ABSTRACT

Title of Dissertation: DEVELOPING A STRONG DRIVING TOOLKIT FOR FLOQUET SYSTEMS WITH SUPERCONDUCTING QUBITS

Martin Ritter
Doctor of Philosophy, 2024

Dissertation Directed by: Professor Alicia J Kollár
Department of Physics

Simple systems such as a spin 1/2 particle driven by periodic drives can exhibit surprisingly rich physics. These systems can be described as lattices in phase space, in analogy with spatially periodic systems. By varying the phase and amplitude of the drives one can synthesize arbitrary complex hopping terms where the dimensionality of the effective lattice is set by the number of drives. In this thesis, we explore how to construct a 2-dimensional synthetic lattice with a topological band structure and study the effect of dissipation on the steady state of a strongly driven system.

Topological band structures are well known to produce symmetry-protected chiral edge modes which transport particles unidirectionally. The half-BHZ (Bernevig–Hughes–Zhang) model, defined as a 2-d lattice of two-level systems, exhibits edge modes in the limit that the hopping exceeds the on-site energy splitting. We synthesize this model by coupling a qubit to a cavity and driving it with a large external effective magnetic field. In the limit

of strong driving, the Floquet lattice can be topologically non-trivial, where the analog of a topologically protected edge state is a topologically protected energy pump. We have developed a toolkit for generating and characterizing the large synthetic magnetic fields required to reach the topological regime.

In the second part of the thesis, we explore a surprising result discovered in the process of characterizing the large synthetic fields: the stabilization of Floquet states in the presence of dissipation. While dissipation generally leads to errors in quantum systems, it can also lead to the stabilization of specific states of a driven system. This concept has been extensively studied in the context of optical pumping in atomic systems to drive systems to a pure ground state. Here we show that cavity-induced loss can have the dramatic and unexpected consequence of purifying a mixed state to a dynamically driven Floquet state.

Developing a Strong Driving Toolkit for Floquet Systems with Superconducting Qubits

by

Martin Ritter

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2024

Advisory Committee:

Professor Alicia J. Kollár, Chair
Dr. Benjamin S. Palmer
Professor Steve Rolston
Professor Nathan Schine
Professor Ronald Walsworth, Dean's Representative

Dedication

To all those who wonder *why*

Acknowledgements

Building up a lab from scratch during a pandemic was an experience and I cannot overstate how much I am indebted to so many people for their help, support, creativity, persistence, and companionship. I cannot list all the ways in which people have helped me here but I hope to provide an overview of the truly incredible people I have interacted with in these past six years.

Firstly, I would like to thank my advisor, Alicia Kollár for her support and guidance during my PhD. I have learned to be a *much* better writer through her feedback and learned that sometimes, you just have to try something even if its not perfectly planned out. Even though COVID sent us remote right as we were getting started, our daily zoom calls to discuss theses and develop control code kept us chugging along.

I would like to thank my committee members, Professors Alicia Kollár, Steve Rolston, Nathan Schine, and Dr. Ben Palmer for taking the time to provide me with valuable feedback and pushing me to be a better scientist. In parallel, I would like to thank both the ARCS Metro Washington chapter as well as the Laboratory for Physical Sciences for their generous fellowships which allowed me to do much of the work presented in this thesis.

I would also like to thank my lab mates for making coming to work that much more fun, especially when the measurements were not working. As the first two students in the lab, Maya and I formed a great team breaking up the challenge into bite-size pieces. I focused on the software and hardware infrastructure while Maya developed all the fabrication recipes, her wirebonding skills are truly legendary and I knew that I could always count on her to fix my devices. Building up a lab is just as much about the infrastructure

as is it about the lab culture. I would like to thank Zhiyin and Theo for starting up the bike gang and slowly converting the rest of the lab to the "gotta go fast" mantra, Kellen for all of our extremely insightful discussions on the metro about the lab, research, and life in general, Won and Billy for making our daily runs to the dining hall a consistent event. I had the privilege of mentoring three excellent undergraduates who joined our lab: Jeffrey, Theo, and Ruthie who all pushed me to understand the physics of cQED better. I would also like to thank IBK, Ben Cochran for their contributions to the lab as well as Rodney Snyder and Sandesh Kalantre who taught Maya and I so much about fabrication in the early days.

I would not have stuck around for so long if it weren't for the amazing group of people I met at the University of Maryland (to list some): Adam, Jake and Shaalini, Evan and Lindsay, Nick, Zach and Emily, Stefano and Grace, Chris, Dhruvit, Alex and Em, Matt, and Jared and Morgan. From starting a flag football IM dynasty to exploring the national parks around Washington DC, from board game nights to "deep" philosophical discussions after steak and whiskey nights, from lightning gelato runs to weekend walks to Zeke's coffee, from signing up to triathlons to yearly Friendsgiving, you made all these adventures possible.

Finally, I would like to thank my family for always being there for me. Even though we were usually an ocean away, I looked forward to our weekly Skype calls to just chat and catch up. Lastly, I would not be the person I am today without my wonderful partner in crime, Allie. Your steadfast support and almost endless patience as I was away in Maryland kept me motivated even when the measurements were not cooperating.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Since the invention of the transistor nearly 70 years ago, drastic improvements in computing power have radically increased the range of problems that can be practically studied. In particular, the increase in computational power has enabled the fully digital simulation of physical systems to incredible accuracy levels. One of the most striking examples of this computational might, is the increase in weather forecast accuracy: we can predict the weather 5 days out with a higher accuracy today than single-day forecasts from 30 years ago. However, not all systems can be efficiently simulated on a classical architecture. In particular, *quantum mechanical systems* composed of entangled particles display an exponential complexity in the computational resources required to model them.

This naturally begs the question: why not use a quantum system to perform quantum simulation? This idea was first presented by Feynman in 1982 and, for a decade, drew very little serious consideration as the task of building such a computer seemed completely out of reach with problems ranging from coherently controlling the qubits to correcting errors arising from decoherence to building systems with enough qubits to do anything interesting. Dramatic improvements in quantum technology starting with trapped ions [1], followed by superconducting qubits [2], then, in the last decade, neutral atoms [3] have revolutionized the field and started making serious inroads towards a fully digital *quantum computer*.

Additionally, the quantum nature of this new proposed computer allows it to solve problems that are computationally inefficient for classical computers. Famously, the problem of factoring numbers, the backbone of modern RSA cryptography, can be efficiently solved on a quantum computer with Shor's algorithm [4]. The existence of Shor's algorithm has lead to the search for other so-called *quantum advantage* algorithms which offer an exponential speedup over their classical counterparts.

However, current quantum systems are still far from realizing a *practical* digital quantum computer due to the relatively small number of qubits and the difficulties surrounding error correction. In some ways, the current quantum computing landscape is reminiscent of the early days of classical computing where the theory foundation for a new computational paradigm exists but the hardware is still in its infancy. While a fully digital system can in principle simulate *any* system, specialized hardware designed to solve a *specific* problem can be a much more efficient way to solve it. In the early days of digital computing, analog computers were often used for ballistics calculations, to solve partial differential equations (Water Integrator), or even to simulate the hydrodynamics of the North sea (Deltar, Ishiguro storm surge computer). The same principle can be applied for quantum mechanical systems: to simulate a specific system, we do not require a fully digital computer, just a system that reproduces the dynamics we wish to study.

In our lab, we use superconducting circuits to study the different physics arising in lattice systems.

1.1 Experimental Platform

Our experimental platform consists of superconducting circuits to simulate atomic and photonic modes described with circuit quantum electrodynamics. Circuit quantum electrodynamics (cQED) has emerged in the last decade as one of the premier quantum computing and simulation platforms due to its fast, high fidelity gates and rich tunability of circuit

topologies for Hamiltonian engineering [5, 6]. The basic building blocks for a cQED system are a qubit and a microwave photon. Both the qubit and the photonic mode are implemented using superconducting circuits which, while macroscopic in size and consisting of billions of individual atoms, behave as simple quantum mechanical systems. The photonic mode is typically implemented as a microwave resonator made from a coplanar waveguide cavity. The qubit mode is effectively an artificial atom whose energy level structure can be tailored by changing the underlying circuit topology. The most commonly used qubit, the transmon qubit [7], is implemented using a capacitor and a non-linear inductor and displays an anharmonic oscillator level structure which isolates the lowest two levels. Other qubit designs such as the fluxonium qubit (first developed in [8]) or the zero-pi qubit [9], include additional inductances and capacitances to sculpt a much richer level structure and can provide protection from decoherence. The interaction between the qubit and photon is described by the Jaynes-Cummings Hamiltonian which has been extensively studied in the context of atoms trapped in optical cavities. Circuit QED experiments have light-matter coupling that is orders of magnitude higher than cavity QED experiments allowing for a wider range of systems to be studied.

cQED systems are typically analyzed using simple circuit elements such as capacitors and inductors along with Josephson junctions which imbue the circuit with nonlinearity. Harmonic oscillators, qubits, and coupling structures are decomposed into their circuit representation; the resulting circuit can then be directly lithographically defined on a superconducting thin film allowing for virtually any planar circuit topology to be realized. In addition to freedom in circuit topology, the control electronics used in the microwave domain allow for direct wave synthesis leading to a high degree of control over qubit pulses for gates and time varying Hamiltonians.

In this thesis, we consider two kinds of simulated systems: photonic lattices (studied in detail in [10]) in which a spatial array of resonators is used to synthesize effective band structures and Floquet lattices where the action of strong periodic drives leads to an

effective lattice description.

Photonic lattices Photonic lattices are realized by arranging microwave resonators in a regular network and coupling their endpoints. In contrast to regular spatial lattices where the "site" is a 0-dimensional point and the interaction strength depends on the distance between sites, these microwave photonic lattices have extended "sites" and point-like interactions. This inversion of scales allows for the synthesis of non-Euclidean lattice tilings on Euclidean substrates. In particular, the freedom to deform the photonic resonator without affecting its resonant frequency or coupling strength to neighboring sites has been used to create a true representation of a hyperbolic lattice on a chip. These lattices can be used to study the effect of non-commuting spatial symmetries as well as study other simulation paradigms involving hyperbolic spin systems.

In addition to non-Euclidean tilings, the extended nature of the lattice sites leads to the emergence of flat bands in the band structure. Flat bands have long been of interest due to their large scale degeneracy and non-trivial ground states. Adding qubits to the photonic lattice allows for a measurement of the band structure and photonic modes. By tuning the qubit frequency into resonance with a particular lattice band and mapping out the photon-mediated qubit-qubit interactions we can study the effect of different band types (flat, linear, quadratic). This is studied in Chapter 5 where we develop a tunable coupler prototype to provide in-situ tunability of both the qubit frequency and the qubit-resonator coupling allowing us to isolate specific bands in a photonic lattice.

Floquet Systems While spatial lattices offer a large range of possible band structures, the native capacitive and inductive coupling are limited to real hopping terms. Magnetic fields (real or synthetic) have long been a source of topology in materials with early demonstrations such as the integer quantum hall effect. In addition to topological bands, magnetic fields are also the source of virtually all non-reciprocal devices such as circulators and isolators, which are a fundamental component of experimental setups. These fields can be synthesized on a lattice by adding a complex phase on hopping terms so that the phase

enclosed around a plaquette is non-trivial. This is usually achieved either by parametric modulation of the coupling between sites or by using magnetic materials to directly break time-reversal symmetry. However, this can also by realized in a Floquet lattice arising from strong, periodic driving of a qubit. The full amplitude and phase control of the microwave drives applied on the qubit allow us to synthesize lattices with complex hopping terms.

Historically, the half-BHZ model is considered as a minimal model for a 2-dimensional lattice with topological bands. This model consists of a square lattice of two level systems where the topological regime is achieved when the hopping terms exceed the on-site level splitting. This topological regime leads to the generation of chiral, uni-directional propagating modes. By synthesizing this model in the Floquet regime, the signature of non-trivial topology is an energy pump with a quantized rate of energy exchange between drives. This work has recently been extended to "boosting" where replacing one of the classical drives with a quantum state leads to a translation in the Fock basis of the underlying state. In Chapter 7, we develop a framework to synthesize the required fields on a driven transmon.

In the process of characterizing the driving fields, we discovered a persistent oscillation in the qubit population extending for much longer than any of the experimental decoherence timescales and discuss it in Chapter 8. While driven-dissipative systems have long been studied in the context of initial state preparation and stabilization, they usually involved either a tailored dissipator, which can be difficult to implement experimentally, or lead to a lab frame ground state such as in optical pumping. The system we studied in Chapter 7 consists of a qubit adiabatically following a rotating field coupled to a lossy cavity. At long times, the system is driven to Floquet ground state (i.e. "aligned" with the rotating field). This result lead to a generalized theoretical treatment extending the stabilization of Floquet states outside of the adiabatic regimes shown in the experiment opening up a large class of problems to be studied in other platforms.

1.2 Thesis Overview

This dissertation starts by introducing the theoretical framework for circuit QED in Chapter 2. We introduce the Jayne-Cummings Hamiltonian, circuit quantization (in which an electrical circuit is mapped onto an effective Hamiltonian), the transmon qubit, and qubit state readout.

In Chapter 3, we provide a full overview of the measurements used to characterize and control qubits. This includes spectroscopy, coherent control, qubit readout, and qubit characterization.

Chapter 4 introduces the experimental hardware used for measurements. We discuss the basic operating principles of the dilution fridge, provide an overview of the microwave equipment, and briefly discuss the control code architecture developed for the lab.

In Chapter 5, we develop a tunable coupler to selectively interact with a specific band in a photonic lattice without introducing defects due to frequency shifts in the resonator mode. We present a proof of concept experimental implementation and discuss a circuit model to reduce the frequency shift on the resonator mode.

Chapter 6 provides a brief introduction to the synthetic fields used in the final portion of the thesis. We discuss how to synthesize effective magnetic fields on a transmon using an amplitude modulated drive and frequency modulation on the qubit.

Chapters 7 and 8 develop an experimental framework to synthesize the required fields and study the dynamics of a qubit driven with a rotating magnetic field.

Chapter 2

Circuit Quantum Electrodynamics

In this chapter, I will lay out the building blocks for all our devices and present the simple theoretical framework underlying circuit quantum electrodynamics (cQED) [2, 11]. At their core, all of our devices can be modeled as harmonic (photonic modes) and anharmonic (atomic/ qubit modes) quantum oscillators. To build up a theory model for a given design, all one has to do is decompose the resulting circuit into these two types of quantum systems and compute the various couplings between the different subsystems to build up the system Hamiltonian. The rich topology available to circuit models allows for a wide range of Hamiltonians to be constructed using only a few circuit elements: capacitors, inductors, and a Josephson junction. Over the last two decades of development in the field, these three simple building blocks have led to the emergence of a veritable zoo of qubit designs, novel amplifiers, and exotic coupling schemes between photonic and atomic modes. A full description of these various improvements is outside the scope of this chapter; however, the interested reader will find the following publications a fruitful starting point [5, 6, 11, 12]. I will start with a theory overview of the Jaynes-Cummings Hamiltonian before setting up the circuit quantization framework to build Hamiltonians from electrical circuits. This will lead us to discuss the coplanar waveguide resonator and transmon as photonic and qubit modes respectively.

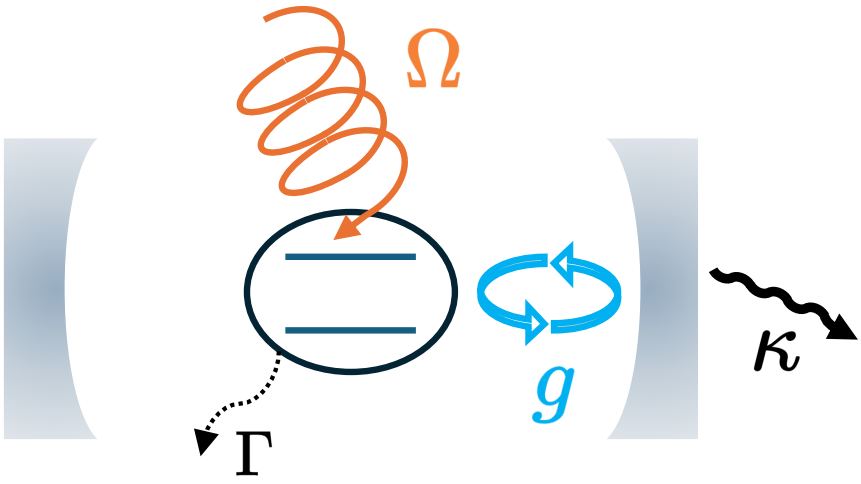

Figure 2.1: The simple picture we will consider for the theory section of this chapter: an atom approximated as a two level system(circled pair of lines) is coupled to a cavity (illustrated with the two curved mirrors) with a strength g. Both the atom and the cavity can experience dissipative processes indicates by the wavy lines leaving the system (Γ corresponding to depolarization and decoherence of the atomic state and κ corresponding to photon losses). The atom can also be driven with a resonant drive of strength Ω for control.

2.1 The Jaynes-Cummings Hamiltonian

The Jaynes-Cummings (JC) Hamiltonian was developed to describe the interaction between an atom and a photonic mode near resonance(schematically shown in Fig. 2.1 [13]). The atomic mode is usually approximated as a two-level system with transition frequency ω_q and the photonic mode is modeled as a simple quantum harmonic oscillator (QHO) mode with resonance frequency ω_c. The interaction between the atomic and photonic mode is $\overrightarrow{d} \cdot \overrightarrow{E}$ where we need to provide a quantum description for each term. Quantizing the electric field from the cavity mode allows us to rewrite $\overrightarrow{E}$ as $\mathcal{E}_{RMS}(a + a^\dagger)$ where $\mathcal{E}_{RMS}$ is the root-mean squared field of a single photon in the cavity mode. The dipole operator for a two level system is given by $d(\sigma_- + \sigma_+)$, where d is the matrix element of the position operator $|\langle g| e\hat{x} |e\rangle|$ and $\sigma_\pm$ are the usual raising and lowering operators for a two-level system. Combining all the terms, we obtain the general form of the JC Hamiltonian:

$$H = \omega_c \left(a^\dagger a + \frac{1}{2}\right) + \frac{\omega_q \sigma_z}{2} + d\mathcal{E}_{RMS}(\sigma_- + \sigma_+)(a + a^\dagger) \tag{2.1}$$

we can simplify this expression by removing terms that violate energy conservation to first order. Intuitively, we expect the de-excitation of the atom (σ_-) to be associated with the creation of a photon ($a^\dagger$) and vice versa. We can also explicitly show this by going into the interaction frame with the following unitary transformation $U = e^{it(\omega_c a^\dagger a + \omega_q \frac{\sigma_z}{2})}$. The resulting Hamiltonian H' is given by:

$$H' = UHU^\dagger + i\dot{U}U^\dagger \tag{2.2}$$

$$= g\left(e^{it\omega_q\frac{\sigma_z}{2}}(\sigma_- + \sigma_+)e^{-it\omega_q\frac{\sigma_z}{2}}\right)\left(e^{it\omega_c a^\dagger a}(a + a^\dagger)e^{-it\omega_c a^\dagger a}\right) \tag{2.3}$$

where we've simplified the diagonal terms $\propto a^\dagger a, \sigma_z$ as they commute with the transformation U. The coupling term can be simplified using the Baker-Campbell-Hausdorff (BCH) formula for exponentiated operators

$$e^X Y e^{-X} = \sum_m \frac{1}{m!}[X, [X, ..., [X, Y]...]] = \sum \frac{1}{m!}[X, Y]_m. \tag{2.4}$$

Recall $[a^\dagger a, a] = -a$ and $[\sigma_z, \sigma_\pm] = 2\sigma_\pm$ and after some algebra we can simplify H':

$$H' = g(e^{i\omega_q t}\sigma_+ + e^{-i\omega_q t}\sigma_-)(e^{i\omega_c t}a^\dagger + e^{-i\omega_c t}a) \tag{2.5}$$

which oscillates at two frequencies: $|\omega_c + \omega_q|$ and $|\omega_c - \omega_q|$. The Rotating Wave Approximation (RWA) assumes that the rapidly oscillating terms average out to zero relative to the slow oscillations at $|\omega_c - \omega_q|$. Connecting this to our intuitive approach, we can see that the difference frequency occurs for a de-excitation of one system and excitation of the other keeping the total number of excitations constant. Note, the RWA is only valid in the regime where $|\omega_c - \omega_q| \ll |\omega_c + \omega_q|$ and $g \ll \omega_c, \omega_q$ i.e. the detuning between the atom and the cavity is not so large that the sum and difference frequencies are comparable and g is small enough that higher order terms are suppressed. Writing our final Hamiltonian in the

interaction picture:

$$H' = g(e^{i(\omega_c-\omega_q)t}a^\dagger\sigma_- + e^{i(\omega_q-\omega_c)t}a^\dagger\sigma_-) \tag{2.6}$$

which can be rotated back to the lab-frame Hamiltonian using $U^\dagger$.

This effectively reduces the Hamiltonian to constant particle number sectors ($\langle a^\dagger a\rangle + \langle\sigma_z\rangle$) set by the initial state preparation of the system. The simplified Hamiltonian now reads:

$$H = \omega_c a^\dagger a + \frac{\omega_q}{2}\sigma_z + g(a\sigma_+ + a^\dagger\sigma_-) \tag{2.7}$$

where g is $d\mathcal{E}_{\mathcal{RMS}}$ is the coupling strength in units of frequency. This Hamiltonian can be diagonalized into 2×2 sectors for each number of excitations. The dressed state energies are:

$$E_{\overline{-,n}} = n\omega_c - \frac{1}{2}\sqrt{(\omega_c-\omega_q)^2 + 4ng^2} \tag{2.8}$$

$$E_{\overline{+,n}} = n\omega_c + \frac{1}{2}\sqrt{(\omega_c-\omega_q)^2 + 4ng^2} \tag{2.9}$$

with corresponding eigenvectors:

$$|-,n\rangle = \cos\left(\tfrac{\theta_n}{2}\right)|g,n\rangle - \sin\left(\tfrac{\theta_n}{2}\right)|e,n-1\rangle \tag{2.10}$$

$$|+,n\rangle = \cos\left(\tfrac{\theta_n}{2}\right)|e,n-1\rangle + \sin\left(\tfrac{\theta_n}{2}\right)|g,n\rangle \tag{2.11}$$

where $|\pm\rangle$ are the eigenvectors for upper and lower energy branches respectively, $\tan(\theta_n) = \frac{2g\sqrt{n}}{\Delta}$, $\Delta = \omega_c - \omega_q$ is the detuning between the qubit and the cavity, and n is the number of photons in the cavity. Near resonance, $\Delta \to 0$ leading to $\theta_n \approx \pi/2$, the two dressed states are the symmetric and antisymmetric combination of $|g,n\rangle$ and $|e,n-1\rangle$ and are split by $2g\sqrt{n}$, the coupling strength. An atom-cavity system initialized in the $|e,0\rangle$ state will spontaneously undergo oscillations, called vacuum Rabi oscillations, between $|g\rangle$ and $|e\rangle$ as the excitation is exchanged between the atomic component ($|e,0\rangle$) and the photonic component ($|g,1\rangle$).

Far from resonance, we can expand the trigonometric functions to first order to get the corrections to the bare $|g, n\rangle$ and $|e, n-1\rangle$ states:

$$|\overline{g, n}\rangle = |g, n\rangle - \frac{g\sqrt{n}}{\Delta}|e, n-1\rangle \tag{2.12}$$

$$|\overline{e, n-1}\rangle = |e, n-1\rangle + \frac{g\sqrt{n}}{\Delta}|g, n\rangle \tag{2.13}$$

where we have used $\overline{|\quad\rangle}$ to indicate that the dressed states are very close to the bare, uncoupled states.

In the limit of a "good" cavity ($\kappa < g$), the cavity spectral lineshape will change the spontaneous emission of the qubit, either enhancing or suppressing the decay rate. Qualitatively, the electromagnetic spectrum of the noise inside the cavity inducing spontaneous emission in the atom is sharply peaked at the cavity resonance frequency but suppressed everywhere else. For a far-detuned cavity-atom system, this can increase the lifetime by a few orders of magnitude relative to the free-space lifetime [14, 15]. This effect is called the *Purcell effect* and we can derive a simple intuitive model using the dispersive eigenstates defined above: if the cavity is lossy with rate κ, a system initialized in $\overline{|e, 0\rangle}$ will have its photonic portion decay overtime, slowly driving it to the $|g, 0\rangle$ state. We can compute the effective qubit decoherence rate due to the cavity by squaring the photonic contribution (to extract the probability of the system being in the $|g, 1\rangle$ state) and multiplying it by the photon decay rate:

$$\Gamma_{purcell} \approx \frac{g^2}{\Delta^2}\kappa \tag{2.14}$$

where κ is the cavity linewidth and $\frac{g^2}{\Delta^2}$ is the overlap with the photonic mode. The Purcell loss sets an upper limit on the coherence time of the atom in the cavity. An atom whose lifetime is limited by this cavity-induced decay rate is said to be *Purcell limited* [16]. This can be important to know when trying to increase qubit coherence times as the dominant issue might not be the qubit itself but instead the linewidth of the cavity it is coupled to.

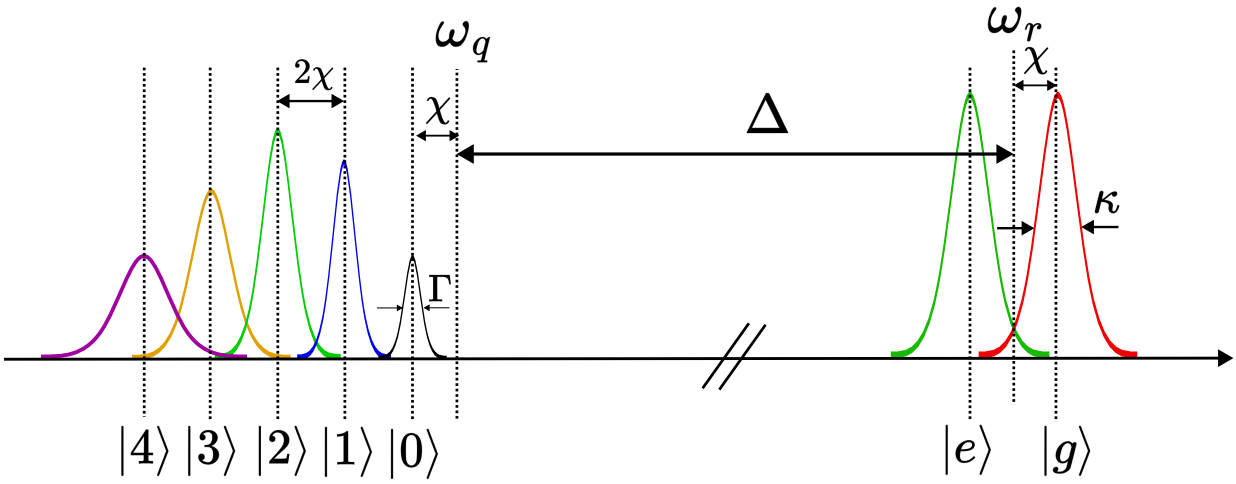

Figure 2.2: In the dispersive regime, both the qubit and the cavity experience frequency shifts based on the state of the other. The qubit line (on the left) with base linewidth Γ gets split into a set of lines corresponding to the photon number distribution of the cavity state. At low photon numbers and $\chi > \Gamma$, the qubit is in the *photon-number resolved* regime where individual transitions corresponding to specific $|n\rangle$ are resolvable. In the large $|n\rangle$ limit, all that is visible is an overall shift and broadening of the qubit line (due to the Stark shift and photon number broadening). The cavity (with linewidth κ) also is split into two new lines corresponding to $|e\rangle$ and $|g\rangle$ qubit states.

2.1.1 Dispersive Readout

In the previous section, we introduced the Jaynes-Cummings Hamiltonian and one might reasonably wonder why we should couple our qubit to a cavity, especially when considering the Purcell-induced decay added by the cavity. While not immediately obvious in the more general description above, the key lies in the fact that the cavity resonance frequency is now dependent on the *qubit state* or equivalently that the qubit resonance frequency has become sensitive to the *photon number* in the cavity as shown in Fig. 2.2. Our original coupling Hamiltonian read as follows:

$$H = \omega_c a^\dagger a + \frac{\omega_q}{2}\sigma_z + g(a^\dagger \sigma_- + a\sigma_+) \tag{2.15}$$

where we have omitted the 1/2 energy offset for the cavity term and dropped the energy non-conserving terms $a\sigma_-, a^\dagger\sigma_+$. In the *dispersive limit*, $\Delta = \omega_c - \omega_q \gg g$, we can rewrite

the Hamiltonian as follows [5, 6]:

$$H = \omega_c a^\dagger a + \frac{\omega_q}{2}\sigma_z + \chi a^\dagger a \sigma_z + \frac{\chi}{2}\sigma_z \tag{2.16}$$

where $\chi = \frac{g^2}{\Delta}$ is the dispersive shift. Note, this expression for χ is only valid for two-level systems and we will discuss correction to this term in Sec. 2.5.5. The third term, containing photonic and qubit operators, can be thought of as a Stark shift on the qubit frequency, and the bare $\chi\sigma_z$ term is a Lamb shift due to vacuum fluctuations. Depending on how we group terms, we can highlight a qubit-dependent frequency shift on the cavity frequency or a photon-number shift on the qubit frequency:

$$H = (\omega_c + \chi\sigma_z)a^\dagger a + \frac{(\omega_q + \chi)}{2}\sigma_z \tag{2.17}$$

$$H = \omega_c a^\dagger a + (\omega_q + 2\chi a^\dagger a + \chi)\frac{\sigma_z}{2} \tag{2.18}$$

If the cavity linewidth κ is smaller than the dispersive shift χ, we can infer the state of the qubit by measuring the resonance frequency of the cavity [2, 17](shown schematically in Fig. 2.2). In addition to qubit state readout, this dispersive interaction makes it possible to perform measurements of the cavity photon number distribution (provided the qubit linewidth is narrower than 2χ) [18] which we explored in 6.4 and to control the cavity state with photon-number conditioned qubit control pulses [19]. Although our devices only use the photonic mode as a tool for qubit state readout, using the photonic mode itself as a quantum resource with the qubit as an auxiliary control system has proven to be a very fruitful research direction [20].

2.2 Numerical simulations and talking to theorists

In the previous sections, we discussed the Jaynes-Cummings Hamiltonian as a simple description of an atom-light system with parameters such as ω_c, g, ω_q, etc. In the following

chapters, we will discuss how we measure those in a real system, but I want to take a moment to discuss the bane of physics: factors of 2 and 2π. Namely, when discussing a particular experimental system with theorists or even just when trying to simulate the system on a platform like QuTiP [21] or `scqubits` [22, 23], the units of the various parameters in the Hamiltonian become critical. We typically casually mention that the qubit frequency is 4.8 GHz, the coupling is 13 MHz, the Rabi rate is 100 MHz or the cavity linewidth is 30 kHz but how exactly do those numbers translate to the parameters ω_q, g, Ω, and κ?

It will be useful to define the standard form of the Hamiltonian we will be assuming our theory collaborators will be using:

$$H = (\omega_c - \omega_d)(a^\dagger a + 1/2) + (\omega_q - \omega_d)\frac{\sigma_z}{2} + g(a^\dagger \sigma_- + a\sigma_+) + \frac{\Omega}{2}\left(\cos(\phi)\sigma_x + \sin(\phi)\sigma_y\right) \quad (2.19)$$

where $\omega_{c,d,q}$ are the resonant frequencies for the cavity, applied drive, and qubit respectively, g is the qubit-cavity coupling strength, and Ω is the Rabi rate induced on the qubit by the applied drive with ϕ indicating the phase of the drive. Note, while we have set $\hbar = 1$ for notation simplicity, *all these quantities come with a factor of* 2π *when doing numerics.* For example, Rabi oscillations in the qubit population with a period of 20 ns correspond to a Rabi rate of 50 MHz *however*, the parameter Ω in the Hamiltonian above is $2\pi\times$ 50 MHz.

Dissipation is added to this simple Hamiltonian model using the standard Lindblad-form master equation [24–26]:

$$\frac{\partial \rho}{\partial t} = -i[H_0, \rho] + \sum_j \gamma_j D[\theta_j]\rho \quad (2.20)$$

where $D[\theta_j]\rho = \theta_j \rho \theta_j^\dagger - \frac{1}{2}\left(\theta_j^\dagger \theta \rho + \rho a_j^\dagger a\right)$ is the Lindblad super-operator for operator θ_j and γ_j is the associated loss term. For our typical systems consisting of a qubit coupled to a cavity, we consider three loss mechanisms: qubit depolarization (i.e. decay from $|e\rangle$ to $|g\rangle$, also called longitudinal relaxation), qubit dephasing (i.e. decay of *coherences* ρ_{ge}, ρ_{eg}, also

called transverse relaxation), and cavity leakage (the rate at which *energy* leaks out of the cavity). The collapse operators a_j for each of these dissipation channels can be modeled as $\sqrt{\Gamma}\sigma_-$ and $\sqrt{\frac{\gamma}{2}}\sigma_z$ for longitudinal and transverse noise respectively. and $\sqrt{\kappa}a$ for cavity leakage [25]. The longitudinal and transverse relation rates (Γ, γ) are given by $1/T_1$ and $1/T_2$ for the qubit and are implicitly in units of $2\pi \times$ Hz. The cavity relaxation rate κ can be extracted from the full width half maximum (FWHM) of the cavity intensity or from a ring-down measurement of the leakage voltage. In the first case, $\kappa = 2\pi \times$ FWHM is the correct expression; a useful conversion to remember is: a 1 MHz wide cavity has a 120 ns lifetime, not 1 μs. In the second case, $\kappa = 2/T_{cav}$ where T_{cav} is the exponential decay envelope of the ring-down voltage. The factor of 2 arises from κ being defined as the *energy leakage rate* which scales as the *power* and not the voltage.

2.3 Coplanar Waveguide Resonators

Our first building block is the coplanar waveguide resonator (CPW resonator), which implements the photonic mode in our circuits [27]. A coplanar waveguide is a 2D structure that consists of a metal center trace of width w separated by a gap g from the ground plane on each side on top of a substrate as illustrated in Fig. 2.3, reminiscent of a coaxial cable. To minimize losses, we use a superconductor for the ground plane (tantalum) and a high-dielectric constant material for the substrate (sapphire). The CPW is modeled as a lossless transmission line (although we will revisit this assumption in the following section) with inductance and capacitance per unit length (L_l , C_l) set by the geometry of the waveguide (in particular, the gap to center pin ratio) [28]:

$$L_l = \frac{\mu_0}{4}\frac{K(k_0')}{K(k_0)} \tag{2.21}$$

$$C_l = 4\epsilon_0\epsilon_{eff}\frac{K(k_0)}{K(k_0')} \tag{2.22}$$

where ϵ_{eff} is the effective dielectric constant, $k_0 = \frac{w}{w+2g}$, $k_0' = \sqrt{1-k_0^2}$, and $K(x)$ is the complete elliptic integral of the first kind.

A CPW resonator is achieved by terminating the CPW with a short or an open to ground on both ends. Just as with wind instruments, the termination type and length of the waveguide set the resonant frequency for the fundamental mode of the trace:

$$f_{res} = \frac{v}{nL} \tag{2.23}$$

where $v = \frac{c}{\sqrt{\epsilon_{eff}\mu_{eff}}}$ is the wave speed in the system, L is the length, and n=2 for open-open or closed-closed terminations and n=4 otherwise. For $\lambda/2$ resonators (n=2), the CPW is typically terminated with a pair of open circuits leading to voltage anti-nodes at the ends of the resonator, shown schematically in Fig. 2.4 with associated harmonics. $\lambda/4$ resonators (n=4) consist of an open and a short end with a voltage and current anti-node at those ends respectively. Coupling to the rest of the system (for example a feedline for read-out) is achieved by bringing the resonator into proximity with another CPW trace. This coupling can be either inductive or capacitive depending on the location on the resonator close to the coupling site. Voltage anti-nodes (such as at the open end of a $\lambda/2$ resonator) are predominantly capacitive coupling sites, whereas current anti-nodes (such as at the shorted end of a $\lambda/4$ or the middle of a $\lambda/2$ resonator) are predominantly inductive. However, this approximation is only valid for coupling sites with spatial dimensions much smaller than the wavelength of the resonant mode. Outside of this regime, both coupling types will contribute and require a more extensive analysis incorporating the spatial mode profile of the current and voltage.

For typical devices on sapphire substrates, $\mu_{eff} \approx 1$ and $\epsilon_{eff} \approx 5.6$ (where the effective dielectric constant can be taken as the average of the sapphire dielectric constant $\epsilon = 11.5$ and air $\epsilon = 1$). Note, the dielectric constant of sapphire is anisotropic with ϵ ranging from 9.4 perpendicular to the C-axis (going *into* the substrate) to 11.5 parallel to the C-axis

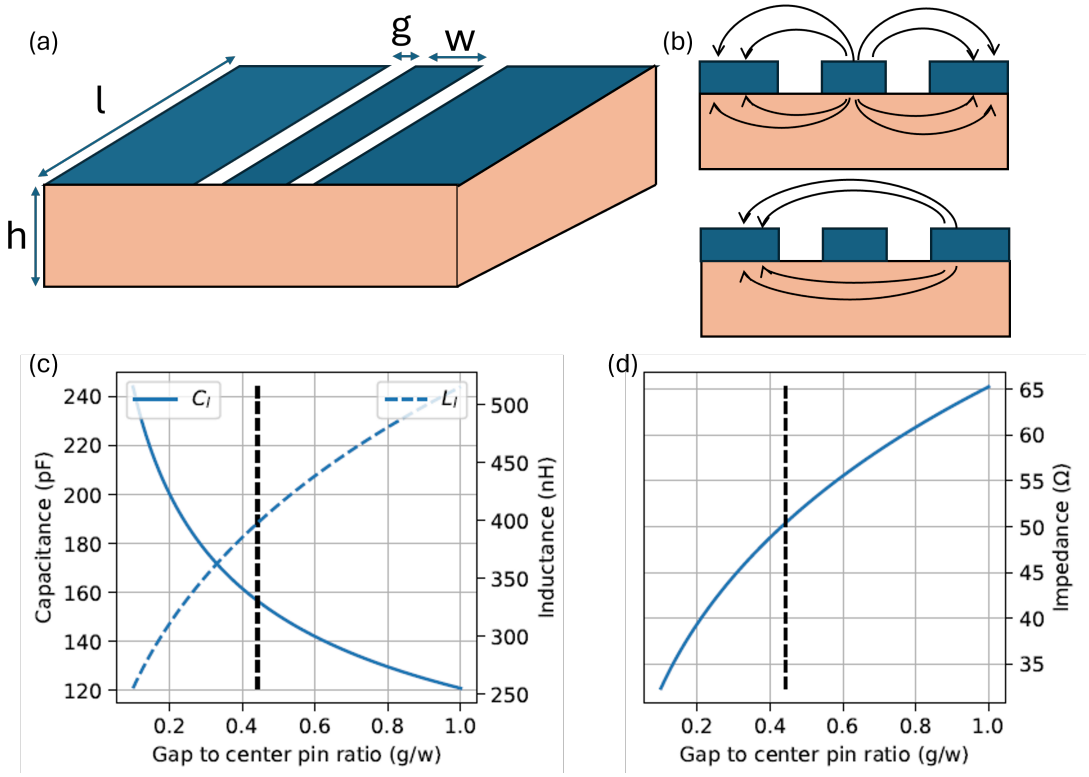

Figure 2.3: The Coplanar Waveguide. (a) Geometry for a standard waveguide: the substrate thickness is set by the substrate choice (530μm sapphire in our case), the center pin width and gap to the ground plane set the characteristic inductance, capacitance, and impedance of the transmission line (see panels (c)-(d) for plots of C,L, and Z vs the gap to center pin ratio). (b) Electric field profiles for the CPW mode (top) and parasitic slotline mode (bottom). In the CPW mode, the field lines go from the center pin to the two ground planes symmetrically. Narrower CPW lines will cause a tighter bending of the field lines and concentrate more field at the interfaces leading to increased losses. While not pictured here, adding a lossy material below the substrate (like a normal metal) will also add loss (discussed more in Chapter 4). The slotline mode is a parasitic mode arising from an imbalance between the ground planes and completely bypasses the center pin. It is suppressed by adding grounding wirebonds between the ground planes regularly along the length of the transmission line (essentially conducting bridges over the center pin) as shown in Fig. 2.4. (c)-(d) Electrical properties of the CPW as a function of gap to center pin ratio. The apparent intersection in the L-C plot is an artifact of the dual y-axis plotting method. Our devices operate with $w = 20\mu$m and $g = 8.372\mu$m.

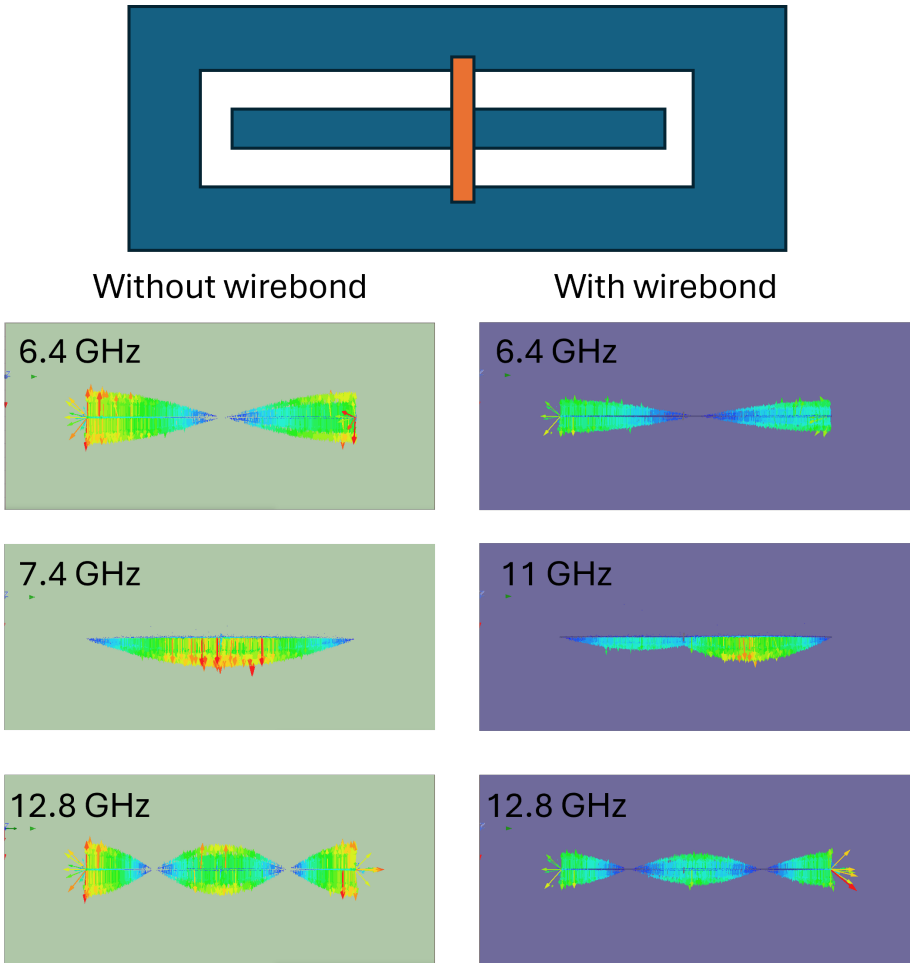

Figure 2.4: Modes in a half-wave CPW resonator without and with a wirebond (orange) in the center. The wirebond acts as a bridge between the two ground planes and does not connect to the center pin. The wirebond serves to short together the ground planes to suppress parasitic modes in the CPW. By placing the wirebond near the center of the resonator, we shift the resonant frequency of the slot line mode from 7.4 GHz to 11 GHz. The wirebond *only* suppresses slot-line modes as evidenced by the unchanged 3rd harmonic of the CPW mode in the lowest row.

(parallel to the surface of the substrate). In the 4-10 GHz range, the main operating range for cQED-based devices, this leads to $\lambda/2$ resonators with lengths of up to 15 mm which becomes impractical when attempting to design devices with many resonators. Fortunately, CPW resonators can be *meandered* on the chip without affecting their resonant frequency (although higher order effects do slightly shift the frequency at the 0.1% level). This allows for compact layouts as well as some interesting applications in non-euclidean geometries [29].

In the following sections, I will quickly show how to transform this extended physical resonator into first an equivalent lumped element circuit composed of inductors, capacitors, and resonators and finally how to model this system using the usual quantum mechanical description (i.e. the circuit Hamiltonian).

2.3.1 The LCR resonator

To represent the CPW resonators described in the previous section, we use the parallel LCR resonator circuit. This circuit consists of a parallel combination of an inductor, a capacitor, and a resistor where the inductor and the capacitor form an oscillating circuit and the resistance introduces loss, damping out the oscillations. The response of this circuit in the frequency domain can be extracted by examining the impedance $Z(\omega)$:

$$Z(\omega) = \left(\frac{1}{R} + j\omega C + \frac{1}{j\omega L}\right)^{-1} \tag{2.24}$$

$$Z(\omega) = \frac{1}{1 + \frac{R}{j\omega L}\left(1 - \frac{\omega^2}{\omega_0^2}\right)} \tag{2.25}$$

where the j is the usual electrical circuit complex component ($j = -i$) and $\omega_0 = 1/\sqrt{LC}$ is the resonance frequency of the circuit. Near resonance, $\omega \approx \omega_0$, we can simplify the expression by using the following approximation: $\omega^2 - \omega_0^2 = (\omega + \omega_0)(\omega - \omega_0) \approx 2\omega\delta\omega$,

giving:

$$Z(\omega) = \frac{R}{1 + \frac{2jR}{L\omega_0}\frac{\delta\omega}{\omega_0}} = \frac{R}{1 + 2jQ\frac{\delta\omega}{\omega_0}} \tag{2.26}$$

The quantity $R/L\omega_0$ (equivalently, $\omega_0 RC$) is defined as the Q-factor of the resonator and is dimensionless. Just as with mechanical or optical resonances, the Q-factor is the ratio between the stored power (in the LC portion of the circuit) to the dissipated power in one cycle (through the resistive portion of the circuit).

We can extract a lumped element description of the CPW resonator by comparing the impedance and collecting terms (following the description in [30]). The input impedance of a transmission line of length l with characteristic impedance Z_0 and termination Z_L is:

$$Z(l) = Z_0 \frac{Z_L + jZ_0 \tanh(\gamma l)}{Z_0 + jZ_L \tanh(\gamma l)} \tag{2.27}$$

where $\gamma = \alpha + j\beta$ is the propagation constant with $\alpha \approx \frac{1}{2}\left(\frac{R_l}{Z_0} + G_l Z_0\right)$ and $\beta \approx \omega\sqrt{L_l C_l}$. For a $\lambda/2$ resonator, the load impedance Z_L is an *open* so we can simplify the expression for $Z(l)$:

$$Z(l) = Z_0 \frac{1}{j\tanh(\alpha l + j\beta l)} = Z_0 \frac{1 + j\tan(\beta l)\tanh(\alpha l)}{\tanh(\alpha l) + \tan(\beta l)} \tag{2.28}$$

which we can expand around $\omega = \Delta\omega + \omega_0$. In resonance, $l = \lambda/2$, which allows us to rewrite $\beta l = \pi + \frac{\pi\Delta\omega}{\omega_0}$. For low loss, $\tanh(\alpha l) \approx \alpha l$, and we can finally rewrite the expression without trig functions:

$$Z_{\lambda/2} = \frac{Z_0}{\alpha l + j\frac{\pi\Delta\omega}{\omega_0}} \tag{2.29}$$

where we have dropped the term $\propto \alpha l \frac{\pi\Delta\omega}{\omega_0}$ in the numerator as we've assumed $\alpha l \ll 1$. By comparing terms with the expression for the LRC resonator we can identify the equivalent

resistance, inductance, capacitance, and impedance for a $\lambda/2$ resonator:

$$R = \frac{Z_0}{\alpha l} \tag{2.30}$$

$$C = \frac{\pi}{2\omega_0 Z_0} = \frac{C_l l}{2} \tag{2.31}$$

$$L = \frac{1}{\omega_0^2 C} = \frac{2L_l l}{\pi^2} \tag{2.32}$$

$$Z = \frac{2Z_0}{\pi} \tag{2.33}$$

intuitively, we can understand the difference between the total capacitance/inductance/resistance/impedance (defined as $l * C_l$ for the capacitance for example) and the effective values as a weighted sum which takes into account the spatial profile of the resonant mode. Near the voltage nodes, the effect of the capacitance is smaller than the inductance, and vice versa at the antinodes.

2.3.2 Coupling a CPW resonator to the environment

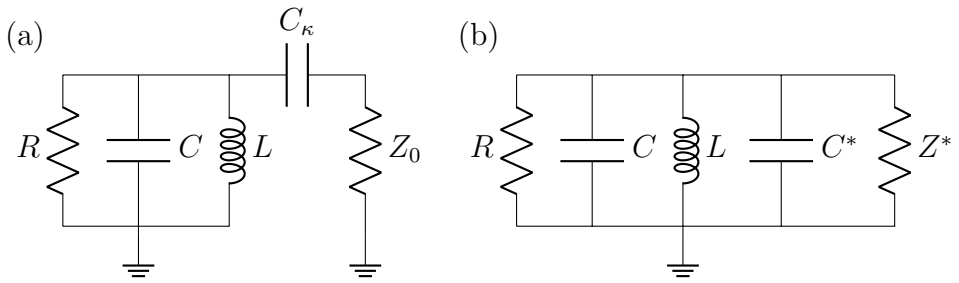

Figure 2.5: Circuit models for the coupled LCR resonator. (a) C_κ is the coupling capacitance to the environment (i.e. our measurement apparatus) and Z_0 is the characteristic impedance of the transmission lines (50Ω). (b) Equivalent full parallel circuit to the one in (a). R^*, C^* are found by setting the parallel impedance configuration equal to the series configuration (see text).

To finalize our discussion of resonators before we jump into the quantized description, we will consider the case of a capacitively coupled resonator. As we will discuss later in Chapter 4, the resonator has to be coupled to the environment for measurement. Here we will analyze the effect of adding a coupling capacitance on both the resonance frequency

and the quality factor of the resonance. This analysis closely follows that presented in [27], but is reproduced here with some helpful comments.

Our starting point is the simple LCR circuit described in the previous section capacitively coupled to a 50Ω load as shown in Fig. 2.5. The bare LCR circuit has a resonant frequency $\omega_0 = 1/\sqrt{LC}$ and Q factor $\omega_0 RC$. By representing the coupling section (C_c, Z_0) with the equivalent parallel circuit (R^*, C^*) as shown in (b), we can compute the new resonant frequency as well as Q-factor in terms of the new quantities R^*, C^* (derived below):

$$\omega_0' = \frac{1}{\sqrt{L(C + C^*)}} \tag{2.34}$$

$$Q_T = \omega_0 (R||R^\star)(C + C^\star) \tag{2.35}$$

where we have simply written ω_0' and Q_T in terms of the total effective resistance and capacitance and the $||$ notation indicates the parallel configuration. For the circuits that we consider, $C \gg C^*$ allowing us to simplify the ω_0' expression:

$$\omega_0' = \frac{1}{\sqrt{LC}} \times \frac{1}{\sqrt{1 + \frac{C^\star}{C}}} \approx \omega_0 \left(1 - \frac{C^\star}{2C}\right) \tag{2.36}$$

where we can see that the addition of coupling has the effect of *lowering* the resonant frequency of the circuit. The Q-factor is typically broken up into two components: Q_{int} and Q_{ext} which represent the internal losses and losses to the environment (i.e. our measured signal!) respectively. We can represent the total Q-factor as a the sum of the two loss rates:

$$\frac{1}{Q_T} = \frac{1}{Q_{int}} + \frac{1}{Q_{ext}} \tag{2.37}$$

$$= \frac{1}{\omega_0 \left(C + C^*\right)} \left(\frac{1}{R} + \frac{1}{R^*}\right) \tag{2.38}$$

$$= \frac{1}{\omega_0 C R} + \frac{1}{\omega_0 C R^*} \tag{2.39}$$

where we have used $C \gg C^*$ to simplify the expressions. We can immediately identify

$Q_{int} = \omega_0 CR$ and $Q_{ext} = \omega_0 CR^*$.

Now that we have computed all our corrections in terms of C^* and R^*, we need to compute the equivalent parallel circuit in terms of the original quantities C_κ and Z_0. This is done by setting the parallel impedance (Z_p) equal to the series impedance (Z_s) it is replacing:

$$Z_s(\omega) = Z_0 + \frac{1}{jC_\kappa \omega} \tag{2.40}$$

$$Z_p(\omega) = (j\omega C^* + 1/R^*)^{-1} = \frac{R^* - j\omega C^* R^{*2}}{1 + (\omega R^* C^*)^2} \tag{2.41}$$

which also serve as definitions for R^* and C^*. By comparing the real and imaginary components, we can calculate R^* and C^*:

$$Z_0 = \frac{R^*}{1 + (\omega R^* C^*)^2} \tag{2.42}$$

$$\frac{1}{\omega C_\kappa} = \frac{\omega C^* R^{*2}}{1 + (\omega R^* C^*)^2} \tag{2.43}$$

while this expression looks quite daunting to invert for R^* and C^*, the following tricks allow for a relatively simple process: taking the ratio of the expressions allows us to relate $Z_0 \omega C_\kappa$ to $(\omega C^* R^*)^{-1}$ and assuming $Z_0 \omega C_\kappa \ll 1$ simplifies the final expressions:

$$R^* = \frac{Z_0}{(\omega C_\kappa Z_0)^2} \tag{2.44}$$

$$C^* \approx C_\kappa \tag{2.45}$$

where we can see that in the parallel circuit version, the load impedance is transformed to a much larger equivalent parallel resistance, scaling as $1/C_\kappa^2$. Putting it all together, we can compute the coupling Q-factor, Q_{ext}, as:

$$Q_{ext} = \omega_0 CR^* = \frac{\pi}{2} \frac{1}{(\omega_0 C_\kappa Z_0)^2} \tag{2.46}$$

where we have used $\omega_0 C = \frac{\pi}{2Z_0}$ for the lumped element representation of a $\lambda/2$ CPW resonator.

With the circuit description of CPW resonators laid out, we can now transition to mapping our electrical circuits to Hamiltonians.

2.4 Circuit Quantization

Here I will quickly sketch out the idea behind circuit quantization, starting with building up the circuit Lagrangian using generalized charge and flux variables, using the standard Legendre transform to build the Hamiltonian from the circuit Lagrangian, and finally working through an example with coupled LC oscillators.

2.4.1 The Quantum Harmonic Oscillator

As a starting point, we will quickly revisit the quantum harmonic oscillator (QHO), possibly the most widely used starting point in physics to model linear systems. Fundamentally, the QHO describes the dynamics of any system with perfectly linear behavior (i.e. independent with respect to the scale of the dynamics) arising from a quadratic potential well. The Lagrangian for a simple spring-mass system (mass m, spring constant k) is given by:

$$\mathcal{L} = T - V = \frac{1}{2}mv^2 - \frac{1}{2}kx^2 \tag{2.47}$$

where we have simply computed the kinetic and potential energy of the system. While we can immediately compute the equations of motion for this system, $\mathcal{L}$ is not an intuitive quantity as it does not measure the total energy in the system. The system Hamiltonian, H, representing the total energy in the system, can be obtained from the Legendre transform

of the variables v and x to the conjugate variables p and q:

$$H = \sum_i p_i \dot{q}_i - \mathcal{L} \tag{2.48}$$

where p and q are defined as:

$$p = \frac{\partial \mathcal{L}}{\partial \dot{q}} = m\dot{q} \tag{2.49}$$

$$q = x \tag{2.50}$$

leaving us with the Hamiltonian H:

$$H = m\dot{q}^2 - \frac{1}{2}m\dot{q}^2 + \frac{1}{2}kq^2 = \frac{1}{2m}p^2 + \frac{1}{2}kq^2 \tag{2.51}$$

While this machinery might seem overly complicated for a simple spring-mass oscillator whose energy one could have simply written down, it becomes particularly well suited for more complex systems involving multiple coupled coordinates or for whom the momentum p definition is not as trivial as mv (such as for electromagnetic systems).

The coordinates p, q are the so-called *conjugate* coordinates and satisfy the following relation: $\{q, p\} = \delta$ where $\{\}$ denotes the Poisson bracket. While the Hamiltonian has many fascinating properties in classical physics (such as finding constants of motion), we will follow the standard steps to go from classical to quantum representations:

- Convert p, q to operators $\hat{p}, \hat{q}$
- Replace the Poisson bracket with the commutator: $\{p, q\} = -1 \Rightarrow [p, q] = -i\hbar$

Leading to the standard QHO Hamiltonian:

$$H = \frac{1}{2m}\hat{p}^2 + \frac{1}{2}k\hat{q}^2 = \frac{1}{2m}\hat{p}^2 + \frac{1}{2}m\omega^2\hat{x}^2 \tag{2.52}$$

where we've used $\omega = \sqrt{\frac{k}{m}}$ and rewritten the general coordinate q as x by convention. Finally, we define raising and lowering operators $a, a^\dagger$ to transform our system into the standard second quantized form:

$$\hat{a} = \sqrt{\frac{m\omega}{2\hbar}}\left(\hat{x} + \frac{i}{m\omega}\hat{p}\right) \qquad \hat{x} = \sqrt{\frac{\hbar}{2m\omega}}\left(a^\dagger + a\right) \tag{2.53}$$

$$\hat{a}^\dagger = \sqrt{\frac{m\omega}{2\hbar}}\left(\hat{x} - \frac{i}{m\omega}\hat{p}\right) \qquad \hat{p} = i\sqrt{\frac{\hbar m\omega}{2}}\left(a^\dagger - a\right) \tag{2.54}$$

for which $a, a^\dagger$ satisfy $[a^\dagger, a] = 1$. Using the substitution and the commutation relation (and setting $\hbar = 1$), we finally obtain:

$$H = -\frac{\omega}{4}(a^\dagger - a)^2 + \frac{\omega}{4}(a^\dagger + a)^2 \tag{2.55}$$

$$H = \omega\left(a^\dagger a + \frac{1}{2}\right) = \omega\left(\hat{n} + \frac{1}{2}\right) \tag{2.56}$$

where $a^\dagger a$ is the *number operator* $\hat{n}$ with eigenstates $|n\rangle$, the Fock states. In this second quantized form, the energy of any level is simply $n\omega$ (ignoring the constant energy offset), clearly indicating that the *energy per excitation* is constant.

While this detour might seem like an extended review of undergraduate quantum mechanics, the machinery developed in building up the Hamiltonian from the Lagrangian and in defining the raising and lowering operators will prove invaluable when analyzing systems with more degrees of freedom.

2.4.2 The circuit Lagrangian

As we discussed in the previous section, the starting point to the system Hamiltonian is the circuit Lagrangian (see [11, 12] for a detailed treatment). For electrical circuits, one might consider voltage V and current I as natural conjugate variables; however, using a different set of conjugate variables makes the analysis of the Josephson junction simpler.

Instead, circuits are defined in terms of generalized charge Q and flux ϕ:

$$\phi(t) = \int_{-\infty}^{t} V(t')dt' \qquad Q(t) = \int_{-\infty}^{t} I(t')dt' \tag{2.57}$$

we can now compute the energy stored in a capacitive or inductive element composing our circuit by integrating the power:

$$E_{cap} = \int IV_{cap}dt' = \int \dot{Q}\frac{Q}{C}dt' = \frac{Q^2}{2C} = \frac{CV^2}{2} = \frac{C\dot{\phi}^2}{2} \tag{2.58}$$

$$E_{ind} = \int I_{ind}Vdt' = \int \frac{\phi}{L}\dot{\phi}dt' = \frac{\phi^2}{2L} \tag{2.59}$$

where we have use the usual expressions for the voltage and current across a capacitor and inductor respectively.

A simple LC circuit consisting of a capacitor and an inductor in parallel has the following Lagrangian:

$$\mathcal{L} = \frac{C\dot{\phi}^2}{2} - \frac{\phi^2}{2L} \tag{2.60}$$

where we have used the $\dot{\phi}$ representation of the capacitive energy to put the Lagrangian in standard form: $\mathcal{L}(q, \dot{q})$. Using the formalism defined above, we can immediately compute the conjugate variable Q:

$$Q = \frac{\partial \mathcal{L}}{\partial \dot{\phi}} = C\dot{\phi} \tag{2.61}$$

and write the Hamiltonian:

$$H = \frac{Q^2}{2C} + \frac{\phi^2}{2L} \tag{2.62}$$

By comparing this Hamiltonian to the mass-spring system, we can make the following

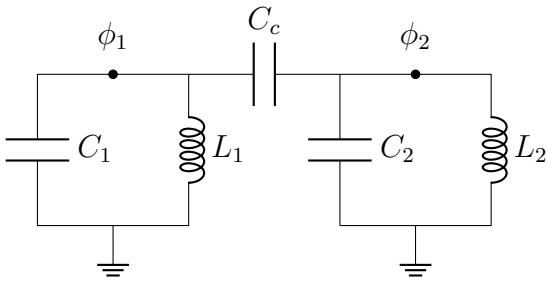

Figure 2.6: Coupled LC resonators. The generalized fluxes (ϕ_1, ϕ_2) are used to compute the circuit Lagrangian and find the equivalent Hamiltonian representation.

identification: $C \to m$ and $L \to k^{-1}$:

$$\hat{a} = \sqrt{\frac{\omega C}{2\hbar}}\left(\hat{\phi} + \frac{i}{\omega C}\hat{Q}\right) \qquad \hat{\phi} = \sqrt{\frac{\hbar}{2\omega C}}\left(a^\dagger + a\right) \tag{2.63}$$

$$\hat{a}^\dagger = \sqrt{\frac{\omega C}{2\hbar}}\left(\hat{\phi} - \frac{i}{\omega C}\hat{Q}\right) \qquad \hat{Q} = i\sqrt{\frac{\hbar\omega C}{2}}\left(a^\dagger - a\right) \tag{2.64}$$

Replacing the various variables in the ladder operator method, we obtain exactly the same final Hamiltonian: $H = \omega\left(a^\dagger a + \frac{1}{2}\right)$ where ω is now $\frac{1}{\sqrt{LC}}$ which is the familiar resonance frequency of an LC oscillator.

2.4.3 Coupled LC oscillators

To show the power of the generalized Lagrangian method, we will compute the Hamiltonian for a system of coupled LC oscillators shown in Fig. 2.6. Each *node* in the circuit is assigned a flux ϕ_i where particular care must be taken when an inductive loop is formed (as an external applied flux ϕ_{ext} will add a term to the inductive energy). In our example, we consider the case of a pair of capacitively coupled LC resonators with coupling capacitance C_c. We start by simply listing the kinetic and potential energy contributions from each circuit element:

$$\mathcal{L} = \frac{C_1\dot{\phi_1}^2}{2} + \frac{C_2\dot{\phi_2}^2}{2} + \frac{C_c(\dot{\phi_1} - \dot{\phi_2})^2}{2} - \frac{\phi_1^2}{2L_1} - \frac{\phi_2^2}{2L_2} \tag{2.65}$$

as usual, we compute the conjugate variables Q_1 and Q_2:

$$Q_1 = C_1\dot{\phi}_1 + C_c(\dot{\phi}_1 - \dot{\phi}_2) = (C_1 + C_c)\dot{\phi}_1 - C_c\dot{\phi}_2 \tag{2.66}$$

$$Q_2 = C_2\dot{\phi}_2 + C_c(\dot{\phi}_2 - \dot{\phi}_1) = (C_2 + C_c)\dot{\phi}_2 - C_c\dot{\phi}_1 \tag{2.67}$$

which can be inverted to find $\dot{\phi}_1$ and $\dot{\phi}_2$:

$$\dot{\phi}_1 = \frac{1}{C_\Sigma}\left[(C_2 + C_c)Q_1 + C_cQ_2\right] \tag{2.68}$$

$$\dot{\phi}_2 = \frac{1}{C_\Sigma}\left[(C_1 + C_c)Q_2 + C_cQ_1\right] \tag{2.69}$$

where we have defined $C_\Sigma = C_1C_2 + C_1C_c + C_2C_c$ which is the determinant of the capacitance matrix defined in the canonical transformation for $Q \to \dot{\phi}$. Using the previous equations, we can write the Hamiltonian:

$$H = \frac{1}{C_\Sigma}\left(\frac{C_2 + C_c}{2}Q_1^2 + \frac{C_1 + C_c}{2}Q_2^2 + C_cQ_1Q_2\right) + \frac{\phi_1^2}{2L_1} + \frac{\phi_2^2}{2L_2} \tag{2.70}$$

which we can write in terms of normal modes ω_1, ω_2 coupled through C_c:

$$H = \left[\frac{Q_1^2}{2(C_1 + C_c)} + \frac{\phi_1^2}{2L_1}\right] + \left[\frac{Q_2^2}{2(C_2 + C_c)} + \frac{\phi_2^2}{2L_2}\right] + \frac{C_c}{C_\Sigma}Q_1Q_2 \tag{2.71}$$

$$H = \omega_1 a^\dagger a + \omega_2 b^\dagger b + g(a^\dagger + a)(b^\dagger + b) \tag{2.72}$$

where $\omega_{1,2} = \frac{1}{\sqrt{L_{1,2}(C_{1,2}+C_c)}}$ and $g = \frac{C_c}{2\sqrt{C_1C_2}}\sqrt{\omega_1\omega_2}$ (where we have assumed that $C_c \ll C_{1,2}$). This method allows us to directly convert a circuit to a set of coupled QHOs with coupling strength g.

2.5 Making a qubit

Our next building block is the qubit. The fundamental requirements for a qubit are as follows (paraphrasing the well known DiVincenzo criteria [31]): well defined states which can be isolated from the rest of the system (labeled $|0\rangle, |1\rangle$), state initialization (i.e. the ability to reliably prepare the qubit in a starting state), full control of the qubit state with universal gates, long coherence times, and the ability to measure the state. To satisfy these requirements with an electrical circuit, we will slightly tweak the LC oscillator to introduce some anharmonicity and isolate a 2-level subspace that we can use for the qubit. The transmon qubit consists of a large capacitor shunted by a Josephson junction, a lossless nonlinear inductance.

2.5.1 The Josephson junction as a nonlinear inductor

The Josephson effect, proposed in 1962 by overachieving graduate student Brian Josephson, is the appearance of a persistent current between superconducting leads separated by a tunnel barrier (usually implemented with an insulator) with no applied voltage [32–34]. The current flowing through the junction arises from the tunneling of Cooper pairs between the two superconducting leads and is modulated by the phase difference between the leads of the junction. The voltage-current relationships for Josephson junctions (also known as the Josephson relations) are:

$$I(t) = I_c \sin(\phi(t)) \tag{2.73}$$

$$\frac{\partial \phi}{\partial t} = \frac{2eV(t)}{\hbar} \tag{2.74}$$

where I_c is the critical current (set in our case by the superconducting metal used, the barrier thickness, and the geometry of the junction), ϕ is the phase difference across the junction, $2e$ is the charge carried by a Cooper pair, and V is the voltage applied across the junction. These simple constitutive equations allow the Josephson junction (JJ) to act

as a perfect voltage to frequency converter (a fixed applied voltage will cause the phase to linearly ramp inducing a sinusoidal current between the leads, the AC Josephson effect) and inversely as a perfect frequency to voltage converter where a time varying phase leads to a fixed voltage between the leads of the junction (the inverse AC Josephson effect). We can also model the junction as a non-linear inductance by using the Josephson relations and the definition of inductance: $V = L\frac{dI}{dt}$:

$$\frac{dI}{dt} = I_c \cos(\phi(t))\frac{\partial \phi}{\partial t} = I_c \frac{2e}{\hbar}\cos(\phi(t))V(t) = L^{-1}V(t) \tag{2.75}$$

where $L(\phi) = \frac{\hbar}{2eI_c\cos(\phi)} = \frac{L_J}{\cos(\phi)}$ is the effective inductance of the junction. We can use the formalism defined earlier to find the equivalent inductive energy for a junction, defined as E_J:

$$E = \int IV dt = \int I_c \sin(\phi(t))\frac{\hbar}{2e}\frac{\partial \phi}{\partial t}dt = -I_c\frac{2e}{\hbar}(1-\cos(\phi)) = E_J(1-\cos(\phi)) \tag{2.76}$$

where E_J is referred to as the *Josephson energy* of the junction. To compare the JJ to a regular inductor, we can expand the cosine term around $\phi = 0$ to compare to a linear inductance: $E = E_J(\frac{\phi^2}{2} - \frac{\phi^2}{24})$.

2.5.2 The Transmon

The transmon [7] is fundamentally an LC oscillator with a non-linear inductance that provides anharmonicity to isolate a two-level subspace in the energy levels. It consists of a large capacitor (E_C) shunted by a Josephson junction (E_J), as shown in Fig. 2.7, and has the following Hamiltonian:

$$H = 4E_C(\hat{n} - n_0)^2 - E_J\cos\left(\hat{\phi}\right) \tag{2.77}$$

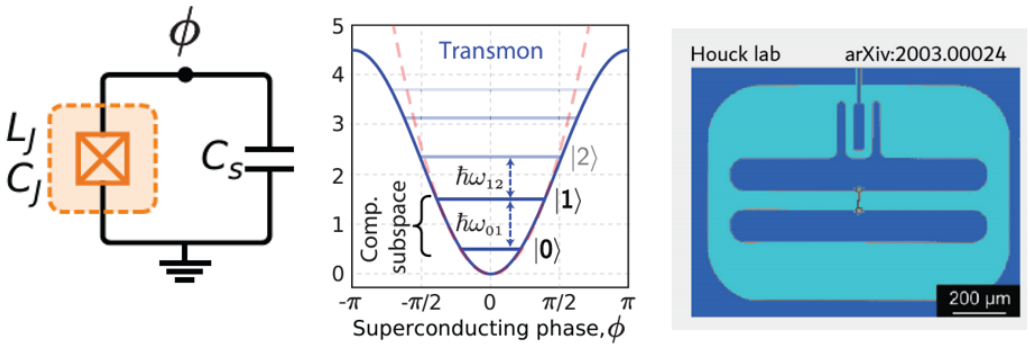

Figure 2.7: The transmon qubit. Left: Circuit representation of the transmon as a Josephson junction shunted by a capacitor acting as a weakly nonlinear LC resonator. Middle: Energy ladder for the transmon with the harmonic oscillator for reference (dashed red line). Due to the cosine potential, higher levels see a "softening" of the potential reducing the energy splitting of consecutive levels which gives a negative anharmonicity. Our computation subspace consists of the $|0\rangle$ and $|1\rangle$ states with transition energy ω_{01}. Right: Physical implementation of a transmon qubit. The capacitor is implemented with the two large metal pads in the center of the picture (the qubit is surrounded by a superconducting ground plane) with the Josephson junction connecting the two pads (faint line in the middle of the image). The first two panels are adapted from [5] and the right panel is adapted from [35].

where n_0 is the offset charge, $E_C = \frac{e^2}{2C_\Sigma}$ is the *charging energy* and $E_J = I_c\frac{2e}{\hbar}$ is the *inductive energy*. This Hamiltonian can be obtained from the previous derivation for the basic LC oscillator and replacing $\hat{Q}$ with $2e\hat{n}$ to cast the Hamiltonian in terms of purely non-dimensional quantities $\hat{n}$ and $\hat{\phi}$. The formal solutions to the energy levels of a transmon can be obtained from the Mathieu equations [7]:

$$E_m(n_0) = E_C a_{2[n_g+k(m,n_g)]}(-E_J/E_C) \tag{2.78}$$

where $a_\nu(q)$ is Mathieu's characteristic value and $k(n,m)$ denotes an ordering of eigenvalues. It will be useful to define two quantities: the charge dispersion ϵ_m and the anharmonicity α. ϵ_m measures the sensitivity of the energy levels to fluctuations in the offset charge n_0 on the island. α is defined as the energy difference between the 0-1 transition and the 1-2 transition. In Ref. [7], Koch et al. develop some guiding asymptotic behavior for ϵ_m and

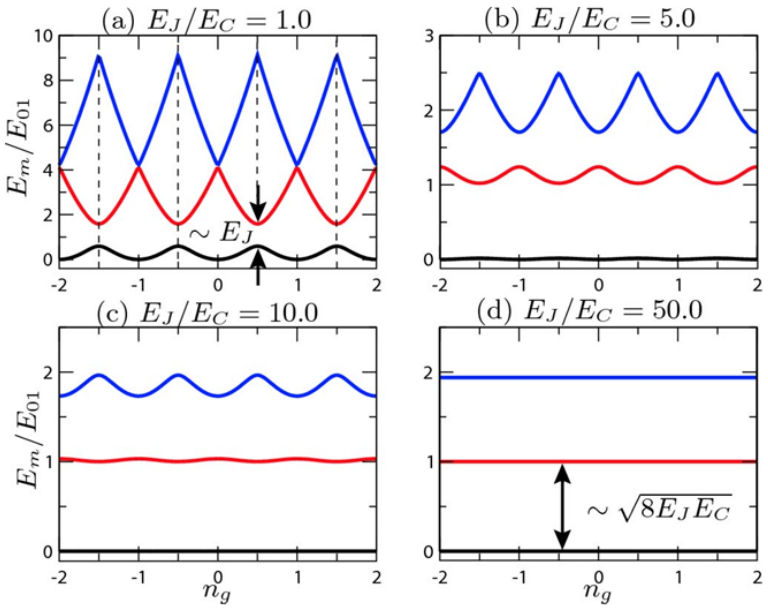

Figure 2.8: Transmon energy levels as a function of E_J/E_c and gate offset charge (n_g). At low ratios (panel (a)), the levels correspond to multiple parabolas centered at $n_g = n$ split by E_J when they intersect. The charge dispersion (ripple in the energy levels as a function of n_g) is large relative to the energy splitting with a single *charge sweetspot* at $n_g = 1/2+n$ where the energy is first order insensitive. As the ratio increases, (b)-(c), the bands flatten out and the anharmonicity α decreases slightly. In the transmon regime, $E_j/E_c > 50$, the bands are effectively flat, almost completely removing the charge dispersion on the energy levels while also significantly decreasing α (from a GHz scale in the cooper pair box to a couple hundred MHz in the transmon regime). Figure reproduced from [7].

α as a function of E_C and E_J:

$$\epsilon_m = (-1)^m E_c \frac{2^{4m+5}}{m!}\sqrt{\frac{2}{\pi}}\left(\frac{E_J}{2E_c}\right)^{\frac{m}{2}+\frac{3}{4}} e^{-\sqrt{8E_JE_C}} \propto e^{-\sqrt{8E_J/E_C}} \tag{2.79}$$

$$\alpha_r = \frac{\alpha}{\omega_{01}} \propto -\left(\frac{8E_J}{E_C}\right)^{-1/2} \tag{2.80}$$

showing that ϵ_m is exponentially suppressed at high E_J/E_C while α only polynomially decreases.

To build more intuition, we will examine two regimes that have been historically considered in the development of superconducting qubits: the Cooper-pair box regime ($E_C > E_J$) and the transmon regime ($E_J \gg E_C$).

The Cooper-pair box was one of the very first superconducting qubits used by the field

and effectively consisted of a Josephson junction capacitively coupled to a gate voltage for control of the offset charge n_0. Because $E_C > E_J$, the energy is dominated by the number of Cooper pairs on the island and is parabolic for each value of n_0. The E_J term induces coupling between the different charge number states and leads to the avoided crossings between $|n\rangle$ and $|n+1\rangle$ at the $n_0 = 1/2 + m$ points called the *charge degeneracy point*, shown in Fig. 2.8(a). The qubit was operated at this degeneracy point with the two levels consisting of the symmetric and anti-symmetric combinations of $|n\rangle$ and $|n+1\rangle$. However, frequency jumps due to changes in the offset charge (from fluctuations on the control lines) made this qubit difficult to operate and led to very short coherence times.

