

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Metrology of Quantum Control and Measurement in Superconducting Qubits

Permalink

<https://escholarship.org/uc/item/0g29b4p0>

Author

Chen, Zijun

Publication Date

2018

Peer reviewed|Thesis/dissertation

UNIVERSITY of CALIFORNIA
Santa Barbara

**Metrology of Quantum Control and Measurement in Superconducting
Qubits**

A dissertation submitted in partial satisfaction of the
requirements for the degree of

Doctor of Philosophy

in

Physics

by

Zijun Chen

Committee in charge:

Professor John Martinis, Chair

Professor Mark Sherwin

Professor Wim van Dam

January 2018

The dissertation of Zijun Chen is approved:

Professor Mark Sherwin

Professor Wim van Dam

Professor John Martinis, Chair

January 2018

Copyright © 2018
by Zijun Chen

To my parents, Ling Chen and Xiaofen Tang,
and to my fiancée, Claire McLellan

Acknowledgements

First and foremost, I would like to thank the teachers who trained me to be a physicist. Phil Sullivan was my first physics teacher, and his ability to get high schoolers excited about physics is uncanny. His emphasis on deduction and learning from experiments was crucial to my physics education. I couldn't have asked for a better undergraduate advisor than Dan Dahlberg, who gave me the freedom to independently work on interesting problems. Through working in his lab, I learned a plethora of experimental skills that have remained valuable in my graduate work, yet Dan always made sure I was well rounded and doing well in classes and outside of school. Thank you for four years of excellent mentoring.

And of course, John Martinis is the man who made this thesis happen. John has a unique combination of big-picture vision, attention to detail, and a wide breadth of knowledge that just might get a quantum computer built. As my graduate advisor, John has been nothing but supportive in my efforts, and I especially appreciate his guidance on my presentations and papers. Furthermore, his emphasis on doing things the right way and understanding everything about our experimental systems has really built a group culture that makes it a pleasure to do research. Thank you for everything you've done these past five years, and I'm looking forward to continuing our efforts for years to come.

Daniel Sank was one of the first people I met in the Martinis lab, and his excitement about the lab's work was a big part of why I joined the group. He was also my first mentor within the group, and ever since then I've learned a tremendous amount from him about quantum mechanics, computer programming, and electronics. I know I'm occasionally a stubborn student - thank you for sticking with it and making the lab smarter and better.

No matter the problem, I know Tony Megrant has a solution. Through booster seats, timelines of sadness, and midnight cooldowns, we've been through it all, and I'm still in awe sometimes of how your mind never stops working. Thank you for always having my back, and congratulations on your newborn girl!

Nobody in the lab is more efficient than Julian Kelly, and I strive to match his capacity for completing even the most complicated projects. We've had a knack for collaboration - the airbridge and leakage benchmarking projects were originally his ideas that he passed on to me, which I managed to turn into papers with his help. Thank you for helping me be successful, and here's to even more triumphs in the future.

Brooks Campbell and Andrew Dunsworth were the other two graduate students in the Martinis Lab from my year, so we've been through it all together. Brooks never fails to project anything but calm - even when his wedding was literally on fire. He's also the master of wirebonding and other skills which have bailed me out many times. Andrew has been my best friend in the group, and when things are going rough, I can always count on him. There's probably no one in the group I enjoy working with more, and his skills in the cleanroom are crucial for the group. Thank you both for the good memories, and good luck with the last thesis push!

Chris Quintana was the only graduate student to join after Andrew and I, but we all quickly learned how capable he was. Nobody in the lab works harder, and his dedication

to research is evident in the high quality of his papers and the sheer volume of his presentations. Thank you for many great conversations over the years about physics and beyond.

Amit Vainsencher never fails to give good advice, whether it be about electronics, complex fabrication processes, or just life in general. Thank you for always being there to help when something goes wrong - I'm sure I'll be needing even more of it in the future.

Peter O'Malley is the one who got away. Every week, something will come up and I think, man, we could really use Peter's help right now. Thank you for your excellent research and programming work, and I hope life in Boston is treating you well.

I swear Evan Jeffrey knows everything. I have yet to witness him *not* have an important insight when asked a question over a wide range of topics. Thank you for bringing your knowledge and wisdom to our group.

Ted White's level-headed and cheerful demeanor makes working on any project better. Thank you for being a constant in the crazy time that is graduate school, and also for making awesome amplifiers that let us do our experiments.

Josh Mutus also makes awesome amplifiers, and helped train me early on in the cleanroom. Thank you for being the one person in the lab that I can talk to about sports.

Jim Wenner is the heart (and the logistical master) of the Martinis group. Thank you for buying snacks, making sure we all sign up for March Meeting, and just generally being a nice person.

Ben Chiaro was my office mate for a long time, and has persevered through some tough projects. To do so with two kids is even more admirable. Thank you for all of your hard work.

No one has a knack for physical intuition like Charles Neill, a quality that has helped him pull off some of the most complicated experiments in the lab from start to finish. Thank you for always being a great person to bounce ideas off of.

I haven't had many chances to work with Yu Chen, a fellow Golden Gopher, but I wish I did. His focus and determination are unmatched, and his progress on the fluxmon with Chris has been astounding. Thank you for always looking out for the graduate students, and for commiserating with me about Chinese food.

Pedram Roushan is a man who knows how to get things done and get papers published. Thank you for letting me work on the gmon project, where I gained some valuable experience with multiqubit systems.

Austin Fowler is the man behind the surface code, and I still remember when he was just a voice dropping wisdom from far away in Australia. Thank you for bringing us the architecture to build a quantum computer.

Rami Barends has been instrumental in devising and executing some of our group's most important experiments. Thank you for your contributions, and also for always promptly reading my papers.

I took over Erik Lucero's desk (and later on, one of his projects) when I first joined

the lab, but didn't meet the legend himself for a few years. Thank you for your continual optimism and excitement about what we do.

I also met Matthew Neeley relatively late in my graduate career, but have already learned so much from him. Thank you for making my code better, and for always having a solution to the thorniest programming problems.

To the engineers, physicists, and staff who have joined the Google group in recent years - Paul, Marissa, Brian, Rob, Trent, Kunal, Craig, and Alessandra - you all have already made significant contributions and made the group better. Thank you for bringing us fresh perspectives, and I look forward to working with all of you more closely in the future.

Physics is very much about collaboration between experiment and theory, and I've been lucky to work with some brilliant theorists. I want to especially thank Alexander Korotkov and Mostafa Khezri from UC Riverside, who worked with Dan and I on the tricky and surprising readout transitions problem, but always managed to make their findings intuitive and digestible.

Many support and administrative staff members were crucial to my graduate career: the reliable and knowledgeable cleanroom engineers in the Nanostructures Cleanroom Facility; likewise for the CNSI cleanroom staff, especially Peter Duda; the administrative staff at CNSI, especially Holly Woo, who is always good for a piece of life advice; and the administrative staff at the UCSB Physics Department, especially Jennifer Farrar, who tirelessly works to make our graduate careers easier and better. Thank you all for making our research possible.

Thank you also to Professors Mark Sherwin and Wim van Dam for taking the time to serve on my thesis committee, especially after having to reschedule because of the wildfires.

To the many music teachers and conductors I've had in my life, thank you for giving me the gift of playing piano and viola. Sorry, Jim and Adam, that I wasn't able to play your concerts while I wrapped this thing up!

I wouldn't be here today without my parents' love and support. My dad was the first person to get a Ph.D. in my family, and I strive to be as smart, patient, and reliable as him every day. My mother was in fact my first math teacher, and has never stopped trying to make me a better person. Thank you both for raising me to be where I am now; my success today is very much yours' as well.

Finally, thank you Claire, for being the nicest person I know, and for being there through thick and thin. There were lots of great things about coming to UCSB, but getting to be with you was by far the best of them all. Here's to finally finishing grad school, and getting to spend the rest of our lives together!

Curriculum Vitæ

Zijun Chen

Education

- 2017 Doctor of Philosophy, Physics, University of California, Santa Barbara
- 2012 Bachelor of Science, Physics and Mathematics, University of Minnesota, Minneapolis, Minnesota
- 2008 Roseville Area High School, Roseville, Minnesota

First author publications

“Measurement-induced State Transitions in a Superconducting Qubit: Beyond the Rotating Wave Approximation”, Daniel Sank, Zijun Chen, Mostafa Khezri, J. Kelly, R. Barends, Y. Chen, A. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Mutus, M. Neeley, P. Roushan, A. Vainsencher, T. White, B. Campbell, B. Chiaro, A. Dunsworth, C. Neill, P. J. J. O’Malley, C. Quintana, J. Wenner, Alexander N. Korotkov, John M. Martinis *Physical Review Letters* Vol. 117, 190503 (2016)

“Measuring and Suppressing Quantum State Leakage in a Superconducting Qubit”, Zijun Chen, Julian Kelly, Christopher Quintana, R. Barends, B. Campbell, Yu Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Lucero, E. Jeffrey, A. Megrant, J. Mutus, M. Neeley, C. Neill, P. J. J. O’Malley, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, A. N. Korotkov, John M. Martinis *Physical Review Letters* Vol. 116, 020501 (2016)

“Fabrication and Characterization of Aluminum Airbridges for Superconducting Microwave Circuits”, Zijun Chen, Anthony Megrant, Julian Kelly, R. Barends, J. Bochmann, Yu Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, J.Y. Mutus, P.J.J. OMalley, C. Neill, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T.C. White, A.N. Cleland, John M. Martinis *Applied Physics Letters* Vol. 104, 052602 (2014)

“A simple technique to measure the magnetic susceptibility of liquids”, Zijun Chen, Joshua Ellis, E. D. Dahlberg *Review of Scientific Instruments* Vol. 83, 095112 (2012)

“Deformation of Water by a Magnetic Field”, Zijun Chen, E. D. Dahlberg *The Physics Teacher* Vol. 49, 14446 (2011)

Abstract

Metrology of Quantum Control and Measurement in Superconducting Qubits

by

Zijun Chen

Quantum computers have the potential to solve problems which are classically intractable. Superconducting qubits present a promising path to building such a computer. Recent experiments with these qubits have demonstrated the principles of quantum error correction, quantum simulation, quantum annealing, and more. Current research with superconducting qubits is focused on two primary goals: creating a fully fault tolerant logical qubit out of many physical qubits using surface code error correction, and demonstrating an exponential speedup over any classical computer for a well-defined computational problem. To achieve either of these goals requires high precision control of three components: single qubit gates, two qubit gates, and qubit measurement. In this thesis, we use randomized benchmarking to characterize single qubit gates with 99.95% fidelity and two qubit gates with 99.5% fidelity in superconducting transmon qubits. In addition, we use standard decoherence measurements as well as newly developed extensions of randomized benchmarking to determine the limiting sources of error. Finally, we explore the surprisingly complicated dynamics of measuring the transmon state through a cou-

pled resonator, and show that fully understanding this process requires breaking a few "standard" assumptions.

Contents

1	Introduction to Quantum Computing	1
1.1	Processing Information using Quantum Mechanics	2
1.1.1	What is the Power of Quantum Computing?	7
1.2	Applications of Quantum Computing	7
1.2.1	Shor’s Algorithm	8
1.2.2	Grover’s Algorithm	9
1.2.3	Quantum Simulation	10
1.3	Two Approaches to Quantum Computing Research	11
1.3.1	Quantum Error Correction	11
1.3.2	Quantum Supremacy	16
1.3.3	Outlook	19
1.4	Qubit Implementations	20
1.4.1	Trapped Ions	20
1.4.2	Semiconductor Qubits	21
1.4.3	Superconducting Qubits	22
1.5	Summary and Thesis Outline	23
2	Superconducting Transmon Qubits	25
2.1	LC Circuit	26
2.2	Transmon Hamiltonian	29
2.3	Driving	33
2.4	Coupling	40
2.4.1	Oscillator-Oscillator Coupling	40
2.4.2	Oscillator-Qubit Coupling	43
2.4.3	Qubit-Qubit Coupling	49
2.5	Frequency Control	51
2.6	Typical Transmon Parameters	52
2.6.1	Frequency	53
2.6.2	Control Coupling	53
2.6.3	Readout Coupling	54
2.6.4	Qubit-Qubit Coupling	55
2.6.5	Summary	55

3	Superconducting Qubit Fabrication	57
3.1	Basic Superconducting Qubit Fabrication	57
3.1.1	Thin Film Deposition	59
3.1.2	Patterning	60
3.1.3	Etching	62
3.1.4	Junction Fabrication	63
3.2	Fabricating Superconducting Airbridges	65
3.3	Contributions	71
4	Experimental Infrastructure for Superconducting Qubits	72
4.1	Cryogenic	72
4.1.1	Cyrostat	72
4.1.2	Sample Mounting	75
4.1.3	Wiring	76
4.1.4	Readout Chain	78
4.2	Room Temperature Electronics	80
4.2.1	High Frequency Control	81
4.2.2	Readout	83
4.2.3	Low Frequency Flux Control	83
4.3	Software	84
4.3.1	LabRAD	84
4.3.2	Hardware Servers	85
4.3.3	Sequencer	86
4.3.4	Gates and Experiments	87
4.3.5	Automated Calibrations	88
5	Single Qubit Calibration	89
5.1	Introduction	89
5.2	Resonator Spectroscopy	91
5.3	Qubit Spectroscopy	94
5.4	Rabi	96
5.5	Readout Frequency	98
5.6	Readout Amplitude	100
5.7	Ramsey	102
5.8	Pulse Corrections	104
5.9	Rabi, Revisited	107
5.10	XY-Z Timing	109
5.11	Bias Amplitude to Frequency Function	109
6	Single Qubit Benchmarking	112
6.1	A Brief Review of Randomized Benchmarking	112
6.2	Clifford Based Randomized Benchmarking of Xmon Qubits	116
6.3	Purity Benchmarking	120

