<small>ABSTRACT</small>

Controlling Photons in Superconducting Electrical Circuits
Blake Robert Johnson
2011

Circuit quantum electrodynamics (circuit QED) is a system that allows for strong coupling between microwave photons in transmission line cavities and superconducting qubits or artificial atoms. While circuit QED is often studied in the context of quantum information processing, it also provides an attractive platform for performing quantum optics experiments on-a-chip, because the level of control and coupling strengths available in circuit QED opens a vast array of possibilities for the creation, manipulation, and detection of quantum states of light. In this thesis, the extension of circuit QED to two cavities is examined, including design issues for cavities with very different Q-factors, and a new qubit design is proposed that couples to both cavities. The qubit-cavity interaction, while providing much of the utility of circuit QED, also introduces additional qubit relaxation. A powerful formalism for calculating this energy decay due to the classical admittance of the electromagnetic environment is presented in the context of circuit QED. Measurements of a wide range of samples validate this theory as providing an effective model for relaxation. A new circuit element, called the 'Purcell filter', is introduced and demonstrated to decouple the relationship between cavity Q and qubit relaxation. Finally, a new method for performing quantum non-demolition measurements of microwave photons is demonstrated.

Controlling Photons in Superconducting Electrical Circuits

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Blake Robert Johnson

Dissertation Director: Professor Robert J. Schoelkopf

May 2011

Contents

List of Figures

for my family
and in memory of Michael Tinkham

Acknowledgements

My physics Ph. D. would not have been possible without several people, the first of whom is my advisor, Rob Schoelkopf, who convinced me to come to Yale. Rob tolerated my pursuit of (occasionally tangential) directions in hardware or software related to the experiments. While these explorations sometimes slowed me down, they also gave me a better understanding of the work of RSL. I continue to be amazed at Rob's experimental intuition, and his skill in presenting and telling a story to convey understanding.

I would like to thank Steve Girvin and Alexandre Blais, who taught me circuit QED when I was still a first year graduate student. Steve has a particular gift for communicating with experimentalists, rapidly understanding the details of an experiment. His suggestions for new knobs to turn frequently yield interesting results.

Thank you also to Michel Devoret, who always found time to help me see our efforts from a broader perspective. His honest feedback dramatically improves every project with which he comes in contact.

Several people at Harvard contributed greatly to my desire to pursue a physics PhD. Few figure so prominently among them as Michael Tinkham, who offered me a position in his lab. There, I learned the ropes of low-temperature measurements, and came to appreciate both his encyclopedic knowledge of superconductivity as well his characteristic wit.

Working at Yale has been a pleasure because of the community of people gathered on the 4th floor of Becton. It is difficult to understate the value of having colleagues who are also friends. The open-door atmosphere of the 4th floor, where one can drop into discussions with anyone at anytime, helped tremendously to grease the wheels whenever I was stuck, whatever I was stuck on (hardware, software, fab, and theory all come to mind). Of these people, I must particularly single out Luigi Frunzio for his expert guidance in fabrication.

He is always enthusiastically ready to tackle any project, and willing to figure out how to get it done on the time table of the person asking for his help.

I owe a great scientific debt to ANDREW HOUCK and DAVE SCHUSTER, who have been continuous sources of ideas and advice. They have also provided a seemingly infinite source of good cheer, being as proficient in the fine art of goofing off as they are in physics—no small feat. These two proved to me that it is possible to be scientifically productive while having fun.

When issues of a more theoretical nature threatened to halt an experiment, LEV BISHOP and JAY GAMBETTA frequently saved the day. LEV has been particularly helpful in the preparation of this manuscript, willing to repay his efforts to get me to procrastinate with equal measures of assistance in reading drafts and helping me work through thorny questions.

My time in New Haven would have been dramatically less enjoyable without the LOST WEDNESDAY crew. Despite the loss of sleep suffered when Phyllis' dessert experiments extended into the late hours of Tuesday night (the epic battle of the macaroons will be forever remembered!), hanging out with LW friends always served to recharge my batteries. Bad movies and mediocre television just aren't the same without their regular banter.

I find it funny that JERRY CHOW and I barely knew each other at Harvard, despite working in labs that were literally across the hall. JERRY is now one of my closest friends and a valued colleague. His energy fueled many of my own efforts, even when we were not working on the same experiment. It was wonderful having someone there everyday with whom I could talk about anything—frustrations as well as successes.

There were several people who played a prominent role in my formative years before I knew that I would go into physics. Among them are my high school physics teacher, KAREN PHILLIPS, who planted the initial seed of the idea that the study of physics was not mutually exclusive with an interest in computers. Of course, I would not have even gotten that far without the support of my parents, DEBBIE and BOB, who encouraged and cultivated my interests, coached *Odyssey of the Mind* and carted me to and from the science museum, *MITY*, and the U of MN, not to mention sending me to excellent schools. Some of these activities began when I followed my sister, KELLY, as she did the things that she found fun. Later, her competitive spirit taught me to make the most of the tools given to me.

Finally, without PHYLLIS, I might still be playing whiffle ball instead of finishing this thesis. Her love and support are my foundation, giving me the assurance and confidence I need to venture into uncharted waters.

Publication list

This thesis is based in part on the following published articles:

1. D. I. Schuster, A. A. Houck, J. A. Schreier, A. Wallraff, J. Gambetta, A. Blais, L. Frunzio, J. Majer, B. R. Johnson, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, "Resolving photon number states in a superconducting circuit," *Nature* **445**, 515–518 (2007).

2. A. A. Houck, D. I. Schuster, J. Gambetta, J. A. Schreier, B. R. Johnson, J. M. Chow, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, "Generating single microwave photons in a circuit," *Nature* **449**, 328–331 (2007).

3. J. Majer, J. M. Chow, J. M. Gambetta, J. Koch, B. R. Johnson, J. A. Schreier, L. Frunzio, D. I. Schuster, A. A. Houck, and A. Wallraff, "Coupling superconducting qubits via a cavity bus," *Nature* **449**, 443–447 (2007).

4. J. A. Schreier, A. A. Houck, J. Koch, D. I. Schuster, B. R. Johnson, J. M. Chow, J. Gambetta, J. Majer, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, "Suppressing charge noise decoherence in superconducting charge qubits," *Phys. Rev. B* **77**, 180502 (2008).

5. A. A. Houck, J. A. Schreier, B. R. Johnson, J. M. Chow, J. Koch, J. Gambetta, D. I. Schuster, L. Frunzio, M. H. Devoret, and S. M. Girvin, "Controlling the spontaneous emission of a superconducting transmon qubit," *Phys. Rev. Lett.* **101**, 80502 (2008).

6. J. M. Chow, J. M. Gambetta, L. Tornberg, J. Koch, L. Bishop, A. A. Houck, B. R. Johnson, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, "Randomized benchmarking and process tomography for gate errors in a solid-state qubit," *Physical Review Letters* **102**, 090502 (2009).

7. L. DiCarlo, J. M. Chow, J. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, "Demonstration of two-qubit algorithms with a superconducting quantum processor," *Nature* **460,** 240–249 (2009).

8. M. D. Reed, B. R. Johnson, A. A. Houck, L. DiCarlo, J. M. Chow, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf, "Fast reset and suppressing spontaneous emission of a superconducting qubit," *Appl. Phys. Lett.* **96**, 203110 (2010).

9. B. R. Johnson, M. D. Reed, A. A. Houck, D. I. Schuster, L. S. Bishop, E. Ginossar, J. M. Gambetta, L. DiCarlo, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, "Quantum non-demolition detection of single microwave photons in a circuit," *Nature Physics* **6**, 663–667 (2010).

CHAPTER 1

Introduction

T HE past decade has witnessed a remarkable convergence of advances in two distinct fields of low-energy physics. In superconducting circuits, there is an increasing body of evidence demonstrating that macroscopic quantities, like the voltage on a wire, can behave quantum mechanically [1–4]. This was rather surprising given that voltages are determined by an ensemble of many ($> 10^{10}$) particles, and such large numbers of particles have typically been thought to behave classically. The evidence for macroscopic quantum coherence in superconductors lead to the development of many flavors of 'artificial atoms', which are electrical circuits that have discrete quantum energy levels, analogous to natural atoms. These circuits are often referred to as 'superconducting qubits' because of their natural applications to quantum computing. There has been rapid progress in using these circuits for this purpose, with several qubit designs, including the quantronium [5], transmon [6] and fluxonium [7], now routinely having decoherence times $\sim 10,000$ times longer than the first superconducting qubits. Furthermore, multiple qubits can be wired together [8, 9], and one can run simple quantum algorithms on processors with a few qubits [10–12].

The other advance occurred in the atomic physics community, where atoms have been strongly coupled to microwave cavities in an architecture known as cavity quantum electro-dynamics (cQED) [13]. Photons normally interact very weakly with atoms, but by trapping photons in a cavity, the effective interaction becomes much stronger because the photons have

many chances to be absorbed by the atoms. The field of cQED has seen extraordinary success performing increasingly complex manipulations of atoms and cavities, with state-of-the-art experiments demonstrating counting of the number of photons stored in the cavity, and watching the decay of those photons as they leave one at a time [14].

In 2003, Steve Girvin, Rob Schoelkopf, and co-workers realized that these two advances could be combined in a completely electrical structure on a chip. They replaced the natural atoms and 3D cavities of standard cQED architecture with superconducting qubits and planar transmission line resonators. Thus, the field of *circuit* QED was born.

The operation of circuit QED devices is sufficiently different from other solid-state systems that when describing these systems it is often easier to use terminology from the field of quantum optics. We call our qubits 'artificial atoms' and the quantized energies stored in the currents and voltages in our cavities we call 'photons'. Some members of the community have resisted this language; however, a photon is still a photon whether its wavelength is a micron or a centimeter, and our superconducting 'qubits' display rich, multi-level structure which is reminiscent of natural atoms. The beauty of these devices, however, is that they are fully *engineered* quantum systems. Unlike natural atoms with their God-given level structure, dipole moments, and so forth, these superconducting circuits have properties chosen by the experimenter. Even better, these circuits can be designed such that their parameters are adjustable with various 'knobs'. The price one pays for this marvelous flexibility is that these same knobs are also channels for energy decay and decoherence, whereby quantum information stored in these systems is lost to the surrounding environment. Part of this thesis will address understanding these effects and engineering new ways to overcome them. This is of principal importance toward using these circuits to build practical quantum computers. However, quantum computing is but one application of quantum control and measurement. It turns out that the use of quantum optics language was not solely a convenience, but that these same devices are powerful components for doing real quantum optics experiments on-a-chip. Some already demonstrated examples are discussed in the theses of David Schuster and Lev Bishop [15, 16], including number splitting [17], production of single-photons on demand [18], and multi-photon transitions in the Jaynes-Cummings ladder [19].

The principal focus of this thesis will be understanding and developing these solid-state quantum optics experiments. These experiments seek to create, manipulate, and detect quantum states of light. This naturally begs the question: what is quantum light? This seemingly simple question turns out to be remarkably difficult to answer precisely. One can
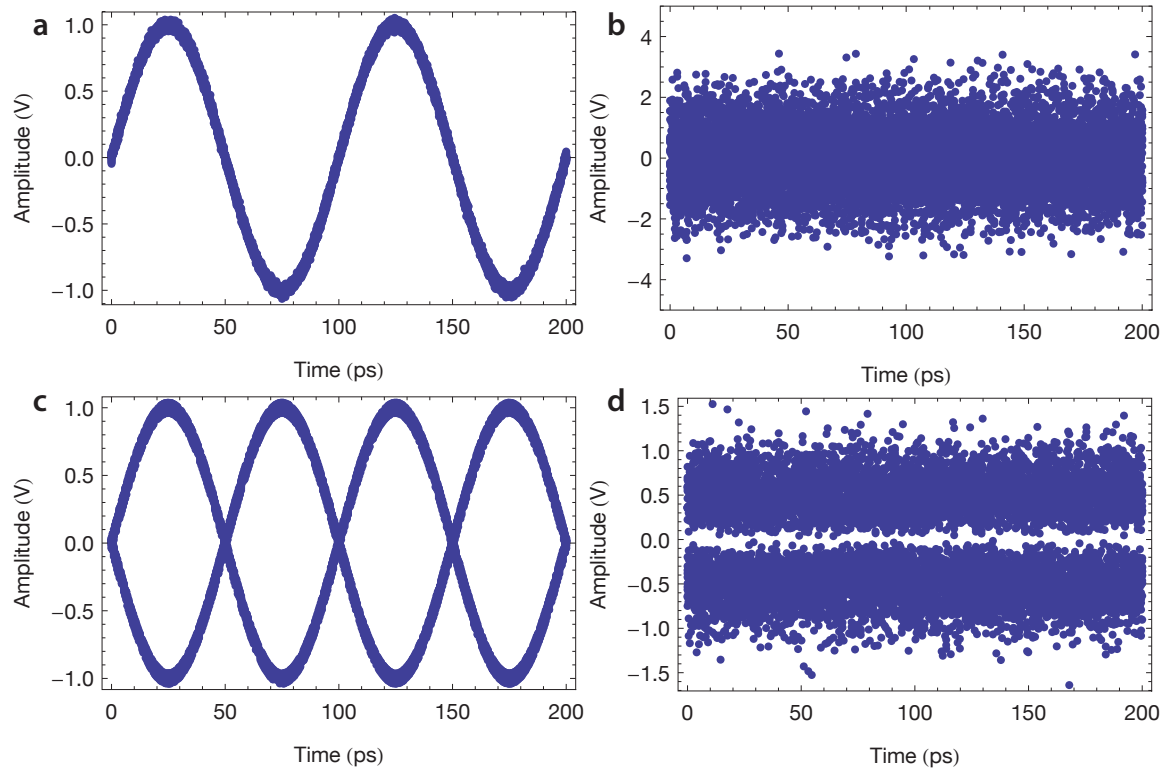
Figure 1.1: Viewing light on an accumulation oscilloscope. Cartoon sketches of what different kinds of light might look like on an accumulation oscilloscope. **a** Coherent light, like from a laser or a microwave generator, **b** thermal light, like sunlight, **c** a 'cat state', or superposition of coherent states with opposite phases, and **d** a single photon 'Fock state', like the light produced by the decay of a single atom.

develop a certain amount of intuition, though, by considering what various kinds of light would look like if we could directly capture the field voltage on the screen of an oscilloscope.

To be more specific, suppose that we have access to an ensemble of identically prepared states of light traveling down a transmission line, and that our oscilloscope is uncoupled from the line until the precise moment in which we sample the voltage on the wire. We repeat many such measurements for various delays and accumulate the measured voltages on the screen of the scope. Figure 1.1 shows a cartoon of the oscilloscope screen for various ensembles of classical and quantum light sources. The first image (a) results from a coherent state, which describes the light produced by a laser or a microwave generator. As viewed on the oscilloscope, it appears as a sine wave with a well-defined average amplitude and phase. The width of the line is not a deficiency of the oscilloscope, but rather is caused by quantum fluctuations that are always present. Another light source is thermal light (b), such

as light from the sun, an incandescent light bulb, or any black body at non-zero temperature. On the oscilloscope this light appears as a flat band around zero volts. Though quantum fluctuations are still present, in this case, the width of the line is primarily determined by the temperature of the source. The thermal state has a well-defined root mean square (RMS) amplitude, but a fluctuating phase. We can compare these two *classical* light sources with two other *quantum* ones. The first, shown in (c), is a so-called 'cat state'*. This particular cat is a superposition of coherent states with opposite phases, which on the oscilloscope screen appears as two interwoven sine waves. Unfortunately, we cannot use this measurement to verify the quantum nature of the light, because despite being *consistent* with a cat state, the quantum nature of this cat is actually not visible on the screen. A malicious party could have replaced our cat state ensemble with a classical mixture of the constituent coherent states and it would appear identical for this measurement. Evidently, one needs to build a very different 'oscilloscope' to observe a coherent superposition. The last state is a Fock state or photon number state (d), like what is produced when a single atom decays. Its appearance on the oscilloscope screen bares some resemblance to the thermal state, because like the thermal state, the Fock state has a definite RMS amplitude (power) without a well-defined phase. However, whereas the lack of phase of the thermal state arises from a classical mixture of signals with random amplitudes and phases, the Fock state is a coherent superposition of all phase states. Thus the apparent noise in this measurement is completely quantum. The absence of points at $V = 0$ (for odd-numbered Fock states) is a result of the particular nature of this noise.† The difficulty in distinguishing quantum from classical with such an oscilloscope, however, reveals the need for better tools, several of which have been fully developed in the quantum optics community and which I review in section 2.4.1. For now, I hope you will accept that quantum light is light that has a fundamentally different character than coherent or thermal light.

Putting this issue aside for the moment reveals another important question: why is quantum light interesting? Whereas the development of superconducting qubits has a specific technological goal of building a scalable quantum computer, the applications for quantum light are not nearly so obvious. It is possible that quantum states of light could be useful

---

* This label appears to be used in quantum optics to describe any superposition state of some macroscopic quantity. Thus, there can be amplitude cats, phase cats, and so on. It is also common to refer to the distance between the superposed states as the 'size' of the cat.

† Amazingly enough, if we were to plot the average signal on our oscilloscope for either the thermal, cat, or Fock states, at any point in time the signal would average to zero despite the presence of the light!

for novel communications protocols, such as quantum cryptography [20] which allows two parties to exchange encryption keys that cannot be secretly intercepted. There is also 'squeezed' light which has applications to ultra-precise measurements [21, chapter 8]. These applications require 'flying', or *itinerant* photons, which are photons traveling through vacuum or down a wire. In contrast, for the most part this thesis deals with stationary photons—photons that are trapped in a cavity. Furthermore, these photons are at microwave frequencies, which makes them more difficult to send losslessly over long distances compared to optical photons.* Consequently, there are remaining technological hurdles before the techniques presented here could be used for quantum communication.

Even before these technological hurdles are overcome, there is reason to take interest in systems that store quantum information in stationary photons. Continued progress in quantum information processing has shown increasing sophistication in the control, entanglement, and measurement of fermions (qubits). Meanwhile, bosonic systems (cavities) have fallen behind in terms of these metrics, despite having longer coherence times than many qubits. It may turn out that these bosons are equivalently good constituent elements in a quantum information processor. In the meantime, though, there is incredible power in coupling even a few cavities. Consider that the size of the Hilbert space for N coupled qubits is 2^N, while coupled cavities can have many more accessible levels. The number of states in experimentally realizable cavities is limited by energy thresholds, such as the critical current of the superconducting wires carrying the current. Even considering cavities truncated to just 5 levels, though, if N such cavities of different frequencies are coupled, then the dimension of the resulting Hilbert space is 5^N. One would need five qubits to have a larger Hilbert space than even two such cavities. Consequently, multi-cavity circuit QED opens up a vast area for exploring quantum control and entanglement in spaces with large degrees of freedom while requiring very few 'moving parts'. This makes multi-cavity circuit QED an attractive frontier for further research and growth in quantum information.

The trade-off for this rapid growth is that the full Hilbert space of coupled cavities is not easily accessible. Qubits are needed in order to load excitations into the cavities one at a time. Recently, Strauch *et al.* [22] have extended earlier work by Law and Eberly [23] to

---

* This limitation is partly an issue of economics. Superconducting transmission lines are better than most optical fibers, but the need for cryogenics makes superconducting transmission lines much more expensive. Note that this limitation is not relevant to on-chip communication where the distances are much smaller. Microwave photons can be efficiently routed around superconducting circuits, making them potential useful for communication between 'distant' qubits in a quantum computer.

show that one can prepare an arbitrary state of two cavities with one qubit that is coupled to each. In addition to the resonant SWAP gate recently demonstrated by several groups [24, 25], Strauch's scheme requires a photon-number selective qubit gate of the type first demonstrated in chapter 8. This thesis thus lays the groundwork for this new direction in quantum information.

1.1   Overview of thesis

This thesis largely deals with extensions to the circuit QED architecture for the purpose of creating, detecting, and manipulating quantum light. Before embarking into unexplored waters, we begin in chapter 2 with a review of circuit QED with transmon qubits. This review discusses the resonant and dispersive regimes of the Jaynes-Cummings Hamiltonian before introducing a new 'quasi-dispersive' regime in section 2.2.4. The quasi-dispersive regime presents a rich level structure that can be easily understood in terms of a smooth connection between the resonant and dispersive regimes. It will turn out that this regime is incredibly useful for photon-number-dependent quantum logic, which is used in chapter 8 to perform quantum non-demolition photon measurements. To fully appreciate this, one needs to understand quantum measurements, which are reviewed in section 2.3. Section 2.4.1 recalls powerful mathematical tools for phase space descriptions of cavity states, allowing us to provide a definition of quantum light. Finally, section 2.5 examines recent experiments that create and measure various quantum states of light.

Chapter 3 develops the theory for circuit QED with two cavities. Any circuit involving more than one cavity naturally involves some kind of direct or indirect coupling between the cavities. Section 3.2 discusses the effects of a classical coupling, while section 3.3 and section 3.5 describe quantum coupling mediated by a qubit. The latter section reveals a Kerr-type interaction between the cavities that is tested experimentally in chapter 7. Sometimes the cavity-cavity coupling is not desired, so section 3.4 introduces a modified transmon design, called the 'sarantapede', that serves to couple a single qubit to two cavities while minimizing the indirect coupling through the qubit.

Cavity-qubit coupling provides the useful physics of circuit QED, but it also introduces additional qubit relaxation that cannot be fully described by a single-mode theory. Chapter 4 examines relaxation from the full multi-mode Purcell effect by applying a powerful formalism for calculating relaxation due to the classical admittance of the electromagnetic environment. With this improved understanding of relaxation, we are poised to suggest two ways to

ameliorate it: a 'balanced' qubit design in section 4.1.1 and the 'Purcell filter' in section 4.1.2. The same admittance issues affect the introduction of fast flux bias lines to the circuit for qubit frequency control. Section 4.2 looks at design issues for these flux bias lines, and estimates the additional relaxation and dephasing from adding them. Even after finding a design which minimally affects qubit performance, these control lines do not respond perfectly. Consequently, we apply deconvolution methods in section 4.3 for improving the outputs from the flux control system.

In chapter 5, I describe advances in the fabrication of circuit QED devices, done in collaboration with Luigi Frunzio, for making an array of tranmon designs and for making samples on sapphire substrates. Modifications to the measurement setup required for two cavity experiments are detailed in section 5.2. The following experiments require precise microwave pulses. A limited quantity of expensive vector microwave sources spurred the development of some custom hardware for pulse generation, which is described in section 5.2.3.

After all the build-up we move onto actual experiments in chapter 6, which looks at relaxation in real circuit QED devices. Section 6.1 applies the previously developed classical admittance formalism to measurements on a wide variety of circuit QED samples with transmon qubits. From these data a consistent explanation of qubit relaxation is found in the multi-mode Purcell effect acting in parallel with a constant-Q source that is consistent with dielectric loss. This work required collecting data from a large number of samples, a task which I performed together with Jerry Chow and Joe Schreier, while Andrew Houck had the insight to see the thread connecting these many different devices. In section 6.2, we go on to demonstrate a new circuit element called the 'Purcell filter' that decouples qubit relaxation from cavity decay. The idea for this filter originated with Andrew Houck, though Matt Reed did most of the legwork to carry it to fruition, with assistance from me. The use of this filter as a means to provide qubit reset emerged from experiments that Matt and I did together.

Chapter 8 describes a true quantum optics on-a-chip experiment that demonstrates a quantum non-demolition (QND) method for detecting photons in a cavity. It operates by means of a qubit-photon logic gate that maps information about the number of photons in a cavity onto a qubit state. The idea for this experiment emerged out of many discussions with Andrew Houck, David Schuster, and Jay Gambetta. This chapter shows repeated measurements of single photons using this logic gate, whereby I am able to claim that the method is at least 90% QND.

The thesis ends in chapter 9 with some thoughts on extensions to the photon detection experiment and possibilities for circuit QED with multiple cavities.

Circuit QED and Quantum Optics

Tʜɪs thesis is meant as a follow-up to *Circuit Quantum Electrodynamics, Vols. I* and *II* by David Schuster and Lev Bishop, respectively. Consequently, I expect the reader to be largely familiar with the material contained in those volumes. Nonetheless, I will present a brief overview of some topics covered there in order to establish certain nomenclature, as well as to refer you to the relevant sections of those volumes to learn more about topics with which you are less familiar.

2.1   Cooper-pair Box and Transmon

By now, there are many established flavors of superconducting quantum bits (qubits). It used to be that superconducting qubits could be easily classified as charge, flux, or phase qubits based upon the quantity in each that served as a good quantum number. However, the community as a whole has increasingly moved toward qubits like the quantronium, transmon, capacitively shunted flux qubit, and fluxonium, which cannot easily be put into one of these categories, but instead exist in intermediate regimes between charge, flux, and phase. It is worth noting that all of these devices are really many-level devices. They are called "qubits" because all have sufficient anharmonicity between levels that some pair of levels can be individually addressed. Thus, they can be operated as two-level systems, though often is
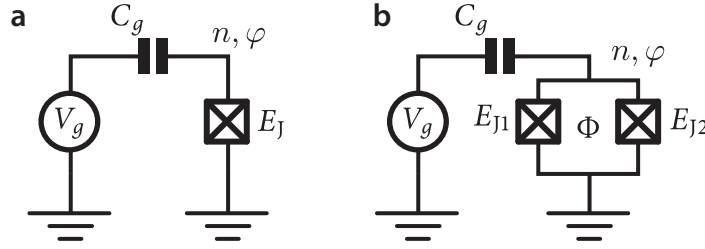
Figure 2.1: The Cooper-pair box. a The standard Cooper-pair box is a superconducting island coupled to a reservoir by a Josephson junction. The charge on the island is modulated by a gate voltage, V_g, which is capacitively coupled to the island via C_g. **b** In the split Cooper-pair box, the single junction is replaced with a SQUID, allowing the effective Josephson energy of the two junction loop to be tuned by the flux, Φ.

easier to think of these devices as *artificial atoms*, in reference to their engineered atom-like level structures.

In my own graduate student career I have participated in the development and characterization of the transmon, which though topologically identical to a charge qubit, has basis states which are largely localized in phase. This gives the transmon the remarkable advantage of being immune to charge noise, which is a persistent problem for devices fabricated on substrates, without increased sensitivity to flux or critical current noise. The transmon transition energies can also be tuned *in situ* by application of a local magnetic field. Slow drifts of field are not a major concern if the devices are sufficiently shielded, so the transmon can be stably operated at a chosen frequency for many days at a time.

Cooper-pair box

In order to describe transmons quantitatively, it is useful to start with the Cooper-pair box (CPB). The CPB is a simple circuit consisting of an island connected to a reservoir by a Josephson junction (see figure 2.1(a)). The island is also capacitively coupled to a voltage source, V_g, which can modulate the electrostatic energy of charges stored on the island. The device energies are determined by two parameters, the Josephson energy, E_J, for Cooper-pairs to tunnel across the junction, and the charging energy, E_C, which is the energy cost of bringing an additional electron onto the island from infinitely far away. The Hamiltonian describing this circuit is

$$H = 4E_C(\hat{n} - n_g)^2 - E_J \cos \hat{\varphi}, \tag{2.1}$$

where $\hat{n}$ is the operator for the number of charges on the island, $n_g = C_g V_g/2e$ is the applied gate voltage, V_g, expressed in units of Cooper-pairs, and $\hat{\varphi}$ is the gauge-invariant phase which is equal to the time integral of the voltage across the junction.* A useful modification of this circuit adds an additional Josephson junction in parallel with the first, shown in figure 2.1(b). In this configuration, the effective Josephson energy of the pair of junctions can be modulated by the flux Φ penetrating the loop formed by the junctions. The Hamiltonian for such a loop can be written, following the spanning tree procedure described in [26] and [16, chapter 2], as

$$H = -E_{J1} \cos \varphi - E_{J2} \cos(\varphi + 2\pi\Phi/\Phi_0), \tag{2.2}$$

where E_{J1} and E_{J2} are the Josephson energies of the junctions, φ is the node phase of the island, and $\Phi_0 = 2e/h$ is the magnetic flux quantum. With trigonometric identities, this can be rewritten as [15, 16]

$$H = -E_{J\Sigma} \cos\left(\frac{\pi\Phi}{\Phi_0}\right)\sqrt{1 + d^2 \tan^2\left(\frac{\pi\Phi}{\Phi_0}\right)} \cos(\varphi - \varphi_0), \tag{2.3}$$

where $E_{J\Sigma} = E_{J1} + E_{J2}$ is the total Josephson energy, $d = (E_{J1} - E_{J2})/(E_{J1} + E_{J2})$ describes the asymmetry between junctions, and φ_0 is given by $\tan(\varphi_0 + \pi\Phi/\Phi_0) = d\tan(\pi\Phi/\Phi_0)$. When the junction asymmetry is small, typically $d \sim 0.05$–0.1, the two junction loop behaves like a single junction with a flux-dependent Josephson energy

$$E_J(\Phi) \simeq E_{J\Sigma} \cos(\pi\Phi/\Phi_0). \tag{2.4}$$

The residual asymmetry presents a term proportional to $d\sin(\pi\Phi/\Phi_0)$, which opens a channel for relaxation because it allows coupling between states of opposite parity. This is discussed in the section on flux bias lines in section 4.2.

When the charging energy is much larger than the Josephson energy, $E_C \gg E_J$, the CPB eigenstates are essentially charge states. In this regime, the first term of (2.1), $4E_C(n - n_g)^2$, gives rise to parabolic energy levels as a function of the gate charge n_g, and the Josephson energy lifts the degeneracy between charge states at half-integer values of n_g. At these values of gate charge, the transition energies are also first-order insensitive to fluctuations in n_g; consequently, this is called a "sweet-spot" [5]. The energy spectrum of the CPB is highly

---

* I have used hats over the variables in (2.1) to emphasize that the charge on the island and phase across the junction are quantum variables. Having established this, I will drop the hats in the remainder of the thesis.
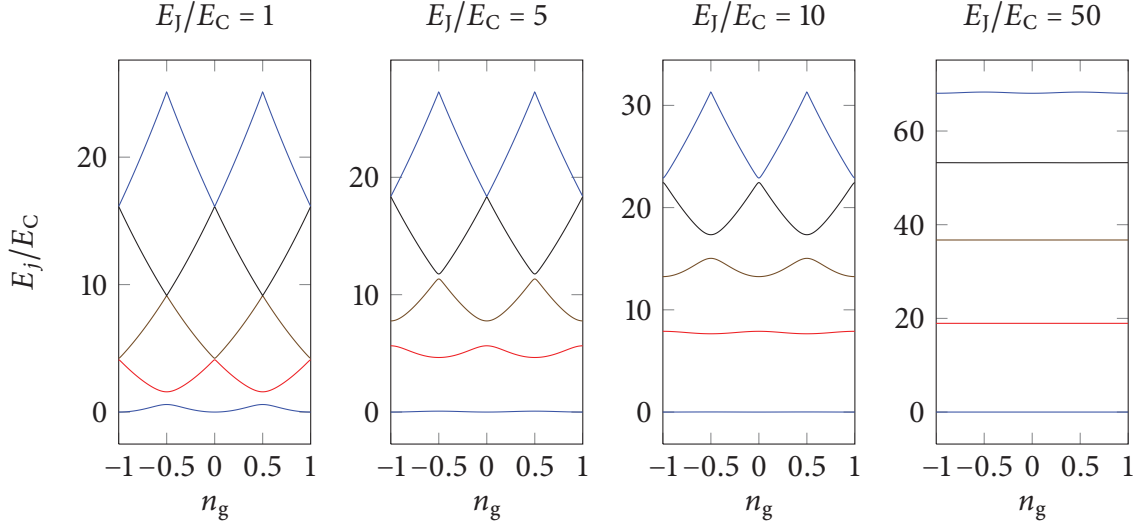
Figure 2.2: Charge dispersion. CPB energy levels, E_j, for the lowest 5 charge states of (2.1) are shown in units of E_C. Energy bands are shown for E_J/E_C = 1, 5, 10, and 50. At small E_J/E_C, the system is in the charge regime and the bands have a parabolic shape. The avoided crossing between the two lowest levels at $n_g = \pm 0.5$ is a 'sweet spot' where the transition energy, E_{01} is first-order insensitive to fluctuations in the gate charge. As E_J/E_C increases the bands flatten and the device enters the transmon regime. Reproduced from [6, 16].

anharmonic, so one can construct a two level approximation that treats the lowest two levels of the CPB as an effective spin-1/2. To find out more about this procedure, see [15, section 3.2].

Transmon

By increasing the ratio of energies E_J/E_C one enters a rather different regime. When the Josephson energy is the dominant energy scale, $E_J \gg E_C$, then the device becomes a weakly anharmonic oscillator. To achieve this ratio of energies, it is sufficient to keep a similar E_J and add a large capacitance in parallel with the junction in the CPB circuit. The increase of E_J/E_C causes the charge bands of figure 2.2 to flatten. In fact, defining the charge dispersion as the spread in transition energy between adjacent energy levels,

$$\epsilon_m = E_{m,m+1}(n_g = 1/2) - E_{m,m+1}(n_g = 0) \tag{2.5}$$

where $E_{ij} = E_j - E_i$ is the energy difference between levels i and j, then the residual charge dispersion decreases exponentially with increasing E_J/E_C. In fact, [6]

$$\epsilon_m \sim E_C \exp(-\sqrt{8E_J E_C}). \tag{2.6}$$

Consequently, even at moderate values of E_J/E_C, the charge dispersion can be suppressed to less than 1 kHz, making the transmon effectively immune to charge noise.

Whereas the charge dispersion is exponential in E_J/E_C, the remaining anharmonicity is only algebraic in E_J/E_C. The anharmonicity is defined as the difference between adjacent transition energies. For operation as a qubit, the relevant levels are the bottom three, so it makes sense to define

$$\alpha = E_{12} - E_{01}. \tag{2.7}$$

As $E_J/E_C \to \infty$, the anharmonicity asymptotically approaches $\alpha \simeq -E_C$. A typical transmon has $E_C \simeq 300$ MHz, which provides sufficient anharmonicity for fast manipulation of the transmon state. Thus, one can treat the transmon as a qubit for quantum information processing. However, unlike a CPB that has relative anharmonicity, $\alpha_r = \alpha/E_{01}$, that is greater than 1 at the charge sweet spot, the higher transmon levels often play a important role in qubit-qubit or qubit-cavity coupling.

Since we will need to refer frequently to these higher levels, from this point forward I will label the transmon states with the "alphabetical" order: g, e, f, h, etc.* This will help in avoiding confusion when we proceed to label joint qubit-cavity states, where I will use numbers to label cavity states.

2.2 Circuit Quantum Electrodynamics

Cavity quantum electrodynamics (cavity QED) describes the interaction of matter with light trapped in an cavity. Typically, light interacts very weakly with matter, but by trapping the light in a cavity, this interaction can be amplified many times such that a single photon in the cavity will be absorbed and re-emitted by the atom many times before the photon escapes from the cavity.

---

* The somewhat odd choice allows referring to the ground and excited states as $|g\rangle$ and $|e\rangle$ while wrapping a mostly alphabetical order around it.

In 2004, Rob Schoelkopf, Steve Girvin, and others realized that the techniques of cavity QED could be used with superconducting qubits and transmission line cavities, creating the field of circuit QED. The CPB and transmon eigenstates have effective dipole moments which allow these qubits to couple to an electric field. If one of these qubits (treating it, for the moment, as a two level system) is placed in a 1D transmission line cavity, the combined system is described by the Hamiltonian (see [27] and [15, sections 3.3 and 4.3.4])

$$H = \hbar\omega_r(a^\dagger a + 1/2) + \frac{\hbar}{2}\omega_q\sigma_z + \hbar g(a + a^\dagger)\sigma_x, \tag{2.8}$$

where ω_r and ω_q are the cavity and qubit frequencies, respectively, and g parameterizes the interaction strength. Writing $\sigma_x = \sigma_+ + \sigma_-$ and taking the rotating wave approximation (RWA) to throw away terms that do not conserve energy (like $a^\dagger\sigma_+$), we arrive at the Jaynes-Cummings Hamiltonian

$$H = \hbar\omega_r a^\dagger a + \frac{\hbar}{2}\omega_q\sigma_z + \hbar g(a\sigma_+ + a^\dagger\sigma_-). \tag{2.9}$$

The Jaynes-Cummings Hamiltonian has been used to effectively model the physics in a wide range of experiments in cavity QED with both natural and artificial atoms. The simple form of the coupling means that the Hamiltonian can be written in a 2×2 block-diagonal form. This allows the Hamiltonian to be solved analytically giving the *dressed* eigenstates

$$|n, +\rangle = \cos(\theta_n)|n - 1, e\rangle + \sin(\theta_n)|n, g\rangle, \tag{2.10a}$$

$$|n, -\rangle = -\sin(\theta_n)|n - 1, e\rangle + \cos(\theta_n)|n, g\rangle, \tag{2.10b}$$

and eigenenergies

$$E_0 = -\frac{\hbar\Delta}{2}, \tag{2.11a}$$

$$E_{n,\pm} = n\hbar\omega_r \pm \frac{\hbar}{2}\sqrt{4g^2n + \Delta^2}, \tag{2.11b}$$

where $\Delta = \omega_q - \omega_r$ is the detuning between the qubit and cavity, and θ_n is given by

$$\tan(2\theta_n) = \frac{2g\sqrt{n}}{\Delta}. \tag{2.12}$$

So far, we have not said anything about dissipation. In practice, photons which enter the cavity will eventually decay, either leaving by traveling out one of the "mirrors", or the photons
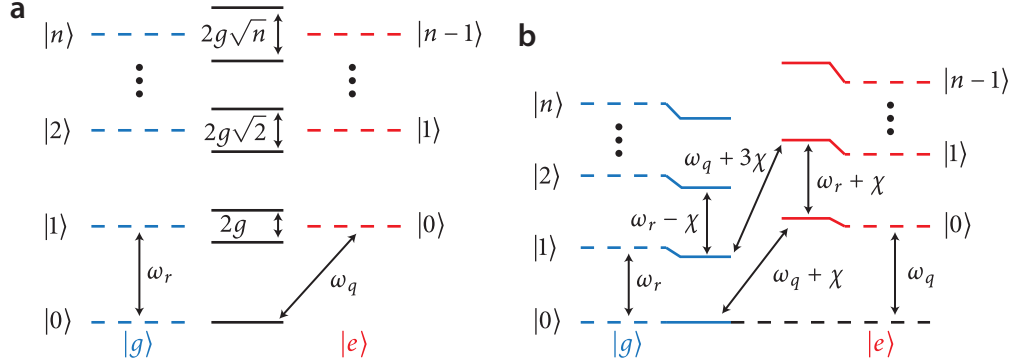
Figure 2.3: Resonant and Dispersive Jaynes-Cummings Regimes. a Resonant regime when $\omega_r = \omega_q$. The qubit-cavity coupling lifts the degeneracy such that the splitting is $2g\sqrt{n}$. **b** Dispersive regime with $\omega_q > \omega_r$. In this regime the photon-like transitions are $\omega_r \pm \chi$, depending on the qubit state. The qubit-like transitions are also photon-number dependent, $\omega_q + (2n+1)\chi$.

can dissipate by radiative or dielectric loss. The sum of all these processes is characterized by a rate, κ. Similarly, energy stored in the qubit can also decay, and we characterize this by the rate, γ. When the coupling strength is larger than these decay channels, *i.e.* $g > \gamma, \kappa$, then the system is in the *strong coupling* regime. In this regime, even a single excitation in the qubit or cavity exerts a strong influence on the other. These effects are explored in the next sections.

2.2.1   Resonant regime

Returning to the case of a simple two-level system coupled to a cavity, we can understand the Jaynes-Cummings Hamiltonian through various regimes. When the qubit and cavity are in resonance, $\omega_r = \omega_q$, then $\theta_n = \pi/4$ and the eigenstates and eigenenergies from (2.10) and (2.11) are

$$|n, \pm\rangle = (|n - 1, e\rangle \pm |n, g\rangle) / \sqrt{2}, \tag{2.13a}$$

$$E_{n,\pm} = n\hbar\omega_r \pm \hbar g\sqrt{n}. \tag{2.13b}$$

This situation is depicted schematically in figure 2.3(a). The eigenstates are equal superpositions of qubit and cavity states with the coupling lifting the degeneracy such that the splitting is $2\hbar g\sqrt{n}$, where n is the number of photons in the cavity. This $\sqrt{n}$ scaling has been observed spectroscopically up to $n = 2$ in [28] and up to $n = 6$ in [19]. It can also be probed in the time-domain, because a bare qubit or cavity state which is brought into resonance with the
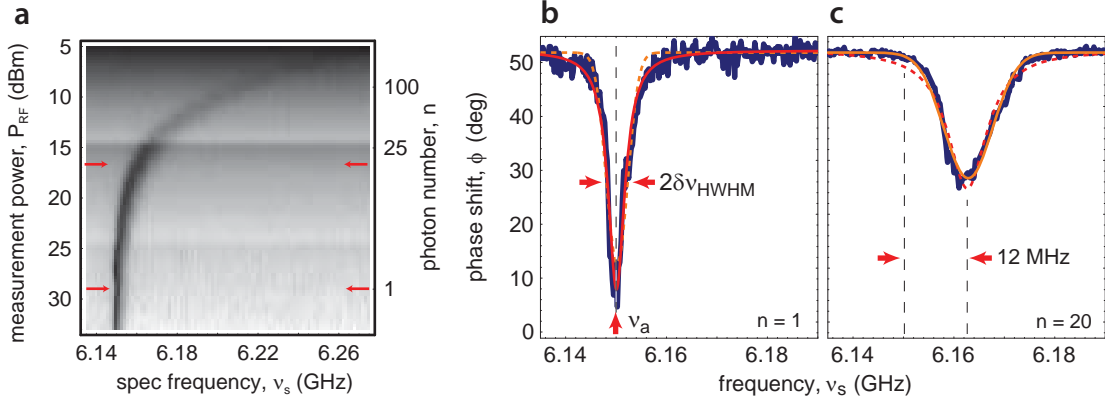
Figure 2.4: AC-Stark effect. a Spectroscopy vs. measurement drive power. The color scale is the phase shift of a drive at the cavity frequency. **b** A slice at low power ($n = 1$), shows a Lorentzian lineshape for the qubit. **b** At higher drive power ($n = 20$), the qubit response is shifted and has a broader Gaussian shape. In each plot, the red and orange traces are Lorentzian and Gaussian fits, respectively. Reproduced from [15].

other will undergo Rabi oscillations between qubit and cavity states at the rate $2g\sqrt{n}$. This technique was used to create and measure photon states in [24, 25, 29–31].

2.2.2 Dispersive regime

When the qubit and cavity are far detuned they can no longer directly exchange energy. Instead, they interact by a *dispersive* interaction that slightly modifies the energy levels. When $g/\Delta \ll 1$, the JC Hamiltonian can be expanded in powers of (g/Δ). The result to second order is [27]

$$H' \approx \hbar \left(\omega_r + \chi \sigma_z \right) a^\dagger a + \frac{\hbar}{2} \left(\omega_q + \chi \right) \sigma_z, \tag{2.14}$$

where $\chi = g^2/\Delta$. This has the form of a cavity with a qubit-state dependent frequency, $\omega_r' = \omega_r \pm \chi$ and a qubit that is slightly renormalized by the *Lamb shift*, $\omega_q' = \omega_q + \chi$. The dispersive interaction is depicted schematically in figure 2.3(b). The qubit-state dependent cavity frequency provides a means to perform non-destructive readout of the qubit state, which is used for qubit measurement throughout this thesis.

AC-Stark effect

If we group the terms of (2.14) differently, we obtain

$$H' \approx \hbar \omega_r a^\dagger a + \frac{\hbar}{2} \left(\omega_q + \chi + 2\chi a^\dagger a \right) \sigma_z. \tag{2.15}$$

In this form it is apparent that the qubit frequency depends on the number of photons in the cavity. This *AC-Stark effect* shifts the qubit frequency proportionally to $n = a^\dagger a$, the number of photons in the cavity. Fluctuations in the photon number dephase the qubit and broaden the linewidth. For instance, if the cavity is populated with a coherent state, then the cavity state is described by a Poisson distribution of photon numbers, with fluctuations that scale like $\sqrt{n}$. Figure 2.4 shows a measurement of the AC-Stark effect in a circuit QED system with a CPB qubit. At low measurement power, there are very few photons in the cavity, so the qubit lineshape is Lorentzian with a width, γ, determined by relaxation and dephasing. At higher measurement power, the qubit shifts to higher frequency. Since the shift per photon, 2χ, is smaller than the linewidth, the peaks for the individual populated photon numbers are not discernible, instead, smearing together into a Gaussian lineshape. For further discussion of this behavior, see [15, section 8.3] as well as [32, 33].

Number splitting

When the Stark shift per photon gets larger than the qubit and cavity relaxation rates, *i.e.* $2\chi > \gamma, \kappa$, the system enters the strong dispersive regime. Here, the Stark shift exerts a sufficiently strong pull on the qubit that a single photon entering the cavity shifts the qubit by more than a linewidth. In this case, when the cavity is populated with a coherent or thermal distribution, the quantized nature of the light is directly manifest in the qubit spectrum. Since the Stark shift is $2\chi = g^2/\Delta$, increasing the shift requires increasing the coupling strength or decreasing the detuning. In the work of [17], the Stark shift was made larger by using a transmon qubit, which, because of its larger size, tends to have a larger coupling strength. Figure 2.5 shows measurements of this sample for various cavity populations. In figure 2.5(a), the cavity is populated with coherent states of increasing amplitude. Since the n^{th} photon in the cavity decays at a rate $n\kappa$, the n^{th} photon number peak inherits this decay rate such that its linewidth is $\gamma_n = \gamma + n\kappa$. Consequently, as the amplitude of the coherent field is increased, the spectrum evolves from sharp distinguishable peaks to a broad, nearly Gaussian, spectrum.

The height of the peaks also reveals information about the cavity population. When
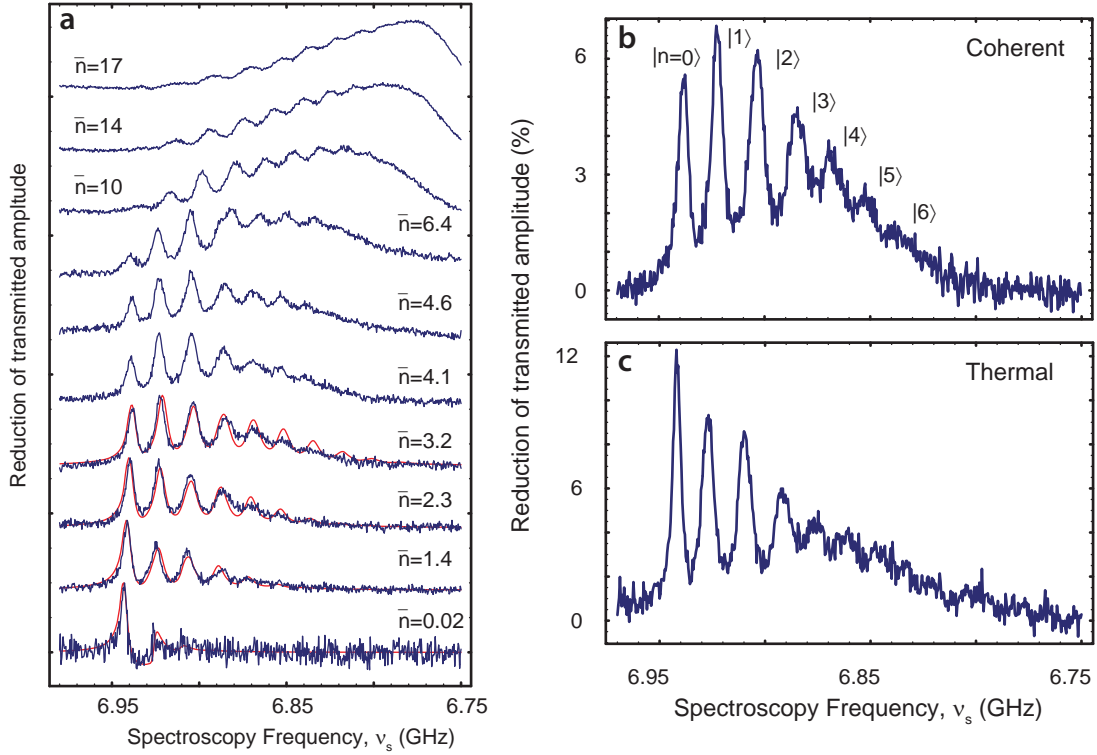
Figure 2.5: Number splitting. Spectroscopy of a transmon coupled to a cavity. **a** Spectroscopy with a coherent drive populating the cavity. **b** Coherent population. **b** Thermal population. Reproduced from [15].

the spectroscopic drive weakly perturbs the qubit state, the resulting phase shift is directly proportional to the population of a particular photon number state. Thus, the height of the peaks correspond with the statistics of the populating field. In figure 2.5(b and c), the cavity is populated with the same average photon number of $\bar{n} \sim 3$, but in one case the populating tone is a coherent field and in the other it is a thermal field. For the coherent population the peaks heights are clearly non-monotonic and are consistent with a Poisson distribution: $P(n) = e^{-\bar{n}}\bar{n}^n/n!$, whereas the thermally populated cavity has monotonic peaks that are consistent with a Bose-Einstein distribution: $P(n) = \bar{n}/(\bar{n}+1)^n$.

This experiment shows how a coupled qubit-cavity system serves as a photon statistics analyzer. However, it is not ideal. In this single cavity experiment, the same cavity is used for population and qubit readout. The same photons that shift the qubit frequency also carry information about the qubit state when they leave the cavity. This re-use of photons makes the system nonlinear, causing squeezing of the cavity states at high drive powers. Furthermore,

this re-use means that information is acquired about the photon state at the same rate as the photons leave the cavity. In chapter 8, I will show how to extend this work into a fast, QND, photon detector.

2.2.3   Generalized Jaynes-Cummings Hamiltonian

As mentioned previously, the transmon has a nearly harmonic level spectrum. Consequently, to effectively model the coupling of a transmon to a cavity, the Jaynes-Cummings Hamiltonian must be generalized to include higher transmon levels. These other states also couple to the electric field in the cavity, so we obtain [6]

$$H = \hbar\omega_r a^\dagger a + \hbar \sum_j \omega_j |j\rangle \langle j| + \hbar \sum_i g_{i,i+1}(a^\dagger |i\rangle \langle i+1| + \text{h.c.}), \tag{2.16}$$

where $\hbar g_{ij} = 2\beta e V^0_{\text{rms}} \langle i|n|j\rangle$, V^0_{rms} is the RMS of the zero-point voltage fluctuations in the cavity, and $\beta = C_g/C_\Sigma$ is the voltage division ratio determined by the gate capacitance compared to the total capacitance. In this expression, I have dropped transmon-photon couplings, g_{ij}, for non-nearest-neighbor transmon states because the matrix elements, $\langle i|n|j\rangle$, rapidly vanish as $E_J/E_C \to \infty$ for $j \neq i \pm 1$. As before, one can make a dispersive approximation when the transmon levels are far detuned from the cavity transition. Making a two-level approximation of the transmon gives an expression like (2.14), but now with [6]

$$\chi = \frac{g_{ge}^2 \alpha}{\Delta(\Delta + \alpha)}, \tag{2.17}$$

where α is the transmon anharmonicity, and $\Delta = \omega_{ge} - \omega_r$. Since $\alpha < 0$, when the detuning is large compared to the anharmonicity, the transmon χ has the opposite sign compared to the CPB.

2.2.4   Quasi-dispersive regime

The dispersive regime is convenient because the Hamiltonian separates into qubit and cavity-like terms, with corrections that are linear in the number of excitations added to the system. The price of working dispersively, though, is that most of the effects are small (there is a factor $(g/\Delta)^2$ in front of everything because of the perturbative expansion). Consequently, in situations which benefit from larger energy shifts, it is advantageous to work in a quasi-dispersive or nearly-resonant regime where $0.1 < g/\Delta < 1$.
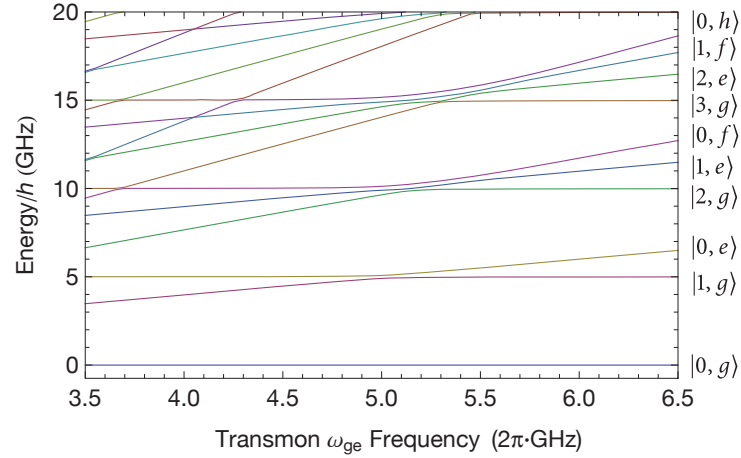
Figure 2.6: Energy levels of a transmon qubit coupled to a cavity. Numerical calculation of the energy levels of a transmon qubit coupled to a cavity, plotted as a function of the $g \to e$ transmon transition frequency in the absence of coupling. Colors are a guide to the eye, and do not identify particular states. Parameters used to make this plot are $g/2\pi = 75\,\text{MHz}$, $\omega_r/2\pi = 5.0\,\text{GHz}$, and $E_C/h = 250\,\text{MHz}$. Levels that are mostly flat, like $|1, g\rangle$, are cavity-like levels. Sloped lines are transmon-like levels, where lines of greater slope are higher up the transmon ladder. Due to the transmon's negative anharmonicity, higher states of the transmon intersect cavity levels at positive detunings of the lowest transmon transition, $\Delta = \omega_{ge} - \omega_r > 0$.

A simple way to understand the quasi-dispersive regime is that it smoothly connects the resonant and dispersive limits which we have just examined. In the resonant limit, the transition energy between states with opposite parity is $\sim 2\hbar g\sqrt{n}$, and in the dispersive limit, these transitions are $2\hbar g^2/\Delta$. Consequently, as the qubit comes into resonance with the cavity, the transition energies grow and become non-linear in the photon number, n. This picture becomes a little more complicated when we replace the qubit with a transmon, because the higher levels of the transmon are resonant with the cavity at various small detunings, Δ_{ge}, of the lowest transmon transition and the cavity.

Consequently, we need to return to calculating energy levels from the full Hamiltonian (2.16). This is relatively straightforward with Mathematica using the transmon package developed by Lev Bishop [16], and can be done efficiently as long as the Hilbert space is truncated to contain a modest number (100–1000) of energy levels. It is useful to consider a few example energy level diagrams to develop some intuition about the qualitative features of this regime.

