

THE UNIVERSITY OF CHICAGO

EFFICIENT CONTROL AND TOMOGRAPHY IN
MULTIMODE CIRCUIT QUANTUM ELECTRODYNAMICS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
DEPARTMENT OF PHYSICS

BY
KEVIN HE

CHICAGO, ILLINOIS
AUGUST 2023

Table of Contents

List of Figures	v
List of Tables	viii
Acknowledgments	ix
Abstract	xi
1 Introduction	1
1.1 Quantum hardware and scaling up to a quantum computer	1
1.2 Circuit Quantum Electrodynamics (cQED)	2
1.3 Overview of Thesis	5
2 3D Circuit Cavity System Hardware	7
2.1 3D Transmon	7
2.1.1 Transmon Background	7
2.1.2 Transmon Fabrication	9
2.2 3D Multimode Cavity	11
2.2.1 Cavity Background	11
2.2.2 Cavity Fabrication	16
2.3 Coupling the Transmon and Multimode Cavity	18
2.4 Auxiliary Setup	20
2.4.1 Cryogenic Hardware and Wiring	20
2.4.2 Control Electronics	21
3 Calibration of the System Hamiltonian	24
3.1 Readout Cavity	24

3.2	Transmon Qubit	26
3.3	Multimode Cavity	37
3.3.1	Initial Cavity Parameters	37
3.3.2	Sideband Experiments	42
3.3.3	Cavity Single Photon Experiments	47
3.3.4	Cavity Beamsplitter Experiments	53
4	Cavity Photon Blockade	58
4.1	Blockade Derivation	58
4.2	Generating the Blockade	62
4.3	Multimode Blockade	66
5	Quantum Optimal Control in 3D cQED	75
5.1	Overview of Optimal Control Algorithms	75
5.2	Optimal Control of a Transmon and Cavity Mode	81
5.2.1	Developing the System Hamiltonian	82
5.2.2	Calibration Experiments	84
5.2.3	Preparing Fock states	87
5.3	Optimal Control with Photon Blockade	91
5.4	Multimode Optimal Control	94
6	State Characterization and Tomography	98
6.1	Wigner Tomography	98
6.2	Generalized Multimode Wigner Tomography	104
6.3	Tomography of W states	106
6.4	Efficient Tomography Sampling: DEMESST	114
6.4.1	Outlining the DEMESST Sampling Method	115

6.4.2	DEMESST Derivation	118
6.4.3	Experimental and Simulated Results	122
6.4.4	W^2 Tomography Sampling	130
7	Conclusion	133
	References	135

List of Figures

2.1	Transmon and cavity circuits and potentials	8
2.2	Fabricated transmon device	10
2.3	Field profiles of a non-tapered and tapered multimode flute cavity	15
2.4	Steps in the flute cavity fabrication method	17
2.5	Schematic of the coupled 3D transmon and cavity system	20
2.6	System wiring diagram	22
3.1	Resonator spectroscopy and readout-storage cross-Kerr	25
3.2	Calibration of a transmon qubit	28
3.3	Readout fidelity histogram experiment	30
3.4	Qubit coherence experiments	31
3.5	Calibration of transmon $ e\rangle - f\rangle$ levels	33
3.6	DRAG pulse calibration and implementation	36
3.7	Initial cavity calibrations	39
3.8	Cavity self-Kerr calibration	41
3.9	$ f0\rangle - g1\rangle$ sideband experiments	45
3.10	Cavity calibrations that utilize single photon state preparation	48
3.11	Multimode cavity dispersive shift and Kerr values	49
3.12	Cavity coherence experiments	51
3.13	Multimode cavity coherences and Q s	52
3.14	Cavity beamsplitter calibration experiments	56
4.1	Photon blockade infidelity	61
4.2	Blockade calibrations	63
4.3	Blockade preparation of $ 1\rangle$	64

4.4	Blockade preparation of $ 2\rangle$	65
4.5	Sideband χ dressing	67
4.6	Hilbert space partition implemented by multimode photon blockade	69
4.7	Demonstrating the blockade with spectroscopy	70
4.8	W state preparation Rabi oscillations	72
4.9	Blockading $n = 3$ for 2 modes	73
5.1	Schematic of GRAPE algorithm	80
5.2	Qubit and cavity drive calibrations	85
5.3	Cavity drive Fourier transform comparison	86
5.4	QOC Fock state preparation on a cavity mode	89
5.5	QOC Fock state preparation results	90
5.6	Universal control of a qudit using optimal control	93
5.7	QOC multimode preparation of $ 11\rangle$	96
6.1	Wigner tomography examples	100
6.2	Sets of optimized Wigner tomography displacements	101
6.3	Parity measurement bandwidth	103
6.4	Parity measurement times for multiple modes	105
6.5	2-mode W state Wigner tomography	108
6.6	Reconstructed 2-mode W state density matrix	109
6.7	3-mode W state Wigner tomography	110
6.8	Reconstructed 3-mode W state density matrix	111
6.9	Reconstructed 4-mode W state density matrix	112
6.10	Entanglement witness for W states of 2-4 modes.	113
6.11	DEMESST tomography sampling method schematic and theoretical scaling.	116
6.12	DEMESST tomography sequence and sampling points	123

6.13	DEMESST and W_{inv} W state density matrix reconstructions	124
6.14	Fidelities and matrix distances for DEMESST and W_{inv} sampling methods vs. measurement number	126
6.15	Simulated matrix distances for vs. measurement number for DEMESST and W_{inv}	128
6.16	Self-consistency verification for the DEMESST method	129
6.17	Experimental results for the W^2 tomography sampling method	131

List of Tables

3.1	Transmon and readout system parameters	34
3.2	Multimode cavity system parameters	53

Acknowledgments

Pursuing a PhD would not have been possible without the support of many people, some of whom I will highlight now. First, I would like to thank my advisor Dave Schuster for cultivating a friendly, collaborative lab environment and for his constant mentorship. Regardless of whatever engineering or physics problem I might run into, Dave always offers creative solutions and provides unique ideas for exciting things to try next. I am grateful to have been able to experience graduate school in his lab with other Schuster Lab members. Starting from the beginning, I would like to thank Yao Lu, Akash Dixit, and Tanay Roy for taking the time to help teach me when I was a new graduate student, and for always being reliable resources for questions about the clean room and our device fabrication processes. Fang Zhao, Kan-Heng Lee, and Alexander (Sasha) Anferov have also helped improve the clean room experience and make it more efficient. I would like to recognize Gabrielle Roberts for sharing in the new student experience and making it less overwhelming, and along with Margaret Panetta organizing many of the social events that helped the lab bond, connect, and have fun. I would like to acknowledge Helin Zhang and Chunyang Ding for being flexible with their timeline and allowing us to share a fridge efficiently. I would also like to acknowledge Liang Jiang and his students, Ming Yuan and Yat Wong, for their theory support and countless discussions that led to productive work and eventual results. I would like to recognize the multimode project members: Srivatsan (Vatsan) Chakram, Andrew Oriani, Ankur Agrawal, and Riju Banerjee. In particular, Andrew helped teach me how to use the software required to design and simulate new devices, and always ensured that the lab as a whole followed best practices for safety and productivity. Ankur's creativity allowed us to overcome several hardware challenges and effectively share a single device. I would like to especially thank Vatsan for being essentially another advisor for me and for being a phenomenal teacher, source of inspiration, and role model for what it means to be a scien-

tist. I have learned so much from working together. Personally, I would like to acknowledge my family: Mom (Sophia), Dad (Gregory), and Joanne. Their encouragement, help, and sacrifice have made it possible for me to pursue this level of education and be the best I can be. I would also like to thank Evan Li for being someone that I could always talk to about anything.

Abstract

Continued improvements to superconducting architectures have placed them as one of the leading quantum systems. In particular, circuit quantum electrodynamics (cQED) combines the strong interactions and controllability of superconducting circuits with the high coherence times of 3D cavities. These features readily allow for the implementation a hardware-efficient quantum processor, even with only a single nonlinear device. In this thesis, we explore one such system that consists of a superconducting transmon qubit coupled to a 3D multimode flute cavity that supports 9 modes with qubit-mode cooperativities exceeding 1 billion and mode lifetimes ~ 2 ms. First, we describe the system and demonstrate how to characterize it, determining the parameters necessary to effectively generate and manipulate quantum states. Next, we expand the cQED toolbox by extending state preparation techniques like photon blockade and quantum optimal control to multiple modes, implementing qudits and a pure multimodal interaction that can engineer the modes' Hilbert space. We use this interaction to easily generate entangled W states, which we demonstrate for up to 5 modes. Finally, to accurately characterize our complex states, we develop a generalized Wigner tomography method that functions even with imperfect parity measurements. We also introduce a new Wigner tomography sampling method whose required number of observations scales polynomially with subspace size, avoiding the traditional exponential scaling with mode number for certain states and thereby allowing for much more efficient state reconstruction. The developments presented in this work enables further study of complicated quantum states, especially multimode ones, with applications in quantum simulation, quantum optics, and quantum information.

Chapter 1

Introduction

1.1 Quantum hardware and scaling up to a quantum computer

Development of a quantum computer is one of the frontiers in contemporary technology. This device promises computational capabilities that cannot be achieved with classical computers, and has motivated intensive scientific investigation into quantum phenomena. The fundamental building block of a quantum computer is the qubit, analogous to a classical bit, with the main distinction being that instead of being exactly 0 or 1, the qubit can be in any superposition between them. This can be experimentally implemented on any hardware that supports a two-level system, such as atoms, superconducting circuits, trapped ions, spins, and diamond vacancy centers, to name just a few. While the platforms may vary, the long-term end goals are the same: to implement a quantum processor capable of demonstrating quantum advantage, and to use that device to accomplish useful operations beyond the capabilities of classical computers. Although the full realization of this technology is still quite a ways out, steps are being taken every day to improve our understanding and ability to manipulate quantum systems.

One category of challenges in developing quantum technology is effectively scaling up the size of quantum systems. Like with a classical computer, we will need many qubits to implement quantum algorithms. However, each additional qubit introduces extra sources of error. On a related note, figuring out the geometry of how to physically contain or place multiple qubits in a single device while retaining the ability to control all of them can be another source of difficulty. Furthermore, effectively controlling the qubits while avoiding

unwanted stray couplings or interactions between them is another obstacle. These are large, overarching problems that are only a small subset of the issues that will be overcome on the way to developing a working quantum computer. With addressing some of these challenges as motivation, throughout this thesis we will discuss a quantum device that offers a potential avenue for efficiently scaling up the size of a quantum system.

1.2 Circuit Quantum Electrodynamics (cQED)

The specific hardware platform that we utilize in this thesis is superconducting circuits and cavities. The study of the combination of these two types of devices is commonly referred to as circuit quantum electrodynamics (cQED). This will be discussed in more detail in Chapter 2, but we provide a brief overview here.

As the name suggests, superconducting circuits are the quantum analog of classical circuits, and are essentially artificial atoms that can implement a qubit. This is accomplished by combining a Josephson junction, which can be modeled as a nonlinear inductor, with a capacitive element. The nonlinearity of the junction shifts the energy levels of this LC oscillator so that they are no longer harmonic or equally spaced. Instead, there is some anharmonicity that allows us to target a single transition, such as the one between the ground state and first excited state, to implement a two-level system that defines the qubit. These artificial atoms are typically on the scale of hundreds of nanometers, and are operated at microwave frequencies (\sim GHz), though there are some prominent superconducting qubits such as fluxonium or millimeter wave designs that deviate from these parameters.

Superconducting qubits have several advantages that have made them one of the leading platforms for quantum computing. These features are the ability to create custom device geometries and arbitrarily set the qubit transition frequency. In general, we can fabricate these artificial atoms in whatever configuration and with whatever parameters we want,

providing a great amount of flexibility in circuit design and optimization. Additionally, by modifying the circuit layout, we can control the interaction strength or coupling of these qubits with other components in our system. This can be used to engineer specific setups such as chains or grids of qubits, and adjust or tune the coupling between each qubit pair to obtain almost completely arbitrary connectivity that is limited only by the amount of physical space available. Creating our own artificial atoms leads to these degrees of freedom in the parameters of our quantum system and the Hamiltonian that defines it, offering adjustable large anharmonicities, a wide range of tunability, and controllable interaction strengths. These features make superconducting qubits an excellent candidate for creating manipulable, large scale quantum devices.

The other main component in superconducting quantum systems is superconducting cavities. Unlike qubits, they are intrinsically linear, and can be treated as a classical LC harmonic oscillator. Nevertheless, they are still very useful in quantum systems, as they are commonly used to couple to qubits and protect them from the environment and error sources such as thermal noise. This is accomplished by modifying the measurement procedure to observe the state of the cavity, which will be affected by the state of the qubit, instead of directly measuring the qubit itself. This indirect readout process exposes the qubit to fewer sources of error that are on resonance with its transition frequency and suppresses the magnitude of such errors, thereby increasing the lifetimes or coherences of quantum information stored in the device. For most effective operation, the qubit and cavity will often be off-resonant with one another, such that the detuning is much larger than the coupling between them. For ease of control, they will both usually still be in the microwave frequency range.

Another application of superconducting cavities is their ability to act as a quantum memory or storage. Although they are intrinsically classical devices, by coupling to a qubit, they can inherit some of its nonlinearity and support non-classical, quantum states. By engineer-

ing these resonators to have electromagnetic fields that live primarily in a lossless cavity volume, we can create cavities with photon lifetimes that are an order of magnitude or more longer than superconducting qubits, significantly enhancing the amount of time we have to manipulate and work with quantum states prepared in them. Additionally, the large Hilbert spaces of these cavities can be used for quantum error correction by encoding information in specific states where certain types of errors are easily detectable and correctable. This can further push the coherence of quantum information stored in these systems to increasingly longer times. Finally, perhaps the most important feature with regards to scaling up a quantum system is the hardware efficiency potential of superconducting cavities, in particular 3D cavities. These waveguide resonators can support a large forest of modes or cavity sites at different microwave frequencies on a single physical device, and can be engineered to create a naturally multiplexed architecture that only requires a few control lines. Like superconducting qubits, these cavities are highly customizable, and we can design them for whatever readout or storage parameters we desire.

By combining these components, we leverage the controllability of superconducting qubits with the long lifetimes and hardware efficiency of superconducting cavities to create a highly customizable, scalable quantum processor system. We couple our artificial atom with 3D resonator modes through a dipole antenna. On this platform, we utilize a single qubit to simultaneously and universally control many cavity modes, with the ability to easily generate many-body and higher-order interactions between the modes. Besides using the cavity to encode the state of the qubit as described earlier, we use this device for experiments in quantum information processing, quantum simulation, and quantum metrology. As a final note, treating these devices as modules and combining multiple copies of them offers a promising avenue for increasing the size of an overall quantum system and scaling up to a quantum computer.

1.3 Overview of Thesis

In this thesis, we start in Chapter 2 by introducing our quantum hardware system and the associated control electronics and cryogenic setup. With a number of hardware innovations, we are able to effectively operate a coupled superconducting transmon qubit and 3D multimode flute cavity with lifetimes in excess of $100 \mu\text{s}$ and 2 ms , respectively, and with cooperativities $> 10^9$. In Chapter 3, we will describe the experiments necessary to fully calibrate the parameters of this quantum system. This includes developing full control over the qubit and cavity, as well as several operations that involve both components and can be used to effectively engineer quantum states. After finishing the calibrations that allow us to precisely control our hardware, in Chapter 4 we will discuss photon blockade, one specific control tool that we apply to manipulate our system and generate interactions between photons that live in distinct cavity modes. With this interaction, we can easily partition the multimode cavity’s bosonic Hilbert space based on total photon number. We apply this tool to prepare entangled multimode W states—states with a single photon shared between multiple cavity modes—with only constant drives on the qubit and cavity. For more complicated controls, we turn to Chapter 5 on quantum optimal control (QOC), the general term for numerically optimizing pulses that achieve a target quantum state or gate. We summarize two particular methods, GRAPE and PICO, and demonstrate QOC by preparing Fock states on a variety of system Hamiltonians, such as a qubit and single cavity mode, a qubit and multiple cavity modes, and a qutrit with a Hamiltonian engineered by the blockade. Finally, in Chapter 6 we describe how we efficiently measure the state of our quantum cavities with Wigner tomography. We describe a technique that allows us to perform the tomography even in imperfect conditions, which we apply to simultaneously measure multiple modes. We then present a tomography sampling method with improved efficiency that scales only polynomially rather than exponentially with the subspace size and thus number of modes for some states. This is

accomplished by utilizing information about the expected state in a chosen basis. We apply this tomography to our multimode W states and observe the expected improvement with mode number. This method enables us to maintain the tractability of state measurement, even for increasingly large multimode states. Ultimately, these control techniques and measurement developments facilitate the future generation and characterization of progressively more complicated quantum states.

Chapter 2

3D Circuit Cavity System Hardware

A number of engineering and design innovations have gone into the construction of our quantum device and hardware system [1, 2]. This chapter will discuss the processes that we used to create our transmon and multimode cavity, as well as provide a summary of the cryogenics and control electronics required to successfully perform quantum experiments.

2.1 3D Transmon

2.1.1 Transmon Background

The transmon is one of the many types of superconducting circuits that utilize Josephson junctions to form a nonlinear device [1, 3]. The design consists of a single junction with a large shunting capacitance, forming a circuit depicted in Fig. 2.1 that can be described by the Hamiltonian:

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

where E_C and E_J are the charging and Josephson energies, \hat{n} and $\hat{\phi}$ are the normalized charge and phase operators across the junction, and n_g is the offset gate charge. Due to the cosine potential $E_J \cos \hat{\phi}$, the transmon will have energy levels that are not equally spaced, in contrast to the harmonic oscillator levels of a purely quadratic potential, as shown in Fig. 2.1. The main advantage of the transmon compared to other types of nonlinear devices is that it is insensitive to charge noise, i.e. fluctuations in n_g . In particular, the sensitivity to charge noise is suppressed exponentially by the ratio E_J/E_C , and for the transmon, $1 \ll E_J/E_C \ll 5 \times 10^4$ [1]. The first inequality condition guarantees the charge noise suppression, while the second ensures that the nonlinearity between energy levels (anharmonicity) $\alpha \simeq -E_C$ is not




Figure 2.1: **Transmon and cavity circuits and potentials.** Due to the cosine potential, the transmon (black) has nonlinear, unevenly spaced energy levels. This is in contrast to the harmonic oscillator cavity (orange), which has a quadratic potential with evenly spaced levels. The transmon portion is outlined by the gray dashed box, and features a nonlinear Josephson junction with a shunting capacitance. The cavity portion is outlined by the dashed orange box, and can be represented as an LC oscillator.

so small as to inhibit or prevent qubit operations. Additionally, the qubit frequency between the ground state $|g\rangle$ and first excited state $|e\rangle$ is $\omega_{ge} \approx \sqrt{8E_J E_C}$. The first few energy levels $|g\rangle, |e\rangle, |f\rangle$ are represented in Fig. 2.1.

In our device, we targeted an anharmonicity $\alpha \approx 150$ MHz and qubit frequency $\omega_{ge} \approx 5$ GHz. To create an appropriate design and determine the physical dimensions that would accomplish this (such as the size of the junction and its coupling pads), we used simulation software like Ansys HFSS (High Frequency Structure Simulator) or Python packages like PyEPR [4, 5]. These simulations are usually able to get us within roughly 20% of our targets, and provide a good starting point for fabrication of the device, which will be described in detail in the following subsection. With that level of accuracy in mind, it was often useful to vary the target design size when creating different copies of the device in the same fabrication run (i.e. changing the junction dimensions across different copies of the chip on the same wafer) for the best chances of creating a usable device within our target parameter range. This could also help compensate for potential changes between fabrication runs, such as fluctuations in instrument parameters or human error and deviations during the several days long fabrication process.

2.1.2 Transmon Fabrication

Our transmon qubit was fabricated on a 430 μm thick C-plane (0001) sapphire wafer with a 50.8 mm diameter. First, the wafer was cleaned with organic solvents (Toluene, Acetone, Methanol, Isopropanol, and DI water) in an ultrasonic bath to remove contamination, with frequencies varied at 104, 72, and 40 kHz, in that order, for 1 minute and 40 seconds each, resulting in 5 minutes of sonication total in each solvent. Varying the frequencies was meant to target and break up particles of different sizes. Afterwards, the wafer was annealed at 1200 $^{\circ}\text{C}$ for 1.5 hours in atmosphere. To reach the target temperature, the furnace required several ramp up steps: one going to 700 $^{\circ}\text{C}$ at a rate of 350 $^{\circ}\text{C}/\text{hr}$, step two going to 1000 $^{\circ}\text{C}$ at a rate of 250 $^{\circ}\text{C}/\text{hr}$, and step three going to 1200 $^{\circ}\text{C}$ at a rate of 100 $^{\circ}\text{C}/\text{hr}$. After the annealing, the furnace turned off and passively cooled until roughly 50 $^{\circ}\text{C}$, at which point the wafer was removed. Prior to deposition of a base aluminum film in a Plassys Evaporator, the wafer underwent a second clean with organic solvents (Toluene, Acetone, Methanol, Isopropanol, and DI water) in an ultrasonic bath with the same parameters as the initial clean.

The junction was made out of aluminum using a combination of optical and electron-beam lithography. The base layer of the device, which includes the capacitor pads for the transmon, consists of 120 nm of Al deposited via electron-beam evaporation at 1 $\text{\AA}/\text{s}$. The pads and chip features were defined via optical lithography using AZ MiR 703 photoresist and exposure by a Heidelberg MLA150 Direct Writer. After writing, the resist was developed for 1 minute in AZ MIF 300 1:1. The features were then etched in a Plasma-Therm inductively coupled plasma (ICP) etcher using chlorine based etch chemistry (30 sccm Cl_2 , 30 sccm BCl_2 , 10 sccm Ar). This was followed by a second layer of optical patterning and thermal evaporation of 50 nm of Au for the alignment marks used in the electron-beam lithography. The resist was subsequently removed by leaving the wafer in 80 $^{\circ}\text{C}$ N-Methyl-2-pyrrolidone (NMP) for at least 4 hours, and specifically for our device, 19 hours into the next day.




Figure 2.2: **Fabricated transmon device.** The final rectangular chip is a clear $430 \mu\text{m}$ thick sapphire substrate with two $120 \mu\text{m}$ thick aluminum coupling pads and the Manhattan style transmon junction. The area of the junction is roughly $170 \times 170 \text{ nm}$, as shown in the Scanning Electron Microscope (SEM) image shown in the red zoomed in box. The vertical line of aluminum far from the transmon and coupling pads was used to confirm proper alignment of the transmon chip in the whole device.

The junction mask was defined through electron-beam lithography of a bi-layer resist (MMA-PMMA) in the Manhattan pattern using a Raith EBPG5000 Plus Electron-Beam Writer, with overlap pads that allow for direct galvanic contact to the optically defined capacitors. After the lithography, the resist stack was developed for 1.5 minutes in a 120 mL solution of 3 parts IPA and 1 part DI water. Before the aluminum junction deposition in the Plassys Evaporator, the overlap regions on the pre-deposited capacitors were milled in-situ with an argon ion mill for 1 minute and 15 seconds to remove the native oxide. The junction was then deposited with a three-step electron-beam evaporation and oxidation process. First, an initial 35 nm layer of Al was deposited at 1 nm/s at an angle of 29° relative to the normal of the substrate, azimuthally parallel to one of the fingers in the Manhattan pattern. Next, the junction was exposed for 12 minutes to 20 mBar of a high-purity mixture of Ar and O_2 in an 80:20 ratio to allow the first layer to grow a native oxide. Finally, a second 120 nm layer of Al was deposited at 1 nm/s at the same 29° angle relative to the normal of the substrate, but azimuthally orthogonal to the first layer of Al.

After evaporation and before liftoff of the resist layers, the individual junction chip was cut out from the wafer using a Disco DAD3240 Dicing Saw. Prior to the dicing, a layer of AZ 1518 photoresist was applied to the wafer to protect the now exposed junctions. After cutting, the chip was left in 80°C NMP for 1 hour and 20 minutes to liftoff the MMA, PMMA, and AZ 1518 resists. To further clean the chip, it was transferred to a fresh beaker of 80°C NMP for 15 minutes, and then a beaker of IPA for 5 minutes. Finally, the chip was dried with a nitrogen gun and stored next to an ion-producing fan to avoid electrostatic discharge of the junction. An image of the final transmon chip and zoom in on the transmon is shown in Fig. 2.2. Besides the junction, the chip also contains two aluminum coupling pads that are used to couple the transmon to the 3D readout and storage cavities, which will be discussed in the following section. We call the transmon 3D because it couples to the 3D cavities through antenna pads, and there is no metal base layer on the transmon chip, such that besides the junction and pads, the chip consists solely of the bare sapphire substrate.

2.2 3D Multimode Cavity

2.2.1 Cavity Background

A single mode of a 3D superconducting cavity can be represented by an LC oscillator circuit. An example is depicted in Fig. 2.1. We can write the Hamiltonian of this circuit as

$$\hat{H} = \frac{\hat{\Phi}^2}{2L} + \frac{\hat{Q}^2}{2C}, \quad (2.2)$$

where Φ is the flux through the inductor and Q is the charge stored on the capacitor. We can rewrite Φ and Q in terms of the annihilation and creation operators \hat{a} and \hat{a}^\dagger as

$$\hat{\Phi} = \sqrt{\frac{\hbar Z_0}{2}}(\hat{a} + \hat{a}^\dagger), \quad \hat{Q} = i\sqrt{\frac{\hbar}{2Z_0}}(\hat{a}^\dagger - \hat{a}), \quad (2.3)$$

where $Z_0 = \sqrt{L/C}$ is the characteristic impedance and the operators satisfy the commutation relations $[\hat{a}, \hat{a}^\dagger] = 1$ and $[\hat{Q}, \hat{\Phi}] = i\hbar$. We can now rewrite our original Hamiltonian in Eqn. (2.2) in terms of \hat{a} and \hat{a}^\dagger , which gives

$$\hat{H} = \hbar\omega_c \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right), \quad (2.4)$$

where $\omega_c = 1/\sqrt{LC}$ is the frequency of the LC oscillator. We re-zero the ground state energy $\hbar\omega_c/2$, so that we can drop the $1/2$ from Eqn. (2.4). From this expression, we see that we have equally spaced harmonic oscillator energy levels at each nonnegative integer value of $\hat{N} = \hat{a}^\dagger \hat{a}$, as shown in Fig. 2.1. This generalizes to the multimode case as a sum of multiple copies of this single mode Hamiltonian, with different frequencies ω_{cj} and operators \hat{a}_j for each mode.

We implement our readout and multimode storage cavities using the seamless flute method [2, 6], which involves drilling two offset sets of cylindrical holes on either end of a block of material. The resulting roughly rectangular cavity volume is set by the overlap between the holes from opposite ends. This allows for highly controllable cavity parameters, such as mode frequencies and mode dispersion. In general, a rectangular cavity will support multiple transverse electric (TE_{ijk}) and transverse magnetic (TM_{ijk}) modes which are sorted and labeled according to their field profiles and frequencies. We control which modes are dominant and the lowest energy mode by changing the dimensions of the cavity. For example, we choose the dimensions of the readout cavity such that only one (TE_{101}) mode is supported in our range of operating frequencies (< 8 GHz), so that we can effectively treat it as a single mode cavity. Before discussing the multimode storage cavity, we address a possible concern with this flute method, the presence of the many holes required to form the cavity volume. These holes act as waveguides that couple to the external environment, providing a potential source of loss. However, we will now show that this is not a large effect.

For a single circular waveguide, the propagation constant is:

$$\beta_{11} = \sqrt{k^2 - \left(\frac{p'_{11}}{r}\right)^2}, \quad (2.5)$$

where we utilize the fact that we are considering the dominant TE₁₁ mode, and r is the radius of the waveguide, p'_{11} is the first root of $J'_1(x)$, and $k = \omega\sqrt{\mu\epsilon}$. Here, $J'_1(x)$ indicates the derivative of a Bessel function of the first kind, ω is the frequency of the waveguide mode, and ϵ and μ are the permittivity and permeability. Note that β_{11} is imaginary for frequencies below the cutoff frequency $\omega_c = p'_{11}/(r\sqrt{\mu\epsilon})$. Consequently, the external coupling of a mode with frequency $\omega < \omega_c$ will be exponentially suppressed with the length of the waveguide, i.e. the hole depth. Therefore, we choose a hole radius r such that ω_c is much larger than the frequencies of our cavity modes, causing the holes to be evanescent and their external coupling not to be a limiting factor on the internal quality factor of our cavity. The effect of each hole on the cavity is a decrease in the mode frequency or effective increase in mode volume [7], which can be calculated and compensated [6].

The main advantage of the flute method is the elimination of seams and thus seam loss. This has historically been a difficult engineering problem, as it is much easier to remove material to form a cavity volume by splitting it into two or more parts, which will then have a seam when put back together. Seam loss can be caused by surface imperfections at the seam or the formation of a resistive metal oxide layer between the cavity parts. We can model the seam as a boundary with conductance G , or conductance per unit length $g = G/l$. The resulting quality factor is $Q_{seam} = g/y_{seam}$, where

$$y_{seam} = \frac{\int_l |\vec{H}_{||}|^2 dl}{\omega\mu \int_V |\vec{H}|^2 dV} \quad (2.6)$$

is the seam admittance for a magnetic field \vec{H} [7, 8]. With the flute method, we avoid this

source of loss entirely. Additionally, the flute method's evanescent holes do not introduce radiative losses as long as we stay well below the cutoff frequency ω_c , as the imaginary propagation constant will cause the radiated energy to also be imaginary.

We now proceed to the storage cavity. For the multimode flute, we utilize the TE_{10m} modes of the rectangular cavity formed by the overlapping evanescent holes. These modes will have frequencies

$$\begin{aligned} f_{10m} &= \frac{c}{2\sqrt{\mu\epsilon}} \sqrt{\left(\frac{1}{h}\right)^2 + \left(\frac{0}{w}\right)^2 + \left(\frac{m}{l}\right)^2} \\ &= \frac{c}{2\sqrt{\mu\epsilon}} \sqrt{\left(\frac{1}{h}\right)^2 + \left(\frac{m}{l}\right)^2}, \end{aligned} \tag{2.7}$$

where $w < h \ll l$ are the width, height, and length of the cavity. We now encounter two challenges. To have the maximum number of usable modes in our multimode cavity, we would like the modes to be equally spaced, but currently, for the lowest order modes m , the mode frequency spacing $f_{10(m+1)} - f_{10m}$ is changing. Additionally, the electric field profile of each mode m will have $m-1$ nodes at different locations along the length of the cavity, which may make it difficult to find a suitable spot to place the transmon to effectively couple it to many modes. Fortunately, we can handle both of these issues simultaneously by adjusting the depth of the holes that make up the cavity. In particular, if we gradually decrease the height of a portion of the cavity by reducing the overlap between the sets of opposite holes, we replace cavity volume with evanescent hole volume. This changes the field profile, since it can now only evanescently decay into the section with reduced overlap. The cavity frequency remains roughly the same, as it is balanced by the decrease in frequency caused by the holes and the increase in frequency caused by a shorter effective propagation length. We choose the height adjustment to achieve a linear mode spacing, which we now derive. Expanding




Figure 2.3: **Field profiles of a non-tapered and tapered multimode flute cavity.** (a) Electric field profile of the fundamental mode of a non-tapered cavity. (b) Electric field profile of the 9th mode of a non-tapered cavity. The field of each higher frequency mode contains an additional node. (c) Electric field profile of the fundamental mode of a tapered cavity. Along the length x of the cavity, the depth of the holes changes such that the amount of overlap and effective height of the cavity decreases as x^2 . The shape of the field changes when compared to a non-tapered cavity, and is contained primarily in the region with larger cavity height. (d) Electric field profile of the 9th mode of a tapered cavity. (e) Comparison of the mode frequencies of a non-tapered (red, dashed) and tapered (blue, solid) cavity. The mode spacing of the non-tapered cavity changes, while the tapered cavity has the desired equally spaced modes.

