Quantum feedback for measurement and control

by

Leigh Martin

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:

Professor Irfan Siddiqi, Co-chair
Professor Birgitta Whaley, Co-chair
Professor Hartmut Häffner
Professor Eric Neuscamann

Summer 2019

Quantum feedback for measurement and control

Abstract

Quantum feedback for measurement and control

by

Leigh Martin

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Irfan Siddiqi, Co-chair

Professor Birgitta Whaley, Co-chair

The standard quantum formalism introduced at the undergraduate level treats measurement as an instantaneous collapse. In reality however, no physical process can occur over a truly infinitesimal time interval. A more subtle investigation of open quantum systems lead to the theory of continuous measurement and quantum trajectories, in which wave function collapse occurs over a finite time scale associated with an interaction. Within this formalism, it becomes possible to ask many new questions that would be trivial or even ill-defined in the context of the more basic measurement model. In this thesis, we investigate both theoretically and experimentally what fundamentally new capabilities arise when an experimental apparatus can resolve the continuous dynamics of a measurement. Theoretically, we show that when one can perform feedback operations on the timescale of the measurement process, the resulting tools provide significantly more control over entanglement generation, and in some settings can generate it optimally. We derive these results using a novel formalism which encompasses most known quantum feedback protocols. Experimentally, we show that continuous measurement allows one to observe the dynamics of a system undergoing simultaneous non-commuting measurements, which provides a reinterpretation of the Heisenberg uncertainty principle. Finally, we combine the theoretical focus on quantum feedback with the experimental capabilities of superconducting circuits to implement a feedback controlled quantum amplifier. The resulting system is capable of adaptive measurement, which we use to perform the first canonical phase measurement.

To my parents, my step parents, and my amazing sister Willow.

Contents

List of Figures

List of Tables

Acknowledgments

The fact that research is not an isolated endeavor is one of the most important lessons I learned from my time at Berkeley. The quality of one's research, and the joy of doing it are a direct result of one's collaborators, colleagues and friends. I feel incredibly fortunate to have worked with such a talented, curious and kind group of people. I first wish to thank my advisors, Irfan Siddiqi and Birgitta Whaley for taking a risk on me by supporting my joint work in theory and experiment. The chance to work with both research groups has been an amazing and irreplaceable opportunity. I wish to thank both of them for providing the perfect balance of encouragement and critique, and guidance and freedom. In my experimental work, I am deeply grateful for the guidance of and collaboration with Shay Hacohen-Gourgy and Emmanuel Flurin, who showed me the ropes of experimental work, helped me discover my flaws and strengths, and never turned down an chance to discuss a crazy idea (no matter how sure they were that it was wrong!). I am also indebted to Mollie Schwartz, who helped give me the opportunity to work in Irfan's group and introduced the field to me. Her warmth and encouragement made an enormous difference in embarking on a new path in research. In my theory work, I am indebted to the guidance and sharp intuition of Felix Motzoi, whose initial suggestion for a project carried me through a PhD's worth of theory research, as well as Mohan Sarovar, whose mentorship gave me confidence and stability. I also wish to thank Mahrud Sayrafi, Sissi Wang, Yitian Chen, Song Zhang and Yuxiao Jiang for the meetings all over Berkeley while we carried out our joint projects. I looked forward and enjoyed each and every one of these discussions, which consistently took us in exciting and unexpected directions. In my experimental work, I especially wish to thank Vinay Ramasesh, who taught me the importance of cordial competition and open communication, and William Livingston, who always helped me see the light in the darkness of challenge or my own stubbornness. Many of the general insights that I attempt to convey in this thesis are of their making. I greatly appreciate Machiel Blok's support, and all of the late afternoons spent tossing around ideas (also my apologies to Esther Blok for the countless times that I made Machiel late). I also wish to thank Sydney Schreppler, Kevin O'Brien, John Mark Kreikebaum, Andrew Eddins, David Toyli and Norman Yao for their support, collaborations and friendship. Finally, an enormous thank you to my parents, step parents, sister and friends in Berkeley, whose support was everything.

Chapter 1

The Building Blocks of Quantum Devices: Qubits, Harmonic Oscillators and Quantum Optics

1.1 Introduction

The ability to trap and control individual electrons, atoms and molecules turned the most bizarre thought experiments of quantum mechanics into tangible, thoroughly verified demonstrations. Once-contentious interpretational issues have been explored at the level of single, well-controlled degrees of freedoms. The double-slit experiment has been performed on single electrons sent and detected one-by-one, confirming wave-particle duality [6, 7]. Quantum superposition of increasingly macroscopic objects is becoming routine. In the last few years, an experiment spanning an entire university campus succeeded in violating Bell's inequality and many loopholes that could have preserved a more classical way of thinking[52]. In short, modern atomic, molecular and optical experiments have demanded a very different way of thinking about physics.

One of the strangest and most subtle differences between classical and quantum mechanics is measurement. Measurement bridges the conceptual and operational gap between wave functions and classical probabilities via the Born rule. At the most basic level, the Born rule is stated as an instantaneous collapse of the wave function into an eigenstate determined by the detection apparatus. In reality however, no process in nature is truly instantaneous. Atoms spontaneously emit via a continuous interaction with the electromagnetic vacuum, for example. Furthermore, no measurements are perfectly sharp, meaning that they do not determine precisely one measurement outcome with perfect certainty. Position measurements are limited by the de Broglie wavelength of the probe, and even discrete-outcome measurements like the readout of a digital memory are corrupted by some amount of noise. On the surface, neither of these inevitable constraints are captured by the Born rule's most basic statement.

(a)

(b)
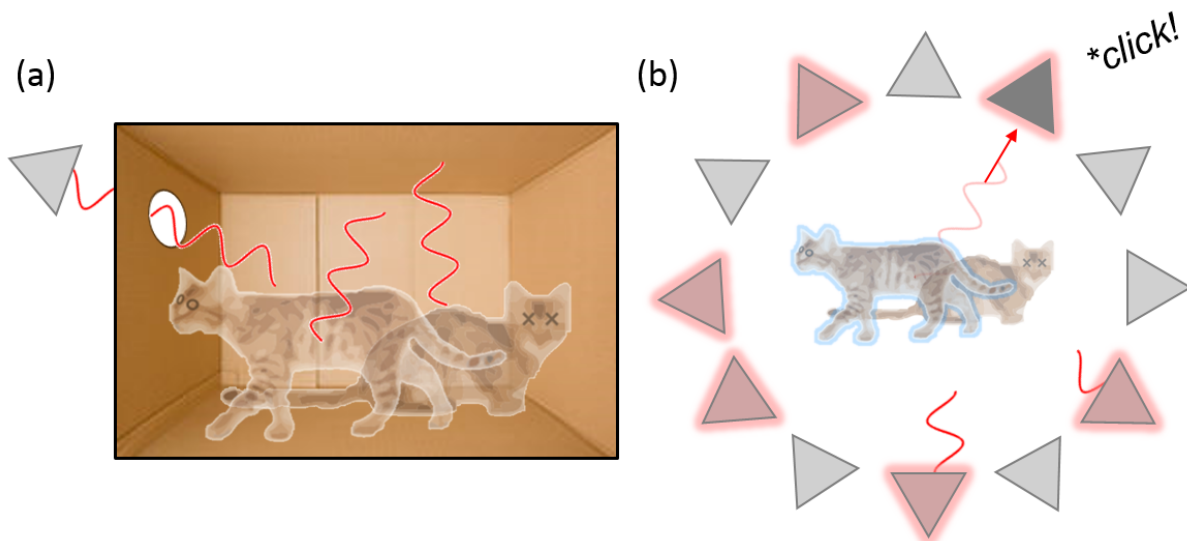
Figure 1.1: (a) An imperfectly isolated quantum system quickly shares information with its environment. An external detector typically acquires only a small fraction of the total available information, and hence sees no residual coherence during the measurement process. (b) In highly controlled systems such as superconducting circuits, one can arrange that the system interacts most strongly with our detector, allowing one to resolve wave function collapse as a continuous-time process.

When applied and interpreted carefully, the Born rule leads to a natural generalization of instantaneous projective measurements that meshes consistently with the above limitations. It predicts that wave function collapse is actually a continuous process, and hence wave function coherence is perturbed rather than destroyed over sufficiently short timescales. This modification forces the notion of a measurement rate dictated by the interaction strength between the system and its environment. The resulting dynamics are called quantum trajectories, which represent the finest-grained picture of a measurement. In most atomic, molecular and optical (AMO) systems, trajectories are exquisitely difficult to observe. The fundamental challenge is one of efficiency; the coupling strength of a quantum system to its environment is likely greater than that to a detector in our possession. For example, if we monitor the spontaneous emission from an atom with a localized detector, there is a good chance that we will miss the majority of emitted photons. Consequently, by the time we have enough information about the system to to say anything definitive, the environment at large has already collapsed the wave function for us.

In recent years, many foundational quantum phenomena have also been observed in artificial quantum systems like circuits. In much the same way that atomic energy levels consist of orbiting electrons, oscillating and circulating currents in electric circuits can also exhibit quantized energy levels. Such artificial atoms offer a greater range of control, as both their energy levels and their environment can be tailored. For superconducting circuits, they

can be made to interact most strongly with a 1-dimensional waveguide, so that all information emitted by the system can be collected by a single detector. Combined with innovations in quantum-limited amplifiers, these systems recently enabled the first experimental observation of quantum trajectories[90], though important experimental advances along these lines had already been achieved with Rydberg atoms[42].

With the ability to observe quantum trajectories comes many new experimental capabilities. Indeed, the full scope of these new capabilities is still as much a subject of theoretical discovery as experimental realization. The scope of opportunities is vast. Despite the apparent abstractness of quantum mechanics, the proposed applications of coherent quantum technology ranges from computation to astrophysical observation, and many derive from the ability to enhance measurement devices with uniquely quantum effects.

Quantum trajectories not only resolve measurement at its finest scales, but also provide the opportunity to act on a system in unique ways. This thesis asks what fundamentally new capabilities arise from continuous measurement, and experimentally implements two foundational examples. One basic example is the ability to measure non-commuting observables simultaneously. In the most basic statement of the Born rule, such a task is ill-defined. However realistic systems interact with their environments via multiple degrees of freedom, and hence this process is occurring constantly in nature.

The other new operation afforded by quantum trajectories is *quantum feedback*, whereby one acts upon a system during the measurement process. From a theoretical vantage point, we consider tasks that are commonly achieved probabilistically with measurement, such as entanglement generation, and show that they can be made to operate deterministically. In a few cases, we are able to prove that the resulting protocols are optimal. In the final experimental work, we apply feedback not to a system that we wish to control but rather to the measurement apparatus itself. By changing the measurement basis during wave function collapse, we implement the first canonical phase measurement.

Outline of this thesis Theses are a primary route of knowledge transfer within fields and outline the necessary background for research. To this end, we include somewhat more introductory material than average, in hopes that the majority of content becomes accessible to an undergraduate who has taken upper-division quantum mechanics. Chapters and sections begin with basic, general calculations and end with a qualitative overview of important material outside the scope of our presentation. To balance familiar and specialized material, we often give a slightly unconventional presentation when reviewing standard topics. Chapter 1 starts with basic quantum mechanics to introduce some concepts that are often skipped in an undergraduate course, like density matrices, frame transformation and the rotating wave approximation. Taking a slightly more information-based approach, we walk through the basics of qubit manipulation and control at a level that keeps the material relevant to all AMO-like systems. Exercises are included in the first three chapters as a way to flag concepts that are easily proven but divert from the flow of presentation. The interesting results are typically given in the problem statement, so the reader can view them as extensions of the text without working through them if desired.

Chapter 2 turns to the experimental system of this work, superconducting circuits. We

overview at a high level the essential elements of superconductivity that arise in the field, and then jump into the quantization of lumped-element and beyond-lumped-element circuits. After explaining the basic classes of qubits, we take our first foray into open quantum systems in our introduction to parametric amplifiers. Chapter 3 builds a foundation for the theoretical contribution of this thesis, describing weak measurement, quantum trajectories and decoherence. Chapter 4 describes our work on quantum feedback, outlining methods for designing quantum feedback protocols. We apply the formalism to the task of remote entanglement generation and show how one can make ordinarily stochastic processes deterministic in realistic systems.

Chapter 5 presents the first experimental work of this thesis. We explain the details of achieving high quantum efficiency measurements in circuit QED, an essential ingredient for quantum feedback. We then describe a novel circuit QED measurement technique, which we use to observe the dynamics of a system under simultaneous non-commuting measurements. As non-commuting observables are necessary to observe non-classical effects, the experiment brings a qualitatively new platform to the study of quantum foundations and their consequences, one that closely resembles the interactions found in nature. As a first step in this direction, we show theoretically and experimentally how the continuous disturbance induced by measurement is bounded by a modern reformulation of the Heisenberg uncertainty principle.

Chapter 6 fully merges the experimental and theoretical works, applying quantum feedback to a parametric amplifier to implement a canonical, or 'ideal' phase measurement. The broader significance of the result is that quantum feedback can turn a standard measurement device into something much more general, and hence able to observe a larger class of physical properties. Finally, in chapter 7 we return to the feedback protocols of chapter 4, proving optimality of a number of feedback protocol and discussing the implications.

1.2   Qubits and Harmonic Oscillators

There are two common approaches to quantum physics. The standard approach transitions from classical mechanics to the Schrödinger equation, emphasizing wave mechanics and only ending with atomic physics and other discrete systems. For review, we briefly detour through a less common but potentially more modern approach, which begins with quantum information in order to build up the framework, and then connects to atoms (and circuits) later. In order to avoid excessive abstraction, we focus on the dynamics that are commonly involved in experimental qubit characterization: Rabi oscillations, Ramsey oscillations and phenomena associated with the Jaynes-Cummings Hamiltonian. The advantage of the abstraction of this chapter is that the material is relevant to most physical implementations of a qubit, including trapped ions, cold neutral atoms, nitrogen vacancy centers, nuclear spins, quantum dots and of course superconducting circuits.

Qubits in abstraction Qubits are the smallest and most fundamental objects in quantum information. A qubit is any system that may be described by a two-dimensional Hilbert

space, either because it only has two available states (such as an electron spin) or because it is somehow prevented from leaving a two-dimensional subspace (for instance by only driving transitions between a single pair of states in a multi-level system).

$$|\psi_q\rangle = \alpha|0\rangle + \beta|1\rangle \qquad (1.1)$$

is a general qubits state, where α and β are complex and normalized such that $|\alpha|^2 + |\beta|^2 = 1$. According to the Born rule, if one measures an operator $\mathcal{O}$ (which for a qubit is a 2×2 Hermitian matrix), then the possible states after measurement are the eigenvectors of $\mathcal{O}|\lambda\rangle$, which are typically labeled by their corresponding eigenvalues λ. In general, the probability for a given measurement outcome λ to occur is then given by

$$P(\lambda|\psi_q) = |\langle\lambda|\psi_q\rangle|^2. \qquad (1.2)$$

Note that as P is unchanged if we replace $|\psi_q\rangle$ with $e^{i\phi}|\psi_q\rangle$, states that differ only by a global phase are physically indistinguishable.[1] A primary focus of this thesis is the physical process underlying measurement, which we describe beginning in chapter 3.

The direct relationship between wave functions and probabilities has lead many to suggest that at its heart, quantum mechanics is as much a generalization of mathematical probability theory as a physical theory. As we will see shortly, the inherent richness of complex numbers allow for interference between distinct possibilities via time (Hamiltonian) evolution. However the wave function by itself does not incorporate classical uncertainty, so strictly speaking, it is more a distinct case than a generalization of a probability. The most general case must include quantum systems in classically uncertain states. For instance, one might have randomly prepared states $|\psi_i\rangle$ with probabilities P_i such that $0 \leq P_i \leq 1$ and $\sum_i P_i = 1$. It is convenient to define the *density matrix* for this scenario as

$$\rho \equiv \sum_i P_i|\psi_i\rangle\langle\psi_i|, \qquad (1.3)$$

which again applies for any quantum system, not just a qubit. At first glance, one might think that the density matrix does not suffice to completely characterize a quantum state; given a density matrix, there may be many possible ways to decompose it in terms of pure states $|\psi_i\rangle$ and probabilities P_i. Surprisingly, if two probabilistic ensembles of pure states yield the same density matrix, then they are physically indistinguishable! To show this, we observe how the density matrix allows one to restate the Born rule so as to take into account the randomness of state preparation. Noting that $|\lambda\rangle\langle\lambda|$ is the density matrix associated with an eigenstate of $\mathcal{O}$, we define the measurement outcome probability as

$$P(\lambda|\rho) \equiv \mathrm{Tr}[|\lambda\rangle\langle\lambda|\rho] = \sum_i P_i\langle\lambda|\psi_i\rangle\langle\psi_i|\lambda\rangle = \sum_i P_i P(\lambda|\psi_i) \qquad (1.4)$$

---

[1]One also needs the fact that $|\psi_q\rangle$ and $e^{i\phi}|\psi_q\rangle$ evolve identically under the Schrödinger equation, which follows from its linearity.

where we have used the cyclic property of the trace ($\mathrm{Tr}[AB] = \mathrm{Tr}[BA]$) and the fact that $\mathrm{Tr}[a] = a$ if a is a scalar. Eq. (1.4) is just a weighted average of the measurement outcome probabilities for each possible state preparation $|\psi_i\rangle$, so the right hand side reproduces the expected probability based on Eq. (1.2). A similar argument shows that $\langle \mathcal{O} \rangle = \mathrm{Tr}[\mathcal{O}\rho]$, where the expectation value of an operator is defined over its eigenvalues *i.e.* as $\langle \mathcal{O} \rangle \equiv \sum_\lambda P(\lambda)\lambda$. As the measurement probabilities are uniquely determined by ρ, ρ completely characterizes the state of a system.

Exercise 1 *Show that $\rho = \rho^\dagger$, that its trace always equals 1, and that the eigenvalues of the density matrix are greater than or equal to 0 (i.e. that it is positive semidefinite). Any matrix satisfying these three properties is a valid density matrix.*

Exercise 2 *Pure states: show that if $\mathrm{Tr}[\rho^2] = 1$, then ρ is rank 1 and thus can be written in the form $|\psi\rangle\langle\psi|$. If $\mathrm{Tr}[\rho^2] < 1$, ρ is referred to as a mixed state.*

Returning to the specific example of a qubit, the most general qubit state is conveniently expressed in terms of the Pauli matrices as

$$\rho_q = \frac{\sigma_0 + x\sigma_x + y\sigma_y + z\sigma_z}{2} \tag{1.5}$$

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

where σ_0 is the 2×2 identity matrix and x, y and z must be real to keep ρ Hermitian. Using the identity $\mathrm{Tr}[\sigma_i\sigma_j] = 2\delta_{i,j}$, one can show the following useful facts:

1. $x = \langle\sigma_x\rangle \equiv \mathrm{Tr}[\sigma_x\rho]$, and likewise for y and z.

2. The purity of ρ_q as per exercise 2 is $(1 + r^2)/2$ where $r = \sqrt{x^2 + y^2 + z^2}$. Thus ρ is pure if and only if $r = 1$.

3. From 2, it follows that ρ_q is a valid density matrix as per exercise 1 if and only if $0 \leq r \leq 1$.

As the Pauli matrices together with σ_0 span the vector space of 2×2 Hermitian matrices, every possible qubit state can be represented uniquely by specifying x, y and z. These facts justify the representation of a qubit as a 3-vector, called the *Bloch vector*; x, y and z are measurable and fully specify ρ_q. In many physical systems, the Bloch vector corresponds to a physical orientation in space, such as the magnetic dipole moment of an electron. Although the Pauli matrices generalize nicely to the Gell-Mann matrices for larger systems, the statements analogous to 2 and 3, even for a qutrit (3-level system), are somewhat more complicated (see Eqs (13-18) of [45]).

Schrödinger equation for wave functions and density matrices The Bloch vector offers a convenient visual representation for intuiting the dynamics of a qubit. A physical

qubit has an associated Hamiltonian, which governs its dynamics. For an isolated qubit, the only term in the Hamiltonian is an intrinsic energy splitting between the ground and excited states, conventionally taken to be along the σ_z axis *i.e.* $H_0/\hbar = \omega_q \sigma_z/2$. The Schrödinger equation is

$$\frac{\partial |\psi\rangle}{\partial t} = -\frac{i}{\hbar} H_0 |\psi\rangle \implies \frac{\partial \rho}{\partial t} = -\frac{i}{\hbar}[H_0, \rho]. \tag{1.6}$$

The version for density matrices may be derived by taking the time derivative of Eq. (1.3). As the Hamiltonian is already diagonalized, the Schrödinger equation decouples into two first order differential equations, one for α and one for β. The general solution for an initially pure state is

$$|\psi_q(t)\rangle = \alpha_0 e^{-i\omega_q t/2}|0\rangle + \beta_0 e^{i\omega_q t/2}|1\rangle \implies \rho(t) = \begin{pmatrix} |\alpha_0|^2 & \alpha_0 \beta_0^* e^{-i\omega_q t} \\ \alpha_0^* \beta_0 e^{i\omega_q t} & |\beta_0|^2 \end{pmatrix} \tag{1.7}$$

where we have taken slightly unfortunate but ubiquitous convention that $|1\rangle$ designates the ground state. Thus under the evolution of H, the relative phase between $|0\rangle$ and $|1\rangle$ varies, but the populations in each state do not. In the density matrix, these observations are apparent in the off-diagonal and diagonal elements, respectively.

The dynamics of the Bloch vector are shown in Fig. 1.2a. The Bloch vector precesses around the σ_z axis at a frequency ω_q. The Bloch vector of a mixed state is the average Bloch vector of each pure state in the underlying statistical ensemble, so the result applies to mixed states as well. The alignment between the rotation axis and the Hamiltonian illustrates another useful property of the Bloch vector. As the Hamiltonian of an isolated qubit may always be expressed in terms of the Pauli matrices, we can also associate to it a vector. As is evident from the solution Eq. (1.7), the vector associated with H_0 is simply the rotation axis. Furthermore, the vector associated with ρ and H_0 transform identically under a change of (Hilbert space) basis, so a solution for $H \propto \sigma_z$ is in fact a general solution for all time-independent Hamiltonians. For example, $H \propto \sigma_y$ would generate rotations about the y axis instead.

Rabi oscillations and rotating frames To control the state of a qubit, and in particular to change the populations in $|0\rangle$ and $|1\rangle$, we need to generate rotations about an axis other than σ_z. To do so, we must somehow couple our qubit to the outside world. The simplest control method is to drive oscillations with a classical control field. Most physical qubits couple to external fields via some kind of dipole moment. Electron spins couple to magnetic fields, atomic energy levels to electric fields etc. The resulting Hamiltonian is called the Rabi model, and the dynamics are called Rabi oscillations.

It is convenient to express the dipole operator σ_x in terms of the ladder operator $\sigma = (\sigma_x - i\sigma_y)/2$, which satisfies $\sigma^2 = 0$, $\sigma|0\rangle = |1\rangle$, and $\sigma^\dagger|1\rangle = |0\rangle$.[2] Parameterizing the unitless

---

[2]We will justify σ_x as a dipole operator in chapter 2. Why we single out σ_x and not σ_y is a matter of convention, but nevertheless carries some physical meaning. We will see how this works when we consider physical qubits in the following chapter.
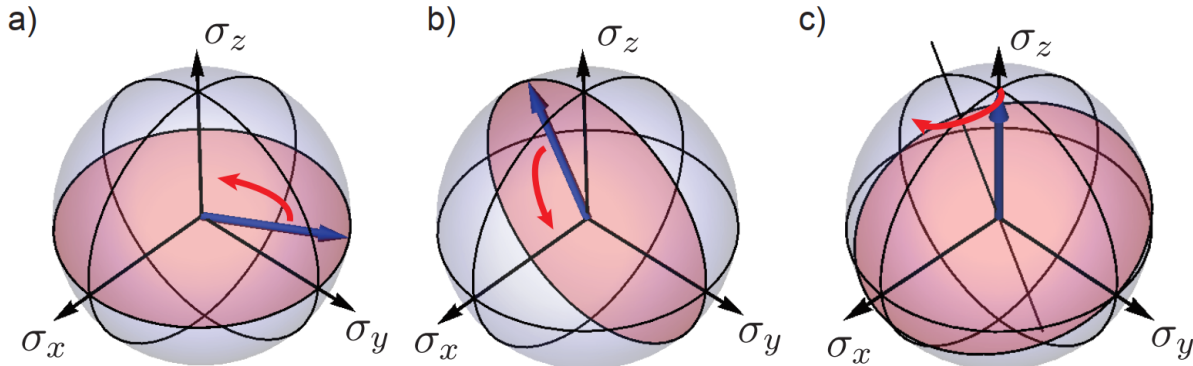
Figure 1.2: (a). Precession of a qubit in the lab frame. The direction of rotation may be obtained by the right-hand rule, and the angular frequency is ω_q. (b). Rabi oscillations about the σ_x axis ($\delta = 0$) in the rotating frame. The angular frequency is $\Omega_R = g\alpha_0/2$. (c). Rabi oscillations for a detuned drive ($\Delta > \Omega_R$), shown in the frame of the frame of the classical drive field ω_d.

classical control field as $\alpha(t)$, the qubit Hamiltonian with a dipolar coupling is

$$H/\hbar = \omega_q \frac{\sigma_z}{2} + g\alpha(t)(\sigma + \sigma^\dagger) \tag{1.8}$$

where g parameterizes the coupling strength.

Typically $g\alpha \ll \omega_q$, so the first term dominates the dynamics. However we understand the dynamics induced by the first term already and would like to focus on the second term. The first term can be eliminated by going into a *rotating frame*, which is similar to a rotating coordinate system in classical mechanics. To do so, we define a transformed state $|\tilde{\psi}(t)\rangle = U(t)|\psi(t)\rangle$, where $U(t)$ is a time-dependent unitary matrix that acts as a basis transformation. Substitution of $|\psi(t)\rangle = U^\dagger(t)|\tilde{\psi}(t)\rangle$ into the Schrödinger equation yields

$$\boxed{\begin{aligned}\frac{\partial|\tilde{\psi}(t)\rangle}{\partial t} &= -\frac{i}{\hbar}\tilde{H}|\tilde{\psi}(t)\rangle \\ \tilde{H} &= U(t)HU^\dagger(t) - i\hbar U(t)\dot{U}^\dagger(t) \\ &= e^{-iXt}He^{iXt} + \hbar X\end{aligned}} \tag{1.9}$$

where in the last line we have defined $U(t) = \exp(-iXt)$, which is guaranteed to be unitary if X is Hermitian. By substitution into the original Schrödinger equation, one can show that $U(t)|\psi(0)\rangle$ is a solution to the Schröodinger equation when $H/\hbar = X$, so that $U(t)$ acts as a time evolution operator. The above transformation is a key method with many applications, and we will use it frequently. The first term of $\tilde{H}$ acts as a basis transformation of the Hamiltonian, while the second term is somewhat analogous to a Coriolis potential.

Intuitively, we wish to transform into a basis that corotates with the dynamics induced by H_0, which is accomplished by setting $X = -H_0/\hbar$. As H_0 is diagonal, it is straightforward
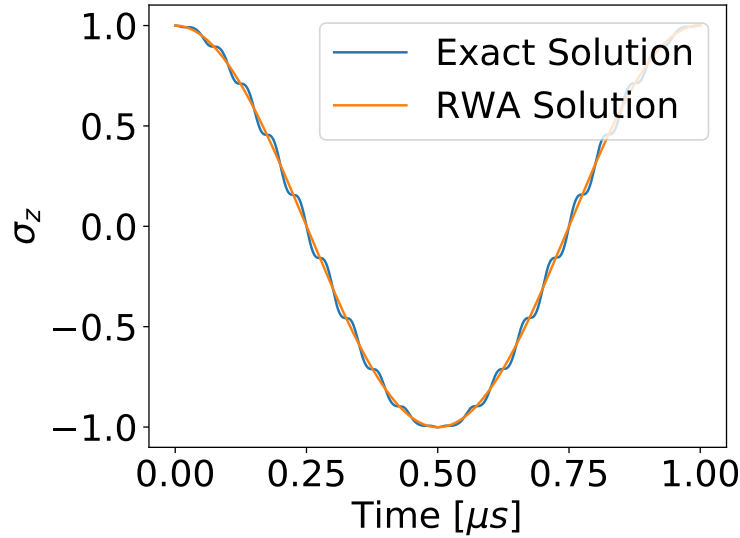
Figure 1.3: A plot of the z component of the Bloch vector vs. time with $\tilde{H}$ and $\tilde{H}_{\text{RWA}}$. The exact solution oscillates rapidly around the approximate solution but never deviates far from it. Parameters for the graphic are $\Omega_R/2\pi = 0.5$ MHz and $\omega_q/2\pi = 20$ MHz. In most realistic systems, the separation in scales between Ω_R and ω_q is even greater.

to exponentiate X (for example via a Taylor series) and calculate $\tilde{H}$

$$\tilde{H}/\hbar = g\alpha(t)(\sigma e^{-i\omega_q t} + \sigma^\dagger e^{i\omega_q t}) \tag{1.10}$$

To control the qubit, one can create a resonant interaction by letting $\alpha(t) = \alpha_0 \cos(\omega_d t + \delta)$ with the drive frequency $\omega_d = \omega_q$

$$\tilde{H}/\hbar = \frac{g\alpha_0}{2}(\underbrace{\sigma e^{i\delta} + \sigma^\dagger e^{-i\delta}}_{\text{co-rotating terms}} + \underbrace{\sigma e^{-2i\omega_q t - i\delta} + \sigma^\dagger e^{2i\omega_q t + i\delta}}_{\text{counter-rotating terms}}). \tag{1.11}$$

The Schrödinger equation is not simple to solve if we consider all four terms. Fortunately, solution is straightforward when only the co-rotating terms are considered, which is called the *rotating wave approximation*. In Fig. 1.3, we plot numerical solutions to the Schrödinger equation with and without counter-rotating terms, which agree very well.

To see why the counter-rotating terms contribute negligibly to the dynamics, first consider the solution in the absence of the counter-rotating terms, which we call $|\tilde{\psi}_{\text{RWA}}(t)\rangle$. $|\tilde{\psi}_{\text{RWA}}(t)\rangle$ changes on the time scale of $1/\alpha g$, which is slow compared to ω_q. To see if it is a good approximate solution, we compute the first order change in the wave function due to the counter-rotating terms by integrating the Schrödinger equation with only those terms present (*i.e.* using time-dependent perturbation theory). The result is

$$|\tilde{\psi}(t)\rangle - |\tilde{\psi}_{\text{RWA}}(t)\rangle \approx \int_0^t \frac{g\alpha_0}{2} \left[\sigma e^{-2i\omega_q t' - i\delta} + \sigma^\dagger e^{2i\omega_q t' + i\delta} \right] |\tilde{\psi}_{\text{RWA}}(t')\rangle dt'. \tag{1.12}$$

As $|\tilde{\psi}(t)\rangle$ is very nearly constant over the timescale $1/\omega_q$, the integrand oscillates rapidly and therefore the integral vanishes to good approximation. Thus adding in the counter-rotating terms has little effect on the solution, and we may replace $\tilde{H}$ with

$$\tilde{H}_{\text{RWA}}/\hbar = \frac{\Omega_R}{2}\sigma_\delta \tag{1.13}$$
$$\sigma_\delta \equiv \sigma e^{i\delta} + \sigma^\dagger e^{-i\delta} = \sigma_x \cos(\delta) + \sigma_y \sin(\delta)$$

where $\Omega_R = g\alpha_0$ is called the Rabi frequency. As $\tilde{H}_{\text{RWA}}$ is equivalent to H_0 up to a basis change, it still induces oscillations, except this time about the σ_δ axis. The dynamics are shown in Fig. 1.2b.

Exercise 3 *In the case that the control field is detuned from the qubit resonance frequency, one can make $\tilde{H}_{RWA}$ time independent by instead choosing $X = -i\omega_d\sigma_z/2$. Show that now $\tilde{H}_{RWA}/\hbar = \Delta\sigma_z/2 + \Omega_R\sigma_\delta/2$, where $\Delta = \omega_q - \omega_d$. Diagonalize this Hamiltonian to show that the Rabi frequency is $\Omega_R = \sqrt{\Omega_R^2 + \Delta^2}$. See Fig. 1.2c.*

Exercise 4 *Using the solution to exercise 3, let $|\tilde{\psi}(0)\rangle = |1\rangle$ and suppose that Δ sweeps slowly from a large negative value to a large positive value. What is the final state? Recall (or read about) adiabatic theorem.*

It is interesting to pause and contrast these dynamics to those of a classical bit. If a classical bit is prepared in an equal mixture of 0 and 1, no reversible transformation can take it to a definite state. By comparison, a qubit in the state $(|1\rangle + |0\rangle)/\sqrt{2}$ is indistinguishable from a classical mixture if one performs σ_z measurements, but can be converted to $|1\rangle$ (or $|0\rangle$) alone under the action of $\tilde{H}_{\text{RWA}}$! Such a transformation would map $|1\rangle \to (|1\rangle + |0\rangle)/\sqrt{2}$ and $|0\rangle \to (|1\rangle - |0\rangle)/\sqrt{2}$, which interfere to yield $|1\rangle$. At the level of quantum information, this interference is no different than the interference of a particle travelling through a pair of slits to yield a diffraction pattern. The Hilbert space is just smaller.

Driving Rabi oscillations is a primary technique in experimental quantum information. Under the evolution of $\tilde{H}_{\text{RWA}}$, the qubit rotates endlessly about the σ_δ axis, but fortunately, the dynamics are still exactly solvable when the drive amplitude α_0 is time-dependent. The exact same rotation occurs, except at a time-dependent rate. This is most easily seen in basis that diagonalizes the Hamiltonian $H \to \Omega_R(t)/2\sigma_z$. The solution to the Schrödinger equation becomes

$$|\tilde{\psi}_{\text{RWA}}(t)\rangle = \alpha_0 \exp\left(-i\int_0^t \Omega_R(t')dt'/2\right) + \beta_0 \exp\left(i\int_0^t \Omega_R(t')dt'/2\right), \tag{1.14}$$

which can be verified by direct substitution into the Schrödinger equation. Therefore the qubit's rotation angle accumulates proportionally to the drive strength. Eq. (1.14) is still of course only an approximate solution for $\tilde{H}$, and the rotating wave approximation can fail if $\Omega(t)$ changes too rapidly. If $\alpha_0(t)$ changes significantly over a period of $2\omega_q$, we can no

longer guarantee that the integral Eq. (1.12) will vanish. Although a closed-form solution is not be available for arbitrary time dependence of δ and ω_d, the intuition that $\tilde{H}$ acts as a rotation always holds, offering a way to induce arbitrary qubit dynamics. Rabi oscillations measured on a real superconducting qubit are shown in Fig. 1.4a.

Qubit pulses and characterization The above result gives us the tool to induce arbitrary time dependence on the state. However, just as classical digital computation is simpler to work with than analog, quantum information tasks also simplify if we impose discreteness. One typically uses Rabi oscillations to implement only $\pi/2$ and π rotations. The obvious way to implement a π pulse is to set Ω_R to some maximum value and then turn it off again after the correct time interval. The resulting boxcar pulse is shown in Fig. 1.4b. However, due to the presence of sharp edges, the Fourier transform of this pulse has support (is nonzero) out to very high frequencies, which leads to several difficulties in practice. In the present model, the only issue would be a breakdown of the rotating wave approximation. Additional issues arise in realistic settings. Firstly, physical qubits often contain other unused energy levels, and a broadband pulse can easily drive undesired transitions. Furthermore, if $\alpha(t)$ represents an electromagnetic signal travelling through a medium, a wide-bandwidth pulse is more susceptible to distortion. A better choice would be the Gaussian pulse envelope also shown in Fig. 1.4b. Although it also has support out to high frequencies, the spectral power decays faster for frequencies away from ω_d (exponential decay as opposed to $\mathcal{O}(1/\omega)$). Furthermore, the pulse bandwidth may be reduced arbitrarily by increasing the pulse duration.

π and $\pi/2$ pulses suffice for most qubit characterization, which is an integral part of experimental quantum information science. The two most basic characterization tools aside from Rabi oscillations are T_1 and Ramsey oscillation measurements. T_1, the excited state lifetime, is measured by applying a π pulse to a qubit and waiting for it to decay back to its ground state. Just as in radioactive processes, the excited state probability typically follows an exponential decay. An experimental measurement of T_1 is shown in Fig. 1.4c. A general and widely applicable model of qubit decay is presented in chapter 3.

Ramsey measurements are a somewhat more subtle but also more useful tool, as they allow measurement of three quantities: T_2 and T_2^*, which characterize how well the qubit maintains a superposition, and the precise qubit frequency. As we use Ramsey measurements to determine the precise qubit frequency, we should analyze the case where ω_q is not already known to the experimenter, and thus the drive frequency ω_d may not exactly match it. We saw in exercise 3 that the rotating frame Hamiltonian in this case is[3]

$$\tilde{H}_{\mathrm{RWA}}/\hbar = \Delta\sigma_z/2 + \Omega_R(t)\sigma_\delta/2. \tag{1.15}$$

The sequence for a Ramsey experiment consists of two $\pi/2$ pulses with a variables waiting time τ in between and a measurement at the end, as shown in Fig. 1.4d. We assume that

---

[3]Eq. (1.15) is written in the frame of ω_d. One can also write it in the frame of ω_q, in which case the $\Delta\sigma_z/2$ term does not appear but δ becomes time dependent. This picture of a time-varying rotation axis is equally valid and intuitive for understanding Ramsey oscillations.
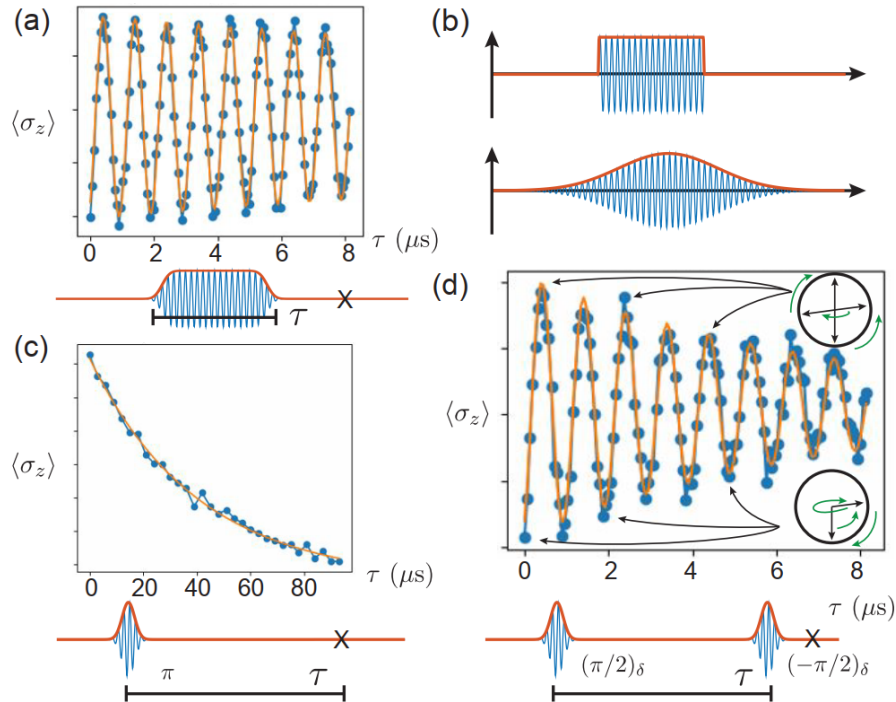
Figure 1.4: (a) Rabi oscillations. A best-fit line and the corresponding experimental pulse sequence implemented to measure them. To measure a single point on this graph, one drives the qubit for a duration τ and then measures in the σ_z basis. Each point is an average $\langle \sigma_z \rangle$ over many runs. Measurement occurs at the 'X'. (b) The electric field profile and envelope of a square and a Gaussian pulse. (c) Similar to (a) but a T_1 measurement. (d) Ramsey oscillation and the associated pulse sequence. Inset Bloch spheres visualize the qubit dynamics at various points on the trace. The oscillation frequency is set by the qubit-pulse detuning, and the exponential decay is the coherence time T_2^*. Best fit indicates an exponential decay time of 7.3 μs.

$\Omega_R(t)$ is large compared to $\Delta \equiv \omega_q - \omega_d$ for most of the duration of the pulse. Thus during the pulses, Eq. (1.15) is well approximated by Eq. (1.13) and the pulses implement the same rotations that they would on resonance.

During the waiting interval τ, the qubit precesses around the σ_z axis at the mismatch frequency Δ. When $\tau = 0$, the $\pi/2$ pulses occur one after the other and qubit ends in the excited state. The same is true if $\tau\Delta$ is an integer multiple of 2π, as illustrated in Fig. 1.4d. On the other hand, if $\tau\Delta$ is an odd multiple of π, then the second pulse rotates the qubit from the equator to the ground state instead. In general, the excited state population oscillates at a rate Δ, allowing us to measure Δ with high precision. Note that the oscillations only occur because during the waiting interval, the cosine term in $\alpha(t) = \alpha_0(t) \cos(\omega_d t + \delta)$ continues to oscillate at the same frequency ω_d. Experimentally, this means that we need a classical

oscillator, called a *local oscillator* to keep time throughout each run of the experiment, even when we are not applying pulses.

Again we have only considered an isolated qubit, but like T_1 measurements, much of their utility applies primarily to imperfect qubits. The primary imperfection in most qubits manifests as an apparent instability in ω_q. For example, as both electron spins and loops of current couple to magnetic fields, many types of qubits, from superconducting qubits to trapped ions have their frequencies changed by an external magnetic field. A fluctuating external field will cause the qubit to rotate randomly about the σ_z axis, so that the average Bloch vector gets shorter for long τ (see exercise 5). T_2^* is the exponential decay time of the Bloch vector length, which manifests as decay of the Ramsey oscillations. Such a decay is visible in Fig. 1.4d. As measurement destroys coherence, Ramsey oscillations also decay if an external system acquires information about the qubit's energy.

Exercise 5 *Semiclassical model of Ramsey oscillations: Show that if Δ is constant but randomly drawn from a zero mean Gaussian distribution with variance σ for each measurement, then the off-diagonal elements of ρ and hence the Ramsey profile decay as $e^{-\sigma^2 t^2}$. If instead Δ changes during each measurement, then the relative phase between $|0\rangle$ and $|1\rangle$ undergoes a random walk. In this case, show that the Ramsey decay is now exponential. What happens in each of these cases if you add an echo pulse? For both parts, use the fact that $\langle e^{i\phi}\rangle = e^{-\langle\phi^2\rangle}$ if ϕ is a zero-mean Gaussian distributed random variable.*

Harmonic oscillators Having introduced the basics of qubits, we now turn to the other extreme of infinite dimensional systems, specifically the harmonic oscillator. The harmonic oscillator often serves as a bridge between classical and quantum mechanics. Although one initially develops it to treat a particle in a parabolic potential, it is ubiquitous in quantum mechanics, also modelling LC circuits, phonons, and even photons[**39**], as we will see in the following chapters.

Focusing on a particle for concreteness, the wave function now forms a complex probability distribution over the position variable x. States of definite position are Dirac delta functions $|\psi_{x_0}(x)\rangle \propto \delta(x - x_0)$, which are eigenstates of the position operator $\hat{x} = x$ (loosely speaking). The de Broglie relation $p = \hbar k$ motivates identifying plane waves $|\psi_k(x)\rangle \propto e^{ikx}$ as states of definite momentum, which in turn motivates the definition of the momentum operator $\hat{p} = -i\hbar\partial/\partial x$ (as $\hat{p}|\psi_k(x)\rangle = p|\psi_k(x)\rangle$). With these operators in hand, it is natural to take the Hamiltonian of the harmonic oscillator to be $H_0 = p^2/2m + kx^2/2$, where $k = m\omega^2$ is the spring constant defining the parabolic potential and ω is the classical oscillation frequency (we drop the hats to avoid clutter).

The time-independent Schrödinger equation may be solved with the help of Hermite polynomials, and the ground state of H_0 is a Gaussian with variance $\sqrt{\hbar/m\omega}$. However, Dirac introduced a much simpler algebraic method for finding the energy spectrum. Introducing the latter operator

$$a \equiv \frac{m\omega}{2\hbar}\left(x + \frac{i}{m\omega}p\right) \qquad (1.16)$$

one can express the Hamiltonian in terms of $a^\dagger a$

$$a^\dagger a = \frac{m\omega}{2\hbar}\left(x^2 + \frac{p^2}{(m\omega)^2} + \frac{i}{m\omega}[x,p]\right) = \frac{1}{\hbar\omega}H_0 - \frac{1}{2} \tag{1.17}$$

$$\implies H_0 = \hbar\omega\left(a^\dagger a + \frac{1}{2}\right)$$

where we have used the fact that $[x,p] = i\hbar$. Using this commutation relation, one can also show that

$$[a, a^\dagger] = 1 \tag{1.18}$$
$$[a, a^\dagger a] = a$$
$$a^\dagger|n\rangle = \sqrt{n+1}|n+1\rangle$$
$$a|n\rangle = \sqrt{n}|n-1\rangle$$

where $|n\rangle$ is the eigenvector of $a^\dagger a$ with non-negative integer eigenvalue n. These eigenvectors are also eigenvectors of the Hamiltonian, and are in one-to-one correspondence with the Hermite polynomial solutions. The associated energies are $\hbar\omega(n+1/2)$. Notice that the ladder operators a are just like the ladder operators σ for a qubit. Both raise and lower a state among different energy levels.

Exercise 6 *If you need the review, derive Eqs. (1.18) and use them to argue that n is an integer greater than zero. Recall that $\langle\psi|a^\dagger a|\psi\rangle \geq 0$ for all $|\psi\rangle$ as it is the square of the state $a|\psi\rangle$.*

The harmonic oscillator eigenstates have progressively larger variance in x and p (which correspond to **E** and **B** for photons, Q and I for an LC oscillator or lattice displacement for phonons) as we go up in energy. However, these states are also time-independent and have zero mean in x and p, which does not line up with our classical intuition of a vibrating mass. To recover these classical states, we study the effect of a time-dependent force by adding $F(t)x = \epsilon_d(t)(a + a^\dagger)$ to the Hamiltonian. The problem is exactly analogous to the driven qubit problem that we have already encountered, and is solved by similar methods. We will make use extensive use of the Baker-Campbell-Hausdorff (BCH) formula and its corollaries, which we state upfront for future reference

$$e^X e^Y = \exp\left(X + Y + \frac{1}{2}[X,Y] + \frac{1}{12}([X,[X,Y]] + [Y,[Y,X]]) + ...\right) \tag{1.19}$$

$$e^{X+Y} = e^X e^Y \exp\left(-\frac{1}{2}[X,Y]\right)\exp\left(\frac{1}{6}(2[Y,[X,Y]] + [X,[X,Y]])\right) ...$$

$$e^X Y e^{-X} = Y + [X,Y] + \frac{1}{2!}[X,[X,Y]] + \frac{1}{3!}[X,[X,[X,Y]]] + ...$$

Subsequent terms involve fourth and higher order commutators of X and Y. The higher order terms are only easily guessed for the last equation. The first identity is the standard BCH formula. The second is called the Zassenhaus formula, which forms the basis for a quantum simulation algorithm called the Trotter expansion. Although they look formidable, they often truncate or lead to simple closed-form expressions, particularly for the harmonic oscillator.

To treat the driving term, we use Eq. (1.9) just as before, setting $X = -H_0/\hbar$ in that formula. The rotating frame Hamiltonian is

$$\tilde{H} = e^{iH_0 t/\hbar} H e^{-iH_0 t/\hbar} + \hbar X = H_0 - H_0 + \epsilon_d(t)(a + a^\dagger) \tag{1.20}$$
$$= \epsilon_d(t)(ae^{-i\omega t} + a^\dagger e^{i\omega t})$$

where we have used the third BCH formula to evaluate

$$\exp(i\omega a^\dagger a t) a \exp(-i\omega a^\dagger a t) = \sum_{n=0}^\infty \frac{1}{n!} \underbrace{[i\omega a^\dagger a, [i\omega a^\dagger a, [...[i\omega a^\dagger a, a]...]}_{n} \tag{1.21}$$
$$= \sum_n \frac{(i\omega t)^n}{n!}(-1)^n a = ae^{-i\omega t}.$$

A simple solution arises because the nested commutators like $[a^\dagger a, [a^\dagger a, a]] = -[a^\dagger a, a] = a$ always simplify to $\pm a$. (Note that we could have also used this formula to calculate $e^{i\omega_q \sigma_z t/2} \sigma e^{-i\omega_q \sigma_z t/2}$ in Eq. (1.10)).

Setting $\epsilon_d(t) = 2\epsilon_0 \cos(\omega t + \delta)$ again yields a resonant interaction

$$\tilde{H}_{\text{RWA}} = \epsilon_0(ae^{i\delta} + a^\dagger e^{-i\delta}) \tag{1.22}$$

where we have dropped counter-rotating terms for the same reason that applied in the qubit case. Solving for the driven dynamics is only slightly more involved than the analogous qubit Hamiltonian Eq. (1.13) with the help of the BCH formulas. If the system starts in its ground state $|0\rangle$, then the state after time t is formally

$$|\psi(t)\rangle = U(t)|0\rangle = \exp(-iH_0 t/\hbar)|0\rangle \tag{1.23}$$
$$= \exp(\alpha a^\dagger - \alpha^* a)|0\rangle$$
$$\alpha \equiv -i\epsilon_0 e^{-i\delta} t/\hbar$$

We have dropped tildes, since we will only work in the rotating frame through the rest of this section. We recognize a potential instance to apply the second BCH formula (the Zassenhaus formula), letting $X = \alpha a^\dagger$ and $Y = \alpha^* a$. As $[a^\dagger, a] = -1$, all higher-order commutators vanish, leaving us with only the first three terms

$$|\psi(t)\rangle = e^{\alpha a^\dagger} e^{-\alpha^* a} e^{-|\alpha|^2/2}|0\rangle = e^{-|\alpha|^2/2} e^{\alpha a^\dagger}|0\rangle \tag{1.24}$$
$$= e^{-|\alpha|^2/2} \sum_n \frac{\alpha^n}{n!} a^{\dagger n}|0\rangle = e^{-|\alpha|^2/2} \sum_n \frac{\alpha^n}{\sqrt{n!}}|n\rangle$$
$$\equiv |\alpha\rangle$$

The fact that $a|0\rangle = 0$ leads to substantial simplification and is responsible for the second equality. $|\alpha\rangle$ is known as a *coherent state*, a family of non-orthogonal states parameterized by the complex displacement scalar α. Unlike the qubit case, where $|0\rangle$ periodically returns to itself under driving, $\tilde{H}$ continually couples $|\psi\rangle$ to higher and higher energy levels, spreading the state out over the ladder of states as it climbs. The coherent states are fundamental to quantum optics and have numerous important properties:

- Coherent states are eigenstates of a. Specifically $a|\alpha\rangle = \alpha|\alpha\rangle$. It follows that $\langle a + a^\dagger\rangle/2 = \mathrm{Re}[\alpha]$ and $\langle a - a^\dagger\rangle/2i = \mathrm{Im}[\alpha]$, so that α represents the mean position of the particle in phase space (note that from Eq. (1.16), $a+a^\dagger$ and $(a-a^\dagger)/i$ are proportional to x and p respectively).

- In the lab frame, coherent states rotate in phase space just like classically displaced states in classical phase space. $|\alpha(t)\rangle = |\alpha(0)e^{-i\omega t}\rangle$.

- Like the ground state $|0\rangle$, $|\alpha\rangle$ is a minimum uncertainty state, meaning that $\Delta x \Delta p$ saturates the Heisenberg uncertainty principle. This relatively low quantum uncertainty partially explains why these states are semiclassical.

- $|\langle n|\alpha\rangle|^2$ is a Poisson distribution in n with mean $|\alpha|^2$.

- As is the case for $|0\rangle$, $|\langle \psi_x|\alpha\rangle|^2$ and $\langle \psi_k|\alpha\rangle$ are both Gaussian distributions with the same variances as those of $|0\rangle$.

- Coherent states are not orthogonal, but satisfy $\langle \alpha|\beta\rangle = \exp(\alpha^*\beta - |\alpha|^2/2 - |\beta|^2/2)$ (which yields $|\langle \alpha|\beta\rangle|^2 = e^{-|\alpha-\beta|^2}$). While this means that they cannot be used as an orthonormal basis, they nevertheless form an overcomplete basis satisfying $\int d^2\alpha |\alpha\rangle\langle\alpha| = \pi$.

We will make use of these properties enough that they will hopefully become intuitive and easy to remember.

In deriving the expression for α, we also introduced a very important operator called the *displacement operator*

$$D[\alpha] \equiv \exp(\alpha a^\dagger - \alpha^* a) \tag{1.25}$$

which also has properties worth noting:

- $D^\dagger[\alpha] = D[-\alpha] = D^{-1}[\alpha]$

- The composition of two displacements is also a displacement. Specifically, $D[\alpha]D[\beta] = e^{(\alpha\beta^* - \alpha^*\beta)/2}D[\alpha + \beta]$ (this follows from the first BCH formula, the one that we haven't used so far). The global phase is physically irrelevant except in the case where the displacement depends on the state of another quantum system. In this case, it can lead to entanglement between the oscillator and that external system.
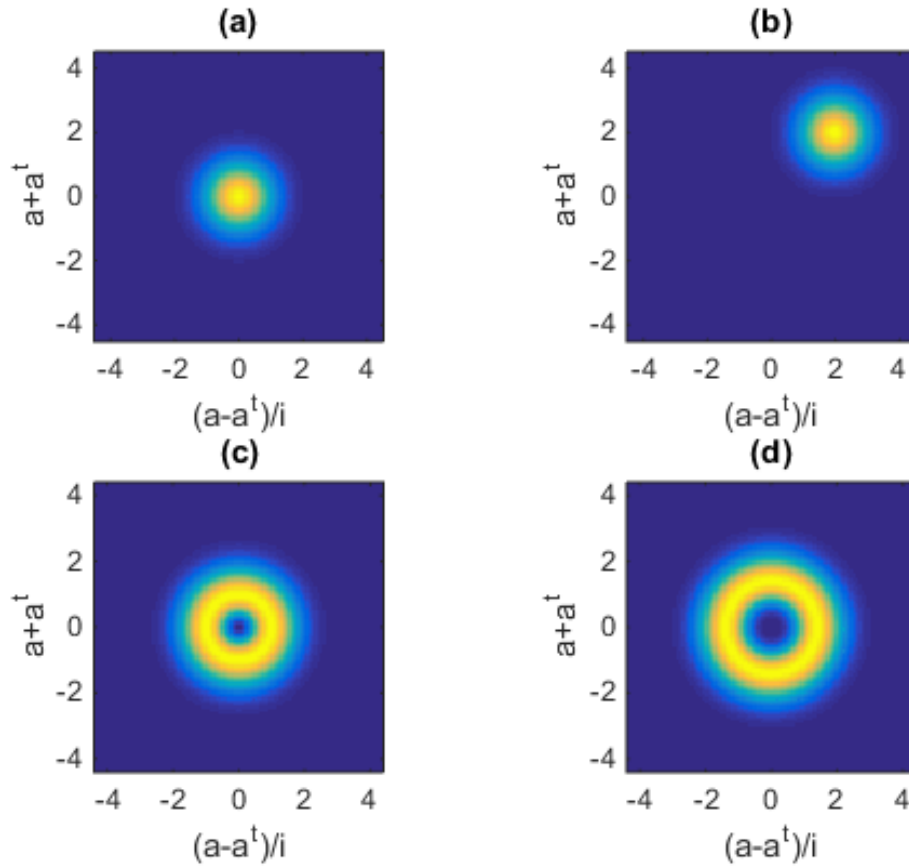
Figure 1.5: Q functions for (a) the ground/vacuum state (b) a coherent state (c) the $n = 1$ Fock state (d) the $n = 2$ Fock state.

- A sequence of displacements that maps a closed path in the complex plane may nevertheless impart a global phase. This phase is e^{2iA}, where A is the area in phase space enclosed by the path[141].

- Just as the Pauli matrices form a complete orthonormal basis under the Hilbert-Schmidt norm, so do the displacement operators.

- $D^\dagger[\alpha]aD[\alpha] = a + \alpha$. Just as we have made heavy use of rotating frames by exponentiating H_0, one can use this identity to work in displaced frames. It is often useful to choose α to displace the oscillator state to its ground state. We will see several examples of this chapter 3.

Note that just as a general qubit drive always implemented a rotation of some kind, $D[\alpha]$ always implements a displacement.

So far, we have covered almost all of the same material for both qubits and harmonic oscillators. The main thing we lack on the oscillator is a simple state visualization like the Bloch vector. We have saved this for last because the most convenient visualizations all rely on coherent states. The first is called the *Husimi Q representation*

$$Q(\alpha) \equiv \frac{1}{\pi}\langle\alpha|\rho|\alpha\rangle. \tag{1.26}$$

$Q(\alpha)$ looks somewhat like a probability distribution, and indeed it does represent the probability of outcome α if one were to measure in the coherent state basis. The only catch is that $|\alpha\rangle$ are not orthogonal states, so orthogonal states do not necessarily have orthogonal Q representations. Q is nevertheless quite a useful quantity, in particular because it depicts the state in phase space as a 2D image, similar to classical phase space. It is also invertible, meaning that one can recover ρ uniquely from $Q(\alpha)$ (though see exercise 7 for an important caveat). Finally, the Q function can be measured using heterodyne detection or phase-preserving amplification, as we will see in the next chapter. The Q functions for several quantum states are plotted in Fig. 1.5.

As we will only use the Q representation for the purpose of visualization, we refer the reader to [39] chapter 3 for a more extensive treatment. Also covered are the equally important Glauber-Sudarshan P and Wigner representations.

Exercise 7 *Cat states: Compute the Q function for $|\psi\rangle \propto |\beta\rangle + |-\beta\rangle$. Don't worry about overall constant factors like normalization. How would the calculation differ if we replaced $|\psi\rangle$ with classical mixture of the same two coherent states ($\rho \propto |\beta\rangle\langle\beta| + |-\beta\rangle\langle-\beta|$)? Show that the difference between these two Q functions scales as $e^{-|\beta|^2}$, indicating that they are difficult to distinguish visually. The Wigner function is much better for this purpose.*

1.3 Quantum Optics: The Jaynes-Cummings Hamiltonian

Qubits and harmonic oscillators form the basis for most of quantum information. To use them in interesting ways, we must combine them and allow them to interact. Of the three possible combinations (qubit-qubit, qubit-oscillator and oscillator-oscillator), the qubit-oscillator system offers the richest dynamics from a quantum optics standpoint. The oscillator expands the number of possible states, while the qubit provides a nonlinearity that is necessary for universal control of the system. To see the power of this combination, note that while the qubit Hamiltonian considered above allowed us to prepare any qubit state, the cavity Hamiltonian was only capable of producing coherent states. This limitation is a fundamental impediment to quantum information processing in harmonic oscillators arising from their linearity; it is not possible to drive the $|i\rangle \leftrightarrow |j\rangle$ transition without also driving the $|i+1\rangle \leftrightarrow |j+1\rangle$ transition. In addition to providing the required nonlinearity for full control, qubit-oscillator interactions also provide a medium through which qubits interact in most physical systems,

and we will see examples of how to generate a two-qubit gate using a harmonic oscillator as
an intermediary in future chapters.

So far we have only considered isolated quantum systems. To treat interactions, we need
a way to model multiple objects with a single wave function, which is accomplished using the
tensor product. This generalization is perhaps most clear for continuous variable states like
harmonic oscillators. Given the states $\psi_1(x)$ and $\psi_2(x)$ of two different particles, the wave
function of the joint state is simply $\psi_1(x_1)\psi_2(x_2)$. Once again, the connection to classical
probabilities is evident; the joint probability distribution of two independent probability dis-
tributions is also given by the product. Furthermore, the measurement outcome probabilities
satisfy $P(x_1, x_2) = |\psi_1(x_1)|^2 |\psi_2(x_2)|^2 = P_1(x_1)P_2(x_2)$. Finally, just as correlated probability
distributions $P(x_1, x_2)$ cannot be written in separable form, there also exist quantum states
$\psi(x_1, x_2)$ that cannot be written as a single tensor product. These are precisely the entangled
states, which play a fundamental role in quantum information and computation.

For discrete systems, the tensor product construction is essentially the same as the above,
except treating the continuous variables x_i as discrete indices ($x_i = 0, 1$ for a qubit, for
example). This notation would suggest treating the composite system $\psi(x_1, x_2)$ as a matrix,
or in the case of more than two subsystems, as a tensor. Although such notation has its
uses, one typically prefers to treat the wave function as column vector for compatibility with
the Schrödinger equation. To do so, we simply stack the columns of $\psi(x_i, x_j)$ on top of
one another. The general construction may be concisely defined for any operators A and B
(column vectors included) as

$$A \otimes B = \begin{bmatrix} A_{1,1}B & A_{1,2}B & \dots & A_{1,n}B \\ A_{2,1}B & \ddots & & A_{2,n}B \\ \vdots & & & \vdots \\ A_{m,1}B & A_{m,2}B & \dots & A_{m,n}B \end{bmatrix}. \tag{1.27}$$

Note that each entry in the above table is a block of the matrix $A \otimes B$, and the vertical and
horizontal dimensions of the whole are therefore the product of the dimensions of A and B.
This particular arrangement is useful because it satisfies $(A \otimes B) \times (C \otimes D) = (A \otimes C) \times (B \otimes D)$. This relation allows us to evolve composite systems with matrix multiplication just as
we did with isolated systems (*i.e.* two isolated systems evolve as $|\psi_{1,2}(t)\rangle = (U_1(t)|\psi_1(0)\rangle) \otimes (U_2(t)|\psi_2(0)\rangle) = (U_1(t) \otimes U_2(t))(|\psi_1(0)\rangle \otimes |\psi_2(0)\rangle)$, meaning that the time evolution operator
for a composite state is simply the tensor product of the original time evolution operators).

Exercise 8 *As an example of when it's useful to think of composite quantum states as
tensors instead of column vectors, the singular value decomposition ($A = U\Sigma V^\dagger$ where Σ is a
diagonal matrix and U and V are unitary) of the matrix $\psi_{i,j}$ yields a useful construction called
the Schmidt decomposition: $|\psi\rangle = \sum_i \Sigma_{i,i}|i_1\rangle \otimes |i_2\rangle$ (the first $\sum$ is a summation symbol, while
the second is a diagonal component of the matrix Σ). Look up the Schmidt decomposition in
Nielsen and Chuang[93]. This decomposition is a simple example of a tensor network, which
includes useful state representations like matrix product states[101].*
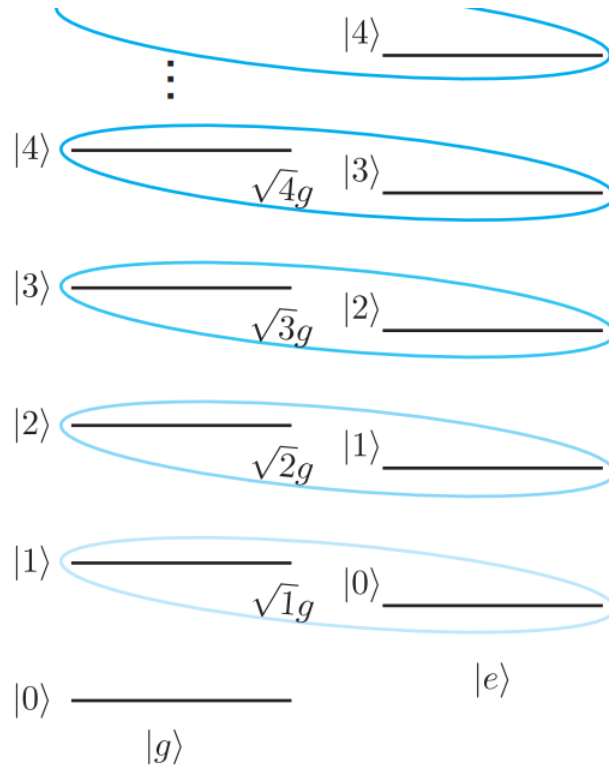
Figure 1.6: The level structure for the Jaynes Cummings Hamiltonian. In the rotating wave approximation, only pairs of levels with the same total excitation number interact, as illustrated with the blue ellipses. Due to the increasing second moment of states as we move up the harmonic oscillator ladder, the coupling strength increases with n.

The Jaynes Cumming Hamiltonian is the most basic form of a qubit coupled to a harmonic oscillator. From an applied standpoint, the coupled qubit-oscillator models countless physical systems, from light-matter interactions to lattice defect states coupled to phonons

$$H/\hbar = \omega_q \frac{\sigma_z}{2} + \omega_o a^\dagger a + g(a + a^\dagger)(\sigma + \sigma^\dagger). \tag{1.28}$$

Each operator acts on only one tensor factor of the composite system, so implicitly $a \to a \otimes I_q$, $\sigma \to I_o \otimes \sigma$ *etc.* Like the single qubit and oscillator models seen already, the essence of the coupling term is a dipole approximation. The new ingredient here is that we treat the dipole driving term as another quantum system.

We can perform calculations with H most easily if we choose a frame in which it is still time independent. This works only if we rotate the qubit and oscillator frames at the same frequency *i.e.* $U(t) = \exp(i(\omega\sigma_z/2 + \omega a^\dagger a))$. Both $\omega = \omega_q$ and $\omega = \omega_o$ are useful choices, so we go with the latter. The Jaynes Cummings Hamiltonian in the rotating wave

approximation is

$$\tilde{H}/\hbar = \Delta\frac{\sigma_z}{2} + g(a^\dagger\sigma + a\sigma^\dagger) \tag{1.29}$$

where $\Delta = \omega_q - \omega_o$. Many sources justify dropping counter-rotating terms $a\sigma$ and $a^\dagger\sigma^\dagger$ because they appear to violate energy conservation, exciting or deexciting the oscillator and the qubit simultaneously[4]. However the Hamiltonian *defines* the energy of a system, so this argument does not hold. What can be said is that if $g \ll \omega_o, \omega_q$, then the interaction acts like a perturbation, and we expect that the energies not to change much when we add it in. $a\sigma$ and $a^\dagger\sigma^\dagger$ change the bare energies by a lot, so we might already expect their contributions to be small. Loosely speaking, time-energy uncertainty allows counter-rotating terms to create or destroy two excitations, as long as these excitations don't last long. This suggests that the contribution of these terms is non-zero, but cancels out on average, as we saw when justifying the rotating wave approximation with time-dependent perturbation theory. This gives us a physical argument for why these terms probably do not contribute. While these arguments hold for many physically reasonable systems, we will see an example in chapter 3 (single- and multi-quadrature measurements) in which these terms play a significant role in the dynamics.

Exercise 9 *The solution for two coupled harmonic oscillators resembles that of the Jaynes-Cummings Hamiltonian. Let $H/\hbar = \omega_a a^\dagger a + \omega_b b^\dagger b + g(a^\dagger b + ab^\dagger)$. The coupling term can be derived from the energy of two masses sharing a common spring: $k(x_a - x_b)^2/2$. Show that H can be diagonalized by introducing the normal mode operators $\tilde{a} = \alpha a + \beta b$ and $\tilde{b} = \alpha b - \beta a$ with $H/\hbar = \tilde{\omega}_a\tilde{a}^\dagger\tilde{a} + \tilde{\omega}_b\tilde{b}^\dagger\tilde{b}$ and $\alpha^2 + \beta^2 = 1$ (chosen so that $[\tilde{a}, \tilde{a}^\dagger] = 1$, $[\tilde{a}, \tilde{b}^\dagger] = 0$ etc.). You don't have to solve for these parameters unless you want to do exercise 10.*

Exercise 10 *Show that $\alpha = \cos(\theta/2)$ and $\beta = \sin(\theta/2)$ with $\theta = \arctan(2g/\Delta)$ and $\Delta = \omega_a - \omega_b$. Show that the normal mode frequencies satisfy $\tilde{\omega}_a + \tilde{\omega}_b = \omega_a + \omega_b$ and $\tilde{\omega}_a - \tilde{\omega}_b = \sqrt{(2g)^2 + \Delta^2}$.*

The Jaynes-Cummings Hamiltonian can be solved exactly in the rotating wave approximation. The elimination of $a\sigma$ and $a^\dagger\sigma^\dagger$ terms reduce the problem to that of an effective qubit. To see how, note that the action of H takes the state $|n, g\rangle$ to $|n-1, e\rangle$, but then can only map the latter back to $|n, g\rangle$. Thus each component of the wave function evolves in a two-dimensional subspace of fixed total excitation number, as depicted in Fig. 1.6. We can then break H up into many of these qubit subspaces and solve each individually. Computing the action of H in a subspace amounts to computing a few matrix elements

$$H_n \equiv \begin{pmatrix} \langle e, n-1|H|e, n-1\rangle & \langle e, n-1|H|g, n\rangle \\ \langle g, n|H|e, n-1\rangle & \langle g, n|H|g, n\rangle \end{pmatrix} \tag{1.30}$$
$$= \hbar\begin{pmatrix} \Delta/2 & \sqrt{n}g \\ \sqrt{n}g & -\Delta/2 \end{pmatrix}$$

---

[4]Arguments about energy conservation were irrelevant to previous sections, as H was time-dependent.