Figure 2.6 shows numerically calculated energy levels of a transmon qubit coupled to a cavity. Finding eigenvalues in Mathematica provides the locations of the energy levels, but
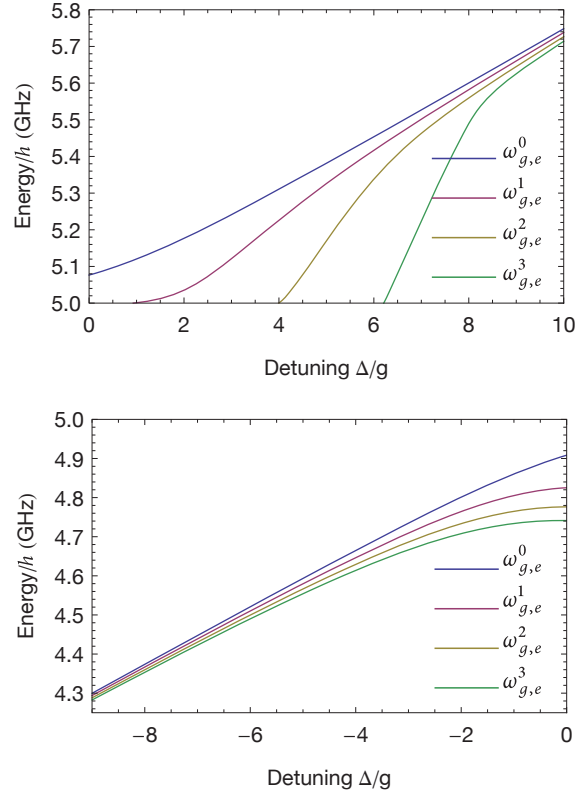
Figure 2.7: Extension of the AC-Stark Shift in the Quasi-dispersive Regime. Calculation of photon-number-dependent transition frequencies for the $g \to e$. The presentation is split for positive (top) and negative (bottom) detunings to allow for state labeling as described in the text. The large difference in apparent photon-number-dependent energy shifts at $\Delta/g = 0$ is due to a change of state labels.

without meaningful labels. One can identify the various transitions with some simple rules, though. First of all, in figure 2.6 the parameter that is changing is the Josephson energy of the transmon; therefore, levels which contain only cavity excitations should be flat because they are independent of the transmon frequency. Conversely, levels containing transmon excitations have slope, and higher levels of the transmon are more strongly dependent on E_J, so the slope of the levels increases with the number of excitations in the transmon. Armed with these rules, one can apply the labels shown on the right of figure 2.6.

   The notion of the AC-Stark shift can be extended into the quasi-dispersive regime by computing the transition energies $\hbar\omega_{g,e}^n = E_{n,e} - E_{n,g}$. Before doing this, however, one needs to decide how to extend the state labeling scheme of the last paragraph into regions where

there are avoided crossings between states. For instance, the correspondence between energy level and label is rather murky in the 3-excitation manifold for $\omega_{g,e}/2\pi$ between 5.2–5.5 GHz. The various choices for how to proceed involve choosing a parameter to adiabatically vary until the system returns to a state where the labeling is "easy". For instance, one possibility is to imagine a means to adiabatically adjust the qubit-photon coupling rate, g, and to find the resulting state when $g \to 0$. This choice is conceptually straightforward, but it is less useful in practice because g is not a parameter which can typically be varied *in situ* in current experiments. Another possibility is to adiabatically vary Δ until $|\Delta|/g \gg 1$. The direction Δ is tuned depends on the particular experimental protocol being considered, but a reasonable choice is to go to more positive (negative) detuning if Δ is initially positive (negative).[*]

This last choice of labeling by adiabatically varying the detuning is used to calculate the AC-Stark effect in the quasi-dispersive regime, shown in figure 2.7. The cases of positive and negative detunings are considered separately because the labeling protocol is only consistent if $\text{sgn}(\Delta)$ is constant. In both cases you see a transition from a region where the Stark shift is approximately linear in the photon number, to regions where it is highly nonlinear. This change is especially evident in the case of positive detunings, where the negative anharmonicity of the transmon implies that higher transmon levels will come into resonance with cavity states as the detuning is decreased. This causes a large shift in the $\omega_{g,e}^n$, a useful feature for spectrally resolving these transitions in a fast, pulsed experiment. However, the negative transmon anharmonicity also creates a proliferation of other transitions which are nearly resonant with the $\omega_{g,e}^n$ when $\Delta > 0$. This spectral crowding complicates experiments involving several-quanta excitations, which are discussed further in chapter 9.

2.3   QND measurements

Undergraduate and graduate quantum mechanics courses often present a very simple model of measurements. They teach that a quantum measurement collapses a system into the eigenstate corresponding to the observed measurement result. However, this idealized view of measurements is not automatically satisfied in practice. In particular, many measurements are destructive, leaving the system in the ground state or an unknown state. For example, a photo-multiplier tube absorbs the photons that it measures. Quantum non-demolition

---

[*] The effect of the direction of the adiabatic change of Δ is particularly apparent when the starting condition is resonance, $\Delta = 0$, where, for instance the highest energy level in the 2-excitation manifold becomes $|0, f\rangle$ when sweeping towards positive detunings, or $|2, g\rangle$ if one sweeps to negative detunings.

(QND) measurements, on the other hand, allow repeated measurements that give the same eigenvalue. The name QND emphasizes that the probe minimally perturbs the system. Note that this is not the same as saying that the probe has no back-action. For example, if a quantum oscillator is in a coherent state, it occupies a superposition of energy eigenstates with a well-defined phase relationship. When one performs a QND measurement of the energy of the oscillator, the state is projected into an energy eigenstate and the phase of the oscillator (the conjugate variable to the energy) is now completely uncertain. Consequently, the QND measurement has perturbed the system. Importantly, however, the probe has not *demolished* the state, and so a subsequent measurement of the oscillator energy will return the same result. This example is particularly relevant to the measurements described in chapter 8, where I present a QND method for measuring the photon number in a cavity. In this section I will give a brief summary of the requirements for QND measurements. For a more in-depth treatment, see Refs. [34, 35].

To adequately describe quantum measurements, one should consider a system + meter model, where the system, S, with Hamiltonian H_S, and meter, M, with Hamiltonian H_M, are coupled by an interaction Hamiltonian, H_{SM}. The total Hamiltonian of the system + meter is

$$H = H_S + H_M + H_{SM}. \tag{2.18}$$

The interaction, H_{SM}, should be such that it couples an observable of the system, O_S, to an observable of the meter, O_M. A necessary and sufficient condition for a QND measurement of O_S is that [34]

$$[O_S, U]|\psi_M\rangle = 0, \tag{2.19}$$

where U is an operator describing the joint evolution of the system and meter, and $|\psi_M\rangle$ is the initial quantum state of the meter. Often it is difficult to find U exactly, so it is common to make the additional assumption that O_S is a constant of the motion, *i.e.*

$$i\hbar\frac{\partial O_S}{\partial t} + [O_S, H] = 0. \tag{2.20}$$

When O_S is also a constant of motion of H_S, (2.20) is equivalent to

$$[O_S, H_{SM}] = 0. \tag{2.21}$$

This last condition is the one most often associated with QND measurements; however, it is

merely a sufficient, not necessary, condition for QND. Nonetheless, all QND measurements discussed in this thesis will approximately satisfy the more stringent (2.21).

Equation (2.21) is an ideal that is never completely satisfied in practice because no experiment implements perfect control on a system. Thus, for real experiments it is useful to also define the *demolition* of a measurement. One sensible definition of demolition is the probability that the measurement process, U, changes the observable, O_S. Thus, for a particular value of $O_S = v$, with system eigenvector $|v\rangle$, the demolition is

$$D = \sum_{u \neq v} |\langle u | U | v \rangle|^2 . \tag{2.22}$$

This definition serves as a figure of merit for comparing QND measurement protocols, as it determines the number of repetitions of a measurement protocol that may be performed before additional measurements do not reveal additional information about an initial state.

2.4   Quantum Optics Background

The last chapters of this thesis primarily deal with creating and detecting non-classical states of light in a cavity. This leads naturally to two specific issues. The first is, what is a 'non-classical' state of light? The second is, how does one go about detecting these states, or how does one perform tomography of a cavity state? It turns out that we can answer both questions by introducing the concept of *quasi-probability distributions*, including a few particular examples which are commonly used in the quantum optics community. What follows is a brief overview that is largely drawn from [21] and [13]. For a more complete discussion, I refer you to those sources.

2.4.1   Representations of cavity modes

A single-mode cavity is a harmonic oscillator. Consequently, it has an equivalent description in terms of a particle with position X and momentum P described by the Hamiltonian

$$H = \frac{P^2}{2m} + \frac{m\omega^2}{2} X^2 . \tag{2.23}$$

A classical description of this oscillator gives a definite value for X and P at any moment in time, but since X and P are conjugate variables, the Heisenberg uncertainty principle prevents complete, simultaneous knowledge of both. Thus, if we represent the state of the
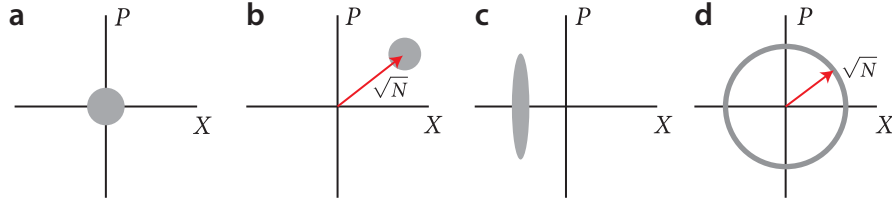
Figure 2.8: Phase space representation of cavity states. a Vacuum state. **b** Coherent state of amplitude $\alpha = \sqrt{N}$. **c** Squeezed state. **d** Fock state of N photons.

oscillator in the X-P plane, the quantum description requires that the state be described by a region of finite volume, rather than with a point.

Figure 2.8 shows some example phase space representations of cavity states.* Figure 2.8(a) and (b) show a vacuum state and a coherent state, respectively. These cavity states are represented as fuzzy disks where the displacement from origin is equivalent to the average amplitude and phase of the field. They are conventionally considered 'classical' because in the limit of large N, the quantum fluctuations are proportionally very small, so the representation becomes nearly point-like. Figure 2.8(c) shows a 'squeezed' state, where the disk is now elongated because the vacuum noise has been reduced in one direction and amplified in the other. Such states are useful because they allow a measurement of one quadrature of a field with a precision that is beyond the 'standard quantum limit'. Consequently, squeezed states are usually considered quantum states, despite having strictly positive Wigner functions, which we will define in a moment. Figure 2.8(d) shows a Fock state (an energy eigenstate of the oscillator) with exactly N photons. Fock states have definite amplitude, but no phase information, so they are circles in phase space. These various states clearly look very different in their phase space representations, but we have not yet given a precise way to distinguish quantum from classical.

In both the quantum and classical cases, the expectation value of any operator O which is a function of the field quadratures can be computed by integrating over phase space with a weighting function $f(X, P)$,

$$\langle O \rangle = \int O(X, P) f(X, P) dX dP. \tag{2.24}$$

For classical states, $f(X, P)$ must be strictly positive, and so it has the meaning of a probability

---

* Due to the 'blob' or lollipop appearance of coherent states, this is sometimes colloquially referred to as 'blob space'.

distribution. For quantum states, however, f can also take on negative values, so it is called a *quasi-probability distribution*. There are several such quasi-probability distributions, the most common of which are the Wigner W and Husimi Q functions.

To my taste, the Q function is actually the easiest to understand. The set of coherent states, α, form a complete basis (actually, an over-complete basis) for the space of cavity states. Thus, if the density matrix ρ describing the cavity state is written in the basis of coherent states, then Q is simply related to the expectation value of the diagonals of the density matrix,

$$Q(\alpha) = \frac{1}{\pi}\langle \alpha | \rho | \alpha \rangle, \tag{2.25}$$

Introducing the displacement operator,

$$D(\lambda) = \exp(\lambda a^\dagger - \lambda^* a), \tag{2.26}$$

which displaces a cavity state by a coherent state $|\lambda\rangle$, we can write,

$$Q(\alpha) = \frac{1}{\pi}\langle 0 | D(-\alpha)\rho D(\alpha) | 0 \rangle. \tag{2.27}$$

This form presents an alternate interpretation that $Q(\alpha)$ is the projection onto the vacuum state when the field is displaced by $-\alpha$. Unlike W, the Q function is strictly positive, though it theoretically contains the same information.*

The Wigner function essentially takes the form of a Fourier transform of the density matrix, ρ. When expressed in the position basis, it is written as[†]

$$W(x, p) = \frac{1}{\pi} \int du\, e^{-2ipu}\langle x + u/2 | \rho | x - u/2 \rangle. \tag{2.28}$$

Since the Fourier transform can be inverted, the Wigner function contains the same information as the density matrix. The Wigner function can be written in another form by introducing the parity operator $\mathcal{P} = \exp(i\pi a^\dagger a)$ that reflects a state about the phase space

---

* The Q function is essentially the convolution of W with a Gaussian. In principle, one can deconvolve a measurement of Q to find W, though in practice the presence of added noise may make it impossible to recover the negative (quantum) features in W by this method. Thus, to verify the quantum nature of a cavity state, it is preferable to measure W directly.

† Due to the usual ambiguity in defining the Fourier transform, there are several conventions for the definition of W. The definition here has $-2/\pi \le W(\alpha) \le 2/\pi$, but other sources define it such that $-2 \le W(\alpha) \le 2$.

origin, then [13]

$$W(\alpha) = \frac{2}{\pi} \, \text{tr} \left[D(-\alpha) \rho D(\alpha) \mathcal{P} \right]. \tag{2.29}$$

Since W is the expectation value of an operator, it is a measurable quantity, and because the parity operator has eigenvalues ± 1, the Wigner function is bounded with $-2/\pi \le W(\alpha) \le 2/\pi$. Continuing by expressing $\mathcal{P}$ in the number basis,

$$\mathcal{P} = \sum_n (-1)^n |n\rangle \langle n|, \tag{2.30}$$

we have

$$W(\alpha) = \frac{2}{\pi} tr \left[\sum_n D(-\alpha) \rho D(\alpha) (-1)^n |n\rangle \langle n| \right], \tag{2.31a}$$

$$= \frac{2}{\pi} \sum_n (-1)^n tr \left[\langle n| D(-\alpha) \rho D(\alpha) |n\rangle \right], \tag{2.31b}$$

$$= \frac{2}{\pi} \sum_n (-1)^n p_\alpha(n), \tag{2.31c}$$

where $p_\alpha(n)$ is the probability to be in Fock state $|n\rangle$ after displacing the field by $-\alpha$. This reveals a prescription for measuring the Wigner function by summing Fock state probabilities. This method has already been used in experiments[24, 31, 36–38] which will be discussed in section 2.5.

The Wigner function has several properties that make it similar to the classical probability distribution discussed previously. In particular, if $W(x, p)$ is integrated along either the x or p direction, the result is the probability distribution in the conjugate variable, $i.e.$

$$P(x) = \langle x | \rho | x \rangle = \int_{-\infty}^{\infty} dp \, W(x, p), \tag{2.32a}$$

$$P(p) = \langle p | \rho | p \rangle = \int_{-\infty}^{\infty} dx \, W(x, p). \tag{2.32b}$$

Furthermore, for any observable $O(x, p)$, which is a *symmetric* function of x and p,

$$\langle O(x, p) \rangle = \int dx dp \, O(x, p) W(x, p). \tag{2.33}$$

These features make the Wigner function analogous to a classical probability distribution. However, certain cavity states, like Fock states, have nodes in their probability distributions.

That is, there exists some x_0 for which $P(x_0) = 0$. By (2.32) this implies that

$$P(x_0) = \int_{-\infty}^{\infty} dp\, W(x_0, p) = 0. \tag{2.34}$$

Consequently, $W(x, p)$ must also take on negative values. One can show [21] that this is only true of non-Gaussian states. Consequently, we can say that having a Wigner function with negative values is a clear signature that a cavity state is non-classical. Some example Wigner functions appear in section 2.5

2.5   Creating and measuring quantum states of light

There has been tremendous progress in the past decade in the preparation and detection of non-classical states of light in cavities. In the following section, I will discuss prior experiments demonstrating preparation and Wigner tomography of quantum states of electromagnetic cavities where a qubit is used to 'measure' the field.* Additional state preparation techniques will be discussed in section 2.5.1. Before diving into cavity experiments, I should mention that the first measurement of a negative Wigner function was actually tomography of the quantized motional state of an ion in a harmonic trap. Leibfried *et al.*[40] used a resonant technique that is analogous to methods employed with photons in electromagnetic cavities [24, 31, 37], so I will focus on these later experiments because we have already developed all the physics needed to understand their workings.

Resonant Rabi oscillations

One Wigner tomography technique uses a resonant interaction between a cavity field and a qubit. The cavity state is prepared with the qubit initially far detuned, then the qubit is brought suddenly into resonance for a time τ. Since the resonant Jaynes-Cummings states of (2.13) are split by $2\hbar g \sqrt{n}$, the qubit and cavity undergo coherent oscillations. When the cavity state is a pure Fock state, these Rabi oscillations are between $|n, g\rangle$ and $|n - 1, e\rangle$ and have a single frequency component. For a general cavity-qubit state written in the Fock basis and initialized with the qubit in $|g\rangle$,

$$|\psi\rangle = \sum_n c_n |n, g\rangle, \tag{2.35}$$

---

* It is also possible to perform Wigner tomography by homodyne detection of the emitted radiation. Such a technique was used in [39] for tomography of squeezed states.
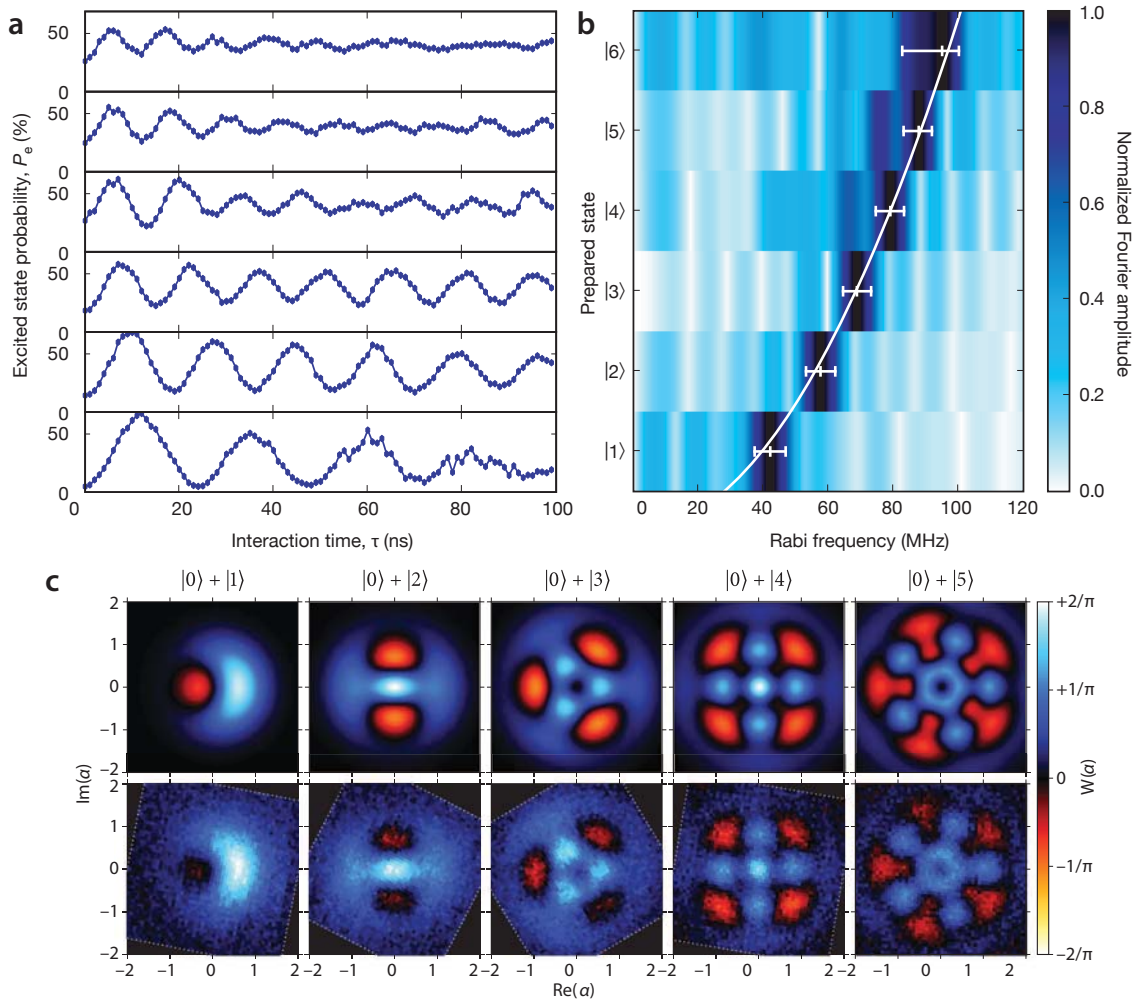
Figure 2.9: Wigner tomography by resonant Rabi interaction. a Time-domain Rabi oscillations of a phase qubit with a resonant cavity for Fock states with $n = 1$–6. **b** Fourier transforms of the time-domain data of **a** show peaks at the expected $2g\sqrt{n}$ frequencies (white curve). **c** Calculated (top) and measured (bottom) Wigner functions for $|0\rangle + |N\rangle$ states for $N = 1$–5. Reprinted from [24] and [31] with minor modifications.

where the $\{c_n\}$ are complex amplitudes, the resonant interaction produces the state

$$|\psi'\rangle(\tau) = \sum_n c_n \left(e^{i(\omega_r - g\sqrt{n})\tau} |n, g\rangle + e^{i(\omega_r + g\sqrt{n})\tau} |n-1, e\rangle \right). \tag{2.36}$$

Performing subsequent qubit read-out thus gives a probability to find the qubit in $|e\rangle$ of

$$P_e(\tau) = \sum_n |c_n|^2 \left(\frac{1 + \cos(2g\sqrt{n}\tau)}{2} \right). \tag{2.37}$$

The Fourier transform of this probability with respect to τ gives peaks at frequencies $2g\sqrt{n}$ with amplitudes $|c_n|^2$, which are the occupation probabilities of the Fock states. Consequently, if one prepares an ensemble of identical states $|\psi\rangle$ and averages the resulting qubit state of the time-domain Rabi oscillations for many τ, the Fourier transform of these oscillations reveals the photon number probabilities. The Wigner function, $W(\alpha)$, is found by displacing the cavity with a coherent state $|\alpha\rangle$ and using (2.31) to compute $W(\alpha)$ with the measured photon number probabilities.

Preparation of cavity states is done with an inverse procedure, where photons are added to the cavity one at a time by preparing the qubit in $|e\rangle$ and bringing it into resonance with the cavity to perform a full SWAP gate. Since the Rabi frequency increases as $\sqrt{n}$, each subsequent SWAP is shorter in duration. Law and Eberly found a deterministic method for synthesizing arbitrary states using just this interaction by varying the initial qubit state before each SWAP [23]. The sequence of steps required to create any state is constructed by starting with the final state, and then finding the operations which remove a photon at a time. The sequence is then inverted to reveal the preparation procedure. See [23] and [31] for details.

Figure 2.9 shows data from two experiments [24, 31] from the Martinis group at UCSB. In these experiments, a superconducting phase qubit is coupled to a high-Q microwave cavity with a coupling rate $g/2\pi \simeq 20\,\text{MHz}$. The researchers prepare Fock states in the cavity for $n = 1$ to 6 and measure the resulting time-domain Rabi oscillations between the phase qubit and cavity, shown in figure 2.9(a). Sinusoidal oscillations of the qubit excited state probability are observed with a frequency that increases with n. The Fourier transform of these data in figure 2.9(b) show peaks at the expected $2g\sqrt{n}$ frequencies, demonstrating the use of the resonant interaction to extract photon number probabilities. In figure 2.9(c), Hofheinz *et al.* proceeded to measure the Wigner functions of $|0\rangle + |N\rangle$ states for $N = 1$–5 by applying coherent microwave pulses at the cavity frequency before measuring the time-domain Rabi

oscillations. The correspondence between the calculated (top) and measured (bottom) $W(\alpha)$ demonstrates that this is an effective technique.

Recently, Fink *et al.* showed a variant of this technique that uses sideband transitions to directly drive Rabi oscillations between a qubit and cavity [25] while they are detuned. In effect, a strong drive replaces qubit frequency excursions. This has been used to create Fock states with $n = 1$ to 5, though they have not yet demonstrated full Wigner tomography.

Despite the clear success of the resonant swap method, it has a substantial weakness in that it demolishes the cavity state. For example, imagine a cavity that is prepared in a Fock state, $|N\rangle$. In any particular realization of a time-domain Rabi experiment, there are many interaction times for which the qubit occupies the excited state with unit probability. For these interaction times, exactly one photon is subtracted from the cavity, leaving the cavity in the state $|N - 1\rangle$. One could modify the procedure to only sample the time-domain Rabi oscillation at intervals that leave the photon in the cavity, *i.e.* sample at $\tau_m = 2\pi m/(2g\sqrt{N})$ for integer m. When m is odd this corresponds to a 2π rotation, causing the qubit $|g\rangle$ state to pick up a global phase factor of -1. This can be detected if the experiment has access to an additional qubit and a two-qubit gate or additional levels of the 'qubit', as was used for QND detection of $|0\rangle$ and $|1\rangle$ using Rydberg atoms in [41]. However, this only works for a particular N and modifies the cavity state when the cavity is prepared in other Fock states.

It is important to realize that the photon demolition of the resonant scheme does not detract from Wigner tomography. Since the Wigner function has the full information of the density matrix, it cannot, even in principle, be measured in a single realization. Rather, it explicitly requires access to an ensemble of identical states in order to average. Therefore, the resonant scheme is sufficient for Wigner tomography. It is not usable, though, to perform continuous monitoring of the evolution of a cavity state. For example, the resonant scheme cannot be used to observe the quantum jumps of the decay of a Fock state in a cavity.

Dispersive Ramsey interferometry

Dispersive Ramsey interferometry is another method which uses a dispersive interaction between a qubit and cavity to perform QND detection of photon number. This scheme uses the AC-Stark shift discussed in section 2.2.2 that causes the qubit frequency to shift linearly with the photon number, $\omega'_q = \omega_q + 2\chi n$. Consequently, one can construct a photon number measurement by performing a Ramsey experiment. For simplicity, let us initially consider the cavity to be in a Fock state, such that the initial state is $|N, g\rangle$ and the qubit
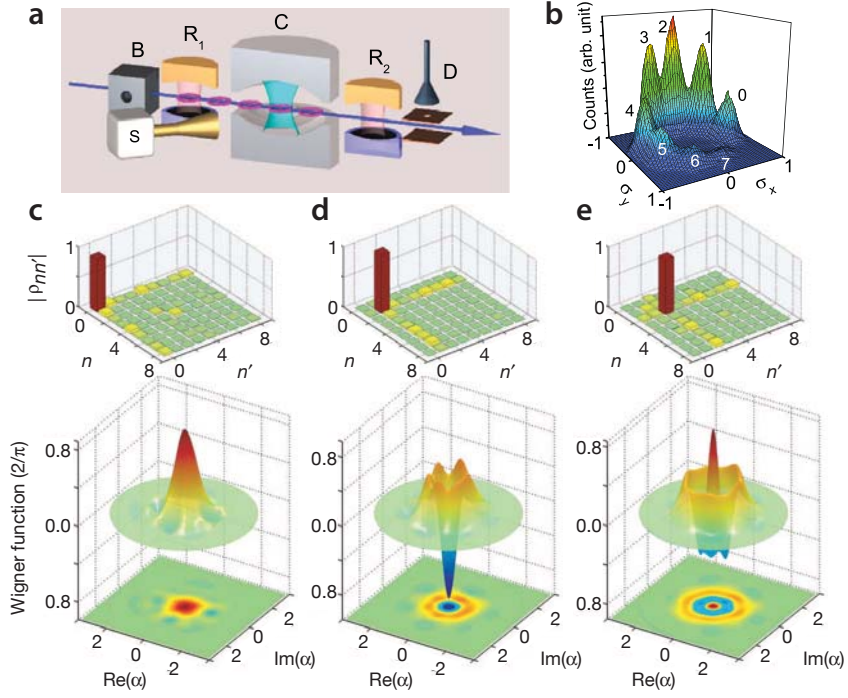
Figure 2.10:   Wigner tomography by dispersive Ramsey interferometry. a Experimental setup for a Ramsey interferometer that probes the state of a microwave cavity, C, by repeated Ramsey experiments on Rydberg atoms produced by the oven, B, and detected with an ionizing detector, D. The two smaller cavities, R_1 and R_2 allow for $\pi/2$ pulses on the qubit formed by the $n = 50$ and $n = 51$ orbital states of the atoms. **b** Measured histogram showing the distribution of phases around the Bloch sphere when the cavity is populated with a coherent state. Peaks corresponding to photon numbers 0–7 are labeled. **c** Measured density matrices and Wigner functions for Fock states with (from left to right) $N = 0, 1,$ and 2. The negative values of the Wigner function for $|1\rangle$ and $|2\rangle$ are clear signatures of the quantum nature of the light. Reprinted from [42] and [38] with minor modifications.

and cavity are far detuned. Then we perform a $X_{\pi/2}$ pulse on the qubit*, taking the state to $(|N, g\rangle - i\, |N, e\rangle)\, /\sqrt{2}$. Next, the AC-Stark interaction is turned on for a fixed time τ by reducing the qubit-cavity detuning or physically passing the qubit through the cavity. Writing the resulting state in a frame rotating at the the bare qubit frequency, ω_q, gives

$$\frac{1}{\sqrt{2}} \left(|N, g\rangle + \mathrm{e}^{i(-\pi/2 + \theta N)} |N, e\rangle \right), \tag{2.38}$$

---

* I use the convenient notation of referring to a rotation, $R_X(\theta)$, as X_θ.

where $\theta = 2\chi\tau$. Consequently, after its interaction with the cavity, the qubit takes on one of a discrete set of phase shifts corresponding to the number of photons in the cavity. This phase can be measured by ensemble averaging the qubit excited state probabilities after applying a second $\pi/2$ pulse around X or Y. The ensemble average probabilities after the second pulses reveal the projections of the qubit state vector along the Y and $-X$ directions of the Bloch sphere, respectively, thus measuring the photon number. * If the cavity is in a superposition of several photon number states, this same procedure gives a distribution of phases from which the photon number probabilities, $p(n)$, can be extracted. Calculation of the Wigner function then proceeds exactly as with the resonant scheme.

Data from a cavity QED experiment from the Haroche group implementing this method is shown in figure 2.10. Their setup sends a stream of Rydberg atoms, atoms excited to the $n \simeq 50$, $l = n - 1$ orbital, through a high-Q Fabry-Pérot cavity. The Ramsey interferometer is formed by a pair of additional low-Q cavities that allow for $\pi/2$ pulses on the atoms before and after they interact with the high-Q cavity, C. Figure 2.10(b) shows a histogram of the average phase shift of an atom passing through the cavity initially prepared in a coherent state. The discrete peaks in the phase clearly show the quantized nature of the light. The cavity field is displaced by means of a microwave source, S, which allowed the researchers to measure Wigner functions for several quantum states of light, such as the Fock states shown in figure 2.10(c).

Unlike the resonant method, this dispersive scheme is QND. This is evident from the form of the interaction between the meter (the atom or qubit) and the system (the photons in the cavity), which in the dispersive limit is $H_{SM} = 2\hbar\chi\sigma_z a^\dagger a$. Consequently, $[H_{SM}, \hbar\omega_r a^\dagger a] = 0$. The QND nature of the Ramsey interferometer method allows it to be used for essentially continuous monitoring of the photon number in the cavity. Since each measured atom carries partial information about the phase, the dynamics of how the cavity states evolves with each measurement is quite complex and beyond the scope of this thesis. To read about that topic, see [14, 42]. It is sufficient to say that the experimental setup of the Haroche group allows thousands of atoms to cross the cavity on the timescale of the decay of a single photon. This allows analyzing 10's or 100's of atoms to converge on a particular photon number, and then watch the discontinuous jumps of the state as the photons decay [14, 44]. Note that this allows the preparation of Fock states through measurement, as the repeated measurement

---

* There has been substantial work developing optimal procedures to minimize the number of measurements needed to determine the phase when the phase can only take on a small number of discrete values. See, for example, [43].

procedure collapses any initial state onto a Fock state. To create a particular Fock state, the cavity is initially populated with a coherent state. This cavity state is then repeatedly measured, and iterations which converge to the wrong state are thrown out. The resulting Fock state is random, coming from the Poisson distribution of photon numbers in the initial coherent state. To increase the probability of collapsing onto the desired state, the amplitude of the initial coherent state is adjusted to maximize the population of the target Fock state.

The are several other methods for Wigner tomography in the dispersive regime. One is a minor modification of the above procedure where the cavity-qubit interaction time, τ, or the detuning, Δ, is changed such that the phase shift per photon is exactly $\theta = \pi$. Then all odd numbered photon number states impart a π phase shift to the qubit, while even number states impart a 2π phase shift. Consequently, one can choose the phase of the second $\pi/2$-pulse in the Ramsey interferometer to map odd numbered photon states to $\sigma_z = -1$ and even numbered photon states to $\sigma_z = +1$. Then,

$$P_e - P_g = \langle \sigma_z \rangle = \sum_n (-1)^n p(n) = \text{tr}\left[\mathcal{P}\,|n\rangle\langle n|\right]. \tag{2.39}$$

Consequently, when the phase shift per photon is π, the Ramsey interferometer maps a measurement of the parity operator onto $\langle \sigma_z \rangle$, allowing a direct measurement of the Wigner function after applying displacement fields. This method has not been used in experiments because, in practice, the phase shifts are not completely linear in the photon number. The Haroche experiments, for instance, have $\Delta/g \simeq 12$, for which there is approximately 3% non-linearity [14]. To remove this source of error, they use smaller phase shifts and calculate the Wigner function from the extracted photon number probabilities.

A qualitatively different method can be achieved by going to the strong dispersive regime where the Stark shift is much larger than the qubit linewidth. In this case, qubit manipulations can be made photon-number dependent by using narrow bandwidth pulses and choosing pulse frequencies to select a single transition. Details of this method and its implementation are presented in chapter 8.

2.5.1 Rapid Adiabatic Passage

The task of cavity-state preparation is one that admits many possible solutions. For full arbitrary state preparation, the most straightforward approach is to use time-domain vacuum Rabi oscillations to transfer excitations from qubit to photon, one at a time using the Law and Eberly protocol [23, 24, 31]. Despite the power of this technique, it is not always the most

efficient means to create a particular state. For instance, if one is interested in synthesizing a pure state with a definite photon number, n, then there are alternate approaches which can create such states in fewer steps. For instance, the post-selection technique used in the Haroche experiments [14, 44] uses repeated measurements to prepare Fock states. In that case, several measurements are required to converge on a number state because each measurement carries information about several Fock states. Chapter 8 demonstrates a new QND detection protocol that is highly selective to particular number states, improving the efficiency of the post-selection method to succeed or fail in a single step. Despite this improvement, the resulting method is still probabilistic.

Rapid adiabatic passage is an efficient, and deterministic method for creating Fock states when coupling a multi-level atom to a cavity. In this approach, an atom-like state is adiabatically transferred into a photon-like state by sweeping through an atom-photon avoided crossing. For instance, for a transmon qubit coupled to a cavity, there are many such crossings, as shown in figure 2.6. To create the state $|n = 1, g\rangle$ one first prepares $|n = 0, e\rangle$ by Rabi driving the transmon. Then E_J is adjusted through a local flux bias to tune the ω_{ge} transition through its avoided crossing with the $n = 1$ state of the cavity. The timing of the sweep is very robust to jitter, so long as the sweep is sufficiently slow as to prevent Landau-Zener tunneling to the other branch of the avoided crossing.

Landau and Zener independently derived an asymptotically exact solution to the tunneling probability for the situation of two coupled levels where the sweep is of infinite duration and the initial and final states are at infinite detuning [45, 46], *i.e.* $\Delta(t) = \Gamma t$. In this case they found that the probability to tunnel from one branch to the other is given by

$$P \simeq \exp\left(-2\pi \frac{(2g)^2}{\Gamma}\right),$$ (2.40)

where $\Gamma = d\Delta/dt$ is the sweep rate. This expression is asymptotically exact in the limit $\Gamma \ll g^2$.

While (2.40) is only exact for the particular idealized case considered by Landau and Zener, it also provides a fairly good approximation in situations where the sweep duration is finite. To check this one can numerically solve the time-dependent Schrödinger equation with the generalized Jaynes-Cummings Hamiltonian,

$$H = \hbar\omega_r a^\dagger a + \hbar \sum_j \omega_j |j\rangle\langle j| + \hbar \sum_i g_{i,i+1}\left(a^\dagger |i\rangle\langle i+1| + \text{h.c.}\right).$$ (2.41)

In order for this to be numerically tractable, one has to choose a finite subset of the full

Figure 2.11: Numerical Simulation of Rapid Adiabatic Passage for Preparing $|1, g\rangle$. a Populations vs time of the upper and lower branches of the $|1, g\rangle$ and $|0, e\rangle$ avoided crossing when starting in lower branch at $t = 0$. States are labeled by the corresponding final state when $\Delta/g \gg 1$ (see section 2.2.4). **b** Final populations of $|1, g\rangle$ and $|0, e\rangle$ vs sweep time Δt. The qubit frequency is adjusted linearly from $\omega_{ge}/2\pi = 4.828$ GHz to 5.4 GHz in time Δt. Dashed line in **b** is the analytical expression (2.40) for an equivalent sweep rate to the finite duration simulations. Parameters (chosen to match the experiment of chapter 8) are: $\omega_r/2\pi = 5.062$ GHz, $E_C/h = 300$ MHz, and $g_{ge}/2\pi = 68$ MHz. Calculation Hilbert space is truncated to 5 cavity and 5 transmon levels. Simulation does not include relaxation.

Hilbert space. Figure 2.11 shows the result of one such numerical calculation keeping 5 cavity and 5 transmon levels. The parameters for the simulation are chosen to match the experiment discussed in chapter 8.* From the time evolution of the populations of the two branches, one sees that there is some dynamics to the evolution—the sweep rates are such that there is some part of the wavefunction that oscillates between the two states. However, since only the final populations matter, it may be possible to use these oscillations to modify the ramp shapes to create faster ramps with a small final probability to be in the other branch. I did not explore this much since the calculation already shows that it is possible to have a very low probability to tunnel with a linear sweep which is fast relative to qubit and cavity decay times. Figure 2.11b shows the final populations as a function of the sweep rate. For comparison, the Landau-Zener formula (2.40) is also plotted as a dashed line. It has similar behavior and is of the same order of magnitude, but does not capture the oscillatory dynamics, which serve to suppress the tunneling rates for the parameters chosen in the simulation.

For crossings of more than 2 levels, (2.40) is of less utility. For instance, consider the case of attempting to prepare $|2, g\rangle$ by rapid adiabatic passage. One might expect to be able to sweep faster when preparing $n = 2$ than when preparing $n = 1$, because the higher photon number splittings in the standard Jaynes-Cummings ladder are $g_{ge}\sqrt{n}$. In the generalized J-C, though, with $2\hbar g_{ge} \sim E_C$, the transmon $|f\rangle$ state comes into resonance with the cavity in the same vicinity as the avoided crossing between $|1, e\rangle$ and $|2, g\rangle$ (see figure 2.6). This $|1, e\rangle$ and $|0, f\rangle$ splitting goes like $g_{ef} \approx g_{ge}\sqrt{2}$. Consequently, the state which connects with $|2, g\rangle$ is actually $|0, f\rangle$. A simulation of this situation is shown in figure 2.12. One can see that very fast sweeps actually tunnel through both crossings to the highest branch. The avoided crossings are bigger, but there are now two of them. So, at least for the simulated parameters, one actually needs to sweep slower in order to stay in the same branch.

This procedure can easily be extended in order to examine state preparation of higher photon number Fock states, so long as one increases the size of the Hilbert space to include the relevant levels.† Investigating whether this method can be turned into a scheme for arbitrary state preparation could be an interesting avenue for further research. Regardless, this method provides a robust method for preparing Fock states without requiring precise timings of flux pulses. This is particularly important when the vacuum Rabi frequency is

---

* The Mathematica code used to generate these plots is in appendix A.

† There probably is a simple method to re-truncate the Hilbert space to only include those energies levels which participate in the relevant avoided crossings.

a

b

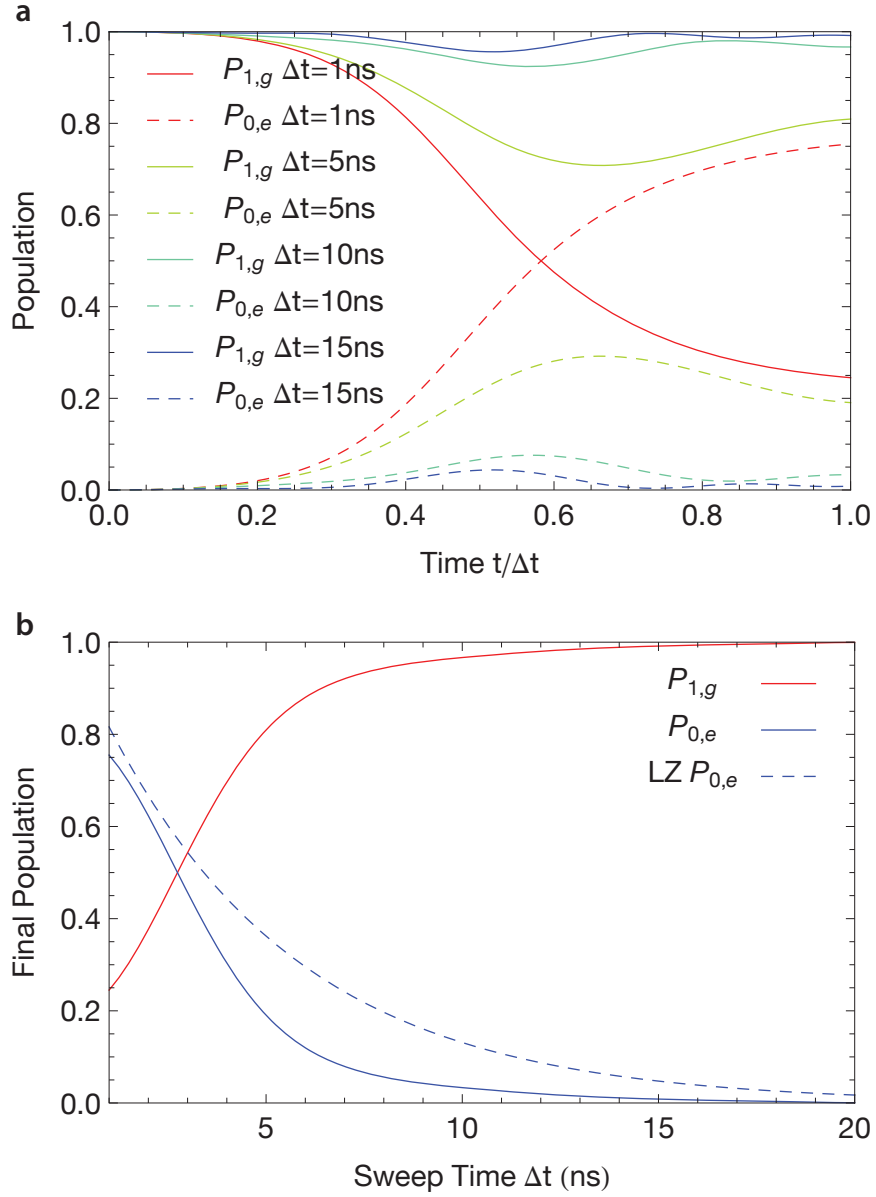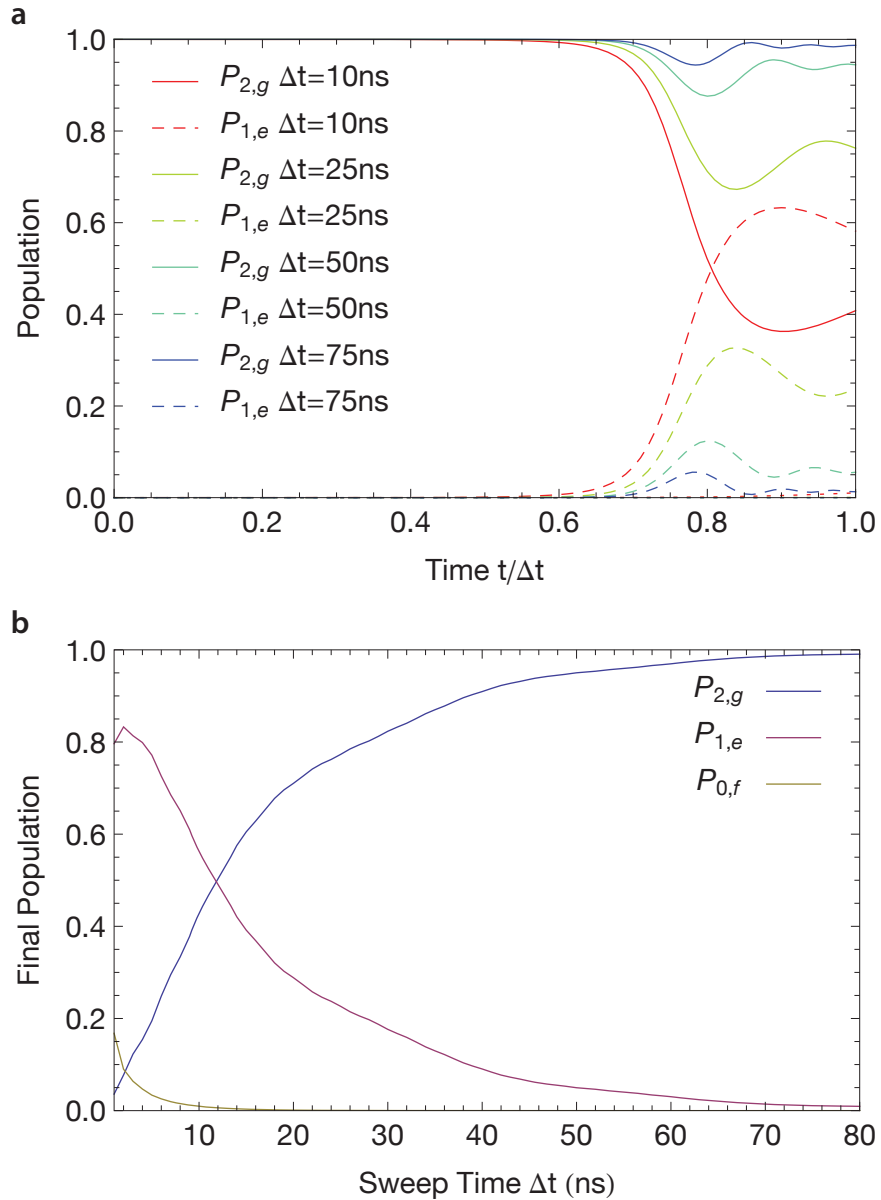**Figure 2.12: Numerical Simulation of Rapid Adiabatic Passage for Preparing $|1, g\rangle$. a** Populations vs time of the 3 branches of the $|2, g\rangle$, $|1, e\rangle$, $|0, f\rangle$ avoided crossing when starting in the lowest branch at $t = 0$. States are labeled by the corresponding final state when $\Delta/g \gg 1$ (see section 2.2.4). $P_{0,f} \approx 0$ for the simulated parameters, so it is not shown. **b** Final populations of $|2, g\rangle$, $|1, e\rangle$, and $|0, f\rangle$ vs sweep time Δt. The qubit frequency is adjusted linearly from $\omega_{ge}/2\pi = 4.828\,\text{GHz}$ to $5.4\,\text{GHz}$ in time Δt. Parameters are the same as in figure 2.11.

very large, where the duration of a SWAP gate may be faster than can be reliably produced with room temperature electronics. This was the driving motivation for using the adiabatic passage technique for the experiment of chapter 8.

Theory of Two-Cavity Architecture

THE theory of circuit QED with one qubit and one cavity has been well documented elsewhere [15, 27, 47]. This chapter will extend the existing formalism to two cavities and one qubit, and explore the consequences of this architecture, such as the cross-Kerr effect and new cavity relaxation channels from coupling to a qubit. I also introduce a modified transmon design that couples to two cavities while introducing minimal indirect capacitive coupling between the cavities.

3.1   The Two Cavity Hamiltonian

The usual Jaynes-Cummings Hamiltonian can be extended to describe two cavities coupled to a single qubit by simply adding additional cavity and cavity-qubit coupling terms. The resulting Hamiltonian is

$$H = \sum_{i=1}^{2} \hbar\omega_i a_i^\dagger a_i + \frac{\hbar}{2}\omega_q \sigma_z + \sum_{i=1}^{2} \hbar g_i(a_i^\dagger \sigma^- + a_i \sigma^+), \tag{3.1}$$

where we neglect, for the moment, direct cavity-cavity coupling terms of the form $\lambda(a_i^\dagger a_j + a_j^\dagger a_i)$.

To get the dispersive (g_1/Δ_1, $g_2/\Delta_2 \ll 1$) two-cavity Hamiltonian, one can extend the

method used in [27] and apply the unitary transformation $H' = UHU^\dagger$, where

$$U = \exp\left[\sum_i \frac{g_i}{\Delta_i}(a_i\sigma^+ - a_i^\dagger\sigma^-)\right]. \tag{3.2}$$

After using the Baker-Hausdorff lemma and a lot of algebra, one obtains (to order g_i^2/Δ_i^2)

$$H \approx \sum_i \hbar\omega_i a_i^\dagger a_i + \frac{\hbar}{2}\left(\omega_q' + \sum_i \frac{2g_i^2}{\Delta_i}a_i^\dagger a_i\right)\sigma_z + \hbar g_{12}\sigma_z(a_1^\dagger a_2 + a_1 a_2^\dagger), \tag{3.3}$$

where $\omega_q' = \omega_q + \sum_i g_i^2/\Delta_i$ is the Lamb-shifted qubit frequency and

$$g_{12} = \frac{g_1 g_2(\Delta_1 + \Delta_2)}{2\Delta_1\Delta_2} \tag{3.4}$$

is the qubit-mediated cavity-cavity coupling rate. In (3.3) one can identify the second set of terms as the AC-Stark shift from photons in each cavity. The final term allows for qubit-state dependent mixing of photons between the cavities. Section 3.3 describes one effect of this mixing, which is to introduce a new mechanism for photon decay in each cavity. For now, however, we want to remove this term in order to diagonalize the Hamiltonian. Therefore, we try an additional transformation of the form

$$U' = \exp\left[\beta(a_1 a_2^\dagger - a_1^\dagger a_2)\right] \tag{3.5}$$

and solve for β. To carry out the expansion it is useful to note that

$$\left[a_1 a_2^\dagger - a_1^\dagger a_2, a_1^\dagger a_1\right] = a_1^\dagger a_2 + a_1 a_2^\dagger \tag{3.6a}$$

$$\left[a_1 a_2^\dagger - a_1^\dagger a_2, a_2^\dagger a_2\right] = -(a_1^\dagger a_2 + a_1 a_2^\dagger) \tag{3.6b}$$

$$\left[a_1 a_2^\dagger - a_1^\dagger a_2, a_1^\dagger a_2 + a_1 a_2^\dagger\right] = 0. \tag{3.6c}$$

Using these commutators, to first order in β one gets

$$H \approx \sum_i \hbar\omega_i a_i^\dagger a_i + \frac{\hbar}{2}\left(\omega_q' + \sum_i \frac{2g_i^2}{\Delta_i}a_i^\dagger a_i\right)\sigma_z + \hbar(g_{12} + \beta\Delta_{12})(a_1^\dagger a_2 + a_1 a_2^\dagger), \tag{3.7}$$

where $\Delta_{12} = \omega_1 - \omega_2$.* If the qubit is in the dressed ground state, one can chose $\beta = -g_{12}/\Delta_{12}$ to arrive at a diagonal Hamiltonian that no longer displays coupling between the cavity modes.

---

* Mariantoni *et al.* derive a similar result through a different path [48].
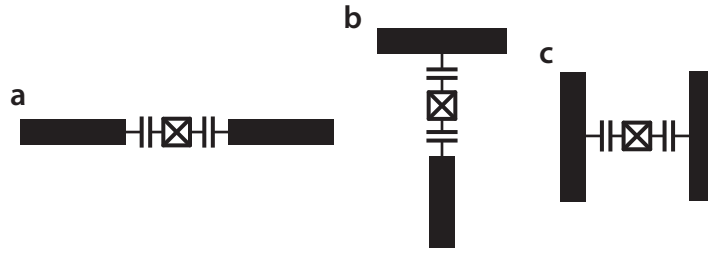
Figure 3.1: Coupling Geometries. a Inline geometry. **b** T geometry. **c** H geometry.

3.2   Separation of Cavity Modes

All implementations of the two-cavity Hamiltonian of (3.1) involve some kind of cavity-cavity coupling. This coupling takes on several different forms: there is direct capacitive coupling, indirect capacitive coupling through a qubit, and the *quantum* coupling that was just derived. The first two are described by the same physics, since both result in an effective capacitance between the cavities. This section examines the effects of this capacitive coupling. Relaxation from the quantum coupling is treated in section 3.3.

Possible geometries for capacitively coupling a qubit to the fundamental or first harmonic of two cavities include: inline, T, or H, as shown in figure 3.1. Only the inline geometry couples to the fundamental mode of each resonator, though, which has the advantage of minimally affecting the qubit decay by the presence of other modes (see section 4.1). All of these geometries involve some kind of direct coupling between the cavities themselves which hybridizes the eigenmodes such that the energy is shared between the two cavities. This is undesirable if one wants cavities with very different decay rates, since increased coupling pushes the frequencies and decay rates of the modes together (in the limit of infinite coupling the modes collapse onto a single mode). Consequently, it is useful to compute the coupling Q of each resonator due to its coupling to the other.