Eqn. (2.7) for small m for the lowest modes, we have, to leading order in h/l :

$$f_{10m} \propto \frac{1}{h} + \frac{m^2 h}{2l^2}. \quad (2.8)$$

We can see that the mode frequency spacing will become linear when the m th cavity mode has an effective propagation length proportional to \sqrt{m} . This is accomplished when we choose $h(x) = h_0 - ax^2$, where $h(x)$ is the hole overlap height at distance x along the length of the cavity. Tapering the cavity quadratically in this way generates the desired constant mode spacing and also allows us to more easily find a suitable location to couple the transmon to the cavity. A comparison between a tapered and non-tapered cavity is shown in Fig. 2.3. We see in Fig. 2.3(e) that the tapered cavity has equally spaced modes, and the tapered field profiles shown in Fig. 2.3(c) and Fig. 2.3(d) are more concentrated on one end of the cavity such that we could place the transmon on that end and effectively couple it to many modes, despite the addition of a node for each successively higher mode frequency. More details and an in-depth discussion of the multimode flute cavity can be found in [2] or Andrew Oriani's thesis [6].

2.2.2 Cavity Fabrication

The readout cavity and multimode storage cavity are machined together from a single rectangular block of 5N5 aluminium (purity > 99.999%). Following the procedure of the seamless flute method [2], offset holes are drilled from either end of the aluminum block, such that the distance between the centers of opposing consecutive holes ($l_0/2 = 3.015$ mm) is less than the hole diameter ($2r = 4.76$ mm). After both sets of holes are drilled, there is an additional reaming step to ensure the accuracy of the final hole dimensions and the hole straightness, as well as to make the machined surfaces more uniform. Finally, a series of honing steps with progressively finer grit are performed on both sets of holes, further smoothing the surface




Figure 2.4: **Steps in the flute cavity fabrication method.** (a) A set of separate holes is drilled from one end of the aluminum block that will become the 3D cavity. (b) An offset set of holes is drilled from the opposite end of the block. The distance between the centers of opposing holes is less than the diameter of the holes, which ensures that there is an overlap volume inside the block that will form the cavity volume. (c) A reaming step removes surface imperfections created during the hole drilling, and ensures the accuracy of the final hole sizes. (d) A series of honing steps with 3 successively finer grits further smooths the surface while removing burrs.

and removing any burrs. A schematic illustrating this process can be found in Fig. 2.4.

After the machining, the cavity undergoes several cleaning and etching steps. First, the cavity is degreased, then sonicated in a series of organic solvents (Toluene, Acetone, Methanol, Isopropanol, and DI water) to remove leftover oils and particle contaminants. The frequency of the sonicating bath is varied for each solvent to target particles of different size. Following these cleaning steps, the cavity is chemically etched to remove defects and impurities caused by the machining process. We use Transene Aluminum Etchant Type A, a combination of phosphoric acid (H_3PO_4) acid, nitric acid (HNO_3), and acetic acid (CH_3COOH), at 50°C . This mixture targets the aluminum oxide instead of the base alu-

minum [9, 10]. The cavity was placed in the etchant while on a hot plate, and initially etched for 2 hours. Then, the etchant was replaced to ensure a more uniform etch rate. This was followed by four 30 minute etching periods with the etchant replaced again between periods. The total etching time was therefore 4 hours, and the total amount of removed material was roughly 100 μm . This removal reduces the effective surface damage suffered by the cavity, such as defects that would decrease the effective mean-free path of Cooper pairs and lead to a longer London-penetration length and greater participation of the supercurrent with lossy magnetic defects [11]. After the etching, the cavity roughness is inspected. We find that the post-etch roughness is much less than the wavelength of the cavity modes, such that deviations do not significantly affect the mode frequencies or cause unwanted scattering.

2.3 Coupling the Transmon and Multimode Cavity

We start by considering the interaction between the transmon and a single cavity or readout mode. The corresponding circuit is represented in Fig. 2.1, with the transmon and cavity capacitively coupled, as shown in purple. This system can be described by the Jaynes-Cummings Hamiltonian,

$$\hat{H}_{JC} = \omega_c \hat{a}^\dagger \hat{a} + \sum_j \omega_j |j\rangle \langle j| + g_j (\hat{a} + \hat{a}^\dagger) (|j\rangle \langle j+1| + |j+1\rangle \langle j|), \quad (2.9)$$

where ω_c is the bare cavity mode frequency, ω_j is the energy of each transmon level and g is the coupling strength between the cavity and transmon. In our system, $g \sim \vec{d} \cdot \vec{E}$ comes from the interaction between the dipole moment of the qubit antenna pad and the electric field of the cavity. We can alternatively express the Hamiltonian as

$$\hat{H}_{JC} = \omega_c \hat{a}^\dagger \hat{a} + \omega_q \hat{b}^\dagger \hat{b} + \alpha \hat{b}^\dagger \hat{b} (\hat{b}^\dagger \hat{b} - 1) + g (\hat{a} + \hat{a}^\dagger) (\hat{b} + \hat{b}^\dagger), \quad (2.10)$$

where now ω_q is the bare $|g\rangle - |e\rangle$ qubit frequency, α is the anharmonicity, and \hat{b} is the annihilation operator acting on the transmon. We operate in the dispersive limit, $|g/\Delta| \ll 1$, where $\Delta = \omega_q - \omega_c$. In this regime, we can eliminate the interaction term $g(\hat{a} + \hat{a}^\dagger)(\hat{b} + \hat{b}^\dagger)$ through a transformation $\hat{H}' = \hat{D}\hat{H}\hat{D}^\dagger$, where $\hat{D} = \exp(\hat{S} - \hat{S}^\dagger)$ with $\hat{S} = (g/\Delta)\hat{a}\hat{b}^\dagger$ [3]. The resulting Hamiltonian to second leading order becomes

$$\hat{H} = \left(\omega_c + \frac{g^2}{\Delta}\right)\hat{a}^\dagger\hat{a} + \left(\omega_{ge} - \frac{g^2}{\Delta}\right)\hat{b}^\dagger\hat{b} + \alpha\hat{b}^\dagger\hat{b}(\hat{b}^\dagger\hat{b} - 1) + \chi\hat{a}^\dagger\hat{a}\hat{b}^\dagger\hat{b} + \frac{k}{2}\hat{a}^\dagger\hat{a}(\hat{a}^\dagger\hat{a} - 1), \quad (2.11)$$

where k is the self-Kerr nonlinearity of the cavity inherited from the transmon and $\chi = \frac{g^2\alpha}{\Delta(\Delta+\alpha)}$ is the dispersive shift between the cavity and transmon. Generalizing to multiple modes, our final Hamiltonian becomes

$$\begin{aligned} \hat{H} = & \omega_{ge}\hat{b}^\dagger\hat{b} + \frac{\alpha}{2}\hat{b}^\dagger\hat{b}(\hat{b}^\dagger\hat{b} - 1) \\ & + \sum_j \omega_j\hat{a}_j^\dagger\hat{a}_j + \chi\hat{a}_j^\dagger\hat{a}_j\hat{b}^\dagger\hat{b} + \frac{k_{jj}}{2}\hat{a}_j^\dagger\hat{a}_j(\hat{a}_j^\dagger\hat{a}_j - 1) + \sum_{j' \neq j} k_{jj'}\hat{a}_j^\dagger\hat{a}_{j'}\hat{a}_j^\dagger\hat{a}_j, \end{aligned} \quad (2.12)$$

where ω_j is the frequency of cavity mode j and $k_{j'j}$ is the cross-Kerr interaction between two cavity modes inherited from their shared coupling to the transmon. Additionally, we note that ω_{ge} and ω_j are corrected from the bare qubit and cavity frequencies by $\mp \sum_j g_j^2/\Delta_j$, but for practical purposes, the corrected values are the only ones accessible in the system, so we treat them as the transmon and cavity frequencies of interest. A schematic illustrating our final 3D transmon-cavity system is shown in Fig. 2.5(a). The two pads connected to the transmon couple it to the readout and storage cavities, as the transmon chip is inserted on one end of both cavities such that the pads overlap with the cavity volumes. An image of a multimode flute cavity is also presented, with mode frequencies visible as the peaks in the transmission spectrum in Fig. 2.5(b). The cavity volume is curved back a few times to allow it to fit in a more compact space, rather than keeping it in a single line, which would be impossible to contain inside our cryogenic hardware. We performed simulations




Figure 2.5: **Schematic of the coupled 3D transmon and cavity system.** (a) The transmon chip (blue) is inserted on one end of the readout cavity (green) and storage cavity (red), and couples to those cavities through their electric fields. An image of an aluminum multimode flute cavity is also shown. The storage cavity wraps around to reduce the amount of space and material required to contain it. The curvature is chosen through simulation by ensuring that the cavity does not interfere with itself and its performance is comparable to a purely linear one. (b) Example frequencies for a multimode flute cavity in such a system are visible as the peaks in the transmission spectrum. The modes are evenly spaced using the tapering techniques previously discussed.

using Ansys HFSS to ensure that this modification did not adversely alter the cavity fields or behavior. For the device used to perform the experiments presented throughout this work, the readout and storage cavities were CNC machined out of the same high-purity 5N5 (99.9995%) aluminum block.

2.4 Auxiliary Setup

2.4.1 Cryogenic Hardware and Wiring

The multimode cavity device is heat sunk to an OFHC copper plate connected to the base stage of a Bluefors LD-400 dilution refrigerator (8–11 mK). The sample is surrounded by a

can containing two layers of μ -metal shielding, with the inside of the inner layer connected to a can made out of copper shim that is painted on the inside with Berkeley black and attached to the copper can lid. A schematic of the cryogenic setup and device wiring is shown in Fig. 2.6, with the base stage highlighted by the yellow shaded region. Each drive line that accesses the device has on it a handmade Eccosorb CR-110 IR low-pass filter. The output from the device is first amplified by a Josephson Parametric Amplifier (JPA) [12, 13]. The JPA is a quantum noise limited amplifier [14–17] which we operate in four-wave mixing (charge pump) mode that gives us ~ 20 dB gain in a bandwidth of 10 MHz [18]. For controlling the JPA, we use a Quantum Machines OPX Ultra-Fast Quantum Controller to implement the JPA pump and an SRS SIM928 voltage/current source to tune the JPA flux.

2.4.2 Control Electronics

The control pulses are digitally synthesized using a four-channel, 64 GSa/s arbitrary waveform generator (Keysight M8195A), resulting in an effective sampling rate of 16 GSa/s for each channel. The channel distribution is as follows: one for strong qubit drives, one for strong direct cavity drives, one for weak qubit and cavity drives, and one for sideband drives. The channel allocated for weak qubit and cavity drives allows for more dynamic range in our drive controls, as the M8195A has 8 bits of amplitude control. All drive channels are input through the readout cavity, except for the strong cavity drive line, which goes to a port that couples directly to the multimode cavity. The combined signals are sent to the device after being attenuated at each of the thermal stages, as shown in Fig. 2.6. The transmitted signal coming out from the readout resonator passes through three cryogenic circulators (thermalized at the base stage) and is amplified using a JPA and an LNF 4–8 GHz HEMT amplifier (anchored at 4K). Between the circulators and the HEMT, the wires are made of NbTi. Outside the fridge, the signal is filtered with a tunable narrow band YIG filter with a bandwidth of 80 MHz and further amplified. The amplitude and phase of the resonator




Figure 2.6: **System wiring diagram.** The schematic, from top to bottom, shows: the control instrumentation, microwave wiring and filtering, and cryogenic setup.

transmission signal are obtained through a homodyne measurement, with the transmitted signal demodulated using an IQ mixer and a local oscillator (Keysight MXG N5183B) at the readout resonator frequency. The homodyne signal is amplified (SRS preamplifier) and recorded using a fast ADC card (Keysight M3102A PXIe 500 MSa/s digitizer).

Chapter 3

Calibration of the System Hamiltonian

With the quantum hardware established in the previous chapter, we now need to operate on the system. The three components of the device are a readout cavity, a multimode storage cavity, and a 3D transmon that bridges the two cavities and couples to them. This chapter will discuss time-domain experiments that can be used to determine the component parameters that are necessary for effective control and measurement. After obtaining values for these parameters, we describe some additional useful experiments and control techniques.

3.1 Readout Cavity

To accurately determine the frequency of the readout cavity, we perform a resonator spectroscopy experiment. This consists of varying the frequency of a drive tone and measuring the response of the readout, which will produce a peak centered at the readout frequency. In the low-power limit, the quality factor of the resonator can be determined from the linewidth κ of the peak, as $Q = \omega/\kappa$. After calibrating a transmon qubit $|g\rangle - |e\rangle$ π pulse, which we will describe in the next section, this experiment can be repeated after first exciting the transmon to the $|e\rangle$ state. The center of the resonator peak will change, corresponding to the dispersive shift χ_r between the qubit and readout.

We can optimize our readout frequency with several considerations. First, perhaps the most important condition is that we want the measured voltage values of $|g\rangle$ and $|e\rangle$ to be as distinct as possible. This minimizes the overlap between the two and the probability of confusing them, thereby improving our readout fidelity. For applications that utilize the transmon second excited state $|f\rangle$ level, it may also be useful to be able to distinguish that value from $|g\rangle$ and $|e\rangle$, although in some cases we can map $|f\rangle$ to either $|g\rangle$ or $|e\rangle$ and avoid




Figure 3.1: **Resonator spectroscopy and readout-storage cross-Kerr.** (a) We sweep the readout frequency and measure the response in I (solid) and Q (dashed), the two quadratures of the measurement result. We perform the experiment when the transmon is in the ground state $|g\rangle$ (red), and when it is in the first excited state $|e\rangle$ (blue). In the above, we have rotated the quadratures such that I roughly corresponds to phase, and Q to magnitude. From the frequency difference between the centers of the $|g\rangle$ and $|e\rangle$ Q peaks (dashed vertical lines), we calculate the dispersive shift of the readout $\chi_r = -1.2$ MHz. Additionally, from the $|g\rangle$ peak (dashed, red), we calculate the quality factor of our readout $Q = 12000$. We choose the readout frequency at the solid black line, as the response changes minimally as the strength of the cross-Kerr interaction between the readout and storage cavities (which is negative) increases. (b) Effect of the cross-Kerr between the readout and storage cavities while measuring $|g\rangle$. As the population $|\alpha|$ of the storage cavity increases, the shift in readout frequency will also increase. From (a), we see that the voltage value corresponding to $|g\rangle$ changes as the readout frequency moves left of the black line, which we observe as a shift from the ideal value of 0.

the need for that, as we will see later in this chapter. The next consideration is that we want to choose a frequency where the readout response is minimally sensitive to the storage

cavity population. As we can see from Eqn. (2.12), due to the coupling with the transmon, there will be a cross-Kerr term between the readout and storage cavities that causes cavity population to shift the readout frequency. To minimize the effects of this, we choose a readout frequency around which the Resonator Spectroscopy responses of $|g\rangle$ and $|e\rangle$ are as flat as possible, as shown in Fig. 3.1. Fortunately, if the storage cavity frequencies are all on the same side of the readout cavity frequency (all lower or all higher), the cross-Kerr will all be in the same direction, so we can place the readout frequency on one end of the relatively flat band. We can also reduce the effective magnitude of the cross-Kerr by lowering the readout power, although this will also likely decrease the base readout fidelity through reducing the contrast between $|g\rangle$ and $|e\rangle$. Depending on the experiment we wish to perform, we give these considerations different weights to achieve the best possible readout parameters.

3.2 Transmon Qubit

We start by roughly estimating the transmon qubit ground to first excited state ($|g\rangle - |e\rangle$) frequency, ω_{ge} , by performing a Pulse Probe Spectroscopy experiment. This is done by sweeping the frequency of a drive tone around the expected value of ω_{ge} and reading out. When the qubit is excited by the tone driving on or near resonance, the readout response will change due to their dispersive interaction. This change is a measurable effect, and will generate a peak around ω_{ge} . The pulse sequence and an example data set are shown in Fig. 3.2(a). We observe a negative peak, although in general the peak can be either positive or negative depending on the $|g\rangle$ and $|e\rangle$ voltage values.

After obtaining a rough estimate of ω_{ge} , we can obtain qubit π pulse and $\pi/2$ pulse parameters with a Rabi experiment. Here, depending on the hardware setup and which is easier, we either vary the length of a qubit drive tone or vary its amplitude before reading out. We expect to observe Rabi oscillations as the qubit is driven between its $|g\rangle$ and $|e\rangle$

states. An example of a Length Rabi experiment (where the length of the drive is varied) is presented in Fig. 3.2(b), with the y-axis already calibrated such that $|g\rangle$ corresponds to 0 and $|e\rangle$ corresponds to 1.

With a $\pi/2$ pulse in hand, we can now measure more the qubit frequency more precisely with a Ramsey experiment. This involves a qubit $\pi/2$ pulse, followed by a variable wait time τ , followed by another $\pi/2$ pulse and readout. With this procedure, if the pulses are exactly on resonance and share the same phase, the qubit will end in $|e\rangle$ regardless of the value of τ . However, if the estimated qubit frequency is not exact (which is very likely the case), the excited part of the qubit state will precess and accumulate some phase before the second $\pi/2$ pulse, resulting in a final state that differs depending on the accumulated phase ϕ . If, for example, $\phi = \pi$, the second $\pi/2$ pulse will put the qubit in $|g\rangle$ instead of $|e\rangle$. Since ϕ depends on τ , as we vary τ , we observe an oscillation at the difference frequency between the frequency of the $\pi/2$ pulses and ω_{ge} . However, if the $\pi/2$ pulses have the same phase, we cannot distinguish whether we are above or below ω_{ge} , since either sign would produce the same frequency oscillation. We avoid this by advancing the phase ϕ_{Ram} of the second $\pi/2$ pulse corresponding to a chosen Ramsey frequency ω_{Ram} , such that $\phi_{\text{Ram}} = \tau\omega_{\text{Ram}}$. We note that with this phase advancing, we can distinguish difference frequencies up to modulo ω_{Ram} , while difference frequencies greater than ω_{Ram} will be aliased. The Ramsey experiment is represented in Fig. 3.2. Like with the Rabi experiment, the y-axis is calibrated to values between 0 and 1. The target Ramsey frequency was $\omega_{\text{Ram}} = 2\pi \times 50$ kHz, which ideally corresponds to 2.5 oscillations over the measurement length of 50 μs . Since the observed oscillation frequency is slightly off, we calculate the frequency difference and use that to more accurately estimate the real qubit frequency. From this measurement, we also determine the decay-limited phase coherence $T_2^* = 160\mu\text{s}$.

After obtaining a more accurate qubit frequency, it is often useful to repeat the Rabi experiment to calibrate more accurate qubit pulses, followed by repeating the Ramsey ex-




Figure 3.2: **Calibration of a transmon qubit.** (a) Pulse sequence and example data for a Pulse Probe Spectroscopy experiment. We use this experiment to obtain a rough estimate of the transmon qubit frequency, ω_{ge} . In this case, the pulse sequence is just a drive with varying frequency around the expected value of ω_{ge} followed by measuring. The fit to a Gaussian peak is shown in red, with the center frequency marked by the dashed vertical line. (b) A Rabi experiment with varying drive pulse length. Driving at the estimated qubit frequency ω_{ge} , we observe oscillations that we fit (red) to determine the times for a $\pi_{ge}/2$ and π_{ge} pulse, marked by the vertical left and right dashed lines. The drive pulse shape has an approximate Gaussian envelope that goes $\pm 2\sigma$ around the mean, so the total pulse length is 4σ . (c) A Ramsey experiment, which we use to obtain a more accurate estimate of ω_{ge} . The phase of the second $\pi/2$ pulse is advanced corresponding to $\phi = \tau\omega_{\text{Ram}}$. The oscillation gives us a correction to ω_{ge} and an estimate of the decay-limited phase coherence T_2^* .

periment again and iterating until reaching a stable point. In this process, to avoid aliasing in the Ramsey experiment, we initially sweep a shorter range of τ values with smaller time steps to allow for a larger ω_{Ram} . In further iterations, we increase the time step and lower ω_{Ram} to obtain a more precise value for ω_{ge} .

We now have a well-calibrated qubit. This allows us to perform a Histogram measurement and determine our readout fidelity. Here, we measure the distribution of many singleshot I and Q voltage values, first with the qubit in $|g\rangle$, and then in $|e\rangle$, as shown in Fig. 3.3. We then calculate a singleshot readout fidelity based on the overlap between the two distributions. This fidelity serves as a metric that we can use to optimize our readout parameters, such as the length and amplitude of the readout pulse. Furthermore, this measurement allows us to determine an angle by which to rotate our I and Q values during data processing such that all information about the qubit state is transferred entirely into one of those quadratures. We then only have to observe that channel, instead of both, as shown in Fig. 3.3(b) and (c), where all information about the state has been rotated into I. In addition to our readout fidelity of 0.82, another useful result produced by this measurement is the voltage values that correspond to $|g\rangle$ and $|e\rangle$, which allows us to linearly transform our measured voltage values into the range between 0 (for $|g\rangle$) and 1 (for $|e\rangle$).

To judge the performance of our transmon qubit, we measure its coherences by performing T_1 and Ramsey Echo experiments. The pulse sequences and example data are presented in Fig. 3.4(a) and (b). For the T_1 experiment, we excite the qubit into the $|e\rangle$ state using a π pulse and observe an exponential decay by varying the wait time before we readout. The Ramsey Echo is very similar to the Ramsey experiment, but with the addition of a qubit π pulse in the middle of the wait time between the two $\pi/2$ pulses. Any slow noise effects that occur before the π pulse will have the opposite effect after the π pulse, thereby cancelling out if the wait times are equal. In our case, the π pulse cancels any noise originating from the $\pi/2$ pulses being slightly inexact, as well as any other slow noise that could occur during the




Figure 3.3: **Readout fidelity histogram experiment.** (a) Data showing each measured single shot I/Q voltage corresponding to the transmon in $|g\rangle$ (red), $|e\rangle$ (blue), and $|f\rangle$ (yellow). The resulting blobs are fit to Gaussians with means and standard deviations shown in the right plot. For these readout parameters, $|g\rangle$ and $|e\rangle$ are separated, but $|e\rangle$ and $|f\rangle$ are highly overlapping and thus harder to distinguish. (b) Rotated version of the raw data such that the line between the means of the Gaussian fits of $|g\rangle$ and $|e\rangle$ is cast to the I axis. This ensures that the information that distinguishes those states is contained entirely in I, so we only have to look at that quadrature during our measurements. (c) Histograms showing the binned shots. As expected, the Q histograms are entirely overlapping, reflecting the fact that we have rotated all state information to the I quadrature. (d) Readout fidelity obtained from the overlap between the histograms. This is calculated by subtracting the cumulative distribution function (CDF) of the $|g\rangle$ histogram from the CDF of the $|e\rangle$ histogram. This data gives a final readout fidelity of 0.82 for the I quadrature.




Figure 3.4: **Qubit coherence experiments.** We present pulse sequences and example data. (a) T_1 experiment that measures the decay lifetime of the qubit. Red dashed lines mark $1/e$ and the fit (red) $T_1 = 118\mu\text{s}$. (b) Echo experiment that measures the qubit $T_2 = 189\mu\text{s}$. Comparing to a Ramsey, we have $T_2 > T_2^*$ due to no longer being T_1 -limited. (c) Qubit temperature experiment. By using the transmon $|f\rangle$ level, we fit to the initial $|e\rangle$ population (blue) compared to $|g\rangle$ (black). The contrast is 0.018, corresponding to a qubit temperature of 54 mK.

sequence. We measure the coherences of our system to be $T_1 = 108 \pm 10\mu s$, $T_2^* = 140 \pm 13\mu s$, and $T_2 = 180 \pm 14\mu s$.

Now that the $|g\rangle - |e\rangle$ space is well-characterized, we can perform similar experiments on the next excited state of the transmon, the $|f\rangle$ level and the $|e\rangle - |f\rangle$ space. The measurement sequences are very similar to the ones previously discussed, but with the addition of an initial $|g\rangle - |e\rangle$ π pulse to start the qubit in $|e\rangle$, as well as final pulses that cast the result to the $|g\rangle$ and $|e\rangle$ subspace, if desired. Starting with an EF Pulse Probe Spectroscopy experiment, we roughly measure the frequency of the $|e\rangle - |f\rangle$ transition, and therefore the anharmonicity $\alpha = \omega_{ef} - \omega_{ge}$ of the transmon. In this experiment, the final transmon state will be in a superposition of $|e\rangle$ and $|f\rangle$. If distinguishing between those two states is difficult, we can add a π_{ge} right before the measurement to cast $|e\rangle$ to $|g\rangle$, so that we instead need to differentiate between $|g\rangle$ and $|f\rangle$.

To continue calibrating the transmon, we perform and iterate on EF Rabi and EF Ramsey experiments to get α more precisely and calibrate π_{ef} and $\pi/2_{ef}$ pulses. With these pulses, we can measure the coherences of the $|f\rangle$ level. Additionally, with a calibrated π_{ef} pulse, we can now fully cast a superposition of $|e\rangle$ and $|f\rangle$ to a superposition of $|g\rangle$ and $|e\rangle$, by applying a π_{ge} pulse followed by a π_{ef} pulse at the end of experiments (right before measuring). This can be useful if some pair of $|g\rangle$, $|e\rangle$, and $|f\rangle$ are difficult to distinguish, and allows us to perform these multilevel transmon experiments while only having to optimize a single readout space, such as $|g\rangle$ and $|e\rangle$. Alternatively, knowing π_{ef} pulse parameters allows us to add the $|f\rangle$ level to the Histogram experiment and simultaneously optimize the readout to distinguish all 3 of $|g\rangle$, $|e\rangle$, and $|f\rangle$.

Furthermore, with the EF transmon pulses, we are able to perform a Qubit Temperature experiment. This allows us to determine how excited the qubit is when we are not operating on it. This consists of two successive EF Rabi experiments: the first (E1) a normal EF Rabi, and the second (E2) without an initial π_{ge} pulse, as shown in Fig. 3.4(c). The contrast of the




Figure 3.5: **Calibration of transmon $|e\rangle - |f\rangle$ levels.** The primary differences between these experiments and those for the $|g\rangle - |e\rangle$ space are the addition of an initial π_{ge} pulse to start the transmon in $|e\rangle$, and optional additional π_{ge} and π_{ef} pulses right before measuring that cast the transmon back to the $|g\rangle - |e\rangle$ space. (a) Pulse sequence and example data for an EF Pulse Probe Spectroscopy experiment. We use this experiment to obtain a rough estimate of the transmon qubit frequency, ω_{ef} , marked by the dashed vertical line. (b) An EF Rabi experiment with varying pulse length. We observe oscillations that we fit (red) to determine the drive times for a $\pi_{ef}/2$ and π_{ef} pulse, marked by the vertical left and right dashed lines. Like in the GE Rabi, we use pulses with approximately Gaussian envelopes with total length 4σ . (c) An EF Ramsey experiment, which we use to obtain a more accurate estimate of ω_{ef} . We also measure the coherence of the $|f\rangle$ level to be $T_2^* = 70\mu$ s.

Parameter	Hamiltonian/Liouvillian Term	Quantity	Value
Transmon frequency	$\omega_q \hat{b}^\dagger \hat{b}$	$\omega_q/(2\pi)$	4.97 GHz
Transmon anharmonicity	$\frac{\alpha}{2}(\hat{b}^\dagger \hat{b})(\hat{b}^\dagger \hat{b} - 1)$	$\alpha/(2\pi)$	-143 MHz
Readout frequency	$\omega_r a_r^\dagger a_r$	$\omega_r/(2\pi)$	7.79 GHz
Readout dispersive shift	$\chi_r a_r^\dagger a_r e\rangle \langle e $	$\chi_r/(2\pi)$	-1.2 MHz
Transmon $ e\rangle \rightarrow g\rangle$ relaxation	$\frac{1}{T_1^q} (1 + \bar{n}) \mathcal{D}[g\rangle \langle e]$	T_1^q	$108 \pm 10 \mu\text{s}$
Transmon $ g\rangle - e\rangle$ dephasing	$(\frac{1}{T_2^q} - \frac{1}{2T_1^q}) \mathcal{D}[e\rangle \langle e]$	T_2^q	$180 \pm 14 \mu\text{s}$
Readout linewidth	$\kappa_r \mathcal{D}[a_r]$	$\kappa_r/(2\pi)$	0.52 MHz
Transmon thermal population	$\frac{\bar{n}}{T_1^q} \mathcal{D}[e\rangle \langle g]$	\bar{n}	$1.8 \pm 0.4 \%$

Table 3.1: **Transmon and readout system parameters.**

first experiment corresponds to the initial population of $|g\rangle$, while the second corresponds to the initial population of $|e\rangle$. The qubit temperature is given by

$$k_B T_{ge} = \frac{\hbar \omega_{ge}}{\log(R)}, \quad (3.1)$$

where R is the ratio of the E1 to E2 oscillation magnitudes. From the temperature, we calculate the base population of the $|e\rangle$ state as:

$$P_e = \left(\exp\left(\frac{\hbar \omega_{ge}}{k_B T_{ge}}\right) + 1 \right)^{-1}. \quad (3.2)$$

In our system, we find a qubit temperature of 54 ± 3 mK, corresponding to an initial $|e\rangle$ state population of $1.8 \pm 0.4\%$.

We could continue these calibrations to higher transmon levels, such as the next (third) excited state, the $|h\rangle$ level. In principle, we could go further to excited states that are as high as we desire until limited by hardware or coherence constraints, as we expect the lifetimes of the n th excited state to be n times shorter than those of $|e\rangle$. A summary of our system's transmon and readout cavity parameters can be found in Table 3.1.

For high fidelity operations in the $|g\rangle-|e\rangle$ subspace, it is important to minimize populating the $|f\rangle$ level off-resonantly. It is therefore helpful to consider the frequency components of our drive pulse, and to choose a pulse shape that minimizes any unwanted components. For this reason, we avoid using a constant amplitude envelope (square) drive tone, as the Fourier transform of a square pulse will be the absolute value of a sinc function, $|\sin(x)/(x)|$, which has an infinite number of peaks and significant off-resonance features. We instead prefer to use a drive tone with a Gaussian envelope, as its Fourier transform will be another Gaussian with exponentially suppressed off-resonant contributions. However, an exact Gaussian would be a pulse of infinite length, so we instead use an approximate Gaussian shape with the amplitude envelope

$$a(t) = a_0 \exp\left(-\frac{(t-2\sigma)^2}{2\sigma^2}\right) \quad (3.3)$$

where $0 \leq t \leq 4\sigma$ is the time during which the drive pulse is on and a_0 is the maximum drive amplitude. This pulse goes $\pm 2\sigma$ around the mean and thus captures 95% of the exact Gaussian's integrated amplitude. We choose $\pm 2\sigma$ to gain most of the benefits of the Gaussian pulse envelope while also balancing its duration. One of the trade-offs of having a non-square pulse envelope is that drives of the same magnitude will necessarily require more time, and we do not want the drive to take so long to where additional decoherence effects outweigh the benefits of the envelope.