This effective Hamiltonian is just like Eq. (1.15) with $\delta = 0$. We have not yet written down the fully general solution, so let's do that now by computing the eigenvectors and eigenvalues

$$\frac{E_{\pm,n}}{\hbar} = \pm\frac{1}{2}\sqrt{\Delta^2 + n(2g)^2}$$

$$|+,n\rangle = \cos(\theta_n/2)|e,n-1\rangle + \sin(\theta_n/2)|g,n\rangle$$

$$|-,n\rangle = \cos(\theta_n/2)|g,n\rangle - \sin(\theta_n/2)|e,n-1\rangle$$

$$\tan(\theta_n) \equiv -\frac{2\sqrt{n}g}{\Delta}. \tag{1.31}$$

Note the similarity to the results of exercises 9 and 10.

There are two interesting limits to consider in the solution to Jaynes-Cummings. The simplest limit is $\Delta \to 0$, in which case the eigenvectors become equally weighted between $|n,g\rangle$ and $|n-1,e\rangle$ and the effective qubit Hamiltonian Eq. (1.30) reduces to the resonant Rabi Hamiltonian Eq. (1.13). If we start the oscillator in its ground state and the qubit in its excited state (the $n=1$ manifold), then the excitation swaps coherently between the qubit in the cavity at an angular frequency of 2g. These oscillations are called vacuum Rabi oscillations.

The other regime, and the last piece of physics that we introduce quantitatively in this chapter is $\Delta \gg g$, called the dispersive regime. In this case, we can treat the off-diagonal elements of Eq. (1.30) as a perturbation to the detuning terms $\pm\Delta/2$. We could use a Taylor expansion of Eq. (1.31), but instead let's use second-order perturbation theory. Breaking H_n into its diagonal and off-diagonal parts H_0 and H' respectively, the new eigenvalues and eigenvectors are

$$E'_{e,n-1} \approx \frac{\Delta}{2} + \frac{|\langle g,n|H'|e,n-1\rangle|^2}{\Delta/2 - (-\Delta/2)} \tag{1.32}$$

$$= \frac{\Delta}{2} + \frac{ng^2}{\Delta}$$

$$E'_{g,n} \approx -\frac{\Delta}{2} - \frac{ng^2}{\Delta}$$

$$|e,n-1\rangle' \approx |e,n-1\rangle + \frac{\langle g,n|H'|e,n-1\rangle}{\Delta/2 - (-\Delta/2)}$$

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

$$|g,n\rangle' \approx |g,n\rangle - \frac{g}{\Delta}|e,n-1\rangle.$$

The qubit and cavity states become 'dressed' by one another and pick up a small energy shift called the dispersive shift. The energy shift $\pm ng^2/\Delta$ takes a form that can be easily reexpressed as a Hamiltonian, called the dispersive Hamiltonian

$$H_{\text{Dispersive}} = \chi a^\dagger a \sigma_z \tag{1.33}$$

$$\chi \equiv \frac{g^2}{\Delta},$$

which is technically in the dressed basis $|g, n\rangle'$, $|e, n-1\rangle'$. Written in this way, we see that the dispersive Jaynes-Cummings Hamiltonian leads to a qubit-state-dependent frequency shift of the cavity, or a cavity-state-dependent shift of the qubit frequency. This form of coupling is quite useful in practice. For instance, one can measure measure the state of the qubit by measuring the frequency of a harmonic oscillator to which it is coupled. This is one of many reasons why the dispersive Hamiltonian arises extensively in circuit QED.

Exercise 11 *An alternative and powerful way to derive Eq. (1.33) is called the Schrieffer-Wolff transformation. Using the Baker-Campbell-Hausdorff formula at second order $e^X H e^{-X} = H + [X, H] + [X, [X, H]]/2$ and $X = (a\sigma^\dagger - \sigma a^\dagger)g/\Delta$, derive Eq. (1.33) from Eq. (1.28). In general, one can use this transformation to diagonalize a Hamiltonian $H_0 + H'$ to first order in H', provided we can find an X such that $[X, H_0] = -H'$.*

1.4 Applications of Quantum Systems

It is no coincidence that the vast scientific progress of the 20th century coincided with rapid technological development. Without quantum theory, the transistor, LED, laser, MRI and many forms of atomic scale microscopy would have remained ideas of science fiction. Despite the magnitude of recent developments, there is reason to believe that some of the most significant quantum technologies are yet to come. Just as a motor is an electronic device that performs an essentially mechanical task, most 20th century technologies based on quantum mechanics still perform fundamentally classical tasks, such as a transistor switching a current or a laser generating an oscillating electric field. Devices that operate on truly quantum degrees of freedom may offer comparably novel capabilities, analogous to classical computers, CCDs and the internet. The full capacity of these technologies, and in many cases even proofs that they supersede the best possible classical implementation are unknown. The field is still wide open.

The experimental and theoretical tools used in this thesis both have a vast range of potential applications. The field of quantum information is still in its infancy, and we only touch on a small range of its potential scope. Although this thesis does not focus specifically on quantum computation, the basic concepts of that field are essential for our work and for quantum information science in general. Below we outline the main anticipated categories of technology that are expected to be impacted by quantum information, and then give a short, high-level overview of quantum computation, referring the reader to reference [**93**] and other references of this section for details.

Quantum simulation: The classical representation of a quantum wave function grows exponentially as degrees of freedom are added to it. While methods like density functional theory (DFT) and the density matrix renormalization group (DMRG) render many systems tractable, a classical computer's ability to model a general quantum system is fundamentally limited. Richard Feynman pointed out that this exponential growth should make quantum systems powerful simulators of other quantum systems. Long-term ambitions for this field

include understanding and developing high-temperature superconductors, understanding and designing chemicals and simulating high-energy physics processes, for example with lattice gauge theories.

Quantum Computation: Inspired by Feynman's insight, researchers developed models of computation that involve quantum mechanics at their core, allowing for superposition and entanglement. In the early 90s, evidence rapidly accumulated that a computer operating on quantum degrees of freedom could be fundamentally more powerful than a Turing machine, in violation of the extended Church-Turing Thesis. Perhaps the most striking example is the realization by Peter Shor that a quantum computer can factor numbers about as efficiently as any computer can multiply them, which is exponentially faster than the fastest known classical algorithm. For review of the potential near and far-term applications, see [105], though the field is moving so quickly that there were significant updates within a year of its publication[131].[5]

Novel quantum devices Quantum devices have many potential applications beyond computation. One important example is quantum cryptography, which uses quantum correlations to guarantee security and secrecy. Along these lines, researchers have also proposed unforgeable quantum money and certifiable random number generators. In a different vein, a general class of advantages comes from quantum metrology. Whereas classical statistical uncertainties on N repeated measurements scale as $1/\sqrt{N}$, quantum measurements can be made to scale as $1/N$[40]. This discovery pointed out the importance of a quantity called the quantum Fisher information, which predicts the ultimate capabilities of a measurement device via the quantum Cramér-Rao bound. There is hope, and in some cases existing demonstrations showing that molecular imaging, magnetometry, interferometry (such as the use of squeezing at LIGO), clocks and very-long-baseline interferometry for telescopes could all be improved with quantum effects.

Classical technologies at their quantum limits Finally, as classical technologies approach their quantum limits, quantum physics will be required to design them. This is already happening in classical computing, where transistors are approaching the atomic scale. Photonic computation and communication will also become more quantum as signals approach the single-photon level.

While many of the above capabilities do not superficially relate to quantum computation, quantum information nevertheless forms a foundation and common language for understanding them. All of the above examples (though in particular the first three) are naturally described in terms of quantum circuits. For example, the fundamental limits of interferometry may be understood in terms of an algorithmic primitive called the quantum phase estimation algorithm[148]. In subsequent chapters, we will understand open systems, quantum trajectories and quantum feedback as quantum circuits. Therefore an understanding of quantum computation is beneficial for understanding the ultimate limits of quantum technology in general.

A brief overview of quantum computation In the current understanding, quantum

---

[5]See also https://www.scottaaronson.com/blog/?p=3880

algorithms gain their advantage over classical computers through two essential primitives. The first is *Grover's search algorithm[46]*, which provides an intrinsic advantage in searching an unsorted database. This database can be thought of as a function $f(x)$ which maps integer inputs to integer outputs in some fashion (x might be encoded by its binary representation over a register of qubits). To find an entry x for which $f(x)$ produces some desired output y, a classical computer must resort to querying each value of f individually until the right x is found. Thus if f is defined over N values of x, then the average number of queries before success is $N/2$ (denoted $\mathcal{O}(N)$). Surprisingly, a quantum computer can find the right value in $\mathcal{O}(\sqrt{N})$ by passing multiple values of x in superposition on each query! While this necessitates implementing $f(x)$ in the quantum computer (so we cannot use this trick to find a book at a disorganized bookstore), it offers a quadratic speedup for an enormous range of problems.

It is often said that a quantum computer achieves its speedup by computing on all possible inputs in superposition. While it is true that a quantum computer can evaluate f on all values of x by preparing each qubit of x in a superposition, the output state is also a superposition over all possible answers. As measurement collapses the wave function, we have no way to access all of these answers in a single run. If we could, or if we could guarantee that the measurement outcome would correspond to y, then such an algorithm would achieve an enormous speedup on the unstructured search problem instead of quadratic. The actual intuition behind Grover's search algorithm is somewhat more subtle, though it is satisfyingly straightforward to understand quantitatively and visually. We refer the reader to reference [**93**] for details.

If only quadratic speedups were available, quantum computation might not be worth the trouble. However quantum systems appear to be exponentially more difficult to simulate than classical, hinting at something more powerful that Grover's search. The second algorithmic primitive for computation is the *quantum Fourier transform*[10], which provides the exponential speedup found in Shor's factoring algorithm. The quantum Fourier transform exploits the fact that the discrete Fourier transform is itself a unitary operation, and therefore can be applied directly to a wave function. It allows a quantum computer to find the period of a function in only a few function calls, and forms the basis for another important routine called quantum phase estimation mentioned above. Again we refer the reader to reference [**93**] for more.

The above algorithms may be implemented efficiently using single and two-qubit operations. Which gates will be available and the connectivity between qubits will ultimately be determined by the physical implementation of the computer. A surprising range of physical systems provide viable, equally powerful models of computation. For example, given photon counters single photon states as inputs, linear optics (beam splitters and phase shifters) suffice to perform any quantum algorithm*[61]*. This construction is particularly remarkable because many of the basic operations upon which it relies fail with high probability. Another significant example is measurement-based computation, in which no single-or two-qubit gates are allowed[**95**]. Provided a particular initial state, called a cluster state, universal computation is implemented by sequentially and adaptively measuring different parts of the
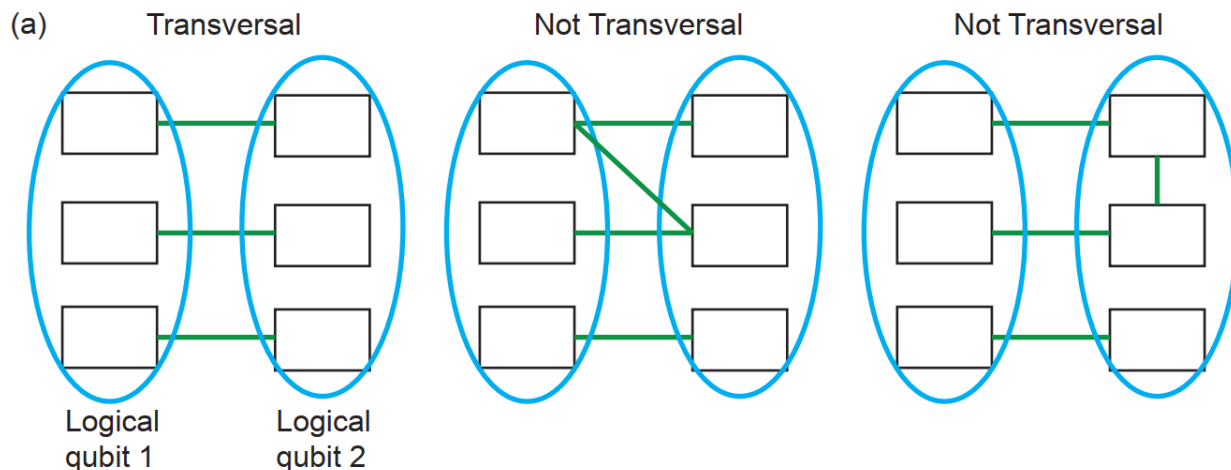
Figure 1.7: (a) Graphical illustration of two logical qubits each encoded into 3 physical qubits (black boxes). A transversal gate operates as shown. (b) A non-transversal gate. (c) Another non-transversal gate.

computer. Finally, it is possible to encode the output of any quantum circuit into the ground state of a local Hamiltonian H_P (meaning that only nearby qubits are coupled). By smoothly varying H from a simple, non-interacting Hamiltonian to H_P, one can efficiently implement any algorithm[2]. Adiabatic theorem ensures that the system remains in its ground state at all times, and the discreteness of the original quantum circuit appears to be lost.

When we consider quantum computation in realistic systems, the continuous nature of the wave function poses a challenge. Classical models of computation are digital, and dissipative processing like latching and copying ensure resilience to noise. Neither of these tricks work in quantum systems, where dissipative interactions destroy coherence and copying is forbidden via the no-cloning theorem[93]. In the early days of quantum computing, researchers worried that even the smallest errors in computation would compound exponentially and eventually ruin a realistic device.

An important conceptual shift in noisy quantum computation was the insight that although errors are continuous, there always exists an equivalent error model that is discrete (we show this result in chapter 3 exercise 23). For example, if a qubit undergoes a random σ_z rotation of arbitrary angle, we can equally well model this process as random rotations by only 0 or π. With a discrete set of errors to worry about, quantum error correcting codes were devised to correct them one-by-one. A single 'logical' qubit can be encoded in multiple physical qubits using entanglement, and errors can be corrected with nonlocal measurements. While encoding introduces overhead, an important result called the threshold theorem showed that this scaling is favorable, and that errors can be suppressed to arbitrary precision without so much overhead that they null the benefits of quantum computation.

There is error correction and then there is fault tolerance. Error correction means that we can hold a quantum state in an encoded memory for longer than if we were to store it in

a single physical qubit. Fault tolerance means that we can perform computations on these qubits while preserving their coherence, and in particular ensuring that a single error does not spread and pollute the rest of the system. A standard method to ensure fault tolerance is to implement gates transversely. In the context of an error correcting code, a transversal gate between two encoded logical qubits only allows for pairwise gates between the physical qubits. This is best illustrated with pictures, as in Fig. 1.7.

Unfortunately, restricting to transversal gates limits the range of operations that may be performed on a logical qubit. Eastin-Knill theorem[31] shows that transversal gates cannot be universal[29], meaning that not all quantum algorithms can be implemented with the available gate set. There are many proposed routes to circumvent this impediment to universal, fault-tolerant quantum computation. One proposed route is magic state distillation, in which one imperfectly prepares a state that cannot be prepared fault-tolerantly and then uses fault tolerant gates and measurements to purify it to a near-perfect state[13]. Magic states also have the remarkable property that one can be 'consumed' to perform gates that otherwise could not be implemented fault tolerantly. Another route to fault tolerant quantum computation are topological quantum states like the surface code. Most, but not all proposals for topological computing also require magic state distillation.

1.5 Recommended Reading

Perhaps the most daunting aspect of starting work in a new field of research is the vast body of potentially relevant literature. Graduate courses provide overviews and calculational tools for many fields, but rarely cover even slightly narrower topics like quantum electrodynamics, let alone superconducting circuits. While the main goal in writing this thesis is to bridge this gap, no two research projects require the same knowledge base and no two students come in with the same background. To provide a more flexible introduction, below is a list of particularly useful books and papers along with recommended sections, roughly sorted with the most accessible texts at the top. The usefulness of these texts depends on one's interest of course, but these are chosen for their breadth and clarity.

- **Quantum Computation and Quantum Information** *Nielsen and Chuang*[**93**]: Conversational and perfect for self-learning, this is the single resource that I recommend to everyone entering the field. Part I chapter 1 presents a broad overview of quantum information science, and chapter 2 contains a nice summary of quantum theory from a more informational perspective. Part II, quantum computing, is a great place to learn about algorithms, circuits and the abstract details of experimental implementation. Part III chapters 8-10 is a must-read for anyone interested in understanding decoherence, open quantum systems and quantum error correction.

- **Introductory Quantum Optics** *Gerry and Knight*[**39**]: When we deal with light quantum mechanically, we are studying quantum electrodynamics (QED). While quantum field theory is a class well-worth taking, not everyone has the time or desire to

trudge through Peskin and Schroeder. Fortunately, QED is a particularly accessible
subset of QFT that doesn't require most of its machinery. *Gerry and Knight* covers
the full quantum theory of light at an undergraduate level. Chapter 1 alone is well
worth reading.

- **Quantum Optics** *Scully and Zubairy*[120]. There are many great resources available
  on quantum optics, but this one strikes a particularly nice balance between pedagogy
  and completeness. While the Jaynes Cummings Hamiltonian is never referenced by
  name, it is developed in the first few pages of chapter 6 and studied extensively. An-
  other comparable resource with a very different approach is Haroche's *Exploring the
  Quantum: Atoms, Photons and Cavities*, which is perhaps a bit more accessible[51].

- **Microwave Engineering** *Pozar*[104]: A standard reference, you may find this to be a
  piece of cake after *Jackson*! It provides a practically useful introduction to the theory of
  standard microwave components like resonators (chapter 6), splitters, circulators, and
  directional couplers (chapter 7) while emphasizing essential concepts like impedance
  matching (chapter 2), scattering matrices (4.3) and reciprocity. Chapter 3 is useful for
  design of superconducting devices. See [87] for a more advanced and complete text on
  the subject.

- **Introduction to Superconductivity** *Tinkham*[135]: Provides a concise but fairly
  broad introduction to the major concepts of superconductivity. Most of the book is
  dedicated to phenomenological models, with one chapter dedicated to an accessible
  first-principles derivation of the Bardeen-Cooper-Schrieffer (BCS) model of supercon-
  ductivity. It is difficult to outline what knowledge is necessary for research on super-
  conducting circuits as the field grows. A few terms to keep an eye out for on more
  specialized topics include quasiparticles and their dissipation, vortices, Andreev bound
  states and the concept of conduction channels. Not all of these areas are covered in
  detail in *Tinkham*, but are commonly discussed in the literature.

- **Introduction to Quantum Noise, Measurement and Amplification** *Clerk et
  al.*[20]: A Review of Modern Physics article that covers many experimentally relevant
  calculations for quantum amplifiers and general measurement theory. The appendices
  are perhaps the most useful, explaining many standard calculations tools and their
  application. You'll see this referenced all over the place.

- **Quantum Noise** *Gardiner and Zoller*[37]: A dense but great book on quantum
  optics with a strong emphasis on open quantum systems. It's a standard reference
  if you're embarking on a serious quantum optics calculation. Chapter 3 derives the
  quantum Langevin equation, chapter 5 derives the master equation and another handy
  trick called quantum regression theorem. Some topics of this book have simpler, more
  modern descriptions elsewhere, such as the SLH formalism for modeling very general
  networks of quantum systems[22].

The reader is advised to recognize the above reference numbers to see their applicability in the following introduction. We have also put in *italic* citations of particular significance for the field. It is hoped that the reader will get a big-picture overview of the field by reading enough of these papers to understand the main results, without working through the details or proofs unless desired. For learning concepts in detail, we highlight particularly useful or pedagogical references in **bold** to aid self-study.

Finally, there is no substitute for reading current research papers to get a sense of the open questions in the field and how people are trying to solve them. With multiple potentially relevant papers being published every day, it is easy to become overwhelmed. The following tips may be useful for managing the deluge:

- You do not have time to read everything, and you don't have to. Choose which papers to read based on their importance to your work or how interesting they sound to you. If a paper piques your interest but is not directly related to your project, it's not a waste of time to go through it. Reading papers that interest you is far more likely to inspire your future work than diligently reading only the papers that are relevant to your current project.

- If you decide to open a paper, spend as much or as little time as is useful. For a crucial or foundational papers that are directly related to your research, it can be worth it to work through derivations and detailed results. If you only want to spend half an hour (or five minutes) on a paper, try to extract the essence of their results rather than get bogged down in details and derivations. What are the main results of the experiment? Why was it challenging? What are the most important pieces of intuition behind a theory paper?

- Sometimes the most important part of a paper are the introduction and references therein. This section is a great place to get a sense of the most important results in a field.

- Set time aside every week to read papers. Make it a part of your routine. Short term demands can make it difficult to fit in, but the long-term payoff is substantial.

- Papers aren't the only way to learn about research. More and more seminar and conference talks are recorded on YouTube, and explanations are often better than those in the corresponding manuscripts.

Chapter 2

Introduction to circuit QED

2.1 Introduction

It is somewhat remarkable that we can even consider encoding a qubit in a superconducting current. Most systems that have been used as qubits encode information in a single degree of freedom, like the orbital state of an electron in an atom or the spin of a nucleus. In contrast, superconducting circuits encode information in the collective motion of a billions of billions of electrons. Coherently encoding information in such a system is analogous to the famous Schrödinger's Cat thought experiment, in which a system is put in a superposition of two macroscopically distinct states. In fact, the original proposal for observing coherence in a superconducting circuit was motivated not by quantum information, but rather by testing if quantum mechanics remained valid for macroscopic systems*[64]*. This work suggested such a system "would probably be as near as we are likely to get to a laboratory version of Schrödinger's Cat," and that preserving its coherence long enough to measure it would "likely to present serious, though possibly not insuperable, difficulties." These difficulties indeed yielded to experimentalists several years later when the macroscopic coherence of the phase difference φ across a junction was demonstrated in reference *[82]*. Since then, steady progress has lead to qubits with near-millisecond coherence times and large-scale efforts to do quantum computation and simulation. However even in the most applied work, the original goal of demonstrating macroscopic coherence is never far away. Efforts to improve qubit coherence and implement progressively larger error correcting codes will push quantum theory to untested limits.

In this chapter, we outline the essential features of circuit quantization and superconducting qubits. We begin with a broad, qualitative overview of superconductivity, highlighting the features of the theory that are relevant for qubit design. This section includes a quantitative derivation of the Josephson effect, the circuit element that makes quantum information processing possible. Section 2.3 describes the process of circuit quantization and illustrates the connections between Josephson junctions and inductors. We do not cover the general quantization of a circuit with multiple degrees of freedom in much depth, but instead pro-

vide overview of the available methods for treating them (including a high-level overview of black-box quantization). Section 2.4 summarizes the known sources of decoherence in superconducting qubits and common qubit designs invented to circumvent them. In section 2.5, we take our first excursion into open quantum systems, introducing input-output theory in the context of microwave amplification.

2.2   Superconductivity: Broad Brush Strokes

In general, electrons in a solid are strongly interacting due to Coulomb forces. However we typically study solids using a band structure theory, which maps out electron energy as a function of its momentum while completely ignoring electron-electron interactions. The surprising success of the band-structure way of thinking derives from the Fermi liquid, or Landau liquid model. In this model, it is shown that electron-electron interactions dress the single-electron states to give them a different effective mass (think of a particle moving through molasses). Most importantly, the result is again a non-interacting model which the structure of the Fermi sea. The Landau liquid model is correct to all orders in perturbation theory, but this does not guarantee that it never fails. In particular, Cooper showed that the Fermi sea is unstable if there is an attractive interaction between electrons; an arbitrarily small binding energy between electrons of equal and opposite momentum destabilizes the Fermi sea and leads to bound pairs of electrons that are not captured at any order in perturbation theory[135]. A net attractive force can arise between electrons via interactions with the underlying lattice. These bound pairs, composed of two fermions, behave like a boson and condense in a manner similar (though not identical) to a Bose Einstein condensate. The resulting state is the Bardeen-Cooper-Schrieffer (BCS) state of superconductivity, named for its discoverers.

Superconductivity plays several crucial roles in circuit QED. The obvious role is the elimination of resistance, which allows oscillations in an LC circuit to persist for a long time, and DC currents to persist almost indefinitely in macroscopic samples.[1] Another useful consequence of superconductivity is that the electron condensate behaves like a rigid fluid of particles that move as one collective degree of freedom. It costs a finite amount of energy, superconducting gap E_g, to 'break' electrons out of this state. Thus we are able to treat the *quantum state* of a chunk of metal as if its only degree of freedom were the net current flowing through it, which allows us to encode and safely store quantum states in a macroscopic material. Finally, when a superconducting wire is interrupted with a thin insulating barrier, a surprising phenomenon called the Josephson effect imbues this simple circuit element with the properties of a dissipationless non-linear inductor. This non-linearity turns out to be necessary to make qubits and amplifiers. For the sake of completeness, we give a qualitative description of superconductivity to introduce the essential concepts that

---

[1]AC currents dissipate because superconducting currents have finite inductance. Therefore an AC current induces a non-zero voltage gradient inside and around the material, which causes the flow of currents that are not in the superconducting phase or are exterior to the superconductor, and thus have finite resistance.

often arise in circuit QED. We then cover the Josephson effect in more detail, as it is the essential ingredient for making a qubit.

Conventional superconductors, those that may be modeled by a phonon-mediated electron-electron interaction and the standard BCS theory, may be characterized by three physical parameters: the gap E_g, the London coherence length λ_p and the Pippard coherence length ξ, all of which are temperature dependent. The gap $E_g(T)$ is the minimum energy required to break a single Cooper pair. The superconductor will not absorb energy below E_g, which provides robustness to thermal fluctuations and noise. The gap also sets the temperature at which the metal begins to superconduct by the relation $T_c = E_g(0)/3.528k_b$, where T_c is called the critical temperature. $E_g(T_c) = 0$ and increases to its maximal value at $T = 0$. When energetic photons, phonons or other radiation sources deposit energy above the gap break a Cooper pair, the result is a quasiparticle. Quasiparticles move throughout the material and lead to dissipation. In particular, if a quasiparticle tunnels across a Josephson junction in a qubit, it can add or remove energy from the qubit state, leading to spurious relaxation and excitation[123]. At typical operating temperatures, a thermal distribution should leave no quasiparticles in a block of metal the size of the Earth![2] However, much higher quasiparticle concentration is observed in practice, and they are a major source of decoherence[123, 122], though it may be possible to eliminate them through improved shielding. Quasiparticles have been observed in circuit QED both through their affect on qubits, and by trapping them in a superconducting nanowires[66].

While the BCS model ultimately provided a physical explanation of superconductivity, phenomenological models still play an important role in its understanding. The London and Pippard coherence lengths are best understood using *Landau-Ginzburg theory*, a phenomenological model that was discovered before BCS. The essential ingredient of Landau-Ginzburg theory is the replacement of the macroscopic many-electron wave function with an unnormalized wave function $\psi(x)$, and the Schrödinger equation with the non-linear Landau-Ginzburg equation. The local density of superconducting electrons is given by $n_s = |\psi(x)|^2/n$ where n is the total electron density. The Landau-Ginsberg equations predict that superconducting currents are rigid, in the sense that they can only support curl-free flow *i.e.* $\nabla \times J(\mathbf{r}) = 0$. This leads to the *Meisner effect*, in which magnetic fields are completely expelled from a superconductor. However the magnetic field cannot go from a finite value to zero perfectly at the interface between the vacuum and the metal, as this would require an infinite current density at the surface. The London penetration depth gives the length scale over which an externally imposed magnetic field decays near an interface. In general, $\lambda_L = \sqrt{m^*/\mu_0 n_s e^{*2}}$, where n_s is the bulk superconducting current density and m^* and e^* are respectively the effective mass and charge of the current carriers (so $e^* = 2e$ for Cooper pairs).

The phase of ψ is the physical phase in the BCS wave function, and relates to the Pippard

---

[2]T_c for aluminum is 1 K, so $E_g = 3.528k_b$. Operating at 30 mK, the Boltzmann suppression factor is $e^{-E_g/k_bT} \approx 10^{-51}$. This wins out against the charge carrier density of aluminum multiplied by the volume of the Earth $((2.1 \times 10^{29}\mathrm{m}^{-3})(1.1 \times 10^{21}\mathrm{m}^3))$. This calculation is somewhat optimistic, since we are neglecting that there is a continuum of quasiparticle states above the ground state, but you get the gist.

coherence length.[3] This phase has a physical importance similar to that of an ordinary wave function. For instance, spatial variation of the phase is associated with momentum and leads to current flow. However a superconductor can only support oscillations up to a certain spatial frequency before the associated energy exceeds the gap. This enforces a distance scale over which the superconducting phase remains coherent, which is the Pippard coherence length ξ. One can think of ξ as the effective size of the Cooper pairs. ξ scales inversely with T_c, so that superconductors with a large gap have a shorter coherence length. It is possible for thermal fluctuations to change the phase coherence between two regions of superconductor by 2π, particularly in constrictions that are narrow compared to the coherence length. This process is called an incoherent phase slip, and can lead to dissipation of a supercurrent. Coherent phase slips were recently observed[5], which are phase slips arising from tunneling instead of thermal activation.

ξ and λ_L are independent parameters, and one can have $\xi > \lambda_L$ or $\lambda_L > \xi$. This comparison turns out to be enormously consequential for the properties of a superconductor. The qualitative difference is so significant that superconductors with $\xi/\lambda_L > \sqrt{2}$ are called *type I* while those with $\xi/\lambda_L < \sqrt{2}$ are called type II. This ratio determines how the material responds to external magnetic fields. For both type I and type II, the Meissner effect breaks down at sufficiently strong fields, and the material returns to its normal metal state. In a type I material, there is an energy penalty to having a normal metal region next to a superconducting region. Loosely speaking, this arises because the Cooper pairs are larger than λ_L, and thus are exposed to magnetic field. Consequently, if the material goes normal, it tends to go normal in large regions, so as to minimize the surface area between normal and superconducting. In contrast, type II have a negative surface energy, so that any penetrating magnetic field will fragment into many tiny normal regions called vortices. Each vortex contains a quantum unit of flux $\phi_0 = h/2e$ and is surrounded by a circulating current that screens this flux from the bulk.[4] As type II superconductors can support some magnetic field penetration without going normal, they have higher critical fields and are more practical for creating superconducting magnets. In principle wandering vortices pose a threat to superconducting qubit lifetimes, though they can be pinned in place by a number of mechanisms.

Josephson junctions We now turn to a quantitative description of the most important superconductivity effect for circuit QED. At first glance, a superconductor that is interrupted by an insulating barrier looks like an unlikely candidate for an inductor, which we normally associate with coils of wire and a significant amount of energy stored in a magnetic field. Indeed, when this problem was considered in the early days of superconductivity, it

---

[3]For those who have worked through BCS before, the phase of ψ is simply the phase factor that appears in the BCS wave function $\psi_{\text{BCS}} = \prod_k (u_k + e^i \varphi v_k c_{k\uparrow}^\dagger c_{-k\downarrow}^\dagger)|0\rangle$. It is often ignored in derivations, but the important feature is that it is forced to have the same value in every term of the product, so that it becomes a macroscopic parameter.

[4]One can imagine that the condensate picks up an Aharonov-Bohm phase as it travels around a loop enclosing a magnetic field. For ψ to be single-valued, the minimum phase Aharonov-Bohm phase is 2π, which limits the amount of flux that can penetrate to ϕ_0.

was assumed that single-electron tunneling across the barrier would dominate and nothing too unusual would happen. This expectation changed drastically when Josephson showed that coherent tunneling of Cooper pairs were not only significant, but would also generate surprising electrical properties resembling those of an inductor.

To model a tunnel junction, we begin by considering two superconducting regions separated by a thin insulating barrier. In superconducting circuits, this is often a layer of aluminum oxide that is grown on aluminum. Anticipating that current will end up flowing between them, we parameterize the quantum state of the system according to the number of Cooper pairs on each island as $|m\rangle = |N_L - m, N_R + m\rangle$. The tunneling of Cooper pairs is described by the Hamiltonian[41, 118]

$$H_T = -\frac{1}{2}E_J \sum_{m=-\infty}^{\infty} |m\rangle\langle m+1| + |m+1\rangle\langle m|. \tag{2.1}$$

E_J may be calculated by modelling electron tunnelling microscopically. It may also be estimated from a physical device by measuring the junctions room temperature resistance R_T and using the *Ambegaokar-Baratoff relation* $E_J = h\Delta/8e^2 R_T$, where Δ is the superconducting gap.

H_T translates Cooper pairs across the junction. Its eigenstates must be a uniform superposition of all values of m, otherwise acting H_T on the state would change it. The eigenstates are

$$|\varphi\rangle = \frac{1}{\sqrt{2\pi}} \sum_{m=-\infty}^{\infty} e^{im\varphi}|m\rangle \tag{2.2}$$

The above states are physically unreasonable, as they contain states with an arbitrarily large charge imbalance. We could remedy this by adding Coulomb attraction to H_T, which would energetically suppress states of large m. However our aim is eventually to do lumped element circuit analysis, so we defer treatment of the Coulomb energy until we consider capacitance.

The above basis diagonalizes H_T, allowing us to write it in terms of a phase operator

$$H_T = -E_J \cos(\hat{\varphi}), \quad \hat{\varphi} = \int_0^{2\pi} \varphi|\varphi\rangle\langle\varphi|d\varphi. \tag{2.3}$$

The phase φ turns out to be the difference in superconducting phases across the junction, the one that we discussed in reference to Landau-Ginzburg theory. Note that $\hat{\varphi}$ is a somewhat unusual operator, as its definition is ambiguous. Taking the integration bounds to be from $-\pi$ to π would yield the same Hamiltonian but a different operator $\hat{\varphi}$. This ambiguity will be a bit of a nuisance, but poses no fundamental problem if we are careful.

Exercise 12 *Verify that $|\varphi\rangle$ is an eigenstate of H_T with eigenvalue $-E_J \cos(\varphi)$.*

Exercise 13 *Show that $e^{i\hat{\varphi}} = \sum_{m=-\infty}^{\infty} |m\rangle\langle m+1|$ and $e^{-i\hat{\varphi}} = \sum_m |m+1\rangle\langle m|$. This phase operator does not suffer from ambiguity in its definition. Note the similarity to the results of exercise 12.*

H_T allows us to compute the current flowing through the junction using the Heisenberg equations of motion. Defining the Cooper pair number operator $\hat{n} = \sum_{m=-\infty}^{\infty} m|m\rangle\langle m|$ lets us derive the current as

$$I \equiv 2e\frac{d\hat{n}}{dt} = \frac{2ie}{\hbar}[H_T, \hat{n}] = \frac{E_J e}{i\hbar}\sum_m (|m\rangle\langle m+1| - |m+1\rangle\langle m|) \tag{2.4}$$

$$= \frac{2eE_J}{\hbar}\sin(\hat{\varphi}).$$

Just as a phase gradient in a wave function leads to momentum (recall that momentum eigenstates are $e^{ixp/\hbar}$), a finite change in phase across the junction leads to a flow of current. Indeed, the Josephson junction can be thought of as a form of nonlinear kinetic inductance.[5] Eq. (2.4) is important enough to merit the title of the *first Josephson relation*. Note that as sin is bounded between -1 and 1, the amount of current that may flow through a Josephson junction is also bounded by the *critical current* $I_c = 2eE_J/\hbar$. If one manages to force more current than I_c through the junction, the superconducting phase breaks down and a finite resistance develops.

To understand the electrical properties of the Josephson junction, we need to understand the dynamics of $\hat{\varphi}$. In isolation, $\hat{\varphi}$ commutes with the Hamiltonian and is therefore a constant of motion. To see any time dependence, we need to see how the junction responds to an externally applied voltage. An applied voltage across the junction adds a term $\hat{Q}V = -2eV\hat{n}$ to H_T. V could be a classical scalar or an operator that commutes with $\hat{n}$.

We could compute the Heisenberg equation of motion for φ directly, but the result is less clean due to the artificial choice of boundaries in its definition Eq. (2.3). It is cleaner to first work with the periodic operator $e^{i\hat{\varphi}}$ from exercise 13. Both $\cos(\hat{\varphi})$ and $e^{i\hat{\varphi}}$ are diagonal in the $|\varphi\rangle$ basis and therefore commute, which leaves

$$\frac{de^{i\hat{\varphi}}}{dt} = \frac{i}{\hbar}[-2eV\hat{n}, e^{i\hat{\varphi}}] \tag{2.5}$$

$$= \frac{2ieVe^{i\hat{\varphi}}}{\hbar}.$$

using the result of exercise 13. $e^{i\hat{\varphi}}$ is not Hermitian and therefore not observable, but we can use the chain rule to recover an equation of motion for $\hat{\varphi}$

$$\frac{de^{i\hat{\varphi}}}{dt} = ie^{i\hat{\varphi}}\frac{d\hat{\varphi}}{dt} = \frac{2ieVe^{i\hat{\varphi}}}{\hbar} \tag{2.6}$$

$$\frac{d\hat{\varphi}}{dt} = \frac{2eV}{\hbar}$$

---

[5]Kinetic inductance is inductance arising from the kinetic energy of the electrons themselves, as opposed to energy arising from the magnetic field generated by current flow. It is only non-negligible in high-mobility conductors like superconductors, and becomes more significant in narrow constrictions such as nanowires (imagine fluid flow forced through a narrow constriction, so that kinetic energy increases).

which is known as the *second Josephson relation.* Combining these relations, we immediately see some bizarre effects. For constant V, $\hat{\varphi}$ grows linearly in time and the current through the junction oscillates at a frequency of $2eV/h$ ($2e/h = 1/\phi_0 = 483.597.011$ MHz/μV). Thus the current spends half of its time flowing against the externally applied voltage! This phenomenon, called the *AC Josephson effect* has several important applications in physics. As frequency may be measured with extremely high precision by counting oscillations over a long time interval, it provides one of the most precise methods to measure voltage. The Josephson relations will break down if the AC Josephson frequency exceeds the superconducting gap divided by $\hbar$, at which point the oscillation are so high frequency that they generate quasiparticles.

Exercise 14 *It is often stated that $[\hat{\varphi}, \hat{n}] = i$, even when $\hat{n}$ is discrete as in Eq. (2.1). Show that if A has countably many eigenvalues (so that it may be written $\sum_n a_n |n\rangle\langle n|$), then $[A, B]$ is hollow in the eigenbasis of A i.e. all the diagonal elements are zero. Argue therefore that $[\hat{\varphi}, \hat{n}] = i$ is impossible*

Exercise 15 *As we will see in the next section, it is often useful or even necessary to treat charge as continuous rather than discrete, in which case $[\hat{\varphi}, \hat{n}] = i$ is perfectly valid. Use this commutation relation to derive the first and second Josephson relations. A useful lemma is that if two operators A and B commute with their commutator $[A, B]$, then $[A, f(B)] = [A, B]f'(B)$.*

For our purposes, Eq. (2.3) will provide the main starting point for quantum circuits. However the Josephson junction's most important application for our purposes, its properties as a nonlinear inductor, are easily shown from the Josephson relations. For a linear inductor, $V = L\dot{I}$, so to cast the relations in this form, we take the time derivative of Eq. (2.4)

$$\dot{I} = \frac{d\hat{\varphi}}{dt} I_c \cos\hat{\varphi} = \pm\frac{2eV}{\hbar}\sqrt{I_c^2 - I^2} \tag{2.7}$$
$$V = \pm\frac{\hbar}{2e\sqrt{I_c^2 - I^2}}\dot{I}.$$

As written, the sign is ambiguous and must be determined from the sign of $\cos\hat{\varphi}$. The possibility for a negative inductance arises for the same reason that current flows against V, and is a manifestation of hysteresis. The voltage across the junction is proportional to the derivative of the current, but the proportionality constant itself depends on the current. To zeroth order in I, the Josephson inductance is

$$L_J = \frac{\hbar}{2eI_c} = \frac{\hbar^2}{4e^2 E_J} \tag{2.8}$$

and behaves as a linear inductor when I is small relative to I_c. For a reasonable value like $E_J/\hbar = 12$ GHz, $\hbar/2eI_c \approx 100$nH, which is quite large for a device that might only be 50×50 nm!

(a)   (b)   (c)   (d)
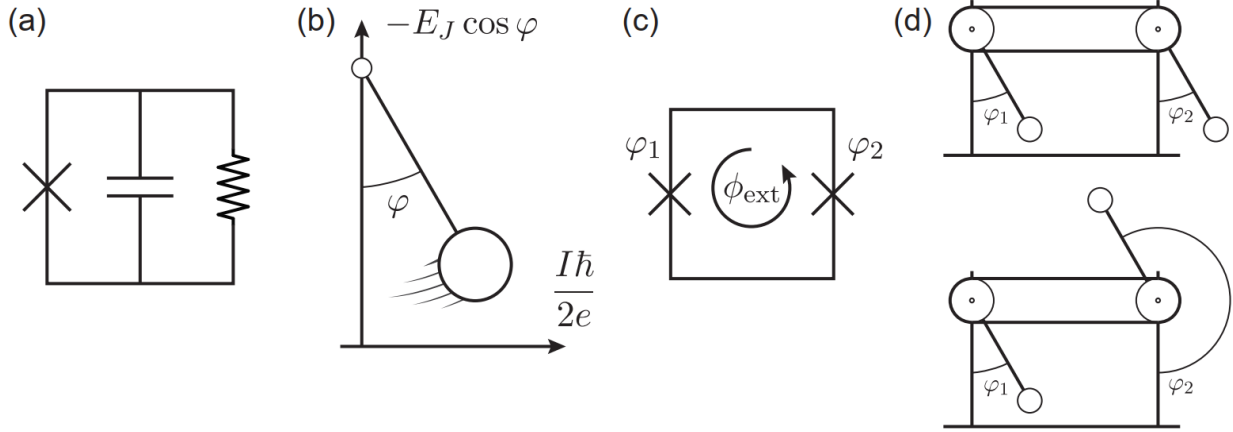
Figure 2.1: (a) A Josephson junction (typically denoted as an X on a circuit diagram) in parallel with a capacitor and a resistor. (b) A mass coupled to a pendulum, which has the same Hamiltonian as (a). A SQUID threaded with an external magnetic flux ϕ_{ext}. (d) A physical analog of the SQUID for $\phi_{\text{ext}} = 0$ (top) and $\phi_{\text{ext}} = \phi_0/2$ (bottom). In the latter case, the gravitational potential of the two masses cancel, effectively removing the junction from the circuit.

As the Josephson junction comes up so much in circuit QED, it is worthwhile to spend a little more time developing intuition for its dynamics. For this purpose, there are two physical analogs of H_T. The first comes from noting that Eq. (2.1) looks like the kinetic energy energy, or hopping term for a 1 dimensional tight binding model Hamiltonian, where m labels lattice positions. Eq. (2.2) shows that φ plays the role of momentum. The elegant feature of this analogy comes in when we add an external voltage. The additional energy term $-2eV\hat{n}$ is equivalent to the application of a constant force Fx. Under this force, φ increases until it hits the edge of the Brillouin zone $\pm\pi$, at which point it Bragg reflects, entering from the opposite side of the Brillouin zone. The resulting Bloch oscillations offer one visualization of the AC Josephson effect.

The second analogy requires us to think of a Josephson junction embedded in a parallel LCR circuit, as in Fig. 2.1a. This simple circuit will form the basis for our two simplest qubits and the most common microwave amplifier. It is also necessarily the effective circuit for any realistic Josephson junction. As a Josephson junction consists of two superconductors separated by a dielectric, it is as much a capacitor as it is a junction. Thus as we noted earlier, we must include the Coulomb energy in any realistic model of a junction. Given the standard formula for a capacitor's energy, the resulting Hamiltonian is

$$H = \frac{(2e)^2}{2C}\hat{n}^2 - E_J\cos(\varphi). \tag{2.9}$$

We defer the mathematical treatment of a resistor for later, but for now note qualitatively that it should tend to damp the system to a state of zero current ($\varphi = 0$ or $\varphi = \pi$). Fig.

2.1 suggests interpreting φ as the angular coordinate of a pendulum, and $-E_J \cos(\varphi)$ as its gravitational potential energy. In contrast to the first analogy, the capacitance term now plays the role of kinetic energy. Kinetic energy is quantized due to the periodic nature of φ, just as in a quantum rigid rotor. Thus we can think of this JCR circuit as a nonlinear resonator, just like a pendulum.

If the capacitance in the above circuit is the intrinsic capacitance of the Josephson junction, then the resulting resonance frequency is called the *Josephson plasma frequency*. This resonance can interfere with normal circuit operation if it is close enough to primary operating frequencies. In practice, our microwave circuits will have an operating frequency much lower than the plasma frequency.

The Josephson oscillator is an extremely rich physical system, and there are many important phenomena that can be derived from the above equations that we do not consider here, such as Shapiro steps and RSFQ logic. The pendulum analogy carries through to many standard examples. The last that we consider in detail is the superconducting quantum interference device, known as a SQUID. A SQUID is simply two Josephson junctions in parallel, as depicted in Fig. 2.1c. If the junctions are identical, the effective Hamiltonian is simply $2H_T$. Things get interesting when we apply a magnetic field through the loop, which sets up a circulating current in the loop via the Meissner effect. We can use Maxwell's equations to find the effect on each of the junctions. Labelling the two junctions as 1 and 2 and using Stoke's theorem, we find

$$\nabla \times E = -\frac{\partial B}{\partial t} \tag{2.10}$$
$$\oint E \cdot dl = -\frac{\partial \phi_{\text{ext}}}{\partial t} = V_1 + V_2 = \frac{\hbar}{2e}(\dot{\varphi}_1 + \dot{\varphi}_2)$$

where ϕ_{ext} is the magnetic flux passing through the loop. The above relation forces a boundary condition between φ_1 and φ_2

$$\varphi_1 + \varphi_2 + \frac{2\pi}{\phi_0}\phi_{\text{ext}} = 0 \tag{2.11}$$

where $\phi_0 = h/2e$ is the flux quanta. Using a basic trigonometric identity, we can reduce the Hamiltonian of two junctions in parallel to that of a single junction with effective E_J

$$E_J(\phi_{\text{ext}}) = \cos\left(\frac{\pi\phi_{\text{ext}}}{\phi_0}\right) \tag{2.12}$$

for the case that $E_{J,1} = E_{J,2}$.

There are many ways to reason about the SQUID intuitively. Eq. (2.12) arises from the interference between two phase-shifted cosine terms in the Hamiltonian. We can interpret this interference as literal interference between Cooper pairs tunneling simultaneously across both junctions in parallel. When $\phi_{\text{ext}} = \phi_0/2$, the two channels interfere destructively, completely inhibiting the flow of current. A slightly more abstract model is shown in Fig.
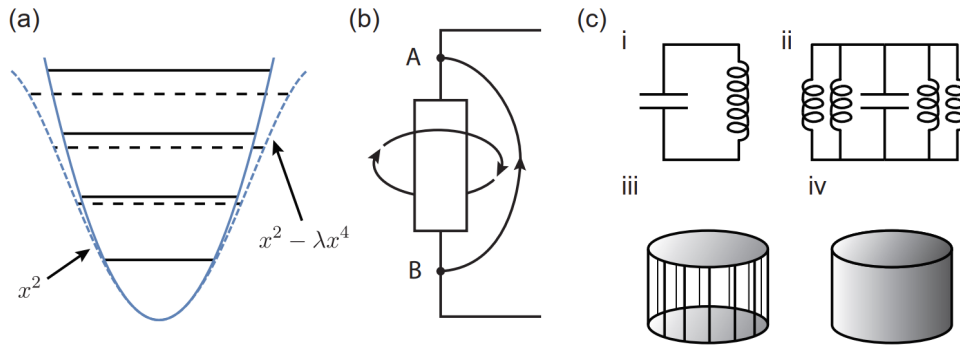
Figure 2.2: (a) A parabolic potential with equally spaced transitions (solid lines) and a quadratic-plus-quartic potential with unequal spacing. Transitions may be addressed individually in the anharmonic case, as they now differ in frequency. (b) Illustration of the integrals used to define branch flux (vertical curved path) and branch charge (loop) of a general circuit element. (c) progressively shrinking an LC circuit until it becomes a microwave cavity.

2.1d. The fixed phase relationship between φ_1 and φ_2 amounts to locking the relative phase of two pendula with a pulley[3]. $\phi_{\text{ext}} = \phi_0/2$ corresponds to putting the two masses exactly out of phase, so that the gravitational (Josephson) energies cancel.

The SQUID has is a ubiquitous device in superconducting electronics. Its magnetic field sensitivity makes it the most sensitive magnetometer ever developed. It also provides a way to tune the resonant frequency of a JC circuit, at the expense of making it sensitive to ambient magnetic field noise. One can play a trade-off between tunability and noise susceptibility by unbalancing $E_{J,1}$ and $E_{J,2}$, so that the cancellation becomes incomplete. The SQUID also shows up unintentionally in large-area junctions. If the flux threading a subregion of the junction becomes comparable to ϕ_0, then we must treat that single junction as many junctions in parallel, each with their own threaded flux. This makes even single Josephson junctions potentially sensitive to flux noise.

2.3   Quantization of a circuit

We have developed a Hamiltonian for the the Josephson junction, the essential ingredient for a superconducting qubit. This nonlinearity will allow us to turn a harmonic LC oscillator into a nonlinear JC oscillator, which in turn lets us selectively drive a single transition in the ladder of states as in Fig. 2.2a. However before we describe superconducting qubits in further detail, it is important to understand how to quantize a general circuit in a bit more detail.

Classical LC circuits like the one shown in Fig. 2.2c (i) behave like classical harmonic oscillators, and there are a number of ways to quantize them of varying levels of generality.

We take a few shortcuts to get to the essential physics as quickly as possible, and refer to other references at the end of this section. We begin with the Hamiltonian for a classical LC circuit

$$H_{\mathrm{LC}} = \frac{1}{2}LI^2 + \frac{Q^2}{2C} \tag{2.13}$$

where I is the current through the inductor and Q is the charge stored on the capacitor. If we were to begin with a Lagrangian for the circuit, then we could immediately identify canonically conjugate coordinates by the relation $p = \partial L/\partial \dot{q}$, and then quantize by promoting p and q to operators satisfying $[q, p] = i\hbar$. To get them from a Hamiltonian, we must find which choice of coordinates reproduces the correct classical equations of motion.[6] We know the classical equation of motion that we seek to reproduce is $\ddot{Q} = \omega^2 Q$ with $\omega = 1/\sqrt{LC}$. In analogy with the harmonic oscillator, we expect the conjugate coordinates to be proportional to I and Q. Let us name our canonically conjugate coordinates $q = \beta Q$ and (suggestively) $\phi = \alpha I$, where α and β are constants to be determined. Hamilton's equations of motion give

$$\dot{q} \equiv \frac{\partial H}{\partial \phi} = \frac{1}{\alpha}LI \tag{2.14}$$

$$\dot{\phi} \equiv -\frac{\partial H}{\partial q} = -\frac{Q}{\beta C}.$$

Combining these by substitution, we find

$$Q\frac{L}{C}\frac{1}{\alpha^2\beta^2} = \ddot{Q}. \tag{2.15}$$

which shows that $\alpha\beta = L$. Alternatively, we could have simply observed that $\dot{Q} = I$, which also implies $\alpha\beta = L$. β by itself is equivalent to choosing a coordinate system, so we are free to take $\beta = 1$ so that $q = Q$. Thus to quantize the oscillator, we take $[\phi, Q] = i\hbar$ with $\phi = LI$. Representing these coordinates in terms of ladder operators, we have

$$\phi = \phi_{\mathrm{ZPF}}(a + a^\dagger) \tag{2.16}$$

$$Q = -iq_{\mathrm{ZPF}}(a - a^\dagger)$$

$$\phi_{\mathrm{ZPF}} = \sqrt{\frac{\hbar Z_{\mathrm{eff}}}{2}}$$

$$Q_{\mathrm{ZPF}} = \sqrt{\frac{\hbar}{2Z_{\mathrm{eff}}}}$$

---

[6]In general, both of these procedures can lead to ambiguities in quantization; classical operators commute, so there may be operator ordering ambiguities when we try to assign a quantum operator to a product of classical operators. However if there is one set of coordinates p_i that only appear in the Lagrangian to second order, then partial derivatives of these terms lead to first-order canonical variables which cannot have ordering ambiguities. We will only deal with linear capacitors, which provides such a guarantee. Quantization of circuits with nonlinearities in both Q and I is an area of ongoing research, so it is surprisingly subtle.

where $Z_{\text{eff}} = \sqrt{L/C}$ is the characteristic impedance of the oscillator.[7] The prefactors ϕ_{ZPF} and Q_{ZPF} quantify the amplitude of vacuum fluctuations in each coordinate of the circuit, as $\langle 0|\phi^2|0\rangle = \phi_{\text{ZPF}}^2$ and likewise for Q. These quantities play an important role the the ultimate behavior of a circuit, and show up in black box quantization.

We might hope that the linear inductor treated in this section can be related to the Josephson inductance of the previous section, and indeed our naming schemes of the variables ϕ and φ are intentionally suggestive. The second Josephson relation connected φ to V. The corresponding relation between ϕ and V for an inductor is the standard relation

$$V = L\frac{dI}{dt} = \frac{d\phi}{dt}. \tag{2.17}$$

Thus Eq. (2.6) is also of the form of an inductor voltage relation if we take $\phi = \varphi\phi_0/2\pi$ for the junction, where $\phi_0 = h/2e$. With this definition of ϕ for a junction, $V = d\phi/dt$ becomes a restatement of the Josephson relations. If we combine the canonical commutation relation with this identification, then we can loosely write $[\hat{\varphi}, \hat{n}] = i$, as we used in exercise 15.

Eq. (2.17) is such a useful relation that it pays to turn things around and define ϕ not as IL or $\varphi\phi_0/2\pi$, but instead as

$$\phi \equiv \int_{-\infty}^{t} V(t')dt' \tag{2.18}$$

Defined as such, ϕ is called the *branch flux*, and it is sufficiently general to be defined for virtually any circuit element. Equally useful is the *branch charge*

$$Q = \int_{-\infty}^{t} I(t')dt'. \tag{2.19}$$

Though it might seem trivial, this definition is useful because it applies to any circuit element, not necessarily just a capacitor. Both the branch flux and branch charges may be defined from the fields generated by (and external to) a circuit element, rather than the detailed dynamics of the charges and currents

$$V(t) = \int_{A}^{B} \vec{E} \cdot d\vec{l} \tag{2.20}$$

$$I(t) = \frac{1}{\mu_0} \oint \vec{B} \cdot d\vec{l}$$

where the contour integral encircles only the circuit element in question and A and B are its terminals as in Fig. 2.2b. These relations, along with the associated branch variables allow one to quantize a very general circuit element[140].

---

[7]We know that $\langle 0|H|0\rangle = \hbar\omega/2$ for any harmonic oscillator. The third and fourth equations of Eq. (2.16) follow if we assume that this zero-point energy is divided evenly between Q and ϕ, as suggested by Virial theorem.

Some resonators in superconducting circuits are true lumped element devices, but it is also common to use microwave cavities. To quantize these, there is a nice argument due to Feynman that directly relates an LC circuit to an electromagnetic resonator. Imagine that we wish to increase the resonant frequency of our LC circuit. $\omega = 1/\sqrt{LC}$, so we should shrink our capacitors and inductors as much as possible. Inductance can be decreased by placing multiple inductors in parallel, as in Fig. 2.2c ii. Each inductance can be decreased by eliminating loops of wire until each is just a straight length of wire made as short as possible, as in Fig. 2.2c iii. Once each wire is so short that it spans only the gap on the capacitor, then all that is left to do is to add more 'inductors' until the interior space of the capacitor is completely enclosed, which results in a cavity! Thus the only distinction between a lumped element circuit and a cavity resonator is that an LC circuit stores electric and magnetic fields in spatially separate regions. Both may be quantized in the same way. This argument also works the other way around. Eq. (2.20) lets us define and quantize lumped element circuits via the electomagnetic fields that they generate, which is a common strategy.

So far, we have quantized LC and LJ oscillators, two of the simplest superconducting circuits. However there is a strange difference between the two. For the LC circuit, Q and ϕ are analogous to x and p, and therefore have their eigenvalues in $\mathbb{R}$. The treatment is not sensitive to the fundamental discreteness of charge. Capacitive energy depends on the spatial distribution of charge, which is continuous even through electrons come in integer amounts. In contrast, our treatment of the Josephson junction restricted the eigenvalues of φ to be between 0 and 2π, and therefore $\hat{n}$ only had integer eigenvalues. The periodicity of φ is also not easily circumvented, as it corresponds to the superconducting phase in the BCS wave function.

The resolution comes from the second Josephson relation Eq. (2.5) and the definition of branch flux. Different circuit elements can have different Hilbert spaces, discrete or continuous, but we can nevertheless couple them by computing how they generate and respond to potential differences. As a lumped element, the Josephson junction couples only to Cooper pairs located very close to the tunnel barrier, but coupling to the capacitor and capacitive energy favors Cooper pairs spread out over the entire superconducting island. The net effect is a device that behaves as if its charge operator is discrete, and the junction mediates tunneling of delocalized Cooper pairs from one side to the other.

We can also interpolate smoothly from a JC circuit to an LC circuit, which has the non-periodic current-phase relationship $I = \phi/L$. Suppose we replace a single Josephson junction with Josephson energy E_J with N junctions in series, each with an increased coupling energy NE_J. As $L_J \propto 1/E_J$, the effective linear inductance is the same as that of the original junction by Eq. (2.8). For the branch flux of the entire network of junctions in series, we have $\phi_{\text{tot}} = \sum_i \phi_i$, so $\varphi_i = 2\pi\phi_{\text{tot}}/N\phi_0$ for each junction individually. Thus the 2π periodic symmetry is broken to a larger period of $2\pi N$. Furthermore, as we take $N \to \infty$, each φ_i becomes very small, so that the nonlinear current-phase relationship for each junction

individually becomes linear via a Taylor expansion

$$I = \frac{2eNE_J}{\hbar} \sin\left(\frac{2\pi\phi_{\text{tot}}}{N\phi_0}\right) \tag{2.21}$$

$$\approx \frac{\phi_{\text{tot}}}{L_J}.$$

Thus the chain of junctions in series becomes an inductor, and our JJ...JC circuit becomes an LC circuit. Qualitatively, it is as if charge becomes continuous because Cooper pairs can now delocalize over the chain.

We have reconciled how some circuits can have discrete charging energy while others have continuous, but what happens when we combine them? This difficulty arises if we put an inductor and a junction in parallel. The inductor breaks the periodic symmetry of the circuit, so that $\phi = 0$ and $\phi = \phi_0$ are now physically distinct states with different amounts of energy. To extend the pendulum analogy that we used to understand the SQUID, it is as if the junction pendulum is connected by a pulley to a torque spring that winds as the pendulum rotates. When shunted by a capacitor, this circuit becomes that of a relatively modern qubit called the capacitively shunted flux qubit.

Other methods of circuit quantization The circuits considered in this section give relatively simple examples because the kinetic and potential energies are stored in a single capacitor and a single inductor respectively. Most circuits will not be so simple, and more general methods will be useful to quantize them. However the essential ingredient that we have used here, that a classical resonance may be replaced with a quantum harmonic oscillator, carries over to more general methods. One is a standard procedures to go from a general lumped-element circuits to a quantum Hamiltonian via methods borrowed from classical circuit analysis[140]. This method does not require one to first diagonalize the circuit, and can handle a circuit with nonlinear inductors or nonlinear capacitors (but not both at the same time). Most of the subtlety lies in writing down a valid classical Hamiltonian for the lumped-element circuit and recognizing independent degrees of freedom.

While this method is important for understanding circuit QED, it is not the most common way to analyze quantum circuits in practice. The current standard is called *black-box quantization*[86], which derives a quantum Hamiltonian from numerical finite-element simulations of a circuit. Standard numerical packages such as COMSOL and HFSS can diagonalize a linear electrodynamics problems into a set of resonant modes. To account for the presence of Josephson junctions, one takes the quadratic approximation of the Josephson potential Eq. (2.8) and places lumped element inductors at the locations of the junctions in simulation. The resulting eigenmodes with frequencies ω_i are then quantized simply by writing down

$$H_{\text{linear}} = \hbar \sum_i \omega_i a_i^\dagger a_i. \tag{2.22}$$

The full Hamiltonian reintroducing nonlinear terms of the Josephson potential is then

$$H = H_{\text{linear}} - \sum_j E_J \left(\cos(\hat{\varphi}_j) - \frac{1}{2}\hat{\varphi}_j^2 \right) \tag{2.23}$$

$$\hat{\varphi}_j = \frac{2\pi\hat{\phi}_j}{\phi_0}.$$

The $\hat{\phi}_j$ operators, one for each Josephson junction, may be expressed as linear combinations of the current quadrature of the finite element eigenmodes

$$\hat{\phi}_j = \sum_i \phi_{ij}^{\text{ZPF}}(a_i^\dagger + a_i) \tag{2.24}$$

where ϕ_{ij}^{ZPF} are the zero point fluctuations from Eq. (2.16), except that now the current quadrature of mode i is divided among the junctions and the electromagnetic energy of other linear parts of the circuit. The ϕ_{ij}^{ZPF} quantities are linearly related to the energy participation ratio, which may be extracted from the finite element simulations[86]. Black-box quantization is an ongoing area of research, and continual improvements are being made to its accuracy and computational efficiency. For further details, we refer the reader to the PyEPR github repository that implements black-box quantization.[8]

2.4 Artificial Atoms

In the previous section, we wrote down some simple circuits and quantized them. We now look at the dynamics of these circuits and learn how they may be operated as qubits. Much of the subtlety of qubit design lies in understanding and minimizing decoherence, so we begin by overviewing the primary sources of it.

Decoherence mechanisms First and foremost, superconducting circuits must remain cold to approach the idealized quantum behavior described above. Aside from working well below the superconducting critical temperature, low temperatures allow for thermal initialization of a qubit to its ground state. As typical superconducting circuits operate at frequencies of a few GHz, millikelvin temperatures are required to get low excited state population (1 GHz corresponds to a temperature of about 48 mK). These extremely low temperatures may be achieved with a commercially available device called a dilution refrigerator. The basic working principle of a dilution refrigerator is described in Fig. 2.3. Low temperatures ensure that nearby degrees of freedom at the qubit frequency are also in their ground state, which mitigates many potential sources of decoherence.

As mentioned at the beginning of this chapter, resistive loss is finite at non-zero frequencies even in a superconductor. Finite inductance leads to potential differences across the superconductor, which in turn drives current through the fraction of electrons that are

---

[8]https://github.com/zlatko-minev/pyEPR

not superconducting. Fortunately at low temperatures, the residual non-superconducting electron density decreases exponentially with temperature ($1 - n_s \propto e^{E_g/k_b T}$), and qubits tend to operate far below the superconducting critical temperature. However dissipation in other nearby materials such as the substrate onto which a qubit is fabricated and (in particular) surface interfaces pose major sources of loss. Dielectric loss is mitigated by using high-resistance substrates like silicon and sapphire, and by designing qubits that store more of their electromagnetic energy in the nearby vacuum[85, 102]. Some experiments have even etched away the substrate in the vicinity of small features where electric field is concentrated[19]. A good deal of qubit fabrication development lies in minimizing these sources of loss by making and testing high-Q resonators.[9]

Finally, microwave circuits can lose energy by radiating microwave photons. One must take care to place the qubit in an isolated environment with no resonances near the qubit frequency. Qubits are place in conducting enclosures that prevent dipole emission into free space. Finite element solvers like COMSOL and HFSS are used to identify electromagnetic resonances of a circuit and its enclosure to eliminate parasitic resonances. Even the substrate can support resonances that harm qubit coherence. These can be mitigated with through-silicon vias that gap these modes up to higher frequencies.

Quasiparticles constitute another major source of qubit decoherence. While in principle they can lead to finite resistance and hence could be characterized by making and testing resonators, their primary effect occurs when one tunnels across a Josephson junction. This is because the junction supports a large potential difference. When a quasiparticle tunnels across the junction, its kinetic energy can change by the qubit energy level spacing and changing the qubit state. Curiously, quasiparticles have been observed to excite the qubit more often than they lead to decay, indicating that they do not follow an equilibrium Boltzmann distribution. One way to mitigate quasiparticles is to enhance the rate at which they are dissipated by putting a normal metal in contact with the qubit[108]. More commonly one attempts to limit how many are generated in the first place by placing infrared filters (such as Eccosorb) on microwave control lines and placing qubits in a sealed light-tight can, as in Fig. 2.4a (hermetic sealing is typically not possible, due to the need to evacuate air from the can during cool-down). It is also important to block radiation from hotter portions of a dilution refrigerator from entering the low temperature region. A commonly used method is to place aluminum tape over holes in various stages and check for light leaks using a flashlight, though it is unclear if this method is sufficient.

Loss due to resistance and quasiparticles has decreased steadily with improved fabrication and filtering. A decoherence source that has been more difficult to mitigate directly is charge fluctuations. The exact source is somewhat unclear, but it appears that the charges in the dielectric substrate or the junction fluctuate, so that the qubit experiences a gradually changing electric field. Charge fluctuations can be modelled by changing the capacitive

---

[9]$Q = \omega/\kappa$ is called the quality factor of a resonator. It measures the number of oscillations the resonator will undergo before the oscillations damp out due to loss.
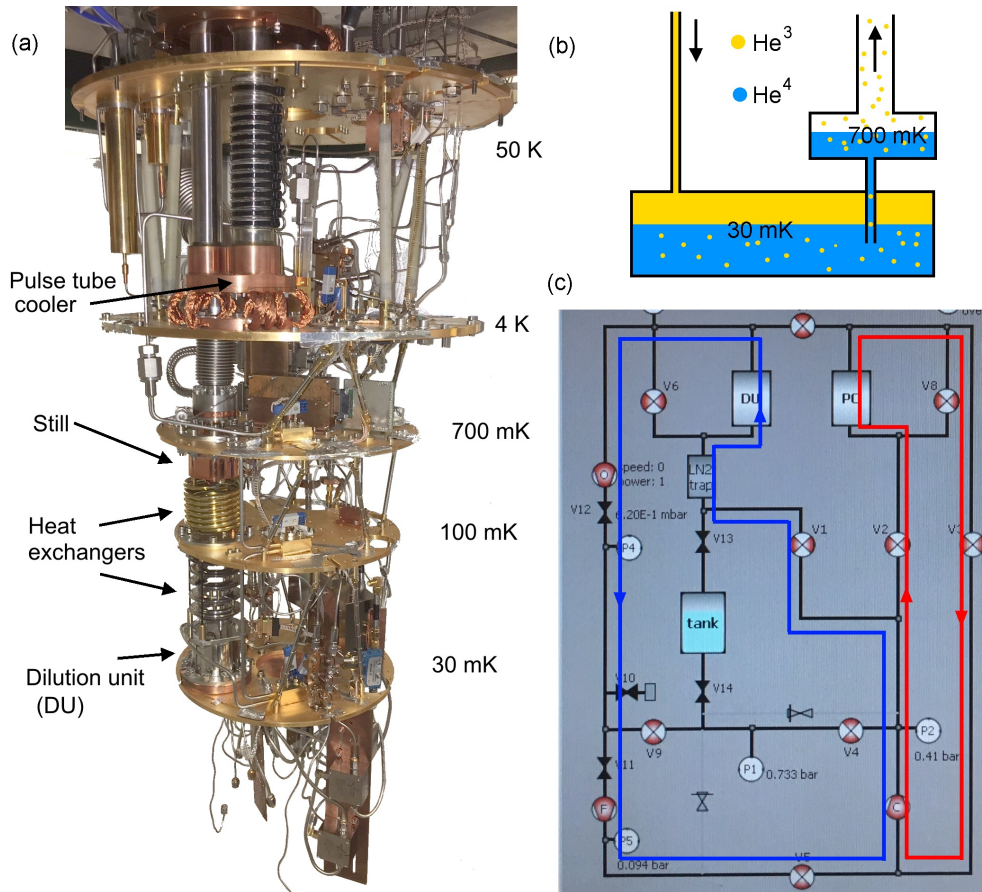
Figure 2.3: (a) An opened dilution refrigerator. Stages are cooled by 3 different mechanism. The 4 K plate is cooled by expansion cooling of Helium via a pulse tube cooler. A circuit of a He^3/He^4 mixture is pumped into the dilution unit and then pumped out of the still with a turbo pump (not pictured). The turbo pump cools the still via evaporative cooling of liquid He^3, which has a higher vapour pressure than He^4. (b) The principle of dilution refrigeration, which is analogous to evaporative cooling. He^4 is bosonic, and has a smaller de Broglie wavelength than He^3 due to its higher mass. Both of these properties allow a smaller interatomic spacing between He^4 atoms, which in turn leads to a lower energy via the Van der Waals force. Therefore at low temperatures, He^3 and He^4 phase-separate like oil and water, as depicted. However a He^3 atom can lower its energy by entering the He^4 phase. Due to the Pauli exclusion principle, He^3 atoms entering the He^4 phase must occupy increasingly higher energy states. Eventually the Fermi energy exceeds the energy gained by entering the He^4 phase. This balance occurs when the He^4 phase contains 6.4% He^3, even at zero temperature! This allows finite cooling power all the way down to $T = 0$ (unlike evaporative cooling, where the vapor pressure goes to zero below the boiling point). (c) Circulation of helium during the initial precooling of the fridge (red), used to go from room temperature to about 10 K using the pulse tube, and dilution cooling (blue).
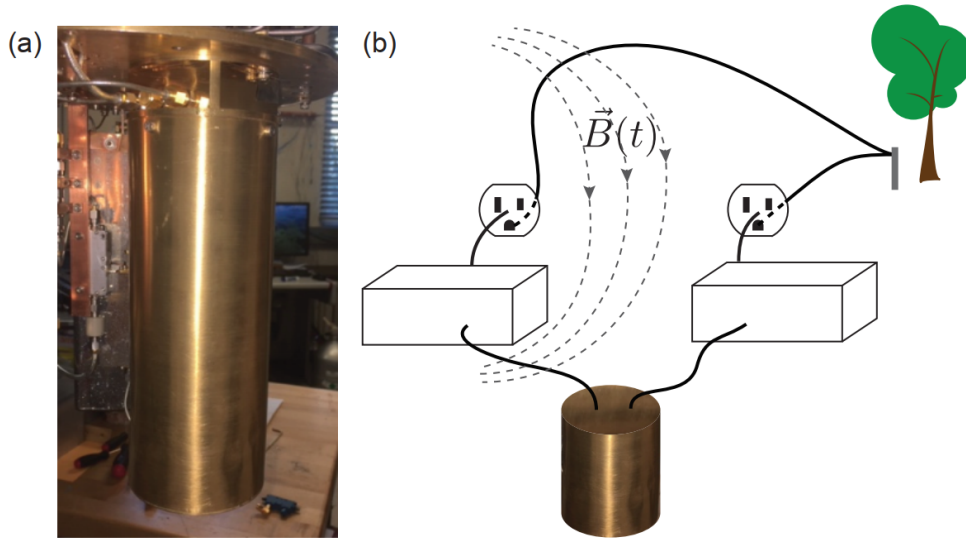
Figure 2.4: (a) A light-tight can for housing superconducting circuits. DC and microwave control lines enter through tightly sealed bulkheads. A venting hole is necessary to allow evacuation before cooling to mK temperatures. (b) A classic example of a ground loop. Ground loops can be shrunk by running all equipment from the same breaker and ensuring electrical isolation.

energy to account for a randomly varying charge offset n_g

$$H_{\text{cap}} = \frac{2e^2}{C}(n - n_g)^2. \tag{2.25}$$

n_g can fluctuate on both short and long timescales, and it is not at all uncommon for n_g to shift by one or more units of charge over minutes to hours. It is commonly modelled to have a $1/f$ power spectrum. The outlook for eliminating charge fluctuations is somewhat unclear. Surface charges are a notoriously difficult problem in many quantum systems, and are one of the primary limiting factors for trapped ion qubits and near-surface NV qubits. Fortunately, superconducting qubits can be designed to be insensitive to n_g as we will see. Much of the success of superconducting qubits has come from the transmon qubit, which is designed to be insensitive to charge noise.