Figure 3.2 shows the most basic capacitive coupling between two cavities without a qubit. For each cavity, one can calculate Q_{couple} given capacitive coupling to 50 Ω and to the other
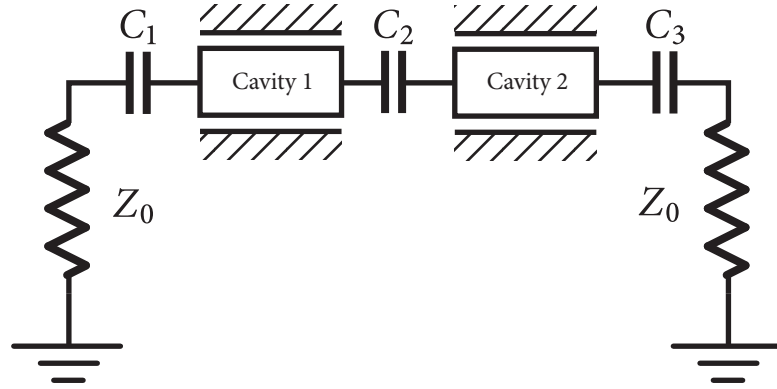
Figure 3.2: Schematic of coupled cavities with fundamental resonance frequencies ω_1 and ω_2, and characteristic impedance Z_0.
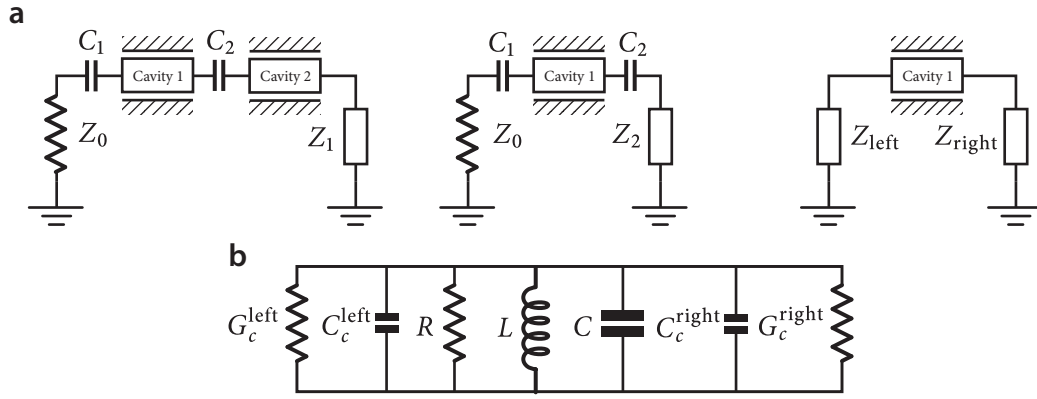
Figure 3.3: Reducing the coupled cavities circuit to calculate Q_{couple} for cavity 1.

cavity. We proceed by reducing the circuit, following figure 3.3(a):

$$Z_1 = \frac{1}{i\omega C_3} + Z_0 \tag{3.8a}$$

$$Z_2 = Z_0 \frac{Z_1 + iZ_0 \tan(\pi\omega/\omega_2)}{Z_0 + iZ_1 \tan(\pi\omega/\omega_2)} \tag{3.8b}$$

$$Z_{\text{right}} = \frac{1}{i\omega C_2} + Z_2 \tag{3.8c}$$

$$Z_{\text{left}} = \frac{1}{i\omega C_1} + Z_0. \tag{3.8d}$$

To find the Q of cavity 1, we approximate the $\lambda/2$ transmission line resonator as a parallel RLC circuit near its resonance frequency, where the damping is determined by the total

parallel conductance (see [49, chapter 6]):

$$G_c^{\text{left}} = \text{Re}(1/Z_{\text{left}}) = \frac{\omega^2 C_1^2 Z_0}{1 + \omega^2 C_1^2 Z_0^2} \approx (\omega C_1)^2 Z_0 \tag{3.9a}$$

$$i\omega C_c^{\text{left}} = \text{Im}(1/Z_{\text{left}}) = \frac{i\omega C_1}{1 + \omega^2 C_1^2 Z_0^2} \approx i\omega C_1 \tag{3.9b}$$

$$G_c^{\text{right}} = \text{Re}(1/Z_{\text{right}}) \tag{3.9c}$$

$$i\omega C_c^{\text{right}} = \text{Im}(1/Z_{\text{right}}) \tag{3.9d}$$

$$Q_{\text{couple}} = \frac{\omega_1 C}{G_c^{\text{left}} + G_c^{\text{right}}} = \frac{\pi}{2Z_0(G_c^{\text{left}} + G_c^{\text{right}})}. \tag{3.9e}$$

In the limit of weak coupling, $\omega C Z_0 \ll 1$, the imaginary part of the admittance, Y_{left}, is purely capacitive with $C_c^{\text{left}} = C_1$, pulling the resonance frequency of cavity 1 down. The coupled Q depends on the real part of the total admittance, which we can separate into contributions from the coupling at the input port (C_1) and from the other cavity. Thus, Q_{couple} splits into two terms

$$1/Q_{\text{couple}} = 1/Q_{\text{left}} + 1/Q_{\text{right}}, \tag{3.10}$$

$$Q_{\text{left}} = \frac{\pi}{2Z_0 G_c^{\text{left}}} \approx \frac{\pi}{2(\omega_1 C_1 Z_0)^2}, \tag{3.11}$$

$$Q_{\text{right}} = \frac{\pi}{2Z_0 G_c^{\text{right}}}. \tag{3.12}$$

The full expression for G_c^{right} is not particularly illuminating, but I can simplify the expression for the special case $\omega_1/\omega_2 = n/2$ (for n odd). Then,

$$Z_{\text{right}} = \frac{1}{i\omega C_3} + \frac{Z_0^2}{Z_1}, \tag{3.13a}$$

$$G_c^{\text{right}} = \frac{1}{Z_0} \frac{(\omega^2 C_2 C_3 Z_0^2)^2}{(1 - \omega^2 C_2 C_3 Z_0^2)^2 + (\omega C_3 Z_0)^2} \approx (\omega^2 C_2 C_3)^2 Z_0^3. \tag{3.13b}$$

In the last line I have used the fact that typical coupling capacitors are 1–100 fF, and cavity frequencies are on the order of 5 GHz. This gives $\omega C Z_0 \sim 10^{-1}$–$10^{-3}$, so I can drop such terms in the denominator of (3.13).

Alternatively, away from an $n\lambda/4$ resonance one can approximate cavity 2 as a parallel LC oscillator with $L = 2Z_0/(\pi\omega_2)$ and $C = \pi/(2\omega_2 Z_0)$. Keeping terms of the same order as
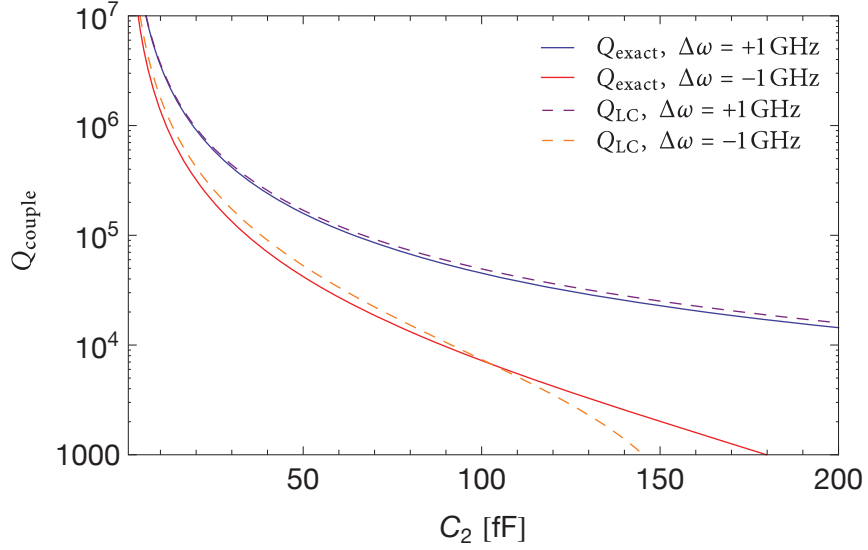
Figure 3.4: Coupled Q vs Inter-cavity Coupling. For this calculation, $\omega_1/2\pi = 5\,\text{GHz}$ and $C_3 = 20\,\text{fF}$. When $\Delta\omega = \omega_1 - \omega_2$ is $+1\,\text{GHz}$, this corresponds to $Q_{\text{ex}} \sim 2,500$ for cavity 2. When $\Delta\omega = -1\,\text{GHz}$, cavity 2 has $Q_{\text{ex}} \sim 1,100$. It is apparent that the LC oscillator approximation of (3.14) works fairly well as long as $C_2 \lesssim 5C_3$.

above, one has

$$G_c^{\text{right}} \approx \frac{(\omega^2 C_2 C_3)^2 Z_0^3}{\pi Z_0 (C_2 + C_3)\frac{\omega^2 - \omega_2^2}{\omega^2} + \frac{\pi^2}{4}\left(\frac{\omega^2 - \omega_2^2}{\omega\omega_2}\right)^2}. \tag{3.14}$$

Roughly speaking, this is the product of the conductances one would calculate for C_2 and C_3 directly coupled to $50\,\Omega$, so it is clear that the presence of a cavity in between C_2 and C_3 does provide some protection.

A comparison of the exact and approximate expressions for Q_{couple} is shown in figure 3.4. From this one can see that building cavities with vastly different Qs requires that the coupling between the cavities remains small. For instance, one might like to have one cavity with a Q which is limited by internal (dielectric) losses. For CPW resonators on silicon or sapphire, present techniques can achieve $Q_{\text{int}} \sim 10^5$–10^6. If this same cavity is coupled to a $Q \sim 1,000$ cavity, the coupling between the cavities should be less than $C_2 \lesssim 20\,\text{fF}$. In section 3.4, I explore the consequences of this for qubit design.

3.3 Qubit-mediated cavity relaxation

In addition to the classical relaxation channel of coupled cavities which was just derived, coupling two cavities to the same qubit also introduces a quantum relaxation channel through the hybridization of the cavity-qubit-cavity eigenmodes. Reference [27] calculates the eigenmodes of a single cavity-qubit system in the dispersive regime*. Writing the dispersive eigenstates in the uncoupled basis,

$$\overline{|g,1\rangle} \simeq -(g/\Delta)\,|e,0\rangle + |g,1\rangle, \tag{3.15a}$$

$$\overline{|e,0\rangle} \simeq |e,0\rangle + (g/\Delta)\,|g,1\rangle. \tag{3.15b}$$

This gives corresponding relaxation rates of

$$\kappa' = \kappa_\gamma + \kappa = (g/\Delta)^2\gamma + \kappa, \tag{3.16a}$$

$$\gamma' = \gamma_\kappa + \gamma = (g/\Delta)^2\kappa + \gamma, \tag{3.16b}$$

where γ and κ are the bare relaxation rates of the qubit and cavity, respectively. The κ_γ and γ_κ terms of (3.16) describe the single-mode Purcell effect. For two cavities, we apply the transformations U and U' from equations (3.2) and (3.5) to the uncoupled eigenstates. To order g_i/Δ_i and g_{12}/Δ_{12} this gives

$$\overline{|0,0,e\rangle} = U'U\,|0,0,e\rangle \simeq |0,0,e\rangle - (g_1/\Delta_1)\,|1,0,g\rangle - (g_2/\Delta_2)\,|0,1,g\rangle, \tag{3.17a}$$

$$\overline{|1,0,g\rangle} \simeq |1,0,g\rangle - (g_{12}/\Delta_{12})\,|0,1,g\rangle + (g_1/\Delta_1)\,|0,0,e\rangle, \tag{3.17b}$$

$$\overline{|0,1,g\rangle} \simeq |0,1,g\rangle + (g_{12}/\Delta_{12})\,|1,0,g\rangle + (g_2/\Delta_2)\,|0,0,e\rangle. \tag{3.17c}$$

Consequently, the approximate qubit and cavity relaxation rates are

$$\gamma' = \gamma + (g_1/\Delta_1)^2\kappa_1 + (g_2/\Delta_2)^2\kappa_2, \tag{3.18a}$$

$$\kappa_1' = \kappa_1 + (g_1/\Delta_1)^2\gamma + (g_{12}/\Delta_{12})^2\kappa_2, \tag{3.18b}$$

$$\kappa_2' = \kappa_2 + (g_2/\Delta_2)^2\gamma + (g_{12}/\Delta_{12})^2\kappa_1. \tag{3.18c}$$

When a qubit is coupled to two cavities, there is a Purcell effect due to each cavity which increases the qubit decay rate. Each cavity also inherits a complementary relaxation due to its

---

* This assumes that the qubit couples only to a single mode of the cavity. The effects of higher modes will be studied in section 4.1.
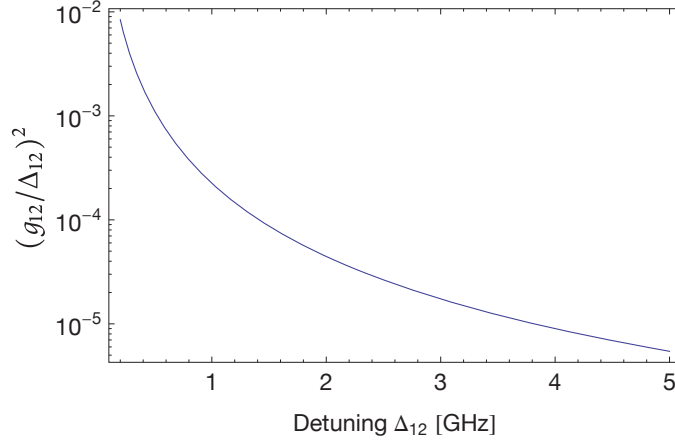
Figure 3.5: Qubit-mediated cavity-cavity hybridization. The line shows the hybridization fraction $(g_{12}/\Delta_{12})^2$ when $g_1/2\pi = g_2/2\pi = 100\,\mathrm{MHz}$, and $\Delta_1/2\pi = 1\,\mathrm{GHz}$ ($\Delta_2 = \Delta_1 + \Delta_{12}$). The cavity mode hybridization limits the ratio of κ's for the cavities. This is a fairly weak effect, even for the moderately large qubit coupling strengths used here. However, as cavity technology improves to the point where one can make $Q \gtrsim 10^7$, this effect will become a limitation.

hybridization with the qubit. Besides these effects which one might expect from single-cavity cQED, one also sees that the qubit-mediated cavity coupling (derived in section 3.1) leads to a hybridization of the cavity modes. Consequently, each cavity inherits some relaxation from the other cavity.

Figure 3.5 shows this hybridization vs. the detuning between the cavities for typical qubit coupling rates and detunings. The ratio $(g_{12}/\Delta_{12})^2$ is fairly small for the plotted set of parameters. However, one can already see that this limits the ratio Q_1/Q_2. For example, for a cavity-cavity detuning $\Delta_{12}/2\pi = 1\,\mathrm{GHz}$ the ratio of Q's cannot be more than 10^3–10^4 for the parameters used in the figure. The particular number strongly depends on the value of g_{12}, and thus on g_1 and g_2 through (3.4).

The qubit-mediated cavity-cavity coupling can also be used as a resource. A direct capacitive coupling between the cavities would enter (3.3) as another term proportional to $a_1^\dagger a_2 + a_1 a_2^\dagger$. When the qubit is in the ground state ($\sigma_z = -1$), these cavity coupling terms have opposite signs. Consequently, one could adjust qubit detunings to make g_{12} cancel the direct capacitive coupling, effectively switching it off. Mariantoni *et al.* have proposed using this mechanism to create a quantum mechanical switch [48].
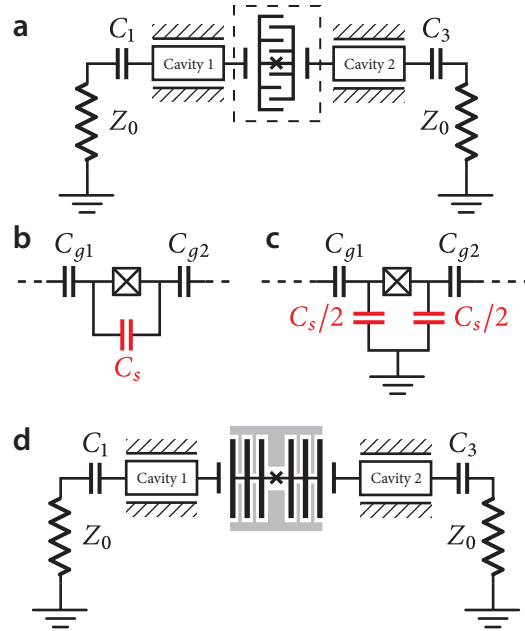
Figure 3.6: Capacitance network for a transmon qubit coupled to two cavities. a The most straightforward way to couple a single transmon qubit to two cavities. **b** The associated capacitance network of components in the dashed box of (**a**). **c** An alternative to the standard design replaces the large shunt capacitance between the islands of the transmon with capacitance to ground. **d** Possible layout of the "sarantapede" transmon. Pictures of real devices are shown in figure 5.1 and figure 8.1.

3.4   The Sarantapede Qubit

Considerations of cavity coupling has practical import to the design of a qubit which can couple to two cavities. In particular, placing a standard transmon between the cavities (see figure 3.6(a)) results in a capacitance network like figure 3.6(b). For a typical transmon design, the shunt capacitor C_s is quite large, in the range of 25–50 fF. The series combination of this shunt capacitance with the coupling capacitors, C_{g1} and C_{g2}, results in a substantial coupling capacitance C_2 between the cavities. As shown in section 3.2, this level of coupling is in the region where the direct capacitance between the cavities begins to limit the Qs. One can ameliorate this problem to a certain extent by modifying the capacitance network of the transmon.

One possibility is to shrink the size of the shunt capacitor C_s. However, in order for a Cooper pair box qubit to be in the transmon regime, it must have a large E_J/E_C ratio [15]. When C_s gets smaller, it raises the charging energy, E_C. Consequently, to maintain a constant
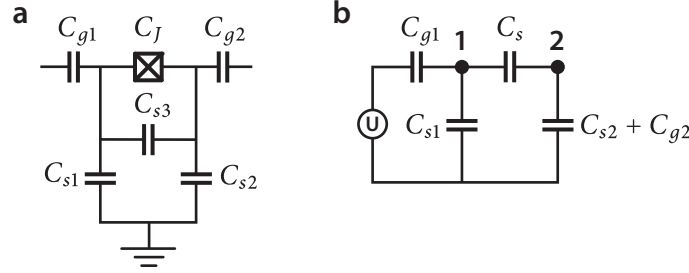
Figure 3.7: Relevant capacitance network for a sarantapede-style transmon qubit. a The local capacitance network for a "sarantapede" qubit. **b** Reduced capacitance network for finding the sarantapede coupling strength. $C_s = C_J + C_{s3}$.

E_J/E_C ratio would require a corresponding increase in E_J. However, there is a simultaneous design constraint that the device should have a ω_{ge} transition frequency which is comparable to the 5–10 GHz cavities which we use. Since $\hbar\omega_{ge} \approx \sqrt{8E_JE_C} - E_C$, increasing both E_J and E_C would push ω_{ge} above the typical frequencies of the cavities.

Another possibility is to replace the large shunt capacitor between the islands of the transmon with capacitors to ground (see figure 3.6(c)). Since each island now has its own shunt capacitor, each capacitor can be half the size and still give the same charging energy for the device. When computing the effective cavity coupling through this network, however, these capacitors add in series. So, the resulting inter-cavity coupling is reduced by a factor of 4. This is not an enormous gain, but it does offer some protection.

One possible 2D layout of this kind of transmon has interdigitated capacitors between the transmon islands and the surrounding ground planes, as shown in figure 3.6(d). An initial iteration of this insect-like qubit design had 40 "legs"; consequently, Luigi Frunzio nicknamed it the "sarantapede" because *saranta* is the Greek word for *forty*.

To find the charging energy, $E_C = e^2/2C_\Sigma$, for a sarantapede qubit coupled to two cavities, one should consider a capacitance network like the one shown in figure 3.7(a). This differs somewhat from the usual network that is considered for a transmon because, in this case, we can neglect the capacitances between the left island and right cavity, and *vice versa*[*]. Then the C_{gi} and C_{si} on each side add in parallel with the direct shunt capacitors C_J and C_{s3}, resulting

---

[*] A Maxwell simulation shows that other capacitances are of the order 10–100 fF, while the capacitances we neglect are much less than 1 fF.

in

$$C_\Sigma = C_J + C_{s3} + \frac{(C_{g1} + C_{s1})(C_{g2} + C_{s2})}{C_{g1} + C_{s1} + C_{g2} + C_{s2}}. \tag{3.19}$$

To find the coupling rates g_i of (3.1), we calculate the voltage drop $\beta_1 = (V_1 - V_2)/U$ across islands 1 and 2 of the qubit when a voltage U is applied to one of the cavities, as shown in figure 3.7(b) for cavity 1. The large capacitance to ground of the transmission line resonator of cavity 2 causes C_{g2} to look like an effective coupling to ground; consequently, it adds in parallel with C_{s2}. Following the procedure in [15, section 5.2] and [50, appendix 1] we consider only the reduced capacitance matrix of the qubit islands 1 and 2,

$$C = \begin{pmatrix} C_{g1} + C_{s1} + C_s & -C_s \\ -C_s & C_{g2} + C_{s2} + C_s \end{pmatrix}, \tag{3.20}$$

where $C_s = C_J + C_{s3}$. To find β_1, we apply a polarization charge to island 1: $\tilde{Q} = (C_{g1}U, 0)$. Then, solving for $V = (V_1, V_2) = C^{-1}\tilde{Q}$ and β_1 we get

$$V = \frac{C_{g1}U}{(C_{g1} + C_{s1})(C_{g2} + C_{s2}) + C_s(C_{g1} + C_{s1} + C_{g2} + C_{s2})} \begin{pmatrix} C_{g2} + C_{s2} + C_s \\ C_s \end{pmatrix}, \tag{3.21a}$$

$$\beta_1 = \frac{C_{g1}(C_{g2} + C_{s2})}{(C_{g1} + C_{s1})(C_{g2} + C_{s2}) + C_s(C_{g1} + C_{s1} + C_{g2} + C_{s2})}. \tag{3.21b}$$

Converting this to a coupling rate requires multiplying by a transmon matrix element, which is best done numerically. The expression is (see [15, section 4.3.4])

$$g_1 = 2\beta_1\omega_1\sqrt{\frac{2Z_0}{R_k}}\langle g|\hat{n}|e\rangle, \tag{3.22}$$

where $R_k \approx 25.8\,\text{k}\Omega$ is the resistance quantum and the matrix element is calculated for $\omega_{ge} = \omega_1$.

3.5   Self- and Cross-Kerr Effects

Thus far I have only looked at the dispersive two cavity Hamiltonian to second order in (g_i/Δ_i). However, since the cavity-qubit coupling causes the cavity and qubit states to hybridize, one might also expect that the cavities would inherit some anharmonicity from the qubit (*i.e.* the

cavities should become more qubit-like). To see this effect, one must expand the dispersive Hamiltonian to fourth order.

Before going there, however, it is simpler to start by returning to the usual Jaynes-Cummings case of a single cavity and qubit. Here, the dispersive Hamiltonian can be diagonalized exactly [51], giving

$$H/\hbar = \omega_r a^\dagger a + \frac{\omega_q}{2}\sigma_z - \frac{\Delta}{2}\left(1 - \sqrt{1 + 4\lambda^2 N_q}\right)\sigma_z, \tag{3.23}$$

where $\lambda = g/\Delta$ and $N_q = a^\dagger a + |e\rangle\langle e|$ is the number of excitations in the system. Expanding to third order in λ, one finds

$$H/\hbar \simeq (\omega_r + \zeta)a^\dagger a + \frac{1}{2}\left[\omega_q + 2\chi(a^\dagger a + 1/2)\right]\sigma_z + \zeta(a^\dagger a)^2\sigma_z, \tag{3.24}$$

where $\chi = g^2(1 - \lambda)/\Delta$ and $\zeta = -g^4/\Delta^3$. This last term causes a photon-number dependent shift of the cavity frequency, known as the self-Kerr effect.

One should expect to find an analogous effect for two cavities, but where the cavity frequencies also depend on the photon number in the *other* cavity. Unfortunately, the two-cavity Hamiltonian of (3.1) does not share the simple 2×2 block diagonal form with the single-cavity case. This prevents one from finding an exact diagonalization which can be expanded. Instead, I will compute the corrections to system energy to 4th order, where I start with the initial Hamiltonian and perturbation $H = H_0 + \lambda V$ where

$$H_0 = \sum_i \hbar\omega_i a_i^\dagger a_i + \frac{\hbar}{2}\omega_q\sigma_z, \tag{3.25a}$$

$$\lambda V = \sum_i \hbar g_i(a_i^\dagger\sigma^- + a_i\sigma^+). \tag{3.25b}$$

If your education was like mine, then you probably only learned how to do perturbation theory to second order in your quantum class. Fortunately, you can look up the expressions for higher orders in many references, including Wikipedia. Defining the perturbed energies as $E_n = E_n^{(0)} + \lambda E_n^{(1)} + \lambda^2 E_n^{(2)} + \cdots$, one has

$$E_n^{(3)} = \sum_{i,j\neq n} \frac{V_{ni}V_{ij}V_{jn}}{E_{ni}E_{nj}} - \sum_{i\neq n} \frac{V_{nn}V_{in}}{E_{ni}^2}, \tag{3.26a}$$

$$E_n^{(4)} = \sum_{i,j,k \neq n} \left[\frac{V_{ni}V_{ij}V_{jk}V_{kn}}{E_{ni}E_{nj}E_{nk}} - \frac{|V_{ni}|^2 |V_{nk}|^2}{E_{ni}^2 E_{nk}} - V_{nn}\frac{V_{ni}V_{ij}V_{jn}}{E_{ni}E_{nj}^2} \right.$$
$$\left. - V_{nn}\frac{V_{ni}V_{ik}V_{kn}}{E_{nk}E_{ni}^2} + V_{nn}^2 \frac{|V_{nk}|^2}{E_{nk}^3} \right], \quad (3.26b)$$

where $V_{ij} = \langle i | V | j \rangle$ and $E_{ij} = E_i^{(0)} - E_j^{(0)}$. Computing the matrix elements is relatively straightforward because V only connects states which differ by one excitation, *i.e.*

$$\langle n_1, n_2, g | V | j_1, j_2, e \rangle = \delta_{j_1+1}^{n_1}\delta_{j_2}^{n_2}\tilde{g}_1\sqrt{n_1 + 1} + \delta_{j_1}^{n_1}\delta_{j_2+1}^{n_2}\tilde{g}_2\sqrt{n_2 + 1}, \quad (3.27)$$

$$\langle n_1, n_2, e | V | j_1, j_2, g \rangle = \delta_{j_1-1}^{n_1}\delta_{j_2}^{n_2}\tilde{g}_1\sqrt{n_1} + \delta_{j_1}^{n_1}\delta_{j_2-1}^{n_2}\tilde{g}_2\sqrt{n_2}, \quad (3.28)$$

where $\tilde{g}_i = g_i/\lambda$. Consequently, all the terms in (3.26) proportional to V_{nn} vanish. In addition, the first term of (3.26a) involves a product of an odd number of matrix elements, where the initial and final states are the same. However, an odd number of applications of V cannot return to the initial state, so the entire 3rd order correction term, $E_n^{(3)}$, is zero.

With respect to the 4th order corrections, the cross-Kerr terms only involve combinations of intermediate states in the sum of (3.26) which change both cavity photon numbers. Keeping only such terms, one obtains

$$\lambda^4 E_{n_1,n_2,g}^{(4)} = g_1^2 g_2^2 \left(\frac{n_1(n_2 + 1)}{\Delta_1^2 \Delta_{12}} - \frac{(n_1 + 1)n_2}{\Delta_2^2 \Delta_{12}} + \frac{n_1 n_2}{\Delta_1^2 \Delta_2} + \frac{n_1 n_2}{\Delta_1 \Delta_2^2} \right). \quad (3.29)$$

Collecting terms proportional to $n_1 n_2$, this simplifies to

$$\lambda^4 E_{n_1,n_2,g}^{(4)} \approx \frac{2 g_1^2 g_2^2 (\Delta_1 + \Delta_2)}{\Delta_1^2 \Delta_2^2} n_1 n_2, \quad (3.30)$$

which gives the cross-Kerr energy shift from population in each cavity. One can easily check that the above procedure reproduces the self-Kerr correction of (3.24) by computing the terms proportional to n_1^2 or n_2^2.

The Electromagnetic Environment and Fast Qubit Control

C ONTROL and readout of superconducting qubits necessitates interactions between qubits and other circuits. This coupling, which provides control and measurement, also introduces channels for energy relaxation and dephasing. Consequently, it is crucial to understand how these processes work in order to design circuits that do not impair the performance of qubits. This chapter presents a general formalism for calculating energy relaxation from coupling to an electromagnetic environment modeled as a classical admittance. This theory is not new, but it is presented here in the context of circuit QED with transmons, where it allows us to generalize the Purcell effect to a model that includes all higher harmonics of a cavity. Using this theory, a new circuit element, the 'Purcell filter', is proposed to mitigate the Purcell effect.

We then turn to the design of a fast flux bias line (FBL) to control qubit frequencies on nanosecond timescales. Issues of relaxation and dephasing from this control knob are examined. With appropriate design choices, this FBL minimally affects qubit lifetimes. Even with these issues resolved, control systems sometimes respond in unexpected ways, so the final sections apply deconvolution methods to correct the output from the FBL.

4.1   Multi-mode Purcell Effect

In section 3.3 I considered the relaxation of a qubit coupled to one or two single-mode cavities. The cavity does not just support a single electromagnetic mode, though, but also all higher harmonics of the fundamental mode. A single-mode description gives a fairly good understanding of the observed relaxation in cQED experiments at small qubit-cavity detunings, where one mode of the cavity(ies) dominates. At larger detunings, however, the presence of higher harmonics of the transmission line resonator matters. In fact, one might expect that these higher harmonics play a significant role because the coupling, g_{eg}^n, to the n^{th} mode increases with the mode number, $g_{eg}^n = g_{eg}^0 \sqrt{n+1}$. Furthermore, the input and output capacitors of the cavity act as frequency-dependent mirrors, such that the decay rate of the n^{th} harmonic, $\kappa_n = (n+1)^2 \kappa$, is faster than the fundamental.

There are ongoing efforts to provide a full quantum mechanical description of the multi-mode Jaynes-Cummings ladder, which are, at present, incomplete because the calculations of quantities like the coupling rate between two qubits lead to divergent sums.* Presumably this can be fixed by truncating the sums to a finite number of modes, but we do not yet have a good understanding of an appropriate choice for this cut-off. Nonetheless, there is a quite elegant formalism which treats the dissipation of macroscopic degrees of freedom using the classical admittance of the environment. As I will show experimentally in chapter 6, this provides a very good model for understanding transmon relaxation over a wide range of parameters.

The formalism which I consider here was developed by Esteve *et al.* following earlier work by Leggett[52, 53]. In their work, Esteve *et al.* consider an equation of motion for a single degree of freedom X in a potential $V(X)$ and subject to friction through coupling to many degrees of freedom in a energy reservoir. The general equation of motion for X is

$$\ddot{X} + \hat{K}\{X\} = -\frac{\partial}{\partial X} V(X) \tag{4.1}$$

where $\hat{K}$ is some linear operator which is given by the properties of the energy reservoir. They treat this system quantum mechanically by representing $\hat{K}\{X\}$ as an infinite set of series LC oscillators chosen to match the spectral density of the environment[26], as shown in

---

* For instance, if you use non-degenerate perturbation theory, in principle there are an infinite number of non-zero matrix elements which contribute to qubit-qubit coupling via virtual exchange with every cavity mode. So, the resulting sum is $\sum_n g_1^n g_2^n (\Delta_1^n + \Delta_2^n)/\Delta_1^n \Delta_2^n \approx \sum_n (-1)^n (\sqrt{n}g_1)(\sqrt{n}g_2) n(\Delta_1 + \Delta_2)/n^2 \Delta_1 \Delta_2$, which is proportional to $\sum_n (-1)^n$.
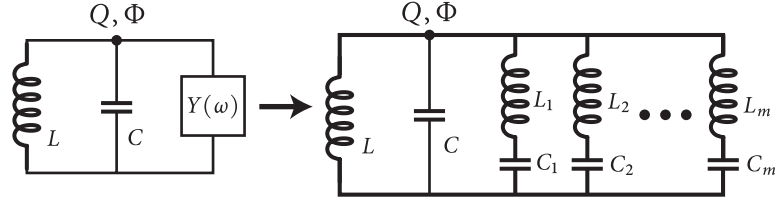
Figure 4.1: Quantum LC oscillator with an environment $Y(\omega)$. Circuit diagram for a harmonic oscillator with node coordinates Q and Φ and an admittance $Y(\omega)$. The right hand side shows the equivalent Caldeira-Leggett representation of $Y(\omega)$ in terms of an infinite set of LC oscillators[26].

figure 4.1. Having a Hamiltonian description for X, they proceed to calculate the energy shifts of the eigenstates due to coupling to the reservoir. The imaginary part of the energy shift is identified as $-1/(2\tau_n)$, where τ_n is the lifetime of level n. For a harmonic oscillator, $V(X) = 1/2\omega_0^2 X^2$, they calculate

$$1/\tau_n = -n\,\mathrm{Im}\big[K^*(\omega_0)\big]/\omega_0. \tag{4.2}$$

This result (4.2) is particularly useful for calculating the relaxation of a transmon, because the transmon is nearly harmonic (see section 2.1). In fact, if we expand the cosine term in the usual CPB Hamiltonian (2.1) to second order, we have

$$H \approx 4E_C \left(\frac{\partial}{\partial\phi}\right)^2 + \frac{1}{2}E_J\phi^2 \tag{4.3}$$

where $2e\phi/\hbar$ is the phase difference across the Josephson junction. This is exactly the Hamiltonian of a quantum LC oscillator provided one makes the correspondence $C \leftrightarrow 1/(8E_C)$ and $L \leftrightarrow 1/E_J$. When this LC oscillator is coupled to an environment with admittance $Y(\omega)$ (see figure 4.1), the resulting circuit has an equation of motion of the form (4.1)[52] with $X = C(2e/\hbar)^2\phi$ and $K(\omega) = i\omega Y(\omega)/C$.

Consequently, the decay time of the n^{th} harmonic oscillator level is

$$\tau_n = \frac{\omega_0}{n\,\mathrm{Im}\big[i\omega_0 Y(\omega_0)/C\big]} = \frac{C}{n\,\mathrm{Re}\big[Y(\omega_0)\big]}. \tag{4.4}$$

Thus, one arrives at the remarkably simple result that the decay of the first excited state is given by an RC time constant where $1/R$ is the real part of the classical admittance of the coupled environment.

To illustrate a few additional issues with typical cQED circuits, I will explicitly compute
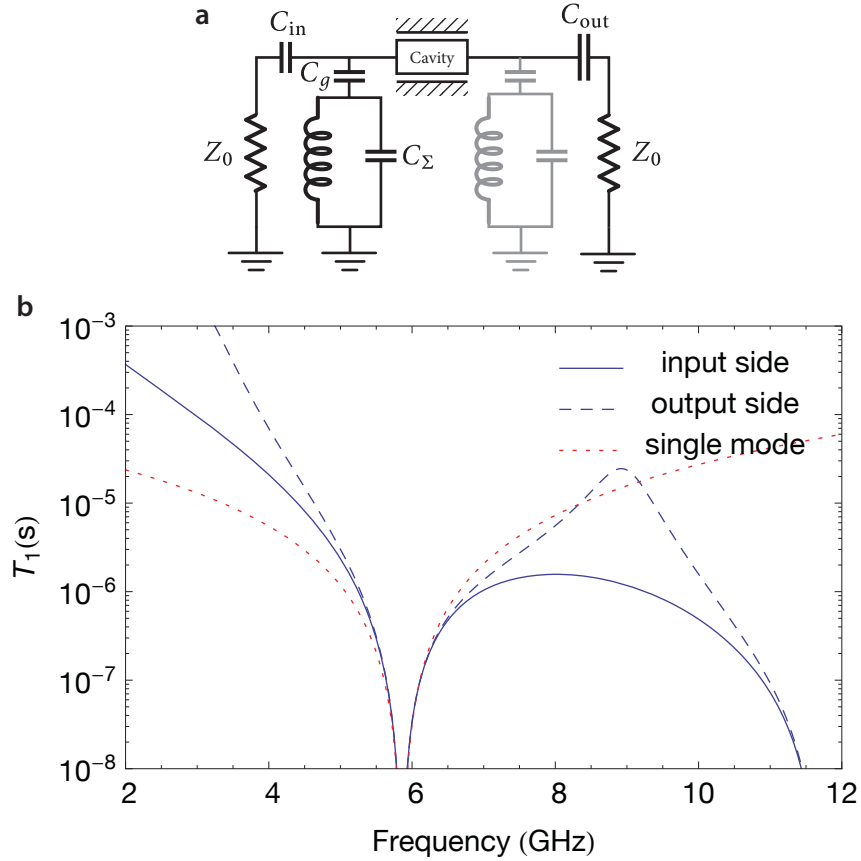
Figure 4.2: Single cavity multi-mode Purcell effect. a Circuit model for a qubit coupled to a single cavity with asymmetric input and output couplers, C_{in} and C_{in}. To couple to a voltage antinode of the fundamental resonance, the qubit is placed near the input (black) or output (grey). **b** The calculated T_1 for the qubit at the input side (solid) or output side (dashed). The specifics of how the transmission line modifies the admittance of the input and output circuits causes different relaxation for these positions. For comparison, the calculated T_1 from the simple single-mode theory is also plotted (dotted). The single-mode theory overestimates the relaxation when the qubit is below the first cavity mode, but underestimates the relaxation when the qubit is above. Parameters used for this calculation are: $C_{in} = 6\,\mathrm{fF}$ ($\kappa_{in}/2\pi = 0.5\,\mathrm{MHz}$), $C_{out} = 27\,\mathrm{fF}$ ($\kappa_{in}/2\pi = 10\,\mathrm{MHz}$), $C_g = 7.7\,\mathrm{fF}$ ($g/2\pi = 100\,\mathrm{MHz}$), $C_\Sigma = 65\,\mathrm{fF}$ ($E_C/2\pi = 300\,\mathrm{MHz}$), and $\omega_r/2\pi = 6.0\,\mathrm{GHz}$.

the decay rate for several examples. In each case, I want to reduce the total circuit to one of the form shown in figure 4.1. In practice, carrying out these calculations can be a bit tedious. One can do it by hand (not recommended) or with a computer algebra system like Mathematica. In Rob Schoelkopf's lab, we frequently use packages like Microwave Office because it allows one to drop circuit elements onto a diagram and calculate admittances with sliders for seeing real-time updates to relaxation as parameter values are adjusted. This is incredibly useful, especially when designing circuits like the Purcell filter which is discussed in section 4.1.2 and section 6.2.

Figure 4.2(a) shows a circuit diagram for a typical single-cavity cQED experiment where the cavity is probed in transmission. In a transmission configuration, it is often advantageous to use asymmetric coupling capacitors so that photons in the cavity will preferentially leak out toward the amplification chain on the output side. However, the asymmetry of the coupling capacitors causes qubits to relax at different rates depending on where they are placed in the cavity. Transmon qubits are dipole coupled, so for maximal coupling strength, they are placed near voltage antinodes. For the fundamental mode of a transmission line cavity, this implies that the transmon is placed near either coupling capacitor (the black and grey harmonic oscillators of figure 4.2(a) illustrate this placement). The calculated relaxation for these two positions is shown in figure 4.2(b). Notice that when the qubit is at the "input side" (solid line), and thus farther from the large output capacitor, it actually has greater Purcell-induced relaxation. The transmission line transforms the admittance of the output circuit in a way which is unfavorable for dissipation. For both qubit positions, though, the calculated Purcell T_1 has a strong asymmetry for qubit frequencies above and below the cavity frequency. One should expect this feature because as the qubit frequency increases, its effective coupling to higher modes increases. Surprisingly, the higher modes actually reduce dissipation compared to the single-mode theory (dotted line) when the qubit frequency is below the fundamental mode.

Another relevant example in the context of this thesis is the multi-mode Purcell effect for two-cavity circuits, as shown in figure 4.3 for a high-Q cavity at 5 GHz and a low-Q cavity at 7 GHz. Having seen the equivalent effect in a single-cavity circuit, the T_1 vs. frequency profile is not too surprising—there are simply more locations where the T_1 drops sharply. One should be a little careful in interpreting the depths of these sharp drops. An important requirement for using the admittance formalism of Esteve *et al.* is that the degree of freedom should only be weakly coupled to any particular mode of the environment. This is certainly not the case when the qubit is resonantly coupled to a cavity mode. In fact, in the resonant
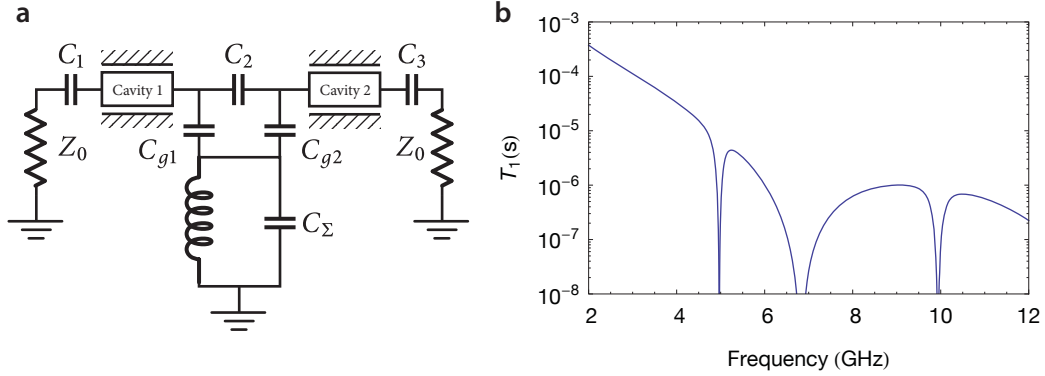
Figure 4.3: Two cavity multi-mode Purcell effect. a Circuit model for a qubit coupled to two cavities. The qubit is coupled at the ends of the cavities opposite to their respective inputs to protect from dissipation. **b** The calculated multi-mode Purcell T_1 for the circuit shown in **a**. Parameters are similar to those of the experiment of chapter 8. They are: C_1 = 2.5 fF ($\kappa_1/2\pi$ = 0.05 MHz), C_2 = 0 fF (no direct coupling), C_3 = 30 fF ($\kappa_3/2\pi$ = 20 MHz), C_{g1} = 6.2 fF ($g_1/2\pi$ = 100 MHz), C_{g2} = 10.2 fF ($g_2/2\pi$ = 100 MHz), C_Σ = 65 fF ($E_C/2\pi$ = 300 MHz), $\omega_1/2\pi$ = 5.0 GHz, and $\omega_2/2\pi$ = 7.0 GHz.

Jaynes-Cummings model the eigenstates are equal superpositions of cavity and qubit with a total decay rate given by the average of the individual qubit and cavity decay rates, *i.e.* $1/T_1 = (\gamma + \kappa)/2$. Throughout this discussion, I only treat cavity-induced decay, which is equivalent to assuming that $\gamma = 0$. Consequently, in resonance, the decay rate should be $\kappa/2$. In figure 4.3(b), I take $\kappa_1/2\pi$ = .05 MHz which gives $T_1 \sim 6\,\mu$s. Thus, the dips at 5 GHz and 10 GHz should have minima of $\sim 10^{-6}$ s. A more complete model would take this into account.

With that caveat aside, in figure 4.3(b) I have calculated the relaxation for the case where the low-Q cavity is at a higher frequency than the high-Q cavity. This is a preferable arrangement for experiments which operate with the qubit frequency in-between the cavity frequencies, since the asymmetry of the relaxation means that the Purcell T_1 would be smaller if the cavity frequencies were swapped. As we have seen in section 2.2.4, however, there are situations where it may be more convenient to work at negative detunings from a high-Q cavity (e.g. the quasi-dispersive regime is simpler below a cavity for qubits with negative anharmonicity). Consequently, it would be useful to have a means to modify the admittance at some frequency or range of frequencies. This will be the topic of section 4.1.2 and section 6.2.

4.1.1   Balanced Transmons

Given that the admittance of the environment seen by a qubit can have a dramatic effect on its T_1, one should be very careful to eliminate coupling to chip or package modes which add in parallel to the admittance of the CPW cavity. In particular, the CPW transmission lines frequently used in RSL designs can also support a slotline mode between the ground planes. This mode is not interrupted by a capacitor, so it is directly coupled to ~ 50 Ω. Furthermore, the 'standard moose' transmon design* is asymmetric with respect to the center pin and ground planes of the CPW transmission line, so it couples fairly strongly to the slotline mode.

There are several ways to go about fixing this. One is to try to suppress the slotline mode by ensuring that the top and bottom metallizations of the CPW are tightly grounded together. In part, this is the purpose of the densely packed wire bonds around the periphery of the chip which are common practice in RSL. Unfortunately, these wire bonds are very far away on the scale of the CPW gap, so they are not very effective at suppressing the mode. Crossovers that connect at the edge of the CPW gap would work better, and are actually used in designs by the NIST and UCSB groups. However, fabrication of these crossovers requires a dielectric spacer material which typically has a detrimental effect on cavity and qubit T_1 from dielectric losses†.

Another possibility is to modify the transmon design such that it couples symmetrically to the top and bottom ground planes. Figure 4.4 shows the relevant geometry and capacitance network. The symmetric solution equates $C_{a1} = C_{a3}$ and $C_{b1} = C_{b3}$. This can be done geometrically, as is the case for the sarantapede design, where the two transmon islands have equal capacitance to each ground plane. This is not the only solution, though. In particular, since the capacitance matrix has a form similar to a Wheatstone bridge we can look for a solution that exploits the relationships between the arms of the diamond. One way to proceed is to apply a voltage between islands 1 and 3, and solve for the capacitance relation that ensures the voltage drop $V_a - V_b = 0$. Since 1 and 3 are held at fixed voltages, we consider only the reduced capacitance matrix, C, for the floating islands: A, B, and 2. That is,

$$C = \begin{pmatrix} C_{ab} + C_{a1} + C_{a2} + C_{a3} & -C_{ab} & -C_{a2} \\ -C_{ab} & C_{ab} + C_{b1} + C_{b2} + C_{b3} & -C_{b2} \\ -C_{a2} & -C_{b2} & C_{a2} + C_{b2} + C_{12} + C_{23} \end{pmatrix}. \quad (4.5)$$

---

* see zoology of transmon qubits in section 5.1.1.

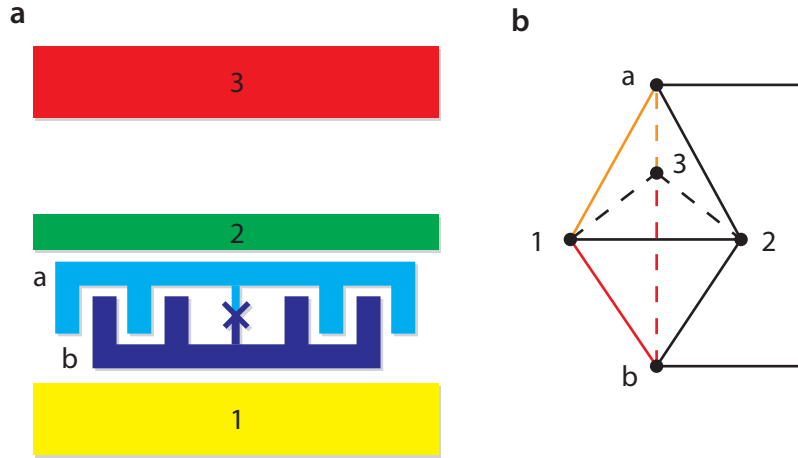† There has been some success in making air-gap crossovers which might eliminate this problem [54].

Figure 4.4: Capacitance network for transmon coupled to CPW. a In the standard geometry of coupling a transmon to a CPW transmission line, there are five pieces of metal, the two transmon islands (labeled A and B), the CPW center pin (2), and the two ground planes (1 and 3). By virtue of being closer to the bottom ground plane, A and B typically are more strongly coupled to 1. **b** The capacitance network for this set for this geometry represented as a diamond, where every edge represents a capacitor. 'Balancing' the transmon means that a voltage between 1 and 3 does not create a voltage drop between A and B. This can be done simply by symmetrically coupling A and B to 1 and 3, *i.e.* the edges shown in orange and red should be equal. There is a also a bridge solution, given in the text.

Applying a voltage, V_3, between 1 and 3 is equivalent to applying a polarization charge $\tilde{Q} = V_3(C_{a3}, C_{b3}, C_{23})$. The bridge solution is found by solving for $V = C^{-1}\tilde{Q}$ and setting $V_a = V_b$. Since the capacitances between the CPW elements, C_{12}, C_{23}, and C_{13}, are much larger than the capacitances to the transmon islands, we can replace $C_{12} = C_{23} = C'$ and $C_{13} = \epsilon C'$, where ϵ is a proportionality factor of order 1. Taking the limit $C' \to \infty$ gives

$$C_{a2}C_{b1} + C_{a3}C_{b2} + 2C_{a3}C_{b1} = C_{a2}C_{b3} + C_{a1}C_{b2} + 2C_{a1}C_{b3}. \tag{4.6}$$

Figure 4.5 shows two example realizations of 'balanced' transmon designs with minimal coupling to the slotline mode. Both designs approximately satisfy (4.6) by increasing C_{a1} and C_{b2} while decreasing C_{a2} and C_{b1}. As a side effect of the capacitance modifications, these designs tend to have a smaller g_{ge} because of a reduced voltage division ratio, $\beta \sim 0.08$. This results in $g_{ge} \sim 50$–100 MHz, as opposed to $g_{ge} \sim 150$–200 MHz for the 'typical' transmon design. Consequently, these designs have utility beyond reduced coupling to the slotline mode for experiments which require a smaller coupling rate. Unfortunately, devices employing
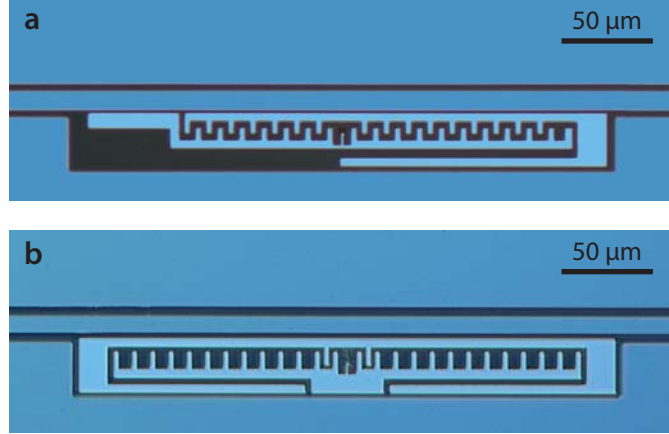
Figure 4.5: Balanced transmon designs. The bridge solution of (4.6) can be satisfied by increasing the capacitance of the upper island to the lower ground plane and the lower island to the center pin. **a** The 'crocodile' transmon used in cQED167 achieves this by wrapping around asymmetrically, whereas **b** the 'moose with helmet' transmon from cQED186 is mirror symmetric in one direction.

these designs have not shown a substantial change in T_1 compared to standard transmon designs, suggesting that the slotline coupling is not the dominant source of relaxation in current circuit QED experiments.

4.1.2 Purcell Filter

Equation (4.4) typically implies a trade-off between the cavity decay rate, κ, and T_1, because fast cavities present a fairly large Re[$Y(\omega)$] to the qubit. This constrains certain design goals, such as optimizing the readout fidelity, because one often wants T_1 to be as large as possible while simultaneously maximizing the information collection rate by increasing κ. The signal to noise ratio for the dispersive cQED read-out is given by a product of the information collection rate and the integration time divided by the number of noise photons of the amplifier. For a signal at the frequency of maximum distinguishability in phase and integrating for a time T_1, the SNR is [15]

$$\mathrm{SNR} = \sin^2(\theta_0)\frac{n\hbar\omega_r\kappa T_1}{k_b T_N} \tag{4.7a}$$

$$= \frac{4\chi^2 T_1}{\kappa}\frac{n\hbar\omega_r}{k_b T_N}, \tag{4.7b}$$

Figure 4.6: Purcell filter. a Circuit model for a filter to make $Y(\omega)$ purely reactive at the qubit frequency. A $\lambda/4$ open-circuited transmission line stub effectively shorts out the 50 Ω envi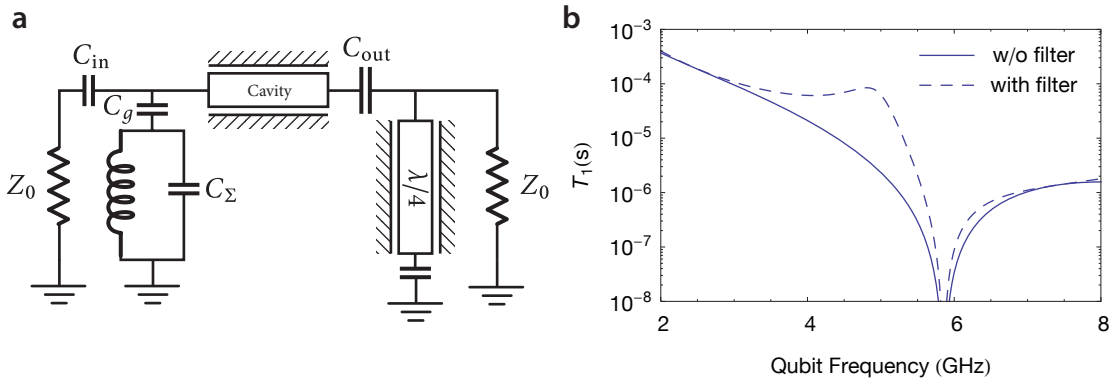ronment at the quarter-wave frequency. **b** Calculated T_1 using the admittance formalism of section 4.1 with and without the filter. Parameters are the same as those used in figure 4.2 with the addition of $v_p/(2\pi\lambda) = 5\,\text{GHz}$. The filter creates a large band around 5 GHz where the T_1 is significantly improved compared to the circuit without the filter.

where θ_0 is the phase difference between the signals when the qubit is in $|g\rangle$ and $|e\rangle$, and the last line follows in the small phase shift limit where $\sin(\theta_0)^2 \approx (2\chi/\kappa)^2$. Consequently, the SNR depends on the ratio of T_1 to κ. The Purcell effect links these two parameters, but since the qubit and cavity frequencies are typically detuned, one can look for a way to modify $\text{Re}[Y(\omega_q)]$ without changing $Y(\omega_r)$ by introducing additional circuit elements to the network of figure 4.2.

One method for achieving this is to place a $\lambda/4$ open-circuited transmission line stub* outside of the output capacitor, C_{out}, as shown in figure 4.6(a). A $\lambda/4$ transmission line acts as an impedance transformer, so by choosing $v_p/\lambda = \omega_q/2\pi$, the stub shorts out the 50 Ω environment at ω_q. The calculated T_1 for this circuit is shown in figure 4.6(b) for a filter centered around 5 GHz. The filter creates a region of qubit frequencies of approximately 10% bandwidth where the T_1 is significantly improved compared to the circuit without the filter. The realization and measurement of such a filter is presented in section 6.2.

4.2   Flux Bias Lines

Initial cQED experiments relied upon large bandwidth control of the gate charge on the CPB to do fast manipulations of the qubit frequency. The difficulty in controlling the low-frequency

---

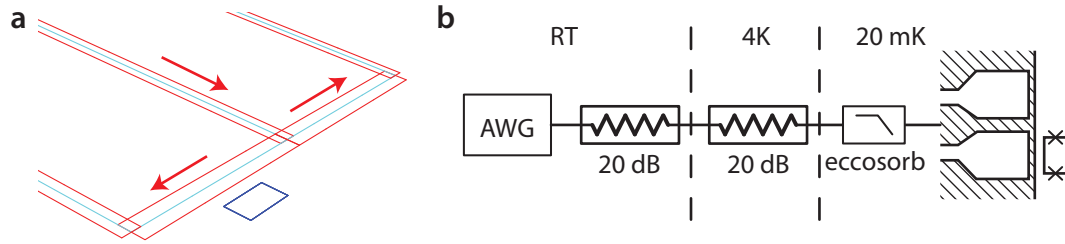* A $\lambda/2$ short-circuited stub would work as well, but it takes up more space on the chip.

