Quantum Trajectories of a Superconducting Qubit

By

Steven Joseph Weber

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Irfan Siddiqi, Chair
John Clarke
Sayeef Salahuddin

Fall 2014

ProQuest LLC.
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 - 1346

Quantum Trajectories of a Superconducting Qubit

Abstract

Quantum Trajectories of a Superconducting Qubit

by

Steven Joseph Weber

Doctor of Philosophy in Physics

University of California, Berkeley

Irfan Siddiqi, Chair

In quantum mechanics, the process of measurement is intrinsically probabilistic. As a result, continuously monitoring a quantum system will randomly perturb its natural unitary evolution. An accurate measurement record documents this stochastic evolution and can be used to reconstruct the quantum trajectory of the system state in a single experimental iteration. We use weak measurements to track the individual quantum trajectories of a superconducting qubit that evolves under the competing influences of continuous weak measurement and Rabi drive. We analyze large ensembles of such trajectories to examine their characteristics and determine their statistical properties. For example, by considering only the subset of trajectories that evolve between any chosen initial and final states, we can deduce the most probable path through quantum state space. Our investigation reveals the rich interplay between measurement dynamics, typically associated with wavefunction collapse, and unitary evolution. Our results provide insight into the dynamics of open quantum systems and may enable new methods of quantum state tomography, quantum state steering through measurement, and active quantum control.

Contents

List of Figures

Acknowledgments

I am incredibly grateful to everyone who has helped to make my time in Berkeley so rewarding.

Thank you Irfan, for the opportunity to work in your lab and for all of your guidance over the years. We all owe a great deal of our success to your thoughtful, patient, and shrewd leadership. You have created an environment that is simultaneously challenging and nurturing, an ideal setting for a graduate student.

Thanks to everyone in QNL, past and present, for all of the support, inspiration, conversation, and trips to House of Curries. I would particularly like to thank Kater Murch for sharing his boundless enthusiasm and close guidance throughout my entire graduate career, and for carefully reading each chapter of this thesis as I wrote it and offering thoughtful feedback. I would also like to thank Dan Slichter for investing his time and energy to help me get off to a good start at QNL.

Thanks to my friends in Berkeley whose constant companionship and easygoing attitudes have made it such a joy to live here.

To my father, who is always supportive and eager to listen to me attempt to explain my research.

To my mother who accepted my decision to move 3000 miles away from home instead of 120. I'm deeply grateful for all of the support over the years, and especially for when she showed up in Berkeley during the peak of my thesis-writing burnout and took me out for a restorative meal at Chez Panisse.

To my brother Andrew, my closest compatriot.

To my sister Sara, it was always a joy to visit you in San Francisco.

To everyone else who has shared their time, guidance, and friendship with me over the years.

And most importantly, I am immeasurably thankful for my wife Katie. Words can't express the joy that she brings to my life. I deeply admire her strength, her honesty, and her loyalty. She was willing to maintain our relationship when I moved across the country for grad school, and afterwards she found every available opportunity to make it out here so that we could be together. She has always supported me though the highs and the lows of grad school, and I can't even imagine the person I would be without her.

Chapter 1

Introduction

The initial observations which motived the theory of quantum mechanics, and later the wide body of results which helped to solidify the theory, were based on experiments performed on large ensembles of quantum systems, such as atoms and photons. In quantum theory, the basic mathematical object describing the state of a physical system, the wavefunction $|\psi\rangle$, evolves deterministically in time with dynamics governed by the Schrödinger equation. When observing some physical property of the system, the wavefunction describes the *probability* that a measurement of an individual quantum system within an ensemble of identically prepared systems will yield a particular measurement result.

While the founders of quantum mechanics were unable to experiment with single atoms or photons, they nevertheless were eager to apply the laws of quantum mechanics to 'thought experiments' on individual systems. These thought experiments reveal a number of strange properties of quantum mechanics, which seem to violate our naive intuition based on classical physics. Perhaps the most striking example of the strange predictions is the property of entanglement, first discussed in a thought experiment by Einstein, Podolsky, and Rosen [1], which predicts that two spatially separated particles can exist in a collective superposition state and that a measurement of one particle will affect the state of the other.

For over half a century after the conception of quantum mechanics, it remained impossible to directly investigate these foundational questions in actual experiments on individual quantum systems. Nevertheless, the theory of quantum mechanics as formulated in the 1920's and 30's has dramatically enhanced our understanding of nature. To quote the recent Nobel laureate Serge Haroche [2],

> Most of the fundamental and technological advances of the last century, which make our life so different from our great-grandparents', are due to the deep understanding of Nature brought about by the quantum revolution. If the theory appears strange, it is mainly because we try to describe it with words of our everyday life, which are adapted to the properties of macroscopic objects. Even if quantum concepts are necessary to understand in depth the electric conductivity of metals, the superfluidity of liquid helium or the colour of the sky, these macroscopic phenomena are not 'strange' because they can be described with usual

> words, which is not the case for an ion in a trap of a photon in a cavity. A distinction must thus be made between the 'microscopic' quantum strangeness directly displayed in thought experiments and the apparent plainness of the macroscopic physics, which do not violate our common sense in spite of its quantum substrate.

Haroche's experiments with Rydberg atoms in electromagnetic cavities are one example of a setting in which, thanks to tremendous technological advances over the past thirty years, it is now possible to investigate and to measure individual quantum systems. Others such systems include trapped ions, crystal defects, electron spins confined by nano-fabricated structures, and superconducting circuits. All of these systems can be used to experimentally probe the quantum strangeness of the microscopic world, which was previously relegated to the realm of theoretical musing.

The experiments discussed in this thesis use superconducting circuits to study the process of measurement of a quantum system. The details of how the state of a microscopic system is mapped onto a classical measuring apparatus are intimately related to the question of how the classical world emerges from quantum mechanics. This question was first posed in the famous thought experiment of Schrödinger [3], where a probabilistic microscopic event (such as the decay of a radioactive isotope) determines whether a cat in an isolated box will live or die. The main goal of Schrödinger's thought experiment was to point out the ambiguity in defining at what point a quantum measurement actually occurs. Can the cat exist in a macroscopic superposition of dead and alive until the box is opened by an experimenter? A number of different interpretations of quantum mechanics have been formulated which attempt to answer this question, including the many worlds interpretation [4] and the consistent histories interpretation [5], both of which bypass the need to delineate exactly where a quantum measurement occurs, relying instead on the concept of decoherence to explain why classical objects do not appear to exist in superposition states.

However, our experiments are concerned less with the interpretation of quantum measurement, which would be very difficult to address in a meaningful way using current experimental capabilities, and rather with the process of measurement itself. In particular, we will investigate how the state of a quantum system evolves *during* the course of a measurement. As we will discuss in detail, for example in sections 1.2, 2.3, and chapter 9, this topic is of both fundamental and practical interest for quantum information processing.

1.1 Superconducting qubits

In classical computing, information is stored in two-state systems, or bits, with the two states labeled as '0' and '1'. In quantum mechanics, a two-level system, or qubit, can exist in a superposition of its basis states $|0\rangle$ and $|1\rangle$. The qubit state can be represented by the wavefunction

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle\,, \tag{1.1}$$

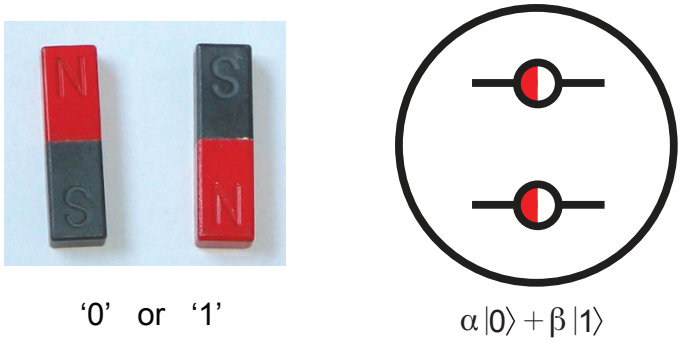

Figure 1.1: Classical and quantum bits. A classical bit (left panel), depicted here as a bar magnet, can be found in one of two states, for example, the magnet pointing either up or down. These two states are labeled '0' and '1'. A quantum two-level system, or qubit (right panel), can exist in a superposition of its basis states $|0\rangle$ and $|1\rangle$. Magnet photo adapted from http://allaboutmagnets.wikispaces.com.

where α and β are complex coefficients[1] which satisfy the condition $|\alpha|^2 + |\beta|^2 = 1$. Qubits have been realized using a variety of physical systems, including a spin-1/2 particle in a magnetic field, two isolated electronic energy levels of an atom, and the two polarizations of a photon. In this thesis, we will focus on qubits which are engineered from superconducting circuits.

As we will discuss in detail in chapter 3, a superconducting circuit is a macroscopic object which nevertheless possess a well-defined quantum degree of freedom and, with the proper choice of circuit parameters and operating conditions, can be experimentally addressed for control and measurement. The first superconducting qubit was realized in 1999 [6], when coherent oscillations between two charge states were observed, lasting for only ~ 1 ns. In the 15 years since then, there has been tremendous progress in understanding the physical mechanisms which limit coherence times and designing qubits which are less sensitive to these decohrence mechanisms. State-of-the-art qubit experiments have now demonstrated coherence times in excess of 100μs.

Ultimately, the goal of quantum information processing is to create an integrated system

[1]Without loss of generality (ignoring a physically insignificant global phase) we can also express $|\psi\rangle$ in terms of two angles θ and ϕ, which define a point on the surface of a sphere (known as the Bloch sphere), by defining $\alpha = \cos(\theta/2)$ and $\beta = e^{i\phi}\sin(\theta/2)$. If we know with certainty that a qubit is in a state $|\psi\rangle$, then we say that the qubit is in a 'pure state'. If instead there is some classical uncertainty as to what pure state $|\psi\rangle$ the system is in, it is said to be in a mixed state. A mixed state defines a point inside the Bloch sphere, and is typically represented by a density matrix ρ. A general qubit state can be described by three real numbers X, Y, and Z bounded by the range [-1,1], which give the coordinates of the state vector within the Bloch sphere. It is also common practice to map the qubit state onto the state of a spin-1/2 particle by defining $X \equiv \langle\hat{\sigma}_x\rangle$, $Y \equiv \langle\hat{\sigma}_y\rangle$, and $Z \equiv \langle\hat{\sigma}_z\rangle$.

in which a large number ($\gtrsim 10^4$) of long-lived qubits can be individually prepared and manipulated, made to interact in a known manner, and faithfully measured. In addition to improvements in coherence times, there has been a great deal of recent progress in developing robust architectures to control, couple, and read out qubits. The field is rapidly scaling up from initial demonstrations of two-qubit [7, 8, 9, 10, 11] and three-qubit [12] entanglement and and basic two-qubit algorithms [13, 14]. For example, the group of John Martinis has recently developed a processor consisting of five nearest-neighbor coupled qubits [15], which can be individually controlled with 99.92% fidelity, jointly controlled (via two-qubit gates) with 99.4% fidelity, and individually measured[16] with a fidelity of 99%.

Despite the improvements in coherence times, they will most likely remain short compared to the human timescales over which one would like to be able to run a quantum algorithm. To address this concern, a number of quantum error correction protocals have been developed [17] which would allow for fault tolerant quantum computing by encoding the state of one ideal, or 'logical', qubit redundantly within a set of imperfect, or 'physical', qubits. Each error correction scheme has a certain threshold for fault tolerance, which sets requirements for coherence times, the rates and fidelities of of qubit operations and measurements, and the number of physical qubits required to represent one logical qubit. One error correction scheme which has received considerable attention in recent years is the surface code [18, 19, 15], which has a less stringent error threshold and connectivity requirements than other proposals, at the expense of requiring a large number of physical qubits to encode each logical qubit.

There is also a great deal of interest in searching for problems that could be solved using qubits without the need to develop general-purpose quantum computer. In principle, by relaxing some of the architectural requirements, it may be possible to scale up more quickly to a system that is large enough to perform computations that are classically intractable. One example of such a procedure is quantum annealing, which may be useful for solving a wide range of optimization problems. In fact, the company D-Wave has already developed an sold a quantum annealer composed of 512 superconducting qubits. However, due to limitations in their choice of architecture, serious doubts have been raised about whether D-Wave's particular brand of quantum annealer will be useful for achieving a quantum computational speedup [20]. Nevertheless, it remains possible that a different style of quantum annealer may prove to a a valuable computational tool, and current theoretical efforts are underway to search for new types of problems which can be mapped onto the Hamiltonian of a quantum annealer. Another area of research which may bear fruit well before a quantum computer can be built is quantum simulation. Here, the main advantage is that only a modest number of qubits are required (~ 20) to access a parameter space which is inaccessible to classical simulations. Reference [21] demonstrates the richness of the physics that can be simulated when the experimenter has complete control over even just two qubits and their interaction.

In addition to their future application as a computational tool, superconducting qubits also provide immediate utility as a rich testbed for the fundamental physics of quantum optics, quantum measurement, and quantum feedback control. A wide range of experiments have made use of the constantly improving toolbox of superconducting qubits to investigate

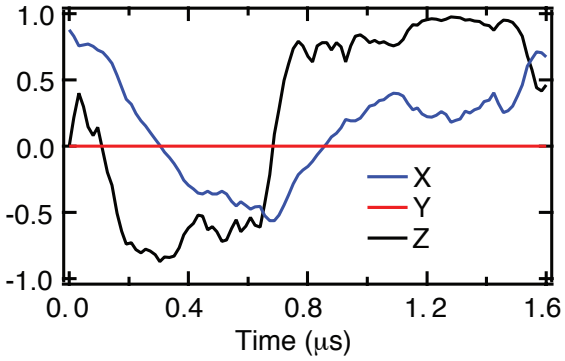

Figure 1.2: An example quantum trajectory of a superconducting qubit. The qubit is initially prepared along the equator of the Bloch sphere and evolves randomly under measurement. Using the measurement record, we can infer the time evolution of the qubit state $|\psi(t)\rangle$, which is depicted in terms of the components of the state vector, $X = \langle\hat{\sigma}_x\rangle$, $Y = \langle\hat{\sigma}_y\rangle$, and $Z = \langle\hat{\sigma}_z\rangle$.

topics such as quantum jumps [22, 23], quantum bath engineering [24], the interaction of squeezed light and matter [25], and the single-photon Kerr effect [26]. Along the same lines, this thesis will describe in detail two recent experiments from our group [27, 28] which track the quantum trajectories that the qubit state takes as it evolves in response to a continuous measurement.

1.2 Quantum trajectories

Textbook quantum mechanics describes the time evolution of isolated or closed quantum systems, whose dynamics are governed by the Schrödinger equation. In reality, no quantum system is completely isolated from its environment. This fact has been recognized from early on, a notable example being Fermi's golden rule, which calculates the rate atomic spontaneous emission based on coupling to the fluctuations of the electromagnetic vacuum. However, a sophisticated treatment of open quantum systems has emerged over only the past quarter century[29], motivated by experimental progress in the ability to control and measure individual quantum systems.

Consider a qubit initially prepared in an arbitrary superposition state $|\psi(t = 0)\rangle$. The length of time that the qubit can exist in a superposition state is determined by how strongly it interacts with its environment. This interaction entangles the quantum state with the inherent fluctuations of the environment. If these fluctuations are not measured, the environment can be viewed as a source of noise, causing random evolution of the qubit from an initially pure state into a statistical mixture— a process known as decoherence. However, by accurately measuring the environment, an observer can maintain complete knowledge of

the the qubit state. The quantum system remains in a pure state, and its time evolution is described by a 'quantum trajectory' $|\psi(t)\rangle$ which is determined by the measurement outcome. Figure 1.2 shows an example quantum trajectory of a superconducting qubit initially prepared along the equator of the Bloch sphere.

The measurement quantum efficiency η, which we formally define in section 2.2.3, is a measure of how well we can monitor a quantum system's environment. If we can perfectly track all of the environmental degrees of freedom which couple to a qubit, then $\eta = 1$, and the qubit will remain in a pure state as it evolves under measurement. However, if $\eta < 1$, we can only maintain partial knowledge of the qubit state, and the quantum trajectory will describe the density matrix $\rho(t)$. In the following section we will introduce the Josephson parametric amplifier, a tool which has enabled us to achieve quantum efficiencies of order unity (≈ 0.5), an essential prerequisite for the observation of quantum trajectories.

We present a historical overview of quantum trajectories in section 2.3.1. For now, we would like to emphasize that although the theory of quantum trajectories was developed in the quantum optics community in the 1990's and employed as a tool in a wide range of numerical simulations, quantum trajectories have only been investigated in a handful of experiments. The main result of this thesis is to examine quantum trajectories of a superconducting qubit in detail[27, 28] and to *experimentally verify* that we can accurately reconstruct the quantum trajectory from our measurement signal. From the perspective of quantum information technology, these experiments demonstrate the great extent to which we understand the process of measurement in our system and are useful for benchmarking experimental parameters such as the quantum efficiency which are essential to applications in measurement-based feedback and error correction.

1.3 Josephson parametric amplifiers

Superconducting qubits are typically designed to operate at GHz frequencies, which means that the energy of a single qubit excitation is much smaller than the thermal fluctuations of the room temperature electronics that are needed to record a measurement result. In order to perform a measurement with high quantum efficiency, it is necessary to encode the qubit state in a signal which is much larger than this noise floor. Typically, this process will require some sort of amplification.

In addition to increasing the magnitude of a signal, amplifiers will also add noise. In fact, quantum mechanics dictates that a minimum amount of noise must be added in any amplification process [30]. An amplifier which adds only the minimum required noise is said to be quantum-limited and can be used to realize an ideal quantum measurement with $\eta = 1$. However, state-of-the-art commercial low-noise amplifiers, which are made from high electron mobility transistors (HEMTs) and can be operated at 4 Kelvin, add considerably more noise than the quantum limit, and can only be used to achieve $\eta \sim 0.01$. Therefore, a more sensitive pre-amplifier is needed in order to overcome the added noise of the HEMT amplifier.

In recent years, superconducting parametric amplifiers based on the non-linear Josephson inductance have emerged as an effective tool for realizing qubit measurements with high quantum efficiency. They operate on the same principle as optical parametric oscillators [31], which have been widely used in the quantum optics community for many decades. A pump tone is applied to the amplifier, which periodically modulates some parameter of the system. This modulation couples energy from the pump tone into a signal mode at a different frequency, thus providing amplification. Superconducting parametric amplifiers were first considered in 1975 [32] and developed in the 1980s [33], but have not achieved quantum-limited noise performance and widespread use until the past five years [34, 35, 36], due to the growing interest in quantum information processing with superconducting circuits.

1.4 Thesis overview

Chapters 2-4 present the theoretical framework and historical background which lay the foundation for our experimental results. The thesis begins in with an introduction to the topic of continuous quantum measurement in chapter 2. We will describe the the type of measurements necessary to continuously monitor a qubit and explain the difference between quantum jumps and diffusive trajectories. We will also derive the theoretical results which we will later use to reconstruct quantum trajectories from experimental measurement records.

Chapter 3 provides a brief introduction to the field of superconducting qubits, focusing on describing the physical picture for our qubit measurements, which rely on a circuit quantum electrodynamics architecture. Chapter 4 discusses the theory and basic operating principles of the Josephson parametric amplifier and introduces the amplifierperformance metrics which are important to our experiments.

Then, in chapter 5, we move on to discuss the details of our experimental setup. We provide a detailed description of our qubit and amplifier designs, measurement techniques, and calibration experiments.

Chapter 6 will demonstrate how we use this experimental setup to track individual quantum trajectories of a superconducting qubit and to experimentally verify that they correctly describe the conditional qubit state. In chapter 7, we will analyze distributions of quantum trajectories to gain insight into the qubit dynamics under measurement.

In chapter 8, we will investigate what limits the quantum efficiency of our qubit measurements, which will lead to a discussion of the connection between squeezing and amplifier efficiency.

Finally, we will conclude in chapter 9 by discussing extensions of our work to quantum trajectories of multi-qubit systems and future applications in continuous quantum error correction.

1.5 Summary of key results

In the work presented in this thesis, we use weak measurements to continuously monitor a microwave cavity containing a superconducting qubit, and track the individual quantum trajectories of the system. These are the first experiments to successfully track diffusive quantum trajectories in a solid state system[27, 28]. Furthermore, these are the first experiments on any system to use quantum state tomography at discrete times along the trajectory to verify that we have faithfully tracked the qubit state.

To gain insight into qubit state dynamics under measurement, we examine ensembles of quantum trajectories. For example, by analyzing sub-ensembles of trajectories that end at an arbitrarilly chosen final state, we determine the most probable path connecting two points in quantum state space, finding our results to be in good agreement with theoretical predictions based on an action principle for continuous quantum measurement[37]. We also track quantum trajectories in the presence of a Rabi drive, which allows us to investigate the rich interplay between measurement dynamics, typically associated with wavefunction collapse, and unitary evolution of the quantum state as described by the Schrödinger equation. Finally, our experiments offer a precise means of characterizing the quantum efficiency of our measurement chain, providing a valuable tool for amplifier development. Our experiments highlight the great depth to which we understand the process of measurement in our system, and may inform future efforts in measurement-based quantum feedback for state steering and continuous quantum error correction.

Chapter 2

Continuous quantum measurement

In textbook quantum mechanics, measurement is typically treated as an instantaneous non-unitary process, through which a quantum system is projected into an eigenstate of the measured observable with a probability given by Born's rule. In reality, measurements are never truly instantaneous, but occur over some finite timescale determined by the details of the interaction between the measured system and the measurement apparatus. It follows naturally that a general treatment of quantum measurement must consider a gradual and continuous process of wavefunction collapse.

In this chapter, we discuss the principle of continuous quantum measurement from both a conceptual and a theoretical perspective. We place particular emphasis on the metrics which are used to describe the process of measurement in our qubit experiments. For a thorough and well written introduction to the general theory of continuous quantum measurement, see reference [38].

2.1 Quantum non-demolition measurement

This section introduces the general class of measurements required to continuously extract information from a quantum system, which are known as quantum non-demolition (QND) measurements. We will begin by reviewing a textbook quantum measurement of a qubit, which we will use to illustrate the concept of measurement backaction. We will then discuss the type of backaction required to realize a QND measurement and describe a general QND measurement setup. Finally, we will introduce the formal requirements for QND measurement.

2.1.1 Measurement backaction

Consider a qubit, initially prepared in the superposition state

$$|\psi_i\rangle = |+x\rangle \equiv \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle), \tag{2.1}$$

as pictured on the Bloch sphere in figure 2.1. For now, we will consider an ideal projective measurement, also known as a von Neumann measurement, in the qubit's energy eigenbasis (the $\hat{\sigma}_z$ basis). In this basis, the eigenstates are the qubit states $|0\rangle$ and $|1\rangle$, with the corresponding eigenvalues of $+1$ and -1, respectively[1]. Since the qubit is in an equal superposition of $|0\rangle$ and $|1\rangle$, it is impossible to predict the result of a single measurement. The probabilities of measuring each eigenvalue are given by Born's rule:

$$P(+1) = |\langle 0 \,|\psi_i\rangle|^2 = 1/2 \tag{2.2}$$

$$P(-1) = |\langle 1 \,|\psi_i\rangle|^2 = 1/2, \tag{2.3}$$

and after a measurement the qubit will remain in the eigenstate corresponding to the measured eigenvalue. If we were to perform two measurements back-to-back, then the second measurement will result in the same outcome as the first measurement, with a probability of 1. Thus, by measuring the qubit state we have improved our ability to predict the result of future measurements. This process of information acquisition can also be thought of as a *reduction* of an initial probability distribution for our measurement outcomes[39].

In this example of an ideal projective qubit measurement, we can identify two distinct types of measurement backaction. The first type is backaction in $\hat{\sigma}_z$, associated with the acquisition of qubit state information. The measurement changes $\langle\hat{\sigma}_z\rangle$ from 0 to either 1 or -1. However, the measurement also transforms $\langle\hat{\sigma}_z\rangle$ from 1 to 0. By measuring in the $\hat{\sigma}_z$ basis, we sacrifice our ability to predict the result of a future $\hat{\sigma}_x$ measurement. Because $\hat{\sigma}_z$ and $\hat{\sigma}_x$ are non-commuting observables, this type of backaction is required by Heisenberg's uncertainty principle and is sometimes referred to as Heisenberg backaction. Stated in formal language, the act of acquiring information about one observable will necessarily perturb its canonically conjugate observable. As discussed in reference [39], Heisenberg backaction is the primary concept that distinguishes quantum measurement from classical measurement.

For the case of an ideal qubit measurement, notice that the the Heisenberg backaction in $\hat{\sigma}_x$ doesn't affect the dynamics of the measured observable, $\hat{\sigma}_z$. While this condition sounds trivial in this example, it is essential to our ability to predict the results of repeated measurements and, in the limit of continuous measurement, to our ability to continuously extract information from the system. A measurement which cases no backaction on the measured observable beyond the usual backaction associated with the acquisition of information, is called a quantum non-demolition (QND) measurment. While an ideal qubit measurement is clearly QND, many other types of measurements, such as a position measurement, are fundamentally non-QND.

Consider, for example, an electron in the ground state of a harmonic oscillator potential, with an initial position wavefunction given by $\psi_i(x) \propto \mathrm{Exp}[-m\omega x^2/2\hbar]$. It is easily shown that the initial position uncertainty $\Delta x_i = \sqrt{\hbar/2m\omega}$ and the initial momentum uncertainty

[1]By this definition, the ground state will point up on the Bloch sphere. We have chosen this definition to be consistent with nuclear magnetic resonance terminology. If a spin is sitting in an external magnetic field which points up, the lower energy spin state (the ground state) will also point up.

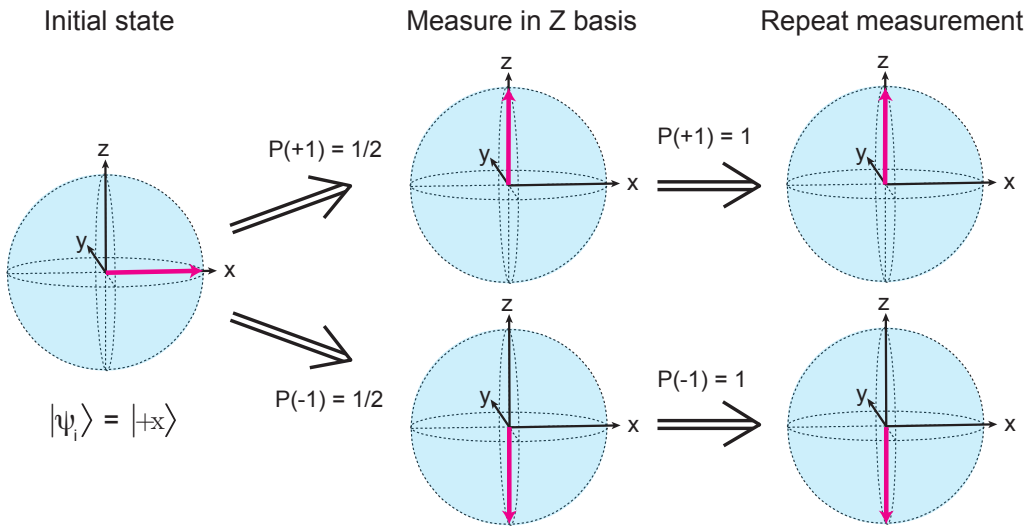

Figure 2.1: Ideal projective qubit measurement. A qubit, initially prepared in an equal superposition of $|0\rangle$ and $|1\rangle$, is measured in the $\hat{\sigma}_z$ basis, resulting in either the $|0\rangle$ state or the $|1\rangle$ state, each with probability 1/2. On repeated measurement, the qubit will remain in eigenstate corresponding to the result of the first measurement.

$\Delta p_i = \sqrt{\hbar m\omega/2}$, and thus the initial state has minimum uncertainty, $\Delta x_i \Delta p_i = \hbar/2$. Therefore, if we perform a measurement of x with imprecision $\Delta x_m \ll \Delta x_i$, then the uncertainly principle requires a corresponding increase in Δp.

This backaction in momentum, which will then randomly perturb the time evolution of x, places a fundamental limit on our ability to use repeated measurements to monitor position as a function of time. This limit, known as the *standard quantum limit*, has arisen in a variety of experimental settings ranging from nano-mechanical resonators [40, 41] to gravitational wave detectors [42, 39].

The standard quantum limit is an excellent example of how, when using non-QND measurements, there are fundamental limits on our ability to continuously extract information from a system and to track the time evolution of a quantum state. For QND measurements, however, such limits do not exist.

2.1.2 Indirect measurement

The most straightforward way one could image performing a quantum measurement is to directly couple the measured quantum system to a classical measuring device. For example, an avalanche photodiode can be used to detect the presence of a single photon. In practice, the classical measurement device is often composed of many noisy degrees of freedom, which cause significantly more backaction than dictated by the uncertainty principle.

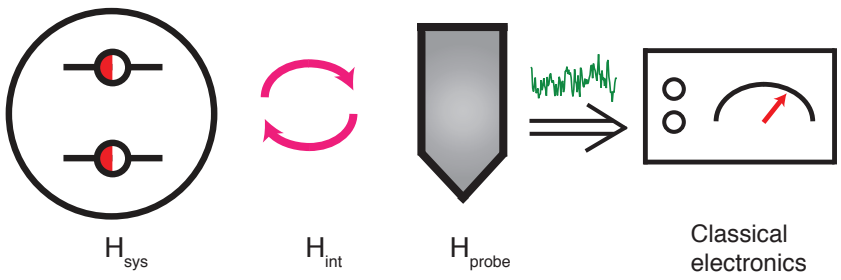

Figure 2.2: Schematic representation of an indirect quantum measurement. A quantum system described by a Hamiltonian H_{sys} is coupled to a probe described by H_{probe} via an interaction Hamiltonian H_{int}. Then the state of the probe is detected using classical electronics.

An avalanche photodiode absorbs the measured photon, completely destroying its quantum state and making a repeated measurement impossible.

An effective method to reduce detector backaction is to isolate the measured system from the classical measuring device, thereby realizing an *indirect measurement*. As depicted in figure 2.2, the measured quantum system couples quantum mechanically to an auxiliary quantum system known as the probe. Together, the measured system and the probe can be described by the Hamiltonian

$$H = H_{sys} + H_{probe} + H_{int}, \tag{2.4}$$

where H_{sys} and H_{probe} describe the uncoupled system and probe, and H_{int} describes their interaction. In order to separate free system evolution from measurement, it is useful to have *in situ* control of H_{int}.

An indirect measurement can be conceptualized as a two step process. In the first step, the system undergoes unitary interaction with the quantum probe, which is initially prepared in a known quantum state. In the second step, the probe interacts with a classical measurement device. In the ideal case where this second step contributes nothing to the measurement error, the only backaction on the system results from the initial quantum uncertainty of the probe and, given a knowledge of the interaction Hamiltonian, is completely specified by the measurement result. We will return in section 2.2.3 to discuss non-idealities in this second step, which determine the quantum efficiency of a measurement.

This two step procedure can be generalized to describe a continuous measurement if it occurs over an infinitesimally short time step and is repeated continuously. In this case, information is continuously extracted from the probe as it interacts with the measured system.

An indirect measurement will be QND, provided that the probe is not affected by more than one of any set of non-commuting observables of the measured system. Indirect measurement schemes have been used to implement QND measurements in a variety of experimental

settings, including optical interferometry [43, 44], photon counting [45], and superconducting qubit readout [46, 47].

2.1.3 Formal criteria for QND measurement

Consider an indirect measurement of a general coordinate q. Then, a necessary and sufficient condition for a QND measurement [39] is given by

$$[q, U] = 0, \tag{2.5}$$

where U is the unitary operator that generates joint evolution of the measured system and the probe over the full measurement time. In practice, U is often difficult to calculate, so a more restrictive, sufficient condition for QND measurement is often used. Assuming there is no explicit time dependance of the operator q, this condition is given by

$$[q, H] = 0. \tag{2.6}$$

By definition, $[q, H_{probe}] = 0$, and if q is conserved during free evolution of the measured system, then $[q, H_{sys}] = 0$. Thus, by equation (2.4), the condition given in equation (2.6) also implies

$$[q, H_{int}] = 0. \tag{2.7}$$

This condition will satisfy equation (2.5) regardless of the measurement duration, and is therefore a useful guideline for constructing a continuous quantum measurement.

2.2 Partial and projective measurements

As discussed in section 2.1.1, an ideal projective measurement is a type of QND measurement. However, one can also a perform non-projective, or partial, QND measurement. To illustrate this concept, consider the indirect measurement of a qubit, initially prepared in the state

$$|\psi_i\rangle_{sys} = \frac{1}{\sqrt{2}}(|0\rangle_{sys} + |1\rangle_{sys}). \tag{2.8}$$

Suppose that the qubit interacts briefly with an auxiliary qubit, which acts as the probe, such that after the interaction the system and probe qubits are in the entangled state

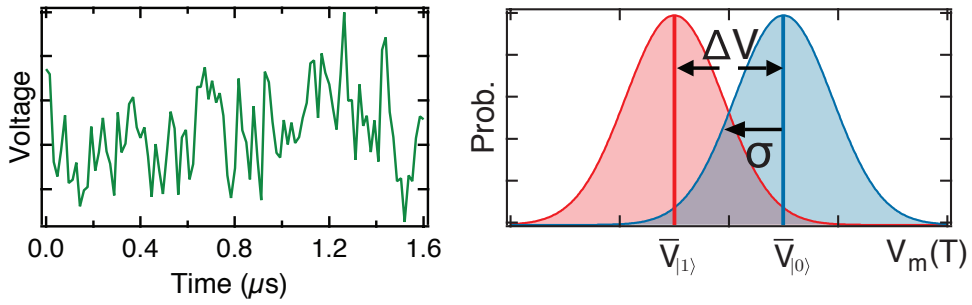

Figure 2.3: Continuous voltage measurement. The left panel shows an example of a measured voltage which fluctuates in time. The right histograms of integrated measurement voltages for a qubit prepared in the ground (blue) and excited (red) states.

$$\begin{aligned} |\Psi\rangle \propto & [(1+\epsilon)|0\rangle_{sys} + (1-\epsilon)|1\rangle_{sys}] \otimes |0\rangle_{probe} \\ & + [(1-\epsilon)|0\rangle_{sys} + (1+\epsilon)|1\rangle_{sys}] \otimes |1\rangle_{probe} . \end{aligned} \tag{2.9}$$

If a projective measurement is then performed on the probe qubit, the state of the system qubit will be rotated from its initial state in a way that depends on the measurement result. For example, if the probe is detected in ground state, then the system qubit will be left in the state

$$|\psi_f\rangle_{sys} \propto (1+\epsilon)|0\rangle_{sys} + (1-\epsilon)|1\rangle_{sys} . \tag{2.10}$$

If $\epsilon \ll 1$, this procedure can be though of as a partial measurement of the system qubit which drives it slightly toward the ground state. By repeating a sequence of such partial measurements, the qubit will eventually be projected into one of its eigenstates.

2.2.1 Characteristic measurement timescale

In experiments, all physical processes occur over some finite timescale. In this section, we will define the characteristic timescale over which a quantum measurement occurs. We will consider the case of an indirect measurement of a qubit, as pictured in figure 2.2. Suppose that the probe system is some property of the electromagnetic field which fluctuates in time (due to the quantum fluctuations, which we will discuss in section 3.1) and can be recorded as a noisy classical voltage signal. For simplicity, we will assume that the average value of

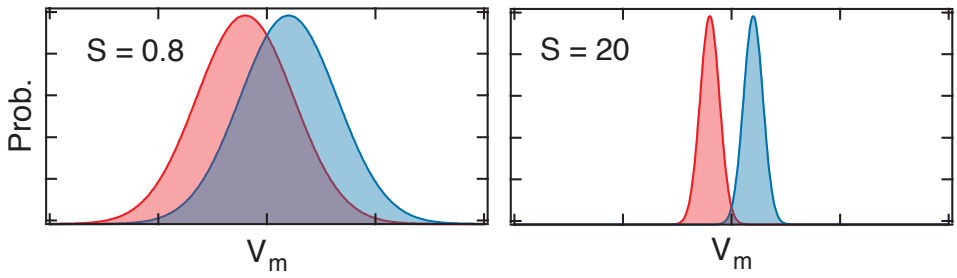

Figure 2.4: Partial and projective measurements histograms. The left panel shows the ground and excited state histograms for a partial measurement with a signal-to-noise ratio $S = 0.8$. The right panel shows histograms for a projective measurement, with $S = 20$, corresponding to 25 times longer measurement time than the partial measurement.

this signal is proportional to the qubit state. As we will show in chapter 3, this picture directly corresponds to the measurement process for our qubit experiments.

Suppose that a measurement voltage of $\overline{V}_{|0\rangle}$ ($\overline{V}_{|1\rangle}$) corresponds to the qubit being in the ground (excited) state. Figure 2.3 shows an example of the voltage measured as a function of time for an individual experimental iteration. Using this trace, we can compute the time-averaged voltage signal

$$V_m(T) \equiv \frac{1}{T}\int_0^T V(t)\,dt, \tag{2.11}$$

for each experimental iteration. The right panel of figure 2.3 shows histograms of $V_m(T)$, for an arbitrary measurement time T. Assuming the voltage fluctuations are uncorrelated, the histograms are Gaussian and can be characterized by their standard deviation σ. If the qubit is prepared in the ground state, the histogram will be centered at $\overline{V}_{|0\rangle}$ (shown in blue), and if the qubit is prepared in the excited state the histogram will be centered at $\overline{V}_{|1\rangle}$ (shown in red). We will define this voltage difference to be $\Delta V \equiv \overline{V}_{|0\rangle} - \overline{V}_{|1\rangle}$.

Note that when the ground and excited state histograms overlap, it is impossible to completely determine the qubit state in an individual measurement. In this case, a measurement will only contain partial information about an unknown qubit state. An important feature of these histograms is that their variance, σ^2, is inversely proportional to the measurement time T. Their center voltage, however, does not depend on T. Thus, as T increases, it becomes easier to distinguish between the ground and excited states. After a sufficiently long measurement, the histograms will be narrow enough to distinguish between the ground and excited states with a high degree of confidence in an individual experimental iteration.

As a convenient metric for the timescale over which a projective measurement occurs, we will define τ to be the value of T for which the ground and excited state histograms are

separated by twice their standard deviation, $\Delta V = 2\sigma$. We will also define the signal-to-noise ratio, or strength, of a measurement to be

$$S \equiv \left(\frac{\Delta V}{\sigma}\right)^2 = \frac{4T}{\tau}. \tag{2.12}$$

As shown in figure 2.4, a partial measurement occurs when $T \lesssim \tau$, and a projective measurement occurs when $T \gg \tau$. We can also define a measurement rate

$$\Gamma_{meas} = 1/\tau. \tag{2.13}$$

Here, Γ_{meas} can be thought of as the rate at which information is extracted from the measured system[2].

2.2.2 Measurement fidelity

In a single experimental iteration, a projective qubit measurement should yield a binary result. We extract this binary result from the continuous measurement voltage V_m by choosing a threshold voltage between $\overline{V}_{|1\rangle}$ and $\overline{V}_{|0\rangle}$. If V_m is below (above) the threshold, then we record a measurement result of 1 (0).

The single shot measurement fidelity, F, describes our ability to faithfully resolve the qubit state in a single experimental iteration. It is defined by subtracting the error probabilities $P(1|q_{|0\rangle})$, the probability of measuring 1 when the qubit is prepared in the state $|0\rangle$, and $P(0|q_{|1\rangle})$, the probability of measuring 0 when the qubit is prepared in the state $|1\rangle$, from unity[3]:

$$F \equiv 1 - P(1|q_{|0\rangle}) - P(0|q_{|1\rangle}). \tag{2.14}$$

The amount of overlap between qubit histograms, set by T/τ, determines our ability to distinguish between qubit states, and therefore sets an upper bound on the measurement fidelity, sometimes referred to as the separation fidelity, F_s [16][4]. When $T = \tau$ ($S = 4$), then $F_s = 68\%$. In the right panel of figure 2.4, where $S = 20$, $F_s = 97\%$.

F can be reduced from its maximal value of F_s for a variety reasons. The most common culprit is qubit energy relaxation during the measurement, which limits the length of time

[2]Note that this definition of the measurement rate is consistent with reference [48], but differs by a factor of two from the definition used in [49]. Our choice of definition will become clear when we discuss our experimental measurement setup in later chapters.

[3]Note that reference [16] defines the ground state and excited state fidelities separately as, $F_{|0\rangle} \equiv 1 - P(1|q_{|0\rangle})$ and $F_{|1\rangle} \equiv 1 - P(0|q_{|1\rangle})$. If one were to average these two values, the result would be a more optimistic value of the fidelity than the standard definition of equation (2.14).

[4]Note that in this thesis, F_s is defined according to the same convention as equation (2.14), while reference [16] uses the more optimistic definition.

over which one can gain information about the initial qubit state. Another source of error arises if the measurement is not entirely QND and induces transitions between qubit states. In section 5.4.3, we will discuss the trade-offs between these different sources of error and how to choose the optimal values of T and τ for a given experimental setup.

2.2.3 Quantum efficiency

The measurement fidelity quantifies how well we can use a projective measurement to determine the pre-measurement state of a quantum system. In the case of a partial measurement, we may instead wish to prepare the system in a known initial state, and then determine the state of the system *after* the measurement (or, in the continuous case, *during* measurement). For this purpose, it is necessary to consider the measurement backaction, both in the measured observable and in its conjugate observables. The *quantum efficiency* of a measurement describes how close it comes to ideal Heisenberg-limited backaction.

As discussed in section 2.1.2, an indirect measurement can realize Heisenberg-limited backaction, provided that the classical measurement of the probe contributes nothing to the overall measurement error. However, in practice, it is often difficult to perform a faithful measurement of the probe. In this case, some of information that the probe extracts from the quantum system will be lost to external (unmeasured) degrees of freedom, and the rate is reduced from its ideal value. We define the quantum efficiency of a measurement to be

$$\eta_m \equiv \frac{\Gamma_m}{\Gamma_{m,ideal}}, \tag{2.15}$$

such that in the ideal case $\eta_m = 1$. If η_m is reduced, then a longer measurement time is required to achieve a given separation fidelity. This additional measurement time will also cause additional backaction in the conjugate observable, in this case the phase of the superposition state. In an individual experimental iteration, this additional backaction is random and unknown, and therefore reduces our knowledge of the qubit state after a partial measurement.

