\phi_i-(\phi_q+\Phi_e)}{L} + \frac{L_i}{L}I_b - I_{0q}\sin\frac{2\pi\phi_q}{\Phi_0} = C_q\ddot{\phi}_q. \tag{2.40}$$

Before moving to the Lagrangian, it is worth showing that these equations of motion impose sensible constraints for given bias conditions. In steady state, there is no kinetic energy contribution and the right side of both equations vanish. Summing both equations and moving negative terms to the right hand side I get

$$I_b = I_{0i}\sin\frac{2\pi\phi_i}{\Phi_0} + I_{0q}\sin\frac{2\pi\phi_q}{\Phi_0}. \tag{2.41}$$

This is the requirement that the bias current be divided between the two junctions. If I then substitute this back into either equation of motion, I get a second condition that must be satisfied:

$$\phi_i - \phi_q = \Phi_e - L_iI_{0i}\sin\frac{2\pi\phi_i}{\Phi_0} + L_qI_{0q}\sin\frac{2\pi\phi_q}{\Phi_0}. \tag{2.42}$$

The LI terms on the right represent contributions to the flux in the loop from each arm. This constraint represents the flux-phase relationship for the dc SQUID [64], the source of the quantum interference of the SQUID.

Moving from the equations of motion in this case to the Lagrangian is slightly

more complicated than before. The extra layer of complication stems from the inductor giving cross terms in the equations of motion, where the equation of motion for ϕ_i depends on ϕ_q. If the Lagrangian contains a term

$$-\frac{(\phi_i - (\phi_q + \Phi_e))^2}{2L}, \tag{2.43}$$

the derivatives of this show up with the correct sign to give the inductive contributions to each equation of motion. For more involved circuits, however, identifying nice separations of this sort is tricky.

A more mechanical process involves differentiating the inductive parts of the equation of motion by each of the node variables again. Since these inductive parts come from $\partial\mathcal{L}/\partial\phi_i$, this second differentiation gives $\partial^2\mathcal{L}/\partial\phi_i\partial\phi_j$. This is a symmetric matrix, and from it we can easily construct the dependence of $\mathcal{L}$ on all quadratic pairings $\phi_i\phi_j$. Diagonal terms enter with a factor of 1/2, since the squared component gives a 2 on its first differentiation. If we also include the expansion of the junction terms, the matrix we get this way is actually $-U_{\text{harm}}$, the potential matrix of the harmonic approximation around a local equilibrium. I use this in Section 2.5 in deriving approximations to the dc SQUID Hamiltonian.

The Lagrangian of the dc SQUID is:

$$\begin{aligned}\mathcal{L} = &\frac{C_i\dot{\phi}_i^2}{2} + \frac{C_q\dot{\phi}_q^2}{2} + I_b\left(\frac{L_q}{L}\phi_i + \frac{L_i}{L}\phi_q\right); \\ &+ \frac{I_{0i}\Phi_0}{2\pi}\cos\frac{2\pi\phi_i}{\Phi_0} + \frac{I_{0q}\Phi_0}{2\pi}\cos\frac{2\pi\phi_q}{\Phi_0} - \frac{(\phi_i - (\phi_q + \Phi_e))^2}{2L}.\end{aligned} \tag{2.44}$$

Because our lab focused on using dc SQUID phase qubits, I examine this system in more detail in Section 2.5. For quantum mechanical calculations, I want the Hamiltonian, so I convert the Lagrangian to the Hamiltonian in the usual way [65]. Choosing the ϕ variables as our canonical coordinates, the conjugate momenta

associated with them are

$$p_i = \frac{\partial \mathcal{L}}{\partial \dot{\phi}_i} = C_i \dot{\phi}_i; \tag{2.45}$$

$$p_q = \frac{\partial \mathcal{L}}{\partial \dot{\phi}_q} - C_q \dot{\phi}_q. \tag{2.46}$$

So

$$\begin{aligned} \mathcal{H} &= \sum p_n \dot{\phi}_n - \mathcal{L}; \\ &= \frac{p_i^2}{2C_i} + \frac{p_q^2}{2C_q} - I_b \left(\frac{L_q}{L} \phi_i + \frac{L_i}{L} \phi_q \right) \\ &- \frac{I_{0i}\Phi_0}{2\pi} \cos \frac{2\pi\phi_i}{\Phi_0} - \frac{I_{0q}\Phi_0}{2\pi} \cos \frac{2\pi\phi_q}{\Phi_0} + \frac{(\phi_i - (\phi_q + \Phi_e))^2}{2L}. \end{aligned} \tag{2.47}$$

In converting to junction phases γ instead of node fluxes ϕ, the momenta pick up a factor of $(\Phi_0/2\pi)^2$:

$$\begin{aligned} \mathcal{H} = &\frac{p_i^2}{2\left(\frac{\Phi_0}{2\pi}\right)^2 C_i} + \frac{p_q^2}{2\left(\frac{\Phi_0}{2\pi}\right)^2 C_q} - \frac{I_b\Phi_0}{2\pi} \left(\frac{L_q}{L} \gamma_i + \frac{L_i}{L} \gamma_q \right) \\ &- \frac{I_{0i}\Phi_0}{2\pi} \cos\gamma_i - \frac{I_{0q}\Phi_0}{2\pi} \cos\gamma_q + \left(\frac{\Phi_0}{2\pi}\right)^2 \frac{(\gamma_i - \gamma_q - \left(\frac{2\pi\Phi_e}{\Phi_0}\right))^2}{2L}. \end{aligned} \tag{2.48}$$

2.4 Circuits with loss: decoherence and master equations

2.4.1 The Lindblad form

The study of open quantum systems is very interesting, although quite technical [66–69]. The challenge is that quantum mechanics, with its unitary time development, is inherently non-dissipative. One must be careful with ad hoc approaches to including the dissipation. Louisell gives the example of creation and annihilation operators in the Heisenberg picture. These pick up a time dependence $\exp(-iEt/\hbar)$. One could imagine including an imaginary part of the energy, but

then the commutation relation for the creation and annihilation operators becomes

$$[\hat{a}, \hat{a}^\dagger] = e^{-\Gamma t} \tag{2.49}$$

Physically, one must include not only dissipation but fluctuations as well, and these fluctuations serve to enforce the uncertainty principle.

This is typically done by introducing the environment as a large reservoir of modes that can couple to the system, usually linearly (in their influential paper, Caldeira and Leggett argue that the linear coupling is extremely general [53].) This provides a place for the energy of the system to move to; Louisell makes the analogy to a classical system of coupled oscillators. Given long enough, this energy would eventually return to the system of interest. However, as the number of oscillators in the bath increases, the spacing of the modes grows smaller, leading to longer recurrence timescales. In the limit of an infinite bath, the recurrence never occurs.

This is a very elegant solution to the dissipation issue. The price paid is that the problem is now infinitely larger than its original form. However, almost all of the information about the environment is extraneous. We have no way of measuring most aspects of the environment and therefore don't care about it. We can recover the relevant dynamics of the system by averaging over the bath variables, subject to assumptions about the correlations of the bath. In quantum information, this process is described as tracing over the environment to yield the *reduced density matrix*. Unlike the dynamics of the full system, the equation of motion for the reduced density matrix can be non-unitary. This equation of motion is referred to as the master equation.

The most important physical assumption made in deriving the master equation is the Markov approximation. A process is Markovian if its time evolution depends only on the state at the present and not any earlier times. This is often

colloquially summarized as the requirement that the system have no memory. Under this assumption, it can be shown (again, consult [66–69] for more detail) that the most general form for a master equation is the Lindblad form:

$$\frac{d\rho(t)}{dt} = -\frac{\imath}{\hbar}[\mathcal{H}, \rho(t)] + \sum \left(\hat{L}\rho(t)\hat{L}^\dagger - \frac{1}{2}(\rho(t)\hat{L}^\dagger\hat{L} + \hat{L}^\dagger\hat{L}\rho(t)) \right), \tag{2.50}$$

where the sum is over all Lindblad operators L corresponding to dissipation.

2.4.2 Two-level master equation: the Bloch equation

As an application of this, I derive the Bloch equation in Appendix A. Originally developed in the context of NMR [70], it is now widely used to describe resonance in all sorts of two-level systems. In particular, it is an equation of motion for the density matrix of a driven qubit. In the derivation, I switch from ρ to a different representation of the density matrix χ^* which yields a simpler equation:

$$\frac{d\chi^*}{dt} = \begin{pmatrix} -\imath\frac{\Omega}{2}(\chi^*_{10} - \chi^*_{01}) + \Gamma_1\chi^*_{11} & -\imath\frac{\Omega}{2}(\chi^*_{11} - \chi^*_{00}) - (\Gamma_2 + \imath\Delta)\chi^*_{01} \\ -\imath\frac{\Omega}{2}(\chi^*_{00} - \chi^*_{11}) - (\Gamma_2 - \imath\Delta)\chi^*_{10} & \imath\frac{\Omega}{2}(\chi^*_{10} - \chi^*_{01}) - \Gamma_1\chi^*_{11} \end{pmatrix}. \tag{2.51}$$

This form is useful because the time dependence has been eliminated on the right hand side. This makes it possible to consider the spectroscopic limit, where $t \to \infty$ and $d\chi^*/dt = 0$. The spectroscopic lineshape arises from solving for the population in this steady state as a function of the detuning. Since the rotating frame transformations leave the populations unchanged,

$$\rho_{11} = \chi^*_{11} = \frac{1}{2}\frac{\Omega^2\Gamma_2}{\Gamma_1(\Gamma_2^2 + \Delta^2) + \Omega^2\Gamma_2} \tag{2.52}$$

Figure 2.9 shows the lineshape for several different drive powers. There are a couple of important features. First, as the power increases, the maximum value of

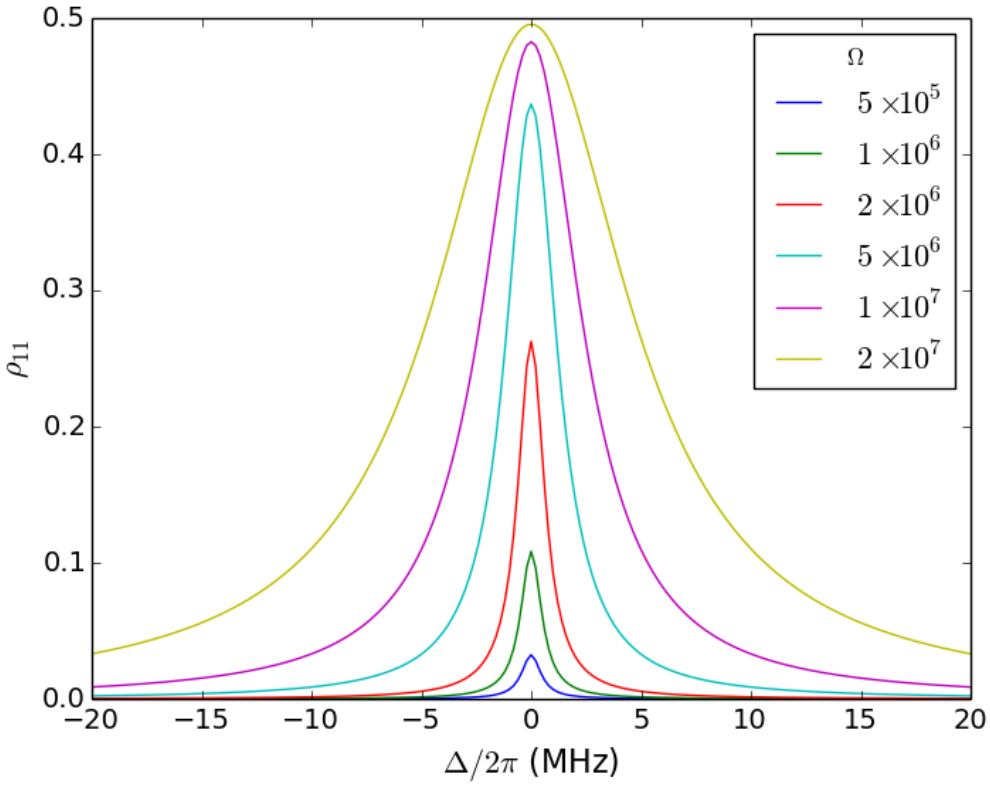

Figure 2.9: Spectral lineshape of Equation 2.52 for several different choices of Ω. The coherence times are $T_1 = 220$ ns, $T_\phi = 1100$ ns.

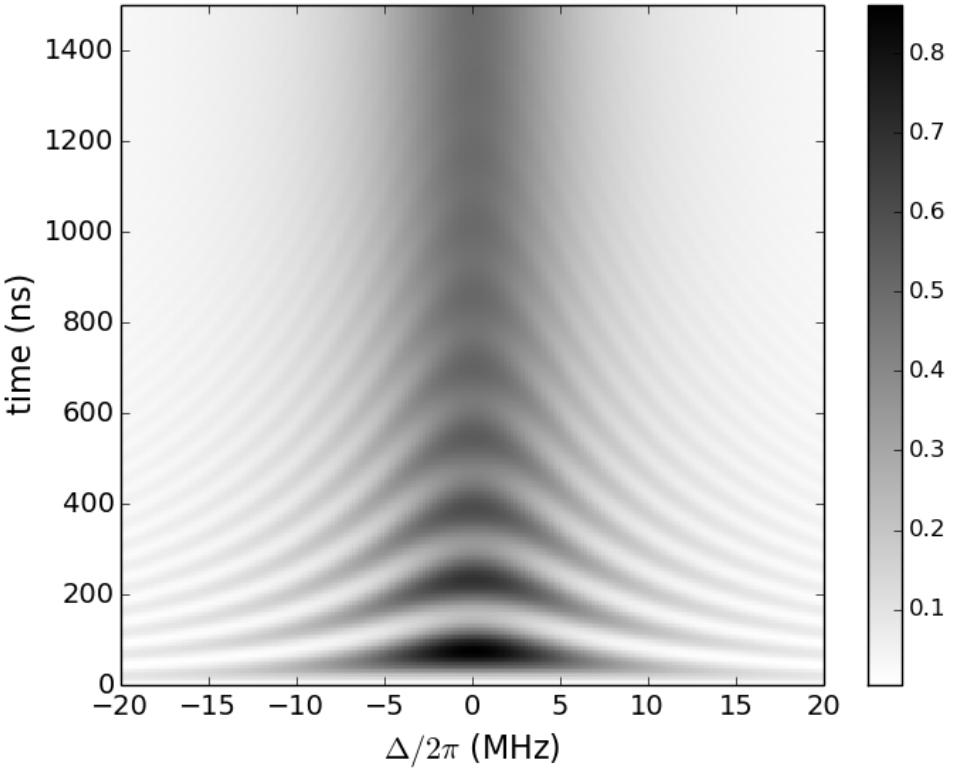

Figure 2.10: Rabi oscillations for a transition with the same parameters as Figure 2.9. At each detuning, $\Omega = 4 \times 10^7$.

ρ_{11} on resonance approaches a saturation point of 1/2. This is because the transition rate induced by the microwaves is equal for both excitation and relaxation. Even in the regime where this transition rate dominates the T_1 processes, it still can't lead to a population inversion. Also, as the power increases, the width of the peak broadens. From (2.52), the full width at half maximum of the peak is

$$\text{FWHM} = 2\sqrt{\Gamma_2^2 + \Omega^2 \frac{\Gamma_2}{\Gamma_1}}. \tag{2.53}$$

Using the same parameters as Figure 2.9, I examine the response of the system immediately after the drive is turned on, assuming that the system begins in its ground state. Figure 2.10 is a color plot of the population ρ_{11} versus detuning and time. The Rabi oscillations are slowest on resonance, speeding up symmetrically as

the detuning increases. The faster Rabi frequency off resonance leads to a greater number of visible oscillations before they decay.

A common visualization for two-level systems is the Bloch sphere. An arbitrary pure state of a two-level system can be expressed as

$$\cos(\frac{\theta}{2})\,|0\rangle + \sin(\frac{\theta}{2})e^{\imath\phi}\,|1\rangle\,, \tag{2.54}$$

where $0 \leq \theta \leq \pi, 0 \leq \phi \leq 2\pi$. These two angles describe the surface of a sphere. $|0\rangle$ and $|1\rangle$ correspond to the north and south poles of the Bloch sphere, and the equator represents all possible equal superpositions of the two. More generally, the mixed states of the reduced density matrix correspond to points within the Bloch sphere.

2.4.3 T_1 in circuits: the role of filtering

The Bloch equations show how the qubit evolves given the coherence times T_1 and T_2, but does not determine where those parameters come from. Given a particular source of noise, perturbation theory and Fermi's golden rule furnish an estimate of the relaxation time T_1. I work out the T_1 calculation in the context of a single junction to make the structure of the argument clear. This argument draws substantially from presentations in [71, 72].

The potential for a single junction is

$$U_{\text{sing}} = -\frac{\Phi_0}{2\pi}\left(I_0 \cos\gamma + I_b\gamma\right). \tag{2.55}$$

I consider the noise as a perturbation

$$V_{\text{sing}}(\tau) = -\frac{\Phi_0}{2\pi} I_n(\tau)\gamma(\tau), \tag{2.56}$$

where $I_n(\tau)$ is a time-dependent random current noise. I use time-dependent perturbation theory to see how this noise affects the wavefunction. I work in the interaction representation, where wavefunctions evolve in time according to the perturbation and operators evolve in time according to the unperturbed Hamiltonian.

First order perturbation theory yields

$$|\Psi(t)\rangle = |\Psi(0)\rangle - \frac{i}{\hbar}\int_0^t d\tau\ \hat{V}_{\text{sing}}(\tau)|\Psi(0)\rangle. \tag{2.57}$$

For the relaxation rate, I assume $|\Psi(0)\rangle = |1\rangle$ and find the rate of transition back to $|0\rangle$. Taking the inner product of this wavefunction with the ground state yields

$$\begin{aligned}\alpha_0 &= -\frac{i}{\hbar}\int_0^t d\tau\ \langle 0|\hat{V}_{\text{sing}}(\tau)|1\rangle \\ &= -\frac{i}{2e}\int_0^t d\tau\ \langle 0|\gamma(\tau)|1\rangle I_n(\tau),\end{aligned}$$

where I have used the definition of the flux quantum to simplify. In the interaction representation

$$\gamma(\tau) = e^{-iH_0t}\hat{\gamma}e^{iH_0t}; \tag{2.58}$$

this lets me rewrite (2.58) as

$$\alpha_0 = -\frac{i}{2e}\langle 0|\hat{\gamma}|1\rangle\int_0^t d\tau\ e^{-i\omega_{01}\tau}I_n(\tau). \tag{2.59}$$

Taking the absolute square of this yields the ground state probability:

$$P_0(t) = \frac{|\langle 0|\hat{\gamma}|1\rangle|^2}{4e^2}\int_0^t\int_0^t d\tau_1\ d\tau_2\ e^{i\omega_{01}(\tau_1-\tau_2)}\langle I_n(\tau_1)I_n(\tau_2)\rangle, \tag{2.60}$$

where I have also taken an ensemble average of the noise currents.

The averaged product of the I_n's is the autocorrelation function of the current noise, and the overall form of our expression looks very similar to a Fourier trans-

form of the autocorrelation function. The Wiener-Khinchin theorem [72] says that this Fourier transform is just the noise spectral density. So, I use a trick to get the resemblance to be more exact. The central argument is that we expect to be interested in time scales much longer than the correlation time of the noise. This is an excellent approximation for thermal noise, which is white until very high frequencies and is thus approximately delta correlated. I also assume that the noise source is stationary, meaning that the correlation function depends on t_1 and t_2 only through the time difference $t_2 - t_1$.

I make a change of variables. Let $\tau = \tau_1 - \tau_2$, so $d\tau = -d\tau_2$. Then

$$P_0 = -\frac{|\langle 0|\gamma|1\rangle|^2}{4e^2} \int_0^t \int_{\tau_1}^{\tau_1 - t} d\tau_1 \, d\tau \, e^{i\omega_{01}\tau} \langle I_n(\tau_1) I_n(\tau_1 - \tau)\rangle. \tag{2.61}$$

I use the negative sign to flip the order of integration in the $d\tau$ integral. We can also exploit the stationary noise to replace the autocorrelation with one shifted in time by $\tau - \tau_1$, yielding

$$P_0 = \frac{|\langle 0|\gamma|1\rangle|^2}{4e^2} \int_0^t \int_{\tau_1 - t}^{\tau_1} d\tau_1 \, d\tau \, e^{i\omega_{01}\tau} \langle I_n(\tau) I_n(0)\rangle. \tag{2.62}$$

Now, since $\tau_1 < t$, the τ integral crosses 0. Since we expect noise with a very short correlation time, the $d\tau$ integral only makes significant contributions near the origin. Thus, it is an excellent approximation to extend the limits of integration to infinity:

$$P_0 = \frac{|\langle 0|\gamma|1\rangle|^2}{4e^2} \int_0^t \int_{-\infty}^{\infty} d\tau_1 \, d\tau \, e^{i\omega_{01}\tau} \langle I_n(\tau) I_n(0)\rangle. \tag{2.63}$$

In this form, I can now apply the Wiener-Khinchin theorem

$$\mathcal{S}_{II}(\omega_{01}) = \int_{-\infty}^{\infty} d\tau e^{i\omega_{01}\tau} \langle I_n(\tau) I_n(0)\rangle, \tag{2.64}$$

where $\mathcal{S}_{II}(\omega_{01})$ is the power spectral density of the current noise at frequency ω_{01}. With this simplification the probability (2.60) becomes

$$\begin{aligned} P_0 &= \frac{|\langle 0|\gamma|1\rangle|^2}{4e^2}\int_0^t d\tau_1\, \mathcal{S}_{II}(\omega_{01}) \\ &= \frac{t|\langle 0|\gamma|1\rangle|^2}{4e^2}\mathcal{S}_{II}(\omega_{01}). \end{aligned}$$

The time derivative of the probability yields the transition rate

$$\Gamma_{\text{stim}} = \frac{|\langle 0|\gamma|1\rangle|^2}{4e^2}\mathcal{S}_{II}(\omega_{01}). \tag{2.65}$$

The transition rate (2.65) characterizes stimulated emission. T_1 corresponds to spontaneous emission. However, thermodynamics lets me calculate spontaneous emission given the stimulated emission rate. The stimulated emission rate is equal to the rate at which the noise stimulates transitions from 0 to 1. However, to achieve thermal equilibrium, the rates must obey detailed balance; because of the population disparity, the rate down must be larger than the rate up by a Boltzmann factor to compensate. This determines the spontaneous emission rate $\Gamma_1 = 1/T_1$.

Since the energy difference between $|0\rangle$ and $|1\rangle$ is $\hbar\omega_{01}$,

$$\begin{aligned} \Gamma_{\uparrow} e^{\frac{\hbar\omega_{01}}{kT_q}} &= \Gamma_{\downarrow} \\ \Gamma_{\text{stim}} e^{\frac{\hbar\omega_{01}}{kT_q}} &= \Gamma_{\text{stim}} + \Gamma_{\text{spon}} \\ \Gamma_{\text{stim}}\left(e^{\frac{\hbar\omega_{01}}{kT_q}} - 1\right) &= \Gamma_{\text{spon}}, \end{aligned}$$

where T_q denotes the temperature at the qubit. The current spectral density for a thermal noise source is

$$\mathcal{S}_{II}(\omega_{01}) = \frac{2\hbar\omega_{01}}{R}\frac{1}{e^{\frac{\hbar\omega_{01}}{kT_n}} - 1}, \tag{2.66}$$

where now T_n is the temperature of the noise source. This is usually written as being larger by a factor of 2; this comes from using the symmetrized expression

where we also consider the contributions from the negative frequency. We're left with a general expression for the spontaneous emission rate:

$$\Gamma_{\text{spon}} = \frac{2\hbar\omega_{01}}{R}\frac{|\langle 0|\gamma|1\rangle|^2}{4e^2}\frac{e^{\hbar\omega_{01}/kT_q}-1}{e^{\hbar\omega_{01}/kT_n}-1}. \tag{2.67}$$

In the case where the noise source and the qubit are at the same temperature, the last term reduces to 1; this is the case I consider for now. Quoting Xu [71], the matrix element for the single junction case is $2e^2/(\hbar\omega_{01}C)$, yielding a $T_1 = RC$ for the $|1\rangle$ to $|0\rangle$ transition.

This argument shows that reducing $\mathcal{S}_{II}(\omega_{01})$ can improve T_1. The simplest way to reduce this noise is through filtering. A filter diverts some fraction of the incoming current, preventing it from reaching the qubit. The reduction in the noise spectral density gives a smaller Γ_{stim}, and since $\Gamma_1 \propto \Gamma_{\text{stim}}$, this improves T_1. Functionally, reduced current noise density looks as if it is coming from a larger resistor. In considering the loss, the correct quantity to use is not Re(Z), but 1 / Re (Y); this comes from the parallel topology.

This gives us a method for handling more complicated circuits. By considering how additional circuit elements transform the admittance presented to the qubit, I can examine the anticipated effect on T_1.

2.5 The dc SQUID phase qubit

I now consider the design and analysis of the dc SQUID phase qubit. This approach to phase qubits was originally proposed by Martinis et al. [24] in 2002. Variations on the dc SQUID qubit have been the central focus of our lab for several years, including the BBC2 qubit that I study in this thesis.

Figure 2.5 is a schematic of the dc SQUID. The labeling of the SQUID arms as i and q, previously unexplained, arises from their intended role in this circuit. I

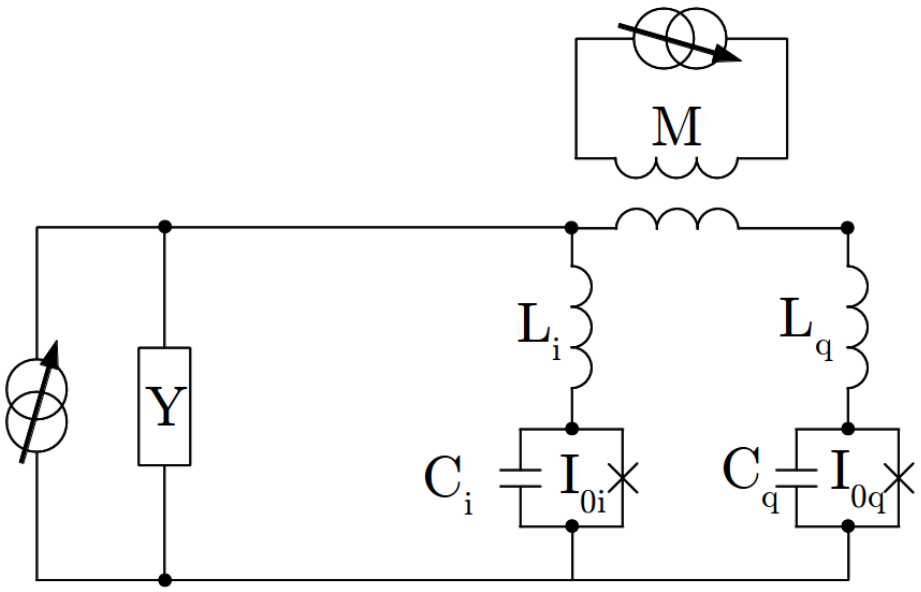

Figure 2.11: Schematic of the dc SQUID phase qubit, including bias line admittance Y. The isolation junction filters noise from Y from reaching the qubit junction. Voltage tunneling events are measured at V.

call I_{0q} the qubit junction, as it represents an intended single junction phase qubit. I_{0i} is the isolation junction. As seen in Section 2.4.3, reducing current noise to the qubit increases T_1. The dc SQUID design is motivated by this idea of filtering. The appeal of this approach is that it is broadband, capable of filtering noise both at low frequencies (improving contributions to T_2 from inhomogeneous broadening) and at the transition frequency ω_{01} (improving T_1.)

The filtering mechanism provided by the dc SQUID is an inductive current divider. By attaching the bias leads to the SQUID asymmetrically, I make the contributions of the loop inductance on the two arms very different. In particular, connecting the bias leads very close to the isolation junction yields $L_i \ll L_q$. Since the impedance of an inductive element is linear in L and current divides between the two arms in inverse proportion to their impedance, this means that current from the bias leads will predominantly travel through the isolation arm (which is where the name comes from.)

To see how much this improves T_1, I consider a simpler version of the divider that excludes the isolation junction[3]. I also assume $Y = 1/R$ real. Then the transformed impedance and admittance are

$$Z_t = \imath\omega L_q + \frac{\imath\omega L_i R}{R + \imath\omega L_i}; \tag{2.68}$$

$$Y_t = \frac{R + \imath\omega L_i}{\imath\omega L_q(R + \imath\omega L_i) + \imath\omega L_i R}; \tag{2.69}$$

$$\mathrm{Re}Y_t = \frac{R\omega^2 L_i^2}{(\omega^2 L_i L_q)^2 + \omega^2 R^2 (L_i + L_q)^2}; \tag{2.70}$$

$$\frac{1}{\mathrm{Re}Y_t} = \frac{\omega^2 L_q^2}{R} + \left(\frac{L_q + L_i}{L_i}\right)^2 R. \tag{2.71}$$

Given a typical 50 Ω line impedance, for frequencies at the qubit transition frequency

[3]The chief effect of the junction is to increase the impedance in the isolation arm near the resonance of the isolation junction; thus, the performance of the divider is markedly worsened near the isolation resonance.

or lower the second term of (2.71) dominates. Thus the loss to the leads decreases roughly quadratically with the ratio of the large inductance to the small inductance. If the leads are one of the major sources of dissipation, then this improves T_1 dramatically.

The advantage to putting a junction in the isolation arm is that a pure inductor would short circuit the qubit junction. We would have an rf SQUID, and as mentioned in Section 2.3.2, voltage state tunneling does not occur in an rf SQUID. This makes the voltage switch measurement technique impossible. With a junction in the isolation arm, the SQUID can support a non-zero voltage across the qubit junction. The closed loop also gives us the option of applying flux bias to the SQUID. This is a very helpful tool for biasing the qubit, since by design most of the bias current ends up going through the isolation junction instead.

2.5.1 Modeling the dc SQUID phase qubit

The preceding discussion assumes that the qubit junction does indeed behave like a single junction phase qubit. The current biased Josephson junction, while presenting some subtleties, is a comparatively simple system. Much of its dynamics are captured by a cubic approximation to the potential around the minimum of one of the wells. The natural question is how well this sort of analysis translates to the dc SQUID phase qubit.

As derived in Section 2.3.3, the Hamiltonian for the dc SQUID is given by

$$\begin{aligned}\mathcal{H} = &\frac{p_i^2}{2\left(\frac{\Phi_0}{2\pi}\right)^2 C_i} + \frac{p_q^2}{2\left(\frac{\Phi_0}{2\pi}\right)^2 C_q} - \frac{I_b\Phi_0}{2\pi}\left(\frac{L_q}{L}\gamma_i + \frac{L_i}{L}\gamma_q\right) \\ &- I_{0i}\cos\gamma_i - I_{0q}\cos\gamma_q + \left(\frac{\Phi_0}{2\pi}\right)^3 \frac{(\gamma_i - \gamma_q - \left(\frac{2\pi\Phi_e}{\Phi_0}\right))^2}{2L}.\end{aligned} \tag{2.72}$$

To simplify this Hamiltonian, I begin by making the cubic approximation to the cos

terms, discarding all fourth order and higher terms in the Taylor expansion of cos:

$$\begin{aligned}\cos\gamma &= \cos\left(\gamma_* + \delta\gamma\right), \\ &= \cos\gamma_* - (\sin\gamma_*)\delta\gamma - (\cos\gamma_*)\frac{\delta\gamma^2}{2!} + (\sin\gamma_*)\frac{\delta\gamma^3}{3!}\end{aligned} \tag{2.73}$$

I am interested in the potential near a well, and thus for small deviations from a potential minimum. By definition, this occurs where all first order terms in the potential vanish. Given a particular choice of $(\gamma_{i*}, \gamma_{q*})$, I assume that I have chosen I_b and Φ_e such that $(\gamma_{i*}, \gamma_{q*})$ is in fact the minimum. This allows me to discard all linear terms in the potential, not just those arising from the cos expansion. This leads to the cubic expansion of the dc SQUID Hamiltonian

$$\begin{aligned}\mathcal{H}_{\text{cub}} = &\frac{p_q^2}{2m_q} + \frac{p_i^2}{2m_i} + \left(\frac{\Phi_0}{2\pi}\right)^2\left(\frac{1}{2L_{qj}} + \frac{1}{2L}\right)\delta\gamma_q^2 \\ &+ \left(\frac{\Phi_0}{2\pi}\right)^2\left(\frac{1}{2L_{ij}} + \frac{1}{2L}\right)\delta\gamma_i^2 - \left(\frac{\Phi_0}{2\pi}\right)^2\frac{1}{L}\delta\gamma_q\delta\gamma_i \\ &- \frac{1}{6}\frac{\Phi_0}{2\pi}\left(I_{0q}\sin\gamma_{q*}\ \delta\gamma_q^3 - I_{0i}\sin\gamma_{i*}\ \delta\gamma_i^3\right),\end{aligned}$$

where $L_{ij} = \Phi_0/2\pi\cos\gamma_{i*}$ is the Josephson inductance (2.13) of the isolation junction, arising from the linearization of the washboard potential

In 2008, Mitra et al. [37] considered this Hamiltonian as a starting point for a description of a dc SQUID phase qubit. A particular set of qubit parameters were chosen to match a then-recent device [73]. For this device, $I_{0q} = 17.75\mu\text{A} > I_{0i} = 6.40\mu\text{A}$. The unbiased Josephson inductance is $330\text{pH}\mu\text{A}/I_0$; for these critical currents the Josephson inductances are tens of picohenries. This is much smaller than the loop inductance $L \equiv L_i + L_q \approx 3.4nH$.

Given these parameters, certain approximations are made to simplify the Hamiltonian:

- The current and the flux are assumed to be simultaneously linearly ramped

such that no net current is sent through the isolation junction. This minimizes the Josephson inductance of the isolation junction, allowing the highest possible isolation.

- Because the isolation junction is unbiased, $\sin\gamma_i^*$ vanishes and so there is no $\delta\gamma_i^3$ contribution to the potential. All of the relevant nonlinearity is assumed to come from the qubit junction alone.

- This leaves two terms in the Hamiltonian that differ from two uncoupled harmonic oscillators — the $\delta\gamma_q^3$ term and the $\delta\gamma_i\delta\gamma_q$ term. These contributions are both treated as perturbations to the Hamiltonian of two uncoupled harmonic oscillators.

The resulting theory agreed well with observations on qubits like the one described. However, the parameter regime of more recent dc SQUID phase qubits is much different. In particular, while loop inductances are still large, the junction areas and critical currents are much smaller. Early phase qubit designs used large Josephson junctions for two reasons. First, the characteristic phase qubit behavior requires $E_J \gg E_C$. Large junctions have bigger capacitances (and thus smaller E_C) and bigger critical currents (thus larger E_J.) Second, it was thought that the dominant decoherence mechanism in phase qubits would be fluctuations in the junction critical current [74]. The relative importance of such fluctuations is less the larger the critical current of the junction.

Subsequent work on phase qubits [75, 76] showed that dielectric loss was a more important loss mechanism. Unexpected avoided crossings in the qubit spectroscopy suggested a specific loss mechanism of coupling between the qubit and other two-level systems, likely in the junction barrier. Such two-level systems are well-known features in amorphous systems like the aluminum oxide junction dielectric layer [77]. This work on phase qubits suggested that existing materials could be improved

by using smaller junctions, decreasing the number of relevant two-level systems by reducing the volume of dielectric. To recoup the necessary E_J/E_C ratio, the junction is shunted by a capacitor using a less lossy dielectric [78].

Following the move to smaller junctions, the assumptions of the model of Mitra et al. are no longer valid. Because smaller critical currents correspond to larger Josephson inductances, recent designs typically have $I_{0i} \gg I_{0q}$. In this regime, it is no longer possible to perform voltage tunneling measurement while the isolation junction remains unbiased; I discuss this in Section 6.3.2. Thus cubic contributions in $\delta\gamma_i$ reappear in the Hamiltonian.

More importantly, the perturbative approach previously adopted breaks down. When both junctions are biased, there is usually a place where the resonances of the two independent junctions are degenerate. Failing to account for this in the perturbation theory leads to very poor numerical results, even relatively far away from the degeneracy. Beyond that, with low critical currents the biased junctions can have Josephson inductances comparable to the loop inductance. In this regime, the coupling term $\delta\gamma_q\delta\gamma_i$ becomes comparable in size to the individual junction contributions. To call this a perturbation is inaccurate, and suggests many orders of perturbation theory would be necessary for useful results.

The underlying idea of perturbative corrections to a harmonic basis is a good one. The chief problem is that the choice of each junction as an independent harmonic oscillator is the wrong harmonic basis for discussing the problem. A better choice of basis includes the coupling exactly at the quadratic level and uses the normal modes of the SQUID as the harmonic basis.

2.5.2 Coupled junctions: normal mode analysis

The normal modes of the SQUID represent solutions where all of the coordinates have a common sinusoidal time dependence. A general result from classical

mechanics [65] states that any homogeneous second-degree potential (i.e. containing only terms like $q_i q_j$) can be transformed to a coordinate system where all of the coupling terms vanish.

To get the equations of motion from a Hamiltonian, I use Hamilton's equations[4]:

$$\dot{q}_n = \frac{\partial \mathcal{H}}{\partial p_n}; \tag{2.74}$$

$$-\dot{p}_n = \frac{\partial \mathcal{H}}{\partial q_n}. \tag{2.75}$$

In the cubic Hamiltonian (2.74), the canonical coordinate q_n is $\delta\gamma_n$. Plugging this into Hamilton's equations for the two coordinates,

$$\dot{\delta\gamma}_q = \frac{p_q}{m_q}; \tag{2.76}$$

$$\dot{\delta\gamma}_i = \frac{p_i}{m_i}; \tag{2.77}$$

$$-\dot{p}_q = \left(\frac{\Phi_0}{2\pi}\right)^2 \left(\frac{1}{L_{qj}} + \frac{1}{L}\right) \delta\gamma_q - \left(\frac{\Phi_0}{2\pi}\right)^2 \frac{1}{L}\delta\gamma_i; \tag{2.78}$$

$$-\dot{p}_i = \left(\frac{\Phi_0}{2\pi}\right)^2 \left(\frac{1}{L_{ij}} + \frac{1}{L}\right) \delta\gamma_i - \left(\frac{\Phi_0}{2\pi}\right)^2 \frac{1}{L}\delta\gamma_q, \tag{2.79}$$

where $m_n = (\Phi_0/2\pi)^2 C_n$. I use the first two equations to return to two coupled second-order equations rather than four first-order:

$$-C_q \ddot{\delta\gamma}_q = \left(\frac{1}{L_q^*}\right) \delta\gamma_q - \frac{1}{L}\delta\gamma_i; \tag{2.80}$$

$$-C_i \ddot{\delta\gamma}_i = \left(\frac{1}{L_i^*}\right) \delta\gamma_i - \frac{1}{L}\delta\gamma_q, \tag{2.81}$$

[4]Where variables are subscripted as n, this denotes the choice of any appropriate coordinate from the problem — for this Hamiltonian, the isolation or qubit coordinates.

where I have defined the parallel inductances

$$\frac{1}{L_n^*} \equiv \frac{1}{L} + \frac{1}{L_{nj}}. \tag{2.82}$$

To get the normal modes, assume a sinusoidal time dependence:

$$\delta\gamma_q = \alpha_q e^{i\omega t}; \tag{2.83}$$

$$\delta\gamma_i = \alpha_i e^{i\omega t}, \tag{2.84}$$

where the α's are complex. Substituting this into equations (2.80) and (2.81) and gathering terms, I get the matrix equation:

$$\left(\mathbf{V} - \omega^2\mathbf{T}\right)\vec{\alpha} = \mathbf{0}, \tag{2.85}$$

where

$$\mathbf{V} = \begin{pmatrix} \frac{1}{L_q^*} & -\frac{1}{L} \\ -\frac{1}{L} & \frac{1}{L_i^*} \end{pmatrix}; \tag{2.86}$$

$$\mathbf{T} = \begin{pmatrix} C_q & 0 \\ 0 & C_i \end{pmatrix}; \tag{2.87}$$

$$\vec{\alpha} = \begin{pmatrix} \alpha_q \\ \\ \alpha_i \end{pmatrix}. \tag{2.88}$$

Equation 2.85 has nontrivial solutions only when $\det(\mathbf{V} - \omega^2\mathbf{T}) = 0$. Let $\lambda = \omega^2$. Then I get the quadratic equation in λ

$$\lambda^2 - \left(\frac{1}{L_q^* C_q} + \frac{1}{L_i^* C_i}\right)\lambda + \left(\frac{1}{L_q^* C_q}\right)\left(\frac{1}{L_i^* C_i}\right) - \left(\frac{1}{LC_q}\right)\left(\frac{1}{LC_i}\right) = 0. \tag{2.89}$$

Rewriting this with the substitutions

$$\lambda_q = \frac{1}{L_q^* C_q}; \tag{2.90}$$

$$\lambda_i = \frac{1}{L_i^* C_i}; \tag{2.91}$$

$$\lambda_L^2 = \frac{1}{L^2 C_i C_q}; \tag{2.92}$$

yields

$$\lambda^2 - \lambda(\lambda_q + \lambda_i) + \lambda_q \lambda_i - \lambda_L^2 = 0. \tag{2.93}$$

To gain insight, consider two possible limits: $L \to 0$ and $L \to \infty$. In the first limit, the parallel inductances $L_i^* \to L$, so the constant term vanishes, yielding the two solutions

$$\begin{aligned} \lambda &= 0, \\ \lambda &= \frac{1}{L}\left(\frac{1}{C_q} + \frac{1}{C_i}\right). \end{aligned} \tag{2.94}$$

The first is a dc mode; the second has the frequency associated with the series combination of the loop inductance and both junction capacitances. Physically these modes make sense. Since in this limit the loop inductance looks like a short relative to the junction inductances, the dc mode moves both capacitor voltages together in phase, not driving any current through the loop inductance. The high frequency mode corresponds to the capacitors moving out of phase and charge and discharge each other through the loop inductance.

In the $L \to \infty$ limit, the parallel inductances $L_i^* \to L_{ij}$. Thus many of the factors in (2.93) now look like the bare junction frequencies ω_q^2, ω_i^2. And in fact, solving the quadratic equation, we find that $\lambda = \omega_q^2$ and $\lambda = \omega_i^2$ are the acceptable modes. Here, the huge impedance presented by the loop inductance relative to the Josephson inductance means that the capacitors are completely decoupled and

oscillate at their own natural frequencies. Since as mentioned in Section 2.5.1 this was the parameter regime studied in the earlier paper by Mitra et al., we see that the earlier theory is the independent junction limit of this one.

For intermediate L cases — and to make a more quantitative estimate of the impact of the loop inductance — I can solve (2.93):

$$\begin{aligned}\lambda &= \frac{\lambda_q + \lambda_i}{2} \pm \frac{1}{2}\sqrt{(\lambda_q + \lambda_i)^2 - 4\lambda_q\lambda_i + 4\lambda_L^2}, \\ &= \frac{\lambda_q + \lambda_i}{2} \pm \frac{1}{2}\sqrt{(\lambda_q - \lambda_i)^2 + 4\lambda_L^2}.\end{aligned} \tag{2.95}$$

If $\lambda_L^2 \ll (\lambda_q - \lambda_i)^2$, then the square root can be expanded to give the approximate eigenvalues:

$$\lambda_Q = \lambda_q + \frac{\lambda_L^2}{\lambda_q - \lambda_i}; \tag{2.96}$$

$$\lambda_I = \lambda_i - \frac{\lambda_L^2}{\lambda_q - \lambda_i}. \tag{2.97}$$

So far, those are just the frequencies associated with the normal modes. The modes themselves are the vectors that solve the generalized eigenvalue problem

$$\mathbf{V}\vec{\alpha} = \omega^2\mathbf{T}\vec{\alpha} \tag{2.98}$$

for the correct eigenvalues. These eigenvectors serve as column vectors in a basis transformation matrix for moving from the normal modes to the junction coordinates:

$$\delta\gamma_q = A_{qQ}\eta_Q + A_{qI}\eta_I \tag{2.99}$$

$$\delta\gamma_i = A_{iQ}\eta_Q + A_{iI}\eta_I, \tag{2.100}$$

where I have denoted the normal coordinates η_n. To emphasize the difference be-

tween normal coordinates and junction coordinates, I will refer to the qubit and isolation normal modes using capital letters in the subscripts.

The transformation matrix $\mathbf{A}$ is normalized such that

$$\mathbf{A}^T\mathbf{T}\mathbf{A} = 1, \tag{2.101}$$

where 1 indicates the identity matrix. This normalization represents a rescaling of the masses of the problem — which are the weights in the kinetic energy matrix $\mathbf{T}$ — to 1. This is very convenient and is the most popular choice of normalization convention for normal mode calculations. For asymmetrical junctions, it has the counterintuitive property that the new normal coordinates do not appear orthogonal when plotted in terms of $\delta\gamma_i$ and $\delta\gamma_q$.

It is helpful to know the relative strength of the two junction coordinates in the qubit mode in the weakly coupled approximation. Picking the qubit mode frequency λ_Q from (2.96):

$$\begin{aligned}
\frac{1}{L_q^*}A_{qQ} - \frac{1}{L}A_{iQ} &= \left(\lambda_q + \frac{\lambda_L^2}{\lambda_q - \lambda_i}\right) C_q A_{qQ}, \\
\lambda_q A_{qQ} - \frac{1}{LC_q}A_{iQ} &= \left(\lambda_q + \frac{1}{L^2 C_i C_q(\lambda_q - \lambda_i)}\right) A_{qQ}, \\
A_{iQ} &= \frac{1}{LC_i(\lambda_q - \lambda_i)}A_{qQ}.
\end{aligned} \tag{2.102}$$

Substituting for λ_q and λ_i, this factor becomes

$$\begin{aligned}
\frac{A_{iQ}}{A_{qQ}} &= \frac{1}{LC_i}\frac{1}{\frac{1}{LC_q} + \frac{1}{L_{qj}C_q} - \frac{1}{LC_i} - \frac{1}{L_{ij}C_i}}, \\
&= \frac{L_{ij}}{\frac{C_i}{C_q}\left(L_{ij} + \frac{L_{ij}}{L_{qj}L}\right) - L_{ij} - L}.
\end{aligned} \tag{2.103}$$

Performing a similar calculation for the ratio of qubit and isolation associated with

the new isolation junction mode, I get

$$\begin{aligned} A_{qI} &= \frac{1}{LC_q(\lambda_q - \lambda_i)} A_{iI}, \\ \frac{A_{qI}}{A_{Ii}} &= \frac{C_i}{C_q}\frac{A_{iQ}}{A_{qQ}}. \end{aligned} \tag{2.104}$$

In the units where the masses are 1, the conjugate momenta are now just the time derivative of the phase, $\varrho_I = \eta_I'$. The transformed normal mode Hamiltonian is

$$\mathcal{H}_{\text{norm}} = \varrho_Q^2 + \frac{1}{2}\omega_Q^2\eta_Q^2 + \varrho_I^2 + \frac{1}{2}\omega_I^2\eta_I^2. \tag{2.105}$$

The coupling has been absorbed into the coordinates so that now the Hamiltonian is just that of two noninteracting harmonic oscillators. This is the correct approximate basis for including the nonlinear perturbations.

To make the calculation of matrix elements easier, I convert to the ladder operator formalism. Using the usual translation from position and momenta to creation and annihilation operators,

$$\eta_Q = \sqrt{\frac{\hbar}{2m_Q\omega_Q}}\left(\hat{a}_Q^\dagger + \hat{a}_Q\right) \tag{2.106}$$

$$\varrho_Q = \sqrt{\frac{\hbar m_Q\omega_Q}{2}}\left(\hat{a}_Q^\dagger - \hat{a}_Q\right). \tag{2.107}$$

Because of the normal coordinate transform, $m_Q = 1$. The mass scaling is incorporated into the transform matrix $\mathbf{A}$.

After applying these scalings to the system, we end up in what I will call the "natural" normal coordinates for the system, which I will denote

$$\zeta_Q \equiv \frac{\eta_Q}{\eta_{Q0}} = \frac{\eta_Q}{\sqrt{\frac{\hbar}{2m_Q\omega_Q}}}. \tag{2.108}$$

If we define a transformation matrix

$$\mathbf{B} = \begin{pmatrix} \eta_{Q0} A_{qQ} & \eta_{I0} A_{qI} \\ \eta_{Q0} A_{iQ} & \eta_{I0} A_{iI} \end{pmatrix} \tag{2.109}$$

then $\delta\vec{\gamma} = \mathbf{B}\vec{\zeta}$ will be the final linear transform between the original junction coordinates and the natural normal coordinates. The reason for using the natural normal coordinates is that in the ζ system, the wavefunctions for the harmonic oscillator states adopt the familiar form:

$$\Psi_{Qn} = \sqrt{\frac{1}{2^n n!}}\, \pi^{-\frac{1}{4}} e^{\frac{-\zeta_Q^2}{2}} H_n(\zeta), \tag{2.110}$$

where $H_n(\zeta)$ is the nth Hermite polynomial. Using $\mathbf{B}^{-1}$, I transform these back into the $\delta\gamma$ coordinate system and plot them superposed on the potential.

This is done in Figure 2.12. Observe how the lobes of the excited state are stretched out along the long axis of the potential well. In a single junction approximation, with the qubit junction on the vertical axis, the lobes of the excited state would be oriented along the y-axis. This is an excellent visual way of demonstrating the preferability of the normal modes for analyzing the system.

This construction of the normal mode basis provides the harmonic background for perturbation theory. I now reincorporate the cubic terms as a perturbative shift. Expressing the harmonic oscillator terms as number operators,

$$\mathcal{H}_{\text{nat}} = \hbar\omega_Q \left(\hat{n}_Q + \frac{1}{2}\right) + \hbar\omega_I \left(\hat{n}_I + \frac{1}{2}\right) - \frac{1}{6}\frac{\Phi_0}{2\pi}\left(I_{0q} \sin\gamma_{qm}\ \delta\gamma_q^3 + I_{0i} \sin\gamma_{im}\ \delta\gamma_i^3\right). \tag{2.111}$$

Before applying $\mathbf{B}$ to convert this to its final form, it is worth commenting on the apparent scale of the perturbation. The cubic terms are each proportional to the Josephson energies of the two junctions, E_{jq} and E_{ji}. Since the normal modes are

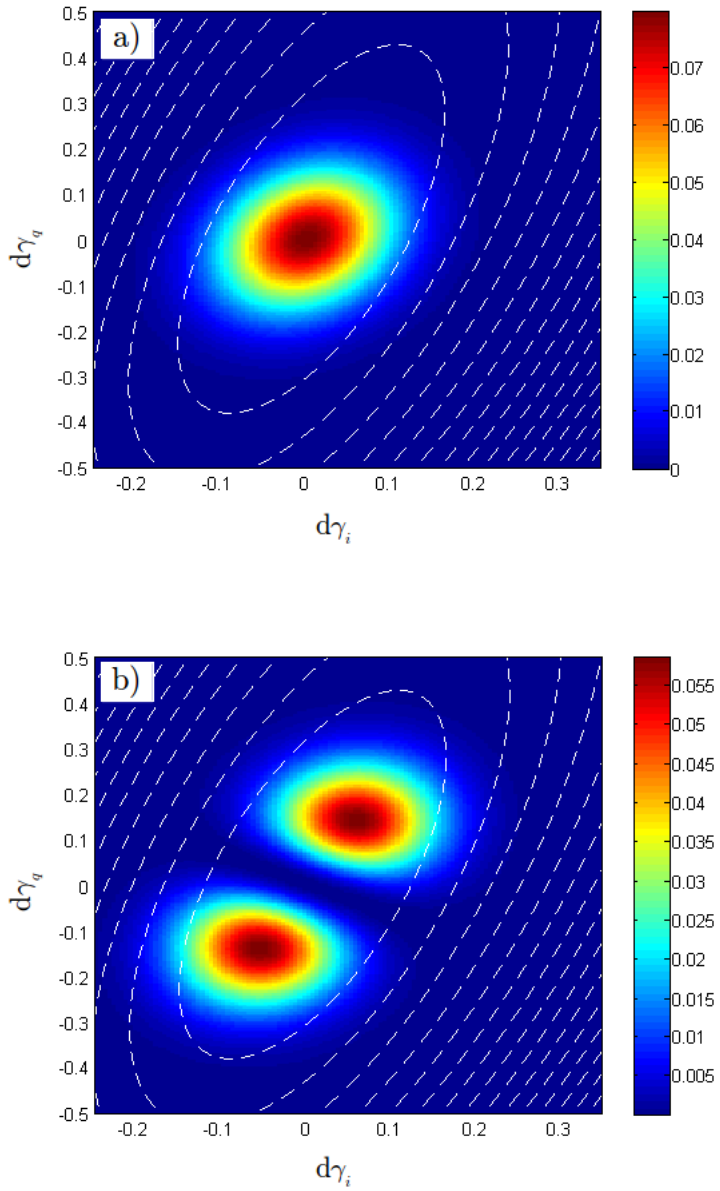

Figure 2.12: Plot of the probability density of the a) ground and b) first excited states of the qubit normal mode. The white dashed lines are contour lines of the SQUID potential.

of basically the same order of frequency as the junction modes, we can use what we know about phase qubits to estimate this scale difference quickly. The energy of the unbiased plasma oscillations of a bare junction can be written as $\sqrt{8E_J E_C}$, so dividing through by this factor suggests that the relative weight of the perturbation terms goes like roughly $\sqrt{\frac{E_J}{8E_C}}$. For phase qubits, we expect the ratio of these energy scales to be quite large — figures like 10^4 are not unreasonable — and so it seems that our perturbation is much bigger than the quadratic it perturbs. However, it turns out that for most bias conditions the phase scaling comes to the rescue again, and in transforming from $\delta\gamma$ to ζ we do indeed get a pretty small correction to the original quadratic.