6.4	Benchmarking Errors vs Pulse Length	123
6.5	Coherence Measurements	126
6.6	Conclusion	131
6.7	Contributions	132
7	Measuring and Calibrating for Two State Errors	133
7.1	Gate Errors Due to Non-Computational Energy Levels	133
7.2	Test Device	136
7.3	Measuring Total Gate Error and Leakage Error	137
7.4	Pulse Detuning to Minimize Phase Error	139
7.5	Optimized Microwave Pulses	142
7.6	Leakage vs Pulse Length	144
7.7	Conclusions	147
7.8	Contributions	148
8	Two Qubit Calibration and Benchmarking	149
8.1	The Controlled Phase Interaction in Coupled Transmons	150
8.2	Calibrating Non-Idealities in Qubit Flux Control	154
8.2.1	Correcting for Flux Control Distortions	156
8.2.2	Multiqubit Timing	158
8.2.3	Flux Crosstalk	161
8.3	Calibrating the Control Z	162
8.3.1	Conditional Phase	164
8.3.2	Leakage	166
8.3.3	Iteration and Phase Compensation	166
8.4	Benchmarking the CZ Gate	167
8.5	Conclusions	172
8.6	Contributions	172
9	Readout Induced Qubit Transitions	174
9.1	Measurement Induced State Transitions	175
9.1.1	Motivation	175
9.1.2	Experimental Observations	175
9.1.3	Theory of non-RWA Transitions	180
9.1.4	Comparing Theory to Experiment	181
9.1.5	Summary	185
9.2	Induced Transitions in a Driven Qubit	185
9.2.1	The AC Stark Effect with $\omega_r < \omega_q$	187
9.2.2	Revisiting the JC Ladder	188
9.2.3	Comparison to Theory	191
9.2.4	Multi Qubit Transitions	197
9.3	Conclusions	198
9.4	Contributions	199

10 Conclusion	200
A Measurements of Airbridge Circuits	202
A.1 Resistance and Critical Current	202
A.2 Resonators With Airbridge Groundplane Shunts	204
A.2.1 Shift in Resonant Frequency	205
A.2.2 Loss	207
A.3 Participation Ratio of the Airbridge Interface	209
A.3.1 Center Trace Airbridges	213
Bibliography	215

Chapter 1

Introduction to Quantum

Computing

The theory of quantum mechanics is a foundation of modern physics, and successfully describes a myriad of physical phenomena, from high energy particle physics to the behavior of solid state materials near absolute zero. Quantum mechanics has also had a profound impact on computing. For example, quantum mechanics was crucial to the development of semiconductor physics which led to the invention of the transistor [1], the bedrock of modern computing. Once esoteric quantum effects such as quantum tunnelling are also becoming increasingly important as logic and storage grow more dense [2, 3, 4]. However, the primary paradigm for computing has remained the same since the days of vacuum tubes, and is based on classical, binary bits of information. Quantum computing is the idea that by harnessing the power of quantum mechanics, we

can fundamentally change the way we store and process and information, and potentially out perform classical computation.

1.1 Processing Information using Quantum Mechanics

The fundamental unit of quantum information is the quantum bit, or *qubit*. A qubit is a quantum system with two states, which we denote as $|0\rangle$ and $|1\rangle$. Information is encoded in the amplitudes for each state

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

where α and β are complex. The vector $|\psi\rangle$ is also known as a *wavefunction*. Information is extracted from $|\psi\rangle$ by *measuring* the qubit, and the probability of measuring $|0\rangle$ is $|\alpha|^2$, and likewise for $|1\rangle$ and β . After measurement, the wavefunction is said to be collapsed into the state that was measured, with subsequent measurements always yielding the same value. Since the combined measurement probability must be 1, α and β must satisfy $|\alpha|^2 + |\beta|^2 = 1$.

Equation 1.1 states that a qubit may exist in a *superposition* where it is neither definitely $|0\rangle$ or $|1\rangle$. However, a qubit is more than just a probabilistic bit of information. For example, the states $(|0\rangle + |1\rangle)/\sqrt{2}$ and $(|0\rangle - |1\rangle)/\sqrt{2}$ are fundamentally different states despite yielding the same measurement probabilities, and both are different from the state $(|0\rangle + i|1\rangle)/\sqrt{2}$.




Figure 1.1: The Bloch sphere, representing a single qubit.

To better understand the information encoded in a qubit, we consider a qubit's degrees of freedom. Since α and β are complex, they each contain two degrees of freedom. The normalization of total probability imposes one constraint. Additionally, we ignore a second degree of freedom which scales $|\psi\rangle$ by a global complex phase factor that does not have practical consequences for a single qubit. We then rewrite Eqn. 1.1 in the following manner:

$$|\psi\rangle = \cos \frac{\theta}{2} + e^{i\phi} \sin \frac{\theta}{2}, \quad (1.2)$$

and pair it with the geometrical construction shown in Fig. 1.1, which is known as the *Bloch sphere* [5]. The qubit is represented as a vector from the origin to a point on the surface of the sphere. Information is encoded in the azimuthal and equatorial angles of the vector. The poles of the sphere correspond to $|0\rangle$ and $|1\rangle$, while the equator of the sphere corresponds to 50/50 superposition states. However, note that the sphere is

symmetric with respect to rotation about any axis that passes through the origin. For example, we can rotate the sphere such that the poles now correspond to $(|0\rangle + |1\rangle)/\sqrt{2}$ and $(|0\rangle - |1\rangle)/\sqrt{2}$. In an experiment, this action is equivalent to measuring the qubit in a different *basis*, and in this way, we can resolve the difference between different superposition states.

To bring a qubit from one state to another, we evolve it under a Hamiltonian H according to Schrodinger's equation.

$$H|\psi\rangle = i\hbar \frac{d|\psi\rangle}{dt} \quad (1.3)$$

In quantum computing, it is often useful to assume that the Hamiltonian is zero during "static" operation. Transformations are achieved by turning on a control Hamiltonian for a discrete amount of time. The total transformation during this time is given by solving the Schroedinger equation ¹

$$U = \exp [iHt/\hbar] \quad (1.4)$$

where U is known as a *propagator*, and is also referred to as a *quantum gate* in the context of quantum computing. Importantly, U is *unitary*, which implies that quantum gates are reversible. Unlike classical gates such as AND, the input to a known quantum gate can always be constructed given the output.

For a single qubit, quantum gates can also be represented on the Bloch sphere. All

¹If H is time dependent, then the expression for U is more complex and involves a time-ordered integral.

single qubit unitaries can be written as

$$U = \exp \left[-\frac{i\Theta \vec{S} \cdot \vec{n}}{2} \right] \quad (1.5)$$

where Θ is an angle, \vec{n} is a length 3 direction vector, and $\vec{S} = (S_x, S_y, S_z)$ is the vector of Pauli matrices

$$S_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad (1.6)$$

$$S_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad (1.7)$$

$$S_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (1.8)$$

In the Bloch sphere, the action of U is to rotate the vector $|\psi\rangle$ about the axis \vec{n} by the angle Θ .

Next, we consider systems of multiple qubits. A simple multiqubit state might consist of two independent qubits with wavefunctions $|\psi_1\rangle = \alpha_1|0\rangle + \beta_1|1\rangle$ and $|\psi_2\rangle = \alpha_2|0\rangle + \beta_2|1\rangle$. The total wavefunction is given by the outer product

$$|\psi_{\text{total}}\rangle = |\psi_1\rangle \otimes |\psi_2\rangle \quad (1.9)$$

$$= \alpha_1\alpha_2|00\rangle + \beta_1\alpha_2|01\rangle + \alpha_1\beta_2|10\rangle + \beta_1\beta_2|11\rangle. \quad (1.10)$$

Such a state is called a *product state*. However, quantum mechanics allows other other

multiqubit states such as

$$|\psi_{\text{total}}\rangle = \frac{|01\rangle + |10\rangle}{\sqrt{2}}. \quad (1.11)$$

This state fundamentally cannot be factored into separate wavefunctions for the two qubits. In other words, information is contained not only in the state of individual qubits, but also in the correlations between the qubits. As a consequence, a system of N qubits requires not $2N$ amplitudes to describe measurement probabilities, but 2^N to fully capture all of the joint probabilities. This phenomenon is known as *quantum entanglement*, and is one of the primary resources for quantum computing since a collection of qubits can store and manipulate an amount of information that is exponential in the number of qubits. Entanglement can be generated by applying multiqubit quantum gates. For example, the CNOT, which we express in the ordered basis ($|00\rangle, |01\rangle, |10\rangle, |11\rangle$) as

$$U_{\text{CNOT}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad (1.12)$$

flips the state of one of the qubits conditional on the state of the other qubit. If applied to the product state $(|01\rangle + |11\rangle)/\sqrt{2}$, we obtain the entangled state from Eqn. 1.11.

Having established the basic concepts of quantum computing, we can now describe how a simple program for a quantum computer might run. We take a collection of qubits and initialize them into a known state. We then apply a series of single and two (or possibly more) qubit gates to the system, which is known as a *quantum circuit*. Finally, we measure the qubits to extract some information from the output wavefunction.

1.1.1 What is the Power of Quantum Computing?

Naively, superposition and entanglement appear to be powerful tools which should readily give exponential speedups over classical computation. A oft-repeated trope is that a quantum computer could prepare a superposition of N inputs using $\log_2 N$ qubits, act on all of them simultaneously, and thus reduce an exponential number of function calls down to one. However, this program will not provide any speedup because each measurement of the final wavefunction only yields one bitstring. There, it will take an exponential number of measurements to fully characterize the output wavefunction. A more subtle hint that superposition and entanglement are not sufficient is the fact that quantum circuits which contain only Clifford gates (which we will discuss in Chapter 7) can generate quantum states which are highly entangled, and yet these circuits can be efficiently simulated on a classical computer [6, 7].

Quantum algorithms which provide a speed up over classical algorithms must generally exploit a particular problem's structure in some way, and we will see in the next section how this is done in some well known cases. The computational complexity of a generic quantum algorithm is still a highly debated topic, and drives much of the theoretical and experimental research today.

1.2 Applications of Quantum Computing

In this section, we will discuss three applications of quantum computing and their general strategies for achieving speedup over classical algorithms.

1.2.1 Shor's Algorithm

Shor's algorithm factors a number into its prime factors, and was first discovered by Peter Shor in 1994 [8]. Prime factorization is an important problem because many modern encryption protocols rely on the assumption that prime factoring large numbers is an intractable problem. However, unlike the best known classical algorithms which scale nearly exponentially with the number of digits, Shor's algorithm scales polynomially. Thus, Shor's algorithm shows that quantum computing could provide exponential speedups in problems with real-world consequences, and has historically been a primary motivator for funding quantum computing research.

Given a number N which we want to factor, Shor's algorithm works by first constructing the function

$$f(x) = a^x \bmod N \tag{1.13}$$

where a is a random integer smaller than N . Then, factoring N can be reduced to finding the period of f , which is the integer r such that

$$f(x+r) = f(x) \tag{1.14}$$

for all integers x . Classically, we might evaluate f for all $x < N$, then compute the discrete Fourier transform which converts the vector of the outputs of f to a frequency domain representation. However, even just storing N values is classically expensive. With a quantum computer, we can prepare a superposition of the outputs of f using $n = \log_2 N$ qubits, then apply the *quantum Fourier transform* (QFT) which is analogous to the discrete Fourier transform for quantum amplitudes.

The key to Shor’s algorithm is that the QFT can be computed much more efficiently than the discrete Fourier transformation. In general, accessing the information in the output of the QFT still requires an exponential number of measurements. However, since our input vector is periodic by the nature of the problem, measuring the output state will yield (with high probability) a result which informs us of the period of the input. Many quantum algorithms, such as quantum phase estimation [9], take advantage of this period finding property of the QFT to achieve speedups over classical algorithms.

1.2.2 Grover’s Algorithm

Grover’s algorithm [10] is used to invert functions of the form

$$f(x) = \begin{cases} 1 & \text{if } x = \bar{x} \\ 0 & \text{else} \end{cases} \quad (1.15)$$

where \bar{x} is a member of a discrete set of N possible inputs to f . Such a function might represent searching through a list of values for one which satisfies a certain property.

Grover’s algorithm works by using *amplitude amplification*. From an equal superposition of N states constructed using $n = \log_2 N$ qubits, Grover’s algorithm iteratively and selectively increases the amplitude of $|\bar{x}\rangle$. Given enough iterations, the final measured state will be $|\bar{x}\rangle$ with high probability. The number of iterations required for Grover’s algorithm scales as \sqrt{N} , which is a quadratic speed up over a brute force classical search.

While the speed up is not exponential, a quadratic speedup is still substantial if N is sufficiently large, and Grover’s algorithm may be used to help speed up brute force searches for hard problems. Importantly, the scaling for Grover’s algorithm has been

proven to be optimal and definitively better than any algorithm which does not use quantum resources.

1.2.3 Quantum Simulation

The concept of quantum computing was first proposed by Richard Feynman [11], largely in response to the difficulty of simulating physically interesting quantum systems with classical computers. The general idea of quantum simulation is to initialize a known quantum state, evolve it under a Hamiltonian which mimics a physical system of interest, then measure observables of the final wavefunction. Proposals for interesting problems to simulate include high temperature superconductivity [12, 13], quantum phase transitions [14, 15, 16], chemical reactions [17, 18, 19], and even more exotic phenomenon such as Hawking radiation [20].

