

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Continuous Feedback on Quantum Superconducting Circuits

Permalink

<https://escholarship.org/uc/item/2s508385>

Author

Livingston, William

Publication Date

2021

Peer reviewed|Thesis/dissertation

Continuous Feedback on Quantum Superconducting Circuits

by

William Livingston

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, Chair

Professor Birgitta Whaley

Assistant Professor Norman Yao

Spring 2021

Continuous Feedback on Quantum Superconducting Circuits

Copyright 2021
by
William Livingston

Abstract

Continuous Feedback on Quantum Superconducting Circuits

by

William Livingston

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Irfan Siddiqi, Chair

Quantum measurement theory describes the dynamics of a quantum system when it transfers information to an observer. Within the framework of continuous measurement, this outflow of information and resultant back-action occur at a finite rate. Continuously measured systems do not instantaneously collapse to eigenstates, but rather undergo stochastic evolution highly correlated with the observer's received information. When the measurement rate is on the same timescale as the processing speed of classical control electronics, we can perform feedback on a quantum system during its collapse process. We experimentally realize such a protocol by using an adaptively controlled quantum amplifier to implement a canonical phase measurement on a flying photon. Continuous measurements can also be used as a tool in quantum error correction. Error correction typically requires discrete rounds of measurement, often using entangling gates and ancilla qubits. In contrast, continuous error correction is implemented via always-on measurements of parity syndromes. Using direct parity measurements of pairs of qubits in a three-qubit system along with a custom classical controller, we implement the first continuous quantum error correction code.

To my family

Contents

Contents	ii
List of Figures	iv
List of Tables	xii
1 Introduction	1
1.1 Quantum measurement and feedback	1
1.2 Quantum computation and error correction	2
1.3 Superconducting qubits	3
1.4 Thesis overview	4
2 Circuit Quantum Electrodynamics	5
2.1 Classical resonators	5
2.2 Quantum circuits	16
2.3 Dispersive coupling	19
2.4 Black box quantization	22
3 Quantum Measurement	31
3.1 Discrete maps	31
3.2 Continuous measurement	34
3.3 Qubit readout with a cavity	47
4 Canonical Phase Measurement with Adaptive Feedback	62
4.1 Background on canonical phase measurements	62
4.2 Implementation of the canonical phase measurement	65
4.3 Outlook	77
5 Continuous Quantum Error Correction	78
5.1 Background on quantum error correction	78
5.2 Experimental demonstration	84
5.3 Outlook and conclusion	97

A Making a Continuous Feedback Controller	99
A.1 FPGA basics	99
A.2 A board for control	104
A.3 Writing an AWG for feedback	107
A.4 Computer interface	111
Bibliography	113

List of Figures

2.1	Lumped element oscillators. (a) Parallel LC oscillator with no loss. (b) LCR oscillator. (c) Resonator connected to a transmission line in a reflection geometry. A small input capacitance weakens the coupling to the transmission line.	7
2.2	Physical layout for two varieties of resonators with a common measurement feed-line. The left resonator is primarily a lumped element design. A meandering wire at the top acts as an inductor and the interdigitated fingers at the bottom form the capacitor. This element supports one low frequency mode and has a large frequency gap to the next resonance. On the right is a distributed element resonator, which supports many modes roughly evenly spaced in frequency. The design is that of a quarter wave (open on one end and terminated on the other).	10
2.3	(a) 4-port directional coupler with weak scattering ϵ . (b) Hanger resonator geometry. Port 3 is terminated in an open, represented by an x. Port 4 is connected to a resonant circuit and, for illustration purposes grounded on the far end. (In principle, there could be additional coupling to ground.)	12
2.4	Resonator in transmission. The central red section is the resonator. The capacitors act as beamsplitters which partially transmit and partially reflect signals with coupling constant ϵ_1 and ϵ_2 . The incoming and outgoing waves at each port are shown in blue. Although the transmission lines and the resonator often have the same impedance, they do not have to. Effects of mismatched port impedances can be absorbed into the coupling constants.	13
2.5	Summary of typical styles of resonators. The scattering parameters and resonator geometry are shown, with the resonator in red. In the middle are plots of complex sample scattering parameters for different internal kappas of the resonators. The bottom plots show the phase and Log Mag responses of the resonators as a function of detuning from resonance. Notice in that for the resonator in reflection, as the system transitions from $\kappa_i > \kappa_e$ to $\kappa_e < \kappa_i$, the phase response starts to wrap a full 2π	14
2.6	Fitting resonance data. The resonator is measured in a hangar geometry, with measured S_{21} parameters in light blue. The green dashed line is the fit guessed from the linear regression, and the red solid line is the result of the full regression using the parameters found in the linear fit as an initial guess.	16

2.7	Transmon qubit. (a) The transmon consists of a capacitor and a Josephson junction indicated by a box with an x. The junction allows a jump in BCS phase from φ_a in purple to φ_b in pink. (b) Physical images of a transmon. i) The transmon's footprint is primarily capacitor pads. A coupling resonator enters at the top right and a control line enters from the bottom. ii) Leads connecting the capacitor pads to a junction. iii) Josephson junction formed from a thin layer of aluminum oxide between two layers of aluminum. (c) Energy potential in the phase basis for $E_J/E_C = 50$, with the first six eigenfunctions plotted. Dotted lines represent the periodicity of the potential and eigenfunctions as they are defined over a phase variable.	20
2.8	Transmon in brown coupling to a resonator in black through a capacitor with coupling strength g	21
2.9	Parameters extracted for a coupled resonator-transmon system with a single junction as a function of Josephson inductance (L_J). The “Symplectic” method takes the quartic approximation in the transformed frame and the “Sym w/ND” method diagonalizes the two excitation subspace using the quartic (and quadratic) terms in the transformed frame. The bare frequencies and participation ratios for the resonator and qubit respectively are (6.379 GHz, .02) and (5.751 GHz, .965). These values correspond to a subsystem of a fabricated parity chip with a tunable qubit and a resonator of frequency 6.314 GHz. After tuning the qubit so that the system match the numerical-BBQ expected resonator-qubit detuning, we measure qubit anharmonicity α and dispersive shift χ . This single data point is plotted as a red star. A more comprehensive comparison of theory to experiment over a wide range of participation ratios may be found in [32].	30
3.1	Interaction of a quantum system and a transmission line. (a) An example of an interacting system: a cavity couples to a transmission line with decay constant γ . (b) Coupling between the system and the waveguide can be modeled as a series of harmonic oscillators which serially interact with the cavity each for time ΔT . The outgoing modes can then be measured by a detector.	34
3.2	Simulating decay under photodetection measurement. (a) Single shot decay of a cavity ($S = a$) starting in a Fock state, and starting in a coherent state. (b) Decay of a qubit ($S = \sigma$) starting in a pure state with initial population $\langle \sigma^\dagger \sigma \rangle \approx .96$ under different trajectories of a photon loss. The dashed line shows the ensemble average decay.	37
3.3	Decay under homodyne measurement. (a) Single shot decay of a cavity ($S = a$) starting in a Fock state, and starting in a coherent state. (b) Decay of a qubit ($S = \sigma$) starting in a pure state with initial population $\langle \sigma^\dagger \sigma \rangle \approx .96$ under different trajectories. For an individual trajectory, the state stays pure for the duration of the measurement.	42
3.4	Sample Wiener processes. Such processes occur in many contexts, including when integrating vacuum noise coming down a transmission line.	43

3.5	Cavity states separating during readout. The centers of the circles represent the coherent states α_j . Dynamics of the cavity states associated with each qubit eigenvector are classical in that they remain coherent states throughout the measurement. The labeled vectors on each cavity state are the different components of $\dot{\alpha}_j$. The cavity drive \mathcal{E} moves both states in the same direction, the dispersive term χ rotates the states in opposite directions, and the dissipative term κ pulls the states back towards the ground state at the center of the IQ plane. In steady state (not shown), the vectors acting on each coherent state add to zero. Maximal qubit state information is obtained by measuring along the real axis.	52
3.6	Simulation of a qubit dephasing from a measurement. Parameters used are $\chi = \kappa = 1$ MHz. (a) In-phase (real) and quadrature (imaginary) response of the cavity to a square drive pulse in the rotating frame of the drive. Only α_0 dynamics are shown, but since the drive term \mathcal{E} is real, under the cavity equations of motion in Eq. 3.78a, $\alpha_1 = \alpha_0^*$ for all time. The dotted line shows the drive \mathcal{E} in units of μs^{-1} . (b) Measurement-induced (Γ_ϕ) and irreversible dephasing ($\tilde{\Gamma}_\phi$) of the qubit during the readout pulse. The measurement induced dephasing dips below zero in the final transient, indicating coherence revival of the ρ . Inset: integrated difference $\Gamma_\phi - \tilde{\Gamma}_\phi$ remains positive throughout the readout and asymptotes to zero.	53
3.7	Geometry of backaction on the Bloch sphere. The effects of dephasing are shown in pink and pull the state towards the Z-axis. The effects of backaction are shown in blue and are tangent to the Bloch sphere (for a pure state). (a) In the case where the measurement axis is aligned with information quadrature, measurement backaction disturbs the state (partially) along the z axis as shown by the direction of the blue arrows. Assuming perfect efficiency ($\Gamma_m = 2\Gamma_\phi$), the stochastic motion repurifies the state. (b) In the case where the measurement axis is perpendicular to the information quadrature, phase backaction rotates the state in the XY plane to repurify the state. The dashed arc represents the initial XY radius of the state.	56
3.8	Simulation of readout separation using the solution in Eq. 3.94 with $\Gamma_m = 1$ μs . The red histogram shows the distribution of integrated voltages as a function of time for the qubit initialized in the excited state. The blue trace shows the same for the qubit initialized in the ground state. The separation of the histograms grows faster than their standard deviations, allowing for good readout.	57

4.1 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 FPGA. (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. 67

4.2 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 ($\phi(t) = 0$). The state only propagates in the plane of the measurement axis. (b) A single heterodyne trajectory ($\phi(t) = \omega_{\text{het}} t$). 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. 68

4.3 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 the optimal amplification phase is $\phi_{\text{opt.}} = \theta(t) + \pi/2$. Adaptive-dyne significantly outperforms the heterodyne and comes close to the ideal phase by $T = 13\mu\text{s}$. The difference $\phi_{\text{opt.}} - \phi$ is cut to lie on the interval $[-\pi/2, \pi/2]$. (C) Distribution of trajectories at $t = 10\mu\text{s}$. Due to suppression of photon-number back-action, adaptivedyne trajectories cluster in a ring at late times. In contrast, heterodyne trajectories would form a roughly spherical shape inside the Bloch sphere. (D) Statistics of the phase back-action $d\theta$ for adaptivedyne and heterodyne. On average, the phase back-action is significantly larger for adaptivedyne, indicating a stronger measurement of phase. 69

4.4 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 blue distribution for Heterodyne is behind the other histograms. The distribution is significantly narrower for adaptivedyne, indicating suppression of this information channel. 70

4.5	Pulse sequence and associated calibrations. (A) Pulse sequence of each shot of the experiment. Vertical axis represents the amplitude of each pulse, with the exception of the cooling sideband, in which the vertical axis is the quantitative cooling rate. Cavity-resonant pulses are shown in red, cavity sideband pulses in blue and qubit-resonant pulses in green. The discoloration and distortion of the central qubit pulse represent cooling-sideband-dependent frequency and amplitude modulation, as calibrated in (C) and (D) respectively. Effects are exaggerated for visual clarity. (B) Measurement of the bath engineering cooling rate versus sideband amplitude. (C) Measurement of the Stark shift induced by the sideband as a function of sideband amplitude (D) Sideband amplitude required to drive 20 MHz Rabi oscillations as a function of the sideband amplitude.	72
4.6	Internal logic block diagram for the FPGA. (A) Photon signal demodulation. The demodulation phase ϕ' is determined by the JPA's pump phase. The demodulation amplitude $P(t)$ is given by Eq. 4.16. (B) JPA pump frequency selection, advanced each trigger of the AWG: 1) adapting JPA frequency proportional to the incoming signal; 2) replaying JPA frequency from the previous AWG trigger; 3) heterodyning using a fixed JPA frequency. (C) Instantaneous JPA frequency f_{pump} increments the JPA pump phase ϕ and determines the instantaneous amplitude A through the gain calibration shown in Fig. 4.7. The real and imaginary parts of this pump tone are sent to DAC0 and DAC1 respectively for single sideband modulation. (D) The JPA pump frequency is halved, accumulated, and delayed to determine the demodulation phase ϕ' .	75
4.7	Gain calibration curve for the JPA. (A) Major axis variance (arb. units) of amplified vacuum as a function of pump frequency and amplitude. The red line represents a contour of constant variance, the gain curve. (B) Major axis variance of amplified vacuum along the gain curve with the JPA pump on as opposed to off. (C) Angle in radians of the amplified vacuum's major axis along the gain curve. An electrical delay is calibrated to keep this curve flat across frequency.	76
5.1	Density matrix of the three qubit bit-flip code in the logical basis. The logical subspace is shown in green and has syndrome 00. Error spaces comprise the three orange blocks with definite syndrome values other than 00. Physical X flips lead to transitions amongst the syndrome subspaces. Purple transitions directly out of the logical subspace move the logical state to an error subspace, but do not affect the logical qubit. Red transitions between two error spaces induce an X flip in the logical subspace in addition to changing subspaces.	82
5.2	Ancilla based three qubit error correction. Stabilizer values are copied to two ancillary qubits, which are then measured. Based on the measured syndrome, correction pulses are applied.	84

5.3	Full Parity Detection. (a) Three qubits in two cavities, with each cavity implementing a full parity measurement. Lower right: ideal phase responses of a coherent tone reflected off each cavity for different qubit states. The parity probe tones are centered on the odd-parity resonances. The phase space (IQ) plots show the ideal steady state reflected tone for the shown qubit configuration. Dashed circles are centered on all possible steady state responses. (b) Micrograph of the superconducting chip with three transmons and two joint readout resonators.	85
5.4	Error Correction. (a) Sample experimental voltage traces of the controller correcting induced bit flips. With no errors, both voltages remain positive. When an error occurs, one or both of the voltages flip and the cross thresholds, triggering the controller to send a corrective π pulse to bring the system back to the codespace. (b) Voltage responses to an induced flip on Q_0 with (blue) and without (red) feedback. Bold lines are averages and light lines are sample individual traces.	86
5.5	Characterizing the time to correct an error. a , Histogram of time between an induced error and the correction pulse for each of the qubits, normalized such the integral of the probability distribution $P_{flip}(t)$ gives the detection probability. Dashed lines indicate the dark count rates for each error type. b , Probability of detecting certain flip sequences given a flip on Q_0 preceding a flip on Q_2 . The green region is the probability of the controller correctly detecting a Q_0 flip and then a Q_2 flip. The red region is the probability of the controller detecting a Q_1 flip, resulting in a logical error. The dotted line indicates the dead time, when these two probabilities are equal. c , Population decay of the excited logical state, $ 101\rangle$, of the odd-odd subspace with and without feedback. With feedback on, the lifetime of the logical basis state is longer than that of an individual bare qubit.	88
5.6	Population lifetimes for each of the protected subspaces with and without feedback.	89

5.7	Preservation of quantum coherence. (a) Distinguishability of various state pairs in steady state readout for each measurement tone. Pairs of states in the yellow region differ in one or both of their parities. Pairs of states in the green region share their parities. Dashed lines indicate theoretically predicted distinguishability of the even eigenstates. (b) Relative state coherence after preparing a logical $ +X_L\rangle$ state in each of the plotted parity subspaces, applying parity measurement tones without feedback, and flipping one of the qubits. Coherences are normalized to results from the same sequence without the measurement tones applied. Error bars are statistical uncertainty from repeated runs of the measurement. Dashed lines indicate predicted relative dephasing due to an odd to even parity flip on each, both, or no resonators. (c) Sample coherences from preparing a logical $ +X_L\rangle$ state in the odd-odd (OO) subspace, applying an error pulse, and letting the controller correct the error. Coherences are reconstructed by time bins set by the time it takes to correct the error. Oscillations due to static ZZ coupling are visible.	90
5.8	Cryogenic wiring diagram. The Josephson parametric amplifiers (JPAs) operate in reflection, and additionally have off chip coils not shown. The JPAs also provide narrow-band gain, so when the readout chains are combined at room temperature, the combined noise at each cavity frequency is dominated by the noise amplified by that cavity's JPA. Each superconducting coil has its leads connected by a small piece of copper wire, forming a low frequency (< 1 Hz) RL filter with the coil. The room temperature wiring is also shown, but with linear elements (attenuators, amplifiers, filters, isolators) removed.	93
A.1	Inside an FPGA. (a) FPGA wire routing is shown in blue. At the intersections of the wires, there are are switches to route the signals in a configurable manner. Example processing blocks such as adders (square) and multiplexers (trapezoids) are shown in the middle of the routing. (b) Example of a "programmed" FPGA. Two bits are added together and the result is used to select one of two remaining inputs to be an output.	100
A.2	Sample hierarchy of HDL source files, which eventually get compiled onto the FPGA. Modules have inputs and outputs, and can instantiate other modules as submodules. The wires connecting to the outside get mapped to physical FPGA pins as described in a user constraint file not shown here.	101
A.3	Using Matlab simulink to write for an FPGA. Wires are made between blocks in a GUI and blocks are can be inspected to tailor their behavior.	103

A.4 Data flow through the system. Acquired data is streamed back to the computer through fast lanes using off-chip RAM as buffering memory. Fast lanes are additionally used to transfer sequencing instructions and pulse shapes to the FPGA. Extra configuration data (not to be confused with the bitstream) can be sent to FPGA registers over the slow lanes. These lanes are also used to read back status information. The FPGA Sequencer interacts sends and receives data to and from the DACs and ADCs respectively. Processed data is sent back to the computer. The ISE computer (used for bitstream compilation) programs the FPGA with the bitstream over a JTAG connector.	106
---	-----

List of Tables

2.1	Resonator-transmon Hamiltonian with dipole coupling arranged into blocks of equal excitation number, truncated to the $M \leq 2$ subspaces. Greyed out coupling terms are number non-conserving. These terms are dropped since they connect energy levels which are separated on the energy scale of the bare modes.	21
5.1	Qubit parameters	92
5.2	Resonator parameters	92

Acknowledgments

I am incredibly grateful to have been able to spend my graduate studies working in the Quantum Nanoelectronics Laboratory. I first thank my advisor, Irfan Siddiqi for his adept leadership and ability to bring together very talented groups of people. He has an excitement for teaching and it has been an unexpected pleasure to learn from him not only the physics of the field, but also a myriad of other topics as well: professional presentation, restaurant selection, and moving heavy machinery to name a few. Irfan always encourages us to think about what's next, and it has been exciting to see the lab change and grow over the time I have been here.

I would like to thank Birgitta Whaley and Norman Yao for taking the time to be on my thesis committee and to Holger Müller for being on my qual committee.

I have had the opportunity and pleasure to work and interact with a great number of people in the lab. I thank Allison Dove, who was willing to teach a young graduate student the subtle arts of fabrication. I thank Shay Hacohen-Gourgy and Emmanuel Flurin for their tireless humor as well as sparking many an interesting discussion about physics and everything else. I wish to thank Sydney Schreppler, James Colless, Kevin O'Brien and Machiel Blok for their simultaneously fun and professional attitudes in addition to their constant advice. I commend their incredible efforts in constructing new lab spaces and bringing the lab into the multi-qubit, planar chip era. I would like to give extra thanks to Kevin for passing on some of his extensive intuition for electromagnetic design, as well as for his uncanny ability to trawl through old literature for good ideas. I express thanks to Machiel in particular for his mentorship and guidance in my thesis work. Machiel often had deep understandings of physical concepts, and was always open to conversation.

While working in the lab, I overlapped with an amazing set of graduate students. I thank Andrew Eddins for his calm resolve with assurances that everything would probably be okay, and Dar Dahlen for helping the lab out of the coding dark ages. Thanks to Vinay Ramasesh for always having an insatiable curiosity about everything, and being able to stoke creativity when I was bogged down or stuck in a loop. Big thanks to Leigh Martin, for changing my graduate trajectory towards the path of quantum measurement. Leigh was an incredible lab partner, teaching me how to make hard decisions and always move an experiment forward. I would like to extend similarly large thanks to John Mark Kreikebaum, whose methodical and relentless approach to everything he touched has benefited the lab and myself enormously. His dedication to bringing a new cleanroom online and improving every aspect of fab has certainly helped make this work possible. I also thank Brad Mitchell and Marie Lu for many great discussions and their contagious positivity.

In the later years, it has been wonderful to have a new set of postdocs and graduate student to work and drink tea/coffee with: Ravi Naik, Gerwin Koolstra, Alexis Morvan, Jean-Loup Ville, Kasra Nowrouzi, Archan Banerjee, Noah Stevenson, Trevor Chistolini, Akel Hashim, and Larry Chen. Thanks to you all and it will be exciting to see where you go.

Thank you to those keeping the physics department running, especially to Carlos Bustamante and Anthony Vitan, and Stephen Raffel.

I am also thankful for my friends throughout the years and for my family. To my parents Pam Martin and Robert Livingston, thank you for, among countless things, giving me the desire to learn. To my sister Laura Livingston, thank you for being a lifelong role model.

Chapter 1

Introduction

Over the last century, physicists have worked towards finer control of quantum systems. Much of modern technology from transistors to lasers to MRIs relies on concepts from quantum mechanics. However, most of these applications relate to material properties at a macroscopic scale and rely on gross control over quantum ensembles. They consist of many degrees of freedom interacting strongly with an outside environment. Within the last few decades, it has become increasingly feasible to not just engineer bulk properties, but also exhibit control over small quantum systems. These systems can be made highly isolated from their environment, and contain a limited number of degrees of freedom. Scientists have significant control over the unitary and dissipative dynamics these systems experience. Once confined to the thought experiments of theorists, concepts such as entanglement are now regularly realized in physical laboratories. Indeed, an experiment violating Bell's inequality has verified quantum entanglement while closing all except the most contrived loopholes to get around its existence [1]. We are now at the point where we can conduct research into how to use this new quantum technology.

1.1 Quantum measurement and feedback

Recent technological advances have allowed for detailed investigation into the nature of quantum measurement. Measurement plays the fascinating role in quantum mechanics of simultaneously providing information and generating true randomness. On one hand, making measurements on a system can provide information about its prior state. On the other, even with full knowledge of the state, most measurements are only predictable with probabilistic outcomes. Indeed, for any given state, it is possible to construct a measurement whose outcome is completely uncertain and dictated by pure randomness. An intimately related core feature of quantum measurement is that the act of state collapse as dictated by the Born rule: measurement will change the state of the system, with this change perfectly correlated to the outcome (up to inefficiencies). From the bare statement of the Born rule, state collapse appears as an instantaneous process. However, physical measurements are performed with

finite interaction strengths, so are not truly instantaneous. Instead, the act of measurement can be considered to be smeared over time.

Using the Born rule to study the temporal dynamics of measurement leads to the field of quantum trajectories, also known as continuous measurement. Introduced as a numerical tool to allow theorists to simulate dissipation with fewer computer resources [2, 3, 4], it was difficult to measure these trajectories experimentally. The main challenge to studying quantum trajectories in a laboratory setting is that, although the information leaving the system may be perfectly correlated with collapse dynamics, this information can be difficult to collect. For example, consider observing the fluorescence of an atom. We would need to completely surround the atom with photodetectors just to have a chance at collecting most of the emitted photons. Without access to a substantial portion of these photons, the information needed to reconstruct a trajectory is lost.

With proper care and engineering, scientists have been able to isolate systems from the environment to the degree where they can capture a large percentage of the emitted information. For example, superconducting qubits can be coupled to a one dimensional waveguide acting as their primary window to the outside world. Such a platform has been used to experimentally track state collapse using quantum trajectories [5, 6]. Indeed, multiple measurements can be tracked at the same time [7], even if those measurements do not commute [8]!

One compelling aspect of studying quantum trajectories is that since information can be extracted as the state is collapsing, it can be used to feed back on the system before total collapse. Such feedback allows one to generate deterministic entanglement [9] and to indefinitely stabilize the coherence of a Rabi-driven qubit [10]. In the first experimental work of this thesis, we employ feedback on a measurement device during a continuous measurement to realize a canonical phase measurement.

1.2 Quantum computation and error correction

In the last few years, the technology to isolate, control, and measure qubits has improved to the point where it is has become conceivable to build large interconnected systems of qubits. Unlike the large ensembles of the past, these newer platforms are vastly more reconfigurable and programmable, with the ability to support controllable multipartite entanglement. Such technology can be used to emulate a variety of Hamiltonians under the model of “it takes a quantum system to simulate a quantum system.” One can construct and simulate exotic phases of matter, and inspect individual components in great detail. Being able to individually address every qubit and tailor qubit-qubit interactions also opens the door to digital quantum computation. Therein lies the tantalizing possibility of performing certain computations faster than could be done on a classical computer. The most famous example of such an algorithm is Shor’s algorithm, a method to factorize large numbers, with the cryptographic implication of essentially breaking RSA. However, a machine with such capabilities is still quite a distant goal.

A large hurdle in the path of quantum computation is the compounding effects of errors as the complexity of devices and algorithms grow. Unlike in a classical computer, information in a quantum computer is inherently fragile. Interactions with the environment decohere the computational state, and imperfections in applied unitaries accumulate over the course of a circuit. Although this fragility could be a devastating blow to quantum computation, the advent of error correction [11] has possibly turned the impossible into merely the extraordinarily difficult. Further developments demonstrated that if one could lower error rates sufficiently, then under certain assumptions error correction could be made fault tolerant [12]. In other words, with the ability to produce an arbitrary number of “good enough” qubits, one could run an algorithm arbitrarily well.

A key ingredient to quantum error correction is checking whether or not errors have occurred and correcting any that have. In order to check for errors, one has to measure certain joint properties of the qubits. When considered in standard stabilizer error correction, these measurements are treated as an instantaneous process which occur periodically during a quantum circuit. However, as previously described, measurement is inherently a continuous process. In fact, if we can construct a direct measurement of the correct operators, we can leave the measurement on at all times and implement a continuous error correction code. In the second experimental work of this thesis, we demonstrate such a protocol on three qubits to protect them against bit flip errors.

1.3 Superconducting qubits

Quantum mechanics is one of the most well tested theories, and is used as a framework underpinning most, if not all, physical processes. This universality means that there are a range of systems which can be used to construct qubits and process quantum information. Natural choices for qubits include photon polarization and spin-1/2 particles (such as electrons trapped in a quantum dot [13]). Similarly to spin-1/2 systems, certain crystal defects such as nitrogen-vacancy (NV) centers in diamonds have addressable electronic and nuclear spin degrees of freedom [14]. One can alternatively use two levels of multi-level system as a qubit so long as transitions can be selectively driven between the two chosen levels. Such systems include the electronic states of trapped ions [15] and neutral atoms [16], where the atomic potentials provide a strong non-linearity in the electronic level spacing. In most of the above examples, the physical sizes of the qubits is on the scale of atoms with a correspondingly small interaction cross-section.

An intriguing alternative for processing quantum information is to engineer qubits out of superconductors. Circuit elements made from superconductors are nearly loss-less in the microwave domain. They can be used to construct microwave resonators which can hold high quality bound photons. With the addition of Josephson junction, a non-linear circuit element, these resonators can be turned into non-linear oscillators whose lowest two levels can be treated as qubits. Superconducting circuits are typically quite large with respect to their atomic scale peers, which helps them to strongly couple to other circuit

elements. The variety of coupling strengths as well as the amount of qubit connectivity are extremely flexible, as these properties can be engineered when designing a microwave circuit. A particular style of superconducting qubit, the transmon [17], is the workhorse of both this thesis and (currently) of the superconducting field more broadly. However, there are other superconducting qubits varieties which are being explored as alternatives or complimentary partners to the transmon. These include capacitively shunted flux qubits [18, 19], fluxonium [20, 21], and $0\text{-}\pi$ qubits [22, 23].

1.4 Thesis overview

In chapter 2, we describe the classical physics associated with microwave resonators. Resonators are the basic elements of readout and can be additionally used to characterize the quality of our fabrication processes. Understanding microwave resonator physics is crucial when designing and measuring superconducting qubits. We then introduce the qubit used throughout the thesis, the transmon, and describe its coupling to a resonator. We include a primer on how to view this interaction through black-box quantization, a useful technique for simulating parameters of a physical chip.

In chapter 3, we provide background on continuous measurement. We go through the generalization of the quantum measurement beyond the standard von Neumann description. We then derive the continuous time limit of measuring the emitted field of a quantum system decaying into a transmission line. In doing so, we connect the ideas of decoherence and measurement. We additionally introduce the concept of stochastic calculus, a useful descriptor of quantum trajectories. Finally, we discuss how to use a resonator to dispersively measure a qubit and what backaction such a measurement has on the qubit state.

In chapter 4, we present our first work on integrating continuous measurement and feedback. We measure a photonic wave-packet while constantly adapting the quadrature operator measured by our Josephson parametric amplifier. By doing so, we are able to implement a canonical phase measurement, demonstrating that feedback can extend our measurement capabilities.

In chapter 5, we demonstrate the first experimental realization of a continuous quantum error correction code. We start the chapter with a review of how information is stored in a quantum error correction and how it is protected against errors. We also describe the ancilla-based approach of performing discrete error correction. In the main body of the chapter, we describe how engineer a direct full ZZ parity measurement between two qubits. Using two of these measurements, we implement a three qubit continuous error correction code. Under this code, we are able to correct bit flips as they are detected and, in doing so, extend the logical T_1 lifetimes of our logical qubit past the lifetimes of the comprising qubits.

Chapter 2

Circuit Quantum Electrodynamics

Quantum electrodynamics (QED) describes the interaction between light and matter. Atoms in free space have consistent coupling to incoming and outgoing photons. In order to enhance or suppress these interactions, atoms can be placed inside of optical cavities which bunch the local density of optical modes near the cavity resonances. The physics of these modified interactions between atoms and light in a cavity is known as cavity quantum electrodynamics (cQED).

When making microwave superconducting qubits, their coupling to the environment is highly controllable. Through variations in coupling capacitance and inductance, one can engineer a range of interaction strengths between these qubits and microwave cavities. In fact, their interaction strengths can be made relatively much larger than those in their tiny atomic counterparts. These microwave qubits became known as artificial atoms due to their anharmonic energy levels, and their interactions with cavities became known as circuit quantum electrodynamics (another cQED).

2.1 Classical resonators

One of the most basic elements of a quantum microwave circuit is the microwave resonator. Resonators can be used to readout the state of qubits, mediate interactions between qubits, and even act as qubits themselves [24, 25]. However, in the absence of a non-linear circuit element or a non-linear microwave drive, resonators will respond classically (i.e. they will be describable using a coherent state). Even in this mode of linear response, resonators are useful for characterizing material properties such as microwave loss and probing the local electromagnetic environment. When just considering the linear response of a resonator, it is convenient to measure them in continuous wave, where a single frequency tone is sent into a microwave network and recollected to reconstruct a scattering parameter. By fitting this scattering parameter as a function of frequency, one can extract information about the resonator such as the center frequency, strength of coupling to the drive line, and internal loss [26]. The exact details of this resonance feature will depend on the coupling geometry,

but here we present the structure of the three most common couplings (reflection; hanger; and transmission) as well as other details about measuring resonators.

Resonator impedance

To begin, we determine the impedance of a resonant circuit uncoupled to a microwave network. The continuous wave dynamics of a resonator can be captured by considering a parallel LC circuit as shown in Fig. 2.1a. The impedance of this circuit from probe to ground is:

$$Z_{LC}(\omega) = \left(\frac{1}{i\omega L} + i\omega C \right)^{-1} \quad (2.1)$$

This circuit has a resonance feature at $\Omega_0 = 1/\sqrt{LC}$. Near this resonance ($|\omega - \Omega_0| \ll \Omega_0$), the impedance is approximately

$$Z_{LC}(\omega) = \frac{1}{2iC(\omega - \Omega_0)}. \quad (2.2)$$

This equation is in the form of a Möbius transformation, which are functions $z \rightarrow w$ of the form

$$w = \frac{az + b}{cz + d}. \quad (2.3)$$

These transformations have many nice properties, chiefly that they map every circle in the complex plane to another circle in the complex plane (with straight lines being counted as circles going through complex infinity). For example, $Z_{LC}(\omega)$ maps the real line in complex frequency space (ω) to the imaginary line in complex impedance space (Z_{LC}). Additionally, Möbius transformations form a group under composition, so that the composition of two Möbius transformations is a third Möbius transformation. The matrix representation of the Möbius group allows for formulaic composition and is useful when doing these sorts of calculations. As a trivial example, consider adding a resistor in parallel to the LC oscillator as shown in Fig. 2.1b. This new impedance, $Z_{LCR} = (1/Z_{LC} + 1/R)^{-1}$ is a Möbius transformation of Z_{LC} , and is thus also a Möbius transformation of ω :

$$Z_{LCR}(\omega) = \frac{1}{2iC(\omega - \Omega_0) + R^{-1}} \quad (2.4a)$$

$$= \frac{1}{2iC(\omega - \tilde{\Omega}_0)} \quad (2.4b)$$

Here, we define $\tilde{\Omega}_0 = \Omega_0 + i\kappa_i/2$ as the general complex resonance frequency, with κ_i called the internal linewidth of the oscillator. In our LCR example, $\kappa_i = 1/RC$. In physical systems, internal losses comes from many mechanisms, such as dielectric and radiative losses. These losses are also often specified as a dimensionless quality factor, $Q_i = \Omega_0/\kappa_i$. Note that this form of $Z_{LCR}(\omega)$ still requires $|\omega - \Omega_0| \ll \Omega_0$, meaning that it is most useful in the limit where $Q \gg 1$.

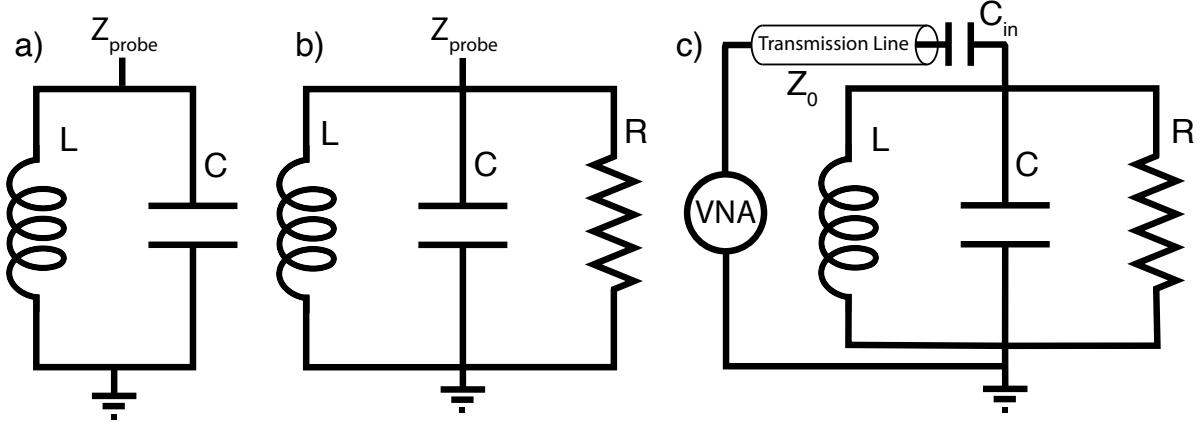


Figure 2.1: Lumped element oscillators. (a) Parallel LC oscillator with no loss. (b) LCR oscillator. (c) Resonator connected to a transmission line in a reflection geometry. A small input capacitance weakens the coupling to the transmission line.

Resonators in reflection

If we want to perform real measurements on a resonator, we need to couple it to a measuring device, usually as part of a microwave scattering network. In the simplest case, we have a resonator at the end of a lossless transmission line with characteristic impedance Z_0 . We can then use an impedance-matched Vector Network Analyzer (VNA) to measure the reflected scattering parameter Γ of voltage waves sent down the transmission line. This Γ can be written in terms of the impedance of the load Z_L and the impedance of the line [27]:

$$\Gamma(Z_L) = \frac{Z_L - Z_0}{Z_L + Z_0} \quad (2.5)$$

To add another free parameter (which will end up controlling coupling strength to the network) to the model, we take the load to be Z_{LCR} in series with a coupling impedance Z_C , such that $Z_L = Z_{LCR} + Z_C$. This coupling impedance is often taken to be a capacitor as shown in Fig. 2.1c, but we will assume the coupling impedance is more generically purely reactive and not very frequency dependent within the bandwidth of the resonance. We then get an equation for Γ ,

$$\Gamma(Z_{LCR}) = \frac{Z_{LCR} + Z_C - Z_0}{Z_{LCR} + Z_C + Z_0}. \quad (2.6)$$

Notice that Γ is a Möbius transformation of Z_{LCR} , meaning that we can write $\Gamma(\omega)$ as a Möbius transformation as well.

$$\Gamma(\omega) = \frac{Z_c - Z_0}{Z_c + Z_0} \left(\frac{\omega - (1/2iC(Z_0 - Z_c) - \tilde{\Omega}_0)}{\omega + (1/2iC(Z_0 + Z_c) - \tilde{\Omega}_0)} \right) \quad (2.7)$$

While the coefficients may look complicated, we can define new variables to capture some essential details. We will use the fact that, under the given conditions, $\overline{Z_0 + Z_c} = Z_0 - Z_c$. We can then rewrite the above in a new form, defining ϕ and v in the process.

$$\Gamma(\omega) = e^{i\phi} \left(\frac{\omega + v - \tilde{\Omega}_0}{\omega + \bar{v} - \tilde{\Omega}_0} \right) \quad (2.8)$$

Since the line impedance is positive, $\Im[v]$ is guaranteed to be positive, and we define the external linewidth as $\kappa_e = 2\Im[v]$. The resonance in the scattering parameter gets shifted with respect to the uncoupled resonator due to the coupling impedance. However, if the coupling impedance is high (for example because of a small capacitor), this effect will be small. Defining $\omega_0 = \Omega_0 - \Re[v]$, we now have

$$\Gamma(\omega) = e^{i\phi} \left(\frac{\omega - (\omega_0 + \frac{i}{2}(\kappa_i - \kappa_e))}{\omega - (\omega_0 + \frac{i}{2}(\kappa_i + \kappa_e))} \right) = e^{i\phi} \left(\frac{\omega - \tilde{\omega}_z}{\omega - \tilde{\omega}_0} \right). \quad (2.9)$$

We define $\tilde{\omega}_z$ and $\tilde{\omega}_0$ as seen directly in Eq. 2.9. Recognizing this equation is in the form of a Möbius transformation as described in Eq. 2.3, we see that $\Gamma(\omega)$ forms a circle in the complex scattering plane. Furthermore, if $\kappa_e \gg \kappa_i$, then $|\Gamma| \approx 1$, indicating that little energy is lost in the reflection. If $\kappa_e \ll \kappa_i$, then the resonance is a very small circle, as the resonator is dissipating faster than it can be pumped. Scattering parameters for different κ_e, κ_i regimes are shown in Fig. 2.5.

Waveguide resonators

The behavior of a lumped element resonator (schematically shown Fig. 2.1c and visually shown Fig. 2.2, left) can be replicated by a section of transmission line acting as a distributed element resonator (Fig. 2.2, right). In contrast to the theoretical lumped element circuit model which supports a single electromagnetic mode, a theoretical distributed element resonator supports many modes roughly even spaced apart.¹ As an example, we will discuss a resonator terminated in a short to ground through a very small resistance to provide internal loss. The resonator has impedance Z_0 . The shorted end of the resonator has the reflection coefficient $\Gamma \approx -1 + \epsilon_i$, where the deviation from -1 is due to the resistor. Looking at the shorted end of the resonator from the coupling end of the resonator, the reflection coefficient is $\Gamma_{res} = \Gamma e^{2i\omega\tau}$, where τ is the electrical delay of the resonator, the time for a wave to propagate from one end of the resonator to the other. The impedance of this resonator is then:

¹In practice, even lumped element circuits act as distributed element resonators at high enough frequencies; capacitors have stray inductance and inductors have stray capacitance. However, with proper design, a physical "lumped element" resonator can have a large frequency gap between the lowest mode and higher order modes.

$$Z = Z_0 \frac{1 + \Gamma e^{2i\omega\tau}}{1 - \Gamma e^{2i\omega\tau}}. \quad (2.10)$$

Resonance conditions occur near poles of this function (when $\Gamma e^{2i\omega\tau} \approx 1$), where we can define the resonance condition as $\omega = \Omega_0^{(n)} \equiv (2n + 1)\pi/2\tau$ for integer n . Such conditions imply that the resonator supports multiple modes, with the lowest mode being a quarter wavelength. Near one of these resonances, we can expand the impedance in $\omega - \Omega_0^{(n)}$.

$$Z^{(n)} = Z_0 \frac{2}{2i(\omega - \Omega_0^{(n)})\tau + \epsilon_i} \quad (2.11a)$$

$$= \frac{Z_0}{i\tau} \frac{1}{\omega - \Omega_0^{(n)} - i\epsilon_i/2\tau} \quad (2.11b)$$

$$= \frac{Z_0}{i\tau} \frac{1}{\omega - \Omega_0^{(n)} - i\kappa_i/2} \quad (2.11c)$$

$$= \frac{Z_0}{i\tau} \frac{1}{\omega - \tilde{\Omega}_0^{(n)}} \quad (2.11d)$$

We thus again get to the impedance of a resonator with complex poles $\tilde{\Omega}_0^{(n)} = \Omega_0^{(n)} + i\kappa_i/2$, similar to Eq. 2.4b. If we couple the resonator for a transmission line in reflection, we will get the same form of reflection parameter as in Eq. 2.9.

Resonators terminating a network

More generally, suppose that we have some scattering network, S_{ij} , where i, j enumerate the network's N ports. We can terminate one of the network ports with Z_{LCR} and come up with a reduced scattering matrix, T_{kl} , where k, l enumerate the reduced network's $N - 1$ ports. To do so we write the \mathbf{S} matrix in block format with l, k running from 0 to $N-1$ and separating out the N -th index. We then apply the condition that $a_0 = \Gamma b_0$, where Γ is a generic reflection parameter.

$$\begin{bmatrix} b_k \\ b_N \end{bmatrix} = \begin{bmatrix} S_{kl} & S_{kN} \\ S_{Nl} & S_{NN} \end{bmatrix} \begin{bmatrix} a_l \\ a_N \end{bmatrix} \quad (2.12)$$

$$a_N = \Gamma (S_{Nl}a_l + S_{NN}a_N) \quad (2.13a)$$

$$a_N = \frac{\Gamma S_{Nl}}{1 - \Gamma S_{NN}} a_l \quad (2.13b)$$

$$b_k = S_{kN}a_N + S_{kl}a_l = \left(\frac{\Gamma S_{kN}S_{Nl}}{1 - \Gamma S_{NN}} + S_{kl} \right) a_l \quad (2.14)$$

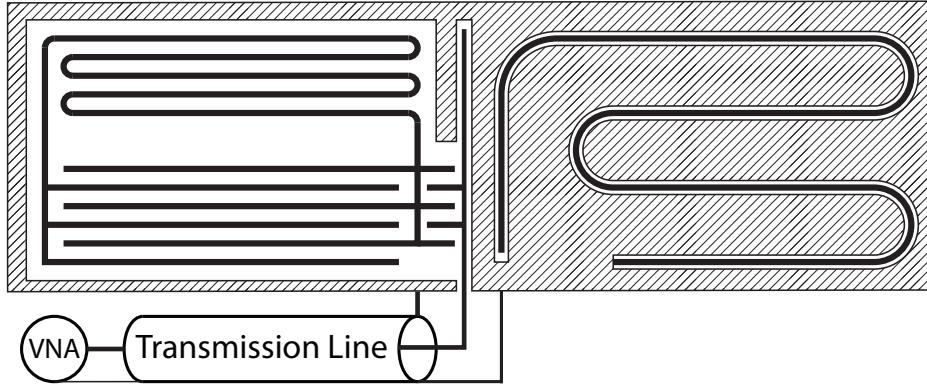


Figure 2.2: Physical layout for two varieties of resonators with a common measurement feedline. The left resonator is primarily a lumped element design. A meandering wire at the top acts as an inductor and the interdigitated fingers at the bottom form the capacitor. This element supports one low frequency mode and has a large frequency gap to the next resonance. On the right is a distributed element resonator, which supports many modes roughly evenly spaced in frequency. The design is that of a quarter wave (open on one end and terminated on the other).

We then identify the the matrix T_{kl} as the transformation such that $b_k = T_{kl}a_l$.

$$T_{kl} = \frac{\Gamma S_{kN} S_{Nl}}{1 - \Gamma S_{NN}} + S_{kl} \quad (2.15)$$

Now we consider Γ to be the reflection parameter of the Z_{LCR} resonator and assume that the scattering parameters of the original network are roughly constant over the bandwidth of the resonance. Since $T_{kl}(\Gamma)$ and $\Gamma(\omega)$ are Möbius transformations, the scattering parameters $T_{kl}(\omega)$ are Möbius transformations as well, so will show up as circles in the complex scattering plane.