Figure 4.7:   Flux bias line setup. a Geometry used for estimating the mutual inductance between the bias line and the SQUID loop, calculated using Biot-Savart and assuming uniform current sheets. Here, the SQUID loop is approximated as a 3 μm × 8 μm rectangle that is 1 μm away from the edge of a 2 μm × 20 μm bias line. **b** Schematic showing the connection between room temperature control electronics and the chip. The 20 dB attenuator at room temperature attenuates noise from the AWG and effectively transforms the voltage output into a current output. A 20 dB attenuator at 4 K serves to thermalize the line. At the base temperature of the dilution refrigerator, a series combination of a reflective 1 GHz low-pass filter and an eccosorb filter attenuate high frequency noise and thermalize the line.

charge environment was an impetus for the development of the transmon. However, the move to the transmon temporarily involved a loss of this control, since the first few levels of the transmon have very little sensitivity to gate charge. Instead, the transmon frequency is modified by changing the flux which penetrates the SQUID loop that connects the islands. In the first transmon experiments, this field was provided by a large superconducting bore magnet outside the vacuum can of the dilution refrigerator. This magnet has significant inductance (> 1 Henry), and the normal metal leads driving it have ~ 100 Ω resistance, leading to an L/R time in the millisecond range. Furthermore, the field that the bore magnet provides is global, so it simultaneously changes the frequencies of all qubits on a chip. Adding local on-chip flux bias lines can fix both of these problems, by allowing individual tuning of the qubit frequencies on nanosecond timescales.

The basic design approach pursued at Yale, shown in figure 4.7(b) uses a 50 Ω transmission line which is terminated in a short near the transmon SQUID loop. The line is placed off-center, such that one arm of the short couples more strongly to the SQUID loop. Maintaining 50 Ω in the bias line makes it compatible with commercial arbitrary waveform generators for producing fast pulses. The short termination allows one to bring the bias current very close to the qubit SQUID loop, at the price of a factor of two in the current since the current is split in two directions at the short. In addition, the short termination causes the entire chip (in fact, for our setup it is the entire frame of the fridge) to serve as the return path. This means

that the current travels through a rather uncontrolled path, which may be the cause of the anomalous time-domain response which will be discussed in section 4.3.

To have complete control of the qubit frequency, the flux bias line must have sufficient mutual inductance with the SQUID loop in order to thread one flux quantum without requiring an excessive amount of current on the line. The current limits are set by many factors, including current sourcing ability of room temperature electronics, the critical current of the narrowest section of the short, and the available cooling power at the dilution refrigerator stages where the line is thermalized with resistors. For the setup shown in figure 4.7(b), an RF voltage pulse of amplitude V produces a current

$$I = 2 \cdot \frac{1}{100} \cdot \frac{V}{50\,\Omega} = \frac{V}{2.5\,\mathrm{k\Omega}}, \tag{4.8}$$

where the factor of 2 comes from the short circuit termination and the factor of 1/100 is the *voltage* division ratio for two 20 dB attenuators. The maximum amplitude pulse one can produce with a Tektronix AWG5014 is 10 V, which by (4.8) gives a current of 4 mA. This implies that the flux bias line should, at minimum, have a mutual inductance of 250 Φ_0/A. With sufficiently quiet electronics, one could remove the attenuator at room-temperature to get an extra factor of 10 in current.

The primary design variables for a flux bias line are the distance between the short and loop, and the area of the loop. Typically, the SQUID loop areas have been kept fairly small, on the order of 3 μm × 8 μm, in order to limit sensitivity to flux noise and crosstalk. Rather counterintuitively, measurements of flux noise in SQUID loops has shown a "universal" value, independent of loop size [55]. Nonetheless, increasing the loop size does increase the mutual inductance with other loops on the chip, increasing "cross-talk" between different flux bias lines. One can estimate the mutual inductance, M, by treating the short as an infinite wire. This gives the appropriate intuition that, since the field goes as $1/r$, M should be related to $\log(r_1/r_2)$, where r_1 and r_2 are the distances to the near and far edges of the loop, respectively. Consequently, M is only weakly dependent on the distance between the short and loop. To get a better estimate, one can use Biot-Savart to numerically calculate the fields from a simplified geometry, shown in figure 4.7(a). This calculation* gives a mutual inductance of ~ 160 Φ_0/A. The observed inductive coupling for such a loop is ~ 800 Φ_0/A. This difference is likely due to "flux focusing", where the Meisner effect in the surrounding superconductors concentrates the magnetic field in the gaps of the CPW. The focusing effect is difficult to account for precisely

---

* This work was done in collaboration with Hannes Majer.

for any given geometry. Consequently, the numerical calculations can only be relied upon for a rough estimate.

Relaxation

One might worry that coupling a large bandwidth ($\sim 1\,\mathrm{GHz}$, limited by cold filters) $50\,\Omega$ line to the qubit would shorten the qubit T_1 or T_2. This noise couples in two different ways: i) it directly modulates the Josephson energy of the SQUID loop, and ii) it modifies the admittance seen by the effective LC circuit. Koch *et al.* calculate the relaxation from channel i) using the fluctuation-dissipation relation. They separate the applied flux into a constant term plus noise, $\Phi = \Phi_e + \Phi_n$. Then the transmon Hamiltonian becomes $H \rightarrow H + \Phi_n \hat{A}$, with [6]

$$\hat{A} = \left.\frac{\partial H}{\partial \Phi}\right|_{\Phi_e} = E_J \frac{\pi}{\Phi_0}\left[\sin\left(\frac{\pi \Phi_e}{\Phi_0}\right)\cos\hat{\phi} - d\cos\left(\frac{\pi \Phi_e}{\Phi_0}\right)\sin\hat{\phi}\right], \tag{4.9}$$

where d is the junction asymmetry. Then the fluctuation-dissipation theorem gives [6]

$$\Gamma_1 = \frac{1}{T_1} = \frac{1}{\hbar^2}\langle e|\hat{A}|g\rangle^2 M^2 S_{I_n}(\omega_{ge}). \tag{4.10}$$

The matrix element in this expression has terms proportional to $\langle e|\cos\hat{\phi}|g\rangle$ and $\langle e|\sin\hat{\phi}|g\rangle$. Since the transmon basis states are nearly harmonic oscillator states, $|g\rangle$ and $|e\rangle$ have opposite parity in the charge basis. Consequently, the first matrix element vanishes, and the remaining term is due to the asymmetry between the two junctions of the SQUID loop. Junction fabrication tests with 10's of junctions on a chip typically show a spread of 5–10% in room temperature junction resistance, so it is reasonable to assume a similar spread in the E_J's of the two junctions when cold. If the bias line is appropriately filtered and thermalized, the spectral density of the noise is the quantum noise of a zero-T resistor, $S_{I_n}(\omega) = 2\Theta(\omega)\hbar\omega/R$. Figure 4.8 shows the dependence of (4.10) on the junction asymmetry, the mutual inductance, and the static flux bias. For parameters corresponding to real samples, $d = .1$ and $M = 800\,\Phi_0/\mathrm{A}$, the T_1 due to this coupling is $\gtrsim 10\,\mathrm{ms}$.

Channel ii) can be calculated using the admittance formalism (section 4.1) with estimates of the capacitive coupling, C_c, between the flux bias circuit and the transmon, as well as the inductance of the short, L_s. The capacitance can be simulated in Maxwell by modifying the simulated geometry to introduce gaps in the short circuit termination, as shown in figure 4.9(a), to separate the flux bias line center pin from the ground plane. This simulation
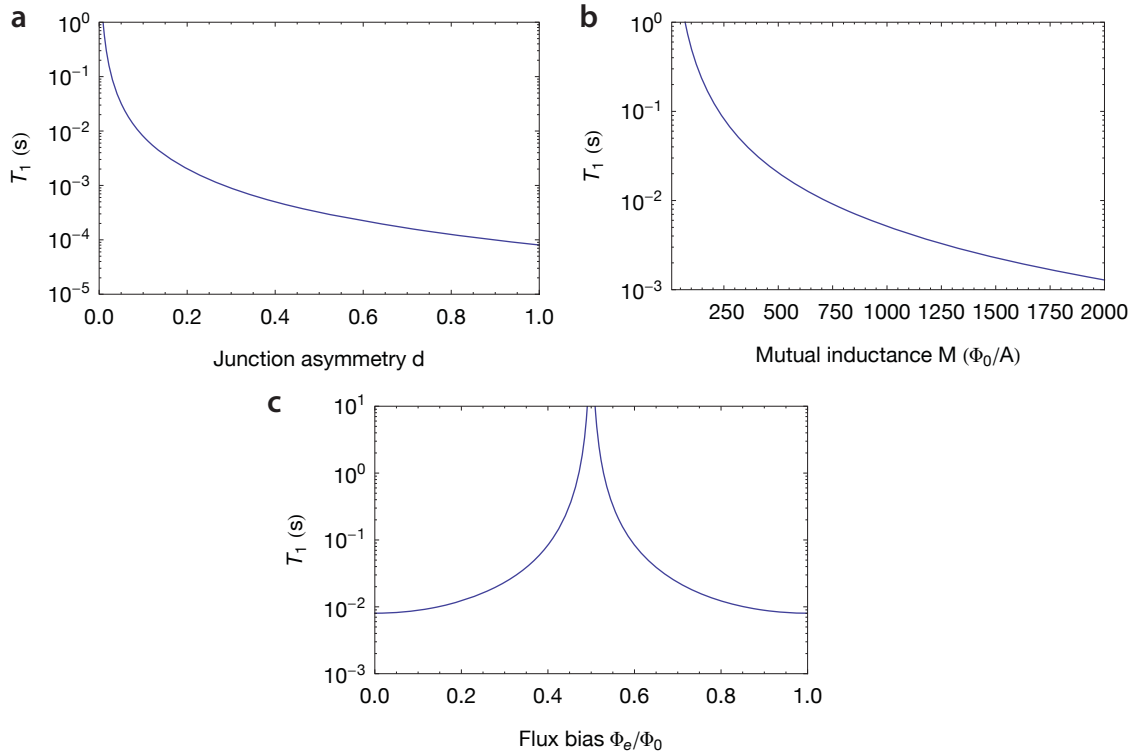
Figure 4.8: Relaxation from inductive coupling to SQUID loop. a T_1 vs. junction asymmetry, d. **b** T_1 vs. mutual inductance, M. **c** T_1 vs. flux bias Φ_e.

Figure 4.9: Relaxation from charge coupling to transmon. a Maxwell geometry for calculating the capacitive coupling to the flux bias line. Gaps are introduced in the short circuit termination in order to separate t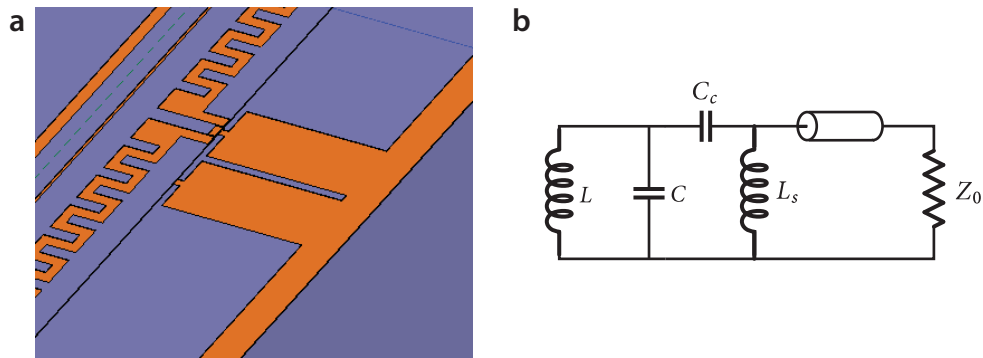he center pin of the flux bias line from the ground plane. **b** Equivalent circuit for modeling the admittance $Y(\omega)$.
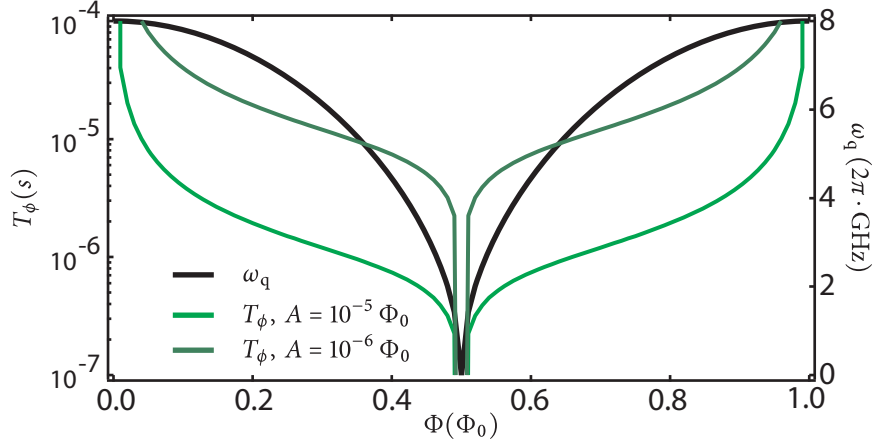
Figure 4.10: T_ϕ **due to flux noise.** The inferred dephasing from (4.12) is plotted as a function of flux bias, Φ, for $1/f$ noise spectra with amplitudes $A = 10^{-5}\,\Phi_0$ (light green) and $A = 10^{-6}\,\Phi_0$ (dark green). Reprinted from [47].

gives a capacitance $C_c \sim 3\,\text{fF}$. Similar simulations give $L_s \sim 20\,\text{pH}$. Calculating the real part of the admittance for this circuit (shown in figure 4.9(b)) gives

$$Y(\omega) = \frac{Z_0 \omega^4 L_s^2 C_c^2}{\omega^2 L_s^2 + Z_0^2(1 - \omega^2 L_s C_c)^2} \approx \frac{1}{Z_0}\left(\frac{\omega}{\omega_s}\right)^4, \tag{4.11}$$

where $\omega_s = 1/\sqrt{L_s C_c}$. Plugging this into (4.4) results in $T_1 > 100\,\mu\text{s}$, though due to the ω_s^4 dependence, this estimate is extremely sensitive to the values of C_c and L_s (an increase in L_s or C_c by a factor of 2 decreases T_1 by a factor of 4). For instance, if $C_c = 10\,\text{fF}$ and $L_s = 40\,\text{pH}$, then T_1 at 10 GHz would be $\sim 1\,\mu\text{s}$. Consequently, it is rather important to limit the capacitive coupling and inductance of the flux bias line.*

Dephasing

The current-noise induced fluctuations of E_J also cause ω_{ge} fluctuations, which causes dephasing. Since $\omega_{ge} \propto \sqrt{\cos \Phi_e/\Phi_0}$, the sensitivity of ω_{ge} to flux noise is periodic in the externally applied flux. Analogous to the charge "sweet-spots" in the CPB, there are flux sweet-spots where the transmon is first order insensitive to flux noise. Ref. [6] calculates the effect of flux noise at and away from the sweet spot, assuming a $1/f$ noise spectrum, $S_\Phi(\omega) = 2\pi A^2/\omega$.

---

* Jerry Chow's thesis contains further simulations of the admittance of the entire chip using Sonnet, which reveals that the presence of the bias lines allows an additional "wiggle-waggle" mode on the chip. This can be suppressed with wirebonds and a new package design that eliminates empty volume.[47]

Away from the sweet spot, the dephasing time is given by

$$T_\phi = \frac{\hbar}{A}\left|\frac{\partial E_{ge}}{\partial \Phi}\right|^{-1} = \frac{\hbar}{A}\frac{\Phi_0}{\pi}\left(2E_C E_J \left|\sin\frac{\pi\Phi_e}{\Phi_0}\tan\frac{\pi\Phi_e}{\Phi_0}\right|\right)^{-1/2}. \tag{4.12}$$

Consequently, the dephasing from flux noise depends on the magnitude of A. Figure 4.10 is a plot of dephasing vs. flux bias, Φ, showing the dependence of T_ϕ on the derivative of the qubit frequency. In the absence of additional coupling, various experiments have shown a "universal" value of $A \sim 10^{-5}$–$10^{-6}\,\Phi_0$ from local flux sources [55]. Finding the source of this noise, and eliminating it, is a crucial requirement for improving T_2 times in superconducting qubits. Noise which travels down the flux bias line adds on top of this background, such that the total noise spectrum is $S_\Phi(\omega) = M^2 S_I(\omega) + S_\Phi^{\mathrm{univ.}}(\omega)$. One should filter the flux bias line such that this added noise is negligible compared to the "universal" background. Assuming that the room temperature electronics has a $1/f$ voltage noise spectrum, $S_V(\omega) = 2\pi A_V^2/\omega$, then for the observed mutual inductance, $M = 800\,\Phi_0/\mathrm{A}$ and the voltage to current conversion factor of (4.8), this implies that

$$(800\,\Phi_0/\mathrm{A})^2 \frac{A_V^2}{(2.5\,\mathrm{k}\Omega)^2} \ll (10^{-6}\,\Phi_0)^2 \Rightarrow A_V \ll 3 \times 10^{-6}\,V. \tag{4.13}$$

The Yokogawa 7651 programmable DC source has a noise spectrum which depends on the selected output range. On its 10 V range setting it is specified to have $50\,\mu$V peak-to-peak voltage noise from DC to 10 Hz. This corresponds to $A_V \sim 5 \times 10^{-6}\,\mathrm{V}/\sqrt{\mathrm{Hz}}$, which is right at the level where one should expect to see some effect from the added noise. Less information is available for the Tektronix 520 and 5014 arbitrary waveform generators, which are used in RSL for fast flux control. Other members of the lab have reported observing a lower T_2 when biasing a qubit with a Tek AWG compared to biasing with a Yokogawa. This has led to a modification of the schematic of figure 4.7(b), where a bias tee is used at room temperature to allow DC flux control with a Yokogawa and fast flux control with a Tek AWG. Regardless, the noise of current electronics is likely already limiting observed T_2's in transmon experiments.

4.3   Classical Control Theory (Deconvolution)

After carefully designing this local control line to manipulate qubit frequencies on fast time scales, the natural question is: how well does it work? The answer, it turns out, is pretty well, but with a few annoyances. The photon detection experiment discussed in chapter 8

puts some extraordinary demands on the accuracy and precision of the flux control. In particular, that experiment requires that the qubit frequency is changed by ~ 1.5 GHz in 10's of nanoseconds, and that the local flux is sufficiently precise to dial-in a specific qubit frequency to within 1–2 MHz. From the perspective of the room temperature electronics, this means that voltage steps need to be correct to a part in 10^3. Unfortunately, when attempting to carry this out, I ran into problems with both the room temperature electronics (in this case a Tektronix AWG5014) and with the time-domain response of the flux bias circuit (meaning the entire path from the input at the top of the dilution refrigerator, through the filtering, onto the chip, and out whatever path the current takes through experiment ground). This section serves to document these problems as well as their resolution, which takes the common form of basic classical control theory and deconvolution of linear response. I will introduce the general concept, and then show how it is applied to the TekAWG and flux bias circuit.

In classical control theory, one uses the response of a system to modify its input until the desired response is obtained. The response, $o(t)$, of a *linear* system to an arbitrary input, $i(t)$, is determined by its transfer function, $h(t)$, through the relation

$$o(t) = \int_{-\infty}^{\infty} h(t - \tau) i(\tau) \mathrm{d}\tau. \tag{4.14}$$

If the system is causal, then $h(t) = 0$ for $t < 0$, and the integral has a natural upper limit of t. In general, one cannot make this assumption as there exist linear devices which are a-causal. For instance, many digital electronics devices apply a-causal linear filters to their inputs, resulting in a-causal transfer functions. An "ideal" system has δ-function-like transfer function, such that the output is identical to the input, perhaps with some gain factor. However, most systems, at minimum, have a finite bandwidth which causes $h(t)$ to have some width. The transfer function of a system can be measured by suitable choice of input, $i(t)$. For instance, if the input is itself a δ-function, $i(t) = \delta(t)$, then (4.14) gives

$$O(t) = \int_{-\infty}^{\infty} h(t - \tau) \delta(\tau) \mathrm{d}\tau = h(t). \tag{4.15}$$

Unfortunately, δ-functions are hard to produce experimentally in practice. A more convenient choice is given by a step function input, $I(t) = \Theta(t)$. Then,

$$O(t) = \int_{-\infty}^{\infty} h(t - \tau) \Theta(\tau) \mathrm{d}\tau = \int_{0}^{\infty} h(t - \tau) \mathrm{d}\tau \Rightarrow O'(t) = h(t). \tag{4.16}$$

Once $h(t)$ is known, one can modify the system input to get the desired output. That is, I

want to find the inverse transfer function $h^{-1}(t)$ such that

$$\tilde{i}(t) = \int_{-\infty}^{\infty} h^{-1}(t - \tau)i(\tau)\mathrm{d}\tau \tag{4.17}$$

produces $o(t) = h(t) \circ \tilde{i}(t) = i(t)$.

Equation (4.14) is a convolution, so in the Fourier domain,

$$O(\omega) = H(\omega)I(\omega), \tag{4.18}$$

where I use uppercase letters to denote the Fourier transform of a function. Then the inverse transfer function can be simply computed by division, $H^{-1}(\omega) = 1/H(\omega)$, and transformed back into the time domain. This procedure fails if $H(\omega)$ vanishes for some ω. Fortunately, for many situations of interest, this is not the case.

A further detail of note is that the expressions above are all for continuous functions, but in practice the inputs and outputs are sampled at discrete time steps. The deconvolution method still works, but integrals become sums, and Fourier transforms become discrete Fourier transforms. The details of appropriate choices of 2π's and units in numerical derivatives are given in the recipes of the next sections, where the above method is applied to the Tektronix AWG5014 and the flux bias circuit.

Tektronix AWG5014

The Tektronix AWG5014 is a very fast (1.2 GS/s), multi-channel digital-to-analog converter. It is ideally suited for testing large bandwidth communications devices, like cellphones and Wi-Fi routers. For these applications, speed is paramount, while accuracy and precision are of lesser concern. Consequently, the TekAWG does not automatically match the required specifications for driving a flux bias line in some experiments. The typical response of an analog output channel for a step function input is shown in figure 4.11(a). This response displays several anomalous features. The first is that after the initial sharp rise which travels 99% toward the programmed value in under 2 ns, there follows a slow (> 1 μs) rise until the output reaches the final value. Stranger yet, the output of the AWG anticipates the step with a small 0.5% rise ~ 200 ns *before* the step. This second feature suggests that the AWG applies some kind of digital filtering to its output.[*] These two features scale linearly with the step size, so it is likely that they can be fixed through deconvolution methods. It is important to

---

[*] Somewhat ironically, this same hardware could be used to fix these issues.
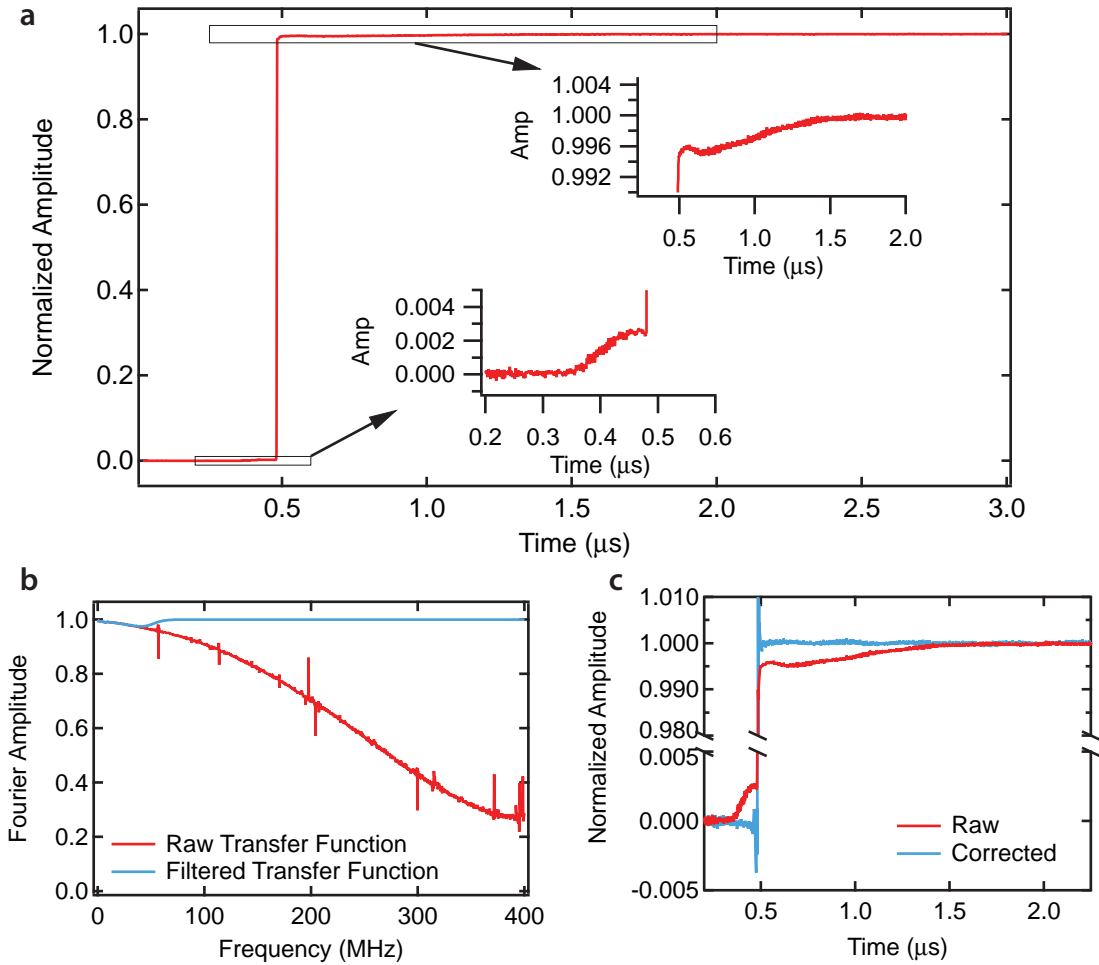
Figure 4.11: Correcting the output of the Tektronix AWG5014. a The issue. **b** The raw and filtered transfer functions. **c** Applying the deconvolution kernel to fix the response.

remember, however, that the TekAWG is not a completely linear device. As alluded to in Refs. [47, 56], the TekAWG also has additional anomalous behavior on the sub-10 ns timescale that is non-linear, and consequently cannot be corrected with deconvolution. Thus, we need to modify the deconvolution procedure to only deal with the linear part of the response function.

The following recipe can be used to find a deconvolution kernel for the linear response of a Tektronix AWG5014:

1. Measure the response, r_i of a step function input at the same sample rate as the desired output (usually $1 \, GS/s$), making sure to sample the regions before and after the step

edge. Subtract any DC offset and scale the data such that the step goes from 0 to 1 in normalized units.

2. Compute the transfer function h_i by doing a numerical derivative, $h_i = r_i - r_{i-1}$. Do *not* divide by the sampling interval, Δt, so that h_i is in units of Δt.

3. Rotate the points of h_i such that h_0 is the beginning of the step edge. Compute the FFT of h_i, $H(f_i)$.

4. The resulting spectrum, shown in red in figure 4.11(b), is noisy. We do not want to introduce additional noise in the output, so fit the real and imaginary parts of $H(f_i)$ to n^{th} order polynomials where n is some small number, like 5–20. Construct $H_{\text{smooth}}(f_i)$ by pulling from the data for the low frequency points and from the polynomial fits for the high frequency points, *i.e.*

$$H_{\text{smooth}}(f_i) = \begin{cases} H(f_i) & f_i < 10\,\text{MHz} \\ \text{fit of } H(f_i) & f_i \geq 10\,\text{MHz} \end{cases}, \tag{4.19}$$

to ensure that the relevant data remains while smoothing out the rest of the spectrum.

5. The spectrum still contains high frequency information, including the roll-off from the $\sim 300\,\text{MHz}$ analog bandwidth of the AWG. Apply a filter which removes this high frequency information by pushing $H(f_i)$ to 1. The particular form of this filter does not matter. One such choice is

$$H_{\text{smooth, filtered}}(f_i) = H_{\text{smooth}}(f_i) \cdot F(f_i), \tag{4.20}$$

where

$$F(f_i) = 1 + \frac{1}{2}\left(\frac{1}{H_{\text{smooth}}(f_i)} - 1\right)\left(\tanh\left(\frac{f_i - f_0}{\sigma}\right) + 1\right), \tag{4.21}$$

with $f_0 \approx 50\,\text{MHz}$ and $\sigma \approx 20\,\text{MHz}$. The resulting spectrum is shown in blue in figure 4.11(b).

6. Find the inverse, $G(f_i) = 1/H_{\text{smooth, filtered}}(f_i)$, and perform an inverse FFT.

7. Finally, rotate the result such that the point representing $t = 0$ is the central point in the list.

This recipe produces a deconvolution kernel which works quite well in practice, improving the accuracy of programmed voltage levels by a factor of 5–10. Figure 4.11(c) shows the result of applying this kernel to a step function input. One can see that the step is now much flatter—the anomalous features before and after the step edge are gone. Importantly, subsequent steps no longer depend on the previous 1–2 μs of output. However, the step edges now display some overshoot due to the gain of the inverse transfer function at intermediate frequencies. This can be removed by slightly rounding pulse edges before applying the deconvolution kernel.

Flux bias circuit

With the response of the room temperature electronics under control, we can move on to examining the response of the rest of the flux bias circuit. Characterizing this circuit is not as straightforward as hooking up a fast oscilloscope, since we are interested in the local flux on the chip. Instead, the qubit itself can be used as a local magnetic field probe.*

Figure 4.12(a) shows a schematic of an experimental protocol to measure the time-domain flux response. A voltage pulse of width w and height h is applied to the bias line, then the qubit frequency is tracked as the local flux returns to its original value. Ideally, the width w should be chosen such that the flux bias reaches steady-state. In this case, the procedure is equivalent to testing the step function response of the bias line. To track the qubit frequency, a weak ~ 20 ns spectroscopy pulse is scanned in frequency for each value of delay from the trailing edge of the flux pulse. The resulting spectrum, shown in figure 4.12(b), displays a rather broad Gaussian profile because of the large bandwidth of the spectroscopy pulse. Nonetheless, the center frequencies of these Gaussians can be extracted with high precision, giving the time domain response of the qubit frequency after the step edge, shown in figure 4.12(c).

This method has some limitations. First of all, it requires a frequency sweep for every measured delay. This is sufficiently time consuming that one typically needs to examine a small frequency window, which is equivalent to zooming in on the flat section of the step response, rather than recording the entire step. Second, this method has a much restricted bandwidth compared to the flux bias control. The time resolution of the qubit frequency measurement is roughly equivalent to the spectroscopy pulse duration. One can use a shorter pulse, but this leads to a broader Gaussian response, and thus greater uncertainty in the

---

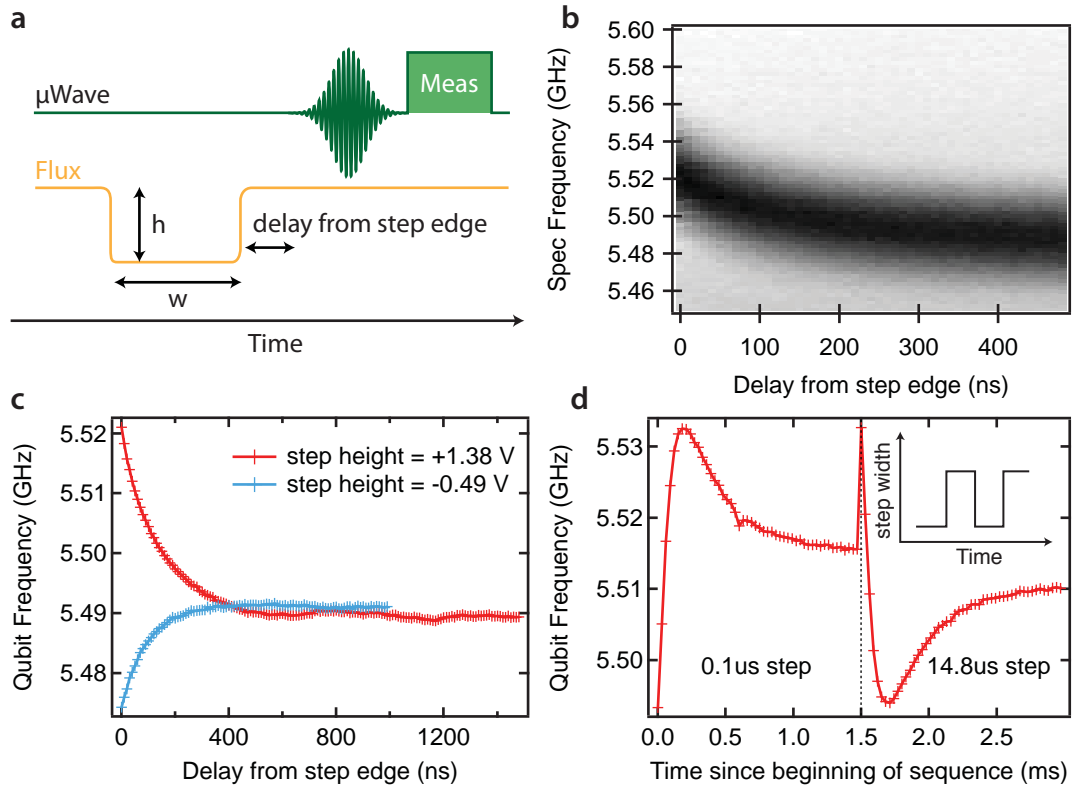* The supplement of Ref. [24] describes a similar method for charactering flux response.

Figure 4.12: Flux bias line response. a "Qubit scope" experiment pulse sequence. **b** Measured homodyne phase shift as a function of spectroscopy frequency and delay from the trailing edge of the flux pulse. **c** The center frequencies of data such as **b** for various flux pulse heights. **d** The flux bias line also has a millisecond timescale behavior which can be observed by alternating between repeated flux scope experiments of fixed delay but different step widths or heights.

center frequency. Therefore, deconvolving the linear response will require resampling or interpolation at the AWG sampling rate.

 Nonetheless, this "qubit scope" method is sufficient to reveal a number of flux bias response features. First of all, the data of figure 4.12(c) show the qubit response to voltage steps which change the qubit frequency by ~ +1 GHz and −500 MHz, respectively. In each case, the qubit response has an exponential tail of roughly 3% amplitude and a time constant $\tau \approx 100$ ns. This behavior consistently scales with the amplitude of the flux bias step, which is a good indication that it is a linear effect which can be corrected by deconvolution. Unfortunately, the flux bias line also has a much slower response behavior which is not apparent in figure 4.12(c). If the "qubit scope" procedure is repeated many times for one step width, w_1, and then switched to a different width, w_2, in a square wave pattern, figure 4.12(d) shows that the qubit frequency
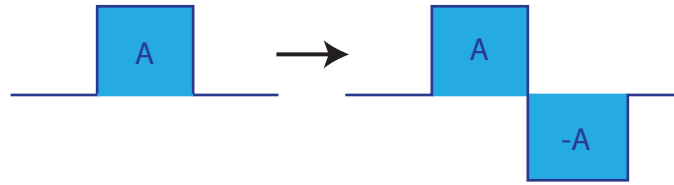
Figure 4.13: Flux balancing. The slow flux bias behavior is compensated by replacing current pulses that enclose a positive area A with a pair of pulses with the second mirrored about the DC value, such that the average current integrated over the entire pulse sequence is equal to the DC current. This second pulse does not need to immediately follow the first, so long as the time between the pulses is short compared to the slow flux bias response time.

actually changes despite nominally measuring the same thing many times. What is really happening is that the average value of the current on the bias line is different for the two pulse widths, and the bias line responds to this change in average value on a timescale of roughly a millisecond. This response points to the existence of a large effective inductance, L_{eff} in the bias circuit. The origin of this inductance is currently unknown, but fortunately, the L_{eff}/R time for the slow response is 2–3 orders of magnitude larger than a typical experiment. Consequently, we can employ a trick and reprogram any sequence of flux bias pulses to maintain a constant average value of current, as shown in figure 4.13. This prevents the flux in L_{eff} from changing when looping through a sequence of pulses.

The fast time scale behavior can be deconvolved using a similar recipe to the one presented for the Tektronix AWG5014, but with a few modifications. First, due to the windowing and sampling issues discussed above, I find that it is easier to use an effective transfer function model than it is to try to convert qubit frequencies back into voltages at the output of AWG. Consequently, instead of starting with the measured step response, I use a model of an instantaneous rise with a residual exponential tail, where the amplitude and time constant are given by fits to the measured "qubit scope" response. One can proceed as before, but omitting steps 4 and 5 since they are not necessary with the idealized model. The resulting qubit frequency control is precise to 2–3 MHz. This is good enough for the experiment of chapter 8, but leads to substantial phase errors in extended pulse sequences. Consequently, it will be necessary to find and remove the sources of non-ideal flux bias response as experiments grow in duration and complexity.

Device Fabrication and Experiment Setup

T HE realization of a circuit QED experiment involves the coordination of many elements. From the perspective of smallest to largest size, this begins with the electrical design of the qubits and cavities. These circuits are fabricated on a silicon or sapphire substrate, with optical lithography used to define the resonator and bias ports, and e-beam lithography to define the qubits. The chip is mounted in a package which consists of a PC board that plugs into a sample holder. The sample holder, in turn, is mounted on the base plate of a dilution refrigerator, and filtered, large-bandwidth transmission lines connect the sample ports to room temperature electronics. The electronics include modulated microwave sources, fast pattern generators to modulate the input signals, amplifiers, filters, and a fast analog-to-digital converter to digitize the output signals. All of this is controlled by custom software written in LabVIEW and Mathematica. Much of this infrastructure is independent of the particular experiment, and so many of the details regarding fabrication and experiment setup can be found in [15, 47]. Rather than repeating that material, I will describe changes and developments to the circuit QED infrastructure that are new to this work. The principal developments include: a modified e-beam lithography process to make devices on sapphire substrates, a sample holder which allows for up to 8 control ports, filters for flux bias lines, microwave setup to facilitate reflection measurements, pulse modulation techniques with external I/Q mixers, and a domain specific language for programming pulse sequences.

5.1 Circuit QED devices

All devices presented in this thesis are made with a fairly simple process that uses optical lithography to define large features and e-beam lithography for small features. The basic process is as follows: niobium is DC sputter deposited on a silicon or corundum (sapphire) substrate. Then two layers of photoresist are spun and the large features, including the resonators, flux bias lines, and input and output ports, are defined with contact optical lithography, after which the Nb is dry etched by fluorine-based reactive ion etching. The wafer is then diced into individual chips before two layers of e-beam resist are spun on a single chip. E-beam lithography defines the qubit geometry, and aluminum is deposited with a double-angle e-beam evaporation step to create Al-AlO_x-Al Josephson junctions by the Dolan bridge technique. The details of this process for silicon substrates are presented in [15]. In the following sections I will describe two developments to this process. The first is a proliferation of qubit designs, the reasoning behind several of which are presented in chapter 3 and chapter 4. The other development is a change to fabricating devices on corundum substrates. This change is motivated by loss tangent measurements which suggest that corundum presents less intrinsic loss than silicon for high-Q devices [57, 58].

5.1.1 Transmon zoo

The history of circuit QED with transmon qubits has seen the evolution of the original transmon design to suit different needs. Consequently, there are now many different 'species' of transmon in the wild. Figure 5.1 presents many of these designs along with tongue-in-cheek names derived from similarities in appearance to various animals. Figure 5.1(a) is the original transmon design that has $E_C/2\pi \simeq 300\,\text{MHz}$ and coupling strength $g_{ge}/2\pi \sim 150$–$250\,\text{MHz}$ to the CPW mode, but that also has significant undesired coupling to the spurious slotline mode. Figure 5.1(b and c) are variations on 'balanced' transmons (see section 4.1.1) that attempt to reduce this spurious coupling. The constraint of fitting these designs into the same $300\,\mu\text{m} \times 20\,\mu\text{m}$ slot required reducing the coupling to the desired mode to $g_{ge}/2\pi \sim 50$–$100\,\text{MHz}$. For some experiments, such as the strong driving of the vacuum Rabi measurement of [19], a smaller g_{ge} is desired, and so these balanced designs are also useful in such cases. Figure 5.1(d) is made for a deeper $300\,\mu\text{m} \times 50\,\mu\text{m}$ slot in order to make a larger shunting capacitor. This decreases $E_C/2\pi$ to $\sim 200\,\text{MHz}$, which is useful for maintaining large E_J/E_C ratios at smaller ω_{ge}. Figure 5.1(e and f) are 'flux bias' transmons that increase the size of
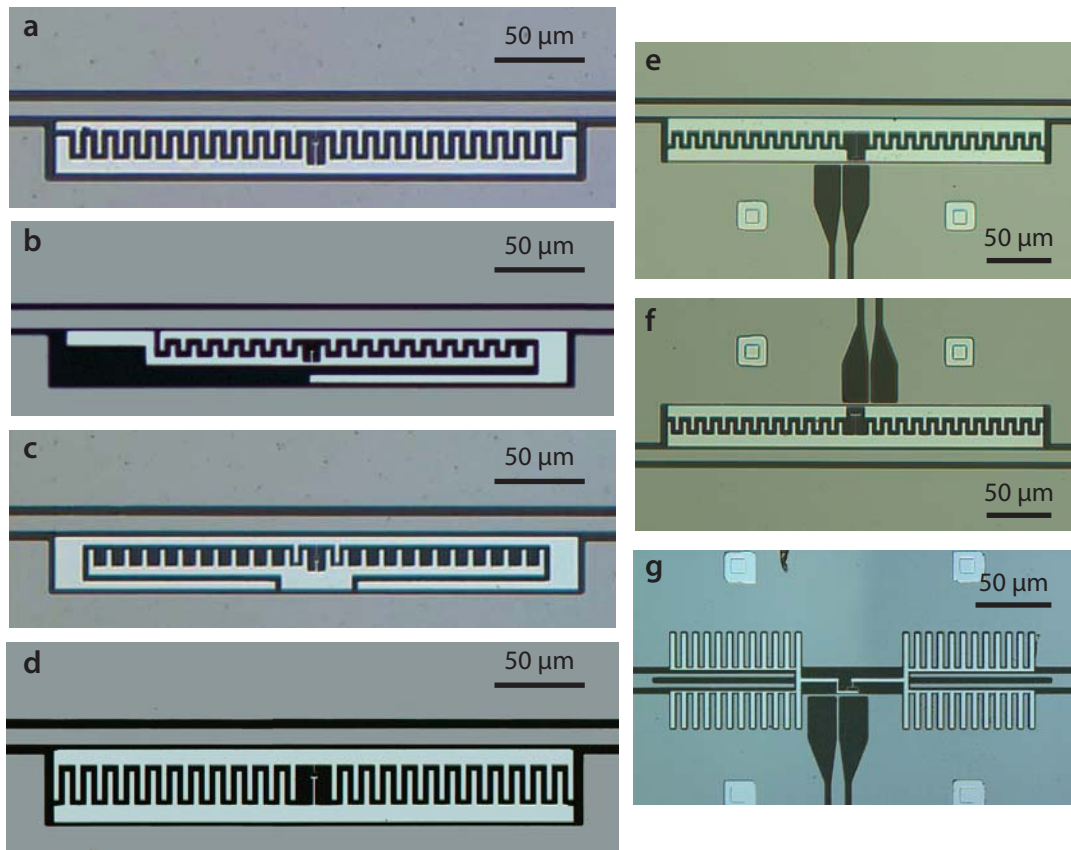
Figure 5.1: Optical images of transmon designs. a The original 'standard moose' transmon. **b** The 'crocodile' balanced transmon. **c** The 'moose with helmet' transmon. **d** The 'deep slot' or 'long antler moose' transmon. **e** and **f** 'Flux bias' transmons with the SQUID loop at the top or bottom. **f** The 'sarantapede' transmon.

the SQUID loop and bring it to the edge of the slot nearest to the short on the flux bias line. Since (e) is below the center pin and (f) is above, in order to fabricate Josephson junctions for (e) and (f) simultaneously and with nearly-identical E_J, the SQUID loop design in (f) had to be modified slightly. Figure 5.1(g) shows a 'sarantapede' qubit, which edge-couples to the center pins of two separate resonators (the forks on the left and right) while minimizing the capacitance between the resonators (see section 3.4).

5.1.2   Qubit fabrication on sapphire substrates

Electron-beam (e-beam) lithography requires accurate control of a focused beam of electrons. Almost all of the charge delivered to a small area of resist is due to 'secondary electrons', which are electrons that scatter off the substrate. For conducting substrates, the remaining electrons are carried away to the edges of the substrate and into grounding clips, which also serve to hold the chip in-place. On highly insulating substrates, however, the charge is trapped near the surface, causing the substrate to accumulate localized pockets of charge. This is problematic because these trapped charges can deflect the beam, making it impossible to precisely define small features.

Lightly doped silicon with resistance greater than $10 \, \mathrm{k\Omega \, cm}$ is sufficiently conductive at room temperature to not affect e-beam lithography. Corundum substrates, on the other hand, do not work with an unmodified e-beam process. Luigi Frunzio and I developed a simple fix by depositing a thin (~ 10 nm) layer of metal on top of the resist. In our case, since we were already using Al in the deposition system, we used e-beam evaporation to cover the MMA/PMMA resist stack with 12–13 nm of Al. This layer of metal provides an effective channel to carry away excess charge once the sample is clipped down to the SEM stage. After the device is exposed, the Al top layer is removed before development. We use a common photoresist developer, *Shipley* Microposit MF-312, which is also an effective wet etch for Al. After the Al is removed, we continue with the standard MMA/PMMA development process of MIBK:IPA followed by double-angle e-beam evaporation of aluminum in the *Plassys* system. These steps modify the e-beam resist such that exposure dose must be increased 10–20% compared to the silicon process.

Initial tests of junctions made on R-plane corundum displayed an anomalous aging effect. Usually, the diffusion of oxygen into the surrounding aluminum of the Al-AlO$_x$-Al junction is a self-limiting process that evolves on a timescale of a few hours. In practice this means that the normal resistance, or R_N, of Josephson junctions made on silicon increases by ~ 20–30% over the course of a few days after the deposition and lift-off process is finished. After this, the resistance stops increasing. This junction 'aging' is sufficiently repeatable that it can be accounted for when choosing junction area and opacity, the latter being determined by the duration and pressure of the oxidation step of the deposition process. In contrast, initial aging tests of junctions made on R-plane corundum continually increased in resistance for weeks without stopping, as shown in figure 5.2. This is undesirable because the device properties
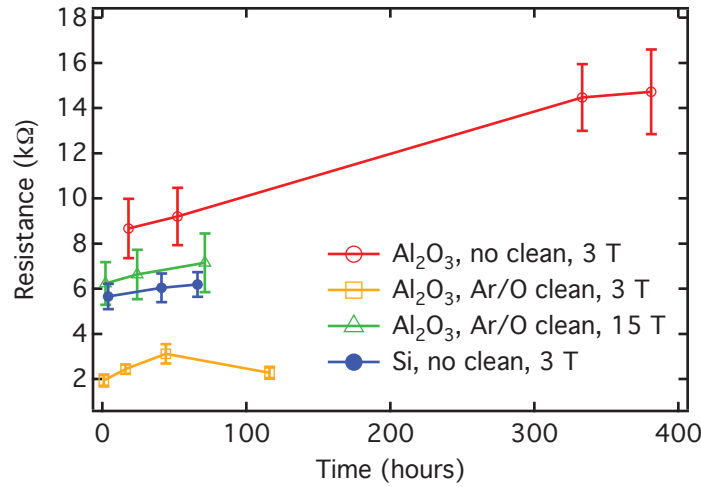
Figure 5.2: Josephson junction aging on silicon and corundum. Junctions made on corundum substrates without an additional cleaning step (red) display a runaway effect with aging that appears not to be self-limiting after 15 days. By adding an argon/oxygen plasma cleaning step, the junction aging is controlled (orange), but results in junctions with much lower resistance than given by the usual recipe on silicon (blue). Increasing the pressure in the oxidation step to 15 Torr gives junctions with similar transparency (green) to the silicon recipe. Error bars are the standard deviation in measured resistance across ~ 15 devices.

are then largely determined by the time between fabrication and cool-down, which is not always predictable.

The source of this run-away aging remains a mystery; however, Luigi and I did find a way to prevent it with a modified deposition recipe. We added a gentle argon/oxygen ion milling/cleaning step prior to depositing the base layer of aluminum. With this added step, junction aging on corundum displayed similar behavior to aging on silicon (see figure 5.2). Later tests have shown that this recipe also works with c-plane corundum, though we have never tried aging tests on c-plane corundum without the additional cleaning step. For the full e-beam process and deposition recipe, see appendix C.

5.2   Measurement setup

Performing microwave measurements on superconducting samples at 20 mK requires extreme care in the design of the cryogenic measurement setup. Filtering, attenuation, thermal anchoring, cable choice—all of these play a role in separating the signals which are intended to interact with the samples from the thermal background provided by the room temper-
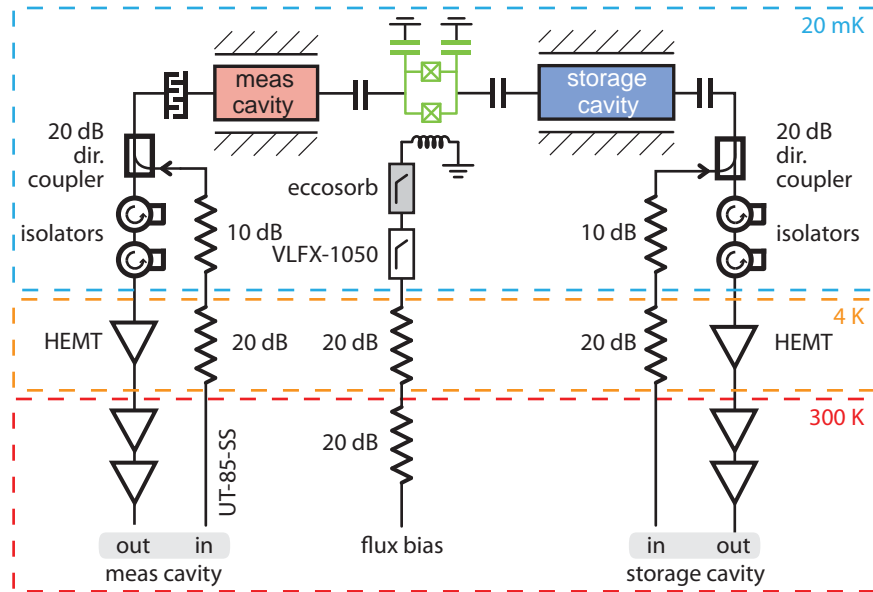
Figure 5.3: Schematic of cryogenic measurement setup. The two cavity experiments have 5 ports: each cavity has an input and output port, provided by directional couplers, plus there is an additional port for the qubit flux bias line.

ature and intermediate temperatures that contact the cables connected to the sample. The experiments discussed in this thesis all used a cryogenic measurement setup similar to the design discussed in [15, chapter 6]. This design was modified for the two-cavity experiment of chapter 8 to allow for amplification and the measurement of the reflected signal off of the input ports of each cavity, and for fast control of the qubit frequency via the flux bias line.

There are two common microwave tools for measuring reflected signals: circulators and directional couplers. The first is a 3-port, non-reciprocal device. In cryogenic experiments, circulators are frequently found between the cold sample and the amplifier which is thermalized to 4 K. In this use, the third port of the circulator is terminated with 50 Ω, such that the resulting device acts like a one-way valve for microwave photons. Then signal photons from the sample pass through with little-to-no attenuation, while noise photons from the amplifier are heavily attenuated ($\sim 20\,$dB) when passing through in the opposite direction. The effective 2-port device formed when the third port is terminated is sometimes called an *isolator*. In fact, the components shown in front of the HEMT amplifiers in figure 5.3 are *Pamtech* broadband 4–12 GHz isolators (model CWJ1019), where the third port is internally terminated, so it is not accessible to the experimenter. In a standard circulator, however, the third port can also be used to send signals the other way, providing a means to separate incoming and
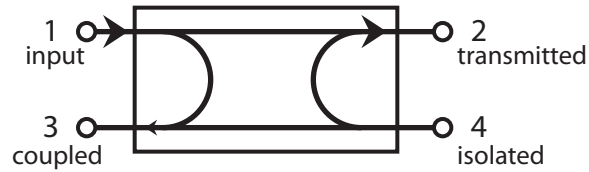
Figure 5.4: A directional coupler is a four port device that splits a signal between multiple outputs. A signal injected at port 1 is largely transmitted to port 2, with a much smaller fraction of the signal going to port 3 (the coupled port). Ideally, no signal arrives at port 4 (the isolated port). There is nothing special about the labeling of ports in the diagram: signal can be injected into any port. For instance, if port 4 is the input, then 3 is the transmitted port, 2 is the isolated port, and 1 is the coupled port.

outgoing signals to and from the device. Circulators are typically very narrow-band devices; the broadband variety used in these experiments are very expensive, costing several thousand dollars each. Fortunately, for the purpose at hand, a much less expensive directional coupler will suffice.

A directional coupler, schematically illustrated in figure 5.4, is a reciprocal 4-port device that splits a signal sent into any port onto two outputs. The division ratio is typically very asymmetrical, such that the strongly coupled (or transmitted) port receives a nearly un-affected signal, while the other port gets a copy which is attenuated by 10–20 dB. The ratio of this coupling, and the bandwidth over which the component works, are the primary specifications of any directional coupler. The couplers shown in figure 5.3 are 20 dB models. When strong attenuation of one direction of photons is acceptable, a directional coupler can be used to measure reflected signals by connecting the sample and amplifier to a pair of strongly coupled ports, and using the port which weakly couples to the sample port for the drive. A 50 Ω terminator (from XMA) is connected to the fourth port, which is strongly coupled to the drive. It serves as an effective sink for the majority of the power sent down the 'input' cable. For sufficient attenuation* of thermal photons coming from higher temperature stages, one typically needs ~ 30 dB of attenuation of the input at the mixing chamber. Since the directional coupler already provides 20 dB of attenuation, an additional 10 dB attenuator is sufficient.

The other modification of the cryogenic setup is the addition of the flux bias line, which requires a very different kind of filtering than the microwave lines. Whereas the input and

---

* 'Sufficient attenuation' is application specific, but in this case we designed for a thermal photon number, $n_{th} < 0.01$.

output ports for the cavities need to transmit 4–12 GHz signals, the flux bias line only requires sufficient bandwidth to be fully controlled by a Tektronix AWG520 or AWG5014, which fast arbitrary waveform generators. Each instrument is specified to output signals at 1 GS/s; however, both devices have ∼ 300 MHz of analog bandwidth, which is significantly below the Nyquist frequency of 500 MHz. The flux bias lines also control the DC flux seen by qubit, so these lines should have bandwidth DC–500 MHz. Unfortunately, the usual tool for filtering DC lines in cryogenic experiments, the copper powder filter, has a much lower cutoff frequency of a few 10's of MHz. Fortunately, recent work by Daniel Santavicca [59] has produced a new kind of filter, called an Eccosorb filter, which is discussed below in section 5.2.2. This absorptive filter gives attenuation of more than 30 dB above 4 GHz, which is likely sufficient to filter out thermal noise. Before learning about Eccosorb filters, we used 'chocolate' filters, which are modified powder filters that pack copper powder around an embedded microstrip transmission line, rather than the long coil of wire of the traditional powder filter*. These filters have cutoff frequencies ∼ 500 MHz, but the attenuation in the stopband is rather weak, with a slope on the order of 10 dB/octave. To supplement the chocolate filter, a *Minicircuits* VLFX-1050 low-pass filter was used in series. The *Mini-Circuits* filter has a much sharper cutoff, but is a reflective filter. Consequently, it was placed with the chocolate filter in-between it and the sample to prevent forming a cavity. The Eccosorb filters have a much steeper stopband slope of nearly 20 dB/octave, but we have not tried it without the Mini-Circuits filter, so it is not known if the Eccosorb filter alone is sufficient.