Converting to the final form for the perturbation involves quite a bit of algebra. To clean up the notation a bit, I label the prefactors F:

$$F_q = -\frac{1}{6}\frac{\Phi_0}{2\pi}\left(I_{0q}\sin\gamma_{qm}\right); \tag{2.112}$$

$$F_i = -\frac{1}{6}\frac{\Phi_0}{2\pi}\left(I_{0i}\sin\gamma_{im}\right). \tag{2.113}$$

Then from the definition (2.109) of $\mathbf{B}$,

$$\begin{aligned}\mathcal{H}_{\text{pert}} = &\left(F_q B_{qQ}^3 + F_i B_{iQ}^3\right)\zeta_Q^3 + 3\left(F_q B_{qQ}^2 B_{qI} + F_i B_{iQ}^2 B_{iI}\right)\zeta_Q^2\zeta_I \\ &+ 3\left(F_q B_{qQ} B_{qI}^2 + F_i B_{iQ} B_{iI}^2\right)\zeta_Q\zeta_I^2 + \left(F_q B_{qI}^3 + F_i B_{iI}^3\right)\zeta_I^3.\end{aligned} \tag{2.114}$$

Using the commutation properties of ladder operators, I expand the ζ terms:

$$\zeta_Q^3 = (\hat{a}_Q^\dagger)^3 + 3\hat{a}_Q^\dagger(\hat{n}_Q + 1) + 3\hat{a}_Q\hat{n}_Q + (\hat{a}_Q)^3; \tag{2.115}$$

$$\zeta_Q^2\zeta_I = (\hat{a}_Q^\dagger)^2\hat{a}_I^\dagger + (\hat{a}_Q)^2\hat{a}_I^\dagger + (2\hat{n}_Q + 1)\hat{a}_I^\dagger \tag{2.116}$$

$$+ (\hat{a}_Q^\dagger)^2\hat{a}_I + (\hat{a}_Q)^2\hat{a}_I + (2\hat{n}_Q + 1)\hat{a}_I; \tag{2.117}$$

$$\zeta_I^2\zeta_Q = (\hat{a}_I^\dagger)^2\hat{a}_Q^\dagger + (\hat{a}_I)^2\hat{a}_Q^\dagger + (2\hat{n}_I + 1)\hat{a}_Q^\dagger \tag{2.118}$$

$$+ (\hat{a}_I^\dagger)^2\hat{a}_Q + (\hat{a}_I)^2\hat{a}_Q + (2\hat{n}_I + 1)\hat{a}_Q; \tag{2.119}$$

$$\zeta_I^3 = (\hat{a}_I^\dagger)^3 + 3\hat{a}_I^\dagger(\hat{n}_I + 1) + 3\hat{a}_I\hat{n}_I + (\hat{a}_I)^3. \tag{2.120}$$

There is minimal simplification to be performed here. There are two terms linear in $\hat{a}_Q^\dagger$ and $\hat{a}_Q$ (and their I mode analogues) but that appears to be it.

The perturbation theory is useful for calculations addressing specifically nonlinear effects. The most important of these for typical qubit operation is what is generally referred to as the anharmonicity, the difference $\omega_{01} - \omega_{12}$, but two-photon processes are another area where the perturbative wavefunctions are useful. While the perturbation theory also imposes some corrections to the harmonic energies, in practice these are usually fairly small. For many applications, the harmonic approximation gives a good picture of where the qubit frequency will be.

2.5.3 Normal mode theory and the PB9 device

I originally developed the normal modes model in my work on applying Josephson bifurcation amplification (JBA) [79, 80] to a capacitively shunted dc SQUID phase qubit. The first practical application of the new model was to the PB9 qubit measured by Tony Przybysz and Rangga Budoyo. The work surrounding this device was the basis for Tony's thesis [28]. I will discuss some of the particular highlights and implications of the model that Tony touched upon.

Figure 2.13 shows two optical images of the PB9 device. The first image shows

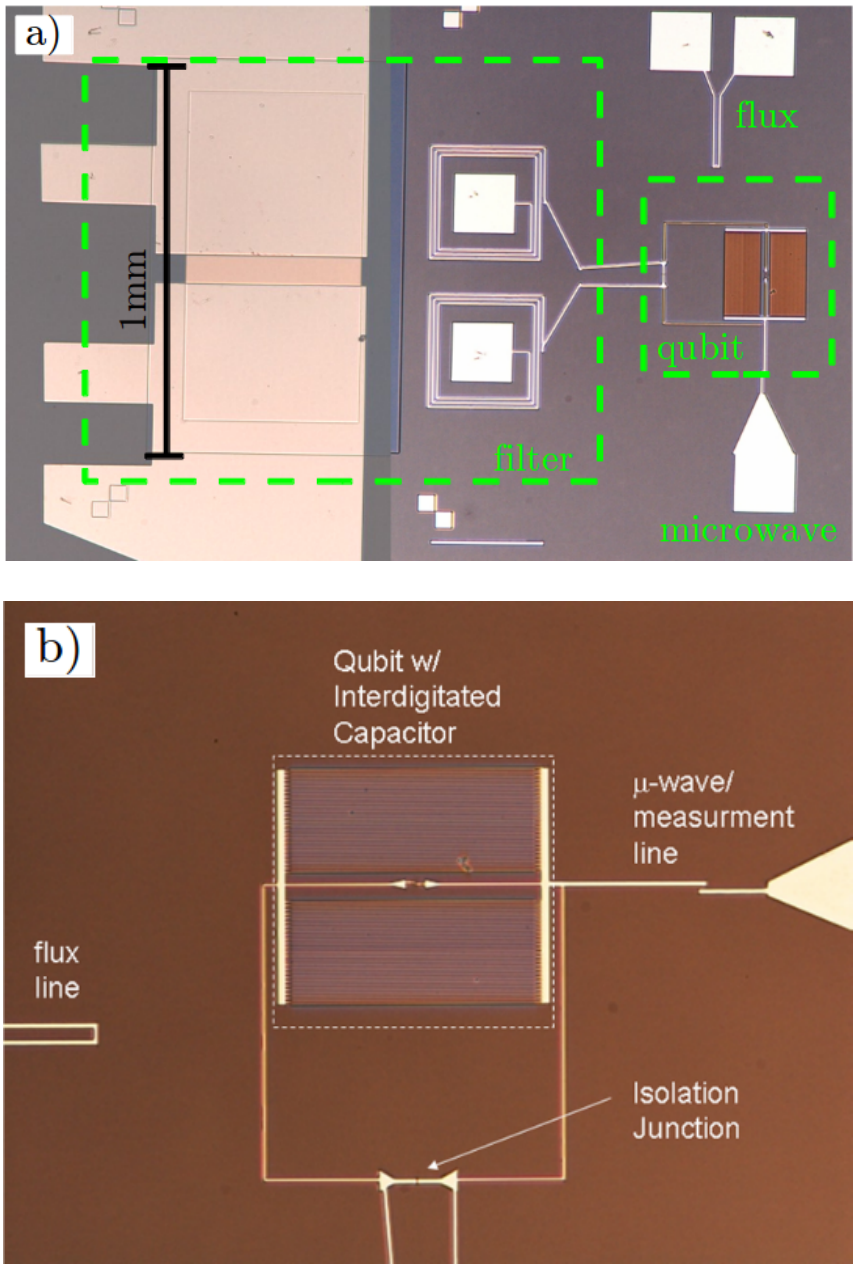

Figure 2.13: Optical images of the PB9 device. a) Image of the whole device; most of what is seen is the LC filter. b) Close-up of the qubit region.

the whole device. The leftmost features are the contact pads, which also form plates for two series parallel-plate capacitors. From there, wirebonds (not photographed) connect the plates to the contact pads inside of the spiral inductors. These spiral inductors are then connected directly to the qubit. In the upper right, the flux bias coil is visible. In the bottom right is the microwave bias line. This is weakly capacitively coupled to a finger protruding from the dc SQUID, allowing microwaves to enter the SQUID for qubit control.

The second image shows a close-up of the SQUID loop. At the bottom, the leads from the spiral inductors can be seen joining up with the SQUID loop. The short distance between where these leads join the SQUID gives the inductive imbalance $L_i \ll L_q$. This short section of SQUID loop is interrupted by a relatively large area Josephson junction, formed by double angle evaporation[5]. On the other side of the SQUID loop is a smaller Josephson junction I_{0q}. This junction is shunted by a large interdigitated capacitor, boosting E_J/E_C ratio and using the very low loss sapphire substrate as a dielectric.

One of the main signs of the failure of the independent junction approximation was difficulty in recovering plausible values for the junction capacitances given the observed frequencies. According to Przybysz, the best choice of capacitance for a single junction fit was 500 fF, a factor of 3 less than the anticipated design value. Similar problems occurred in my earlier JBA samples. Intuitively, the issue is that the relevant inductance for calculating the frequency of the qubit mode is the parallel combination of the loop inductance $L_i + L_q$ with L_{jq}. For inductors of comparable size, this parallel combination is a factor of two smaller than either component. With small junctions and high bias, the reduction in effective inductance can be even greater because the parallel combination will be dominated by the loop inductance. This decreasing inductance leads to an increase in the frequency for

[5]Double angle evaporation is discussed in Chapter 4.

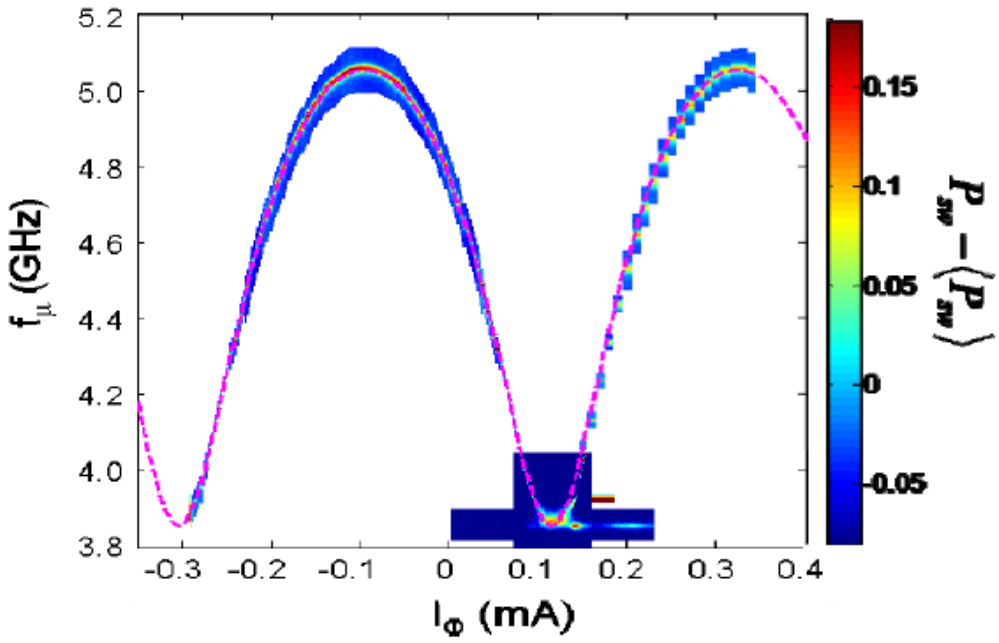

Figure 2.14: Plot of experimental spectroscopy data from the PB9 device with a fit to the normal mode model. The model is slightly adjusted to include parasitic inductance from the interdigitated shunt capacitor. The plot is from Anthony Przybysz's thesis [28].

a given capacitance. If we do not account for the change in inductance, we must postulate an excessively low capacitance to capture the frequency correctly.

Figure 2.14 shows spectroscopy data from PB9. Bias current is held fixed; the x-axis represents change in the applied flux. The y-axis is the frequency of applied microwaves. The colorbar measures the enhancement in measured probability of tunneling to the voltage state at given bias conditions, our proxy for the qubit state. Overlaid on the color plot as a magenta dashed line is the prediction from a modified normal mode theory, given the relevant parameters from the SQUID. Although the agreement of the unmodified theory was good, there were regions where it was off a little bit. The modification consisted of incorporating parasitic inductance from the interdigitated shunt capacitor[6].

The spectroscopic data also shows regions where the frequency is extremized

[6]The techniques of Section 2.2 can handle such a corrected model directly at the cost of more complexity. Przybysz's approach basically amounts to tuning the qubit inductance slightly.

and insensitive to small variations in the applied flux. This leads to a "sweet spot" [81] where dephasing due to flux noise is suppressed. In flux and charge qubits, sweet spots are a well-known and often exploited technique. In phase qubits, sweet spots are novel. In the traditional single junction phase qubit, where

$$\omega_q = \sqrt{\frac{1}{L_J C}} \propto \sqrt{\cos\gamma} = \left(1 - \sin^2\gamma\right)^{\frac{1}{4}} = \left(1 - \left(\frac{I_b}{I_0}\right)^2\right)^{\frac{1}{4}}, \tag{2.121}$$

the only point where $\partial\omega_q/\partial I_b = 0$ is at zero bias. However, at zero bias, low anharmonicity and difficulty making voltage based measurements usually prevent phase qubit use. rf SQUID phase qubits have the same behavior in this regard. Thus, the appearance of sweet spots in a phase qubit was quite surprising. Chapter 10 of Przybysz's thesis [28] examines the sweet spot behavior of the normal modes model in more detail, including the surprising prediction that the normal modes model exhibits a sweet spot for critical current noise.

2.6 Summary

In this chapter, I reviewed the process of quantizing superconducting circuits. I began with a discussion of Josephson junctions, which provide the necessary nonlinearity to limit superconducting circuits to two levels. I then discussed how nodal analysis is used to generate Hamiltonians of circuits. Using this technique, I derived the Hamiltonian for the dc SQUID. I discussed how filtering improves lifetime T_1 and how this suggested the idea of a dc SQUID phase qubit. Lastly, I discussed my model for the dc SQUID phase qubit where the nonlinearity is a perturbation to the normal modes of the SQUID, and mentioned the successes of this model with our earlier qubit PB9.

Chapter 3

Coupling resonators to qubits: the Jaynes-Cummings model

Recent years have seen extensive interplay between condensed matter and AMO physics. Each field contributes both techniques and interesting problems to the other, offering opportunities to study these problems in completely new ways. In AMO, for example, classic solid-state ideas inspired photonic crystals [82], using Bloch's theorem and periodic structures to control light through local band gaps. Considerable research effort has been spent on dilute gases of atoms in optical lattices [83, 84]. Counter-propagating laser beams create periodic potentials that host a small number of atoms. These systems are inherently very clean compared to solid state systems and afford great control over interaction strengths. Thus optical lattices are a fantastic testbed for exploring interesting limits of condensed matter systems, such as superfluid-Mott insulator transitions or BEC-BCS crossover.

Superconducting quantum information has drawn heavily from the AMO playbook. A superconducting qubit can be thought of as an artificial atom with properties that can be tailored for specific applications. As such, many tricks used to manipulate conventional atoms find applications to superconducting qubits as well. Sideband cooling — pumping to a third level with faster relaxation to the ground state — allowed active cooling in a flux qubit [85]. A superconducting qubit was used as the gain medium in a single atom maser [86]. Electromagnetically induced transparency and coherent population trapping have been studied by several groups [87–89]. There has recently been renewed interest in superconducting parametric amplifiers [90–92] motivated by the goal of quantum limited measurement. These

have been used to produce squeezed states [93, 94]. A more comprehensive recent review of AMO ideas in superconducting QC is [95].

The most influential contribution of AMO physics to superconducting qubits has been "circuit QED." This is the superconducting analogue of cavity QED [96], where atoms are sent through a high Q microwave cavity. While in the cavity, the atoms interact with a single electromagnetic field mode. By limiting the field to a single mode, one can describe the atom-light interaction using a surprisingly simple model originally due to Jaynes and Cummings. I begin by discussing the Jaynes-Cummings Hamiltonian and then describe circuit QED in more detail. Lastly, I discuss an extension of the dc SQUID model developed in Section 2.5. This new model includes an on-chip LC filter. The filter reduces high frequency noise on the bias line to improve the qubit lifetime T_1. After deriving the Hamiltonian for the filter plus SQUID model, I show how to derive an approximation to the full Hamiltonian with a Jaynes-Cummings Hamiltonian, with special attention paid to extracting the filter-qubit coupling g.

3.1 Jaynes-Cummings model

The Jaynes-Cummings model was proposed 50 years ago in a discussion of semiclassical vs. quantum theories of the maser [97]. It is a simple, exactly solvable model of the interaction between a single mode of the electromagnetic field and a two-level atom. Following the usual quantization of electromagnetic fields [98], the field mode is represented as a harmonic oscillator. The coupling between the two systems is electric dipole, further simplified by making a rotating wave approximation (RWA). As discussed in Section 2.4.2, the RWA discards high frequency terms. For Jaynes-Cummings, this means discarding coupling terms where both the atom and the field mode gain or lose quanta, since the time dependence of the creation or annihilation operators leads to contributions at the sum of the two frequencies in

that situation.

The Jaynes-Cummings Hamiltonian is

$$\mathcal{H}_{\mathrm{JC}} = \hbar\omega_r \hat{a}^\dagger \hat{a} + \hbar\omega_q \hat{\sigma}_z + \imath\hbar g \left(\hat{a}^\dagger \hat{\sigma}_z^- - \hat{a}\hat{\sigma}_z^+\right). \tag{3.1}$$

The subscript choices reflect the language I will use in referring to these two subsystems: the resonator and the qubit.

What makes this Hamiltonian tractable is a special property of the coupling term. Since the coupling term creates excitations in one subsystem while destroying them in the other, the total excitation number is conserved by the Hamiltonian. Because the qubit has only two levels, this results in the total Hamiltonian being a collection of 2x2 Hamiltonians corresponding to different numbers of excitations. This structure is visually apparent when the first levels of $\mathcal{H}_{\mathrm{JC}}$ are written in matrix form:

$$\mathcal{H}_{\mathrm{JC}} = \begin{array}{c} \\ |0g\rangle \\ |0e\rangle \\ |1g\rangle \\ |1e\rangle \\ |2g\rangle \\ \vdots \end{array} \begin{array}{c} \begin{array}{cccccc} |0g\rangle & |0e\rangle & |1g\rangle & |1e\rangle & |2g\rangle & \dots \end{array} \\ \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & \hbar\omega_q & -\imath\hbar g & 0 & 0 & \dots \\ 0 & \imath\hbar g & \hbar\omega_r & 0 & 0 & \dots \\ 0 & 0 & 0 & \hbar\omega_q + \hbar\omega_r & -\imath\sqrt{2}\hbar g & \dots \\ 0 & 0 & 0 & \imath\sqrt{2}\hbar g & 2\hbar\omega_r & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \end{array}, \tag{3.2}$$

where I have chosen the basis $(|0g\rangle, |0e\rangle, |1g\rangle, |1e\rangle, \dots)$. g and e denote the ground or excited state of the qubit and the number is the number of photons in the resonator.

Equation 3.2 illustrates another important feature of $\mathcal{H}_{\mathrm{JC}}$. The raising and lowering operators $\hat{a}^\dagger$ and $\hat{a}$ pick up factors of $\sqrt{n}$ when operating on states with $n-1$ ($\hat{a}^\dagger$) or n ($\hat{a}$) photons in the resonator. Thus, the effective coupling strength depends

on the resonator photon number. This effect produces many of the characteristic features of the Jaynes-Cummings model.

Because it breaks down into discrete blocks, one diagonalizes $\mathcal{H}_{\mathrm{JC}}$ by diagonalizing the individual blocks. The ground state $|0g\rangle$ is already diagonal, so all that remains is diagonalizing the subspace with $n > 0$ total excitations. In this subspace,

$$\mathcal{H}_n = \begin{pmatrix} \hbar\omega_q + (n-1)\hbar\omega_r & -\imath\sqrt{n}\hbar g \\ \imath\sqrt{n}\hbar g & n\hbar\omega_r \end{pmatrix}. \tag{3.3}$$

I define the detuning $\Delta \equiv \omega_q - \omega_r$ and rewrite $\mathcal{H}_n$ as

$$\frac{1}{\hbar}\mathcal{H}_n - n\omega_r = \begin{pmatrix} \Delta & -\imath\sqrt{n}g \\ \imath\sqrt{n}g & 0 \end{pmatrix}, \tag{3.4}$$

where to simplify the form of the matrix I have adjusted the zero of energy. Finding the eigenvalues of this matrix and adding the energy $n\hbar\omega_r$ back,

$$\frac{1}{\hbar}E_{n\pm} = n\omega_r + \frac{\Delta}{2} \pm \frac{1}{2}\sqrt{\Delta^2 + 4ng^2}. \tag{3.5}$$

There are two important limits to consider. The first is when the qubit and resonator are in resonance, $\Delta \to 0$. The coupling g lifts the degeneracy, yielding energies $E_n = n\hbar\omega_r \pm \sqrt{n}g$ and eigenstates $\frac{1}{\sqrt{2}}(|n, g\rangle \pm |n-1, e\rangle)$. States with n excitations form a ladder of paired states at harmonics of the resonator frequency. The energy spacing of states in a pair increases as the square root of n.

The other important limit is the dispersive limit, where the detuning is large relative to the coupling. Since the coupling depends on the excitation number n, the quantitative description of the dispersive limit is $2\sqrt{n}g/\Delta \ll 1$. In this limit,

one can perform a Taylor expansion of the radical in the eigenenergies:

$$\begin{aligned}\frac{1}{\hbar}E_{n\pm} &= n\omega_r + \frac{\Delta}{2} \pm \frac{1}{2}\sqrt{\Delta^2 + 4ng^2}, \\ &\approx n\omega_r + \frac{\Delta}{2} \pm \frac{\Delta}{2}\left(1 + \frac{2ng^2}{\Delta^2}\right).\end{aligned} \tag{3.6}$$

Expanding the plus and minus cases,

$$E_{n+} = n\hbar\omega + \hbar\Delta + \hbar\frac{ng^2}{\Delta}; \tag{3.7}$$

$$E_{n-} = n\hbar\omega - \hbar\frac{ng^2}{\Delta}. \tag{3.8}$$

(Note that I assume $\Delta > 0$; a very similar analysis goes through for negative Δ with the role of the plus and minus states swapped.) The new eigenstates are of the form $a_\pm |n-1, e\rangle + b_\pm |n, g\rangle$. With the Taylor expanded eigenenergies, the eigenstates satisfy

$$b_+ = \frac{\sqrt{n}g}{\Delta}a_+, \tag{3.9}$$

$$a_- = -\frac{\sqrt{n}g}{\Delta}b_- \tag{3.10}$$

In words, in the dispersive limit the mixing of qubit states and resonator states is weak. The eigenstates with energy E_{n+} are approximately $|n-1, e\rangle$; those with energy E_{n-} are approximately $|n, g\rangle$.

This weak mixing suggests considering the Jaynes-Cummings Hamiltonian as a perturbation on the original qubit-resonator basis. How does this perturbation affect the qubit and resonator transitions? Since the coupling strength depends on the number of photons in the resonator, I expect the qubit transition to depend on this photon number. From the discussion following (3.10), the state with excited qubit and n photons in the resonator has energy $E_{n+1,+}$. The ground state qubit

with n photons in the resonator has energy $E_{n,-}$. Subtracting these yields

$$\begin{aligned}\hbar\omega_{qn} &= E_{n+1,+} - E_{n-}, \\ &= \left(n\hbar\omega_r + \hbar\omega_q + \hbar\frac{(n+1)g^2}{\Delta}\right) - \left(n\hbar\omega_r - \hbar\frac{ng^2}{\Delta}\right), \\ &= \hbar\omega_q + \hbar\frac{(2n+1)g^2}{\Delta}.\end{aligned} \tag{3.11}$$

Thus, the coupling between the resonator and the qubit leads to a shift in the frequency of the qubit. Each photon in the resonator contributes an energy shift $2\hbar g^2/\Delta$. Since the dispersive limit assumed that $2g/\Delta \ll 1$, this is a small fraction of the coupling $\hbar g$.

A similar calculation shows that the resonator frequency moves depending on the state of the qubit. When the qubit is excited, adding one photon to the resonator takes an energy

$$\begin{aligned}\hbar\omega_{re} &= E_{n+1,+} - E_{n+}, \\ &= \left(n\hbar\omega_r + \hbar\omega_q + \hbar\frac{(n+1)g^2}{\Delta}\right) - \left((n-1)\hbar\omega_r + \hbar\omega_q + \hbar\frac{ng^2}{\Delta}\right), \\ &= \hbar\omega_r + \frac{\hbar g^2}{\Delta}.\end{aligned} \tag{3.12}$$

When the qubit is in the ground state, adding one photon to the resonator takes an energy

$$\begin{aligned}\hbar\omega_{rg} &= E_{n+1,-} - E_{n-}, \\ &= \left((n+1)\hbar\omega_r - \hbar\frac{(n+1)g^2}{\Delta}\right) - \left(n\hbar\omega_r - \hbar\frac{ng^2}{\Delta}\right), \\ &= \hbar\omega_r - \frac{\hbar g^2}{\Delta}.\end{aligned} \tag{3.13}$$

The difference between these two energies is the

$$\hbar\omega_{re} - \hbar\omega_{rg} = \frac{2\hbar g^2}{\Delta}. \tag{3.14}$$

Thus the resonator frequency sees the same $2g^2/\Delta$ "photon shift" — since the qubit is a two-level system, though, at most one photon can be in the qubit. Experimental methods for exploiting both types of frequency shift have been highly fruitful in superconducting qubits.

3.2 Circuit QED

Circuit QED (cQED) using superconducting qubits was first proposed by Blais et al. [30] in 2004. Cavity QED requires atoms and cavities. In [30], the cQED atom was a superconducting charge qubit. The cavity analogue is a 1-D microwave coplanar waveguide resonator. The resonator is formed by interrupting the center conductor of the waveguide in two places, giving a weak capacitive coupling. The impedance mismatch at these breaks reflects incident microwaves. Like in a Fabry-Perot etalon, the resonant frequencies are those at which standing waves can be formed. By making the capacitance of the two breaks different, Blais et al. created a preferred output direction for the resonator. The qubit and resonator are coupled through the electric field. To give a strong coupling, the charge qubit was originally placed between the center conductor and ground plane of the coplanar waveguide at an anti-node of the resonator mode.

cQED offers some advantages over traditional cavity QED. Atomic cavity QED experiments use a fixed microwave cavity that atoms travel through. This gives the atoms only a limited amount of time for interaction with the cavity field. In cQED, the qubit is permanently located in the cavity. This allows for much longer atom-cavity interactions and thus more complicated manipulation of quantum states. By working with a 1-D resonator with a small gap, cQED circuits can store the photonic field energy over a small mode volume, therefore yielding strong fields. Further, superconducting qubits can be engineered to have dipole moments much larger than atoms. The combination of these two effects lets cQED realize the "strong-coupling

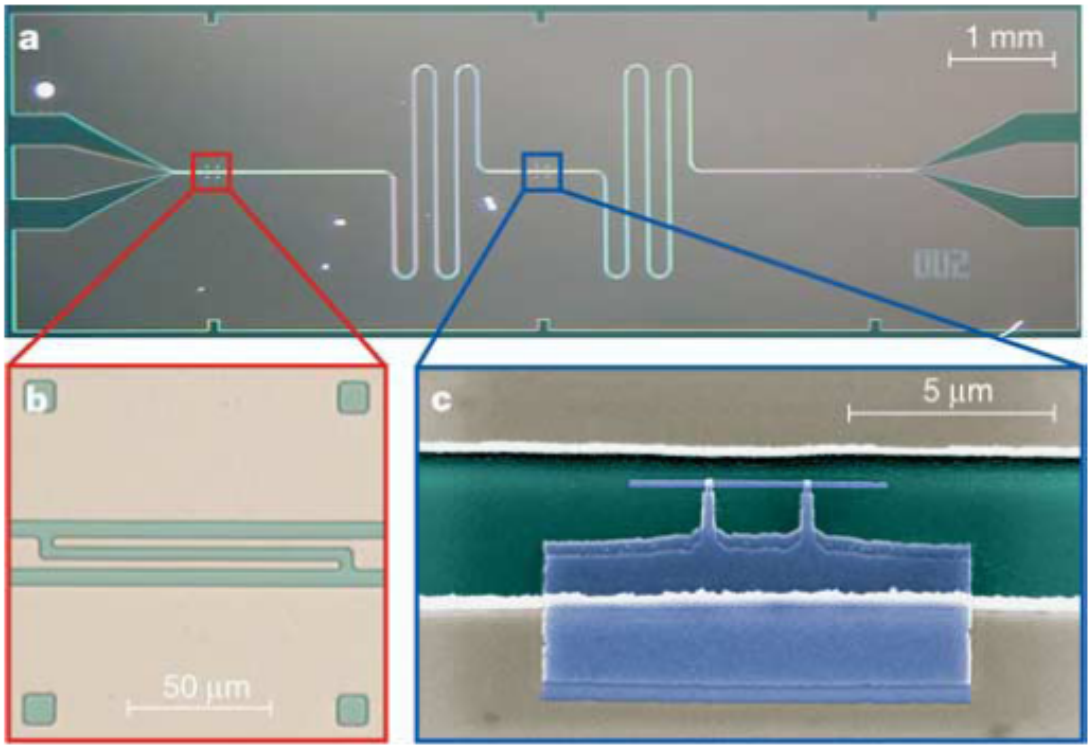

Figure 3.1: Images of an early circuit QED setup, taken from Wallraff et al. [99]. a) shows the chip; the major feature visible is the meandering coplanar waveguide. b) shows the capacitive coupling at one end of the resonator. c) is a false-color image showing the charge qubit, a Cooper pair box.

regime" where the qubit frequency shift (3.11) exceeds the linewidth. In this regime, peaks corresponding to distinct photon numbers in the resonator are visible.

Shortly after the proposal, an experimental realization of circuit QED was demonstrated by Wallraff et al. [99]. Figure 3.1 shows an image of their chip. To achieve the desired resonator frequency, the coplanar waveguide meanders for extra length. The capacitive couplings defining the cavity are formed by aluminum fingers on the center conductor. The qubit used in the original experiment was a Cooper pair box charge qubit. By adjusting the gate voltage applied to the box, the qubit frequency could be tuned in and out of resonance with the resonator. When tuned to resonance the coupling broke the degeneracy of the two states and led to a splitting of the resonator line. This splitting was detected using homodyne detection of a

microwave probe signal transmitted through the resonator.

While this was a notable accomplishment, more direct applicability to superconducting QC research was demonstrated in a pair of follow-up papers [100, 101]. These established the usefulness of circuit QED in the dispersive limit as a tool for qubit readout. Since the state of the qubit changes the frequency of the resonator as shown in (3.1), the phase of the transmitted microwave signal depends on the qubit state. Measuring the qubit then consists of averaging the microwave signal long enough to determine this phase.

This style of measurement is very attractive for several reasons. Unlike in the switching measurement described in Section 2.1, the qubit always remains superconducting. This allows for much faster experimental repetition rates, as no cooling period is necessary following each measurement. More significantly, at high measurement powers the measurement can be made projective, an essential element for quantum error correction. Low power measurements access the intriguing regime of weak quantum measurement [102, 103]. Communicating with the qubit solely through the resonator offers opportunities for frequency multiplexing. For example, a single feedline could contain resonators at several different frequencies, each of which contains qubits. This sort of reduction in wiring load is essential for operating significant numbers of qubits in a dilution refrigerator.

Another important result was the resolution of individual photon number states in the resonator by Schuster et al. [104] The resonator was driven at a fixed probe frequency. A second, spectroscopic tone was used to pump the qubit. From (3.11), the qubit frequency depended on the photon occupancy of the resonator. However, a classical microwave drive does not consist of a single photon number (Fock) state, but instead a coherent state that superposes many different Fock states. Each Fock state corresponded to a different resonant frequency for the qubit. Using the state-dependent resonator frequency (3.1), the spectroscopic lines corresponding

to each Fock state were visible in the response of the probe beam. By operating the probe beam at different powers — thereby changing the relative population in the Fock states — the height of the different Fock state peaks in the data changed, confirming that they were due to the photon shifts.

Following this paper, the Fock states of the resonator were seen as a possibly useful tool in their own right. Using the different coupling of each Fock state to the qubit, Hofheinz et al. [105] demonstrated the formation of Schrödinger cat states in a resonator coupled to a phase qubit. Since the resonators typically have no junctions and weak coupling to the microwave line, they are often long-lived relative to the qubits. This suggests an architecture where resonators are used both to store qubits and to couple qubits, transferring information from resonator to qubit only when necessary for operations [106].

One particularly influential trend in recent cQED experiment is the use of 3d cavities instead of planar ones. Paik et al. [107] demonstrated an order of magnitude enhancement in the lifetime of a transmon qubit by coupling it to a 3d microwave cavity. The underlying idea was to decrease the relative importance of fields near surfaces. By using a 3d cavity with greater mode volume, the fields were reduced and loss to surface defects decreased. Using careful processing and improved materials, it seems much of this performance gain can still be captured in planar systems [108].

3.3 Filtered dc SQUID

3.3.1 T_1 improvement from filter

In Section 2.4.3, I discussed how a dc SQUID phase qubit can reduce qubit loss by filtering high frequency current noise from the bias leads. Our dc SQUID qubits have had short T_1 times compared to competing superconducting approaches. As the bias leads represent one of the major architectural differences, adding more

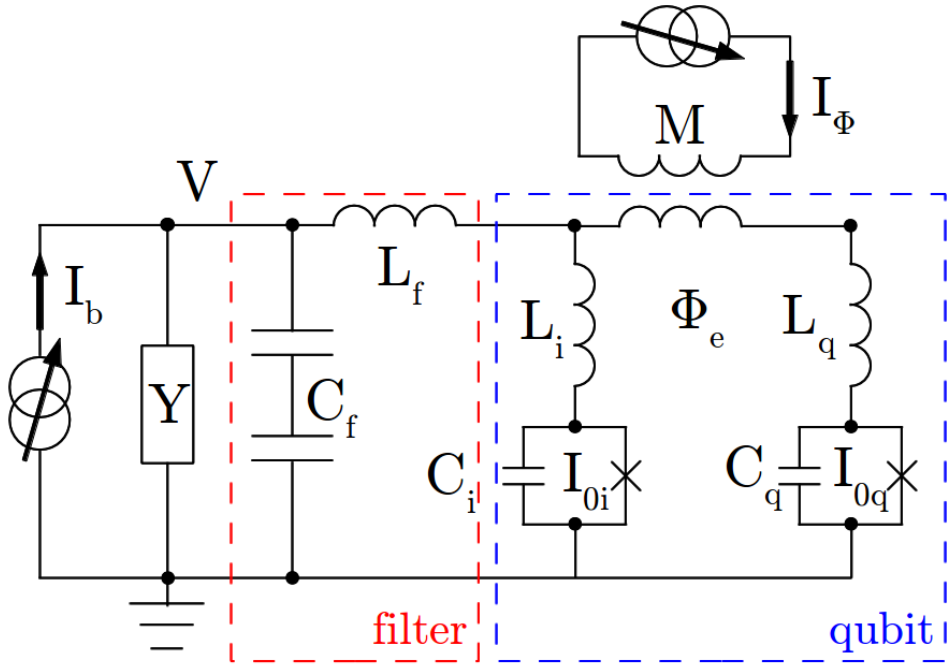

Figure 3.2: Schematic of the LC filtered dc SQUID. The filter and SQUID are indicated in the schematic.

filtering to the bias line was a logical change. Thus, we began working with dc SQUID phase qubits featuring an on-chip LC low pass filter on the bias line.

Figure 3.2 shows a schematic of a dc SQUID with an additional low-pass LC filter on the bias line. As discussed in Section 2.4.3, the loss to the bias leads depends on the admittance $Y(\omega)$. The LC filter increases this admittance, reducing loss to the bias leads and potentially improving T_1. Modeling the admittance as $Y(\omega) = 1/R$, looking at the bias line from the SQUID we see a series combination of L_f and the parallel combination of R and C_f:

$$Z_{\text{line}} = \imath\omega L_f + \frac{R}{1 + \imath\omega C_f R}; \tag{3.15}$$

$$Y_{\text{line}} = \frac{1 + \imath\omega C_f R}{\imath\omega L_f + R(1 - \omega^2 L_f C_f)}. \tag{3.16}$$

Denoting $L_f C_f \equiv 1/\omega_f^2$,

$$\begin{aligned} \mathrm{Re}Y_{\mathrm{line}} &= \frac{R\left(1-\omega^2/\omega_f^2\right)+R\omega^2/\omega_f^2}{R^2\left(1-\omega^2/\omega_f^2\right)^2+\omega^2L^2}; \\ \frac{1}{\mathrm{Re}Y_{\mathrm{line}}} &= R\left(1-\frac{\omega^2}{\omega_f^2}\right)^2+\frac{\omega^2L^2}{R}. \end{aligned} \tag{3.17}$$

In the dc limit $\omega \to 0$, the line impedance remains R. This is sensible, as the LC filter is low-pass and doesn't block dc signal. More important is the behavior for $\omega \gg \omega_f$. In this regime, $1/\mathrm{Re}Y_{\mathrm{line}} \approx (\omega/\omega_f)^4 R$. A filter with $\omega_Q = 10\omega_f$ leads to a 10^4 improvement in T_1; a filter with $\omega_Q = 30\omega_f$ gets almost a 10^6 improvement. An important note is that this is the admittance seen by the SQUID; thus, it leaves out the filtering effects of the SQUID itself on the qubit. Stacking the admittance transformations from both the filter and SQUID should basically eliminate loss to the bias line.

Several qubits in our lab used these LC filters [28, 29]. These qubits were much longer lived than their predecessors, with lifetimes improved by an order of magnitude. Given the apparent effectiveness, we kept the filter in the design for the BBC2 qubit I study in this thesis. The microwave measurements on BBC2 (the subject of Chapter 7) suggested the LC filter might be interacting with the qubit as more than just a filter. Given a microwave resonator (the filter) coupled to a qubit, it is natural to ask whether the Jaynes-Cummings physics of cQED might apply to our filter plus SQUID system.

In the rest of this chapter, I use the techniques of Chapter 2 to derive a Hamiltonian for the full filter plus SQUID system. Since the Jaynes-Cummings model is well understood, I show how to reduce the full Hamiltonian to an effective Jaynes-Cummings model for the filter and SQUID. In Section 7.3 I examine the consistency of this model with the measured microwave data from the BBC2 qubit.

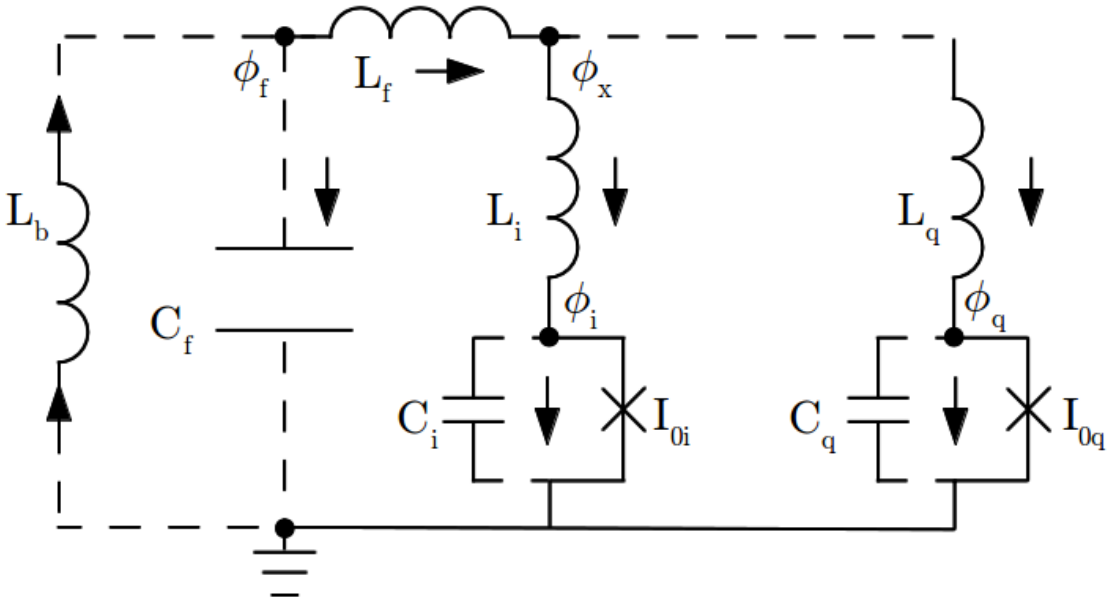

Figure 3.3: Decomposition of the LC filtered dc SQUID into spanning tree and closure branches.

3.3.2 Filtered SQUID Hamiltonian

Figure 3.3 decomposes the filtered LC SQUID schematic into the appropriate spanning tree. There are four active nodes, yielding four equations of motion:

$$-\frac{\phi_x - \phi_f}{L_f} + \frac{\phi_i - \phi_x}{L_i} + \frac{\phi_q - \phi_x + \Phi_e}{L_q} = 0; \tag{3.18}$$

$$-\frac{\phi_f + \Phi_b}{L_b} + \frac{\phi_x - \phi_f}{L_f} = C_f \ddot{\phi}_f; \tag{3.19}$$

$$-\frac{\phi_i - \phi_x}{L_i} - I_{0i} \sin \frac{2\pi\phi_i}{\Phi_0} = C_i \ddot{\phi}_i; \tag{3.20}$$

$$-\frac{\phi_q - \phi_x + \Phi_e}{L_q} - I_{0q} \sin \frac{2\pi\phi_q}{\Phi_0} = C_q \ddot{\phi}_q. \tag{3.21}$$

As with the dc SQUID calculation in Section 2.3.3, there is one extraneous node with no capacitance. I eliminate this node:

$$\phi_x(\frac{1}{L_f}+\frac{1}{L_i}+\frac{1}{L_q}) = \frac{\phi_f}{L_f}+\frac{\phi_i}{L_i}+\frac{\phi_q+\Phi_e}{L_q}; \tag{3.22}$$

$$\phi_x\frac{\mathfrak{L}^2}{L_fL_iL_q} = \frac{\phi_f}{L_f}+\frac{\phi_i}{L_i}+\frac{\phi_q+\Phi_e}{L_q}; \tag{3.23}$$

$$\phi_x = \frac{1}{\mathfrak{L}^2}(L_iL_q\phi_f + L_fL_q\phi_i + L_fL_i(\phi_q+\Phi_e)), \tag{3.24}$$

where I have defined

$$\mathfrak{L}^2 = L_iL_q + L_fL_q + L_fL_i. \tag{3.25}$$

The other equations contain three terms that depend on ϕ_x:

$$\frac{\phi_x-\phi_f}{L_f} = \frac{1}{L_f\mathfrak{L}^2}((L_iL_q-\mathfrak{L}^2)\phi_f + L_fL_q\phi_i + L_fL_i(\phi_q+\Phi_e); \tag{3.26}$$

$$\frac{\phi_x-\phi_i}{L_i} = \frac{1}{L_i\mathfrak{L}^2}(L_iL_q\phi_f + (L_fL_q-\mathfrak{L}^2)\phi_i + L_fL_i(\phi_q+\Phi_e); \tag{3.27}$$

$$\frac{\phi_x-(\phi_q+\Phi_e)}{L_q} = \frac{1}{L_q\mathfrak{L}^2}(L_iL_q\phi_f + L_fL_q\phi_i + (L_fL_i-\mathfrak{L}^2)(\phi_q+\Phi_e)). \tag{3.28}$$

In each term, the extra inductance out front divides through cleanly. Plugging these back in and taking the infinite bias inductor limit, I get the reduced equations of motion

$$I_b - \frac{(L_q+L_i)\phi_f - L_q\phi_i - L_i(\phi_q+\Phi_e)}{\mathfrak{L}^2} = C_f\ddot{\phi}_f; \tag{3.29}$$

$$-\frac{-L_q\phi_f + (L_f+L_q)\phi_i - L_f(\phi_q+\Phi_e)}{\mathfrak{L}^2} - I_{0i}\sin\frac{2\pi\phi_i}{\Phi_0} = C_i\ddot{\phi}_i; \tag{3.30}$$

$$-\frac{-L_i\phi_f - L_q\phi_i + (L_f+L_i)(\phi_q+\Phi_e)}{\mathfrak{L}^2} - I_{0q}\sin\frac{2\pi\phi_q}{\Phi_0} = C_q\ddot{\phi}_q. \tag{3.31}$$

The filter should have no effect on dc signals. Since the current and flux-phase constraints for the dc SQUID (described in Section 2.3.3) are based on dc behavior, I expect the same constraints to hold in the filter plus SQUID model. This provides

a powerful tool for checking the consistency of the two Hamiltonians. The current constraint requires that the current through both junctions sum to the bias current. At dc, the right hand side of the equations of motion vanish. Summing them shows the current constraint (2.41) is satisfied.

Showing that the flux-phase constraint (2.42) also holds is more involved. Again, in the dc limit,

$$\phi_f = \frac{\mathfrak{L}^2}{L}(I_b + \frac{L_q}{\mathfrak{L}^2}\phi_i + \frac{L_i}{\mathfrak{L}^2}(\phi_q + \Phi_e), \tag{3.32}$$

where $L \equiv L_i + L_q$. Plugging into the third reduced equation of motion,

$$I_{0q}\sin\frac{2\pi\phi_q}{\Phi_0} + \frac{L_i + L_f}{\mathfrak{L}}(\phi_q + \Phi_e) - \frac{L_i}{L}(I_b + \frac{L_q}{\mathfrak{L}^2}\phi_i + \frac{L_i}{\mathfrak{L}}(\phi_q + \Phi_e)) - \frac{L_f}{\mathfrak{L}^2}\phi_i = 0; \tag{3.33}$$

$$\frac{L_i + L_f - \frac{L_i^2}{L}}{\mathfrak{L}^2}(\phi_q + \Phi_e) - \frac{\frac{L_i L_q}{L} + L_f}{\mathfrak{L}^2}\phi_i = \frac{L_i}{L}I_b - I_{0q}\sin\frac{2\pi\phi_q}{\Phi_0}. \tag{3.34}$$

I can make some useful simplifications:

$$L_i + L_f - \frac{L_i^2}{L} = \frac{\mathfrak{L}^2}{L}; \tag{3.35}$$

$$\frac{L_i L_q}{L} + L_f = \frac{\mathfrak{L}^2}{L}. \tag{3.36}$$

Substituting these in, I get

$$\phi_i - \phi_q = \Phi_e + L_q I_{0q}\sin\frac{2\pi\phi_q}{\Phi_0} - L_i I_{0i}\sin\frac{2\pi\phi_i}{\Phi_0}. \tag{3.37}$$

As expected, we have arrived at the dc SQUID phase flux condition. This consistency check suggests the equations of motion for the filter plus SQUID model are correct.

As in the dc SQUID, I extract the coefficients of the inductive parts of the

Lagrangian by differentiating the equations of motion once more. Skipping ahead to the construction of the Lagrangian,

$$\begin{aligned}\mathcal{L} = & C_f\frac{\dot{\phi}_f^2}{2} + C_i\frac{\dot{\phi}_i^2}{2} + C_q\frac{\dot{\phi}_q^2}{2} \\ & + I_b\phi_f + \frac{I_{0i}\Phi_0}{2\pi}\cos\frac{2\pi\phi_i}{\Phi_0} + \frac{I_{0q}\Phi_0}{2\pi}\cos\frac{2\pi\phi_q}{\Phi_0} \\ & - \frac{L_i+L_q}{2\mathfrak{L}^2}\phi_f^2 - \frac{L_f+L_q}{2\mathfrak{L}^2}\phi_i^2 - \frac{L_f+L_i}{2\mathfrak{L}^2}(\phi_q+\Phi_e)^2 \\ & + \frac{L_q}{\mathfrak{L}^2}\phi_f\phi_i + \frac{L_i}{\mathfrak{L}^2}\phi_f(\phi_q+\Phi_e) + \frac{L_f}{\mathfrak{L}^2}\phi_i(\phi_q+\Phi_e).\end{aligned} \tag{3.38}$$

The capacitive terms are kinetic energy; everything else is $-U$. As in Section 2.3.3, I transform to the Hamiltonian for the filtered SQUID. Switching from junction fluxes ϕ_n to junction phases $\gamma_n \equiv 2\pi\phi_n/\Phi_0$[1], I find the full Hamiltonian

$$\begin{aligned}\mathcal{H} = & \frac{p_f^2}{2m_f} + \frac{p_i^2}{2m_i} + \frac{p_q^2}{2m_q} - \frac{I_b\Phi_0}{2\pi}\gamma_f - \frac{I_{0i}\Phi_0}{2\pi}\cos\gamma_i - \frac{I_{0q}\Phi_0}{2\pi}\cos\gamma_q \\ & + \left(\frac{\Phi_0}{2\pi}\right)^2\left[\frac{L_i+L_q}{2\mathfrak{L}^2}\gamma_f^2 + \frac{L_f+L_q}{2\mathfrak{L}^2}\gamma_i^2 + \frac{L_f+L_i}{2\mathfrak{L}^2}\left(\gamma_q+\frac{2\pi\Phi_e}{\Phi_0}\right)^2\right] \\ & - \left(\frac{\Phi_0}{2\pi}\right)^2\left[\frac{L_q}{\mathfrak{L}^2}\gamma_f\gamma_i - \frac{L_i}{\mathfrak{L}^2}\gamma_f\left(\gamma_q+\frac{2\pi\Phi_e}{\Phi_0}\right) - \frac{L_f}{\mathfrak{L}^2}\gamma_i\left(\gamma_q+\frac{2\pi\Phi_e}{\Phi_0}\right)\right],\end{aligned} \tag{3.39}$$

3.3.3 Reduction to Jaynes-Cummings model

The Hamiltonian (3.39) is unwieldy. Recasting it as an approximate Jaynes-Cummings Hamiltonian (3.1) makes it easier to examine how the filter affects the qubit quantum mechanically. Of particular interest is the effective coupling strength g in such a model.

There are several steps in the reduction:

- The Hamiltonian (3.39) involves three coordinates, but the Jaynes-Cummings Hamiltonian couples two one-dimensional systems. Eliminating one of the coordinates is the first step. Following the analysis of the dc SQUID in Sec-

[1] Although ϕ_f does not correspond to a junction, it is convenient to make this switch anyway.

tion 2.5, I find the normal modes of the SQUID degrees of freedom.

- For the parameters of the BBC2 qubit, the SQUID modes are very widely separated (9 and 100 GHz). The 100 GHz mode is frozen out and unlikely to be driven by any of the signals I used. To reflect this, I discard the high frequency mode from the Hamiltonian.

- Following the previous step, the Hamiltonian is now in the form of two coupled harmonic oscillators — the filter and the qubit mode of the SQUID. To compare to the Jaynes-Cummings model, I assume that the harmonic approximation is close to the qubit frequency and treat the qubit mode as a two-level system. As in Section 2.5.2 cubic terms in the SQUID coordinates can be kept for perturbative corrections to the harmonic qubit approximation.

- The coupling terms are of the form $\gamma_i\gamma_f$ and $\gamma_q\gamma_f$. After replacing these coordinates with normal coordinates, I discard the high frequency mode, leaving a coupling term $\gamma_Q\gamma_f$. I then replace the γ operators with their ladder operator equivalents and make the rotating wave approximation to discard terms that create or destroy excitations in both subsystems. The coefficient of the remaining terms is $\hbar g$.