Resonators in hanger geometry

When a transmission line is weakly coupled to a resonator on a third port as shown in Fig. 2.3b, this is know as a hanger resonator. One way to construct a hanger resonator is to use a directional coupler as shown in Fig. 2.3a and leave port 3 open while attaching resonator to port 4. Although this resonator may be quarter-wave, half-wave, lumped element, or other, we will simply assume that the impedance Z_{res} is in a standard form, with $A > 0$ being an arbitrary constant.

$$Z_{res} = -iZ_0 A \frac{1}{\omega - \tilde{\Omega}_0} \quad (2.16)$$

This element has the reflection coefficient

$$\Gamma = \frac{-\omega - Ai + \tilde{\Omega}_0}{\omega - Ai - \tilde{\Omega}_0}. \quad (2.17)$$

Assuming the directional coupler has a small coupling ϵ between weakly coupled ports, its scattering matrix \mathbf{V} is approximately [27]

$$\mathbf{V} = \begin{bmatrix} 0 & 1 - \frac{\epsilon^2}{2} & i\epsilon & 0 \\ 1 - \frac{\epsilon^2}{2} & 0 & 0 & i\epsilon \\ i\epsilon & 0 & 0 & 1 - \frac{\epsilon^2}{2} \\ 0 & i\epsilon & 1 - \frac{\epsilon^2}{2} & 0 \end{bmatrix}. \quad (2.18)$$

Reducing the matrix using an open on port 4 and Eq. 2.15, we get a 3-port matrix \mathbf{T} :

$$\mathbf{T} = \begin{bmatrix} 0 & 1 - \frac{\epsilon^2}{2} & i\epsilon \\ 1 - \frac{\epsilon^2}{2} & -\epsilon^2 & i\epsilon \\ i\epsilon & i\epsilon & 1 - \epsilon^2 \end{bmatrix}. \quad (2.19)$$

We then use the general reflection coefficient Γ in Eq. 2.17 to terminate port 3 to form a 2-port network \mathbf{S} , and calculate the resultant S_{12} .

$$S_{12} = \frac{(T_{23}T_{31} - T_{21}T_{33})\Gamma + T_{21}}{-T_{33}\Gamma + 1} \quad (2.20a)$$

$$= \frac{\omega - \tilde{\Omega}_0}{\omega - (\tilde{\Omega}_0 + iA\epsilon^2/2)} \quad (2.20b)$$

$$= \frac{\omega - \tilde{\Omega}_0}{\omega - \tilde{\omega}_0} = \frac{\omega - (\omega_0 + \frac{i}{2}\kappa_i)}{\omega - (\omega_0 + \frac{i}{2}(\kappa_i + \kappa_e))} \quad (2.20c)$$

Here, the external linewidth is $\kappa_e = A\epsilon^2$, and $\tilde{\omega}_0 = \omega_0 + i(\kappa_i + \kappa_e)/2$. Plots of this scattering function are shown in Fig. 2.5. When the probe tone is far off resonance, S_{12} is near unity, as the resonator does not get populated. Interestingly, on resonance in the $\kappa_i \rightarrow \infty$ limit, the transmission goes to zero. This is the effect of the resonator being highly excited and emitting a coherent tone which destructively interferes with the transmitted signal. In a real experiment, there are other effects such as impedance mismatches in the line which modify the S_{12} circle [26]. In these cases, $\Re[\tilde{\Omega}_0]$ gets shifted with respect to $\Re[\tilde{\omega}_0]$, resulting in a circle that still goes through unity in the scattering plane, but is tilted away from the real axis. Formally, this shift is introduced by making κ_e complex.

Resonators in transmission

As a final example of basic resonator-waveguide couplings, consider a resonator in transmission, as shown in Fig. 2.4. In this geometry, there is no direct coupling from the input port of

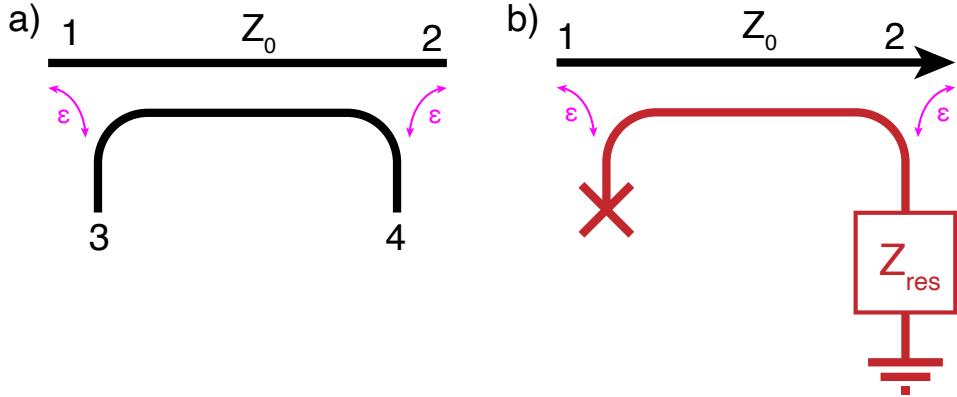


Figure 2.3: (a) 4-port directional coupler with weak scattering ϵ . (b) Hanger resonator geometry. Port 3 is terminated in an open, represented by an x. Port 4 is connected to a resonant circuit and, for illustration purposes grounded on the far end. (In principle, there could be additional coupling to ground.)

the resonator to the output port. This network has two high reflective ports (parameterized by transmission coefficients ϵ_1, ϵ_2) joined by a section of lossy transmission line (parameterized by loss ϵ_i). We solve the following scattering equations for the output scattering parameter d_o .

$$\begin{bmatrix} a_o \\ b_o \end{bmatrix} = \begin{bmatrix} 1 - \frac{\epsilon_1^2}{2} & i\epsilon_1 \\ i\epsilon_1 & 1 - \frac{\epsilon_1^2}{2} \end{bmatrix} \begin{bmatrix} a_i = 1 \\ b_i \end{bmatrix} \quad (2.21a)$$

$$\begin{bmatrix} a_c \\ d_o \end{bmatrix} = \begin{bmatrix} 1 - \frac{\epsilon_2^2}{2} & i\epsilon_2 \\ i\epsilon_2 & 1 - \frac{\epsilon_2^2}{2} \end{bmatrix} \begin{bmatrix} c_i \\ d_i = 0 \end{bmatrix} \quad (2.21b)$$

$$b_i = c_o e^{i\omega\tau} (1 - \epsilon_i) \quad (2.21c)$$

$$c_i = b_o e^{i\omega\tau} (1 - \epsilon_i) \quad (2.21d)$$

We then get an equation for the transmission $S_{21} = d_o$ as a function of frequency, which we can expand out near resonance conditions, $e^{i\omega\tau} \approx e^{-i\omega\tau}$, where the resonance frequencies are $\omega_0^{(n)} = n\pi/\tau$. Depending on if n is even or odd, the scattering parameter will have a different

sign.

$$S_{21} = \frac{-\epsilon_1 \epsilon_2}{e^{-i\omega\tau}(1 + \epsilon_i) - \left(1 + \frac{\epsilon_1^2 + \epsilon_2^2}{2}\right) e^{i\omega\tau}(1 - \epsilon_i)} (-1)^n \quad (2.22a)$$

$$\approx \frac{\epsilon_1 \epsilon_2}{2i\tau} \frac{1}{\omega - \omega_0 + \frac{\epsilon_1^2 + \epsilon_2^2 + 4\epsilon_i}{4i\tau}} (-1)^n \quad (2.22b)$$

$$= \frac{-i\sqrt{\kappa_1 \kappa_2}}{\omega - (\omega_0 + \frac{i}{2}(\kappa_1 + \kappa_2 + \kappa_i))} (-1)^n = \frac{-i\sqrt{\kappa_1 \kappa_2}}{\omega - \tilde{\omega}_0} (-1)^n \quad (2.22c)$$

Here the resonator has three decay paths: internal ($\kappa_i = 2\epsilon_i/\tau$); out port 1 ($\kappa_1 = \epsilon_1^2/2\tau$); and out port 2 ($\kappa_2 = \epsilon_2^2/2\tau$). Example scattering parameters are shown in Fig. 2.5. We see that transmission goes to zero far away from the resonance, and is maximal on resonance. An interesting note is that when $\kappa_1 = \kappa_2$ and $\kappa_i = 0$, the resonator transmits with unit efficiency ($|S_{21}| = 1$). Also notice that for a transmission resonator, a priori knowledge of the external couplings would be required to extract the internal linewidth.

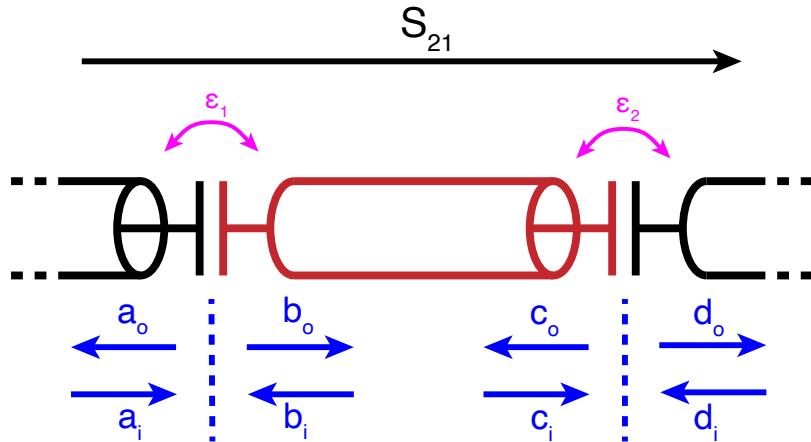


Figure 2.4: Resonator in transmission. The central red section is the resonator. The capacitors act as beamsplitters which partially transmit and partially reflect signals with coupling constant ϵ_1 and ϵ_2 . The incoming and outgoing waves at each port are shown in blue. Although the transmission lines and the resonator often have the same impedance, they do not have to. Effects of mismatched port impedances can be absorbed into the coupling constants.

Fitting resonators

As mentioned in section 2.1, the scattering parameters of a resonator follow a particular functional form regardless of the coupling geometry, namely a Möbius transformation. With

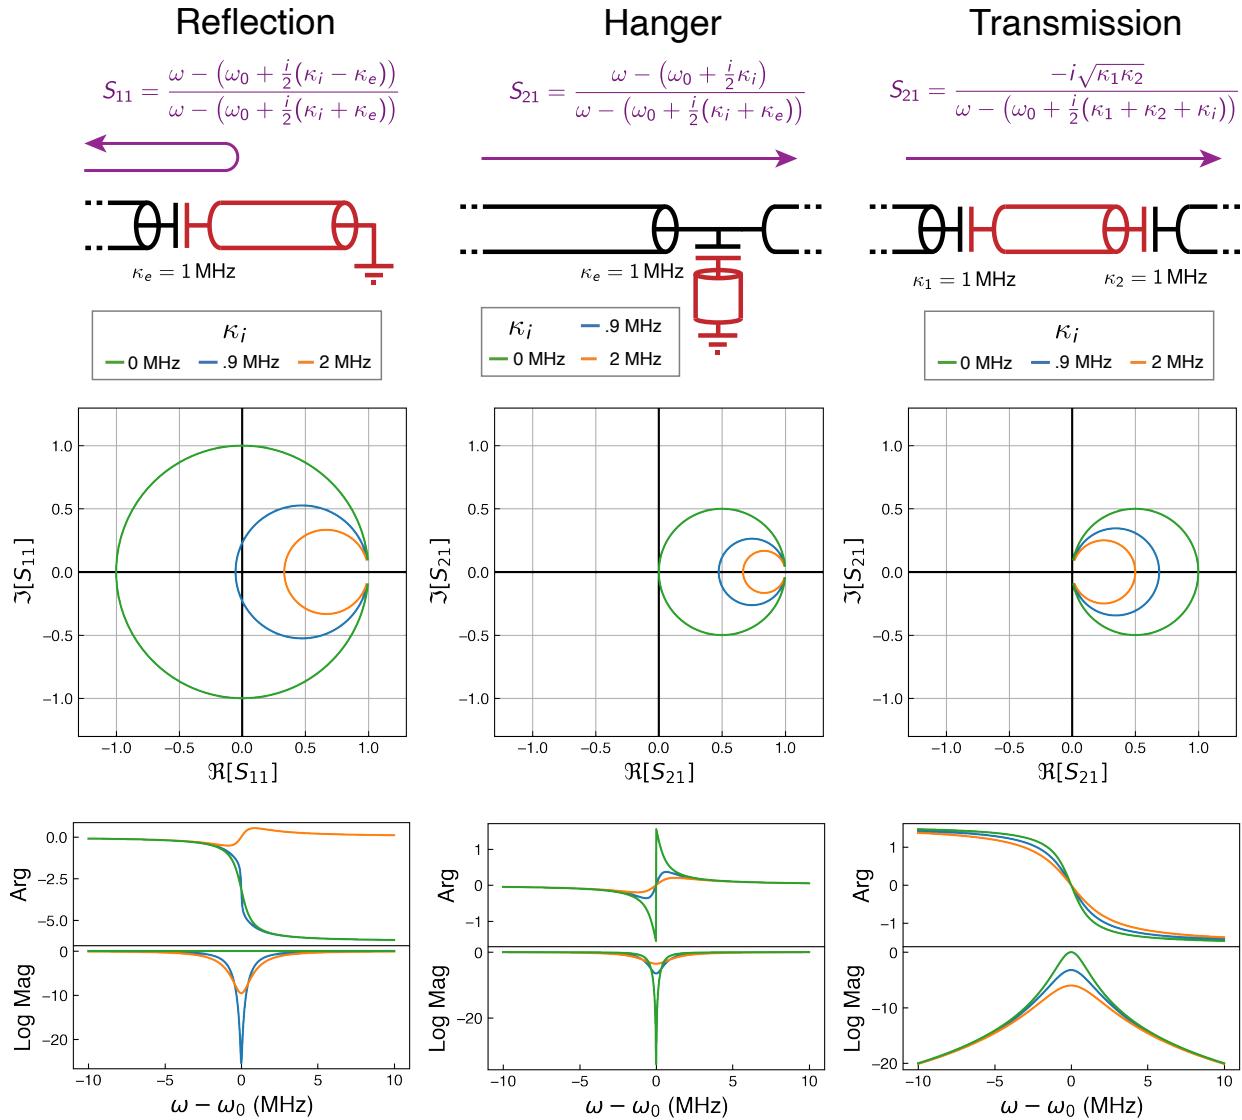


Figure 2.5: Summary of typical styles of resonators. The scattering parameters and resonator geometry are shown, with the resonator in red. In the middle are plots of complex sample scattering parameters for different internal kappas of the resonators. The bottom plots show the phase and Log Mag responses of the resonators as a function of detuning from resonance. Notice in that for the resonator in reflection, as the system transitions from $\kappa_i > \kappa_e$ to $\kappa_e < \kappa_i$, the phase response starts to wrap a full 2π .

knowledge of the coupling geometry and the parameters of this transformation, one can determine the resonance frequency and the linewidths. Furthermore, if the coupling geometry allows, the internal loss of the resonator can be separated from the external loss of the resonator. This internal loss can then be used to estimate various loss channels of the resonator and the material it is made from. The first step in this analysis process is fitting the scattering parameters to a model. If we scan over the resonator with a list of frequencies ω and get back a list of corresponding scattering parameters s , then we use least squares regression to minimize the function,

$$\sum_i |S(\omega_i) - s_i|^2. \quad (2.23)$$

In addition to the standard Möbius form, we include an additional electrical delay parameter τ into the model:

$$S(\omega) = \frac{a\omega + b}{c\omega + d} e^{i\omega\tau} \quad (2.24)$$

In practice, this delay is mostly calibrated out of the measurements such that $1/\tau$ in the resonator fit should be large as compared to the sampled bandwidth.

In order to get the regression to converge, we need good initial guesses of the complex parameters a, b, c, d (while assuming τ is close to zero). We note that, in theory, a measured scattering datum (ω_i, s_i) would satisfy $(cs_i\omega_i + ds_i = a\omega_i + b)$. We also note that a, b, c, d can be uniformly scaled without changing $S(\omega)$. Given that c will generally be non-zero we set $c = 1$ without loss of generality. Motivated by these insights, we solve a linear system of equations to get initial guesses of the fit parameters.²

$$\begin{bmatrix} \omega_0 & 1 & -s_0 \\ \vdots & \vdots & \vdots \\ \omega_N & 1 & -s_N \end{bmatrix} \begin{bmatrix} a_0 \\ b_0 \\ d_0 \end{bmatrix} = \begin{bmatrix} s_0\omega_0 \\ \vdots \\ s_N\omega_N \end{bmatrix} \quad (2.25)$$

After solving the linear equation, we can use these (a_0, b_0, d_0) as a starting point for the full regression in Eq. 2.23 along with the initial guesses $\tau_0 = 0$. The linewidths for a particular resonance geometry can then be found with the parameters of the fully regressed data. An example of this strategy working on resonator data is shown in Fig. 2.6.

Error bars on fit parameters are very useful, particularly when making claims about a high internal quality factor. In order to get error bars on desired coordinates like $\mathbf{q}_{\text{new}} = \{Q_i, Q_e, \omega_0, \tau, \dots\}$ (the other coordinates depend on the geometry), the covariance matrix calculated in the coordinates of the fit ($\mathbf{q}_{\text{old}} = \{\Re[a], \Re[b], \Re[d], \Im[a], \Im[b], \Im[d], \tau\}$) must be transformed using the Jacobian of the transformation from the original coordinates to the new ones ($\partial\mathbf{q}_{\text{new}}/\partial\mathbf{q}_{\text{old}}$). Alternatively, one can explicitly write a fitting routine for a specific resonance geometry and directly get error bars from this fit.

²One useful numerical trick to prevent singularities in this initial linear regression is to offset all the frequency values such that the center frequency is at 0 Hz, and then re-add this offset to the fit frequencies.

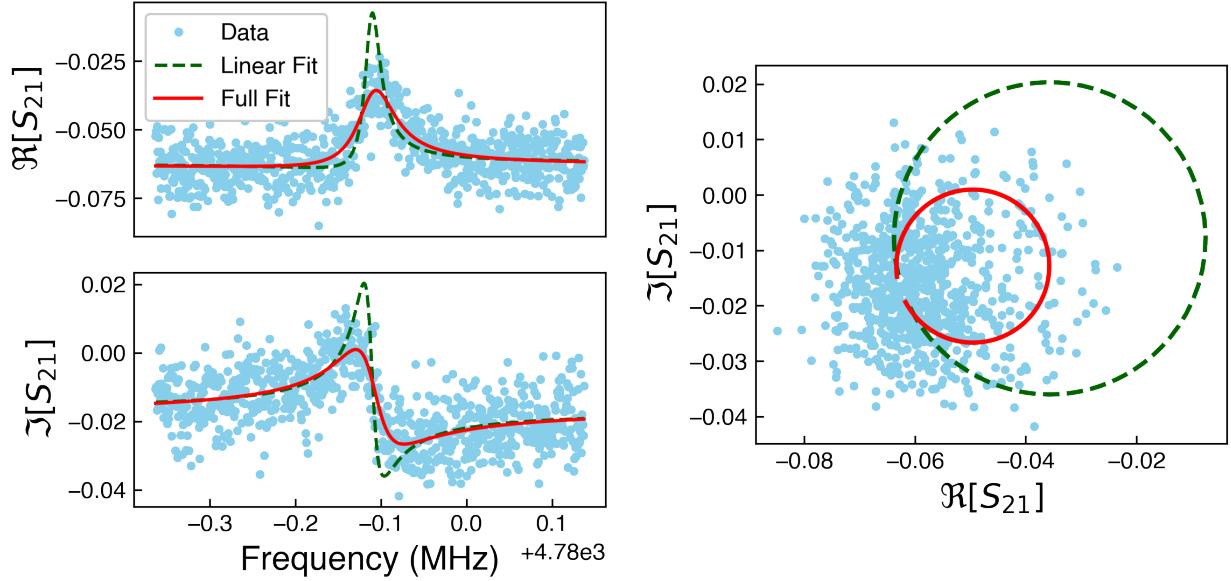


Figure 2.6: Fitting resonance data. The resonator is measured in a hangar geometry, with measured S_{21} parameters in light blue. The green dashed line is the fit guessed from the linear regression, and the red solid line is the result of the full regression using the parameters found in the linear fit as an initial guess.

2.2 Quantum circuits

So far we have discussed the classical linear response of resonators. However, in order to treat a resonator in a quantum manner, we need to quantize the circuit. Furthermore, in order to experimentally use microwave oscillators as quantum bits, we need to add a non-linearity which breaks the level spacing degeneracy. We will first summarize the Hamiltonian and supporting fields involved in a harmonic LC oscillator, and then write down the Hamiltonian for our qubit degree of freedom, the transmon.

Quantum resonators

Continuing the discussion of resonators, we will quantize their fields and write down their quantum Hamiltonian. An excellent in-depth discussion of this topic for a more general electromagnetic network can be found in a review [28], which is an update of notes from Les Houches. Here, we will present a brief summary of results as they apply to a single resonator with inductance L and capacitance C in parallel. Foreseeing the introduction of Josephson junctions as non-linear elements into our circuits, we use a generalized flux variable Φ to

describe the state of the circuit.

$$\Phi(t) = \int_{-\infty}^t V(t') dt' \quad (2.26)$$

Although such a flux can be defined across any circuit element, here Φ denotes the generalized flux across the inductor. In analogy to a mass on a spring, we chose generalized flux to act analogously to the position variable. In this analogy, the charge Q across capacitor is analogous to momentum.

$$Q(t) = \int_{-\infty}^t I(t') dt' \quad (2.27)$$

In a more general network, one may have a complex set of circuit elements, each with an associated flux and charge variable. However, one can use Kirchoff's laws to constrain these and reduce the physical degrees of freedom; in our LC example, we only have one degree of freedom despite having two circuit elements. The potential energy U and kinetic energy T of this degree of freedom are

$$U = \frac{1}{2L} \Phi^2, \quad (2.28a)$$

$$T = \frac{C}{2} \dot{\Phi}^2, \quad (2.28b)$$

resulting in the Lagrangian $\mathcal{L} = T - U$,

$$\mathcal{L} = \frac{C}{2} \dot{\Phi}^2 - \frac{1}{2L} \Phi^2. \quad (2.29)$$

One can verify that the conjugate momentum p to Φ is indeed charge across the capacitor,

$$p = \frac{\partial \mathcal{L}}{\partial \dot{\Phi}} = C \dot{\Phi} = CV(t) = Q(t), \quad (2.30)$$

as well as that the Lagrangian gives the correct equations of motion:

$$\ddot{\Phi} = \frac{1}{LC} \Phi. \quad (2.31)$$

Constructing a classical Hamiltonian from the Lagrangian through the standard transformation,

$$H_{classical} = \frac{Q^2}{2C} + \frac{\Phi^2}{2L}, \quad (2.32)$$

one can promote Φ and its conjugate variable Q to quantum operators using the commutation relation

$$[\hat{\Phi}, \hat{Q}] = i\hbar. \quad (2.33)$$

We can now write down these coordinates in terms of ladder operators, so that $\hat{H} = \omega_0 a^\dagger a$, with the resonator's impedance as $Z = \sqrt{L/C}$.

$$\hat{\Phi} = \sqrt{\frac{\hbar Z}{2}} (a + a^\dagger) \quad (2.34a)$$

$$\hat{Q} = -i \sqrt{\frac{\hbar}{2Z}} (a - a^\dagger) \quad (2.34b)$$

While the flux and charge were originally written in terms of a lumped element model, these same coordinates can be used to describe distributed, multimode systems as well. It is often helpful to think of the quantum mode as being the electromagnetic field confined by boundary conditions set by resonator's geometry rather than the current flowing in the conductor. That being said, when dealing superconducting cavities, there is extra kinetic inductance which can contribute to potential energy not stored in the magnetic field. Similarly, Josephson junctions will store potential energy not calculable from just magnetic fields.

Transmon qubits

The qubits used in this thesis and the most popular superconducting qubit at the time of writing is the transmon qubit [17]. These oscillators are weakly non-linear oscillators, such that their energy levels are not evenly spaced apart. This allows an external microwave drive to selectively drive the $|0\rangle \leftrightarrow |1\rangle$ transition without exciting the system to high levels. Transmons are constructed by taking a lumped element LC circuit and replacing the inductor with a Josephson junction. For a full derivation of the Hamiltonian, see the original paper on transmons [17]. Before introducing the Hamiltonian, we give motivation on how Josephson junctions may be viewed as non-linear inductors. A Josephson junction connects two islands of superconductors, with BCS phase φ_a and φ_b respectively. The phase difference $\varphi = \varphi_a - \varphi_b$ is proportional to the generalized flux Φ used to describe the resonator field, as can be seen by the first Josephson relation, with Φ_0 as the magnetic flux quantum:

$$\dot{\varphi} = \frac{V}{\Phi_0}, \quad (2.35a)$$

$$\varphi = \frac{1}{\Phi_0} \int_{-\infty}^t V(t') dt' \bmod 2\pi = \frac{\Phi}{\Phi_0} \bmod 2\pi. \quad (2.35b)$$

Note that the phase is restricted from 0 to 2π as it corresponds to a phase difference in the BCS wavefunction. This phase jump across the junction plays the role of changing the junction's inductance as voltage is applied. To see this, consider the second Josephson relation and differentiate it:

$$I = I_0 \sin(\varphi), \quad (2.36a)$$

$$\dot{I} = I_0 \cos(\varphi) \dot{\varphi} = \frac{I_0}{\Phi_0} \cos(\varphi) V. \quad (2.36b)$$

Comparing to the standard inductor's $V = L\dot{I}$, we have the relation

$$V = \frac{\Phi_0}{I_0 \cos(\varphi)} \dot{I}. \quad (2.37)$$

So far, although φ has represented the phase jump of the BCS phase, it has been treated as a definite non-quantum variable. Indeed one can write down a classical potential energy needed to twist the BCS phase by φ ,

$$U_J = E_J(1 - \cos(\varphi)), \quad (2.38)$$

where $E_J = \Phi_0 I_0 / 2\pi$ is the Josephson energy. However, when promoted to operator status, φ becomes canonically conjugate to the number of Cooper pairs on each island n such that \hat{n} and $\hat{\varphi}$ satisfy the commutation relation $[\hat{n}, \hat{\varphi}] = i$.³ Adding a capacitance in parallel with the junction as shown in Fig 2.7a, the full Hamiltonian for the transmon is then

$$\hat{H} = 4E_C(\hat{n} - n_g)^2 - E_J \cos(\hat{\varphi}). \quad (2.39)$$

Here, the charging energy $E_C = e^2/2C$ is given by the energy penalty of moving a Cooper pair (charge $2e$) across the capacitance C . The gate charge n_g is an offset charge present if there is an external classical electric potential which biases the steady state charge on the capacitor. The energy levels of this Hamiltonian are described by Mathieu functions [31], and are shown in Fig. 2.7 for $n_g = 0$. Although this system is more generally called the Cooper pair box, in the limit $E_J \gg E_C$, the level splittings of the system become exponentially less sensitive in n_g . The system becomes highly localized in φ and much less sensitive to charge noise in the local environment. We call devices in such a limit transmons and, by expanding the cosine term in the Hamiltonian to fourth order, can approximate their n -th energy level to be [17]

$$E_n = \sqrt{8E_c E_J} \left(n + \frac{1}{2} \right) - E_c \left(\frac{1}{2}n^2 + \frac{1}{2}n + \frac{1}{4} \right). \quad (2.40)$$

2.3 Dispersive coupling

An important coupling in cQED is that of a transmon to a resonator, as shown in Fig. 2.8. Resonators can facilitate multiqubit gates and are the primary method of reading out the state of a transmon. The relevant interaction has the form of a dipole coupling with strength g .⁴ We will approximate the transmon as an anharmonic oscillator with anharmonicity α ,

³There is a subtlety in this relation in that it requires $n \gg 0$ [29, 30]. The same subtlety arises in chapter 4 when discussing phase measurements.

⁴Although it is often convenient and useful to think of the interaction as an electric dipole-dipole coupling of the qubit-transmon, there are circumstances where this intuition breaks down for systems of multiple junctions when determining dispersive shifts. Such an example is explained in Appendix B of [32].

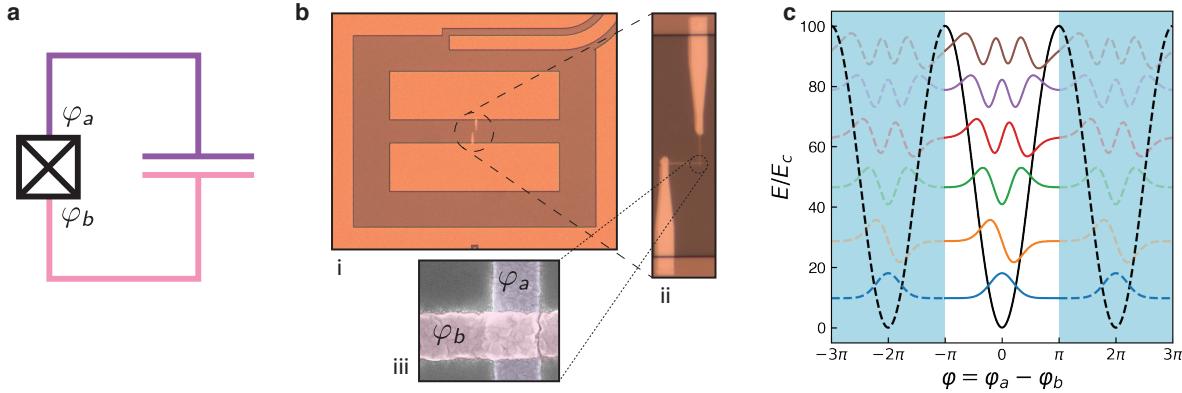


Figure 2.7: Transmon qubit. (a) The transmon consists of a capacitor and a Josephson junction indicated by a box with an x. The junction allows a jump in BCS phase from φ_a in purple to φ_b in pink. (b) Physical images of a transmon. i) The transmon's footprint is primarily capacitor pads. A coupling resonator enters at the top right and a control line enters from the bottom. ii) Leads connecting the capacitor pads to a junction. iii) Josephson junction formed from a thin layer of aluminum oxide between two layers of aluminum. (c) Energy potential in the phase basis for $E_J/E_C = 50$, with the first six eigenfunctions plotted. Dotted lines represent the periodicity of the potential and eigenfunctions as they are defined over a phase variable.

lowering operator b , and frequency ω_q . We let the resonator have lowering operator a and frequency ω_r . The coupled Hamiltonian can be written as

$$H = \omega_r a^\dagger a + \omega_q b^\dagger b + \frac{\alpha}{2} (b^\dagger)^2 b^2 + g(a + a^\dagger)(b + b^\dagger). \quad (2.41)$$

In order to use the resonator for qubit readout, the system is taken into the dispersive regime[17, 33], meaning that the qubit resonator detuning $\Delta = \omega_q - \omega_r$ is much greater than the dipole coupling ($\Delta \gg g$). Using second-order perturbation theory, we will renormalize the qubit and cavity eigenstates and acquire adjusted energy levels. The dipole coupling term contains two types of terms: number conserving terms (like $a^\dagger b$) and number non-number conserving terms (like ab). We make another approximation that the qubit-resonator detuning is smaller than the bare frequencies ($\omega_q, \omega_r \gg \Delta$). Although this is not strictly necessary, we can use this approximation to drop the number non-conserving terms from the perturbation. We are then left with a coupling which only has support within subspaces of total excitation number, M . We can then apply perturbation theory with each subspace with M excitations and extract relevant effects. Note that the $M = 0$ subspace comprises just the ground state and is thus unaffected.

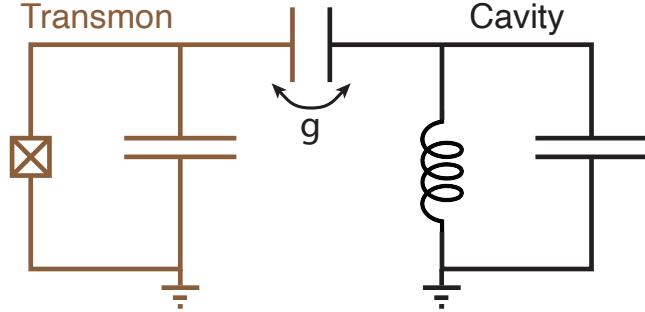


Figure 2.8: Transmon in brown coupling to a resonator in black through a capacitor with coupling strength g .

H	$\langle 00 $	$\langle 01 $	$\langle 10 $	$\langle 02 $	$\langle 11 $	$\langle 20 $
$ 00\rangle$	0			$\sqrt{2}g$		$\sqrt{2}g$
$ 01\rangle$		ω_q	g			
$ 10\rangle$		g	ω_r			
$ 02\rangle$	$\sqrt{2}g$			$2\omega_q + \alpha$	$\sqrt{2}g$	0
$ 11\rangle$				$\sqrt{2}g$	$\omega_q + \omega_r$	$\sqrt{2}g$
$ 20\rangle$	$\sqrt{2}g$			0	$\sqrt{2}g$	$2\omega_r$

Table 2.1: Resonator-transmon Hamiltonian with dipole coupling arranged into blocks of equal excitation number, truncated to the $M \leq 2$ subspaces. Greyed out coupling terms are number non-conserving. These terms are dropped since they connect energy levels which are separated on the energy scale of the bare modes.

The $M = 1$ subspace includes two bare states, $|01\rangle$ and $|10\rangle$, where the first index is the number of excitations in the resonator and the second is the number of excitations in the qubit. The approximate dressed states and frequencies are

$$|01\rangle' \approx |01\rangle + \frac{g}{\Delta} |10\rangle, \quad \omega'_{01} \approx \omega_q + \frac{g^2}{\Delta}, \quad (2.42a)$$

$$|10\rangle' \approx |10\rangle - \frac{g}{\Delta} |01\rangle, \quad \omega'_{10} \approx \omega_r - \frac{g^2}{\Delta}. \quad (2.42b)$$

In the $M = 2$ subspace, there are three states: $|02\rangle$, $|11\rangle$, $|20\rangle$. For simplicity, we just

write down the energy perturbations, and not the state perturbations:

$$\omega'_{02} \approx 2\omega_q + \alpha + \frac{2g^2}{\Delta + \alpha} \quad (2.43a)$$

$$\omega'_{11} \approx \omega_q + \omega_r + \frac{2g^2}{\Delta} - \frac{2g^2}{\Delta + \alpha} = \omega_q + \omega_r + \frac{2g^2\alpha}{\Delta(\Delta + \alpha)} \quad (2.43b)$$

$$\omega'_{20} \approx 2\omega_r - \frac{2g^2}{\Delta} \quad (2.43c)$$

Using these dressed energies, we can rewrite the dispersive Hamiltonian as:⁵

$$H_d = \omega'_{10}a^\dagger a + \omega'_{01}b^\dagger b + \frac{\alpha'}{2}(b^\dagger)^2b^2 - 2\chi a^\dagger ab^\dagger b \quad (2.44)$$

If we let $\sigma_z = \Pi(1 - 2b^\dagger b)\Pi$, where Π projects onto the subspace of zero or one transmon excitation, we find the Hamiltonian of a resonator dispersively coupled to a transmon as

$$\Pi H_d \Pi = (\omega'_{10} - \chi)a^\dagger a - \frac{\omega'_{01}}{2}\sigma_z + \chi a^\dagger a\sigma_z. \quad (2.45)$$

We can use the above energies to write down an expression for χ :

$$\chi = -\frac{\omega'_{11} - \omega'_{01} - \omega'_{10}}{2} = -\frac{g^2\alpha}{\Delta(\Delta + \alpha)}. \quad (2.46)$$

We can see that, under the dispersive approximation the gXX like term turns into a χZZ term.⁶ One can imagine performing a similar transformation on a true qubit-resonator system, a resonator-resonator system, and a true qubit-qubit system. A true qubit has infinite anharmonicity, so a true qubit-resonator system has a dispersive coupling constant of g^2/Δ . A resonator has no anharmonicity, so, as one would expect, there is no dispersive shift in a resonator-resonator system. If we consider a true qubit-qubit system, the levels $|02\rangle$ and $|20\rangle$ do not exist, leaving $|11\rangle$ as the only state in the $M = 2$ subspace. Again, in this case, there is no dispersive shift.⁷

2.4 Black box quantization

When designing quantum circuits, one is interested in efficiently simulating the dispersive couplings between different transmon and resonator modes. One method is to simulate the

⁵Here we use the original (undressed) lowering/raising operators for simplicity in presentation. Effects from using the dressed lowering/raising operators are suppressed by $\mathcal{O}(g/\Delta)$.

⁶Here we refer to X loosely as a field operator such as $\sigma_- + \sigma_+$ or $a + a^\dagger$ and Z as a number-like operator such as σ_z or $a^\dagger a$.

⁷Even if one considers the full Hamiltonian of two true qubits coupled with an $g\sigma_x\sigma_x$ interaction, including the non-number conserving term, there is no dispersive shift.

direct coupling g and use Eq. 2.46 to calculate χ . An alternative way to simulate the effects of dispersive coupling between modes is to derive the couplings and energy levels through black box quantization (BBQ)[34, 32]. One can go further than finding just the couplings between a single transmon and a single resonator; one can find multiple couplings in a multi-mode system. The goal of this method is to simulate all the relevant modes of the system in a classical eigenmode finite-element electromagnetic simulation using lumped-element inductors in place of the non-linear Josephson junctions. The fraction of the mode energy stored in the junction can then be used in conjunction with the simulated frequencies and junction inductances to extract relevant parameters about the quantum system. In particular, the coefficients of ZZ -like ($a_i^\dagger a_j^\dagger a_i a_j$ -like) interactions can be found: anharmonicities of qubits and self-Kerr of resonators when $i = j$; cross-Kerr between resonators, dispersive shift between a resonator and a qubit, and static ZZ couplings between two qubits when $i \neq j$. In principle, higher order interactions may be found as well. In BBQ, resonators and qubits modes are set on equal footing, with the only difference being the fraction of the mode energy stored in the junctions. Indeed, Josephson parametric amplifiers (reviewed in [35]), which have a much weaker anharmonicity than transmons, may also be analyzed using this method.

Calculating participation ratios

A paper by Minev et al. [32] presents the full theory behind setting up the black box equations, with most details in the paper's appendix, but here we will present a brief summary. Consider a system with modes indexed by i and junctions indexed by α . The phase across each junction α is a dipole-like element,

$$\phi^{(\alpha)} = \sum_i \gamma_i^{(\alpha)} (a_i + a_i^\dagger). \quad (2.47)$$

We have assumed that phases are the sum of field operators, where a_i are the ladder operators of as yet undefined modes, which we will refer to as the bare modes. We have also assumed that the $\gamma_i^{(\alpha)}$ are real. The first step in the BBQ procedure is to calculate the coefficients $\gamma_i^{(\alpha)}$. We assume that Hamiltonian can be broken up into a piece H_{EM} determined by linear electromagnetics, and a term H_{JJ} from the Josephson junctions.

$$H = H_{EM} + H_{JJ} \quad (2.48a)$$

$$H_{EM} = \sum_{ij} E_{ij}^{EM} (a_i + a_i^\dagger)(a_j + a_j^\dagger) \quad (2.48b)$$

$$H_{JJ} = \sum_{\alpha} E_J^{(\alpha)} (1 - \cos(\phi^{(\alpha)})) \quad (2.48c)$$

Here E_{ij}^{EM} are energy coupling terms which will not explicitly be used in the analysis, but are included to demonstrate the functional form of H_{EM} . If we truncate H_{JJ} to quadratic order in $\phi^{(\alpha)}$ to form $H_{JJ,L}$ and add in H_{EM} , we get a linear Hamiltonian H_L with vacuum

energy E_0 . Since H_L consists of purely linear circuit elements, the modes may be found with a classical eigenmode solver. We claim these simulated modes correspond to the quantum ladder operators a_i , with frequency f_i found in the simulation.

$$H_L = \sum_i h f_i a_i^\dagger a_i + E_0 \quad (2.49a)$$

$$= H_{EM} + H_{JJ,L} \quad (2.49b)$$

$$= H_{EM} + \frac{1}{2} \sum_{\alpha} E_J^{(\alpha)} (\phi^{(\alpha)})^2 \quad (2.49c)$$

$$\quad (2.49d)$$

We now find the energy stored in junction α when a single mode i of the system is excited, with colons indicating the normal ordering of the enclosed operator:

$$E_{1i}^{(\alpha)} = \frac{1}{2} E_J^{(\alpha)} \langle 1_i | :(\phi^{(\alpha)})^2: | 1_i \rangle \quad (2.50a)$$

$$= E_J^{(\alpha)} |\gamma_i^{(\alpha)}|^2 \quad (2.50b)$$

For consistency with [32], we define the participation ratio $p_i^{(\alpha)}$ as ratio of $E_{1i}^{(\alpha)}$ to the total inductive energy of the mode. Note that the total inductive energy of the mode is half of the total energy of the mode. We thus have:

$$p_i^{(\alpha)} = 2 \frac{E_{1i}^{(\alpha)}}{h f_i} = 2 \frac{E_J^{(\alpha)} |\gamma_i^{(\alpha)}|^2}{h f_i} \quad (2.51)$$

$$|\gamma_i^{(\alpha)}|^2 = \frac{h f_i}{2 E_J^{(\alpha)}} p_i^{(\alpha)} \quad (2.52)$$

We can take the $\gamma_i^{(\alpha)}$ values to be real and can calculate their absolute magnitudes, which leaves only a sign degree of freedom undefined in Eq. 2.52. This sign represents the flow of current through the junction and is extractable from simulations. In the simplest case, the sign comes into play when two modes have support on two junctions with one mode having aligned current flow and the other having anti-aligned current flow.

BBQ parameter extraction

Once the participation ratios have been calculated, the full non-linear Hamiltonian H_{BBQ} can be written down by replacing the linearized junction in H_L with the full junction Hamiltonian, H_{JJ} :

$$H_{BBQ} = H_L + H_{JJ} - H_{JJ,L} \quad (2.53a)$$

$$= \sum_i h f_i a_i^\dagger a_i + \sum_{\alpha} E_J^{(\alpha)} \left(1 - \cos(\phi^{(\alpha)}) - \frac{1}{2} (\phi^{(\alpha)})^2 \right) \quad (2.53b)$$

There are multiple ways of approximating the eigenmodes and eigenenergies to these equations. Here we will cover three methods: quartic approximation; numerical diagonalization; and symplectic approximation. The first two are presented in [32] while the third represents a middle ground.