To illustrate the concept of quantum efficiency more concretely, we will return to the example of an indirect measurement in which the state of a qubit is encoded in a fluctuating voltage signal. In this case, Γ_m is proportional to S, and thus $\eta_m = S/S_{ideal}$. There are two different mechanisms through which S can be reduced from its ideal value– either ΔV is reduced, or σ is increased. We will explain in detail in chapter 9 how, for our qubit measurement setup, a reduction in ΔV results from cryogenic losses and can be expressed as a collection efficiency η_{col}, while an increase in σ results from added noise in the amplification chain and can be expressed as an amplification efficiency η_{amp}. Therefore, we can express the measurement quantum efficiency as

$$\eta_m = \left(\frac{\Delta V}{\Delta V_{ideal}}\right)^2 \left(\frac{\sigma}{\sigma_{ideal}}\right)^2 = \eta_{col}\eta_{amp}. \tag{2.16}$$

Any additional dephasing due to system's interaction with its environment over the course of a measurement will also reduce our knowledge of the qubit state. It is convenient to describe the environmental contribution[5] to the total quantum efficiency as η_{env}. Then, the total quantum efficiency is expressed as

$$\eta = \eta_m \eta_{env} \tag{2.17}$$

As we will discuss in detail in later chapters, achieving a high quantum efficiency in superconducting qubit measurements is essential to our ability to track individual quantum trajectories.

2.2.4 Positive operator-valued measures

This section will briefly touch upon the language of positive operator-valued measures (POVMs), which are used to formally describe the process of partial measurement. For a more thorough pedagogical introduction to this topic, see reference [38].

A projective (von Neumann) measurement can be described by a set of projection operators $\{P_n = |n\rangle \langle n|\}$, where each n corresponds to one of the possible measurement outcomes. If a quantum system is initially prepared in a state $\rho_i = |\psi\rangle \langle \psi|$, then a measurement outcome n will will occur with probability

$$P(n) = \mathrm{Tr}[P_n \rho P_n], \tag{2.18}$$

and will leave the system in the state

$$\rho_f = |n\rangle \langle n| = \frac{P_n \rho P_n}{\mathrm{Tr}[P_n \rho P_n]}. \tag{2.19}$$

In a similar manner, it is possible to construct a more general set of measurement operators, $\{\Omega_m\}$, with the constraint $\sum_m \Omega_m^\dagger \Omega_m = \hat{I}$, which can describe any measurement, partial or projective, such that the mth measurement outcome occurs with probability

$$P(m) = \mathrm{Tr}[\Omega_m \rho \Omega_m^\dagger], \tag{2.20}$$

leaving the system in the state

$$\rho_f = \frac{\Omega_m \rho \Omega_m^\dagger}{\mathrm{Tr}[\Omega_m \rho \Omega_m^\dagger]}. \tag{2.21}$$

[5]Here, we are referring to the *qubit's* environment. That is, all of the degrees of freedom which couple to the qubit except for the probe.

To describe the qubit measurements discussed in the preceding sections, we can define a set of POVMs, $\{\Omega_V\}$, where V is a continuous parameter and $\int_{-\infty}^{\infty} \Omega_V^\dagger \Omega_V \, dV = \hat{I}$. These operators can be expressed as a weighted sum the qubit state projection operators, which for a measurement time Δt take the form

$$\Omega_V = \frac{1}{\mathcal{N}}[e^{-2k\eta_m \Delta t(1-V)^2} \left|0\right\rangle \left\langle 0\right| + e^{-2k\eta_m \Delta t(-1-V)^2} \left|1\right\rangle \left\langle 1\right|], \tag{2.22}$$

where $\mathcal{N}$ is a normalization factor, $k\eta_m \equiv 1/4\tau$ parametrizes the measurement strength (the measurement operator is given by $\sqrt{k}\sigma_z$), and the 1 and -1 in the exponents are the eigenvalues of the ground and excited states, respectively. The index V corresponds to the integrated measurement voltage V_m, provided that the histograms for the ground and excited state are centered around $V_m = 0$ and V_m is scaled such that $\Delta V/2 = 1$. The probability of each measurement result is given by a sum of two Gaussian histograms, weighted by the ground and excited state populations,

$$P(\Omega_V) = \mathrm{Tr}(\Omega_V \rho \Omega_V^\dagger) = P(|0\rangle)e^{-4k\eta_m \Delta t(1-V)^2} + P(|1\rangle)e^{-4k\eta_m \Delta t(-1-V)^2}. \tag{2.23}$$

Note that in the case of perfectly efficient measurement ($\eta_m = 1$), these POVMs fully describe the measurement backaction, including both the backaction in $\hat{\sigma}_z$ corresponding to the acquisition of qubit-state information and the backaction $\hat{\sigma}_x$ required by Heisenberg's uncertainty principle. In the case of inefficient detection ($\eta_m < 1$), these POVMs will still correctly describe the backaction in $\hat{\sigma}_z$. However, the undetected portion of the measurement signal will also produce backaction, with a strength of $k(1 - \eta_m)$. Since we have no way of tracking this backaction in an individual experimental iteration, we must average over all possible outcomes, which leads to decoherence.

In section 2.3.3, we will use these POVMs to construct the stochastic master equation for a qubit measurement.

2.3 Quantum trajectories

In classical dynamics, a trajectory describes the path that an object takes as it travels through space. More generally, it describes the time evolution of an object's state. By analogy, a quantum trajectory describes how a quantum system evolves in state space. In a closed quantum system, dynamics are governed by the Schrödinger equation. This evolution is deterministic, and thus the resulting trajectories are similar to classical trajectories. On the other hand, the quantum trajectories of an open quantum system, one which interacts with its noisy environment, will evolve stochastically. If an open quantum system starts in a known quantum state, and if we perfectly monitor the fluctuations of its environment, then we can track the state of the system, conditioned on an individual measurement record. In other words, if we understand the backaction of an individual measurement, then this

understanding can be used to update our knowledge of the quantum state after (and during) measurement.

As we will discuss in chapter 9, an understanding of quantum trajectories is essential to applications in measurement-based quantum feedback and control [50]. As experimental quantum information processing is just now entering the age of error correction [51], quantum trajectories may remain relevant to experimental efforts for a long time to come.

This section will begin by discussing the history of quantum trajectories, tracing their evolution from a theoretical tool to an experimental reality. Then we will introduce the general microscopic theory for continuous quantum measurement, which yields a stochastic master equation. Finally, we will discuss an alternative, phenomenological approach to quantum measurement, known as the quantum Bayesian approach [52, 53], which, in certain experimental settings is equivalent to the microscopic model [54] and is particularly convenient when applied to our experiments.

2.3.1 A brief history of quantum trajectories

Techniques to model quantum sources of light date back to the development of the laser and the advent of quantum optics in the 1960's. Perhaps the most well known treatment of open quantum systems is the master equation approach, which describes the time evolution of the system density matrix, $\rho(t)$, and can be used, for example, to model coherent properties of the light emitted from a laser. However, such techniques are unable to describe the instantaneous state of the quantum source itself. Progress in the field of atomic cavity quantum electrodynamics (CQED) motivated a more sophisticated treatment of open quantum systems. The concept of quantum trajectories was first developed in the in the early 1990's as a theoretical tool to model continuously monitored quantum emitters [29, 55, 56]. As Howard Carmichael states in reference [29],

> The words "quantum trajectory" refer to the path of a stochastic wavefunction that describes the state of an optical source, conditioned at each instant on a history of classical stochastic signals realized at ideal detectors monitoring the fields radiated by the source.

For the next decade, quantum trajectories were used primarily in the quantum optics community, as a theoretical tool for numerical simulations of the ensemble behavior of open quantum systems[55, 57]. Typically, the master equation of an open quantum system cannot be solved analytically, and thus numerical solutions are often necessary. For a Hilbert space of dimension N, the density matrix ρ consists of N^2 real numbers, and the computational time required to solve for its time evolution through the master equation scales as N^4 [50]. In contrast, the pure quantum state $|\psi(t)\rangle$ of an individual quantum trajectory can be described by N complex numbers ($2N$ real numbers, at most). By generating an ensemble of M stochastic quantum trajectories, the density matrix can be approximated by their ensemble average, $\mathrm{E}[|\psi(t)\rangle\langle\psi(t)|]$. The time to complete this procedure scales as MN^2, which is much

faster than solving the master equation, provided that $M \ll N^2$. This approach is sometimes referred to as 'quantum Monte Carlo trajectories' [2], because it is often employed in Monte Carlo simulations.

Despite widespread theoretical use, quantum trajectories have only been investigated in a handful of experiments, due in part to the difficulty of performing highly efficient continuous quantum measurements. The earliest experiments to continuously monitor individual quantum systems were in the regime of strong measurement, where the system is quickly projected into an eigenstate of measurement, destroying any information about the phase of a coherent superposition. In such experiments, it is possible to track the 'quantum jumps' between eigenstates [58]. Cavity quantum electrodynamics (CQED) experiments with Rydberg atoms have explored the weak measurement regime, tracking the quantum trajectories of a cavity field as it collapses from a coherent state into a photon number eigenstate [59]. Other CQED experiments have used a cavity probe to continuously track the position of individual Cesium atoms [60]. Quantum trajectories were first considered for solid state systems in the context of a quantum dot qubit monitored in real time by a quantum point contact charge sensor. The equations for tracking the conditional quantum state of the qubit were derived first by Korotkov, using a phenomenological model [52]. Shortly thereafter, equivalent equations were independently derived from a microscopic model [61]. In 2007, the conditional measurement dynamics of a quantum dot were investigated experimentally [62]. More recently, quantum trajectory theory has been used to solve for the conditional evolution of a continuously monitored superconducting qubit [48, 53]. These results, when combined with recent advances in nearly-quantum-limited parametric amplifiers which can be used to achieve highly efficient qubit readout, have enabled a detailed investigation of measurement backaction[63, 64], and set the stage for the quantum trajectories experiments discussed in this thesis [27, 28, 11, 65].

2.3.2 Quantum jumps vs. diffusive trajectories

It is convenient to consider quantum trajectories of a continuously monitored system in the two limiting cases of strong and weak measurement. To make this discussion explicit, we will first extend the discussion of section 2.1.1 describe the process of continuous measurement.

Consider a continuous QND measurement of a qubit, lasting for a total measurement time T and with a measurement strength characterized by the timescale τ. The continuous measurement signal $V(t)$ is broken up into discrete time steps of width Δt and recorded as a sequence $[V_0, V_1, ...V_i, ...V_{n-1}]$, where $n = T/\Delta t$ and

$$V_i = 1/\Delta t \int_{t_i}^{t_i+\Delta t} V(t)\, dt. \tag{2.24}$$

In the limit where $\tau \lesssim \Delta t$, each time step will constitute a (nearly) projective measurement of the qubit. In the absence of free evolution (which here we mean to be any

non-measurement dynamics), the measurements in subsequent time steps will continue to project the qubit into the eigenstate corresponding to the initial measurement result. However, in the case that free evolution occurs on a time scale less than T and greater than τ, this free evolution will manifest itself as discontinuous jumps in the detector signal, corresponding to a changes in the quantum system's state. These 'quantum jumps' can result, for example, from spontaneous emission or from unitary system evolution (such as a Rabi drive), and occur stochastically on an average timescale, τ_{jump}. Quantum jumps were first observed in experiments with trapped ions in the 1980's [58] but have only recently been observed in solid state systems [22]. For a comprehensive account of quantum jumps in superconducting qubits, see Daniel Slichter's PhD. thesis [66].

If the measurement rate is reduced, such that $\Delta t \ll \tau$, then the measurement in each time step will only slightly perturb the qubit state. The measurement record can be used to track the state of the qubit as it is gradually projected by the weak measurement. This gradual random process is often referred to as 'diffusive' evolution, and is most easily modeled when uncontrolled free evolution (such as spontaneous emission) occurs on a timescale much slower than the total measurement time T. The results presented in this thesis will focus on such diffusive quantum trajectories.

Note that in the quantum jump limit, the the qubit is always in an eigenstate of measurement, and therefore quantum coherence plays no role in the dynamics. Therefore, the quantum efficiency of the measurement is only relevant in that it contributes to τ. Any extra backaction in the phase of the qubit state due to a non-ideal measurement efficiency is irrelevant to jump-like dynamics. In contrast, diffusive trajectories track the qubit state populations *and* the phase of a superposition state. Therefore, near-unity quantum efficiency is an essential prerequisite to observing diffusive quantum trajectories.

To recap, quantum jumps are most readily observed when

$$\tau \lesssim \Delta t < \tau_{jump} < T, \tag{2.25}$$

and diffusive trajectories are most readily observed when[6]

$$\Delta t \ll \tau < T < \tau_{jump}. \tag{2.26}$$

A complete description of any real open quantum system should account for both jump-like and diffusive behavior. For superconducting qubits in a cavity QED architecture, a full theory for the conditional state of qubit under measurement was worked out in reference [48]. They derive a stochastic master equation for the system and discuss the crossover between jump-like and diffusive trajectories. In the following section we will introduce the general concept of stochastic master equations.

[6]With the notable exception that when τ_{jump} is set predominantly by a known unitary evolution, in which case diffusive quantum trajectories can also be observed when $\Delta t < \tau, \tau_{jump} < T$. In chapter 7 we will demonstrate how to reconstruct such trajectories in the presence of a Rabi drive.

2.3.3 Stochastic master equations

Here, we will provide an introduction to stochastic master equations (SMEs). For a more thorough pedagogical introduction to the topic, see reference [67], which considers the case of qubit measurements, and [38], which derives a more general form of the SME that applies to (continuous) measurements of a continuous variable. The SME for superconducting qubit measurements in a cavity QED architecture is derived in reference [48]. Here, we will derive an SME, following mostly the derivation of reference [38], but adapting it for the type of qubit measurements discussed in section 2.2.1.

As in the previous section, we will consider a qubit which is continuously monitored, yielding a measurement record consisting of a sequence of measurement results V_i, each of which occurs over a time step Δt. Formally, each of these measurements can be represented by a POVM, Ω_{V_i}. In equation (2.23), we showed that the probability of acquiring a measurement result corresponding to the measurement result V_i is given by a sum of two Gaussian distributions, weighted by the qubit state populations. In the limit of weak measurement ($\Delta t \ll \tau$), this distribution can be replaced by a single Gaussian centered about the expectation value of σ_z,

$$P(\Omega_{V_i}) \approx e^{-4k\eta_m \Delta t (V_i - \langle\sigma_z\rangle)^2}. \tag{2.27}$$

This equation highlights the fact that V_i can be thought of as simply a noisy estimate of $\langle\sigma_z\rangle$. In fact, it can be written as a sum of $\langle\sigma_z\rangle$ and a stochastic quantity,

$$V_i = \langle\sigma_z\rangle + \frac{\Delta W}{\sqrt{8k\eta_m}\Delta t}, \tag{2.28}$$

where ΔW is a zero-mean, Gaussian random variable with variance Δt[38].

Just as we re-expressed $P(\Omega_{V_i})$ in the limit of weak measurement, we can also re-express the expression for Ω_{V_i} found in equation (2.22) as

$$\Omega_{V_i} \propto e^{-2k\eta_m \Delta t (V_i - \sigma_z)^2} \tag{2.29}$$

As an initial step to deriving the stochastic equation of motion for the system, let us consider the change in the qubit state $|\psi(t)\rangle$ after measuring for a time Δt resulting in a measurement outcome V_i, represented by the operator Ω_{V_i}. For simplicity, we will start by assuming that $\eta_m = 1$. The resulting qubit state is given by

$$|\psi(t+\Delta t)\rangle \propto \Omega_{V_i} |\psi(t)\rangle \tag{2.30}$$

$$\propto e^{-2k\Delta t (V_i - \sigma_z)^2} |\psi(t)\rangle \tag{2.31}$$

$$\propto e^{-2k\Delta t \sigma_z^2 + \sigma_z [4k\langle\sigma_z\rangle \Delta t + \sqrt{2k}\Delta W]} |\psi(t)\rangle . \tag{2.32}$$

Expanding this expression to first order in Δt, and keeping terms to second order in ΔW,[7] we find

$$|\psi(t+\Delta t)\rangle \propto \{1 - 2k\Delta t\sigma_z^2 + \sigma_z[4k\langle\sigma_z\rangle\Delta t + \sqrt{2k}\Delta W + k\sigma_z(\Delta W)^2]\}\,|\psi(t)\rangle\,. \tag{2.33}$$

We then take the continuum limit, where $\Delta t \to 0$, by setting $\Delta t = dt, \Delta W = dW$, and $(\Delta W)^2 = dt$, which results in

$$|\psi(t+dt)\rangle \propto \{1 - [k\sigma_z^2 - 4k\sigma_Z\langle\sigma_z\rangle]dt + \sqrt{2k}\sigma_z dW\}\,|\psi(t)\rangle\,. \tag{2.34}$$

By normalizing $|\psi(t+dt)\rangle$, keeping only terms up to order dt and $(dW)^2$ in the normalization factor, and then expressing $|\psi(t+dt)\rangle = |\psi(t)\rangle + d\,|\psi\rangle$, we arrive at the differential equation

$$d\,|\psi\rangle = \{-k(\sigma_z - \langle\sigma_z\rangle)^2 dt + \sqrt{2k}(\sigma_z - \langle\sigma_z\rangle)dW\}\,|\psi(t)\rangle\,. \tag{2.35}$$

This differential equation is known as the *Stochastic Schrödinger equation* (SSE). It describes evolution of the qubit state in a time interval dt, given that a measurement result $dV = \langle\sigma_z\rangle dt + dW/\sqrt{8k}$ is acquired during that interval. By integrating over a measurement record $dV(t)$, the SSE can be solved to reconstruct the conditional evolution of the quantum state, given by the quantum trajectory $|\psi(t)\rangle$.

Note that the SSE only applies to a pure state ψ. When generalized to instead describe the density operator ρ, we arrive at the stochastic master equation (SME) [38, 68]

$$\begin{aligned} d\rho &= (d\,|\psi\rangle)\,\langle\psi| + |\psi\rangle\,(d\,\langle\psi|) + (d\,|\psi\rangle)(d\,\langle\psi|) \\ &= -k[\sigma_z[\sigma_z,\rho]]dt + \sqrt{2k}(\sigma_z\rho + \rho\sigma_z - 2\langle\sigma_z\rangle\rho)dW. \end{aligned} \tag{2.36}$$

Finally, we will return to the case of inefficient detection. As we discussed in section 2.2.4, inefficient detection can be modeled as two concurrent measurement processes– an observed measurement of strength $k\eta$ and an unobserved measurement of strength $k(1-\eta)$ (such that the total measurement strength is $k\eta + k(1-\eta) = k$). There will be two separate contributions to the measurement backaction,

$$dV_1 = \langle\sigma_z\rangle dt + dW_1/\sqrt{8k\eta} \tag{2.37}$$

$$dV_2 = \langle\sigma_z\rangle dt + dW_2/\sqrt{8k(1-\eta)}, \tag{2.38}$$

[7] A rough explanation for this choice can be made by dimensional analysis. ΔW has units of $\sqrt{t}$, and therefore second order terms in ΔW are as significant as first order terms in Δt. A more complete justification of this process requires the use of stochastic calculus, and is explained in reference [38].

which will lead to the following SME:

$$\begin{aligned} d\rho = &- k[\sigma_z[\sigma_z,\rho]]dt \\ &+ \sqrt{2k\eta}(\sigma_z\rho + \rho\sigma_z - 2\langle\sigma_z\rangle\rho)dW_1 \\ &+ \sqrt{2k(1-\eta)}(\sigma_z\rho + \rho\sigma_z - 2\langle\sigma_z\rangle\rho)dW_2. \end{aligned} \tag{2.39}$$

We recover the observer's knowledge of the qubit state by averaging over all possible outcomes of the unobserved measurement dW_2. Since ρ and dW_2 are statistically independent, the average value of ρdW is zero, and we can ignore the last term in equation (2.39), leading to the SME

$$d\rho = -k[\sigma_z[\sigma_z,\rho]]dt + \sqrt{2k\eta}(\sigma_z\rho + \rho\sigma_z - 2\langle\sigma_z\rangle\rho)dW_1. \tag{2.40}$$

If we consider the limit of $\eta = 0$, we recover the standard non-stochastic master equation describing qubit decoherence

$$\frac{d\rho}{dt} = -k[\sigma_z[\sigma_z,\rho]]. \tag{2.41}$$

2.3.4 Bayesian state update

Although the stochastic master equation provides an elegant and complete description of the conditional dynamics of an open quantum system, it is often difficult to solve due the fact that it is nonlinear[8] in ρ. Using the appropriate approximations, the SME describing qubit measurements can be solved and applied to our experiments [65]. However, for the work described in this thesis, we choose to analyze our measurement records using a simpler approach known as the quantum Bayesian formalism[52, 53]. This approach requires a phenomenological assumption regarding the effect of measurement on the qubit coherence, but nevertheless it correctly predicts the diffusive quantum trajectories in our experiments. As we will discuss in chapter 7, the experimental verification of these trajectories is a major result of this thesis.

Consider a qubit initially prepared in the state $\rho(t=0)$. The idea is to apply Bayes rule of conditional probabilities,

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}, \tag{2.42}$$

to update our knowledge of the qubit state conditioned on an integrated measurement result V_m, as defined in equation (2.11). From Bayes rule, we have

[8]Note that the last term in equation 2.36 is proportional to $\langle\sigma_z\rangle\rho$. Since $\langle\sigma_z\rho\rangle$ depends on ρ, this term is nonlinear in ρ.

$$P(0|V_m) = \frac{P(V_m|0)P(0)}{P(V_m)}, \text{ and} \tag{2.43}$$

$$P(1|V_m) = \frac{P(V_m|1)P(1)}{P(V_m)}. \tag{2.44}$$

Here, the initial probabilities for finding the qubit in the ground or excited state are given by $P(0) = \rho_{00}(t=0)$ and $P(1) = \rho_{11}(t=0)$. The probabilities of measuring a particular value of V_m when the qubit is prepared in the ground or excited state are Gaussian distributions of width σ centered around $\pm\Delta V/2$. Taking the ratio of the two conditional probabilities in equation (2.43), we find that

$$\frac{\rho_{11}(t)}{\rho_{00}(t)} = \frac{\rho_{11}(0)}{\rho_{00}(0)} \frac{\exp[-(V_m(t) + \Delta V/2)^2/2\sigma^2]}{\exp[-(V_m(t) - \Delta V/2)^2/2\sigma^2]}. \tag{2.45}$$

Making use of the formula $\langle\sigma_z\rangle = \rho_{00} - \rho_{11}$, and the definition $S \equiv \Delta V^2/\sigma^2$, for a qubit initially in an equal superposition of $|0\rangle$ and $|1\rangle$ we find that

$$\langle\sigma_z\rangle(t) = \tanh\left(\frac{V_m(t)S}{2\Delta V}\right). \tag{2.46}$$

Note that so far we have only used a classical rule of conditional probabilities to determine how qubit populations (i.e. diagonal components of the density matrix) evolve under measurement. In the case where the phase of the qubit state is unaffected by the measurement (Heisenberg-limited backaction), then the resulting off-diagonal components of the density matrix are completely specified by the the change in the diagonal components and our knowledge of the initial state:

$$\rho_{01}(t) = \rho_{01}(0) \frac{\sqrt{\rho_{00}(t)\rho_{11}(t)}}{\sqrt{\rho_{00}(0)\rho_{11}(0)}}. \tag{2.47}$$

Korotkov refers to this non-unitary backaction as "spooky" (or quantum) backaction[53].

To account for an inefficient measurement, we can introduce and additional dephasing term $e^{-\gamma t}$ to the off-diagonal components of ρ, characterized by a dephasing rate γ, which we will define in chapter 6 in terms of cavity QED parameters and the the measurement efficiency η_m.

We would also like to account for a specific type of non-Heisenberg backaction[9], the motivation for which will become clear when we discuss the details of our measurement

[9]By non-Heisenberg backaction, we mean backaction in the conjugate variable beyond what is required by the uncertainty principle, and of a different origin.

apparatus. It is possible to devise a measurement that results in a stochastic shift Δ_ϕ in the phase of the qubit superposition. Korotkov refers to this effect as "realistic" backaction, since its effect can be described by a unitary transformation. This backaction can be accounted for phenomenologically by simply applying a rotation to the off-diagonal components of the density matrix.

As in the preceding section, instead of treating integrating the full measurement record, we could instead choose to break up the measurement record into short intervals of length Δt, and apply the Bayesian state update procedure sequentially, using the updated state after each time step as the initial state for the following time step. In the case where the qubit is simultaneously measured and driven, such a treatment will become necessary.

Chapter 3

Superconducting qubits and circuit QED

This chapter provides an introduction to superconducting qubits, with an emphasis on the transmon qubit used in the experiments covered in this thesis. We also introduce basic principles of the circuit quantum electrodynamics architecture, in which a qubit is coupled to a microwave frequency cavity, and explain how the cavity field can be used to indirectly probe the qubit state, realizing a QND measurement.

These topics are covered in great depth in other works, including Daniel Slichter's and David Schuster's theses [66, 69], as well as in a number of comprehensive review papers [70, 71, 72, 73]. Therefore, here we will focus primarily on the basic results necessary to provide a physical description of our qubit measurements.

3.1 Quantization of electromagnetic fields and electrical circuits

To realize a qubit requires an individual quantum system that can be manipulated and measured by an experimenter. It follows naturally that microscopic quantum systems such as trapped ions and nuclear spins are good candidates for qubits. Indeed, some of the earliest advances in quantum computing were performed using nuculear magnetic resonance, and to this day trapped ions remain very promising due their long coherence times and high fidelity gate operations. But why should we expect a macroscopic object like an electrical circuit to behave as quantum as an individual quantum system?

The answer to this question is very similar to the justification for the quantization of the electromagnetic field, a well known phenomenon which is fundamental to the fields of quantum electrodynamics and quantum optics [31]. The basic procedure is to find the system Hamiltonian H by expressing the total energy in terms of the system parameters, and then to choose the appropriate generalized coordinates q_k and momenta p_k which obey the Hamilton equations of motion,

$$\frac{\partial H}{\partial q_k} = -p_k \tag{3.1}$$

$$\frac{\partial H}{\partial p_k} = \dot{q_k}. \tag{3.2}$$

Any q_k and p_k which satisfy this relationship, are said to be 'canonical' variables and can be treated as quantum mechanical operators $\hat{p}_k$ and $\hat{q}_k$, which obey the commutation relation $[\hat{p_k}, \hat{q_k}] = i\hbar$.

For the electromagnetic field, the energy can be expressed as a sum of infinitely many modes indexed by the letter k, each oscillating at a frequency ω_k. Each mode can be described by the generalized coordinate $[\hat{X_1}]_k$ and momentum $[\hat{X_2}]_k$, which happen to represent the quadrature amplitudes[1], or sinusoidal and cosinusoidal components of an electromagnetic wave oscillating at frequency ω_k. To be explicit, we can express the electric field of a wave oscillating at frequency ω_k as

$$E(t) \propto \mathrm{Re}[\tilde{A}(t)e^{i\omega_k}]\,, \text{ where} \tag{3.3}$$

$$\tilde{A}(t) = [\hat{X_1}]_k(t) + i[\hat{X_2}]_k(t). \tag{3.4}$$

Furthermore, the Hamiltonian for the electromagnetic field can be expressed as a sum of harmonic oscillators

$$H = \sum_k \hbar\omega_k(a_k^\dagger a_k + 1/2), \tag{3.5}$$

by defining for each mode

$$a \equiv \frac{\hat{X_1} + \hat{X_2}}{2}\,, \text{ and} \tag{3.6}$$

$$a^\dagger = \frac{\hat{X_1} - \hat{X_2}}{2} \tag{3.7}$$

(from here on we will drop the k subscripts and assume we are working with a single mode). X_1 and X_2 are dimensionless conjugate observables and obey the commutation relation

$$[\hat{X_1}, \hat{X_2}] = 2i. \tag{3.8}$$

[1]Here, $\hat{X_1}$ and $\hat{X_2}$ refer to dimensionless quadrature amplitudes. This is a common definition used in the quantum optics literature, but sometimes $\hat{X}$ and $\hat{P}$ are used instead, which are often defined as dimension-full quantities. In microwave electronics, the preferred nomenclature for the quadrature amplitudes is I and Q.

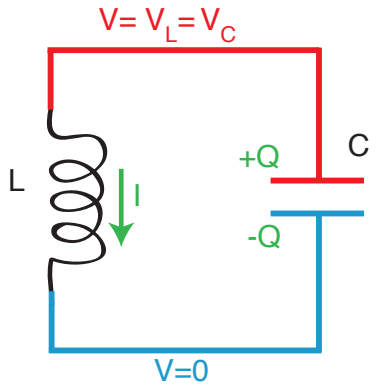

Figure 3.1: An LC resonator. There is a an equal voltage drop across the inductor and the capacitor ($V_L = V_C$), and a positive current is defined to go against the direction of a positive voltage.

Because the quadrature amplitudes are non-commuting observables, they obey the uncertainty relation

$$\Delta X_1 \Delta X_2 \geq 1. \tag{3.9}$$

This uncertainty relation implies that the quadrature amplitudes of the electromagnetic field are constantly fluctuating. These fluctuations are known as vacuum fluctuations, and can be thought of as the zero-point fluctuations of the harmonic oscillator used to model each mode.

We will now follow a similar procedure to quantize a simple circuit, the parallel LC resonator depicted in figure 3.1. There is an equal voltage drop across the inductor and the capacitor ($V_L = V_C$), which can be expressed both in terms of the inductance and the capacitance:

$$V_C = Q/C, V_L = L\dot{I}, \tag{3.10}$$

where Q is the charge on the capacitor, I is the current through the inductor, and $I = -\dot{Q}$. The total energy in the resonator can be expressed as

$$E = \frac{1}{2}LI^2 + \frac{1}{2}CV^2. \tag{3.11}$$

By defining the branch flux Φ to be the time integral of the voltage across the inductor (setting V to 0 at $t = -\infty$),

$$\Phi = \int_{-\infty}^{t} V(t')\, dt', \tag{3.12}$$

we can express the Hamiltonian as

$$H = \frac{\Phi^2}{2L} + \frac{Q^2}{2C}. \tag{3.13}$$

We then find that

$$\frac{\partial H}{\partial \Phi} = \Phi/L = I = -\dot{Q}, \text{ and} \tag{3.14}$$

$$\frac{\partial H}{\partial Q} = Q/C = L\dot{I} = \dot{\Phi}. \tag{3.15}$$

Thus, Φ and C satisfy equations (3.1), and can therefore be replaced with the operators $\hat{\Phi}$ and $\hat{C}$, which satisfy the commutation relation

$$[\hat{\Phi}, \hat{Q}] = i\hbar. \tag{3.16}$$

Then, as with a mode of the electromagnetic field, the Hamiltonian for an electrical resonator maps onto that of a quantum harmonic oscillator, and the resonator field can be expressed in terms of creation and annihilation operators a and $a^\dagger$, taking the same form as a single mode of equation (3.5), with a resonance frequency $\omega = 1/\sqrt{LC}$. For a formal and more thorough treatment of this derivation, see reference [74].

3.2 Superconducting qubits

We have demonstrated that, in principle, an electrical circuit can be quantized, but of course, most circuits will not function as qubits. What, then, are the basic requirements for making a qubit from an electrical circuit? First of all, the qubit needs to be cooled to its quantum ground state. This is essential for state preparation and to prevent unwanted thermal transitions between states. The temperature required for ground state cooling will depend on the qubit transition frequency (i.e. the resonance frequency of the circuit). As the ratio of occupation between the ground and first excited state is given by the Boltzmann factor $\mathrm{Exp}[\hbar\omega_{01}/k_B T]$, if $\omega_{01} = 6$ GHz, where $\hbar\omega_{01}$ is the energy difference between the ground and excited states, then the temperature requirement for ground state cooling is $T \ll \hbar\omega_{01}/k_B \approx 300\,\mathrm{mK}$.

The second essential requirement to realize a qubit is anharmonicity. In atomic physics, the standard method to manipulate the the state of an atom is to send in laser light at

a frequency ω_{01}, which will drive resonant transitions between the $|0\rangle$ and $|1\rangle$ states[2]. The quantized LC resonator depicted in figure 3.1 will have equally spaced energy levels, separated in energy by $\hbar\omega$. This means that microwave light at a frequency ω can drive transitions between all levels. Therefore, it is impossible to choose any two levels to manipulate in isolation. If we wish to drive closed transitions within a two-level qubit subspace, then we need to make the level spacing unequal, or *anharmonic*. This is achieved by replacing the linear inductor with a *non-linear* inductance. The absolute anharmonicity α and relative anharmonicity α_r of a quantum circuit can be defined as

$$\alpha \equiv E_{12} - E_{01}, \;\; \alpha_r \equiv \alpha / E_{01}, \tag{3.17}$$

where the energy eigenstates are labeled $[0, 1, 2, ...]$, and $E_{ij} \equiv \hbar\omega_{ij}$. As we will discuss later, α determines the rate at which we can resonantly drive the $0-1$ transition without exciting higher energy levels.

The final requirement to make an electrical circuit behave as a qubit is low dissipation. Dissipation sets the maximum timescale over which a qubit will remain in a coherent superposition. Since the anharmonicity sets the maximal speed of qubit gate operations, the ration of anharmonicity to loss will determine how may coherent gate operations can be performed over the qubit lifetime. Note that in the derivation in section 3.1 we ignored loss by considering only an LC resonator, rather than an RLC resonator. In a classical circuit, energy is dissipated as current flows across the resistor, such that any energy stored in the resonator will eventually decay. The same is true for a qubit. In the frequency domain, dissipation leads to a broadening of the transition frequencies ω_{ij}. In the time domain, loss is characterized by the energy relaxation time T_1, the average exponential timescale over which a qubit prepared in the excited state will decay to the ground state. Superconducting metals are a natural choice of material for building quantum circuits, because of their low loss, which allows for long T_1 times. In fact, as we will discuss in appendix C, the predominant source of loss in superconducting qubits does not come from the superconducting components themselves, but from other sources such as surface dielectric loss at the interface between superconducting films and their substrates, non-equilibrium quasiparticles, and vortices of trapped flux.

We will now turn our attention to the low loss non-linear circuit component which allows us to achieve anharmonicity. This element, known as a Josephson junction[75], consists of a thin (and narrow) insulating barrier connecting two superconducting leads. The current I flowing through the junction and the voltage V across the junction are described by the Josephson relations

[2]While atomic physics experiments often work with states other than the ground and first excited states, any two arbitrary states can nevertheless be labeled $|0\rangle$ and $|1\rangle$.

$$I = I_0 \sin(\delta) \tag{3.18}$$

$$V = \frac{\Phi_0}{2\pi}\dot{\delta}, \tag{3.19}$$

where I_0 is the critical current, Φ_0 is the magnetic flux quantum $h/2e$, and δ is the difference between the phases of the superconducting order parameter on each side of the junction.

By combining these relations and making use of the definition of inductance, $V = L\dot{I}$, we can express the Josephson inductance as

$$L_J = \frac{\Phi_0}{2\pi I_0 \cos\delta} \equiv \frac{L_{J0}}{\cos\delta}, \tag{3.20}$$

where we have defined $L_{j0} \equiv \Phi_0/2\pi I_0$. Because of the $\cos(\phi)$ dependence, L_j will depend on the the current passing through the junction, and therefore the Josephson junction will behave as a nonlinear inductor.

Furthermore, we can realize a *tunable* nonlinear inductance by connecting two Josephson junctions in a loop of superconducting metal. This circuit element is known as a superconducting quantum interference device (SQUID)[76]. In the limit where the loop inductance is negligible compared to the Josephson inductance, a SQUID can be modeled as a single junction with a critical current I_c which depends on the magnetic flux Φ_{ext} threading the SQUID loop:

$$I_c(\Phi_{ext}) = 2I_0 \left|\cos\left(\frac{\pi\Phi_{ext}}{\Phi_0}\right)\right|, \tag{3.21}$$

where I_0 is the critical current of a single junction.

As depicted in figure 3.2a, a general superconducting qubit can be modeled as a Josephson inductance L_j in parallel with a capacitance C_Σ and a linear inductance L. Here, C_Σ refers to the parallel combination of the intrinsic junction capacitance and any additional shunt capacitance. To find a quantum mechanical description of a qubit, we need write down its Hamiltonian, by expressing the energy across each circuit component. Let us first focus our attention on the Josephson junction. By expressing the energy in the junction U as the time integral of IV from negative infinity to a time t (assuming zero energy at $t = -\infty$), and using the Josephson relations, we arrive at [66]

$$U = E_j(1 - \cos\phi), \tag{3.22}$$

where the Josephson energy scale $E_j \equiv \Phi_0 I_0/2\pi = \hbar I_0/2e$. This cosinusoudal potential can be tilted by applying an external phase bias across the junction (using either a bias current or a bias flux), leading to the double-welled potential depicted in figure 3.2b. This

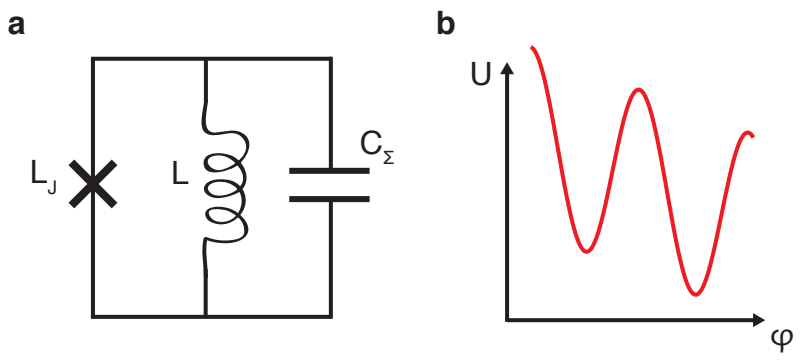

Figure 3.2: Superconducting qubits. a) A general schematic of a superconducting qubit, a Josephson inductance L_j in parallel with a capacitance C_Σ and a linear inductance L. b) A schematic of the double well potential which arises when the cosine Josephson potential is tilted by applying an external phase bias.

tilted potential gives rise to a variety of qubit designs, including the phase qubit, and the fluxonium. However, as the experiments presented in this thesis deal only with an unbiased qubit, the transmon, we will only need to focus our attention on a single well of the potential given in equation (3.22).

3.3 The transmon qubit

The transmon is a robust and simple qubit design which has yielded some of the longest reported energy relaxation times [77, 78, 79]. This section will offer a brief conceptual overview of the transmon, focusing on the basic parameters needed to describe its behavior. A thorough theoretical description of the transmon can be found in reference [80].

A circuit schematic for the transmon is depicted in figure 3.3a. The basic components are a Josephson junction, described by and an inductive energy scale E_J and a capacitance C_J, and an additional shunt capacitance C_B. The transmon is coupled to an external voltage V_g, through a gate capacitance C_g. The total capacitance shunting the junction is given by $C_\Sigma = C_B + C_J + C_g$. We will define $\hat{n}$ to be the number of Cooper pairs (superconducting charge carriers with charge $2e$) on the superconducting 'island', the area in the top righthand portion of figure 3.3a. The value of $\hat{n}$ can only change when a Cooper pair tunnels across the Josephson junction. We can also express the gate voltage in units of Cooper pair number by defining the gate number n_g through the relation $V_g = 2en_g/C_g$. The voltage across C_Σ can be expressed as $2e(\hat{n} - n_g)/C_\Sigma$. Therefore, the total capacitive energy $C_\Sigma V^2/2$, is equal to $4E_C(\hat{n} - n_g)^2$, where we have defined $E_C = e^2/2C_\Sigma$. Combining this with equation (3.22) and dropping the constant offset term, we arrive at the Hamiltonian

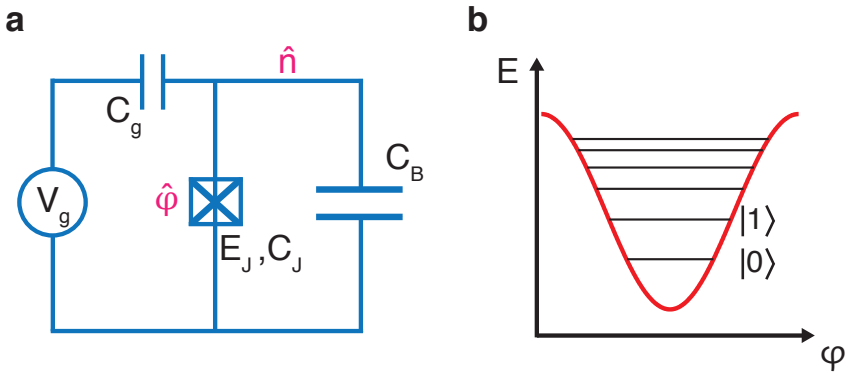

Figure 3.3: Transmon qubit. a) Circuit schematic for a transmon qubit. A Josephson junction, characterized by a capacitance C_j and a Josephson energy E_j, is shunted by an additional capacitance C_B. It is coupled to an external voltage V_g, through a gate capacitance C_g. The canonically conjugate operators for this circuit are the the number of cooper pairs on the island, $\hat{n}$, and the superconducting phase drop across the junction, $\hat{\phi}$. b) A single well of the cosinusoidal potential. For typically transmon parameters, there are roughly 5-10 energy levels allowed within the well.

$$H = 4E_C(\hat{n} - n_g)^2 - E_J\cos\hat{\phi}. \tag{3.23}$$

As explained in reference [80], the energy eigenvalues of this Hamiltonian can be expressed analytically in terms of the Mathieu functions (which are built in to Mathematica and easily computed), and are given by

$$E_m(n_g) = E_C a_{2[n_g+k(m,n_g]}(-E_J/2E_C), \tag{3.24}$$

where m is a whole number that indexes the eigenstates, $a_r(q)$ is Mathieu's characteristic value, and $k(m, n_g)$ is a sorting function explained in the appendix of reference [80].

This general solution to the qubit Hamiltonian applies regardless of our choice of the circuit parameters E_J and E_C. When $E_J/E_C \sim 1$, the eigenenergies depend strongly on n_g, and therefore the qubit is very susceptible to charge noise. However, as E_J/E_C is increased, this dependence is exponentially suppressed. Typically, transmons operate with $E_J/E_C \sim 100$, and the effect of n_g is so insignificant that it is not necessary to control it. Therefore, transmons require no explicit reference to ground, which makes their design very simple to implement.

For typical transmon parameters, the potential allows for $\sim 5-10$ bound eigenstates, as depicted in figure 3.3b. The qubit frequency ω_{01} is given by the energy difference between the $m = 0$ and $m = 1$ levels. A qubit with $\omega_{01}/2\pi \approx 6\,\text{GHz}$ and $E_J/E_C = 100$, would require $E_J/h \approx 20$ GHz and $E_C/h \approx 200$ MHz. The energy levels (and hence the qubit frequency)

can be made tunable by replacing the Josephson junction with a SQUID and applying an external magnetic field.

The one downside to the transmon is its limited anharmonicity, which is approximately given by

$$\alpha \simeq -E_C, \text{ and } \alpha_r \simeq -(8E_J/E_C)^{-1/2}. \tag{3.25}$$

A typical value of $\alpha = -200$ MHz is at least an order of magnitude less than for other qubit designs such as the flux qubit. Therefore, it is difficult to perform fast gate operations on the transmon without causing population leakage to higher levels[81], although such as DRAG pulses [82] have been developed to reduce this this effect. Such leakage, even at the level of a few percent, is a serious concern for fault tolerant quantum computing schemes such as the surface code[19, 15].