The normal mode calculation proceeds as in Section 2.5.2. I begin by assuming an equilibrium point $(\gamma_f^*, \gamma_i^*, \gamma_q^*)$. I then make small expansions around that point. Let $\delta\gamma_n \equiv \gamma_n - \gamma_n^*$. Since $\dot{\delta\gamma}_n = \dot{\gamma}_n$, the momentum terms remain unchanged. Leaving junctions unexpanded for now to simplify the bookkeeping, and discarding

constant terms, I rewrite the Hamiltonian in terms of $\delta\gamma$s as:

$$\begin{aligned}
\mathcal{H} = & \frac{p_f^2}{2m_f} + \frac{p_i^2}{2m_i} + \frac{p_q^2}{2m_q} \\
& - \frac{I_b\Phi_0}{2\pi}\delta\gamma_f - \frac{I_{0i}\Phi_0}{2\pi}\cos\gamma_i - \frac{I_{0q}\Phi_0}{2\pi}\cos\gamma_q \\
& + \left(\frac{\Phi_0}{2\pi}\right)^2\left[\frac{L_i+L_q}{2\mathfrak{L}^2}\delta\gamma_f^2 + \frac{L_f+L_q}{2\mathfrak{L}^2}\delta\gamma_i^2 + \frac{L_f+L_i}{2\mathfrak{L}^2}\delta\gamma_q^2\right] \\
& + \left(\frac{\Phi_0}{2\pi}\right)^2\left[\frac{L_i+L_q}{2\mathfrak{L}^2}2\delta\gamma_f\gamma_f^* + \frac{L_f+L_q}{2\mathfrak{L}^2}2\delta\gamma_i\gamma_i^* + \frac{L_f+L_i}{2\mathfrak{L}^2}2\delta\gamma_q\left(\gamma_q^* + \frac{2\pi\Phi_e}{\Phi_0}\right)\right] \\
& - \left(\frac{\Phi_0}{2\pi}\right)^2\left[\frac{L_q}{\mathfrak{L}^2}\delta\gamma_f\delta\gamma_i + \frac{L_i}{\mathfrak{L}^2}\delta\gamma_f\delta\gamma_q + \frac{L_f}{\mathfrak{L}^2}\delta\gamma_i\delta\gamma_q\right] \\
& - \left(\frac{\Phi_0}{2\pi}\right)^2\left[\frac{L_q}{\mathfrak{L}^2}(\gamma_i^*\delta\gamma_f + \gamma_f^*\delta\gamma_i) - \frac{L_i}{\mathfrak{L}^2}\left(\left(\gamma_q^* + \frac{2\pi\Phi_e}{\Phi_0}\right)\delta\gamma_f + \gamma_f^*\delta\gamma_q\right)\right. \\
& \left. - \frac{L_f}{\mathfrak{L}^2}\left(\left(\gamma_q^* + \frac{2\pi\Phi_e}{\Phi_0}\right)\delta\gamma_i + \gamma_i^*\delta\gamma_q\right)\right].
\end{aligned} \tag{3.40}$$

In preparation for expanding around the minimum of the SQUID, for the moment I neglect all terms with any dependence on $\delta\gamma_f$. The terms linear in $\delta\gamma$ set the equilibrium point, but if I assume a particular equilibrium $(\gamma_i^*, \gamma_q^*, \gamma_f^*)$, then they may be discarded. The resulting Hamiltonian is:

$$\begin{aligned}
\mathcal{H}_{\text{SQUID}} = & \frac{p_i^2}{2m_i} + \frac{p_q^2}{2m_q} \\
& - \frac{I_{0i}\Phi_0}{2\pi}\cos(\gamma_i^* + \delta\gamma_i) - \frac{I_{0q}\Phi_0}{2\pi}\cos(\gamma_q^* + \delta\gamma_q) \\
& + \left(\frac{\Phi_0}{2\pi}\right)^2\left(\frac{L_f+L_q}{2\mathfrak{L}^2}\delta\gamma_i^2 + \frac{L_f+L_i}{2\mathfrak{L}^2}\delta\gamma_q^2 - \frac{L_f}{\mathfrak{L}^2}\delta\gamma_i\delta\gamma_q\right),
\end{aligned} \tag{3.41}$$

I now make an expansion of the cosines, again discarding the linear terms:

$$\cos(\gamma_i^* + \delta\gamma_i) = (1 - \frac{\delta\gamma_i^2}{2})\cos\gamma_i^* - (\delta\gamma_i - \frac{\delta\gamma_i^3}{6})\sin\gamma_i^*. \tag{3.42}$$

Keeping only the harmonic term from these expansions for now yields a "simple

harmonic SQUID" model:

$$\begin{aligned}\mathcal{H}_{\mathrm{SHS}} = &\frac{p_i^2}{2m_i} + \frac{p_q^2}{2m_q} \\ &+ \left(\frac{\Phi_0}{2\pi}\right)^2 \left(\frac{\delta\gamma_i^2}{2L_{ji}} + \frac{\delta\gamma_q^2}{2L_{jq}}\right) \\ &+ \left(\frac{\Phi_0}{2\pi}\right)^2 \left(\frac{L_f + L_q}{2\mathfrak{L}^2}\delta\gamma_i^2 + \frac{L_f + L_i}{2\mathfrak{L}^2}\delta\gamma_q^2 - \frac{L_f}{\mathfrak{L}^2}\delta\gamma_i\delta\gamma_q\right).\end{aligned} \tag{3.43}$$

From this Hamiltonian I build the kinetic energy matrix

$$\mathbf{T} = \left(\frac{\Phi_0}{2\pi}\right)^2 \begin{pmatrix} C_i & 0 \\ 0 & C_q \end{pmatrix}, \tag{3.44}$$

and the potential energy matrix

$$\mathbf{V} = \left(\frac{\Phi_0}{2\pi}\right)^2 \begin{pmatrix} \frac{L_f+L_q}{\mathfrak{L}^2} + \frac{1}{L_{ji}} & -\frac{L_f}{\mathfrak{L}^2} \\ -\frac{L_f}{\mathfrak{L}^2} & \frac{L_f+L_i}{\mathfrak{L}^2} + \frac{1}{L_{jq}} \end{pmatrix}. \tag{3.45}$$

Although the current constraint and flux-phase condition (3.37) were the same as for an unfiltered dc SQUID, the filter is seen to affect the frequencies of the SQUID normal modes. This is because the normal modes also consider the possibility of ac signals, where the filter capacitance shunting the SQUID can come into play. This allows L_f to combine with the SQUID inductances.

In practical calculations, I typically solve for the modes and basis transforms numerically. Following that, the rest of the calculation of g is fairly straightforward. After the normal mode transformation, the new approximate Hamiltonian near the

equilibrium is

$$\begin{aligned}\mathcal{H} = &\frac{p_f^2}{2m_f} + \left(\frac{\Phi_0}{2\pi}\right)^2 \frac{L_i + L_q}{2\mathfrak{L}^2}\delta\gamma_f^2 \\ &+ p_I^2 + \omega_I^2\delta\gamma_I^2 + p_Q^2 + \omega_Q^2\delta\gamma_Q^2 \\ &- \left(\frac{\Phi_0}{2\pi}\right)^2 \left[\frac{L_q}{\mathfrak{L}^2}\delta\gamma_f\delta\gamma_i + \frac{L_i}{\mathfrak{L}^2}\delta\gamma_f\delta\gamma_q\right],\end{aligned} \tag{3.46}$$

where the mass of the properly scaled normal coordinates is 1. The capital letters denote normal coordinates, the lowercase ones the actual isolation and qubit phases. Replacing the junction coordinates with the proper linear combinations of mode coordinates, I get

$$\begin{aligned}\mathcal{H} = &\frac{p_f^2}{2m_f} + \left(\frac{\Phi_0}{2\pi}\right)^2 \frac{L_i + L_q}{2\mathfrak{L}^2}\delta\gamma_f^2 \\ &+ p_I^2 + \omega_I^2\delta\gamma_I^2 + p_Q^2 + \omega_Q^2\delta\gamma_Q^2 \\ &- \left(\frac{\Phi_0}{2\pi}\right)^2 \left[\frac{L_q}{\mathfrak{L}^2}\delta\gamma_f\left(A_{iI}\delta\gamma_I + A_{iQ}\delta\gamma_Q\right) + \frac{L_i}{\mathfrak{L}^2}\delta\gamma_f\left(A_{qI}\delta\gamma_I + A_{qQ}\delta\gamma_Q\right)\right].\end{aligned} \tag{3.47}$$

Stripping all terms involving the isolation mode, this leaves me with

$$\begin{aligned}\mathcal{H} = &\frac{p_f^2}{2m_f} + \left(\frac{\Phi_0}{2\pi}\right)^2 \frac{L_i + L_q}{2\mathfrak{L}^2}\delta\gamma_f^2 \\ &+ p_I^2 + \omega_I^2\delta\gamma_I^2 + p_Q^2 + \omega_Q^2\delta\gamma_Q^2 \\ &- \left(\frac{\Phi_0}{2\pi}\right)^2 \left[\frac{L_q}{\mathfrak{L}^2}\delta\gamma_f\left(A_{iI}\delta\gamma_I + A_{iQ}\delta\gamma_Q\right) + \frac{L_i}{\mathfrak{L}^2}\delta\gamma_f\left(A_{qI}\delta\gamma_I + A_{qQ}\delta\gamma_Q\right)\right].\end{aligned} \tag{3.48}$$

Finally, I can replace $\delta\gamma_n$ with appropriate creation and annihilation operators:

$$\delta\gamma_f = \sqrt{\frac{\hbar}{2m_f\omega_f}}(\hat{a}_r^\dagger + \hat{a}_r); \tag{3.49}$$

$$\delta\gamma_Q = \sqrt{\frac{\hbar}{2\omega_Q}}(\hat{a}_Q^\dagger + \hat{a}_Q), \tag{3.50}$$

where

$$\omega_f = \sqrt{\frac{1}{L_f^* C_f}}; \tag{3.51}$$

$$L_f^* = \frac{\mathfrak{L}^2}{L_i + L_q}. \tag{3.52}$$

In the limit where $L_f \gg L_i, L_q$, L_f^* reduces to L_f. $\delta\gamma_Q$ has already been scaled to have a mass of 1, so there is no mass term in the scaling factor for it. When I make this substitution, the first two lines of (3.48) collapse into the usual number operator representation of harmonic oscillators. I also make the rotating wave approximation, discarding the coupling terms $\hat{a}_r^\dagger \hat{a}_Q^\dagger$ and $\hat{a}_r \hat{a}_Q$. These terms pick up frequency contributions at the sum $\omega_Q + \omega_f$, whereas the other terms pick up frequencies at $|\omega_Q - \omega_f|$. Following this approximation, I have arrived at the Jaynes-Cummings form:

$$\mathcal{H} = \hbar\omega_r \hat{a}_r^\dagger \hat{a}_r + \hbar\omega_Q \hat{a}_Q^\dagger \hat{a}_Q + \hbar g \left(\hat{a}_r^\dagger \hat{a}_Q + \hat{a}_Q^\dagger \hat{a}_r \right), \tag{3.53}$$

where

$$g = -\sqrt{\frac{1}{2m_f \omega_f \omega_Q}} \left(\frac{\Phi_0}{2\pi} \right)^2 \left[\frac{L_q}{\mathfrak{L}^2} A_{iQ} + \frac{L_i}{\mathfrak{L}^2} A_{qQ} \right]. \tag{3.54}$$

There are some slight differences between (3.1) and (3.53). The coupling in the SQUID Jaynes-Cummings model is bilinear in γ, whereas it is a momentum coupling in the original Jaynes-Cummings Hamiltonian. This is unimportant, as the absolute values of the off-diagonal matrix elements are the same, leaving the eigenvalues unchanged. More significant is that I am using harmonic oscillator states to describe the qubit rather than treating it as a true two-level system. When the qubit has high anharmonicity, treating the two lowest lying states of the qubit mode as a true two-level system is a good approximation and the mapping to the Jaynes-Cummings model is sensible.

Instead of only considering the normal modes of the SQUID part of (3.39), I

could instead find the normal modes of the entire Hamiltonian to approximate the qubit dynamics. This sort of approach is sometimes used in circuit QED to avoid issues with multi-mode Jaynes-Cummings models [42]. The advantage of the Jaynes-Cummings approach is that it is a well-understood Hamiltonian with predictable consequences parameterized by g. Further, for weak coupling the separation of the filter and the qubit is a clearer, more physically transparent description of the system than a single harmonic mode.

3.4 Summary

In this chapter, I reviewed the physics of the Jaynes-Cummings model and discussed its relevance for superconducting quantum computing through the circuit QED approach. I then discussed modifying the dc SQUID phase qubit of Section 2.5 by including an LC filter on the bias line. Using the techniques of Chapter 2, I found the Hamiltonian for the lumped element circuit model of the filter plus SQUID. I then reduced this Hamiltonian to an effective Jaynes-Cummings model, where the SQUID is the qubit and the filter is the resonator.

Chapter 4

Fabrication details for filtered dc SQUID qubit

One of the most important topics in superconducting quantum information experiments is the effect of materials on coherence times. A series of influential experiments suggested that dielectric loss from discrete two-level systems is an important contributor to dissipation effects [75, 76]. This contribution is especially important in phase qubits with their larger Josephson junctions. This suggests a natural road to T_1 improvement. By working with as little dielectric as possible — reducing junction size, limiting wiring crossover — and using low-loss substrates like sapphire, we expect substantial gains in qubit lifetime.

Our lab's earliest work in phase qubits used devices fabricated by the superconducting foundry Hypres. They use a multilayer niobium process with a great deal of silicon oxide insulation between layers. Further, the smallest junction they can make is roughly 2 μm by 2 μm. The Hypres qubits had very short lifetimes. Although separating the dielectric loss from contributions due to the design is challenging, it was clear that we could get much cleaner qubits by building them ourselves.

For several years, our lab has fabricated qubits using aluminum on sapphire wafers. To reduce the presence of extraneous dielectric, as much of the qubit wiring as possible is contained in a single layer of aluminum. To produce the junctions themselves, we use thermal double angle evaporation. In this chapter, I describe the double angle evaporation technique, since its requirements were a major element of the design of BBC2. I then describe the fabrication of BBC2 which I performed jointly with Rangga Budoyo.

4.1 Double angle evaporation

Our lithography processes use the patterned resist as a negative. We remove the exposed areas, leaving a stencil where aluminum is deposited. For the majority of the wiring this is all that is needed. Josephson junctions, though, require two layers of aluminum separated by a thin oxide layer. A second deposition directly above the first does not help; this would result in a junction shaped like the entire pattern.

The solution is to pattern the resist with a suspended bridge near where the junction is desired. The resist bridge casts a shadow on the substrate and yields a break in the deposited aluminum. The location of the shadow depends on the angle of deposition of the aluminum. By rotating the substrate, this angle can be changed between evaporations. Two evaporations are made at well-chosen angles. Far away from the resist bridge, both layers are present. Near the resist bridge, the two shadows of the bridge define regions where only one layer is present. In between the shadows is a small region where both layers are present. This layer is the desired Josephson junction. Images of this process for specific designs are shown in Figures 4.1 and 4.2.

We use two layers of resist to achieve the suspended bridge. The bottom resist develops at a different rate than the layer above it. With appropriate processing we can use this to undercut the top layer. By developing long enough to dissolve the bottom layer completely underneath a resist bridge in the top layer, we leave an exposed region on the substrate where junctions will be formed.

4.1.1 Some technical details

Double angle evaporation is a general framework that allows for different specific ways of fabricating junctions. Translating the size of elements in the litho-

graphic pattern to the actual junction size requires knowledge of the pattern. I discuss the geometry of the two types of resist bridge used in the BBC2 design.

Perhaps the simplest junction is one that results from a bridge interrupting a straight line, as seen in Fig. 4.1. The bridge is suspended above the desired junction location on the substrate. Ideally the rotation axis would run through the bridge, perpendicular to the patterned line. The two angles then project the shadow of the bridge in opposite directions along the line. For certain choices of angle, there will be a junction defined between the two shadows. The width of the junction is the width of the patterned line. The length of the junction is given by

$$t_b \left(\tan\theta_1 + \tan\theta_2\right) - w,$$

where t_b is the thickness of the undercut layer, θ_i is the angle of the ith deposition measured from the normal to the substrate, and w is the width of the resist bridge.

This bridge works well for the larger junction of BBC2. It is less reliable for the small junction. Long stretches of thin line are delicate and prone to breakage. Additionally, as (4.1.1) shows, the size of junctions formed this way are sensitive to several parameters. In making junctions with dimensions less than the ≈ 1 μm thickness of the undercut layer, small inaccuracies in θ_1 and θ_2 can lead to significant relative changes in junction size or even a failure to form the junction.

For greater consistency in making BBC2's small junction, we employed a different pattern featuring one straight finger and one L-shaped one, with contact being made between the straight finger and the arm of the L perpendicular to that finger. This L bridge approach is illustrated in Fig. 4.2. If θ_1 is steep enough, the straight finger is projected underneath where the L-shaped finger is patterned. The second evaporation is performed at an angle θ_2 near vertical[1].

[1]In a sample using only L bridge junctions,vertical deposition would be fine. If the sample uses both types of bridge, as BBC2 did, then deviating from vertical helps with the fabrication of the other style of junction.

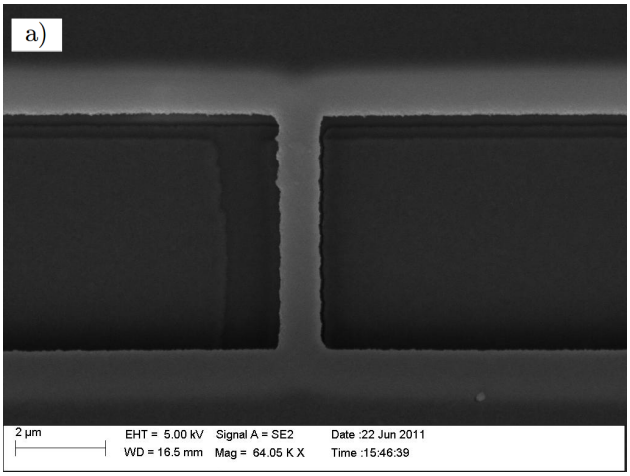

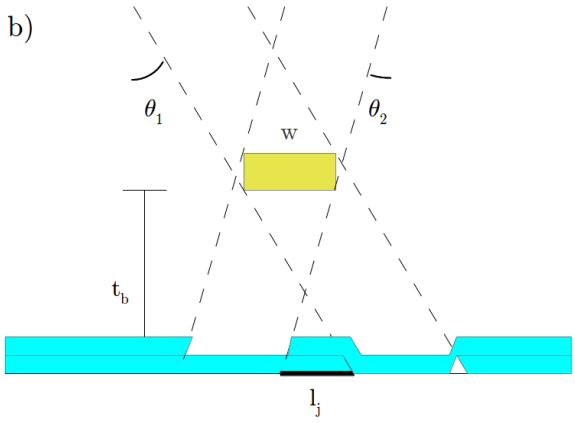

Figure 4.1: a) SEM image of a straight line bridge taken following e-beam development and aluminum deposition but before resist lift-off. b) Side view of the double angle process. The yellow layer is the top layer resist, separated from the wafer by the undercut resist thickness t_b. The cyan layers are the deposited aluminum. The bold line indicates the location of the Josephson junction.

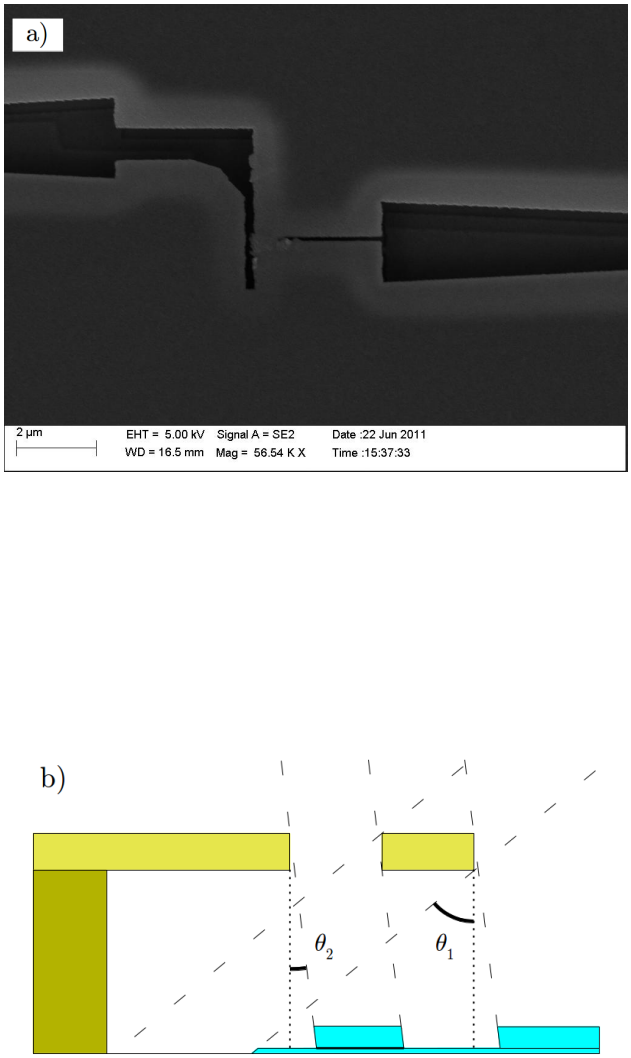

Figure 4.2: a) SEM image of an L bridge taken following e-beam development and aluminum deposition but before resist lift-off. b) Side view of double angle on an L bridge. Unlike in Fig. 4.1, resist is present in both layers on one side of the bridge. This is a scale drawing with the correct angles of the process used for the BBC2 device.

The advantage of this approach is that provided full overlap occurs, the dependence of the junction size on the angles and resist parameters is much weaker. The junction size is determined instead by the width of the lithographically patterned fingers. One potential complication of the L bridge approach is the resist sidewall on the L finger side. Since the L finger does not have a line continuing in the same direction as the evaporation, the bottom layer of resist remains there. A sufficiently steep angle can end up plating out on the sidewall of this resist and be removed during the development process. This can lead to ragged edges on the leads around the junction, or possibly the tearing of the junction itself during liftoff.

4.1.2 Evaporator setup

For junction depositions, we used the cryo-pumped thermal evaporator in the thin-film deposition laboratory in CNAM. The evaporator has four water-cooled electrodes in a cross formation at the bottom. For double angle evaporations, we typically chose two opposing electrodes. Tungsten wire coil boats were then attached to these electrodes and filled with aluminum shot.

We experienced occasional problems with these boats breaking during aluminum evaporation. To minimize boat burnout, we found that gradual heating of the element at the beginning worked well. The potentiometer on the evaporator was set to 2.3 (the threshold for turning on at all), at which point the current through the boat would typically be 10 to 20 amperes. After one minute at this setting, the pot was adjusted to 2.4 and held for another minute, then to 2.5 for 30 seconds. Following this, the pot was adjusted until we achieved an acceptable deposition rate, typically 9-10 Å per second.

The evaporator is a shared facility, used by many people for several different processes. Specific capabilities are associated with particular lids for the evaporator. For double angle evaporation, we used a lid that has an ion mill attachment and a

rotatable sample stage. For whole wafer processing (e.g. the initial photolithographic steps), we directly mounted the wafer to the stage using spring clips or vacuum grease. The electron beam lithography for BBC2 was done after the wafer was diced and reduced to individual chips ≈ 1 mm on a side. To do evaporations with these chips, we attached two or three at a time with spring clips to a chip carrier screwed to the rotating stage.

We oriented the evaporator lid so that the stage was almost directly above the electrodes[2]. Both bridge designs discussed assume that the rotation axis is oriented along the bridge. To make sure we accomplished this we used a consistent alignment of the BBC2 chips relative to the sample carrier. The stage rotation was controlled by a mechanical dial on the evaporator lid. After mounting the sample, the stage was rotated to its first angle and the position of the dial noted. We then rotated to the second angle and marked the dial position there as well. To reduce angle inaccuracies from backlash in the dial, the dial measurement was always performed at the end of a rotation before releasing the dial.

Between the two evaporation steps, we oxidized the bottom aluminum layer to create the junction barrier. A valve on the ion mill controller adjusts the rate at which oxygen enters the system, and a valve on the lid itself allows oxygen into the vacuum chamber. Both pressure and time can be controlled to determine total exposure $E = P \times t$. Empirically, the critical current scales like $E^{-0.4}$ [109]. Generalizing to a time dependent pressure, we integrated over the periods where oxygen is being admitted to or pumped out of the vacuum chamber. Since the pressure varied roughly linearly during these periods of changing pressure, a simple approximation is to add half the time ramping up or pumping down.

[2]Displacements from vertical lead to corrections to the apparent angle of the stage.

4.2 Fabrication of BBC2

Experimentally my thesis focused on the BBC2 qubit. I discuss the design of BBC2 in Chapter 6; here, I discuss the fabrication process in detail. BBC2 was built jointly by Rangga Budoyo and myself. Although there was some flexibility in the division of labor, typically I performed the preparation of the e-beam resist and the e-beam lithography while Rangga performed most steps involving evaporation. Vitaley Zaretskey provided extensive help with the operation of the JEOL SEM at LPS as well as confocal imaging of earlier versions of the sample to diagnose problems in the fabrication process.

The BBC2 device has on-chip LC filtering requiring a large capacitor using a silicon nitride dielectric. The qubit sits in a cutout in the ground plane of a coplanar waveguide. The ground plane has many small square holes to reduce vortex motion. BBC2 also features an interdigitated capacitor shunting one of the junctions. The capacitor fills much of a region several hundred microns on a side and has almost 100 fingers — it represented the bulk of the time spent during electron beam writing.

Our process for the fabrication of BBC2 was a modified version of a process designed by Tony Przybysz for his PB9 sample [28]. The main difference is in the choice of resists. Tony used MMA as his undercut layer; I chose LOR 10B. The appeal of LOR is a more predictable undercut dynamic; the rate at which MMA is removed by the developer is dosage dependent, whereas the rate for LOR is determined by the time and temperature of the bake following spincasting. This predictability simplified the process of determining useful e-beam dosages for the lithography. Except where specified, the fabrication process was done using tools in the UMD Fab Lab.

The process is as follows:

- We started with a 3 inch sapphire wafer purchased from Kyocera. The wafers

were 430 μm thick and polished on one side in a c-plane orientation. The wafer was given an initial cleaning in a series of solvent baths — first acetone, then methanol, then isopropanol (IPA).

- Rangga performed an initial aluminum deposition. This aluminum layer provided the bottom plate of the filter capacitor C_f. No vias were provided to this layer. Instead, C_f was designed as two capacitors in series with a shared bottom plate. This allowed all electrical contacts to be made on the top level of deposited aluminum. The typical thickness of the bottom layer was 50-60 nm. I do not have the thickness used for BBC2, but I did not expect any dependence of the qubit performance on this thickness.

- We performed the first photolithographic step, patterning the capacitor backplate using the clean room's mask aligner. All of our photolithography steps use Shipley 1813 as the photoresist. S-1813 is a positive resist, meaning that development removes exposed regions. To define the capacitor backplates, we used a photomask with clear field and dark pattern, leaving the desired backplates as the only unexposed region. Most of the features we define through photolithography are relatively gross, so extensive optimizing of the photo steps was not necessary to get good results.

- We developed using a metal ion free developer such as MF-319 or MF-CD-26 (for BBC2 we used CD-26). After this step, all of the aluminum on the wafer is exposed except at the capacitor backplates.

- We used a wet aluminum etchant to remove the uncovered aluminum on the wafer. The MF developers can also etch aluminum, although considerably more slowly. Following the etch the remaining resist was removed by soaking in heated Remover PG. Any residue that remained following the Remover PG

bath was removed with an oxygen plasma clean in the reactive ion etch (RIE) tool.

- For the capacitor dielectric, we deposited a layer of silicon nitride. This step was performed by our collaborators at LPS using a PE-CVD process developed to optimize loss tangent [31]. Nitride is a popular choice for dielectric material in the community because it is less lossy than aluminum or silicon oxides and is not too difficult to grow. We deposited 120 nm of silicon nitride. Going much thinner than this poses risks of pinholes and shorts; there is also a danger of the film breaking on the aluminum backplates already deposited if it is too thin.

- Rangga performed a second photolithography step to pattern the capacitor dielectric. The mask for this step is very similar to the mask for the first step but with the pattern extended by 15 μm in each direction. This pattern leaves enough nitride to cover the capacitor backplates, with a safety margin on each side to reduce shorts caused by misalignment of the top wiring layer.

- The remaining nitride was removed using RIE. The recipe is listed on the RIE in the Fab Lab and uses SF_6. 45 seconds was typically enough to remove the nitride.

- Rangga performed a third photolithography step. In this layer are the ground plane, flux lines for qubit bias, the microwave transmission line, the filter inductors L_f, and the top layer of C_f. Unlike the previous lithography steps, here we performed the lithography first and then deposited material afterward. Thus, for this step we used a photomask with clear pattern and dark field.

- Following development, Rangga performed another aluminum deposition of 50-60 nm. Both aluminum depositions are single angle, since the wiring of

these layers contains no junctions.

- Excess aluminum is lifted off in an acetone bath placed in a sonicator. The sample is oriented at an angle facing down. We take this step because the ground plane features many 10 μm by 10 μm holes where small squares of aluminum are removed to reduce vortex motion in the ground plane. If these squares redeposit on the wafer they pose a risk of shorts, especially in the region where the interdigitated capacitor is placed. Doing liftoff in the sonicator was very successful at preventing that.

- Rangga performed a fourth and final photolithography step. In this step gold alignment marks are deposited. We used gold because the contrast between aluminum features and the sapphire wafer is very poor when the electron microscope is configured for lithography. A large cross located above the microwave transmission line is the primary alignment mark. Four smaller markers are placed in the corners of a box around the location where the device is written.

- A thin (1-2 nm) chromium layer was deposited in the thermal evaporator to provide an adhesion layer for the gold. The gold deposition followed afterward. Using two different electrodes in the evaporator, both evaporations were performed without exposing the wafer to air. Following acetone liftoff of the chromium and gold, the wafer was again cleaned in Remover PG. At this point, the wafer was ready for the e-beam resist stack.

- Following a solvent cleaning, I spun a layer of LOR 10B on the wafer at 4000 rpm for 45 seconds. Since the LOR is quite viscous, I ran the spinner manually at about 50 rpm for 3-4 seconds while pouring LOR on the wafer. This dispersed the LOR evenly over the wafer. I then accelerated to 4000 rpm over the course of 3-4 more seconds. At this spin speed, the specified

thickness of LOR 10B is just under 1 μm; I did no measurements to confirm this thickness. The sample was baked in an oven at 180 °C for 10 minutes, aiming for an undercut rate of about 35 nm / sec. LOR produces little whiskers of resist on the edge of the wafer. Although these can be removed with edge bead remover, since the outer edges of the wafer were generally unpatterned or sacrificial I did not do this.

- I spun PMMA-950-A4 at 4000 rpm for 45 seconds on the wafer. Immediately following the spinning, I put the wafer on a hot plate to bake at 150 °C for 3.5 hours. The specified thickness of the resist layer under these conditions is about 200 nm. While baking, I covered the wafer with a Pyrex vessel to limit accretion of particulates during the lengthy bake.

- I thermally evaporated a 20 nm layer of aluminum on the resist stack. This aluminum anti-charging layer disperses the charge buildup on the sapphire wafer during the e-beam write. Initially I brought the wafer back to the thermal evaporator in CNAM for this step. However, I encountered difficulties with this process. Further in the process, acetone was used to remove a protective layer of photoresist applied before dicing. This later step resulted in craters in the anti-charging layer, shown in Fig. 4.3. The culprit was dust on the wafer — accumulated after leaving the cleanroom environment — causing breaks in the anti-charging layer. These breaks allowed acetone to penetrate and remove the e-beam resist underneath the anti-charging layer. When I switched to immediately transferring the wafer to the Kim building thermal evaporator following the PMMA bake, I no longer observed the craters.

- I applied a protective layer of S-1813 resist to prepare the wafer for dicing. The dicing saw uses a blade cooled by continuously spraying water on the blade and wafer. The protective resist layer shielded the resist stack from the

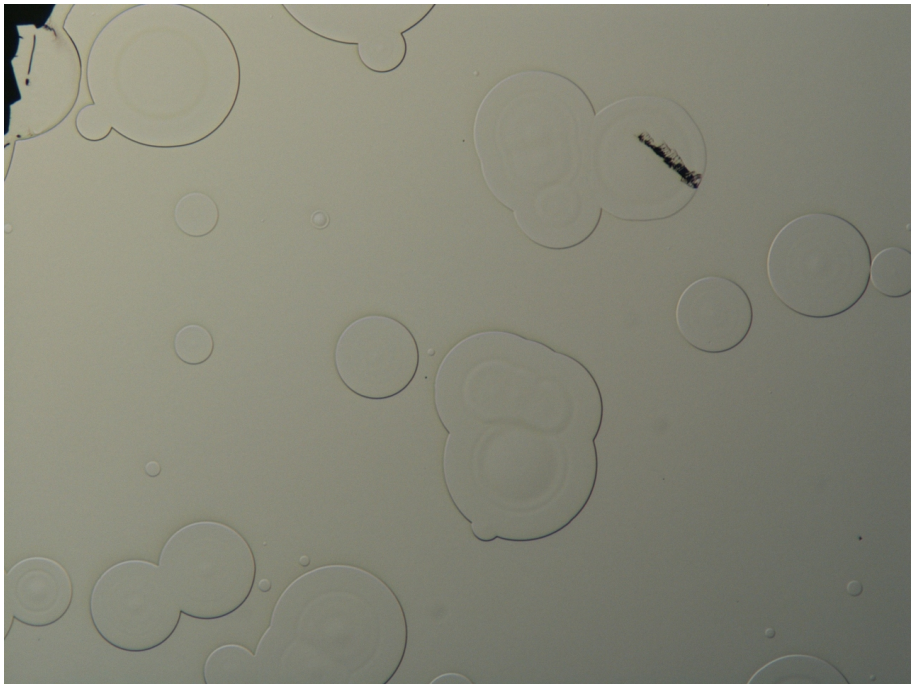

Figure 4.3: Photo of chip with cratered anti-charging layer.

water. After spinning the S-1813, I baked at 150 ° C for 20 minutes.

- I diced the wafer into individual chips for the e-beam process. The mask contained devices with differing amounts of on-chip LC filtering — half the chips had one stage, and half had two. These two types of devices had two different chip sizes: 650 μm by 750 μm, and 650 μm by 850 μm. Following dicing, I dried the wafer and stored it in the clean room.

- For e-beam writing, I removed individual chips from the wafer. The protective layer of S-1813 was removed with a 5 minute soak in acetone followed by an IPA rinse. I then took the chips to LPS, where I used a JEOL 6500 system equipped with the Nabity pattern generation system (NPGS) for doing e-beam writes. The write was done in four steps, with a pause following the first three to adjust the microscope magnification and beam current:

  – In the first step I wrote the small Josephson junction. This was done at a magnification of 900x with a typical beam current of 17 pA. I used an

area dose of 450 μC/cm^2. This step took 1 minute.

  - In the second step I wrote the large Josephson junction. This was done at a magnification of 90x, beam current of 17 pA, and area dose 450 μC/cm^2. This step also took 1 minute.
  - The third step wrote the inductor loop and interdigitated capacitor and was the longest step in the process. The magnification was 90x, the beam current $\approx$ 150 pA, and the area dose 600 μC/cm^2. This step took 15 minutes.
  - The final step wrote the contact pads and wires connecting them to the SQUID loop. The magnification was 45x, the beam current 700 pA and the area dose 450 μC/cm^2. This step took 10 minutes.

- I developed the e-beam resist. First, I placed the chip in an MF-CD-26 bath to remove the aluminum anti-charging layer. This took about 5 minutes; I pulled the chip when visual inspection showed the aluminum to be removed. After rinsing with water and drying, I placed the chip in a bath of MIBK/IPA 1:3, a PMMA developer recommended for high resolution. I gently agitated in the MIBK/IPA bath for 75 seconds and then rinsed in an IPA bath briefly before drying. To produce the undercut, I then put the chip in the MF-CD-26 bath with gentle agitation. The undercut rate depends on how the LOR 10B was baked when first applied. This step required careful timing. There must be sufficient undercut to clean the area under the junction bridges, but too much undercut causes collapse of device features. This collapse was especially common in the fingers of the interdigitated capacitor. With chips from the BBC2 wafer I would typically develop for 20 or 21 seconds before rinsing with water and then IPA. If visual inspection showed more development was necessary I would develop for 1-2 seconds more and then rinse and check again.

- Rangga performed double angle aluminum deposition in the CNAM thermal evaporator using the ion mill lid. The first deposition used electrode 1, with a background pressure of 1.6×10^{-6} torr. 55 nm of aluminum was deposited according to the crystal monitor[3] at $\theta_1 = 50°$. We oxidized the first layer for 10 minutes. The second deposition used electrode 3, with a background pressure of 7×10^{-7} torr. 120 nm of aluminum was deposited at an angle of $\theta_2 = 7°$. Using a thin first layer and making the low-angle deposition second resolved an issue with cracks forming near the junctions (see Fig. 4.4.)

- I lifted off the excess aluminum and remaining resist stack. I started with an acetone bath to remove most of the aluminum, and then used Remover PG to eliminate as much resist residue as possible. Following the Remover PG bath, I rinsed the chip in IPA and dried with nitrogen.

- I mounted the chip in an aluminum sample box described further in Section 5.1.1. I attached the chip to the box with a very small drop of GE varnish. To spread the varnish under the chip, I pressed the chip down using tweezers at the chip corners.

- After the varnish hardened and the chip was mechanically secure, I made aluminum wirebonds between the circuit board and the chip using a Kulicke-Soffa Model 4523 wedge bonder. To reduce the risk of shorting out the junction with a static discharge, I attached an SMA shorting cap to the bias line. The flux line and bias line were both grounded to the on-chip aluminum ground plane. For good grounding, I made $\approx$ 20 wirebonds around the perimeter of the chip connecting the on-chip ground plane to the ground plane of the printed circuit board.

[3]The angle of the monitor does not change with the sample. The effective depth of the deposited layer will be lower at high angles.

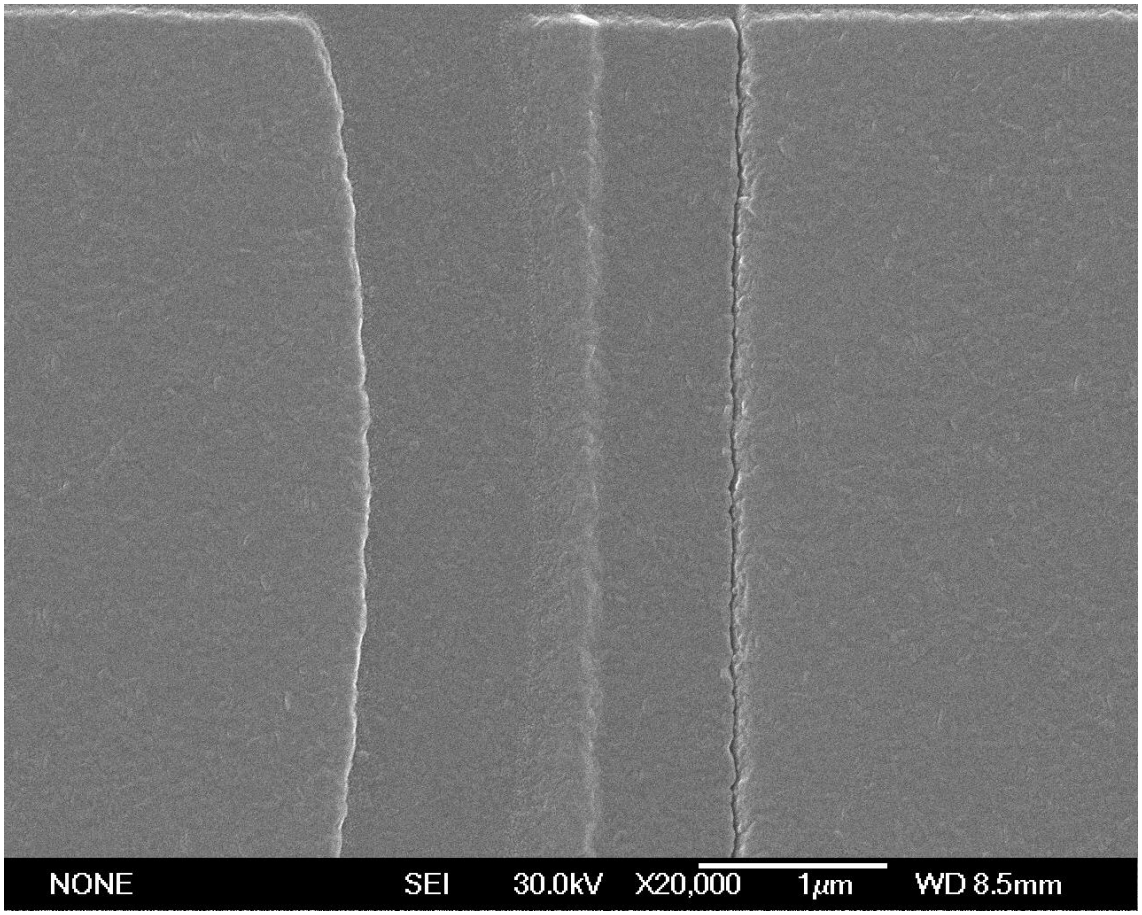

Figure 4.4: SEM image of the big junction from an early attempt at the BBC2 design. To the right of the junction a large crack is visible.

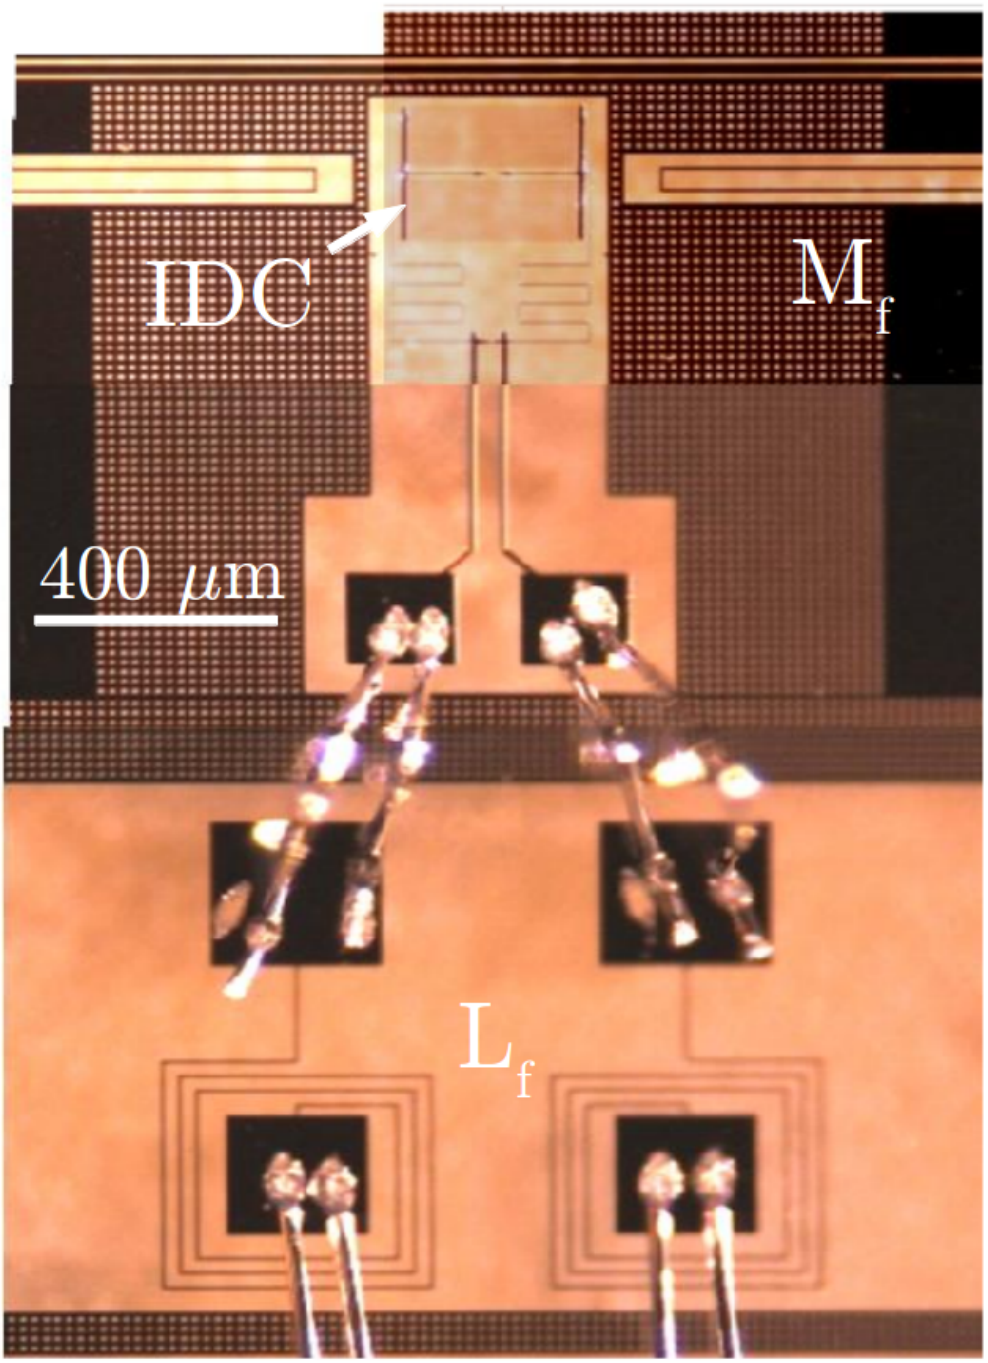

Figure 4.5: Photo of the BBC2 device following wirebonding. This is a composite of several photos. The wirebonds at the bottom of the image lead to the top plates of the filter capacitor and the contact pads where the system is grounded.

Fig. 4.5 shows an image of the wirebonded BBC2 device, stitched together from several different photographs. The visible bonds connect C_f to L_f and then to the qubit.

4.3 Summary

In this chapter, I discussed the fabrication of the BBC2 qubit. After describing double angle evaporation, I explained how it works for the geometry of the BBC2 junction patterns. I reviewed the PMMA/LOR process I used for the lithography.

Chapter 5

Experimental setup

Superconducting qubit experiments are conducted at millikelvin (mK) temperatures. There are several reasons for this. Most obviously, the reliance on superconducting circuits demands temperatures below T_c to achieve superconductivity. Operating at $T \ll T_c$ dramatically lowers the thermal quasiparticle background. This reduces a major source of dissipation, possibly yielding longer-lived qubits. Another important point is that most experiments rely on qubit initialization in the ground state through thermal relaxation (although some consideration has been given to supplementing this with active cooling in longer lived samples [85]).

The most common approach for achieving mK temperatures is the dilution refrigerator. Our studies on the BBC2 device were performed in an Oxford Instruments Model 200 wet fridge. This refrigerator was used for several generations of experiments; descriptions of it can be found in earlier theses from the group [25–29, 71, 110]. Maintenance and changing experimental needs demanded some modifications to the refrigerator. I begin this chapter by describing the refrigerator setup for the measurements I performed. I then describe the setup of the instrumentation, followed by a section briefly describing some measurement techniques.

5.1 Refrigerator setup

Most of the dilution refrigerator is just infrastructure supporting the cooling of a small mixing chamber (so called because the He-3/He-4 phase separation driving the cooling process occurs there.) The dilution unit is separated from a cryogen

bath (nitrogen during initial pre-cools, then helium) by a vacuum can, sealed with an indium O-ring. At the top of the dilution unit is a reservoir known as the "1K pot." The pot is connected to the helium bath through a siphon, with a flow regulated by two needle valves adjusted from above the refrigerator. The pot fills continuously with liquid helium and cools to ≈ 1.5 K through pumping. This temperature is sufficient to liquefy a refrigerant mixture ("the mix") of He-3/He-4 that is circulated in a closed cycle through the refrigerator. The volume of mix is set so that the liquid surface occurs in the still, another reservoir in the refrigerator. Since He-3 can be cooled through pumping to 300 mK [111], pumping at the still cools the mix through the phase separation temperature 700 mK. The mix separates into He-3 rich and He-3 dilute ($\approx 6\%$ concentration at base temperature) phases. Pumping at the still drives these phases out of equilibrium, and the movement of He-3 across the phase barrier to restore it has a latent heat associated with it that provides cooling down to millikelvin temperatures. A more in-depth explanation of the process is available in Lounasmaa [112].

In our refrigerator, the base temperature as normally operated is around 20 mK. This limit is set by conduction through the wiring to the mixing chamber, as well as the heat load of the continuously cycling mixture. This base temperature is only achieved at the mixing chamber; other stages of the fridge have their own characteristic temperatures. The qubit is mounted in an aluminum sample box thermally anchored to the mixing chamber. The mixing chamber is surrounded by a brass radiation shield anchored to the still flange. The Stefan-Boltzmann law states that radiated blackbody power is proportional to T^4. Since the still temperature 300 mK is an order of magnitude less than the vacuum can temperature 4.2 K, this represents four orders of magnitude less radiated thermal power.

Figure 5.1 is a photograph of the dilution unit. The long finger at the bottom is bolted to the mixing chamber. The finger is made of OFHC copper for good

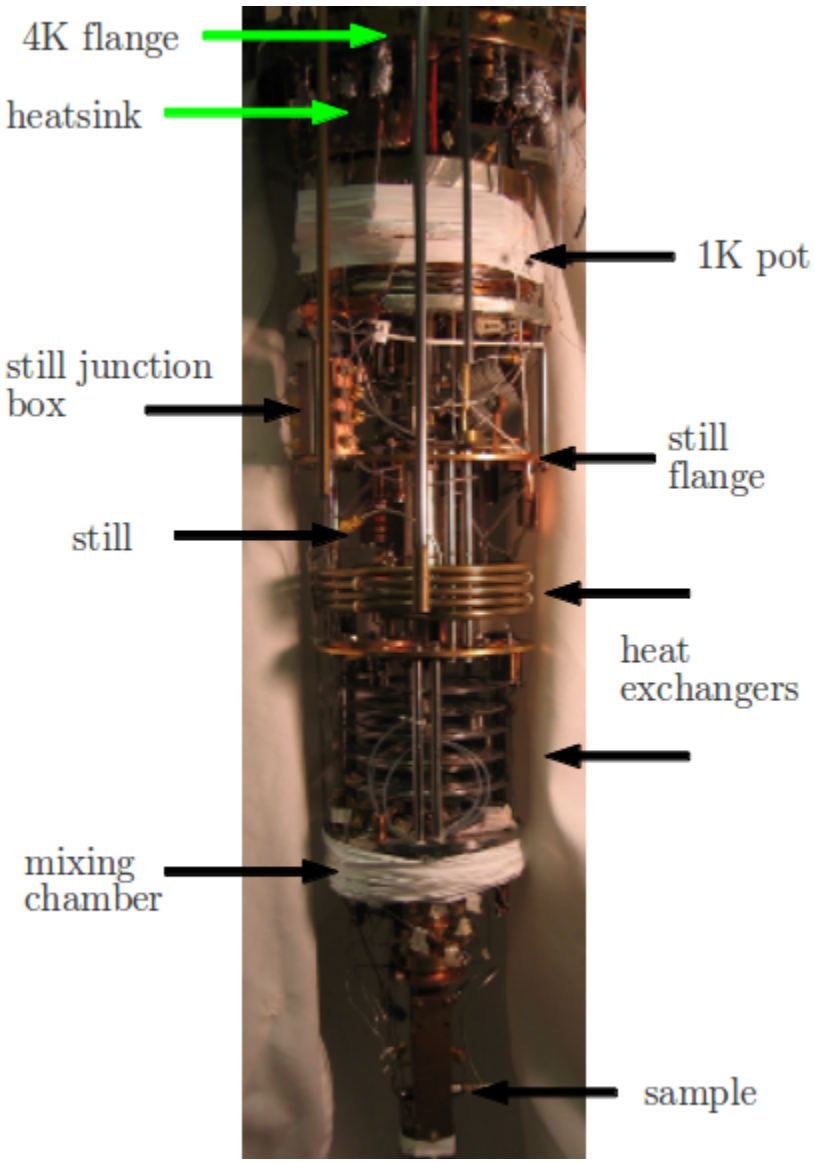

Figure 5.1: Photo of the dilution unit. The IVC is mounted to the 4K flange; the radiation shield is mounted to the still flange.

thermal conductivity. The finger holds the sample box as well as two sets of SMA connectorized filters, seen in the long boxes in Figure 5.1. These filters, described in Sudeep Dutta's thesis [25], consist of lumped element low-pass LC filters rolling off at 10 MHz followed by copper powder filters to eliminate high frequency signal bypassing the lumped element filters. By placing these at base temperature, right before the sample, the filters simultaneously remove most of the high temperature thermal radiation while also generating comparatively little of it themselves.

During a cooldown, the major maintenance associated with the refrigerator was periodic transfers of liquid helium. During the BBC2 measurements, I observed fluctuations in the flux offset of the SQUID following a transfer. This was observed as a change in the bias conditions necessary to reach a particular frequency, or a change in the location of histogram peaks if run with swept currents. The size of the shift was typically of order $\sim 0.1\Phi_0$. It is unclear what the source of this shift is. At points in the helium transfer the blowoff becomes rather strong, so mechanical disturbance of the refrigerator is possible. On the other hand, while there was some warming according to the mixing chamber thermometry, it was quite modest.

One possible explanation is vortex movement from residual field. I strive to minimize the background field as the sample and box go superconducting. A mu-metal shield is raised around the dewar after reaching 4 K but before condensing the mixture and moving to 1 K. The entire refrigerator is operated in a shielded room to provide extra isolation. Typically the shielded room would be depowered and closed during cooldowns. For the BBC2 experiments, an electrical heater ran during the first two cooldowns. The heater mitigated a helium leak in the L ports that showed up during transfers as the top of the fridge cooled[1]. It is possible this led to a larger background field during the superconducting transition and led to trapped vortices.