In principle, any physical parameter of a qubit circuit can fluctuate, including E_J or E_C, though these do not appear to change much in practice. The last source of decoherence we consider here is magnetic flux noise, which is in a sense dual to charge noise.[10] As we saw previously, a pair of Josephson junctions in a loop serves as an exquisitely sensitive magnetometer, which is great for sensing but terrible for a qubit! A qubit that contains a SQUID as its inductive element can be tuned by a magnetic field passing through the loop,

---

[10]Note that in our analysis of the SQUID, φ appeared in the argument of the Josephson energy, much as n_g appears in the argument of the capacitive energy.

which is sometimes worth the additional susceptibility to flux noise. Flux noise, like charge noise, often follows a $1/f$ power spectrum, and recent evidence suggests that it arises largely from diamagnetic impurities like oxygen. It can also come from wandering superconducting vortices in type II superconductors, which can be pinned in place by adding holes to any large-area metal regions (which is called 'waffling').[11] An insidious source of flux noise is ground loops; if magnetic field fluctuations from the laboratory can drive current through various parts of the system at large, these currents can generate magnetic fields near the qubit. An example of a ground loop is shown in Fig. 2.4b. Since most qubits derive their inductance from Josephson junctions, the rest of the device tends to have very low mutual inductance with its surroundings. Therefore non-tunable qubits are insensitive to flux noise. However some qubit designs such as the flux qubit necessarily have loops containing Josephson junctions and are therefore sensitive to flux noise by design. Typically there exists some kind of sweet spot where the first-order sensitivity (*i.e.* first derivative of the qubit energy vs. flux) goes to zero.

Cooper pair box Perhaps the simplest qubit is the *Cooper pair box*, which is a Josephson junction and a capacitor in parallel as in Fig. 2.1a. Including the possibility of a charge offset n_g, the Hamiltonian is

$$H = 4E_C(n - n_g)^2 - E_J\cos(\varphi) \tag{2.26}$$
$$E_C = \frac{e^2}{2C}$$

where C is the capacitance, which includes any capacitance of the junction itself (which is in parallel with the lumped element capacitance of the circuit). E_C is the capacitive energy associated with one electron's worth of charge displaced in the capacitor. The factor of four arises because $\hat{n}$ counts Cooper pairs, each of which contains two electrons.

In the Cooper pair box, $E_C > E_J$. A large E_C energetically favors states with small amounts of charge. As $\hat{n}$ and $\hat{\varphi}$ do not commute, such states are highly delocalized in $\hat{\varphi}$ by the uncertainty principle. If we interpret the Cooper pair box circuit as a pendulum, as we did in Fig. 2.1, then it is as if the mass is so small that it is completely delocalized, constantly spinning much like an electron bound to a nucleus. As $\hat{n}$ behaves like the (angular) momentum operator, the gate charge n_g behaves like a magnetic field, so that the gate charge creates a Zeeman splitting between clockwise and counterclockwise motion.

Exercise 16 *If the Hamiltonian for a charged particle in a magnetic field is unfamiliar to you, show that Hamilton's equations of motion applied to $H = (\vec{p} - e\vec{A})^2/2m + eV(x)$ reproduces the Lorentz force. Note the similarity between this Hamiltonian and Eq. (2.26).*

As the Cooper pair box eigenstates are localized in charge, they are more easily thought of in that basis The energy levels of the qubit correspond to displacing different numbers

---

[11]Vortices can also lead to resistive dissipation, as they experience a Lorentz in the presence of current flow. As vortex motion can experience scattering, they offer a way to damp current flow in a superconductor. Type II superconductors would not have conductivities much better than ordinary metals without flux pinning.
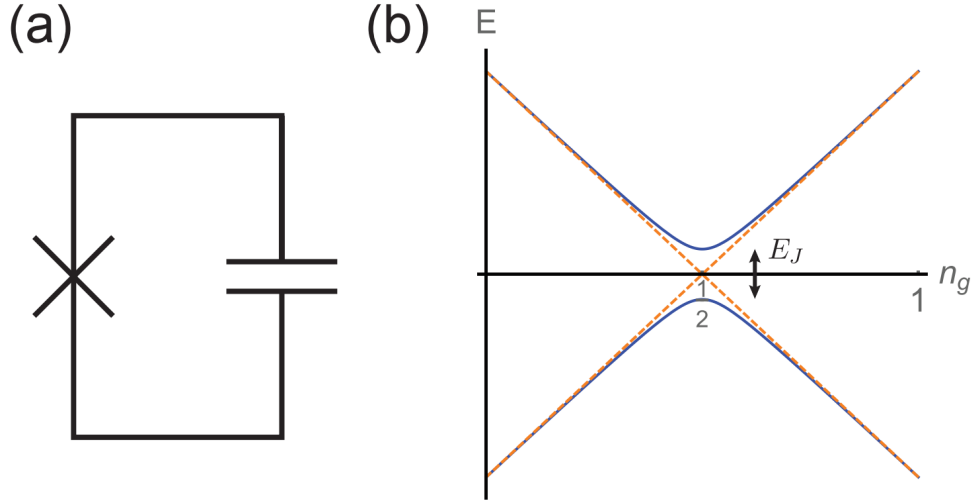
Figure 2.5: (a) The lumped-element circuit for a Cooper pair box (b) The energy levels of a Cooper pair box in the 2-level approximation. The sweet spot occurs at $n_g = 1/2$.

of individual Cooper pair across the junction. A state with $\hat{n}$ close to n_g are energetically favored. Supposing $0 < n_g < 1$, it becomes relatively straight-forward to think of the Cooper pair box as a qubit, with the qubit subspace given by states of 0 or 1 Cooper pairs displaced across the junction. The junction couples these two states, so that the actual eigenstates are superpositions of $n = 0$ and $n = 1$. Computing the Hamiltonian matrix elements on these states, we can reduce the Cooper pair box Hilbert space to that of a qubit, with an effective Hamiltonian[119]

$$H = 2E_C\sigma_z(1 - 2n_g) - E_J\frac{\sigma_x}{2} \tag{2.27}$$

where we have dropped a scalar energy offset. Eq. (2.27) is the Rabi Hamiltonian of chapter 1. The eigenstates of H are the bare qubit states, analogous to the lab frame Hamiltonian from chapter 1. The Cooper pair box can be controlled by applying an external electric field, which adds a time-dependent term proportional to σ_z. As per the Rabi model discussed in chapter 1, oscillating this external field at the qubit frequency can drive transitions between eigenstates, allowing for qubit control.

To consider the coherence properties of the Cooper pair box, the eigenstates of H are plotted in Fig. 2.5c. There exists a 'sweet spot' at $n_g = 1/2$ where the transition frequency is first-order insensitive to fluctuations in the value of n_g, which is an avoided crossing between the $n = 0$ and $n = 1$ states split by the junction energy. The junction mediates rapid tunneling between the $|n = 0\rangle$ and $|n = 1\rangle$ states, and the new eigenstates are instead equal superpositions of $|0\rangle$ and $|1\rangle$ with different relative phases. Usefully, decreasing sensitivity to charge noise by operating at the sweet spot does not eliminate sensitivity to an oscillating control field.

From the standpoint of decoherence, this smearing out of the $|n = 0\rangle$ and $|n = 1\rangle$ states hides information about the qubit state from the environment. When the eigenstates are unequal superpositions of these states, they have different expectation values of charge and hence are easily distinguished by anything nearby with a dipole moment. Any form of measurement constitutes a source of decoherence, and thus these states are quickly dephased. At the sweet spot, the expectation value of charge is the same for both qubit states. To learn about the qubit state, the environment must be able to distinguish the 0 and 1 charge states faster than the period of oscillation $\hbar/E_J$. Otherwise the measurement outcome washes out over both states and the environment obtains no information of the qubit's state.

The Cooper pair box derives its stability from E_J at the sweep spot, so fluctuations in this parameter set a fundamental limit on its coherence. In practice however, fluctuations in n_g are severe, such that not only do second-order effects come into play, but also that it can be difficult to maintain the qubit near its sweet spot for long periods of time. However note that the larger the value of E_J, the greater the splitting and hence the wider the range over which the energy bands remain flat.

The transmon qubit The above reasoning suggests that we stand to lose little by further increasing E_J. as we increase e_j, states will be forced to localize near $\varphi = 0$, and thus will delocalize beyond the $n = \{0, 1\}$ subspace. this change of roles favors visualizing the eigenfunctions in the φ basis, like the mass on a pendulum. the transmon regime $e_j \gg e_c$ corresponds to a region in which gravity is strong enough that the pendulum is well-localized around $\varphi = 0/$major[62].

The eigenstates and eigenvalues of the transmon (and by extension the Cooper pair box) can be found exactly in terms of Mathieu functions. These eigenfunctions are plotted in Fig. 2.6a for $E_J/E_C = 50$. As the state is localized near the bottom of the cosine potential, the lowest-energy eigenstates resemble those of a harmonic oscillator. We can use this observation to find a closed-form expression for the approximate eigenenergies, which is useful for building intuition. For states near $\varphi = 0$, we can expand the cosine in a power series of φ in the Hamiltonian

$$H \approx 4E_C(\hat{n} - n_g)^2 + E_J\frac{\hat{\varphi}^2}{2!} - E_J\frac{\hat{\varphi}^4}{4!}. \tag{2.28}$$

The Hamiltonian is no longer 2π periodic in ϕ, so we should replace the charge and phase operators with their continuous counterparts Eq. (2.16) (using $Q = 2e\hat{n}$ and $\phi = 2\pi\varphi/\phi_0$). This approximation is equivalent to our approximation that φ is localized. Eigenstates are so delocalized in $\hat{n}$ that the discreteness of charge can be neglected.

Expanding out the capacitive energy term, we recognize the gate charge as a drive term $8E_C n_g \hat{n}$. From our discussion of coherent states and displacement operators in chapter 1, we see that we can eliminate this offset charge by entering a displaced frame with $U = \exp(-in_g\hat{\phi})$ (using $U(\hat{n} - n_g)U^\dagger = \hat{n}$). This is our first hint of reduced sensitivity to charge noise. We can also see that reducing charge noise sensitivity does not prevent us from driving the qubit. U transforms out the static offset $\bar{n}_g$. If $n_g(t)$ oscillates about the mean, then H contains a residual term $8E_C(n_g(t) - \bar{(n)}_g)\hat{n}$. Fortunately the $1/f$ power spectrum of
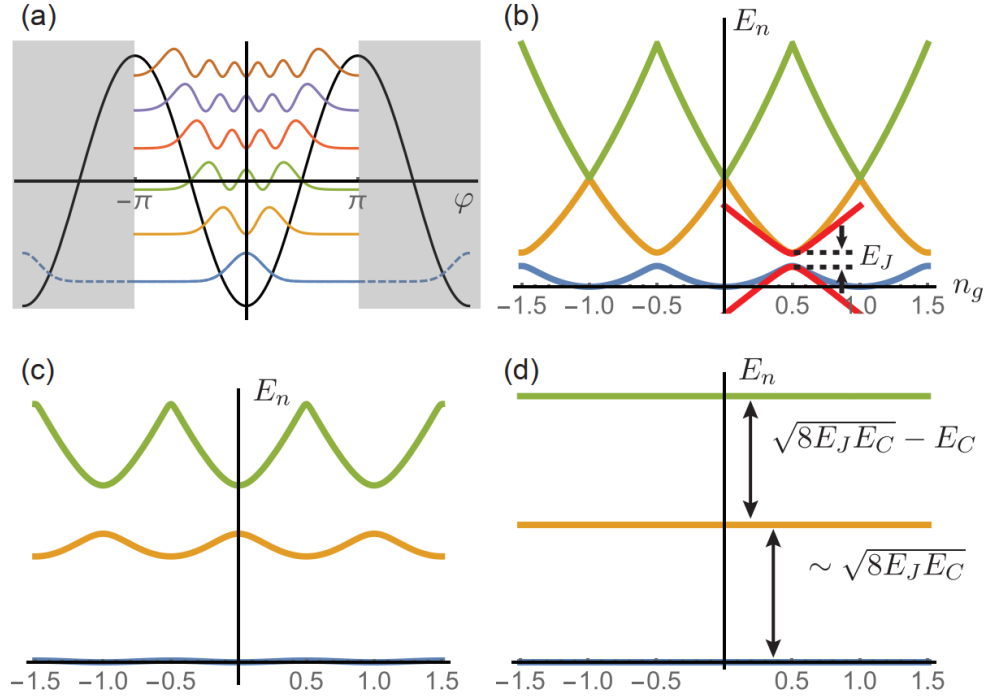
Figure 2.6: (a) Exact eigenstates of the transmon for $E_J/E_C = 50$, plotted in the φ basis. The ground state is well-localized around $\varphi = 0$, and is unlikely to tunnel into an adjacent well of the cosine potential. (b) Energy of the first three transmon/Cooper pair box eigenstates as a function of n_g for $E_J/E_C = 1/2$. Approximate energies of the first two levels from Fig. 2.5 are replotted for comparison. (c) Same but with $E_J/E_C = 5$. (d) Same but with $E_J/E_C = 50$, at which point the energies depend very weakly on n_g. Figure based on [62].

$n_g(t)$ becomes negligible near the qubit frequency where the rotating wave approximation would allow n_g-driven transitions, but we can generate such a term at will by applying a time-dependent electric field to the qubit.

Converting Eq. (2.28) into ladder operators yields a harmonic oscillator with a non-linear correction

$$\sqrt{8E_J E_C}(a^\dagger a + 1/2) - \frac{E_C}{12}(a + a^\dagger)^4 \tag{2.29}$$

Computing the approximate eigenenergies is then a straightforward application of first-order perturbation theory

$$E_n \approx \sqrt{8E_J E_C}(m + 1/2) - \frac{E_C}{12}\langle m|(a + a^\dagger)^4|m\rangle \tag{2.30}$$
$$= \sqrt{8E_J E_C}(m + 1/2) - \frac{E_C}{4}(2m^2 + 2m + 1)$$

where we have labeled the eigenstates with m to avoid confusion with charge $(\hat{n})$. Crucially, n_g does not appear in the spectrum, indicating insensitivity to charge noise. To use our system as a qubit, we require the $0 \leftrightarrow 1$ and $1 \leftrightarrow 2$ transitions are well-separated in frequency, so that we can preferentially address the lowest one. The relevant figure of merit is the relative anharmonicity

$$\alpha_r \equiv \frac{E_{12} - E_{01}}{E_{01}} = \frac{E_C}{\sqrt{8E_J E_C} - E_C} \tag{2.31}$$
$$\approx \sqrt{8E_C/E_J}.$$

Thus as we enter the transmon regime by increasing E_J/E_C, it becomes more difficult to address the lowest-level transition without inadvertently driving higher transitions. In practice we require longer, spectrally narrower control pulses as the anharmonicity decreases.

Eq. (2.31) indicates a trade-off between anharmonicity and charge noise-induced decoherence. We also saw that the above treatment cannot estimate the effect of charge noise. To understand how charge noise comes in, it pays to return to the pendulum analogy. The charge offset acts as a magnetic field applied parallel to the axis of rotation. Classically, the applied field generates a radial Lorentz force, which has no effect on the dynamics. However quantization of angular momentum enforces boundary conditions on the allowed energy values. This boundary condition is effected by the applied magnetic field, and leads to the Zeeman splitting (via the Aharonov-Bohm phase). Thus a charge offset can only affect the energy spectrum if the probability for a 2π excursion in φ is non-negligible; if $\psi(\pm\pi) = 0$, then the periodic boundary condition is trivially satisfied and the magnetic field has no effect.

There are numerous ways to compute the charge noise sensitivity of the energy spectrum (called charge dispersion)[62, 41]. Here we merely aim to estimate how the effect scales with the ratio E_J/E_C. Given the above argument, we expect charge noise to scale with the probability for the state to tunnel by 2π into an adjacent well of the cosine potential. We can estimate the probability for this occurrence using the harmonic oscillator wave functions. Making the correspondence $x \leftrightarrow \hat{\varphi}$ and $p \leftrightarrow \hat{n}$, we have

$$|\psi_0(x)|^2 \propto \exp\left(-\frac{m\omega x^2}{\hbar}\right) \tag{2.32}$$
$$\rightarrow \exp\left(-\sqrt{\frac{E_J}{8E_C}}\varphi^2\right).$$

Thus the probability to tunnel by a distance $\varphi = 2\pi$ decreases *exponentially* with $\sqrt{E_J/8E_C}$. If we consider energy levels above $n = 0$, the standard deviation of $|\psi(\varphi)|^2$ scales as $\sqrt{n}$,[12] so that higher levels are more susceptible. These scaling factors match the results found using

---

[12]$E \propto n^2$, and the classical turning point at which all energy is potential energy occurs when $E = m\omega^2 x_{\max}^2/2$. Thus $x_{\max} \propto \sqrt{n}$.
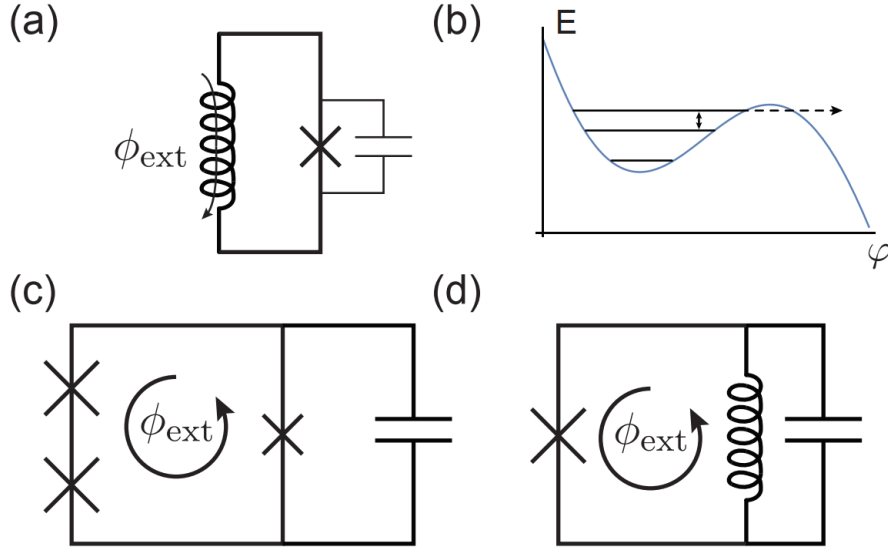
Figure 2.7: (a) Circuit for a phase qubit, which operates with a large flux bias ϕ_{ext}. Circuit capacitance is provided by the intrinsic capacitance of the junction. (b) Readout method of the phase qubit. $|e\rangle$ state is mapped to the third level, which tunnels out of a single well of the cosine potential and down the parabolic potential of the inductor, producing a macroscopic readout signal. (c) The capacitively shunted flux qubit (d) Fluxonium qubit. Inductance is provided by a chain of Josephson junctions in series.

the WKB approximation and the analytic solution in terms of Mathieu functions. In Fig. 2.6b-d we plot the energies of the lowest few transmon eigenenergies as a function of n_g for three values of the ratio E_J/E_C. Fig. 2.6d plots the bands for $E_J/E_C = 50$, a common value that balances the trade-off well enough to make charge dispersion negligible. For a qubit with a resonant frequency of 6 GHz, the absolute anharmonicity is $(E_{12} - E_{01})/h \approx 200$ MHz, allowing for control pulses lasting on the order of 10s of nanoseconds.

Flux and phase qubits In the following chapters, we only work with transmon qubits, which have been the workhorse of the field for some time. Coherence times have improved steadily since its introduction, and experiments have yet to hit any fundamental limits. Nevertheless one of the major appeals of the qubit is its simplicity as a circuit. As fabrication and qubit control has improved, researches continue to revisit old designs and explore new parameter regimes. The full range of circuits is relatively unexplored, and the coming years will likely bring many new forms of superconducting circuits with unanticipated properties. We now qualitatively outline the other major classes of qubits.

The most well-explored class of qubits outside of the transmon are flux qubits, which modify the inductive part of the potential in some way, but are otherwise similar to the transmon in their topology. An early iteration is the phase qubit[83] shown in Fig. 2.7. The flux potential consists of a parabola from the inductor modulated by the cosine potential

of the Josephson junction. In these qubits, $E_J \gg E_C$ to the extent that the qubit is very nearly harmonic. The small capacitive energy greatly reduces sensitivity to charge noise. To operate as a qubit, the circuit's inductance is given a large current bias, which offsets the parabolic portion of the potential away from $\phi = 0$ and thus introduces a large asymmetry in the potential as seen in Fig. 2.7.

Phase qubits have a somewhat unusual readout mechanism shown in Fig. 2.7b. The flux bias is chosen so that only a small number of states are bound. To read out the state, one pulses the $1 \leftrightarrow 2$ transition. The 2 state rapidly tunnels into the quasi-continuum of the parabolic potential, which creates a large change in current and hence a macroscopically large voltage across the junction via the second Josephson relation.

The *flux qubit*[88] is another major type of qubit that has garnered increasing attention recently after a period of relatively little use. The circuit for a modern incarnation, the capacitively shunted flux qubit[153] is shown in Fig. 2.7c. In the simplest picture of this circuit, the φ potential consists of a sum of $\cos(\varphi)$ and a $\cos(\varphi/2)$ terms. The relative phase between these terms can be modulated by applying an external flux, just as in the SQUID. In reality, Josephson junctions have intrinsic capacitances, so that the network of junctions has more than one degree of freedom and a 2D Josephson potential is useful for understanding the full picture.

A further incarnation is the fluxonium qubit[76] shown in Fig. 2.7d, which contains a very large inductance generated by a chain of Josephson junctions in series. Both the flux and fluxonium qubits can have large inductive energies while preserving large nonlinearities. The effective potential of a fluxonium qubit is shown in Fig. 2.8. Not only does the potential deviate substantially from a parabola, but the resulting level structure can be tuned by the external flux. Nearly degenerate states to not couple strongly to each other, so one can generate a fairly long-lived qubit in Fig. 2.8a. The higher state can decay to either of the two lower states. This level configuration is called a Λ system, which is the basis for many fundamental AMO effects like electromagnetically induced transparency (EIT) and stopped light. At another flux bias point like Fig. 2.8c, the two higher levels become degenerate, which then both decay to a common lower state. This configuration is called a V system. An interesting effect that arises in this setting is quantum beats, in which spontaneous emission from the two levels can interfere, leading to coherent oscillations.

2.5   Microwave amplification and detection

Input-output theory Inductors, capacitors and Josephson junctions form the essential ingredients of superconducting qubits. To understand readout and amplifiers quantum mechanically, we need a formalism to treat non-classical input and output fields. This section provides our first treatment of dissipation, in which a degree of freedom interacts irreversibly with our system. For amplifiers, it is convenient to work in the Heisenberg picture, and the resulting formalism is called input-output theory. We treat dissipation in the Schrödinger picture in chapter 3.
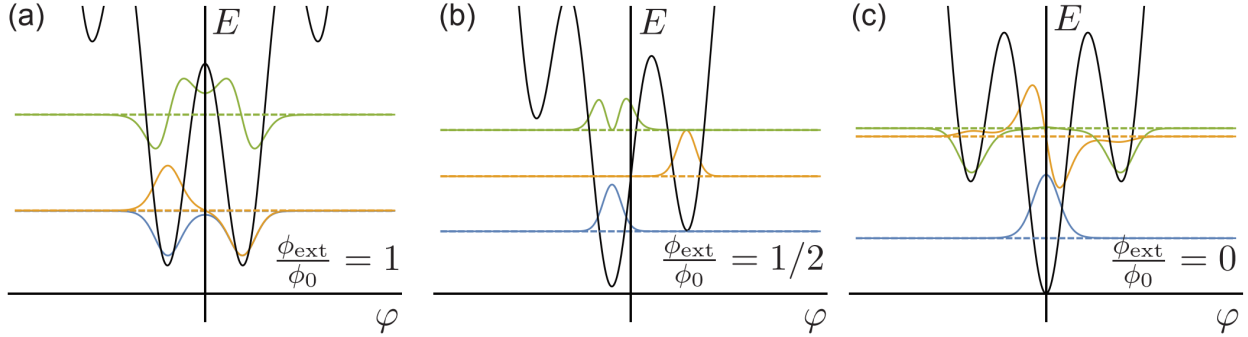
Figure 2.8: Potential for the fluxonium qubit, neglecting junction capacitances so that the φ potential is 1D. Plots also include the lowest 3 eigenvectors placed vertically at their corresponding eigenenergies, derived numerically. Physical parameters match those from [76], but with twice as large a value of E_J to enhance the degeneracies. Flux bias shifts the cosine portion of the potential relative to the parabolic portion. (a) Flux bias which leads to a Λ-type system (b) Ladder system (c) V-type system.

Several scenarios in circuit QED involve a system coupled to a continuum of harmonic oscillators. A qubit radiating into free space couples to an infinite number of normal modes of the vacuum. An infinitely long transmission line also has an infinite number of normal modes, so anything coupled to even a 1D waveguide can be described in roughly the same way. We can treat these systems very general with the following Hamiltonian

$$H = H_{\text{sys}} + \sum_k \hbar\omega_k b_k^\dagger b_k - i\hbar(f_k a^\dagger b_k - f_k^* a b_k^\dagger). \tag{2.33}$$

Here k is an index that labels modes b_k by their wave vector. Each mode has energy $\hbar\omega_k$. H_{sys} is the system of interest, such as a qubit or amplifier, and a is a system operator. The only assumption that we make about the system is that $[a, a^\dagger] = 1$, so that Eq. (2.33) represents swap coupling to the modes b_k like the Jaynes Cummings Hamiltonian. While this precludes H_{sys} being an ideal qubit, all of the qubits that we will consider in circuit QED are actually nonlinear bosonic oscillators, so this commutation relation is be satisfied.

We depict Eq. (2.33) loosely in Fig. 2.9a. We wish to find the dynamics for the system operator a in the Heisenberg picture, and ultimately to look at the relationship between incident and outgoing fields, which both consist of excitations of the b_k modes. We know from basic wave analysis that a continuum of modes can be decomposed into travelling wave packets that are spatially localized. In a 1D waveguide in the absence of dispersion, we expect these temporal wave packets to be time independent, except at the instant when they overlap with our system. In what follows, we recover this intuition mathematically by solving for the dynamics induced by this instantaneous interaction, largely following [20].

We begin with the Heisenberg equations of motion for b_k

$$\dot{b}_k = \frac{i}{\hbar}[H, b_k] = -i\omega_k b_k + f_k^* a \qquad (2.34)$$

This is an inhomogeneous first order differential equation in which the system acts as a drive on each b_k mode. The standard solution is

$$b_k(t) = e^{-i\omega_k(t-t_0)}b_k(t_0) + \int_{t_0}^t dt' e^{-i\omega_k(t-t')} f_k^* a(t') \qquad (2.35)$$

where the first term gives time evolution based on initial conditions and the second term is the b_k mode Greens function convolved with the system operator and represents driving of the transmission line by the system. With the full solution of the b_k mode in hand, we next find the equations of motion for a and substitute our solution

$$\dot{a} = \frac{i}{\hbar}[H_{\text{sys}}, a] - \sum_k f_k b_k \qquad (2.36)$$

$$= \frac{i}{\hbar}[H_{\text{sys}}, a] - \sum_k e^{-i\omega_k(t-t_0)}b_k(t_0)$$

$$- \sum_k |f_k|^2 \int_{t_0}^t e^{-i\omega_k(t-t')}a(t')dt'$$

The second term is the particular linear combination of the b_k modes that drives the system at time t, which we recognize as the operator for a temporally localized wave packet. We label this operator $b_{\text{in}}[t]$, using brackets instead of parenthesis to emphasize that t is a label like k; the $b_{\text{in}}[t]$ operators have no time dependence except at the moment that they interact with the system.[13] The third term simplifies to a δ function if we assume make the *Markov approximation* that f_k is independent of k. We will have more to say about it in chapter 3. Applying these simplifications yields

$$= \frac{i}{\hbar}[H_{\text{sys}}, a] - \kappa \int_{t_0}^t \delta(t-t')a(t')dt' - \sqrt{\kappa}b_{\text{in}} \qquad (2.37)$$

$$b_{\text{in}}[t] \equiv \frac{f}{\sqrt{\kappa}} \sum_k e^{-i\omega_k(t-t_0)}b_k(t_0)$$

where we have defined the rate κ so that $\sum_k |f|^2 \exp(-i\omega_k(t-t')) = \kappa\delta(t-t')$, which naturally takes the density of states into account. While it may seem unusual that our b_{in}

---

[13]This is not to say that the wave packet associated with $b_{\text{in}}[t]$ does not change in time. It propogates at the group velocity and may also disperse as dictated by ω_k. The point is rather that the excitation of the mode, described by a superposition of different numbers of photons, remains constant as it propogates.

operators have units of $\sqrt{\text{Hz}}$, the result appears natural when we compute the commutation relation

$$[b_{\text{in}}[t], b_{\text{in}}^\dagger[t']] = \delta(t - t') \tag{2.38}$$

where we have used the fact that $[b_k, b_{k'}^\dagger] = \delta_{k,k'}$. The unusual units arise because the δ function has the inverse units of its argument. Physically, $b_{\text{in}}^\dagger b_{\text{in}}$ has units of Hz because it quantifies power, just as $a^\dagger a$ measures energy.

The second term of Eq. (2.37) has an interesting interpretation. It arose from the system driving the b_k modes, and then computing in turn how this contribution to the b_k drives the system. The integral of half of a δ function is $1/2$, so the second term simplifies and we arrive at the important *input-output relation*

$$\dot{a} = \frac{i}{\hbar}[H_{\text{sys}}, a] - \frac{\kappa}{2}a(t) - \sqrt{\kappa}b_{\text{in}}[t]. \tag{2.39}$$

We see that this 'back-reaction' term is actually damping, or radiative decay. The system's amplitude *i.e.* its quadrature operators $a \pm a^\dagger$ damp at a rate of $\kappa/2$, just like a harmonic oscillator under the influence of friction.

Just as we derived a relation for the input field, one can derive a similar relation for the output field by solving Eq. (2.34) using the advanced Greens function instead of the retarded Greens function (the state of each b_k at time t is determined by the future evolution of $a(t')$ just as well as by the past). The end result is

$$\dot{a} = \frac{i}{\hbar}[H_{\text{sys}}, a] + \frac{\kappa}{2}a - \sqrt{\kappa}b_{\text{out}}[t] \tag{2.40}$$

$$b_{\text{out}}[t] = \frac{f}{\sqrt{\kappa}}\sum_k e^{-i\omega_k(t-t_1)}b_k(t_1)$$

where t_1 is in the distant future just as t_0 was taken to be in the distance past. Eq. (2.40) is a sort of time-reversed version of Eq. (2.39). Subtracting Eq. (2.39) from Eq. (2.40) yields a relation between the input and output fields

$$b_{\text{out}}[t] = b_{\text{in}}[t] + \sqrt{\kappa}a(t). \tag{2.41}$$

The physical interpretation is that b_{out} consists of input modes reflected off of the system and an additional field radiated by the system itself, as depicted in Fig. 2.9a. In general, the b_{out} modes could constitute a reflected or transmitted signal, depending on whether the b_k operators describe a single-sided transmission line as depicted in the figure or a transmission line (or higher-dimensional region) that extends from the system in all directions. The mathematical description of both scenarios is essentially the same, though we will mostly encounter the single-sided case in circuit QED.

Eq. (2.41) and Eq. (2.39) are extremely useful, offering a simple way to treat many dissipative systems. In this chapter, we will use them to derive and understand parametric
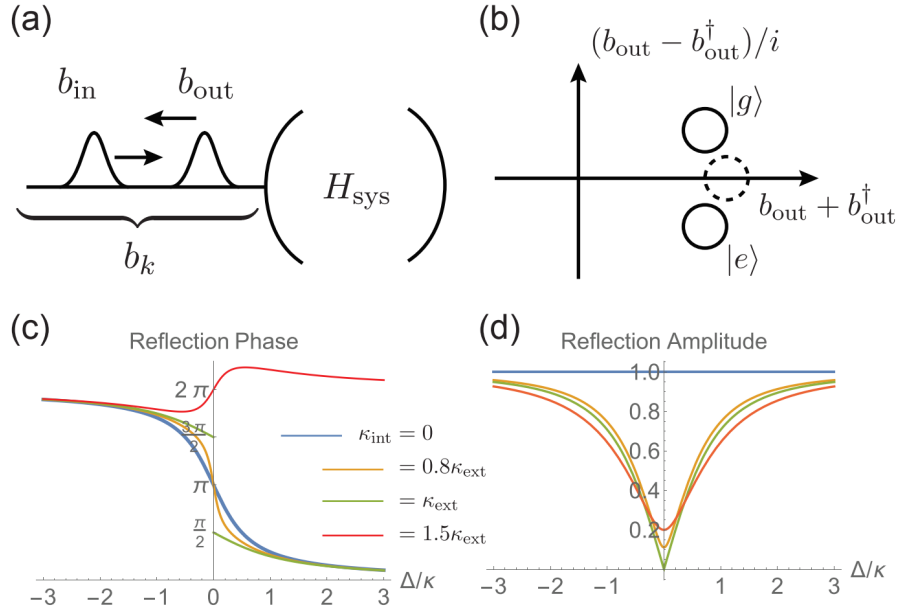
Figure 2.9: (a) A transmission line coupled to a cavity or other system. In input-output theory, we break the degrees of freedom on the transmission line into ingoing and outgoing modes. (b) Coherent state representation of qubit readout. Dashed circle represents in the input state, while the solid circles represent the reflected output states conditioned on the qubit state. (c) Reflected phase of a coherent state incident on a cavity as a function of detuning from the cavity resonance. Total phase shift across the resonance decreases as we increase the internal loss. (d) Amplitude of the reflected signal, which goes to zero at the critical coupling point $\Delta = 0$, $\kappa_{\text{int}} = \kappa_{\text{ext}}$

amplifiers, though they are ubiquitous in circuit QED and well worth committing to memory! Note that some conventions replace b_{in} with $-b_{\text{in}}$.

Before turning to amplifiers, we first describe an important application that motivates microwave amplification in the first place: qubit readout. Qubit readout is also a simple application of the input-output relations, though some of its aspects are better treated with the formalism of the next chapter. Suppose we have a qubit dispersively coupled to a cavity as in chapter 1. In the rotating frame, the system Hamiltonian is

$$H_{\text{sys}}/\hbar = \chi a^\dagger a \sigma_z. \tag{2.42}$$

This equation applies to a transmon when we restrict it to lie in its lowest two energy levels. One can use this Hamiltonian to perform a *quantum non-demolition* measurement of the qubit state. This term means that the measurement has no other action on the qubit but to project it into an eigenstate of the measurement operator, which in this case will be σ_z. To see how this works, let's write down the input-output relations for the cavity based on

Eq. (2.42)

$$\dot{a} = -i\chi a \sigma_z - \frac{\kappa}{2}a - \sqrt{\kappa}b_{\text{in}}[t]. \tag{2.43}$$

As in most circuit QED applications, we take the $b_{\text{in/out}}$ modes to describe excitations of a 1D transmission line, so that k is a scalar rather than a vector.

Suppose the qubit is in either $|g\rangle$ or $|e\rangle$, and b_{in} is in a coherent state. In this case, the system reduces to an off-resonantly driven cavity. We know from chapter 1 that the cavity should remain in a coherent state at all times.[14] Let us assume that the state of the system at time t is $|\psi\rangle = |g/e, \alpha_{g/e}(t), \epsilon\rangle$, where $\alpha_{g/e}$ is the coherent state amplitudes of a conditioned on the qubit state and ϵ is proportional to the coherent state amplitude of $b_{\text{in}}[t]$. Taking the expectation value of Eq. (2.43) yields

$$\frac{d}{dt}\langle a \rangle = \dot{\alpha}_{g/e} = \pm i\chi\alpha_{g/e} - \frac{\kappa}{2}\alpha_{g/e} - \sqrt{\kappa}\epsilon. \tag{2.44}$$

Note that putting b_{in} in a coherent state has the same effect as adding a term proportional to $i(a^\dagger - a)$ to H_{sys}, which explains how we generate this term in an experiment. We can easily solve for the steady-state cavity displacement by taking $\dot{\alpha}_{g/e} = 0$, and then using this solution to compute the output field

$$\alpha_{g/e} = \frac{\sqrt{\kappa}\epsilon}{\pm i\chi - \kappa/2} \tag{2.45}$$

$$\langle b_{\text{out}} \rangle = \epsilon\left(1 + \frac{\kappa}{\pm i\chi - \kappa/2}\right)$$

$$= \epsilon\frac{\chi^2 - \kappa^2/4 \mp i\chi\kappa}{\chi^2 + \kappa^2/4}.$$

Thus the reflected phase of the output field depends on the qubit state. Fig. 2.9b depicts the input and output fields in the IQ plane. The observed phase shift is a completely classical effect, analogous to an off-resonantly driven mass on a spring. When driven far below its resonance frequency, the mass moves in phase with the drive, whereas it responds π out of phase when driven above resonance. Interference between the input and output fields leads to a full 2π phase shift as the drive tone is swept across the cavity resonance, as plotted in Fig. 2.9d (the thick blue curve only). For qubit readout, the dispersive coupling shifts the cavity resonance conditioned on the qubit state, which generates a qubit-state-dependent phase shift of the drive.

Although b_{out} is not Hermitian, and hence not observable, all one has to do is measure $(b_{\text{out}} - b_{\text{out}}^\dagger)/i$ (or actually $\int_0^t dt'(b_{\text{out}}[t'] - b_{\text{out}}^\dagger[t'])/i$ to get an appreciable signal-to-noise ratio;

---

[14]Strictly speaking, we have not yet shown that this is the case once we add a coupling to an external system, however it turns out to be the case. It is a basic and essential fact of coherent states that they are mapped to coherent states under the action of Hamiltonians that are linear or quadratic ladder operators and contain no *squeezing terms* $a_i^\dagger a_j^\dagger$ or $a_i a_j$.

more on that in chapter 3), which picks out the imaginary component of the above expression. By performing many measurements of $(b_{\text{out}} - b_{\text{out}}^\dagger)/i$, we can infer $\langle (b_{\text{out}} - b_{\text{out}}^\dagger)/i \rangle$, from which we can infer $\langle \sigma_z \rangle$.

If we can reach a signal-to-noise ratio much greater than 1 then we can also infer the specific measurement outcome for each shot of the measurement. This more demanding task is called *single-shot readout*, which is important because it lets us observe correlations between simultaneous measurements of different qubits, and the quantum trajectories of a single qubit (much more on that later). We can rarely measure anything noiselessly, particularly a microwave signal composed of a handful of photons. Any noise on the signal can mask the difference between the coherent state signals associated with $|e\rangle$ and $|g\rangle$, which prevents us from being sure of the measurement outcome. We might hope that we can increase ϵ, or else the duration of the measurement, but both have their limitations. Readout duration is limited by the excited state lifetime of the qubit, T_1. The intercavity photon number $|\alpha|^2$ can only go so high before the dispersive approximation on which Eq. (2.45) is based breaks down. Even before this critical photon number is reached, the readout power b_{in} can start driving unwanted transitions in the system[115]. To achieve a signal-to-noise ratio greater than one at our room temperature electronics, we need a low noise, low temperature microwave amplifier.

Exercise 17 *Consider a cavity initially in a coherent state α. If all of the $b_{in}[t]$ modes are in the vacuum, show that the coherent state amplitude decays as $\alpha(t) = \alpha(0)e^{-\kappa t/2}$. If you work in the rotating frame, then $H_{sys} = 0$.*

Exercise 18 *Lossy cavities and critical coupling: Realistic cavities have some internal loss in addition to their coupling to an external control line. We can model this loss with another decay rate κ_{int} and an additional set of input and output modes. Show that when this additional loss channel is present, the reflected coherent state is*

$$\langle b_{out} \rangle = \epsilon \frac{(\kappa_{int} - \kappa)/2 + i\Delta}{(\kappa_{int} + \kappa)/2 + i\Delta} \tag{2.46}$$

where Δ is the detuning between the cavity resonance and the drive. The amplitude and phase of this function are plotted in Fig. 2.9c and d. Note that perfect transmission to the loss channel occurs when $\kappa = \kappa_{int}$. This scenario is called critical coupling, and is a simple example of impedance matching.

Parametric amplifiers Amplifiers necessarily have inputs and outputs, and thus are naturally described by input-output theory. While the most common classical amplifiers have seperate input and output connections, the simplest possible amplifier has only a single port that serves as both its input and output. We call these reflection-mode amplifiers, and they are often based on driven nonlinear cavities. In what follows, we first describe how to generate a *squeezing Hamiltonian* using a weakly nonlinear JC circuit. We then show that our new favorite Hamiltonian ideally amplifies b_{in}.
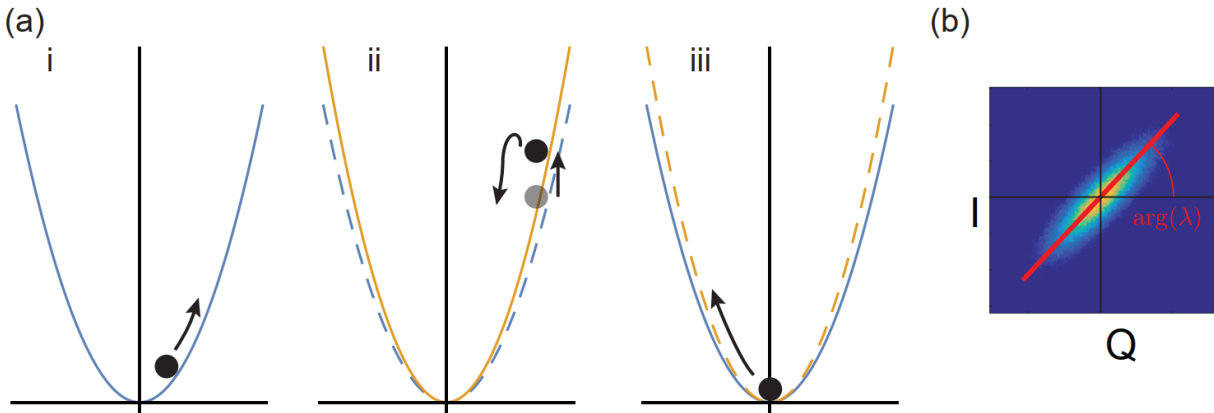
Figure 2.10: (a) Graphical illustration of parametric amplification. Ball gains potential energy at the top of the parabolic potential. (b) A squeezed state, which is the output of a quantum-limited amplifier.

The simplest way to implement a squeezing Hamiltonian technically requires no non-linearity at all, though in practice we rely on the Josephson nonlinearity to generate it. The basic concept is explained in Fig. 2.10a, and can be visualized as a mechanical harmonic oscillator for all intents and purposes. We embue ourselves with the ability to change the spring constant, which amounts to changing the potential energy by a multiplicative factor. Normally the mass oscillates back and forth at a frequency ω in the potential. Suppose that at moment that the mass arrives at the classical turning point, we increase the spring constant slightly, so that the potential rises up as in Fig. 2.10b ii. Although at rest, the system has gained a small amount of energy. Furthermore, we can increase the overall amplitude of oscillation if we return the spring constant to its original value once the mass returns to $x = 0$. If we continuously modulate the potential in this way, which is at twice the resonant frequency ω, the amplitude of oscillation increases exponentially! This simple analogy, not so different from a kid on a swingset (the kind with rigid poles), describes parametric amplification.

Exercise 19 *Using the same logic as above, argue that the motion will be deamplified if the mass oscillates $\pi/2$ out of phase relative to the scenario of Fig. 2.10.*

The quantum version of a parametric amplifier is straight-forward to implement in a JC circuit. A SQUID allows us to periodically vary E_J, which will serve as our time-dependent spring constant. This method of driving a parametric amplifier is called flux pumping. As we saw in the previous section, we can ignore the discreteness of $\hat{n}$ in a JC circuit when $E_J \gg E_C$. We also saw that the non-linearity of the junction becomes negligible when E_J is large enough to confine the wave function to the bottom of the cosine potential. While a small nonlinearity causes problems for a transmon qubit, it lines up nicely with our classical

model of a parametric amplifier. Thus taking $E_J \gg E_C$ beyond the transmon limit lets us approximate

$$H = 4E_C n^2 - E_J(t)\cos\varphi \approx 4E_C n^2 + E_J(t)\frac{\varphi^2}{2} \tag{2.47}$$

Motivated by our classical paramp, let $E_J(t) = E_{J,0} + \epsilon \cos(2\omega t + \delta)$ for some constant small ϵ that we won't worry about too much. A standard application of ladder operators converts the time-independent terms of H to $\hbar\omega a^\dagger a$, leaving us with

$$H/\hbar = \omega a^\dagger a + 2\lambda \cos(2\omega t + \delta)(a + a^\dagger)^2 \tag{2.48}$$

where we have absorbed ϵ and ϕ_{ZPF} into a new constant λ. At this point, the reader may recognize another potential application of the rotating wave approximation, just like our treatment of the driven harmonic oscillator in chapter 1. The main difference is that we now have terms like a^2. Entering a rotating frame applies a phase factor $e^{-i\omega t}$ to each a individually, which transforms the whole Hamiltonian to

$$H/\hbar = \lambda(e^{2i\omega t + i\delta} + e^{-2i\omega t - i\delta})(a^2 e^{-2i\omega t} + a^\dagger a + aa^\dagger + a^{\dagger 2} e^{2i\omega t}) \tag{2.49}$$
$$\approx \lambda^* a^2 + \lambda a^{\dagger 2}$$

where we have absorbed a constant $e^{i\delta}$ into lambda, making it complex.

Eq. (2.49) is known as a squeezing Hamiltonian, because time evolution under H leads to squeezed states like that depicted in Fig. 2.10b. In that figure, we start with the vacuum state $|\psi(t)\rangle = |0\rangle$, evolve for a fixed amount of time, and then plot the Q distribution of the final state. As H contains only even powers of a and $a^\dagger$, photons are only created in pairs. Physically, the nonlinearity of the SQUID is splitting single pump photons of energy $2\hbar\omega$ into two cavity photons of energy $\hbar\omega$.

Although real parametric amplifiers have input and output ports, we can already see strong hints of amplification from Eq. (2.49). Just as the displacement operator that we associated with $a + a^\dagger$ generated all coherent states, we can write down a squeezing operator by exponentiating H

$$S(z) = \exp\left[\frac{1}{2}\left(z^* a^2 - z a^{\dagger 2}\right)\right]. \tag{2.50}$$

In analogy with the displacement operator identity $D^\dagger(\alpha)aD(\alpha) = a + \alpha$, the squeezing operator satisfies

$$S^\dagger(z)aS(z) = a\cosh|z| - e^{i\arg(z)}a^\dagger \sinh|z| \tag{2.51}$$

which is a standard application of the Baker-Campbell-Hausdorff formula Eq. (1.19). Eq. (2.51) is useful because it lets us compute what happens to the the amplitude quadrature operators

after evolution under H

$$S^\dagger(z)I_{\arg(z)}S(z) = e^{-|z|}I_{\arg(z)} \tag{2.52}$$
$$S^\dagger(z)Q_{\arg(z)}S(z) = e^{|z|}Q_{\arg(z)}$$
$$I_\phi \equiv ae^{-i\phi/2} + a^\dagger e^{i\phi/2}$$
$$Q_\phi \equiv (ae^{-i\phi/2} - a^\dagger e^{i\phi/2})/i$$

using $\cosh(r) \pm \sinh(r) = e^{\pm r}$. I_ϕ and Q_ϕ are quadrature operators rotated in phase space by ϕ, which are exponentially amplified and deamplified respectively just like our classical model of Fig 2.10a. Curiously, the squeezed quadrature I_ϕ will have fluctuations less than that of the vacuum state. The Heisenberg uncertainty principle is saved by the fact that the conjugate Q_ϕ has larger fluctuations by the same factor, conserving the product of the uncertainties. Most importantly, Eq. (2.52) clearly shows that our squeezing Hamiltonian acts as a linear amplifier of the Q_ϕ quadrature. The enhanced fluctuations of the output field correspond to amplified vacuum fluctuations.

Exercise 20 *One can also generate a squeezing Hamiltonian by driving a nonlinear cavity resonantly, which is called single pumping. Starting in the rotating frame with $H = \Delta a^\dagger a + \lambda(a^\dagger a)^2$, show that one can generate Eq. (2.49) by applying a strong drive i.e. $a \to a + \alpha$ and setting $\Delta = -4\lambda|\alpha|^2$. You will need to drop third and forth order terms in a, which amounts to assuming that quantum fluctuations are small. What is the physical interpretation for this kind of pumping?*

Now that we recognize our pumped cavity as a linear amplifier, let's add an input/output mode to see how it amplifies an externally applied field[33]. Just as we did with qubit readout, we will first solve for the internal system dynamics using Eq. (2.39), and then substitute the result into Eq. (2.41) to find a relation between the input and output fields. Taking our squeezing Hamiltonian Eq. (2.49) for H_{sys} in Eq. (2.39) gives

$$\dot{a} = -i\lambda a^\dagger - \frac{\kappa}{2}a^\dagger - \sqrt{\kappa}b_{\text{in}} \tag{2.53}$$

We could solve for the steady state a as before, but as the dynamics are more interesting, let us derive the full time-dependent solution this time around. As Eq. (2.53) is linear, we can solve for $a(t)$ by Fourier transforming it

$$\int_{-\infty}^{\infty} \dot{a}(t)e^{i\Delta t}dt = -i\lambda \int_{-\infty}^{\infty} a^\dagger(t)e^{i\Delta t}dt - \frac{\kappa}{2}\tilde{a}(\Delta) - \sqrt{\kappa}\tilde{b}_{\text{in}}(\Delta) \tag{2.54}$$

$$= -i\Delta\tilde{a}(\Delta) = -\lambda\left[\int_{-\infty}^{\infty} a(t)e^{-i\Delta t}dt\right]^\dagger - \frac{\kappa}{2}\tilde{a}(\Delta) - \sqrt{\kappa}\tilde{b}_{\text{in}}(\Delta)$$

$$= -i\lambda\tilde{a}^\dagger(-\Delta) - \frac{\kappa}{2}\tilde{a}(\Delta) - \sqrt{\kappa}\tilde{b}_{\text{in}}(\Delta)$$

The Fourier transform converts the time derivative into multiplication by $-i\Delta$, which can be shown using integration by parts. We have used Δ as our frequency variable to remind ourselves that we work in a rotating frame, so $\Delta = 0$ corresponds to fields oscillating at the paramp resonance frequency ω. The appearance of $\tilde{a}(-\Delta)$ is of fundamental importance; nonlinearity of the paramp symmetrically mixes modes of frequency $\pm\Delta$. If the signal of interest is at a frequency Δ, then $-\Delta$ is called the idler mode.

Using Eq. (2.54) and its complex conjugate evaluated at $-\Delta$, we can solve for $\tilde{a}(\Delta)$ in terms of $b_{\text{in}}(\Delta)$ and $b_{\text{in}}^{\dagger}(-\Delta)$. Some straightforward algebraic manipulations gives

$$\tilde{a}(\Delta) = \sqrt{\kappa} \left(\frac{i\lambda \tilde{b}_{\text{in}}^{\dagger}(-\Delta) - (\kappa/2 - i\Delta)\tilde{b}_{\text{in}}(\Delta)}{(\kappa/2 - i\Delta)^2 - |\lambda|^2} \right). \tag{2.55}$$

With the cavity dynamics in hand, we can compute the input-output relation for a general parametric amplifier

$$\boxed{\begin{aligned} \tilde{b}_{\text{out}}(\Delta) &= -G_S \tilde{b}_{\text{in}}(\Delta) + \underbrace{ie^{i\arg(G_S)}\sqrt{|G_S|^2 - 1}}_{G_I} \tilde{b}_{\text{in}}^{\dagger}(-\Delta) \\ G_S &= \frac{\kappa^2/4 + \Delta^2 + |\lambda|^2}{(\kappa/2 - i\Delta)^2 - |\lambda|^2} \end{aligned}} \tag{2.56}$$

where G_S and G_I are called the signal and idler gains. Note the close similarity between Eq. (2.56) and Eq. (2.51). In particular, G_S and G_I satisfy $|G_S|^2 - |G_I|^2 = 1$ in analogy with the identity $\cosh^2(z) - \sinh^2(z) = 1$. Thus at least for $\Delta = 0$, the relationship between G_S and G_I implies that we can recognize $\tilde{b}_{\text{out}}(0)$ as a squeezed copy of $\tilde{b}_{\text{in}}(0)$ (up to a minus sign, which we recognize as the π phase shift caused by reflection). For $\Delta \neq 0$, $\tilde{b}_{\text{in}}(\Delta) \neq \tilde{b}_{\text{in}}(-\Delta)$, which breaks down the exact correspondence between Eq. (2.56) and Eq. (2.51). In this case, $\tilde{b}_{\text{out}}$ is an example of a *two-mode squeezed state*; vacuum fluctuations from the idler mode $\tilde{b}_{\text{in}}(-\Delta)$ are mixed into the amplified signal, which contributes noise. In the large gain limit $\lambda \to \kappa/2$, $G_S \approx G_I$, indicating that ideally half of the signal consists of amplified idler vacuum fluctuations.

The noise added by an amplifier is one of its most important figures of merit, and as hinted above, there are two important cases to consider. When $\Delta \neq 0$, the input quadratures are both amplified by G_S, but in the limit $G_S \gg 1$, substantial signal from the $-\Delta$ mode is mixed in. Even if the idler mode is in its ground state, vacuum fluctuations pollute the signal, so that the maximum possible signal-to-noise ratio is $1/2$. This mode of operation is called phase-preserving amplification, because the signal is amplified faithfully aside from the added noise. The situation differs when $\Delta = 0$. In this case, the signal and idler modes are the same mode, and Eq. (2.56) matches up with the squeezing relation Eq. (2.51). Just as in Eq. (2.52), there is one now one amplified quadrature (with gain of $2G_S$ instead of G_S) and one squeezed quadrature. This situation, called phase-sensitive gain, does not add any noise.

What determines if an amplifier operates in phase-sensitive or phase-preserving mode? Realistic signals have finite bandwidth, and so we cannot assign an exact frequency to them, but only a center frequency. In reality, the distinction between phase-sensitive and phase-preserving is largely a matter of perspective. If the Fourier component of the signal that is 2Δ away from its center frequency is far out of the bandwidth of interest, then we would call the operation phase-sensitive, and the input idler mode to contribute noise. Conversely, if the signal of interest has time dependence on the time scale of $1/2\Delta$, then we resolve that time dependence as a (potentially rotating) single-mode squeezing of the input field.

One more important property of ideal parametric amplifiers is an intrinsic trade-off between gain and bandwidth, which we now derive. In line with Eq. (2.52), the phase-sensitive gain of an amplifier may be computed as

$$G_{\text{PS}}(\Delta) = |G_S| + |G_I| = \frac{\kappa/2 + \lambda + i\Delta}{\kappa/2 - \lambda - i\Delta} \tag{2.57}$$

$$|G_{\text{PS}}(\Delta)|^2 = \frac{\kappa/2 + \lambda)^2 + \Delta^2}{(\kappa/2 - \lambda)^2 + \Delta}$$

The second line is power gain, which is also often labelled G by convention. In the large gain limit $\lambda \approx \kappa/2$ and for in-band frequencies ($\Delta < \kappa$), we can drop the Δ^2 term in the numerator. This simplification gives

$$|G_{\text{PS}}|^2 \approx \frac{|G_{\text{PS}}(0)|}{1 + \dfrac{\Delta^2}{\Delta_{\text{BW}}^2}} \tag{2.58}$$

$$\Delta_{\text{BW}} = \kappa/2 - \lambda \approx \frac{\kappa}{G_{\text{PS}}(0)}.$$

Eq. (2.58) is a Lorenzian distribution with bandwidth Δ_{BW}. As Δ_{BW} scales inversely with the peak of the gain profile, the gain-bandwidth product is constant $|G_{\text{PS}}|\Delta_{\text{BW}} = \kappa$ is constant as a function of gain.

The appearance of a Lorentzian profile and the gain bandwidth trade-off are essential and ubiquitous predictions of input-output theory and driven resonators in general (note another appearance of a Lorentzian in the first line of Eq. (2.45) when we considered qubit resonance). Both can be circumvented if we violate the Markov approximation made at the beginning of this section, allowing the cavity-transmission line coupling to be frequency dependent. Indeed, parametric amplifiers without a gain-bandwidth trade-off have been made by introducing frequency dependence to the transmission line impedance[111].

Realistic parametric amplifiers Just like qubits, experimental parametric amplifiers do not behave exactly like their idealized theoretical counterparts. Here we treat the nonideealities that matter most for tuning up a real paramp, which are higher-order nonlinearities and internal loss. As we saw in exercise 20, another way to generate a squeezing Hamiltonian is to apply a strong resonant drive to a weakly nonlinear cavity. This mode of operation is called single-pumping. The essence of the single-pumping (and double-pumping, with the

help of the rotating wave approximation) is that under a large displacement $a \rightarrow a + \alpha$, the nonlinear term $(a^\dagger a)^2$ expands out to $\alpha^2 a^{\dagger 2} + \alpha^{*2} a^2 + 4|\alpha|^2 a^\dagger a$ plus higher-order terms that may be neglected at first order.[15] Without these corrections (and the detuning term $a^\dagger a$ that may be cancelled by detuning the pump), the above treatment applies exactly, leading to parametric amplification. Higher-order corrections become non-negligible at high gain or high input power, when the fluctuations of a about its mean α become large. As Eq. (2.56) predicts infinite gain at finite pump powers, high-order corrections inevitably affect realistic devices.

It is easiest to understand nonlinear effects in the case of single pumping, and there are few essential differences between that and flux pumping. We parameterize the nonlinear cavity Hamiltonian as

$$H = \Delta a^\dagger a + \frac{\zeta}{2} a^{\dagger 2} a^2 \tag{2.59}$$

where Δ is the detuning between the pump and the cavity resonance and ζ parameterizes the nonlinearity, which is called the Kerr nonlinearity in this context. Including finite internal loss of the cavity, the resulting input-output relation for the cavity using Eq. (2.39) is

$$\dot{a} = -i\Delta a - i\frac{\zeta}{2} a^\dagger a^2 - \frac{\kappa}{2} a - \sqrt{\kappa_{\text{ext}}} b_{\text{in}} - \sqrt{\kappa_{\text{int}}} b_{\text{in,int}} \tag{2.60}$$

where $\kappa = \kappa_{\text{ext}} + \kappa_{\text{int}}$. To include the effects of loss within the cavity, we have included another mode $b_{\text{in,int}}$, which provides an alternative decay channel. Assuming that this additional mode is in its ground state, the expectation value of the above expression leads to an equation of motion for the coherent displacement of the cavity

$$\dot{\alpha} \approx -i\frac{\zeta}{2} \alpha^* \alpha^2 - \left(\frac{\kappa}{2} + i\Delta \right) \alpha - \epsilon \tag{2.61}$$
$$\epsilon \equiv \sqrt{\kappa_{\text{ext}}} \langle b_{\text{in}} \rangle.$$

We are approximating the intercavity state as a coherent state, so the above represents a semiclassical approximation that neglects quantum fluctuations. A linear response calculation about the classical solution can recover more details than what we work out below, such as gain in the nonlinear regime[155]. We simplify this calculation further by only considering the steady-state behavior, applying the condition $\dot{\alpha} = 0$. Writing $\alpha = e^{-i\phi}|\alpha|$, we have

$$-i\frac{\zeta}{2}|\alpha|^3 - \left(\frac{\kappa}{2} + i\Delta \right) |\alpha| = \epsilon e^{i\phi}. \tag{2.62}$$

As the left-hand side contains only odd powers of $|\alpha|^2$, we can convert it to a third degree polynomial in $|\alpha|^2$ if we multiply both sides by their complex conjugates

$$\frac{\zeta^2}{4}|\alpha|^6 + \zeta\Delta|\alpha|^4 + \left(\frac{\kappa^2}{4} + \Delta^2 \right) |\alpha|^2 = |\epsilon|^2. \tag{2.63}$$
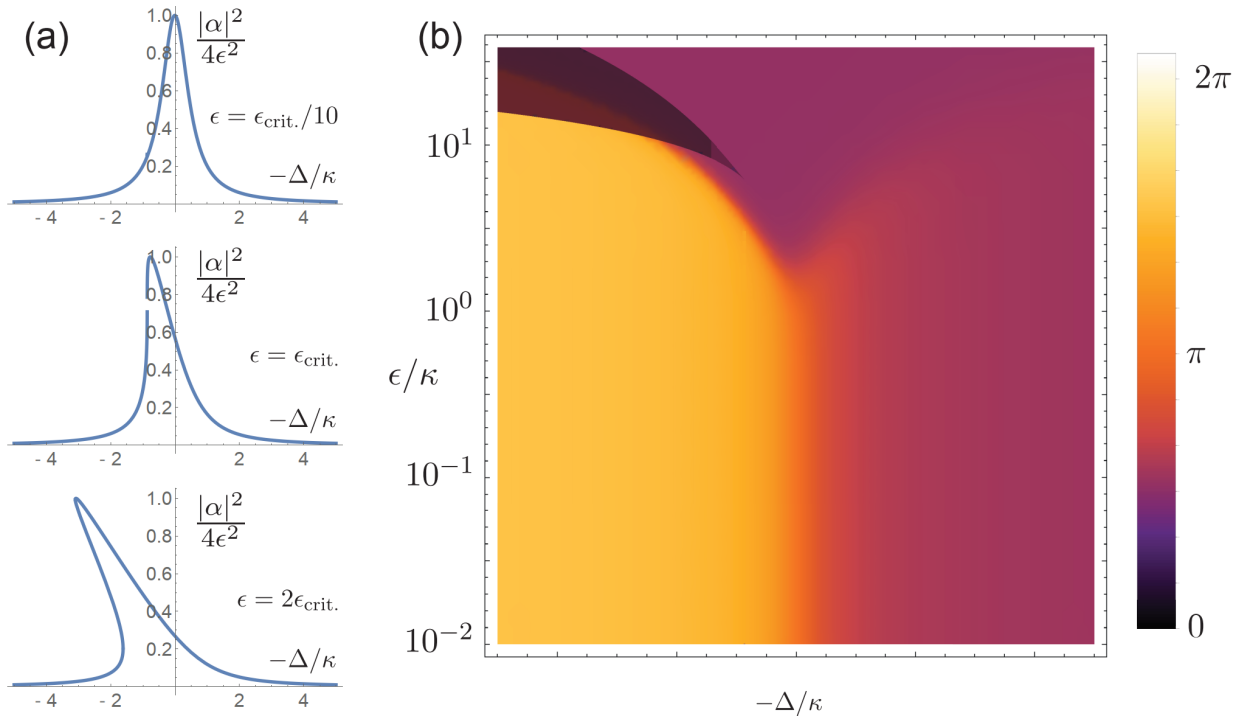
Figure 2.11: Behavior of a parametric amplifier when including the effects of the Kerr non-linearity. (a) Intercavity field amplitude as a function of pump detuning for different drive powers. Multivaluedness of the bottom plot indicates bistability. (b) Phase of the internal state as a function of pump power and frequency. Region in which multiple solutions exist is shown with a darker color.

This equation has three roots, two of which may be complex.

When there is only one real root (such as when $\zeta = 0$, so that we recover the linear cavity result), then we have only one valid solution, as $|\alpha|^2$ cannot be complex. As we increase ζ or ϵ, one finds a regime in which all three roots become real. To understand the physical significance of these additional roots, consider the intercavity power as a function of drive power and frequency, as plotted in Fig. 2.11a. At low powers, the Kerr nonlinearity is negligible and the intercavity field follows a Lorentzian distribution as a function of detuning. With increasing power, the resonance shifts to the left (given negative ζ, as we would have for a cosine potential) and distorts slightly. At some critical power $\epsilon_{\text{crit}} = \sqrt{2\kappa^3/3\sqrt{3}|\zeta|}$, the left portion of the Lorentzian becomes infinitely steep, indicating an instability. Past this power, there are drive frequencies at which all three roots are real. One of these roots corresponds to an unstable state, but the other two are stable, indicating bistability of the device. Thus before we achieve infinite gain, the system bifurcates.

---

[15]Yet another common method is double-pumping, in which two off-resonant tones detuned by $\pm\Delta_{\text{sideband}}$ create the same effective Hamiltonian.

With the full solution for $|\alpha|$ on hand, we can easily compute the complex phase of α using Eq. (2.62). In Fig. 2.11b, we plot the internal cavity phase as a function of drive amplitude and detuning. This plot is useful to measure experimentally when characterizing real devices, and may be generated using only a vector network analyzer. The dark region in the upper-left hand corner shows where bistability occurs. When tuning up a paramp in practice, it is important to work away from the bifurcation regime, and hence away from the point of peak gain. A vertical cut that maximizes gain tends to be preferable, as power sources tend to fluctuation. A dynamical analysis of the above shows that the gain profile becomes distinctly non-Lorentzian above the point of maximum gain[127].

There are other nonidealities that can affect the performance of a realistic device besides the Kerr nonlinearity. Internal cavity loss essentially dropped out of the above analysis, as results only depended on the total loss rate κ. The actual output field may be recovered from the input-output relation Eq. (2.41), which shows that we simply lose some of the internal field. Due to interference effects, phase shifts of the output field as a function of frequency will be less than 2π in the presence of internal loss, as in Fig. 2.9c. Another common nonideality is frequency dependence of κ, which can arise from impedance variations in the output line or external resonant modes. This can lead to different gain bandwidths as the amplifier is tuned in frequency (with a static flux bias), and also to non-Lorentzian behavior even at low powers.

The cavity-based parametric amplifier is one of many ways to amplify microwave signals. A recent development is the travelling wave parametric amplifier (TWPA), in which the nonlinear cavity from above is replaced with a nonlinear transmission line consisting of hundreds or thousands of Josephson junctions[74, 98]. The gain mechanism is the same as above, and intuitively but loosely one can think of a TWPA as a line of cavity-based amplifiers chained one after the other. An important subtlety of the device is that 'sub-amplifiers' must work in concert, so that the same quadrature is amplified at each portion of the device. This requires the pump to remain in phase with the signal as the two propagate down the transmission line. Due to strong nonlinearities, the pump experiences much more dispersion than a far-detuned signal. This dispersion is cancelled out by introducing many resonators along the TWPA near the pump frequency, a technique known as phase matching. Much recent work has also focused on measuring different field operators than the field quadratures, such as photon number[63, 65]. A major contribution of this thesis is the introduction of a phase detector, which measures an observable conjugate to photon number. We describe this result in chapter 6.

Chapter 3

Open Quantum Systems and Measurement Theory

The Hamiltonians of the last chapters are necessarily approximations, as they neglect interactions with the rest of the universe. In classical physics, such interactions present no great difficulty; think of including air resistance and friction when modeling the motion of a projectile. On the other hand, as quantum mechanics is fundamentally a theory of information [92] (hence the mantra 'information is physical'[1]), we might expect dramatically different behavior when an external system 'learns' about the state of our qubit. This behavior is the subject of this chapter, and its applications are the main subject of later chapters.

Undergraduate courses in quantum mechanics almost always omit the study of open quantum systems, which creates the impression that decoherence, weak measurement and quantum trajectories are as mysterious as the measurement problem itself. In reality however, these phenomena follow directly from the basic axioms of quantum theory, and are completely independent of the more philosophical issues regarding the interpretation of wavefunction collapse. The underlying connections are quite simple to state; weak measurement of a degree of freedom (position, spin, photon number *etc.*) results when one interacts that degree of freedom with another system, and then applies a standard projective measurement to that secondary system. Decoherence is just weak measurement by the environment, or to be more precise, by a system that is inaccessible to the experimenter. A quantum trajectory of a state is the state of that system as it undergoes a continuous series of infinitesimally short weak measurements. Importantly, all of these phenomena involve exchange of information.