3.4 Cavity quantum electrodynamics

Cavity quantum electrodynamics (CQED) [83] is a field of atomic physics which studies the interaction between an individual atom (or a collective quantum state of an ensemble of atoms) and a single discrete mode of light confined to an electromagnetic cavity. Such a system provides a valuable tool for studying the fundamental quantum mechanics of light-matter interactions and for engineering quantum states [2]. In this section, we will explain how a a CQED architecture can be implemented using superconducting circuits [84, 85]. This architecture, known as circuit quantum electrodynamics (circuit QED) provides a fascinating macroscopic reallization of atomic physics experiments. Due to flexibility in the choice of circuit parameters and stronger coupling strengths than can be achieved in atomic CQED experiments, circuit QED offers a powerful testbed for quantum optics. For example, as we will discuss in chapter 8, we recently used a circuit QED experiment [25] to investigate the interaction between squeezed light and matter, a long sought result which has remained elusive in atomic physics experiments. Furthermore, perhaps most importantly, the circuit QED architecture provides a straightforward method of realizing indirect QND measurements of the state of a superconducting qubit [84, 80]. The primary purpose of this section will be to provide a physical picture for such measurements.

3.4.1 The Jaynes-Cummings Hamiltonian

One reason why CQED provides such a rich testbed for quantum optics is that the atom-cavity system can be described by a simple theoretical framework known as the Jaynes-Cummings model [31], which, under some straightforward approximations, is analytically solvable. As depicted in figure 3.4, this model considers a single cavity mode of frequency ω_r which couples with a strength g to a two-level atom (or qubit) with a transition frequency ω_q. Additionally, the cavity decays at a rate κ and the qubit can decay into environmental

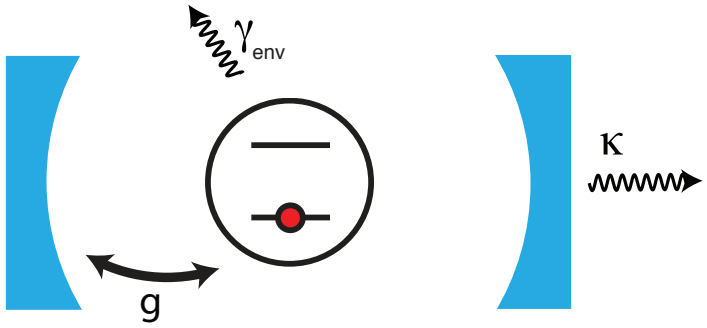

Figure 3.4: Cavity quantum electrodynamics. A qubit couples to a single mode of an electromagnetic cavity at a rate g. The cavity mode decays at a rate κ, and the qubit can decay into environmental degrees of freedom at a rate γ_{env}.

(non-cavity) degrees of freedom at a rate γ_{env}. Ignoring the decay terms, the Hamilton for the qubit cavity system can be written as a sum of the qubit energy, the cavity energy, and a dipole interaction term H_{int}

$$H = H_q + H_c + H_{int} \tag{3.26}$$

$$= \frac{1}{2}\hbar\omega_q\sigma_z + \hbar\omega_r(\hat{a}^\dagger\hat{a} + 1/2) + \hbar g(\hat{a} + \hat{a}^\dagger)(\hat{\sigma}_+ + \hat{\sigma}_-), \tag{3.27}$$

where $\hat{a}^\dagger$ and $\hat{a}$ are the cavity creation and annihilation operators, and $\hat{\sigma}_+$ and $\hat{\sigma}_-$ are the qubit raising and lowering operators given by $(\hat{\sigma}_x \pm i\hat{\sigma}_y)/2$. This expression can be simplified by invoking the rotating wave approximation, which allows us to ignore the interaction terms that don't conserve excitation number, resulting in the Jaynes-Cummings Hamiltonian

$$H_{JC} = \frac{1}{2}\hbar\omega_q\sigma_z + \hbar\omega_r(\hat{a}^\dagger\hat{a} + 1/2) + \hbar g(\hat{a}\hat{\sigma}_+ + \hat{a}^\dagger\hat{\sigma}_-). \tag{3.28}$$

Here, the interaction term allows a single quantum of energy to be exchanged between the qubit and the cavity. It is a straightforward matter to diagonalize this Hamiltonian and solve for the energy eigenstates. The transition frequencies between energy eigenstates are depicted in figure 3.5 as a function of qubit frequency. The dashed lines show the tranition frequencies for the uncoupled qubit (red) and cavity (green). When the qubit-cavity detuning $\Delta \equiv \omega_q - \omega_r$ is large compared to g, there is are two distinct transition frequencies for the coupled system, corresponding to the qubit and the cavity, which are only slightly shifted from their bare (uncoupled) values. However, as Δ approaches zero, the two transition frequencies will remain separated by $2g$. This avoided level crossing is known as the "vacuum Rabi splitting". In this resonant regime, the energy eigenstates are a hybrid superposition of

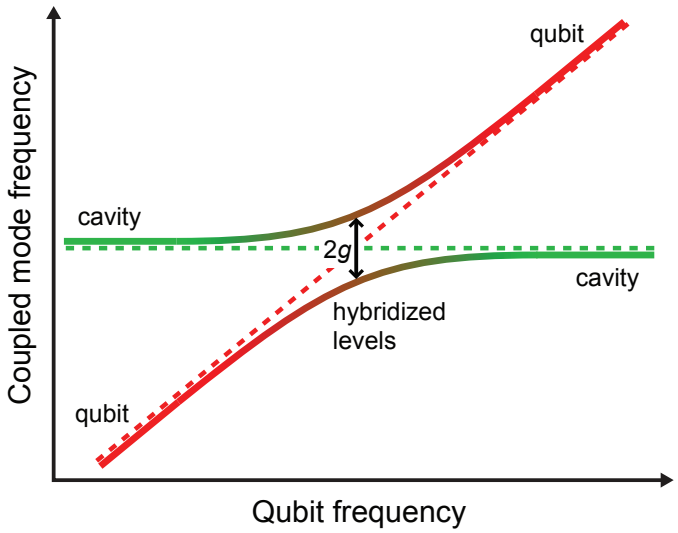

Figure 3.5: Avoided crossing of the Jaynes-Cummings Hamiltonian. Dashed lines indicate the uncoupled ('bare') qubit and cavity frequencies, ω_a and ω_r. Solid lines represent the eigenenergies of the coupled qubit-cavity system, also known as 'dressed states'. Figure adapted from reference [66].

qubit and cavity states. An experiment that we will discuss in chapter 8 makes use of these hybridized levels. However, all other experiments discussed in this thesis take place in the large detuning or *dispersive* regime, which provides a straightforward mechanism for QND measurement.

3.4.2 Dispersive measurements

In the dispersive limit, where $|\Delta| \gg g$, the qubit and the cavity do not directly exchange energy. However, they still interact, such that the cavity field can be used to probe the qubit state, realizing an indirect measurement. To gain some insight into the dispersive qubit-cavity interactions, we can expand H_{int} to second order in g/Δ, arriving at the approximate dispersive interaction term

$$H_{int,disp} = -\frac{\hbar g^2}{\Delta}\hat{a}^\dagger \hat{a}\hat{\sigma}_z. \tag{3.29}$$

Note that here the $\hat{\sigma}_+$ and $\hat{\sigma}_-$ terms have been replaced by $\hat{\sigma}_z$. Therefore, this Hamiltonian clearly satisfies the formal criteria for a QND measurement of $\hat{\sigma}_z$, given by equations (2.6) and (2.7).

To illustrate the physical mechanism for such a QND measurements, we can rearrange the full qubit-cavity Hamiltonian in the dispersive limit:

$$H_{disp} = \frac{1}{2}\hbar\omega_q\hat{\sigma}_z + \hbar\left(\omega_c + \frac{g^2}{\Delta}\hat{\sigma}_z\right)(\hat{a}^\dagger\hat{a} + \frac{1}{2}). \tag{3.30}$$

The second term in the Hamiltonian looks like the energy of a cavity mode with a resonance frequency that depends on, σ_z. Depending on the qubit state, the cavity acquires a frequency shift of $\pm\chi$, where $\chi \equiv (\omega_{|0\rangle} - \omega_{|1\rangle})/2 = g^2/\Delta$. Therefore, by measuring some property of the cavity that depends on its frequency, we should be able to infer the state of the qubit.

Note that thus far we have the Jaynes-Cummings Hamiltonian of an ideal two-level qubit. However, as we discussed in section 3.3, the transmon qubit has limited anharmonicity, and therefore higher energy levels need to be taken into account. The Hamiltonian describing a multi-level qubit coupled to a single cavity is referred to as the generalized Jaynes-Cummings Hamiltonian. As we explain in Appendix A, it is straightforward to diagonalize this Hamiltonian numerically. In the dispersive limit, the energy eigenstates can be labeled $E_{m,n}$, where m represents the transmon level and $n = \hat{a}^\dagger\hat{a}$ describes the number of photons in the cavity. Note that when we excite the cavity with a coherent tone, the internal fields will be a superposition of photon number states, characterized by an average photon number $\bar{n} = \langle\hat{a}^\dagger\hat{a}\rangle$ and a phase representing the angle of the coherent state in the X_1–X_2 plane.

Each transmon level will contribute to the total shift of cavity frequency, which for small n can be defined as $\chi_{\text{tot}} \equiv (E_{0,n+1} - E_{0,n}) - \omega_r \approx g_{01}^2/\Delta$, where g_{01} is the coupling strength between the cavity and the 0–1 qubit state transition. To perform a dispersive measurement,

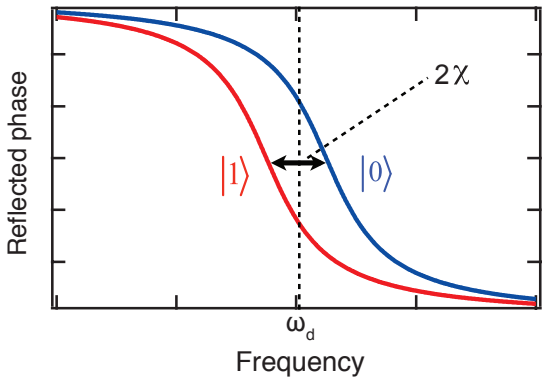

Figure 3.6: Qubit state-dependent phase shift. Cavity phase response as a function of frequency. In the dispersive limit of CQED, the cavity frequency depends on the qubit state. Therefore, the phase response sill be shifted depending on whether the qubit is prepared in the ground (blue) or excited (red) state. If we choose to measure at a frequency $\omega_d = (\omega_{|0\rangle} + \omega_{|1\rangle})/2$, then a measurement of the phase difference in the internal cavity field for the two qubit states is given by $\Delta\theta = 4\chi/\kappa$.

we want to use a measurement of the cavity to determine whether the transmon is in the ground or the excited state. Therefore, we want to consider only the cavity frequency shift corresponding to whether the qubit is in the ground or excited state, which is defined as

$$2\chi \equiv \omega_{|0\rangle} - \omega_{|1\rangle} = \frac{1}{\hbar}[(E_{1,n+1} - E_{1,n}) - (E_{0,n+1} - E_{0,n})]. \tag{3.31}$$

Note that for the rest of this thesis, we will assume that we are working with a multi-level transmon, and χ will refer to the expression defined in equation (3.31). Reference [66] gives an approximate expression for χ in terms of transmon and CQED parameters. It can be determined more accurately using the exact numerical solution discussed in appendix A. However, the most precise experimental method of determining χ is an AC Stark shift measurement, which we will describe in section 5.4.6.

We have established that we can perform an indirect QND measurement of the qubit by measuring frequency of the cavity, but how do we measure the frequency of the cavity? In general, we can send a signal of frequency ω to a cavity and measure reflected or transmitted signal. As depicted in figure 3.6, as we sweep the signal frequency of over the cavity resonance frequency ω_0 we observe a shift in the phase θ of the reflected or transmitted signal. For a transmission measurement the phase response is given by

$$\theta = \arctan\left[\frac{2}{\kappa}(\omega_0 - \omega)\right]. \tag{3.32}$$

This expression also applies to the phase of the internal cavity field, and results in a phase shift of 180° as the signal frequency crosses the resonance. For a reflection measurement,

there is a phase shift of 360°. As shown in firgure 3.6, the qubit state-dependent frequency shift will lead to a different phase response depending on whether the qubit is prepared in the ground or the excited state. If we choose to measure at a frequency $\omega_d = (\omega_{|0\rangle} + \omega_{|1\rangle})/2$, then a a measurement of the phase difference in the internal cavity field for the two qubit states is given by $\Delta\theta = 4\chi/\kappa$.

3.5 Signal to noise ratio for dispersive measurments

In the previous section, we found that we can infer the state of a qubit in a CQED architecture by measuring the phase of a signal that reflects off of the cavity. However, in reality, as we discussed in section 3.1, the instantaneous phase of any signal is never perfectly well-defined due to the quantum fluctuations of the electromagnetic field. As illustrated in figure 3.7, these fluctuations will limit our ability to resolve the qubit state-dependent phase shift.

Consider a measurement signal initially aligned in the X_1 quadrature. In the small phase shift limit ($\chi/\kappa \ll 1$), the information about the phase shift will be contained entirely in the X_2 quadrature. If the X_2 quadrature of the reflected measurement signal is averaged for a time T and recorded as a measurement voltage V_m, the probability of finding a particular value of V_m is given by a Gaussian distrubition, as depicted in the bottom panel of figure 3.7. These are the very same measurement histograms discussed in section 2.2.1. For a given phase shift, the corresponding shift in X_2 will scale linearly with the amplitude of the measurement tone. Then, since the fluctuations of a coherent tone are independent of its amplitude, the separation between the histograms for the qubit prepared in the gound and excited state will increase as $\sqrt{\bar{n}}$. Their width will decrease as we average for longer time (i.e. as T increases).

Recall from section 2.2.1 that our ability to resolve the ground an excited states, described by the histogram separation fidelity F_s, is determined by the signal to noise ratio (SNR), or strength, of a measurment. As we discussed in section 2.2.3, after the measurement signal leaves the cavity, it can experience a reduction in SNR due to a non-ideal quantum efficiency ($\eta_m < 1$). Thus, our ability to resolve the ground and excited states, will depend on the phase shift $\Delta\theta = 4\chi/\kappa$, the average photon number of the measurement tone $\bar{n}$, the measurement time T, and the measurement efficiency η_m.

Quantitatively, we can express the SNR of a CQED qubit measurement in terms of the standard deviation of the measurement histograms $\sigma = \sqrt{G/\kappa\eta T}$ and the separation between the histograms, which in the limit of small χ/κ is given by $\Delta V_{\text{opt}} = \sqrt{G}\Delta X_2 = 8\chi\sqrt{G\bar{n}}/\kappa$, where the subscript in ΔV_{opt} specifies that we are considering the histogram separation along the optimal quadrature, in this case X_2, and $\sqrt{G}$ is the voltage gain of the amplification chain. This gives us

$$S \equiv \left(\frac{\Delta V_{\text{opt}}}{\sigma}\right)^2 = \frac{64\chi^2\bar{n}\eta_m T}{\kappa}. \tag{3.33}$$

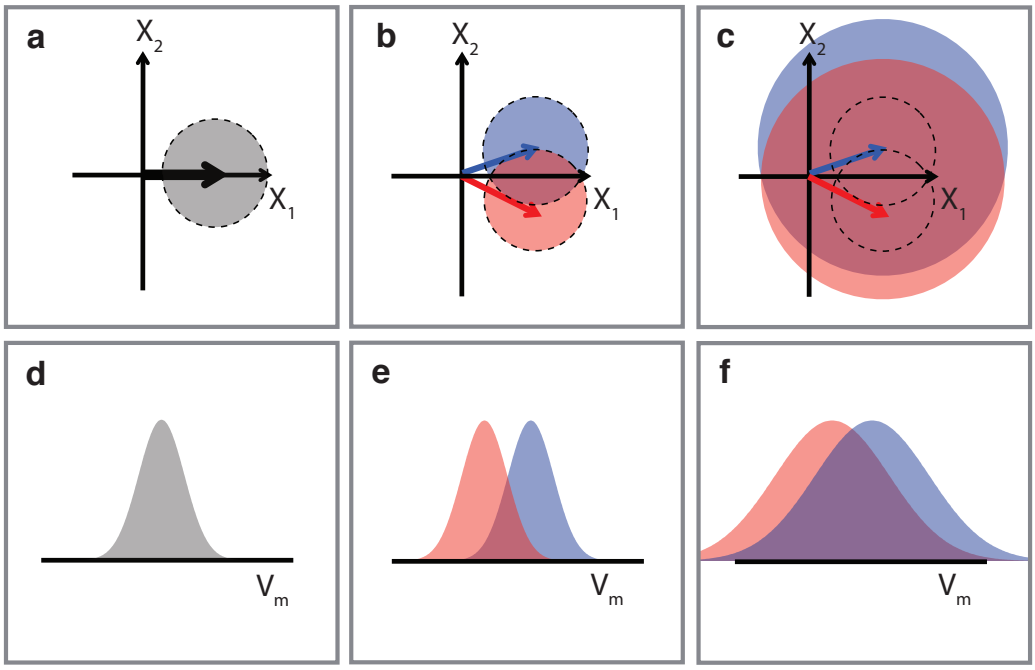

Figure 3.7: Dispersive CQED measurement signal. Panels (a) and (d) represent the measurement tone before it reaches the cavity. Panels (b) and (e) show the signal after reflecting off of the cavity and acquiring a qubit state-dependent phase shift. After the passing through an inefficient amplification chain, the SNR is reduced, as depicted in panels (c) and (f).The top panel depicts the measurement signal in quadrature space, with blue and red circles representing the width of the Gaussian-distributed quantum fluctuations of the measurement tone when the qubit is prepared in the ground and excited states, respectively. The bottom panel depicts histograms that result when the X_2 quadrature of the reflected signal is integrated for a time T and recorded as a measurement voltage V_m.

In chapter 6, we use this expression for S along with the Bayesian state update equations discussed in section 2.3.4 to track the quantum trajectories of a CQED measurement.

Chapter 4

Parametric amplifiers and squeezing

This chapter introduces the basic operating principles of Josephson parametric amplifiers, or paramps for short. Paramps have emerged at the forefront of measurement science for superconducting qubits, and thus there is a large body of research regarding their theory, design, and operation. These topics are covered in great depth in a number of theses [86, 66, 87, 88, 89] and extended papers [35, 90]. Rather than replicating these works, the aim of this chapter is to present a brief and simple picture of the basic paramp operating principles required to understand the results of this thesis.

4.1 Amplification and the quantum limit

4.1.1 Phase preserving and phase sensitive amplification

The purpose of an amplifier is to increase the amplitude of a signal by a factor of $\sqrt{G}$, where G is the power gain of the amplifier. As depicted in figure 4.2, there two main types of amplification: phase-preserving and phase-sensitive. For now, we will consider ideal (quantum-limited) amplification in the limit of large gain.

Suppose that we would like to amplify an input signal defined by an amplitude $|\alpha|$ and a phase ϕ, for example, the signal depicted in quadrature-space in figure 4.2a. The noise of this input signal is characterized by amplitudes ΔX_1 and ΔX_2. For a coherent state, $\Delta X_1 = \Delta X_2 = 1$. As shown in figure 4.2b, a phase-preserving amplifier will increase the amplitude of the signal by a factor of $\sqrt{G}$, without affecting the phase ϕ. It will also increase the amplitude of the noise, yielding

$$(\Delta X_{1,2})_{out} = \sqrt{G}\sqrt{\Delta X_{1,2} + 1}. \tag{4.1}$$

Note that extra noise, of equal amplitude to the vacuum fluctuations (or equivalently, half a photon of noise power) is added to the input noise before amplification by a factor of $\sqrt{G}$. This added noise is a direct result of Heisenberg's uncertainty principle. To noiselessly

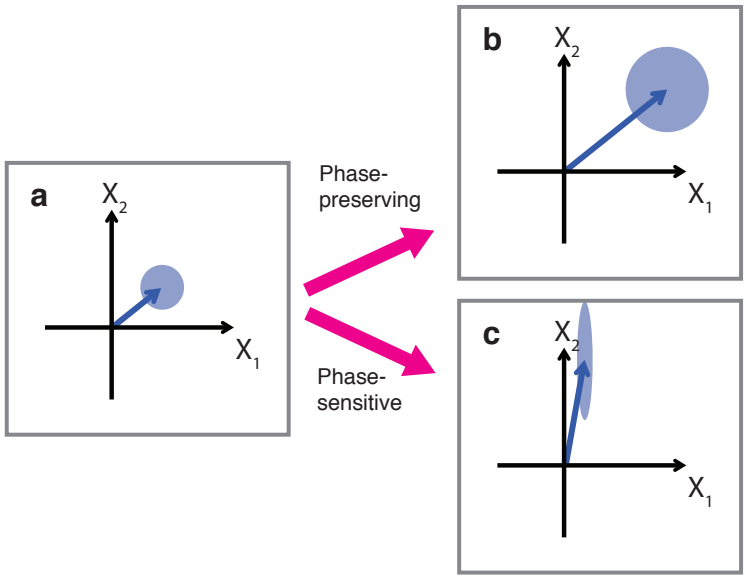

Figure 4.1: Phase-preserving and phase-sensitive amplification. a) An input signal pictured in the $X_1 - X_2$ plane. b) Under phase-preserving amplification, the amplitude and the the fluctuations of the signal are increased evenly in both quadratures. c) Phase-sensitive amplification along the X_2 quadrature. The signal amplitude and fluctuations are increased along X_2 quadrature and decreased along the X_1 quadrature.

amplify both quadratures of a signal would require simultaneously knowing the value of both quadrature amplitudes, which would violate the uncertainty principle. Therefore, an extra half photon of noise power must be added, which sets the quantum limit on amplifier noise performance. The quantum limits of amplification were discussed in detail by Caves [30] in the 1980s and are derived in a straightforward manner in section 3.2 of Slichter's thesis [66].

In phase-sensitive amplification, only one quadrature of the signal is amplified, and the other is de-amplified. An example of phase-sensitive amplification along the X_2 quadrature is shown in figure 4.2c. The X_2 component of the signal is increased by a multiplicative factor of $2\sqrt{G}$, where G is the power gain of a phase-preserving amplifier, and the X_1 component is reduced by a the same factor[88]. The amplitudes of the output noise are given by

$$(\Delta X_1)_{out} = \frac{1}{2\sqrt{G}}\Delta X_1 \tag{4.2}$$

$$(\Delta X_2)_{out} = 2\sqrt{G}\Delta X_2. \tag{4.3}$$

Note here that no extra noise is added prior to amplification. By de-amplifying one quadrature of a signal, it is possible to noiselessly amplify the other quadrature. Also note that

$$(\Delta X_1)_{out}(\Delta X_2)_{out} = \Delta X_1 \Delta X_2, \tag{4.4}$$

which means if the input is a minimum uncertainty state (such as a coherent state), then after ideal phase-sensitive amplification the output will also be a minimum uncertainty state.

4.1.2 Squeezing

Figure 4.2 depicts the case where the input state is just the vacuum fluctuations ($\bar{n} = 0$). After phase-sensitive amplification of the X_2 quadrature, the fluctuations of the output state in the X_1 quadrature will be reduced below the vacuum level. Such a state is referred to as a *squeezed vacuum state*, and X_1 is referred to as the squeezed quadrature. Note that in order for ΔX_1 to be reduced below the vacuum level ΔX_2 must be increased, as required by Heisenberg's uncertainty principle.

It is convenient to parametrize electromagnetic noise in terms of the autocorrelations of the creation and annihilation operators, by defining N and M such that [91]

$$\langle a^\dagger(t+\tau)a(t)\rangle = N\delta(\tau) \tag{4.5}$$

$$\langle a(t+\tau)a(t)\rangle = M\delta(\tau), \tag{4.6}$$

where $\delta(\tau)$ is the Dirac delta function. Qualitatively, N describes the amplitude of the fluctuations, and M describes their asymmetry in quadrature-space. For any state which is wider in the X_2 direction than the X_1 direction, we find that

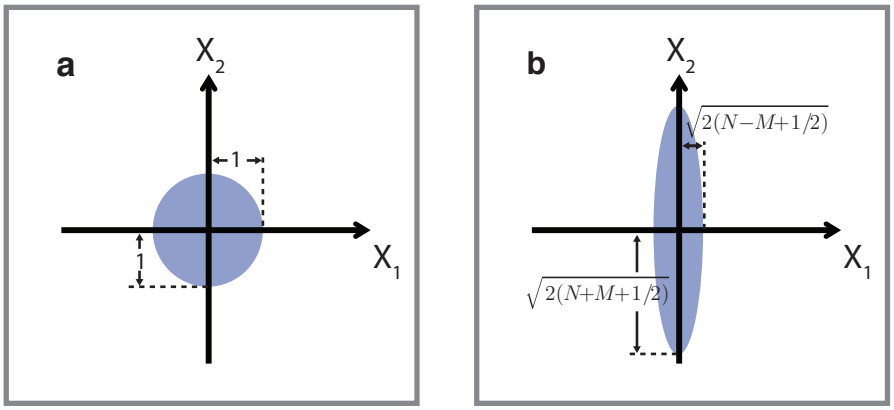

Figure 4.2: Squeezed vacuum fluctuations. a) Input vacuum state, with $\Delta X_1 = \Delta X_2 = 1$. b) After ideal phase-sensitive amplification along the X_2 quadrature, the output signal is a squeezed vacuum state, with $(\Delta X_1)_{out} = (1/2\sqrt{G})\Delta X_1 = \sqrt{2(N - M + 1/2)}$ and $(\Delta X_2)_{out} = 2\sqrt{G}\Delta X_2 = \sqrt{2(N + M + 1/2)}$.

$$\Delta X_1 = \sqrt{2(N - M + 1/2)} \tag{4.7}$$

$$\Delta X_2 = \sqrt{2(N + M + 1/2)}. \tag{4.8}$$

As shown in figure 4.3, for any state of a single mode of the EM field, the fluctuations can be parametrized by M and N. For the vacuum state, $N = M = 0$. Classical states, where the fluctuations in both quadratures are above the vacuum level, occur when $N > M$. Squeezed states only occur when

$$N < M \leq \sqrt{N(N + 1)}. \tag{4.9}$$

Here, the inequality is a manifestation of Heisenberg's uncertainty principle. When $M = \sqrt{N(N + 1)}$ (as in the case of the output of a quantum-limited phase-sensitive amplifier), the state is said to be an ideal squeezed state. In reality, the process of amplification is never completely ideal. However, if a non-ideal phase-sensitive amplifier adds a small amount of noise, it can still produce a (non-ideal) squeezed state, as pictured in the white region (green ellipses) of figure 4.3.

4.1.3 Added noise

Most amplifiers are far from reaching the quantum limit of noise performance. The best commercially available amplifiers, which are based on high electron-mobility transistors

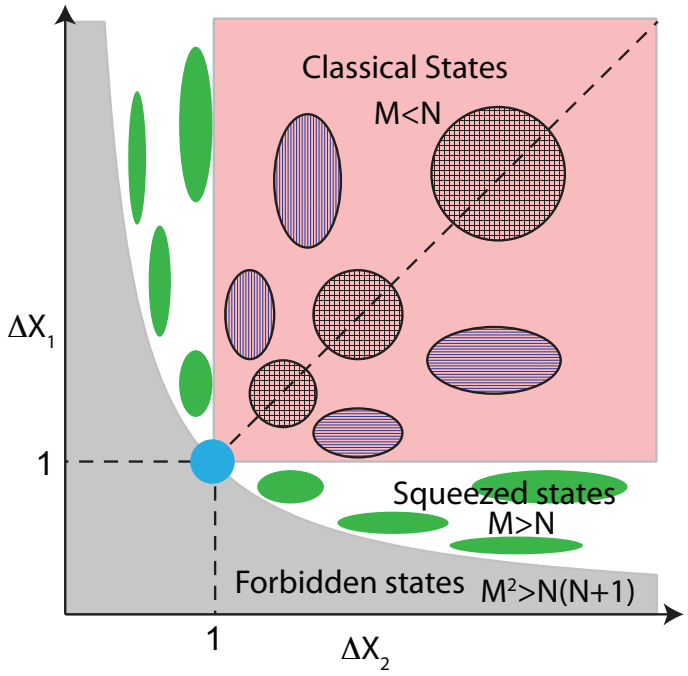

Figure 4.3: Schematic representation of electromagnetic fluctuations. For the vacuum state, $\Delta X_1 = \Delta X_2 = 1$. The red region represents states with classical noise, for which the fluctuations in both quadratures are greater than or equal to the vacuum fluctuations. The white region (and green ellipses) represent squeezed vacuum states, for which the fluctuations in one quadrature are reduced below the vacuum level. The grey region depicts states which are forbidden by Heisenberg's uncertainty principle. Adapted from reference [92].

(HEMTs) and can operate at 4 Kelvin, will add many tens of photons of noise to a signal. This added noise is typically characterized as an effective noise temperature T_{n}. On the other hand, the nearly-quantum-limited amplifiers that we will discuss in the following section are typically limited to $\sim$20$-$30 dB of gain, which is not enough for a small qubit measurement signal to overcome the noise of the room temperature electronics used to record the measurement result. Therefore, achieving a qubit measurement with high quantum efficiency will require multiple stages of amplification. For a clear and thorough description of how to compute the noise temperature of an amplification chain, see section 3.1 of Slichter's thesis [66].

4.2 Parametric amplifiers

The basic operating principle of a parametric amplifier (paramp) is that by periodically modulating some parameter of a system it is possible to transfer energy from the modulation tone, or pump, to a signal at a different frequency. In quantum optics, a crystal with a nonlinear index of refraction is used as the amplifier system, such that a coherent pump tone will periodically modulate the index of refraction. An analogous device can be realized in an electrical circuit by modulating the impedence. In a general parametric process, a pump tone at a frequency ω_p couples two modes at frequencies ω_{sig} and ω_{id} referred to as the signal and the idler, and transfers energy to both of them. Depending on the details of the amplifier, there are two possible mixing processes with different frequency relationships between the various modes. In 'three-wave mixing', the relationship is given by

$$\omega_p = \omega_{\mathrm{sig}} + \omega_{\mathrm{id}}, \tag{4.10}$$

while in 'four-wave mixing'

$$\omega_p + \omega_p = \omega_{\mathrm{sig}} + \omega_{\mathrm{id}}. \tag{4.11}$$

In either case, if the signal and idler modes are degenerate ($\omega_{\mathrm{sig}} = \omega_{\mathrm{id}}$), then ω_{sig} will be an integer multiple of ω_p. Thus, if we apply an input signal at ω_{sig}, there will exist a well-defined relationship between the phase of the pump and the signal. Therefore, only the component of the signal which is in-phase with the pump will be amplified. In other words, a degenerate parametric amplifier will provide phase-sensitive amplification for signals at ω_{sig}.

4.2.1 Josephson parametric amplifiers

To realize a parametric amplifier in a superconducting circuit, we need to modulate its impedance by varying a circuit element in time. The most convenient choice of circuit element to modulate is the Josephson inductance L_J, which varies approximately as $I(t)^2$,

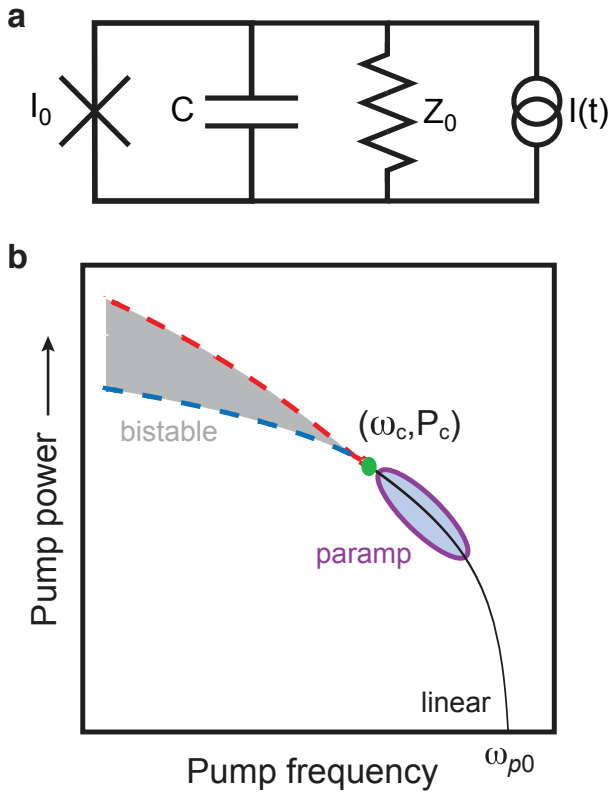

Figure 4.4: a) Circuit diagram for a Josephson parametric amplifier. b) Resonance frequency as a function of pump frequency and pump power. Adapted from reference [66].

the square of the current passing though the junction. Since $I(t)^2$ is always positive, if the junction is driven with a tone at a frequency ω_d, its inductance will be modulated at a frequency $2\omega_d$.

A Josephson parametric amplifier is realized by embedding the junction in a resonant circuit as pictured in figure 4.4a. This circuit is a nonlinear resonator just like the transmon, but with a much smaller anharmonicity, such that it effectively behaves as a classical nonlinear oscillator. There are two complementary approaches to the theory of the Josephson paramp which are treated in the literature.

The first method, detailed in Vijay's, Castaellanos-Beltran's, and Slichter's theses [86, 89, 66], applies a circuit model for the paramp, using Kirchoff's laws to write down a differential equation for $\delta(t)$. Approximating the $\sin\delta(t)$ term as $\delta(t) - \delta(t)^3/6$, we retrieve the equation of motion of a classical duffing oscillator, which can be solved for $\delta(t)$ in the presence of a periodic drive at a frequency ω_d.

The alternate approach, considered by Castaellanos-Beltran [89] and worked out in depth by Eichler [90, 87], is to write down the Hamiltonian for the paramp and to use input-output

formalism [31] to solve for the dynamics of the quantum mechanical input, output, and internal resonator fields. A starting point for this approach is the approximate Hamiltonian for the paramp circuit

$$H_{\text{JPA}} = \hbar\omega_0 a^\dagger a + \hbar\frac{K}{2}(a^\dagger)^2 a^2, \tag{4.12}$$

where K is the effective Kerr nonlinearity, which can be expressed in terms of circuit parameters.

The key qualitative feature of the parmp circuit is that, unlike a linear resonator, the resonance frequency will depend on the drive power, up until a bifurcation point beyond which there exists two stable solutions, as depicted in figure 4.4b. The parametric amplification regime occurs at powers just below the critical power (the power of the bifurcation point). Here, when the input signal is in phase with the pump, the phase of the output signal will depend strongly on the amplitude of the input signal, providing a mechanism for amplification. Similarly, when the input signal is in quadrature (90° out of phase) with the pump, the phase of the output signal will be insensitive to the amplitude of the input signal, providing a mechanism for squeezing.

4.2.2 Strongly driven Kerr Hamiltonian

Here, we present a brief outline of a model for the paramp developed by our theory collaborators in the group of Alexandre Blais [93]. This model can be used for numerical simulations of the paramp dynamics, which are useful for predicting the squeezing performance. A similar model was previously applied to a more complex circuit in reference [94].

We start from the undriven Kerr Hamiltonian, equation (4.12), and an additional drive Hamiltonian

$$H_d = \epsilon_d \mathrm{e}^{-i\omega_d t} a^\dagger + \epsilon_d^* \mathrm{e}^{i\omega_d t} a, \tag{4.13}$$

where ω_d is the drive frequency, and ϵ_d is the drive amplitude. The total Hamiltonian for the driven Kerr system can be re-expressed as

$$H_{sys} = \Delta_d a^\dagger a + \frac{K}{2}(a^\dagger)^2 a^2 + \epsilon a^\dagger + \epsilon^* a, \tag{4.14}$$

where $\Delta_d \equiv \omega_0 - \omega_d$. Assuming ϵ is large compared to the amplitude of the vacuum fluctuations, we can transform to a displaced frame $a \to \alpha + d$ where α is the strong classical field and d represents the quantum fluctuations of the field. Then, choosing the phase of the rotating frame such that α is real, we can rewrite the Hamiltonian as

$$H_{sys}^{(4)} = \tilde{\Delta} d^\dagger d + \frac{\lambda}{2}(d^{\dagger 2} + d^2) + \mu(d^{\dagger 2} d + d^\dagger d^2) + \frac{K}{2} d^{\dagger 2} d^2, \tag{4.15}$$

where we have defined $\tilde{\Delta} \equiv \Delta_d + 2|\alpha|^2 K$, $\lambda \equiv |\alpha|^2 K$, and $\mu = \alpha K$. The superscript $\cdot^{(4)}$ indicates that this Hamiltonian is fourth-order in $d, d^\dagger$ [1]. By ignoring the higher order terms, we recover the quadratic Hamiltonian

$$H_{sys}^{(2)} = \tilde{\Delta} d^\dagger d + \frac{\lambda}{2}(d^{\dagger 2} + d^2). \tag{4.16}$$

The Lindblad-type master equation which describes the dynamics of the internal cavity field ρ is given by [94]

$$\dot{\rho} = -i\,[H_{sys}, \rho] + \kappa \mathcal{D}[d]\rho, \tag{4.17}$$

where here κ is the damping rate of the paramp (its linewidth in the linear regime), and $\mathcal{D}[d]\rho \equiv (2d\rho d^\dagger - d^\dagger d\rho - \rho d^\dagger d)/2$. The steady-state solution to equation (4.17) can be found numerically for both the quadratic and the quartic Hamiltonians. The output field can be solved by input-output formalism, through the relation $d_{\text{out}} = \sqrt{\kappa} d - d_{\text{in}}$ [95]. Assuming that the input field is just the vacuum fluctuations, then the quadratic Hamiltonian will produce ideal squeezed vacuum in the output mode d_{out} [95, 91]. In chapter 8, we will discuss how higher order terms in $H_{sys}^{(4)}$ can contribute to non-ideal squeezing.

4.3 Paramp performance

Here we will briefly outline the basic considerations of paramp performance: gain, bandwidth, and dynamic range. For a more in-depth treatment, the fundamental limitations and design trade-offs for paramps are eloquently discussed in section 1.3.3 of Flavius Schackert's thesis [88].

Because the paramp is a resonant structure, it will amplify signals within a finite frequency bandwidth, B, about its resonance frequency. In the linear regime, where no amplification occurs, the bandwidth is determined by the quality factor of the resonator $Q = \omega_0/\kappa$, which is set by the choice of circuit parameters. As the drive power is increased toward the critical power, the gain increases and the resonance steepens, decreasing the amplification bandwidth. As derived in reference [66], the bandwidth will decrease as the amplitude gain:

$$B\sqrt{G} \propto \frac{1}{Q}. \tag{4.18}$$

Thus, by decreasing Q, it is possible to increase both the gain and the bandwidth independently. However, as Q is decreased, larger pump powers will be required to reach a given

[1] The Hamiltonian could be extended to higher order by adding an additional term to the un-driven Hamiltonian, $(K'/3)a^{\dagger 3}a^3$, where K' is the cubic Kerr constant, as demonstrated in reference [94].

gain point. If the required pump power is too large, it will excite higher-order junction nonlinearities, preventing stable paramp operation. This constraint leads to the approximate requirement that

$$Qp \gtrsim 5, \tag{4.19}$$

where $p = L_J/L_{\text{tot}}$ is the participation ratio of the Josephson inductance to the total inductance.

It is also important to consider what happens when the signal power becomes comparable to the pump power. In this case, the amplification process will deplete the pump, causing a reduction in the gain. The signal power at which the gain is reduced by 1 dB is know as the '1 dB compression point' and determines the dynamic range of the amplifier. For typical design parameters, the dynamic range of Josephson parametric amplifiers is limited to around -120 dBm, although this number can be improved upon by using arrays of Josephson junctions [90].

4.4 Paramps in modern qubit experiments

Although the basic operating principles and design considerations for Josephson paramps are well understood, paramp development remains an active field of research. This is due to the fact that pramps are now widely used in qubit measurements setups, and any further device improvements will allow for better measurements. In particular, many state-of-the-art qubit experiments[15, 19] rely on a combination of dispersive measurements and paramps to achieve fast high-fidelity QND measurements, which are required for fault-tolerant quantum computing schemes such as the surface code. Three areas of active paramp development include: (i) increased dynamic range, (ii) increased bandwidth, and (iii) integration.

Ultimately, the measurement fidelity is limited by a trade-off between separation fidelity and qubit energy relaxation during measurement. As the measurement time is increased, the separation fidelity increases, but the errors due to T_1 will increase. As we discussed in section 3.5, increasing the power of a dispersive measurement tone will increase the measurement rate, allowing for a higher fidelity in a given integration time. However, there are two limitations on how much measurement power can be used. The first concern is that if the internal cavity field goes beyond the critical photon number $n_{\text{crit}} \approx \Delta^2/4g^2$ at which the dispersive approximation breaks down, then the measurement will no longer be fully QND. In this case, it will drive transitions between qubit states, and thus reduce the measurement fidelity. The second is that too large a measurement tone can saturate the dynamic range of the paramp, reducing the gain and thus the measurement fidelity. In fact, dynamic range is currently the limiting factor for state-of-the-art qubit measurements[16].

Looking forward towards multi-qubit architectures, it is desirable to measure multiple qubits on a single measurement line, with different frequency measurement tones for each qubit, using the a single paramp. Thus, large bandwidth is desirable, in order to amplify

signals over a wide frequency range. Without sacrificing gain, paramp bandwidth can be improved by using an impedance transformer and other more aggressive methods of impedance engineering [96]. Another promising avenue is to use traveling wave parametric amplifiers (TWPAs), which are currently under development here at Berkeley and will allow for substantially better dynamic range and bandwidth than resonant paramps [97, 98, 99].

Finally, there are also ongoing efforts to integrate paramps on-chip with qubits [100, 101]. An integrated device would eliminate the need for microwave components such as cables and circulators between the qubit and the paramp, which currently contribute loss and limit the measurement quantum efficiency.

Chapter 5

Experimental setup

5.1 Qubit designs

5.1.1 3D transmon qubit

In the experiments presented in this thesis, we use transmon qubits that are designed to couple to the the electric field of a three-dimensional microwave cavity [77]. The cavity is a rectangular waveguide cavity made from copper (or aluminum), and we choose to work with the fundamental TE_{101} cavity mode. As depicted in figure 5.1, the qubit consists of a Josephson junction which shunts together two aluminum paddles. The transmon capacitance is formed by the self-capacitance of these paddles, which couple through the sapphire (or high-resistivity silicon) substrate and the vacuum. The paddles also form a dipole antenna, which couples the qubit to the electromagnetic vacuum. In free space, this radiative coupling would severely limit the qubit's T_1 time. However, when the qubit is placed in a cavity it is isolated from the electromagnetic continuum. Instead, it will couple to the cavity mode, and it's radiative decay rate will be set by the Purcell effect, which in the dispersive limit is given by

$$T_{1,\text{Purcell}} = \frac{\Delta^2}{\kappa g^2}. \tag{5.1}$$

Note that this is the expression for the Purcell effect due to a single cavity mode, and is not entirely accurate in the case of a multi-mode cavity such as the waveguide cavity used in our experiments. One way to predict the actual Purcell T_1 time is to use a 3D EM simulator to solve for the multi-mode Lorentzian cavity transmission profile. It turns out that the multi-mode Purcell T_1 time is increased beyond the single-mode result when the qubit is below the fundamental cavity mode and decreased when the qubit is above the fundamental mode.