Quartic approximation

The first way to approximate the eigenenergies of H_{BBQ} is to expand the cosine term of the Hamiltonian out to quartic order and directly match terms to the dispersive couplings as they would appear in a generalized expansion. Instead of expanding the cosine and then normal ordering, we will use an identity to normal order the cosine and then expand to quartic order. Since normal ordering results in higher order terms partially reducing to lower order terms as one commutes a_i and a_i^\dagger , normal ordering before expansion gives a slightly different result than the converse (which is derived in [32]). The identity[36] we will use is

$$e^{i\lambda(a+a^\dagger)} = e^{-\lambda^2/2} :e^{i\lambda(a+a^\dagger)}:, \quad (2.54)$$

$$\cos(\lambda(a + a^\dagger)) = e^{-\lambda^2/2} :\cos(\lambda(a + a^\dagger)):. \quad (2.55)$$

Looking at a single cosine term in H_{BBQ} , we expand while only keeping terms that conserve excitation number within every mode. This approximation assumes the system is in the dispersive regime where the mode frequencies are well separated as compared to their couplings. This prevents an excitation in one mode from hopping to another mode as well as preventing particle gain or loss. Expanding out to fourth order, we get

$$\cos(\phi^{(\alpha)}) \approx \aleph^{(\alpha)} \left(1 - \sum_i |\gamma_i^{(\alpha)}|^2 a_i^\dagger a_i + \frac{1}{4} \sum_{i \geq j} (4 - 3\delta_{ij}) |\gamma_i^{(\alpha)}|^2 |\gamma_j^{(\alpha)}|^2 a_i^\dagger a_j^\dagger a_i a_j \right), \quad (2.56)$$

where the normalization constant $\aleph^{(\alpha)}$ is defined as

$$\aleph^{(\alpha)} = \exp \left[-\frac{1}{2} \sum_i |\gamma_i^{(\alpha)}|^2 \right]. \quad (2.57)$$

We also expand $H_{JJ,L}$ out in terms of $\gamma_i^{(\alpha)}$, dropping the constants and non-particle conserving terms:

$$H_{JJ,L} \approx \sum_{\alpha,i} E_J^{(\alpha)} |\gamma_i^{(\alpha)}|^2 a_i^\dagger a_i \quad (2.58)$$

We now use this correction to look for quadratic terms $H_{JJ} - H_{JJ,L}$ not present in the linear Hamiltonian H_L . We first look for Lamb shifts using the quadratic terms,

$$H_{Lamb} = (H_{JJ} - H_{JJ,L})_{quad} = \sum_{\alpha,i} E_J^{(\alpha)} (\aleph^{(\alpha)} - 1) |\gamma_i^{(\alpha)}|^2 a_i^\dagger a_i. \quad (2.59)$$

In the transmon and resonator regime, $E_J \gg hf_i$, typically making $|\gamma_i^{(\alpha)}|^2 \ll 1$, such that $\aleph^{(\alpha)}$ is close to one. Under this assumption we get

$$H_{Lamb} \approx \frac{1}{2} \sum_{\alpha,ij} E_J^{(\alpha)} |\gamma_i^{(\alpha)}|^2 |\gamma_j^{(\alpha)}|^2 a_i^\dagger a_i. \quad (2.60)$$

We do the same approximations for the quartic terms in Eq. 2.56, which we label as ZZ terms. Again, we also present the equation with and without the assumption that $\aleph^{(\alpha)} \approx 1$:

$$H_{ZZ} = (H_{JJ} - H_{JJ,L})_{quart} = -\frac{1}{4} \sum_{\alpha, i \geq j} (4 - 3\delta_{ij}) E_J^{(\alpha)} \aleph^{(\alpha)} |\gamma_i^{(\alpha)}|^2 |\gamma_j^{(\alpha)}|^2 a_i^\dagger a_j^\dagger a_i a_j \quad (2.61a)$$

$$\approx -\frac{1}{4} \sum_{\alpha, i \geq j} (4 - 3\delta_{ij}) E_J^{(\alpha)} |\gamma_i^{(\alpha)}|^2 |\gamma_j^{(\alpha)}|^2 a_i^\dagger a_j^\dagger a_i a_j \quad (2.61b)$$

We then compare these terms to the the functional form of the Lamb shifts Δ_i , the anharmonicities α_i (distinct from junction α), and the dispersive shifts χ_{ij} , where

$$H_{Lamb} + H_{ZZ} = \sum_i \hbar \Delta_i a_i^\dagger a_i + \sum_i \frac{\hbar \alpha_i}{2} (a_i^\dagger)^2 a_i^2 + \sum_{i>j} 2\hbar \chi_{ij} a_i^\dagger a_i a_j^\dagger a_j \quad (2.62)$$

and then match terms in the $\aleph^{(\alpha)} \approx 1$ limit:

$$\hbar \Delta_i \approx \frac{1}{2} \sum_{\alpha,j} E_J^{(\alpha)} |\gamma_i^{(\alpha)}|^2 |\gamma_j^{(\alpha)}|^2 \quad (2.63a)$$

$$\hbar \alpha_i \approx -\frac{1}{2} \sum_{\alpha} E_J^{(\alpha)} |\gamma_i^{(\alpha)}|^4 \quad (2.63b)$$

$$\hbar \chi_{ij} \approx -\frac{1}{2} \sum_{\alpha} E_J^{(\alpha)} |\gamma_i^{(\alpha)}|^2 |\gamma_j^{(\alpha)}|^2 \quad (2.63c)$$

Careful! Note that the χ_{ij} derived here have a different sign convention than χ in Eq. 2.46 due to the difference of sign convention between Eq. 2.45 and Eq. 2.62.

Numerical diagonalization

Although the quartic approximation gives elegant formulas, in practice these tend to not give accurate results for the coupling. One alternative is to extract the energy levels of H_{BBQ} numerically and use these energies to calculate couplings, as implemented in the python package pyEPR[37]. For the simplest approach to diagonalization, one truncates the Hilbert space of each mode a_i to a set number of levels and expands the junctions' cosine terms to set order in $\phi^{(\alpha)}$. By diagonalizing the matrix, the lowest energy state is the calculated vacuum. The single excitation subspace will consist of renormalized modes which have large overlap

with the bare simulated modes. Similarly, doubly excited states may be identified by their large overlap with the doubly excited states of the linear Hamiltonian. Once the states are matched, the dispersive couplings may be calculated by computing differences in energies. It is worth noting that the size of the Hilbert space being diagonalized is exponential in the number of modes. When simulating couplings for the continuous error correction project, we used twelve levels per mode with a cosine truncation at order 8. Simulating five modes would give a Hilbert space of dimension $12^5 \approx 250,000$, which is prohibitively large to reasonably simulate. Instead, one can run the diagonalization with a reduced set of modes that have the highest impact on the desired coupling parameter. For example, using three modes with twelve levels each reduces the Hilbert space to a more reasonable 1,728 dimensions.

Symplectic transformation

While exploring ways to reduce the computational demands of numerical computation, we found a method to increase the accuracy of a quartic-like approximation to H_{BBQ} . In this approach, we map the bare ladder operators $\{a_i, a_i^\dagger\}$ into new ladder operators $\{b_i, b_i^\dagger\}$ using a symplectic transformation to preserve the proper commutation relations. We will discuss how to find this transformation later, but it is chosen such that the only quadratic terms of H_{BBQ} in the new basis are of the form $b_i^\dagger b_i$. The cosine terms can then be expanded and coupling terms can be identified as done in the quartic approximation above. To start, we map $\{a_i, a_i^\dagger\}$ to $\{b_i, b_i^\dagger\}$ using the symplectic transform represented as matrix T [38, 39]. Restricting ourselves to real matrices, we can split T into submatrices A and B :

$$\begin{bmatrix} \vec{b} \\ \vec{b}^\dagger \end{bmatrix} = T \begin{bmatrix} \vec{a} \\ \vec{a}^\dagger \end{bmatrix} = \begin{bmatrix} A & B \\ B & A \end{bmatrix} \begin{bmatrix} \vec{a} \\ \vec{a}^\dagger \end{bmatrix} \quad (2.64)$$

The matrix T obeys the relation

$$T^{-1} = \Sigma T^\top \Sigma \quad (2.65)$$

where Σ is defined block-wise as

$$\Sigma = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & -\mathbf{I} \end{bmatrix}. \quad (2.66)$$

We then find the transformed versions of the field coefficients $\gamma_i^{(\alpha)}$ in the new frame and label these new field coefficients $\xi_i^{(\alpha)}$. If we consider the ladder operators to transform covariantly, then $\gamma_i^{(\alpha)}$ transform contravariantly to preserve $\phi^{(\alpha)} = [\vec{\gamma}^{(\alpha)} \vec{\gamma}^{(\alpha)}] \cdot [\vec{a} \vec{a}^\dagger] = [\vec{\xi}^{(\alpha)} \vec{\xi}^{(\alpha)}] \cdot [\vec{b} \vec{b}^\dagger]$.

$$\begin{bmatrix} \vec{\xi}^{(\alpha)} \\ \vec{\xi}^{(\alpha)} \end{bmatrix} = \Sigma T \Sigma \begin{bmatrix} \vec{\gamma}^{(\alpha)} \\ \vec{\gamma}^{(\alpha)} \end{bmatrix} = \begin{bmatrix} (A - B) \vec{\gamma}^{(\alpha)} \\ (A - B) \vec{\gamma}^{(\alpha)} \end{bmatrix} \quad (2.67)$$

$$\vec{\xi}^{(\alpha)} = (A - B) \vec{\gamma}^{(\alpha)} \quad (2.68)$$

With these new field coefficients we want to find the quadratic terms of $H_{JJ} - H_{JJ,L}$. Putting the cosine term in its normal order and defining the the normalization coefficients in the new frame as $\mathfrak{N}'^{(\alpha)} = \exp \left[-\frac{1}{2} \sum_i |\xi_i^{(\alpha)}|^2 \right]$, we express this part of the Hamiltonian as

$$(H_{JJ} - H_{JJ,L})_{quad} = \frac{1}{2} \sum_{\alpha} E_J^{(\alpha)} (\mathfrak{N}'^{(\alpha)} - 1) \sum_{ij} \xi_i^{(\alpha)} \xi_j^{(\alpha)} (b_i b_j + b_i b_j^\dagger + b_j b_i^\dagger + b_i^\dagger b_j^\dagger) \quad (2.69a)$$

$$= \frac{1}{2} \sum_{\alpha} E_J^{(\alpha)} (\mathfrak{N}'^{(\alpha)} - 1) \begin{bmatrix} \vec{b} & \vec{b}^\dagger \end{bmatrix} \begin{bmatrix} \vec{\xi}^{(\alpha)} \otimes \vec{\xi}^{(\alpha)} & \vec{\xi}^{(\alpha)} \otimes \vec{\xi}^{(\alpha)} \\ \vec{\xi}^{(\alpha)} \otimes \vec{\xi}^{(\alpha)} & \vec{\xi}^{(\alpha)} \otimes \vec{\xi}^{(\alpha)} \end{bmatrix} \begin{bmatrix} \vec{b} \\ \vec{b}^\dagger \end{bmatrix}, \quad (2.69b)$$

up to additive constants from commutation relations. Here we have used “ \otimes ” to symbolize the outer product. We define the matrix Q to be

$$Q = \frac{1}{2} \sum_{\alpha} E_J^{(\alpha)} (\mathfrak{N}'^{(\alpha)} - 1) \vec{\xi}^{(\alpha)} \otimes \vec{\xi}^{(\alpha)}. \quad (2.70)$$

Similarly, we transform H_L , which only includes uncoupled harmonic oscillator terms in the original basis, to the new basis. Here, we let ω be a diagonal matrix whose entries correspond to the frequencies of the bare modes.

$$H_L = \frac{\hbar}{2} \begin{bmatrix} \vec{a} & \vec{a}^\dagger \end{bmatrix} \begin{bmatrix} \mathbf{0} & \omega \\ \omega & \mathbf{0} \end{bmatrix} \begin{bmatrix} \vec{a} \\ \vec{a}^\dagger \end{bmatrix} \quad (2.71a)$$

$$= \frac{\hbar}{2} \begin{bmatrix} \vec{b} & \vec{b}^\dagger \end{bmatrix} \Sigma T \Sigma \begin{bmatrix} \mathbf{0} & \omega \\ \omega & \mathbf{0} \end{bmatrix} \Sigma T^\top \Sigma \begin{bmatrix} \vec{b} \\ \vec{b}^\dagger \end{bmatrix} \quad (2.71b)$$

$$= \frac{\hbar}{2} \begin{bmatrix} \vec{b} & \vec{b}^\dagger \end{bmatrix} \begin{bmatrix} -(A\omega B^\top + B\omega A^\top) & A\omega A^\top + B\omega B^\top \\ A\omega A^\top + B\omega B^\top & -(A\omega B^\top + B\omega A^\top) \end{bmatrix} \begin{bmatrix} \vec{b} \\ \vec{b}^\dagger \end{bmatrix} \quad (2.71c)$$

We define the matrices S and R to be

$$S = A\omega A^\top + B\omega B^\top, \quad (2.72a)$$

$$R = A\omega B^\top + B\omega A^\top, \quad (2.72b)$$

leading to the fully transformed Hamiltonian under the renormalized modes.

$$(H_{BBQ})_{quad} = \frac{\hbar}{2} \begin{bmatrix} \vec{b} & \vec{b}^\dagger \end{bmatrix} \begin{bmatrix} Q - R & Q + S \\ Q + S & Q - R \end{bmatrix} \begin{bmatrix} \vec{b} \\ \vec{b}^\dagger \end{bmatrix} \quad (2.73)$$

If the correct transformation matrix T is chosen, then the quadratic terms of H_{BBQ} will only be those of decoupled harmonic oscillators. In other words, $Q - R$ will be all zero and $Q + S$ will be diagonal and will consist of the Lamb shifted eigenfrequencies. In order to find the correct transformation T , one can numerically optimize the n^2 coefficients for generators

of the symplectic transform with n modes to reach these conditions. Expanding the cosine term out in this new basis, the coupling coefficients can be matched as has been done for the quartic case. Alternatively, one can numerically diagonalize the two-excitation subspace in this new basis using the Hamiltonian terms of the form $b_i^\dagger b_j^\dagger b_k b_l$ as well as the quadratic terms. In the b basis, the two-excitation states do not mix with the single-excitation states to first order.

When choosing a BBQ extraction method, there is some trade off between speed of the BBQ computation and parameter accuracy. If one is including a full finite element solver in an optimization loop, then numerical diagonalization (incorporating only critical modes and junctions) is likely the best option. If one is using many modes or exploring parameter space of frequencies and participation ratios, then using a symplectic method is a good way to get preliminary results. Fig. 2.9 demonstrates the results of BBQ for a single junction resonator-qubit system as a function of the junction inductance for a constant bare simulated frequency. All methods perform similarly in the highly linear case of a low Josephson inductance, but the quartic approximation quickly becomes inaccurate for realistic transmon regimes. Note that this comparison is for a particular coupling and frequency regime. It does not represent a general comparison of the methods in more exotic system regimes such as the straddling regime where the resonator frequency is between ω_{01} and ω_{12} of the qubit.

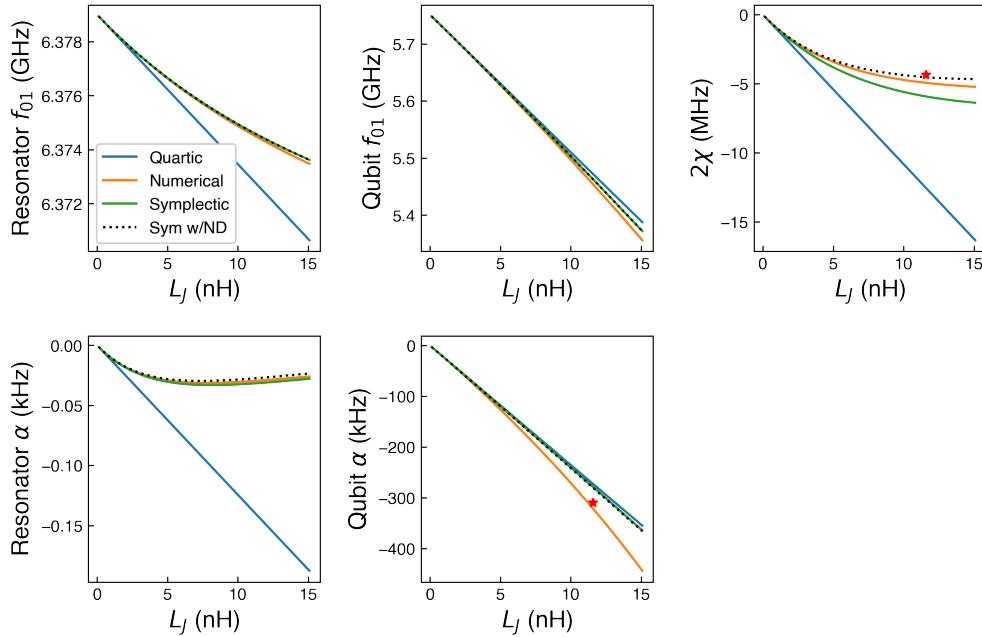


Figure 2.9: Parameters extracted for a coupled resonator-transmon system with a single junction as a function of Josephson inductance (L_J). The “Symplectic” method takes the quartic approximation in the transformed frame and the “Sym w/ND” method diagonalizes the two excitation subspace using the quartic (and quadratic) terms in the transformed frame. The bare frequencies and participation ratios for the resonator and qubit respectively are (6.379 GHz, .02) and (5.751 GHz, .965). These values correspond to a subsystem of a fabricated parity chip with a tunable qubit and a resonator of frequency 6.314 GHz. After tuning the qubit so that the system match the numerical-BBQ expected resonator-qubit detuning, we measure qubit anharmonicity α and dispersive shift χ . This single data point is plotted as a red star. A more comprehensive comparison of theory to experiment over a wide range of participation ratios may be found in [32].

Chapter 3

Quantum Measurement

Quantum mechanics is a theory of information. The dynamics of a quantum systems are often broken up into the two strictly distinct processes of unitary evolution and projective measurement. Whereas unitary evolution dictates the flow of information within a quantum system, projective measurement describes the irreversible transfer of information to an observer. Although these two processes are theoretically sufficient to determine the evolution of any quantum system, measurements in experimental setups are better described as a combination of unitary evolution and projective measurements. This description encompasses a larger class of generalized measurements representing the action of measurement on a system over a finite time. Additionally, we can use these generalized measurements to describe continuous measurements, wherein the observer gets information about a system at a finite rate. Such continuous measurements describe the internal dynamics of measurement, as finite interaction strengths limit how fast information can be transferred.

3.1 Discrete maps

We will begin our measurement discussion by describing measurements as discrete actions on a quantum system. Here we will discuss measurement on the level of mapping density matrices to density matrices, without discussing the internal, infinitesimal dynamics of how the measurements are made.

Projective measurement

When making a measurement on a quantum system, two things happen: the state of the system is disturbed and information about how the system was disturbed is picked up by the observer. We can define a projective measurement M through a set of complete orthonormal

quantum projectors and classical meter labels indexed by λ :

$$M = \left\{ (\hat{\Pi}_\lambda, \lambda) \right\} \quad (3.1a)$$

$$\sum_\lambda \hat{\Pi}_\lambda = I \quad (3.1b)$$

$$\hat{\Pi}_\lambda \hat{\Pi}_{\lambda'} = \delta_{\lambda, \lambda'} \hat{\Pi}_\lambda \quad (3.1c)$$

When the measurement is applied to a system represented by the density matrix ρ , the meter reports the outcome λ with probability

$$p_\lambda = \text{Tr}[\rho \hat{\Pi}_\lambda]. \quad (3.2)$$

and, conditional on this outcome the state evolves as

$$\rho \mapsto \frac{\hat{\Pi}_\lambda \rho \hat{\Pi}_\lambda}{\text{Tr}[\rho \hat{\Pi}_\lambda]}. \quad (3.3)$$

The evolution of the density matrix is non-linear due to the denominator being dependent on ρ , which is a general feature of measurements so long as one is observing the classical meter. If the observer performs a measurement, but does not look at the meter, their knowledge of the state may be captured in a density matrix which is the weighted sum of density matrices over all possible outcomes:

$$\rho \mapsto \sum_\lambda \frac{\hat{\Pi}_\lambda \rho \hat{\Pi}_\lambda}{p_\lambda} p_\lambda = \sum_\lambda \hat{\Pi}_\lambda \rho \hat{\Pi}_\lambda. \quad (3.4)$$

In this unconditional case, the density matrix evolution stays linear. In the section on continuous measurement, we will similarly see that measuring without looking at the outcomes results in linear evolution, but measuring and looking at the outcomes results in a non-linear evolution.

Generalized measurement: POVMs

One can perform a larger class of discrete measurements on a system than just projective measurements. These measurements are known as positive operator-valued measures (POVMs). To model POVMs, we include a quantum meter which is able to interact with the system we wish to perform a measurement on. For this section, we assume that the observer is perfectly efficient at extracting information from the meter, such that any information leaving the system is captured by the observer. The meter and system start out in a factorized state. For simplicity, we will additionally consider each of them to be in a pure state (with the meter in $|m\rangle$ and the system in $|\psi_i\rangle$):

$$|\Psi_i\rangle = |m\rangle \otimes |\psi_i\rangle. \quad (3.5)$$

We then apply a joint unitary on the two systems, resulting in the (possibly) entangled state

$$|\Psi_f\rangle = U |\Psi_i\rangle. \quad (3.6)$$

Finally, the meter is projectively measured with outcome λ . To prevent possible future backflow of information from the meter to the system, we want this projection to completely disentangle the meter from the system for all possible initial states of the system. We ensure this condition by making the corresponding projector $\hat{\Pi}_\lambda$ project onto one state: $\hat{\Pi}_\lambda = |\lambda\rangle\langle\lambda|$. The unnormalized final state of the unitary and projection is

$$|\psi_f^{(\lambda)}\rangle = (\langle\lambda| \otimes I) U |\Psi_i\rangle. \quad (3.7)$$

Indeed, we can capture the system process conditional on outcome λ in the operator

$$\Omega_\lambda = (\langle\lambda| \otimes I) U (|m\rangle \otimes I). \quad (3.8)$$

If we now generalize the system's initial state $|\psi_i\rangle$ to the density matrix ρ , the probability of outcome λ is

$$p_\lambda = \text{Tr}[\Omega_\lambda^\dagger \Omega_\lambda \rho], \quad (3.9)$$

resulting in a nonlinear state evolution when normalized. We present this action in both the density matrix and state vector formats:

$$\rho \mapsto \frac{\Omega_\lambda \rho \Omega_\lambda^\dagger}{\text{Tr}[\Omega_\lambda^\dagger \Omega_\lambda \rho]} \quad (3.10a)$$

$$|\psi_i\rangle \mapsto \frac{\Omega_\lambda |\psi_i\rangle}{\sqrt{\langle\psi_i|\Omega_\lambda^\dagger \Omega_\lambda |\psi_i\rangle}} \quad (3.10b)$$

These equations are similar in form to the direct projective measurements in Eq. 3.3, but the $\{\Omega_\lambda\}$ operators are not necessarily projectors and instead satisfy the weaker condition

$$\sum_\lambda \Omega_\lambda^\dagger \Omega_\lambda = I. \quad (3.11)$$

As in the case of projective measurements, if observer does not look at the results of the meter, they again recover a linear map for the density matrix:

$$\rho \mapsto \sum_\lambda \frac{\Omega_\lambda \rho \Omega_\lambda^\dagger}{p_\lambda} p_\lambda = \sum_\lambda \Omega_\lambda \rho \Omega_\lambda^\dagger. \quad (3.12)$$

This unconditional map is known as a Kraus map, with the set of operators $\{\Omega_\lambda\}$ called the Kraus operators. This formulation indeed captures any possible unconditional dynamical map of a quantum system (including unitaries) [40], under the condition that the system is factorizable from any other Hilbert space at the start of the map.

3.2 Continuous measurement

Moving away from discrete measurements, we wish to model the infinitesimal processes of continuous measurement. In order to perform a continuous measurement, an observer must have a constant stream of information coming in from the quantum system under observation. In a laboratory setting, this stream is often encoded in an outgoing electromagnetic field, so we will consider a system interacting with a series of harmonic oscillators flying past the system as shown in Fig. 3.1 [[41, 42]]. In a microwave experiment, these waves travel through a transmission line. We assume there is a system operator S which allows the system to interact with each of these probing oscillators in turn for a short period of time ΔT . Although we leave S generic, it is often the lowering operator of a cavity or a qubit. Such a process models reflection off of circuit elements such the cavities as shown in Fig 5.3. We label each flying oscillator τ with an annihilation operator B_τ , where $\tau \in \{0, \Delta T, 2\Delta T, \dots\}$. During the period from τ to $\tau + \Delta T$, the system interacts with cavity τ under a Hamiltonian

$$H_\tau = ig(SB_\tau^\dagger - S^\dagger B_\tau) \quad (3.13)$$

where we have taken g to be real. Such a Hamiltonian can arise using the most standard way of coupling cavities or a transmons together through dipole moments. Indeed, when S is the lowering operator of a qubit, this Hamiltonian is the rotating Jaynes-Cummings for a qubit and oscillator on resonance.

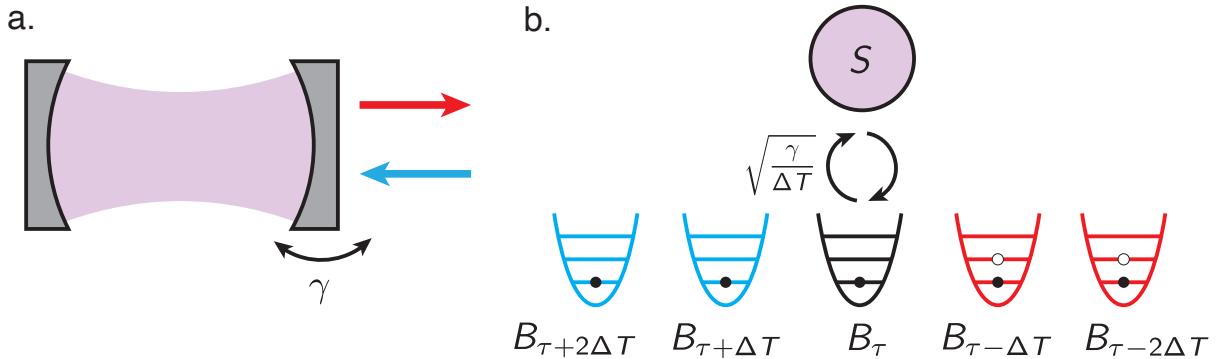


Figure 3.1: Interaction of a quantum system and a transmission line. (a) An example of an interacting system: a cavity couples to a transmission line with decay constant γ . (b) Coupling between the system and the waveguide can be modeled as a series of harmonic oscillators which serially interact with the cavity each for time ΔT . The outgoing modes can then be measured by a detector.

Before finding the dynamics this Hamiltonian generates, we take a closer look at g , and in particular how g scales as we shrink the time interval ΔT . Here, we give a plausibility

argument rather than a full proof. (See [43] for a more full derivation.) Consider that B exists as a mode traveling down in a transmission line with temporal width ΔT with a corresponding spacial width $\Delta x \propto \Delta T$. The electric field operator on the transmission line scales like

$$E \propto \frac{1}{\sqrt{\Delta x}} (B + B^\dagger) \quad (3.14)$$

The g in the system-probe coupling, which has the form $\mu \cdot E$, then scales like $1/\sqrt{\Delta x}$, resulting in a scaling of $1/\sqrt{\Delta T}$. We can then write g by including a new frequency γ which is independent of ΔT :

$$g = \sqrt{\frac{\gamma}{\Delta T}} \quad (3.15)$$

We now use the original Hamiltonian in Eq. 3.13 to generate a unitary which we truncate at order ΔT .

$$U_\tau = e^{-iH_r\Delta T} \quad (3.16a)$$

$$\approx 1 + \sqrt{\gamma\Delta T}(SB_\tau^\dagger - S^\dagger B_\tau) - \frac{\gamma}{2}(SB_\tau^\dagger - S^\dagger B_\tau)^2\Delta T \quad (3.16b)$$

Note that to get to order ΔT , we had to expand out the exponential to second order. This second order expansion is not only important because of the physical considerations of g 's scaling, but actually necessary from an informational perspective. If the scaling of g in ΔT were different, the rate of information coming out of the system would either diverge or go to zero in the limit that $\Delta T \rightarrow 0$.

To see the measurement's effect on the system, we need to first decide both how to initialize the probe and how we want to measure the probe. For simplicity will assume that each of the probe oscillators is initialized in the ground state. We will then consider two cases of measurement bases corresponding to common lab bases: the Fock basis, and the quadrature basis.

Photon counting

Our first choice of bases is the Fock basis, which can be implemented using a photon counter. Starting with the unitary in Eq. 3.16b, we perform the partial trace over the probe with initial probe state $|0\rangle$ and final probe state $|i\rangle$, $i \in \{0, 1, 2\}$ to recover the Kraus operators Ω_i on the system associated with each outcome. Tracing over the probe with any other final state will result in zero, and thus does not result in a Kraus operator. Indeed, we find below that in the $\Delta T \rightarrow 0$ limit, the chance of measuring $|2\rangle$ in a finite time also goes to zero, so photon counting may only result in zero or one photons.

$$\Omega_0 = \langle 0|U|0\rangle = 1 - \frac{\gamma}{2}S^\dagger S\Delta T \quad (3.17a)$$

$$\Omega_1 = \langle 1|U|0\rangle = \sqrt{\gamma\Delta T}S \quad (3.17b)$$

$$\Omega_2 = \langle 2|U|0\rangle = -\frac{\gamma}{2}S^2\Delta T \quad (3.17c)$$

From these Kraus operators, we then get the probabilities of a given outcome, with expectation values taken with respect to the system's density matrix:

$$p_0 = \langle \Omega_0^\dagger \Omega_0 \rangle = 1 - \gamma \langle S^\dagger S \rangle \Delta T \quad (3.18a)$$

$$p_1 = \langle \Omega_1^\dagger \Omega_1 \rangle = \gamma \langle S^\dagger S \rangle \Delta T \quad (3.18b)$$

$$p_2 = \langle \Omega_2^\dagger \Omega_2 \rangle = \mathcal{O}(\Delta T^2) \quad (3.18c)$$

As promised, the integrated probability of getting two photons goes to zero as ΔT goes to zero while the integrated probability of getting a single photon in a finite time is between zero and one. Using Eq. 3.12, we derive the unconditional evolution of the system density matrix ρ ,

$$\Delta \rho = \sum_{\lambda} \Omega_{\lambda} \rho \Omega_{\lambda}^\dagger - \rho \quad (3.19a)$$

$$= \left(-\frac{1}{2} \{ S^\dagger S, \rho \} + S \rho S^\dagger \right) \gamma \Delta T \quad (3.19b)$$

$$d\rho = \gamma \left(S \rho S^\dagger - \frac{1}{2} \{ S^\dagger S, \rho \} \right) dt \quad (3.19c)$$

$$= \gamma \mathcal{D}[S] \rho dt, \quad (3.19d)$$

where in the final two lines we take the infinitesimal limit $\Delta T \rightarrow dt$ express the decoherence in Lindbladian form. Here \mathcal{D} is a linear map of ρ with parameter S defined as

$$\mathcal{D}[S]\rho \equiv S \rho S^\dagger - \frac{1}{2} \{ S^\dagger S, \rho \}, \quad (3.20)$$

which is a common form of a linear map generating irreversible decoherence in a system [44].

To derive the equations of motion for the system conditional on the measurement, we first note that in the limit of short time, the probability of getting a single count in a single time bin tends to zero. Indeed, if a count comes out of the system, it will happen at a finite time. Up to that point, we can consider the system evolving under the unnormalized equation

$$d\rho = \Omega_0 \rho \Omega_0^\dagger - \rho = -\frac{\gamma}{2} [S^\dagger S, \rho] dt \quad (3.21a)$$

$$\dot{\rho} = -\frac{\gamma}{2} [S^\dagger S, \rho]. \quad (3.21b)$$

Although one could normalize the density matrix at every time step to get a normalized state, it is helpful both computationally and theoretically to keep the state unnormalized. By doing so, the trace of the density matrix dictates the probability that the system has so far not undergone any jumps. To determine whether a jump occurs, the experimentalist looks at their photo detector to see if a click has happened at a particular time. In contrast, the theorist simulating such a system chooses a random number p uniformly from 0 to 1. If the trace of the evolved density matrix ever falls below p , a jump event occurs. Since,

under perfect photodetection, the density matrix remains pure, it is often preferable to evolve an unnormalized wave function (which we denote with an overhead tilde) according to the equation

$$\Delta |\tilde{\psi}\rangle = \Omega_0 |\tilde{\psi}\rangle - |\tilde{\psi}\rangle = -\frac{\gamma}{2} S^\dagger S |\tilde{\psi}\rangle \Delta T, \quad (3.22a)$$

$$|\dot{\tilde{\psi}}\rangle = -\frac{\gamma}{2} S^\dagger S |\tilde{\psi}\rangle. \quad (3.22b)$$

This form allows one to keep track of fewer numbers in a computer's memory, though at the expense of having to run a simulation multiple times over different trajectories, each of which constitutes an experimentalist's set of detection times or theorist's probabilities p_i depending on the number of decay events one wishes to consider. Once a decay event does occur, and the photon counter clicks, the system undergoes a discrete jump:

$$\rho \mapsto \frac{\Omega_1 \rho \Omega_1^\dagger}{\text{Tr}[\Omega_1^\dagger \Omega_1 \rho]} = \frac{S \rho S^\dagger}{\text{Tr}[S^\dagger S \rho]} \quad (3.23)$$

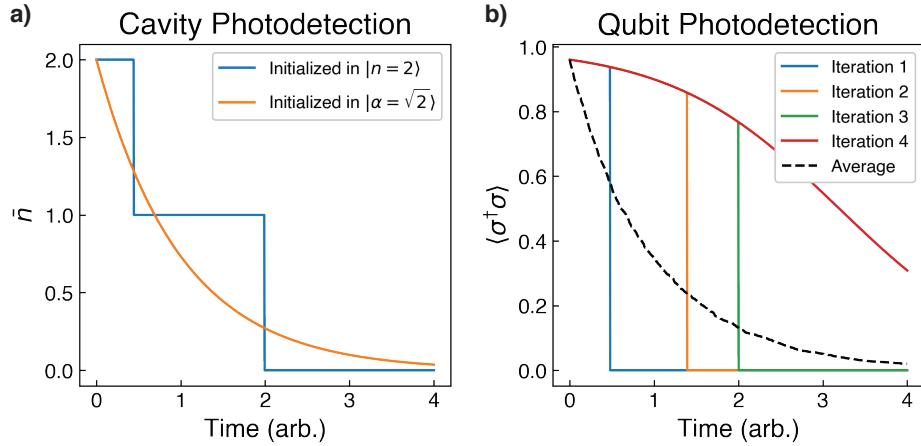


Figure 3.2: Simulating decay under photodetection measurement. (a) Single shot decay of a cavity ($S = a$) starting in a Fock state, and starting in a coherent state. (b) Decay of a qubit ($S = \sigma$) starting in a pure state with initial population $\langle \sigma^\dagger \sigma \rangle \approx .96$ under different trajectories of a photon loss. The dashed line shows the ensemble average decay.

Example trajectories of decay under photodetection are shown in Fig. 3.2 for cavities ($S = a$) and qubits ($S = \sigma$) under decay operators. Note that for a cavity's coherent state, the jump events themselves do not actually affect the state other than a global phase since coherent states are eigenstates of the annihilation operator. However, this phase could be picked up if the cavity is prepared in a cat state such as $|\alpha\rangle + |-\alpha\rangle$, in which case a decay

event flips the relative phase of the two coherent component states. As in the cavity case, a decaying qubit ($S = \sigma$) is slowed by the act of collecting no photons, and a pure state will decay on the surface of the Bloch sphere. If a photon is detected, then the qubit immediately jumps to the $|0\rangle$ state.

Homodyne measurement

In the microwave domain, photon detection remains relatively rare. A more common measurement is to readout the quadrature basis of the probe (“in-phase” I , “quadrature” Q , or some linear combination thereof). Such a measurement is known as homodyne and can be implemented by a variety of devices. In this work, we use Josephson parametric amplifiers (JPAs) [45].¹ By pumping a JPA with a tone phase-locked to the measurement tone and choosing a relative phase offset, a chosen quadrature can be amplified. For our analysis, we will measure in the I quadrature, in which we construct the operator R :

$$R = B + B^\dagger. \quad (3.24)$$

This quadrature measurement is analogous to measuring dimensionless position in a mechanical oscillator, with a normalization such that $\langle 0|R^2|0\rangle = 1$ instead of $\frac{\hbar}{2m\omega}$. Measuring in this basis, the output is a continuous variable r with associated Kraus operator

$$\Omega_r = \langle r|U|0\rangle \quad (3.25a)$$

$$= \langle r|0\rangle + \langle r|1\rangle \sqrt{\gamma\Delta T}S - \frac{\gamma}{2}\Delta T \left(\langle r|0\rangle S^\dagger S - \langle r|2\rangle \sqrt{2}S^2 \right) \quad (3.25b)$$

$$= \langle r|0\rangle \left[1 + r\sqrt{\gamma\Delta T}S - \frac{\gamma}{2}\Delta T (S^\dagger S - (r^2 - 1)S^2) \right]. \quad (3.25c)$$

Here, $\langle r|0\rangle = (2\pi)^{-1/4} \exp[-r^2/4]$ is a Gaussian representing the probe’s ground state expressed in the position basis. In the last line, we have used the Hermite polynomials to express the harmonic oscillator’s first and second excited states in terms of it’s ground state. We also drop out the $(r^2 - 1)\Delta T$ term for convenience since, as we will later learn, even in conditional evolution this term does not affect the system dynamics. Letting the probability of measuring r if the probe is in the ground state be

$$P_0(r) = |\langle r|0\rangle|^2 = \frac{1}{\sqrt{2\pi}} e^{-r^2/2}, \quad (3.26)$$

we are then left with the operator

$$\Omega_r = \sqrt{P_0(r)} \left[1 + r\sqrt{\gamma\Delta T}S - \frac{\gamma}{2}\Delta T S^\dagger S \right] \quad (3.27)$$

¹For design considerations of JPAs, see [35].

which governs the measurement dynamics of the system. We can then rederive the unconditional evolution of the system under a homodyne measurement:

$$\Delta\rho = \int_{-\infty}^{\infty} dr \Omega_r \rho \Omega_r^\dagger - \rho \quad (3.28a)$$

$$= \int_{-\infty}^{\infty} dr P_0(r) \left[r \sqrt{\gamma \Delta T} (S\rho + \rho S^\dagger) + \gamma \Delta T r^2 S \rho S^\dagger - \frac{\gamma}{2} \Delta T (S^\dagger S \rho + \rho S^\dagger S) \right] \quad (3.28b)$$

$$= \int_{-\infty}^{\infty} dr P_0(r) \left[\gamma \Delta T r^2 S \rho S^\dagger - \frac{\gamma}{2} \Delta T (S^\dagger S \rho + \rho S^\dagger S) \right] \quad (3.28c)$$

$$= \int_{-\infty}^{\infty} dr P_0(r) \left[\gamma \Delta T (r^2 - 1) S \rho S^\dagger - \frac{\gamma}{2} \Delta T (S^\dagger S \rho + \rho S^\dagger S - 2 S \rho S^\dagger) \right] \quad (3.28d)$$

$$= -\gamma \Delta T \left(\frac{1}{2} \{S^\dagger S, \rho\} - S \rho S^\dagger \right) \quad (3.28e)$$

$$d\rho = \gamma \left(S \rho S^\dagger - \frac{1}{2} \{S^\dagger S, \rho\} \right) dt. \quad (3.28f)$$

Here we have used the orthonormality of the Hermite polynomials under the Gaussian weight of $P_0(r)$ to do all of our integrations. As we would hope, the unconditional density matrix evolution matches Eq. 3.19c, the unconditional evolution when using a photodetector.

To determine the conditional evolution of the system, we determine the probability of measuring a particular outcome r as

$$p_r = \langle \Omega_r^\dagger \Omega_r \rangle_\rho = P_0(r) \left[1 + r \sqrt{\gamma \Delta T} \langle S + S^\dagger \rangle_\rho + \gamma \Delta T (r^2 - 1) \langle S^\dagger S \rangle_\rho \right] \quad (3.29a)$$

$$\approx P_0(r) \left[1 + r \sqrt{\gamma \Delta T} \langle S + S^\dagger \rangle_\rho \right]. \quad (3.29b)$$

Here, we again drop the $(r^2 - 1)\Delta T$ term even though it is of order ΔT (as we will justify in a digression on stochastic calculus). We also explicitly label the system expectation values with ρ such that $\langle A \rangle_\rho \equiv \text{Tr}[A\rho]$. We then want to determine the mean and variance of r , where we use $\langle A \rangle_r = \int A p_r dr$:

$$\langle r \rangle_r = \int_{-\infty}^{\infty} dr P_0(r) \left[r + r^2 \sqrt{\gamma \Delta T} \langle S + S^\dagger \rangle_\rho \right] = \sqrt{\gamma \Delta T} \langle S + S^\dagger \rangle_\rho \quad (3.30a)$$

$$\langle r^2 \rangle_r = \int_{-\infty}^{\infty} dr P_0(r) \left[r^2 + r^3 \sqrt{\gamma \Delta T} \langle S + S^\dagger \rangle_\rho \right] = 1 \quad (3.30b)$$

Here we start to introduce the finite time equivalent of a concept from stochastic calculus: the stochastic increment ΔW . ΔW is a random variable drawn from a Gaussian with zero mean and variance ΔT . At every time-step, r also has the statistics of a (shifted and scaled) Gaussian, so every measured r has a corresponding ΔW . Therefore, if one has a sequence $\{r_\tau\}$, one can turn it into a sequence $\{\Delta W_\tau\}$ and vice versa. Explicitly this relation is

$$r \sqrt{\Delta T} = \sqrt{\gamma} \langle S + S^\dagger \rangle_\rho \Delta T + \Delta W. \quad (3.31)$$

One can verify that this correctly gives us $\langle \Delta W \rangle_r = 0$ and $\langle \Delta W^2 \rangle_r = \Delta T$. Note that $\Delta W = \mathcal{O}(\sqrt{\Delta T})$, which is important when taking approximations to the correct order in Taylor expansions.

Under a normalized time step and assuming $\Delta W^2 = r^2 \Delta T = \Delta T$ (equivalent to dropping $(r^2 - 1)\Delta T$), the density matrix propagates as

$$\rho \mapsto \frac{\Omega_r \rho \Omega_r^\dagger}{p_r} = \frac{P_0(r) [\rho + r\sqrt{\gamma \Delta T}(S\rho + \rho S^\dagger) - \frac{\gamma}{2} \Delta T (S^\dagger S \rho + \rho S^\dagger S - 2r^2 S \rho S^\dagger)]}{P_0(r) [1 + r\sqrt{\gamma \Delta T} \langle S + S^\dagger \rangle_\rho]} \quad (3.32a)$$

$$= \frac{\rho + r\sqrt{\gamma \Delta T}(S\rho + \rho S^\dagger) + \gamma \Delta T \mathcal{D}[S]\rho}{1 + r \langle r \rangle_r} \quad (3.32b)$$

$$\approx \left(\rho + r\sqrt{\gamma \Delta T}(S\rho + \rho S^\dagger) \right) (1 - r \langle r \rangle_r + r^2 \langle r \rangle_r^2) + \gamma \Delta T \mathcal{D}[S]\rho \quad (3.32c)$$

$$\approx \left(\rho + r\sqrt{\gamma \Delta T}(S\rho + \rho S^\dagger) \right) \left(1 - \langle r \rangle_r \frac{\Delta W}{\sqrt{\Delta T}} \right) + \gamma \Delta T \mathcal{D}[S]\rho \quad (3.32d)$$

$$\approx \rho + \sqrt{\gamma} \left(S\rho + \rho S^\dagger - \langle S + S^\dagger \rangle_\rho \rho \right) \Delta W + \gamma \Delta T \mathcal{D}[S]\rho \quad (3.32e)$$

$$\Delta \rho = \sqrt{\gamma} \left(S\rho + \rho S^\dagger - \langle S + S^\dagger \rangle_\rho \rho \right) \Delta W + \gamma \mathcal{D}[S]\rho \Delta T \quad (3.32f)$$

Although some algebra steps have been skipped, we have calculated the density matrix propagation as a function of the probe's measured quadrature. Note that due to the normalization, the propagation is no longer linear in ρ . To finish the finite time-step treatment of measuring a decaying system in homodyne, we need to transform r into a classical measured voltage. As described in Eq. 3.14, the electric field and hence the voltage at the detector scales like $(B + B^\dagger)/\sqrt{\Delta T}$. Using Eq. 3.24, the quantum measured voltage will be of the form

$$\hat{V} = A \frac{B + B^\dagger}{\sqrt{\Delta T}} = A \frac{R}{\sqrt{\Delta T}} \quad (3.33)$$

resulting in a classical voltage

$$V = A \frac{r}{\sqrt{\Delta T}} \quad (3.34)$$

where A is a constant scaling factor in units of Volts/ $\sqrt{\text{s}^{-1}}$ characterizing the physical noise power in a certain bandwidth. Using Eq. 3.31, we write down the measured voltage in terms of ΔW .

$$V \Delta T = A \left(\sqrt{\gamma} \langle S + S^\dagger \rangle_\rho \Delta T + \Delta W \right). \quad (3.35)$$

This finite-time equation is often appropriate in a laboratory setting as one generally integrates voltage for some finite time ($\Delta T \approx 10 - 100$ ns for a typical superconducting microwave circuit applications), in which case V is the average voltage during this interval, and $A = \sqrt{\Delta T \langle \Delta V^2 \rangle_{\Delta W}}$, where $\langle \Delta V^2 \rangle_{\Delta W}$ is the variance of the average voltage in a

time window ΔT .² Since the value of A as measured at the detector is dependent on the amplification chain, we set $A = 1$ for the remainder of the chapter.

Although so far we have worked in the limit of finite time steps and in so doing been able to take liberties with our notation, this conditional evolution can be taken to the infinitesimal time limit. To do so we take $\Delta T \rightarrow dt$ and $\Delta W \rightarrow dW$. Here, dW represents a stochastic increment, which we will briefly discuss after presenting Eq. 3.32f and Eq. 3.35 in their infinitesimal form. We also introduce the non-linear operator $\mathcal{H}[S]$,

$$\mathcal{H}[S]\rho \equiv S\rho + \rho S^\dagger - \langle S + S^\dagger \rangle_\rho \rho. \quad (3.36)$$

Setting $A = 1$, we present the stochastic master equation for our original system in a compact form:

$$d\rho = \gamma \mathcal{D}[S]\rho dt + \sqrt{\gamma} \mathcal{H}[S]\rho dW \quad (3.37a)$$

$$Vdt = \sqrt{\gamma} \langle S + S^\dagger \rangle_\rho dt + dW. \quad (3.37b)$$

Different references on the stochastic master equation have different conventions on how to scale V , as well as whether or not to absorb $\sqrt{\gamma}$ into the operator S . We present the equations in their current form since the units are all captured by γ and voltage fluctuations are normalized without reference to the coupling strength of the system. Under the chosen normalization, if one measures the average voltage ($V_{avg} = \frac{1}{T} \int_0^T Vdt$) over time T , then the variance of V_{avg} with respect to different realizations of dW is

$$\text{Var}(V_{avg}) = \left\langle \left(\frac{1}{T} \int_0^T dW \right)^2 \right\rangle_{dW} = \frac{1}{T}. \quad (3.38)$$

We additionally reiterate that the initial state of the transmission line is the vacuum state for the previous analysis. This is a good assumption for a cavity in transmission. However, when a cavity is in a reflection geometry, the emitted signal is combined with the reflected tone used to drive the cavity state. Since this input tone is usually a coherent state, the emitted voltage is simply displaced by the reflected voltage, adding a constant offset to the dt term, but not changing the dW term.³

As with the case for photon counting, a pure state evolving under the stochastic master equation stays pure, so one can derive a similar stochastic equation for the wave function. Conceptually the derivation is the same as for the density matrix evolution in that

$$|\psi\rangle \mapsto \frac{\Omega_r |\psi\rangle}{\sqrt{\langle \psi | \Omega_r^\dagger \Omega_r | \psi \rangle}} \quad (3.39)$$

²For experimentalists determining the scaling factor A , check the bandwidth of your amplifiers when characterizing the noise. In order to measure in homodyne, one often uses a Josephson parametric amplifier, which offers limited bandwidth resulting in temporal voltage correlations. Effectively, an experimental calculation of A depends on ΔT when ΔT is less than the inverse bandwidth of the amplifier.