As long as we know how to apply the postulates of quantum mechanics, open quantum systems should present no great difficulty. However it is important to have precise statements of these postulates. Another topic that is often left somewhat nebulous is the question of *when* one should apply the measurement postulate. What constitutes a measurement? We take the following operational answer:

---

[1]It is tempting to elaborate on this notion further, but I couldn't do this phrase justice better than Scott Aaronson's blog post here https://www.scottaaronson.com/blog/?p=3327.

> *Projective measurement may be said to occur, and hence one may apply the Born rule, once complete information of the system's state has been irreversibly transferred to another system.*

Here, irreversible could be in the thermodynamic sense, or as the result of an experimentalist directing a measurement probe to and then away from their system. We take the notion of complete information transfer to have the following definition, which although rather technical-sounding is nevertheless intuitive. There should exist a complete, orthogonal set of system states that each map the environment to one of a set of its own orthogonal states under the measurement interaction. This condition implies that there exists a hypothetical projective measurement on the environment that can perfectly distinguish the original system state. If furthermore the interaction preserves the system state, then the second part of the Born rule applies, namely that the post-measurement state is an eigenvector of the observable. In essence, we can apply the Born rule if the environment has effectively already applied a projective measurement for us.

Notice that our answer immediately begs the question of when a process becomes irreversible.[2] However as long as the assumption of complete information transfer will hold at all future times, we can apply the measurement postulate whenever we like, so there is no need to ascribe a moment at which measurement occurs. Thus this ambiguity is actually a feature of the definition. It allows us skirt the ultimately philosophical question of when wave function collapse occurs and focus on the operational question of when the Born rule will agree with experimental observation.

The above discussion applies to a rather narrow range of physical processes. The Born rule as we stated it in chapter 1 pertains to projective measurements that leave the system in an eigenstate of the measurement operator. However there are plenty of processes that qualitatively resemble a measurement but satisfy neither of these properties. For example, an excited state of an atomic system that decays and emits a photon offers an obvious way to measure whether the atom is in that excited state or its ground state; wait and see if you observe a photon. This process leaves the system in its ground state regardless of the measurement outcome, in contradiction with the Born rule as typically stated. More subtly, this measurement is technically not even projective. Spontaneous emission follows an exponential decay, so formally one would have to wait an infinite amount of time to be sure that the measurement outcome is 'ground state' if no photon were observed. While projective measurement can at least be recovered in the limit of infinite waiting time, our present definition refuses to acknowledge that information is acquired continuously in this process, and thus that it is a sort of measurement at shorter times too. In the following

---

[2]Imagine reversing time by somehow replacing H with $-H$ in the future light cone of the system. Most interpretations of quantum mechanics predict that such a process would revert the system to its original state, in contradiction with the irreversible change induced by applying the measurement postulate. However we know from statistical mechanics that such a process is extremely unlikely to occur. Regardless, the point is not that this projection has physical reality, which would be a matter of interpretation, but rather that after such an irreversible process, one is free to apply the measurement postulate.

section, we will use the Born rule as a launching point to develop a more general definition of measurement that encompasses most realistic physical processes.

3.1 The POVM Formalism and Weak Measurement

Measurement and unitary evolution are usually thought of as disparate types of dynamics. However given the perspective that measurement is merely a transfer of information from one system to another, one might hope to unify them into a single equation. This is accomplished with the positive operator valued measure (POVM) formalism, which is a somewhat unfortunate choice of name given the simplicity of the underlying concept.

The free evolution of a closed system is described by the time-independent Schrödinger equation

$$\rho(t) = U\rho(0)U^\dagger \tag{3.1}$$
$$U \equiv e^{-iHt}$$

(Throughout this chapter, we work in a system of units that has $\hbar = 1$, a common convention that amounts to measuring energy in units of frequency). Projective measurement of an operator X changes the state according to

$$\rho(t, i) = \frac{|i\rangle\langle i|\rho(0)|i\rangle\langle i|}{\text{Tr}[|i\rangle\langle i|\rho(0)|i\rangle\langle i|]}. \tag{3.2}$$

The operator $|i\rangle\langle i|$ is a projection operator into the ith eigenspace of X. We can unify equations 3.1 and 3.2 into one expression by writing

$$\rho(t, i) = \frac{\Omega_i\rho(0)\Omega_i^\dagger}{\text{Tr}[\Omega_i\rho(0)\Omega_i^\dagger]} \tag{3.3}$$
$$P(i) = \text{Tr}[\Omega_i\rho(0)\Omega_i^\dagger]$$

Unitary evolution: $\Omega_1 = U, \quad \Omega_1^\dagger\Omega_1 = \hat{I}$

Projective measurement: $\Omega_i = |i\rangle\langle i|, \quad \sum_i \Omega_i = \hat{I}$

where the latter condition is necessary to ensure that the probability of all measurement outcomes sums to 1. The Ω_i are known as *Kraus operators*, and they are capable of modeling a much wider range of physical processes than the above two. Note that the condition

$$\boxed{\sum_i \Omega_i^\dagger\Omega_i = I} \tag{3.4}$$

is satisfied for both unitary evolution and projective measurement. We now set out to show that Eq. (3.4) in a sense *defines* a valid operation on a system. The formal version of this result is known as Kraus' theorem.

If we wish to use Eq. (3.3) to describe the widest possible range of dynamics, we need to treat the possibility that the system interacts with an external system. For example, suppose that one introduces an auxiliary system and then applies a joint unitary on the whole. To ultimately write down an equation of motion for the system alone, we need a way to discard the auxiliary system. One way to accomplish this is with a projective measurement on the external system alone, which yields

$$\rho(t,i) \otimes |i\rangle\langle i| = \frac{(I_{\text{sys}} \otimes |i\rangle\langle i|)U(\rho(0) \otimes |\psi\rangle\langle\psi|)U^\dagger(I_{\text{sys}} \otimes |i\rangle\langle i|)}{\text{Tr}[(I_{\text{sys}} \otimes |i\rangle\langle i|)U(\rho(0) \otimes |\psi\rangle\langle\psi|)U^\dagger(I_{\text{sys}} \otimes |i\rangle\langle i|)]}. \tag{3.5}$$

We know the state of a system after measurement occurs, so we can simply discard it by replacing the projection operator $I_{\text{sys}} \otimes |i\rangle\langle i|$ with $I_{\text{sys}} \otimes \langle i|$

$$\rho(t,i) = \frac{(I_{\text{sys}} \otimes \langle i|)U(\rho(0) \otimes |\psi\rangle\langle\psi|)U^\dagger(I_{\text{sys}} \otimes |i\rangle)}{\text{Tr}[(I_{\text{sys}} \otimes \langle i|)U(\rho(0) \otimes |\psi\rangle\langle\psi|)U^\dagger(I_{\text{sys}} \otimes |i\rangle)]} = \frac{\Omega_i \rho(0)\Omega_i^\dagger}{\text{Tr}[\Omega_i \rho(0)\Omega_i^\dagger]} \tag{3.6}$$

$$\Omega_i = (I_{\text{sys}} \otimes \langle i|)U(I_{\text{sys}} \otimes |\psi\rangle).$$

Importantly, Ω_i is an operator that acts on the system alone and satisfies Eq. (3.4). The above Ω_i describes precisely the scenario of a weak measurement. If U maps $|\psi\rangle$ to a different state $|i\rangle$ for each system subspace, then Eq. (3.6) describes a projective measurement on the system via the auxiliary system. We will see a concrete example of such measurements in the next section.

Exercise 21 *Eq. (3.6) does not exhaust all possible processes involving interaction and measurement. Show that the following processes also yield sets of Kraus operators Ω_i that satisfy Eq. (3.4): 1. Replace U in Eq. (3.6) with a global projective measurement on the whole. 2. Include in Eq. (3.6) a unitary U_i on the system that depends on the measurement outcome i. The latter scenario describes quantum feedback.*

We have succeeded in introducing an external system $|\psi\rangle$ to the dynamics, but only if we make sure to measure it once we are done. The key to circumventing this constraint is to realize that if someone else were to measure it for us, we would never be able to tell by looking at the system alone; such a measurement acts as the identity on our system. Thus to develop a mathematical representation, we are free to assume that someone else performed a measurement for convenience. As we do not know the outcome of this measurement, the final state of the system is an average over all possible measurement outcomes, weighted by their probabilities

$$\rho(t) = \sum_i \frac{\Omega_i \rho(0)\Omega_i^\dagger}{\text{Tr}[\Omega_i \rho(0)\Omega_i^\dagger]}P(i) = \sum_i \Omega_i \rho(0)\Omega_i^\dagger. \tag{3.7}$$

The above is known as a *Kraus map*, and the corresponding process is called a *quantum operation*. It is notable that by averaging over measurement outcomes, we have restored

linearity in ρ. Non-linearity only arises when we consider the state conditioned on a measurement outcome. Eq. (3.7) is an example of decoherence, and shows that such processes may be modeled with a linear equation just like the Schrödinger equation. Kraus' theorem, and a significant related result called Choi's theorem, state that all operations that map valid density matrices to valid density matrices may be written in the form Eq. (3.7)[93],[3] with one important exception discussed shortly. Thus the POVM formalism allows us to model almost all physical processes, whether they include measurement or not.

Despite Choi's and Kraus' theorems, there is still one type of process that cannot be modeled by the POVM formalism. By the way we have treated it, the transfer of information from the system is implicitly irreversible, as the measurement permanently alters the auxiliary system. Thus although Eq. (3.7) can model an initial interaction with an external system, it would be incorrect to use Eq. (3.7) again to model a subsequent interaction with that same system if in reality no measurement were applied in between. Such a scenario could transfer information back to the system and even reverse the decoherence induced by the first interaction. Such processes are called *non-Markovian*, and arise when a system interacts with a system with which it is already entangled.[4] Conversely, an interaction in which such a back-flow of information is not possible is called Markovian. The latter processes will form the basis for our model of continuous measurement, the master equation as developed in Sec. 3.3.

Exercise 22 *Eq. (3.6) assumed that the introduced auxiliary system was in a pure state. Use Eq. (1.3) ($\rho \equiv \sum_i P_i |\psi_i\rangle\langle\psi_i|$) to generalize to the mixed state case. How many Kraus operators does one need to express this process?*

Exercise 23 *Non-uniqueness of the Kraus representation: Show that two Kraus maps $\sum_i \Omega_i \rho \Omega_i^\dagger$ and $\sum_i \Omega_i' \rho \Omega_i'^\dagger$ are equivalent if there exists a unitary matrix U such that $\Omega_i' = \sum_j U_{ij}\Omega_j$. Note that for every Kraus map, there exists an equivalent set of 'canonical' Kraus operators satisfying $Tr[\Omega_i\Omega_j^\dagger] = \delta_{i,j}$ [93].*

The process of discarding an external system via measurement is ubiquitous in the treatment of open quantum systems. The introduction of a hypothetical measurement is important enough to deserve its own name; the *partial trace* Tr_B on a composite system ρ_{AB} discards subsystem B, leaving us with a density matrix ρ_A

$$\rho_A \equiv \text{Tr}_B[\rho_{AB}] \equiv \sum_i (I_A \otimes \langle i|_B)\rho_{AB}(I_A \otimes |i\rangle_B) \tag{3.8}$$

The same cancellation that occurred in Eq. (3.7) has occurred in Eq. (3.8), again resulting in a linear equation. The partial trace generalizes the standard trace operation from linear

---

[3]Such a valid map is linear and completely positive. The latter condition means that it maps positive semidefinite matrices to positive semidefinite matrices, even if one were to append an auxiliary system to ρ and then apply the map to the composite system. See the discussion of completely positive maps around Eq. (8.43) of [93].

[4]These operations also need not be completely positive.

algebra, and satisfies many corresponding identities. For example, given an operator $\mathcal{O}_B$ that only acts as the identity on subsystem A, we have $\text{Tr}_B[\mathcal{O}_B \rho_{AB}] = \text{Tr}_B[\rho_{AB} \mathcal{O}_B]$. Furthermore, for any operator that may be written as $\mathcal{O}_A \otimes \mathcal{O}_B$, we have $\text{Tr}_B[\mathcal{O}_A \otimes \mathcal{O}_B] = \mathcal{O}_A \text{Tr}[\mathcal{O}_B]$. The latter condition implies that if ρ_{AB} is a separable *i.e.* not entangled, then the partial trace does nothing but undo the tensor product operation $\rho_{AB} = \rho_A \otimes \rho_B$ and return ρ_A. In general if ρ_{AB} is entangled between subsystems A and B, then the purity of ρ_A is lower than that of ρ_{AB}. This illustrates a fundamental fact about entanglement, that an entangled state looks like a classical mixture if we only have access to a subsystem.

3.2 A Toy Model of Weak Measurement

To make the notion of weak measurement concrete, we now turn to our first example. Although the models that we write down are initially motivated by simplicity and pedagogy, the second leads neatly to a model of continuous measurements and can be derived from realistic physical models. By the end of this section, we will have a stochastic master equation and a basic understanding of the stochastic calculus that underpins it. Continuing with the presentation style of the first two chapters, we begin by stating up front the ultimate physical models and their solutions. With intuition for the dynamics in hand, we derive them from a general model of an open quantum system in the following section, and then apply it to relevant examples in circuit QED.

Consider a pair of qubits, one playing the role of the system and the other playing the role of the auxiliary. The projective measurement limit is likely the most familiar case, so we begin with that. Let us initialize the auxiliary qubit in $|g\rangle$ and apply a CNOT gate with the system qubit as the control. The CNOT unitary flips the auxiliary qubit to $|e\rangle$ if the system qubit is in $|e\rangle$, so to read out the state of the system qubit, we should measure the auxiliary in the $\{|g\rangle, |e\rangle\}$ basis. To calculate the Kraus operators associated with this process, we use Eq. (3.6)

$$\Omega_g = (I_{\text{sys}} \otimes \langle g|)U_{\text{CNOT}}(I_{\text{sys}} \otimes |g\rangle) = |g\rangle\langle g| \tag{3.9}$$
$$\Omega_e = (I_{\text{sys}} \otimes \langle e|)U_{\text{CNOT}}(I_{\text{sys}} \otimes |g\rangle) = |e\rangle\langle e|$$
$$U_{\text{CNOT}} = |g\rangle\langle g| \otimes \sigma_0 + |e\rangle\langle e| \otimes \sigma_x$$

As expected, measurement of the auxiliary qubit implements a projective measurement on the system qubit.

To apply a weak measurement, we need a unitary interaction that does less than completely transfer the system state to the auxiliary qubit. A generic Hamiltonian suited for the job resembles the Jaynes-Cummings Hamiltonian

$$H = i\epsilon(c_{\text{sys}}\sigma^\dagger - c_{\text{sys}}^\dagger\sigma). \tag{3.10}$$

In what follows, we treat the qubit degree of freedom as the auxiliary, and let c_{sys} be any system operator, not necessarily just a ladder operator. For our 'toy model' of weak measurement $c_{\text{sys}} = |e\rangle\langle e|$, we obtain $H = -\epsilon|e\rangle\langle e|\sigma_y$, the Hamiltonian for a controlled σ_y rotation.

If we apply this Hamiltonian for a duration $\Delta t = \pi \epsilon$, then we recover the CNOT gate of our first example. If we apply H for a shorter duration, then we recover a weak measurement of $|e\rangle\langle e|$, or in general the operator c_{sys}. Another classic example to bear in mind is $c_{\text{sys}} = \sigma$, which turns out to model a spontaneous emission process (swapping of an excitation from system into an environment).

To see how this works, we compute the unitary generated by H and plug it into Eq. (3.6). In the limit of a very weak interaction/measurement, we can expand U to second order in the interaction time Δt

$$U(\Delta t) = \exp(-iH\Delta t) \approx I + \epsilon\Delta t(c\sigma^\dagger - c^\dagger\sigma) + \frac{(\epsilon\Delta t)^2}{2}(c\sigma^\dagger - c^\dagger\sigma)^2 \qquad (3.11)$$

$$= I + \epsilon\Delta t(c\sigma^\dagger - c^\dagger\sigma) + \frac{(\epsilon\Delta t)^2}{2}(cc^\dagger|e\rangle\langle e| - c^\dagger c|g\rangle\langle g|)$$

where we have dropped the 'sys' subscript to prevent clutter. For the moment, imagine that Δt is small but finite. To compute a weak measurement POVM from Eq. (3.11), we have two choices to make: the initial state of the auxiliary qubit and the auxiliary measurement basis after interaction. For our toy model $c = |e\rangle\langle e|$, if we align the initial state along the σ_y axis then $U(t)$ does nothing and no measurement occurs. Conversely, starting it out in the σ_x, σ_z plane maximizes the amount of information transferred to the environment, regardless of where we put it. Therefore without much loss of generality, we can simplify Eq. (3.11) and the resulting POVM further by starting the auxiliary qubit in $|g\rangle$

$$\Omega_i = (I_{\text{sys}} \otimes \langle i|)U(\Delta t)(I_{\text{sys}} \otimes |g\rangle) \qquad (3.12)$$

$$= \langle i|g\rangle I_{\text{sys}} + \epsilon\Delta t\langle i|e\rangle c - \frac{(\epsilon\Delta t)^2}{2}\langle i|g\rangle c^\dagger c.$$

If we bear in mind the example of spontantous emission, this simplification amounts to starting the environmental degree of freedon in the vacuum. Simple as it looks, the above POVM will carry us through the rest of this chapter, from the toy models below to a stochastic master equation that models countless realistic systems.

Let us consider the dynamics under various measurement bases $\{|i\rangle\}_{i=1,2}$. If we measure in the $|i\rangle = |e/g\rangle$ basis, the two possible measurement outcomes produce qualitatively different results. The least likely measurement outcome $|i\rangle = |e\rangle$ causes a discrete jump in the system, just like a projective measurement

$$\Omega_e = \epsilon\Delta tc. \qquad (3.13)$$

In our toy model, this just collapses the system to $|e\rangle$. As only $|\psi_{\text{sys}}\rangle = |e\rangle$ leads to a rotation of the auxiliary, it makes sense that detecting a rotation of the auxiliary qubit into $|e\rangle$ implies with 100% certainty that the system must be in $|e\rangle$ as well. However the probability for this measurement outcome is quite small. The most likely measurement outcome is $|i\rangle = |g\rangle$, as this aligns with the initial state of the auxiliary qubit. In this case, the system is only

changed infinitesimally, and the Kraus operator is

$$\Omega_g = I - \frac{(\epsilon\Delta t)^2}{2}c^\dagger c \tag{3.14}$$

For $c = |e\rangle\langle e|$, Ω_g pushes the system slightly toward the ground state. Physically, this arises because this 'null' measurement outcome still provides information about the system, hinting that it is in $|g\rangle$.

Eq. (3.13) and Eq. (3.14) generally describe a *quantum jump process*. If c is a ladder operator for the system such that $c + c^\dagger$ gives the system's dipole moment, then they model detection of spontaneous emission with a photon counter, as we will see in the next section. They also provide the mathematically simplest model of an open quantum system in which we do *not* measure the auxiliary, so that we must use the average evolution equation Eq. (3.7). The system evolves continuously under Eq. (3.14), punctuated by random jumps from Eq. (3.13). In some scenarios, such as in the quantum Zeno effect described in the next chapter, the jump probability is negligible and we can propagate under Eq. (3.14) alone. It is convenient to introduce the non-Hermitian

$$H_{\text{eff}} = -i\frac{\gamma}{2}c^\dagger c, \tag{3.15}$$

which generates Ω_g (though with a more physical linear time dependence instead of quadratic). In the limit that the jump probability $P_e = \text{Tr}[\Omega_e\rho(t)\Omega_e^\dagger]$ is small, one can add H_{eff} to a closed system's Hamiltonian to model coupling to an environment.

Things change significantly when we measure in a basis perpendicular to the initial auxiliary qubit state. All measurement axes aligned with the equator of the Bloch sphere can be parameterized as $\sigma_\delta = \sigma_x \cos(\delta) + \sigma_y \sin(\delta)$, like we defined in Eq. (1.13). The eigenvectors of σ_δ can be parameterized as $|i\rangle = (\pm e^{-i\delta},\ 1)^\intercal/\sqrt{2}$, which gives the POVM

$$\Omega_\pm = \frac{1}{\sqrt{2}}\left(I_{\text{sys}} \pm \epsilon\Delta t e^{-i\delta}c - \frac{(\epsilon\Delta t)^2}{2}c^\dagger c\right) \tag{3.16}$$

Unlike quantum jumps, both measurement outcomes change the state only infinitesimally. Once again, we can visualize the action of our POVM by considering $c = |e\rangle\langle e|$. Things will be slightly cleaner if we take $c = |e\rangle\langle e| - I/2 = \sigma_z/2$, which maintains the interaction strength but symmetrizes the action of H on the system. For $\delta = 0$, the measurement axis σ_x lies in the plane of rotation, so our projective measurement is sensitive to the direction of rotation. Due to the presence of $I_{\text{sys}}/\sqrt{2}$ in both Kraus operators, both measurement outcomes occur with roughly equal probability. Depending on the system qubit state, the auxiliary qubit is rotated one way or the other, slightly biasing the measurement outcome probabilities away from 50/50. $\Omega_\pm$ kicks the qubit ever so slightly toward or away from the poles of the Bloch sphere, as we might expect from the partial acquisition of information.

When $\delta = \pi/2$, the measurement axis aligns with the axis of rotation. This situation is peculiar. On one hand, the projective measurement gains no information about how the

auxiliary qubit has been rotated, which does not look like a measurement. On the other hand, we know that some information has been transmitted from system to environment, and therefore something must happen. The resolution of this apparent paradox is that $\Omega_\pm$ is unitary in this case. Specifically, $\Omega_\pm = \exp(\pm i\epsilon\Delta t\sigma_z/2)$ to second order, indicating that the POVM generates a random rotation about the σ_z axis. As we do not learn about the state, the qubit does not get pushed toward a measurement eigenstate, but the measurement still has a random effect on the system.

Continuous measurement, stochastic calculus and the master equation So far, we have worked with a somewhat unnatural model in which we interact the system and environment for a fixed amount of time, stop the interaction, and then measure the environment. As our notational use of ϵ and Δt suggests, the above treatment is poised to provide a continuous model of measurement. It is tempting to simply take the $\Delta t \to 0$ limit, but the resulting theory runs into technical difficulties that will become clear once we move down the right path. A hint of trouble comes already if we try to compute the average evolution using Eq. (3.7) in this limit, which results in $\rho(dt) = \rho(0)$ ($\dot{\rho} = 0$) and hence no decoherence despite the measurement. To get a sensible theory, we need to keep terms to second order in $\epsilon\Delta t$. One ostensibly naive if unexpected way to keep second order terms in the $\Delta t \to 0$ limit is to take

$$\epsilon = \sqrt{\frac{\gamma}{\Delta t}} \tag{3.17}$$

where γ is some kind of rate. Although this seemingly pathological substitution leads to a divergent interaction Hamiltonian, it actually turns out to lead to the correct model of continuous measurement, both abstractly and phenomenologically. Physically, this divergence arises because our system interacts with an infinite number of degrees of freedom. It is essentially the same divergence that leads to the δ function commutator of the $b_{\text{in/out}}$ operators in input-output theory. In fact, we will essentially end up *identifying* the environment operator $\sqrt{\gamma/\Delta t}\sigma$ with b_{in} in the next section.

It turns out that we have already derived a continuous version of the quantum jump process. As the reader may have noticed, Eq. (3.14) leads to Eq. (3.15) when we substitute Eq. (3.17). The more subtle case and much of the interesting physics comes in when we take the continuum limit of Eq. (3.16)

$$\Omega_\pm = \frac{1}{\sqrt{2}}\left(I_{\text{sys}} \pm \sqrt{\gamma\Delta t}e^{-i\delta}c - \frac{\gamma\Delta t}{2}c^\dagger c\right). \tag{3.18}$$

If only the first and last terms were present, we could simply replace Δt with dt and write down a corresponding differential equation (which would be Eq. (3.15)). Unfortunately, the second term would lead to a divergent derivative in our differential equation as we take $dt \to 0$.

If only one of the two possible measurement outcomes ever occurred, then we would indeed be stuck with non differentiable dynamics and no sensible continuum limit. Fortunately the

randomness of the measurement outcome comes to the rescue, and the salvaged theory leads us to stochastic calculus. To see how we arrive at something sensible and finite, first consider the simplest version of such a stochastic problem, which resembles Brownian motion

$$W(t + \Delta t) = W(t) \pm \sqrt{\Delta t} \tag{3.19}$$
$$W(0) = 0$$

where the $+$ and $-$ cases occur with equal probability. We can compute the variance of $W(T)$ using central limit theorem.[5] If the interval T is broken up into $n = T/\Delta t$ time steps, then the variance at time T is $\sum_{i=1}^{n} \Delta t = T$ regardless of the value of Δt! In other words the $\Delta t^{1/2}$ dependence is just what we needed for there to be a well-defined continuum limit. As $\Delta t \to 0$, $W(t)$ is known as a *Wiener process*, and exhibits a striking self-similarity as we zoom in, as depicted in Fig. 3.1. To think of W as the position of a particle undergoing Brownian motion, the $\Delta t \to 0$ limit corresponds to simultaneously reducing the mean free path and increasing the thermal velocity so that the dynamics over a finite time interval preserve the same macroscopic diffusion process.

In the true continuum limit, the Wiener process is conventionally thought of as an integral

$$W(t) = \int_0^t dW \tag{3.20}$$

where dW is a zero-mean Gaussian distributed random variable with variance dt. We will not formally develop stochastic calculus here, but rather state the essential features and show how to use it in practice. For an informal explanation, see section 5B of [**58**], and for a rigorous but manageable derivation, see in [99]. For all intents and purposes, we can loosely take $dW = \pm\sqrt{dt}$. This handwaving equality suggests an odd property that turns out to be rigorously true, which is that

$$dW^2 = dt, \tag{3.21}$$

The above relation is called the Ito rule, and it is a much stricter statement than our assertion that $\langle dW^2 \rangle = dt$. This equality holds because the fluctuations in dW^2 about its mean dt turn out to scale as dt^2[**58**], which equals zero exactly in the $dt \to 0$ limit.

Using the Ito rule, we can derive or at least motivate one of the most useful equations in stochastic calculus, known as Ito's lemma. Suppose we have a twice-differentiable function $f(X)$ that we wish to evaluate $f(X + dX)$ where X is a stochastic variable

$$dX_t = \mu_t dt + \sigma_t dW. \tag{3.22}$$

---

[5]If $\{X_i\}$ is a set of n (well-behaved) independent random variables each with variance σ_i and mean μ_i, then in the limit $n \to \infty$, $\sum_{i=1}^{n}(X_i - \mu_i) \to \left(\sum_{i=1}^{n}\sigma_i^2\right)^{1/2} N(0,1)$ where $N(0,1)$ is a zero-mean Gaussian distribution with unit variance.
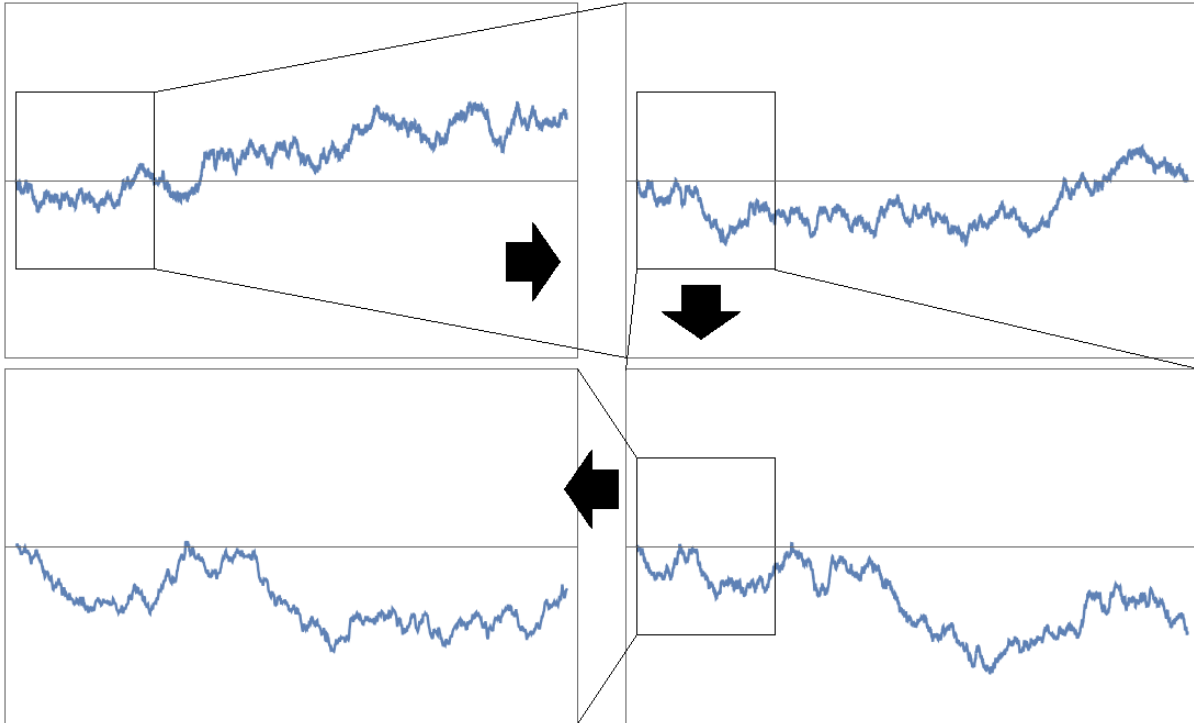
Figure 3.1: Sequential zoom-ins of a Wiener process. Each panel zooms in by a factor of 4 in the horizontal axis and a factor of $\sqrt{4} = 2$ in the vertical axis. Statistically, each panel is indistinguishable from the others.

By Taylor expanding $f(X + dX)$ about X to first order in dt and second order in dW, we might guess that

$$f(X + dX) \equiv f(X) + df \tag{3.23}$$

$$df = f'(X)dX_t + \frac{f''(X)}{2}dX_t^2$$

$$= \sigma_t f'(X)dW + \left(\mu_t f'(X) + \frac{\sigma_t^2}{2} f''(X) \right) dt$$

where we have used the Ito rule and dropped higher-order terms containing $dW\,dt$ or dt^2 in the third equality. Eq. (3.23) is a generalization of the chain rule of ordinary calculus, and we will have many occasions to demonstrate its usefulness in what follows.

As we will have little else to say about the mathematics underlying stochastic calculus, we note that the above formalism is but one of its two forms. As we will continue to see, Ito calculus enables quick-and-dirty calculations in which we multiply differentials as if this is a sensible thing to do (soon we will make use of a generalized product rule $(f + df)(g + dg) = fg + f\,dg + g\,df + df\,dg$). While such manipulations would make any mathematician cringe, they are quite useful in practice, making Ito calculus surprisingly straightforward to work

with. An alternative formulation is Stratonovich calculus, which is also quite useful in physics applications[**37**].

With a rough idea of how to take continuum limits of stochastic processes, let us return to Eq. (3.18) make use of the Wiener process. The potentially most obvious way to apply it is to simply replace $\pm\Delta t$ with dW. Anticipating a differential (or at least difference) equation, we write down the state update as

$$|\tilde{\psi}(t+dt)\rangle = \Omega_{dW}|\psi(t)\rangle = \sqrt{P(dW)}\left(I_{\text{sys}} + \sqrt{\gamma}e^{-i\delta}cdW - \frac{\gamma}{2}c^\dagger cdt\right)|\psi(t)\rangle.$$

(3.24)

Eq. (3.24) is called the linearized stochastic Schrödinger equation, and its solution generates our first example of a quantum trajectory. We have included a tilde over ψ to remind us that it is not normalized, and replaced the $1/\sqrt{2}$ normalization factor as we no longer have a 2-outcome measurement.[6] $|\tilde{\psi}(t+dt)\rangle$ represents a physical state that could be generated by measurement. Unfortunately due to our naive substitution of $\pm\sqrt{\Delta t} = dW$, the stochastic variable does not generate $|\tilde{\psi}(t+dt)\rangle$ with the correct statistics; we should have $P(\tilde{\psi}(t+dt)|\psi(t)) = \langle\tilde{\psi}(t+dt)|\tilde{\psi}(t+dt)\rangle$, but we have already fully specified dW as a zero-mean Gaussian random variable and hence specified the statistics of $|\tilde{\psi}(t+dt)\rangle$. Nevertheless, Eq. (3.24) is quite useful because although the dynamics that it generates do not follow the correct probability distribution implied by the Born rule, it nevertheless generates a *possible* trajectory. Furthermore, it is linear, which enables closed-form solutions in many cases. Finally, we can still use the norm of $|\tilde{\psi}\rangle$ to compute the probability of a given trajectory occurring, as Ω_{dW} is like any other Kraus operator.

To write down an equation of motion with the correct statistics, we seek an Ito random variable with the same statistics as an actual measurement outcome. The probabilities for each of the two possible measurement outcomes from Eq. (3.18) are

$$P(\pm) = \langle\psi(t)|\Omega_\pm^\dagger\Omega_\pm|\psi(t)\rangle \tag{3.25}$$
$$= \frac{1}{2}\left[1 \pm \sqrt{\gamma\Delta t}\langle ce^{-i\delta} + c^\dagger e^{i\delta}\rangle\right].$$

We wish to replace the random variable $\Delta r \equiv \pm\Delta t$ appearing in Eq. (3.18) with an Ito random variable. The above probability distribution implies that

$$\langle\Delta r\rangle = \sqrt{\Delta t}[P(+) - P(-)] = \sqrt{\gamma}\Delta t\langle ce^{-i\delta} + c^\dagger e^{i\delta}\rangle \tag{3.26}$$
$$\text{var}(\Delta r) = \langle\Delta r^2\rangle - \langle\Delta r\rangle^2 = \Delta t + \mathcal{O}(\Delta t^2).$$

---

[6]To justify Eq. (3.24) more carefully, the composition $\Omega_{\pm\pm...\pm} = \prod_{i=1}^N \Omega_\pm = \frac{1}{2^{N/2}}(I + (\sum_i \pm_i\sqrt{\Delta t})M - \frac{1}{2}M^\dagger M\delta t)$, where we have defined $\delta t = N\Delta t$ and dropped cross terms of the form $(\pm_i 1)(\pm_j 1)M^2\Delta t$ as the variance of the sum of these terms limits to zero as $\Delta t \to 0$. We can label the possible Kraus operators as $\Omega_{\delta W}$ with $\delta W = 2n\sqrt{\Delta t}$ and $n = -N/2...N/2$ counts the number of '+'s and '−'s. For each value of n, there are $\binom{N}{N/2+n}$ identical $\Omega_{\pm\pm...\pm}$ that we can group together, so we have $\Omega_{\delta W} = \frac{1}{2^{N/2}}\frac{N!}{(N/2+n)!(N/2-n)!}(I + \delta W M - \frac{N}{2}M^\dagger M\delta t)$. The prefactor limits to a Gaussian as $N \to \infty$, and we recover the main result.

Suppose we sum Δr over many time steps, but keep Δt small enough that $\langle ce^{-i\delta} + c^\dagger e^{i\delta} \rangle$ remains constant over the longer time interval. Central limit theorem implies that the sum limits to a Gaussian distributed random variable with the above mean and variance. Thus in the $\Delta t \to 0$ limit, we can make the replacement

$$\Delta r = \pm \Delta t \tag{3.27}$$
$$\to dr = \sqrt{\gamma} \langle ce^{-i\delta} + c^\dagger e^{i\delta} \rangle dt + dW$$

Substituting dr into Eq. (3.18) yields a stochastic Schrödinger equation with the right statistics

$$|\bar{\psi}(t+dt)\rangle = \left(I_{\text{sys}} + \sqrt{\gamma} e^{-i\delta} c \, dr - \frac{\gamma}{2} c^\dagger c \, dt \right) |\psi(t)\rangle \tag{3.28}$$

where the bar again indicates that ψ is unnormalized but different from $\tilde{\psi}$.

The last step to take is to compute the normalization prefactor. As the right hand side of Eq. (3.28) is infinitesimally closed to being normalized, we can normalize it using Ito's lemma. First, the norm of $\bar{\psi}$ is

$$\langle \bar{\psi}(t+dt)|\bar{\psi}(t+dt)\rangle = 1 + \langle M + M^\dagger \rangle dW + \langle M + M^\dagger \rangle^2 dt \tag{3.29}$$
$$M \equiv \sqrt{\gamma} e^{-i\delta} c$$

where we have defined M for the sake of brevity. We normalize by multiplying by the inverse square root of Eq. (3.29). It is this normalization factor that may be readily computed using Ito's lemma. Let $f(\epsilon) = 1/\sqrt{1+\epsilon}$ and take ϵ to be the infinitesimal part of Eq. (3.29), then the normalization factor is

$$f(\epsilon) = 1 - \frac{\epsilon}{2} + \frac{3\epsilon^2}{8} \tag{3.30}$$
$$= 1 - \frac{1}{2} \langle M + M^\dagger \rangle dW - \frac{1}{8} \langle M + M^\dagger \rangle dt.$$

which is simply a Taylor expansion of f to first order in dt and second order in dW. The normalized state, and hence our stochastic Schrödinger equation is the product $f(\epsilon)|\bar{\psi}(t+dt)\rangle$, which when expanded out becomes

$$|\psi(t+dt)\rangle = |\psi(t)\rangle + (M - \frac{1}{2}\langle M + M^\dagger \rangle)|\psi(t)\rangle \, dW \tag{3.31}$$
$$- \frac{1}{2} \left[M^\dagger M - M\langle M + M^\dagger \rangle + \left(\frac{\langle M + M^\dagger \rangle}{2} \right)^2 \right] |\psi(t)\rangle \, dt$$

The above expression is a bit messy, but it does simplify reasonably well if we take M to be Hermitian, in which case

$$|\psi(t+dt)\rangle = |\psi(t)\rangle + (M - \langle M \rangle)|\psi(t)\rangle \, dW - \frac{1}{2}(M - \langle M \rangle)^2 |\psi(t)\rangle \, dt \tag{3.32}$$

.

The above equations give a reasonably complete model for weak and continuous measurement. The one important aspect of open systems that we have not covered fully is decoherence. We did see how decoherences arises from a discrete weak measurement process in Eq. (3.7), but this model does not realistically capture a qubit coupled to its environment, which is almost always a continuous process. To perform the statistical average over all possible measurement outcomes as we did in Eq. (3.7), we need to work with density matrices. The stochastic master equation associated to Eq. (3.31) above may be computed by expanding $\rho(t + dt) = |\psi(t + dt)\rangle\langle\psi(t + dt)|$ out to first order in dt and second order in dW. After a bit of algebra, the result is

$$d\rho = \mathcal{D}[M]\rho \, dt + \mathcal{H}[M]\rho \, dW \tag{3.33}$$
$$\mathcal{D}[M]\rho \equiv M\rho M^\dagger - \frac{1}{2}(M^\dagger M\rho + \rho M^\dagger M)$$
$$\mathcal{H}[M]\rho \equiv M\rho + \rho M^\dagger - \langle M + M^\dagger\rangle\rho$$
$$dr = \langle M + M^\dagger\rangle dt + dW.$$

We have included the expression for dr to emphasize that its form is intimately tied to the form of the stochastic master equation.

Eq. (3.33) contains a great deal of physical intuition. Note that if the argument of $\mathcal{H}$ is antihermitian then it reduces to a commutator. We saw this feature before in Eq. (3.16) when we took $\delta = \pi/2$, in which case the Kraus operator became unitary. In both cases, we identify types of measurements that look like application of a random Hamiltonian to the system, which is the model of decoherence that we used in chapter 1 when considering Ramsey measurements. As we will show below, when M is Hermitian, Eq. (3.33) describes a QND measurement, and $\mathcal{H}$ tends to push ρ toward an eigenstate of M. In the general case, M can lead to a non-QND process like spontaneous emission, which inexorably pushes ρ toward a certain state or set of states. As the next exercise shows, the second term $\mathcal{D}$ leads to damping of the off-diagonal elements of ρ, a process that we know happens in the long-time limit when Eq. (3.33) should reduce to a projective measurement. For QND measurements, $\mathcal{D}$ does not change the diagonal elements of the density matrix. Thus only the stochastic $\mathcal{H}$ term corresponds to acquisition of information.

Exercise 24 *Let M be Hermitian, in which case it has an eigenvector decomposition $M = \sum_i m_i|i\rangle\langle i|$. Ignoring the $\mathcal{H}$ term of Eq. (3.33), expand the $\mathcal{D}$ term in the eigenbasis of M and show that $\dot{\rho}_{ij} = -(m_i - m_j)^2/2$. Note that this implies that the diagonal elements of ρ are constant on average, the hallmark of a QND measurement.*

There are many conventions for Eq. (3.33) in the literature, but they are all equivalent as long as one modifies $d\rho$ and dr together. Our convention is rather simple because we absorb all unitful quantities into M, so that it has units of $\sqrt{\text{frequency}}$. All other conventions may be obtained by adding a prefactor to M and then pulling it out of $\mathcal{D}$ and $\mathcal{H}$. As

dr labels a measurement outcome, we are also free to multiply it by a non-zero constant without changing the physics. However the ratio between the dt and dW terms is physical, and quantifies an intrinsic signal-to-noise ratio derived from quantum fluctuations of the auxiliary qubit *i.e.* the Poisson fluctuations in measurement outcomes.

To understand decoherence from Eq. (3.33), we can simply average over all possible measurement outcomes dW

$$d\rho = \mathcal{D}[M]\rho dt. \tag{3.34}$$

Note that we recover a linear equation, just like what happened in Eq. (3.7). The more experienced reader may have recognized $\mathcal{D}$ as the Lindblad dissipation term. Indeed, Markovian open systems take the form of Eq. (3.34), though possibly with more terms. It is an incredibly general model of open systems, and appears routinely in both artificial and natural settings.

Exercise 25 *Show that the jump process of Eq. (3.15) also leads to Eq. (3.34) when we average over the jump and no-jump evolution with the correct probabilities.*

From the standpoint of quantum trajectories, Eq. (3.34) also lets us model loss of the signal, such as if a portion of the measurement signal dr is washed out by classical noise or dissipated before detection. We can model these scenarios as a system undergoing two measurements, one of which we detect and the other we don't. Recall that M contains a factor of $\sqrt{\gamma}$. If we break a single measurement of rate γ into two measurements of unequal strength $M_1 = \sqrt{\eta\gamma}c$ and $M_2 = \sqrt{(1-\eta)\gamma}c$, then combining Eq. (3.33) with M_2 and Eq. (3.34) with M_1 gives

$$\begin{aligned} d\rho &= (\mathcal{D}[\sqrt{1-\eta}M]\rho + \mathcal{D}[\sqrt{\eta}M]\rho)dt + \mathcal{H}[\sqrt{\eta}M]\rho\, dW \\ &= \mathcal{D}[M]\rho\, dt + \sqrt{\eta}\mathcal{H}[M]\rho\, dW \end{aligned} \tag{3.35}$$
$$dr = \sqrt{\eta}\langle M + M^\dagger\rangle + dW.$$

η is called the *quantum efficiency*, which is a crucial figure of merit in continuous measurement experiments. When $\eta = 0$, Eq. (3.35) reduces to Eq. (3.34), indicating that a zero-efficiency measurements is just a decoherence process.

Solving the stochastic master equation The stochastic master equations have two primary uses, one experimental and one theoretical. Theoretically, they allow us to predict possible quantum trajectories for a system and compute the average dynamics. The simplest way to compute solutions is with Monte-Carlo simulation, in which one generates dW using a random number generator and then computes $\rho(t+dt) = \rho(t)+d\rho$ iteratively. There are also numerous modifications of this procedure with better numerical stability such as [110] and the analytic solution obtained below. One can also calculate $P(\rho)$ directly by converting the master equation to a *Fokker-Planck equation*, for which there is a standard procedure. The experimental use of the master equations is the computation of a predicted state $\rho(t, r(t))$

from a sequence of measurement outcomes dr. In this case, one computes dW from dr rather than from a random number generator.

In either case, it can be extremely useful to have an analytic solution to the master equation. Eq. (3.24) provides an equivalent linear equation that generates the same trajectories as the full nonlinear stochastic Schrödinger equation, which is harder to solve directly. If M is Hermitian, it is possible to solve in complete generality. Among its many uses, the solution will let us verify that these stochastic differential equations generate projective QND measurements in this case.

Solving the stochastic Schrödinger equation amounts to computing a finite-time Kraus operator from a product of infinitesimal Kraus operators. This computation becomes easier if we identify an exponential that generates Eq. (3.24), as exponentials turn multiplication into addition. By Ito's lemma, one can verify by direct computation that the following exponential is equivalent to Eq. (3.24)

$$\Omega_{dW} = \sqrt{P(dW)} \exp\left(MdW - \frac{1}{2}(M^\dagger M dt + M^2)dt\right) \tag{3.36}$$
$$= \sqrt{P(dW)} \exp(MdW - M^2 dt).$$

In the second line we have used our assumption that $M = M^\dagger$. As $[M, M^2] = 0$, the product of many copies of Eq. (3.36) is easily computed

$$\Omega_r = \left(\frac{1}{\sqrt{2\pi T}} e^{-r^2/2T}\right)^{1/2} \exp(Mr - M^2 T) \tag{3.37}$$
$$= \frac{1}{(2\pi T)^{1/4}} \exp\left(-\frac{(r - 2MT)^2}{4T}\right)$$

where T is the time interval over which we compute our solution, $r = \int_0^T dW$ is the measurement record, and we have computed the prefactor using an argument similar to that of footnote 6 on page 80. Ω_r characterizes the full solution to the master equation in the QND case. The solution demonstrates that the action of measurement over a finite time interval is only determined by the integral of the measurement record, so that r uniquely determines everything.

As we have kept track of normalization, Eq. (3.37) allows us to compute the probability for a given measurement outcome r to occur.[7] Given a state ρ written in the eigenbasis of M as $\rho = \sum_{ij} \rho_{ij} |i\rangle\langle j|$, we find

$$P(r) = \text{Tr}[\Omega_r \rho \Omega_r^\dagger] \tag{3.38}$$
$$= \frac{1}{\sqrt{2\pi T}} \sum_i \rho_{ii} \exp\left(-\frac{(r - 2m_i T)^2}{2T}\right)$$

---

[7]Recall the issue with Eq. (3.24) was that it did not reproduce the correct statistics if we took dW to be a Gaussian random variable with zero mean. Here we are taking dW and its integral r to be unspecified, so that we may determine their probability distributions from Ω_r.
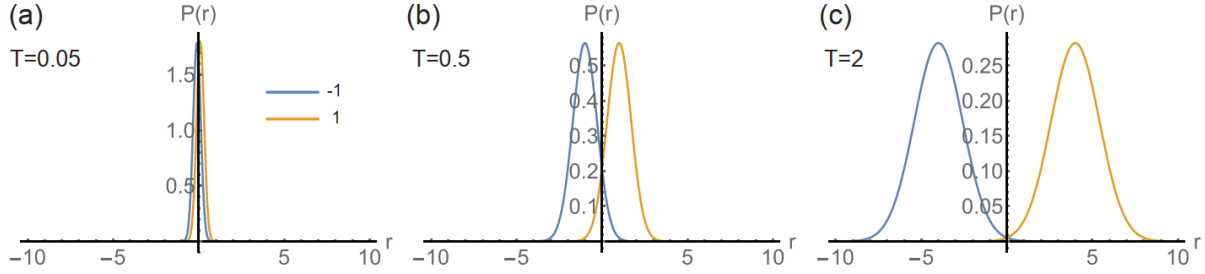
Figure 3.2: Probability distribution of the integrated measurement record dr for three different time intervals. The signal increases faster than the noise as a function of integration interval, so that the distributions become more distinguishable.

where m_i are the eigenvalues of M. Satisfyingly, the probability distributions are a series of Gaussians weighted by the classical probabilities ρ_{ii}, indicating that the measurement outcomes derive from the diagonal elements of the density matrix as expected. Note also that when ρ is an eigenstate of M, ρ and hence $\langle M + M^\dagger \rangle$ are constant and the probability distribution of $\int_0^T dr$ defined in Eq. (3.33) defined previously coincides with the above.

The above solution interpolates nicely between weak and projective measurement. We plot the probabilities $P(r)$ in Fig. 3.2 for $M = \sigma_z$ (so ignoring units for the moment). In the small T limit, the Gaussians defined by Eq. (3.38) overlap, indicating that an inconclusive measurement outcome is likely. Likewise, all diagonal elements of Ω_r are non-zero. As T increases, the separation between the Gaussians scales as T while the standard deviation of each Gaussian only grows as $\sqrt{T}$. In the long time limit, the overlap between the Gaussians becomes negligible, and the measurement outcome is decisive with high probability. In this same limit, when r is likely near $2m_iT$, Ω_r limits to a projection operator, recovering the projective measurement limit.

We conclude with a few general remarks on master equations. Some of the most important properties of Eq. (3.34) are mathematical and explain its generality in describing open quantum systems. Just as the Kraus representation is in a sense the most general possible quantum evolution, Eq. (3.34) is related to the most general form of Markovian continuous quantum evolution. An important generalization of Eq. (3.34) that occurs for example in systems coupled to multiple baths is

$$\dot{\rho} = -i[H, \rho] + \sum_{i,j} \gamma_{i,j} \left(F_i \rho F_j^\dagger - \frac{1}{2} \left[F_j^\dagger F_i \rho + \rho F_j^\dagger F_i \right] \right) \tag{3.39}$$

where F_i form an orthonormal matrix basis ($\mathrm{Tr}[F_i F_j] = \delta_{ij}$) and $\gamma_{i,j}$ is called the Kossakowski matrix[14][43, 68]. If $\gamma_{i,j}$ is positive semidefinite, then Eq. (3.39) generates positive semidefinite state evolution and hence is physically valid. Conversely, all completely positive generators can be put in this form. For finite-dimensional systems, we can diagonalize $\gamma_{i,j}$ to

find

$$\dot{\rho} = -i[H, \rho] + \sum_i \gamma_i \left(A_i \rho A_i^\dagger - \frac{1}{2} \left[A_i^\dagger A_i \rho + \rho A_i^\dagger A_i \right] \right) \tag{3.40}$$

$$= -i[H, \rho] + \sum_i \gamma_i \mathcal{D}[A_i]\rho$$

where there are no particular restrictions on A_i, but $\gamma_i \geq 0$ by positivity. Thus although we derived our master equations for a specific measurement model, they are completely general, and essentially define valid Markovian dynamics. These equations, along with the Kraus representation, have the interesting property that they only make systems less distinguishable, and hence tend to erase information[14, 93].

3.3 Markovian Baths, Continuous Measurement and the Stochastic Master Equations

Now that we have a framework for treating partial measurements, we return to the physical systems considered in the first two chapters in hopes of finding a model of continuous measurement. As an excitation may leave the qubit at one time but reenter it later time, the Jaynes Cummings Hamiltonian models an example of a non-Markovian bath. This would be a bad model of a qubit measurement, as it violates our principle of irreversible information transfer. We can get around this difficulty by adding more oscillators. For a qubit coupled to two harmonic oscillators each with half the strength, it will take twice as long for a single excitation to leave and return. In the limit of an infinite number of oscillators, a lost excitation never returns and the bath becomes memoryless, or Markovian. This system yields a continuous model of measurement that is more physically natural than the discrete operations considered earlier in this chapter.

Despite the suspicious appearance of infinities, this model underlies the entire theory of open quantum systems. It directly and to excellent approximation models countless physical systems, such as a circuit coupled to a transmission line, resistive loss, an atom coupled to the electromagnetic vacuum, a mechanical mode coupled to a bath of phonons in a solid and a nitrogen vacancy center coupled to a bath of defects or magnetic impurities. When we assume that the bath of oscillators is unobserved, we obtain the standard form of the master equation, which is the most general equation of motion for a memoryless decoherence process. If we measure the bath, then we obtain a model for continuous measurement that yields the theory of quantum trajectories and feedback. The resulting theory and the steps used to derive it mirror our previous derivation of input-output theory in many respects.

We begin with a generic system-bath Hamiltonian containing a continuum of harmonic

oscillators

$$H = H_{\text{sys}} + H_{\text{bath}} + H_{\text{int}} \tag{3.41}$$

$$H_{\text{bath}} = \int_0^\infty \frac{d\omega}{2\pi} \omega \, a^\dagger[\omega]a[\omega]$$

$$[a[\omega], a^\dagger[\omega']] = 2\pi\delta(\omega - \omega').$$

As in chapter 2, we use the bracket notation $a[\omega]$ to emphasize that ω indexes an infinite set of operators a. We reserve the notation $\mathcal{O}(\omega)$ to denote the Fourier transform of a single time dependent operator $\mathcal{O}(t) = U(t)^\dagger \mathcal{O} U(t)$ in the Heisenberg picture. Note that as the Hamiltonian has units of ω, $a[\omega]$ has units of $\sqrt{\text{time}}$. We use a more general interaction Hamiltonian that contains the multi-oscillator Jaynes-Cummings Hamiltonian as a special case

$$H_{\text{int}} = i \int_0^\infty \frac{d\omega}{2\pi} \sqrt{\gamma(\omega)} \left(ca^\dagger[\omega] - c^\dagger a[\omega] \right) \tag{3.42}$$

where the coupling strength parameter γ has units of frequency and c could be any system operator. Going into the interaction picture yields

$$H_{\text{int}}^I \equiv H_I(t) = e^{i(H_{\text{sys}}+H_{\text{bath}})t} H_{\text{int}} e^{-i(H_{\text{sys}}+H_{\text{bath}})t} \tag{3.43}$$

$$= i \int_0^\infty \frac{d\omega}{2\pi} \sqrt{\gamma(\omega)} \left(c(t)a^\dagger[\omega]e^{i\omega t} - c^\dagger(t)a[\omega]e^{-i\omega t} \right)$$

where $c(t) \equiv e^{iH_{\text{sys}}t} c \, e^{-iH_{\text{sys}}t}$. Recall that $H_{\text{sys}}^I = H_{\text{bath}}^I = 0$, so Eq. 3.43 fully characterizes the system.

At this point, we have another opportunity to simplify things using the rotating wave approximation. In the eigenbasis of H_{sys}, $[c(t)]_{i,j} = e^{-i\Omega_{i,j}t}c_{i,j}$ where $\Omega_{i,j} = \nu_j - \nu_i$ is the transition frequency from eigenstate $|i\rangle$ to eigenstate $|j\rangle$ and ν_i are eigenvalues of H_{sys} (see Fig. 3.3a). We can divide up the problem into groups of degenerate transitions Ω by defining a set of time-independent operators c_Ω such that

$$c(t) = \sum_\Omega c_\Omega e^{-i\Omega t}. \tag{3.44}$$

For example in the case of a harmonic oscillator, the sum would take $\Omega = \nu$, 2ν, 3ν *etc.* as illustrated in Fig. 3.3a. A nonlinear system like a transmon would have its transitions subdivided further. By the rotating wave approximation, terms with $\Omega < 0$ may be dropped, so that $c(t)$ and c_Ω are triangular matrices.[8] One might worry that the rotating wave approximation falls apart for $\Omega = 0$, which occurs for diagonal elements of $c(t)$ and for matrix elements corresponding to degenerate transitions. Diagonal elements correspond to

---

[8]Whether c is upper or lower triangular depends on whether the we work in a basis in which the diagonal elements of H_{sys} are written in ascending order (as is the convention for a harmonic oscillator) or descending order (as is conventional for a qubit), respectively.
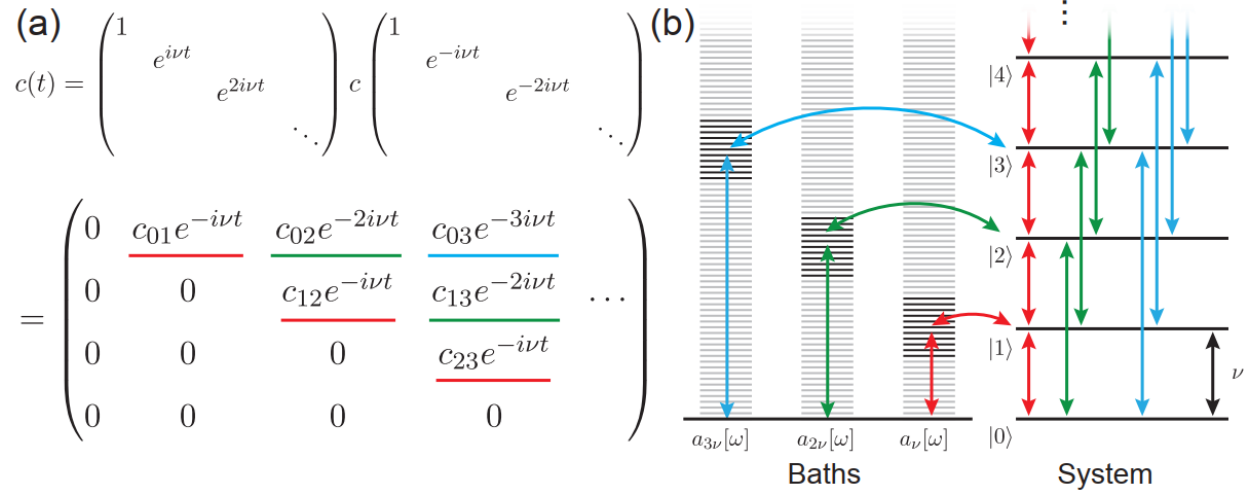
Figure 3.3: (a) Illustration of $c(t)$ for a harmonic oscillator with frequency ν. Terms further from the diagonal drive multi-excitation transitions, cost more energy and hence couple to different components of the bath. For a non-linear system like a transmon, the $0 \leftrightarrow 1$ and $1 \leftrightarrow 2$ transitions could differ in frequency, and hence effectively couple to different baths as well. Elements of $c(t)$ below the diagonal have already been dropped via the rotating wave approximation. (b) Splitting the bath up into multiple independent baths, again illustrated for the harmonic case. We replace one bath coupled to all three transitions with three baths, each of which only couples to one transition. The net result is the same as one bath coupled to all three transitions.

permanent dipole moments, which are rare. Degenerate energy levels are often decoupled from each other and therefore have $c_\Omega = 0$. For example, atomic dipole transition strength typically scale as Ω^3.[9] Regardless, one must be careful if these special cases arise.

If we assume that the values of $\Omega_{i,j}$ are well separated relative to the coupling strengths $\gamma(\Omega)c_\Omega$, then again by the rotating wave approximation, each set of degenerate transitions $i \leftrightarrow j$ couples to a different set of bath operators $a[\omega \approx \Omega_{i,j}]$. Since the bath Hamiltonian is non-interacting, we can completely decouple these groups of degenerate transitions by thinking of the as around each Ω as coming from an independent copy of H_{bath} with operators $a_\Omega[\omega]$.

$$H_I(t) = \sum_\Omega i \int_0^\infty \frac{d\omega}{2\pi} \sqrt{\gamma(\omega)} \left(c_\Omega a_\Omega^\dagger[\omega] e^{i(\omega-\Omega)t} - c_\Omega^\dagger a_\Omega[\omega] e^{-i(\omega-\Omega)t} \right) \tag{3.45}$$

See Fig. 3.3. As in this equation each bath only couples to one set of system transitions, the overall dynamics will agree with those of Eq. (3.43).

At this point, the only approximation we have made is the rotating wave approximation. To make the system tractable, we make a further series of assumptions that are essential to

---

[9]To understand why, look up the Einstein A and B coefficients. There is a very general argument for this scaling based on the thermodynamics of black body radiation.

eliminate memory effects, just like we did in input-output theory. Firstly, we assume that $\sqrt{\gamma(\omega)}$ is constant over the relevant bandwidth (a few $\pm\gamma\Omega$) around each system transition Ω. This is called the quasimonochromatic approximation. Secondly, we add even more harmonic oscillators to our model by extending the ω integral to $-\infty$. By the rotating wave approximation, these additional terms do not participate in the dynamics. However their inclusion helps formalize the notion that the bath is uncorrelated with itself in time, by letting us parameterize $H_I(t)$ in terms of temporal mode operators $b_\Omega[t]$

$$H_I(t) = \sum_\Omega i\sqrt{\gamma(\Omega)} \left(c_\Omega b_\Omega^\dagger[t] - c_\Omega^\dagger b_\Omega[t] \right) \tag{3.46}$$

$$b_\Omega[t] \equiv \int_{-\infty}^\infty \frac{d\omega}{2\pi} a_\Omega[\omega] e^{-i(\omega-\Omega)t}$$

$$[b_\Omega[t], b_\Omega^\dagger[t']] = \delta(t - t').$$

Notice that t plays a funny dual role in the above equation. It indexes the time-independent operators $b_\Omega[t]$, but also specifies the functional time dependence of $H_I(t)$.[10] Physically, this tells us that at each time, the system interacts with a different temporal mode $b[t]$, just like the b_{in} modes from input-output theory. If an excitation swaps into this mode, it never returns, so the system is memoryless. This is the essence of the Markov approximation.

By the commutation relations above, $b[t]$ behaves somewhat like a δ function, so Eq. (3.46) does not make mathematical sense unless it is integrated over a finite time interval. $H_I(t)$ does not commute with itself at different times, so one cannot compute a time evolution operator using the standard relation $U(\Delta t) = \exp(-i \int_0^{\Delta t} H(t) dt)$. Fortunately this expression hold approximately if Δt is short compared to the timescale over which the system changes, $1/\gamma$. On the other hand for the rotating wave approximation to make sense, we should have $\Delta t \gg 1/\Omega$. This establishes a hierarchy of time scales $1/\gamma \gg \Delta t \gg 1/\Omega$. With these constraints in mind, we can write the time evolution operator over a discrete time interval Δt as

$$U(t_n, t_{n+1}) \approx \exp\left(-i \int_{t_n}^{t_{n+1}} H(t) dt\right) = \exp\left(\sum_\Omega \sqrt{\gamma(\Omega)\Delta t} [c_\Omega b_{\Omega,n}^\dagger - c_\Omega^\dagger b_{\Omega,n}]\right) \tag{3.47}$$

$$b_{\Omega,n} \equiv \frac{1}{\sqrt{\Delta t}} \int_{t_n}^{t_{n+1}} b_\Omega[t] dt$$

$$[b_{n,\Omega}, b_{n',\Omega'}] = \delta_{n,n'}, \delta_{\Omega,\Omega'}$$

---

[10]In reality there are two equally valid perspectives regarding time dependence of b. In the view we take in the main text, we change basis and think of $b[t]$ as acting on a single harmonic oscillator indexed by t. In the perspective that $b[t]$ is time-dependent, then it acts on the entire tensor-product Hilbert space of all ∞ harmonic oscillators. The point is that in the latter perspective, we work in the temporal mode basis and $b[t]$ factorizes as $I_0 \otimes I_{dt} \otimes I_{2dt}... \otimes b_t \otimes ....$ where b acts on the single temporal mode labeled by t. These perspectives just boil down to where the implicit identity operators should appear, so the reader is free to take whichever perspective is most intuitive. However in contrast, $a[\omega]$ was actually time-independent, as we pulled out its time dependence $e^{-i\omega t}$ and made it explicit.

where we have normalized the $b_{\Omega,n}$ operators so that they satisfy the standard commutation relations for a harmonic oscillator. The appearance of $\sqrt{\Delta t}$ may look like a mathematical oddity, but it is also the crucial ingredient in modeling continuous measurements as we saw in the previous section.

The reader may have noticed that Eq. (3.47) connects up directly with the previous section deriving the stochastic master equation, specifically Eq. (3.12) at the beginning of the derivation. The main catch is the representation of the environment as a harmonic oscillator $b_{\Omega,n}$ instead of a qubit. As long as the environment harmonic oscillator starts out in its ground state, this distinction vanishes when we consider measurement. Consider expanding out the argument of the exponential in Eq. (3.47). Unlike the qubit case, we find terms like $\Delta t c_\Omega^2 b_{\Omega,n}^{\dagger 2}$ that do not vanish even when we put the $b_{\Omega,n}$ mode in its ground state. Via the Born rule, this leads to an $\mathcal{O}(\Delta t^2)$ probability of detecting two excitations in the $b_{\Omega,n}$ mode, which limits to exactly zero and hence may be neglected. This argument also applies to the cross terms like $c_\Omega c_{\Omega'} b_{\Omega,n}^\dagger b_{\Omega',n}^\dagger$; the probability to observe two excitations in two different modes at the exact same instant is also negligible, and only the leading-order $\sqrt{\Delta t}$ terms contribute to detection. Taking Eq. (3.47) and running it through the derivations of the previous section leads to multiple copies of Eq. (3.33), one for each mode

$$d\rho = \sum_\Omega \gamma(\Omega)\mathcal{D}[c_\Omega]\rho \, dt + \sqrt{\eta\gamma(\Omega)}\mathcal{H}[c_\Omega e^{-i\phi(t)}]\rho \, dW_\Omega \tag{3.48}$$

$$dr_\Omega = \sqrt{\eta\gamma(\Omega)}\langle c_\Omega e^{-i\phi(t)} + c_\Omega^\dagger e^{-i\phi(t)}\rangle + dW_\Omega$$

If we think of Eq. (3.48) as coming from the measurement of a single bath, then dr_Ω is the measurement outcome on that larger bath that has been spectrally filtered around Ω. If the modes are spectrally filtered before performing a measurement, then it is possible to apply a different ϕ_Ω measurement on each part of the bath as well.

The above derivation connects the stochastic master equation to physical measurements on realistic systems. In our qubit model, we derived it from measurements of $\sigma e^{-i\phi} + \sigma^\dagger e^{i\phi}$. Replacing σ with b, we recognize these measurements as quadrature detection, for example of a travelling electromagnetic wave. Thus the detection phase ϕ could be the phase of the pump of a parametric amplifier. An alternative method to perform such an electric field measurement is to interfere the signal with a strong coherent state on a beam splitter, resulting in heterodyne or homodyne measurement[147].

We have avoided a number of important and interesting subtleties in the above derivation that are worth noting.

1. We assumed that the system transitions Ω were either perfectly degenerate or well-separated in frequency. When $|\Omega_i - \Omega_j| \sim \gamma$, interference between decay pathways can lead to a beating effect, which occurs only if the final states are the same[**120**].

2. We have asserted a portion Markov assumption just by writing the interaction Hamiltonian as we did. Normally the coupling term would consist of an integral over the

spatial extent of the system, since each part of it couples to the bath. The first Markov approximation assumes that the system is much smaller than the characteristic interaction time γ multiplied by the time it takes for information to propagate within the bath (so c for an electromagnetic field). This assumption is related to the Lamb Dicke regime from atom trapping and cooling.

3. Many of the assumptions we made to simplify the math, but could in principle be omitted to obtain the same result. For example, interacting baths can still be Markovian. As long as information quickly spreads in the bath, it is unlikely to return to the system and Eq. (3.48) is likely to hold in some form.

4. We can also modify the state of the bath and still recover a Markovian master equation. There are standard master equations for illumination with thermal and even squeezed light[37, 147]. The latter case introduces pair-wise photon correlations, which in a sense reintroduces the $c^{\dagger 2}b^2$ terms that we have dropped.

5. A good deal of physics can be extracted by going slightly beyond the Markov approximation. If the coupling strength $\gamma(\Omega)$ is relatively flat over a small bandwidth and we have the ability to tune the energy levels of our system, then we can map out $\gamma(\omega)$, the excitation spectrum of the bath and even its density of states (which we have assumed to be flat when performing our dk integral with no k-dependent prefactor)[20]. If the density of states goes to zero around one of the system's transition frequencies, then bath modes become spatial exponential decay instead of travelling waves, and we expect the formation of photonic bound states[69].

Exercise 26 *As a simple application of the master equation, let $d\rho = \frac{1}{T_1}\mathcal{D}[\sigma]\rho dt$, which is the master equation for unobserved spontaneous emission from a qubit. Show that ρ_{ee} decays to zero at a rate $1/T_1$ and ρ_{eg} decays at half that rate. This shows that $T_2 \leq 2T_1$, a basic fact of NMR and qubit spectroscopy.*

3.4 Quantum Non-Demolition Measurement and Circuit QED Readout

The previous section gave a general physical derivation of the master equation. However there is one important case that it does not cover. The derivation required the diagonal elements of M to be zero in the system's energy eigenbasis, so that the system-environment coupling comes in at finite frequency. While this constraint is satisfied in many realistic situations, it does not hold for the dispersive Hamiltonian $H = \chi\sigma_z a^{\dagger}a$ and the resulting qubit readout mechanism described in section 2.5 of chapter 2. The reason is that a cavity is not a Markovian bath. It has is a finite correlation time set by the linewidth κ. Equivalently, the density of states is Lorentzian, whereas the master equation derived from a flat density of states.