Due to the Gaussian envelope only being approximate, it is possible that the drive component at an unwanted frequency like the $|e\rangle-|f\rangle$ transition will not be exponentially suppressed, and in the worst case may be at a local maximum of the the Fourier spectrum. This is because the small jump from 0 to the start of the approximate Gaussian will still produce features similar to the sinc function. To mitigate this and further improve the fidelity of our transmon pulses, we can implement Derivative Removal by Adiabatic Gate (DRAG) pulses [19–21]. These pulses operate by adding the derivative of the primary Gaussian pulse




Figure 3.6: **DRAG pulse calibration and implementation.** (a) Calibration experiment used to determine the DRAG coefficient β . We measure the effect of changing β on two individual qubit experiments: the first, a π_y pulse followed by a $\pi_x/2$ pulse (purple), and the second, a π_x pulse followed by a $\pi_y/2$ pulse (black). The two experiments have opposite error syndrome and thus opposite slope, and intersect at a population of 0.5, the expected value for a π pulse followed by a $\pi/2$ pulse. The optimal value obtained from the intersection is $\beta = -1.15$. (b) Fourier transform comparison for the DRAG pulse (orange) and the original pulse (blue). At the qubit frequency ω_{ge} (blue right dashed line), the two pulses have the same magnitude. However, at ω_{ef} (left right dashed line), the magnitude of the DRAG pulse is roughly an order of magnitude less than the original pulse, demonstrating that the DRAG pulse is accomplishing the goal of reducing unwanted driving of the $|f\rangle$ level.

to its orthogonal quadrature, and help correct for phase errors induced by higher transmon levels. Specifically, for our qubit drive pulses, we modulate the envelope in Eqn. (3.3) by $\cos(\omega_{ge}t + \phi)$ for pulse phase ϕ . Adding the orthogonal derivative part gives the total DRAG pulse

$$\begin{aligned} \Omega_{\text{DRAG}}(t) = & a_0 \exp\left(-\frac{(t-2\sigma)^2}{2\sigma^2}\right) \cos(\omega_{ge}t + \phi) \\ & + \beta a_0 \frac{-(t-2\sigma)}{\sigma^2} \exp\left(-\frac{(t-2\sigma)^2}{2\sigma^2}\right) \sin(\omega_{ge}t + \phi), \end{aligned} \quad (3.4)$$

where β is a constant DRAG coefficient. We calibrate β experimentally following the procedure outlined in [21]. We sweep the value of β while performing two experiments: the first

with a qubit π_x pulse followed by a $\pi_y/2$, and the second with a qubit π_y followed by a $\pi_x/2$, where each pulse is of the form in Eqn. (3.4) and there is a 90° phase difference between the x and y pulses. As β changes, the effective error will change proportionally. Since we have chosen two sets of pulses with opposite error syndrome, these experiments will produce two lines that intersect at the optimal value of β . Experimental data obtained from this procedure is shown in Fig. 3.6(a). As expected, the fitted lines intersect near a qubit population of 0.5, the ideal expected value corresponding to a π pulse followed by a $\pi/2$ pulse. In Fig. 3.6(b), we see the effect of the DRAG pulse compared to the original one by looking at the Fourier transforms of the pulses. At the qubit frequency ω_{ge} (blue right dashed line), the magnitude of the Fourier transform is the same for the two pulses. More importantly, at ω_{ef} (red left dashed line), the magnitude of the DRAG pulse is roughly an order of magnitude less than the original, demonstrating that the DRAG pulse is serving its intended purpose. The above discussion also applies when we are operating in the $|e\rangle - |f\rangle$ subspace and want to avoid driving the $|g\rangle - |e\rangle$ transition. The calibration sequence is similar, except with preparing the qubit in $|e\rangle$ and acting π_{ef} or $\pi_{ef}/2$ pulses rather than π_{ge} or $\pi_{ge}/2$ ones.

3.3 Multimode Cavity

3.3.1 Initial Cavity Parameters

One method for roughly measuring the frequency ω_c of a cavity mode is by using a similar approach to the Pulse Probe Spectroscopy, but with a drive sweeping around the cavity frequency instead of the qubit frequency. Due to the cross-Kerr between the readout and storage cavity mode, populating the storage cavity mode will change the frequency of the readout and thereby alter its response at the original readout frequency. Alternatively, with a well calibrated qubit, we can perform a Resolved Cavity Frequency Spectroscopy experiment,

where the drive around the expected cavity frequency is followed by a narrow-bandwidth, resolved qubit π pulse before the readout measurement. Due to the dispersive interaction $\chi a^\dagger ab^\dagger b$, populating the cavity will shift the qubit frequency and take the resolved π pulse off resonance. This experiment produces an observable peak, as the qubit ends in $|e\rangle$ for drive frequencies far from ω_c , and ends at least partially in $|g\rangle$ for drive frequencies near ω_c . This is depicted in Fig. 3.7(a). In our system, we prefer to use this latter method, as the dispersive interaction (\sim MHz) is a much larger effect than the cross-Kerr (\sim kHz). For best results and the most easily observed peak, the π pulse should be slow enough such that its bandwidth is $< \chi$.

With an estimate of ω_c , we proceed to estimating the dispersive shift χ and the strength of our cavity drive by perform a Photon Number Resolved Spectroscopy experiment. This experiment is shown in Fig. 3.7(b). We first apply a cavity drive and then sweep the frequency of a slow, narrow bandwidth qubit π pulse. The π pulse response will change at frequencies $\omega_{ge} + n\chi$, when the final qubit state will be $|e\rangle$ instead of $|g\rangle$. This will produce photon number resolved peaks, allowing us to estimate χ from the frequency difference between the peak centers, particularly the 0 (vacuum) and 1 photon peak [22]. A rough estimate of the cavity self-Kerr can be obtained in this manner by calculating the difference frequency between the 2 and 1 photon peak and the 1 and 0 photon peak, although the difference can be small (\sim kHz) and difficult to resolve with this approach. Additionally, simply driving the cavity in this manner will produce a coherent state α where the distribution of the cavity photon populations (peak heights) follows a Poisson distribution [22, 23]. By fitting to this distribution $P(\alpha, n) = |\alpha|^{2n} e^{-|\alpha|^2} / n!$, we measure the magnitude $|\alpha|$ of the coherent state and thus our drive strength. Even if χ is small and the photon number peaks are not well resolved, we can still obtain an estimate of χ by fitting to the peak shape [24].

Analogously to the case of the transmon, we measure ω_c more precisely by performing a Cavity Ramsey experiment. Much like the qubit Ramsey, here we apply a small displacement




Figure 3.7: **Initial cavity calibrations.** We present pulse sequences, with cavity pulses in orange and qubit pulses in black, as well as example data. (a) We roughly estimate the cavity frequency ω_c through Resolved Cavity Frequency Spectroscopy. The π pulse should be slow relative to the qubit-cavity dispersive shift χ , so that the frequency change caused by that interaction produces an observable peak from taking the qubit off resonance. (b) Photon Number Resolved Qubit Spectroscopy populates the cavity, allowing us to see individual photon number peaks when sweeping the frequency of the slow π_{ge} probe pulse. We fit the coherent state magnitude $|\alpha|$ by comparing to a Poisson distribution as marked by the blue dashed line. For this state, $|\alpha| = 1.01$. We also estimate $\chi = -1.09$ MHz as the distance between peak centers. (c) The Cavity Ramsey experiment lets us determine ω_c more accurately. The displacements are sufficiently small so that primarily only the $|0\rangle$ and $|1\rangle$ states are populated. The resulting oscillation contrast will be lower, but this will approximate a two-level system, allowing us to perform the Ramsey.

α to the cavity, wait for a variable time τ , and apply a small displacement with the same magnitude but opposite phase $-\alpha$ to the cavity before adding a resolved π pulse at the cavity 0 photon peak (ω_{ge}) and measuring, as shown in Fig. 3.7(c). We choose a small value of α such that the majority of the cavity population remains in the 0 and 1 photon subspace. Although this reduces the maximum possible contrast as probed by the π pulse, it ensures that the cavity approximates a two-level system, despite being a bosonic oscillator. Similarly to the qubit Ramsey, if the initial displacement pulse is exactly equal to ω_c , then in the frame of the cavity rotating at ω_c , the cavity state will always be restored back to the vacuum state regardless of τ . However, if the displacement is slightly off resonance, phase $\phi = \tau\delta$ will accrue, where δ is the difference between ω_c and the frequency of the displacement drive. Varying τ will produce an oscillation with frequency δ , which we fit to obtain ω_c precisely. Like for the qubit, in this process we advance the phase of the second cavity displacement in the experiment corresponding to a chosen $\omega_{\text{Ram},c}$. This removes any ambiguity in the sign of δ , as the oscillation will now have frequency $\omega_{\text{Ram},c} + \delta$. After obtaining this more precise value for ω_c , it is useful (again like in the case of the qubit) to iterate the Photon Number Resolved Qubit Spectroscopy experiment and this Cavity Ramsey experiment to improve the precision of the parameters until reaching a stable point.

To measure the higher-order self-Kerr k of the cavity, we repeat the Cavity Ramsey experiment for many displacement amplitudes. As the amplitudes increase, the resulting oscillations will be increasingly nonlinear as a result of the larger photon number cavity states becoming populated and introducing additional frequencies, and in particular those set by the self-Kerr [25, 26]. The final population of the qubit can be fit to the expression:

$$P_0(\alpha, t) = |\exp(-|\alpha|^2) \sum_n \frac{1}{n!} |\alpha|^{2n} \exp(-itn(\omega_c - kn/2))|^2 \quad (3.5)$$

which we use to extract the best fit value of k [27]. This calibration measurement is shown




Figure 3.8: **Cavity self-Kerr calibration.** The spectrum is obtained by performing Cavity Ramsey experiments with different displacement magnitudes $|\alpha|$ and wait times τ . As $|\alpha|$ and τ increase, the comb-like structure becomes increasingly nonlinear as higher cavity photon levels become significantly populated and introduce additional frequencies and thus phase accrual rates. These effects are reflective of the self-Kerr k of the cavity, and allow us to fit the spectrum and determine $k = 2\pi \times -2.9$ kHz.

in Fig. 3.8. As $|\alpha|$ increases and higher photon numbers of the cavity become populated, the wait times τ at which the cavity mostly returns to $|0\rangle$ have increasingly narrow ranges, producing the observed comb-like spectrum. This narrowing and the tilting of the fingers of the comb as τ increases reflect the effect of the self-Kerr k . For larger τ , the different photon numbers of the cavity have more time for their accrued phases to become disparate. This also causes the observed effect of the maximum amplitude of the fingers decreasing, as the final displacement can no longer fully restore the state to $|0\rangle$. From this spectrum, we can see that the values of $|\alpha|$ that work best for using the Cavity Ramsey to accurately determine ω_c is roughly $|\alpha| \lesssim 0.4$, as the horizontal slice remains approximately linear in that region.

We determine the base population of our cavity mode through a Cavity Temperature experiment. Here, we perform two successive qubit Rabi experiments with resolved drive

tones: the first on resonance with ω_{ge} , and the second at $\omega_{ge} + \chi$. As with the other cavity experiments, the drive is sufficiently weak with narrow bandwidth ($\ll \chi$) such that when at one of those frequencies, it will not drive the other. The oscillation magnitude of the first of these experiments will correspond to the population of $|0\rangle$, while the second will correspond to the population of $|1\rangle$. We use Eqn. (3.1) to determine the cavity temperature, and the cavity population will be given by

$$P_{cav} = \exp\left(-\frac{\hbar\omega_{cav}}{k_B T_{cav}}\right). \quad (3.6)$$

An example of the two oscillations that make up the cavity temperature experiment is shown in Fig. 3.12(c), and the cavity populations \bar{n}_j for the modes j in our system can be found in Table 3.2.

3.3.2 Sideband Experiments

With knowledge of the cavity and transmon frequencies, we can now calibrate a sideband type interaction between the two. This is a four-wave mixing process that transfers excitations between the transmon and cavity [28–30]. We derive this interaction in our system as follows: the Hamiltonian of the device with readout cavity (\hat{b}), transmon (\hat{c}), and storage cavity (\hat{a}) is:

$$\begin{aligned} \hat{H} = & \omega_{ge}\hat{c}^\dagger\hat{c} + \frac{\alpha}{2}(\hat{c}^\dagger\hat{c})(\hat{c}^\dagger\hat{c} - 1) + \omega_r\hat{b}^\dagger\hat{b} + \omega_c\hat{a}^\dagger\hat{a} + g_c(\hat{a} + \hat{a}^\dagger)(\hat{c} + \hat{c}^\dagger) + g_r(\hat{b} + \hat{b}^\dagger)(\hat{c} + \hat{c}^\dagger) \\ & + \epsilon \cos(\omega_d t)(\hat{b} + \hat{b}^\dagger), \end{aligned} \quad (3.7)$$

where g_c and g_r are the coupling between the transmon and storage cavity and transmon and readout cavity, respectively, and the final term is a drive input through the readout. Excluding the drive term, we can diagonalize the Hamiltonian and go from these bare states

into dressed states:

$$\begin{aligned} \hat{H}_{\text{dressed}} = & \tilde{\omega}_{ge} \hat{c}^\dagger \hat{c} + \tilde{\omega}_r \hat{b}^\dagger \hat{b} + \tilde{\omega}_c \hat{a}^\dagger \hat{a} + \frac{\alpha}{2} \left[\left(\hat{c}^\dagger + \frac{g_c}{\Delta_c} \hat{a}^\dagger + \frac{g_r}{\Delta_r} \hat{b}^\dagger \right) \left(\hat{c} + \frac{g_c}{\Delta_c} \hat{a} + \frac{g_r}{\Delta_r} \hat{b} \right) \right]^2 \\ & + \epsilon \cos(\omega_d t) \left(\hat{b} + \hat{b}^\dagger + \frac{g_r}{\Delta_r} (\hat{c} + \hat{c}^\dagger) \right), \end{aligned} \quad (3.8)$$

where $\Delta_j = \omega_j - \omega_{ge}$ is the difference between the respective cavity frequency and the qubit frequency, and we include terms up to 4th order in g/Δ . We now go into the rotating frame at each of the system components' frequencies, $\hat{U} = e^{-i\hat{H}_0 t}$, with $\hat{H}_0 = \tilde{\omega}_{ge} \hat{c}^\dagger \hat{c} + \tilde{\omega}_r \hat{b}^\dagger \hat{b} + \tilde{\omega}_c \hat{a}^\dagger \hat{a}$.

The resulting Hamiltonian is:

$$\begin{aligned} \hat{H}' = & \frac{\alpha}{2} \left(\left(\hat{c}^\dagger e^{i\omega_{ge} t} + \frac{g_c}{\Delta_c} \hat{a}^\dagger e^{i\omega_c t} + \frac{g_r}{\Delta_r} \hat{b}^\dagger e^{i\omega_r t} \right) \left(\hat{c} e^{-i\omega_{ge} t} + \frac{g_c}{\Delta_c} \hat{a} e^{-i\omega_c t} + \frac{g_r}{\Delta_r} \hat{b} e^{-i\omega_r t} \right) \right)^2 \\ & + \epsilon \cos(\omega_d t) \left(\hat{b} e^{-i\omega_r t} + \hat{b}^\dagger e^{i\omega_r t} + \frac{g_r}{\Delta_r} (\hat{c} e^{-i\omega_{ge} t} + \hat{c}^\dagger e^{i\omega_{ge} t}) \right). \end{aligned} \quad (3.9)$$

Focusing on the drive term, we make a rotating wave approximation

$$\begin{aligned} \hat{H}_d(t) = & \epsilon \cos(\omega_d t) \left(\hat{b} e^{-i\omega_r t} + \hat{b}^\dagger e^{i\omega_r t} + \frac{g_r}{\Delta_r} (\hat{c} e^{-i\omega_{ge} t} + \hat{c}^\dagger e^{i\omega_{ge} t}) \right) \\ & \approx \frac{\epsilon}{2} \left(\hat{b} e^{i\delta_r t} + \frac{g_r}{\Delta_r} \hat{c} e^{i\delta_{ge} t} + \text{c.c.} \right), \end{aligned} \quad (3.10)$$

where $\delta_j = \omega_d - \omega_j$ and we have dropped the terms oscillating quickly at the sum frequencies $\omega_d + \omega_r$ and $\omega_d + \omega_{ge}$. Now, we apply the displacement transformation $\hat{D} = e^{\xi(t)\hat{b} + \eta(t)\hat{c} - \text{c.c.}}$, which will act on the annihilation operators as $\hat{D}\hat{b}\hat{D}^\dagger = \hat{b} + \xi(t)$ and $\hat{D}\hat{c}\hat{D}^\dagger = \hat{c} + \eta(t)$. Under this transformation,

$$\begin{aligned} \hat{H}' & \rightarrow \hat{D}\hat{H}'\hat{D}^\dagger - i\hat{D}\dot{\hat{D}}^\dagger \\ & = (\hat{H}' - \hat{H}_d(t)) \left[\hat{c} + \eta(t), \hat{a}, \hat{b} + \xi(t) \right] + \frac{\epsilon}{2} \left(\hat{b} e^{i\delta_r t} + \frac{g_r}{\Delta_r} \hat{c} e^{i\delta_{ge} t} + \text{c.c.} \right) \\ & \quad - i \left(\dot{\xi} \hat{b}^\dagger - \dot{\xi}^* \hat{b} \right) - i \left(\dot{\eta} \hat{c}^\dagger - \dot{\eta}^* \hat{c} \right), \end{aligned} \quad (3.11)$$

where the brackets indicate that we replace \hat{c} , \hat{a} , and \hat{b} with the terms inside the brackets in the expression for \hat{H}' , excluding $\hat{H}_d(t)$, which is not affected. In other words, we make those substitutions and only include the first line of Eqn. (3.9). By choosing $\xi(t)$ and $\eta(t)$, we can cancel out the drive term $\hat{H}_d(t)$. In particular, using Eqn. (3.11), we accomplish this by choosing:

$$\xi(t) = \frac{\epsilon}{2\delta_r} e^{-i\delta_r t}, \quad \eta(t) = \frac{\epsilon}{2\delta_{ge}} \frac{g_r}{\Delta_r} e^{-i\delta_{ge} t}. \quad (3.12)$$

With these definitions, the transformed Hamiltonian becomes

$$\begin{aligned} \hat{H}'' = \frac{\alpha}{2} & \left(\left((\hat{c}^\dagger + \eta^*(t)) e^{i\omega_{ge} t} + \frac{g_c}{\Delta_c} \hat{a}^\dagger e^{i\omega_c t} + \frac{g_r}{\Delta_r} (\hat{b}^\dagger + \xi^*(t)) e^{i\omega_r t} \right) \right. \\ & \left. \left((\hat{c} + \eta(t)) e^{-i\omega_{ge} t} + \frac{g_c}{\Delta_c} \hat{a} e^{-i\omega_c t} + \frac{g_r}{\Delta_r} (\hat{b} + \xi(t)) e^{-i\omega_r t} \right) \right)^2. \end{aligned} \quad (3.13)$$

We further simplify this expression using the definitions of $\xi(t)$ and $\eta(t)$ to arrive at our final Hamiltonian:

$$\begin{aligned} \hat{H}'' = \frac{\alpha}{2} & \left(\left(\hat{c}^\dagger e^{i\omega_{ge} t} + \frac{g_c}{\Delta_c} \hat{a}^\dagger e^{i\omega_c t} + \frac{g_r}{\Delta_r} \hat{b}^\dagger e^{i\omega_r t} + \frac{\epsilon}{2} \frac{g_r}{\Delta_r} \left(\frac{1}{\delta_r} + \frac{1}{\delta_{ge}} \right) e^{i\omega_d t} \right) \right. \\ & \left. \left(\hat{c} e^{-i\omega_{ge} t} + \frac{g_c}{\Delta_c} \hat{a} e^{-i\omega_c t} + \frac{g_r}{\Delta_r} \hat{b} e^{-i\omega_r t} + \frac{\epsilon}{2} \frac{g_r}{\Delta_r} \left(\frac{1}{\delta_r} + \frac{1}{\delta_{ge}} \right) e^{-i\omega_d t} \right) \right)^2. \end{aligned} \quad (3.14)$$

We can now match terms to see what interactions we can generate and what their expected strengths are. For example, a resonant $|f0\rangle - |g1\rangle$ sideband is possible when we choose $\omega_d = \omega_f - \omega_c$, where $\omega_f = \omega_{ge} + \omega_{ef}$ is the frequency difference between the $|g\rangle$ and $|f\rangle$ levels of the transmon. Furthermore, we can calculate the rate of the interaction by computing the coefficient of the $\hat{a}^\dagger \hat{c}^2$ term. Other similar examples of sideband interactions are $|h0\rangle - |e1\rangle$ and $|e0\rangle - |g2\rangle$.

The following discussion will focus on the $|f0\rangle - |g1\rangle$ sideband, although similar proce-




Figure 3.9: $|f_0\rangle - |g_1\rangle$ **sideband experiments**. Transmon pulses are shown in black/gray and sideband pulses are purple. (a) Sideband Frequency Spectroscopy is used to determine the frequency of the sideband interaction $\omega_{f_0-g_1} \approx 2\omega_c - (2\omega_{ge} + \alpha)$, marked by the dashed vertical line. The pulse sequence has an optional π_{ef} pulse at the end that would cast the final state from the $|g\rangle$ and $|f\rangle$ space back to $|g\rangle$ and $|e\rangle$. (b) Sideband Rabi experiment. The dashed vertical lines indicate the $\pi/2$ time and π time of the oscillation, which prepare $(|0\rangle + |1\rangle)/\sqrt{2}$ and $ket1$ in the cavity, respectively. The sideband dephases the transmon by off-resonantly populating the readout cavity, producing the noticeable oscillation decay despite the relatively short time scale.

dures will work for other sideband types. We experimentally determine the resonant sideband frequency by performing a Sideband Frequency Spectroscopy experiment, as shown in Fig. 3.9(a). The measurement sequence involves initializing the transmon in the $|f\rangle$ state by applying a π_{ge} and π_{ef} pulse, then sweeping the frequency of a drive tone around the expected sideband frequency and reading out. This expected frequency is $\omega_{f_0-g_1} \approx 2\omega_{ge} + \alpha - \omega_c$,

corresponding to the frequency difference from the transmon transitioning from $|g\rangle$ to $|e\rangle$, then $|e\rangle$ to $|f\rangle$, minus a cavity excitation (in our system $2\omega_{ge} + \alpha > \omega_c$). When the tone is on resonance, it will drive the qubit from $|f\rangle$ to $|g\rangle$, generating an observable change. If $|f\rangle$ and $|g\rangle$ are not distinguishable, we can apply an additional π_{ef} pulse after the drive tone so that the final qubit state is instead in a superposition of $|e\rangle$ and $|g\rangle$.

After obtaining the sideband frequency, we perform a Sideband Rabi experiment. Analogously to the Rabi experiment on the transmon, here we first prepare the transmon in $|f\rangle$ before varying the length or amplitude of the sideband drive tone. We observe oscillations as the final state of the qubit moves between $|f\rangle$ and $|g\rangle$ (or $|e\rangle$ and $|g\rangle$ if we apply a π_{ef} pulse after the drive). By fitting to the period of this oscillation, we determine the drive time that corresponds to putting the transmon fully in $|g\rangle$, and thus prepares a single photon in the cavity. This experiment is presented in Fig. 3.9(b). With a well-calibrated sideband, we can easily create Fock states in the cavity. This is a useful tool that expands our experimental capabilities. For example, preparing the cavity Fock state $|1\rangle$ enables us to perform additional experiments that determine our cavity parameters more precisely and measure its coherences, as we will see in the next section. We can prepare higher Fock states by sequentially applying sideband drives of the form $|fn\rangle - |g(n+1)\rangle$. After obtaining a π_{f0-g1} pulse, we then calibrate the $|f1\rangle - |g2\rangle$ transition, and continue similarly to higher Fock states $|n\rangle$. Furthermore, since each of these transitions has a slightly different frequency due to the dispersive interaction between the transmon and the cavity, we can start with a $\pi_{f0-g1}/2$ pulse, then drive the $|1\rangle$ part of the state to $|4\rangle$ without affecting $|0\rangle$. This prepares $(|0\rangle + |4\rangle)/\sqrt{2}$, which is one of the logical states of the binomial code [31, 32, 28, 33] (along with $|2\rangle$, which is also easily prepared with sideband drives) with applications in quantum error correction. These particular states can be generated even more immediately with $|en\rangle - |g(n+2)\rangle$ sidebands instead of $|fn\rangle - |g(n+1)\rangle$, demonstrating the flexibility of this kind of interaction.

3.3.3 Cavity Single Photon Experiments

The experiments in this section utilize the ability to prepare a single photon in the cavity. There are a number of techniques that can accomplish this, such as the $|f0\rangle - |g1\rangle$ sideband described in the previous section, or through the use of photon blockade or optimal control, which will be described in later chapters. Regardless of the method, we assume that we have the ability to initialize the transmon-cavity state $|g1\rangle$.

First, we measure the dispersive shift χ more precisely by performing a χ_{ge} Ramsey experiment. Here, we prepare $|g1\rangle$, then perform a Ramsey measurement on the qubit, as shown in Fig. 3.10(a). We advance the phase of the second $\pi/2$ pulse by $\phi_{ge} = \tau(\omega_{\text{Ram}} + \chi_{ge,\text{est}})$ for wait time τ , where $\chi_{ge,\text{est}}$ is our current estimate for χ_{ge} and ω_{Ram} is a chosen artificial Ramsey frequency. If $\chi_{ge,\text{est}} = \chi_{ge}$, we observe an oscillation between $|g\rangle$ and $|e\rangle$ with frequency ω_{Ram} . Otherwise, any deviation from that frequency will give us a correction to $\chi_{ge,\text{est}}$. Similarly to the Ramsey with just the transmon, we repeat this experiment multiple times for increasingly larger ranges of τ and correspondingly lower ω_{Ram} to measure χ_{ge} accurately. The χ_{ge} values for our multimode cavity modes are shown in Fig. 3.11(a), indexed from 0–8 from low to high frequency for the 9 modes that lie between the qubit ($\omega_{ge} = 4.97$ GHz) and readout cavity ($\omega_r = 7.79$ GHz) frequencies. The frequencies of these modes are given in Table 3.2.

By slightly modifying this procedure, we also calibrate a Parity Measurement. As the name suggests, this sequence detects whether a cavity mode has an even or odd number of photons [27, 34], and has a variety of applications in state preparation, measurement, and error correction [35–37]. In the context of this thesis, it is a crucial component of Wigner tomography [38], which we will expand upon in Chapter 6. The Parity Measurement consists of 2 qubit $\pi/2$ pulses with π relative phase separated by a wait time $\tau_{\text{parity}} = \pi/(2\chi)$. These choices cause the qubit to end in $|e\rangle$ if there is an odd number $2n + 1$ of photons in the cavity,




Figure 3.10: **Cavity calibrations that utilize single photon state preparation.** The initialized state before the pulses shown (qubit, black and cavity, orange) is assumed to be $|g1\rangle$. (a) A χ_{ge} Ramsey experiment allows us to accurately determine the dispersive shift term. Like with other Ramsey experiments, the offset frequency between the expected and observed oscillation gives us the correction to χ_{ge} . (b) Parity Measurement Calibration. We expect the parity time to be $\tau = \pi/(2\chi)$. However, the actual time may be slightly different. The pulse sequence is similar to the χ_{ge} Ramsey, but the phase of the second $\pi/2$ pulse is 0 (the same as the initial one), rather than being advanced proportional to τ . The shown data gives a parity time of 602 ns, compared to the $\chi_{ge} = 2\pi \times -0.359$ MHz. (c) Cavity self-Kerr (right) and cross-Kerr (left) calibrations. The probe π pulse is at the single photon frequency of mode j , and produces peaks centered at the cross-Kerr and self-Kerr values.




Figure 3.11: **Multimode cavity dispersive shift and Kerr values.** (a) χ_{ge} values between the qubit and each cavity mode (red crosses, left axis) and the corresponding coupling g 's to the qubit (blue circles, right axis) are shown. Modes are indexed from 0–8 in order of increasing frequency, and are the 9 modes between our qubit ($\omega_{ge} = 4.97$ GHz) and readout ($\omega_r = 7.79$ GHz). (b) Matrix of self-Kerr and cross-Kerr values of our multimode cavity, which are all on the order of a few kHz. We observe that the matrix is nearly symmetric, as we would ideally expect

as the qubit will accrue a phase $2\pi(2n+1)\chi/(2\chi) \equiv \pi$ during the wait time that cancels with the relative phase. On the other hand, the qubit will end in $|g\rangle$ if there is an even number $2n$ of photons in the cavity, as the accrued phase will be $2n\pi \equiv 0$. We could flip these results by setting the relative phase between the two qubit pulses to be 0 instead of π . In practice, we observe τ_{parity} directly by performing a Parity Measurement Calibration. Emulating the above procedure, after preparing a photon in the cavity, we sweep the wait time between two $\pi/2$ pulses with relative phase π , which is shown in Fig. 3.10(b). We observe an oscillation between $|g\rangle$ and $|e\rangle$, and can calculate τ_{parity} from its π pulse time.

Just as how we accurately determine χ_{ge} , we use a similar measurement to precisely measure χ_{ef} and χ_{gf} , the dispersive shift interactions between the $|e\rangle - |f\rangle$ and $|g\rangle - |f\rangle$ levels of the transmon. For χ_{ef} , we prepare a photon in the cavity, then perform an EF Ramsey measurement on the transmon. Similarly, for χ_{gf} , we initialize $|g1\rangle$, then perform a GF Ramsey measurement. This consists of a $\pi_{ge}/2$ pulse, π_{ef} pulse, varied wait time τ , π_{ef} pulse, and $\pi_{ge}/2$ pulse with advanced phase, so that the oscillation with respect to τ

provides a correction to χ_{gf} . If desired, we could perform similar calibrations for higher levels of the transmon ($|h\rangle$ and above) and precisely determine any of the transmon-cavity dispersive interaction terms.

Considering the multimode nature of the cavity, we now measure the cross-Kerr interaction $k_{ij}a_i^\dagger a_i b_j^\dagger b_j$ between the cavity modes. This Cross-Kerr Calibration experiment involves preparing a photon in mode i , then sweeping the frequency of a weak drive tone near the original frequency ω_j of mode j before performing a resolved π pulse on the qubit. We shift the center frequency of this resolved π pulse by $\chi_{ge,j}$, the dispersive shift of mode j , so that when mode j is unpopulated, the final qubit state will be $|g\rangle$. The result of this experiment will be a peak whose center frequency is shifted from ω_j , the frequency of cavity mode j , by k_{ij} , allowing us to determine the value of the cross-Kerr coefficient k_{ij} . Performing this experiment when $i = j$ gives us another way to determine the value of the self-Kerr. These Kerr calibrations are shown in Fig. 3.10(c), and the measured values are presented in Fig. 3.11(b). Due to the experimental procedure, the cross-Kerr calibration will produce positive peaks, as mode j is initially unpopulated, so the resolved π_{ge} pulse at $\omega_{ge} + \chi_j$ will have no effect until it becomes populated near ω_j . Conversely, the self-Kerr calibration will produce negative peaks, as now $i = j$, the mode is initially populated and the qubit will be in $|e\rangle$ when the drive tone frequency is far from ω_j . In this case, when the cavity is driven, its population in $|1\rangle$ will shrink (moving either to $|0\rangle$ or $|2\rangle$ and higher Fock states depending on the phase of the drive), reducing the final transmon $|e\rangle$ population. Determining the self-Kerr in this way produces consistent results with the Cavity Drive Ramsey approach shown in Fig. 3.8.