5.2.1   Sample holders

As frequently discussed in chapter 4 and chapter 6, careful control of the electromagnetic environment is crucial to protect qubits from additional relaxation channels. This same type of care must be taken in designing the package to mount a circuit QED sample. The 'coffin' sample holders presented in [15] are a very simple rectangular design, that display a clean electromagnetic environment up to ∼ 12 GHz. For experiments with more than one qubit or cavity, however, a new package is required that provides more than two ports. The 'octobox' sample holder, shown in figure 5.5, is a new design that I developed with David Schuster†. It

---

* The name 'chocolate' filter comes from the light-brown appearance of the copper powder and epoxy mixture. I would not recommend eating it, though.

† Continuing the tradition of naming sample packages after sea creatures, the name references the octopus-like appearance of the package when it is closed up and has cables hanging from all eight ports.
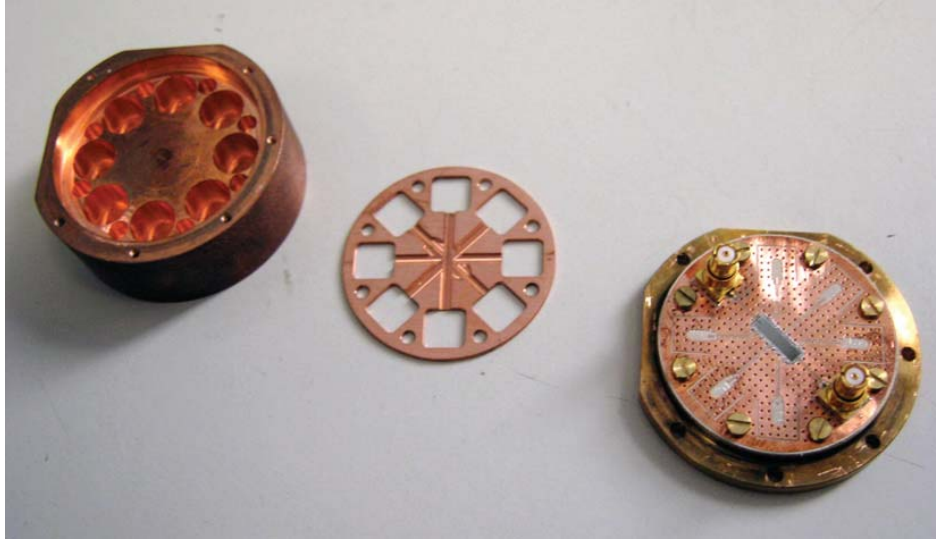
Figure 5.5: Octobox sample holder. The 'octobox' is a compact quasi-cylindrical design for circuit QED devices with up to 8 ports. In this design iteration, the lid (left) fills up the majority of the volume, except for cylindrical gaps for the Rosenberger SMP launchers and the 'bullets' which are intermediaries in SMP connections. The shim (center) and PC board completely occupy the remaining volume, helping to suppress spurious modes. The base (right) with mounted PC board and sample fits into the lid with a unique orientation determined by the flats on two sides.

is a very compact quasi-cylindrical design with eight ports that can be connected to a sample. The size and shape of the box were motivated by an effort to design a package which could be used in a hot-swap system where samples could be mounted and un-mounted at the base of a dilution refrigerator while the fridge was running. The cylindrical package allows it to fit inside a line-of-sight port on an *Oxford* Kelvinox or Triton dilution refrigerator. The flats on two sides define a unique orientation for blind mounting. While some effort has been spent designing this hot-swap system, it turns out that the new cryogen-free, or 'dry', dilution refrigerators can be turned around sufficiently quickly that there is less demand for a hot-swappable package.

In the current design iteration, the octobox consists of three separate pieces: a lid, a base, and a shim. The shim is a recent addition to the package design from Leo DiCarlo after a careful study of the mode structure in a previous iteration of the design. It is a copper spacer that is mounted flush with the top of the PC board. Trenches are milled over the CPW traces and sample pocket to prevent shorting out any traces or wire bonds. Otherwise, the shim fills the gap between the PCB and the lid to suppress forming Helmholtz resonator and
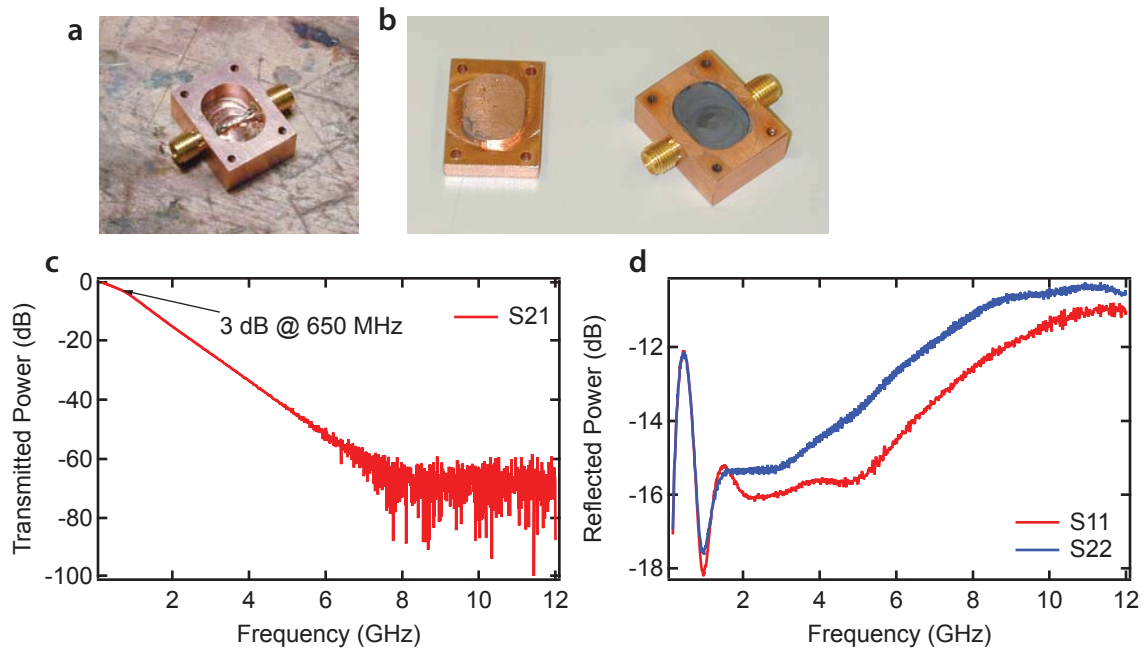
Figure 5.6: Eccosorb filter. The filter is based upon a simple geometry of a microstrip embedded in a lossy dielectric [59]. **a** Filter box before adding the CD-124 epoxy, showing the suspended copper strip. **b** The box after the epoxy has cured and has been milled down so that the lid will mount flush with the top of the box. The filter can be allowed to cure with the lid on as well, as long as one takes care to avoid filling the mounting holes with epoxy. **c, d** The transmission and reflection measurements of the filter at room temperature taken with a network analyzer. This particular filter has a 3 dB point at 650 MHz and is well matched to 50 Ω as evidenced by the S_{11} and S_{22} which are below −10 dB over the range 0.5–12 GHz.

whispering gallery-type modes in the cylindrical space. With the lid and shim filling almost all empty volume, the octobox sample holder functions quite well, delivering performance similar to the previous coffin box but with eight ports. See [47, chapter 5] for discussion of parasitic modes in the octobox and the use of wire bonds for improved performance.

5.2.2   Eccosorb filters

The Eccosorb filter has recently been developed by Daniel Santavicca [59] to fill a gap in the space of existing cryogenic filters. The metal powder filter is the workhorse for transport measurements requiring bandwidth from ~ DC–1 MHz, but building absorptive filters that transmit radio frequencies while still providing sufficient attenuation of microwaves is extremely challenging due to the difficulty in controlling impedances over many orders of

magnitude in frequency. In particular, most lumped element filters have resonances at frequencies in the nominal stopband, which are undesirable for qubit experiments. In contrast, distributed element filters based on microstrip transmission lines can have well-controlled impedances far into the microwave without resonances. The addition of a lossy dielectric material to the transmission line provides increasing attenuation as the frequency increases. In the Eccosorb filter, the dielectric is a commercially available magnetically-loaded silicone or epoxy material which is produced by *Emerson & Cummings*. The Eccosorb dielectric provides much greater attenuation per octave in the stopband than similar filters made with copper or stainless steel powders.

Construction of the filter is straightforward (see figure 5.6). Its properties are determined by the dimensions of the box, with the width and height controlling the impedance, and the length setting the 3 dB frequency. I followed the recipe for the CD-124 epoxy material described in [59], designing a box of length 0.5 inches to get a cut-off frequency above 500 MHz. A piece of copper foil is soldered to two push-mounted SMA connectors before the volume of the box is filled with CD-124. The epoxy is cured overnight at 75° C, and then the surface is milled such that the lid mounts flush to the box. Room temperature measurements of this filter with a network analyzer show that it has a 3 dB frequency of ~ 650 MHz and that it provides nearly 20 dB/octave attenuation in the stopband. The device is also well-matched to 50 Ω as evidenced by $S_{11}, S_{22} < -10$ dB. This particular filter has not been separately measured cold, though similar devices in [59] displayed an increase of the 3 dB frequency of roughly 20 % when cooled to 4 K.

One of the concerns with using magnetic materials is that they could have very long thermal relaxation times, providing a heat source that could warm up the sample. This issue was investigated in [60], where the authors show that filters made with the silicone Eccosorb material quickly thermalize to the base temperature of a dilution refrigerator. I have not done equivalent measurements with the CD-124 epoxy, but having used these filters in many experiments, I can report no evidence of spurious heating from these filters.

5.2.3 Pulse Generation

The fast single-qubit manipulation is achieved by modulating the in-phase (I) and quadrature-phase (Q) components of a microwave carrier signal at the qubit frequency. The usual way this is done in Schoelkopf lab (as in [10, 56, 61, 62]) is to use the *Agilent* E8267D vector microwave generator which has a built-in IQ mixer. This is a remarkable instrument, providing highly
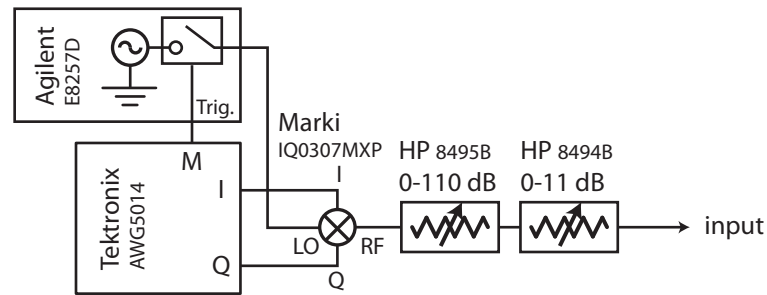
Figure 5.7: Microwave pulse generation setup. An Agilent E8257D microwave source drives the LO port of an IQ mixer. Baseband signals (pulse envelopes) are generated using two channels of a Tektronix AWG5014, which is a fast arbitrary waveform generator. Step attenuators control the output power of the pulses.

linear control of both quadratures with $1\,\mathrm{GHz}$ of IF bandwidth, carrier suppression above $80\,\mathrm{dB}$, and 0.1 degree phase adjustments of the baseband signals. Unfortunately, the cost of the 'vector' feature is prohibitively expensive, adding roughly $20,000 to the price of the generator. Consequently, I endeavored to find an alternative solution.

The pulse generation setup employed in the experiment of chapter 8 is shown in figure 5.7. The internal IQ mixer of the *Agilent* E8267D is replaced with a *Marki* IQ0307MXP IQ mixer. This external mixer can perform nearly as well as the internal mixer of the 8267D by making a few modifications. First, the output of the mixer is linear in the voltage applied at the I and Q inputs as long as the mixer is appropriately biased. Since mixers derive their bias from the power supplied on the LO, this input power must be held constant and within the designed range of the mixer. For the Marki IQ0307MXP, this means that the LO power is fixed between $+12$–$14\,\mathrm{dBm}$. To achieve control of the output power over many decades, step attenuators are placed at the output (the RF port) of the IQ mixer. Component mixers also tend to have poor carrier leakage performance, meaning that even when $0\,\mathrm{V}$ is applied at both I and Q inputs, there is significant power emitted at the RF port. For the Marki mixer, this carrier leakage is roughly $-20\,\mathrm{dB}$. This can be greatly improved by performing fine adjustments of the I and Q voltages to zero the carrier leakage. In practice, this was done by connecting the RF output port to a spectrum analyzer and searching over the I and Q offset voltages on the TekAWG until a minimum output power was found. This method is sufficient to achieve carrier suppression on the order of $60\,\mathrm{dB}$. The limiting factor is the precision of the voltage offsets in the DAC outputs of the TekAWG, which provide a minimum step size of $1\,\mathrm{mV}$.

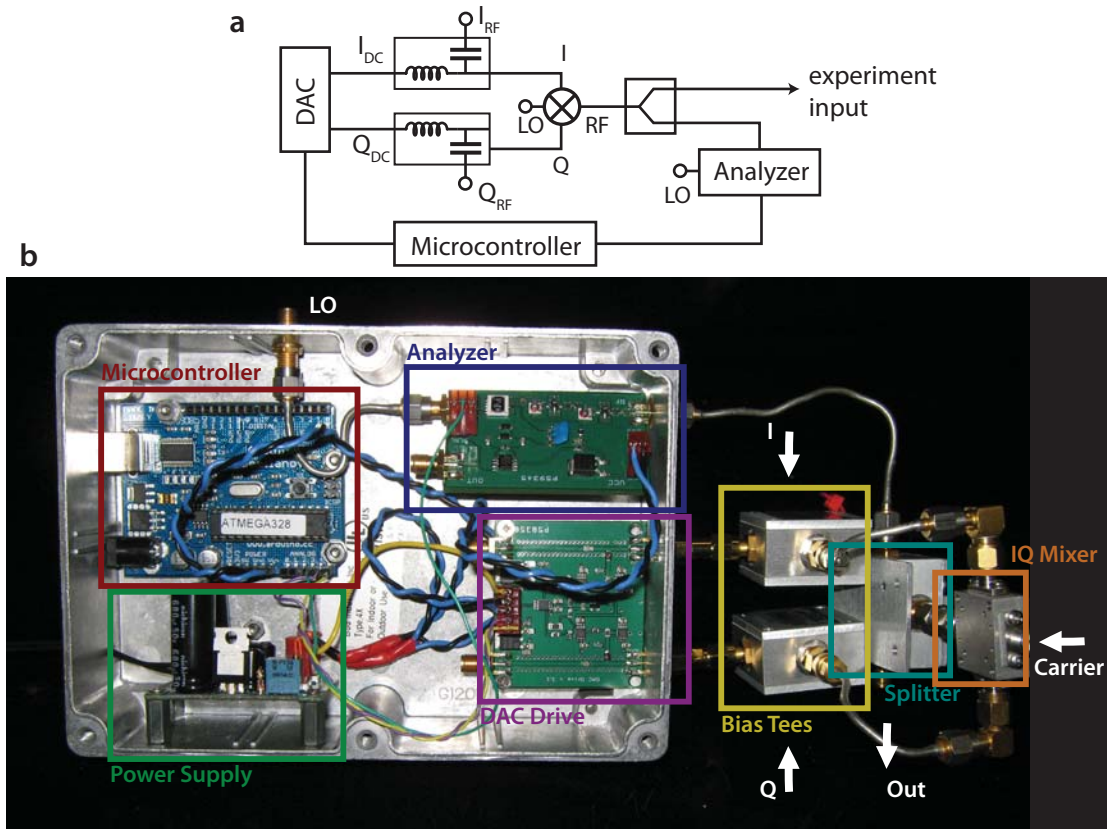Somewhat annoyingly, the I and Q voltages that optimally suppress the carrier leakage

Figure 5.8: Custom electronics for IQ mixer calibration. a Mixer calibrator schematic. **b** Photo of first-generation electronics. In this iteration, the on-board bias tees on the 'DAC Drive' board suffered from several resonances, so external Mini-Circuits bias tees are used instead. Future versions of the system will integrate the bias tees.

are not constant throughout the LO bandwidth of the mixer. In fact, the LO bandwidth over which the carrier leakage is below $-40\,\text{dB}$ is only $\sim 10\,\text{MHz}$*. Consequently, the I and Q offsets should be re-optimized every time the LO frequency is changed. The drift of the offsets is much better; I was able to operate for severals days at a time at the same LO frequency without observing a change in the carrier suppression of more than a few dB.

In an effort to automate and improve the carrier suppression process, I built some custom electronics together with Jerry Chow and Brian Vlastakis, shown in figure 5.8. The electronics offload the function of providing DC offset voltages onto a dedicated 'DAC drive' board that uses an *Analog Devices* digital-to-analog converter to provide two outputs from $\pm 125\,\text{mV}$ in

---

* I have not studied this systematically, but I make this claim off of similar measurements done by Nick Masluk with another IQ mixer. It is certainly the case that when the LO frequency is changed by more than 100 MHz that the optimal I and Q offsets change dramatically.
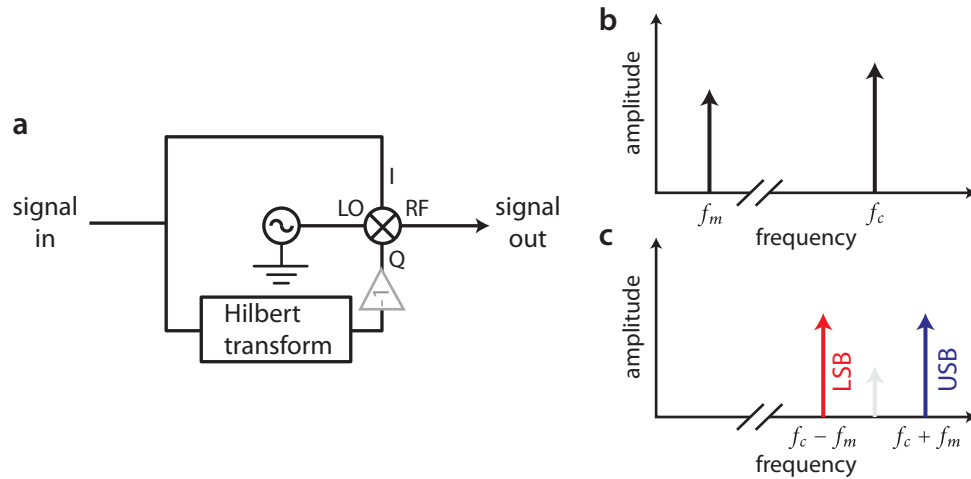
Figure 5.9: Single sideband modulation. a Schematic for a single sideband modulator. The signal is sent to the I port of an IQ mixer, while the Hilbert transform of the signal goes to the Q port. An optional inverter controls the sideband selection. **b** The input signal is at a frequency f_m and it is multiplied by a carrier of frequency f_c. **c** At the output of the SSB modulator, only the upper sideband (USB) or lower sideband (LSB) is present. The LSB is selected by adding the inverter between the Hilbert transformer and the Q channel of the mixer.

100 μV steps. Bias-tees with RF bandwidth 0.1–1000 MHz allow the addition of small DC voltages to the modulation pulses defining fast qubit rotations. A separate board serves as a basic spectrum analyzer to examine the leaked power with a 1 MHz bandwidth around the LO frequency. An *Arduino* microcontroller, based off the *Atmel* ATMEGA328, provides a USB interface to a computer, as well as feedback control to optimize the I and Q offsets by minimizing the leaked power with a Nelder-Mead downhill simplex search. This project is still a work in progress, but initial tests have already shown that this package can achieve carrier suppression of ~ 80 dB.

Single sideband modulation

Carrier suppression can also be mitigated by means of single-sideband modulation (SSB), which is like amplitude modulation (AM) but with the modulated signal either entirely above or below the carrier frequency. SSB allows the pulse output to be centered around a frequency that is shifted away from the LO carrier frequency, meaning that any residual carrier leakage is now off-resonance with the driven system. For sufficiently large detunings, this heavily suppresses the effects on a qubit system from the carrier. As shown in figure 5.9(a), SSB is implemented with a simple circuit, where a signal, $s(t)$, is directly multiplied by a carrier of

frequency f_c by the I port of an IQ mixer, while the Hilbert transform of s, $\hat{s}(t)$, is multiplied by a 90 degree phase-shifted carrier on the Q port. The result is that the signal is shifted up in frequency by f_c, which is equivalent to producing only the upper sideband (USB) in an AM system (see figure 5.9(b and c)). To produce the lower sideband (LSB), $\hat{s}(t)$ is multiplied by −1.

One can think of the Hilbert transform as applying a +90° phase shift to the negative frequency components of a signal, and a −90° phase shift to the positive frequency components. The additional 90° phase shift from the Q quadrature causes the positive frequency components to add in-phase while the negative frequency components cancel. Actually calculating the Hilbert transform of an arbitrary signal is a difficult problem numerically. However, we can use the fact that

$$H(\sin(\omega t)) = -\cos(\omega t), \tag{5.1}$$

$$H(\cos(\omega t)) = \sin(\omega t), \tag{5.2}$$

to approximately calculate $\hat{s}(t)$ when the signal takes the simple form

$$s(t) = \text{Re}[m(t)]\cos(\omega_m t) + \text{Im}[m(t)]\sin(\omega_m t), \tag{5.3}$$

where $m(t)$ is the pulse envelope phasor and ω_m is the desired frequency shift of the carrier. In fact, if $m(t)$ is slow compared to $\cos(\omega_m t)$, which is to say that the Fourier transform, $M(\omega)$, vanishes for $\omega \geq \omega_m$, then by Bedrosian's theorem [63],

$$\hat{s}(t) = \text{Re}[m(t)]\sin(\omega_m t) - \text{Im}[m(t)]\cos(\omega_m t). \tag{5.4}$$

When $M(\omega)$ has frequency content above ω_m, (5.4) is only approximately valid. When it does apply, the resulting I and Q channel signals are

$$I_\pm(t) = \text{Re}[m(t)]\cos(\omega_m t) \pm \text{Im}[m(t)]\sin(\omega_m t), \tag{5.5}$$

$$Q_\pm(t) = \text{Im}[m(t)]\cos(\omega_m t) \mp \text{Re}[m(t)]\sin(\omega_m t), \tag{5.6}$$

where $I_+(t)$ and $Q_+(t)$ are the signals for selecting the USB and the other pair select the LSB.

Pattern generation

Producing the baseband signals for an experiment involving many qubit operations and flux pulses, like the photon detection experiment discussed in chapter 8, is a non-trivial task. In

order to simplify this process, I extended the work started by Jerry Chow [47] and developed a domain-specific language for constructing pulse sequences from simple building blocks indicated by symbols, like 'X90p' for a $\pi/2$ rotation around $+X$. Example usage and a full code listing are presented in appendix B.

CHAPTER 6

Purcell Effect, Purcell Filter, and Qubit Reset

C OHERENCE poses the most important challenge for the development of a superconducting quantum computer. As the dephasing time T_2^* can never exceed twice the relaxation time T_1, it is the relaxation time which ultimately sets the limit on qubit coherence. Although T_2^* turned out to be small compared to T_1 in the earliest superconducting qubits [4], steady progress over the last decade has significantly reduced this gap [5, 64–67], such that it is now possible for qubits to exhibit homogenous broadening ($T_2^* \simeq 2T_1$) at parameter sweet-spots [67]. Therefore, understanding relaxation mechanisms is becoming critical to further improvements in both T_1 and T_2^*. This chapter will apply the classical admittance technique, presented in section 4.1, to understand relaxation in circuit QED experiments. The following section will show how the classical admittance can be modified to protect the qubit from relaxation with a 'Purcell filter', which was examined theoretically in section 4.1.2. Excellent agreement over a wide range in frequency is found between measured qubit relaxation times and the predictions of a circuit model for the filter. Using fast (nanosecond time-scale) flux biasing of the qubit, we demonstrate in-situ control of qubit lifetime over a factor of 50, allowing for implementation of qubit reset with 99.9% fidelity in 120 ns.

6.1   Multi-mode Purcell Effect

One of the main advantages of superconducting qubits is their strong interaction with the wires of an electrical circuit, making their integration with fast control and readout possible and allowing for large, controllable couplings between widely separated qubits [68]. The large coupling also implies a strong interaction between the qubits and their electromagnetic environment, which can lead to a short T_1. However, careful control of the coupling to the environment has been shown to allow prevention of circuit dissipation [26, 52]. Relaxation times have been studied in a wide variety of superconducting qubits, created with different fabrication techniques, and measured with a multitude of readout schemes. Typically, values of T_1 vary strongly from sample to sample as they can depend on many factors including materials, fabrication, and the design of both readout and control circuitry. In some instances a separation of these components has been achieved [69–72], but typically it is difficult to understand the limiting factors, and T_1 often varies strongly even among nominally identical qubit samples.

Here, we demonstrate that in a circuit QED architecture, where qubits are embedded in a microwave transmission line cavity[27, 64], transmon qubits have reproducible and understandable relaxation times. Due to the simple and well-controlled fabrication of the qubit and the surrounding circuitry, involving only two lithography layers and a single cavity for both control and readout, we are able to reliably understand and predict qubit lifetimes. This understanding extends to a wide variety of different qubit and cavity parameters. We find excellent agreement between theory and experiment for seven qubits over two orders of magnitude in relaxation time and more than an octave in frequency. The relaxation times are set by either spontaneous emission through the cavity, called the Purcell effect [73], or a shared intrinsic limit consistent with a lossy dielectric. Surprisingly, relaxation times are often limited by electromagnetic modes of the circuit which are far detuned from the qubit frequency. In the circuit QED implementation studied here, the infinite set of cavity harmonics reduces the Purcell protection of the qubit at frequencies above the cavity frequency.

Generally, any discrete-level system coupled to the continuum of modes of the electromagnetic field is subject to radiative decay. By placing an atom in a cavity, the rate of emission can be strongly enhanced [73, 74] or suppressed [75–77], depending on whether the cavity modes are resonant or off-resonant with the emitter's transition frequency. This effect is named after E. M. Purcell [73], who considered the effect of a resonant electrical circuit

on the lifetime of nuclear spins. Suppression of spontaneous emission provides effective protection from radiative qubit decay in the dispersive regime, where qubit and cavity are detuned [27]. Specifically, the Purcell rate for dispersive decay is given by $\gamma_\kappa = (g_{ge}/\Delta)^2 \kappa$, where g_{ge} denotes the coupling between the lowest transmon levels and cavity mode, Δ their mutual detuning, and κ the average photon loss rate.

The suppression and enhancement of decay rates can alternatively be calculated within a circuit model. For concreteness, we treat the system as a harmonic oscillator capacitively coupled to an arbitrary environment with impedance $Z_0(\omega)$, see figure 6.1(a). This approximation is appropriate because of the nearly harmonic level-structure of the transmon. The coupling circuit may be reduced to an effective dissipative element, as in figure 6.1(b). Specifically, replacing the coupling capacitor C_g and the environment impedance Z_0 by an effective resistor $R = 1/\operatorname{Re}[Y(\omega)]$, one finds[26, 52] that the T_1 is given by RC, where C is the qubit capacitance, as discussion in section 4.1. Choosing a purely resistive environment, $Z_0 = 50\,\Omega$, yields a decay rate $\gamma \simeq \omega^2 Z_0 C_g^2/C$. If instead we couple to a parallel LRC resonator, the calculated radiation rate can be reduced to that of the atomic case, $\gamma_\kappa = (g_{ge}/\Delta)^2 \kappa$, thus reproducing the Purcell effect.

The qualitative features of the Purcell effect are apparent in measurements of T_1, shown for 3 qubits in figure 6.2, measured with a dispersive readout by varying a delay time between qubit excitation and measurement [67, 78]. Near the cavity resonance at 5.2 GHz, spontaneous emission is Purcell-enhanced and T_1 is short. Away from resonance, the cavity protects the qubit from decay and the relaxation time is substantially longer than expected for decay into a continuum. However, at detunings above the cavity frequency, the measured T_1 deviates significantly from the single-mode Purcell prediction. This deviation can be directly attributed to the breakdown of the single-mode approximation.

The cavity does not just support a single electromagnetic mode, but also all higher harmonics of the fundamental mode. This has a striking impact on relaxation times. At first glance, it would appear that the effects of higher modes could be ignored when the qubit is close to the fundamental frequency and detuned from all higher modes. However, the coupling g_n to the n^{th} mode of the cavity increases with mode number, $g_n = g_0\sqrt{n+1}$. In addition, the input and output capacitors act as frequency-dependent mirrors, so that the decay rate of the n^{th} harmonic, $\kappa_n = (n+1)^2\kappa$, is larger than that of the fundamental. As a result, higher modes significantly contribute to the qubit decay rate, and the simple single-mode quantum model turns out to be inadequate for understanding the T_1 of the system.
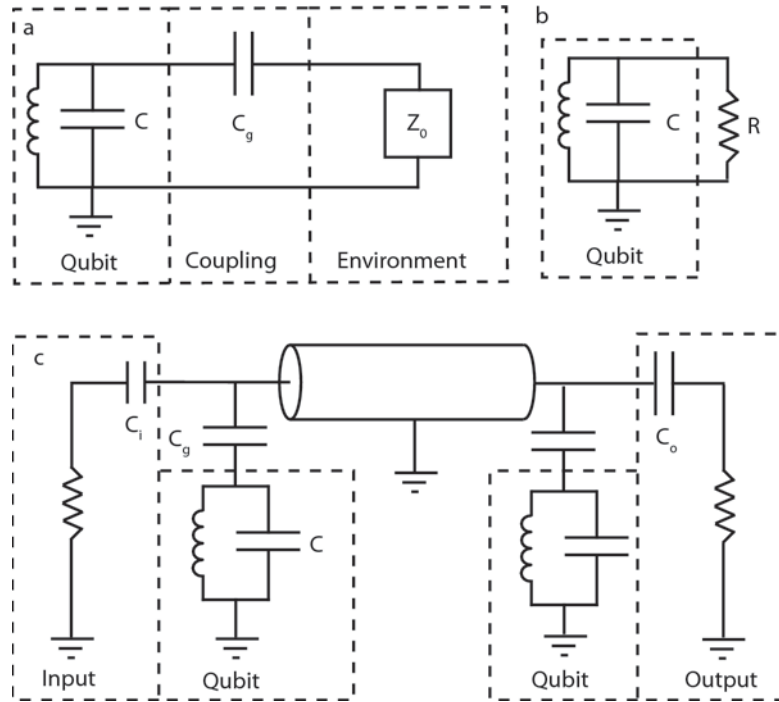
Figure 6.1: Circuit model of qubit relaxation. a Generalized model for a qubit coupled to an environment. **b** Reduced model of dissipation. The coupling capacitor and environment impedance are replaced by an effective resistance $R = 1/\operatorname{Re}[Y(\omega)]$, where $Y(\omega)$ is the admittance of the rest of the circuit seen by the qubit. The T_1 for the qubit is RC, where C is the qubit capacitance. **c** Full circuit diagram. Qubits are capacitively coupled to either end of a transmission line cavity. Both the input and output of the cavity are connected to a 50 Ω environment. The cavity is asymmetric in the sense that the input capacitance is smaller than the output capacitance.

Here, we follow an alternative route of calculating T_1, using the admittance formalism of section 4.1, and show that this accurately reproduces the measured T_1. The full calculation includes a transmission line cavity rather than a simple LRC resonator, see figure 6.1(c). The results from this are shown in figure 6.2, and reveal two striking differences as compared to the single-mode model: First, there is a strong asymmetry between relaxation times for qubit frequencies above (positive detuning) and below (negative detuning) the fundamental cavity frequency. While the single-mode model predicts identical relaxation times for corresponding positive and negative detunings, T_1 can be two orders of magnitude shorter for positive detunings than for negative detunings in the circuit model. Second, the circuit model shows a surprising dependence of T_1 on the qubit position in the cavity. While qubits located at opposite ends of the cavity have the same T_1 within the single-mode model, the
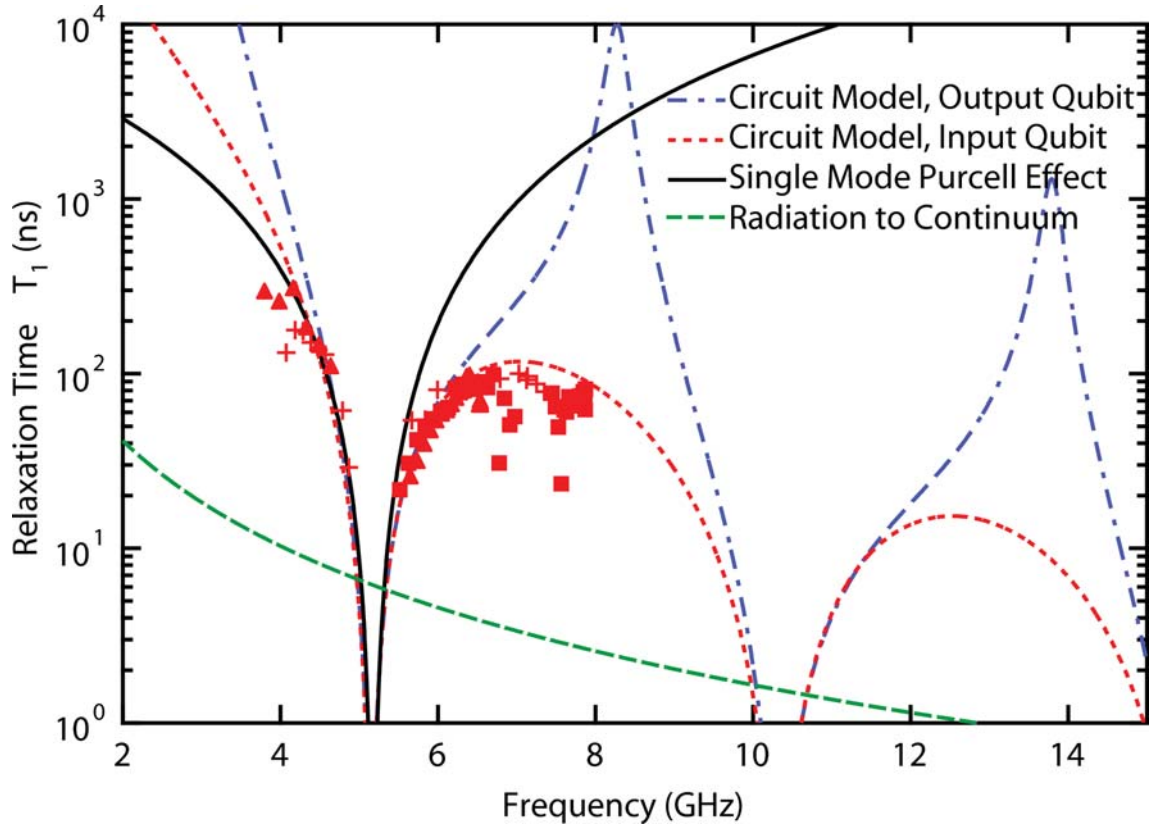
Figure 6.2: Comparison of circuit and single-mode models of relaxation. Spontaneous emission lifetimes into a single-mode cavity are symmetric about the cavity frequency, while within the circuit model lifetimes below the cavity are substantially longer than above. The measured T_1 for three similar qubits deviates substantially from the single-mode prediction, but agrees well with the circuit model. The expected decay time for radiation into a continuum is shown for comparison.

circuit model correctly captures the asymmetry induced by the differing input and output coupling capacitors and leads to vastly different T_1. The circuit model accurately resolves the discrepancy between the experimental data and the single-mode model, see figure 6.2.

The predictive power of the circuit model extends to all of our transmon qubits. Here, we present T_1 measurements on a representative selection of seven qubits. The qubits were fabricated on both oxidized high-resistivity silicon and corundum substrates, and coupled to microwave cavities with various decay rates and resonant frequencies. Table 1 provides parameters for each of the seven qubits. Qubits are fabricated via electron beam lithography and a double-angle evaporation process (25 nm and 80 nm layers of aluminum), while cavities

ID	Res.	ω_r (GHz)	κ (MHz)	g (MHz)	Pos.
1	Al on Si	5.17	44	107	In
2L	Al on Si	5.19	33	105	In
2R	Al on Si	5.19	33	105	Out
3L	Nb on Corun	6.69	40	166	In
3R	Nb on Corun	6.69	40	50	Out
4L	Nb on Corun	6.905	0.7	150	In
4R	Nb on Corun	6.905	0.7	55	Out

Table 6.1: Qubit parameters. Sample 1 is a single-qubit sample, all others are two-qubit samples. The Res. column indicates material and substrate for the cavity. The Pos. column indicates the position of the qubit at the input or output end of the cavity.

are fabricated by optical lithography with either lifted-off Al or dry-etched Nb on a Si or corundum substrate as discussed in Refs. [15, 79] and section 5.1.

Predictions from the circuit model are in excellent agreement with observed qubit lifetimes (see figure 6.3), up to a $Q \simeq 70,000$ for qubits on sapphire. The agreement is valid over more than two orders of magnitude in qubit lifetime and more than an octave of frequency variation. We emphasize that the circuit model does not correspond to a fit to the data, but rather constitutes a prediction based on the independently measured cavity parameters ω_r and κ, and the coupling g_{ge}.

In the qubits on silicon, coherence times of no more than 100 ns are observed above the cavity resonance, far below predictions from the single-mode model, but consistent with the circuit model. Initially, this caused concern for the transmon qubit: it appeared as if the transmon solved the $1/f$-noise dephasing problem for charge qubits, but introduced a new relaxation problem [9, 17, 18]. However, with the circuit model of relaxation, it is now clear that the 100 ns limit originated from the surprisingly large spontaneous emission rate due to higher cavity modes. By working at negative detunings instead, it is possible to achieve long relaxation times, here observed up to $4 \, \mu$s.

All qubits on corundum substrates reach a shared intrinsic limit of $Q \simeq 70,000$ when not otherwise Purcell limited. The constant-Q frequency dependence of the intrinsic limit ($T_1 \propto 1/\omega$) is suggestive of dielectric loss as the likely culprit. While the observed loss tangent $\tan \delta \sim 10^{-5}$ is worse than previous high-power measurements on corundum, it is not unreasonable depending on the type and density of surface dopants that might be present [57, 80]. The overall reproducibility of the intrinsic limit gives hope that future experiments may isolate its cause and reveal a solution. It is instructive to re-express the relaxation times
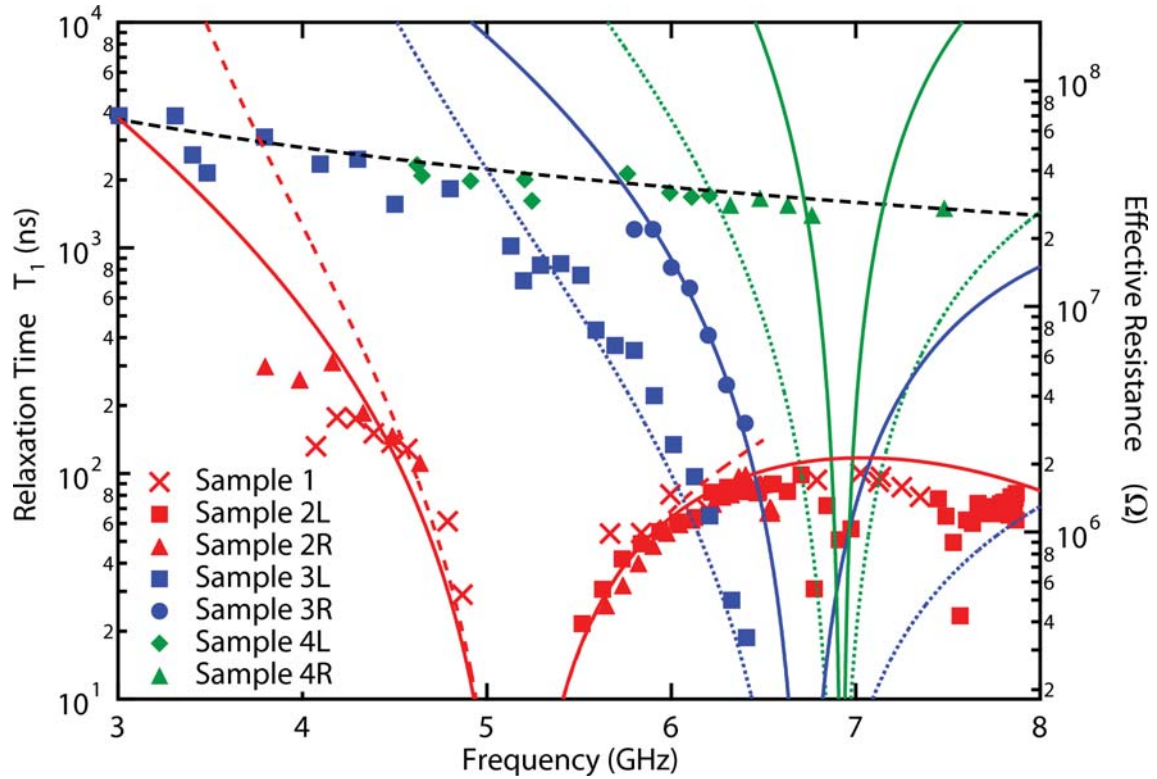
Figure 6.3: Relaxation times for seven superconducting qubits. Predictions for qubit lifetime based on the circuit model (colored lines) agree well with observed relaxation times (points). Solid lines represent predictions for input side (L) qubits, while dashed lines correspond to output side (R) qubits. All sapphire qubits (blue and green) reach the same common intrinsic limits (black line), with lifetimes limited to a constant $Q \sim 70,000$. Some deviation is seen in the lowest frequency silicon qubits, though it is unclear if this is an intrinsic limit. Qubit lifetimes are accurately predicted over a wide range of frequencies and more than two orders of magnitude in time.

in terms of a parasitic resistance, see figure 6.3. Note that here a T_1 of a microsecond roughly corresponds to a resistance of 20 MΩ. To build more complex circuits with still longer T_1, all dissipation due to parasitic couplings must be at the GΩ level.

Transmon qubits benefit greatly from the increased relaxation times, as they are insensitive to $1/f$-charge noise, the primary source of dephasing in other charge qubits. As a result, coherence is limited primarily by energy relaxation and transmons are nearly homogeneously broadened ($T_2^* \simeq 2T_1$). Improvements in T_1 thus translate directly into improvements in dephasing times T_2^*. This is demonstrated in figure 6.4, showing a comparison of relaxation and dephasing times. Here, T_2^* is measured in a pulsed Ramsey experiment and without
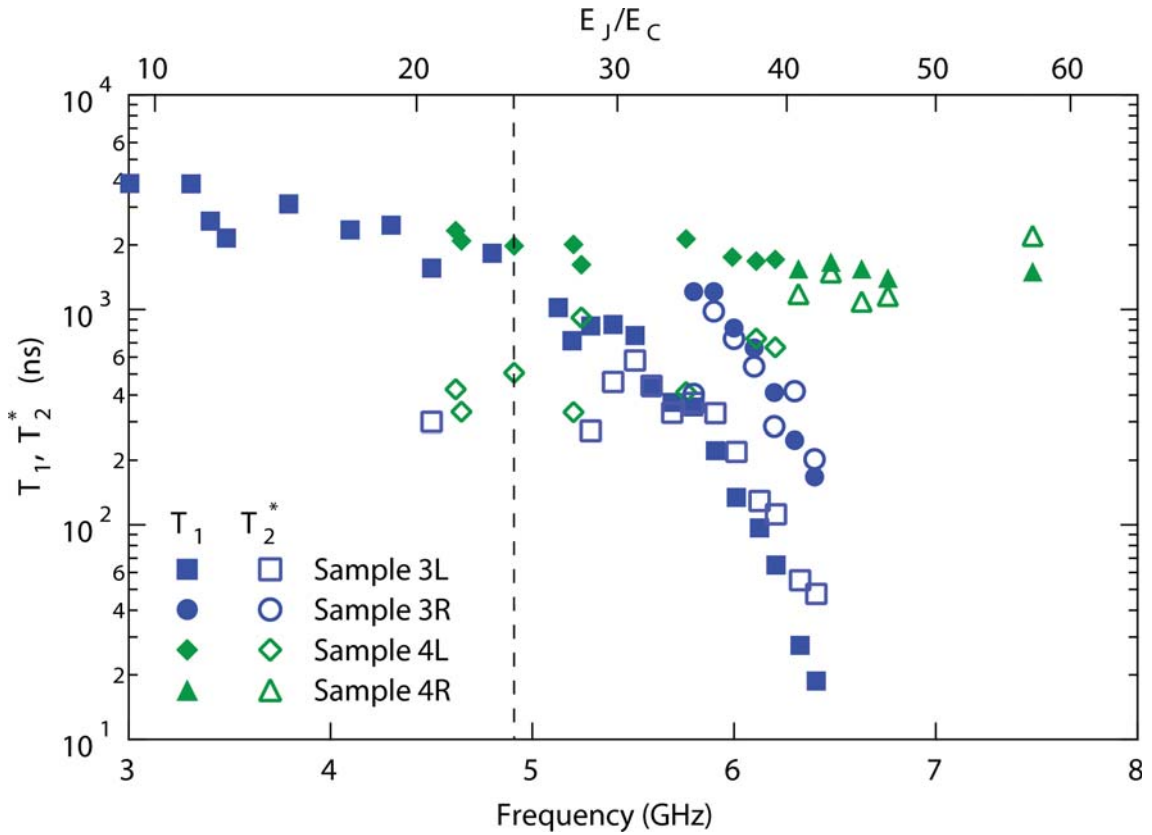
Figure 6.4: Dephasing times for four sapphire qubits. Measured dephasing times for each of the four sapphire qubits are nearly homogenously broadened, with T_2^* (open symbols) similar to T_1 (closed symbols) over a wide range of frequencies, even away from the flux sweet spot (the maximum frequency for each qubit). Charge noise, is suppressed exponentially in the ratio of Josephson to charging energies E_J/E_C (top axis), tuned along with qubit frequency (bottom axis) by changing an applied magnetic field. For small E_J/E_C charge noise dephasing is relevant and causes short T_2^*. Onset of significant charge noise is indicated by the dashed vertical line.

echo [67]. The gain in coherence time is most striking in samples with a higher-frequency cavity, $\omega_r/(2\pi) \sim 7$ GHz, where it is easier to operate at negative detunings and attain long T_1. In all these samples, we observe consistently long dephasing times of nearly a microsecond, with the largest T_2^* exceeding two microseconds without echo.

There are two main effects determining the observed dependence of T_2^* on the qubit frequency. First, away from the maximum frequency for each qubit, i.e. the flux sweet spot [5], the sensitivity to flux noise increases. This can cause additional inhomogeneous broadening. Despite this, T_2^* remains close to two microseconds, even away from the flux sweet spot. Second, tuning the qubit frequency via E_J directly affects the ratio of Josephson to charging energy, E_J/E_C, which dictates the sensitivity to charge noise. At low qubit frequencies, the qubits regain the charge sensitivity of the Cooper Pair Box, thus explaining the strong drop in dephasing times seen in figure 6.4.

Future improvements in T_2^* require further improvements in T_1. The accurate modeling of relaxation processes will be essential as quantum circuits become more complicated. In particular, the addition of multiple cavities and individual control lines may introduce accidental electromagnetic resonances. As we have shown here, even far off-resonant modes of a circuit can have a dramatic impact on qubit lifetimes. However, with careful circuit design, it should be possible not only to avoid additional accidental resonances, but to utilize the circuit model of relaxation to build filters to minimize dissipation. The next section presents one such method for effectively decoupling qubit relaxation from decay through the cavity.

6.2 Purcell Filter and Qubit Reset

In this section, we introduce a design element for cQED termed the 'Purcell filter', which protects a qubit from spontaneous emission while maintaining strong coupling to a low-Q cavity. We demonstrate an improvement of qubit T_1 by up to a factor of 50 compared to predicted values for an unfiltered device with the same $\kappa/2\pi \approx 20$ MHz. Combining the large dynamic range of almost two orders of magnitude in T_1 with fast flux control, we then demonstrate fast qubit reset to 99% (99.9%) fidelity in 80 ns (120 ns).

In cQED, qubits are generally sufficiently detuned to have suppressed relaxation, but T_1 can still be limited by decay through the cavity. As qubit lifetime is of paramount importance in quantum computing [56], a means of further inhibiting radiative decay is desirable.
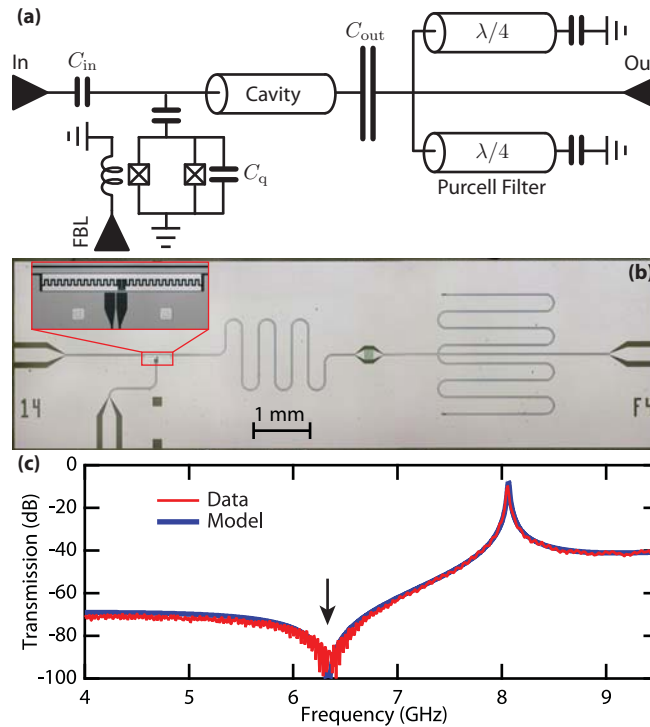
Figure 6.5: Design, realization, and diagnostic transmission data of the Purcell filter. **a** Circuit model of the Purcell-filtered cavity design. The Purcell filter, implemented with twin $\lambda/4$ open-circuited transmission-line stubs, inhibits decay through C_{out} near its resonance ω_f. **b** Optical micrograph of the device with inset zoom on transmon qubit. Note the correspondence of the circuit elements directly above in **a**. **c** Cavity transmission measured at 4.2 K and comparison to the circuit-model prediction. The Purcell filter shorts out the 50 Ω output environment at ω_f, yielding a 30 dB drop in transmission (arrow). A circuit model involving only the parameters C_{in}, C_{out}, ω_c, and ω_f shows excellent correspondence.

The Purcell decay rate can be significantly reduced by increasing either the cavity quality factor Q or the detuning Δ between the qubit (ω_q) and cavity (ω_c) frequencies, but these solutions have unwelcome implications of their own. For example, reducing the cavity decay rate $\kappa = \omega_c/Q$ can diminish qubit readout fidelity [81] because fewer signal photons are collected in a qubit lifetime. A large κ is also beneficial for resetting a qubit to its ground state by bringing it near to the cavity resonance and exploiting the Purcell-enhanced decay rate. Increasing Δ similarly has adverse effects on readout fidelity and applications that exploit large state-dependent frequency shifts [10, 17, 61]. A better solution would improve qubit T_1 independent of the cavity Q, leaving its optimization up to other experimental concerns.

The filter works by exploiting the fact that the qubit and cavity are typically far detuned.

We can therefore modify the qubit's electromagnetic environment (e.g. the density of photon states at ω_q) without, in principle, affecting the cavity Q or resonant transmission. The relationship between qubit T_1 due to spontaneous emission and admittance Y of the coupled environment is

$$T_1^{\text{Purcell}} = \frac{C}{\text{Re}[Y(\omega_q)]}. \tag{6.1}$$

The relevant admittance network with the filter is shown in figure 6.5(a). The previous section demonstrated that (6.1) accurately models the observed T_1^{Purcell} when all modes of the cavity are taken into account in the calculation of Y. As the relationship holds for any admittance, this decay rate can be controlled by adjusting Y with conventional microwave engineering techniques. In particular, by manipulating Y to be purely reactive (imaginary-valued) at ω_q, T_1^{Purcell} diverges and the Purcell decay channel is turned off. This solution decouples the choice of cavity Q from the Purcell decay rate as desired, and, as we will see, has the advantage of using only conventional circuit elements placed in an experimentally convenient location.

We implement the Purcell filter with a transmission-line stub terminated in an open circuit placed outside the output capacitor C_{out} [figure 6.5(a)]. The length of this stub is set such that it acts as a $\lambda/4$ impedance transformer to short out the 50 Ω environment at its resonance frequency ω_f. We choose C_{out} to be much larger than the input capacitor, $C_{\text{in}} \approx C_{\text{out}}/15$, to ensure that the qubit would be overwhelmingly likely to decay through C_{out}. The Purcell filter eliminates decay through this channel, leaving only the negligible decay rate through C_{in}. The combined total capacitance $C_{\text{tot}} \approx 80$ fF results in a small cavity Q. We use two identical stubs above and below the major axis of the chip [Fig. 6.5(b)] to keep the design symmetric in an effort to suppress any undesired on-chip modes. The cavity resonates at $\omega_c/2\pi = 8.04$ GHz, the filter at $\omega_f/2\pi = 6.33$ GHz, and a flux bias line (FBL) is used to address a single transmon qubit [67] with a maximum frequency of 9.8 GHz, a charging energy $E_C/2\pi$ of 350 MHz, and a resonator coupling strength $g_{ge}/2\pi$ of 270 MHz. Transmission through the cavity measured at 4.2 K was compared with our model to validate the microwave characteristics of the device [figure 6.5(c)]. There is a dip corresponding to inhibited decay through C_{out} at ω_f. The predicted and measured curves are also qualitatively similar, lending credence to the circuit model. This method provided a convenient validation before cooling the device to 25 mK in a dilution refrigerator.