Quantum simulation generally comes in two flavors. In a *digital quantum simulation*, the evolution is divided into discrete time steps (a process called Trotterization [21]), and the propagator for each time step is approximated using the available single and two qubit gates on the hardware. *Analog quantum simulations* is a more direct approach in which the system Hamiltonian is tailored to approximate the Hamiltonian of interest, so that the desired evolution occurs naturally. Proof of principle experiments for both digital and analog simulations have been performed.

While quantum simulation is of great interest to physicists, caution should be exercised when claiming that quantum simulation will provide exponential speed up in any

given problem. Initialization, evolution, and measurement may all turn out to have significant overhead, especially when considering the presence of errors and noise and the precision required for the simulation results.

1.3 Two Approaches to Quantum Computing Research

The applications we have discussed promise significant progress for solving important problems. However, realizing these applications is beyond the capabilities of current quantum devices. In this section, we look at two approaches to quantum computing research today. The first approach is to build a scalable (to millions of qubits) device with high enough fidelity to implement *quantum error correction*, which is seen as the path to ultimately implementing substantial quantum algorithms such as Shor's algorithm. The second approach is to build a smaller device - but still larger than any universal quantum computer that exists today - with sufficient fidelity to perform a specific computational task better than today's classical computers.

1.3.1 Quantum Error Correction

Qubits are fundamentally analog devices. Information is encoded in complex amplitudes which have no restrictions other than an overall normalization. Unlike classical bits which are binary and digital, qubits are highly susceptible to noise. Real world systems will

always have noise, either from the surrounding environment or from the very instruments used to control the qubits. How can we protect qubits from these sources of error?

One of the breakthroughs in quantum information theory was the development of *quantum error correction* (QEC). Like many classical error correction schemes, QEC relies on encoding information from a single qubit into a multiqubit system. However, the destructive nature of measurement prevents us from directly copying one qubit's information onto another ². Instead, entanglement is used to construct multiqubit states which will leave telltale signs if an error occurs. For example, in the simplest possible code [23], a single logical qubit is encoded in the joint entangled state of three physical qubits

$$|0\rangle_L = |000\rangle \tag{1.16}$$

$$|1\rangle_L = |111\rangle \tag{1.17}$$

$$|\psi\rangle_L = \alpha|000\rangle + \beta|111\rangle \tag{1.18}$$

If one of the three qubits flips from $|0\rangle$ to $|1\rangle$ or vice versa, then we will measure a joint state that is neither $|000\rangle$ or $|111\rangle$, such as $|011\rangle$. Furthermore, the measured joint state can tell us which of the qubits flipped bits, e.g. the first qubit in our example.

This three qubit repetition code has one obvious drawback - we cannot distinguish one qubit flipping from two qubits flipping. Even worse, we get no error detection at all if all three qubits flip. One solution to this problem is to simply add more physical

²This feature of quantum mechanics is known as the no-cloning theorem [22].

qubits to our logical qubit - the more qubits we have, the more bit flips we can tolerate before logical failure. However, a more fundamental problem is that the repetition code does not protect us from errors in which our logical qubits acquires an unwanted phase shift between α and β . In fact, because measuring the joint state of the qubits collapses each qubit into an eigenstate of Z (i.e. $|0\rangle$ or $|1\rangle$), the measurement destroys any phase information encoded in the qubits. We could instead encode the logical qubit as

$$|0\rangle_L = |+++ \rangle \tag{1.19}$$

$$|1\rangle_L = |-- \rangle \tag{1.20}$$

where $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ and $|-\rangle = (|0\rangle - |1\rangle)/\sqrt{2}$. We would then rotate the qubits prior to measurement such that we measure in the X basis. Our logical qubit is now protected against phase flips, but at the expense of bit flip protection.

Many QECs have been developed which address these shortcomings [25, 26, 27, 28, 29], and one of the primary codes driving research today is the surface code [30, 24]. The surface code is based on topological error correcting codes developed by Bravyi and Kitaev [31]. A single logical qubit consists of many qubits arranged into a 2D checkerboard pattern, where half of the qubits are *data qubits* and the other half of the qubits are *measure qubits* which are interleaved between the data qubits. The measure qubits further alternate between measure-X and measure-Z qubits. The surface code proceeds in a repetitive cycle. On each cycle, each measure-X (measure-Z) qubit is entangled with each of its four neighbors in the X (Z) basis with four successive CNOT (CZ) gates. After

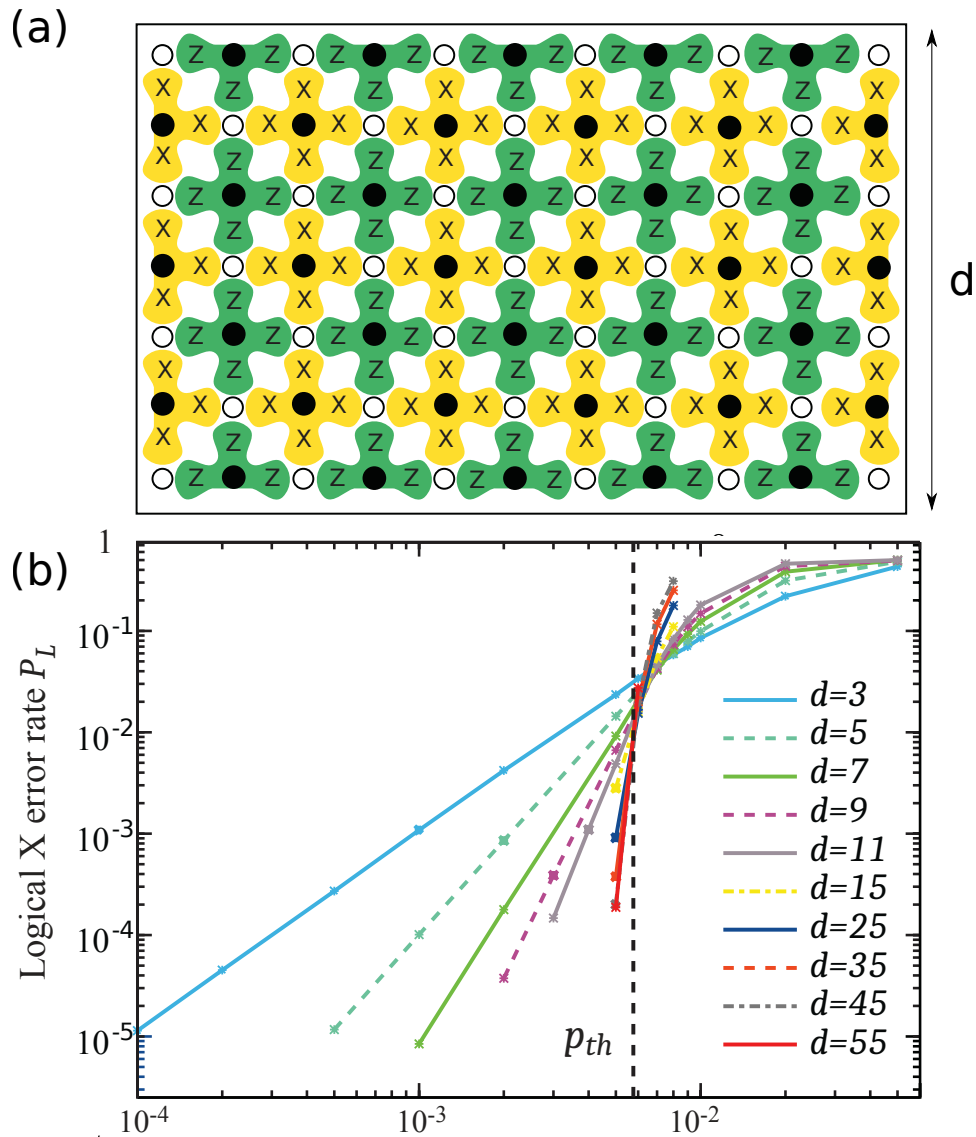


Figure 1.2: Surface code logical qubit. (a) White circles represent data qubits, while black qubits represent measure qubits. Half of the measure qubits entangle with their neighbors in the Z basis, while the other half entangle in the X basis. (b) Logical error rate versus physical error rate per qubit per step, for different sizes of logical qubits. Below the threshold error rate $p_{th} = 0.57\%$, the error rate decreases exponentially. Adapted with permission from Ref. [24]

the entanglement step, the measure qubits are measured. The result of the measurement corresponds to an eigenvalue of the joint XXXX (ZZZZ) operator on the four data qubits, which is known as a stabilizer.

Even though we do not measure the data qubits directly, the stabilizer measurement still collapses the data qubits into an eigenstate of the joint measurement basis since they are entangled with the measure qubit. However, the data qubits can simultaneously be eigenstates of both XXXX and ZZZZ because the joint operators commute even though the single qubit operators do not. This property means that we can extract both phase and bit information from the data qubits without destroying either. Every measurement cycle of the surface code collapses the data qubits into a highly entangled, simultaneous eigenstate of many joint operators, and the pattern of measurement outcomes tells us which eigenstate the data qubits are in. If the measurement outcome changes on any given cycle, we have detected an error, and different types and locations of errors will leave different signatures in the measurement. Given this knowledge, we can then correct the errors in post processing.

As with the repetition code, a surface code logical qubit has greater protection from errors with increasing number of qubits. A second important factor is the per cycle error rate, which encapsulates errors in single qubit gates, two qubit entangling gates, and measurement. The logical error rate as a function of these cycle rate and number of qubits is shown in Fig.1.2(b). Remarkably, the surface code has a fault tolerance threshold which is akin to a phase transition. Below this threshold, error suppression

becomes exponentially more effective with decreasing error rate.

A primary focus in quantum computing research today is building systems which are below the error threshold, are scalable, and will maintain low error rates as they scale. The first milestone in this research path would be to demonstrate a logical qubit as a quantum memory which would have bit flip and phase flip rates that is lower than its constituent systems.

1.3.2 Quantum Supremacy

Error correction will be necessary to do general quantum computing or to run complex algorithms such as Shor's algorithm. However, it is estimated that millions of physical qubits will be required to build enough logical qubits to do useful computations [24], while the state of the art today is tens of qubits. Nonetheless, these smaller devices can be used to answer an important question - will a quantum computer work at even a modest scale, and if so, can it do anything faster than a classical computer?

An affirmative answer to this question would demonstrate *quantum supremacy*³, a term coined by John Preskill [33]. Quantum supremacy requires two ingredients: a well defined computational task which can be performed on a quantum computer with an exponential speedup over any classical computer, and hardware capable of running the algorithm to complete this task. There are a few approaches to the computational task [34, 35]. We will discuss the approach of Ref [32], which is to sample the output of a

³This term is somewhat controversial. Other proposed terms are quantum advantage, quantum edge, or somewhat jokingly, quantum relevancy.

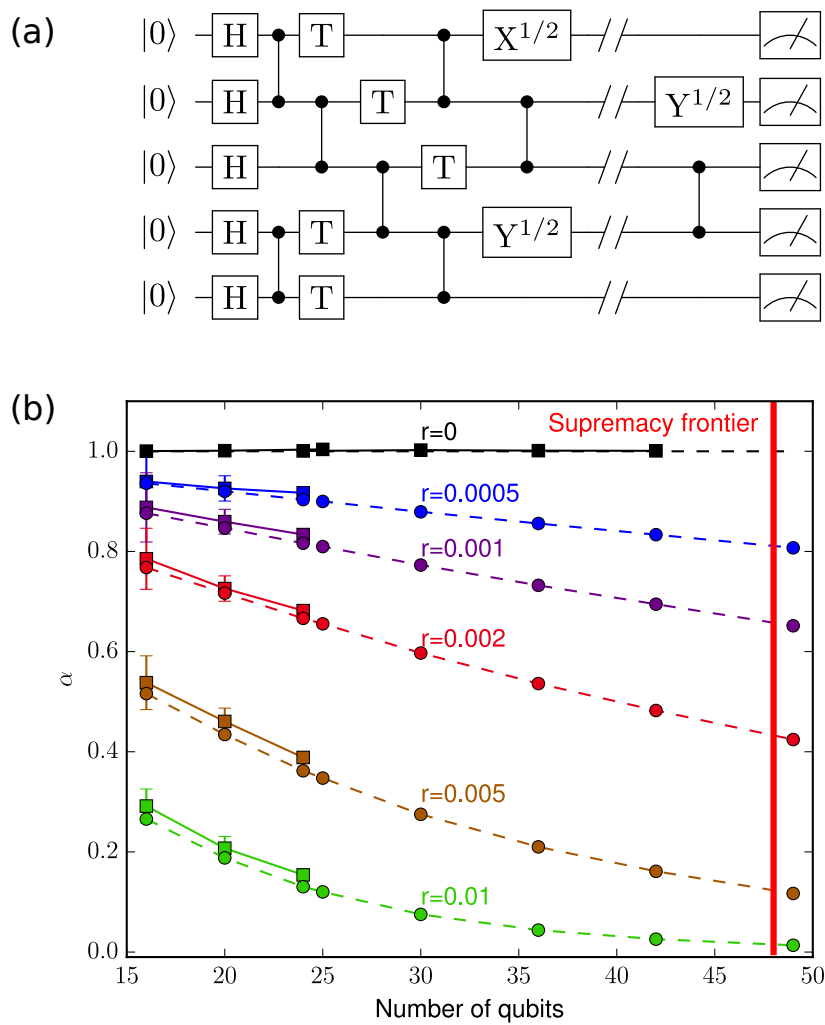


Figure 1.3: (a) An example random circuit on a 1D array of qubits. (b) Cross-entropy fidelity (α) versus number of qubits in a 2D grid, for a range of two qubit error rates, r . At 49 qubits, current classical computations can no longer accurately simulate the evolution of the system. Adapted with permission from Ref. [32]

random quantum circuit.