Finally, we measure the coherences of our cavity with Cavity T_1 and Cavity T_2 Ramsey experiments. For the T_1 experiment, after preparing $|g1\rangle$, we wait for a time τ and perform a resolved π pulse on the single photon peak at frequency $\omega_{ge} + \chi$. We observe an exponential decay based on the T_1 time. For the T_2 Ramsey, we instead prepare an equal superposition of $|0\rangle$ and $|1\rangle$ in the cavity, $(|0\rangle + |1\rangle)/\sqrt{2}$ (for example, by driving for half the length that




Figure 3.12: **Cavity coherence experiments.** (a) Cavity T_1 . Our maximum measurable τ is limited by our AWG memory and its pulse implementation; nevertheless, we observe an exponential decay that we fit to obtain $T_1 = 2.136$ ms. This corresponds to a mode quality factor of $Q = 76.7$ million. (b) Cavity T_2 . The $\pi_{cav}/2$ pulses indicate preparing the state $(|0\rangle + |1\rangle)/\sqrt{2}$ for the first pulse, then driving that pulse with opposite phase for the second pulse. The decay of the oscillation gives us $T_2 = 2.91$ ms. (c) Cavity Temperature experiment. We fit to the initial $|1\rangle$ population (blue) compared to $|0\rangle$ (black) by driving the transmon with narrow bandwidth pulses at those two number resolved frequencies. To help resolve the $|1\rangle$ oscillation, we use the fact that the oscillation frequency should be the same as that of $|0\rangle$. The contrast is 0.005, corresponding to a cavity temperature of 56 mK.




Figure 3.13: **Multimode cavity coherences and Q s.** (a) T_1 's (red circles) and T_2 's (blue squares) are roughly 2–3 ms, and have errorbars obtained from fitting to either the exponential T_1 decay or the T_2 oscillation decay. (b) Corresponding quality factors Q , calculated as $Q = \omega_c T_1$.

produced the single photon), wait for a variable time τ , then act that equal superposition pulse again with an advanced phase before applying a resolved qubit π pulse at $\omega_{ge} + \chi$ and measuring. Like the transmon Ramsey experiments, this will produce an oscillation whose amplitude decays exponentially depending on the T_2 of the cavity. Examples of these experiments are presented in Fig. 3.12. For the data shown, we obtain $T_1 = 2.14 \pm 0.05$ ms and $T_2 = 2.9 \pm 0.3$ ms. The coherences for all 9 of our cavity modes with frequencies between our qubit and readout frequencies are shown in Fig. 3.13(a). The T_1 's are mostly around 2 ms, while the T_2 's range from 2–3 ms. We also present the corresponding mode quality factors in Fig. 3.13(b), which primarily range from 70–95 million.

Parameter	Quantity	Values
Frequency (GHz)	$\omega_m/(2\pi)$	5.461, 5.716, 5.965, 6.223, 6.469, 6.727, 6.989, 7.247, 7.51
Dispersive shift (MHz)	$\chi_m/(2\pi)$	see Fig. 3.11
Self-Kerr (kHz)	$k_m/(2\pi)$	"
Cross-Kerr	$k_{mn}/(2\pi)$	"
Relaxation (ms)	T_1^m	see Fig. 3.13
Dephasing (ms)	T_2^m	"
Cavity Population (%)	\bar{n}	0.44, 0.62, 0.30, 0.39, 0.93, 0.91, 1.01, 0.99, 0.44

Table 3.2: **Multimode cavity system parameters.**

3.3.4 Cavity Beamsplitter Experiments

We are able to generate beamsplitter type interactions in our multimode cavity of the form:

$$\hat{H}_{BS} = g(t)\hat{a}_1^\dagger\hat{a}_2 + g^*(t)\hat{a}_1\hat{a}_2^\dagger. \quad (3.15)$$

This control has applications in manipulating coupled system components for quantum information processing [39–41], quantum simulation [42, 43], and studies of entanglement [44]. Based on the form of Eqn. (3.14), if we consider the readout cavity as another storage cavity mode and add a second drive tone, a cavity swap is possible when $\omega_{d2} - \omega_{d1} = \omega_{c2} - \omega_{c1}$. This interaction is specifically derived in [45, 46], which we summarize in the following. Starting with the Hamiltonian for two cavity modes coupled to a transmon with two drive tones, we have

$$\begin{aligned} \hat{H} = & \omega_1\hat{a}_1^\dagger\hat{a}_1 + \omega_2\hat{a}_2^\dagger\hat{a}_2 + \omega_{ge}\hat{c}^\dagger\hat{c} - E_J(\cos(\hat{\phi}) + \frac{\hat{\phi}^2}{2}) \\ & + \epsilon_1 \cos(\omega_{d1}t)(\hat{a}_1 + \hat{a}_1^\dagger) + \epsilon_2 \cos(\omega_{d2}t)(\hat{a}_2 + \hat{a}_2^\dagger), \end{aligned} \quad (3.16)$$

where $\hat{\phi}$ is the phase across the junction and can be expressed as

$$\hat{\phi} = \phi_{a_1}(\hat{a}_1 + \hat{a}_1^\dagger) + \phi_{a_2}(\hat{a}_2 + \hat{a}_2^\dagger) + \phi_c(\hat{c} + \hat{c}^\dagger), \quad (3.17)$$

where ϕ_j is the zero-point fluctuation of the phase for each system component. For example, for the transmon, $\phi_c^4 \sim 2E_C/E_J$. We then perform the rotating frame transformation at the system frequencies, $\hat{U} = e^{-i\hat{H}_0 t}$ with $\hat{H}_0 = \omega_1 \hat{a}_1^\dagger \hat{a}_1 + \omega_2 \hat{a}_2^\dagger \hat{a}_2 + \omega_{ge} \hat{c}^\dagger \hat{c}$, followed by the displacement transformation $D = \exp((\xi_1 e^{-i\omega_{d1} t} + \xi_2 e^{-i\omega_{d2} t}) \hat{c})$ that takes $\hat{c} \rightarrow \hat{c} + \xi_1 e^{-i\omega_{d1} t} + \xi_2 e^{-i\omega_{d2} t}$, with effective drive strengths $\xi_j = \epsilon_j / (\omega_{ge} - \omega_{dj})$. These transformations are similar to the ones that we performed for the sideband earlier in this chapter. We also consider the drives to be sufficiently off resonant, such that the detunings in the denominators of ξ_j are much greater than the linewidth of the transmon transitions, so that the transmon is only virtually excited by the drives.

We now obtain the beamsplitter Hamiltonian by expanding the $\cos(\hat{\phi})$ in Eqn. (3.16) to 4th order with the rotating wave approximation and $\omega_{d2} - \omega_{d1} = \omega_2 - \omega_1$. The remaining non-rotating terms are:

$$\hat{H}_{\text{final}} = \hat{H}_{\text{Stark}} + \hat{H}_{\text{Kerr}} + \hat{H}_{\text{int}} \quad (3.18)$$

$$\hat{H}_{\text{Stark}} = -E_J \phi_c^4 (|\xi_1|^2 + |\xi_2|^2) \hat{c}^\dagger \hat{c} = -2\alpha (|\xi_1|^2 + |\xi_2|^2) \hat{c}^\dagger \hat{c} \quad (3.19)$$

$$\begin{aligned} \hat{H}_{\text{Kerr}} = & -E_J (\phi_{a_1}^2 \phi_{a_2}^2 \hat{a}_1^\dagger \hat{a}_1 \hat{a}_2^\dagger \hat{a}_2 + \phi_{a_1}^2 \phi_c^2 \hat{a}_1^\dagger \hat{a}_1 \hat{c}^\dagger \hat{c} + \phi_{a_2}^2 \phi_c^2 \hat{a}_2^\dagger \hat{a}_2 \hat{c}^\dagger \hat{c}) \\ & - \frac{E_J}{4} (\phi_{a_1}^4 \hat{a}_1^\dagger \hat{a}_1 \hat{a}_1 \hat{a}_1 + \phi_{a_2}^4 \hat{a}_2^\dagger \hat{a}_2 \hat{a}_2 \hat{a}_2 + \phi_c^4 \hat{c}^\dagger \hat{c} \hat{c} \hat{c}) \end{aligned} \quad (3.20)$$

$$\hat{H}_{\text{int}} = -E_J \phi_c^2 \phi_{a_1} \phi_{a_2} (\xi_1 \xi_2^* \hat{a}_1^\dagger \hat{a}_2 + \xi_1^* \xi_2 \hat{a}_1 \hat{a}_2^\dagger). \quad (3.21)$$

Here, \hat{H}_{Stark} describes the shift in transmon frequency caused by the drives, \hat{H}_{Kerr} describes the cavity mode nonlinearities inherited from the transmon, which are independent of the drives and whose calibration was discussed earlier in this chapter, and \hat{H}_{int} describes our desired beamsplitter interaction. We can rewrite the coupling strength in Eqn. (3.15) as

$$g(t) = E_J \phi_a \phi_b \phi_c^2 \xi_1 \xi_2 = \sqrt{\chi_1 \chi_2} \xi_1 \xi_2, \quad (3.22)$$

where χ_j is the dispersive shift of mode j . From the expression for ξ_j , we see that the

beamsplitter rate is dependent on the detuning from the transmon, with drive tones closer to the qubit frequency ω_{ge} producing stronger rates. Therefore, to maximize the rate, we choose ω_{ge} to be approximately the center of our two drive tones, such that $\omega_{d1} \approx \omega_{ge} + \delta_{21}$ and $\omega_{d2} \approx \omega_{ge} - \delta_{21}$, where $\delta_{21} = \omega_2 - \omega_1$. We find that these choices of ω_{d1} and ω_{d2} are sufficiently far away from ω_{ge} and ω_{ef} to generate the beamsplitter interaction without unwanted excitation of the transmon's $|e\rangle$, $|f\rangle$, or higher levels.

We calibrate the beamsplitter using experiments similar to the Sideband and Photon Number Resolved Spectroscopy Experiments. First, we need to determine the frequencies of our drive tones with a Cavity Swap Frequency Scan. For this, we prepare either a single photon or a coherent state in the first cavity mode, as the beamsplitter will work on both of these states. Coherent states are more easily prepared by driving the cavity, but a single photon will produce a sole population peak with maximum contrast. After the preparation, we turn on the two drive tones, holding one of the frequencies fixed while sweeping the second one. We then apply a resolved qubit π pulse at the frequency corresponding to the prepared single photon peak ($\omega_{ge} + \chi_1$), or whatever photon number is most populated for the coherent state. The single photon version of this experiment is shown in Fig. 3.14(a). When the two tones are on resonance with the beamsplitter interaction, we observe a peak corresponding to a decrease in population of the original mode and increase in the population of the second mode. After determining the frequencies of the tones that work best for the beamsplitter, we find the swap time through a Cavity Swap Rabi experiment, where we prepare either initial state in the first cavity, turn on the two tones for a variable amount of time, apply the resolved π pulse at the photon number frequency of interest, and measure. For example, for a single photon, we observe complementary oscillations for π pulses applied at $\omega_{ge} + \chi_1$ and $\omega_{ge} + \chi_2$, as shown in Fig. 3.14(b) for two adjacent modes in our cavity. We use the beamsplitter to generate interactions between modes and transfer states between mode pairs. Like the sideband interaction, this gives us another way to control our multimode cavity.




Figure 3.14: **Cavity beamsplitter calibration experiments.** (a) Cavity Swap Frequency Spectroscopy. We sweep the frequency of one of the beamsplitter drives while observing the population of the prepared photon state. The population will be moved from the observed mode to the swap mode when the drive is on resonance (dashed vertical line), producing a negative peak. This resonant frequency is what we then set as the frequency of the swept drive. (b) Cavity Swap Rabi experiment. Starting with a photon in mode index 3 of our multimode cavity (red), we transfer the photon to mode 4 (blue) using the beamsplitter. The swap time is indicated by the dashed vertical line. We also observe the population of $|0\rangle$ (black) as an indicator for the excitation of the transmon.

In this chapter, we have demonstrated the methods that we use to determine the interaction strengths and parameters of our coupled transmon, readout cavity, and multimode storage cavity system. With relatively straightforward on/off drive pulses, we have fully calibrated our transmon and cavity Hamiltonian (Eqn. (2.12)) and characterized the performance of our devices by measuring their coherences. Using this system information, we have

also implemented pulse shaping techniques like DRAG pulses to improve the fidelity of our operations, as well as control techniques like four-wave mixing sidebands and beamsplitter interactions that can be used to engineer the cavity's quantum state. Looking ahead, in future chapters we discuss additional, more complicated control and measurement techniques that rely on the information obtained here.

Chapter 4

Cavity Photon Blockade

Photon blockade is a useful technique for engineering a cQED system. By applying a constant qubit drive at a frequency that targets a specific cavity photon number, we modify the cavity's energy level spacing and effectively partition its Hilbert space [47, 48, 27]. The blockade can also be thought of as a measurement of the photon number, and is intimately related to Quantum Zeno Dynamics [49–57]. This chapter will describe how to implement photon blockade and demonstrate some of its applications in controlling a single mode or multiple modes simultaneously.

4.1 Blockade Derivation

We start by considering the system of a single cavity mode coupled to a qubit. In the lab frame, this has the Hamiltonian

$$\hat{H}_{\text{lab}} = \omega_{ge} |e\rangle \langle e| + \omega_c \hat{a}^\dagger \hat{a} + \frac{k}{2} \hat{a}^\dagger \hat{a} (\hat{a}^\dagger \hat{a} - 1), \quad (4.1)$$

where we assume the energy of the ground state $|g0\rangle$ is 0, ω_{ge} and ω_c are the qubit and cavity mode frequencies, χ is the dispersive shift interaction, and k is the self-Kerr of the cavity. We perform the rotating frame transformation at the qubit and cavity frequencies with $\hat{U} = e^{-it\hat{H}_0}$, where $\hat{H}_0 = \omega_{ge} |e\rangle \langle e| + \omega_c \hat{a}^\dagger \hat{a}$, and add in drives on the qubit and cavity, so that

$$\hat{H} = \chi \hat{a}^\dagger \hat{a} |e\rangle \langle e| + \frac{k}{2} \hat{a}^\dagger \hat{a} (\hat{a}^\dagger \hat{a} - 1) + \{\Omega(t) |g\rangle \langle e| + \epsilon(t) \hat{a} + \text{c.c.}\}, \quad (4.2)$$

where $\Omega(t)$ and $\epsilon(t)$ are the qubit and cavity drive terms, respectively. When choosing a photon number $|n_0\rangle$ to blockade, we apply a qubit blockade drive at frequency $\omega_{ge} + n_0\chi$.

With this in mind, we can express $\Omega(t) = \Omega e^{-in_0\chi t}$. We now go into the frame of the drive and eliminate the time-dependence of this $\Omega(t)$ term by applying the frame transformation $\hat{U} = e^{-i\chi|e\rangle\langle e|n_0t}$, which gives

$$\tilde{H} = \chi(\hat{a}^\dagger\hat{a} - n_0)|e\rangle\langle e| + \frac{k}{2}\hat{a}^\dagger\hat{a}(\hat{a}^\dagger\hat{a} - 1) + \{\Omega|g\rangle\langle e| + \epsilon(t)\hat{a} + \text{c.c.}\}. \quad (4.3)$$

We further manipulate this Hamiltonian by individually diagonalizing each photon number subspace ($|gn\rangle, |en\rangle$). The blockade drive is resonant with $|gn_0\rangle - |en_0\rangle$, splitting those levels by 2Ω and equally mixing them. For Fock states $|n\rangle$ above or below $|n_0\rangle$, the dressing between $|gn\rangle$ and $|en\rangle$ is proportional to $\Omega/(\chi(n - n_0))$, to leading order in Ω/χ . We take the leading order in this case because the blockade Hamiltonian \tilde{H} is only valid in the regime where $\epsilon\sqrt{n_0} < \Omega < \chi$. The first of these conditions limits the maximum cavity drive strength to prevent leakage to $|\widetilde{gn_0}\rangle, |\widetilde{en_0}\rangle$ and past the blockaded level $|n_0\rangle$, while the second ensures that the qubit drive selectively blockades the $|gn_0\rangle - |en_0\rangle$ transition and minimally affects transitions corresponding to other photon numbers. The blockade can act as an upper or lower boundary (so that the Hamiltonian applies for $n < n_0$ or $n > n_0$). In the remainder of this derivation and in our experiments, we choose $n < n_0$. The Hamiltonian can be rewritten in terms of these dressed states, as

$$\begin{aligned} \tilde{H} - \chi(\hat{a}^\dagger\hat{a} - n_0)|e\rangle\langle e| - \frac{k}{2}\hat{a}^\dagger\hat{a}(\hat{a}^\dagger\hat{a} - 1) &= \sum_{n < n_0} \xi(t) \left(\sqrt{n+1} |\widetilde{gn}\rangle \langle \widetilde{g(n+1)}| + \text{c.c.} + \dots \right) \\ &+ \sum_{n < n_0} \sqrt{\chi^2(n - n_0)^2/4 + \Omega^2} (|\widetilde{en}\rangle \langle \widetilde{en}| - |\widetilde{gn}\rangle \langle \widetilde{gn}|), \end{aligned} \quad (4.4)$$

where $\xi(t)$ is the effect of $\epsilon(t)$ on the dressed states. In the above, we have dropped the drive terms that couple the dressed ground and excited states $|\widetilde{gn}\rangle$ and $|\widetilde{en}\rangle$, as these terms are off-resonant and suppressed by Ω/χ . The physics of the blockade can be approximated within a truncated Hilbert space that involves only the dressed transmon ground state and

not the excited one, described by the following Hamiltonian:

$$\begin{aligned}
\hat{H}_{\text{trunc}} &\approx \sum_{n < n_0} \left(\left[\frac{\kappa}{2} n(n-1) + \chi(n-n_0) - \sqrt{\chi^2(n-n_0)^2/4 + \Omega^2} \right] |\widetilde{g\bar{n}}\rangle \langle \widetilde{g\bar{n}}| \right. \\
&\quad \left. + \xi(t) \left(\sqrt{n+1} |\widetilde{g\bar{n}}\rangle \langle \widetilde{g(\bar{n}+1)}| + \text{c.c.} \right) \right) \\
&\approx \sum_{n < n_0} \left(\left[\frac{k}{2} n(n-1) - \frac{\Omega^2}{\chi(n-n_0)} \right] |\widetilde{g\bar{n}}\rangle \langle \widetilde{g\bar{n}}| \right. \\
&\quad \left. + \xi(t) \left(\sqrt{n+1} |\widetilde{g\bar{n}}\rangle \langle \widetilde{g(\bar{n}+1)}| + \text{c.c.} \right) \right).
\end{aligned} \tag{4.5}$$

We can see that the dressed cavity Fock state energies are corrected by the original self-Kerr k of the cavity, and by a new photon number dependent Stark shift induced by the blockade drive. Overall, the effect of the blockade drive is to hybridize the $|g, n_0\rangle, |e, n_0\rangle$ levels and shift them by $\approx \pm\Omega$, thereby engineering the energy spacing of the cavity mode.

Despite no direct occupation of the transmon, state preparation fidelity using photon blockade is limited by transmon decay and dephasing (in the limit of much smaller cavity loss). This infidelity is due to a combination of leakage to the dressed $|n_0\rangle$ state and subsequent decay via the transmon,

$$\mathcal{E}_{\text{leakage}} = \frac{\epsilon}{\Omega^2 T^q}, \tag{4.6}$$

and Purcell decay of the cavity states from additional dressing caused by the off-resonant blockade drive,

$$\mathcal{E}_{\text{Purcell}} = \frac{\Omega^2}{\epsilon \chi^2 T^q} \tag{4.7}$$

for the $|n_0 - 1\rangle$ state. Here, $T^q = \min(T_1^q, T_2^q)$ denotes the limiting qubit decoherence time. From Eqn. (4.6), Eqn. (4.7), and considering the regime where the blockade Hamiltonian is valid, we choose ϵ and Ω such that $\epsilon \ll \Omega \ll \chi$. However, we do not want ϵ to be too small, as that will force longer state preparation times with larger decoherence errors. In the case




Figure 4.1: **Photon blockade infidelity.** We plot the infidelity due to the leakage and Purcell decays from Eqns. 4.6 and 4.7, as well as cavity decay $\mathcal{E}_{\text{cav}} = 1/(\epsilon T_1^{\text{cav}})$, as a function of the blockade cavity drive strength ϵ for various values of the blockade dressing drive strength Ω . In this plot, all frequencies are in linear units. The dashed line indicates the infidelity of SNAP gates, $\sim 1/(\chi T^q)$. The dominant loss effect for lower values of Ω is the cavity decay, which increases the minimum infidelity (blue, orange, green). For larger values of Ω (red, purple, brown), the minimum infidelity asymptotes to a value set by the leakage and Purcell losses.

of blockading $|n_0 = 2\rangle$, optimizing the cavity drive strength results in a minimum infidelity of $\sim 1/(\chi T^q)$ when preparing $|1\rangle$ (blockade cavity π pulse). This infidelity is comparable to the error (primarily from qubit decoherence) in other state preparation techniques, like optimized SNAP gates [58–61], which is indicated by the dashed black line in Fig. 4.1. There, we also show calculated infidelities for different values of ϵ and Ω when including intrinsic cavity decay $\mathcal{E}_{\text{cav}} = 1/(\epsilon T_1^{\text{cav}})$. The effect of cavity decay will generally be larger for state preparation via photon blockade than SNAP due to the weaker cavity drives and correspondingly longer pulse times. For lower values of Ω , cavity decay is the dominant loss source, while for larger values of Ω , the leakage and Purcell losses limit the minimum infidelity. In practice, another consideration is that the cavity drive strength ϵ sets the rate of state preparation. To perform state operations quickly, we want ϵ to be as large as possible

without adversely affecting the overall blockade infidelity. In some cases, such as performing sequential gates, this may mean that we sacrifice a small amount of fidelity for a larger ϵ . However, this is primarily a future consideration and is not an issue for the experiments presented in the remainder of this chapter.

4.2 Generating the Blockade

To implement photon blockade on a single cavity mode, we perform several necessary calibration experiments. These are presented in Fig. 4.2. From Eqn. (4.5), we see that the blockade drive generates a Stark shift on the cavity drive frequency. We determine the new cavity drive frequency by performing a Blockade Cavity Spectroscopy experiment, where we turn on the blockade tone and sweep the frequency of a simultaneous cavity drive before applying a resolved qubit pi pulse and measuring. The result of this experiment is shown in Fig. 4.2(a). In this case, we probe the $|1\rangle$ level, which starts empty and can only be populated when we successfully drive the cavity. The center of the resulting peak, which for a square drive pulse will be the absolute value of a sinc function, will give us the desired cavity drive frequency. In Fig. 4.2(b), we see that the photon blockade succeeds in disallowing cavity population beyond the chosen photon number (in this case $|n_0 = 2\rangle$). Rather than generating a coherent state, a constant weak cavity drive will instead generate a Rabi oscillation between the two accessible Fock states $|0\rangle$ and $|1\rangle$. In (c), we show the experimentally measured cavity Stark shift and blockade splittings, which are consistent with the expressions from Eqn. (4.4) and Eqn. (4.5).

As mentioned in the previous section, one application of photon blockade is to prepare a single photon $|1\rangle$ in the cavity by applying the blockade at $|n_0 = 2\rangle$. To calibrate this, we perform a Blockade Rabi experiment. Similarly to the qubit Rabi experiment, we vary the time for which the simultaneous blockade and cavity drive tones are on before performing a




Figure 4.2: **Blockade calibrations.** (a) A Blockade Cavity Spectroscopy Experiment gives the new cavity frequency after it is Stark shifted by the blockade dressing drive. We probe the $|n = 1\rangle$ level while sweeping the frequency of the cavity drive, and adjust the cavity frequency based on the center of the peak (black dashed line). (b) A Blockade Rabi experiment demonstrates the effect of photon blockade, in this case for a blockade at $|2\rangle$. With the constant blockade dressing drive and a constant weak cavity drive, we are unable to populate the cavity Fock states at and beyond the blocked level. (c) An expanded cavity spectrum showing the energy splitting of the $|2_{\pm}\rangle$ levels due to the blockade as well as the Stark shift of the $|0\rangle - |1\rangle$ transition (central vertical blue line). The theoretical curves for the Rabi split $|1\rangle - |2_{\pm}\rangle$ transition are indicated by the dashed black lines, and asymptotically approach Ω .

resolved qubit π pulse at the frequency of the photon number resolved peak that we want to measure. If we start in the ground state $|g0\rangle$, we observe an oscillation between $|0\rangle$ and $|1\rangle$, and the π time will correspond to the preparation of $|1\rangle$. This is shown in Fig. 4.3. In (a), we provide a diagram illustrating the transmon-cavity energy levels and the effect of the blockade at $|n_0 = 2\rangle$ to shift the $|2\rangle$ level. In (b), we present Blockade Rabi data for a few different photon numbers, with the blockade π time marked by the vertical line. In (c), we present Wigner tomography of the cavity at the π time, which is a phase space representation of the cavity state [62] and will be explained more thoroughly in Chapter 6.

We use a similar procedure to prepare $|2\rangle$ and higher Fock states. After preparing $|1\rangle$, we turn off the blockade at $|2\rangle$, and instead blockade $|0\rangle$ and $|3\rangle$ simultaneously, as shown in Fig. 4.4(a). Now, performing a Blockade Rabi and driving the cavity like in Fig. 4.4(b) will generate an oscillation between $|1\rangle$ and $|2\rangle$, as those are the only reachable states in the




Figure 4.3: **Blockade preparation of $|1\rangle$.** (a) Energy level diagram corresponding to the blockade at $|2\rangle$. The photon number splitting $\chi = 2\pi \times -1.09$ MHz allows us to implement a dressing drive of strength $\Omega = 2\pi \times 108$ kHz that targets only the $|g2\rangle - |e2\rangle$ transition, hybridizing those states and generating a splitting by roughly Ω , as indicated by the green dashed lines. Then, a weak cavity drive (red squiggly arrow) can drive from the $|0\rangle$ state to $|1\rangle$, but not $|1\rangle$ to $|2\rangle$, creating the blockade. (b) Blockade Rabi data for the blockade at $|2\rangle$. Because there are only 2 accessible states $|0\rangle$ and $|1\rangle$, a constant cavity drive $\epsilon = 2\pi \times 20$ kHz produces an oscillation over time with higher photon numbers unpopulated. The drive time $\tau = 25\mu\text{s}$ that prepares $|1\rangle$ in the cavity is marked by the dashed vertical line. (c) Wigner tomography of the state prepared at the vertical line in (b). This measurement gives us a state preparation of $|1\rangle$ with fidelity 0.99 ± 0.005 .

blockaded space. There will therefore be a drive time that corresponds to exactly preparing $|2\rangle$ in the cavity. Before performing the Blockade Rabi, we need to redo the Blockade Cavity Spectroscopy experiment to find the cavity frequency in this situation, as the two blockade dressing drives at different frequencies will produce a different overall Stark shift. Nevertheless, we could continue this process for successively higher $|n\rangle$, where we first prepare $|n-1\rangle$, turn on two blockade tones at $|n-2\rangle$ and $|n+1\rangle$, and drive the cavity into $|n\rangle$.

Additionally, we can implement qudits by applying a blockade tone at $|d\rangle$. For the simplest case of $d=2$, we achieve universal control just with the cavity drive [63]. For larger d , we can achieve universal control in the blockaded qudit subspace of $\{|0\rangle, \dots, |d-1\rangle\}$ as well, although rather than just the on/off cavity drive, we need to use more complicated optimal control pulses [27, 64, 65]. We discuss this in more detail in Chapter 5. Overall, photon blockade allows us to easily and effectively control the state of a cavity mode. With just




Figure 4.4: **Blockade preparation of $|2\rangle$.** (a) Energy level diagram showing a simultaneous blockade at $|0\rangle$ and $|3\rangle$. The blockade drives $\Omega = 2\pi \times 108$ kHz (green) target the $|g0\rangle - |e0\rangle$ and $|g3\rangle - |e3\rangle$ selectively due to photon number splitting $\chi = 2\pi \times -1.09$ MHz. Starting from $|g1\rangle$, a cavity drive $\epsilon = 2\pi \times 20$ kHz (red squiggly arrow) can only access the $|g2\rangle$ state. (b) Blockade Rabi oscillation when starting in $|1\rangle$ while blockading $|0\rangle$ and $|3\rangle$. Since the only two accessible states are $|1\rangle$ and $|2\rangle$, we observe a Rabi oscillation between those two states when applying a constant cavity drive ϵ in the presence of the two blockade dressing drives. All the other photon numbers are unpopulated. The $|2\rangle$ state is prepared in the cavity at the time indicated by the dashed vertical line. (c) Wigner tomography of the state prepared at the vertical line in (b). Compared to $|1\rangle$, it has an additional node or fringe, which is a feature of cavity Fock states. We are able to prepare $|2\rangle$ using this approach with fidelity 0.985 ± 0.009 .

a constant qubit drive, we partition the accessible space of Fock states. Applying different or multiple blockade tones with cavity drives in succession gives us even greater control of the cavity state, which we can use to prepare arbitrary Fock states $|n\rangle$ or Fock state superpositions.

4.3 Multimode Blockade

We now move beyond a single mode and apply photon blockade to multiple modes. The blockade acts as an many-body interaction across M modes that we can utilize for quantum state preparation. Generating this interaction requires shifting the energies of multiple multimode states. Due to its photon number selectivity, the implementation and effectiveness of the blockade will depend on the dispersive shifts χ_j of each of the modes. We first consider the case where they are equal, $\chi_j = \chi$ for all modes j . In this case, a drive targeted at photon number n will blockade all of the corresponding multimode Fock states. For example, for $n = 3$ and $M = 2$ modes, the energy levels of $|03\rangle, |12\rangle, |21\rangle$, and $|30\rangle$ will all be shifted. If we start in the vacuum or ground state, $|00\rangle$, the only accessible states will be $|00\rangle, |01\rangle, |10\rangle, |02\rangle, |11\rangle$ and $|20\rangle$. In this case, the multimode blockade works essentially the same as for a single mode. A single drive tone and frequency will partition the effective Hilbert space by addressing the different energy levels that are all on resonance. If desired, we could also apply multiple of these resonant drives to further adjust the reachable states and restrict the space to some subset of photon numbers.

Even if the χ_j are not initially the same, if they are close, we can use χ -engineering techniques to make them effectively equal. For example, we can accomplish this by applying off-resonant sideband drives [66]. By driving a sideband near the $|h0\rangle - |e1\rangle$ transition with a detuning δ , we Stark shift the level $|e1\rangle$ inversely proportional to δ [28, 67]. The drive also slightly shifts the other energy levels involved in calculating the dispersive shift, as $\chi = \omega_{e1-g1} - \omega_{e0-g0}$. By choosing the detuning, its sign, and the strength of the drive, we induce an adjustment to one χ_j and make it match another $\chi_{j'}$. An example of this χ -engineering is shown in Fig. 4.5. This can work well for low mode number M and low photon number n ; however, each additional mode will require an additional sideband drive, and the drives will affect each other and complicate the calibration process. Consequently,




Figure 4.5: **Sideband χ dressing.** By applying a drive near the $|h0\rangle - |e1\rangle$ transition with a chosen detuning, we Stark shift the transmon-cavity energy levels and thereby modify the dispersive shift by $\delta\chi$. This can be used to match the χ between multiple modes when they are initially different. In practice, we often choose larger detunings at the cost of smaller $\delta\chi$ to minimize odd effects near resonance. Dashed lines are shown at 0 detuning and 0 shift.

while sideband χ -engineering is a useful technique, we do not utilize it to equalize our χ_j in our experiments.

We now consider the case where the χ_j are not all equal. The difference here is that a blockade tone that is on resonance with one mode will now be off resonance for another. One approach to implement the blockade at photon number n would be to apply a drive tone at each distinct frequency $\omega_{ge} + n\chi_j$. However, as we can see from Eqn. (4.5) and Eqn. (4.7), each of these drives will produce a Stark shift on the cavity drive, and each additional drive tone will contribute to the blockade infidelity. These effects will limit the effectiveness of the blockade and our state preparation, especially as the number of modes and distinct χ_j increases. Furthermore, having many drive tones that are close in frequency will produce beats that may have unwanted effects.