While a cavity does not act as a Markovian bath for a qubit, the finite linewidth of a cavity does derive from a Markovian bath of some kind, such as the transmission line considered in the context of input-output theory. We can understand this system using the above formalism, which gives us a starting point to understand circuit QED readout. Starting with Eq. (3.35) with $M = ae^{-i\phi}$, the stochastic master equation for a qubit coupled to an open cavity in the rotating wave approximation is

$$d\rho = -i[H, \rho]dt + \kappa\mathcal{D}[a]\rho\ dt + \sqrt{\eta\kappa}\mathcal{H}[ae^{-i\phi}]\rho\ dW \qquad (3.49)$$

$$H = \omega_{\text{Cav}}a^\dagger a + \omega_q\frac{\sigma_z}{2} + g(a^\dagger\sigma + a\sigma^\dagger) + \epsilon^*(t)a$$

$$dr = \langle ae^{-i\phi} + a^\dagger e^{i\phi}\rangle dt + dW.$$

We can transform the above equation into a rotating frame using the standard techniques. As we saw in chapter 1, the Jaynes-Cummings Hamiltonian preserves the qubit state if $\Delta \gg g$, so qubit readout works best in this regime. In this case, we can make the dispersive approximation, which replaces H with $\chi a^\dagger a\sigma_z$. From Eq. (1.32), we saw that the Jaynes-Cummings interaction dresses qubit and cavity excitations. Exercise 11 of that section implies that the ladder operators should therefore be dressed as well, necessitating the replacement $a \to a - [\sigma a^\dagger, a]g/\Delta = a + \sigma g/\Delta$. In the interaction picture, this operator becomes $\tilde{a} = ae^{-i\omega_{\text{Cav}}t} + \sigma e^{-i\omega_q t}\frac{g}{\Delta}$, and Eq. (3.49) becomes

$$d\rho = -i[H_{\text{Disp}}, \rho]dt + \kappa\mathcal{D}[\tilde{a}]\rho\ dt + \sqrt{\eta\kappa}\mathcal{H}[\tilde{a}e^{-i\phi}]\rho\ dW \qquad (3.50)$$

$$\approx -i[H_{\text{Disp}}, \rho]dt + \kappa\mathcal{D}[a]\rho\ dt + \sqrt{\eta\kappa}\mathcal{H}[ae^{-i(\phi+\omega_{\text{Cav}}t)}]\rho\ dW$$

$$+ \kappa\frac{g^2}{\Delta^2}\mathcal{D}[\sigma]\rho\ dt + \sqrt{\eta\kappa\frac{g^2}{\Delta^2}}\mathcal{H}[\sigma e^{-i(\phi+\omega_q t)}]\rho\ dW$$

$$dr = \langle ae^{-i(\phi+\omega_{\text{Cav}}t)} + a^\dagger e^{i(\phi+\omega_{\text{Cav}}t)}\rangle dt + \frac{g}{\Delta}\langle\sigma e^{-i(\phi+\omega_q t)} + \sigma^\dagger e^{i(\phi+\omega_q t)}\rangle dt + dW.$$

$$H = \chi a^\dagger a\sigma_z$$

In going from the first to the second line, we have expanded the $\mathcal{D}$ term and dropped cross terms like $a\rho\sigma^\dagger e^{i(\omega_q-\omega_{\text{Cav}})t}$ via the rotating wave approximation. However we cannot drop fast-rotating terms that are proportional to dW. dW has no correlation time, so it is equivalent to white noise when Fourier transformed. As white noise contains a non-zero spectral component at all frequencies, it leads to a resonant effect even if multiplied by a fast rotating term.

Eq. (3.50) contains a great deal of physics and is worth taking a minute to digest. The second line contains the first-order effects evident from our starting point: cavity decay and measurement-induced back-action arising from observing the emitted field. The third line predicts qubit decay at a rate $\kappa g^2/\Delta^2$, which is known as the *Purcell effect*. There is a corresponding measurement back-action term $\mathcal{H}[\sigma]$ on the qubit, which arises from observing the photons that it emits via the cavity. The last line shows that the signal emitted by the cavity contains two terms, one oscillating at the cavity frequency and one oscillating at the qubit frequency. These terms correspond to cavity and qubit photons respectively.

Purcell decay harms qubit readout and coherence, so we are better off working with a small value of g/Δ. Purcell decay can also be eliminated via a Purcell filter, which decreases the density of electromagnetic states near the qubit frequency[124]. Either way, we work in a regime in which the Purcell decay terms can be neglected. In chapter 6, we will see an example of quantum trajectories derived from spontaneous emission that resembles the above.

In dropping the qubit terms of Eq. (3.50), it seems like we are back to square one and further from our goal of describing qubit readout. To recover qubit readout, we need to solve for the qubit-conditioned dynamics of the cavity. It turns out that we can solve for these dynamics exactly if we neglect Purcell decay[36, 27]. Suppose for a moment that the qubit is in an eigenstate of σ_z. In this case, we know from our example application of input-output theory in chapter 2 that the cavity remains in a coherent state at all times. If we add a drive term into the dispersive Hamiltonian, then the steady-state cavity displacement is given by Eq. (2.45). For notational convenience, it pays to define an operator representation of this coherent displacement

$$\hat{\alpha} \equiv \frac{\sqrt{\kappa}\epsilon}{i\hat{\mathcal{O}} - \kappa/2} \qquad (3.51)$$

$$\hat{\mathcal{O}} \equiv \chi\sigma_z$$

where we have defined $\hat{\mathcal{O}}$ so that our results may be applied to systems larger than a qubit (in the general case one would take $H_{\mathrm{Disp}} = \hat{\mathcal{O}}a^\dagger a$). If the qubit is in a superposition, then the cavity coherent displacement depends on the qubit state, and the general solution written in the eigenbasis of $\hat{\mathcal{O}}$ is

$$|\psi(t)\rangle = \sum_i \psi_i |i\rangle \otimes |\alpha_i\rangle \qquad (3.52)$$

$$\rho(t) = \sum_{ij} \rho_{ij} |i\rangle\langle j| \otimes |\alpha_i\rangle\langle\alpha_j|.$$

This solution assumes that the qubit state is not affected by decay or coherent evolution that does not commute with $\hat{\mathcal{O}}$.

An ideal readout would effectively only depend on the qubit state and not the cavity. Now that we know the cavity dynamics, we can eliminate them with a partial trace. Starting with Eq. (3.50) in the $g/\Delta \to 0$ limit, we can immediately eliminate the $\mathcal{D}[a]$ term

$$\mathrm{Tr}_{\mathrm{Cav}}[\mathcal{D}[a]\rho] = \sum_{ij} \rho_{ij} |i\rangle\langle j| \mathrm{Tr}[\mathcal{D}[a]|\alpha_i\rangle\langle\alpha_j|] = 0. \qquad (3.53)$$

That the trace is zero follows from the fact that $\mathcal{D}$ preserves the norm of the state and therefore must be traceless. The dispersive interaction term is non-zero but simplifies nicely

using $\hat{\alpha}$

$$\mathrm{Tr}_{\mathrm{Cav}}[[a^\dagger a \hat{\mathcal{O}}, \rho]] \tag{3.54}$$

$$= \sum_{ij} \rho_{ij} \mathrm{Tr}_{\mathrm{Cav}}[\hat{\mathcal{O}} a^\dagger a, |i\rangle\langle j| \otimes |\alpha_i\rangle\langle\alpha_j|]$$

$$= \sum_{ij} (\hat{\mathcal{O}}_{ii} - \hat{\mathcal{O}}_{jj}) \alpha_j^* \alpha_i \langle\alpha_j|\alpha_i\rangle \rho_{ij} |i\rangle\langle j|$$

$$= \hat{\alpha}[\hat{\mathcal{O}}, \rho^Q]\hat{\alpha}^\dagger$$

$$\rho^Q \equiv \mathrm{Tr}_{\mathrm{Cav}}[\rho] = \sum_{ij} \rho_{ij} \langle\alpha_j|\alpha_i\rangle |i\rangle\langle j|.$$

We have used the fact that the partial trace allows cyclic permutation on operators that act only on the degrees of freedom that we trace out. Finally, the $\mathcal{H}[ae^{-i(\phi+\omega_{\mathrm{Cav}}t)}]$ term simplifies to a familiar form. If we replace ϕ with $\phi - \omega_{\mathrm{Cav}}t$, we can cancel the explicit time dependence, which corresponds to measuring a fixed quadrature of the field in the rotating frame. Tracing out the cavity then gives

$$\mathrm{Tr}_{\mathrm{Cav}}[\mathcal{H}[ae^{-i\phi}]\rho] \tag{3.55}$$

$$= \sum_{ij} \rho_{ij}|i\rangle\langle j| \, \mathrm{Tr}[ae^{-i\phi}|\alpha_i\rangle\langle\alpha_j| + |\alpha_i\rangle\langle\alpha_j|a^\dagger e^{i\phi} - \langle ae^{-i\phi} + a^\dagger e^{i\phi}\rangle|\alpha_i\rangle\langle\alpha_j|]$$

$$= \sum_{ij} \rho_{ij}|i\rangle\langle j|\langle\alpha_j|\alpha_i\rangle \left[\alpha_i e^{-i\phi} + \alpha_j^* e^{i\phi} - \left(\sum_k \rho_{kk}(\alpha_k e^{-i\phi} + \alpha_k^* e^{i\phi})\right)\right]$$

$$= \mathcal{H}[\hat{\alpha}e^{-i\phi}]\rho^Q$$

Thus the stochastic master equation for the system alone is

$$d\rho^Q = -i\hat{\alpha}[\hat{\mathcal{O}}, \rho^Q]\hat{\alpha}^\dagger \, dt + \sqrt{\eta\kappa}\mathcal{H}[\hat{\alpha}e^{-i\phi}]\rho^Q \, dW. \tag{3.56}$$

We did not derive the above equation by tracing out a Markovian bath, so in general the commutator term cannot be cast into the Lindblad form $\mathcal{D}[...]\rho$[27]. Fortunately the qubit case is especially simple. The diagonal elements of ρ are not changed by the dt term, and the off-diagonal terms are complex conjugates of one another since ρ must be Hermitian. By computing this off-diagonal matrix element, we can match terms and derive an equivalent Lindblad form. The $\mathcal{H}$ term also simplifies nicely in the case of a qubit, leaving us with

$$d\rho^Q = -i\frac{\Delta_{AC}}{2}[\sigma_z, \rho^Q] + \frac{\Gamma_D}{2}\mathcal{D}[\sigma_z]\rho^Q \, dt + \sqrt{\eta\kappa}\mathcal{H}\left[\frac{\alpha_e - \alpha_g}{2}e^{-i\phi}\sigma_z\right]\rho^Q \, dW \tag{3.57}$$

$$\Gamma_D = 2\chi\mathrm{Im}[\alpha_g\alpha_e^*]$$

$$\Delta_{AC} = 2\chi\mathrm{Re}[\alpha_g\alpha_e^*].$$

Δ_{AC} is an AC Stark shift resulting from the net coherent displacement of the cavity, and Γ_D is the dephasing rate induced by measurement.
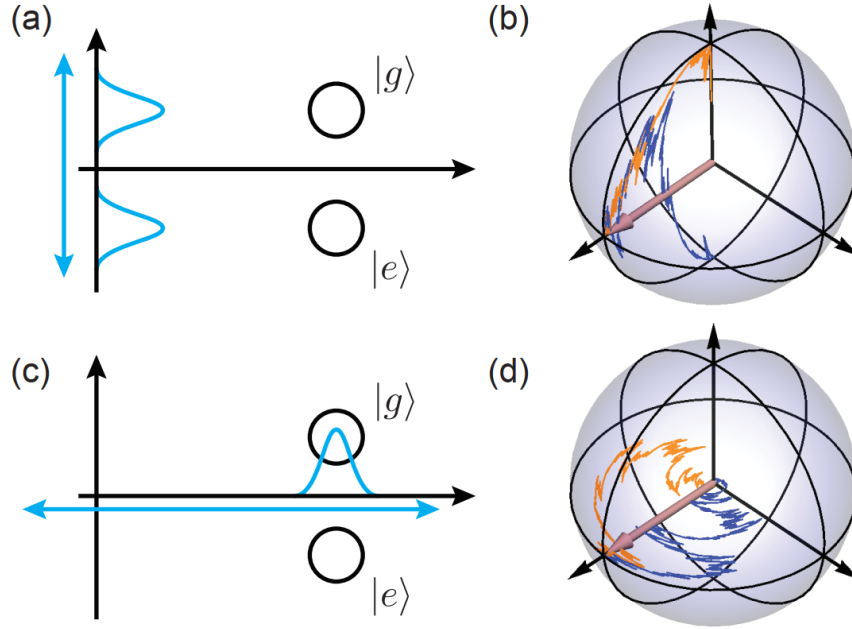
Figure 3.4: Simulated qubit trajectories via disperisive readout ($\eta = 0.4$). (a) Measurement axis (blue) aligned with the quadrature containing qubit state information. Resulting trajectories (b) move toward the eigenstates of the measurement operator. (c) Measurement axis aligned perpendicular to the information-containing quadrature. Although qubit state information is hidden, measurement outcomes to the right end of the distribution project the cavity into a state with higher photon number. (d) Resulting quantum trajectories. Measurement-induced photon number fluctuations lead to stochastic precession around the measurement axis.

Although our starting point assumed that the α_is had reached their steady-state values, our derivation follows through identically even if they change in time. Eq. (3.56) and Eq. (3.57) are valid generally. If we plug in the steady state values for $\alpha_{g/e}$, we find $\Gamma_D \geq 0$. However during transients, it is entirely possible to have Γ_D become negative, indicating non-Markovian (and non-completely positive) dynamics. This unusual situation can arise when the cavity states α_e and α_g are moving closer together, so that the joint qubit-cavity state Eq. (3.52) goes from being entangled to being separable. Negative Γ_D increases the diagonal elements of ρ^Q, as we would generally expect when information flows from the environment back into a system. This classic example of non-Markovian physics is quite similar to the collapse and revival phenomena observed in the Jaynes-Cummings Hamiltonian[120].

The $\mathcal{H}$ term of Eq. (3.57) has an elegant physical interpretation. Its dependence on the difference $\alpha_e - \alpha_g$ shows that the term scales with the amount of qubit information stored in the cavity. If the argument of $\mathcal{H}$ is purely real, then $\mathcal{H}[\sigma_z]$ pushes the qubit up or down along the σ_z axis as we would expect when we acquire qubit state information. This scenario corresponds to measuring the field quadrature that is aligned with $\alpha_e - \alpha_g$ in the complex

plane, and is mathematically identical to the $\delta = 0$ case of Eq. (3.16). When the argument of $\mathcal{H}$ is purely imaginary, the term reduces to a commutator $[\sigma_z, \rho^Q]dW$ *i.e.* stochastic rotations about the σ_z axis. This is exactly analogous to the $\delta = \pi/2$ case of Eq. (3.16) when the Kraus operator became unitary. The unitary dynamics of these equations now has an elegant physical explanation; we are now measuring orthogonal to the quadrature that contains qubit state information, but the measurement still learns something about the intercavity amplitude.

Chapter 4

Measurement as a Control Resource

Quantum trajectories provide a much more nuanced picture of decoherence and measurement. Essentially all forms of observations may be described by weak measurement when analyzed at sufficiently short time scales. Thus the dynamics considered in chapter 3 apply almost universally. However an important result that we derived is that the long-time limit reproduces standard projective measurements, and thus our improved understanding does not initially lead to many new capabilities. Quantum feedback, the process of acting on a system during the measurement process, changes this picture dramatically.

This chapter describes a series of applications of continuous measurement to quantum control. The first section describes a method for control with measurement alone using the Zeno effect. While the scheme involves applying a coherent drive during the measurement process, the way in which we apply it does not depend on the measurement outcome. The remaining sections describe actual quantum feedback, in which we drive the system in a way that depends on the sequence of measurement outcomes. Of particular significance is the PaQS feedback equations Eqs. (4.23,4.24) and their generalization Eqs. (4.36-4.40). These equations encapsulate and generalize many of the existing measurement-based feedback protocols in the literature and lead to some surprising new results. The final section is also of interest in light of the many developments in remote entanglement generation. The sections of this chapter depend on each other only loosely and may be read in any order.

4.1 Control With Measurement Alone: Multi-Qubit Gates Using the Zeno Effect

The Zeno effect As we have seen, measurement damps the diagonal elements of the density matrix when it is written in the eigenbasis of the measurement operator. This damping can actually prevent coherent transitions between measurement eigenstates. To see this effect, consider the solution to the Rabi model $H = \Omega_R \sigma_y / 2$ of chapter 1, first without any

measurement

$$|\psi(t)\rangle = \cos\left(\frac{\Omega_R t}{2}\right)|0\rangle + \sin\left(\frac{\Omega_R t}{2}\right)|1\rangle \tag{4.1}$$

$$\rho(t) = |\psi(t)\rangle\langle\psi(t)| = \frac{1}{2}\begin{pmatrix} 1 + \cos(\Omega_R t) & \sin(\Omega_R t) \\ \sin(\Omega_R t) & 1 - \cos(\Omega_R t) \end{pmatrix}.$$

At $t = 0$, the system starts out in $|0\rangle$. The off-diagonal elements of ρ grow linearly in time to first order, while the population in $|1\rangle$ only grows quadratically. Thus off-diagonal elements build up first, and only after they are significant does population transfer between $|0\rangle$ and $|1\rangle$ begin. Intuitively, this indicates that $|\psi\rangle$ must pass through a state of large coherence like $|\psi\rangle \propto |0\rangle + |1\rangle$ in order to reach $|1\rangle$. If the off-diagonal elements are suppressed by a strong measurement of σ_z with a rate $\Gamma \gg \Omega_R$, then the system will find itself frozen in the $|0\rangle$ state despite the coherent drive.

Exercise 27 *The Zeno effect may also be understood in the context of instantaneous projective measurements. Suppose one performs N projective measurements of σ_z while driving a qubit for a duration T. Show that the probability to remain in $|0\rangle$ is $P_{success} \approx \left(1 - \left(\frac{\Omega_R T}{2N}\right)^2\right)^N$. Show that this limits to one as $N \to \infty$, implying again that the state becomes pinned where it starts.*

The Zeno gate In this section, we describe a novel application of the quantum Zeno effect, aimed at demonstrating its utility for control. In this work and throughout a large fraction of this chapter, we will rely on *degenerate measurements*, which cannot distinguish between certain subspaces, and hence preserve coherence within them. In the context of the Zeno effect, a degenerate measurement allows the system to evolve within a subspace. The state is only frozen in the sense that it is trapped to lie within this subspace, which allows for much richer behavior than in standard Zeno physics. Curiously, the Zeno-restricted dynamics can show new behavior that did not appear under the unmeasured Hamiltonian. This modification of the dynamics represents the concept of Zeno dynamics, which is increasingly considered for quantum control and Hamiltonian engineering in quantum technology[117, 106]. In particular, Zeno dynamics can theoretically transform a trivial quantum system (such as an array of qubits evolving under an easily solved Hamiltonian) into a universal quantum computer[15]. This result, however relies on arbitrary pulse-shaping, which makers it more similar to analog computation and is less amenable to error correction and standard compilation methods. In what follows, we instantiate such ideas to a more explicit model of gate-based quantum computation. We will refer to these gates as Zeno gates, which act as two- or multi-qubit Cphase gates. The gate relies only on the ability to drive a single qubit and assumes no qubit-qubit coupling. Strong measurement of the correct subspace turns this trivial, non-interacting system into one with an entangling Hamiltonian. The required measurement can be implemented in circuit QED systems using standard dispersive measurement.

The Zeno gate acting on three qubits is locally equivalent to a Toffoli gate and hence universal for quantum computation when combined with single-qubit operations[126]. Due to the probabilistic nature of measurement, the Zeno effect can fail, which reduces gate fidelity as it did in the single qubit case. We show that if the measurement channel implementing the Zeno effect is monitored, the gate can be heralded, which yields a probabilistic gate of higher fidelity. Despite its finite error rate, this heralded implementation of the gate can also be used for efficient universal computation, for instance by growing cluster states and then performing adaptive measurement[94, 9].

We depict the system under consideration in Fig. 4.1a. We consider N 3-level systems such as transmons, with no direct interaction. Our computational sub-space is formed by the lowest two levels of each qutrit. The energy levels are labeled $|g\rangle$, $|e\rangle$ and $|f\rangle$ in order of increasing energy as usual. For simplicity we start with only two qutrits. Our goal will be to implement an entangling gate in the qubit subspace $\{|gg\rangle, |ge\rangle, |eg\rangle, |ee\rangle\}$ by driving $|e\rangle \leftrightarrow |f\rangle$ transitions. The system will always end in the computational subspace, either because we drive full 2π rotations or because the transition is blocked by the Zeno effect.
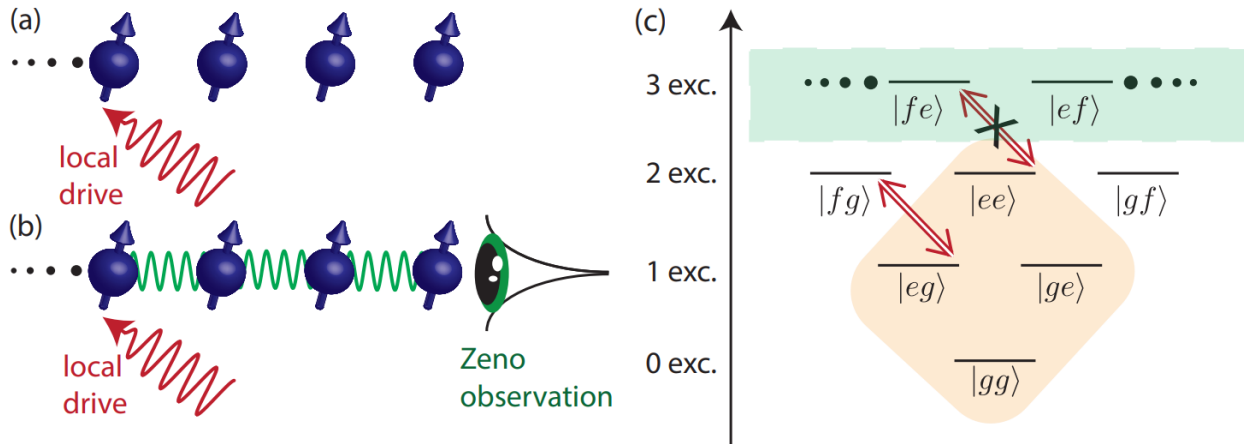
Figure 4.1: (a) Isolated spins with local drives but no spin-spin interaction. (b) Measurement converts local drives into an entangling gate. (c) Energy level diagram for two three-level systems, explaining the basic principle of the Zeno gate. The drive Hamiltonian H induces the transitions drawn in red. Measurement of P blocks the $|fe\rangle$ state, as indicated in green. The combined action of H and P induces a drive on the first qubit conditioned on the second qubit being in $|g\rangle$. Under a full 2π rotation, a geometric phase is imparted on the $|eg\rangle$ level, which constitutes an entangling gate within the computational subspace (orange).

We first consider the ideal scenario when utilizing the Zeno effect, in which one repeatedly applies the following infinite strength projective measurement. The required measurement interrogates whether the system has exactly 3 excitations, and cannot distinguish between
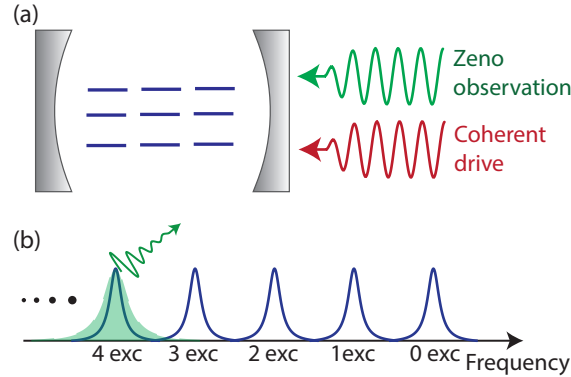
Figure 4.2: Cavity QED implementation of the Zeno gate. (a) Several three-level atoms are coupled dispersively to a high-finesse cavity with leakage rate κ. The cavity is driven by a classical field ω_d near its resonance. (b) Cavity density of states as a function of frequency, plotted for all qubit states in the computational basis. This plot can also be interpreted as the intercavity field as a function of drive frequency. Due to the nature of the dispersive Hamiltonian, cavity resonances cluster in groups of equal qubit excitation number, facilitating measurement of P.

any other subspaces

$$P = 1 - |fe\rangle\langle fe| - |ef\rangle\langle ef|. \tag{4.2}$$

In the limit considered, strong measurement of P prevents the system from entering the three-excitation subspace. As shown in Fig. 4.2, such a measurement can be implemented with a pair of transmon qubits dispersively coupled to a resonator. Driving the cavity resonance associated with the $|fe\rangle$ state implements the above measurement, assuming that it is degenerate with $|ef\rangle$ The gate turns out to be insensitive to the measurement's effect on the $|ef\rangle$, so this degeneracy need not be engineered in practice. During the measurement, we also drive the $|e\rangle \leftrightarrow |f\rangle$ on the first qubit alone

$$H = i\frac{\Omega_R}{2}(|e\rangle\langle f| - |f\rangle\langle e|) \otimes I. \tag{4.3}$$

The above operations, when applied simultaneously yield a unitary map that is locally equivalent to a Cphase gate when applied for a time duration $t = 2\pi/\Omega_R$. Figure 4.1c depicts the basic concept. If the systems starts in the computational subspace (orange diamond), then H alone drives the $|eg\rangle \leftrightarrow |fg\rangle$ and $|ee\rangle \leftrightarrow |fe\rangle$ transitions. However the latter transition is blocked by the Zeno effect under continuous measurement of P, so $|ee\rangle$ is left untouched. Although under a full 2π rotation, the $|eg\rangle$ component of the wave function is also mapped to itself, it picks up a geometric phase of π in the process. This acquired phase is analogous and equivalent to the global phase acquired by a single qubit undergoing a 2π rotation. Thus the net operation applies a π phase shift to the $|eg\rangle$ component of the state

while leaving all other components unchanged. This map is equivalent to a Cphase gate when conjugated by a pair of π pulses on the first qubit before and after the gate.

To see this effect explicitly, consider the action of H projected into the 2-excitation-or-less subspace via P[35]

$$H_{\text{Zeno}} = PHP$$
$$= i\frac{\Omega_R}{2}P(|e\rangle\langle f| - |f\rangle\langle e|) \otimes \underbrace{(|g\rangle\langle g| + |e\rangle\langle e| + |f\rangle\langle f|)}_{I}P$$
$$= i\frac{\Omega_R}{2}(|eg\rangle\langle fg| - |fg\rangle\langle eg|). \tag{4.4}$$

This Hamiltonian drives Rabi oscillations on the $|e\rangle \leftrightarrow |f\rangle$ levels of the first system conditioned on the second system lying in the $|g\rangle$ state, which corresponds to the transition diagram 4.1. As it only acts non-trivially within a 2-dimensional subspace, we can compute $U_{\text{Zeno}} = \exp(-itH_{\text{Zeno}})$ using Euler's formula

$$U_{\text{Zeno}} = (I - \Pi_{eg,fg}) + \Pi_{eg,fg}\cos(\Omega_R t/2) \tag{4.5}$$
$$+ (|eg\rangle\langle fg| - |fg\rangle\langle eg|)\sin(\Omega_R t/2)$$
$$= I - 2\Pi_{eg,fg}$$

where we have defined the operator $\Pi_{eg,fg} \equiv |eg\rangle\langle eg| + |fg\rangle\langle fg|$ to project into the subspace affected by H_{Zeno} and we have chosen $t = t_{\text{final}} = 2\pi/\Omega_R$ for the second equality. Eq. (4.5) is the unitary for a controlled phase gate that applies a π phase conditioned on occupying the $|eg\rangle$ level, as desired. Note that we can also generate gates with any phase between 0 and π. Any geometric phase may be attained by subtending less than a hemisphere of the Bloch sphere when driving the $|e\rangle \leftrightarrow |f\rangle$ transition. This allows one to generate infinitesimal gates, for example to generate Trotterized evolution of a desired Hamiltonian.

U_{Zeno} generalizes straightforwardly to N qubits. The relevant measurement operator and Hamiltonian are now

$$P_N = 1 - |fe...e\rangle\langle fe...e| \tag{4.6}$$
$$H_N = i\frac{\Omega_R}{2}(|e\rangle\langle f| - |f\rangle\langle e|) \otimes I \otimes ... \otimes I.$$

where we have only written the terms of P_N that are relevant to the dynamics. As before, we compute H_{Zeno} by conjugating H_N with P_N. For notational compactness, we write H_{Zeno} by specifying its action only on the relevant subspace. For all $x_i \in \{g, e\}$, we have

$$H_{\text{Zeno},N}|gx_2...x_N\rangle = 0 \tag{4.7}$$
$$H_{\text{Zeno},N}|ex_2...x_N\rangle = \begin{cases} 0 & : x_2...x_N = e...e \\ -i|fx_2...x_N\rangle & \text{otherwise} \end{cases}$$
$$H_{\text{Zeno},N}|fx_2...x_N\rangle = \begin{cases} 0 & : x_2...x_N = e...e \\ i|ex_2...x_N\rangle & \text{otherwise} \end{cases}$$

$H_{\text{Zeno},N}$ acts non-trivially on 2^{N-1} uncoupled qubit subspaces $\{|ex_2...x_N\rangle, |fx_2...x_N\rangle\}$. By the same calculation as Eq. (4.5), every component of the wave function of the form $|ex_2...x_N\rangle$ except $|ee...e\rangle$ picks up a π phase, so that

$$U_{\text{Zeno},N} = \exp(-i2\pi H_{\text{Zeno},N}) = I - 2\Pi_{ex_2...x_N \neq ee...e}. \tag{4.8}$$

Again, any phase between 0 and 2π may be applied by subtending the corresponding solid angle on the Bloch sphere of each subspace.

$U_{\text{Zeno},N}$ is locally equivalent to an N-body Cphase gate, which in turn is locally equivalent to an N-qubit Toffoli gate. To implement a Cphase gate, simply apply another 2π rotation of H_N, this time without measurement. Now all kets of the form $|ex_2...x_N\rangle$ acquire a π phase, this time including $|ee...e\rangle$. All phases imprinted by this second operation cancel with those of $U_{\text{Zeno},N}$ except that on $|ee...e\rangle$, resulting in an N-body Cphase gate. Furthermore, one may imprint an arbitrary phase on $|ee...e\rangle$ by subtending different solid angles on the first and second application of H_N. The full π Cphase gate becomes a Toffoli gate when we conjugate any one qubit with a Hadamard gate before and after application of the Cphase gate.

A significant challenge of measurement-based control is the inevitable introduction of randomness. Although in principle an infinite-strength measurement yields perfect, deterministic control, realistic implementations of the Zeno effect are of finite strength. We now analyze the effects of finite measurement rate. Continuous measurement of an operator P combined with coherent evolution under Hamiltonian H is described by the master equation

$$\frac{d\rho}{dt} = -i[H, \rho] + \Gamma \mathcal{D}[P]\rho \tag{4.9}$$

where Γ is the measurement rate and $\mathcal{D}[P]\rho \equiv P\rho P^\dagger - (P^\dagger P\rho + \rho P^\dagger P)/2$ is the standard Lindblad dissipator term that models coupling to a Markovian bath[37]. Note that in the conventions used, measurement damps off-diagonal elements of ρ at a rate of $\Gamma/2$. Eq. (4.9) averages over all possible measurement outcomes weighted according to their probabilities, so it takes into account the possibility that the system may enter the Zeno subspace that we are attempting to block.

When Γ is large but finite, the primary source of infidelity is population transfer into the $|fe\rangle$ state. When this occurs, the wave function collapses entirely to $|fe\rangle$ due to measurement and the gate fidelity is zero. As the gate acts as a Rabi drive within the 2-dimensional subspace $\{|ee\rangle, |fe\rangle\}$, we can use a qubit model to estimate the probability of failure. Mapping this subspace to a single qubit $\tilde{\rho}$, the master equation analogous to Eq. (4.9) is

$$\frac{d\tilde{\rho}}{dt} = -i[\tilde{H}, \tilde{\rho}] + \Gamma \mathcal{D}[|0\rangle\langle 0|]\tilde{\rho} \tag{4.10}$$

$$\tilde{H} = \Omega_R \frac{\sigma_y}{2}.$$

An exact solution is readily available. If the qubit is initialized in $|1\rangle$ and allowed to undergo a full 2π rotation ($t = 2\pi/\Omega_R$), then the probability to find the qubit in $|0\rangle$ is approximately

$$\tilde{P}_0 \approx \frac{1 - e^{-4\pi\Omega_R/\Gamma}}{2} \tag{4.11}$$

where we have dropped terms of order Ω^2 and terms that are exponentially small in Γ (see reference [79] for details). Eq. (4.11) gives the probability of failure in a single qubit Zeno model. To derive the corresponding probability of failure for the full Zeno gate, we simply multiply by the probability to find the system in $|ee\rangle$, which is $\rho_{ee,ee}$. In this first-order estimate of the gate fidelity, we assume that if the Zeno projection succeeds in preventing a transition to $|ef\rangle$, then the fidelity is 1. Putting these results together yields a first-order expression for the fidelity as a function of the initial state

$$F^{(1)}_{\text{finite } \Gamma} \equiv 1 - \rho_{ee,ee}(0)P_0 \tag{4.12}$$
$$\approx 1 - \rho_{ee,ee}(0)\frac{1 - e^{-4\pi\Omega_R/\Gamma}}{2}$$

which yields 1 in the ideal Zeno limit.

Eq. (4.12) represents the gate fidelity as a function of the initial state. An equally useful figure of merit is gate fidelity averaged over all pure input states, denoted $\bar{F}$. By direct comparison, it can be shown that Eq. (4.12) agrees with the average fidelity if one takes the initial state to be $(|gg\rangle + |ge\rangle + |eg\rangle + |ee\rangle)/2$[79]. The associated density matrix has $\rho_{ee,ee}(0) = 1/4$, which yields

$$\bar{F}^{(1)}_{\text{finite } \Gamma} \approx \frac{1}{8}\left(7 + e^{-4\pi\Omega_R/\Gamma}\right). \tag{4.13}$$

A comparison between the exact Zeno gate fidelity on the above-mentioned initial state (which may also be solved for exactly, though the solution is complicated) and the average fidelity estimated using Eq. (4.12) are plotted in Fig. 4.3a in blue. The agreement is good for large values of Γ/Ω_R. At smaller values, our assumption of unit gate fidelity conditioned on success of the Zeno measurement breaks down. Furthermore, even at large values of Γ/Ω_R, convergence of fidelity to 1 is quite slow.

The above calculation assumes a unit fidelity gate if the measurement outcome indicates successful exclusion from the $|fe\rangle$ state. To go beyond this approximation, we calculate the gate fidelity post-selected on realizations in which the measurement outcome is always so. The corrected fidelity also represents the experimentally accessible scenario of post-selection. The post-selected fidelity converges to one much faster than the unconditional fidelity already calculated.

There exist many equivalent ways to unravel Eq. (4.9), which represent different physical implementations of the measurement[147]. For calculational ease, we assume a quantum jump model, in which the $|fe\rangle$ state decays to a continuum of states that we continuously monitor. If we never register population in these auxiliary states, then the system evolves under the following non-Hermitian Hamiltonian

$$H_{\text{no jump}} = i\frac{\Omega_R}{2}(|e\rangle\langle f| - |f\rangle\langle e|) \otimes I + i\frac{\Gamma}{2}|fe\rangle\langle fe| \tag{4.14}$$

Note that one propagates the post selected state under $H_{\text{no jump}}$ only using the Schrödinger equation, not the master equation Eq. (4.9). One can recover Eq. (4.9) by evolving the state with Eq. (4.14) and randomly applying quantum jumps to $|fe\rangle$ with the right statistics.
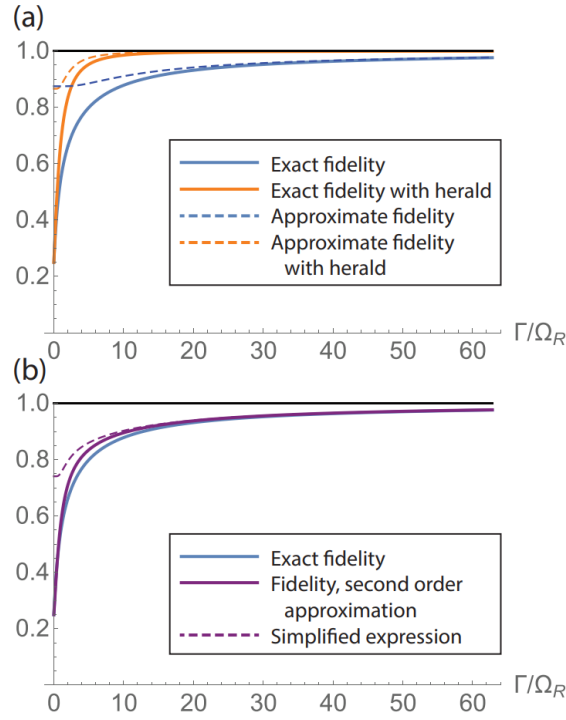
Figure 4.3: Gate fidelity as a function of the unitless ratio Γ/Ω_R, which must diverge to reach the ideal Zeno limit. Fidelity is calculated using the initial state $(|gg\rangle + |ge\rangle + |eg\rangle + |ee\rangle)/2$, which is known approximately or in some cases exactly to coincide with the average fidelity. (a) Fidelity for the unheralded (blue) and heralded (orange) gate. Solid blue and orange lines indicate the exact results obtained by solutions of Eq. (4.9) and Eq. (4.14) respectively, while the blue and orange dashed lines respectively plot the approximate expressions $F_{\text{finite }\Gamma}$ (Eq. (4.12)) and F_{herald} (Eq. (4.16)). (b) Comparison of the exact fidelity (blue line, same as in (a)) with second-order approximations to it. The solid purple line plots $F^{(2)}_{\text{finite }\Gamma}$ using exact solutions inj Eq. (4.17), while the dashed purple line plots Eq. (4.18), in which we have dropped some terms to yield a concise expression.

The above Schödinger equation may be exactly solved. The solution is algebraically almost identical to that of Eq. (4.10), and similar simplifications may be made by dropping terms[79]. The main difference is that we now solve for a wave function instead of a density matrix. Setting $t = 2\pi/\Omega_R$, we find

$$\psi_{gg}(t) = \psi_{gg}(0) \tag{4.15}$$
$$\psi_{ge}(t) = \psi_{ge}(0)$$
$$\psi_{eg}(t) = -\psi_{eg}(0)$$
$$\psi_{ee}(t) \approx e^{-\pi\Omega_R/\Gamma}\psi_{ee}(0)$$

where we have also assumed that $\psi_{fe}(0) = 0$, which yields $\psi_{fe}(t) = 0$. One also finds a non-zero solution for $\psi_{fe}(t)$. However this undesired population can be eliminated by continuing measurement for a time that is large compared to $1/\Gamma$, which exponentially damps ψ_{fe}. Thus we take ψ_{fe} to be zero.

The heralded gate fidelity is simply the overlap squared of the above state with the initial state under application of an ideal gate as per Eq. (4.5). The only subtlety is that the post-selected wave function above must first be normalized, since evolution under a non-Hermitian Hamiltonian does not preserve the norm. The resulting heralded gate fidelity is

$$F_{\text{herald}} \approx \frac{1 - |\psi_{ee}(0)|^2(1 - e^{-\pi\Omega_R/\Gamma})}{\sqrt{1 - |\psi_{ee}(0)|^2(1 - e^{-2\pi\Omega_R/\Gamma})}} \tag{4.16}$$

where we write '$\approx$' because we have used the approximate solutions of Eq. (4.15). Due to post-selection, F_{herald} is non-linear in $|\psi\rangle$, and thus the usual methods of calculating the average fidelity fail. $\bar{F}_{\text{herald}}$ is well approximated by taking $|\psi(0)\rangle = (|gg\rangle + |ge\rangle + |eg\rangle + |ee\rangle)/2$ in Eq. (4.16), just as was the case for the unheralded gate in Eq. (4.12) and Eq. (4.13)[79].

F_{herald} plays a dual role. First interpreting it in the context of a heralded implementation of the gate, we plot the heralded fidelity in orange alongside the unheralded fidelity Fig. 4.3a. The fidelity converges to unity orders of magnitude more quickly. In this setting $\bar{F}_{\text{finite }\Gamma}^{(1)}$ may be interpreted as the success probability, or the fraction of the time in which measurement indicates to jump to $|fe\rangle$.

Secondly, Eq. (4.16) provides a second-order correction to the unheralded fidelity. As F_{herald} explicitly removes the infidelity calculated in $F_{\text{finite }\Gamma}^{(1)}$, we can compute the total fidelity by subtracting from 1 the infidelity associated to each effect individually. The result is

$$F_{\text{finite }\Gamma}^{(2)} \equiv 1 - (1 - F_{\text{finite }\Gamma}^{(1)}) - (1 - F_{\text{herald}}) \tag{4.17}$$

$$\approx \frac{1 - |\psi_{ee}(0)|^2(1 - e^{-\pi\Omega_R/\Gamma})}{\sqrt{1 - |\psi_{ee}(0)|^2(1 - e^{2\pi\Omega_R/\Gamma})}} + \frac{|\psi_{ee}(0)|^2}{2}\left(1 - e^{-4\pi\Omega_R/\Gamma}\right) \tag{4.18}$$

We plot this fidelity against the exact fidelity in Fig. 4.3. We also include the fidelity calculated using the exact solutions to Eq. (4.9) and Eq. (4.14), which are omitted in the

main text for brevity but appear in ref. [79]. The three fidelity curves agree well, indicating that we have quantified the main sources of infidelity in an intuitive, closed-form expression.

The infidelity estimated in $F_{\text{finite }\Gamma}^{(1)}$ has a clear physical interpretation of leakage into the $|fe\rangle$ state. Similarly, the infidelity calculated in F_{herald} has a simple physical explanation. Returning to Eq. (4.15), observe that the only deviation from an ideal gate is damping of the ψ_{ee} component of the wave function. This damping is a result of information acquisition. Only ψ_{ee} population can lead to population in $|fe\rangle$, so if we do not measure $|fe\rangle$ population, then we can infer a lower likelihood to find the system in $|ee\rangle$. This damping is analogous the case of an atom prepared in a superposition of its excited and ground states, which decays to its ground state even if it does not emit a photon.

4.2 General Quantum Feedback Protocols

In this section, we calculate the locally optimal feedback for a measurement-based control system[156]. We leave the measurement and feedback operators completely general. Our optimization technique guarantees that the optimal feedback coefficients are given by ratios of two polynomials that are each linear in ρ. Their computation does not involve computing inverse functions, and so they may be found analytically. A key simplifying assumption is that ρ (or $\bar{\rho}$, in the case of Markovian feedback) stays close to the locally optimal path at each time step. We check the validity of this assumption during numerical propagation of the master equation and rotate the system to the locally optimal path if necessary. With the goal of state preparation in mind, we use fidelity as our cost function throughout this section, though the calculation works with any figure of merit that is linear in the density matrix, such as the expectation value of any observable.

Suppose we have a state ρ at time t, and that we wish to reach some target state $|\psi_T\rangle$. If we take fidelity to be our cost function, then our aim becomes finding the rotation that maximizes the fidelity after a single round measurement and feedback. We take feedback to consist of applying some Hamiltonian H, so that we may parameterize the feedback operation as $U(\theta) = \exp(-iH\theta)$. The fidelity after a single round of measurement and feedback is then given by

$$\mathcal{F}_{t+dt}(\theta) = \langle\psi_T|\rho^c(\theta)|\psi_T\rangle, \tag{4.19}$$
$$\rho^c(\theta) \equiv U(\theta)(\rho + d\rho)U^\dagger(\theta).$$

Our local optimality condition is given by

$$\mathcal{G} \equiv \frac{\partial\mathcal{F}_{t+dt}(\theta_{\text{opt.}})}{\partial\theta_{\text{opt.}}} = \left[U'(\theta_{\text{opt.}})(\rho + d\rho)U^\dagger(\theta_{\text{opt.}}) + h.c.\right] = -i\langle\psi_T|[H, \rho^c(\theta_{\text{opt.}})]|\psi_T\rangle = 0 \tag{4.20}$$

where $\theta_{\text{opt.}}$ is the locally optimal rotation that we wish to calculate. For systems much more complicated than a single qubit, the above equation often lacks a closed-form solution.

However, by definition our protocol has applied the locally optimal rotation at the previous time step. In this case, since $d\rho$ is infinitesimal, then typically $\theta_{\text{opt.}}$ will be infinitesimal as well (we deal with possible exceptions below). Thus we can parameterize it as

$$\theta_{\text{opt.}} = A_1(t)dW + A_2(t)dt. \tag{4.21}$$

Inserting this parameterization into the definition of $U(\theta)$ and expanding to $2nd$ order in dW according to Ito's lemma[99] yields

$$U(\theta) = I - iA_1 H dW - \left(iA_2 H + \frac{1}{2}A_1^2 H^2\right)dt. \tag{4.22}$$

We now substitute $d\rho$ for the measurement stochastic master equation (Eq. (3.35)), which yields

$$
\begin{aligned}
\rho^c = & \rho + \mathcal{D}[M]\rho dt + \sqrt{\eta}\mathcal{H}[M]\rho dW - iA_1[H, \rho]dW + A_1^2 \mathcal{D}[H]\rho dt \\
& - i[H, \sqrt{\eta}A_1(M\rho + \rho M^\dagger) + A_2\rho]dt
\end{aligned}
\tag{4.23}
$$

The above equation is simply the Wiseman-Milburne equation[145] with general coefficients. It serves as the equations of motion for a system undergoing continuous measurement, proportional feedback (via the $A_1[H, \rho]$ term) and evolution under the time-dependent Hamiltonian $A_2 H$. The $\mathcal{D}[H]$ term arises from random, noisy application of H under feedback, which explains why it is present even if the measurement efficiency goes to zero. The last term derives from correlation between the measurement back-action and the feedback operation applied as a result, so it is most readily interpreted as the feedback term.

Eq. (4.23) is quite useful on its own and worth nothing. To find the locally optimal feedback operation, we substitute ρ^c into Eq. 4.20 and solve $\mathcal{G} = 0$ order by order in dW. Despite the multitude of terms, it is straight-forward to solve for A_1 and A_2 in complete generality. The only assumption we make is that the optimal rotation was applied at the immediately preceding time step as mentioned above, so that $\mathcal{F}_t(\theta)$ is maximized at $\theta = 0$ as required for consistency with the assumption that $\theta_{\text{opt.}}$ is infinitesimal. It implies that $\partial\mathcal{F}_t(\theta)/\partial\theta|_{\theta=0} = -i\langle\psi_T|[H, \rho]|\psi_T\rangle = 0$, so that the first term may be dropped. Terms proportional to dW yield a linear equation in A_1 which is easily solved. Once A_1 is known, terms proportional to $dW^2 = dt$ yield another linear equation, this time for A_2. The final result in full form is

$$A_1 = \frac{-i\sqrt{\eta}\langle\psi_T|[H, M\rho + \rho M^\dagger]|\psi_T\rangle}{\langle\psi_T|[H, [H, \rho]]|\psi_T\rangle} \tag{4.24}$$

$$A_2 = \frac{-\langle\psi_T|[H, i\mathcal{D}[M]\rho + \sqrt{\eta}A_1[H, M\rho + \rho M^\dagger] + iA_1^2\mathcal{D}[H]\rho]|\psi_T\rangle}{\langle\psi_T|[H, [H, \rho]]|\psi_T\rangle}$$

A_1 and A_2 are directly related to the proportional feedback gain and constant Hamiltonian amplitudes that would be applied experimentally. Using the expression for the measurement record dr, the locally optimal feedback rotation is given by

$$\theta_{\text{opt.}} = \sqrt{\eta}A_1 dr + (A_2 - \sqrt{\eta}A_1 \text{Tr}[\rho(t)(M + M^\dagger)])dt \tag{4.25}$$

so that the first term corresponds to proportional feedback gain and the second corresponds to a constant Hamiltonian drive that does not depend on the measurement record.

So far, we have assumed that the optimal angle is infinitesimal. However, equation Eq. (4.24) only guarantees that the solution $\theta_{\text{opt.}}$ is a local extremum and does not guarantee that it is necessarily a maximum. A sufficient condition for it to be a local maximum is that the second derivative of the fidelity function evaluated at $\theta_{\text{opt.}}$ be negative

$$\left. \frac{\partial^2 \mathcal{F}_{t+dt}(\theta)}{\partial \theta^2} \right|_{\theta=\theta_{\text{opt.}}} = -\langle \psi_T | [H_F, [H_F, \rho^c]] | \psi_T \rangle \Big|_{\theta=\theta_{\text{opt.}}} < 0. \tag{4.26}$$

Failure of this test, *i.e.*, when the second derivative is positive, suggests the presence of a local minimum from the infinitesimal solution. Then we will need a large (*i.e.*, non-infinitesimal) rotation, which we compute by maximizing the fidelity over the entire angular range.

It should be noted that A_1 and A_2 can in principle become singular. However as $\rho^c_{t-dt}|_{\theta_{\text{opt.}}} = \rho_t$, the denominator diverges only when the second derivative test failed at the previous time step (compare Eq. (4.26)) to the denominator of Eq. (4.24)). Thus this divergence is typically prevented by the global search described above. A special case is $[H_F, \rho_t] = 0$, in which case feedback has no effect on the state, so that we may simply set $\theta_{\text{opt.}} = 0$.

To simulate this form of feedback in practice, we assume that the controller chooses the rotation angle $\theta_{\text{opt.}}$ that ensures a global maximum of $\mathcal{F}_{t=0}$ at the initial time step. During evolution of the state, the above protocol typically continues to pick $\theta_{\text{opt.}}$ as the global maximum of $\mathcal{F}_t$ and thus maintains the system on a locally (time-)optimal trajectory. However even if Eq. (4.26) remains negative, it is possible that the nearest local maximum of $\mathcal{F}_t(\theta)$ can fail to be the global maximum. The only way to catch such instances is to occasionally undertake a brute-force maximization of $\mathcal{F}$ and to thereby check whether the local maximum identified by equation Eq. (4.24) is also a global maximum. In practice, such global maximization procedures are often unnecessary; table 4.2 shows many combinations of M and H that allow Eq. (4.24) and Eq. (4.26) to reproduce the indicated feedback protocols established in the control literature. Global searches are only required when indicated by equation Eq. (4.26). We explore some of these applications further in the next section.

The functions A_1 and A_2 are dependent on both the initial and target states, as well as on the state at time t. Dependence on the current state implies implicit dependence on the full measurement record, yielding a potentially non-Markovian feedback protocol in general. Non-Markovian protocols are more difficult to implement in practice. The feedback controller must either calculate $\rho(t)$ in real time or perform an exponential amount of precomputation to determine the optimal action for all possible measurement records, as illustrated in Fig. 4.4a. However we can also use the feedback master equation to simulate a memoryless controller, which is only capable of implementing Markovian feedback. We illustrate the basic concept in Fig. 4.4b. At the initial time step, many possible measurement outcomes dW are possible, leading to many possible conditioned states (red lines). A memoryless controller knows that it has applied feedback at the previous time step, but does not recall

Feedback protocol	Measurement operator (M)	Feedback Hamiltonian (H)	Target state $(\lvert\psi_T\rangle)$
Adaptive phase measurement[146]*	σ	σ_z (Heisenberg picture)	$(\lvert 0\rangle + i\lvert 1\rangle)/\sqrt{2}$
Rapid qubit purification[57]*	σ_z	σ_y	$(\lvert g\rangle + \lvert e\rangle)/\sqrt{2}$
Half-parity Bell state preparation[77]*	$\sigma_{z,1} + \sigma_{z,2}$	$\sigma_{y,1} + \sigma_{y,2}$	$(\lvert eg\rangle + \lvert ge\rangle)/\sqrt{2}$
Full-parity Bell state preparation [53, 78]*	$\sigma_{z,1}\sigma_{z,2}$	$\sigma_{x,1}$	$(\lvert gg\rangle + i\lvert eg\rangle$ $+i\lvert eg\rangle + \lvert ee\rangle)/2$
N-qubit Dicke states [133, 129, 143, 156]	$\sum_i \sigma_{z,i}$	$\sum_i \sigma_{y,i}$	N-qubit Dicke state with n excitations
N-qubit GHZ states*	$\sigma_{z,i} - \sigma_{z,j}$ for all i,j	$\sigma_{y,i} - \sigma_{y,j}$ for all i,j	N-qubit GHZ state
Hong-Ou-Mandel Bell state*	$i(\sigma_1 + \sigma_2)$, $\sigma_1 - \sigma_2$	$\sigma_{x,1} + \sigma_{x,2}$ $\sigma_{y,1} - \sigma_{y,2}$	$(\lvert gg\rangle + \lvert ee\rangle)/\sqrt{2}$

Table 4.1: Summary of some measurement-based feedback protocols that may be derived from Eq. (4.24) or its generalization in the following section. Protocols in which perfect noise cancellation (cancellation of dW terms) can occur are marked with an asterisk. Note that for N-qubit GHZ states, noise cancellation only occurs for $N = 3$.
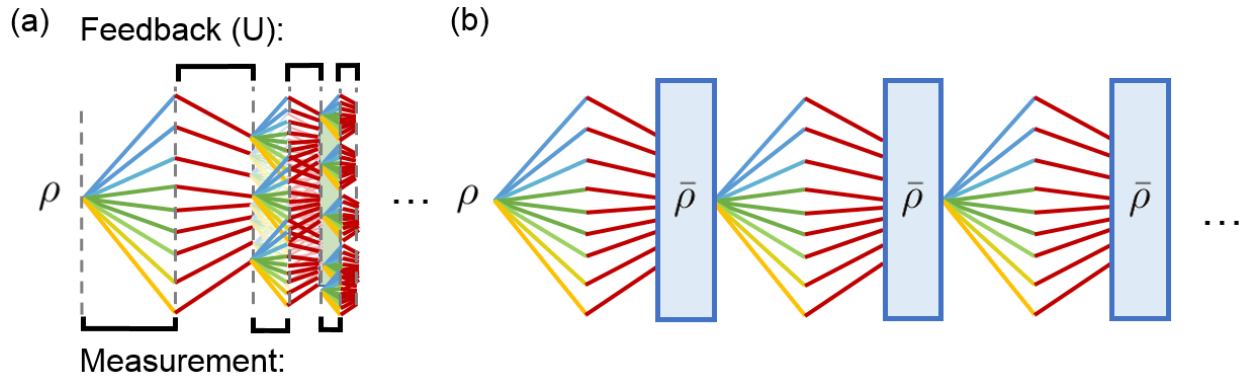
Figure 4.4: (a) A cartoon illustration of a non-Markovian feedback protocol. The number of possible states, and hence the number of potentially distinct feedback operations needed in response scale exponentially in time. (b) The behavior of a forgetful feedback controller, which only applies controls based on the evolution of the average state and the most recent measurement outcome.
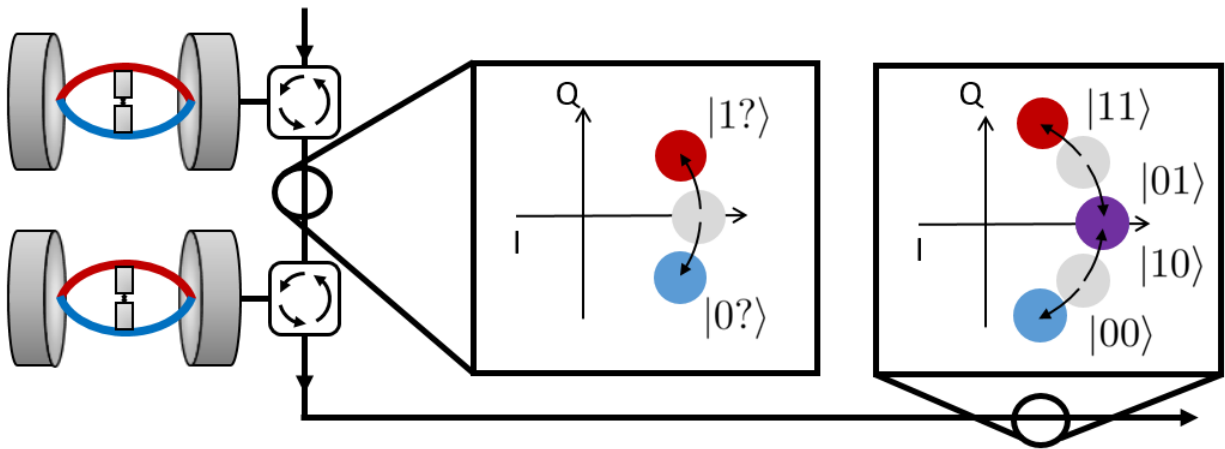
Figure 4.5: A graphical illustration of the bounce-bounce remote entanglement experiment performed in [109]. A coherent measurement tone interacts sequentially with two cavities, each housing a superconducting transmon qubit. The final coherent state picks up two conditional phase shifts, and cannot distinguish between the $|01\rangle$ and $|10\rangle$ states.

the specific measurement outcome. It therefore only has access to the average state $\bar{\rho}$, which may be computed by averaging over *i.e.* dropping the dW term in Eq. (4.23). At the next time step, the feedback controller computes $\theta_{\text{opt.}}$ based on $\bar{\rho}$ instead of ρ, which amounts to replacing ρ with $\bar{\rho}$ in Eq. (4.24). Thus we may generate a Markovian feedback protocol simply by dropping the dW terms in Eq. (4.23) and propagating the state. We call such a protocol average sense locally optimal, or ASLO for short.

In essence, ASLO protocols exploit the subjectivity of ρ, so that we can use different calculations to simulate different states of knowledge of our feedback controller. An ASLO feedback controller is memoryless in the same way as a general Markovian open system. The ASLO concept emphasizes that dW is also a subjective quantity. The measurement outcome dr is objective, but the right hand side $\sqrt{\eta}\langle M + M^{\dagger}\rangle dt + dW$ contains an expectation value that differs depending on if we use ρ or $\bar{\rho}$. To reproduce the same observable dr no matter the state of knowledge of the observer, dW must also differ based on which we use.

4.3   Bi- and Multipartite Entanglement Generation

Quantum feedback has many potential application, particularly as it can be implemented in almost any measurement process, at least in principle. A natural starting point to apply it is in situations where measurement is already known to be a useful control resource and to see if feedback offers any fundamentally new capabilities.

One of the most important applications of quantum measurement from a control standpoint is remote entanglement generation. Entanglement occurs naturally and even automat-

ically in interacting systems. Two atoms in close proximity will hybridize, and two dipole-coupled transmon qubits form Bell states when resonantly exchanging energy. A much more difficult task is to create entanglement between non-interacting or even remote systems. The standard method involves interference between signals emitted by the two systems. If the signals are indistinguishable, so that we cannot tell from which system it came, then we can implement a *degenerate measurements*. An example of such an experiment performed in superconducting circuits is shown in Fig. 4.5[109]. In the scheme, a measurement signal interacts sequentially with two different cavities separated by over a meter of cable. The cavities are dispersively coupled to superconducting qubits, so the coherent tone picks up a phase shift conditioned on the qubit states. If the phase shifts are equal in magnitude, then the final output signal cannot distinguish the $|01\rangle$ and $|10\rangle$ states. Performing a homodyne measurement of this output field implements measurement of the operator[89]

$$M = \sqrt{\frac{\Gamma}{2}} \frac{\sigma_{z,1} + \sigma_{z,2}}{2}. \tag{4.27}$$

The above operator is called a half-parity measurement, and it is degenerate in the $|01\rangle$, $|10\rangle$ subspace. This can be used to probabilistically generate entanglement by first preparing the separable uniform superposition state $|\psi_0\rangle = \frac{1}{2}(|00\rangle + |01\rangle + |10\rangle + |11\rangle)$ and then projectively measuring M. Half of the time, the measurement outcomes ± 1 occur (dropping unitful quantities) and we collapse into the separable states $|00\rangle$ or $|11\rangle$. The other half of the time, the system collapses into the $|01\rangle$, $|10\rangle$ subspace while retaining its coherence within that subspace. The result is a Bell state $|\psi^+\rangle \equiv \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)$. The above process, and many like it, are intrinsically probabilistic. Even if implemented ideally, the success rate cannot exceed 50%, which means that the unheralded state $\bar{\rho} = \frac{1}{4}|00\rangle\langle 00| + \frac{1}{4}|11\rangle\langle 11| + \frac{1}{2}|\psi^+\rangle\langle\psi^+|$ has no entanglement.

The above procedure can be made to produce a Bell state deterministically if we add in continuous feedback during the measurement process. Although not originally derived in this way, we can use Eq. (4.24) to derive the globally optimal protocol for this task. All we have to do is choose a suitable feedback Hamiltonian. First and foremost, we must restrict our feedback Hamiltonian to be local, as nonlocal Hamiltonians are simply not available when there is no interaction. Feedback should correct for measurement perturbations toward the undesired states $|00\rangle$ and $|11\rangle$. We can rotate between these states and $|\psi_0\rangle$ with a $\pi/2$ rotation of $H = (\sigma_{y,1} + \sigma_{y,2})/2$, and from $|\psi_0\rangle$ we can probabilistically prepare the target state. It seems reasonable to guess that infinitesimal rotations of H could steer the measurement outcome toward the desired subspace of M continuously, so we take it as our feedback Hamiltonian. Initializing the system in $|\psi_0\rangle$ and taking $\eta = 1$ for now, the state evolves as

$$|\psi(t)\rangle = \frac{1}{\sqrt{2}}\left(e^{-\Gamma t/4}|\phi^+\rangle + \sqrt{2 - e^{-\Gamma t/2}}|\psi^+\rangle\right) \tag{4.28}$$

where $|\phi^+\rangle \equiv \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. We have $|\psi(0)\rangle = |\psi_0\rangle$ and $|\psi(t \gg 1/\Gamma)\rangle = |\psi^+\rangle$, indicating that we can deterministically prepare an entangled state from a separable state using feedback.

One can derive Eq. (4.28) by taking general coefficients and substituting a state of this form into the feedback master equation. This ansatz amounts to assuming that feedback maintains the $0 \leftrightarrow 1$ symmetry of the initial state. A detailed derivation is provided in [77], and also in chapter 7 where we prove global optimality of the resulting protocol. For now, we note a striking feature of the result, which is that $|\psi(t)\rangle$ only depends on time, and not on any function of the measurement record. This occurs because the stochastic terms of the master equation, one coming from measurement back action and one from feedback, cancel exactly. This cancellation is somewhat surprising, given that feedback is nonlocal while measurement is not. We remark on the generality of this phenomenon at the end of this section.
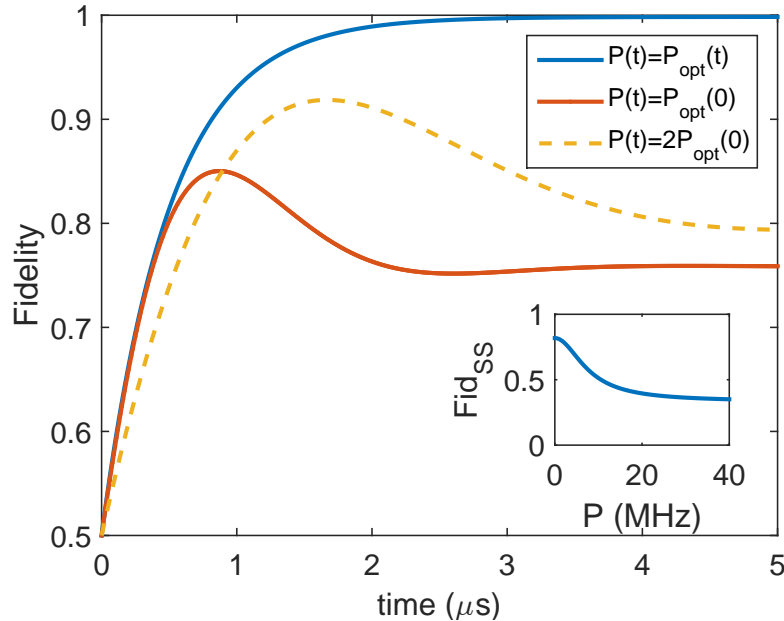
Figure 4.6: Performance of the half-parity measurement under locally optimal feedback (blue curve) and suboptimal protocols that enforce P to be fixed as a function of time. Fidelity is with respect to the target state $|\psi^+\rangle = (|01\rangle + |10\rangle)/\sqrt{2}$. Inset shows the steady state fidelity for fixed P. The highest achievable value is $9/11$.

In Fig. 4.6 we plot the fidelity as a function of time under the optimal feedback protocol. Feedback has $A_2 = 0$ and maintains $\langle M + M^\dagger \rangle = 0$, so it requires only proportional feedback. We label the proportional feedback coefficient P, as is standard in control theory. Fig. 4.6 also plots the fidelity when P is held constant at a few fixed values. Constant feedback is a common strategy in quantum control, as it is often argued that the steady state solution is the most important in application. However in this case, the fidelity does not reach 1 for any fixed value of P, indicating that a time-dependent strategy is essential for satisfactory performance. Performance in the presence of loss is given in [77], which confirms that the

protocol could be successfully applied with the experimental parameters already achieved in [109].

Dicke state generation An obvious extension of the above result is to more than two qubits. There are two particularly simple generalizations of the Bell state to N qubits: Dicke states and GHZ states. We begin with the Dicke states, which turn out to be easier to prepare given the measurement operators that tend to be accessible in nonlocal or weakly interacting systems. An n-excitation Dicke state of N qubits is defined to be

$$|N,n\rangle = \frac{1}{\sqrt{\binom{N}{n}}} \Sigma_{P \in S_N} P(|0\rangle^{\otimes(N-n)} \otimes |1\rangle^{\otimes n}) \tag{4.29}$$

where P is an operator belonging to the permutation group S_N on N qubits. $|N,n\rangle$ is a uniform superposition over all states with the same number of excitations, and is also known as a spin squeezed state. They have proposed applications in a wide range of sensing protocols, including very long baseline interferometry[44] and Heisenberg-limited measurement sensitivity. When $n = 1$, we have the usual W state.

$$|W\rangle = \frac{1}{\sqrt{N}}(|10\cdots0\rangle_N + |01\cdots0\rangle_N + \cdots + |00\cdots1\rangle_N) \tag{4.30}$$

Several previous works have demonstrated deterministic Dicke state preparation with feedback, but these protocols were state-based and hence non-Markovian[133, 129, 143]. It is interesting to ask how well a Markovian protocol performs at the same task, particularly since dynamical state estimation becomes exponentially more challenging for larger systems. The most straightforward generalization of our two-qubit protocol is to simply add more operators to M and H

$$M_N = \sigma_{z,1} + \sigma_{z,2} + \cdots + \sigma_{z,N} \tag{4.31}$$
$$H_N = \sigma_{y,1} + \sigma_{y,2} + \cdots + \sigma_{y,N}.$$

We have dropped units, which amounts to working in a unit system in which $\Gamma = 2$. H_N is still local as required. M_N can be measured by appending more qubits in cavities to the scheme outlined in Fig. 4.5, or by coupling many qubits dispersively to the same cavity. The latter method (without continuous feedback) has been applied to generate spin squeezing in cold neutral atoms coupled to an optical cavity[26].

Once again, direct application of the feedback equations of section 4.2 produces excellent results. We plot the asymptotic fidelity in Fig. 4.7 for many values of N and n. To simulate up to 100 qubits, we work in a symmetry-reduced subspace spanned by $|N,n\rangle$, which shrinks the Hilbert space exponentially. Unlike in the two-qubit case, the average state $\bar{\rho}$ does not remain pure under ASLO feedback. This indicates that the stochastic terms do not cancel in Eq. (4.23), and therefore the unaveraged dynamics (*i.e.* conditioned on the entire measurement record) depend on the entire measurement record. Curiously, the final fidelity remains above 94% despite the fact that the system does not take a predictable path through

Hilbert space, and hence an ASLO feedback controller does not know the true state. The efficacy appears to be due in part to the highly symmetric nature of the problem. The symmetry reduction to $|N, n\rangle$ has removed all degeneracy from the measurement operator M_N, so that the measurement outcome now uniquely determines the state.
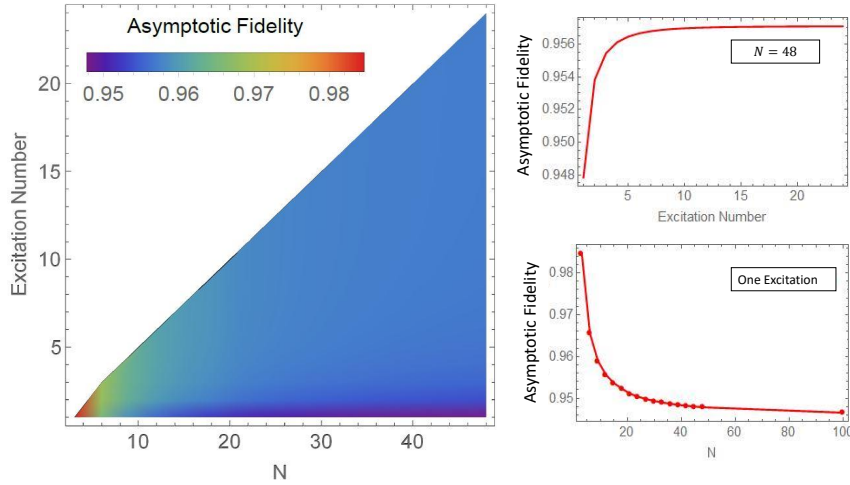
Figure 4.7: Fidelity with respect to the nth Dicke state consisting of N (N takes values from 3 to 48 in increment of 3 qubits under the ASLO (Markovian) protocol. Insets show cutouts along $N = 48$ (top) and $n = 1$ (bottom), where the latter includes also the N=100 case.

From a practical standpoint, it should be noted that although the ASLO protocol does not achieve unit fidelity beyond $N = 2$, one can still produce unit-fidelity states by adding a final projective measurement. Although the success probability under the ASLO protocol is less than one, it has been significantly enhanced by feedback.

GHZ states and multiple measurements Generation of GHZ states using the above method poses new challenges and leads to an interesting generalization of the results of section 4.2. The GHZ state is defined as

$$|GHZ\rangle \equiv \frac{1}{\sqrt{2}}(|000\rangle + |111\rangle). \tag{4.32}$$

with the obvious generalization to more than 3 qubits. As we have argued above, measurement operators that are linear combinations of single-body observables like σ_z are easiest to implement in remote and weakly interacting settings. The GHZ state is an eigenvector of operators of the form

$$M = 2\sigma_{z,1} - s\sigma_{z,2} - (1 - s)\sigma_{z,3}. \tag{4.33}$$

However the above operator does not satisfy the permutation symmetry of the target state, leaving us with no obvious symmetry reduction and no way to break the degeneracy of M.