A random quantum circuit is implemented on an $N \times N$ 2D grid of qubits by first initializing all of the qubits in superposition states, then alternating between rounds of 2 qubit entangling gates and single qubit gates ⁴. The goal is to generate a quantum state where the 2^{N^2} amplitudes of the total wavefunction are completely randomly distributed. Randomization requires that every qubit has the opportunity to be entangled with every other qubit, and this can be done if the number of two qubit entangling rounds is roughly at least N . For a large enough system, the final random quantum state will exhibit an *exponential distribution* over its measurement probabilities ⁵. We cannot experimentally characterize the entire distribution because the number of possible bitstrings is intractable. However, the exponential distribution implies that certain measurement outcomes are far more likely to occur than others, and if our device is working correctly, the bitstrings which we sample should be biased towards those in the high probability part of the distribution.

Separately, a classical computer is also tasked with sampling from the same random quantum state. Because the random quantum circuit has no structure, the only general algorithm is to simulate the entire quantum circuit, a task that scales exponentially with the number of qubits ⁶. For a sufficiently large number of qubits - currently around 50

⁴Importantly, some of the gates should not be Clifford gates, and this can be achieved by using T gates, i.e. $\pi/4$ rotations about Z .

⁵This observation about random quantum states is related to work by Wigner on energy levels and couplings in complex nuclear systems [36]. It can formally be obtained via random matrix theory [37], and is also referred to as the Porter-Thomas distribution [38].

⁶It has not been definitively proven that random quantum sampling is exponentially difficult for a classical computer. However, if this was not the case, there would be interesting and unexpected consequences for computational complexity theory [39]

for state of the art algorithms on a powerful supercomputer - this computation becomes impossible, giving us a path to demonstrating quantum supremacy.

However, up until we reach this threshold number of qubits, the classical simulation will be necessary to confirm that the quantum device is behaving as expected. Because the quantum supremacy device will likely be too small to use error correction, we can tolerate at most one total error in the execution of the entire random circuit. Figure 1.3(b) shows the cross entropy - a measure of the sampling fidelity of the quantum device - versus the error rate per two qubit entangling gate, the most error prone part of the quantum circuit. In order to comfortably reach quantum supremacy we would like to two qubit error rates below 5×10^{-3} on a device with at least 50 qubits.

1.3.3 Outlook

The error correction and supremacy proposals we have discussed are different in some important ways. For example, the surface code operates on a repetitive measurement cycle, while quantum supremacy would only require one-shot experiments with one measurement. The error requirements for quantum supremacy are somewhat more stringent, while the surface code is particularly sensitive to leakage to non-computational states. However, the two approaches are fundamentally similar in that they require high fidelity control over a system of qubits with at least a 2D grid of connectivity.

1.4 Qubit Implementations

The possibility of solving classically intractable problems has led to many experimental advances in controllable quantum systems. In this section we will take a non-exhaustive look at some of the popular approaches to building a qubit systems today.

1.4.1 Trapped Ions

Nature readily gives us stable, discrete quantum systems - atoms. Atoms can be isolated by first ionizing them, then trapping them in an electromagnetic field within an ultra-high vacuum environment [40, 41, 42]. Historically, trapped ions were first used as highly stable clocks because the energy levels of the ions were highly stable and predictable [43, 44]. These properties also lend themselves to high fidelity single qubit manipulation. The qubit basis is formed using two of the ion's energy levels, chosen based on a balance of stability and ease of manipulation and measurement. Single qubits can be controlled using microwaves [45, 46] or lasers [47]. The qubit state is measured by driving a transition which causes the ion to fluoresce conditional on the qubit state [48]. Two qubit interactions are generally mediated by the Coulomb force between two ions, and high fidelity two qubit gates have been achieved by conditionally driving a vibrational mode between the two qubits [47].

The challenge in building a trapped ion quantum computer is scaling up the number of qubits. One focus of current research is to microfabricate traps on chip to reduce the footprint of the trap [49, 50]. However, a single trap can only support a finite number of

ions, and a full scale quantum computer will likely require connecting multiple separate traps with photonic interconnects [51, 52].

1.4.2 Semiconductor Qubits

A more recent approach to building qubits is to create artificial atom-like objects in solid state materials. The hope is that these qubits can be created in large quantities by leveraging modern semiconductor fabrication technology, allowing the technology to scale. DiVencenzo himself proposed that qubits could be created out of quantum dots [53], which are nanometer sized confinements created by electrically gating a 2D or 1D semiconductor structure. The dots are electrically tuned so that they contain a stable electron configuration, and the spin of the electron(s) can then be used as a qubit. Control is performed using standard electron spin resonance (ESR) techniques [54], while readout is achieved using a technique called spin-to-charge conversion [55].

Qubits can also be created by introducing dopants and defects into semiconductors. The nitrogen-vacancy (NV) center in diamond [56] is a prime example of this type of quantum system. The NV center consists of a substitutional nitrogen paired with a vacancy defect in a diamond carbon lattice, and the result is a spin-1 system which contains both optical transitions and a tunable microwave transition. Remarkably, the NV center retains its quantum properties even at room temperature. Other popular defect systems include vacancy defects in silicon carbide [57, 58], and phosphorus donors in silicon [59].

Two qubit interactions remain a challenge for spin qubits because the interactions between single electrons is inherently weak, though progress has been made in executing two qubit gates on quantum dots. Another difficulty in scaling up spin qubit systems is that the qubits are often very sensitive to their microscopic environment. Thus, spin qubits may vary between manufacturing runs and even from cooldown to cooldown. Aside from quantum computing, spin and defect qubits show promise for other practical quantum applications, such as magnetic field sensing [60, 61], and quantum key distribution [62].

1.4.3 Superconducting Qubits

Trapped ions and spin qubits derive their quantum properties from fundamental particles of nature. Superconducting qubits lie on the other end of the spectrum - they are macroscopic and man-made objects which exhibit quantum behavior. Superconducting qubits are radio frequency circuits which are fabricated out of a superconducting metal such as aluminum or niobium. When cooled to sub-kelvin temperatures inside a dilution refrigerator, the metals become superconducting, which has two important consequences. First, a metal in a superconducting state has zero resistance, and low dissipation is a requirement for quantum coherence since quantum mechanics is reversible. Second, electrons in a superconducting metal are bound into pairs called Cooper pairs, and because Cooper pairs are bosons, they can condense into the same quantum ground state. Therefore, the electrical properties of a superconducting circuit, such as the charge or phase, can be considered to be single, macroscopic quantum properties of the circuit.

Superconducting qubits have a number of attractive qualities for quantum computing. Because they are man-made circuits, superconducting qubits are not constrained by natural constants such as atomic transitions or dipole moments. The qubit transition frequency and qubit-qubit couplings can be designed by simply varying capacitances and inductances in the circuit, and can even be tuned in situ with electrical signals. Like semiconductor qubits, superconducting qubits can also leverage existing fabrication technology to make large quantities of qubits. Somewhat counterintuitively, the macroscopic nature of superconducting qubits is actually an advantage for scaling up, because the larger size affords more space for routing control and measurement wires.

Historically, the primary disadvantage of superconducting qubits is their short coherence times. The tunability and large size of superconducting qubits makes them highly susceptible to environmental noise and materials defects. However, qubit coherence has improved dramatically in recent years, thanks to improvements in materials quality and in the design of the qubits. Of particular note is the introduction of the transmon qubit by researchers at Yale [63], which solved an outstanding problem of sensitivity of qubit devices to charge noise.

1.5 Summary and Thesis Outline

Quantum computing exploits the features of quantum mechanics to store and manipulate information, and could potentially solve certain problems faster than classical computing. Moreover, precisely controlling a quantum system with many degrees of freedom would

in and of itself be an important milestone for experimental physics. Current research in quantum computing is focused on scaling up the number of qubits while maintaining high fidelities in order to achieve fault tolerance for error correction and quantum supremacy. There are many approaches to building quantum systems, each with their own merits and drawbacks.

Superconducting qubits have seen particularly exciting progress in recent years. They have improved to the point that their two qubit fidelities are among the best of any implementation [64, 65]. Furthermore, recent devices indicate that these fidelities can be maintained when increasing the number of qubits, enabling proof of principle experiments in error correction [66, 67]. This thesis will focus on the implementation of single qubit gates, two qubit gates, and readout in superconducting qubits. These components are the basis for any quantum algorithm, and all of them need to be high fidelity operations to achieve fault tolerance or quantum supremacy. In Chapter 2, we will review the design of the transmon circuit . Then, in Chapter 3, we will survey the techniques used to fabricate these circuits, and in Chapter 4, we will focus on the infrastructure we use control superconducting qubits. In Chapter 5, we will detail the calibration procedure for single qubit gates. In Chapters 6 and 7, we will benchmark the result of our single qubit calibration procedure using a variety of metrological tools. In Chapter 8, we will calibrate two qubit gates then apply the same metrological tools to characterize two qubit gate fidelity. Finally, in Chapter 9, we will explore the complicated dynamics of dispersive readout and the limits it places on high fidelity qubit measurement.

Chapter 2

Superconducting Transmon Qubits

In this chapter, we introduce the basic building block for our experiments, the transmon qubit. We start with a review of quantum mechanics as it applies to electrical oscillators, then apply similar concepts to derive the energy level structure for a transmon. Next, we discuss how to couple quantum circuits, both to the outside world for control and measurement, and to other quantum devices. Finally, we look at the constraints on various device parameters such as frequencies and control couplings, and list the device parameters that we typically use in experiments. The primary sources for this chapter were Daniel Sank's thesis [68], notes by Steve Girvin [69] and Michel Devoret [70], and the original paper by Koch et al. describing the transmon [63].




Figure 2.1: A simple LC oscillator.

2.1 LC Circuit

To understand how quantum mechanics applies to electrical circuits, we begin with the oscillating LC circuit shown in Fig. 2.1. We identify the potential energy of the circuit as the energy stored in the capacitor which has charge Q , and kinetic energy as due to the current \dot{Q} flowing in the inductor

$$U = \frac{Q^2}{2C} \tag{2.1}$$

$$T = \frac{L\dot{Q}^2}{2}. \tag{2.2}$$

We then write the Lagrangian:

$$\mathcal{L} = T - U \tag{2.3}$$

$$= \frac{L\dot{Q}^2}{2} - \frac{Q^2}{2C} \tag{2.4}$$

and after applying the Euler-Lagrange equation, we confirm that the resulting equations of motion match our expectation from applying Kirchoff's laws

$$\frac{\partial \mathcal{L}}{\partial Q} - \frac{\partial}{\partial t} \frac{\partial \mathcal{L}}{\partial \dot{Q}} = 0 \quad (2.5)$$

$$\frac{Q}{C} + L\ddot{Q} = 0 \quad (2.6)$$

$$\omega^2 Q + \ddot{Q} = 0. \quad (2.7)$$

where $\omega = 1/\sqrt{LC}$ is the resonance frequency of the circuit. The conjugate coordinate to Q is

$$\frac{\partial \mathcal{L}}{\partial \dot{Q}} = L\dot{Q} = \Phi, \quad (2.8)$$

where Φ is the flux through the inductor. We then arrive at the Hamiltonian for the LC circuit

$$H = \Phi\dot{Q} - \mathcal{L} = \frac{\Phi^2}{2L} + \frac{Q^2}{2C} \quad (2.9)$$

which is the familiar Hamiltonian for a harmonic oscillator.

So far, we have chosen the charge Q as our "position" coordinate. Thus, in an analogous mechanical oscillator, Φ would represent the momentum, and L and C would represent the mass and inverse of the spring constant, respectively. However, Eqn. 2.9 is symmetric in the position and momentum coordinates, and we will see later on that it is more convenient to choose Φ as the position coordinate. Earlier, we interpreted Φ as the flux threading the inductor. Alternatively, we may define Φ as the integral of the voltage

at a node in the circuit

$$\Phi(t) = \int^t V(\tau) d\tau. \quad (2.10)$$

This quantity is known as the *node flux* or *branch flux*. Starting with Φ as the position coordinate, we would reverse the position of the minus sign in the Lagrangian in Eqn. 2.4, but would eventually arrive at the same Hamiltonian as Eqn. 2.9.

The LC oscillator has all of the familiar properties of a harmonic oscillator. The two conjugate coordinates follow the canonical commutation relation

$$[\Phi, Q] = i\hbar \quad (2.11)$$

and the Hamiltonian can be rewritten in terms of raising and lowering operators

$$H = \hbar\omega \left(a^\dagger a + \frac{1}{2} \right) \quad (2.12)$$

$$a = \sqrt{\frac{1}{2\hbar Z}} (\Phi + iZQ) \quad (2.13)$$

$$a^\dagger = \sqrt{\frac{1}{2\hbar Z}} (\Phi - iZQ) \quad (2.14)$$

where a^\dagger and a are the raising and lowering operators, and $Z = \frac{1}{\sqrt{LC}}$ is the characteristic impedance of the circuit. We can also rewrite the charge and flux operators in terms of




Figure 2.2: The transmon qubit.

the raising and lowering operator

$$Q = i\sqrt{\frac{\hbar}{2Z}}(a^\dagger - a) \quad (2.15)$$

$$= iQ_{\text{zpf}}(a - a^\dagger) \quad (2.16)$$

$$\Phi = \sqrt{\frac{\hbar Z}{2}}(a^\dagger + a) \quad (2.17)$$

$$= \Phi_{\text{zpf}}(a + a^\dagger) \quad (2.18)$$

where $Q_{\text{zpf}} = \sqrt{\langle 0|Q^2|0\rangle} = \sqrt{\hbar/2Z}$ is the zero point fluctuation of the charge operator, and $\Phi_{\text{zpf}} = \sqrt{\hbar Z/2}$ is likewise the zero point fluctuation of the phase operator.