Instead, we can achieve the desired effects of the blockade with a single drive tone, assuming the χ_j are sufficiently close to each other. There are several criteria that must be

satisfied for a blockade drive at $\omega_d = \omega_{ge} + \delta\omega_b$. As in the single-mode case, the drive must be weak enough to satisfy $\Omega \ll \delta\omega_b - |\chi_j|$ for all modes j , so as to avoid dressing states with a different total photon number. On the other hand, the drive must also be strong enough to simultaneously blockade all multimode Fock states with a total of n photons. This requires $\Omega > |\delta\chi| = n|\chi_i - \chi_j|$ across all pairs of modes $\{i, j\}$. In our system, this criteria can be satisfied for all mode pairs that are nearest neighbours in frequency for at least $n = 3$; the mode pairs with the most disparate $\delta\chi$'s can still satisfy this condition for $n = 3$, while modes with more similar χ_i, χ_j can satisfy it for larger n . Additionally, even if a target multimode state is off resonant from the blockade drive, it can still be disallowed based on the drive-induced Stark shift $\Omega^2/(4\Delta_d)$, where Δ_d is the detuning between the blockade drive and the state. Driving any of the cavity modes with sufficiently weak drives ($\epsilon < \min(\Omega, \Omega^2/(4\Delta_d))$) in the presence of the blockade interaction results in constrained multimode dynamics within the space of states with total photon number $< n$ when starting from the vacuum state. This way, as shown in Fig. 4.6 for 3 modes and $n = 2$, we have implemented a hyper-plane of blocked multimode Fock states with the same total photon number, whose energies are all shifted, defining the multimode many-body interaction. This interaction acts when a state would normally be driven into the blocked hyper-plane, and effectively partitions the space of states that are accessible by cavity drives weaker than the interaction strength.

One class of states that we prepare with our multimode photon blockade is W states. These are entangled states of the form

$$|W_M\rangle = \frac{1}{\sqrt{M}} \sum_{j=1}^M e^{i\phi_j} |0..01_j0...0\rangle, \quad (4.8)$$

that are an equal superposition of a single photon in each of the M modes. To create W states, we apply a single blockade drive tone targeted at photon number $n = 2$. We choose a drive frequency offset from the qubit frequency by the average dispersive shifts of




Figure 4.6: **Hilbert space partition implemented by multimode photon blockade.** In this diagram, each gray point represents a multimode Fock state, and each axis corresponds to increasing photon number in one of the modes. For the case of blockading $n = 2$ on 3 modes, all states with total photon number 2 marked by the green hyperplane have their energies shifted. For a sufficiently weak cavity drive, when starting from the vacuum ground state ($|000\rangle$) the only remaining accessible states are $|000\rangle$ (orange), $|001\rangle$ (purple), $|010\rangle$ (blue), and $|100\rangle$ (red).

combinations of adding 2 photons, i.e.,

$$\omega_d = \omega_{ge} + \frac{2}{M} \sum_{j=1}^M \chi_j. \quad (4.9)$$

The resulting blockaded hyper-plane is shown in Fig. 4.6 for 2 photons and 3 modes. The blockaded states (in green) will be of similar form for other numbers of modes, but with different numbers of 0's. With this blockade tone, simultaneously driving the cavity modes with equal strengths for an appropriate time prepares W states.

We first demonstrate the multimode blockade interaction with qubit spectroscopy. In the absence of the blockade drive, independent weak drives in each mode produce coherent states with occupation in the two photon levels $|200\rangle$, $|020\rangle$ or $|002\rangle$ (and higher), as shown in




Figure 4.7: **Demonstrating the blockade with spectroscopy.** Photon number resolved spectroscopy data is shown for 3 modes with and without a blockade at $n = 2$. (a) When the blockade is off, driving any of the 3 modes produces classical coherent states. In particular, there is visible population at each of the 2-photon resolved peaks (green). (b) The blockade is applied at the frequency corresponding to the vertical green line, with dressing drive strength indicated by the shaded green band. Now, population of the 2-mode Fock states is disallowed, and we instead prepare an entangled 3-mode W state. Dashed horizontal lines indicate population of $1/3$, which we expect for that state. We attribute the peaks being slightly below the ideal $1/3$ to a combination of imperfect state preparation (fidelity < 1) and decay of the transmon during the resolved π_{ge} probe pulse, as due to the need for bandwidth narrow enough to distinguish between χ_j , we required a probe pulse of length $10.8 \mu\text{s}$, a sizable portion of the transmon T_1 time.

Fig. 4.7(a). However, in the presence of the blockade, the states with more than 1 photon have no occupation, as shown in Fig. 4.7(b). Here, 3 modes are driven simultaneously with equal strengths for the time required to prepare a W state, resulting in peaks only at $|100\rangle$, $|010\rangle$ and $|001\rangle$. Using the peak heights from spectroscopy, we measure the population oscillations as a function of time in Fig. 4.8 for 2–5 modes. We see that most of the population remains in the 0 and 1 photon multimode subspace, and that the populations are consistent with a

W state at the times when the vacuum state (orange circles) is emptied. By observing the return of the vacuum state oscillation, we see that the W state preparation fidelity decreases with increasing mode number, as it becomes increasingly difficult to stay in the regime where the blockade Hamiltonian is valid with a single drive tone when there are more χ_j 's that are increasingly disparate. Nevertheless, these W states were created using some of the modes with most different χ_j in our system, so we expect that their fidelities represent a worst case, and could be improved by using different sets of modes with more similar χ_j . The reason why we chose a set with more different χ_j is that their individual populations are more easily distinguished in spectroscopy experiments.

Continuing to larger n , the next case we examine is blockading the $n = 3$ photon level for 2 modes. Compared to the blockade at $n = 2$, we have increased the size of the allowed Hilbert space and the number of blockaded multimode Fock states. The maximum frequency difference between levels that the blockade will need to dress is $n(\chi_2 - \chi_1)$, where χ_2 and χ_1 are the dispersive shifts of the 2 modes of interest. An implementation of this blockade is shown in Fig. 4.9. As expected, none of the levels with photon number ≥ 3 are populated, and there are now 6 total 0-photon, 1-photon, and 2-photon levels in the accessible Hilbert space. Additionally, unlike the case of $n = 2$, driving with constant, equal cavity drives does not result in a time at which the original ground state population is restored. Instead, we observe that the two modes begin with similar population evolutions, but after hitting and interacting with the blockade, those evolutions begin to deviate due to the modes' different χ . In general, with these equal cavity drive strengths, there will be no time for which the state fully returns to vacuum. Nevertheless, the ability to implement this kind of interaction could be useful as a tool for preparing certain kinds of quantum states across multiple modes. In combination with different blockade parameters, changing the cavity drive strengths such as by making them unequal can produce completely different entangled states that are all still contained in a desired subspace of photon number levels.




Figure 4.8: **W state preparation Rabi oscillations.** We are able to effectively prepare multimode entangled W states in up to 5 modes by applying multimode photon blockade at $n = 2$. In all plots, the values are normalized such that the beginning population of the vacuum state is 1; this compensates for any decay during the long $10.8 \mu\text{s}$ qubit probe pulse. For M modes, the dashed line at $1/M$ indicates the expected population of each of the single-photon multimode Fock states when the vacuum state (orange circles) is emptied. (a) 2-mode W state. (b) 3-mode W state. (c) 4-mode W state. (d) 5-mode W state. In this case, the rising populations of the $n \geq 2$ levels (gray) indicate that the single blockade drive tone is struggling to forbid all of those states.




Figure 4.9: **Blockading $n = 3$ for 2 modes.** (a) Spectroscopy of a state created by the blockade. There are now 6 allowed levels that can be populated, and levels beyond those are disallowed. The blockade drive frequency is marked by the vertical green line, with dressing drive strength indicated by the shaded green band. (b) Populations over time of the 2-mode Fock states. The evolution of the states of the 2 modes begin similarly, before deviating after interacting with the blockade around the time when the $|00\rangle$ evolution reaches its minimum. The dashed line indicates the time at which the spectroscopy in (a) was performed.

In summary, we have applied photon blockade to simultaneously dress multiple cavity modes, creating a pure many-body interaction. The system remains linear and non-interacting until we reach the blockade number, at which point the modes interact strongly. Through this interaction, they can become entangled and, in the case of a blockade at $n = 2$, create W states. For blockades at larger n , we can similarly prepare entangled states between modes with population at higher multimode Fock states. Using only simple on/off cavity

and qubit dressing drives, we are able to use the blockade to partition the accessible Hilbert space of our bosonic modes, which can be leveraged for state preparation and generating entanglement. Overall, we have contributed to the multimode cQED toolbox with a scheme to realize these higher-order, exotic blockade interactions and state preparation techniques that generalize to arbitrary numbers of modes.

Chapter 5

Quantum Optimal Control in 3D cQED

Optimal control is a method to engineer the state of a quantum system. In this technique, control pulses are numerically optimized with a given Hamiltonian to produce a desired unitary gate or target quantum state. There are several approaches to solve these optimization problems, such as Gradient Ascent Pulse Engineering (GRAPE) [64, 65, 68], trajectories [69], Selective Number-dependent Arbitrary Phase (SNAP) pulses [58–61], Krotov for closed systems [70, 71], proprietary approaches like Q-CTRL’s Boulder Opal [72], and recently developed methods like Echoed Conditional Displacement (ECD) gates [73] and Pade Integrator COLlocation (PICO) [74], to name a few. There are a steadily growing number of these distinct quantum optimal control (QOC) algorithms that each claim different benefits. Rather than attempt to go into detail about every one of these methods, which could fill a book in itself, in this chapter we highlight some of the more common features shared between these algorithms. We also describe how we implement optimal control on our transmon and multimode cavity cQED system for a few different effective Hamiltonians.

5.1 Overview of Optimal Control Algorithms

This section will provide an overview of the structure of a QOC problem and its mathematical formulation. Afterwards, we give a brief summary of some of the different ways to solve the problem, and highlight some of the important considerations in the process.

A general QOC problem can be viewed as an optimization process that is subject to dynamics described by a Hamiltonian (or Liouvillian). The optimization is performed with respect to some set of time-dependent control fields $u_k(t)$, and aims to minimize deviations from a target state or target unitary. Assuming the action of the k th control on the system

is given by \hat{H}_k , we can express the total Hamiltonian with the controls as

$$\hat{H} = \hat{H}_0 + \sum_k u_k(t) \hat{H}_k, \quad (5.1)$$

where H_0 is the static or intrinsic part of the system. As part of the optimization, we choose a time interval $[0, T]$ such that T is the maximum duration of the controls. To make the problem tractable, we then discretize this period into N time steps of size Δt to numerically model the evolution of the state of the system, $|\psi(t_j)\rangle$ or $\rho(t_j)$, and the controls $u_k(t_j)$. If we have a total of M real controls, this will become an optimization problem over $M \cdot N$ real numbers. Depending on the framework of the problem, the form of the objective function that quantifies the optimization can vary.

We consider two main types of objectives for the optimal control problem: a target state or a target unitary. These represent a state optimization and a unitary optimization problem, respectively. State problems are more efficient in situations where we only care about how the controls acts on one or a few states, while unitary problems are most appropriate for implementing gates. For example, if we seek a set of controls that prepares $|3\rangle$ when starting in $|0\rangle$ in a cavity mode, and we do not care about what the controls do to $|1\rangle$ or $|2\rangle$, we should frame this as a state problem. On the other hand, if we are trying to implement a CNOT gate between two qubits, we should set that up as a unitary problem, as we care about the action of the controls on every basis state. In state optimization, the objective function is given by the infidelity between the target state and the final state generated by the controls \vec{u} , and can be expressed as

$$C(\psi_N(\vec{u})) = 1 - |\langle \psi_T | \psi_N(\vec{u}) \rangle|^2, \quad (5.2)$$

where ψ_T is the ideal target state and ψ_N is the state produced at time T after evolving under Eqn. 5.1 with \vec{u} . Similarly, for a unitary problem, the objective function for a target

\hat{K}_T is given by

$$C(\hat{K}_N) = 1 - \frac{1}{D} |\text{tr}(\hat{K}_T \hat{K}_N)| \quad (5.3)$$

where D is the size of the Hilbert space of the problem. We numerically compute ψ_N or \hat{K}_N by propagating the controls over the total time interval using our discretized time steps [65]. The evolution up to time $t_n = n\Delta t$ is given by the propagator

$$\hat{K}_n = \hat{U}_n \hat{U}_{n-1} \dots \hat{U}_1 \hat{U}_0, \quad (5.4)$$

where $\hat{U}_n = \exp(-i\hat{H}_n\Delta t)$ is the propagator for the time interval $[t_n, t_{n+1} \equiv t_n + \Delta t]$ and

$$\hat{H}_n = \hat{H}_0 + \sum_k u_{k,n} \hat{H}_k \quad (5.5)$$

is the discretized version of Eqn. (5.1) and is set by the values of the controls over that same n th time interval. Thus, we can express our final optimization problem as

$$\underset{\vec{u}_{1:N-1}}{\text{minimize}} \quad J(\vec{u}_{1:N-1}) = C(\psi_N(\vec{u}_{1:N-1})) \text{ or } C(\hat{K}_N(\vec{u}_{1:N-1})), \quad (5.6)$$

depending on whether we have a target state or unitary.

In practice, it is often necessary to impose additional constraints on the optimization problem. These could be due to hardware limitations, such as a maximum possible drive amplitude, or physics considerations, such as smoothness or Fourier components of the drives and ensuring that the conditions required for the modeled Hamiltonian remain satisfied. The controls should be producible in the lab and ideally avoid large amplitude fluctuations and rapid variations so that they are more robust to noise in the control instrumentation [65]. One way to impose these soft constraints is to add terms to the cost function $J(\vec{u}_{1:N-1})$, such that solutions that rely on any undesirable control features are penalized. For example, to

reduce the amount of fluctuation in the optimized controls, we apply the cost function

$$C_1(\vec{u}) = \sum_{k,n} |u_{k,n} - u_{k,n-1}|^2, \quad (5.7)$$

which acts as a penalty to the first derivative of the controls and thus suppresses the amount of rapid variations that they have. Another useful cost function serves to limit the population of some subset of chosen forbidden states in the Hilbert space. For each forbidden state $|\psi_F\rangle$, this can be written as

$$C_2(\vec{u}) = \sum_n |\langle \psi_F | \psi_n \rangle|^2. \quad (5.8)$$

This can be used to avoid occupation of higher energy, potentially lossier states during the optimized time interval. This is beneficial in reducing the chance of a spontaneous relaxation or decoherence event occurring while the controls are active. Perhaps most importantly, this cost is crucial for accurately modeling Hilbert spaces of large or infinite dimension, such as a bosonic cavity mode. In those cases, it is necessary to truncate the space at some level to model it with finite computational resources. However, the highest level in the model will then have an infinite anharmonicity, even though this is not really the case. To ensure the accuracy of our model and prevent the optimizer from utilizing this artificial nonlinearity, we forbid population of the highest level or two. Finally, depending on the problem type, each of these cost functions can be scaled by different weights to have more or less emphasis when performing the optimization.

Another way to implement constraints is to add inequalities or equations to the QOC problem. We reformulate the problem as a trajectory optimization [69], given by

$$\begin{aligned} & \underset{\vec{x}_{1:N-1}, \vec{u}_{1:N-1}}{\text{minimize}} && J(\vec{x}_{1:N-1}, \vec{u}_{1:N-1}) \\ & \text{subject to} && \vec{g}_n(\vec{x}_n, \vec{u}_n) \leq \vec{0} \quad \forall n \\ & && \vec{h}_k(\vec{x}_n, \vec{u}_n) = \vec{0} \quad \forall n, \end{aligned} \quad (5.9)$$

where \vec{g} and \vec{h} are vectors of constraints and \vec{x}_n is the state vector at time step n . We generally treat \vec{x} and \vec{u} as augmented state and control vectors, so that they contain all the relevant variables in the optimization problem, not just the quantum states and controls [69]. An example of an important inequality constraint is imposing a maximum amplitude for each of the controls, which can depend on hardware capabilities or may be informed by approximations that were made in developing the system Hamiltonian. Equality constraints can be used to force the beginning and end of the controls to 0 to ensure smoothness and prevent sudden jumps on either side. Constraints of this type can also be used to implement bandwidth limitations by forcing certain Fourier components to 0. This can be used to restrict unwanted transitions, such as to the higher excited states $|f\rangle$ or $|h\rangle$ of a transmon. Besides these examples, the ability to impose constraints of this form and the previously discussed cost functions offer a large amount of flexibility in QOC problems that gives the best chance of producing controls that will operate successfully on quantum hardware.

As mentioned previously, there are a multitude of different approaches to solving QOC problems. We highlight two different ones here: GRAPE [64, 65] and PICO [74]. As suggested by its name, GRAPE is a gradient-based method. Starting with an initial guess for the controls, in the first iteration of the optimization, the final state and intermediate states are computed using the propagator given by Eqn. (5.4). With the state information, we compute the total cost function J and update the controls via the method of steepest descent, given by

$$\vec{u}_{\text{new}} = \vec{u} - \beta \nabla_{\vec{u}} J, \quad (5.10)$$

where β is a constant that is typically referred to as the learning rate. We iterate on this process, continually updating the controls \vec{u} until we reach a desired value of J or the change in \vec{u} becomes negligible. The choice of β can vary for each iteration, and typically decreases with iteration number to assist with steady convergence to a minimum. This is shown




Figure 5.1: **Schematic of GRAPE algorithm.** Starting with an initial guess for the controls (gray circles), GRAPE modifies the controls by some learning rate times the gradient of the cost function J . In both plots presented, the rate decreases with each iteration (individual arrow). (a) A relatively small initial learning rate results in a smooth convergence to the minimum value of J . (b) A relatively large initial learning rate can still produce the same solution, although the value of the cost function may jump around slightly more. If the learning rate is too large or monotonically increases, the value of J will struggle or fail to converge. Some adaptive schemes occasionally increase the learning rate to generate faster convergence or help with escaping a local minimum.

schematically in Fig.5.1 for two different initial values of β . One simple form for β is an exponential decay, $\beta(m) = \beta_0 e^{-m/c}$ for some constant c for each iteration number m . There are also a number of schemes that adaptively determine $\beta(m)$, with some examples being L-BFGS-B [75] or ADAM [76]. These adaptive approaches can help the optimizer converge in fewer iterations and reduce the chance of getting stuck in a local minimum, which is a regular concern with gradient-based algorithms. GRAPE falls into the class of indirect solvers, in that the optimization treats the final state as a function of the controls such that only the controls (and not the states) are the decision variables.

On the other hand, PICO is a direct solver. In addition to the controls, it also treats the states as variables. Consequently, we can evaluate the objective or total cost function J more directly, as the quantum states are built into the optimization process and do not need to be computed with a propagator, like in indirect methods. One notable difference of direct methods is that because the states are variables, the dynamics of the system are

not strictly imposed during the solve, but are instead enforced as equality constraints in the problem [74]. The numerical solver can thus violate the dynamics during intermediate iterations before satisfying them at convergence, offering an increased flexibility that can improve the efficiency of the solve [69]. Similarly, the initial guess for the state can be unfeasible. Like in GRAPE, updates occur based on a learning rate β times the derivatives or gradients of J with respect to all the decision variables, like in Eqn. (5.10). The difference is that now the states are included in addition to the controls. Additionally, due to the possibly nonphysical starting state, the algorithm is less dependent on the quality of the initial guess and is less susceptible to getting stuck in local minima. At the end of the optimization, we must confirm that the final solution indeed satisfies the dynamics defined by the system Hamiltonian.

Now that we have summarized the mathematical formulation of a QOC problem and presented some examples of QOC solvers, in the following sections we consider some example problems and effective Hamiltonians that we implement on our cQED hardware. The main solvers that we use are the GRAPE package developed in [65] and the PICO solver from [74].

5.2 Optimal Control of a Transmon and Cavity Mode

Arguably the simplest system that we could start with is just a transmon. However, if we ignore pulse features that are sometimes desired, like frequency robustness or bandwidth limitations [69], universal control of a transmon can be effectively implemented with constant on/off drives. Numerically optimized pulses can still be very useful in these systems for speeding up operations or reducing sensitivity to different types of noise, but in many cases, like ours, it is not strictly necessary for preparing target states or gates. The next most straightforward system that we consider on our hardware is a transmon coupled to a single cavity mode. Due to the cavity mode being a bosonic oscillator with nearly equal

transitions between all energy levels (photon numbers), implementing arbitrary operations is much harder. There are some analytical approaches for quantum state preparation, like applying transmon-cavity sidebands as discussed in Chapter 3, photon blockade as discussed in Chapter 4, or SNAP pulses [58, 59]. Nevertheless, numerical optimal control techniques can improve upon these by finding faster pulses or determining the drive coefficients needed to reach a target state in those schemes. The remainder of this section will discuss the coupled transmon-cavity system and QOC that we implement on it.

5.2.1 Developing the System Hamiltonian

As part of setting up a QOC problem, we must first determine the Hamiltonian that accurately models the real system while also being as efficient as possible to help speed up the optimization process. We first consider a system with a transmon coupled to a single cavity mode. Starting with the Hamiltonian in Eqn. (2.12), we reduce to a single mode and add drive terms on our transmon and cavity mode to get:

$$\begin{aligned} \hat{H} = & \omega_{ge} \hat{b}^\dagger \hat{b} + \frac{\alpha}{2} \hat{b}^\dagger \hat{b} (\hat{b}^\dagger \hat{b} - 1) + \omega_c \hat{a}^\dagger \hat{a} + \chi \hat{a}^\dagger \hat{a} \hat{b}^\dagger \hat{b} + \frac{k}{2} \hat{a}^\dagger \hat{a} (\hat{a}^\dagger \hat{a} - 1) \\ & + \epsilon_q(t) (\hat{b} + \hat{b}^\dagger) + \epsilon_c(t) (\hat{a} + \hat{a}^\dagger), \end{aligned} \quad (5.11)$$

where $\epsilon_q(t)$ and $\epsilon_c(t)$ are complex, time-dependent drives on the transmon and cavity, respectively. Next, we go into the rotating frame at the transmon and cavity frequencies by applying the transformation $\hat{U} = e^{-i\hat{H}_0 t}$, with $\hat{H}_0 = \omega_{ge} \hat{b}^\dagger \hat{b} + \omega_c \hat{a}^\dagger \hat{a}$. This reduces the maximum frequencies in our system from a few GHz (ω_{ge}, ω_c) to a few hundred MHz (α). If we model the transmon as a two-level system, which is often still accurate, the maximum system frequency reduces even further to only a few MHz (χ) or tens of MHz depending on the maximum allowed values for $\epsilon_q(t), \epsilon_c(t)$. With lower maximum system frequencies, we can retain the accuracy of our model with a larger simulated time step for the optimal control

pulse, reducing the total number of steps required and thus the complexity of the problem. This also reduces the amount of real time needed to generate the pulse. The Hamiltonian in this frame becomes:

$$\begin{aligned}\hat{H}_{rot} = & \frac{\alpha}{2}\hat{b}^\dagger\hat{b}(\hat{b}^\dagger\hat{b} - 1) + \chi\hat{a}^\dagger\hat{a}\hat{b}^\dagger\hat{b} + \frac{k}{2}\hat{a}^\dagger\hat{a}(\hat{a}^\dagger\hat{a} - 1) \\ & + \epsilon_q(t)(\hat{b}e^{-i\omega_{ge}t} + \hat{b}^\dagger e^{i\omega_{ge}t}) + \epsilon_c(t)(\hat{a}e^{-i\omega_c t} + \hat{a}^\dagger e^{i\omega_c t}).\end{aligned}\quad (5.12)$$

Now, it is useful to express the drive terms as $\epsilon_q(t) = \epsilon_q^*(t) \cos(\omega_{ge}t + \phi_q(t))$ for the qubit and $\epsilon_c(t) = \epsilon_c^*(t) \cos(\omega_c t + \phi_c(t))$ for the cavity mode. With these substitutions, we rewrite our rotating frame Hamiltonian as:

$$\begin{aligned}\hat{H}_{rot} = & \frac{\alpha}{2}\hat{b}^\dagger\hat{b}(\hat{b}^\dagger\hat{b} - 1) + \chi\hat{a}^\dagger\hat{a}\hat{b}^\dagger\hat{b} + \frac{k}{2}\hat{a}^\dagger\hat{a}(\hat{a}^\dagger\hat{a} - 1) \\ & + \frac{\epsilon_q^*(t)}{2}(\hat{b}e^{i\phi_q(t)} + \hat{b}^\dagger e^{-i\phi_q(t)}) + \frac{\epsilon_c^*(t)}{2}(\hat{a}e^{i\phi_c(t)} + \hat{a}^\dagger e^{-i\phi_c(t)}),\end{aligned}\quad (5.13)$$

where we have expanded the cosines and used the rotating wave approximation to drop terms rotating at $2\omega_{ge}$ or $2\omega_c$, assuming $\phi_q(t) \ll \omega_{ge}$ and $\phi_c(t) \ll \omega_c$. We rearrange this Hamiltonian by expanding the $\phi_j(t)$ exponentials and grouping by $\cos(\phi_j(t))$ or $\sin(\phi_j(t))$ to get our final Hamiltonian,

$$\begin{aligned}\hat{H}_{rot} = & \frac{\alpha}{2}\hat{b}^\dagger\hat{b}(\hat{b}^\dagger\hat{b} - 1) + \chi\hat{a}^\dagger\hat{a}\hat{b}^\dagger\hat{b} + \frac{k}{2}\hat{a}^\dagger\hat{a}(\hat{a}^\dagger\hat{a} - 1) \\ & + \epsilon_{q,x}(t)(\hat{b} + \hat{b}^\dagger) + \epsilon_{q,y}(t)i(\hat{b}^\dagger - \hat{b}) + \epsilon_{c,x}(t)(\hat{a} + \hat{a}^\dagger) + \epsilon_{c,y}(t)i(\hat{a}^\dagger - \hat{a}),\end{aligned}\quad (5.14)$$

where $\epsilon_{j,x}(t) = \frac{\epsilon_j^*(t)}{2} \cos(\phi(t))$ and $\epsilon_{j,y}(t) = -\frac{\epsilon_j^*(t)}{2} \sin(\phi(t))$. We have a static Hamiltonian given by the first line of Eqn. (5.14), and a total of four drive terms that represent the real and imaginary quadratures of the transmon and cavity drives. We use these four drives to run any QOC problem in this rotating frame. Furthermore, this gives us a prescription for

returning from the rotating frame back to the lab frame:

$$\begin{aligned}\epsilon_q(t) &= \epsilon_{q,x}(t) \cos(\omega_{ge}t) - \epsilon_{q,y}(t) \sin(\omega_{ge}t), \\ \epsilon_c(t) &= \epsilon_{c,x}(t) \cos(\omega_c t) - \epsilon_{c,y}(t) \sin(\omega_c t).\end{aligned}\tag{5.15}$$

In the above, we have used the fact that we are in the limit where $\phi_j(t) \ll \omega_{ge}, \omega_c$, and we have applied cosine and sine sum-to-product formulas. Before we apply the Hamiltonian given by Eqn. (5.14) to generate optimal control pulses and test them on our hardware, we must first perform several calibration experiments to ensure that we accurately generate the pulses.

5.2.2 Calibration Experiments

To effectively implement optimal control pulses, we require precise knowledge of our drive lines and drive parameters. In our system, and for the pulses that will be discussed later in this chapter, we have two primary system components that we drive and manipulate, and therefore need to calibrate: our transmon and storage cavity.

First, we need to convert our dimensionless AWG amplitudes to a drive strength in frequency units, and vice versa. We measure this conversion with the experiments shown in the insets in Fig. 5.2. For the transmon, we vary the AWG amplitude in a series of Rabi experiments and observe the resulting oscillation frequencies. These directly reflect the transmon drive strength, and are shown in Fig. 5.2(a). The cavity calibration is slightly more complicated. Due to being a bosonic oscillator, we do not observe an oscillation between two energy levels when driving it. Instead, we vary the time and amplitude of a cavity drive pulse and measure the resolved photon number populations that it produces. These populations will follow a Poisson distribution that we fit to obtain the magnitude of the coherent state α produced in the cavity, as described in Chapter 3.3.1 and by Eqn. (3.5) with $t = 0$. We




Figure 5.2: **Qubit and cavity drive calibrations.** In both cases, the relationship is slightly nonlinear, especially at low values, reflecting the bit resolution of the AWG. We develop an interpolation function to convert from AWG amplitudes to drive strengths (green lines). (a) Qubit drive calibration. Drive strengths corresponding to each AWG amplitude are obtained by performing a Rabi experiment (inset) at that amplitude and fitting to the resulting oscillation frequency. (b) Cavity drive calibration. Drive strengths corresponding to each AWG amplitude are obtained by performing a series of spectroscopy experiments (inset) that allow us to fit the photon number populations to a Poisson distribution. In each series, we sweep the length of the cavity drive with a fixed amplitude, and the slope of the resulting line gives us the value of the drive strength for that amplitude.

perform this measurement at 4 distinct time steps for each fixed AWG amplitude, as shown in Fig. 5.2(b). Fitting to the slope of the line produced by those 4 points gives us the magnitude of the cavity drive strength that corresponds to that amplitude.

We repeat these experiments for multiple AWG amplitudes instead of just one. Ideally, we could measure the drive strength corresponding to a single amplitude and fit to a proportional relationship that should give us the full transfer function. However, in Fig. 5.2 we see that




Figure 5.3: **Cavity drive Fourier transform comparison.** The measured points (blue) were collected by sweeping the carrier frequency of the cavity drives, with the AWG amplitude to drive strength transfer function measured at the resonant cavity frequency. The data is consistent with the simulated pulse (green line), indicating that the response of our drive lines is frequency-independent near our cavity mode frequency.

the fits are not proportional or even linear, and instead have some significant curvature, especially at low AWG amplitudes. The reason for this behavior is the finite bit resolution of the AWG. In our case, the M8195A AWG has an 8-bit amplitude resolution, with 1 bit reserved for the sign. The remaining 7 bits therefore cover the entirety of the 0-1 amplitude magnitude range, giving us a bit resolution of $1/2^7 = 1/128$. Rounding and the way the AWG handles values near these bit resolutions produces the observed nonlinear behavior. Additionally, it can be useful to regularly redo these full amplitude calibrations, as the calibration tends to drift over time.

The previous calibration experiments involved resonantly driving the qubit and cavity. However, most optimal control solves produce drive pulses that have have significant off-resonant contributions that are crucial to their performance. Consequently, we must also measure our experimental drive response off resonance. Ideally, the response at our device will be the same regardless of the drive frequency, so that we can continue to use our existing fixed-frequency calibration. However, depending on the hardware wiring and filtering, the

response may be slightly frequency-dependent. Our calibration experiments involve varying the modulation frequency of the qubit or cavity pulses obtained from QOC. We apply each set individually, so that when calibrating the qubit we do not apply the cavity pulses, and vice versa. After the pulse, we observe the qubit or cavity population and compare with simulation. An example of this measurement for the cavity is shown in Fig. 5.3. Like in the previous calibration, the coherent state magnitudes $|\alpha|$ were obtained from fitting the photon number populations to Poisson distributions. We compare the Fourier transform of the ideal pulse to the measured magnitudes from sweeping the cavity carrier frequency. Discrepancies can be attributed to the aforementioned effects of a slightly frequency-dependent response, or small systematic errors like AWG quantization that originate from the finite bit resolution of our control electronics. Nevertheless, the results are generally in good agreement, indicating that our fixed-frequency AWG amplitude calibrations will work well for our optimal control pulses.