In practice, this choice of measurement operator does not seem amenable to GHZ state generation with continuous feedback.[1]

One way to reintroduce permutation symmetry to the problem is to allow for multiple measurement operators. The GHZ state is an eigenstate of the following operators

$$M_{(i,j)} = \frac{\sigma_{z,i} - \sigma_{z,j}}{2}, \ i \neq j \tag{4.34}$$

which are just like the half-parity measurement operators but are instead degenerate in the $|00\rangle$, $|11\rangle$ subspace of each i, j pair of qubits.

To find the locally optimal protocol, we need to generalize the expression for $\theta_{\text{opt.}}$ to handle multiple simultaneous measurements M_i and feedback operations H_j. In general there is no preferred pairing between them, so that the ith measurement outcome may affect how we apply the jth feedback Hamiltonian. This forces us to rederive the feedback master equation with a more general feedback unitary

$$U = \exp\left(-i\sum_{ij} A_{ij}(\rho)H_j dW_i - i\sum_i B_i(\rho)H_i dt\right) \tag{4.35}$$

$$= I - i\sum_{ij} A_{ij}H_j dW_i - \left[i\sum_i B_i H_i + \frac{1}{2}\sum_{ijk} A_{ij}A_{ik}H_j H_k\right] dt$$

where A_{ij} and B_i are analogous to A_1 and A_2 respectively, and our goal is to find their locally optimal values. The presence of cross terms in the above modifies the resulting feedback master equation

$$\rho(t + dt) = \rho + \sum_i \left[\mathcal{D}[M_i]\rho \ dt + \sqrt{\eta_i}\mathcal{H}[M_i]\rho \ dW_i - i\sum_j A_{ij}[H_j, \rho]dW_i - iB_i[H_i, \rho]dt\right.$$

$$\tag{4.36}$$

$$\left. - i\sum_j \sqrt{\eta_i}A_{ij}[H_j, M_i\rho + \rho M_i^\dagger]dt + \sum_{jk} A_{ij}A_{ik}[H_k\rho H_j - \frac{1}{2}(H_j H_k\rho + \rho H_j H_k)]dt\right]$$

In contrast to Eq. (4.23), the last term is not in Lindblad form due to the presence of cross terms, though it may be cast into Lindblad form by defining the modified feedback operators $\tilde{H}_i = \sum_j H_j$. For simplicity, we assume that the control Hamiltonians commute pairwise, though the end result is essentially unmodified if one relaxes this assumption[81].[2] The locally optimal feedback coefficients must satisfy $\langle\psi_T|[H_\alpha, d\rho]|\psi_T\rangle = 0$, this time for all α.

---

[1]This includes attempts that involve beyond-locally optimal strategies and strategies based on entanglement monotones such as the tangle. See [156] for details.

[2]In general, one must have that the H_i form a Lie algebra *i.e.*, that the vector space formed by H_i is closed under commutation.

Dealing first with the dW terms, we find

$$\sum_i \left[\sum_j -iA_{ij} \underbrace{\langle \psi_T|[H_\alpha, [H_j, \rho]]|\psi_T \rangle}_{c_{j\alpha}} + \underbrace{\sqrt{\eta_i}\langle \psi_T|[H_\alpha, M_i\rho + \rho M_i^\dagger]|\psi_T \rangle}_{a_{i\alpha}} \right] dW_i = 0 \qquad (4.37)$$

The solution is evident if we rewrite the expression in matrix form with the help of a and c, which are state-dependent

$$-iAc + a = 0 \quad \implies \quad A = -iac^{-1} \qquad (4.38)$$

In general, one should use the Moore-Penrose pseudoinverse to handle the case in which c is not invertible[81]. With a solution for A in hand, we can solve for the B coefficients by collecting $\mathcal{O}(dt)$ terms from $\langle \psi_T|[H_\alpha, d\rho]|\psi_T \rangle = 0$

$$b_\alpha \equiv \langle \psi_T|[H_\alpha, \sum_i \left[\mathcal{D}[M_i]\rho - i\sum_j \sqrt{\eta_i}A_{ij}[H_j, M_i\rho + \rho M_i^\dagger] \right. \qquad (4.39)$$

$$\left. + \sum_{jk} A_{ij}A_{ik}[H_k\rho H_j - \frac{1}{2}(H_j H_k \rho + \rho H_j H_k)] \right]]|\psi_T \rangle$$

$$-i\sum_i B_i c_{i\alpha} + b_\alpha = 0 \quad \implies \quad \vec{B} = -i\vec{b}c^{-1}$$

treating $\vec{b}$ and $\vec{B}$ as row vectors. Eq. (4.38) and Eq. (4.39) generalize Eq. (4.24). They even allow for the unusual situation in which we apply simultaneous non-commuting measurements, a situation that we consider experimentally in the following chapter.

A significant advantage of this generalization is the extension to multiple feedback Hamiltonians, even if there is only one measurement operator (note that we did not have to assume that the number of measurement and feedback operators need not coincide). In the previous formulation, we had to guess the best feedback Hamiltonian based on intuition, whereas now we can let H_i be a Hamiltonian basis for the system and let the equations of motion determine the locally optimal Hamiltonian for us. As before, we should discard Hamiltonians that commute with the state. The condition to have found a local maximum as opposed to a minimum or saddle point is that the Hessian matrix

$$-\langle \psi_T|[H_\alpha, [H_j, \rho^c]]|\psi_T \rangle \qquad (4.40)$$

should be negative definite.

To generate GHZ states, we need to choose a set of feedback operators. For simplicity, we limit ourselves to local σ_y rotations. c can fail to be invertible if there are to many H_is, and we have learned the importance of symmetry, so some care is required in choosing them. One obvious choice of Hamiltonian basis is $H_i = \sigma_{y,i}$, one for each qubit. However, this basis

does not enforce the $|11...1\rangle \leftrightarrow |00...0\rangle$ symmetry of the GHZ state and the $M_{(i,j)}$s. A better choice would be a set orthogonal to $\sum_i \sigma_{y,i}$, such as

$$H_i = \frac{\sigma_{y,i} - \sigma_{y,i+1}}{2}, \quad i = 1...N - 1 \tag{4.41}$$

where N is the number of qubits. Although this basis appears to break permutation symmetry, any Hamiltonian of the form

$$H_{(i,j)} = \frac{\sigma_{y,i} - \sigma_{y,j}}{2} \tag{4.42}$$

can be written as a linear combination of them. Thus the locally optimal solution can retain permutation symmetry (and should, unless the protocol explicitly breaks symmetry, which is possible and would happen in any case).

GHZ generation under ASLO feedback is plotted in Fig. 4.8. The 2-qubit case reproduces the half-parity measurement under the transformation $\sigma_{z,2} \leftrightarrow -\sigma_{z,2}$, while the remaining results $N = 3$ through $N = 8$ are novel. Although we use Eq. (4.41) as our feedback Hamiltonians, the locally optimal protocol takes linear combinations of the form Eq. (4.42), so that each $H_{(i,j)}$ is paired with its corresponding $M_{(i,j)}$. We speed up the numerics by taking this observation into account, so that many terms in Eqs. (4.36-4.40) may be set to zero. Perhaps most interestingly, the $N = 3$ case reaches exactly unit fidelity, and the purity of the average state remains 1 throughout the feedback process, as in the example at the beginning of this section. One can derive an analytic solution in this case. We suspect that the protocol is globally optimal, but have not yet taken the time to prove this potential result.
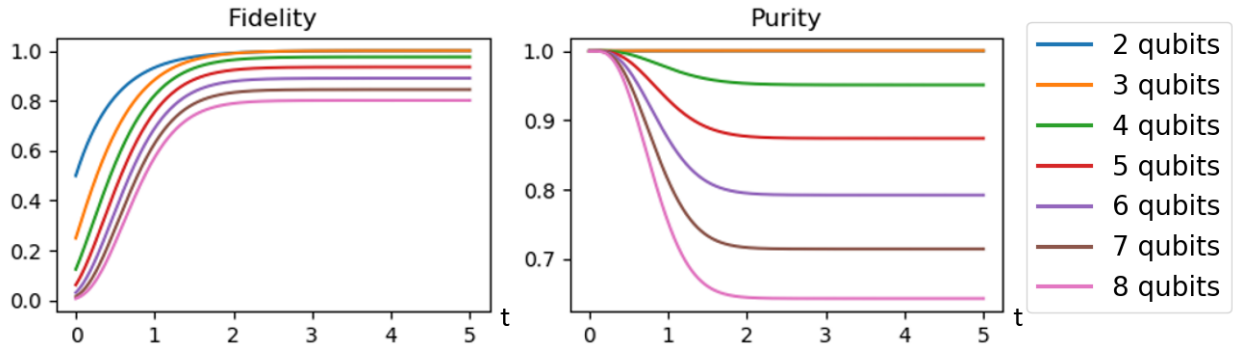
Figure 4.8: Fidelity and purity of the average state under the feedback protocol described in the text. Purity remains exactly 1 for $N = 2$ and $N = 3$.

While we have focused on application to entanglement generalization, the full scope of Eqs. (4.36-4.40) extends to virtually any application of measurement-based feedback. The ability of these equations to reproduce all existing feedback protocols that we have tried them on and the ease with which they generate new protocols suggests that there is much more to

be gained from them. An interesting future direction would be to characterize more broadly the types of states that may be generated efficiently using these results. Entangled state generation in the context of Markovian dissipation has yielded interesting general results in other contexts[59, 134], and would be natural to apply and extend here. An important direction for future applied work would be to classify what states can be prepared deterministically (under a suitable definition) using local feedback and measurement operators that are linear combinations of local observables.

Deterministic evolution We have seen a number of intriguing examples in which feedback perfectly eliminates the randomness of the measurement process. These cases are highlighted with asterisks in table 4.2. All of these protocols are globally optimal for the task at hand, except for the 3-qubit GHZ protocol that we nevertheless conjecture to be optimal. They also yield analytic solutions in all cases, further indicating that they are a natural class of protocols in continuous measurement.

It is straightforward to introduce a strict necessary condition for whether or not deterministic evolution is possible given fixed measurement and feedback operators. This condition also provides a simple tool for finding solutions to the Wiseman-Milburne feedback master equation. Suppose that we apply proportional feedback of the form $U = \exp(-iH(Pdr+cdt))$ for time-dependent constants c and P. Applying U to the stochastic Schrödinger equation yields a pure-state equation of motion for $|\psi\rangle$ valid for $\eta = 1$

$$|\psi(t+dt)\rangle = \Big[I + (M - m - iPH)dW - \frac{1}{2}(M^\dagger M - 2Mm + m^2 + P^2 H^2)dt \qquad (4.43)$$
$$- iH(P(M+m)+c)dt\Big]|\psi(t)\rangle$$

where $m \equiv \langle M + M^\dagger\rangle/2$. We should allow for the solution to be deterministic up to a global phase, so we make the substitution

$$H \to H + V_0. \qquad (4.44)$$

For a deterministic solution, the dW terms of the stochastic Schrödinger equation must cancel *i.e.*

$$[M - m - iPH - iPV_0]|\psi\rangle = 0. \qquad (4.45)$$

This is a nonlinear constraint on $|\psi\rangle$, as m depends on the state. It can only be satisfied if $|\psi\rangle$ is a right eigenvector of $M - iPH$. To see that this necessary condition is also sufficient, compute the corresponding eigenvalue λ by taking an expectation value

$$\langle\psi|[M - iPH]|\psi\rangle = \lambda = \langle M\rangle - iP\langle H\rangle \qquad (4.46)$$

and then compute the left-hand side of Eq. (4.45)

$$[M - m - iPH - iPV_0]|\psi\rangle = \Big[\langle M\rangle - iP\langle H\rangle - \frac{\langle M + M^\dagger\rangle}{2} - iPV_0\Big]|\psi\rangle \qquad (4.47)$$
$$= \Big[\frac{\langle M^\dagger - M\rangle}{2} - iP\langle H\rangle - iPV_0\Big]|\psi\rangle$$

As $\langle M^\dagger - M \rangle$ is purely imaginary, we can always find a V_0 such that the above expression is zero. Thus feedback can exactly cancel the measurement-induced noise on $|\psi\rangle$ if and only if $|\psi(t)\rangle$ is a right eigenvector of $M - iPH$. As P is a free parameter, we can check if this condition is possible for any value of it.

With the stochastic terms cancelled, the next step is to see if the deterministic part of the equations of motion can maintain the above condition for a finite time interval. Given the above result, we can search for such solutions within each eigenspace of $M - iPH$ separately (the eigenspaces will typically evolve continuously as a function of P). Let $\{|v_i(P(t))\rangle\}$ be a set of orthonormal eigenvectors associated with a chosen eigenspace. To obtain a deterministic solution, we must have

$$|\psi(t)\rangle = \sum_i c_i(t)|v_i\rangle. \tag{4.48}$$

Now consider the deterministic part of the equation of motion. Using the above expansion, we have

$$\frac{d|\psi\rangle}{dt} = \sum_i \dot{c}_i(t)|v_i\rangle + c_i(t)\frac{dP}{dt}\frac{d|v_i\rangle}{dP} \tag{4.49}$$

$$= \left[\frac{1}{2}(M^\dagger M - 2Mm + m^2 + P^2 H^2)dt - iH(P(M+m) + c)dt\right]|\psi(t)\rangle.$$

As $m = \mathrm{Re}(\lambda)$ and $\lambda(t)$ is fixed by working within a fixed eigenspace, the above equation of motion is linear in $|\psi\rangle$, which significantly simplifies the finding of a solution. One can use the above procedure to identify candidate solutions, which appear to be plentiful and relatively unexplored. For example, the 3-qubit GHZ state protocol yields an analytic solution under the above procedure.

Chapter 5

Simultaneous Measurement of Non-Commuting Observables

In the previous chapters, we discussed a significant generalization of the more basic projective, or Von Neumann description of measurement. Continuous measurement and POVMs lead to a great deal of novel physics that would not even be well defined in the standard formalism. We now turn to our first experimental application of this body of research. We revisit one of the most elementary and fundamental predictions of quantum mechanics, the Heisenberg uncertainty principle.

The Heisenberg microscope thought experiment and the resulting uncertainty principle were among the earliest realizations of the importance of non-commuting observables in quantum mechanics. As it became clear that the basic properties of a system cannot be simultaneously determinate, doubt arose as to whether or not it was even possible to simultaneously measure them. Indeed, in the Von Neumann interpretation of measurement as instantaneous collapse of the system into an observable's eigenstate, the notion of simultaneous non-commuting measurements is ill-defined, because a system cannot simultaneously exist in the eigenstates of two non-commuting observables even mathematically.

A more accurate and subtle understanding arose in the following decades. Researchers came to realize that physical measurements occur over a finite timescale, and that interactions with multiple baths can indeed lead to non-trivial combinations of measurements. In fact, because most systems couple to their environments in multiple different ways, they are consequently measured in non-commuting bases in their natural settings. For example, charge fluctuations, phonons and spin baths all potentially form Markovian reservoirs, and can couple via different system operators. A ubiquitous example is the duo of dephasing (T_2) and relaxation (T_1), induced by the non-commuting operators σ_z and σ respectively. However, despite the fact that this process is happening in all realistic systems, researchers have not resolved multiple non-commuting dissipation channels as the measurements that they are until the following work.

In the experiment of this chapter, we begin with a transmon qubit dispersively coupled to two modes of a 3D cavity. These two modes act as two separate baths for our system.

However because our experiment relies on the dispersive coupling to two cavity modes, we natively only couple to σ_z. Thus having two modes yields simultaneous measurement of two commuting observables. In order to access non-commuting observables, we devised a scheme that lets us change the measurement axis, which we call a single-quadrature, or stroboscopic measurement. We implement this scheme by adapting a technique from optomechanics known as a back-action evading measurement[20]. We also describe a natural generalization to arbitrary quantum systems, which we term a multi-quadrature measurement.

In section 5.3, we apply the above technique to measure two non-commuting observables of a qubit. We derive a novel connection to the uncertainty principle that governs the dynamics of the state by enforcing a lower bound on the measurement-induced disturbance. Consequently, as a function of the measurement non-commutativity, the dynamics transition smoothly from standard wave function collapse to isotropic, persistent diffusion. Although evolution of the state now differs drastically from that of a conventional measurement, information about both non-commuting observables is extracted by keeping track of the time ordering of the measurement record. We use maximum likelihood techniques to perform a novel form of quantum state tomography that does not involve alternating measurements in different bases.

Before going into experimental details specific to this work, we begin in section 5.1, outlining the crucial steps and techniques for implementing high-quantum efficiency measurements in circuit QED. This section is important for any quantum trajectory, quantum feedback or adaptive measurement experiment, including the experiment of chapter 6. The following sections describe work first presented in [48], with an elaborated description of single-quadrature measurements and the introduction of multi-quadrature measurements (not described in the original work). In section 5.3, we present the bulk of the experimental work, which includes a novel connection between measurement back-action and the uncertainty principle. A more lengthy and technical calculation, the polaron transformation deriving our effective stochastic master equation, is relegated to the last section, section 5.4.

5.1 Achieving high quantum efficiency measurements

In principle, continuous measurement are implemented in all experiments at some level. While they may be difficult to resolve temporally in some physical systems, the bandwidths of microwave and optical detectors easily reach the GHz range, and time resolutions for some photodetectors are measured in picoseconds. Loss presents a much more serious problem. The difficulty of a realistic quantum trajectories experiment is that by the time an experimentalist's detector has acquired a finite amount of information about the system in question, the environment at large has likely already learned far more. Once complete information has been transferred to the environment, then the residual coherence that makes quantum trajectories and feedback interesting has been destroyed. Thus the primary experimental parameter limiting the observation of quantum trajectories is quantum efficiency. Only if a large fraction of the total amount of information content emitted by the measured system is
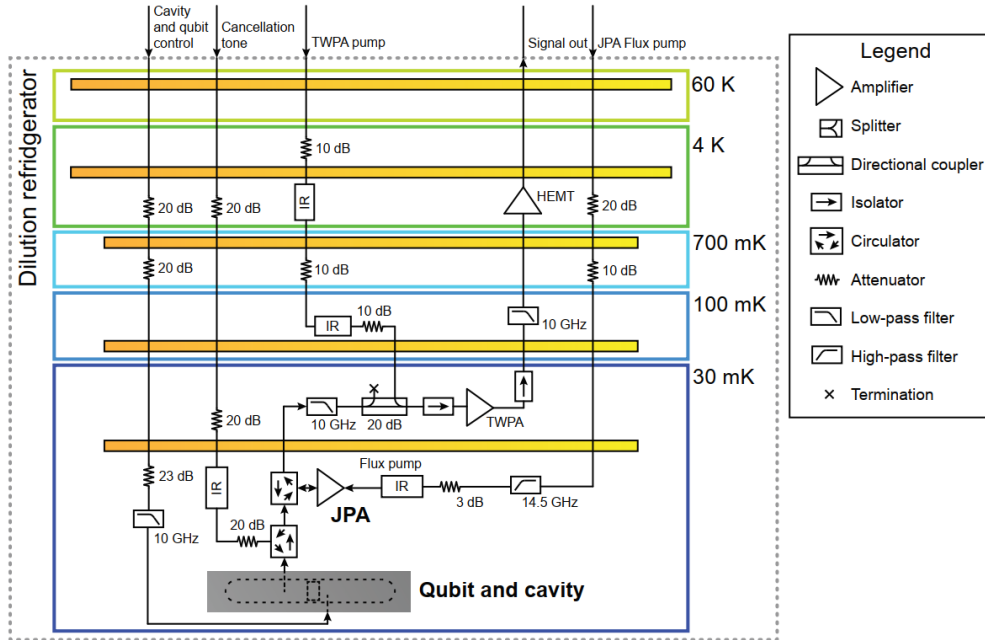
Figure 5.1: Wiring diagram for the experiment outlined in chapter 6, which captures the essential features of a high quantum efficiency setup. The JPA amplifies in reflection and is pumped and feedback-controlled via the flux pump line. IR filters are Eccorsorb or similar.

in our hands can we experimentally observe quantum trajectories and their consequences.

Fig. 5.1 shows the low-temperature wiring diagram for a quantum trajectories experiment. The diagram is actually for the experiment of the following chapter, which is less cluttered because it contains only one measurement chain instead of two. The first critical aspect of the setup is the assortment of input line attenuators distributed throughout the refrigerator. To maintain low temperatures at the qubit and parametric amplifier, we must filter out blackbody radiation from higher temperature stages. At microwave frequencies, we eliminate power with standard microwave attenuators. To eliminate infrared power, light-tight sealing and Eccosorb filters are essential[12].

To determine the required level of attenuation, we can use equipartition theorem to estimate the amount of thermal power travelling through the lines. Over a fixed bandwidth at high temperature, $E \propto k_B T$, so it is sensible to ignore prefactors between power and temperature and measure of blackbody radiation power in units of Kelvin.[1] At low temperatures (relative to the photon quanta energy), we will tend to overestimate the total power, which is better for making conservative estimates. The power emitted by an attenuator consists of the attenuated incident thermal power plus the blackbody radiation emitted by

---

[1]If absolute or total integrated powers are needed, it is important to note that the T^4 power dependence of the Stefan Boltzmann law is modified when we consider a 1-dimensional system like a microwave transmission line.

the attenuator itself. Taking the top right attenuator of Fig. 5.1 as an example, it is held at a temperature of 4 K and receives 300 K noise at its input. 20 dB = ×100, so the emitted power is $T_{\text{out}} = T_{\text{in}}/100 + 99T_{\text{atten.}}/100 \approx 7$ K. Note that this only applies if the attenuator is thermally anchored to the 4 K stage, which requires a metallic clamp around the portion of the attenuator that dissipates power. Noise power at the base-temperature circuits should be comparable to the base plate temperature of 30 mK. The energy scale of a 1 GHz photon is 48 mK, so GHz-scale circuits can be quite sensitive to thermal power. For example, thermal population of a readout resonator is often a dominant source of dephasing[142].

Control line attenuation can be large, as we may compensate for it by increasing the control power during operation. We also naturally gain the potentially weak coupling of the control line to the device as an additional source of attenuation, as well as the spectral filtering of a cavity if one is present. Achieving proper thermal isolation for the output line is more challenging. Attenuation is not an option if we wish to attain high quantum efficiency, and we also require strong coupling of the output line to the device itself (relative to the control line coupling strength, so that information does not leak out of the input ports). The output line in Fig. 5.1 is filtered with circulators, wide-band isolators and low-pass filters (such as K&L brand). The presence of amplifiers further restricts what components must be used. For instance, ideally a single circulator would suffice between the qubit and the JPA, but two often appear to be necessary in practice. The amplifier emits a substantial amount of power over its bandwidth, and reflections can interfere with ideal operation. The HEMT (high electron mobility transistor amplifier) dissipates enough power that it must remain at the 4 K stage cooled by the pulse tube. Aluminum does not superconduct at this temperature, so unlike other lines, one must use another material such as niobium for the connection.

In moving a microwave signal from 30 mK to room temperature, amplifiers are absolutely necessary. We can model a linear amplifier as an ideal amplifier preceded by a noise source. Each must amplify the signal above the noise level of the subsequent device, so that we can eventually digitize the signal with a room temperature analog-to-digital converter. A standard way to characterize the performance of an amplifier is a signal-to-noise (SNR) improvement measurement using a spectrum analyzer. Data from a low-frequency JPA are shown in Fig. 5.2a. To measure the SNR improvement, we apply a weak coherent tone to the amplifier with it powered off and record the signal power and noise power. We do not require a calibrated power measurement here, though it is important to use a noise marker to measure the noise power quantitatively.[2] We then measure the signal and noise powers again with the amplifier on. Noise rise is defined to be the ratio of noise powers, and gain is of course the ratio of signal powers. SNR improvement is the ratio of these ratios. At

---

[2]Spectrum analyzers work by filtering the input signal over a narrow bandwidth and then plotting the transmitted power as a function of the filter center frequency. They are calibrated to report the absolute power assuming that the input field is infinitely narrow in frequency. As noise is actually broadband, the detected signal power depends on the transmission profile of the filter. A noise marker calibrates out the filter profile and reports the noise level in units of power per unit bandwidth = energy, which is the appropriate set of units for noise.
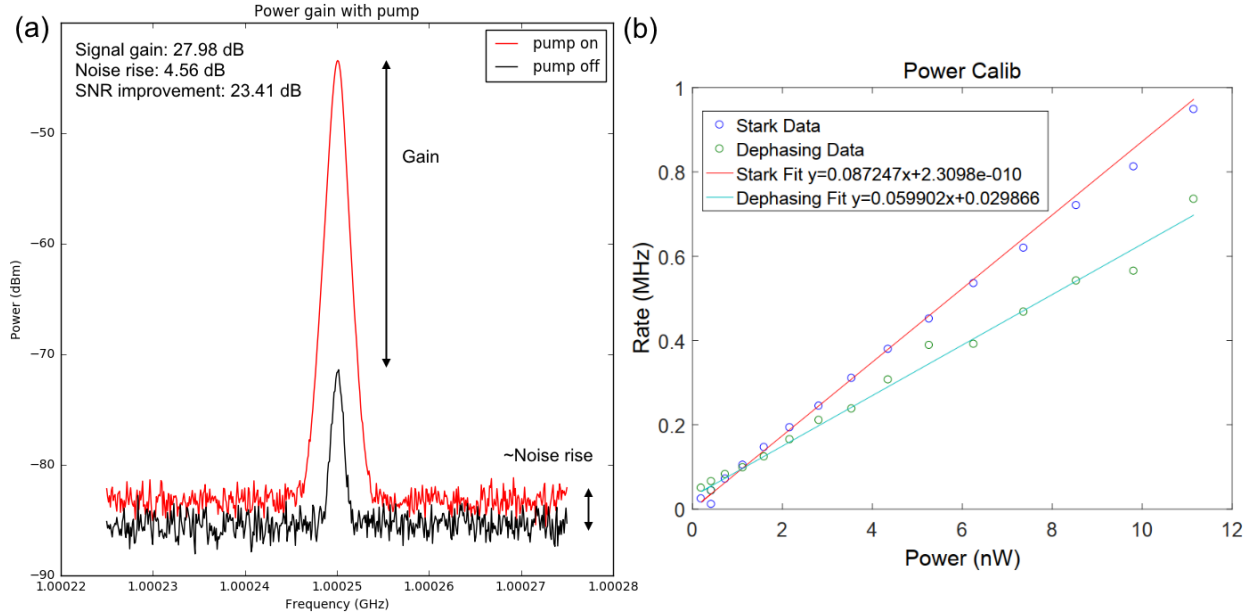
Figure 5.2: (a) Spectrum analyzer data from an SNR improvement measurement of a single-pumped JPA. Axes plot power, and hence are not accurate for noise power (which would have units of mW/Hz, not mW. A weak signal is applied at 1.00025 GHz, which is by necessity detuned from the pump frequency of 1 GHz. The SNR improvement is unusually high, indicating loss or a high noise temperature at the HEMT. Data courtesy of Akel Hashim. (b) χ and photon number calibration data. Stark shift and dephasing rate fits taken from Ramsey measurements are fit to a calibrated power axis. The slopes of the Stark shift and dephasing rate curves are $2\chi/A$ and $8\chi^2/\kappa A$ respectively (assuming Δ_{AC} is measured as an angular frequency. If not, we must include a factor of 2π in Eq. (5.2)).

low gain, the SNR increases with increasing gain, but the measured noise is dominated by devices after the amplifier and remains constant. At high gain, the dominant noise source becomes noise amplified by the amplifier itself, so gain and noise rise increase in tandem. To achieve high quantum efficiency, we should work in the latter regime.

SNR improvement measurements provide a way to check the sufficiency of an amplifier, but do not provide absolute metrics on device performance. A key figure of merit is noise temperature, which is the effective temperature at the input, assuming that all output noise arises from ideally amplified input noise from a hypothetical source. For this we need some kind of calibrated power source. A simple method is to source the amplifier with an impedance matched termination at a known temperature. In practice however, we would like to measure the amplifier performance while it is connected to something more interesting. Fortunately, if the measurement chain is connected to a qubit in a cavity, then the latter system provides a convenient calibrated power source via the dispersive Jaynes Cummings Hamiltonian. The essential ingredients for this process are that $H = \chi\sigma_z a^\dagger a$ provides con-

version between photonic energy and qubit frequency, and the cavity linewidth κ provides a conversion between internal cavity energy and output power. Once we know the output power from the cavity, we can convert the measured power at a spectrum analyzer to absolute power units referenced to the amplifier input. For example, if we read a power of -40 dBm at the spectrum analyzer but know that the true input power is $\kappa\langle a^\dagger a \rangle = -128$ dBm, then this provides a conversion factor to infer the effective noise power at the amplifier input from a room temperature noise power measurement. In reality, some noise is due to loss and subsequent amplifiers, but ideally noise from the amplifier closest to the source dominates. Due to the uncertainty principle, the lowest possible noise power at a given frequency is $\omega \times \hbar/2$, the power per unit bandwidth of vacuum fluctuations.

κ is easy to measure, as it is the full-width half-maximum of the cavity transmission profile (multiplied by 2π), in accordance with input-output theory. The challenge of the above measurement is the precise determination of χ and $\langle a^\dagger a \rangle$. If $\chi > \kappa$, then the qubit resonance splits into many resonances each separated by 2χ, one for each value of the intercavity photon number $a^\dagger a$ with non-negligible population. By driving the cavity with a probe tone matched to the cavity resonance (conditioned on the qubit being in its ground state) and then sweeping another tone over the qubit resonance, we can observe these peaks. When the qubit tone hits a qubit resonance, the cavity frequency is shifted, which in turn shifts the phase of the cavity tone. We can detect the resulting phase shift in the cavity tone if we apply it using a vector network analyzer (VNA). The distribution of peaks and their relative heights is determined by the coherent amplitude of the cavity tone via Eq. (1.24), and $\langle a^\dagger a \rangle = |\alpha|^2$.

The procedure is somewhat more involved when $\chi < \kappa$, as the large cavity bandwidth masks the discrete photon-number peaks[138]. To measure χ, we linearly sweep the power of a cavity-resonant tone and monitor the resulting qubit Stark shift and dephasing rate with a Ramsey measurement. Using Eq. (3.57) and the standard input-output relation for a cavity ($\dot{\alpha} = -i\Delta\alpha - \kappa\alpha/2 + \epsilon$), one can derive the following expressions for the Stark shift and dephasing rates in the limit $\chi \ll \kappa$

$$\Delta_{\mathrm{AC}} = 2\chi\bar{n} \tag{5.1}$$

$$\Gamma_D = \frac{8\chi^2\bar{n}}{\kappa}$$

where $\bar{n} = |\alpha_{g/e}|^2$ is the intercavity photon number, which takes the same value regardless of the qubit state. The general expressions for any χ, κ and cavity-drive detuning Δ are straightforward to derive as well. Experimentally, we can measure Δ_{AC} and Γ_D for any $\bar{n}$ by fitting the Ramsey data to

$$\langle \sigma_z \rangle = e^{-\Gamma_D \tau} \cos(\Delta_{\mathrm{AC}}\tau + \delta) \tag{5.2}$$

where τ is the delay between $\pi/2$ pulses in the Ramsey measurement. By fitting the linear relations Eq. (5.1) together, we can extract χ and a calibration factor A between actual power P applied from a room temperature source and the intercavity photon number $\bar{n}$ (*i.e.*

$P = A\bar{n}$). A provides the calibration factor needed to produce a known power at the input of an amplifier coupled to the cavity.

Quantum efficiency The above procedures offer a starting point for debugging unexpected features in a high quantum efficiency measurement chain. Quantum efficiency itself is the ratio of the measurement-induced dephasing rate to the rate at which information is acquired. Formally, if we write the master equation for a qubit with $M = \sqrt{\Gamma_D/2}\sigma_z$ as

$$d\rho = \mathcal{D}[M]\rho\, dt + \sqrt{\eta}\mathcal{H}[M]\rho\, dW \tag{5.3}$$
$$= \frac{\Gamma_D}{2}\mathcal{D}[\sigma_z]\rho\, dt + \sqrt{\frac{\Gamma_M}{2}}\mathcal{H}[\sigma_z]\rho\, dW$$

then the quantum efficiency is typically defined as

$$\eta \equiv \frac{\Gamma_M}{\Gamma_D}. \tag{5.4}$$

Notice that η is not quite a signal-to-noise ratio, but rather a ratio of rates.

Γ_D measures the total rate at which information is extracted from the system by everything else in the universe, and we already know how to estimate it using a Ramsey measurement. Given the convention above, it is also the rate at which the off-diagonal elements of the density matrix $\rho_{e,g}$ and $\rho_{g,e}$ decay, which can be measured by fitting a Ramsey trace to Eq. (5.2). Measured in this way, Γ_D also includes damping due to the intrinsic qubit dephasing time T_2^*. One can compute the true dephasing rate as $\Gamma_D = \Gamma_{D,\text{fit}} - 1/T_2^*$, or measure Γ_D at multiple readout powers ($\bar{n}$) to measure and subtract off the zero-intercept.

We can infer Γ_M from the signal-to-noise ratio of the integrated measurement record

$$dr = \sqrt{\eta}\langle M + M^\dagger\rangle dt + dW \tag{5.5}$$
$$= \sqrt{2\Gamma_M}\langle\sigma_z\rangle dt + dW$$

If we integrate the signal over a fixed time interval τ, we obtain two Gaussians with variances $\sigma^2 = \tau$ and means $\mu_{g/e} = \pm\sqrt{2\Gamma_M}\tau$ conditioned on the prepared qubit state. The measurement can then be inferred with

$$\Gamma_M = \frac{(\mu_e - \mu_g)^2}{8\tau\sigma^2}. \tag{5.6}$$

This formula is a ratio of signal amplitudes, and so any arbitrary multiplicative factor between the numerical value of the signal and dr cancels. In practice, linear fits in τ are advisable, as finite measurement bandwidth introduces offsets.[3]

---

[3]On timescales short compared to the measurement signal correlation time, dW no longer approximates its statistics well, and the variance scales as τ^2 instead of τ.

5.2 Single and multi-quadrature measurements

As described in chapter 2, dispersive readout is the standard method of readout in circuit QED. Here we present an alternative scheme that provides several novel and useful capabilities. Firstly, it enables dynamical changing of the measurement axis even during measurement[47], a useful capability for adaptive measurement and quantum state tomography. Secondly, it provides a method to implement longitudinal $(a + a^\dagger)$ as opposed to transverse $(a^\dagger a)$ coupling to a resonator without modification of the underlying circuit. This enables enhancement of qubit readout with squeezed light [32], and stabilizer measurements for the Gottesman-Kitaev-Preskill qubit encoded in a cavity[34, 136]. Finally, we show how our scheme has a built-in method for rapidly and deterministically depopulating the cavity after readout, enabling complete measurement in a duration shorter that the inverse cavity bandwidth. In this chapter, we will be primarily be interested in the ability to change the measurement axis in-situ.

Conveniently, the following measurement scheme derives from the dispersive Jaynes Cummings Hamiltonian, meaning that we can implement it in the standard circuit QED setup. The basic concept involves Rabi driving the qubit fast enough that the dispersive coupling term averages out and effectively decouples. We 'reactivate' it by applying sidebands to the cavity, which restore a resonant coupling in a different basis. For a qubit, it allows us to measure any operator orthogonal to the Rabi drive basis. For a multi-level system like a transmon or a set of qubits, it enables measurement of system operators which are *hollow* in the basis of the drive Hamiltonian, meaning that the diagonal elements of the matrix are zero. By Schur-Horn theorem, any traceless matrix can be cast into hollow form by change of basis. As adding a multiple of the identity to a measurement operator does not change its eigenbasis, it follows that our scheme can measure any system operator.

An important advantage of multi-quadrature measurements is that they can also be applied concurrently, as in the present experiment. The only restriction on the set of possible simultaneous measurements is the requirement that all measured operators must be hollow in the same basis. We note that an extension of the scheme presented in reference [139] can be applied to measure in the eigenbasis of the drive Hamiltonian, thus providing access to an operator with a non-zero diagonal. Combining these schemes allows for a form of continuous tomography in which the full set of tomographic measurements can be applied simultaneousely. Even without this scheme however, the hollowness constraint is not severe. Many important sets of measurement operators satisfy this property, including the parity measurements $\{XX, YY\}$, the repetition code stabilizers $\{XXI, IXX\}$ and the five-qubit code stabilizers $\{XZZXI, IXZZX, XIXZZ, ZXIXZ\}$ capable of protecting against single bit- and phase-flip errors on the encoded logical qubit.

Single-quadrature measurements We begin with a detailed derivation of a single-quadrature measurement (abbreviated SQM), and then outline the multi-quadrature (MQM) generalization. The system described in this chapter involves two single-quadrature measurements applied to the same qubit. The experimental setup, shown in Fig. 5.3a consists of a Rabi-driven 3D transmon qubit coupled to the two lowest frequency modes of a 3D super-
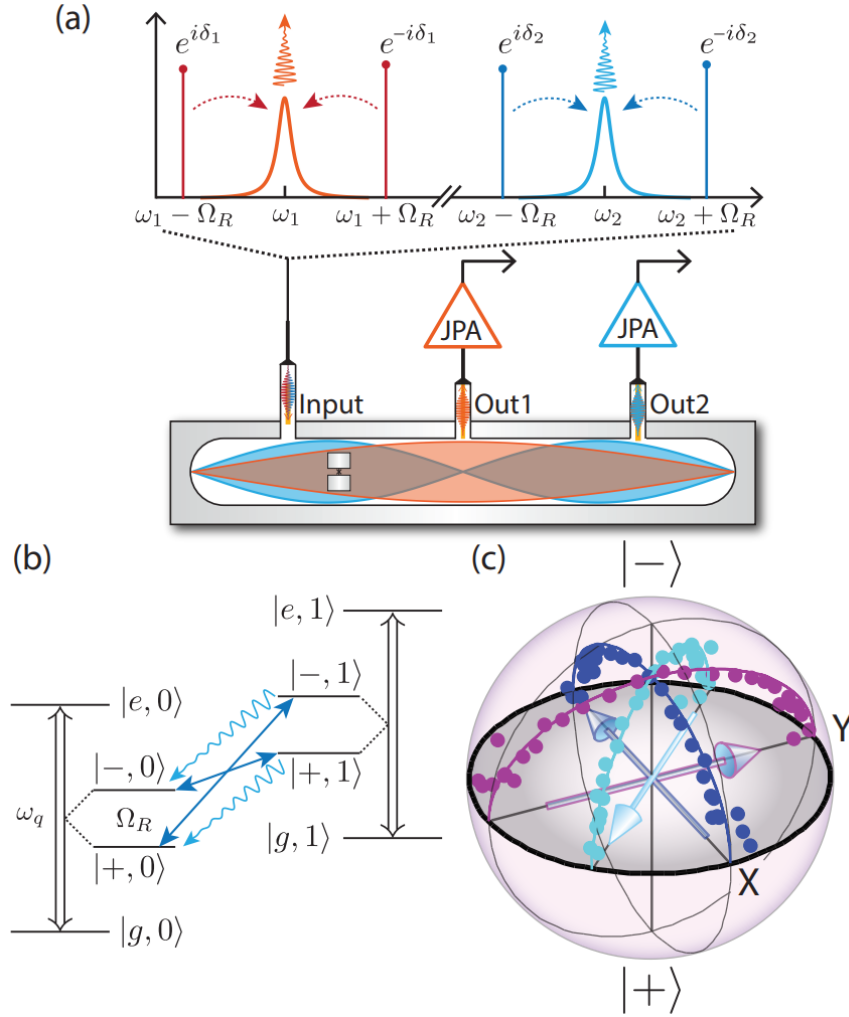
Figure 5.3: **Experimental implementation of a multimode single quadrature measurement (SQM)** (a) A transmon qubit in an aluminum cavity. Via the input port, we drive Rabi oscillations at frequency Ω_R and apply $\pm\Omega_R$ sideband tones to each of the two lowest cavity modes. Each output is monitored using a lumped-element Josephson parametric amplifier (LJPA)[18]. Coupling is designed so that $\sim$90% of each signal is routed to its corresponding LJPA. This system yields two separate measurement channels, each with a controllable measurement basis. (b) Dressed state picture of the SQM scheme for one of the cavity modes. Sideband tones drive transitions indicated by the solid lines, and undulating lines represent cavity decay, which we detect. (c) Tomographic reconstruction after a 1 μs SQM applied only to the lower mode, showing collapse along three separately chosen measurement axes σ_{δ_1}, $\delta_1 = \{0, \pi/4, \pi/2\}$. Circles plot tomographic data and lines give theory based on a dephasing rate of $\Gamma_1/2\pi = 122$ kHz and quantum efficiency of $\eta_1 = 0.41$.

conducting cavity. The key aspect of an SQM are the cavity sidebands depicted above the
cavity. Altogether, the system is described by two of the standard dispersive Hamiltonian
with time-dependent drives

$$H = H_q + H_{\text{drive}} + \sum_{i=\{1,2\}} (\omega_{c,i} + \chi_i \sigma_z) a_i^\dagger a_i \tag{5.7}$$

$$H_q = \frac{\omega_q}{2} \sigma_z + \Omega_R (\sigma + \sigma^\dagger) \cos(\omega_d t)$$

$$H_{\text{drive}} = \sum_{i=\{1,2\}} \epsilon_i(t) a_i^\dagger + \epsilon_i(t)^* a_i$$

where $\epsilon_i(t) \in \mathbb{C}$ is the coherent drive that will represent the SQM sideband drives. Details of
physical parameters are given in [48]. For now we note that $\Omega_R \gg \kappa$ is an important physical
constraint. It prevents sidebands from populating the cavity significantly, and ensures that
an application of the rotating wave approximation holds in what follows.

An intuitive description of the SQM is depicted in Fig. 5.3b. Suppose we wish to create
an effective longitudinal interaction like $\sigma_x(a + a^\dagger) = (\sigma + \sigma^\dagger)(a + a^\dagger)$. Such a Hamiltonian is
typically not relevant in nature because of the rotating wave approximation, which averages
out the σa and $\sigma^\dagger a^\dagger$ terms. However it is possible to generate a resonant interaction if both
the qubit and cavity operators have commensurate time dependence. To create this situation,
we drive the qubit at a frequency Ω_R. In the rotating frame of the cavity, the sideband tones
have a time dependence given by the sideband detuning. If we set this detuning to be $\pm \Omega_R$,
then we can expect a resonant interaction. As visualized in Fig. 5.3b, the drive Hamiltonian
leads to dressed states with an effective energy splitting of Ω_R in the rotating frame of the
qubit. The lower $(\omega_c - \Omega_R)$ simultaneously excites the cavity and removes energy from the
qubit, taking it from the $|+\rangle$ to the $|-\rangle$ dressed states. This process sounds just like the
Jaynes-Cumming interaction $\sigma a^\dagger + \sigma^\dagger a$, so we are back to where we started. However we can
also add an upper sideband $(\omega_c + \Omega_R)$ that simultaneously creates or destroys qubit and cavity
excitations $\sigma^\dagger a^\dagger + \sigma a$. Thus the combination of an upper and a lower sideband gives us the the
longitudinal coupling that we wanted. Moreover the relative phase between these two terms
can be controlled by varying the relative phase of the sidebands, which yields a controllable
measurement axis. In what follows, we justify the above intuition mathematically.

If the cavity decay channels κ_i are monitored with parametric amplifiers with quantum
efficiencies η_i, the qubit+cavity dynamics are modelled by the following stochastic master
equation

$$d\rho = -i[H, \rho]dt + \sum_{i=\{1,2\}} \kappa_i \mathcal{D}[a_i]\rho \, dt + \sqrt{\kappa_i \eta_i} \mathcal{H}[a_i e^{i\phi_i}]\rho \, dW_i \tag{5.8}$$

where ϕ_i is the amplification axis of the phase sensitive amplifier used to monitor mode i.
For the derivation, it suffices to consider just one of the two cavity modes. To measure a
single quadrature of the qubit, we drive the cavity with a pair of sidebands detuned by $\pm \Omega_R$
from the cavity resonance, and with phases $\pm(\delta - \text{atan}(\kappa/2\Omega_R))$. The second term in the

phase cancels the lag between the external drive and the internal cavity field. As $\Omega_R \gg \kappa$, it can be ignored in our experiment. We choose this drive $\epsilon(t) = -i\bar{a}_0\sqrt{\Omega_R^2 + \kappa^2/4}\sin(\Omega_R t + \delta - \mathrm{atan}(\kappa/2\Omega_R))e^{-i\omega_c t}$ so that the cavity internal displacement due to the sidebands takes the relatively simple form

$$\bar{a}(t) = \bar{a}_0 \cos(\Omega_R t + \delta)e^{-i\omega_c t} \;\rightarrow\; \bar{a}_0 \cos(\Omega_R t + \delta). \tag{5.9}$$

The right arrow indicates that we have transformed into the interaction picture with respect to the cavity. We include the complicated prefactor in $\epsilon(t)$ to compensate for the cavity-drive detuning, and soon leads to simplifying cancellations that justify Eq. (5.9). We take $\bar{a}_0$ to be real for simplicity.

   To understand the effects of the cavity drive, we transform into a displaced frame in which the cavity is in the vacuum (neglecting for a moment the cavity-qubit interaction). This time dependent unitary transformation is $U(t) = D[-\bar{a}(t)] = \exp(\bar{a}^* a - \bar{a}a^\dagger)$. As we noted in chapter 1 around Eq. (1.25), this transformation maps $a \rightarrow a + \bar{a} = D[-\alpha]aD[-\bar{a}]^\dagger$, which greatly simplifies calculations. We first transform the cavity dissipation terms, in direct correspondence with Eq. (1.9)

$$\kappa\mathcal{D}[a]\rho \rightarrow U(t)\kappa\mathcal{D}[a + \bar{a}]\rho = \kappa\mathcal{D}[a]\rho - i[H_{\mathrm{dis}}, \rho] \tag{5.10}$$

$$H_{\mathrm{dis}} \equiv -\frac{i\kappa}{2}(\bar{a}a^\dagger - \bar{a}^* a)$$

$$\sqrt{\kappa\eta}\mathcal{H}[ae^{i\phi}]\rho \rightarrow \sqrt{\kappa\eta}\mathcal{H}[ae^{i\phi}]\rho.$$

Eq. (1.9) also tells us that we must add a term $H' = -iU(t)\dot{U}^\dagger(t) = i(\dot{\bar{a}}a - \dot{\bar{a}}a^\dagger)$ to the Hamiltonian. A bit of algebra shows that H_{dis} and H' cancel the drive term H_{drive} from Eq. (5.7), which completely eliminates cavity drive terms from the Hamiltonian. In this sense, $\bar{a}$ is a steady state, which confirms that Eq. (5.9) gives the internal cavity displacment. Transforming the remaining terms of H for one cavity mode into the interaction picture (*i.e.* transforming the Hamiltonian by $\exp[i\omega_c a^\dagger a t + i\omega_d \sigma_z t/2]$), applying the rotating wave approximation, and then applying the displacement transformation above yields

$$\begin{aligned}
H_{\mathrm{int}} &= \chi a^\dagger a \sigma_z + H_{q,\mathrm{int}} \tag{5.11}\\
&= \chi[(a^\dagger + \bar{a}^*)(a + \bar{a})]\sigma_z + H_{q,\mathrm{int}}\\
&= \chi\left[\bar{a}_0\cos(\Omega_R t + \delta)(a^\dagger + a) + \frac{\bar{n}_0}{2} + \bar{n}_0\frac{\cos(2\Omega_R t + 2\delta)}{2} + a^\dagger a\right]\sigma_z + H_{q,\mathrm{int}}\\
H_{q,\mathrm{int}} &= \frac{1}{2}[(\omega_q - \omega_d)\sigma_z + \Omega_R\sigma_x]
\end{aligned}$$

where we have defined $\bar{n}_0 = \bar{a}_0^2$. Choosing the Rabi drive to be resonant (*i.e.* $\omega_q - \omega_d = -\chi\bar{n}_0$, which accounts for the Stark shift term), we diagonalize the qubit drive term of the Hamiltonian by going into the Hadamard frame ($\sigma_z \leftrightarrow \sigma_x$), and then going into a frame rotating at the Rabi frequency ($\exp[i\Omega_R\sigma_z t/2]$). These transformations eliminate $H_{q,\mathrm{int}}$ and

map σ_z to $\sigma e^{-i\Omega_R t} + \sigma^\dagger e^{i\Omega_R t}$, so the Hamiltonian becomes

$$H_{\text{q-frame}} = \frac{\chi \bar{a}_0}{2} (e^{i\Omega_R t + i\delta} + e^{-i\Omega_R t - i\delta})(a^\dagger + a)(\sigma e^{-i\Omega_R t} + \sigma^\dagger e^{i\Omega_R t}) \tag{5.12}$$

$$+ \chi \left[\bar{n}_0 \frac{\cos(2\Omega_R t + 2\delta)}{2} + a^\dagger d \right] (\sigma e^{-i\Omega t} + \sigma^\dagger e^{i\Omega_R t}).$$

Dropping terms which rotate at Ω_R or $2\Omega_R$, we are left with

$$H_{\text{q-frame}} = \frac{\chi \bar{a}_0}{2} (a^\dagger + a)(\sigma e^{i\delta} + \sigma^\dagger e^{-i\delta}) = \tilde{g} \sigma_\delta (a^\dagger + a)$$

$$\sigma_\delta \equiv \cos(\delta)\sigma_x + \sin(\delta)\sigma_y$$

$$\tilde{g} \equiv \frac{\chi \bar{a}_0}{2} \tag{5.13}$$

which is a qubit-state-dependent cavity drive Hamiltonian. In the same way that a qubit-state-dependent phase shifts induced by $\sigma_z a^\dagger a$ yields a σ_z measurement of the qubit, Eq. (5.13) yields a measurement of σ_δ. We defer the remainder of this calculation to section 5.4, though note that the methods of section 3.4 would also yield a valid stochastic master equation. The most important detail of the result is the measurement rate

$$\Gamma = 2\chi^2 \bar{a}_0^2 / \kappa. \tag{5.14}$$

Eq. (5.13) also yields a simple method to turn off the measurement quickly. In dispersive readout, one typically waits a duration longer than $1/\kappa$ after turning off the coherent cavity drive. There are methods to depopulate the cavity coherently after readout, but the simplest scheme suffers from the fact that the optimal depopulating displacement depends on the measurement outcome. In an SQM, the $\bar{a}_0$ prefactor gives us a way to invert the sign of H and hence undo its action. A large pulse can depopulate the cavity relatively quickly. This scheme is particularly convenient for the MQM scheme, which benefits from a smaller κ.

One minor experimental imperfection arises which is worthwhile considering theoreticall. In practice, we generate the cavity sideband tones by amplitude-modulating a cavity-resonant tone that we call the local oscillator at a frequency Ω_R using an arbitrary waveform generator and a diode-based microwave mixer. In principle, amplitude modulation creates two sidebands at $\omega_c \pm \Omega_R$ and eliminates the local oscillator completely. In practice, diode mixers must be nulled with a DC bias. This nulling is rarely perfect, which leads to a small coherent tone resonant with the cavity in addition to $\epsilon(t)$. Adding this LO leakage to Eq. 5.9 gives $\bar{\alpha}_0 \cos(\Omega_R t + \delta) + \bar{a}_{\text{LO}}$, and the effective Hamiltonian would instead be

$$H_{\text{q-frame}} = \tilde{g} \sigma_\delta (a^\dagger + a + 2\text{Re}[\bar{a}_{\text{LO}}]). \tag{5.15}$$

Thus LO leakage induces unwanted rotations about the σ_δ axis.

To demonstrate the above scheme, and in particular control of the measurement axis, we perform a single SQM using the lower cavity mode. Further details of the experimental
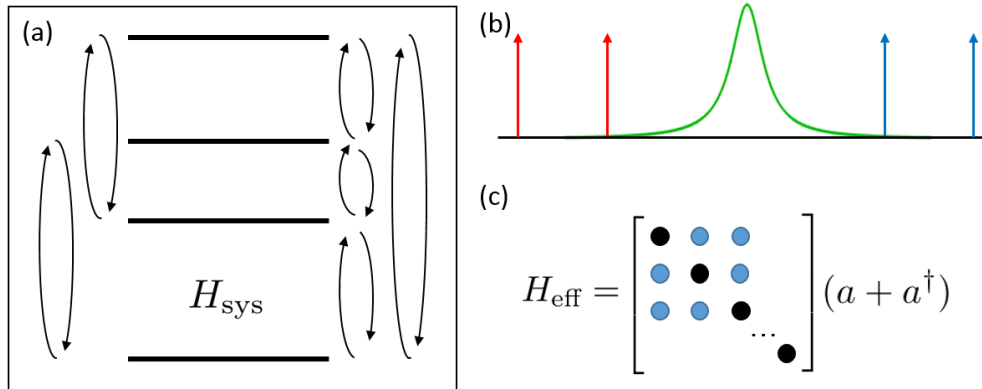
Figure 5.4: (a) Possible transitions that may be driven and stroboscopically read out via a multi-quadrature measurement in a 3-level system. (b) Red- and blue-detuned sidebands applied to a cavity. The sideband detuning is greater than the cavity linewidth. (c) A schematic illustration of the multi-quadrature measurement Hamiltonian. Only off-diagonal elements of the coupling operator are time-dependent and therefore may be addressed by a sideband.

setup are given in [48]. We prepare the $|-\rangle$ state and read out for 1 μs, which is comparable to the inverse measurement time of 1.3μs. Because we stop the on the timescale of the rate of collapse, the system exhibits residual coherence. We observe this residual coherence by including one of six possible pulses and a projective measurement using a standard dispersive readout of σ_z in the lab frame at the end of each experimental sequence. These pulses, called quantum state tomography, allow us to effectively measure σ_x and σ_y in addition to σ_z, which enables reconstruction of the density matrix. Fig. 5.3c shows the results of these tomographic measurements for three different measurement axis orientations. We compute the data points by using the stochastic master equation to predict ρ for each experimental run, clustering these datasets into groups with similar final states, and then plotting the Bloch vector for each subset. The final states lie along 3 arcs for the three measurement axes chosen. The final state is somewhat impure because the measurement efficiency is less than 1. Note that the three measurements applied in Fig. 5.3c come from three separate experiments, so we have not yet applied simultaneous non-commuting measurements.

Multi-quadrature measurements Before describing the dynamics of simultaneous non-commuting measurements, we turn to the concept of a multi-quadrature measurement. We called the above measurement single-quadrature because it singles out one quadrature σ_δ of the qubit Rabi oscillations. This arose from driving the sideband transitions depicted in Fig. 5.3. For an N-level system, there are up to $N(N+1)/2$ distinct transitions, each of which can be driven and coupled to the cavity via sidebands (see Fig. 5.4a and b). The picture is somewhat complicated by the fact that these transitions interact with one another, but we can nevertheless handle this system analytically in the general case.

Consider a completely general N-level system dispersively coupled to a cavity, written in

the cavity frame and the qubit drive frame

$$H = H_\chi a^\dagger a + H_d \tag{5.16}$$

where H_d is an external drive Hamiltonian on the N-level system and H_χ is a system operator representing the dispersive coupling to the cavity. In writing Eq. (5.16) to be time-independent, we have implicitly taken H_χ to be diagonal in the system energy eigenbasis, as is usually the case in practice (for example $H_\chi = \chi\sigma_z$ for a qubit). Note that the coherent system drive term H_d could include off-resonant drive terms, so it need not be hollow (see for example Eq. (1.15) with $\Delta \neq 0$). We now transform H into the eigenbasis of H_d, so that $VH_dV^\dagger = \Lambda$ is diagonal. In this basis, it is simple to go into the frame of the drive Hamiltonian, as $U(t) = \exp[i\Lambda t]$ is also diagonal. In this frame, H becomes

$$H = U(t)VH_\chi V^\dagger U^\dagger(t)a^\dagger a. \tag{5.17}$$

As U is time-dependent, it will couple to sideband tones applied on the cavity. Much of the first derivation of this section amounted to showing that we can put a resonator in a specified time-dependent coherent state using coherent drives. We take this fact for granted and begin by assuming a specific time-dependence of the internal cavity state $\bar{a}(t)$. We now transform into the displaced frame

$$H = U(t)VH_\chi V^\dagger U^\dagger(t)(a + \bar{a}(t))^\dagger(a + \bar{a}(t)) \tag{5.18}$$
$$= Q(t)(a + \bar{a}(t))^\dagger(a + \bar{a}(t))$$
$$Q_{ij}(t) \equiv \sum_{ij}(VH_\chi V^\dagger)_{ij}e^{-i(\Lambda_{jj} - \Lambda_{ii})t}$$

We suppose that the dressed state transition frequencies $\Lambda_{jj} - \Lambda_{ii}$ are non-zero and nondegenerate, as would typically be the case for a generic H_d. In this case, each pair of off-diagonal elements of Q_{ij} rotate at a different frequency, and thus may be addressed individually. Under a few assumptions on H_d, we can measure the hollow, Hermitian operator X using the following choice of sideband tones

$$\bar{a}(t) = \sum_{i \neq j}\left|\frac{X_{ij}}{Q_{ij}}\right|\cos((\Lambda_{jj} - \Lambda_{ii})t + \delta_{ij}) \tag{5.19}$$

where $\delta_{ij} = \arg(X_{ij})$. If $Q_{ij} = 0$, then the singularity in the above expression means that $X_{ij} = 0$. Eq. (5.19) leads to the following effective Hamiltonian under the rotating wave approximation

$$H_{\text{eff}} = X(a + a^\dagger) \tag{5.20}$$

The primary assumption going into Eq. (5.20) is that $VH_\chi V^\dagger$ is hollow, so that $Q(t)$ has no time-independent matrix elements and we can therefore drop $Q(t)|\bar{a}(t)|^2$ and $Q(t)a^\dagger a$. We

also assume that the sideband frequencies work out such that the crossterms within $|\bar{a}(t)|^2$ do not lead to additional resonances with $Q(t)$.

To check hollowness experimentally, one can apply H_d without sideband tones, apply a standard dispersive readout and then look for state-dependent cavity frequency shifts via the term $Q(t)a^\dagger a$. We can also derive an appropriate H_d with the following necessary and sufficient condition for H_χ to be hollow in the eigenbasis of H_d

$$\text{Tr}[H_d^n H_\chi] = 0 \quad \text{for } n = 0, ...N - 1 \tag{5.21}$$
$$\Lambda_{ii} \neq 0$$
$$\Lambda_{ii} \neq \Lambda_{jj}.$$

The second constraint can be satisfied by adding an inconsequential multiple of the identity to H_d, and we already assumed the third constraint in order to ensure that all matrix elements of Q can be addressed with a sideband of non-zero frequency. This formula lets us solve a linear system of homogeneous polynomial equations to find an appropriate H_d given a fixed H_χ. In practice, one can generate many solutions to Eq. (5.21) based on any additional experimental constraints, and then pick a solution that also satisfies the second and third constraints.

To derive Eq. (5.21), note that if we write it in the eigenbasis of H_d, then $\text{Tr}[\Lambda^n V^\dagger H_\chi V] = \text{Tr}[\Lambda^n Q(0)] = 0$ is trivially satisfied if Q is hollow. To prove the reverse implication, rewrite it as

$$\text{Tr}[\Lambda^n Q(0)] = \sum_i \Lambda_{ii}^n Q_{ii}(0) = 0 \tag{5.22}$$

This equation implies that $Q_{ii}(0) = 0$ if

$$\det \begin{pmatrix} 1 & 1 & \dots & 1 \\ \Lambda_1 & \Lambda_2 & \dots & \Lambda_N \\ \vdots & & & \\ \Lambda_1^{N-1} & \Lambda_2^{N-1} & \dots & \Lambda_N^{N-1} \end{pmatrix} = \prod_{1 \leq i < j \leq N-1} (\Lambda_j - \Lambda_i) \neq 0. \tag{5.23}$$

The above matrix is called a Vandermonde matrix, and it is clearly invertible if (and only if) the Λ_is are distinct.

For the purposes of measuring non-commuting observables, Eq. (5.21) also provides a way to check if a set of operators $\{X_m\}$ can be made hollow in the same basis, which determines if we measure them simultaneously. One searches for a Hermitian matrix H_d satisfying Eq. (5.21) simultaneously for each X_m in place of H_χ. We are unaware of any simpler condition that can perform the same test. This is in contrast to the question of whether a set of operators can be made diagonal in the same basis, which is possible if and only if they commute. Nevertheless Eq. (5.21) is usable for small systems to which we might end up applying an MQM.

5.3 Dynamics under incompatible measurements and a connection to the uncertainty principle

To simultaneously measure two non-commuting observables, we apply the SQM to the two lowest modes of the cavity housing our transmon qubit. Each pair of sidebands sets a measurement axis, and thus we are able to control them independently. Because non-commutativity leads to dependence on the time ordering of measurement outcomes, the integral of the output signal is no longer sufficient to uniquely determine the qubit state as it was in the previous section. Nevertheless, the combined measurement records encode complete information, and we are able to track quantum trajectories of an initially known state. We calculate the trajectories not using the master equation, but instead using the single-measurement solution to the master equation over a short (16 ns) time interval and alternatingly applying them iteratively. Both of these objects are treated in the next section, but are faimiliar to the reader from chapter 3. This procedure is more numerically stable and guarantees positivity of the density matrix. We validate this theory using quantum state tomography and the post-selection process described in the previous section. Figure 5.5a shows two such traces with their tomographic validation from data in which the measurement axes are set to be orthogonal, $\sigma_{\delta_1} = \sigma_x$ and $\sigma_{\delta_2} = \sigma_y$.

Due to the disturbances introduced by measurement incompatibility, extraction of an initially unknown quantity, such as a Hamiltonian parameter or system observable, requires use of the *combined* measurement records and their full time orderings. In particular, estimation methods must rely not only on the statistics of the measurement records, but also on some estimate of this disturbance. We encode this information into a single Kraus operator $\Omega_{r_1(0;T),r_2(0;T)}$ associated with each measurement record. To compute it, we use a Kraus operator $\Omega_{r_1(t),r_2(t)}$ that generates the master equation (Eq. (5.31) derived in the following section) and build up the full Kraus operator from the measurement outcomes integrated over $\Delta t = 16$ ns intervals

$$\Omega_{r_1(0;T),r_2(0;T)} = \Omega_{r_1(T),r_2(T)}\Omega_{r_1(T-\Delta t),r_2(T-\Delta t)}...\Omega_{r_1(0),r_2(0)} \tag{5.24}$$

Using Eq. (5.24), we can compute the probability for an observed record $\{r_i(0;T)\}$ to occur as $P(r_1, r_2) = \text{Tr}[\Omega_{r_1,r_2}\rho\Omega_{r_1,r_2}^\dagger]$. By maximizing this probability for all measurement records over all possible initial states, we perform maximum likelihood estimation of the initial state. Figure 5.5b shows reconstruction of sixteen state preparations again in the case of orthogonal measurement axes. Agreement within the confidence interval demonstrates that despite the complicated dynamics induced by competing observables, our scheme performs as a measurement, and that it extracts information about σ_x and σ_y simultaneously.

A state undergoing a non-commuting measurement exhibits dynamics beyond those of usual wavefunction collapse. We directly observe this evolution by measuring the probability distributions of the resulting density matrices. Figure 5.6a shows the steady state probability distributions. When the measurement axes align, the system collapses to one of two points at the poles of the Bloch sphere as expected for commuting measurement operators. When
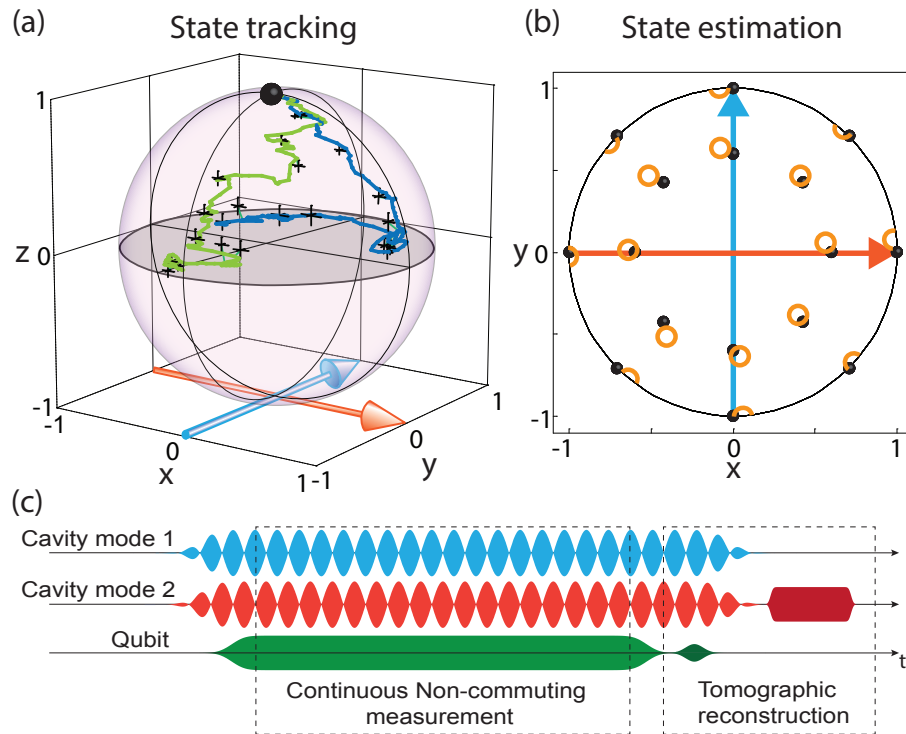
Figure 5.5: **Validation of simultaneous non-commuting measurements** (a) Reconstruction of two quantum trajectories initialized at state $|-\rangle$ (black sphere) and their associated tomographic validation. Arrows indicate axis of each measurement. Tomography points are taken every 200 ns, and are linked to the corresponding trajectory time point with a green line. All trajectories ending within ± 0.11 around a given point along the plotted trajectory are used for the corresponding tomography reconstruction. Error bars are generated from statistical uncertainty arising in qubit readout. (b) Estimation of 16 initially unknown state preparations (black spheres) using non-commuting measurements. Orange circles mark 95% confidence intervals, generated by applying maximum likelihood estimation to $\sim 10,000$ trajectories each. (c) Pulse sequence, showing cavity sideband tones and Rabi drive. Tomographic reconstruction consists of a tomography qubit pulse and projective readout.
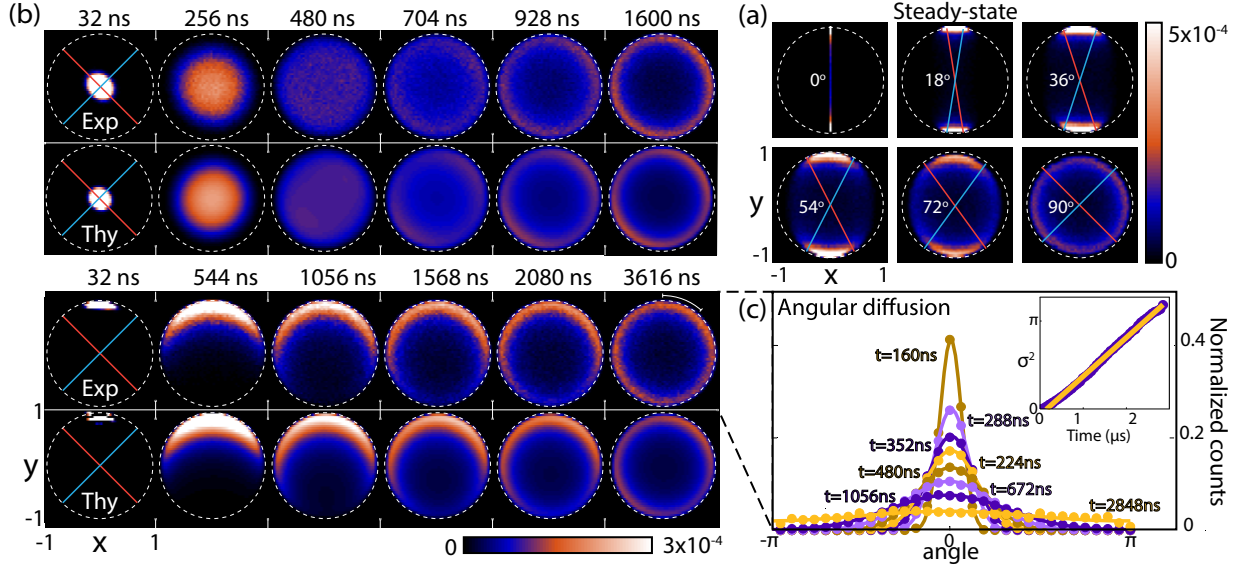
Figure 5.6: **Probability distribution of the density matrix** (a) in the steady-state
as function of angle, between the axes, demonstrating loss of collapse for non-commuting
observables. (b) as function of time for perpendicular measurement axes. We prepare a
mixed state for observing the radial dynamics (top), and a state with purity of $P = 0.89$
for observing the azimuthal dynamics (bottom). Upper row of each is experimental data,
and bottom row is theoretical comparison derived from the Fokker-Planck equation (see
methods). (c) Angular probability distributions for states within a ring of inner radius 0.86
and outer radius 0.92, showing that dynamics match those of a random walk. Points are
normalized counts and lines are fits to normal distributions convolved with a 2π periodic
Dirac comb. Inset – variance as a function of time (violet) and linear fit (yellow) yielding a
slope of 1.4 μs^{-1}. The expected slope for a perfect random walk in our system is 1.5 μs^{-1}
with the main source of uncertainty due to measurement rate uncertainty of about 10%.

the axes are separated by a finite angle less than 90 degrees, the state does not collapse to
any point, but rather localizes to regions of finite size defined by the measurement axes. A
salient feature is that when the axes are orthogonal and hence maximally incompatible, the
location of the measurement axes no longer leaves any imprint on the state evolution. The
distinct regions merge and we lose all notion of collapse.