To mitigate the effects of offset charge noise, the transmon qubit reverses the energy hierarchy and uses a large E_J/E_C ratio (typically in the 50-100 range). In this new hierarchy, the Hamiltonian is dominated by the cosine term leading to localized wavefunctions in the phase basis and delocalization in the number basis leading to an exponential suppression of the charge dispersion ϵ_m as seen in panel (d) of Fig. 2.8 and Eq. 2.79. The following derivation follows the typical arguments presented in circuit QED reviews [5–7]. Expanding the cosine term allows us to rewrite the Hamiltonian as a Duffing oscillator:

$$H = 4E_C\hat{n}^2 - E_J\cos(\phi) = 4E_c\hat{n}^2 + E_J\frac{\hat{\phi}^2}{2} - E_J\frac{\hat{\phi}^4}{4!} \tag{2.81}$$

where the first two terms are the usual QHO terms for which we can define creation and annihilation operators a and $a^\dagger$:

$$\hat{n} = \frac{a - a^\dagger}{\sqrt{2}}\left(\frac{E_J}{8E_C}\right)^{1/4} \tag{2.82}$$

$$\hat{\phi} = \frac{a + a^\dagger}{\sqrt{2}}\left(\frac{8E_C}{E_J}\right)^{1/4} \tag{2.83}$$

allowing us to rewrite the Hamiltonian as:

$$H = \omega(a^\dagger a + \frac{1}{2}) - \frac{E_C}{12}(a^\dagger + a)^4 \tag{2.84}$$

where $\omega = \sqrt{8E_J E_C}$. By treating the second term as a perturbation on the usual QHO Hamiltonian (and computing the on-site energy corrections $\langle n| E |n\rangle$), the corrections to the QHO energy levels are given by:

$$E_m = m\omega - \frac{E_C}{12}(6m^2 + 6m + 3) \tag{2.85}$$

where we have dropped the 1/2 constant term from the QHO energy. By computing the energy difference between consecutive transitions, α is:

$$f_{m,m-1} = E_m - E_{m-1} = \omega_{01} - E_C m \tag{2.86}$$

$$\alpha = f_{m+1,m} - f_{m,m-1} = -E_C \tag{2.87}$$

The decoupling of α from E_J allows us to independently adjust the qubit frequency, ω, and anharmonicity, α, to match the design criterion. Typical values of α range from 150 MHz to 250 MHz and ω ranges from 2-6 GHz although can go higher or lower depending on the device. Higher precision for α can be obtained by keeping the ϕ^6 term in the cosine before performing the perturbation theory calculation.

2.5.3 The Tunable Transmon

The tunable transmon allows us to vary the resonance frequency *in situ* which is very useful for a variety of applications: studying frequency dependent effects (such as loss and noise spectrum [36]), for gates ([37–39]), to vary the interaction strength between qubits or between qubits and resonators ([40, 41] and which we explore in the later chapters of this thesis), and much more. To achieve tunability, we replace the single Josephson junction

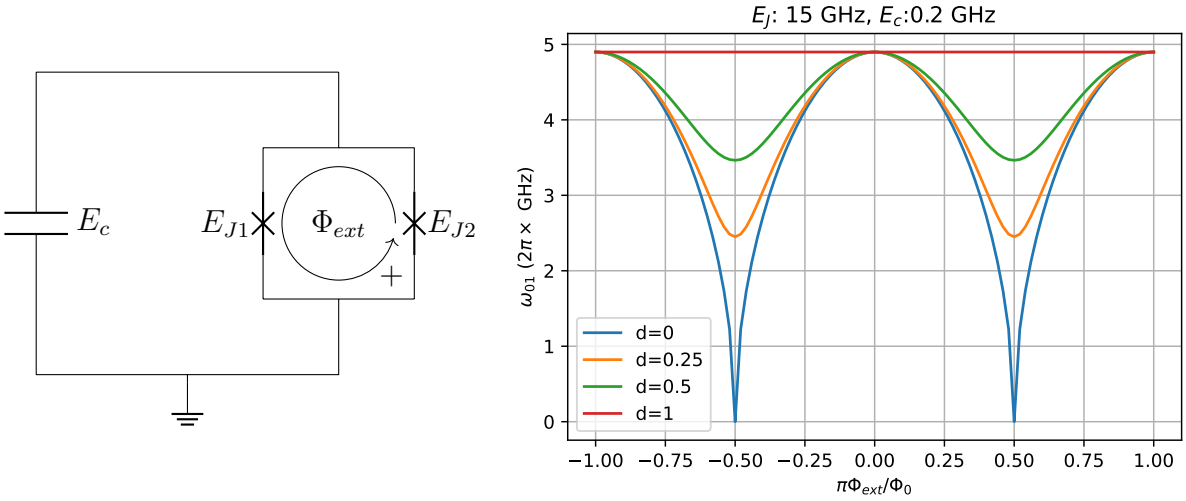

Figure 2.9: The tunable transmon replaces the single Josephson junction with a SQUID loop composed of two JJs with asymmetry d ($d = \frac{E_{J_2} - E_{J_1}}{E_{J_1} + E_{J_2}}$). The transition frequency varies periodically as a function of the applied flux Φ_{ext} with the modulation contrast set by the junction asymmetry. For symmetric junctions, the transition frequency approaches 0 (although the simple model used to plot the curves is no longer valid as we exit the transmon regime). Devices presented in this thesis operate with $d \approx 0.4 - 0.6$ which allows for some tunability and leads to the appearance of a lower sweetspot at $\Phi_{ext} = 0.5\Phi_0/\pi$.

with a pair of junctions connected in a loop configuration called a superconducting quantum interference device (SQUID) [34]. This causes interference between the two possible paths and leads to a phase-dependent effective junction energy E_J. By applying an external magnetic field through the loop, we can vary the interference phase to tune the junction energy and therefore the qubit resonance frequency (recall $\omega_{ge} = \sqrt{8E_J E_C}$).

To calculate the effective junction energy E_{J_T} as a function of the externally applied flux Φ_{ext} the two junction energies are added up and fluxoid quantization (i.e. the flux going through a loop has to be an integer) is used to express the phases across the junctions in terms of the external flux:

$$E_{J_T} = E_{J_1} \cos(\phi_1) + E_{J_2} \cos(\phi_2) \tag{2.88}$$

$$\phi_1 - \phi_2 = 2\pi n + 2\pi \frac{\Phi_{ext}}{\Phi_0} \tag{2.89}$$

where the phases for the second line are signed according to the direction the junctions

are traversed (in this case we chose a counter-clockwise path around the loop), as shown in Fig. 2.9, and $\Phi_0 = \frac{h}{2e}$ is the magnetic flux quantum. It will be useful to express $\phi_{1,2}$ in terms of the average and difference:

$$E_{J_T} = E_{J_1}\cos\left(\phi + \frac{\pi\Phi_{ext}}{\Phi_0}\right) + E_{J_2}\cos\left(\phi - \frac{\pi\Phi_{ext}}{\Phi_0}\right) \tag{2.90}$$

which can be rewritten using some trig identities, namely:

$$\cos(a+b) = \cos(a)\cos(b) - \sin(a)\sin(b) \tag{2.91}$$

$$A\cos(\phi) + B\sin(\phi) = \sqrt{A^2+B^2}\cos(\phi+\delta) \quad \text{with} \quad \tan(\delta) = -\frac{B}{A} \tag{2.92}$$

using the previous identities, Eq. 2.90 simplifies to:

$$E_{J_T} = E_{J_\Sigma}\left|\cos\left(\frac{\pi\Phi_{ext}}{\Phi_0}\right)\right|\sqrt{1+d^2\tan^2\left(\frac{\pi\Phi_{ext}}{\Phi_0}\right)}\cos(\phi-\phi_0) \tag{2.93}$$

where $E_{J_\Sigma} = E_{J_1} + E_{J_2}$, $d = \frac{E_{J_2}-E_{J_1}}{E_{J_1}+E_{J_2}}$, and $\tan(\delta) = d\tan\left(\frac{\pi\Phi_{ext}}{\Phi_0}\right)$. For fixed external flux Φ_{ext}, ϕ_0 can be gauged away leaving us with an effective junction with a tunable junction energy $E_J(\Phi_{ext}) = E_{J_\Sigma}|\cos\left(\frac{\pi\Phi_{ext}}{\Phi_0}\right)|\sqrt{1+d^2\tan^2(\frac{\pi\Phi_{ext}}{\Phi_0})}$. The asymmetry factor d controls the depth of modulation strength on E_J with perfectly symmetric junctions allowing for a very wide E_J range from approximately 0 to $2E_J$. However, for many applications, d is kept finite to reduce the sensitivity to flux noise away from the flux sweetspots ($\Phi_{ext} = \frac{n\Phi_0}{2}$).

2.5.4 Calculating the qubit-resonator coupling

By replacing the atom with a transmon qubit and confining the cavity mode to a 1D waveguide (the CPW resonator), circuit QED systems drastically increase the light-matter interaction strength. The dipole strength of the transmon is enhanced by its macroscopic dimensions: instead of looking at the dipole of an electron going between two atomic orbitals, we are considering a dipole made of a Cooper pair oscillating between two metal

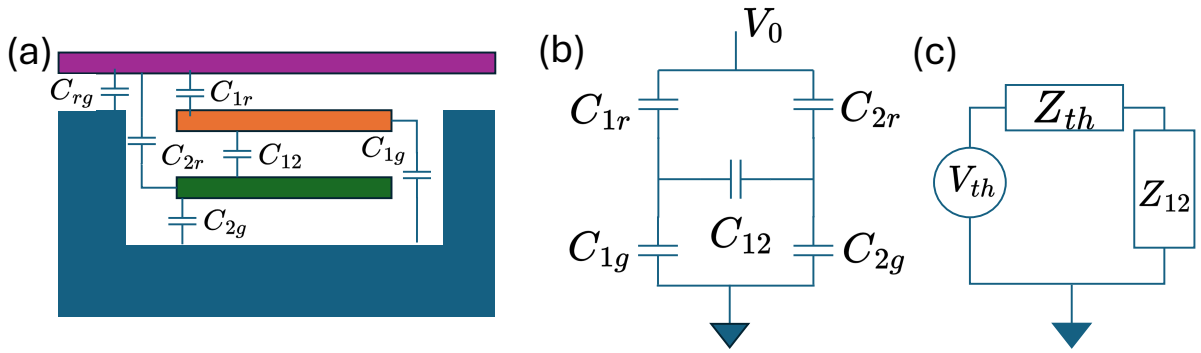

Figure 2.10: Capacitive network used to calculate the qubit-resonator coupling. (a) Schematic representation of the resonator (center pin in purple and ground plane in blue) and the qubit (capacitor defined by the orange and green blocks). The capacitance values between the various pieces of metal are extracted using the Maxwell EM solver (more on that later). (b) Circuit representation of the network in the schematic (ignoring the small center-pin to ground capacitance from the resonator for simplicity). To compute the coupling factor β, we calculate the ratio of the voltage across C_{12} (which is the voltage across the junction) to the resonator voltage V_0. This network is also used to find the total effective capacitance of the transmon C_Σ. (c) Thevenin equivalent circuit of the two-port circuit in (b).

pads separated by a few tens of microns. Recall, the qubit-cavity coupling in the JC Hamiltonian is given by $\vec{d} \cdot \vec{E}$ which, for the quasi 1D system consisting of a transmon in a CPW cavity, reduces to [5, 7]:

$$\hbar g = 2e\hat{n}\beta V_{rms}(a + a^\dagger) \tag{2.94}$$

where $2e\hat{n}$ serves as the dipole strength, V_{rms} is the voltage drop between the resonator center pin and the ground plane, and β is the coupling factor corresponding to the ratio of the resonator voltage to the voltage drop across the Josephson junction.

The typical transmon-resonator coupling scheme is shown in Fig. 2.10, with the relevant capacitances highlighted. Two main corrections need to be calculated for this coupled system: the total effective capacitance in parallel with the junction and the net voltage drop across the qubit capacitor and junction due to a field on the center pin of the resonator.

The effective capacitance C_Σ, is calculated by looking at the total capacitance between

the nodes labeled 1 and 2 in Fig. 2.10. The dominant contributions to the total capacitance are C_{12}, the direct capacitance between the qubit paddles, and C_{ig} the capacitance from each paddle to ground. The full capacitance includes the addition capacitance to the resonator center pin, however, this capacitance is much smaller and muddies the simple analytical expression. By simplifying the circuit network, we extract $C_\Sigma = C_{12} + \frac{C_{1g}C_{2g}}{C_{1g}+C_{2g}}$.

To compute the *coupling factor* β, we calculate the ratio of the voltage between nodes 1 and 2 to the total voltage drop between the center pin of the resonator and ground. Here we consider the voltage divider network shown in Fig. 2.10(b), which can be analyzed using the Thevenin equivalent impedance and voltage [42]. This specific configuration is also called a Wheatstone bridge and can be used to measure the impedance of an unknown load. Z_{th} is computed by calculating the impedance between nodes 1 and 2 and removing the voltage sources leaving:

$$Z_{\text{th}} = (Z_1 + Z_2)||(Z_3 + Z_4) = \frac{(Z_1 + Z_2)(Z_3 + Z_4)}{Z_1 + Z_2 + Z_3 + Z_4} \tag{2.95}$$

where the $||$ notation indicates parallel impedances. V_{th} is calculated by replacing the impedance between 1 and 2 with an open circuit and computing the voltage difference between the nodes. Here the voltages at 1 and 2 are set by the voltage dividers made up by Z_1/Z_3 and Z_2/Z_4 respectively. The final voltage difference between 1 and 2 is given by:

$$V_{th} = V_1 - V_2 = V_0 \left(\frac{Z_3}{Z_1 + Z_3} - \frac{Z_4}{Z_2 + Z_4} \right) \tag{2.96}$$

We can now replace the full impedance network with the Thevenin equivalent circuit and calculate V_{12} as the result of the voltage divider:

$$V_{12} = V_{th} \frac{Z_{12}}{Z_{12} + Z_{th}} = V_0 \left(\frac{Z_3}{Z_1 + Z_3} - \frac{Z_4}{Z_2 + Z_4} \right) \frac{Z_{12}}{Z_{12} + Z_{th}} = \beta V_0 \tag{2.97}$$

We can now consider a few limiting cases using typical values for the various capacitances

in the circuit:

- The capacitances to ground are much larger than the capacitances to the resonator center pin: C_{1r}, $C_{2r} \ll C_{1g}$, $C_{1g} \rightarrow Z_{1,2} \gg Z_{3,4}$
- The capacitances from the paddles to ground are roughly equivalent: $C_{1g} \approx C_{2g} \rightarrow Z_3 \approx Z_4$
- The direct capacitance between 1 and 2 is much larger than the capacitance to ground: $C_{12} \gg C_{1g}, C_{2g} \rightarrow Z_{12} \ll Z_{1g}, Z_{2g}$

Simplifying the earlier expression for β, we obtain a much simpler expression that we can use as a guiding design rule:

$$\beta \approx \frac{Z_{12}}{2}\left(\frac{1}{Z_1} - \frac{1}{Z_2}\right) = \frac{C_{1r} - C_{2r}}{2C_{12}} \tag{2.98}$$

Here we can see that for large coupling β, one has to maximize the relative difference of the capacitance to the center pin with respect to the direct capacitance between the paddles. In practice, achieving more than $\beta = 0.3$ is difficult as these capacitance cannot be independently tuned. However, $\beta = 0.3$ designs can achieve coupling values up to 150-200 MHz, which is sufficient for most designs.

2.5.5 Dispersive shift for the transmon qubit

To finalize our discussion of the coupled qubit-cavity system we have to revisit our expression for the dispersive shift derived at the beginning of this chapter. In the simple JC case, the atom is considered to be a two level system (with effectively infinite anharmonicity). However, the transmon is explicitly an n-level system whose lowest two levels we use to perform quantum operations. Each transition in the transmon will couple to the cavity leading to multiple dispersive shifts to each level which must be summed up to calculate the net shift. In particular, resonances between a specific qubit transition and the cavity lead

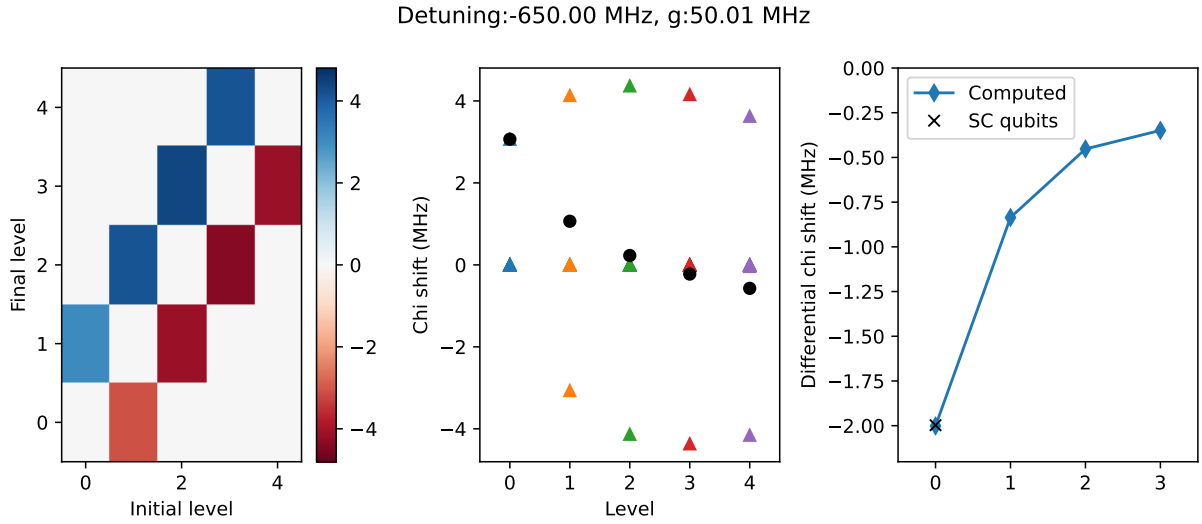

Figure 2.11: Left: Energy shifts on level i due to coupling to level j computed using Eq. 2.99 using ScQubits for the transmon numerics [22, 23]. The total chi shift for each level is found by summing corresponding column in the matrix (middle) where the colored dots show the energy shifts due to coupling to all the other levels (traversing the plot on the left vertically) and the black circles show the total shift (sum). The differential dispersive shift (i.e. the frequency shift in the cavity frequency for the qubit state in $|i\rangle$ or $|j\rangle$) is computed by taking the difference between consecutive levels (right).

to a breakdown in the perturbative approach used to calculate the individual chi shifts. The generic expression for the energy shift on each level can be obtained from second order perturbation theory [7](see [8] for an excellent overview of energy shifts for arbitrary atoms):

$$\chi_\alpha = g^2 \sum_{\beta \neq \alpha} |n_{\alpha\beta}|^2 \frac{2\nu_{\alpha\beta}}{\nu_{\alpha\beta}^2 - \nu_0^2} \tag{2.99}$$

where g is the coupling (here written without the matrix element $n_{\alpha,\beta}$ for notational easy, the generalized coupling appearing in the JC model would be $g|n_{\alpha,\beta}|$), $n_{\alpha,\beta}$ and $\nu_{\alpha,\beta}$ are the charge matrix element and transition frequency between levels α, β, and ν_0 is the cavity frequency. Fig. 2.11 shows the computed chi shifts for each level using the full transmon Hamiltonian to compute the charge matrix element. For the transmon qubit, this expression simplifies greatly as the charge operator only couples consecutive elements with a harmonic oscillator type scaling: $\langle n+1|\,\hat{n}\,|n\rangle = \sqrt{n+1}$ (where here confusingly, $\hat{n}$ denotes the charge

operator and not the typical number operator for harmonic oscillators). The ground state chi shift is unchanged (to first order, a more complete calculation replaces the harmonic oscillator expression derive above with the full transmon raising and lowering operators):

$$\chi_g = \frac{(g_{ge})^2 2\nu_{ge}}{\nu_{ge}^2 - \nu_0^2} \tag{2.100}$$

where $g_{ge} = gn_{ge}$. The excited state now has two contributions: one from the $|g\rangle$ level and another from the $|f\rangle$ level:

$$\chi_e = \frac{(g_{eg})^2 2\nu_{eg}}{\nu_{eg}^2 - v_0^2} + \frac{(g_{ef})^2 2\nu_{ef}}{\nu_{ef}^2 - \nu_0^2} \tag{2.101}$$

which can be simplified using $g_{ef} = \sqrt{2}g_{ge}, \nu_{ge} = -\nu_{eg}, \nu_{ef} = -(\nu_{eg} - E_c)$. Calculating the differential dispersive shift $\chi_{eg} = \chi_e - \chi_g$:

$$\chi_{eg} = \frac{4g_{eg}^2\nu_{eg}}{\nu_{eg}^2 - \nu_0^2} - \frac{4g_{eg}^2(\nu_{eg} - E_c)}{(\nu_{eg} - E_c)^2 - \nu_0^2} \tag{2.102}$$

and using the expansion $\nu_{eg}^2 - \nu_0^2 \approx 2\nu_{eg}(\nu_{eg} - \nu_0)$ for the denominators, we finally obtain:

$$\chi_{eg} = -\frac{2g^2 E_c}{\Delta(\Delta - E_c)} \tag{2.103}$$

where $\Delta = \nu_{ge} - \nu_0$ is the detuning between the qubit and the cavity. We can observe two interesting features: the dispersive shift is *suppressed* by a factor $\frac{E_c}{\Delta - E_c}$ and acquires a secondary pole at $\Delta = E_c$ where the cavity enters resonance with the $e - f$ transition. The frequency range between $\Delta = 0$ and $\Delta = E_c$ is called the *straddling* range where the cavity frequency is between the $g - e$ and $e - f$ qubit transition and should be avoided. In general, placing the qubit $g - e$ transition *below* the cavity transition avoids spurious first order resonances between the cavity and higher-lying states as all the higher states have a lower frequency than the $g - e$ transition. Fig. 2.12 shows the exact dispersive shift using

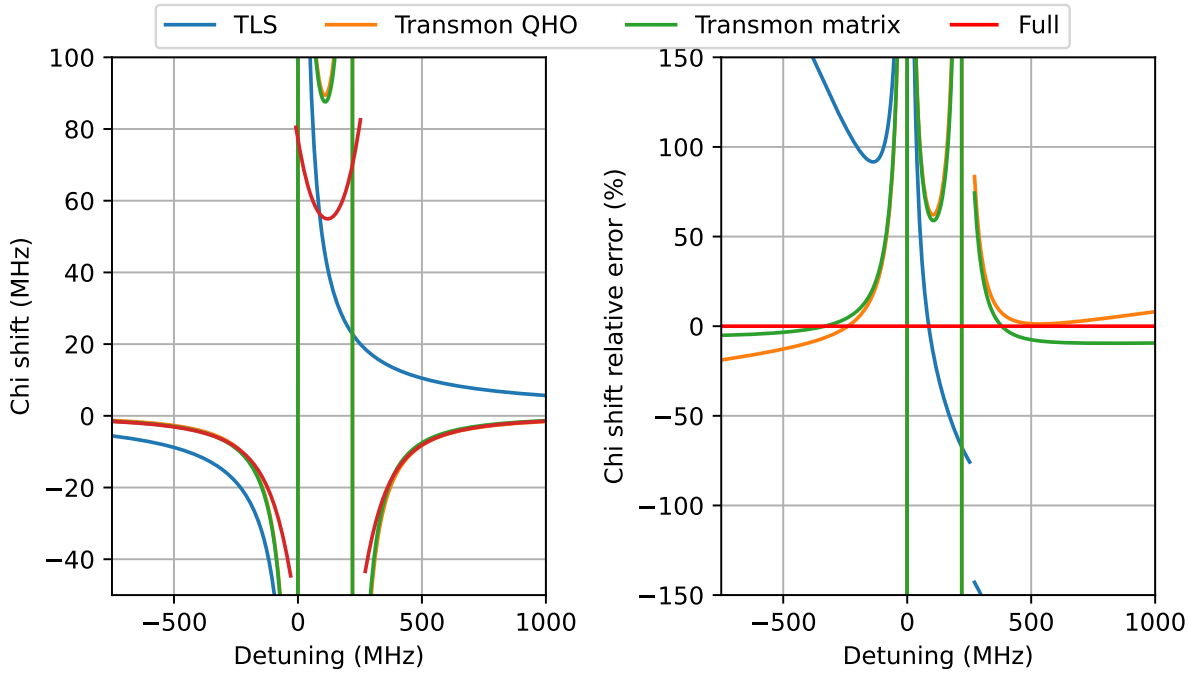

Figure 2.12: Comparing different approximations to calculate the chi shift. Left: chi shift values as a function of detuning between the qubit and cavity, Right: relative error for the computed chi shift to the more exact value incorporating the full coupled Hamiltonian as provided by ScQubits [22, 23]. The two-level system (TLS) approximation (used in the JC section of this chapter) dramatically over estimates the dispersive shift over the entirety of the detuning range. Adding in the contribution of the $|f\rangle$-state from the weak QHO model (transmon QHO, Eq. 2.103) of the transmon improves the agreement to within 10-20% which is improved by incorporating the exact matrix element on the $e-f$ transition instead of using the QHO approximation (transmon matrix, Eq. 2.101)

the full transmon Hamiltonian with comparisons to the calculated dispersive shift as well as the two level approximation.

Chapter 3

Readout and Coherent Control

In this section, we will learn how to perform measurements and apply control signals on the qubit. We will start by laying out how we calibrate a standard qubit device (both for fixed frequency qubits as well as flux tunable ones) using both continuous-wave (CW) measurements and pulsed measurements and briefly discuss qubit population normalization before moving on to coherent control of both the qubit and the cavity.

3.1 Initial device characterization

So you've finally made a device and cooled it down, congrats! Now, how can we check that it is working correctly?

The initial set of measurements is designed to quickly determine if the device is working properly. For a standard device consisting of a qubit coupled to one or more resonators, this consists of establishing that the readout resonator is at the correct frequency, that the qubit is functional, and, for tunable devices, that the frequency of the qubit can be changed in situ. Once these parameters have been established, depending on the use case, more extensive characterization can be undertaken.

3.1.1 Finding and characterizing the cavity

For circuit QED architectures, everything starts and ends with the readout cavity. As described in the previous chapter, this usually consists of a CPW resonator in a $\lambda/2$ or $\lambda/4$ configuration coupled to a feedline used to send signals to and from the chip. To establish notation in the electrical engineering way, we will label the input port as "Port 1" and the output port as "Port 2". We distinguish two types of coupling sketched out with typical responses in Fig. 3.1. The first is the "hanger" geometry where the feedline is uninterrupted between the input and output ports of the chip and the resonator *hangs* off the feedline. In this configuration, the resonator acts as a notch filter, cutting out transmission from port 1 to 2 at the resonance frequency. The second is the "transmission" geometry where the input and output of the resonator are coupled to stubs leading to the ports off the chip. The *transmission* geometry only allows for signals to propagate from port 1 to 2 on resonance acting as a band pass filter. For those more familiar with optics terms, the hanger geometry acts like a *reflection* measurement with a circulator, whereas the transmission geometry gives us the *transmitted* signal through the cavity.

The notch-filter behavior of the hanger geometry allows for multiple resonators with different frequencies to be coupled to the same feedline. This multiplexing is extremely useful for fast prototyping of qubits as many qubit-readout resonator pairs can be placed on the same chip and measured using only a single feedline. In addition to multiplexing, the broadband nature of the feedline guarantees a measurable signal over a wide frequency range allowing for independent tuning of possible amplifiers on the output side without requiring a full chip characterization. Finally, because the response of a hanger resonator is an interference between the signal coming from the resonator and the base signal in the feedline, it is a *self-referenced* measurement allowing for an absolute measurement of the losses in the resonator independent of losses in the lines leading to and from the device.

On the other hand, the transmission style geometry provides filtering of the various non-readout pulses applied to the cavity (for example, qubit control pulses are applied to

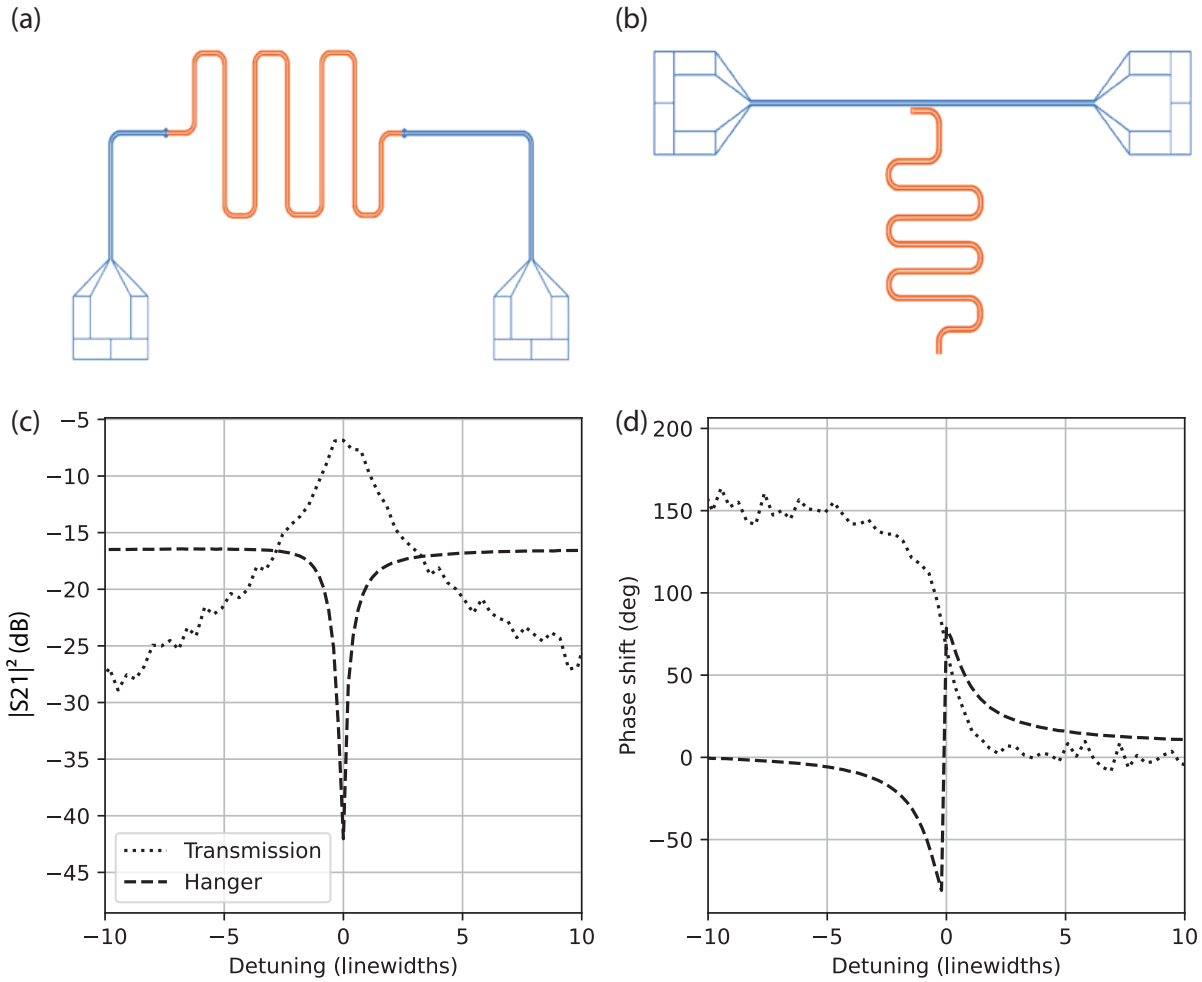

Figure 3.1: The two coupling geometries used for readout resonators: transmission style (a) and hanger style (b) where the feedline is blue and the resonator is shown in orange. The magnitude (c) and phase (d) of response for each resonator is plotted as a function of detuning from the resonance frequency scaled by the linewidth. These measurements were taken with with different powers (as they occurred on different devices) leading to the difference in signal to noise ratio between the hanger and transmission style measurements.

the cavity and typically have much higher powers than readout pulses), which can reduce saturation effects in the output line amplifiers. Additionally, using asymmetric input and output coupling strengths, one can achieve a more directional emission of the readout signal, increasing the readout signal to noise ratio. In contrast to the self-referenced nature of the hanger geometry, transmission-style resonators are unable to separate transmission losses from all other loss sources on the lines (input attenuation, packaging losses, etc.) without a dedicated line calibration.

In both cases, we perform a S_{21} measurement using the Vector Network Analyzer (VNA) to find the resonance frequency of the cavity. S_{21} is a fancy electrical engineering term to say that we are measuring the amplitude and phase at port 2 relative to the signal injected at port 1. The VNA sweeps the frequency of the probe tone injected into port 1 and measures the *linear* response (i.e. measures the output signal at the same frequency) at port 2 while using an internal phase and amplitude reference to compute the gain and phase shift acquired by the signal going through the device. Our R&S VNA has over 100dB dynamic range (with a maximum at 140dB for the narrowest bandwidth measurements) allowing it to measure a power ratios of 1 part in 10 billion (because power scales as V^2, this only maps to a part in 100 thousand for voltages)!

The typical S21 response for a hanger resonator is:

$$S_{21}(f) = 1 - \frac{Q/Q_c}{1 + 2iQ\frac{\delta f}{f_0}} \tag{3.1}$$

where Q is the total quality factor of the resonator, Q_c is the coupling quality factor, f_0 is the resonance frequency, and $\delta f = (f - f_0)$ is the detuning from resonance. From the coupling quality factor, as well as the total quality factor, we can extract Q_i (using $Q^{-1} = Q_c^{-1} + Q_i^{-1}$) to estimate the internal losses due to interfaces and other impurities. A more realistic model that incorporates impedance mismatches and other effects of the imperfect microwave environment is discussed in Appendix A. While a complete fit would

give exact values for f_0, Q, and Q_c for quick estimates we can use a few simple heuristics. The resonance frequency f_0 occurs at the minimum of S_{21} and the total quality factor is extracted from the ratio of f_0 to the linewidth Δf which is approximated by the bandwidth of the 3 dB points (measured relative to the *baseline*). For highly overcoupled devices ($Q_i > Q_c$), the internal quality factor Q_i can be estimated from the contrast in the S_{21} curve: $\frac{\max - \min}{\min} \times Q$ [30] (for typical VNA-based S_{21} measurements given in dB instead of linear units, the expression becomes $\frac{\max(S21) - \min(S_{21})}{20} \times Q$).

For transmission style resonators, S_{21} looks very similar:

$$S21(f) = \frac{Q/Q_c}{1 + 2iQ\frac{\delta f}{f_0}} \tag{3.2}$$

except we lose the interference term. If we model the input/output lines as applying some phase factor ϕ and attenuation α, we can see that it becomes impossible to separate attenuation from quality factor for transmission measurements whereas the self-referencing of the hanger measurement makes it insensitive to loss (as we are measuring a dip *relative* to the baseline instead of an absolute signal):

$$S21(f) = e^{\alpha l} e^{i\phi} \left(1 - \frac{Q/Q_c}{1 + 2iQ\frac{\delta f}{f_0}}\right) \tag{3.3}$$

$$S_{21}(f) = e^{\alpha l} e^{i\phi} \left(\frac{Q/Q_c}{1 + 2iQ\frac{\delta f}{f_0}}\right) \tag{3.4}$$

where ϕ is typically the phase acquired due to linear dispersion in the transmission lines ($\phi = 2\pi l/\lambda$ where l is the electrical length and λ is the wavelength).

In either case, once the cavity frequency is known (usually by just taking the minimum (maximum) for hanger (transmission) style resonators), we can start looking for a qubit. To quickly identify whether the qubit works, we can perform a power sweep on the cavity. At high power, the cavity decouples from the qubit leading to a frequency shift of the cavity to its bare resonance frequency [2] (Fig. 3.2 shows a typical power scan measurement). A

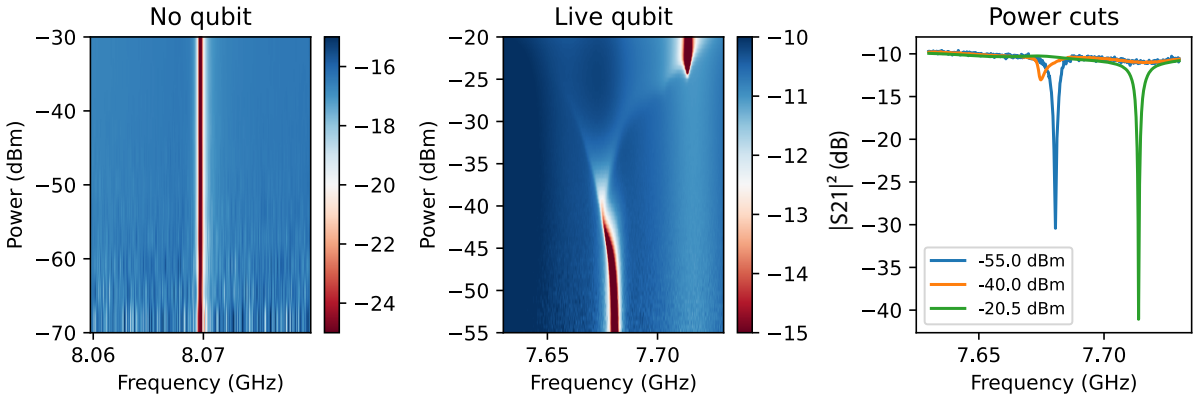

Figure 3.2: Transmission power scans for a simple cavity (left) and a cavity-qubit system (middle) with associated power cuts shown for reference (right), the colormap displays $|S_{21}|^2$(dB). The cavity readout power is set *below* the power at which the cavity decouples from the qubit (around -45 dBm for the data shown here)

power dependence on the cavity frequency is typically a good indicator that there is a qubit in it, although it really only tells us that there is something non-linear in the system.

3.1.2 Finding the qubit

To find the qubit resonance frequency we perform two-tone spectroscopy, that is, we monitor S21 at the cavity frequency and sweep a probe tone around the expected qubit frequency. As discussed in Chap. 2, the cavity resonance frequency depends on the state of the qubit, and so exciting the qubit should be visible if we monitor transmission at the original frequency. The resonance frequency we found in the cavity measurement corresponds to the $|g\rangle$ state on the qubit as the cryogenic environment leads to an almost perfect ground state initialization. A continuous drive near the qubit frequency will drive it to a mixed state, leading to a bimodal cavity transmission spectrum (the combination of the $|g\rangle$ and $|e\rangle$ resonances). By plotting $S_{21}(\omega_g)$ (i.e. S_{21} at the base cavity resonance frequency as a function of the applied tone frequency) the qubit will appear as a peak (dip) for hanger (transmission) style resonators. In addition to finding the ω_{ge} frequency, we can also get an estimate of the anharmonicity α by increasing the drive power. At high enough powers,

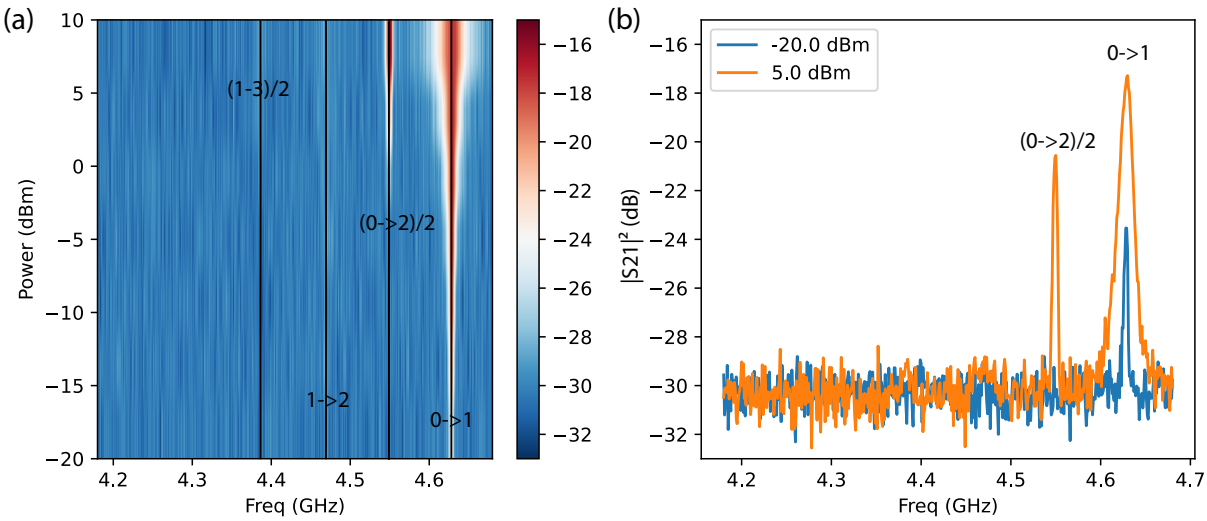

Figure 3.3: Qubit drive power scan in two-tone spectroscopy. At low powers only the $|g\rangle - |e\rangle$ transmission is visible (around 4.62 GHz) but higher order photon transitions appear at higher power. The faint line visible at low power for the left panel (around 4.48 GHz) indicates residual $|e\rangle$-state population from imperfect thermalization.

a peak at the two photon resonance from $|0\rangle$ to $|2\rangle$ will appear ($\omega_{gf}/2$). A typical power scan for a qubit device is shown in Fig. 3.3 with the 0-1 transition and 0-2 transitions identified. For a properly thermalized qubit, only transitions from 0 to n should be visible as all higher states should be empty. In particular, due to the transmon acting as a harmonic oscillator, transmissions between levels further than 1 index apart are suppressed and only the (0→n)/n lines appear as multi-photon transitions. A good diagnostic of poor thermalization is the appearance of lines at the 1-2 and 2-3 transition frequencies.

3.1.3 Flux tunable devices

For flux tunable devices, the above procedures need to be repeated over an entire flux quantum to identify the frequency tuning range of the qubit. For faster measurements, the cavity can be used to identify the applied current/voltage to flux conversion. Although the cavity itself is not tunable, by coupling it to the qubit, it acquires a weak flux dependence (recall, $\chi \propto g^2/\Delta$). For some devices, the qubit crosses the cavity leading to an avoided crossing which can be used to calculate the coupling to the cavity (see Chapter 5). Varying

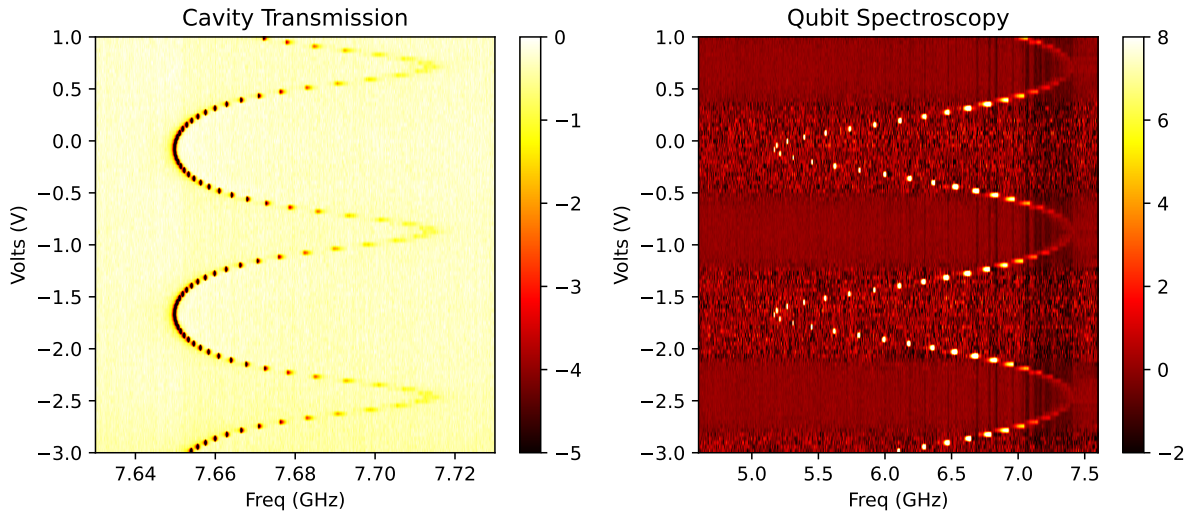

Figure 3.4: Transmission and spectroscopy scans for a flux tunable qubit. Due to coupling to the qubit, the readout cavity acquires a small degree of flux sensitivity. Near resonance, the cavity system becomes strongly hybridized with the qubit turning it into a near single photon source. This reduces the output signal drastically.

the control parameter (voltage or current) will lead to periodic modulation of the cavity frequency which can be used to identify the conversion to applied flux in units of flux quanta. The qubit frequency can then be mapped as a function of applied flux by first performing a transmission measurement to find the cavity resonance frequency before doing two-tone spectroscopy. Due to the potentially large dynamic range in detuning between the qubit and cavity, different readout powers might be required as the qubit frequency changes although the power used to measure the flux periodicity in the earlier transmission measurements is a good place to start (for instance, for qubits approaching the readout cavity frequency, the readout power would need to be *lowered* to avoid saturating the system and decoupling the qubit from the cavity). Some example data is shown in Fig. 3.4 for a tunable transmon clearly showing the periodic behavior of the flux bias.

3.1.4 CW measurement summary

The goal of the CW measurements is to quickly identify wether a device performs correctly and to extract most of the design parameters of the device. The cavity readout frequency

and power are extracted from pulsed transmission measurements, the qubit ge frequency, anharmonicity, and thermalization are extracted from the two-tone spectroscopy power scan, and current or voltage sweeps for flux tunable devices allow us to convert from laboratory units (A or V) to flux on the qubit Φ/Φ_0. Finally, the qubit-cavity coupling strength can also be extracted from a measurement of the χ-shift on the cavity or from an explicit avoided crossing (if the qubit crosses resonance, the resulting crossing can be fit to extract g).

3.2 Pulsed Measurements

Once the basic characteristics of the device are measured using CW measurements, we can switch to pulsed measurements to start performing actual gates and to treat the system as an actual coherent quantum system instead of just a thermal mixture of coupled and anharmonic oscillators.

3.2.1 Pulsed Transmission

Just as for the CW measurements, the first step is finding the cavity resonant frequency. While the CW calibration described above already provides frequencies for both the readout and qubit, small frequency shifts occur when switching from continuous measurements to pulsed measurements, in particular on the qubit frequency (as the qubit and cavity drives no longer overlap and the qubit can fully relax to the $|g\rangle$ state between measurements). The pulsed transmission measurement also provides a quick way to optimize the readout settings as we switch from the VNA to the pulsed set up. To perform pulsed measurements, we use an arbitrary waveform generator (AWG) to modulate the output of the generator sending signals to the dilution fridge. The qubit control pulses require both amplitude and frequency or phase control (although for basic calibration, only amplitude modulation is required) which is performed with IQ modulation, described in the following section. The

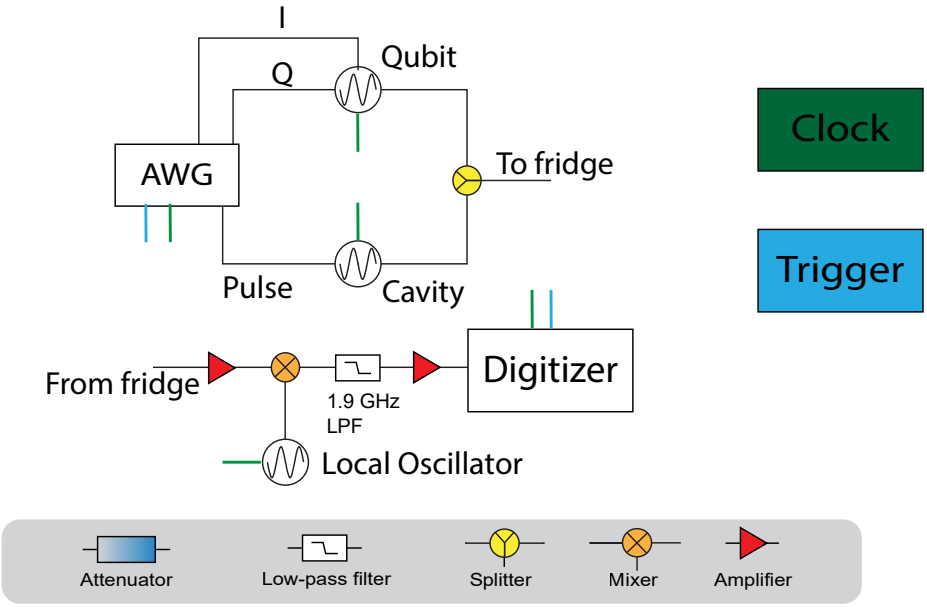

Figure 3.5: Schematic representation of the hardware required for pulsed measurements. On the input side, an arbitrary waveform generators produces the pulses for the qubit and cavity drives. On the output, the cavity signal is downconverted before digitization. The trigger generator (light blue) sends a pulse at the beginning of every measurement to synchronize the AWG and the digitizer. An atomic clock (dark green) is used to provide a common frequency reference to all the equipment.

return signal is then downconverted and digitized on an analog to digital converter card (ADC). A trigger generator is used to synchronize the AWG and digitizer and sets the experimental repetition rate. A simplified wiring diagram is shown in Fig. 3.5 with the full detailed diagram left to Chap. 4.

A measurement consists of measuring the phase and amplitude of the cavity pulse at the resonance frequency. While in principle it is possible to directly digitize signals in the 5 to 9 GHz frequency range used by our readout cavities, these digitizers are extremely expensive and completely overkill. By multiplying our cavity pulse with another RF tone (and filtering out the other mixing products), we can downconvert it to a lower frequency which we can directly digitize. This process is called downconversion or heterodyning and is used virtually everywhere in communications (optical fibers, radio, WiFi signals, etc.). The same process can be used in reverse to upconvert a *baseband* containing information to a frequency used for data transmission.

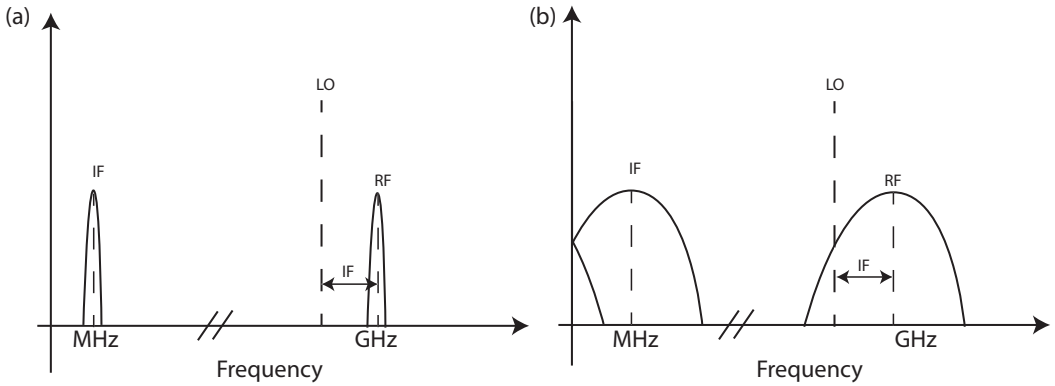

Figure 3.6: Downconversion of signals from a high frequency baseband to a lower frequency baseband. (a) Typical configuration where a signal at frequency RF is mixed with a local oscillator detuned by an offset frequency IF which shifts the RF spectrum down to the detuning frequency IF. (b) When the detuning between the local oscillator and RF signal is smaller than the bandwidth of the incoming RF spectrum, some of the frequency components (below the LO frequency in this case) will get mirrored back onto the spectrum when shifted down to IF.

A primer on up/down conversion, IQ points and IQ modulation

The product of two oscillatory signals at frequencies $\omega_{1,2}$ generates sum and difference frequencies:

$$A\sin(\omega_1 t+\phi)*\sin(\omega_2 t)=\frac{A}{2}\left[\cos((\omega_1-\omega_2)t+\phi)-\cos((\omega_1+\omega_2)t+\phi)\right]. \tag{3.5}$$

By filtering the output with a low-pass or high-pass filter we can isolate either the sum or the difference frequency corresponding to up or downconversion of our signal. In general, a signal with frequency components centered around ω_0 multiplied by a *local oscillator* signal at frequency ω_{LO} will be shifted to two bands centered at $|\omega_0-\omega_{LO}|$ and $|\omega_0+\omega_{LO}|$. Some care must be taken when downconverting signals with a wide bandwidth baseband to avoid "mirroring" of the spectrum at $\omega=0$ as shown in Fig. 3.6. The intermediate frequency, $\omega_{IF}=\omega_0-\omega_{LO}$, signal still contains all the amplitude and phase information of the original microwave signal and can now be further processed or digitized. This process can also be reversed to generate modulated signals at frequencies much higher than the base bandwidth of the waveform generator by mixing the baseband *up* to the desired carrier frequency with

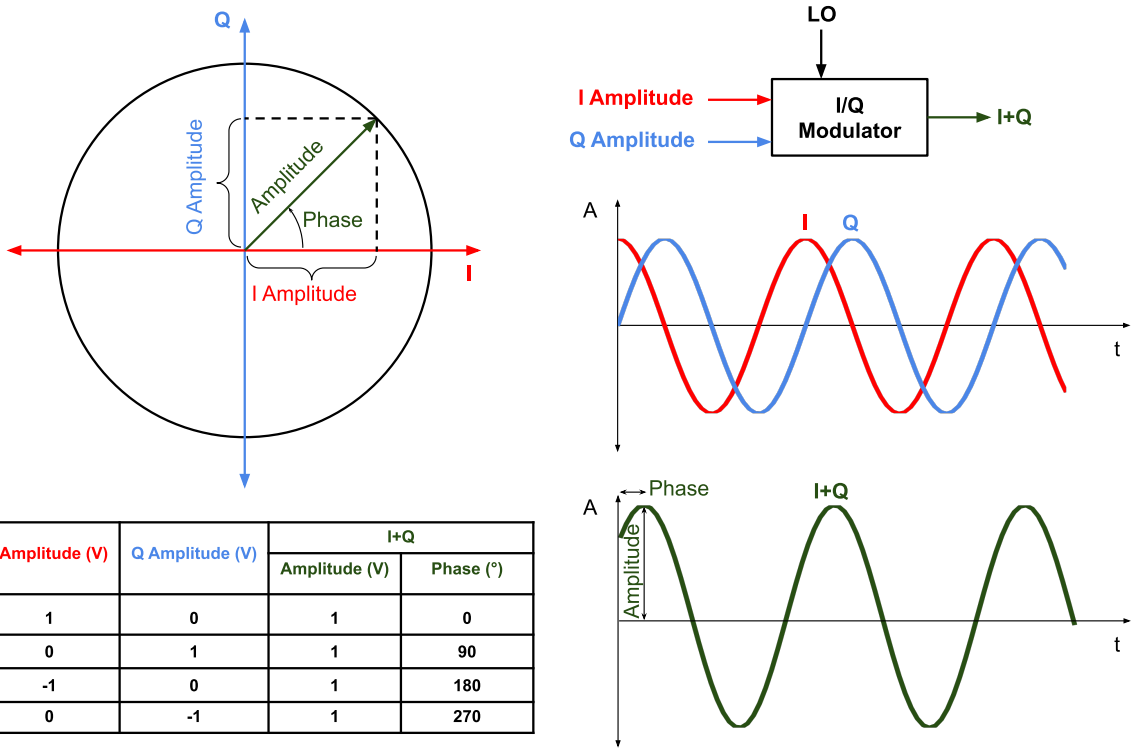

I Amplitude (V)	Q Amplitude (V)	I+Q	
		Amplitude (V)	Phase (°)
1	0	1	0
0	1	1	90
-1	0	1	180
0	-1	1	270

Figure 3.7: Visualization of IQ signals on the unit circle (top) and for time dependent signals (right panels). Figure credit: Vigneshdm1990 (`https://commons.wikimedia.org/wiki/File:IQ_phasor_diagram.svg`,"IQ phasor diagram", `https://creativecommons.org/licenses/by-sa/4.0/legalcode`

an LO at the frequency of interest (typically 4-7 GHz in our applications). While a simple mixer as described above can perform up and down conversion, it comes with the drawback that we lose half the amplitude of the signal we wanted to the other sideband which we also need to eliminate. We will see that combining in-phase (I) and quadrature (Q) signals with an IQ mixer gives us full control over which sideband is generated and even the amplitude and phase of each sideband.

While a signal is typically defined with an amplitude and phase (which we can think of as a polar representation of the phasor), it is actually simpler, from a modulation standpoint, to use a Cartesian representation of the signal (I and Q):

$$A\cos(\omega t+\phi) = A\left[\cos(\omega t)\cos(\phi) - \sin(\omega t)\sin(\phi)\right) = I\cos(\omega t) - Q\sin(\omega t) \tag{3.6}$$

where

$$I = A\cos(\phi) \qquad A = \sqrt{I^2 + Q^2} \tag{3.7}$$

$$Q = A\sin(\phi) \qquad \phi = \arctan\left(\frac{Q}{I}\right) \tag{3.8}$$

are the usual polar to Cartesian conversions, see Fig. 3.7 for a visual representation. While the distance between two points is independent of the representation (polar vs Cartesian), using IQ points removes the distinction between changes mostly in amplitude or phase making the data analysis more straightforward. We can compute the I and Q components of a signal by multiplying it by cosine and sine at the frequency of oscillation:

$$A\cos(\omega t + \phi)\cos(\omega t) = A\left[\cos(\omega t)^2\cos(\phi) - \sin(\omega t)\cos(\omega t)\sin(\phi)\right] \tag{3.9}$$

$$= \frac{A}{2}\left[(1 + \cos(2\omega t))\cos(\phi) - \sin(2\omega t)\sin(\phi)\right] \tag{3.10}$$

$$\overset{LPF}{\Rightarrow} \frac{I}{2} \tag{3.11}$$

$$A\cos(\omega t + \phi)\sin(\omega t) = A\left[\cos(\omega t)\sin(\omega t)\cos(\phi) - \sin(\omega t)^2\sin(\phi)\right] \tag{3.12}$$

$$= \frac{A}{2}\left[\sin(2\omega t)\cos(\phi) + (1 - \cos(2\omega t))\sin(\phi)\right] \tag{3.13}$$

$$\overset{LPF}{\Rightarrow} \frac{Q}{2} \tag{3.14}$$

by filtering the output with a low-pass filter ($\overset{LPF}{\Rightarrow}$) or averaging over a period of oscillation we recover the I and Q components of the original signal. This is effectively how the downconversion process in an IQ mixer works: the incoming signal is split into two branches which are mixed with the cos and sin of the LO generating the I and Q signals at the output ports.

Arbitrary Modulation with IQ signals

The IQ mixer can also be used to perform arbitrary amplitude and phase modulation on a microwave tone. We will examine three different kinds of modulation although, fundamentally, the only limitations on the achievable output waveform are the bandwidth of the components and IQ imperfections (which will be discussed at the end of this section). The simplest form of modulation is simply amplitude modulation. For example, by modulating I we directly vary the output amplitude of the RF waveform:

$$s(t) = I(t)\cos(\omega t) = (1 + \beta\cos(\omega_m t))\cos(\omega t) \tag{3.15}$$

where β is the modulation depth and ω_m is the modulation frequency. Note, for pure amplitude modulation, β should stay below 1 to avoid imprinting phase modulation on the signal. The exact same principle applies when using the Q channel for modulation except for a 90 degree phase shift in the carrier frequency. The spectrum for amplitude modulation is characterized by sidebands on either side of the carrier at frequencies $\omega - \omega_m$ and $\omega + \omega_m$ with amplitudes set by $\beta/2$ relative to the main carrier power.

By varying I and Q out of phase, we perform phase modulation. Phase modulation and frequency modulation are effectively equivalent as they are related by a derivative/ integral relationship. Consider the RF signal $A\cos(\gamma)$, where $\gamma = \omega t + \phi$ for the simple fixed-frequency case, while in this case we directly identify the frequency of oscillation as ω, in general, we can define the frequency as the derivative of the phase argument in the cos:

$$\omega = \frac{\partial\gamma}{\partial t} \tag{3.16}$$

which can have arbitrary time dependence. In general, we desire to generate a well defined time dependent frequency for our signal (for example ramping the frequency between two points) however, the I and Q components are defined as functions of ϕ, the phase of the signal: $I = \cos(\phi), Q = \sin(\phi)$. To convert frequency modulation to phase modulation,

we simply integrate the desired frequency trajectory. Let's consider three simple cases: a constant frequency shift of the carrier frequency, a frequency chirp, and sinusoidal frequency modulation.

A constant frequency shift on the carrier is used to rapidly change the frequency of the carrier within an experimental cycle (faster than the instrument can change its frequency setting in the actual hardware). As we discussed earlier, a fixed frequency signal has a phase proportional to ωt which means that linearly ramping the phase of the carrier changes the frequency by a fixed amount (positive ramps shift ω_0 to $\omega_0 + \omega_{IF}$ and negative ramps shift to $\omega_0 - \omega_{IF}$). This is easily achieved with $I = \cos(\omega_{IF} t)$ and $Q = \sin(\omega_{IF} t)$ where the upper or lower sideband is generated depending on the sign in front of the Q component. For the frequency chirp, $\omega_{IF}(t) = \alpha t + \beta$ where α and β are the frequency sweep rate and initial frequency offset. Here the required modulation phase is simply $\frac{\alpha t^2}{2} + \beta t$ which we can again plug into our expressions for I and Q in the modulator. Finally, some applications require periodic frequency modulation (see Chap. 7 for example applications). This modulation is specified by two parameters: the frequency deviation, $\Delta\omega$ (i.e. how far the frequency is shifted at the extrema of the modulation signal) and the modulation frequency ω_m. The required phase modulation is obtained by integrating the periodic frequency:

$$\phi(t) = \int \Delta\omega \cos(\omega_m t) = \frac{\Delta\omega}{\omega_m} \sin(\omega_m t) \tag{3.17}$$

where the prefactor is sometimes referred to as β, the phase modulation depth.

In principle, any frequency/amplitude modulation can be achieved with IQ modulation by following the method outlined above giving us exquisite control over the microwaves signals used to drive our quantum systems provided the bandwidth of the modulator is sufficient.

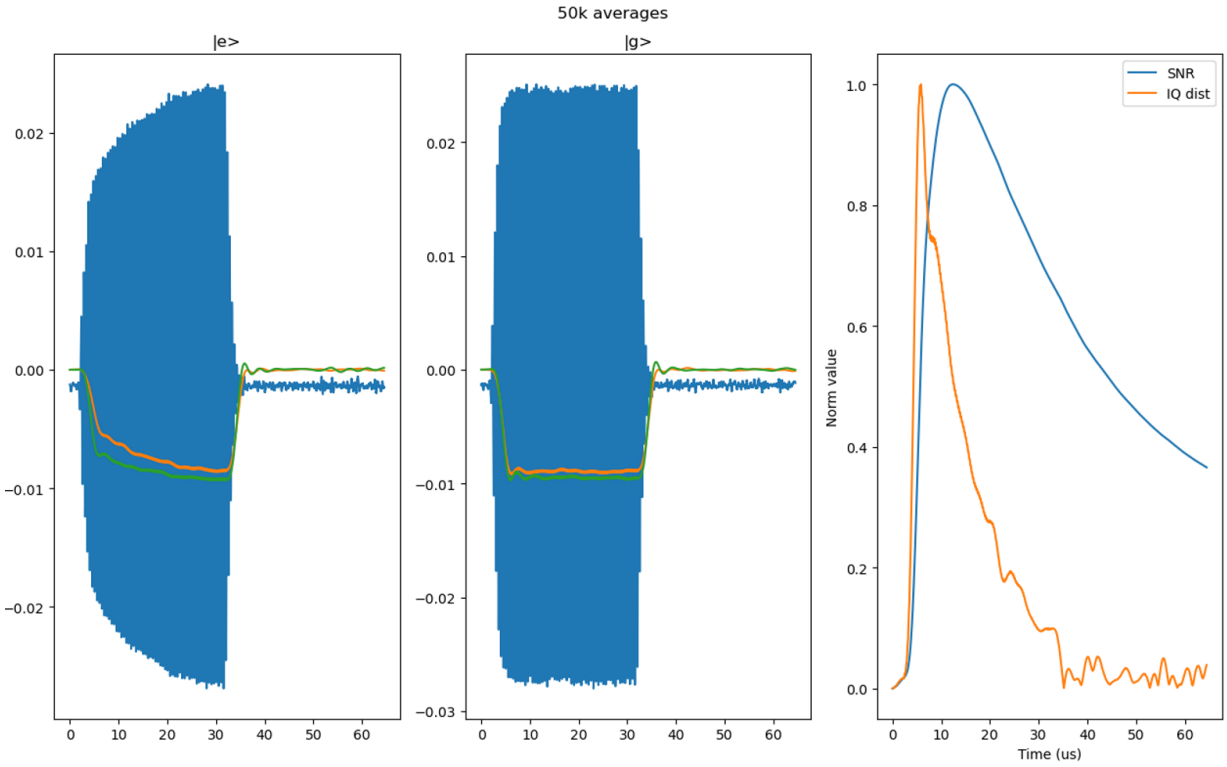

Figure 3.8: Cavity response (at ω_{IF}) for the qubit in $|g\rangle$ vs $|e\rangle$ with the extracted IQ components overlaid for reference. In the first two panels, the blue trace is the raw voltage acquired by the DAC card and the green/ orange traces correspond to the filtered I and Q signals. The third panel shows the IQ distance between the two traces showing a rapid rise as the cavity initially fills with photons before decaying as the qubit $|e\rangle$-state population decays back to g. The signal to noise ratio (defined as the integral of the signal divided by the characteristic noise calculated from the variance in the traces after the pulse) is also plotted showing a clear peak in integration time roughly corresponding to T_1 (around 10-15 μs for this device).

3.2.2 Readout

Now that the reader is hopefully a bit more familiar with IQ points, we can describe our readout with the typical jargon used in the field. We perform cavity state readout by measuring the I and Q components of an RF pulse sent to the device. To reduce the required bandwidth of the ADC card, we first downconvert the cavity signal from 5-7 GHz to a 2 MHz IF signal which is then filtered, amplified and digitized. We then perform digital downconversion (DDC) to extract the I and Q components of the pulse. A typical pulse response is shown in Fig. 3.8 showing the DDC process. By sweeping the probe frequency we can find the resonance frequency of the cavity, just as with the VNA in the CW measurement section.

For all the subsequent measurements, the qubit is assumed to be initialized in the $|g\rangle$ state. In pulsed measurement mode, the repetition rate of the experiments should be set such that it exceeds 5-10 T_1 (or at least a best guess as to the expected T_1 value of the qubit). This guarantees that the qubit will have decayed back down the a pure $|g\rangle$ state before the start of the following measurement.

3.2.3 Pulsed qubit spectroscopy

We can find the qubit frequency by pulsing the qubit and measuring the response of the cavity. As we've discussed earlier, the resonance frequency of the cavity shifts based on the qubit. A resonant frequency shift of the cavity can appear as mostly a magnitude signal ($\chi > \kappa$), mostly a phase signal ($\chi < \kappa$), or both ($\chi \approx \kappa$). By computing the IQ distance between two measurements, both amplitude and phase shifts are converted to a 1-dimensional number instead of having to account for the two shifts separately. In pulsed spectroscopy, we pulse the qubit with a 10 to 20 μs pulse immediately followed by a readout pulse of length 5 to 10 μs. At this stage, it is important to use a *long* qubit pulse to drive it to a somewhat mixed state and avoid accidentally driving the qubit perfectly back to the ground state with an $2n\pi$ pulse. Additionally, a long qubit pulse allows for a

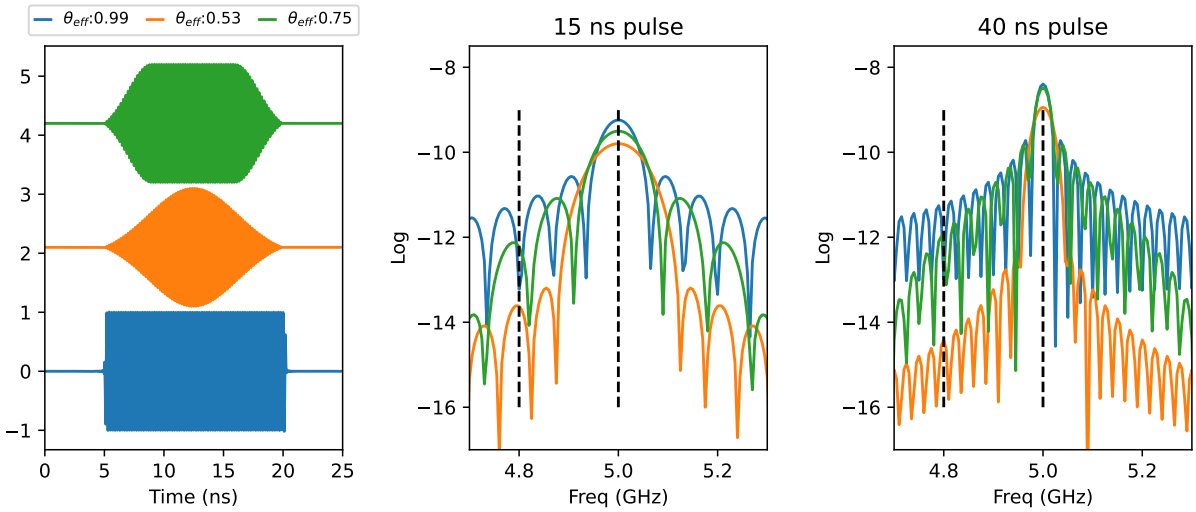

Figure 3.9: Comparison of different pulse shapes. (a) Time traces for square pulses (blue), gaussian pulses (orange), and gaussian-ramp square pulses (green) with their effective rotation angle for a given drive power (θ_{eff}). (b), (c) Fourier spectrum for the pulses in panel (a) for two different pulse lengths showing change in spectral width as a function of pulse shape and length.

higher resolution frequency sweep as the spectral width of the pulse is narrow. The qubit frequency is extracted by fitting a Lorentzian to the IQ distance vs frequency trace acquired (although simply picking the maximum amplitude is sufficient for a first calibration pass).

While in principle one could use square pulses to drive the qubit, this leads to spurious excitations outside of the computational subspace due to the high frequency content in the pulse. Fig. 3.9 shows the Fourier transform for square, gaussian, and gaussian square pulse shapes of length 50 ns. By rounding the square pulse shape, we drastically reduce power at the $|e\rangle - |f\rangle$ frequency which is only $\alpha \approx 250$ MHz away from the resonance frequency. While the gaussian pulse shape has the narrowest frequency spread it has a much lower effective drive rate due to its rapid fall-off in the time domain. Our qubit gates consist of a square wave of variable length with a fixed length gaussian ramp truncated at $\pm 2\sigma$ with $\sigma = 2.1$ ns.

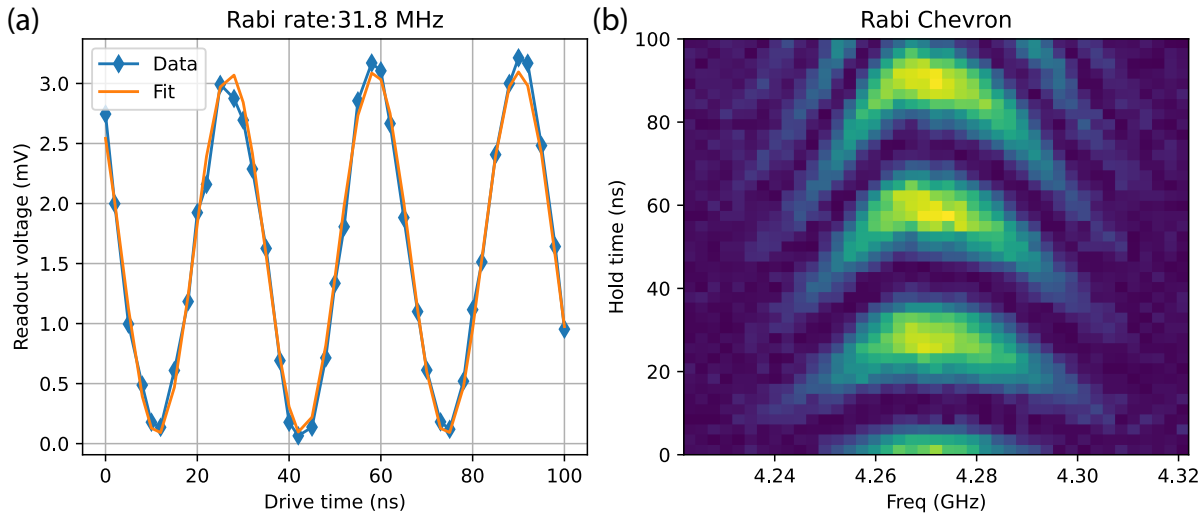

Figure 3.10: (a) As the drive length on the qubit increases, the qubit oscillates between $|g\rangle$ (high voltage) and $|e\rangle$ (low voltage) where the period of oscillation gives the Rabi rate. (b) By sweeping the drive frequency we can map out the *Rabi chevron* showing increased Rabi rates with smaller oscillation contrasts as expected from Eq. 3.19.

3.2.4 Rabi oscillations and π-pulses

Once the qubit frequency is found, we can calibrate the first of our gates: the π-pulse. While all the previous measurements only required mixed qubit states to extract a resonance frequency, we are now switching to coherent operations on the qubit state. The π-pulse flips the state of the qubit from $|g\rangle$ to $|e\rangle$ and once calibrated allows us to normalize the readout voltage to convert to qubit population. Recall the Hamiltonian for a qubit driven by an electric field at its resonant frequency:

$$H = \frac{\Omega}{2}\sigma_x. \tag{3.18}$$

For a qubit initialized in the $|g\rangle$ state, the excited state probability will oscillate as $\sin(\Omega t)$ where Ω is the Rabi rate. The excitation probability is acquired by performing repeated measurements and averaging the results. By varying the length of the qubit-drive pulse, we map out these oscillations and identify the π-pulse as the first extremum. Fig. 3.11(a) shows an example curve where the π-pulse time indicated with the vertical line. One can

also keep the pulse time fixed and vary the linear amplitude of the driving field to obtain similar oscillations for which the first extremum corresponds to π-pulse again.

In addition to calibrating the π-pulse, this measurement is also used to calibrate the strength of the driving field as the Rabi rate corresponds to the frequency of oscillation. The Rabi rate is proportional to the *electric field* of the driving tone so doubling the Rabi rate corresponds to a quadrupling of the driving power (6 dB) or a doubling of the I/Q amplitude from the AWG generator.

Off resonance, the effective Hamiltonian acquires a term proportional to σ_z reducing the amplitude of oscillations and increasing the effective Rabi rate as shown by:

$$H = \frac{\Delta}{2}\sigma_z + \frac{\Omega}{2}\sigma_x \tag{3.19}$$

where Δ is the detuning between the qubit frequency and the applied drive. In this case, the excited state probability is $P(e) = \frac{\Omega^2}{\Omega^2+\Delta^2}\cos\left(\sqrt{\Delta^2+\Omega^2}t\right)$ where $\sqrt{\Delta^2+\Omega^2}$ is the effective Rabi rate. Fig. 3.10(b) shows a colormap of the oscillations as a function of drive detuning displaying the characteristic chevrons.

Finally, the rotation axis about which the qubit state rotates is set by the phase of the applied microwave drive. By convention, we associate the I component with a σ_x rotation and Q with a σ_y rotation.