Prior to the invention of these '3D transmons', readout cavities for transmons were typically fabricated on-chip with the qubit, either as a lumped element LC resonaor or as a

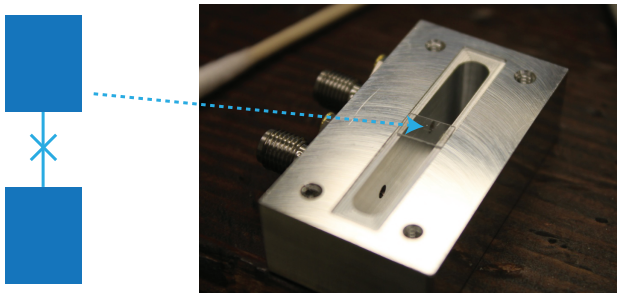

Figure 5.1: 3D transmon qubit. A diagrammatic representation for the 3D transmon qubit (left), which consists of a Josephson junction which shunts together two aluminum paddles. The paddles capacitively couple through vacuum (and the substrate) to form the transmon capacitance, and they also act as a dipole antenna, which couples to the electromagnetic field of a microwave frequency waveguide cavity made from aluminum (left).

distributed element coplanar waveguide resonator. In such 'planar transmons', the transmon capacitance is typically formed from interdigitated capacitors. The electric fields of interdigitated capacitors are concentrated near the surface of the substrate and metal structures, and T_1 times are typically limited by surface dielectric loss. In 3D transmons, the electric fields are predominantly contained in the bulk substrate and the cavity vacuum, and therefore the surface participation is substantially lower.

Using the same materials and fabrication procedures described in Slichter's thesis [66], we are able to improve transmon T_1 times from $1\,\mu$s to $30\,\mu$s by moving from a planar to a 3D design. Energy relaxation times for state-of-the-art 3D transmons, for example, those made at Yale and IBM, exceed $100\,\mu$s. However, ultimately it is likely that it will be easier to fabricate complex multi-qubit architectures using planar qubits and cavities than with a 3D architecture. T_1 times for planar transmons have also improved in recent years, thanks to design improvements to reduce the effect of surface dielectric loss[102] and materials improvements which have decreased the amount of surface dielectric loss [103, 78].

5.1.2 Choice of sample parameters for weak measurements

As discussed in section 2.3.2, in order to observe diffusive quantum trajectories we need to work in the weak measurement regime, where the measurement timescale is much slower than the timestep over which we record our measurement results, $\Delta t \ll \tau$. As discussed in section 3.5, in the dispersive limit of circuit QED the qubit-state-dependent phase shift of the internal cavity field is given by $\Delta\theta = 4\chi/\kappa$. While it is possible in principle to perform weak measurements for any choice of χ/κ, it is most straightforward to do so in the limit of $\chi/\kappa \ll 1$, for three reasons. First of all, in the small χ/κ limit the separation between measurement histograms for the ground and excited states, ΔV, is directly proportional

to $\Delta\theta$, which simplifies our analysis and allows us to apply equation (3.33). The second reason is that choosing a small χ/κ simplifies our calibration of χ, as discussed in section 5.4.6. Finally, a smaller χ/κ results in a slower measurement rate for a given choice of $\bar{n}$ and measurement time[1].

The transmon qubit used for the quantum trajectory experiments presented in this thesis has $E_c/h \sim 200$ MHz and $E_J/h \sim 11$ GHz, resulting in a qubit frequency of $\omega_q/2\pi = 3.999$ GHz. The qubit couples to the fundamental mode of a copper waveguide cavity at $\omega_r = 6.8297$ GHz, with a strength $g \approx 90$ MHz, which results in a dispersive shift of $\chi/2\pi = -0.52$ MHz. The cavity decay rate $\kappa/2\pi = 10.86$ MHz is set by adjusting the length of a metal pin which extends from a coaxial microwave connector into a small hole in the side of the cavity. For this choice of sample parameters $\chi/\kappa = 0.048$, which puts us well within the small phase shift limit where equation (3.33) is valid.

Qubit energy relaxation is another concern, as it can limit projective measurement fidelity, disrupt diffusive trajectories, and contribute to environmental dephasing. Therefore, we want to achieve as high of a T_1 time as possible. For the sample parameters detailed in the pervious paragraph, Purcell relaxation is negligible and the measured T_1 time of $30\,\mu$s is set by other environmental loss mechanisms. Since all of the measurements in this thesis take place over a timescale less than $2\,\mu$s, energy relaxation during measurement will have only a small effect on the observed trajectories. Finally, it is also important to minimize environmental dephasing. This has less to do with the choice of sample parameters, and more to do with the embedding environment. We have found, for example, that high frequency noise traveling down the microwave lines from the cryogenic switches can adversely affect T_2^* times. Such noise can be reduced by using absorptive filters at the base of the fridge. For the experiments discussed in this thesis, we measured T_2^* times ranging from $15-20\,\mu$s.

5.2 Measurement setup

5.2.1 Basic setup, paramp parameters, double pump

Our basic measurement setup is illustrated in figure 5.2. A measurement tone reflects off of a 3D cavity which embeds a transmon qubit. A directional coupler is then used to combine the measurement tone with the paramp pump. The pump and the amplfied signal both reflect off of the paramp and continue to further stages of amplification.

The paramp consists of a two-junction SQUID formed from $2\,\mu$A Josephson junctions, shunted by 3 pF of capacitance formed by aluminum pads on top of a SiN_x dielectric and a Niobium ground plane. The paramp couples differentially to a hybrid microwave launch. Details on paramp and launch designs can be found in Slichter's thesis [66].

In figure 5.3 we provide a detailed diagram of the microwave setup used to measure individual quantum trajectories in references [27] and [28] which we will refer to throughout

[1] Although, this reason is not of fundamental significance, as one could also reduce the measurement rate by reducing $\bar{n}$.

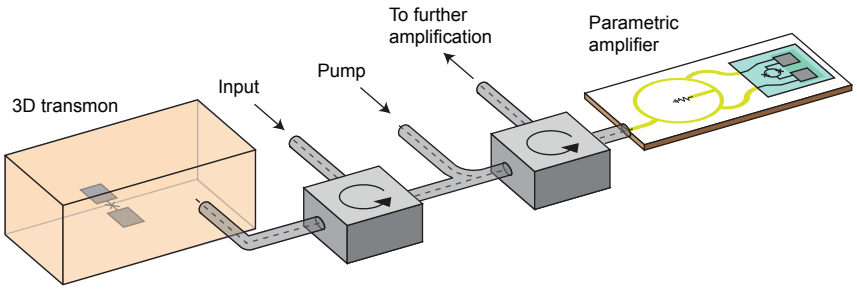

Figure 5.2: Basic measurement setup. A measurement tone reflects off of a 3D cavity which embeds a transmon qubit. It is then combined with a pump tone before reflecting of a parametric amplifier and then moving on to further stages of amplification.

this chapter.

Next, we discuss the technique that we use to bias up the paramp. To achieve phase-sensitive amplification, the signal and the effective pump must be degenerate in frequency. The most straightforward method to do so is to use a single microwave signal generator to produce both the pump tone and the measurement tone at the same frequency. This method is sometimes referred to as 'single pumping'. The primary problem with this approach is that there is only limited isolation between the paramp and the cavity, which means that a part of the strong pump tone at the cavity frequency will reach the cavity and influence the qubit. As we will discuss in chapter 8, it is desirable to use as few circulators as possible, since each circulator has a small amount of loss which will reduce the collection efficiency of the measurement.

For these reasons, it is desirable to achieve a gain profile centered around the measurement frequency without applying a pump tone at that frequency. The method that we used to achieve this goal is sometimes referred to as 'double pumping'. As depicted in figure 5.4, a pump tone at the measurement frequency ω_m is modulated at a frequency ω_{SB}, which generates two sidebands at $\omega_m + \omega_{SB}$ and $\omega_m - \omega_{SB}$. By choosing the appropriate mixer offset voltages, we can eliminate any output signal at ω_m. The overall power level is set by a variable attenuator. The two sidebands are then sent to the paramp, where they are mixed due to the paramp non-linearity, generating a virtual pump at ω_m. The theoretical details of this process are described in reference [104].

To maximally distinguish between the qubit states, we choose a measurement frequency that is halfway between the cavity frequencies for the qubit in the gound and and excited states,

$$\omega_m = \frac{1}{2}(\omega_{|0\rangle} + \omega_{|1\rangle}). \tag{5.2}$$

We used the following procedure to bias the paramp. First, using a single pump, we tune

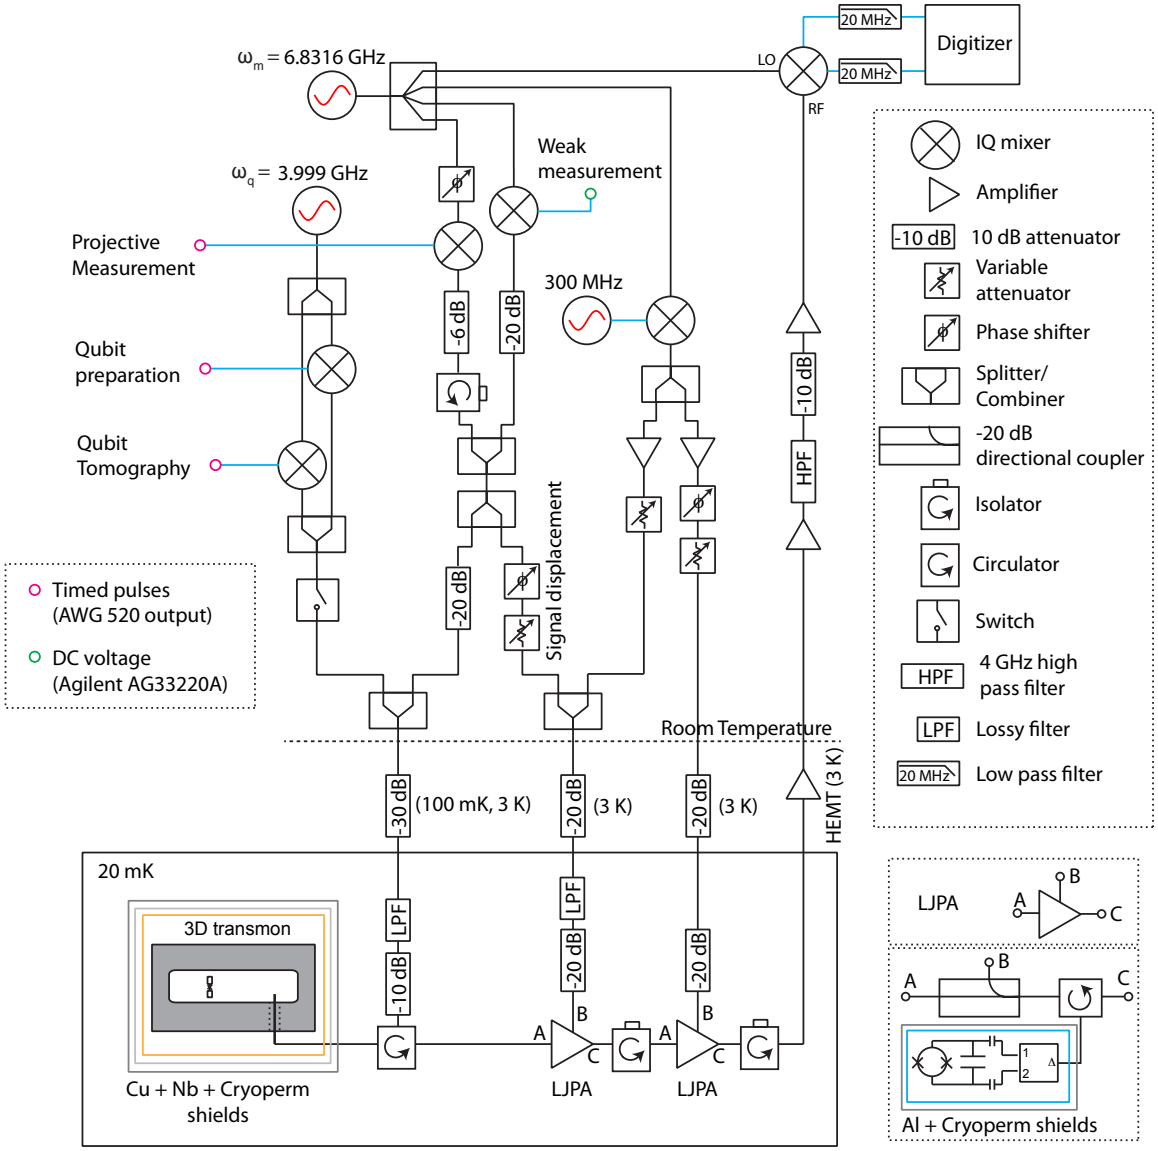

Figure 5.3: Schematic of the full measurement setup for the experiments in reference [27]. The setup for [28] is nearly identical, but only uses one paramp. For visual clarity, additional DC offsets to null mixers are not pictured.

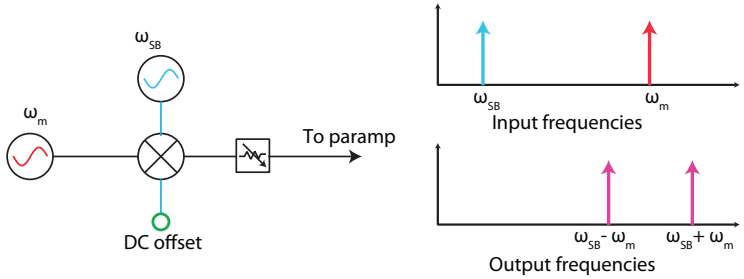

Figure 5.4: Double-pumping setup. A mixer is used to generate sidebands above and below measurement frequency, at $\omega_m + \omega_{SB}$ and $\omega_m - \omega_{SB}$, and to eliminate any signal at ω_m. The pump power is set using a variable attenuator. The two pump tones are then sent to the paramp, where they will combine to form a virtual pump at ω_p.

the pump power and the flux through the SQUID loop (and hence the paramp frequency) to achieve a clean gain profile with approximately 20 dB of gain centered about ω_m, as described in detail in section 6.2 of Slichter's thesis [66]. We save a trace of the gain profile, and then switch to a double pumping setup. The goal is to reproduce the same clean gain profile that we achieved from single pumping without applying any pump power at ω_p. In order to accomplish this, the powers of the two sidebands needs to be balanced. This is not a trivial task, as the sidebands will experience different amounts of frequency-dependent attenuation as they travel through the coaxial lines leading to the paramp. Through trial and error, we found that we were able to achieve a clean gain profile with a modulation frequency $\omega_{SB} = 300$ MHz. In other experiments, we sometimes found that it was necessary to use additional techniques to balance the sideband powers. Since the higher frequency sideband will usually experience greater attenuation, it can help to add a high-pass filter with a rolloff frequency in between the two sidebands, and to fine-tune the amount of attenuation by varing ω_{SB}. A more precise method to balance the sidebands would be to apply modulations to both the I and Q ports of an IQ mixer, so that the relative sideband amplitudes are set by the relative phase of the modulation tones.

Another technique to eliminate the pump tone at the measurement frequency is to directly modulate the the paramp frequency by applying a periodic flux tone through the SQUID loop at $2\omega_p$. This method, known as flux-pumping, closely resembles an ideal three-wave mixing process. We choose to use double-pumping instead of flux-pumping in our experiments simply as a matter of convenience, since flux-pumping requires an additional microwave control line to modulate the flux.

5.2.2 Signal displacement and choice of amplified quadrature

Consider a dispersive measurement in the limit of small χ/κ, as pictured in figure 5.5. A measurement tone initially aligned along the X_1 quadrature reflects off of the cavity and acquires a qubit-state-dependent phase shift, resulting in a shift in the X_2 quadrature of the reflected signal. The amplitude fluctuations of the measurement tone are primarily in the X_1 quadrature, and provide a record of the fluctuating intra-cavity photon number. In general, we can choose to perform phase-sensitive amplification of the measurement tone along an arbitrary axis in the X_1-X_2 plane, with corresponding de-amplification along the orthogonal axis. If we choose to amplify along X_2 (fig. 5.5b), we acquire qubit state information and squeeze the photon number fluctuations. If we choose instead to amplify along X_1 (fig. 5.5), we acquire information about the intra-cavity photon number and squeeze away the qubit state information.

However, prior to amplification, the majority of the amplitude of the measurement tone remains in the X_1 quadrature and contains no information about the qubit. We sometimes refer to this useless part of the measurement tone as the 'dumb signal'. If dumb signal is sent to the paramp, it will need to be de-amplified (fig. 5.5b) or amplified (fig. 5.5c), depending on which quadrature we choose to amplify. The dumb signal can contribute to paramp compression, especially when amplifying along X_1. In addition to affecting gain, compression also leads to non-linear amplification. It is easiest to interpret our measurement signal in the regime of linear amplification, where the signal is proportional to the amount of measurement backaction. For these reasons, it is desirable to eliminate the dumb signal before amplification.

The dumb-signal can be removed by simply applying a coherent tone with the appropriate phase and amplitude to displace the measurement signal back to the origin of the $X_1 - X_2$ plane. This procedure is depicted in figure 5.6. The measurement and pump tones are generated at a frequency ω_m using a single microwave generator. The measurement pulse is shaped using a mixer, and then split. One branch is attenuated and then sent to the cavity to perform a qubit measurement. The second branch, referred to as the signal displacement, passes through a variable phase shifter and attenuator, and then combines with the paramp pump. After the measurement tone reflects off of the cavity, it recombines with the signal displacement tone (see full schematic, section **??**). The phase shifter and attenuator are tuned to displace the measurement signal back to the origin of the $X_1 - X_2$ plane prior to amplification. Because the signal and the signal-displacement tones pass through only passive, linear (except for the qubit) components before recombining, once the dumb signal cancellation is tuned up it will work regardless of the amplitude, phase, or shape of the measurement pulse.

5.2.3 Cascading paramps

In the experiments reported in reference [27], we found that instead of using a single paramp, we could achieve a slight improvement in quantum efficiency by cascading two

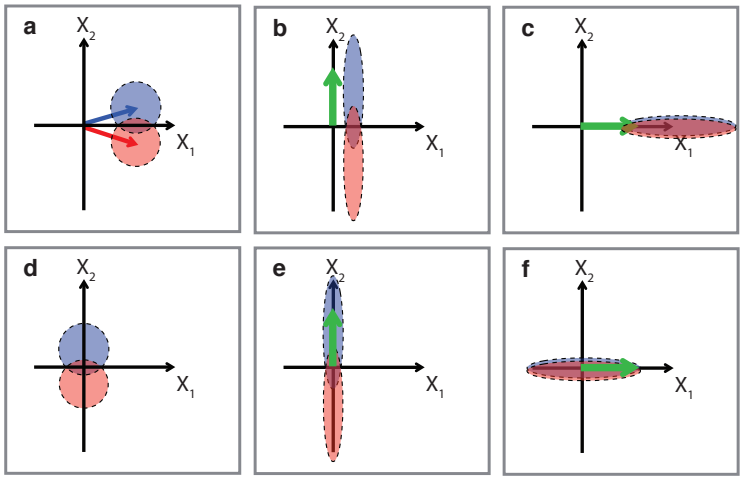

Figure 5.5: Signal displacement. a) A measurement signal after reflecting off of the cavity and acquiring a quit-state-dependent phase shift. b) The signal after phase-sensitive amplification along the X_2 quadrature. c) The signal after instead amplifying along the X_1 quadrature. d) The initial measurement signal from (a), displaced to the center of the X_1-X_2 plane prior to amplification. e) The displaced signal from (d) after amplification along the X_2 quadrature. f) The displaced signal after amplification along the X_1 quadrature.

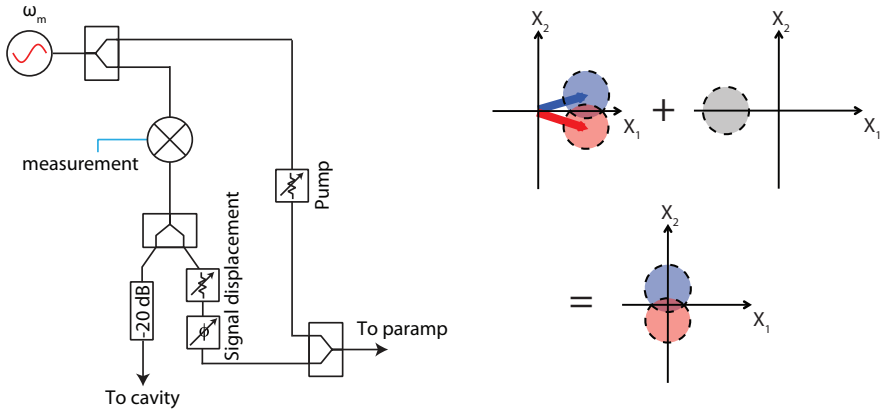

Figure 5.6: Signal displacement cancellation setup. The measurement and pump tones are generated at a frequency ω_m using a single microwave generator. The measurement pulse is shaped using a mixer, and then split. One branch is attenuated and then sent to the cavity to perform a qubit measurement. The second branch passes through a variable phase shifter and attenuator, which are tuned to displace the measurement signal back to the origin of the X_1-X_2 plane prior to amplification.

paramps back-to-back, as depicted in the full experimental setup in section **??**. The first paramp was biased for 10 dB of gain, and the second for 15 dB. With this setup we measured $\eta_m = 0.49$, as opposed to our best single paramp value of $\eta_m = 0.44$.

As this is only a small improvement, and this is not a topic that we studied systematically, it's difficult to pinpoint the exact cause. One explanation is that by using two paramps, we were able to achieve a better overall gain-bandwidth product. To achieve a high quantum efficiency, it is important that the bandwidth is larger than the cavity linewidth. This is beacuse the fluctuations of the measurement signal after reflecting off of the cavity, which contain information about the qubit state, will have a correlation timescale of $1/\kappa$. In order to collect all of the qubit state information, the paramp must be capable of reacting faster than that timescale. Another possible explanation, which we will return to in chapter 8, is that the paramps may perform more ideally when operated at lower gain.

5.3 Measurement techniques

5.3.1 Generating pulse sequences

Microwave tones for the trajectories experiments were produced by two commercial signal generators (Agilent E8257D), one set to the qubit frequency, and one set to the measurement frequency. A low-frequency generator (HP 8657B) was used to produce the 300 MHz IF signal for double-pumping. The tones at the qubit and measurement frequencies are split and then sent to the input (LO) port of commercial IQ mixers (Marki IQ0307 and IQ4509). The amplitude and phase of the signal leaving the mixer is determined by the voltages applied to the I and Q ports. By applying time-dependent voltages to the I and Q ports, we can shape the pulses leaving the mixer.

As depicted in figure 5.3, we have two separate measurement mixers, one for weak measurements, and one for projective readout. The weak measurement is a continuous tone that remains on for the entire experimental sequence. Its amplitude and phase are set by DC voltages (sourced by Agilent AG33220A generators), which can be set in software. The projective readout and all qubit manipulations are pulsed. We set the offset voltages on the I and Q ports to minimize the signal leakage when the pulses are 'off'. We are typically able to achieve an on/off ratio of 60 dB. To further improve the isolation, qubit pulses are gated by a fast switch (ZASW-2-50DR+).

The timing and pulse shaping are controlled by an arbitrary waveform generator (Tektronix AWG520), which can create sequences with 1 giga-sample per second resolution. The AWG has two main outputs and an additional four marker outputs, all of which were used in our experiments.

5.3.2 Homodyne detection

In order to analyze our measurement signal, we would like to record it's average value in quadrature-space (or equivalently, its amplitude and phase) as a function of time. This

can be accomplished by comparing the measurement signal to a reference signal, or local oscillator (LO), oscillating at the same frequency. To achieve good phase locking, the LO is split off of the same generator used to produce the measurement tone. The LO and the measurement tone are sent to the LO and RF ports of an IQ mixer where they are multiplied, and the output signals on the I and Q ports will represent the two quadrature amplitudes of the demodulated signal. These signals are then amplified using ultra-low-noise voltage amplifiers (OPA847s).

5.3.3 Data acquisition

To reduce the noise bandwidth and to exclude any oscillations due to our double-pump sidebands, each branch of the demodulated signal passes through a 20 MHz low pass filter. The signals are then sent to an ADC (Alazar ATS9870) which digitally records the measurement voltage in 4 ns segments[2], which we then down-sample in software into 16 ns time-steps.

Since we are performing phase-sensitive amplification, only one quadrature of the measurement signal will contain information about the qubit. Therefore, one could rotate the phase of the LO so that all of the measurement signal is contained in one channel. Instead, in our experiments we rotated the two digitizer channels in software to achieve the same effect. This approach worked well for us, but in retrospect it is not advisable, because unequal gain for the two channels of the demodulation circuit can lead to inaccuracies in software rotation. It is better to manually adjust the phase of the LO.

5.4 Calibration experiments

5.4.1 Spectroscopy, Rabi, Ramsey, and T_1

The first step to setting up our experiments is to characterize the cavity and the qubit in the frequency domain. Using a vector network analyzer, we measure the cavity response, which we fit to extract the resonance frequency and the linewidth κ. We take data both at low probe power ($\bar{n} \ll 1$) which gives us $\omega_{|0\rangle}$, and at high powers ($\bar{n} \gg n_{crit}$), where the qubit is saturated (or 'punched out') and no longer affects the cavity resonance, which gives us the bare cavity frequency ω_r. Then, to get a rough sense for the qubit frequency, we perform two-tone spectroscopy. This procedure involves probing the cavity phase on resonance with the VNA, while scanning the frequency of a spectroscopy tone across the ω_q. In the experiments discussed in this thesis, we operate the qubit at a fixed frequency, so the only parameters to consider are the powers of the probe and spectroscopy tones. At low enough powers, we can fit the peak to get an initial estimate for ω_q and the qubit's spectral linewidth. At higher spectroscopy power, a sharp peak appears at one half the frequency of

[2]The Alazar is capable of digitizing in 1 ns segments, but as our cavity bandwidth is roughly 16 ns and we aren't using any advanced pulse shaping, it is unnecessary for us to sample that quickly.

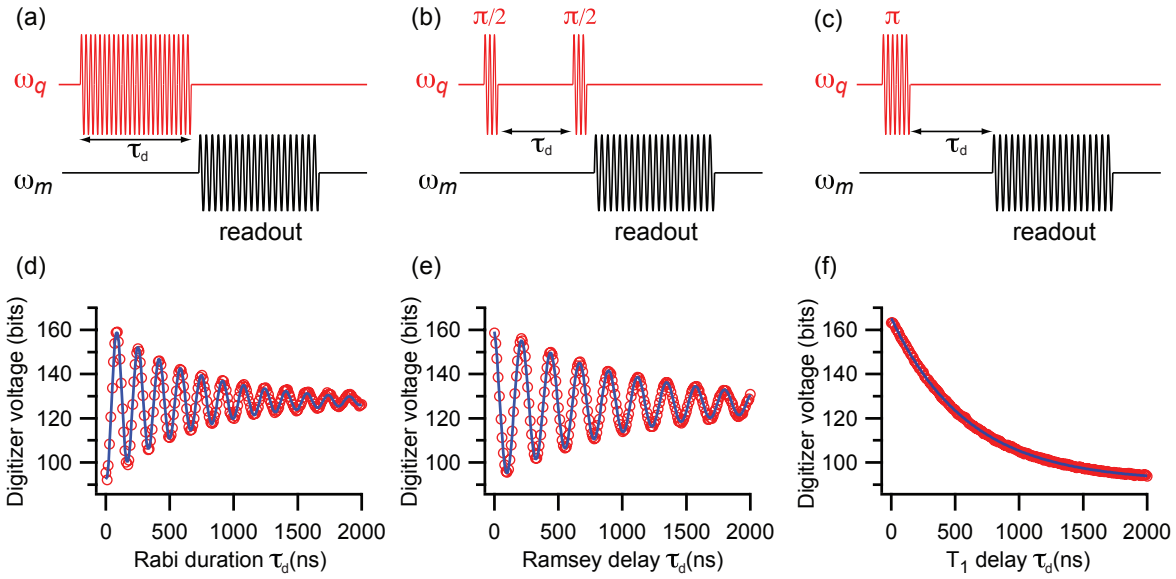

Figure 5.7: Rabi, Ramsey, and T_1 measurements. Panels (a), (b), and (c) show timing diagrams for qubit (red) and measurement (black) pulses. Panels (d), (e), and (f), show the averaged measurement digitizer voltage from a large ensemble of experimental iterations as a function of the delay time τ_d. Adapted from reference [66].

the qubit 0−2 transition, which, when combined with knowledge of ω_q, provides information about the qubit anharmonicity. For a detailed description of two-tone qubit spectroscopy, see chapter 7 of Schuster's thesis [69].

Next, we move on to time-domain measurements, which allow us to characterize qubit parameters more precisely. Based on the spectroscopic measurements, we can start with a rough estimate of the optimal measurement frequency $\omega_m = \omega_{|0\rangle} + \chi$ (recall that χ is negative in our experiments). After biasing up the parmp for gain at ω_m, we then proceed to set up a projective measurement. We apply a continuous drive at ω_q to create an incoherent superposition of qubit states, and then we send in a sequence of repeated $2\,\mu$s measurement pulses at ω_m, separated by enough time for the qubit to relax back to an incoherent superposition state between each measurement. We integrate each measurement signal for a time Δt, and histogram the results. The relative phase between the measurement tone and the paramp will set the axis of amplification in the X_1-X_2 plane. We tune this relative phase to achieve maximally separated measurement histograms for the ground and excited qubit states, corresponding to amplification along X_2 (see figure 5.5). We also tune the measurement power and Δt, until we can see well-separated measurement histograms. In section 5.4.3, we will discuss how we optimize the measurement fidelity.

For our first time-domain characterization, we measure Rabi oscillations which allow us

to roughly calibrate qubit pulse amplitudes. As pictured in figure 5.7a, we excite the qubit at ω_q for a variable time τ_d, and then we perform a projective measurement, as described above. Rabi oscillations are an ensemble measurement, which means that for each value of τ_d, we average together the measurement result from many experimental iterations, which gives us an average digitizer voltage proportional to $\langle\hat{\sigma}_z\rangle$. The drive at ω_q causes the qubit state to oscillate around the Bloch sphere, which results in an oscillating digitizer voltage as a function of τ_d, as pictured in figure 5.7d. The oscillations decay on a characteristic timescale which depends on T_1, T_2^*, and the noise spectral density at the oscillation frequency. The functional form for this Rabi decay time is given in reference [66].

For a given qubit drive power, the amount of time required to rotate the qubit state by an angle of π is given by one half of the oscillation period. A pulse of that duration is known as a π-pulse, and can be used to swap the qubit from the ground to the excited state. A pulse of half the duration of a π pulse is known as a $\pi/2$, and will transform the ground state into a maximal superposition state[3]. In our experiments, we have built experimental sequences to work with a 32 ns π-pulse (and a 16 ns $\pi/2$-pulse). We use Rabi oscillations for an initial pulse calibration, by varying the amplitude of the qubit drive until the oscillation period is roughly 64 ns. As discussed in reference [66], the Rabi oscillation frequency also depends on the detuning of the qubit drive from the actual qubit frequency. Therefore, to precisely calibrate our qubit pulses we must first precisely calibrate the qubit frequency.

A Ramsey measurement, depicted in figure 5.7b, can be used to characterize both the qubit frequency and the dephasing time T_2^*. Our measurement sequence begins with a $\pi/2$ pulse to prepare the qubit along the equator of the Bloch sphere with a phase ϕ, followed by free qubit evolution for a variable time τ_d. Then we apply a second $\pi/2$ pulse and perform projective measurement. Averaging together an ensemble of measurement results, we will see a maximal readout voltage (here corresponding to the qubit in the excited state), if at the time of the second $\pi/2$ pulse the qubit phase is ϕ. If instead the phase is $\phi + \pi$, the second $\pi/2$ pulse will return the qubit to the ground state, resulting in a minimum readout voltage. During free evolution, in a reference frame rotating at the drive frequency of the qubit pulses the qubit state will rotate around the equator of the Bloch sphere at a frequency given by the difference between the drive frequency and the actual qubit frequency. If the drive is on resonance, we will see no oscillations, and only an exponential decay in readout voltage with a decay constant set by the dephasing time T_2^*. If the drive frequency of our qubit pulses is detuned from the qubit frequency, then we will observe a digitizer voltage which oscillates at the detuning frequency, as pictured in figure 5.7b. Then, our experimental procedure for determining the qubit frequency is to intentionally detune the drive from the expected qubit frequency by 1 MHz (in order to see oscillations that can be easily fit), perform a Ramsey experiment, fit the Ramsey oscillations to determine the detuning frequency, and adjust the drive frequency accordingly.

At this point, we can also characterize the qubit's energy relaxation by performing a T_1

[3]However, due to transients, an optimal $\pi/2$ pulse will not necessarily last for exactly half the duration of an optimal π pulse. This becomes a concern when trying to use very fast pulses, or when trying to achieve very high state pulse fidelities and very precise tomography.

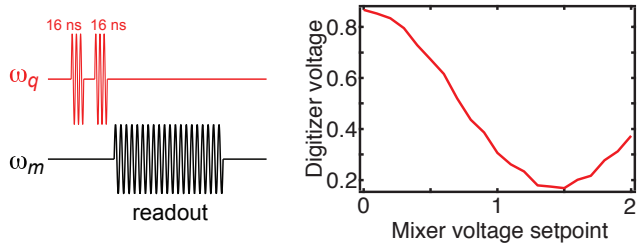

Figure 5.8: Pulse calibration. We apply two 16 ns pulses with a few ns seperation, and then read out the qubit state. We scan the voltage setpoint which determines the qubit pulse power and look for the maximum response in the digitizer voltage, which corresponds to the qubit being prepared in the excited state after two 16 ns $\pi/2$ pulses.

measurement, as depicted in figure 5.7c. We apply a π pulse to bring the qubit into the excited state, and then wait for a variable amount of time τ_d, over which the qubit may or may not undergo spontaneous emission and relax back to the ground state. After performing many measurements for each τ_d, the average measurement signal will decay exponentially, with a timescale given by T_1.

5.4.2 Pulse calibration

After using a Ramsey measurement to find the qubit frequency, we can perform a pulse calibration. The phase of our qubit pulses will determine about which axis in the x-y plane of the Bloch sphere we will drive rotations. For the experimental sequences that we'll use in the trajectories experiments, we need to calibrate $\pi/2$ pulses about both the $\hat{x}$ and $-\hat{y}$ axes. For simplicity we use rectangular pulse envelopes filtered by 100 MHz low pass filters on the output of the AWG. The pulse speed and fidelity can be improved and leakage to higher transmon levels can be limited by using pulse shaping techniques [82], but by simply using 16 ns rectangular pulses we achieve pulse errors below a few percent, which is sufficient for our purposes and not a dominant source of experimental uncertainty.

As illustrated in figure 5.8, we calibrate a 16 ns $\pi/2$ pulse by applying two 16 ns pulses separated in time by a few ns and then perform a projective measurement of the qubit state. We scan the voltage setpoint of the gating pulse that we send to the mixer to generate the qubit pulse and fit for the maximum response in the digitizer voltage, which corresponds to the qubit being prepared in the excited state.

5.4.3 Projective measurement fidelity and multi-level readout

As we discussed in section 2.2.2, to optimize the projective measurement fidelity we need to choose the optimal measurement strength and integration time for a given set of qubit parameters. The measurement timescale τ and the integration time T will determine the

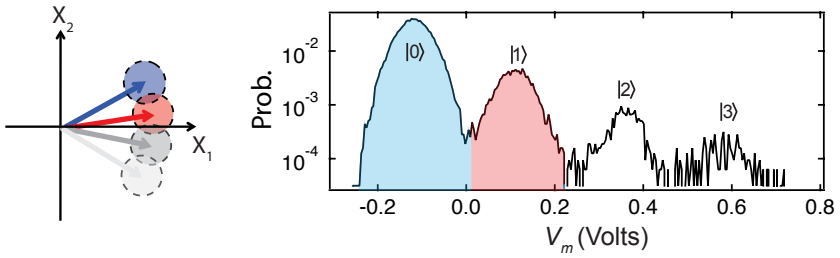

Figure 5.9: Single-shot, multi-state quantum non-demolition measurement. In the left panel, the readout signals corresponding to multiple qubit states are depicted schematically in the X_1-X_2 plane. In the right panel, the graph shows the experimental distribution of digitizer voltage V_m corresponding to the amplified X_2 quadrature. Note that the y axis is plotted on a log scale. In the regime of small χ/κ several states of the transmon are visible.

separation fidelity[4] F_s. Decreasing τ (increasing the measurement rate) will increase the separation ΔV between the histograms, therefore increasing F_s. Increasing T will decrease the width of the measurement histograms, also increasing F_s.

If we try to increase the separation fidelity too far by increasing $\bar{n}$ or or T, then the total measurement fidelity will eventually be limited by qubit-state transitions during measurement. The most obvious mechanism for state transitions is environment-induced transitions from T_1 processes. Additionally, as the measurement rate increases the dispersive approximation gradually breaks down and the measurement becomes less QND, which results in an increasing rate of measurement-induced qubit-state transitions[105]. As the integration time increases, the measurement errors from both of these transition mechanisms will increase.

A further consideration when biasing up our projective measurement is the ability to distinguish between multiple qubit levels. In the limit of small χ/κ, the phase shifts corresponding to the second excited state of the qubit, as well as even higher excited states, are still small enough to resolve by phase sensitive amplification. Figure 5.9 shows an experimental histogram for a projective measurement from reference [27], where we are able to distinguish between four different qubit states with high single-shot fidelity. If we allow our qubit to thermalize between measurements, then such a measurement histogram is an excellent method to calibrate the qubit temperature.

In our trajectories experiments, we tuned up the projective measurement by hand, by tuning the measurement rate and integration time and observing the effect on the measurement histograms. We know that the integration time is too long when we start to see a flat region between the ground and excited state histograms, which indicates qubit state transitions during measurement. Once we have tuned up our measurement histograms, in order to determine the qubit state from a single-shot measurement we need to choose a threshold

[4]Recall from section 2.2.2 that F_s is defined to be the maximal possible measurement fidelity F for a given T and τ.

voltage V_{th}. For the histogram in figure 5.9, a measurement outcome below V_{th} (blue) is counted as the ground state, and an outcome above V_{th} and below the threshold for the second excited state (red) is counted as the excited state. From an ensemble of repeated experimental iterations, we can use the single-shot measurement results to determine $\langle\hat{\sigma}_z\rangle$. We define the number of iterations found in the ground (excited) state as N_0 (N_1), and ignore any iterations that were found in higher qubit levels. Then,

$$\langle\hat{\sigma}_z\rangle = \frac{N_0 - N_1}{N_0 + N_1} \tag{5.3}$$

Note that by this definition the ground state corresponds to $\langle\hat{\sigma}_z\rangle = 1$, and the excited state corresponds to $\langle\hat{\sigma}_z\rangle = -1$, and the ground state sits at the north pole of the Bloch sphere[5]. We can characterize our measurement fidelity by separately preparing the qubit in the ground and excited states and measuring $\langle\hat{\sigma}_z\rangle$. In reference [27], our choice of τ and T for our projective measurement to resulted in $S = 42$, corresponding to a separation fidelity $F_s = 99.9\%$. Due to state transitions during measurement, we achieved a total measurement fidelity of $F \sim 95\%$. This measurement fidelity could be improved by more carefully optimizing τ and T (for example by reducing T so that the measurement errors and separation errors are better balanced), by increasing the qubit T_1, and by employing measurement pulse shaping and more sophisticated signal processing to account for the cavity bandwidth.

5.4.4 Heralded state preparation

Once we have access to a high fidelity single-shot measurement, we can use the measurement to prepare the ground state with high fidelity. This is accomplished by a technique known as heralded state preparation [106]. At the beginning of an experimental sequence, we perform a projective measurement. In any further analysis, we only consider the sub-ensemble of experimental iterations where the initial measurement resulted in the ground state. Within this heralded sub-ensemble, the qubit is prepared in the ground state with high fidelity. This technique is useful if there is a substantial equillibrium thermal excited state population, and also if we would like to repeat the experiment faster than it takes for the qubit to re-equillibrate. Whenever possible, we also perform a projective measurement at the end of the sequence and post-select on the $\{|0\rangle, |1\rangle\}$ subspace by ignoring any experimental iterations where the qubit is found to be in the second excited state or higher.

5.4.5 Quantum state tomography

Quantum state tomography is a procedure which uses ensemble measurements to reconstruct the state of a quantum system. A general (mixed) qubit state is completely specified

[5]Recall that we chose this definition to be consistent with nuclear magnetic resonance terminology. If a spin is sitting in an external magnetic field which points up, the lower energy spin state (the ground state) will also point up.

by three quantities, $\langle\hat{\sigma}_x\rangle, \langle\hat{\sigma}_y\rangle,$ and $\langle\hat{\sigma}_z\rangle$. We can measure $\langle\hat{\sigma}_z\rangle$ by simply performing an ensemble of projective measurements. By applying a $\pi/2$ pulse about the $\hat{y}$ (respectively, $-\hat{x}$) axis before measuring $\hat{\sigma}_z$, we perform an effective measurement of $\langle\hat{\sigma}_x\rangle$ (respectively, $\langle\hat{\sigma}_y\rangle$).

Through this simple procedure, we can fully reconstruct the qubit density matrix to within the accuracy of our qubit pulses and measurement fidelity. If there is a concern about the accuracy of the qubit pulses, one could instead perform quantum state tomography for a range of different pulse amplitudes and angles, thus mapping out the qubit state in a 2D rotation space as described in reference [107]. An imperfect measurement fidelity is less significant than pulse errors, as it will simply rescale the tomography results and can be corrected for with a multiplicative factor.