5.2.3 Preparing Fock states

On our system of a transmon and a single cavity mode, we primarily use QOC to prepare Fock states. This class of states has a clear form of increasing complexity (increasing photon number), which allows us to test how QOC performs on successively more difficult target states. Preparation of higher Fock states enhances the emission of cavity photons, and has applications in axion dark matter detection [77, 78, 18].

One important consideration in the QOC setup is the number of cavity levels to model. We prefer to include as few levels as possible to improve the runtime of the solve, but we must retain enough levels to accurately reflect the system. In particular, we need to model more levels if we utilize stronger cavity drives that can quickly populate the cavity. For cavity drive strengths on the order of a few MHz, we typically model 15-20 levels, such that the expected population of higher levels beyond that is negligible. After generating a set of

controls, we can verify this behavior by simulating the state evolution with a larger number of levels. This single simulation will be much less costly than including the greater number of levels throughout the entire QOC problem.

In our optimization, the additional cost functions that we add to the objective are a penalty on the slope of the controls with respect to time (in the form given in Eqn. (5.7)), a constraint forcing the endpoints to be zero, and a penalty on population of the highest modeled cavity photon number (in the form given in Eqn. (5.8)). The first two conditions improve the smoothness of the controls and thus the accuracy with which our control electronics can reproduce them, while the last condition improves the accuracy of our model by ensuring that the optimizer does not take advantage of the artificial infinite anharmonicity of the highest cavity level.

We present QOC pulses that prepare Fock states up to $|4\rangle$. These pulses are shown in Fig. 5.4. For generating the $|1\rangle$ pulse, we use the PICO algorithm [74], while for the higher Fock states $|2\rangle - |4\rangle$ we use the GRAPE package from [65]. The solutions are fundamentally different, as the pulses for $|1\rangle$ in Fig. 5.4(a) and (b) were found with a stronger maximum cavity drive strength than the other Fock states (c)–(h). This produces a qualitatively different solution where the cavity drive strength fluctuates more and the qubit drive oscillates slightly less. For increasing photon numbers from $|2\rangle - |4\rangle$, the higher Fock states have increasingly more complex features and additional fluctuations in the pulse, reflecting the difficulty in preparing those states. Even though PICO and GRAPE are vastly different algorithms, (the former is direct and the latter is indirect), they are both able to find solutions to the QOC problem with final simulated fidelities near 99%. PICO was able to find a shorter time solution, in roughly 981 ns instead of 2000 ns, although that may partially be from preparing a lower photon number state and having a minimum-time function built in, which our GRAPE package did not. Implementing these pulses produces the results shown in Fig. 5.5. With a resolved spectroscopy experiment, we measure the fidelities of Fock states $|1\rangle - |4\rangle$ prepared




Figure 5.4: **QOC Fock state preparation on a cavity mode.** We show controls and transmon and cavity populations over time for pulses that prepare (a),(b) $|1\rangle$; (c),(d) $|2\rangle$; (e),(f) $|3\rangle$, and (g),(h) $|4\rangle$.




Figure 5.5: **QOC Fock state preparation results.** A resolved spectroscopy experiment allows us to fit peak heights corresponding to the fidelities of our QOC pulses that prepare $|1\rangle - |4\rangle$. We are able to produce $|1\rangle$ (blue circles) with fidelity 0.988 ± 0.011 , $|2\rangle$ (orange squares) with fidelity 0.983 ± 0.015 , $|3\rangle$ (green triangles) with fidelity 0.962 ± 0.024 , and $|4\rangle$ (pink diamonds) with fidelity 0.950 ± 0.021 . Vertical dashed lines indicate the χ photon number resolved peaks for $|0\rangle - |5\rangle$ from right to left.

by our QOC pulses to be 0.988 ± 0.011 , 0.983 ± 0.015 , 0.962 ± 0.024 , and 0.950 ± 0.021 , respectively. The corresponding simulated fidelities are 0.997, 0.993, 0.994, and 0.974, respectively. We attribute the discrepancies to slight calibration errors in the self-Kerr parameter k , slight drive calibration changes over time, and AWG quantization errors. In particular, in Fig. 5.5 the pulse for $|4\rangle$ has a visible nonzero baseline, indicating that the qubit is still partially excited, rather than being fully returned to $|g\rangle$. The decreasing fidelities can also be seen in the progressively lower peak heights.

As a final note for this discussion, there are other meaningful QOC problems to solve besides Fock states. Another class of problems that QOC could be used for is quantum error correction encodings. For example, to encode a qubit state using the binomial code [31, 32, 28, 33], we would want controls that simultaneously take $|g0\rangle \rightarrow |g\rangle \otimes (|0\rangle + |4\rangle)/\sqrt{2} \equiv |0\rangle_L$ and $|e0\rangle \rightarrow |g2\rangle \equiv |1\rangle_L$. Additionally, to correct a photon loss during this encoding, we would want another pulse that performs $|g3\rangle \rightarrow |0\rangle_L$ and $|g1\rangle \rightarrow |1\rangle_L$. QOC can determine controls that implement these operations, giving it applications in quantum error correction. Beyond

that, QOC can be applied to really any problem, and can offer shorter-time solutions that minimize decay and decoherence effects in quantum systems.

5.3 Optimal Control with Photon Blockade

Besides the bare undressed transmon-cavity Hamiltonian, we perform optimal control in the dressed frame generated with photon blockade. In this case, we start with the Hamiltonian in Eqn. (4.3), which we reproduce here. For a single blockade tone that targets the photon transition $|n_0\rangle$, our system is defined by

$$\hat{H} = \chi (\hat{a}^\dagger \hat{a} - n_0) |e\rangle \langle e| + \frac{k}{2} \hat{a}^\dagger \hat{a} (\hat{a}^\dagger \hat{a} - 1) + \{\Omega |g\rangle \langle e| + \epsilon(t) \hat{a} + \text{c.c.}\}, \quad (5.16)$$

where Ω is the strength of the dressing drive and $\epsilon(t)$ is the time-dependent drive that we give to the optimizer. We rewrite this Hamiltonian as:

$$\hat{H} = \chi \hat{a}^\dagger \hat{a} |e\rangle \langle e| + \delta\omega_b |e\rangle \langle e| + \frac{k}{2} \hat{a}^\dagger \hat{a} (\hat{a}^\dagger \hat{a} - 1) + \{\Omega |g\rangle \langle e| + \epsilon(t) \hat{a} + \text{c.c.}\}, \quad (5.17)$$

where $\delta\omega_b = \omega_b - \omega_{ge} \equiv -n_0\chi$ is the difference between the qubit and blockade drive frequencies. In contrast to the previous transmon and cavity case, we now only have a cavity drive. We apply the blockade dressing drive on the qubit, but its magnitude is time-independent, so it is not part of the optimization process (outside of its effect on the system Hamiltonian). Additionally, we must ensure that we remain in the regime where the blockade is valid during the optimal control process, which requires $\epsilon_{\max} \ll \Omega \ll \chi$, where ϵ_{\max} is the maximum magnitude of $\epsilon(t)$.

Like before, to help perform the optimal control, we solve the problem in the rotating frame and split $\epsilon(t)$ into its real and imaginary quadratures, $\epsilon_x(t)(\hat{a} + \hat{a}^\dagger)$ and $\epsilon_y(t)i(\hat{a}^\dagger - \hat{a})$, and treat the two drives independently in the optimization process. However, if we want to

keep the same ϵ_{\max} from the lab frame, we must adjust the maximum drive amplitudes of these quadratures to $\epsilon_{\max}/\sqrt{2}$. Similarly, to reconstruct the lab frame pulse, we apply the modulation given by Eqn. (5.15).

One application of optimal control in this scheme is to generate qudits in our cavity modes. Since photon blockade partitions the accessible Hilbert space, we can easily implement a space with d levels. The remaining challenge, and the reason why optimal control is necessary, is to demonstrate universal control on those d levels. For this purpose, we use optimal control to generate pulses that take cavity Fock states $|0\rangle \rightarrow |1\rangle, \dots, |0\rangle \rightarrow |d-1\rangle$ to control all d levels of the qudit. For $d = 2$, we can accomplish this with a constant cavity drive and do not need more complicated optimal control pulses. We present the next highest case of a qutrit ($d = 3$) in Fig. 5.6. In Fig. 5.6(a), we show the energy level diagram that corresponds to the blockade at $|n_0 = 3\rangle$, restricting the accessible Hilbert space to three levels. We show pulses that take $|0\rangle \rightarrow |1\rangle$ (Fig. 5.6(c) and (g)) and $|0\rangle \rightarrow |2\rangle$ ((d) and (h)), as well as a constant cavity drive ((b) and (f)) for comparison. The pulses that prepare the Fock states have nontrivial, numerically determined shapes that generate $|1\rangle$ and $|2\rangle$ with fidelities 0.953 ± 0.022 and 0.965 ± 0.022 , respectively. Additionally, the maximum amplitudes of the lab frame drive pulses will be roughly $\epsilon_{max} = 2\pi\sqrt{20^2 + 20^2} \approx 2\pi \times 28.3$ kHz, which satisfies our required condition that $\epsilon_{max} \ll \Omega \ll \chi$, as $\Omega = 2\pi \times 108$ kHz and $\chi = 2\pi \times -1.08$ MHz. The Wigner tomography plots in Fig. 5.6(f)–(h) are a phase space representation of the cavity state, and will be discussed further in Chapter 6. The pulses that prepare $|1\rangle$ and $|2\rangle$ each produce a state that is symmetric around the origin, which is a feature of Fock states, while the uniform pulse prepares some superposition of all three photon number states that make up the qutrit ($|0\rangle, |1\rangle, |2\rangle$). Thus, optimal control combined with photon blockade allows for universal control of qutrits implemented on our cavity modes. This can be extended further to higher dimensional qudits. We have therefore added a class of pulses to the optimal control tool box in cQED.




Figure 5.6: **Universal control of a qudit using optimal control.** In the pulse plots ((b)–(h)), the uniform pulse is presented in the black boxes, |1> is in the purple boxes, and |2> is in the orange ones. (a) Energy level diagram for photon blockade of $|n_0 = 3\rangle$. The blockade drive with Rabi strength Ω resonant with $|g3\rangle - |e3\rangle$ causes those two levels to hybridize and shift by $\pm\Omega$ (dashed green lines). A resonant cavity drive (red arrows) drives the cavity in the qutrit space formed by the three levels of $|0\rangle, |1\rangle, |2\rangle$, with $|3\rangle$ remaining off-resonant and unoccupied. Higher photon numbers are also inaccessible. To implement the blockade, we apply $\Omega = 2\pi \times 108$ kHz with $\chi = 2\pi \times -1.08$ MHz. (b)–(d) The two quadratures of the pulses for a uniform cavity drive and the two optimal control pulses that prepare |1> and |2> from $|0\rangle$. The optimal control pulses are generated using the GRAPE algorithm. (e) Spectroscopy of the population evolution over time for the 3 prior pulses. Dashed lines indicate the qubit frequencies corresponding to $n = 0 - 3$ photons. The thick horizontal green line near the top of the plots marks the frequency of the blockade tone, which is equal to the dispersive shift frequency corresponding to 3 photons. (f)–(h) Wigner tomography over time of the uniform drive and optimal control pulses shown in (b)–(d), in the same order. The dashed green circle indicates the location (in phase space) of the blockade drive. The blockade acts as a wall that constrains the region of allowed occupation, with the states prepared by interference from the resulting reflections. We prepare |1> and |2> with fidelities 0.953 ± 0.022 and 0.965 ± 0.022 , respectively. |1> and |2> have tomography plots that are symmetric around the origin, which is a feature of Fock states, while the state produced by the uniform pulse has a more complicated plot caused by being some superposition of all three of the qutrit states $|0\rangle, |1\rangle$, and $|2\rangle$. For these experiments, we used cavity mode index 3 of our system.

Due to the blockade, optimal control in this scheme is noticeably different from the previous case of an undressed transmon and cavity mode. For one, our pulses take much longer ($25 \mu\text{s}$ instead of $\approx 1 \mu\text{s}$), as we are limited by the blockade regime condition $\epsilon \ll \Omega \ll \chi$. This causes our prepared states to have slightly lower fidelities, as there is more time for decay and decoherence effects. On the other hand, the optimization process itself can be performed much more quickly, as the blockade forbids cavity population beyond the targeted photon number. Even if we include an extra level or two to allow for leakage past the blockade, we only need to model 5 total cavity levels, compared to 15 for the undressed cavity mode. Additionally, we only have the cavity drive and no transmon drive. This decreases the size of the Hilbert space that we need to accurately model the system dynamics, and reduces the complexity of the QOC problem, allowing us to obtain solutions more quickly. Some of this improvement is lost to modeling a longer pulse time, but we often have coarser time steps because of the lower dominant frequency in the system, which is set by χ for the blockade, but is often set by the optimized pulse drive strengths for the undressed Hamiltonian.

QOC succeeds in preparing quantum states even with the engineered Hamiltonian implemented by photon blockade, reflecting the fact that it is a numerical process that can be applied to any quantum system. In the next section, we will move from this relatively small blockaded Hilbert space with two independent controls to a more complicated multimode problem that has both more controls and a larger problem size.

5.4 Multimode Optimal Control

In this section, we extend our existing optimal control discussion to a larger multimode system. Generalizing the transmon and single cavity mode Hamiltonian in Eqn. (5.14) to

multiple modes, we have

$$\begin{aligned} \hat{H}_{rot} = & \frac{\alpha}{2} \hat{b}^\dagger \hat{b} (\hat{b}^\dagger \hat{b} - 1) + \sum_j \left[\chi_j \hat{a}_j^\dagger \hat{a}_j \hat{b}^\dagger \hat{b} + \frac{k_j}{2} \hat{a}_j^\dagger \hat{a}_j (\hat{a}_j^\dagger \hat{a}_j - 1) + \sum_{i \neq j} k_{ij} \hat{a}_i^\dagger \hat{a}_i \hat{a}_j^\dagger \hat{a}_j \right] \\ & + \epsilon_{q,x}(t) (\hat{b} + \hat{b}^\dagger) + \epsilon_{q,y}(t) i (\hat{b}^\dagger - \hat{b}) + \sum_j \left[\epsilon_{c_j,x}(t) (\hat{a}_j + \hat{a}_j^\dagger) + \epsilon_{c_j,y}(t) i (\hat{a}_j^\dagger - \hat{a}_j) \right], \end{aligned} \quad (5.18)$$

where k_{ij} is the cross-Kerr interaction between modes. Besides the cross-Kerr and the transmon anharmonicity, we have multiple copies of the cavity mode that we discussed earlier, each with their own χ_j and κ_j , and a corresponding pair of drive quadratures to optimize. For M modes, we have a total of $2M + 2$ drives that we treat independently to implement our target operation.

One of the main challenges in multimode optimal control is the problem size. With each additional cavity mode, we need to model another set of 10 or more mode levels, increasing the dimension of the Hilbert space by at least an order of magnitude. In particular, that dimension will scale exponentially with the number of modes, rapidly increasing the computational complexity of the QOC problem. As a result, we have limited our hardware implementation to a low photon count target state and relatively weak cavity drives that allow us to model fewer levels. In principle, with sufficient computational resources or time, we could continue on to solve QOC problems in even larger multimode spaces. Another possible approach would be to develop QOC algorithms that can better handle these large matrices by taking advantage of sparsity or some other feature to efficiently perform the matrix calculations involved in the optimization.

The simplest multimode state that we prepare is the 2-mode Fock state $|11\rangle$. Even though this state is not entangled, it serves as a demonstration of how optimal control can realize this state in a different manner than, for example, sequentially preparing a photon in one mode followed by another. To generate the control pulse, we use the GRAPE package in [65]. The controls produced by solving the QOC problem and the resulting data are shown in Fig. 5.7.




Figure 5.7: **QOC multimode preparation of $|11\rangle$** . (a) Resolved spectroscopy of the state prepared by QOC pulses. The fidelity extracted from the fit peak height is 0.976 ± 0.009 . The dashed vertical black line indicates the frequency shift corresponding to $\chi_1 + \chi_2$, the rightmost dashed vertical gray line indicates the original qubit frequency, and the red and blue dashed lines indicate the frequency shifts of the two individual modes, i.e. multiples of χ_1 and χ_2 . (b) Control pulses that prepare $|11\rangle$ from $|00\rangle$. There are 6 total, corresponding to the real and imaginary quadratures of 3 drives: one on the transmon, and one on each of the two cavity modes. (c) State evolutions over time for the transmon and each of the two cavity modes.

With these pulses, we are able to produce the $|11\rangle$ state with fidelity 0.976 ± 0.009 , which agrees with the simulated fidelity of 0.983. We chose a total control time of $4\mu\text{s}$, although this is likely not the minimum time solution, since none of the controls are limited by their maximum bounds, as shown in Fig. 5.7(b). In Fig. 5.7(c), we see that the state preparation

is not strictly sequential, as there is simultaneous population of various Fock states in both of the cavity modes.

Ultimately, we have demonstrated the flexibility of optimal control to operate on a variety of system Hamiltonians. They all involve a nonlinear device coupled to one or more bosonic oscillator modes. We implemented the optimized pulses on our superconducting cQED system, but the same procedure can be applied to other hardware platforms as well. QOC is an ever-expanding field with new algorithms regularly being developed and applied to increasingly complex problems, serving as a promising avenue for efficiently implementing quantum states and gates.

Chapter 6

State Characterization and Tomography

In previous chapters, we have demonstrated several quantum state preparation and control techniques on a 3D cQED system. An important related aspect is accurate state reconstruction. This is generally referred to as quantum state tomography (QST), and is a fundamental part of quantum information processing. Here, we focus on QST of one or more cavity modes. We have already seen some figures that feature one method for cavity QST, Wigner tomography. This chapter will go into detail about how we use Wigner tomography for faithful cavity state reconstruction.

6.1 Wigner Tomography

Wigner tomography provides a way to measure the density matrix of a cavity through displacements and parity measurements, allowing for complete characterization of the oscillator's quantum state [38]. The Wigner operator is defined as

$$\hat{\mathcal{W}}(\alpha) = \frac{2}{\pi} \hat{\mathcal{D}}(\alpha) \hat{\Pi} \hat{\mathcal{D}}(-\alpha), \quad (6.1)$$

where $\hat{\mathcal{D}}(\alpha) = e^{\alpha \hat{a}^\dagger - \alpha^* \hat{a}}$ is the displacement operator for complex α and $\hat{\Pi} = e^{i\pi \hat{a}^\dagger \hat{a}}$ is the parity operator for a single mode. If $\alpha = 0$, then $\hat{\mathcal{W}}(\alpha) \propto \hat{\Pi}$. We can express the Wigner operator in the Fock basis [62, 79] as

$$\hat{\mathcal{W}}(\alpha)(|n\rangle \langle m|) = \frac{2}{\pi} e^{-2|\alpha|^2} \sqrt{\frac{m!}{n!}} (-1)^{m-1} (2\alpha)^{n-m} L_m^{n-m}(4|\alpha|^2), \quad (6.2)$$

where L_m^{n-m} is a generalized Lagurre polynomial. Measurements of the Wigner operator on a state ρ can be expressed as

$$x = \text{Tr}(\hat{\mathcal{W}}(\alpha)\rho). \quad (6.3)$$

We must perform multiple measurements of this form for different values of α in order to reconstruct the state and determine the underlying ρ . However, even before that, we can use the results of a Wigner tomography measurement x to make conclusions about the state. For example, a completely classical coherent state $|\alpha_{\text{prep}}\rangle$ will produce a positive value of x for any displacement α , while non-classical Fock states, such as the single photon state $|1\rangle$, will have regions where x will be negative, as shown in Fig. 6.1. In general, if the Wigner operator is negative anywhere, the underlying state must be quantum. Also, the result of the Wigner tomography with zero displacement ($\alpha = 0+0i$) alternates between its maximum and minimum values for increasing Fock states $|0\rangle - |3\rangle$, corresponding to the parity measurement flipping between 1 and -1 based on the even/odd number of photons in the cavity.

We now describe how to use Wigner tomography to identify a prepared state ρ . The measurement sequence for a single α involves preparing ρ , then acting $\hat{\mathcal{W}}(\alpha)$ by displacing the cavity by α , followed by a parity measurement $\hat{\Pi}$. By repeating this procedure for multiple displacements, we go from a set of Wigner tomography measurements \vec{x} to a reconstructed state through the matrix inversion technique in [79], which we summarize here. We can express the results of our measurements as

$$\vec{x} = \mathcal{M}|\rho\rangle\rangle, \quad (6.4)$$

where $|\rho\rangle\rangle$ is the vectorized form of ρ (matrix flattened into a vector) and \mathcal{M} is the matrix representing all the measurement operators. In our case, \mathcal{M} will correspond to the Wigner operators of each displacement α up to the maximum dimension D of our Hilbert space, so its total size will be $N \times D^2$, where N is the number of α 's that we measure. If N is exactly




Figure 6.1: **Wigner tomography examples.** State labels are shown above each of the plots. The x and y axes are the real and imaginary parts of the cavity displacement α . For purely classical coherent states $|\alpha\rangle$, the Wigner tomography is strictly positive, and the center of the Gaussian phase space distribution (red) is located at coordinates $(\text{Re}[\alpha], \text{Im}[\alpha])$. On the other hand, Fock states are non-classical, and thus have regions of negativity (blue). The value of the tomography at the center of each plot corresponds to the result of a parity measurement of the state, and therefore alternates between the maximum and minimum values for even/odd Fock states.

equal to the number of elements in ρ ($N = D^2$), we can directly invert this equation. However, the resulting state matrix may not satisfy physicality constraints (unit trace and positive semidefinite eigenvalues). Additionally, we would like to have the freedom to perform more measurements, as more observations should improve the reconstruction accuracy. Then, an exact solution does not exist, and we instead need to solve a maximum likelihood problem to minimize the error $\|\mathcal{M}|\rho\rangle\rangle - \vec{x}\|$ over all physical ρ . This requires solving the expanded problem:

$$\begin{pmatrix} \mathcal{M}^T \mathcal{M} & |I\rangle\rangle \\ \langle\langle I| & 0 \end{pmatrix} \begin{pmatrix} |\rho_{tr1}\rangle\rangle \\ \lambda \end{pmatrix} = \begin{pmatrix} \mathcal{M}^T \vec{x} \\ 1 \end{pmatrix}, \quad (6.5)$$




Figure 6.2: **Sets of optimized Wigner tomography displacements.** Generated sets of Wigner displacements minimize the condition number κ of the ideal measurement matrix \mathcal{M} . The Hilbert space dimension is truncated at different maximum photon numbers N_{max} , corresponding to maximum displacement magnitudes of roughly $\sqrt{N_{max}}$ marked by the dashed circles. Larger N_{max} require more points to achieve the same κ . The sets are $N_{max} = 7$ with 75 displacements and $\kappa = 1.51$ (orange circles), $N_{max} = 4$ with 75 displacements and $\kappa = 1.04$ (purple squares), and $N_{max} = 3$ with 19 displacements and $\kappa = 1.12$ (black diamonds).

where $|I\rangle\rangle$ is the vectorized identity matrix and λ is a Lagrange multiplier. The first equation from the matrix sets up the minimization problem, while the second enforces the unit trace constraint. We multiply both sides by the inverse of the leftmost matrix to solve for the ρ_{Tr1} matrix, which might not be positive semidefinite. To find the closest (minimum 2-norm distance) physical ρ to our ρ_{Tr1} , we apply the method in [80] where the eigenvalues of ρ_{Tr1} are sorted from high to low and any negative eigenvalues have their weight distributed equally to the others until they are all nonnegative. This way, the trace is preserved and we ultimately produce a physical ρ .

To reduce the total number of measurements required to determine ρ , we optimize the set of α that we choose with the techniques presented in [79]. The optimal set will minimize

the condition number κ of the measurement matrix \mathcal{M} , where

$$\kappa(\mathcal{M}) = \left| \frac{\lambda_{max}(\mathcal{M})}{\lambda_{min}(\mathcal{M})} \right| \quad (6.6)$$

is the ratio of the maximum to minimum eigenvalues of \mathcal{M} , so that $\kappa \geq 1$. The condition number bounds the error amplification, so that initial error ϵ is magnified to no more than $\kappa\epsilon$. Consequently, minimizing the condition number minimizes the error and makes the tomography most efficient [81]. Examples of this optimization are shown in Fig. 6.2. The ideal distribution of measurement displacements depends on the maximum photon number of the state that we expect to observe; that is, states with population at higher Fock states require larger displacements to accurately reconstruct, and need more measurements to achieve the same condition number.

The two components we need to implement the Wigner operator are a well-calibrated parity measurement and a cavity drive with controllable amplitude and phase, which we discussed in Chapters 3 and 5. We need phase control of the cavity drive to implement complex values of α , and we use the qubit to perform the parity measurement. As a reminder, the parity measurement consists of 2 qubit $\pi/2$ pulses with π relative phase separated by a wait time $\tau_{\text{parity}} \approx \pi/(2\chi)$. An additional calibration that can be useful is a Parity Bandwidth Calibration, which is shown in Fig. 6.3. In this experiment, we prepare some state in the cavity, then perform two subsequent $\pi_{ge}/2$ pulses with no intermediate wait time. For one case, we keep them the same phase so that the ideal final state of the qubit is $|e\rangle$, and for the other case, we give them opposite phase so that the qubit should end in $|g\rangle$. However, due to the dispersive interaction between the qubit and cavity, the qubit frequency will shift depending on the cavity population. This may force the $\pi/2$ pulses in the parity measurement off resonance so that they no longer act as true $\pi/2$ pulses. The effect of this will be a reduction in the range of possible averaged parity measurement results. For




Figure 6.3: **Parity measurement bandwidth.** As we increase the magnitude $|\alpha|$ of a coherent state prepared in mode j of the cavity, the qubit frequency will shift by a proportional amount dictated by the dispersive shift χ_j . This can take the $\pi/2$ pulses in the parity measurement off resonance. Then, two subsequent $\pi/2$ pulses with the same phase (blue up arrows), which should ideally always leave the qubit in $|e\rangle$, may leave the qubit partially in $|g\rangle$. Similarly, two $\pi/2$ pulses with opposite phase (orange down arrows) may leave the qubit partially in $|e\rangle$. The y-axis is calibrated such that $|e\rangle$ corresponds to 1 and $|g\rangle$ to 0. To correct this effect, we project the reduced range back to the maximum range of 0 to 1 during the data analysis. This calibration will change depending on both the prepared state and the cavity displacement in the Wigner tomography sequence, as they both contribute to the overall frequency shift.

example, rather than having limits of 0 and 1 (which is later projected to a range of $\pm 2/\pi$ for the Wigner operator), the limits might be reduced to 0.1 to 0.9 for exactly even or odd parity, respectively. This may introduce a source of systematic error to the experiment results. This is especially relevant for Wigner tomography measurements, as we prepare some state that we want to measure in the cavity, then drive it further by a displacement α to enact the Wigner operator. This range reduction can be partially alleviated by using a faster (and thus larger bandwidth) $\pi/2$ pulse. However, we also want to avoid driving too strongly and populating the transmon $|f\rangle$ level. These effects must be balanced to implement the most accurate parity measurement and thus the best Wigner tomography measurement.

6.2 Generalized Multimode Wigner Tomography

We extend Wigner tomography to multiple modes in a seemingly straightforward way. The multimode Wigner operator for M modes is:

$$\hat{\mathcal{W}}(\vec{\alpha}) = \left(\frac{2}{\pi}\right)^M \hat{\mathcal{D}}(\vec{\alpha}) \hat{\Pi} \hat{\mathcal{D}}(-\vec{\alpha}), \quad (6.7)$$

where now $\hat{\mathcal{D}}(\vec{\alpha}) = e^{\sum_j \alpha_j \hat{a}_j^\dagger - \alpha_j^* \hat{a}_j} \equiv \otimes_j \hat{\mathcal{D}}(\alpha_j)$ is the multimode displacement operator and $\vec{\alpha}$ is a vector of complex displacements for each cavity mode. Similarly, $\hat{\Pi}$ is now the joint parity operator $\hat{\Pi} = e^{i\pi \sum_j \hat{a}_j^\dagger \hat{a}_j} = e^{i\pi \sum_j N_j}$. The main challenge in performing multimode tomography is the joint parity measurement. Due to the Ramsey style of our parity measurement sequence, adding a photon in mode m causes the qubit to accrue a phase $\theta_m = \chi_m \tau_R$. For a single mode, we can always choose $\tau_R = \pi/\chi_m$ such that $\theta_m = \pi$ to realize the parity operator. However, in the multimode case, since our χ_m 's are different, there is in general no wait time that corresponds to a joint parity measurement. This is shown for three of the modes in our system in Fig. 6.4. Previous tomography schemes have addressed this issue by using sideband χ -engineering [66], as discussed in Chapter 4, or through the use of higher transmon levels [82]. However, while χ -engineering can change the effective χ_m of some of the modes to match them, applying detuned drives near the $|h0_m\rangle - |e1_m\rangle$ transition could interfere with the system state, as the $|e1_m\rangle$ level will likely be populated during the joint parity measurement. Another concern in our case was that the sideband drives were input through our readout cavity. As a result, the large drive strengths needed to sufficiently shift the χ_m resulted in off-resonantly populating the readout to a level that significantly dephased the qubit ($T_2^* \sim 1 - 6\mu\text{s}$) through their dispersive interaction. Utilizing $|f\rangle$ or other higher transmon levels also did not work in our system. The storage cavity mode frequencies were all located between our qubit and readout cavity frequency, resulting in the $\chi_{ge,m}, \chi_{gf,m}$, and




Figure 6.4: **Parity measurement times for multiple modes** Three modes (orange circles, purple squares, black diamonds) adjacent to each other in our multimode system have different dispersive interactions with the transmon χ_j and thus different parity measurement times, as listed in the legend. The experiment sequence involves preparing a photon in the cavity mode, then performing two $\pi/2$ pulses separated by a wait time τ_R , as described in Chapter 3. The parity measurement times are the π times of the oscillations, and are marked by the dashed vertical lines.

$\chi_{ef,m}$ interaction terms all having the same behavior. That is, if two storage modes $\{i, j\}$ satisfied $|\chi_{ge,i}| > |\chi_{ge,j}|$, then they also satisfied $|\chi_{gf,i}| > |\chi_{gf,j}|$ and $|\chi_{ef,i}| > |\chi_{ef,j}|$. Thus, there was no solution for time t_1 spent in the $|g\rangle - |e\rangle$ space and t_2 in the $|g\rangle - |f\rangle$ or $|e\rangle - |f\rangle$ that would satisfy $t_1|\chi_{ge,i}| + t_2|\chi_{gf,i}| = t_1|\chi_{ge,j}| + t_2|\chi_{gf,j}|$ or $t_1|\chi_{ge,i}| + t_2|\chi_{ef,i}| = t_1|\chi_{ge,j}| + t_2|\chi_{ef,j}|$. We needed to develop a different approach.

To perform our multimode tomography, we relax the condition that the joint parity measurement be perfect, in the sense that each photon in any of the modes contributes phase π to the operator. Instead, we allow different modes to have different phase contributions. We use the fact that a qubit Ramsey experiment for a time τ_R corresponds to the measurement of a known generalized operator $\hat{\Theta} = \cos(\sum_m \theta_m \hat{N}_m)$ acting on the cavity modes, where $\theta_m = \chi_m \tau_R$ need not equal π . We still prefer to keep the angles as close to π as possible, as they offer the largest qubit Ramsey contrast and thus provide the most information about

the state. In the extreme worst case, if $\theta_m = 2\pi$ for one of the modes, this operator would give us no information about the state of that mode. Values of θ_m near π will also be the least sensitive to errors in the calibration of θ_m . Additionally, if $\theta_m = \pi$ for all the modes, this generalized operator reduces to the usual perfect joint parity operator. Replacing the joint parity operator with this one, our generalized multimode Wigner operator becomes

$$\hat{\mathcal{W}}(\vec{\alpha}, \vec{\theta}) = \left(\frac{2}{\pi}\right)^M \hat{\mathcal{D}}(\vec{\alpha}) \hat{\Theta} \hat{\mathcal{D}}(-\vec{\alpha}). \quad (6.8)$$

One potential downside of this change is that the components of this Wigner operator are no longer generalized Laguerre polynomials. However, they can still be computed numerically following the formula given in Eqn. (6.8) and using the definitions of the displacement operator $\hat{\mathcal{D}}(\vec{\alpha})$ and $\hat{\Theta}$. We use this generalized Wigner operator to perform the multimode Wigner tomography measurements presented in the next section.