2.2 Transmon Hamiltonian

Up to now, we have studied a linear LC circuit where the energy levels are evenly spaced with energy difference $\hbar\omega$. This linear energy structure is not suitable for performing quantum computation, because no two energy levels can be addressed by classical controls without also addressing other transitions in the ladder. To obtain a nonlinear energy level structure, we require a nonlinear circuit element. Fortunately, a nonlinear superconducting circuit element exists: the Josephson junction [71]. The Josephson junction

(JJ) consists of two superconducting leads connected by a weak link, typically an thin (few nanometer) insulating barrier. A JJ is governed by the Josephson relations:

$$V(t) = \frac{\hbar}{2e} \frac{\partial \phi}{\partial t} \quad (2.19)$$

$$I(t) = I_c \sin \phi(t) \quad (2.20)$$

where ϕ is the phase difference across the junction and I_c is the critical current of the junction, above which the junction acts like a normal resistor. By combining Eqn. 2.10 and Eqn. 2.20, we see that ϕ is related to the branch flux by $\phi = 2e/\hbar\Phi = 2\pi\Phi/\Phi_0$, where Φ_0 is the flux quantum. As current flows through the junction, it accumulates energy, which we can find by integrating the power

$$U = \int_0^t V(t)I(t)dt \quad (2.21)$$

$$= \frac{\hbar}{2e} \int I_c \sin \phi \frac{\partial \phi}{\partial t} dt \quad (2.22)$$

$$= \frac{\hbar}{2e} \int I_c \sin \phi d\phi \quad (2.23)$$

$$= -E_J \cos \phi. \quad (2.24)$$

We find that the energy stored in the junction is proportional to the cosine of the phase difference, with an energy scale $E_J = \Phi_0 I_c / 2\pi$, also called the Josephson energy.

We now replace the linear inductor of our LC circuit with the nonlinear Josephson

inductance as shown in Fig. 2.2, and obtain the following Hamiltonian for the transmon

$$H = \frac{Q^2}{2C} - E_J \cos \frac{2\pi\Phi}{\Phi_0} \quad (2.25)$$

which is also commonly written as

$$H = 4E_c n^2 - E_J \cos \phi \quad (2.26)$$

where $n = Q/2e$ is the number of Cooper pairs, $E_c = e^2/2C$ is the charging energy, and the variables n and ϕ have the commutation relation $[\phi, n] = 1$.

To understand the transmon Hamiltonian, we expand the cosine term assuming that fluctuations in ϕ are small

$$\cos\phi \approx 1 - \phi^2 + \phi^4/24 + \dots \quad (2.27)$$

Ignoring the constant term, to first order we obtain the harmonic oscillator Hamiltonian

$$H \approx 4E_c n^2 + E_J \phi^2 \quad (2.28)$$

where the energy levels form an evenly spaced ladder with energy spacing $\sqrt{8E_J E_c}$. Next, we consider the effect of the quartic term perturbatively by writing it in terms of harmonic oscillator ladder operators

$$\phi = \left(\frac{2E_c}{E_J} \right)^{1/4} (a^\dagger + a) \quad (2.29)$$

$$-E_J \frac{\phi^4}{24} = -\frac{E_c}{12} (a^\dagger + a)^4. \quad (2.30)$$

We now apply this correction term to the m th first order oscillator basis state

$$\Delta E_m = -\frac{E_C}{12} \langle m | (a^\dagger + a)^4 | m \rangle \quad (2.31)$$

$$= -\frac{E_C}{12} (6m^2 + 6m + 3) \quad (2.32)$$

$$E_m = m\sqrt{8E_J E_C} - \frac{E_C}{12} (6m^2 + 6m + 3). \quad (2.33)$$

Finally, we compute the difference between neighboring energy levels

$$E_m - E_{m-1} = -mE_C. \quad (2.34)$$

We find that the energy spacing decreases as we go up the ladder in energy. Thus, we can use the ground and first excited states of the energy ladder as our two level qubit, which will have energy

$$E_{10} = \sqrt{8E_J E_C} - E_C, \quad (2.35)$$

and the difference between E_{10} and E_{21} is simply $E_{21} - E_{10} = -E_C$, the charging energy.

This quantity is also known as the anharmonicity, η .

The anharmonicity sets a speed limit on how fast we can control our qubit - faster control pulses are more spectrally broad, and if the width of the pulse is larger than η in frequency space, we can induce undesired transitions to the second excited state. Based on this property, we might be tempted to make E_C as large as possible by making the capacitance as small as possible. However, up to now we have neglected a term in the transmon Hamiltonian in Eqn. 2.26 due to external voltages on the capacitor, which we




Figure 2.3: An LC oscillator coupled to a voltage drive.

write as follows

$$H = 4E_c (n - n_g)^2 - E_J \cos \phi \quad (2.36)$$

where n_g is known as the gate charge. This gate charge may be due to external voltage biases applied by the experimenter, but can also come from charge fluctuations in the environment. Fluctuations in the gate charge lead to fluctuations in the energy levels and are thus a source of decoherence. In fact, $1/f$ charge noise was the dominant source of decoherence in early superconducting qubit devices.

The key to the design of the transmon is that the sensitivity of the energy levels to the gate charge is proportional to $e\sqrt{8E_J/E_C}$ - in other words, decreasing the charging energy by adding additional capacitance to the circuit exponentially decreases the sensitivity to charge noise. Thus, picking transmon parameters is largely a matter balancing the ratio of E_J/E_C so that our device is sufficiently anharmonic but insensitive to charge noise.

2.3 Driving

We now turn to the mechanisms for controlling and coupling transmon qubits. Because transmons are nearly harmonic, we will use the LC oscillator as a stand-in for under-

standing the basic physics, then truncate down to the two lowest levels to obtain the Hamiltonian for a qubit. First consider the circuit shown in Fig. 2.3, consisting of a voltage source with output impedance Z_d coupled to an LC circuit by a capacitor C_d . If there is no voltage on the driving circuit, the Lagrangian and Hamiltonian of the total circuit is

$$\mathcal{L} = \frac{1}{2}C\dot{\Phi}^2 + \frac{1}{2}C_d\dot{\Phi}^2 - \frac{\Phi^2}{2L} \quad (2.37)$$

$$H = \frac{Q^2}{2C_\Sigma} + \frac{\Phi^2}{2L} \quad (2.38)$$

and the conjugate variable to Φ is now slightly modified to be $Q = C_\Sigma\dot{\Phi}$, where $C_\Sigma = C + C_d$ is the total capacitance to ground. Next, we add a time dependent voltage $V_d(t)$ due to the driving circuit

$$\mathcal{L}_{\text{driven}} = \frac{1}{2}C_\Sigma\dot{\Phi}^2 + \frac{1}{2}C_d(V_d(t) - \dot{\Phi})^2 - \frac{\Phi^2}{2L} \quad (2.39)$$

$$= \mathcal{L} + \frac{1}{2}C_dV_d^2 - C_dV_d(t)\dot{\Phi}. \quad (2.40)$$

We can ignore the $\frac{1}{2}C_dV_d(t)^2$ term because it does not contain Φ or $\dot{\Phi}$. We now recompute the conjugate to Φ

$$\tilde{Q} = C_\Sigma\dot{\Phi} - C_dV_d(t) \quad (2.41)$$

and express the Hamiltonian as a function of \tilde{Q} and Φ

$$H_{\text{driven}} = \tilde{Q}\dot{\Phi} - \mathcal{L}_{\text{driven}} \quad (2.42)$$

$$= \frac{1}{2}C_{\Sigma}\dot{\Phi}^2 + \frac{\Phi^2}{2L}. \quad (2.43)$$

Next we invert Eqn. 2.41 to get $\dot{\Phi}$ as a function of \tilde{Q} , and substitute into Eqn.2.43 to get

$$H_{\text{driven}} = \frac{\tilde{Q}^2}{2C_{\Sigma}} + \frac{\Phi^2}{2L} + \frac{C_d}{C_{\Sigma}}V_d\tilde{Q} + \frac{C_d^2V_d(t)^2}{2C_{\Sigma}}. \quad (2.44)$$

Again, we can neglect the last term which does not include either coordinate. We see that the driving term couples to the momentum of the oscillator. We now treat this term as a small perturbation on the undriven Hamiltonian and assume $\tilde{Q} \approx Q$. Then, using Eqn. 2.18, we write the Hamiltonian in terms of raising and lowering operators

$$H_{\text{driven}} = \hbar\omega a^{\dagger}a + \frac{1}{2} - i\frac{C_dV_d(t)}{C_{\Sigma}}Q_{\text{zpf}}(a - a^{\dagger}). \quad (2.45)$$

Having derived the Hamiltonian for an oscillator, we can truncate down to the lowest two levels to get an approximate Hamiltonian for a two level qubit. In converting the oscillator Hamiltonian to a qubit Hamiltonian, we use the following relations between

raising and lowering operators and the qubit operators in the Pauli basis

$$a \rightarrow \sigma^- \quad (2.46)$$

$$a^\dagger \rightarrow \sigma^+ \quad (2.47)$$

$$a + a^\dagger \rightarrow \sigma_x \quad (2.48)$$

$$i(a^\dagger - a) \rightarrow \sigma_y \quad (2.49)$$

$$1 - 2a^\dagger a \rightarrow \sigma_z \quad (2.50)$$

and obtain the following qubit Hamiltonian

$$H_{\text{qubit}} = H_0 + H_d \quad (2.51)$$

$$H_0/\hbar = -\frac{\omega}{2}\sigma_z \quad (2.52)$$

$$H_d = \Omega f(t)\sigma_y, \quad (2.53)$$

where $V_d = V_0 f(t)$ and Ω is known as the *Rabi frequency*, and is given by

$$\Omega = \frac{C_d V_0 Q_{\text{zpf}}}{\hbar(C + C_d)}. \quad (2.54)$$

First, consider the static part of the Hamiltonian, H_0 . The time evolution of a state $|\psi(t)\rangle$ is given by the solution to Schrodinger's equation for a time independent

Hamiltonian

$$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = H_0 |\psi(t)\rangle \quad (2.55)$$

$$|\psi(t)\rangle = \exp\left[\frac{-iH_0 t}{\hbar}\right] |\psi(0)\rangle \quad (2.56)$$

$$= U |\psi(0)\rangle \quad (2.57)$$

where U is the *propagator*, a unitary operator describing the evolution of the system. Often, we are interested in only the effect the driving, time dependent part of the Hamiltonian. In other words, we want the time evolution of the system after accounting for the effect of U . We consider a new time dependent basis

$$|\Theta(t)\rangle = U^\dagger |\psi(t)\rangle, \quad (2.58)$$

which removes the time dependence of $|\psi\rangle$. The time evolution of this new basis under the Hamiltonian $(H_0 + H_d)$ is

$$i\hbar \frac{\partial}{\partial t} |\Theta(t)\rangle = i\hbar \left(\frac{\partial U^\dagger}{\partial t} |\psi(t)\rangle + U^\dagger \frac{\partial}{\partial t} |\psi(t)\rangle \right) \quad (2.59)$$

$$= -H_0 U U^\dagger |\Theta(t)\rangle + U^\dagger (H_0 + H_d) U |\Theta(t)\rangle \quad (2.60)$$

$$= U^\dagger H_d U |\Theta(t)\rangle \quad (2.61)$$

where we used the fact that U commutes with H and $\partial U^\dagger / \partial t = (iH_0 / \hbar) U$. This procedure is known as moving to the "rotating frame" or the "interaction picture". From Eqn. 2.53,

we now have

$$H' = \hbar\Omega f(t)e^{-i\omega t\sigma_z/2}\sigma_y e^{i\omega t\sigma_z/2} \quad (2.62)$$

$$= \hbar\Omega f(t) \begin{pmatrix} 0 & -ie^{-i\omega t} \\ ie^{i\omega t} & 0 \end{pmatrix} \quad (2.63)$$

$$= i\hbar\Omega f(t) (e^{i\omega t}\sigma^+ - e^{-i\omega t}\sigma^-). \quad (2.64)$$

Now, suppose $f(t)$ is a continuous wave drive on resonance with the qubit

$$f(t) = \sin \omega t + \phi \quad (2.65)$$

$$= -i (e^{i(\omega t + \phi)} - e^{-i(\omega t + \phi)}) / 2. \quad (2.66)$$

Substituting into Eqn. 2.64, and ignoring terms which rotate at 2ω ¹, we have

$$H'/\hbar = -\frac{\Omega}{2} (e^{-i\phi}\sigma^+ + e^{i\phi}\sigma^-) \quad (2.67)$$

$$= -\frac{\Omega}{2} (\cos \phi \sigma_x - \sin \phi \sigma_y). \quad (2.68)$$

In the Bloch sphere picture, an on resonance drive causes the state to rotate about the equator of the sphere, with the axis of rotation determined by the phase of the drive. To

¹We ignore these terms because they oscillate quickly and average out to zero.