³For an example where the initial state of the transmission line is not a coherent state, and instead a squeezed state, see [46].

for the normalized case and

$$|\tilde{\psi}\rangle \mapsto \Omega_r |\tilde{\psi}\rangle \quad (3.40)$$

for the unnormalized case. Then one would express r in terms of ΔW to get the equation of motion in standard form which can be combined with Eq. 3.37b for the full evolution of the state and measurement outcome. Without derivation (see [47, 48] for details), this state update equation is

$$d|\tilde{\psi}\rangle = \left(-\frac{\gamma}{2} S^\dagger S dt + \sqrt{\gamma} S dW \right) |\tilde{\psi}\rangle \quad (3.41)$$

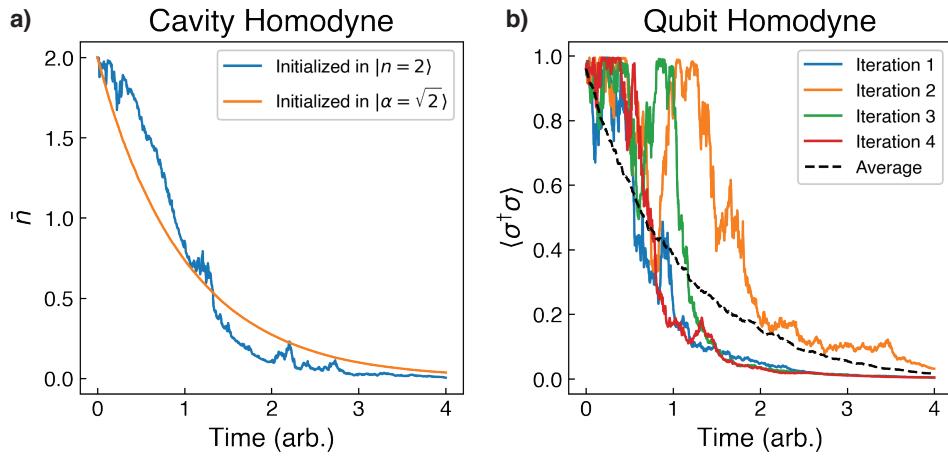


Figure 3.3: Decay under homodyne measurement. (a) Single shot decay of a cavity ($S = a$) starting in a Fock state, and starting in a coherent state. (b) Decay of a qubit ($S = \sigma$) starting in a pure state with initial population $\langle \hat{\sigma}^\dagger \hat{\sigma} \rangle \approx .96$ under different trajectories. For an individual trajectory, the state stays pure for the duration of the measurement.

Sample trajectories of decay under homodyne detection are shown in Fig. 3.2. Similarly to the photon detection case, the stochastic measurement dynamics do not affect a coherent state. One can confirm this by plugging a coherent state into Eq. 3.36 and finding $\mathcal{H}[a]|\alpha\rangle\langle\alpha|=0$. In the case of a qubit, the stochastic evolution will disturb Z , and in the process also disturb at least one of the quadratures X or Y . If we measure in the X quadrature while also knowing $\text{Tr}[\rho Y] = 0$, then the stochastic dynamics will only affect the X and Z expectation values of the system.

Digression on stochastic calculus

In the previous section, we used the random variable ΔW and then abruptly changed it into the quantity dW . Although this turns out to be a reasonable thing to do, understanding the properties of dW and how to manipulate formulas involving dW requires the use of principles

from stochastic calculus. Although we cover some basic principles here, see [49, 50] for more rigorous introductions. Central to stochastic calculus is the concept of a Wiener process, $W(t)$, which can be thought of as the continuous limit of a random walk. Each $W(t)$ is an element from a set of Wiener processes $\{W\}$ all obeying the same set of rules. Given any time interval $[t_1, t_2]$ and W , the quantity $\Delta W = W(t_2) - W(t_1)$ must satisfy the conditions $\langle \Delta W \rangle_W = 0$ and $\langle \Delta W^2 \rangle_W = t_2 - t_1$. The process also is independent in time in the sense that for two non-overlapping intervals ΔT_1 and ΔT_2 , their stochastic increments ΔW_1 and ΔW_2 are independent ($\langle \Delta W_1 \Delta W_2 \rangle_W = 0$). Examples of such $W(t)$ are shown in Fig. 3.4. One can write $W(t)$ as the integral of the quantity dW , which is the limiting infinitesimal of ΔW :

$$W(t) = \int_0^t dW \quad (3.42)$$

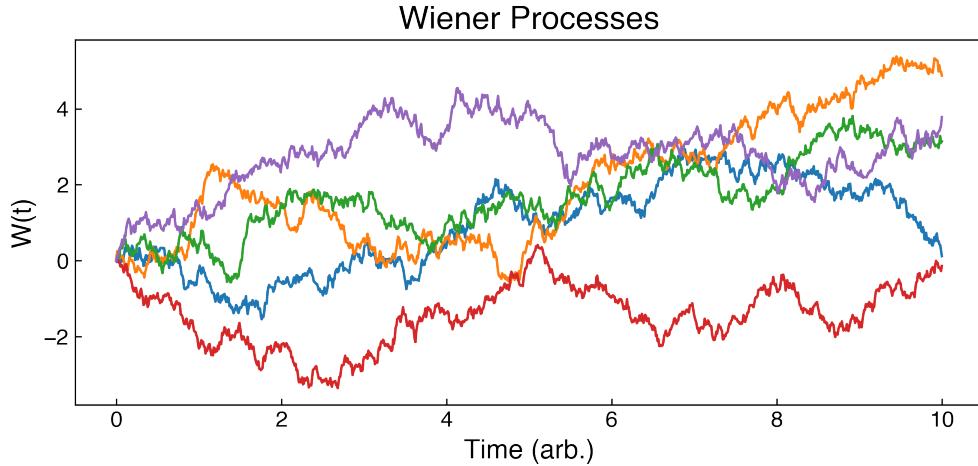


Figure 3.4: Sample Wiener processes. Such processes occur in many contexts, including when integrating vacuum noise coming down a transmission line.

In the finite interval case, $\langle \Delta W^2 \rangle_W = \Delta T$, but in the infinitesimal limit, an even stronger condition holds where the expectation brackets are unneeded:

$$dW^2 = dt \quad (3.43)$$

This may be seen by considering the variance on the variance of ΔW

$$\Delta \Delta W = \langle (\Delta W^2)^2 \rangle_W - \langle \Delta W^2 \rangle_W^2 = 3\Delta T^2 - \Delta T^2 = 2\Delta T^2 \quad (3.44)$$

when summing $\Delta \Delta W$ calculated over increasingly smaller partitions of a time interval, $\Delta \Delta W$ vanishes. We have actually used this property in the previous section when judiciously dropping terms of order $(r^2 - 1)\Delta T$.

Another technique we unwittingly borrowed from stochastic calculus was Taylor expanding to second order when using the chain rule. For example if we want to find variations in $G(X)$ where

$$dX(t) = a(t)dt + b(t)dW, \quad (3.45)$$

then we would expand G and use $dX^2 = b^2(t)dt$ to find

$$dG = G'(X)dX + \frac{G''(X)}{2}dX^2 \quad (3.46a)$$

$$= \left(G'(X)a(t) + \frac{G''(X)}{2}b^2(t) \right) dt + b(t)G'(X)dW \quad (3.46b)$$

This expansion is known as Itô's lemma, and underpins the Itô formulation of stochastic calculus. Although this expansion may seem straightforward, there are also less intuitive aspects of stochastic calculus. In particular, when solving equations of the form

$$dX(t) = a(X, t)dt + b(X, t)dW, \quad (3.47)$$

where a and/or b depend on X , integration is tricky both analytically and numerically, and one needs to take care to use the correct definition of the integral. An alternative formulation of stochastic calculus uses a slightly different integral known as the Stratonovich integral, but whose differential form does not follow Itô's lemma. We will not discuss this form, but it is often used in place of or in conjunction with the Itô form.

Finally, we briefly introduce the concept of multiple Wiener increments acting as the basis vectors for a vector space. If one has a vector of independent Wiener increments $d\vec{W} = (dW_1, \dots, dW_n)$, then multiplication forms an inner product:

$$dW_i dW_j = \delta_{ij} dt \quad (3.48)$$

Among other things, this means that one can transform a vector of Wiener increments into other bases.

Measurement inefficiency

Although the stochastic master equation in Eq. 3.37 is valid for a perfectly efficient measurement where all the information leaving the qubit makes its way to the observer, in reality physical measurements are noticeably lossy. Physically, one can conceptualize this process in two distinct ways. In the first case, one could imagine two simultaneous measurements on the system, in which case the system would propagate as one state update equation with two voltage records. Such inefficiencies can model internal losses of a cavity, which can be thought of decay through an unmonitored port. The equations of motion for this process are:

$$d\rho = \sum_{i=1}^2 \gamma_i \mathcal{D}[S]\rho dt + \sqrt{\gamma_i} \mathcal{H}[S]\rho dW_i \quad (3.49a)$$

$$V_i dt = \sqrt{\gamma_i} \langle S + S^\dagger \rangle dt + dW_i \quad (3.49b)$$

Consider that we only have access to V_1 . If $\gamma = \gamma_1 + \gamma_2$ is the total measurement strength of both measurements, we call the fraction $\eta = \gamma_1/\gamma$ our quantum efficiency. Since we don't have access to dW_2 , we remove this term from the master equation and are left with

$$d\rho = \gamma \mathcal{D}[S]\rho dt + \sqrt{\eta\gamma} \mathcal{H}[S]\rho dW_1 \quad (3.50a)$$

$$V_1 dt = \sqrt{\eta\gamma} \langle S + S^\dagger \rangle dt + dW_1 \quad (3.50b)$$

Note the stochastic term in 3.50a shrinks as the quantum efficiency drops, so it cannot "combat" the decohering Lindbladian as well. Additionally, the average voltage magnitude drops while the noise on the voltage stays at the same level, resulting in worse SNR for a given total dephasing rate.

Equivalently, one can consider inefficiency based on a single outgoing voltage being put through a beam splitter which transmits a fraction η of the power. We will consider the other input of the beamsplitter to be the vacuum state, though if this input is in a thermal state, it will add extra noise and further lower the efficiency. This beamsplitter model of quantum efficiency is realized in physical systems where lossy lines and circulators act as beam splitters with the environmental temperature acting as the second input. In this model, our initial master equation is

$$d\rho = \gamma \mathcal{D}[S]\rho dt + \sqrt{\gamma} \mathcal{H}[S]\rho dW_1 \quad (3.51a)$$

$$V dt = \sqrt{\eta} (\sqrt{\gamma} \langle S + S^\dagger \rangle dt + dW_1) + \sqrt{1-\eta} dW_2 \quad (3.51b)$$

One can verify that $dW = \sqrt{\eta}dW_1 + \sqrt{1-\eta}dW_2$ is indeed a Wiener process itself. Unfortunately, we can't simply plug $dW_1 = dW/\sqrt{\eta} - dW_2\sqrt{(1-\eta)/\eta}$ into Eq. 3.51a and drop dW_2 because dW is correlated with dW_2 . Instead, we need to treat (dW_1, dW_2) as a vector, which we transform to a new basis (dW, dV) using a unitary matrix:

$$\begin{bmatrix} dW \\ dV \end{bmatrix} = \begin{bmatrix} \sqrt{\eta} & \sqrt{1-\eta} \\ -\sqrt{1-\eta} & \sqrt{\eta} \end{bmatrix} \begin{bmatrix} dW_1 \\ dW_2 \end{bmatrix} \quad (3.52)$$

We invert the transformation and find that

$$dW_1 = \sqrt{\eta}dW - \sqrt{1-\eta}dV \quad (3.53)$$

with independent Wiener increments dW and dV . We can now plug this result into Eq. 3.51a and drop the dV term to get

$$d\rho = \gamma \mathcal{D}[S]\rho dt + \sqrt{\eta\gamma} \mathcal{H}[S]\rho dW \quad (3.54a)$$

$$V dt = \sqrt{\eta\gamma} \langle S + S^\dagger \rangle dt + dW \quad (3.54b)$$

which is in agreement with Eq. 3.50.

Heterodyne detection

A standard alternative to homodyne measurement is heterodyne measurement. In this configuration, one measures both quadratures of the incoming field simultaneously. This can be accomplished theoretically by putting a beam splitter on the outgoing field and measuring one of the arms in the I quadrature and measuring the other arm in the Q quadrature [47]. When the cavity's output field passes through the beam splitter, it is combined with vacuum noise on the second input port. This process prevents both quadratures from being measured with the same SNR as a single quadrature without a beam splitter, maintaining the Heisenberg uncertainty principle. Motivated by this picture, we can write down the equations of motion for heterodyne:

$$d\rho = \gamma \mathcal{D}[S]\rho dt + \sqrt{\frac{\eta\gamma}{2}} \mathcal{H}[S]\rho dW_I + \sqrt{\frac{\eta\gamma}{2}} \mathcal{H}[iS]\rho dW_Q \quad (3.55a)$$

$$V_I dt = \sqrt{\frac{\eta\gamma}{2}} \langle S + S^\dagger \rangle dt + dW_I \quad (3.55b)$$

$$V_Q dt = \sqrt{\frac{\eta\gamma}{2}} \langle i(S - S^\dagger) \rangle dt + dW_Q \quad (3.55c)$$

Here, as before, ρ is the system density matrix. A similar effect can be achieved by directly amplifying both quadratures of the emitted field using a phase preserving amplifier such as a Josephson traveling wave parametric amplifier (TWPA) [51]. The amplified signal can now be interpreted as a complex voltage $\tilde{V} = V_I + iV_Q$ with each V encoded in a quadrature of the signal.⁴ Although one inevitably adds noise in each quadrature during amplification since the two quadratures do not commute, this effect is informationally compensated by having information in both quadratures. Notice that if one ignores one of the measured quadratures, the quantum efficiency appears to be cut in half. If one is interested in a problem such as dispersive qubit readout (as described in the next section), the state discrimination information is always in one quadrature, so it is often loosely said that the maximum quantum efficiency for readout with a heterodyne detector is 50%. However, this is a matter of interpretation, so it is good to always check and be clear when talking about efficiency.

If one has a homodyne detector such as a JPA, it can be converted into a heterodyne detector by detuning the pump frequency from the signal frequency. In this manner, it will alternately amplify the I quadrature and the Q quadrature. So long as this detuning is fast compared to the dynamics of the system, this procedure implements a good approximation to a heterodyne detector.

⁴Although we implicitly worked in the rotating frame of the resonator or qubit, the physical voltage is oscillating at a characteristic frequency ω of the system. This physical voltage is real, but can be decomposed into $V = V_I \cos(\omega t + \phi) + V_Q \sin(\omega t + \phi)$ where the phase ϕ is set by a reference oscillator. It is mathematically convenient to represent this voltage as a complex quantity $\tilde{V} = V_I + iV_Q$.

3.3 Qubit readout with a cavity

In a typical superconducting system, the qubits are not directly coupled to a readout bus, but are instead dispersively coupled to a cavity with strength χ as in Eq. 2.45. The readout cavity is then driven with a classical drive \mathcal{E} , which is generally time dependent (possibly with a detuned frequency). We present this Hamiltonian here in a rotating frame such that there is no non-qubit dependent $a^\dagger a$ term.

$$H = \chi \sigma_z a^\dagger a + i (\mathcal{E}(t) a^\dagger - \mathcal{E}^*(t) a) \quad (3.56)$$

To make the dynamics tractable here, we assume that the qubit is undriven and does not undergo decay. In this case we can write down a density matrix ansatz for the qubit-cavity system where each of the qubit basis states $|j\rangle$ is associated with a coherent state $|\alpha_j\rangle$:

$$\tilde{\rho} = \sum_{jk} \tilde{\rho}_{jk} |j\alpha_j\rangle \langle k\alpha_k| = \sum_{jk} \tilde{\rho}_{jk} |j\rangle \langle k| \otimes |\alpha_j\rangle \langle \alpha_k| \quad (3.57)$$

We now aim to find equations of motion for $\tilde{\rho}_{jk}$ and α_j under the master equation

$$d\tilde{\rho} = -i[H, \tilde{\rho}]dt + \kappa D[a]\tilde{\rho}dt + \sqrt{\eta\kappa}H[a]\tilde{\rho}dW. \quad (3.58)$$

To help find the dynamics, we use an identity for the variations of a coherent state $|\alpha\rangle$:

$$d|\alpha\rangle = d \left(e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle \right) \quad (3.59a)$$

$$= d \left(e^{-|\alpha|^2/2} e^{\alpha a^\dagger} |0\rangle \right) \quad (3.59b)$$

$$d|\alpha\rangle = (-\Re[\alpha^* d\alpha] + a^\dagger d\alpha) |\alpha\rangle \quad (3.59c)$$

to expand our ansatz in differential form $d\tilde{\rho}$:

$$d\tilde{\rho} = \sum_{jk} |j\rangle \langle k| \left[(d\tilde{\rho}_{jk} - \Re[\alpha_j^* d\alpha_j + \alpha_k d\alpha_k^*] \tilde{\rho}_{jk}) |\alpha_j\rangle \langle \alpha_k| + \tilde{\rho}_{jk} (d\alpha_j a^\dagger |\alpha_j\rangle \langle \alpha_k| + |\alpha_j\rangle \langle \alpha_k| a d\alpha_k^*) \right]. \quad (3.60)$$

Along with the dynamics of the full density matrix $\tilde{\rho}$, we will consider the dynamics of the reduced density matrix of the qubit

$$\rho = \text{Tr}_{cav}[\tilde{\rho}] = \sum_{jk} \rho_{jk} |j\rangle \langle k| \langle \alpha_k | \alpha_j \rangle, \quad (3.61)$$

where we trace out the cavity resulting in the relation

$$\rho_{jk} = \tilde{\rho}_{jk} \langle \alpha_k | \alpha_j \rangle. \quad (3.62)$$

Although ρ is similar to $\tilde{\rho}$, the off diagonal elements are suppressed by a factor of $\langle \alpha_0 | \alpha_1 \rangle$ which will have consequences in how we account for dephasing of the qubit.

Under the ansatz in Eq. 3.57, the qubit-cavity state is completely represented by the 2x2 density matrix $\tilde{\rho}_{jk}$ (or alternatively the reduced density matrix ρ_{jk}) and two complex numbers (α_0, α_1) representing the coherent states. Going through each term in Eq. 3.58, we find their individual contribution to the dynamics of α_j , $\tilde{\rho}_{jk}$, and ρ_{jk} by expanding the term with our ansatz (Eq. 3.57) and matching the result with the ansatz's differential form (Eq. 3.60). We find that all terms preserve the form of the ansatz. After determining their individual contributions, we present the full dynamics.

Dispersive interaction Combining the dispersive term from the equations of motion and our ansatz, we get

$$-i\chi[\sigma_z a^\dagger a, \tilde{\rho}] = -i\chi \sum_{jk} \tilde{\rho}_{jk} |j\rangle \langle k| \otimes [(-1)^j \alpha_j a^\dagger | \alpha_j \rangle \langle \alpha_k | + | \alpha_j \rangle \langle \alpha_k | a (-1)^k \alpha_k^*]. \quad (3.63)$$

Matching terms of Eq. 3.63 to the differentiated ansatz Eq. 3.60, we find

$$-i\chi \tilde{\rho}_{jk} (-1)^j \alpha_j a^\dagger |j\alpha_j\rangle \langle k\alpha_k| = \tilde{\rho}_{jk} d\alpha_j a^\dagger |j\alpha_j\rangle \langle k\alpha_k| \quad (3.64a)$$

$$0 = (d\tilde{\rho}_{jk} - \Re[\alpha_j^* d\alpha_j + \alpha_k d\alpha_k^*] \tilde{\rho}_{jk}) |j\alpha_j\rangle \langle k\alpha_k| \quad (3.64b)$$

leading to the dynamics:

$$d\alpha_j = (-1)^j (-i\chi \alpha_j) dt \quad (3.65a)$$

$$d\tilde{\rho}_{jk} = \Re [-(-1)^j |\alpha_j|^2 i - (-1)^k |\alpha_k|^2 i] \tilde{\rho}_{jk} dt = 0 \quad (3.65b)$$

As expected, the coherent states rotate at different frequencies, but the combined qubit-cavity system does not dephase. We can also look at the reduced qubit dynamics obtained by tracing Eq. 3.63 over the cavity state:

$$d\rho_{jk} = i\chi \alpha_j \alpha_k^* [(-1)^k - (-1)^j] \langle \alpha_k | \alpha_j \rangle \tilde{\rho}_{jk} dt \quad (3.66a)$$

$$d\rho_{10} = 2i\chi \alpha_1 \alpha_0^* \rho_{10} dt \quad (3.66b)$$

$$= (2i\chi \Re[\alpha_1 \alpha_0^*] \rho_{10} - 2\chi \Im[\alpha_1 \alpha_0^*] \rho_{10}) dt \quad (3.66c)$$

$$d\rho_{00} = d\rho_{11} = 0 \quad (3.66d)$$

In this reduced basis, we see the effects of the AC stark shift on the qubit inducing a rotation of the qubit coherence (first term in Eq. 3.66c), as well as the effects of dephasing (second term in 3.66c). Note that the sign of the dephasing term is dependent on whether the coherent states associated with each basis state are moving away from each other or closer together. This can be shown by looking at the time derivative of the $|\alpha_0 - \alpha_1|^2$:

$$d|\alpha_1 - \alpha_0|^2 = d|\alpha_0|^2 + d|\alpha_1|^2 - d(\alpha_1 \alpha_0^*) - d(\alpha_0 \alpha_1^*) \quad (3.67a)$$

$$= -2\Re[d(\alpha_1 \alpha_0^*)] = -2\Re[d\alpha_1 \alpha_0^* + \alpha_1 d\alpha_0^*] \quad (3.67b)$$

$$= -4\chi \Re[i\alpha_1 \alpha_0^*] dt \quad (3.67c)$$

$$= 4\chi \Im[\alpha_1 \alpha_0^*] dt \quad (3.67d)$$

We can see that when the distance between the coherent states is increasing, $4\chi\Im[\alpha_1\alpha_0^*]$ is positive, resulting in an effective dephasing of the reduced qubit density matrix. On the other hand, when the distance between the coherent states is decreasing (due to the dispersive interaction), there is non-Markovian recoherence of the reduced qubit density matrix.

Cavity drive The cavity drive term is not qubit dependent, so we only need to look at the dynamics of the cavity part of the full density matrix:

$$-i [i(\mathcal{E}a^\dagger - \mathcal{E}^*a), \tilde{\rho}] = -i [i(\mathcal{E}a^\dagger - \mathcal{E}^*a), |\alpha_j\rangle\langle\alpha_k|] \otimes \sum_{jk} \tilde{\rho}_{jk} |j\rangle\langle k|, \quad (3.68a)$$

$$-i [i(\mathcal{E}a^\dagger - \mathcal{E}^*a), |\alpha_j\rangle\langle\alpha_k|] = -(\mathcal{E}^*\alpha_j + \mathcal{E}\alpha_k^*) |\alpha_j\rangle\langle\alpha_k| + \mathcal{E}a^\dagger |\alpha_j\rangle\langle\alpha_k| + |\alpha_j\rangle\langle\alpha_k| a\mathcal{E}. \quad (3.68b)$$

Matching terms of Eq. 3.68 to the differentiated ansatz Eq. 3.60, we find

$$\mathcal{E}\tilde{\rho}_{jk}a^\dagger |j\alpha_j\rangle\langle k\alpha_k| = \tilde{\rho}_{jk}d\alpha_j a^\dagger |j\alpha_j\rangle\langle k\alpha_k|, \quad (3.69a)$$

$$-(\mathcal{E}^*\alpha_j + \mathcal{E}\alpha_k^*)\tilde{\rho}_{jk} |j\alpha_j\rangle\langle k\alpha_k| = (d\tilde{\rho}_{jk} - \Re[\alpha_j^*d\alpha_j + \alpha_k d\alpha_k^*]\tilde{\rho}_{jk}) |j\alpha_j\rangle\langle k\alpha_k|. \quad (3.69b)$$

The effect of the cavity drive is to displace the coherent states in a qubit-independent manner:

$$d\alpha_j = \mathcal{E} dt \quad (3.70a)$$

$$d\tilde{\rho}_{jk} = (-(\mathcal{E}^*\alpha_j + \mathcal{E}\alpha_k^*) + \Re[\mathcal{E}\alpha_j^* + \mathcal{E}^*\alpha_k]) \tilde{\rho}_{jk} dt \quad (3.70b)$$

$$= i\Im[\mathcal{E}^*(\alpha_k - \alpha_j)]\tilde{\rho}_{jk} dt \quad (3.70c)$$

Additionally, the off-diagonal elements of the density matrix pick up a geometric phase from the act of translating different coherent states with respect to one another. In the reduced density matrix of just the qubit, this phase disappears, and in fact when tracing out the cavity in 3.68, one is left with

$$d\rho = 0 \quad (3.71)$$

Cavity dissipation Similarly to the displacement term, the dissipation term is not qubit dependent, so it is sufficient to consider its action on the cavity.

$$\kappa\mathcal{D}[a]\tilde{\rho} = \kappa\mathcal{D}[a] (|\alpha_j\rangle\langle\alpha_k|) \otimes \sum_{jk} \tilde{\rho}_{jk} |j\rangle\langle k|, \quad (3.72a)$$

$$\kappa\mathcal{D}[a] (|\alpha_j\rangle\langle\alpha_k|) = \kappa\alpha_j\alpha_k^* |\alpha_j\rangle\langle\alpha_k| - \frac{\kappa}{2}\alpha_j a^\dagger |\alpha_j\rangle\langle\alpha_k| - \frac{\kappa}{2} |\alpha_j\rangle\langle\alpha_k| a\alpha_k^* \quad (3.72b)$$

Extracting the equations of motion in a similar manner as above, we get

$$d\alpha_j = -\frac{\kappa}{2}\alpha_j dt \quad (3.73a)$$

$$d\tilde{\rho}_{jk} = \left(\kappa\alpha_j\alpha_k^* - \frac{\kappa}{2}\Re[|\alpha_j|^2 + |\alpha_k|^2] \right) \tilde{\rho}_{jk} dt \quad (3.73b)$$

$$= \left(-\frac{\kappa}{2}|\alpha_j - \alpha_k|^2 + i\kappa\Im[\alpha_j\alpha_k^*] \right) \tilde{\rho}_{jk} dt \quad (3.73c)$$

As expected, we see that the dissipation term damps the cavity field with time constant $\kappa/2$. If α_0 and α_1 have different phases, there is an additional geometric phase from the $i\kappa\Im[\alpha_j\alpha_k^*]\tilde{\rho}_{jk}$ term associated with the dissipation in Eq. 3.73c. However this phase is only relevant in the full qubit-cavity system. Tracing out the cavity in Eq. 3.72, we are again left with no dynamics on the reduced qubit subspace:

$$d\rho = 0 \quad (3.74)$$

Stochastic update The final term to consider is the stochastic update term. We assume that we are measuring the outgoing wave in the I quadrature i.e., along the real axis $V \propto B + B^\dagger$. To mathematically change the axis, one would change the phase of a in the $\mathcal{H}[a]$ operator. Since this term includes the variation dW , one might expect that the ansatz expansion for $d\tilde{\rho}$ (Eq. 3.60) does not go to high enough order, and that we should consider second order variations as per Itô's lemma. However, when we consider what a higher order variation of $d\tilde{\rho}$ would look like, contributions from a non-zero $d|\alpha\rangle$ would result in terms including $a^\dagger|\alpha_j\rangle\langle\alpha_k|$ and $(a^\dagger)^2|\alpha_j\rangle\langle\alpha_k|$. In contrast, when we look at the form of the equations of motion, we find

$$\sqrt{\eta\kappa}\mathcal{H}[a]\tilde{\rho} dW = \sqrt{\eta\kappa}\mathcal{H}[a] (|\alpha_j\rangle\langle\alpha_k|) dW \otimes \sum_{jk} \tilde{\rho}_{jk} |j\rangle\langle k|, \quad (3.75a)$$

$$\sqrt{\eta\kappa}\mathcal{H}[a] (|\alpha_j\rangle\langle\alpha_k|) dW = \sqrt{\eta\kappa} \left(\alpha_j + \alpha_k^* - \langle a + a^\dagger \rangle_{\tilde{\rho}} \right) |\alpha_j\rangle\langle\alpha_k| dW \quad (3.75b)$$

which is strictly proportional to $|\alpha_j\rangle\langle\alpha_k|$. Since there are no terms of Eq. 3.75b proportional to $a^\dagger|\alpha_j\rangle$ or $(a^\dagger)^2|\alpha_j\rangle$, the coherent states are not affected by the stochastic term, resulting in $d\alpha_j = 0$. The only non-zero equation of motion we need to extract is for $d\tilde{\rho}_{jk}$:

$$d\tilde{\rho}_{jk} = \sqrt{\eta\kappa} \left(\alpha_j + \alpha_k^* - \langle a + a^\dagger \rangle_{\tilde{\rho}} \right) \tilde{\rho}_{jk} dW \quad (3.76a)$$

$$= \sqrt{\eta\kappa} (\alpha_j + \alpha_k^* - \tilde{\rho}_{00}(\alpha_0 + \alpha_0^*) - \tilde{\rho}_{11}(\alpha_1 + \alpha_1^*)) \tilde{\rho}_{jk} dW \quad (3.76b)$$

$$= \sqrt{\eta\kappa} (\Re[(\alpha_0 - \alpha_1)i^{j-k}]i^{j+k} + \Re[\alpha_1 - \alpha_0](\tilde{\rho}_{00} - \tilde{\rho}_{11})) \tilde{\rho}_{jk} dW \quad (3.76c)$$

One can verify the third equality in Eq. 3.76 by explicitly plugging in the various combinations for i and j . We present the equation in this form to explicitly see that the stochastic term only depends on the differential state $\alpha_1 - \alpha_0$. Additionally, since the coherent states themselves have no motion under this term, we find that the reduced qubit state mirrors the dynamics of the the full density matrix:

$$d\rho_{jk} = \frac{\rho_{jk}}{\tilde{\rho}_{jk}} d\tilde{\rho}_{jk} \quad (3.77)$$

Full dynamics In computing the equation of motion for our the parameters in our ansatz Eq. 3.57, we have shown that the dynamics preserve the form of the ansatz. We present

the evolution of the coherent cavity states, the full density matrix, and the reduced qubit density matrix in one series of equations. We additionally add in the dynamics of the voltage record computed from Eq. 3.50b for completeness. The voltage provided here only includes the emitted field from the cavity, and no reflected tone, which can be achieved by measuring the cavity in transmission.

$$\dot{\alpha}_j = \left(-(-1)^j i\chi - \frac{\kappa}{2} \right) \alpha_j + \mathcal{E} \quad (3.78a)$$

$$\begin{aligned} d\tilde{\rho}_{jk} = i & \left(\Im[\mathcal{E}^*(\alpha_k - \alpha_j)] + \kappa \Im[\alpha_j \alpha_k^*] \right) \tilde{\rho}_{jk} dt \\ & - \frac{\kappa}{2} |\alpha_j - \alpha_k|^2 \tilde{\rho}_{jk} dt \end{aligned} \quad (3.78b)$$

$$+ \sqrt{\eta\kappa} (\Re[(\alpha_0 - \alpha_1)i^{j-k}]i^{j+k} + \Re[\alpha_1 - \alpha_0](\tilde{\rho}_{00} - \tilde{\rho}_{11})) \tilde{\rho}_{jk} dW$$

$$\begin{aligned} d\rho_{jk} = 2i\chi\alpha_j\alpha_k^*(1 - \delta_{jk})\rho_{jk} dt \\ + \sqrt{\eta\kappa} (\Re[(\alpha_0 - \alpha_1)i^{j-k}]i^{j+k} + \Re[\alpha_1 - \alpha_0](\rho_{00} - \rho_{11})) \rho_{jk} dW \end{aligned} \quad (3.78c)$$

$$V dt = 2\sqrt{\eta\kappa} (\tilde{\rho}_{00}\Re[\alpha_0] + \tilde{\rho}_{11}\Re[\alpha_1]) dt + dW \quad (3.78d)$$

Although the reduced qubit system is not directly coupled to a Markovian bath, and thus was not guaranteed to have dynamics expressible in Lindbladian form, in this case it is possible:

$$d\rho = -i\frac{\Delta_{AC}}{2} [\sigma_z, \rho] dt + \frac{\Gamma_\phi}{2} \mathcal{D}[\sigma_z]\rho dt + \sqrt{\frac{\eta\tilde{\Gamma}_\phi}{2}} \mathcal{H}[\sigma_z e^{i\delta}]\rho dW \quad (3.79a)$$

$$V dt = \left(2\sqrt{\eta\kappa}\Re[\alpha_0 + \alpha_1] + \sqrt{2\eta\tilde{\Gamma}_\phi} \langle \sigma_z \rangle \cos(\delta) \right) dt + dW \quad (3.79b)$$

In these update equations we have defined the following quantities:

$$\Delta_{AC} = 2\chi\Re[\alpha_1\alpha_0^*] \quad (3.80a)$$

$$\Gamma_\phi = 2\chi\Im[\alpha_1\alpha_0^*] \quad (3.80b)$$

$$\tilde{\Gamma}_\phi = \frac{\kappa}{2} |\alpha_1 - \alpha_0|^2 \quad (3.80c)$$

$$\delta = \arg[\alpha_0 - \alpha_1] \quad (3.80d)$$

Δ_{AC} is the AC Stark shift of the qubit, a shift in the qubit frequency dependent on the number of photons in the cavity. Γ_ϕ and $\tilde{\Gamma}_\phi$ are measurement-related dephasing rates associated with ρ and $\tilde{\rho}$ respectively. We will discuss these dephasing rates more thoroughly in the next section. δ is the angle of the difference phasor $\alpha_0 - \alpha_1$ to the measurement axis (in our case the real axis), and sets the type of backaction which will be seen on the qubit. The first term in the voltage update equation, Eq. 3.79b, generates deterministic evolution regardless of the qubit state. In steady state, this term simply amounts to a constant voltage offset.

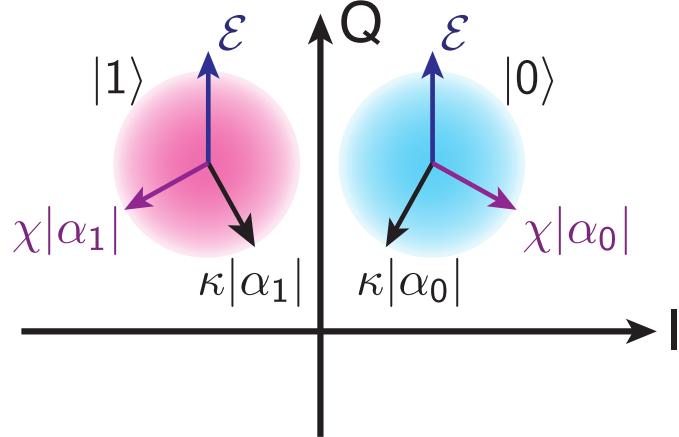


Figure 3.5: Cavity states separating during readout. The centers of the circles represent the coherent states α_j . Dynamics of the cavity states associated with each qubit eigenvector are classical in that they remain coherent states throughout the measurement. The labeled vectors on each cavity state are the different components of $\dot{\alpha}_j$. The cavity drive \mathcal{E} moves both states in the same direction, the dispersive term χ rotates the states in opposite directions, and the dissipative term κ pulls the states back towards the ground state at the center of the IQ plane. In steady state (not shown), the vectors acting on each coherent state add to zero. Maximal qubit state information is obtained by measuring along the real axis.

Measurement induced dephasing

Unlike the standard form of the master equation (Eq. 3.37), the dynamics of the reduced qubit system in Eq. 3.79 included two distinct dephasing rates: Γ_ϕ and $\tilde{\Gamma}_\phi$. The former appears in the dissipative term and is known as the measurement induced dephasing, while the latter appears in the stochastic term, which we will call the irreversible dephasing. The measurement induced dephasing represents dephasing as seen in the reduced qubit system. As shown in the analysis of Eq. 3.66c, it is not strictly positive, since collisions of the coherent states can result in coherence revivals. The irreversible dephasing is more commonly reported in terms of a quantity known as the measurement rate $\Gamma_m = 2\eta\tilde{\Gamma}_\phi$, where η is the quantum efficiency. The irreversible dephasing rate is strictly positive since decay of the coherent states is an irreversible process.

When the cavity field is in steady state and on resonance with the cavity, the coherent states are $|\alpha_j\rangle$ with

$$\alpha_j = \frac{2\mathcal{E}}{4\chi^2 + \kappa^2} (\kappa - (-1)^j 2i\chi) = \mathcal{N} (\kappa - (-1)^j 2i\chi). \quad (3.81)$$

This results in equal dephasing rates Γ_ϕ and $\tilde{\Gamma}_\phi$:

$$\Gamma_\phi = 2\chi\Im[\alpha_1\alpha_0^*] = 2\chi|\mathcal{N}|^2\Im[(\kappa + 2i\chi)^2] = 8|\mathcal{N}|^2\chi^2\kappa, \quad (3.82a)$$

$$\tilde{\Gamma}_\phi = \frac{\kappa}{2}|\alpha_1 - \alpha_0|^2 = |\mathcal{N}|^2\frac{\kappa}{2}|4i\chi|^2 = 8|\mathcal{N}|^2\chi^2\kappa. \quad (3.82b)$$

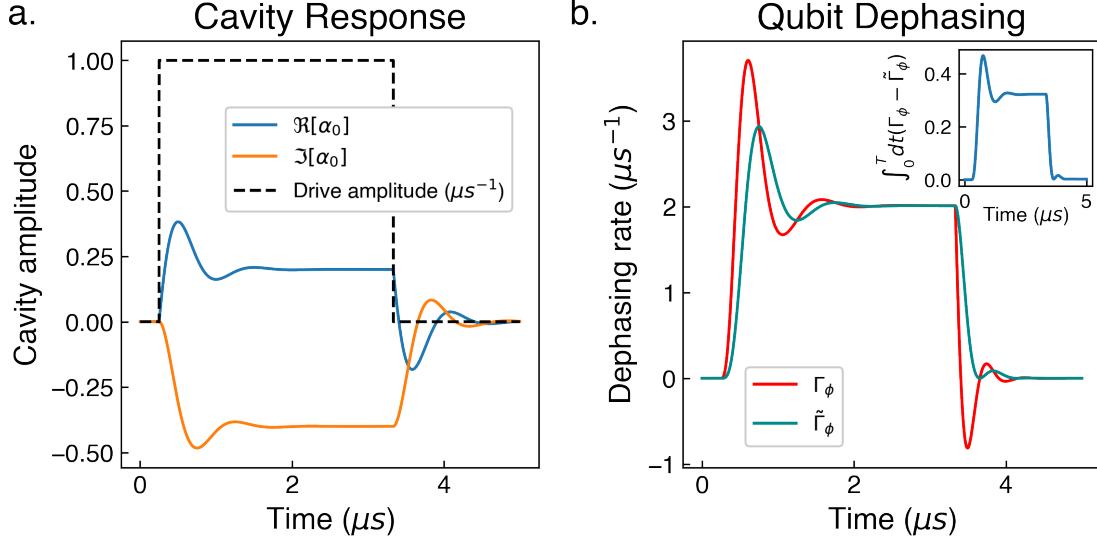


Figure 3.6: Simulation of a qubit dephasing from a measurement. Parameters used are $\chi = \kappa = 1$ MHz. (a) In-phase (real) and quadrature (imaginary) response of the cavity to a square drive pulse in the rotating frame of the drive. Only α_0 dynamics are shown, but since the drive term \mathcal{E} is real, under the cavity equations of motion in Eq. 3.78a, $\alpha_1 = \alpha_0^*$ for all time. The dotted line shows the drive \mathcal{E} in units of μs^{-1} . (b) Measurement-induced (Γ_ϕ) and irreversible dephasing ($\tilde{\Gamma}_\phi$) of the qubit during the readout pulse. The measurement induced dephasing dips below zero in the final transient, indicating coherence revival of the ρ . Inset: integrated difference $\Gamma_\phi - \tilde{\Gamma}_\phi$ remains positive throughout the readout and asymptotes to zero.

Although we do not show it here explicitly, these rates are also equal when the cavity fields are in steady state, but in a different rotating frame (i.e. when the drive tone is detuned from the cavity). Indeed, we will show the two measurement rates can only differ during cavity transients. More generally, we can discuss the integrated induced dephasing rates deriving from each of these terms. Ignoring the measurement record so that we can drop the stochastic terms in Eq. 3.78, the integrated dynamics of the coherences $|\tilde{\rho}_{10}|$ and

$|\rho_{10}|$ are obtained from

$$|\tilde{\rho}_{10}(T)| = |\tilde{\rho}_{10}(0)| \exp \left[- \int_0^T \tilde{\Gamma}_\phi dt \right]. \quad (3.83a)$$

$$|\rho_{10}(T)| = |\rho_{10}(0)| \exp \left[- \int_0^T \Gamma_\phi dt \right]. \quad (3.83b)$$

$$(3.83c)$$

Using our knowledge of the relation (Eq. 3.77) between $\tilde{\rho}$ and ρ , we can write

$$\int_{t_0}^{t_1} (\Gamma_\phi - \tilde{\Gamma}_\phi) dt = \ln \left| \frac{\langle \alpha_0(t_0) | \alpha_1(t_0) \rangle}{\langle \alpha_0(t_1) | \alpha_1(t_1) \rangle} \right|. \quad (3.84)$$

In a typical experiment of duration T , the cavity drive \mathcal{E} starts and ends at zero, as seen in Fig. 3.6. The coherent states representing the cavity response similarly start and (asymptotically) end in the ground state. As a result, the right hand side of Eq. 3.84 is zero, implying that the integrated dephasing of the system over a readout cycle is the same whether calculated using Γ_ϕ or $\tilde{\Gamma}_\phi$. From an entanglement based viewpoint, during the readout itself, the qubit becomes entangled with the cavity, at a rate characterized by Γ_ϕ . The cavity transfers its entanglement with the qubit to the transmission line as the cavity decays, characterized by the rate $\tilde{\Gamma}_\phi$. After the cavity has rung down (after about 5 μ s in Fig. 3.6a), the cavity state factorizes from the qubit, so that the qubit now only shares entanglement with the outgoing wave on the transmission line. By measuring this emitted tone, all the information extracted from qubit is retrieved (up to measurement inefficiency).

Measurement backaction

We now look at the effects of the stochastic term on the reduced qubit system, also known as the backaction. It is now useful to properly introduce the measurement rate [52] of the system,

$$\Gamma_m = 2\eta\tilde{\Gamma}_\phi = \eta\kappa|\alpha_0 - \alpha_1|^2. \quad (3.85)$$

We first look at the dynamics of the diagonal elements of the density matrix. Since $\rho_{11} = 1 - \rho_{00}$, it is sufficient to consider the dynamics of ρ_{00} :

$$d\rho_{00} = 2\sqrt{\eta\kappa}\Re[\alpha_0 - \alpha_1]\rho_{00}\rho_{11} dW \quad (3.86a)$$

$$= 2\sqrt{\Gamma_m} \cos(\delta)\rho_{00}\rho_{11} dW. \quad (3.86b)$$

This equation represents the measurement backaction on the qubit. This backaction is only realized if the coherent states α_0 and α_1 have some amount of distinguishability when projected onto the measurement quadrature. In the steady state of the measurement operation, the qubit will find a minimal variation location, i.e., where $d\rho_{00} = 0$. The two steady states in this case are the qubit basis states $|0\rangle$ and $|1\rangle$. Since there are no other dynamics of

the total system to push the state away, these states are the long time steady states of the system. We can also use this equation to argue for why the measurement rate deserves its name [53]. The entropy of the distribution $P(Z|\rho)$ is

$$S = -\rho_{00} \ln(\rho_{00}) - \rho_{11} \ln(\rho_{11}) \quad (3.87)$$

Using Itô's lemma, we get the differential form of S

$$dS = -d\rho_{00}(\ln(\rho_{00}) + 1) - d\rho_{11}(\ln(\rho_{11}) + 1) - \frac{(d\rho_{00})^2}{\rho_{00}} - \frac{(d\rho_{11})^2}{\rho_{11}} \quad (3.88)$$

Knowing from Eq. 3.86b that $d\rho_{00} = -d\rho_{11}$ only includes a stochastic term we can find the expected decrease in entropy from doing a measurement on the state by looking at the second order variation of dS .

$$\langle dS \rangle_W = -\frac{(d\rho_{00})^2}{\rho_{00}} - \frac{(d\rho_{11})^2}{\rho_{11}} \quad (3.89a)$$

$$= -4\Gamma_m \cos^2(\delta) \rho_{00} \rho_{11} dt \quad (3.89b)$$

When we are measuring fully in the information quadrature (defined as the quadrature parallel to the difference phasor $\alpha_1 - \alpha_0$, i.e., $\delta = 0$) and have no information about the Z axis of the qubit ($\rho_{00} = \rho_{11} = .5$), the expected rate of decrease in entropy for $P(Z|\rho)$ is the measurement rate ($\langle \dot{S} \rangle_W = -\Gamma_m$).