5.4.6 AC Stark shift calibration

In dispersive cavity QED, the ac Stark shift is the shift in the qubit frequency when the cavity is populated with a coherent tone. Thus far, we have thought of the interaction term in the dispersive Jaynes-Cummings Hamiltonian, $H_{int} = -\hbar\chi\hat{a}^\dagger\hat{a}\hat{\sigma}_z$, as a qubit-state dependent shift in the cavity frequency. However, it equivalently results in a cavity-photon-number-dependent qubit frequency. A coherent state is a superposition of photon number eigenstates characterized by an average photon number $\bar{n}$. The qubit frequency shift due to a coherent cavity state is proportional[6] to $\bar{n}$, and given by[108]

$$\Delta\omega = 2\chi\bar{n}. \tag{5.4}$$

In addition to this frequency shift, a coherent state cavity occupation will also cause the a qubit superposition state to dephase, due to the quantum fluctuations of the cavity field. The fluctuations in both amplitude and the phase will contribute equal amounts of dephasing, but through different mechanisms. The fluctuations in amplitude correspond to fluctuations in $\bar{n}$, which will cause the qubit frequency to fluctuate. Qubit frequency fluctuations are equivalent to random qubit motion along the equator of the Bloch sphere, which results in ensemble dephasing. On the other hand, the phase of the cavity field is related to $\hat{\sigma}_z$. Fluctuations in the cavity phase will cause random qubit motion along the meridian of the Bloch sphere, which also leads to ensemble dephasing. The total qubit dephasing due to a coherent cavity occupation is referred to as 'measurement-induced dephasing', and occurs at a rate of Γ_{MID}. The measurement-induced dephasing rate is given by[108, 109]

$$\Gamma_{MID} = \frac{8\chi^2\bar{n}}{\kappa}. \tag{5.5}$$

[6]Here we consider only the linear ac Stark shift, where the $\Delta\omega$ is proportional to $\bar{n}$. For very large photon number, and in certain parameter regimes of the Jaynes-Cummings Hamiltonian, higher order terms in $\bar{n}$ can also contribute to the frequency shift.

A Ramsey measurement can act as a sensitive probe of both $\Delta\omega$ and Γ_{MID}. We apply a weak continuous measurement to the cavity which will remain on throughout the free evolution period of the Ramsey measurement. We intentionally increase the qubit pulse carrier frequency from the qubit frequency by $\delta\omega/2\pi \sim 1$ MHz in order to see Ramsey oscillations. We scan the power of the measurement tone and record the Ramsey oscillations for each measurement power. By fitting the Ramsey oscillations, we extract the Ramsey oscillation frequency, which is equal to $\delta\omega + \Delta\omega$, and the ensemble decay rate, $\Gamma = \Gamma_{MID} + 1/T_2^*$. Since $\bar{n}$ is proportional to the measurement power, both $\Delta\omega$ and Γ will be linear in $\bar{n}$. We plot these two lines and fit for their slopes, from which we can solve for χ/κ.

Since we already know κ precisely from our cavity measurements, the AC Stark shift is an excellent method of calibrating χ. Once we know χ and $\Delta\omega$ for each measurement power setpoint, then we also have a calibration for $\bar{n}$ at each setpoint.

5.4.7 Measurement strength and quantum efficiency calibration

We can determine the signal-to-noise ratio $S \equiv (\Delta V_{\text{opt}}/\sigma)^2$ for a measurement of length T by separately preparing the ground and excited states and fitting the resulting measurement voltage histograms to Gaussians functions[7]. From S, we can also compute the characteristic measurement time $\tau = 4T/S$. As we can see from equation (3.33) ($S = 64T\chi^2\bar{n}\eta_m/\kappa$), S will depend linearly on $\bar{n}$. Since we know $\bar{n}$ and κ from the AC Stark shift calibration, the slope of S vs $\bar{n}$ will determine the measurement quantum efficiency η_m. Figure 5.10 plots experimental results for S as a function of $\bar{n}$ for the sample parameters used in reference [27]. The solid line shows the expected dependence for $\eta_m = 0.49$. Once we have determined η_m, the total quantum efficiency is given by $\eta \equiv \eta_m\eta_{env}$, where[53]

$$\eta_{env} = \frac{1}{1 + \kappa/8\chi^2\bar{n}T_2^*}. \tag{5.6}$$

In reference [28], we repeat this procedure for a number of different measurement times T, from which we can calibrate τ and η independently for each measurement strength. We can plot S vs T, and fit to a line to extract the slope dS/dT. Then the measurement timescale is given by

$$\tau = 4\left(\frac{dS}{dT}\right)^{-1}. \tag{5.7}$$

Then, to complete the η calibration, we simply measure the ensemble decay rate Γ using an $\hat{x}$ tomography measurement and fitting $\langle\hat{\sigma}_x\rangle \propto e^{-\Gamma T}$. Recall from the pervious section that both the environmental dephasing and the measurement induced dephasing contribute to Γ. The total quantum efficiency is given by

[7] Recall that ΔV_{opt} is the separation between the ground and excited state measurement histograms when we have chosen to amplify the quadrature which provides the maximal qubit-state information.

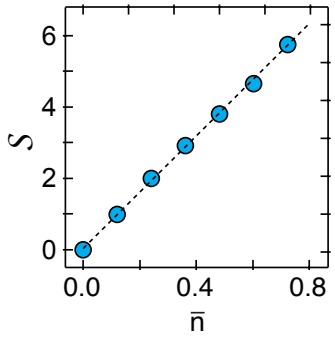

Figure 5.10: S vs $\bar{n}$. The solid line depicts the expected dependence for $\eta_m = 0.49$.

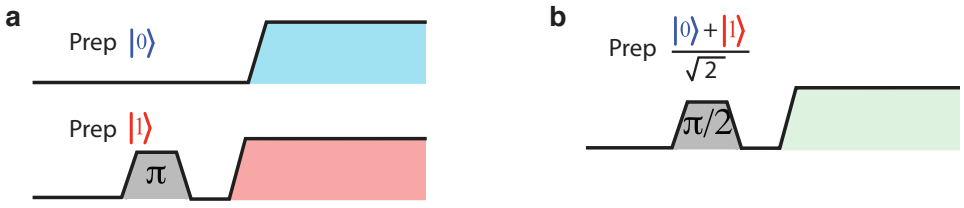

Figure 5.11: Measurement calibration sequences. a) We calibrate the measurement rate τ by preparing the qubit in the ground and excited states followed by a weak measurement. Using the full measurement record, we can extract the width of the measurement histograms as a function of integration time. b) In order to calibrate the total quantum efficiency η, the only additional piece of information that we need is the ensemble dephasing rate Γ. We could get this value from a Ramsey measurement, but it is more convenient (and less susceptible to drifts) if we use the ensemble decay of the coherence in our actual trajectories dataset, which begins by preparing a superposition state and is followed by a weak measurement.

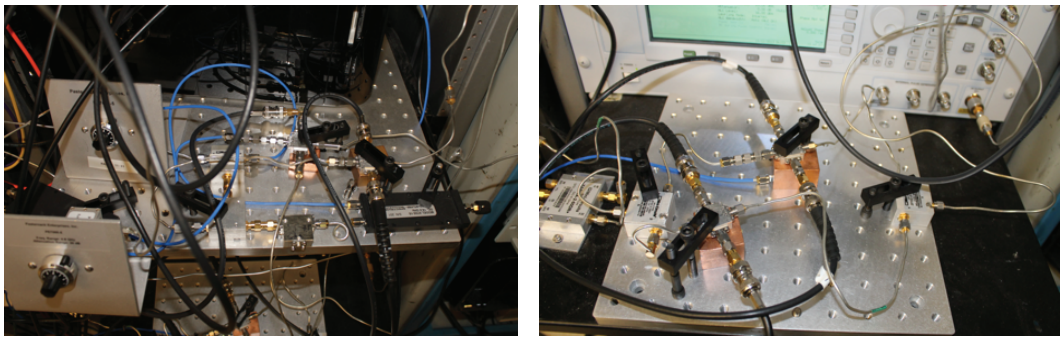

Figure 5.12: Mounting room temperature microwave components to optical breadboards helps with temperature and mechanical stability and is a convenient way to lay out the experimental setup.

$$\eta = \frac{1}{2\tau\Gamma}. \tag{5.8}$$

5.4.8 Experimental stability

After calibrating our experiment, we need the parameters to remain stable for long enough to take a complete dataset. One concern is the stability of room temperature components, especially the mixers. Mixer offset voltages tend to drift on the timescale of ~ 1 day. It is also possible for the mixer response to drift, requiring a recalibration of qubit pulses. We have found that anchoring microwave components to the optical breadboards, as pictured in figure 5.12, helps with temperature stability and reduces drifts.

Another concern is the stability of the qubit parameters. T_1 and T_2^* times are known to fluctuate by at least 10%, as discussed in the supplemental information of reference [77], but since our experiments take place on timescales much shorter than T_1 and T_2^* these fluctuations are insignificant. The qubit frequency is usually very stable over the course of many days. However, as pictured in figure 5.13 on some cooldowns the qubit frequency is highly unstable, randomly jumping around on a timescale of ~ 10 minutes. This is believed to be the result of a two-level system (TLS) defect in the junction, which is only active at the qubit frequency during some cooldowns. This problem will often go away if we simply cycle the fridge above the superconducting transition temperature for aluminum and cool back down.

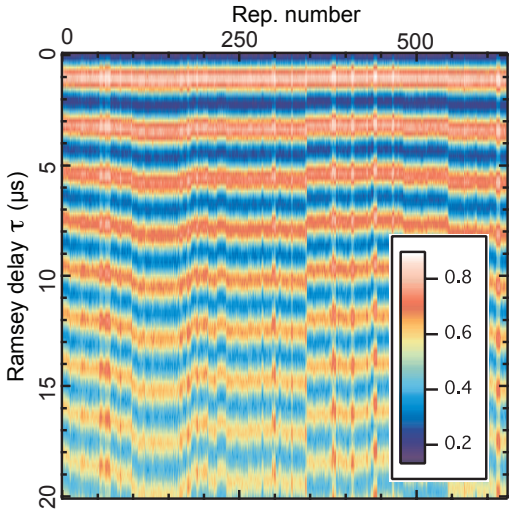

Figure 5.13: Repeated Ramsey measurement. Ramsey measurement repeated over several hours. Color scale indicates the average digitizer voltage for each Ramsey delay time τ. In this measurement, a defect would occasionally come in and out of interaction with the qubit, causing a shift in the qubit frequency, as evidenced by a shift in the Ramsey oscillation frequency.

Chapter 6

Tracking individual quantum trajectories

In this chapter, we show how we use the measurement signal track individual quantum trajectories of the qubit state. We then explain how we use conditional quantum state tomography to verify that we have faithfully tracked the trajectories.

6.1 Quadrature-dependent measurement backaction

Consider a measurement tone initially pointing in the X_1 direction, which reflects off of a cavity and acquires a qubit-state-dependent phase shift, as illustrated in figure 6.1. The quantum fluctuations of the measurement tone will disturb the qubit state, causing an ensemble dephasing rate of $\Gamma_{MID} = 8\chi^2\bar{n}/\kappa$. As we mentioned in section 5.4.6, amplitude and phase fluctuations of the measurement tone cause two separate dephasing mechanisms. Amplitude fluctuations are related to fluctuations in the intra-cavity photon number, which leads to a fluctuating AC Stark shift. In an individual measurement, the integrated AC Stark shift will cause a unitary rotation of the qubit state about the $\hat{z}$ axis of the Bloch sphere. In reference [53], Korotkov refers to this as the 'realistic' backaction, because it can be explained by a simple physical mechanism, without considering the process of quantum measurement. In an ensemble measurement, this random backaction will cause dephasing at a rate which we define to be $\Gamma_{unitary}$.

On the other hand, the phase of the measurement signal is related to the qubit state in the $\hat{\sigma}_z$ basis. In an individual measurement, the integrated phase fluctuations will determine the amount of measurement backaction, which rotates the qubit state toward the poles of the Bloch sphere. Korotkov refers to this as the 'spooky' backaction. Recall from chapter 2 that measurement backaction can be though of as two processes: the acquisition of information about $\hat{\sigma}_z$ and a concurrent reduction in $\hat{\sigma}_x$, as required by the Heisenberg's uncertainty principle. In an ensemble measurement, the Heisenberg backaction in $\hat{\sigma}_x$ will cause dephasing at a rate which we define to be Γ_{Heis}. In the limit of small χ/κ, the amplitude of the measurement is determined by the X_1 quadrature and the phase by X_2. Thus, if we initially prepare the qubit in the the state $\hat{\sigma}_x = +1$, then fluctuations in X_1 will drive rotations

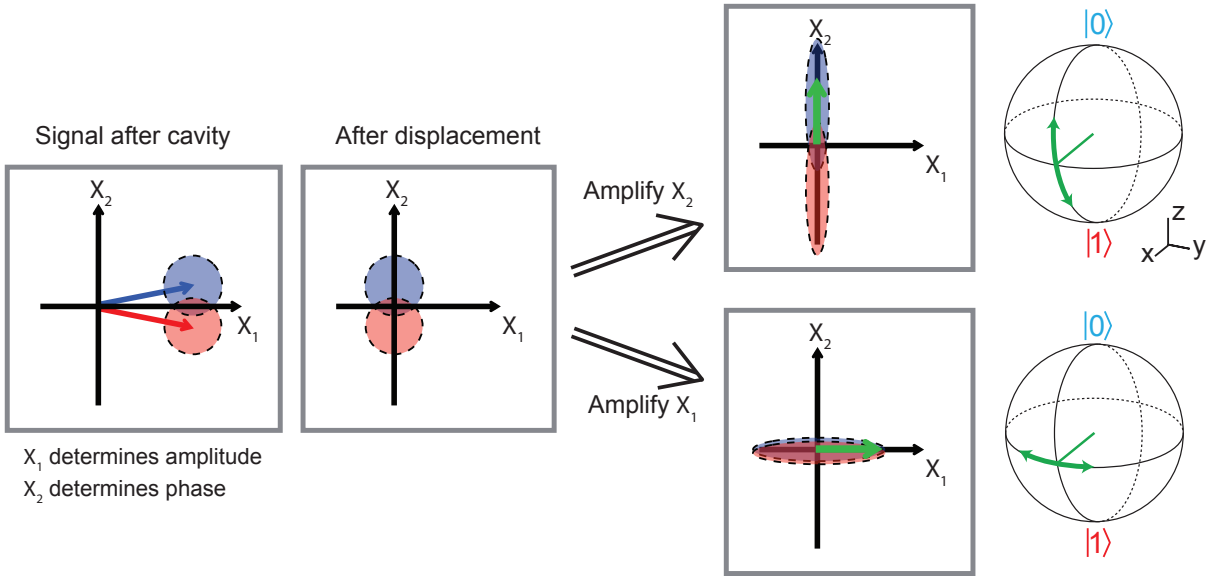

Figure 6.1: 'Z' and 'ϕ' measurements. A measurement tone, initially pointing along the X_2 quadrature, reflects off of the cavity and acquires a qubit-state-dependent phase shift. Its amplitude is determined by the value of X_1 and its phase by X_2. After the signal is displaced back to the origin of the X_1-X_2 plane, we can choose which quadrature of the signal to amplify. If we amplify the X_2 quadrature, which contains qubit-state information, then the measurement backaction will cause the qubit state to move up and down the meridian of the Bloch sphere. If instead we amplify the X_1 quadrature, which contains information about the fluctuating intra-cavity photon number, the qubit state will move along the equator of the Bloch sphere.

along the equator of the Bloch sphere, and fluctuations in X_2 will drive rotations along the meridian, as pictured in figure 6.1.

After the measurement tone leaves the cavity, we have a choice of how to amplify it. Importantly, the qubit ensemble decay rate Γ_{MID} will be the same regardless of how we choose to process the measurement signal, as required by causality. However, our choice of amplification scheme will have a significant impact on the backaction of an *individual* measurement. If we choose to perform phase-sensitive amplification of the X_2 quadrature, which contains information about $\hat{\sigma}_z$, then we will also squeeze the fluctuations of the X_1 quadrature. In the case of quantum-limited amplification, we cannot measure the photon number fluctuations, and therefore $\Gamma_{unitary} = 0$. Thus, $\Gamma_{Heis} = \Gamma_{MID}$, and the rate of information extraction $\Gamma_{meas} \equiv 1/\tau$ is given by $\Gamma_{meas} = 2\Gamma_{Heis} = 2\Gamma_{MID}$ [1]. We will refer to

[1]Recall that our definition of the measurement rate Γ_{meas} is consistent with reference [48], but differs by a factor of two from the definition used in [49], which instead defines $\Gamma_{meas} = \Gamma_{MID}$ when $\eta = 1$. An additional point of confusion is that reference [49] refers to the measurement induced dephasing as Γ_ϕ, a

this choice of amplification axis as a '$z-$measurement'.

If instead we perform phase-sensitive amplification of the X_1 quadrature, which contains information about $\bar{n}$, then we will will squeeze away the $\hat{\sigma}_z$ information contained in X_2. In this case, there is no measurement backaction, $\Gamma_{meas} = 2\Gamma_{Heis} = 0$, and thus $\Gamma_{unitary} = \Gamma_{MID}$. We will refer to this choice of amplification axis as a '$\phi-$measurement', since the the backaction affects the phase ϕ of a qubit superposition state.

The last case to consider is phase-preserving amplification, where both backaction mechanisms contribute equally to the ensemble dephasing, $\Gamma_{MID} = \Gamma_{Heis} + \Gamma_{unitary} = 2\Gamma_{Heis}$. Therefore, $\Gamma_{meas} = \Gamma_{MID}$, which is half as fast as the measurement rate for phase-sensitive amplification of X_2. This can be conveniently understood as the result of the extra half photon of noise which is added in quantum-limited phase-preserving amplification. It is important to note that this added noise is fundamental to the process of measuring both quadratures, and therefore does not enforce any limit on the quantum efficiency of phase-preserving amplification, provided that the measurement signal in both quadratures is accounted for. If we were to ignore the information in the X_1 quadrature, then the unitary backaction would limit the quantum efficiency to 0.5. Reference [63] investigates the backaction of CQED measurements where the measurement signal is amplified by a phase-preserving Josephson parametric converter (JPC). The work presented in the thesis is restricted to phase-sensitive amplification.

In the case of phase-sensitive amplification, we could in principle chose the amplification axis after the measurement signal leaves cavity. This choice determines the type of backaction that the measurement will have on the qubit, and could be used to steer the qubit evolution. The concept of 'steerability' was developed by Howard Wiseman [110, 111] as a way to quantify the correlations between two quantum objects through which a choice of how to measure one will affect the state of the other. They derive an 'EPR steering' criterion, which is more a more strict indication of non-classical correlations than entanglement and less strict than a violation of Bell's inequality. While it would certainly be interesting to investigate steering in our system, especially if we can further improve η, it should be noted that there is still some debate[2] as to whether EPR steering can be applied to our system, since we can choose between measurement backaction in $\hat{\sigma}_z$ and unitary backaction from photon number fluctuations, rather than choosing between two different bases of quantum measurement backaction.

label which we will reserve in this thesis for the pure dephasing rate, which will become important in chapter 8.

[2]This is my current understanding based on a personal conversation with Jay Gambetta.

6.2 Correlations between measurement outcomes and the qubit state

6.2.1 Bayesian state update

In this section, we will consider the backaction of an individual measurement. We will first consider the case where we prepare the qubit in an initial state $\sigma_x = +1$ and perform a measurement with an average cavity photon number $\bar{n}$ for a length of time T. For a 'z-measurement', the signal-to-noise is given by equation (3.33), $S = \left(\frac{\Delta V}{\sigma}\right)^2 = 64\chi^2\bar{n}\eta_m T/\kappa$. For each experimental iteration, we will use the time-averaged measurement outcome V_m to compute the qubit state after measurement. In other words, we want to find the qubit state conditioned on the measurement outcome V_m, which we can express in terms of the conditional Bloch sphere components $x \equiv \langle\hat{\sigma}_x\rangle|V_m$, $y \equiv \langle\hat{\sigma}_y\rangle|V_m$, and $z \equiv \langle\hat{\sigma}_z\rangle|V_m$[3].

In section 2.3.4, we used a Bayesian analysis to derive the state update equations for a $z-$measurement, which we will denote with the superscript 'z', which are given by

$$z^z = \tanh\left(\frac{V_m S}{2\Delta V}\right), \tag{6.1}$$

$$x^z = \sqrt{1-(z^z)^2}e^{-\gamma T}, \tag{6.2}$$

$$y^z = y(0) = 0, \tag{6.3}$$

where $\gamma = 8\chi^2\bar{n}(1-\eta_m)/\kappa + 1/T_2^*$ is the rate of decay of qubit coherence due to measurement inefficiency and environmental dephasing. γ can also be expressed in terms of the total dephasing rate $\Gamma \equiv \Gamma_{MID} + 1/T_2^*$, as $\gamma = \Gamma - 1/2\tau$. The first term on the right-hand side of equation (6.2) expresses the Heisenberg backaction, the fact that x will decrease as we acquire information about z.

If we instead perform a $\phi-$measurement, the values of x and y are periodic in the accumulated qubit phase shift, and are given by [53, 27]

$$x^\phi = \cos\left(\frac{SV_m}{2\Delta V}\right)e^{-\gamma T}, \tag{6.4}$$

$$y^\phi = -\sin\left(\frac{SV_m}{2\Delta V}\right)e^{-\gamma T}, \tag{6.5}$$

$$z^\phi = z(0) = 0, \tag{6.6}$$

where the superscript 'ϕ' denotes a $\phi-$measurement. Note that the dephasing rate γ due the unaccessible part of the measurement signal is the same regardless of our choice of amplification axis.

[3]The conditional Bloch sphere components can be equivalently defined in terms of the qubit density matrix as $x \equiv \mathrm{tr}[\rho\hat{\sigma}_x]$, $y \equiv \mathrm{tr}[\rho\hat{\sigma}_y]$, and $z \equiv \mathrm{tr}[\rho\hat{\sigma}_z]$.

6.2.2 Conditional quantum state tomography

Conditional quantum state tomography is a method that we can use to verify that we can accurately account for the backaction of an individual measurement. For a given measurement outcome V_m, we can perform quantum state tomography on the sub-ensemble of experimental iterations with similar measurement outcomes, in the range $V_m \pm \epsilon$, where $\epsilon \ll \Delta V$. For superconducting qubits, this technique was first introduced by Hatridge in reference [63], which considers the case of phase-preserving amplification.

Our experimental procedure for conditional quantum state tomography is illustrated in figure 6.2a. We prepare the state $\hat{\sigma}_x = +1$ by heralding the ground state and applying a $\pi/2$ pulse about the $-\hat{y}$ axis of the Bloch sphere. We then weakly measure the qubit state for $1.8\,\mu$s and record the measurement outcome V_m. Finally, we an optional tomography pulse and perform a projective measurement. We repeat the experiment $\sim 10^5$ times for each orientation of tomography pulse ($\pi/2$ about $\hat{y}$, $\pi/2$ about $-\hat{x}$, and no pulse). For each experimental iteration, we record V_m and the binary projective measurement result. Tomography results for matching V_m are averaged together for determine x, y, and z. Panels (c) and (d) show the results for a $z-$ and $\phi-$ measurements, respectively. Dashed lines show theoretical results based on independently calibrated sample and measurement parameters. Solid lines show the results of conditional quantum state tomography. The excellent agreement between theory and experiment provides evidence that have a firm understanding of the correlations between the measurement outcome and the qubit state and that we can accurately compute the qubit state after an individual measurement.

6.3 Quantum trajectories

6.3.1 Calculating individual trajectories

So far, we have demonstrated the ability to reconstruct the qubit state after a measurement of fixed time T. There is nothing special about our choice of $T = 1.8\,\mu$s, and in principle our state update equations should work for any choice of T. In fact, we can divide a continuous measurement record $V(t)$, such as the one displayed in figure 6.3b, into successive segments and perform a state update for each timestep. The integrated measurement signal can be written as a string $\{V_m(t_0), V_m(t_1), V_m(t_2), ...\}$, where $\Delta t \equiv t_{i+1} - t_i = 16$ ns[4]. At each time point, $V_m(t_i)$ can be used to infer the qubit state. The strings composed of $\{x(t_i)\}$, $\{y(t_i)\}$, and $\{z(t_i)\}$ define the quantum trajectory. In the limit of small Δt, we can think of these strings as continuous quantum trajectories, $x(t), y(t), \text{and} z(t)$. Figure 6.3c shows the quantum trajectory corresponding to the measurement record in panel (b).

[4]Recall from chapter 2 that $V_m(t)$ is defined as $V_m(t) \equiv (1/t)\int_0^t V(t)\,dt$,, and thus each $V_m(t_i)$ represents the full time-averaged measurement signal up until time t_i.

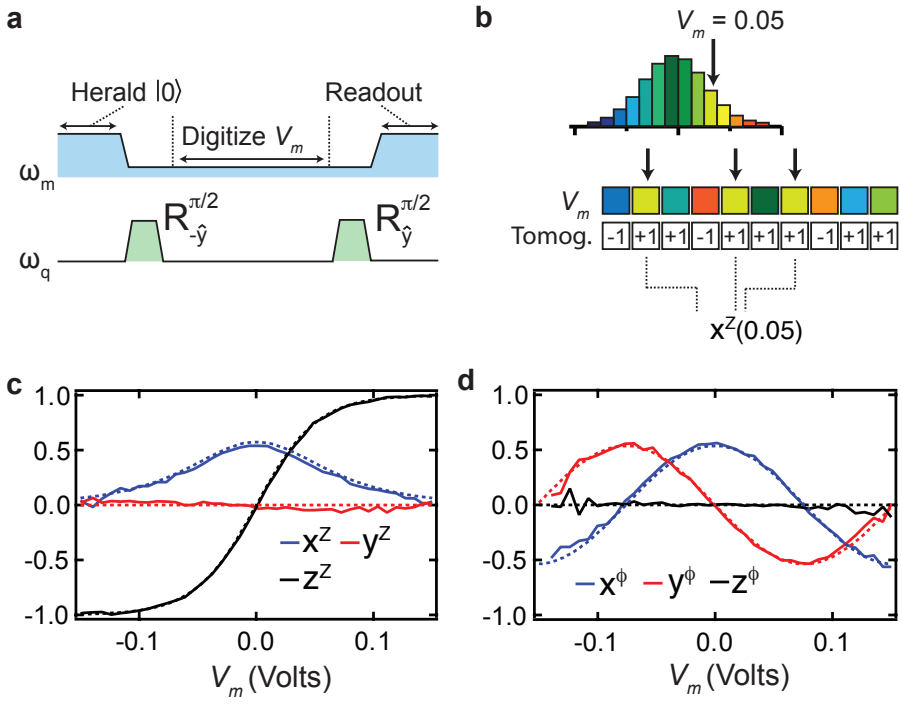

Figure 6.2: Conditional quantum state tomography. a) Pulse sequence for reconstructing x^z. We prepare the qubit along $\hat{x}$ by heralding the ground state and applying a $\pi/2$ pulse about the $-\hat{y}$ axis. We then weakly measure the qubit state for 1.8 μs and finally measure the projection of the qubit state along $\hat{x}$ by applying another qubit rotation and a projective qubit measurement. We repeat the experiment many times, recording the weak measurement outcome V_m and the binary projective measurement result for each experimental iteration. b) Tomography correlation procedure. The measurement outcomes are divided into discrete bins and histogramed, and different measurement values are indicated as different colors. Boxes indicate the measurement outcome, V_m, for each experimental iteration, and the associated tomography result is indicated as ± 1. Tomography results for matching V_m are averaged together to determine x^z. c) Tomography results versus V_m for a $z-$measurement with $\bar{n} = 0.4$. The dashed lines are independent theory curves based on equations (6.1)-(6.3) for $\eta_m = 0.49$ and $S = 3.15$, where $\gamma = 2.8\times10^5\,\mathrm{s}^{-1}$. d) Tomography results for a $\phi-$measurement for $\bar{n} = 0.46$. The dashed lines are independent theory curves based on equations (6.4)-(6.6) for $\eta = 0.49$ and $S = 3.62$, where $\gamma = 3.1\times10^5\,\mathrm{s}^{-1}$.

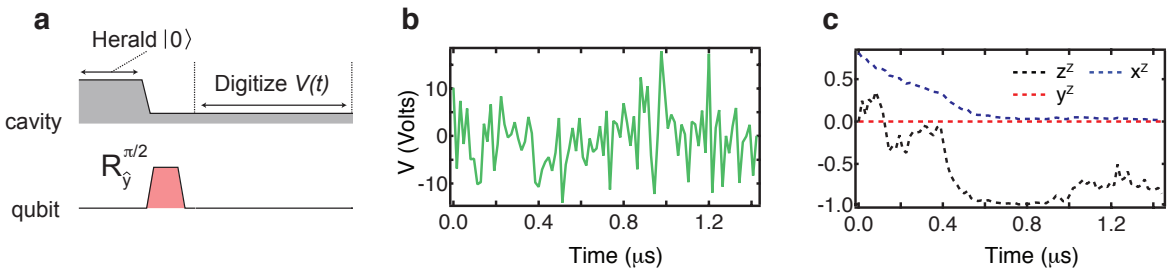

Figure 6.3: Calculating individual quantum trajectories. a) Pulse sequence used to generate individual quantum trajectories. b) A measurement signal $V(t)$ from an individual experimental iteration, binned in timesteps of $\Delta t = 16$ ns. c) Using the integrated measurement signal for each timestep, we calculate the quantum trajectory for the measurement signal shown in panel (b).

6.3.2 Tomographic reconstruction

Just as we used conditional quantum state tomography to validate our calculation of the post-measurement qubit state for a single measurement time T, we can also use conditional quantum state tomography to validate a quantum trajectory. Suppose that we use an individual measurement signal $\widetilde{V}(t)$ of length T to generate a target quantum trajectory specified by $\{\tilde{x}(t_i)\}, \{\tilde{y}(t_i)\}$, and $\{\tilde{z}(t_i)\}$), from the integrated measurement signal $\{\widetilde{V}_m(t_i)\}$. Here, the tildes indicate that we have selected a specific target trajectory which we will attempt to validate by performing conditional tomography at each time stop of the trajectory.

To do this we need to perform a separate set of experiments for each time t_i and for each orientation of tomography pulse. Thus, to reconstruct a quantum trajectory of length $T = 1.6\,\mu$s with a timestep $\Delta t = 16$ ns requires an experimental sequence consisting of $3 \times 101 = 303$ different experiments, each of which is repeated $\sim 10^5$ times. For each experimental iteration, we record the measurement record $V(t)$, and the binary result of the final projective measurement. For each timestep, we tomographically reconstruct the conditional quantum state $(x(t_i), (y(t_i), z(t_i))$, using only the subset of experimental iterations where $V_m(t_i)$ is within the range of $\tilde{V}_m(t_i) \pm \epsilon$. Note that because of the fact that the measurements in in each timestep commute with each other, we do not need to match the exact measurement signal, just the integrated signal $V_m(t_i)$. Experimental results are depicted in figure 6.4, which clearly demonstrate our ability to track an individual quantum trajectory.

6.3.3 Trajectories under driven evolution

We can also reconstruct quantum trajectories under driven evolution, where we simultaneously measure at a rate $1/\tau$ and apply a qubit drive characterized by the Rabi frequency Ω, as pictured in figure 6.5. The Rabi evolution is described by the Hamiltonian

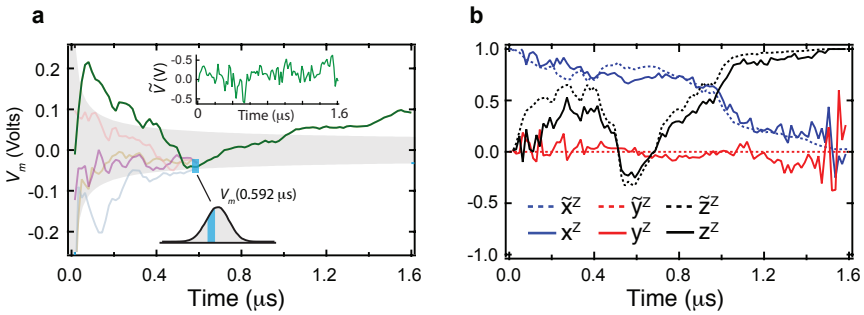

Figure 6.4: Trajectory validation procedure. a) Time-averaged measurement signal $\widetilde{V}_m(t_i)$ (dark green) corresponding to the individual measurement record displayed in the inset. The grey region indicates the standard deviation of the distribution $V_m(t)$. Measurement traces that converge to an integrated value within the blue matching window are used to reconstruct, through conditional quantum state tomography, the trajectory at that time point. In lighter colors, we display a few of the measurement traces that contribute to the reconstruction at $0.592\,\mu$s. b) Dotted lines indicate the quantum trajectory obtained from the green measurement signal. Solid lines indicate the tomographically reconstructed quantum trajectory based on the sub-ensemble of measurements at each time-step that are within the matching window of the original measurement signal.

$$H_R = \hbar\frac{\Omega}{2}\hat{\sigma_y}. \tag{6.7}$$

For a given measurement strength (and thus a known Γ), we can determine Ω by fitting the ensemble evolution to a Lindblad equation with arbitrary Rabi drive: $\dot{x}(t) = -\Gamma x(t) + \Omega z(t)$, $\dot{z}(t) = -\Omega x(t)$. With the initial state (x_0, z_0), these equations have an analytic solution

$$x(t) = e^{\Gamma t/2}\left(x_0\cos(\lambda t) - \frac{\Gamma x_0 - 2\Omega z_0}{2\lambda}\sin(\lambda t)\right) \tag{6.8}$$

$$z(t) = e^{\Gamma t/2}\left(z_0\cos(\lambda t) - \frac{\Gamma z_0 - 2\Omega x_0}{2\lambda}\sin(\lambda t)\right) \tag{6.9}$$

where $\lambda = \sqrt{\Omega^2 - (\Gamma/2)^2}$.

Note that that the Rabi Hamiltonian H_R does not commute with the measurement Hamiltonian, which is proportional to σ_z. Thus, to compute the conditional qubit state $(x(t), y(t), z(t))$ at a time t, we need to make use of the entire measurement signal up until t, not just the integrated signal $V_m(t)$. Note that from here on we have dropped the subscript i from t_i for clarity. Also, for the remainder of this thesis we will restrict ourselves to only z-measurements, and therefore drop the superscript z for clarity.

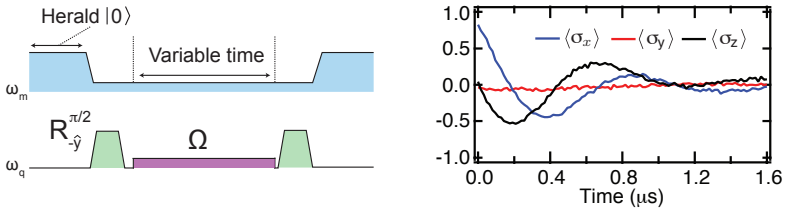

Figure 6.5: Driven evolution. a) Pulse sequence for which the qubit evolves under the competing influences of a weak measurement tone and Rabi drive characterized by the Rabi frequency Ω. b) Ensemble time-evolution, which we fit to determine Ω. For data shown here, $\Omega/2\pi = 1.08$ MHz and $\tau = 315$ ns.

If we work with sufficiently small timesteps $\Delta t \ll \tau, 1/\Omega$ then we can use a sequential two-step procedure to update the qubit state after each timestep, as illustrated schematically in figure 6.6. For each timestep from an arbitrary time t to $t+\Delta t$, we preform a state update based on our knowledge of the state at time t, $\rho(t)$, and on the measurement outcome $V(t)$, which is defined as the measurement signal integrated from time t to $t+\Delta t$. We first apply a unitary rotation to account for the Rabi drive

$$\rho'_{01} = \rho_{01} + \frac{\Omega}{2}(\rho_{00} - \rho_{11})\Delta t \tag{6.10}$$

$$\rho'_{11} = \rho_{11} + \frac{\Omega}{2}(\rho'_{01} + \rho'_{10})\Delta t \tag{6.11}$$

where ρ_{00}, ρ_{01}, ρ_{10}, and ρ_{11} are the matrix elements of $\rho(t)$. With the input values ρ'_{01} and ρ'_{11}, we next apply a Bayesian update based on $V(t)$, using a slight adaptation of equations (2.45) and (2.47):

$$\rho_{11}(t+\Delta t) = \frac{(\rho'_{11}\rho'_{00})\exp(-4V(t)\Delta t/\tau\Delta V)}{1+(\rho'_{11}\rho'_{00})\exp(-4V(t)\Delta t/\tau\Delta V)} \tag{6.12}$$

$$\rho_{01}(t+\Delta t) = \rho'_{01}\frac{\sqrt{(1-\rho_{11}(t+\Delta t))\rho_{11}(t+\Delta t)}}{\sqrt{(1-\rho'_{11})\rho'_{11}}}\mathrm{e}^{-\gamma\Delta t}. \tag{6.13}$$

Using this update procedure, we can generate a target trajectory $(\tilde{x}(t), \tilde{y}(t), \tilde{z}(t))$ from an individual measurement record. To verify that we have accurately tracked the state of the system, we follow the same reconstruction procedure discussed for un-driven trajectories in the previous section. The only difference is that now we must perform the state update procedure for each experimental iteration. For each time t, we only include runs where $x(t) = \tilde{x}(t) \pm \epsilon$ and $z(t) = \tilde{z} \pm \epsilon$ in the tomographic reconstruction. Note for a trajectory to be

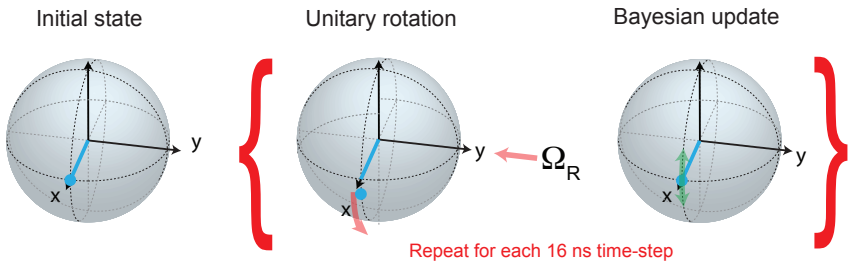

Figure 6.6: Two-step state update. After preparing the qubit in a known initial state, we propagate the qubit state in time steps of $\Delta t = 16$ ns. For each time t, we calculate the qubit state at time $t + \Delta t$ by appling a unitary rotation to account for Rabi evolution and then a Bayesian update based on the measurement result obtained between t and $t + \Delta t$.

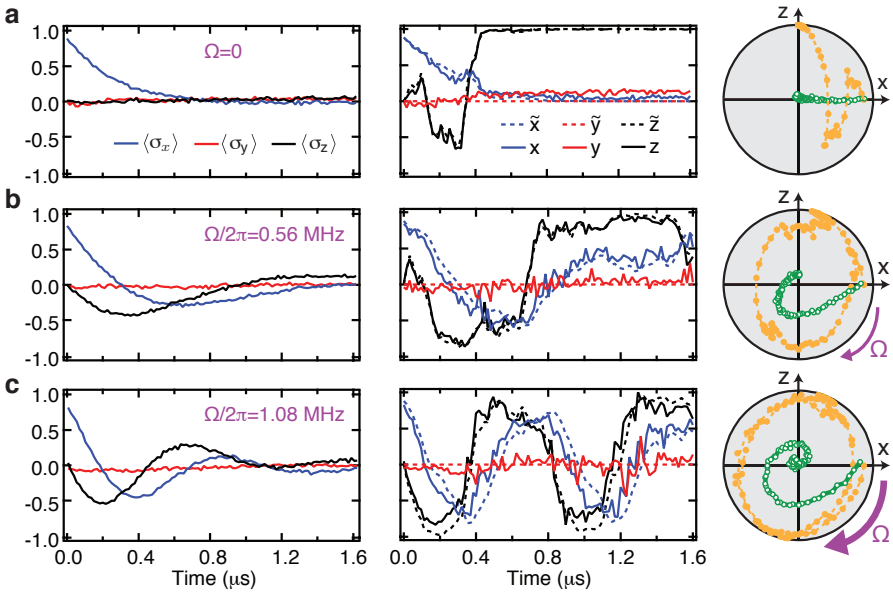

Figure 6.7: Driven quantum trajectories. The left panels depict the full ensemble evolution. The middle panels display individual quantum trajectories (dashed curves) and the corresponding tomographic reconstructions (solid curves). On the right, we plot individual trajectories (orange) and the ensemble averages (green) in the x-z plane of the Bloch sphere. Rows (a), (b), and (c) correspond to different values of the Rabi drive $\Omega/2\pi = 0$ MHz, 0.56 MHz, and 1.08 MHz, respectively. In this figure, $\tau = 315$ ns and $\Gamma = 3.85 \times 10^6 \mathrm{s}^{-}1$.

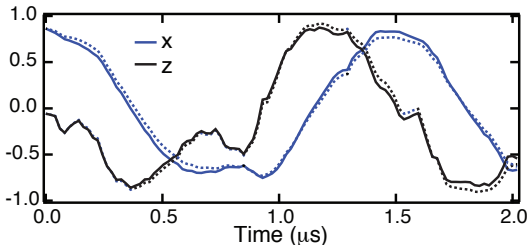

Figure 6.8: Comparison of the quantum trajectory calculated using an SME approach (solid) and the Bayesian update procedure (dashed line). The two techniques give very similar results. Adapted from reference [65].

used in the tomographic reconstruction at time t, it only needs to match the target trajectory at that one time. It's not necessary for the complete trajectory to match. Figure 6.7 shows reconstructed quantum trajectories for $\Omega/2\pi = 0$, 0.56, and 1.08 MHz, with $\tau = 315$ ns, $\eta = 0.44$, and a matching window of $\epsilon = 0.03$. For each Rabi frequency, we repeated the experiment 5000 times for each of 101 time points and each of the 3 tomography orientations. Therefore, for each Rabi frequency we performed a total of $3 \times 101 \times 5000 = 1.5 \times 10^6$ experimental iterations. Note that despite our modest quantum efficiency, the trajectories remain remarkably coherent long after the full ensemble has decohered.

Our experiments are the first on any system to use quantum state tomography at discrete times along a quantum trajectory to verify that we have faithfully tracked the qubit state. The excellent agreement between individual trajectories and their tomographic reconstruction is indicative of our detailed understanding of the measurement process in our system. In order to achieve this level of agreement, we needed to very carefully calibrate the Ω, τ, and the voltage offset of the measurement signal. We also made efforts to reduce the effect of transients (such as the ringdown of the herald pulse) on the measurement signal.

6.3.4 Comparing Bayesian trajectories to SME trajectories

In the experiments presented in this thesis, we chose to use a Bayesian state update procedure because it is simple to apply to our system. However, we could have achieved similar results using state update equations derived from a stochastic master equation, as demonstrated in reference [65]. Figure 6.8 compares quantum trajectories generated from a single measurement signal using the two different approaches, showing that they are nearly identical.

Chapter 7

Ensembles of trajectories

7.1 What can we learn from quantum trajectories?

Investigating quantum trajectories opens a window into the measurement dynamics of open quantum systems. As Korotkov states in reference [53],

> Understanding the qubit evolution in the process of measurement is important for developing an intuition, which is useful in many cases, in particular in designing various schemes of the quantum feedback.