We measured the qubit T_1 as function of frequency and found it to be in excellent agreement with expectations. T_1 is well modeled by the sum of the Purcell rate predicted by our
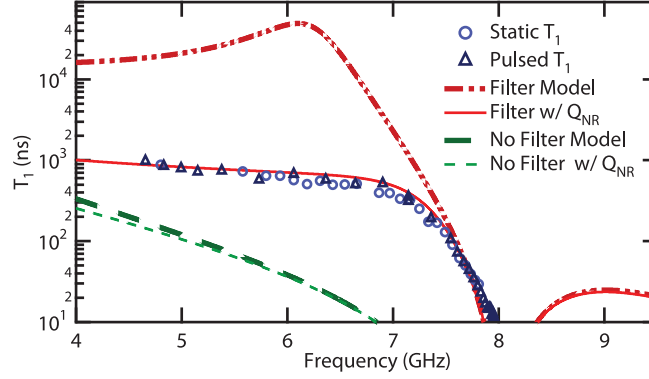
Figure 6.6: Qubit T_1 as a function of frequency measured with two methods, and comparison to various models. The first method is a static measurement (circles): the qubit is excited and measured after a wait time τ. The second (triangles) is a dynamic measurement: the qubit frequency is tuned with a fast flux pulse to an interrogation frequency, excited, and allowed to decay for τ, and then returned to its operating frequency of 5.16 GHz and measured. This method allows for accurate measurement even when T_1 is extremely short. Measurements using the two methods show near perfect overlap. The top dashed curve is the predicted T_1^{Purcell}, while solid curve includes also non-radiative internal loss with best-fit $Q_{\mathrm{NR}} = 2\pi f\, T_1^{\mathrm{NR}} \approx 27,000$. The two lower curves correspond to an unfiltered device with the same C_{in}, C_{out}, and ω_{c}, with and without the internal loss. In this case, the Purcell filter gives a T_1 improvement by up to a factor of ~50 (6.7 GHz).

filtered circuit model and a non-radiative internal loss $Q_{\mathrm{NR}} \approx 27,000$ (Fig. 6.6). The source of this loss is a topic of continuing research, though some candidates are surface two level systems [58, 82], dielectric loss of the tunnel barrier oxide [70] or corundum substrate, and non-equilibrium quasiparticles [83]. It is also unknown why this sample followed a lower Q_{NR} value than the group of samples studied in the last section. Nonetheless, this model contains only the fit parameter Q_{NR} combined with the independently measured values of g_{ge}, E_{C}, ω_{c}, ω_{f}, C_{in}, and C_{out}. An improvement to T_1 due to the Purcell filter was found to be as much as a factor of 50 at 6.7 GHz by comparison to an unfiltered circuit model with the same parameters. This would be much greater in the absence of Q_{NR}. The device also exhibits a large dynamic range in T_1—about a factor of 80 between the longest and shortest times measured.

This range in T_1 can be a challenge to quantify because measurements made at small detunings, where T_1 is a few tens of nanoseconds, have a very low SNR. This issue was avoided through the use of fast flux control. For measurements at small Δ, the qubit is pulsed to the detuning under scrutiny, excited and allowed to decay, then pulsed to 5.16 GHz where
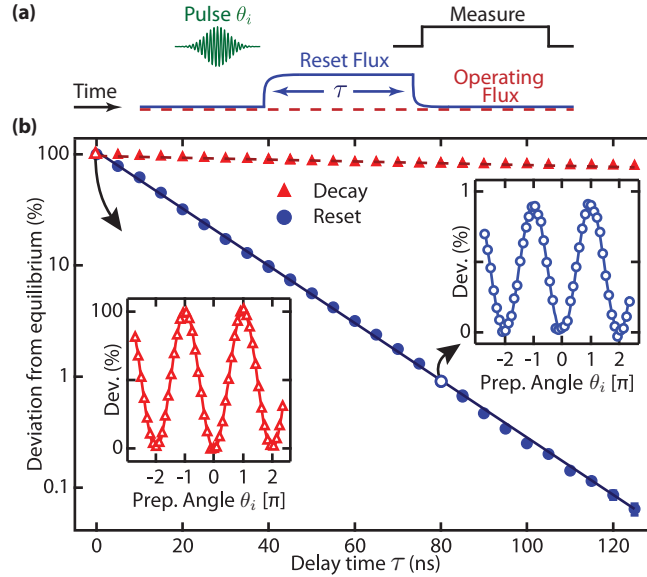
Figure 6.7: Fast qubit reset. a Schematic of a pulse sequence used to realize a qubit reset and characterize its performance. The fidelity of reset was quantified using a modified Rabi oscillation scheme. The qubit is first rotated around the x-axis by an angle θ at the operating frequency of 5.16 GHz and then pulsed into near resonance with the cavity (solid line) or left at the operating frequency (dashed line) for a time τ. The state of the qubit is measured as a function of θ and τ after being pulsed back to 5.16 GHz. **b** The Rabi-oscillation amplitude as a function of τ, normalized to the amplitude for τ=0. This ratio gives the deviation of the qubit state from equilibrium. Curves are fit to exponentials with decay constants of 16.9 ± 0.1 ns and 540 ± 20 ns respectively. Insets: Measured Rabi oscillations for τ=0 (lower left) and τ=80 ns (top right). The vertical scales differ by a factor of 100.

measurement fidelity is higher, and interrogated. In the cases where the qubit is nearly in resonance with the cavity, the T_1 is actually so short that it constitutes an interesting resource.

The ability to reset, or quickly cool a qubit to its equilibrium state on demand, is an important capability with a diverse set of applications. Using a qubit to make repeated measurements of a coupled system, for example, requires resetting the qubit between interrogations [84]. Similarly, experiment repetition rates can be greatly enhanced when they are otherwise limited by T_1. Fast reset is also vital for measurement-free quantum error correction [85]. In this scheme, an error syndrome is encoded in two ancilla qubits and conditionally corrected using a three qubit gate. The ancillas, which now hold the entropy associated with the error, are then reset and reused. The Purcell filter is an ideal element with which to demonstrate reset as it allows for a relatively short reset time through the use of a low-Q cavity without limiting T_1 at the operating frequency.

The efficacy of reset in this device is readily quantified using a modified Rabi oscillation sequence, described in figure 6.7(a). Each experiment measures the degree to which the qubit is out of equilibrium after some reset time τ; the protocol is insensitive to any equilibrium thermal population of the qubit. The non-equilibrium population is found to exhibit pure exponential decay over three orders of magnitude. The qubit can be reset to 99.9% in 120 ns or any other fidelity depending on τ. The sequence is also performed with the qubit remaining in the operating frequency during the delay to demonstrate the large dynamic range in T_1 available in this system. In the case of multi-qubit devices, it is possible that this reset process would affect other qubits coupled to the same bus, but this issue could be avoided by using separate coupling and reset cavities.

6.3   Chapter Summary

This chapter shows that the classical admittance formalism is an effective means to understanding relaxation in real devices. It also shows how effective modeling of a circuit admittance can lead to the development of ways to manipulate that admittance to improve relaxation. In particular, we showed how the Purcell filter allows for the use of low-Q cavities without adversely affecting qubit T_1. This ability is well-suited for in-situ qubit reset, a prerequisite for measurement-free quantum error correction and other applications. Hopefully, the concise understanding of spontaneous emission lifetimes in our system, and the reproducibility of intrinsic lifetimes will open pathways for a systematic exploration of limits on relaxation and coherence.

Measurement of the Self- and Cross-Kerr Effects in Two-Cavity Circuit QED

Tʜɪs chapter presents preliminary investigations into the presence of self-Kerr and cross-Kerr effects in two-cavity circuit QED, which were examined theoretically in section 3.5. Observation of these effects presents strong evidence that the experimental two-cavity system is well described by the Hamiltonian of (3.1). The device used for these measurements, shown in figure 7.1, has cavities at frequencies $\omega_1/2\pi$ = 6.473 GHz and $\omega_2/2\pi$ = 6.937 GHz, with associated decay rates $\kappa_1/2\pi$ = 0.71 MHz and $\kappa_2/2\pi$ = 2.89 MHz. Both cavities are coupled to a sarantapede qubit, with $E_C/h \approx 275$ MHz. Unfortunately, the maximum E_J for the device was such that the qubit did not cross either cavity. The maximum qubit frequency, as seen in the spectroscopy map of figure 7.2(a), was $\omega_{ge}/2\pi$ = 5.62 GHz, meaning that it was impossible to directly measure g_1 or g_2 from the vacuum Rabi splittings with each cavity. Consequently, other methods were used to infer g_1 and g_2.

Near the qubit's maximum frequency, pulsed spectroscopy displayed number splitting [17] when cavity 1 was populated with a coherent state, as shown in figure 7.2(b and c). One can use this spectrum to find g_1. Precise fitting of number splitting data is notoriously tricky, in practice requiring slow numerical simulations (see refs. [33] and [17]). However, the distance between peaks, which is much easier to extract, is $2\chi_1 + \delta_1$, where $\delta_1 = \omega_d - \omega_1$ is the detuning between the populating drive and the cavity frequency. Since $\chi_1 = g_1\alpha/\Delta_1(\Delta_1 + \alpha)$, where
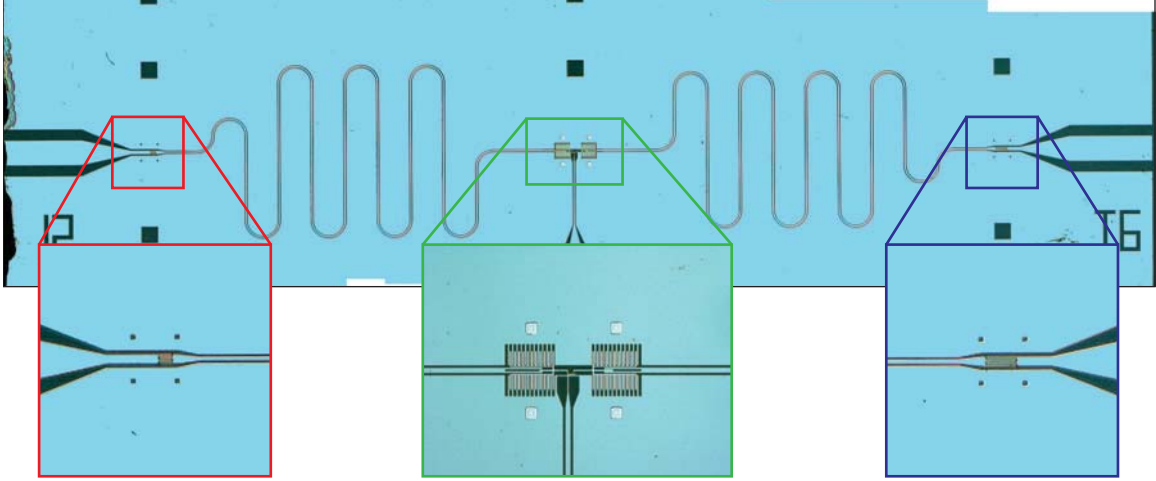
Figure 7.1: Optical image of cQED274. The chip has two CPW cavities with frequencies $\omega_1/2\pi = 6.473$ GHz and $\omega_2/2\pi = 6.937$ GHz. Coupling capacitors controlling the decay rates $\kappa_1/2\pi = 0.71$ MHz and $\kappa_2/2\pi = 2.89$ MHz are shown in red and blue. A transmon qubit, shown in green, is coupled to each cavity with rates g_1 and g_2.

$\Delta_1 = \omega_q - \omega_1$ is the qubit-cavity detuning and α is the transmon anharmonicity, knowledge of χ_1 allows one to find g_1. When driving at $\omega_d/2\pi = 6.4775$ GHz, the distance between peaks is $\chi_1 + \delta_1 \simeq 5$ MHz. I estimate the drive-cavity detuning, δ_1, from reflection measurements of the cavity with the transmon in the ground state, giving $\delta_1/2\pi = 1$ MHz. This implies that $g_1/2\pi = 82 \pm 5$ MHz. The estimated error in g_1 is determined by uncertainty in the precise values of δ_1 and χ_1.

At no detuning was the qubit close enough to the higher frequency cavity to display number splitting when cavity 2 was populated. However, χ_2 is also directly observable in the qubit-state-dependent cavity shift, which can be observed in a time-domain experiment as done in [61]. In this type of measurement, the qubit is prepared in $|e\rangle$ and the reflected cavity response is scanned to find the shifted cavity frequency. With this method and the qubit at its maximum frequency, I measured $\chi_2/\pi \simeq 2.1$ MHz, implying that $g_2/2\pi = 85 \pm 5$ MHz. Both inferred coupling strengths are within 10% of the values calculated using Maxwell simulations of the capacitance matrix and (3.21).

The Kerr effect was measured for three different qubit frequencies $\omega_{ge}/2\pi = (4.84, 5.08, 5.61)$ GHz, and thus three different detunings to each cavity, Δ_1 and Δ_2. For each choice of qubit frequency, the response of cavity 1 was probed in reflection with a weak tone ($\langle n \rangle < 1$) as a function of the power of a second populating tone at $\omega_d \simeq \omega_1$ or $\omega_d \simeq \omega_2$. Traces for
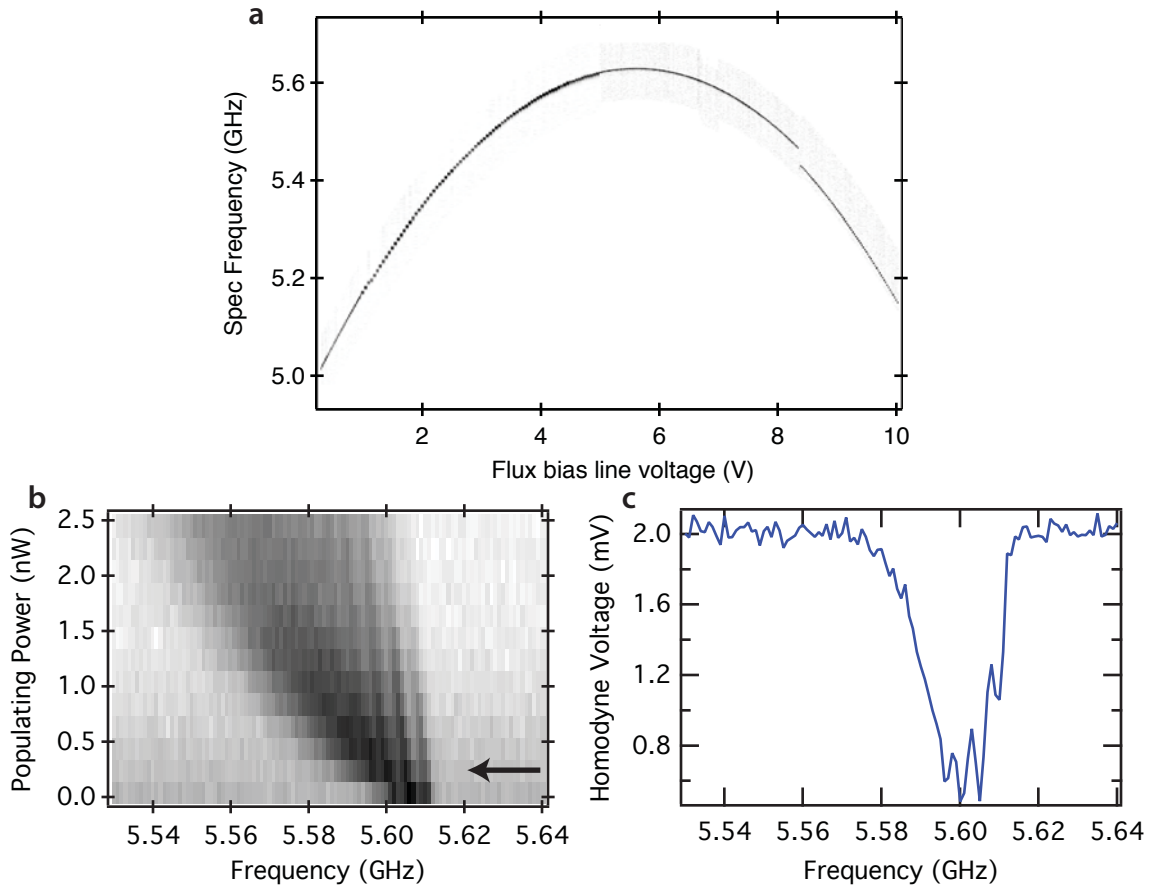
Figure 7.2: Spectroscopy of cQED274. a Tracked qubit spectroscopy vs. voltage applied to the
flux bias line. The color scale shows the reflected heterodyne phase. The discontinuity near
8.2 V is a flux jump which is not repeatable. **b** Near the maximum transmon frequency, pulsed
spectroscopy vs. drive power at $\omega_d/2\pi = 6.4775$ GHz that populates cavity 1. **c** A cut of **b** at
drive power of 0.25 nW (black arrow), showing a number split spectrum with ~ 5 MHz peak
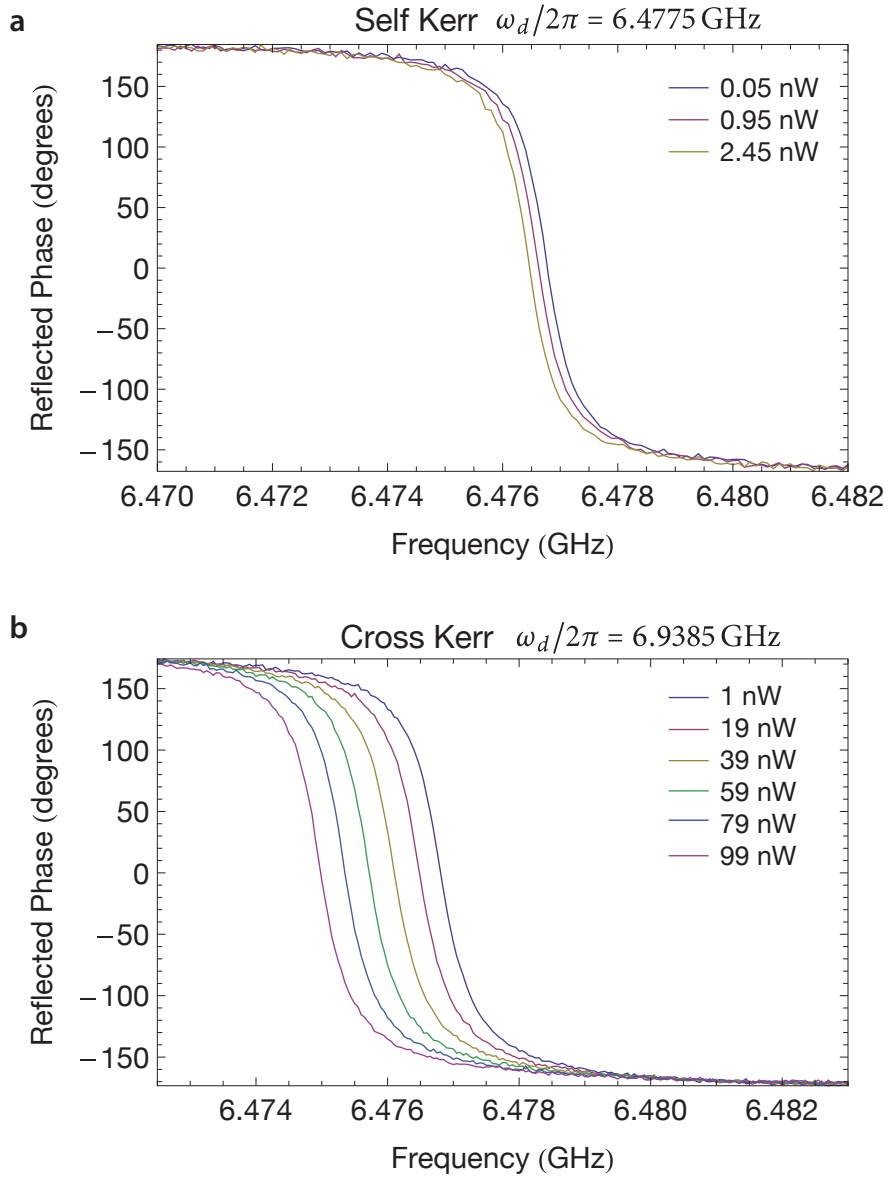separation.

Figure 7.3: Cavity Response vs Self- and Cross-Power. a Reflected cavity phase response vs
frequency for various drive powers at $\omega_d/2\pi = 6.4775$ GHz. **b** Reflected cavity phase response
vs frequency for various drive powers at $\omega_d/2\pi = 6.9385$ GHz. In both **a** and **b**, the center
frequency of the response shifts down to lower frequencies as the drive power is increased. The
S-shaped curve with a ~ 360 degree excursion is typical of an over-coupled harmonic oscillator.
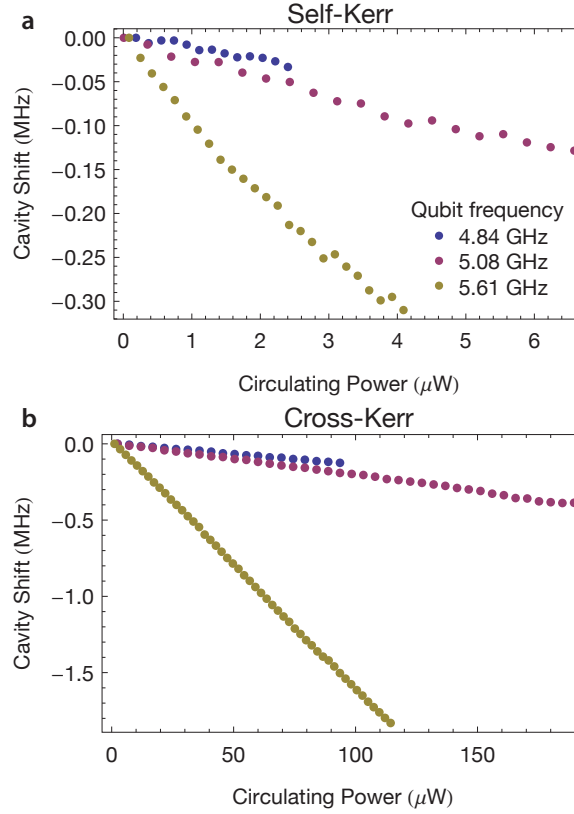
Figure 7.4: Self- and Cross-Kerr Effect. a, b Shift of cavity center frequency vs. circulating power when the qubit is at 4.87 (blue), 5.08 (magenta) or 5.61 GHz (gold). In **a**, the drive frequency is $\omega_d = (6.4735, 6.474, 6.4775)$ GHz, whereas in **b** the drive is $\omega_d = (6.937, 6.937, 6.9385)$ GHz. A linear relationship is evident between the cavity shift and the circulating powers, a clear signature of the self- and cross-Kerr effects.

$\omega_{ge}/2\pi = 5.61$ GHz are shown in figure 7.3. For all displayed choices of populating power, the response shows a 360 degree phase shift as the signal frequency passes through resonance, consistent with an over-coupled harmonic oscillator. Each trace can be fit to the function [15]

$$\arg\left[S_{11}(\omega)\right] = 2\arctan(\omega - \omega_0, \kappa/2), \tag{7.1}$$

to extract the cavity resonance frequency ω_0 and linewidth κ, where the factor of 2 comes from the additional 180 degree phase shift imparted by reflection, and the two argument arctangent gives $\arctan(x, y) = \arctan(y/x)$, taking into account the quadrant of the point (x, y).

A plot of the center frequency shifts of cavity 1 as a function of the circulating power in
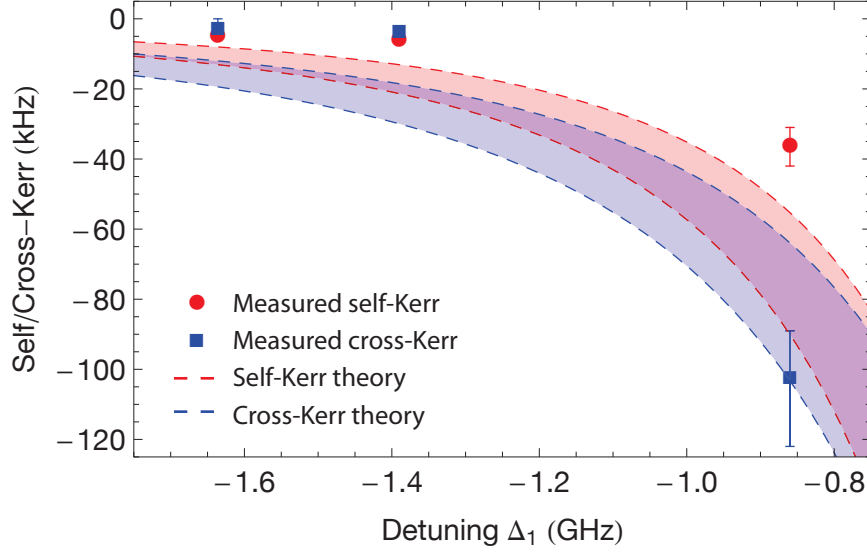
Figure 7.5: Comparison of measured self- and cross-Kerr effect with theory. Measured values of the self- and cross-Kerr are inferred from figure 7.4 by using the AC-Stark shift to calibrate the photon number in each cavity. The error bars are determined by the uncertainties in the photon number calibration. Theory bands are plotted using (3.24) and (3.30), where the upper and lower bounds come from the uncertainty in g_1 and g_2. The theory predicts the correct sign and order of magnitude of the effects.

cavities 1 and 2 is shown in figure 7.4. For an over-coupled cavity, the circulating power is solely determined by the loaded cavity Q and a Lorentzian filter factor from the detuning of the drive from the cavity resonance frequency, *i.e.*

$$P_{\text{cir}} = \frac{QP_{\text{in}}}{1 + \delta^2/(\kappa/2)^2}, \tag{7.2}$$

where $\delta = \omega_d - \omega_0$, and Q, κ, and ω_0 refer to the relevant cavity. It is readily apparent, both when populating at $\omega_d \simeq \omega_1$ and ω_2, that the resonance frequency is shifted down proportionally to the populating drive power. In order to compare the Kerr effects with the theoretical predictions, one should rescale the data of figure 7.4 in terms of the average cavity populations $\langle n_1 \rangle$ and $\langle n_2 \rangle$. For this, the best tool to use is the AC-Stark effect. As described in detail in [15, 32, 33], for small average photon numbers and $\chi \ll \gamma$, where γ is the qubit linewidth, the qubit frequency depends linearly on the photon number with $\Delta\omega_{ge} = 2\chi\bar{n} = 2\chi P_{\text{in}}/P_{\text{1photon}}$. Consequently, knowledge of χ allows one to extract $P_{\text{1photon}} = \hbar\omega_0\kappa$ by measuring the qubit frequency shift as a function of P_{in}.

Using the extracted g_i and fitting the cavity shifts to a line, I obtain the self- and cross-

Kerr effects reported in figure 7.5. Uncertainty in the extracted device parameters prevents a detailed numerical comparison between data and theory; nonetheless, these initial measurements show the self- and cross-Kerr effects to be of the right sign and order of magnitude as the predictions of (3.24) and (3.30).

This preliminary experiment serves to demonstrate some basic features of the two-cavity architecture. In particular, all the single-cavity circuit QED physics is still present, including dispersive qubit readout, the AC-Stark effect, and number splitting. The 'sarantapede' geometry also functions as a standard transmon that is coupled to each cavity. Furthermore, these initial measurements of the cross-Kerr effect give credence to the Hamiltonian (3.1) and open the door to using cross-Kerr for photon detection, a method which is discussed in chapter 9. The next chapter uses the two-cavity architecture to demonstrate a different photon detection method that is based on qubit-photon quantum logic.

Quantum Non-Demolition Detection of Single Microwave Photons in a Circuit

T HOROUGH control of quantum measurement is key to the development of quantum information technologies. Many measurements are destructive, removing more information from the system than they obtain. Quantum non-demolition (QND) measurements allow repeated measurements that give the same eigenvalue [35]. They could be used for several quantum information processing tasks such as error correction [86], preparation by measurement [87], and one-way quantum computing [88]. Achieving QND measurements of photons is especially challenging because the detector must be completely transparent to the photons while still acquiring information about them [14, 33]. Recent progress in manipulating microwave photons in superconducting circuits [18, 24, 31] has increased demand for a QND detector which operates in the gigahertz frequency range. Here we demonstrate a QND detection scheme which measures the number of photons inside a high quality-factor microwave cavity on a chip. This scheme maps a photon number, n, onto a qubit state in a single-shot via qubit-photon logic gates. We verify the operation of the device for $n = 0$ and 1 by analyzing the average correlations of repeated measurements, and show that it is 90% QND. It differs from previously reported detectors [14, 17, 24, 29, 31] because its sensitivity is strongly selective to chosen photon number states. This scheme could be used to read-out the state of a photon-based memory in a quantum computer.

Several teams have engineered detectors which are sensitive to single microwave photons by strongly coupling atoms (or qubits) to high-Q cavities. This cQED architecture can be used in various ways to detect photons. One destructive method measures quantum Rabi oscillations of an atom or qubit *resonantly* coupled to the cavity [24, 29, 31]. The oscillation frequency is proportional to $\sqrt{n}$, where n is the number of photons in the cavity, so this method essentially measures the time-domain swap frequency.

Another method uses a *dispersive* interaction to map the photon number in the cavity onto the phase difference of a superposition of atomic states $(|g\rangle + e^{i\phi}|e\rangle)/\sqrt{2}$. Each photon number n corresponds to a different phase ϕ, so repeated Ramsey experiments [14] can be used to estimate the phase and extract n. This method is QND, because it does not exchange energy between the atom and photon. However, since the phase cannot be measured in a single operation, it does not extract full information about a particular Fock state $|n\rangle$ in a single interrogation. Nonetheless, using Rydberg atoms in cavity QED, remarkable experiments have shown quantum jumps of light and the collapse of the photon number by measurement. [14, 44]

Here we report a new method which implements a set of programmable controlled-NOT (CNOT) operations between an n-photon Fock state and a qubit, asking the question "are there exactly n photons in the cavity?" A single interrogation consists of applying one such CNOT operation and reading-out the resulting qubit state. To do this we use a quasi-dispersive qubit-photon interaction which causes the qubit transition frequency to depend strongly on the number of photons in the cavity. Consequently, frequency control of a pulse implements a conditional π rotation on the qubit – the qubit state is inverted if and only if there are n photons in the storage cavity. To ensure that this is QND, the qubit and storage cavity are adiabatically decoupled before performing a measurement of the qubit state.

To realize this method we extend circuit-based cavity QED [68] by coupling a single transmon qubit [6, 67] simultaneously to *two* cavities. In this scheme, the transmon qubit is used to interrogate the state of one cavity, then the transmon is measured with the other cavity. This separation of functions allows one cavity to be optimized for coherent storage of photons (high-Q) and the other for fast qubit readout (low-Q). Related work by Leek *et al.* [25] realized a single transmon coupled to two modes of a single cavity, where the two modes were engineered to have very different quality factors. A schematic of the two-cavity device is shown in Fig. 8.1(a). The cavities are realized as Nb coplanar waveguide resonators with $\lambda/2$ resonances at $\omega_s/2\pi = 5.07$ GHz and $\omega_m/2\pi = 6.65$ GHz, respectively. The cavities are engineered, by design of the capacitors C_s and C_m, to have very different decay rates
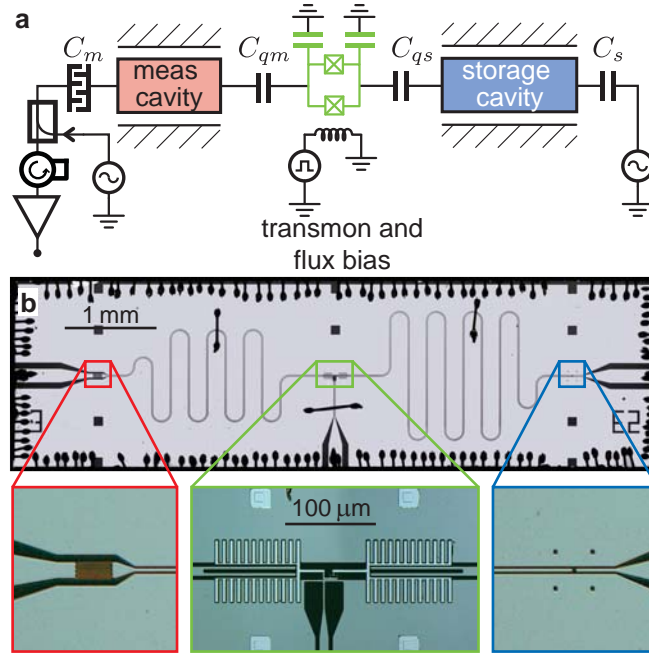
Figure 8.1: Circuit schematic and cQED291 device. a, Circuit schematic showing two cavities coupled to a single transmon qubit. The measurement cavity is probed in reflection by sending microwave signals through the weakly coupled port of a directional coupler. A flux bias line allows for tuning of the qubit frequency on nanosecond timescales. **b**, Implementation on a chip, with $\omega_m/2\pi = 6.65\,\text{GHz}$ measurement cavity on the left and its large coupling capacitor (red), and $\omega_s/2\pi = 5.07\,\text{GHz}$ storage cavity on the right with a much smaller coupling capacitor (blue). A transmon qubit (green) is strongly coupled to each cavity, with $g_s/2\pi = 70\,\text{MHz}$ and $g_m/2\pi = 83\,\text{MHz}$. It has a charging energy $E_C/2\pi = 290\,\text{MHz}$ and maximal Josephson energy $E_J/2\pi \approx 23\,\text{GHz}$. At large detunings from both cavities, the qubit coherence times are $T_1 \approx T_2 \approx 0.7\,\mu\text{s}$.

($\kappa_s/2\pi = 50\,\text{kHz}$ and $\kappa_m/2\pi = 20\,\text{MHz}$) so that the qubit state can be measured several times per photon lifetime in the storage cavity. By having these cavities at different frequencies, the fast decay of the readout cavity does not adversely affect photons in the storage cavity. A transmon qubit is end-coupled to the two cavities, with finger capacitors controlling the individual coupling strengths ($g_s/2\pi = 70\,\text{MHz}$ and $g_m/2\pi = 83\,\text{MHz}$). The usual shunt capacitor between the transmon islands is replaced with capacitors to the ground planes to reduce direct coupling between the cavities. Additionally, a flux bias line [10] allows fast, local control of the magnetic field near the transmon. This facilitates manipulations of the detunings $\Delta_s = \omega_{g,e} - \omega_s$ and $\Delta_m = \omega_{g,e} - \omega_m$ between the transmon and cavities, where we use the convention of labeling the transmon states from lowest to highest energy as (g, e, f, h, ...).
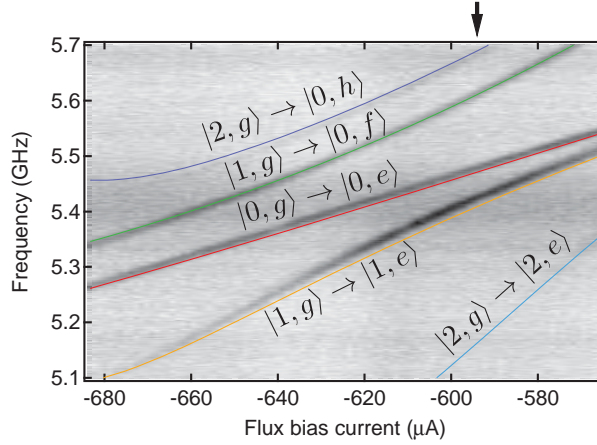
Figure 8.2:   Pulsed spectroscopy with coherent state in storage cavity ($\langle n \rangle \approx 1$) vs. qubit-cavity detuning $\Delta_s = \omega_{g,e} - \omega_s$. Calculated transition frequencies are overlaid in color. Red and orange lines are the $|g\rangle \leftrightarrow |e\rangle$ transitions of the qubit when $n = 0$ and 1, respectively. Transitions to higher transmon levels ($|f\rangle$ and $|h\rangle$) are visible because of the small detuning. The arrow indicates the flux bias current used during the CNOT operations.

To achieve high photon number selectivity of the CNOT operations, there must be a large separation between the number-dependent qubit transition frequencies. To obtain this, we use small detunings ($\Delta_s / g_s < 10$) between the qubit and storage cavity. Figure 8.2 shows spectroscopy in this quasi-dispersive regime as a function of flux bias when the storage cavity is populated with a coherent state ($\langle \hat{n} \rangle \sim 1$). A numerical energy-level calculation is overlaid, showing the positions of various transitions. We define $\omega_{g,e}^n$ as the photon number-dependent transition frequency $|n, g\rangle \rightarrow |n, e\rangle$. Other transitions, such as $|2, g\rangle \rightarrow |0, h\rangle$, are allowed due to the small detuning. Fortunately, we also see that the separation between $\omega_{g,e}^0$ and $\omega_{g,e}^1$ grows rapidly to order $\sim 2g = 140\,\text{MHz}$ as the qubit approaches the storage cavity.

To test the photon meter, we generate single photons in the storage cavity with an adiabatic protocol. Our method uses the avoided crossing between the $|0, e\rangle$ and $|1, g\rangle$ levels to convert a qubit excitation into a photon. The preparation of a photon begins with the qubit detuned below the storage cavity ($\Delta_s \simeq -3g_s$), where we apply a π-pulse to create the state $|0, e\rangle$. We then adiabatically tune the qubit frequency through the avoided crossing with the storage cavity, leaving the system in the state $|1, g\rangle$. The sweep rate is limited by Landau-Zener transitions which keep the system in $|0, e\rangle$. The preparation protocol changes the qubit frequency by 600\,MHz in 50\,ns, giving a spurious transition probability less than 0.1% (calculated with a multi-level numerical simulation). This protocol actually allows for the creation of arbitrary superpositions of $|0, g\rangle$ and $|1, g\rangle$ by changing the rotation angle of
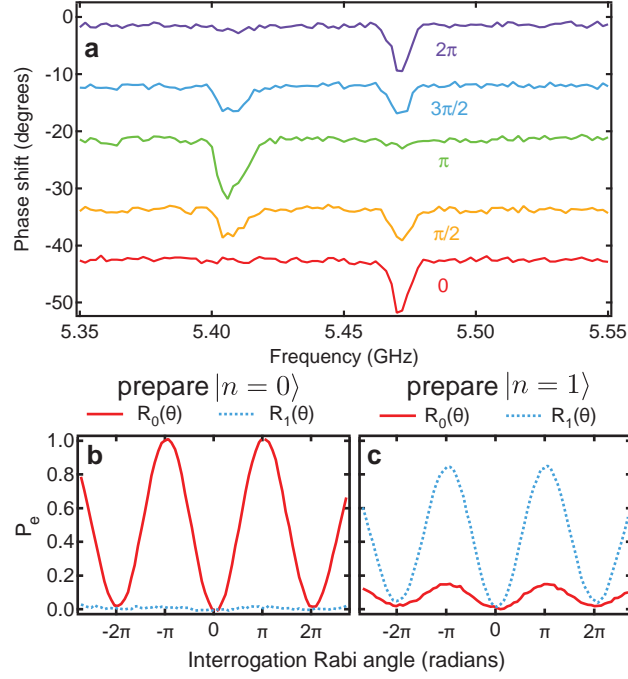
Figure 8.3: Single photon preparation and CNOT selectivity. a, Pulsed spectroscopy vs. Rabi angle of preparation pulse, showing the reflected phase of a pulse at the measurement cavity frequency after a ~ 80 ns pulse near the qubit frequency. Traces are offset vertically for clarity and labeled with the rotation angle of the control pulse used in the preparation step. The dips correspond to $\omega^0_{g,e} \approx 5.47$ GHz and $\omega^1_{g,e} \approx 5.41$ GHz, respectively. **b** and **c**, Rabi driving the qubit transitions after preparing $|n = 0\rangle$ (**b**) and $|n = 1\rangle$ (**c**). The red (blue) traces show the measured qubit excited state probability after applying an interrogation Rabi pulse with varying angle at $\omega^0_{g,e}$ ($\omega^1_{g,e}$). The residual oscillation of $R_1(\theta)$ in **c** is mostly due to preparation infidelity.

the initial pulse. For example, if we use a $\pi/2$-pulse, after the sweep we end up in the state $(|0, g\rangle + e^{i\phi} |1, g\rangle)/\sqrt{2}$, where ϕ is determined by the rotation axis of the $\pi/2$-pulse. One could also use a resonant swap scheme, which has been successfully used to create Fock states [31] up to $|n = 15\rangle$. The method used here has the advantage of being very robust to timing errors.

After the photon is prepared in the storage cavity, the qubit frequency is adjusted such that $\Delta_s/g_s \simeq 5$. At this detuning, the separation between $\omega^0_{g,e}$ and $\omega^1_{g,e}$ is ~ 65 MHz. In Fig. 8.3(a), we show pulsed spectroscopy at this detuning for several rotation angles of the initial preparation pulse. We observe well-resolved dips in the reflected phase of a pulsed signal sent at the measurement cavity frequency. The locations of these dips correspond to

the qubit transition frequencies for $n = 0$ ($\omega_{g,e}^0$) and $n = 1$ ($\omega_{g,e}^1$), and the relative heights match expectations from the different preparation pulse rotations (e.g. a $\pi/2$-pulse results in equal height signals).

To show selective driving of these transitions, we perform Rabi experiments at $\omega_{g,e}^0$ and $\omega_{g,e}^1$ for the cases where we prepare $|0, g\rangle$ and $|1, g\rangle$. In each experiment we ensemble average measurements of the resulting qubit state after further decoupling the qubit from the storage cavity. For the $|0, g\rangle$ case [Fig. 8.3(b)] there is a large amplitude oscillation when the drive is at $\omega_{g,e}^0$ [red, $R_0(\theta)$] and almost no oscillation when the drive is at $\omega_{g,e}^1$ [blue, $R_1(\theta)$]. When we prepare $|1, g\rangle$ the situation is reversed [Fig. 8.3(c)]; however, in this case the residual oscillation of $R_0(\theta)$ (red) is substantial due to small errors in the preparation of $|1, g\rangle$ associated with the initial rotation of the qubit and, more importantly, the ~ 10% probability of energy decay during the subsequent adiabatic sweep through the cavity.

The responses $R_i(\theta)$ are a result of driving $\omega_{g,e}^i$ and the far off-resonant drive of $\omega_{g,e}^j$, where $j \neq i$. The cross-talk is seen in the small residual oscillation of $R_1(\theta)$ in Fig. 8.3(b). Section 8.0.1 derives a method for extracting a selectivity and preparation fidelity from these data, giving a selectivity ≥ 95% for both interrogations and a preparation fidelity of $|\langle n = 1|\psi\rangle|^2 \approx 88\%$. These numbers were confirmed by doing equivalent experiments over a range of preparation pulse rotation angles between 0 and 2π (not shown).

If π-pulses are used in the interrogation step, measurement results of the average qubit state directly correlate with the probability of being in the states $|n = 0\rangle$ or $|n = 1\rangle$. Details of the scaling needed to do this transformation when the selectivity is < 100% are presented in section 8.0.1. These are the desired CNOT operations of the photon meter. If we now insert a variable delay before interrogating, we find that P_0 (P_1), the probability of being in $|n = 0\rangle$ ($|n = 1\rangle$) decays exponentially towards 1 (0), as shown by the red (orange) trace in Fig. 8.4(b). The decay constant of $T_1 \simeq 3.11 \pm 0.02 \, \mu$s agrees with the linewidth of the storage cavity, $1/\kappa_s = 1/(2\pi \, 50 \, \text{kHz}) = 3.2 \pm 0.1 \, \mu$s, measured in a separate, low power ($\bar{n} \sim 1$) reflection experiment.

Strong QND measurements are projective, such that if the measurement observable commutes with the Hamiltonian, the system will remain in an eigenstate of both operators between measurements. Consequently, comparing the results of successive interrogations provides a mechanism to test whether a particular protocol causes additional perturbations on the system. Here, we only compare ensemble average results, because the single-shot qubit readout fidelity for the device is ~ 55%. This is sufficient to reveal processes which change the
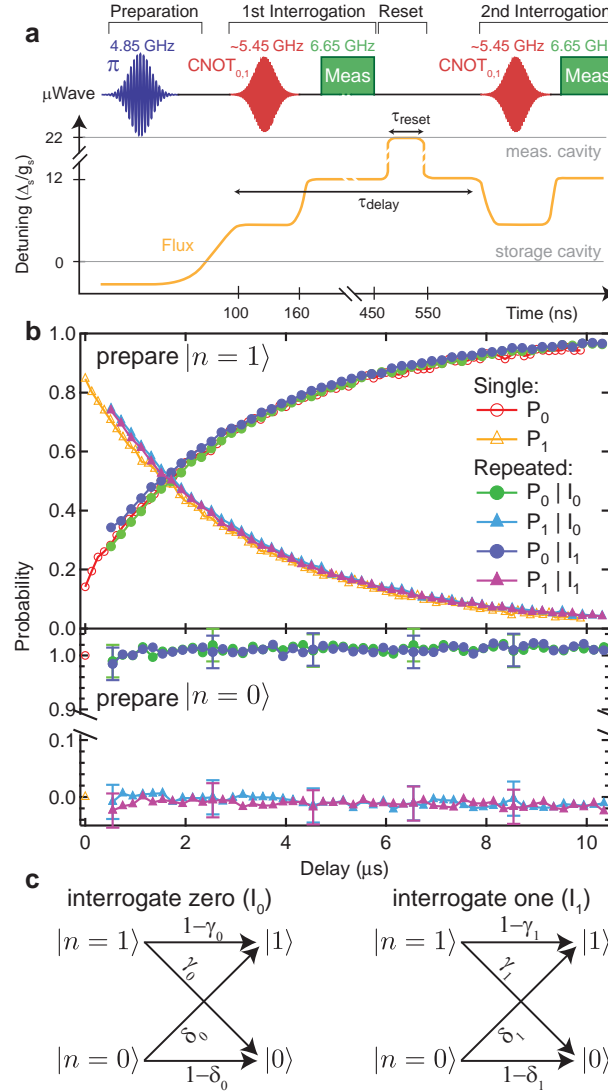
Figure 8.4: Repeated measurements of photons. a, Experiment protocol. A microwave pulse and adiabatic sweep load a single photon into the storage cavity in the preparation step. This photon is interrogated repeatedly by number-selective CNOT gates on the qubit, followed by adiabatic decoupling, qubit readout, and reset. **b,** Single and repeated interrogation after preparing $|n = 1\rangle$ (top) or $|n = 0\rangle$ (bottom), ensemble averaged over $\sim 50,000$ iterations. The near-perfect overlap between single and repeated results demonstrate that the protocol is highly QND. **c,** Transition probability diagrams for the interrogate $n = 0$ (I_0) and interrogate $n = 1$ (I_1) processes. We extract γ_0 (γ_1) = 1 (10) $\pm$ 3% and δ_0 (δ_1) = 7 (3) $\pm$ 3%.

photon number, and technical improvements to interrogation speed or qubit readout fidelity should allow for real-time monitoring of the photon state.

The protocol cannot be repeated immediately, though, because the first interrogation may leave the qubit in the excited state. To circumvent this problem, we use the fast decay rate of the measurement cavity to cause the qubit to spontaneously decay into the 50 Ω environment. The "reset" protocol brings the qubit into resonance with the measurement cavity for a time, τ_{reset} = 50 ns, which is sufficient to reset the qubit with probability ~ 98%, as described in section 6.2 and in [89].

After resetting the qubit, we can interrogate a second time. The full protocol for a repeated interrogation sequence is shown in Fig. 8.4(a). The combination of a CNOT_0 (CNOT_1), a qubit measurement, and a qubit reset define an interrogation process I_0 (I_1). Data for the four possible combinations of interrogating $|n = 0\rangle$ and $|n = 1\rangle$ are shown in Fig. 8.4(b) as a function of delay between the first and second interrogations. The data are ensemble averaged over all results from the first interrogation, so we do not observe projection onto number states. Instead, we again observe exponential decay, where the result of the second measurement is essentially indistinguishable from the first, indicating that the interrogation is highly QND.

Small deviations between first and second interrogations stem from finite photon lifetime in the storage cavity and non-QND processes which cause transitions to other photon numbers [Fig. 8.4(c)]. A single interrogation takes ~ 550 ns, which accounts for the time gap between the single and repeated interrogation points in Fig. 8.4(b). Recording the second interrogation results for different delays allows us to subtract the effect of photon T_1 and calculate the transition probabilities for the I_0 and I_1 processes [90]. In principle, I_0 and I_1 can cause transitions to photon numbers outside of the $n \in \{0, 1\}$ manifold; however, the absence of statistically significant deviations from $P_0 + P_1 = 1$ suggests that any such effects are negligible. Instead, we consider only transitions from $|n = 1\rangle \to |n = 0\rangle$ or $|n = 0\rangle \to |n = 1\rangle$, characterized by the probabilities γ_i and δ_i, respectively, where i references the interrogation process I_0 or I_1. We observe γ_0 (γ_1) = 1 (10) ± 3% and δ_0 (δ_1) = 7 (3) ± 3%, demonstrating that this protocol is highly QND.

The protocol presented here is a fast and highly QND measurement of single photons, which we believe can be extended to detect higher photon numbers. It should be possible to demonstrate the projective nature of the interrogation and create highly non-classical states of light via post selection, and eventually with higher fidelity readout one could observe quantum jumps of light in a circuit.

8.0.1   Measured Voltage Scaling

When the interrogation selectivity is less than 100%, we need to account for undesired rotations to correctly calculate the state probabilities from the measured voltages. The details of our calibration procedure follow.

If we prepare $|n = 0\rangle$ or $|n = 1\rangle$ at time $t = 0$, when we interrogate at some later time there is an additional probability p_d to decay, giving the density matrices

$$\rho^0 = |g\rangle \langle g| \otimes |0\rangle \langle 0|, \tag{8.1a}$$

$$\rho^1 = |g\rangle \langle g| \otimes \{p_d |0\rangle \langle 0| + (1 - p_d) |1\rangle \langle 1|\}, \tag{8.1b}$$

where ρ^i indicates preparing state $|i\rangle$ at $t = 0$. We can model the interrogation pulses as operations which act on ρ^i:

$$U_0 = R_y(\pi) \otimes |0\rangle \langle 0| + R_y(\epsilon') \otimes |1\rangle \langle 1|,$$

$$U_1 = R_y(\epsilon) \otimes |0\rangle \langle 0| + R_y(\pi) \otimes |1\rangle \langle 1|,$$

$$U_I = \mathbb{1},$$

where ϵ and ϵ' are small angles. After interrogation, the integrated homodyne response is

$$W_r^n = V_g + \Delta V \cdot \text{tr}(\Pi_e U_r \rho^P U_r),$$

where $n \in \{0, 1\}$ is the Fock state of the cavity, $r \in \{0, 1, I\}$, V_g (V_e) is the voltage measured when the qubit is in $|g\rangle$ ($|e\rangle$), $\Delta V = V_e - V_g$, and $\Pi_e = |e\rangle \langle e|$.

By abusing the notation slightly and treating ϵ, ϵ' as probabilities rather than rotation angles, we can calculate the W_r^n

$$W_I^0 = V_g,$$

$$W_0^0 = V_g + \Delta V,$$

$$W_1^0 = V_g + \Delta V \cdot \epsilon,$$

$$W_0^1 = V_g + \Delta V (\epsilon'(1 - p_d) + p_d),$$

$$W_1^1 = V_g + \Delta V ((1 - p_d) + \epsilon p_d).$$

We measure these five voltages in calibration experiments and invert the equations to find the parameters V_g, ΔV, p_d, ϵ, and ϵ'. Note that this does not require perfect preparation fidelity because the model includes decay between preparation and interrogation p_d which

will also capture any fixed preparation infidelity. This gives the selectivities, $(1-\epsilon)$ and $(1-\epsilon')$, as well as the preparation fidelity, $(1-p_d)$.

An unknown mixture of $n = 0$ and $n = 1$ is characterized by a single probability p,

$$\rho = |g\rangle\langle g| \otimes \{p\,|0\rangle\langle 0| + (1-p)\,|1\rangle\langle 1|\},$$

which produces the responses

$$W_0^\rho = V_g + \Delta V\,(p + \epsilon'(1-p)),$$
$$W_1^\rho = V_g + \Delta V\,(\epsilon \cdot p + (1-p)).$$

This leads to a simple rescaling to transform W_0^ρ and W_1^ρ into P_0 and P_1

$$P_0 = \frac{W_0^\rho - (V_g + \Delta V \epsilon')}{\Delta V(1-\epsilon')},$$
$$P_1 = \frac{W_1^\rho - (V_g + \Delta V \epsilon)}{\Delta V(1-\epsilon)}.$$

8.0.2   Error Estimate

The primary challenge in these experiments is obtaining sufficiently accurate and precise control of the qubit frequency to do high-fidelity operations. The narrow bandwidth pulses used in the CNOT operations means that even a few MHz error in frequency control results in a significant rotation error. We use deconvolution techniques similar to those described in the supplement of Ref. 31 and in section 4.3; however, the flux bias current response function drifts on a time scale of about one day, making it difficult to eliminate all classical control errors. Even after applying corrections, there is a remaining spread of $2-3\,\mathrm{MHz}$ in the qubit frequencies over the various realizations of preparation to interrogation delay. This translates into a $2-3\%$ error in the probability to find the qubit in $|e\rangle$ after applying a conditional π-pulse. The errors bars reported in the lower panel of Fig. 8.4(b) are due to this systematic error.

Conclusions and Outlook

THE demonstration of QND photon detection presented in chapter 8 is only the first step in photon creation and detection experiments in two-cavity circuit QED. That work can be expanded upon in a number of directions, many of which I believe will be fruitful.

Faster, higher-fidelity readout

Since the photon detection work of chapter 8, recent experiments in Schoelkopf Lab have demonstrated a new readout technique that uses the non-linearity of the Jaynes-Cummings ladder to selectively drive the cavity to a 'bright' state with a large average photon number [91, 92]. The power needed to drive this state is ~ 50, 000 larger than the power used for the usual linear dispersive readout with only a few photons. The result is a fast, high fidelity (~ 90%) qubit readout. This technique is likely not QND to the qubit, but the photon measurement protocol of chapter 8 does not require the qubit readout to be QND, since the qubit state is reset after each measurement.* Use of this new measurement protocol will shorten qubit readout time, increasing the number of QND measurements that can be done in a cavity T_1, and increase the information gained per measurement. This might open the

---

* A bigger unknown is whether it will take longer to reset the qubit since it is believed that the high-power readout also highly excites the transmon.
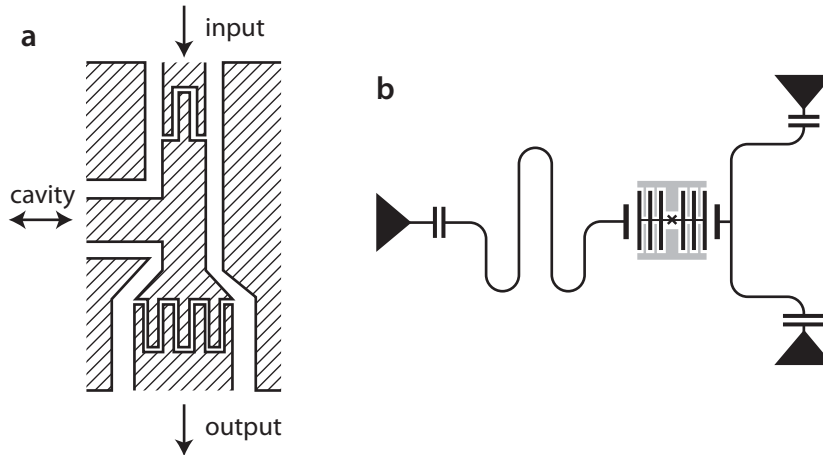
Figure 9.1: Transmission readout geometries. a A single-ended transmission coupler which asymmetrically couples a cavity (that is extended to the left) to two ports. The cavity is driven through the top port, and the amplifier connects to the bottom port. **b** In this variation, the transmon is coupled to the first harmonic of the readout cavity in the center of the mode.

way to experiments analogous to the observation of quantum jumps in Rydberg atom cQED [14, 44].