3.2.5 Readout normalization

To convert the readout amplitude to a qubit population we can use the results of the previous calibration:

$$P(e) = \frac{\mathbf{I}\vec{\mathbf{Q}}_{d-g} \cdot \mathbf{I}\vec{\mathbf{Q}}_{e-g}}{||IQ_{e-g}||^2} \tag{3.20}$$

where $IQ_{g,e}$ correspond to reference IQ points with the qubit prepared in $|g\rangle$ vs $|e\rangle$ and IQ_d is the IQ point of the sample to be normalized. This corresponds to taking the projection of the data IQ vector from $|g\rangle$ onto the $g-e$ IQ vector and normalizing to obtain a number

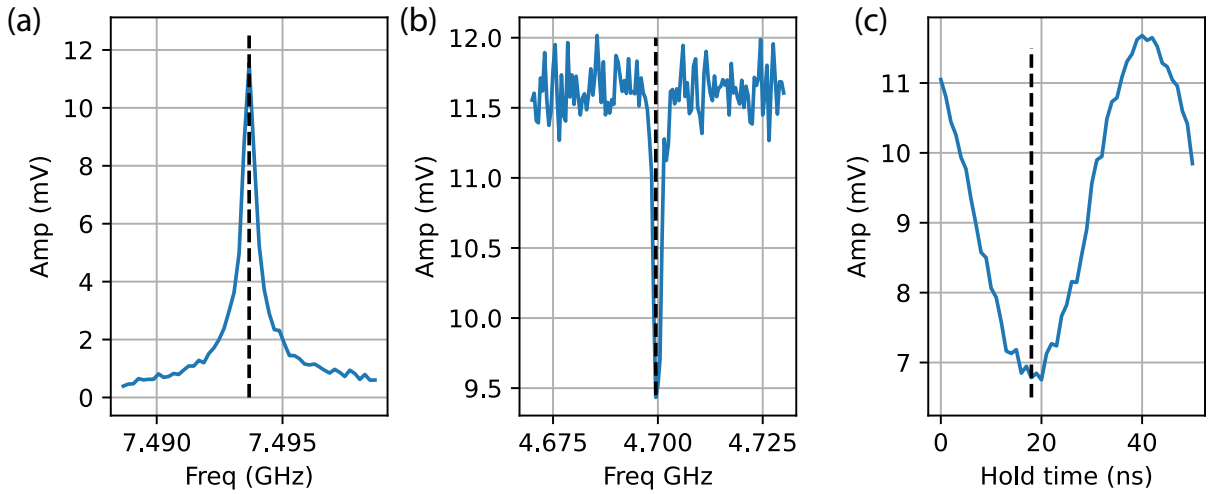

Figure 3.11: Tune up procedure for a typical device: (a) the resonant frequency of the cavity is found, (b) the qubit frequency is found using the long pulse spectroscopy method described in the text, (c) the π-pulse is tuned up by sweeping the length of an applied pulse on the qubit and finding the first extrema (here a minimum as the cavity is measured in transmission meaning the $|e\rangle$-state corresponds to a *low* signal).

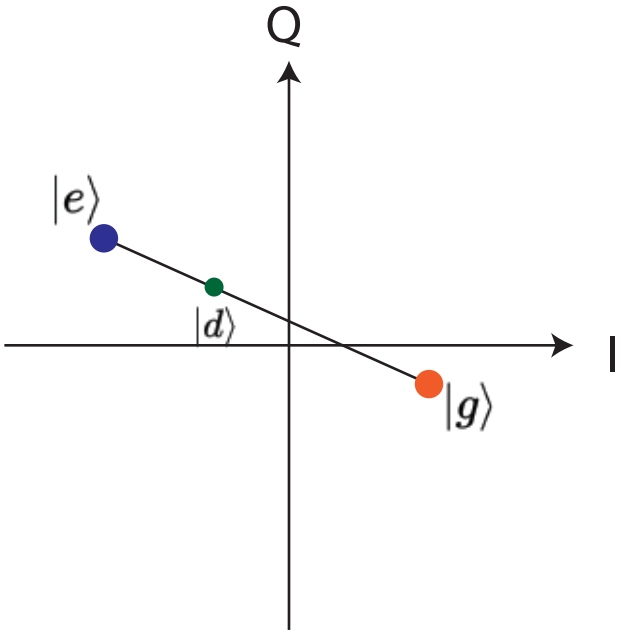

Figure 3.12: By acquiring reference IQ point for $|g\rangle$ and $|e\rangle$ qubit state preparation, we can convert the outcome of a specific result $|d\rangle$ into a fractional population between g and e.

between 0 and 1 as shown in Fig. 3.12.

3.2.6 T_1, T_2, T_{echo}

With a π-pulse in hand, we can now complete our characterization of the qubit by measuring its coherence times [5]. The depolarization time, T_1, corresponds to lifetime of the excited state and is measured by simply π-pulsing the qubit and waiting a variable time τ before readout (as shown in Fig. 3.13(a)). T_1 is extracted from an exponential fit of $Ae^{-t/T1} + B$ to the data. For normalized data, $A = 1$ and $B = 0$, leaving only T_1 as a fit parameter. For an accurate fit, the total evolution time should be $\approx 5T_1$ to provide enough long time data to "pin" the fit. If T_1 is found to be significantly larger than expected, the measurement repetition rate might need to be reduced to allow the qubit to decay back to $|g\rangle$ before the following measurement. In some cases, if the qubit T_1 is significantly longer than expected relative to the time between measurements, the contrast of the measurements will be low as the qubit state becomes effectively random at the beginning of every measurement.

The dephasing time, T_ϕ, corresponds to the lifetime of a superposition of $|g\rangle$ and $|e\rangle$. For a perfect qubit (without phase noise), the coherence time, T_2, is only limited by T_1 as the $|e\rangle$ state portion of the superposition decays to $|g\rangle$ at a rate $e^{-t/2T_1}$. Phase noise (due to fluctuations in the energy splitting of the qubit) is an additional loss channel leading the coherence time to be:

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi} \tag{3.21}$$

where T_ϕ is the characteristic lifetime due to dephasing processes. To measure the Ramsey dephasing time, T_2^*, the qubit is initialized in a superposition with a $\pi/2$ pulse before waiting a time τ and applying a second $\pi/2$ to bring the qubit back to the σ_z basis, a Ramsey sequence. On resonance, this leads to a decaying exponential going to $P(e) = 0.5$. However, if the qubit drive is not perfectly in resonance with the $g - e$ transition, the state will precess during the wait time τ leading to oscillations in the final output state. The

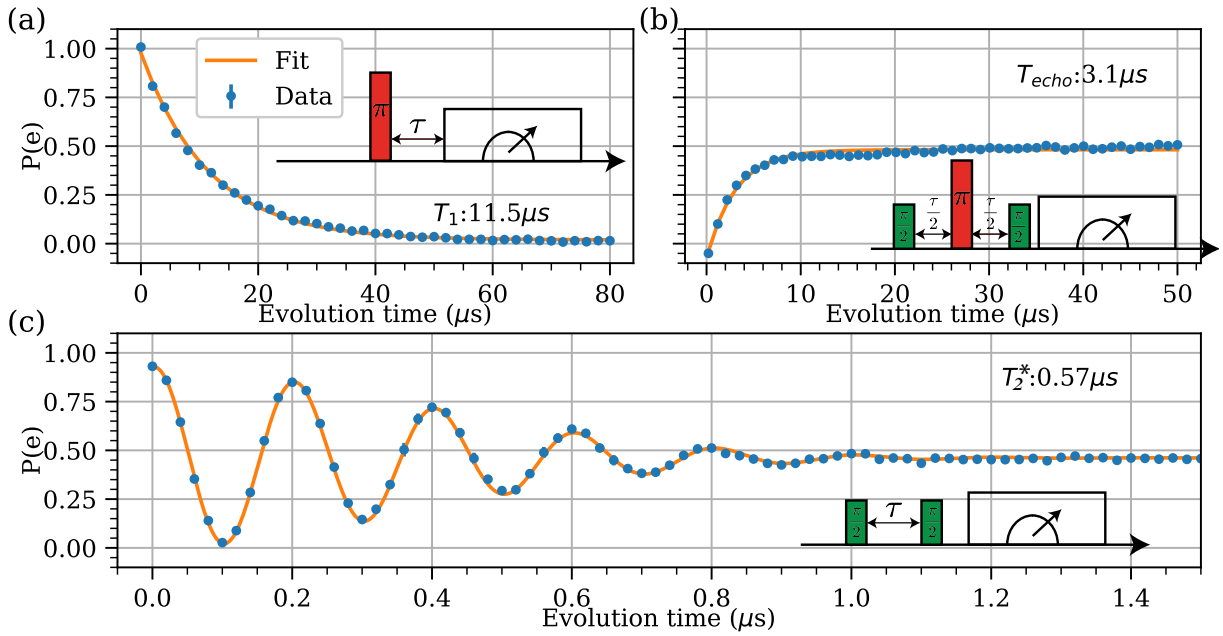

Figure 3.13: Coherence times for the qubit with their associated pulse sequences (red pulses correspond to π-pulses, green are $\pi/2$ pulses, and the long rectangular block at the end is the measurement). (a) T_1 measurement showing an exponential decay of the $|e\rangle$ population as the time between the π-pulse and measurement is increased. (b) T_2 echo measurement where a π-pulse is inserted in between the two $\pi/2$ pulses to remove the effects from DC qubit frequency noise. (c) T_2 measurement using two $\pi/2$ pulses to measure the qubit *coherence* time (i.e. the time for which a superposition of $|g\rangle$ and $|e\rangle$ maintains its phase) under free evolution. The insets show the corresponding pulse sequences.

damped oscillatory response can be fit to extract the decay rate ($1/T_2$) and the qubit-drive detuning (δ) if the two scales are sufficiently separated (as shown in Fig. 3.13(c)).

In particular, near critical damping ($1/T_2 \approx \delta$), the two scales cannot be separated leading to a poor estimation of either parameter. Two methods are used to avoid this issue: detuning the qubit drive frequency deliberately to induce a controlled precession of the qubit or rotating the phase of the second $\pi/2$ pulse to virtually detune the drive. While the direct detuning method is simpler as it only requires shifting the drive frequency, it incurs a cost in the final contrast of the oscillations if the required detuning is large. For example, for short lived-devices with $T_2 < 1-5\mu$s, the required separation of scales leads to 5-10 MHz detuning values which can reduce the contrast of the oscillations as the π-pulses are no longer well calibrated. On the other hand, the virtual detuning method allows for any oscillation frequency without loss of contrast but is more sensitive to drive imperfections due to IQ imbalances. Fig. 3.14 shows the effect of small IQ imbalances (essentially different Rabi rates for X vs Y rotations) on the long-time signal: persistent small amplitude oscillations arising from a imperfectly calibrated X-Y π-pulses. The tomography section will discuss how to calibrate out these effects.

Finally, the spin-echo time, T_{echo}, is an NMR term which characterizes the low-frequency noise in the qubit. It consists of a sequence of three pulses: $\pi/2$, π, $\pi/2$ with equal wait times between both segments on either side of the π pulse. A different way to think about the dephasing process from the T_2 measurement is to consider an ensemble of qubits with a spread of frequencies centered about ω_{ge} whose behavior we average together. Each qubit will precess at a different rate relative to the drive leading to the averaged magnetization (which controls the contrast of the oscillations) to decay over time. However, by adding a π-pulse in the middle of the evolution time, we can rephase all the qubits again for the second measurement pulse. Consider two qubits with frequencies $\omega_{01} \pm \delta$: after an evolution time τ, the lower frequency qubit will seem to be "lagging" by $\delta\tau$ and the higher frequency qubit will be "leading" by δt. The π-pulse flips the state on the Bloch

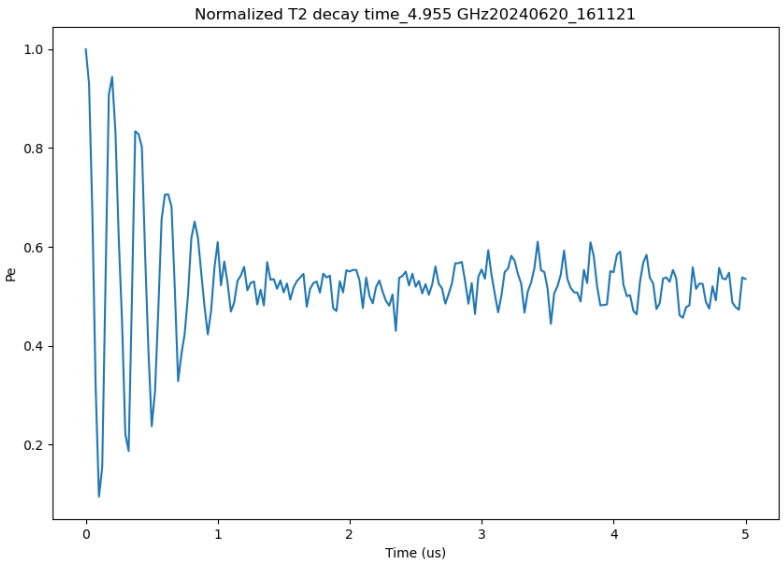

Figure 3.14: T_2 Ramsey measurement using a variable phase on the second pulse. At long times, small offsets on the IQ modulation ports of the RF generator lead to persistent oscillations long after the coherence has decayed away due to over/ under rotations as the second $\pi/2$-pulse phase is rotated.

sphere leading to the lower frequency qubit "leading" and the higher frequency qubit lagging. After another evolution time τ, the accumulated precession phases cancel out for both qubits leading to a full rephasing of the entire collection of qubits. In contrast to the T_2 measurement, all pulses are on resonance with the qubit as the π-pulse removes the oscillation in $P(e)$. Here the qubit state will start in $|g\rangle$ instead of $|e\rangle$ and grow to the 50-50 mixed state exponentially as shown in Fig. 3.13(b).

A single π-pulse can compensate the effect of a low frequency qubit frequency noise source (i.e. the noise is essentially constant from shot to shot) but will amplify noise with frequency $2/\tau$. This method can be used to perform noise spectroscopy by increasing the number of π-pulses between the two $\pi/2$ pulses [36]. This more general sequence is called the CPMG sequence and is characterized by the following spacing between π-pulses: $\tau_i = \frac{\tau(i-0.5)}{N}$ where N is the total number of pulses and i is the index of the ith π-pulse (starting at 1) [5, 43, 44].

3.2.7 Ramsey spectroscopy

Ramsey spectroscopy uses the pulse sequence defined in the T_2 discussion above to more precisely calibrate the qubit frequency. The oscillation frequency of $P(e)$ is set by the qubit-drive detuning allowing for a very accurate measurement of the qubit frequency for long-lived qubits. Fitting many periods of oscillation allows for a very accurate measurement of the detuning (recall the usual Fourier transform arguments: $\delta f \propto 1/T$) which can be used to extract the qubit frequency from the drive frequency. Two measurements are required to determine the absolute frequency of the qubit as $\delta = |\omega_q - \omega_d|$ doesn't give the sign of the detuning.

For fixed frequency qubits, once the qubit frequency has been calibrated to high precision, it can be useful to perform a secondary calibration routine to optimize the readout power and frequency and qubit π-pulse. For readout optimization, we perform a cavity transmission measurement with the qubit in $|g\rangle$ vs $|e\rangle$ and pick the point with the highest IQ contrast. While in principle this method also works for frequency tunable qubits, drifts in the flux bias point lead to marginal returns as the readout and qubit frequency need to be frequently calibrated.

3.3 Qubit pulse tuneup

A quantum algorithm might consist of thousands of gates applied consecutively before performing state readout requiring high-fidelity qubit pulses. However, at high pulse fidelities, the error on the final state after a single pulse can be difficult to measure accurately as the error is combined with the standard measurement noise in the readout. Error amplification sequences are used to increase the signal to noise ratio on the pulse infidelity. The typical gate set used for single qubit operations is π pulses and X, Y gates (defined as $\pi/2$ pulses about either axis). Errors in these pulses consist of over or under rotation, incorrect rotation axes due to small residual detuning (from imperfect qubit spectroscopy or the

Stark shift) or leakage to the $|f\rangle$ state, or IQ imbalances in the qubit generator leading to different Rabi rates for X vs Y pulses.

3.3.1 Pulse amplitude calibration

For fixed pulse duration, the amplitude of the drive pulse sets the rotation angle of the qubit. Over or under rotations due to miscalibrated pulses can quickly add up and the qubit in an effectively random state after N pulses. To fine tune the pulse amplitude, we perform an odd number of $\pi/2$ pulses and measure the qubit population. A $\pi/2$ pulse is used instead of a π-pulse as the qubit population is linearly sensitive to the rotation angle near the equator but only quadratically sensitive at the poles. A slightly over or under rotating pulse will have a characteristic growing oscillatory response as the accrued phase error grows with the number of pulses (see Fig. 3.18(c)). The optimal pulse amplitude corresponds to a flat response without any oscillations.

3.3.2 Stark Shift and DRAG Pulses

At large drive amplitudes, required for fast gates, the qubit no longer behaves strictly like a two-level system and the higher levels lead to errors in the computation basis. In particular, when the drive strength Ω approaches α, the $|f\rangle$-state will lead to a shift in the $g-e$ frequency, detuning the drive. We can estimate this effect by considering the driven 3-level Hamiltonian with a drive at the ω_{ge} frequency:

$$H = \begin{pmatrix} 0 & \frac{\Omega}{2} & 0 \\ \frac{\Omega}{2} & 0 & \frac{\sqrt{2}\Omega}{2} \\ 0 & \frac{\sqrt{2}\Omega}{2} & -\alpha \end{pmatrix} \tag{3.22}$$

where the first 2×2 block is the usual driven two level qubit and α is the anharmonicity. The off-diagonal elements between $|e\rangle$ and $|f\rangle$ have the $\sqrt{n}$ scaling expected from the

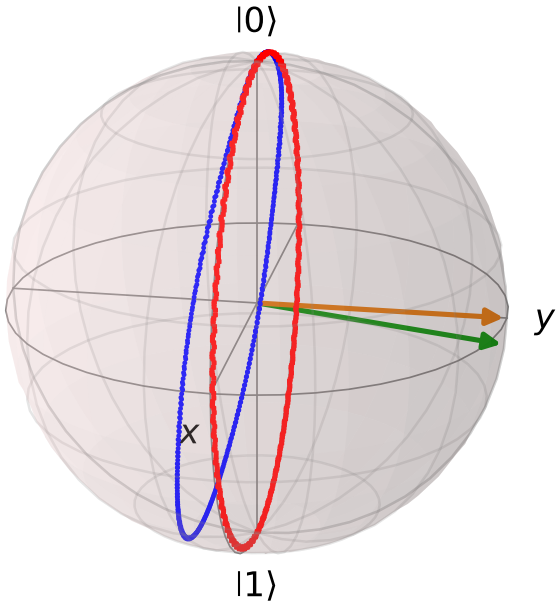

Figure 3.15: Bloch sphere visualization of the drive induced Stark shift on the qubit. The ideal qubit evolution is plotted in red corresponding to a pure Rabi drive on the y-axis (orange arrow). However, at large drive strengths relative to the qubit anharmonicity, the coupling to higher levels shifts the $g-e$ transition out of resonance with the drive leading to a "tilt" of the rotation axis (green arrow associated blue circle).

creation/annihilation operators in a harmonic oscillator. We can compute the shift on the $g-e$ transition by first diagonalizing the lower 2×2 block to compute the new eigenvalues due to the off-resonant drive on $e-f$:

$$E_{\pm} = \frac{\alpha}{2}\left(1 \pm \sqrt{1+\frac{2\Omega^2}{\alpha^2}}\right) \approx -\frac{\Omega^2}{2\alpha}, \alpha + \frac{\Omega^2}{2\alpha} \tag{3.23}$$

where we've expanded for $\frac{\Omega}{\alpha} < 1$. We can see that the $|e\rangle$-state energy gets shifted *down* by $\frac{\Omega^2}{2\alpha}$. This detuning leads to a tilt in the driving field as the drive is no longer perfectly resonant. This effect can be seen by plotting the $\sigma_{x,y}$ projections of a qubit driven with a nominally pure X drive (see Fig. 3.16): as the driving field amplitude increases, the rotation goes from purely in the x-z plane to a plane tilted relative to the z-axis. By shifting the driving frequency by the Stark shift ($\frac{\Omega^2}{2\alpha}$), we can compensate this error almost perfectly up to $\frac{\Omega}{\alpha} \approx 0.3 - 0.4$.

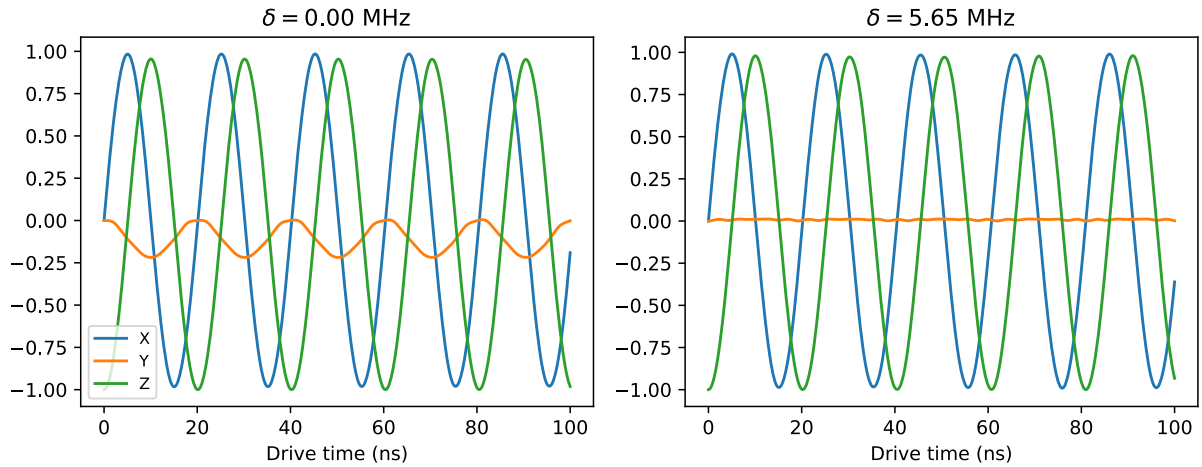

Figure 3.16: Stark shift induced from a $\Omega = 50$ MHz drive on a qubit with $\alpha = 220$ MHz anharmonicity. The blue,orange, and green traces show the qubit projections onto the X,Y,Z axes as a resonant microwave drive is applied on the qubit. As shown in the previous figure, in the ideal case, the qubit state should rotate in the XZ plane with no component in the Y direction, however, the drive induced stark shift slightly detunes the qubit g-e transition frequency leading to a tilt as shown in (a). This is corrected by shifting the drive frequency such that it is resonant with the new g-e transition frequency as shown in (b) where the oscillation in Y completely vanish. Simulation performed using QuTiP [21]

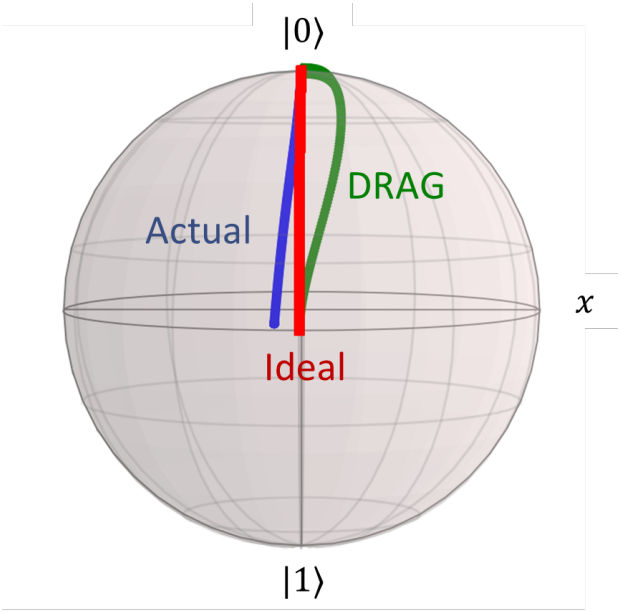

Figure 3.17: Trajectory of a qubit state on the Bloch sphere under a $\pi/2$ pulse. The Stark shift induced by the drive causes a small error shown with the blue trace. By applying a time dependent detuning, the DRAG pulse compensates the drive induced Stark shift visible as a tilt in the time evolution of the uncompensated pulse. Figure reproduced from [45]

The Derivative Removal by Adiabatic Gate protocol [46, 47] removes the stark shift and reduces the power at the $e - f$ frequency to reduce leakage and phase errors in the $g - e$ subspace. DRAG pulses apply the derivative of the pulse on the Q channel (for an X pulse):

$$\Omega_{DRAG}(t) = \Omega_x(t) + i\frac{\beta}{\alpha}\frac{d\Omega_x(t)}{dt} \tag{3.24}$$

where β is a scaling factor and $\Omega_x(t)$ is the qubit drive pulse. The derivative term will change the phase of the drive (effectively changing the frequency) which cancels out the first order effects from the Stark shift. Fig. 3.17 compares the trajectory of a qubit state under an unmodulated pulse and with a DRAG pulse showing a clear improvement in the final state fidelity.

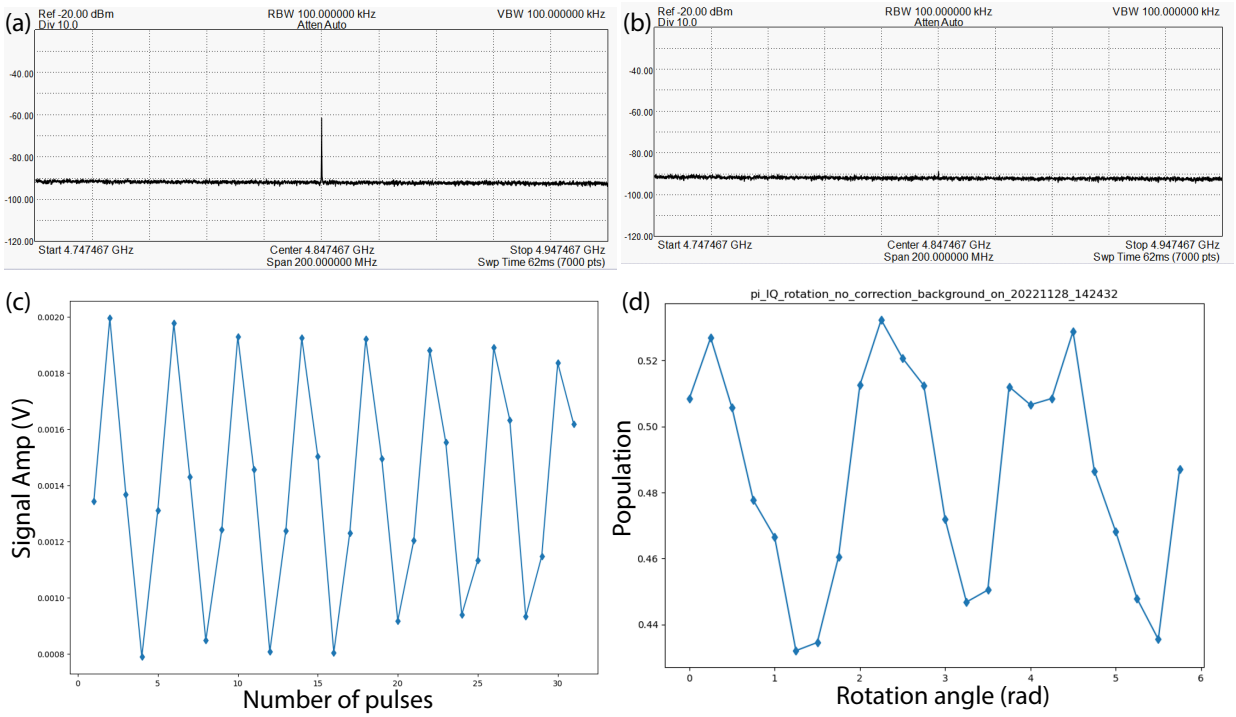

Figure 3.18: Characterization of different sources of qubit pulse errors. Top two panels show the effect of calibrating the IQ leakage terms on the qubit generator. Without the impairment, the qubit drive *leaks* through even for nominally 0 applied voltage on I and Q leading to spurious rotations. Bottom left panel, (c), shows the effect of slightly miscalibrated $\pi/2$ pulses where a sequence of pulses is applied to the qubit and the resulting oscillation indicates the sign of the error. In this case, reference $n\pi$ pulses are also acquired to provide a visual reference for the full scale contrast of the qubit readout amplitude. Bottom right, (d), shows qubit rotation errors as a function of drive angle. The generator IQ modulator introduces some small imbalances in the resulting drive amplitude which are mapped onto the resulting qubit population after a $\pi/2$ pulse.

3.3.3 IQ Impairments

One last source of error comes from imperfections in the IQ modulation of the qubit drive. In particular, offsets and amplitude imbalances lead to different Rabi rates for X vs Y and X vs -X pulses which all need to be calibrated out. Nominally, these effects can be modeled as follows:

$$\Omega(t) = (I(t) + I_0)\cos(\omega_q t) + (\gamma Q(t) + Q_0)\sin(\omega_q t) \tag{3.25}$$

where I_0, Q_0 are the IQ offsets and γ is the amplitude imbalance (i.e. the ratio of the output power in Q vs I).

The IQ offsets lead to additional rotation even for a zero amplitude pulse (recall, the qubit rotation angle is set by the *integral* of the qubit drive so a small amplitude, always on, drive can lead to a significant rotation error) as well as an asymmetry for +X and -X pulses. These offsets can be calibrated out by using a spectrum analyzer, setting I and Q to zero (so that there is no nominal power), and adjusting the IQ offsets on the generator to minimize the power at the carrier frequency (as shown in Fig. 3.18(a) and (b) showing carrier leakage before and after IQ offset compensation). By using digital pulse modulation on the generator in addition to IQ modulation, we can completely remove the leakage outside of the pulse envelope as the generator output is fully shut off.

Amplitude imbalance, i.e. different Rabi rates for X and Y pulses, can be characterized by rotating the azimuthal phase of a $\pi/2$ pulse and mapping the excited state fraction as shown in Fig. 3.18(d). We can then compensate the pulse amplitude as a function of azimuthal angle to maintain a perfect $\pi/2$ pulse regardless of the rotation axis. Note, to avoid combining errors, the IQ offsets should be canceled out before performing amplitude balancing. Residual offsets will appear in a non-symmetric excited state population for X vs -X and Y vs -Y, effectively tilting the rotation axis.

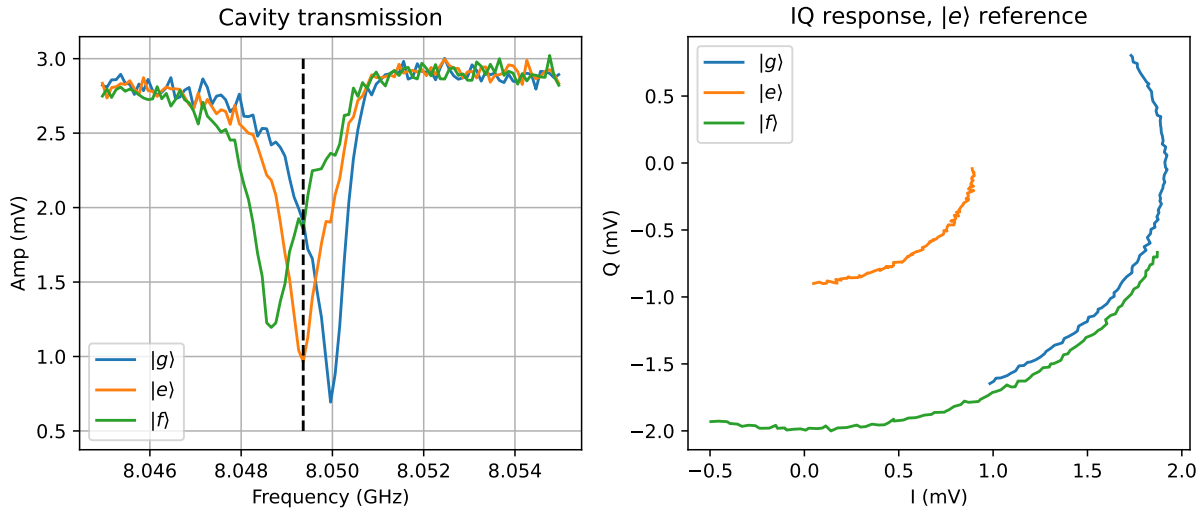

Figure 3.19: $|f\rangle$-state spectroscopy. Left: cavity transmission spectra for the qubit initialized in $|g\rangle, |e\rangle, |f\rangle$ showing three different resonance frequencies. By placing the cavity monitor tone (dashed line) on the $|e\rangle$-state transition, all three states have distinguishable IQ points shown in the right panel. Right, IQ response for the qubit in different states. The circular smearing of the IQ points arises from a slow phase drift between the cavity generator and the local oscillator but does not prevent population extraction (as long as the reference points are acquired *with* the data traces and not collected all at once at the beginning or end of the full measurement).

3.4 $|f\rangle$ state readout

While we usually only focus on and work with the lowest two levels of the transmon, reading out the $|f\rangle$-state is sometimes necessary. This can either be to increase the Hilbert space directly, going from a qubit to a qutrit, measuring leakage out of the g,e subspace, or even model some non-Hermitian dynamics (by post-selecting evolutions where the qubit stayed in the $e - f$ manifold [48]). The $|f\rangle$ state induces a dispersive shift on the cavity, just like the $|g\rangle$ and $|e\rangle$ states, which we can use for readout. Fig. 3.19 shows the cavity response for states prepared in $|g\rangle$ vs $|e\rangle$ vs $|f\rangle$ and we can immediately see that using the $|g\rangle$-state cavity frequency for our probe is only going to resolve g vs not-g instead of the states individually. By parking the readout frequency near the $|e\rangle$ state we can cleanly separate the response for all three states. Note, even though the cavity amplitude response is roughly the same for $|g\rangle$ and $|f\rangle$ at this probe frequency, they have opposite phases as

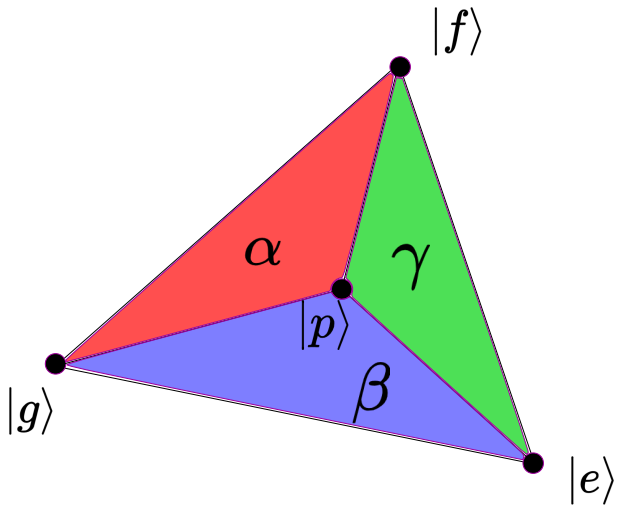

Figure 3.20: Probability simplex of the qubit population. Any point inside the triangle is a valid combination of g, e, f population fractions. The analysis in the text provides a way to map the point $|p\rangle$ into a weighted sum of g,e, and f. α, β, γ indicate the $|e\rangle, |f\rangle, |g\rangle$ populations.

they are on different sides of the resonance.

The process of tuning up a π-pulse for the $e - f$ transition is very similar to the $g - e$ π-pulse except that we prepend all the pulses with a π-pulse to excited the qubit to $|e\rangle$ before performing spectroscopy and Rabi oscillations. This can be done with a single generator using frequency modulation on the qubit drive to vary the microwave frequency between the $g - e$ and $e - f$ transition frequencies or with two generators with basic pulse modulation.

By acquiring reference IQ points for $|g, e, f\rangle$ we can map the output signal to a population distribution using barycentric coordinates on the probability simplex defined by $|g, e, f\rangle$ (i.e. the combination of P_g, P_e, P_f that gives the measured data point inside the triangle defined by all the possible combinations satisfying $P_g + P_e + P_f = 1$). Consider the point IQ_p inside the triangle defined by points IQ_g, IQ_e, IQ_f (see Fig. 3.20), the population distribution satisfies: $P_g * IQ_g + P_e * IQ_e + P_f * IQ_f = IQ_p$ with $P_g + P_e + P_f = 1$ to preserve normalization. We can express the vector $\vec{\mathbf{GP}}$ as a linear combination:

$$\vec{\mathbf{GP}} = \alpha\vec{\mathbf{GE}} + \beta\vec{\mathbf{GF}} \tag{3.26}$$

where α, β are the $|e\rangle$ and $|f\rangle$ state populations and $\vec{\mathbf{ab}}$ indicates the vector in IQ space from point a to b. This equation can be simply inverted to extract α, β using Cramer's rule for linear equations (first expression):

$$\alpha = \frac{\mathrm{GF}_y\mathrm{GP}_x - \mathrm{GF}_x\mathrm{GP}_y}{\mathrm{GE}_x\mathrm{GF}_y - \mathrm{GE}_y\mathrm{GF}_x} = \frac{(\vec{\mathbf{GF}} \times \vec{\mathbf{GP}}) \cdot \hat{z}}{(\vec{\mathbf{GE}} \times \vec{\mathbf{GF}}) \cdot \hat{z}} = \frac{Area(GFP)}{Area(GEF)} \tag{3.27}$$

$$\beta = \frac{\mathrm{Ge}_y\mathrm{GP}_x - \mathrm{Ge}_x\mathrm{GP}_y}{\mathrm{GE}_x\mathrm{GF}_y - \mathrm{GE}_y\mathrm{GF}_x} = \frac{(\vec{\mathbf{GE}} \times \vec{\mathbf{GP}}) \cdot \hat{z}}{(\vec{\mathbf{GE}} \times \vec{\mathbf{GF}}) \cdot \hat{z}} = \frac{Area(GEP)}{Area(GEF)} \tag{3.28}$$

where we have re-expressed α, β as the ratio of triangles to highlight a more geometric interpretation of the problem. Finally, we can express P_g, P_e, P_f as $(1 - \alpha - \beta), \alpha, \beta$ to force normalization of the qubit population. Note that points near the edges of the simplex can lead to small but negative values for the population in the state living on the opposite vertex which must be truncated to 0.

3.5 State tomography and arbitrary state preparation

To fully characterize the state of the quantum state of a qubit, we need to perform multiple measurements in different bases. The typical representation of a quantum state is a vector pointing on the Bloch sphere with the south pole being the the $|g\rangle$ state and the north pole being the $|e\rangle$ state (note, this is exactly the opposite convention typically used by theorists but it serves as a useful intuitive resource). An arbitrary, pure state lives on the surface of the Bloch sphere with quantum gates consisting of rotations about the XYZ axes (XY through microwave drives, Z by detuning the qubit or the drive from each other). By specifying the X, Y, Z projections of the vector we can reconstruct the state uniquely. The density matrix of the qubit can be written as:

$$\rho = \frac{1}{2}\left(I + \vec{\mathbf{r}} \cdot \vec{\sigma}\right) \tag{3.29}$$

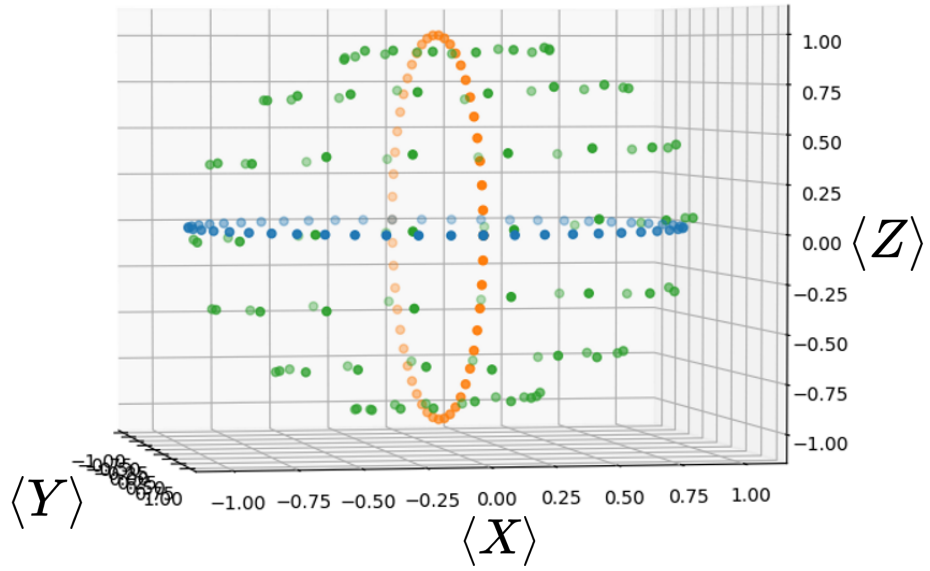

Figure 3.21: Arbitrary state preparation on the Bloch sphere. The orange and blue circles indicate the great circles in the XY and XZ planes for visual reference. Each green point corresponds to a state prepared to a specified polar and azimuthal angle. The small residual tilt along the x-axis indicates an potential offset error in the IQ calibration (as $+X$ and $-X$ have opposite errors).

where I is the identity matrix, $\vec{\mathbf{r}}$ is the Bloch vector, and $\vec{\sigma}$ is the Pauli vector. By computing $\langle \sigma_i \rangle$ (where i loops over the X,Y,Z bases) we are effectively computing the XYZ coordinates of the qubit. The natural basis for state measurements is σ_z from the dispersive qubit-readout coupling. The X and Y projections are measured by applying a $\pi/2$ rotation right before readout to rotate the X, Y axes onto the Z axis. In this representation, pure states will have $\|\vec{r}\| = 1$ whereas mixed states have $\|\vec{r}\| < 1$. To obtain physical density matrices the Bloch vector can be normalized if it's length is larger than one due to noise (or poor calibration). More sophisticated methods like Maximum Likelihood Estimation or Bayesian analysis can be used to obtain more accurate representations of the density matrix however we did not require them for our analysis.

In addition to being able to perform state tomography, we might want to be able to initialize out qubit anywhere on the surface of the Bloch sphere. The state $|g\rangle$ can be rotated to any point on the Bloch sphere by specifying the polar and azimuthal angles

required to reach it. The polar angle is set by the area of the qubit pulse (usually by changing the amplitude of the pulse) whereas the azimuthal angle is set by the qubit drive phase. Fig. 3.21 shows a set of state initializations spanning the surface of the Bloch sphere followed by state tomography to verify the state. When performing these calibrations, we are sensitive to so-called SPAM errors: State Preparation And Measurement. Over/under rotations due to slight impairments in the microwave drives lead to both incorrect states as well as errors in the $\sigma_{x,y}$ projections as those rely on calibrated $\pi/2$ pulses. Normalization errors due to imperfect reference $|g\rangle$ and $|e\rangle$ state measurements propagate to all projection measurements leading to unphysical density matrices (Tr > 1 or negative diagonal elements).

Chapter 4

Experimental Apparatus

In this chapter, we will go over the experimental apparatus used for all the measurements and experiments described in the following three chapters. At the core of our measurement set up is the dilution fridge, cooling the qubits to 10 mK and isolating them from the environment. Both the input and output lines to/from the qubit need to be carefully set up to avoid bringing in unwanted noise and destroying the coherence of the devices. We will briefly discuss the physical wiring of the fridge as well as the guiding principles behind the choices we made. The qubit packaging is another critical component for high coherence devices which we will discuss. We can then switch gears and review the microwave equipment used to coherently control the qubit and perform measurements. Finally, we will discuss the control code architecture and the general experimental script template used for measurements.

4.1 The dilution fridge

Visit any superconducting lab and the first object you will notice is the loud cylinder sitting in the corner with cables coming out of it. This is the dilution fridge in which qubits are cooled to milli-Kelvin temperatures. The fridge serves two fundamental purposes: physically cooling down the qubits to make them superconducting as well as reducing

thermal photons at the resonant frequency of the qubit.

As we discussed in Chapter 2, our qubits consist of resonant modes in an electrical circuit. These circuits are fabricated using a superconducting metal and a high quality substrate to reduce resistive losses in the material and increase the qubits' lifetime. For all the qubits we fabricate in the lab, we use tantalum (Ta) which has a critical temperature of 4.4 K. Below this temperature, electrons in the metal combine into Cooper pairs leading to a dissipationless DC supercurrent. However, not all electrons pair up leaving so-called quasi-particles (unpaired electrons) which behave dissipatively [34]. The thermal equilibrium quasi-particle density in a superconductor at temperature $T < \Delta/k_B$, is exponentially dependent on the temperature [49]:

$$n_{qp}^{eq}(T) \propto \sqrt{\Delta T} e^{\frac{-\Delta}{k_B T}} \tag{4.1}$$

where Δ is the superconducting gap of the metal and T is the temperature of the sample. Operating a superconductor far below the critical temperature exponentially suppresses quasi-particles, removing a dissipation channel from the qubit. In addition to quasi-particles, spurious two level systems (produced by small defects at the interfaces between the metal layer and substrate or impurities in the substrate itself) can also lead to reductions in the coherence time of the qubit due to coupling to a random array of spins. The dynamics of these spins can be frozen out by cooling the environment to a temperature much colder than the energy splitting. The exact interplay between these two forms of qubit decoherence losses is still an active field of research which the interested reader can learn more about in [49, 50].

In addition to materials considerations, cooling down the qubit environment also changes the spectrum of the blackbody radiation the qubit sees. Any body at temperature T will emit a spectrum of radiation given by Planck's law (giving the spectral energy

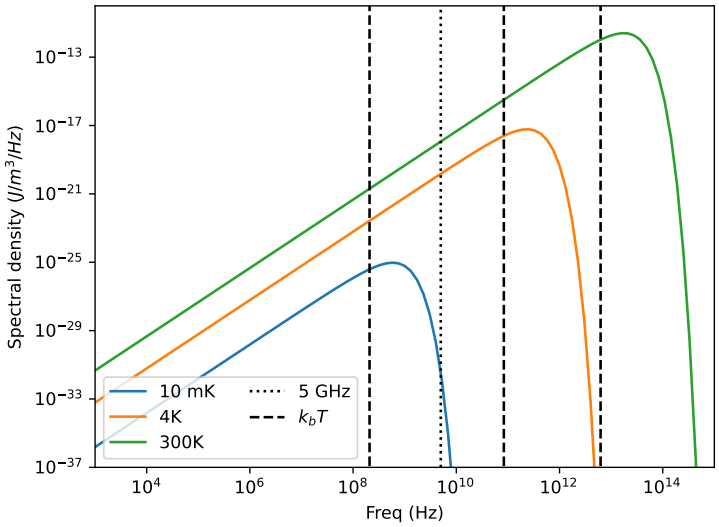

Figure 4.1: Blackbody radiation spectral intensity at a few different temperatures using Planck's Law. We can clearly see that as the temperature increases, the overall radiated power increases as well as the peak frequency. k_bT/h is shown in dashed lines for each temperature clearly approximating the turnover point in the emission spectrum. The dotted line shows the typical qubit operating frequency, 5 GHz, indicating an exponential suppression of blackbody radiation at the base temperature of the fridge.

density):

$$u_v(\nu, T) = \frac{8\pi h\nu^3}{c^2} \frac{1}{e^{\frac{h\nu}{k_B T}} - 1} \tag{4.2}$$

where T is the temperature of the blackbody, ν is the frequency of emitted light, c is the speed of light and h, k_B are Planck's constant and the Boltzmann constant respectively. k_BT sets the energy scale for the system exponentially suppressing emission of photons with energy significantly above this value. The peak emission frequency (i.e. the maximum of u_v) is given by Wien's displacement law (by setting the derivative of u_b equal to zero):

$$\nu_p = bT \tag{4.3}$$

where $b = 0.058789257\ \text{THz} \cdot \text{K}^{-1}$. We can build some intuition by calculating the peak frequency for a few different reference temperatures: room temperature (300 K, 17 THz), liquid Helium (4 K, 235 GHz), base temperature in a dilution fridge (10 mK, 600 MHz).

Figure 4.1 shows the spectral density for these reference temperatures highlighting the density in the 4 to 8 GHz range.

4.1.1 Principle of operation

We've just discussed how a cold environment is critical for the qubit, but how does the dilution fridge achieve this? The most common cryogen is liquid helium with a boiling point around 4 K which is still orders of magnitude higher than the required 10 mK. There are three main ways cryogenic equipment operates to cool down a load: cyclic expansion of a working fluid (pulse tubes), evaporation (1K pots), and endothermic phase separation (dilution units). The workhorse of the dilution unit is the pulse tube which uses periodic compression and expansion of a working gas (in our case helium) to remove heat from the load. Fundamentally, it relies on spatially separating out the expansion part of the gas (where heat is absorbed from the load) and the compression part of the gas (where heat is dumped to the environment) using pulses of gas that periodically travel between the hot and cold ends of the system. The pulse tube is the source of the periodic pulsing sound heard in any lab and provides the majority of the cooling power to the fridge, bringing it down to 4 K (where the helium would start condensing).

While modern designs achieve temperatures reaching 3 K, a different method is required to achieve sub Kelvin temperatures. A straightforward method to achieve this is evaporative cooling where heat is extracted from a load by evaporating a working fluid off of it (which is exactly how our bodies cool themselves with sweat). For cryogenic systems, a small vessel (usually called a 1K pot) is filled with Helium and connected to a vacuum pump which reduces the pressure above the fluid surface. This leads to a decrease in the vapor pressure of the fluid increasing the evaporation rate. Helium molecules with higher energy than the average energy will preferentially evaporate off leading to a net cooling of the fluid over time. In principle, the final temperature is set by pump pressure and the thermal engineering of the feedline connected to the pot to keep it filled. However, the

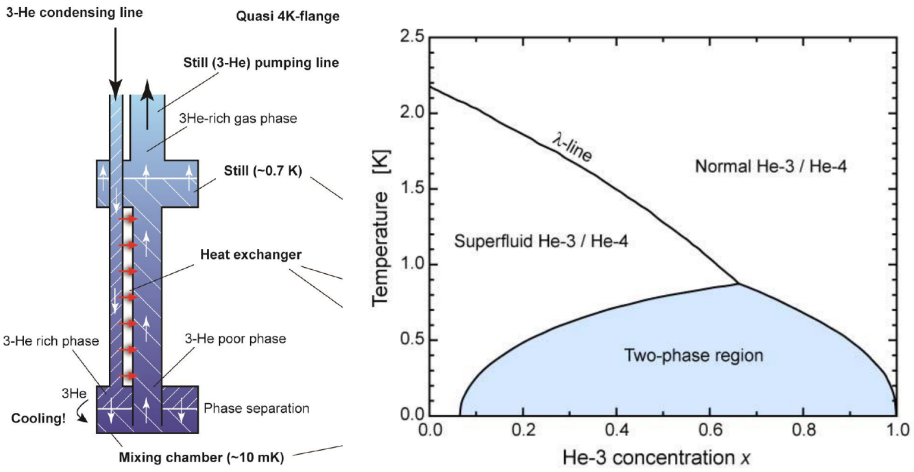

Figure 4.2: Schematic representation of the principle of dilution cooling. Left, shows the He-3 circulation loop highlighting the main components: phase separation in the mixing chamber leading to cooling power, preferential pumping on the dilute He-3 phase up to the still with heat exchange with the incoming He-3, evaporation of He-3 from dilute phase to close circulation loop. Right, phase diagram of a He-3/He-4 mixture as a function of temperature and He-3 concentration. In the two phase region, the mixture separates into a He-3 rich and a He-3 poor phase by absorbing energy from the environment. Figure adapted from Bluefors LD400 users manual.

typical working fluid is He-4 which becomes a superfluid around 2.2 K which creeps up the tube until it reaches a temperature where it becomes normal again and evaporates there. This limits the 1K pot to temperatures around 1 K (hence the name) requiring yet another trick to achieve mK temperatures.

The key operative principle of dilution fridges is the energy absorbed by a He-3/He-4 mixture. Mixing enthalpy is a thermodynamic term indicating the energy required or release when two substances are combined. An endothermic mixing enthalpy ($\Delta E < 0$) leads to the required energy being removed from the environment, cooling it by ΔE. In particular, a He-3/He-4 mixture below 1 K will separate into two immiscible phases: an almost pure He-3 phase and a He-3 Fermi liquid in a He-4 superfluid (dilute He-3 phase). The dilute phase is pumped to the *still* which is heated to 700-800 mK, boiling off pure He-3 leading to more He-3 diluting from the He-3 rich to the He-3 poor phase in the *mixing chamber*. The boiled off He-3 gas is then recirculated and condensed back into the mixing chamber completing the cycle as shown in Fig. 4.2. There is no fundamental limit

to the lower temperature achievable through dilution; however, the cryogenic engineering becomes very challenging for temperatures below 10 mK. For example, the heat conduction between He-3 and the heat exchanger used to cool down the gas before it enters the mixing chamber scales as T^4 (the Kapitza resistance) meaning that a 10x reduction in the operating temperature is accompanied by a 10,000 increase in the surface area required for the same heat conduction.

Our dilution fridge is a Bluefors LD400 (see Fig. 4.3) which uses a pulse tube for the main cooling power and a dilution unit to get the experimental sample space down to 10 mK. The fridge consists of a few different "plates" which are kept at different temperatures. Each plate is associated with a full thermal shield leading to a Russian nesting doll structure with the innermost shield at the base temperature. The typical break point temperatures are 50 K, 4 K, 800 mK (still), 10 mK (mixing chamber) with an additional plate between the 800 mK and 10 mK plate called the quasi-cold plate. These thermal shields help reduce the radiated power from the higher temperature shell to the colder stage, decreasing the additional thermal load on that stage allowing it to operate more efficiently. In particular, the 50 K stage breaks up the temperature drop from room temperature into 300 K to 50 K and then 50 K to 4 K drastically reducing the radiated power between the shields (recall, the radiated power goes as $T_H^4 - T_C^4$ where T_H, T_C are the temperatures of the hot and cold side). The still and mixing chamber are part of a closed He-3 circulation loop which is thermally coupled to the main pulse tube cryogenic unit. A typical cool down will start with running the pulse tube for 10-15 hours bringing the full system down to approximately 4 K. The He-3/He-4 mix is then forced into the fridge at high pressure, over 30 minutes, through a flow impedance leading to condensation in the mixing chamber. Once all the mix is condensed, the still is heated up to drive the He-3 circulation loop driving the system all the way down to 10 mK in a few hours.

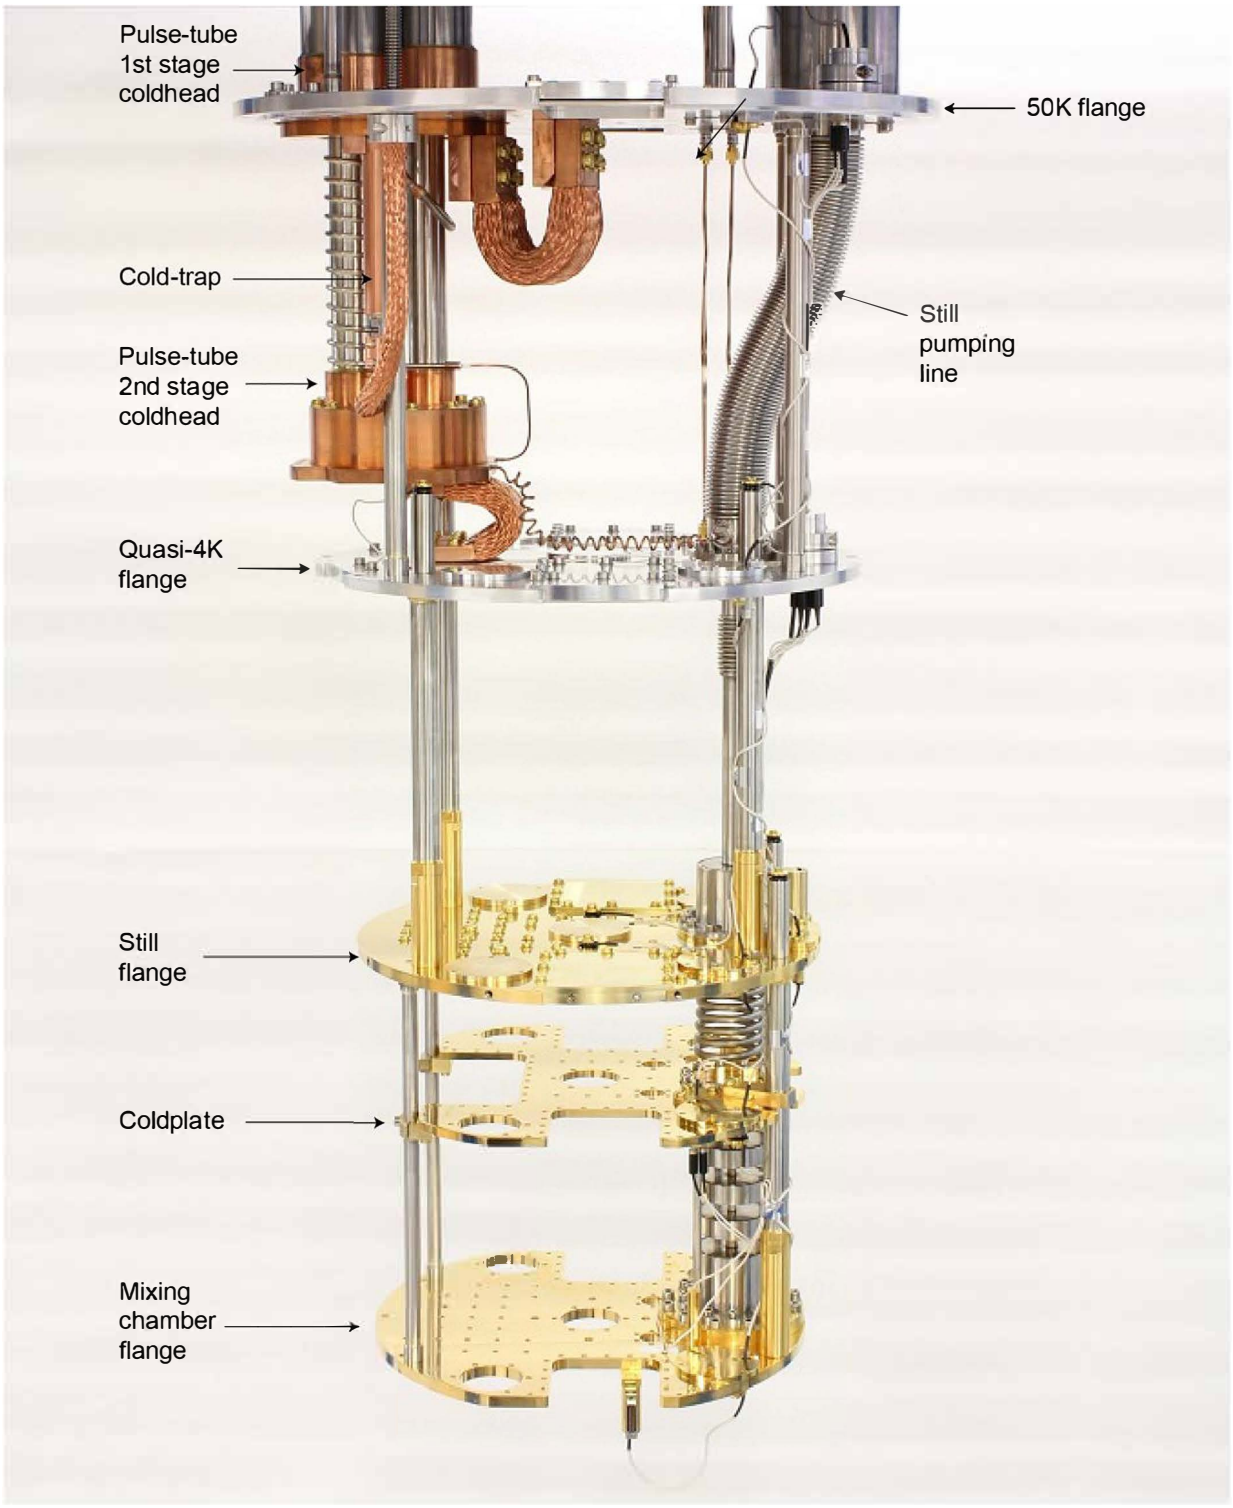

Figure 4.3: Picture of the full fridge "chandelier". The lower temperature flanges are gold plated to reduce radiative heating from the higher flanges. Picture extracted from Bluefors users manual.

4.1.2 Thermal noise mitigation

While ideally all the stages in the dilution fridge would be completely isolated, both qubit control and readout require direct electrical connections from the room temperature equipment to the base plate in the fridge. This presents a challenge in two ways: firstly, cables thermally couple consecutive stages leading to heating from the upper stage if proper care isn't taken; secondly, electrical noise from the microwave generators can lead to additional microwave photons at the base plate.

Electronics at non-zero temperature generate noise from the thermal motion of the charge carriers called Johnson-Nyquist noise:

$$\overline{V_n^2} = 4k_B T R \Delta f \tag{4.4}$$

where V_n is the noise voltage, k_B is the Boltzmann constant, T, R are the temperature and resistance of the object respectively, and Δf is the bandwidth over which the signal is acquired. This noise is extremely broadband and uniform up to the THz frequency range. It also sets the noise floor for any power measurement in a given bandwidth: for example, a 1 Hz bandwidth measurement has a noise floor of -174 dBm indicating that signals approaching this power become indistinguishable from thermal noise in the system. This noise source is the key reason why sensitive measurement equipment is operated either at cryogenic temperatures (for example with LN2 cooled electronics) or with a narrow bandwidth (effectively all telecommunications applications use narrow band filters to suppress the noise) or both for particularly sensitive applications.

By placing attenuators in the fridge, we can suppress the thermal noise from the room temperature electronics and compensate the decreased signal amplitude by increasing the input power. In addition to attenuators, we use lossy stainless steel cables to connect the different temperature stages together to increase the attenuation and decrease the thermal coupling between consecutive stages. Ideally, all the attenuation would be placed on

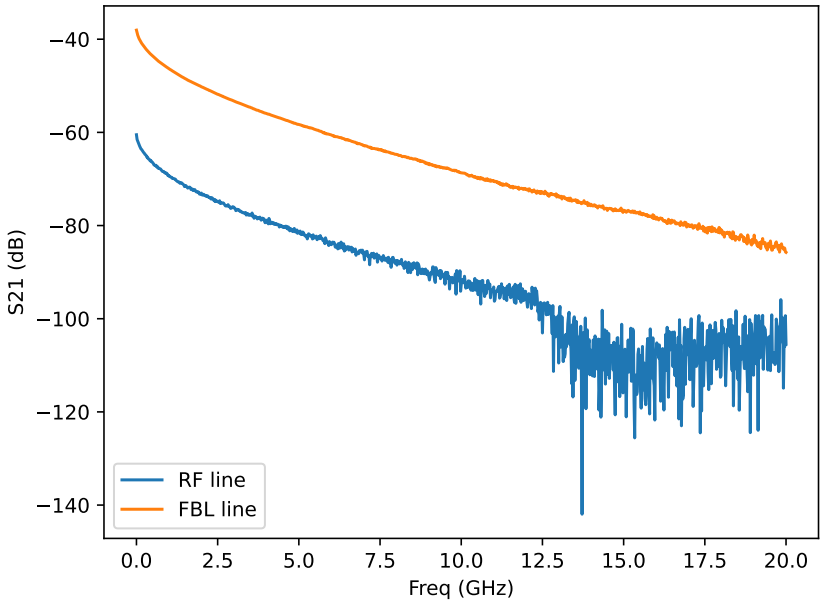

Figure 4.4: Line attenuation acquired at room temperature for an RF and a FBL type line. The down turn in the RF data around 12 GHz arises from the K&L filter used to filter out IR radiation. At large attenuation, the incoming signal approaches the VNA noise floor which can be seen by the noisier trace at high frequencies.

the mixing plate as that would introduce the minimal amount of thermal noise from the environment (a simple model for an attenuator is a beam splitter where one of the port inputs in the environmental vacuum). However, the dissipated heat from the attenuators would overwhelm the cooling power at the base plate. Instead, we spread out the attenuation between all the plates to reduce the final heat load on the mixing chamber while (a) maintaining the total required attenuation and (b) utilizing the higher cooling power of the upper plates to dissipate most of the heat [51]. Our particular configuration uses 61 dB of attenuation split up as follows: 1 dB at 50 K, 20 dB at 4 K, 10 dB at the still, 10 dB at the quasi-cold plate, and a final 20 dB at the mixing chamber. The additional attenuation from the lossy steel coax cables leads to a total attenuation of around 100 dB in the 4-8 GHz range. An example attenuation curve acquired at room temperature is shown in Fig. 4.4.

In addition to signal attenuation, we also use filters to remove high-frequency photons (15+GHz into the infrared). These photons will not directly couple to the qubit but can

cause heating effects on the chip or even break Cooper pairs leading to an avalanche of quasi-particles. To suppress these photons, we use K&L filters with a cutoff frequency at 12 GHz. Particular care has to be taken when considering filters to cut out high-frequency signals as the packaging and design of the filter can limit the upper frequency range it will provide attenuation for. For example, a 100 MHz low pass filter from Mini-Circuits (SLP-100+) is only rated up to 400 MHz on the datasheet and will not be guaranteed to suppress signals at 5 GHz.

4.1.3 Amplifiers and Circulators

On the output side we use amplifiers to increase the signal amplitude enough to make it measurable at room temperature. We use high electron mobility transistors (HEMTs) placed at the 4 K stage to amplify the signal by 40 dB before it reaches room temperature. A high-quality amplifier is *crucial* to preserve the initial signal to noise ratio of the signal coming from the chip. A basic model for an amplifier is:

$$V_{out} = GV_{in} + N \tag{4.5}$$

where $V_{in,out}$ are the input and output voltages, G is the gain (in linear units), and N is the noise added by the amplifier. Chaining multiple amplifiers together leads to the Friis equation:

$$V_{out} = \prod_i G_i V_{in} + \sum_i \prod_{j=0}^{N-i} G_i N_i \tag{4.6}$$

where G_i and N_i are the gain and added noise of each amplification stage. We can see that the added noise from the very first amplifier gets multiplied by all the subsequent gains making it the dominant source of noise at the end of the chain. We can compute the signal to noise ratio of the full amplifier chain (composed of three amplifiers in this example):

$$SNR_o = \frac{G_1 G_2 G_3 V_{in}}{G_1 G_2 G_3 N_i + G_2 G_3 N_1 + G_3 N_2 + N_3} \tag{4.7}$$

where N_i is the original noise amplitude in the signal. The noise factor F is defined as the ratio of the input SNR to the output SNR and is always greater than 1. An amplifier with high gain and low added noise will have a noise factor close to one whereas a poor amplifier will have a higher noise factor. We can rearrange the expression above to express the noise factor of an amplifier chain as a function of the individual noise factors F_i:

$$F_T = \frac{SNR_i}{SNR_o} = 1 + \frac{N_1}{G_1 N_i} + \frac{N_2}{N_i G_1 G_2} + \frac{N_3}{N_i G_1 G_2} \tag{4.8}$$

$$= F_1 + \frac{F_2 - 1}{G_1} + \frac{F_3 - 1}{G_1 G2} \tag{4.9}$$

where the effect of the subsequent amplifiers is suppressed by the accumulated gain before them. This model can be extended to model attenuation/ general loss by simply using G values less than one. To preserve as much of the input SNR as possible, the loss before the first amplifier must be minimized and the first amplification stage must have the lowest noise factor. We achieve this by using superconducting coaxial cables between the qubit and the HEMTs at the 4 K stage, reducing transmission losses while operating the HEMTs at cryogenic temperatures to minimize the added thermal noise on the input of the amplifier.

The final components at the base plate are the isolators and circulators. While we discussed an attenuation based scheme to shield the qubit from room temperature radiation on the input side, such a scheme will not work on the output as we cannot attenuate our output signal (as discussed in the previous section, attenuation before the first amplifier drastically reduces the overall achievable SNR). This is particularly important between the base plate and the 4 K stage as the amplifiers can send 4 K noise down to the base plate. To address this issue, we use non-reciprocal devices based on time reversal symmetry breaking induced by a magnetic field. A circulator is a 3-port device arranged in a ring shape around a large ferrite bead. The magnetic field from the bead breaks the symmetry between clockwise and counter-clockwise propagation modes leading to a preferred signal propagation direction. The ratio between the clockwise and counter-clockwise propagation

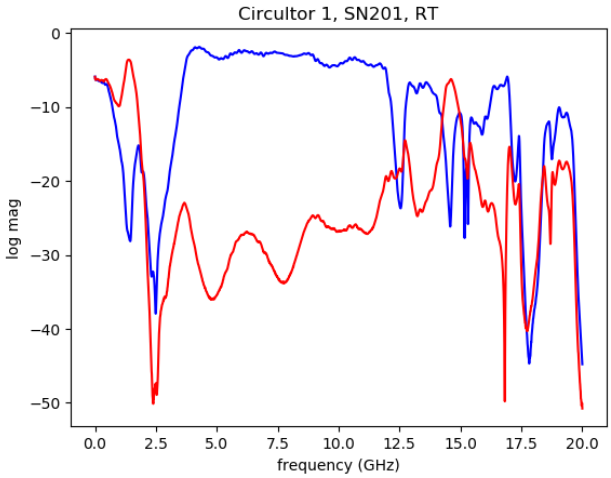

Figure 4.5: Forward (blue) and reverse (red) signal gain through an isolator acquired at room temperature on a Raditek isolator. In the operating bandwidth (4-12 GHz), there is a consistent 20+ dB isolation.

amplitude is called the *isolation* (typically expressed in dB) and can be thought of as a one way attenuation between the ports. The isolator is a two port device made by terminating one of the outputs of a circulator.

4.1.4 Wiring diagram

The current wiring configuration for the Bluefors dilution fridge used in all the experiments presented in this thesis is in Fig. 4.6. In addition to the regular RF input (8) and output lines (4) used to drive the qubit and cavity as well as perform readout, we also have flux bias lines (FBL, 8 total) used to tune components in the devices. These FBLs are configured with heavy filtering to mitigate flux noise on the qubit but do not require as much attenuation as the operating bandwidth is much lower than the qubit transition frequency (DC-5 MHz). These lines are split into "aggressive" (4) and "mild" (4) indicating the total attenuation on the line (38 and 32 dB of attenuation respectively).

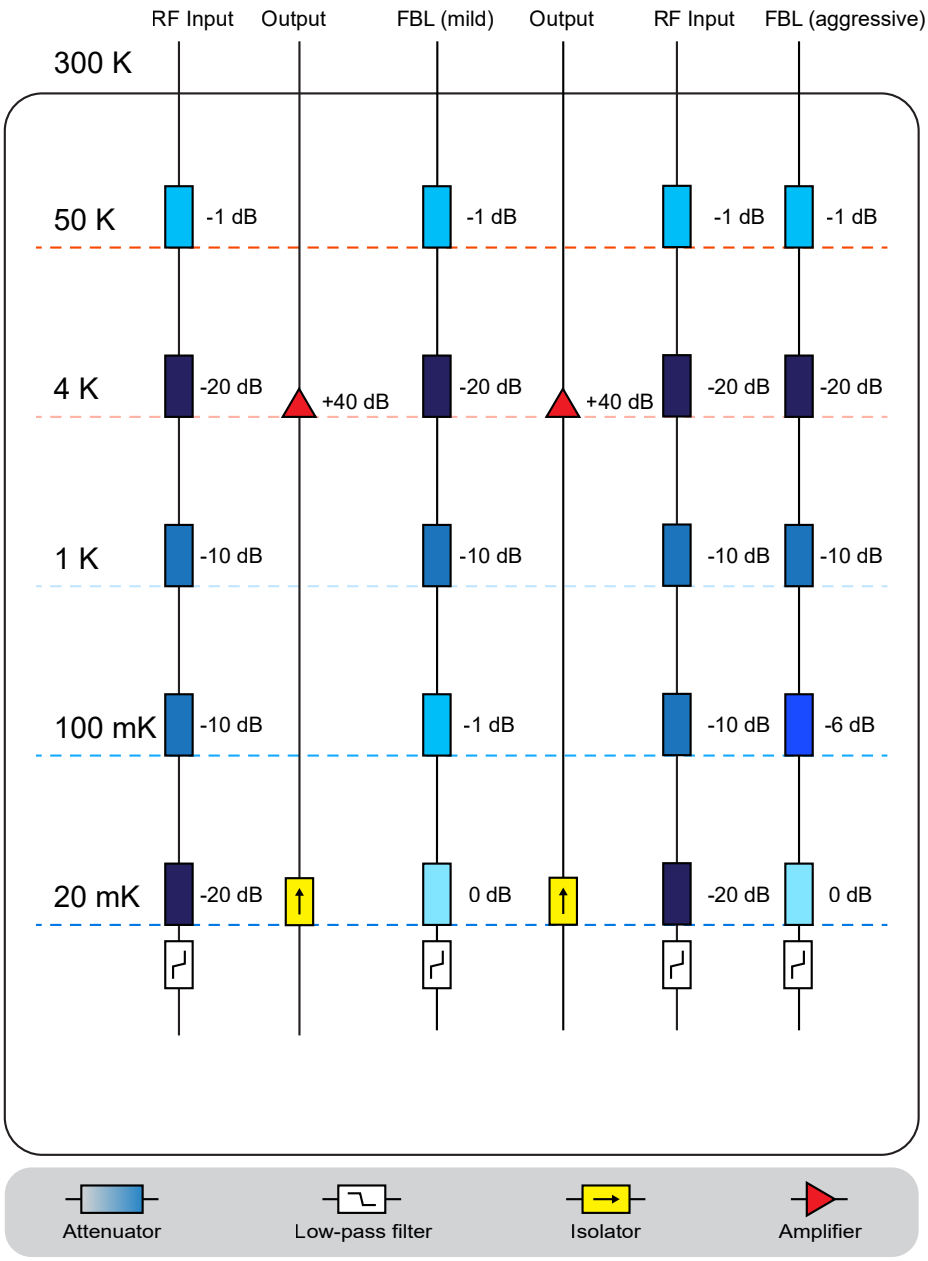

Figure 4.6: Fridge wiring configuration for a two input, two output device with two tunable features using FBLs. On the RF lines, the low pass filters are K&L 12.5 GHz (6L250-00089) tubular filters, the flux bias lines use a combination of VLFX80+ and SLP5+ for a 5 MHz bandwidth. The amplifiers are Amplitech (APT5-04000800-0710-D6). The isolators are Quinstar (QCI-G0401202AS) isolators with an operating frequency range of 4 to 12 GHz.

4.2 Qubit packaging

The qubit packaging serves three main roles: shielding the device from stray magnetic fields, thermalizing it to the base plate, and interfacing it with the connectors used to send control signals to it. Poor packaging typically results in decreased coherence times and parasitic microwave modes.

The innermost layer of packaging is the wirebonding of the device to a breakout printed circuit board (PCB) with soldered SMA connectors. This is the same method used in the fabrication of integrated circuits (ICs) where a silicon die containing all the electrical components is packaged into a plastic shell with pins used to interface with the device. Three different kinds of wirebonds are used for the device: launchpad connectors, ground plane wirebonds, and on-chip wirebonds. The launchpad wirebonds connect the control lines to the PCB center pin trace and typically consist of three to four wirebonds to reduce the electrical resistance. The ground plane wirebonds serve to connect the device ground plane to the PCB ground with a very low impedance to reduce ground loops and other small voltage offsets between the device and the rest of the electrical control signals as these can introduce noise. Finally, the on-chip wirebonds serve to unify the ground plane on the chip itself by jumping across the CPW resonators defined on the chip. This prevents parasitic chip modes between the different ground plane islands as well as suppresses the slot-line mode in the CPW resonators themselves. Fig. 4.7 shows a packaged transmon device with all the required wirebonds.