[1]The leak was caused by slightly mismatched O-rings installed to replace old ones that were no longer resilient.

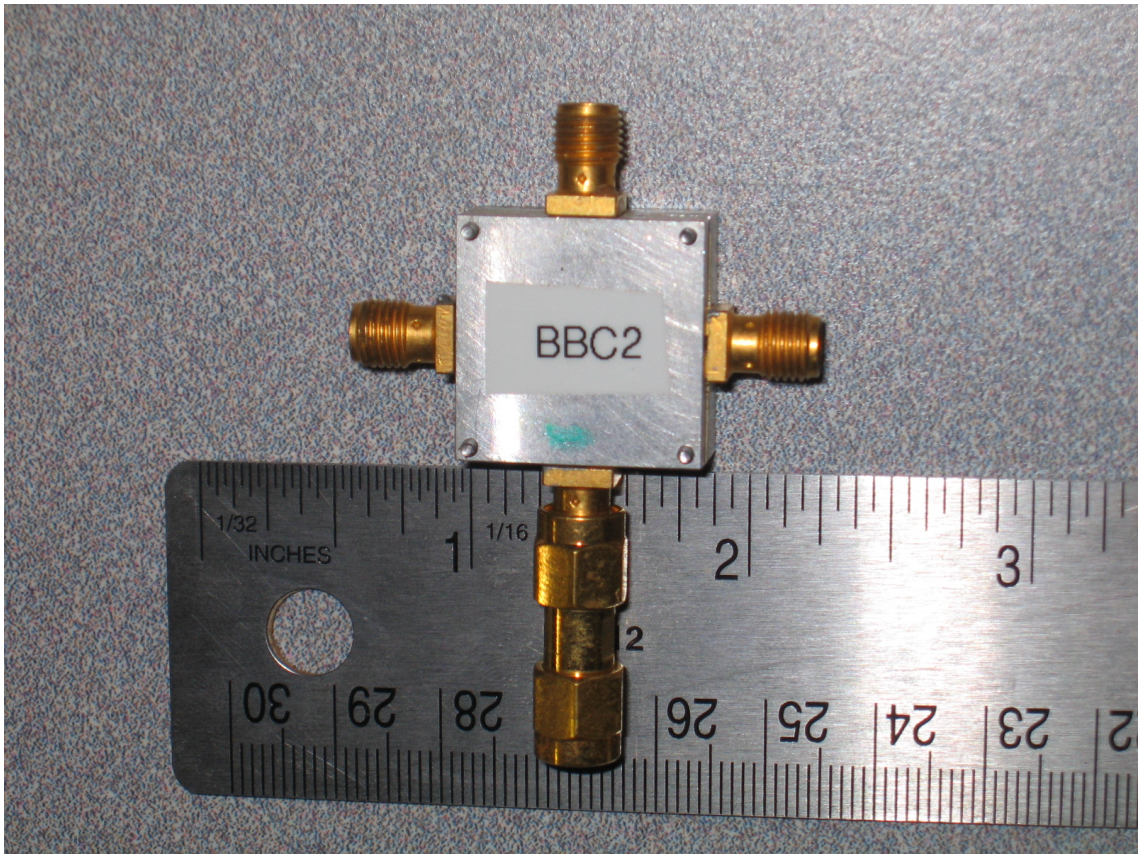

Figure 5.2: Photo of sample box used for the BBC2 measurements.

5.1.1 Sample box

For the BBC2 measurements, I used an aluminum sample box provided by Ben Palmer's group at LPS. The sample box is photographed in Figure 5.2. The box consisted of two square pieces: a lid, and the piece where the sample was mounted. The mounting piece has a custom printed circuit board (PCB) featuring four tapered coplanar waveguides (CPWs). On each edge of the box, an SMA connector was soldered to couple signals to the CPWs. A hole was milled out of the center of the CPB so the sample could be placed inside. In Section 4.2 I discuss the process of mounting and wirebonding the sample.

The idea of the aluminum sample box was to provide additional magnetic screening for the device using the Meissner effect. Surprisingly, though, recent experiments [113] have obtained substantially better results by using copper sample

boxes instead of aluminum. It is conjectured that this is because of the better thermalization afforded by using a normal metal instead of a superconductor. Similarly, painting sample boxes with an absorbing material [114], or potting the sample in absorbing epoxy [115], appear to improve T_1 by absorbing thermal radiation.

5.1.2 Wiring setup

In the BBC2 device, I use three lines that pass through the filter bank for low frequency applications. One line provides the bias current. A second line monitors the voltage of the sample to measure switches. The third line flux biases the sample. The bias and voltage lines are joined by a tee after the filter bank, entering the sample box on a common line. This allows four-wire measurements of the qubit voltage. So two of the ports of the sample box are used for low-frequency lines.

While the impetus for the BBC2 design was reducing the overcoupling of the microwave line measured in earlier devices [28], another appeal of the coplanar microwave feedline was the prospect of extracting information about the qubit through the microwave line. For example, the continuous feedline allowed me to measure the dc inductive coupling of the microwave line to the sample. A longer term goal was dispersive microwave readout of the qubit state a la circuit QED (see Section 3.2). To accommodate this requires microwave lines in and out of the sample box. The input and output lines are not symmetric. While both lines must prevent noise from reaching the qubit, the output line uses directional attenuation to shield the qubit without degrading the desired signal. The input line uses thinner stainless steel coax and attenuators. The output line uses thicker coax to reduce dissipation and cryogenic microwave isolators to provide ≈ 40 dB of directional attenuation.

The BBC2 experiment involved some rewiring of the fridge, along with repairs to some of the older parts. What follows is a catalog of the final setup for the five relevant lines pictured in Figure 5.3:

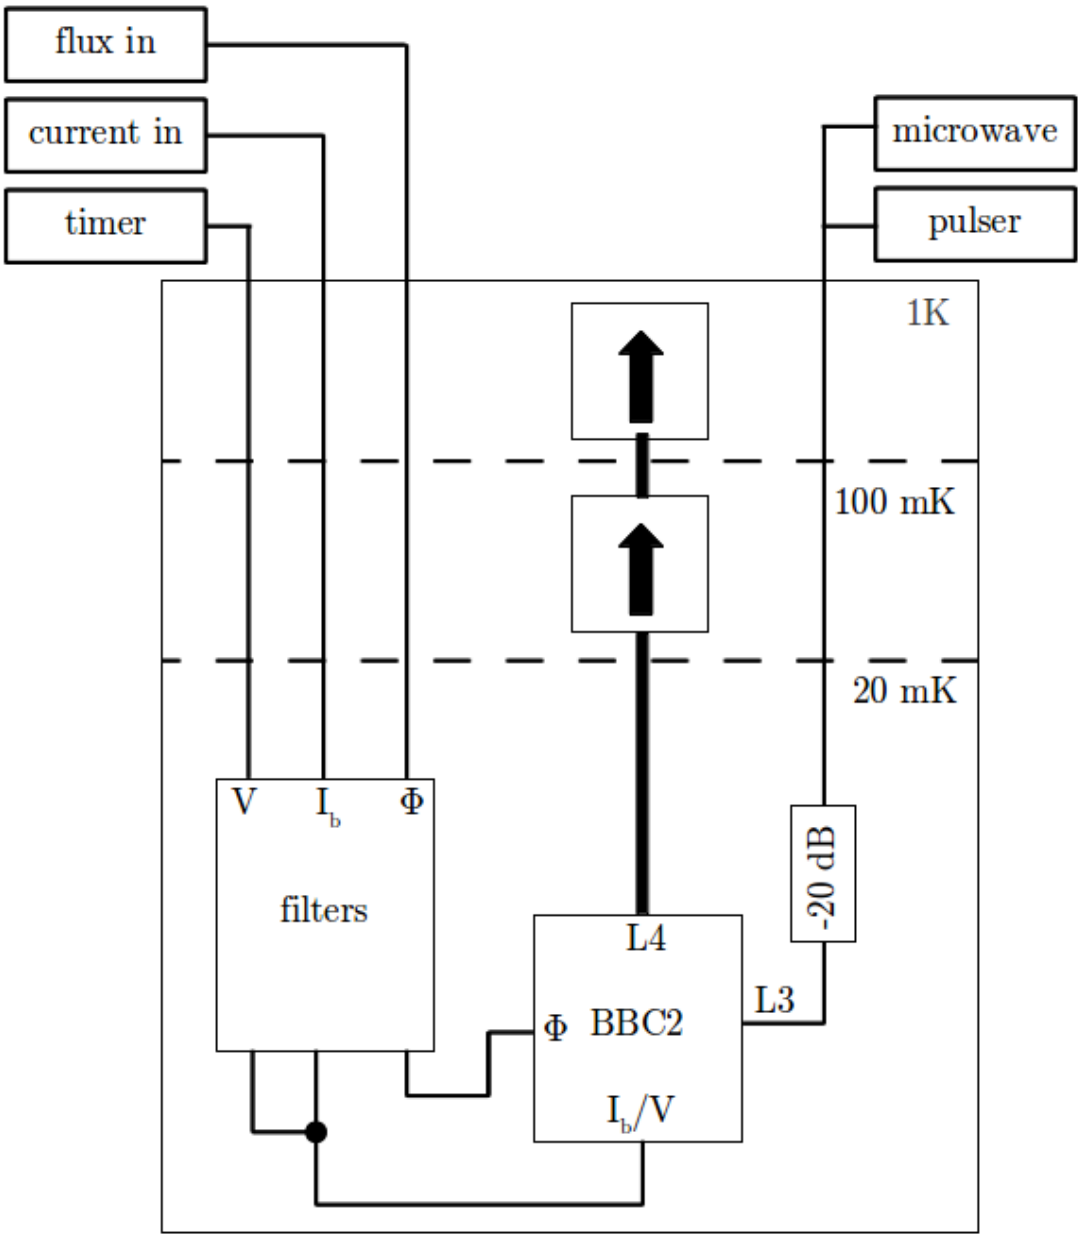

Figure 5.3: Simplified schematic drawing of the experimental wiring in the refrigerator for the BBC2 experiment. The left hand side shows the low-frequency wiring; the right hand side the microwave wiring. The microwave output line is left open atop the refrigerator.

- The bias current line I_b entered the fridge through the junction box on the top on the port 1B. This line goes continuously from the top of the fridge to the bottom and has not been changed in years.

- The remaining four lines that I used all entered through four distinct ports labeled L1 through L4 on the left side of the fridge. Each port has a threaded steel feedthrough. Female SMA connectors are soldered to brass flanges for rubber O-ring seals. All four of these lines are heat sunk to 4 K through a copper foil mounted to the brass flange the vacuum can is mounted on.

- L1 was used as the flux line. A length of UT-85 stainless steel inner and outer coax (abbreviated SS/SS) goes through the foil heatsink and is tied to the 1 K pot exterior. There it is connected to UT-85 flexible cable made of silver plated copper ("flexi"). This short link is attached to the outer edge of a copper heatsink bolted to the 1 K pot. This heatsink has two SMA feedthroughs soldered to it. On the other side of the heatsink, a UT-20 cable winds down to the junction box at the still where it is heatsunk again. The heatsinking in the still junction box consists of an enamel wire GE varnished to the wall of the copper box. L1 enters the still junction box in the second port from the top. Across the junction box is home-built superconducting coax, with stainless steel tubing and copper-plated niobium respectively as the outer and inner conductors, insulated with Teflon tubing. This coax continues down to the copper powder filter setup, entering powder filter 1, and is finally attached to the sample box through a short length of UT-85 flexi.

- L2 was used as the voltage line. UT-85 SS/SS extends from room temperature to the copper heatsink at the 1 K pot. From the heatsink to the junction box, UT-34 SS/SS coax is used, entering the junction box on the third port from the top. An unknown UT-20 cable goes from the junction box down

to copper powder filter 4, passing through a couple of other stages of non-breakout heatsinking.

- L3 was used as the microwave input line. Since we desire high frequency signals on this line, it is not run through the powder filters. The entire line is UT-34 SS/SS with two interruptions. There are connectors at the still, allowing for the introduction of an attenuator at this point if one wishes. (During the BBC2 run, we did not use an attenuator at this location.) At the mixing chamber, a 20 dB SMA attenuator was used. At 10 GHz, the cable is specified to have a loss of 10 dB per foot at room temperature. At cryogenic temperatures the loss will be less, but likely still sufficient at the qubit transition frequency.

- L4 was the microwave output line. UT-85 SS/SS runs from room temperature through the foil heatsink. To attach to the isolator, a connector was introduced to allow a simple right angle bend. Another length of UT-85 SS/SS runs to the output port of a PAMtech cryogenic isolator. The isolator is in-band from 4-8 GHz and provides a minimum of 18 dB of directional isolation. The isolator is mounted on a bracket anchored to the 1 K pot. From the input port of this isolator, a length of UT-85 SS/SS runs to a second, identical isolator anchored to the flange below the continuous heat exchanger. To facilitate mounting, a short length of UT-85 flexi goes from the input port of this isolator to a length of UT-85 SS/SS running down to the sample box.

  As discussed in Chapter 7, the BBC2 experiment was performed over the course of three cooldowns. The described setup for L4 — which is how the fridge is presently configured — was inserted between the first and second cooldown. On the first cooldown, the cable from the sample box all the way up through both isolators was present, but instead of continuing to room temperature, a 50 ohm termination was attached to the output of the higher

temperature isolator.

Switching from a terminator to complete transmission through the fridge offered some immediate advantages. It allowed me to measure the mutual inductance between the microwave line and the qubit, as described in Section 6.3.2. Surprisingly, it also greatly simplified the measurement pulse setup, allowing "dc" pulses instead of pulses on a 1 GHz carrier as described in Section 5.3.3. It is not clear why this would be; if anything, the opposite seems more likely since L4 was left open during typical operation.

5.2 Electronics setup

5.2.1 Instrumentation

Figure 5.4 shows a block diagram of the measurement setup. I will first summarize the flow and then discuss specific instrumentation details. A function generator produces a clock signal that starts each individual measurement. This clock starts programmed current and flux waveforms. These are sent respectively to the bias and flux lines through low-pass filtered ports in the screen room wall. A sync pulse from the current bias function generator starts a timer. The sync pulse also triggers the digital delay generators used for measurement pulses and gating the microwave sources. After some time, the sample switches to the voltage state, sending a pulse to the voltage line. This pulse is amplified by a homemade amplifier chain and then stops the timer. The time t recorded by the timer is the data point for this repetition. An actual measurement consists of thousands of repetitions, building a histogram of switching times.

Here is the instrumentation in further detail:

- The master clock is an Exact model 628 function generator mounted on the refrigerator instrument rack. It puts out a 7.5 V inverted square wave with

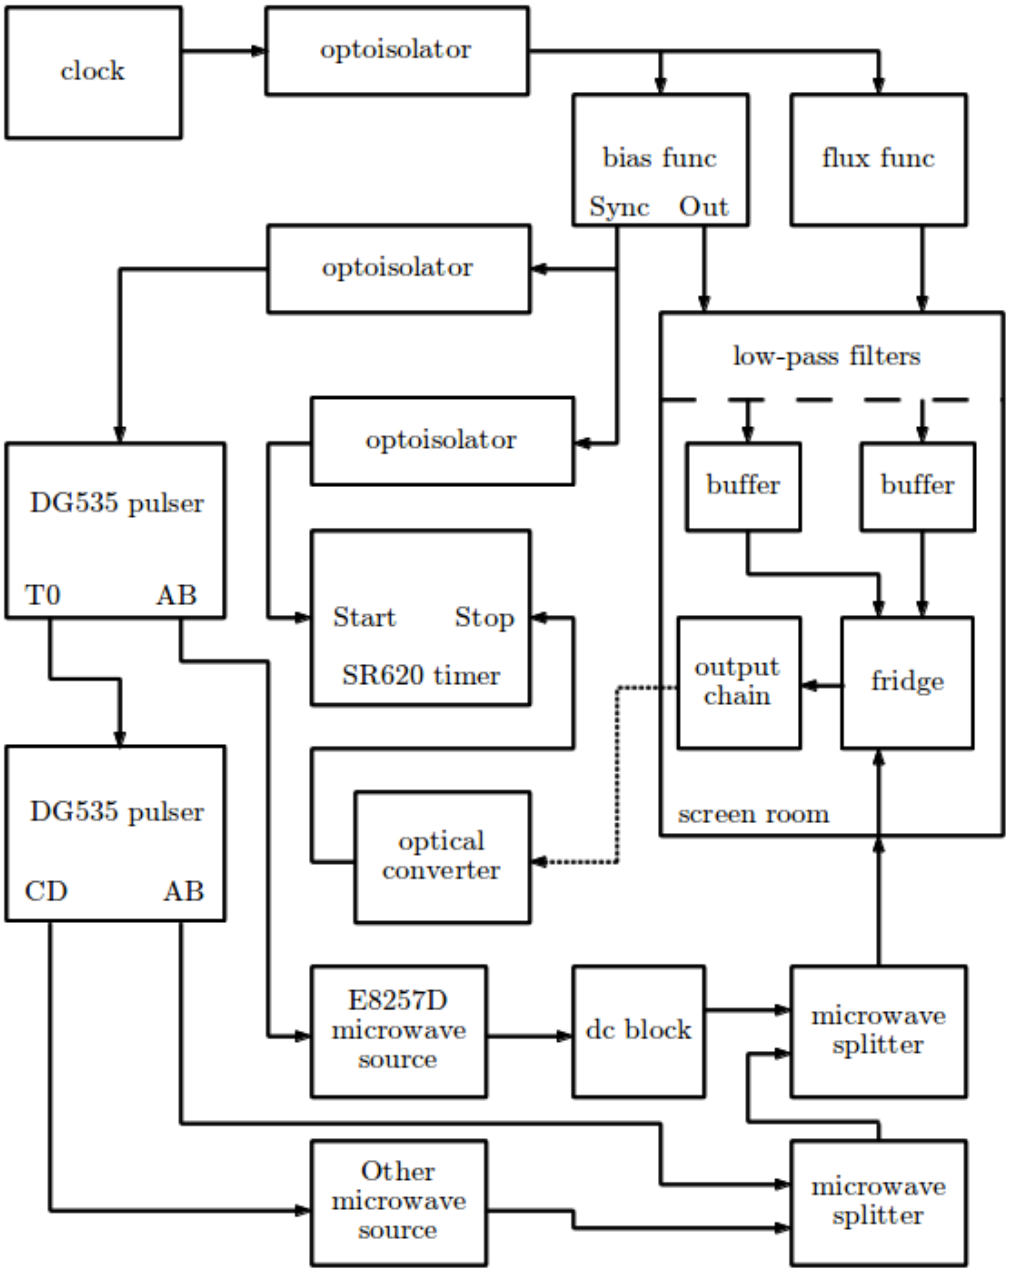

Figure 5.4: Block diagram of the measurement electronics setup. The fine dotted line represents an optical fiber connecting the Schmitt trigger in the output chain to an optical converter outside the screen room.

4.7 V offset at 400 Hz, the clock rate of the experiment. The output from the function generator comes through a 6N137 optoisolator (labeled CLOCK) powered by a 12 V lead acid cell.

- From the optoisolator, the clock signal triggers the Agilent 33220A bias and Agilent 33120A flux function generators. The clock signal is split by a BNC T junction located on the external trigger port of the bias function generator. The output of the flux function generator is connected to a BNC port on the screen room wall labeled "Data 1" or "Flux." The bias function generator is T'ed at the output port, with the waveform being monitored on a Tektronix scope running on wall power. The other side connects to the BNC port labeled "Data 2" or "Bias." These ports both contain fairly aggressive low-pass filtering to smooth out the digitization in the waveform. The -3 dB point is at a few kHz, although there is a resonant enhancement near 20 kHz.

- The sync port on the bias function generator is also T'ed. The two arms of the T connect to two more battery powered optoisolators labeled "START" and "PULSER." For most experiments, "CLOCK" and "PULSER" were run off a shared 12 V cell, with "START" getting its own dedicated cell.

- The "START" optoisolator leads to the start trigger on a Stanford Research Systems SR620 timer. The timer has adjustable trigger points for both the start and stop ports. By choosing the trigger value where the pulse from the optoisolator is at its steepest, I minimize the impact of voltage noise on the uncertainty in the time of the waveform start. I set the start trigger to 2.09 V and the stop trigger to 1.25 V.

- The "PULSER" optoisolator leads to the external trigger port (set to 50 ohms) on a Stanford Research Systems DG535 digital delay generator. The delay generator was used to send the short measurement pulse down the microwave

line and also to gate the microwave source. The DG535 has four separate output channels. These can be turned on at specific times relative to the initial trigger. Two special output ports, labeled "A Π B" and "C Π D", will put out pulses between the turn-on times for the two channels listed in the port name. Only these ports — which I abbreviate AB and CD — gave control over the start and stop time of the signal, so AB and CD are the channels I used for sending pulses. Since both the measurement pulses and the microwave gate require SMA connections, I used BNC-SMA adapters directly on the BNC ports of the DG535.

The output ports can output at standard voltages such as TTL or ECL, but for measurement pulses I used the adjustable voltage mode, which allowed voltages between 0.10 V and 4.00 V in steps of 0.01 V. As I discuss in Section 7.2, adjusting the pulse height is essential for measurement. I usually found the 0.01 V increment too large, since it led to large changes in the pulsed escape rate. I addressed this problem by putting a 50 Ω SMA attenuator on the line. For the BBC2 sample, I used a 5 dB Minicircuits attenuator.

Because only two outputs are useful for pulsing, I use two DG535 boxes. The second one is triggered from the T0 output of the first (i.e. a sync signal turned on when it is triggered.) The AB output from the first DG535 is used to gate an Agilent E8257D microwave source. The second DG535 gates a second HP microwave source from the CD output. The AB output of the second DG535 is used to source the measurement pulse; it is routed through a broadband microwave splitter (called splitter 1) on the front of the HP microwave source.

- The Agilent E8257D source was the primary source I used for qubit control. It is gated by the AB channel of the first DG535. Gating the source allowed me to turn the microwaves on and off at set points in the bias waveforms. In

particular, when ramping to a fixed bias plateau I could turn the microwaves off to avoid exciting unintended transitions. In addition to the hardware gate, I could turn the microwave output on or off through software. The RF output of the E8257D goes through an inside-outside dc block to another broadband microwave splitter (splitter 2). The measurement pulse and output from the HP microwave source (used for two-photon experiments) are combined through splitter 1. Splitter 2 combines the output of splitter 1 with the E8257D output and sends to the refrigerator through an SMA bulkhead in the screen room wall.

- The bias, flux and voltage lines all feature custom built electronics run off batteries inside the screen room. To protect the sample from spikes in current that may occur when devices are turned on or off, all three lines have a grounding switch as the last element before entering the fridge. These switches have BNC input and outputs. The center pins of the two BNC connectors are soldered together. The throw on the switch is open when the switch is set to allow measurements, and shorted to ground when the sample is grounded.

  Before entering the fridge, the bias and flux lines go through custom-built unity gain op amp buffers. These present a high input impedance to the function generators. This keeps the wall filters from loading the generators and breaks ground loops in the circuit. The buffers are built using an Analog Devices AMP03 op amp powered by 15 V chip regulators in the box; each buffer is powered with two series 12 V lead acid cells. To make the function generators behave more like current sources, resistors are placed on the output of the buffers. I used a 100 kΩ resistor for the bias line and a 1 kΩ resistor for the flux line. The relatively weak inductive coupling of the flux coil requires higher currents; mitigating possible heating by these currents is why much of the flux line is superconducting coax.

When the SQUID tunnels to the voltage state, the voltage increases until it corresponds to the electrostatic energy necessary to break a Cooper pair. For an aluminum device, this corresponds to about 400 μV. Substantial amplification is needed to bring this to the $\sim$ 1 V order of magnitude needed for triggering. Following the voltage line switch, I use a two stage amplifier chain built for our lab's earliest qubit experiments [71, 110]. The first amplifier is a low-noise amplifier using 16 JFETs as the active element. The second stage amplifier is ac coupled and can be switched to run inverted. This allows the amplifier chain to be used to measure voltage pulses of either sign, which could arise by changing the sign of the bias current. The second stage amplifier outputs to a Schmitt trigger with both electrical and optical outputs. The electrical output is only a debugging mechanism; during experiments, an optical fiber carries the switching data out of the screen room. All three devices in this chain are powered using low noise 15 V power supplies which are in turn powered by lead acid batteries.

- The pulse on the optical fiber is converted back into an electrical signal by a box outside of the screen room. This box is run off a custom built 12 V wall-powered power supply. Finally, this signal is sent to the stop port on the SR620 timer. Receipt of this pulse signals the collection of one data point. The raw data for the experiment is the series of stop times on the timer, typically collected in a histogram. With the next cycle of the master clock, the measurement repeats.

In addition to the measurement electronics, there is also an instrument rack with tools for operating the refrigerator. Resistance thermometers are mounted at various stages in the fridge; the mixing chamber in particular uses a calibrated ruthenium oxide thermometer. The resistances were measured using a Picowatt AVS-47 active resistance bridge. A heater controller sourced specified power to

wire-coil heaters at the mixing chamber and still. A separate circuit was available to heat the 1K pot. All heaters were attached to a temperature controller, although I did not use this in measuring BBC2. During cooldowns the resistance bridge was attached to the computer through a special GPIB link to allow logging of the temperature data.

5.2.2 Computer control

Most data collection was automated using LabView code on a Windows XP machine. Communications to the various instruments were performed using GPIB. The central program is named "OneSR." OneSR sends the SR620 timer a request to buffer a user-specified number of data points (typically 1000). After the switches are recorded, the data is sent to the computer before the next set of events is taken. OneSR has a flag that can be set by other code to pause the data gathering. This is typically used to change measurement settings (e.g. microwave frequency, pulse height.) It also has code to produce a live histogram of the data gathered so far, a graph showing where the last 100 counts occurred, the tunneling rate inferred from the histogram, and the average switch location of the data in each buffer.

Given the large amount of data (and concomitant amount of time) needed, the data automation was necessary. However, the GPIB chain leads to a complicated set of connections between all of the instruments and the relatively noisy computer. The optoisolators and buffer amplifiers discussed in Section 5.2.1 break ground loops on the signal ports. The GPIB connections do not employ any of these safeguards (although some commercial solutions for optoisolated GPIB exist).

5.2.3 Grounding issues

My initial measurements on BBC2 suggested there were electrical issues with the experimental setup. I will show the data suggesting the problem and discuss my

approach to fixing it.

Figure 5.5 plots some of my first spectroscopic data from BBC2. This measurement used the swept measurement technique discussed in Sections 5.3.2 and 7.1. The x-axis is the applied bias current I_b. The y-axis is the frequency of the continously applied microwave drive. Higher frequencies were taken at higher powers to counteract the extra dissipation in the microwave line at higher frequencies. The color represents the number of counts measured in a particular time bin.

Whenever I measured microwave effects, I took several hundred counts with the microwaves on ("microwave") followed by several hundred counts with the microwaves off ("background"). Since the background rate can change with time due to changes in the bias (intended or through drift), this ensured that I compared microwave data to the relevant background. Figure 5.5a) shows the microwave data, Figure 5.5b) the background. In both plots, near 5700 MHz the character of the histogram changes. Below 5700 MHz the histogram is broader and bimodal; above 5700 MHz, the histogram collapses down to the earlier of those two peaks. What is striking is that this apparent frequency dependence occurs even in the background data, where no microwaves are being applied.

The change from unimodal to bimodal histogram behavior also showed hysteresis. Once I crossed the threshold on the microwave source where the transition occurred, the sample would continue to produce unimodal histograms until I reverted to a different setting on the microwave source. This held true even when I grounded the sample for a couple hours before measuring again. Further testing revealed that the shift was not a property of the applied frequency, but rather the applied power. When the microwave source was operated at an output power of -5dBm or higher, the histogram was in the unimodal state, regardless of the microwave frequency. The appearance of a frequency dependence from Figure 5.5 arose because I had used high power to counteract the loss at high frequencies.

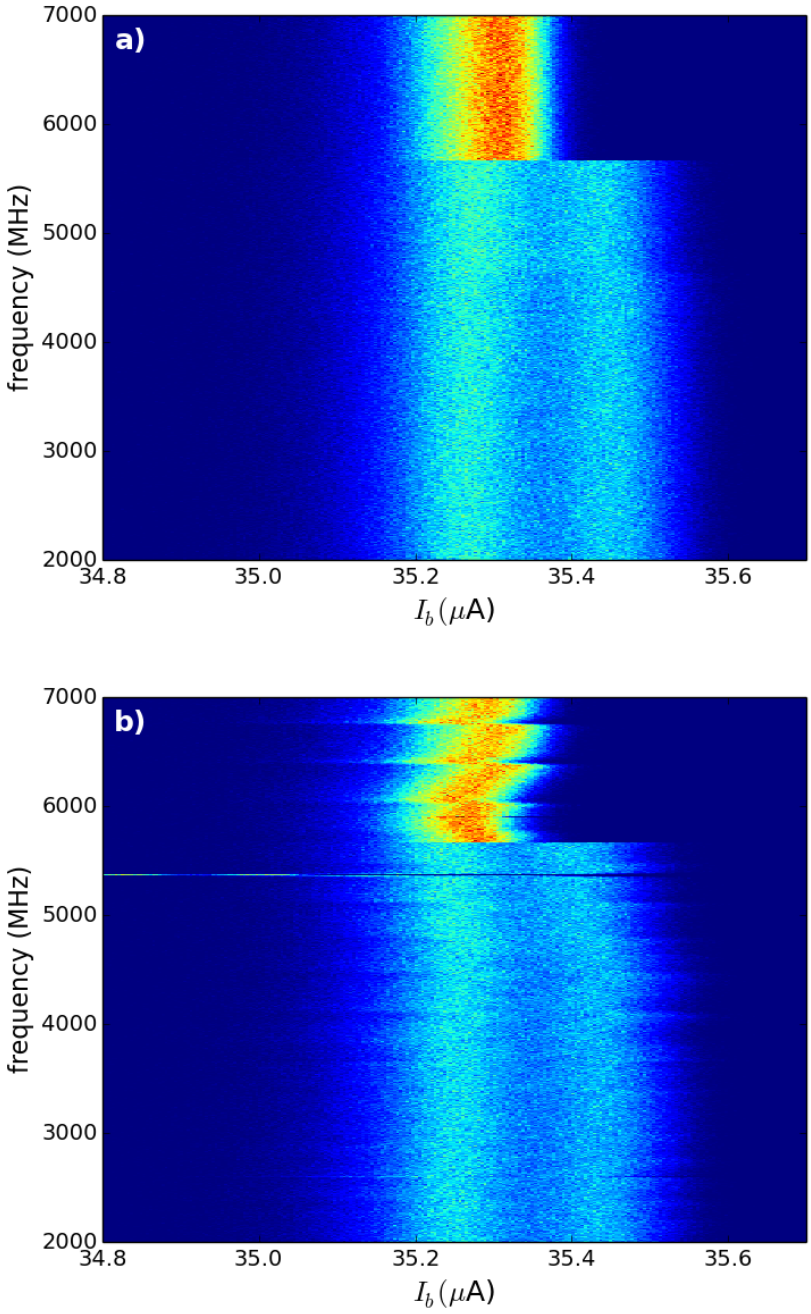

Figure 5.5: a) Measured histograms at several microwave frequencies and powers. b) Histograms taken during the same run with microwaves turned off.

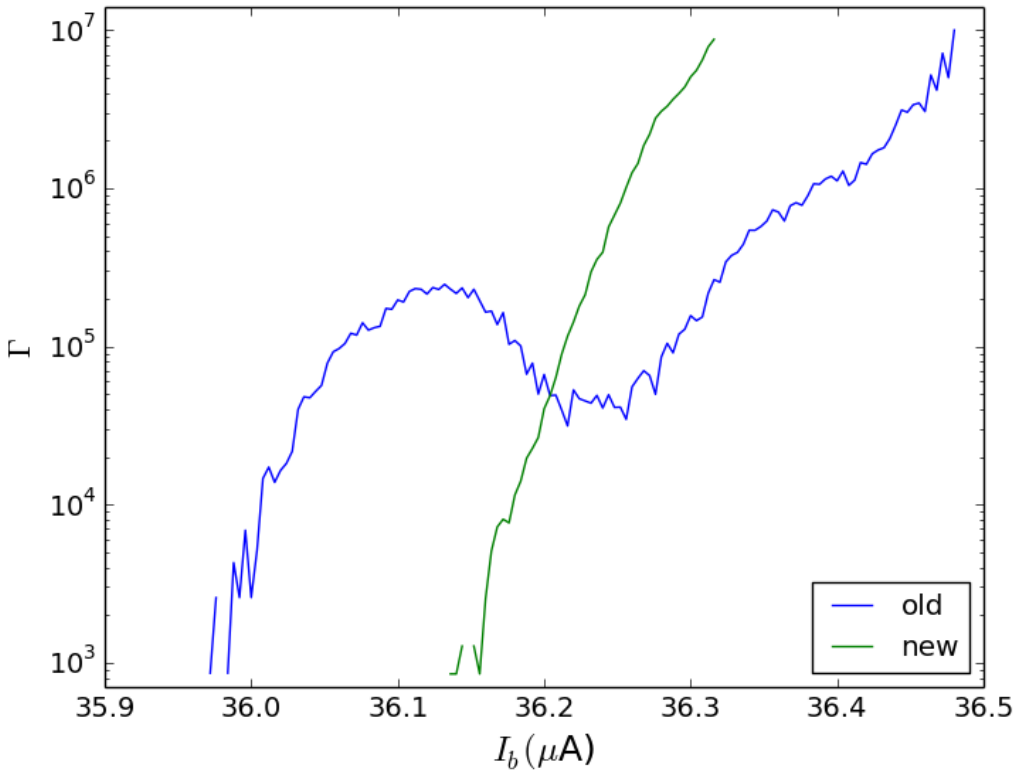

Figure 5.6: Calculated escape rates from histogram data. The two curves correspond to two different electrical setups.

The persistence of the phenomenon when microwaves were turned off, along with a lack of frequency dependence in the power at which the histogram transition occurred, suggested the histogram transition did not depend on the properties of the microwaves at the sample. The next thing to check for was a room temperature problem. Unplugging the microwave input line, either at the input atop the fridge or at the screen room wall, appeared to eliminate the problem. Introducing the inside/outside dc block to the microwave line appeared to give good results. However, for these swept measurements the DG535 pulsers were not being used. Introducing the pulsers to the measurement circuit caused the problem to return. At this point, I removed the GPIB connection to the microwave source; this appeared to make the issue go away.

After considerable tinkering with the setup, I finally arrived at a configuration that gave the quietest looking histograms. Figure 5.6 plots the tunneling rate calculated from the measured histograms for two different wiring setups. The x-axis is the applied bias current I_b, and the y-axis is now the tunneling rate in events per second. An advantage of plotting escape rates instead of the histograms is that modest changes in histogram shape are often more clear as features in the escape rate.

The "old" data of Figure 5.6 shows the escape rate with no microwaves applied in the original experimental setup. In this setup, most of the instruments were plugged in through strips attached to two isolation transformers. The exceptions were the computer and the thermometry rack, which were plugged into receptacles on the screen room that had been recently installed for safety reasons. The old escape rate contains a broad hump; this is how the bimodal histogram shows up in the escape rate.

The new setup reorganized the electronics so that the computer was now plugged in to one of the isolation transformers and the thermometry GPIB connection removed altogether during measurement. In particular, the computer, the microwave source, the timer and the pulser were all on the same power strip attached to the isolation transformer. I also reordered the configuration of GPIB cables. Compared with the old setup in Figure 5.6, the escape rate is much cleaner. The most likely source of improvement was changing where the computer got its power from; later measurements suggested the thermometry GPIB was not the primary source of trouble.

While these corrections were enough to perform interesting measurements on the sample, the obvious question is whether I eliminated all of the problems. Indeed, it is very likely that some issues remained. Two phenomena in particular suggested trouble. In Section 5.1, I discussed the microwave output line entering the fridge

through the port L4. During most measurements, this port was left open. When I put a 50 Ω terminator at L4, the histograms became extremely noisy, suggesting a coupling of noise from the fridge ground through the terminator to the center line on the coax. I also observed occasional dark counts on the timer, where the stop port on the timer would trigger despite the amplifier chain being turned off.

5.3 Measurement techniques

All of my data on BBC2 used the voltage switching readout. However, I used two different approaches to generating switches. One method — "swept measurement" — linearly ramped current and flux biases up to and through the critical current of the SQUID. Near this critical point, the tunneling rate Γ to the voltage state is very high and switches are easily visible. The second method — "pulsed measurement" — instead ramped to lower values of current and flux and then plateaued. To enhance the escape rate, I sent a short measurement pulse to the qubit. In this section I will discuss aspects of these measurement techniques; further details are available in Chapter 7. With both approaches, I began the flux ramp with a sinusoidal current to "flux shuffle" the SQUID, a technique for speeding data collection [116, 117].

5.3.1 Flux shuffling

For values of the parameter $\beta \equiv 2\pi L I_0/\Phi_0 > 1$, there are multiple possible stable equilibria for the SQUID. Intuitively, trapped flux quanta in the SQUID produce loop currents. If the loop current added by introducing another flux quantum does not exceed the critical currents, then the SQUID is stable there as well. Large critical currents or big inductances (and thus small loop currents) make this type of current rearrangement easier. I refer to these multiple equilibria as flux states. Since

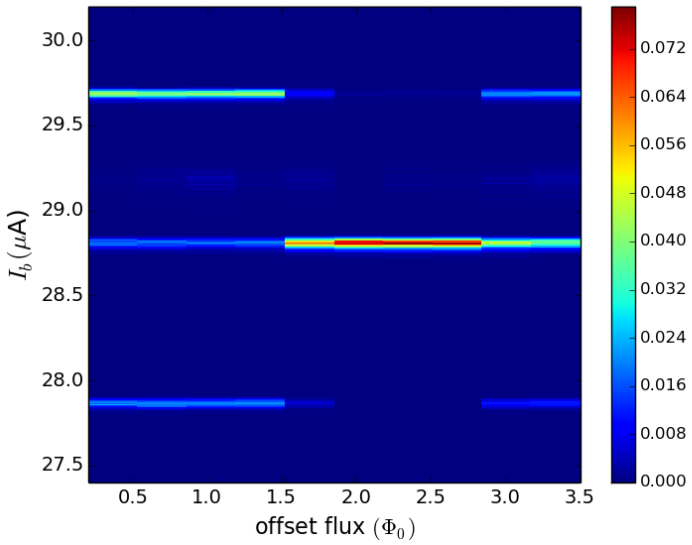

Figure 5.7: Flux shuffles at fixed amplitude but varying offset.

each of the flux states is at different bias conditions, an experiment set up assuming one flux state will not typically work very well in the others. Without special preparation, the qubit starts in a random flux state, leading to wasted counts.

Flux shuffling reduces the amount of data needed by preparing the system in a chosen flux state. The flux states correspond to the wells remaining after the magnetic energy associated with the inductance bends the washboard into a parabola. Applying flux tilts these wells, deepening those on one side of the parabola while destabilizing those on the other. As a flux state turns unstable and vanishes, the SQUID retraps in a new flux well. By carefully choosing the amplitude and offset of the sinusoidal flux shuffling signal, one can arrange so that only one flux state is stable at all times during the waveform. As the flux cycles back and forth, the other flux states empty out, but this one flux state continues to fill. After several cycles of the shuffling signal, the probability of being in the chosen state approaches 1.

Figure 5.7 is a plot from a scan of flux shuffling parameters. All of the shuffles in this plot were taken at a single amplitude of the shuffling waveform; the full

scan consists of numerous plots like this. The x-axis is the offset of the flux shuffle waveform. The y-axis is the applied bias current. The color is the fraction of all switches for a given offset occurring in a particular bin. At most offsets, three distinct flux states are visible. For offsets between 2 and 2.5 Φ_0, the shuffle is effective and the three flux states collapse to the central flux state.

5.3.2 Swept measurement

The key feature of swept measurement is that, for at least a portion of the measurement, the bias and flux currents are both linearly ramped. As shown in Equation 2.13, the Josephson inductances of the SQUID junctions are sensitive to the current passing through them. Thus, adjusting the bias leads to changes in the frequency of the qubit. This allows me to apply continuous wave microwaves[2] through the whole ramp. When ω_{01} moves through the applied frequency ω with the changing ramp, the escape rate is enhanced. By mapping where this enhancement occurs for a series of drive frequencies ω, I establish the spectroscopy of the qubit.

Figure 5.8 shows a timing diagram for the various sources in a swept measurement. Following the flux shuffle, both the flux and bias currents are linearly ramped. The bias ramp has a small negative offset. Since the junctions are strongly hysteretic, this helps ensure that the SQUID retraps to the zero-voltage state following each shot. The amplitude of the bias ramp is chosen to exceed the maximum critical current of the SQUID. This guarantees that each waveform will eventually produce a voltage pulse. The data is collected into a histogram of switching times. This histogram is then used to calculate an escape rate following the technique of Fulton and Dunkleberger [118]. That calculation is described in Section 7.1.

The ratio of the ramp currents in the flux line and the bias line is not chosen arbitrarily. It is dictated by the current-flux characteristic of the SQUID, discussed

[2]As discussed in Section 5.2.3, half of the trials have the microwaves turned off to establish background data.

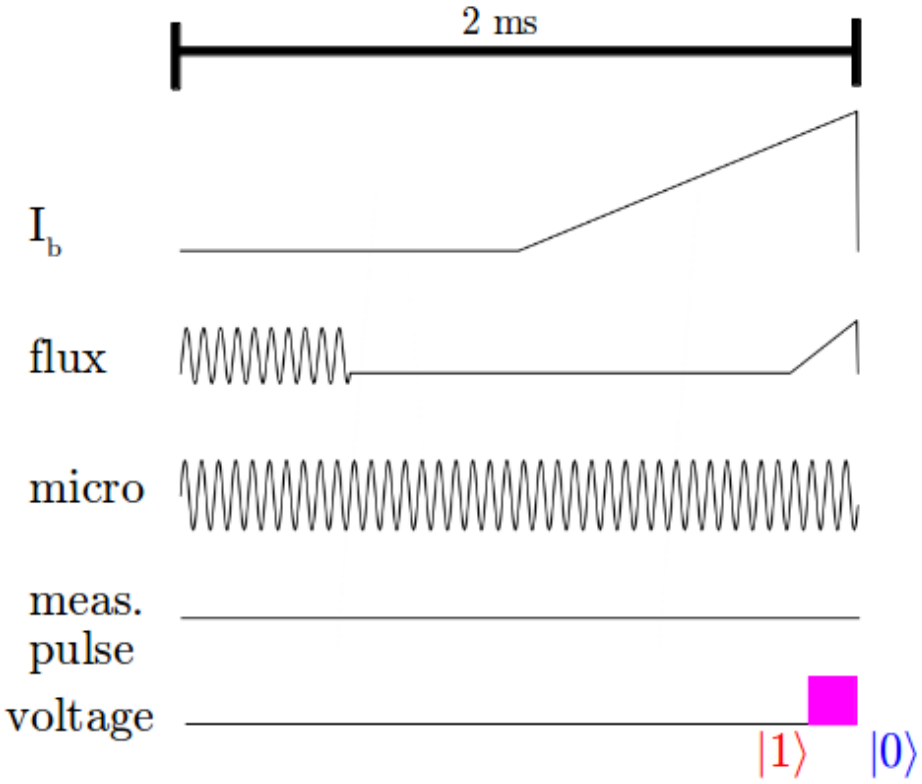

Figure 5.8: Timing diagram for a single shot of swept measurement. The bias current I_b and flux are linearly ramped. Microwaves are applied continuously; no measurement pulse is used. Voltage pulses are measured in the high-bias region near the end of the ramp.

in Section 6.3.2. The arms of the current-flux characteristic correspond to switches associated with the two different junctions of the SQUID. Since our design tries to emphasize the more well- isolated, smaller qubit junction, I am interested in switches that occur on the correct arm of the characteristic. Ramping both the current and the flux — a "double ramp" — avoids switching on the undesirable arm. For the BBC2 device, the qubit arm was the one which switched with decreasing current as flux was increased.

Earlier samples used double ramps where the flux and bias were linearly related over the entire ramp region. As seen in Figure 5.8, I did not double ramp this way. The critical current of the big SQUID junction was almost an order of magnitude larger than most of the lab's recent samples. This larger current on the bias ramp means that a fixed linear relationship between flux and bias would require a flux current peaking at around 6 mA. At this much current, dissipation on the flux line leads to heating in the refrigerator.

Instead, I followed a more complicated trajectory on the current-flux characteristic. This path is shown in Figure 5.9. The x-axis is the applied flux; the y-axis is the applied bias current. The dashed line is a cartoon showing the path my chosen ramps followed. I started applying only bias, moving upward on the characteristic plot[3]. As I approached the region where switches are visible, I began the flux ramp. This technique helps with heating in two ways. Since I ramped over a shorter region, the maximum current I applied on the flux line was considerably smaller. The overall duty cycle of the flux waveform was also reduced by a factor of 5. This let me measure switches on the qubit branch without heating.

The advantage of the swept measurement is that crossing the resonance anywhere in the measurable region (i.e. where the escape rate is high enough for the qubit to switch) yields microwave enhancement in the escape rate. Figure 5.10

[3]The bias line includes some mutual inductance to the SQUID loop, meaning that any linear bias ramp will also include some small flux contribution. I typically ignore this.

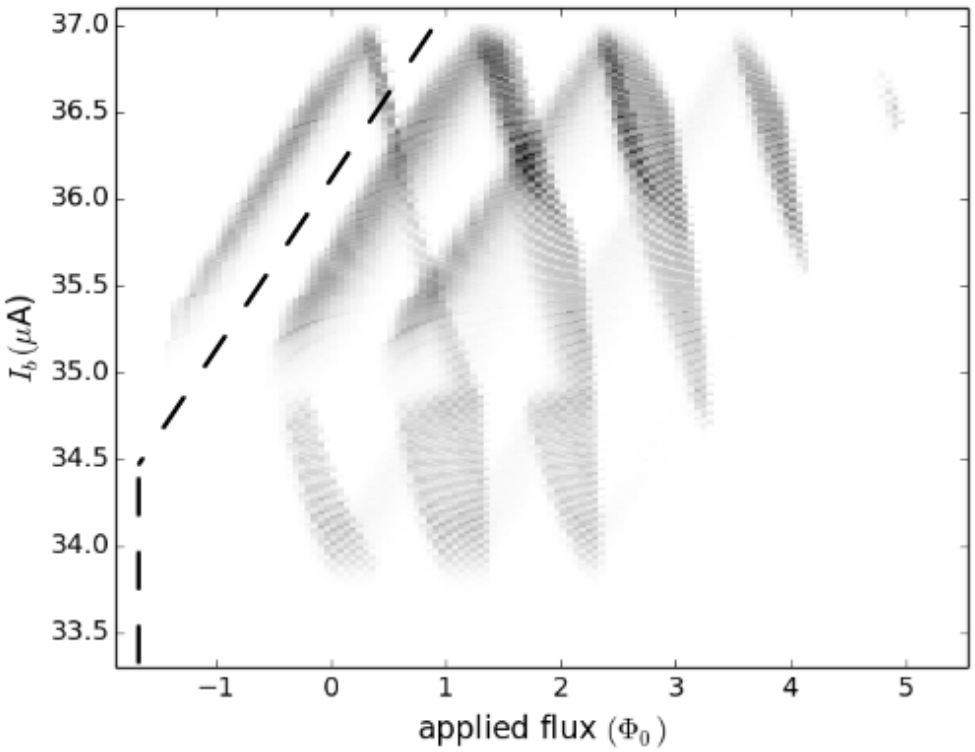

Figure 5.9: Current-flux characteristic for the qubit. The dashed line is the path traced in flux and bias by the ramps shown in Figure 5.8.

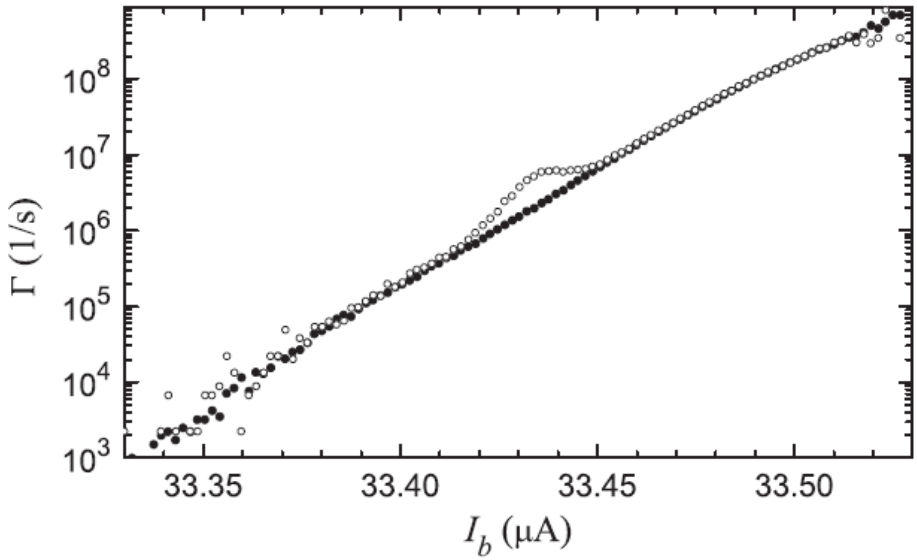

Figure 5.10: Example of resonant escape rate enhancement in the swept measurement. This plot is taken from Figure 6.5b in Sudeep Dutta's thesis [25].

shows this microwave enhancement in data taken from Sudeep Dutta's thesis [25]. The white circles are the microwave data; the black circles are the background data. The microwave peak occurs at the bias I_b such that the qubit resonance ω_{01} matches the applied microwave frequency ω. To map the system spectroscopically, only the microwave frequency needs to be scanned. Compared to pulsed measurement, far less data is required for the spectroscopy.

While this convenience is very nice, the swept measurement also has several disadvantages. Since I can only measure the SQUID in places where it tunnels, I can only perform swept measurements when the qubit is biased close to I_c. At high bias, the Josephson inductance (2.13) changes rapidly with small fluctuations, yielding large $d\omega_{01}/dI_b$. This sensitivity causes a substantial reduction in T_2 as discussed in Section 2.4.2 and manifests as "inhomogeneous broadening" [56, 119]. The swept waveforms also mean that ω_{01} is continuously changing. For short times the effect is minor, but long-lived coherent oscillations become complicated when the microwave drive has a time-sensitive detuning. Lastly, as I discuss in Section 7.1, for BBC2 in particular the microwave enhancement proved difficult to calculate because of the shape and location of the peak.

5.3.3 Pulsed measurement

Pulsed measurement addresses all of those problems. The key difference with pulsed measurement is that the current and flux bias are brought to some fixed level and then held constant during the course of the experiment. To measure the qubit state, I use a brief current pulse to quickly bring the system back to strong bias. The pulse height is chosen so that during the pulse the SQUID tunnels with high probability if and only if it is in the excited state.

Figure 5.11 shows the timing diagram for the pulsed readout. The waveforms begin with a flux shuffle. On the bias and flux lines, the linear ramps begin some-

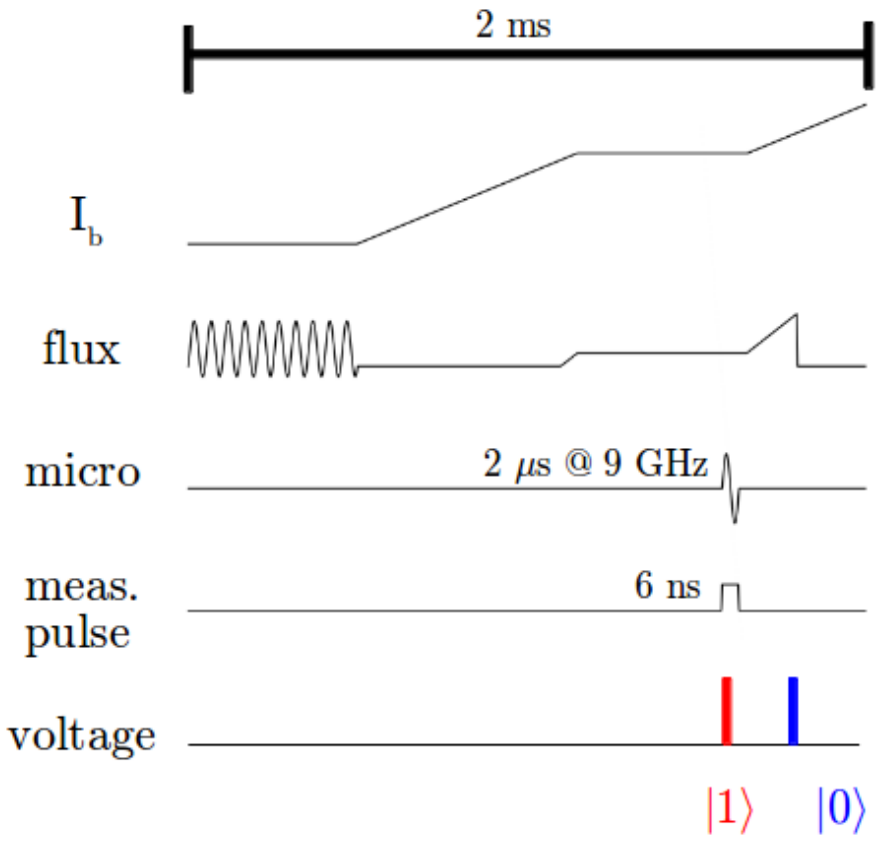

Figure 5.11: Timing diagram for a single shot of pulsed measurement. The ramps of Figure 5.8 are interrupted and stretched out to a long plateau. Microwave power is applied only in a small region of the plateau. Using measurement pulses I can measure the qubit state at any point in the plateau. Following the plateau, there is a short ramp to force a switch.

what earlier, allowing room for a long plateau. By waiting at the plateau value for hundreds of microseconds, ringing introduced by filters on the bias and flux lines can settle down to yield steady operating conditions. Near the end of the plateau, microwave power is applied, typically for 2-3 μs. As I discuss in Section 7.2.1, by adjusting the timing of the measurement pulse relative to the microwave turn-on I can perform three different types of microwave measurement. The waveform finishes by ramping flux and current to force a switching event on every shot. This effectively records the total number of repetitions in the data, which is necessary to calculate the switch probability P_{sw} in response to a pulse.