Figure 2.4: Two coupled LC oscillators.

determine the rotation frequency, let us assume $\phi = 0$. Then, we have

$$U' = e^{-iH't/\hbar} \quad (2.69)$$

$$= \begin{pmatrix} \cos \Omega t/2 & i \sin \Omega t/2 \\ i \sin \Omega t/2 & \cos \Omega t/2 \end{pmatrix}. \quad (2.70)$$

If the qubit begins in the ground state, then at time t the state of the qubit is

$$U'|0\rangle = \cos \frac{\Omega t}{2} |0\rangle + i \sin \frac{\Omega t}{2} |1\rangle. \quad (2.71)$$

The qubit completes a full cycle on the Bloch sphere when $\Omega t = 2\pi$, and hence the rate of rotation is $\Omega/2\pi$, the Rabi frequency.

2.4 Coupling

2.4.1 Oscillator-Oscillator Coupling

Next we turn to two coupled LC circuits as shown in Fig. 2.4. We first consider only the kinetic part of the Lagrangian

$$T = \frac{1}{2}C_1\dot{\Phi}_1^2 + \frac{1}{2}C_2\dot{\Phi}_2^2 + \frac{1}{2}C_g(\dot{\Phi}_2 - \dot{\Phi}_1)^2 \quad (2.72)$$

$$= \frac{1}{2}(C_1 + C_g)\dot{\Phi}_1^2 + \frac{1}{2}(C_2 + C_g)\dot{\Phi}_2^2 - C_g\dot{\Phi}_2\dot{\Phi}_1 \quad (2.73)$$

Equation 2.73 can be rewritten in matrix form as

$$T = \frac{1}{2}\dot{\Phi}^T C \dot{\Phi} \quad (2.74)$$

$$\Phi = \begin{pmatrix} \Phi_1 \\ \Phi_2 \end{pmatrix} \quad (2.75)$$

$$C = \begin{pmatrix} C_1 + C_g & -C_g \\ -C_g & C_2 + C_g \end{pmatrix} \quad (2.76)$$

The conjugate charges to the vector Φ are

$$Q = \begin{pmatrix} \frac{\partial \mathcal{L}}{\partial \dot{\Phi}_1} \\ \frac{\partial \mathcal{L}}{\partial \dot{\Phi}_2} \end{pmatrix} \quad (2.77)$$

$$= \begin{pmatrix} (C_1 + C_g)\dot{\Phi}_1 - C_g\dot{\Phi}_2 \\ -C_g\dot{\Phi}_1 + (C_2 + C_g)\dot{\Phi}_2 \end{pmatrix} \quad (2.78)$$

$$Q = C\dot{\Phi} \quad (2.79)$$

$$C^{-1}Q = \dot{\Phi} \quad (2.80)$$

Using the fact that transposing a matrix product reverses the multiplication order, we can now rewrite Eqn. 2.73 as

$$T = \frac{1}{2}Q^T C^{-1}Q, \quad (2.81)$$

and it follows that the total Hamiltonian is

$$H = \frac{1}{2}Q^T C^{-1}Q + \frac{\Phi_1^2}{2L_1} + \frac{\Phi_2^2}{2L_2} \quad (2.82)$$

The inverse capacitance matrix is

$$C^{-1} = \frac{1}{C_1 C_2 + C_1 C_g + C_2 C_g} \begin{pmatrix} C_2 + C_g & C_g \\ C_g & C_1 + C_g \end{pmatrix} \quad (2.83)$$

and the charge portion of the Hamiltonian can be written as

$$H_Q = \frac{Q_1^2}{2C'_1} + \frac{Q_2^2}{2C'_2} + \frac{Q_1 Q_2}{C'_g} \quad (2.84)$$

$$(2.85)$$

where the new capacitances are

$$C'_1 = C_1 + \frac{C_2 C_g}{C_2 + C_g} \quad (2.86)$$

$$C'_2 = C_2 + \frac{C_1 C_g}{C_1 + C_g} \quad (2.87)$$

$$C'_g = \frac{C_1 C_2 + C_1 C_g + C_2 C_g}{C_g}. \quad (2.88)$$

Using Eqn. 2.18, we can write the coupling Hamiltonian as

$$H_g = \frac{Q_{\text{zpf},1} Q_{\text{zpf},2}}{C'_g} (a_1^\dagger - a_1)(a_2 - a_2^\dagger) \quad (2.89)$$

$$= \frac{\hbar}{2\sqrt{Z_1 Z_2} C'_g} (a_1^\dagger - a_1)(a_2 - a_2^\dagger). \quad (2.90)$$

We use the fact that $Z = 1/\omega C$, and fold the prefactor into a new quantity, the coupling

strength g

$$H_g = \hbar g(a_1^\dagger - a_1)(a_2 - a_2^\dagger) \quad (2.91)$$

$$g = \sqrt{\omega_1 \omega_2} \frac{\sqrt{C_1' C_2'}}{2C_g'} \quad (2.92)$$

$$= \sqrt{\omega_1 \omega_2} \frac{C_g}{2\sqrt{(C_1 + C_g)(C_2 + C_g)}}. \quad (2.93)$$

2.4.2 Oscillator-Qubit Coupling

Having derived the coupling term for two oscillators, we now consider the case of an oscillator coupled to a qubit. We will first take the standard mathematical approach to solving this system, then consider a graphical approach which allows us to easily account for the higher states of the transmon.

From Eqn.2.93, we identify the ladder operators a and a^\dagger with the lowering and raising operators σ^- and σ^+ , and rewrite the coupling Hamiltonian as

$$H_g = \hbar g(a^\dagger - a)(\sigma^- - \sigma^+) \quad (2.94)$$

$$H_g = \hbar g(a^\dagger \sigma^- + a \sigma^+ - a^\dagger \sigma^+ - a \sigma^-) \quad (2.95)$$

The third and fourth term do not conserve the total number of excitations in the oscillator-qubit system, and in the *rotating wave approximation*, these terms are ne-

glected. The total Hamiltonian is now

$$\frac{H}{\hbar} = \omega_r a^\dagger a - \frac{1}{2} \omega_q \sigma_z + g(a^\dagger \sigma^- + a \sigma^+). \quad (2.96)$$

Equation 2.96 is the Jaynes-Cummings Hamiltonian [72], which is frequently used in quantum optics to describe the interaction between an atom and a resonant cavity. The important parameters in the Jaynes-Cummings Hamiltonian are the detuning $\Delta = \omega_q - \omega_r$ and the coupling strength g . We will mostly work in the dispersive limit where $|\Delta| \gg |g|$. In this limit, the coupling term of the Hamiltonian can be found via perturbation theory to be

$$\frac{H_g}{\hbar} \approx -\frac{g^2}{\Delta} \sigma_z a^\dagger a \quad (2.97)$$

$$= \chi \sigma_z a^\dagger a, \quad (2.98)$$

where $\chi = -g^2/\Delta$ is known as the dispersive shift.

The effect of Eqn 2.98 can be seen by rewriting the total Hamiltonian as

$$\frac{H}{\hbar} = \omega_r a^\dagger a - \frac{1}{2} (\omega_q - 2\chi a^\dagger a) \sigma_z \quad (2.99)$$

where we find that the frequency of the qubit decreases by 2χ for every excitation in the resonator. This effect is commonly known as the *AC Stark effect*. Alternatively, we can




Figure 2.5: The first three levels of the Jaynes-Cummings ladder for a transmon.

write

$$\frac{H}{\hbar} = (\omega_r + \chi\sigma_z)a^\dagger a - \frac{1}{2}\omega_q\sigma_z \quad (2.100)$$

and observe that the frequency of the resonator depends on the state of the qubit, with the difference in resonator frequency between the two qubit states equal to 2χ .

The Jaynes Cummings Ladder for a Transmon

To visualize the levels of the transmon and resonator, we use a construction called the Jaynes-Cummings (JC) ladder. The JC ladder is built from the joint states of the qubit and resonator, which we denote by $|q, n\rangle$ where q is the number of excitations in the qubit and n is the photon occupation of the resonator. For this example, we will assume that $\omega_r > \omega_q$ and thus $\Delta < 0$.

At the bottom of the ladder is the joint ground state $|0,0\rangle$. Above the ground state, either the qubit can be excited with joint state $|1,0\rangle$, or the resonator can be excited with joint state $|0,1\rangle$. These two states comprise the first "rung" of the ladder since these two states are close in energy. The next rung contains the joint states with two excitations, which include $|0,2\rangle$ and $|1,1\rangle$, as well as $|2,0\rangle$ since the transmon has higher excited states. The ladder continues upward with increasing excitation number, in principle to infinity since the resonator has no limit on its number of bound states.

Next, we consider how the states within the ladder are coupled. In our ladder picture, the effect of the rotating wave approximation is to restrict couplings to only occur within the each rung of the ladder, under the assumption that states in different rungs are far detuned. For this reason, these rungs are also referred to as *RWA strips*. Furthermore, because the transmon is nearly harmonic, we restrict couplings to only occur between neighboring transmon states.

We also have to determine the coupling strengths between states within each rung. The nominal coupling strength g corresponds to the coupling between $|1,0\rangle$ and $|0,1\rangle$. Moving up the ladder, additional excitations enhance the couplings by \sqrt{m} where $m = n + q$ is the number of excitations for each state in the rung. Here we have assumed that adding a resonator excitation or adding a transmon excitation are equivalent from the standpoint of coupling strength. However, the transmon is not completely harmonic so the \sqrt{q} scaling of the charge matrix element is only an approximation.

The effect of each coupling is to pairwise repel each state's energy by the square of

the coupling strength divided by the detuning between the states. For example, in the single excitation rung of the JC ladder, the detuning is

$$E_{|1,0\rangle} - E_{|0,1\rangle} = \omega_q - \omega_r = \Delta. \quad (2.101)$$

Therefore, each state is repelled by

$$\overline{E_{|0,1\rangle}} - E_{|0,1\rangle} = -(\overline{E_{|1,0\rangle}} - E_{|1,0\rangle}) = \frac{g^2}{\Delta}, \quad (2.102)$$

where the overbar indicates that the energy is an *eigenenergy*, as distinguished from *bare energies* which do not take the coupling into account. Equation 2.102 is also known as the Lamb shift [73].

In the two excitation rung, we gain a $\sqrt{2}$ factor in coupling strength so the repulsion appears to be doubled. However, we must also take into account the repulsion due to $|2,0\rangle$. To first order, we can independently add these repulsions to find the eigenenergy for $|1,1\rangle$

$$\overline{E_{|1,1\rangle}} - E_{|1,1\rangle} = \frac{2g^2}{\Delta} - \frac{2g^2}{\Delta + \eta} \quad (2.103)$$

$$(2.104)$$

where the minus sign is due to the fact that $|2,0\rangle$ is below $|1,0\rangle$ and thus repels in the opposite direction compared to $|0,1\rangle$.

We now have all of the tools to compute the quantities of interest in our system. The

dispersive shift is the difference in resonator frequency conditional on the qubit state,

$$2\chi = \omega_r^{|1\rangle} - \omega_r^{|0\rangle} \quad (2.105)$$

$$= (\overline{E_{|1,1\rangle}} - \overline{E_{|1,0\rangle}}) - (\overline{E_{|0,1\rangle}} - \overline{E_{|0,0\rangle}}) \quad (2.106)$$

$$= \left(\omega_r - \frac{g^2}{\Delta} + \frac{2g^2}{\Delta} - \frac{2g^2}{\Delta + \eta} \right) - \left(\omega_r + \frac{g^2}{\Delta} \right) \quad (2.107)$$

$$= \frac{2g^2}{\Delta} - \frac{2g^2}{\Delta + \eta} \quad (2.108)$$

$$= \frac{-2g^2\eta}{\Delta^2(1 + \eta/\Delta)}. \quad (2.109)$$

In the limit of $\eta \rightarrow \infty$, we have $\chi \rightarrow -g^2/\Delta$ giving us our two level formula, while $\eta \rightarrow 0$ yields $\chi \rightarrow 0$, matching our intuition for two coupled oscillators.

We can also compute the AC Stark shift, which is the shift in qubit frequency when one resonator photon is added

$$\text{AC Stark shift} = (\overline{E_{|1,1\rangle}} - \overline{E_{|0,1\rangle}}) - (\overline{E_{|1,0\rangle}} - \overline{E_{|0,0\rangle}}) \quad (2.110)$$

$$= \left(\omega_q - \frac{g^2}{\Delta} - \frac{2g^2}{\Delta} + \frac{2g^2}{\Delta + \eta} \right) - \left(\omega_q - \frac{g^2}{\Delta} \right) \quad (2.111)$$

$$= \frac{2g^2}{\Delta} - \frac{2g^2}{\Delta + \eta} \quad (2.112)$$

$$= 2\chi. \quad (2.113)$$

As before, we find that the AC Stark shift is equivalent to the dispersive shift. Furthermore, since increasing the resonator photon number increases the coupling strength by \sqrt{n} , the AC Stark shift of the qubit should be linear with the photon number to first

order.