By tomographically reconstructing individual quantum trajectories, we have proven that we can accurately track the qubit state over the course of any individual measurement. Our results provide a sound footing for future efforts to use this understanding to implement various quantum feedback schemes. However, beyond a proof-of-principle, how quantum trajectories will specifically inform future developments in quantum control remains an open question. Nevertheless, our experiments are helpful for building an intuition for how the qubit state is most likely to evolve during measurement.

While it is difficult to build much intuition by looking at a single noisy trajectory, a natural step is to consider ensembles of quantum trajectories. Qualitatively, ensembles of quantum trajectories offer a convenient method of visualizing complex measurement dynamics. In this chapter, we examine distributions of trajectories for a variety of different parameter regimes, investigating the interplay between measurement backaction and unitary dynamics. Furthermore, we discuss an example of a quantitative question that can investigated experimentally only once we have access to these trajectory distributions: what is the most probable path through quantum state space connecting an initial state $|\psi_i\rangle$ and a final state $|\psi_f\rangle$ in a given time T? In reference [28], we address this question in both experiment and theory. Here, we present an extended discussion of this work.

7.2 The optimal path through quantum state space

7.2.1 Action principle for continuous quantum measurement

Here, we will address the problem of finding the most probable quantum trajectory connecting an initial state $|\psi_i\rangle$ at an time $t = 0$ with a final state $|\psi_f\rangle$ at time T. One straightforward approach to this problem would be to solve the stochastic master equation (SME) numerically a large number of times and perform statistical analysis on the sub-ensemble of trajectories which end in $|\psi_f\rangle$ at time T. In this section, we introduce a different approach which is based on an action principle for continuous quantum measurement introduced by our theory collaborators in reference [112].

Classical stochastic processes, for example, the activation of a particle over a potential barrier by a random applied force, are often analyzed by introducing a canonical phase-space structure and minimizing the action subject to certain constraints [113, 114, 115, 116]. For continuous quantum measurement, a similar situation arises when we consider only the subset of trajectories which end up in a particular final state. While it therefore seems logical to apply an action principle such problems, it is not obvious how to construct an appropriate phase-space description of continuous quantum measurement.

Reference [112] solves this problem by constructing a doubled quantum state space, where the generalized coordinates are the components of the state vector, and the generalized momenta are auxiliary dynamical parameters which impose the conditional measurement backaction. This technique relies on a path integral approach to continuous quantum measurement[1], which they use to construct a joint probability density function (PDF) for all measurement outcomes and corresponding trajectories connecting an initial and a final state. From this PDF, they construct a stochastic action, which is extremized to find the most probable path. The most likely path is the solution to a set of ordinary (non-stochastic) differential equations, and can be solved using standard techniques.

In the limit of no post-selection, this approach is consistent with the known formulation based on stochastic master equations. However, for certain types of problems, it offers several advantages over using the full SME. This formalism naturally incorporates pre- and post-selection as boundary conditions, and enables the derivation of statistical quantities which may be difficult to calculate using the SME. While the doubled state space increases the computational complexity beyond that of the standard master equation, the equations are still simpler to solve numerically than simulating a large number of SMEs and then post-selecting for the desired final state.

The action principle derived in reference [112] is a general theory, which can be applied to a wide range of quantum systems. For concreteness, here we will derive the action principle for the specific case of the measurements presented in this thesis: we initially prepare a qubit in the state $x_I \approx 1$ and then perform a continuous $z-$measurement at a rate $1/\tau$ while

[1]Path integral approaches to continuous quantum measurement have been previously investigated by a number of different authors, including Mensky [117] and Caves [118], but reference [112] is the first to use them in conjunction with post selection to formulate an action principle

concurrently driving the qubit about the $-y$ axis of the Bloch sphere at a Rabi frequency Ω. In this case, the y component of the Bloch vector will always be zero, and the qubit state is described by the coordinate[2] $q = (x, z)$.

We consider a sequence of unitless measurement results $\{r_j\} = \{r_0, r_1, ..., r_{n-1}\}$ where $r_j = 2V_j/\Delta V$, at times $\{t_j = j\Delta t\}$ for $j = 0, 1, ..., n-1$, and its corresponding set of qubit states $\{q_j\}$. Here, V_j is defined as the measurement signal integrated from time t_j to $t_j + \Delta t$. We write a joint probability density function of all measurement results $\{r_j\}$, the quantum states $\{q_j\}$, and the chosen final state q_F, conditioned on the initial state q_I as

$$P(\{q_j\}\{r_j\}, q_F|q_I) = \delta^2(q_0 - q_I)\delta^2(q_n - q_F)\prod_{j=0}^{n-1} P(q_{j+1}, r_j|q_j), \tag{7.1}$$

where the delta functions impose the initial and final boundary conditions (the delta functions are squared because here q has two components, x and z). The conditional probabilities on the right hand side of equation (7.1) can be factored into products of two terms, $P(q_{j+1}, r_j|q_j) = P(q_{j+1}|q_j, r_j)P(r_j|q_j)$, which results in

$$P(\{q_j\}\{r_j\}, q_F|q_I) = \delta^2(q_0 - q_I)\delta^2(q_n - q_F) \times \left(\prod_{j=0}^{n-1} P(q_{j+1}|q_j, r_j)P(r_j|q_j)\right). \tag{7.2}$$

Our goal is to find an appropriate action S such that we can express $P(\{q_j\}\{r_j\}, q_F|q_I) \propto \int \mathcal{D}p\ \mathrm{e}^S$, where $\mathcal{D}p$ is an integral measure or 'path differential', which we will define below. The first step is to express the probability density functions $P(q_{j+1}|q_j, r_j)$ and $P(r_j|q_j)$ explicitly.

Because the qubit state at any time t_{j+1} is updated deterministically from q_j and r_j, the probability density function $P(q_{j+1}|q_j, r_j)$ is a delta function which imposes the state update equations. We introduce the shorthand $\mathcal{L}[q_j, r_j]$ to describe the state update procedure, such that

$$q_{j+1} = q_j + \mathcal{L}[q_j, r_j]\Delta t. \tag{7.3}$$

Then, we can express the state update as the delta function

$$P(q_{j+1}|q_j, r_j) = \delta^2(q_{k+1} - q_k - \mathcal{L}[q_j, r_j]\Delta t). \tag{7.4}$$

We can express this delta function in a Fourier-transformed form by introducing a set of dummy variables $\{p_j\} = \{(p_j^x, p_j^z)\}$ for $j = -1, 0, ...n$, which gives us

[2]Recall from section 6.2.1 that a general single qubit state can be represented as $q = (x, y, z)$, where $x \equiv \mathrm{tr}[\rho\hat{\sigma}_x]$, $y \equiv \mathrm{tr}[\rho\hat{\sigma}_y]$, and $z \equiv \mathrm{tr}[\rho\hat{\sigma}_z]$.

$$P(q_{j+1}|q_j, r_j) = \left(\frac{1}{2\pi i}\right)^2 \int_{-\infty}^{\infty} d^2p_j \exp[-ip_j(q_{j+1} - q_j - \mathcal{L}[q_j, r_j]\Delta t)]. \tag{7.5}$$

We now turn our attention to the conditional distribution $P(r_j|q_j)$ which describes the probability of obtaining a measurement result r_j in a time interval Δt for a given q_j and takes the form

$$P(r_j|q_j) = \sqrt{\frac{\Delta t}{2\pi\tau}} \left(\frac{1+z_j}{2} \mathrm{e}^{-\frac{\Delta t}{2\tau}(r_j-1)^2} + \frac{1-z_j}{2} \mathrm{e}^{-\frac{\Delta t}{2\tau}(r_j+1)^2}\right). \tag{7.6}$$

Note that in this case the exponent is first order in Δt. We arrive at the full probability distribution function $P(\{q_j\}\{r_j\}, q_F|q_I)$ by combining equations (7.2), (7.5), and (7.6). The terms $\prod_j d^2p_j$ constitute the path differential $\mathcal{D}p$. The full probability distribution function is a product of exponentials, and by converting to a sum of exponents we arrive at the following expression for the action S:

$$S = -p_{-1} \cdot (q_0 - q_I) - p_n \cdot (q_n - q_F) + \sum_{j=0}^{n-1} \{-p_j \cdot (q_{j+1} - \mathcal{L}[q_j, r_j]) + \ln P(r_j|q_j)\} \tag{7.7}$$

$$\begin{aligned} &= -B + \int_0^T \mathrm{d}t[-p_x\dot{x} - p_z\dot{z} + p_x(-\gamma x + \Omega z - xzr/\tau) \\ &\quad + p_z(-\Omega x + (1-z^2)r/\tau) - (r^2 - 2rz + 1)/2\tau] \end{aligned} \tag{7.8}$$

where we have introduced B as a short-hand for the first two terms in equation (7.7). We note that, in equation (7.8), we have replaced Δt with $\mathrm{d}t$ and taken the time-continuum limit $\mathrm{d}t \to 0$ and written the action explicitly for our qubit measurement case with the state update equations (6.12) and (6.13). We have also used shortened notation for the variables, for example $x = x(t) \equiv \lim_{\mathrm{d}t\to 0}\{x_0, x_1, ..., x_n\}$. To obtain the most likely path, we then extremize the action in equation (7.8) over the variables q, p, and r and obtain the following ordinary differential equations[3]:

$$\dot{x} = -\gamma x + \Omega z - xzr/\tau \tag{7.9}$$

$$\dot{z} = \Omega x + (1-z^2)r/\tau \tag{7.10}$$

$$\dot{p_x} = +\gamma p_x + \Omega p_z + p_x zr/\tau \tag{7.11}$$

$$\dot{p_z} = -\Omega p_x + (p_x x + 2p_z z - 1)r/\tau, \tag{7.12}$$

[3]We leave the derivations of these equations as an exercise for the mathematically inclined reader. If you are interested in discussing these derivations in detail, then talk to Vinay.

where $r = z + p_z(1-z^2) - p_x xz$ and the forced boundary conditions are $x(t=0) = x_I, z(t=0) = z_I, x(t=T) = x_F, z(t=T) = z_F$. Here, (x, z) is the optimal path and r is the optimal dimensionless measurement signal connecting the states q_I and q_F in a time T. The solution to these nonlinear equations admits four constants of motion, which permits the imposition of both initial (x_I, z_I) and final (x_F, z_F) boundary conditions.

The equations have a simple analytic solution $(\bar{x}, \bar{z})$ for $\Omega = 0$. For an initial state $(x_I = 1, z_I = 0)$, the solution is given by

$$\bar{x}(t) = \mathrm{e}^{-\gamma t}\mathrm{sech}(\bar{r}t/\tau) \tag{7.13}$$

$$\bar{z}(t) = \tanh(\bar{r}t/\tau) \tag{7.14}$$

$$\bar{r}(t) = (\tau/T)\tanh^{-1}(z_F). \tag{7.15}$$

For the driven case, where $\Omega \neq 0$, our theory collaborators in the group of Andrew Jordan solve the equations numerically using a shooting method.

7.2.2 Most likely time

In addition to finding the most likely path taken between pre- and post-selected states in a fixed time, a complementary problem in quantum control is to find the optimal waiting time between an initial state and a target final state. Here, we consider the undriven case, where $\Omega = 0$. While the most likely time taken to travel between two states can be derived from a path integral approach, we give a simpler derivation here based on the probability distribution of the time-averaged measurement signal $V_m = (1/n)\sum_{j=0}^{n-1} V_j$.

In the undriven case, the z coordinate of the qubit on the Bloch sphere at any time T is solely determined by V_m. Given an initial coordinate (z_I), we would like to calculate the distribution of the final z coordinate (z_F) at any time T. This distribution, $P(z_F, z_I)$, can be derived from probability distribution for the measurement outcome $P(V_m|z_I)$, given by:

$$P(V_m \,|\, z_I) = P(V_m \,|\, 0)\frac{1+z_I}{2} + P(V_m \,|\, 1)\frac{1-z_I}{2} \tag{7.16}$$

$$= \sqrt{\frac{1}{2\pi\sigma^2}}\left(\frac{1+z_I}{2}e^{-\frac{1}{2\sigma^2}(V_m - \Delta V/2)^2} + \frac{1-z_I}{2}e^{-\frac{1}{2\sigma^2}(V_m+\Delta V/2)^2}\right), \tag{7.17}$$

where the variance of the voltage signal measured in a time Δt is $\sigma^2 = \Delta V^2\,\tau/4\,\Delta t$. We change variables from V_m to the final z-component z_F using the expression $V_m = (\tau\Delta V/2T)[\tanh^{-1}(z_F) - \tanh^{-1}(z_I)]$. We obtain the differential measure $\mathrm{d}V_m = (\tau/2T)[\Delta V/(1-z_F^2)]\mathrm{d}z_F$. The probability density function of z_F given z_I can be computed via a relation $P(V_m \,|\, z_I)\mathrm{d}V_m = P(z_F \,|\, z_I)\mathrm{d}z_F$,

$$P(z_F|z_I) = \frac{\sqrt{\frac{\tau}{2\pi T}}}{(1-z_F^2)}\exp\left\{-\frac{T}{2\tau}(\bar{r}^2+1) + \frac{1}{2}\ln\left(\frac{1-z_I^2}{1-z_F^2}\right)\right\}, \tag{7.18}$$

where $\bar{r} \equiv (\tau/T)\tanh^{-1}[(z_F - z_I/1 - z_I z_F)] = (\tau/T)[\tanh^{-1}(z_F) - \tanh^{-1}(z_I)]$. For the case where the initial state is $x = +1$, $(z_I = 0)$, the probability density function simplifies to

$$P(z_F|z_I = 0) = \frac{\sqrt{\frac{\tau}{2\pi T}}}{(1 - z_F^2)^{\frac{3}{2}}} \exp\left\{-\frac{T}{2\tau} - \frac{\tau}{2T}(\tanh^{-1} z_F)^2\right\}. \tag{7.19}$$

We then compute the most likely time T_{opt} for which $P(z_F|z_I)$ is largest for the fixed values of z_I and z_F. By maximizing the probability function $P(z_F|z_I)$ with respect to T, we obtain,

$$T_{\text{opt}} = \tau\left(\frac{\sqrt{1 + 4\,\bar{\gamma}^2} - 1}{2}\right), \tag{7.20}$$

where $\bar{\gamma} \equiv \tanh^{-1}\left(\frac{z_F - z_I}{1 - z_I z_F}\right)$.

7.2.3 Schrödinger bridges

Aside from qubit measurements, most-likely path problems occur in a variety of other physical settings, and are typically treated through the formalism of rare-event statistics. For example, reference [115] considers the statistics of charge transport between nodes of a stochastic network. Here, we describe a similar problem in classical probability theory which investigates the most likely behavior of large deviations from an empirical distribution, and is of historical interest for its connection to alternative interpretations of quantum mechanics. We will introduce the problem as originally posed by Schrödinger in 1931 [119] and discussed in reference [120].

A cloud of N Brownian particles is observed at an initial time t_0 to have a density distribution $\rho_0(\mathbf{r})d\mathbf{r}$, as illustrated for 2 dimensions in figure 7.1. The particles undergo diffusive evolution, and at a later time t_1 the system is observed to have a a distribution $\rho_1(\mathbf{r})d\mathbf{r}$. Suppose that $\rho_1(\mathbf{r})$ differs considerably from the distribution expected by the law of large numbers. Then, the particles must have been transported in an unlikely way. But out of all the possible but unlikely trajectories which transform ρ_0 into ρ_1, which is the *most likely*? The solution to this problem is know as the Schrödinger bridge. As initially guessed by Schrödinger and later proved by Fortet [121], there exists a unique solution to the general Schrödinger bridge problem.

Schrödinger also noted that the bridge from ρ_1 to ρ_0 is simply the time reversal of the bridge from ρ_0 to ρ_1, or in his own words[4]: "Abnormal states have arisen with high probability by an exact time reversal of a proper diffusion process." Furthermore, Schrödinger suspected that there was some connection between the time-reversal of the bridge and time reversal in quantum mechanics, which led him to title his paper "On the reversal of natural laws"

[4]Based on a translation quoted in reference [120].

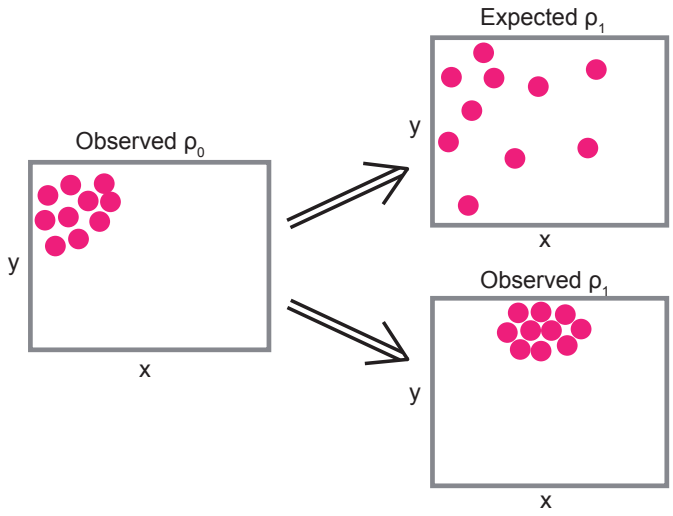

Figure 7.1: Schrödinger bridge in two dimensions. A cloud of Brownian particles is observed at time t_0 in an initial configuration $\rho_0(x, y)$. After diffusive evolution, it is observed at a later time t_1 to be in a configuration $\rho_1(x, y)$ which differs considerably from the distribution expected by the law of large numbers.

and has in inspired later attempts to reformulate quantum theory in terms of stochastic mechanics [122, 123]. In fact, a number of papers have investigated the Schrödinger bridge problem for quantum systems [120, 124, 125]. For both the classical and the quantum case, the solution to Schrödinger's bridge is typically approached by variational methods[120].

Regardless of the role of stochastic mechanics as a potential interpretation of quantum mechanics, the quantum Schrödinger bridge may provide a valuable theoretical tool for steering a quantum system from one state into another, as discussed in reference [126]. Advances in quantum state steering could find application in a wide range of research fields including the control of molecular dynamics [127] and quantum error correction.

7.2.4 Pre- and Post-selection

In order to experimentally investigate the most likely path between an initial and a final state, we use the techniques of pre- and post-selection to create from an ensemble of experiments a sub-ensemble in which a system begins and ends in a particular state. We have already introduced in chapter 5 the idea of pre-selection for heralded state preparation, and post-selection to eliminate experimental iterations where the qubit state left the $0-1$ subspace.

In the experiments presented in this chapter we use projective pre-selection (a herald) followed by a $\pi/2$ pulse to prepare the qubit in initial state ($x_I \approx 1, z_I \approx 0$), just as we did in the previous chapter. In fact, here we use the same datasets which were used to

tomographically validate the trajectories shown in figure 6.7. After computing the quantum trajectory for each experimental iteration, we can then choose to post-select on the subset of trajectories which end near a chosen final state (x_F, z_F). Note that unlike a projective postselection, which separates the full ensemble into a discrete number of sub-ensembles (two for a qubit), here we can chose to post-select on any qubit state. The number of possible post-selections is only set by the size of our post-selection window around (x_F, z_F), which is not a fundamental limit, but instead depends on the total number of experimental iterations, the probability of reaching a given final state, and the desired size of the post-selected ensemble.

7.3 Distributions of trajectories

Here, we examine distributions of quantum trajectories. We perform 10^5 iterations of each experiment, with a measurement duration of $1.424\,\mu$s. For each experiment, we construct the quantum trajectory by finding x and z for every time step, using equations (6.12) and (6.13). Qualitatively, distributions of trajectories allow us to visualize complex qubit measurement dynamics. We also analyze distributions of pre- and post-selected trajectories, from which we can extract statistical properties such as the most probable path through state space connecting initial and final states.

7.3.1 Un-driven trajectories

In figure 7.2a,b,c we show the full ensemble of quantum trajectories for $\Omega = 0$, starting from the state $(x_I = 0.97, Z_I = 0)$. A few example trajectories are plotted in color. The greyscale histograms represent the values of x and z at each time point, binning with a bin size of 0.02. For visual clarity, the greyscale shading is normalized independently for each time point, such that the most frequent value is 1 at each time point. Here, $\tau = 1.25\,\mu$s and $\Gamma = 0.94 \times 10^6\,\mathrm{s}^{-1}$.

Note that by by end of the $1.424\,\mu$s measurement, most but not all of the trajectories have been projected into the ground or excited state. For very short times essentially none of the trajectories are projected, illustrating the fact that it takes time for the qubit state to diffuse. Also, notice that the x component of the trajectories always remains non-negative, as expected for an undriven z-measurement. The upper limit on x decays exponentially in time a rate $\gamma = \Gamma - 1/2\tau$, as a result of the measurement inefficiency.

In panels (d) and (e), we show trajectories post-selected on a final state $z_F = -0.85 \pm 0.03$. Here x_F is completely specified by z_F, because $\Omega = 0$. The most likely paths in z and x predicted from the theory equations (7.13) and (7.14) are displayed as yellow dashed curves. The experimentally determined most likely path is plotted in magenta, with the standard deviations shown as magenta bands.

To find the experimental most likely path, we first define the closeness of any two trajectories (labeled a and b) as a time-average of the Euclidean distance:

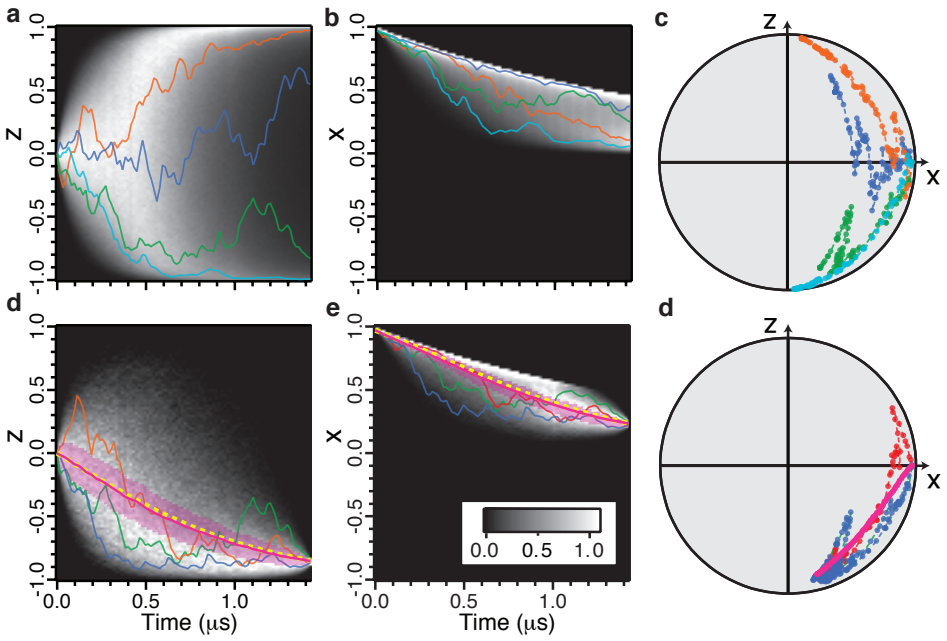

Figure 7.2: Distribution of un-driven trajectories. Here, the measurement duration $T = 1.424\,\mu$s, $\tau = 1.25\,\mu$s, and $\Gamma = 0.94 \times 10^6\,\mathrm{s}^{-1}$. a), b) Greyscale histograms of all measured z (a) and x (b) trajectories, beginning from state $(x_I = 0.97, z_I = 0)$. Example trajectories are shown in color, and plotted in panel (c) on the x-y plane of the Bloch sphere. d), e) Histograms of the sub-ensemble of z (d) and x (e) trajectories conforming to the boundary condition $z_F = -0.85 \pm 0.03$. The experimentally determined most likely trajectories are shown as magenta curves, with their standard deviations shown by the magenta bands. The most likely paths in z and x predicted from the theory are shown as yellow dashed curves, and other example trajectories are shown in different colors and plotted on the x-y plane of the Bloch sphere in panel (f).

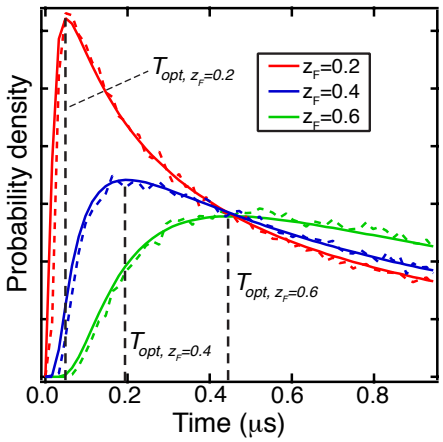

Figure 7.3: Distributions of transit times between initial and final states, for $\Omega = 0$. The theoretical probability density functions $P(z_F|z_I = 0)$ plotted as functions of time T (solid curves) along with experimental data (dashed curves) with $\tau = 1.25\,\mu$s. The red, green, and blue curves correspond to different final states $z_F = 0.2, 0.4$, and 0.6, respectively. The optimized times T_{opt} for the three cases are shown as the vertical black dashed lines.

$$D \equiv (1/n) \sum_{j=0}^{n-1} \sqrt{(x_a(t_j) - x_b(t_j))^2 + (z_a(t_j) - z_b(t_j))^2} \tag{7.21}$$

where $t_j = j\Delta t$. For a qubit, the Euclidean distance is proportional to the trace distance between two quantum states $\rho_a(t_j)$ and $\rho_b(t_j)$. We compute the distance D between all possible pairs of trajectories within the post-selected sub-ensemble. We then search for the N trajectories that have a minimum average distance to all others, and average them to obtain an estimate of the most likely path. In our experiments, there are a total approximately 10^3 trajectories in the sub-ensemble, and we choose $N \sim 10^2$, resulting in a smooth estimate of the most likely path which is still very different from the total (sub-ensemble) averaged trajectory. The magenta bands represents the standard deviations in x and z of the group of N trajectories for each time step. As shown in figure 7.2, the experimentally determined most likely path closely approximates the theoretical most likely path given by equations (7.13), (7.14), and (7.15). We expect that the experimental curves will converge to smooth theory curves in the limit of an infinite ensemble trajectories.

In figure 7.3, we show the distributions $P(z_F|z_I = 0)$ as a function of the post-selection time T for $z_F = 0.2, 0.4, 0.6$. The experimental distributions show good agreement with theoretical predictions of equation (7.19), verifying that we can indeed use equation (7.20) to predict the optimal measurement time connecting arbitrary initial and final states.

7.3.2 Driven trajectories

In the top panels of figure 7.4 we show the full ensemble of driven trajectories for $\Omega/2\pi = 1.08$ MHz, starting from the state $(x_I = 0.88, Z_I = 0)$. A few example trajectories are plotted in color. The greyscale histograms are scaled just as in the previous section. Here, $\tau = 315\,\mu$s, $\Gamma = 3.85 \times 10^6\,\mathrm{s}^{-1}$. The averaged is shown in black and matches the expected ensemble Rabi oscillations.

Note that initially, due the Rabi drive, the trajectories are preferentially projected toward the excited state (z = -1). After half of a Rabi period, they are preferentially projected toward the ground state. As you can see from the sample trajectories and from the overall distribution, here we are in an intermediate parameter regime between quantum jumps and diffusive trajectories. If we were to further increase the measurement rate such that $\tau \sim \Delta t \ll 1/\Omega$, then the Rabi drive would induce random quantum jumps between the eigenstates of z, as discussed in section 2.3.2.

In the lower panels of figure 7.4, we post-select on the final state $(z_F = 0.7, x_F = -0.29)$, with a post-selection window of ± 0.08. This final state was chosen to give us a large enough post-selected ensemble to perform a meaningful statistical analysis. As the time duration between the boundary conditions is increased from $t_1 = 0.464\,\mu$s to $t_2 = 0.944\,\mu$s and then to $t_3 = 1.424\,\mu$s, the most likely trajectory connecting the initial and final states changes drastically but is well described by the theory. For short post-selection times, the trajectories are tightly clustered around the most likely trajectory. For longer post-selection times there greater spread mid-trajectory, but the trajectories converge at the end to conform to the final boundary condition. For t_3, the most likely trajectory begins to resemble a Rabi oscillation, but with greater coherence than the decaying oscillations of the full ensemble.

7.3.3 Weak functions

In addition to the most likely path, the solution of equations (7.9)-(7.12) also gives us r, which describes the optimal detector response for moving the quantum system to the target state after a given time. We compare these optimal signals to the conditioned average detector signals in figure 7.5, which we refer to as a 'weak functions'. The post-selection allows the conditioned average detector signal r to exceed the usual range of $[-1, 1]$ for z. This is analogous to weak values [128, 129, 130] and time-continuous generalizations [64, 131] which can also lie outside their eigenvalue range. While experiments on weak values usually consider an additional projective post-selection measurement, here we simply use the conditioned state after a continuous weak measurement for our post-selection.

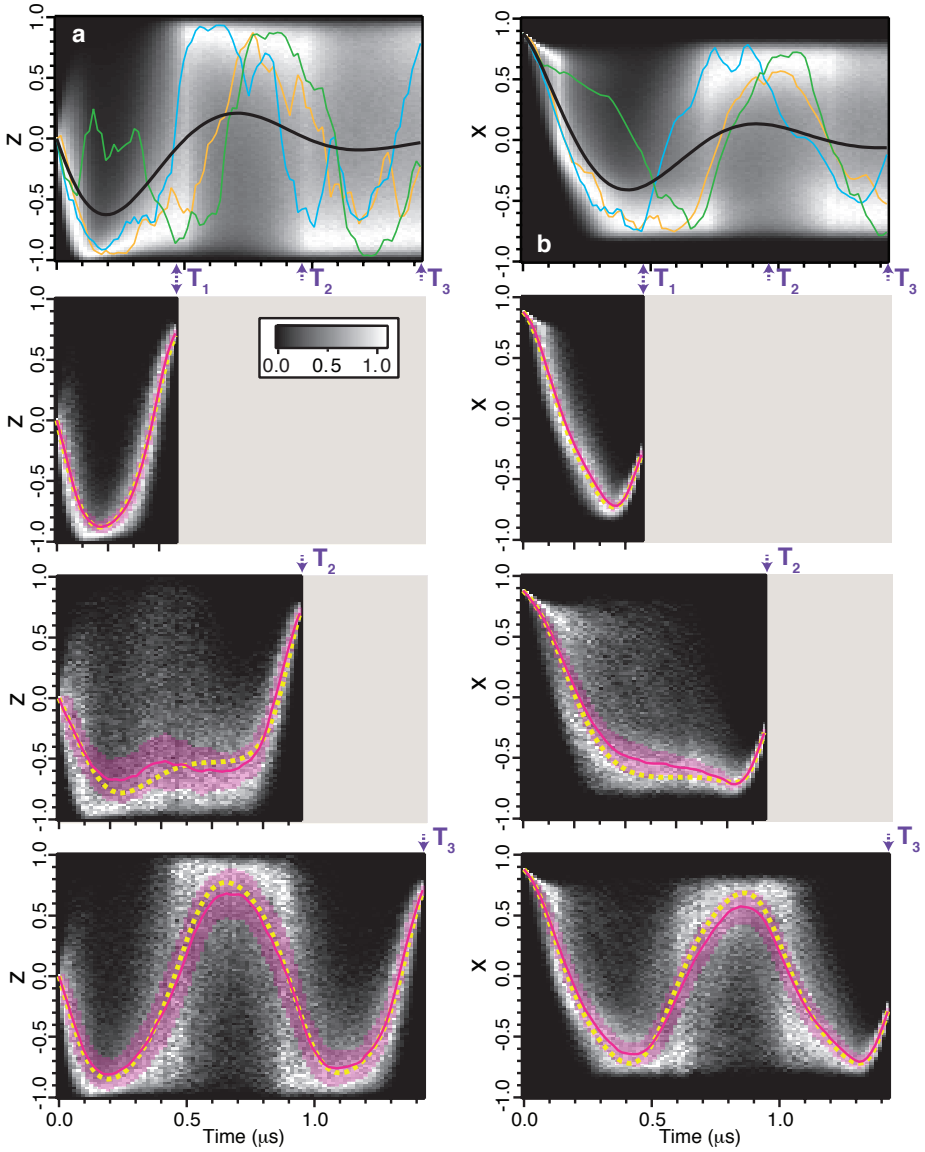

Figure 7.4: Distribution of driven trajectories. Here, $\tau = 315$ ns, $\Gamma = 3.85 \times 10^6\,\mathrm{s}^{-1}$, and $\Omega/2\pi = 1.08$ MHz, and the measurements begin at state $(x_I = 0.88, z_I = 0)$. a), b) (top) Greyscale histograms for z (a) and x (b) with example trajectories shown in color and with the average trajectory shown in black. In lower panels of (a) and (b), we post-select on the final state $(z_F = 0.7, x_F = -0.29)$, with a post-selection window of ± 0.08, at final times $T_1 = 0.464\,\mu$s, $T_2 = 0.944\,\mu$s, and$T_3 = 1.424\,\mu$s. Solid magenta curves show the experimental most likely path, and the yellow dashed curves are numerical solutions to (7.9)-(7.12). The standard deviations of the experimentally determined most likely paths are shown by magenta bands.

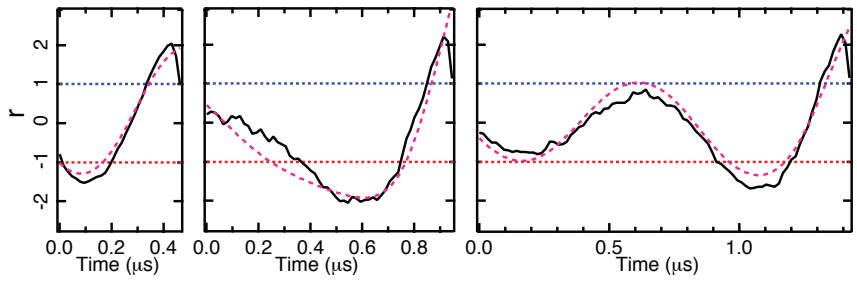

Figure 7.5: Weak functions. The optimal detector signals (r; dashed lines) connecting an initial state ($x_I = 0.88, z_I = 0$) to a final state ($z_F = 0.8, x_F = -0.29$), at times $0.464\,\mu$s (left) $0.944\,\mu$s (middle) and $1.424\,\mu$s (right) for $\tau = 315$ ns, $\Gamma = 3.85\times10^6\,\mathrm{s}^{-1}$, and $\Omega/2\pi = 1.08$ MHz are compared against the conditional average signal (weak functions; black lines). Adapted from reference [28].

Chapter 8

Quantum efficiency and squeezing

In this chapter, we investigate the quantum efficiency of our measurement chain. First, we characterize the different types of inefficiency in our experiments and discuss their possible causes. We then turn our attention to one particular mechanism which may be limiting our quantum efficiency– imperfect squeezing. We discuss different methods of characterizing squeezing at microwave frequencies, focusing on a method that we developed which uses qubit decay rates to detect squeezing.

8.1 Quantum efficiency budget

As we discussed in section 5.4.7, in the process of setting up a quantum trajectories experiment we perform a sensitive calibration of the quantum efficiency. In reference [27], we achieved a measurement efficiency of $\eta_m = 0.49$. The inefficiency due to environmental dephasing, $\eta_{env} = (1+\kappa/8\chi^2\bar{n}T_2^*)^{-1}$ depends on the measurement rate (which is proportional to $\bar{n}$), as does the total quantum efficiency $\eta = \eta_m\eta_{env}$. For the trajectories shown in reference [27], where $\bar{n} = 0.4$, $\eta_{env} = 0.91$ and thus $\eta = 0.45$.

For a given set of qubit parameters, η_{env} is well understood and can be improved by increasing the measurement rate. For example, in reference [28], with $\tau = 315$ ns ($\bar{n} = 1.8$), $\eta_{env} = 0.98$. What limits η_m is a more subtle question, and will be the focus of this chapter. As discussed in chapter 2, we can factor η_m into separate collection and amplification efficiencies

$$\eta_m = \eta_{col}\eta_{amp} = 0.49, \tag{8.1}$$

where $\eta_m = 0.49$ from reference [27] represents the best quantum efficiency measured in our lab to date. This number is comparable to some of the highest reported values from other labs, including reference [132] which achieved $\eta_m = 0.35$ for measuring photon number states with Rydberg atoms and references [63] and [133], which measured superconducting qubits with $\eta_m = 0.55$ and 0.67, respectively, using a Josephson parametric converter. Reference

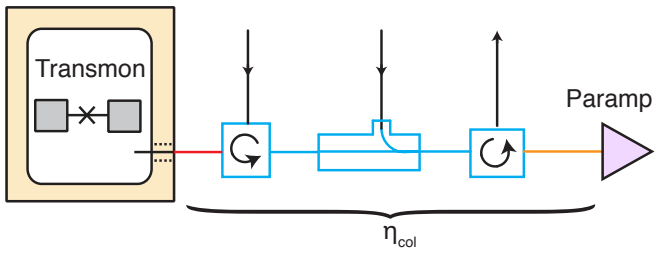

Figure 8.1: Collection efficiency. We calibrated the losses through various branches of the measurement chain to be 0.08 dB (red), 1.13 dB (blue), and 0.3 dB (orange).

[47] achieved a η_m = 0.85-0.94 for an ancilla-based qubit measurement, at the expense of increased measurement complexity which would make it difficult to implement continuous measurements.

To better understand what limits η_m in our setup, we performed a set of calibration experiments to determine the collection efficiency η_{col}, which is set by losses in the microwave components between the cavity and the paramp. We used a cryogenic transfer switch to compare transmission loss through various branches of the measurement setup against a calibrated through. We measured 0.08 dB of loss in the cable from the cavity to the first circulator, 1.13 dB of loss through two circulators and a directional coupler (and any cables or barrels in between), and 0.3 dB of loss in the cable leading to the paramp, as depicted in figure 8.1. These add up to a total of 1.41 dB of loss, which gives us a collection efficiency of $\eta_{col} = 0.72$. Therefore, $\eta_{amp} = \eta_m/\eta_{col} = 0.68$. The collection efficiency could be improved slightly by decreasing cable lengths and carefully testing cables, but it will be difficult to improve by much, because each circulator is known to have $0.3 - 0.5$ dB of loss. While in principle it could be possible to use only one circulator, it is difficult to do in practice because proper paramp operation requires sufficient isolation from the the cavity.

For the remainder of this chapter, we will focus on what may be limiting η_{amp} to 0.68. Ideal parametric amplification should be quantum-limited[30], allowing for $\eta_{amp} = 1$. There are a number of possibilities for what may be reducing the amplification efficiency below this ideal value. One possibility is that the gain may be too low to overcome the added noise of the HEMT. While 20 dB of gain in the signal mode should be sufficient, in practice, when we were working at 20 dB of gain we found that we needed to optimize the HEMT bias settings to achieve the best quantum efficiency, which means that we were working with barely enough gain to overcome the HEMT noise. However, for the cascaded paramp setup in reference [27] where we achieved our highest efficiency of $\eta_m = 0.49$, we were operating with 25 dB of combined paramp gain, so we should be able to rule out insufficient gain as a potential limiting factor. Another possible problem could be insufficient paramp bandwidth, as discussed in section 5.2.3, but once again we can rule out this concern for the experiments

in reference [27]. Yet another possibility is loss in the paramp. For our paramps, the losses are dominated by the dielectric losses in the SiN_x capacitors, limiting the quality factor to $Q \sim 5000$ at low powers (the quality factor improves at higher powers due to the saturation of loss from two-level system defects). For a paramp with $Q = 10$, the effective Q increases by $\sqrt{G}$ to 100 at 20 dB of gain, setting a lower limit on the contribution of paramp loss to η_{amp} at 0.98. Furthermore, it is likely that the paramp pump will saturate this loss, reducing its effect on the quantum efficiency. A additional possibility is that the paramp could be hotter than the 20 mK plate that it's attached to, in which case it would add thermal noise to a measurement signal. While we didn't directly test this hypothesis, and therefore cannot rule it out, we do know that our qubit was well-themalized, and we might expect that the paramp was as well.

However, there is another possible explanation for our limited amplifier efficiency which is more fundamentally related to the paramp device physics. In order to perform ideal phase-sensitive amplification, we also need to achieve ideal squeezing of the unmeasured quadrature. This is essential, because the fluctuations of the unmeasured quadrature are correlated with the qubit state, as discussed in section 6.1. In the following sections, we will discuss theoretical limitations of the amount of squeezing produced by our paramps and our efforts to characterize squeezing.

8.2 Paramp nonlinearity and squeezing

In this section, we outline recent analysis from our theory collaborators in the group of Alexandre Blais [93]. As this is still a work in progress, rather than presenting their results, we will simply provide a qualitative picture the connection between paramp nonlinearity and squeezing.

Our collaborators have performed numerical simulations of the strongly driven Kerr Hamiltonian introduced in section 4.2.2, finding that the amount of squeezing that a paramp can generate depends on the device nonlinearity, which is expressed as a dimensionless parameter $K/(2\kappa)$, where K is the Kerr nonlinearity and κ is the paramp damping rate (i.e. the bandwidth in the linear regime).

They find, as expected [91], the quadric driven Kerr Hamiltonian $H_{sys}^{(2)}$ defined in equation (4.16) generates ideal squeezing of the paramp output field, such that[1] $M = \sqrt{N(N+1)}$. However, when the higher order terms of the quartic Hamiltonian $H_{sys}^{(4)}$ defined in equation (4.15) are included, the squeezing is no longer ideal, and $M < \sqrt{N(N+1)}$. Preliminary analysis of the internal cavity field suggests that that the squeezing degrades as the nonlinearity increases, and that nearly ideal squeezing can be achieved when $K/(2\kappa) < 10^{-4}$. Interestingly, for a typical lumped element Josephson parametric amplifier with negligible geometric inductance, a single SQUID, and direct coupling to the launch (no coupling capac-

[1]Recall that the squeezing parameters N and M describe the amplitude and the asymmetry of the quantum fluctuations and were defined in section 4.1.2 in terms of the autocorrelations of the creation and annihilation operators.

itor), $K/(2\kappa) = (\pi/2)(R/R_q)$ [134], where R is the effective shunt resistance and $R_q = h/e^2$ is the quantum of resistance. Therefore the dimensionless nonlinearity is independent of κ and the SQUID critical current. For such a device with $R = 50\,\Omega$, the nonlinearity $K/(2\kappa) = 4 \times 10^{-3}$.

The nonlinearity can be reduced by adding linear inductance in series with the SQUID, thereby reducing the junction participation ratio, $p = L_J/L_{tot}$, where L_J is the inductance of the SQUID. However, recall from section 4.3 that we cannot reduce p without sacrificing bandwidth, as $Qp \gtrsim 5$ is required for stable paramp performance. It is possible to reduce the nonlinearity without sacrificing bandwidth by instead replacing the single SQUID with a series array of N approximately identical SQUIDs (each with a critical current N times as large as the for the original single-squid device). The contribution of each SQUID to the nonlinear term in the circuit equation scales as the cube of the voltage across that SQUID, and thus the total nonlinearity, found by summing over the array, scales as $N/N^3 = 1/N^2$. In reference [135], a paramp with an array of $N = 30$ SQUIDs was able to generate at least 12 dB of squeezing. Currently, we are working to experimentally characterize squeezing as a function of paramp nonlinearity and to test whether squeezing is affecting the quantum efficiency of our measurement setup.