If we now consider the stochastic dynamics of ρ_{10} ,

$$(d\rho_{10})_{stochastic} = \sqrt{\eta\kappa} (\Re[\alpha_1 - \alpha_0](\rho_{00} - \rho_{11}) + i\Im[\alpha_1 - \alpha_0]) \rho_{10} dW \quad (3.90a)$$

$$= -\sqrt{\Gamma_m} (\cos(\delta) \langle \sigma_z \rangle + i \sin(\delta)) \rho_{10} dW \quad (3.90b)$$

we see that both the real and imaginary parts of the displacement $\alpha_0 - \alpha_1$ contribute to its evolution. If $\alpha_0 - \alpha_1$ is aligned with the measurement axis ($\delta = 0$), then ρ_{10} is disturbed in magnitude. Alternatively, if the measurement axis is orthogonal to $\alpha_0 - \alpha_1$, then ρ_{10} will stochastically evolve in phase and cause the qubit to rotate around the z axis. In the latter case, the stochastic portion of the evolution just amounts to random unitary evolution and does not move the state towards an eigenvector.

When including a stochastic term in a master equation, there is the associated dissipation term, in this case $\Gamma_\phi/2 \mathcal{H}[\sigma_z]\rho$. For all angles of δ , assuming that $\Gamma_m = 2\Gamma_\phi$, the combination of the stochastic term and the dissipative term will result in the qubit stochastically evolving on the surface [54]

$$\frac{1}{\beta^2} (x^2 + y^2) + z^2 = 1 \quad (3.91)$$

where $\{x, y, z\}$ are the three coordinates of the Bloch sphere and β is Bloch coordinate radius which is determined by the lack of purity in the system. We visualize this interplay of dissipation and backaction for the case of an initially pure state in Fig. 3.7. During a measurement, β will shrink due to the qubit's entanglement with the cavity, but its value can be revived by letting the cavity ring down. More significantly in practice, realistic quantum efficiencies (on the order of .5) will irreversibly shrink β .

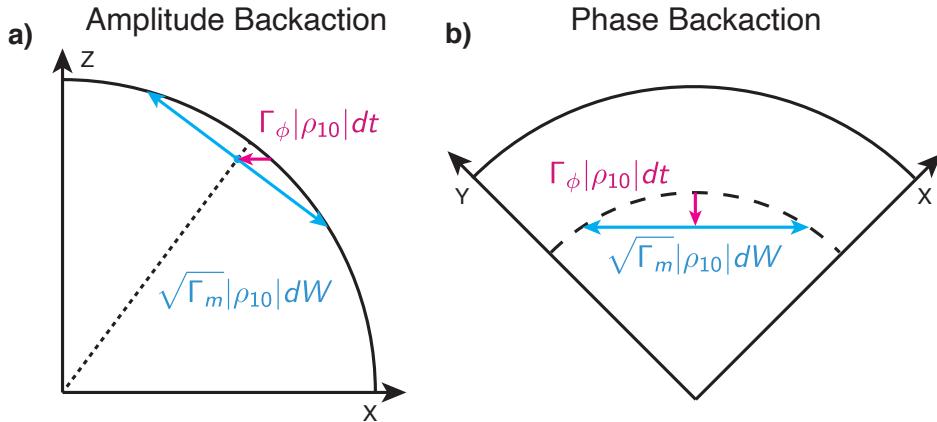


Figure 3.7: Geometry of backaction on the Bloch sphere. The effects of dephasing are shown in pink and pull the state towards the Z-axis. The effects of backaction are shown in blue and are tangent to the Bloch sphere (for a pure state). (a) In the case where the measurement axis is aligned with information quadrature, measurement backaction disturbs the state (partially) along the z axis as shown by the direction of the blue arrows. Assuming perfect efficiency ($\Gamma_m = 2\Gamma_\phi$), the stochastic motion repurifies the state. (b) In the case where the measurement axis is perpendicular to the information quadrature, phase backaction rotates the state in the XY plane to repurify the state. The dashed arc represents the initial XY radius of the state.

Voltage readout

When considering the effects of backaction in an experimental setting, the voltage record plays the crucial role of determining the stochastic increment dW , also known as the “unravelling” of the trajectory. In this section we will look at the dynamics of this record under the condition that $\delta = 0$ such that the measurement axis is aligned with the informational quadrature. First, we will assume that the measurement rate is constant and that the voltage is offset to have zero mean when $\langle \sigma_z \rangle_\rho = 0$. Although these assumptions do not qualitatively change the results in this section, they prevent undue stretching and offset factors. We use $z = \langle \sigma_z \rangle_\rho$ as the qubit coordinate and define the quantity q to be the integrated voltage

$$dq = \sqrt{\Gamma_m} z dt + dW. \quad (3.92)$$

We also recast the equations of motion for the qubit z-axis from Eq. 3.86b in terms of z :

$$dz = \sqrt{\Gamma_m} (1 - z^2) dW. \quad (3.93)$$

If the qubit is in one of the eigenstates of σ_z , then the qubit is stationary and the integrated voltage is

$$q_\pm(t) = \pm \sqrt{\Gamma_m} t + W_t \quad (3.94)$$

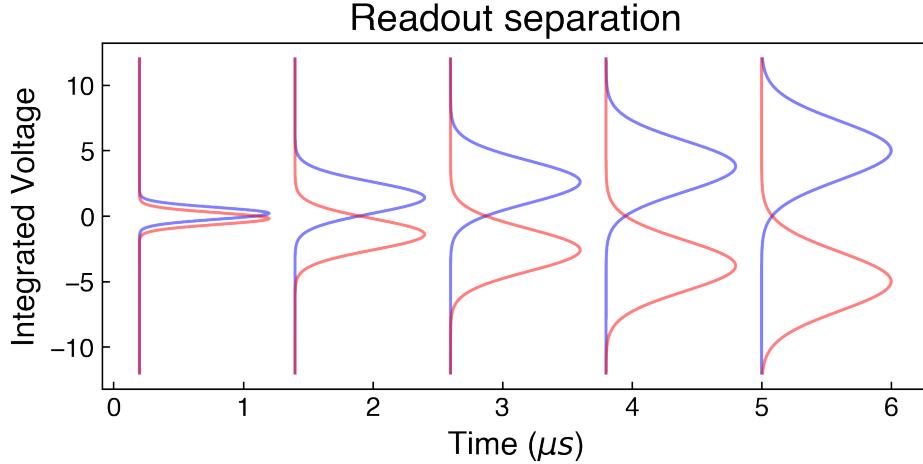


Figure 3.8: Simulation of readout separation using the solution in Eq. 3.94 with $\Gamma_m = 1 \mu\text{s}$. The red histogram shows the distribution of integrated voltages as a function of time for the qubit initialized in the excited state. The blue trace shows the same for the qubit initialized in the ground state. The separation of the histograms grows faster than their standard deviations, allowing for good readout.

where q_{\pm} represents the trajectory with the qubit being in the $+1$ or -1 eigenstate. Due to the stochastic nature of the trajectory, q_{\pm} is distributed normally with mean $\pm\sqrt{\Gamma_m}t$ and variance t . (Here we assume that we start the measurement at $t = 0$ such that $\langle W_t^2 \rangle = t$.) Since the standard deviation of q rises with \sqrt{t} and the mean rises proportional to t , the q_+ and q_- distributions separate, allowing the two states to be distinguishable as shown in Fig. 3.8. If the qubit is not initially in an eigenstate, we still see that the average integrated voltage distributions still separate conditional on the final measurement result, which can thereby be guaranteed to be ± 1 . We can write the conditional average integrated voltage as $\mathbf{E}[q_{\pm}|z_{t \rightarrow \infty}]$ where the expectation value is taken over noise realizations W and note by symmetry that

$$\mathbf{E}[q|z_{t \rightarrow \infty} = 1] = -\mathbf{E}[q|z_{t \rightarrow \infty} = -1]. \quad (3.95)$$

We can use this fact to cast this quantity into a more manageable form. For the rest of the section, expectation values unless otherwise stated are taken with respect to the stochastic process W .

$$\mathbf{E}[q|z_{t \rightarrow \infty} = 1] = \langle q(t)z(\infty) \rangle \quad (3.96a)$$

$$= \langle q(t)z(t) \rangle + \int_t^{\infty} \langle q(t)dz(\tau) \rangle \quad (3.96b)$$

$$= \langle qz \rangle \quad (3.96c)$$

Since $q(t)$ is uncorrelated with noise processes after time t and $\langle dz \rangle = 0$, the integral in Eq. 3.96c drops out. We are left with the correlator $\langle qz \rangle$, whose dynamics we can readily determine:

$$d\langle qz \rangle = \langle q \, dz \rangle + \langle z \, dq \rangle + \langle dq \, dz \rangle \quad (3.97a)$$

$$= \sqrt{\Gamma_m} \langle z^2 \rangle dt + \sqrt{\Gamma_m} (1 - \langle z^2 \rangle) dt = \sqrt{\Gamma_m} dt \quad (3.97b)$$

We thus see $\langle qz \rangle$ grow linearly with time just like $\langle q_{\pm} \rangle$. This means that the average trajectories starting from an arbitrary state but post-selected on a final measurement outcome are the same as the average trajectories when the qubit starts in an eigenstate.

Since continuous measurements are inherently noisy, we wish to characterize the signal to noise ratio of the integrated voltage. To find the noise on q conditional on a particular final outcome for an arbitrary initial z , we find the dynamics of $\langle q^2 \rangle$:

$$d\langle q^2 \rangle = 2\langle q \, dq \rangle + \langle dq^2 \rangle \quad (3.98a)$$

$$= 2\sqrt{\Gamma_m} \langle qz \rangle dt + dt \quad (3.98b)$$

$$= 2\langle qz \rangle d\langle qz \rangle + dt = d(\langle qz \rangle^2) + dt \quad (3.98c)$$

$$\langle q^2 \rangle = \langle qz \rangle^2 + t \quad (3.98d)$$

Similarly to the case where the qubit started out in an eigenstate, we have a signal power $S = \langle qz \rangle^2 = \Gamma_m t^2$ and a noise power $N = \langle q^2 \rangle - \langle qz \rangle^2 = t$. Dividing the two, we get a signal to noise ratio of

$$SNR = \Gamma_m t \quad (3.99)$$

Experimentally, the SNR can be extracted by preparing a qubit in $|0\rangle$ and $|1\rangle$ and extracting sets of integrated voltages $\{q_+\}$ and $\{q_-\}$. The SNR can then be computed from the means μ_{\pm} and standard deviations σ_{\pm} of these traces:

$$SNR = \frac{(\mu_+ - \mu_-)^2}{(\sigma_+ + \sigma_-)^2}. \quad (3.100)$$

Now that we have solved some basic voltage dynamics, we wish to answer the following question: given that we have a particular integrated readout voltage, what is the probability that the qubit is in $|0\rangle$ or $|1\rangle$? Given a prior distribution of a qubit's z-state (described by the qubit's initial z_i) and the integrated voltage q at time T , we can construct a posterior probability $P(0_f|q)$ of the qubit being in the ground state. In these probability functions,

we abbreviate $|0_i\rangle$ and $|0_f\rangle$ as 0_i and 0_f for notational clarity.

$$P(0_f|q) = \frac{P(q|0_i)P(0_i)}{P(q|0_i)P(0_i) + P(q|1_i)P(1_i)} \quad (3.101a)$$

$$= \left(1 + \frac{P(q|1_i)}{P(q|0_i)} \frac{P(1_i)}{P(0_i)}\right)^{-1} \quad (3.101b)$$

$$= \left(1 + \frac{1 - z_i}{1 + z_i} \exp\left(\frac{(q - \sqrt{\Gamma_m}T)^2 - (q + \sqrt{\Gamma_m}T)^2}{2T}\right)\right)^{-1} \quad (3.101c)$$

$$= \left(1 + \frac{1 - z_i}{1 + z_i} \exp(-2\sqrt{\Gamma_m}q)\right)^{-1} \quad (3.101d)$$

$$= \left(1 + \exp\left(-2\left(\sqrt{\Gamma_m}q + \tanh^{-1}(z_i)\right)\right)\right)^{-1} \quad (3.101e)$$

$$= \left(1 + \exp\left(-2\left(\sqrt{\Gamma_m}q + \tanh^{-1}(z_i)\right)\right)\right)^{-1} \quad (3.101f)$$

Briefly defining the suggestive quantity $\theta = \sqrt{\Gamma_m}q + \tanh^{-1}(z_i)$, we now write out the expected final z state of the qubit which we notate as \tilde{z} , where the tilde indicates that this value is calculated through a Bayesian method and not by integration of the stochastic equations.

$$\tilde{z} = 2P(0_f|q) - 1 \quad (3.102a)$$

$$= \frac{1 - e^{-2\theta}}{1 + e^{-2\theta}} = \tanh(\theta) \quad (3.102b)$$

$$= \tanh\left(\sqrt{\Gamma_m}q + \tanh^{-1}(z_i)\right) \quad (3.102c)$$

Since $q = 0$ at time zero, as expected we find that at time zero, $\tilde{z} = z_i$. Furthermore, we can use Itô's lemma to find the stochastic equations of motion for \tilde{z} . Without loss of generality, we let $z_i = 0$, so Eq. 3.102c reduces to

$$\tilde{z} = \tanh\left(\sqrt{\Gamma_m}q\right) \quad (3.103)$$

Then,

$$d\tilde{z} = \sqrt{\Gamma_m} \operatorname{sech}^2\left(\sqrt{\Gamma_m}q\right) dq - \Gamma_m \tanh\left(\sqrt{\Gamma_m}q\right) \operatorname{sech}^2\left(\sqrt{\Gamma_m}q\right) dq^2 \quad (3.104a)$$

$$= \sqrt{\Gamma_m} (1 - \tilde{z}^2) \left(\sqrt{\Gamma_m} z dt + dW - \sqrt{\Gamma_m} \tilde{z} dt \right) \quad (3.104b)$$

We can then self-consistently set $\tilde{z} = z$ and verify we recover Eq. 3.93. This gives us the nice result that the probability distribution for Z is the same whether calculated through a Bayesian method or through a stochastic equation.⁵ Furthermore, this leaves us with a

⁵Although we do not cover the approach here, there is an elegant method to generate full state update equations using a purely Bayesian method [55]. In summary, the on-diagonal elements of the density matrix are updated based on the voltage traces according to a Bayesian rule, while the magnitudes of the off-diagonal elements roughly "go along for the ride" in the form $|\rho_{01}| \propto \sqrt{\rho_{00}\rho_{11}}$ up to measurement efficiency.

direct mapping between q and z [5]:

$$z = \tanh\left(\sqrt{\Gamma_m}q + \tanh^{-1}(z_i)\right) \quad (3.105)$$

As a caveat to this equation, a real system includes the effects of T_1 decay at rate γ , so the above equation only represents a short time approximation. Indeed, most results of this section would need reanalysis to account for the effects of T_1 , which one could approximate (up to possibly breaking the ansatz in Eq. 3.57 by mixing coherent states) through an additional decay term in Eq. 3.93:

$$(dz)_{T_1} = -\gamma z \, dt + \sqrt{\Gamma_m}(1 - z^2) \, dW \quad (3.106)$$

SNR and quantum efficiency

The calculated measure of SNR in Eq. 3.100 can be used in conjunction with a Ramsey experiment to extract the quantum efficiency of the measurement. In a typical Ramsey experiment, one prepares $|+X\rangle$ and then measures the X and Y component of the qubit to extract the qubit polarization $\rho_{10} = \langle\sigma\rangle_\rho = x + iy$. The polarization decays with rate $\gamma_\phi + \Gamma_\phi$ where γ_ϕ is the intrinsic dephasing rate. By setting up a Ramsey experiment with the initialization pulse and the tomography pulse a set time T apart, the net dephasing as a function of the applied readout pulse \mathcal{E} is

$$|\rho_{10}[\mathcal{E}]| = |\rho_{10}(0)| \exp\left[\int_0^T \gamma_\phi \, dt\right] \exp\left[\int_0^T \Gamma_\phi[\mathcal{E}] \, dt\right]. \quad (3.107)$$

Since $\int \Gamma_m dt = 2\eta \int \Gamma_\phi dt$, we can write the decay of ρ_{10} in terms of the SNR and a newly defined effective quantum efficiency $\bar{\eta}$ (which we are careful not to conflate with the actual quantum efficiency η):

$$|\rho_{10}[\mathcal{E}]| \propto \exp\left[-\frac{SNR[\mathcal{E}]}{2\bar{\eta}}\right]. \quad (3.108)$$

By varying the amplitude or duration of the readout pulse, one can then measure both ρ_{10} and the SNR for different readout strengths to extract $\bar{\eta}$. $\bar{\eta}$ is dependent on the readout pulse shape and, as we will see in the next paragraph, how we integrate the signal traces.

At this point we want to be a bit more careful and consider the effects of a time dependent measurement rate $\Gamma_m(t)$. Specifically, we can increase our SNR by replacing q with q_k , which we take to be a integral of the voltage weighted with some kernel $k(t)$:

$$q_k = \int_0^t k(\tau) V(\tau) d\tau. \quad (3.109)$$

Using this modified coordinate q_k , our new SNR dependent on kernel k is

$$SNR_k = \frac{\langle q_k z \rangle^2}{\langle q_k^2 \rangle - \langle q_k z \rangle^2} = \frac{\left(\int_0^t \sqrt{\Gamma_m} k \, dt\right)^2}{\int_0^t k^2 dt}. \quad (3.110)$$

From inspection and treating the integrals as inner products ($\langle a, b \rangle = \int ab \, dt$), we can maximize the SNR by choosing $k(t) \propto \sqrt{\Gamma_m(t)}$. In this case, the SNR is maximized at

$$SNR_{max} = \int_0^T \Gamma_m d\tau = \int_0^T 2\eta \Gamma_\phi d\tau. \quad (3.111)$$

Plugging SNR_{max} into Eq. 3.108 and comparing to Eq. 3.107, we find that $\bar{\eta} = \eta$. Therefore, $\bar{\eta}$ sets a lower bound on the quantum efficiency and is acceptable to report as an achievable quantum efficiency.

Chapter 4

Canonical Phase Measurement with Adaptive Feedback

Much of modern metrology and communication technology encodes information in traveling electromagnetic waves, typically encoded in amplitude or phase. Current hardware can perform near-ideal measurements of photon number or field amplitude, in the sense that such devices are limited by material, manufacturing, and experimental capabilities. In contrast, to date no device exists that can even in principle perform an ideal phase measurement on an incoming electromagnetic wave. Nevertheless, phase is used in a wide variety of applications from telecom to the detection of gravitational waves.

In this chapter, we describe the implementation of a single-shot canonical phase measurement on a traveling single photon wave packet using a standard homodyne (amplitude) detector in conjunction with a fast feedback loop. This method provides a more accurate estimate of the phase than heterodyne detection, which is the best non-adaptive technique for determining the phase [56]. These results demonstrate an important capability for optical quantum computing, and show that quantum feedback can both enhance the precision of a detector and enable it to measure new classes of physical observables. We use a transmon qubit and a cavity to generate the traveling wave packet in the microwave regime, which is then amplified using a JPA and digitized for processing. By doing this processing on a field programmable gate array (FPGA), we can adjust the axis of the JPA amplification in a closed feedback loop, thus implementing an adaptive protocol. We first give some introductory material in section 4.1 and then in section 4.2, we present the experimental work which first appeared in [57].

4.1 Background on canonical phase measurements

We wish to define what it means to measure phase in the context of quantum mechanics. To simplify the explanation, we will ignore the temporal duration of the incoming photon and treat the incoming photon as a state in the Hilbert space of a harmonic oscillator. Consider

the classical definition of the phase of an oscillator. The set of classical states in a cavity are the coherent states $|\alpha\rangle$ which can be labeled by $\alpha = |\alpha|e^{i\theta}$. Here, $|\alpha|$ represents the amplitude of the state while θ represents the phase we wish to measure. However, if we try to write down an observable operator that can correspond to a measurement of the phase, we run into trouble. Since we consider phase to be conjugate to photon number [58], we could try to write down a phase operator in the form:

$$\hat{\varphi} = \int_0^{2\pi} \varphi |\varphi\rangle \langle \varphi| d\varphi \quad (4.1a)$$

$$|\varphi\rangle = \frac{1}{\sqrt{2\pi}} \sum_{n=-\infty}^{\infty} e^{i\varphi n} |n\rangle \quad (4.1b)$$

Here, $|\varphi\rangle$ has uniform support on the number basis, and $\hat{\varphi}$ is a Hermitian operator. However, there are multiple problems with defining phase as being a Hermitian operator and attempts to do so have a surprisingly long history [29, 30]. A minor issue is the ambiguity of choosing integration bounds such they cover all angles. By changing the bounds, one changes the weights on each projector by offsets of $2\pi n$. However, in practice, φ enters into physical equations in cyclic terms such as $\cos(\varphi)$, in which case offsets in the definition do not matter. A more glaring issue is that the definition of $|\varphi\rangle$ in Eq. 4.1b includes negative photon number states, an obvious non-physicality. When dealing with states which only have support on large photon-number states, one can use this definition of phase since the extension to negative Fock states will not affect simple calculations. In this situation, one can use an approximate commutation relation $[\hat{n}, \hat{\phi}] = -i$ and the uncertainty relation $\Delta n \Delta \phi \geq 1/2$. Indeed, this uncertainty relation's equality is satisfied for coherent states in the large $|\alpha|$ limit [59]. However, when dealing with coherent states with small $|\alpha|$ and other low photon number states, these relations no longer hold. One could try to salvage the phase operator by keeping only the positive n terms in the sum in $|\varphi\rangle$, but this results in $\{|\phi\rangle\}$ forming a non-orthogonal basis [30].

Instead of using the traditional Hermitian observable method, we can use the POVM formalism to describe measurements resulting in collapse into states which are non-orthogonal. Here, we define the Kraus operators of the canonical phase measurement. We omit negative number states and for ease of normalization, we truncate the Hilbert space to a maximal Fock state $N - 1$:

$$\Omega_\phi = \sqrt{N} |\phi\rangle \langle \phi| \quad (4.2a)$$

$$|\phi\rangle = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} e^{i\phi n} |n\rangle \quad (4.2b)$$

One can verify that these operators satisfy the POVM normalization with I_N being the identity matrix of dimension $N \times N$:

$$\int_0^{2\pi} \Omega_\phi^\dagger \Omega_\phi \frac{d\phi}{2\pi} = \frac{1}{2\pi} \sum_{n,m=0}^{N-1} \int_0^{2\pi} e^{i(n-m)\phi} |n\rangle \langle m| d\phi = I_N \quad (4.3)$$

Interestingly, we never needed to take N to infinity here, meaning that this operator is valid even in finite dimensions. One consequence of using this definition is that, assuming one could perform a non-destructive canonical phase measurement (one such that $\rho \rightarrow \Omega \rho \Omega^\dagger / \text{Tr}[\rho \Omega^\dagger \Omega]$), there is no steady state obtained upon repeated phase measurements. Additionally, once the measurement has been performed, the system is left in a pure state $|\phi\rangle$ that is completely known to the observer. As a consequence, there is no extra measurement one could do on the system to improve one's knowledge of the phase of the initial state. Indeed, performing a canonical phase measurement on a single instance of the system is the optimal way to extract phase information [60]. In tradeoff, the canonical phase measurement gives no information on the conjugate variable of photon number. Indeed, one can construct an upper bound on how well one can extract phase information from a particular measurement scheme based on how much photon number information is gained by the measurement [61].

When discussing how well phase can be extracted from a system, it is useful to introduce the concept of the Holevo variance [62]. The Holevo variance has similar properties to the standard variance, but is defined on a periodic probability distribution (or equivalently a probability distribution on a circle). Supposing one has a distribution $P(\theta)$ on $\theta \in [0, 2\pi]$, the Holevo variance is defined as

$$V_H[P] = \frac{1}{|\langle e^{i\theta} \rangle_P|^2} - 1. \quad (4.4)$$

The Holevo variance is a natural way to quantify variations of an unbiased estimator on a circle. In the limit that P is a uniform distribution, the Holevo variance diverges. In the limit of a distribution sharply peaked on a small range of angles (away from a phase cut), the Holevo variance reduces to the standard variance. Additionally, the uncertainty principle $\Delta\phi\Delta n \geq 1/2$ holds if the variance used for the phase variable is the Holevo variance [62].

Canonical phase measurement on Two Levels

Given that phase is a continuous variable, and physical states have support on a finite subset of Fock states, one cannot extract phase information perfectly with a single instance of the unknown state. This fact becomes increasing relevant for states with low photon numbers. In our experiment, we used two Fock states of a traveling electromagnetic mode, so it is useful to consider features of phase measurements on a qubit system. To compress the temporal/spatial dimension of the waveguide, we define the Fock state wavepackets (polychromatic Fock states [63, 64]).

$$|1\rangle = \int \sqrt{u(t)} e^{-i\omega t} a_t^\dagger |0\rangle dt \quad (4.5)$$

where $|0\rangle$ is the vacuum state of the waveguide, $u(t)$ is the mode shape of the photon (which integrates to 1), and ω is the central frequency of the wavepacket. a_t^\dagger is the creation

operator for an output mode labeled by time t in the input-output formalism [65] with the commutation relations $[a_t, a_{t'}^\dagger] = \delta(t - t')$.¹ We then treat $|0\rangle$ and $|1\rangle$ as the basis vectors of a qubit system.

Suppose one wants to encode maximal information in the phase of the qubit. To do so, one would prepare the initial state

$$|\Theta\rangle = \frac{1}{\sqrt{2}} (|0\rangle + e^{i\Theta} |1\rangle) \quad (4.6)$$

and then perform a measurement to try to recover this phase as accurately as possible. At the end of the measurement protocol, one has a guess θ of the phase. Depending on the type of measurement performed, one will have different distributions for $P(\theta|\Theta)$. In the case of a canonical phase measurement, one finds that

$$P(\theta|\Theta) = \langle \Theta | \Omega_\theta^\dagger \Omega_\theta | \Theta \rangle = 1 + \cos(\theta - \Theta). \quad (4.7)$$

(Note that the normalization here is with respect to the measure $d\theta/2\pi$.) In absence of a canonical phase measurement, one can estimate the phase using other standard measurement techniques such as heterodyne or homodyne (with a photon counter, one would not recover any phase information!) Under these other measurements, the posterior distribution of θ will have the form

$$P_F(\theta|\Theta) = 1 + F \cos(\theta - \Theta). \quad (4.8)$$

where F is a number from 0 to 1, and can be used as a metric for how well one is performing the canonical phase measurement. Under the form of this distribution, there is a direct map from F to V_H , as can be verified by calculating the Holevo variance of Eq. 4.8:

$$V_H = \frac{4}{F^2} - 1. \quad (4.9)$$

In the worst case (photon detection), $F = 0$, and in the ideal case, $F = 1$. The best method to estimate phase using a non-adaptive protocol is to do heterodyne detection on the incoming photon, which has $F \approx .89$ [61]. A homodyne detector does even worse, with $F \approx .80$.

4.2 Implementation of the canonical phase measurement

To implement the canonical phase measurement, we apply quantum feedback to a Josephson parametric amplifier, adaptively changing its measurement basis 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[56]. We verify implementation of a canonical phase measurement using the entanglement between the emitted photon and its source. By confirming smaller stochastic variations

¹It is often convenient to go into a rotating frame with creation operators b_t^\dagger where $b_t^\dagger = e^{-i\omega t} a_t^\dagger$.

in the emitter's Z axis (which is correlated to the photon number), we confirm that acquisition of photon number information is suppressed. The system surpasses heterodyne detection by $15 \pm 2\%$ in Holevo variance.

As shown in Fig. 4.1A, our system consists of a transmitter, which encodes a variable Θ_{true} into the phase of a single-photon electromagnetic wavepacket, and a receiver, which uses a continuous feedback protocol to guess this phase in a single shot. A transmon embedded in a 3D aluminum cavity acts as the transmitter. We use coherent bath engineering [66] of this artificial atom to generate our photonic state, a technique which yields more control over the rate of photon emission than does direct spontaneous decay. To implement this scheme, we Rabi drive our qubit at $\Omega_R/2\pi = 20$ MHz, which creates an effective low-frequency qubit. Simultaneously, we apply a cavity sideband at $\omega_{\text{cav.}} + \Omega_R$, where $\omega_{\text{cav.}} = 7.3918$ GHz is the cavity resonance frequency. As shown in Fig. 4.1B, 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, we tune the sideband-induced coupling to generate a photon with a flat modeshape (Fig. 4.1C), which greatly ameliorates the detrimental effects of feedback delay at the receiver[61]. 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 JPA flux-pumped at twice its resonance frequency and controlled by a field-programmable gate array (FPGA), which serves as a classical feedback controller (Fig 4.1A). 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[67] (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[56]. 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. 4.1E). 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. 4.1F), 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

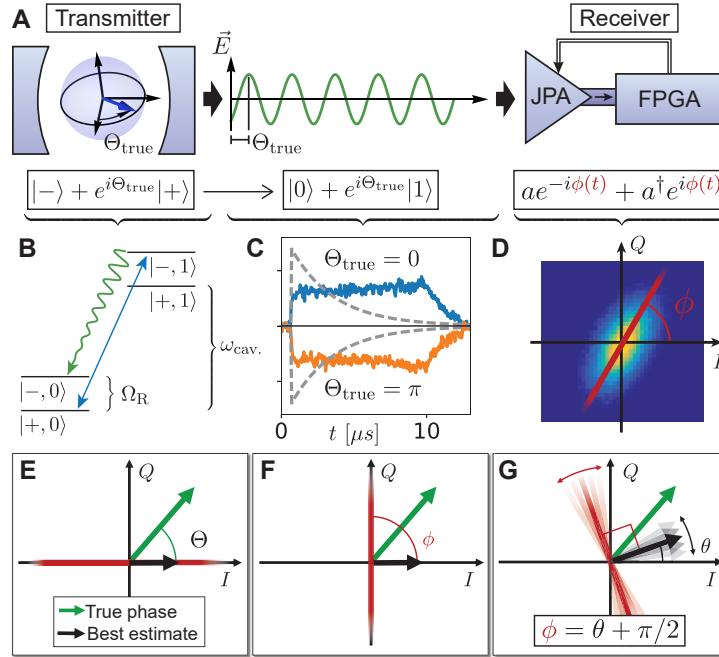


Figure 4.1: Experimental implementation. (a) Atom in a cavity, with phase Θ_{true} encoded into its dipole moment. The atom decays and emits a photon into a 1D waveguide with phase encoded into the electric field as shown. The JPA receives the photon and measures an amplitude quadrature selected by the FPGA. (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.

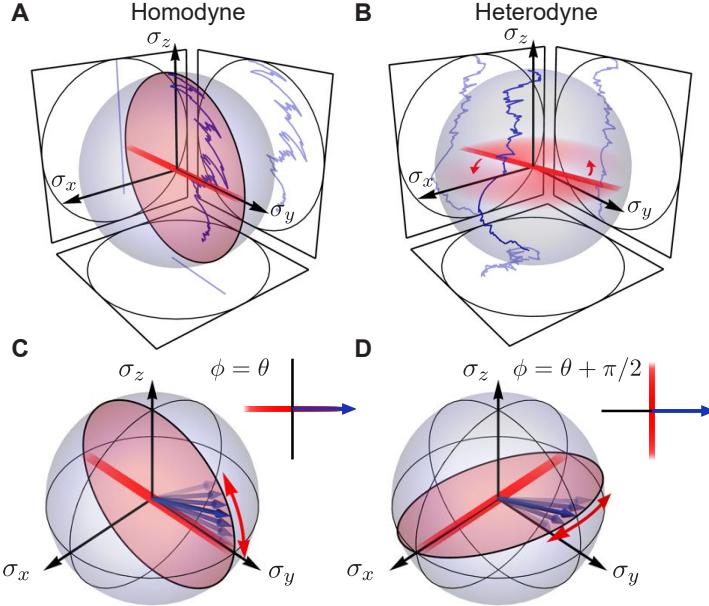


Figure 4.2: Measurement back-action and quantum trajectories. Coordinate axes are chosen so that the atom decays to $\sigma_z = -1$. (a) A single homodyne quantum trajectory ($\phi(t) = 0$). The state only propagates in the plane of the measurement axis. (b) A single heterodyne trajectory ($\phi(t) = \omega_{\text{het}} t$). 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.

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. 4.1G). 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 photon's state based on the measurement signal starting with no prior information *i.e.*, it should track the quantum trajectories of the system[68, 6] given an initially maximally mixed state. We begin by observing and verifying quantum trajectories [48] for homodyne ($\phi(t) = 0$) and heterodyne ($\phi(t) = \omega_{\text{het}} t, \omega_{\text{het}}/2\pi = 0.5 \text{ MHz}$) detection[6]. Example trajectories are plotted in Fig. 4.2A,B and tomographically validated in [57]. 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

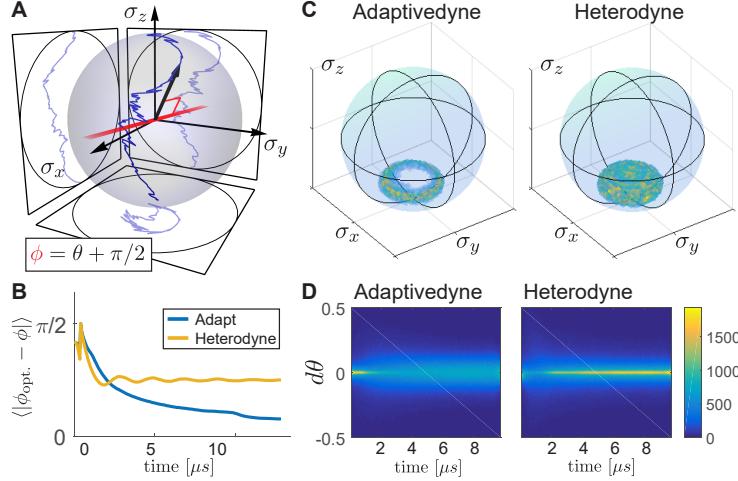


Figure 4.3: Back-action and measurement validation. (A) A single adaptive-dyne quantum trajectory. The red right-angle bracket emphasizes orthogonality between the measurement axis and the state. (B) Quality of tracking for heterodyne and adaptivedyne, where the optimal amplification phase is $\phi_{\text{opt.}} = \theta(t) + \pi/2$. Adaptivedyne significantly outperforms the heterodyne and comes close to the ideal phase by $T = 13\mu\text{s}$. The difference $\phi_{\text{opt.}} - \phi$ is cut to lie on the interval $[-\pi/2, \pi/2]$. (C) Distribution of trajectories at $t = 10\mu\text{s}$. Due to suppression of photon-number back-action, adaptivedyne trajectories cluster in a ring at late times. In contrast, heterodyne trajectories would form a roughly spherical shape inside the Bloch sphere. (D) Statistics of the phase back-action $d\theta$ for adaptivedyne and heterodyne. On average, the phase back-action is significantly larger for adaptivedyne, indicating a stronger measurement of phase.

phase measurement. Because an ideal phase measurement acquires maximal phase information and no photon-number information, it maximally disturbs the atomic dipole phase while minimally disturbing the atomic excitation probability. This effect is directly visible in the quantum trajectories, as illustrated conceptually in Fig. 4.2C 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. 4.2C). 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. 4.2D). Using this knowledge, 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. 4.3. Fig. 4.3A shows a single quantum trajectory under adaptivedyne detection, in which $\phi(t)$ is continuously adapted by the feedback controller. Fig. 4.3B shows the difference between the ideal quadrature phase

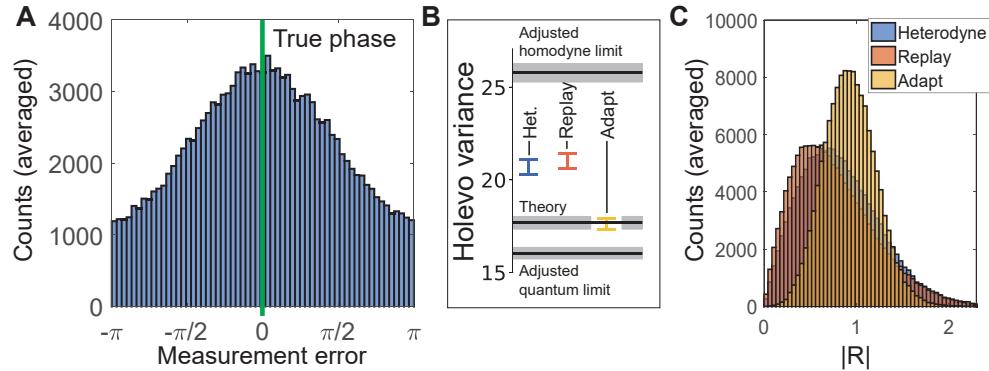


Figure 4.4: Phase-estimation performance. (A) Histogram of the difference between the measurement outcome and the true phase *i.e.* $\theta(T) - \Theta_{\text{true}}$ (B) Performance is evaluated by computing the Holevo variance of this distribution. Quantum limit (bottom black line) homodyne limit (top 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 blue distribution for Heterodyne is behind the other histograms. The distribution is significantly narrower for adaptivedyne, indicating suppression of this information channel.

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. 4.3D, with the heterodyne detection case included for comparison. It can be seen that the phase back-action $d\theta$ is significantly larger for adaptivedyne detection. Fig. 4.3C shows the ensemble statistics of the state at $t = 10 \mu\text{s}$. As observed in [6], 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[61] estimates the phase of the photon by computing the following quantity

$$R = \int_0^T e^{i\phi(t)} \sqrt{u(t)} V(t) dt \quad (4.10)$$

where $u(t)$ is the photon mode shape, and T is the duration of each experimental run and $V(t)$ 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

$\theta(T) = \arg(R)$, regardless of the function $\phi(t)$ [61]. Fig. 4.4A plots a histograms of this best estimate for adaptivedyne detection, which exhibits the $\cos(\theta - \Theta_{\text{true}})$ dependence expected theoretically[61].

We compare the performance of adaptivedyne and heterodyne detection by plotting the Holevo variance of each underlying distribution in Fig. 4.4B. We also include data for what we term replay detection; in each replay shot, $\phi(t)$ from the preceding 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%. 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, both of which directly reduce the contrast F , increasing the Holevo variance. We calibrate the adjusted quantum limit from the hypothetical homodyne Holevo limit of ~ 5.28 and, as calculated from $F = \sqrt{\pi/2}$, and the measured efficiency and photon purity. The canonical phase measurement comes significantly closer to this adjusted quantum limit than any other scheme, and is 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. 4.4. 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[61].

Single photon generation, mode shape control and feedback

To generate our photon, we use the bath engineering scheme first described in [66]. We work in a dressed frame of the qubit set by an applied Rabi drive of $\Omega_R/2\pi = 20$ MHz. As the qubit undergoes hundreds of coherent oscillations during photon wavepacket emission, it is not possible to produce a phase-stable photon unless the Rabi frequency is actively stabilized[8]. We maintain $\Omega_R/2\pi$ to within ± 2 kHz standard deviation by periodically running a pair of sequences in which we Rabi drive the qubit for $T_{\text{Rabi}} \approx 4 \mu\text{s}$, measure the qubit state and then correct the Rabi drive amplitude based on the measurement result. By timing the measurement so that $\langle \sigma_z \rangle = 0$ at the end of the calibration Rabi drive, we ensure maximum sensitivity to small drifts in Ω_R . We measure at two different points in time separated by half a period, since that a shift in Ω_R leads to a differential shift in $\langle \sigma_z \rangle$ between the two time points, as opposed to readout drift, which shifts them in the same direction. The measurement time is chosen to satisfy a trade-off between sensitivity and maximum tolerable frequency drift before slipping to another period of the Rabi oscillation.

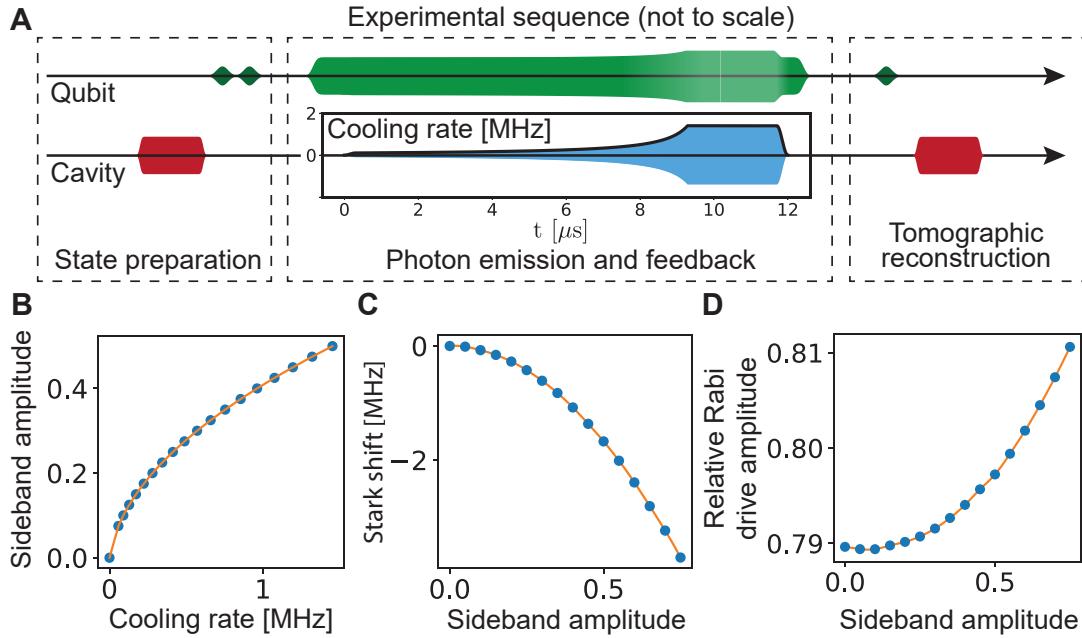


Figure 4.5: Pulse sequence and associated calibrations. (A) Pulse sequence of each shot of the experiment. Vertical axis represents the amplitude of each pulse, with the exception of the cooling sideband, in which the vertical axis is the quantitative cooling rate. Cavity-resonant pulses are shown in red, cavity sideband pulses in blue and qubit-resonant pulses in green. The discoloration and distortion of the central qubit pulse represent cooling-sideband-dependent frequency and amplitude modulation, as calibrated in (C) and (D) respectively. Effects are exaggerated for visual clarity. (B) Measurement of the bath engineering cooling rate versus sideband amplitude. (C) Measurement of the Stark shift induced by the sideband as a function of sideband amplitude (D) Sideband amplitude required to drive 20 MHz Rabi oscillations as a function of the sideband amplitude.

If the qubit is measured N times, then the uncertainty in the measured Rabi frequency is $1/(2\pi\sqrt{N}T_{\text{Rabi}})$ while the maximum tolerable frequency drift is $1/8T_{\text{Rabi}}$.

The full experimental sequence is shown in Fig. 4.5A. We first focus on the generation of a photon with the desired flat mode shape, which requires a time-dependent cooling rate $\gamma(t)$. To implement these dynamics, we measure the induced cooling rate as a function of sideband amplitude as shown in Fig. 4.5B. The sideband also induces a Stark shift on the qubit, which in turn changes the Rabi drive amplitude that achieves $\Omega_{\text{R}}/2\pi = 20$ MHz. These calibrations are plotted in Fig. 4.5C and D respectively. To emit the photon, we first ramp up the Rabi drive, and then apply the time-dependent sideband drive. As the sideband amplitude changes, we adjust the Rabi drive frequency and amplitude according to Fig. 4.5 C and D respectively. The result is the flat photon shown in Fig. 1A, which is

highly phase-stable.