I used two different approaches to generate measurement pulses. The straightforward approach involves attaching the AB output of a DG535 directly to the splitters on the microwave sources. Channel B is set to be 6-10 ns after Channel A, and the voltage of the pulse is adjusted to control the pulse current. The pulse is sent down the microwave line and transmitted to the qubit through the coplanar waveguide feedline.

In the second approach, I used a 6 ns pulse from the DG535 to modulate a 1 GHz carrier signal from the second microwave source. The carrier signal was applied to the local oscillator (LO) channel of an Eclipse Microwave 1020 doubly-balanced ring mixer. The pulse from the DG535 was applied to the IF channel, and the output on the RF channel of the mixer was sent to the splitter on the front of the Agilent source. Since the mixer demands relatively high power, even with strong isolation on the mixer an undesirable amount of LO power reaches the RF channel. To combat this, I used another channel of the DG535 to gate the 1 GHz microwave source in a small window surrounding the desired pulse. Figure 5.12 shows an oscilloscope trace of one of these pulses amplified by a MiniCircuits microwave amplifier. The most significant part of the pulse is the 7-8 ns high amplitude region at the beginning, which is followed by a small secondary pulse. Since the escape rates depend

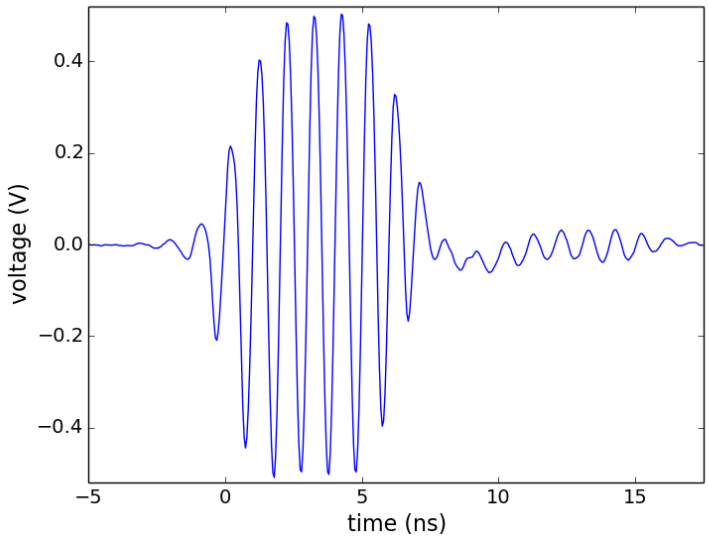

Figure 5.12: Oscilloscope trace of the modulated 1 GHz pulse. The main pulse body is 7-8 ns long.

exponentially on bias I do not expect the secondary pulse to produce switches.

I tried these two different pulse techniques because of the grounding problems discussed in Section 5.2.3. For the earliest measurements on BBC2, the microwave output line did not extend all the way to room temperature but stopped at the isolator mounted at 1K. Using the initial setup with the grounding issues, direct DG535 pulses were useful for measurement. However, after changing the setup to fix the grounding issue, the DG535 pulses did not appear to produce escapes even at very high voltages. I then tried the modulated 1 GHz pulses and found they worked. Curiously, after the change to L4 described in Section 5.1.2 the situation reversed. This stresses the need in future experiments to fully characterize the microwave lines from dc to 10 or 20 GHz.

5.4 Summary

In this chapter I reviewed the experimental setup for my measurements on BBC2. In the refrigerator, the important changes from earlier qubits were the sample

box and the rewiring of the microwave lines L3 and L4. I reviewed the adjustments I made to the electrical setup to eliminate the grounding issues demonstrated by my initial measurements on BBC2. Finally, I discussed how I biased the qubit and two different techniques for inducing switches.

Chapter 6

BBC2 qubit: design and initial characterization

Our lab has studied many generations of dc SQUID phase qubits. We achieved considerable improvement in the performance of our qubits by moving to designs featuring small junctions with an external capacitive shunt, as previously studied by the Martinis group [78]. While we first used parallel plate capacitors with a silicon nitride dielectric to accomplish this [29], interdigitated capacitors are substantially less lossy still.

This chapter focuses on the design and characterization of the BBC2 dc SQUID phase qubit. To explain the design, it is helpful to review some features of PB9 [28], an earlier dc SQUID phase qubit studied by our group. PB9 shunted the qubit junction with a large interdigitated capacitor (IDC). At its best operating points, PB9 offered the best T_1 times we have observed. However, at 350 ns, T_1 still fell well short of the > 1 μs design goal. I show a key plot from [28] suggesting why PB9's T_1 was so short and explain how the design of BBC2 accounts for this. Finally, I show some of the initial measurements I performed to characterize the BBC2 qubit.

6.1 PB9 and redesign impetus

With improvement in T_1 being the principal design goal, the most significant question is where the energy stored in the qubit can go. Unfortunately, this question is still not well understood. The most commonly considered possibilities are coupling to individual microscopic local quantum systems, dissipation from quasiparticles, or coupling to electromagnetic radiation. Two-level systems, manifesting as avoided

level crossings in spectroscopic data [75], are certainly present in junction barriers. The long lifetimes of qubits designed to minimize the importance of surface effects [107, 108] suggest two-level systems at surfaces are also important.

Similarly, there are many different ways for the qubit to radiate that need to be considered. Most (if not all) superconducting qubits are measured while in small metal boxes, presumably limiting straightforward dipole radiation from the qubit. However, recent work suggests that even though these sample boxes are operated at dilution refrigerator temperatures, thermal radiation continues to be a problem [114, 115, 120] through the generation of quasiparticles. This is an important insight, suggesting that engineering better shielding around the qubit can lead to improved performance.

On chip, we believe the dominant form of radiation is coupling to the leads we use to measure and control the system. On the PB9 device, there were three lines to couple to: the dc current bias lines, the flux bias line, and the microwave line. The discussion of the filtered dc SQUID in Section 3.3 explained how filtering reduced loss to the current bias line. For the flux and microwave lines, the goal was to minimize dissipation by making the coupling to the lines sufficiently weak. Thus, PB9's flux line was made relatively short and far away from the SQUID loop. The microwave was coupled with a short overlapping segment of two nearby striplines, presenting a very small capacitance. Photos of PB9 are shown in Figure 2.13.

Given all these couplings, it is natural to ask which one was the limiting factor for PB9's T_1. Figure 6.1 plots two separate measurements as a function of bias. The blue points are measurements of T_1 at the different bias conditions, with T_1 on the left y-axis. The dashed line is the inverse of the coupling between the qubit and the microwave line at the same bias points. This coupling is inferred through Rabi measurements. Since the Rabi frequency depends on the detuning and power of the drive signal, by operating at zero detuning one can use the Rabi frequency as

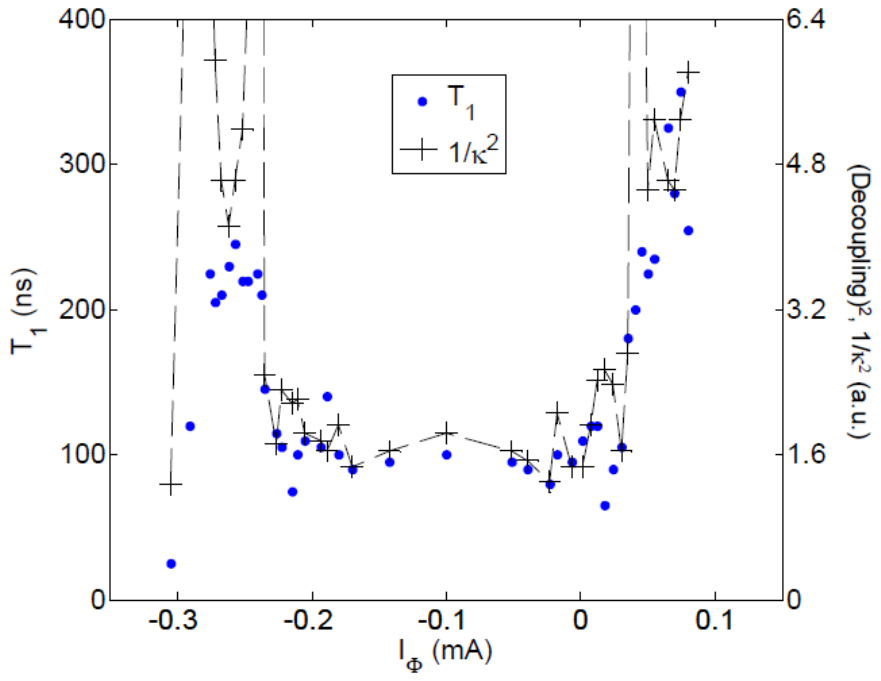

Figure 6.1: Plots of T_1 and inverse microwave couplings from the PB9 device as bias conditions are varied. This plot is from Tony Przybysz's thesis [28].

a power meter for the microwaves arriving at the qubit.

It is thus perhaps a misnomer to refer to this as the coupling. Since the qubit frequency is changing with the bias conditions, it is entirely possible that the coupling is roughly constant and the microwave line itself has wide variation in the power transmitted from room temperature. Whether the issue is coupling or the transmissivity of the line, in both cases the effect on the lifetime is the same. The current noise power spectral density $S_{II}(\omega_{01})$ transmitted from the microwave line to the qubit varies with bias conditions. As shown in Section 2.4.3, more noise means less T_1.

The importance of Figure 6.1 is that T_1 and the inverse of the coupling track each other remarkably well as the bias is varied. This is especially notable at low T_1 where the microwave coupling is strongest. At very weak couplings, the T_1 saturates, suggesting that some other dissipation channel became more important. Still, this plot was a very promising lead, suggesting that a more appropriately coupled

microwave line would yield immediate T_1 improvement at many bias conditions.

Since the intended coupling of the microwaves was already very small, we believed it was likely that they were coupling to the qubit in some way other than the intended capacitive coupling. To get a more controlled coupling, we decided to use an on-chip coplanar waveguide (CPW). This approach is quite popular in superconducting qubits. For us, particular inspiration came from the circuit QED architecture [30] and the work our colleagues at LPS had done on coupling lumped-element resonators to CPWs [31, 32, 121, 122] . Our goal for BBC2 was to make the microwave coupling more predictable and well-controlled. Rangga Budoyo designed the original version of this qubit [123], using Microwave Office to calculate the intended coupling.

6.2 BBC2 design

Figure 6.2 is a photograph of the qubit and its immediate environs from an early version of the BBC2 device. The CPW center conductor is highlighted in red. The qubit sits in a cavity within one side of the ground plane of the CPW. With this placement, the qubit should have limited exposure to microwaves entering in places other than the CPW. Not all elements of the qubit are visible here; to see elements of the circuit such as the LC filter and the flux lines, refer to Figure 4.5.

Given some of the complications introduced by the normal modes model discussed in Section 2.5.2, one of the goals for this qubit was to choose parameters closer to the independent junction regime. A quick heuristic for this is to compare the loop inductance $L = L_i + L_q$ to the Josephson inductance L_{qj} of the qubit junction; I choose the qubit junction since it, being smaller, has the larger Josephson inductance of the two junctions. The ratio is

$$\frac{L}{L_{qj}} = \frac{2\pi L I_{0q}}{\Phi_0} \equiv \beta_L. \tag{6.1}$$

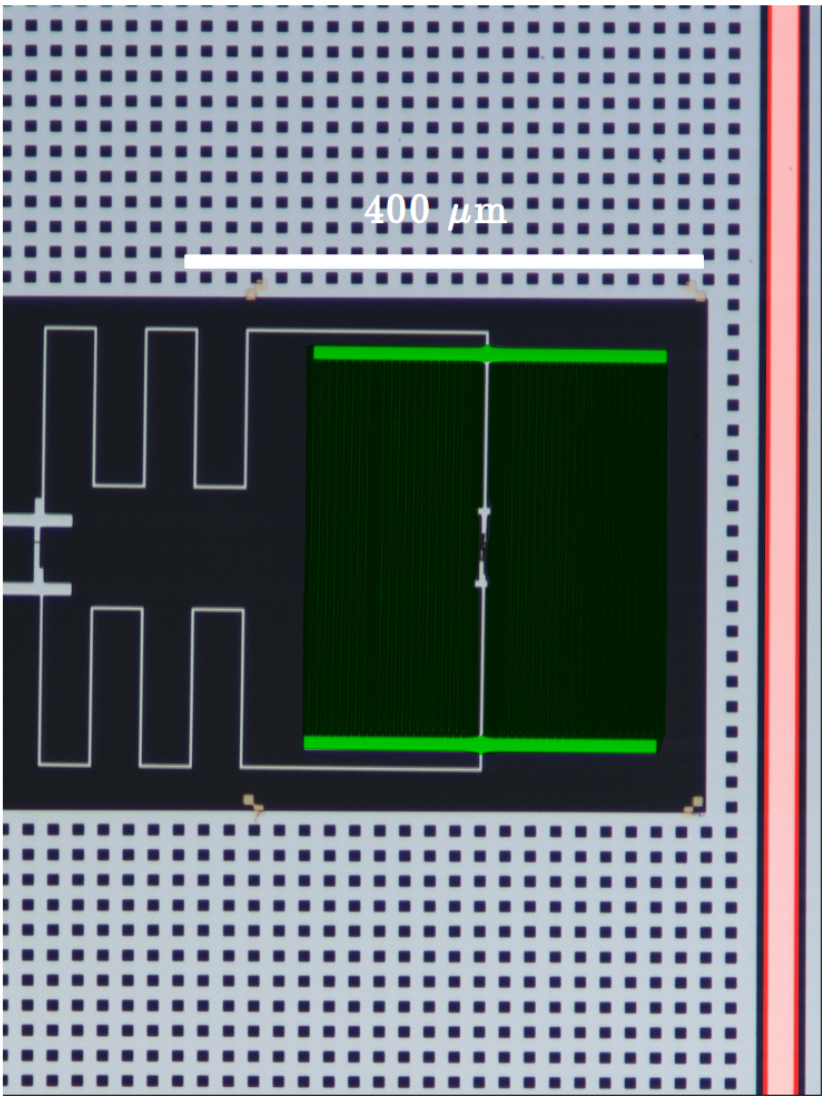

Figure 6.2: Photo of the qubit region from a test fabrication of the BBC2 design. Visible features include the meandering qubit inductor L_q and the IDC C_c (colored green in the image). The LC filter and contact pads are to the left of the image. The CPW center conductor is colored red.

When β_L is large compared to unity, then we expect the coupling to be relatively unimportant and the independent junction picture is valid. With a qubit junction critical current of 1 μA and a loop inductance $L \approx 2$ nH, this leads to a ratio of 2π. This estimate is for the unbiased junction; at higher bias currents the coupling becomes more significant. However, since the dependence is $\sqrt{(1-(I_b/I_0)^2)}$, even at a bias of 87 percent of the critical current $L/L_{qj} = \pi$. Thus the independent junction approximation seems potentially reasonable for these parameters, although we must still examine the impact of the normal modes model on the lifetime.

As discussed in Section 2.4.3, T_1 can be considered the RC time constant of our junction. That model considered a single junction and a single noise source. Here I extend that analysis to include the SQUID and multiple sources of noise. The analysis leading to Eq. (2.67) remains mostly intact. All that changes is I use a more general perturbation operator instead of γ:

$$\Gamma = \frac{2\hbar\omega_{01}}{R} \frac{\left|\langle 0|\,\hat{X}\,|1\rangle\right|^2}{\hbar^2} \frac{e^{\hbar\omega_{01}/kT_q}-1}{e^{\hbar\omega_{01}/kT_n}-1}, \tag{6.2}$$

where $\hat{X}$ is the perturbation operator appropriate for the type of noise we're considering. I will be assuming that $T_q = T_n$ for the various types of noise sources for now. Given that assumption, I separate Γ into two contributions

$$\Gamma = S_{II}(\omega_{01})\kappa_X, \tag{6.3}$$

where

$$S_{II} = \frac{2\hbar\omega_{01}}{R_X}; \tag{6.4}$$

$$\kappa_X = \frac{\left|\langle 0|\,\hat{X}\,|1\rangle\right|^2}{\hbar^2}. \tag{6.5}$$

This partition is helpful for physical intuition, as it distinguishes the contributions from the relevant matrix elements from those arising from filtering effects that result in less current (or flux) noise reaching the SQUID in the first place.

Further simplification results from considering the types of operators $\hat{X}$ I am interested in. From the normal mode analysis, these operators will be some combination of $\hat{\gamma}_i$ and $\hat{\gamma}_q$. All I care about are the qubit mode components of these operators, which have matrix elements involving the harmonic oscillator rescaling term $\sqrt{\hbar/2\omega_{01}}$. The missing mass in this term is buried in the basis transform to the normal coordinates given by A_{iQ} and A_{qQ}, as mentioned in the discussion following Eq. (2.107). The harmonic oscillator rescaling cancels the frequency dependence in S_{II}; the mass rescaling remaining in κ_X corresponds to an effective capacitance. Thus filtering extends T_1 by giving a larger R_X; smaller matrix elements correspond to a larger effective C_X.

Each noise source has its own characteristic dissipation rate. I assume the noise sources are uncorrelated; this results in a total decay rate that is simply the sum of the individual decay rates. While there are many plausible dissipation channels, the ones I examined are dielectric loss in the junctions, radiation through the various leads, and 1/f flux noise.

6.2.1 Dielectric loss

Dielectric loss in the junctions is expected to be one of the major sources of loss. In the original model for this [76], the mechanism of loss was posited to be individual two-level defects in the junction barrier. By making a smaller junction, fewer defects would remain in the junction and the loss would be less significant. Martinis et al. reported a loss tangent[1] of 1.6×10^{-3} for AlOx[2] tunnel barriers. Since individual

[1]The loss tangent is the ratio of the resistive part of the impedance to the reactive part; greater loss tangent means more loss.

[2]Since the exact stoichiometry is unknown, AlOx is a customary notation for the junction barrier oxide.

discrete fluctuators were visible in the qubit spectrum, it is an interesting question whether this loss tangent can be extrapolated to junctions sufficiently small so that individual fluctuators are not visible in the spectroscopy. However, by assuming that it can, I make a worst-case estimate of the loss associated with the junction dielectric.

Making the junction smaller does reduce or eliminate defects, decreasing dissipation (and thus *increasing* the shunt resistance R) linearly with junction area A. However, since this also (neglecting edge effects) shrinks the junction capacitance by a factor of $1/A$, we do not expect a gain in T_1. The only way to win is to use a less lossy capacitor dielectric. One approach to this problem involves better junction barriers, epitaxially grown [124–126]. Some groups have achieved considerable success this way. However, this style of device fabrication requires more involved methods, equipment and expertise.

The approach we took followed Steffen et al. [78] by employing a low-loss shunt capacitance. To achieve very low loss, we used an interdigitated capacitor (IDC) design instead of a silicon nitride parallel plate shunt capacitor. Lumped element resonator measurements using identical materials performed at LPS [32] suggested we would achieve a loss tangent of 5×10^{-6}, almost three orders of magnitude better than amorphous AlOx. By making the IDC large relative to the junction capacitance, we can improve the overall dissipation considerably.

The loss tangent of a capacitor is

$$\tan\delta \equiv \frac{\mathrm{Re}Y}{\mathrm{Im}Y} = \frac{1}{\omega RC}. \tag{6.6}$$

With two capacitors C_c and C_j in parallel, I can get the effective loss tangent from (6.6) using the parallel combination of the Rs and Cs. The combined resistances

give

$$\begin{aligned} \frac{1}{R_C} &= \frac{1}{R_{Cj}} + \frac{1}{R_{Cc}}; \\ &= \omega \left(C_j \tan \delta_j + C_c \tan \delta_c \right). \end{aligned} \tag{6.7}$$

Since the combined capacitance is just $C_j + C_c$, dividing both sides by $\omega(C_j + C_c)$ yields the loss tangent:

$$\tan \delta_e = \frac{C_j}{C_j + C_c} \tan \delta_j + \frac{C_c}{C_j + C_c} \tan \delta_c. \tag{6.8}$$

Thus, the loss tangent for the combined junction-IDC system is an average of the loss tangent for the two separate capacitors, weighted by the capacitances C_c and C_j. This calculation assumes there is insignificant parasitic inductance in the IDC, as otherwise the admittances add in a more complicated fashion.

For a single junction, this RC picture is quite useful. For multiple junctions, it is less clear to me how to account for this type of loss. In principle, it would seem that any resistive noise associated with the isolation junction is transferred through the loop inductance L to the qubit junction. I ignore effects of this kind; I expect that noise generated from a lossy capacitor has a much more significant impact on the junction directly shunted by that capacitor. So, I will consider the local loss tangents to directly give an RC timescale a la the single junction worked out in Section 2.4.3. The isolation junction still contributes to the loss in this picture, since in the normal modes theory the normal coordinate of the qubit mode has a γ_i component.

The qubit junction was designed to have an area of 0.175 μm^2. For the sought critical current density, a reasonable estimate for the specific capacitance of the junction is 50 fF / μm^2 [127]. Thus I expect C_{qj} = 7-10 fF. The IDC value can be

estimated from a microstrip formula derived using conformal transformations

$$C_c = \frac{\epsilon \times 10^{-3}}{18\pi} \frac{K(k)}{K'(k)} (n-1)L \text{ pF}. \tag{6.9}$$

The formula is from [128], which explains the elliptic integral terms like $K(k)$ which depend on the aspect ratio of the fingers in a complicated way. More simply, C_c also depends linearly on the number of interdigitated fingers and the length of those fingers. The effective dielectric constant

$$\epsilon \approx \frac{\epsilon_d + 1}{2} + \frac{\epsilon_d - 1}{2} \frac{1}{\sqrt{1 + 12h/W}} \tag{6.10}$$

depends on the ratio of the height h of the capacitor fingers above the ground plane to the width W of those fingers. The closest conducting surface under the strips is the sample box, thus $h \geq 430$ μm, while the fingers themselves are 0.7 μm wide. Thus the effective dielectric constant is just an average of the constants for vacuum and sapphire = 10.8.

The design capacitance $C_c \approx 1pF$, using around 90 fingers 350 μm long. To get a low charging energy and the best effective loss tangent possible, the goal was to make C_c as big as feasible. As seen in Figure 6.2, the IDC already takes up a great deal of chip real estate. Equation 6.9 suggests that more fingers would increase C_c. However, the undercut needed for double angle evaporation (see Section 4.1 limits how tightly the fingers can be packed. If too much resist is cut away underneath the resist spacers between fingers, the IDC pattern gets washed away in development. For BBC2, I settled on fingers 0.7 μm wide separated by 2.3 μm for reliability. Combining the design capacitances C_c and C_qj with their expected loss tangents gives the effective loss tangent $\tan\delta_c + \tan\delta_j/100 = 5\times10^{-6} + 1.6\times10^{-5} = 2.1\times10^{-5}$. Then

$$R_{dq} = \frac{1}{\omega_0 (C_j + C_c) \tan\delta_e}. \tag{6.11}$$

Next is the filter capacitor. This is made with silicon nitride, expected to have a loss tangent one to two orders of magnitude better than junction barriers [31]. In addition, the matrix element entering into T_1 calculations will be small. From (3.10), the resonator component of the dispersive limit Jaynes-Cummings wavefunction is proportional to (g/Δ), a quantity I show is quite small in BBC2 in Chapter 7. Thus the matrix element $\langle 0|\,\hat{\gamma}_f\,|1\rangle$ will also have this g/Δ weighting, which is then squared in calculating the transition probability needed for Γ. I expect loss to the filter capacitor dielectric to be unimportant and I do not consider it further.

The isolation junction loss is also suppressed, but not as heavily as the filter. From Equation (2.103) we see that when $C_i \ll C_q$ the weight of $\langle 0|\,\gamma_i\,|1\rangle$ is a factor of approximately $L_{ij}/L_{ij}+L$ less than $\langle 0|\,\hat{\gamma}_q\,|1\rangle$[3]. Thus the decay rate associated with dielectric loss to the isolation junction goes as $(L_{ij}/L_{ij}+L)^2$. With no applied bias, this contribution is quite small. However, as greater bias is applied, L_{ij} increases. Since the loss tangent for the big junction is just the bare loss tangent for a junction, at high isolation bias the loss through the isolation junction can become significant.

6.2.2 Radiation to leads

The other source of loss considered in the original design was through radiation to the various leads. There are three lines to consider: the current bias line, the dc flux coil, and the coplanar microwave line. The Fermi's golden rule analysis applied to the bias line in Section 2.4.3 applies to the other lines as well. The most important point in the analysis is the consideration of the current noise spectral density S_{II} presented by each line. Assuming 50 Ω noise sources, this is just a measure of the coupling to the various lines. The other consideration is the appropriate matrix element for each line. This can be determined from the Hamiltonian. The original design document for BBC2 used the dc SQUID Hamiltonian (2.48), which I use here

[3]Numerically I do not make this approximation and instead weight the escape rate by $(A_{iQ}/A_{qQ})^2$ explicitly.

as well.

- For the bias line, the impedance transform arguments of Section 3.3.1 lead to an increase in R_b from the LC filter (3.17):

$$R_b \approx \left(\frac{\omega_Q}{\omega_f}\right)^4 \times 50\Omega. \tag{6.12}$$

Earlier estimates of T_1 in the dc SQUID phase qubit considered the current division coming from the SQUID as well. However, by using the full Hamiltonian I can handle this more carefully by looking at the relevant matrix element arising from small fluctuations in the bias current. The bias current enters in one term in (2.48):

$$-\frac{I_b\Phi_0}{2\pi}\left(\frac{L_q}{L}\gamma_i + \frac{L_i}{L}\gamma_q\right). \tag{6.13}$$

The operator that I_b weights is therefore the desired perturbation operator for the bias line $\hat{X}_b$:

$$\hat{X}_b = -\frac{\Phi_0}{2\pi}\left(\frac{L_q}{L}\hat{\gamma}_i + \frac{L_i}{L}\hat{\gamma}_q\right). \tag{6.14}$$

Replacing the γs with the appropriate normal mode ladder operators, including the rescaling factor $\sqrt{\hbar/2\omega_Q}$, makes the matrix element of $\hat{X}_b$ easy to compute:

$$\langle 0|\,\hat{X}_b\,|1\rangle = -\sqrt{\frac{\hbar}{2\omega_Q}}\frac{\Phi_0}{2\pi}\left(\frac{L_q}{L}A_{iQ} + \frac{L_i}{L}A_{qQ}\right). \tag{6.15}$$

Taking the absolute square of, I get:

$$\begin{aligned}\kappa_b &\equiv \frac{\left|\langle 0|\,\hat{X}_b\,|1\rangle\right|^2}{\hbar^2}, \\ &= \frac{1}{2\hbar\omega_Q}\left(\frac{\Phi_0}{2\pi}\right)^2\left(\frac{L_q}{L}A_{iQ} + \frac{L_i}{L}A_{qQ}\right)^2.\end{aligned} \tag{6.16}$$

- The flux line couples to the SQUID by introducing flux in the SQUID loop.

The external flux $\Phi_e = M_f I_f$, so that $S_{ff} = M_f^2 S_{II}$. The relevant term from (2.48) is

$$\left(\frac{\Phi_0}{2\pi}\right)^2 \frac{\left(\gamma_i - \gamma_q - (2\pi\Phi_e/\Phi_0)\right)^2}{2L}. \tag{6.17}$$

The term quadratic in Φ_e yields no matrix element between 0 and 1, so we're left with the operator

$$-\frac{\Phi_0}{2\pi}\frac{\Phi_e}{L}\left(\hat{\gamma}_i - \hat{\gamma}_q\right) = -\frac{\Phi_0}{2\pi}\frac{M_f I_f}{L}\left(\hat{\gamma}_i - \hat{\gamma}_q\right). \tag{6.18}$$

To compare to the bias line, I keep the inductance terms with the current. When I replace that with the equivalent expression in terms of S_{II}, I get a coupling (similar to the filter factor from the bias line) of $(M_f/L)^2$. This scaling is very intuitive. It comes from a current formed by the inductance L to cancel out the flux induced by the flux line inductance M. While this type of cancellation is not always possible without exceeding the critical currents of one of the junctions and hopping to another flux well of the SQUID, for the small signals associated with these noise calculations this analysis is valid. Thus

$$R_f = \left(\frac{L}{M_f}\right)^2 \times 50\Omega. \tag{6.19}$$

$\hat{X}_f$ is given by

$$\hat{X}_f = \frac{\Phi_0}{2\pi}\left(\hat{\gamma}_i - \hat{\gamma}_q\right). \tag{6.20}$$

As in (6.2.2), the matrix element is given by replacing γ's with appropriate normal mode weights:

$$\langle 0|\, \hat{X}_b\, |1\rangle = \sqrt{\frac{\hbar}{2\omega_Q}}\frac{\Phi_0}{2\pi}\left(A_{iQ} - A_{qQ}\right); \tag{6.21}$$

this matrix element gives

$$\begin{aligned}\kappa_f &\equiv \frac{\left|\langle 0|\,\hat{X}_b\,|1\rangle\right|^2}{\hbar^2} \\ &= \frac{1}{2\hbar\omega_Q}\left(\frac{\Phi_0}{2\pi}\right)^2 (A_{iQ} - A_{qQ})^2 \,.\end{aligned} \tag{6.22}$$

Unlike the bias line, the noise affects the qubit and bias junctions equally. This makes sense physically, as the circulating current induced in the loop by flux noise gives the same current in both junctions.

- The microwave line presents some challenges because the coupling mechanism is not obvious. In the design process, the behavior at relevant microwave frequencies was modeled using Microwave Office, a 2.5D field solver. Replacing the junctions in the SQUID with inductances gives the SQUID a harmonic resonance at the anticipated qubit frequency. The behavior of this resonance in response to fields on the microwave line can be modeled and the total quality factor Q of the resonance determined. The internal quality factor Q_i can be controlled in the model, leaving the external quality factor Q_e from loss to the microwave line as the unknown determining the total quality factor. This inferred Q_e can be used to get an estimate to the lifetime T_1 that would arise if all loss was through the microwave line.

  While this is sufficient for design purposes, it does not make clear whether the coupling is inductive, capacitive, or a mixture of both. Given the layout — where the microwave line runs parallel to the SQUID loop, separated from it by a ground plane — an inductive coupling seems physically reasonable. The inductance of the microwave line to the sample can be modeled and easily measured for dc currents. If we assume inductive coupling, κ_{flux} applies for this line as well; the only variation from the flux line is that S_{II} is now rescaled

by $(M_\mu/L)^2$ instead of using the flux line mutual inductance. This is how I approximate the dissipation from the microwave line.

6.2.3 1/f noise

Although it was not considered in the original design document, it is important to consider the effect of 1/f noise in flux on dissipation. 1/f noise occurs in many condensed matter systems, and while there are general arguments that suggest how it might appear [129, 130], the specific cause is still not understood. In SQUIDs this behavior has been observed for almost 30 years [131], and with the advent of quantum computing research has become of interest once more. The prevailing theories at the moment suggest that 1/f noise is coming from surface behavior, such as local polarized spin fluctuations [132].

The low-frequency nature of 1/f noise makes it particularly important for considerations of dephasing [133]. It is less clear if it is important for relaxation, which concerns the presence of noise at the qubit frequency. Typical rms values of 1/f flux noise , extracted from measured spectral power densities, are of the order of 1-10 $\mu\Phi_0/\sqrt{\text{Hz}}$ at 1 Hz. One of the findings in [131] was that this number is weakly sensitive to most aspects of the SQUID. Some more recent work [134] suggests that having a large inductance with small wires results in larger flux noise. This is consistent with measurements I made of SQUIDs of similar size to BBC2. Thus, I will assume 1/f flux noise at the higher end of this range.

I can compare this level of flux noise to the spectral density of flux noise from the flux line, assuming that the noise on the flux line is Johnson noise from a 50 Ω resistor at the fridge base temperature of 20 mK. (In practice, the noise on the flux line is likely to be substantially warmer than this, suggesting this calculation overestimates the importance of 1/f noise.) Picking some representative numbers,

Table 6.1: Design values of significant BBC2 parameters. Other than C_{jq}, which I have given a larger estimate for, these are taken from the original design document [123].

Parameter	Value	Units
qubit area	.1	μm^2
isolation area	2.5	μm^2
I_{0q}	.4	μA
I_{0i}	7	μA
L_i	.02	nH
L_q	2	nH
L_f	16	nH
M_f	1	pH
M_μ	2	pH
C_{ji}	50	fF
C_{jq}	10	fF
C_{idc}	1	pF
C_f	50	pF

in a 1 Hz bandwidth at a frequency of 8 GHz I get:

$$\begin{aligned} \frac{1}{f}(10\mu\Phi_0)^2 &< \frac{4k_BT}{R}M_\phi^2, \\ \frac{1}{8\times10^9}(2.06\times10^{-20})^2 &< \frac{4(1.3\times10^{-23})(2\times10^{-2})}{50}(5\times10^{-12})^2, \\ 5.3\times10^{-50} &< 5.5\times10^{-49}. \end{aligned} \tag{6.23}$$

Thus, perhaps surprisingly, with small variations in the numbers it is possible that flux noise could be a limiting factor to T_1. A caveat to this analysis is that recent measurements [135] suggest the $1/f$ flux noise may cut off over 1 GHz.

6.2.4 T_1: putting it all together

Table 6.1 contains design values for the various parameters of BBC2. These were chosen to achieve a design $T_1 > 1$ μs. Rather than specific numbers, guidelines were used to choose certain parameters. To minimize dielectric loss we wanted to

make the qubit junction very small. We decided we could fabricate a junction with area of 0.1-0.2 μm^2 reliably. The critical current estimate I_{0q} then comes from the typical critical current densities we achieved using our double angle evaporation fabrication process. L_q was chosen to use a loop size consistent with previous dc SQUID phase qubits like PB9; given L_q, the area of I_{0i} was chosen so β_L (see Equation 6.1) would be large. L_i was chosen as small as possible to maximize the isolating effect of the SQUID. C_c was made as large as possible to minimize the effective loss tangent of the qubit junction. The mutual inductances M_f and M_μ were chosen to be the smallest values that seemed practical.

My design value estimates include a somewhat larger value of C_{jq} than Rangga originally considered, corresponding to estimates taken from [127] for the likely capacitance given our critical current density. Since most of the capacitance is dominated by the IDC anyway, this has little effect on the dynamics of the system. But, compared to the original estimate of 3 fF, it does lead to a substantially higher effective $\tan\delta_e$. This accounts for most of the difference in the T_1 predictions of that document and this thesis.

Budoyo's discussion in [123] is modeled after Przybysz's treatment in [28]. Both use the idea of characterizing all of the noise sources by the effective R they provide when considered as Johnson noise sources. The total T_1 is then obtained by adding all of the resistances in parallel. This assumes that there is a single correct capacitance to use for getting T_1. The arguments involving κ presented above suggest this is not quite right. Instead, I find the loss rate Γ associated with each process and use the uncorrelated noise assumption to give the total loss rate as the sum of all of the individual loss rates. T_1 is then just the inverse of $\Gamma_{\rm tot}$.

Table 6.2 shows the results of the calculations performed in this section given the design parameters for BBC2. The exact numbers depend on the bias conditions, with the frequencies and transformation weights A_{iQ} and A_{qQ} calculated numeri-

Table 6.2: Estimated T_1 from the considered noise channels, given the original design parameters of Table 6.1.

Source of dissipation	T_1 (μ s)
Qubit dielectric	1.4
Isolation dielectric	10
Bias line	1.7×10^4
Flux line	220
Microwave line	56
$1/f$ flux noise	64
Total T_1	1.2

cally. The values in the table were calculated with junction phases corresponding to $I_i/I_{0i} = .85, I_q/I_{0q} = .9$, yielding $f_Q \approx 5.5$ GHz. These choices are arbitrary, but represent substantial bias in both junctions, which was typically the way these qubits were operated.

Of the noise sources, dielectric losses depend the most on the qubit frequency. This is because $\tan\delta$ sets a Q for the qubit. This means that if dielectric loss dominates, we expect the oscillation to last for a fixed number of cycles instead of a fixed time. At higher qubit frequency f_Q, the same number of cycles occurs quicker and thus the decay time T_1 is shorter.

Probably the biggest questions regarding T_1 are whether this list of noise sources is sufficiently comprehensive and whether the dielectric loss in the small junction is as significant as I assumed. The former concern would suggest this list is too optimistic, the latter too pessimistic. Table 6.2 suggests dielectric loss in the qubit junction is the chief limiting factor. Perhaps surprisingly, the second most significant noise source is loss in the isolation junction, despite the low weight of γ_i in the qubit normal mode.

6.3 Initial characterization

Given this analysis, it seemed that BBC2 should be capable of hitting the goal of 1 μs T_1. Finding out whether we succeeded in this was the initial goal of the experimental measurements; I discuss the T_1 measurements in Chapter 7. Before starting on those microwave measurements, I performed dc measurements of the sample to get in-situ values for several of the important device parameters.

6.3.1 I-V diagram

While almost all of my measurements use the switching readout scheme discussed in Chapter 5, the exception is the initial I-V curve. In practice, I did not use the I-V curve to extract much information about the sample. However, it is a useful first measurement because the supercurrent branch is visible even at very low current. Switching based measurements can't be seen until I_b approaches the critical current of the SQUID. If this critical current is unknown, as with a new device, starting with switching based characterization creates the possibility of blowing out the junctions by sourcing too much current. Thus, I started measurements with an I-V curve for safety reasons.

To perform the I-V measurement, I normally did a four wire measurement using the bias and voltage lines. A sinusoidal current waveform is sourced using the bias function generator, typically at a frequency between 1 and 100 Hz. To measure the voltage, I attach a Stanford Research Systems amplifier to the voltage line. Since the custom amplifier chain is ac-coupled and we are interested in looking at I-V curves with low-frequency current drive, the normal amplifier setup is unhelpful here. The current sourced is measured by monitoring the resistance across the resistor following the unity gain buffer on the bias line. Since this resistor is usually well over 99% of the resistance on the bias line, I often just measure the voltage on

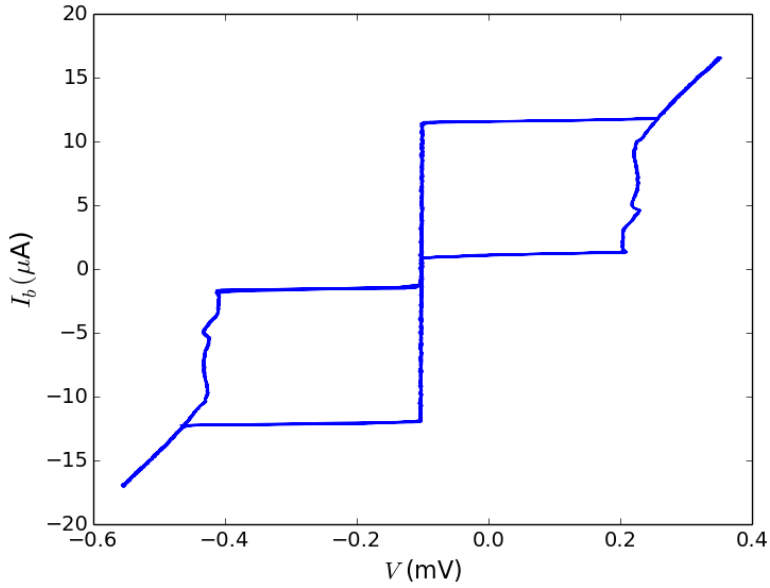

Figure 6.3: I-V curve from the BBC device.

the input of the resistor and use it as the voltage difference.

The current and voltage measurements are then sent to two channels on a battery powered Tektronix oscilloscope. The oscilloscope takes some number of samples (at higher source frequencies, 256; at lower source frequencies, usually fewer) and averages them. I then stop the sampling and save the resulting waveforms.

Figure 6.3 shows the I-V diagram for BBC, the immediate precursor to the BBC2 device. (BBC had a qubit junction that proved too small and flux coupling that was too weak, leading to very weak flux modulation.) The I-V looks very similar to the cartoon shown in Figure 2.4. There is still a supercurrent branch with zero voltage and a resistive branch in either direction. The horizontal lines connecting the supercurrent branch and the resistive branches show the critical current of the SQUID and the retrapping current, where the voltage relaxes back to zero following a switch. The ratio of the resistive branch height to the supercurrent branch height is smaller than expected. There are also some curious distortions of the resistive branches at low currents. The origin of these is likely heating of the small junctions.

There is additional information that can be gleaned from the I-V diagram. There appears to be some nonzero offset voltage on the voltage line typically. We believe this is a thermoelectric effect associated with the joining of dissimilar metals in the coaxial cables and connectors. By looking at the corners of the branches (i.e. at the effective critical current, or at the retrapping currents on the voltage branch) we could make some inferences about the noise in our system. In practice, I have not explored this much, since coherent operation of the qubit gives an even better measure of the noise.

6.3.2 I-Φ diagram: current-flux characteristic

The defining feature of SQUIDs that gives them their name is the quantum interference displayed as the amount of flux in the loop is changed. From the dc SQUID equations of motion I derived the condition (2.42):

$$\phi_i - \phi_q = \Phi_e - L_i I_{0i} \sin \frac{2\pi\phi_i}{\Phi_0} + L_q I_{0q} \sin \frac{2\pi\phi_q}{\Phi_0}. \tag{6.24}$$

Converting from junction fluxes to junction phases by multiplying through by $2\pi/\Phi_0$,

$$\gamma_i - \gamma_q = \frac{2\pi}{\Phi_0} \left(\Phi_e - L_i I_{0i} \sin\gamma_i + L_q I_{0q} \sin\gamma_q\right). \tag{6.25}$$

This is known as the flux-phase condition for the dc SQUID.

Physically, the requirement of fluxoid quantization in the loop constrains the phase difference between the two junctions, influencing how currents can divide between them. By changing the amount of flux in the loop, we change the effective critical current of the SQUID. Notice that by adding an extra flux quantum of external flux we just shift the phase differences by 2π and the equation is still satisfied. Thus, we expect the modulation of the SQUID currents by external flux to occur with a periodicity of the flux quantum.

This suggests that by measuring the critical current of the SQUID for many fluxes, we can use the periodicity of the measurement to determine how much current yields a particular flux at the SQUID, thereby establishing the mutual inductance of various lines. In fact, this measurement of the current-flux characteristic — henceforth referred to as the I-Φ measurement — can tell us quite a bit more. It allows us to infer the critical currents of both junctions, as well as the parameters L_i and L_q. Thus, the I-Φ measurement produces many of the key parameters necessary to characterize the SQUID.

I will show how this works by considering how the current constraint derived in (2.41)

$$I_b = I_{0i} \sin \gamma_i + I_{0q} \sin \gamma_q \tag{6.26}$$

and the flux-phase condition (6.25) interact. Specifically, suppoes that for $\Phi_e = \Phi$ the critical current is $I_c = I_b$. The question is how changes in these quantities are related. A small change in the critical current leads to small changes in the junction phases:

$$\begin{aligned} I_c + \delta I_c &= I_{0i}(\sin \gamma_i + \delta\gamma_i \cos \gamma_i) + I_{0q}(\sin \gamma_q + \delta\gamma_q \cos \gamma_q); \\ \delta I_c &= I_{0i} \cos \gamma_i \delta\gamma_i + I_{0q} \cos \gamma_q, \end{aligned} \tag{6.27}$$

where I move from the first to the second step by subtracting the known previous solution. Similarly,

$$\delta\Phi = \frac{\Phi_0}{2\pi} \left(\delta\gamma_i - \delta\gamma_q + \frac{2\pi L_i I_{0i}}{\Phi_0} \cos \gamma_i \delta\gamma_i - \frac{2\pi L_q I_{0q}}{\Phi_0} \cos \gamma_q \delta\gamma_q \right). \tag{6.28}$$

I can combine these two results using the chain rule:

$$\begin{aligned}\frac{\delta I_c}{\delta\Phi} &= \frac{\delta I_c}{\delta\gamma_i}\frac{\delta\gamma_i}{\delta\Phi} + \frac{\delta I_c}{\delta\gamma_q}\frac{\delta\gamma_q}{\delta\Phi},\\ &= \frac{2\pi\alpha_i\cos\gamma_i}{L_i} * \left(\frac{\Phi_0}{2\pi}\right)\frac{1}{1+\alpha_i\cos\gamma_i} - \frac{2\pi\alpha_q\cos\gamma_q}{L_q} * \left(\frac{\Phi_0}{2\pi}\right)\frac{1}{1+\alpha_q\cos\gamma_q},\\ &= \frac{\alpha_i\cos\gamma_i}{L_i(1+\alpha_i\cos\gamma_i)} - \frac{\alpha_q\cos\gamma_q}{L_q(1+\alpha_q\cos\gamma_q)},\end{aligned} \tag{6.29}$$

where

$$\alpha \equiv \frac{2\pi I_{0n} L_n}{\Phi_0} \tag{6.30}$$

is a coefficient representing how much flux is associated with phase changes in a given junction.

To understand (6.29), it is helpful to consider the case where one junction is heavily biased and the other is not. This amounts to saying that when the SQUID switches, it is largely because the critical current of one junction was exceeded. With this assumption, the junction responsible for the switching will have a phase very close to $\pi/2$, leading to the vanishing of the relevant cos term. Then we need only consider a single term of (6.29). For example, when the isolation junction is near its critical current and $\gamma_i \approx \pi/2$,

$$\frac{\delta I_c}{\delta\Phi} = -\frac{\alpha_q\cos\gamma_q}{L_q(1+\alpha_q\cos\gamma_q)}. \tag{6.31}$$

When γ_q is weakly biased, the cos changes very little with small changes in phase. Thus for large stretches, the I-Φ measurement will yield critical currents that vary linearly with changes in flux, with the slope being related to the inductance of one arm of the SQUID. Because of the association with a particular junction, I will refer to the two possibilities as the isolation branch and the qubit branch. The difference in size between L_i and L_q leads to marked difference in the steepness of

the two branches; (6.29) also shows that we expect the branches to have slopes of opposite sign. Given this argument about what the branches mean, we expect the branches join at the place where both junctions are near switching. The critical current at this joint is just the maximum critical current $I_{0i} + I_{0q}$ that the SQUID can support.

This picture is useful for building an intuitive feel for the I-Φ plot. For performing fits to measured I-Φ characteristics, a more exact calculation is employed. The critical current for a given flux represents the maximum bias current that can run through the SQUID without switching. Since the bias current is just the sum of the currents through the two junctions, we need only maximize that. The trick is to use a Lagrange multiplier to incorporate the flux-phase condition; this was first done in a paper by Tsang and van Duzer [136]. Denoting the Lagrange multiplier λ, what we want to maximize is:

$$I_{0i}\sin\gamma_i + I_{0q}\sin\gamma_q + \lambda\left(\gamma_q - \gamma_i + \frac{2\pi\Phi_e}{\Phi_0} + \alpha_q\sin\gamma_q - \alpha_i sin\gamma_i\right). \tag{6.32}$$

.

Differentiating by the three variables yields

$$\frac{\partial I_c}{\partial\gamma_i} = I_{0i}\cos\gamma_i^* - \lambda\left(1 + \alpha_i\cos\gamma_i^*\right); \tag{6.33}$$

$$\frac{\partial I_c}{\partial\gamma_q} = I_{0q}\cos\gamma_q^* + \lambda\left(1 + \alpha_q\cos\gamma_q^*\right); \tag{6.34}$$

$$\frac{\partial I_c}{\partial\lambda} = \gamma_q - \gamma_i + \frac{2\pi\Phi_e}{\Phi_0} + \alpha_q\sin\gamma_q - \alpha_i\sin\gamma_i. \tag{6.35}$$

Equating these to zero, I can use the first two equations to eliminate two of the

variables:

$$\cos\gamma_i^* = \frac{\lambda}{I_{0i} - \lambda\alpha_i}; \tag{6.36}$$

$$\lambda = -\frac{I_{0q}\cos\gamma_q^*}{1 + \alpha_q\cos\gamma_q^*}. \tag{6.37}$$

This yields an expression for the critical current of the SQUID in terms of the phase of a single junction:

$$I_c = I_{0q}\sin\gamma_q^* + I_{0i}\sin\arccos\left(\frac{1}{\frac{I_{0i}}{I_{0q}\cos\gamma_q^*} + \frac{I_{0i}}{I_{0q}}\alpha_q + \alpha_i}\right). \tag{6.38}$$

A similar expression arises if we instead describe everything in terms of γ_i. As $\gamma_q^* \to \pi/2$, the arccos term also goes to $\pi/2$.

The result of this is that the extrema of I_c correspond to choices where both junctions carry maximal current. If the currents are in the same direction, this gives $I_{0i} + I_{0q}$. If the currents oppose, then $I_{0i} - I_{0q}$ is the result. Thus, by looking at the highest and lowest currents where switches are visible, I extract both of these quantities and thus the critical current of both junctions.

Figure 6.4 is the result of the simplest approach I took to measuring I-Φ characteristics. I used the I-V measurement to pick a total current a few microamps above the critical current of the SQUID. This guaranteed that a switch would occur every time I applied the bias waveform, as discussed in Section 5.3.2. I would then apply a fixed dc flux to the SQUID using the flux line and collect 5000 switching events for that flux.

After doing this for a succession of fluxes, the results are plotted as seen here. The y-axis is I_b, or the voltage across the bias resistor divided by its resistance. (This quantity is actually measured during I-V measurements or current calibrations, but not during qubit measurements.) The x-axis is the voltage applied at the flux

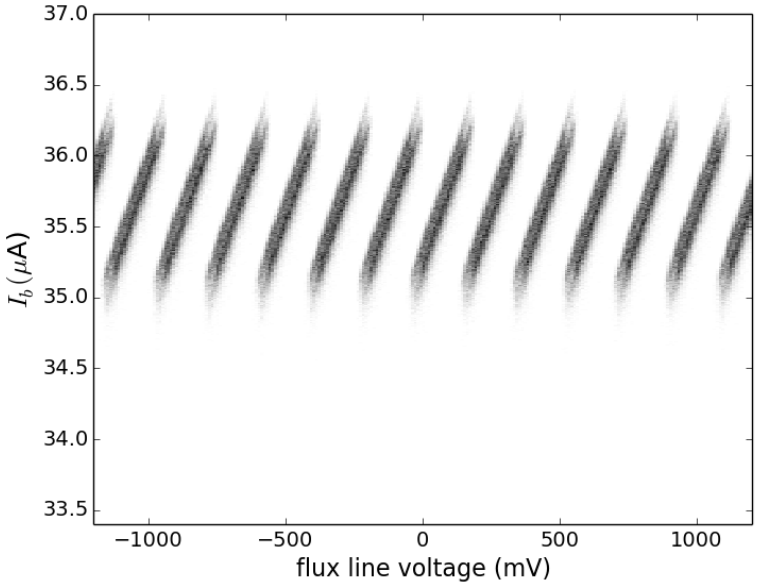

Figure 6.4: I-Φ plot from the initial cooldown of BBC2. In this plot, only I_b is ramped. The true value of the applied flux Φ_e is inferred from this plot.

function generator, which is linearly related to the current sourced and thus the external flux Φ_e. The greyscale colormap corresponds to the number of voltage switches recorded at each setting; darker points correspond to more switches.

The easiest quantity to extract from this plot is the mutual inductance M_f. Since the I-Φ is periodic with a period of one flux quantum, the periodicity of the diagram immediately shows how much current is necessary to yield one flux quantum at the SQUID. The same technique is used to measure M_μ by running a dc current down the microwave line. Somewhat more care is necessary for the microwave measurement, as the presence of attenuators in the microwave line means that sourcing solely down the input microwave line would lead to much of the current running to ground through an attenuator. This would lead to a considerable underestimate of the microwave line inductance, as the current actually reaching the on-chip microwave line would be dramatically overstated. To avoid this shorting problem, for the microwave measurement the BNC connection on the function generator is broken out to connect directly to the two center pins.