Figure 5.6b shows probability distributions as a function of time for this canonical case.
Starting in a mixed state, we see that the state purifies isotropically to a mean steady-state
radius given by $r = \sqrt{\eta}$. High quantum efficiency in our system allows us to observe the
azimuthal dynamics by preparing a state of purity $P = 0.89$, the most likely steady-state
purity. As predicted by Ruskov *et al.* [114] in the case of unit quantum efficiency, competition
between the maximally incompatible observables leads to diffusive motion given by a uniform
random walk. Figure 5.6c shows the angular distribution as a function of time, which agrees
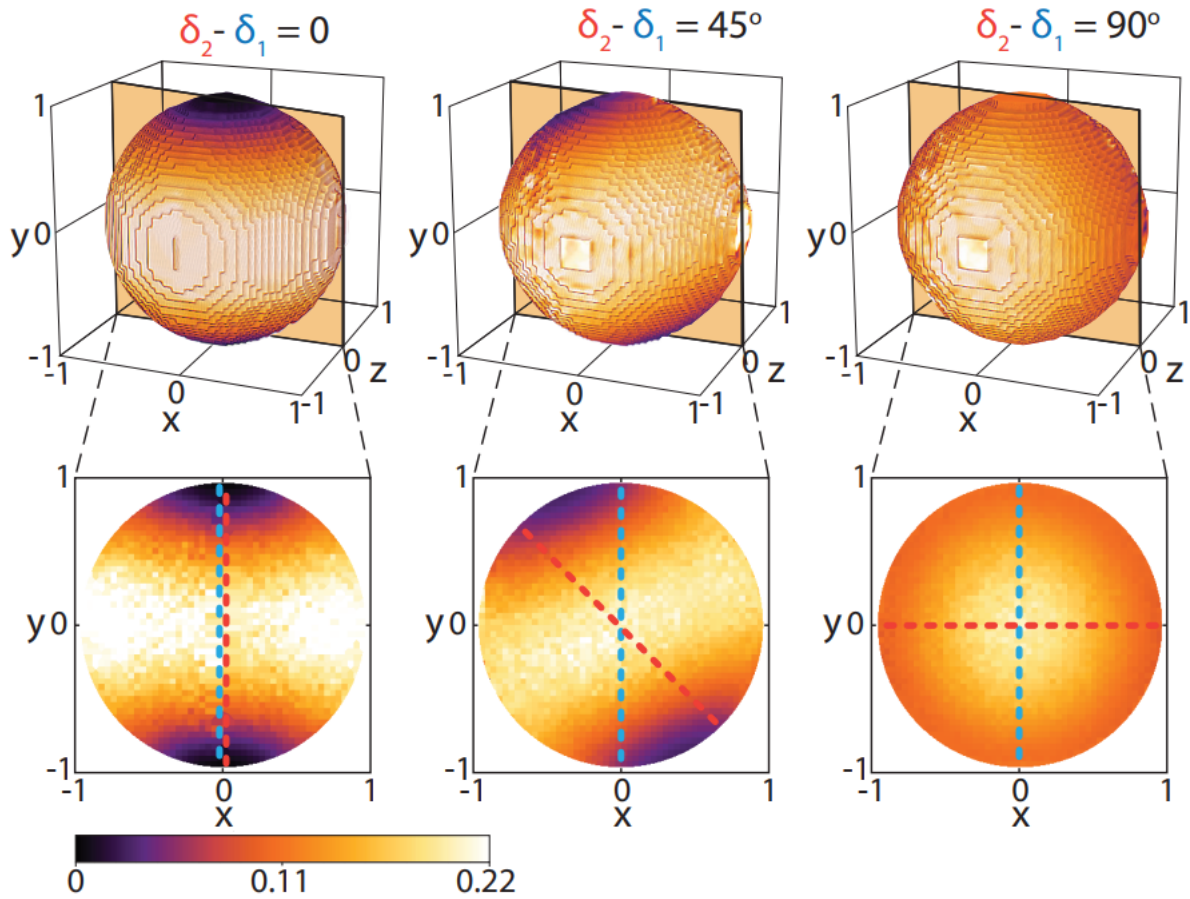
Figure 5.7: **Magnitude of the quantum back-action.** The amplitude of disturbance induced by measurement $(\mathrm{Tr}[d\rho^\dagger d\rho])$ at each point on the Bloch sphere for three relative measurement angles $\delta_2 - \delta_1$. (top) Disturbance at the surface of the Bloch sphere, which is bounded from below by the uncertainty principle. (bottom) A slice through the $z = 0$ plane. Dotted lines mark measurement axes. For non-commuting observables $(\delta_2 - \delta_1 \neq 0)$, no point of zero disturbance exists within the entire volume of the Bloch sphere, indicating that the state must diffuse indefinitely.

quantitatively with this prediction.

The diffusive behavior seen in Fig. 5.6c suggests that even once the probability distributions have reached steady state, the system continues to evolve. To quantify this measurement-induced disturbance, we plot the norm of the state change $d\rho$ over an interval of 64 ns versus position on the Bloch sphere in Fig. 5.7. When the measurements are compatible, two points on the Bloch sphere exhibit zero disturbance, which indicate mutual eigenstates of the measurements and hence points of collapse. No such points exist when the measurements are incompatible, implying that the state diffuses indefinitely. If we calculate

this disturbance directly from the stochastic master equation for the system, we find the following relation valid for all initially pure states

$$\text{Tr}[d\rho^\dagger d\rho] = (\Delta\sigma_{\delta,1}^2 \Gamma_1 \eta_1 + \Delta\sigma_{\delta,2}^2 \Gamma_2 \eta_2)dt \tag{5.25}$$
$$\geq |\langle[\sigma_{\delta_1}, \sigma_{\delta_2}]\rangle|\sqrt{\eta_1\eta_2\Gamma_1\Gamma_2}dt$$

which holds for any Hermitian operators of any system. The right hand side of the equality closely resembles the original Heisenberg uncertainty relation, but contains the sum of the variances instead of the product. Unlike the latter, the sum can be bounded by a stronger inequality which is never trivial for non-commuting measurements[73]. This shows that the persistent diffusion observed in Fig. 5.6c is a universal consequence of the uncertainty principle and can be quantitatively derived from it.

There are several interesting applications of our work that we have not yet explored. First, with a modest improvement in quantum efficiency, one could purify a quantum state faster with non-commuting observables than with a single measurement[113, 25]. The ability to purify a state faster using quantum trajectories is a fundamental prediction of the field that has yet to be realized[57]. This enhancement scales favorably with system size, suggesting an application for MQMs as well.

Second, existing work on quantum foundations has focused on testing the validity of essential features of quantum mechanics, such as contextuality and the various bounds on error and disturbance. Our work presents the possibility of exploring the consequences such concepts in realistic systems[97, 72], which tend to interact continuously with their environments via multiple non-commuting decoherence channels. Error-disturbance bounds would provide a natural starting point, as we have separately extracted state information in Fig. 5.5b and quantified state disturbance in Fig. 5.7, but have not yet investigated the intrinsic trade-off between these concepts.

5.4 The Polaron Transformation

Eq. (5.13) largely justifies thinking of SQMs as measurements. It couples the qubit to the cavity via a Hermitian operator, and we saw a very similar Hamiltonian lead to a QND measurement in the dispersive case. To go the rest of the way, we need to treat the cavity as an open quantum system and derive the stochastic master equation that we relied upon in the previous section. The master equation for a cavity which is damped at rate κ and monitored with quantum efficiency η is

$$d\rho = -i[H_{\text{q-frame}}, \rho]dt + \kappa\mathcal{D}[a]\rho\,dt + \sqrt{\kappa\eta}\mathcal{H}[ae^{i\phi}]\rho\,dW. \tag{5.26}$$

As before, we treat the cavity modes one at a time. If the cavity starts in the ground state when this Hamiltonian is turned on, then it remains in a coherent state at all times. To derive an effective master equation for the qubit, we must eliminate the cavity entirely from

Eq. (5.26). This may be accomplished by first applying a transformation which displaces
the cavity to its ground state, which is known as the polaron transformation[36]

$$U = |e_\delta\rangle\langle e_\delta|D[\alpha_e] + |g_\delta\rangle\langle g_\delta|D[\alpha_g] \tag{5.27}$$

$$\dot{\alpha}_{e/g} = -i\epsilon_{e/g} - \frac{\kappa}{2}\alpha_{e/g} \qquad \epsilon_{e/g} = \pm\tilde{g}$$

where D is the cavity displacement operator Eq. (1.23), $|e_\delta/g_\delta\rangle$ are the eigenstates of σ_δ and
$\alpha_{e/g}$ is the cavity displacement conditional on the qubit state. Because the second line is the
classical equation of motion for a resonantly driven cavity with drive rate ϵ and damping κ,
this transformation maps the cavity to its ground state at all times, allowing us to trace it
out. Transforming back from the polaron frame to the qubit frame, we find that the effective
stochastic master equation for the qubit is

$$d\rho = \mathcal{L}\rho\, dt + \frac{\sqrt{\Gamma_M}}{2}\mathcal{H}[\sigma_\delta]\rho\, dW - i\frac{\sqrt{\Gamma'_M}}{2}[\sigma_\delta, \rho]dW \tag{5.28}$$

$$\mathcal{L}\rho = -i\frac{B}{2}[\sigma_\delta, \rho] + \frac{\Gamma}{2}\mathcal{D}[\sigma_\delta]\rho$$

$$B = 2\tilde{g}\mathrm{Re}(\alpha_e + \alpha_g) \qquad \Gamma = -2\tilde{g}\mathrm{Im}(\alpha_e - \alpha_g)$$

$$\sqrt{\Gamma_M} = \sqrt{\kappa\eta}|\beta|\cos(\phi - \theta_\beta) \qquad \sqrt{\Gamma'_M} = \sqrt{\kappa\eta}|\beta|\sin(\phi - \theta_\beta)$$

$$\beta = \alpha_e - \alpha_g$$

where $\theta_\beta = \arg(\beta)$ is the angle of the cavity displacement axis β in the IQ plane and Γ is
the dephasing rate. Substituting the equations of motion for the cavity Eq. (5.27) into the
remaining expressions of Eq. (5.28), we find

$$d\rho = \frac{4\tilde{g}^2}{\kappa}(1 - e^{-\kappa t/2})\mathcal{D}[\sigma_\delta]\rho\, dt + \sqrt{\frac{4\tilde{g}^2}{\kappa}}\eta(1 - e^{-\kappa t/2})\mathcal{H}[e^{i\phi}\sigma_\delta]\rho\, dW \tag{5.29}$$

$$\implies \Gamma = \frac{8\tilde{g}^2}{\kappa} = \frac{2\chi^2\bar{n}_0}{\kappa}.$$

We align the LJPA amplification axis to the axis of displacement arising in Eq. (5.13), so that
$\phi = \theta_\beta$. Application of sidebands to another mode of the cavity does not change the above
derivation except to add an additional measurement also modeled by Eq. (5.29). Neglecting
time-scales of order $1/\kappa$, measurement of two observables σ_{δ_1} and σ_{δ_2} is modeled by

$$d\rho = \sum_{i=\{1,2\}} \frac{\Gamma_i}{2}\mathcal{D}[\sigma_{\delta_i}]\rho\, dt + \sqrt{\frac{\Gamma_i\eta_i}{2}}\mathcal{H}[\sigma_{\delta_i}]\rho\, dW \tag{5.30}$$

$$r_i dt = \langle\sigma_{\delta_i}\rangle dt + \frac{dW_i}{\sqrt{2\eta_i\Gamma_i}}$$

where r_i is the measurement signal at time t, normalized appropriately. Equation (5.30) can
also be converted to a Fokker Planck equation, which propagates probability distributions of

ρ. We use the latter to generate theory plots for Fig. 5.6. This stochastic master equation is generated by the following measurement operator

$$\Omega_{r_1(t),r_2(t)} = \exp\left[\sum_{i=1,2} -\frac{\Gamma_i \eta_i}{2} \left(r_i(t) - \sigma_{\delta,i}\right)^2 dt\right] \tag{5.31}$$

$$\rho(t+dt) = \mathcal{E}_{1-\eta_i} \frac{\Omega\rho(t)\Omega^\dagger}{\mathrm{Tr}[\Omega\rho(t)\Omega^\dagger]}$$

where $\mathcal{E}_{1-\eta_i}$ is a superoperator which models dephasing due to finite quantum efficiency. To assure positivity of the state when dt is taken to be finite, we use Eq. (5.31) to numerically propagate quantum trajectories and also to calculate probabilities for maximum likelihood reconstruction.

It is natural to ask if the above calculation generalizes to multi-quadrature measurements. It turns out that for systems larger than a qubit, the polaron transformation does not necessarily lead to a Markovian master equation[27]. Nevertheless, the methods of chapter 3, section 3.4 can be used to show that the effective dynamics still damp off-diagonal elements and project into an eigenstate of X. Thus we can still use MQM measurements for quantum trajectories and feedback.

Chapter 6

Adaptive Measurements and the Canonical Phase Measurement

Much of modern metrology and communication technology encodes information in electromagnetic waves, typically as an amplitude or phase. While current hardware can perform near-ideal measurements of photon number or field amplitude, to date no device exists that can even in principle perform an ideal phase measurement, as we noted in chapter 2. However phase is no less fundamental than amplitude and power, and underlies a huge range of technological applications, from FM radio and optical communication to metrological interferometry such as gravitational wave detection and even long-baseline telescopes.

In this chapter, we describe the implementation of a single-shot canonical phase measurement on a one-photon wave packet. As defined below, the canonical phase measurement defines the quantum-mechanically ideal POVM for measurement of phase. The measurement surpasses the standard method of heterodyne detection and is optimal for single-shot phase estimation. Our system adaptively changes the measurement basis of a parametric amplifier (operated in phase-sensitive mode) during photon arrival. Changing the measurement basis during the measurement process constitutes a form of quantum feedback, and provides much more control to optimize sensitivity to phase. We use a superconducting qubit to emit our photon signal, which allows us to validate the detector's performance by tracking the quantum state of the source. This gives us an opportunity to put to practical use the physics of quantum trajectories.

These results are part of a broader effort to understand and demonstrate the full capabilities of adaptive measurement. In chapter 4, we saw many examples in which quantum feedback offered fundamentally enhanced control over state preparation as compared to what can be achieved with measurement alone. In these cases, we applied feedback operations to the device that we wished to control, such as a register of qubits. The canonical phase measurement gives an example in which quantum feedback fundamentally enhances the capabilities of a measurement device. In particular, our use of feedback converts a standard JPA (with a TWPA as backup) into a theoretically ideal phase detector. In contrast to control applications, the protocol works by changing the measurement basis (and hence POVM)

in response to the incoming signal. In the broadest sense, we can think of adaptive measurement as a kind of 'POVM-generating machine,' which widens the range of observables that may be accessed by a detector.[1] The following work demonstrates that quantum feedback can both enhance the precision of a detector and enable it to measure new classes of physical observables. Excitingly, the full theoretical capabilities of adaptive measurement are not yet known.

Even the definition of phase as an observable in quantum mechanics is surprisingly subtle, and has lead to a protracted debate that took decades to resolve[70, 71, 130, 96]. In chapter 2, we saw the definition of a somewhat suspicious phase operator

$$\hat{\varphi} = \int_0^{2\pi} \varphi |\varphi\rangle\langle\varphi| \tag{6.1}$$

$$|\varphi\rangle = \frac{1}{\sqrt{2\pi}} \sum_{n=-\infty}^{\infty} e^{i\varphi n} |n\rangle.$$

As φ is 2π periodic, the limits of the integration could just as well have been chose as $-\pi$ and π or any other 2π interval, and therefore $\hat{\varphi}$ has no unique definition. Fortunately, this ambiguity posed no problem for the physically relevant operators of that section, such as $\cos(\hat{\varphi})$, but it does raise some subtleties in the discussion of phase measurements.

Things get even more complicated when one considers a truncated version of the above Hilbert space, such as a harmonic oscillator. This case is perhaps the most important from an applied standpoint, as it models the phase of an electromagnetic wave. As the lower bound on n is now zero instead of $-\infty$, $|\varphi\rangle$ no longer forms an orthonormal basis. As the eigenbasis of a Hermitian operator is necessarily orthonormal, this implies that there is no Hermitian operator corresponding to the phase basis $|\varphi\rangle$, and thus the standard measurement formalism cannot apply here.

Fortunately, the POVM formalism introduced in chapter 3 provides the generalization of measurement that naturally works with phase. The set of Kraus operators

$$\Omega_\phi = |\phi\rangle\langle\phi| \tag{6.2}$$

$$|\phi\rangle \equiv \frac{1}{\sqrt{2\pi}} \sum_{n=0}^{\infty} e^{i\phi n} |n\rangle$$

satisfy the required normalization condition

$$\int_0^{2\pi} \Omega^\dagger \Omega d\phi = \langle\phi|\phi\rangle \int_0^{2\pi} |\phi\rangle\langle\phi| d\phi = 2\pi \langle\phi|\phi\rangle I. \tag{6.3}$$

This relation also holds if we further truncate the Hilbert space down to a qudit, meaning that it defines a phase measurement on any finite system as well.

---

[1]Thanks to John Preskill for coining this term.

There are several appealing features of the canonical phase basis. Firstly, if a signal is encoded to a state by applying phase shifts, the canonical phase measurement is optimal when no prior information on the phase is available, as is the case in communication[49]. Secondly, although for many reasons we cannot write an expression like $[\hat{\phi}, \hat{n}] = i$, there is a clear though unique sense in which ϕ is canonically conjugate to $\hat{n}$. Phase is a 2π periodic variable, so the standard definition of variance does not make sense if we wish to write down a Heisenberg uncertainty relation like $\Delta\phi^2 \Delta n^2 \geq 1/4$. However there is a generalization of variance that is compatible with periodic variables, called the Holevo variance

$$\Delta\phi^2 \equiv |\langle e^{i\phi}\rangle|^{-2} - 1. \tag{6.4}$$

For a uniform distribution over ϕ, $\Delta\phi^2 = \infty$. Furthermore in the limit of a narrowly peaked distribution, Eq. (6.4) reproduces the standard definition of variance. Most importantly, if we use this definition of variance, then the above uncertainty is principle exact[11]. Although ϕ is not an operator, ϕ labels the measurement outcome Ω_ϕ, and thus the number-phase uncertainty relation is exactly analogous to the Heisenberg uncertainty relation. Intuitively, this implies a bizarre property of the canonical phase measurement, which is that it is completely insensitive to the incident power. An ideal phase detector would be fundamentally incapable of distinguishing the vacuum from a bright, phase-incoherent light, but it could optimally detect the coherent phase of a laser beam.

Canonical phase measurements fit into the broader context of quantum phase estimation[**103**]. Here we focus on performing an optimal *single-shot* phase measurement, meaning that given a single copy of a state to which we have applied some phase shift $e^{i\phi a^\dagger a}$, we wish to optimally estimate ϕ. This formulation underlies more intricate versions of the phase estimation problem, such as estimating ϕ given multiple copies of the same state, or given the ability to choose the input state to the phase shift operation. The latter enables Heisenberg-limited scaling of the measurement uncertainty, for example by applying $e^{i\phi a^\dagger a}$ to multiple subsystems of an entangled state[40].

In the absence of an instrument capable of implementing a canonical phase measurement, heterodyne detection, in which one measures a rapidly varying quadrature of the input, serves as the standard technique for estimating the phase of an unknown signal. Several schemes have surpassed heterodyne detection [154, 144, 4, 54], however these protocols also acquire undesired photon number information, and thus cannot reach the quantum limit or implement a canonical phase measurement. In this work, we implement a feedback-controlled quantum-limited amplifier which dynamically updates its amplitude measurement in response to the incident field. When the system continuously optimizes this measurement basis for phase sensitivity, it implements a canonical phase measurement on an incoming single-microwave-photon state[146]. We verify implementation of a canonical phase measurement using the entanglement between the emitted photon and its source, which allows us to confirm that acquisition of photon number information is suppressed. The system surpasses heterodyne detection by $15 \pm 2\%$. The first section of this chapter present the experimental work, and the second section derives the theory behind the experiment.

These results are part of a broader effort to understand quantum feedback applied to adaptive measurement, which is still an exploratory direction of research even theoretically. Given a set of POVMs that may be implemented by a given detector, it is not known how to construct the set of POVMs that may be implemented by adaptive measurement. In the following work, we start with the set of POVMs generated by infinitely short homodyne measurements in an arbitrary basis and generate the canonical phase measurement. The result is one of many examples given in this thesis in which quantum feedback enhances experimental capabilities in the presence of linearity constraints. In deterministic remote entanglement generation, we required quantum feedback to circumvent the impossibility of a linear optics Bell measurement[16]. We speculate that there are many further examples in which quantum feedback can improve a linear optics experiment or other system with severe constraints on the set of POVMs that may be implemented. Conversely, when strong nonlinearities are already available, the potential gains of adaptive measurement narrow substantially. A formal piece of evidence for this suggestion is Neumark's dilation theorem, which shows that *any* POVM may be implemented by applying a unitary operation on the system plus some auxiliary system and then applying a projective measurement on the latter, as in Eq. (3.5). This implies that there exists a digital quantum circuit that can implement any POVM, circumventing any possible need for adaptive measurement. This result certainly does not obviate the need for adaptive measurement. Quantum computers are a long way off, and it will not necessarily be possible to load the quantum state of the system that we wish to measure into it.

6.1 Implementation and Verification of a Canonical Phase Measurement

As shown in Fig. 6.1a, our system consists of a transmitter, which encodes a variable Θ_{true} into the phase of a single-photon electromagnetic signal, and a receiver, which uses a continuous feedback protocol to guess this phase in a single shot using an adaptive feedback protocol. A superconducting transmon qubit[62] embedded in a 3D aluminum cavity acts as the transmitter. We use coherent bath engineering [91] of this artificial atom to generate our photonic state, which yields more process control than direct spontaneous decay. To implement this scheme, we Rabi drive our qubit at $\Omega_{\text{R}}/2\pi = 20$ MHz, which creates an effective low-frequency qubit. Simultaneously, we apply a cavity sideband at $\omega_{\text{cav.}} + \Omega_{\text{R}}$, where $\omega_{\text{cav.}}$ is the cavity resonance frequency. As shown in Fig. 6.1b, the sideband drives a transition from the $|+, 0\rangle$ state to $|-, 1\rangle$ state, where $|\pm\rangle \equiv (|e\rangle \pm i|g\rangle)/\sqrt{2}$ are the dressed states of qubit under driving and $0, 1$ count the number of photons in the cavity. The cavity then decays, emitting a photon and leaving the system in the $|-, 0\rangle$ state, which is not affected by the sideband. We ensure that the cavity decay rate is fast compared to the sideband-induced coupling, so that the qubit's effective decay rate from $|+\rangle$ to $|-\rangle$ is limited by the sideband amplitude. By modulating the sideband amplitude during photon emission,
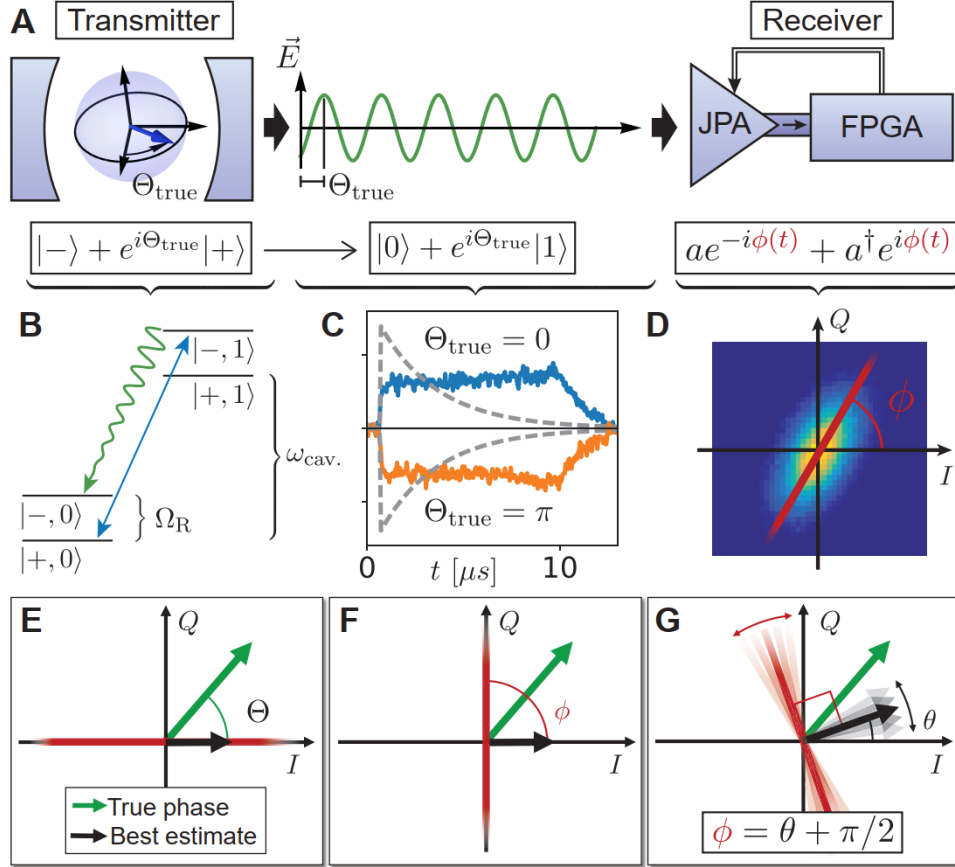
Figure 6.1:   Experimental implementation. (a) Atom in a cavity, with phase Θ_{true} encoded
into its dipole moment. The atom decays and emits a photon into a 1D waveguide with
phase encoded into the electric field as shown. The JPA receives the photon and measures
an amplitude quadrature selected by the FGPA. (b) Sideband cooling scheme to emit photon.
Sideband converts a qubit excitation to a cavity excitation, which is then emitted as a single
photon at the cavity frequency. (c) Measured mode shape (E-field envelope) of emitted
photon. Dashed line shows mode shape if constant cooling rate were used instead. (d)
Output of JPA. Signal is amplified along measurement axis ϕ and squeezed along the other.
(e-g) Estimating and tracking state by changing measurement basis. Receiver attempts to
maintain the phase measurement condition $\phi = \theta + \pi/2$. See text for details.

we generate a photon with a flat modeshape (Fig. 6.1c), which greatly ameliorates the
detrimental effects of feedback delay at the receiver[28]. The details of mode shape control
are derived in the following section. To encode the phase Θ_{true}, we prepare the qubit in a
superposition state of the form $(|-\rangle + e^{i\Theta_{\text{true}}}|+\rangle)/\sqrt{2}$, which decays by emitting the photonic
state $(|0\rangle + e^{i\Theta_{\text{true}}}|1\rangle)/\sqrt{2}$.

Our receiver consists of a Josephson parametric amplifier (JPA) (like that of chapter 2
pumped at twice its resonance frequency by a field-programmable gate array (FPGA), which
serves as a classical feedback controller (Fig 6.1a). To maintain high measurement bandwidth
for quantum feedback, we operate the JPA at a relatively low gain of 6 dB, which yields a
gain bandwidth of 45 MHz, and follow it with a traveling wave parametric amplifier[74] (not
shown) to boost the signal strength and maintain a quantum efficiency of $\eta = 0.4$. The JPA
measures field amplitude via the quantum mechanical quadrature operator $ae^{-i\phi(t)} + a^\dagger e^{i\phi(t)}$,
where a is the quantum mechanical annihilation operator of the incident field and $\phi(t)$ is the
instantaneous phase of the parametric pump.

To perform a canonical phase measurement on the incident field, the feedback controller
continuously adapts the measurement axis $\phi(t)$ as the photon arrives at the receiver[146].
The measurement axis is chosen to maximize the acquisition of phase information as follows.
Before the photon reaches the JPA, the receiver has no information and therefore chooses
ϕ arbitrarily. Upon arrival of a portion of the photon, the JPA detects a small positive (or
negative) fluctuation, which then informs the system that the true phase is likely oriented
along (or opposite) the measurement axis (Fig. 6.1e). At this point, any further measurement
in this basis interrogates the amplitude of the incident field and thus yields undesired photon
number information. Ideally, the system would then rotate the measurement axis by 90
degrees (Fig. 6.1f), so that a small deviation between the current best estimate of the phase
$\theta(t)$ and the true phase Θ_{true} would be detectable as a positive or negative fluctuation in the
signal. As the photon continues to arrive, the feedback controller gains more information
and updates the phase $\phi(t)$ to maximize sensitivity to phase (Fig. 6.1g). If the phase
measurement condition $\phi(t) = \theta(t) + \pi/2$ is maintained at all times, then the system acquires
no photon number information and implements a canonical phase measurement.

To track the best estimate of the phase, the feedback controller must continuously update
its best guess of the atom's state based on the measurement signal starting with no prior
information *i.e.* it should track the quantum trajectories of the system[90, 17] given an
initially maximally mixed state. We begin by observing and verifying quantum trajectories
for homodyne ($\phi(t) = 0$) and heterodyne ($\phi(t) = \omega_{\text{het.}}t$, $\omega_{\text{het.}}/2\pi = 0.5$ MHz) detection[17].
Example trajectories are plotted in Fig. 6.2a,b and tomographically validated in [80]. These
data allow us to characterize measurement back-action and check consistency with theory.
The stochastic component of the back-action always lies in the plane of the instantaneous
measurement basis, as is clear from the homodyne data.

The presence of back-action not only governs how to adapt the measurement axis $\phi(t)$, but
also offers a method to independently validate the receiver's implementation of a canonical
phase measurement. Because an ideal phase measurement acquires maximal phase infor-
mation and no photon-number information, it maximally disturbs the atomic dipole phase
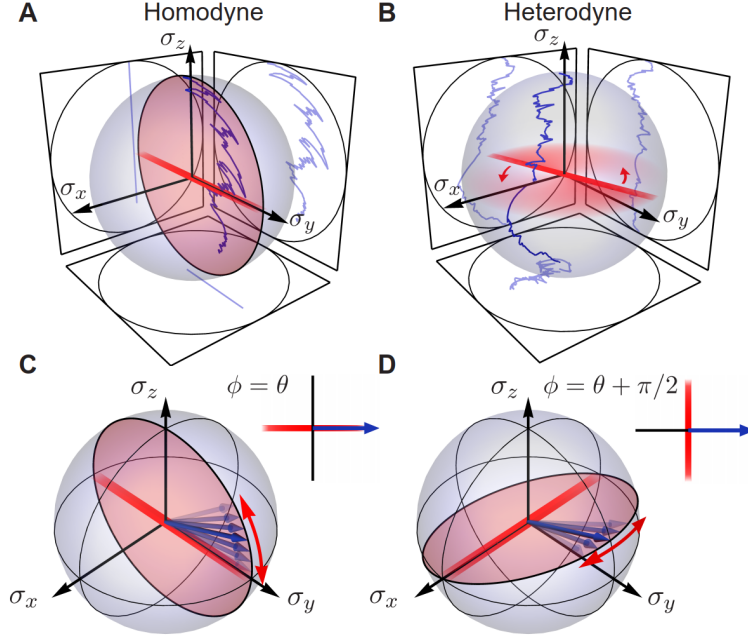
Figure 6.2: Measurement back-action and quantum trajectories. Coordinate axes are chosen so that the atom decays to $\sigma_z = -1$. (a) A single homodyne quantum trajectory ($\phi(t) = 0$). State only propagates in the plane of the measurement axis. (b) A single heterodyne trajectory ($\phi(t) = \omega_{\mathrm{het.}}t$). The qubit is initialized in $|+\rangle$ for both trajectories. (c) Amplitude back-action, which occurs when the measurement axis (red line) is aligned to the best estimate of the state (blue arrow). (d) Phase back-action, which occurs when the phase measurement condition is satisfied.

while minimally disturbing the atomic excitation probability. This effect is directly visible in the quantum trajectories, as illustrated conceptually in Fig. 6.2c and d. When the measurement axis is aligned with the best estimate of the phase ($\phi = \theta$), the resulting acquisition of amplitude information manifests as a random disturbance of the qubit state along the axis of decay (Fig. 6.2c). Conversely, when the phase measurement condition is satisfied ($\phi = \theta + \pi/2$), then only the phase of the qubit state is subject to noise (Fig. 6.2d). In this way, we can verify the performance of our receiver by characterizing the dynamics of the transmitter. This capability is uniquely quantum, and arises from entanglement between the atom and its emitted photon.

We show the results of this verification scheme in Fig. 6.3. Fig. 6.3a shows a single quantum trajectory under adaptivedyne detection, in which $\phi(t)$ is continuously adapted by the feedback controller. Fig. 6.3b shows the difference between the ideal quadrature phase and the measured phase, which shows that the feedback controller approximately maintains the phase measurement condition. To interpret the dynamics, we plot the ensemble statistics of the phase back-action as a function of time in Fig. 6.3d, with the heterodyne detection case
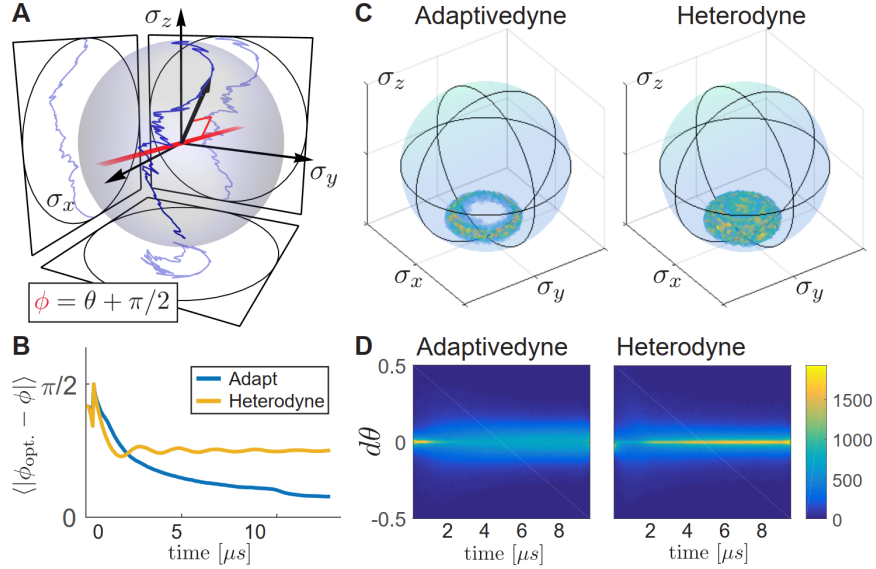
Figure 6.3: Back-action and measurement validation. (a) A single adaptive-dyne quantum trajectory. The red right-angle bracket emphasizes orthogonality between the measurement axis and the state. (b) Quality of tracking for heterodyne and adaptivedyne, where $\phi_{\text{opt.}} = \theta(t) + \pi/2$. Adaptivedyne significantly outperforms the heterodyne and comes close to the ideal phase by $T = 13\mu s$. The difference $\phi_{\text{opt.}} - \phi$ is cut to lie on the interval $[-\pi/2, \pi/2]$. (c) Distribution of trajectories at $t = 10\mu s$. Due to suppression of photon-number back-action, adaptivedyne trajectories cluster in a ring at late times. (d) Statistics of the phase back-action $d\theta$ for adaptivedyne and heterodyne. On average, the phase back-action is significantly larger for adaptivedyne.

included for comparison. It can be seen that the phase back-action $d\theta$ is significantly larger for adaptivedyne detection. Fig. 6.3c shows the ensemble statistics of the state at $t = 10$ μs. As observed in [17], the quantum trajectories of a decaying atom evolve on a spherical shell that shrinks deterministically to the south pole of the Bloch sphere. Due to the suppression of back-action along the decay axis, adaptivedyne trajectories are further confined, exhibiting something closer to a ring-like structure. This feature presents an unambiguous signal that our system approximately implements a canonical phase measurement.

A canonical phase measurement should outperform heterodyne detection in estimating the phase Θ_{true}. To verify superior performance, we prepare our qubit in one of 8 equally spaced points along the equator of the Bloch sphere. From each shot, the receiver optimally[28] estimates the phase of the photon by computing the following quantity

$$R = \int_0^T e^{i\phi(t)} \sqrt{u(t)} r(t) dt \tag{6.5}$$

where $u(t)$ is the photon mode shape, and T is the duration of each experimental run and
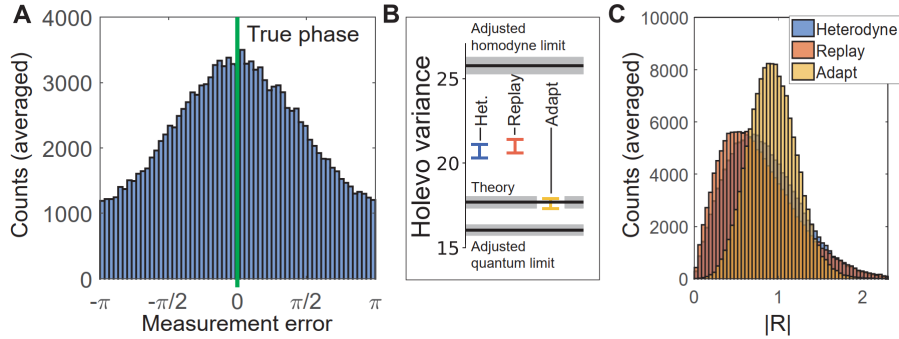
Figure 6.4: Phase-estimation performance. (a) Histogram of the difference between the measurement outcome and the true phase *i.e.* $\theta(T) - \Theta_{\text{true}}$ (b) Performance is evaluated by computing the Holevo variance of this distribution. Quantum limit (bottom black line) homodyne limit (top black line) and absolute theory prediction based on feedback delay are inferred from the performance of heterodyne, with corresponding error bars shown as gray rectangles. (c) Distribution of the amplitude information. The distribution is significantly narrower for adaptivedyne, indicating suppression of this information channel.

$r(t)$ is the measurement signal read out from the JPA normalized such that its variance is dt. The best estimate of the photon's phase in a single shot is given by the complex argument $\theta(T) = \arg(R)$. Fig. 6.4a plots a histograms of this best estimate for adaptivedyne detection, which exhibits the $\cos(\theta - \Theta_{\text{true}})$ dependence expected theoretically[28].

We compare the performance of adaptivedyne and heterodyne detection by plotting the Holevo variance of each underlying distribution in Fig. 6.4b. We also include data for what we term replay detection, in which $\phi(t)$ from an adaptivedyne shot of the experiment is replayed instead of feeding back based on the current signal. In this way, we can confirm that it is the correlations between $\phi(t)$ and the state that yield enhanced performance, rather than the independent statistics of $\phi(t)$. For additional confirmation, we independently measure the signal-to-noise ratio of our amplifier chain for heterodyne and adaptivedyne detection and verify that it remains the same to well within 1%[80]. Heterodyne and replay perform equally well, and are both significantly surpassed by a canonical phase measurement implemented via adaptivedyne detection. Adaptivedyne does not reach the quantum limited Holevo variance of 3 due to a combination of loss, qubit decoherence and feedback delay. However from our heterodyne data we infer an adjusted quantum limit given our quantum efficiency and purity of the emitted photon, as well as the hypothetical homodyne limit. The canonical phase measurement comes significantly closer to this adjusted quantum limit than any other scheme, limited almost entirely by feedback delay.

We infer the sensitivity of each scheme to photon-number information from the distributions of $|R|$, which are shown in Fig. 6.4. The distributions for heterodyne and replay are almost identical, while the adaptivedyne histogram is substantially narrower, indicating that the latter is less sensitive to this undesired information[28].

The above experiment gives an important example of the power of quantum feedback applied to adaptive measurement. There are a number of tangible, immediate applications to quantum information. Firstly, the implementation of quantum feedback on a detector is known to allow enhanced readout of superconducting circuits[116]. Furthermore, the ability to perform a canonical phase measurement enables linear-optics preparation of the $|0\rangle + |1\rangle$ photonic state, which is a major experimental challenge of single-rail linear optics quantum computing[107]. More broadly, it is known that adaptive measurements are universal[100], meaning that many relatively simple measurement devices augmented with quantum feedback can perform any measurement allowed by quantum mechanics. Thus our extension of a standard amplitude measurement device to a theoretically ideal phase measurement represents one of potentially many possible directions for future work in this field.

6.2 Adaptive Measurements and Quantum Feedback

We now derive the feedback protocol that enables the implementation of a canonical phase measurement in the $|0\rangle$, $|1\rangle$ manifold by solving the stochastic master equation for the source. Before diving into the math, we outline where we are going with this calculation and what we hope to extract from it. A general photonic state with a well-defined phase in the subspace considered can be written as

$$|\psi\rangle = \frac{1}{\sqrt{2}} \left[|\vec{0}\rangle + e^{i\Theta} \int_{-\infty}^{t} \sqrt{u(s)} a^{\dagger}[s] |\vec{0}\rangle \right] \tag{6.6}$$

where $|\vec{0}\rangle$ denotes the electromagnetic vacuum and $u(t)$ is a mode shape that square-integrates to 1. In solving the stochastic master equation for the artificial atom with a time-dependent cooling rate $\gamma(t)$, we will derive the atomic excitation probability as a function of time, which fixes $u(t)$ by conservation of energy. We can also use a solution to the master equation to determine the POVM implemented by the experiment. Given a measurement record $r(t)$ occurring experimentally, the unnormalized stochastic Schrödinger equation produces an unnormalized final state $|\tilde{\psi}(\infty)\rangle$. We can use this to compute probability for $r(t)$ to occur given any arbitrary initial state, which in turn lets us determine the measurement operator $E[r] \equiv \Omega_r^{\dagger}\Omega$ associated to $r(t)$

$$P[r(t)|\psi(0)] = \langle\tilde{\psi}(\infty)|\tilde{\psi}(\infty)\rangle = \text{Tr}(E[r(t)]|\psi(0)\rangle\langle\psi(0)|) \tag{6.7}$$

By matching terms for different initial states $|\psi(0)\rangle$, we can compute $E[r(t)]$ and confirm that adaptive measurement leads to the canonical phase measurement. In principle we can also use the solution to verify quantum trajectories, though in practice we use numerics to take decoherence and finite measurement efficiency into account.

To solve for the dynamics under adaptive measurement, we use the full master equation to derive the required $\gamma(t)$, which also lets us calculate the optimal feedback strategy for a given photon mode shape. Many of these results are also derived in [28], but we include a variant

of that derivation here for completeness. As photon loss does not affect the decay dynamics or the best estimate of the phase, we assume $\eta = 1$ for this analysis. We also neglect other forms of decoherence, which have a negligible effect on the decay dynamics. This allows us to perform our computations with a pure state, so we begin with the unnormalized stochastic Schrödinger equation for an atom observed via homodyne detection[147], which provides a state update from the acquisition of an infinitesimal amount of information via homodyne detection of the atom's spontaneous emission

$$\frac{d}{dt}|\tilde{\psi}\rangle = \left[-\frac{1}{2}\gamma(t)\sigma^\dagger \sigma + \sqrt{\gamma(t)}e^{-i\phi(t)}\sigma r(t) \right]|\tilde{\psi}\rangle. \tag{6.8}$$

$r(t)$ is the measurement record, $\sigma = |-\rangle\langle+|$ and $|\tilde{\psi}\rangle$ is the unnormalized pure state describing the state of the atom[2]. As in most references, we have dropped the normalization prefactor $\sqrt{P(dW)}$ that we included in Eq. (3.24). This normalization factor turns out to be unimportant for reasons that will be clarified later. If we write $|\tilde{\psi}\rangle$ as $|\tilde{\psi}\rangle = c_-|-\rangle + c_+|+\rangle$, then the equation of motion for c_+ is

$$\frac{dc_+}{dt} = -\frac{1}{2}\gamma(t)c_+ \quad \rightarrow \quad c_+(t) = c_+(0)e^{-\frac{1}{2}\int_0^t \gamma(s)ds} \tag{6.9}$$

where we have assumed that $\gamma(t < 0) = 0$. Recall that our system decays from $|+\rangle$ to $|-\rangle$. Although in general Eq. (6.8) does not preserve the norm of $|\tilde{\psi}\rangle$, one nevertheless derives the correct equation of motion for the average population from the above in the absence of measurement. The result is $d|c_+|^2/dt = \gamma(t)|c_+|^2$, which coincides with the expectation based on a standard rate equation for decay of the excited state population. We identify the mode shape with the instantaneous emitted intensity, assuming for a moment that the atom was initialized with $c_+ = 1$, $c_- = 0$

$$u(t) \equiv \gamma(t)|c_+|^2 = \gamma(t)e^{-\int_0^t \gamma(s)ds}. \tag{6.10}$$

Notice that $u(t)$ integrates to 1 for any $\gamma(t)$. If we demand a flat mode shape so that $u(t)$ is constant, then $\gamma(t) = 1/(\tau - t)$, where $\tau = 10\mu s$ parameterizes the photon's duration. As $\gamma(t)$ diverges at $t = \tau$, we set a maximum cooling rate of 1.4 MHz and cool at this maximum rate for several microseconds longer than τ, such that more than 99% of the excited state population has decayed by $T = 13\mu s$. The constant γ portion when $\gamma/2\pi = 1.4$ MHz coincides with the portion of the photon that decays exponentially, as can be seen in 6.1C.

Now that we have developed the necessary tools for emitting a flat photon, we derive the optimal feedback protocol given our photon. The equations of motion for c_- determine the best estimate of the phase

$$\frac{dc_-}{dt} = c_+\sqrt{\gamma(t)}e^{-i\phi(t)}r(t) \quad \rightarrow \quad c_-(t) = c_-(0) + c_+(0)\int_0^t e^{-i\phi(s)}\sqrt{u(s)}r(s)ds. \tag{6.11}$$

---

[2]It is a slight abuse of notation to write Eq. (6.8) as a differential equation, as $r(t)$ is actually an unbounded stochastic quantity, but this convention is nevertheless common in the literature.

Notice the similarity between Eq. (6.11) and Eq. (6.5) of the previous section. For feedback, we wish to compute the best estimate of the atomic dipole phase at time t assuming that the controller initially has no information about the phase. This best estimate coincides with the best estimate for the phase of the emitted photon after that time. To compute it, we note that the dynamics are trivial if the system is initialized in $|-\rangle$, so that the dipole phase evolution of the zero-knowledge mixed state $\rho_0 = (|-\rangle\langle-| + |+\rangle\langle+|)/2$ is entirely determined by the dynamics of the second term. Again taking $c_+ = 1$, $c_- = 0$, the dipole phase is given by the relative complex phase between c_+ and c_-^*. The complex phase of c_+ remains constant, so the dipole moment phase is simply

$$\theta(t) = \arg(R), \quad R(t) \equiv \int_0^t e^{i\phi(s)} \sqrt{u(s)} r(s) ds \tag{6.12}$$

in agreement with Eq. (6.5). In principle, Eq. (6.5) and the phase measurement condition $\phi(t) = \theta(t) + \pi/2$ define the optimal protocol. For ease of implementation, this protocol may be further simplified by solving for the absolute value and complex argument of R individually as follows. If the controller maintains the phase measurement condition, then we have $\exp(i\phi(s)) = iR/|R|$. Making this substitution and differentiating with respect to t yields

$$dR = i\frac{R}{|R|} \sqrt{u(t)} r(t) dt. \tag{6.13}$$

To compute a differential equation for $|R|$, one must be aware that $r(t)$ is a random variable. $r(t)dt$ is unbounded, and the standard chain rule of differential calculus must be replaced with Ito's lemma, which looks like the chain rule but expanded to higher order like a Taylor series. As $r(t)$ is normalized to have a variance dt, $(r(t)dt)^2 = dt$ and we have

$$d|R|^2 = u(t)dt \quad \rightarrow \quad |R(t)|^2 = \int_0^t u(s)ds. \tag{6.14}$$

Thus the time evolution of $|R|$ is deterministic. Substituting this solution into Eq. (6.13) yields

$$dR = iRP(t)r(t)dt, \quad P(t) \equiv \sqrt{\frac{u(t)}{\int_0^t u(s)ds}}. \tag{6.15}$$

Finally, we use Ito's lemma one more time to compute the differential increment of $\theta = \arg(R) = \text{Im}[\log(R)]$

$$d\theta = \text{Im}(d\log(R)) = \text{Im}\left[iP(t)r(t)dt + \frac{P(t)^2}{2}dt\right] = P(t)r(t)dt. \tag{6.16}$$

As $d\phi(t) = d\theta(t)$, Eq. (6.16) states that the instantaneous angular frequency of the measurement axis is proportional to the measurement outcome. Thus in the limit that the feedback

delay is small, the process of computing a quantum trajectory and then calculating the optimal phase may be reduced to applying proportional feedback. We implement this feedback law in the FPGA, as described in [80].

In the above, we have maximized phase sensitivity using an intuitive argument to justify $\phi(t) = \theta(t) + \pi/2$. To show that this indeed leads to a canonical phase measurement, we return to Eq. (6.7), substituting our solution Eq. (6.9) and Eq. (6.11)

$$P[r] \propto \langle \psi(\infty)|\psi(\infty)\rangle = |c_-(0) + c_+(0)R(\infty)^*|^2 \tag{6.17}$$
$$= |c_-(0)|^2 + c_-(0)c_+^* R(\infty) + c_-(0)^* c_+(0) R^*(\infty) + |c_+(0)|^2 |R(\infty)|^2$$

using $c_+(\infty) = 0$. We have written $\propto$ instead of $=$ because we dropped a normalization prefactor in Eq. (6.8). We can then compute the measurement operator by matching terms

$$P[r] = \text{Tr}[E[r]|\psi(0)\rangle\langle\psi(0)|] \propto \text{Tr}\left[\begin{pmatrix} E_{--} & E_{-+} \\ E_{+-} & E_{++} \end{pmatrix} \begin{pmatrix} |c_-(0)|^2 & c_-(0)c_+^*(0) \\ c_-^*(0)c_+(0) & |c_+(0)|^2 \end{pmatrix}\right] \tag{6.18}$$
$$= E_{--}|c_-(0)|^2 + E_{-+}c_-^*(0)c_+(0) + E_{+-}c_-(0)c_+^*(0) + E_{++}|c_+(0)|^2$$
$$\implies E(r) \propto \begin{pmatrix} 1 & R^*(\infty) \\ R(\infty) & |R(\infty)|^2 \end{pmatrix}.$$

We have $|R(\infty)|^2 = 1$ by Eq. (6.14), so

$$P[r] \propto 1 + \cos(\arg(c_-^*(0)c_+(0)) - \theta(\infty)) \tag{6.19}$$

where $\theta(\infty) = \arg(R(\infty))$ is by definition the measurement outcome. This expression coincides with the probability distribution for a canonical phase measurement. On might worry that we have dropped something in the proportionality factor $P(dW_1)P(dW_2)...$ that biases the measurement result away from Eq. (6.19) on average. However it is easily shown that this prefactor (often called P_{ost} for the ostensible statistics of the measurement) leads to a uniform distribution in $\arg(R^*)$, so the results are not affected.

Chapter 7

Optimal measurement-based control

In previous chapters, we have developed several methods for finding measurement-based feedback protocols, and used them to derive explicit control strategies. We have mostly concerned ourselves with finding protocols that achieve a certain task, and in some cases applied substantial constraints on the system in order to narrow the search space to something tractable. While the ASLO and symmetry restrictions lead to tractable calculations, we have no reason to expect that these simple solutions couldn't be improved upon by a less restrictive approach.

The following sections address this concern by proving that some of the presented protocols are theoretically optimal. In proving these results, we derive a powerful method for deriving measurement-based feedback protocols on bipartite systems, which was the technique originally used to derive the Hong-Ou-Mandel based feedback protocol mentioned in chapter 4.[1]

To date only a limited number of measurement-based quantum feedback protocols have been established to be globally optimal, owing in part to the non-linear nature of the problems. A feedback scheme for qubit purification was discovered by Jacobs [57] and its optimality properties were studied in a number of subsequent works [150, 149, 132, 67]. A protocol for rapid purification of a qudit was shown to exist in [21], and constructed explicitly in [125]. Upper bounds on qudit and multiqubit purification speedups are known [24], but proving global optimality for all situations remains an open challenge except for the qutrit case, which was solved in [125]. A globally optimal protocol for rotating a monitored qubit to a desired state was given in [128], but the precise way in which it is optimal is somewhat unnatural and several basic questions remain open.

Regarding feedback for entanglement generation, some optimality results are known for linear quadratic Gaussian systems [121, 75, 38]. Feedback protocols for enhancing the rate

---

[1] We have also generalized this method to tripartite systems based on the associated Schmidt decomposition[1]. However there are now five entanglement invariants instead of one, which appears to limit the usefulness of this technique. The resulting coupled, nonlinear stochastic differential equations did not provide any novel protocols, though they may provide a way to prove optimality of the GHZ protocol presented in chapter 4.

of entanglement generation using half-parity and full-parity measurements were given in [77] and [53] respectively, but without proof of optimality. Both protocols generate entanglement faster than is possible with measurement alone. They are also of interest because they drive the system along a deterministic path in Hilbert space; feedback is chosen to exactly cancel the randomness introduced by measurement, allowing for deterministic entanglement generation. This is particularly significant for the half-parity measurement, which has only a 50% success rate in the absence of feedback.

7.1 Optimal Entanglement Generation

In this section, we prove that the half-parity measurement protocol given in [77] and chapter 4 is globally optimal for several practically relevant tasks listed in the following section. We also consider entanglement generation under a continuous full-parity measurement $X_F = \sigma_{z1}\sigma_{z2}$, which behaves like the half-parity measurement except that it is also degenerate in the $\{|00\rangle, |11\rangle\}$ subspace. For the full-parity measurement, we derive a protocol that is related to that given in [53] but which is designed for initially pure states. Elaborating on the connection between rapid entanglement generation and rapid purification first discussed in [53], we derive a natural mapping between this two-qubit system and measurement of an effective qubit, in which the concurrence [151] exactly maps to this qubit's Bloch vector length. We then use this mapping to prove several optimality results for the proposed full-parity protocol. We also show that our full-parity protocol is optimal among all measurement-based entanglement generation schemes that acquire single qubit information at a fixed rate. This result sets a useful standard by which the effectiveness of an entanglement generation protocol may be assessed.

We assume a controller has the ability to apply feedback only in the form of local unitary (LU) rotations on the measured qubits, so that measurement is the only source for entanglement generation. This restriction substantially simplifies experimental implementation and is essential when the qubits are remotely separated. The protocol for the full and half-parity measurements X_F and X_H are given below, and the half-parity protocol is also given in chapter 4. We name these P_F and P_H respectively. Each protocol involves applying proportional feedback, *i.e.*, local unitary rotations with a rotation angle that is proportional to the most recent measurement outcome. In this section, we prove the following global optimality properties:

Half-parity measurement $X_H = (\sigma_{z1} + \sigma_{z2})/2$:
P_H is globally optimal for the following tasks.

1. Max. concurrence goal: P_H maximizes the expectation value of the entanglement (quantified using the concurrence $\mathcal{C}$ [151]) reached at a chosen stopping time T.

2. Max. fidelity goal: Same as above but for fidelity $\mathcal{F}$ [60] with respect to any target state which is pure and maximally entangled.

3. Min. time goal for concurrence: P_H minimizes the expected time to reach a desired concurrence $\mathcal{C}_{\text{Threshold}}$, so long as $\mathcal{C}_{\text{Threshold}} \leq 1/\sqrt{2}$.

4. Min. time goal for fidelity: Same as above but for a desired fidelity $\mathcal{F}_{\text{Threshold}}$, so long as $\mathcal{F}_{\text{Threshold}} \leq (1 + \sqrt{2})/\sqrt{8}$

Full-parity measurement $X_F = \sigma_{z1}\sigma_{z2}/2$:

1. Max. concurrence goal: P_F is globally optimal.

2. Max. Fidelity goal: P_F is globally optimal

3. Min. time goal ($\mathcal{C}$ or $\mathcal{F}$): Measuring without feedback is globally optimal.

The min. time results for concurrence also apply to any monotonic function of $\mathcal{C}$, such as entanglement of formation. Finally, we show that P_F is optimal for the max. concurrence goal among all protocols that acquire single qubit information at a fixed rate. A precise statement of this result is given with its proof.

In what follows, we make the simplifying assumptions that $\eta = 1$ and that the initial state is pure. In the absence of additional decoherence, the state remains pure at all times, which allows us to use a stochastic Schödinger equation instead of the corresponding master equation

$$\psi(t + dt) = \left[-\frac{1}{2}(X - \langle X \rangle)^2 dt + (X - \langle X \rangle)dW(t) \right]|\psi(t)\rangle, \tag{7.1}$$

where we have set the measurement rate $\Gamma = 2$ for simplicity, a convention that we will retain in what follows.

Our objective is to quantify how entanglement changes under arbitrary feedback protocols for binary systems, so that we can identify which are optimal. We characterize entanglement using the concurrence, which for pure states is defined as [151]

$$\mathcal{C} \equiv |\langle \psi^* | \sigma_y \otimes \sigma_y | \psi \rangle| \tag{7.2}$$

with $\langle \psi^* |$ the complex conjugate of $\langle \psi |$, or equivalently the transpose of $|\psi\rangle$. Like all valid entanglement measures, $\mathcal{C}$ is invariant under local unitary rotations, so our allowed feedback operations leave it unchanged. For the time being, we take $\mathcal{C}$ to be our figure of merit.

Since all bipartite pure states with the same concurrence are equivalent up to LU operations [93], it is possible to parameterize any pure state in terms of $\mathcal{C}$ and single qubit rotations. Such a parameterization is useful in this context because feedback can directly control the latter quantities. Therefore we can model feedback to directly set these qubit rotation parameters to desired values without specifying the Hamiltonian necessary to prepare the resulting state. The Schmidt decomposition provides an explicit example of such a parameterization [93]. By expressing the Schmidt coefficients in terms of the concurrence,

we can write a general two-qubit state as

$$\psi(\mathcal{C}, \theta_1, \theta_2, \phi_1, \phi_2, \gamma_1, \gamma_2) \tag{7.3}$$

$$= U_1 \otimes U_2 \left[\sqrt{\frac{1 + \sqrt{1 - \mathcal{C}^2}}{2}} |00\rangle - \sqrt{\frac{1 - \sqrt{1 - \mathcal{C}^2}}{2}} |11\rangle \right]$$

$$U_i \equiv \exp(-i\gamma_i\sigma_z/2)\exp(-i\sigma_y\theta_i/2)\exp(-i\phi_i\sigma_z/2),$$

where we have written U_i in terms of the Euler angles $\{\phi_i, \theta_i, \gamma_i\}$. For convenience in subsequent calculations, we break the local unitaries into symmetric and antisymmetric rotations by defining $\theta \equiv (\theta_1 + \theta_2)/2$, $\Delta\theta \equiv (\theta_1 - \theta_2)/2$, and likewise for ϕ and γ. The final expression does not depend on $\Delta\phi$ because the state in brackets is invariant under antisymmetric rotations about σ_z. This separation of variables into an LU-invariant quantity depending on $\mathcal{C}$ and LU rotations depending on $\{\theta_i, \phi_i, \gamma_i\}$ motivates an analogy with parameterization of a qubit in terms of a Bloch vector $\vec{r}$ in spherical coordinates; $\mathcal{C}$ is analogous to $|\vec{r}|$, which is invariant under unitary operations, while $\{\theta_i, \phi_i, \gamma_i\}$ can be set arbitrarily using Hamiltonian feedback like θ and ϕ. When considering the full-parity measurement, we will find that this analogy admits an explicit mapping in which $\mathcal{C}$ and r obey the same equations of motion under measurement.

To study how entanglement changes under measurement, we substitute Eq. (7.3) into Eq. (7.1) and compute the concurrence of the resulting state $\psi(t + dt)$. It is not necessary to compute how the five angles parameterizing ψ evolve, since we will model the controller to set them according to some feedback protocol. The computation is further simplified by the fact that σ_z rotations commute with X_F and X_H, so that the resulting equations of motion do not depend on γ or $\Delta\gamma$. We first focus on the half-parity measurement X_H. As is detailed in the appendix of [78], application of Ito's lemma yields

$$dC = \begin{cases} 2\mathcal{C}\sqrt{1 - \mathcal{C}^2}uv \ dW \\ \quad + [(v^2 - u^2)w - \mathcal{C}(v^2 + u^2)]dt & | \ \mathcal{C} > 0 \\ |v^2 - u^2| \ dt & | \ \mathcal{C} = 0 \end{cases} \tag{7.4}$$

where we have defined the control variables $u = \cos(\theta)$, $v = \cos(\Delta\theta)$ and $w = \cos(2\phi)$.

In what follows, we implicitly assume that at any time, the controller can instantaneously set the angles θ, $\Delta\theta$ and ϕ to any desired value. This is equivalent to assuming that the controller can implement any local unitary on both qubits with no time delay. For many systems such as superconducting qubits, single qubit rotations can be performed much faster than the measurement time scales, which makes this approximation appropriate. The assumption of zero time delay can be satisfied as long as the propagation delay between the qubits is small relative to the inverse measurement rate. However, this restriction can be relaxed for the full-parity protocol, as it will turn out that the protocol can be implemented by applying

feedback on only one of the two qubits. Thus, Eq. (7.4) is an equation of motion for the concurrence of the state under arbitrary feedback protocols that may be specified by choosing a particular, set $u(t, \mathcal{C})$, $v(t, \mathcal{C})$ and $w(t, \mathcal{C})$.

Although there is no reason *a priori* that a locally optimal strategy is also globally optimal, it often turns out to be so in practice. This will be the case in the results that we prove here. The locally optimal protocol maximizes the expectation value of the concurrence at time $t + dt$. To find it, one simply chooses the values of u, v and w to maximize the dt term of Eq. (7.4). This occurs for $\{u = 0, v = 1, w = 1\}$ with all values of γ and $\Delta\gamma$ allowed.[2] We henceforth refer to this protocol as P_H. The resulting equations of motion under this set of control parameters may be easily solved:

$$d\mathcal{C} = (1 - \mathcal{C})dt \quad \implies \quad \mathcal{C}(t) = 1 - (1 - \mathcal{C}(0))e^{-t}. \tag{7.5}$$

The evolution of the state under this feedback protocol may be computed by substituting this solution and the above control values into Eq. (7.3). Setting $\mathcal{C}(0) = 0$ for simplicity, the state evolution under feedback is thus

$$\psi(t) = \frac{1}{2} \begin{bmatrix} e^{-i\gamma}\sqrt{e^{-t}} \\ e^{-i\Delta\gamma}\sqrt{2 - e^{-t}} \\ e^{i\Delta\gamma}\sqrt{2 - e^{-t}} \\ e^{i\gamma}\sqrt{e^{-t}} \end{bmatrix}. \tag{7.6}$$

For $\gamma = \Delta\gamma = 0$, this state evolution exactly matches the solution to the feedback protocol given in [77] and chapter 4, indicating that they are the same feedback protocol.

The two methods that we have used to obtain this result differ substantially. The method of chapter 4 parameterizes feedback operations according to what an experimental controller would actually implement. Being a proportional feedback strategy, Eq. (4.23) calculates $d\rho$ by computing state updates due to infinitesimal measurement and feedback. In contrast, the above method applies infinitesimal measurement to determine the change in entanglement, but otherwise discards all information about the state at each time step. In a sense, Eq. (7.4) is an equation of motion for an equivalence class of states, and the state is parameterized not by the components of ψ, but rather by the unitary operations that one would perform to map $\alpha|00\rangle - \beta|11\rangle$ (the representative element singled out in Eq. (7.3)) to the physical state. Each method has strengths and weaknesses. PaQS allows for treatment of mixed states, average evolution and an arbitrary Hilbert space, while the present method allows us to more easily maximize over all possible protocols, including those with discontinuous evolution. The latter feature is necessary for optimality proofs, where we essentially need to maximize over all possible protocols.

---

[2]$\{u = 1, v = 0, w = -1\}$ is also a solution, but since it is actually equivalent to the first if one makes the transformation $\gamma \to \gamma - \pi/2$, $\Delta\gamma \to \Delta\gamma + \pi/2$, we ignore it. Solutions corresponding to $u = -1$ or $v = -1$ are similarly redundant.

Equation (7.6) can be more easily understood by writing down the states for $t = 0$ and $t \to \infty$. Up to a global phase, these states are

$$\psi(0) = \frac{1}{2} \left(|0\rangle + e^{i\gamma_1} |1\rangle \right) \otimes \left(|0\rangle + e^{i\gamma_2} |1\rangle \right)$$

$$\psi(\infty) = \frac{1}{\sqrt{2}} \left(|01\rangle + e^{2i\Delta\gamma} |10\rangle \right). \tag{7.7}$$

Thus this protocol involves preparing both qubits in a separable state polarized along some axis of the equator of their respective Bloch spheres. Feedback deterministically projects the qubits into a Bell pair with an arbitrary relative phase.

Before proving global optimality results for the above protocol, we first repeat the above analysis for the full-parity measurement. Following similar steps to those above but now for X_F, the equations of motion for the concurrence can be shown to be

$$dC = \begin{cases} (1 - C^2)\big[(u^2 - v^2)w\, C\, dW \\ \qquad + (u^2 - v^2)^2(1 - w^2)dt/2C\big] & | \, C \neq 0 \\ (u^2 - v^2)dW & | \, C = 0, \end{cases} \tag{7.8}$$

A minor technicality in the derivation detailed in the appendix of [78] has forced us to allow C to take on negative values in general. However since $\psi(C)$ only depends on C^2, this fact presents no further difficulty. We shall simply interpret $|C|$ as the concurrence instead of C.

For the full-parity measurement, two distinct sets of locally optimal parameters emerge: $\{u = 0, v = 1, w = 0\}$ and $\{u = 1, v = 0, w = 0\}$. The resulting state evolutions are equivalent up to a complex conjugation of ψ, so the underlying dynamics are essentially identical. We may therefore only consider the protocol $P_F : \{u = 0, v = 1, w = 0\}$ and ignore the other solution. As for the half-parity case, the equations of motion for $C(t)$ are deterministic and again yield an analytic expression with easy solution:

$$dC = \frac{1 - C^2}{2C} dt \implies C(t) = \pm\sqrt{1 - [1 - C(0)^2]e^{-t}}. \tag{7.9}$$

Validity of this solution at $C(0) = 0$ is most easily established by deriving equations of motion for $C(t)^2$ at $C = 0$ and showing that the solutions coincide. State evolution under this protocol is given by

$$\psi(t) = \frac{1}{\sqrt{8}} \begin{bmatrix} e^{-i(\gamma+\pi/4)}\left(\sqrt{1 - e^{-t/2}} - i\sqrt{1 + e^{-t/2}}\right) \\ e^{-i(\Delta\gamma+\pi/4)}\left(\sqrt{1 - e^{-t/2}} + i\sqrt{1 + e^{-t/2}}\right) \\ e^{i(\Delta\gamma-\pi/4)}\left(\sqrt{1 - e^{-t/2}} + i\sqrt{1 + e^{-t/2}}\right) \\ e^{i(\gamma-\pi/4)}\left(\sqrt{1 - e^{-t/2}} - i\sqrt{1 + e^{-t/2}}\right) \end{bmatrix}. \tag{7.10}$$
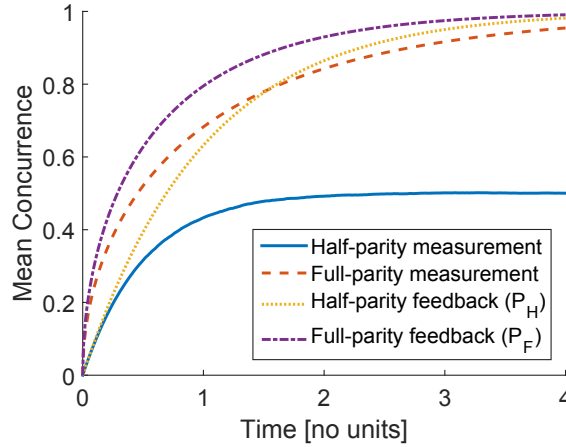
Figure 7.1: Expectation value of the concurrence while applying feedback or by applying measurement alone. Using the method of linear quantum trajectories [55], concurrence under full-parity measurement without feedback can be shown to take the form $\mathcal{C}(t) = \text{erf}(\sqrt{t/2})$. Concurrence under half-parity measurement without feedback is calculated numerically by averaging 10000 trajectories.

Again, taking early- and late-time limits of the solution gives some insight into the induced dynamics. Here

$$\psi(0) = \frac{1}{2}\big(|0\rangle + e^{i\gamma_1}|1\rangle\big) \otimes \big(|0\rangle - ie^{i\gamma_2}|1\rangle\big) \tag{7.11}$$

$$\psi(\infty) = \frac{1}{2}\big(e^{-i\gamma}|00\rangle + e^{-i\Delta\gamma}|01\rangle + ie^{i\Delta\gamma}|10\rangle - ie^{i\gamma}|11\rangle\big),$$

where as before we have dropped a global phase from the states. The optimal initial state is again to prepare both qubits in an equal superposition of $|0\rangle$ and $|1\rangle$. Somewhat counterintuitively, the final state produced by feedback is not an eigenstate of the measurement operator. This is somewhat analogous to Jacobs' purification speedup protocol, in which the state is maintained to be in an equal superposition of the measurement eigenstates [57]. Application of Jacobs' protocol to an encoded qubit led the authors of [53] to a method for converting a classically correlated mixed state into a maximally entangled state. The protocol presented here performs the analogous task for an initially pure state. $\mathcal{C}(t)$ under application of P_F coincides with that given in [53], as shown in [78]. Below we establish an even more precise connection between P_F and Jacobs' rapid qubit purification protocol.