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. These results are also derived in [61], but we include a variant 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 coherence, 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[47], which provides a state update from the acquisition of an infinitesimal amount of information via homodyne detection of the atom's spontaneous emission

$$d|\tilde{\psi}\rangle = \left[-\frac{1}{2}\gamma(t)\sigma^\dagger\sigma + \sqrt{\gamma(t)}e^{-i\phi(t)}\sigma V(t) \right] |\tilde{\psi}\rangle dt. \quad (4.11)$$

$V(t)$ is the measurement voltage, $\sigma = |-\rangle\langle+|$ and $|\tilde{\psi}\rangle$ is the unnormalized pure state describing the state of the atom. If we write $|\tilde{\psi}\rangle$ as $|\tilde{\psi}\rangle = c_-|-\rangle + c_+|+\rangle$, then the equation of motion for c_+ is

$$\frac{dc_+}{dt} = -\frac{1}{2}\gamma(t)c_+ \rightarrow c_+(t) = c_+(0)e^{-\frac{1}{2}\int_0^t \gamma(s)ds} \quad (4.12)$$

where we have assumed that $\gamma(t < 0) = 0$. Recall that our system decays from $|+\rangle$ to $|-\rangle$. Although in general Eq. 4.11 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 the atom was initialized with $c_+ = 1$, $c_- = 0$

$$u(t) \equiv \gamma(t)|c_+|^2 = \gamma(t)e^{-\int_0^t \gamma(s)ds}. \quad (4.13)$$

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 $\tau = 10 \mu\text{s}$ 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 $T = 13\mu\text{s}$. The $\gamma(t)$ used experimentally is shown in Fig. 4.5A. The flat portion when $\gamma/2\pi = 1.4$ MHz coincides with the portion of the photon that decays exponentially, as can be seen in Fig. 4.1C.

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

$$\frac{dc_-}{dt} = c_+\sqrt{\gamma(t)}e^{-i\phi(t)}V(t) \rightarrow c_-(t) = c_-(0) + c_+(0) \int_0^t e^{-i\phi(s)}\sqrt{u(s)}V(s)ds. \quad (4.14)$$

Notice the similarity between Eq. 4.14 and Eq. 4.10. 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_-^* . As the complex phase of c_+ remains constant, the dipole moment phase is simply

$$\theta(t) = \arg(R), \quad R(t) \equiv \int_0^t e^{i\phi(s)} \sqrt{u(s)} V(s) ds \quad (4.15)$$

in agreement with Eq. 4.10 of the main text. In principle, Eq. 4.10 and the phase measurement condition $\phi(t) = \theta(t) + \pi/2$ define the optimal protocol. For ease of implementation, this protocol may 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

$$dR = i \frac{R}{|R|} \sqrt{u(t)} V(t) dt. \quad (4.16)$$

To compute a differential equation for $|R|$, one must be aware that $V(t)dt$ is a stochastic increment, so one must apply Itô's lemma when using the chain rule. As $V(t)dt$ is normalized to have a variance dt , $(V(t)dt)^2 = dt$ and we have

$$d|R|^2 = R^*dR + R dR^* + dR^*dR = u(t)dt \quad (4.17a)$$

$$|R(t)|^2 = \int_0^t u(s) ds. \quad (4.17b)$$

Thus the time evolution of $|R|$ is deterministic. Substituting this solution into Eq. 4.16 yields

$$dR = iRP(t)V(t)dt, \quad P(t) \equiv \sqrt{\frac{u(t)}{\int_0^t u(s) ds}}. \quad (4.18)$$

Finally, we use Itô's lemma one more time to compute the differential increment of $\theta = \arg(R) = \text{Im}[\log(R)]$

$$d\theta = \text{Im}(d\log(R)) = \text{Im} \left[iP(t)V(t)dt + \frac{P(t)^2}{2}dt \right] = P(t)V(t)dt. \quad (4.19)$$

As $d\phi(t) = d\theta(t)$, Eq. 4.19 states that the instantaneous angular frequency $\dot{\theta}$ of the measurement axis is proportional to the measurement outcome $V(t)$ with a constant of proportionality $P(t)$ set by the temporal shape of the photon wavepacket. 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 the next section.

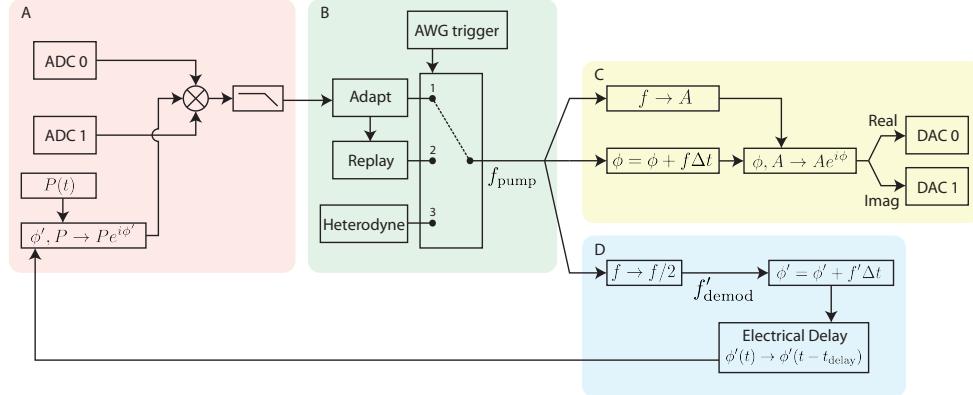


Figure 4.6: Internal logic block diagram for the FPGA. **(A)** Photon signal demodulation. The demodulation phase ϕ' is determined by the JPA's pump phase. The demodulation amplitude $P(t)$ is given by Eq. 4.16. **(B)** JPA pump frequency selection, advanced each trigger of the AWG: 1) adapting JPA frequency proportional to the incoming signal; 2) replaying JPA frequency from the previous AWG trigger; 3) heterodyning using a fixed JPA frequency. **(C)** Instantaneous JPA frequency f_{pump} increments the JPA pump phase ϕ and determines the instantaneous amplitude A through the gain calibration shown in Fig. 4.7. The real and imaginary parts of this pump tone are sent to DAC0 and DAC1 respectively for single sideband modulation. **(D)** The JPA pump frequency is halved, accumulated, and delayed to determine the demodulation phase ϕ' .

FPGA Implementation

As shown schematically in Fig. 4.6, we use an Innovative Integration X6-1000M FPGA board to control the flux pump tone for the JPA and to digitize the down-converted photon signal. The JPA pump is generated using two on-board 1 Gsample/s digital-to-analog converters (DACs) generating tones at 210 MHz, and an external I/Q mixer to perform single sideband modulation. The LO for this mixer is the frequency-doubled cavity LO, so that the JPA pump is twice the cavity frequency and phase locked to the photon. On the FPGA input side, the photon signal is down-converted to 105 MHz by an external mixer using the cavity LO, and is sampled by the FPGA board's analog-to-digital converters (ADCs) at 1 GSa/s. Inside the FPGA, the signal is further demodulated to DC and filtered (Fig. 4.6A). During adaptive feedback, the pump tone is continuously detuned from the baseline pump frequency of 210 MHz by an amount proportional to the instantaneous value of the demodulated signal quadrature according to Eq. 4.19. The FPGA board is the same as described in Appendix A, though the logic firmware is different.

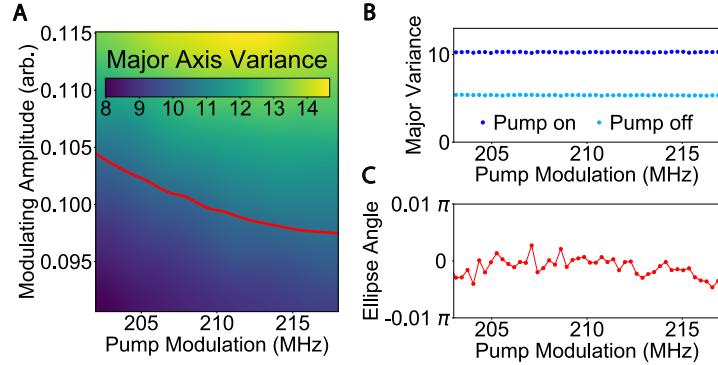


Figure 4.7: Gain calibration curve for the JPA. (A) Major axis variance (arb. units) of amplified vacuum as a function of pump frequency and amplitude. The red line represents a contour of constant variance, the gain curve. (B) Major axis variance of amplified vacuum along the gain curve with the JPA pump on as opposed to off. (C) Angle in radians of the amplified vacuum's major axis along the gain curve. An electrical delay is calibrated to keep this curve flat across frequency.

Since the JPA pump phase is continuously changing, we use the pump's instantaneous output phase delayed by the 374 ns electrical delay of the feedback cycle to determine the amplified quadrature of the input (Fig. 4.6D). To ensure that the correct quadrature is read out, we perform a separate calibration in which we sweep the pump frequency from 202 MHz to 218 MHz. We see less than a 0.5 degree variation between the pump frame delayed through the FPGA and the amplified quadrature measured from the JPA. We also ensure that the pump gain stays constant over this frequency band using the lookup table shown in Fig 4.7, which determines the pump amplitude for a given pump frequency. Using the major axis variance of the squeezed vacuum ellipse as a proxy for gain, we intersperse the calibration of this lookup table throughout our measurements to compensate for slow changes in the pump chain, which occur due to room temperature amplifier drift.

Alongside the feedback mode, the FPGA has two other modes: replay and heterodyne. In replay mode, the output pump waveform is identical to the waveform of the previous adaptive pump waveform; the JPA undergoes the same rotations as it had during adapting, but the pump phase is no longer correlated with the estimated photon phase. In heterodyne mode, the JPA pump frequency is detuned from 210 MHz by a constant heterodyne frequency, 0.5 MHz. During our measurements, the board cycles through adaptive, replay, and heterodyne, changing modes on each trigger of the AWG (Fig. 4.6B).

4.3 Outlook

Several avenues remain for future work in this area. Firstly, we have optimized the system for fair comparison between heterodyne and adaptivedyne detection in order to be sure that the observed improvements arise from feedback alone. A system that is optimized for adaptivedyne could easily yield further improvements. Quantum efficiencies as high as 80% have been demonstrated in circuit QED[69], and similar improvements could be achieved by increasing amplifier gain or adding low-loss or on-chip circulators[70]. Integrating low-temperature electronics closer to the amplifier could also significantly reduce feedback latency, which would yield immediate gains in the phase estimation efficiency.

Our system has several immediate applications to quantum information and computation. Firstly, the implementation of quantum feedback on a detector is known to allow enhanced readout of superconducting circuits[71]. 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[72]. More broadly, it is known that adaptive measurements are universal[73], meaning that many relatively simple measurement devices augmented with quantum feedback can perform any generalized measurement allowed by quantum mechanics (see section 3.1). Thus our extension of a standard amplitude measurement device to an ideal phase measurement represents a more general and exciting direction for future research.

Chapter 5

Continuous Quantum Error Correction

Quantum computing offers tantalizing prospects of performing certain algorithms faster than possible on a classical computer. However, quantum processors are inherently noisy and susceptible to decoherence from the environment and noise on applied unitaries. As one performs quantum circuits of greater gate depth, these errors accumulate until the signal to noise ratio at the circuit's output is vanishingly small. Although one can try to increase the quality of qubits and the quality of control over said qubits until the effects of errors are negligible, realistic systems do not appear to be conducive to such efforts. An alternative to making better and better qubits is to make the qubits good enough, and then cleverly encode information in these qubits so that is protected against common sources of errors. By detecting these error before they have a chance to impact the protected information, we can theoretically perform longer circuits in the protected space. Under the standard error correction formalism, the measurement (though possibly imperfect) is regarded as an instantaneous operation. As we have seen in the previous chapters though, measurement is not instantaneous. We can use the dynamics of measurements to turn quantum error correction to a continuous process. We experimentally demonstrate this capability using a three qubit bit-flip code.

5.1 Background on quantum error correction

We first give a brief primer on quantum error correction. Error correction is an extremely rich field, and there are many extensive introductory resources [74, 75, 40] for more detailed explanations and a wider range of concepts. Codes of particular interest are surface codes[76], bosonic codes[77], and holographic codes[78]. Also interesting is the Holy Grail of fault tolerant quantum error correction whereby when one's physical qubits and gates get good enough, one can make computation better by simply throwing more qubits at it[79]. However, we will not dive into these depths. Instead we present a recap of the basic concepts of quantum error correction.

Quantum information is inherently fragile. The No-Cloning theorem states that there is

no physical operation which performs the map

$$|\psi\rangle \rightarrow |\psi\rangle |\psi\rangle. \quad (5.1)$$

While this theorem is beautiful from a theoretical and informational standpoint, it is frustrating for the experimentalist trying to protect their quantum state. In classical error correction, we can simply make copies of a classical bit string to have informational redundancy. Indeed, if we think about even a supposedly atomic bit of information stored in a capacitor or on a magnetic disk, such a bit comprises enormous numbers of electrons storing the same bit of classical information. In quantum systems, such an encoding does not appear so naturally. Anything that interacts with a quantum state can siphon away information, or at the very least disturb the state.

As another complication, quantum states are analog in the sense that they are parameterized by continuous complex coefficients. The operations we apply to quantum systems are similarly analog and susceptible to imperfections. For example, if we try to perform a gate which brings a qubit to the state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, we may end up with the state $|\psi'\rangle = (\alpha + \epsilon_\alpha)|0\rangle + (\beta + \epsilon_\beta)|1\rangle$ for some small complex ϵ_α and ϵ_β . These over-rotations can add up over the course of an algorithm and produce errors.

In order to protect quantum information, we can construct a quantum error correction code. Such a code consists of a small logical quantum system embedded inside a larger physical system. Often times, both the logical system and the physical system are thought of in terms of qubits. In this case, a small number of qubits k is mapped onto to a larger number of qubits n . If \mathcal{H} is the Hilbert space of the full system (dimension 2^n), we can decompose \mathcal{H} into two important bases:

$$\text{Physical Basis} \quad \mathcal{H} = \bigotimes_{i=1}^n \mathcal{Q}_i^P \quad (5.2a)$$

$$\text{Logical Basis} \quad \mathcal{H} = \bigoplus_{\vec{s}} \mathcal{H}_{\vec{s}} \quad (5.2b)$$

Here, $\{\mathcal{Q}_i^P\}$ are physical qubits labeled by physical number i and $\{\mathcal{H}_{\vec{s}}\}$ are the protected subspaces of the code, labeled by $\vec{s} \in \{0, 1\}^{\otimes(n-k)}$ representing an error syndrome. Each element of \vec{s} represents a particular error having occurred. Every $\mathcal{H}_{\vec{s}}$ is of dimension 2^k and could be used as a subspace to protect quantum information. Without loss of generality, we take $\mathcal{H}_{\vec{s}=\vec{0}}$ to be the logical subspace. The logical state stays in this subspace until an error occurs. The other protected subspaces of definite error syndrome are known as error spaces of the code. After a detectable error, the logical state moves from the logical subspace to a corresponding error subspace.

In the context of quantum error correction, qubits are a two-state quantum degree of freedom (i.e. a Hilbert space of dimension 2). The physical qubits are usually taken to be spatially local, and are generally associated with the construct which holds the qubit degree of freedom. For example, transmons, atoms, ions, etc. are loosely referred to as qubits,

whereas in the context of quantum error correction, it is more accurate to say that these physical objects have a local degree of freedom which holds a qubit of information. The logical subspace (and error subspaces) consists of degrees of freedom which are spread across many physical qubits.

Looking at the form of Eq. 5.2b, it is not immediately obvious that the logical qubits are protected any more-so than a physical qubit. In particular, if $|\psi\rangle$ is the desired k -qubit state to protect, then the equivalent logical state on the full Hilbert space can be decomposed like $|\psi\rangle \otimes |\text{gauge qubits}\rangle$ and not $|\psi\rangle \otimes |\psi\rangle \otimes \dots \otimes |\psi\rangle$. If all operations in the Hilbert space were equally as likely to occur, then the logical qubit would be no better at protecting information than a physical qubit. This leads into two basic requirements for quantum error correction: errors in the physical basis should be local and uncorrelated; and operations on the logical subspace should be non-local in the physical basis.¹

When an error which the code protects against (say a local error on a physical qubit), occurs on a state in the logical subspace, the syndrome label \vec{s} changes and the state moves to a new protected subspace. As long as the number of errors is low, the original logical state can be recovered by measuring which error subspace the state is in, and applying corrective actions to move it back to the original logical subspace. If too many errors occur, they can eventually affect the logical state as will be described for the three qubit code in the next section.

Three qubit bit-flip code

To give a concrete example, consider the 3-qubit bit-flip code used in this work. Although the code does not protect against all single qubit errors, it demonstrates the basic principles of error correction. The logical states of the code are given by

$$|\bar{0}\rangle = |000\rangle \quad (5.3a)$$

$$|\bar{1}\rangle = |111\rangle, \quad (5.3b)$$

where kets are taken to be in the physical basis. The different protect subspaces can be specified by a set of measurement operators $\{M\}$ known as stabilizers of the code. The stabilizers have eigenvalues of ± 1 , and have constant value within each protected subspace. If we consider the state $|\psi\rangle_{\vec{s}}$ to be a k -qubit state $|\psi\rangle$ encoded in the logical subspace $\mathcal{H}_{\vec{s}}$, then we can define a minimal set of stabilizers $\{M_i\}$ such that

$$M_i |\psi\rangle_{\vec{s}} = (-1)^{s_i} |\psi\rangle_{\vec{s}}. \quad (5.4)$$

In the case of the 3-qubit code, the base stabilizers are

$$M_1 = Z_0 Z_1 \quad (5.5a)$$

$$M_2 = Z_1 Z_2 \quad (5.5b)$$

¹These conditions are not hard and set. Error correction codes can handle some spatially and temporally correlated errors, though these errors decrease code performance. Additionally, these correlations must die off with distance at some reasonable rate.

where Z_i are the Pauli Z operators in the physical basis.² These two stabilizers are not unique. Indeed, we can use these stabilizers to generate a third stabilizer $M_3 = M_1 M_2 = Z_0 Z_2$. However, measuring M_3 will give information degenerate with information from the first two stabilizers, so is not strictly necessary. By measuring the two stabilizers, we project the system into a single definite logical or error subspace, discretizing any errors which result in state support on multiple protected subspaces.

The 3-qubit code protects against single qubit bit-flips. To see this, consider the action of an X_0 error in the Hilbert space of the syndrome qubits. X_0 anti-commutes with M_1 and commutes with M_2 . Therefore, if the state $|\psi\rangle_{\vec{s}}$ is an eigenvector of all the stabilizers, then $X_0 |\psi\rangle_{\vec{s}}$ is also an eigenvector of the stabilizers, though with the eigenvalue of M_1 flipped. Similarly, an X_2 error flips the eigenvalue of M_2 , and an X_1 error flips the eigenvalues of both M_1 and M_2 . We consider the set $\{E_i = X_i\}$ to be the possible errors on system detectable by the code, labeling the errors with E_i since the following results do not depend on the errors being X errors. We also let $|\bar{a}\rangle$ and $|\bar{b}\rangle$ be basis states in the logical subspace. To see that a set of errors are not only detectable but correctable, we need to make sure that every error-state $E_i |\bar{a}\rangle$ is distinguishable from every other error state $E_j |\bar{b}\rangle$ for $a \neq b$:

$$\langle \bar{a} | E_i^\dagger E_j | \bar{b} \rangle = 0, \quad a \neq b. \quad (5.6)$$

We would also like to ensure that we can apply a correct error inversion for every error that occurs on a given logical state. Although the condition

$$\langle \bar{a} | E_i^\dagger E_j | \bar{a} \rangle = \delta_{ij} \quad (5.7)$$

is clearly sufficient to achieve distinguishability of errors occurring on a single state, it turns out that only a weaker condition is necessary[74]:

$$\langle \bar{a} | E_i^\dagger E_j | \bar{b} \rangle = C_{ij} \delta_{ab}, \quad (5.8)$$

where C_{ij} is a Hermitian matrix. We have already seen that a Pauli X error acting on a single qubit, does not commute with the two stabilizers. One can also show that the combination of any two X errors ($X_i X_j$) also does not commute with the stabilizers; Acting $X_i X_j$ on a state in the logical subspace either takes the system out of the logical subspace or preserves the state, thus satisfying Eq. 5.8.

Now that we have defined our logical states and seen how errors affect them, we would like to define the basis vectors of the error subspaces. These error spaces do not have immediately well defined bases since arbitrary SU(2) rotations in these subspaces do not affect stabilizer values or the ability of the code to correct errors. To fix their bases, we ensure that when a correctable error occurs on the logical subspace, the logical qubit basis vectors do not change labeling other than switching to a new \vec{s} . In the three qubit code using the logical subspace defined in Eq. 5.3, the logical qubit's $|0\rangle$ and $|1\rangle$ states in each $\mathcal{H}_{\vec{s}}$ are

$$\{|\bar{0}\rangle, X_0 |\bar{0}\rangle, X_1 |\bar{0}\rangle, X_2 |\bar{0}\rangle\} = \{|000\rangle, |100\rangle, |010\rangle, |001\rangle\}, \quad (5.9a)$$

$$\{|\bar{1}\rangle, X_0 |\bar{1}\rangle, X_1 |\bar{1}\rangle, X_2 |\bar{1}\rangle\} = \{|111\rangle, |011\rangle, |101\rangle, |110\rangle\}. \quad (5.9b)$$

²For consistency with the experimental work, we use zero indexing for the physical qubits.

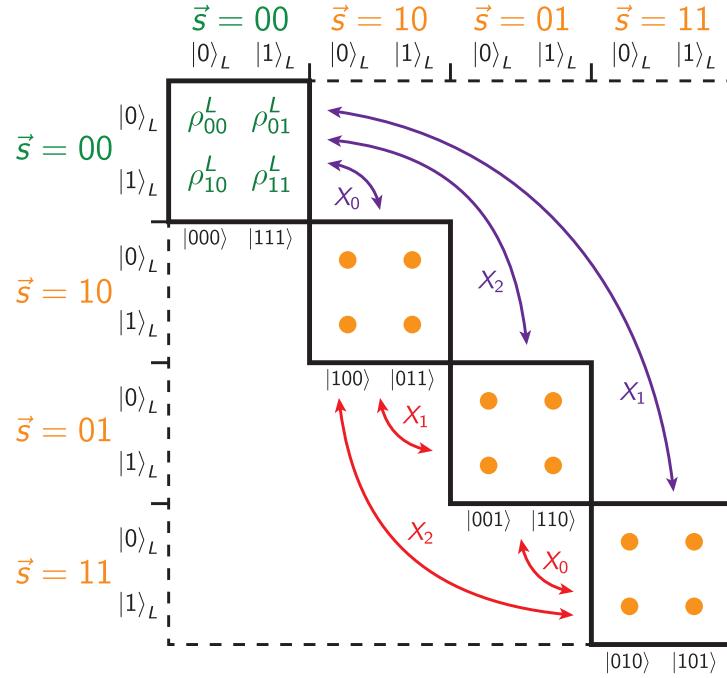


Figure 5.1: Density matrix of the three qubit bit-flip code in the logical basis. The logical subspace is shown in green and has syndrome 00. Error spaces comprise the three orange blocks with definite syndrome values other than 00. Physical X flips lead to transitions amongst the syndrome subspaces. Purple transitions directly out of the logical subspace move the logical state to an error subspace, but do not affect the logical qubit. Red transitions between two error spaces induce an X flip in the logical subspace in addition to changing subspaces.

This choice of basis vectors is dependent on which syndrome subspace was taken as the codespace ($\vec{s} = \vec{0}$). In this frame, transitions among the subspaces are shown in Fig. 5.1. If we had chosen one of the error spaces as the codespace, we would construct a basis such that transitions out of that chosen subspace would always preserve the logical qubit.

To perform operations in our logical qubit subspace, we need unitary operators which commute with all of the stabilizers, and anti-commute with each other. These operators will be identified as the \bar{X} and \bar{Z} of the logical state and in general these logical operators are not unique [75]. With this in mind we can choose

$$\bar{X} = X_0 X_1 X_2, \quad (5.10a)$$

$$\bar{Z} = Z_0 I_1 I_2 \quad (5.10b)$$

to be the logical operators for the three qubit code. One can also verify that these operators act as expected in the logical subspace given the labellings $|\bar{0}\rangle$ and $|\bar{1}\rangle$. Note that \bar{X} has

weight 3 and that \bar{Z} has weight 1. This is expected as we know that the code will detect (not correct³) X errors up to weight 2 and correct X errors of weight 1, but cannot detect Z errors at all.

Measurement in error correction

In the formulation of standard error correction, logical gates are interspersed with rounds of error correction. In each correction round, the stabilizers are measured, and then any detected errors are corrected.⁴ This process discretizes errors by forcing the system into one of the protected subspaces. As an example, suppose that we start with the state $|\bar{\psi}\rangle = \alpha |\bar{0}\rangle + \beta |\bar{1}\rangle$ which undergoes a small rotation around X_0 :

$$|\bar{\psi}\rangle \mapsto (1 + i\epsilon X_0) |\bar{\psi}\rangle \quad (5.11)$$

Before measuring the stabilizer, the state has support on both the codespace and one of the error spaces. By measuring with the $M_1 = Z_0 Z_1$ stabilizer, the system is projected definitively into either the codespace or the error space. The outcome of the correction is the same as if we had instead corrected a probabilistic X_0 gate (i.e. where $\rho \mapsto \rho + \epsilon^2 X_0 \rho X_0$). To see how this correction process can extend the lifetime of a logical state, we can consider each error channel X_i as having some probability p of occurring between rounds of error correction. We consider zero or one errors as correctable and two or three errors as a logical error. Assuming we can perform a perfect syndrome measurement, the probability of the state being correctable is

$$P_{corr} = (1 - p)^3 + 3p(1 - p)^2 = 1 - \mathcal{O}(p^2), \quad (5.12)$$

so, for low error rates, the probability of being correct increases when a round of error correction is applied.

In an experimental setting, there are different ways to measure stabilizers. A crucial aspect of stabilizers is that they are inherently joint-qubit measurements. In order to measure $Z_0 Z_1$ for use as a stabilizer, one cannot simply measure Z_0 and Z_1 individually, as doing so will destroy the coherence of the logical state. To preserve local readout, the most common way to measure stabilizers is to map their value onto the ancillary qubits in the ancilla's Z basis, as shown in Fig. 5.2. These ancilla are then projectively measured along the Z basis to determine the stabilizers' values.

An alternative method is to perform a direct joint readout on all qubits involved in the stabilizer measurement. This method avoids the need for entangling gates between ancilla

³Using the stabilizer measurements, a weight 2 error on two of the qubits is indistinguishable from a weight 1 error on the remaining qubit; if a weight 2 error such as $X_0 X_1$ occurs, this will be decoded as an X_2 error. If a “corrective” X_2 pulse is applied to the system, the net action of error and “correction” is $X_0 X_1 X_2$: a logical error.

⁴As an alternative to active error correction, the logical subspace can be reassigned to be the current measured error subspace. Under this protocol, the logical basis states of each protected subspace are reassigned from their original definitions in Eq. 5.9.

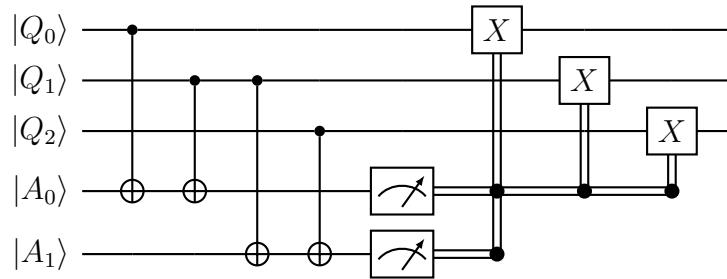


Figure 5.2: Ancilla based three qubit error correction. Stabilizer values are copied to two ancillary qubits, which are then measured. Based on the measured syndrome, correction pulses are applied.

and data qubits (the physical qubits of the code). In our experiment, we implement a direct measurement of the ZZ stabilizers using a coherent microwave tone to probe a joint readout resonator. Using this method of direct stabilizer measurement opens the possibility of performing continuous error correction [47, 80]. In contrast to discrete error correction, where rounds of error correction are periodically performed, continuous error correction is always on and reading out information about the error syndromes.

5.2 Experimental demonstration

In this experimental work we implement a continuous error correction protocol for the three qubit code. The continuous measurement is monitored by an FPGA controller that actively corrects errors as they are detected. Using this method, we achieve an average bit-flip detection efficiency of up to 91%. Furthermore, we use the protocol to increase the relaxation time of the protected logical qubit by a factor of 2.7 over the relaxation times of the bare comprising qubits. Our results showcase resource-efficient stabilizer measurements in a multi-qubit architecture and demonstrate how continuous error correction codes can address challenges related to ancilla-based errors in realizing a fault-tolerant system.

Demonstrations of discrete quantum error correction have been exhibited using various physical systems such as ion traps[81, 82, 83], defects in diamonds[84], and superconducting circuits[85, 86, 87, 88, 89, 90]. Continuous measurements have been used to study the dynamics of wavefunction collapse and, with the addition of classical feedback, to stabilize qubit trajectories and correct for errors in single qubit dynamics[10, 91, 92]. In systems of two or more qubits, direct measurements of parity can be used to prepare entangled states through measurement[93, 94, 95, 96, 97, 9]. When applied to error correction, continuous parity measurements eliminate the cycles of discrete error correction and remove the need for entangling gates[98, 99]. Errors are detected on a rolling basis, with the measurement rate as the primary limitation to how quickly errors are detected. In this experimental section,

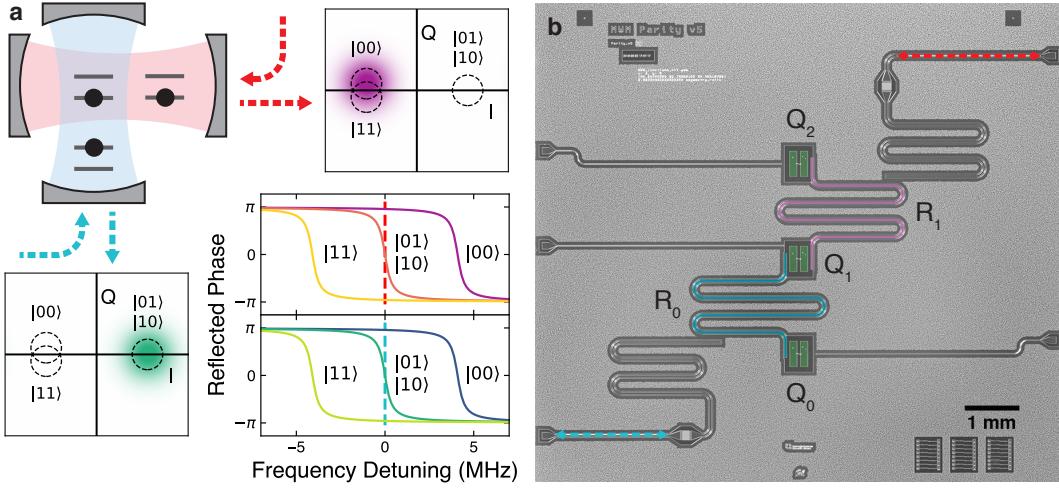


Figure 5.3: Full Parity Detection. (a) Three qubits in two cavities, with each cavity implementing a full parity measurement. Lower right: ideal phase responses of a coherent tone reflected off each cavity for different qubit states. The parity probe tones are centered on the odd-parity resonances. The phase space (IQ) plots show the ideal steady state reflected tone for the shown qubit configuration. Dashed circles are centered on all possible steady state responses. (b) Micrograph of the superconducting chip with three transmons and two joint readout resonators.

we use two continuous parity measurements and associated filtering[100] to correct bit-flip errors while maintaining logical coherence.

We realize our code in a superconducting architecture using three transmons as the physical qubits. As depicted in Fig. 5.3, we implement the ZZ parity measurements using two pairs of qubits coupled to joint readout resonators[101, 9]. The resonators are coupled to the qubits with the same dispersive coupling χ_i with i indexing the resonator, thereby making the resonator reflection response when the two qubits are in $|01\rangle$ identical to the response when the qubits are in $|10\rangle$. For each resonator, we set the parity probe frequency to be at the center of this shared odd parity resonance. To approximately implement a full parity measurement, we make the line-width κ_i (636 kHz, 810 kHz) of each resonator smaller than their respective dispersive shift χ_i (2.02 MHz, 2.34 MHz). When the qubit pair is in either $|00\rangle$ or $|11\rangle$, the resonance frequency is sufficiently detuned from the odd parity probe tone to keep the cavity population low and the reflected phase responses for the two even states nearly identical as demonstrated in the section on steady-state dephasing. After reflecting a parity tone off a cavity, the signal is amplified by a JPA in phase-sensitive mode aligned with the informational quadrature.

We implement the three qubit repetition code using two ZZ parity measurements as

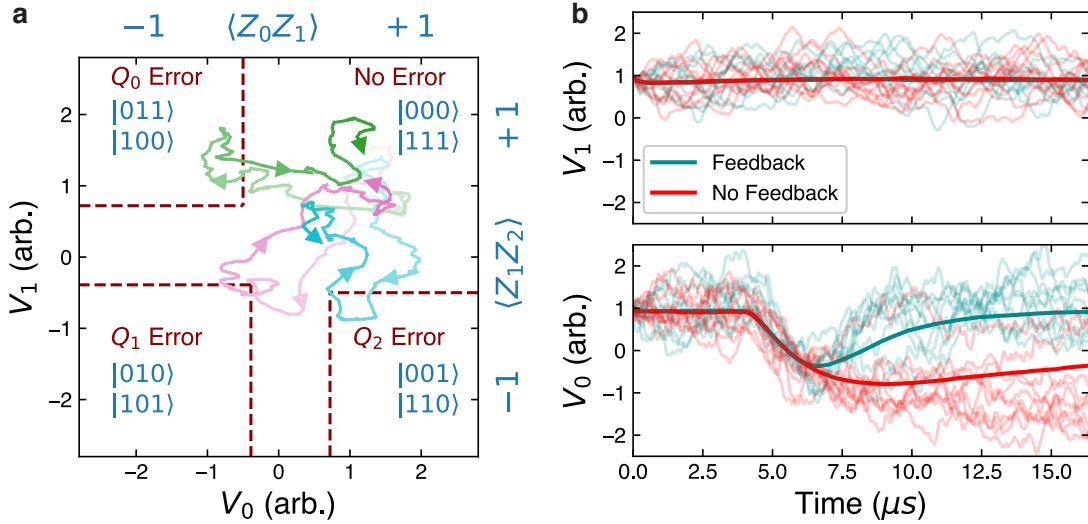


Figure 5.4: Error Correction. (a) Sample experimental voltage traces of the controller correcting induced bit flips. With no errors, both voltages remain positive. When an error occurs, one or both of the voltages flip and the cross thresholds, triggering the controller to send a corrective π pulse to bring the system back to the codespace. (b) Voltage responses to an induced flip on Q_0 with (blue) and without (red) feedback. Bold lines are averages and light lines are sample individual traces.

stabilizers: Z_0Z_1 and Z_1Z_2 . The codespace can be any of the four subspaces with definite stabilizer values, but, as in the previous section, we choose the subspace with positive parity values $(+1, +1)$ for simplicity. This choice of codespace is spanned by the logical code states $|\bar{0}\rangle = |000\rangle$ and $|\bar{1}\rangle = |111\rangle$. The three remaining stabilizer values identify error subspaces in which a qubit has a single bit-flip (X) error relative to the codespace. A change in parity heralds that the logical state has moved to a different subspace with a different logical state encoding.

Ideal strong measurements of both code stabilizers project the logical state into either the original codespace or one of the error spaces, effectively converting analog errors to correctable digital errors. In contrast, measurements with a finite rate of information extraction, like the homodyne detection used in this experiment, result in the qubit state undergoing stochastic evolution such that the logical subspaces are invariant attractors [102]. The observer receives noisy voltage traces with mean values that are correlated to stabilizer eigenvalues and variances that determine the continuous measurement collapse timescales. Monitoring both parity stabilizers in this manner suppresses analog drifts away from the logical subspaces, while providing a steady stream of noisy information to help identify and correct errors that do occur.

Previous work has shown that Bayesian filtering is optimal at extracting parity information from such noisy voltage traces[103, 100]. Here, we implement a simpler technique with performance theoretically comparable to that of the Bayesian filter while using fewer resources on our FPGA controller[100]. We first filter the incoming voltage signals with a 1536 ns exponential filter to reduce the noise inherent from measuring our system with a finite measurement rate (0.40 MHz) and call this signal $V_i(t)$ for resonator i . We offset and normalize $V_i(t)$ such that $\langle V_i(t) \rangle = -1$ corresponds to the system being in an odd parity state, and $\langle V_i(t) \rangle = 1$ corresponds to the system in an even parity state. Here we have defined expectation values as averaging over all possible noise realizations. As shown in Fig. 5.4a, we monitor the trajectories of V_i for signatures of bit-flips using a thresholding scheme[100, 104, 105]. Supposing we prepare an even-even parity state, a bit-flip on one of the outer qubits is detected when one of the signals goes lower than a threshold $\Theta_1 = -0.50$ while the other signal stays above another threshold, $\Theta_2 = 0.72$. A flip of the central qubit is detected when both signal traces fall below a threshold $\Theta_3 = -0.39$. These thresholds are numerically chosen based on experimental trajectories to maximize detection efficiencies of flips while minimizing dark counts and misclassification errors due to noise. When a thresholding condition is met, the controller sends out a corrective π -pulse to the qubit on which the error was detected, with threshold to qubit decoding shown in Fig. 5.4a. The controller also performs a reset operation on the voltage signals in memory to reflect the updated qubit state. As shown in Fig. 5.4b, when a deterministic flip is applied to the $|000\rangle$ state, the system is reset back to $|000\rangle$ faster with feedback than through natural T_1 decay.

To characterize the code, we first check the ability of the controller to correct single bit-flips. We prepare the qubits in $|000\rangle$ and apply the parity readout tones for 16 μ s. After 4 μ s of readout to let the resonators settle, we apply a π -pulse to one of the qubits, inducing a controlled error. We record if and when the controller detects the error and sends out a correction pulse. Errors are successfully detected on Q_0 with 90% efficiency, Q_1 with 86% efficiency, and Q_2 with 91% efficiency. The primary source of inefficiency is T_1 decay bringing the qubits back to ground before detection can happen. On average, the controller corrects an error 3.1 μ s to 3.4 μ s after the error occurs, with the full time distribution shown in Fig. 5.5a. We also characterize a dark count rate for each flip variety by measuring the rate at which the controller detects a qubit flip after preparing in the ground state (3.4, 1.0, 4.0) ms^{-1} . In comparison, the thermal excitation rates for each qubit are estimated to be (1.8, 1.0, 2.0) ms^{-1} .

We next investigate the dominant source of logical errors while running the code: two bit flips occurring in quick succession. When two different qubits flip close together in time relative to the inverse measurement rate, the controller may incorrectly interpret the signals as an error having occurred on the unflipped qubit. The controller then flips this remaining qubit, resulting in a logical error. For continuous error correction, this effect results in a time after an error occurs we call the dead time, when a following error cannot be reliably corrected. To characterize this behavior, we prepare the system in the ground state and apply two successive bit-flips with varying times between the pulses. We then check if the controller responds with the right sequence of correction pulses. In Fig. 5.5b, we show the

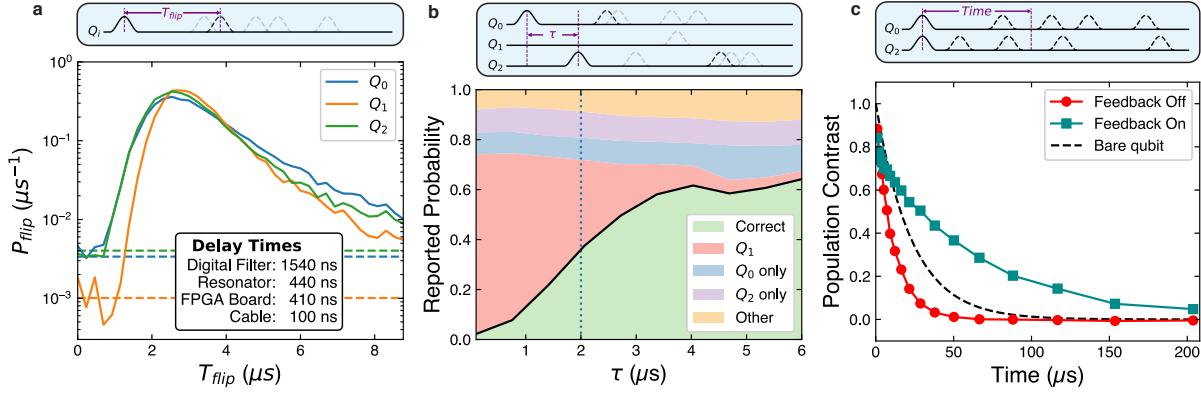


Figure 5.5: Characterizing the time to correct an error. **a**, Histogram of time between an induced error and the correction pulse for each of the qubits, normalized such the integral of the probability distribution $P_{flip}(t)$ gives the detection probability. Dashed lines indicate the dark count rates for each error type. **b**, Probability of detecting certain flip sequences given a flip on Q_0 preceding a flip on Q_2 . The green region is the probability of the controller correctly detecting a Q_0 flip and then a Q_2 flip. The red region is the probability of the controller detecting a Q_1 flip, resulting in a logical error. The dotted line indicates the dead time, when these two probabilities are equal. **c**, Population decay of the excited logical state, $|101\rangle$, of the odd-odd subspace with and without feedback. With feedback on, the lifetime of the logical basis state is longer than that of an individual bare qubit.

controller's interpretation of successive flips on Q_0 and Q_2 as a function of time between them. We mark the dead time at the point where the probability of a logical error crosses the probability of successfully correcting the state. Among the possible pairs and orderings of two qubit errors, the dead times vary from 1.6 μs to 2.6 μs .

Although the code is designed to correct bit-flip errors, the code will also protect the logical computational basis states against qubit decay, extending the T_1 lifetimes of the logical system beyond that of the bare qubits. As opposed to a bit-flip, a qubit decaying loses any coherent phase of the logical state, and the system will be corrected to a mixed state with the same probability distribution in the computational basis as the initial state. For example, the state $\frac{1}{\sqrt{2}}(|0_L\rangle + |1_L\rangle)$ undergoing a qubit decay and correction will be restored as the density matrix $\frac{1}{2}(|0_L\rangle\langle 0_L| + |1_L\rangle\langle 1_L|)$. In the long time limit of active feedback, the system will reach a steady state described by a mixed density matrix with the majority of population (87-99.6%) in the selected codespace. The T_1 of a codespace is defined by the exponential time constant at which population of computational basis states in the codespace approach this steady state. The different codespaces of different parities have different T_1 decay times, with the longest decay time of 66 μs associated with the odd-odd subspace, as shown in Fig. 5.5c. The lifetimes of other subspaces are shown in Fig. 5.6. The

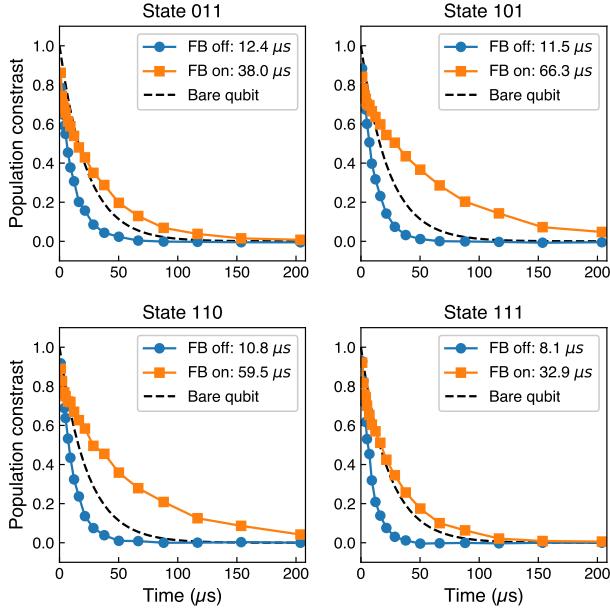


Figure 5.6: Population lifetimes for each of the protected subspaces with and without feedback.

shortest lifetime, 32 μ s, is associated with the even-even subspace, since the higher energy level in this codespace has three bare excitations and the lower energy has no excitations. In comparison, the bare T_1 values of the bare qubits range from 20 μ s to 24 μ s, making the logical qubit excited life 2.7 times longer than that of a bare qubit.

Although phase errors are not protected against by this code, an ideal implementation of a bit-flip code should not increase their occurrence rate. However, with our physical realization of continuous correction, we induce extra dephasing in the logical subspace through three primary channels: continuous dephasing due to the measurement tone; dephasing when going from an odd parity subspace to an even parity subspace; and dephasing related to static ZZ interactions intrinsic to the chip design.