The encapsulation of our device uses two copper parts (referred to as *puck* and *lid* respectively) for thermalization and mechanical mounting inside the fridge. To guarantee good thermal contact between the qubit and the rest of the system, we use the crash-pack method in which the qubit is sandwiched between the PCB and the puck with screws. By slightly sanding down the puck so that the qubit chip is no longer perfectly flush with the surface of the puck but slightly sticks out, we can generate a large contact force between the puck and the qubit after tightening the screws resulting is a good thermal contact. However,

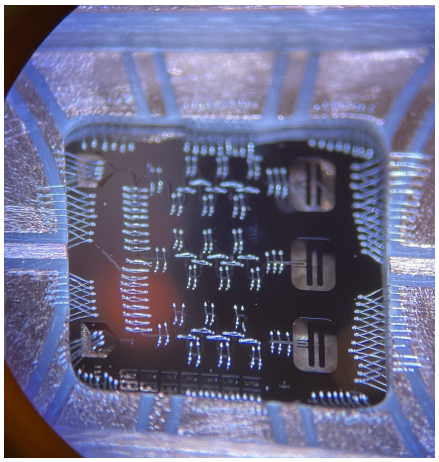

Figure 4.7: Picture of a fully wirebonded device, photo credit: Maya Amouzegar

it is important to not sand down the puck too much to avoid mechanically distorting the edge of the PCB during tightening. In addition to thermalization, the puck and lid also provide some microwave shielding by forming a metal box which acts as a high-pass filter with a cutoff frequency above the operating frequencies. We have three standard device sizes with associated PCB, puck, and lid assemblies: 7×7 mm, 10×10 mm, and 1×1 inch. All the devices used in this thesis use the 7×7 mm device size as the design doesn't require more space.

The packaged device is placed in a magnetic shield composed of two cans: an outer μ-metal can providing magnetic shielding at room temperature and an inner superconducting can made of aluminum. While the inner superconducting can will eventually fully expel all magnetic fields due to the Meisner effect, magnetic fields present during the cooldown will be pinned in vortices as the aluminum goes superconducting. The μ-metal can is made with a high magnetic susceptibility material which strongly reduces the fields deep inside the can without requiring a superconducting transition. In addition to magnetic shielding, the cans also provide some thermal radiation shielding from higher temperature stages of the fridge as they form a light-tight assembly around the qubit.

The final device package assembly is shown in Fig. 4.8. It consists of an adapter rod connecting the can lid to the base plate of the dilution fridge, another adapter rod

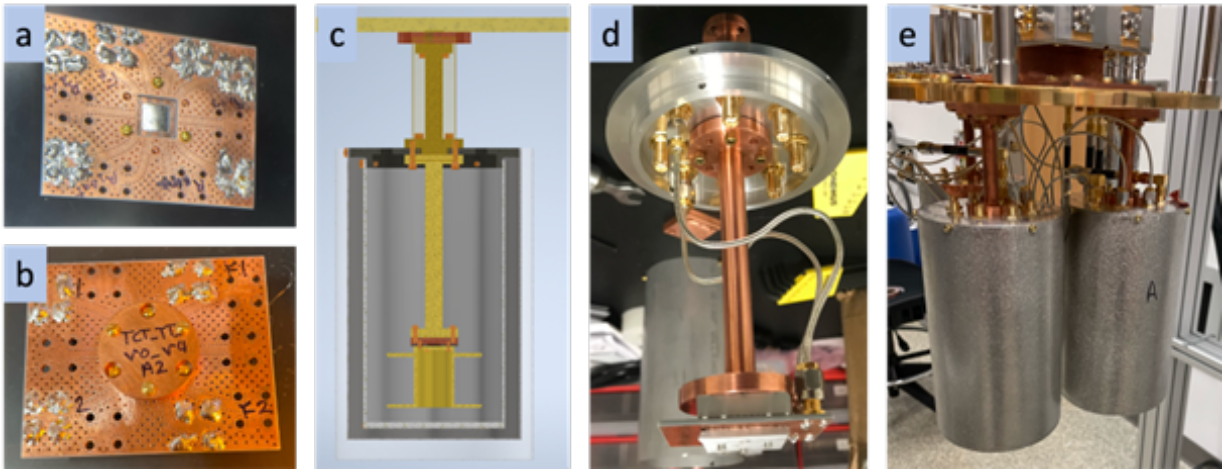

Figure 4.8: Device packaging [10]. (a) The device is placed in the puck which is attached to the PCB with 3 screws to provide thermal contact. Each launch pad on the chip is associated with a SMA connector which is soldered to the PCB (groupings of 4 solder blobs). (b) After wirebonding, the lid is placed on the qubit to protect it from debris and provide some microwave shielding. (c) Inventor assembly for the packing. An external magnet (lower cylinder) can be added to provide an additional form of flux biasing. (d)-(e) Pictures during dvice mounting in the fridge. Figure reproduced from [10]

on the bottom of the lid to connect to the qubit device, and the two concentric magnetic shield cans. All copper parts at the base plate are machined from oxygen free high thermal conductivity (OFHC) copper to provide good thermal contact. The two adapter rods are placed in direct thermal contact by drilling a hole through the can lid and machining a small stub sticking out of the qubit adapter rod to contact the fridge adapter rod.

4.3 Hardware

The hardware used for all of our experiments can be broken up into four main categories: microwave drives and shaping, flux bias sources, triggering and clocking, and digitization and signal processing. We'll briefly discuss the equipment used for each category and show a typical measurement configuration for qubit experiments. The idea here is to give each category an abstract model as well as a hardware implementation example so that the reader can improve and expand the existing set up with a solid understanding of each block.

4.3.1 Microwave drives and shaping

As we discussed in the previous chapter (Chap. 3), both state readout and qubit control pulses use microwave drives in the 2-10 GHz range. In addition, qubit pulses require full IQ modulation to shape both the frequency and amplitude of the drive for high fidelity operations. While direct synthesis boards are emerging as a viable options for full control of the microwave waveform (see QICK board [52]), our setup uses the traditional combination of an arbitrary waveform generator (AWG) with microwave generators. The AWG has two output types per channel: an analog *wave* (used for I or Q modulation) and a digital *marker* (used for pulse modulation). A qubit pulse requires two wave channels for full IQ control and one marker channel for leakage suppression. The cavity readout pulse only requires a marker channel allowing for each qubit/readout resonator pair to be controlled by a pair of AWG channels. The generators used for qubit pulses require both pulse and IQ modulation capability, also called vector modulation in the electrical engineering world, whereas the cavity generators only require a pulse modulation option. We use Rohde-Schwarz SGS100A vector generators for the qubit and Holzworth HS9000B generators for the cavity. Note, for cavity readout, we require two generators: the cavity pulse and the local oscillator (LO) used to mix down the signal coming back from the fridge. The qubit pulses are generated on a Zurich Instruments HDAWG with four output channels (each with a wave and marker output).

4.3.2 Flux bias sources

The tunability of a device in-situ is achieved by threading a magnetic field (flux) through a SQUID loop on the chip. This is achieved either with an external magnet (consisting of a coil of wire attached to the bottom of the device) or with an on-chip flux bias line (consisting of a shorted trace near the SQUID). In both cases, the applied flux is generated by a current flowing through the element near the SQUID. This can be achieve by applying a voltage or a current on the room temperature end of the line and calibrating out the

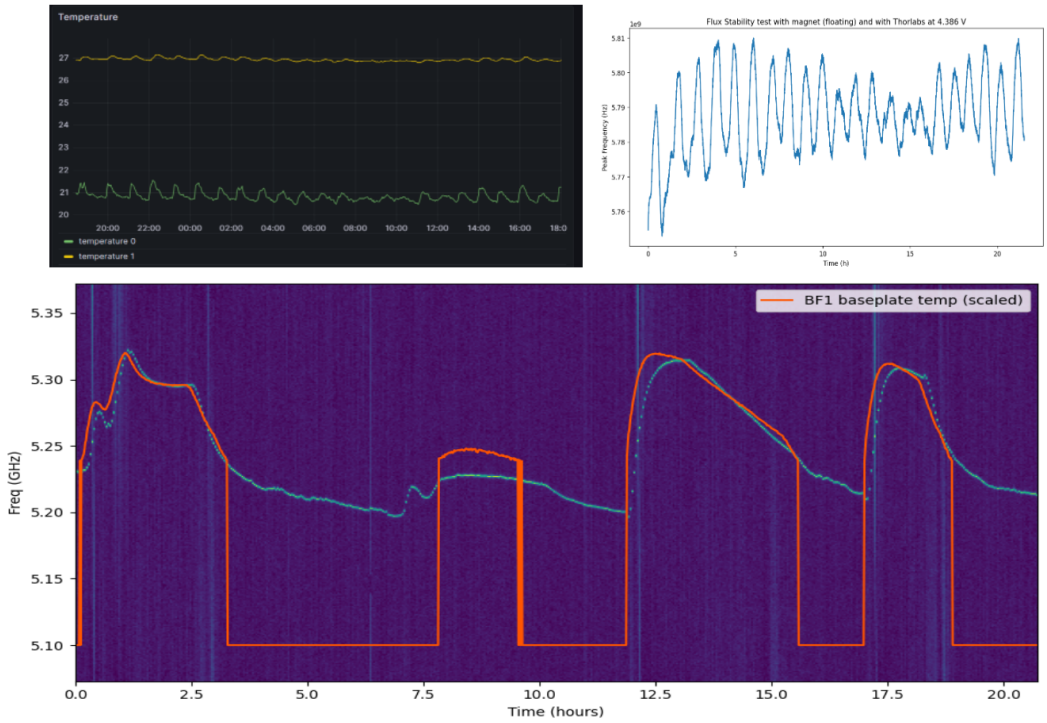

Figure 4.9: Qubit frequency drifts as a function of environmental temperature changes (top left panel is the temperature in the lab) and base plate temperature (bottom). The top right plot shows the qubit frequency tracking the lab temperature (left) over 20 hours. In the bottom plot, the qubit frequency (green) tracks the base plate temperature (orange line) with clear shifts during heating and cooling periods.

applied input to flux conversion (see Chapter 3 for details). Stanford Research System DC205 precision voltage sources and Keithley 2400 current supplies were used for voltage and current bias modes respectively. The external magnet is used for DC operating points with a 10 kHz low-pass filter to reduce the noise bandwidth reaching the qubit. Here the windings of the magnet act as a large inductor increasing the filtering effect. The on-chip flux bias line bandwidth is set by the requirements of the particular protocol, but for the experiments presented later in this thesis, it was set to 5 MHz.

In addition to frequency noise on the qubit, flux-tunable devices suffer from drifts as the temperature of the lab and the dilution fridge changes over time, as shown in Fig. 4.9. This is most likely due to a change in the current division both inside and outside the fridge as the components inside the filters and attenuators have a slight temperature dependence. In particular, we identified 1 kHz low-pass filters from Thorlabs (EF110) as

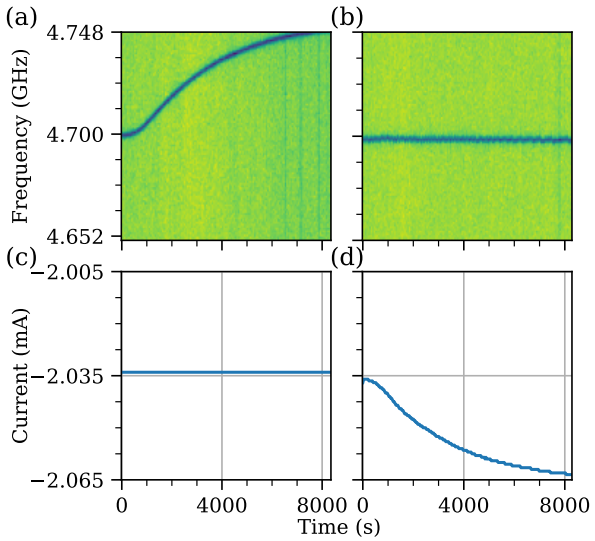

Figure 4.10: Stabilization of the qubit frequency using feedback on the applied current for flux biasing. (a) Qubit drift over 2-3 hours as the fridge base plate temperature is increased by applying an external voltage to another device without the servo adjusting the bias current (as shown in (c) with a constant current). (b) With the servo turned on, the qubit frequency remains steady while the current (d) is modified by the servo. Figure reproduced from [53].

having a large thermal coefficient due to the large inductor used to achieve the low cutoff frequency. We observed periodic oscillations in the qubit frequency which synchronized with the HVAC temperature oscillations in the lab. This environmental drift was strongly reduced by replacing the EF110 filter with an EF120 (10 kHz cutoff) filter and operating the magnet with a current supply (using a second FBL to connect the other terminal of the magnet to the return of the current supply). In addition to lab temperature induced drifts, we also observed drifts in the qubit frequency as the base plate temperature changed (for example due to another measurement). This drift could not be compensated by the room temperature set up and required periodic qubit frequency measurements and current adjustments to maintain the qubit at the desired set point as shown in Fig. 4.10. The details of the current servo will be laid out in Chapter 7.

4.3.3 Triggering and clocking

We use a master trigger generator (Keysight 33500B) to initiate each measurement in an experiment. In particular, measurements are always performed in the same section of the trigger period to avoid spurious signals due to the trigger signal itself. To that end, all schedules defined on the AWG are *backdated* from a specific measurement time (i.e. a T_1 measurement will sweep the position of the π excitation pulse backwards relative to the measurement pulse instead of performing the more intuitive sweep of the measurement position with a fixed pulse position). The repetition rate of the experiment should not exceed $1/(5T_1)$ to avoid starting a measurement with the qubit in the $|e\rangle$ state due to residual excitation. To maintain a fixed heterodyne beat note phase between measurement shots, the trigger rate should be commensurate with $IF = f_{cav} - f_{LO}$. If this condition is not satisfied, consecutive measurements will have a random phase and average out destructively to zero.

To guarantee a common absolute frequency basis, all the oscillatory sources are referenced to a central high precision frequency reference. We use a rubidium atomic clock from SRS (FS725) which outputs a 10 MHz reference clock signal. For experiments using multiple SGS generators, we found that clocking one to the atomic clock and the other to a 1 GHz reference signal from the first generator produced much better phase stability between the generators. In addition, the Holzworth generators have a known frequency drift between channels due to a small and constant absolute frequency error. See Fig. 4.11 for some examples.

4.3.4 Digitization and signal processing

Finally, the measurement is performed with an Acquiris SA220P digitizer card. This digitizer allows for on-board averaging which drastically reduces the amount of data read from the card as all the shots are collapsed into a single trace (per channel). The acquisition

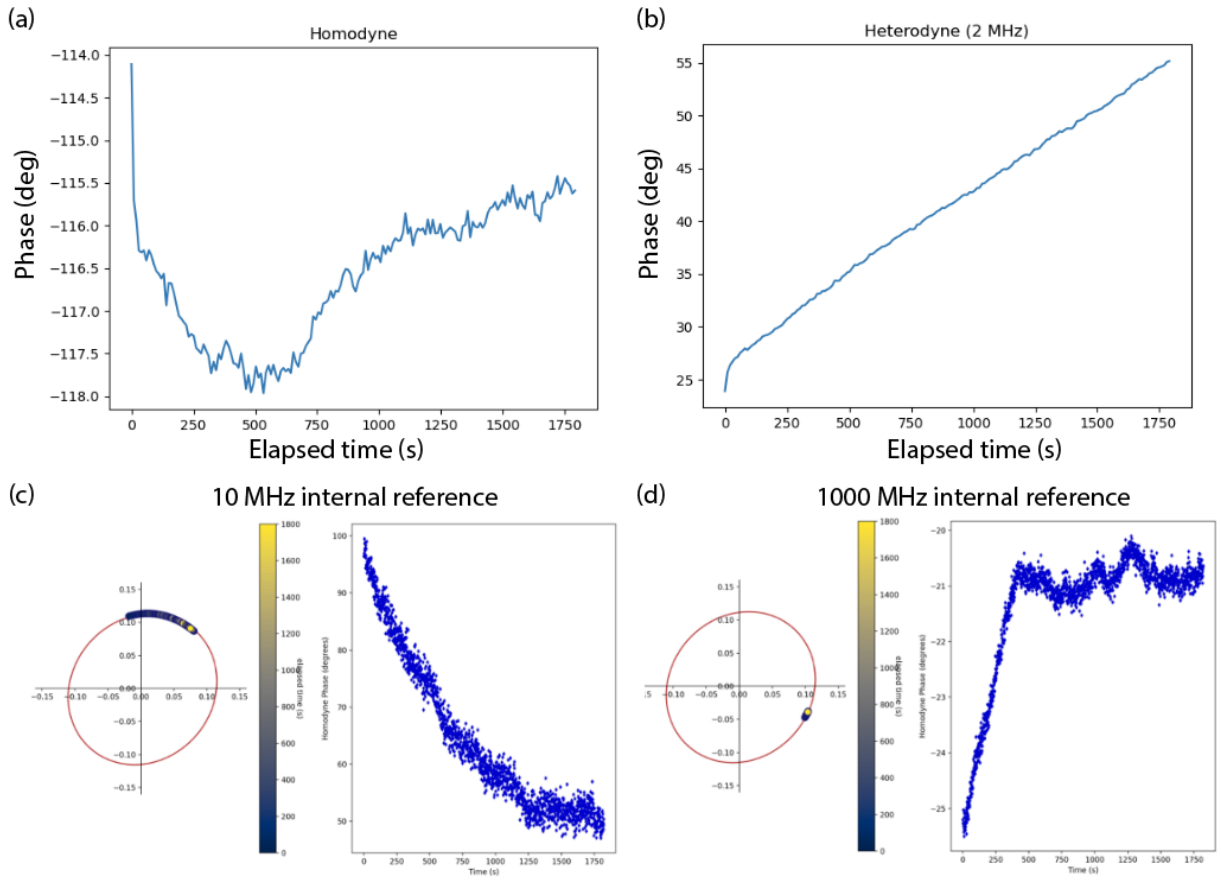

Figure 4.11: Frequency stability between Holzworth channels (top) and SGS sources (bottom). (a) Homodyne phase between two channels of the Holzworth set to the same frequency showing effectively no phase drift over 30 minutes. (b) When using a 2 MHz frequency difference between the channels, the phase of the beat note drifts by a consistent amount indicating a small residual frequency offset between the channels. The phase drift corresponds to an approximately 200 μHz offset in the desired beat frequency. (c) Homodyne stability for a pair of SGS generators referenced with a 10 MHz clock signal. Both the IQ point (left plot) and beat note phase (right) are shown over a 30 minute measurement showing a significant phase drift. (d) When a 1 GHz reference between the SGS generators is used, the beat note phase stays constant outside of an initial settling time. Both (c) and (d) were acquired in an office during COVID instead of the temperature controlled lab which most likely worsened the performance (especially the initial phase drift in panel (d) which was made worse by performing the calibration right after turning on the generator).

trigger is offset from the main experimental trigger to capture the result of the measurement pulse from the AWG. Originally, an IQ mixer was used to perform downconversion of the cavity signal requiring two digitizer channels for a single qubit (as both I and Q need to be digitized and downconverted for a full SNR measurement). However, by replacing the IQ mixer with a double balanced mixer we only require a single channel (as we obtain the IQ response from digital downconversion).

A typical configuration for the room temperature side of the output line consists of a band pass filter around the cavity frequency (consisting of a high and low pass filter combination with cutoff frequencies as close to the resonance frequencies as is available), a high quality RF amplifier (Amplitech APT5-04000800-0710-D6, 50 dB), a mixer/ down-converter to shift the cavity tone down to the measurement IF (Marki IQ4509 LXP or ZX05-14-S+), and a low-pass filter/DC amplifier configuration (SRS445) to filter out the higher sidebands and compensate for the mixer conversion loss.

4.3.5 Full measurement setup

Fig. 4.12 shows a full experimental hardware setup with associated part numbers.

4.4 Control code and scheduler

The last section we will discuss is the software control system developed when the lab was initially set up. A full, detailed documentation of the code is outside the scope of this thesis (the control code can be found on the lab GitHub account with documentation in progress), however we will discuss the general coding philosophy and key elements defining the control code system. The control code is organized into three main blocks: instrument drivers, schedule files, and experimental control scripts. We will briefly discuss each in turn and finish with a discussion on the main utility functions developed for data processing and analysis.

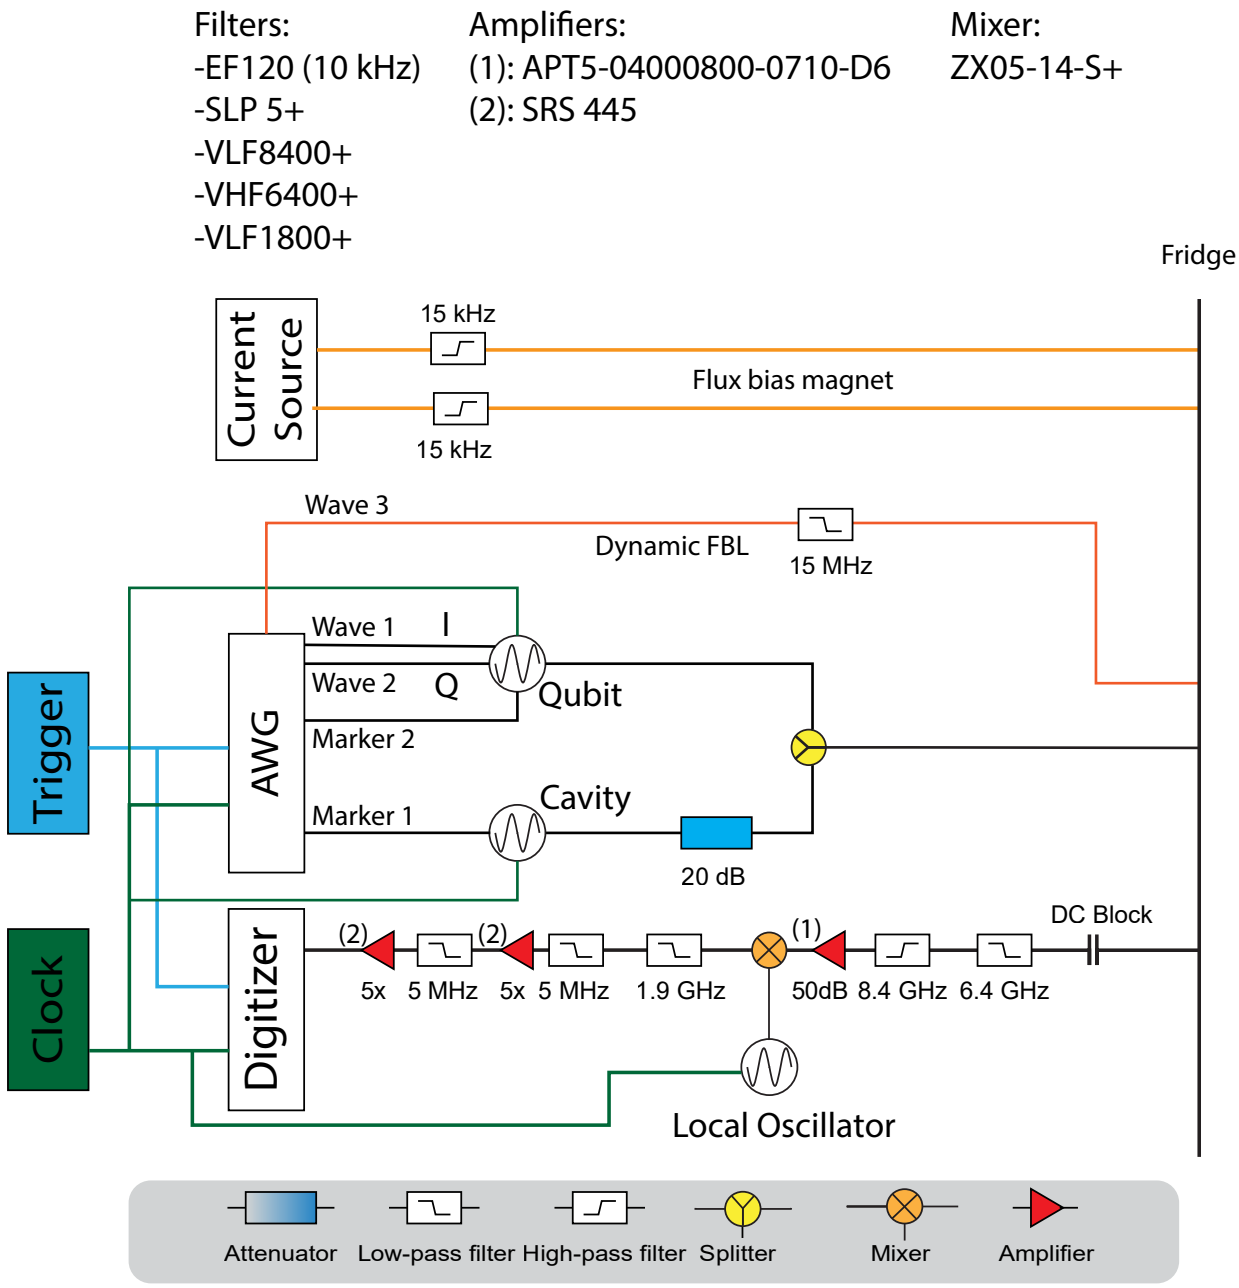

Figure 4.12: Wiring diagram for pulsed measurements.

4.4.1 Instrument drivers

Each instrument is controlled from the measurement computer through its driver which establishes a connection to the instrument, reads and sets individual settings, and in the cases of more complicated instruments like the VNA, AWG, or digitizer, configures batches of settings together for specific measurements. In all cases, the expected behavior of an instrument is:

```
instrument.prop = setting
instrument.module.prop = setting
```

where all the settings are accessed using the "dot" operator (with more detailed settings organized in sub modules as discussed in the SCPI instruments section later). The driver also should be able to acquire a "snapshot" of all the instrument settings to be saved along with the data. The driver implements *useful* error handling, where the error is returned to the user instead of being silently ignored. Finally, the driver should also enact some safeguards to prevent damage to the equipment. This is particularly relevant for the voltage and current supplies used for flux biasing where two safety measures are implemented: the maximum output range is limited (10 V for voltage supplies, 10 mA for current supplies) to avoid accidentally damaging the cryogenic equipment and changes to the bias points are applied through ramps instead of sudden jumps.

The communication with all the instruments uses the VISA (virtual instrument software architecture) protocol [54] which implements an intermediate layer between the user and the instrument to abstract away the low-level differences in communication protocol based on the connector used (USB, Ethernet, serial, etc.). **pyvisa** links to the VISA framework for use in python scripts. This protocol uses a unique VISA address for each instrument which specifies the communication protocol, some unique identifier (IP address, USB port, vendor ID, etc.), and a device. Some example addresses for equipment used in the lab:

- **PXI23::0::0::INSTR** (the address for the digitizer which uses a PCIe slot on the computer)
- **TCPIP0::rssgs100a110739::inst0::INSTR** (here we initialized an Ethernet device identified through it's vendor, model, and serial number)
- **TCPIP0::HS9002B-181::9760::SOCKET** (the Holzworth devices use a direct "SOCKET" type of connection instead of a standard "INSTR" requiring the specification of a communication port, 9760 in this case)
- **TCPIP0::192.168.10.139::inst0::INSTR** (the simplest, most general Ethernet VISA address simply uses the IP address of the device on the network)
- **ASRL5::INSTR** (asynchronous serial communication with the instrument on port 5)
- **USB0::0x0957::0x2507::MY58000681::0::INSTR** (USB instruments are defined in terms of a vendor ID, product ID, and model number).

Some of these addresses are fixed (defined by the product ID or name) whereas others can be dynamic (if the USB port changes or the IP address changes) requiring an update to the configuration script if a power outage occurs or the setup is modified.

Each instrument has a native driver specified by the manufacturer which we wrap in a custom python class to present the user with standard instrument behavior to help the modularity of experiments. For example, setting the frequency or output power of a microwave generator might be different for different manufacturer but at the user level, both commands are simply `generator.freq = 5e9`. There are three main types of native instrument drivers (in increasing order of work required for a wrapper): SCPI based instruments, python API based instruments, and DLL based instruments. The framework described above can also be used to create *virtual* instruments which are composed of multiple instruments. For example, the cavity generator and associated local oscillator generator

used to perform readout are combined into a single higher level instrument. This allows for simpler control of the cavity generator as a single frequency is specified and the LO is moved to maintain the specified IF frequency.

SCPI instruments

The simplest instruments implement the SCPI (standard commands for programmable instruments) protocol [55] to set and read settings. The SCPI protocol specifies a standard set of commands that all instruments should implement (for example `*RST` resets the instrument to a configuration specified by the manual or `*IDN?` reads the instrument identifier string) as well as a general format for reading and writing settings. The SCPI protocol is particularly simple to parse as it utilizes a hierarchical organization for setting groups: for example, the commands to control the output amplitude, phase, and frequency of the SGS100A generator all stem from the `SOURCE` module. Reading and writing settings use the standard `inst.query()` and `inst.write()` commands implemented by pyvisa where the same resource string is used in both cases except it is appended with a `?` for queries and the value of the setting for write. While one could write a python object that specifies properties for each command string we desire to use, this is inefficient and requires a lot of boiler plate code.

The `SCPIinst` class implements a template that all SCPI based instruments inherit from and converts a dictionary of property names (and groupings for more complex instruments) into an object whose settings are class attributes (queried and set using the standard python dot operator). It also implements a standard `settings` which loops through all the defined settings and returns a dictionary to save alongside the data. Finally, `SCPIinst` also implements error checking after each command using either the standard SCPI error command or the instrument specific error command if specified. The code shown below is an example implementation of the SCPIinst object for the SGS100A generators. The commands are grouped into a `commandlist` dictionary which can be further subdivided

into as many dictionaries if needed. The dictionary keys become the object attribute once initialized (for example, the I leakage value for the IQ modulator can be accessed using `generator.IQ.Ileak`). The functions in the "core" dictionary are top level commands which can be accessed directly from the instrument itself. This generator also provides some utility functions which set multiple settings allowing for more complex behavior and flexibility to match behavior between instruments if a slightly different set of settings is required to reach the same configuration.

```
class SGS100A(SCPIinst):
    errcmds             = {}
    errcmds['error']    = 'SYST:ERR?'
    errcmds['serror']   = 'SYST:SERR?'

    commandlist = {}
    commandlist['core']    = {}
    commandlist['IQ']      = {}
    commandlist['Ref']     = {}
    commandlist['LO']      = {}
    commandlist['RefOut'] = {}

    core = {}
    core['Output'] = 'OUTPut:STATe'
    core['Power']  = 'SOURce:POWer'
    core['Phase']  = 'SOURce:PHASe'
    core['Freq']   = 'SOURce:FREQuency'
    core['Offset'] = 'SOURce:FREQuency:OFFSet'

    IQ = {}
    IQ['Mod']    = 'SOURce:IQ:STATe'
    IQ['Imp']    = 'SOURce:IQ:IMPairment'
    IQ['Ileak'] = 'SOURce:IQ:IMPairment:LEAKage:I'
    IQ['Qleak'] = 'SOURce:IQ:IMPairment:LEAKage:Q'

    Pulse = {}
    Pulse['Mod']       = 'SOURce:PULM:STATe'
    Pulse['Source']    = 'SOURce:PULM:SOURce'
    Pulse['Polarity'] = 'SOURce:PULM:POLarity'
    Pulse['Trig_out'] = 'CONNector:TRIGger:OMODe'

    Ref = {}
```

```
Ref['Source']     = 'SOURce:ROSCillator:SOURce'
Ref['Frequency'] = 'SOURce:ROSCillator:EXTernal:FREQuency'

LO = {}
LO['Source'] = 'SOURce:LOSCillator:SOURce'

RefOut = {}
RefOut['Source']     = 'CONNector:REFLo:OUTPut'
RefOut['Frequency'] = 'SOURce:ROSCillator:OUTPut:FREQuency'

commandlist['core']   = core
commandlist['IQ']     = IQ
commandlist['Pulse'] = Pulse
commandlist['Ref']    = Ref
commandlist['LO']     = LO
commandlist['RefOut'] = RefOut

def __init__(self, address):
    self.instrument_type = 'SGS'

    super().__init__(address, self.commandlist, self.errcmds)

def enable_IQ(self, Ileak=None, Qleak=None):
    self.IQ.Mod = 'On'
    self.IQ.Imp = 'On'
    if Ileak is not None and Qleak is not None:
        self.IQ.Ileak = Ileak
        self.IQ.Qleak = Qleak

def disable_IQ(self):
    self.IQ.Mod = 'Off'

def enable_pulse(self, source='EXT', polarity='NORM',
                 trig_conn='PEMSource'):
    self.Pulse.Mod    = 'On'
    self.Pulse.Source = source
    self.Pulse.Polarity = polarity
    self.Pulse.Trig_out = trig_conn

def disable_pulse(self):
    self.Pulse.Mod = 'Off'
```

Python API based instruments (HDAWG)

Some instruments arrive with a prebuilt python API which drastically reduces the required overhead to incorporate them into the existing code base. For the HDAWG, Zurich Instruments also provides a web-based interface to control the instrument while logging the API calls made when changing settings. This feature was extensively used to build the HDAWG object which grouped settings and commands into the functional blocks we required (`awg`, `trigger`, `channel`, etc.).

DLL based instruments (Acqiris Driver)

Finally, some experiments only provide precompiled dll libraries to control the instruments. These libraries can easily be imported and used in a `C++` script but require intensive wrapping in python to be usable. As the compiled functions are typically written in a C based language, the python wrapper needs to implement C-typed variables instead of relying on the flexibility native to python. Each individual function requires its own python equivalent wrapper with specified input and return types.

While the Acqiris driver used this dll format and early driver prototypes implemented the function wrapping described above, technical issues in the implementation in python led to the company providing us with a python wrapper. While this higher-level interface allowed for a more object-oriented approach in python, setting validation checks performed in the `C` driver required settings to be applied in batches which is performed by the `card.SetParams()` function in the python driver. This means that until that function is called, changes to settings on the card object in software *remain* in software.

4.4.2 Scheduler

The last instrument-like object that we will discuss is the scheduler object. The experiment schedule consists of all the pulses required for a particular measurement (qubit drives, flux pulses, readout measurements). The scheduler consists of analog and digital channels

(implemented as numpy arrays) which correspond to the AWG wave and marker ports respectively. An experimental sequence is then defined by adding *pulse* objects to the analog channels (controlling the IQ modulator on the generator) or by defining an associated *marker* window in the digital channel (controlling the pulse modulation portion). The schedule can then be compiled into a format usable by the HDAWG and uploaded to run the experiment. The scheduler also generates a plot of all the defined traces (Fig. 4.13) to allow for debugging for more complicated sequences.

A second generation of scheduler is currently being developed (and used for the last two experiments discussed in this thesis) to make the scheduler act like the other instruments defined above. This allows a schedule parameter (like the pulse length in a Rabi measurement or the delay time in a T_1 measurement) to be varied with the exact same command as varying the frequency of a generator making the experimental scripts much simpler as they can implement the generalized `sweep(instrument, parameter, range)` function.

4.4.3 Experiment template

Each experimental script consists of three main parts: instrument initialization, parameter sweep, and data plotting and logging. The experiment has an associated *driver* script which runs the experiment and allows the user to change the parameters of the measurement (probe power, pulse length, etc.). All measurements use the same function signature: `experiment_type(instruments, settings)`. To improve readability and reduce the number of settings visible to the user when running an experiment, we separate out the settings into two blocks: `exp_globals` and `exp_settings`. `exp_globals` contains all the high-level configuration settings which are not modified on a run to run basis (the data saving folder, attenuation configuration, trigger rate, etc.). `exp_settings` on the other hand groups all the settings relevant to the particular measurement (qubit frequency, cavity probe power, parameter sweep range, etc.). The instruments dictionary collects the list of instruments

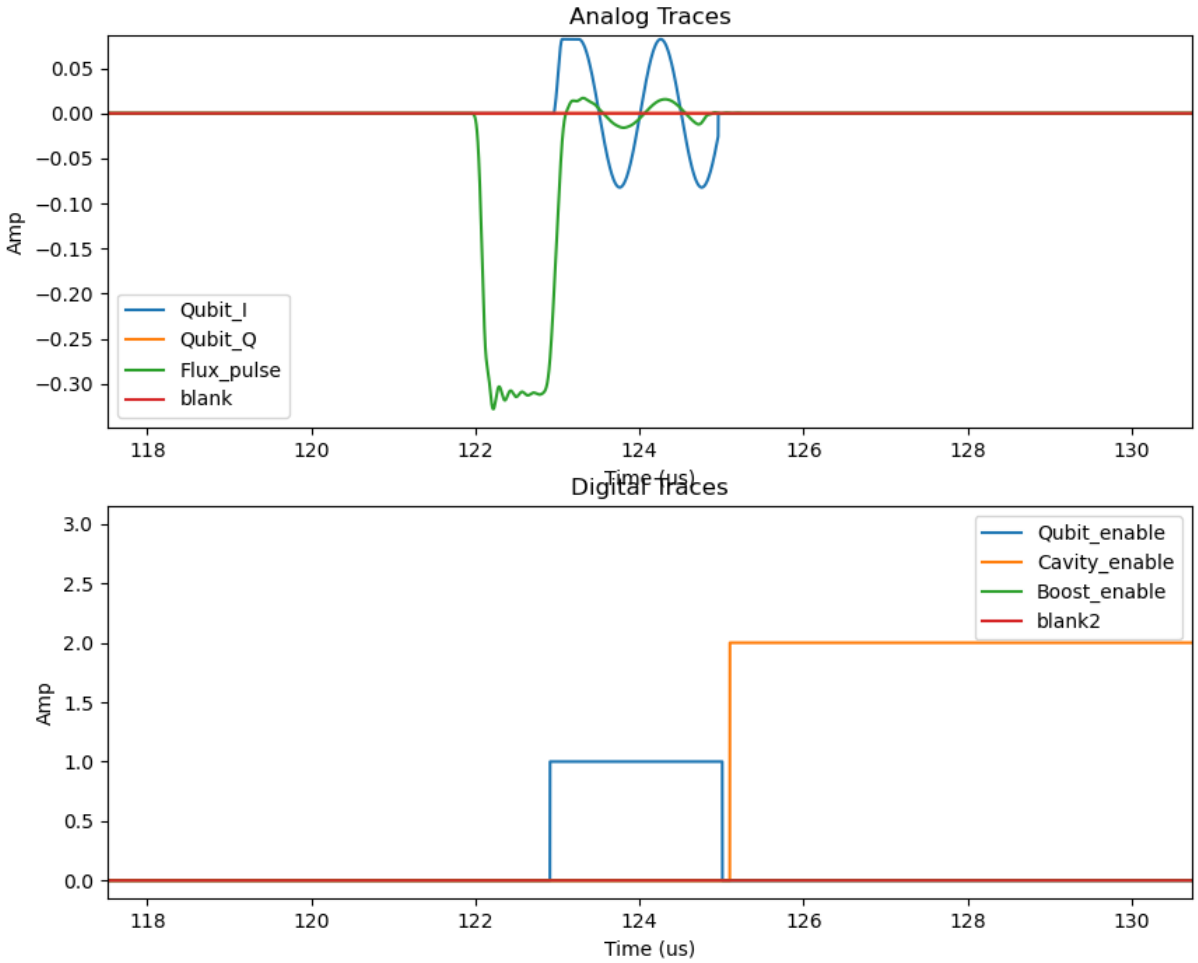

Figure 4.13: Example schedule for the topological pumping sequences. The qubit control (blue lines in both panels) consists of an analog pulse (top) bracketed by a digital marker (bottom). Measurements are implemented with a simple digital pulse sent to the cavity generator. The flux pulse is sent directly to the fridge to the on-chip flux bias line for dynamic qubit frequency modulation. "0" for the AWG is set by the experimental trigger with the measurement position (125 μs) being fixed relative to it.

used for the particular measurement. It is also used during the data saving phase to gather a "snapshot" of the hardware configuration.

The core of any measurement is the loop over the sweep parameter (for example qubit frequency in a pulsed spectroscopy measurement): the configuration is updated to the new measurement point (either updating an instrument setting or by changing the schedule), the measurement is acquired on the digitizer by averaging N shots, the measurement results are converted to an IQ point using digital down conversion, and reference $|g\rangle$ and $|e\rangle$ traces can be taken to convert the output IQ to a population. The `sweep()` function described in the scheduler section attempts to simplify most scripts to an arbitrary instrument initialization followed by any parameter sweep available to the current configuration. This lead to the development of a second iteration of the scheduler object as well as the use of virtual instruments for the readout generators (combining two SGS generators into a single `readout_generator` object which automatically updates the LO frequency if the cavity frequency is changed).

Chapter 5

Tunable Coupler

In this chapter, we will discuss the project to develop a tunable qubit-cavity coupling element based on the tunable qubit-qubit coupler developed by the Martinis group [41]. We will start by framing this coupler in the context of photonic lattices, provide a simplified circuit model, and discuss the design constraints for our particular application. We will then discuss the results from a proof of concept device and provide some possible modifications and extensions to the concept.

5.1 Background

The goal of this project is to develop a tunable coupler between a qubit and the modes of a photonic lattice. By tuning the qubit near resonance with a band of interest and lowering the coupling strength to reduce the interaction with the other bands (in particular if there are other bands nearby in frequency space), we can selectively isolate a particular band to simulate different dynamics. In particular, coupling a pair of qubits (in spatially separated lattice sites) to the same band will have different effective interactions depending on the curvature of the band (flat, linear, quadratic). The coupler studied in this project additionally attempts to minimize the so-called on-diagonal shift on the lattice site (i.e. the superconducting resonator) to avoid introducing defects in the lattice at the qubit coupling

locations.

5.1.1 Photonic lattices

Lattices consist of a periodic array of *sites* connected by *edges*. In materials, these sites consist of atoms and the edges are realized through an electromagnetic interaction between the atomic orbitals. The fundamental building block of a lattice consists of a *unit cell* whose connection graph dictates the properties of the lattice [56, 57]. Changing the atomic arrangement, i.e. changing the unit cell, dramatically changes the properties of the resulting lattice. For example, both diamonds and graphite are pure carbon lattices with very different properties due to their different unit cells: the tetrahedral lattice of diamond makes it transparent and the hardest material on earth while the hexagonal lattice of graphite makes it a dark material soft enough to make marks when writing on paper.

To compute the electrical properties of a lattice we will use the tight-binding model in which the electrons are localized on the grid defined by the atoms but can tunnel to neighboring sites[56]:

$$H = \Sigma_i \Delta_i a_i^\dagger a_i - \Sigma_{<i,j>} t_{ij} a_i^\dagger a_j + h.c. \tag{5.1}$$

where $a_i, a_i^\dagger$ are the annihilation and creation operators for an electron at site i,Δ_i is the on-site energy of the electron at site i, t_{ij} is the tunneling rate between sites i and j, and the sum over i, j only accounts for nearest neighbor tunneling. Each individual atom has the exact same orbital energy levels, however, the coupling between sites leads to *hybridization* between the orbitals which splits the energy levels. In an infinite lattice, the hybrized orbitals form a continuous *band* as the energy splitting between discretized levels becomes vanishingly small. *Band gaps* occur when consecutive bands do not overlap in energy leading to a range of energies forbidden in the lattice. The width of a band-gap determines the photon frequencies that the material can absorb, which has far-ranging effects on the behavior of the material making it transparent or opaque and conductive or insulating. The

band structure of a material is the full representation of the energy-momentum relationship for the electrons in a material. In general, these bands have complicated relationships between E and k but they can be Taylor expanded near extrema into three main types: quadratic bands ($E \propto \alpha k^2$), linear bands ($E \propto \alpha k$), and flat bands ($E = E_0$). In all three cases, we can build intuition by computing the group velocity (i.e. the speed of propagation of an excitation in the lattice):

$$v_g = \frac{\partial \omega}{\partial k} \tag{5.2}$$

where ω is the frequency (recall $E = \hbar\omega$) and k is the momentum of the excitation. The group velocities for each band edge type are $2\alpha k$, α, 0 for the quadratic, linear, and flat bands respectively. Excitations in quadratic bands therefore behave like massive particles with an effective mass given by α^{-1} whereas linear bands generate photon-like (i.e. massless) excitations which propagate at a fixed velocity α. Flat bands have no group velocity indicating localized modes which do not propagate through the lattice. In addition to localization, flat bands are massively degenerate as all excitations have exactly the same energy. For non-bipartite graphs such as the Kagome lattice with anti-ferromagnetic interactions, this can lead to the emergence of massively entangled ground states due to spin frustration.

The photonic lattices studied in the group consist of arrays of CPW resonators capacitively coupled to each other at their end-points [59–61]. In contrast to regular lattices, the *sites* in the photonic lattices are now extended objects and the *edges* of the lattice are point-like objects. The extended nature of the sites leads to effective lattices that are line graphs of the layout lattice which have unique properties studied in [62, 63]. In addition, the resonance frequency of the CPW resonators is set only by the length of the cavity (to first order [10]), allowing the resonators to be coiled or meandered effectively decoupling the distance between the end points from the resonance frequency. This inversion of dimensionality (extended sites and point-like interactions) along with the freedom to modify the shape of the CPW resonator without affecting the photonic mode frequency allow for

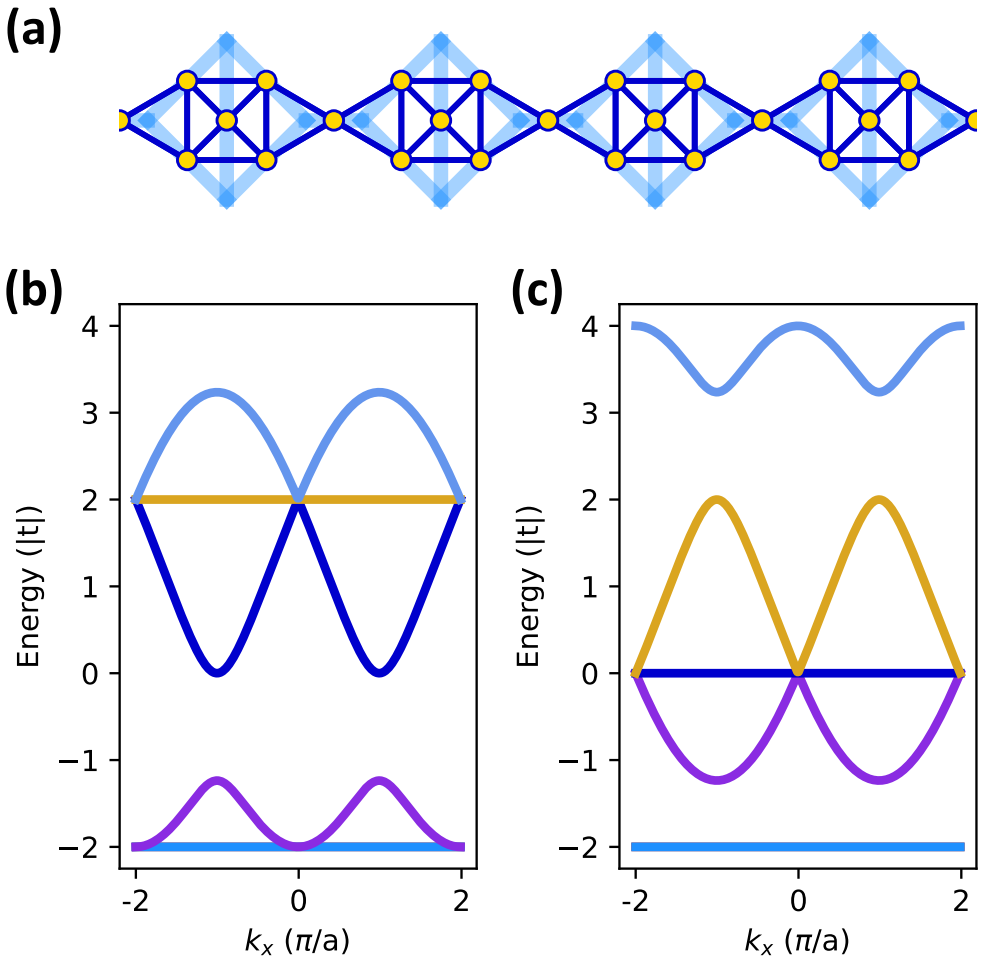

Figure 5.1: (a) Quasi 1D lattice with diamond unit cells. The CPW resonators are represented by the light blue lines and the effective lattice is given by the yellow points (*sites*) and dark blue lines (*edges*). (b), (c) Band structure for the half-wave and full-wave modes of the CPW lattice. In the full-wave modes, (c), the flat band at -2$|t|$ is isolated from the rest of the lattice allowing for its dynamics to be studied independently. Linear and quadratic bands are also present in this lattice which can be tuned to with a tunable qubit. Figure reproduced from [58]

a decoupling between the true graph representation of the lattice and the physical embedding on a chip. For example, embedding a curved-space lattice on a flat 2D substrate leads to displacements of the sites when embedding which cause the edges to stretch or shorten. The extended nature of the lattice sites in a CPW lattice allow this distortion to be absorbed in the exact path of the resonators (for example, if the lattice points are pushed together, the CPW can be meandered more tightly such that the end points are close). This freedom to embed non-Euclidean lattices on a flat substrate has been used to study a hyperbolic lattice [29]. In addition to interesting non-Euclidean layouts, the extended nature of the lattice sites in CPW photonic lattices also leads to interesting band structures. In particular, they reliably lead to the formation of flat bands due to their line graph nature [62, 63]. All three band types (quadratic, linear, flat) can be present in a CPW lattice allowing for the different nature of lattice excitations to be studied with a singular device [64]. An example unit cell with a gapped flat band (i.e. a flat band fully separated from all other energy bands) is shown in Fig. 5.1 and extensively studied by the group [10, 58].

5.1.2 Tunable coupling

Two-qubit operations, such as entanglement, require coupling between qubits and have a gate speed set by the coupling strength. However, for complex qubit processor architectures with many qubits, always-on coupling between neighboring qubits leads to unwanted crosstalk. This problem is fairly easily solved for atom based architectures (such as trapped ions or neutral atoms) as the atoms can be moved around to dynamically adjust their coupling strength ([1, 3]). Superconducting qubit based architectures however do not have that flexibility and require a different approach. The coupling strength is typically modulated in one of two ways: flux biasing a frequency tunable coupling element [41, 65] or parametrically driving a coupling circuit with fixed coupling strength at the frequency difference between the qubits to be coupled [66, 67].

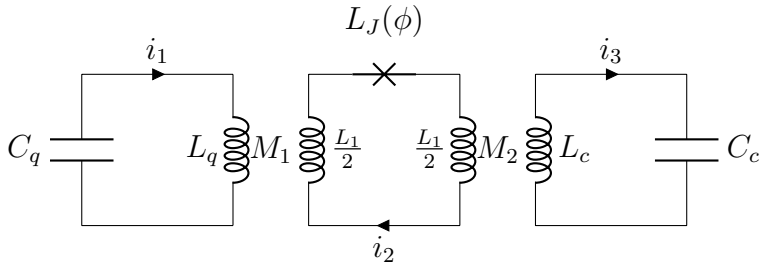

Figure 5.2: Circuit representation for the tunable coupler, currents are labeled positive in the clockwise direction. The mutual inductances M_1, M_2 can be independently tuned by changing the coupling geometry (see subsection 5.2). The qubit and cavity are modeled as LC resonators for a first order analysis.

We develop an inductive tunable coupler derived from the design described in [41, 68] using a Josephson junction as a tunable inductor to vary the coupling strength. The coupler consists of a DC SQUID inductively coupled to both the cavity and the qubit. By varying the flux across the junction in the SQUID, we vary the effective inductance of the loop:

$$L_{loop} = L_1 + L_J(\phi) = L_1 + \frac{L_{j0}}{\cos(\phi)} \tag{5.3}$$

where ϕ is the normalized flux (Φ_{ext}/Φ_0), L_1 is the linear inductance of the loop, and L_{j0} is the nominal inductance at zero applied flux. Near $\phi = \pm\pi/2$, the inductance of the junction diverges, effectively creating an open circuit in the RF squid loop which in turn decouples the qubit from the cavity.

5.1.3 Circuit Model

The basic circuit representation of the coupler is presented in Fig. 5.2. Both the cavity mode and qubit mode are modeled as LC oscillators to simplify the analysis. We can use Kirchhoff's voltage law (for AC currents) in each loop to solve for the normal modes in the

circuit:

$$\frac{I_1}{j\omega C_q} - j\omega M_1 I_2 + j\omega L_q I_1 = 0 \tag{5.4}$$

$$j\omega(L_1 + L_J)I_2 - j\omega M_1 I_1 - j\omega M_2 I_3 = 0 \tag{5.5}$$

$$j\omega L_c I_3 - j\omega M_2 I_2 + \frac{I_3}{j\omega C_c} = 0 \tag{5.6}$$

where L_q, C_q are the equivalent inductance and capacitance for the qubit mode, L_c, C_c are the equivalent capacitance for the cavity mode, and M_1, M_2 are the mutual inductance between the qubit and the DC squid and the cavity and the RF squid respectively. I_2 can be eliminated from the expressions above and after simplification the current relations read:

$$\begin{pmatrix} \frac{-1}{\omega^2 C_q} + \left(L_q - \frac{M_1^2}{L_T}\right) & \frac{-M_1 M_2}{L_T} \\ \frac{-M_1 M_2}{L_T} & \frac{-1}{\omega^2 C_c} + \left(L_c - \frac{M_2^2}{L_T}\right) \end{pmatrix} \begin{pmatrix} I_1 \\ I_3 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \tag{5.7}$$

where $L_T = L_1 + L_J$ is the total inductance of the coupling loop. The normal modes can be found by setting the determinant of the 2×2 matrix equal to zero. We can already identify two kinds of effects: an *on-site* shift in the resonance frequency of the modes (the correction to the effective inductance $\frac{-M_i^2}{L_T}$), as well as an off-diagonal term which leads to coupling between the modes ($\frac{M_1 M_2}{L_T}$). Dividing each row by the effective inductance $L_i - \frac{M_i^2}{L_T}$ and defining the modes $\omega_{q'}, \omega_{c'}$ as the new resonance frequencies of the qubit and cavity modes respectively gives:

$$\omega_{q'} = \frac{1}{\sqrt{C_q(L_q - \frac{M_1^2}{L_T})}} \tag{5.8}$$

$$\omega_{c'} = \frac{1}{\sqrt{C_c(L_c - \frac{M_2^2}{L_T})}}. \tag{5.9}$$

The off-diagonal elements are now scaled by the effective inductance of each circuit: $g_i = \frac{M_1 M_2}{L_T L_{i,eff}}$. In this notation, the matrix expression simplifies to:

$$\begin{pmatrix} \omega_{q'}^2 - \omega^2 & g_1\omega^2 \\ g_2\omega^2 & \omega_{c'}^2 - \omega^2 \end{pmatrix} \tag{5.10}$$

where g_1, g_2 can be different if the circuit inductances are different for the qubit or cavity modes. We compute the determinant:

$$(\omega_{q'}^2 - \omega^2)(\omega_{c'}^2 - \omega^2) - \omega^4 g_1 g_2 = 0 \tag{5.11}$$

$$\omega_{q'}^2\omega_{c'}^2 - \omega^2(\omega_{q'}^2 + \omega_{c'}^2) + \omega^4\left(1 - g_1 g_2\right) = 0 \tag{5.12}$$

which is a quadratic equation for ω^2 whose solution provides the two new coupled mode frequencies. Solving for $\omega_\pm^2$ (the two roots of the quadratic equation) yields:

$$\omega_\pm^2 = \frac{(\omega_{q'}^2 + \omega_{c'}^2) \pm \sqrt{(\omega_{q'}^2 + \omega_{c'}^2)^2 - 4\omega_{q'}^2\omega_{c'}^2(1 - g_1 g_2)}}{2(1 - g_1 g_2)} \tag{5.13}$$

$$= \frac{(\omega_{q'}^2 + \omega_{c'}^2) \pm \sqrt{(\omega_{q'}^2 - \omega_{c'}^2)^2 + 4\omega_{q'}^2\omega_{c'}^2 g_1 g_2}}{2(1 - g_1 g_2)} \tag{5.14}$$

which can be simplified in two limiting cases: resonance ($\omega_c' = \omega_{q'} = \omega_0$) and the far detuned limit ($|\omega_{q'} - \omega_{c'}| \gg 1$). On resonance, we recover the original modes symmetrically split from their bare resonance by a constant:

$$\omega_\pm^2 = \omega_0^2 \frac{1 \pm \sqrt{g_1 g_2}}{1 - g_1 g_2} = \omega_0^2 \frac{1}{1 \pm \sqrt{g_1 g_2}} \approx \omega_0^2\left(1 \pm \sqrt{g_1 g_2}\right) \tag{5.15}$$

$$\omega_\pm = \omega_0 \sqrt{1 \pm \sqrt{g_1 g_2}} \approx \omega_0 \pm \frac{M_1 M_2}{2L_T\sqrt{L_{c'} L_{q'}}}\omega_0 \tag{5.16}$$

whereas the off-resonance case does not simplify as cleanly:

$$\omega_{\pm}^2 = \frac{(\omega_{q'}^2 + \omega_{c'}^2) \pm |\omega_{q'}^2 - \omega_{c'}^2| \sqrt{1 + \frac{4\omega_{q'}^2 \omega_{c'}^2}{(\omega_{q'}^2 - \omega_{c'}^2)^2} g_1 g_2}}{2(1 - g_1 g_2)} \tag{5.17}$$

$$\omega_{+}^2 = \omega_{q'}^2 \frac{1 + \frac{\omega_{c'}^2}{(\omega_{q'}^2 - \omega_{c'}^2)} g_1 g_2}{1 - g_1 g_2} \Rightarrow \omega_{+} \approx \omega_{q'} \tag{5.18}$$

$$\omega_{-}^2 = \omega_{c'}^2 \frac{1 - \frac{\omega_{q'}^2}{(\omega_{q'}^2 - \omega_{c'}^2)} g_1 g_2}{1 - g_1 g_2} \Rightarrow \omega_{-} \approx \omega_{c'} (1 + g_1 g_2) \tag{5.19}$$

where we have assumed $\omega_{c'} \gg \omega_{q'}$ to simplify the final expressions.

Using the on-resonance results, we can identify the coupling strength (which depends on the total inductance in the coupler loop) and on-diagonal shifts for each oscillator as:

$$g = \frac{M_1 M_2}{2L_T \sqrt{L_c' L_q'}} \omega_0 = \frac{g_0}{L_1 + L_J(\phi)} \tag{5.20}$$

$$\delta\omega_i \approx -\frac{M_i^2}{2L_T L_i} \omega_0 \tag{5.21}$$

which can be semi-independently adjusted by keeping $M_1 M_2$ constant but varying the ratio of M_1 to M_2. By designing an asymmetric coupling scheme, we can shift most of the on-diagonal shift on the qubit mode and reduce the effects of the inductive loading on the resonator which we discuss in section 5.2.1.

Finally, recall that the inductance of the coupling loop, L_T, is composed of a linear component (L_1) and a flux tunable component ($L_J = \frac{L_{J0}}{\cos(\phi)}$). The coupling strength g can then be modulated by varying the applied flux in the loop, going from 0 to $\pm g_0$ where g_0 is the coupling at zero applied bias.

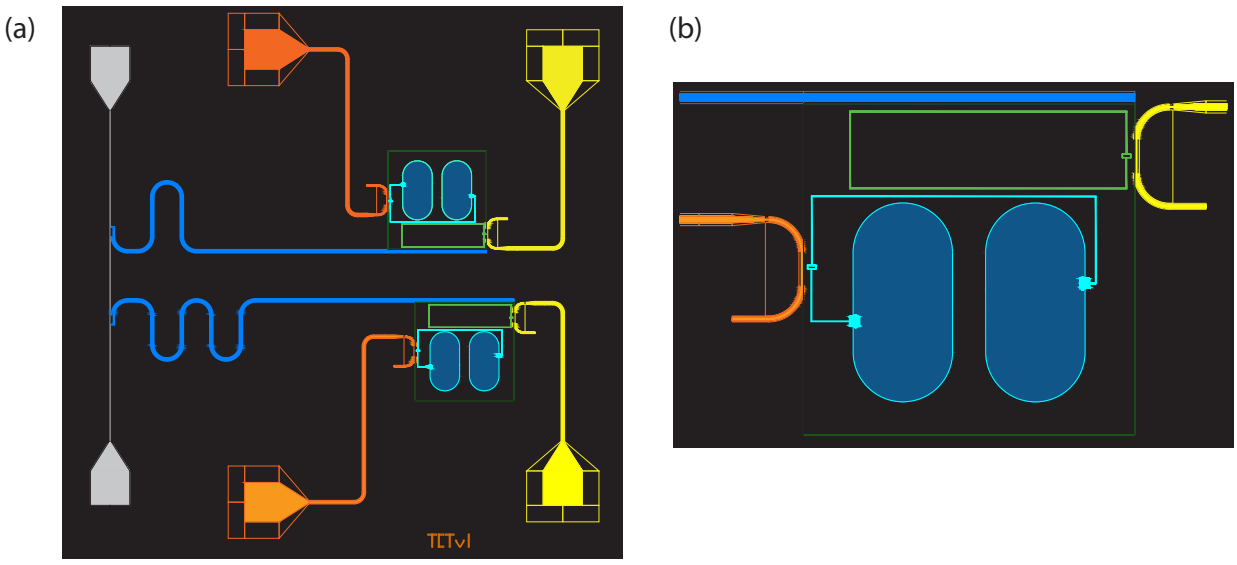

Figure 5.3: (a) Full chip CAD for the tunable coupling device with two qubit-resonator pairs for redundancy. The qubit (dark blue capacitor paddles and cyan inductive trace) is flux biased with the orange trace and coupled to the inductive loop (dark green) by running the trace parallel to the loop. The coupled flux bias line (yellow) applies an external flux to the inductive loop to vary the effective junction inductance and modify the resulting qubit-cavity coupling. The cavity (dark blue trace) is a $\lambda/4$ CPW resonator coupled to a feedline (gray) for hanger style readout.

5.2 Device Layout

The first prototype design is shown in Fig. 5.3 where the cavity and qubit are symmetrically coupled by the DC SQUID. The cavity is implemented with a $\lambda/4$ CPW resonator and the qubit is fabricated as a tunable transmon. To minimize the direct capacitive coupling between the qubit and cavity (which is fixed as opposed to the inductive coupling), the qubit is placed near the shorted end of the cavity and oriented parallel to it such that the field from the cavity is perpendicular to the field between the qubit capacitor paddles. The coupler element is a simple trace of superconducting tantalum with a Josephson junction for flux tunability. The qubit and coupler flux bias lines are placed on opposite sides of the qubit to minimize direct crosstalk. The large physical footprint of the coupler and qubit portions of the circuit will be challenging for implementation in a photonic lattice architecture, however, they provide simpler fabrication for a proof of concept device while

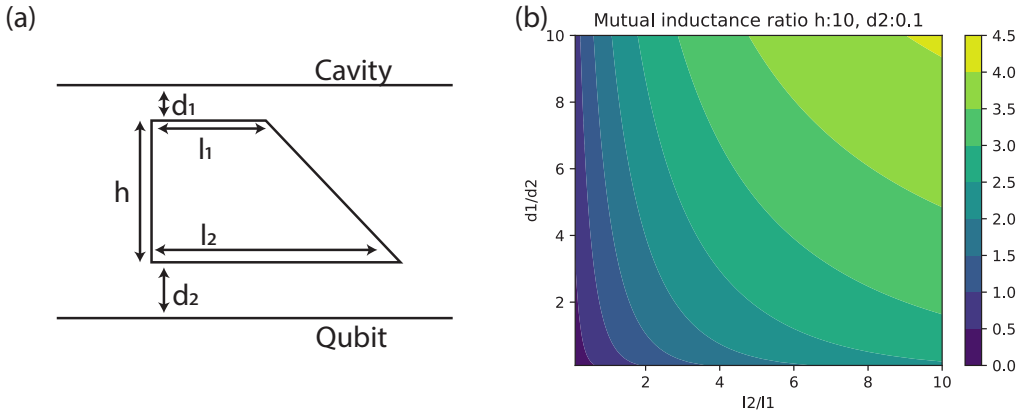

Figure 5.4: (a) Schematic for asymmetric inductive coupling to a common loop (where the inductive coupling is modeled as a wire coupled to an inductive loop). By adjusting l_1, l_2 and d_1, d_2 one can achieve arbitrarily large asymmetries between M_{cc} and M_{qc} as shown in panel (b)

still implementing the general coupling scheme we desire to explore.

While not implemented in the proof-of-concept prototype, asymmetric mutual inductances for the qubit and cavity to the coupler can be achieved by using a non-rectangular shaped coupling loop as shown in Fig. 5.4(a). As discussed in the circuit model section, the coupling strength between the qubit and cavity is proportional to the *product* of the mutual inductances whereas the on-diagonal frequency shift is proportional to the square of the mutual inductance to the coupler. Developing an asymmetric coupler would, in principle, allow us to shift the majority of the on-diagonal shift on the qubit mode, removing the frequency shift on the resonator mode.

The mutual inductance between a wire and a loop is given by:

$$M = \frac{\Phi}{I} \tag{5.22}$$

where Φ is the total flux in the loop ($\int \vec{B} \cdot d\vec{A}$) and I is the current flowing through the wire. The mutual inductance for a rectangular loop with length l (parallel to the wire) and

width h (away from the wire) placed a distance d from the wire is:

$$M = \frac{\int_d^{d+h} \frac{\mu_0 I}{2\pi x} l dx}{I} = \frac{\mu_0 l}{2\pi} \ln\left(\frac{d+h}{d}\right) \tag{5.23}$$

showing a linear dependence on the length of the loop parallel to the wire and only a logarithmic dependence on the height of the loop. By using a trapezoidal loop such that the section parallel to the cavity side is shorter (l_1)than the section parallel to the qubit side (l_2), we can achieve different mutual inductances to the coupler:

$$M_{cc} = \frac{\mu_0}{2\pi} \int_{d_1}^{d_1+h} \frac{1}{x} \left(l_1 + (x-d)\frac{l_2 - l_1}{h}\right) dx \tag{5.24}$$

$$= \frac{\mu_0}{2\pi} \left(\left[l_1 + \frac{d_1(l_1 - l_2)}{h}\right] \ln\left(\frac{d_1+h}{d_1}\right) + (l_2 - l_1)\right) \tag{5.25}$$

$$M_{qc} = \frac{\mu_0}{2\pi} \int_{d_2}^{d_2+h} \frac{1}{x} \left(l_2 + (x-d)\frac{l_1 - l_2}{h}\right) dx \tag{5.26}$$

$$= \frac{\mu_0}{2\pi} \left(\left[l_2 + \frac{d_2(l_2 - l_1)}{h}\right] \ln\left(\frac{d_2+h}{d_2}\right) + (l_1 - l_2)\right) \tag{5.27}$$

where M_{cc}, M_{qc} are the mutual inductances from the coupler to the cavity and qubit respectively. By adjusting the ratio of the trapezoid lengths (l_1, l_2) and the relative distance from the coupler to the cavity/ qubit large asymmetries in mutual inductances can be achieved as shown in Fig. 5.4(b) where the ratio M_{qc}/M_{cc} is plotted as a function of the length and distance ratios.

5.2.1 Circuit Simulation

The circuit model was simulated in Microwave Office to verify the analytical expressions derived above. In particular, this simulation was used to compute the coupling strength and on-diagonal shift due to the coupler for both the symmetric and asymmetric mutual inductance cases. The circuit modeled is presented in Fig. 5.5 with the addition of coupling capacitors to measurement ports to extract the resonant frequency. Three different cases are

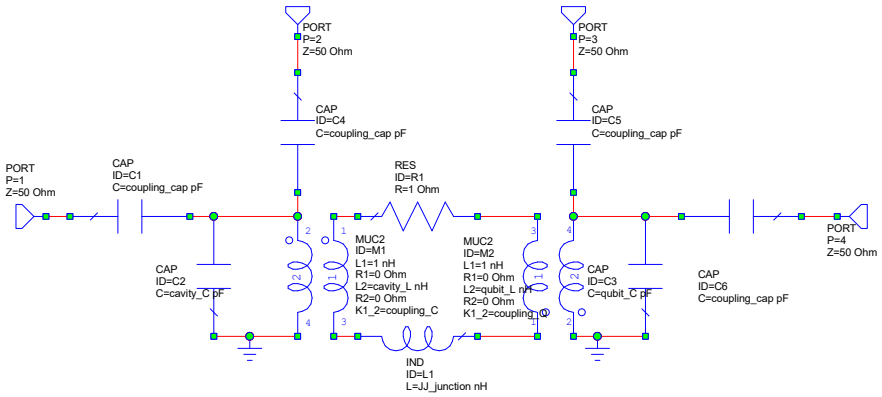

Figure 5.5: Circuit model used for Microwave Office simulations. The coupling capacitors allow for an S_{21} type measurement to find the resonant frequency of the LC circuits. A resistor is added to the inductive loop to avoid a fully inductive loop which can cause issues in the simulation.

presented: symmetric coupling on resonance (Fig. 5.7), asymmetric coupling with the *bare* frequencies on resonance (Fig. 5.8(b)), and asymmetric coupling with the *shifted* frequencies on resonance (Fig. 5.8(f)).

Fig. 5.7 shows the coupling strength as the Josephson junction inductance is varied in the coupling loop. As expected, the coupling strength is proportional to $1/L_J$ with an offset set by the linear inductance (in this circuit model, this stems from the inductance in the mutual inductance loops, a total of 2 nH). A comparison between the extracted splitting and the splitting given by Eq. 5.20 is shown in Fig. 5.6 showing good agreement with a small offset at large L_j values due to imperfect matching of the circuit parameters.

Fig. 5.8 (a) and (b) show the modes for symmetric and antisymmetric coupling respectively. In both cases, the total coupling is kept constant however, in the asymmetric case (panel (b)), the coupling from the qubit to the coupler is 4 times larger than the coupling from the cavity to the coupler. Panels (c) and (d) show the on-diagonal frequency shifts due to the coupling loop loading the main resonant circuits extracted by turning the coupling to the second circuit off (i.e. to get the on-diagonal shift for the qubit, the coupling from the cavity to the coupler is set to 0 and vice-versa). As expected, the large coupling

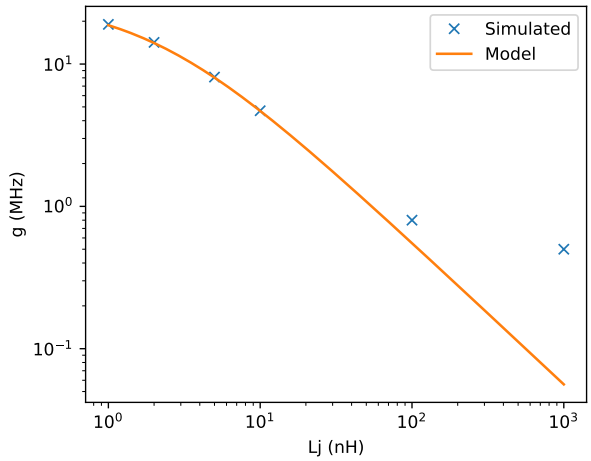

Figure 5.6: Coupling between the cavity and qubit as a function of the loop inductance (with "simulated" data extracted from Fig. 5.7). The model curve is generated from the expression in Eq. 5.20 showing good agreement until small constant frequency offsets come in.

asymmetry allows for most of the on-diagonal shift to occur on the qubit instead of the cavity. Finally, the on-diagonal shift on the qubit is corrected by flux biasing it (in the circuit model this amounts to changing the qubit frequency) so that it is back on resonance with the cavity before turning the coupler back on (panel (f)) showing a symmetric split of the resonant modes with the same coupling strength as the symmetric coupler circuit. By comparing the shifted frequencies (from panels (c) and (d)) to the un-corrected coupled frequencies (panel (b)) we can calculate the coupling between the modes as g^2/Δ where Δ is the detuning between the shifted modes and g is extracted from the symmetric case.

On a technical note, the mutual inductance is implemented in Microwave Office as a pair of inductors with a coupling coefficient (K_{ij}) from which the mutual inductance M_j can be calculated as: $M_j = K_{ij}\sqrt{L_i * L_j}$.

5.3 Experimental Results and Discussion

Here we will briefly discuss the experimental results from the first prototype with symmetric coupling. This prototype (CAD shown in Fig. 5.3) consisted of a tunable transmon, the

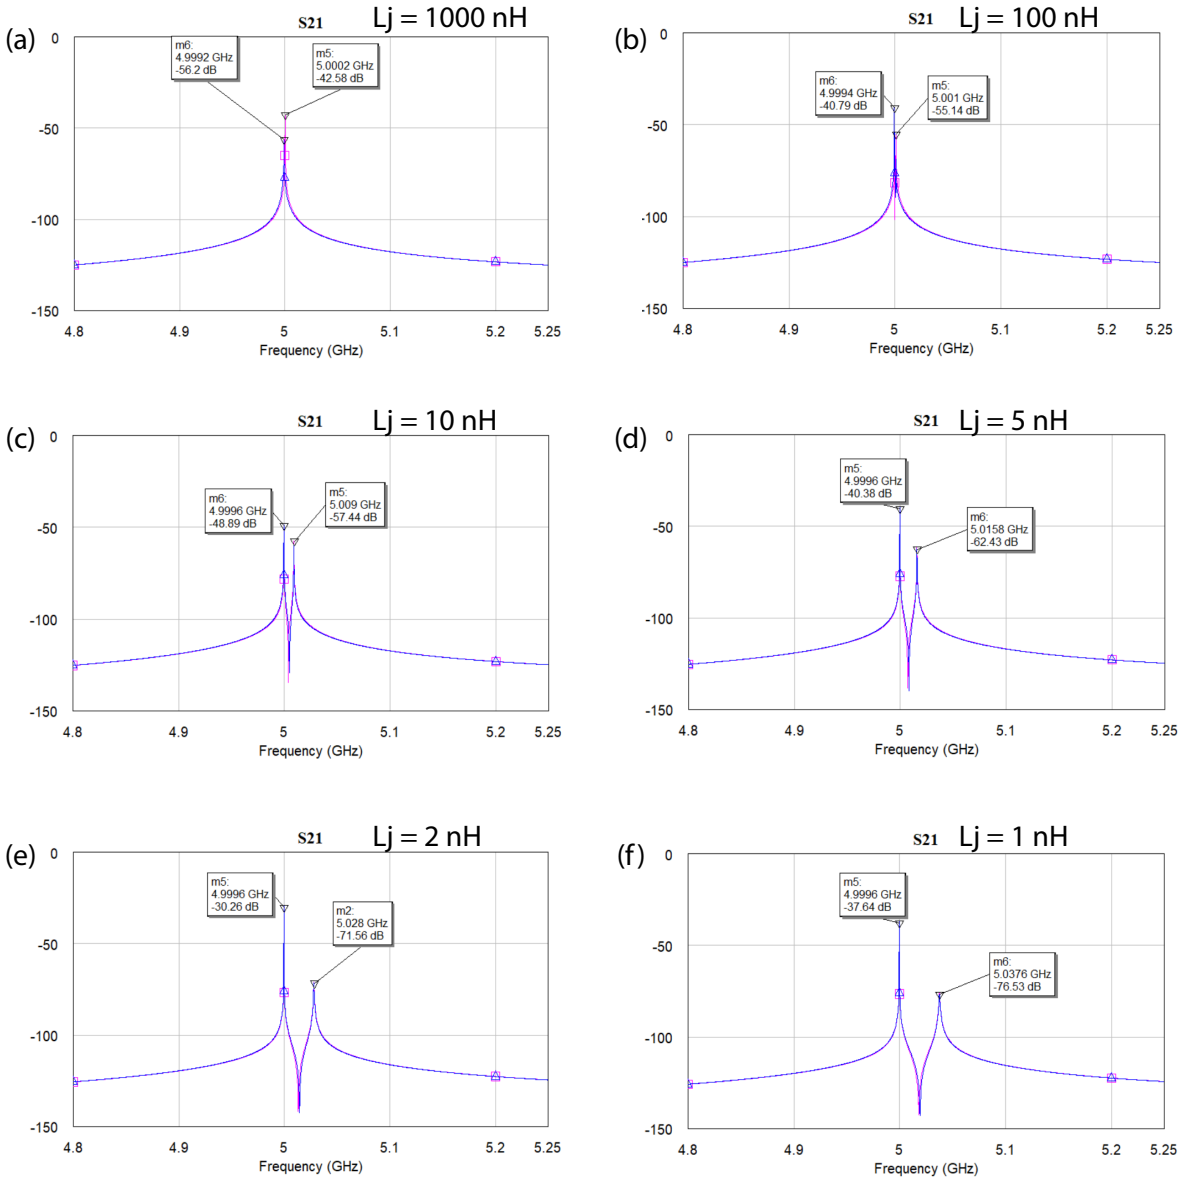

Figure 5.7: S_{21} simulated curves from the circuit displayed in Fig. 5.5, both qubit and cavity modes are coupled in a "transmission" style which displays a band pass filter response. The inductance in the loop is varied to display the variable tuning from effectively "off" (panel (a)) to strongly coupled (panel (d)).