6.3 Tomography of W states

The main category of states for which we demonstrate our multimode Wigner tomography is W states. These entangled states generalize to any number of modes, as we can see from Eqn. (4.8). We prepare these states using photon blockade and constant cavity drives for up to 5 modes, as discussed in Chapter 4. Here, we present Wigner tomography of these states for up to 4 modes.

We start with the simplest case of a 2-mode W state. This can also be thought of as a Bell state or dual rail state, and has the form

$$|W_2(\phi)\rangle = \frac{1}{\sqrt{2}} (|01\rangle + e^{i\phi} |10\rangle), \quad (6.9)$$

for some angle ϕ . We control the ϕ with which we prepare this state by changing the phase

of one of the cavity drives. To generate a set of Wigner displacements to measure, we apply the methods discussed previously in this chapter to minimize the condition number of the measurement matrix corresponding to a single mode. We choose an $N_{max} = 3$ to allow for leakage past the blockade in the W state preparation, and a total of 19 distinct displacements, as shown by the black points in Fig. 6.2. Then, we take every combination of those displacements over 2 modes, giving us $19^2 = 361$ total pairs of displacements at which to perform the Wigner tomography. The results of the tomography are shown in Fig. 6.5. Since we have 4 quadratures that are the real and imaginary components of the displacement for either of the 2 modes, we have a total of $\binom{4}{2} = 6$ different 2D tomography slices. Using the results of these measurements, we reconstruct the density matrix of our prepared state, as shown in Fig. 6.6, which also compares the reconstructed state to the ideal and simulated states. The simulations include the effects of qubit and cavity decay and dephasing, as well as blockade preparation errors. For simplicity, we present the absolute value of ρ rather than the separate real and imaginary parts, which can be found in Fig. 6.5. We prepare a 2-mode W state with a fidelity of $98.4 \pm 0.5\%$ with $\phi = -0.222$, which is consistent with the simulated fidelity of 98.7%. The main sources of error are leakage outside the blockaded subspace and slightly unequal cavity state populations caused by their dispersive shifts.

We now continue to the 3-mode W state. We can explicitly express this state with phases on its state components as:

$$|W_3(\phi_1, \phi_2)\rangle = \frac{1}{\sqrt{3}} (|100\rangle + e^{i\phi_1} |010\rangle + e^{i\phi_2} |001\rangle). \quad (6.10)$$

We perform measurements at the same set of optimized displacements as in the single and 2-mode cases. This gives us a total of $19^3 = 6859$ total Wigner tomography displacement vectors. The results of the 3-mode tomography are shown in Fig. 6.7, from which we also determine the phases of each multimode Fock state component of the W state. Since there




Figure 6.5: **2-mode W state Wigner tomography.** We show the real and imaginary parts of the density matrix of the reconstructed 2-mode state. $\phi = -0.222$ for the state $(|10\rangle + e^{i\phi}|01\rangle)/\sqrt{2}$ is determined from the maximum fidelity projection, and is shown in the inset with the dashed red line marking the fit value. The corresponding state fidelity is 0.984 ± 0.005 . We also present the 6 orthogonal 2D slices of the two-mode Wigner function, with all combinations of the real and imaginary quadratures of both modes. Maximum values of the Wigner tomography range from $\pm(2/\pi)^2$.

are the real and imaginary parts of the displacements for each of the 3 modes, we have a total of 6 unique quadratures and thus $\binom{6}{2} = 15$ different 2D slices of the phase space. There are two qualitatively different patterns: a ring of slightly positive values around a negative center, similar to the Wigner tomography of a single photon ($|1\rangle$), and a tilted positive-negative-positive fringe symmetric around a rotated line. The exact quadratures that produce each of these plots will change depending on the phase values ϕ_1, ϕ_2 . Reconstructing the state with our Wigner tomography measurements gives us a 3-mode W state fidelity of 0.945 ± 0.012 and $\phi_1 = -2.497, \phi_2 = -2.980$. This agrees with the simulated fidelity of 0.954. A comparison of the prepared, simulated, and ideal W states is shown in Fig. 6.8. Like in the 2-mode case, we show the absolute value of the density matrix rather than the individual real and imaginary parts, which are shown in Fig. 6.7. The majority of the deviation between the ideal and experimentally prepared states is caused by one of the modes, the one indexed




Figure 6.6: **Reconstructed 2-mode W state density matrix.** Absolute value of the 2-mode W state density matrix. Populations are represented with colors ranging from red to blue. The multimode Fock states that should ideally be populated have red labels. Dashed black boxes indicate the ideal W state populations, while yellow boxes show the simulated populations. Simulations use the experimental drive parameters and account for the effects of cavity and transmon decoherence and decay, as well as blockade preparation errors.

by $|001\rangle$, which is likely due to the fact that the χ_j of that mode is farthest away from the frequency of the blockade tone. We also see this effect reflected in the simulation.

For 4 modes, we again use the set of 19 displacements with optimized condition number. This continues to scale exponentially with the number of modes, requiring $19^4 = 130321$ measurements to scan these displacements across all 4 modes. This exponential increase in the size of the problem and the number of required measurements is a feature of the typical formulation for state reconstruction using Wigner tomography [83–85]. In the next section,




Figure 6.7: **3-mode W state Wigner tomography.** We show the real and imaginary parts of the density matrix of the reconstructed two-mode state. The values of $\phi_1 = -2.497$ and $\phi_2 = -2.980$ of the prepared state $(|100\rangle + e^{i\phi_1}|010\rangle + e^{i\phi_2}|001\rangle)/\sqrt{3}$ are determined from the maximum fidelity projection, and is shown in the inset with fit values marked by the dashed red lines. The corresponding state fidelity is 0.945 ± 0.012 . We also present the 15 orthogonal 2D slices of the two-mode Wigner function, with all combinations of the real and imaginary quadratures of the three modes. Which mode quadratures produce the two qualitatively different plots (rings vs. fringes) depends on the values of ϕ_1 and ϕ_2 . Since there are 3 modes, the maximum values of the tomography range from $\pm(2/\pi)^3$.




Figure 6.8: **Reconstructed 3-mode W state density matrix.** Absolute value of the 3-mode W state density matrix. The multimode Fock states that should ideally be populated have red labels. Dashed black boxes indicate the ideal W state populations, while yellow boxes show the simulated populations obtained with the drive parameters used in the experiment after accounting for blockade preparation errors and the effects of cavity and transmon decoherence and decay.

we will see how to reduce this to a polynomial scaling with mode number for our W states. Nevertheless, for now, because of the increasingly large number of required measurements, although we demonstrate preparing W states of up to 5 modes with spectroscopy data, we only perform Wigner tomography on up to 4 modes. Even for 4 modes, because this number of measurements can be prohibitive to collect in a reasonable amount of time, we instead choose a subset of 37000 of these 130321 4-tuples of displacements to measure the Wigner tomography. This subset is chosen randomly from the total ordered list and inspected




Figure 6.9: **Reconstructed 4-mode W state density matrix.** Absolute value of the 4-mode W state density matrix. The multimode Fock states that should ideally be populated have red labels. Dashed black boxes indicate the ideal W state populations, while yellow boxes show the simulated populations obtained with the drive parameters used in the experiment after accounting for blockade preparation errors and the effects of cavity and transmon decoherence and decay. We are able to produce the 4-mode W state with a fidelity of 0.906 ± 0.018 and reconstructed angles $\phi_1 = -1.57$, $\phi_2 = -2.36$, and $\phi_3 = 0.78$.

afterwards to ensure that we are effectively sampling the entire set of 19 displacements for all 4 modes, and not missing some of those displacements or requiring all 19^3 measurements before moving on to the next displacement for one of the modes. This reduction leads to larger error bars in our final measurement, but we are still able to obtain a reasonable result.

In particular, with this subset, we are able to reconstruct our prepared state and measure the fidelity of our 4-mode W state as 0.906 ± 0.018 . The reconstructed angles are $\phi_1 = -1.57$, $\phi_2 = -2.36$, and $\phi_3 = 0.78$ when writing the W state as $W_4 = (|1000\rangle + e^{i\phi_1} |0100\rangle +$




Figure 6.10: **Entanglement witness for W states of 2-4 modes.** The value of the witness over time for 2 modes (purple) is presented over time from the reconstructed density matrices from Wigner tomography measurements (circles) and from simulation (line). For 3 (orange, star) and 4 (green, triangle) modes, we perform the tomography only at the time corresponding to the W state preparation, although we plot the simulated value (lines) over the whole time range. The measured values agree with the simulation for all cases.

$e^{i\phi_2} |0010\rangle + e^{i\phi_3} |0001\rangle)/\sqrt{4}$. The absolute value of the density matrix is shown in Fig. 6.9. The fidelity is lower than the 3-mode and 2-mode cases, as we would expect, since the χ_j 's of the four modes span a larger range than for lower mode numbers, which could result in a less effective blockade of one of the modes by the dressing drive tone. We see this in the reconstructed density matrices, as the elements corresponding to one of the modes deviates more from the ideal values than the others. In the 3-mode case, this is the rightmost mode index, while in the 4-mode case, this is the leftmost mode index. This is consistent with what we expect, as the modes' χ_j change monotonically, so the least effectively blockaded χ_j should be on either extreme relative to the blockade dressing tone.

We verify that we have indeed generated entangled W states by measuring an entanglement witness [86, 87]. For multimode states, this metric measures the amount of entangle-

ment generated. For an M -mode W state, this is given by

$$\mathcal{W}_M = \frac{M-1}{M} - \mathcal{F}_M, \quad (6.11)$$

where $\mathcal{F}_M = \text{Tr}(|W_M\rangle\langle W_M|\rho)$ is the overlap fidelity between our prepared state and the ideal M -mode W state. Entangled states will have $\mathcal{W}_M < 0$. We show this in Fig. 6.10 for two modes over the course of one W state preparation oscillation, and three and four modes at the time when the 3-mode or 4-mode W state is prepared. We can see that we have indeed generated our entangled W states in the regions where the points are negative, and that our measured values of the witness are in good agreement with the simulated values for all three mode numbers, which are obtained from evaluating Eqn. (6.11) for the simulated states over time. Consequently, our multimode Wigner tomography allows us to characterize our prepared states and quantitatively test their agreement with ideal entangled W states.

6.4 Efficient Tomography Sampling: DEMESST

As we have seen, Wigner tomography provides full state information, but in its usual formulation requires a number of measurements that scales exponentially with the Hilbert space size and thus the number of modes [88, 89]. Matrix inversion-based methods like the one we used in the previous section require at least D^2 observations, where $D = N^M$ is the dimension of the Hilbert space for M modes with population in up to N levels. This produces the exponential scaling with mode number M that we seek to avoid. In this section, we present a Wigner tomography sampling method that scales polynomially with subspace size and thus the number of modes for certain states, like our W states. In the following discussion, we refer to the inversion-based Wigner tomography method that scales exponentially as W_{inv} . We primarily use it as a baseline for comparison.

A number of proposed theoretical methods are able to extract multimode state information while circumventing the observation number scaling exponentially with mode number. These include techniques that apply additional unitaries between modes as part of the measurement process [90], perform local measurements with polynomial post-processing [90], make use of ancillary modes [91], or employ operators based on the ability to implement excitation counting [92]. However, none of these approaches are based solely on Wigner tomography. We add to this space of efficient measurement schemes with an approach based on Wigner tomography.

6.4.1 Outlining the DEMESST Sampling Method

Here, we introduce the Direct Extraction of (Density) Matrix Elements from Subspace Sampling Tomography (DEMESST) method and describe how it can demystify quantum states with a sampling requirement that scales polynomially with the dimension of the subspace [93]. For certain multimode states, the total measurement number will therefore depend polynomially on the number of modes, rather than exponentially. This can be especially advantageous when a state lives in a subspace of interest that is much smaller than the full space.

The procedure for DEMESST is as follows: first, the Wigner function corresponding to a matrix element in the chosen basis, such as the multimode Fock state basis, is normalized to a probability distribution based on its absolute value, so that negative regions have positive weight. Then, displacement vectors to apply to the bosonic cavity modes are sampled from the resulting cumulative distribution function (CDF) by uniformly selecting a value between 0 and 1 and inverting to find the corresponding angular and radial values of the displacements. This calculation can be performed efficiently by utilizing Laguerre functions and their inverses. Additionally, each displacement vector will have the original sign of its Wigner function value preserved, so that if the Wigner function was negative at that point, the final measured value will be multiplied by -1. Observing a set of these will provide an




Figure 6.11: **DEMESTT tomography sampling method schematic and theoretical scaling.** (a) Schematic representing the advantage of the DEMESTT method. Rather than sampling an entire multimode Hilbert space (3D axes), if a state lives in some number of subspaces (blue plane), the sampling can be done on each of those instead. The \hat{O} basis operators can be of the form $\hat{O}_{\vec{n}, \vec{n}'} \propto |\vec{n}\rangle \langle \vec{n}'| + h.c.$ for basis states $|\vec{n}\rangle, |\vec{n}'\rangle$. This can improve the overall efficiency of the sampling, especially for states with support across large numbers of modes. (b) Number of measurements required for the DEMESTT (purple, circles) and W_{inv} (orange, squares) methods to reach a 90% state reconstruction fidelity on W states of up to 7 modes, assuming perfect state preparation. W_{inv} scales exponentially with the size of the Hilbert space (and therefore the number of modes M), while DEMESTT scales only polynomially.

estimate of the chosen density matrix element. Repeating this for multiple elements will thus produce the density matrix of the prepared state.

The DEMESTT tomography method scales polynomially for states with a known maximum excitation number, rather than exponentially with the number of modes. This is accomplished by leveraging that information, and rather than sampling displacements corresponding to all possible basis states (such as the multimode Fock basis), only sampling points corresponding to those that are expected to support the state. For example, if a 3-mode state has a maximum of 2 photons, we can reconstruct the unknown state by measuring the matrix elements associated with the basis of the subspace, namely $|000\rangle, |001\rangle, |011\rangle, |002\rangle$, and permutations (as schematically illustrated in Fig. 6.11(a)). These matrix elements can be obtained by sampling methods similar to direct fidelity estimation [88, 94], and without

introducing bias. This way, we avoid extracting irrelevant information about unnecessarily high photon number states (such as $|111\rangle$ or $|012\rangle$ in this example), thereby improving the efficiency and lowering the total measurement number required for an accurate result. Although each set of measurements estimates a single basis element of the multimode density matrix, if that basis has polynomial size, the overall number of observations required for an accurate state reconstruction is similar to the W_{inv} approach for 2 modes, while demonstrating a noticeable improvement at larger mode numbers (≥ 3).

Furthermore, measuring independent density matrix elements can reduce the size of the sampling problem to a lower number of modes or smaller dimension when some of the modes are in their vacuum state. We accomplish this by projecting out those unpopulated modes. For example, for the 3-mode Fock state $|001\rangle\langle 010|$, we can project out the first mode and reduce the measurement to the 2-mode state $|01\rangle\langle 10|$. This allows us to reduce the total number of observations required to estimate each of the density matrix elements in our multimode W state reconstructions.

Another advantage of the DEMESST sampling method when compared to W_{inv} is its self-consistency. In particular, individual density matrix elements for any multimode state can be measured independently, without needing to choose a cutoff maximum photon number or Hilbert space size that could subject the reconstructed state to inversion errors. This eliminates the risk of obtaining an inaccurate tomography result if, for example, the prepared state contains population beyond the space spanned by the chosen basis for the W_{inv} sampling. We demonstrate this effect later in this section, in Fig. 6.16.

For different mode number M , we simulate Wigner tomography measurements of M -mode W states. W states are excellent candidates for testing our tomography sampling methods because they have a clearly defined photon number and are irreducible multimode states that generalize straightforwardly to any number of modes. The theoretical performance of DEMESST and W_{inv} is shown in Fig. 6.11(b). Assuming perfect state preparation, we

compare the number of observations required to accurately reconstruct the W state with 90% fidelity. In contrast with the exponential scaling of W_{inv} , the DEMESST methods scales polynomially with M , an advantage which becomes clear at large M . For two modes, W_{inv} actually performs better due to some overhead required for the DEMESST approach. However, for larger M , DEMESST requires fewer measurements to converge to the same level of fidelity, and scales much more efficiently than W_{inv} .

6.4.2 DEMESST Derivation

Here, we derive the advantageous polynomial scaling with mode number of DEMESST [93]. We consider a system with M modes and maximum total photon number N between those modes. This restricts the dimension of the Hilbert space to $\binom{M+N}{N}$. We focus on the scaling of the sampling overhead (number of samples required) vs. mode number M , in the limit where M is much larger than $2N$, and show that this overhead scales polynomially vs. M with bounded photon number N , demonstrating the efficiency of the DEMESST approach as M increases.

Let \hat{O} be the basis operator whose Wigner function we measure. Like for the multimode Fock basis, we assume that \hat{O} can be expressed in one of the forms:

$$\begin{aligned}
O_{\vec{n},\vec{n}} &= |\vec{n}\rangle \langle \vec{n}|, \\
O_{\vec{n},\vec{n}'}^R &= \frac{|\vec{n}\rangle \langle \vec{n}'| + |\vec{n}'\rangle \langle \vec{n}|}{\sqrt{2}} \quad (\vec{n} \neq \vec{n}'), \\
O_{\vec{n},\vec{n}'}^I &= i \frac{|\vec{n}\rangle \langle \vec{n}'| - |\vec{n}'\rangle \langle \vec{n}|}{\sqrt{2}} \quad (\vec{n} \neq \vec{n}').
\end{aligned} \tag{6.12}$$

Here, $|\vec{n}\rangle = \otimes_{m=1}^M |n_m\rangle$ and $|\vec{n}'\rangle = \otimes_{m=1}^M |n'_m\rangle$ are Fock basis states that satisfy $\sum_{m=1}^M n_m \leq N$ and $\sum_{m=1}^M n'_m \leq N$. For a system with M modes and maximum total photon number N , we have $\binom{M+N}{N}^2$ of these operators. One essential observation is that, when $M > 2N$, for any (\vec{n}, \vec{n}') pair, there are at least $(M - 2N)$ elements in the set $S \equiv S_{(\vec{n}, \vec{n}')} = \{m | n_m = n'_m = 0\}$.

We also denote $\bar{S} = \{1, 2, \dots, M\} \setminus S$. Because of this, we can decompose the basis operator:

$$\hat{O} = \left(\bigotimes_{m \in S} |0\rangle \langle 0|_m \right) \otimes \hat{O}_{\bar{S}}, \quad (6.13)$$

where $\hat{O}_{\bar{S}}$ has support on the modes with index $m \in \bar{S}$. The number of elements in \bar{S} is no greater than $2N$, and independent of M . We now denote

$$\widehat{\mathcal{W}}_{\rho}(\vec{\alpha}, \vec{\theta}) = \text{Tr} \left[\hat{\mathcal{D}}(\vec{\alpha}) \cos \left(\phi(\vec{\alpha}) + \sum_m \theta_m \hat{N}_m \right) \hat{\mathcal{D}}(-\vec{\alpha}) \rho \right], \quad (6.14)$$

such that $\widehat{\mathcal{W}}$ is the generalized Wigner function from Eqn. (6.8) without the $2/\pi$ factors and an additional phase ϕ that is acted on ρ . Now, the generalized Wigner function of such an operator \hat{O} satisfies

$$\widehat{\mathcal{W}}_{\hat{O}}(\vec{\alpha}, \vec{\theta}) = \left(\prod_{m \in S} \widehat{\mathcal{W}}_{|0\rangle\langle 0|}(\alpha_m, \theta_m) \right) \cdot \widehat{\mathcal{W}}_{\hat{O}_{\bar{S}}}(\vec{\alpha}_{\bar{S}}, \vec{\theta}_{\bar{S}}), \quad (6.15)$$

where $\vec{\alpha}_{\bar{S}}, \vec{\theta}_{\bar{S}}$ contain those elements in $\vec{\alpha}, \vec{\theta}$ whose mode index $m \in \bar{S}$. Again, $\widehat{\mathcal{W}}_{\hat{O}_{\bar{S}}}(\vec{\alpha}_{\bar{S}}, \vec{\theta}_{\bar{S}})$ is independent of M .

The expectation value of our state ρ with a basis operator \hat{O} is given by $\text{Tr}(\rho \hat{O})$. We want to determine the number of measurements required to obtain a precise estimate of $\text{Tr}(\rho \hat{O})$. We can express this quantity [88] as:

$$\begin{aligned} \text{Tr}(\rho \hat{O}) &= \int \frac{d^{2M} \vec{\alpha}}{\pi^M} \widehat{\mathcal{W}}_{\rho}(\vec{\alpha}, -\vec{\theta}) \widehat{\mathcal{W}}_{\hat{O}}(\vec{\alpha}, \vec{\theta}) \\ &= C_M \int d^{2M} \vec{\alpha} \widehat{\mathcal{W}}_{\rho}(\vec{\alpha}, -\vec{\theta}) \widehat{\mathcal{W}}_{\hat{O}}(\vec{\alpha}, \vec{\theta}), \end{aligned} \quad (6.16)$$

where $C_M = \prod_{m=1}^M (2(1 - \cos \theta_m)/\pi)$ is a prefactor that comes from using the generalized Wigner function instead of the typical ideal one. We rewrite this in the form of a probability

distribution as

$$\text{Tr}(\rho\hat{O}) = C_M \int d^{2M}\vec{\alpha} p(\vec{\alpha}) \frac{\widehat{\mathcal{W}}_{\hat{O}}(\vec{\alpha}, \vec{\theta})}{p(\vec{\alpha})} \widehat{\mathcal{W}}_{\rho}(\vec{\alpha}, -\vec{\theta}). \quad (6.17)$$

We can model a measurement as a set of binomial outcomes $A_j^{(k)} \in \{1, -1\}$, where k denotes indices of different displacement vectors $\vec{\alpha}$ and j denotes repetitions at each displacement. Then, we obtain an estimate of $\text{Tr}(\rho\hat{O})$ by measuring an average value of $C_M \frac{\widehat{\mathcal{W}}_{\hat{O}}(\vec{\alpha}^{(k)}, \vec{\theta})}{p(\vec{\alpha}^{(k)})} A_j^{(k)}$. For the DEMESST method, we choose

$$p(\vec{\alpha}) = \frac{|\widehat{\mathcal{W}}_{\hat{O}}(\vec{\alpha}, \vec{\theta})|}{Z_{\hat{O}}}, \quad (6.18)$$

where $Z_{\hat{O}} = \int d^{2M}\vec{\alpha} |\widehat{\mathcal{W}}_{\hat{O}}(\vec{\alpha}, \vec{\theta})|$. Therefore, we average $C_M Z_{\hat{O}} A_j^{(k)}$ over sampling vectors $\vec{\alpha}^{(k)}$ and the possible binomial outcomes from the qubit measurement for each sampling vector. In the limit where we do one qubit measurement per $\vec{\alpha}^{(k)}$, we can use Hoeffding's inequality to estimate the number of samples N_{spl} required to reach an accuracy ϵ_1 with probability $1 - \delta_1$ to be

$$P\left(\left|\frac{C_M Z_{\hat{O}}}{N_{\text{spl}}} \sum_{k=1}^{N_{\text{spl}}} A^{(k)} - \text{Tr}(\rho\hat{O})\right| \geq \epsilon_1\right) \leq \delta_1 \quad (6.19)$$

when

$$N_{\text{spl}} \geq \left\lceil \frac{2C_M^2 Z_{\hat{O}}^2}{\epsilon_1^2} \ln(2/\delta_1) \right\rceil. \quad (6.20)$$

We can see that in general, $N_{\text{spl}} \propto (C_M Z_{\hat{O}})^2$, so we focus on that quantity. We have

$$\begin{aligned} C_M Z_{\hat{O}} &= \left(\prod_{m=1}^M \frac{2(1 - \cos \theta_m)}{\pi} \right) \int d^{2M}\vec{\alpha} |\widehat{\mathcal{W}}_{\hat{O}}(\vec{\alpha}, \vec{\theta})| \\ &= \prod_{m \in S} \left(\frac{2(1 - \cos \theta_m)}{\pi} \int d^2\vec{\alpha} |\widehat{\mathcal{W}}_{|0\rangle\langle 0|}(\alpha_m, \theta_m)| \right) \cdot C_{\bar{S}} \int d^{2|\bar{S}|}\vec{\alpha}_{\bar{S}} |\widehat{\mathcal{W}}_{O_{\bar{S}}}(\vec{\alpha}_{\bar{S}}, \vec{\theta}_{\bar{S}})| \\ &= 2^{M-|\bar{S}|} C_{\bar{S}} Z_{O_{\bar{S}}}, \end{aligned} \quad (6.21)$$

where $C_{\bar{S}} = \prod_{m \in \bar{S}} [2(1 - \cos \theta_m)/\pi]$. Since $\hat{O}_{\bar{S}}$ is supported on at most $2N$ modes, which is

independent of M when $M > 2N$, the only M -dependence in $C_M Z_O$ comes from the 2^M factor. However, this still scales exponentially. To resolve this, we introduce the projection operator $\hat{P}_S = \otimes_{m \in S} |0\rangle\langle 0|_m$ and denote $\rho_{\bar{S}} = \text{Tr}_S(\rho \hat{P}_S)$, where $\text{Tr}_S(\bullet)$ indicates the partial trace over all modes with $m \in S$. Since $\hat{O}_{\bar{S}}$ is supported on the modes contained in \bar{S} , we have $\text{Tr}(\rho \hat{O}) = \text{Tr}(\rho_{\bar{S}} \hat{O}_{\bar{S}})$. Here, $\rho_{\bar{S}}$ and $\hat{O}_{\bar{S}}$ are entirely supported on modes in \bar{S} , which contains at most $2N$ elements. Thus, we can reduce the problem from estimating $\text{Tr}(\rho \hat{O})$ to instead estimating $\text{Tr}(\rho_{\bar{S}} \hat{O}_{\bar{S}})$. In the experiment, we use projection to the transmon $|f\rangle$ level to restrict the cavity state to $\rho_{\bar{S}}$. Now, we again use Hoeffding's inequality, which gives us

$$P \left(\left| \frac{C_{\bar{S}} Z_{\hat{O}_{\bar{S}}}}{N_{\text{spl}}} \sum_{k=1}^{N_{\text{spl}}} A^{(k)} - \text{Tr}(\rho_{\bar{S}} \hat{O}_{\bar{S}}) \right| \geq \epsilon_2 \right) \leq \delta_2 \quad (6.22)$$

when

$$N_{\text{spl}} \geq \left\lceil \frac{2C_{\bar{S}}^2 Z_{\hat{O}_{\bar{S}}}^2}{\epsilon_2^2} \ln(2/\delta_2) \right\rceil. \quad (6.23)$$

When N is bounded, $C_{\bar{S}} Z_{\hat{O}_{\bar{S}}}$ is independent of mode number M when $M > 2N$. Therefore, N_{spl} in Eqn. (6.23) scales as

$$N_{\text{spl}} \sim \mathcal{O}_M \left(\frac{f(N)}{\epsilon_2^2} \ln(2/\delta_2) \right), \quad (6.24)$$

where \mathcal{O}_M indicates that we focus only on the scaling over M in the large M limit, and $f(N) = 2C_{\bar{S}}^2 Z_{\hat{O}_{\bar{S}}}^2$ is a function that depend solely on N and the specific form of \hat{O} from Eqn. (6.12).

To reconstruct the density matrix, we must perform the above sampling for all operators \hat{O} with form given in Eqn. (6.12). Because the dimension of the Hilbert space is restricted to $\binom{M+N}{N}$, we have at most $\binom{M+N}{N}^2$ of these operators. Consequently, the total measurement number N_{spl} will be scaled by this factor, which depends polynomially on M . Therefore, we are able to fully reconstruct our state without an exponential sampling requirement with the

number of modes, and instead only a polynomial one. In the next section, we demonstrate this experimentally on W states from 2–4 modes, which satisfy the $M > 2N$ condition for 3 and 4 modes.

6.4.3 Experimental and Simulated Results

The tomography sampling methods use the same generalized Wigner tomography sequence shown in Fig. 6.12(a), where each cavity mode is displaced before performing a generalized parity measurement on the transmon [27]. For the DEMESST method, depending on the number of modes and the state or density matrix element being measured, there may also be a conditional pi pulse followed by an ef pi pulse on the transmon. These pulses perform the projection operation that reduces the size of the sampling problem. Besides these possible pi pulses, the measurement is the same as our previously discussed generalized Wigner tomography. In Fig. 6.12(b), we show some of the points that were sampled for the two methods for one of the modes. The locations in the displacement phase space where the points are most concentrated depends on the basis state being measured for the DEMESST approach, while the W_{inv} attempts to measure the entire basis at once and thus has points more heavily concentrated around rings corresponding to Fock state numbers.

We compare the state reconstruction results obtained from the DEMESST method with W_{inv} , which we present first and use as a baseline. Due to the time difference between performing these experiments and the ones presented previously, the fidelity numbers are slightly different. Using the W_{inv} approach, we measure W state fidelities of 0.966 ± 0.005 , 0.949 ± 0.004 , and 0.912 ± 0.007 for the 2-, 3-, and 4-mode W states, respectively. These are consistent with or very close to the simulated fidelities obtained from modeling state preparation with the blockade, which are 0.971, 0.956, and 0.912 for the 2-, 3-, and 4-mode cases, respectively. The simulations include error sources such as transmon or cavity decoherence and blockade imperfections such as leakage outside of the blockaded subspace.




Figure 6.12: **DEMESST tomography sequence and sampling points.** (a) Wigner tomography pulse sequence for the DEMESST sampling method. Initial cavity displacements and a final generalized multimode parity measurement implement the tomography, while optional conditional pi pulses (gray) project target modes out of the transmon $|g\rangle - |e\rangle$ subspace for the DEMESST approach. An additional angle ϕ can be applied between the $\pi/2$ pulses of the parity measurement to project it onto the real axis. (b) Cavity displacement plots for the W_{inv} (orange) and DEMESST (purple) sampling methods. The W_{inv} has ring features corresponding to measurement of larger Fock states, while DEMESST point density depends on the basis state being measured.

For the DEMESST approach, we reconstruct the density matrices for 2–4 modes by measuring every multimode Fock state with up to 2 photons, but with at most one in a single mode. For 2 modes, these are the Wigner functions corresponding to states of the form: $\{|i\rangle\langle j|, j \in \{00, 01, 10, 11\}\}$. The basis states for larger mode numbers are similar, but padded with more zeros. We measure sets of states of this form, instead of just single photon states, to capture state preparation imperfections like leakage outside of the blocked subspace. These observations directly provide the values of each density matrix element. From the density matrix, we obtain the phase angles ϕ_j of our prepared W states by calculating the phase angle value that best matches the resulting data. These angles are checked




Figure 6.13: **DEMESST and W_{inv} W state density matrix reconstructions.** Absolute value of the final density matrices determined using the DEMESST tomography sampling method (top row) and the W_{inv} method (bottom row) for W states of (a) 2, (b) 3, and (c) 4 modes. The results for DEMESST and W_{inv} are in good agreement.

to verify that they match with the ones obtained from the W_{inv} approach. These fits allow us to obtain our final fidelity numbers.