The first source of excess dephasing is measurement-induced dephasing, where the dephasing rate is proportional to the distinguishability of different qubit eigenstates under the measurement[106]. Distinguishability is measured as $D_{m,n}^{(i)} = \left| \alpha_{|m\rangle}^{(i)} - \alpha_{|n\rangle}^{(i)} \right|^2$ where $|m\rangle$ and $|n\rangle$ are different basis states of the two qubits coupled to resonator i , and $\alpha^{(i)}$ is the resonator's associated coherent state[106]. By tuning the qubit frequencies, the dispersive shifts of the system are calibrated such that $D_{01,10}^{(i)}$ are close to zero. The parity measurement distinguishability ($D_{01,11}^{(i)} \approx D_{01,00}^{(i)}$) determines the measurement-induced dephasing rate of the code. Due to finite χ/κ , the even subspaces are not perfectly indistinguishable, with the theoretical distinguishability ratio $D_{00,01}^{(i)}/D_{00,11}^{(i)} \approx 4(\chi_i/\kappa_i)^2$. We use this formula to

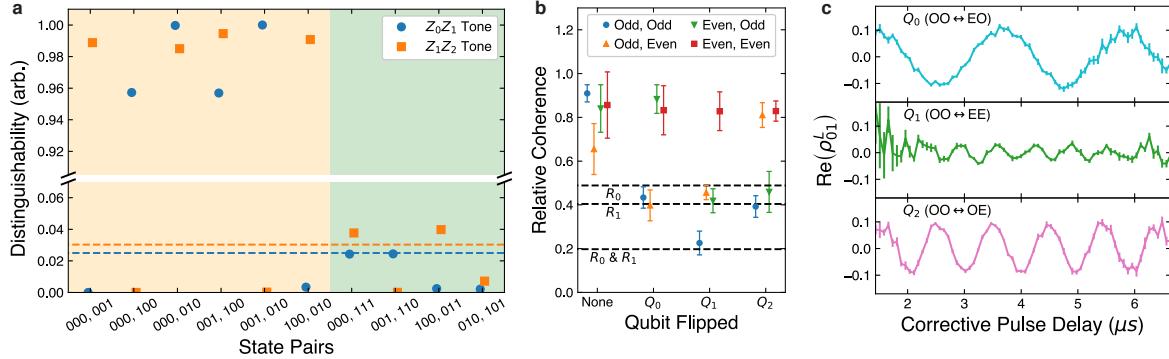


Figure 5.7: Preservation of quantum coherence. (a) Distinguishability of various state pairs in steady state readout for each measurement tone. Pairs of states in the yellow region differ in one or both of their parities. Pairs of states in the green region share their parities. Dashed lines indicate theoretically predicted distinguishability of the even eigenstates. (b) Relative state coherence after preparing a logical $|+X_L\rangle$ state in each of the plotted parity subspaces, applying parity measurement tones without feedback, and flipping one of the qubits. Coherences are normalized to results from the same sequence without the measurement tones applied. Error bars are statistical uncertainty from repeated runs of the measurement. Dashed lines indicate predicted relative dephasing due to an odd to even parity flip on each, both, or no resonators. (c) Sample coherences from preparing a logical $|+X_L\rangle$ state in the odd-odd (OO) subspace, applying an error pulse, and letting the controller correct the error. Coherences are reconstructed by time bins set by the time it takes to correct the error. Oscillations due to static ZZ coupling are visible.

calculate this ratio to be 40 and 33 for resonator 0 and 1 respectively. We plot the measured distinguishability of various state pairs in Fig. 5.7a, and find agreement with these predicted values as well as low distinguishability between eigenstates of odd parity. Using the experimental distinguishability and the known measurement rate, we calculate an maximum induced even-even dephasing rate from a single tone to be

The second source of excess dephasing occurs when a pair of qubits switches from an odd parity state to an even parity state. When two qubits coupled to one of the resonators have odd parity, the resonator is resonantly driven by the measurement tone and thus reaches a steady state with a larger number of photons as compared to when the qubits have even parity. If one of these qubits undergoes a bit-flip while the system is in an odd parity state, the resonator frequency shifts and the system undergoes excess dephasing as the resonator rings down to the steady state for the even subspace. The coherence of the logical state is expected to contract by a factor of $e^{-\bar{n}}$, with \bar{n} being the steady state photon number of a resonator when its qubits are in an odd parity state. From the steady state dephasing rates and the resonator parameters, we independently estimate the photon number in each

resonator to be $\bar{n}_0=0.7$ and $\bar{n}_1=0.6$. To measure this effect, we prepare a 3-qubit logical encoding of an X -eigenstate, $|+X_{L'}\rangle = \frac{1}{\sqrt{2}}(|0_{L'}\rangle + |1_{L'}\rangle)$, where L' is one of the four possible logical encodings. With the measurement tone on, but without feedback, we apply a pulse on one (or none) of the qubits, taking the state to a different (or the same) codespace, L . We then tomographically reconstruct the magnitude of the coherence in the new codespace, $|\rho_{01}^L|$, as shown in Fig. 5.7b. The coherences are normalized to the $|\rho_{01}^L|$ generated by same experiment with the measurement tones off. The system demonstrates significantly less coherence when one of the parities changes from odd to even than when the parities change from even to odd, with reasonable agreement to the expected dephasing.

The third source of excess dephasing is related to static ZZ interactions among the qubits and the uncertainty in timing between when a bit-flip error occurs and when the correction pulse is applied. Performing a Ramsey sequence on Q_i while Q_j is either in the ground or excited state, we measure the coefficients of the system's intrinsic ZZ Hamiltonian, $H_{ZZ} = \frac{1}{2} \sum_{i \neq j} \beta_{ij} Z_i Z_j$. Since the three qubits are in a line topology, with the joint readout resonators also acting as couplers, there is significant coupling between Q_0 and Q_1 ($\beta_{01} = 0.49$ MHz) and between Q_1 and Q_2 ($\beta_{12} = 1.05$ MHz) while there is almost no coupling between Q_0 and Q_2 ($\beta_{02} < 2$ kHz). Due to this coupling, the definite parity subspaces have different energy splittings: In the rotating frame of the qubits, the odd-odd, odd-even, even-odd, and even-even subspaces have logical energy splittings of 0, β_{12} , β_{01} , and $\beta_{01} + \beta_{12}$ respectively. When a bit-flip occurs, the system jumps to an error space and precesses at the frequency of that error space until being corrected by the controller. Since the time from the error flip to the correction pulse is generally unknown, the state can be considered to have picked up a random unknown relative phase. The net dephasing ζ_{zz} can be calculated by averaging the potential phases over the probability distribution of time, T , it takes to correct an error: $e^{i\phi - \zeta_{zz}} = \langle e^{iT\Delta\beta} \rangle_T$ with $\Delta\beta$ being the energy difference between codespace and error space. Using the distributions in Fig. 5.5a and known $\Delta\beta$, we compute ζ_{zz} to be from 2.5 to 5.7 depending on the codespace and the qubit flipped. Although we don't observe this dephasing directly, we perform an experiment to capture this effect. For each of the codespaces, we prepare a $|+X_L\rangle$ state in the odd-odd codespace and induce a bit-flip error while the feedback controller is active. After 6 μ s, we perform tomography on all three qubits and note the time at which the correction pulse occurred. We then reconstruct the logical coherence element ρ_{01}^L of the density matrix conditional on time it took the controller to apply the correction pulse. As shown in Fig. 5.7c, we observe oscillations with frequency corresponding to the effects of ZZ coupling.

Design and fabrication

The microwave properties of the chip were simulated in Ansys high-frequency electromagnetic-field simulator (HFSS), and dispersive couplings were simulated using the energy participation method with the python package pyEPR[32]. Resonators, transmission lines, and qubit capacitors were defined by reactive ion etching of 200 nm of sputtered niobium on a silicon wafer. Al-AlOx-Al Josephson junctions were added using the bridge-free "Manhattan

style" method[107]. The junctions were then galvanically connected to the capacitor paddles through a bandaid process[108]. The middle qubit is fixed frequency, and the outer two qubits are tunable with a tuning range of 260 MHz and 220 MHz. Wire bonds join ground planes across the resonators and bus lines. We summarize chip frequencies and typical qubit lifetimes in the tables below.

	Q_0	Q_1	Q_2
Frequency (MHz)	5355	5182	5392
Anharmonicity (MHz)	307	310	310
T_1 (μ s)	22	23	23
T_2^* (μ s)	18	26	20
T_2^{echo} (μ s)	31	31	35

Table 5.1: Qubit parameters

	R_0	R_1
Frequency (MHz)	6314	6405
κ (kHz)	636	810
χ (MHz)	2.02	2.34
Quantum Efficiency	0.62	0.56

Table 5.2: Resonator parameters

Measurement setup

A wiring diagram of our experimental setup is show in Fig. 5.8. The Josephson Parametric Amplifiers (JPAs) are fabricated with a single step using Dolan bridge Josephson junctions. They are flux pumped at twice their resonance frequency, providing narrow-band, phase-sensitive amplification. The signals are further amplified by two cryogenic high-electron-mobility transistor (HEMT) amplifiers, model LNF4_8. In the output chain for resonator 0, we include a TWPA between the JPA and the HEMT to operate that JPA at a lower gain. Infrared filters on input lines are made with an Eccosorb dielectric. The outer qubits are flux tuned with off-chip coils. The FPGA board provides full control of the qubits and readout of the resonators. The arbitrary waveform generator controls the cavity and JPA drives, as well as triggering the FGPA. The JPA modulation tone is split with one branch phase shifted before both go into an IQ mixer to modulate a reference RF tone on a single sideband.

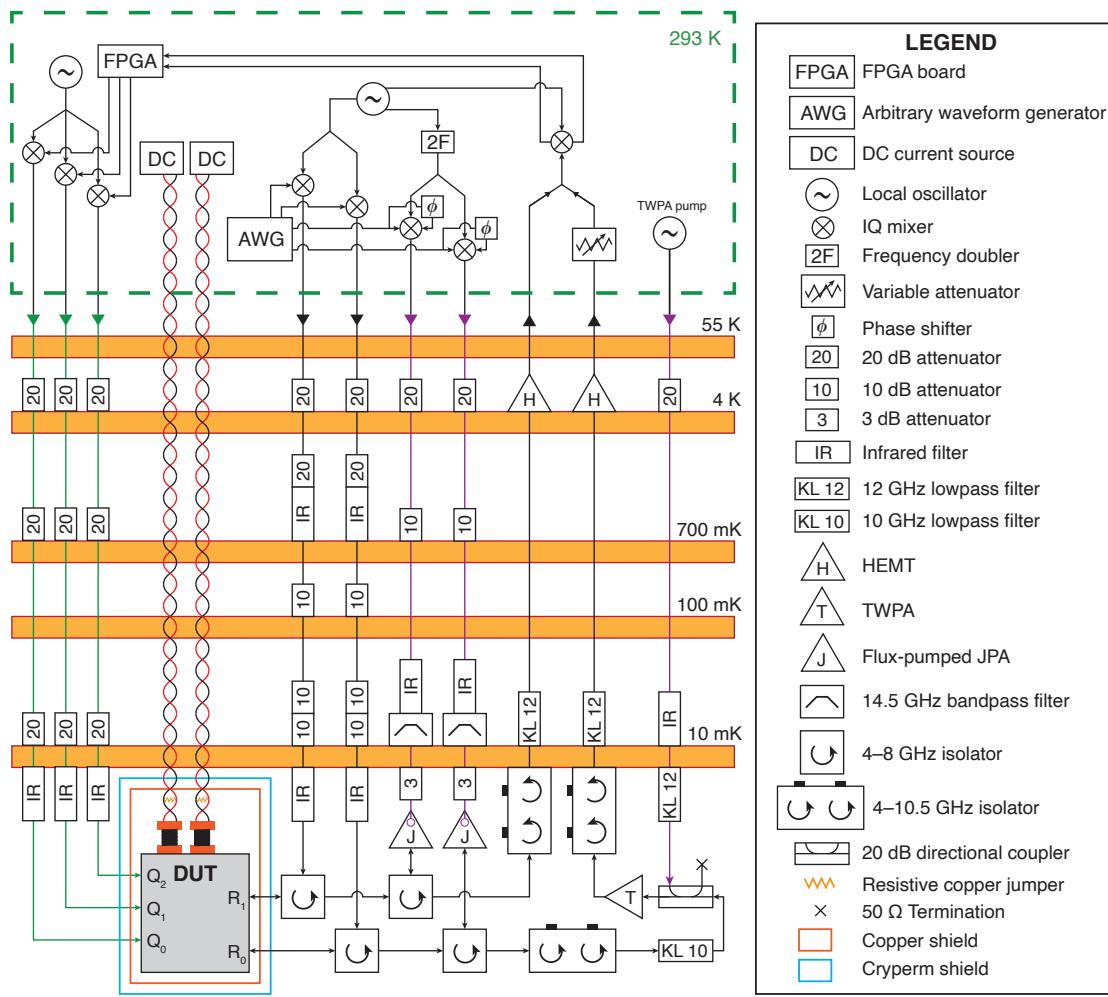


Figure 5.8: Cryogenic wiring diagram. The Josephson parametric amplifiers (JPAs) operate in reflection, and additionally have off chip coils not shown. The JPAs also provide narrow-band gain, so when the readout chains are combined at room temperature, the combined noise at each cavity frequency is dominated by the noise amplified by that cavity's JPA. Each superconducting coil has its leads connected by a small piece of copper wire, forming a low frequency (< 1 Hz) RL filter with the coil. The room temperature wiring is also shown, but with linear elements (attenuators, amplifiers, filters, isolators) removed.

FPGA logic

The FPGA board we used for the feedback is an Innovative Integration X6-1000M board. For technical details on the FPGA code, see Appendix A. We programmed a custom pulse generation core to drive qubit pulses and to demodulate and filter incoming readout signals. A control unit parses instructions loaded in an instruction register. These instructions may include 1) putting a specified number of pulse commands into a queue to await pulse timing; 2) resetting a pulse timer keeping track of time within a sequence while incrementing a trigger counter; and 3) resetting the pulse timer, the trigger counter, and the instruction pointer. When a pulse instruction enters the timing queue, it waits until a specified time and is then sent to one of three different possible locations. The first possible location is a pulse library where the instruction points to a complex pulse envelope of a given duration, which is then modulated by one of three CORDIC sine/cosine generators and sent to the correct digital to analog converter (DAC). These pulses are sent down one of three qubit control lines. The second possible location is to one of the CORDIC sine/cosine generators, where the instruction will increment the phase of the generator by a specified argument, thus implementing Z rotations in the qubit frame. The third location is a demodulation core, which, similarly to the qubit pulse block, plays a complex waveform from memory for a specified duration. This waveform is then multiplied against the complex incoming readout signals and low-pass filtered with a 32 ns exponential filter to generate the signal V_i^{DC} for feedback as well as to readout projective measurements.

When the feedback control unit is active, it takes V_i^{DC} , applies a secondary 1536 ns ns exponential filter/accumulator to further reduce the noise, and then continuously checks these traces (V_i) against the threshold conditions for an error to have been detected (as shown in Fig. 5.4a). When an error is detected, the controller injects instructions for a corrective pi pulse into the pulse generation unit. Any voltage V_i which went across a threshold is then inverted as to not trip further corrective pulses. After a delay such that the corrective pulse has taken effect on chip, the V_i^{DC} is inverted before being accumulated into the running average V_i . In conjunction with the previous inversion of V_i , this effectively resets the feedback controller while avoiding interpreting the corrective pulse as another error.

The board's I/O comprises the PCIe slot for exchanging data with the computer and the ADC/DACs on the analog front-end. The FPGA can stream from multiple sources to the computer along 4 data pipelines. The primary sources are V_i^{DC} and a list of timestamped pulse commands. The timing of any corrective pulses can be obtained from this second source. Further data sources include raw ADC voltages, raw DAC voltages, and V_i , which are only used as diagnostics. On the analog front-end, there are two ADCs running at 1 GSa/s which take in the IF readout signals from the I and Q ports of an IQ mixer, treating the two ADC inputs as the real and imaginary parts of a complex signal. To drive the three qubit lines, there is one DAC running at 1 GSa/s and, due to board constraints, two DACs running at 500 MSa/s.

Optimizing filter parameters

To optimize the threshold values, we prepare the ground state and then flip either one or none of the qubits while taking parity traces (V_i^{DC}). In post processing, we filter the traces with the same exponential filter as on the FPGA to recreate V_i , and classify the resultant traces according to whether or not they pass the different thresholds registering as a qubit flip. We use these data to construct a confusion matrix $P_{ij} = P(i|j)$, the probability of classifying a trace as a flip on i given a preparation flip j , where $i, j \in (\text{None}, 0, 1, 2)$. The thresholds were chosen to minimize $\sum_{ij} (P_{ij} - \delta_{ij})^2$.

Tomographic reconstruction

We use the parity resonators to perform qubit tomography and measure the various lifetimes of the qubits. However, due to the nature of the parity condition, not all states are distinguishable by this measurement. To perform tomography, we use single qubit pulses to map each three-qubit Pauli eigenstate to $|000\rangle$ and then measure both resonators on their $|00\rangle$ resonance [109, 110]. We then measure the probability that full qubit system is in the ground state, which corresponds to reading out both resonators as 0. We additionally include data into the tomography analysis if one of the resonators reads out 1 and the other reads out 0, since we know the final state to be in either $|100\rangle$ or $|001\rangle$ depending on which resonator reads 1. Using this information, we construct partial Pauli expectation values such as $\langle X^+Y^-I \rangle$, with P^+, P^- being the plus and minus projectors for a particular Pauli P such that $P = P^+ - P^-$. We then apply readout correction on these probabilities to mitigate the effects of readout infidelity. From this corrected data taken over many tomographic sequences, we can reconstruct full Pauli expectation values such as $\langle XYI \rangle$. When reconstructing logical coherences, we only measure in the X and Y bases. When reconstructing populations, we only measure in the Z basis.

Ramsey heralding

Qubits 0 and 2 demonstrate a strong temporal bistability in qubit frequency, with a splitting of about 80 kHz and a typical switching time on the order of .1–10 s. When taking data to reconstruct logical coherences, we include five extra sequences in our AWG sequence table, each consisting of five repeated restless Ramsey measurements with free precession times of 6 μ s. With a typical initial sequence length of 64 and a repetition rate of 100 μ s, the qubit’s frequency state is sampled every 7 ms, allowing us to herald data runs to only include data from runs when the qubits have a particular frequency.

Steady state dephasing

Here we derive relative dephasing rates for two qubits in a dispersive parity measurement using a classical analysis of the resonator steady states. The measurement dephasing rate

is proportional to the distinguishability of resonator responses when the coupled qubits are in different eigenstates[52]. We set the probe frequency on resonance with the cavity when qubits are in the single-excitation subspace and assume that $\chi \gg \kappa$. We also assume the external cavity coupling is much larger than the internal cavity loss, so the cavity responds with the following scattering parameter:

$$S(f_0 := \chi \langle Z_0 + Z_1 \rangle) = \frac{-2f_0 + i\kappa}{-2f_0 - i\kappa} \quad (5.13)$$

Odd parity states are perfectly indistinguishable. The distinguishability between states of opposite parity is

$$D_{01,00} = |S(0) - S(2\chi)|^2 = \left| (-1) - \frac{-4\chi + i\kappa}{-4\chi - i\kappa} \right|^2 \approx 4 \quad (5.14)$$

With $z \equiv 4\chi + i\kappa$, the distinguishability between the two even parity states is:

$$\begin{aligned} D_{11,00} &= |S(-2\chi) - S(2\chi)|^2 = \left| \frac{4\chi + i\kappa}{4\chi - i\kappa} - \frac{-4\chi + i\kappa}{-4\chi - i\kappa} \right|^2 \\ &= \left| \frac{z}{\bar{z}} - \frac{\bar{z}}{z} \right|^2 = \left| \frac{z^2 - \bar{z}^2}{|z|^2} \right|^2 = \left| \frac{(z + \bar{z})(z - \bar{z})}{|z|^2} \right|^2 \\ &= \left(\frac{(8\chi)(2\kappa)}{16\chi^2 + \kappa^2} \right)^2 \approx \left(\frac{\kappa}{\chi} \right)^2 \end{aligned} \quad (5.15)$$

From these equations, we get the following relative dephasing (Γ) and measurement rates (Γ^m) between states of different parity and states of even parity:

$$\frac{\Gamma_{01,00}}{\Gamma_{11,00}} = \frac{\Gamma_{01,00}^m}{\Gamma_{11,00}^m} \approx \frac{4\chi^2}{\kappa^2} \quad (5.16)$$

Dynamic dephasing

When the resonator is not at steady state, one can have significantly increased dephasing rates after a parity flip. Here we will consider the effect of a bit flip error taking an odd parity qubit state to an even parity state while the parity measurement is on. In this case, the measurement tone is on resonance with the cavity and the cavity field will initially be in a steady state α_0 . When the qubit parity is flipped from odd to even, the cavity evolves as two copies, one for each even parity basis state (α_{00} and α_{11}). As a simplifying approximation, we assume the measurement tone is turned off at the moment the parity changes as to capture just the transient dynamics. There are two equivalent methods as shown in chapter 3 to calculate the net dephasing ζ . The first can be obtained by integrating the rate at which information leaves the cavity, $\Gamma_\phi^m = \frac{\kappa}{2} |\alpha_{00} - \alpha_{11}|^2$. The second can be obtained by integrating the rate at which the cavity dephases the qubit, $\Gamma_\phi = 4\chi \text{Im}[\alpha_{00}\alpha_{11}^*]$, with 4χ

being the frequency difference between the $|00\rangle$ resonance and the $|11\rangle$ resonance. Here we use the second method to simplify the calculation. We work in the rotating frame of the odd-parity resonance and define $k \equiv \kappa/2 - 2i\chi$ to get two cavity equations, one associated with each basis state:

$$\begin{aligned}\dot{\alpha}_{00} &= \left(2\chi i - \frac{\kappa}{2}\right) \alpha_{00} \\ \dot{\alpha}_{11} &= \left(-2\chi i - \frac{\kappa}{2}\right) \alpha_{11}\end{aligned}\tag{5.17}$$

$$\begin{aligned}\alpha_{00}(t) &= \alpha_0 e^{-kt} \\ \alpha_{11}(t) &= \alpha_0 e^{-\bar{k}t}\end{aligned}\tag{5.18}$$

$$\begin{aligned}\zeta &= \int_0^\infty 4\chi \operatorname{Im} [\alpha_{00}\alpha_{11}^*] dt = 4\chi \operatorname{Im} \left[\int_0^\infty \alpha_{00}\alpha_{11}^* dt \right] \\ &= |\alpha_0|^2 4\chi \operatorname{Im} \left[\int_0^\infty e^{-2kt} dt \right] \\ &= |\alpha_0|^2 4\chi \operatorname{Im} \left[\frac{1}{2k} \right] = |\alpha_0|^2 4\chi \operatorname{Im} \left[\frac{2\bar{k}}{|2k|^2} \right] \\ &= |\alpha_0|^2 \frac{16\chi^2}{\kappa^2 + 16\chi^2} \approx |\alpha_0|^2\end{aligned}\tag{5.19}$$

Therefore, the magnitude of the final coherence between $|00\rangle$ and $|11\rangle$, $|\rho_{00,11}^f|$, will be dephased from the initial coherence between $|01\rangle$ and $|10\rangle$, $|\rho_{01,10}^i|$:

$$|\rho_{00,11}^f| = e^{-\zeta} |\rho_{01,10}^i| = e^{-|\alpha_0|^2} |\rho_{01,10}^i|\tag{5.20}$$

5.3 Outlook and conclusion

Our experiment demonstrates the power of continuous measurements, demonstrating active feedback on multiple multipartite measurement operators. We demonstrate the feasibility of using continuous measurements in quantum error correction. Furthermore, the protocol as implemented is compatible with existing superconducting qubit architectures so can in principle be combined with ancilla-based error correction methods. Future improvements could be made by reducing spurious decoherence effects through novel implementations of continuous parity measurements[111, 112], engineering systems with a lower static ZZ coupling[113] or tunable gate coupling, and using additional feedback to reduce the effects of measurement induced dephasing[114].

Beyond the technological improvements to the existing devices, the ability to continuously measure XZ or XX style stabilizers would allow for the creation of continuous error

correction codes which correct for Z errors as well as X errors. When added to the knowledge that non-commuting continuous observables can be simultaneously measured, one could implement a continuous Bacon-Shor Code [104]. Similar to the work in this paper such a code would only require two-qubit measurement stabilizers.

The canonical phase measurement demonstrates the power of implementing feedback during continuous measurement. At its heart, this experiment required some notion of quantum state tracking. The implemented tracking was relatively easy, as the update equation was linear in the measured voltage, but other systems require more complicated update equations, such as when qubit Rabi drive terms are added in at rates comparable to the readout resonator's linewidth[47]. Linear feedback has been used to stabilize a driven state in a particular parameter regime[115], but FPGAs offer more processing flexibility when the linear regime breaks down. Even in the continuous error correction experiment, there are inherently non-idealities in linear filtering even without considering unwanted effects such as cavity ringing. Neural networks have been shown to be able to track quantum states [116] and are able to capture a wide range of dynamics. Implementing such an architecture on an FPGA could prove useful in more generalized real-time feedback protocols.

Continuous measurements could also be used to probe questions more grounded in physics [48]. With its flexible nature, cQED allows for the simulation of many quantum systems. Collective phenomena such as superradiance, in which atoms get synchronized via a common mode to emit coherently could be observed under continuous measurement [117]. One could observe how trajectories change when these collective systems cross a phase boundary. With the improvement of microwave photon detectors [118], one can additionally consider unravelling these trajectories using photodetectors as well as homodyne and heterodyne detectors. It remains to be seen what dynamics could be generated in these larger system under the influence of feedback.

Appendix A

Making a Continuous Feedback Controller

The electronics controlling quantum circuits need to be fast enough to operate on the timescales of the qubits, parallel enough to route and process multiple streams of data, and deterministic enough to make sure that every pulse sent to or received from the quantum chip happens exactly when expected. Field Programmable Gate Arrays (FPGAs) are the answer to this call. FPGAs are digital integrated circuits allowing configurability of digital logic to a point just shy of having to make an entirely new piece of silicon.

A.1 FPGA basics

FPGAs are enormously complex devices, as can be verified by looking at an FPGA datasheet. We will not get into the details of all the individual features, though you are encouraged to actually look at a datasheet. Instead we hope to convey a framework in which to think about FPGAs. There are two and a half main ingredients in an FPGA: the routing fabric, the processing blocks, and I/O interfaces.¹ The routing fabric does just as its name suggests and routes signals around the chip. The fabric consists of wires periodically going through switchboard-like structures which can be reconfigured after power-up to electrically connect different wires together. When programming (i.e. configuring or flashing) an FPGA, the main role of the bitstream (i.e. configuration file) is to set the correct wiring on these switchboards.

The second set of components is the processing blocks. These are atomic digital circuits built into the silicon of the FPGA. They have input wires where signals come in and output wires where processed signals come out. The primary components are known as logic blocks

¹The I/O blocks count as half mostly because they almost fall under the category of processing blocks. However they are special in that they form the boundary between internal FPGA logic and the external environment of the PCB. They can have advanced features such as integrated (de)serialization, voltage control, and encoding/decoding logic.

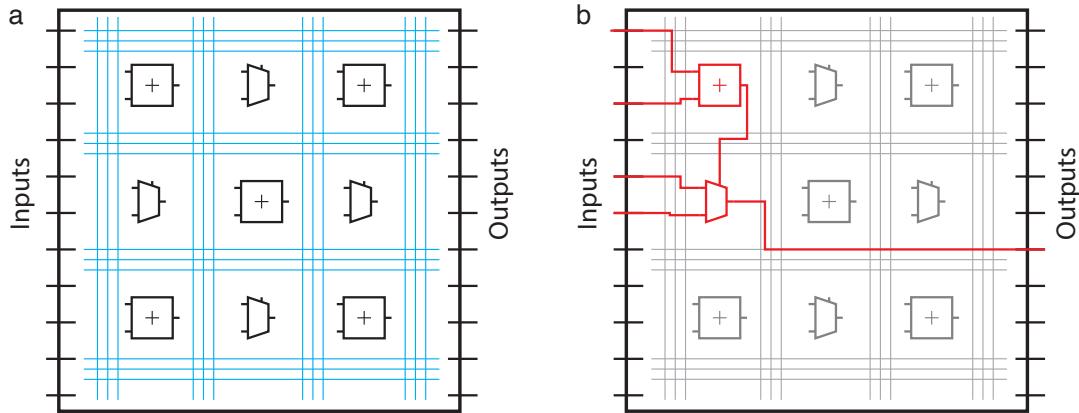


Figure A.1: Inside an FPGA. (a) FPGA wire routing is shown in blue. At the intersections of the wires, there are are switches to route the signals in a configurable manner. Example processing blocks such as adders (square) and multiplexers (trapezoids) are shown in the middle of the routing. (b) Example of a “programmed” FPGA. Two bits are added together and the result is used to select one of two remaining inputs to be an output.

and comprise look-up tables, multiplexers (muxs), and registers. Logic blocks take up most of the physical space on the FPGA, numbering in the tens or hundreds of thousands on a single device. Additionally, FPGAs can have more specialized processing blocks such as dedicated multiplier and RAM modules. These blocks also have run-time static configurations which are set when loading the bitstream onto the FPGA.

A model of an FPGA is shown in Fig. A.1. The programmer’s job is to route incoming signals from the outside world to internal blocks which process the data, and then route the results back out of the FPGA. This is accomplished through a hardware description language (HDL) such as VHDL or verilog, which describes these connections, though at a much higher level than the atomic processing blocks! That being said, the source code of an FPGA mirrors the framework of the FPGA itself in that they define modules² with inputs and outputs. As visualized in Fig. A.2, there is a top level module for the entire device which contains submodules which can in turn contain more submodules, and so on. For example, if someone else has programmed a module to increment every clock cycle, then it can be incorporated into a module using the following VHDL code:

```
inst_counter : Counter
port map(
    Clock      => sys_clk,
```

²Often these modules (aka blocks or cores) are called IP blocks, which often gets shortened to just IP. IP stands for intellectual property, with IP blocks referring to proprietary blocks.

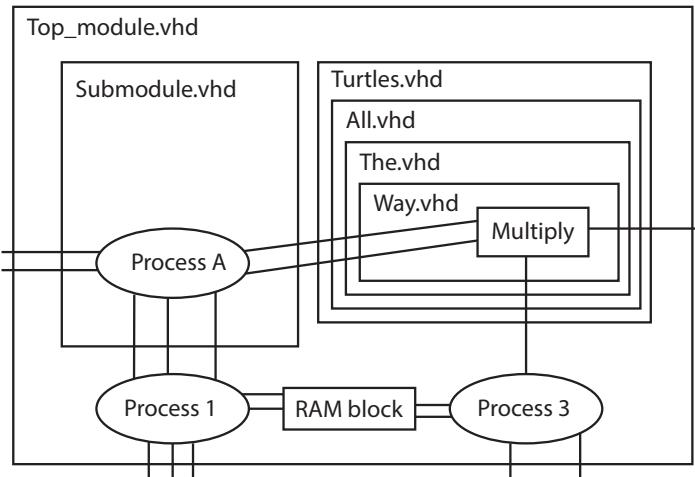


Figure A.2: Sample hierarchy of HDL source files, which eventually get compiled onto the FPGA. Modules have inputs and outputs, and can instantiate other modules as submodules. The wires connecting to the outside get mapped to physical FPGA pins as described in a user constraint file not shown here.

```

active      => enable ,
reset       => rst ,
count_out   => counter
);

```

Here, we have made an instance of a counter and hookup its input wire to the system clock (sys_clk), and hookup its output wire to a wire called ‘counter’ which we can now use for something else. We also add two extra inputs, an enabling wire “enable” and reset wire “rst”. Modules also contain elements called processes. These contain the information to update the values of wires using conditional logic, allowing for step by step computation. For example, if we could not find a module for counting, we could make it ourselves inside a process:

```

process(sys_clk) begin
  if (rising_edge(sys_clk)) then
    if (rst = '1') then
      counter <= x"00";
    elsif (enable = '1')
      counter <= counter + 1;
    else
      counter <= counter;
  end if;
end process;

```

Processes are all instantiated independently of one another, allowing for mass parallelization. However, there are limits as to how complex a single process can get. All the logic inside a process must get accomplished within a clock cycle. If the logic becomes too complex (say one tries to add sequentially add five or six 16-bit numbers together in series), then the time for this procedure to have a stable output on all wires might be longer than a clock cycle. This leads to unexpected and possibly non-deterministic behavior. Fortunately, there are simulation tools which can tell if your code will break this timing constraint.

FPGAs, as with other digital integrated circuits, run on clocks, which synchronize logic throughout the device. Clocks are dedicated signal lines which periodically toggle between on and off. These signals can be generated with dedicated clock circuitry on the FPGA or, as with this board, can be generated off-chip. Ideally, on the rising edge of a clock, all wires have a stable value. When the clock rises, elements called registers change their output state to match their input state, and then hold that value. The logic is then in an unstable state, with the old outputs not valid for the new inputs, but quickly stabilizes as the new logic propagates. If all the timing constraints are met, then the logic will be stable when the next clock cycle rolls around and the race begins again. One thing to bear in mind is that FPGAs can run multiple clocks, where logic running on a particular clock is said to be in a clock domain. Trying to "cross domains" can lead to unpredictable behavior. For example, trying to pass a multi-bit value over a clock domain can cause some bits to lag or lead by a clock cycle in the new domain, leading to nonsensical results. The proper way to do this transfer is through first in, first out (FIFO) buffers, which keeps data synchronized at the cost of some small time delay.

Writing, simulating, and compiling

The code for this thesis was written partially in VHDL, and partially in a graphical interface using Matlab Simulink. Simulink is a Matlab tool for diagrammatic programming (drawing wires between processing boxes). Although Simulink's main purpose is to simulate models in control systems, this framework overlaps nicely with describing data flow in FPGA logic (also wires and boxes). Using the System Generator for DSP (digital signal processing) plugin from Xilinx, one can create and simulate modules which can then be compiled into HDL files and integrated into the rest of the VHDL codebase.

Using this method of coding has some advantages and disadvantages. The main upside is that it is very easy to visualize at a high level how data is moved around, and is very beginner friendly. It is easy to trace the path of any signal. Testing is quite straightforward, though I never tried automating tests. The downsides of using Simulink are that conditional logic is difficult to directly implement, and, since simulink files are binary, version control is non-transparent. Finite state machines (FSMs), a critical concept in FPGA programming, are a mess if done purely diagrammatically, but straightforward in VHDL. Although there is supposedly a way to include VHDL code blocks into Simulink, I found this to always crash

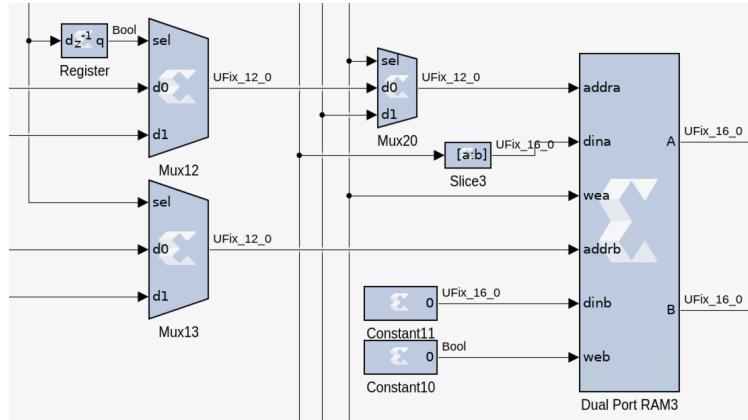


Figure A.3: Using Matlab simulink to write for an FPGA. Wires are made between blocks in a GUI and blocks are can be inspected to tailor their behavior.

the program.³ When writing FSMs in Matlab, I used a variation of Matlab m-code, which is less preferable than VHDL. As a minor point, Simulink models can only have one clock.

In order to synthesize a design (i.e. compile it to an FPGA), one also needs a user constraint file. This maps the input/output wires of the top level HDL file onto the physical pins of the FPGA. This file also contains information on what speed the various clocks are expected to run at, which allows the compilation tools to know when to stop optimizing circuit placement (because timing is met), or throw a warning (because timing is not met). The constraint file can also limit where certain modules are allowed to be physically placed on the FPGA. It is an enormously complex global optimization task to try to shove all the code onto the physical logic cells in a giant uniform mess. By adding area constraints, all logic cells in the constrained modules are physically bundled together within a discrete areas. As a result, the compilation is more likely to succeed.

Simulation and testing are extremely important! Starting with the HDL and simulink files, there is an entire stack of compilers which transform your code into the final bitstream, the configuration file which is actually uploaded to the FPGA at run time.⁴ Compiling everything from source to bitstream for the works in this thesis took on the order of 1-2 hours. Additionally, once the code is on the FPGA, it is inside a black box, with limited opportunity to peek inside. Once you compile, you want to be sure that it will run properly. As mentioned above, Simulink can run a simulation for the code written in Simulink. VHDL code can be simulated in software such as ModelSim, though for the VHDL in this thesis, we primarily used the simulator built into the Xilinx design suite (ISE).

³This error occurred in Matlab R2013b, which is recommended by the manufacturer of the FPGA board.

⁴This upload is done every time the host computer is power-cycled, and doesn't add any time to the booting process. Once the FPGA is programmed, loading the soft configuration of the internal waveform generator and feedback parameters is near instantaneous.

Although ISE has now been replaced by Xilinx's Vivado software, since we are using a Virtex-6 chip, we are tied to using ISE for all of our compilation steps.⁵ This compilation process on a secondary computer with ISE installed. As previously mentioned, there are a list of steps to go through for compilation: synthesis, translation, mapping, place & routing, and bitstream generation. One can also use an internal ISE GUI called PlanAhead to visualize and potentially modify the area constraints in the user constraints file. After placing & routing, the FPGA configuration is completely determined. One can then check if all the timing constraints are met. If they are, then this mid-level code is ready to be further compiled into a bitstream and sent to the FPGA. Otherwise, one can determine "how bad" the timing violation is. If there are a small number of small violations, the map compilation step can be run again with a different seed in order to insert some randomness into the compilation. (The compilation is otherwise deterministic.) If the violations are large, or the reseeding does not help, then either the area constraints need to be modified, or extra delay is needed in critical time paths. It is helpful to consult PlanAhead, where one can check where all the logic elements from a particular module were placed on the chip. If they are far from where it would make sense for them to be or if they are scattered throughout the chip, the constraints file can be modified to mitigate these effects.

Once the bitstream is generated, it needs to be loaded onto the FPGA. With the board used for this thesis, the loading was done over a JTAG connector with the Impact tool inside ISE. Once loaded onto the FPGA, then the only way to check internal logic is if you had the foresight to put virtual probes inside your code. These were done in this project using ChipScope IP cores. When certain conditions are triggered, these cores will record and transmit data back to the computer over JTAG. Although everything should work as in the simulations, this is still an invaluable debugging tool. Although the simulations are accurate, sometimes the stimuli are not as expected, or something gets driven into a strange configuration.

A.2 A board for control

An FPGA alone is not enough to drive transmons and readout their state. The quantum system needs analog inputs and outputs, and the FPGA is entirely digital. The FPGA therefore needs to be connected to analog to digital converters (ADCs) and digital to analog converters (DACs). The FPGA also should be able to stream data back to the computer on the order of GB/s. Given that this is a very challenging task, we opted to buy a commercial board with these components all integrated on a single module. Our only task is then to reprogram the FPGA in order to implement our custom arbitrary waveform generator (AWG) and feedback code. Here we discuss the board we chose to implement our code on.

⁵Open source compilation chains exist, but were not used for this project.

The integrated solution

The board we used for both of the projects in this thesis was the X6-1000M board from Innovative Integration. This board is built around a Xilinx Virtex-6 SX475T FPGA. The FPGA is connected to two ADC chips (TI ADS5400) and two DAC chips (TI DAC5682Z). The ADC chips digitize signals at 1 GS/s with 12-bit resolution. The 16-bit DAC chips each has two outputs, and can be run with one output running at 1 GS/s, or with both outputs independently effectively running at 500 MS/s such that data rate from the FPGA to the DAC chip is always 1 GS/s.⁶

The board also pairs with a bridging board to plug into a high speed PCIe slot on a computer. The board additionally has RAM chips to act as large FIFO buffers to make sure no data is lost as it is transferred to the computer. As configured for the AWG, there are four “fast channels” streaming back to the computer, each using one of these buffering memory banks. These fast lanes use the Velocia/Vita packet protocol provided by Innovative Integration and described in the user manual [119]. In addition, there is a slower channel over the PCIe designed to configure the “Wishbone” registers on the board. These operate on a slower timescale and are not intended for streaming data or configuring large blocks of FPGA RAM. Both the fast and the slow data lanes interface with the host computer using a custom DLL based on Innovative Integration’s Malibou framework to expose these data transfer protocols to python. A schematic of data flow and the various coding pieces are shown Fig. A.4.

Along with the physical board, which comes as a turnkey device, we bought the FPGA logic source code, called the Framework Logic [120]. It contains the basic functionality to talk to all of the other chips on the board as well as to the computer. As a result, we did not for the most part have to deal with finicky hardware bugs, and could focus on implementing our applications. The FPGA code we wrote for this project was built on top of this Framework Logic.

Clock talk

The timing of the DACs and the ADCs is tightly controlled by hardware design of the X6 module. An external phase locked loop provides each of the four A/D chips with a clock line and also feeds a local clock on the FPGA. This means that the front end logic communicating for each A/D chip is on a separate local clock running at 250 MHz. Although these clocks are phase locked, the relative clock phases are skewed. There is another phase-locked 250 MHz input which we use as a global clock to run our sequencing code. Much of the rest of

⁶When running the DAC in 500 MS/s, the chip and outputs are still physically running at 1 GS/s, but the incoming data is from the FPGA is interpolated so effective the bandwidth is 500 MS/s. The interpolation unfortunately also adds an extra 88 ns of delay. One could in theory run the entire DAC chip on a 500 MS/s clock while still supply 1 GS/s of data from the FPGA for a latency of 72 ns, but this proved more challenging to implement. Additionally, if run in interpolation mode, the output can be mode-shaped to put the output power in the 250-500 MHz band.

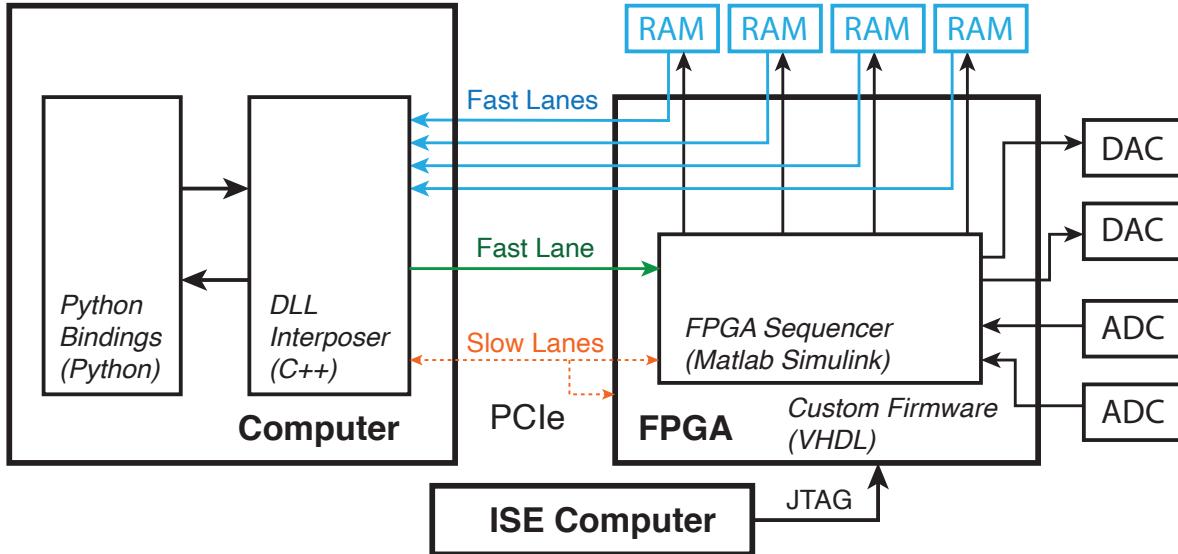


Figure A.4: Data flow through the system. Acquired data is streamed back to the computer through fast lanes using off-chip RAM as buffering memory. Fast lanes are additionally used to transfer sequencing instructions and pulse shapes to the FPGA. Extra configuration data (not to be confused with the bitstream) can be sent to FPGA registers over the slow lanes. These lanes are also used to read back status information. The FPGA Sequencer interacts with the FPGA to send and receive data to and from the DACs and ADCs respectively. Processed data is sent back to the computer. The ISE computer (used for bitstream compilation) programs the FPGA with the bitstream over a JTAG connector.

the FPGA code is run on a global 260 MHz clock, allowing data to be shuttled around and eventually sent back to the computer.

The first task when processing incoming ADC data is to synchronize the data coming from both ADCs. We use FIFOs to bring the samples from the local clock to the global clock and another buffer to ensure they are synchronized with each other. Data sent from the sequencer to the DACs are similarly sent into FIFOs to cross into the DAC local clocks. Once the samples are on this clock, the Framework Logic ensures that they are synchronized to each other and to an external trigger provided at the start of every sequence. Additional latency is artificially added if the DAC is run in single output mode at 1 GSa/s⁷ so that it is synchronized when the other DAC is run in 500 MSa/s mode.