In order to use this 'bright-state' readout, the two-cavity circuit design should be modified to probe the readout cavity in transmission. To see why, consider that in an ideal reflection measurement, the reflected amplitude is always the same. The 'bright' state might be visible in the phase of the reflected signal, but even though the high-power readout excites the cavity to a large photon number $n \sim 1,000$, an even larger number of photons ($\sim 50,000$) reflect off the input coupler without entering the cavity. Consequently, the bright state signal adds with a very large background, making it difficult to distinguish. Furthermore, when the cavity switches to the bright state, it responds at the bare cavity frequency. For typical circuit QED parameters, the frequency shift of the cavity is substantially more than the cavity linewidth. The reflected phase also undergoes a 2π rotation when probing the cavity response through its resonance frequency. Consequently, for most choices of drive frequencies, the reflected phase actually returns to the same value!

Fortunately, modifying the two-cavity circuit to probe the readout cavity in transmission is a relatively minor change. One possibility leaves intact the basic geometry of coupling to each cavity at the far end from the sarantapede transmon, but replaces the coupler on the readout cavity with something like the sketch shown in figure 9.1(a). The idea here is to achieve the same kind asymmetric coupling that is typical in single-cavity circuit QED
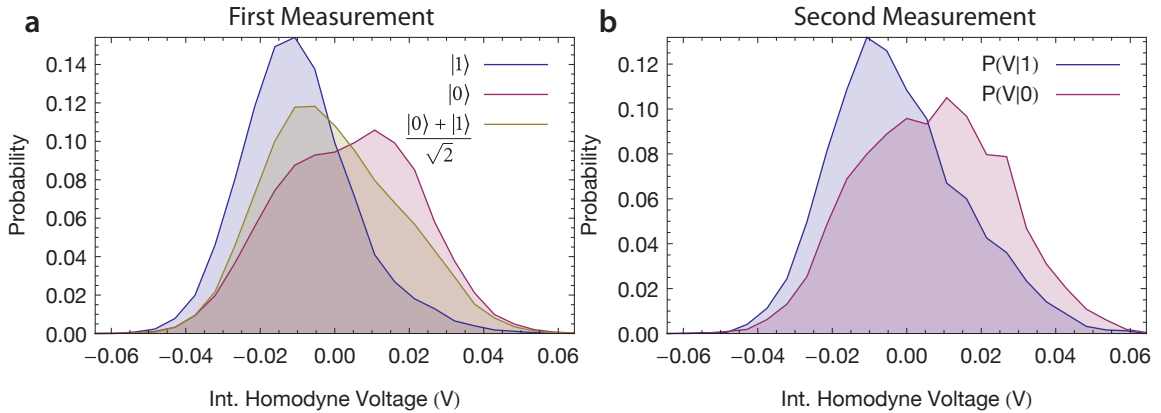
Figure 9.2: Single shot histograms of photon readout. a Histograms of the integrated homo-
dyne voltage after preparing $|0\rangle$ (red), $|1\rangle$ (blue), and $(|0\rangle + |1\rangle)/\sqrt{2}$ (yellow) after performing
a $CNOT_0$ gate, showing the voltage probability in 20,000 shots. The integration time τ, is
roughly similar to the transmon T_1; consequently, the $|0\rangle$ state histogram has an extra bump
corresponding to qubit relaxation. The histogram of the superposition state looks very much
like the average of the 0 and 1 histograms. The *photon* readout fidelity is ~ 40%. **b** Histograms
of a second measurement conditioned on the result of the first with thresholds $V_0 = 0.02$ and
$V_1 = -0.02$.

devices, but just place the two finger capacitors on the same end of the cavity. Clearly, this
kind of design gives a strong channel to directly couple the input and output ports, which
might be undesirable. Another possibility is to couple the transmon to the first harmonic
of the readout cavity by placing it in the center, as shown in figure 9.1(b). This gives better
isolation between the input and output ports, but presents a larger classical admittance to
the qubit since the qubit frequency is necessarily above the first mode. Consequently, this
second design would benefit from the additional of a Purcell filter as well.

State Projection

Already, the linear, dispersive qubit measurement gives a single-shot qubit readout fidelity of
50–60% when using fast pulses to prepare the qubit in $|e\rangle$. The longer pulses necessary for
the selective CNOT operations result in a reduced *photon* readout fidelity of ~ 40%, as shown
in figure 9.2(a), simply because CNOT pulse is a larger fraction of the qubit T_1. Even with this
fidelity, one can observe evidence of state projection in repeated measurements. With Brian
Vlastakis, we have begun exploring single-shot repeated measurements of photons using
the existing cQED291 sample. The three histograms in figure 9.2(a) show the distribution of
integrated homodyne voltage after a $CNOT_0$ gate for $|0\rangle$ (red), $|1\rangle$ (blue), and $(|0\rangle + |1\rangle)/\sqrt{2}$

(yellow) states in the storage cavity. Since the $CNOT_0$ gate rotates the qubit into the $|e\rangle$ state when there are no photons in the cavity, the red histogram displays an extra 'bump' due to qubit relaxation during the integration period*. The histogram for the superposition state displays the expected average of the other two histograms.

Since the $CNOT_0$ gate entangles the qubit and cavity states, readout of the qubit state causes projection of the cavity state. Even with mediocre single-shot fidelity, it is possible to observe evidence of this state collapse with repeated measurements. Given a particular first measurement result, one can assign that result to the $|0\rangle$ or $|1\rangle$ distribution with a probability which depends on the value of the measurement. Voltages which come from the tails of the distributions can be assigned to one state or the other with much higher confidence than voltages in the center of the distributions where it is nearly equally likely to correspond to either state. If one insists on keeping all measurement results, the best one can do is to set a threshold at the point where the two distributions intersect. However, with the given measurement fidelity, this will give a relatively high probability to misidentify the state. Instead, we set two thresholds, V_0 and V_1, and say that all voltages greater than V_0 belong to the $|0\rangle$ distribution, and voltages less than V_1 belong to the $|1\rangle$ distribution. Results between V_0 and V_1 are discarded as 'unknown'. Given these first assignments of '0' and '1', we analyze the resulting distributions of second measurements for these two bins. These second measurement histograms are shown in figure 9.2(b) for $V_0 = 0.02$ and $V_1 = −0.02$. One sees that the sorting from the first measurement splits the histograms into distinct shapes that closely match the results for measurements of $|0\rangle$ and $|1\rangle$ states. In theory this separation can be improved as the thresholds are pushed further out into the tails, at the cost of throwing out increasing amounts of data. Brian is working on ways to more efficiently collect this data for further study of many repeated measurements.

Readout of higher photon numbers

The photon detection experiment of chapter 8 demonstrated QND readout of the $n = 0$ and 1 photon states. A fair criticism of that work is that when limited to a two-state manifold, the bosonic character of the cavity system is not apparent. Consequently, one would like to extend the detection system to higher photon numbers. This is also a requirement for Wigner tomography using the photon number probability counting method presented in section 2.5.

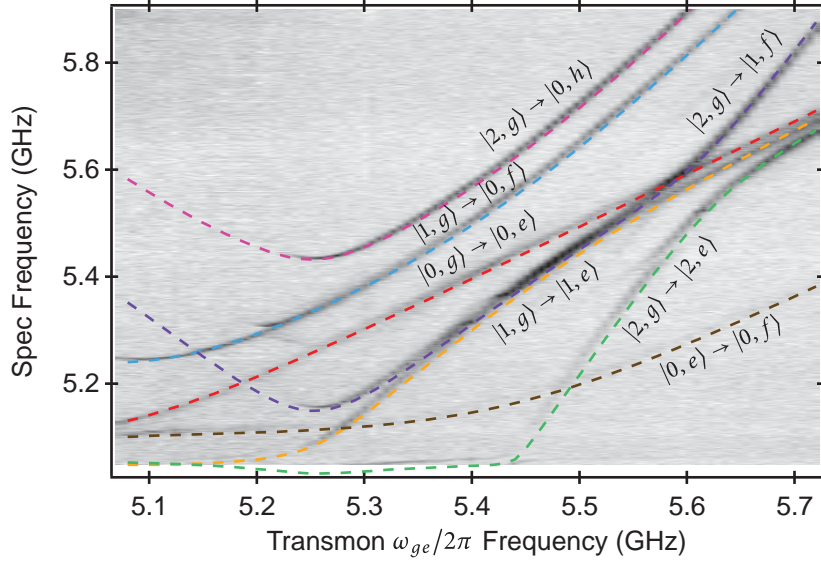* See [81] for more information on protocols for readout using weak continuous measurement.

Figure 9.3: Quasi-dispersive spectroscopy of $|\psi\rangle \simeq |2, g\rangle$. Spectroscopy after initially preparing $|0, f\rangle$ with the qubit below the storage cavity, and then ramping through the avoided crossing in 100 ns before doing pulsed spectroscopy vs. flux bias. The x-axis has been rescaled to match the predicted uncoupled $\omega_{ge}/2\pi$ frequency of the transmon. This spectroscopy map reveals a host of transitions, some of which have small detunings from transitions we would like to use, such as ω_{ge}^1. Dashed lines are calculated numerically using the generalized Jaynes-Cummings Hamiltonian for a single cavity and transmon. The spectrum closely matches the calculation except for avoided crossings that are visible in the ω_{ge}^2 (green) and $|1, g\rangle \to |0, f\rangle$ (turquoise) lines. The source of these avoided crossings is unknown.

In this vein, I have begun initial investigations into preparing and detecting higher photon numbers with the same cQED291 sample used in the previous experiment.

To extend to $n = 2$, one needs to find the appropriate transition to drive for the CNOT_2 gate. One way to do this would be to populate the storage cavity with a coherent state and then perform quasi-dispersive spectroscopy. The disadvantage of this approach is that population of higher photon numbers has a multiplicative effect on the number of allowed transitions. To minimize the chance of mis-identifying these transitions, I pursued a different approach of performing pulsed spectroscopy after attempting to prepare only the $|n = 2\rangle$ state. As discussed in section 2.5.1, the $|2, g\rangle$ state can be prepared by rapid adiabatic passage by first driving to $|0, f\rangle$ and then sweeping through the avoided crossing with the cavity. Figure 9.3 shows pulsed spectroscopy vs. flux bias after such a preparation sequence with a 100 ns sweep. Since the spectroscopy pulse length is 300 ns and preparation takes 100 ns, we also expect to see transitions from populating $|0, g\rangle$ and $|1, g\rangle$. As figure 9.3 shows, even in this reduced

complexity situation, there are many allowed transitions. Overlaid on top of the data are numerically calculated transition frequencies and labels identifying them using the scheme of section 2.2.4.*

Unfortunately, in addition to the desired ω_{ge}^n transitions(red, orange, and green), there are other transitions like $|2, g\rangle \rightarrow |1, f\rangle$ (purple) that are nearly overlapping. There is another set of well-separated number-dependent transitions at higher frequencies (light blue and pink); however, these are not-photon number preserving, so they cannot be used for the CNOT gates. Driving these transitions with π-pulses can be used for destructive readout, though. Furthermore, these transitions could be used in a modified protocol akin to [41] where 2π rotations are done on these transitions. If the transmon is initially prepared in a superposition of levels, such a rotation will cause one of the states to pick up a minus sign. For instance, if the transmon is initially prepared in $(|g\rangle + |f\rangle)\sqrt{2}$ and the $|2, g\rangle \rightarrow |0, h\rangle$ transition is driven with a 2π-pulse, then the transmon will end up in $(-|g\rangle + |f\rangle)/\sqrt{2}$ if the cavity is in the $|n = 2\rangle$ state. This minus sign can be detected by inverting the transmon preparation procedure. Such a scheme adds substantial of complexity, however.

An easier path forward may be to swap the cavity frequencies such that the storage cavity is at a higher frequency than the measurement cavity. This allows storage cavity state to be read out at negative detunings. As discussed in section 2.2.4, the ω_{ge}^n transitions have nearly linear shifts at negative detunings. Furthermore, the negative anharmonicity of the transmon causes the undesired transitions to have larger negative shifts, so they do not overlap the ω_{ge}^n transitions. As with the discussion of transmission readout in the previous section, having the low-Q readout cavity negatively detuned from the qubit has adverse consequences on relaxation. Consequently, as with the previous case, I would recommend adding a Purcell filter for this situation.

Cross-Kerr readout

The cross-Kerr effect discussed in section 3.5 presents an alternative method to implement QND photon detection. Instead of mapping photon states to qubit states, in this protocol the qubit remains in the ground state, and the presence of photons in the storage cavity is detected through the qubit-mediated frequency shift (cross-Kerr) of the readout cavity. In principle, one should be able to continuously monitor the phase of a reflected signal from the

---

* It turns out that such a spectroscopy map is a good way to find system parameters like g and E_C, since the position of these transitions in the quasi-dispersive regime depends quite sensitively on these parameters.
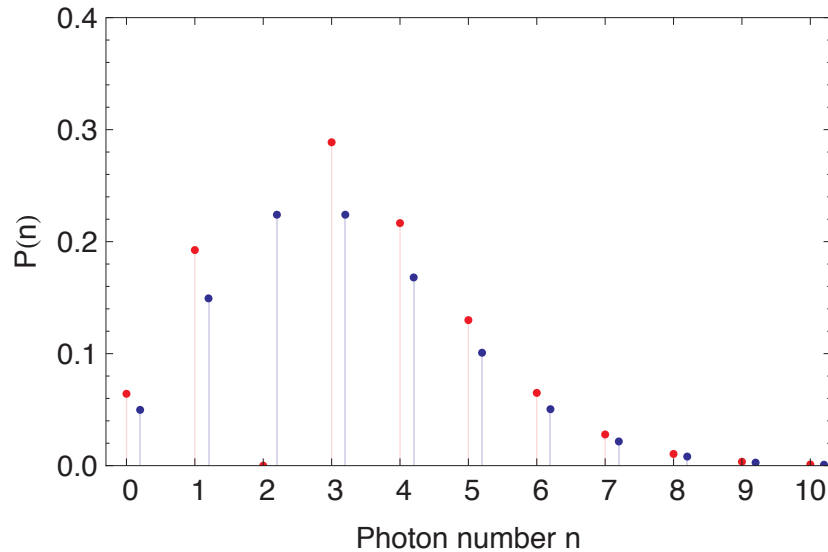
Figure 9.4: Coherent state minus a Fock state. Plot of the conditional probabilities (red) to occupy Fock states $n = 0$–10 when starting with a coherent state with $\bar{n} = 3$ and observing a negative result from probing the occupation of $n = 2$. For comparison, the standard Poisson distribution is shown in blue.

readout cavity to detect the presence of photons in the storage cavity. In practice, however, the cross-Kerr effect is fairly week, leading to a substantial trade-off between cavity Q and readout speed. Initial investigations by Eran Ginossar assuming weak driving ($\bar{n} \sim 10$) of a moderate Q ($\sim 10,000$) cavity showed insufficient SNR to convincingly detect photon jumps given the noise of the following HEMT amplifier. If the HEMT amplifier is replaced with a lower noise device, like a DC-SQUID amp [93] or a JPC [94], this scheme might become feasible. Another alternative is to attempt using the cross-Kerr effect as the source of nonlinearity for the bright state readout already demonstrated with transmons [91, 92]. Given that the 'size' of the nonlinearity required for this readout is not yet well understood, it is unclear if this will work. One possibility to increase the size of the effect is to adiabatically bring the qubit into resonance with the storage cavity before turning on the high power readout. Since the $|\pm\rangle = (|1, e\rangle \pm |0, g\rangle)/\sqrt{2}$ states present different dispersive shifts to the readout cavity, one might expect this to work similarly to the bright state readout of the transmon.

State preparation

In addition to using single-shot measurements to demonstrate the projective nature of the photon readout, one can also prepare quantum states by post selection on measurement results. For example, the standard way to prepare Fock states in the Rydberg atom experiments is to initially populate the cavity with a coherent state, then do Ramsey experiments on a series of atoms (~ 100 of them) as they pass through the cavity until the phase shift imparted by the cavity is projected onto a definite value, at which point the cavity is in a Fock state. With the photon number selectivity of the qubit-photon CNOT gates, though, one can do something more unusual such as prepare a state like a coherent state 'minus' a Fock state, *i.e.* $|\psi\rangle = (I - \Pi_n)/\sqrt{p_n}|\alpha\rangle$, where $\Pi_n = |n\rangle\langle n|$ is the projector onto Fock state $|n\rangle$ and $p_n = |\langle n|\alpha\rangle|^2$ is the prior probability to be in state $|n\rangle$. This state is simply created with a single measurement by conditioning on a negative result after performing a $CNOT_n$ gate. The resulting photon number distribution, shown in figure 9.4, is very strange looking, as it is very Poisson-like, but a single photon number 'knocked out'. The selectivity of the CNOT gates allows very efficient conditional generation of such states.

With highly selective gates one can also tailor the information gained per measurement by using different photon-number dependent gates. For instance, instead of doing a π-pulse on a single ω_{ge}^n transition, one did π/N-pulses on N transitions, measurement of the final qubit state reveals joint information about all N number states. In fact, it is in principle possible to construct a photon readout which is sensitive to an arbitrary sum of diagonal elements of the density matrix (written in the Fock state basis). This can be used to create a broader class of quantum states by post-selection.

Of course, Hofheinz *et al.* [31] have already demonstrated the Law and Eberly protocol [23] for deterministic generation of arbitrary cavity states. Given this, one might wonder if state preparation is a 'solved' problem. As I have just argued, though, for certain quantum states there are much more efficient state preparation protocols, so long one accepts a conditional rather than deterministic method. Consequently, we could ask if there is a way to get the best of both worlds by combining deterministic preparation to certain initial states followed by QND measurement. There is a sense in which the methods are incompatible because resonant swap requires a small coupling g to accurately control the swap time, while the selective QND detection requires large g to get sufficient separation of the number-dependent qubit transitions. Consequently, to be able to do swaps and QND detection, one either needs to have two qubits coupled to the storage cavity, one with a large g and the other with a small
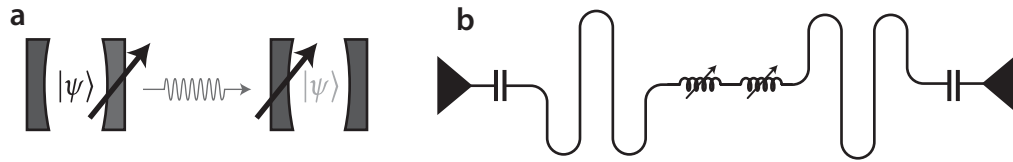
Figure 9.5:   State transfer with tunable mirrors. a A quantum state of one cavity can be transferred to another by making cavities with tunable mirrors. The coupling rate of the left cavity is increased while it is decreased on the right cavity. If the timing is done correctly, the photons are completely captured in the right cavity. **b** A sketch of an implementation using CPW cavities by using inductive rather than capacitive coupling. Each tunable inductor represents a SQUID loop with two Josephson junctions.

g. Or, one should design a means to have a tunable coupling strength between qubit and cavity. In fact, the TIE Fighter or Three Island Qubit (TIQ) proposed in [16] with David Schuster, and also independently developed by Jay Gambetta and Andrew Houck, has just such a tunable coupling. Consequently, I think it would be extremely attractive to use such a qubit for cavity state preparation.

Finally, moving beyond state preparation of a single cavity, one can begin to think about preparation of arbitrary states of two or more cavities. Strauch *et al.* [22] have extended the Law and Eberly protocol for arbitrary state preparation of two cavities. Their scheme requires exactly the number-dependent qubit transitions which are available in circuit QED with large coupling strengths. One way to realize this protocol with transmon qubits is to create a three cavity, two qubit device where two of the cavities are high-Q for storing the prepared state, and the third cavity is low-Q for qubit readout. Both qubits are strongly coupled to the readout cavity, but one qubit has weak coupling to the storage cavities for performing swaps, while the other qubit is strongly coupled for the number dependent operations. Building a device with this many cavities and qubits will require understanding some of the issues with scaling circuit QED to larger systems. Consequently, this type of experiment is a useful stepping stone for pushing in that direction.

Quantum state transfer

Continuing the idea of state preparation in multiple cavities, on might consider possibilities to distribution prepared states among a network of cavities. One method proposed (in the optical domain) in [95] suggests using cavities with tunable mirrors, as shown in cartoon form in figure 9.5(b). The idea is to have time-domain control of the decay rates of the cavities,

such that one effectively has valves that can be opened or closed to controllably trap photons in a cavity or let them leak out. In a chain of such cavities, one can prepare a state in one cavity, and then send that state to another cavity by slowly increasing the decay rate while simultaneously decreasing the decay rate of the next cavity in the chain. Said differently, one starts with the 'valve' on the first cavity fully closed and the valve on the next cavity fully open. The valve of the first cavity is slowly opened while the next valve is slowly closed. Noh shows that under certain conditions, this provides perfect state transfer.

I do not know how to make a tunable capacitor; however, tunable inductors are easy to make in superconducting circuits—they are just SQUID loops with two Josephson junctions. Inductors make perfectly good 'mirrors' for CPW cavities, so one might imagine implementing this idea with a design like figure 9.5(b). One also probably needs qubits coupled to each cavity for state preparation and readout, respectively. For this to work, one first needs to show that it is possible to make high-Q cavities with tunable couplers. After that, verifying that it works requires some combination of the state preparation and detection techniques discussed in this thesis. Consequently, this is a difficult experiment, but it is an exceedingly simple way of thinking about quantum state distribution in a network of cavities.

Final thoughts

I hope to have convinced you that circuit QED is a field with abundant possibilities for producing, manipulating, and detecting quantum light. The work of this thesis has added several new tools to the circuit QED toolbox, including a two-cavity architecture with a qubit that couples to multiple cavities, fast control methods, improved understanding of relaxation, and a QND photon detection protocol. And yet, I hope that this is just the beginning of a vast array of increasingly sophisticated experiments involving networks of cavities and qubits with the ability to create and shuttle quantum light around a chip at-will. The challenges of building these experiments are formidable, but the promise of new understanding of entanglement and potential technologies are sufficiently exciting that I expect people will find ways of overcoming these challenges in eager anticipation of what is to come.

❦

Bibliography

1. M. H. Devoret, J. M. Martinis, and J. Clarke, "Measurements of macroscopic quantum tunneling out of the zero-voltage state of a current-biased Josephson junction," *Physical Review Letters* **55**, 1908–1911 (1985). Cited on page 15.

2. J. Clarke, A. Cleland, M. Devoret, D. Esteve, and J. Martinis, "Quantum mechanics of a macroscopic variable: the phase difference of a Josephson junction," *Science* **239**, 992 (1988). Cited on page 15.

3. V. Bouchiat, D. Vion, P. Joyez, D. Esteve, and M. Devoret, "Quantum coherence with a single Cooper pair," *Physica Scripta* **1998**, 165 (1998). Cited on page 15.

4. Y. Nakamura, Y. Pashkin, and J. Tsai, "Coherent control of macroscopic quantum states in a single-Cooper-pair box," *Nature* **398**, 786–788 (1999). Cited on pages 15 & 105.

5. D. Vion, A. Aassime, A. Cottet, P. Joyez, H. Pothier, C. Urbina, D. Esteve, and M. Devoret, "Manipulating the quantum state of an electrical circuit," *Science* **296**, 886 (2002). Cited on pages 15, 24, 105 & 113.

6. J. Koch, T. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. Devoret, S. Girvin, and R. J. Schoelkopf, "Charge-insensitive qubit design derived from the Cooper pair box," *Phys. Rev. A* **76**, 042319 (2007). Cited on pages 15, 25, 26, 32, 78, 80 & 127.

7. V. E. Manucharyan, J. Koch, L. I. Glazman, and M. H. Devoret, "Fluxonium: Single Cooper-pair circuit free of charge offsets," *Science* **326**, 113–116 (2009). Cited on page 15.

8. M. Steffen, M. Ansmann, R. C. Bialczak, N. Katz, E. Lucero, R. Mcdermott, M. Neeley, E. M. Weig, A. N. Cleland, and J. M. Martinis, "Measurement of the entanglement of

two superconducting qubits via state tomography," *Science* **313**, 1423 (2006). Cited on page 15.

9. J. Majer, J. M. Chow, J. M. Gambetta, J. Koch, B. R. Johnson, J. A. Schreier, L. Frunzio, D. I. Schuster, A. A. Houck, and A. Wallraff, "Coupling superconducting qubits via a cavity bus," *Nature* **449**, 443–447 (2007). Cited on pages 15 & 110.

10. L. DiCarlo, J. M. Chow, J. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, "Demonstration of two-qubit algorithms with a superconducting quantum processor," *Nature* **460**, 240–249 (2009). Cited on pages 15, 99, 114 & 128.

11. L. DiCarlo, M. D. Reed, L. Sun, B. R. Johnson, J. M. Chow, J. M. Gambetta, L. Frunzio, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf, "Preparation and measurement of three-qubit entanglement in a superconducting circuit," *Nature* **467**, 574–578 (2010). Cited on page 15.

12. M. Neeley, R. C. Bialczak, M. Lenander, E. Lucero, M. Mariantoni, A. D. O'Connell, D. Sank, H. Wang, M. Weides, J. Wenner, Y. Yin, T. Yamamoto, A. N. Cleland, and J. M. Martinis, "Generation of three-qubit entangled states using superconducting phase qubits," *Nature* **467**, 570–573 (2010). Cited on page 15.

13. S. Haroche and J. Raimond, *Exploring the Quantum: Atoms, Cavities, and Photons*. Oxford University Press, 2006. Cited on pages 15, 37 & 40.

14. C. Guerlin, J. Bernu, S. Deléglise, C. Sayrin, S. Gleyzes, S. Kuhr, M. Brune, J.-M. Raimond, and S. Haroche, "Progressive field-state collapse and quantum non-demolition photon counting," *Nature* **448**, 889–893 (2007). Cited on pages 16, 46, 47, 48, 126, 127 & 137.

15. D. I. Schuster, *Cavity Quantum Electrodynamics*. PhD thesis, Yale University, 2007. Cited on pages 16, 24, 25, 27, 29, 30, 31, 53, 61, 63, 74, 89, 90, 94, 96, 110, 123 & 124.

16. L. S. Bishop, *Circuit Quantum Electrodynamics: Vol. II*. PhD thesis, Yale University, 2010. Cited on pages 16, 24, 25, 33, 144 & 155.

17. D. I. Schuster, A. A. Houck, J. A. Schreier, A. Wallraff, J. Gambetta, A. Blais, L. Frunzio, J. Majer, B. R. Johnson, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, "Resolving photon number states in a superconducting circuit," *Nature* **445**, 515–518 (2007). Cited on pages 16, 30, 110, 114, 119 & 126.

18. A. A. Houck, D. I. Schuster, J. Gambetta, J. A. Schreier, B. R. Johnson, J. M. Chow, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, "Generating single microwave photons in a circuit," *Nature* **449**, 328–331 (2007). Cited on pages 16, 110 & 126.

19. L. S. Bishop, J. M. Chow, J. Koch, A. A. Houck, M. H. Devoret, E. Thuneberg, S. M. Girvin, and R. J. Schoelkopf, "Nonlinear response of the vacuum Rabi resonance," *Nature Physics* **5**, 105–109 (2009). Cited on pages 16, 28 & 90.

20. C. H. Bennett, F. Bessette, G. Brassard, L. Salvail, and J. Smolin, "Experimental quantum cryptography," *Journal of Cryptology* **5**, 3–28 (1992). Cited on page 19.

21. D. Walls and G. Milburn, *Quantum Optics*. Springer, 2006. Cited on pages 19, 37 & 41.

22. F. W. Strauch, K. Jacobs, and R. W. Simmonds, "Arbitrary control of entanglement between two superconducting resonators," *Phys. Rev. Lett.* **105**, 050501 (2010). Cited on pages 19 & 144.

23. C. Law and J. Eberly, "Arbitrary control of a quantum electromagnetic field," *Physical Review Letters* **76**, 1055–1058 (1996). Cited on pages 19, 43, 47 & 143.

24. M. Hofheinz, E. M. Weig, M. Ansmann, R. C. Bialczak, E. Lucero, M. Neeley, A. D. O'Connell, H. Wang, J. M. Martinis, and A. N. Cleland, "Generation of Fock states in a superconducting quantum circuit," *Nature* **454**, 310–314 (2008). Cited on pages 20, 29, 40, 41, 42, 43, 47, 86, 126 & 127.

25. P. J. Leek, M. Baur, J. M. Fink, R. Bianchetti, L. Steffen, S. Filipp, and A. Wallraff, "Cavity quantum electrodynamics with separate photon storage and qubit readout modes," *Phys. Rev. Lett.* **104**, 100504 (2010). Cited on pages 20, 29, 44 & 127.

26. M. H. Devoret, *Quantum Fluctuations: Les Houches Session LXIII*. Elsevier, 1997. Cited on pages 24, 67, 68, 106 & 107.

27. A. Blais, R. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, "Cavity quantum electrodynamics for superconducting electrical circuits: an architecture for quantum computation," *Physical Review A* **69**, (2004). Cited on pages 27, 29, 53, 54, 59, 106 & 107.

28. J. M. Fink, M. Göppl, M. Baur, R. Bianchetti, P. J. Leek, A. Blais, and A. Wallraff, "Climbing the Jaynes-Cummings ladder and observing its nonlinearity in a cavity QED system," *Nature* **454**, 315–8 (2008). Cited on page 28.

29. M. Brune, F. Schmidt-Kaler, A. Maali, J. Dreyer, E. Hagley, J. M. Raimond, and S. Haroche, "Quantum Rabi oscillation: A direct test of field quantization in a cavity," *Phys. Rev. Lett.* **76**, 1800–1803 (1996). Cited on pages 29, 126 & 127.

30. T. Meunier, S. Gleyzes, P. Maioli, A. Auffeves, G. Nogues, M. Brune, J.-M. Raimond, and S. Haroche, "Rabi oscillations revival induced by time reversal: A test of mesoscopic quantum coherence," *Physical Review Letters* **94**, 010401 (2005). Cited on page 29.

31. M. Hofheinz, H. Wang, M. Ansmann, R. C. Bialczak, E. Lucero, M. Neeley, A. D. O'Connell, D. Sank, J. Wenner, and J. M. Martinis, "Synthesizing arbitrary quantum states in a superconducting resonator," *Nature* **459**, 546–549 (2009). Cited on pages 29, 40, 41, 42, 43, 47, 126, 127, 130, 135 & 143.

32. D. I. Schuster, A. Wallraff, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. M. Girvin, and R. J. Schoelkopf, "AC Stark shift and dephasing of a superconducting qubit strongly coupled to a cavity field," *Physical Review Letters* **94**, 123602 (2005). Cited on pages 30 & 124.

33. J. Gambetta, A. Blais, D. I. Schuster, A. Wallraff, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, "Qubit-photon interactions in a cavity: Measurement induced dephasing and number splitting," *Phys. Rev. A* **74**, 042318 (2006). Cited on pages 30, 119, 124 & 126.

34. V. B. Braginsky and F. Y. Khalili, *Quantum Measurement*. Cambridge University Press, 1992. Cited on page 36.

35. V. B. Braginsky and F. Y. Khalili, "Quantum nondemolition measurements: the route from toys to tools," *Rev. Mod. Phys.* **68**, 1–11 (1996). Cited on pages 36 & 126.

36. G. Nogues, A. Rauschenbeutel, S. Osnaghi, P. Bertet, M. Brune, J.-M. Raimond, S. Haroche, L. Lutterbach, and L. Davidovich, "Measurement of a negative value for the Wigner function of radiation," *Physical Review A* **62**, 054101 (2000). Cited on page 40.

37. P. Bertet, A. Auffeves, P. Maioli, S. Osnaghi, T. Meunier, M. Brune, J. Raimond, and S. Haroche, "Direct measurement of the Wigner function of a one-photon Fock state in a cavity," *Physical Review Letters* **89**, 200402 (2002). Cited on pages 40 & 41.

38. S. Deleglise, I. Dotsenko, C. Sayrin, J. Bernu, M. Brune, J.-M. Raimond, and S. Haroche, "Reconstruction of non-classical cavity field states with snapshots of their decoherence," *Nature* **455**, 510–514 (2008). Cited on pages 40 & 45.

39. D. Smithey, M. Beck, M. Raymer, and A. Faridani, "Measurement of the Wigner distribution and the density matrix of a light mode using optical homodyne tomography: Application to squeezed states and the vacuum," *Physical Review Letters* **70**, 1244–1247 (1993). Cited on page 41.

40. D. Leibfried, D. Meekhof, B. King, C. Monroe, W. Itano, and D. J. Wineland, "Experimental determination of the motional quantum state of a trapped atom," *Physical Review Letters* **77**, 4281–4285 (1996). Cited on page 41.

41. G. Nogues, A. Rauschenbeutel, S. Osnaghi, M. Brune, J. Raimond, and S. Haroche, "Seeing a single photon without destroying it," *Nature* **400**, 239–242 (1999). Cited on pages 44 & 141.

42. M. Brune, J. Bernu, C. Guerlin, S. Deleglise, C. Sayrin, S. Gleyzes, S. Kuhr, I. Dotsenko, J.-M. Raimond, and S. Haroche, "Process tomography of field damping and measurement of Fock state lifetimes by quantum nondemolition photon counting in a cavity," *Physical Review Letters* **101**, 240402 (2008). Cited on pages 45 & 46.

43. S. Haroche, M. Brune, and J.-M. Raimond, "Measuring photon numbers in a cavity by atomic interferometry: optimizing the convergence procedure," *J. Phys. II France* **2**, 659–670 (1992). Cited on page 46.

44. S. Gleyzes, S. Kuhr, C. Guerlin, J. Bernu, S. Deléglise, U. B. Hoff, M. Brune, J.-M. Raimond, and S. Haroche, "Quantum jumps of light recording the birth and death of a photon in a cavity," *Nature* **446**, 297–300 (2007). Cited on pages 46, 48, 127 & 137.

45. L. Landau, "Zur Theorie der Energieubertragung. II," *Phys. Z. Sowjetunion* **2**, 46–51 (1932). Cited on page 48.

46. C. Zener, "Non-adiabatic crossing of energy levels," *Proceedings of the Royal Society London, Series A* **137**, 696 (1932). Cited on page 48.

47. J. M. Chow, *Quantum Information Processing with Superconducting Qubits*. PhD thesis, Yale University, 2010. Cited on pages 53, 80, 84, 89, 98 & 104.

48. M. Mariantoni, F. Deppe, A. Marx, R. Gross, F. Wilhelm, and E. Solano, "Two-resonator circuit quantum electrodynamics: A superconducting quantum switch," *Phys. Rev. B* **78**, 104508 (2008). Cited on pages 54 & 60.

49. D. M. Pozar, *Microwave Engineering*. Wiley, 3rd ed., 2005. Cited on page 57.

50. P. Lafarge, *Macroscopic charge quantization in metallic nanostructures*. PhD thesis, Université Paris VI, 1993. Cited on page 63.

51. M. Boissonneault, J. M. Gambetta, and A. Blais, "Dispersive regime of circuit QED: Photon-dependent qubit dephasing and relaxation rates," *Physical Review A* **79**, 013819 (2009). Cited on page 64.

52. D. Esteve, M. H. Devoret, and J. M. Martinis, "Effect of an arbitrary dissipative circuit on the quantum energy levels and tunneling of a Josephson junction," *Phys. Rev. B* **34**, 158–163 (1986). Cited on pages 67, 68, 106 & 107.

53. A. Leggett, "Quantum tunneling in the presence of an arbitrary linear dissipation mechanism," *Phys. Rev. B* **30**, 1208–1218 (1984). Cited on page 67.

54. K. Cicak, D. Li, J. A. Strong, M. S. Allman, F. Altomare, A. J. Sirois, J. D. Whittaker, J. D. Teufel, and R. W. Simmonds, "Low-loss superconducting resonant circuits using vacuum-gap-based microwave components," *Appl. Phys. Lett.* **96**, 093502 (2010). Cited on page 72.

55. F. Wellstood, C. Urbina, and J. Clarke, "Low-frequency noise in DC superconducting quantum interference devices below 1 K," *Applied Physics Letters* **50**, 772 (1987). Cited on pages 77 & 81.

56. J. M. Chow, J. M. Gambetta, L. Tornberg, J. Koch, L. Bishop, A. A. Houck, B. R. Johnson, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, "Randomized benchmarking and process tomography for gate errors in a solid-state qubit," *Physical Review Letters* **102**, 090502 (2009). Cited on pages 84, 99 & 113.

57. V. Braginsky and V. Panov, "Superconducting resonators on sapphire," *IEEE Transactions on Magnetics* **15**, 30–32 (1979). Cited on pages 90 & 110.

58. A. O'Connell, M. Ansmann, R. Bialczak, M. Hofheinz, N. Katz, E. Lucero, C. McKenney, M. Neeley, H. Wang, and E. Weig, "Microwave dielectric loss at single photon energies and millikelvin temperatures," *Applied Physics Letters* **92**, 112903–112903 (2008). Cited on pages 90 & 116.

59. D. F. Santavicca and D. E. Prober, "Impedance-matched low-pass stripline filters," *Meas. Sci. Technol.* **19**, 087001 (2008). Cited on pages 96, 98 & 99.

60. D. H. Slichter, O. Naaman, and I. Siddiqi, "Millikelvin thermal and electrical performance of lossy transmission line filters," *Appl. Phys. Lett.* **94**, 192508 (2009). Cited on page 99.

61. J. Chow, L. DiCarlo, J. Gambetta, A. Nunnenkamp, L. Bishop, L. Frunzio, M. Devoret, S. Girvin, and R. Schoelkopf, "Detecting highly entangled states with a joint qubit readout," *Physical Review A* **81**, 062325 (2010). Cited on pages 99, 114 & 120.

62. J. M. Chow, L. DiCarlo, J. M. Gambetta, F. Motzoi, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, "Optimized driving of superconducting artificial atoms for improved single-qubit gates," *Physical Review A* **72**, 040305 (2010). Cited on page 99.

63. E. Bedrosian, "A product theorem for Hilbert transforms," *RAND Memorandum* (1962). Cited on page 103.

64. A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf, "Strong coupling of a single photon to a superconducting qubit using circtui quantum electrodynamics," *Nature* **431**, 162–167 (2004). Cited on pages 105 & 106.

65. P. Bertet, I. Chiorescu, G. Burkard, K. Semba, C. Harmans, D. Divincenzo, and J. Mooij, "Dephasing of a superconducting qubit induced by photon noise," *Physical Review Letters* **95**, 257002 (2005). Cited on page 105.

66. M. Steffen, M. Ansmann, R. Mcdermott, N. Katz, R. C. Bialczak, E. Lucero, M. Nee-ley, E. M. Weig, A. N. Cleland, and J. M. Martinis, "State tomography of capacitively shunted phase qubits with high fidelity," *Physical Review Letters* **97**, 050502 (2006). Cited on page 105.

67. J. A. Schreier, A. A. Houck, J. Koch, D. I. Schuster, B. R. Johnson, J. M. Chow, J. Gam-betta, J. Majer, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, "Sup-pressing charge noise decoherence in superconducting charge qubits," *Phys. Rev. B* **77**, 180502 (2008). Cited on pages 105, 107, 113, 115 & 127.

68. R. J. Schoelkopf and S. M. Girvin, "Wiring up quantum systems," *Nature* **451**, 664–669 (2008). Cited on pages 106 & 127.

69. P. Bertet, I. Chiorescu, G. Burkard, K. Semba, C. J. P. M. Harmans, D. P. DiVincenzo, and J. E. Mooij, "Relaxation and dephasing in a flux-qubit," arXiv:cond-mat/0412485v1 (2004). Cited on page 106.

70. J. M. Martinis, K. Cooper, R. Mcdermott, M. Steffen, M. Ansmann, K. D. Osborn, K. Cicak, S. Oh, D. Pappas, R. W. Simmonds, and C. Yu, "Decoherence in joseph-son qubits from dielectric loss," *Physical Review Letters* **95**, 210503 (2005). Cited on pages 106 & 116.

71. M. Neeley, M. Ansmann, R. C. Bialczak, M. Hofheinz, N. Katz, E. Lucero, A. O'Connell, H. Wang, A. N. Cleland, and J. M. Martinis, "Transformed dissipa-tion in superconducting quantum circuits," *Phys. Rev. B* **77**, 180508 (2008). Cited on page 106.

72. O. Astafiev, Y. Pashkin, Y. Nakamura, T. Yamamoto, and J. Tsai, "Quantum noise in the Josephson charge qubit," *Physical Review Letters* **93**, 267007 (2004). Cited on page 106.

73. E. Purcell, "Spontaneous emission probabilities at radio frequencies," *Phys. Rev.* **69**, 681 (1946). Cited on page 106.

74. P. Goy, J. M. Raimond, M. Gross, and S. Haroche, "Observation of cavity-enhanced single-atom spontaneous emission," *Phys. Rev. Lett.* **50**, 1903–1906 (1983). Cited on page 106.

75. D. Kleppner, "Inhibited spontaneous emission," *Phys. Rev. Lett.* **47**, 233–236 (1981). Cited on page 106.

76. R. G. Hulet, E. S. Hilfer, and D. Kleppner, "Inhibited spontaneous emission by a Ryd-berg atom," *Phys. Rev. Lett.* **55**, 2137–2140 (1985). Cited on page 106.

77. W. Jhe, A. Anderson, E. A. Hinds, D. Meschede, L. Moi, and S. Haroche, "Suppres-sion of spontaneous decay at optical frequencies: Test of vacuum-field anisotropy in confined space," *Phys. Rev. Lett.* **58**, 1497 (1987). Cited on page 106.

78. A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, "Approaching unit visibility for control of a superconducting qubit with dispersive readout," *Physical Review Letters* **95**, 060501 (2005). Cited on page 107.

79. L. Frunzio, A. Wallraff, D. Schuster, J. Majer, and R. Schoelkopf, "Fabrication and characterization of superconducting circuit QED devices for quantum computation," *IEEE Transactions on Applied Superconductivity* **15**, 860– 863 (2005). Cited on page 110.

80. V. B. Braginsky, V. S. Ilchenko, and K. S. Bagdassarov, "Experimental observation of fundamental microwave absorption in high-quality dielectric crystals," *Physics Letters A* **120**, 300–305 (1987). Cited on page 110.

81. J. M. Gambetta, W. Braff, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, "Protocols for optimal readout of qubits using a continuous quantum nondemolition measurement," *Physical Review A* **76**, 012325 (2007). Cited on pages 114 & 139.

82. A. Shnirman, G. Schön, I. Martin, and Y. Makhlin, "Low- and high-frequency noise from coherent two-level systems," *Phys. Rev. Lett.* **94**, 127002 (2005). Cited on page 116.

83. J. M. Martinis, M. Ansmann, and J. Aumentado, "Energy decay in superconducting Josephson-junction qubits from nonequilibrium quasiparticle excitations," *Phys. Rev. Lett.* **103**, 097002 (2009). Cited on page 116.

84. B. R. Johnson, M. D. Reed, A. A. Houck, D. I. Schuster, L. S. Bishop, E. Ginossar, J. M. Gambetta, L. DiCarlo, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, "Quantum non-demolition detection of single microwave photons in a circuit," *Nature Physics* **6**, 663– 667 (2010). Cited on page 117.

85. N. D. Mermin, *Quantum Computer Science*. Cambridge University Press, New York, 2007. Cited on page 117.

86. A. M. Steane, "Simple quantum error-correcting codes," *Phys. Rev. A* **54**, 4741–4751 (1996). Cited on page 126.

87. R. Ruskov and A. N. Korotkov, "Entanglement of solid-state qubits by measurement," *Phys. Rev. B* **67**, 241305 (2003). Cited on page 126.

88. R. Raussendorf and H. J. Briegel, "A one-way quantum computer," *Phys. Rev. Lett.* **86**, 5188–5191 (2001). Cited on page 126.

89. M. D. Reed, B. R. Johnson, A. A. Houck, L. DiCarlo, J. M. Chow, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf, "Fast reset and suppressing spontaneous emission of a superconducting qubit," *Appl. Phys. Lett.* **96**, 203110 (2010). Cited on page 133.

90. A. Lupaşcu, S. Saito, T. Picot, P. C. D. Groot, C. J. P. M. Harmans, and J. E. Mooij, "Quantum non-demolition measurement of a superconducting two-level system," *Nature Physics* **3**, 119–123 (2007). Cited on page 133.

91. M. D. Reed, L. DiCarlo, B. R. Johnson, L. Sun, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf, "High fidelity readout in circuit quantum electrodynamics using the Jaynes-Cummings nonlinearity," *Physical Review Letters* **105**, 173601 (2010). Cited on pages 136 & 142.

92. L. S. Bishop, E. Ginossar, and S. M. Girvin, "Response of the strongly-driven Jaynes-Cummings oscillator," *Physical Review Letters* **105**, 100505 (2010). Cited on pages 136 & 142.

93. L. Spietz, K. Irwin, and J. Aumentado, "Superconducting quantum interference device amplifiers with over 27 GHz of gain-bandwidth product operated in the 4–8 GHz frequency range," *Appl. Phys. Lett.* **95**, 092505 (2009). Cited on page 142.

94. N. Bergeal, F. Schackert, M. Metcalfe, R. Vijay, V. Manucharyan, L. Frunzio, D. E. Prober, R. J. Schoelkopf, S. M. Girvin, and M. H. Devoret, "Phase-preserving amplification near the quantum limit with a josephson ring modulator," *Nature* **465**, 64–68 (2010). Cited on page 142.

95. C. Noh, *Broadband Teleportation and Entanglement in Cascaded Open Quantum Systems*. PhD thesis, Univeristy of Auckland, 2009. Cited on page 144.

Mathematica code for Landau-Zener simulations

NUMERICAL code for solving the transmon–cavity time-dependent Schrödinger equation
follows. This calculation makes use of Lev Bishop's transmon package, which is listed in
the appendix of his thesis [16]. The first part constructs the full transmon–cavity Hamiltonian
for a finite number of levels. T. Felbinger's qmatrix package is used for convenience, to manage
the book keeping of writing operators that span composite Hilbert spaces. After constructing
the Hamiltonian, the code numerically solves the time-dependent Schrödinger equation as
the transmon E_J is changed. In order to analyze the results of the numerical solution, the
code also finds the instantaneous eigenstates for all values of E_J considered in the sweep.

Initialize

```
SetOptions[{Plot, ListLinePlot, ListPlot},
   {PlotRange → All, Frame → True, Axes → False, PlotRangePadding → None}];

$HistoryLength = 0;

<< qmatrix.m

Needs["transmon`"]

<< "/Users/bjohnson/Documents/Physics/calculations/Plot Exporting Options.m"

MediumPlots

<< "/Users/bjohnson/Documents/Physics/calculations/Legend.m"
```

Hamiltonian

Get transmon energies from Lev Bishop's transmon package.

```
{ef, gf} = makeinterp[.2, 20, 8, {2, 400, 5}];
```

Simulate 5 transmon and 5 cavity levels

```
nq = 5;
nc = 5;

setSystem[sys = {q1, cav}];
setModeType[#, {bosonic, nq}] & /@ {q1};
setModeType[cav, {bosonic, nc}];

{ao, ado} = {matrix[op[a, cav]], matrix[op[ad, cav]]};
HQfull[q_?modeQ, Ej_, Ec_, g_] :=
  matrix[
    DiagonalMatrix[Table[ef[i][Ej, Ec], {i, 0, nq - 1}]], {ket[q], bra[q]}] +
   g (ao + ado) ** matrix[Table[gf[i, j][Ej, Ec], {i, 0, nq - 1}, {j, 0, nq - 1}],
      {ket[q], bra[q]}];

HQrwa[q_?modeQ, Ej_, Ec_, g_] :=
  matrix[
    DiagonalMatrix[Table[ef[i][Ej, Ec], {i, 0, nq - 1}]], {ket[q], bra[q]}] +
   g (# + hc[#] &) @ (ao ** matrix[DiagonalMatrix[
        Table[gf[i, i + 1][Ej, Ec], {i, 0, nq - 2}], -1], {ket[q], bra[q]}]);

HQnrwa[q_?modeQ, Ej_, Ec_, g_] :=
 matrix[
   DiagonalMatrix[Table[ef[i][Ej, Ec], {i, 0, nq - 1}]], {ket[q], bra[q]}] +
  g (ao + ado) ** (# + hc[#] &) @ (matrix[DiagonalMatrix[
        Table[gf[i, i + 1][Ej, Ec], {i, 0, nq - 2}], -1], {ket[q], bra[q]}])

HH[wc_,
  {Ej1_, Ec1_, g1_}] :=
 wc ado ** ao + HQ[q1, Ej1, Ec1, g1]
```

```mathematica
orderm = {ket[cav], ket[q1], bra[cav], bra[q1]};
orderv = {ket[cav], ket[q1]};
tomatrix[m_?properMatrixQ] :=
  Developer`ToPackedArray@flatten[reorderMatrix[m, orderm]][[1]];
tovector[v_?properMatrixQ] := Developer`ToPackedArray@
  flatten[reorderMatrix[v, orderv]][[1]]

{H0full, H0rwa, H0nrwa} = Table[FullSimplify@tomatrix[HH[wc, {Ej1, Ec1, g1}]],
  {HQ, {HQfull, HQrwa, HQnrwa}}];

Block[{Ec1 = 0.317`, Ec2 = 0.2958`, Ej1 = 28.457`,
  g1 = 0.199`, g2 = 0.1833`, oo2 = -1.14`, oo = -1.035`,
  wc = 6.8934`, x0b = 20.5`, x0 = 2.87`, Ej2 = 42.465`},
 HermitianMatrixQ /@ {H0full, H0rwa, H0nrwa}]

{True, True, True}

barevec[c_, i_] := basisKet[cav, c + 1] ** basisKet[q1, i + 1];
closest[vecs_?(MatrixQ[#, NumericQ] &), v_?(VectorQ[#, NumericQ] &)] :=
  Ordering[Abs[v.vecsᵀ], -1][[1]];
```

Now define the valvec[ωc, Ej, Ec, g] function to return the eigenvalues and eigenvectors of the coupled transmon-cavity system.

```mathematica
Block[{
  H0this = H0full,
  wc, Ej1, Ec1, g1, comp, nl},
 With[{
   H0 = H0this,
   H0sp = SparseArray@H0this,
   sp = {wc, Ej1, Ec1, g1}},
  With[{ssp = Sequence @@ sp,
    DH = Transpose[D[SparseArray@evalinterp@H0, {sp}], {2, 3, 1}]},
   With[
    {C0 = Developer`ToPackedArray@
       evalinterp[H0sp /. HoldPattern@SparseArray[__, {__, a_}] :> a],
     DC0 = Developer`ToPackedArray[DH /.
       HoldPattern@SparseArray[__, {__, a_}] :> a]},
    valvec =
      Function[
       Evaluate@sp,
       Module[{eval, evec, oo},
        {eval, evec} = Eigensystem[H0];
        oo = Ordering[eval];
        eval = (# - #[[1]]) &@eval[[oo]];
        evec = evec[[oo]];
        evec = # / Norm@# & /@ evec;
        {eval, evec}
        ]];

   ]]]]

EjFromEc[Ec_, v_] := (Ec + v)^2 / (8 Ec)

Spectrum = Table[
  valvec[5.05, EjFromEc[0.25, v], 0.25, .075] // First, {v, 4.75, 6.5, .05}];
```

```
ListLinePlot[Spectrumᵀ, PlotRange → {4, 28}, DataRange → {4.75, 6.5},
  FrameLabel → {"Transmon ω₀₁ Frequency (GHz)", "Energy/ℏ (GHz)"},
  BaseStyle → {Thick}, LabelStyle → {FontSize → 14}]
```
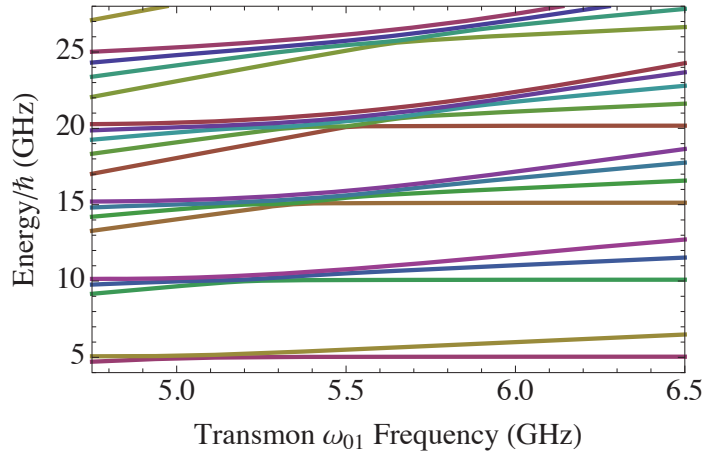

Landau - Zener

The Landau-Zener simulation numerically solves the time-dependent Schrödinger equation as the transmon Ej is changed in time. In order to view results, we also construct a function to return the instantaneous eigenvectors at any time point during the sweep.