2.4.3 Qubit-Qubit Coupling

Using Eqn. 2.93 again, we can now replace both sets of ladder operators with Pauli operators to get the Hamiltonian for qubit-qubit coupling

$$H_g = \hbar g (\sigma_1^+ - \sigma_1^-) (\sigma_2^- - \sigma_2^+) \quad (2.114)$$

$$= \hbar g \sigma_{y1} \otimes \sigma_{y2} \quad (2.115)$$

Next, we move to the rotating frame, which we saw how to do for σ_y in Eqn. 2.64

$$H_g = -\hbar g (e^{i\omega_1 t} \sigma_1^+ - e^{-i\omega_1 t} \sigma_1^-) (e^{i\omega_2 t} \sigma_2^+ - e^{-i\omega_2 t} \sigma_2^-) \quad (2.116)$$

$$= -\hbar g (e^{i(\omega_1+\omega_2)t} \sigma_1^+ \sigma_2^+ + e^{-i(\omega_1+\omega_2)t} \sigma_1^- \sigma_2^- - e^{i(\omega_1-\omega_2)t} \sigma_1^+ \sigma_2^- - e^{i(\omega_2+\omega_1)t} \sigma_1^- \sigma_2^+) \quad (2.117)$$

Assuming ω_1 and ω_2 are relatively close, we again use RWA to ignore the high frequency $\omega_1 + \omega_2$ terms to obtain

$$H_g = \hbar g (e^{i(\omega_1-\omega_2)t} \sigma_1^+ \sigma_2^- + e^{i(\omega_2+\omega_1)t} \sigma_1^- \sigma_2^+) \quad (2.118)$$

and finally, if $\omega_1 = \omega_2$, we have

$$H_g = \hbar g (\sigma_1^+ \sigma_2^- + \sigma_1^- \sigma_2^+) \quad (2.119)$$

which is often written as

$$H_g = \hbar g (\sigma_x \otimes \sigma_x + \sigma_y \otimes \sigma_y). \quad (2.120)$$

To understand the effect of the Hamiltonian in Eqn. 2.120, we write it in matrix form with basis states ordered as $|00\rangle, |01\rangle, |10\rangle, |11\rangle$.

$$H_g/\hbar = g \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad (2.121)$$

We see that on-resonance coupling results in a Hamiltonian which swaps the $|01\rangle$ and $|10\rangle$ states. In other words, H_g swaps excitations between the two qubits. Note that within the $|01\rangle, |10\rangle$ subspace, the coupling Hamiltonian resembles to σ_x and is analogous to the driven single qubit Hamiltonian in Eqn. 2.68 with $\phi = 0$ and $\Omega/2 \rightarrow g$. Thus, without having to do any more work, we see that the frequency of swapping is $g/4\pi$.

2.5 Frequency Control

We will see in Chapter 8 that being able to tune the frequency of the transmon is crucial for executing the type of entangling gate used in the UCSB Xmon scheme. We saw in Eqn. 2.35 that the energy of the qubit transition in the transmon is determined by the parameters E_C and E_J . To dynamically change the transition energy, we must dynamically change either the capacitance or Josephson inductance of the transmon circuit. Fortunately, the Josephson inductance can be tuned by using a superconducting quantum interference device (SQUID).

A SQUID consists of two Josephson junctions in parallel inside a superconducting loop. For simplicity, suppose the two junctions have the same critical current and Josephson energy, E_J . Then, the energy in the two junctions is

$$U = -E_J(\cos \phi_1 + \cos \phi_2) \tag{2.122}$$

$$= -2E_J \cos \frac{\phi_1 + \phi_2}{2} \cos \frac{\phi_1 - \phi_2}{2} \tag{2.123}$$

where ϕ_1 and ϕ_2 are the phase differences across the two junctions. When the SQUID is threaded by an external flux Φ_{ext} , a screening current will flow to cancel this flux, because the integrated phase around a superconducting loop must be an integer multiple of the flux quantum, Φ_0 . The screening current creates a difference in the phases across the two junctions, which must obey

$$\phi_2 - \phi_1 = 2\pi n + \frac{2\pi\Phi_{\text{ext}}}{\Phi_0}. \tag{2.124}$$

where n is an integer. Substituting back into Eqn. 2.123, we have

$$U = -2E_J \left| \cos \frac{\pi \Phi_{\text{ext}}}{\Phi_0} \right| \cos \phi \quad (2.125)$$

$$(2.126)$$

where $\phi = (\phi_1 + \phi_2)/2$ is the effective phase difference across the SQUID. The absolute value is a result of the fact that if $|\Phi_{\text{ext}}| > \Phi_0/2$, the screening current will prefer to flow in the opposite direction, reversing the phase difference. Thus, we see that the energy of the SQUID is effectively the energy of a single junction with a tunable E_J that goes as the cosine of the external flux. As a function of the external flux, the transmon $|0\rangle \rightarrow |1\rangle$ frequency is

$$\hbar\omega_{10} = \sqrt{8E_C E_{J,\text{max}} \left| \cos \frac{\pi \Phi_{\text{ext}}}{\Phi_0} \right|} - E_C. \quad (2.127)$$

2.6 Typical Transmon Parameters

In this section we will enumerate the typical parameters used in the UCSB Xmons. Before we do so, we first go over some of the design considerations when choosing transmon parameters.

2.6.1 Frequency

The frequency of the transmon must be high enough to avoid thermally populating the qubit. The excited state population is given by the usual Boltzmann formula

$$P_{\text{thermal}} = \frac{e^{-\hbar\omega_{10}/k_bT}}{1 + e^{-\hbar\omega_{10}/k_bT}} \quad (2.128)$$

where we have ignored the presence of the higher states of the transmon. Typical measurements have found that despite a 10 mK environment in the dilution refrigerator, qubits tend to thermalize instead to 30 mK. To attain a population of 0.1%, the frequency should be at least 4 GHz. On the other end of the spectrum, the transmon frequency should lie comfortably within the available range of off-the-shelf microwave components, so transmons are typically operated in the 4-8 GHz range.

2.6.2 Control Coupling

When coupling the capacitive drive line and SQUID bias line to the qubit, we want as strong a coupling as possible without damping the qubit. For the capacitive drive line, the loaded quality factor of the qubit due to coupling to the drive is given by [68]

$$Q = \frac{C}{RC_d^2\omega} \quad (2.129)$$

where R is the impedance of the drive line and is always 50 Ohms, C_d is the coupling capacitance, and ω is the frequency of the qubit.

For the SQUID bias, we are primarily concerned with electrical noise in the SQUID dephasing the qubit ². We express noise as a single sided power spectral density S_V which

²Asymmetry in the SQUID's junctions can also cause damping, though this effect is usually small.

has units of V^2/Hz . The flux noise in the SQUID due to voltage noise on the bias line is

$$S_{\Phi}(\omega) = S_V(\omega) (M/R)^2 \quad (2.130)$$

where S_{Φ} is typically expressed in units of Φ_0^2/Hz and R is the impedance of the bias line. We can then convert this flux fluctuation to a frequency fluctuation of the qubit

$$S_{\omega_{10}} = \left(\frac{d\omega_{10}}{d\Phi} \right)^2 S_{\Phi}(\omega). \quad (2.131)$$

The conversion from $S_{\omega_{10}}$ to a dephasing time depends on the shape of the S_V spectrum. Assuming white noise so that $S_{\omega_{10}} = S_0$, the dephasing time is given by $T_2^* = 4/S_0$, yielding

$$\frac{1}{T_2^*} = \frac{1}{4} \left(\frac{d\omega_{10}}{d\Phi} \frac{M}{R} \right)^2 S_V \quad (2.132)$$

2.6.3 Readout Coupling

For readout, there are three important parameters: the coupling of the readout resonator to its drive, which is commonly denoted by κ ; the detuning of the readout resonator from the qubit, Δ ; and the coupling between the qubit and the resonator g_{res} . The resonator-drive coupling determines the readout speed, which we want as fast as possible to mitigate any T_1 decay during the readout. However, this coupling can also lead to damping of the qubit through the *Purcell effect*. The Purcell decay is given by

$$\Gamma_{\text{Purcell}} = \kappa \left(\frac{g_{\text{res}}}{\Delta} \right)^2. \quad (2.133)$$

This qubit decay rate can be suppressed by introducing a bandpass filter between the resonator and drive line [74].

The qubit-resonator coupling and detuning factor into both the Purcell decay rate and as we saw earlier, the dispersive shift of the resonator. To get the largest difference in resonator response between the $|0\rangle$ and $|1\rangle$ qubit states, the dispersive shift should match the linewidth of the resonator, and assuming the linewidth of the resonator is dominated by the drive coupling, we have

$$\chi \approx \kappa. \tag{2.134}$$

In principle, the resonator frequency may be placed either above or below the qubit. We will more carefully examine the consequences of this choice in Chapter 9.

2.6.4 Qubit-Qubit Coupling

The qubit-qubit coupling determines the speed at which entangling interactions can occur between qubits. However, in a fixed coupling system like our capacitively coupled transmons, the only way to turn off the interaction is to detune the qubits apart. Typically, this means that some qubits must be operated in a region of high flux sensitivity and increased dephasing. Therefore, the coupling is bounded by the maximum qubit-qubit detuning that we can tolerate in our system.

2.6.5 Summary

Table 2.1: Typical UCSB Xmon Parameters

Parameter	Common Symbol	Typical Value
Qubit capacitance	C	80 fF
Charging energy	E_C	240 MHz
Junction critical current (combined for SQUID)	I_C	40 nA
Josephson inductance	L_J	8 nH
Josephson energy	E_J	20 GHz
Impedance	Z	320 Ω
Charge zero point fluctuation	Q_{zpf}	2.5 e
Phase zero point fluctuation	Φ_{zpf}	0.06 ϕ_0
Qubit drive coupling capacitance	C_c	30 aF
Qubit bias mutual inductance	M	2 pH
Qubit frequency	f_{10}	5-7 GHz
Anharmonicity	η	240 MHz
Resonator-Qubit coupling	g_{res}	100 MHz
Resonator-Qubit detuning	Δ	1 GHz
Resonator decay rate	κ	$2\pi / 30$ ns
Qubit-Qubit coupling	g	15 MHz

Chapter 3

Superconducting Qubit Fabrication

In this chapter, we will briefly review the techniques used to fabricate our transmon qubits, then discuss more specifically how we fabricate 3 dimensional wiring for superconducting circuits.

3.1 Basic Superconducting Qubit Fabrication

As we discussed in Chapter 1, one of the appeals of superconducting qubits is that they can be fabricated using techniques developed for integrated semiconductor circuits. However, traditional IC techniques cannot be simply transferred to superconducting qubits, because superconducting qubits must exist in a low loss microwave environment in order for the qubits to maintain coherence. Typically, semiconductor ICs processes use amorphous oxides such as silicon dioxide as an insulator and dielectric, but these oxides have loss tangents of order 10^{-3} [75, 76]. At qubit frequencies of 4 – 6 GHz, prevalent use these

oxides would suppress the coherence times to only a few hundred nanoseconds, far too short for high fidelity operation. Hence, modern superconducting qubits are typically fabricated in a geometry where the electric energy of the capacitor is concentrated in either vacuum or a low loss intrinsic semiconductor substrate.

The geometry typically used in the UCSB devices is the co-planar waveguide (CPW) geometry [77], where a conductive center trace is surrounded by two ground planes. In this geometry, most of the electric field energy is stored in either the high quality crystalline substrate, or vacuum. However, even in this geometry we are not completely protected from the effects of lossy, amorphous materials. Wenner et. al. [78] showed through simulation that up to 1% of the electric energy resides in the interfaces between the metal film, substrate, and vacuum, and these interfaces tend to be dominated by native surface oxides and residues leftover from fabrication. Experiments have also found that superconducting microwave devices with larger areas and volumes have lower loss [79, 80, 81] which is consistent with a model where thin interfaces are the dominant source of loss in these devices. Thus, ensuring clean interfaces throughout the fabrication process is the critical challenge in fabricating superconducting qubits.

As devices become more complex, process development in superconducting qubits has also become focused on reliability and scalability. At times, scalability, reliability, and coherence are conflicting requirements, and a fabrication process will inevitably have to compromise between these requirements.

3.1.1 Thin Film Deposition

The first step in fabricating superconducting qubits is selecting a superconducting film and a semiconducting substrate. The traditional film used for superconducting devices is aluminum, because aluminum can be easily and controllably oxidized to form the insulator for a Josephson junction. In addition, superconducting resonators fabricated from aluminum on both sapphire [80] and silicon substrates have been demonstrated to have low loss. Other materials systems that have shown promise include titanium nitride (TiN) [82, 81] and niobium titanium nitride (NbTiN) [83]. The current standard at UCSB is to use aluminum on silicon, which has been demonstrated to have high coherence without requiring the use of a complex molecular beam epitaxy (MBE) system. In addition, unlike sapphire, silicon can be easily micromachined, which opens up possibilities for using the third dimension in routing circuit elements.

Having selected our film and substrate, we next turn to the deposition method. As we discussed earlier, the interface between the aluminum film and silicon substrate should be as clean as possible for low loss. Prior to deposition, we clean the silicon substrate by sonicating it in acetone and isopropanol, followed by dipping in heated piranha solution (a mixture of sulfuric acid and hydrogen peroxide) to clean off any organic residue, and finally by dipping in hydrofluoric acid to remove the native oxide of silicon. We then immediately transfer the substrate to our deposition chamber, a high vacuum electron beam evaporation system. E-beam evaporation is a physical vapor deposition technique in which a target of high purity material is bombarded with high energy electrons (typi-

cally 10 keV) in high vacuum (typically 10^{-7} torr in our system). Under these conditions, the target material melts then evaporates into a gaseous phase and coats the entire chamber including the substrate. We typically deposit 100 nm of aluminum at 1 nm/s.