⁷Gigasample per second

A.3 Writing an AWG for feedback

The first goal of the FPGA sequencer is to send qubit control pulses out via the DACs, after which the pulses are unconverted and sent to the qubit chip. The waveforms coming out of the DAC will be an envelope function (such as a Gaussian) which has modulated a digital local oscillator inside a “pulse block.” This digital local oscillator frequencies are controlled by “phase accumulators.” The waveform envelopes are stored in a bank of internal RAM, and a “control unit” dictates when a particular waveform is to be played. The second goal of the sequencer is to readout the state of the qubits. An external AWG sources the readout pulses which are modulated, reflected off the quantum chip, and demodulated to intermediate frequency (IF). This IF signal is read in by the ADCs and further demodulated to DC on the FPGA board. The FPGA demodulation is performed similarly to the envelope modulation: a bank of internal RAM holds pulses which are multiplied against the incoming signal inside the “readout block” when the control unit dictates. The ultimate goal is then to link input and output, using a feedback module to analyze the incoming data and determining when a correction pulse is needed.

Control unit

The control unit for this project is relatively simple. An instruction memory holds commands for the control unit’s finite state machine (the CUFSM) as well as timed commands for the pulse blocks, control blocks, and phase accumulators. The instruction memory is 46 bits wide and 16,384 deep. A 12-bit trigger counter counts the number of external triggers received where each trigger represents a new set of pulses to be run. A 16-bit timer increments every clock cycle to keep track of time. The control unit enqueues timed commands into one of ten FIFOs (for ten destinations) which wait until the timer matches the release time and the trigger counter matches the pulse’s trigger number.

The format of CUFSM instructions is

Unused	<i>Arguement</i>	Command
29 bits	14 bits	3 bits

The command bits contain a number from 0 to 5 with definitions are as follows:

- 0: Enqueue the next *Arguement* commands in the instruction memory as timed commands.
- 1: Unused. (historical)
- 2: Jump to the address in instruction memory as specified in *Arguement*.⁸
- 3: Return to the address instruction memory just after the last jump command was called.
- 4: Reset the timer, and wait for the next trigger
- 5: Reset the timer, reset the trigger, and send the instruction pointer back to zero.

⁸The jumping commands (3 and 4) were to be used as a compression technique, but never put into practice since we almost never ran up against the depth of the instruction memory.

The format of a timed command is

Block Command	Time (in ns)	Destination
24 bits	18 bits	4 bits

Destination is one of 10 locations (4 DACs [via 2 pulse blocks], 3 phase accumulators, 2 readout blocks, and 1 feedback signal). *Time* is the number of ns since the time of the last timed command for the next command to launch. Since the FPGA is on a 4 ns clock, the last two bits of *Time* cannot be used to compare with the timer. Instead, these two bits are tacked onto the end of the block command (making it 26 bits) before it is enqueued.⁹

Pulse block

There are two pulse blocks implemented, each associated with a DAC chip. The pulse blocks hold the envelope waveforms in RAM. To supply a sample on every nanosecond, 4 samples per clock cycle need to be looked up. Additionally, the envelopes are complex with 16-bit resolution for both the real and imaginary components. In total, there are 8 numbers to look up every clock cycle,¹⁰ so 8 banks of RAM are used, each with depth 4,096. Therefore, the sum time of all the stored pulse is 16,384 ns, which was virtually never a limiting factor.

The pulse block can be used in either single output mode for a sampling rate of 1 GSa/s, or in dual output mode for two outputs each with 500 MSa/s sampling rate. When sending loading samples to the FPGA, the sample data is sent using the following format, with subscripts referring to nanoseconds:

Single output	Im ₃	Im ₂	Im ₁	Im ₀	Re ₃	Re ₂	Re ₁	Re ₀
Dual output	Im ₂ ^{dac1}	Im ₂ ^{dac0}	Im ₀ ^{dac1}	Im ₀ ^{dac0}	Re ₂ ^{dac1}	Re ₂ ^{dac0}	Re ₀ ^{dac1}	Re ₀ ^{dac0}

We implemented a barrel shifter to dictate where the first of the 4 samples should start to ns accuracy, or 2-ns accuracy in dual output mode.

Each pulse pulse block also contains 4 DDS generators, which generate the sine and cosine components of the complex modulation from one of the phase accumulators. The complex envelope is multiplied against the complex modulation, after which the real part is taken and sent off to the DAC.

The block command for the pulse block (coming from the control unit) has the following structure:

Duration (ns)	Instruction Address	LO Choice
10 bits	12 bits	2 bits

⁹Only the pulse blocks and phase accumulator can use this extra 2 bits for ns resolution. If a pulse block is run in dual mode (2 outputs at 500 MHz), then only the first of the two bits is used. If these two bits are non-zero in a command to another block, such as a readout block, the behavior will not be as desired, and will cause unexpected rotations in readout. This is only protected against in software.

¹⁰If the end of one pulse happens “off a clock cycle” and another one starts immediately, two numbers need to be looked up from the same RAM, but the RAM used is dual read, so that is ok.

LO choice selects one of the three phase accumulators, thus choosing the frequency to modulate against. The instruction address gives the starting location of the pulse envelope. The duration is the duration of the pulse in ns.

Phase accumulator

There are three phase accumulators implemented, each associated with a qubit and its modulation frequency. The frequency is set before the experiment by changing a register dictating how much phase to accumulate each clock cycle. The block outputs 4 phases per clock cycle, each corresponding to a particular nanosecond offset. The control unit can send block commands to the phase accumulators to implement virtual Z gates [121]. These commands take the following form:

Unused	Phase shift
8 bits	16 bits

where the phase shift has the range $[0, 1)$ representing 0 to 2π , with resolution of 2^{-16} . The phase accumulators themselves use 30 bits of resolution for a future-proof 230 mHz of resolution.

Readout block

There are two readout blocks, each associated with a readout frequency multiplexed on the same RF line. When the reflected readout tone comes out of the cryostat at 6 GHz or so, it is demodulated using an IQ mixer, with the I quadrature feeding ADC 0 and the Q quadrature feeding ADC 1. By treating the two inputs as the real and imaginary parts of a complex signal, the IF tone is selectively at a positive or negative frequency component in frequency space. Similarly to the pulse block, the readout block has 8 banks of RAM containing the real and imaginary parts of the complex conjugate of the desired signal to extract within a 4 ns window. Again, each RAM is 16 bits wide and 4096 samples (16,384 ns) deep. Unlike the pulse blocks, there is no extra modulation step: the signal stored in RAM is already modulated and ready to multiply against the incoming data.¹¹ After the multiplication (i.e. demodulation), the signal is filtered with a sliding exponential filter as an anti-aliasing device. This filter can be configured to have a time constant of 0, 32, 64, or 128 ns. This filtered data is then sent out from the block for feedback analysis or streaming back to the computer.

The block command for the readout block (coming from the control unit) has the following structure:

¹¹Since the stored waveform is already modulated, we need to play tricks to readout longer than 16 μ s. We can do this by storing a chunk of readout waveform twice as large as an “atomic readout” and then replaying parts of this waveform having precomputed a starting phase for each section of readout. This requires the frequency to be higher than an inverse atomic integration time, and incommensurate with the 250 MHz FPGA clock.

Unused	Duration (clks)	Instruction Address
2 bits	10 bits	12 bits

The instruction address gives the starting address in memory for the pulse. The duration is given in clock cycles, since the readout block does not have barrel shifters implemented for ns resolution.

Feedback block

The feedback block is specialized for the continuous error correction protocol. It takes in data from the readout block and applies an exponential filter of a much longer time constant (in increments of 256 ns) to cut down on stochastic noise. This filter is reinitialized to a programmable preset value when the feedback ends. The pigpen thresholding scheme described in Chapter 5 is then applied. When the filtered trace enters one of the “error pens”, the thresholds are reflected around the correct axes in voltage space such that the voltages are no longer in an error pen. After a “loop delay,” the time it takes for the feedback pulse to have taken effect, the voltage would start trending back to the original voltage quadrant, which is now an error pen. To prevent this, after a loop delay, the correct voltage trace(s) is inverted before going into the stochastic filter. This means that even though the true voltage flips after a correction pulse, the effective voltage as seen by the thresholding scheme is unchanged in sign.

When the feedback detects an error, it sends out a signal to the control unit encoding which qubit needs correction. There is a special block inside the control unit which injects predefined block commands immediately to every pulse block. Only the relevant pulse blocks get actionable commands,¹² at which point they send out the corrective pi pulse.

The feedback block gets block commands from the control unit in the form:

Unused	Duration (clks)	Initial R ₀ Parity	Initial R ₁ Parity
6 bits	16 bits	1 bit	1 bit

The initial R₀ and R₁ parities tell the feedback block whether the parities are initially even (0) or odd (1). This value tells the stochastic filter the correct initial value as well as how to set up the pigpen thresholds correctly. The duration represents the duration of the feedback in clock cycles.

Data streaming

The RAM blocks of the control unit, pulse blocks, and readout blocks are loadable through the use of Velocia/Vita packets, the fastest and most efficient way to transfer data to the FPGA. Other small configuration numbers for the FPGA such as qubit modulation frequency and feedback thresholds are configured through the Wishbone registers.

¹²Due to crosstalk cancellation, the pi pulses on Q₀ and Q₂ have weight on both drive lines 0 and 2.

There are up to four output data channels to stream back to the computer simultaneously over the Velocia/Vita channels. The channels have the following possible streaming sources:

Source \ Channel	0	1	2	3
Raw ADC 0	X			
Raw ADC 1		X		
Raw DAC 0			X	
Raw DAC 1				X
ADC demod/filtered 0	X		X	
ADC demod/filtered 1		X		X
Command unit data	X			X
Feedback filtered data		X	X	

Raw ADC data allows streaming from a particular ADC on the board. Raw DAC is the raw signal sent to the DAC chip. Both of these raw values are used for diagnostic purposes and not on a regular basis. The ADC demodulated data is the output of the readout block (V_i^{DC} in chapter 5), which is further binned and averaged. Typically, this averaging bin size was chosen to be 32 ns, allowing for efficient streaming while still being able to resolve continuous dynamics. Choosing command unit data streams back time-stamped versions of every block command coming from the command unit. This is useful for diagnostics but primarily used to get timing data as to when feedback pulses occurred. The feedback filtered data is the heavily filtered signal trace used for feedback (V_i in chapter 5) and is used as a diagnostic to verify the accuracy of the filter.

A.4 Computer interface

Data streaming from the FPGA can be very demanding on bandwidth. The FPGA board uses direct memory access (DMA) to transfer data directly to a dedicated block of computer RAM, which acts as a computer-side memory buffer for streaming data. This process is controlled by a dynamic-link library (DLL) written in C++ based on the Malibou framework provided by Innovative Integration. This code includes board initialization procedures and access to Wishbone registers as well as the ability to stream data.

The main purpose of the DLL is to expose board functionality as “quickly” as possible to python and minimize the amount of code written in C++. Basic functionality such as opening the board and accessing the registers were wrapped and exposed in the DLL. Large data transfers, such as sending data to configure the AWG RAM, are accomplished by creating the data in numpy arrays, which can then be accessed by the DLL layer and transferred to the board. When receiving data, the python layer initializes numpy arrays, and as data comes into the computer, the DLL layer copies the data from the computer’s DMA RAM buffer to these arrays.

Python interface

The lowest level python module is the `II_Board_Interface` of `ii_fpga.py`. This provides the wrapper class for the functions in the DLL. For the parity project, this is sub-classed by the `Parity_Board` in `ii_parity_board.py`, which has code specific to the implementation of the AWG sequencer on the FPGA. The primary functions are

- `open`: creates a connection to the board.
- `start_streaming`: starts the ADCs and DACs (though does not output any waveforms or acquire data). Usually only needs to be run after opening the board.
- `load_config`: loads the a configuration dictionary with all data to run a sequence.
- `start_acquire`: starts the FPGA listening for triggers.
- `acquire`: blocks python until all the expected data is acquired or a timeout is hit. Returns the acquired data.
- `close`: closes the board in a clean manner. Nothing bad happens if the python session ends without closing the board.¹³

In practice, the `Parity_Board` object is instantiated as part of the initialization of a `Parity_Experiment` object, which also has control over the other instrumentation including the external AWG which controls the FPGA trigger. The `Parity_Experiment.acquire_ii` is then typically used to acquire data.

The configuration dictionary for the `Parity_Board` is typically generated by compiling a sequence. First an “AMiable Sequence” is generated using the `Sequence` object in `amiable_sequencer.py`. This sequence is then compiled through `seq_utils.standard_compile`, which uses functions in `ii_pulse_parser` to generate a compiled sequence for the external AWG as well as a configuration dictionary for the FPGA and `RO_slices`, which is related to the FPGA readout. This is all the configuration technically needed to load and run an experiment. However, we have one more convenience layer, the `FPGA_Measurement` object in `fpga_measurement.py`. The output of `standard_compile` can be used to initialize an `FPGA_Measurement` object which can be made to hold metadata for the sequence as well as an `RO_map` object, which holds information on how to process data coming from the FPGA. This `RO_map` maps demodulation elements of the AMiable Sequence to an `fpga_measurement.RO_settings` object, which dictates how the data will be post-processed. This `FPGA_Measurement` is then loaded into the parity experiment using `Parity_Experiment.set_measurement`. Examples of sequences are in `seq_utils.py`.

¹³The MemDrv Controller, an Innovative Integration applet, may need resetting before restarting.

Bibliography

- [1] Bas Hensen et al. “Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres”. In: *Nature* 526.7575 (2015), pp. 682–686.
- [2] H J Carmichael. *An open systems approach to quantum optics*. Springer, 1993.
- [3] Jean Dalibard, Yvan Castin, and Klaus Mølmer. “Wave-function approach to dissipative processes in quantum optics”. In: *Phys. Rev. Lett.* 68 (5 Feb. 1992), pp. 580–583. DOI: 10.1103/PhysRevLett.68.580. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.68.580>.
- [4] C. W. Gardiner, A. S. Parkins, and P. Zoller. “Wave-function quantum stochastic differential equations and quantum-jump simulation methods”. In: *Phys. Rev. A* 46 (7 Oct. 1992), pp. 4363–4381. DOI: 10.1103/PhysRevA.46.4363. URL: <https://link.aps.org/doi/10.1103/PhysRevA.46.4363>.
- [5] K. W. Murch et al. “Observing single quantum trajectories of a superconducting quantum bit”. In: *Nature* 502.7470 (Oct. 2013), pp. 211–214. ISSN: 1476-4687. DOI: 10.1038/nature12539. URL: <https://doi.org/10.1038/nature12539>.
- [6] Philippe Campagne-Ibarcq et al. “Observing quantum state diffusion by heterodyne detection of fluorescence”. In: *Physical Review X* 6.1 (2016), p. 011002.
- [7] Q. Ficheux et al. “Dynamics of a qubit while simultaneously monitoring its relaxation and dephasing”. In: *Nature Communications* 9.1 (May 2018), p. 1926. ISSN: 2041-1723. DOI: 10.1038/s41467-018-04372-9. URL: <https://doi.org/10.1038/s41467-018-04372-9>.
- [8] Shay Hacohen-Gourgy et al. “Quantum dynamics of simultaneously measured non-commuting observables”. In: *Nature* 538.7626 (2016), p. 491.
- [9] D Ristè et al. “Deterministic entanglement of superconducting qubits by parity measurement and feedback”. In: *Nature* 502 (2013), p. 350.
- [10] R. Vijay et al. “Stabilizing Rabi oscillations in a superconducting qubit using quantum feedback”. In: *Nature* 490.7418 (Oct. 2012), pp. 77–80. ISSN: 1476-4687. DOI: 10.1038/nature11505. URL: <https://doi.org/10.1038/nature11505>.
- [11] Peter W. Shor. “Scheme for reducing decoherence in quantum computer memory”. In: *Phys. Rev. A* 52 (4 Oct. 1995), R2493–R2496. DOI: 10.1103/PhysRevA.52.R2493. URL: <https://link.aps.org/doi/10.1103/PhysRevA.52.R2493>.

- [12] David P. DiVincenzo and Peter W. Shor. “Fault-Tolerant Error Correction with Efficient Quantum Codes”. In: *Phys. Rev. Lett.* 77 (15 Oct. 1996), pp. 3260–3263. DOI: 10.1103/PhysRevLett.77.3260. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.77.3260>.
- [13] Xin Zhang et al. “Semiconductor quantum computation”. In: *National Science Review* 6.1 (Dec. 2018), pp. 32–54. ISSN: 2095-5138. DOI: 10.1093/nsr/nwy153. eprint: <https://academic.oup.com/nsr/article-pdf/6/1/32/30336105/nwy153.pdf>. URL: <https://doi.org/10.1093/nsr/nwy153>.
- [14] Marcus W. Doherty et al. “The nitrogen-vacancy colour centre in diamond”. In: *Physics Reports* 528.1 (2013). The nitrogen-vacancy colour centre in diamond, pp. 1–45. ISSN: 0370-1573. DOI: <https://doi.org/10.1016/j.physrep.2013.02.001>. URL: <https://www.sciencedirect.com/science/article/pii/S0370157313000562>.
- [15] Colin D. Bruzewicz et al. “Trapped-ion quantum computing: Progress and challenges”. In: *Applied Physics Reviews* 6.2 (2019), p. 021314. DOI: 10.1063/1.5088164. eprint: <https://doi.org/10.1063/1.5088164>. URL: <https://doi.org/10.1063/1.5088164>.
- [16] Mark Saffman. “Quantum computing with neutral atoms”. In: *National Science Review* 6.1 (Sept. 2018), pp. 24–25. ISSN: 2095-5138. DOI: 10.1093/nsr/nwy088. eprint: <https://academic.oup.com/nsr/article-pdf/6/1/24/30336094/nwy088.pdf>. URL: <https://doi.org/10.1093/nsr/nwy088>.
- [17] Jens Koch et al. “Charge-insensitive qubit design derived from the Cooper pair box”. In: *Phys. Rev. A* 76 (4 Oct. 2007), p. 042319. DOI: 10.1103/PhysRevA.76.042319. URL: <https://link.aps.org/doi/10.1103/PhysRevA.76.042319>.
- [18] Fei Yan et al. “The flux qubit revisited to enhance coherence and reproducibility”. In: *Nature Communications* 7.1 (Nov. 2016), p. 12964. ISSN: 2041-1723. DOI: 10.1038/ncomms12964. URL: <https://doi.org/10.1038/ncomms12964>.
- [19] Jeffrey Scott Birenbaum. *The C-shunt flux qubit: a new generation of superconducting flux qubit*. 2014.
- [20] Vladimir E. Manucharyan et al. “Fluxonium: Single Cooper-Pair Circuit Free of Charge Offsets”. In: *Science* 326.5949 (2009), pp. 113–116. ISSN: 0036-8075. DOI: 10.1126/science.1175552. eprint: <https://science.sciencemag.org/content/326/5949/113.full.pdf>. URL: <https://science.sciencemag.org/content/326/5949/113>.
- [21] Long B. Nguyen et al. “High-Coherence Fluxonium Qubit”. In: *Phys. Rev. X* 9 (4 Nov. 2019), p. 041041. DOI: 10.1103/PhysRevX.9.041041. URL: <https://link.aps.org/doi/10.1103/PhysRevX.9.041041>.
- [22] Peter Brooks, Alexei Kitaev, and John Preskill. “Protected gates for superconducting qubits”. In: *Phys. Rev. A* 87 (5 May 2013), p. 052306. DOI: 10.1103/PhysRevA.87.052306. URL: <https://link.aps.org/doi/10.1103/PhysRevA.87.052306>.

- [23] András Gyenis et al. “Experimental Realization of a Protected Superconducting Circuit Derived from the $0-\pi$ Qubit”. In: *PRX Quantum* 2 (1 Mar. 2021), p. 010339. DOI: 10.1103/PRXQuantum.2.010339. URL: <https://link.aps.org/doi/10.1103/PRXQuantum.2.010339>.
- [24] François Mallet et al. “Single-shot qubit readout in circuit quantum electrodynamics”. In: *Nature Physics* 5.11 (Nov. 2009), pp. 791–795. ISSN: 1745-2481. DOI: 10.1038/nphys1400. URL: <https://doi.org/10.1038/nphys1400>.
- [25] A. Grimm et al. “Stabilization and operation of a Kerr-cat qubit”. In: *Nature* 584.7820 (Aug. 2020), pp. 205–209. ISSN: 1476-4687. DOI: 10.1038/s41586-020-2587-z. URL: <https://doi.org/10.1038/s41586-020-2587-z>.
- [26] M. S. Khalil et al. “An analysis method for asymmetric resonator transmission applied to superconducting devices”. In: *Journal of Applied Physics* 111.5 (2012), p. 054510. DOI: 10.1063/1.3692073. eprint: <https://doi.org/10.1063/1.3692073>. URL: <https://doi.org/10.1063/1.3692073>.
- [27] David M Pozar. *Microwave engineering*. John wiley & sons, 2011.
- [28] Uri Vool and Michel Devoret. “Introduction to quantum electromagnetic circuits”. In: *International Journal of Circuit Theory and Applications* 45.7 (June 2017), pp. 897–934. ISSN: 0098-9886. DOI: 10.1002/cta.2359. URL: <http://dx.doi.org/10.1002/cta.2359>.
- [29] Paul Adrien Maurice Dirac and Niels Henrik David Bohr. “The quantum theory of the emission and absorption of radiation”. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 114.767 (1927), pp. 243–265. DOI: 10.1098/rspa.1927.0039. eprint: <https://royalsocietypublishing.org/doi/pdf/10.1098/rspa.1927.0039>. URL: <https://royalsocietypublishing.org/doi/abs/10.1098/rspa.1927.0039>.
- [30] Michael Martin Nieto. “Quantum phase and quantum phase operators: some physics and some history”. In: *Physica Scripta* T48 (Jan. 1993), pp. 5–12. DOI: 10.1088/0031-8949/1993/t48/001. URL: <https://doi.org/10.1088/0031-8949/1993/t48/001>.
- [31] M. Devoret, Andreas Wallraff, and J.M. Martinis. “Superconducting Qubits: A Short Review”. In: (Dec. 2004). URL: <https://arxiv.org/abs/cond-mat/0411174>.
- [32] Zlatko K. Minev et al. *Energy-participation quantization of Josephson circuits*. 2020. arXiv: 2010.00620 [quant-ph].
- [33] Daniel Huber Slichter. *Quantum jumps and measurement backaction in a superconducting qubit*. University of California, Berkeley, 2011.
- [34] Simon E. Nigg et al. “Black-Box Superconducting Circuit Quantization”. In: *Phys. Rev. Lett.* 108 (24 June 2012), p. 240502. DOI: 10.1103/PhysRevLett.108.240502. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.108.240502>.

- [35] Andrew Wilson Eddins. *Superconducting Circuits for Quantum Metrology with Non-classical Light*. University of California, Berkeley, 2017.
- [36] P. Blasiak et al. “Combinatorics and Boson normal ordering: A gentle introduction”. In: *American Journal of Physics* 75.7 (July 2007), pp. 639–646. ISSN: 1943-2909. DOI: 10.1119/1.2723799. URL: <http://dx.doi.org/10.1119/1.2723799>.
- [37] Zlatko K. Minev and Zaki Leghtas. *pyEPR*. URL: <https://github.com/zlatko-minev/pyEPR>.
- [38] Arvind et al. “The real symplectic groups in quantum mechanics and optics”. In: *Pramana* 45.6 (Dec. 1995), pp. 471–497. ISSN: 0973-7111. DOI: 10.1007/BF02848172. URL: <https://doi.org/10.1007/BF02848172>.
- [39] Gerardo Adesso, Sammy Ragy, and Antony R. Lee. “Continuous Variable Quantum Information: Gaussian States and Beyond”. In: *Open Systems & Information Dynamics* 21.01n02 (2014), p. 1440001. DOI: 10.1142/S1230161214400010. eprint: <https://doi.org/10.1142/S1230161214400010>. URL: <https://doi.org/10.1142/S1230161214400010>.
- [40] Michael A Nielsen and Isaac Chuang. *Quantum computation and quantum information*. American Association of Physics Teachers, 2002.
- [41] Todd A. Brun. “A simple model of quantum trajectories”. In: *American Journal of Physics* 70.7 (2002), pp. 719–737. DOI: 10.1119/1.1475328. eprint: <https://doi.org/10.1119/1.1475328>. URL: <https://doi.org/10.1119/1.1475328>.
- [42] Joshua Combes, Joseph Kerckhoff, and Mohan Sarovar. “The SLH framework for modeling quantum input-output networks”. In: *Advances in Physics: X* 2.3 (2017), pp. 784–888. DOI: 10.1080/23746149.2017.1343097. eprint: <https://doi.org/10.1080/23746149.2017.1343097>. URL: <https://doi.org/10.1080/23746149.2017.1343097>.
- [43] Dario Cilluffo et al. “Collisional picture of quantum optics with giant emitters”. In: *Phys. Rev. Research* 2 (4 Oct. 2020), p. 043070. DOI: 10.1103/PhysRevResearch.2.043070. URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.2.043070>.
- [44] M. A. Nielsen and I. L. Chuang. *Quantum computation and quantum information*. Cambridge University Press, 2010.
- [45] MA Castellanos-Beltran et al. “Amplification and squeezing of quantum noise with a tunable Josephson metamaterial”. In: *Nature Physics* 4.12 (2008), p. 929.
- [46] A. Eddins et al. “Stroboscopic Qubit Measurement with Squeezed Illumination”. In: *Phys. Rev. Lett.* 120 (4 Jan. 2018), p. 040505. DOI: 10.1103/PhysRevLett.120.040505. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.040505>.
- [47] H. M. Wiseman and G. J. Milburn. *Quantum measurement and control*. Cambridge University Press, 2009.

- [48] Leigh S. Martin. *Quantum feedback for measurement and control*. 2020. arXiv: 2004.09766 [quant-ph].
- [49] Lawrence C Evans. *An introduction to stochastic differential equations*. Vol. 82. American Mathematical Soc., 2012.
- [50] Bernt Øksendal. “Stochastic differential equations”. In: *Stochastic differential equations*. Springer, 2003, pp. 65–84.
- [51] C. Macklin et al. “A near-quantum-limited Josephson traveling-wave parametric amplifier”. In: *Science* 350.6258 (2015), pp. 307–310. ISSN: 0036-8075. DOI: 10.1126/science.aaa8525. eprint: <https://science.sciencemag.org/content/350/6258/307.full.pdf>. URL: <https://science.sciencemag.org/content/350/6258/307>.
- [52] Jay Gambetta et al. “Quantum trajectory approach to circuit QED: Quantum jumps and the Zeno effect”. In: *Phys. Rev. A* 77 (1 Jan. 2008), p. 012112. DOI: 10.1103/PhysRevA.77.012112. URL: <https://link.aps.org/doi/10.1103/PhysRevA.77.012112>.
- [53] A. A. Clerk et al. “Introduction to quantum noise, measurement, and amplification”. In: *Rev. Mod. Phys.* 82 (2 Apr. 2010), pp. 1155–1208. DOI: 10.1103/RevModPhys.82.1155. URL: <https://link.aps.org/doi/10.1103/RevModPhys.82.1155>.
- [54] Alexander N. Korotkov. “Quantum Bayesian approach to circuit QED measurement”. In: (Nov. 2011), p. 11. arXiv: 1111.4016.
- [55] Alexander N. Korotkov. *Quantum Bayesian approach to circuit QED measurement*. 2011. arXiv: 1111.4016 [quant-ph].
- [56] H. M. Wiseman. “Adaptive phase measurements of optical modes: going beyond the marginal Q distribution”. In: *Phys. Rev. Lett.* 75 (1995), p. 4587.
- [57] Leigh S. Martin et al. “Implementation of a canonical phase measurement with quantum feedback”. In: *Nature Physics* 16.10 (Oct. 2020), pp. 1046–1049. ISSN: 1745-2481. DOI: 10.1038/s41567-020-0939-0. URL: <https://doi.org/10.1038/s41567-020-0939-0>.
- [58] Paul Busch et al. “Are number and phase complementary observables?” In: *Journal of Physics A: Mathematical and General* 34.30 (2001).
- [59] Charles H. Townes. “1964 Nobel lecture: Production of coherent radiation by atoms and molecules”. In: *IEEE Spectrum* 2.8 (1965), pp. 30–43. DOI: 10.1109/MSPEC.1965.6501319.
- [60] M J W Hall and I G Fuss. “Quantum phase detection and digital communication”. In: *Quantum Optics: Journal of the European Optical Society Part B* 3.3 (June 1991), pp. 147–167. DOI: 10.1088/0954-8998/3/3/002. URL: <https://doi.org/10.1088/0954-8998/3/3/002>.

- [61] Nicola Dalla Pozza, Howard M Wiseman, and Eleanor H Huntington. “Deterministic preparation of superpositions of vacuum plus one photon by adaptive homodyne detection: experimental considerations”. In: *New Journal of Physics* 17.1 (2015), p. 013047.
- [62] D. W. Berry et al. “How to perform the most accurate possible phase measurements”. In: *Phys. Rev. A* 80 (5 Nov. 2009), p. 052114. DOI: 10.1103/PhysRevA.80.052114. URL: <https://link.aps.org/doi/10.1103/PhysRevA.80.052114>.
- [63] Lev Landau and Rudolf Peierls. “Quantenelektrodynamik im konfigurationsraum”. In: *Zeitschrift für Physik* 62.3-4 (1930), pp. 188–200.
- [64] U. M. Titulaer and R. J. Glauber. “Density Operators for Coherent Fields”. In: *Phys. Rev.* 145 (4 May 1966), pp. 1041–1050. DOI: 10.1103/PhysRev.145.1041. URL: <https://link.aps.org/doi/10.1103/PhysRev.145.1041>.
- [65] C. W. Gardiner and M. J. Collett. “Input and output in damped quantum systems: Quantum stochastic differential equations and the master equation”. In: *Phys. Rev. A* 31 (6 June 1985), pp. 3761–3774. DOI: 10.1103/PhysRevA.31.3761. URL: <https://link.aps.org/doi/10.1103/PhysRevA.31.3761>.
- [66] KW Murch et al. “Cavity-assisted quantum bath engineering”. In: *Physical Review Letters* 109.18 (2012), p. 183602.
- [67] Chris Macklin et al. “A near-quantum-limited Josephson traveling-wave parametric amplifier”. In: *Science* 350.6258 (2015), pp. 307–310.
- [68] K W Murch et al. “Observing single quantum trajectories of a superconducting quantum bit”. In: *Nature* 502 (2013), p. 211.
- [69] A Eddins et al. “High-efficiency measurement of an artificial atom embedded in a parametric amplifier”. In: *Physical Review X* 9.1 (2019), p. 011004.
- [70] Joseph Kerckhoff et al. “On-chip superconducting microwave circulator from synthetic rotation”. In: *Physical Review Applied* 4.3 (2015), p. 034002.
- [71] Mohan Sarovar and K Birgitta Whaley. “Adaptive homodyne phase discrimination and qubit measurement”. In: *Phys. Rev. A* 76 (2007), p. 052316.
- [72] TC Ralph, AP Lund, and HM Wiseman. “Adaptive phase measurements in linear optical quantum computation”. In: *Journal of Optics B: Quantum and Semiclassical Optics* 7.10 (2005), S245.
- [73] Ognyan Oreshkov and Todd A Brun. “Weak measurements are universal”. In: *Physical Review Letters* 95.11 (2005), p. 110409.
- [74] John Preskill. “Lecture notes for Physics 219: Quantum computation”. In: *Caltech Lecture Notes* (1999).
- [75] Daniel Gottesman. “An Introduction to Quantum Error Correction”. In: (Apr. 2000), p. 15. URL: <https://arxiv.org/abs/quant-ph/0004072>.

- [76] Austin G. Fowler et al. “Surface codes: Towards practical large-scale quantum computation”. In: *Phys. Rev. A* 86 (3 Sept. 2012), p. 032324. DOI: 10.1103/PhysRevA.86.032324. URL: <https://link.aps.org/doi/10.1103/PhysRevA.86.032324>.
- [77] Daniel Gottesman, Alexei Kitaev, and John Preskill. “Encoding a qubit in an oscillator”. In: *Phys. Rev. A* 64 (1 June 2001), p. 012310. DOI: 10.1103/PhysRevA.64.012310. URL: <https://link.aps.org/doi/10.1103/PhysRevA.64.012310>.
- [78] Fernando Pastawski et al. “Holographic quantum error-correcting codes: toy models for the bulk/boundary correspondence”. In: *Journal of High Energy Physics* 2015.6 (June 2015), p. 149. ISSN: 1029-8479. DOI: 10.1007/JHEP06(2015)149. URL: [https://doi.org/10.1007/JHEP06\(2015\)149](https://doi.org/10.1007/JHEP06(2015)149).
- [79] P. W. Shor. “Fault-tolerant quantum computation”. In: *Proceedings of 37th Conference on Foundations of Computer Science*. 1996, pp. 56–65. DOI: 10.1109/SFCS.1996.548464.
- [80] Kurt Jacobs. *Quantum measurement theory and its applications*. Cambridge University Press, 2014.
- [81] Philipp Schindler et al. “Experimental Repetitive Quantum Error Correction”. In: *Science* 332.6033 (2011), pp. 1059–1061. ISSN: 0036-8075. DOI: 10.1126/science.1203329. eprint: <https://science.sciencemag.org/content/332/6033/1059.full.pdf>. URL: <https://science.sciencemag.org/content/332/6033/1059>.
- [82] V. Negnevitsky et al. “Repeated multi-qubit readout and feedback with a mixed-species trapped-ion register”. In: *Nature* 563.7732 (Nov. 2018), pp. 527–531. ISSN: 1476-4687. DOI: 10.1038/s41586-018-0668-z. URL: <https://doi.org/10.1038/s41586-018-0668-z>.
- [83] Norbert M. Linke et al. “Fault-tolerant quantum error detection”. In: *Science Advances* 3.10 (2017). DOI: 10.1126/sciadv.1701074. eprint: <https://advances.sciencemag.org/content/3/10/e1701074.full.pdf>. URL: <https://advances.sciencemag.org/content/3/10/e1701074>.
- [84] J. Cramer et al. “Repeated quantum error correction on a continuously encoded qubit by real-time feedback”. In: *Nature Communications* 7.1 (May 2016), p. 11526. ISSN: 2041-1723. DOI: 10.1038/ncomms11526. URL: <https://doi.org/10.1038/ncomms11526>.
- [85] J. Kelly et al. “State preservation by repetitive error detection in a superconducting quantum circuit”. In: *Nature* 519.7541 (Mar. 2015), pp. 66–69. ISSN: 1476-4687. DOI: 10.1038/nature14270. URL: <https://doi.org/10.1038/nature14270>.
- [86] Nissim Ofek et al. “Extending the lifetime of a quantum bit with error correction in superconducting circuits”. In: *Nature* 536.7617 (Aug. 2016), pp. 441–445. ISSN: 1476-4687. DOI: 10.1038/nature18949. URL: <https://doi.org/10.1038/nature18949>.

- [87] Christian Kraglund Andersen et al. “Repeated quantum error detection in a surface code”. In: *Nature Physics* 16.8 (Aug. 2020), pp. 875–880. ISSN: 1745-2481. DOI: 10.1038/s41567-020-0920-y. URL: <https://doi.org/10.1038/s41567-020-0920-y>.
- [88] C. C. Bultink et al. “Protecting quantum entanglement from leakage and qubit errors via repetitive parity measurements”. In: *Science Advances* 6.12 (2020). DOI: 10.1126/sciadv.aay3050. eprint: <https://advances.sciencemag.org/content/6/12/eaay3050.full.pdf>. URL: <https://advances.sciencemag.org/content/6/12/eaay3050>.
- [89] Diego Ristè et al. “Real-time processing of stabilizer measurements in a bit-flip code”. In: *npj Quantum Information* 6.1 (Aug. 2020), p. 71. ISSN: 2056-6387. DOI: 10.1038/s41534-020-00304-y. URL: <https://doi.org/10.1038/s41534-020-00304-y>.
- [90] Roman Stricker et al. “Experimental deterministic correction of qubit loss”. In: *Nature* 585.7824 (Sept. 2020), pp. 207–210. ISSN: 1476-4687. DOI: 10.1038/s41586-020-2667-0. URL: <https://doi.org/10.1038/s41586-020-2667-0>.
- [91] P. Campagne-Ibarcq et al. “Persistent Control of a Superconducting Qubit by Stroboscopic Measurement Feedback”. In: *Phys. Rev. X* 3 (2 May 2013), p. 021008. DOI: 10.1103/PhysRevX.3.021008. URL: <https://link.aps.org/doi/10.1103/PhysRevX.3.021008>.
- [92] G. de Lange et al. “Reversing Quantum Trajectories with Analog Feedback”. In: *Phys. Rev. Lett.* 112 (8 Feb. 2014), p. 080501. DOI: 10.1103/PhysRevLett.112.080501. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.112.080501>.
- [93] Rusko Ruskov and Alexander N. Korotkov. “Entanglement of solid-state qubits by measurement”. In: *Phys. Rev. B* 67 (24 June 2003), p. 241305. DOI: 10.1103/PhysRevB.67.241305. URL: <https://link.aps.org/doi/10.1103/PhysRevB.67.241305>.
- [94] B. Trauzettel et al. “Parity meter for charge qubits: An efficient quantum entangler”. In: *Phys. Rev. B* 73 (23 June 2006), p. 235331. DOI: 10.1103/PhysRevB.73.235331. URL: <https://link.aps.org/doi/10.1103/PhysRevB.73.235331>.
- [95] Nathan S. Williams and Andrew N. Jordan. “Entanglement genesis under continuous parity measurement”. In: *Phys. Rev. A* 78 (6 Dec. 2008), p. 062322. DOI: 10.1103/PhysRevA.78.062322. URL: <https://link.aps.org/doi/10.1103/PhysRevA.78.062322>.
- [96] N. Roch et al. “Observation of Measurement-Induced Entanglement and Quantum Trajectories of Remote Superconducting Qubits”. In: *Phys. Rev. Lett.* 112 (17 Apr. 2014), p. 170501. DOI: 10.1103/PhysRevLett.112.170501. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.112.170501>.

[97] Areeya Chantasri et al. “Quantum Trajectories and Their Statistics for Remotely Entangled Quantum Bits”. In: *Phys. Rev. X* 6 (4 Dec. 2016), p. 041052. DOI: 10.1103/PhysRevX.6.041052. URL: <https://link.aps.org/doi/10.1103/PhysRevX.6.041052>.

[98] Charlene Ahn, Andrew C. Doherty, and Andrew J. Landahl. “Continuous quantum error correction via quantum feedback control”. In: *Phys. Rev. A* 65 (4 Mar. 2002), p. 042301. DOI: 10.1103/PhysRevA.65.042301. URL: <https://link.aps.org/doi/10.1103/PhysRevA.65.042301>.

[99] Joseph Kerckhoff et al. “Physical model of continuous two-qubit parity measurement in a cavity-QED network”. In: *Phys. Rev. A* 79 (2 Feb. 2009), p. 024305. DOI: 10.1103/PhysRevA.79.024305. URL: <https://link.aps.org/doi/10.1103/PhysRevA.79.024305>.

[100] Razieh Mohseninia et al. “Always-On Quantum Error Tracking with Continuous Parity Measurements”. In: *Quantum* 4 (Nov. 2020), p. 358. ISSN: 2521-327X. DOI: 10.22331/q-2020-11-04-358. URL: <https://doi.org/10.22331/q-2020-11-04-358>.

[101] Kevin Lalumière, J. M. Gambetta, and Alexandre Blais. “Tunable joint measurements in the dispersive regime of cavity QED”. In: *Phys. Rev. A* 81 (4 Apr. 2010), p. 040301. DOI: 10.1103/PhysRevA.81.040301. URL: <https://link.aps.org/doi/10.1103/PhysRevA.81.040301>.

[102] Howard M. Wiseman and Gerard J. Milburn. *Quantum Measurement and Control*. Cambridge University Press, 2009. DOI: 10.1017/CBO9780511813948.

[103] Hideo Mabuchi. “Continuous quantum error correction as classical hybrid control”. In: *New Journal of Physics* 11.10 (Oct. 2009), p. 105044. DOI: 10.1088/1367-2630/11/10/105044. URL: <https://dx.doi.org/10.1088/1367-2630/11/10/105044>.

[104] Juan Atalaya, Alexander N. Korotkov, and K. Birgitta Whaley. “Error-correcting Bacon-Shor code with continuous measurement of noncommuting operators”. In: *Phys. Rev. A* 102 (2 Aug. 2020), p. 022415. DOI: 10.1103/PhysRevA.102.022415. URL: <https://link.aps.org/doi/10.1103/PhysRevA.102.022415>.

[105] J. Atalaya et al. *Continuous quantum error correction for evolution under time-dependent Hamiltonians*. 2020. eprint: [arXiv:2003.11248](https://arxiv.org/abs/2003.11248).

[106] Jay Gambetta et al. “Qubit-photon interactions in a cavity: Measurement-induced dephasing and number splitting”. In: *Phys. Rev. A* 74 (4 Oct. 2006), p. 042318. DOI: 10.1103/PhysRevA.74.042318. URL: <https://link.aps.org/doi/10.1103/PhysRevA.74.042318>.

[107] A. Potts et al. “CMOS compatible fabrication methods for submicron Josephson junction qubits”. In: *IEE Proceedings - Science, Measurement and Technology* 148.5 (2001), pp. 225–228. DOI: 10.1049/ip-smt:20010395.

- [108] A. Dunsworth et al. “Characterization and reduction of capacitive loss induced by sub-micron Josephson junction fabrication in superconducting qubits”. In: *Applied Physics Letters* 111.2 (2017), p. 022601. DOI: 10.1063/1.4993577. eprint: <https://doi.org/10.1063/1.4993577>. URL: <https://doi.org/10.1063/1.4993577>.
- [109] J. M. Chow et al. “Detecting highly entangled states with a joint qubit readout”. In: *Phys. Rev. A* 81 (6 June 2010), p. 062325. DOI: 10.1103/PhysRevA.81.062325. URL: <https://link.aps.org/doi/10.1103/PhysRevA.81.062325>.
- [110] Mollie Elisheva Schwartz. “Engineering Dissipation to Generate Entanglement Between Remote Superconducting Quantum Bits”. PhD thesis. UC Berkeley, 2016.
- [111] Baptiste Royer, Shruti Puri, and Alexandre Blais. “Qubit parity measurement by parametric driving in circuit QED”. In: *Science Advances* 4.11 (2018). DOI: 10.1126/sciadv.aau1695. eprint: <https://advances.sciencemag.org/content/4/11/eaau1695.full.pdf>. URL: <https://advances.sciencemag.org/content/4/11/eaau1695>.
- [112] David P DiVincenzo and Fırat Solgun. “Multi-qubit parity measurement in circuit quantum electrodynamics”. In: *New Journal of Physics* 15.7 (July 2013), p. 075001. DOI: 10.1088/1367-2630/15/7/075001. URL: <https://iopscience.iop.org/article/10.1088/1367-2630/15/7/075001>.
- [113] A. Kandala et al. *Demonstration of a High-Fidelity CNOT for Fixed-Frequency Transmons with Engineered ZZ Suppression*. 2020. arXiv: [arXiv:2011.07050](https://arxiv.org/abs/2011.07050) [quant-ph].
- [114] A. Frisk Kockum, L. Tornberg, and G. Johansson. “Undoing measurement-induced dephasing in circuit QED”. In: *Phys. Rev. A* 85 (5 May 2012), p. 052318. DOI: 10.1103/PhysRevA.85.052318. URL: <https://link.aps.org/doi/10.1103/PhysRevA.85.052318>.
- [115] R. Vijay et al. “Stabilizing Rabi oscillations in a superconducting qubit using quantum feedback”. In: *Nature* 490 (2012), p. 77.
- [116] E. Flurin et al. “Using a Recurrent Neural Network to Reconstruct Quantum Dynamics of a Superconducting Qubit from Physical Observations”. In: *Phys. Rev. X* 10 (1 Jan. 2020), p. 011006. DOI: 10.1103/PhysRevX.10.011006. URL: <https://link.aps.org/doi/10.1103/PhysRevX.10.011006>.
- [117] S. Hacohen-Gourgy and L. S. Martin. “Continuous measurements for control of superconducting quantum circuits”. In: *Advances in Physics: X* 5.1 (2020), p. 1813626. DOI: 10.1080/23746149.2020.1813626. eprint: <https://doi.org/10.1080/23746149.2020.1813626>. URL: <https://doi.org/10.1080/23746149.2020.1813626>.
- [118] John Mark Kreikebaum. “Superconducting Qubit Enabled Single Microwave Photon Detection”. PhD thesis. UC Berkeley, 2020.
- [119] Innovative Integration. *X6-1000M User’s Manual*. 2016. URL: www.innovative-dsp.com.

- [120] Innovative Integration. *X6-1000M FrameWork Logic Guide*. 2016. URL: www.innovative-dsp.com.
- [121] David C. McKay et al. “Efficient Z gates for quantum computing”. In: *Phys. Rev. A* 96 (2 Aug. 2017), p. 022330. DOI: 10.1103/PhysRevA.96.022330. URL: <https://link.aps.org/doi/10.1103/PhysRevA.96.022330>.