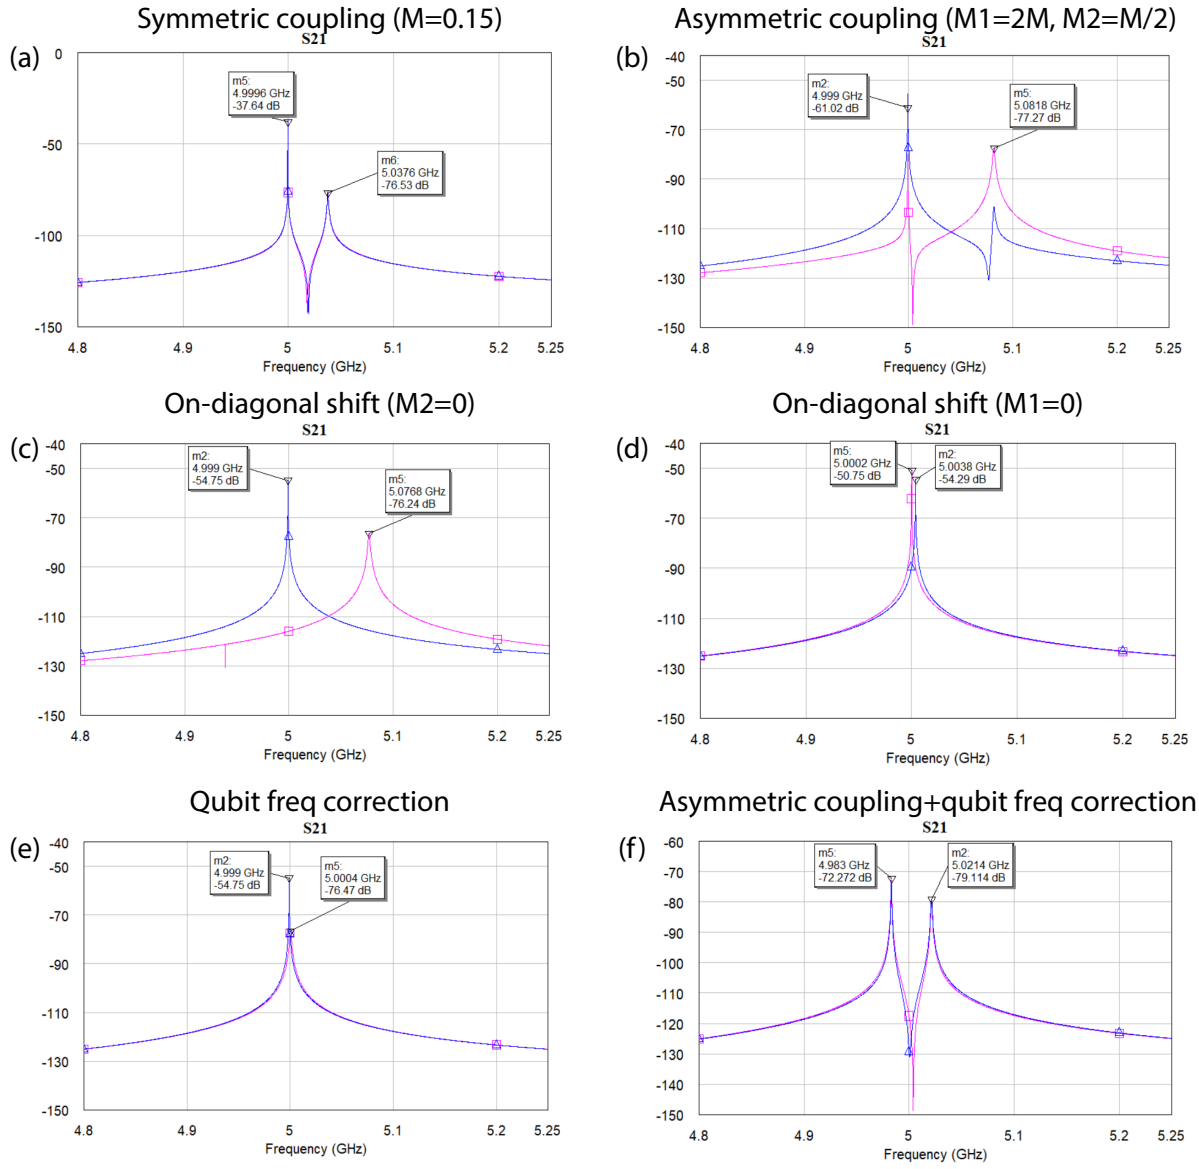

Figure 5.8: Comparing symmetric and anti-symmetric coupling schemes. (a)-(b) coupled modes using equal vs asymmetric coupling strengths for 18 MHz coupling (corresponding to a 0.238 nH mutual inductance M_i with a 1.59 nH circuit inductance L_i). In the asymmetric case, $M_1 = 0.477$ nH and $M_2 = 0.119$ nH. (c)-(d) Measurement of the on-diagonal shifts for the cavity (blue) and qubit modes (pink). (e)-(f) After compensating the qubit on-diagonal shift (e), the effective coupling between the modes is symmetric and equal to the symmetric coupling scheme (f)

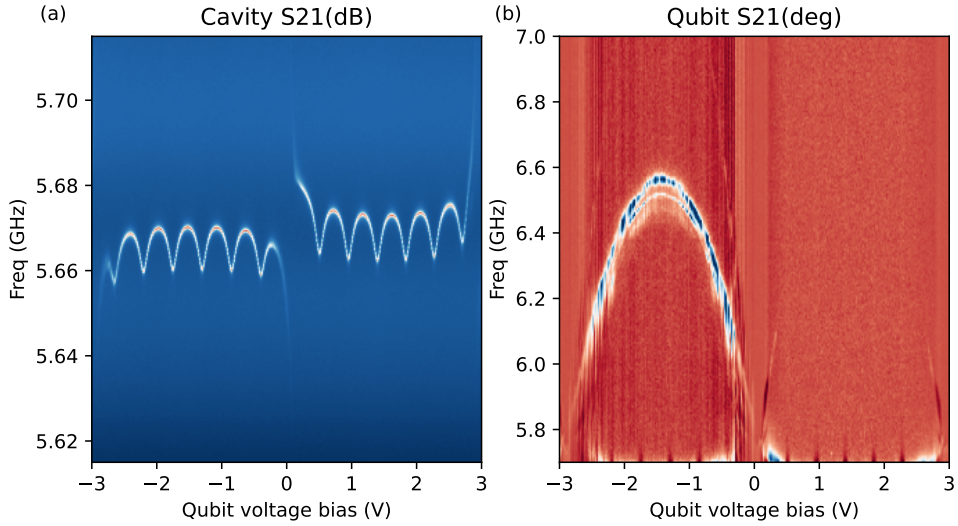

Figure 5.9: (a) Cavity spectroscopy showing a fast periodic frequency modulation due to coupler flux changing rapidly as well as a slower modulation as the qubit frequency changes and goes through resonance near 0 applied voltage. (b) Qubit spectroscopy showing a maximum frequency of 6.6 GHz.

inductive coupler loop, and a quarter-wave CPW resonator (with a duplicate on the same chip for redundancy). Both qubits used a symmetric SQUID loop (for maximum tunability) with $E_{J,max} = 20$ GHz and $E_c = 240$ MHz, leading to a maximum qubit frequency of 6.6 GHz. Both resonators were placed below the qubit frequency at 4.75 GHz and 5.6 GHz respectively so that the qubit crossing the cavity would lead to an avoided crossing which could be used to extract the coupling strength. The coupling loop was a rectangular trace 2 μm wide, 830 μm long, and 230 μm tall. The loop was placed 30 μm away from the cavity and 20 μm from the qubit trace, leading to a small asymmetry in the coupling strength. Finally, the coupler used the same sizing for the SQUID for the tunable inductance as the qubit giving $L_j \approx 8$ nH. Both the qubit and coupler were biased with a SRS DC205 voltage supply with a 1 kHz low-pass filter.

Our first measurement is a flux qubit spectroscopy scan to extract the volt to flux conversion on the qubit as shown in Fig. 5.9. The qubit tunes from below 4.5 GHz to 6.6 GHz and crosses the cavity leading to a clear avoided crossing. However, in addition to biasing the qubit, the qubit FBL line also couples to the inductive loop leading to a

	Volts/flux (qubit FBL)	Volts/flux (coupler FBL)	Offset
Qubit	5.437	-2000	-1.467
Coupler	-0.42	0.393	0.143

Table 5.1: Crosstalk calibration

large crosstalk visible as the fast modulation of the cavity frequency as the qubit is tuned through resonance.

5.3.1 Crosstalk correction

To properly decouple the coupler and qubit applied fluxes we need to calibrate out the crosstalk between the two flux bias lines. In general, we can represent this crosstalk as a linear transformation from applied voltages to fluxes on the chip:

$$\begin{pmatrix} \phi_q \\ \phi_c \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot \begin{pmatrix} V_q - V_{q0} \\ V_c - V_{c0} \end{pmatrix} \tag{5.28}$$

where ϕ_q (ϕ_c) is the effective flux on the qubit (coupler), a (d) is the volt to flux conversion for the direct qubit (coupler) modulation, b (c) is the crosstalk term, and V_q (V_c) it the applied voltage on the qubit (coupler) FBL with offset V_{q0} (V_{c0}). These coefficients can be extracted by extracting the periodicity in the frequency of the cavity or qubit as a function of applied voltage bias to either the coupler or qubit flux line as shown in the panels of Fig. 5.10. The extracted calibration parameters are shown in Table 5.1 where the *inverse* of the entry corresponds to a, b, c, or d. With the full crosstalk calibration in hand, we can now sweep the qubit and coupler fluxes independently using the inverse of Eq. 5.28 (since we are specifying fluxes and calculating the required applied voltage).

5.3.2 Tunable coupling?

To extract the coupling strength between the qubit and cavity, we pick a value for the coupler flux where 0 corresponds nominally to maximum coupling and 0.5 corresponds

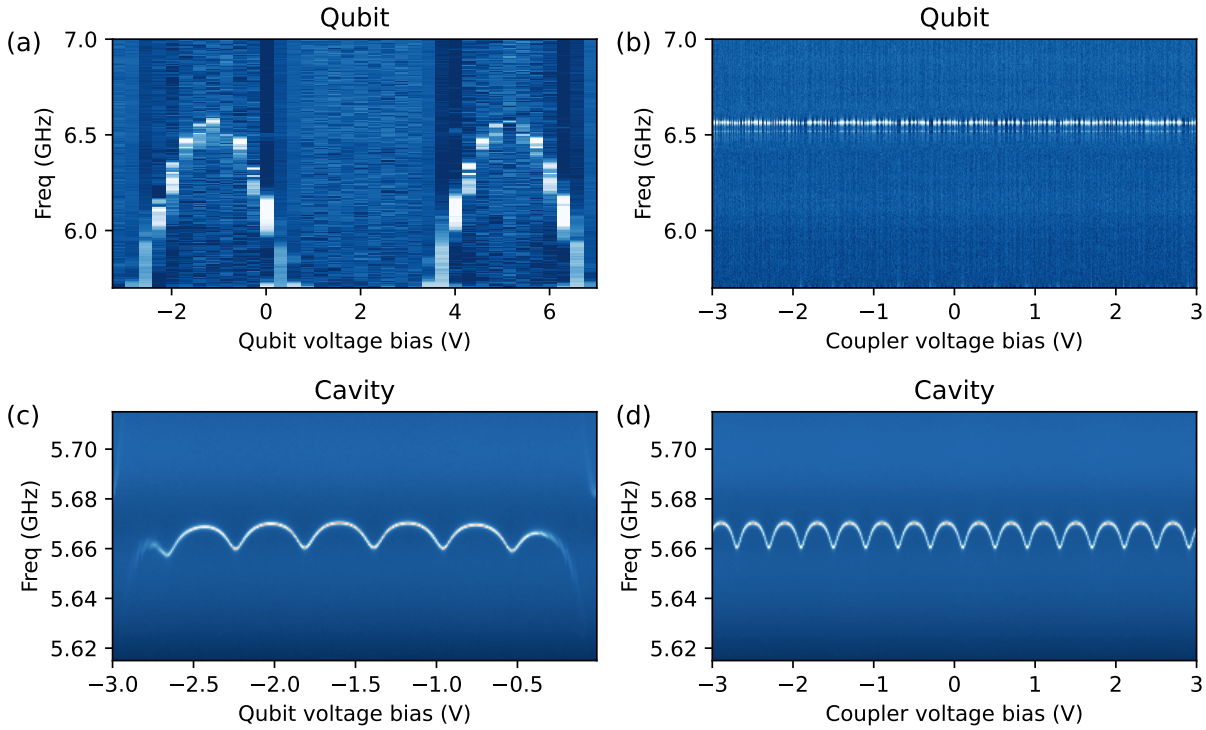

Figure 5.10: Qubit-coupler flux calibration. The diagonal panels (a), (d) show the intended tuning of the qubit and coupler using the qubit and coupler flux bias lines respectively. The off-diagonal panels show the crosstalk between the flux bias lines indicating a large crosstalk from the qubit FBL to the coupler but effectively no inverse crosstalk.

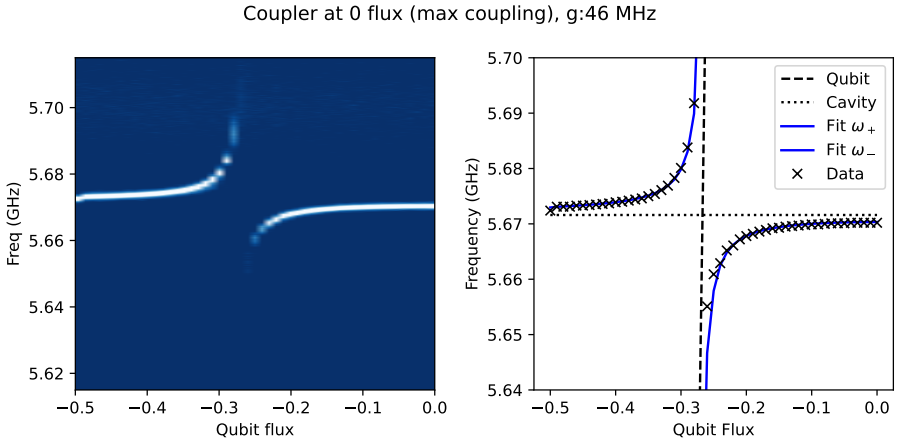

Figure 5.11: Avoided crossing between the qubit and cavity. (a) Typical S21 data showing the break in the cavity transmission as the qubit passes through resonance. (b) Fit to cavity frequency for the upper and lower branch using the eigenvalues of the Hamiltonian defined in Eq. 5.29.

to no coupling (recall, the junction inductance is proportional to $1/\cos(\phi\pi)$ when ϕ is expressed in units of flux quanta). The qubit is then ramped through resonance with the cavity and the width of the avoided crossing gives the coupling strength. We fit a simple model to the data to extract a coupling strength using the standard coupled oscillator Hamiltonian:

$$H = \begin{pmatrix} f_Q & g \\ g & f_C \end{pmatrix} \tag{5.29}$$

where f_Q, f_C are the *bare* qubit and cavity frequencies respectively and g is the coupling strength. The qubit frequency is fit to a modified tunable transmon incorporating the additional linear inductance arising from the long trace required for the mutual inductance:

$$E_J(\phi) = \frac{1}{\frac{L_J}{\cos(\phi)} + L_1} \tag{5.30}$$

$$\omega_q = \sqrt{8E_cE_J(\phi)} - E_c \tag{5.31}$$

where L_J is the effective inductance of the transmon SQUID loop and L_1 is the linear inductance from the superconducting trace. This effectively reduces the strength of the non-linearity, "flattening" the qubit spectrum near integer flux as shown in Fig. 5.12. The bare cavity frequency is extracted at each coupler flux setting, to include the on-diagonal shift due to the inductive loading on the resonator, as the mean resonant frequency. Fig. 5.11 shows an example fit to the avoided crossing data with the bare and coupled states for reference.

Fig. 5.13 shows the cavity frequency as the qubit is ramped through for different values of the coupler strength indicating strong coupling, around 40-45 MHz, for all values of the coupler flux with a small variation near $\pm$ 0.5, dipping down to 38 MHz. The small tuning range could arise from two sources: a large linear inductance in the coupling loop which would reduce the maximum tuning value (recall, $g \propto \frac{g_0}{L_1+L_J}$ so a large L_1 will suppress the achievable g), and large, static capacitive coupling to the resonator which would set

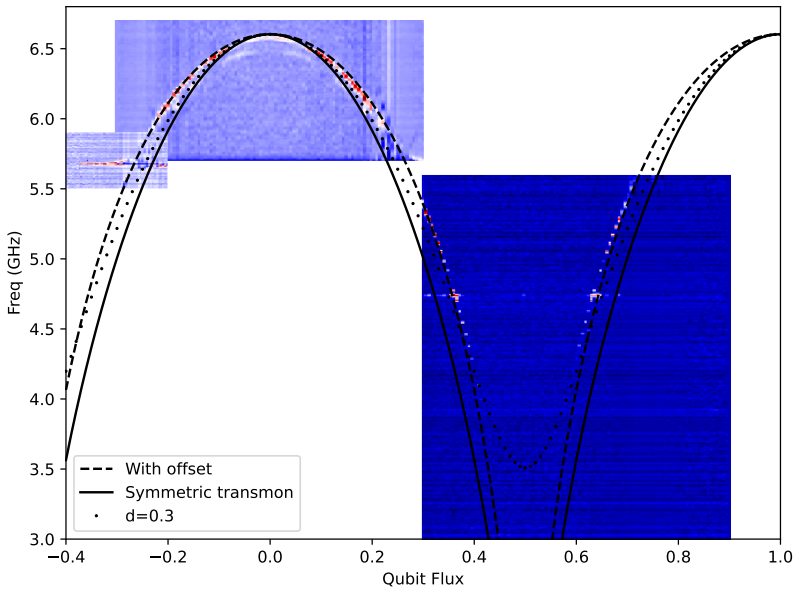

Figure 5.12: Qubit spectroscopy and fits. The solid line and dotted lines use the standard tunable transmon frequency (with the dotted line using a junction asymmetry d of 0.3 to attempt to fit the data). The dashed line incorporates the effect of the additional linear inductance on the qubit spectrum according to Eq. 5.30. For the first two fits, $E_J = 24$ GHz, $E_C = 240$ MHz; in the modified fit, $E_J = 35$ GHz and $E_L = 77$ GHz corresponding to 4.65 nH and 2.1 nH respectively

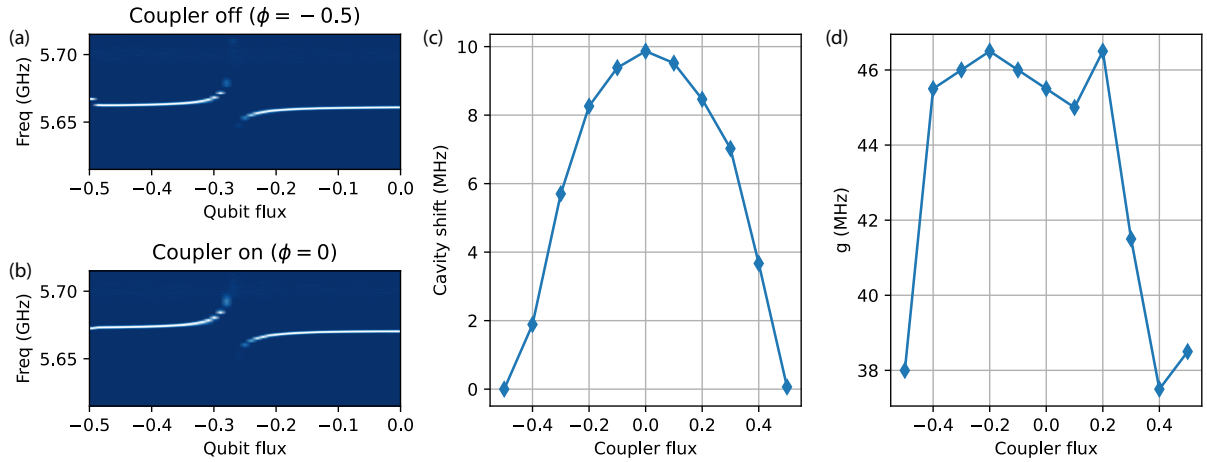

Figure 5.13: (a),(b) Avoided crossing for the coupler in the nominally "off" (flux=0.5) and "on" positions. (c) On-diagonal shift on the cavity due to the inductive loading on the resonator. (d) Extracted coupling strength g as a function of coupler flux bias.

the minimum coupling strength between the qubit and resonator.

This capacitive coupling stems from the extended nature of the coupler relative to the wavelength of the modes: while in the lumped element approximation for the coupler the capacitive coupling should be zero due to both the qubit location near a voltage node *and* the orientation of the qubit paddles perpendicular to the voltage drop from the center pin, a more realistic model for the electric field in the resonator provides a possible source for residual coupling. The mode profile of the cavity resonance is shown in Fig. 5.14 with the location of the capacitor paddles highlighted. The large spacing between the paddles of almost 1 mm to achieve the desired mutual inductance with the loop leads to an approximately 15% difference in the voltage on the center pin seen by the capacitor paddles.

5.3.3 Conclusions

While the prototype developed in this project did not show the expected dynamic range of coupling due to residual linear inductance damping out the maximum coupling strength along with large residual capacitive coupling, the methods outlined in this chapter should

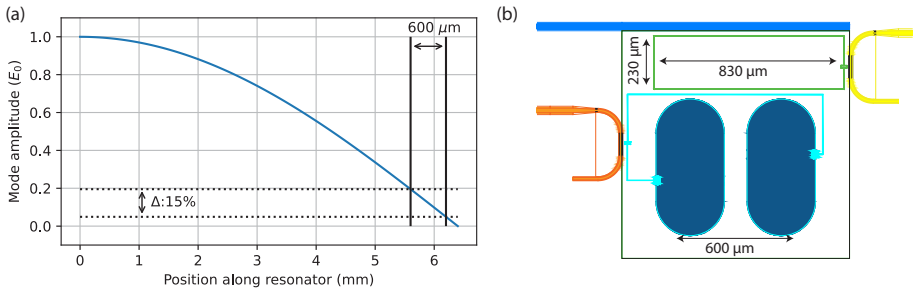

Figure 5.14: Effect of extended coupling site. (a) Field mode of the cavity resonance as a function of the location along the resonator. The location of the qubit capacitor pads are labeled showing a 15% difference of full scale in the voltage at each location. (b) Annotated CAD for reference in (a)

provide a good starting point for a next iteration. In particular, by making the traces in both the qubit and coupling elements *wider*, the linear inductance can be reduced which should improve the dynamic range of the coupler as well as remove the linear correction to the qubit inductance (which also reduces the anharmonicity [41]). A different qubit design, such as Xmon-style qubits with a singular capacitor paddle [69], should also reduce the capacitive coupling to the resonator as the capacitor paddle can be placed further away from the resonator with just the inductive short to ground (through the Josephson junction) in proximity to the cavity and the coupler loop.

For a successful implementation in a full photonic lattice. the physical footprint of the coupler would have to be reduced. In particular, the asymmetric coupling scheme presented in 5.2 requires a large separation of scales to achieve appreciable differences in the mutual inductance while maintaining the total coupling strength.

In the process of developing this tunable coupler prototype, we also found that it generates a convenient form of tunable resonator with minimal effects on the quality factor. Since the coupling loop is physically separate from the resonator trace, imperfections in the junction (such as increased loss at the interfaces and low power handling) are not transferred to the resonator. Fig. 5.15 shows the extracted cavity frequency and quality factor as a function of applied bias indicating a 10 MHz tuning range.

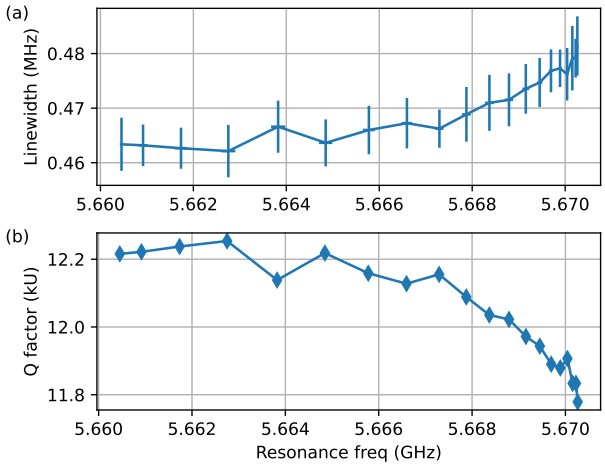

Figure 5.15: Cavity linewidth and Q-factor over tuning range

Chapter 6

Synthetic Fields for Transmons

In this chapter, we will discuss how to synthesize the effective magnetic fields used in the following two chapters. We will start by revisiting the Jaynes-Cummings Hamiltonian with an applied drive and show how to map to an effective magnetic field in a rotating frame. Then we will discuss how the coupling between the transmon and the resonator can also lead to an effective magnetic field proportional to the number of photons in the resonator.

6.1 The driven transmon

The transmon qubit, while mathematically approximated as a spin 1/2 particle, does not have an inherent magnetic moment that could be used to magnetically interact with externally applied fields. However, by abstracting away from physical magnetic fields applied with coils and instead considering the effect the different field components have on a quantum system, we can expand our definitions of a magnetic field. For instance, in atomic physics, the Z-component of the magnetic field typically sets the quantization axis and leads to energy shifts on the different atomic levels. Although for real atoms this shift is due to the magnetic interaction between the electron or nuclear spin with the applied magnetic field, in our more abstract description of a magnetic field, anything leading to a *on-diagonal*, or pure energy shift, will be considered a Z-magnetic field. On the other

hand, X or Y magnetic fields lead to hybridization between energy levels, and systems initialized in eigenstates of an applied Z-magnetic field will undergo oscillations as the applied Hamiltonian is not diagonal in the Z-basis. In the more abstract picture, fields with *off-diagonal* elements in the energy eigenbasis of the system (i.e. the eigenstates of the Z-component of the applied Hamiltonian) are considered X-fields (for "real" fields, X and Y fields correspond to a particular spatial orientation, although the choice of reference is somewhat arbitrary once the quantization axis is chosen).

Let's revisit the transmon Hamiltonian (for now uncoupled from a resonator) in the two-level approximation with an applied classical microwave drive:

$$H = \omega_q \frac{\sigma_z}{2} + \Omega \cos(\omega_d t)\sigma_x \tag{6.1}$$

where the first term describes the bare qubit levels separated by ω_q and the second term arises due to the microwave drive applied on the qubit. In particular, this form arises from the dipole coupling between the qubit and the applied electric field (recall, $\vec{d} \cdot \vec{E}$) where the field operators have been replaced with the expectation value of the classical coherent state Ω (for a detailed treatment of the transformation from the fully quantized electric field description to this semi-classical approximation, see [70]). We can immediately identify that modulating the qubit frequency (for example, by threading flux through the SQUID loop of a tunable qubit) is mathematically equivalent to an effective Z-magnetic field in the Bloch sphere of the qubit state. To determine the effect of the applied microwave drive on the qubit, we can go to the rotating frame of the applied drive:

$$U = e^{i\omega_d t} \tag{6.2}$$

$$H' = UHU^\dagger + i\hbar \frac{\partial U}{\partial t} U^\dagger \tag{6.3}$$

$$H' = \Delta \frac{\sigma_z}{2} + \Omega \frac{\sigma_x}{2} \tag{6.4}$$

where $\Delta = \omega_q - \omega_d$ and we have dropped terms rotating at $2\omega_d$ and $\omega_d + \omega_q$ as the rapid oscillations average to zero (see Chap. 2 for a more detailed derivation). In this new frame, when the applied microwave drive is on resonance with the qubit (i.e. $\Delta = 0$), the effect of the microwave drive is again mathematically equivalent to a magnetic field in the X-direction acting on the Bloch sphere of the qubit. Here the phase of the applied drive sets the particular axis about which the qubit evolves. By controlling Δ and Ω (i.e. the qubit-drive detuning and the drive amplitude) we can synthesize effectively arbitrary fields in the X-Z plane of the qubit Bloch sphere.

The astute reader will have noticed that for this simple system, there are actually *two* ways to generate an effective Z-magnetic field on the transmon: tuning the qubit frequency (directly changing ω_q) and detuning the applied microwave drive. Mathematically, both of these methods generate exactly the same effective field (although with opposite signs as $\Delta = \omega_q - \omega_d$); however, for the applications presented in the next two chapters, only qubit frequency modulation will be used. This arises from two main sources: the frequency-dependent loss in the microwave transmission from the generator to the qubit due to variable line attenuation, the drive resonator line shape, variable gain in the amplifiers, etc. makes the X and Z components of the field no longer independently tunable (as Ω now depends on the applied drive frequency); additionally, this symmetry between qubit and drive detuning is only applicable when they are the only two components on the device. For example, in Chapter 7, a secondary cavity coupled to the qubit breaks this ambiguity: if the drive frequency is modulated, the detuning between the qubit and the secondary cavity is time-independent, whereas for qubit frequency modulation, this detuning acquires an explicit time dependence.

The fields considered in the following chapters are all circularly polarized magnetic

fields of the form:

$$\vec{B} = \begin{pmatrix} B_x \cos(\omega_{mod} t) \\ 0 \\ B_z \sin(\omega_{mod} t) \end{pmatrix} \tag{6.5}$$

where B_x, B_z are the magnitudes of the applied field and ω_{mod} is the rotation rate of the field. For slowly rotating fields ($|B| \gg \omega_{mod}$), the qubit state will adiabatically follow the applied field. We study the adiabatic following regime in more detail in Chapter 7 showing the breakdown in adiabatic following at small applied fields. In the adiabatic following regime, there are two new eigenstates of the time-dependent Hamiltonian: $|+\rangle, |-\rangle$ with energies $\pm|B|/2$ corresponding to the state *aligned* or *anti-aligned* with the applied magnetic field. In Chapter 8, we show how coupling one of these states to a lossy environment can lead to the indefinite stabilization of the other state.

6.2 Rotating field protocol

In this section, we will describe the explicit implementation of a rotating field on a tunable transmon qubit used for the following two chapters. In particular, we will discuss how to properly initialize the state into an eigenstate of the field and how to measure the final state in both the bare qubit basis and in the eigenbasis of the rotating field.

6.2.1 Initialization

To guarantee good adiabatic following of the rotating field, the initialization of the system is critical. In particular, misalignment between the applied field and the state will lead to Larmor precession about the applied field. Intuitively, we can think of the magnetic field as applying a torque on the qubit state: when the state is perpendicular to the applied field (for example, initialized in $|g\rangle$ with a field applied on the X axis), the qubit undergoes full contrast Rabi oscillations; in contrast, a state perfectly aligned with the magnetic field

will not oscillate and will simply remain in the state. This is typically generalized in the time dynamics of a state under a Hamiltonian:

$$\frac{\partial \rho}{\partial t} = -i[H, \rho] \tag{6.6}$$

where ρ is the density matrix and H is the applied Hamiltonian. If the state *commutes* with the Hamiltonian, it will not pick up any time dependence (outside of decoherence which is not modeled here).

For our rotating fields, we initialize the state in three steps: (1) detune the qubit from the working frequency (typically -200 MHz) which effectively acts as a large negative Z magnetic field for which the $|g\rangle$ state is an eigenstate, (2) ramp up the microwave drive (at the working frequency) which slightly rotates the field in the XZ plane by an angle $\tan(\theta) = \frac{B_x}{B_z}$, (3) ramp the qubit back into resonance which rotates the qubit state from $|g\rangle$ to the equator of the Bloch sphere. The advantage of this preparation scheme is that the qubit initially starts in a Z eigenstate ($|g\rangle$ typically but $|e\rangle$ can be prepared by simply adding a π-pulse) allowing us to ramp a large Z-magnetic field without affecting the state. This large field then suppresses the effect of the X field turning on which keeps the qubit aligned with the applied field. By ramping the Z field back to 0 slowly, the angle between the state and the qubit remains small which simply drags the state along instead of introducing additional nutation. Either eigenstate (aligned or anti-aligned) of the field can be prepared with the addition of an optional π-pulse at the beginning of the initialization sequence.

6.2.2 Rotating field

Once the qubit is initialized, the field is then rotated by sinusoidally varying the amplitude of the microwave drive and the qubit frequency with a 90-degree offset (as shown in Eq. 6.5). The dynamics of the rotating field are mapped out by varying the length of the applied drive before performing readout either in the bare qubit basis (for σ_z, σ_x measurements) or

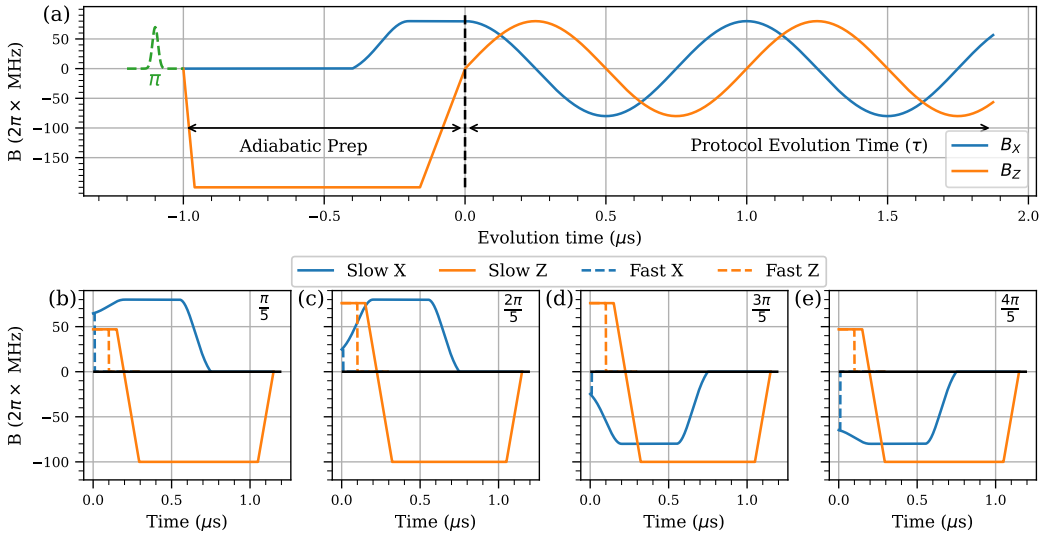

Figure 6.1: (a) State preparation and Hamiltonian evolution sections of the protocol. An optional π pulse followed by a ramp to a large Z-field (orange trace) adiabatically prepares the qubit in one of the two eigenstates of a large Z-field. Increasing the strength of the X-field (blue trace) before bringing the Z-field to zero allows the qubit state to be adiabatically prepared along the equator of the Bloch sphere. To map out the effect of the rotating magnetic field, we evolve under the driving Hamiltonian for a variable amount of time τ before ramping out for measurement. Sample $B(t)$ for ramping out at different points in the drive cycle are shown in (b)-(e). In the slow ramp out protocol (solid lines), the X-field is ramped to its maximum value before the Z-field is ramped down to a fixed reference point. This particular order is chosen to minimize spurious diabatic transitions induced by the cavity as the qubit goes through resonance. This scheme adiabatically rotates the qubit and maps the population in the instantaneous basis at time τ to the population in the z basis. The fast ramp out (dashed lines) consists of a fast shutoff of the X-field while holding the Z-field constant before shutting off the Z-field. The buffer between X and Z shutoffs mitigates the effect of distortions and delays from the 5 MHz low pass filter on the flux-bias line. This filter also reduces the maximum rate of change of the Z-field leading to only an approximate diabatic/instantaneous measurement of the population in the z basis at $t = \tau$. Figure reproduced from the Appendix of [71]

in the eigenbasis of the rotating field (effectively measuring $|+\rangle, |-\rangle$).

6.2.3 Measurement

In this section we'll describe the two measurement schemes used to characterize the dynamics of the rotating field which are referred to as the fast ramp out measurement (which effectively measures in the qubit basis) and the slow ramp out measurement (effectively measuring in the instantaneous basis of the rotating field).

The fast ramp out measurement consists of a rapid snap-off of the applied field before measuring the qubit state. In this basis, the measurements will oscillate at ω_{mod} as the qubit rotates around on the Bloch sphere. In the ideal case, this would be a perfectly diabatic shut-off and the qubit state would simply freeze in the state aligned with the field directly before the shutoff. By measuring σ_z and σ_x we can then verify the time evolution of the state as it follows the applied field. However, due to the low-pass filtering on the flux bias line to reduce flux noise reaching the qubit, the Z-component of the field has a finite response time reducing the speed at which the field can be shut-off. To avoid accidentally rotating the state while the fields shut off, we keep the Z-field on for a buffer time of 100 ns while the X-field is turned off. This guarantees that the X field is completely turned off before the Z field is ramped to 0. While the additional buffer will lead to precession of the qubit state before measurement, the *projection* of the state onto the z-axis will remain unchanged. For measurements in the X-basis, the buffer is applied on the X drive instead while the Z drive is ramped off to apply precession about the X axis instead which again leaves the *projection* of the state unchanged. Some example traces for the fast ramp out measurement (and σ_z measurements) are shown in panels (b)-(e) of Fig. 6.1

On the other hand, the slow ramp out measurement performs a measurement in the rotating basis of the applied field. In this basis, the oscillatory dynamics vanish and only the decoherence of the states is visible. To perform this measurement, we rotate the qubit state back to a fixed point on the Bloch sphere before measuring the state. In our case,

we ramp the field back to pointing towards the $|g\rangle$ state on the Bloch sphere. Assuming that the rotation back to the measurement does not introduce spurious transitions, this effectively maps the $|-\rangle, |+\rangle$ basis of the rotating field onto the bare qubit basis $|g\rangle, |e\rangle$. To reliably rotate the state back to the measurement point, we first ramp the X-field to its maximal value before ramping the Z field to a large negative detuning. The X-drive is then turned off before ramping the qubit back to the operating frequency. This sequence guarantees a large applied field on the qubit at all times during the rotation and minimizes sudden changes in direction of the field which would lead to diabatic transitions. Some example traces of the ramp out at different phases of the rotating field are shown in panels (b)-(e) of Fig. 6.1 in the solid traces.

Both measurements are equivalent and we expect that the slow ramp out measurement effectively describes the *envelope* of the fast ramp out measurement. Since the slow-ramp out protocol contains all the information about the state (the fast ramp out protocol oscillations do not convey any additional information about the state of the qubit), it is a lot more efficient to measure in this basis as we can sample the dynamics with fewer points without incurring aliasing effects. A comparison between the two measurement schemes is shown in Fig. 6.2 and the following two chapters will delve into the details of the adiabatic following regime (Chapter 7) and the stabilization of one of the eigenstates of the driving field (Chapter 8).

6.3 Effective Magnetic Fields in the JC Hamiltonian

Another source of effective magnetic fields for transmons is coupling to a resonator. If we re-examine the Jaynes-Cummings Hamiltonian:

$$H = \omega_c a^\dagger a + \omega_q \frac{\sigma_z}{2} + g(a^\dagger \sigma^- a\sigma^+) \tag{6.7}$$

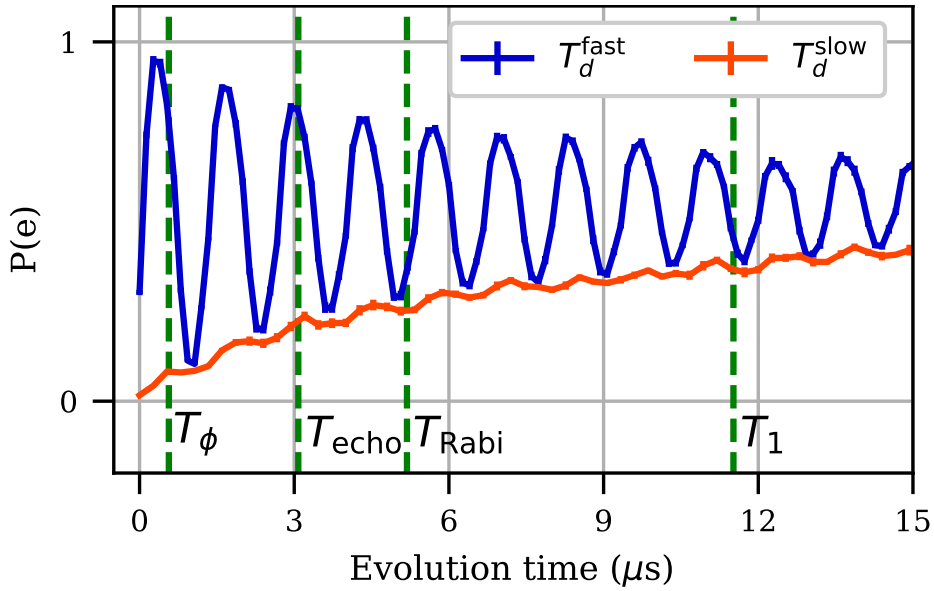

Figure 6.2: Comparing the fast and slow ramp out measurements. As expected, the fast measurement shows the time dynamics of the state with a damped oscillatory response as the qubit decays whereas the slow measurement is effectively performing a measurement in the rotating frame of the qubit, removing the oscillatory response. The qubit coherence timescales are shown in the figure indicating that the field following time exceeds the typical phase coherence times of the qubit. Figure reproduced from [53].

where $\omega_{c,q}$ are the usual cavity and qubit frequencies respectively and g is the coupling between the cavity and qubit modes. In the interaction picture ($U = e^{\omega_c a^\dagger a + \omega_q \frac{\sigma_z}{2}}$), the coupling becomes:

$$H' = g(a^\dagger \sigma^- e^{i\Delta t} + a\sigma^+ e^{-i\Delta t}) \tag{6.8}$$

where $\Delta = \omega_c - \omega_q$. In the dressed basis ($|g, n\rangle, |e, n-1\rangle$), the matrix representation of the Hamiltonian becomes:

$$H = \begin{pmatrix} 0 & g\sqrt{n}e^{i\Delta t} \\ g\sqrt{n}e^{-i\Delta t} & 0 \end{pmatrix} \tag{6.9}$$

where the $\sqrt{n}$ factor arises from the creation/ annihilation operators in the coupling term. In this representation, the coupling to the cavity acts as an effective magnetic field in the X-Y plane. Again, we can build intuition by considering the on-resonance case: if the

qubit and cavity are coupled and on resonance, they will coherently exchange excitations. A qubit initially prepared in the $|e\rangle$ state in an empty cavity will emit into the cavity driving the system to the $|g,1\rangle$ state. This photon can be reabsorbed by the qubit leading to oscillations in the qubit population at rate $g\sqrt{n}$. Just as with our discussion of synthetic fields on the transmon earlier, this interaction is equivalent to a magnetic field in the X-Y plane. For non-zero detuning values, this magnetic field rotates in the XY plane at a rate Δ.

6.4 Coherent state preparation

In this section we briefly discuss how to calibrate the size of a coherent state in a cavity using qubit spectroscopy. Recall the dispersive Hamiltonian for a qubit-cavity system:

$$H = \omega_c a^\dagger a + \omega_q \frac{\sigma_z}{2} + \chi a^\dagger a \sigma_z \tag{6.10}$$

where as usual, $\omega_{c,q}$ are the cavity and qubit frequencies, and χ is the dispersive shift. So far we have used the dispersive shift to perform qubit state readout by measuring the cavity, however, we can also use it to measure the photon number distribution in the cavity:

$$\omega_q' = \omega_q + 2\chi a^\dagger a. \tag{6.11}$$

In the photon number splitting regime ($\chi > \kappa$), the qubit spectrum will form a series of peaks split by 2χ as shown in Fig. 6.3. To extract the photon number distribution, we fit a sum of gaussians to the spectrum. The population in a given photon number is given by the ratio of the area of the gaussian to the total area under the curve:

$$p_i = \frac{\int g_i}{\sum[data]} \tag{6.12}$$

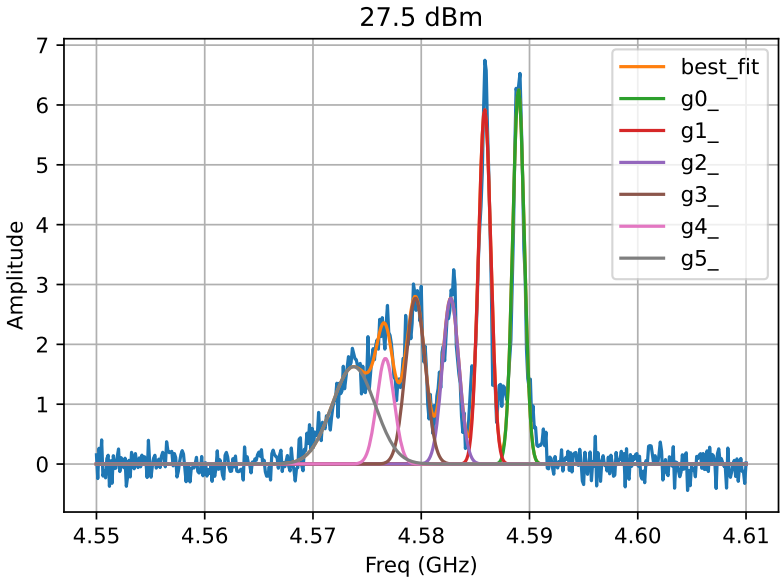

Figure 6.3: Photon number splitting in the qubit frequency. Each peak corresponds to a photon number (0 to 5 going right to left) and is separated from the next by 2χ. The relative photon distribution is calculated by computing the ratio of a given peak to the total area under the data curve.

where g_i is the gaussian for the photon number of interest. The average photon number can then be computed as the weighted sum:

$$< n >= \sum_i n_i \cdot i \tag{6.13}$$

where i is the number of photons and n_i is the population in that state.

To build a conversion factor between applied cavity power and average photon number, we sweep the power and perform qubit spectroscopy at each power as shown in Fig. 6.4(a). At each power, we compute the photon number distribution (shown in (b)) and the average photon number (c). The photon number grows linearly as a function of the applied power but starts to diverge at higher power as the peaks start to blur together making the population extraction more difficult.

For this measurement, we use a device with two cavities on it: a regular cavity to perform qubit readout (far detuned from the qubit) and another cavity near the qubit

frequency to store the coherent state. Qubit spectroscopy is acquired at low qubit power to reduce the linewidth and obtain more clearly separated peaks. For the data presented in this section, the coherent state cavity has a resonant frequency of 5.05 GHz. The qubit is tuned to 4.59 GHz to achieve a large χ shift which increases the separation between the photon number peaks.

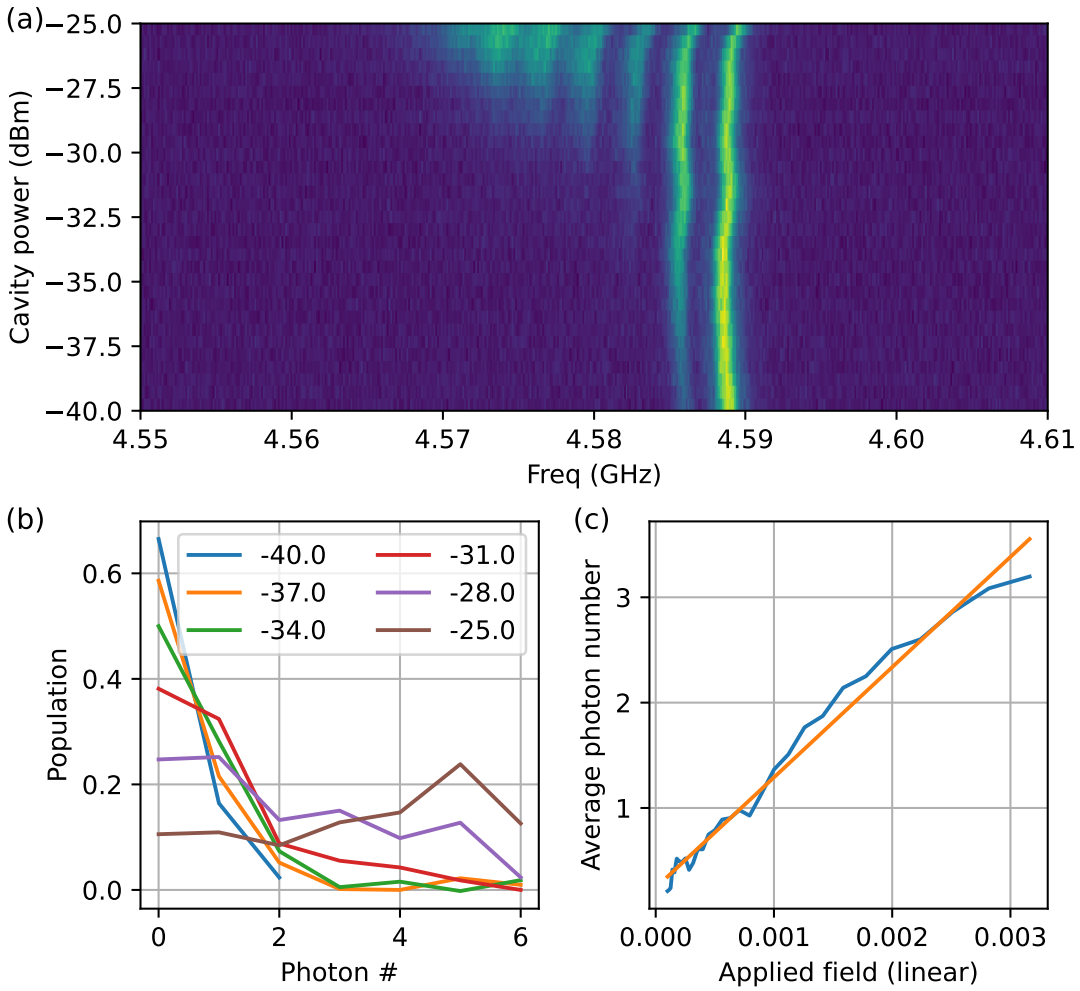

Figure 6.4: Coherent state as a function of applied cavity power. (a) Qubit spectroscopy vs applied cavity power. As the power increases, higher and higher photon number peaks appear. (b) Photon number distribution extracted from a multi-gaussian fit to the qubit spectroscopy data. At low powers (blue), only the 0 ad 1 states have significant population. As the power increases, the distribution shifts to higher photon numbers. (c) Average photon number computed from the distributions in panel (b).

Chapter 7

Topological Pumping

This chapter is broken up into two sections: a brief introduction to topological pumping through the lens of a superconducting qubits platform followed by an in preparation publication characterizing the synthetic magnetic fields required for the topological pump.

7.1 An experimentalist's introduction

7.1.1 Introduction and Motivation

Preparing the right initial state for a quantum computing protocol or sensing application with high fidelity can be an expensive operation in terms of initialization time or control sequence complexity. Of particular interest for this chapter is the preparation of large $|n\rangle$ Fock states. These highly quantum states allow for better-than-standard-quantum limit measurements [72] in a similar way as squeezed states improve the performance of sensitive measurements like those performed by LIGO [73]. Protocols have been developed utilizing the large native light-matter interaction strength achievable in circuit QED to prepare arbitrary photon states inside of cavities [74–76]. While these protocols can quickly generate high fidelity Fock states, they require precise calibration and timing with increasing precision as the prepared Fock state grows in size. Our boosting protocol utilizes a topologically

protected invariant, the Chern number, to remove sensitivity to amplitude and phase noise in the driving fields. Larger Fock states can be prepared by running the Hamiltonian for longer without requiring fine calibration of the protocol parameters.

7.1.2 Topological pumping

A qubit driven by a periodic time dependent Hamiltonian can be described using the Floquet formalism [77]. In this formalism, the time periodicity is analogous to a lattice in frequency space. The dimensionality of this new lattice can be increased by adding more drives on the qubit. The amplitude, frequency and phase of these drives relative to each other allow us to explore a rich set of physics in this synthetic space. In this space, hopping between sites corresponds to the absorption/ emission of a photon of a particular drive with the hopping strength set by the coupling between the driven qubit and drive. The lattice sites are n-level systems, set by the number of levels used in the qubit, living on a (n_1, n_2) grid indicating the number of photons from each drive. In particular, a well studied system is a spin driven by two incommensurate drives [78, 79]. The resulting 2D lattice exhibits a topological band structure when the hopping rate (i.e. the qubit-drive coupling) exceeds the on-site energy (defined as the frequency of the drives). This regime is similar to the half Bernevig–Hughes–Zhang (BHZ) model [80]:

$$H = \vec{h} \cdot \vec{\sigma} \text{ with: } \vec{h} = \begin{pmatrix} A_x \cos(k_x) \\ A_y \cos(k_y) \\ \Delta + \sin(k_x) + \sin(k_y) \end{pmatrix} \tag{7.1}$$

where $\vec{\sigma}$ represents the $\sigma_{x,y,z}$ components of the two-level system and $\vec{h}$ is an effective magnetic field living in the abstract 3D space spanned by the σ matrices. One can also represent this Hamiltonian in the two-level Bloch sphere picture by expanding the dot

product:

$$H = \begin{pmatrix} \Delta + \sin(k_x) + \sin(k_y) & A_x \cos(k_x) - iA_y \cos(k_y) \\ A_x \cos(k_x) + iA_y \cos(k_y) & -\Delta + \sin(k_x) + \sin(k_y) \end{pmatrix} \tag{7.2}$$

This simple model represents a 2D lattice with two hopping strengths (A_x, A_y) and a pair of orbitals on each lattice site split by an energy Δ. For $|\Delta| < 2$, the resulting energy bands acquire a non-zero Chern number (defined as the integral of the Berry curvature over the full Brillouin zone). Analogously, we can define a Chern number for the pseudo-bands forming in the dual space of our Floquet lattice, where the phase of the drives ϕ_1, ϕ_2 take the role of momentum in the BHZ model (k_x, k_y). As the qubit evolves under the two drives, if the drives are incommensurate, it will sample the *entirety* of the ϕ_1, ϕ_2 phase space effectively integrating it out leading to behavior dictated by the Chern number. In "number" space, the local Berry curvature in the currently sampled region of the Brillouin zone maps to an effective magnetic field piercing the n_1, n_2 plane where the direction of the field is set by the sign of the Berry curvature. In addition to this magnetic field, an electric field is applied on the lattice in the $\vec{\omega} = (\omega_1, \omega_2)$ direction as a result of the drives acting on the qubit. The qubit will undergo Bloch oscillations in the $\vec{\omega}$ direction but will also accrue a displacement *perpendicular* to this vector due to the effective magnetic field piercing the plane as shown in Fig. 7.1. Over time, the qubit will accrue a displacement proportional to the average or integral of the Berry curvature i.e. the Chern number. For the topological regime, the Chern number is non-zero resulting in a net power transfer between the drives:

$$[P_{1\to 2}] = C\frac{\omega_1 \omega_2}{2\pi} \tag{7.3}$$

where C takes the values ± 1 depending on the sign of Δ. This power transfer is quantized in the time-averaged sense with a photon from ω_1 being converted to ω_2 every period of the first drive.

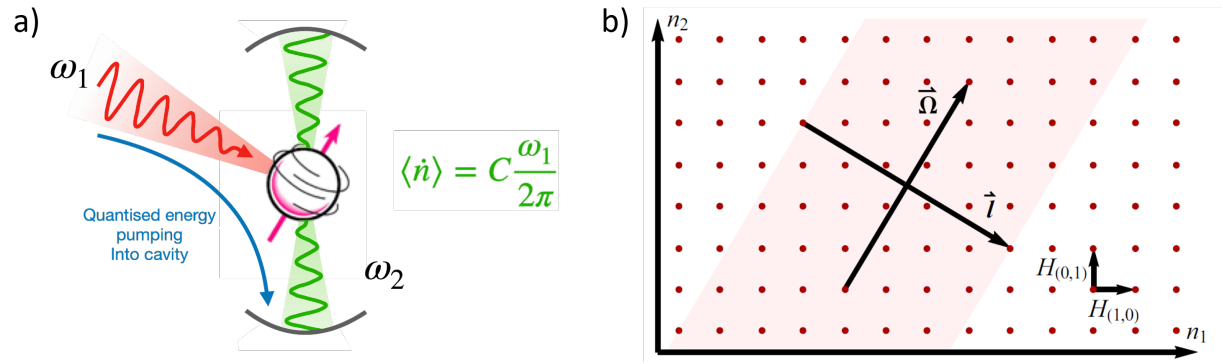

Figure 7.1: a) Schematic of protocol, the red drive pumps photons into the green cavity through the qubit, b) lattice/ Floquet representation of the driving protocol: in photon number space of the two drives, an effective electric field in the n_1, n_2 direction combined with a Berry curvature in the underlying band structure of the energy levels results in an effective "force" in a direction perpendicular to the electric field. The applied electric field $\vec{\Omega}$ results in Bloch oscillations without net transport. Fig. adapted from [78]

7.1.3 Coherent Boosting

"Boosting" in Fock space

While in general only the average pumping rate is quantized, recent work [81] has shown that there exist specific times, "quasi-periods", where the population in the pumped drive can be precisely defined as:

$$N_1 = C \times nT_2$$

where nT_2 is a multiple of the second driving period. Since this protocol acts on the number basis of the pumped drive one can define "boosting" as a translation in Fock space maintaining the quantum coherences of the initial signal:

$$|n\rangle \Rightarrow |n+m\rangle$$

With this protocol, it is possible to initialize a "small" (defined as a small $|n\rangle$) state in the cavity and pump it to a large state without relying on complicated initialization schemes. Since the boosting proceeds at the rate the the second drive, the only resource required is the appropriate evolution time to achieve the desired state. The simulation results and Wigner

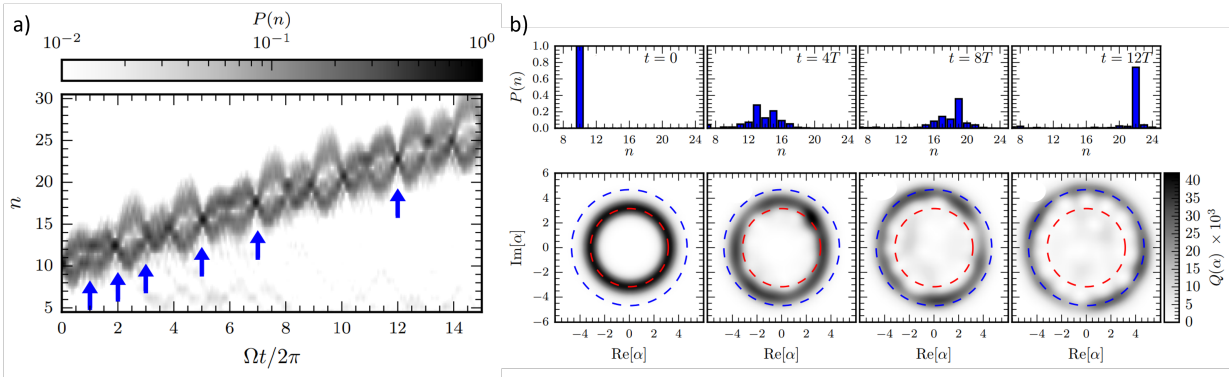

Figure 7.2: Boosting of Fock states. (a) Photon number distribution as the system is driven by Ω (the external field). At the rephasing times (blue arrows), the number distribution narrows back to effectively a single photon number state. (b) Cross-sections of the photon number distribution (top) and phase-space representation of the cavity photon state. Over 12 periods, the photon state gets shifted from n=10 to n=22 with fairly good overlap with the n=22 pure Fock state. Figure reproduced from [81]

space representation of the boosted Fock state state are shown in Fig. 7.2 illustrating the potential high fidelity of the boosting protocol.

Archimedes screw picture

A simple, experimentalist, picture of this effect is the Archimedes screw. Every revolution, the screw transfers water between reservoirs with the direction set by the "helicity" or winding of the screw (clockwise or counter clockwise). This picture was studied early on in the context of the Thouless pump [82, 83] where two electrodes are connected by a sinusoidally varying potential. Each electrode is modelled as a simple reservoir filled with electrons up to the Fermi level . As the phase of the potential winds, if the Fermi level is higher than the minimum of the potential, electrons will enter the potential minimum from the reservoir and be dragged along with the potential. Since the number of electrons in each level below the Fermi level is quantized, each pumping cycle transfers a fixed number of electrons between the reservoirs which is independent of the specifics of the level structure of the reservoirs. As long as there are no level crossings at the Fermi level, the pumping rate will stay the same making it topologically invariant.

The following sections are from [53]. Qianao (Billy) and I collected the experimental data, David Long and Anushya Chandran provided the theory support, and Alicia Kollár provided extensive experimental support and guidance. All authors contributed to the manuscript.

7.2 Driven systems with synthetic topology

The experimental toolkit developed in this work will make it possible to realize and directly observe a *topological photon pump*. In this section, we review the theoretical description of these pumps, and establish the experimental requirements that must be met in order to observe photon pumping.

Topological photon pumps (also called topological energy pumps or topological frequency converters) were first proposed in qubit systems with strong, adiabatic (slow) driving by classical fields [78]. The Hamiltonian for this system is

$$H_B = \frac{1}{2}\vec{\sigma} \cdot \vec{B}(t), \tag{7.4}$$

where $\vec{\sigma}$ is a vector of qubit Pauli operators, and $\vec{B}(t)$ is the vector of their time-dependent coefficients. We refer to $\vec{B}(t)$ as a magnetic field, regardless of whether it is implemented through the coupling of a real magnetic field to the magnetic moment of a spin-1/2 particle.

When the field $\vec{B}(t)$ is a superposition of two orthogonally oriented circularly polarized drives and a static field, the qubit can preferentially absorb energy from one driving field and emit into the other field. Anticipating the model we realize experimentally, we say the qubit absorbs from a drive with (angular) frequency $\omega_{\text{mod}} = 2\pi/T_{\text{mod}}$ and emits into a drive with frequency Δ. The average power pumped by the qubit is quantized when $\omega_{\text{mod}}/\Delta \notin \mathbb{Q}$ is irrational, even though there is no exactly resonant conversion process between quanta of energy ω_{mod} and Δ.

It was recognized in Ref. [84] that the same phenomenon occurs when one (or both)

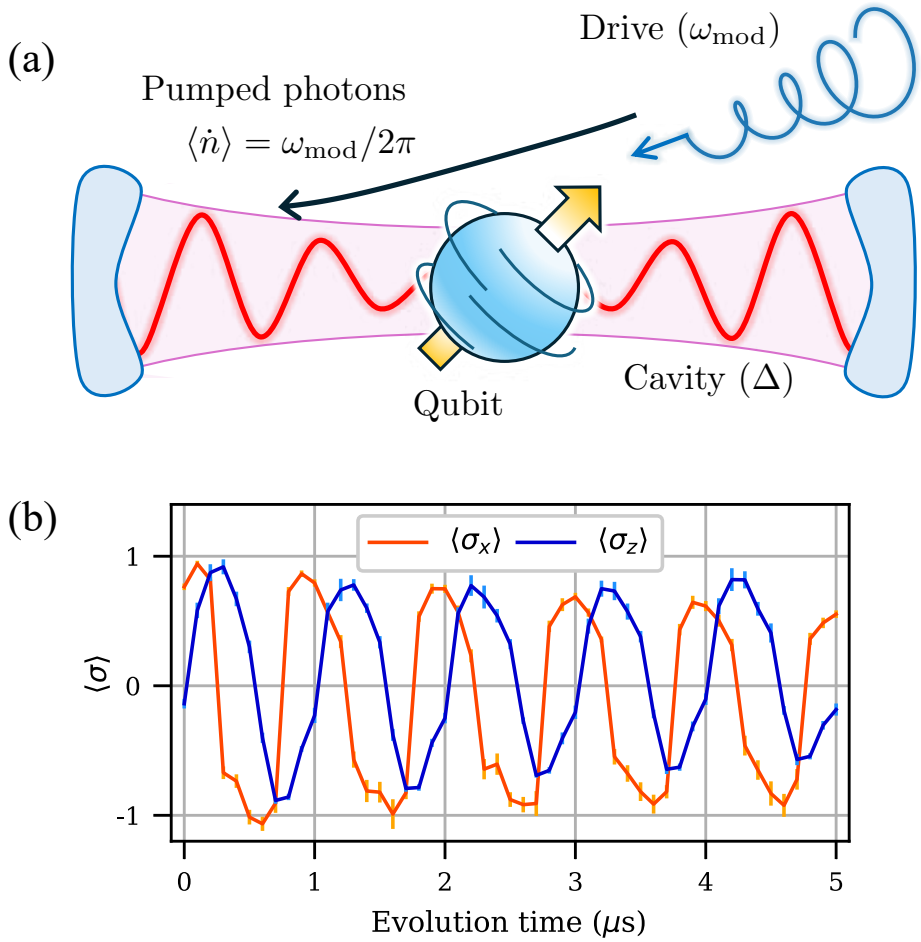

Figure 7.3: The topological frequency pump. (a) A qubit is driven with two oscillatory fields: an external drive (blue) and a cavity mode (red). In the topological regime, the qubit absorbs energy from the drive and pumps energy into the cavity. For the experimental implementation, the external drive is the combination of a static x field and circularly polarized magnetic field in the xz plane. (b) Measurement of the qubit following the rotating field from part (a). The qubit adiabatically follows the applied field as seen by the 90-degree phase shift between the x and z projections of the state vector.

of the classical drives is replaced by the quantized mode of an oscillator, and in Ref. [85] it was shown that the topological pumping mechanism can even *boost* non-classical states of a cavity. That is, an initial unentangled state of the qubit and cavity of the form

$$|\psi(0)\rangle = |s\rangle \otimes \sum_n c_n |n\rangle \tag{7.5}$$

may be subjected to a time-periodic drive on the qubit which, at special times T_N, produces a time-evolved state

$$|\psi(T_N)\rangle \approx |s\rangle \otimes \sum_n c_n |n + h_N\rangle \,. \tag{7.6}$$

Above, h_N is an integer equal to the number of periods of the drive $T_N \omega_{\rm mod}/2\pi = T_N/T_{\rm mod}$. The topological pump acts to translate the cavity state in Fock space. This can, for instance, enable the preparation of highly excited Fock states provided that Fock states with a smaller number of photons can be reliably prepared.

A Hamiltonian which achieves topological energy pumping between a classical drive and a quantum cavity with annihilation operator a is

$$H_B^+ = \frac{1}{2}\vec{\sigma}\cdot\vec{B}(t) + \Delta\, a^\dagger a + g\,(a^\dagger\sigma^- + a\sigma^+), \tag{7.7}$$

where

$$\vec{\sigma}\cdot\vec{B}(t) = B_0\{\sin(\omega_{\rm mod}t)\sigma_z + [m + \cos(\omega_{\rm mod}t)]\sigma_x\} \tag{7.8}$$

is a periodic drive on the qubit, m is a dimensionless parameter, $\sigma_{x,y,z}$ are Pauli matrices, and $\sigma^\pm = (\sigma_x \pm i\sigma_y)/2$.

There is a simple toy picture for topological pumping of Fock states with the Hamiltonian in Eq. (7.7) when we set $m = 1$ and neglect evolution of the cavity when it is not

resonant with the qubit. At the beginning of the protocol ($t = 0$) the cavity is initialized in an accessible Fock state $|n\rangle$ and the qubit is prepared in the instantaneous excited state $|\phi_+\rangle$ of $\vec{\sigma}\cdot\vec{B}(0)$, which is aligned in the x direction on the Bloch sphere if $B_0 > 0$. The qubit and cavity are far off-resonant, and the qubit adiabatically follows the rotating external field $\vec{B}(t)$ to the state aligned along $+z$, denoted by $|e\rangle$, until $t \approx T_{\text{mod}}/2$. Near $t \approx T_{\text{mod}}/2$, the external field functions as a ramp, slowly driving the qubit through an avoided level crossing with the cavity, causing an adiabatic transition from the $|e, n\rangle$ state to the state $|g, n+1\rangle$ (where $|g\rangle$ is the ground state) as the cavity absorbs a qubit excitation. At the end of this ramp, the field on the qubit is aligned along the $-z$ direction, so that the qubit is again in an excited state of $\vec{\sigma}\cdot\vec{B}$. The external field then slowly returns to its large initial value, so that the qubit again adiabatically follows the field and ends the Floquet period in an excited state of $\vec{\sigma}\cdot\vec{B}(T_{\text{mod}}) = \vec{\sigma}\cdot\vec{B}(0)$. The state of the system thus returns to its initial conditions, except that the cavity now has $n+1$ photons. The protocol is then repeated at a fixed rate $1/T_{\text{mod}}$, pumping an average of one photon per period into the cavity.

More generally, topological pumping can be understood as Thouless pumping [86, 87] in the Fock basis. This is a topological effect which requires adiabaticity of the drives and the existence of a nonvanishing Chern number C in the Fock-basis tight-binding model, which is effectively extended by an additional *synthetic dimension* due to the periodic drive [78]. In the toy picture, a non-zero C simply verifies that the external field passes close enough to zero for the qubit to exchange excitations with the cavity [$C = \pm 1$ in the model of Eq. (7.7)]. The sign of the Chern number indicates whether the state loses an excitation to the cavity, or absorbs an excitation from the cavity, so that the average photon pumping rate is

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t \langle \dot{n}(t')\rangle \mathrm{d}t' = \frac{C}{T_{\text{mod}}}. \tag{7.9}$$

In order to realize pumping, we must operate in the adiabatic regime, which implies some restrictions on experimental parameters. The frequencies ω_{mod} and Δ must be much smaller than the external field strength B_0 and the qubit-cavity coupling $g\sqrt{n}$. In the cavity

or circuit QED context, it is most natural to realize this requirement in a rotating frame, so that Δ is a detuning between a large lab frame cavity frequency, ω_c, and the frequency of the rotating frame, ω_d. That is, $\Delta = \omega_c - \omega_d$. Similarly, the mean rotating frame qubit frequency is given by $\omega_q - \omega_d$, where ω_q is the mean lab frame qubit frequency. As Eq. (7.7) has a mean qubit frequency of zero, we set $\omega_d = \omega_q$. The Hamiltonian (7.7) would then come with counter-rotating terms of frequency $\omega_q + \omega_c$, which can be ignored when this frequency is large compared to other scales. Our experiment uses a transmon qubit, which has a relatively low anharmonicity α. In order to treat the transmon as having two levels, we need that this α is also much larger than any Hamiltonian parameter in Eq. (7.7), which is a much stronger constraint than that imposed by $\omega_q + \omega_c$. In summary, operating in the adiabatic regime requires that

$$\frac{1}{T_{\text{coh}}} \ll \Delta, \omega_{\text{mod}} \ll B_0, g\sqrt{n} \ll \alpha, \omega_q + \omega_b, \tag{7.10}$$

where T_{coh} is the coherence time of the system and sets the lower limit for Δ and ω_{mod}. Furthermore, the requirement of non-vanishing Chern number imposes [88]

$$B_0^2(1-m)^2 < g^2 n < B_0^2(1+m)^2. \tag{7.11}$$

We see that the photon occupation number n cannot be too small or large for pumping to occur. We refer to the range of parameters specified by Eq. (7.11) as the *topological regime*.

In the remainder of the article, we develop the experimental tools to engineer the Hamiltonian (7.7) in the adiabatic regime. We describe how to generate a strong, coherent, and slowly varying field $\vec{B}(t)$ in the rotating frame satisfying $B_0 \gg g$, and verify adiabatic following of the qubit in the absence of the boost cavity and show that the relevant coherence time for the system is approximately T_1 and is not limited by the relatively short T_2 time.

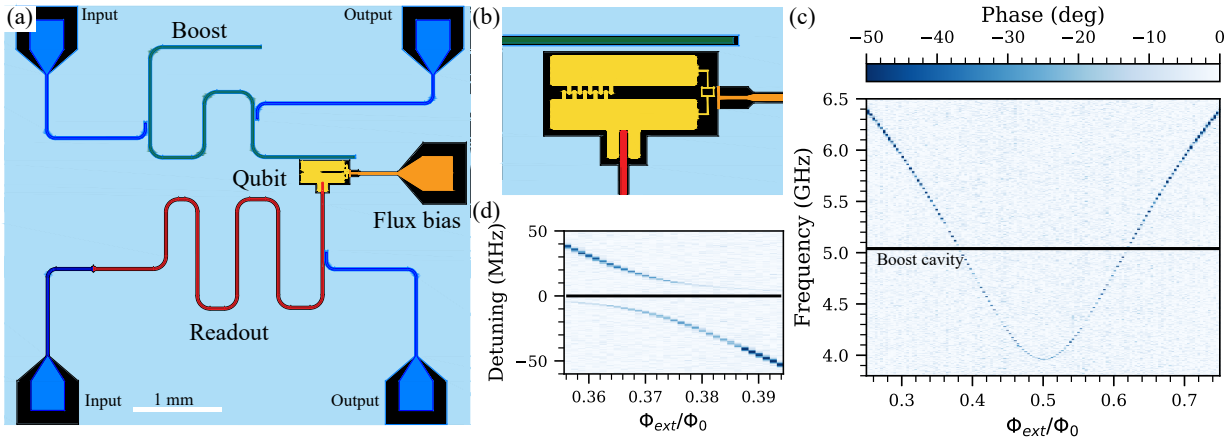

Figure 7.4: Device layout and qubit spectroscopy. (a) Full device layout: readout cavity in red ($\lambda/2$ CPW resonator), tunable transmon in yellow with associated flux bias line (orange), a secondary, low-loss cavity used for the full topological pumping protocol (green, unused in this experiment), and an external magnet for DC bias (not shown). (b) Zoom in on the qubit layout. (c) Qubit spectroscopy as a function of externally applied flux (in units of Φ_0, the flux quantum). The qubit frequency is measured using dispersive readout and the colorbar indicates the phase shift on the readout cavity. The qubit is tunable over a 2.5 GHz range and can be brought into resonance with the boost cavity (black line at 5.04 GHz). The resulting avoided crossing is shown in (d) where the detuning is measured relative to the boost cavity frequency.

7.3 Experimental Implementation and Device Design

In this section. we show that the desired Hamiltonian [Eq. (7.7)] can be realized in the circuit QED platform and present a device which satisfies the required hierarchy of scales shown in Eq. (7.10).