Our reconstructed density matrices are shown in Fig. 6.13. We plot the absolute values of the density matrix elements so that we can investigate a single matrix for each combination of tomography method and mode number. The two methods are in good agreement, with the largest visible deviation being in the 3-mode case for Fock basis elements with nonzero population in the second (middle) cavity mode. Nevertheless, the Frobenius norm distances between the final matrices obtained from the two methods is low, at 0.05 for the 2-mode case, 0.22 for the 3-mode case, and 0.30 for the 4-mode case. These are all below the corresponding minimum distances at our maximum total measurement number when splitting our data into independent sets. Consequently, these differences do not significantly affect those results,

which we show in Fig. 6.14.

Besides some slight deviations in the measured populations of individual density matrix elements, the remaining distance between the final reconstructed matrices can be explained by small differences in the fit phase angles of the W states. For the 2-mode W state $|W_2\rangle = (|10\rangle + e^{i\phi} |01\rangle)/\sqrt{2}$, we measure $\phi_D = 0.04$ using the DEMESST method, and $\phi_{\text{inv}} = 0.03$ using the W_{inv} method. In the 3-mode case, for the W state $|W_3\rangle = (|100\rangle + e^{i\phi_1} |010\rangle + e^{i\phi_2} |001\rangle)/\sqrt{3}$, we measure $\phi_{1,D} = -0.19$ and $\phi_{2,D} = 1.57$, while $\phi_{1,\text{inv}} = -0.12$ and $\phi_{2,\text{inv}} = 1.57$. Similarly to what we see in the populations, the deviation is primarily caused by the middle mode. Finally, for 4 modes and $|W_4\rangle = (|1000\rangle + e^{i\phi_1} |0100\rangle + e^{i\phi_2} |0010\rangle + e^{i\phi_3} |0001\rangle)/\sqrt{4}$, we find $\phi_{1,D} = -1.36$, $\phi_{2,D} = -2.90$, and $\phi_{3,D} = 0.60$, while $\phi_{1,\text{inv}} = -1.38$, $\phi_{2,\text{inv}} = -3.02$, and $\phi_{3,\text{inv}} = 0.63$. These angles are obtained by discretely sweeping the ϕ_j values over the full 2π range for each of the modes and determining which set of ϕ_j gives the largest fidelity when compared to an ideal W state with those phases.

We now compare the performance of DEMESST and W_{inv} vs. the number of sampled displacements. The results are shown in Fig. 6.14, in the form of reconstructed fidelity ((a)–(c)) and Frobenius norm matrix distance ((d)–(f)) from the final reconstructed state vs. measurement number for both the DEMESST and W_{inv} sampling methods. The final fidelities using the DEMESST approach at the maximum observation number are 0.96 ± 0.01 , 0.955 ± 0.004 , and 0.911 ± 0.007 for the 2-, 3-, and 4-mode W states, respectively, which are consistent with the W_{inv} results. Error bars are obtained from the results of 10 independent sets of 10 repetitions of tomography measurements for each sampled displacement. In other words, we took 100 observations at each displacement vector, split that into 10 distinct groups, and analyzed each of those data sets independently. Therefore, the number of distinct displacements for the W_{inv} method equals the Total Measurement Number shown on the x-axis divided by 10. For the DEMESST method, the number of distinct displacements is further divided by the number of different measured basis elements. Given a fixed total




Figure 6.14: **Fidelities and matrix distances for DEMESST and W_{inv} sampling methods vs. measurement number.** We sampled approximate entangled W states of 2–4 modes for different numbers of displacements. The top panels (a)–(c) show fidelities vs. an ideal W state for 2–4 modes, with dashed horizontal lines indicating the final converged fidelity obtained from the W_{inv} method. These final fidelities are, for 2–4 modes, 0.966 ± 0.005 , 0.949 ± 0.004 , and 0.912 ± 0.007 for W_{inv} and 0.96 ± 0.01 , 0.954 ± 0.004 , and 0.911 ± 0.007 for DEMESST, and are in good agreement. The bottom panels (d)–(f) show Frobenius norm matrix distances between the state at a given measurement number vs. the final measured state. The convergence rates are close to $1/\sqrt{x}$ or a power of -0.5 , as expected. As the mode number increases, the DEMESST method performs increasingly more efficiently by requiring fewer measurements to reach a given level of convergence or error threshold.

measurement number (number of distinct sets of cavity displacements, times the number of repetitions of each displacement), we would gain the most information from maximizing the number of distinct displacements and measuring each a single time. We choose instead to average each measurement 10 times due to our imperfect readout fidelity. This allows us to minimize our average number while still being able to obtain accurate results. We then choose to repeat this process 10 times to obtain statistics, resulting in a total of 100 averages for each distinct displacement. This choice lets us balance this theoretical maximum of information from singleshoot measurements with our measurement errors. For the 2-mode W

state, the two methods perform similarly, while for 3 and 4 modes, the DEMESST begins to perform better than W_{inv} with faster convergence to the final state fidelity.

This improvement is most evident in the matrix distance comparisons. The distances are computed using the Frobenius norm, with error bars again obtained from 10 independent sets of 10 repetitions for each displacement. The final state density matrix against which the distances are computed is obtained by considering all 100 repetitions at once (Fig. 6.13), which is why the final distances do not completely vanish. The behavior of both sampling methods is very similar for the 2-mode W state. However, for the 3-mode case, DEMESST has noticeably faster convergence vs. total measurement number x , in the sense that the matrix distance d to the final state is smaller, as seen by the fit coefficient to $d = ax^b$ (1.1 ± 0.1 for DEMESST vs. 1.71 ± 0.02 for W_{inv}). This effect is further enhanced in the 4-mode case. The ratio of the W_{inv} to DEMESST fit values scales roughly geometrically, from 0.96 ± 0.02 to 1.55 ± 0.1 to 2.3 ± 0.2 for 2,3, and 4 modes, respectively. This reflects the exponential scaling of W_{inv} compared to the polynomial scaling of DEMESST vs. total measurement number. In all cases, the distances fall off roughly as $x^{-1/2}$, as expected.

To further confirm these results, we perform simulations of the matrix distances shown in Fig. 6.14(d)–(f). In particular, we obtain the density matrix of an imperfect W state prepared by photon blockade, with errors from transmon and cavity decoherence and leakage through the blockade. We then simulate Wigner tomography measurements on that state using the same cavity displacement samples that we applied to our experiment, and reconstruct the density matrix while including readout and bit flip error. These simulated results are shown in Fig. 6.15. As expected, the DEMESST method performs increasingly more efficiently with larger mode number. Overall, the simulated and experimental results are similar, albeit with some slight differences. For example, the measurement number at which the simulated distance reaches roughly 0.1 for 2 modes is slightly less than 10^5 , while we observe a distance slightly above 0.1 at that point number in our data. For the 3-mode case,




Figure 6.15: **Simulated matrix distances for vs. measurement number for DEMESST and W_{inv} .** (a) DEMESST and (b) W_{inv} . The distance is computed as the Frobenius norm between the reconstructed state at a given measurement number and a final simulated W state with state preparation errors from photon blockade and decoherences. Error bars are obtained from repeating the simulation multiple times while including readout bit flip errors. The distances decrease to lower values more quickly for the DEMESST approach, especially for larger mode numbers.

in the experiment, we observe a distance of roughly 0.3 at 10^5 measurements for DEMESST and 0.4 for W_{inv} . These values are close to the simulated distances of roughly 0.25 and 0.3, respectively. For the 4-mode case, we measure a distance of 0.4 at roughly 2×10^5 measurements for DEMESST, compared to a simulated distance of roughly 0.2, and a distance of roughly 1.0 vs 0.8 for W_{inv} at that measurement number. We attribute the discrepancies to fluctuations over time in the readout error that may affect the accuracy of our simulation. This effect is particularly pronounced for the 4-mode case, where more measurements are required, forcing the measurement process to last multiple days. Nevertheless, the trend is clear: the DEMESST method requires fewer measurements for the same level of accuracy as the W_{inv} approach, especially for larger mode numbers.

Additionally, we verify that the DEMESST tomography sampling method leads to self-




Figure 6.16: **Self-consistency verification for the DEMESST method.** We investigate the trace vs. point number for prepared (a) 2-mode, (b) 3-mode, and (c) 4-mode W states. Error bars are shown for every fourth point, as well as the final one. As expected, the traces converge to values near unity. (d) Trace vs. point number when measuring only a 2-mode subspace of a prepared 4-mode W state. Due to only measuring half of the populated space, the trace converges to 0.5. This demonstrates that the DEMESST sampling method is self-consistent and does not depend on the chosen measurement subspace.

consistent measurement results. For this purpose, we investigate the traces of our prepared W states to compare them with unity, a property required for physical states. This test allows us to confirm that our prepared state indeed lives in our chosen measured Hilbert space. The results are shown in Fig. 6.16, in the form of average trace vs. observation number. Like before, the averages are taken over 10 independent sets of 10 measurement repetitions for each sampled displacement in the tomography. We find that in all cases, the observed traces are near one, and get closer with more samples. Deviations from unity for low measurement numbers can be attributed to noise and statistics, while the final traces being slightly less than one can be attributed to imperfect state preparation that produces population outside our measured subspace. We perform a further check by considering only a 2-mode subspace of a prepared 4-mode W state. This is shown in Fig. 6.16(d). As expected, the measured trace converges to a value near 0.5, as we are effectively observing half of the total state. This demonstrates that the DEMESST method does not rely on the choice of measured

space, and can provide accurate results independently. In particular, we can identify when we have measured insufficient basis elements to fully characterize a state, such as when the state lives partially (or entirely) outside our measured space. This is a useful capability that lets us detect when we have made incorrect assumptions about our state preparation.

6.4.4 W^2 Tomography Sampling

As another consistency confirmation, we implement another Wigner tomography sampling method: the W^2 method, which was first proposed in [88]. In this approach, sets of coherent cavity displacements α_i are chosen using rejection sampling. This approach computes the overlap between a prepared state and a desired ideal state ρ_I , and can be used for direct fidelity estimation. For M modes, a cutoff c and a displacement vector $(\beta_1, \dots, \beta_M)$ is randomly sampled from a uniform distribution between 0 and a maximum value of $|\mathcal{W}(\alpha_1, \dots, \alpha_M)|^2 \prod_{i=1}^M |\alpha_i|$, where \mathcal{W} denotes the Wigner function of ρ_I . We keep the vector if $|\mathcal{W}(\beta_1, \dots, \beta_M)|^2 \prod_{i=1}^M |\beta_i| > c$. This ensures that we measure cavity displacements that provide the most information about the state, while also avoiding displacements with large magnitude or Wigner values near zero, which are more susceptible to experimental errors. After measuring a set of n of these vectors, the final overlap fidelity is computed as $\frac{1}{n} \sum_{i=1}^n \mathcal{W}_{\text{exp}}(\vec{\alpha}_i) / \mathcal{W}_{\text{ideal}}(\vec{\alpha}_i)$. The W^2 method can be used in a similar manner to the DEMESST, where the fidelity estimation is performed with respect to multimode Fock state basis elements. Repeating for multiple elements measures a reconstructed density matrix.

Experimentally, we use the W^2 method for direct fidelity estimation of our prepared W states. We set the ideal state to be the multimode W state with ϕ 's determined from the DEMESST and W_{inv} methods. The results are shown in Fig. 6.17. The fidelities at the maximum observation number are 0.972 ± 0.013 , 0.95 ± 0.35 , and 0.90 ± 0.08 for the 2-, 3-, and 4-mode W states respectively. These averages are consistent with the results of the previous DEMESST and W_{inv} methods, with the W_{inv} fidelities indicated by the




Figure 6.17: **Experimental results for the W^2 tomography sampling method.** (a) State reconstruction fidelities for 2- (purple circles), 3- (orange triangles), and 4-mode (blue squares) W states. Horizontal lines indicate the fidelities obtained from the W_{inv} method, which are consistent with the W^2 . (b) Error magnitudes for the fidelities shown in (a). All results follow a roughly $1/\sqrt{x}$ relationship vs. total measurement number, as expected.

horizontal lines in Fig. 6.17(a). The data converges quickly to the expected fidelity obtained from those two approaches, although with large uncertainties, as shown in Fig. 6.17(b). One reason for these errors is the relatively low total measurement number compared with the other methods. However, an odd behavior is that the 3-mode data has much greater uncertainties than even the 4-mode case, when we would expect the uncertainties to increase monotonically with mode number. Some possible explanations for this behavior could be a particularly low readout fidelity during data collection or fluctuations in drive strength during the measurement sequence that modify the effective Wigner operator differently for distinct sets of cavity mode displacements. These might have affected the 3-mode data but

not the 4-mode due to the measurements being performed on different days. Another reason could be the choice in cutoff c , which will introduce some amount of systematic bias that could lead to the massive error bars [88, 93]. All the uncertainties have the expected $1/\sqrt{x}$ dependence on total observation number.

In summary, in this chapter, we have discussed Wigner tomography, a prominent technique for performing QST to determine quantum states. To assist with implementing the measurements on multimode systems with non-uniform parameters, we have developed a generalized version of the Wigner operator that can be applied flexibly. With a focus on multimode states, and using W states as an example class of states, we have introduced the DEMESST sampling method and applied it to multimode Wigner tomography. DEMESST is most appropriate for multimode states that have population contained in a few basis elements of an overall Hilbert space. DEMESST outperforms traditional optimized inversion-based methods like W_{inv} for states of that type, by scaling polynomially rather than exponentially with subspace size and thus mode number. We have demonstrated that improvement on W states of 2–4 modes. Here, we have presented comparisons using the multimode Fock basis on W states, but in principle DEMESST can be applied to different bases that more readily support other states. This tomography method can even be used for direct fidelity estimation by choosing as a basis the intended target state. While Wigner tomography was presented in this work, the method can operate beyond the bosonic Wigner function. Additionally, our methods can work without requiring coupling gates between modes, which will for example be useful for calibrating entangled states over distributed quantum networks. We also implement another Wigner tomography method, the W^2 approach, as a direct fidelity estimator, and confirm that it agrees with DEMESST and W_{inv} . The DEMESST sampling method enables efficient measurement of large multimode states, which will be crucial as the size of quantum hardware systems increases and larger, more complicated states are generated and applied to quantum simulation, bosonic logical state encoding, and error correction.

Chapter 7

Conclusion

In this thesis, we have described a multimode 3D cQED quantum hardware system that lives in a class of promising platforms for researching quantum phenomena. The advantages of this system are long coherence times, high qubit-cavity cooperativities, and a design that is both hardware-efficient and modular which easily lends itself to scaling up and creating even larger systems. In particular, the entire device can be controlled with as few as two control lines: one input and one output. We have discussed how to fully calibrate the qubit and multimode cavity with time-domain measurements that allow us to measure all the relevant system parameters. We have applied this information to implement a number of control techniques for efficient manipulation of our system, enabling us to prepare arbitrary quantum states and generate interactions between different cavity modes. These include four-wave mixing sideband interactions, cavity beamsplitters, multimode photon blockade, and quantum optimal control. In conjunction with these controls, we have developed tomography and measurement approaches with improved efficiency to handle larger, more complicated quantum states.

Moving forward, this platform can serve as a starting point for a number of further improvements. For example, as much as the interaction between qubit and cavity is what allows us to achieve our quantum controls, there are times where we would like to turn it off, such as during the measurement process, or when we do not want the population of some number of cavity modes to affect how we control others. This interaction “switch” can be implemented by including an additional intermediary cavity and a tunable coupling element between our qubit and main multimode cavity, and would facilitate easier control and measurement of a large multimode system. Other innovations could be coupling two or more of these systems to fully demonstrate their modular nature, or increasing the number

of modes manipulated in a single version. Both of these additions would aim to further scale up the size of the system, and would be necessary as control and measurement techniques build on what we have developed in the discussion here and evolve further.

Ultimately, this work is an addition to the burgeoning field of quantum research and quantum information science. Our multimode cQED system can be treated as a quantum processor, with the qubit acting as the central processor and each of the cavity modes acting as memory sites. Ongoing developments that build on this one, while also pushing in other related areas, will strive to further explore quantum science and eventually achieve a useful, working quantum computing device. On the way, there will be a great variety of meaningful results that are compelling in their own right and move the field forward. Quantum science and the exploration of quantum phenomena has and will continue to be an exciting environment for technological development and discovery.

References

- [1] Schuster, D. *Circuit Quantum Electrodynamics*. Ph.D. thesis, Yale University (2007). URL <https://rsl.yale.edu/theses>.
- [2] Chakram, S. *et al.* Seamless high-q microwave cavities for multimode circuit quantum electrodynamics. *Physical Review Letters* **127**, 107701 (2021).
- [3] Koch, J. *et al.* Charge-insensitive qubit design derived from the cooper pair box. *Phys. Rev. A* **76**, 042319 (2007).
- [4] Mineev, Z. *et al.* Energy-participation quantization of josephson circuits. *npj Quantum Information* **7** (2021).
- [5] Mineev, Z., Leghtas, Z. & *et. al.* pyepr: Powerful, automated analysis and design of quantum microwave devices (2020). URL <https://pyepr-docs.readthedocs.io/en/latest/>.
- [6] Oriani, A. *Multimodal and Ultra High-Q Superconducting Niobium Cavities for Circuit Quantum Electrodynamics*. Ph.D. thesis, University of Chicago (2022). URL <https://schusterlab.stanford.edu/publications.html>.
- [7] Reagor, M. *et al.* Quantum memory with millisecond coherence in circuit qed. *Phys. Rev. B* **94**, 014506 (2016).
- [8] Brecht, T. *Micromachined Quantum Circuits*. Ph.D. thesis, Yale University (2017). URL <https://rsl.yale.edu/theses>.
- [9] Reagor, M. *et al.* Reaching 10 ms single photon lifetimes for superconducting aluminum cavities. *App. Phys. Lett.* **102**, 192604 (2013). URL <http://scitation.aip.org/content/aip/journal/apl/102/19/10.1063/1.4807015>.
- [10] Zhou, B. & Ramirez, W. Kinetics and modeling of wet etching of aluminum oxide by warm phosphoric acid. *Journal of The Electrochemical Society* **143**, 619 (1996).
- [11] Tinkham, M. *Introduction to superconductivity* (Courier Corporation, 2004).
- [12] Yamamoto, T. *et al.* Flux-driven josephson parametric amplifier. *Applied Physics Letters* **93**, 042510 (2008).
- [13] Eichler, C. & Wallraff, A. Controlling the dynamic range of a josephson parametric amplifier. *EPJ Quantum Technology* **1**, 2 (2014). URL <https://link.springer.com/article/10.1140/epjqt2>.
- [14] Roy, T. *et al.* Broadband parametric amplification with impedance engineering: Beyond the gain-bandwidth product. *Applied Physics Letters* **107**, 262601 (2015).
- [15] Esposito, M. *et al.* Development and characterization of a flux-pumped lumped element josephson parametric amplifier. *EPJ Web of Conferences* **198**, 00008 (2019).

- [16] Wustmann, W. & Shumeiko, V. Parametric resonance in tunable superconducting cavities. *Physical Review B* **87**, 184501 (2013).
- [17] Zhou, X. *et al.* High-gain weakly nonlinear flux-modulated josephson parametric amplifier using a squid array. *Physical Review B* **89**, 214517 (2014).
- [18] Agrawal, A. *Superconducting Qubit Advantage for Dark Matter (SQuAD)*. Ph.D. thesis, University of Chicago (2022). URL <https://schusterlab.stanford.edu/publications.html>.
- [19] Motzoi, F., Gambetta, J., Reberntrost, P. & Wilhelm, F. Simple pulses for elimination of leakage in weakly nonlinear qubits. *Physical Review Letters* **103**, 110501 (2009).
- [20] Chow, J. *et al.* Optimized driving of superconducting artificial atoms for improved single-qubit gates. *Physical Review A* **82**, 040305 (2010).
- [21] Reed, M. *Controlling Error-Correctable Bosonic Qubits*. Ph.D. thesis, Yale University (2013).
- [22] Schuster, D. *et al.* Resolving photon number states in a superconducting circuit. *Nature* **445**, 515–518 (2007).
- [23] Irish, E. & Schwab, K. Quantum measurement of a coupled nanomechanical resonator – cooper-pair box system. *Physical Review B* **68**, 155311 (2003).
- [24] Gambetta, J. *et al.* Qubit-photon interactions in a cavity: Measurement induced dephasing and number splitting. *Physical Review A* **74**, 042318 (2006).
- [25] Gertler, J. *et al.* Protecting a bosonic qubit with autonomous quantum error correction. *Nature* **590**, 243–248 (2021).
- [26] Combes, J. & Brod, D. Two-photon self-kerr nonlinearities for quantum computing and quantum optics. *Physical Review A* **98**, 062313 (2018).
- [27] Chakram, S. *et al.* Multimode photon blockade. *Nature Physics* **18**, 879–884 (2022).
- [28] Rosenblum, S. *et al.* A cnot gate between multiphoton qubits encoded in two cavities. *Nature communications* **9**, 1–6 (2018).
- [29] Pechal, M. *et al.* Microwave-controlled generation of shaped single photons in circuit quantum electrodynamics. *Physical Review X* **4**, 041010 (2014).
- [30] Naik, R. *et al.* Random access quantum information processors using multimode circuit quantum electrodynamics. *Nature communications* **8**, 1–7 (2017).
- [31] Leghtas, Z. *et al.* Hardware-efficient autonomous quantum memory protection. *Phys. Rev. Lett.* **111**, 120501 (2013). URL <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.111.120501>.

- [32] Michael, M. H. *et al.* New class of quantum error-correcting codes for a bosonic mode. *Phys. Rev. X* **6**, 031006 (2016). URL <https://journals.aps.org/prx/abstract/10.1103/PhysRevX.6.031006>.
- [33] Hu, L. *et al.* Quantum error correction and universal gate set operation on a binomial bosonic logical qubit. *Nature Physics* **15**, 503–508 (2019).
- [34] Sun, L. *et al.* Tracking photon jumps with repeated quantum non-demolition parity measurements. *Nature* **511**, 444–448 (2014).
- [35] Brune, M., Haroche, S., Raimond, J., Davidovich, L. & Zagury, N. Manipulation of photons in a cavity by dispersive atom-field coupling: Quantum-nondemolition measurements and generation of “schrödinger cat” states. *Physical Review A* **45**, 5193 (1992).
- [36] Bultink, C. *et al.* Protecting quantum entanglement from leakage and qubit errors via repetitive parity measurements. *Science Advances* **6** (2020).
- [37] Andersen, C. *et al.* American association for the advancement of science. *npj Quantum Information* **5** (2019).
- [38] Royer, A. Wigner function as the expectation value of a parity operator. *Physical Review A* **15**, 449 (1977).
- [39] Weedbrook, C. *et al.* Gaussian quantum information. *Reviews of Modern Physics* **84**, 621 (2012).
- [40] Kok, P. *et al.* Linear optical quantum computing with photonic qubits. *Reviews of Modern Physics* **79**, 135 (2012).
- [41] Lau, H. & Plenio, M. Universal quantum computing with arbitrary continuous-variable encoding. *Physical Review Letters* **117**, 100501 (2016).
- [42] Sparrow, C. *et al.* Simulating the vibrational quantum dynamics of molecules using photonics. *Nature* **557**, 660–667 (2018).
- [43] Shen, Y. *et al.* Quantum optical emulation of molecular vibronic spectroscopy using a trapped-ion device. *Chemical Science* **9**, 836–840 (2018).
- [44] Islam, R. *et al.* Measuring entanglement entropy in a quantum many-body system. *Nature* **528**, 77–83 (2015).
- [45] Gao, Y. *et al.* Programmable interference between two microwave quantum memories. *Physical Review X* **8**, 021073 (2018).
- [46] Zhang, Y. *et al.* Programmable interference between two microwave quantum memories. *Physical Review A* **99**, 012314 (2019).

- [47] Bretheau, L., Campagne-Ibarcq, P., Flurin, E., Mallet, F. & Huard, B. Quantum dynamics of an electromagnetic mode that cannot contain n photons. *Science* **348**, 776–779 (2015).
- [48] Vrajitoarea, A., Huang, Z., Groszkowski, P., Koch, J. & Houck, A. A. Quantum control of an oscillator using a stimulated josephson nonlinearity. *Nature Physics* **16**, 211–217 (2020).
- [49] Facchi, P., Gorini, V., Marmo, G., Pascazio, S. & Sudarshan, E. Quantum zeno dynamics. *Physics Letters A* **275**, 12–19 (2000).
- [50] Facchi, P. & Pascazio, S. Quantum zeno dynamics: mathematical and physical aspects. *Journal of Physics A: Mathematical and Theoretical* **41**, 493001 (2008).
- [51] Raimond, J.-M. *et al.* Phase space tweezers for tailoring cavity fields by quantum zeno dynamics. *Physical review letters* **105**, 213601 (2010).
- [52] Raimond, J.-M. *et al.* Quantum zeno dynamics of a field in a cavity. *Physical Review A* **86**, 032120 (2012).
- [53] Burgarth, D. K. *et al.* Exponential rise of dynamical complexity in quantum computing through projections. *Nature Communications* **5** (2014).
- [54] Signoles, A. *et al.* Confined quantum zeno dynamics of a watched atomic arrow. *Nature Physics* **10**, 715–719 (2014).
- [55] Schäfer, F. *et al.* Experimental realization of quantum zeno dynamics. *Nature communications* **5**, 1–6 (2014).
- [56] Barontini, G., Hohmann, L., Haas, F., Estève, J. & Reichel, J. Deterministic generation of multiparticle entanglement by quantum zeno dynamics. *Science* **349**, 1317–1321 (2015).
- [57] Patil, Y., Chakram, S. & Vengalattore, M. Measurement-induced localization of an ultracold lattice gas. *Physical review letters* **115**, 140402 (2015).
- [58] Heeres, R. W. *et al.* Cavity state manipulation using photon-number selective phase gates. *Physical review letters* **115**, 137002 (2015).
- [59] Krastanov, S. *et al.* Universal control of an oscillator with dispersive coupling to a qubit. *Physical Review A* **92**, 040303(R) (2015).
- [60] Fasel, T., Krastanov, S., Marquardt, F. & Jiang, L. Efficient cavity control with snap gates. *Bulletin of the American Physical Society* **65**, R16.00003 (2020).
- [61] Kudra, M. *et al.* Robust preparation of wigner-negative states with optimized snap-displacement sequences. *PRX Quantum* **3**, 030301 (2021).

- [62] Cahill, K. & Glauber, R. Density operators and quasiprobability distributions. *Physical Review Journals Archive* **177**, 1882 (1969).
- [63] Schirmer, S. G., Fu, H. & Solomon, A. I. Complete controllability of quantum systems. *Physical Review A* **63**, 063410 (2001).
- [64] Khaneja, N., Reiss, T., Kehlet, C., Schulte-Herbrüggen, T. & Glaser, S. J. Optimal control of coupled spin dynamics: design of nmr pulse sequences by gradient ascent algorithms. *Journal of Magnetic Resonance* **172**, 296–305 (2005).
- [65] Leung, N., Abdelhafez, M., Koch, J. & Schuster, D. Speedup for quantum optimal control from automatic differentiation based on graphics processing units. *Physical Review A* **95**, 042318 (2017).
- [66] Rosenblum, S. *et al.* Fault-tolerant detection of a quantum error. *Science* **361**, 266–270 (2018).
- [67] Zeytinoglu, S. *et al.* Microwave-induced amplitude- and phase-tunable qubit-resonator coupling in circuit quantum electrodynamics. *Physical Review A* **91**, 043846 (2015).
- [68] Ma, W. *et al.* Quantum control of bosonic modes with superconducting circuits. *Science Bulletin* **66**, 1789–1805 (2021).
- [69] Propson, T., Jackson, B., Koch, J., Manchester, Z. & Schuster, D. Robust quantum optimal control with trajectory optimization. *Physical Review Applied* **17**, 014036 (2022).
- [70] Morzhin, O. & Pechen, A. Krotov method for optimal control of closed quantum systems. *Russian Mathematical Surveys* **74**, 851 (2019).
- [71] Goerz, M. *et al.* Krotov: A python implementation of krotov’s method for quantum optimal control. *SciPost Physics* **7**, 080 (2019).
- [72] Ball, H. *et al.* Software tools for quantum control: improving quantum computer performance through noise and error suppression. *Quantum Science and Technology* **6**, 044011 (2021).
- [73] Eickbusch, A. *et al.* Fast universal control of an oscillator with weak dispersive coupling to a qubit. *Nature Physics* **18**, 1464–1469 (2022).
- [74] Trowbridge, A., Bhardwaj, A., He, K., Schuster, D. & Manchester, Z. Direct collocation for quantum optimal control. *arXiv2305.03261* (2023).
- [75] Byrd, R., Lu, P., Nocedal, J. & Zhu, C. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing* **16**, 0916069 (1995).
- [76] Kingma, D. & Ba, J. Adam: A method for stochastic optimization. *arXiv1412.6980* (2014).

- [77] Dixit, A. *et al.* Searching for dark matter with a superconducting qubit. *Physical Review Letters* **126**, 141302 (2021).
- [78] Agrawal, A. *et al.* Quantum enhancement of weak signals using c-qed tools. *in preparation* (2023).
- [79] Reinhold, P. *Controlling Error-Correctable Bosonic Qubits*. Ph.D. thesis, Yale University (2019). URL [https://rsl.yale.edu/sites/default/files/files/RSL_Theses/Reinhold-Thesis%20\(1\).pdf](https://rsl.yale.edu/sites/default/files/files/RSL_Theses/Reinhold-Thesis%20(1).pdf).
- [80] Smolin, J. A., Gambetta, J. M. & Smith, G. Efficient method for computing the maximum-likelihood quantum state from measurements with additive gaussian noise. *Physical review letters* **108**, 070502 (2012).
- [81] Shen, C. *et al.* Optimized tomography of continuous variable systems using excitation counting. *Physical Review A* **94**, 052327 (2016).
- [82] Wang, C. *et al.* A schrödinger cat living in two boxes. *Science* **352**, 1087–1091 (2016).
- [83] Cramer, M. *et al.* Efficient quantum state tomography. *Nature Communications* **1** (2010).
- [84] Paris, M. & Rehacek, J. *Quantum State Estimation* (Springer, 2004).
- [85] Botelho, L. & Vianna, R. Efficient quantum tomography of two-mode wigner functions. *The European Physical Journal D* **74** (2020).
- [86] Siudzinksa, K. & Chruscinski, D. Entanglement witnesses from mutually unbiased measurements. *Scientific Reports* **11**, 22988 (2021).
- [87] Chruscinski, D. & Sarbicki, G. Entanglement witnesses: construction, analysis and classification. *Journal of Physics A: Mathematical and Theoretical* **47**, 483001 (2014).
- [88] da Silva, M., Landon-Cardinal, O. & Poulin, D. Practical characterization of quantum devices without tomography. *Physical Review Letters* **107**, 210404 (2011).
- [89] Mohseni, M., Rezakhani, A. & Aspuru-Guzik, A. Direct estimation of single- and two-qubit hamiltonians and relaxation rates. *Physical Review A* **77**, 042320 (2008).
- [90] quantum state tomography, E. Scalable and programmable phononic network with trapped ions. *Nature Communications* **1** (2010).
- [91] Chen, W. *et al.* Scalable and programmable phononic network with trapped ions. *Nature Physics* **19**, 877–883 (2023).
- [92] Shen, C. *et al.* Optimized tomography of continuous variable systems using excitation counting. *Physical Review A* **94**, 052327 (2016).

- [93] He, K. *et al.* Efficient multimode wigner tomography. *in preparation* (2023).
- [94] Flammia, S. & Liu, Y. Direct fidelity estimation from few pauli measurements. *Physical Review Letters* **106**, 230501 (2011).