- **Slow linear sweep**

```
csol = Block[{
    H0 = H0full,
    wc = 5.062,
    Ec1 = .300,
    g1 = .068,
    ν0 = 4.828,
    νf = 5.4,
    tf = 20},
  Module[{vals, vecs, X0},
   {vals, vecs} = valvec[wc, EjFromEc[Ec1, ν0], Ec1, g1];
   X0 = vecs[[2]];
   NDSolve[{

      ⅈ X'[t] ⩵ 2 π H0.X[t] /. Ej1 → EjFromEc[Ec1, ν0 + (νf - ν0) t/tf],

      X[0] ⩵ X0
      }, X, {t, 0, tf}, MaxSteps → 1 000 000] // First

   ]]

{X → InterpolatingFunction[{{0., 20.}}, <>]}
```

```mathematica
With[{
   H0 = H0full,
   wc = 5.062,
   Ec1 = .300,
   g1 = .068,
   ν0 = 4.828,
   νf = 5.4,
   tf = 20},
  sweepvecs = Function[t,
     Module[{vals, vecs, ν},

        ν = ν0 + (νf - ν0) t/tf;

        {vals, vecs} = valvec[wc, EjFromEc[Ec1, ν], Ec1, g1];

        vecs]
     ];
  ]
```

```mathematica
colors = {ColorData[20, 1], ColorData[20, 2]}
```

```mathematica
{RGBColor[0.745098, 0.0666667, 0.00392157],
 RGBColor[0.901961, 0.196078, 0.129412]}
```

```mathematica
With[{tf = 20},
 ListLinePlot[Table[{Abs[X[t * tf].sweepvecs[t * tf][[2]] /. csol]^2,
      Abs[X[t * tf].sweepvecs[t * tf][[3]] /. csol]^2},
     {t, 0, 1, 1 / 100.}]^T, PlotRange → {0, 1}, DataRange → {0, 1},
    FrameLabel → {"Time t/Δt", "Population |α|^2"}, PlotStyle → colors]
 ]
```
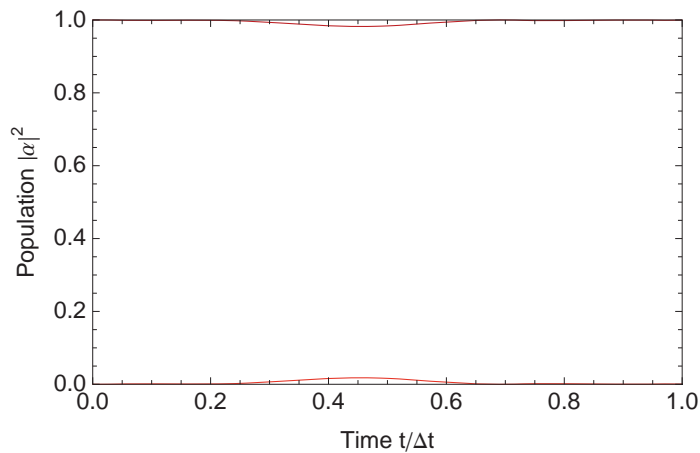

The final probability to stay in the same state is:

```mathematica
With[{tf = 20}, Abs[X[tf].sweepvecs[tf][[2]] /. csol]^2]
```

```mathematica
0.999434
```

Table of sweeps

Now we look at the dependence on the sweep duration.

```
sweepPts = {1, 5, 10, 15};
```

```mathematica
LZplots = Table[

   csol = Block[{

      H0 = H0full,
      wc = 5.062,
      Ec1 = .300,
      g1 = .068,
      ν0 = 4.828,
      νf = 5.4,
      tf = sweepPts[[ii]]},

     Module[{vals, vecs, X0},

      {vals, vecs} = valvec[wc, EjFromEc[Ec1, ν0], Ec1, g1];
      X0 = vecs[[2]];

      NDSolve[{

         ⅈ X'[t] ⩵ 2 π H0.X[t] /. Ej1 → EjFromEc[Ec1, ν0 + (νf - ν0) t/tf],

         X[0] ⩵ X0
         }, X, {t, 0, tf}, MaxSteps → 1 000 000] // First

      ]];

   With[{

      H0 = H0full,
      wc = 5.062,
      Ec1 = .300,
      g1 = .068,
      ν0 = 4.828,
      νf = 5.4,
      tf = sweepPts[[ii]]},

    sweepvecs = Function[t,

       Module[{vals, vecs, ν},

        ν = ν0 + (νf - ν0) t/tf;
        {vals, vecs} = valvec[wc, EjFromEc[Ec1, ν], Ec1, g1];
        vecs]

      ];

     ];
   (*colors={ColorData[20,2*sweeptime/5-1],ColorData[20,2*sweeptime/5]};*)
   colors = {Hue[(ii - 1) / Length[sweepPts]], Hue[(ii - .5) / Length[sweepPts]]};
   color = Hue[.9 * (ii - 1) / Length[sweepPts]];
   With[{tf = sweepPts[[ii]]},
    ListLinePlot[Table[{Abs[X[t * tf].sweepvecs[t * tf][[2]] /. csol]^2,
        Abs[X[t * tf].sweepvecs[t * tf][[3]] /. csol]^2},
       {t, 0, 1, 1 / 100.}]^⊤, PlotRange → {0, 1}, DataRange → {0, 1}, FrameLabel →
      {"Time t/Δt", "Population"}, PlotStyle → {color, {color, Dashed}}]
    ],

   {ii, 1, Length[sweepPts]}];
```

```
mylabels =
  Flatten[{"P₁,g Δt=" <> ToString[#] <> "ns", "P₀,e Δt=" <> ToString[#] <> "ns"} & /@
    sweepPts]
```

```
{P₁,g Δt=1ns, P₀,e Δt=1ns, P₁,g Δt=5ns, P₀,e Δt=5ns,
 P₁,g Δt=10ns, P₀,e Δt=10ns, P₁,g Δt=15ns, P₀,e Δt=15ns}
```

```
fig1a = AddLegend[Show[LZplots], mylabels,
  LegendLabelSide → Right, LegendItemWidth → .5]
```
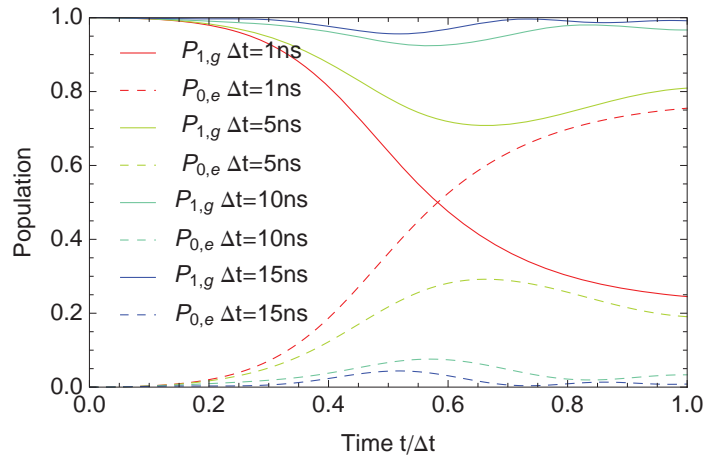
- **Final Probability vs sweep time**
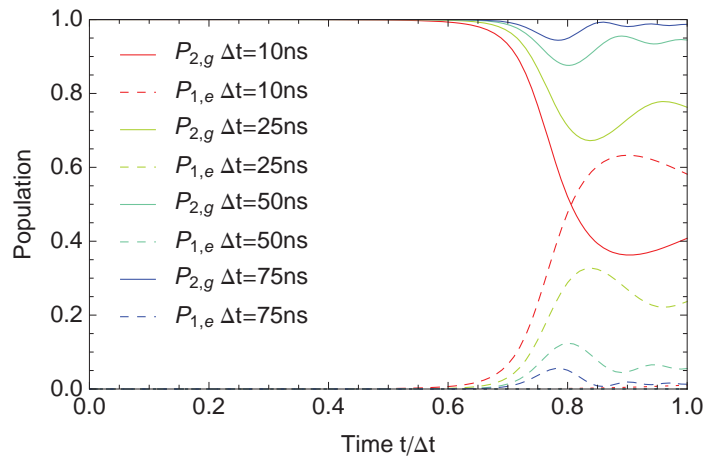
```mathematica
LZprobs = Table[

   csol = Block[{
       H0 = H0full,
       wc = 5.062,
       Ec1 = .300,
       g1 = .068,
       ν0 = 4.828,
       νf = 5.4},
     Module[{vals, vecs, X0},
       {vals, vecs} = valvec[wc, EjFromEc[Ec1, ν0], Ec1, g1];
       X0 = vecs[[2]];
       NDSolve[{

           i X'[t] == 2 π H0.X[t] /. Ej1 → EjFromEc[Ec1, ν0 + (νf - ν0) t/tf],

           X[0] == X0
           }, X, {t, 0, tf}, MaxSteps → 1 000 000] // First

       ]];
   With[{
       H0 = H0full,
       wc = 5.062,
       Ec1 = .300,
       g1 = .068,
       ν0 = 4.828,
       νf = 5.4},
     sweepvecs = Function[t,
         Module[{vals, vecs, ν},

             ν = ν0 + (νf - ν0) t/tf;
           {vals, vecs} = valvec[wc, EjFromEc[Ec1, ν], Ec1, g1];
           vecs]
         ];
     ];
   {Abs[X[tf].sweepvecs[tf][[2]] /. csol]^2,
     Abs[X[tf].sweepvecs[tf][[3]] /. csol]^2}
   , {tf, 1, 20, .25}];
```

```
fig1bsim = ListLinePlot[LZprobsᵀ, PlotRange → {0, 1}, DataRange → {1, 20},
    FrameLabel → {"Sweep Time Δt (ns)", "Final Population"},
    PlotStyle → {Hue[0], Hue[.65]}]
```
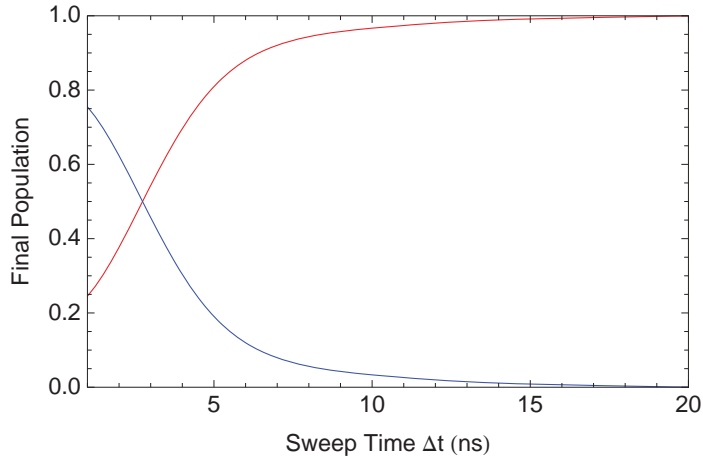

We can compare this result to the analytic Landau-Zener result with the same sweep rate.

```
5.4 - 4.828
```

```
0.572
```

```
fig1theory = Plot[Exp[- (2 g)² / Γ] /. {g → 2 π * .068, Γ → 2 π * 0.572 / Δt},
    {Δt, 1, 20}, PlotRange → {0, 1}, PlotStyle → {Hue[.65], Dashed}]
```
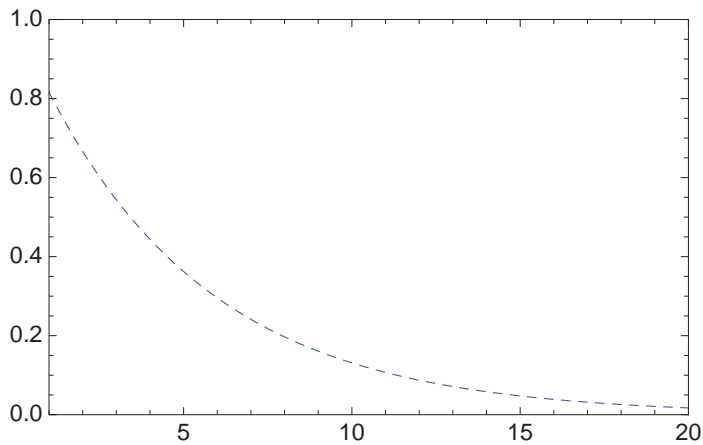
```
fig1b = AddLegend[Show[fig1bsim, fig1theory], {"P₁,g", "P₀,e", "LZ P₀,e"},
    LegendLabelSide → Left, LegendItemWidth → .6, LegendPosition → {Right, Top}]
```
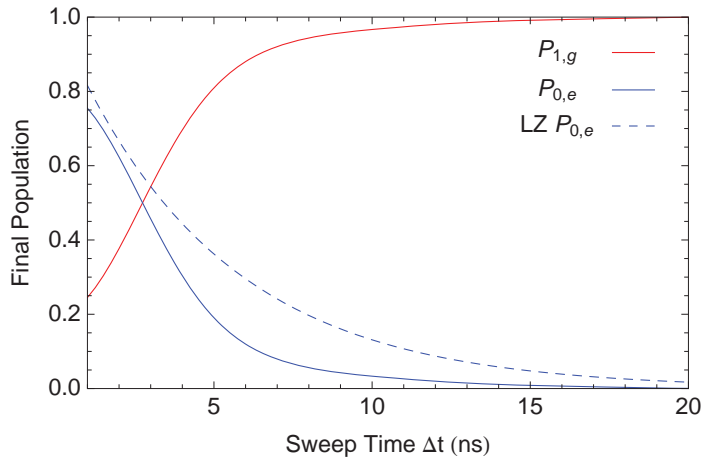


```
fig1 = GraphicsColumn[{fig1a, fig1b}]
```
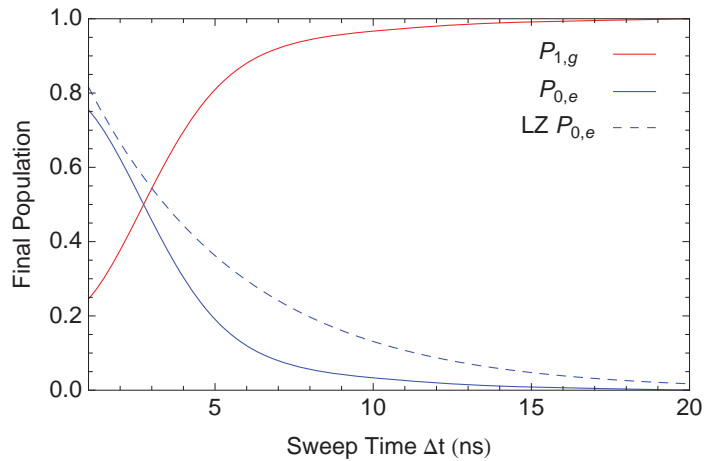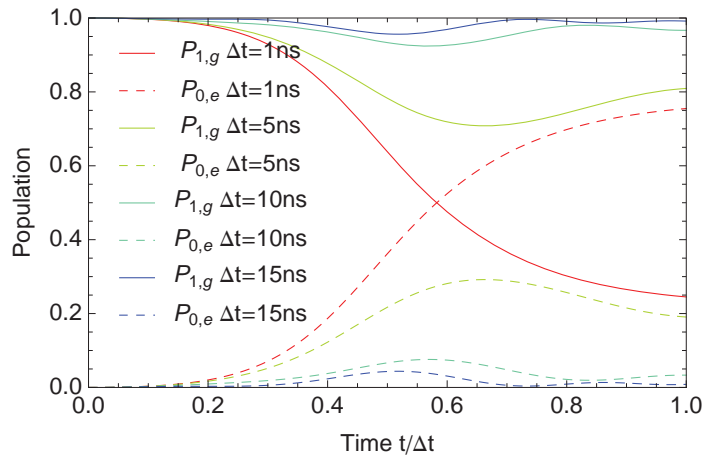
Prepare n=2 (transfer |0,f> to |2,g>)

```mathematica
sweepPts2 = {10, 25, 50, 75};

LZplots2 = Table[
   csol = Block[{
       H0 = H0full,
       wc = 5.062,
       Ec1 = .300,
       g1 = .068,
       v0 = 4.8,
       vf = 5.4,
       tf = sweepPts2[[ii]]},
     Module[{vals, vecs, X0},
       {vals, vecs} = valvec[wc, EjFromEc[Ec1, v0], Ec1, g1];
       X0 = vecs[[4]];
       NDSolve[{
           i X'[t] == 2 π H0.X[t] /. Ej1 → EjFromEc[Ec1, v0 + (vf - v0) t / tf],
           X[0] == X0
         }, X, {t, 0, tf}, MaxSteps → 1 000 000] // First
     ]];
   With[{
       H0 = H0full,
       wc = 5.062,
       Ec1 = .300,
       g1 = .068,
       v0 = 4.8,
       vf = 5.4,
       tf = sweepPts2[[ii]]},
     sweepvecs = Function[t,
         Module[{vals, vecs, v},
           v = v0 + (vf - v0) t / tf;
           {vals, vecs} = valvec[wc, EjFromEc[Ec1, v], Ec1, g1];
           vecs]
       ];
     ];
   color = Hue[.9 * (ii - 1) / Length[sweepPts]];
   With[{tf = sweepPts2[[ii]]},
    ListLinePlot[Table[{Abs[X[t * tf].sweepvecs[t * tf][[4]] /. csol]^2,
          Abs[X[t * tf].sweepvecs[t * tf][[5]] /. csol]^2,
          Abs[X[t * tf].sweepvecs[t * tf][[6]] /. csol]^2},
         {t, 0, 1, 1 / 100.}]^T, PlotRange → {0, 1}, DataRange → {0, 1},
       FrameLabel → {"Time t/Δt", "Population"},
       PlotStyle → {color, {color, Dashed}, {color, Dotted}}]
     ],
   {ii, 1, Length[sweepPts2]}];

mylabels2 =
   Flatten[{"P_{2,g} Δt=" <> ToString[#] <> "ns", "P_{1,e} Δt=" <> ToString[#] <> "ns",
       "P_{0,f} Δt=" <> ToString[#] <> "ns"} & /@ sweepPts2];

mylabels2 = Flatten[
    {"P_{2,g} Δt=" <> ToString[#] <> "ns", "P_{1,e} Δt=" <> ToString[#] <> "ns", None} & /@
     sweepPts2];
```

```
fig2a = AddLegend[Show[LZplots2],
    mylabels2, LegendLabelSide → Right, LegendItemWidth → .5]
```



```
With[{tf = 100}, Abs[X[tf].sweepvecs[tf][[4]] /. csol]]
```

0.998457

- **Final Probability vs Sweep Time**

```mathematica
LZprobs2 = Table[

    csol = Block[{

        H0 = H0full,
        wc = 5.062,
        Ec1 = .300,
        g1 = .068,
        ν0 = 4.828,
        νf = 5.4},

      Module[{vals, vecs, X0},

        {vals, vecs} = valvec[wc, EjFromEc[Ec1, ν0], Ec1, g1];
        X0 = vecs[[4]];

        NDSolve[{

            i X'[t] == 2 π H0.X[t] /. Ej1 → EjFromEc[Ec1, ν0 + (νf - ν0) t/tf],

            X[0] == X0
          }, X, {t, 0, tf}, MaxSteps → 1000000] // First

      ]];

    With[{

        H0 = H0full,
        wc = 5.062,
        Ec1 = .300,
        g1 = .068,
        ν0 = 4.828,
        νf = 5.4},

      sweepvecs = Function[t,

        Module[{vals, vecs, ν},

          ν = ν0 + (νf - ν0) t/tf;
          {vals, vecs} = valvec[wc, EjFromEc[Ec1, ν], Ec1, g1];

          vecs]

        ];

      ];

    {Abs[X[tf].sweepvecs[tf][[4]] /. csol]^2,
      Abs[X[tf].sweepvecs[tf][[5]] /. csol]^2,
      Abs[X[tf].sweepvecs[tf][[6]] /. csol]^2}

    , {tf, 1, 80, 1}];
```
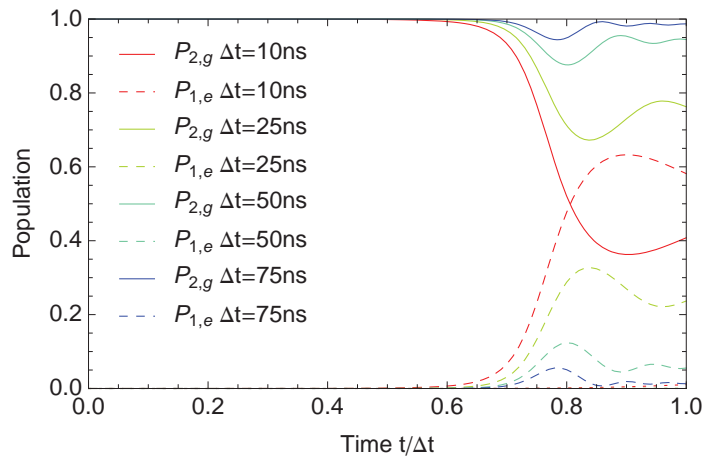
```
fig2b =
  AddLegend[ListLinePlot[LZprobs2ᵀ, PlotRange → {0, 1}, DataRange → {1, 80},
    FrameLabel → {"Sweep Time Δt (ns)", "Final Population"}],
   {"P₂,g", "P₁,e", "P₀,f"}, LegendLabelSide → Left,
   LegendItemWidth → .6, LegendPosition → {Right, Top}]
```
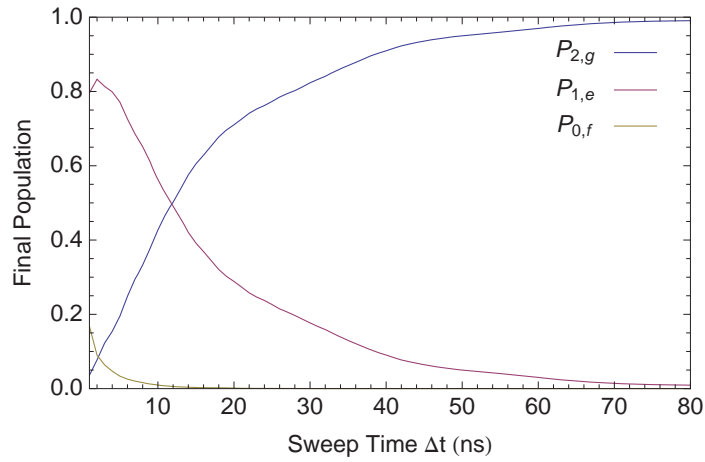
```
fig2 = GraphicsColumn[{fig2a, fig2b}]
```
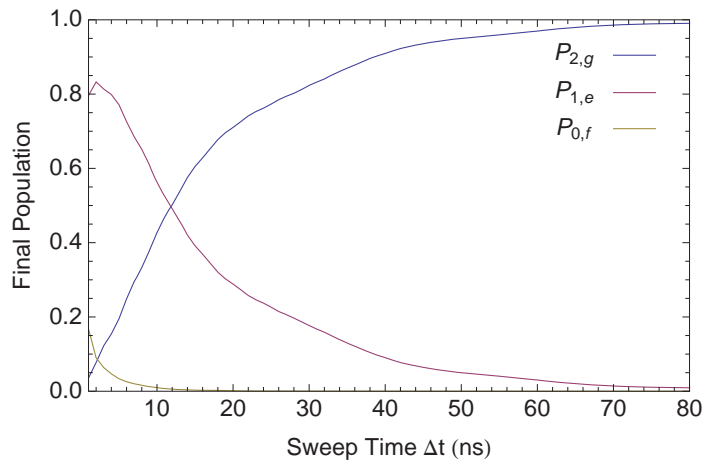

- **Export**

```
figurepath = "/Users/bjohnson/Documents/Physics/Thesis/figures/raw/";
```

```
Export[figurepath <> "LZ-sim1.pdf", fig1]
```

/Users/bjohnson/Documents/Physics/Thesis/figures/raw/LZ-sim1.pdf

```
Export[figurepath <> "LZ-sim2.pdf", fig2]
```

/Users/bjohnson/Documents/Physics/Thesis/figures/LZ-sim2.pdf

Mathematica code for pulse sequence generation

Tʜɪs appendix provides a Mathematica code listing for pulse sequence generation. The method used here specifies a domain-specific language for describing pulse sequences in terms of labels that identify pulses as well as lists that indicate how parameters change in an experiment. Pulses are defined using the function `Pulse[]`. This function takes an argument which is a pulse label, such as Xp or Y90p for π and $\pi/2$ rotations about the positive X- and Y-axes, respectively. The `Pulse[]` function accepts many more arguments (all of which are optional) to specify the amplitude, angle, width, duration (a pulse can exist in a time bin longer than the pulse itself), sigma (for Gaussian or hyperbolic tangent pulses), pulse type (square, Gaussian, etc.), and several other parameters. All of these optional arguments can be specified with a scalar value, or a list of values. A list of `Pulse[]`'s is passed to the `GetSeq[]` function along with a sequence number n, which produces the list of points corresponding to the concatenated pulses, selecting the nth value of any parameter that was specified with a list.

For instance, a Ramsey experiment is defined by two $\pi/2$-pulses about the same axis, with a variable time delay in the middle. To create a Ramsey pulse sequence, one creates a list of delay times (e.g. `DelayTimes = Range[0,1000,10]`), and then constructs the pulse list: `seq = {Pulse[X90p], Pulse[QId, width->DelayTimes], Pulse[X90p]}`, where QId speci-

fies the identity (do nothing) operation. The `seq` list is then passed to `GetSeq[]` to construct the pattern tables. See the last two sections of the code listing for complete examples.

Sequence Functions

```
Options[Pulse] = {amp → 0, width → SpecPulseLength,
    duration → same, angle → 0, sigma → GaussWidth, align → right,
    offset → 0, buffer → 0, delay → 0, pType → Gauss, index → 0,
    MotzoiScale → 0, SSBmod → 0, IQphase → 0, SSBmod2 → 0};
Pulse[p_, OptionsPattern[]] := Module[
  {a, w, dur, L, s, al, o, b, pulseType, pdelay, DACoffset,
   qubitPulses, fluxPulses, measPulses, pf, in, ms, ssbm, ssbm2},
  (* set defaults *)
  a = OptionValue[amp];
  w = OptionValue[width];
  dur = OptionValue[duration]; If[dur === same || dur < w, dur = w];
  L = OptionValue[angle];
  s = OptionValue[sigma];
  al = OptionValue[align];
  o = OptionValue[offset];
  b = OptionValue[buffer];
  pdelay = OptionValue[delay];
  pulseType = OptionValue[pType];
  pf = FullSpecPulse;
  in = OptionValue[index]; (* added by mdr/ldc *)
  ms = OptionValue[MotzoiScale]; (* added by mdr Jan 13, 2010 *)
  ssbm = OptionValue[SSBmod]; (* added by ldc Jan 17, 2010 *)
  ssbm2 = OptionValue[SSBmod2]; (* added by ldc Jan 21, 2010 *)

  qubitPulses = {QId, Xp, Xm, X90p, X90m, X45p, X45m, Yp, Ym, Y90p,
     Y90m, X45p, X45m, Xθ, Yθ, Xs, Ys, X90s, Y90s, Up, Um, U90p, U90m};
  fluxPulses = {ZId, Zf};
  measPulses = {MId, Ms};

  If[MemberQ[qubitPulses, p],
   DACoffset = DACoffsetX;
  ];
  If[MemberQ[fluxPulses, p],
   DACoffset = DACoffsetZ;
  ];
  If[MemberQ[measPulses, p],
   DACoffset = 0;
  ];

  (* Single Qubit Pulses *)
  If[p === QId, {a = o}];
  If[p === X90p, {L = 0, a = ampX90p}];
  If[p === X90m, {L = 0, a = ampX90m}];
  If[p === Y90p, {L = π / 2, a = ampY90p}];
```

```mathematica
If[p === Y90m, {L = π / 2, a = ampY90m}];
If[p === Xp, {L = 0, a = ampXp}];
If[p === Xm, {L = 0, a = ampXm}];
If[p === Yp, {L = π / 2, a = ampYp}];
If[p === Ym, {L = π / 2, a = ampYm}];

(* Single Qubit Arbitrary Rotation Pulses *)
If[p === Xθ, {L = 0}];
If[p === Yθ, {L = π / 2}];

(* Special angles *)
If[p === X45p, {L = 0,  a = ampX90p / 2}];
If[p === Y45p, {L = π / 2,  a = ampY90p / 2}];
If[p === X45m, {L = 0,  a = ampX90m / 2}];
If[p === Y45m, {L = π / 2,  a = ampY90m / 2}];

(* Arbitrary axis angle *)
If[p === Up, {a = ampXp}];
If[p === Um, {a = ampXm}];
If[p === U90p, {a = ampX90p}];
If[p === U90m, {a = ampX90m}];

(* Arbitrary length square pulses *)
If[p === Xs, {L = 0, a = ampXp, pulseType = Square}];
If[p === Ys, {L = π / 2, a = ampYp, pulseType = Square}];
If[p === X90s, {L = 0, a = ampX90p, pulseType = Square}];
If[p === Y90s, {L = π / 2, a = ampY90p, pulseType = Square}];

(* Single Qubit Z-Pulses *)
If[p === ZId, {a = o}];
If[p === Zf, {pf = FluxPulse}];

(* Measurement pulses *)
If[p === Ms, {a = 1; pulseType = Square}];
If[p === MId, {a = 0, pulseType = Square}];

Function[{n}, Module[{angle, amp, offset, width, sigma,
    delay, duration, buffer, quadratures, mScale, ssb, ssb2},
  {angle, amp, offset, width, sigma, delay, duration, buffer,
    mScale, ssb, ssb2} = (If[Head[#] === List, #[[n]], #]) & /@
    {L, a, o, w, s, pdelay, dur, b , ms, ssbm, ssbm2};
  (* adjust delay to align the pulse within its time window *)
  If[duration > width,
   Switch[al,
    right,  delay += 0,
    center, delay += Round[(duration - width) / 2],
    left,   delay += (duration - width)
    ]
```

```
    ];

    quadratures = RotationMatrix[angle].
      pf[amp, width, sigma, pulseType, mScale, ssb, ssb2];

    {offset + PadLeft[Join[N@quadratures[[1]], Delay[delay]], duration],
     offset + PadLeft[Join[N@quadratures[[2]], Delay[delay]], duration],
     DACoffset, duration, buffer}
    (*{phase,amp,offset,width,sigma,delay,duration}*)
    ]
   ]
  ]

Options[GetSeq] = {modf → 0, phaseError → 0, IQscale → 1};
GetSeq[seq_, n_, delay_, fPoint_, cLength_, OptionsPattern[]] :=
 Module[{xpattern, ypattern, xpulse, ypulse,
   offset, duration, buffer, toffset, f, phase, iqs},
  toffset = 0; xpattern = 0; ypattern = 0;
  Scan[
   ({xpulse, ypulse, offset, duration, buffer} = #[n];
     xpattern += MakePattern[xpulse, fPoint + delay - toffset, {}, cLength];
     ypattern += MakePattern[ypulse, fPoint + delay - toffset, {}, cLength];
     toffset += buffer + duration;) &,
   Reverse[seq]];
  f = OptionValue[modf];
  phase = OptionValue[phaseError];
  iqs = OptionValue[IQscale];
  If[f ≠ 0,
   {xpattern, ypattern} = SSB[FastRound[xpattern], FastRound[ypattern], f]];
  If[phase ≠ 0, {xpattern, ypattern} =
    PhaseCorrect[xpattern, ypattern, N[phase]]];
  If[iqs ≠ 1, {xpattern, ypattern} =
    DiagonalMatrix[{iqs, 1}].{xpattern, ypattern}];
  {NormalizePattern[xpattern, offset], NormalizePattern[ypattern, offset]}
 ]

ConcatenateSeqs[seqs__] := MapThread[Join, {seqs}]

Options[AddPulses] = {DACoffset → DACoffsetX};
AddPulses[p__, OptionsPattern[]] :=
 Apply[Plus, {p} - OptionValue[DACoffset]] + OptionValue[DACoffset]

NormalizePattern =
  Compile[{{pat, _Real, 1}, offset}, Round[Clip[pat + offset, {0, 2^14 - 1}]]];

FastRound = Compile[{{pat, _Real, 1}}, Round[pat]];

Convert14to10bit =
  Compile[{{pat, _Real, 1}}, Round[Clip[(pat - 8192) / 2^4 + 512, {0, 1023}]]];
```

```
Options[BalancePattern] = {endLength → 100};
BalancePattern[pat_List,
  offset_?NumberQ, opts : OptionsPattern[BalancePattern]] :=
 Block[{total, amp, numpoints, elength},
  elength = OptionValue[endLength];
  total = Total[pat - offset];
  amp = If[total > 0, -offset + 100, (2^14 - 1) - offset - 100];
  numpoints = Floor@Abs[total / amp];
  Join[
   Take[pat, Length[pat] - numpoints - elength - 1], (* replace numpoints +
    final corrector + elength offset levels at the end *)
   ConstantArray[offset + amp, numpoints],
   {offset - (amp * numpoints + total)},
   ConstantArray[offset, elength]
  ]
 ]
```

Pattern Generation

```
MakePattern::usage =
  "MakePattern[leftPat_,fixedPoint_,rightPat_,totalLength_]:\n
This function makes it easy to create patterns with a fixed point.\n
leftPat is pattern data before fixed point\n
fixedPoint is the location to be fixed (leftPat will be left padded)\n
rightPat is pattern data after the fixed point\n
the whole pattern will
    be padded to the right to a length = totalLength";

MakePattern[leftPat_, fixedPoint_, rightPat_, totalLength_] :=
  PadRight[Join[PadLeft[leftPat, fixedPoint], rightPat], totalLength];
```

Pulse Functions

- **General**

```
FluxPulse::unknown = "Unknown pulse type `1`.";
FluxPulse[amp_, length_, sigma_,
    type_: Square , MotzoiScale_: 0, SSBmod_: 0, SSBmod2_: 0] :=
  Switch[type,
   Gauss, {GaussianPulse[amp, length, sigma], ConstantArray[0, length]},
   Exp, {ExpPulse[amp, sigma, length], ConstantArray[0, length]},
   GaussOn, {GaussOn[amp, length, sigma], ConstantArray[0, length]},
   GaussOnOff,
   {Join[GaussOn[amp, 3 * sigma, sigma], SquarePulse[amp, length - 6 * sigma],
     GaussOff[amp, 3 * sigma, sigma]], ConstantArray[0, length]},
   GaussOff, {GaussOff[amp, length, sigma], ConstantArray[0, length]},
   Square, {SquarePulse[amp, length], ConstantArray[0, length]},
   Tanh,
   {TanhPulse2[amp, sigma, sigma, 2, length], ConstantArray[0, length]},
   Over, {OverPulse[amp, sigma, 1, length], ConstantArray[0, length]},
   Ramp, {RampPulse[amp, length], ConstantArray[0, length]},
   (*TanhB,
   {TanhPulse3[amp,sigma, sigma, 2, length],ConstantArray[0,length]},*)
   _, Message[FluxPulse::unknown, type]
   ];

FullSpecPulse::unknown = "Unknown pulse type `1`.";
FullSpecPulse[amp_, length_, sigma_,
    type_: Gauss , MotzoiScale_: 0, SSBmod_: 0, SSBmod2_: 0] :=
  Switch[type,
   Gauss, {GaussianPulse[amp, length, sigma], ConstantArray[0, length]},
   Exp, {ExpPulse[amp, length, sigma], ConstantArray[0, length]},
   Square, {SquarePulse[amp, length], ConstantArray[0, length]},
   Tanh, {TanhPulse2[amp, sigma, sigma, 3, length], ConstantArray[0, length]},
   Deriv, {GaussianPulse[amp, length, sigma],
    DerivGaussianPulse[MotzoiScale * amp, length, sigma]},
   Motzoi, StarkyPulse[amp , length, sigma, MotzoiScale, SSBmod, SSBmod2],
   GaussOn, {GaussOn[amp, length, sigma], ConstantArray[0, length]},
   GaussOff, {GaussOff[amp, length, sigma], ConstantArray[0, length]},
   GaussOnOff,
   {Join[GaussOn[amp, 3 * sigma, sigma], SquarePulse[amp, length - 6 * sigma],
     GaussOff[amp, 3 * sigma, sigma]], ConstantArray[0, length]},
   _, Message[FullSpecPulse::unknown, type]
   ];
```

- **Motzoi**

```
GetPhase[amp_, n_, σ_, ssb_] := Module[{midpoint, Ex, PhaseVec, FreqVec},
  midpoint = (n + 1) / 2;
  Ex = Table[ Exp[- (x - midpoint)^2 / (2 σ^2) ], {x, 1, n}];

  FreqVec = Table[2 * Pi * Ex[[x]]^2 (amp / amp180)^2 * ssb / 1000 , {x, 1, n}];
  (*instantaneous angular frequency *)
  PhaseVec = Accumulate[FreqVec];
  N[PhaseVec[[n]]] ]
```

```
GetPhaseTable[amp_, n_, σ_, ssb_] := (GetPhase[#, n, σ, ssb] &) /@ amp;
```

```
StarkyPulse[amp_, n_, σ_, scaleD_, ssb_, ssb2_] :=
 Module[{midpoint, Ex, Exdot, PhaseVec, FreqVec, IPhvec, QPhvec, OutVec},
  midpoint = (n + 1) / 2;
  Ex = Table[amp Exp[- (x - midpoint)^2 / (2 σ^2) ], {x, 1, n}];

  Exdot = Table[- amp * (x - midpoint) / σ^2  * Exp[- (x - midpoint)^2 / (2 σ^2) ], {x, 1, n}];
  IPhvec = Ex;
  QPhvec = scaleD * Exdot;
  OutVec = {IPhvec, QPhvec};

  FreqVec = Table[

    2 * Pi * ( (Ex[[x]] / amp180)^2 * ssb + (Exp[- (x - midpoint)^2 / (2 σ^2) ])^2 * ssb2) / 1000,

    {x, 1, n}];
  PhaseVec = Accumulate[FreqVec];

  OutVec = {IPhvec * Cos[PhaseVec] - QPhvec * Sin[PhaseVec],
    QPhvec * Cos[PhaseVec] + IPhvec * Sin[PhaseVec]};

  Round[OutVec] ]
```

```
SimpleMotzoiPulse[amp_, n_, σ_, scaleD_, ssb_] :=
 Module[{midpoint, Ex, Exdot, PhaseVec, IPhvec, QPhvec, OutVec},
  midpoint = (n + 1) / 2;

  Ex = Table[amp Exp[- (x - midpoint)^2 / (2 σ^2)], {x, 1, n}];

  Exdot = Table[- amp * (x - midpoint) / σ^2 * Exp[- (x - midpoint)^2 / (2 σ^2)], {x, 1, n}];

  IPhvec = Ex;
  QPhvec = scaleD * Exdot;
  OutVec = {IPhvec, QPhvec};
  PhaseVec = Table[2 * Pi * ssb / 1000 * x, {x, 1, n}];

  OutVec = {IPhvec * Cos[PhaseVec] + QPhvec * Sin[PhaseVec],
    QPhvec * Cos[PhaseVec] - IPhvec * Sin[PhaseVec]};

  Round[OutVec]]
```

```
MotzoiPulse[quad_, amp_, n_, σ_, order_,
  lambda_, delta_, lever_, doquad_, dophase_] :=
 Module[{rate, prefacA1, prefacA3, prefacA5, prefacB1, prefacB3,
   prefacC2, prefacC4, midpoint, Ex, Exdot, IPhvec, QPhvec,
   DetVec, PhaseVec, OutVec},

  rate = 10^9;

  (* First, prepare the prefactors *)
  prefacA1 = 1;
  prefacA3 = 0;
  prefacA5 = 0;
  prefacB1 = 0;
  prefacB3 = 0;
  prefacC2 = 0;
  prefacC4 = 0;
  If[order ≥ 3, prefacA3 = (lambda^2 - 4) / 2 * (lever / delta)^2];
  If[order ≥ 5,
   prefacA5 = - (13 * lambda^4 - 76 * lambda^2 + 112) / 8 * (lever / delta)^4];
  If[order ≥ 1 && doquad == 1, prefacB1 = - rate / delta / (2 * Pi * 10^6)];
  If[order ≥ 3 && doquad == 1,
   prefacB3 = rate * 33 * (lambda^2 - 2) / 6 * lever^2 / delta^3 / (2 * Pi * 10^6)];
  If[order ≥ 2 && dophase == 1, prefacC2 =
    1 / rate * (lambda^2 - 4) * (lever^2 / delta) * (2 * Pi * 10^6)];
  If[order ≥ 4 && dophase == 1, prefacC4 = -1 / rate *
     (lambda^4 - 7 * lambda^2 + 12) * lever^4 / delta^3 * (2 * Pi * 10^6)];

  (* Second, calculate gaussian pulses and their derivative *)
  midpoint = (n + 1) / 2;

  Ex = Table[amp Exp[- (x - midpoint)^2 / (2 σ^2)], {x, 1, n}];

  Exdot = Table[- amp * (x - midpoint) / σ^2 * Exp[- (x - midpoint)^2 / (2 σ^2)], {x, 1, n}];

  (*
  IPhvec=prefacA1*Ex +prefacA3*Ex^3+prefacA5*Ex^5;
  QPhvec=prefacB1*Exdot+prefacB3*Ex^2*Exdot;
  DetVec=prefacC2*Ex^2+prefacC4*Ex^4;
  PhaseVec=Accumulate[DetVec];
  *)

  IPhvec = prefacA1 * Ex;
  QPhvec = prefacB1 * Exdot;
  DetVec = prefacC2 * Ex^2;
  PhaseVec = Accumulate[DetVec];

  OutVec = Which[quad == 0,
    IPhvec * Cos[PhaseVec * dophase] - QPhvec * Sin[PhaseVec * dophase], quad == 1,
    QPhvec * Cos[PhaseVec * dophase] + IPhvec * Sin[PhaseVec * dophase]];
  ;

  Round[OutVec] ]
```

Pulse Shapes

```
SquarePulse[amp_, n_] := amp × Table[1, {n}];
```

```mathematica
Delay[n_] := Table[0, {ii, 1, n}];

GaussianPulse[amp_, n_, σ_] := Module[{midpoint},
  midpoint = (n + 1) / 2;
  Round[Table[amp Exp[-(x - midpoint)^2/(2 σ^2)], {x, 1, n}]]]

GaussOn[amp_, n_, σ_] := Round[Table[amp Exp[-(x - n)^2/(2 σ^2)], {x, 1, n}]]

GaussOff[amp_, n_, σ_] := Round[Table[amp Exp[-(x)^2/(2 σ^2)], {x, 1, n}]]

BufferPulseGen[list_, zeroLevel_, padding_, reset_, delay_] :=
  Module[{result},
   (*buffer to the left*)
   result = PadRight[ListCorrelate[ConstantArray[1, 1 + padding], list],
     Length[list], (1 + padding) * zeroLevel];
   (*buffer to the right*)
   result = PadLeft[ListConvolve[ConstantArray[1, 1 + padding], result],
     Length[list], (1 + padding)^2 * zeroLevel];
   (*convert to on/off*)
   result = If[# ≠ (1 + padding)^2 * zeroLevel, 1, 0] & /@ result;
   (*keep the pulse high if the
    delay between pulses is less than the reset time*)
   result = Flatten[Module[{l = Length[#]},
         ConstantArray[If[l < reset, 1, First[#]], l]] & /@ Split[result]];
   (*shift to the left by the delay amount*)
   result = PadRight[Drop[result, delay], Length[list]]
  ];
IQBufferPulseGen[Ipts_, Qpts_, zeroLevel_, padding_, reset_, delay_] :=
  BufferPulseGen[#, 0, padding, reset, delay] & /@
   MapThread[If[#1 ≠ zeroLevel || #2 ≠ zeroLevel, 1, 0] &, {Ipts, Qpts}, 2];

DerivGaussianPulse[amp_, n_, σ_] := Module[{midpoint},
  midpoint = (n + 1) / 2;
  Round[Table[- amp (x - midpoint)/σ^2 Exp[-(x - midpoint)^2/(2 σ^2)], {x, 1, n}]]]
```

```mathematica
ExpPulse[amp_, σ_, length_] :=
  Round[Table[amp * Exp[-t / σ], {t, 0, length - 1}]];

TanhPulse2[amp_, σup_, σdown_, numσ_, length_] := Module[{t0, t1},
    t0 = numσ * σup + 1;
    t1 = length - numσ * σdown;
    Round[Table[amp (Tanh[(t - t0) / σup] + Tanh[- (t - t1) / σdown]) / 2,
      {t, 1, Round[length]}]]
  ];

OverPulse[amp_, overA_, overL_, length_] :=
  Round[Join[Table[(1 + overA / 100) * amp, {t, 1, Round[overL]}],
    Table[amp , {t, 1, Round[length - 2 * overL]}],
    Table[(-overA / 100) * amp, {t, 1, Round[overL]}]
    ]];

RampPulse[amp_, length_] :=
 Round[Table[amp * t / length, {t, 1, length}]]

(* Measurement Pulse *)
MeasPulse[MeasPulseLength_, fixedPoint_, cycleLength_] :=
    MakePattern[{}, fixedPoint, SquarePulse[1, MeasPulseLength], cycleLength];
MeasPulseArray[n_, MeasPulseLength_, fixedPoint_, cycleLength_] :=
    Join @@ Table[MeasPulse[MeasPulseLength, fixedPoint, cycleLength], {n}];
```

▪ **Single Sideband Modulation**

```mathematica
(* modf in units of MHz *)
SSB::usage =
  "SSB[Ipat, Qpat, modf] returns the SSB modulated signal at frequency modf
    (in units of MHz when the time step is ns). Use positive modf to get
    the upper sideband and negative modf to get the lower sideband.";
SSB[Ipat_ /; Depth[Ipat] - 1 == 1, Qpat_ /; Depth[Qpat] - 1 == 1, modf_] :=
  Module[{t, f},
    t = Range[0, Length[Ipat] - 1];
    f = modf / 1000.0;
    SSBC[Ipat, Qpat, t, f]
  ];
SSB[Ipat_ /; Depth[Ipat] - 1 == 2, Qpat_ /; Depth[Qpat] - 1 == 2, modf_] :=
  Transpose[MapThread[SSB[#1, #2, modf] &, {Ipat, Qpat}]];
SSBC = Compile[{{Ipat, _Real, 1}, {Qpat, _Real, 1}, {t, _Real, 1}, f},
    {Ipat * Cos[2 π * f * t] + Qpat * Sin[2 π * f * t],
     Qpat * Cos[2 π * f * t] - Ipat * Sin[2 π * f * t]}
    ];
```

▪ **Phase Correction**

```mathematica
PhaseCorrect::usage =
  "PhaseCorrect[Ipat, Qpat, phase] corrects phase errors from I/Q
    mixers. 'phase' is defined in radians as the offset from π/2
    separation between the quadratures. phase > 0 means that the
    separation is < π/2, phase < 0 means that the separation > π/2.";
PhaseCorrect = Compile[{{Ipat, _Real, 1}, {Qpat, _Real, 1}, phase},
    {Ipat - Qpat * Tan[phase], Qpat * Sec[phase]}];
```

Example Usage

```
numsteps = 80;
rabistep = 17 / 2.;
fluxstep = -2 ^ 14 / 80.;
RabiPoints = Range[0, rabistep * (numsteps - 1), rabistep];
AnglePoints = Range[0, 2 π, 2 π / (numsteps - 1)];
fluxPoints = Range[0, fluxstep * (numsteps - 1), fluxstep];

qubitSeq = {
   Pulse[Xθ, amp → RabiPoints],
   Pulse[X90p, align → center, duration → 50],
   Pulse[U90p, angle → AnglePoints],
   Pulse[QId, width → 30]
  };
fluxSeq = {
   Pulse[ZId],
   Pulse[Zf, amp → fluxPoints, width → 50, pType → Square],
   Pulse[ZId],
   Pulse[Zf, amp → fluxPoints, width → 30]
  };

{patTableCh1, patTableCh2, patTableCh3, null} = Transpose[Table[
    Join[
     GetSeq[qubitSeq, ii, SpecPulseDelay, fixedPointCh1, cycleLength],
     GetSeq[fluxSeq, ii, FluxPulseDelay, fixedPointCh1, cycleLength]
    ]
    , {ii, numsteps}]];

plotnum = 70;
ListPlot[{patTableCh1[[plotnum]],
  patTableCh2[[plotnum]], patTableCh3[[plotnum]]}, Joined → True,
 PlotRange → {{6800, 7100}, {0, 2 ^ 14 - 1}}, Axes → False, Frame → True]
```
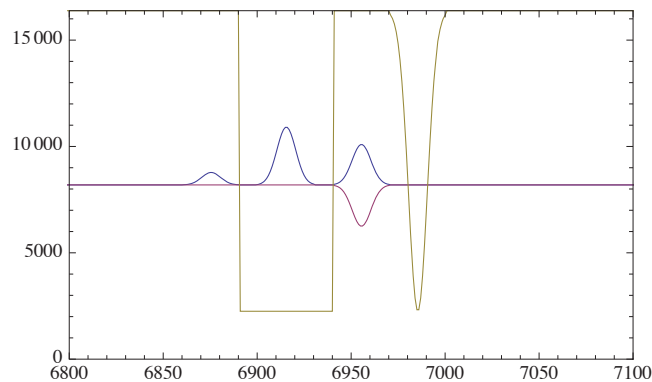
Full Rabi Ramsey Experiment Example

```mathematica
basename = "RabiRamsey";
datapathAWG5014 = RootAWG5014 <> basename <> "_5014\\";
CreateDirectory[datapathAWG5014];
filePrefix5014 = basename <> "_5014";

numsteps = 80;

RabiCenter = 0; RabiStep = 17;
RabiPoints =
   Table[RabiCenter + (ii - numsteps / 2) * RabiStep, {ii, 0, numsteps - 1}];
RamseyStep = 5;
RamseyPoints = Range[0, RamseyStep * (numsteps - 1), RamseyStep];

Exp1QubitSeq = {
    Pulse[Xθ, amp → RabiPoints]
    };
Exp1FluxSeq = {
    Pulse[ZId]
    };

Exp2QubitSeq = {
    Pulse[X90p],
    Pulse[QId, width → RamseyPoints],
    Pulse[X90p]
    };
Exp2FluxSeq = {
    Pulse[ZId],
    Pulse[ZId, width → RamseyPoints],
    Pulse[ZId]
    };

{patTableCh1Exp1, patTableCh2Exp1, patTableCh3Exp1, null} = Transpose[Table[
    Join[
     GetSeq[Exp1QubitSeq, ii, SpecPulseDelay, fixedPointCh1, cycleLength],
     GetSeq[Exp1FluxSeq, ii, FluxPulseDelay, fixedPointCh1, cycleLength]
     ]
     , {ii, numsteps}]];

{patTableCh1Exp2, patTableCh2Exp2, patTableCh3Exp2, null} = Transpose[Table[
    Join[
     GetSeq[Exp2QubitSeq, ii, SpecPulseDelay, fixedPointCh1, cycleLength],
     GetSeq[Exp2FluxSeq, ii, FluxPulseDelay, fixedPointCh1, cycleLength]
     ]
     , {ii, numsteps}]];
```

```mathematica
(* Join the experiments *)

patTableCh1 = ConcatenateSeqs[patTableCh1Exp1, patTableCh1Exp2];
patTableCh2 = ConcatenateSeqs[patTableCh2Exp1, patTableCh2Exp2];
patTableCh3 = ConcatenateSeqs[patTableCh3Exp1, patTableCh3Exp2];
patTableCh4 = Table[ConstantArray[0, 2 * cycleLength], {ii, numsteps}];

(* AWG5014 Markers *)

(* Marker 1 is the measurement buffer *)
markerTable1 = Table[
    MeasPulseArray[2, MeasPulseLength,
     fixedPointCh1 + MeasPulseDelay, cycleLength], {ii, numsteps}];
(* Marker 2 is the spec buffer for Qubit 1 Channels *)
markerTable2 = IQBufferPulseGen[patTableCh1, patTableCh2,
    DACoffsetX, specBuffer, specBufferReset, specBufferDelay];
(* Marker 3 is a trigger *)
markerTable3 = Table[
    PadRight[Join[Delay[1], SquarePulse[1, triggerWidth]], 2 * cycleLength],
    {ii, numsteps}];
(* Marker 4 is the population buffer *)
markerTable4 = Table[ConstantArray[0, 2 * cycleLength], {ii, numsteps}];
(*1-BufferPulseGen[#,DACoffsetZ,specBuffer,specBufferReset,MeasPulseDelay]&/@
  patTableCh3;*)

(* Markers 5-8 are triggers *)
markerTable5 = markerTable6 = markerTable7 = markerTable8 = markerTable3;

thisNum = 60;
ListPlot[
 {patTableCh1[[thisNum]], patTableCh2[[thisNum]], patTableCh3[[thisNum]],
  800 * markerTable1[[thisNum]], 800 * markerTable2[[thisNum]]},
 Joined → True, PlotRange → {{15 750, 17 500}, All}, Axes → False, Frame → True]
```
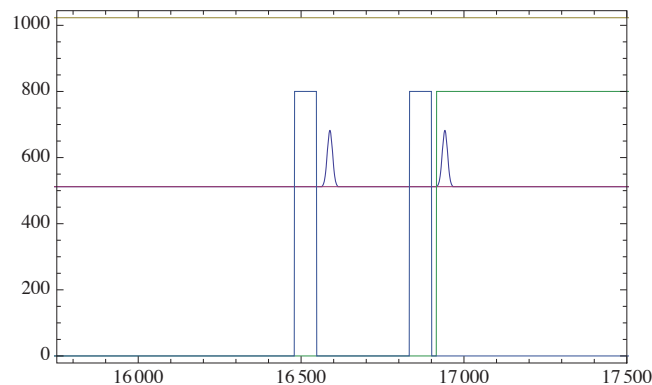


```mathematica
(* Export pulses *)
steps = Range[numsteps];
ExportTekPatternsAWG5014[datapathAWG5014, filePrefix5014,
 patTableCh1, markerTable1, markerTable2, patTableCh2, markerTable3,
 markerTable4, patTableCh3, markerTable5, markerTable6,
 patTableCh4, markerTable7, markerTable8, steps, parameterList]

basename
DateString[]
```

APPENDIX C

Fabrication recipes

This appendix provides the detailed electron beam lithography and deposition recipes.

Table C.1: PMMA/MMA bilayer with aluminum cap

1. Clean chip with **NMP, acetone**, and **methanol** by sonicating in each solvent for 60 seconds.

2. Spin **MMA EL13** at 4000 rpm for 60s

3. Bake at $170°C$ for 60s (under petri dish propped up on slides and using local thermometer to measure temperature)

4. Spin **PMMA 950k A3** at 4000 rpm for 60s

5. Bake at $170°C$ for 30min

6. Load into Playssys deposition system and pump the minimum time allowed to open the load lock to the main chamber (~ 45 minutes).

7. Deposit 12 nm Al at ~ 1 nm/s.

Table C.2: PMMA/MMA development

1. Shake chip in **MF312** for 60s or until aluminum is no longer visible anywhere on the surface

2. Shake chip back and forth in **MIBK:IPA 1:3** for 50s

3. Immediately place chip in **IPA** for 10s

4. Dry with N_2 gas

Table C.3: Electron beam evaporation

1. Pump out chamber to 1.5 x 10^{-7}, takes 4 hours

2. Clean the device with 3.5:1 argon:oxygen plasma with power of 250 mW/cm².

3. With the shutter closed, deposit titanium for 5 min at 0.2 nm/s to lower the chamber pressure

4. Deposit 20 nm aluminum at 1 nm/s

5. Oxidize with 85:15 argon:oxygen at a pressure of 15 Torr for 12 minutes

6. Pump down to 1.5 x 10^{-7} Torr in load lock and 7 x 10^{-8} Torr in chamber

7. Deposit 90 nm aluminum at 1 nm/s with 36 degree tilt

8. Oxidize with 85:15 argon:oxygen at 3 Torr for 10 minutes

9. Vent

Copyright Permissions

- Figure 2.2 reproduced with permission from:
  Jens Koch, Terri Yu, Jay Gambetta, Andrew A Houck, David I Schuster, J Majer, Alexandre Blais, M Devoret, S Girvin, and Robert J Schoelkopf, *Phys. Rev. A* **76,** 042319 (2007).
  Copyright (2007) by the American Physical Society.

- Figure 2.10 reproduced (with minor formatting modifications) with permission from:
  Michel Brune, Julien Bernu, Christine Guerlin, Samuel Deleglise, C Sayrin, Sébastien Gleyzes, Stefan Kuhr, I Dotsenko, Jean-Michel Raimond, and Serge Haroche, *Physical Review Letters* **101,** 240402 (2008).
  Copyright (2008) by the American Physical Society.

- Figure 2.4 reproduced with permission from:
  David I Schuster, Andreas Wallraff, Alexandre Blais, L Frunzio, R.-S Huang, J Majer, Steve M Girvin, and Robert J Schoelkopf, *Physical Review Letters* **94,** 123602 (2005).
  Copyright (2005) by the American Physical Society.

- Figure 6.1 reproduced with permission from:
  Andrew A Houck, J A Schreier, Blake R Johnson, J M Chow, Jens Koch, Jay Gambetta, David I Schuster, L Frunzio, Michel H Devoret, and Steve M Girvin, *Phys. Rev. Lett.* **101,** 80502 (2008).
  Copyright (2008) by the American Physical Society.

- Figure 6.2 reproduced with permission from:
  Andrew A Houck, J A Schreier, Blake R Johnson, J M Chow, Jens Koch, Jay Gambetta, David I Schuster, L Frunzio, Michel H Devoret, and Steve M Girvin, *Phys. Rev. Lett.* **101,** 80502 (2008).
  Copyright (2008) by the American Physical Society.