Our experimental platform consists of a driven tunable transmon coupled to a cavity and is described using the Jaynes-Cummings Hamiltonian:

$$H = \omega_c a^\dagger a + \omega_q \frac{\sigma_z}{2} + g(a^\dagger \sigma^- + a\sigma^+) + 2\Omega \cos(\omega_d t)\sigma_x \tag{7.12}$$

where $\omega_{c,q,d}$ are the frequencies of the cavity, qubit, and drive respectively; a is the bosonic annihilation for the cavity, σ^- is the Pauli operator for the qubit; and Ω, g are the driving strength and qubit-cavity coupling strength respectively. The device design, shown in Fig, 7.4(a), features a qubit coupled to two cavities. One cavity, which we denotes as the

boost cavity, is configured to host the topological pump, whereas a second, lower Q, cavity is used to read out the state of the transmon. This readout cavity is also used to apply drives to the qubit. The qubit frequency is tuned through the application of an external magnetic flux either from an on-chip flux bias line or from an external magnet.

We synthesize the required rotating field in the *rotating frame* of the applied qubit drive [specified by the transformation $U = \exp\{(i\omega_d t[\hat{n} + \frac{\sigma_z}{2}])\}$]. Rewriting the JC Hamiltonian in this frame and dropping rapidly oscillating terms we obtain:

$$H = \Delta a^\dagger a + \delta\frac{\sigma_z}{2} + g(a^\dagger\sigma^- + a\sigma^+) + \Omega\sigma_x \tag{7.13}$$

where $\Delta = \omega_c - \omega_d$ ($\delta = \omega_q - \omega_d$) is the detuning between the cavity (qubit) and the applied drive. By comparing terms to Eq. (7.7), we obtain the required form for δ and Ω to synthesize the topological pump Hamiltonian:

$$\delta(t) = B_0\sin(\omega_{\text{mod}}t),\ \Omega(t) = B_0[m + \cos(\omega_{\text{mod}}t)] \tag{7.14}$$

where B_0 and ω_{mod} are the amplitude and frequency of the effective magnetic field respectively. We refer to the δ term as an effective z-magnetic field and the Ω term as an effective x-magnetic field. Their generation and calibration are described in Section 7.4.

The readout cavity [red trace in Fig. 7.4(a)] is a standard coplanar waveguide half-wave resonator with resonant frequency 7.492 GHz and linewidth 350 kHz. As the qubit microwave drives are also applied through this cavity, it is configured in a transmission geometry to reduce the power reaching the readout amplification chain. The cavity is capacitively coupled to the qubit with coupling strength 90 MHz to facilitate large microwave drive strengths on the qubit.

The tunable transmon [yellow in Fig. 7.4(a) and (b)] consists of a DC SQUID loop in parallel with floating capacitor pads. The SQUID loop is made with asymmetric junctions (with an inductive energy ratio: $E_{J_1} = 3E_{J_2}$), reducing the noise from the flux bias line by

decreasing the frequency tuning range. A relatively large charging energy, E_c, of 240 MHz is used for this qubit to improve the performance at large microwave drive strengths. Qubit spectroscopy as a function of applied external flux is performed using two-tone spectroscopy [89] showing a tuning range from 4 to 6.4 GHz in Fig. 7.4(c). The qubit can be tuned into resonance with the boost cavity. The resulting avoided crossing is shown in Fig. 7.4(d), from which the qubit-boost coupling strength is extracted ($g = 13$ MHz).

The qubit frequency is controlled by two sources of external flux: an external magnet providing the DC operating frequency of the qubit and an on-chip flux bias line [orange in Fig. 7.4(a)] with a 5 MHz bandwidth for AC modulation. The on-chip flux bias line (FBL) is driven with an arbitrary waveform generator to provide full control over the frequency modulation.

The boost cavity [green trace in Fig. 7.4(a)] consists of a $\lambda/4$ CPW resonator at 5.04 GHz with $\kappa = 84$ kHz. It is designed for the full implementation of the topological pump described in the introduction to store the boosted state, but unused in this manuscript. The particular 90-degree orientation of the readout and boost cavity minimizes the direct classical crosstalk between the cavities by introducing an effective mode-mismatch.

7.4 Calibration of B Field Components

Achieving the Hamiltonian in Eq. (7.7) requires a time-dependent effective magnetic field in the x-direction (B_x) and z-direction (B_z). In the sections below, we present the experimental procedure to implement and calibrate B_x and B_z respectively. As the field components originate from different sources, we also develop a method to measure the delay between them. In particular, we synthesize a *circularly* polarized field of the form:

$$\vec{B} = B_0 \begin{pmatrix} \cos(\omega_{\text{mod}} t) \\ 0 \\ \sin(\omega_{\text{mod}} t) \end{pmatrix} \tag{7.15}$$

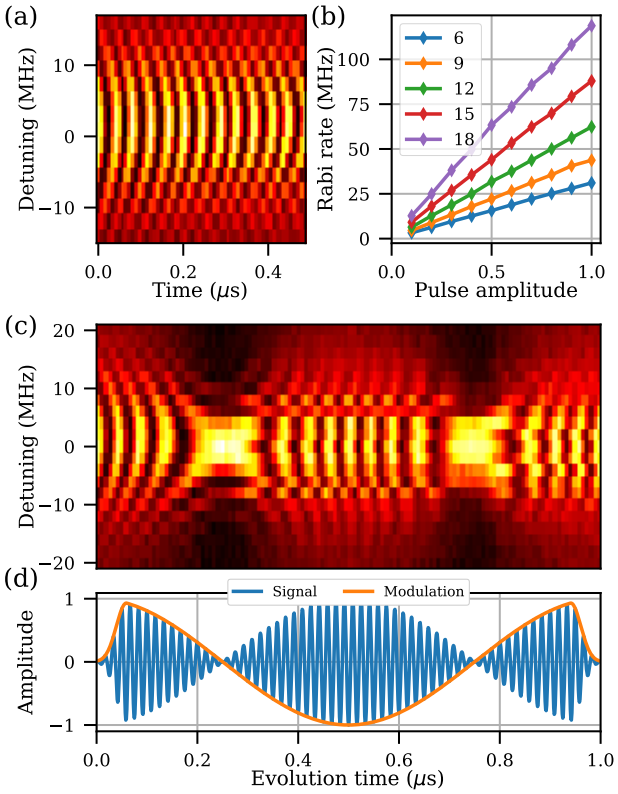

Figure 7.5: Calibration of B_x: (a) Rabi oscillations as a function of qubit-drive detuning and evolution time showing the usual Rabi chevron pattern. (b) Extracted Rabi rate as a function of AWG modulation amplitude for different generator output powers, achieving a peak Rabi rate of 110 MHz (in units of dBm). (c) Rabi chevron for amplitude modulated drive at $\omega_{\text{mod}} = 1$ MHz [modulating waveform shown in (d)]. The amplitude of the drive is directly mapped to the qubit population oscillation frequency as seen by the fact that the zero crossings in (d) are accompanied by a freezing of the qubit dynamics. For clarity, we present an example waveform which a much lower carrier frequency in (d).

with strength B_0 and (angular) rotation rate ω_{mod}. To isolate the performance of the synthesized magnetic field portion of the Hamiltonian in Eq. (7.7), the qubit is detuned from the boost cavity.

7.4.1 Transverse Field (B_x)

We implement an effective B_x by driving the qubit on resonance with a microwave drive (x-drive through the readout cavity. The Hamiltonian for a two level system under resonant drive is given by $H_q = \frac{\Omega}{2}\sigma_x$ where Ω is the Rabi rate and is linearly proportional to the amplitude of the drive. The applied drive leads to oscillations between $|g\rangle$ and $|e\rangle$ at a rate Ω. Measuring the $|e\rangle$-state population as a function of time and drive detuning produces a Rabi chevron pattern as shown in Fig. 7.5(a). As the detuning between the microwave drive and qubit frequency increases, the oscillation frequency increases with a corresponding decrease in the oscillation amplitude as given by:

$$P(e) = \frac{\Omega^2}{\Omega^2 + \Delta^2} \cos\Big(\sqrt{\Omega^2 + \Delta^2}t\Big). \tag{7.16}$$

We calibrate the effective transverse field (B_x) as a function of generator power and modulation amplitude using qubit spectroscopy, as shown in Fig. 7.5(b). For each combination, we perform a standard Rabi measurement to extract the Rabi rate. As expected, the measured Rabi rate scales as the square root of the applied power meaning that a 6 dB increase is required to double the Rabi rate. The maximum Rabi rate of 110 MHz is limited by the output power of our generator and the linewidth of the cavity used to drive the qubit and not by degradation of the qubit response, despite an anharmonicity of $\alpha = 240$ MHz.

The required time-dependence of the transverse field is generated by modulating the amplitude sinusoidally: $\Omega(t) = \Omega_0 \cos(\omega_{\text{mod}}t)$. The modulation is performed using an arbitrary waveform generator (AWG), providing full control over the waveform, connected

to the IQ modulation port of the qubit generator. Simple amplitude modulation is not sufficient as the *sign* of the applied field must change during the modulation period. We verify the time-dependence of the applied field by acquiring a Rabi chevron measurement in the presence of modulation, shown Fig. 7.5(c). The amplitude modulation of the drive [shown in Fig. 7.5(d)] is visible as a time-dependent oscillation frequency in the population. In particular, the zero crossings in the field amplitude are identifiable as locations where the dynamics freeze (at 0.25 and 0.75 μs).

Using the method described above, we can generate both static (DC) and dynamic (AC) effective x-fields on the transmon with amplitudes reaching 110 MHz, significantly exceeding the native boost cavity-qubit coupling of 13 MHz as required by Eq. (7.10).

7.4.2 Longitudinal Field (B_z)

The z-component of the rotating field is implemented by modulating the qubit frequency about a DC set point ω_{q_0}. An external magnet with large dynamic range and low bandwidth is used to set ω_{q_0}, while an on-chip FBL (driven by the AWG) applies a sinusoidal frequency modulation to the qubit frequency, effectively generating

$$B_z(t) = B_0 \sin(\omega_{\text{mod}} t)\sigma_z, \tag{7.17}$$

where B_0 and ω_{mod} are the field strength and rotation modulation rate respectively.

Just as with the B_x calibration, we first calibrate the size of the frequency shift. This is performed using qubit spectroscopy during a long (1 μs) flux pulse as a function of pulse amplitude from the AWG.

The x and z components of the applied field originate from two very different control processes and natively have an uncontrolled delay between them. However, they can be synchronized using two measurement protocols which convert the in-situ delay between the components into qubit-state population. The first is a time-resolved spectroscopy

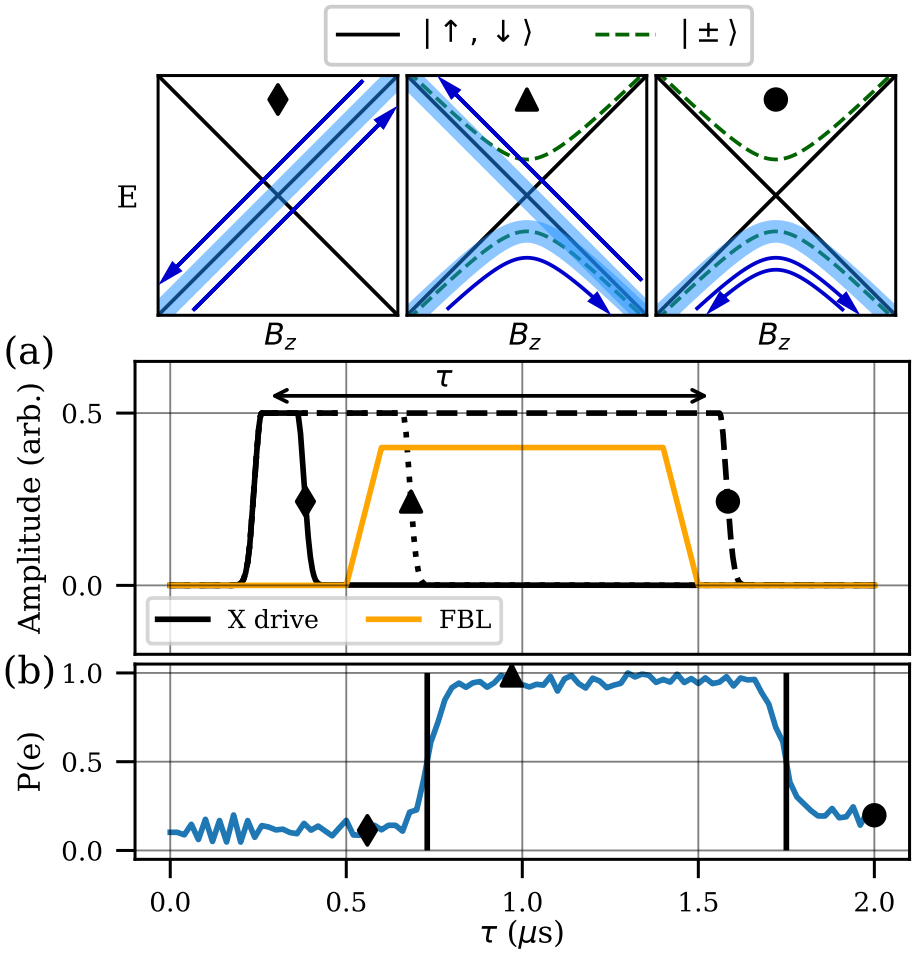

Figure 7.6: Calibration of the delay between x and z components of the applied field using Landau-Zener transitions. (a) Pulse sequence used to determine the delay: a fixed 1 μs flux bias pulse (orange) is applied to the qubit to ramp it through resonance with a microwave drive (black) of variable duration τ. Resonance occurs at the middle of the FBL pulse ramp such that both the starting and ending qubit frequency are far off resonance with the applied drive. (b) Measured qubit population as τ is varied. There is a net qubit population excitation when the microwave drive pulse overlaps with the FBL pulse. The line delay is extracted by comparing the measured time interval of the qubit excitation to the position of the FBL pulse in the schedule. Insets: energy level diagrams for the dressed states ($|g, n\rangle$, $|e, n-1\rangle$, black) and coupled states ($|\pm\rangle$ dashed green) of Eq. (7.18). The markers indicated configurations for which the system undergoes zero (diamond), one (triangle), or two (circle) LZ transitions.

measurement of the qubit to map out the flux bias pulse, and the second is a measurement of Landau-Zener transitions which occur when the qubit is ramped through resonance with an applied drive on. While the first method provides a full picture of the FBL pulse, it is time intensive, requiring a spectroscopic measurement at every time interval. We use the full time-dependent characterization to measure the filter induced distortions on the FBL pulse in Section 7.5. The Landau-Zener (LZ) based measurement on the other hand only requires a single frequency measurement making it significantly faster.

Landau-Zener transitions occur when two coupled energy levels are slowly swept through resonance. Here we use an analogous process in the presence of AC drives, known in this context as adiabatic rapid passage, to convert the time-dependent qubit detuning to a qubit excitation which can be readily measured. Consider the following coupled two-level system Hamiltonian:

$$H = \frac{1}{2}\begin{pmatrix} -B_z & B_x \\ B_x & B_z \end{pmatrix}, \tag{7.18}$$

where B_x is the driving strength and B_z is the detuning between the qubit and the drive. The eigenstates, $|\pm\rangle$, of this Hamiltonian are admixtures of the bare states $|g, n\rangle, |e, n-1\rangle$ with a detuning dependent mixing fraction. As B_z is swept through zero, the character of the ground state changes from $|g, n\rangle$ to $|e, n-1\rangle$. By slowly sweeping the detuning between the levels ($\dot{B}_z \ll B_x^2$) it is possible to perform a high-fidelity population transfer from one level to the other without requiring a π-pulse. The full energy diagram for the LZ transition is shown in the insets of Fig. 7.6, where the bare states are shown as the solid black lines and the coupled states are indicated with the dashed green lines.

Our measurement of the line delay utilizes LZ transitions to locate the times at which the qubit crosses resonance with an applied drive. In particular, we apply a 1 μs square flux bias pulse [shown in orange in Fig. 7.6(a)] of known amplitude and apply an off-resonant drive on the qubit [shown in black in Fig. 7.6(a)]. Schematically, we can write the pulses

as:

$$B_x(t) = (\Omega/2) b_{[0,\tau]}(t) \sigma_x \tag{7.19}$$

$$B_z(t) = \Delta[-1/2 + b_{[\tau_{\rm del}, \tau_{\rm del}+1\,\mu{\rm s}]}(t)] \sigma_z \tag{7.20}$$

where Ω, Δ are the drive strength and qubit detuning respectively, $b_{[\tau_1,\tau_2]}(t)$ is a unit amplitude step function supported in the time interval $[\tau_1, \tau_2]$, and $\tau_{\rm del} > 0$ is the unknown delay between the lines. LZ transition occurs when B_z changes sign and B_x is non-zero. The delay time $\tau_{\rm del}$ can then be found by measuring the qubit population as a function of τ and mapping out the interval $[\tau_{\rm del}, \tau_{\rm del} + 1\,\mu{\rm s}]$ as shown in Fig. 7.6(b).

We consider three qualitatively different regimes for the microwave pulse duration corresponding to zero, one, and two LZ transitions respectively in the insets of Fig. 7.6. At short hold times ($\tau < \tau_{\rm del}$, diamond marker), the microwave drive and FBL pulse are temporally separated leading the qubit to ramp through resonance with no applied drive and stay in the ground state. At intermediate times ($\tau \in [\tau_{\rm del}, \tau_{\rm del} + 1\mu{\rm s}]$, triangle marker), the qubit ramps through resonance with the applied drive on during the rising edge of the FBL pulse, leading to a LZ transition to the $|e\rangle$ state. The falling edge of the FBL pulse occurs with no applied drive, leaving the qubit in $|e\rangle$. Finally, at long hold times ($\tau > \tau_{\rm del} + 1\mu{\rm s}$, circle marker), both the rising and falling edges of the FBL pulse occur with an applied microwave drive leading, to two LZ transitions which bring the state back to $|g\rangle$ at the end of the measurement.

The methods described above provide a reliable way to synthesize z-fields as well as set the phase relationship between the x and z-components of the magnetic field. For the circularly polarized field described in the introduction, the two components are driven with a 90-degree phase shift, although in principle any polarization is achievable with this toolkit.

7.5 Flux bias line filter compensation

In this section, we discuss how we mitigate the ringing that occurs on the flux bias line due to the low-pass filter. This ringing introduces errors in the flux pulse shape applied to the qubit and prevents accurate measurements of σ_x at the end of an experiment as the qubit frequency oscillates during the $\pi/2$ pulse, unless appropriately compensated for.

We use a low-pass filter (minicircuits SLP 5+) on the on-chip flux bias line to reduce the bandwidth of the noise reaching the qubit, increasing its coherence time. To maximize the available bandwidth but minimize the noise power reaching the qubit, the filter used is a sharp elliptical filter with a cutoff frequency at 6 MHz. The filter introduces two artifacts that must be corrected: a group delay which increases drastically near resonance (from 140 ns below 4 MHz to 320 ns near 6 MHz) and a ringing response due to the sharp cutoff. The group delay is calibrated out with the method defined in the main text and assumed to be a constant over the range of frequencies considered ($\omega_{\text{mod}} = 1$ to 5 MHz). To reduce the ringing in the FBL pulses reaching the qubit, we perform a simplified version of compensation which pre-distorts the signal before reaching the filter in a way to minimize the ringing response of the filter.

The required compensation is calculated from the time-domain response of the filter. The response is acquired at room temperature by sending a 10 ns pulse to the filter and measuring the ringing response. This response in shown in Fig. 7.7(a) (blue) and displays two types of features: a primary peak acting as smoothing function on the data and damped oscillatory response which introduces the ringing in the filtered data. Our approach to reduce the ringing is to simply flip the sign of the ringing response. The resulting precompensation filter effectively introduces ripples of the opposite phase in the data, and is shown in orange in Fig. 7.7(a). We simulate the effect of the filter on a square pulse numerically, shown in Fig. 7.7(b). The uncorrected response (green) displays large ringing at transitions in the square pulse whereas the computed predistortion (orange) displays oscillations of the opposite phase. The corrected response (blue) which is the combination

of the two shows a factor of 5 reduction in the ripple amplitude. While more sophisticated methods exist to predistort the signal, this method is computationally simple to apply and provides a large enough ripple suppression for the proof of concept measurements presented here [90]. Finally, the predistortion kernel introduces some damping near the cutoff which is compensated using a small increase to the amplitude of the oscillatory frequency portion of the applied z-field.

The effect of the precompensation is also verified with qubit spectroscopy with and without predistortion, shown in Fig. 7.7(c)-(d). The time dynamics of the FBL are acquired by performing pulsed spectroscopy at each time step with a short pulse (50 ns) to determine the real-time qubit frequency. While the finite duration of the spectroscopy pulse leads to some smearing of the time domain data, the pulses cannot be shortened further without the frequency resolution becoming too low. The effect of the predistortion is visible in the extracted qubit frequency [Fig. 7.7(e)] as a reduction in ringing near the sharp transitions in the FBL pulse as well as a more uniform transition to the oscillatory portion of the protocol. This spectroscopy technique is a simplified version of the cryoscope method described in Ref. [90] in which a Ramsey-type measurement is used to measure the accrued phase for fast flux pulses.

7.6 Transmons in strong synthetic fields

We confirm the simultaneous operation of the synchronized effective B-field components, and adiabatic following of the qubit by measuring the σ_x and σ_z components as a function of evolution time. Data shown in Fig. 7.3(b). The measurement is acquired by evolving the qubit for a variable time under a field of strength B_0 rotating at rate ω_{mod}, before shutting off the fields diabatically and measuring the qubit state. The qubit is initialized in an eigenstate of the field before rotation to reduce nutation about the applied field (see Chapter 6). The x-projection is measured by applying a $\pi/2$-pulse after the field shut-off

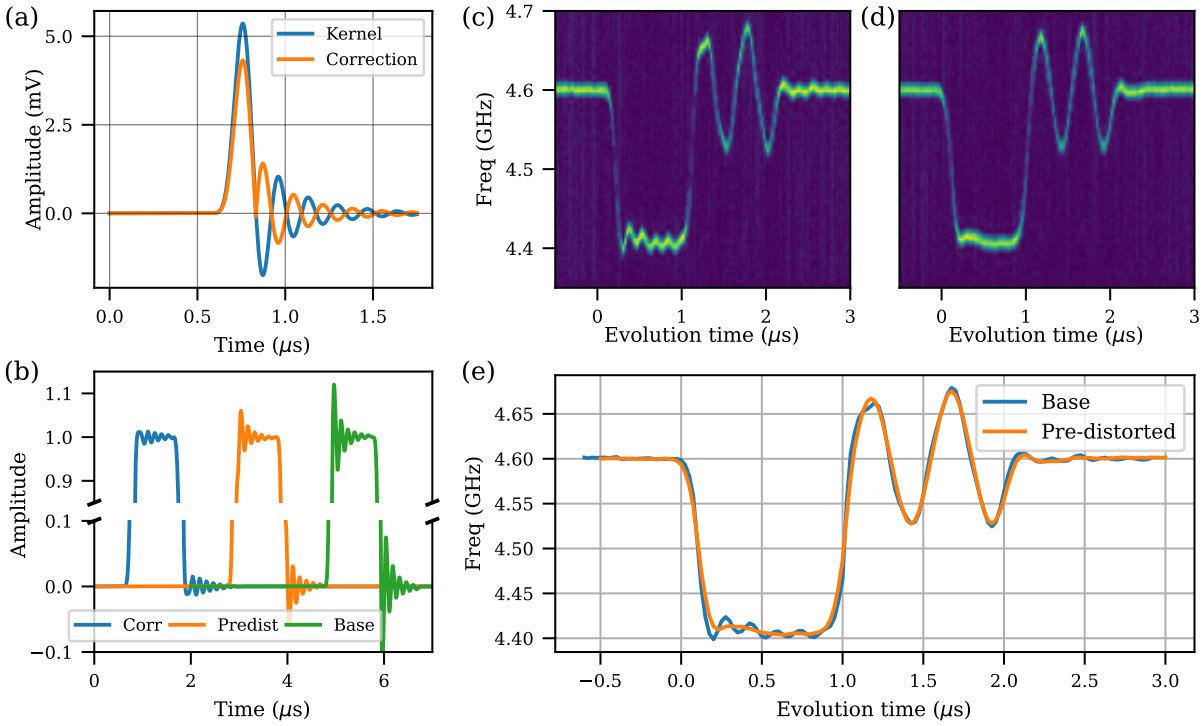

Figure 7.7: Filter compensation: (a) Filter convolution kernel (blue) acquired from the delta function response of an SLP5+ filter at room temperature. The compensation kernel (orange) is computed by flipping the sign of the oscillatory portion. Both kernels are normalized to provide unity gain. (b)Numerical simulation of corrected (blue) and uncorrected (green) square pulses going through the filter showing a 5x reduction in ringing amplitude. The compensation kernel is used to generate the predistorted trace (orange) which is shown *before* going through the filter. (c)-(d) Qubit spectroscopy during the full z protocol without (c) and with (d) pulse predistortion. Ripples after the initial ramp at $t = 0$ are reduced and the amplitude of the oscillatory portion of the schedule becomes more uniform. (e) Qubit frequency extracted from panels (c) and (d).

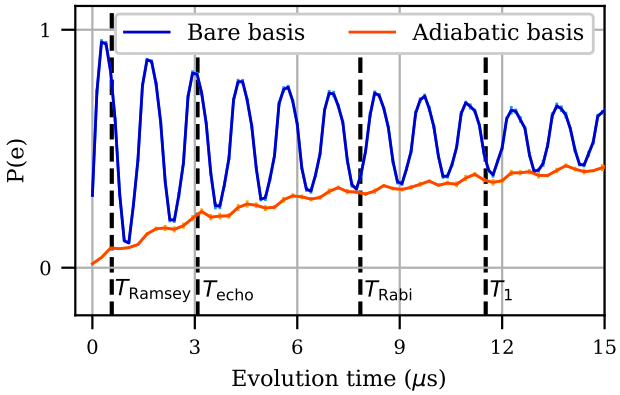

Figure 7.8: Measurement of the field following time (T_d) in the bare qubit basis (blue trace) and in the basis of the rotating field (orange). In both measurements, the field following time ($T_d = 12\ \mu$s) exceeds the other coherence times of the qubit ($T_{\text{Ramsey}}, T_{\text{echo}}, T_{\text{Rabi}}$) and is comparable the qubit lifetime T_1 (11.5 μs). The bare qubit basis measurements are acquired by an approximate diabatic shutoff of the fields and the adiabatic basis measurements are acquired by mapping the field $|\pm\rangle$ states back onto $|g\rangle, |e\rangle$ [71].

to rotate the measurement axis. In this measurement scheme, the x and z-components of the spin oscillate at frequency ω_{mod} with a 90-degree phase shift as expected from a spin locked to $B(t)$ and rotating in the xz plane.

In addition to the measurement described above, we perform complementary measurements in the adiabatic basis of the rotating field (i.e. the instantaneous eigenbasis of H) which directly measure the Floquet-state fidelity as a function of time. This measurement is performed by adiabatically rotating the final state back to a fixed point on the Bloch sphere (in our case back to $-z$) before performing readout [71]. A comparison between the two measurement methods is shown in Fig. 7.8. Both measurements evolve towards a mixed state at long evolution time, with the adiabatic basis measurement tracing out the envelope of the bare qubit basis oscillatory data.

Measurements in either basis (bare qubit basis or adiabatic basis) demonstrate that the field following time (T_d) far exceeds the qubit phase coherence times (T_2, T_{echo}), despite the flux sensitive nature of our architecture, and that T_d is comparable to the polarization time, T_1, as shown in Fig. 7.8. We also measure an improvement compared to the Rabi time which is less sensitive to the high-frequency noise on the qubit transition frequency

but is limited by slow frequency qubit drifts. The long field following time indicates that the adiabatic following is effectively insensitive to z-field noise on the qubit transition, reminiscent of the spin-locking time in NMR-systems [91].

In addition to the inevitable limit set by qubit decoherence, adiabatic following of the field can also be limited by diabatic processes which cause the qubit to no longer follow the applied field. Below we consider only this additional diabatic breakdown and show that for large enough fields, the qubit follows the applied field, without extra nutation, independent of the rotation rate and field amplitude. We measure the field following over a range of field amplitudes (2-80 MHz) and field rotation rates (0.5-5 MHz) in both the bare and adiabatic bases. Adiabatic following is identifiable in the bare basis measurements as a decaying oscillation with frequency ω_{mod}, as seen in the large field traces in Fig. 7.9(a-b) . For the adiabatic basis measurements, ideal following is characterized by a pure exponential response.

The breakdown in adiabatic following at low B_0 or high ω_{mod} is marked by the appearance of oscillations in the adiabatic eigenstate populations as the qubit nutates about the applied field. This can be measured directly via the Fourier spectrum of the adiabatic basis measurements. The degree of breakdown is quantified by the fraction of the spectrum that lies above a cutoff frequency corresponding to the linewidth of the ideal Lorentzian response centered at zero frequency. This metric, termed fractional harmonic content, is defined as follows:

$$F = \frac{\int_{\omega_c}^{\omega_s/2} d\omega |f(\omega)|^2}{\int_{0}^{\omega_s/2} d\omega |f(\omega)|^2}, \tag{7.21}$$

where ω_c is the cutoff frequency ($2\pi \times 167$ kHz), ω_s is the sampling rate of the data, and $|f(\omega)|^2$ is the power spectrum of the data. For ideal following, this metric approaches zero whereas a breakdown in adiabatic following corresponds to a value close to one. We compute this metric over a range of field amplitudes and rotation rates, shown in Fig. 7.9(a), and observe a sharp crossover in F when $B_0 < 2.5\,\omega_{\mathrm{mod}}$ indicating the breakdown of adiabatic following.

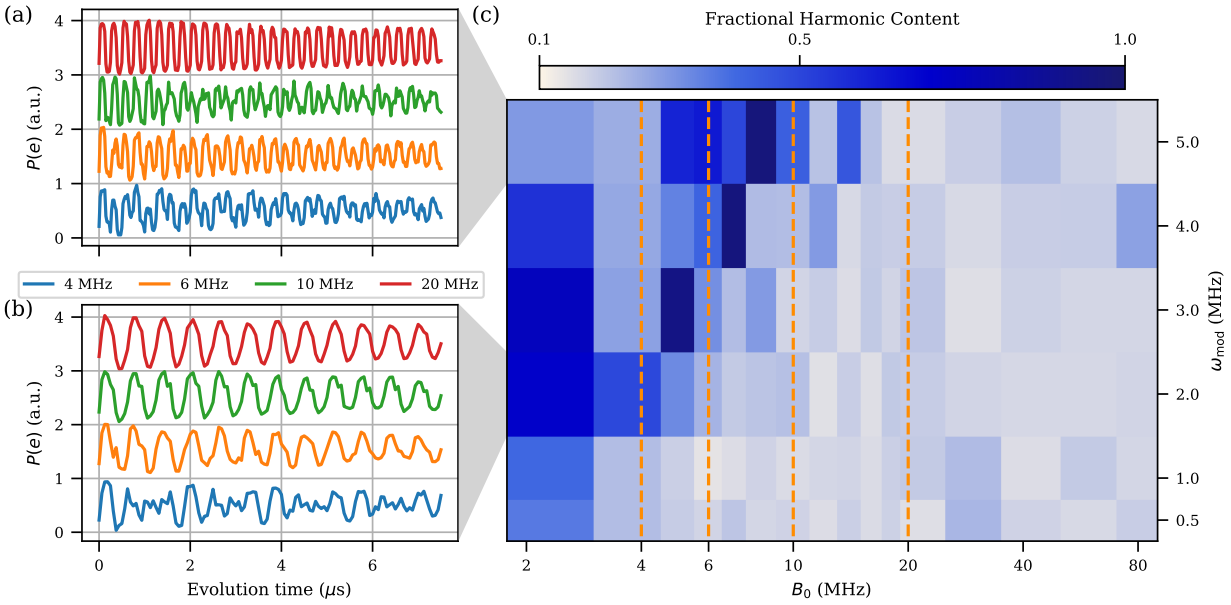

Figure 7.9: Degree of adiabatic following of external field (a) Bare qubit basis measurement versus time with a fast rotation rate $\omega_{\text{mod}} = 5$ MHz and $B_0 = 4$,6,10,20 MHz [shown in the legend of (a), (b)]. At large $B_0 = 20$ MHz, the qubit state is effectively pinned to the rotating field as shown by the sinusoidal oscillations with rate ω_{mod}. As B_0 decreases, the qubit no longer adiabatically follows the rotating field; instead, it rotates with the field while nutating about it, creating multiple frequency components in the traces. (b) Bare qubit basis traces for a slower rotation rate $\omega_{\text{mod}} = 2$ MHz as compared to (a) showing onset of good following around 6 MHz as opposed to the 20 MHz field required for the faster rotation rate in panel (a). (c) Fractional harmonic content [as defined in Eq. (7.21)] as a function of applied field strength and rotation rate measured in the adiabatic basis. Lower values indicate smaller deviations from adiabatic following. For $B_0 > 2.5\ \omega_{\text{mod}}$, the system follows the applied field with no observed dependence on the field size or rotation rate.

7.7 Time dynamics and Fourier spectrum of adiabatic following data

In this section we discuss the breakdown in the adiabatic field following as measured in both the adiabatic basis and the bare basis. To quantify this, we examine the Fourier spectrum of the data and show that the breakdown is accompanied by the emergence of additional frequency components.

We first show the breakdown of adiabatic following for measurements in the adiabatic basis of the rotating field. In this measurement, the instantaneous eigenstates of the rotating field are mapped onto $|g, e\rangle$ before readout. For ideal following, the qubit follows the applied field perfectly leading to an exponential decay to a mixed state as the qubit population decays. For sufficiently large rotation rates (or small applied field strengths), the qubit no longer follows the applied field and instead nutates about the applied field. The transition from adiabatic following to nutation is visible in Fig. 7.10. For very slow rotation rates or large field strengths, the qubit follows the applied field independently of the field amplitude or rotation rate respectively. At higher rotation rates or smaller applied fields, the breakdown in adiabatic following is visible in the emergence of an oscillatory component in the traces.

The breakdown of adiabatic following is more visible in the Fourier transform of the data. For ideal following, the exponential decay leads to a Lorentzian response centered at zero frequency, with a width set by the decay rate. The breakdown regime is identified as deviations from this Lorentzian regime, either through the emergence of peaks at higher frequencies or by a general increase in the high-frequency portion of the spectrum. The frequency spectrum for different rotation rates and field strengths is shown in Fig. 7.11. The adiabatic following regime is clearly identifiable for $B_0 > \ 2.5\ \omega_{\mathrm{mod}}$, as the region in which the data exhibits low Fourier amplitude at higher frequencies. As stated in the main text [Eq. (7.21)], we quantify the adiabatic following breakdown using the fraction of the

spectrum above the cutoff frequency (defined as the width of the ideal Lorentzian response and shown as a dashed red line in Fig. 7.11). The value of each pixel in Fig. 7.9(c) of the main text is the fraction of the spectrum above this cutoff frequency.

The breakdown of adiabatic following is also visible in the bare qubit basis measurements. In bare basis measurements, ideal following is characterized by oscillations at the field rotation frequency ω_{mod}. As with the adiabatic basis measurements, the oscillations damp out to a mixed state as the qubit depolarizes. When the Fourier spectrum only exhibits peaks at ω_{mod} and its harmonics, the qubit is successfully following the applied field. Fourier spectra versus ω_{mod} and B_0 are shown in Fig. 7.12.

7.8 Conclusion

In this article, we have demonstrated an experimental toolkit to synthesize large rotating effective magnetic fields for a transmon. The effective field is synthesized by a combination of amplitude modulation of a resonant microwave drive (x-component) and sinusoidal frequency modulation of the qubit (z-component). We achieved time-dependent fields with peak amplitude exceeding 100 MHz, far exceeding the qubit-boost cavity coupling g of 13 MHz. We also provide a simple method to combine the field components into a controlled rotating field.

We show that, for large fields, the qubit adiabatically follows the applied field by performing measurements both in the bare qubit basis (where the population oscillates at ω_{mod}) and in the adiabatic basis (where the instantaneous eigenstates of the applied field are mapped onto $|g\rangle, |e\rangle$). The field following time (T_d) for exceeds the qubit phase coherence times and Rabi time indicating a relative insensitivity to wide band frequency noise.

Finally, we map out the adiabatic following regime for B_0 (the field amplitude) and ω_{mod} (the field rotation rate) and show that the adiabatic following is insensitive to B_0, ω_{mod}

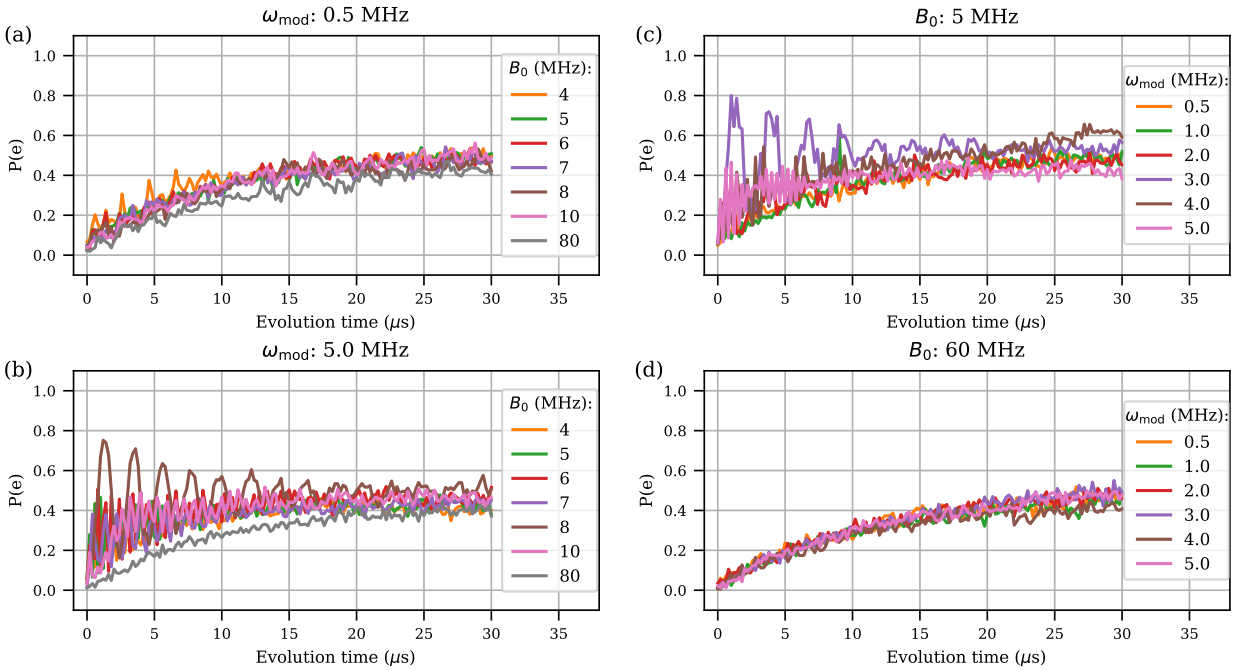

Figure 7.10: Adiabatic basis measurements of the qubit under the rotating field versus B_0 and ω_{mod}. (a)-(b) Each subpanel shows the probability of being misaligned to the applied field versus time and B_0 for fixed rotation rate ω_{mod}. (a) For slow rotation rates ($\omega_{\text{mod}} = 0.5$ MHz), the qubit follows the applied field regardless of the field strength B_0 and exponentially decays to a mixed state with a decay rate independent of field strength. (b) For fast rotation rates ($\omega_{\text{mod}} = 5$ MHz), the breakdown in adiabatic following is visible in the emergence of large contrast oscillations superposed on the exponential decay. These features correspond to nutation about B_0. For $\omega_{\text{mod}} = 5$ MHz, adiabatic following occurs when the field strength is larger than the threshold $B_0 = 10$ MHz. (c)-(d) Each subpanel shows the probability of being misaligned to the applied field versus time and ω_{mod} for fixed applied field strength B_0. (d) At large field strength ($B_0 = 60$ MHz), the qubit adiabatically follows the applied field regardless of rotation rate ω_{mod} and decays exponentially to a mixed state with a decay rate independent of rotation rate ω_{mod}. (c) At small field strength ($B_0 = 5$ MHz), adiabatic following breaks down as the rotation rate increases. For $B_0 = 5$ MHz, adiabatic following occurs when the rotation frequency is smaller than the threshold $\omega_{\text{mod}} = 1$ MHz.

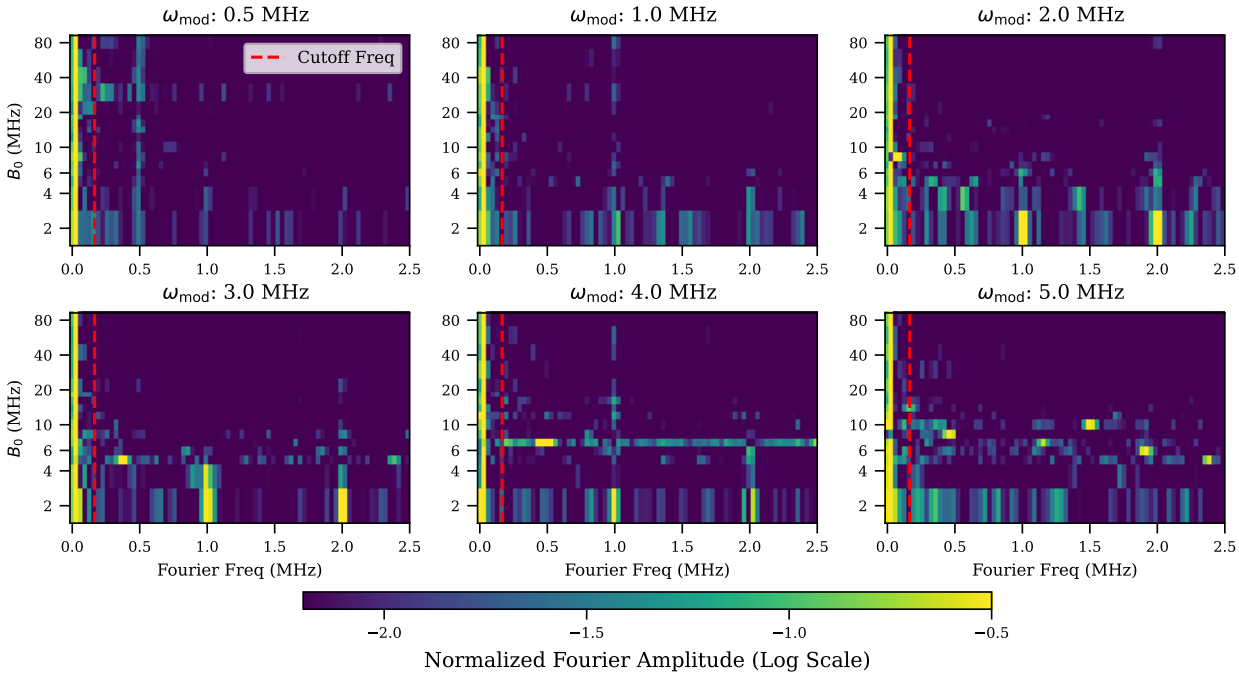

Figure 7.11: Normalized power spectrum of the adiabatic basis measurements. For each rotation rate ω_{mod}, the subpanel contains data of normalized power spectrum versus field strength B_0 by taking the Fourier transform of the traces in Fig. 7.10. In the case of ideal adiabatic following, the Fourier spectrum is a Lorentzian centered at zero frequency. Adiabatic breakdown is then identified as emergence of peaks at higher frequencies or by a general increase in the high-frequency portion of the spectrum. Good adiabatic following is clearly visible for $B_0 > 2.5\,\omega_{\text{mod}}$ as the peak Fourier amplitude is centered at zero frequency with minimum high frequency noise. The red dashed line shows the location of the cutoff used to compute the fractional harmonic content defined in Eq. (7.21) of the main text. Fig. 7.9 in the main text shows the fractional harmonic content of this data versus ω_{mod} and B_0.

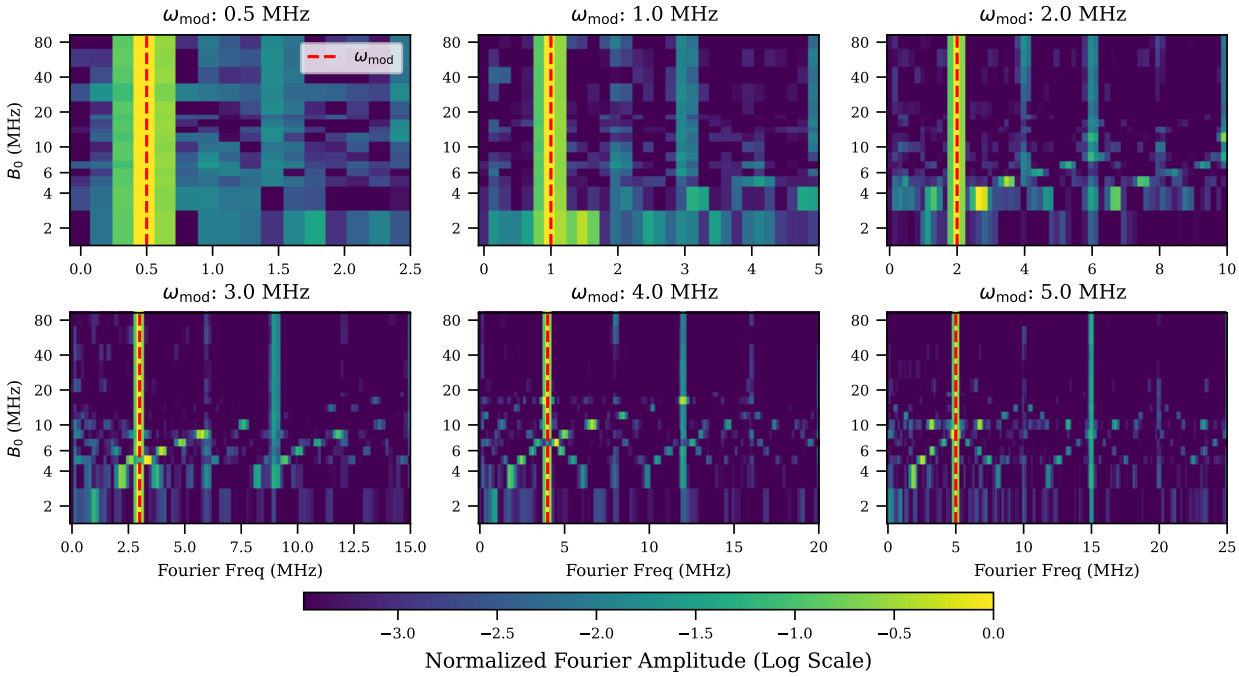

Figure 7.12: Normalized power spectrum of the bare basis measurements. For each rotation rate ω_{mod}, the subpanel contains data of normalized power spectrum versus field strength B_0 by taking the Fourier transform of the traces in Fig. 7.9(a)-(b) [red dashed line indicates ω_{mod}]. With ideal adiabatic following, the bare basis measurement should be an exponentially decaying sinusoidal function with frequency equal to ω_{mod}. This corresponds to large absolute Fourier amplitude at ω_{mod} (and its harmonics) but low amplitudes at other frequencies. Breakdown of adiabatic following is visible as additional features appear at frequencies other than ω_{mod}(or its harmonics). The additional features are due to nutation about the applied field and scale linearly with B_0. Note that the linearity is clearly visible in each subpanel due to logarithmic axes. Similar to Fig. 7.11, good adiabatic following is clearly visible for $B_0 > 2.5\ \omega_{\text{mod}}$ when the additional Fourier spurs disappear.

above the condition $B_0 > 2.5\ \omega_{\text{mod}}$.

The tools developed in this article demonstrate the required components and driving regime needed to implement the topological pump proposed in [85]. In addition, we show that the flux-tunable architecture does not severely limit the performance of the coherent drives as the field amplitudes reduce the sensitivity to flux noise.

Recent work in [71] developed a protocol to further reduce the sensitivity of the pump to T_1 decoherence using a secondary lossy cavity to stabilize the pumping state. With the exception of the additional cavity, the proposed protocol relies on the same configuration of driving fields presented here allowing for another method to experimentally observe the topological energy pump.

7.9 Technical Details

7.9.1 Stabilization of qubit frequency setpoint

Environmental fluctuations (such as changes in the lab temperature or base plate temperature in the dilution fridge) cause the qubit frequency to drift over the course of a measurement. To stabilize the qubit about the mean qubit frequency ω_{q_0} required for the synthetic field, an active servo is implemented through the DC bias line. During each measurement, the servo periodically measures the qubit frequency, and applies a correction to the DC bias current of the magnet to maintain the qubit at ω_{q_0}.

The servo takes in the desired qubit frequency ω_{target} as input. At each adjustment step, it measures the current qubit frequency ω through a standard spectroscopy measurement. The servo computes a required current correction $I_{\text{corr}} = (\omega_{\text{target}} - \omega)/A$, where A is a predetermined scaling factor, and adjusts the DC bias current to $I_{\text{new}} = I_{\text{old}} + I_{\text{corr}}$. A maximum cap of current correction I_{cap} is set such that each step the servo adjusts the DC bias current I_{old} by only $\min\{I_{\text{corr}}, I_{\text{cap}}\}$. The servo will run for a minimum number of adjustment steps N_{min} and stop once the measurement frequency is within tolerance

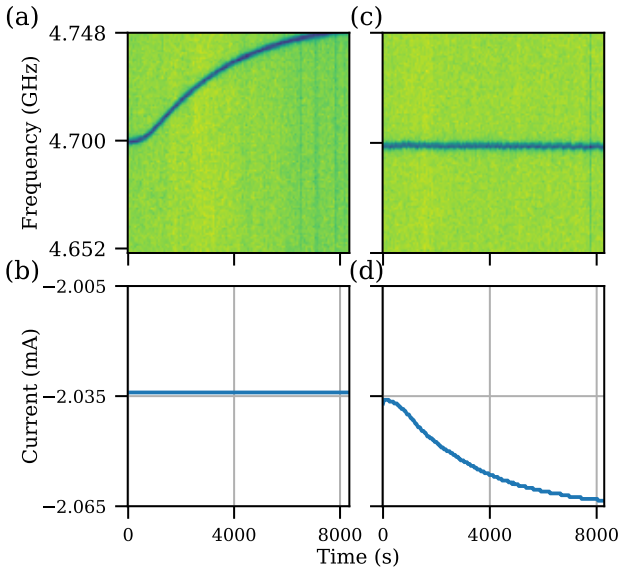

Figure 7.13: Characterization of qubit drift under environmental fluctuations with or without the servo. For these data sets, both the initial qubit frequency and the target frequency are set to 4.7 GHz. (a-b) Qubit spectroscopy and DC bias current as a function of time when the servo is turned off. To simulate environmental fluctuations, we turn on and off another device in the fridge to change base plate temperature. As the other device is turned on over 2 hours, qubit frequency steadily increases from 4.7 GHz to above 4.748 GHz. (c-d) Qubit spectroscopy and DC bias current as a function of time and under the condition but when the servo is turned on. For over two hours, the servo actively decreases the DC bias current to maintain qubit frequency at 4.7 GHz within tolerance $\delta_{\text{freq}} = 4$ MHz.

$\delta_{\text{freq}} < |\omega_{\text{target}} - \omega|$ or when the number of steps reaches a maximum number N_{max}. Fig. 7.13 shows qubit spectroscopy data and corresponding DC bias current data when the servo is off [Fig. 7.13(a)(b)] versus on [Fig. 7.13(c)(d)]. Because the device is sensitive to heat load from other devices in the fridge, the data is taken just turning on another device in the fridge, when the induced effect on the qubit frequency is the most dramatic. In both cases, the initial qubit frequency ω and the target frequency ω_{target} are set to 4.7 GHz. When the servo is off, the DC bias current I stays constant throughout the period and the qubit frequency gradually drifts away from ω_{target}. When the servo is on, it actively adjusts DC bias current I. The qubit frequency ω remains stable around $\omega_{\text{target}} = 4.7$ GHz with a tolerance of $\delta_{\text{freq}} = 4$ MHz.

7.9.2 Qubit Timescales

Qubit decoherence-timescale characterization for the qubit-cavity system is performed with the qubit frequency at 4.7 GHz, and shown in Fig. 7.14. The qubit decay time $T_1 = 11.5$ μs is extracted from an exponential decay fit to the data, shown in Fig. 7.14(a). We perform a Ramsey-style measurement with two consecutive $\pi/2$ pulses to extract the pure dephasing time $T_\phi = 570$ ns, shown in Fig. 7.14(b). To maintain a large separation of scales between the oscillatory and decaying time constants of the Ramsey signal while preserving the signal contrast, we apply a phase rotation to the second $\pi/2$ pulse. Measuring a tunable qubit away from its sweetspot introduces low-frequency noise. Therefore, the pure dephasing time $T_\phi = 570$ ns is extracted by fitting the data with a Gaussian decay envelope [92] instead of a standard exponential decay. We perform Hahn echo measurements and extract $T_{\text{echo}} = 3$ μs from an exponential fit, see Fig. 7.14(c). All data is taken with 100,000 shots by averaging ten acquisitions each with 10,000 shots per time point.

In addition, we extract the decay time of a 40 MHz Rabi oscillation and find $T_{\text{Rabi}} = 7.8$ μs, shown in Fig. 7.14(d). The data is taken by varying the hold time of a on resonance x-drive and is fit with a Gaussian envelope to extract the decay parameter. At high Rabi oscillation frequency, the measurement is sensitive to accumulated phase noise and fluctuations in qubit frequency. We therefore randomize the hold time to transfer phase noise into amplitude noise, take data with 50,000 shots from ten acquisitions with 5,000 shots per time point, and use a Gaussian fit instead of standard exponential to the data. The Rabi oscillations do not quite follow textbook form, likely due to higher levels of the transmon or frequency-dependent decay rates (e.g. due to TLS). The oscillations are not fully symmetric do not asymptote exactly to 50% population, but the Gaussian-envelope fit captures the dominant features, and the time constant is clear, as well as the fact that T_{Rabi} is distinctly shorter than T_d.

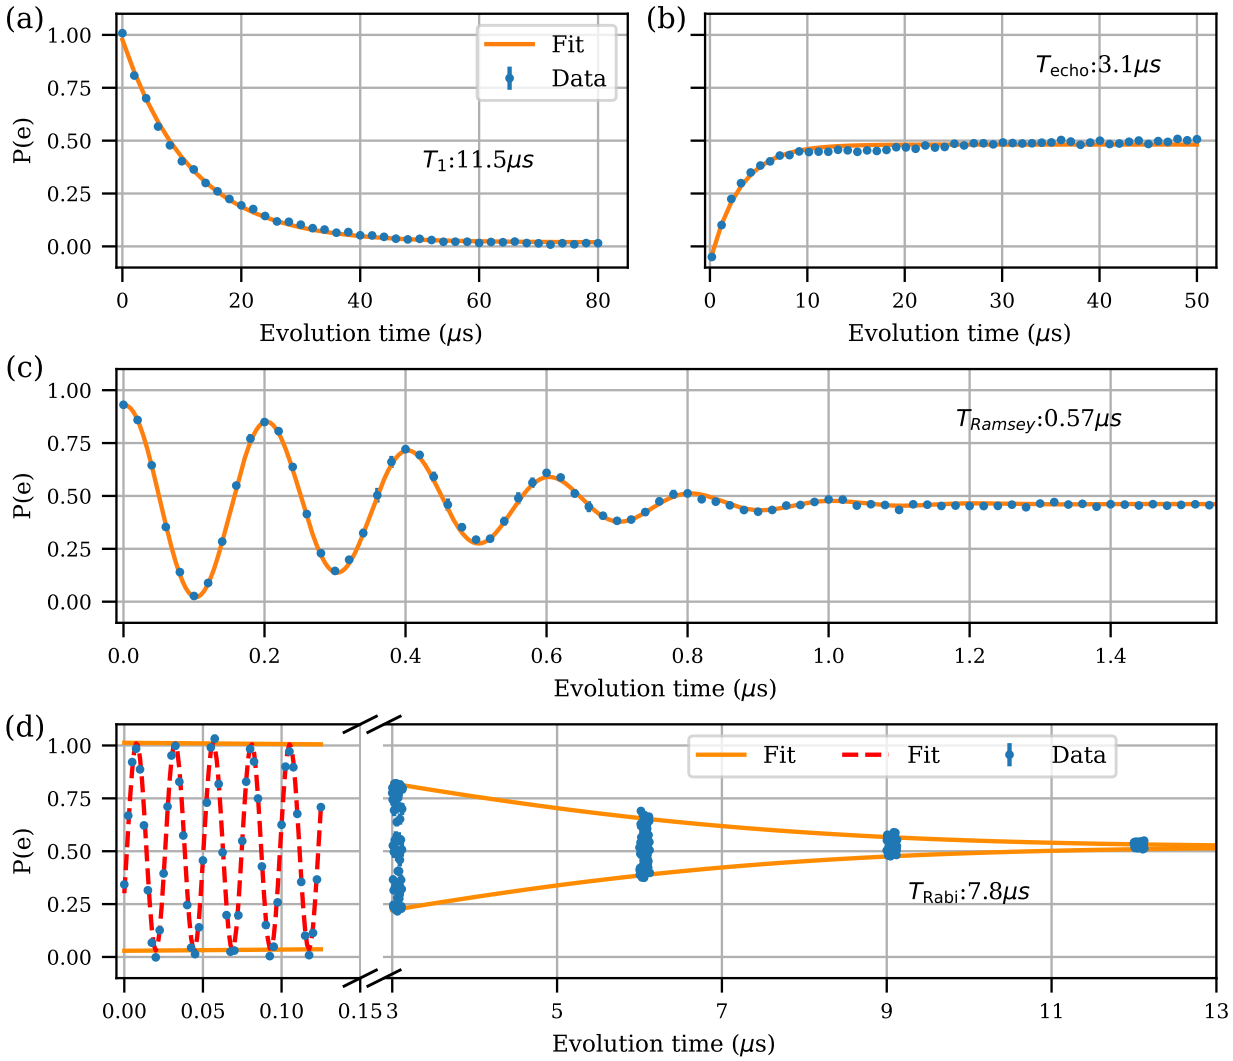

Figure 7.14: Qubit time scales at 4.7GHz. All traces are acquired as the average of 10 traces with 10,000 averaged measurements each. The statistical error bars are computed as the standard deviation between the traces and are smaller than the markers. (a) Excited state population as a function of delay time after qubit excitation. $T_1 = 11.5$ μs is extracted from an exponential fit to the data. (b) Excited state population after a standard Hahn echo sequence. $T_{\text{echo}} = 3.1$ μs is extracted from an exponential fit to the data. (c) Ramsey oscillations of the qubit population. Due to the short lifetime, a virtual rotation of the second $\pi/2$-pulse was used instead of using a detuned pulse. T_{Ramsey} is extracted from a Gaussian fit to the envelope of the data. (d) Rabi oscillations for a 40 MHz drive strength. T_{Rabi} is extracted from a Gaussian fit to the envelope due to phase noise arising from small qubit frequency drifts during the measurement.

Chapter 8

Driven-Dissipative Stabilization of Floquet States

This chapter is modified from [71] which was a collaboration with Anushya Chandran and David Long providing the theory components of the paper, and uses the same experimental device as Chapter 7. Qianao Yue (Billy) performed measurements with me and helped in the figure making process. Alicia Kollár provided help with experimental design and interpretation as well as figure layout and design. All authors contributed to writing and editing the manuscript.

8.1 Introduction

The ability to experimentally control and manipulate quantum systems has rapidly advanced over the last two decades. One indispensable tool in the quantum mechanic's toolkit is Floquet engineering, in which the properties of the system are qualitatively modified by strong periodic driving. It underlies the creation of optical lattices and synthetic band structures [77, 93–97], artificial gauge fields [77, 98–101], topological phases impossible in static systems [102–106], topological charge and photon pumps [78, 79, 81, 84, 107–111], and exotic quantum codes [112], to name a few.

However, Floquet engineering is always limited by heating. There are three conceptually distinct mechanisms for heating. Consider a single degree of freedom with a few energy levels subject to Floquet driving. The stationary states of the system (known as quasienergy states) involve admixtures of *both* ground and excited states. Accounting for dissipation and dephasing in the laboratory frame leads to *both* excitation and de-excitation processes in the quasienergy state basis, and thus hot steady states in this basis. Second, restricting to a few levels is always an approximation in a physical system. Higher order processes often allow the system to absorb energy from the drive and leak out of the intended operating subspace. Third, with many interacting degrees of freedom, the system can absorb energy from the drive and heat up, even if the first two mechanisms are suppressed [113–115].

We present a simple and general purpose dissipative stabilization scheme to mitigate the effects of the first type of (excited-state admixture) heating. Fortuitously, our scheme addresses the other two types of heating as well, at least within certain physical contexts.

Conventional dissipative preparation methods leverage native decay processes to successfully cool into ground states [116, 117] or dark states [118]. However, in the presence of strong drives, laboratory frame decay processes cause *both* excitation and de-excitation. This, combined with the unboundedness of Floquet spectra from above and below, causes conventional autonomous cooling schemes like optical pumping [116] and dark-state engineering [118–120] to fail. Their generalization to the time-dependent case is not known.

One route to stabilizing more general, but still time-independent, states is to engineer new dissipation terms distinct from the native dissipators [120, 121]. For example, in superconducting qubit architectures, two photon loss has been used to stabilize bosonic Kerr-Cat qubits [122–124], while selective photon addition can stabilize code spaces in a Cat code [125]. Collective dissipators can also stabilize target entangled states of atoms [126–132]. In principle, periodic modulation of engineered dissipators could implement spontaneous emission into a specific quasienergy state. Such an approach is however highly

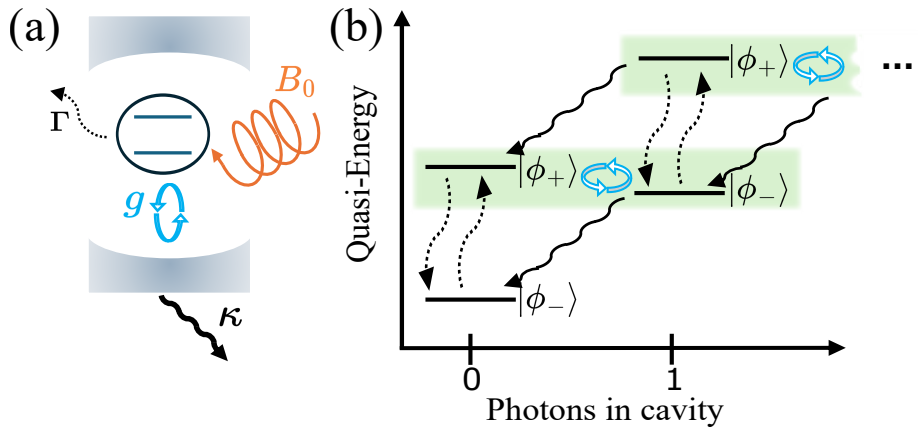

Figure 8.1: **Static dissipation stabilizes Floquet states.** (a) Schematic of the augmented system consisting of a spin with intrinsic dissipation rate Γ (dotted wavy arrow) subject to a strong drive of strength B_0 (orange), coupled to an auxiliary cavity with coupling strength g (light blue) and photonic decay rate κ (solid wavy arrow). (b) Quasienergy level structure of the augmented system. The representative quasienergy states, $|\phi_\pm\rangle$, are shown for $n = 0, 1$ photons in the cavity. Far from resonance, photon decay drives the system to $n = 0$ while preserving the spin state. On resonance however, the cavity coupling g hybridizes spin states with different photon numbers, allowing the state $|\phi_+\rangle$ in the $n = 0$ sector to convert to $|\phi_-\rangle$ and decay. When the spin decay rate Γ is much weaker, the unique steady state remains close to $|\phi_-, n = 0\rangle$.

fine-tuned, requiring knowledge of the target quasienergy state at all times during the drive cycle.

A different route to stabilization is to actively extract heat from the system continuously [133–141] or stroboscopically [142–144] into an auxiliary system. This type of approach has been used to stabilize arbitrary static single-qubit states [138–140] and entangled states of two qubits [141]. These methods have also been extended to Floquet systems in the limit of high-frequency driving, where they stabilize the ground state space of static effective Hamiltonians. Refs. [136, 137] numerically demonstrated such stabilization, while Ref. [144] experimentally cooled into the ground state of an effective many-qubit Hamiltonian by cycling between a Trotterized Hamiltonian evolution step and an auxiliary qubit reset step.

Here we show that, regardless of drive-frequency regime, a static coupling to an unmodulated but lossy auxiliary system can cool a Floquet system into a target quasienergy

state. This method is autonomous, requiring no active feedback, and exploits no microscopic knowledge of the time-dependence of the target quasienergy state.

We present experimental and numerical evidence for stabilization in a strongly driven two-level system coupled to a lossy bosonic mode. The experimental system consists of a strongly modulated transmon qubit in the low-frequency regime coupled to an auxiliary microwave cavity with photon loss [Fig. 8.1(a)]. Numerical simulations treat the combined system within the master equation framework and demonstrate stabilization in all drive regimes.

As an application beyond two-level systems, we develop the theory of a topological photon pump [78, 79, 81, 84, 109] coupled to an auxiliary lossy cavity. The topological photon pump is a strongly driven cavity-QED system with a long-lived photon current of topological origin. We show that the auxiliary cavity brings the value of the photon current closer to its ideal value, and increases the fidelity of high-photon number Fock states that are prepared in the cavity at special times [81].

Unlike dissipation in the Floquet system, dissipation in the auxiliary does not result in heating. In the limit of weak driving, the states of the Floquet system are near-degenerate and couple to the auxiliary symmetrically. However, in the limit where the drive in the Floquet model is strong compared to the coupling to the auxiliary, hybridization between the Floquet states and the auxiliary is dependent on the quasienergy [Fig. 8.1(b)]. When a resonance condition is satisfied, the loss in the auxiliary leads to the state $|d\rangle = |\text{target}(t)\rangle \otimes |\text{vac}\rangle$ being an approximate attractor of the combined dynamics, with the Floquet system in the target quasienergy state $|\text{target}(t)\rangle$ and the auxiliary in its vacuum state. Dissipation and decoherence in the Floquet system drive transitions away from $|d\rangle$, but the loss in the auxiliary pumps the combined system back to $|d\rangle$. As long as the auxiliary is the strongest dissipation channel, the steady state of the primary system has a high degree of purity and significant overlap with $|\text{target}(t)\rangle$.

The remainder of the manuscript is organized as follows. In Sec. 8.2, we describe

the physical setup involving a periodically modulated spin coupled to a lossy cavity, which showcases the autonomous stabilization effect. In Sec. 8.3, we present experimental evidence for stabilization in this physical setup in the low-frequency (adiabatic) regime. We turn to the general theory in Sec. 8.4, and provide numerical evidence for stabilization in all frequency regimes. Lastly, in Sec. 8.6, we show that dissipative stabilization enhances the performance of the topological photon pump [81, 84].

8.2 The Physical Set-up

Consider a minimal Floquet model of a single spin in a rotating magnetic field. The Hamiltonian is,

$$H_B = \frac{1}{2}\vec{\sigma}\cdot\vec{B}(t) = \frac{B_0}{2}\cos(\omega_{\text{mod}}t)\sigma_z + \frac{B_0}{2}\sin(\omega_{\text{mod}}t)\sigma_x, \tag{8.1}$$

with a fixed modulation frequency ω_{mod} and constant magnetic field amplitude B_0. The model happens to be soluble (it reduces to a static problem in the frame rotating about the y-axis with frequency ω_{mod}). This solubility is however irrelevant to what follows, as the augmented system with dissipation is not soluble. See the schematic in Fig. 8.1(a).

The quasienergy states, or Floquet states, are special solutions to the Schrödinger equation with the same periodicity as the drive. They take the form

$$|\psi_\pm(t)\rangle = e^{-i\epsilon_\pm t}\,|\phi_\pm(t)\rangle\,, \tag{8.2}$$

where "±" label the two states of the qubit, $|\phi_\pm(t)\rangle = |\phi_\pm(t+T_{\text{mod}})\rangle$ is the periodic quasienergy state, and $\epsilon_\pm$ is its corresponding quasienergy. The quasienergy states are not unique: the same physical solutions $|\psi_\pm(t)\rangle$ are obtained if we replace $\epsilon_\pm \to \epsilon_\pm + m\omega_{\text{mod}} = \epsilon_\pm^{(m)}$ and $|\phi_\pm(t)\rangle \to e^{im\omega_{\text{mod}}t}\,|\phi_\pm(t)\rangle = \left|\phi_\pm^{(m)}(t)\right\rangle$ for m an integer. Thus, we obtain countably many equivalent copies of the quasienergy states $\left|\phi_\pm^{(m)}(t)\right\rangle$ with quasieneries differing by

integer multiples of the drive (angular) frequency ω_{mod}.

As is the usual convention, we select $\left|\phi_{\pm}^{(0)}(t)\right\rangle$ from among these countably many choices by demanding $|\epsilon_{\pm}^{(0)}| < \omega_{\text{mod}}/2$. [For the model in Eq. (8.1), this choice imposes $\epsilon_{+}^{(0)} = -\epsilon_{-}^{(0)}$.] The interval of quasienergies $\epsilon \in [(m-\frac{1}{2})\omega_{\text{mod}}, (m+\frac{1}{2})\omega_{\text{mod}})$ is called the mth Floquet zone, in analogy to Brillouin zones of quasimomentum. In the intermediate and high frequency regimes, we work with the states $\left|\phi_{\pm}^{(0)}(t)\right\rangle$ belonging to the zeroth zone. In the adiabatic regime of small ω_{mod} however, it is helpful to use representative quasienergy states belonging to different zones (see the discussion at the beginning of Sec. 8.3). In all cases, we refer to the chosen two representative quasienergy states as $|\phi_{\pm}\rangle$, without the Floquet index and with time-dependence implicit, and to their quasi-energies as $\epsilon_{\pm}$.

Static dissipation on the spin [denoted by Γ in Fig. 8.1(a)] generically leads to a mixed steady state, even if the spin is initialized in a quasienergy state of H_B. Depending on the time t within the period, the static dissipation process primarily excites or de-excites the spin in the quasienergy basis; as both processes can happen, the steady state is mixed.

Our main result is that an additional *static* coupling to an auxiliary system with loss can lead to steady states which are pure and which have high overlap with a chosen quasienergy state of H_B. For specificity, the auxiliary is taken to be a lossy bosonic mode, shown schematically as an optical cavity with photonic dissipation rate κ in Fig. 8.1(a). However, the result holds for more general auxiliary systems that have a characteristic energy scale and relax to a unique state, such as a lossy qubit, or a collection of bosonic modes in an energy band.

The augmented spin-cavity system is governed by the Hamiltonian

$$H_B^+ = \frac{1}{2}\vec{\sigma}\cdot\vec{B}(t) + \Delta\, a^\dagger a + g\,(a^\dagger\sigma^- + a\sigma^+), \tag{8.3}$$

where Δ denotes the effective cavity frequency and g denotes the strength of the spin-cavity coupling. The Hilbert space corresponding to H_B^+ is spanned by the states $|\phi_{\pm}, n\rangle$, where

n is the number of photons in the cavity. The pairs of quasienergy states for $n = 0$ and $n = 1$ are shown in Fig. 8.1(b) along with the couplings induced by the coherent light-matter interaction (blue solid arrows) and the photonic dissipation (black wavy arrows). The hierarchy of energy scales that we assume is:

$$\Gamma < \kappa \lesssim g \ll \Delta, |\epsilon_+ - \epsilon_-|. \tag{8.4}$$

This hierarchy is realized when the driving is strong, and the cavity-QED coupling is in the intermediate to strong coupling regime. The representative quasienergies in Eq. (8.4) depend on the drive frequency regime.

Stabilization is a consequence of resonances between states in the augmented system. Dissipative processes in the spin lead to transitions between $|\phi_\pm, n\rangle$, while photon loss reduces the photon number without affecting the spin state, relaxing the system to $|\phi_\pm, 0\rangle$. However, if the cavity frequency Δ is comparable to the quasienergy difference between $|\phi_\pm\rangle$, then there is a resonance between the $|\phi_+, n-1\rangle$ and $|\phi_-, n\rangle$ states. In the hierarchy of scales in Eq. (8.4), the spin-cavity coupling g is the strongest coupling between these two resonant states. As shown in Fig. 8.1(b), it causes the two states to hybridize, opening up a pathway for $|\phi_+, 0\rangle$ to decay to $|\phi_-, 0\rangle$. As there is no $|\phi_+, -1\rangle$ state, the state $|\phi_-, 0\rangle$ becomes an attractor of the photonic dissipation. The much weaker spin dissipation rate Γ simply perturbs this attractor. The augmented spin-cavity system thus achieves autonomous stabilization to an approximate product state, in which the spin is time-dependent and in a single quasienergy state $|\phi_-(t)\rangle$, while the cavity is close to empty.

8.3 Experimental Observation of Stabilization in the Adiabatic Regime

In this section, we present an experimental implementation of the model in Eq. (8.3) in the adiabatic limit and demonstrate stabilization of a quasienergy state using a transmon qubit [5, 6] coupled to a superconducting cavity. We show that, when the resonance condition in Fig. 8.1(b) holds, the transmon is stabilized to a chosen instantaneous eigenstate for timescales that far exceed the native T_1 and T_2 decoherence times of the qubit.

In the adiabatic limit, the quasienergy states are close to instantaneous eigenstates. The representatives $|\phi_\pm\rangle$ are the states most strongly coupled by the cavity and are several Floquet zones apart; indeed, their quasienergy splitting is close to the instantaneous splitting, $|\epsilon_+ - \epsilon_-| \approx B_0$. From Fig. 8.1(b), the resonance condition is then:

$$\delta = \Delta - B_0 \approx 0. \tag{8.5}$$

We refer to the representative quasienergy states and the instantaneous eigenstates of H_B as $|\phi_\pm\rangle$ below.