A curious feature of Figure 6.4 is that only one branch of the I-Φ characteristic is visible. The previous discussion suggested that there should be a second branch visible. Presumably what is happening is that ramping only the bias current leads to transfer between various flux wells without switching to the voltage state. For a given flux, the flux state with the greatest critical current I_c becomes unstable at some lower bias and dumps into another flux state that I measure instead.

In order to extract the currents from the I-Φ characteristic, it is desirable to see all the possible switches on both branches. At the high end, the two branches of the I-Φ join, which would let me fix the sum of critical currents $I_{0i} + I_{0q}$. To find the second branch and make sure that the joint between the branches is visible, I perform a series of I-Φ measurements where both flux and bias current are ramped.

As illustrated in Figure 5.9, the usual choice of flux ramp corresponds to simultaneous current and flux biases following a trajectory on the I-Φ characteristic parallel to the isolation branch. In the independent switching picture where each branch corresponds to a particular branch switching, this technique amounts to using the current and flux simultaneously to bias one junction while leaving the other unchanged. (This is an oversimplification; one could ramp flux and bias together in the same ratio from zero bias, and the big junction must finish with considerable bias to accomplish switching, so clearly it is picking up some current somewhere. It is probably a good approximation in the narrow window near switching, though.)

For getting I-Φ fits, I used a series of different ramps, summed the switches associated with each ramp, and plotted them all on one master plot. Each ramp corresponded to a different peak-to-peak voltage of the flux waveform. An individual dataset consisted of the switching histograms recorded for a particular ramp height over a series of flux offsets.

Plotting this data is more complicated than the fixed-flux I-Φ. When only the current is changing, only the y-axis depends on time. When both current and

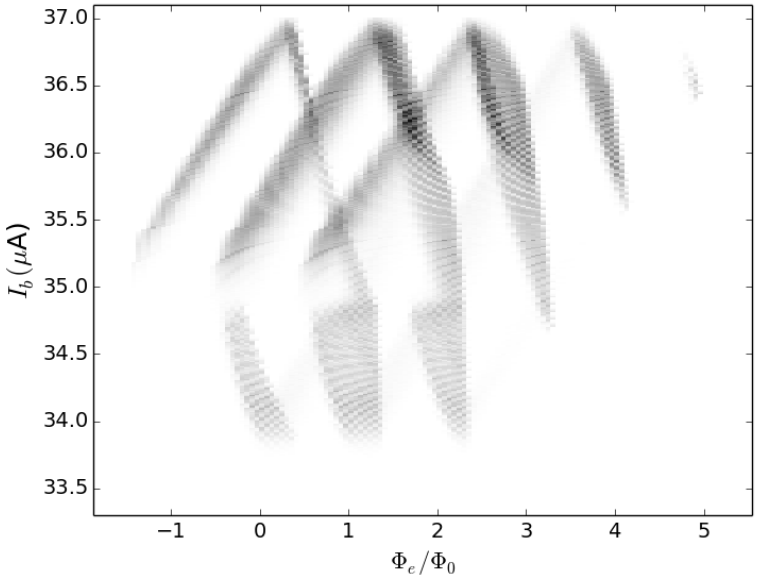

Figure 6.5: Master I-Φ plot from the initial cooldown of BBC2. This plot combines data from many different current-flux ramp ratios onto a single plot.

flux are ramped, the flux gains a time dependence as well, meaning you cannot get accurate results merely by plotting the offsets. The double ramp data is plotted by calculating the flux as a function of time and then rounding to get the appropriate bin. The master I-Φ is plotted in Figure 6.5. The x-axis is the applied flux Φ_e. The y-axis is I_b. As before, darker points indicate more voltage switches.

The master I-Φ plot shows some broadening, particularly for some of the higher flux switches associated with the later ramps. I am unclear on the origin of this broadening, as it is not obvious in the individual component plots (see Figure 6.6 as an example.) It showed up in the master I-Φ plots for each of the three cooldowns in a similar fashion. It might be a numerical artifact where the rounding into different bins affects lines with shallower slopes more strongly.

Figure 6.6 shows a fit of the critical currents predicted by (6.38) to an I-Φ plot. This I-Φ was chosen from of the master I-Φ because it clearly shows the joint of the two branches at high current, giving a useful target for fitting. The fit was performed by calculating (6.38) for the desired fit parameters and visually inspecting

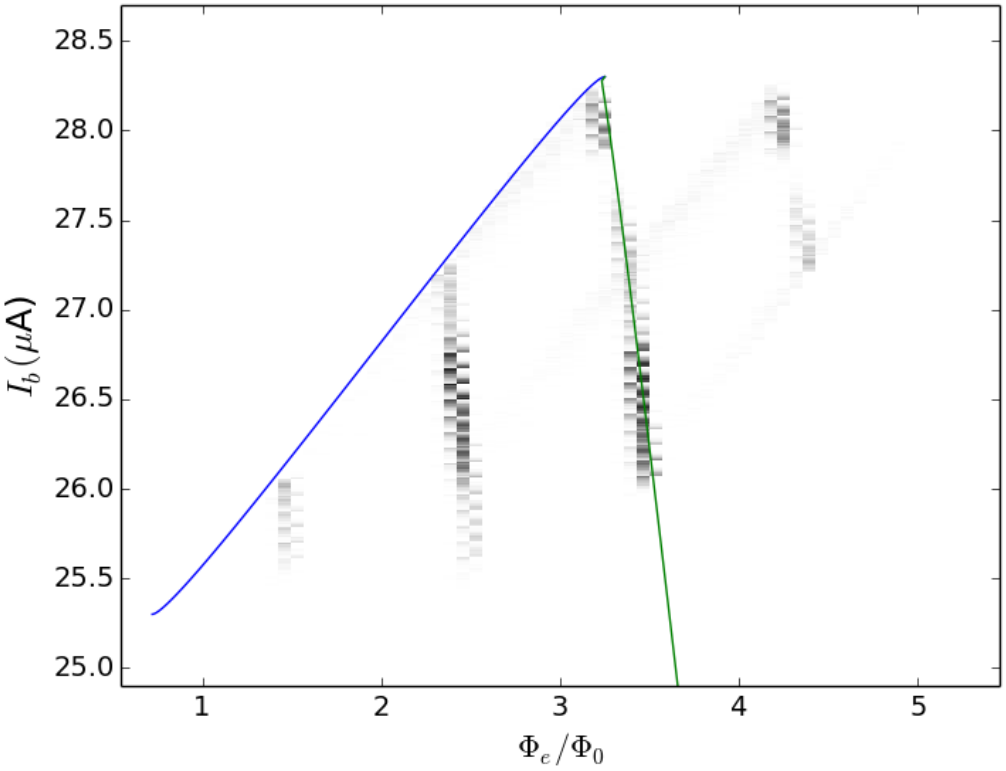

Figure 6.6: Fit of I-Φ plot from the final cooldown of BBC2 to (6.38). This I-Φ used the 1200 mVpp flux ramp dataset from a master I-Φ plot. The fit parameters are $I_{0i} = 26.8$ μA, $I_{0q} = 1.5$ μA, $L_i = 220$ pH, $L_q = 1.4$ nH, $M_f = 15$ nH.

the agreement. The plotted fit parameters are $I_{0i} = 26.8\mu\text{A}, I_{0q} = 1.5\mu\text{A}, L_i = 220$ pH, $L_q = 1.4$ nH, $M_f = 15$ nH.

I will make a few remarks about the fit. The true critical current I_c for a given Φ_e is at least as high as the highest bias current at which switching is measured. The majority of switching events occur before this current as the tunneling rate Γ increases. Thus, on the isolation branch I tried to fit the predicted branches to the highest measured I_b at each Φ_e. On the qubit branch, this was not possible. The discretization of flux into coarse bins is more noticeable, leading to broadening of the qubit branch in I_b. This results in some switches being associated at a higher flux than they ordinarily would, thus landing above the theoretical green qubit branch. The slope of the branch seems reasonable.

One thing that is unclear is the role of mutual inductance from the bias line. In the equations of constraint used to fit the I-Φ, one can include such a mutual inductance by adding it to one L and subtracting it from the other. This effect can always be accommodated simply by changing the L and neglecting M_b altogether, as I have done. However, I do not know if this is true in other applications. Thus it is possible that fitting the inductances in this way may lead to inaccuracies for other applications. Since M_b is likely pretty small, though, and the difference between L_q and L_i substantial, I expect the relative importance of such corrections to be small. Even with such errors, the relative coupling of the junctions is unlikely to be heavily affected.

Potentially more serious is the issue of establishing the current $I_{0i} - I_{0q}$. In the measurements I took, there was no clear signpost that the low current maximum had been reached. To try to estimate where this occurred, I roughly extrapolated the branches to see where they would meet up, using this as a basis for the lower critical current. Taking the smallest visible current from the I-Φ leads to an overestimate of I_{0i} and an underestimate of I_{0q} of the same absolute scale. The relative importance

of this error on I_{0q} can be quite substantial because the qubit junction is so small. It would be desirable to further refine the fitting process to get a more reliable method for finding the bottom of the I-Φ. A suggestion from [137] is to look at the I-Φ curve for the opposite bias sense; apparently the maximum is shifted somewhat in flux, and this gives a good measure of SQUID asymmetry.

6.4 Summary

In this chapter, I discussed the design for the BBC2 sample. I examined the likely sources of dissipation and estimated their contributions to T_1 based on calculations from the full SQUID model described in Chapter 2. I discussed the dc measurements used to do initial characterization of BBC2, particularly measurements of the I-Φ characteristic to extract critical currents and various inductances. This was preparatory work for spectroscopic and coherent measurements on the qubit, which I discuss in the following chapter.

Chapter 7

BBC2 qubit: spectroscopy and lifetime measurements

One of the most obvious requirements for using a qubit is control over its state. In superconducting qubits, this is accomplished by applying microwaves tuned to the resonant frequency f_{01}. To determine this frequency, I performed a series of spectroscopic measurements on the qubit. These measurements examined the long-time response of the qubit to continuous wave microwaves.

Spectroscopy generally serves as a prelude to further investigation of the qubit. Once the transition was identified, to measure the coherence of the qubit I studied Rabi oscillations. These correspond to the time-resolved dynamics of the qubit immediately following microwave turn-on. Further details, like the relative importance of low-frequency noise in dephasing, require more sophisticated pulse sequences. Two common ones are Ramsey spectroscopy and spin echo sequences [81, 119]. I did not pursue these techniques, as the spectroscopy and Rabi data exhibited surprising features that were already quite complex.

I took the data in this chapter from measurements on the BBC2 sample over the course of three different cooldowns. The first cooldown went from April to early June, 2012; the second cooldown from late June to late July; and the third cooldown from late November to early December 2012. Both the first and second cooldowns were interrupted due to unavailability of helium. Between the first and second cooldown, the fridge was opened for several days to complete the rewiring discussed in Section 5.1, then immediately closed and recooled. To reduce aging effects on the junction between the second and third cooldown, the sample was left

in the refrigerator's inner vacuum can, which remained closed and evacuated.

7.1 Initial spectroscopy data

As discussed in Chapter 5, there were two different approaches I used in performing the spectroscopy measurements. The first approach sweeps both current and flux biases of the qubit while applying continuous wave microwaves so that the qubit resonance is tuned through the microwave frequency. This measurement is done for a succession of microwave frequencies, mapping out the frequency of the resonance as the bias changes. I refer to this approach as "swept spectroscopy." The second approach sets the current and flux biases to a fixed value. I would then vary the microwave frequency until I found the qubit resonance. Repeating the measurement with different biases maps the frequency response with changing biases. Since this technique requires a current pulse in order to make the measurement, I will refer to this approach as "pulsed spectroscopy."

On the first cooldown (following the initial work correcting the grounding issues discussed in Section 5.2.3) I performed swept spectroscopy. I chose this method because it requires far less data for an initial exploration of the qubit resonance. Since each individual measurement involves a sweep over biases, the microwave frequency is the only variable that needs to be scanned. Compared to the pulsed measurement, data can be gathered an order of magnitude faster.

Figure 7.1 shows escape rate data from a swept spectroscopy measurement taken on the first cooldown at a frequency of 7450 MHz. One curve shows the escape rate with no microwaves applied; I call this the background escape rate. The other three curves correspond to different applied powers. The x-axis is the current sourced on the bias line I_b; since I used the double ramp discussed in Section 5.3.2, the x-axis is also a proxy for the applied flux. The y-axis is the escape rate, calculated from raw histogram data and plotted on a logarithmic scale. The escape rate Γ is

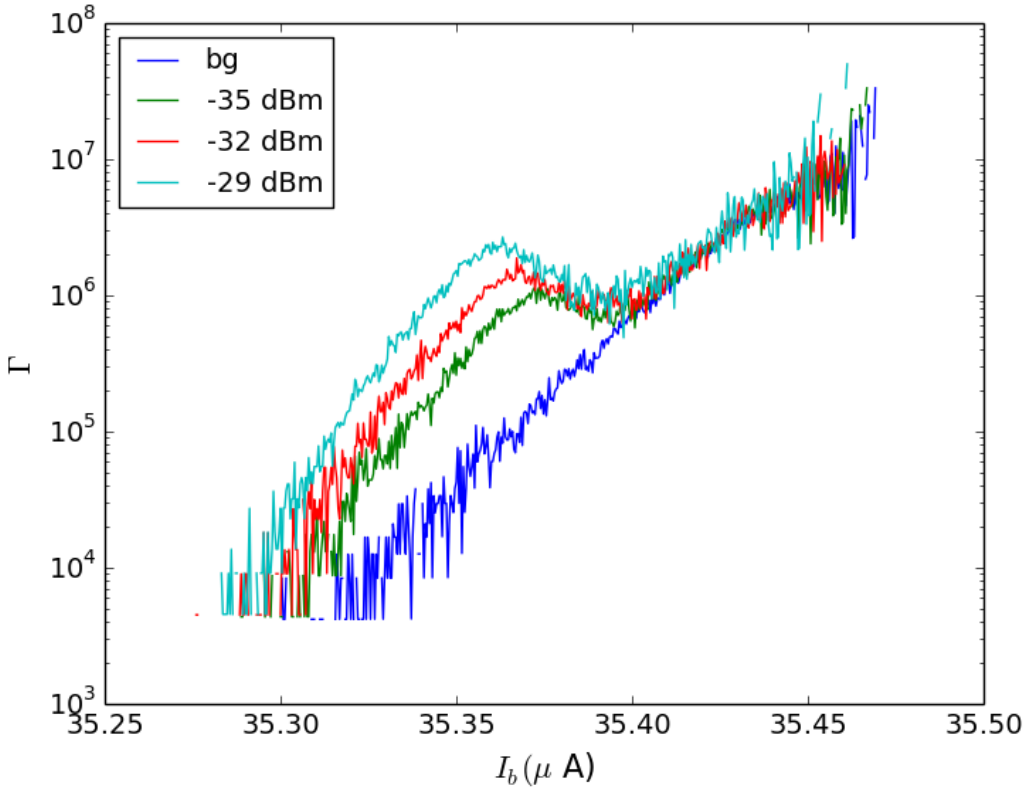

Figure 7.1: Plot of measured qubit escape rates vs. current bias. The background escape rate is plotted, as well as the escape rate with 7450 MHz microwaves at several powers.

defined by

$$\frac{dP(t)}{dt} = -\Gamma(t)P(t), \tag{7.1}$$

where P is the probability that the qubit hasn't escaped yet. dP and P can be extracted from the histogram as follows. Let $S(t)$ be the cumulative number of switches measured up to a time t from the start of the current ramp. I denote the total count of all switches measured as S_∞. Since the swept ramp is designed to exceed the SQUID critical current, a switch is guaranteed on each sweep and S_∞ is just the total number of repetitions. Given this notation,

$$\begin{aligned} P(t) &= \frac{S_\infty - S(t)}{S_\infty}; \\ dP(t) &\approx P(t+dt) - P(t) \\ &= \frac{S(t) - S(t+dt)}{S_\infty}; \\ \Gamma(t) &\approx \frac{S(t+dt) - S(t)}{S_\infty - S(t)} \frac{1}{dt}. \end{aligned} \tag{7.2}$$

Interpreting the final equation, $\Gamma(t)$ is just the fraction of all future switches occurring between times t and $t + dt$, divided by the time interval dt.

For spectroscopic measurements, the escape rate is more meaningful than the histograms. Comparing the background and microwave escape rates Γ_{bg} and Γ_μ, we can find the microwave enhancement $(\Gamma_\mu - \Gamma_{bg})/\Gamma_{bg}$. Γ typically varies by several orders of magnitude over a narrow bias window. The microwave enhancement normalizes this out, whereas direct comparisons of histogram counts or escape rates conceal the resonance at lower biases.

Figure 7.2 shows a color plot of the microwave enhancement for the complete set of spectroscopy data at one power, -35 dBm[1]. The x-axis is the applied bias current, the y-axis is the frequency of microwaves applied, and the color is the base 10 logarithm of the enhancement. For this dataset, the enhancement varies

[1]Unless otherwise specified, powers listed are always the nominal power at the source.

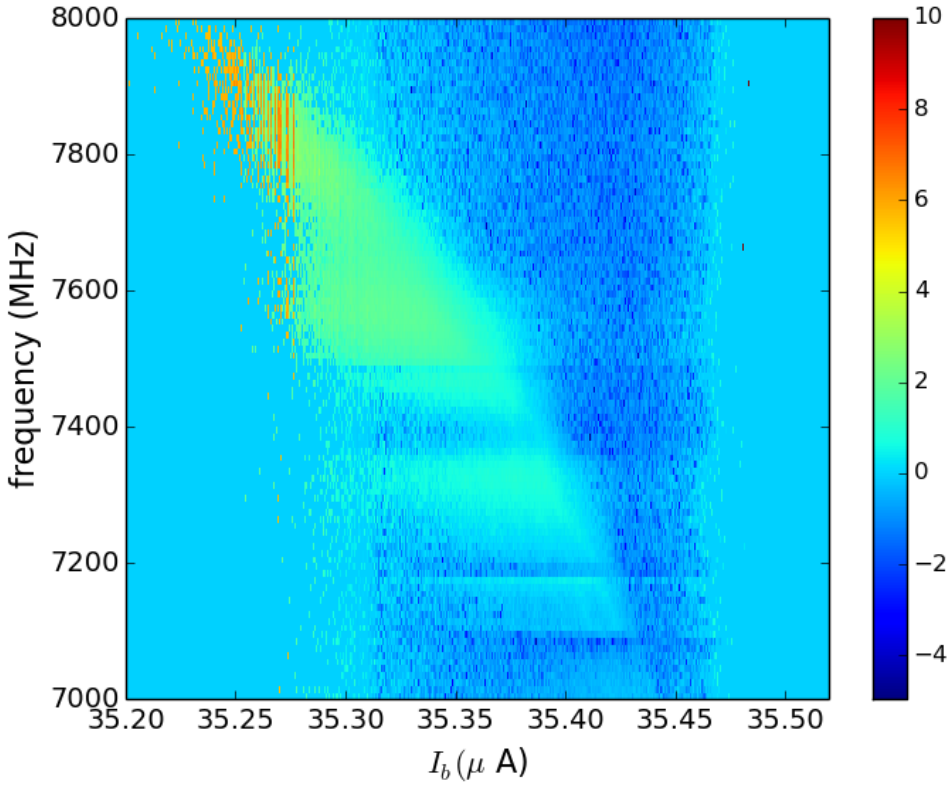

Figure 7.2: Plot of microwave enhancement from swept spectroscopy taken on the first cooldown. Each row of this plot represents enhancements calculated from individual data sets like the one shown in Figure 7.1.

tremendously with bias and frequency. The light green arc pointing toward the upper left corner of the plot is the qubit resonance. Lower bias yields higher f_{01}.

Beyond this broad description, this plot is difficult to use. Because the number of counts in a bin are integers, the smallest nonzero $\Gamma(t)$ that can be measured is $1/(S_\infty dt)$. This is where the low escape cutoff in Figure 7.1 comes from; similarly, one cannot measure an escape rate higher than $1/dt$, giving the high escape cutoff. The low cutoff is problematic, because if at a given bias there are visible microwave escapes but no background escapes, the effective enhancement is infinite. One can "fix" this by pretending any unmeasured escape rate is 1 instead of 0, but as seen in Figure 7.1 this would neglect the considerable variation in the background escape rate, skewing the resulting data.

There are a few ways to improve the enhancement plot:

- One can perform more repetitions at each frequency, increasing S_∞ and extending the low cutoff. This is usually not feasible; extending the escape rate two decades would require 100 times as many counts, which would take days.

- One can work only with microwave peaks that occur entirely within regions with a measured background escape rate. For earlier qubits studied in our lab [25, 71] this worked quite well. However, the "peaks" I measured in swept spectroscopy studies looked comparatively very broad. In Figure 7.1, a sharp shoulder can be seen on the high-bias side where the escape rates collapse back to the background. By contrast, microwave enhancement is present in a very large region on the low-bias side. Limiting the measurement in this way would leave very little spectrum to work with.

- One can extrapolate the background escape rate to lower bias on the basis of the measured data. The method of doing this in the past assumed that Γ had an exponential dependence on the bias current, thus forming a straight

line on these plots. This is a good description of the predictions of the WKB approximation for the tunneling of a single current biased junction. For a dc SQUID, it is less clear that this is the case [37]. For example, in Figure 7.1 some rounding in the escape rate can be seen at high bias.

On top of these practical difficulties, measuring anything at all with the swept approach requires that the system be biased close to the critical current. In these regions, f_{01} is more sensitive to small fluctuations in bias, leading to increased broadening of spectroscopic features. For all of these reasons, I did only a limited amount of swept spectroscopy. However, it still proved useful for establishing a microwave feature with the correct qualitative behavior of the resonance. It also suggested a good frequency range to look at using pulsed spectroscopy.

7.2 Pulsed spectroscopy, Rabi, T_1

Pulsed spectroscopy fixes the microwave enhancement problem. The key idea is that using brief measurement pulses raises the qubit's escape rate to a measurable level. Figure 5.11 shows the timing diagram for a pulsed measurement. Instead of a continuous linear ramp, the bias consists of an early flux shuffle followed by a ramp of current and flux to a set plateau. Once the plateau is reached, I wait to allow any ringing from the bias low-pass filters to settle. Qubit control microwaves are then turned on. By waiting until I have reached the desired operating point to turn on the microwaves, I avoid errors introduced by microwave excitation in the swept region leading up to the bias plateau. Finally, I send a measurement pulse at some time after the turn-on of the control microwaves. The timing of the pulse determines what type of measurement I perform.

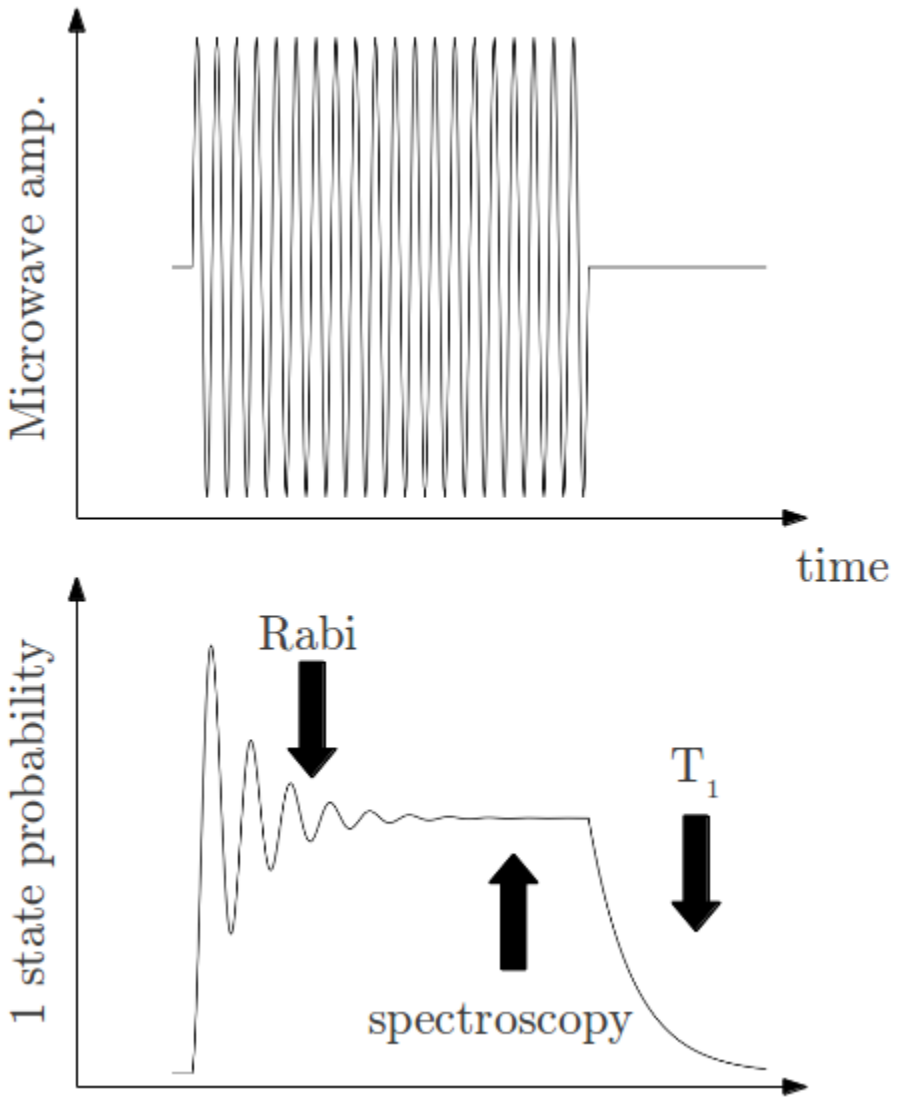

Figure 7.3: Cartoon representation of three microwave measurements I performed. The microwave frequency is much faster than suggested by this image. The arrows correspond to different choices of when to send the measurement pulse.

7.2.1 Preliminaries

Figure 7.3 illustrates the pulse timing of the three main measurements I took. The top plot qualitatively shows the amplitude of the microwave signal versus time; the actual microwave frequency is orders of magnitude faster than the Rabi frequency of the qubit. The bottom plot sketches the expected excited state probability ρ_{11} as a function of time. Before the microwaves are turned on, the qubit is in its ground state. When the microwaves come on, the qubit begins coherent Rabi oscillations. If the microwaves are left on for a time much longer than the Rabi coherence time T', the qubit population reaches some steady state. After the microwaves are turned off, this steady state population relaxes to the ground state.

By choosing when I send the pulse, I can monitor the qubit state in each of these three regions. If I take a series of measurements where the measurement pulse is just before or immediately following the microwave turn-on, I measure the Rabi oscillations. By choosing a measurement pulse several microseconds after the microwave turn-on but before the microwave turn-off, I can monitor the long-term saturation behavior. This corresponds to a spectroscopic measurement. By making a rapid series of measurements following microwave turn-off, I can time resolve the decay of the population and extract T_1. For each pulse location, I typically perform thousands of repetitions to get a good estimate of the probability P_{sw} of measuring a switch.

In addition to the timing of the pulse, the amplitude and shape of the pulse are also variable. This flexibility motivates the pulse technique. Stronger or longer pulses lead to higher escape rates. This allows me to operate at bias plateaus inaccessible to swept measurement, since I can enhance the escape rate by using a bigger pulse. Typically, I only varied the amplitude of the pulse. Pulses that are too short require impractical amplitudes given our equipment and also carry the risk of exciting transitions between states. Pulses that are too long smear out the

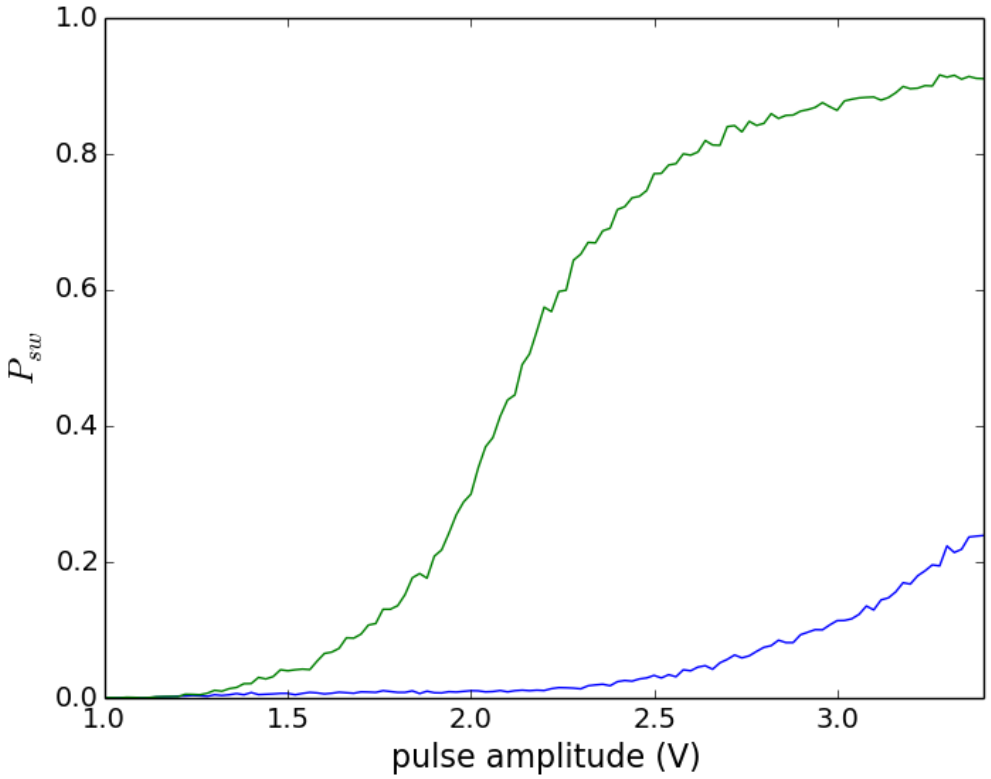

Figure 7.4: s-curve data taken from the first cooldown. With microwaves turned on (green curve), smaller pulses are needed to switch the qubit.

measurement information over a long time interval, making it more complicated to reconstruct the original state. I settled on a pulse time of 6 ns.

Figure 7.4 shows how the escape rate changes as the pulse height is adjusted. Two curves are plotted. One corresponds to the background escape rate, and the other to the escape rate at long time with microwaves on. The x-axis is the height of the microwave pulse, measured as the nominal voltage on the DG535 sourcing the pulse. The y-axis is the probability P_{sw} that the qubit switches during the pulse; this is measured by dividing the total number of counts occuring in the histogram window surrounding the pulse by the total number of measured counts (typically 3000-10000). Given the shape of the resulting data, these sorts of pulse calibration plots are called s-curves. When microwaves are turned on near resonance, the qubit

is excited, leading to an enhancement in the escape rate. This yields an s-curve with greater P_{sw} for a given pulse size.

The s-curve shows that the discriminating power of the measurement depends on the amplitude of the pulse. For sufficiently small or large amplitudes, there is no difference in the escape probability between the microwave and background data. The pulse is either too small to excite any escapes, or sufficiently strong that even without microwaves the system always tunnels. The point of optimal contrast is where the probability difference between the microwave escape rate and the background escape rate is largest.

Ideally, s-curves would be performed for each possible bias condition in order to optimize the measurement. In practice this imposes severe overhead on already lengthy measurements. In addition to the time spent measuring the s-curves themselves, P_{sw} must also be calibrated to recover the excited state probability ρ_{11} from it. This requires assumptions about the state of the qubit at the pulse time where the s-curve is taken. To make accurate estimates of the qubit state requires Rabi oscillations and fits. The process is further complicated by the presence of any extraneous population in higher states of the qubit. This higher state population also shows up as switches but yields errors if treated as $|1\rangle$ state population.

Instead of performing these calibrated s-curves for every measurement, I instead adjusted the pulse so that I achieved a fixed switching probability for the background data. This technique was used by Tauno Palomaki [27] in his work setting up this measurement technique for our qubits. Since the expected difference in the escape rate between $|0\rangle$ and $|1\rangle$ at high bias is $\Gamma_1 \approx 50\Gamma_0$, a background switching probability of 2 percent suggests nearly guaranteed switching of the $|1\rangle$ state. Pulses at this height, therefore, give high contrast between $|0\rangle$ and $|1\rangle$. I typically used a probability of about 4 percent, reasoning that the possible loss of contrast was minor compared to the possible gain if the escape rate ratio Γ_1/Γ_0 fell

short.

For spectroscopy experiments that scan qubit bias as well as the microwaves, I established this measurement pulse amplitude for every second bias in the scan and linearly interpolated to get an appropriate pulse amplitude for the remaining bias points. Using this approach, I was able to measure spectroscopy, Rabi oscillations and T_1 decays in a reasonable amount of time. However, as I discuss in Chapter 8, the measurement process in our recent dc SQUID phase qubits proved very strange. Several measurements seem incompatible with a direct mapping of the measured switching probabilities to the populations in $|0\rangle$ and $|1\rangle$. This makes some aspects of the data — particularly the height of various lines — difficult to interpret. Other features, such as the presence of multiple lines or the coherence timescales, are less affected by this problem. To make clear that the measurement is uncalibrated, I label plots with P_{sw} or the microwave enhancement $P_{sw}^{mic} - P_{sw}^{bg}$ instead of ρ_{11}.

7.2.2 Spectroscopy

Figure 7.5 shows two sets of spectroscopy data taken during the first cooldown. In both plots, the x-axis represents the tuning of some bias parameter; in the top plot, it is the applied flux, while in the bottom plot, it is the applied bias current. The y-axis is the frequency of applied microwaves, and the color scale represents $P_{sw}^{mic} - P_{sw}^{bg}$. Thus, a vertical cut through the color plot represents the microwave response of the qubit at one set of bias parameters and a single frequency.

In both plots, two distinct lines are visible. In Figure 7.5a), the lower frequency line varies relatively strongly as the flux is varied. By contrast, in the current plot, the dependence on the current bias is weak. This behavior is consistent with the lower frequency line being the qubit resonance. Given our design, the qubit mode is expected to be dominated by the behavior of the qubit junction. As discussed in Section 2.5, the strong inductive isolation of the SQUID shunts off most of the

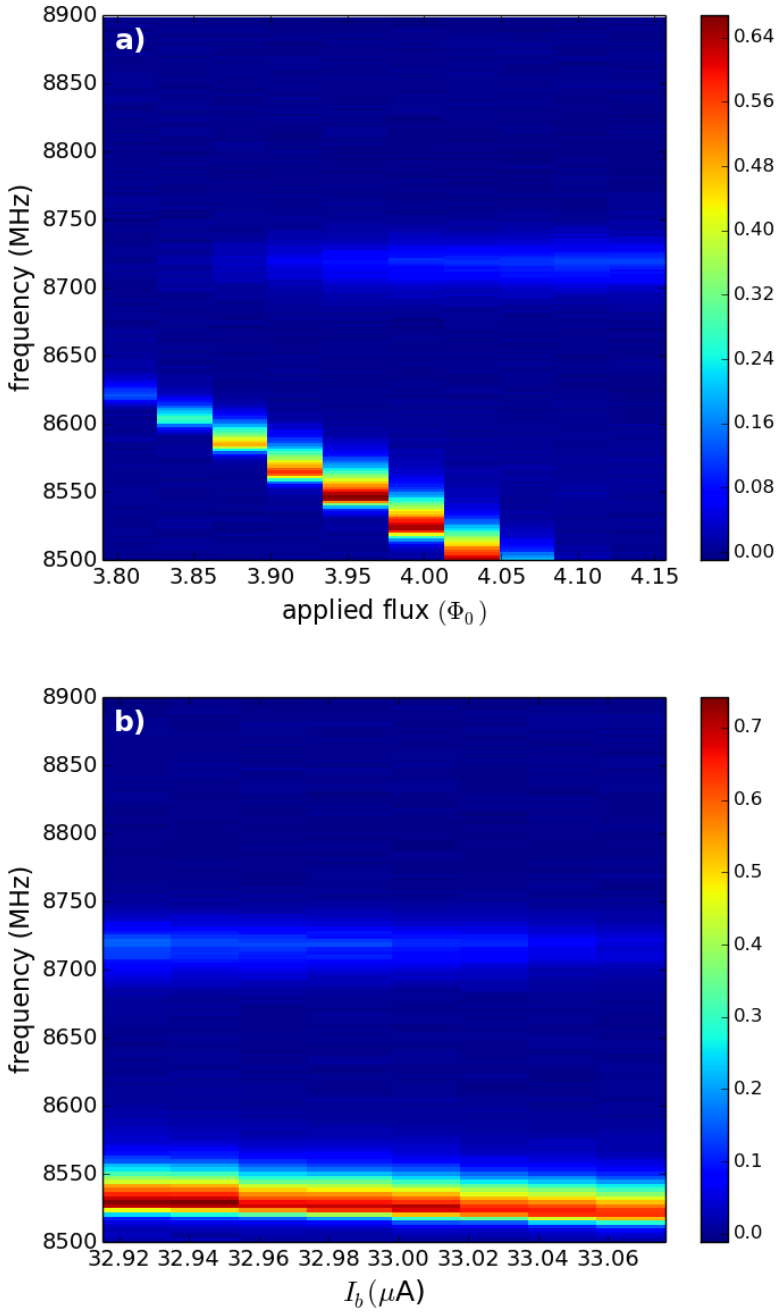

Figure 7.5: Spectroscopic data from the first cooldown. a) Bias held constant while flux varies on the x-axis. b) Flux held constant while bias varies on the x-axis.

current from the bias line, leading to modest dependence of ω_{01} on the bias. Since the flux line affects both junctions equally, we see a much stronger response there.

The higher line shows no noticeable response to either form of bias, suggesting that this is something else. A likely candidate would be a physical resonance in the microwave setup itself, such as a cavity mode of the sample box or a standing wave in a section of microwave line. An argument against this interpretation is that the location of this feature shifted slightly between cooldowns. Although there were some deliberate changes to the microwave wiring, they were quite far removed from the qubit and would have had to go through 40 dB of isolation to affect the qubit. On the other hand, it is possible that thermal cycling of the system may have caused physical changes in the immediate environment of the qubit.

An alternative interpretation of the flat line is that it represents the unbiased frequency of the qubit mode. At first, this seems impossible — I am claiming the low frequency line is the biased qubit mode, which would appear to preclude measuring the unbiased mode at the same time. Because the SQUID supports multiple flux wells, however, it is possible that the qubit might tunnel between flux wells without going to the voltage state. If one well is relatively heavily biased, then the other wells will usually be far from maximally biased. Thus, if the ramping process or application of microwaves leads to occasional tunneling between wells, it may be possible to see the resonance in those new wells. In doing fits later on, I use this interpretation of the flat feature.

One notable feature of the spectral response in this qubit is pictured in Figure 7.6. The x-axis of this figure shows the nominal power of the qubit microwave drive at the source. The y-axis shows the peak (extracted from spectroscopy data at each power) of $P_{sw}^{mic} - P_{sw}^{bg}$. As the applied power increases, the frequency of the resonance decreases.

This sort of behavior is common in nonlinear oscillators. As an example, the

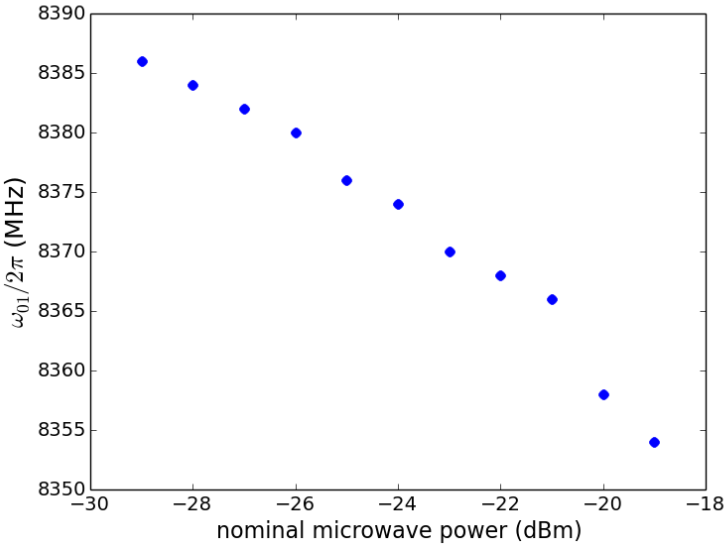

Figure 7.6: Location of the resonance peak as a function of applied microwave power.

Duffing oscillator — which is a good approximation to a weakly biased Josephson junction where the cubic term in the potential is small — has the same qualitative frequency response to increasing power [138]. In fact, this was one of the key ideas underlying the Josephson bifurcation amplifier [79]. What is unusual here is that, by measures such as Rabi oscillations, we do not appear to be driving the system especially hard. On-resonance Rabi frequencies seem comparable to the similar qubits Kwon and Przybysz measured [28, 29], though they do not report observing these power-induced frequency shifts. On the other hand, Dutta et al. [73] reported measurements of the ac Stark effect in strongly driven phase qubits, so this term is not unprecedented.

I repeated the spectroscopy measurements on all of the cooldowns to see if anything had changed. Figure 7.7 shows representative spectroscopy data from the second cooldown. The top plot consists of a full spectroscopy experiment over multiple choices of bias, as in Figure 7.5. The bottom plot is a cut through the indicated flux plateau to clarify what is going on. The spectrum now exhibits smaller, secondary peaks to either side of the main resonance. As the qubit bias is

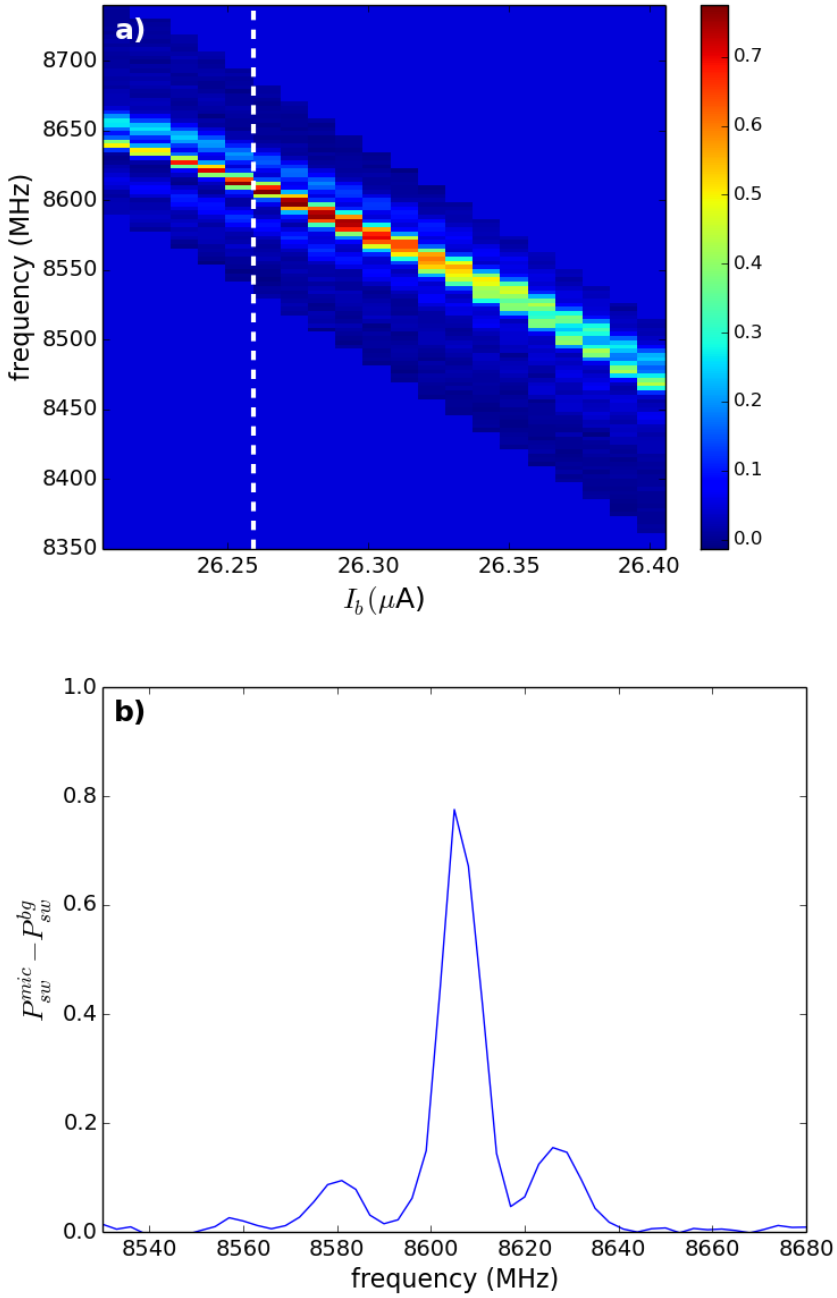

Figure 7.7: Spectroscopy data with sidebands from the second cooldown. a) Flux and bias varied according to the double-ramp extracted from I-Φ measurements. b) A cut of plot a) through the indicated bias plateau.

adjusted, these peaks move as well. Some splitting of the central resonance on one side also appears.

The cause of these sidebands is unclear. One possibility is the natural anharmonicity of the qubit. The scale of this anharmonicity — represented by $(\omega_{01} - \omega_{12})/2\pi$ — can be calculated from the model presented in Chapter 3, either through direct numerical calculation or by using perturbation theory to incorporate the nonlinearities. For many of the measurements on BBC2, the qubit junction was relatively weakly biased. In this limit, a reasonable lower bound on the anharmonicity can be set using a slightly modified single junction approximation. The qubit junction is treated as unbiased and the quartic term in the potential treated as perturbation. To incorporate the SQUID, the loop inductance is added in parallel to the L_{jq} to renormalize the frequency of the qubit. For BBC2, this gives a minimum anharmonicity at zero bias of $(\omega_{01} - \omega_{12})/2\pi \approx 10$ MHz. With stronger bias, this anharmonicity should increase.

Given this estimate of the nonlinearity, the distance between the central peak and the sidebands seems like it might plausibly be associated with the junction nonlinearity. However, since we expect the population of higher levels to be low to begin with, the $0 \rightarrow 1$ line should be most significant. This should be the highest frequency line, not the one in the middle which we typically measure to be the strongest.

Another possibility is Mollow sidebands [139]. These arise when a two-level system is driven sufficiently hard. They appear symmetrically about the resonance, occurring at frequencies $\omega_{01} \pm \Omega$, where Ω is the Rabi frequency. Although we had not observed these sidebands in earlier qubits, they have been observed in transmon qubits [140].

Figure 7.8 shows another set of spectroscopic data taken from the second cooldown. The different colors represent spectroscopy runs taken at different applied

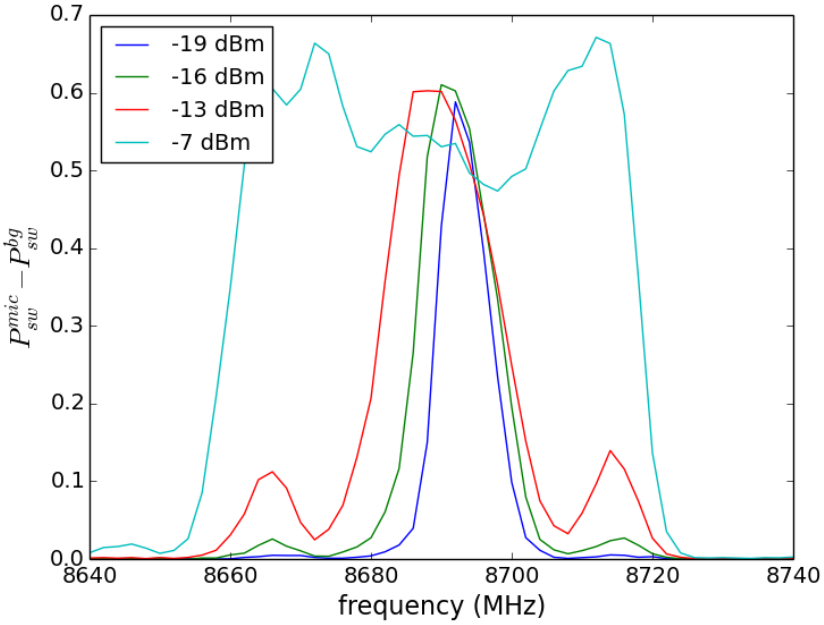

Figure 7.8: Spectroscopic data from the second cooldown at several powers, examining the effect of power changes on the sideband structure.

microwave powers. The x-axis is the applied frequency, and the y-axis is $P_{sw}^{mic} - P_{sw}^{bg}$. Note that while the height of the sidebands depends strongly on power, they do not appear to spread relative to the center peak (in fact, they seem to move less than the central peak itself does.) While the frequency separation of $\approx$ 20 MHz is on par with measured Rabi frequencies, this lack of movement suggests these are not Mollow sidebands.

Figure 7.9 shows that the sideband splitting and visibility (along with the whole peak) can be suppressed by varying only the flux while leaving the bias current constant. Here, the x-axis represents changes in the applied flux during the bias plateau, while the y-axis is the applied frequency and the color is the enhancement of the tunneling probability. This result appears related to some strange properties of the tunneling measurement I observed. If the switches following the measurement pulse are resolved finely enough, tunneling events can be seen at two well-separated points. The sideband behavior appears largely associated with tunneling measured

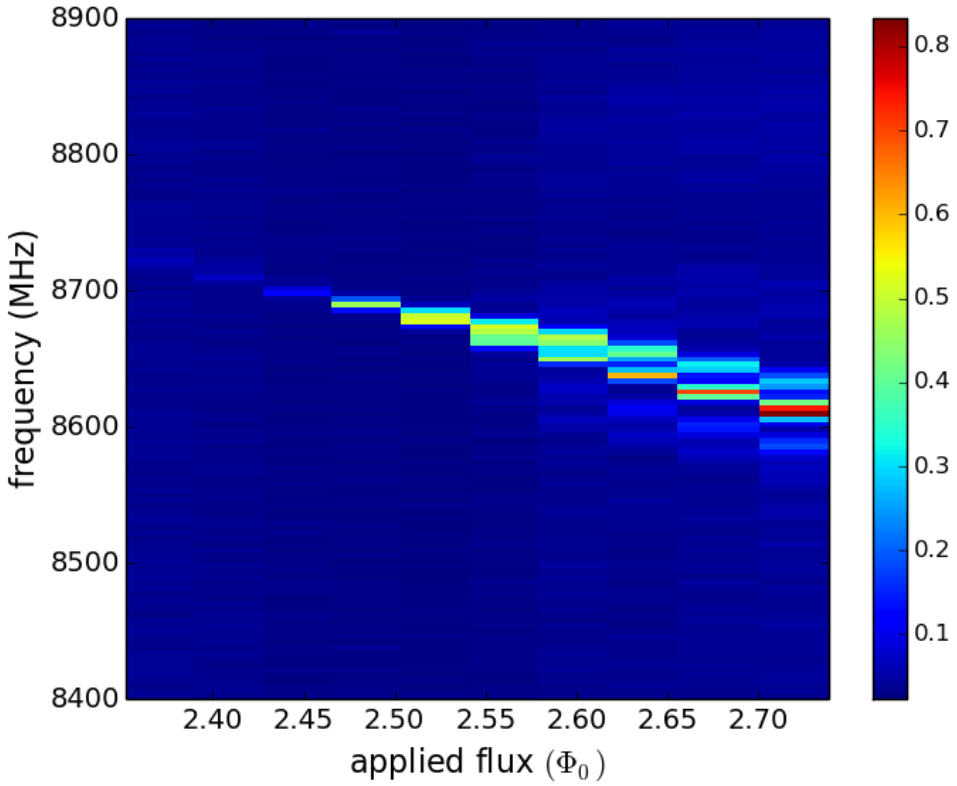

Figure 7.9: Spectroscopic data where the bias is held fixed while the flux is varied. This leads to suppression of the sidebands.