The performance of the half-parity and full-parity feedback protocols are plotted in Fig. 7.1. We also plot the average concurrence without feedback for both measurement operators. We have normalized the measurement operators X_F and X_H so that the dephasing rate on one of the two qubits (tracing out the other qubit and assuming no feedback) are the same under measurement, allowing quantitative comparison the half- and full-parity performance.

It is evident that the full-parity measurement with feedback strictly outperforms all other protocols.

We now set out to prove the optimality results stated above. For the half-parity feedback protocol, we apply the standard verification theorems often used in control theory to check whether a protocol is globally optimal [56, 149]. While one might expect that proving such a result would require knowledge about all allowed strategies, it turns out that knowing the performance of the trial protocol is sufficient, so long as one knows how it behaves for any initial state.

Before using the verification theorems, we first provide a brief summary following Ref. [56]. A common feedback goal is to minimize the expectation value of some cost function after applying feedback for a fixed time interval T. The cost function can depend on the final state of the system as well as the resources used to apply feedback. From the cost function, one defines the cost-to-go $c(P, \mathbf{x}, t)$ where the vector $\mathbf{x}$ represents the system state at time t. Here c is defined to be the expectation value of the cost function at time T, assuming that the system was in the state $\mathbf{x}$ at time t and that the controller used feedback protocol P from t to T. To check optimality, one first generically writes the equations of motion for $\mathbf{x}$ in terms of their deterministic and stochastic parts as

$$d\mathbf{x} = \mathbf{A}(t, \mathbf{v}_P(t), \mathbf{x}(t))\ dt + \mathbf{B}(t, \mathbf{v}_P(t), \mathbf{x}(t))\ dW \tag{7.12}$$

where $\mathbf{v}_P$ are the feedback control settings and parameters specified by protocol P. In the context of quantum feedback, $\mathbf{v}$ can represent Hamiltonian, measurement and dissipation parameters. The procedure can easily be generalized to include multiple noise processes if needed. If $\partial c(P)/\partial t$ and $\partial^2 c(P)/\partial \mathbf{x}^2$ are continuous, then global optimality of P follows from the following two conditions. Firstly, assuming the cost function only depends on the final state, the Hamilton-Jacobi-Bellman equation for the cost-to-go $c(P, \mathbf{x}, t)$, given by

$$G(t, \mathbf{u}, \mathbf{x}, P) \equiv \tag{7.13}$$
$$-\frac{1}{2}\mathbf{B}(t, \mathbf{u}, \mathbf{x})^\top \frac{\partial^2 c(P)}{\partial \mathbf{x}^2}\mathbf{B}(t, \mathbf{u}, \mathbf{x}) - \mathbf{A}(t, \mathbf{u}, \mathbf{x})^\top \frac{\partial c(P)}{\partial \mathbf{x}}$$

and

$$\frac{\partial c(P)}{\partial t} = \max_{\mathbf{u}}[G(t, \mathbf{u}, \mathbf{x}, P)]. \tag{7.14}$$

must be satisfied for all time 0 to T and for all states $\mathbf{x}$. Secondly, $\mathbf{u} = \mathbf{v}_P$ must maximize G for all time and states.

In order to apply the verification theorems, we parameterize the state $|\psi\rangle$ as a vector $\mathbf{x} = \{\mathcal{C}, u, v, w\}$ using Eq. (7.3) and the definitions of u, v and w given below Eq. (7.4). Applying a specific local unitary feedback protocol is equivalent to setting $\mathbf{v} = \{u, v, w\}$ equal to specific functions of t and $\mathcal{C}$.

Half-parity measurement: For the max. concurrence goal (maximizing $\langle \mathcal{C}(T)\rangle$), we define our cost function to be $1 - \mathcal{C}(t)$. Choosing the locally optimal protocol for the half-parity measurement $P_H : \{u = 0, v = 1, w = 1\}$, we can use the analytic solution Eq. (7.5)

to find the cost-to-go when evolving according to P_H as

$$c(P_H, \mathbf{x}, t) = (1 - \mathcal{C}(t))e^{t-T}. \tag{7.15}$$

The Hamilton-Jacobi-Bellman equation uses the following derivatives of c

$$\frac{\partial c}{\partial t} = (1 - \mathcal{C}(t))e^{t-T}, \quad \frac{\partial c}{\partial \mathbf{x}} = \{-e^{t-T}, 0, 0, 0\}, \quad \frac{\partial^2 c}{\partial \mathbf{x}^2} = 0 \tag{7.16}$$

which satisfy the continuity conditions required by the verification theorem. We divide the equations of motion under feedback and half-parity measurement Eq. (7.4) into their deterministic and stochastic parts as

$$d\mathcal{C} = A_\mathcal{C} \; dt + B_\mathcal{C} \; dW$$

$$A_\mathcal{C} = \begin{cases} [(v^2 - u^2)w - \mathcal{C}(v^2 + u^2)]dt & | \; \mathcal{C} > 0 \\ |v^2 - u^2| \; dt & | \; \mathcal{C} = 0 \end{cases} \tag{7.17}$$

$$B_\mathcal{C} = \begin{cases} 2\mathcal{C}\sqrt{1 - \mathcal{C}^2}uv \; dW & | \; \mathcal{C} > 0 \\ 0 & | \; \mathcal{C} = 0 \end{cases}$$

where we have only calculated the $\mathcal{C}$ components of $\mathbf{A}$ and $\mathbf{B}$. Due to the form of G in Eq. (7.13), the other components are unnecessary, since $\partial c/\partial \mathbf{v} = 0$, *i.e.*, our cost function is invariant under the feedback control settings. Substituting Eqs. (7.16) and (7.17) into Eqs. (7.13) and (7.14), we find the condition for P_H to be globally optimal is that it satisfies the maximization condition

$$(1 - \mathcal{C}(t))e^{t-T} \tag{7.18}$$
$$= \max_{\{u,v,w\}}[(v^2 - u^2)w - \mathcal{C}(t)(v^2 + u^2)]e^{t-T}.$$

The maximum occurs for the values of u, v and w specified by P_H, and the equation then is satisfied. This proves that the half-parity protocol P_H is globally optimal for maximizing the concurrence at fixed time T. Because $c(P_H, \mathbf{x}, t)$ is linear in $\mathcal{C}$, one can also show that global optimality follows directly from local optimality in this case [132], but this proof method is not applicable in general.

A similar calculation can be performed in which we test for global optimality with respect to the min. time goal [56] (minimizing the expected time at which $\mathcal{C}$ reaches some desired value $\mathcal{C}_{\text{Threshold}}$, called the expected hitting time). One finds that P_H maximizes G only when $\mathcal{C} \leq 1/\sqrt{2}$, which implies that P_H is not globally optimal in general. If one chooses $\mathcal{C}_{\text{Threshold}} \leq 1/\sqrt{2}$ however, then only the dynamics of the system when $\mathcal{C} \leq 1/\sqrt{2}$ are relevant for determining the expected hitting time. Thus we can instead ask if P_H is globally optimal within the constraint that $\mathcal{C} \leq 1/\sqrt{2}$. Restricted to this parameter space, the

Hamilton-Jacobi-Bellman equation Eqs. 7.13-7.14 is satisfied for all allowed values of $\mathbf{x}$, and P_H maximizes G. This proves that P_H is globally optimal for the min. time goal when $\mathcal{C}_{\text{Threshold}} \leq 1/\sqrt{2}$. As the expected hitting time is the same whether one considers $\mathcal{C}$ or some arbitrary monotonic function $f(\mathcal{C})$ and the corresponding threshold $f(\mathcal{C}_{\text{Threshold}})$, the min. time proof applies to all monotonic functions of $\mathcal{C}$, such as entanglement of formation.

Concurrence has the favorable property of being invariant under the control parameters, which makes it amenable to methods for proving global optimality. For many tasks however, a specific target state is desired, in which case fidelity is a more relevant figure of merit. We now show that the global optimality proofs given above extend to the corresponding fidelity goals.

We take any maximally entangled state $|\Psi\rangle$ to be the target state and define P_H^Ψ to be a variant of P_H which rotates ψ to have the maximal fidelity with respect to Ψ at the final time. For the max. fidelity goal, the final time is simply T. For the fidelity min. time goal, the final time is the earliest time at which $\mathcal{F}_{\text{Threshold}}$ can be reached. For both of these goals, we need to know the fidelity of a given state ψ maximized over all local unitaries. For pure states, this maximal fidelity is uniquely determined by the concurrence of ψ, and is given by [137]

$$\mathcal{F}_{\text{max}} = \frac{\mathcal{C}+1}{2}. \tag{7.19}$$

In the context of entanglement distillation, $\mathcal{F}_{\text{max}}$ is often called the singlet fraction, and measures the usefulness of a general quantum state for the task.

To prove that P_H is globally optimal for the fidelity min. time goal, we note that any optimal protocol must have $\mathcal{F} = \mathcal{F}_{\text{max}}(\mathcal{C})$ at the hitting time. If this were not the case, the protocol could have applied local unitaries to ψ at an earlier time that increase $\mathcal{F}(\psi)$ to $\mathcal{F}_{\text{max}}(\mathcal{C}(\psi))$, and hence achieve an earlier hitting time. As $\mathcal{F}_{\text{max}}$ is a monotonic function of $\mathcal{C}$, this relation suffices to prove optimality of P_H with respect to the fidelity min. time goal when $\mathcal{F}_{\text{Threshold}} \leq \mathcal{F}_{\text{max}}(1/\sqrt{2}) = (1+\sqrt{2})/\sqrt{8}$.

To prove global optimality of the max. fidelity goal, we assume the evolution of the system is well-approximated by a discrete protocol in which one measures for a small but finite duration before applying feedback [132]. As realistic implementations of feedback inevitably suffer from feedback delay and finite bandwidth effects at sufficiently short timescales, the continuum limit may not be a good physical model, and we shall therefore not focus on this here. In the discrete approximation, we can express the expectation value of the fidelity at the stopping time T as an integral over all possible measurement outcomes $\{\mathbf{V}\}$. Suppose that some hypothetical feedback protocol P' is globally optimal for the max. fidelity goal. We write

$$\langle \mathcal{F}(P', T) \rangle = \int \mathcal{F}_{\mathbf{V}}(P', T) p(\mathbf{V}) d\mathbf{V}$$

$$= \int \mathcal{F}_{\text{max}}(\mathcal{C}_{\mathbf{V}}(P', T)) p(\mathbf{V}) d\mathbf{V} \tag{7.20}$$

where $\mathcal{F}_\mathbf{V}(P',T)$ and $\mathcal{C}_\mathbf{V}(P',T)$ are respectively the fidelity and concurrence at time T, assuming protocol P' was applied and measurement outcome $\mathbf{V}$ occurred. $p(\mathbf{V})$ is the probability of $\mathbf{V}$ occurring (note that p implicitly depends on P' and the state evolution, but this is of no consequence for the proof). The last equality follows from the fact that the globally optimal protocol could perform at least as well if it maximizes the fidelity of each possible final state. If it does not do so on a set of non-zero measure, then a better protocol exists which does. We continue by relating the performance of P' to that of P_H^Ψ:

$$
\begin{aligned}
&= \int \frac{\mathcal{C}_\mathbf{V}(P',T)+1}{2}p(\mathbf{V})d\mathbf{V} \\
&\leq \int \frac{\mathcal{C}_\mathbf{V}(P_H^\Psi,T)+1}{2}p(\mathbf{V})d\mathbf{V} \\
&= \int \mathcal{F}_{\max}(\mathcal{C}_\mathbf{V}(P_H^\Psi,T))p(\mathbf{V})d\mathbf{V} = \int \mathcal{F}_\mathbf{V}(P_H^\Psi,T)p(\mathbf{V})d\mathbf{V} \\
&= \langle \mathcal{F}(P_H^\Psi,T)\rangle.
\end{aligned}
\tag{7.21}
$$

The inequality follows from global optimality of P_H with respect to the max. concurrence goal. Note that relating the performance of P' to that of P_H relies on the fact that $\mathcal{F}_{\max}$ is linear in $\mathcal{C}$. The second-to-last equality follows from the fact that at time T, P_H^Ψ rotates ψ to have maximum fidelity with respect to the target state. Together Eqs. (7.20) and (7.21) imply that P_H^Ψ performs at least as well as any potential protocol P', and therefore that P_H^Ψ is globally optimal for the task.

Full-parity measurement: A connection between Jacobs' rapid purification protocol, and rapid entanglement using full-parity measurement was first established in Ref. [53]. To prove optimality of the full-parity protocol P_F, we observe that the dynamics of the concurrence of a two-qubit pure state under full-parity measurement and local feedback are precisely those of the Bloch vector length of a single continuously monitored qubit with feedback. We then use existing optimality results regarding rapid qubit purification [149] to prove the analogous two-qubit results.

To see the correspondence between 2-qubit concurrence $\mathcal{C}$ and the Bloch vector length of a single qubit, consider a qubit undergoing a continuous measurement of $\tilde{\sigma}_z$ at a rate $\tilde{\Gamma}$. We assume that some feedback controller can instantly apply any unitary operation at any time, as we have assumed for the two-qubit case. Parameterizing the qubit as a Bloch vector in spherical coordinates $\{\tilde{r},\tilde{\theta},\tilde{\phi}\}$, one can use the stochastic master equation to derive an equation of motion for the Bloch vector length $\tilde{r}$ as a function of $\tilde{\theta}$ and $\tilde{\phi}$ [67]

$$
d\tilde{r} = (1-\tilde{r}^2)\Big[\frac{\tilde{\Gamma}}{4\tilde{r}}(1-\tilde{u}^2)dt + \sqrt{\frac{\tilde{\Gamma}}{2}}\tilde{u}\ dW\Big]
\tag{7.22}
$$

where $\tilde{u} = \cos(\tilde{\theta})$ and $\tilde{\theta}$ is the angle the Bloch vector makes with the measurement axis $\tilde{\sigma}_z$. We do not derive equations of motion for $\tilde{\theta}$ and $\tilde{\phi}$ because we assume feedback can set them

to their desired values at any time. Making the following identifications

$$\mathcal{C} \to \tilde{r}$$
$$w \to \tilde{u}$$
$$2(u^2 - v^2)^2 \to \tilde{\Gamma}, \tag{7.23}$$

the equation of motion for concurrence under full-parity measurement and feedback, Eq. (7.8) becomes Eq. (7.22). This mapping reveals several interesting features of the dynamics. As observed in Ref. [53], entanglement generation can be turned on and off by rotating the system into one of the decoherence free subspaces of the measurement operator X_F. This is evident from the dependence of $\tilde{\Gamma}$ on u and v, as $u = v = 1$ and $u = v = 0$ correspond to states fully localized to the even and odd parity subspaces, respectively. Of particular interest is the direct mapping between $\mathcal{C}$ and $\tilde{r}$; the concurrence coincides exactly with Bloch vector length of the effective qubit. We provide more details on this effective qubit in the appendix.

Although the effective qubit we consider here is different from that discussed in Ref. [53], rapid entanglement generation corresponds to applying Jacobs' protocol to the effective qubit in both mappings. $\tilde{u} = 0$ is globally optimal for maximizing the linear entropy $\langle \tilde{r}(T) \rangle$ as shown in [149]. Although that work did not consider allowing the measurement rate $\tilde{\Gamma} \leq 2$ to vary, it is straightforward to extend the proofs of [149] to this case by repeating their calculation with $\tilde{\Gamma}$ as a control variable bounded from 0 to 2. Thus the mapping between equations of motion for $\tilde{r}$ and $\mathcal{C}$ implies that P_F is globally optimal for the max. concurrence goal.

It is also the case that not applying feedback to the effective qubit yields the same equations of motion as not applying feedback to the two-qubits. Since it was proved in [149] that not applying feedback is globally optimal for the min. time goal for linear entropy, the analogous result applies to the full-parity measurement of two qubits for the concurrence min. time goal. The arguments used to extend the max. concurrence and the concurrence min. time results of P_H to the corresponding fidelity results (*i.e.* Eqs. 7.19-7.21) also apply to P_F without modification. This completes the proof of the enumerated results given in at the beginning of this section.

Upper-bound on measurement-based protocols: So far, we have focused on optimality given a fixed measurement operator. However, one may ask whether a different measurement operator could offer superior performance using similar resources. In the context of remote entanglement generation, when entanglement is created using some signal degree of freedom as an intermediary (see Fig. 7.2), one could ask whether entanglement is transferred with unit efficiency. Motivated by these questions, we now prove a more general result which sets an upper limit on the entanglement entropy achievable generally under a much larger class of measurement-based protocols. We will find that the bound is only saturated by the full-parity feedback protocol P_F.

We consider any system in which the action of the measurement on qubit 1 is of the form

$$d\rho_1 = (\Gamma_{\text{deph.}}/2)\mathcal{D}[\vec{\sigma} \cdot \hat{n}]\rho_1 \, dt \tag{7.24}$$
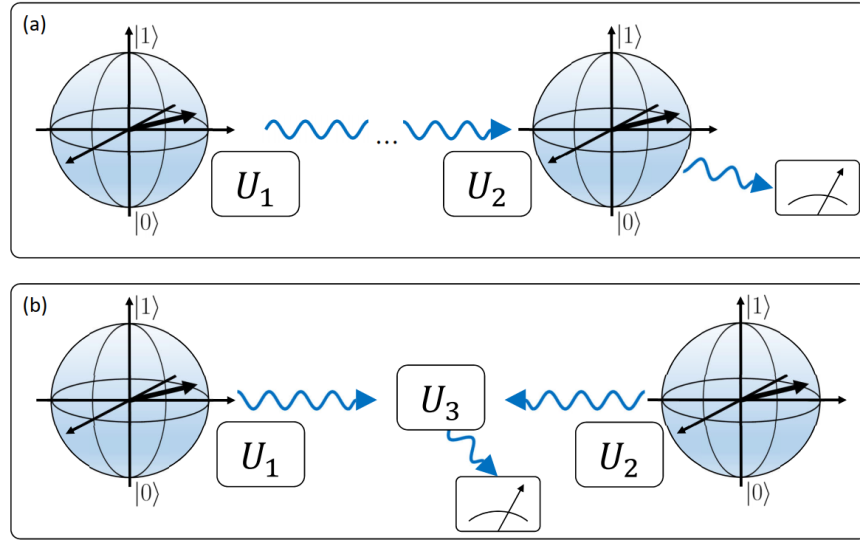
Figure 7.2: Remote 2-qubit measurement implementations. (a) Remote 2-qubit measurement in which the signal degrees of freedom (wavy line) propagate from qubit 1 to qubit 2. U_1 and U_2 represent the interactions between the signal and each qubit, while the dial represents projective measurement of the signal after qubit interactions. These three operations determine the effective 2-qubit measurement operator. (b) Remote measurement based on entanglement swapping. U_3 is the interaction between the incoming signals (such as a beam splitter) that erases which-path information. Note that in both schemes, the rate of entanglement generation is bounded by U_1.

where ρ_1 is the state of qubit 1 unconditioned on the measurement outcome and tracing out qubit 2. Physically, the dephasing rate $\Gamma_{\text{deph.}}$ sets an upper bound on the amount of information that can be extracted from the measurement [20], so this restriction fixes the rate at which information about qubit 1 is transferred to the rest of the system. Thus $\Gamma_{\text{deph.}}$ defines a physically meaningful reference that lets us compare the performance of P_H and P_F to more general protocols. By tracing out qubit 2 in the stochastic master equation, one arrives at Eq. (7.24) with $\Gamma_{\text{deph.}} = 1/2$ for measurement of both X_H and X_F.

The entanglement entropy of a state is given by $E_1 = \text{Tr}[\rho_1 \log_2(\rho_1)]$, the Von Neumann entropy of qubit 1. Note that if the joint state is pure, then this expression also equals the entanglement of formation. In order for the entanglement to increase, the entropy of the subsystem must increase, as it does under the action of Eq. (7.24). As measurement constitutes the only available non-local interaction, Eq. (7.24) fully determines how entanglement may change. At time t, the increase in entropy of ρ_1 is maximized if the state is unbiased with respect to the measurement axis $\vec{\sigma} \cdot \hat{n}$. For example, if $\vec{\sigma} \cdot \hat{n} = \sigma_z$, then optimal states would be of the form $\rho_1 = (x\sigma_x + y\sigma_y + \sigma_0)/2$ where σ_0 is the identity matrix. Assuming this condition is satisfied at all times, the entropy as a function of time may be derived by solving Eq. (7.24) with the initial condition $\rho_1 = (|0\rangle + |1\rangle)(\langle 0| + \langle 1|)/2$. Taking $\Gamma_{\text{deph.}} = 1/2$,

we find

$$E_1(t) = -\frac{1 - \langle \sigma_{x1} \rangle}{2} \log_2 \left[\frac{1 - \langle \sigma_{x1} \rangle}{2} \right] \tag{7.25}$$
$$- \frac{1 + \langle \sigma_{x1} \rangle}{2} \log_2 \left[\frac{1 + \langle \sigma_{x1} \rangle}{2} \right]$$
$$\langle \sigma_{x1} \rangle = e^{-t/2}. \tag{7.26}$$

This Von Neumann entropy sets an upper bound on the entanglement entropy that can be achieved by the action of a given measurement operator (as well as on the entanglement of formation if the joint evolution is pure).

The main result of this section is that by tracing out qubit 2 from Eq. (7.3) and substituting Eq. (7.9), one can show that the entanglement entropy of qubit 1 is exactly Eq. (7.25). For pure states, a similar result can be derived for concurrence, which has a one-to-one relation with entanglement entropy in this case [151]. Thus P_F saturates the bound Eq. (7.25) given by the dephasing rate of X_F, and hence no measurement-based protocol of the form Eq. (7.24) can generate entanglement faster than P_F. It can also be shown that P_H does not saturate this bound, nor does measurement of X_F or X_H without feedback.

This bound may be intuitively be understood by considering a remote implementation of the measurement, as depicted in Fig. 7.2. Eq. (7.24) then governs how much entanglement is generated between qubit 1 and the auxiliary qubit. An ideal entangling protocol would transfer all of this entanglement to qubit 2, so that measurement of the auxiliary does not decrease the entanglement entropy. Evidently only P_F fully transfers the entanglement from the measurement signal to qubit 2. For the other protocols, the remaining entanglement is destroyed when the signal is measured.

Note that if the effect of measurement on qubit 1 yields a dephasing operator that is not normal (*i.e.*, $[X, X^\dagger] \neq 0$), then it cannot be put in the form of Eq. (7.24) and our derivation does not apply. Thus our bound does not apply for quantum-demolition measurements such as a spontaneous emission process on the first qubit *i.e.*, $d\rho_1 = \mathcal{D}[\sigma]\rho_1 \, dt$, though a similar bound could be derived for such cases. An example of the latter is entanglement generation via Hong-Ou-Mandel interference.

7.2   Optimal Hamiltonian-Free Control

Control with measurement alone can serve as a proxy to understand the ultimate advantages of measurement-based feedback and control. By understanding what is possible using arbitrary, potentially adaptive measurements, we gain insight into the limits of quantum feedback and derive a kind of 'quantum speed limit' for measurement-based control. In what follows, we develop a feedback protocol based on some of our experimental work[47] extending the results of chapter 5. The experiment involves using the single-quadrature measurement to control a quantum state. As the axis of a single-quadrature measurement can be changed during the measurement process, we can use the Zeno effect to 'drag' the Bloch vector of a
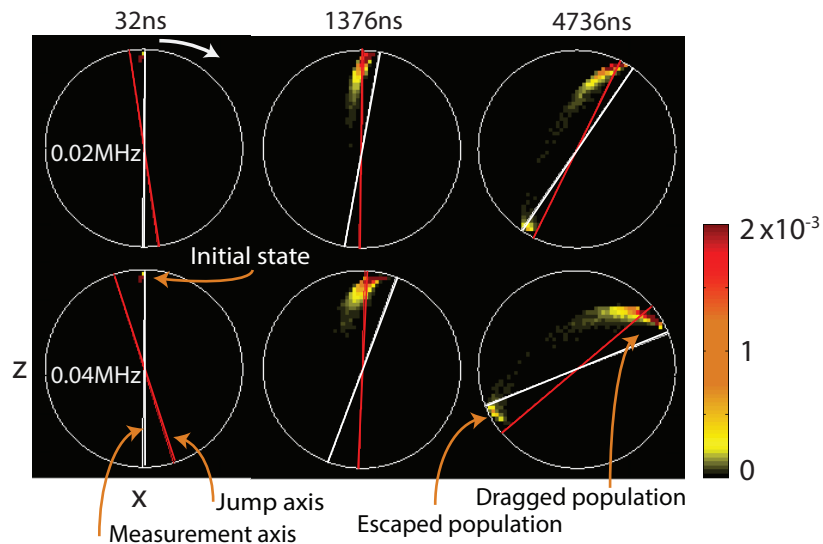
Figure 7.3: Histograms in the XZ plane of the Bloch vector at different time points in the Zeno dragging experiment. State is initialized at the North pole of the Bloch sphere at $t = 0$, and the measurement axis is aligned with the state. Measurement axis is gradually rotated with the angular frequency shown in each pane, which drags the state along with it via the Zeno effect. Red lines show a theoretically calculated 'jump axis' discussed in ref. [47].

qubit along the measurement axis. Experimental data showing this procedure are plotted in Fig. 7.3. In the limit of infinite measurement strength, the Bloch vector remains aligned to the measurement axis, enabling arbitrary state rotations simply by continuously changing the measurement axis. When considering measurements of finite strength, the state has a finite probability to jump to the other pole of the measurement axis, as is visually apparent from the histograms.

Given an upper bound on the measurement strength, there should exist an optimal strategy for minimizing quantum jumps and maximizing the fidelity to the target state. In general, this strategy could involve adaptive measurement, in which we choose the measurement axis based on past measurement outcomes. In the following, we imagine that the feedback controller maintains the measurement axis half way between the current and target states within a fixed plane of the Bloch sphere. For the time being, we restrict the Bloch vector and measurement axis to lie in the XZ plane. A variant of this protocol was given in ref. [8], though the work only considered the case when the current state is infinitesimally close to the target state. In what follows, we calculate the dynamics under this protocol and show that it is globally optimal for control with measurement alone.

We parameterize the state and measurement axis with the angular variables χ and δ

respectively as

$$|\psi\rangle = \begin{pmatrix} \psi_e \\ \psi_g \end{pmatrix} = \begin{pmatrix} \cos(\chi/2) \\ \sin(\chi/2) \end{pmatrix} \tag{7.27}$$

$$M = \cos(\delta)\sigma_z + \sin(\delta)\sigma_x.$$

Given our usual convention, the effective measurement rate is implicitly $\Gamma = 2$. If the initial state is pure and $\eta = 1$, then we can simplify calculations by working with a stochastic Schrödinger equation

$$d|\psi\rangle = \left[-\frac{1}{2}(M - \langle M \rangle)^2 dt + (M - \langle M \rangle)dW \right] |\psi\rangle. \tag{7.28}$$

We do not expect an exact solution to exist for an arbitrary feedback strategy. However things simplify nicely if we maintain the measurement axis halfway between the current and target states. If the target state is $|e\rangle$ (corresponding to $\chi = 0$), then the corresponding feedback strategy is $\delta = \chi/2$. The stochastic equation for ψ_e becomes

$$d\psi_e = (1 - \psi_e^2) \left(\frac{\psi_e}{2} dt + dW \right). \tag{7.29}$$

ψ_g is uniquely determined by ψ_e given our restriction to the XZ plane. Using Ito's lemma, we can derive an equation of motion for $\mathcal{F}_e = \psi_e^2$, the fidelity with respect to the target state

$$d\mathcal{F}_e = (1 - \mathcal{F}_e)dt + 2\sqrt{\mathcal{F}_e}(1 - \mathcal{F}_e)dW. \tag{7.30}$$

As the drift (dt) term is linear in $\mathcal{F}_e$, averaging over dW on both sides yields a very simple equation of motion for $\langle \mathcal{F}_e \rangle$

$$d\langle \mathcal{F}_e \rangle = (1 - \langle \mathcal{F}_e \rangle)dt + 2\langle \sqrt{\mathcal{F}_e}(1 - \mathcal{F}_e) \rangle \langle dW \rangle = (1 - \langle \mathcal{F}_e \rangle)dt \tag{7.31}$$
$$\implies \langle \mathcal{F}_e(t) \rangle = 1 - (1 - \langle \mathcal{F}_e(0) \rangle)e^{-t}.$$

Not only does linearity enable a simple solution to the equations of motion, but it also greatly simplifies the proof of optimality. The cost-to-go for the above protocol is

$$c(t) = (1 - \langle \mathcal{F}_e(t) \rangle)e^{t-T} \tag{7.32}$$

where T is the stopping time as usual. The corresponding Hamilton-Jacobi-Bellman equation is

$$\frac{\partial c}{\partial t} = \max_v \left[-\frac{\partial c_v}{\partial t} A_{e,v} - \frac{1}{2} \frac{\partial^2 c_v}{\partial t^2} B_{e,v}^2 \right] \tag{7.33}$$

The second term drops due to linearity. As $A_{e,v}$ is positive, the above maximization over c_v is then equivalent to the condition that u is locally optimal. It can be shown algebraically

that the above protocol is locally optimal, even if we maximize over all possible orientations of the measurement axis. By Eq. (7.33), it is also globally optimal.

We can also generalize the above feedback protocol to control an arbitrary quantum system. Suppose the current state is $|\psi\rangle$ and we wish to take it to $|\psi_T\rangle$. One can always find a basis in which $|\psi\rangle = \cos(\chi/2)|\psi_T\rangle + \sin(\chi/2)|\psi_T^\perp\rangle$, where $|\psi_T^\perp\rangle$ is orthogonal to $|\psi_T\rangle$. In the multidimensional case, imposing a maximum measurement rate amounts to upper-bounding the difference between the largest and smallest eigenvalues of M. Without loss of generality, let the maximal and minimal eigenvalues be ± 1. Let us orient the measurement operator M such that

$$M = (|\psi_T\rangle\langle\psi_T| - |\psi_T^\perp\rangle\langle\psi_T^\perp|)\cos(\delta) + (|\psi_T^\perp\rangle\langle\psi_T| + |\psi_T^\perp\rangle\langle\psi_T|)\sin(\delta) + \mathcal{O} \qquad (7.34)$$

$$\langle\psi_T|\mathcal{O}|\psi_T\rangle = \langle\psi_T^\perp|\mathcal{O}|\psi_T^\perp\rangle = \langle\psi_T^\perp|\mathcal{O}|\psi_T\rangle = \langle\psi_T|\mathcal{O}|\psi_T^\perp\rangle = 0.$$

In general, there will exist infinitely many unitary changes of basis that put M in this form. The first two terms of M exactly coincide with σ_z and σ_x in Eq. (7.27), and the eigenvectors associated to this subspace are ± 1. As $|\psi\rangle$ has no overlap with $\mathcal{O}$, none of the measurement outcomes associated to the eigenvectors of $\mathcal{O}$ can occur, and the problem reduces to the qubit case considered above. In essence, we have rotated M so that the maximally distinguished eigenvectors are 'in-plane' with the current and target states, effectively maximizing the measurement strength. If this choice of basis (with $\delta = \chi/2$ as before) is locally optimal, then it follows that this protocol is also globally optimal by the same argument as for a single qubit. This result would imply a measurement-based quantum speed limit for an arbitrary quantum system. We have numerical evidence of local optimality, but have not proven the result in full generality.

Chapter 8

Outlook

What new capabilities arise from the ability to resolve quantum measurements as the continuous processes that they are? We have identified and expanded on three potential areas, adaptive measurement, quantum foundations and entanglement generation. In what follows, we give a perspective on what may come next, and introduce one relatively new direction for future research.

A common theme in the work of this thesis is quantum feedback's ability to circumvent the limitations of linear optics. In chapter 4, we saw numerous examples in which feedback enhanced the success rate of entanglement generation in remote or weakly interacting systems. The reason that these experimental systems cannot create entanglement deterministically to begin with goes back to a fundamental result in linear optics that prevents deterministic implementation of a Bell measurement[16]. This theorem forces some fraction of the measurement outcomes to project into separable states. Feedback provides an effective nonlinearity, which allows us to steer the system toward entangled outcomes with up to 100% efficacy. These results provide evidence that there is something fundamental to be gained from applying quantum feedback to these types of problems. We have many examples in which these schemes can be applied in realistic experimental systems. We believe that most of the protocols developed in this thesis could be readily applied in superconducting circuit systems with parameters that have already been demonstrated. Expanding beyond superconducting circuits, we have developed an improved version of the standard Barrett-Kok scheme for remote entanglement generation in atomic and other solid state systems. This improved protocol retains the standard method's robustness to all common experimental imperfections, and will be the subject of a future manuscript.

The present work focused primarily on the preparation of relatively simple entangled states such as Bell, Dicke and GHZ states. Beyond these states, there is ample opportunity to apply the formalism developed in chapter 4 to more advanced systems. In a way, generation of these highly symmetric states states serves mostly as a proof of usefulness of the feedback equations derived here, and we expect them to be capable of far more. Interesting applications could be toward cluster states and tensor network states, where there is a locality structure to consider and exploit in the context of remote entanglement and quantum

networks. Furthermore, just as we see feedback benefiting systems constrained by linearity, we expect the converse to also hold to some extent; systems in which strong nonlinearities are available will likely benefit far less from quantum feedback. We have given some justification for this assertion in previous chapters, for instance by noting that any POVM can be synthesized by a quantum circuit.

Our implementation of a canonical phase measurement provides another example in which feedback fundamentally enhances the capabilities of a linear optics system. Adaptive measurement enables exactly reaching the canonical phase measurement in our experiment, and more generally can enhance the precision of a detector[4, 144]. Much progress has been made in theoretically mapping out the POVMs that may be implemented with continuous measurement[100]. However the question of what POVMs can be implemented given a measurement device and the ability to adapt is still largely an open theoretical problem. This question ties into the more specific question of determining the minimal resources to implement a general canonical phase measurement. Both of these questions would be interesting to address with stochastic control theory methods like the one developed in this thesis. Just as we addressed optimal control for state generation, it should be possible to apply these same techniques to optimal generation of a POVM.

There are a number of potential and known experimental applications of adaptive measurement worth noting. Firstly, the universality of weak measurements for generating arbitrary POVMs can be demonstrated elegantly in superconducting circuits with existing technology[30]. There are also protocols for enhanced state purification[57] and faster readout of qubit registers[23] with forms of adaptive measurement. More broadly, implementation of general quantum computation with linear optics has been a theoretically rich direction for quantum information[61] and a still nascent but promising experimental platform[112]. The canonical phase measurement fills a need in photonic quantum information processing by enabling the generation of arbitrary single-rail photonic qubits[28]. It will be interesting to see if adaptive measurement offers further novel capabilities for the technology moving forward, particularly for state generation.

The third major topic addressed in this thesis is the simultaneous measurement of non-commuting observables. While this novel capability may have applications for metrology and quantum state tomography, we feel that the most exciting future direction is in quantum foundations. When viewed in a fixed basis, quantum mechanics largely reduces to classical mechanics, and interesting phenomena like contextuality and nonlocality do not manifest themselves. Typical experiments probing these properties involve sequential or discrete measurements, but these do not mimic how the phenomena would manifest in a natural setting. In contrast, most physical systems are continuously measured by their environments. Thus the experimental direction here provides an opportunity to merge the study of quantum foundations and the study of decoherence in natural systems. Several works have begun investigating this line of research[72, 97]. As we noted in chapter 5, 'Error-disturbance bounds would provide a natural starting point, as we have separately extracted state information in Fig. 5.5b and quantified state disturbance in Fig. 5.7, but have not yet investigated the intrinsic trade-off between these concepts.'

All existing quantum trajectory experiments have been performed with one or two qubits. Is it possible to scale them to many-body systems? At first glance, the problem does not look scalable. To validate quantum trajectories experimentally, we repeated each experiment hundreds of thousands of times, acquiring statistics for all possible final states and then post-selecting on realizations that lead to the same final state. Quantum state tomography scales exponentially in the number of qubits, but trajectory validation scales exponentially in the number of orthogonal system states, which itself scales exponentially in the number of qubits! Fortunately, there are far less costly ways to test a trajectory model experimentally. One can post-select on states that lead to the same value of some particular operator, and then measure that operator repeatedly, for example. This method does away with all exponential scaling, except in theoretical simulation of the system itself. One can also use feedback to limit the space of possible final states. In a feedback-based control problem, the final number of states is ideally just the target state.

Quantum trajectories in AMO and many-body physics While our focus has been on elucidating new capabilities, quantum trajectories also offer a tool to more deeply understand existing, or even 'canonical' systems. One existing example is the ability to use weak measurements to measure correlators in many body systems[50]. More generally, quantum trajectories should offer both an experimental tool and a unique theoretical perspective. As only a handful of quantum trajectory experiments have been performed, many of the most basic physics effects at the foundation of AMO and many-body physics have only been observed at the ensemble level. Unravelling the decoherence channels in these systems as diffusing quantum trajectories offers a new perspective on even the most basic phenomena.

There are many standard experiments that one could imagine unravelling. The Mollow triplet has been observed in many atomic and artificial atomic systems, but quantum trajectories would allow direct observation of the transitions between dressed states that lead to this characteristic spectrum. Photon absorption and more generally atom-field dynamics in the presence of non-classical light could yield interesting experiments in which one monitors the trajectories of the joint atom-field state. Arguably, the recent observation of quantum jumps as a continuous process is a step in this direction[84]. A particularly interesting potential avenue would be the unravelling of superradiance and superradiant lasing. In the many-body limit, this phenomena may be understood as a nonlinear synchronization process between atoms, which is mathematically analogous to the synchronization of fireflies, for example[152]. At the level of a few-atom system, entanglement between atoms should begin to play a role. Observing entanglement generation and spontaneous synchronization with trajectories would be interesting not only from a fundamental physics perspective, but could also provide new directions for the development of superradiant lasers. Such devices would have unprecedented linewidths, and a trajectories approach could offer practical benefits via feedback.

One broader generalization of quantum trajectories experiments would be to study truly many-body physics. Phase transitions induced by dissipation, such as the superradiant phase transition would be natural candidates. However in principle virtually any phase transition could be continuously monitored by observing the bath with which it exchanges heat, volume,

particles and entropy. While observing the microscopic dynamics in a truly macroscopic material is daunting, much could be learned in cold-atom systems such as quantum gas microscopes, or in systems simulated via superconducting circuits.

Bibliography

[1] Antonio Acín, A Andrianov, L Costa, E Jané, JI Latorre, and Rolf Tarrach. Generalized schmidt decomposition and classification of three-quantum-bit states. *Physical Review Letters*, 85(7):1560, 2000.

[2] Dorit Aharonov, Wim Van Dam, Julia Kempe, Zeph Landau, Seth Lloyd, and Oded Regev. Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM review*, 50(4):755–787, 2008.

[3] Ernesto Altshuler and R Garcıa. Josephson junctions in a magnetic field: Insights from coupled pendula. *American Journal of Physics*, 71(4):405–408, 2003.

[4] Michael A Armen, John K Au, John K Stockton, Andrew C Doherty, and Hideo Mabuchi. Adaptive homodyne measurement of optical phase. *Physical Review Letters*, 89(13):133602, 2002.

[5] OV Astafiev, LB Ioffe, S Kafanov, Yu A Pashkin, K Yu Arutyunov, D Shahar, O Cohen, and JS Tsai. Coherent quantum phase slip. *Nature*, 484(7394):355, 2012.

[6] Roger Bach, Damian Pope, Sy-Hwang Liou, and Herman Batelaan. Controlled double-slit electron diffraction. *New Journal of Physics*, 15(3):033018, 2013.

[7] Roger Bach, Damian Pope, Sy-Hwang Liou, and Herman Batelaan. The double slit experiment performed with electrons, 2016.

[8] Ashkan Balouchi and Kurt Jacobs. Optimal measurement-based feedback control for a single qubit: a candidate protocol. *New Journal of Physics*, 16(9):093059, 2014.

[9] Sean D Barrett and Pieter Kok. Efficient high-fidelity quantum computation using matter qubits and linear optics. *Phys. Rev. A*, 71(6):060310, 2005.

[10] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *SIAM Journal on computing*, 26(5):1411–1473, 1997.

[11] Dominic Berry. Adaptive phase measurements. *arXiv preprint quant-ph/0202136*, 2002.

[12] Jeffrey Scott Birenbaum. *The C-shunt flux qubit: a new generation of superconducting flux qubit*. 2014.

[13] Sergey Bravyi and Alexei Kitaev. Universal quantum computation with ideal clifford gates and noisy ancillas. *Physical Review A*, 71(2):022316, 2005.

[14] Heinz-Peter Breuer, Francesco Petruccione, et al. *The theory of open quantum systems*. Oxford University Press on Demand, 2002.

[15] Daniel Klaus Burgarth, Paolo Facchi, Vittorio Giovannetti, Hiromichi Nakazato, Saverio Pascazio, and Kazuya Yuasa. Exponential rise of dynamical complexity in quantum computing through projections. *Nature communications*, 5:5173, 2014.

[16] John Calsamiglia and Norbert Lütkenhaus. Maximum efficiency of a linear-optical bell-state analyzer. *Applied Physics B*, 72(1):67–71, 2001.

[17] Philippe Campagne-Ibarcq, Pierre Six, Landry Bretheau, Alain Sarlette, Mazyar Mirrahimi, Pierre Rouchon, and Benjamin Huard. Observing quantum state diffusion by heterodyne detection of fluorescence. *Physical Review X*, 6(1):011002, 2016.

[18] MA Castellanos-Beltran, KD Irwin, GC Hilton, LR Vale, and KW Lehnert. Amplification and squeezing of quantum noise with a tunable josephson metamaterial. *Nature Physics*, 4(12):929, 2008.

[19] Y Chu, C Axline, C Wang, T Brecht, YY Gao, L Frunzio, and RJ Schoelkopf. Suspending superconducting qubits by silicon micromachining. *Applied Physics Letters*, 109(11):112601, 2016.

[20] A. A. Clerk, S. M. Girvin, Florian Marquardt, and R J Schoelkopf. Introduction to quantum noise, measurement, and amplification. *Rev. Mod. Phys.*, 82(2):1155–1208, April 2010.

[21] Joshua Combes and Kurt Jacobs. Rapid State Reduction of Quantum Systems Using Feedback Control. *Phys. Rev. Lett.*, 96(1):010504, January 2006.

[22] Joshua Combes, Joseph Kerckhoff, and Mohan Sarovar. The slh framework for modeling quantum input-output networks. *Advances in Physics: X*, 2(3):784–888, 2017.

[23] Joshua Combes, Howard M Wiseman, and Kurt Jacobs. Rapid measurement of quantum systems using feedback control. *Phys. Rev. Lett.*, 100(16):160503, 2008.

[24] Joshua Combes, Howard M Wiseman, Kurt Jacobs, and Anthony J O'Connor. Rapid purification of quantum systems by measuring in a feedback-controlled unbiased basis. *Phys. Rev. A*, 82(2):022307, 2010.

[25] Joshua Combes, Howard M Wiseman, and Andrew James Scott. Replacing quantum feedback with open-loop control and quantum filtering. *Phys. Rev. A*, 81(2):020301, 2010.

[26] Kevin C Cox, Graham P GReve, Joshua M Weiner, and James K Thompson. Deterministic squeezed states with collective measurements and feedback. *Phys. Rev. Lett.*, 116(9):093602, 2016.

[27] Ben Criger, Alessandro Ciani, and David P DiVincenzo. Multi-qubit joint measurements in circuit qed: stochastic master equation analysis. *EPJ Quantum Technology*, 3(1):6, 2016.

[28] Nicola Dalla Pozza, Howard M Wiseman, and Elanor H Huntington. Deterministic preparation of superpositions of vacuum plus one photon by adaptive homodyne detection: experimental considerations. *New Journal of Physics*, 17(1):013047, 2015.

[29] Christopher M Dawson and Michael A Nielsen. The solovay-kitaev algorithm. *arXiv preprint quant-ph/0505030*, 2005.

[30] Justin Dressel, Todd A Brun, and Alexander N Korotkov. Implementing generalized measurements with superconducting qubits. *Physical Review A*, 90(3):032302, 2014.

[31] Bryan Eastin and Emanuel Knill. Restrictions on transversal encoded quantum gate sets. *Physical review letters*, 102(11):110502, 2009.

[32] Andrew Eddins, Sydney Schreppler, David M Toyli, Leigh S Martin, Shay Hacohen-Gourgy, Luke CG Govia, Hugo Ribeiro, Aashish A Clerk, and Irfan Siddiqi. Stroboscopic qubit measurement with squeezed illumination. *Physical review letters*, 120(4):040505, 2018.

[33] Andrew Wilson Eddins. *Superconducting Circuits for Quantum Metrology with Nonclassical Light*. University of California, Berkeley, 2017.

[34] Alec Eickbusch, Steven Touzard, Phillipe Campagne-Ibarcq, Evan Zalys-Geller, Nicholas Frattini, Volodymyr Sivak, Shruti Puri, Mazyar Mirrahimi, Shyam Shankar, and Michel Devoret. Grid states for encoding and stabilizing a logical qubit in superconducting circuits (part 1). *Bulletin of the American Physical Society*, 2019.

[35] Paolo Facchi and Saverino Pascazio. Quantum zeno dynamics: mathematical and physical aspects. *Journal of Physics A: Mathematical and Theoretical*, 41(49):493001, 2008.

[36] Jay Gambetta, Alexandre Blais, Maxime Boissonneault, Andrew A Houck, DI Schuster, and Steven M Girvin. Quantum trajectory approach to circuit qed: Quantum jumps and the zeno effect. *Physical Review A*, 77(1):012112, 2008.

[37] Crispin Gardiner, Peter Zoller, and Peter Zoller. *Quantum noise: a handbook of Markovian and non-Markovian quantum stochastic methods with applications to quantum optics*, volume 56. Springer Science & Business Media, 2004.

[38] Marco G Genoni, Stefano Mancini, and Alessio Serafini. Optimal feedback control of linear quantum systems in the presence of thermal noise. *Phys. Rev. A*, 87(4):042333, 2013.

[39] Christopher Gerry, Peter Knight, and Peter L Knight. *Introductory quantum optics*. Cambridge university press, 2005.

[40] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Quantum metrology. *Physical review letters*, 96(1):010401, 2006.

[41] Steven M Girvin. Superconducting qubits and circuits: Artificial atoms coupled to microwave photons. *Lectures delivered at Ecole d'Eté Les Houches*, 2011.

[42] Sebastien Gleyzes, Stefan Kuhr, Christine Guerlin, Julien Bernu, Samuel Deleglise, Ulrich Busk Hoff, Michel Brune, Jean-Michel Raimond, and Serge Haroche. Quantum jumps of light recording the birth and death of a photon in a cavity. *Nature*, 446(7133):297, 2007.

[43] Vittorio Gorini, Andrzej Kossakowski, and Ennackal Chandy George Sudarshan. Completely positive dynamical semigroups of n-level systems. *Journal of Mathematical Physics*, 17(5):821–825, 1976.

[44] Daniel Gottesman, Thomas Jennewein, and Sarah Croke. Longer-baseline telescopes using quantum repeaters. *Physical review letters*, 109(7):070503, 2012.

[45] Sandeep K Goyal, B Neethi Simon, Rajeev Singh, and Sudhavathani Simon. Geometry of the generalized bloch sphere for qutrits. *Journal of Physics A: Mathematical and Theoretical*, 49(16):165203, 2016.

[46] Lov K Grover. A fast quantum mechanical algorithm for database search. *arXiv preprint quant-ph/9605043*, 1996.

[47] Shay Hacohen-Gourgy, Luis Pedro García-Pintos, Leigh S Martin, Justin Dressel, and Irfan Siddiqi. Incoherent qubit control using the quantum zeno effect. *Physical review letters*, 120(2):020505, 2018.

[48] Shay Hacohen-Gourgy, Leigh S Martin, Emmanuel Flurin, Vinay V Ramasesh, K Birgitta Whaley, and Irfan Siddiqi. Quantum dynamics of simultaneously measured non-commuting observables. *Nature*, 538(7626):491, 2016.

[49] MJW Hall and IG Fuss. Quantum phase detection and digital communication. *Quantum Optics: Journal of the European Optical Society Part B*, 3(3):147, 1991.

[50] Nicole Yunger Halpern, Brian Swingle, and Justin Dressel. Quasiprobability behind the out-of-time-ordered correlator. *Physical Review A*, 97(4):042105, 2018.

[51] Serge Haroche and J-M Raimond. *Exploring the quantum: atoms, cavities, and photons.* Oxford university press, 2006.

[52] Bas Hensen, H Bernien, AE Dréau, A Reiserer, N Kalb, MS Blok, J Ruitenberg, RFL Vermeulen, RN Schouten, C Abellán, et al. Loophole-free bell inequality violation using electron spins separated by 1.3 kilometres. *Nature*, 526(7575):682–686, 2015.

[53] Charles Hill and Jason Ralph. Weak measurement and control of entanglement generation. *Phys. Rev. A*, 77:014305, 2008.

[54] Kohjiro Iwasawa, Kenzo Makino, Hidehiro Yonezawa, Mankei Tsang, Aleksandar Davidovic, Elanor Huntington, and Akira Furusawa. Quantum-limited mirror-motion estimation. *Physical review letters*, 111(16):163602, 2013.

[55] K Jacobs and PL Knight. Linear quantum trajectories: Applications to continuous projection measurements. *Phys. Rev. A*, 57(4):2301, 1998.

[56] K. Jacobs and A. Shabani. Quantum feedback control: how to use verification theorems and viscosity solutions to find optimal protocols. *Contemp. Phys.*, 49(6):435–448, 2008.

[57] Kurt Jacobs. How to project qubits faster using quantum feedback. *Phys. Rev. A*, 67(3), March 2003.

[58] Kurt Jacobs and Daniel A Steck. A straightforward introduction to continuous quantum measurement. *Contemp. Phys.*, 47:279, 2006.

[59] Peter D Johnson, Francesco Ticozzi, and Lorenza Viola. Exact stabilization of entangled states in finite time by dissipative quantum circuits. *Phys. Rev. A*, 96(1):012308, 2017.

[60] Richard Jozsa. Fidelity for mixed quantum states. *J. Mod. Opt.*, 41(12):2315–2323, 1994.

[61] Emanuel Knill, Raymond Laflamme, and Gerald J Milburn. A scheme for efficient quantum computation with linear optics. *nature*, 409(6816):46, 2001.

[62] Jens Koch, M Yu Terri, Jay Gambetta, Andrew A Houck, DI Schuster, J Majer, Alexandre Blais, Michel H Devoret, Steven M Girvin, and Robert J Schoelkopf. Charge-insensitive qubit design derived from the cooper pair box. *Physical Review A*, 76(4):042319, 2007.

[63] Shingo Kono, Kazuki Koshino, Yutaka Tabuchi, Atsushi Noguchi, and Yasunobu Nakamura. Quantum non-demolition detection of an itinerant microwave photon. *Nature Physics*, page 1, 2018.

[64] Anthony J Leggett. Macroscopic quantum systems and the quantum theory of measurement. *Progress of Theoretical Physics Supplement*, 69:80–100, 1980.

[65] Raphaël Lescanne, Samuel Deléglise, Emanuele Albertinale, Ulysse Réglade, Thibault Capelle, Edouard Ivanov, Thibaut Jacqmin, Zaki Leghtas, and Emmanuel Flurin. Detecting itinerant microwave photons with engineered non-linear dissipation. *arXiv preprint arXiv:1902.05102*, 2019.

[66] EM Levenson-Falk, F Kos, R Vijay, L Glazman, and I Siddiqi. Single-quasiparticle trapping in aluminum nanobridge josephson junctions. *Physical review letters*, 112(4):047002, 2014.

[67] Hanhan Li, Alireza Shabani, Mohan Sarovar, and Birgitta K Whaley. Optimality of qubit purification protocols in the presence of imperfections. *Phys. Rev. A*, 87:032334, 2013.

[68] Goran Lindblad. On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics*, 48(2):119–130, 1976.

[69] Yanbing Liu and Andrew A Houck. Quantum electrodynamics near a photonic bandgap. *Nature Physics*, 13(1):48, 2017.

[70] F. London. über die jacobischen transformationen der quantenmechanik. *Zeitschrift für Physik*, 37:915, 1926.

[71] F. London. winkelvariable und kanonische transformationen in der undulationsmechanik. *Zeitschrift für Physik*, 40:193, 1927.

[72] Justin Dressel Luis Pedro García-Pintos. Probing quantumness with joint continuous measurements of non-commuting qubit observables. *arXiv preprint arXiv:1606.07934*, 2016.

[73] Lorenzo MacCone and Arun K. Pati. Stronger uncertainty relations for all incompatible observables. *Phys. Rev. Lett.*, 113(26):1–5, 2014.

[74] Chris Macklin, K O'Brien, D Hover, ME Schwartz, V Bolkhovsky, X Zhang, WD Oliver, and I Siddiqi. A near–quantum-limited josephson traveling-wave parametric amplifier. *Science*, 350(6258):307–310, 2015.

[75] Stefano Mancini and Howard M Wiseman. Optimal control of entanglement via quantum feedback. *Phys. Rev. A*, 75(1):012330, 2007.

[76] Vladimir E Manucharyan, Jens Koch, Leonid I Glazman, and Michel H Devoret. Fluxonium: Single cooper-pair circuit free of charge offsets. *Science*, 326(5949):113–116, 2009.

[77] Leigh Martin, Felix Motzoi, Hanhan Li, Mohan Sarovar, and K Birgitta Whaley. Deterministic generation of remote entanglement with active quantum feedback. *Phys. Rev. A*, 92(6):062321, 2015.

[78] Leigh Martin, Mahrud Sayrafi, and K Birgitta Whaley. What is the optimal way to prepare a bell state using measurement and feedback? *Quantum Science and Technology*, 2(4):044006, 2017.

[79] Leigh S. Martin, Emmanuel Flurin, Song Zhang, Shay Hacohen-Gourgy, Daniel Burgarth, and K. Birgitta Whaley. A multi-qubit quantum gate using the zeno effect. *in preparation*.

[80] Leigh S Martin, William P Livingston, Shay Hacohen-Gourgy, Howard M Wiseman, and Irfan Siddiqi. Implementation of a canonical phase measurement with quantum feedback. Submitted, 2019.

[81] Leigh S Martin and K Birgitta Whaley. Single-shot deterministic entanglement between non-interacting systems with linear optics. *arXiv preprint arXiv:1912.00067*, 2019.

[82] John M Martinis, Michel H Devoret, and John Clarke. Experimental tests for the quantum behavior of a macroscopic degree of freedom: The phase difference across a josephson junction. *Physical Review B*, 35(10):4682, 1987.

[83] John M Martinis, S Nam, J Aumentado, and C Urbina. Rabi oscillations in a large josephson-junction qubit. *Physical review letters*, 89(11):117901, 2002.

[84] ZK Minev, SO Mundhada, Shyam Shankar, Philip Reinhold, Ricardo Gutiérrez-Jáuregui, RJ Schoelkopf, Mazyar Mirrahimi, HJ Carmichael, and MH Devoret. To catch and reverse a quantum jump mid-flight. *arXiv preprint arXiv:1803.00545*, 2018.

[85] ZK Minev, K Serniak, IM Pop, Zaki Leghtas, K Sliwa, Michael Hatridge, Luigi Frunzio, Robert J Schoelkopf, and MH Devoret. Planar multilayer circuit quantum electrodynamics. *Physical Review Applied*, 5(4):044021, 2016.

[86] Zlatko K Minev. Catching and reversing a quantum jump mid-flight. *arXiv preprint arXiv:1902.10355*, 2019.

[87] Carol Gray Montgomery, Robert Henry Dicke, and Edward M Purcell. *Principles of microwave circuits*. Number 25. Iet, 1987.

[88] JE Mooij, TP Orlando, L Levitov, Lin Tian, Caspar H Van der Wal, and Seth Lloyd. Josephson persistent-current qubit. *Science*, 285(5430):1036–1039, 1999.

[89] Felix Motzoi, K Birgitta Whaley, and Mohan Sarovar. Continuous joint measurement and entanglement of qubits in remote cavities. *Phys. Rev. A*, 92(3):032308, 2015.

[90] K W Murch, S J Weber, C Macklin, and I. Siddiqi. Observing single quantum trajectories of a superconducting quantum bit. *Nature*, 502:211, 2013.

[91] KW Murch, U Vool, D Zhou, SJ Weber, SM Girvin, and I Siddiqi. Cavity-assisted quantum bath engineering. *Physical review letters*, 109(18):183602, 2012.

[92] M. a. Nielsen. Conditions for a class of entanglement transformations. *Phys. Rev. Lett.*, 83(2), 1998.

[93] M. A. Nielsen and I. L. Chuang. *Quantum computation and quantum information*. Cambridge University Press, 2010.

[94] Michael A Nielsen. A simple formula for the average gate fidelity of a quantum dynamical operation. *Physics Letters A*, 303(4):249–252, 2002.

[95] Michael A Nielsen. Cluster-state quantum computation. *Reports on Mathematical Physics*, 57(1):147–161, 2006.

[96] Michael Martin Nieto. Quantum phase and quantum phase operators: some physics and some history. *Physica Scripta*, 1993(T48):5, 1993.

[97] Atsushi Nishizawa and Yanbei Chen. Universally valid error-disturbance relations in continuous measurements. *arXiv preprint arXiv:1506.00304*, 2015.

[98] Kevin O'Brien, Chris Macklin, Irfan Siddiqi, and Xiang Zhang. Resonant phase matching of josephson junction traveling wave parametric amplifiers. *Physical review letters*, 113(15):157001, 2014.

[99] Bernt Øksendal. *Stochastic Differential Equations*. Spinger-Verlag, 2003.

[100] Ognyan Oreshkov and Todd A Brun. Weak measurements are universal. *Physical review letters*, 95(11):110409, 2005.

[101] Román Orús. A practical introduction to tensor networks: Matrix product states and projected entangled pair states. *Annals of Physics*, 349:117–158, 2014.

[102] Hanhee Paik, DI Schuster, Lev S Bishop, G Kirchmair, G Catelani, AP Sears, BR Johnson, MJ Reagor, L Frunzio, LI Glazman, et al. Observation of high coherence in josephson junction qubits measured in a three-dimensional circuit qed architecture. *Physical Review Letters*, 107(24):240501, 2011.

[103] Luca Pezze and Augusto Smerzi. Quantum theory of phase estimation. *arXiv preprint arXiv:1411.5164*, 2014.

[104] David M Pozar. *Microwave engineering*. John Wiley & Sons, 2009.

[105] John Preskill. Quantum computing in the nisq era and beyond. *arXiv preprint arXiv:1801.00862*, 2018.

[106] Jean-Michel Raimond, Paolo Facchi, Bruno Peaudecerf, Saverio Pascazio, Clément Sayrin, Igor Dotsenko, Sébastien Gleyzes, Michel Brune, and Serge Haroche. Quantum zeno dynamics of a field in a cavity. *Physical Review A*, 86(3):032120, 2012.

[107] TC Ralph, AP Lund, and HM Wiseman. Adaptive phase measurements in linear optical quantum computation. *Journal of Optics B: Quantum and Semiclassical Optics*, 7(10):S245, 2005.

[108] R-P Riwar, Amin Hosseinkhani, Luke D Burkhart, Yvonne Y Gao, Robert J Schoelkopf, Leonid I Glazman, and Gianluigi Catelani. Normal-metal quasiparticle traps for superconducting qubits. *Physical Review B*, 94(10):104516, 2016.

[109] N Roch, M E Schwartz, F Motzoi, C Macklin, R. Vijay, A W Eddins, A. N. Korotkov, K. B. Whaley, M. Sarovar, and I. Siddiqi. Observation of Measurement-Induced Entanglement and Quantum Trajectories of Remote Superconducting Qubits. *Phys. Rev. Lett.*, 112(17):170501, April 2014.

[110] Pierre Rouchon and Jason F Ralph. Efficient quantum filtering for quantum feedback control. *Physical Review A*, 91(1):012118, 2015.

[111] Tanay Roy, Suman Kundu, Madhavi Chand, AM Vadiraj, A Ranadive, N Nehra, Meghan P Patankar, J Aumentado, AA Clerk, and R Vijay. Broadband parametric amplification with impedance engineering: Beyond the gain-bandwidth product. *Applied Physics Letters*, 107(26):262601, 2015.

[112] Terry Rudolph. Why i am optimistic about the silicon-photonic route to quantum computing. *APL Photonics*, 2(3):030901, 2017.

[113] R. Ruskov, J. Combes, K. Mølmer, and H. M. Wiseman. Qubit purification speed-up for three complementary continuous measurements. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 370(1979):5291–5307, 2012.

[114] Rusko Ruskov, Alexander N. Korotkov, and Klaus Mølmer. Qubit state monitoring by measurement of three complementary observables. *Phys. Rev. Lett.*, 105(10):1–4, 2010.

[115] Daniel Sank, Zijun Chen, Mostafa Khezri, J Kelly, R Barends, B Campbell, Y Chen, B Chiaro, A Dunsworth, A Fowler, et al. Measurement-induced state transitions in a superconducting qubit: Beyond the rotating wave approximation. *Physical review letters*, 117(19):190503, 2016.

[116] Mohan Sarovar and K Birgitta Whaley. Adaptive homodyne phase discrimination and qubit measurement. *Phys. Rev. A*, 76:052316, 2007.

[117] Florian Schäfer, Ivan Herrera, Shahid Cherukattil, Cosimo Lovecchio, Francesco Saverio Cataliotti, Filippo Caruso, and Augusto Smerzi. Experimental realization of quantum zeno dynamics. *Nature communications*, 5:3194, 2014.

[118] J Robert Schrieffer. *Theory of superconductivity*. CRC Press, 2018.

[119] David Isaac Schuster. *Circuit quantum electrodynamics*. Yale University, 2007.

[120] Marlan O Scully and M Suhail Zubairy. *Quantum optics*. AAPT, 1999.

[121] Alessio Serafini and Stefano Mancini. Determination of maximal gaussian entanglement achievable by feedback-controlled dynamics. *Phys. Rev. Lett.*, 104(22):220501, 2010.

[122] K Serniak, S Diamond, M Hays, V Fatemi, S Shankar, L Frunzio, RJ Schoelkopf, and MH Devoret. Direct dispersive monitoring of charge parity in offset-charge-sensitive transmons. *arXiv preprint arXiv:1903.00113*, 2019.

[123] K Serniak, M Hays, G de Lange, S Diamond, S Shankar, LD Burkhart, L Frunzio, M Houzet, and MH Devoret. Hot non-equilibrium quasiparticles in transmon qubits. *arXiv preprint arXiv:1803.00476*, 2018.

[124] Eyob A Sete, Jay M Gambetta, and Alexander N Korotkov. Purcell effect with microwave drive: Suppression of qubit relaxation rate. *Physical Review B*, 89(10):104516, 2014.

[125] Alireza Shabani and Kurt Jacobs. Locally optimal control of quantum systems with strong feedback. *Phys. Rev. Lett.*, 101(23):230403, 2008.

[126] Yaoyun Shi. Both toffoli and controlled-not need little help to do universal quantum computation. *arXiv preprint quant-ph/0205115*, 2002.

[127] Daniel Huber Slichter. *Quantum jumps and measurement backaction in a superconducting qubit*. University of California, Berkeley, 2011.

[128] Srinivas Sridharan, Masahiro Yanagisawa, and Joshua Combes. Optimal rotation control for a qubit subject to continuous measurement. *arXiv preprint arXiv:1211.5617*, 2012.

[129] John K Stockton, Ramon van Handel, and Hideo Mabuchi. Deterministic Dicke-state preparation with continuous measurement and control. *Phys. Rev. A*, 70:022106, 2004.

[130] Leonard Susskind and Jonathan Glogower. Quantum mechanical phase and time operator. *Physics Physique Fizika*, 1(1):49, 1964.

[131] Ewin Tang. A quantum-inspired classical algorithm for recommendation systems. *arXiv preprint arXiv:1807.04271*, 2018.

[132] Colin Teo, Joshua Combes, and Howard M Wiseman. Global versus local optimality in feedback-controlled qubit purification: new insights from minimizing rényi entropies. *New J. Phys.*, 16(10):105010, 2014.

[133] LK Thomsen, S Mancini, and HM Wiseman. Continuous quantum nondemolition feedback and unconditional atomic spin squeezing. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 35(23):4937, 2002.

[134] Francesco Ticozzi, Luca Zuccato, Peter D Johnson, and Lorenza Viola. Alternating projections methods for discrete-time stabilization of quantum states. *IEEE Transactions on Automatic Control*, 63(3):819–826, 2018.

[135] Michael Tinkham. *Introduction to superconductivity*. Courier Corporation, 2004.

[136] Steven Touzard, Alec Eickbusch, Phillipe Campagne-Ibarcq, Evan Zalys-Geller, Nicholas Frattini, Volodymyr Sivak, Shruti Puri, Mazyar Mirrahimi, Shyam Shankar, and Michel Devoret. Grid states for encoding and stabilizing a logical qubit in superconducting circuits (part 2). *Bulletin of the American Physical Society*, 2019.

[137] Frank Verstraete and Henri Verschelde. Fidelity of mixed states of two qubits. *Phys. Rev. A*, 66(2):022307, 2002.

[138] R. Vijay, C Macklin, D H Slichter, S J Weber, K W Murch, R Naik, and A N Korotkov I Siddiqi. Stabilizing Rabi oscillations in a superconducting qubit using quantum feedback. *Nature*, 490:77, 2012.

[139] U. Vool, S. Shankar, S. O. Mundhada, N. Ofek, A. Narla, K. Sliwa, E. Zalys-Geller, Y. Liu, L. Frunzio, R. J. Schoelkopf, S. M. Girvin, and Devoret M. H. Continuous quantum nondemolition measurement of the transverse component of a qubit. *arXiv:1605.08004*, 2011.

[140] Uri Vool and Michel Devoret. Introduction to quantum electromagnetic circuits. *International Journal of Circuit Theory and Applications*, 45(7):897–934, 2017.

[141] Amar C Vutha, Eliot A Bohr, Anthony Ransford, Wesley C Campbell, and Paul Hamilton. Displacement operators: the classical face of their quantum phase. *European Journal of Physics*, 39(2):025405, 2018.

[142] Z Wang, S Shankar, ZK Minev, P Campagne-Ibarcq, A Narla, and Michel H Devoret. Cavity attenuators for superconducting qubits. *Physical Review Applied*, 11(1):014031, 2019.

[143] Jia-Hua Wei, Bo Qi, Hong-Yi Dai, Jian-Hua Huang, and Ming Zhang. Deterministic generation of symmetric multi-qubit dicke states: an application of quantum feedback control. *IET Control Theory & Applications*, 9(17):2500–2505, 2015.

[144] TA Wheatley, DW Berry, H Yonezawa, D Nakane, H Arao, DT Pope, TC Ralph, HM Wiseman, A Furusawa, and EH Huntington. Adaptive optical phase estimation using time-symmetric quantum smoothing. *Physical Review Letters*, 104(9):093601, 2010.

[145] H. M. Wiseman. Quantum theory of continuous feedback. *Phys. Rev. A*, 49:2133, 1994.

[146] H. M. Wiseman. Adaptive phase measurements of optical modes: going beyond the marginal Q distribution. *Phys. Rev. Lett.*, 75:4587, 1995.

[147] H. M. Wiseman and G. J. Milburn. *Quantum measurement and control*. Cambridge University Press, 2009.

[148] Howard M Wiseman, Dominic W Berry, Stephen D Bartlett, Brendon L Higgins, and Geoffrey J Pryde. Adaptive measurements in the optical quantum information laboratory. *IEEE Journal of Selected Topics in Quantum Electronics*, 15(6):1661–1672, 2009.

[149] Howard M Wiseman and Luc Bouten. Optimality of Feedback Control Strategies for Qubit Purification. *Quantum Information Processing*, 7(2-3):71–83, April 2008.

[150] Howard Mark Wiseman and JF Ralph. Reconsidering rapid qubit purification by feedback. *New J. Phys.*, 8(6):90, 2006.

[151] William K. Wootters. Entanglement of Formation of an Arbitrary State of Two Qubits. *Phys. Rev. Lett.*, 80(10):2245–2248, March 1998.

[152] Minghui Xu. Theory of steady-state superradiance. 2016.

[153] Fei Yan, Simon Gustavsson, Archana Kamal, Jeffrey Birenbaum, Adam P Sears, David Hover, Ted J Gudmundsen, Danna Rosenberg, Gabriel Samach, Steven Weber, et al. The flux qubit revisited to enhance coherence and reproducibility. *Nature communications*, 7:12964, 2016.

[154] Hidehiro Yonezawa, Daisuke Nakane, Trevor A Wheatley, Kohjiro Iwasawa, Shuntaro Takeda, Hajime Arao, Kentaro Ohki, Koji Tsumura, Dominic W Berry, Timothy C Ralph, et al. Quantum-enhanced optical-phase tracking. *Science*, 337(6101):1514–1517, 2012.

[155] Bernard Yurke and Eyal Buks. Performance of cavity-parametric amplifiers, employing kerr nonlinearites, in the presence of two-photon loss. *Journal of lightwave technology*, 24(12):5054–5066, 2006.

[156] Song Zhang, Leigh Martin, and K. Birgitta Whaley. Locally optimal measurement-based quantum feedback with application to multi-qubit entanglement generation. *arXiv preprint arXiv:1807.02029*, 2018.