To create an *effective* rotating magnetic field of the form in H_B [Eq. (8.1)], we combine simultaneous frequency modulation of the $g - e$ transition of the transmon and strong microwave driving at the mean $g - e$ transition frequency [81]. The flux bias sets the frequency of the $g - e$ transition [5, 6], and therefore creates an effective B_z. On the other hand, the Rabi rate induced by a *constant* AC microwave drive $A\cos(\omega_d t)$ implements an effective *static* B_x term in a frame rotating at the drive frequency ω_d [145]. To synthesize H_B, we first select an operating point using a DC flux bias $\Phi_{\text{ext}} = \Phi_{\text{static}}$, then set the microwave drive frequency to be resonant with the qubit transition at this bias point, and operate in a rotating frame at this frequency. Rotation of the effective field at frequency ω_{mod} is induced by modulating *both* the flux bias [$\Phi_{\text{ext}} = \Phi_{\text{static}} + \delta\Phi\sin(\omega_{\text{mod}} t)$] and the

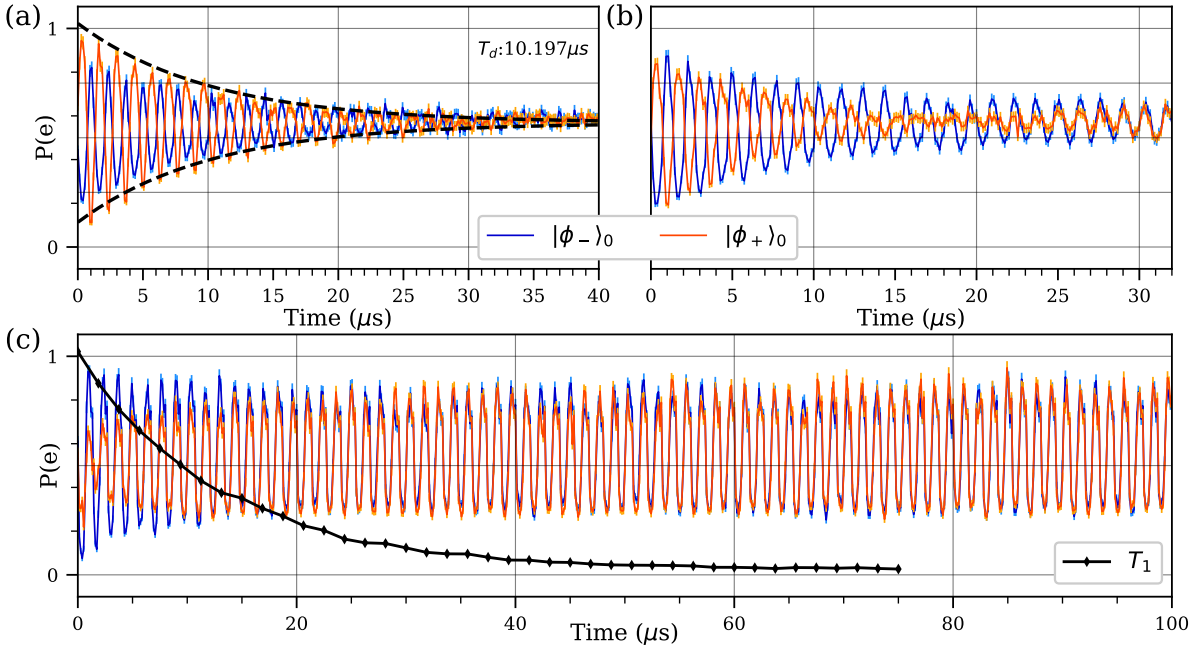

Figure 8.2: **Dissipative stabilization of a quasienergy state in the adiabatic limit.** The panels show the probability of being in the e or $|\uparrow\rangle$ state (obtained from fast ramp-out measurements) after evolution under the Hamiltonian H_B^+ with $B_0 = 2\pi \times 80$ MHz and $\omega_{\text{mod}} = 2\pi \times 0.75$ MHz for three different cavity detunings. The orange (blue) curve is obtained by starting with the qubit aligned (anti-aligned) with the field, which we denote as $|\phi_+\rangle_0$ and $|\phi_-\rangle_0$, respectively. When $g = \kappa = \Gamma = 0$, the curves are sinusoidal with unit amplitude and are exactly out-of-phase. (a) Far from resonance ($\delta = 6.15g = 2\pi \times 80$ MHz), the curves oscillate out-of-phase for all times, but with a decaying amplitude. At long times, both states are close to a 50-50 mixture in the z-basis. Dashed line indicates a fit to the $|\phi_+\rangle_0$ data with extracted time constant $T_d = 10.2$ μs. (b) At intermediate detuning $\delta = 2.31g = 2\pi \times 30$ MHz, the blue trace shows persistent oscillations, while the orange trace first decays to zero, and re-emerges as a curve with the same amplitude and phase as the blue curve. As the steady state contrast is not unity, the stabilized state is a mixture with greater weight on the $|\phi_-\rangle$ state. (c) Closer to resonance, $\delta = 0.77g = 2\pi \times 10$ MHz, the blue curve is the attractor with a high contrast in the steady state, indicating that $|\phi_-\rangle$ is stabilized with high fidelity. The steady-state oscillations persist for far longer than any of the system coherence times, for example, the qubit T_1 decay shown for reference in black. For details on error bar computation, see Section 8.9.3.

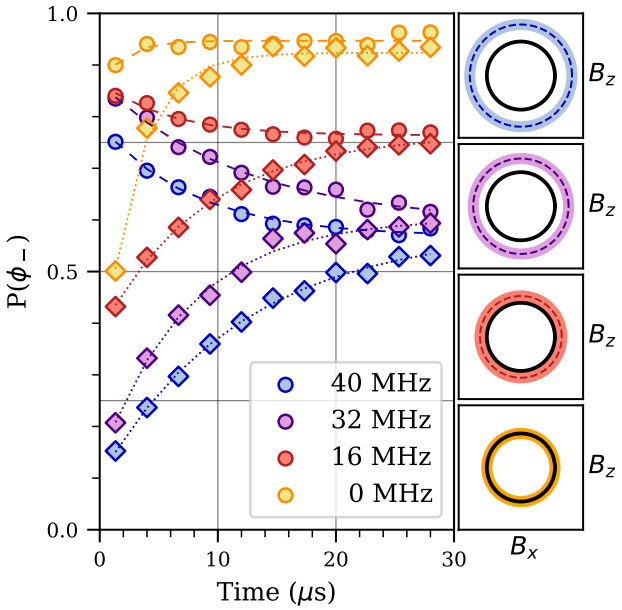

Figure 8.3: **Onset of stabilization in the quasienergy basis in the adiabatic limit.** Measurements of the $|\phi_-\rangle$ population (averaged over two periods) as a function of evolution time. Diamond (circle) markers denote $|\phi_+\rangle$ ($|\phi_-\rangle$) initial state preparation for four different detuning values from the cavity. Curves are labeled by the detuning $\delta/2\pi = (\Delta - B_0)/2\pi$. Side panels show the applied $2\pi \times 80$ MHz rotating magnetic field (solid line) and the $\delta = 0$ ring for each detuning (dashed line). The shaded ring indicates the region of $\pm g = \pm 2\pi \times 13$ MHz around cavity resonance where the qubit states can hybridize strongly with the cavity. As $\delta \to 0$, the system steady state approaches the pure state $|\phi_-\rangle$ with increasing fidelity. Exponential fits to the data are shown in dashed (dotted) lines for preparation in $|\phi_-\rangle$ ($|\phi_+\rangle$).

AC drive [$V_{\text{drive}} = A\cos(\omega_{\text{mod}}t)\cos(\omega_d t)$] in quadrature. See Chapter 6 for more details on the effective rotating field. In order to be deep within the adiabatic regime, we use $B_0 = 2\pi \times 80$ MHz, and modulation frequencies of $\omega_{\text{mod}} = 2\pi \times 0.75$ MHz.

The auxiliary system is a superconducting resonator with effective frequency $\Delta = 2\pi \times 40 - 140$ MHz, linewidth $\kappa = 2\pi \times 84$ kHz, and coupling strength $g = 2\pi \times 13$ MHz to the transmon. The effective frequency Δ is the detuning between the cavity and the *DC frequency set point* of the tunable qubit, and is varied using the latter. The coupling strength g satisfies Eq. (8.4), but is over two orders of magnitude larger than κ. As the time-scale for stabilization is controlled by the photon lifetime κ^{-1}, this allows for a larger time window within which to observe the decay from non-stabilized states.

In principle, decay and dephasing processes in the qubit both contribute to mixing in the Floquet basis. (See Section 8.9.2 for qubit characterization details.) However, the presence of the constant strong drive suppresses the effect of T_2 fluctuations in the qubit frequency, and the qubit follows the drive for $T_d \gg T_2$, even in the absence of stabilization from the cavity. We measure $T_d = 10.2$ μs empirically from the decay of adiabatic following at large cavity detuning, $\delta = 2\pi \times 80$ MHz shown in Fig. 8.2(a). As the lifetime of the cavity $1/\kappa = 1.9$ μs is smaller than T_d, we have the required hierarchy of dissipative scales in Eq. (8.4).

In the adiabatic regime, the resonance condition is alternatively visualized in the $B_x - B_z$ plane. The trajectory traced out by $\vec{B}(t)$ on the plane is a closed curve (black curves, right panels of Fig. 8.3), while the resonance condition $\delta = \Delta - |\vec{B}(t)| = 0$ defines a circle of radius Δ (dotted, right panels of Fig. 8.3). Dissipative stabilization occurs in the vicinity of every intersection point of the two curves. In the simple model in Eq. (8.1), stabilization thus occurs at all times during the protocol, provided $\Delta \approx B_0$. However, this is unnecessary; as long as the trajectory traced out by $\vec{B}(t)$ spends a significant fraction of time within roughly g of the $\delta = 0$ ring, photonic dissipation stabilizes the chosen quasienergy state. The stabilized cavity-state boosting protocol presented in Sec. 8.6 uses this property in an essential manner.

The dynamics of the augmented qubit-cavity system are most transparent if the initial state of the spin coincides with a quasienergy state. In the adiabatic regime, this is an instantaneous eigenstate of H_B at the start of the protocol. Since generating H_B requires strong driving fields, pulsed preparation of the instantaneous eigenstates is challenging, and we instead make use of an adiabatic preparation scheme. An optional π-pulse initializes the qubit in the bare $|g\rangle$ or $|e\rangle$ states, which are then mapped to $|\phi_+\rangle$ and $|\phi_-\rangle$, respectively, by the adiabatic preparation step. After evolution under H_B^+ for a variable amount of time, the final state is mapped back to the bare transmon basis using either a fast ramp out protocol where the fields are shut off as quickly as possible to preserve the instantaneous

qubit populations in the σ_z basis, or by a slow ramp out protocol where the field is rotated to a fixed point on the Bloch sphere to map adiabatic states to the bare qubit states before measurement. See Chapter 6 for details of the full waveforms and the different ramp out protocols.

Both fast ramp-out measurements, shown in Fig. 8.2 and slow ramp-out measurements, shown in Fig. 8.3 display clear signatures of stabilization to the $|\phi_-\rangle$ state when $\delta \approx 0$. When $\delta \gg g$, the differential loss rate between $|\phi_+, 0\rangle$ and $|\phi_-, 1\rangle$ is negligible and the dynamics tends to a (slighly biased) mixed state. However, when $\delta \lesssim g$, significant steady-state coherence in the $|\phi_-\rangle$ state is observed, regardless of whether the initial state is prepared in $|\phi_+\rangle$ or $|\phi_-\rangle$.

Fast ramp-out measurements, such as those shown in Fig. 8.2, approximately measure the population of the state obtained at the end of the evolution under H_B^+ in the σ_z-basis on the qubit. The two instantaneous eigenstates, $|\phi_+\rangle$ and $|\phi_-\rangle$, manifest as $\pm\sin(\omega_{\mathrm{mod}}t)$ oscillations in the probability of qubit excitation after a hold time t. In the far detuned case, the cavity induces negligible preferential loss, and the two states evolve with a π phase shift before eventually decaying to a mixed state due to dissipative processes on the qubit, as shown in Fig. 8.2(a). Near resonance, if the qubit is initially prepared in the $|\phi_-\rangle$ state, it remains close to $|\phi_-\rangle$ indefinitely, leading to persistent oscillations [blue curves in Fig. 8.2(b-c)]. However, qubits initially prepared in $|\phi_+\rangle$ behave qualitatively differently. The initial oscillation damps out, followed by a high-contrast persistent oscillation with the same phase as the $|\phi_-\rangle$ state [orange curves in Fig. 8.2(b-c)]. These steady-state oscillations persist for timescales far exceeding both the atomic and photonic decoherence times of the system. The crossover from an out-of-phase oscillation with the $|\phi_-\rangle$ state to a persistent in-phase oscillation is most clearly visible at intermediate detuning, such as that shown in Fig. 8.2(b) where $\delta \approx 2g$.

The slow ramp-out protocol implements a complementary measurement in the instantaneous eigenbasis of the system. These measurements, shown in Fig. 8.3 provide a

direct measure of the $|\phi_-\rangle$ population at the end of the evolution under H_B^+. Neglecting admixture from the one-photon sector in the cavity and imperfect diabaticity of the ramp out, the value of $\langle\sigma_z\rangle$ obtained from fast ramp-out measurements is approximately $P(\phi_+)\sin(\omega_{\rm mod}t) - P(\phi_-)\sin(\omega_{\rm mod}t)$. The slow ramp out measurements in Fig. 8.3 are therefore related to the contrast of the oscillations in Fig. 8.2. Away from resonance, both initial states decay to a mixed state with a small bias, $P(\phi_-) \approx 1/2$. As $\delta \to 0$, the steady-state approaches a pure state with $P(\phi_-) \approx 1$.

Near resonance, the cavity increases loss during the adiabatic preparation phase, and as a result, the states starting from $|g\rangle$ or $|e\rangle$ end up with significant weight on $|\phi_-\rangle$, even at $t = 0$ (yellow curves in Fig. 8.3). When this weight is near or above 50%, the corresponding fast ramp-out data is not expected to be initially out-of-phase. This is borne out by Fig. 8.2(c).

Section 8.5 presents the detuning dependence of the steady-state $P(\phi_-)$ and the timescale for relaxation to the steady state. There are two notable features. First, in addition to a peak at $\delta = 0$, $P(\phi_-)$ is maximal at $\delta = -B_0/2$ (or $\Delta = B_0/2$). This is a consequence of stabilization by a two photon resonance. Second, peaks in $P(\phi_-)$ go hand-in-hand with dips in the timescale for relaxation to the steady state. This is a consequence of the larger photonic loss rate κ, and not the smaller spin dissipation rate Γ, controlling the decay to the steady state. Both features re-appear in Sec. 8.4.

8.3.1 Onset of stabilization in fast ramp-out measurements

In addition to the slow ramp-out data presented in Fig. 8.3, it is also possible to observe the onset of stabilization in the bare qubit $|g\rangle$, $|e\rangle$ basis using the fast ramp out protocol.

This protocol gives direct access to the time dependence of the instantaneous eigenstates. States prepared in $|\phi_-\rangle$ (using an initial π-pulse before the protocol starts) evolve with a π-phase shift relative to states initialized in $|\phi_+\rangle$ as shown in the blue and orange traces in panels (a.*i*)-(a.*ii*) in Fig. 8.4 respectively. As $\delta = \Delta - B_0$ approaches zero, states

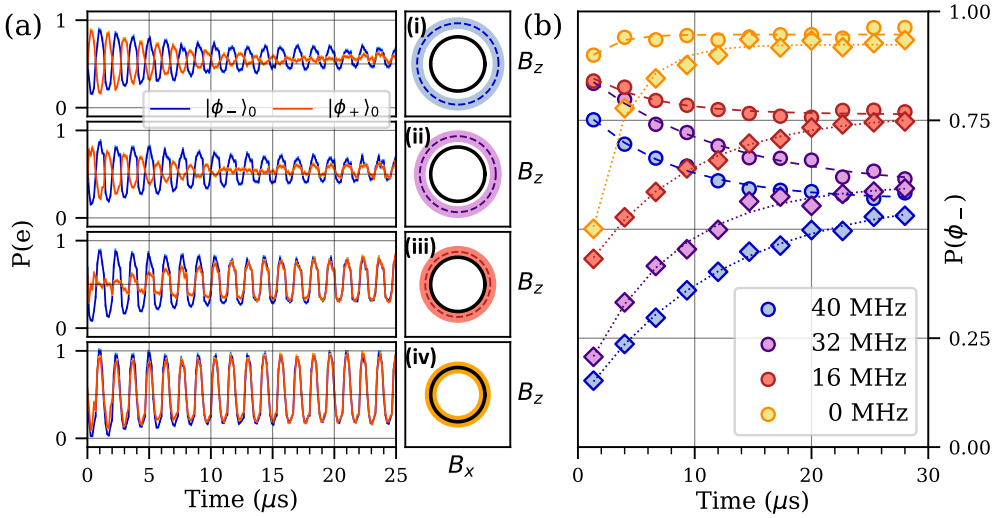

Figure 8.4: **Onset of stabilization.** (a) The probability of being in state $|e\rangle$ as a function of time for different value of minimum detuning to cavity resonance $\delta/2\pi = (\Delta - B_0)/2\pi$ measured using the fast ramp out protocol. The orange (blue) curves are obtained by starting with the qubit aligned (anti-aligned) with the field, which we denote as $|\phi_+\rangle_0$ and $|\phi_-\rangle_0$, respectively. (a.*i*) Far from resonance ($\delta \gg g$), both state initializations damp out to a mixed state while maintaining a π-phase shift. (a.*ii*-a.*iii*) As the detuning δ decreases, the state initially prepared in $|\phi_+\rangle$ damps out. Its signal then grows to a non-zero oscillatory amplitude in phase with that from the $|\phi_-\rangle$ state. (a.*iv*) On resonance, both states are quickly stabilized to $|\phi_-(t)\rangle$. (b) Corresponding data acquired using the slow ramp out protocol, effectively showing the envelope of the oscillatory data in panels (a.*i*-a.*iv*). Curves are labeled by the minimum detuning to cavity resonance $\delta/2\pi$ in MHz. Exponential fits to the data are shown in dashed (dotted) lines for preparation in $|\phi_-\rangle$ ($|\phi_+\rangle$). Insets on the right of (a) show a schematic representation of the resonance condition in the $B_x - B_z$ plane. The solid black line denotes the magnetic field strength and the dashed line denotes the mean qubit-cavity detuning (the colored bands show the $\pm g$ width from the cavity-qubit coupling).

initially prepared in $|\phi_+\rangle$ damp out before being stabilized in the $|\phi_-\rangle$ state, oscillating at high contrast indefinitely. The instantaneous eigenstate data [shown again in panel (b) for convenience], represents the $|\phi_-\rangle$ population for the oscillatory data presented in panels (a.*i*)-(a.*iv*) with 0.5 indicating a perfectly mixed state (i.e. zero oscillation contrast).

The pinch-point, where the oscillation contrast vanishes, observed in the oscillatory data for states prepared initially in $|\phi_+\rangle$ represents the time for the system to reach $P(\phi_-) = 0.5$, not the stabilization time. Due to the adiabatic state preparation time before the evolution under H_B^+ begins and finite ramp out time, configurations near the resonance

condition [as shown in panel (a.*iv*) and in the $\delta = 0$ detuning trace in the adiabatic data in panel (b)] are already mostly in $|\phi_-\rangle$ by the time they are measured even for short evolution times. This is more easily seen in the instantaneous eigenstate data [panel (b)], where the initial value of $P(\phi_-)$ is above 0.5, before quickly growing to roughly 0.9 after a few microseconds.

8.4 Dissipative Stabilization Beyond the Adiabatic Regime

Floquet quasienergy states can be stabilized with static dissipation in any frequency regime. Away from the adiabatic limit, the quasienergy states are no longer close to instantaneous eigenstates. Nevertheless, the mechanism outlined in Sec. 8.2 and in Fig. 8.1(b) stabilizes a target quasienergy state when the right resonance condition is met. In this section, we formalize the heuristic argument in Sec. 8.2, and test its myriad predictions with numerical simulations (Fig. 8.5).

In the strong drive and strong coupling regime [Eq. (8.4)] of the augmented spin-cavity system in Eq. (8.3), a target quasienergy state is stabilized if the resonance condition (approximately) holds for integer m:

$$\Delta = m\omega_{\text{mod}} + \delta\epsilon, \tag{8.6}$$

where $\delta\epsilon = \epsilon_+^{(0)} - \epsilon_-^{(0)}$ is the quasienergy difference in a Floquet zone. Near resonance, the states of the joint system are not close to separable product states. Instead, to leading order in the spin-cavity coupling g, the qubit-cavity system has hybridized quasienergy states

$$|\phi_{\text{qb-cav}}(t)\rangle \approx \frac{1}{\sqrt{2}}\left(\left|\phi_-^{(0)}(t), n+1\right\rangle \pm \left|\phi_+^{(m)}(t), n\right\rangle\right), \tag{8.7}$$

in every photon number sector n, except for one unhybridized state, $\left|\phi_-^{(0)}(t), 0\right\rangle$. Reinstating photon loss, all states in Eq. (8.7) decay (including crucially the states $\left|\phi_+^{(m)}(t), 0\right\rangle$, which

acquire a one-photon character through the hybridization) at a rate set by κ. If κ exceeds the intrinsic qubit decay rate Γ, then this process stabilizes $|\phi_-(t), 0\rangle$ with high fidelity on a time scale set by κ^{-1}.

Figure 8.5 numerically confirms that a representative quasienergy state $|\phi_-(t)\rangle$ is stabilized whenever the resonance conditions are met. We simulate the augmented spin-cavity system with the Hamiltonian in Eq. (8.3) and dissipation using a Floquet-Lindblad master equation. The simulation includes two jump operators: a for photon loss, and σ^- for qubit relaxation to the $|g\rangle$ state. See Section 8.7 for more details. The x-axis varies the frequency of the drive on the qubit, going from adiabatic on the left to high frequency on the right, while the y-axis varies the effective cavity frequency. The color quantifies the period-averaged trace overlap between the steady state of the qubit-cavity system $\rho_{\rm SS}(t)$ and the target quasienergy state:

$$F = \frac{1}{T_{\rm mod}} \int_0^{T_{\rm mod}} \mathrm{d}t \, \mathrm{Tr} \left[\rho_{\rm SS}(t) (|\phi_-(t)\rangle \langle \phi_-(t)| \otimes \mathbb{1}) \right]. \tag{8.8}$$

We immediately see that when the resonance conditions hold for $m = 0, 1, 2$ (dashed lines), the steady state has nearly unit overlap with the target quasienergy states (black color). At generic detunings however, this overlap is close to one-half, as the steady state is close to being maximally mixed between the two quasienergy states of the isolated qubit system. In Section 8.7, we also demonstrate that the timescale for decay to the steady state is shortest, of order κ^{-1}, when the resonance conditions hold. There are several other features visible in Fig. 8.5 that can be simply understood within the resonance-based picture of Fig. 8.1(b). First, the resonance condition need not be perfectly met; the width of the black curves quantifies how close Δ needs to be the dashed curve. This is because the $-$ and $+$ quasienergy states hybridize between number sectors of the cavity as long as the matrix element connecting the states, $|\tilde{H}^{(m0)}_{+-}|$, exceeds their energy difference. A related second observation is that the larger the value of m, the less broad the black curve. That

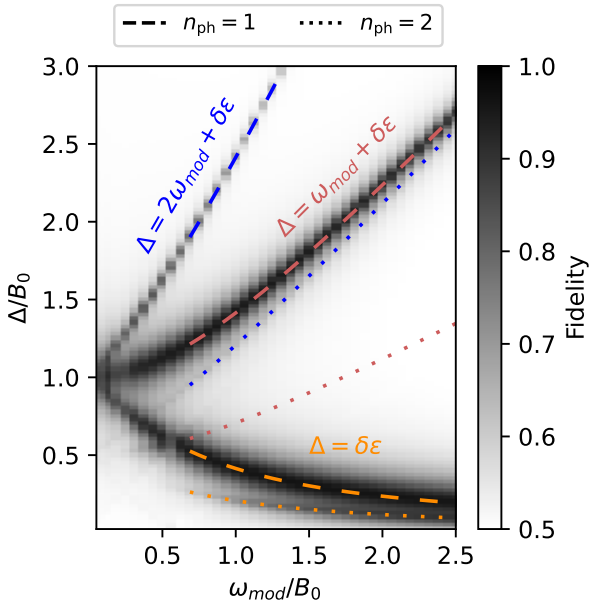

Figure 8.5: **Stabilization to a target quasienergy state in various drive regimes.** The steady state in the driven and dissipative spin-cavity system has nearly unit fidelity with the quasienergy state $|\phi_-\rangle$ [Eq. (8.8)] (indicated by a dark color) when a resonance condition is met [lines, Eq. (8.10)]. The resonances with $m = 0, 1, 2$, and $n_{\rm ph} = 1$ (dashed lines) and $n_{\rm ph} = 2$ (dotted lines) all lead to stabilization, with the $m = 0, n_{\rm ph} = 1$ line being the most robust. *Parameters:* $g/B_0 = 0.05$, $\kappa/B_0 = 0.05$ and $\Gamma = \kappa/20$. See Section 8.7.

is, the resonance condition has to be met more stringently as m increases. This follows from the dependence of $|\tilde{H}^{(m0)}_{+-}|$ on m. The matrix element is given as,

$$\tilde{H}^{(m0)}_{+-} = \frac{1}{T_{\rm mod}} \int_0^{T_{\rm mod}} {\rm d}t \left\langle \phi^{(m)}_+(t), n \right| V \left| \phi^{(0)}_-(t), n+1 \right\rangle , \tag{8.9}$$

which is the time-averaged matrix element of the qubit-cavity coupling $V = g(a^\dagger \sigma^- + a\sigma^+)$ between the quasienergy states in adjacent photon number sector. Taking m to be large causes the integral in Eq. (8.9) to rapidly oscillate and for $|\tilde{H}^{(m0)}_{+-}|$ to become (exponentially) small in m. Further, even when the resonance condition is exactly met, the smallness of $|\tilde{H}^{(m0)}_{+-}|$ takes the system out of the strong coupling regime at large enough m. This changes the rate of decay of the state $|\phi_+, 0\rangle$ to a rate set by Fermi's golden rule, $|\tilde{H}^{(m0)}_{+-}|^2/\kappa$, and significantly slows down the stabilization process. Indeed, for fixed rate of intrinsic spin dissipation Γ, the stabilization would disappear at large enough m as

spin dissipation eventually dominates over the stabilizing process. For the parameters in Fig. 8.5, stabilization disappears for $m \geq 3$ (see Section 8.7 for more details).

A third feature in the figure is that there are faint black lines not given by the resonance condition in Eq. (8.6) (shown as dotted lines). These are also a consequence of dissipative stabilization, but now by *two-photon* resonances. For example, the dotted line in orange is the condition that $|\phi_+, 0\rangle$ and $|\phi_-, 2\rangle$ are resonant. The generalized resonance criterion

$$n_{\text{ph}}\Delta = m\omega_{\text{mod}} + \delta\epsilon, \tag{8.10}$$

with $m \in \mathbb{Z}$ and $n_{\text{ph}} \in \mathbb{Z}^+$ accounts for these faint lines. At first order in g, only those resonances with $n_{\text{ph}} = 1$ are accessible, as the Jaynes-Cummings term couples states which differ by exactly one photon. Indeed, naive perturbation theory suggests that resonances between quasienergy states that differ by n_{ph} photons in the cavity come at order n_{ph} in g. Analogously to the high-order Floquet resonances indexed by m, these higher photon resonances become narrower and and harder to access with increasing n_{ph}. This explains why the numerical simulations find no evidence for stabilization by three photon resonances. Nevertheless, dissipative stabilization from the two-photon resonance is quite robust; we observe stabilization by the two-photon resonance in the experiment as well, see Section 8.5.

We close this section with a discussion of the adiabatic limit in the theoretical framework of this section. Observe that the dashed and dotted lines indicating the one and two photon resonance conditions are not drawn for values below $\omega_{\text{mod}}/B_0 \approx 1$, while the black curves smoothly continue and meet near $\Delta/B_0 \approx 1$. Below $\omega_{\text{mod}}/B_0 \approx 1$, the quasienergies in a single Floquet zone are separated by $\approx \omega_{\text{mod}}$, but the matrix element $|\tilde{H}^{(00)}_{+-}|$ is sufficiently small that the system is not in the strong coupling regime and dissipative stabilization disappears. The matrix element $|\tilde{H}^{(m0)}_{+-}|$ is instead maximal between copies of the

quasienergy states that are nearly B_0 apart in quasienergy, or between Floquet zones separated by $m \approx B_0/\omega_{\text{mod}}$. This reconciles the resonance criteria in Eq. (8.5) and Eq. (8.10), and explains why the darkest black curves meet as $\omega_{\text{mod}}/B_0 \to 0$ at $\Delta/B_0 = 1$. The fainter secondary curves meet at $\Delta/B_0 = 1/2$, corresponding to the two-photon resonance condition in the adiabatic limit.

8.5 Stabilization Rate and Fidelity

In this section, we compare theory with experiment. Specifically, we extract (i) the steady state value of $P(\phi_-)$, and (ii) the stabilization time constant T_{stab} controlling the decay of $P(\phi_+)$, as a function of the detuning δ. Experimental values for both (i) and (ii) are extracted from an exponential fit to the slow ramp out data. For comparison, the dotted lines in Fig. 8.4(b) show the exponential fits for $\delta/2\pi = 0, 16, 32, 40$ MHz. Panels (a) and (b) of Fig. 8.6 show the extracted time constants and steady state $P(\phi_-)$ for detunings spanning in the range of $\pm 2\pi \times 40$ MHz.

The theoretical curves come from a simulation of the Lindblad master equation [21] described in Section 8.7, with jump operators for qubit depolarization, qubit dephasing, and cavity decay. The Hamiltonian parameters of the numerical simulation are set to the experimentally calibrated energy-scales shown in Table 8.1. The strengths of the cavity and qubit decay jump operators are set to $\kappa = 1/T_{\text{cav}}$ and $\Gamma_q = 1/T_1$, using the experimentally measured values shown in Section 8.9.2.

An appropriate treatment of pure dephasing is more challenging because the Lindblad master equation fundamentally relies on a Markov approximation, and, as a result, it only fully captures dephasing in the presence of white noise. The noise spectrum of qubit-frequency fluctuations present in our devices is not white, as evidenced by the Gaussian decay of coherence in Ramsey measurements, shown in Fig. 8.11. Setting the dephasing rate to $1/T_2$, with T_2 obtained from Ramsey measurements, vastly overestimates dephasing

in the presence of the strong drive. We therefore approximate the experiment using an effective pure dephasing rate $1/T_d$, where $T_d = 10.2$ μs is the drive following time shown in Fig. 8.2(a), since this rate more closely captures the effect of the experimental noise spectrum on adiabatic following of the drive field.

Both the experimental and theory traces show a clear feature on resonance ($\Delta - B_0 = 0$) with a peak in the steady state $P(\phi_-)$ accompanied by a steep drop in the time required to reach the steady state. In addition to the single photon resonance between $|\phi_+, 0\rangle$ and $|\phi_-, 1\rangle$, both theory and experiment show and additional two-photon resonance at $\delta = -B_0/2$, where $|\phi_-, 2\rangle$ and $|\phi_+, 0\rangle$ become resonant. This two-photon resonance is a generic feature and persists even outside the adiabatic regime explored in the experiment. Indeed, peaks in the steady state $P(\phi_-)$ associated with two-photon resonances are visible in Fig. 8.5 even at high frequency (indicated by dotted lines).

The theory shows good qualitative agreement with the experimental results. Discrepancies in the exact values of the stabilization time and steady-state $P(\phi_-)$ are likely due to the inability of the Lindblad treatment to capture the colored noise present in our actual experiment. A more complete noise treatment incorporating the experimental noise spectrum is outside the scope of this article.

8.6 Dissipative Stabilization of a Topological Photon Pump

Quantum states can serve as resources for classically difficult tasks, making reliable and stable preparation of these states a useful module in a variety of quantum engineering protocols. The most familiar examples of such resource states are static—a given scheme would call for the preparation of a state $|\psi_{\text{resource}}\rangle$, which does not evolve except when acted on by the protocol [146–150]. However, some schemes explicitly call for time-varying periodic states, $|\psi_{\text{resource}}(t)\rangle$, and utilize this time dependence as part of their dynamical

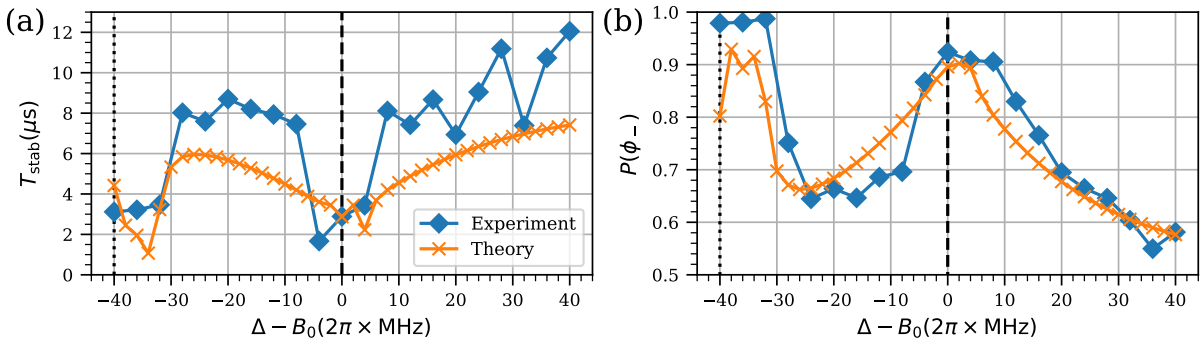

Figure 8.6: **Steady state fidelity and stabilization time in the adiabatic regime.** Extracted stabilization time and steady-state $P(\phi_-)$ as a function of the difference between the detuning Δ of the cavity and the qubit splitting B_0. Blue curves (diamond markers) in panels (a) and (b) show the fit results from an exponential fit to the slow ramp out data, including that shown in Fig. 8.3. On resonance ($\Delta - B_0 = 0$, dashed line), the stabilization time T_{stab} reaches a minimum with corresponding high steady-state $P(\phi_-)$ value. The additional feature at $\Delta - B_0 = -B_0/2$ (dotted line) arises due to a two-photon resonance between $|\phi_-, 2\rangle$ and $|\phi_+, 0\rangle$. Orange curves (x markers) in panels (a) and (b) show the extracted parameters from exponential fits to the Lindblad master equation simulation, which exhibit the same trends as the experimental data.

protocol [78, 79, 81, 84, 109]. These periodic resource states can be stabilized by our methods.

Cavity state boosting [81] is a topological pumping protocol which can be used to coherently translate, or *boost*, the quantum state of a cavity in the Fock basis. This enables the preparation of highly excited non-classical states, such as Fock states, if lower photon number states can be reliably prepared (see Refs. [151–155] for other methods to prepare Fock states). More generally, given an initial cavity state written in terms of Fock states

$$|\psi_0\rangle = \sum_n c_n |n\rangle \,, \tag{8.11}$$

cavity boosting produces evolved states

$$|\psi(MT_{\mathrm{mod}})\rangle \approx \sum_n c_n |n + M\rangle \tag{8.12}$$

at a sequence of predictable times $t = MT_{\text{mod}}$, where T_{mod} is the period of a time-varying resource state and M is an integer.

The required resource state is a strongly varying qubit state that covers the Bloch sphere [78, 79]. When this state is chosen correctly and coupled to the cavity with exchange (Jaynes-Cummings) interactions, the influence of the qubit on the cavity is to effectively implement a Thouless pump in photon number space. On average, every cycle of the qubit state pumps a integer number C of photons into the cavity, leading to an average increase in the photon number with the quantized rate

$$\lim_{t\to\infty} \frac{1}{t} \int_0^t \langle \dot{n}(t') \rangle \mathrm{d}t' = \frac{C}{T_{\text{mod}}}, \tag{8.13}$$

where $\dot{n} = \mathrm{d}n/\mathrm{d}t$ is the average photon number current into the cavity [84]. Henceforth, we take $C = 1$. Further, Ref. [81] showed that at times t^* which are close to being a common period of the cavity and the qubit, $t^* \approx N(2\pi/\Delta_b) \approx MT_{\text{mod}}$, the cavity state rephases into a boosted copy of the initial state [Eq. (8.12)]. The cavity frequency, Δ_b, is defined in Eq. (8.14).

Reference [81] appealed to an adiabatic control scheme for the resource qubit. By implementing a strong, slow drive on the qubit, the back-action of the cavity on the qubit is negligible except when they need to exchange energy. This necessitates an operating regime where the pumping rate—given by $1/T_{\text{mod}}$—is very low, which in turn limits the size of a cavity state that can be prepared in a time- or loss-limited experiment. The introduction of an auxiliary lossy element can stabilize the required resource state, lending additional robustness to the boosting protocol and hence allowing for faster, more reliable preparation of excited cavity states.

More concretely, stabilized boosting can be achieved in a model with two cavities

coupled to a periodically driven qubit,

$$H_{\text{boost}}(t) = \frac{1}{2}\vec{\sigma}\cdot\vec{B}_{\text{boost}}(t) + \Delta_{\text{b}}a_{\text{b}}^{\dagger}a_{\text{b}} + \Delta_{\text{s}}a_{\text{s}}^{\dagger}a_{\text{s}} + g_{\text{b}}(a_{\text{b}}^{\dagger}\sigma^{-} + a_{\text{b}}\sigma^{+}) + g_{\text{s}}(a_{\text{s}}^{\dagger}\sigma^{-} + a_{\text{s}}\sigma^{+}). \quad (8.14)$$

The boost cavity—with frequency Δ_{b}, annihilation operator a_{b}, and coupling g_{b} to the qubit—holds the non-classical state which is to be boosted, and ideally would be lossless. The stabilizing cavity—with label s to differentiate it from the boost cavity—stabilizes the resource qubit. It experiences photon loss with a power decay rate κ_{s} (which we model through Lindblad evolution with a jump operator a_{s}).

We choose a qubit drive

$$\vec{\sigma}\cdot\vec{B}_{\text{boost}}(t) = B_0\left[\max\{0, \sin(\omega_{\text{mod}}t)\}\sigma_x - \cos(\omega_{\text{mod}}t)\sigma_z\right]. \quad (8.15)$$

The vector $\vec{B}_{\text{boost}}(t)$ traces out a semicircle of radius B_0 in the $B_x - B_z$ plane [Fig. 8.7(a)].

Cavity state boosting (and topological photon pumping more broadly [84]) requires the boost cavity energy and drive frequencies to be much smaller than the qubit energy and interaction: $\Delta_{\text{b}}, \omega_{\text{mod}} \ll B_0, g_{\text{b}}\sqrt{n_{\text{b}}}$. Meanwhile, stabilization requires a resonance condition between the qubit quasienergy (which can be replaced by the instantaneous energy near the adiabatic limit) and the stabilizing cavity: $|\Delta_{\text{s}} - B_0| \lesssim g_{\text{s}}$.

In order for the boost cavity to not couple strongly to the stabilizing cavity, and inherit its photon loss, we also demand that they are far detuned: $|\Delta_{\text{s}} - \Delta_{\text{b}}| \gg g_{\text{b}}, g_{\text{s}}$.

Provided these hierarchies of scales are satisfied, there is a simple toy picture for boosting. At the beginning of the protocol ($t = 0$) the cavity is initialized in an accessible Fock state $|n_0\rangle$ and the qubit is prepared in its ground state $|g\rangle$, which is an instantaneous eigenstate of Eq. (8.15). The qubit then adiabatically follows the rotating external field $\vec{B}_{\text{boost}}(t)$ from the negative B_z direction to the positive direction, and ends in the $|e\rangle$ state.

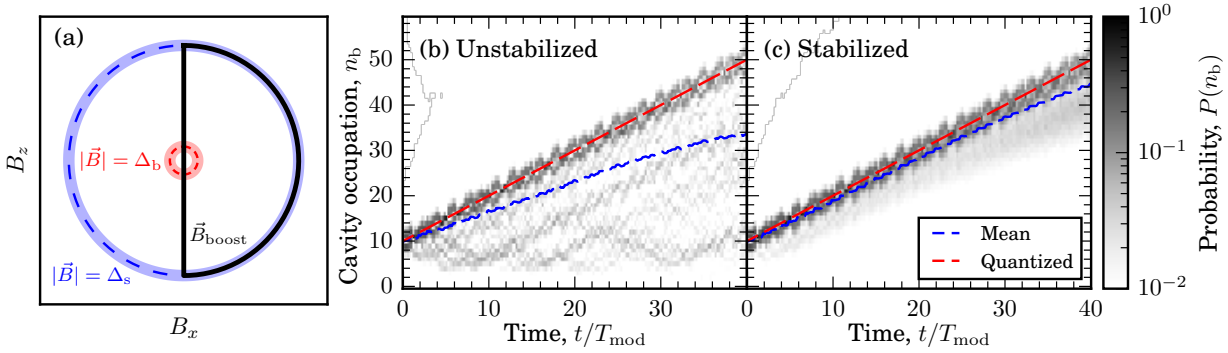

Figure 8.7: **Stabilized cavity state boosting.** (a) The $\vec{B}$ field for the boosting protocol traces out a semicircle of radius B_0 (black) on the $B_x - B_z$ plane with period T_{mod}. On the blue dashed circle, the auxiliary cavity s is resonant with the qubit, while on the red dashed circle, the boost cavity b is resonant with the qubit. (b, c) Plots of the Fock state occupation of the boost cavity ($P(n_{\text{b}}) = \langle n_{\text{b}}| \rho_{\text{b}}(t) |n_{\text{b}}\rangle$ where $\rho_{\text{b}}(t)$ is the reduced density matrix for the boost cavity) without (b) and with (c) the auxiliary cavity as a function of protocol duration. In (b), non-adiabatic excitations of the qubit produce significant population in the cavity at small photon number, indicated by a dark color in the figure. As a result, the expected photon number in the cavity state (blue dashed) is significantly below the quantized value (red dashed). (c) With the auxiliary cavity, non-adiabatic excitations of the qubit are suppressed, resulting in $\langle n_b \rangle$ becoming closer to the quantized value and cavity states that are closer to the desired high photon number Fock state at the special rephasing times t^*. *Parameters:* $B_0/g_{\text{b}} = 20$, $\omega_{\text{mod}}/g_{\text{b}} = 1.5$, $\Delta_{\text{b}}/\omega_{\text{mod}} = (1+\sqrt{5})/2$, $\Delta_{\text{s}}/B_0 = 1$, initial state $|g, n_{\text{b}} = 10, n_{\text{s}} = 0\rangle$; (b) $g_{\text{s}}/g_{\text{b}} = 0$; (c) $g_{\text{s}}/g_{\text{b}} = 1$, photon loss from the lossy cavity with rate $\kappa_{\text{s}}/g_{\text{b}} = 1$ is modeled with a jump operator a_{s} in a Lindblad equation, other dissipative processes are neglected.

As there is a large B_x field on the qubit when B_z crosses Δ_b, the qubit and cavity remain far detuned and unentangled, though the boost cavity state experiences a mean displacement field from the qubit $g_b(a_b^\dagger \langle \sigma^- \rangle + \text{H.c.})$. Then, in the vertical segment of the semicircle protocol, $\vec{B}_{\text{boost}}$ is slowly ramped through an avoided level crossing with the boost cavity, causing an adiabatic transition from the $|e, n\rangle$ state to the state $|g, n+1\rangle$ as the cavity absorbs a qubit excitation. The protocol is then repeated at a fixed rate $1/T_{\text{mod}}$. While the displacement field causes the cavity to deviate from a Fock state at most stroboscopic times, at the special rephasing times MT_{mod} the total displacement averages to zero and an approximate Fock state is recovered.

The introduction of the auxiliary stabilizing cavity suppresses non-adiabatic excitation of the qubit along the curved part of the semicircle protocol through precisely the

same mechanism as Sec. 8.3. Without this, the qubit has a higher probability to flip to the $|g\rangle$ state by the beginning of the vertical segment, and then absorb photons from the cavity by adiabatically following the opposite branch of the avoided crossing, transitioning from $|g, n\rangle$ to $|e, n-1\rangle$.

We numerically verify the quantitative improvement of introducing coupling to the auxiliary lossy cavity in Figs. 8.7(b,c) using a Lindblad master equation simulation. The stabilized boosting protocol produces an average pumping rate which is closer to the expected value of one photon per period [Eq. (8.13)] and a boosted Fock state with a smaller variance, despite the introduction of dissipation. This is shown quantitatively in the following section, where we also discuss details of our numerical simulations.

8.7 Numerical simulations of the augmented spin-cavity and boosting setups

In this section, we describe details of the numerical simulations of driven dissipative dynamics that produced Figs. 8.5 and 8.7. We also present additional data on stabilization rates, the inability of high order resonances to stabilize Floquet states, and the improvement of the topological photon pump when coupled to the auxiliary cavity.

The open cavity-qubit system is modeled within the Floquet Lindblad master equation framework, with jump operators accounting for photon loss and qubit relaxation. Specifically, the equation of motion for the joint cavity-qubit density matrix is

$$\frac{d\rho}{dt} = \mathcal{L}_t[\rho] = -\frac{i}{\hbar}[H_B^+(t), \rho] + \sum_{j=1}^{2} L_j \rho L_j^\dagger - \frac{1}{2}\{L_j^\dagger L_j, \rho\}, \tag{8.16}$$

with $L_1 = \sqrt{\kappa}a$ (the jump operator modeling photon loss), $L_2 = \sqrt{\Gamma}\sigma^-$ (the jump operator modeling qubit loss, or T_1 relaxation), and $\{A, B\} = AB + BA$ being the anticommutator. In Sec. 8.7 we use a similar Lindblad framework in a Hilbert space with two cavities, but

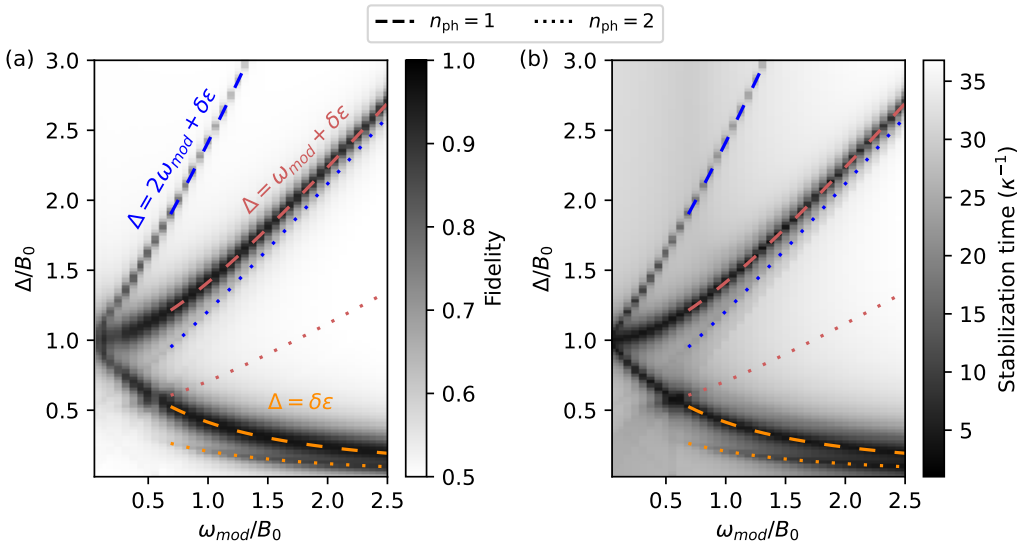

Figure 8.8: **Properties of the steady state of the Lindblad-Floquet master equation.** (a) Reproduced from Fig. 8.5. (b) The time to stabilize into the steady state in units of κ^{-1} (color) at different values of the drive frequency (x-axis) and cavity detunings (y-axis). The high fidelity between the steady state and $|\phi_-\rangle$ in (a) is mirrored in (b) by a short stabilization time, on the order of a few κ^{-1}, of all states to $|\phi_-\rangle$ in (b). When the resonance condition $n_{\rm ph}\Delta = m\omega_{\rm mod} + \delta\epsilon$ is met, the state $\left|\phi_+^{(m)}, 0\right\rangle$ hybridizes with the state $\left|\phi_-^{(0)}, n_{\rm ph}\right\rangle$ and thus decays with a rate set by κ, the cavity linewidth. Far from resonance, the time to reach the steady state is the much longer $1/\Gamma$. *Parameters:* $g/B_0 = 0.05$, $\kappa/B_0 = 0.05$ and $\Gamma = \kappa/20$.

with H_B^+ replaced by $H_{\rm boost}$ and a single jump operator $L_1 = \sqrt{\kappa_{\rm s}} a_{\rm s}$. We solve for $\rho(t)$ given some initial conditions by direct integration in a truncated cavity Hilbert space (we have checked that the state remains far from the truncation throughout the evolution), using the master equation methods of QuTiP [21]. The data displayed in Figs. 8.7, 8.10, and 8.6 is obtained in this way.

Figs. 8.5, 8.8 and 8.9 display features of the driven-dissipative steady state and the approach to it. The steady state is obtained by numerically finding $\rho_\alpha(T_{\rm mod})$ for every member of a basis of initial density matrices $\rho_\alpha(0)$ (labeled by α), and hence constructing a matrix for the superoperator $\mathcal{T}\{e^{\int_0^{T_{\rm mod}} \mathcal{L}_t dt}\}$ in this basis. This is the Floquet superoperator which maps an input density matrix to its time evolution after one period. Diagonalizing this superoperator matrix provides the full spectrum of Floquet modes, and in particular

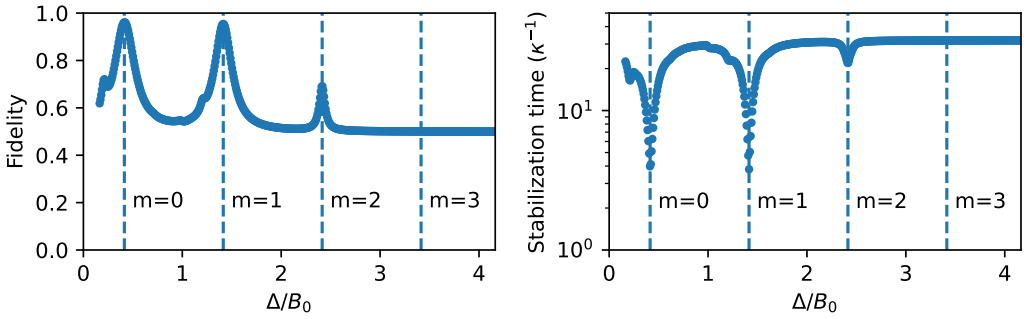

Figure 8.9: **Detuning scan at fixed** $\omega_{\rm mod}/B_0 = 1$**.** Data from Fig. 8.8 for a vertical line cut. The fidelity (time to approach to the steady state) is locally maximum (minimum) when $\Delta = m\omega_{\rm mod} + \delta\epsilon$ for $m = 0, 1, 2$. Stabilization is weak at $m = 2$, and disappears at $m = 3$ because the stabilization rate is smaller than Γ. *Parameters:* $\omega_{\rm mod}/B_0 = 1$, $g/B_0 = 0.05$, $\kappa/B_0 = 0.05$ and $\Gamma = \kappa/20$.

the steady state is identified as the mode with decay rate 0. The behavior of the steady state within a period is found by again integrating the master equation, now using the steady state of the Floquet superoperator as an initial state.

The diagonalization also provides the inverse timescale on which any initial state of the qubit-cavity system relaxes to the steady state, identified as the stabilization rate. This stabilization rate is of the order of κ when the resonance conditions are met, and is otherwise on the scale of Γ, which is significantly smaller (Figs. 8.8, 8.9). As the higher order resonances have smaller stabilization rates, stabilization vanishes for $m \geq 3$, as Fig. 8.9 shows.

Figures 8.7 and 8.10 show simulations of the cavity state boosting protocol obtained from numerically integrating the Lindblad master equation for the total system density matrix. This includes the qubit and two cavities (the boost cavity and stabilizing cavity). The distribution of cavity occupations in the boost cavity is computed as

$$P(n_{\rm b}, t) = {\rm Tr}\left[\rho(t)\,|n_{\rm b}\rangle\langle n_{\rm b}|\right]. \tag{8.17}$$

In the Sec 8.6, it is claimed that the stabilized protocol resulted in a boosted state with lower variance. This is visible in Fig. 8.10, which shows the same data on a larger scale.

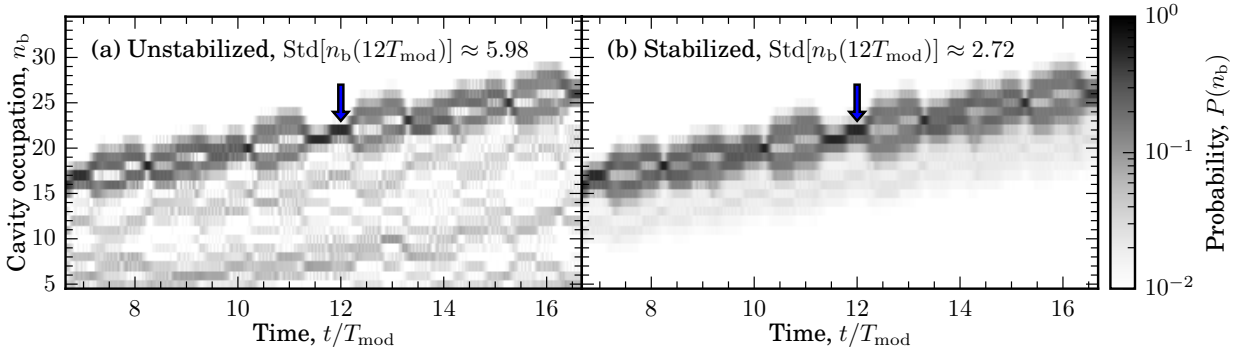

Figure 8.10: **Stabilized cavity state boosting.** Plotting the same data as Figs. 8.7(b-c) on a zoomed-in scale shows the special rephasing times, $t^* = MT_{\text{mod}}$, where the boost cavity becomes close to a Fock state, $\rho_{\text{b}}(t^*) \approx |n_{\text{b}}\rangle\langle n_{\text{b}}|$. One such time is indicated at $M = 12$ by an arrow. The unstabilized protocol (a) shows higher occupation at low photon numbers than the stabilized protocol (b), which results in a standard deviation of n_{b} which is more than a factor of two larger in (a) than in (b) at $t^* = 12T_{\text{mod}}$. *Parameters:* $B_0/g_{\text{b}} = 20$, $\omega_{\text{mod}}/g_{\text{b}} = 1.5$, $\Delta_{\text{b}}/\omega_{\text{mod}} = (1+\sqrt{5})/2$, $\Delta_{\text{s}}/B_0 = 1$, initial state $|g, n_{\text{b}} = 10, n_{\text{s}} = 0\rangle$; (b) $g_{\text{s}}/g_{\text{b}} = 0$; (c) $g_{\text{s}}/g_{\text{b}} = 1$, photon loss from the lossy cavity with rate $\kappa_{\text{s}}/g_{\text{b}} = 1$ is modeled with a jump operator a_{s}.

At, for instance, $t^* = 12T_{\text{mod}}$, the unstabilized protocol produces a state with modal value $n_{\text{b}} = 22$, as does the stabilized protocol. However, the unstabilized state has a standard deviation in n_{b} at this time of $\text{Std}[n_{\text{b}}] \approx 5.98$—more than twice the value of the stabilized state, which has $\text{Std}[n_{\text{b}}] \approx 2.72$. The latter is less than the standard deviation of a coherent state with a mean of 22 photons, $\sqrt{22} \approx 4.69$, so that the stabilized state can outperform a coherent state of the same size in phase estimation measurements [146, 147].

8.8 Conclusion and Outlook

We have demonstrated a simple and fully autonomous scheme to stabilize target states in Floquet systems. To our knowledge, this is the first such scheme which applies *both* with strong driving, for which laboratory-frame ground and excited states are strongly mixed during a Floquet cycle, *and* across all drive frequency regimes, from the low-frequency adiabatic limit all the way up to high drive frequencies exceeding all other energy scales in the problem. Our method relies on the natural static coupling between the Floquet system

of interest and an un-modulated auxiliary system, and on a separation of energy scales which allows a target quasienergy state $|\phi_-(t)\rangle$ to hybridize weakly with the auxiliary compared to all other quasienergy states. Dissipation in the auxiliary then relaxes the combined system toward a steady state with $|\phi_-\rangle$ in the Floquet system and vacuum in the auxiliary.

We have provided combined experimental and theoretical evidence for stabilization in a strongly-driven two-level system coupled to a lossy harmonic oscillator. Experimentally, we study a driven flux-tunable transmon qubit coupled to fixed-frequency microwave resonator, and observe robust steady state stabilization of the $|\phi_-(t)\rangle$ state in the adiabatic regime. Our numerical simulations demonstrate stabilization in all frequency regimes. In addition to this proof-of-principle two-level system example, we numerically studied a two-cavity model in which this stabilization effect improves the performance of a topological photon pump.

The combined experimental and theoretical results demonstrate that the stabilization scheme effectively combats heating processes in Floquet systems which arise due to admixture of excited states with finite decay and dephasing rates. On the other hand, the Hilbert space of a driven two-level system model is too small to exhibit non-trivial heating due to leakage out of the desired operational subspace caused by higher-order drive processes or due to interaction effects with more degrees of freedom, so the results on the phenomenology of H_B^+ do not directly imply robustness against these kinds of heating.

However, there are two aspects of our results that suggest that the same scheme can combat these other kinds of heating. First, our experimental implementation is carried out on a transmon qubit, which is not a true two-level system. In fact, the transmon is *notorious* for being only loosely describable as a two-level system due to its relatively low anharmonicity of $100 - 250$ MHz (or roughly $2\% - 5\%$). The data in Fig. 8.2 shows persistent oscillations for 100 μs and over 100 drive periods without signs of uncompensated heating due to the higher levels of the transmon, despite applied drive strengths which are

roughly one third of the anharmonicity. Second, the stabilized topological photon pump described in Sec. 8.6 uses a second lossy cavity to stabilize a Floquet system consisting of *both* a qubit and a cavity. The presence of autonomous stabilization in both of these cases despite the larger Hilbert space is a strong indicator that it is possible to suppress Floquet heating in general using our simple scheme, rather than the contributions of excited-state lifetime alone. There is also evidence that the scheme can combat heating due to many-body interactions, and can cool into desired ground states of effective Floquet Hamiltonians [136, 156]. However, a full characterization of the capability of autonomous stabilization with a static auxiliary remains an outstanding question.

Optimal stabilization, with the fastest possible stabilization time and highest possible steady-state purity, relies on exact resonances in the augmented spin-cavity system. However, significant stabilization can be achieved for a much wider range of parameters. Supplementary experimental data, shown in Section 8.5, shows that significant stabilization is observed for *all* configurations in the range $|\delta| \lesssim g$. As the transmon decay rate is only a fourth of the photon decay rate, using an even lossier cavity (up to $\kappa \approx g$) would result in stabilization for an even wider range of parameters. Additionally, we find that elliptical drive fields, which are only resonant during part of the drive cycle, also exhibit stabilization. The combination of these two experimental observations is a strong indicator that a single lossy auxiliary degree of freedom can be used to counteract heating into *all* excited states within a range of roughly g of the resonance, and hence systematically remove population from entire targeted Floquet bands. Alternatively, an auxiliary consisting of a band-pass filter, rather than a single cavity, could be used to enhance loss for a large number of Floquet states [137]. Implementations along these lines would give the ability to, for example, cool directly into the ground band of a Floquet engineered lattice [157], rather than relying on adiabatic band mapping, which is generally only possible in synthetic lattice systems.

8.9 Technical Details

8.9.1 Device Fabrication and Design

The device used for the experiment consists of a tunable transmon coupled to two coplanar waveguide (CPW) cavities which we refer to as the readout and main cavities respectively. Table 8.1 shows all the extracted experimental parameters. An on-chip flux-bias line (FBL) along with an external magnet provide AC and DC control of the qubit frequency by changing the flux through the DC SQUID forming the inductive portion of the transmon. Coherent control and readout are performed through the readout cavity, which is far detuned from the operating qubit frequency. The main cavity, much closer in frequency to the qubit operating point, introduces the external loss through hybridization with the qubit.

The device was fabricated on a sapphire substrate with a tantalum (Ta) superconducting ground plane. The qubit paddles and resonators were defined using standard photolithography techniques followed by a wet etch. Finally, the qubit Josephson junctions were fabricated using electron-beam lithography followed by double angle aluminum evaporation and oxidation.

Parameter	Symbol	Value
Qubit g-e frequency	$\omega_q/(2\pi)$	3.9-7.4 GHz
Qubit anharmonicity	$\alpha/(2\pi)$	240 MHz
Qubit-main cavity coupling	$g_m/(2\pi)$	13 MHz
Qubit-readout cavity coupling	$g_r/(2\pi)$	90 MHz
Qubit decay rate	$\Gamma_q/(2\pi)$	13.9 kHz
Readout cavity frequency	$\omega_r/(2\pi)$	7.492 GHz
Readout cavity linewidth	$\kappa_r/(2\pi)$	350 kHz
Main cavity frequency	$\omega_m/(2\pi)$	5.04 GHz
Main cavity linewidth	$\kappa_m/(2\pi)$	84 kHz

Table 8.1: Device Parameters. The qubit and main cavity parameters are acquired with the qubit at 4.9 GHz to isolate the individual decay rates and linewidths. Γ_q is defined as $1/T_1$, the population decay rate of the qubit from the excited state without Purcell loss into the main cavity.

8.9.2 Qubit characterization

We perform standard characterization of the various decoherence timescales in the qubit-cavity system with the qubit at 4.9 GHz, far detuned from the main cavity. Corresponding measurements are shown in Fig. 8.11.

We extract the qubit decay time $T_1 = 11.5$ μs from an exponential fit to the data after initializing the qubit in the $|e\rangle$ and waiting a time τ, with the qubit far detuned from both main cavity. Purcell loss from the main cavity provides a negligible contribution to measurement, so $T_1 = 11.5$ μs, therefore, represents the decay rate of the qubit into channels other than the main cavity.

To find pure dephasing time $T_\phi = 570$ ns, we perform a Ramsey-style measurement where two $\pi/2$ pulses are applied to the qubit with a variable wait time. However, to maintain a large separation of scales between the oscillatory and decaying time constants of the Ramsey signal, we rotate the phase of the second pulse at 5 MHz to apply a virtual detuning between the drive and qubit without loss of contrast. We use a gaussian envelope to compute the dephasing time [5] instead of the regular exponential envelope, due to the low-frequency noise present in tunable frequency qubits away from a sweet spot. We extract $T_{\text{echo}} = 3$ μs from the standard Hahn echo sequence and fit an exponential model to the data.

Finally, the cavity ringdown time, $T_{\text{cav}} = 3.8$ μs, is measured by populating the cavity with a coherent state and fitting the resulting leakage signal after turning off the cavity drive. Since heterodyne detection measures the cavity field *amplitude*, we fit to a damped sinusoidal response. The photon lifetime in the cavity is *half* the ringdown time as it scales with the power in the cavity instead of the field amplitude giving a final lifetime of 1.9 μs.

8.9.3 Experimental error bars and qubit state readout

The qubit state is read out using standard two-tone pulsed spectroscopy on the readout cavity [5, 6]. In this method, $|g\rangle$ and $|e\rangle$ correspond to two different resonant frequencies

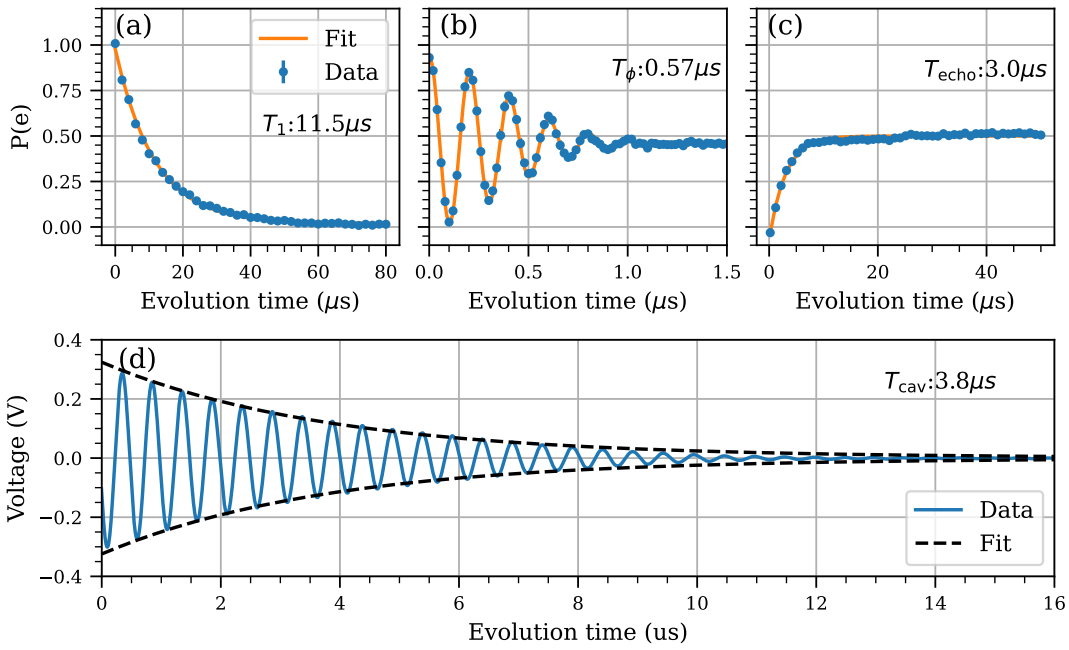

Figure 8.11: **System coherence times.** (a) $|e\rangle$ state population as a function of delay time after qubit excitation. An exponential fit gives $T_1 = 11.5$ μs. (b) Ramsey oscillations using a 5 MHz virtual detuning on the second $\pi/2$-pulse to extract $T_\phi = 570$ ns. (c) $|e\rangle$-state population after a standard Hahn echo sequence, giving $T_{\text{echo}} = 3$ μs. The six-fold increase between T_ϕ and T_{echo} indicates the presence of a large low-frequency noise component in the qubit frequency. (d) Heterodyne measurement of the main cavity field after populating it with a coherent state. The field decay time $T_{\text{cav}} = 3.8$ μs is extracted from a fit to a damped oscillatory response (envelope of the fit is indicated by dashed lines). The photon lifetime corresponds to $T_{\text{cav}}/2$. Presented data is the average of 10 repeated traces, each consisting of 10^4 shots per time point. Error bars are computed as the standard error of the mean of the repeated traces. The error bars are smaller than the marker size for panels (a), (b), (c) and not shown for panel (d).

of the readout cavity. The qubit state can be inferred by measuring the amplitude and phase of a microwave pulse at the cavity frequency, which can be converted to IQ points (the cartesian representation of the sine and cosine components of the cavity response). By taking reference $|g\rangle$ and $|e\rangle$ measurements, we can convert an output voltage to a normalized qubit population:

$$P(e) = \frac{\overrightarrow{IQ}_{\mathrm{data-g}} \cdot \overrightarrow{IQ}_{e-g}}{|\overrightarrow{IQ}_{e-g}|^2}, \tag{8.18}$$

where $\overrightarrow{IQ}_{\mathrm{data-g}}$ is vector-valued difference between the resonator response for the given configuration and the response for $|g\rangle$. The quantity $\overrightarrow{IQ}_{\mathrm{e-g}}$ is the analogous quantity for a reference measurement of a qubit in the $|e\rangle$ state.

For all data in the text, the reference $|e\rangle$ trace was acquired at the beginning of a measurement run (i.e. for a given detuning) by averaging 5×10^4 single shots, and the $|g\rangle$ reference trace was acquired every 10 measurement points as the average of 10^4 shots.

Data points was acquired by averaging 10^4 shots at a 5 kHz repetition rate to avoid spurious $|e\rangle$ population between experimental runs. Due to different readout amplitudes for $|g\rangle$ and $|e\rangle$ outcomes, the measurement uncertainties are not constant and need to be characterized as a function $P(e)$.

We estimate the typical readout error for a given $|g\rangle$-$|e\rangle$ superposition by interpolating from a set of reference uncertainty measurements. These reference uncertainties are computed by taking the standard deviation of 10 repeated measurements of one period of oscillation. All the data presented uses this method as a typical readout uncertainty instead of computing an error bar within each dataset.

The qubit characterization data, shown in Fig. 8.11, is the average of 10 repeated measurements, each consisting of the average of 10^4 single shots (which are not individually recorded). The uncertainty for these measurements is computed as the standard error of the mean ($\mathrm{Std}/\sqrt{N-1}$ with $N = 10$) and is typically smaller than the marker size.

Chapter 9

Conclusion and Outlook

In this thesis, we built up the required hardware and software infrastructure for superconducting qubit experiments. Early work focused on developing a robust code infrastructure for measurements starting with standard instrument drivers and prescribed behaviors as a function of instrument class, an experimental measurement template, and a centralized pulse scheduling object for coherent measurements. In parallel, we also implemented standard qubit characterization measurements and developed a series of calibration procedures to quickly shift the qubit frequency and recalibrate the measurement configuration.

With the measurement infrastructure set up, we explored two main research directions: firstly, a tunable qubit-resonator coupler, and secondly, a topological photon pump using rotating synthetic magnetic fields.

The tunable coupler was ultimately unsuccessful due to large residual capacitive coupling in addition to the tunable inductive coupling which reduced the dynamic tuning range. However, in this project we developed two important technical tools used in other aspects of the lab: the crosstalk calibration procedure used to individually tune the qubit and coupler and the frequency tunable cavity. In addition, we proposed and simulated a circuit model which achieves strong coupling without shifting the resonator frequency due to inductive loading.

For the topological pumping project, we developed the tools to generate a large synthetic rotating field acting on a transmon. This field consists of simultaneous frequency modulation on the qubit and amplitude modulation on a microwave drive at the qubit frequency. In Chapter 7, we describe the procedure to calibrate and synchronize the two field components and demonstrate a large experimental parameter range of achievable fields, with field amplitudes reaching 100 MHz and a 5 MHz rotation rate. We characterize the rotating field and the degree to which the qubit follows the applied field using both measurements in the bare qubit basis as well as measuring in the instantaneous eigenbasis of the rotating field (i.e. measuring whether the qubit is aligned or anti-aligned with the field at any given time).

In the process of characterizing the field for the topological pumping project, we made a genuine experimental discovery: by driving the qubit with the rotating field near (in frequency) a lossy system, we could *indefinitely stabilize* either of the instantaneous eigenstates of the field. While the experimental regime of the rotating field was deeply adiabatic (i.e. the rotation rate was much slower than all the other frequencies in the system), we developed a generalized framework (with our theory collaborators), extending the stabilization to the high-frequency regime. This effect was particularly surprising as it was originally thought that Floquet states (eigenstates of a periodically driven system) were adversely affected by dissipation.

The goal of this thesis is to serve as a springboard for the next generation of students in the lab through the built up infrastructure and especially the synthetic fields developed in the second half of this work.

9.1 Future Work and Outlook

9.1.1 Miniaturized tunable coupler

While the tunable coupler prototype developed in Chapter 5 demonstrated a proof of principle device for tunable qubit-resonator coupling, the large physical footprint required to implement it makes it impractical for use in a photonic lattice. As shown in related work in the lab, [10, 58], the available space for a qubit is very limited for practical implementations of a photonic lattice. Going from a floating qubit design (composed of a pair of metal islands connected by a Josephson junction) to a grounded design (similar to the Xmon [69]) would provide some improvements for the layout constraints. Similarly, the coupling strength can be increased while reducing the physical footprint of the coupler by replacing the floating loop with a current divider circuit as explored in [41]. However, this would require a new approach the reduce the on-diagonal shift on the resonator mode.

9.1.2 Boosting measurement

The topological pump in Chapter 7 has been shown theoretically to be able to translate photon number states in the Fock basis [81]. These boosted Fock states can then be, in principle, as resources for quantum algorithms or sensing schemes. Experimentally, we developed the tools to generate the required fields and thoroughly characterized dynamics of the driven qubit. The existing measurement configuration proved to be unable to directly measure the photon pump due to excessive noise on the output. Using a quantum limited amplifier on the output of the pumping cavity would drastically improve the signal to noise and provide a pathway to measuring the pumping rate directly. The current measurement scheme measures the amplitude of the electric field in the cavity instead of the photon number directly. By measuring the photon number directly, we become less susceptible to frequency variations in the cavity resonance which should also improve the readout.

9.1.3 Non-reciprocal devices

The non-trivial synthetic bands arising in the topological pump in Chapter 7 can also be used to develop new non-reciprocal devices such as circulators and isolators. One promising approach is to add a third drive on the qubit to increase the dimensionality of the system. By arranging the direction of the pump between any pair of drives (pumping from drive A to B or vice versa), we can synthesize an effective chiral circulation of energy between the three drives. As we discussed in Chapter 4, these non-reciprocal devices are a key building block for superconducting measurements, in particular to shield the quantum device from back scatter off of amplifiers used for measurement. Current wideband, commercial devices rely on a ferrite bead for isolation requiring bulky shielding to protect the qubits from the stray magnetic field.

Bibliography

1. Kielpinski, D., Monroe, C. & Wineland, D. J. Architecture for a large-scale ion-trap quantum computer. *Nature* **417,** 709–711. ISSN: 1476-4687. https://doi.org/10.1038/nature00784 (June 2002).
2. Blais, A., Huang, R.-S., Wallraff, A., Girvin, S. M. & Schoelkopf, R. J. Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation. *Phys. Rev. A* **69,** 062320 (June 2004).
3. Bluvstein, D. *et al.* Logical quantum processor based on reconfigurable atom arrays. *Nature* **626,** 58–65. ISSN: 1476-4687. https://doi.org/10.1038/s41586-023-06927-3 (Feb. 2024).
4. Shor, P. W. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Journal on Computing* **26,** 1484–1509. eprint: https://doi.org/10.1137/S0097539795293172. https://doi.org/10.1137/S0097539795293172 (1997).
5. Krantz, P. *et al.* A quantum engineer's guide to superconducting qubits. *Applied Physics Reviews* **6,** 021318. ISSN: 1931-9401. eprint: https://pubs.aip.org/aip/apr/article-pdf/doi/10.1063/1.5089550/16667201/021318_1_online.pdf. https://doi.org/10.1063/1.5089550 (June 2019).
6. Blais, A., Grimsmo, A. L., Girvin, S. M. & Wallraff, A. Circuit quantum electrodynamics. *Reviews of Modern Physics* **93.** ISSN: 1539-0756. http://dx.doi.org/10.1103/RevModPhys.93.025005 (May 2021).
7. Koch, J. *et al.* Charge-insensitive qubit design derived from the Cooper pair box. *Phys. Rev. A* **76,** 042319. https://link.aps.org/doi/10.1103/PhysRevA.76.042319 (4 Oct. 2007).
8. Manucharyan, V. *Superinductance* PhD thesis (Yale University, 2012).
9. Gyenis, A. *et al.* Experimental Realization of a Protected Superconducting Circuit Derived from the 0–π Qubit. *PRX Quantum* **2,** 010339. https://link.aps.org/doi/10.1103/PRXQuantum.2.010339 (1 Mar. 2021).
10. Amouzegar, M. *Photon-Mediated Interactions in Lattices of Coplanar Waveguide Resonators* PhD thesis (University of Maryland, 2024).
11. Vool, U. & Devoret, M. Introduction to quantum electromagnetic circuits. *International Journal of Circuit Theory and Applications* **45,** 897–934. ISSN: 1097-007X. http://dx.doi.org/10.1002/cta.2359 (June 2017).

12. Devoret, M., Huard, B., Schoelkopf, R. & Cugliandolo, L. F. *Quantum Machines: Measurement and Control of Engineered Quantum Systems: Lecture Notes of the Les Houches Summer School: Volume 96, July 2011* ISBN: 9780199681181. https://doi.org/10.1093/acprof:oso/9780199681181.001.0001 (Oxford University Press, June 2014).

13. Jaynes, E. & Cummings, F. Comparison of quantum and semiclassical radiation theories with application to the beam maser. *Proceedings of the IEEE* **51,** 89–109 (1963).

14. Purcell, E. M. Proceedings of the American Physical Society. *Phys. Rev.* **69,** 681. https://link.aps.org/doi/10.1103/PhysRev.69.674.2 (11-12 June 1946).

15. Kleppner, D. Inhibited Spontaneous Emission. *Phys. Rev. Lett.* **47,** 233–236. https://link.aps.org/doi/10.1103/PhysRevLett.47.233 (4 July 1981).

16. Houck, A. A. *et al.* Controlling the Spontaneous Emission of a Superconducting Transmon Qubit. *Phys. Rev. Lett.* **101,** 080502. https://link.aps.org/doi/10.1103/PhysRevLett.101.080502 (8 Aug. 2008).

17. Wallraff, A. *et al.* Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics. *Nature* **431,** 162–167. ISSN: 1476-4687. https://doi.org/10.1038/nature02851 (Sept. 2004).

18. Schuster, D. I. *et al.* Resolving photon number states in a superconducting circuit. *Nature* **445,** 515–518. ISSN: 1476-4687. https://doi.org/10.1038/nature05461 (Feb. 2007).