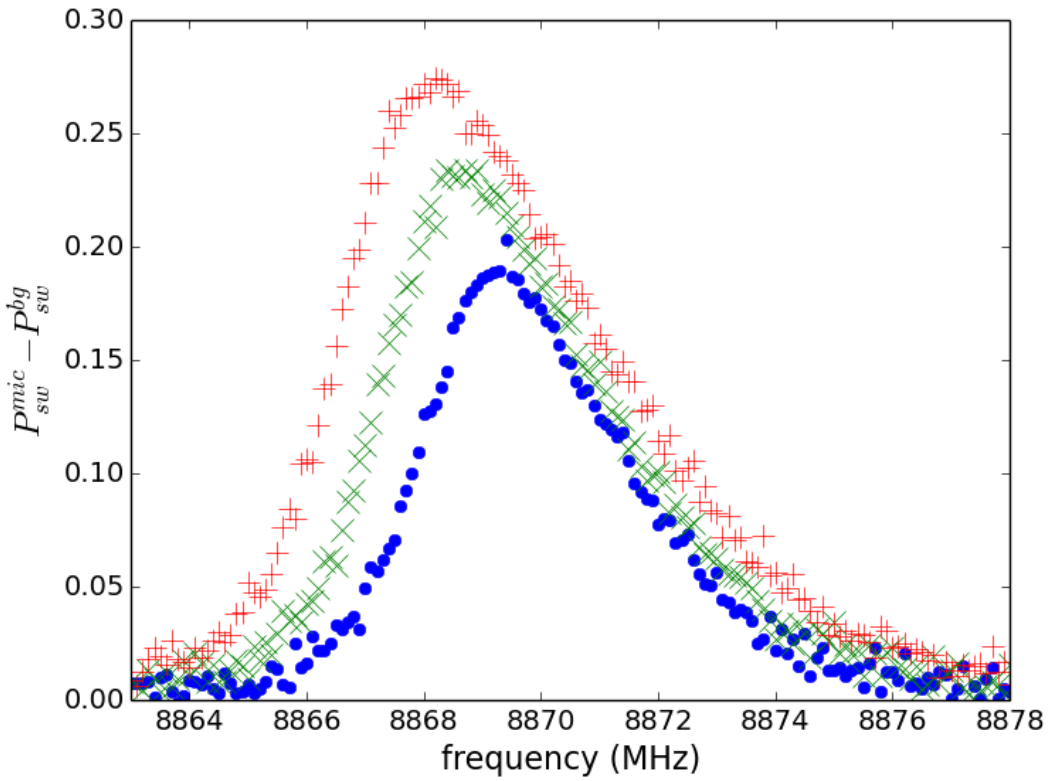

Figure 7.10: Spectroscopic data from the final cooldown, taken in the region where sidebands are suppressed. The three data sets correspond to nominal microwave powers -28, -26 and -24 dBm.

at the second pulse. This aspect of the measurement will be discussed further in Chapter 8.

A promising feature of Figure 7.9 is the sideband free region where a narrow single line is observed. Figure 7.10 shows this line for three different powers at fixed bias. Even with the sidebands eliminated, the shift in the resonance frequency and broadening with power are visible. Since the spectroscopic linewidth is a measure of the coherence of the qubit, in looking for good places to study the qubit I sought the sharpest line possible.

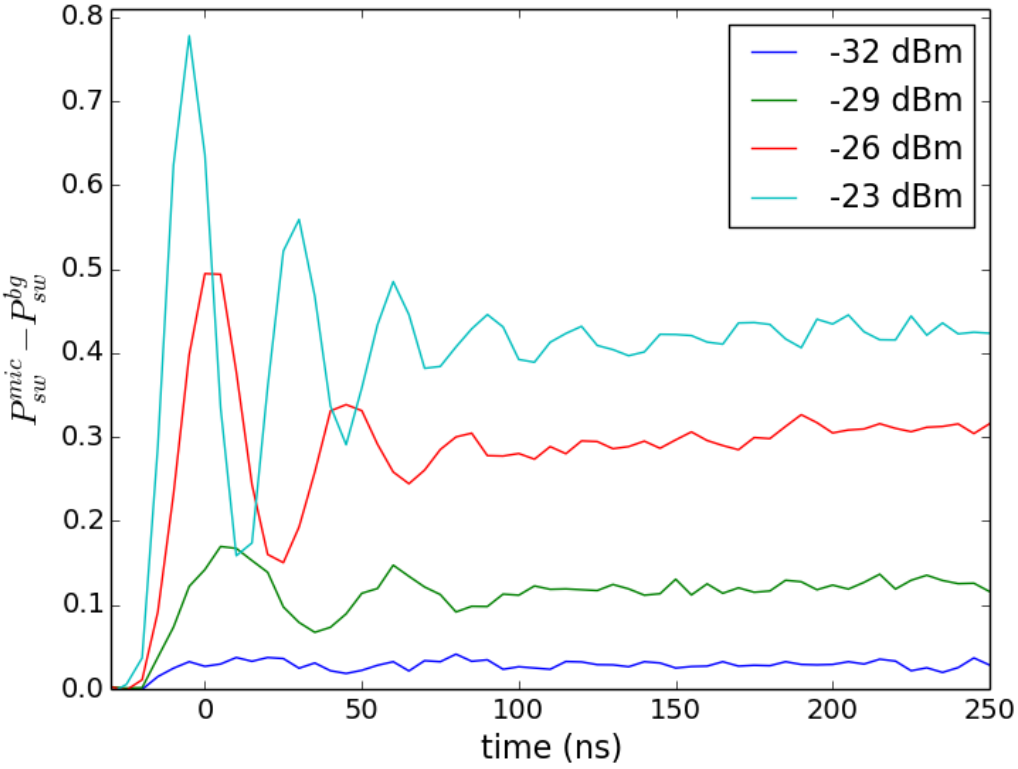

Figure 7.11: Rabi oscillations from the first cooldown. Several different powers are shown.

7.2.3 Qubit coherence: Rabi and T_1

As shown in Figure 7.3, I measured Rabi oscillations by sending the measurement pulses immediately following the microwave turn-on. As mentioned in Section 5.2, to get enough channels to do everything I want to do I use two DG535 sources. One of these gates the microwave source, turning it on for a short window late in the bias plateau as shown in Figure 5.11. The other controls the timing of the measurement pulse. There is usually a small delay between these two boxes; accounting for that, I can scan the measurement pulse over a window surrounding the turn-on of the microwaves.

Figure 7.11 shows some of the earliest Rabi data on BBC2, taken during the first cooldown at a microwave frequency of 8422 MHz. The x-axis plots the time

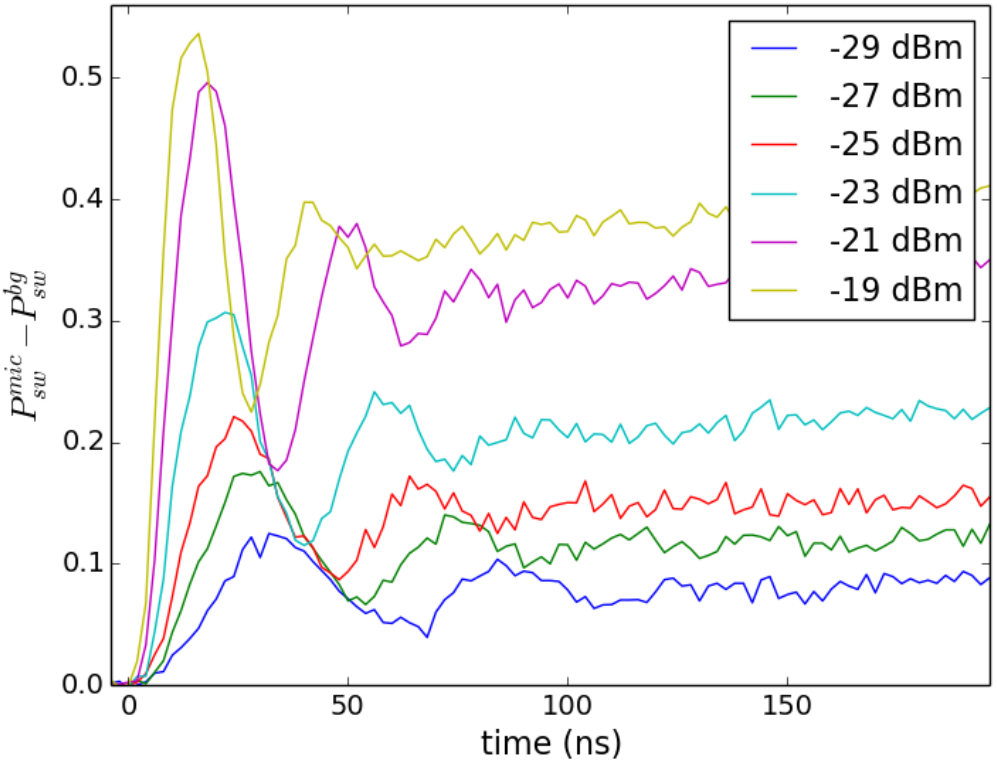

Figure 7.12: More data on Rabi oscillations at various powers. In this data, the frequency corrections of Figure 7.6 are applied; each power is at the corresponding peak frequency.

difference between the DG535 gating the microwave source and the DG535 sending the measurement pulse. The y-axis plots $P_{sw}^{mic} - P_{sw}^{bg}$. Four different powers are shown, each separated by -3 dB.

Qualitatively, this data is somewhat surprising. The Rabi frequency does increase as the power goes up, which is to be expected. What is more unusual is that the apparent timescale over which the Rabi oscillations decay appears to be longer for the higher power signal. One possible explanation for this would be a low frequency noise source. For example, $1/f$ noise would be less at the higher Rabi frequencies.

However, taking Rabi oscillations at the same drive frequency while adjusting the power neglects the frequency shift pictured in Figure 7.6. Using the resonant

frequencies vs. power extracted from that plot, Figure 7.12 shows Rabi oscillations at different powers when the frequencies are adjusted to remain on resonance. Here we still observe the increasing Rabi frequency as the power increases, although the increase is not as large as expected — on resonance, the Rabi frequency is proportional to applied power. However, the number of oscillations visible appears independent of power. With the shorter periods at higher power, this corresponds to reduced Rabi decay time T'.

By contrast, Figure 7.13 shows the results of detuning the Rabi drive from resonance. Figure 7.13a) shows a spectrum taken from the first cooldown at a nominal microwave power of -19 dBm. Using the same bias and power, I performed Rabi oscillations at several frequencies around the resonance. Several of those frequencies are highlighted and the corresponding Rabi oscillations are plotted in Figure 7.13b).

The visibility and lifetime of the Rabi oscillations are much longer if the applied drive is blue detuned relative to the measured resonance. As seen in Figure 2.10, it is not necessarily surprising to see more oscillations when detuned from resonance. However, the asymmetry is noteworthy, as is the longer decay envelope and the unusual shape. The unusual shape, like the sidebands, seems related to the measurement complications of Chapter 8. This behavior occurred on both the first and second cooldowns; this is significant because the method of applying pulses changed between these two cooldowns, suggesting that this behavior was unlikely to be an obvious artifact of the pulsing method. This behavior also occurred in regions where the sidebands were suppressed, as in Figure 7.10.

What might be causing this behavior in the Rabi oscillations? A possible clue is the asymmetry of the lineshape itself. Intuitively, any bias fluctuations leading to small shifts in the resonance will be more significant on the sharper red side of the resonance than the shallow blue side. However, this type of argument suggests that the oscillations should be longest lived on resonance, where the sensitivity to

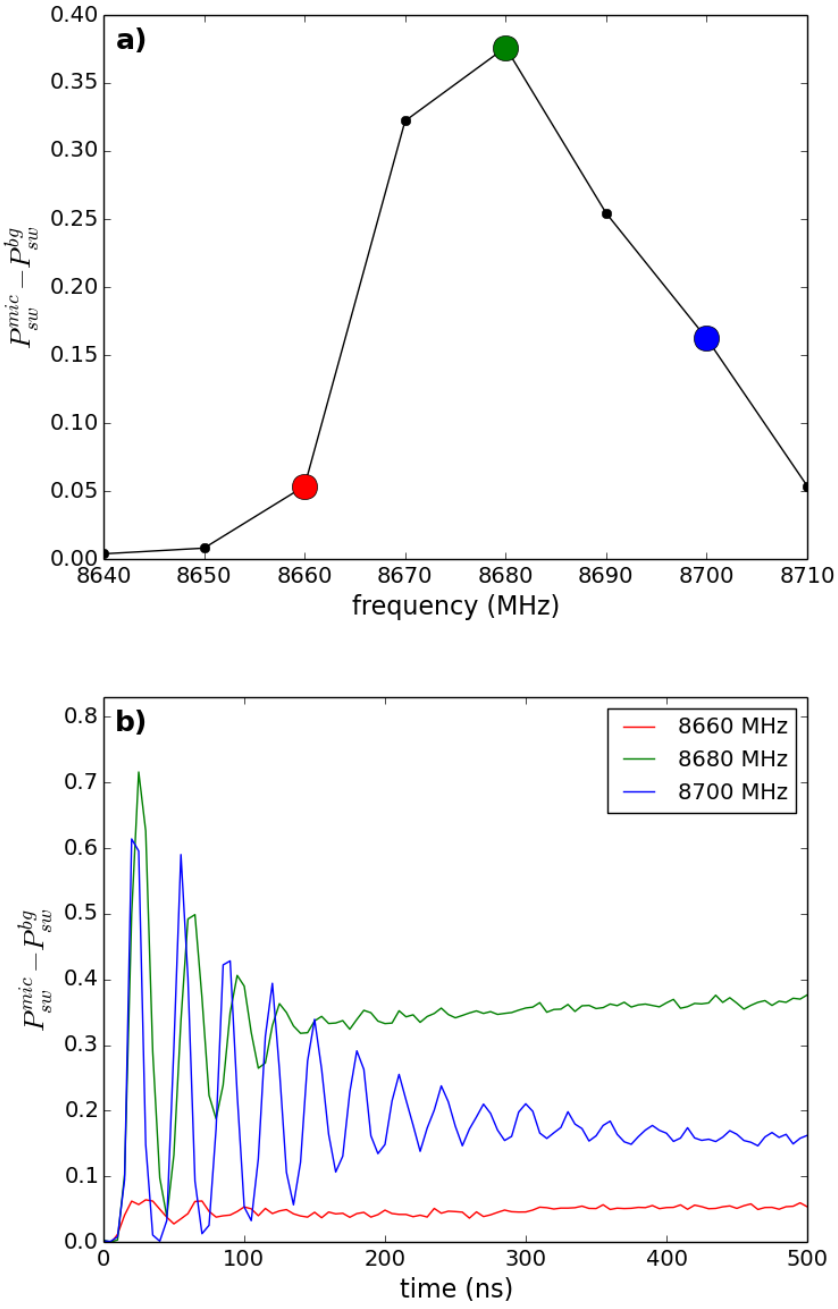

Figure 7.13: a) Spectroscopy data at nominal microwave power -19 dBm. b) Rabi oscillations at -19 dBm and several frequencies; this corresponds to changing the detuning.

fluctuations is weakest. Thus there is presumably something else going on. In the following section, I will discuss another possible explanation.

Figure 7.14 shows some of the best Rabi and T_1 data I measured for BBC2. These were taken at the end of the second cooldown in a region where the sidebands were suppressed. Although true fits of these functions rely on the $|1\rangle$ state population and not just $P_{sw}^{mic} - P_{sw}^{bg}$, barring the involvement of higher levels those two quantities should be linearly related, since the total escape rate is a weighted average of the escape rate for the 0 and 1 states. The two plots were taken at the same bias conditions but at different applied microwave powers. For measuring T_1 this is unimportant as no microwave power is applied.

I fit the T_1 plot with an exponential decay, giving $T_1 = 230$ ns. While substantially shorter than the 1.2 μs estimated in Table 6.2, that estimate used anticipated design values. The actual BBC2 device differed these design values; in particular, the mutual inductances M_f and M_μ proved to be about an order of magnitude bigger than the original design values. In the case of M_f, this was because the original design's M_f proved so low that flux biasing the qubit required excessive current. With M_μ, the design value was based on assumptions needed to interpret microwave simulations of the qubit. Using the measured values for all of the parameters (see Table 7.1 in the next section), I find a revised estimate of $T_1^{est} = 320$ ns. This suggests the noise model of Chapter 6 is reasonable.

The Rabi plot is fitted to a numerical solution of the Bloch equations. Since the measurement is of $P_{sw}^{mic} - P_{sw}^{bg}$ and not ρ_{11}, I apply an overall scaling factor .26 to the numerical solution. The other fitting parameters were the detuning $\Delta/2\pi = 8$ MHz, $\Omega/2\pi = 10$ MHz, $T_1 = 215$ ns and $T_\phi = 1100$ ns. Experimentally, I was aiming for a detuning of 7 MHz, so that number is reasonable. The T_1 is in line with the T_1 extracted directly from the decay measurement. The particularly long T_ϕ is noteworthy; it is the best dephasing timescale that we have found for a qubit

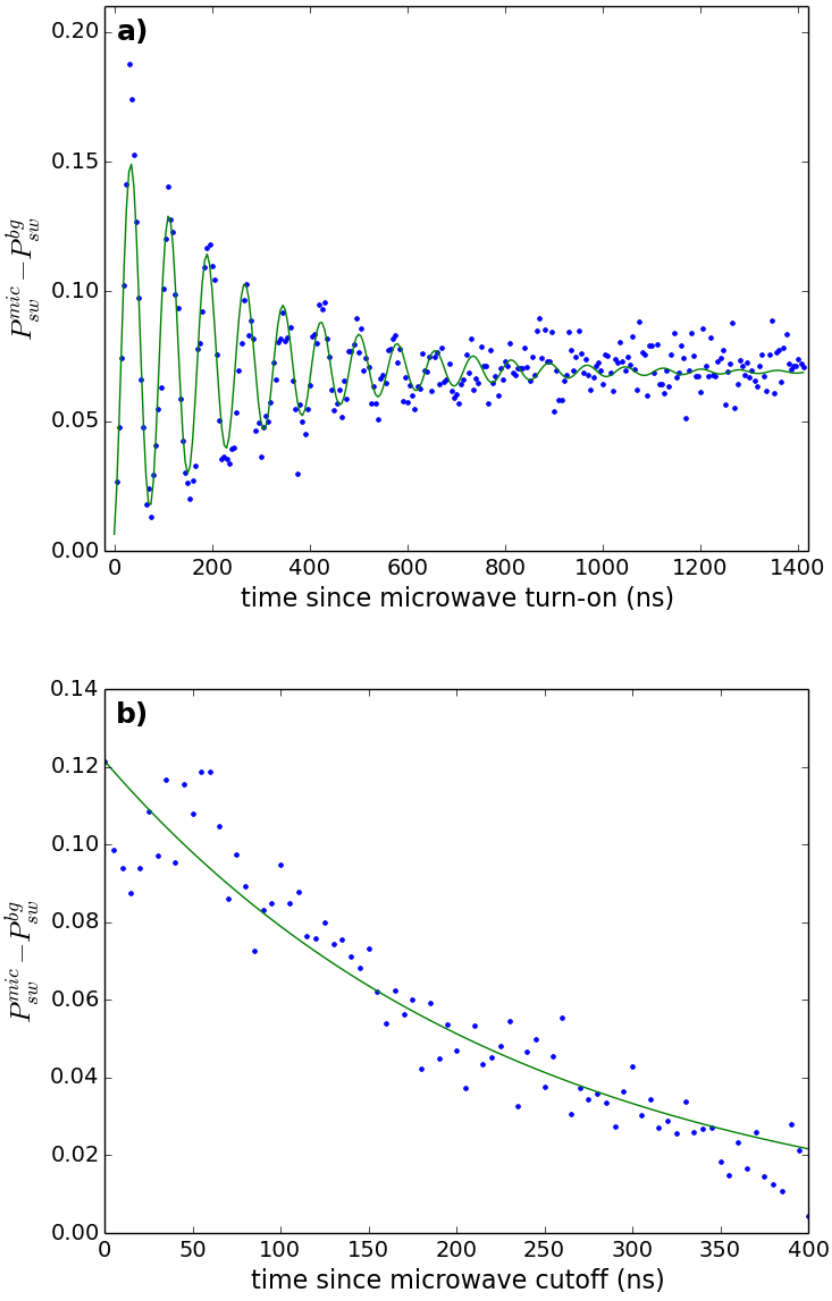

Figure 7.14: Rabi oscillation and T_1 measurements from BBC2. a) Rabi oscillations, fit to a numerical solution to the Bloch equations yielding $T_1 = 215$ ns, $T_\phi = 1100$ ns. b) Qubit relaxation after microwave turn-off, fit to an exponential decay with timescale $T_1 = 230$ns.

in our lab. In Chapter 8 I present a table comparing these times to other dc SQUID phase qubits.

7.3 Effective Jaynes-Cummings model

In Chapter 3, I discussed how the on-chip LC filter in the bias line could be incorporated into an effective Jaynes-Cummings model for the qubit. The impetus for this model was the spectroscopy data of the previous section. The spectroscopic sidebands were suggestive of the photon splitting associated with the coupling to a resonator. While many things could serve as a resonator, the most logical oscillatory modes are the isolation mode of the SQUID and the filter. The isolation mode is at very high frequency relative to the qubit mode and I expect it is frozen out.

Thus, I examined the filter plus SQUID model and sought to determine if the coupling g was sufficient to produce photon splittings of that scale, as well as whether g modulated enough with bias to account for the changing presence of the sidebands. While the model does not yield enough g to account for the sidebands, it does offer an explanation for the asymmetric spectroscopic lineshape. This lineshape derives from the presence of multiple unresolved photon peaks. These photon peaks appear at higher frequency than the qubit resonance because $\omega_f < \omega_q$. In this section I will discuss my model for this behavior, followed by the construction of fits to experimental spectroscopic data.

7.3.1 g and photon shifts

As discussed in (3.11), in the dispersive limit where $g/\Delta \ll 1$, the qubit frequency appears to be shifted by an amount proportional to the number of photons in the resonator. Since $\omega_f \ll \omega_q$, it would take very large coupling to leave the dispersive limit. This suggested redoing some of the spectroscopy to see if individual

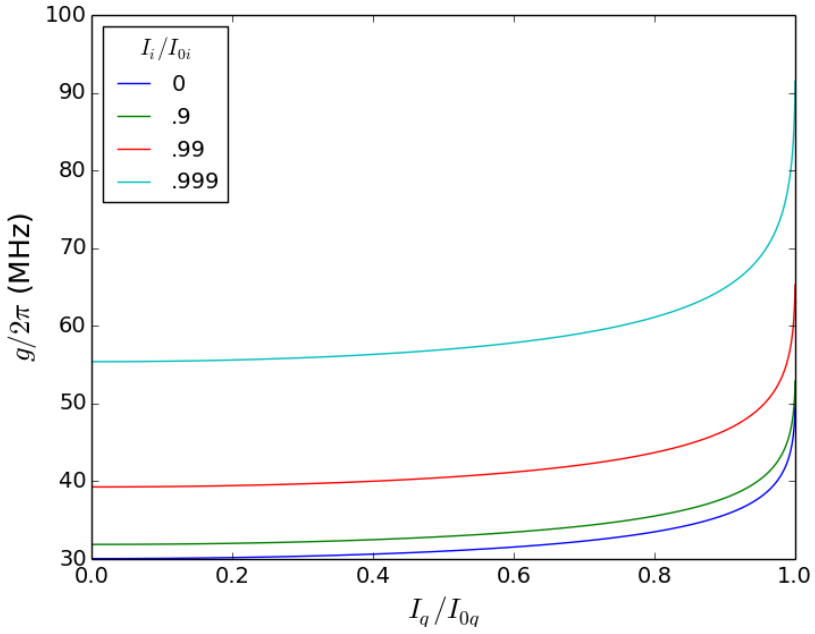

Figure 7.15: Variation of g with junction phases. Each curve represents a particular choice of γ_i.

peaks corresponding to distinct photon numbers could be resolved. This effect, first observed in superconducting qubits by Schuster et al. [104], is one of the hallmarks of Jaynes-Cummings behavior.

Ultimately I was unsuccessful in resolving individual peaks. However, the shape of the spectral lines — broadened on the blue side of the line — is consistent with multiple unresolved photon peaks. The anticipated filter frequency was near 200 MHz, corresponding to a temperature of 10 mK. Given that even the nominal base temperature of the fridge was 20 mK — and the electronic temperature could have been considerably higher — I expected many photon states to be occupied. To see if the broadening could be explained in this way, I developed the model of Chapter 3 to calculate g.

As shown in (3.54), g can be derived from the qubit and filter model:

$$g = -\sqrt{\frac{1}{2m_f\omega_f\omega_Q}}\left(\frac{\Phi_0}{2\pi}\right)^2\left[\frac{L_q}{\mathfrak{L}^2}A_{iQ} + \frac{L_i}{\mathfrak{L}^2}A_{qQ}\right]. \tag{7.3}$$

The various coefficients in this model are determined by the device parameters and the applied flux and current biases. Figure 7.15 shows several plots of $g/2\pi$ for the parameters of BBC2 as the junction biases are varied. There is some dependence on bias. At high biases, the Josephson inductances approach the larger inductances of the problem, L_q and L_f. Since the coupling is determined mostly by the relative importance of L_f, this increase in coupling with bias is unsurprising. However, the effect is pretty small, and over a large range of biases g responds weakly to bias.

Using biases corresponding to the spectral peaks examined during the third cooldown for the presence of photon peaks, Equation 7.3 predicts $g/2\pi \approx 32$ MHz. The size of the expected photon shift is $g^2/\pi\Delta \approx 200$ kHz per photon in the resonator. Much of the final spectroscopy was resolved at 100 kHz, so with sufficiently sharp lines this might have been observable. However, given the measured coherence times, the expected width of spectral peaks is on the order of megahertz. Thus, the model predicts individual photon peaks should not be visible.

Fitting this model to actual data is challenging. One can produce a master equation for the full Jaynes-Cummings system, allowing transitions between both qubit state and photon number in the resonator. The spectroscopic behavior could then be recovered by looking at the long-time response of the system to microwave drive at various frequencies. However, the low filter frequency ω_f suggests that many photon numbers will be relevant. Since the density matrix for an N-dimensional Hilbert space has N^2 elements, each time step would involve solving $O(1000)$ equations. For spectroscopic behavior, we are interested in long-time solutions. Accurate solutions will involve time steps small compared to the natural frequencies of the system; this means that getting to the spectroscopic limit will take thousands or tens of thousands of time steps. For doing spectroscopic fits, this task would be repeated for each frequency and each choice of parameters. All of this represents an infeasible computational task.

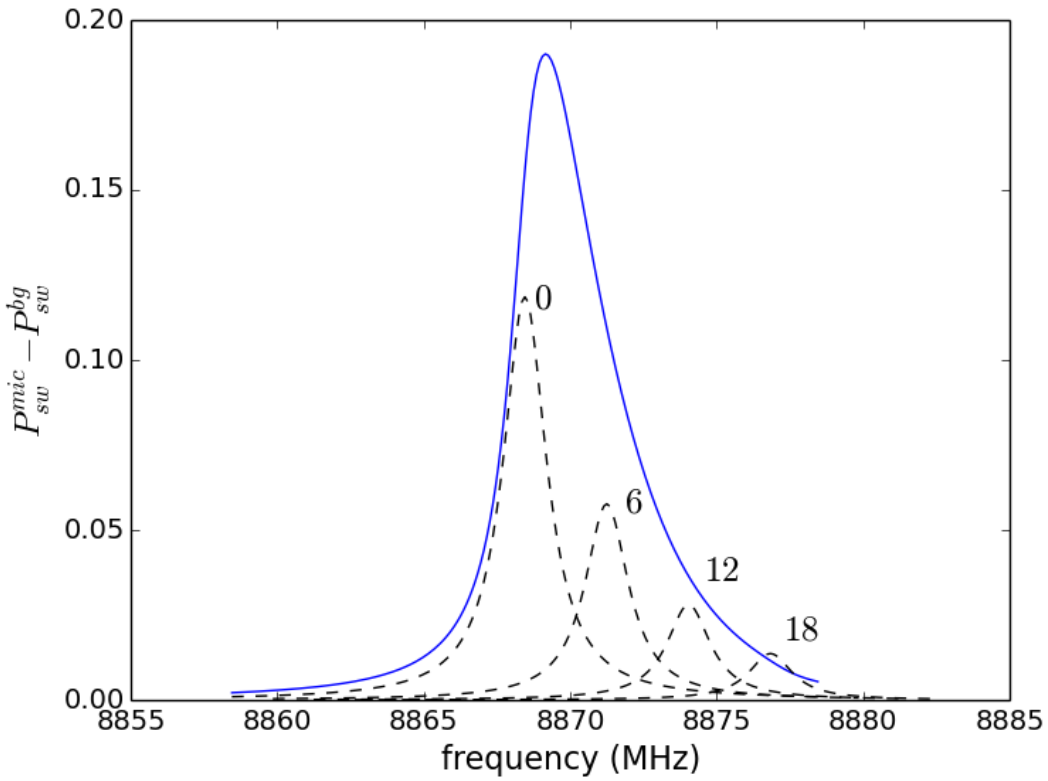

Figure 7.16: Typical lineshape predicted by the effective Jaynes-Cummings model. The blue curve is the sum of all photon peaks; the dotted black curves (blown up for clarity) are the individual photon peaks at 0, 6, 12 and 18 photons.

In the face of this complexity, I opted for a simpler approximation. I modeled the total spectroscopic response as a sum over discrete Lorentzians (see Equation 2.52). The location of the peaks, as well as the size of the photon shift, were determined by the effective Jaynes-Cummings model for BBC2. To determine the relative weights of each photon peak I assumed a thermal distribution of photon states.

Figure 7.16 shows a representative spectral line predicted by this spectroscopic model. The blue peak represents the output of the model; it is the sum of all the discrete photon peaks. The dashed black curves are individual photon peaks. I have increased the height of these peaks by a factor of 3 to make them more visible for this plot. The pictured photon peaks correspond to 0, 6, 12 and 18 photons.

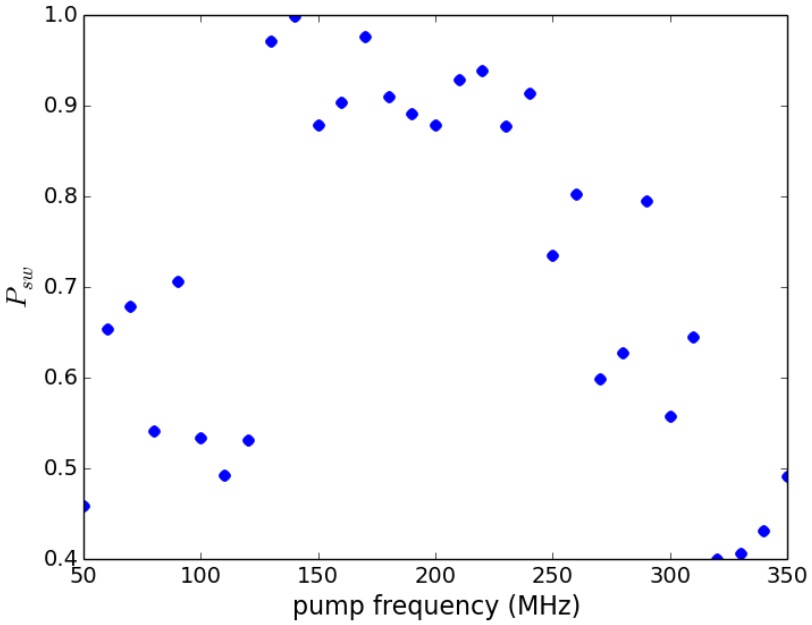

Figure 7.17: Tunneling probability during pulse as low frequency part of two-photon drive is tuned.

I truncate the model by including only photon states where the Boltzmann factor $e^{-\beta n\hbar\omega} > .01$. The particular cutoff is arbitrary.

7.3.2 The rest of the parameters

The I-Φ characteristic does a good job of establishing various inductances associated with the SQUID and the junction critical currents. While this covers most of the required parameters for the model, a few others must still be determined. The normal modes — responsible for the A and ω terms in (7.3) — depend on γ_i and γ_q. Since I know the current and flux being applied, I can use the I-Φ fits to determine a classical solution to the SQUID potential and estimate the phases. I also need to know the capacitance C_c of the IDC shunting the qubit junction, as well as the filter parameters L_f and C_f.

Figure 7.17 shows results from two-photon spectroscopy performed during the third cooldown. Two different frequency microwave drives are applied to the sys-

tem: a low frequency (50-350 MHz) pump intended to excite the filter; and a high frequency probe to drive qubit excitations. For this data, only the low frequency pump is on. The x-axis shows the pump frequency. The y-axis shows the probability of tunneling due to the pulse; note this is elevated compared to usual because this measurement was performed at ≈ 150 mK. The significant feature of this plot is the large region from 130-250 MHz where enhancement in the escape rate is seen. This frequency range is consistent with the design frequency of the filter, suggesting that the enhancement is due to resonant driving of the filter.

To further understand the behavior of the filter, I also studied the microwave response using a network analyzer. To get rid of the complication associated with transmission through the microwave line and most of the SQUID before reaching the filter, I measured S_{11} on the bias line. However, the heavy low-pass filtering on the bias line made attempts to measure this in the fridge impractical. Instead, I performed this measurement by attaching the sample box directly to one port of an Agilent vector network analyzer at room temperature. Ideally, this measurement would have been conducted at low temperature in order to reduce or eliminate the resistance of the various aluminum leads. However, since the geometry of the lines is well specified by the mask, the resistance of the filter wiring can be pretty well estimated using just the deposited thickness.

Figure 7.18 shows a schematic for the model I used to fit the S_{11} data, along with the data itself. In addition to the resistances, the model also includes inductance from the wiring leading up to C_f; this is separate from the inductance L_f following the filter that enters into the BBC2 circuit model. This extra inductance arises from the wirebonds connecting the chip to the printed circuit board in the sample box and from the relatively wide traces from the on-chip contact pads to C_f. The filter inductor uses comparatively long, narrow lines, leading to a significantly higher resistance that unfortunately makes this model not very sensitive to the pre-

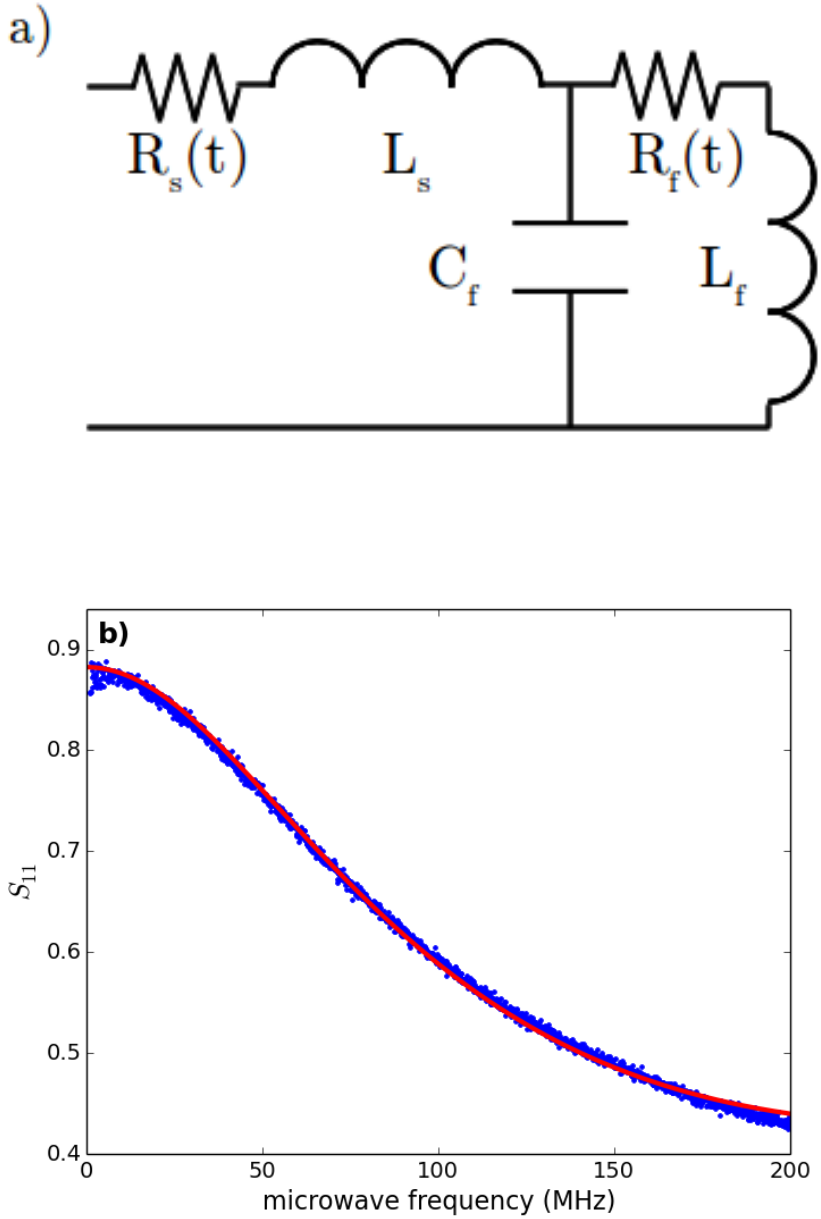

Figure 7.18: a) A schematic of the circuit model used for fitting the S_{11} data from BBC2. The fitting parameters are C_f, t, and L_s. b) The S_{11} data, along with best fit of the model with parameters $t = 70$ nm, $L_s = 11$ nH, $C_f = 33$ pH.

cise value of L_f used. Thus, in performing the fits, I used a fixed value of L_f at 15 nH, leaving me with three parameters: the thickness of deposited aluminum t, the stray inductance L_s, and the filter capacitance C_f.

The second plot in Figure 7.18 shows the S_{11} data along with the fit. This is just the amplitude of the signal reflected from the sample, referenced to the amplitude of the source. For an impedance Z_L hooked up to a $Z_0 = 50\ \Omega$ source, the reflection coefficient is [141]:

$$\Gamma = \frac{Z_L - Z_0}{Z_L + Z_0} \tag{7.4}$$

I made a least squares fit of Γ vs. frequency for the expected Z_L of the filter model pictured above. The resulting parameters were $t = 70$ nm, $L_s = 11$ nH, and $C_f = 33$ pF. The thickness and filter capacitance are both very plausible values. C_f is somewhat smaller than the design value of 50 pF, but the design came from a simple parallel plate estimate, and would also be sensitive to relatively small deviations in the planned thickness of the SiN dielectric.

The last parameter I need to estimate is the IDC capacitance C_c. I get a value for this by matching the predicted frequency of the qubit mode to the measured spectroscopy. Since f_{01} also depends on the junction phases γ_i and γ_q, the capacitance is harder to pin down. Given the estimates for the phases from the classical solution, I find $C_c \approx 1.6$pF. This is somewhat large compared to the design value of 1 pF. The discrepancy may come from stray capacitance to the surrounding ground plane, which does not enter into the microstrip formula (6.9).

7.3.3 Fits to spectroscopic data

While I succeeded in measuring most of the parameters associated with my Jaynes-Cummings spectroscopic model, a few free variables remained to fit:

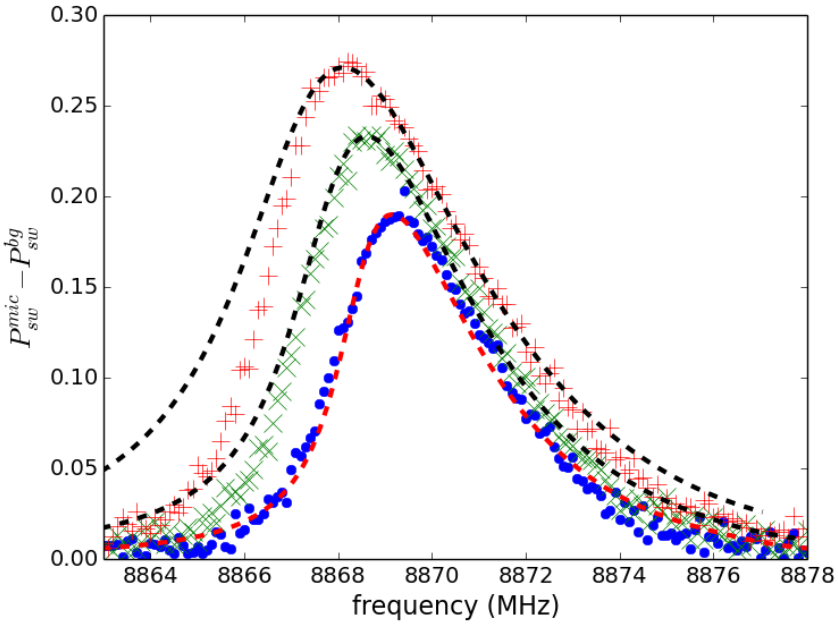

Figure 7.19: Data from Figure 7.10, along with temperature constrained fits to approximate spectroscopy from the effective Jaynes-Cummings model.

- the filter inductance L_f;
- the microwave power reaching the sample, expressed as the angular Rabi frequency at zero detuning Ω;
- the effective temperature of the system T;
- the linewidth of individual photon peaks.

The linewidth is the most complicated issue. A single spectrum can be fit with individual photon peaks assigned some width. For comparing spectra at different powers, it is desirable to have some way of accounting for power broadening of the peaks. Thus, I treat each photon peak according to the spectrum Equation 2.52 predicted by the Bloch equations. To simplify, I assume each photon peak has the same T_{1p} and T_{2p}. Since the T_ϕ I measured for BBC2 was so long, I make the further simplifying assumption that $T_{2p} = 2T_{1p}$. I then use T_{1p} as a fitting parameter.

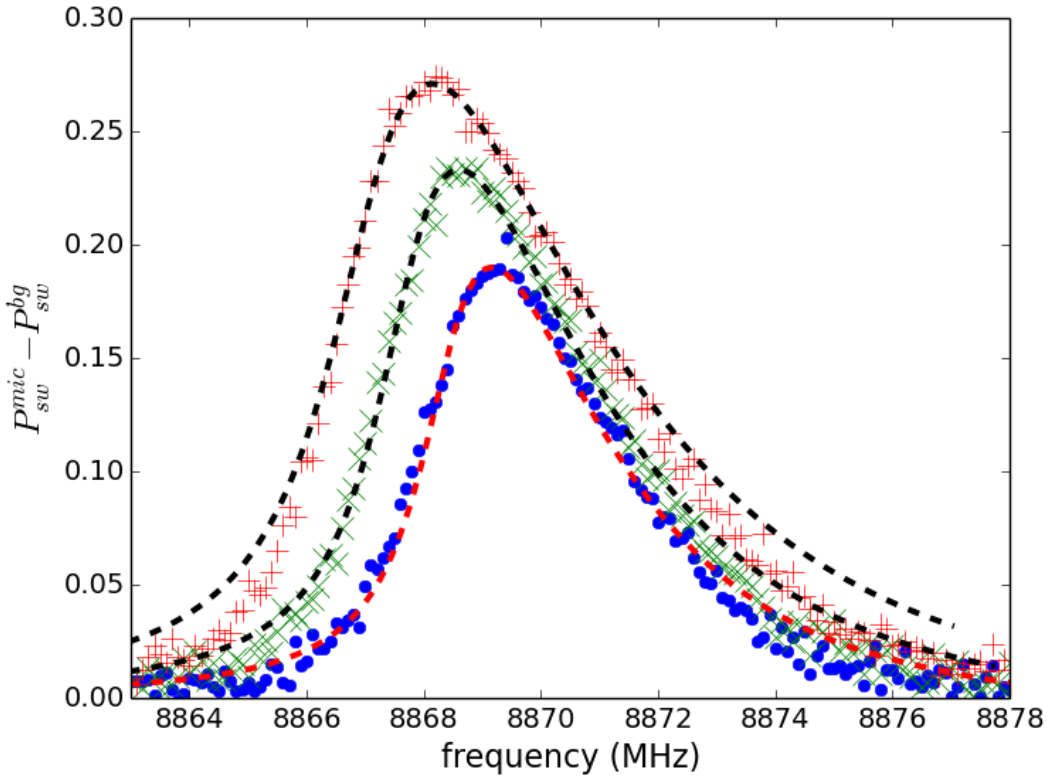

Figure 7.20: Similar to Figure 7.19, but now with temperature allowed to vary between the curves.

Figure 7.19 shows fits of my spectroscopic model to the data of Figure 7.10. The three data sets correspond to microwave powers of -28, -26 and -24 dBm at identical bias conditions. In fitting, I am only concerned with the shape of the spectral lines. The model I use does not capture the shift in frequency with power — since it is only a harmonic approximation — nor does it capture the relative heights of the three lines. I thus applied shifts of 0.7 and 1.4 MHz to the two higher power fits to line up with the peaks, and scaled the fit curves to agree at the peaks.

The lowest power peak was fit using $T = 95$ mK, $T_{1p} = 300$ ns, $\Omega/2\pi = 1.2$ MHz. The higher power peaks were then fit using the same temperature and 2 dB increases in Ω. The agreement on the low frequency side of the resonance is poor. The consistency of the shape of the low frequency side of the data with

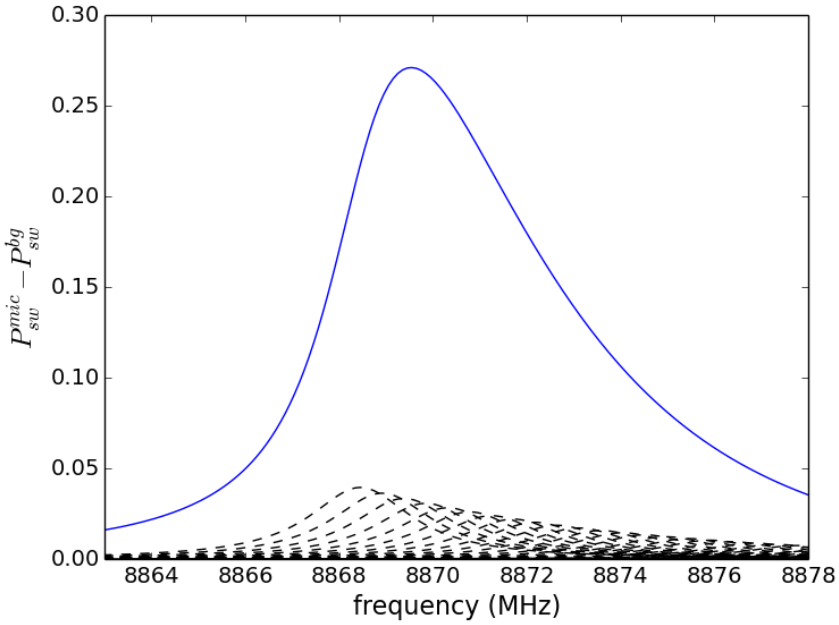

Figure 7.21: Fit to -24 dBm data from Figure 7.20. The blue curve is the sum of the black dotted curves, which are all of the individual photon peaks.

increasing power suggests that power broadening is not yet significant. This makes the increasing width of the peak difficult to capture at fixed temperature.

Increasing the microwave power might also cause heating of the sample. This effect can be included in the fit by allowing different temperatures for the three fits. Figure 7.20 shows the same data, but now with fits where the temperature is allowed to differ. The agreement is much better, suggesting the broadening of the peak with increasing power is primarily due to occupation of higher photon states. Again using a fit temperature of 95 mK for the lowest power, I use temperatures of 115 and 140 mK for the two higher microwave powers. The other fitting parameters were $L_f = 14$ nH, $T_{1p} = 100$ ns, $\Omega/2\pi = 600$ kHz. Compared to the fixed temperature fit, T_{1p} is relatively short, which reduces the significance of power broadening. The plot used in Figure 7.16 was generated using the fit to the -28 dBm data and shows scaled-up versions of the photon peaks at 0, 6, 12, and 18 photons to give a feel for how these peaks affect the lineshape. In Figure 7.21, the -24 dBm data is shown

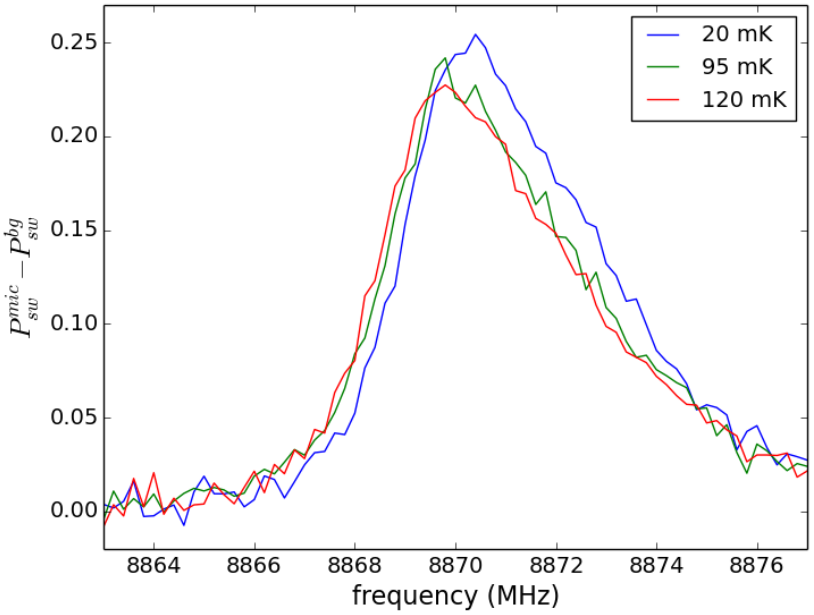

Figure 7.22: Spectroscopic data using -27 dBm microwaves at nominal mixing chamber temperatures 20 mK, 95 mK and 120 mK.

with all of its photon peaks.

As a plausibility check for these fit temperatures, I can look at spectroscopic data taken at fixed power and several different mixing chamber temperatures. Figure 7.22 plots three sets of spectroscopic data taken at 20 mK, 95 mK and 120 mK. These temperatures are the temperature reported by the mixing chamber thermometer. I took the data at elevated temperature by running a heater on the mixing chamber until the temperature roughly stabilized at the new value. The spectra do not appear to depend significantly on the temperature, arguing that an effective temperature above 100 mK is quite plausible.

The photon peak picture also offers a qualitative argument why the Rabi oscillations may be more visible on the blue detuned side of the resonance. Suppose that I am driving the system on the red detuned side of the spectral peak. There I am detuned from all of the individual photon peaks. The further from resonance, the more significant detuning becomes as part of the Rabi frequency. Thus, fluctuations

Parameter	Design value	Measured value
I_{0q} (μA)	.4	1.5
I_{0i} (μA)	7	26.8
L_q (nH)	2	1.4
L_i (nH)	.02	.22
L_f (nH)	16	14
C_q (pF)	1	1.63
C_i (pF)	.05	.1
C_f (pF)	50	33
M_f (pH)	1	15
M_μ (pH)	2	14
T_1 (ns)	320	215-230
T_ϕ (ns)	n/a	1100
$g/2\pi$ (MHz)	n/a	32

Table 7.1: Table of parameters for BBC2.

in the photon number can lead to relatively large shifts in the Rabi frequency. By contrast, if we operate on the blue detuned side of the resonance, then we will be very close to on resonance for one particular photon number. Small fluctuations in the photon number around this point will not produce significant change in the Rabi frequency. Thus, we might expect that the resulting Rabi oscillations would be less washed out.

7.4 Summary

I examined spectroscopy data and measured the coherence times of the BBC2 phase qubit. While the lifetime of 230 ns fell short of the design goal, it was close to the 320 ns estimate from the measured parameters. It is possible that uncertainty in the measured parameters can account for the discrepancy. Other loss mechanisms (e.g. coupling to modes of the sample box, thermal radiation) might also account for the shortfall. Table 7.1 contains the design values, measured device parameters and inferred coherence times of the qubit, as well as the estimated g coupling the

filter and SQUID.

Anomalies in the spectral data suggested the possibility of Jaynes-Cummings physics in BBC2. To address this, I developed the filter plus qubit model of Chapter 3 and applied it to our device. The expected coupling $g/2\pi \approx 32$ MHz is not sufficient to produce resolvable photon peaks at the level of the data I took. However, it does seem capable of explaining the asymmetric lineshape, and offers some hints to the strange dependence of the Rabi decay time T' on detuning.

The evidence for the significance of the effective Jaynes-Cummings model is hardly ironclad. In the following chapter, I will discuss some experiments that could assess the importance of the Jaynes-Cummings physics for these qubits. I will also examine some of the issues surrounding the pulsed measurement, a mystery now in several generations of dc SQUID phase qubits.

Chapter 8

Conclusions and future work

Over the years, the quest for improved T_1 in dc SQUID phase qubits has led to the LC filtered, asymmetric, small junction dc SQUID as the best technique thus far. Using these small junctions placed us in a regime where the dynamics of the full SQUID are relevant. The model I developed in Section 2.5 was successful in explaining the spectroscopic behavior of earlier dc SQUID phase qubits.

The goal for the BBC2 design was to reduce loss to the microwave line by providing a more controlled microwave architecture. Earlier qubit designs in our group coupled microwaves to the qubit using a weakly capacitively coupled microstrip. Another hope for BBC2 was to perform qubit measurements using the transmitted microwave signal. Neither of these goals were well realized in BBC2. The inductance of the microwave line was almost an order of magnitude larger than expected, and the lack of a cryogenic amplifier made transmission measurements infeasible at low microwave powers. Nevertheless, BBC2 did yield the longest Rabi oscillations of any of the dc SQUID phase qubits that our lab has studied. Table 8.1 compares BBC2 to several other dc SQUID qubits. T_1 is comparable, and the T_ϕ is excellent.

Table 8.1: Comparison of T_1 and T_ϕ to other dc SQUID phase qubits.

Qubit	T_1 (ns)	T_ϕ (ns)
Kwon [29]	280	90
Przybysz [28]	350	140
Lecocq et al. [41]	200	240
BBC2	215-230	1100

The BBC2 data showed that there is much about these dc SQUID qubits that remains unclear. The various microwave measurements yielded confusing results, containing features — frequency shifts with power, sidebands, coherence properties depending on detuning — that appear inconsistent with a basic two-state model. Given the previous success of the normal modes model, I developed a circuit model for BBC2 that included the on-chip LC filter. This model, presented in Chapter 3 and compared with experimental data in Chapter 7, makes some predictions that appear consistent with the data, but the agreement is not conclusive.