8.3 Reconstructing a squeezed state

In this section, we will introduce the experimental techniques that can be used to characterize squeezed states at microwave frequencies. The state of a single mode of the electromagentic vacuum can be represented in terms of the statistical moments $\langle a^{\dagger n} a^m \rangle$, where for a completely general state m and n range from 0 to infinity. Gaussian states, which include coherent, thermal, and squeezed states and any combination thereof, can be completely specified by the first and second order moments. For Gaussian states, the squeezing parameters N and M describe the second order moments $\langle a^\dagger a \rangle$ and $\langle a^2 \rangle$. Thus, by measuring the second order moments we can characterize the amount of squeezing, an by measuring the higher order moments we can estimate the non-Gaussianity of a squeezed state.

Several moment-based reconstruction techniques have been developed in the superconducting qubit community to characterize squeezed states of an itinerant microwave field. A standard approach in quantum optics to state reconstruction is to directly measure the field using homodyne or heterodyne detection and reconstruct its statistical moments [136]. The challenge of implementing such techniques at microwave frequencies is that small nonclassical signals are easily swamped by noise for the HEMT amplifiers. Three basic techniques have emerged which approach this problem in different ways. The fist technique uses one paramp to generate squeezing, followed by a second paramp which is used to detect the squeezing [137, 138, 139]. This technique is the most straightforward, in principle, but it requires a careful calibration of the losses between various stages of amplification, and is therefore very challenging to implement. A second method, known as the single path method (SPM) involves successive heterodyne (or homodyne) measurements of signal and noise distri-

butions [140] and measures the correlations between the two branches. The third technique, known as the dual path method (DPM) splits the signal into two separate measurement chains with uncorrelated HEMT noise [141]. The SPM and the DPM are compared in reference [142], and for a sufficiently low system noise temperature they are similarly effective. We are currently using the single-path method to investigate squeezing and non-Gaussianity as a function of gain for various paramp designs and pump configurations.

In the remainder of this chapter, we will discuss an alternative technique that we have developed which uses the radiative decay of qubit coherence to characterize a squeezed field. While this technique can only characterize the field up to second order moments, it offers the advantage that the qubit decay rates are directly related to the squeezing parameters N and M by a simple analytic formula. Furthermore, the qubit acts as an "in situ" detector of squeezing, and therefore this technique doesn't require precise characterization of the measurement chain. While qubits have been used in previous experiments to reconstruct localized nonclassical states of light [143, 26], our experiment was the first to use a qubit to detect an itinerant squeezed state.

8.4 Radiative decay of qubit coherence in squeezed vacuum

This section provides an overview of our efforts, originally presented in reference [25], to measure the effect of squeezed vaccum on the radiative decay rates of a qubit.

8.4.1 Historical background

In 1986, Gardiner predicted in a seminal paper [144] that the radiative decay rates of an atom will be drastically modified in the presence of squeezed vacuum. This paper gave rise to a large body of theoretical works investigating the effect of squeezed light on atomic properties. However, to this day, only a handful of experiments have investigated the interaction between squeezed light and matter [145, 92, 146, 25]. Prior to our work, none of these experiments constituted a direct quantitative demonstration of Gardiner's original prediction. For an atom in free space, it is very challenging to squeeze all of the spacial modes which contribute to an atom's radiative decay. For this reason, it is convenient to work with a cavity QED architecture in the Purcell limited regime, where the atom's radiative decay rate is set by decay of the cavity mode. For a microwave cavity and an artificial atom (transmon qubit), the cavity decays into the continuum modes of a coaxial cable, which can be readily controlled. This 1-D architecture has also been crucial to our quantum trajectories experiments, and was in fact proposed by Gardiner in reference [144] as potential method to realize the physical effects of squeezed vacuum.

8.4.2 Gardiner-Bloch equations

We consider a qubit that interacts solely with broadband squeezed vacuum, and is characterized by a single radiative decay rate, γ_1. The transverse decay rates γ_x and γ_y, and the longitudinal decay rate γ_z are given by the Gardiner-Bloch equations [144]:

$$\langle \dot{\sigma}_x \rangle = -\gamma(N + M + 1/2)\langle \sigma_x \rangle \equiv -\gamma_x \langle \sigma_x \rangle, \tag{8.2}$$

$$\langle \dot{\sigma}_y \rangle = -\gamma(N - M + 1/2)\langle \sigma_y \rangle \equiv -\gamma_y \langle \sigma_y \rangle, \tag{8.3}$$

$$\langle \dot{\sigma}_z \rangle = -\gamma(2N + 1)\langle \sigma_z \rangle - \gamma_1 \equiv -\gamma_z \langle \sigma_z \rangle - \gamma_1, \tag{8.4}$$

where we have omitted the symbol ($\hat{}$) from the Pauli operators for clarity.

In the limit of pure vacuum fluctuations (no squeezing), $N = M = 0$, and we recover the optical Bloch equations which describe radiative decay into ordinary vacuum. In the case of squeezed vacuum, where $M > N$, the transverse relaxation rate along the y axis is reduced from its nominal decay rate. In the case of large squeezing, where $M, N \gg 0$, the relaxation rate approaches zero. The qubit decay rates γ_x, γ_y, and γ_z can be measured with a Ramsey sequence, and γ_1 is determined from a T_1 measurement. Therefore, the Gardiner-Bloch equations offer a direct method to characterize M and N. Measuring γ_y to be less than $\gamma_1/2$, ($T_2^* > 2T_1$) constitutes direct evidence of squeezing.

8.4.3 Observation of the squeezed light-matter interaction

A simplified schematic of our experiment is shown in figure 8.2a. We realize an effective qubit using the ground state and lower energy level of a polariton formed by a transmon resonantly coupled the fundamental mode of a superconducting cavity. The transition frequency of the effective qubit is $\omega_q/2\pi = 5.8989$ GHz with $T_1 = 0.65$ μs set by deliberate coupling to the $50\,\Omega$ environment. The squeezed output of the paramp is connected with coaxial cables to the input port of the cavity. In order to ensure that the amplified bandwidth doesn't overlap with the upper polariton state $|+\rangle$, we added coupling capacitors which isolated the paramp resonator from the $50\,\Omega$ environment. This has the result of reducing the paramp decay rate to $\kappa/2\pi = 60$ MHz, from the value of 250 MHz for the directly coupled paramp used in the quantum trajectories experiments discussed in the previous chapters. In hindsight, we have realized that this design choice lead to increased nonlinearity $K/2\kappa$ of $\sim 10^{-2}$, thereby likely degrading squeezing performance.

Note from figure 8.2c that there exists a well-defined phase relationship between the angle ϕ in the X_1-X_2 plane and the qubit state phase on the Bloch sphere. To characterize the effect of squeezed vacuum on the transverse decay of qubit coherence we perform Ramsey measurements with the qubit prepared at different angles along the equator of the Bloch sphere. We apply a $\pi/2$ pulse to prepare the qubit at an angle ϕ along the equator of the Bloch sphere, allow it to evolve for a variable time t, and then preform a second $\pi/2$ pulse. Finally, we measure the qubit in the σ_z basis.

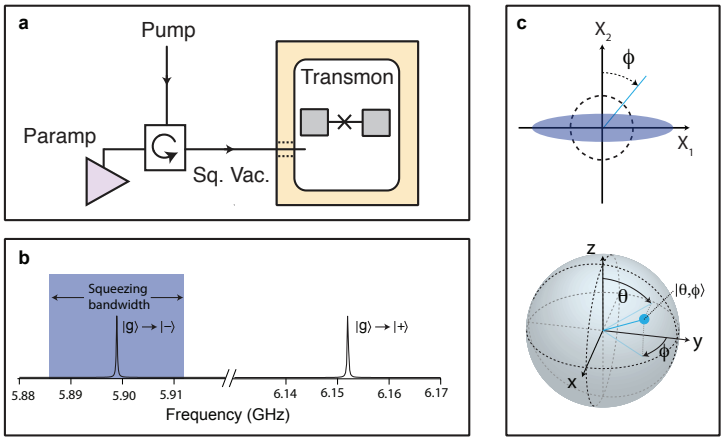

Figure 8.2: Qubit as a squeezing detector. a) A lumped-element Josephson parametric amplifier is used to generate squeezed vacuum that is coupled to the input port of a 3D transmon qubit via a circulator with coaxial cables. b) The resonant strong coupling of the transmon circuit with the 3D cavity results in two polariton states $|+\rangle$ and $|-\rangle$. The bandwidth of the squeezing is centered about the $|g\rangle \rightarrow |-\rangle$ transition frequency and is large compared to the qubit linewidth. c) On resonance, there exists a well-defined phase relationship between an angle ϕ in the X_1-X_2 plane and the qubit state on the Bloch sphere.

.

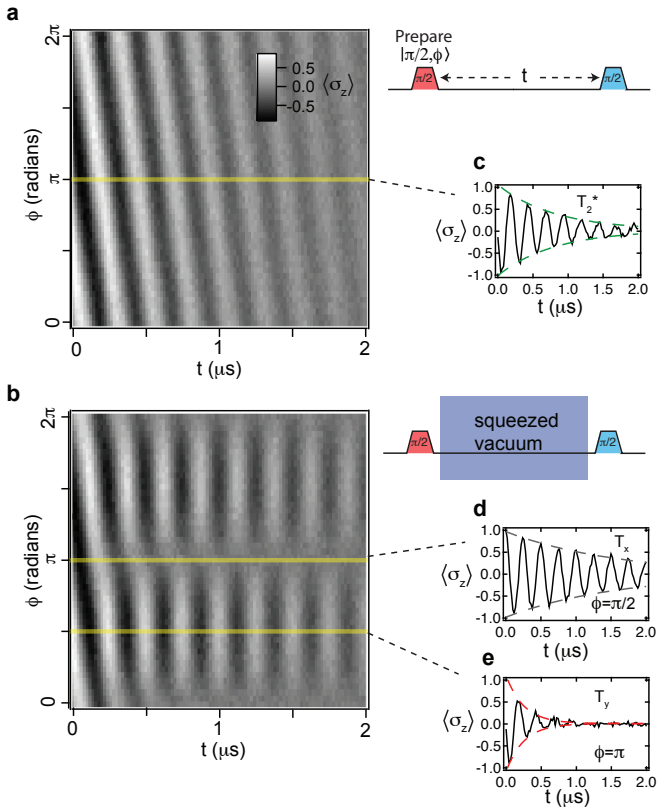

Figure 8.3: Transverse decay in squeezed vacuum. a) Ramsey measurement as a function of angle. The 2D plot displays $\langle\sigma_z\rangle$ as a function of t and ϕ which is characterized by sinusoidal decay with a uniform decay constant T_2^* and phase ϕ. b) The transverse decay into squeezed vacuum was measured by turning the paramp pump on between the qubit pulses. The 2D plot indicates that after rapid decay of coherence along the $\pm\hat{y}$ axes, the resulting coherence along the $\pm\hat{x}$ axes decays with time constant $T_x > T_2^*$. c) This horizontal cut shows a Ramsey measurement for the qubit prepared along the $-\hat{y}$ ($\phi = \pi$) axis with the squeezing off. d,e) These cuts show Ramsey measurements in the presence of squeezed vacuum for the qubit prepared along the $-\hat{y}$ ($\phi = \pi$), and $+\hat{x}$ ($\phi = \pi/2$) axes.

In order to maintain the simple phase relationship between the X_1-X_2 plane and the qubit phase, we need to apply qubit pulses on resonance. However, a standard Ramsey measurement produces no oscillations on resonance, which makes it difficult to accurately fit for the decay rate. However, we can recover oscillations on resonance by modulating the phase of the second $\pi/2$ pulse with respect to the Ramsey delay time t at a frequency ω_{mod}. Figure 8.3a displays $\langle\sigma_z\rangle$ as a function of time and angle with the squeezing turned off; $\langle\sigma_z\rangle$ exhibits exponentially damped, sinusoidal oscillations with a uniform decay time $T_2^* = 1.08$ μs. The T_2^* time is reduced slightly from $2T_1$ due to a small amount of pure dephasing characterized by a time scale $T_\phi = 6.6$ μs. The pure dephasing is not affected by squeezing, and is therefore a major obstacle to the unambiguous detection the squeezed light-matter interaction. Achieving $T_\phi \gg T_2^*$ was a significant technical challenge which we overcame by carefully filtering readout and control lines and shielding our qubit from infrared radiation.

Figure 8.3b displays the results of the Ramsey measurement when the paramp pump was turned on to generate squeezed vacuum for the variable time t between the first and second $\pi/2$ pulses. Here, the amplifier was operated with a power gain of 4 dB. In the presence of squeezed vacuum the transverse decay occurs on two timescales, $T_x \equiv 1/\gamma_x = 1.67$ μs, and $T_y \equiv 1/\gamma_y = 0.28$ μs, that describe the exponential decay of coherence when the qubit was prepared along the $\pm\hat{x}$ and $\pm\hat{y}$ axes respectively. Subtracting the pure dephasing rate from the measured decay rates gives the radiative transverse decay timescales, $\tilde{T}_x = 2.2$ μs, and $\tilde{T}_y = 0.29$ μs. Note that the decay along the $\hat{y}$ axis has increased due to the increased fluctuations in the X_1 quadrature, and the decay along the $\hat{x}$ axis has decreased due to the reduced fluctuations in X_2.

8.4.4 Squeezing as a function of amplifier gain

From the transverse and longitudinal decay rates we can use the Gardiner-Bloch equations to extract N and M, which allow us to tomographically reconstruct, to second order, the Wigner distribution for the itinerant squeezed state generated by the paramp. For the data shown in figure 8.3, we determine $N = 0.88$ and $M = 1.08$, from which the Gaussian variances $\sigma^2_{X_1} = 2(N + M + 1/2)$ and $\sigma^2_{X_2} = 2(N - M + 1/2)$ can be calculated.

To illustrate the power of this detection scheme, we measure the qubit decay rates for different bias conditions of the paramp obtained by changing the power of the pump tones (8.4). The transverse decay rates were measured as depicted in figure 8.3a. As expected, larger paramp gain results in larger amounts of squeezing with an associated increase of T_x and decrease of T_y and T_z. Figure 8.4b displays $M - N$ versus N. The reduction of M from its maximum allowed value, shown as a dashed line, may be attributed to two possible sources; losses in the microwave components between the LJPA and the qubit, and non-ideal performance of the paramp. If we assume that the paramp produces an ideal squeezed state, with $M = \sqrt{N(N+1)}$, then the degradation can be accounted for by an attenuation of the squeezed vacuum by a factor of $\eta = 0.5$, as suggested by the fact that the theory curve for $\eta = 0.5$ shown in figure 8.4b matches with the measured results for samll

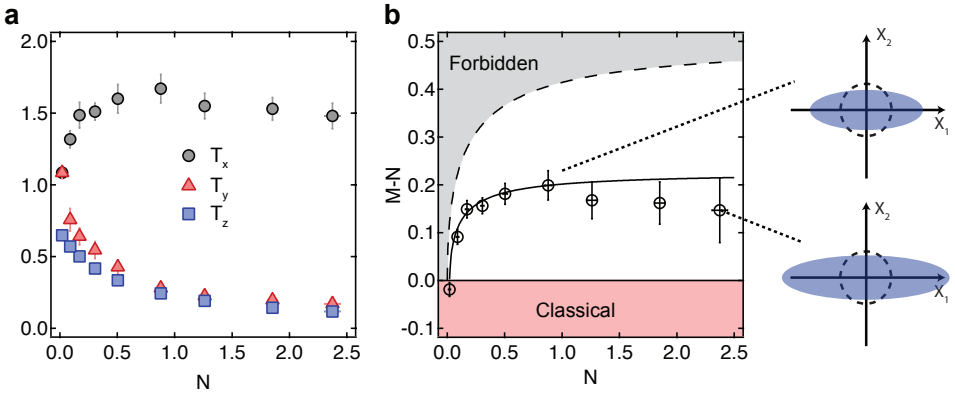

Figure 8.4: Dependence of the qubit decay on paramp bias. a) Measured values for T_x, T_y, T_z for increasing gain of the paramp vs N. Error bars indicate the standard error of the mean based on 10 successive measurements. b) $M-N$ is plotted versus N. The dashed line indicates a minimum uncertainty squeezed state that is expected for ideal squeezing. The solid line is a theory curve indicateing the expected dependence for a quantum efficiency of $\eta = 0.5$. The grey region indicates values of M and N that are forbidden by Heisenberg's uncertainty principle, and the red region indicate values of M and N that correspond to classical states of light. The graphs on the right show the reconstructed squeezed states for $N = 0.88$ (top) and $N = 2.39$ (bottom).

N. Attenuation degrades the squeezed vacuum by absorbing correlated photons thereby making the quadrature fluctuations tend toward the normal vacuum fluctuations. This level of attenuation is roughly consistent with the anticipated insertion loss between the LJPA and the qubit due to the microwave components we used, although based on our trajectories experiments we may expect the that the quantum efficiency could have been as high as 0.72. Regardless, at values of $N > 1$, it appears that the performance of the LJPA may become non-ideal as indicated by the slight reduction of $M-N$ for $N > 1$.

We have demonstrated that, by measuring qubit decay rates at different operating points, a qubit can be used not only as a sensitive detector squeezing, but also to calibrate the losses between the squeezing source and the qubit. While this technique is perhaps more complicated to implement than the SPM and doesn't allow us to characterize higher order moments of the squeezed field, it is an 'in situ' method with high quantum efficiency, and thus doesn't rely on as many sensitive calibrations as other 'ex situ' techniques. Therefore, it still may be a valuable tool, especially if used in conjunction with other methods such as the SPM or DPM.

Chapter 9

Outlook and conclusions

In this thesis, we have explained how to reconstruct individual quantum trajectories of a superconducting qubit and analyzed ensembles of trajectories to gain insight into qubit measurement dynamics. Our experiments provide a clear demonstration of the physics of continuous quantum measurement and help to confirm a detailed understanding of the measurement process in our system. In this chapter, we reflect on the potential role of continuous quantum measurement in future experiments. Already, quantum trajectory experiments have begun to examine fundamental topics such as the roles of time-symmetry in quantum mechanics and the ability to use this symmetry to perform retrodictions[64, 65, 147]. As the field of superconducting qubits is rapidly progressing toward multi-qubit architectures for quantum computing [19, 15] and quantum simulation [21], it is important to consider the role that continuous quantum measurement might play in these settings. Here, we introduce two possible applications of continuous quantum measurement: multi-qubit state estimation and measurement-based quantum feedback.

9.1 Quantum trajectories of multi-qubit systems

9.1.1 Observing the genesis of two-qubit entanglement

While measurement typically destroys quantum coherence, it is in fact possible to design a measurement that projects a multi-qubit system into an entangled state [148, 149]. Such measurement-induced entanglement has been demonstrated in atomic systems [150, 151], nitrogen vacancy centers [152], and for two superconducting qubits occupying a single microwave cavity [10]. Here, we discuss an experiment from our lab where measurement-induced entanglement was generated between two spatially separated qubits [11]. This experiment employed weak measurements, allowing for the observation individual quantum trajectories of the two-qubit system as it is projected into an entangled state.

Here, two transmon qubits couple dispersively to two spatially separated 3D cavities connected by coaxial cables and two circulators. A joint qubit state measurement is performed by applying a measurement tone which sequentially monitors the two cavities. In general,

each cavity will produce a different qubit-state-dependent phase shift $\Delta\phi_1$ and $\Delta\phi_2$, and the four different basis states $|00\rangle, |01\rangle, |10\rangle$, and $|11\rangle$ will correspond to a distinct point in X_1-X_2 plane. However, it is possible to carefully engineer the cavities and the dispersive coupling such that $\Delta\phi_1 = \Delta\phi_2$. In this case, a measurement will be unable to distinguish between the $|01\rangle$ and the $|10\rangle$ states. If the qubit is initially prepared in the equal superposition state $(1/2)(|00\rangle + |01\rangle + |10\rangle + |11\rangle)$, then a measurement outcome corresponding to the odd parity ($|10\rangle$ / $|01\rangle$) subspace will leave the qubit in the Bell state $(1/\sqrt{2})(|01\rangle + |10\rangle)$. In the limit of weak measurement, this experimental setup has been be used to observe quantum trajectories of the two-qubit system as it evolves under measurement from the equal superpositions state toward the state $|00\rangle$, $|11\rangle$, or $(1/\sqrt{2})(|01\rangle + |10\rangle)$ state in an individual experimental iteration.

9.1.2 State estimation

Traditional protocols for quantum state tomography (QST) of a multi-qubit system are based on a series of fixed rotations, each followed by a strong projective measurement. The resource requirements for this type of sequential characterization are prohibitive for large numbers of qubits, and result in errors due to state evolution from the strong measurement pulse and any residual qubit couplings which can not be switched off.

It may be possible to develop tomography protocols based on continuous quantum measurement whose resource requirements scale more favorably with system size than conventional QST protocals. The group of Ivan Deutsch introduced protocols for state estimation based on repeated continuous weak measurements and continuous driving [153, 154]. During driven evolution, a single measurement can probe the system in multiple bases, and it is possible to gain an informationally complete measurement record from a single ensemble of identically prepared qubits. We are currently developing similar techniques which involve simultaneous continuous measurements of multiple superconducting qubits.

In the inverse case, where the qubit state is know but its dynamics are unknown, continuous quantum measurements can be used to estimate an unknown Hamiltonian parameter.

9.2 Quantum feedback

A continuous measurement record, which allows us to track a qubit state in real time, can be incorporated into a feedback loop for a number of applications including state preparation and stabilization and continuous quantum error correction.

9.2.1 State preparation and stabilization

As we discussed in section 5.4.4, projective measurements will purify an initially mixed state, and can be used for state preparation. Without feedback, an arbitrary qubit state can be prepared probabilistically though a herald measurement and unitary qubit rotations[106].

With feedback, the measurement result can be used to control a subsequent qubit rotation, allowing for deterministic fast state-reset protocals [155, 133]. It should also be noted that a number of theory proposals have suggested that quantum feedback could be used to increase the rate of state purification by measurement [156, 157, 158], and that all of these proposals require a quantum efficiency greater than 0.5 [158].

In the case of weak measurement, it is possible prepare arbitrary states by measurement alone, without applying any qubit rotations. Without feedback, state preparation is probabilistic: one simply post-selects an ensemble of trajectories which end up in the desired state. With feedback, it is possible to prepare an arbitrary qubit state deterministically through adaptive measurement, as recently demonstrated with nitrogen vacancy centers[159].

In addition to state preparation, quantum feedback can also be used to stabilize a qubit state or trajectory. In reference [49], we used weak measurements and continuous quantum feedback to stabilize Rabi oscillations of a superconducting qubit. Reference [133] demonstrated that stroboscopic projective measurements and feedback can be used to stabilize an arbitrary trajectory, such as Rabi or Ramsey oscillations. We are currently working to use quantum feedback with the two-qubit setup from reference [11] to stabilize and entangled state.

9.2.2 Continuous quantum error correction

While many quantum error correction (QEC) schemes such as surface codes rely on discrete projective measurements of a syndrome qubit [18, 19, 15], a wide body of QEC proposals are based instead on continuous measurement-based quantum feedback [50]. In these techniques, a single logical qubit is encoded in several physical qubits, and an error syndrome is detected by processing (or 'filtering') a continuous measurement signal. The error signal is used to generate a suitable feedback Hamiltonian in real time.

By tomographically validating individual quantum trajectories, we have demonstrated the ability to correctly 'filter' a measurement signal for one and two qubit systems. A sensible next step is to build a system of several qubits and attempt to correctly filter an error syndrome. Then, the following step would be to feed-back on this error syndrome to realize a single logical qubit whose lifetime exceeds that of its constituent qubits. To realize this goal will require a robust multi-qubit architecture and improvements in the measurement quantum efficiency. Recent experiments from the Martinis group [15] have demonstrated that it is possible to individually measure and control 5-9 qubits in a planar CQED architecture, and as we discussed in chapter 8, we are currently working to improve the quantum efficiency of our qubit measurements. Although there are still formidable challenges to overcome, and while the ultimate utility of measurement based QEC in comparison to other methods remains an open question, it seems that an initial demonstration of measurement based QEC may lie on the horizon.

Bibliography

[1] Einstein, A., Podolsky, B. & Rosen, N. Can Quantum-Mechanical Description of Physical Reality Be Considered Complete? *Physical Review* **47**, 777–780 (1935).

[2] Haroche, S. & Raimond, J. M. *Exploring the Quantum* (Oxford University Press, 2006).

[3] Schrödinger, E. Die gegenwärtige situation in der quantenmechanik. *Naturwissenschaften* **23**, 807–812 (1935).

[4] Everett, H. "Relative State" Formulation of Quantum Mechanics. *Reviews of Modern Physics* **29**, 454–462 (1957).

[5] Griffiths, R. B. Consistent histories and the interpretation of quantum mechanics. *Journal of Statistical Physics* **36**, 219–272 (1984).

[6] Nakamura, Y., Pashkin, Y. A. & Tsai, J. S. Coherent control of macroscopic quantum states in a single-Cooper-pair box. *Nature* **398**, 786–788 (1999).

[7] Steffen, M. *et al.* Measurement of the entanglement of two superconducting qubits via state tomography. *Science* **313**, 1423–5 (2006).

[8] Chow, J. M. *et al.* Detecting highly entangled states with a joint qubit readout. *Physical Review A* **81**, 062325 (2010).

[9] Chow, J. M. *et al.* Simple All-Microwave Entangling Gate for Fixed-Frequency Superconducting Qubits. *Physical Review Letters* **107**, 080502 (2011).

[10] Ristè, D. *et al.* Deterministic entanglement of superconducting qubits by parity measurement and feedback. *Nature* **502**, 350–4 (2013).

[11] Roch, N. *et al.* Observation of Measurement-Induced Entanglement and Quantum Trajectories of Remote Superconducting Qubits. *Physical Review Letters* **112**, 170501 (2014).

[12] Dicarlo, L. *et al.* Preparation and measurement of three-qubit entanglement in a superconducting circuit. *Nature* **467**, 574–8 (2010).

[13] DiCarlo, L. *et al.* Demonstration of two-qubit algorithms with a superconducting quantum processor. *Nature* **460**, 240–4 (2009).

[14] Mariantoni, M. *et al.* Implementing the quantum von Neumann architecture with superconducting circuits. *Science (New York, N.Y.)* **334**, 61–5 (2011).

[15] Barends, R. *et al.* Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature* **508**, 500–3 (2014).

[16] Jeffrey, E. *et al.* Fast Accurate State Measurement with Superconducting Qubits. *Physical Review Letters* **112**, 190504 (2014).

[17] Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge Univ. Press, Cambridge, 2000).

[18] Bravyi, S. & Kitaev, A. Quantum codes on a lattice with boundary. *arXiv* 9811052 (1998).

[19] Chow, J. M. *et al.* Implementing a strand of a scalable fault-tolerant quantum computing fabric. *Nature communications* **5**, 4015 (2014).

[20] Rønnow, T. F. *et al.* Defining and detecting quantum speedup. *Science* **345**, 420 (2014).

[21] Roushan, P. *et al.* Observation of topological transitions in interacting quantum circuits. *Nature* **515**, 241–244 (2014).

[22] Vijay, R., Slichter, D. H. & Siddiqi, I. Observation of Quantum Jumps in a Superconducting Artificial Atom. *Physical Review Letters* **106**, 110502 (2011).

[23] Sun, L. *et al.* Tracking photon jumps with repeated quantum non-demolition parity measurements. *Nature* **511**, 444–448 (2014).

[24] Murch, K. W. *et al.* Cavity-Assisted Quantum Bath Engineering. *Physical Review Letters* **109**, 183602 (2012).

[25] Murch, K. W., Weber, S. J., Beck, K. M., Ginossar, E. & Siddiqi, I. Reduction of the radiative decay of atomic coherence in squeezed vacuum. *Nature* **499**, 62–5 (2013).

[26] Kirchmair, G. *et al.* Observation of quantum state collapse and revival due to the single-photon Kerr effect. *Nature* **495**, 205–9 (2013).

[27] Murch, K. W., Weber, S. J., Macklin, C. & Siddiqi, I. Observing single quantum trajectories of a superconducting quantum bit. *Nature* **502**, 211–4 (2013).

[28] Weber, S. J. *et al.* Mapping the optimal route between two quantum states. *Nature* **511**, 570–573 (2014).

[29] Carmichael, H. J. *An Open Systems Approach to Quantum Optics* (Springer, Berlin, 1993).

[30] Caves, C. Quantum limits on noise in linear amplifiers. *Physical Review D* **26**, 1817–1839 (1982).

[31] Walls, D. F. & Milburn, G. J. *Quantum Optics* (Springer, 2008).

[32] Feldman, M. J., Parrish, P. T. & Chiao, R. Y. Parametric amplification by unbiased josephson junctions. *J. Appl. Phys.* **46**, 4031 (1975).

[33] Yurke, B. *et al.* Observation of parametric amplification and deamplification in a Josephson parametric amplifier. *Physical Review A* **39**, 2519–2533 (1989).

[34] Castellanos-Beltran, M. A. & Lehnert, K. W. Widely tunable parametric amplifier based on a superconducting quantum interference device array resonator. *Applied Physics Letters* **91**, 083509 (2007).

[35] Hatridge, M., Vijay, R., Slichter, D. H., Clarke, J. & Siddiqi, I. Dispersive magnetometry with a quantum limited SQUID parametric amplifier. *Physical Review B* **83**, 134501 (2011).

[36] Bergeal, N. *et al.* Phase-preserving amplification near the quantum limit with a Josephson ring modulator. *Nature* **465**, 64–8 (2010).

[37] Chantasri, A., Dressel, J. & Jordan, A. N. Action principle for continuous quantum measurement. *Physical Review A* **88**, 042110 (2013).

[38] Jacobs, K. & Steck, D. A straightforward introduction to continuous quantum measurement. *Contemporary Physics* **47**, 279 (2006).

[39] Braginsky, V. B. & Khalili, F. Y. *Qunatum Measurement* (Cambridge University Press, 1992).

[40] LaHaye, M. D., Buu, O., Camarota, B. & Schwab, K. C. Approaching the quantum limit of a nanomechanical resonator. *Science (New York, N.Y.)* **304**, 74–77 (2004).

[41] Teufel, J. D., Donner, T., Castellanos-Beltran, M. A., Harlow, J. W. & Lehnert, K. W. Nanomechanical motion measured with an imprecision below that at the standard quantum limit. *Nature nanotechnology* **4**, 820–3 (2009).

[42] Caves, C., Thorne, K., Drever, R., Sandberg, V. & Zimmermann, M. On the measurement of a weak classical force coupled to a quantum-mechanical oscillator. I. Issues of principle. *Reviews of Modern Physics* **52**, 341–392 (1980).

[43] Milburn, G. J. & Walls, D. F. Quantum nondemolition measurements via quadratic coupling. *Physical Review A* **28**, 2065–2070 (1983).

[44] Roch, J. F. *et al.* Quantum-Nondemolition Measurements Using Cold Trapped Atoms. *Phys Rev Lett* **78**, 634–637 (1997).

[45] Nogues, G. *et al.* Seeing a single photon without destroying it. *Nature* **400**, 239–242 (1999).

[46] Wallraff, A. *et al.* Approaching unit visibility for control of a superconducting qubit with dispersive readout. *Phys. Rev. Lett.* **95** (2005).

[47] Groen, J. P. *et al.* Partial-measurement backaction and nonclassical weak values in a superconducting circuit. *Physical Review Letters* **111** (2013).

[48] Gambetta, J. *et al.* Quantum trajectory approach to circuit QED: Quantum jumps and the Zeno effect. *Physical Review A - Atomic, Molecular, and Optical Physics* **77**, 012112 (2008).

[49] Vijay, R. *et al.* Stabilizing Rabi oscillations in a superconducting qubit using quantum feedback. *Nature* **490**, 77–80 (2012).

[50] Wiseman, H. M. & Milburn, G. J. *Quantum Measurement and Control* (Cambridge University Press, 2010).

[51] Devoret, M. H. & Schoelkopf, R. J. Superconducting circuits for quantum information: an outlook. *Science (New York, N.Y.)* **339**, 1169–74 (2013).

[52] Korotkov, A. Continuous quantum measurement of a double dot. *Physical Review B* **60**, 5737–5742 (1999).

[53] Korotkov, A. N. Quantum Bayesian approach to circuit QED measurement. *arXiv:* 1111.4016 (2011).

[54] Korotkov, A. Selective quantum evolution of a qubit state due to continuous measurement. *Physical Review B* **63**, 115403 (2001).

[55] Dalibard, J., Castin, Y. & Mø lmer, K. Wave-function approach to dissipative processes in quantum optics (1992).

[56] Gardiner, C., Parkins, A. & Zoller, P. Wave-function quantum stochastic differential equations and quantum-jump simulation methods (1992).

[57] Schack, R., Brun, T. A. & Percival, I. C. Quantum state diffusion, localization and computation. *J. Phys. A.* **28**, 5401–5413 (1995).

[58] Nagourney, W., Sandberg, J. & Dehmelt, H. Shelved optical electron amplifier: observation of quantum jumps. *Phys. Rev. Lett.* **56**, 2797 (1986).

[59] Guerlin, C. *et al.* Progressive field-state collapse and quantum non-demolition photon counting. *Nature* **448**, 889–93 (2007).

[60] Hood, C. J. The Atom-Cavity Microscope: Single Atoms Bound in Orbit by Single Photons. *Science* **287**, 1447–1453 (2000).

[61] Goan, H.-S., Milburn, G., Wiseman, H. & Bi Sun, H. Continuous quantum measurement of two coupled quantum dots using a point contact: A quantum trajectory approach. *Physical Review B* **63**, 125326 (2001).

[62] Sukhorukov, E. V. *et al.* Conditional statistics of electron transport in interacting nanoscale conductors. *Nature Physics* **3**, 243–247 (2007).

[63] Hatridge, M. *et al.* Quantum back-action of an individual variable-strength measurement. *Science* **339**, 178–81 (2013).

[64] Campagne-Ibarcq, P. *et al.* Observing Interferences between Past and Future Quantum States in Resonance Fluorescence. *Physical Review Letters* **112**, 180402 (2014).

[65] Tan, D., Weber, S., Siddiqi, I., Mø lmer, K. & Murch, K. W. Prediction and retrodiction for a continuously monitored superconducting qubit. *arXiv:* 1409.0510 (2014).

[66] Slichter, D. H. *Quantum jumps and measurement backaction in a superconducting qubit.* Ph.D. thesis, University of California, Berkeley (2011).

[67] Brun, T. A. A simple model of quantum trajectories. *Am. J. Phys.* **70**, 7 (2002).

[68] Belavkin, V. P. *Information, Complexity and Control in Quantum Physics* (Springer-Verlag, 1987).

[69] Schuster, D. I. *Circuit Quantum Electrodynamics.* Ph.D. thesis, Yale University (2007).

[70] Devoret, M. H. & Martinis, J. M. Implementing Qubits with Superconducting Integrated Circuits. *Quant. Inf. Proc.* **3**, 163–203 (2004).

[71] Schoelkopf, R. J. & Girvin, S. M. Wiring up quantum systems. *Nature* **451**, 664 (2008).

[72] Clarke, J. & Wilhelm, F. K. Superconducting quantum bits. *Nature* **453**, 1031–42 (2008).

[73] You, J. Q. & Nori, F. Superconducting circuits and quantum information. *Physics Today* **58**, 42–47 (2005).

[74] Devoret, M. H. Quantum fluctuations in electrical circuits. In Reynaud, S., Giacobino, E. & Zinn-Justin, J. (eds.) *Quantum fluctuations, Les Houches LXIII*, chap. 10. (Elsevier, New York, 1995).

[75] Josephson, B. Possible new effects in superconductive tunnelling. *Physics Letters* **1**, 251–253 (1962).

[76] Clarke, J. & Braginski, A. I. *The SQUID Handbook*, vol. 2 (Wiley, 2006).

[77] Paik, H. *et al.* Observation of High Coherence in Josephson Junction Qubits Measured in a Three-Dimensional Circuit QED Architecture. *Physical Review Letters* **107**, 240501 (2011).

[78] Chang, J. B. *et al.* Improved superconducting qubit coherence using titanium nitride. *Applied Physics Letters* **103**, 012602 (2013).

[79] Barends, R. *et al.* Coherent Josephson Qubit Suitable for Scalable Quantum Integrated Circuits. *Physical Review Letters* **111**, 080502 (2013).

[80] Koch, J. *et al.* Charge-insensitive qubit design derived from the Cooper pair box. *Physical Review A* **76**, 042319 (2007).

[81] Chow, J. M. *et al.* Microwave-activated conditional-phase gate for superconducting qubits. *New Journal of Physics* **15**, 115012 (2013).

[82] Chow, J. M. *et al.* Optimized driving of superconducting artificial atoms for improved single-qubit gates. *Physical Review A* **82**, 040305 (2010).

[83] Berman, P. R., Bates, D. R. & Bederson, B. (eds.) *Cavity Quantum Electrodynamics* (Academic Press, 1993).

[84] Blais, A., Huang, R.-S., Wallraff, A., Girvin, S. & Schoelkopf, R. Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation. *Physical Review A* **69**, 062320 (2004).

[85] Blais, A. *et al.* Quantum-information processing with circuit quantum electrodynamics. *Physical Review A* **75**, 032329 (2007).

[86] Vijay, R. *Josephson Bifurcation Amplifier: Amplifying quantum signals using a dynamical bifurcation.* Ph.D. thesis, Yale University (2008).

[87] Eichler, C. *Experimental Characterization of Quantum Microwave Radiation and its Entanglement with a Superconducting Qubit.* Ph.D. thesis, ETH Zuerich (2013).

[88] Schackert, F. *A Practical Quantum-Limited Parametric Amplifier Based on the Josephson Ring Modulator.* Ph.D. thesis, Yale University (2013).

[89] Castellanos-Beltran, M. A. *Development of a Josephson Parametric Amplifier for the Preparation and Detection of Nonclassical States of Microwave Fields.* Ph.D. thesis, University of Colorado (2010).

[90] Eichler, C. & Wallraff, A. Controlling the dynamic range of a Josephson parametric amplifier. *EPJ Quantum Technology* **1**, 2 (2014).

[91] Gardiner, C. & Zoller, P. *Quantum Noise* (Springer, 2004).

[92] Turchette, Q., Georgiades, N., Hood, C., Kimble, H. & Parkins, A. Squeezed excitation in cavity QED:âĂČExperiment and theory. *Physical Review A* **58**, 4056–4077 (1998).

[93] Boutin, S. & Blais, A. *unpublished* (2014).

[94] Boissonneault, M. *et al.* Back-action of a driven nonlinear resonator on a superconducting qubit. *Physical Review A* **85**, 022305 (2012).

[95] Gardiner, C. & Collett, M. Input and output in damped quantum systems: Quantum stochastic differential equations and the master equation. *Physical Review A* **31**, 3761–3774 (1985).

[96] Mutus, J. Y. *et al.* Strong environmental coupling in a Josephson parametric amplifier. *Applied Physics Letters* **104**, 263513 (2014).

[97] Yaakobi, O., Friedland, L., Macklin, C. & Siddiqi, I. Parametric amplification in Josephson junction embedded transmission lines. *Physical Review B* **87**, 144301 (2013).

[98] Yaakobi, O., Friedland, L., Macklin, C. & Siddiqi, I. Erratum: Parametric amplification in Josephson junction embedded transmission lines. *Physical Review B* **88**, 219904 (2013).

[99] O'Brien, K., Macklin, C., Siddiqi, I. & Zhang, X. Resonant Phase Matching of Josephson Junction Traveling Wave Parametric Amplifiers. *Physical Review Letters* **113**, 157001 (2014).

[100] Khan, S., Vijay, R., Siddiqi, I. & Clerk, A. A. Large gain quantum-limited qubit measurement using a two-mode nonlinear cavity. *arXiv* 1408.0226 (2014).

[101] Narla, A. *et al.* Wireless josephson amplifier. *Applied Physics Letters* **104**, 232605 (2014).

[102] Geerlings, K. *et al.* Improving the quality factor of microwave compact resonators by optimizing their geometrical parameters. *Applied Physics Letters* **100**, 192601 (2012).

[103] Megrant, A. *et al.* Planar superconducting resonators with internal quality factors above one million. *Applied Physics Letters* **100**, 113510 (2012).

[104] Kamal, A., Marblestone, A. & Devoret, M. Signal-to-pump back action and self-oscillation in double-pump Josephson parametric amplifier. *Physical Review B* **79**, 184301 (2009).

[105] Slichter, D. H. *et al.* Measurement-Induced Qubit State Mixing in Circuit QED from Up-Converted Dephasing Noise. *Physical Review Letters* **109**, 153601 (2012).

[106] Johnson, J. E. *et al.* Heralded State Preparation in a Superconducting Qubit. *Physical Review Letters* **109**, 050506 (2012).

[107] Steffen, M. *et al.* State Tomography of Capacitively Shunted Phase Qubits with High Fidelity. *Physical Review Letters* **97**, 050502 (2006).

[108] Schuster, D. *et al.* ac Stark Shift and Dephasing of a Superconducting Qubit Strongly Coupled to a Cavity Field. *Physical Review Letters* **94**, 123602 (2005).

[109] Gambetta, J. *et al.* Qubit-photon interactions in a cavity: Measurement-induced dephasing and number splitting. *Physical Review A* **74**, 042318 (2006).

[110] Wiseman, H. M., Jones, S. J. & Doherty, A. C. Steering, Entanglement, Nonlocality, and the Einstein-Podolsky-Rosen Paradox. *Physical Review Letters* **98**, 140402 (2007).

[111] Wiseman, H. M. & Gambetta, J. M. Are Dynamical Quantum Jumps Detector Dependent? *Physical Review Letters* **108**, 220402 (2012).

[112] Chantasri, A., Dressel, J. & Jordan, A. N. Action principle for continuous quantum measurement. *Physical Review A* **88**, 042110 (2013).

[113] Dykman, M. I., Mori, E., Ross, J. & Hunt, P. M. Large fluctuations and optimal paths in chemical kinetics. *The Journal of Chemical Physics* **100**, 5735 (1994).

[114] Pilgram, S., Jordan, A. N., Sukhorukov, E. V. & Büttiker, M. Stochastic Path Integral Formulation of Full Counting Statistics. *Physical Review Letters* **90**, 206801 (2003).

[115] Jordan, A. N., Sukhorukov, E. V. & Pilgram, S. Fluctuation statistics in networks: A stochastic path integral approach. *Journal of Mathematical Physics* **45**, 4386 (2004).