19. Heeres, R. W. *et al.* Cavity State Manipulation Using Photon-Number Selective Phase Gates. *Phys. Rev. Lett.* **115,** 137002. https://link.aps.org/doi/10.1103/PhysRevLett.115.137002 (13 Sept. 2015).

20. Chou, K. S. *et al.* A superconducting dual-rail cavity qubit with erasure-detected logical measurements. *Nature Physics* **20,** 1454–1460. ISSN: 1745-2481. https://doi.org/10.1038/s41567-024-02539-4 (Sept. 2024).

21. Johansson, J., Nation, P. & Nori, F. QuTiP 2: A Python framework for the dynamics of open quantum systems. *Computer Physics Communications* **184,** 1234–1240. ISSN: 0010-4655. https://www.sciencedirect.com/science/article/pii/S0010465512003955 (2013).

22. Chitta, S. P., Zhao, T., Huang, Z., Mondragon-Shem, I. & Koch, J. *Computer-aided quantization and numerical analysis of superconducting circuits* 2022. arXiv: 2206.08320 [quant-ph]. https://arxiv.org/abs/2206.08320.

23. Groszkowski, P. & Koch, J. Scqubits: a Python package for superconducting qubits. *Quantum* **5,** 583. ISSN: 2521-327X. http://dx.doi.org/10.22331/q-2021-11-17-583 (Nov. 2021).

24. Cohen-Tannoudji, C., Dupont-Roc, J. & Grynberg, G. in *Atom—Photon Interactions* 257–351 (John Wiley & Sons, Ltd, 1998). ISBN: 9783527617197. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/9783527617197.ch4. https://onlinelibrary.wiley.com/doi/abs/10.1002/9783527617197.ch4.

25. Steck, D. *Quantum and Atom Optics* Revision 0.8.3. http://steck.us/teaching.

26. Campaioli, F., Cole, J. H. & Hapuarachchi, H. Quantum Master Equations: Tips and Tricks for Quantum Optics, Quantum Computing, and Beyond. *PRX Quantum* **5,** 020202. https://link.aps.org/doi/10.1103/PRXQuantum.5.020202 (2 June 2024).

27. Göppl, M. *et al.* Coplanar waveguide resonators for circuit quantum electrodynamics. *Journal of Applied Physics* **104,** 113904. ISSN: 0021-8979. eprint: https://pubs.aip.org/aip/jap/article-pdf/doi/10.1063/1.3010859/14109820/113904_1_online.pdf. https://doi.org/10.1063/1.3010859 (Dec. 2008).

28. Watanabe, K., Yoshida, K. & Kohjiro, T. A. Kinetic Inductance of Superconducting Coplanar Waveguides. *Japanese Journal of Applied Physics* **33,** 5708. https://dx.doi.org/10.1143/JJAP.33.5708 (Oct. 1994).

29. Kollár, A. J., Fitzpatrick, M. & Houck, A. A. Hyperbolic lattices in circuit quantum electrodynamics. *Nature* **571,** 45–50. ISSN: 1476-4687. https://doi.org/10.1038/s41586-019-1348-3 (July 2019).

30. Pozar, D. M. *Microwave engineering; 3rd ed.* https://cds.cern.ch/record/882338 (Wiley, Hoboken, NJ, 2005).

31. DiVincenzo, D. P. The Physical Implementation of Quantum Computation. *Fortschritte der Physik* **48,** 771–783. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/1521-3978%28200009%2948%3A9/11%3C771%3A%3AAID-PROP771%3E3.0.CO%3B2-E. https://onlinelibrary.wiley.com/doi/abs/10.1002/1521-3978%28200009%2948%3A9/11%3C771%3A%3AAID-PROP771%3E3.0.CO%3B2-E (2000).

32. Josephson, B. Possible new effects in superconductive tunnelling. *Physics Letters* **1,** 251–253. ISSN: 0031-9163. https://www.sciencedirect.com/science/article/pii/0031916362913690 (1962).

33. Josephson, B. D. The discovery of tunnelling supercurrents. *Rev. Mod. Phys.* **46,** 251–254. https://link.aps.org/doi/10.1103/RevModPhys.46.251 (2 Apr. 1974).

34. Tinkham, M. *Introduction to superconductivity* English. Mineola, N.Y., 1996. http://www.freading.com/ebooks/details/r:download/ZnJlYWQ6OTc4MDQ4NjEzNDcyNzpl.

35. Place, A. P. M. *et al.* New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds. *Nature Communications* **12.** ISSN: 2041-1723. http://dx.doi.org/10.1038/s41467-021-22030-5 (Mar. 2021).

36. Bylander, J. *et al.* Noise spectroscopy through dynamical decoupling with a superconducting flux qubit. *Nature Physics* **7,** 565–570. ISSN: 1745-2481. https://doi.org/10.1038/nphys1994 (July 2011).

37. DiCarlo, L. *et al.* Demonstration of two-qubit algorithms with a superconducting quantum processor. *Nature* **460,** 240–244. ISSN: 1476-4687. https://doi.org/10.1038/nature08121 (July 2009).

38. Barends, R. *et al.* Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature* **508,** 500–503. ISSN: 1476-4687. https://doi.org/10.1038/nature13171 (Apr. 2014).

39. Rol, M. A. *et al.* Fast, High-Fidelity Conditional-Phase Gate Exploiting Leakage Interference in Weakly Anharmonic Superconducting Qubits. *Phys. Rev. Lett.* **123,** 120502. https://link.aps.org/doi/10.1103/PhysRevLett.123.120502 (12 Sept. 2019).

40. Srinivasan, S. J., Hoffman, A. J., Gambetta, J. M. & Houck, A. A. Tunable Coupling in Circuit Quantum Electrodynamics Using a Superconducting Charge Qubit with a V-Shaped Energy Level Diagram. *Phys. Rev. Lett.* **106,** 083601. https://link.aps.org/doi/10.1103/PhysRevLett.106.083601 (8 Feb. 2011).

41. Chen, Y. *et al.* Qubit Architecture with High Coherence and Fast Tunable Coupling. *Phys. Rev. Lett.* **113,** 220502. https://link.aps.org/doi/10.1103/PhysRevLett.113.220502 (22 Nov. 2014).

42. Horowitz, P. & Hill, W. *The Art of Electronics* (Cambridge University Press, 1980).

43. Carr, H. Y. & Purcell, E. M. Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments. *Phys. Rev.* **94,** 630–638. https://link.aps.org/doi/10.1103/PhysRev.94.630 (3 May 1954).

44. Meiboom, S. & Gill, D. Modified Spin-Echo Method for Measuring Nuclear Relaxation Times. *Review of Scientific Instruments* **29,** 688–691. ISSN: 0034-6748. eprint: https://pubs.aip.org/aip/rsi/article-pdf/29/8/688/19287064/688_1_online.pdf. https://doi.org/10.1063/1.1716296 (Aug. 1958).

45. Steffen, Z. *Characterization of Gap-Engineered Josephson Junctions and Gate Fidelities for a Superconducting Qubit* PhD thesis (University of Maryland, 2024).

46. Motzoi, F., Gambetta, J. M., Rebentrost, P. & Wilhelm, F. K. Simple Pulses for Elimination of Leakage in Weakly Nonlinear Qubits. *Phys. Rev. Lett.* **103,** 110501. https://link.aps.org/doi/10.1103/PhysRevLett.103.110501 (11 Sept. 2009).

47. Gambetta, J. M., Motzoi, F., Merkel, S. T. & Wilhelm, F. K. Analytic control methods for high-fidelity unitary operations in a weakly nonlinear oscillator. *Phys. Rev. A* **83,** 012308. https://link.aps.org/doi/10.1103/PhysRevA.83.012308 (1 Jan. 2011).

48. Chen, W., Abbasi, M., Joglekar, Y. N. & Murch, K. W. Quantum Jumps in the Non-Hermitian Dynamics of a Superconducting Qubit. *Phys. Rev. Lett.* **127,** 140504. https://link.aps.org/doi/10.1103/PhysRevLett.127.140504 (14 Sept. 2021).

49. Catelani, G., Schoelkopf, R. J., Devoret, M. H. & Glazman, L. I. Relaxation and frequency shifts induced by quasiparticles in superconducting qubits. *Phys. Rev. B* **84,** 064517. https://link.aps.org/doi/10.1103/PhysRevB.84.064517 (6 Aug. 2011).

50. Martinis, J. M. *et al.* Decoherence in Josephson Qubits from Dielectric Loss. *Phys. Rev. Lett.* **95,** 210503. https://link.aps.org/doi/10.1103/PhysRevLett.95.210503 (21 Nov. 2005).

51. Krinner, S. *et al.* Engineering cryogenic setups for 100-qubit scale superconducting circuit systems. *EPJ Quantum Technology* **6,** 2. ISSN: 2196-0763. https://doi.org/10.1140/epjqt/s40507-019-0072-0 (May 2019).

52. Stefanazzi, L. *et al.* The QICK (Quantum Instrumentation Control Kit): Readout and control for qubits and detectors. *Review of Scientific Instruments* **93,** 044709. ISSN: 0034-6748. eprint: https://pubs.aip.org/aip/rsi/article-pdf/doi/10.1063/5.0076249/19817152/044709_1_online.pdf. https://doi.org/10.1063/5.0076249 (Apr. 2022).

53. Ritter, M. *et al. A strong-driving synthetic magnetic field toolkit for topological Floquet models with superconducting circuits* In preparation.

54. *VPP-4.3: The VISA library* Rev. 7.2.1. IVI Foundation (2024).

55. *SCPI Syntax and Style* 1999.0. SCPI Consortium (May 1999).

56. Girvin, S. M. & Yang, K. *Modern Condensed Matter Physics* (Cambridge University Press, 2019).

57. Ashcroft, N. W. & Mermin, D. N. *Solid State Physics* (Cengage Learning, 1976).

58. O'Brien, K., Amouzegar, M., Lee, W. C., Ritter, M. & Kollár, A. J. *A Superconducting-Circuit Waveguide QED System with Unconventional Band Structures* In preparation.

59. Underwood, D. L., Shanks, W. E., Koch, J. & Houck, A. A. Low-disorder microwave cavity lattices for quantum simulation with photons. *Phys. Rev. A* **86,** 023837. https://link.aps.org/doi/10.1103/PhysRevA.86.023837 (2 Aug. 2012).

60. Carusotto, I. *et al.* Photonic materials in circuit quantum electrodynamics. *Nature Physics* **16,** 268–279. ISSN: 1745-2481. https://doi.org/10.1038/s41567-020-0815-y (Mar. 2020).

61. Houck, A. A., Türeci, H. E. & Koch, J. On-chip quantum simulation with superconducting circuits. *Nature Physics* **8,** 292–299. ISSN: 1745-2481. https://doi.org/10.1038/nphys2251 (Apr. 2012).

62. Kollár, A. J., Fitzpatrick, M., Sarnak, P. & Houck, A. A. Line-Graph Lattices: Euclidean and Non-Euclidean Flat Bands, and Implementations in Circuit Quantum Electrodynamics. *Communications in Mathematical Physics* **376,** 1909–1956. ISSN: 1432-0916. https://doi.org/10.1007/s00220-019-03645-8 (June 2020).

63. Chiu, C. S., Ma, D.-S., Song, Z.-D., Bernevig, B. A. & Houck, A. A. Fragile topology in line-graph lattices with two, three, or four gapped flat bands. *Phys. Rev. Res.* **2,** 043414. https://link.aps.org/doi/10.1103/PhysRevResearch.2.043414 (4 Dec. 2020).

64. John, S. & Wang, J. Quantum electrodynamics near a photonic band gap: Photon bound states and dressed atoms. *Phys. Rev. Lett.* **64,** 2418–2421. https://link.aps.org/doi/10.1103/PhysRevLett.64.2418 (20 May 1990).

65. Wulschner, F. *et al.* Tunable coupling of transmission-line microwave resonators mediated by an rf SQUID. *EPJ Quantum Technology* **3,** 10. ISSN: 2196-0763. https://doi.org/10.1140/epjqt/s40507-016-0048-2 (July 2016).

66. Reagor, M. *et al.* Demonstration of universal parametric entangling gates on a multi-qubit lattice. *Science Advances* **4,** eaao3603. eprint: https://www.science.org/doi/pdf/10.1126/sciadv.aao3603. https://www.science.org/doi/abs/10.1126/sciadv.aao3603 (2018).

67. Yan, F. *et al.* Tunable Coupling Scheme for Implementing High-Fidelity Two-Qubit Gates. *Phys. Rev. Appl.* **10,** 054062. https://link.aps.org/doi/10.1103/PhysRevApplied.10.054062 (5 Nov. 2018).

68. Li, H.-K. *et al.* Tunable coupling between Xmon qubit and coplanar waveguide resonator*. *Chinese Physics B* **28,** 080305. https://dx.doi.org/10.1088/1674-1056/28/8/080305 (Aug. 2019).

69. Barends, R. *et al.* Coherent Josephson Qubit Suitable for Scalable Quantum Integrated Circuits. *Phys. Rev. Lett.* **111,** 080502. https://link.aps.org/doi/10.1103/PhysRevLett.111.080502 (8 Aug. 2013).

70. Chow, J. *Quantum Information Processing with Superconducting Qubits* PhD thesis (Yale University, 2010).

71. Ritter, M., Long, D. M., Yue, Q., Chandran, A. & Kollár, A. J. *Autonomous Stabilization of Floquet States Using Static Dissipation* 2024. arXiv: 2410.12908 [quant-ph]. https://arxiv.org/abs/2410.12908.

72. Dowling, J. P. Quantum optical metrology – the lowdown on high-N00N states. *Contemporary Physics* **49,** 125–143. ISSN: 0010-7514. https://doi.org/10.1080/00107510802091298 (Mar. 2008).

73. Aasi, J., Abadie, J. & Abbott *et al.*, B. P. Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light. *Nature Photonics* **7,** 613–619. ISSN: 1749-4893. https://doi.org/10.1038/nphoton.2013.177 (Aug. 2013).

74. Hofheinz, M. *et al.* Generation of Fock states in a superconducting quantum circuit. *Nature* **454,** 310–314. ISSN: 1476-4687. https://doi.org/10.1038/nature07136 (July 2008).

75. Premaratne, S. P., Wellstood, F. C. & Palmer, B. S. Microwave photon Fock state generation by stimulated Raman adiabatic passage. *Nature Communications* **8,** 14148. ISSN: 2041-1723. https://doi.org/10.1038/ncomms14148 (Jan. 2017).

76. Eickbusch, A. *et al.* Fast universal control of an oscillator with weak dispersive coupling to a qubit. *Nature Physics* **18,** 1464–1469. ISSN: 1745-2481. https://doi.org/10.1038/s41567-022-01776-9 (Dec. 2022).

77. Goldman, N. & Dalibard, J. Periodically Driven Quantum Systems: Effective Hamiltonians and Engineered Gauge Fields. *Phys. Rev. X* **4,** 031027. https://link.aps.org/doi/10.1103/PhysRevX.4.031027 (3 Aug. 2014).

78. Martin, I., Refael, G. & Halperin, B. Topological Frequency Conversion in Strongly Driven Quantum Systems. *Phys. Rev. X* **7,** 041008. https://link.aps.org/doi/10.1103/PhysRevX.7.041008 (Oct. 2017).

79. Crowley, P. J. D., Martin, I. & Chandran, A. Topological classification of quasiperiodically driven quantum systems. *Phys. Rev. B* **99,** 064306. https://link.aps.org/doi/10.1103/PhysRevB.99.064306 (6 Feb. 2019).

80. Bernevig, B. A., Hughes, T. L. & Zhang, S.-C. Quantum Spin Hall Effect and Topological Phase Transition in HgTe Quantum Wells. *Science* **314,** 1757–1761. eprint: https://www.science.org/doi/pdf/10.1126/science.1133734. https://www.science.org/doi/abs/10.1126/science.1133734 (2006).

81. Long, D. M., Crowley, P. J. D., Kollár, A. J. & Chandran, A. Boosting the Quantum State of a Cavity with Floquet Driving. *Phys. Rev. Lett.* **128,** 183602 (18 May 2022).

82. Niu, Q. & Thouless, D. J. Quantised adiabatic charge transport in the presence of substrate disorder and many-body interaction. *Journal of Physics A: Mathematical and General* **17,** 2453. https://dx.doi.org/10.1088/0305-4470/17/12/016 (Aug. 1984).

83. Thouless, D. J. Quantization of particle transport. *Phys. Rev. B* **27,** 6083–6087. https://link.aps.org/doi/10.1103/PhysRevB.27.6083 (10 May 1983).

84. Nathan, F., Martin, I. & Refael, G. Topological frequency conversion in a driven dissipative quantum cavity. *Phys. Rev. B* **99,** 094311. https://link.aps.org/doi/10.1103/PhysRevB.99.094311 (9 Mar. 2019).

85. Long, D. M., Crowley, P. J. D., Kollár, A. J. & Chandran, A. Boosting the Quantum State of a Cavity with Floquet Driving. *Phys. Rev. Lett.* **128,** 183602. https://link.aps.org/doi/10.1103/PhysRevLett.128.183602 (18 May 2022).

86. Thouless, D. J. Quantization of particle transport. *Phys. Rev. B* **27,** 6083–6087. https://link.aps.org/doi/10.1103/PhysRevB.27.6083 (10 May 1983).

87. Citro, R. & Aidelsburger, M. Thouless pumping and topology. *Nature Reviews Physics* **5,** 87–101. ISSN: 2522-5820. https://doi.org/10.1038/s42254-022-00545-0 (2023).

88. Nathan, F., Martin, I. & Refael, G. Topological frequency conversion in a driven dissipative quantum cavity. *Phys. Rev. B* **99,** 094311. https://link.aps.org/doi/10.1103/PhysRevB.99.094311 (Mar. 2019).

89. Blais, A., Grimsmo, A. L., Girvin, S. M. & Wallraff, A. Circuit quantum electrodynamics. *Reviews of Modern Physics* **93.** ISSN: 1539-0756. http://dx.doi.org/10.1103/RevModPhys.93.025005 (May 2021).

90. Rol, M. A. *et al.* Time-domain characterization and correction of on-chip distortion of control pulses in a quantum processor. *Applied Physics Letters* **116,** 054001. ISSN: 0003-6951. https://doi.org/10.1063/1.5133894 (Feb. 2020).

91. Slichter, C. *Principles of Magnetic Resonance* (Springer, 1990).

92. Krantz, P. *et al.* A quantum engineer's guide to superconducting qubits. *Applied Physics Reviews* **6,** 021318. ISSN: 1931-9401. https://doi.org/10.1063/1.5089550 (June 2019).

93. Bloch, I., Dalibard, J. & Zwerger, W. Many-body physics with ultracold gases. *Rev. Mod. Phys.* **80,** 885–964. https://link.aps.org/doi/10.1103/RevModPhys.80.885 (3 July 2008).

94. Bukov, M., D'Alessio, L. & Polkovnikov, A. Universal high-frequency behavior of periodically driven systems: From dynamical stabilization to Floquet engineering. *Advances in Physics* **64,** 139–226 (2015).

95. Oka, T. & Aoki, H. Photovoltaic Hall effect in graphene. *Phys. Rev. B* **79,** 081406. https://link.aps.org/doi/10.1103/PhysRevB.79.081406 (8 Feb. 2009).

96. Rechtsman, M. C. *et al.* Photonic Floquet topological insulators. *Nature* **496,** 196–200. ISSN: 1476-4687. https://doi.org/10.1038/nature12066 (Apr. 2013).

97. Holthaus, M. Floquet engineering with quasienergy bands of periodically driven optical lattices. *Journal of Physics B: Atomic, Molecular and Optical Physics* **49,** 013001. https://dx.doi.org/10.1088/0953-4075/49/1/013001 (Nov. 2015).

98. Spielman, I. Raman processes and effective gauge potentials. English. *Phys. Rev. A* **79** (2009).

99. Lin, Y.-J., Compton, R. L., Jiménez-García, K., Porto, J. V. & Spielman, I. B. Synthetic magnetic fields for ultracold neutral atoms. *Nature* **462,** 628–632. ISSN: 1476-4687. https://doi.org/10.1038/nature08609 (Dec. 2009).

100. Roushan, P. *et al.* Chiral ground-state currents of interacting photons in a synthetic magnetic field. *Nature Physics* **13,** 146–151. ISSN: 1745-2473. eprint: 1606.00077 (2017).

101. Rosen, I. T. *et al. Implementing a synthetic magnetic vector potential in a 2D superconducting qubit array* 2024. arXiv: 2405.00873 [quant-ph]. https://arxiv.org/abs/2405.00873.

102. Rudner, M. S., Lindner, N. H., Berg, E. & Levin, M. Anomalous Edge States and the Bulk-Edge Correspondence for Periodically Driven Two-Dimensional Systems. *Phys. Rev. X* **3,** 031005. https://link.aps.org/doi/10.1103/PhysRevX.3.031005 (3 July 2013).

103. Harper, F., Roy, R., Rudner, M. S. & Sondhi, S. Topology and Broken Symmetry in Floquet Systems. *Annual Review of Condensed Matter Physics* **11,** 345–368 (2020).

104. Wintersperger, K. *et al.* Realization of an anomalous Floquet topological system with ultracold atoms. *Nature Physics* **16,** 1058–1063 (2020).

105. Khemani, V., Moessner, R. & Sondhi, S. L. *A Brief History of Time Crystals* 2019. arXiv: 1910.10745 [cond-mat.str-el]. https://arxiv.org/abs/1910.10745.

106. Else, D. V., Monroe, C., Nayak, C. & Yao, N. Y. Discrete Time Crystals. *Annual Review of Condensed Matter Physics* **11,** 467–499. ISSN: 1947-5462. https://www.annualreviews.org/content/journals/10.1146/annurev-conmatphys-031119-050658 (2020).

107. Thouless, D. J. Quantization of particle transport. *Phys. Rev. B* **27,** 6083–6087. https://link.aps.org/doi/10.1103/PhysRevB.27.6083 (10 May 1983).

108. Citro, R. & Aidelsburger, M. Thouless pumping and topology. *Nature Reviews Physics* **5,** 87–101. ISSN: 2522-5820. https://doi.org/10.1038/s42254-022-00545-0 (2023).

109. Psaroudaki, C. & Refael, G. Photon pumping in a weakly-driven quantum cavity–spin system. *Annals of Physics* **435.** Special issue on Philip W. Anderson, 168553. ISSN: 0003-4916. https://www.sciencedirect.com/science/article/pii/S0003491621001597 (2021).

110. Long, D. M., Crowley, P. J. D. & Chandran, A. Nonadiabatic Topological Energy Pumps with Quasiperiodic Driving. *Phys. Rev. Lett.* **126,** 106805. https://link.aps.org/doi/10.1103/PhysRevLett.126.106805 (Mar. 2021).

111. Nathan, F., Ge, R., Gazit, S., Rudner, M. & Kolodrubetz, M. Quasiperiodic Floquet-Thouless Energy Pump. *Phys. Rev. Lett.* **127,** 166804. https://link.aps.org/doi/10.1103/PhysRevLett.127.166804 (Oct. 2021).

112. Hastings, M. B. & Haah, J. Dynamically Generated Logical Qubits. *Quantum* **5,** 564. ISSN: 2521-327X. https://doi.org/10.22331/q-2021-10-19-564 (Oct. 2021).

113. Lazarides, A., Das, A. & Moessner, R. Equilibrium states of generic quantum systems subject to periodic driving. *Phys. Rev. E* **90,** 012110 (July 2014).

114. Ponte, P., Chandran, A., Papić, Z. & Abanin, D. A. Periodically driven ergodic and many-body localized quantum systems. *Annals of Physics* **353,** 196–204 (2015).

115. Mori, T., Ikeda, T. N., Kaminishi, E. & Ueda, M. Thermalization and prethermalization in isolated quantum systems: a theoretical overview. *Journal of Physics B: Atomic, Molecular and Optical Physics* **51,** 112001. https://dx.doi.org/10.1088/1361-6455/aabcdf (May 2018).

116. Happer, W. Optical Pumping. *Rev. Mod. Phys.* **44,** 169–249. https://link.aps.org/doi/10.1103/RevModPhys.44.169 (2 Apr. 1972).

117. Phillips, W. D. Nobel Lecture: Laser cooling and trapping of neutral atoms. *Rev. Mod. Phys.* **70,** 721–741 (3 July 1998).

118. Fleischhauer, M., Imamoglu, A. & Marangos, J. P. Electromagnetically induced transparency: Optics in coherent media. *Rev. Mod. Phys.* **77,** 633–673. https://link.aps.org/doi/10.1103/RevModPhys.77.633 (2 July 2005).

119. Diehl, S. *et al.* Quantum states and phases in driven open quantum systems with cold atoms. *Nature Physics* **4,** 878–883. ISSN: 1745-2473. eprint: 0803.1482 (2008).

120. Harrington, P. M., Mueller, E. J. & Murch, K. W. Engineered dissipation for quantum information science. *Nature Reviews Physics* **4,** 660–671 (2022).

121. Poyatos, J. F., Cirac, J. I. & Zoller, P. Quantum Reservoir Engineering with Laser Cooled Trapped Ions. *Phys. Rev. Lett.* **77,** 4728–4731 (Dec. 1996).

122. Leghtas, Z. *et al.* Confining the state of light to a quantum manifold by engineered two-photon loss. *Science* **347,** 853–857. https://www.science.org/doi/abs/10.1126/science.aaa2085 (2015).

123. Puri, S., Boutin, S. & Blais, A. Engineering the quantum states of light in a Kerr-nonlinear resonator by two-photon driving. *npj Quantum Information* **3,** 18. ISSN: 2056-6387. https://doi.org/10.1038/s41534-017-0019-1 (Apr. 2017).

124. Grimm, A. *et al.* Stabilization and operation of a Kerr-cat qubit. *Nature* **584,** 205–209. ISSN: 1476-4687. https://doi.org/10.1038/s41586-020-2587-z (Aug. 2020).

125. Gertler, J. M. *et al.* Protecting a bosonic qubit with autonomous quantum error correction. *Nature* **590,** 243–248. ISSN: 0028-0836. eprint: 2004.09322 (2021).

126. Plenio, M. B., Huelga, S. F., Beige, A. & Knight, P. L. Cavity-loss-induced generation of entangled atoms. *Physical Review A* **59,** 2468–2475. ISSN: 1050-2947. eprint: quant-ph/9811003 (1999).

127. Kraus, B. *et al.* Preparation of entangled states by quantum Markov processes. *Physical Review A* **78,** 042307. ISSN: 1050-2947. eprint: 0803.1463 (2008).

128. Verstraete, F., Wolf, M. M. & Cirac, J. I. Quantum computation and quantum-state engineering driven by dissipation. *Nature Physics* **5,** 633–636. ISSN: 1745-2473 (2009).

129. Barreiro, J. T. *et al.* An open-system quantum simulator with trapped ions. *Nature* **470,** 486–491. ISSN: 0028-0836. eprint: 1104.1146 (2011).

130. Lin, Y. *et al.* Dissipative production of a maximally entangled steady state of two quantum bits. *Nature* **504,** 415–418. ISSN: 0028-0836. eprint: 1307.4443 (2013).

131. Shankar, S. *et al.* Autonomously stabilized entanglement between two superconducting quantum bits. *Nature* **504,** 419–422. ISSN: 0028-0836. eprint: 1307.4349 (2013).

132. Brown, T. *et al.* Trade off-free entanglement stabilization in a superconducting qutrit-qubit system. *Nature Communications* **13,** 3994. eprint: 2107.13579 (2022).

133. Home, J. P. *et al.* Complete Methods Set for Scalable Ion Trap Quantum Information Processing. *Science* **325,** 1227–1230. ISSN: 0036-8075. eprint: 0907.1865 (2009).

134. Home, J. P. *et al.* Memory coherence of a sympathetically cooled trapped-ion qubit. *Physical Review A* **79,** 050305. ISSN: 1050-2947. eprint: 0810.1036 (2009).

135. Raghunandan, M., Wolf, F., Ospelkaus, C., Schmidt, P. O. & Weimer, H. Initialization of quantum simulators by sympathetic cooling. *Science Advances* **6,** eaaw9268. eprint: 1901.02019 (2020).

136. Petiziol, F. & Eckardt, A. Cavity-Based Reservoir Engineering for Floquet-Engineered Superconducting Circuits. *Phys. Rev. Lett.* **129,** 233601 (Nov. 2022).

137. Wang, Z. & Safavi-Naeini, A. H. Quantum control and noise protection of a Floquet 0-π qubit. *Physical Review A* **109,** 042607. ISSN: 2469-9926 (2024).

138. Murch, K. W. *et al.* Cavity-Assisted Quantum Bath Engineering. *Physical Review Letters* **109,** 183602. ISSN: 0031-9007. eprint: 1207.0053 (2012).

139. Lu, Y. *et al.* Universal Stabilization of a Parametrically Coupled Qubit. *Phys. Rev. Lett.* **119,** 150502. https://link.aps.org/doi/10.1103/PhysRevLett.119.150502 (15 Oct. 2017).

140. Ma, R. *et al.* A dissipatively stabilized Mott insulator of photons. *Nature* **566,** 51–57 (Feb. 2019).

141. Li, Z., Roy, T., Lu, Y., Kapit, E. & Schuster, D. I. Autonomous stabilization with programmable stabilized state. *Nature Communications* **15,** 6978. ISSN: 2041-1723. https://doi.org/10.1038/s41467-024-51262-4 (2024).

142. Lloyd, S. Universal Quantum Simulators. *Science* **273,** 1073–1078 (1996).

143. Sarovar, M. & Milburn, G. J. Continuous quantum error correction by cooling. *Physical Review A* **72,** 012306. ISSN: 1050-2947. eprint: quant-ph/0501038 (2005).

144. Mi, X. *et al.* Stable quantum-correlated many-body states through engineered dissipation. *Science* **383,** 1332–1337. ISSN: 0036-8075. eprint: 2304.13878 (2024).

145. Walls, D. F. & Milburn, G. J. *Quantum Optics* (Springer, 2008).

146. Giovannetti, V., Lloyd, S. & Maccone, L. Advances in quantum metrology. *Nature Photonics* **5,** 222–229. ISSN: 1749-4893. https://doi.org/10.1038/nphoton.2011.35 (2011).

147. Tóth, G. & Apellaniz, I. Quantum metrology from a quantum information science perspective. *Journal of Physics A: Mathematical and Theoretical* **47,** 424006. https://dx.doi.org/10.1088/1751-8113/47/42/424006 (Oct. 2014).

148. Briegel, H. J., Browne, D. E., Dür, W., Raussendorf, R. & Van den Nest, M. Measurement-based quantum computation. *Nature Physics* **5,** 19–26. ISSN: 1745-2481. https://doi.org/10.1038/nphys1157 (2009).

149. Knill, E. *Fault-Tolerant Postselected Quantum Computation: Schemes* 2004. arXiv: quant-ph/0402171. https://arxiv.org/abs/quant-ph/0402171.

150. Bravyi, S. & Kitaev, A. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Phys. Rev. A* **71,** 022316. https://link.aps.org/doi/10.1103/PhysRevA.71.022316 (2 Feb. 2005).

151. Varcoe, B. T. H., Brattke, S., Weidinger, M. & Walther, H. Preparing pure photon number states of the radiation field. *Nature* **403,** 743–746 (2000).

152. Hofheinz, M. *et al.* Generation of Fock states in a superconducting quantum circuit. *Nature* **454,** 310–314 (2008).

153. Wang, H. *et al.* Measurement of the Decay of Fock States in a Superconducting Quantum Circuit. *Phys. Rev. Lett.* **101,** 240401 (Dec. 2008).

154. Heeres, R. W. *et al.* Implementing a universal gate set on a logical qubit encoded in an oscillator. *Nature Communications* **8,** 94 (2017).

155. Rivera, N., Sloan, J., Salamin, Y., Joannopoulos, J. D. & Soljačić, M. Creating large Fock states and massively squeezed states in optics using systems with nonlinear bound states in the continuum. *Proceedings of the National Academy of Sciences* **120,** e2219208120 (2023).

156. Schaefer, R., Kollár, A. J. & Chandran, A. *Dissipative stabilization of many-body Floquet systems* In prepartion.

157. Seetharam, K. I., Bardyn, C.-E., Lindner, N. H., Rudner, M. S. & Refael, G. Controlled Population of Floquet-Bloch States via Coupling to Bose and Fermi Baths. *Phys. Rev. X* **5,** 041050 (Dec. 2015).

158. Khalil, M. S., Stoutimore, M. J. A., Wellstood, F. C. & Osborn, K. D. An analysis method for asymmetric resonator transmission applied to superconducting devices. *Journal of Applied Physics* **111,** 054510. ISSN: 0021-8979. eprint: https://pubs.aip.org/aip/jap/article-pdf/doi/10.1063/1.3692073/15093308/054510_1_online.pdf. https://doi.org/10.1063/1.3692073 (Mar. 2012).

159. Megrant, A. *et al.* Planar superconducting resonators with internal quality factors above one million. *Applied Physics Letters* **100,** 113510. ISSN: 0003-6951. eprint: https://pubs.aip.org/aip/apl/article-pdf/doi/10.1063/1.3693409/13262211/113510_1_online.pdf. https://doi.org/10.1063/1.3693409 (Mar. 2012).

160. Guan, H. *et al.* Network modeling of non-ideal superconducting resonator circuits. *Superconductor Science and Technology* **33,** 075004. https://dx.doi.org/10.1088/1361-6668/ab89eb (May 2020).

Appendix A

Full scattering model for hanger resonators

In this appendix we briefly discuss a more complete model for the response of hanger-style resonators. While the simple model described in Chapter 3 is sufficient to obtain estimates of the resonant frequency and quality factors, for high-quality factor resonators the microwave environment the hanger is embedded in can distort the resonance.

The base model for the S_{21} of a hanger resonator is given by:

$$S_{21}(f) = 1 - \frac{Q/Q_c}{1 + 2iQ\frac{\delta f}{f_0}} \tag{A.1}$$

where Q, Q_c are the total and coupling quality factors respectively, f_0 is the resonance frequency, and $\delta f = f - f_0$ is the detuning from resonance. We can calculate the internal quality factor as $Q_i^{-1} = Q^{-1} - Q_c^{-1}$. Plotting the real and imaginary components of S_{21} parametrically displays a circle centered on the x-axis (real). The diameter of the circle corresponds to Q/Q_c and resonance occurs at the point closest to the origin as shown in Fig. A.1 where we compare critical coupling ($Q_i = Q_c$) as well as overcoupled devices ($Q_i = 4Q_c$). We perform a simultaneous fit to the real and imaginary components of S_{21} to obtain all the fit parameters. The following heuristics are used to generate a good starting initial guess:

- Compute the amplitude of the signal from the real and imaginary components of S_21 and normalize it to 1
- f_0 is defined as the minimum in the amplitude data
- Q is estimated from the FWHM of the data (approximated as the frequency range for which the amplitude is less than 0.5)
- The coupling factor g is computed as (1-min(S21))/min(S21) from which we can compute an estimate for $Q_i = (1 + g) * Q$
- Finally, we compute Q_c using the relation: $Q_c^{-1} = Q^{-1} - Q_i^{-1}$

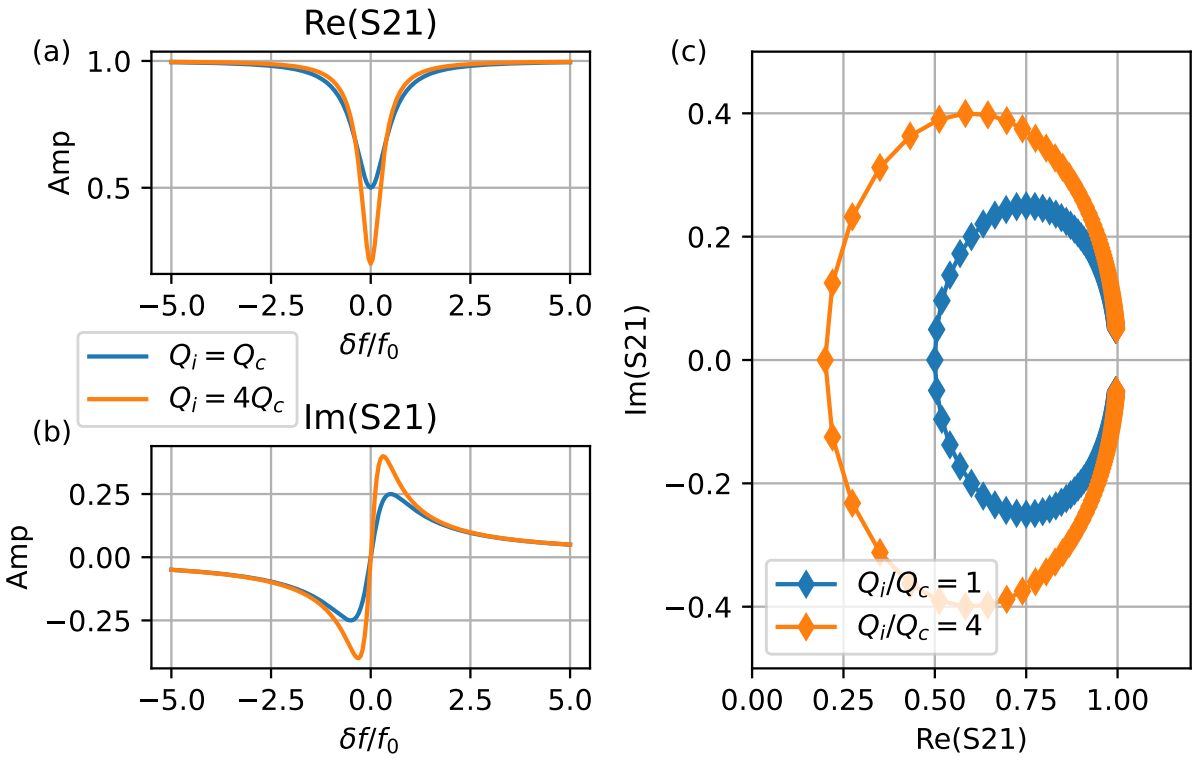

Figure A.1: S_{21} for critically and overcoupled devices. (a),(b) Real and imaginary components of S_{21} as a function of relative detuning (in linewidths) for $Q_i = Q_c$ (critically coupled) and $Q_i = 4Q_c$ (overcoupled). (c) Parametric plot of S_{21}, here the increase in resonator Q for the overcoupled device is visible by an increase in the diameter of the circle.

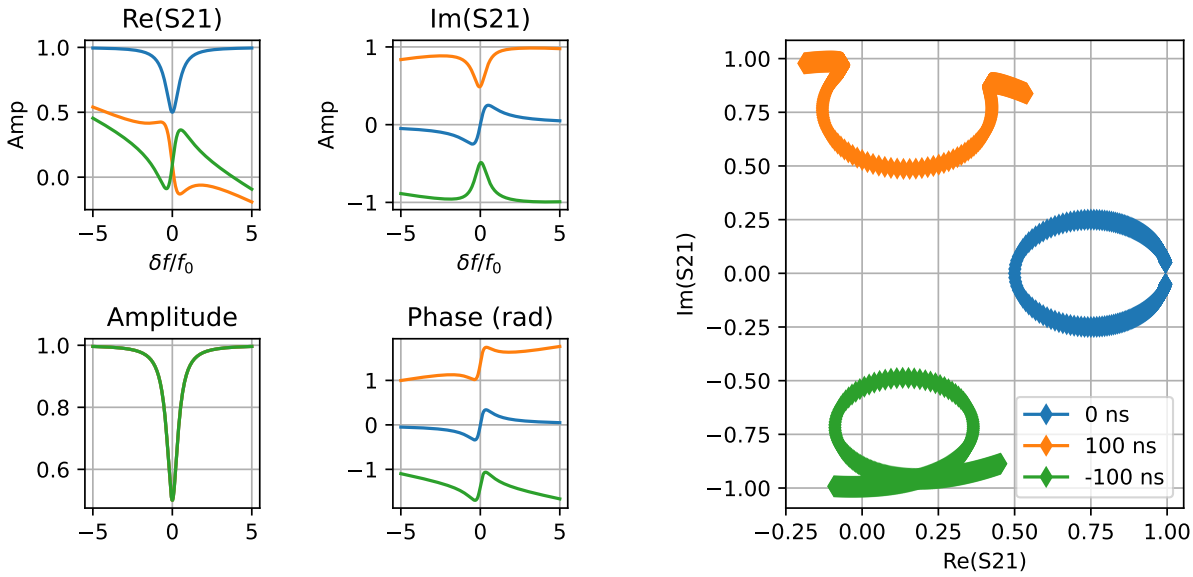

Figure A.2: Effect of line delay (and global phase offsets) on S_{21}. The blue, orange, and green traces correspond to Eq.(A.2) for line delays of 0,100,-100 ns respectively. In the real and imaginary components, the phase delay can be seen as a "tilt" in the real data corresponding to the overall oscillations between the real and imaginary components due to the linear dispersion in the cables. In an amplitude/ phase representation of the data, we can isolate the line delay as the slope of the phase signal (the amplitude is insensitive to the phase wrapping). In the circle plots, the line delay manifests as a distortion of the circle.

However, this simple model does not take into account three important features of a real measurement: the line delay introduced by the physical length of the cables, attenuation or gain of the signal, and impedance mismatches between the resonator and the feedline.

Line delay and signal attenuation can be taken into account by simply adding a complex prefactor to the S_{21} expression:

$$S_{21}(f) = ae^{i2\pi\tau f}\left(1 - \frac{Q/Q_c}{1 + 2iQ\frac{\delta f}{f_0}}\right) \tag{A.2}$$

where a is the gain (or loss) on the round trip signal and τ is the electrical line delay. Figure A.2 shows the effect of the line delay on the S_{21} measurement. Effectively, the line delay causes the real and imaginary components to oscillate (while keeping the amplitude constant) which acts as an additional "winding" in the S_{21} circle plot. Both the line delay and line gain can be calibrated out independently using the following method:

- Compute the amplitude of the signal, since the line delay only adds a phase to the signal, this measurement is insensitive to it. Both the real and imaginary traces can be normalized by the amplitude of the signal far off resonance (typically the maximum

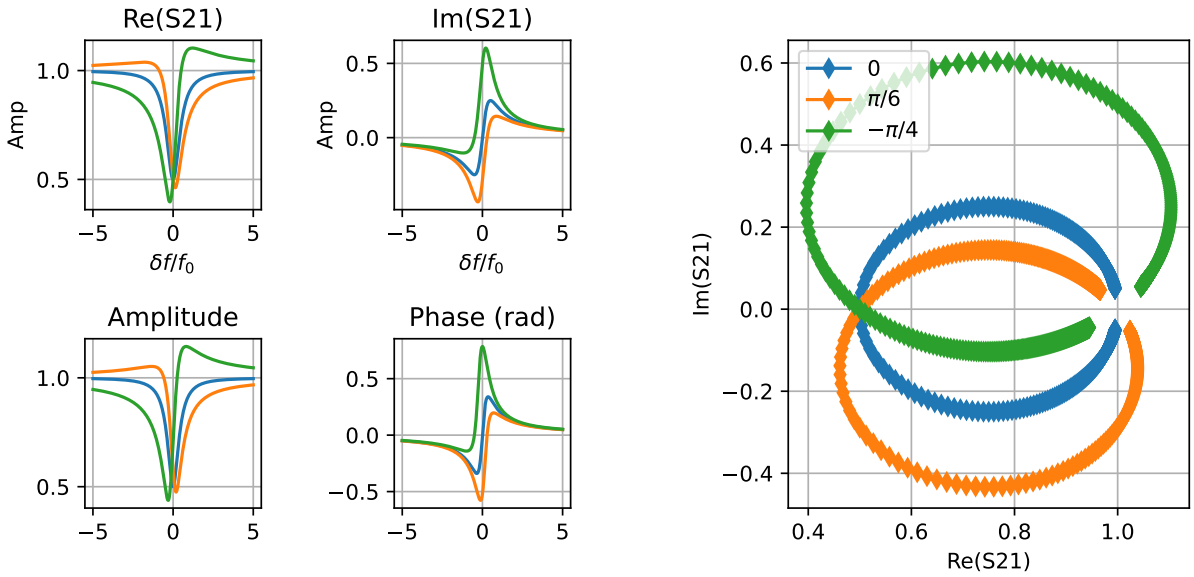

Figure A.3: Effect of an impedance mismatch on the S_{21} circle (using the DCM [158] model to generate the traces). Here a rotation of the circle about (1,0) instead of (0,0) introduces an asymmetry in the amplitude signal.

value)

- Compute the phase of the signal. As shown in Fig. A.2, the line delay introduces a linear phase wrapping whose slope is τ. However, as the resonance introduces a large phase signal, this calibration is better performed with data acquired far off-resonance. The real and imaginary components can then be divided by $e^{i2\pi\tau f}$ to remove the phase wrapping

While the phase delay and offset can cause the S_{21} response to look warped, crucially, it acts as a *rotation* about the origin (i.e. it looks like the circle is simply rotated about the point (0,0)). This leads to the amplitude data being unaffected by the phase delay and does not introduce errors in the fit parameters.

So far, we have assumed that the resonator is perfectly matched to its environment. However, small impedance mismatches can occur either due to the presence of wirebonds on the device or from imperfections in the coupling feedline used for measurements. This impedance mismatch is characterized by an asymmetric amplitude signal on either side of the resonance. Graphically, these mismatches can be thought of as rotations of the circle about the (1,0) point accompanied by an increase in the circle diameter as shown in Fig. A.3.

Multiple different models have been proposed to address the asymmetric nature of the resonances. While all equivalent, they build their effective models using different approaches and are of varying pedagogical quality:

- The Diameter Correction Method (DCM) [158] where the impedance mismatch arises from residual mutual inductance between the hanger and the feedline along with

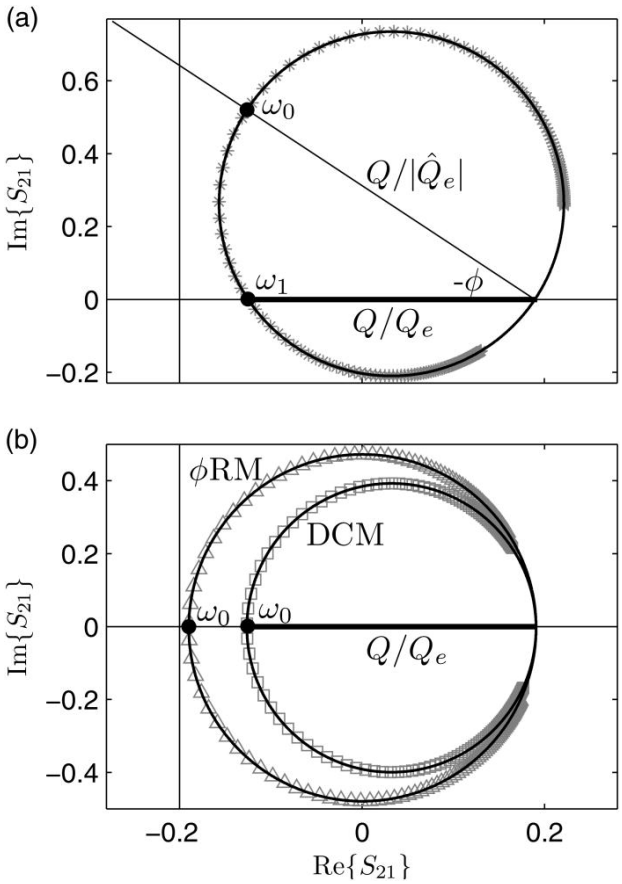

Figure A.4: Diameter correction method. (a) The imbalance in impedances is contained in the complex $\hat{Q}_c$ coupling quality factor which effectively rotates and stretches the resonance circle. (b) Comparison between the DCM method and the previously used "ϕ-rotation method where the circle was simply rotates back to the origin. The DCM method *both* rotates the circle by ϕ and scales the diameter by $\cos(\phi)$. This dual correction addressed a common issue where the calculated value of Q_i was negative which occurs when the circle crosses the y-axis. Figure reproduced from [158]

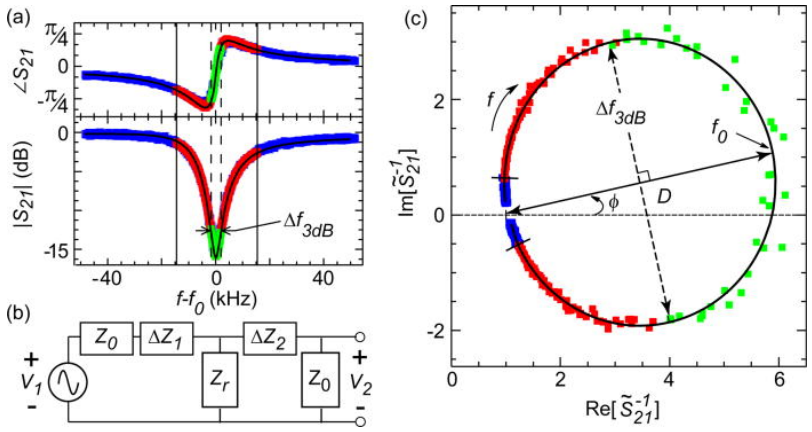

Figure A.5: Inverse S_{21} fitting method. The equivalent circuit model is shown in (b) with the resonator modeled as Z_r and the mismatches modeled as ΔZ_i. (c) Parametric plot of S_{21}^{-1} showing the circle and rotation due to the imbalances. Figure reproduced from [159]

different characteristic impedance measurement ports. The effective expression for S_{21} after normalization and phase delay compensation is: $S_{21} = 1 - \frac{Q/\hat{Q}_c e^{i\phi}}{1+2iQ\frac{\delta f}{f_0}}$ where $Q/\hat{Q}_c$ is the diameter of the circle and $Q_c = \hat{Q}_c \cos(\phi)$ as shown in Fig. A.4. Finally, Q_i is extracted using $1/Q_i = 1/Q - \cos(\phi)/Q_c$.

- The "Inverse" method (INV) [159] where small, inline, impedances are added to the measurement feedline to model imperfections in the coupling. The final expression is simplified by considering the *inverse* of S_{21}: $S_{21}^{-1} = 1 + \frac{Q_i}{Q_c^\star}\frac{1}{1+2iQ_i\frac{\delta f}{f_0}}$ where $Q_c^\star = \frac{Z_0}{Z}Q_c$ and Z is the effective impedance of the feedline (Z_0 for ideal coupling). The circuit model and resulting S_{21}^{-1} circle are shown in Fig. A.5.While this method directly provides access to Q_i, extracting the normalized coupling quality factor Q_c can be difficult.

- The "network" method (S21 network) [160] uses a signal flow approach (shown in Fig. A.6) to calculate the effect of small impedance mismatches on the input and output side. By mapping the coupled resonator to three serial blocks consisting of: input transmission line segment, ideal resonator, output transmission line segment, the final expression for S_{21} is simply obtained from a simplification of the signal flow network. In this model, $S_{21} = \frac{1-\phi}{1+u\phi}$ where $\phi = \frac{Q/Q_c}{1+2iQ\delta f/f_0}$ is the ideal resonator response and u is a complex number which includes all the non-idealities in the impedances. In this simpler picture, the ideal response is transformed by the factor $1 + u\phi$ which rotates and stretches the circle.

All three of these models are implemented in the `S21 helpers` fitting functions. NOTE: while the "inverse" model is stated in terms of S_{21}^{-1}, the fitting functions uses 1/model to preserve the gaussian noise properties of the measurement. Even though an error bar can be calculated for the real and imaginary components of S_{21}^{-1}, this is no longer a gaussian error and standard least-squares fitting routines no longer are guaranteed to produce the correct fit results. An example of the breakdown of the gaussian noise assumption in shown in

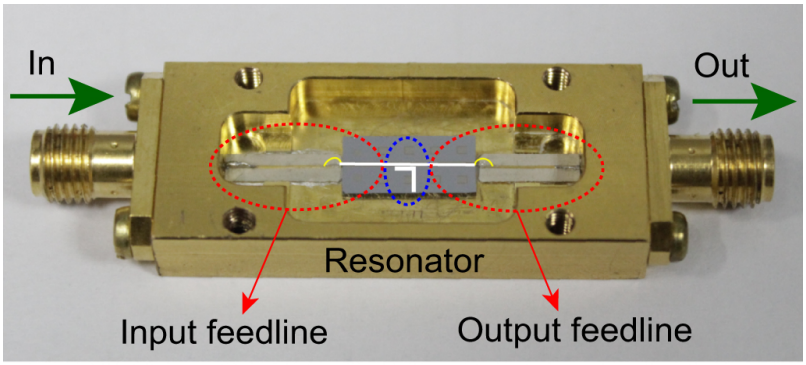

(a)

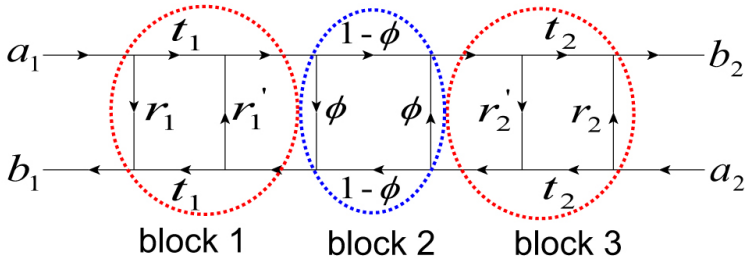

(b)

Figure A.6: Signal flow diagram approach to modeling impedance mismatches. (a) A simple resonator (center) is coupled to input and output ports whose impedance might not be matched (in the transition from the chip to the PCB, due to imperfections in fabrication, etc.)(b) Signal flow graph used to model impedance mismatches. The transmission and reflection properties of the input and output blocks are modeled as t_i, r_i respectively and encode impedance mismatches to the resonator circuit. Figure reproduced from [160]

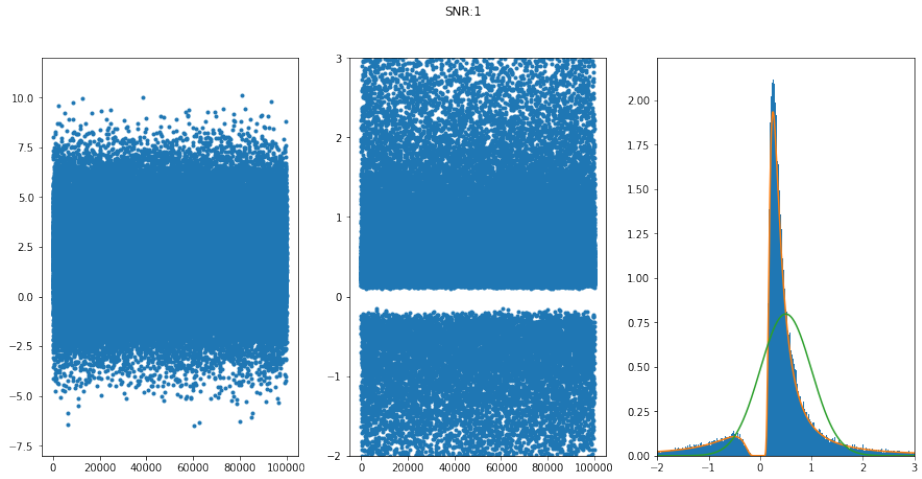

Figure A.7: Left: noisy signal with gaussian distributed error. Middle: inverse of the noisy signal. Right: Histogram of the data points for the inverse. This distribution is clearly non gaussian and displays a large gap near 0

Fig. A.7 where the inverse of a gaussian error signal no longer displays a symmetric profile. This breakdown is particularly noticeable for signals approaching zero as $1/(1+\epsilon)$ and $1/(1-\epsilon)$ can be very different as ϵ approaches one.

Appendix B

Fabrication

These are the standard fabrication recipes used in the lab initially developed by Maya Amouzegar and extended by Won Chan Lee who developed the air bridge recipe (unused for the devices presented in this thesis)

B.1 Dicing

B.1.1 Dicing a 4 Inch Wafer to a 1 Inch Wafer

- Spin-coat a 4-inch wafer (Tantalum on Sapphire) with AZ1518 at 4000 RPM for 1 minute. Make sure the wafer is metal-side up.
- Bake the wafer at 100°C for 4 minutes.
- Stick the wafer metal side up on a blue tape.
- Using the dicing saw to dice the wafer into $1'' \times 1''$ chips using the following parameters:
  - Mode = 60
  - Dim 1 = Dim 2 = 110 mm or larger
  - Index 1 = Index 2 = 25.75 mm
  - Height = 0.05 mm
  - Thickness = 0.6 mm
  - Cut Speed = 0.750 mm/s
  - Increment = 0.125 mm

  Note: It is a good idea to first make a shallow cut on a piece of junk Silicon to make sure the cut is in the right spot on the camera. You can increase the cut speed to 2 mm/s for Silicon.

! Check B.5.1 on how to dice.

B.1.2 Dicing a 1 inch Wafer to a 7mm x 7mm Wafer

- Do the three-solvent clean mentioned in section B.2.1 to remove excess photoresist.
- Spin AZ1518 on the chip at 4000 RPM for 1 minute.
- Bake the chip at 100°C for 4 minutes.
- Use the dicing saw to cut the chip into 7 mm × 7 mm pieces using the following parameters:
  - Mode = 60
  - Dim 1 = Dim 2 = 30 mm
  - Index 1 = Index 2 = 7.35 mm
  - Height = 0.050 mm
  - Thickness = 0.6 mm
  - Cut Speed = 0.750 mm/s
  - Increment = 0.125 mm
- For 10 mm × 10 mm chips change Index 1 = Index 2 = 10.3 mm

B.2 Current Fabrication Recipe for a 7mm x 7mm Wafer

B.2.1 Three Solvent Clean

- Soak chips in 1165 (**Remover PG**) for > 1 hour at 80°C.
- Sonicate in 1165 for 30 seconds.
- Transfer to a clean beaker of 1165 and sonicate for 30 seconds.
- Transfer to acetone while spraying with acetone.
- Transfer to IPA while spraying with IPA.
- Blow dry with N2.

B.2.2 Photolithography

You can either use the mask aligner with a photomask or the Heidelberg Maskless Aligner (MLA 150), which is a direct write system that uses your CAD design. **We now use Heidelberg MLA 150 at the fab lab**.

Exposure Using Heidelberg MLA 150

- Spin AZ1518 on the chip at 4000 RPM for 1 minute. Do this in the clean room under yellow lights. To setup the automated spinning stations near the Heidelberg follow the steps in **B.5.2**.
- Bake the chip for 4 minutes at 100°C.
- If you need to transfer your sample to a different room under blue light, use a dark GelPak and cover it in aluminum foil to avoid unwanted exposure.
- Use MLA150 with 375 nm laser for AZ1518. Check **B.5.3** to setup your designs for photolithography with MLA 150.
- The correct dosage for AZ1518 is still being tested. At the moment, 290 $\mathbf{mJ/cm^2}$ seems to work. 320 $\mathbf{mJ/cm^2}$ is a bit overdosed (features are roughly 18% bigger than they should be).
- Use the correct sample size. For 7 mm x 7 mm samples use the "Maya small 7 x 7" option. Do not place chips that are smaller than 6 mm x 6 mm in MLA 150.
- Center the design manually and carefully.
- Note: Features that contain other features inside them must be in different layers in the CAD design.

Develop Photo Layer

- After exposure, develop in CD-26 MIF for 1 minute.
- Transfer to DI water to stop developing.
- Blow dry with Nitrogen.
- Check under the optical Microscope to make sure your features are of the correct size and shape.

Etching

- Etch in Tantalum Etchant (made mainly of HF) for 21 seconds. (**Use only acid-safe plastic tweezers and plastic beakers**)
  - Do not handle HF unless there is a tube of unexpired Calcium Gluconate Gel in the fume hood.
  - Never use carbon-fiber tweezers in acids. Use only plastic tweezers that are acid-safe.
  - Do not pour HF in glass beakers. HF can melt glass. Use plastic beakers only.
- Transfer to DI water to stop the etch.
- Blow dry with N2.

B.2.3 Prepare for e-Beam Lithography

Three-solvent Clean

- Soak chips in 1165 (**Remover PG**) for $>$ 1 hour at 80°C.
- Sonicate in 1165 for 30 seconds.
- Transfer to a clean beaker of 1165 and sonicate for 30 seconds.
- Transfer to acetone while spraying with acetone.
- Transfer to IPA while spraying with IPA.
- Blow dry with N2.

TAMI Clean

- Sonicate in Toluene, Acetone, Methanol, and IPA for 2 minutes each, in the order mentioned.
- Spray with the next solvent before transferring to the new beaker.
- Blow dry with Nitrogen.

Piranha Clean

- Fill a beaker with 2 parts H2SO4 (Sulfuric Acid).
- Fill a beaker with 1 part H2O2 (Hydrogen Peroxide).
- Slowly and gently pour H2O2 into H2SO4.

! Never mix rapidly. The mixture will get too hot and could shatter your beaker.

! Never do this in reverse. It will blow up!

! **Never use carbon fiber or metal tweezers. Use only plastic ones**.

! **Use glass beakers for Piranha**.

! Always cover the work surface with PIG Pads and wear appropriate PPE (see section B.3.)

- Leave chips in Piranha for 20 minutes.
- Spray chips with DI water in a separate beaker for 30 seconds.
- Immediately spray with IPA, before the DI water dries on your chip.
- Blow dry with N2.

Spinning e-Beam Resist and Evaporating Anti-charging Layer

- Spin MMA at 5000 RPM for 1 minute (**Recipe 7**).
- Bake at 175°C for 2 minutes.
- Spin PMMA at 4000 RPM for 1 minute (**Recipe 5**).
- Bake at 175°C for 30 minutes.

! It is always best to spin resist the same day you are evaporating anti-charging layer. Ideally, you spin resist and evaporate anti-charging layer the evening before your Ebeam write. Essentially, this gives you time to pump down the chamber. You should load your sample before pumping down and evaporate right after. The time closeness or proximity of these steps is crucial because the resist can harden over time and your e-beam dose might not work well.

- Using the sputterer or the Plassys to coat your chip in a 30 nm thick layer of Aluminum. Check **B.5.4** for a complete list of the anticharging process

B.2.4 e-Beam Lithography

e-Beam Write

- Use the Elionix with a 0.5 nA beam current.
- Use 500 μm write fields and 50,000 dots per field.
- Use 1x1 shot pitch.
- For Manhattan style junctions use 1000μC/cm^2 for the full clear layer dose, and 230μC/cm^2 for the undercut layer.
- For Dolan Bridges use 1500μC/cm^2 for the full clear layer dose, 230μC/cm^2 for the undercut layer. **These doses are still being fine-tuned.**

Developing Manhattans

- Remove the anti-charging layer by soaking the chip in AZ300-MIF for 5 minutes.
- Rinse with DI water.
- Blow dry with N2.
- Develop in MIBK:IPA (1:3) for 50 seconds.
- Soak in IPA for 10 seconds.
- Rinse with DI water.
- Blow dry with N2.

Developing Dolans

- Remove the anti-charging layer by soaking the chip in AZ300-MIF for 5 minutes.
- Mix 3 parts of IPA with 1 part of DI Water in a glass bottle.
- Put the mixture in a chiller set to $6°\mathbf{C}$ for 2-3 hours depending on the size of the bottle.
- Fill a large beaker/dish with water and put it in the fridge for 3-4 hours to get it down to $3°\mathbf{C}$. This will be your bath.
- Take the bath out of the fridge, put a clean, empty beaker in it, and wait for the water to get to $5°\mathbf{C}$.
- Put 4 ice cubes in the bath, and pour developer in the clean beaker.
- Develop chip for 2 minutes, while agitating the solution.
- Blow dry with N2.

Note: There is no need to mix a fresh batch of developer every time.

Double-Angle Evaporation for Dolans

- Load your sample in the Plassys, noting the desired evaporation direction. Pump down for $\sim$ 12 hours or more, until the load lock pressure reaches $\sim 9.4e^{-8}$ mBar.
- Perform Argon Ion etch for 45 seconds for each evaporation angle using the following parameters for the IBG: anode voltage 400 V, emitter current 20 mA, acceleration voltage 80 V.
- Evaporate Ti at 0.2 nm/s for 2 minutes.
- Wait for chamber pressure to drop to low e^{-8} again.
- Evaporate 30 nm of Al at 0.5 nm/s with a $25°$ tilt.
- Oxidize using static oxidation for 10 minutes at 2.33 mBar.
- Evaporate 50 nm of Al at $-25°$ tilt and 0.5 nm/s.
- Oxidize at 40 mBar for 20 minutes.

Double-Angle Evaporation for Manhattans

- Load your sample in the Plassys, noting the desired evaporation direction. Pump down for $\sim$ 12 hours or more, until the load lock pressure reaches $\sim 9.4e^{-8}$ mBar.
- Perform Argon Ion etch for 40 seconds for each evaporation angle using the following parameters for the IBG: anode voltage 400 V, emitter current 20 mA, acceleration voltage 80 V.
- Evaporat Ti at 0.2 nm/s for 2 minutes.
- Move planetary to the first evaporation direction and evaporate 30 nm of Al at 0.5 nm/s.
- Oxidize at 2.33 mBar for 10 minutes.
- Move planetary to the second desired evaporation angle. Evaporate 50 nm Al at 0.5 nm/s.
- Oxidize at 40 mBar for 10 minutes.

Lift-Off

- Soak in 1165 at 80°C for 3 hours.
- Spray with IPA to remove excess metal.
- Transfer to a clean 1165 beaker and sonicate for 30 seconds.
- Transfer to IPA while spraying with IPA. Sonicate for 30 seconds.
- Blow dry with N2.

B.2.5 Air Bridge(Optional)

- It's always better to do airbridge fab process (B.2.5 $\sim$ B.2.5) in the same day. To prevent developing residues that blocks proper contact from airbridges and contacts.)

1st Air Bridge Layer

- Spin AZ1518 and stack 5 layers as the following:
  - Spin AZ1518 with 4000 rpm for 60s (No ramping phase).
  - Bake the chip at 100°C for 30s.
  - Repeat the 1, 2 steps 4 times.
  - 5th layer, spin AZ1518 with 4000 rpm, 60s
  - Bake the chip at 100°C for 60s.

- Exposure using Heidelberg MLA 150. Use alignment settings to align automatically during the write. The write settings should be 375 nm, 100%, **2200 mJ/cm^2**.

! Focus the laser on the Ta surface and decrease the focusing $\sim$ 1000.

- Develop using CD-26 MIF for 14 mins. Clean with DI water after.
- Blow dry with N2.
- Reflow the photoresist at 135°C for 1 mins 30s.
- Perform oxygen plasma etching using March Plasma Asher.
  - Put the sample to the lower plate.
  - Start inputting oxygen when the base pressure is 100 mTorr.
  - The settings are Rf power 200 W, oxygen flow 6 sccm(max) and 60 s.

Al deposition

- Load the sample to Denton and fully refill the Aluminium crucible. Pump down until 6.0 e^{-6} torr or lower $\sim$ 1.5 hrs.
- Perform 1 μm Aluminum deposition using Denton with the following parameters: 7.8 KV anode voltage, 0.15 A emitter current (subject to change), deposition rate 10Å/s. Make sure to maintain the deposition rate by increasing the emitter current.

! Do not increase the emitter current too fast. This will lead the crucible to crack.

2nd Air Bridge Layer

! Do it right after the Al deposition.

- Spin AZ1518 with 2000 rpm for 1 min 30s (No ramping phase).
- Bake the chip at 100°C for 1 min 30s.
- Exposure using Heidelberg MLA 150. The write settings are: 375 nm laser, 100 %, and **310 mJ/cm^2**. Use 2nd alignment settings to align automatically during the write.

! Focus the laser on the 1st airbridge layer box and decrease the focusing $\sim$ 2000.

- Develop using CD-26 MIF for 45s. Clean with DI water after.
- Blow dry with N2.

Al Etching

! Read section B.3 for safety before you use Al etchant.

- Perform aluminium etching to remove unnecessary parts.
  - Fill the Transene Al Etchant type A to glass beaker and warm up to 40°C. (∼20 mins)
  - Etch the chip. Etch until the Aluminium layer disappears by eye inspection. There is no definite etching time since it depends on chips.
  - Etch additional 10s after Al layer disappears.
  - During the etching, make sure to gently move the chip to remove bubbles on the surface.
  - Clean with DI water 3 mins.
  - **Gently** blow dry with N2.
- Perform oxygen plasma etching using March Plasma Asher.
  - Put the sample to the lower plate.
  - Start inputting oxygen when the base pressure is 100 mTorr.
  - The settings are Rf power 200 W, oxygen flow 5.5 sccm and two consecutive 5 mins etch.
  - Take 1 min break after the first etch to cool down the sample and the plate. Then perform the second etch with the same settings.
- Soak in 1165 at 80°C for 3 hours.

! Shake the chip **gently** in every 1 hour to remove excess Al layer.

- Transfer to a clean 1165 beaker and wait for 5 mins.
- Transfer to acetone and wait for 2 mins.
- **Gently** spray with IPA to remove excess metal.
- **Gently** blow dry with N2.

! Do not sonicate in any steps.

B.2.6 Packaging

- Probe witness junctions before packaging the chip to make sure E_j is reasonable.
- Sand down a puck to make sure the chip sticks out of its pocket a bit.
- Clean the PCB prior to wire bonding.
- Wire bond using buffer 10 for the Copper to Tantalum bonds, and buffer 11 for Tantalum to Tantalum bonds.

B.3 Standard Operating Procedure and Safety

- Always wear proper PPE when working with chemicals. Wear splash goggles, a lab coat, and gloves.
  - It is always best to double glove.
  - Remove any rings you might be wearing before putting gloves on. Rings can tear gloves.
- Wear **acid-safe gloves**, **apron**, and a **face shield** when working with acids. Be careful handling bottles and beakers when wearing acid-safe gloves, as they are bulky and can affect your fine motor skills.
- Always work in the fume hood when using any chemical. Be particularly careful with Methanol, Remover PG (1165/NMP), and Toluene as they are very harsh and quite toxic.
- Always wash your hands with soap and water after working with chemicals.
- For your safety and in case of emergency, always have someone accompany you when working with acids.
- Do not handle HF unless there is a tube of unexpired Calcium Gluconate Gel in the fume hood.
- Never use carbon-fiber tweezers in acids. Use only plastic tweezers that are acid-safe.
- Do not pour HF in glass beakers. HF can melt glass. Use plastic beakers only.
- Use glass beakers for Piranha.
- Do not put water in Sulfuric Acid. It reacts violently.
- Use PIG Pads to cover your work surface when handling acids.
- When working with solvents, you can simply cover your work surface with Texwipes.
- If you are in doubt about something, ask! Don't forget, you can refer to the safety data sheets that are next to the fume hood as well.

B.4 Waste Bottle Constituents

Three Solvent Clean and MIBK for Developing

- Acetone 30%
- Isopropanol 30%
- N-Methyl Pyrrolidone 30%

- Methyl Isobutyl Ketone 10%
- Check “Ignitable” and “Toxic”.

TAMI Clean

- Toluene 25%
- Acetone 25%
- Methanol 25%
- Isopropanol 25%
- Check “Ignitable” and “Toxic”.

AZ300MIF and DI Water for Removal of Anti-Charging Layer

- Tetramethylammonium Hydroxide 1%
- Water 99%
- Check “Corrosive” and “Toxic”.

Aluminium Etchant Type A and DI water

- Phosphoric Acid 40%
- Nitric Acid 5%
- Acetic Acid 20%
- Water 35%
- Check ”Toxic”.

B.5 How Tos

B.5.1 Dicing

- Turn on air supply
- Pull the valve beside the dicer to release air
- Turn on the vacuum
- Turn on the water
- Turn on the main control (twist the knob clockwise) and the monitor
- Press **RESET**

- Press **SPINDLE** to turn on the spindle
- Press **HOME**, the press **CHUCK ZERO**

Test Dicing

- Stick test material on the dicing stage
- Press **Wafer Lock** to lock the test material
- Pres **Prog**
- Change **height** to 0.475mm. Change **Cut Speed** to 2mm/s if the test material is Silicon. You might need to press the **ENG : MET** button to switch to metric units. All other settings should match the appropriate settings for the main wafer
- Press **Align** to align the cut path and direction
- Press **Single Cut** to cut
- Press **Home** after test dicing is done
- Press **Wafer Release**

Dicing main wafer

- Flip main chip upside down on a clean wipe so that the top or resist part is facing the clean wipe
- Draw the tape on the back of the chip and clear all bubbles
- Tape the wafer on the dicing stage
- Press **Wafer Lock** to lock the test material
- Pres **Prog**
- Change the parameters to the ones listed for the right wafer size
- Press **Align** to align the cut path and direction. Use the **CW/CCW** buttons to adjust the alignment
- Press **Single Cut** to cut
- Press **Home** after test dicing is done
- Press **Wafer Release**
- Turn off air, vacuum, water, and the main control

B.5.2 Spinning Stations

- Click recipe 9 (to customize a spinning recipe). In recipe 9, there will be 3 steps. The first and last steps help to ramp up and ramp down the spinning slowly (slower than the main spinning).
- Select Step 1 (press **STEP TERMINATE** button first before you press the number **1** button). Set speed to 500 by pressing the **SPEED** button then input 500 with the numbers. Make sure you press the **ON or ENTER** button if you had to change the values (sometimes the values on-screen match what you want to use so you can skip to the next step instead of changing the value and pressing **ENTER**). Set time to 1 sec by using the step button (**STEP** button first and set time to 1.0 sec). Press **ENTER** (again, if you changed the on-screen values).
- Select Step 2 (same thing as above) but set speed to 4000 and time/step to 60 secs. Do not forget to press enter if you change the values.
- Select Step 3, this is basically repeating step 1. That is, use speed 500 and step 1.
- Exit from the recipe by selecting Step 0. Step 0 ends recipe. You can check if everything works by testing the spin station with s junk chip.

! UPDATE: Instead of having 3 steps and ramping slowly, Won noticed it is better to do STEP 1: Speed=4000, Ramp=4000, Step=60s and STEP 2: Speed=0, Step=0 (both set to zero to basically make the whole process stop in STEP 1). So now, use 1 step instead of 3 steps.

B.5.3 MLA 150 Setup

- Found a nice tutorial on how to setup your CAD designs for proper exposure in MLA 150. Click the following for the MLA 150 Tutorial

B.5.4 Anticharging

To load your sample into Plassys:

- Vent Load Lock: Select **VENT** under the **LOAD LOCK STATE** to vent the load lock. Open load lock when venting sound stops.
- Tape your chips on the loading stage, make sure your alignment marks are not covered by the tapes. It is better to also not tape the top-right corner of the chip. Grounding pin in the Ellionix system usually goes on the top right corner of the chips.
- Put the loading stage back into Plassys in the right direction. The little pin head on the loading stage should be aligned and close to the gun head.
- Clean the rim of the load lock with IPA to get rid of any extra debris

- Pump Load Lock: Select **PUMP** under the **LOAD LOCK STATE** to pump the load lock. (We usually do the above steps around the evening so Plassys pumps overnight. Before the next steps on the following day, the pressure of load lock should be about or less than **2e-7 mBar**)

To start the anticharging process:

- Write down the pressure of load lock
- Turn on the Ebeam supply. This is located at the bottom of the Plassys system (tagged **EBG Supply**).
- Click **RUN PROCESS** to load your desired anticharging recipe
- Click **Execute** to start the recipe

! During the alert stage(s), check the stage using the shutter in Plassys and make sure everything is centered well. Adjust the position is not centered. Click continue when done.

! During **wait recipe time** or **wait recipe thickness**, write down the values for the **actual current** and **rate control**

- Turn off the EBG supply when recipe is done.
- Vent the load lock to remove your samples. Follow the same steps as above to vent the load lock, return the loading stage, and clean the rim of the load lock.
- Pump the load lock

  d