In this concluding chapter, I consider further work which might illumine the most significant questions raised by BBC2: whether the effective Jaynes-Cummings model is a good description of the qubit; and what is actually happening with the pulsed tunneling measurement.

8.1 Is Jaynes-Cummings important for LC filtered qubits?

While the complex spectroscopic data indicate the need for a more detailed model of the qubit, it is not clear that the Jaynes-Cummings model of Chapter 3 with the LC filter acting as resonator solves these issues. The lumped element effective Jaynes-Cummings model does suggest that, even if Jaynes-Cummings physics does not explain my data, it is still a relevant consideration for the BBC2 design. The coupling is significant enough, and the number of thermal photons expected large enough, that effects on the spectroscopy and Rabi data are to be expected. Prior to this work, our lab considered the filter to be a classical black box, improving lifetimes just by reducing $S_{II}(\omega_{01})$. A better understanding of when that is a good approximation could be useful for the design of other quantum devices.

There are several possible avenues for answering this question, both theoretical and experimental. One very natural approach is to consider replacing the lumped element model with a field analysis of the full circuit. This approach was recently

taken by the group at Yale with their "black box quantization" approach to studying the dynamics of transmon qubits in 3-D cavities [42]. Ironically, the intent of that paper was to move away from Jaynes-Cummings because of the complexities that arise in a multi-mode Jaynes-Cummings calculation. The idea is to extract the normal modes of the system by using a field solver to get the impedance of the circuit design — the modes then correspond to zeroes in the imaginary part of the admittance. Surprisingly, the frequency of the modes is not sensitive to which two points one finds the admittance between. As in my models, junction physics are included by making Taylor expansions of the $\cos(\gamma)$ term arising from the Josephson potential. Since transmons are unbiased, the cubic term vanishes and a quartic term remains.

There are advantages and disadvantages to the field analysis approach. It is a more exact description of the system than the lumped element approximation that I made. However, by comparison, it is far less physically transparent. The coupling behavior is fairly easy to extract from the lumped element picture. In the field model, it seems that we would have difficulties extracting g from the field model similar to those extracting L_μ from the existing simulations. This seems like a beatable problem, but it was one of the reasons that I chose to initially pursue the lumped element picture. One thing the field analysis offers that the lumped element model misses entirely is the possibility of things like $\lambda/4$ modes in the SQUID inductance (considered as a long microstrip) itself.

An area where the analysis could clearly be improved is in considering the response of the system to microwave drive. The model I used for comparisons with the spectroscopy data was very limited. It ignored all possibilities of transitions between photon number states. This seems like a severe omission, particularly in the context of coherent oscillations. As mentioned in Chapter 7, though, the large number of photon states needed renders direct calculation of the master equation

impractical.

One technique for improving the numerical efficiency is the quantum Monte Carlo trajectory approach [142]. Instead of tracking the entire density matrix, this approach (also referred to as just the quantum trajectory method) uses the master equation to get the probability of jumps between various states. At each time step, random numbers are generated to see if the wavefunction will jump between various states. By repeating this for many wavefunctions and averaging, the effective solution to the master equation can be derived. Since tracking the wavefunction in a dimension N Hilbert space only requires N variables rather than the N^2 of the density matrix, there can be considerable savings of time. This technique has been used fruitfully in circuit QED systems [143]; intriguingly, a recent experiment even suggests an experimental realization of such trajectories [144]. Implementing quantum trajectory techniques would be a natural way to examine the model more rigorously and study the microwave response.

Experimentally, the clearest way to demonstrate the correctness of the Jaynes-Cummings approach would be to engineer the system to make the photon states more apparent. I think there are two natural ways to increase the visibility of the photon features. The first is to increase the coupling strength g. When $L_f \gg L_q, L_i$, I can approximate the scaling of g from (3.54) as

$$g \propto \sqrt{\frac{1}{\omega_f}}\frac{1}{\mathfrak{L}^2} \approx L_f^{1/4} L_f^{-1} = L_f^{-3/4}. \tag{8.1}$$

Since the T_1 enhancement of the filter depends on the frequency, I can keep the ω_f roughly constant while increasing g by increasing C_f while proportionally decreasing L_f. Since the photon shift (3.11) is quadratic in g, relatively modest adjustments to the filter parameters could result in substantial increases in photon peak visibility.

Extending the idea further, one possibility would be to design a filter that

used Josephson junctions in the inductive section. Small SQUID loops (i.e. loops small enough so that the self inductance is much smaller than the Josephson inductance) can provide the effect of junctions with tunable critical current and therefore Josephson inductance. By adjusting the inductance, g could be tuned over a wide range without adjusting the qubit bias parameters. This would allow for a thorough, convincing demonstration of the photon peaks, as they could be moved around in a predictable fashion. This scheme does have serious drawbacks. To get good inductance out of the junctions requires many junctions or small junctions. Using many junctions introduces a greater chance of failure in the fabrication process; using small junctions makes the current bias necessary for pulsed measurement difficult to achieve.

While changing g is the most direct attack, another possibility is to increase the filter frequency. The exponential nature of the Boltzmann factor would lead to a big suppression in the number of significant photon states. This would certainly lead to easier modeling problems, and might also simplify the problem of seeing peaks at all. From Table 6.2, we see that the bias line is expected to be a very weak contributor to T_1. However, the dependence of T_1 on the filter frequency is quartic; with a tenfold increase in the filter frequency, loss to the bias line would be one of the most significant expected sources of loss. Moving ω_f up to 1 GHz would dramatically reduce the number of photon states while still giving good loss performance.

8.2 Anomalies in the tunneling measurement

One of the biggest shortcomings of our recent qubit studies is the challenge imposed by the pulsed tunneling measurement. As mentioned earlier, it imposes extra data taking, which hurts doubly because the experimental repetition rate is limited by the need for cooling following the time spent in the voltage state. Yet I

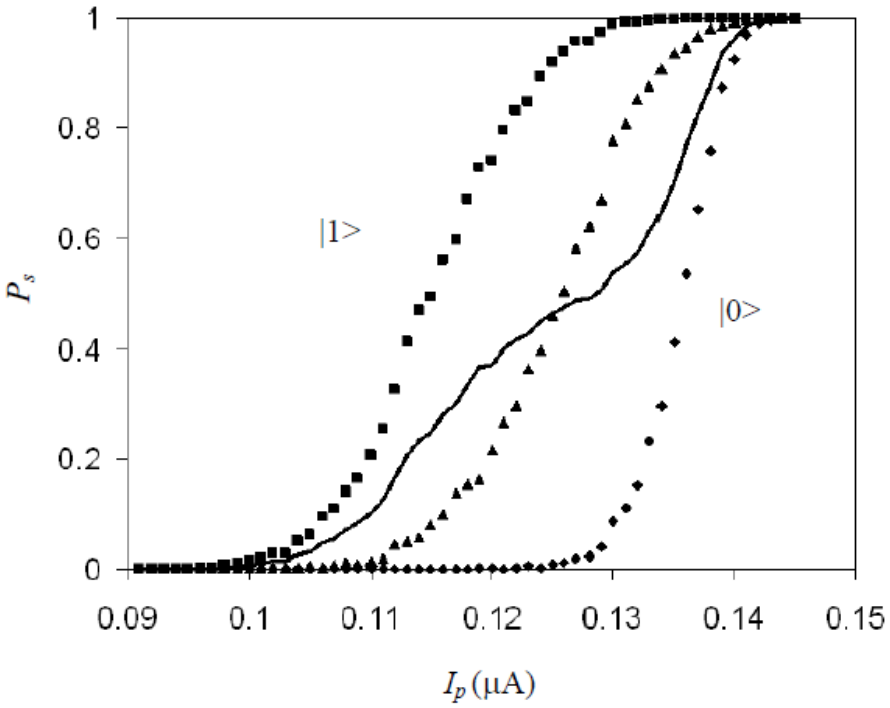

Figure 8.1: Plot of measured s-curves from Hyeokshin Kwon's thesis [29]. The measured s-curve in the superposition state differs from the expected behavior (solid line) of a measurement of the qubit state.

think some of the most interesting questions remaining in these devices come from the behavior of this measurement.

Figure 8.1 shows a plot of s-curves (for a refresher on these, see Section 7.2.1) from Hyeokshin Kwon's thesis [29]. The three collections of points are measured data. The circles are the background, assumed to be the $|0\rangle$ state. From Rabi oscillations, the point of maximum excitation is chosen and considered to be the best estimate of the $|1\rangle$ state, plotted with squares. The triangles are data corresponding to the microwaves being on half as long as needed to excite to the $|1\rangle$ state; this is expected to be an equal superposition of $|0\rangle$ and $|1\rangle$.

The solid line drawn on the plot is the average of the $|0\rangle$ and $|1\rangle$ s-curves. The tunneling measurement is supposed to be a measurement of the energy of the system, since that is what the tunneling rate depends on. Given an equal superposition of

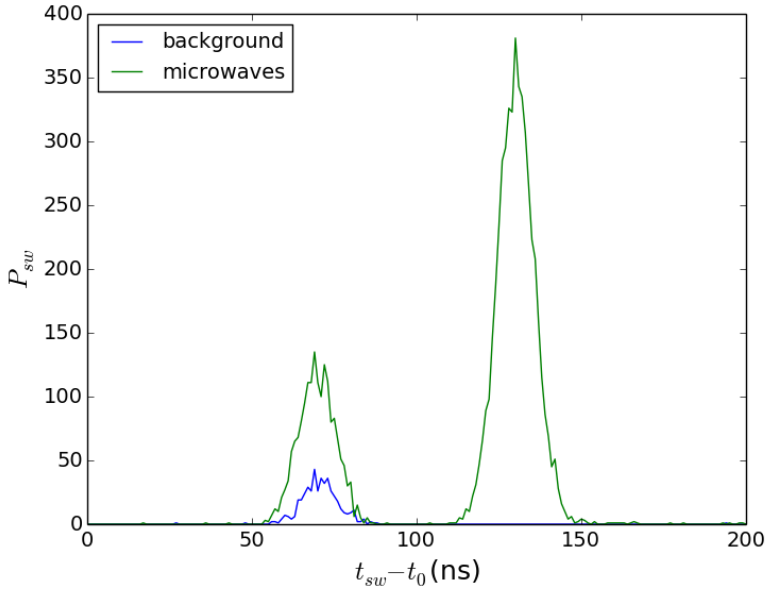

Figure 8.2: Measured switch probability P_{sw} versus time. When microwaves are turned on, an "echo" pulse occurs ≈ 60 ns after the first one.

0 and 1, we expect that the energy measurement would yield 0 half of the time and 1 half of the time, giving the solid line as a result. In fact, this clearly disagrees with the measured data. Despite using a tunneling measurement that seems like it should be an energy measurement, we are measuring something else. This behavior was found in the qubit designs following Kwon's measurements, including in BBC2.

BBC2 produced its own measurement surprises as well. During the second cooldown, there were occasionally some bizarre features in the Rabi and T_1 data. For most of the cooldown, I had been using 100 ns bins to collect the data into histograms. When I switched to using 1 ns bins, I discovered something surprising. Figure 8.2 shows some of this histogram data. The blue curve shows the background response with no microwaves; these escapes were what I used to set the pulse height. The green curve shows the tunneling events occurring when microwaves were on. Not only are there escapes at the location of the pulse itself, there are also escapes at a second time about 60 ns later. I will refer to these latter escapes as the echo and the former as the pulse.

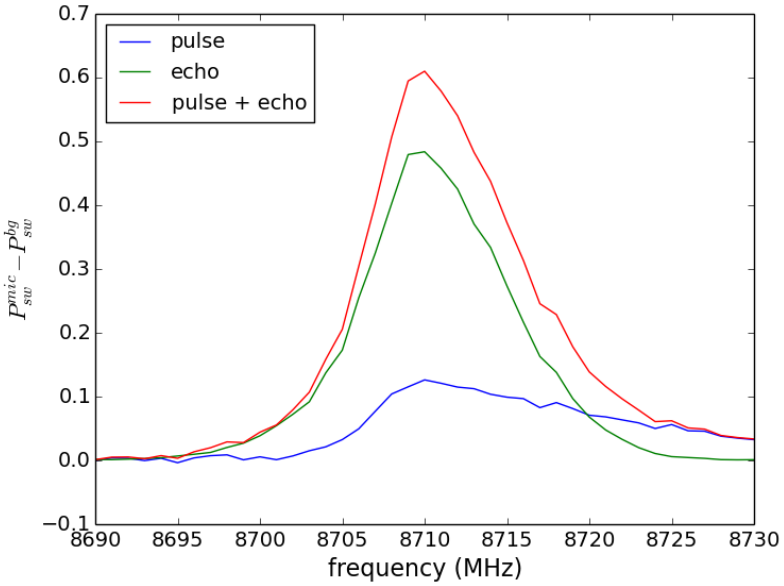

Figure 8.3: Spectroscopic measurement on BBC2 from the second cooldown, broken down into contributions from the pulse location and its echo. The data of Figure 8.2 was taken from this spectroscopy data at 8710 MHz.

Figure 8.3 shows spectroscopy data. The sum curve was what would have been measured with the coarse binning. The other two curves come from considering tunneling events at the pulse and echo locations. Interestingly, much of the asymmetry in the overall lineshape seems to arise from the behavior of the pulse data alone; the echo is much more symmetrically shaped. The spectroscopy data of Figure 7.10 came from considering the response at the pulse location. The sidebands of Figure 7.7 generally appeared to be more visible at the echo location.

During the third cooldown, I noticed that the relative importance of the pulse and echo locations could be adjusted by tuning the bias. Figure 8.4 shows some data illustrating this. Adjusting the bias changed how easy it was to produce tunneling events at both locations. To quantify this, I measured s-curves at many different bias locations, looking for how much voltage was necessary to tunnel half of the time at a given location. The color in these plots is that median value. The horizontal axis is flux; the vertical axis is current. The units of both axes indicate the bin (out

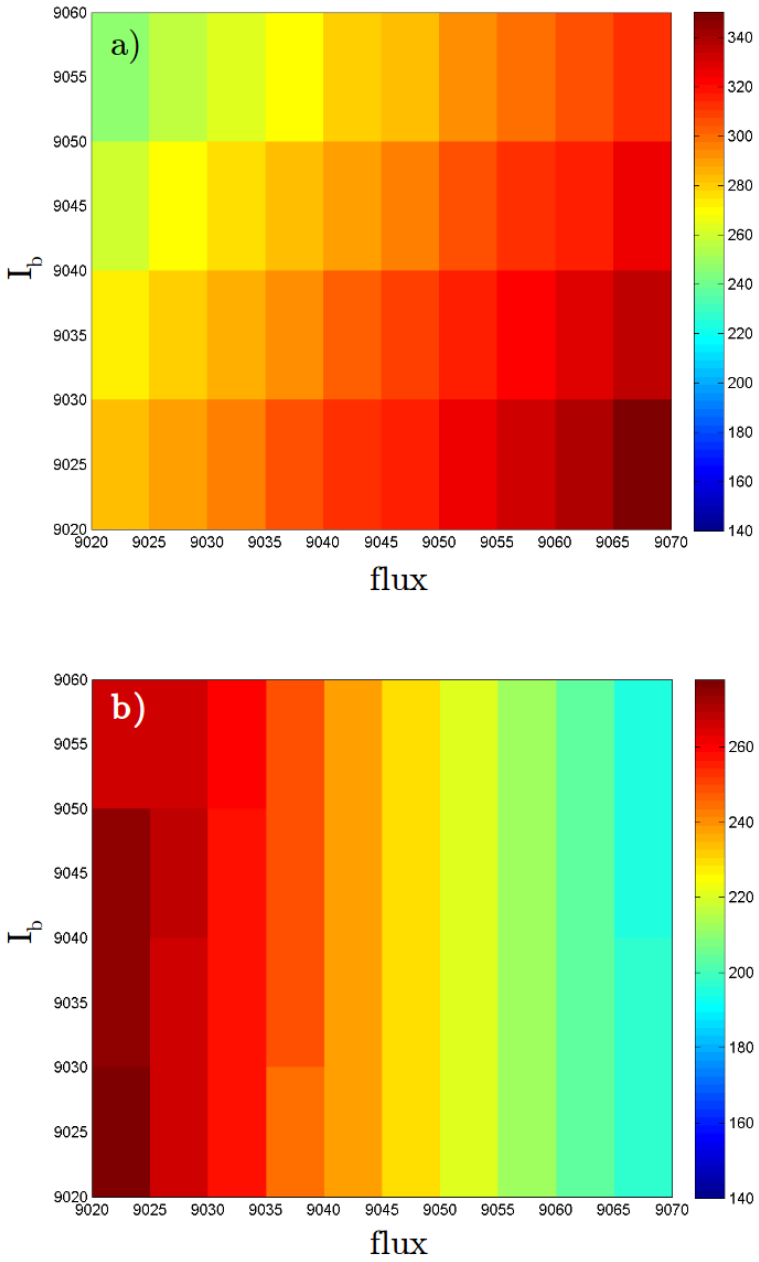

Figure 8.4: Color plots of pulse voltage required to tunnel half of the time vs. current and flux at a) the pulse location, b) the echo location. The units on the axes represent different bins in the current and flux waveforms where the plateau occurs.

of 16000 bins in the waveform) at which the flux or current stop ramping and start to plateau.

The plots show very interesting qualitative behavior. For the pulse location, increasing flux makes it harder to tunnel. Increasing bias current makes it easier to tunnel. By contrast, at the echo location, the tunneling is largely insensitive to the bias current while depending strongly on flux. At the echo, increasing flux makes it easier to tunnel. This behavior is very reminiscent of the intended behavior of the double ramp — applying flux to bias the qubit junction while counteracting the increasing bias current on the isolation junction.

This suggests the idea that the two pulse locations somehow correspond to tunneling dominated by one junction or the other. It is difficult to interpret the data in light of this idea, though. Given that the mode I wished to study is dominated by the qubit junction, one might think that the echo location is the proper one to examine. But, possibly as an artifact of this delay, there are strange aspects to the Rabi and T_1 data corresponding to the echo location. There is also the question of what the delay itself means.

Clearly there are many outstanding issues associated with the pulsed tunneling measurement. Since these issues first became noticeable with Hyeokshin's data — corresponding to when the normal mode analysis of Chapter 2 first became significant — I believe the dynamics of the full circuit are an important piece of the puzzle. I can see some physical arguments for why this might be. The measurement pulse can twist the orientation of the well in addition to tilting it. As Figure 2.12 shows, the expected wavefunctions follow the contours of the well. Thus it is possible that the energy "eigenbasis" (these are resonances, not eigenstates) differs considerably from the original state at the point where the measurement is strongest. A paper by Ashhab et al. [145] explored the implications of measurements that do not commute with the Hamiltonian and found that in certain regimes the measurement result

looked much more like our data. Since the pulse is sufficiently slow (i.e. it lasts for several dozen cycles at the qubit frequency) a sudden approximation is likely invalid; a more careful consideration of the dynamics of the SQUID through the pulse seems worthwhile.

The development of the effective Jaynes-Cummings model raises another question: is the tunneling measurement sensitive to the number of photons? The total energy of the filter plus SQUID system increases as more photons are added. Since the two parts are coupled, it is conceivable that having more photons in the filter could lead to enhanced tunneling from the SQUID. The effect would likely be small — it must depend on g, and dimensional analysis suggests that most likely g/ω_{01} would be the order of the scaling factor associated with extra energy coming from the filter. Since each photon is already about 1/40 the energy of the $|0\rangle$ to $|1\rangle$ transition, this further scaling would lead to modest shifts in the escape rate. Still, I believe this idea deserves more examination.

Given these arguments, I think developing a model of the pulsed measurement in the full model of Chapter 3 would be worthwhile. However, this is a formidable problem. The 1D approach to the measurement uses the tunneling rate extracted from a WKB approximation. The relative simplicity of WKB in 1D arises from having to match wavefunctions at points. In higher dimensions, this matching is over a correspondingly higher dimensional object.

As such, generalizing WKB to multiple dimensions is non-trivial. Caldeira and Leggett [53] use an "instanton" technique adapted from quantum field theory [146]. Schmid found an adaptation of the semiclassical approach [147] that he claimed reproduced much of the instanton behavior; this was heavily influenced by a paper of Banks, Bender and Wu [148] based on the idea of reducing the problem to 1D through the use of a "most probable escape path" in configuration space.

While these approaches seem interesting, they are sufficiently theoretically

involved that I was looking for something more straightforward. Further, since the instantaneous pulse approximation seems poor, some sort of numerical technique seems desirable to study how the SQUID responds during the pulse. The most obvious thing to try is a direct time stepping of the Hamiltonian. However, this is likely to be extremely expensive. Except at very high bias, the lifetime is very long compared even to the period of the qubit. Higher energy states, which the qubit must necessarily evolve into in order to escape, will require even finer timesteps. For this reason, I have not pursued this strategy.

In his thesis [51] Fred Strauch explored the technique of complex scaling (reviewed in [149]). He found that it worked well in reproducing the results of other methods in the case of the 1D washboard. The lifetime of a resonance can be seen as an imaginary part of the energy. As the state evolves in time according to $\exp(iEt/\hbar)$, this imaginary part leads to a decay in the probability. In complex scaling, the coordinates and their momenta are rotated into the complex plane in such a way that the commutator is preserved. Surprisingly, the eigenvalues of the Hamiltonian — including imaginary parts — remain unchanged by this operation. Numerically solving this transformed Hamiltonian then gives the lifetime of each resonance.

This approach sounds very concrete and attractive. In practice, it is still quite complicated. Experimenting on a complex scaled simple harmonic oscillator shows that a relatively large basis is needed to keep even the lowest eigenvalues constant following the scaling procedure. The harmonic basis appears particularly ill suited for capturing tunneling behavior; this makes sense, since it would take very high harmonic levels to get reasonable occupation on the outside of the barrier. A better approach might be to use a discretized position basis, corresponding to delta function-like occupation of the various γ values. The potential operator is diagonal in this basis, whereas the kinetic energy comes from discretizing $\frac{\partial}{\partial\gamma}$. This can be

done to arbitrary order [150], although accuracy to the grid size $d\gamma^4$ is probably fine.

A deeper conceptual issue comes from the multi-well nature of the SQUID. For a range of bias conditions, a state that is localized in a particular well could escape through a barrier either to the voltage state or to a more stable flux well in the SQUID. If the numerical analysis doesn't measure this distinction, then it will give misleading comparisons to our experimental data. While one can exploit flux-state hopping for measurement (e.g. this is how rf SQUIDs are measured), it would require a different set of techniques. Only tunneling to the voltage state registers as a measurement event in our experiments.

8.3 Conclusion

The increasing complexity of both the operation and modeling of dc SQUID phase qubits suggests a need for very broad characterization of prospective devices. This is especially problematic given the limitations of the pulsed switching measurement. Even if the measurement worked completely as expected, the cooling period required following a voltage tunneling event typically limits experimental repetition rates to several hundred Hz. This is orders of magnitude slower than what is possible with dispersive microwave measurements. In practice, the pulsed measurement itself seems to be doing something more complicated than just measuring the state of the qubit, requiring further characterization and better understanding. All of these combine to make exploration of dc SQUID phase qubits a slow process. Unfortunately, the reward for this exploration is also questionable since dc SQUID qubits have coherence times which have not been competitive with other superconducting qubit approaches. Transmons in particular now routinely achieve lifetimes in the tens of nanoseconds. For these reasons, I feel that it is difficult for the dc SQUID to compete with other superconducting qubits.

Nevertheless, I think the questions raised in this chapter are interesting and applicable beyond the context of dc SQUID phase qubits. To me the most interesting aspect of quantum information is how it makes concrete ideas like measurement that were once highly abstract. The advances in quantum control of superconducting circuits during my tenure in this lab have been extremely impressive. Addressing vagaries in the measurement process, the complexities inherent in real devices, and figuring out how seemingly innocuous things like filters might affect the quantum behavior, can only lead to a better sense of what quantum tricks are possible and what is in store for the future.

Appendix A

The Bloch equation

A.1 Two-level system

The Hamiltonian for a driven qubit is

$$\mathcal{H} = \mathcal{H}_0 + V = \begin{pmatrix} E_0 & 0 \\ 0 & E_1 \end{pmatrix} + \begin{pmatrix} 0 & \hbar\Omega\cos(\omega t) \\ \hbar\Omega\cos(\omega t) & 0 \end{pmatrix}. \tag{A.1}$$

For a two-level system, the Lindblad operators are

$$\begin{aligned} \hat{L}_- &= \sqrt{\Gamma_1(1+n)}\sigma_-; \\ \hat{L}_+ &= \sqrt{\Gamma_1 n}\sigma_+, \end{aligned} \tag{A.2}$$

where $\Gamma_1 = 1/T_1$, and n is the Boltzmann factor for the two-level system. This thermal contribution comes from the trace over the bath and is required for the master equation to produce a thermal equilibrium in the case with no drive. In practice, experimental temperatures are usually far below the equivalent temperature of ω_{01}, so approximating the system as being at zero temperature and keeping only $\hat{L}_-$ is reasonable. After evaluating the Lindblad term, the master equation becomes

$$\frac{d\rho}{dt} = -\frac{\imath}{\hbar}[\mathcal{H}, \rho] + \Gamma_1 \begin{pmatrix} \rho_{11} & -\frac{1}{2}\rho_{01} \\ -\frac{1}{2}\rho_{10} & -\rho_{11} \end{pmatrix}. \tag{A.3}$$

Interestingly, even though I have only added dissipation so far, the off diagonal elements of the density matrix are seen to have a decay as well, although at half the rate. One can include a pure dephasing term as well, leading to an overall dephasing timescale

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi}. \tag{A.4}$$

Including this pure dephasing and rewriting with $\Gamma_2 = 1/T_2$, I get

$$\frac{d\rho}{dt} = -\frac{\imath}{\hbar}[\mathcal{H}, \rho] + \begin{pmatrix} \Gamma_1\rho_{11} & -\Gamma_2\rho_{01} \\ -\Gamma_2\rho_{10} & -\Gamma_1\rho_{11} \end{pmatrix}. \tag{A.5}$$

To simplify the equation, we can move to the interaction picture. Classically, this is analogous to moving to a frame that rotates at ω_{01}. I do this by applying the transform

$$\chi = e^{\frac{\imath\mathcal{H}_0 t}{\hbar}} \rho e^{-\frac{\imath\mathcal{H}_0 t}{\hbar}}; \tag{A.6}$$

differentiating this yields

$$\begin{aligned} \frac{d\chi}{dt} &= e^{\frac{\imath\mathcal{H}_0 t}{\hbar}} \frac{d\rho}{dt} e^{-\frac{\imath\mathcal{H}_0 t}{\hbar}} + \frac{\imath}{\hbar} e^{\frac{\imath\mathcal{H}_0 t}{\hbar}} [H_0, \rho] e^{-\frac{\imath\mathcal{H}_0 t}{\hbar}}, \\ &= -\frac{\imath}{\hbar} e^{\frac{\imath\mathcal{H}_0 t}{\hbar}} [V, \rho] e^{-\frac{\imath\mathcal{H}_0 t}{\hbar}} + e^{\frac{\imath\mathcal{H}_0 t}{\hbar}} \begin{pmatrix} \Gamma_1\rho_{11} & -\Gamma_2\rho_{01} \\ -\Gamma_2\rho_{10} & -\Gamma_1\rho_{11} \end{pmatrix} e^{-\frac{\imath\mathcal{H}_0 t}{\hbar}}, \\ &= -\frac{\imath}{\hbar}[V_I, \chi] + D_I, \end{aligned} \tag{A.7}$$

where V_I is the interaction representation of the drive term:

$$V_I = \begin{pmatrix} 0 & e^{-\imath\omega_{01}t}\hbar\Omega\cos(\omega t) \\ e^{\imath\omega_{01}t}\hbar\Omega\cos(\omega t) & 0 \end{pmatrix}, \tag{A.8}$$

and D_I is the interaction representation of the dissipator:

$$\begin{aligned} D_I &= \begin{pmatrix} \Gamma_1\rho_{11} & -e^{-\imath\omega_{01}t}\Gamma_2\rho_{01} \\ -e^{\imath\omega_{01}t}\Gamma_2\rho_{10} & -\Gamma_1\rho_{11} \end{pmatrix} \\ &= \begin{pmatrix} \Gamma_1\chi_{11} & -\Gamma_2\chi_{01} \\ -\Gamma_2\chi_{10} & -\Gamma_1\chi_{11} \end{pmatrix}. \end{aligned} \tag{A.9}$$

The next step is to make the rotating wave approximation (RWA):

$$e^{-\imath\omega_{01}t}\cos(\omega t) = \frac{1}{2}(e^{-\imath(\omega_{01}-\omega)t} + e^{-\imath(\omega_{01}+\omega)t}) \approx \frac{1}{2}e^{-\imath(\omega_{01}-\omega)t}, \tag{A.10}$$

where we assume that we are driving near resonance and can thus ignore the effects associated with the part of cos that is counterrotating at very high frequency. This amounts to replacing the cosine terms with exponentials at the drive frequency. If I denote the detuning $\Delta = \omega - \omega_{01}$ then evaluating the commutator yields

$$[V_I, \chi] = \frac{\hbar\Omega}{2}\begin{pmatrix} e^{\imath\Delta t}\chi_{10} - e^{-\imath\Delta t}\chi_{01} & e^{\imath\Delta t}(\chi_{11} - \chi_{00}) \\ e^{-\imath\Delta t}(\chi_{00} - \chi_{11}) & e^{-\imath\Delta t}\chi_{01} - e^{\imath\Delta t}\chi_{10} \end{pmatrix}; \tag{A.11}$$

including the dissipator, the full equation becomes:

$$\frac{d\chi}{dt} = \begin{pmatrix} -\imath\frac{\Omega}{2}(e^{\imath\Delta t}\chi_{10} - e^{-\imath\Delta t}\chi_{01}) + \Gamma_1\chi_{11} & -\imath\frac{\Omega}{2}e^{\imath\Delta t}(\chi_{11} - \chi_{00}) - \Gamma_2\chi_{01} \\ -\imath\frac{\Omega}{2}e^{-\imath\Delta t}(\chi_{00} - \chi_{11}) - \Gamma_2\chi_{10} & -\imath\frac{\Omega}{2}(e^{-\imath\Delta t}\chi_{01} - e^{\imath\Delta t}\chi_{10}) - \Gamma_1\chi_{11} \end{pmatrix}. \tag{A.12}$$

One final simplification amounts to moving to a frame rotating at the drive frequency

instead of the qubit frequency:

$$\begin{pmatrix} \chi_{00}^* \\ \chi_{01}^* \\ \chi_{10}^* \\ \chi_{11}^* \end{pmatrix} = \begin{pmatrix} \chi_{00} \\ e^{-\imath\Delta t}\chi_{01} \\ e^{\imath\Delta t}\chi_{10} \\ \chi_{11} \end{pmatrix}. \tag{A.13}$$

Then the master equation becomes:

$$\frac{d\chi^*}{dt} = \begin{pmatrix} -\imath\frac{\Omega}{2}(\chi_{10}^* - \chi_{01}^*) + \Gamma_1\chi_{11}^* & -\imath\frac{\Omega}{2}(\chi_{11}^* - \chi_{00}^*) - (\Gamma_2 + \imath\Delta)\chi_{01}^* \\ -\imath\frac{\Omega}{2}(\chi_{00}^* - \chi_{11}^*) - (\Gamma_2 - \imath\Delta)\chi_{10}^* & \imath\frac{\Omega}{2}(\chi_{10}^* - \chi_{01}^*) - \Gamma_1\chi_{11}^* \end{pmatrix}. \tag{A.14}$$

Bibliography

[1] R. P. Feynman, International Journal of Theoretical Physics **21**, 467 (1982).

[2] S. Wiesner, SIGACT News **15**, 78 (1983).

[3] D. Deutsch and R. Jozsa, Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences **439**, 553 (1992).

[4] C. H. Bennett and S. J. Wiesner, Physical Review Letters **69**, 2881 (1992).

[5] P. W. Shor, SIAM Journal on Computing **26**, 1484 (1997).

[6] M. A. Schlosshauer, *Decoherence: and the quantum-to-classical transition* (Springer, Jan. 2007).

[7] P. W. Shor, Physical Review A **52**, R2493 (1995).

[8] B. E. Kane, Nature **393**, 133 (1998).

[9] E. Knill, R. Laflamme, and G. J. Milburn, Nature **409**, 46 (2001).

[10] D. G. Cory, A. F. Fahmy, and T. F. Havel, Proceedings of the National Academy of Sciences **94**, 1634 (1997).

[11] S. Das Sarma, M. Freedman, and C. Nayak, Physical Review Letters **94**, 166802 (2005).

[12] J. I. Cirac and P. Zoller, Physical Review Letters **74**, 4091 (1995).

[13] J. M. Martinis, M. H. Devoret, and J. Clarke, Physical Review Letters **55**, 1543 (1985).

[14] R. Ramos, M. Gubrud, A. J. Berkley, J. R. Anderson, C. Lobb, and F. Wellstood, IEEE Transactions on Applied Superconductivity **11**, 998 (2001).

[15] Y. Nakamura, Y. A. Pashkin, and J. S. Tsai, Nature **398**, 786 (1999).

[16] J. R. Friedman, V. Patel, W. Chen, S. K. Tolpygo, and J. E. Lukens, Nature **406**, 43 (2000).

[17] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Physical Review A **76**, 042319 (2007).

[18] D. P. DiVincenzo and IBM, arXiv:quant-ph/0002077 (2000).

[19] R. P. Feynman and A. J. G. Hey, *Feynman lectures on computation* (Westview Press, July 2000).

[20] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge University Press, Oct. 2000).

[21] F. W. Strauch, P. R. Johnson, A. J. Dragt, C. J. Lobb, J. R. Anderson, and F. C. Wellstood, Physical Review Letters **91**, 167005 (2003).

[22] D. Gottesman, arXiv:0904.2557 (2009).

[23] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Physical Review A **86**, 032324 (2012).

[24] J. M. Martinis, S. Nam, J. Aumentado, and C. Urbina, Physical Review Letters **89**, 117901 (2002).

[25] S. K. Dutta, "Characterization of Josephson devices for use in quantum computation", Ph.D. (University of Maryland, College Park, United States – Maryland, 2006).

[26] H. Paik, "Coherence in dc SQUID phase qubits", Ph.D. (University of Maryland, College Park, United States – Maryland, Sept. 2007).

[27] T. A. Palomaki, "Dc SQUID phase qubit", Ph.D. (University of Maryland, College Park, United States – Maryland, 2008).

[28] A. J. Przybysz, "Reducing decoherence in dc SQUID phase qubits", Ph.D. (University of Maryland, College Park, United States – Maryland, 2010).

[29] H. Kwon, "Quantum coherent dynamics in a DC SQUID phase qubit using an LC filter", Ph.D. (University of Maryland, College Park, United States – Maryland, 2011).

[30] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, Physical Review A **69**, 062320 (2004).

[31] H. Paik and K. D. Osborn, Applied Physics Letters **96**, 072505 (2010).

[32] M. Khalil, F. Wellstood, and K. Osborn, IEEE Transactions on Applied Superconductivity **21**, 879 (2011).

[33] B. K. Cooper, R. P. Budoyo, V. Zaretskey, C. J. Ballard, J. R. Anderson, C. J. Lobb, and F. C. Wellstood, in Proc. SPIE Quantum Information and Computation XI, Vol. 8749 (2013), 87490N.

[34] H. Kwon, A. J. Przybysz, B. K. Cooper, H. Paik, K. D. Osborn, B. S. Palmer, R. Budoyo, J. R. Anderson, C. J. Lobb, and F. C. Wellstood, IEEE Transactions on Applied Superconductivity **21**, 860 (2011).

[35] M. H. Devoret, "Quantum fluctuations in electrical circuits", in *Quantum fluctuations*, edited by S. Reynaud, E. Giacobino, and J. Zinn-Justin (Elsevier, June 1997).

[36] H. B. Callen and T. A. Welton, Physical Review **83**, 34 (1951).

[37] K. Mitra, F. W. Strauch, C. J. Lobb, J. R. Anderson, F. C. Wellstood, and E. Tiesinga, Physical Review B **77**, 214512 (2008).

[38] B. K. Cooper, *Two-junction effects in dc SQUID phase qubit*, APS March Meeting, Mar. 2011.

[39] F. Lecocq, J. Claudon, O. Buisson, and P. Milman, Physical Review Letters **107**, 197002 (2011).

[40] E. Hoskinson, F. Lecocq, N. Didier, A. Fay, F. W. J. Hekking, W. Guichard, O. Buisson, R. Dolata, B. Mackrodt, and A. B. Zorin, Physical Review Letters **102**, 097004 (2009).

[41] F. Lecocq, I. M. Pop, I. Matei, E. Dumur, A. K. Feofanov, C. Naud, W. Guichard, and O. Buisson, Physical Review Letters **108**, 107001 (2012).

[42] S. E. Nigg, H. Paik, B. Vlastakis, G. Kirchmair, S. Shankar, L. Frunzio, M. H. Devoret, R. J. Schoelkopf, and S. M. Girvin, Physical Review Letters **108**, 240502 (2012).

[43] C. C. Grimes, P. L. Richards, and S. Shapiro, Physical Review Letters **17**, 431 (1966).

[44] R. A. Ferrell and R. E. Prange, Physical Review Letters **10**, 479 (1963).

[45] R. P. Feynman, R. B. Leighton, and M. L. Sands, *Feynman lectures on physics 3: quantum mechanics* (Basic Books, Oct. 2011).

[46] D. E. McCumber, Journal of Applied Physics **39**, 3113 (1968).

[47] W. C. Stewart, Applied Physics Letters **12**, 277 (1968).

[48] A. Barone and G. Paternò, *Physics and applications of the Josephson effect* (Wiley, 1982).

[49] V. Ambegaokar and A. Baratoff, Physical Review Letters **10**, 486 (1963).

[50] P. W. Anderson, in *Lectures on the many-body problem. 2*, edited by E. R. Caianello (Acad. Press, 1964).

[51] F. W. Strauch, "Theory of superconducting phase qubits", Ph.D. (University of Maryland, College Park, United States – Maryland, 2004).

[52] A. J. Leggett, Journal de Physiques Colloques **39**, 1264 (1978).

[53] A. Caldeira and A. Leggett, Annals of Physics **149**, 374 (1983).

[54] H. Kramers, Physica **7**, 284 (1940).

[55] R. F. Voss and R. A. Webb, Physical Review Letters **47**, 265 (1981).

[56] C. P. Slichter, *Principles of magnetic resonance* (Springer, 1990).

[57] M. Steffen, J. M. Martinis, and I. L. Chuang, Physical Review B **68**, 224518 (2003).

[58] B. Yurke and J. S. Denker, Physical Review A **29**, 1419 (1984).

[59] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to algorithms* (Mit Press, 2009).

[60] B. S. Deaver and W. M. Fairbank, Physical Review Letters **7**, 43 (1961).

[61] R. Doll and M. Näbauer, Physical Review Letters **7**, 51 (1961).

[62] R. C. Jaklevic, J. Lambe, A. H. Silver, and J. E. Mercereau, Physical Review Letters **12**, 159 (1964).

[63] A. H. Silver and J. E. Zimmerman, Physical Review Letters **15**, 888 (1965).

[64] C. D. Tesche and J. Clarke, Journal of Low Temperature Physics **29**, 301 (1977).

[65] H. Goldstein, C. P. Poole, and J. L. Safko, *Classical mechanics* (Addison-Wesley Longman, Incorporated, 2002).

[66] H. P. Breuer and F. Petruccione, *The theory of open quantum systems* (Oxford University Press, 2002).

[67] C. Gardiner and P. Zoller, *Quantum noise: a handbook of markovian and non-markovian quantum stochastic methods with applications to quantum optics* (Springer, Oct. 2004).

[68] W. H. Louisell, *Quantum statistical properties of radiation* (John Wiley & Sons Canada, Limited, 1973).

[69] N. G. V. Kampen, *Stochastic processes in physics and chemistry* (Elsevier, Aug. 2011).

[70] F. Bloch, Physical Review **70**, 460 (1946).

[71] H. Xu, "Quantum computing with Josephson junction circuits", Ph.D. (University of Maryland, College Park, United States – Maryland, 2004).

[72] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf, Reviews of Modern Physics **82**, 1155 (2010).

[73] S. K. Dutta, F. W. Strauch, R. M. Lewis, K. Mitra, H. Paik, T. A. Palomaki, E. Tiesinga, J. R. Anderson, A. J. Dragt, C. J. Lobb, and F. C. Wellstood, Physical Review B **78**, 104510 (2008).

[74] D. J. Van Harlingen, T. L. Robertson, B. L. T. Plourde, P. A. Reichardt, T. A. Crane, and J. Clarke, Physical Review B **70**, 064517 (2004).

[75] R. W. Simmonds, K. M. Lang, D. A. Hite, S. Nam, D. P. Pappas, and J. M. Martinis, Physical Review Letters **93**, 077003 (2004).

[76] J. M. Martinis, K. B. Cooper, R. McDermott, M. Steffen, M. Ansmann, K. D. Osborn, K. Cicak, S. Oh, D. P. Pappas, R. W. Simmonds, and C. C. Yu, Physical Review Letters **95**, 210503 (2005).

[77] W. A. Phillips, Reports on Progress in Physics **50**, 1657 (1987).

[78] M. Steffen, M. Ansmann, R. McDermott, N. Katz, R. C. Bialczak, E. Lucero, M. Neeley, E. M. Weig, A. N. Cleland, and J. M. Martinis, Physical Review Letters **97**, 050502 (2006).

[79] I. Siddiqi, R. Vijay, F. Pierre, C. M. Wilson, M. Metcalfe, C. Rigetti, L. Frunzio, and M. H. Devoret, Physical Review Letters **93**, 207002 (2004).

[80] R. Vijay, M. H. Devoret, and I. Siddiqi, Review of Scientific Instruments **80**, 111101 (2009).

[81] D. Vion, A. Aassime, A. Cottet, P. Joyez, H. Pothier, C. Urbina, D. Esteve, and M. H. Devoret, Science **296**, 886 (2002).

[82] J. D. Joannopoulos, P. R. Villeneuve, and S. Fan, Nature **386**, 143 (1997).

[83] I. Bloch, J. Dalibard, and W. Zwerger, Reviews of Modern Physics **80**, 885 (2008).

[84] C. Chin, R. Grimm, P. Julienne, and E. Tiesinga, Reviews of Modern Physics **82**, 1225 (2010).

[85] S. O. Valenzuela, W. D. Oliver, D. M. Berns, K. K. Berggren, L. S. Levitov, and T. P. Orlando, Science **314**, 1589 (2006).

[86] O. Astafiev, K. Inomata, A. O. Niskanen, T. Yamamoto, Y. A. Pashkin, Y. Nakamura, and J. S. Tsai, Nature **449**, 588 (2007).

[87] A. A. Abdumalikov, O. Astafiev, A. M. Zagoskin, Y. A. Pashkin, Y. Nakamura, and J. S. Tsai, Physical Review Letters **104**, 193601 (2010).

[88] J. Joo, J. Bourassa, A. Blais, and B. C. Sanders, Physical Review Letters **105**, 073601 (2010).

[89] W. R. Kelly, Z. Dutton, J. Schlafer, B. Mookerji, T. A. Ohki, J. S. Kline, and D. P. Pappas, Physical Review Letters **104**, 163601 (2010).

[90] M. A. Castellanos-Beltran and K. W. Lehnert, Applied Physics Letters **91**, 083509 (2007).

[91] N. Bergeal, F. Schackert, M. Metcalfe, R. Vijay, V. E. Manucharyan, L. Frunzio, D. E. Prober, R. J. Schoelkopf, S. M. Girvin, and M. H. Devoret, Nature **465**, 64 (2010).

[92] M. Hatridge, R. Vijay, D. H. Slichter, J. Clarke, and I. Siddiqi, Physical Review B **83**, 134501 (2011).

[93] M. A. Castellanos-Beltran, K. D. Irwin, G. C. Hilton, L. R. Vale, and K. W. Lehnert, Nature Physics **4**, 929 (2008).

[94] K. W. Murch, S. J. Weber, K. M. Beck, E. Ginossar, and I. Siddiqi, arXiv:1301.6276 (2013).

[95] J. Q. You and F. Nori, Nature **474**, 589 (2011).

[96] S. Haroche and J.-M. Raimond, *Exploring the quantum: atoms, cavities, and photons* (OUP Oxford, Aug. 2006).

[97] E. Jaynes and F. W. Cummings, Proceedings of the IEEE **51**, 89 (1963).

[98] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, *Photons and atoms: introduction to quantum electrodynamics* (Wiley, Aug. 1989).

[99] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf, Nature **431**, 162 (2004).

[100] D. I. Schuster, A. Wallraff, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. M. Girvin, and R. J. Schoelkopf, Physical Review Letters **94**, 123602 (2005).

[101] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Physical Review Letters **95**, 060501 (2005).

[102] M. Hatridge, S. Shankar, M. Mirrahimi, F. Schackert, K. Geerlings, T. Brecht, K. M. Sliwa, B. Abdo, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret, Science **339**, 178 (2013).

[103] J. P. Groen, D. Ristè, L. Tornberg, J. Cramer, P. C. de Groot, T. Picot, G. Johansson, and L. DiCarlo, Physical Review Letters **111**, 090506 (2013).

[104] D. I. Schuster, A. A. Houck, J. A. Schreier, A. Wallraff, J. M. Gambetta, A. Blais, L. Frunzio, J. Majer, B. Johnson, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Nature **445**, 515 (2007).

[105] M. Hofheinz, H. Wang, M. Ansmann, R. C. Bialczak, E. Lucero, M. Neeley, A. D. O'Connell, D. Sank, J. Wenner, J. M. Martinis, and A. N. Cleland, Nature **459**, 546 (2009).

[106] A. Galiautdinov, A. N. Korotkov, and J. M. Martinis, Physical Review A **85**, 042321 (2012).

[107] 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, Physical Review Letters **107**, 240501 (2011).

[108] J. B. Chang, M. R. Vissers, A. D. Córcoles, M. Sandberg, J. Gao, D. W. Abraham, J. M. Chow, J. M. Gambetta, M. Beth Rothwell, G. A. Keefe, M. Steffen, and D. P. Pappas, Applied Physics Letters **103**, 012602 (2013).

[109] A. Kleinsasser, R. Miller, and W. Mallison, IEEE Transactions on Applied Superconductivity **5**, 26 (1995).

[110] A. J. Berkley, "A josephson junction qubit", Ph.D. (University of Maryland, College Park, United States – Maryland, 2003).

[111] F. Pobell, *Matter and methods at low temperatures* (Springer, 1996).

[112] O. V. Lounasmaa, *Experimental principles and methods below 1 k* (Academic Press, 1974).

[113] C. Rigetti, J. M. Gambetta, S. Poletto, B. L. T. Plourde, J. M. Chow, A. D. Córcoles, J. A. Smolin, S. T. Merkel, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen, Physical Review B **86**, 100506 (2012).

[114] R. Barends, J. Wenner, M. Lenander, Y. Chen, R. C. Bialczak, J. Kelly, E. Lucero, P. O'Malley, M. Mariantoni, D. Sank, H. Wang, T. C. White, Y. Yin, J. Zhao, A. N. Cleland, J. M. Martinis, and J. J. A. Baselmans, Applied Physics Letters **99**, 113507 (2011).

[115] A. D. Córcoles, J. M. Chow, J. M. Gambetta, C. Rigetti, J. R. Rozen, G. A. Keefe, M. Beth Rothwell, M. B. Ketchen, and M. Steffen, Applied Physics Letters **99**, 181906 (2011).

[116] T. A. Palomaki, S. K. Dutta, H. Paik, H. Xu, J. Matthews, R. M. Lewis, R. C. Ramos, K. Mitra, P. R. Johnson, F. W. Strauch, A. J. Dragt, C. J. Lobb, J. R. Anderson, and F. C. Wellstood, Physical Review B **73**, 014520 (2006).

[117] V. Lefevre-Seguin, E. Turlot, C. Urbina, D. Esteve, and M. H. Devoret, Physical Review B **46**, 5507 (1992).

[118] T. A. Fulton and L. N. Dunkleberger, Physical Review B **9**, 4760 (1974).

[119] J. M. Martinis, S. Nam, J. Aumentado, K. M. Lang, and C. Urbina, Physical Review B **67**, 094510 (2003).

[120] J. Wenner, Y. Yin, E. Lucero, R. Barends, Y. Chen, B. Chiaro, J. Kelly, M. Lenander, M. Mariantoni, A. Megrant, C. Neill, P. J. J. O'Malley, D. Sank, A. Vainsencher, H. Wang, T. C. White, A. N. Cleland, and J. M. Martinis, Physical Review Letters **110**, 150502 (2013).

[121] Z. Kim, V. Zaretskey, Y. Yoon, J. F. Schneiderman, M. D. Shaw, P. M. Echternach, F. C. Wellstood, and B. S. Palmer, Physical Review B **78**, 144506 (2008).

[122] Z. Kim, B. Suri, V. Zaretskey, S. Novikov, K. D. Osborn, A. Mizel, F. C. Wellstood, and B. S. Palmer, Physical Review Letters **106**, 120501 (2011).

[123] R. P. Budoyo, *Estimate of the coherence time from sources of dissipation and dephasing in a dc-SQUID phase qubit*, unpublished (Oct. 2011).

[124] S. Oh, K. Cicak, J. S. Kline, M. A. Sillanpää, K. D. Osborn, J. D. Whittaker, R. W. Simmonds, and D. P. Pappas, Physical Review B **74**, 100502 (2006).

[125] Y. Nakamura, H. Terai, K. Inomata, T. Yamamoto, W. Qiu, and Z. Wang, Applied Physics Letters **99**, 212502 (2011).

[126] J. S. Kline, M. R. Vissers, F. C. S. da Silva, D. S. Wisbey, M. Weides, T. J. Weir, B. Turek, D. A. Braje, W. D. Oliver, Y. Shalibo, N. Katz, B. R. Johnson, T. A. Ohki, and D. P. Pappas, Superconductor Science and Technology **25**, 025005 (2012).

[127] M. Maezawa, M. Aoyagi, H. Nakagawa, I. Kurosawa, and S. Takada, Applied Physics Letters **66**, 2134 (1995).

[128] K. C. Gupta, *Microstrip lines and slotlines 2nd ed.* (Artech House, Incorporated, Jan. 1996).

[129] P. Dutta and P. M. Horn, Reviews of Modern Physics **53**, 497 (1981).

[130] M. B. Weissman, Reviews of Modern Physics **60**, 537 (1988).

[131] F. C. Wellstood, C. Urbina, and J. Clarke, Applied Physics Letters **50**, 772 (1987).

[132] R. H. Koch, D. P. DiVincenzo, and J. Clarke, Physical Review Letters **98**, 267003 (2007).

[133] F. Yoshihara, K. Harrabi, A. O. Niskanen, Y. Nakamura, and J. S. Tsai, Physical Review Letters **97**, 167001 (2006).

[134] T. Lanting, A. J. Berkley, B. Bumble, P. Bunyk, A. Fung, J. Johansson, A. Kaul, A. Kleinsasser, E. Ladizinsky, F. Maibaum, R. Harris, M. W. Johnson, E. Tolkacheva, and M. H. S. Amin, Physical Review B **79**, 060509 (2009).

[135] V. Zaretskey, "Decoherence and defects in cooper-pair boxes", Ph.D. (University of Maryland, College Park, 2013).

[136] W.-T. Tsang and T. Van Duzer, Journal of Applied Physics **46**, 4573 (1975).

[137] J. Clarke and A. I. Braginski, *The SQUID handbook: fundamentals and technology of SQUIDs and SQUID systems* (John Wiley & Sons, Mar. 2006).

[138] L. D. Landau and E. M. Lifshits, *Mechanics* (Butterworth-Heinemann, Jan. 1976).

[139] B. R. Mollow, Physical Review **188**, 1969 (1969).

[140] M. Baur, S. Filipp, R. Bianchetti, J. M. Fink, M. Göppl, L. Steffen, P. J. Leek, A. Blais, and A. Wallraff, Physical Review Letters **102**, 243602 (2009).

[141] D. M. Pozar, *Microwave engineering* (Wiley, Feb. 2004).

[142] K. Mølmer, Y. Castin, and J. Dalibard, Journal of the Optical Society of America B **10**, 524 (1993).

[143] J. Gambetta, A. Blais, M. Boissonneault, A. A. Houck, D. I. Schuster, and S. M. Girvin, Physical Review A **77**, 012112 (2008).

[144] K. W. Murch, S. J. Weber, C. Macklin, and I. Siddiqi, arXiv:1305.7270 (2013).

[145] S. Ashhab, J. Q. You, and F. Nori, Physical Review A **79**, 032317 (2009).

[146] C. G. Callan and S. Coleman, Physical Review D **16**, 1762 (1977).

[147] A. Schmid, Annals of Physics **170**, 333 (1986).

[148] T. Banks, C. M. Bender, and T. T. Wu, Physical Review D **8**, 3346 (1973).

[149] N. Moiseyev, Physics Reports **302**, 212 (1998).

[150] D. T. Colbert and W. H. Miller, The Journal of Chemical Physics **96**, 1982 (1992).