[116] Sukhorukov, E. V. & Jordan, A. N. Stochastic Dynamics of a Josephson Junction Threshold Detector. *Physical Review Letters* **98**, 136803 (2007).

[117] Mensky, M. Quantum restrictions for continuous observation of an oscillator. *Physical Review D* **20**, 384–387 (1979).

[118] Caves, C. Quantum mechanics of measurements distributed in time. A path-integral formulation. *Physical Review D* **33**, 1643–1665 (1986).

[119] Schrödinger, E. Uber die umkehrung der naturgesetze. *Sitzungsberichte der Preuss Akad. Wissen. Berlin, Phys. Math. Klasse* 144–153 (1931).

[120] Pavon, M. Quantum schrödinger bridges. *Directions in Mathematical Systems Theory and OptimizationDirections in Mathematical Systems Theory and Optimization Lecture Notes in Control and Information Sciences* **286** (2003).

[121] Fortet, R. Résolution d'un système d'equations de m. schrödinger. *J. Math. Pure Appl.* **IX**, 83–105 (1940).

[122] Zambrini, J. Stochastic mechanics according to E. Schrödinger. *Physical Review A* **33**, 1532–1548 (1986).

[123] Levy, B. C. & Krener, A. J. Stochastic mechanics of reciprocal diffusions. *Journal of Mathematical Physics* **37**, 769 (1996).

[124] Pavon, M. & Ticozzi, F. Discrete-time classical and quantum Markovian evolutions: Maximum entropy problems on path space. *Journal of Mathematical Physics* **51**, 042104 (2010).

[125] Georgiou, T. T. & Pavon, M. Positive contraction mappings for classical and quantum Schrodinger systems. *arXiv:* 1405.6650 (2014).

[126] Beghi, A., Ferrante, A. & Pavon, M. How to Steer a Quantum System over a Schrödinger Bridge. *Quantum Information Processing* **1**, 183–206 (2002).

[127] Shapiro, M. & Brumer, P. *Quantum Control of Molecular Processes* (Wiley, 2012).

[128] Watanabe, S. Symmetry of Physical Laws. Part III. Prediction and Retrodiction. *Reviews of Modern Physics* **27**, 179–186 (1955).

[129] Aharonov, Y., Bergmann, P. & Lebowitz, J. Time Symmetry in the Quantum Process of Measurement. *Physical Review* **134**, B1410–B1416 (1964).

[130] Aharonov, Y., Albert, D. & Vaidman, L. How the result of a measurement of a component of the spin of a spin-1/2 particle can turn out to be 100. *Physical Review Letters* **60**, 1351–1354 (1988).

[131] Williams, N. & Jordan, A. Weak Values and the Leggett-Garg Inequality in Solid-State Qubits. *Physical Review Letters* **100**, 026804 (2008).

[132] Sayrin, C. *et al.* Real-time quantum feedback prepares and stabilizes photon number states. *Nature* **477**, 73–7 (2011).

[133] Campagne-Ibarcq, P. *et al.* Persistent Control of a Superconducting Qubit by Stroboscopic Measurement Feedback. *Physical Review X* **3**, 021008 (2013).

[134] Toyli, D. *unpublished* (2014).

[135] Eichler, C., Salathe, Y., Mlynek, J., Schmidt, S. & Wallraff, A. Quantum-Limited Amplification and Entanglement in Coupled Nonlinear Resonators. *Physical Review Letters* **113**, 110502 (2014).

[136] Lvovsky, A. I. Continuous-variable optical quantum-state tomography. *Reviews of Modern Physics* **81**, 299–332 (2009).

[137] Mallet, F. *et al.* Quantum State Tomography of an Itinerant Squeezed Microwave Field. *Physical Review Letters* **106**, 220502 (2011).

[138] Bergeal, N., Schackert, F., Frunzio, L. & Devoret, M. H. Two-Mode Correlation of Microwave Quantum Noise Generated by Parametric Down-Conversion. *Physical Review Letters* **108**, 123902 (2012).

[139] Flurin, E., Roch, N., Mallet, F., Devoret, M. H. & Huard, B. Generating Entangled Microwave Radiation Over Two Transmission Lines. *Physical Review Letters* **109**, 183901 (2012).

[140] Eichler, C. *et al.* Experimental State Tomography of Itinerant Single Microwave Photons. *Physical Review Letters* **106**, 220503 (2011).

[141] Menzel, E. P. *et al.* Path Entanglement of Continuous-Variable Quantum Microwaves. *Physical Review Letters* **109**, 250502 (2012).

[142] Di Candia, R. *et al.* Dual-path methods for propagating quantum microwaves. *New Journal of Physics* **16**, 015001 (2014).

[143] Hofheinz, M. *et al.* Synthesizing arbitrary quantum states in a superconducting resonator. *Nature* **459**, 546–9 (2009).

[144] Gardiner, C. Inhibition of Atomic Phase Decays by Squeezed Light: A Direct Effect of Squeezing. *Physical Review Letters* **56**, 1917–1920 (1986).

[145] Georgiades, N., Polzik, E., Edamatsu, K., Kimble, H. & Parkins, A. Nonclassical Excitation for Atoms in a Squeezed Vacuum. *Physical Review Letters* **75**, 3426–3429 (1995).

[146] Dayan, B., Pe'er, A., Friesem, A. & Silberberg, Y. Two Photon Absorption and Coherent Control with Broadband Down-Converted Light. *Physical Review Letters* **93**, 023005 (2004).

[147] Rybarczyk, T. *et al.* Past quantum state analysis of the photon number evolution in a cavity. *arXiv* 1409.0958 (2014).

[148] Ruskov, R. & Korotkov, A. Entanglement of solid-state qubits by measurement. *Physical Review B* **67**, 241305 (2003).

[149] Lalumière, K., Gambetta, J. M. & Blais, A. Tunable joint measurements in the dispersive regime of cavity QED. *Physical Review A* **81**, 040301 (2010).

[150] Chou, C. W. *et al.* Measurement-induced entanglement for excitation stored in remote atomic ensembles. *Nature* **438**, 828–32 (2005).

[151] Moehring, D. L. *et al.* Entanglement of single-atom quantum bits at a distance. *Nature* **449**, 68–71 (2007).

[152] Bernien, H. *et al.* Heralded entanglement between solid-state qubits separated by three metres. *Nature* **497**, 86–90 (2013).

[153] Silberfarb, A., Jessen, P. & Deutsch, I. Quantum State Reconstruction via Continuous Measurement. *Physical Review Letters* **95**, 030402 (2005).

[154] Smith, A. *et al.* Quantum state tomography by continuous measurement and compressed sensing. *Physical Review A* **87**, 030102 (2013).

[155] Ristè, D., Bultink, C. C., Lehnert, K. W. & DiCarlo, L. Feedback Control of a Solid-State Qubit Using High-Fidelity Projective Measurement. *Physical Review Letters* **109**, 240502 (2012).

[156] Jacobs, K. How to project qubits faster using quantum feedback. *Physical Review A* **67**, 030301 (2003).

[157] Wiseman, H. M. & Ralph, J. F. Reconsidering rapid qubit purification by feedback. *New Journal of Physics* **8**, 90–90 (2006).

[158] Li, H., Shabani, A., Sarovar, M. & Whaley, K. Optimality of qubit purification protocols in the presence of imperfections. *Physical Review A* **87**, 032334 (2013).

[159] Blok, M. S. *et al.* Manipulating a qubit through the backaction of sequential partial measurements and real-time feedback. *Nature Physics* **10**, 189–193 (2014).

[160] Weber, S. J., Murch, K. W., Slichter, D. H., Vijay, R. & Siddiqi, I. Single crystal silicon capacitors with low microwave loss in the single photon regime. *Applied Physics Letters* **98**, 172510 (2011).

[161] O'Connell, A. D. *et al.* Microwave dielectric loss at single photon energies and millikelvin temperatures. *Applied Physics Letters* **92**, 112903 (2008).

[162] Paik, H. & Osborn, K. D. Reducing quantum-regime dielectric loss of silicon nitride for superconducting quantum circuits. *Applied Physics Letters* **96**, 072505 (2010).

[163] Cicak, K. *et al.* Low-loss superconducting resonant circuits using vacuum-gap-based microwave components. *Applied Physics Letters* **96**, 093502 (2010).

[164] Martinis, J. *et al.* Decoherence in Josephson Qubits from Dielectric Loss. *Physical Review Letters* **95**, 210503 (2005).

[165] Shnirman, A., Schön, G., Martin, I. & Makhlin, Y. Low- and High-Frequency Noise from Coherent Two-Level Systems. *Physical Review Letters* **94**, 127002 (2005).

[166] Denis, K. L. *et al.* Fabrication of an Antenna-Coupled Bolometer for Cosmic Microwave Background Polarimetry. In *The Thirteenth International Workshop on Low Temperature Detectors—LTD13*, vol. 1185, 371–374 (AIP Publishing, 2009).

[167] Patel, U. *et al.* Coherent Josephson phase qubit with a single crystal silicon capacitor. *Applied Physics Letters* **102**, 012602 (2013).

[168] Murch, K. W., Weber, S. J., Levenson-Falk, E. M., Vijay, R. & Siddiqi, I. 1/f noise of Josephson-junction-embedded microwave resonators at single photon energies and millikelvin temperatures. *Applied Physics Letters* **100**, 142601 (2012).

[169] Pop, I. M. *et al.* Coherent suppression of electromagnetic dissipation due to superconducting quasiparticles. *Nature* **508**, 369–72 (2014).

[170] Lenander, M. *et al.* Measurement of energy decay in superconducting qubits from nonequilibrium quasiparticles. *Physical Review B* **84**, 024501 (2011).

[171] Córcoles, A. D. *et al.* Protecting superconducting qubits from radiation. *Applied Physics Letters* **99**, 181906 (2011).

[172] 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 (2012).

Appendix A

Solving the generalized Jaynes-Cummings Hamiltonian

In this appendix, we explain how to numerically diagonalize the generalized Jaynes-Cummings Hamiltonian to solve for its energy eigenvalues. The Hamiltonian for a multi-level qubit is given by

$$H/\hbar = \omega_r \hat{a}^\dagger \hat{a} + \sum_{m=0}^{M} \omega_m \left|m\right\rangle \left\langle m\right| + \sum_{m=0}^{M} g_{m,m+1} \left(\left|m\right\rangle \left\langle m+1\right| \hat{a}^\dagger + \left|m+1\right\rangle \left\langle m\right| \hat{a} \right), \qquad \text{(A.1)}$$

where we consider a total of M qubit energy levels indexed by m, with energies $\hbar\omega_m$. We can represent this Hamiltonian as an $(M+1)(N+1) \times (M+1)(N+1)$ matrix, where we choose N to be the maximum photon number. We can express the matrix elements in the form $H_{n',m',n,m} \equiv \langle n', m'|H|n, m\rangle$, where $0 \leq n \leq N$ describes the cavity occupation $a^\dagger a$ and $0 \leq m \leq M$ describes the qubit level. The matrix elements can be expresses as

$$\begin{aligned} H_{n',m',n,m} =& \omega_r n \delta_{n,n'} \delta_{m,m'} + \omega_m \delta_{n,n'} \delta_{m,m'} \\ &+ \sqrt{n+1} g_{m,m-1} \delta_{n',n+1} \delta_{m',m-1} + \sqrt{n} g_{m,m+1} \delta_{n',n-1} \delta_{m',m+1}, \end{aligned} \qquad \text{(A.2)}$$

where $g_{i,j}$ is the coupling strength between levels i and j.

With the help of Oliver Viehmann, I developed a Mathematica code which expresses these matrix elements numerically for a transmon qubit. The energy eigenvalues can be found numerically by simply diagonalizing the Hamiltonian, using the Mathematica function 'eigenvalues'. The raw code, as well as some example energy levels are shown in the following pages. Note that the labeling conventions differ slightly from the preceding discussion, and can be understood by making the substutions $N \to \nu$, $m \to$ xn, $n' \to$ xm, and $n \to$ jn.

```
(* Finds the m'th transmon energy level, given Ej,
Ec, and a flux bias paramatrized by the angle `th'. *)
vark[ng_, m_] := Sum[Mod[Round[2 * ng + ell / 2], 2] *
    (Round[ng] + ell * (-1)^m * Quotient[m + 1, 2]), {ell, {1, -1}}]
ng = .1;
levels[Ej_, Ec_, m_, th_] := With[{Ejrat = Ej / Ec},
    (MathieuCharacteristicA[2 * (ng + vark[ng, m]), -Ejrat * Cos[th] / 2] -
       MathieuCharacteristicA[2 * (ng + vark[ng, 0]), -Ejrat * Cos[th] / 2]) * Ec];

(* Solves the generalized JC Hamiltonian for qubit parameters Ej and Ec,
with M+1 levels, coupled to a resonator of
 frequency ω which can have a maximum of ν photons. *)
GeneralizedJC[M_, ν_, ω_, ge_, Ej_, Ec_, th_] :=
 With[{ωr = ω}, With[{λt = ge}, m = DiagonalMatrix[
       Flatten[
        Table[ωr * jn + levels[Ej, Ec, xn, th]
          , {jn, 0, ν}, {xn, 0, M}]
       ]
      ]]
   Do[Do[If[xn - xm == 1, m[[1 + xm + (M + 1) * (jn + 1), 1 + xn + (M + 1) * jn]] =
        N[Sqrt[xn]] ge * N[Sqrt[jn + 1], 0], {xm, 0, M}, {xn, 0, M}], {jn, 0, ν - 1}];
  Do[Do[If[xm - xn == 1, m[[1 + xm + (M + 1) * (jn - 1), 1 + xn + (M + 1) * jn]] =
        N[Sqrt[xn + 1]] ge * N[Sqrt[jn], 0], {xm, 0, M}, {xn, 0, M}], {jn, 1, ν}]; m
 ]

(* Here, we plot the the energy eigenvalues as a function of flux bias,
for a three-level transmon with a maximum 3 cavity photons *)
Plot[{Eigenvalues[GeneralizedJC[2, 3, 6, .120, 100, .28, phi]]}, {phi, .4 Pi, .46 Pi}]
```

Out[5]=

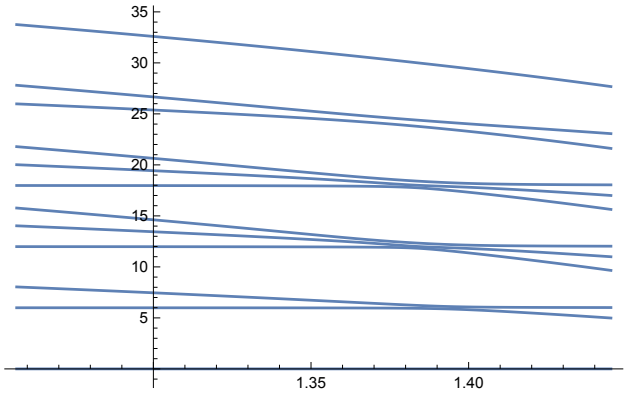

```
(* Here, we generate a list of the eigenvalues for a two-
 level transmon with up to 3 cavity photons, at one particular flux bias. *)
GeneralizedJC[1, 3, 6, .120, 100, .28, .4 * Pi] // Eigenvalues
```

Out[6]= {26.0295, 20.0505, 17.9789, 14.0436, 11.9859, 8.03654, 5.99293, 0.}

```
In[7]:= (* Here, we plot the qubit transition frequencies as a function of flux,
for different cavity photon occupations.  We also
 plot the cavity frequency to show the avoided crossing. *)
Plot[{Eigenvalues[GeneralizedJC[1, 3, 6, .120, 100, .28, phi]][[6]],
  Eigenvalues[GeneralizedJC[1, 3, 6, .120, 100, .28, phi]][[7]],
  Eigenvalues[GeneralizedJC[1, 3, 6, .120, 100, .28, phi]][[4]] -
   Eigenvalues[GeneralizedJC[1, 3, 6, .120, 100, .28, phi]][[7]],
  Eigenvalues[GeneralizedJC[1, 3, 6, .120, 100, .28, phi]][[2]] -
   Eigenvalues[GeneralizedJC[1, 3, 6, .120, 100, .28, phi]][[5]],
  Eigenvalues[GeneralizedJC[1, 3, 6, .120, 100, .28, phi]][[5]] -
   Eigenvalues[GeneralizedJC[1, 3, 6, .120, 100, .28, phi]][[6]],
  Eigenvalues[GeneralizedJC[1, 3, 6, .120, 100, .28, phi]][[3]] - Eigenvalues[
     GeneralizedJC[1, 3, 6, .120, 100, .28, phi]][[4]]}, {phi, .4 * Pi, .47 * Pi}]
```

Out[7]=

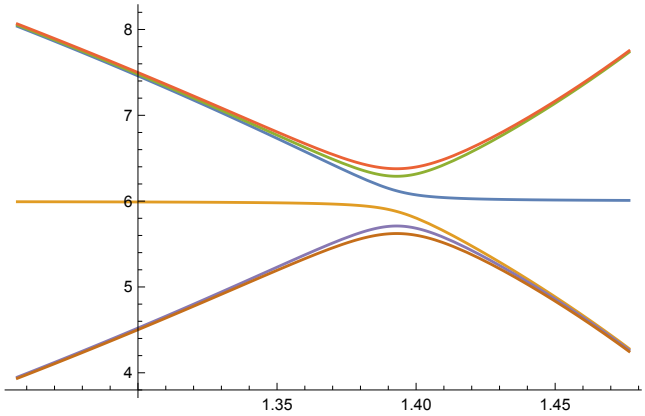

Appendix B

Single crystal silicon capacitors

This section discusses a project from early on in my graduate career, where we developed low-loss parallel plate capacitors using single crystal silicon as the dielectric material. Our discussion closely follows reference [160], where the work was originally published. While this work is not directly related to the quantum trajectories experiments discussed in this thesis, it represents an important part of a concerted effort in our lab and in the broader community to understand the sources of loss and dephasing in superconducting circuits. Such efforts are essential to improving the T_1 and T_2^* times of superconducting qubits. If not for the substantial improvements in coherence times over the course of my PhD, our experiments on quantum trajectories and squeezing would have been prohibitively difficult to implement.

In superconducting resonant circuits, low-loss components are required in order to realize high internal quality factors Q_i. For planar lumped-element resonators and planar qubits the highest values of Q_i are typically achieved using the geometric capacitance of single-layer micro-fabricated superconducting structures such as interdigitated capacitors, on single crystal silicon or sapphire substrates. Using standard electron beam evaporated aluminum or sputtered niobium films, such circuits can achieve Q_i values of several hundred thousand with an average resonator population of one photon [102]. Recently, Q_i values on the order of one million have been demonstrated using using carefully prepared TiN [78] and aluminum grown by molecular beam epitaxy [103]. However, it is difficult to realize capacitances larger than about 1 pF with single-layer capacitors. Higher capacitances can be obtained in a parallel plate geometry using either deposited amorphous dielectrics [161] such as SiO_2, a-Si:H, or SiN_x [162] or vacuum gaps [163]. However, deposited dielectrics exhibit significant loss in the low temperature, low power regime due to the presence of two-level state (TLS) defects [164, 165], limiting their utility for quantum circuits. Current vacuum gap capacitors suffer from loss from necessary support structures and surface oxides, and are on par with the best deposited dielectrics, with a loss tangent $\tan\delta_i = 1/Q_i = 2.2 - 3 \times 10^{-5}$ in the 4-8 GHz band [161, 162, 163]. Crystalline dielectrics, such as silicon can exhibit low intrinsic loss and have been used in the fabrication of superconducting bolometers using a wafer bonding process [166].

We have developed superconducting microwave resonators in a lumped element geometry

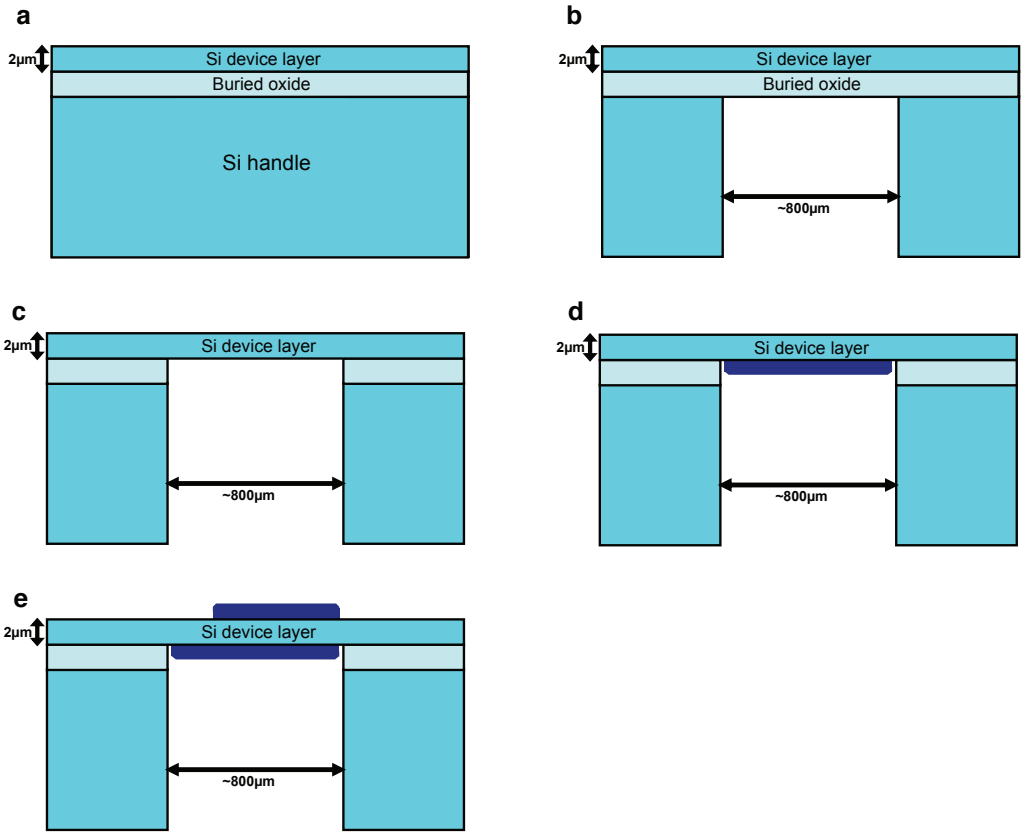

Figure B.1: Fabrication of single crystal silicon capacitors. An SOI wafer (a) is etched from below up to the buried oxide layer (b), which acts as an etch stop. The buried oxide is removed by HF vapor (c), forming a window of suspended high resistivity silicon. We then deposit Aluminum on the bottom (d) and top (e) surfaces of the Si device layer, forming a parallel plate capacitor.

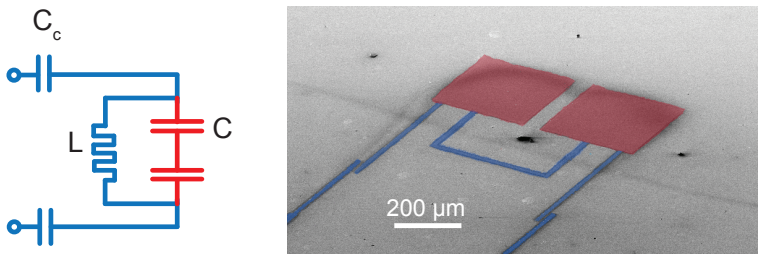

Figure B.2: Resonator design for testing silicon capacitors. The circuit schematic (left) is color-coded to match the a false-colored SEM image (right).

using single crystal silicon dielectric parallel plate capacitors with C > 2 pF. The fabrication process is depicted in figure B.1. The Si dielectric layer was formed from a commercially available (Ultrasil) high resistivity ($\rho > 1$ kω-cm) silicon-on-insulator (SOI) wafer. The wafer consists of a 2 μm device layer which we use as the capacitor dielectric and a 2 μm buried oxide layer on top of a thick ($\sim 400\,\mu$m) silicon handle wafter. The handle wafer was patterned with 8 microns of SPR 220-7 resist, and etched using a deep silicon etch process, with alternating cycles of SF_6 plasma etching and C_4F_8 passivation. The buried oxide layer acted as the etch stop. Both the buried oxide layer and the native oxide on top of the silicon device layer were removed by HF vapor. Ground planes were deposited on the bottom of the device layer using 300 nm of e-beam evaporated aluminum. The top device layer was optically patterned and coated with 100 nm of aluminum.

We fabricated resonators with resonant frequencies between 4.0 and 6.5 GHz by shunting the capacitors of various capacitances ($\sim 2\,\mu$A) with a $L \simeq 600$ pH superconducting meander inductor, as depicted in figure B.2. The resonant response was obtained using a vector network analyzer. The excitation power used to probe the resonators varied between $P_{\rm in} = -94$ and $P_{\rm in} = -157$ dBm, and for the one port (reflection) devices we consider is related to the average photon number on resonance as $\bar{n} = 4P_{\rm in}Q^2/(Q_{\rm ext}\hbar\omega_0^2)$, where ω_0 is the resonant frequency and the total quality factor $Q = (1/Q_i + 1/Q_{\rm ext})^{-1}$. The quality factors Q_i and $Q_{\rm ext}$ are determined by fitting to the resonance curve, as discussed in appendix D.

In figure B.3, we plot the extracted internal quality factor versus excitation power expressed in terms of average photon number for aluminum lumped element resonators in the 4 to 5 GHz range. We note that stripping the native oxide off the top of the device layer prior to patterning the surface increased the resonator Q nearly tenfold. Of the devices tested with this surface treatment, the average intrinsic loss tangent in the low temperature, single average photon regime was $\tan\delta_i = 5.4 \times 10^{-6}$. If we assume that the observed losses are solely due to the dielectric layer, our measurements are consistent with the current reported microwave loss tangent of single crystal silicon under these experimental conditions [161].

Our results indicate that single crystal silicon capacitors are very promising for use in high quality factor superconducting circuits which require > pF of shunting capacitance,

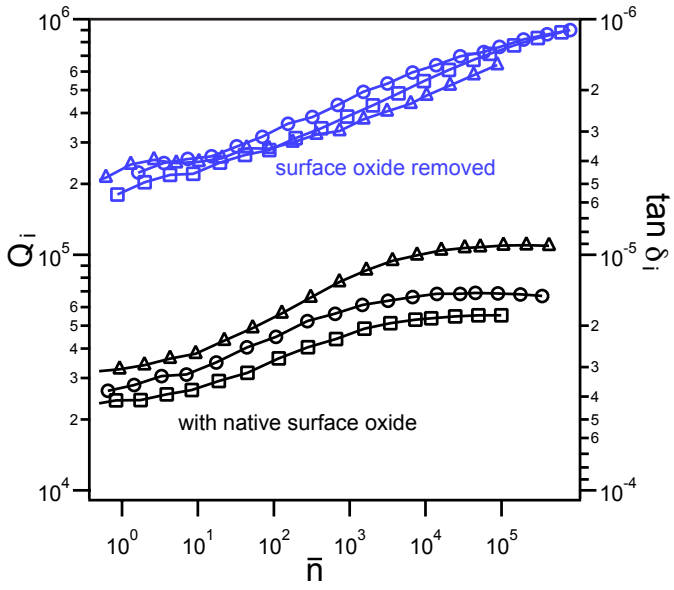

Figure B.3: Q_i vs $\bar{n}$ for several resonators with native oxide present (black) and removed (blue).

such as the phase qubit. In fact, the group of Robert McDermott has recently used single crystal silicon capacitors to realize the longest-lived phase qubits reported to date, with T_1 times of several microseconds [167]. Furthermore, incorporating these elements into classical non-linear resonators consisting of a capacitively shunted Josephson junction has allowed for precise characterization of the microwave loss and 1/f noise [168] in Josephson junctions.

Appendix C

Characterizing microwave frequency loss in Josephson junctions

In this section, we present unpublished results from our efforts to understand microwave frequency loss in Josephson junctions. The current consensus in the superconducting qubit community is that small ($\sim 0.01\,\mu\mathrm{m}^2$) shadow-evaporated $\mathrm{Al/AlO}_x\mathrm{/Al}$ junctions have remarkably low capacitive loss. In fact, recent results on the fluxonium qubit place an upper bound of 3×10^{-7} on the loss tangent of the junction capacitance [169]. However, earlier measurements of phase qubits suggest that large (>10 $\mu\mathrm{m}^2$) aluminum oxide junctions introduce loss by coupling to two level fluctuators in the junction dielectric [164]. Here, we consider an intermediate range of junction areas, from 0.01 to $0.44\,\mu\mathrm{m}^2$.

We present measurements of 4-8 GHz superconducting resonant circuits embedded with Josephson junctions. The junction behaves as a nonlinear inductor, which is connected in series to a linear inductance and shunted by a capacitor (fig. C.1a). The junction also has loss, modeled as a resistor, and a self capacitance (fig. C.1c). Capacitive energy is divided between the junction capacitance, C_j, and the shunting capacitance, C_s, while inductive energy is divided between the Josephson inductance, L_j, and the stray inductance, L_s. Thus, a large area junction with a small shunting capacitance is required to observe loss in the junction dielectric over other sources of resonator loss. Likewise, a high inductive participation ratio $p_j = L_j/(L_s{+}L_j)$ is required to observe loss due to quasiparticles tunneling across the junction.

We characterize junction loss in two different ways: by measuring the quality factor of junction embedded quasi-lumped-element resonators with high p_j (fig. C.1a) and by measuring the relaxation time T_1 of transmon qubits dispersively coupled to 3D cavities (fig. C.1b). In both geometries, we measure shadow evaporated aluminum junctions with a wide range of critical currents and areas. We observe no dependance of loss on junction size, for areas as large as 0.44 μm^2, indicating that our resonator quality factors are not limited by dielectric loss in the junction capacitance. However, junction loss is highly dependent on the embedding geometry, which suggests the influence of non-equilibrium quasiparticles.

Samples were fabricated with a liftoff process on high resistivity silicon in a single electron

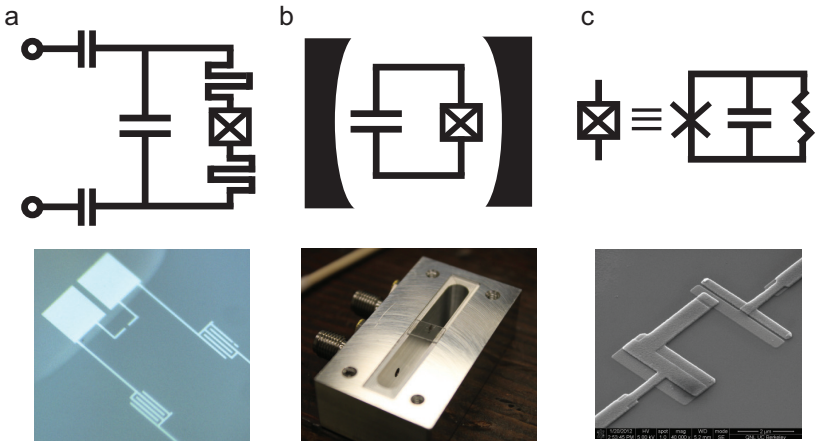

Figure C.1: a) Diagram (top) and optical image (bottom) of a Josephson junction embedded planar resonator. b) Diagram (top) and photograph (bottom) of a transmon qubit coupled to a three dimensional cavity. c) A Josephson junction can be modeled (top) as a nonlinear inductance in parallel with a geometric capacitance and a lossy resistor. (bottom) An SEM image of a typical shadow evaporated aluminum junction.

beam lithography step. Junctions were formed by double angle evaporation of Aluminum, with an intermediate oxidation step. Resonators and qubits were mounted on the mixing chamber stage of a dilution refrigerator and cooled to 25 mK. Resonators were wire bonded to a microwave launch covered with a copper lid, while qubits were enclosed in a 3D cavity. Samples were shielded by successive layers of superconducting and magnetic shields and probed via heavily attenuated coaxial lines.

Large critical current junctions (500 nA - 2.5 μA) were embedded in lumped element resonators based on single crystal silicon dielectric parallel plate capacitors [160] (fig. C.1a). Smaller critical current junctions (130-470 nA) were embedded in quasi-lumped element resonators with interdigitated capacitors. Both single junction and two junction SQUID samples were characterized. Two junction SQUID samples are treated as a junction of twice the area, whose inductance can be tuned by applying an external magnetic flux through the SQUID loop. Samples were probed in reflection using a vector network analyzer. We determine the the participation ratio p_j by fitting the flux dependent resonance frequency. As described in the previous appendix, we extract the internal quality factor Q_{int} of the resonator as a function of $\bar{n}$, the photon number inside the resonator when driven on resonance (fig. C.2).

In order to separate junction loss from other sources of resonator loss, samples were co-fabricated with linear resonators with similar design parameters, which consistently yielded Q values 150,000 for Si capacitors and $>$ 50,000 for interdigitated capacitors. Junction embedded resonators are lossier, with an average $Q_{int} \sim 5000$ at the lowest powers. Given the low participation of the junction capacitance in this geometry, it is likely that this loss

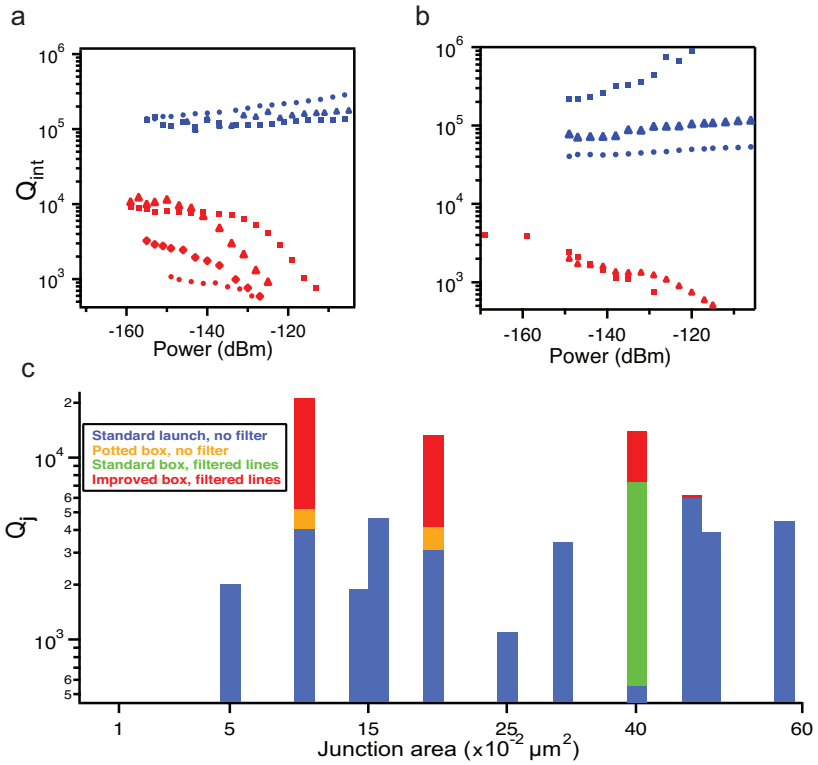

Figure C.2: a,b) Internal quality factor vs power for linear (blue) and Josephson-junction-embedded (red) resonators. Devices in panel (a) used single crystal silicon capacitors, and devices in panel (b) used interdigitated capacitors. c) Q_j vs junction area for junction-embedded resonators. The different colors represent repeated measurements of a given resonator with different shielding and filtering configurations.

is due to quasiparticles. Interestingly, Q_{int} decreases with increasing drive power well below a single photon excitation, where the resonator is still in the linear regime.

We then remeasured the samples while protecting from two possible sources of quasiparticles– stray infrared (IR) radiation [170] and high frequency noise on coaxial input and output lines. Following [171], we protected samples from IR radiation by potting the entire sample box in absorptive material (Ecosorb). We protected samples from high frequency noise with absorptive low pass filters connected directly to the input and output ports of the sample launch. We observed a slight improvement in Q_{int} from IR shielding (fig. C.2c, orange bars) and a more drastic improvement from filtering (fig. C.2c, red and green bars), increasing the average Q_{int} value by roughly a factor of 2.

We can further protect junctions from quasiparticles by confining them to a high quality 3D cavity, which acts as a narrow band filter. Junctions were connected to two paddles with a self-capacitance of $\sim$ 70 fF, forming a transmon qubit which couples capacitively to the TE_{101} mode of the cavity. The unloaded cavity has an internal Q of 2×10^6. Coupling to

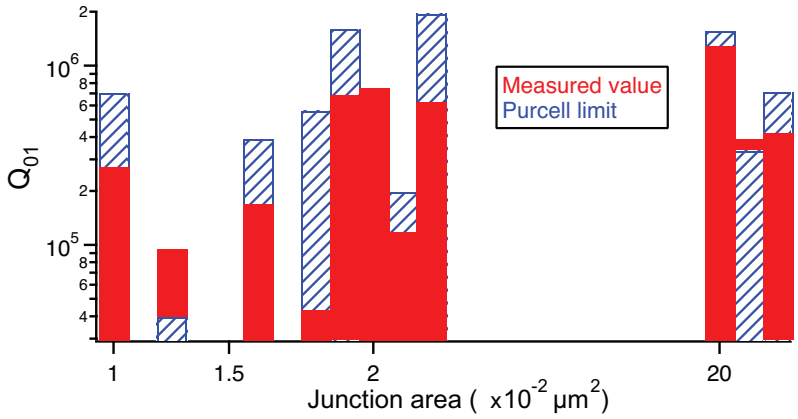

Figure C.3: T1 times versus junction area for 3D transmon samples (solid bars), compared with the single mode Purcell limit (dashed bars)

the external environment was set by the pin length of connectors on cavity input and output ports. The qubit was operated in the dispersive limit, such that the resonant frequency of the cavity is dependent on the state of the qubit. By monitoring the the phase of the cavity on resonance, we can extract information about the state of the qubit. We preform a standard T_1 measurement with a pulsed readout. As shown in figure C.3, T_1 times appear to be limited by cavity decay and generally fall within a factor of 2 of the single mode Purcell limit $T_p = \Delta^2/(\kappa g^2)$. Deviations from this limit can be attributed to the non-Lorentzian cavity line shape at large detunings.

Surprisingly, T_1 times do not decrease with increasing junction area for junctions as large as $0.25\,\mu m^2$, even as the junction capacitance approaches the value of the shunting capacitance. Assuming that $p_j \approx 1$, and attributing all loss to the junction capacitance, we can place an upper limit on the loss tangent of the junction dielectric given by $\tan\delta_j < c_j/[(c_s + c_j)\omega_{01}T_1]$. Based on the subset of qubits with the largest junction areas where the capacitive participation of the junction is roughly 15%, we place an upper limit of $\sim 5\times10^{-6}$ on the dielectric loss tangent of $\sim 0.25\ \mu m^2$ sized AlO_x diffused barrier Josephson junctions.

In conclusion, we observe no dependence of loss on junction area, and a strong dependence on embedding geometry and filtering, suggesting the influence of non-equilibrium quasiparticles. These results suggest that junction dielectric loss does not play a significant role in the energy relaxation of our transmon qubits and that filtering and embedding geometry are important considerations for improving the coherence of superconducting qubits.

Appendix D

Fitting routine for reflection resonators

Here, we present our procedure for fitting a reflection resonance to extract the internal quality factor Q_i, external quality factor Q_{ext}, and resonance frequency $f_0 \equiv \omega_0/2\pi$. The technique presented here only works for resonators with symmetric line shapes, and may be inaccurate for resonators coupled to a feed line. For other fitting techniques see, for example, reference [103] and [172]. Using a vector network analyzer, we measure the reflected signal as a function of excitation frequency, which is recorded with the real and imaginary parts interleaved. When properly scaled and rotated, the real part of the reflected signal is a Lorentzian function of the form

$$\text{Re}(f) = \frac{a}{(f - f_0)^2 + b}, \tag{D.1}$$

and the internal and external quality factors are given by

$$Q_i = \frac{\sqrt{b}\omega_0}{2b - a} \tag{D.2}$$

$$Q_{\text{ext}} = \frac{\sqrt{b}\omega_o}{a}. \tag{D.3}$$

To find the proper rotation of the real and imaginary parts of the reflected signal, we scan the rotation angle and fit $\text{Re}(f)$ to a Lorentzian for each angle. We take the angle with the smallest fitting error to be the proper rotation. See below for the Igor code that we use to implement this fitting procedure.

```
#pragma rtGlobals=1        // Use modern global access method.

//use this for fitting reflection resonator data
//requres real and imaginary waves
//12/1/10 SJW &KWM

function fitsmith(wavein,startf, stopf,ang) //operate on 1 d wave which as interleved real and imaginary pa
wave wavein                                  // set ang to 0 to rotate wave for best angle
variable startf, stopf,ang
wavestats/q wavein
make/o/n=(v_npnts/2) realwave, imwave
realwave = wavein[2*p]
imwave =  wavein[2*p+1]
setscale/i x,startf,stopf, realwave, imwave
fitQs(realwave,imwave,ang)
end

function fitQs(rein,imin,ang) //requires proper wave scaling, procedure rotates the wave by an amount step
//choosing the least squares fit, then scales and fits to a lorentzian and spits out q values
wave rein,imin
variable ang
duplicate/o rein reout, imout, fit2_reout, residwave
make/o/n=1 angwave
variable step =.05
make/o/n=(round(Pi/step)) resid
wave w_coef, w_sigma
if (ang==0)
    variable i
    for(i=0;i<round(Pi/step);i+=1)
        rot(rein,imin,step* i)
        CurveFit/q/NTHR=0 lor reout /D
        fit2_reout =  W_coef[0]+W_coef[1]/((x-W_coef[2])^2+W_coef[3])
        residwave=(fit2_reout-reout)^2
        resid[i] =sum( residwave)
    Endfor
    wavestats/q resid
    angwave[0]=step*x2pnt(resid,v_minloc)
else
    angwave[0]=ang
endif
rot(rein,imin,angwave[0])
CurveFit/q/NTHR=0 lor reout /D
reout/=-w_coef[0]
imout/=-w_coef[0]
Smooth 1, reout
CurveFit/q/NTHR=0 lor reout /D
variable a=w_coef[1]
variable b = w_coef[3]
variable fr =w_coef[2]
```

```
string result
variable qint = sqrt(b)*fr/(2*b-a)
variable  qinterr=sqrt(b*fr^2*w_sigma[1]^2/(2*b-a)^4  +(fr/(2*sqrt(b)*(2*b-a)) - 2*sqrt
variable qext= sqrt(b)*fr/a
variable qexterr=sqrt(b*fr^2*  W_sigma[1]^2/a^4  +fr^2*W_sigma[3]^2/(4*a^2*b))
make/o/n=5 resultwave
resultwave[0]=qint
resultwave[1] = fr
resultwave[2] = qinterr
resultwave[3] = qext
resultwave[4] = qexterr
End

function rot(re,im,ang)  //rotates the real and imaginary parts
wave re,im
variable ang
duplicate/o re reout, imout
reout=  real(Exp(ang*cmplx(0,1))*cmplx(re,im))
imout=  imag(Exp(ang*cmplx(0,1))*cmplx(re,im))
End
```