3.1.2 Patterning

After depositing a thin film on our substrate, we pattern the metal film with the design of our desired circuit. As in a standard IC fabrication process, patterning is typically done using a technique called *photolithography*. First, we layout our design in a CAD program, and have the design etched into a chrome coated plate of quartz known as a photomask. In the cleanroom, we spin an even coating of organic material known as photoresist onto our metal film, sometimes preceded by a primer such as hexamethyl disulfide (HMDS) which promotes adhesion of the photoresist. Next, the resist is baked to drive off solvents in the resist to improve adhesion and prevent bubbling in subsequent processes. The resist is then exposed to ultraviolet light through the photomask. If the photoresist is *positive*, exposed regions of the resist are broken down and become susceptible to dissolution in developer. On the other hand, *negative* resist becomes more strongly bonded when exposed, with the unexposed regions susceptible to developer. Due to the higher resolution of positive resists and that fact that we need to etch small regions of the metal film, we typically use a common positive resist available in the UCSB cleanroom, SPR-955.

After exposure, the resist is baked again to stabilize the resist. The final step in

photolithography is to develop away the exposed regions of the positive resist. The most common developer for SPR-955 is AZ 300 MIF, a developer based on tetramethylammonium hydroxide (TMAH). However, aluminum is easily attacked by many developer solutions, and we found that aluminum on intrinsic silicon in particular can be etched by AZ 300 MIF in the time it takes to develop the photoresist. To prevent undesired etching, a substitute developer AZ DEV 1:1 can be used, which only slightly roughens the aluminum.

We note here that patterning may also be done prior to deposition, using a process called liftoff. After performing photolithography on a bare substrate, metal is deposited and will adhere to the substrate where the resist was developed away. The resist is then placed in a solvent solution which dissolves the remaining photoresist and peels away the metal on top of undeveloped resist. Liftoff is more compatible with e-beam written resists, which often do not have good etch resistance. Additionally, Josephson junctions are almost always fabricated using an e-beam written liftoff process, so the remainder of the transmon design can thus be patterned simultaneously in liftoff. However, by performing photolithography on bare substrate, we leave organic residues which can not be cleaned as aggressively for fear of damaging the resist pattern [84]. The result is that liftoff devices can be lossier unless special techniques are used.

3.1.3 Etching

After photolithography, we now have a resist mask with holes exposing regions of the metal that we want to remove. Etching of the metal can be done in two ways, either by exposure to a chemical solution (known as wet etching) or by bombardment with atoms, or ions from a reactive gas (known as dry etching). Dry etching of aluminum is typically performed using an inductively coupled plasma (ICP) tool, which strikes a plasma in a reactive gas via inductively coupled radio frequency waves. The plasma can then be biased towards the sample by applying a voltage on a set of capacitive plates, leading to an anisotropic (directional) etch. Aluminum can be ICP etched at a high rate using BCL_3Cl_2 , however remnants of the reactive gas will remain after completion of the etch which can lead to undesired etching of the film. Therefore, the chlorine etch is followed by a quench in a fluorine compound such as CF_4 or SF_6 , and then immediate rinsing in deionized water after removal from the ICP chamber. Introducing fluorine into the etch chamber has the side effect of etching the silicon substrate, which may be desirable since it reduces the amount of electric field energy in the aluminum-silicon interface. However, such a process must be carefully developed so that the silicon etch is reproducible and is not too deep as to interfere with subsequent fabrication steps. Following the etch, the photoresist mask is removed by sonicating in heated N-methyl-2-pyrrolidone then in acetone and isopropyl alcohol.

Wet etching of aluminum is also possible, and two wet etch processes are used in the UCSB cleanroom. The standard process uses Transene Aluminum Etchant Type

A [85], which contains a mixture of nitric, phosphoric, and acetic acids and will etch 100 nm of aluminum in a few minutes. As noted before, aluminum on silicon can also be etched by developing the photoresist pattern in a TMAH based developer, combining etching and patterning into a single step. Somewhat surprisingly, the developer etching method has been shown to have reasonably low loss and has successfully been used in qubit devices. Wet etching has a few advantages over dry etching. First, wet etching is less damaging to photoresist and is less likely to leave photoresist residue upon stripping. Second, typical aluminum wet etchants will not etch silicon, which makes it easier to hit target frequencies and impedances. However, wet etching is an isotropic process which is not suitable for small features such as skinny wires for control lines. Because wet etching and dry etching are useful in different contexts, they can both be used in the fabrication process for different parts of the device, with each etch requiring a different patterning step.

3.1.4 Junction Fabrication

After patterning and etching the wiring and capacitor, the final step is to fabricate the Josephson junction. Because the junction consists of a thin dielectric bounded by two conducting metals, it has a high capacitance per area, and we want as small an area as possible to avoid having a significant amount of the transmon's electrostatic energy located in the amorphous junction dielectric. Hence, Josephson junctions for superconducting qubits are typically made using electron beam lithography at the 100 nm

scale using a technique known as the *Dolan bridge* [86]. First, two layers of positive e-beam resist, methylacrylic acid (MAA) and poly(methyl methacrylate) (PMMA) are spun on the sample and baked. Next, the resist is exposed in a dosage pattern such that in certain regions, the MAA is exposed enough to be developed away but the PMMA is not. Upon development in a mixture of methyl isobutyl ketone (MIBK) and IPA, we are left with a suspended “bridge” of PMMA. The Josephson junction is then deposited in an e-beam evaporator in four steps. First, the base aluminum film is milled using high energy argon ions to remove the native oxide to allow for galvanic contact. Next, the aluminum is deposited at an angle perpendicular to the surface, with the PMMA bridge masking a region which creates a break between two metal leads. Then, the deposited aluminum is oxidized in-situ by introducing oxygen into the chamber at a few mTorr for about an hour. Finally, a second layer of aluminum is deposited at an angle that allows for aluminum to be deposited underneath the PMMA bridge, connecting the two metal leads through a Josephson junction. The excess aluminum is then lifted off in NMP, completing the device.

One issue with the “standard” fabrication process described above is damaging of the substrate during the ion mill step. Reference [87] showed that the loss due to argon milling on a silicon substrate is significant, and the concentration of electric fields around the skinny Josephson junction wires exacerbates the problem. One solution is to make galvanic contact *after* the fabrication of Josephson junction. The technique developed in Ref. [87] does so by adding a patterning, ion milling, then lifting off a patch of aluminum

metal on top of the Josephson junction and connecting leads, in effect “bandaging” the two regions of aluminum together. Using this technique introduces an additional lithography and deposition step, and care must be taken to avoid passing too much ion current through the Josephson junction during the ion mill. However, use of the bandaging technique has proven to be crucial to making high coherence aluminum on silicon devices.

3.2 Fabricating Superconducting Airbridges

So far, we have described a basic fabrication procedure which can be used to make planar devices. However, using a purely planar geometry limits the complexity of the device we can make since we cannot cross wires over each other in two dimensions. Additionally, the co-planar waveguide geometry used in our devices ideally supports a symmetric [88] mode where the ground planes on either side of the center trace are held to the same voltage. However, asymmetries and discontinuities in the microwave circuitry can lead to the excitation of parasitic slotline modes [89]. These modes can couple to elements of the circuit such as qubits, and they represent a source of radiation loss and decoherence [90, 91]. In order to suppress these modes and to allow for more complex wiring, we require 3 dimensional wiring such as crossovers or superconducting vias.

In this section, we will illustrate the fabrication and use of superconducting airbridge crossovers. To motivate our use of airbridges, we observe that in past work with superconducting circuits, connections between the different ground planes have been typically




Figure 3.1: (Inset) (a) The slotline mode of a CPW is modelled by removing the center trace. A crossover wire is used to tie together the two planes so that the slotline mode may not propagate. (b) Equivalent transmission line model for the slotline mode shunted by crossovers with an inductance L_S . We obtained a capacitance and inductance per length of $C=140$ fF/mm and $L=450$ pH/m from numerical simulation of a $20\ \mu\text{m}$ gap slotline, giving an impedance of $56\ \Omega$ which is matched by the load. The wirebond and airbridge have an L_S of 1 nH and 10 pH respectively, and are placed at intervals of length ℓ . Main panel: SPICE simulations for 1 mm of the transmission line model, showing that the attenuation due to a single airbridge is more than 20 dB greater than for a wirebond. Ten airbridges per mm can be simply fabricated and gives an attenuation of -150 dB (not shown)

been made using wirebonds. However, with a wire diameter of $25\ \mu\text{m}$ and a typical length of 1 mm, wirebonds have an inductance of order 1 nH and an impedance of $40\ \Omega$ at 6 GHz, making them an ineffective shunt [92]. In comparison, airbridges have 100 times less inductance due to their small size. In order to understand the effect of the crossover impedance on slotline attenuation, we studied a simple transmission line model [93] for the slotline with evenly spaced inductive shunts to ground as shown in Fig. 3.1(a). We simulated in SPICE 1 mm of a transmission line with a terminated load, and varied the number of inductive shunts. As seen in Fig. 3.1, the attenuation per millimeter of the slotline propagation for a single airbridge is two orders of magnitude greater than for one or two wirebonds. This result can be easily understood if we consider only the inductances in the model, which is valid below the cutoff frequency [93]. The inductance of 0.5 mm of the slotline is 0.23 nH, which is smaller than the wirebond inductance but much larger than the inductance of an airbridge. Therefore, in the case of wirebonds, signal will continue propagating down the line rather than flow to ground. Furthermore, while increasing the wirebond density can be difficult and unreliable, increasing the airbridge density can be done by simply copying the structure in the design file. With 10 airbridges per mm, we simulated the attenuation to be -150 dB, implying that the slotline mode does not exist. Recent measurements on qubits built in a CPW architecture also suggest that crossovers are necessary to eliminate slotline modes as a source of loss and crosstalk.

The fabrication process we used for the airbridges follows from earlier work done on




Figure 3.2: (a-f) Fabrication process for superconducting airbridges, with substrate shown in blue, resist in dark red and aluminum in gray. In order, the fabrication steps are: (a) fabrication of CPW base layer, (b) patterning and reflow of photoresist, (c) deposition of aluminum, (d) definition of the bridge using lithography, (e) wet etching of excess aluminum and, (f) release of airbridge. (g) SEM image of airbridges connecting the ground planes of a CPW line and (h) SEM image of airbridges linking together two CPW lines.

kinetic inductance detectors [94], with modifications to adapt the process for an aluminum base layer. We show the process flow and resulting structures in Fig. 3.2. First, we formed the scaffold for the bridge from a $3\ \mu\text{m}$ thick positive photoresist (Megaposit SPR-220-3). The height of the bridge is set by the resist thickness, and photolithography determines the placement and span of the bridge. Throughout the process we used a developer (AZ Dev 1:1) designed to minimize aluminum etching. Away from the bridge area, we did not expose the resist so that it remained as a protective layer and etch stop. Next, we reflowed the resist at 140°C for 3 minutes to form an arch for mechanical stability. We then deposited 300 nm of aluminum in a high vacuum electron beam evaporator to form the bridge layer. Prior to the deposition, we used an in-situ argon ion mill calibrated to remove the native oxide of the base aluminum in order for the bridges to make good electrical contact [79]. The ion mill was operated for 3.5 minutes in 1×10^{-4} mbar of argon, with beam voltage 400V, beam current 21 mA, and beam width 3.2". Using a second layer of patterned $3\ \mu\text{m}$ resist as a mask, we then wet etched (Transene Aluminum Etchant Type A at 30°C) the excess deposited aluminum that is not used to form the bridge. We terminated the etch by visual inspection. When the top layer of aluminum was etched away, the wafer went through a clear change in reflectance from aluminum to the underlying resist layer; the typical etch time was 5 minutes. We continued immersing the wafer in the etchant for 5 seconds after this transition, then immersed the wafer in water for 3 minutes. The termination of the etch is a critical step because the regions around the pads of the bridge are not protected by photoresist during the etch, and can

potentially be etched through, breaking the ground plane. Finally, we stripped both layers of resist in an 80°C heated bath of NMP to release the airbridges.

In developing the process, we initially found a large amount of residual resist remained from the scaffolding layer after stripping in solvents. This residual resist decreased the reliability of our bridges by loading and deforming the bridge arches, and would have contributed a large amount of loss to our circuit. We deduced that the residue consisted of resist cross-linked by ion implantation from the argon ion mill step, a well known problem in semiconductor processing.[95] We were able to mitigate the problem by stripping the resist layers in a downstream oxygen plasma at 150°C for 5 minutes prior to stripping in a solvent bath. The low temperature oxygen plasma acts to burn off the damaged layer of resist.

With this additional cleaning step, we have reliably fabricated bridges over a range of spans from 2 μm to 50 μm . The main sources of bridge failure are factors other than their structural stability such as lithographic errors, and the failure rate is less than 0.1%. We have also tested the bridges in a variety of postprocessing steps, including wafer dicing and fabricating aluminum junctions with a bilayer electron beam resist process; bridges spanning up to 40 μm have been found to survive these steps reliably. We note here that the airbridges generally do not survive sonication.

Having developed a stable airbridge process, we next tested the electrical properties of the airbridges at DC and microwave frequencies. We found that the critical current of the airbridges was of order milliamps and that the microwave loss due to each airbridge

was of order 10^{-8} . For more details on these measurements, see Appendix A.

3.3 Contributions

The second half of this chapter is an adaptation of Ref. [96], for which I fabricated the devices, acquired the data, and wrote the text. Anthony Megrant and Julian Kelly also contributed to the project.

Chapter 4

Experimental Infrastructure for Superconducting Qubits

In this chapter, we describe the infrastructure required to run a multiqubit system in three parts: hardware inside the cryostat, hardware outside the cryostat, and software.

4.1 Cryogenic

4.1.1 Cryostat

The most important tool in the lab is the cryostat used to cool samples to superconducting temperatures. Superconducting Xmon devices are cooled down in a dilution refrigerator (DR) [97] to a nominal temperature of 15 mK. In brief, a dilution refrigerator works as follows:

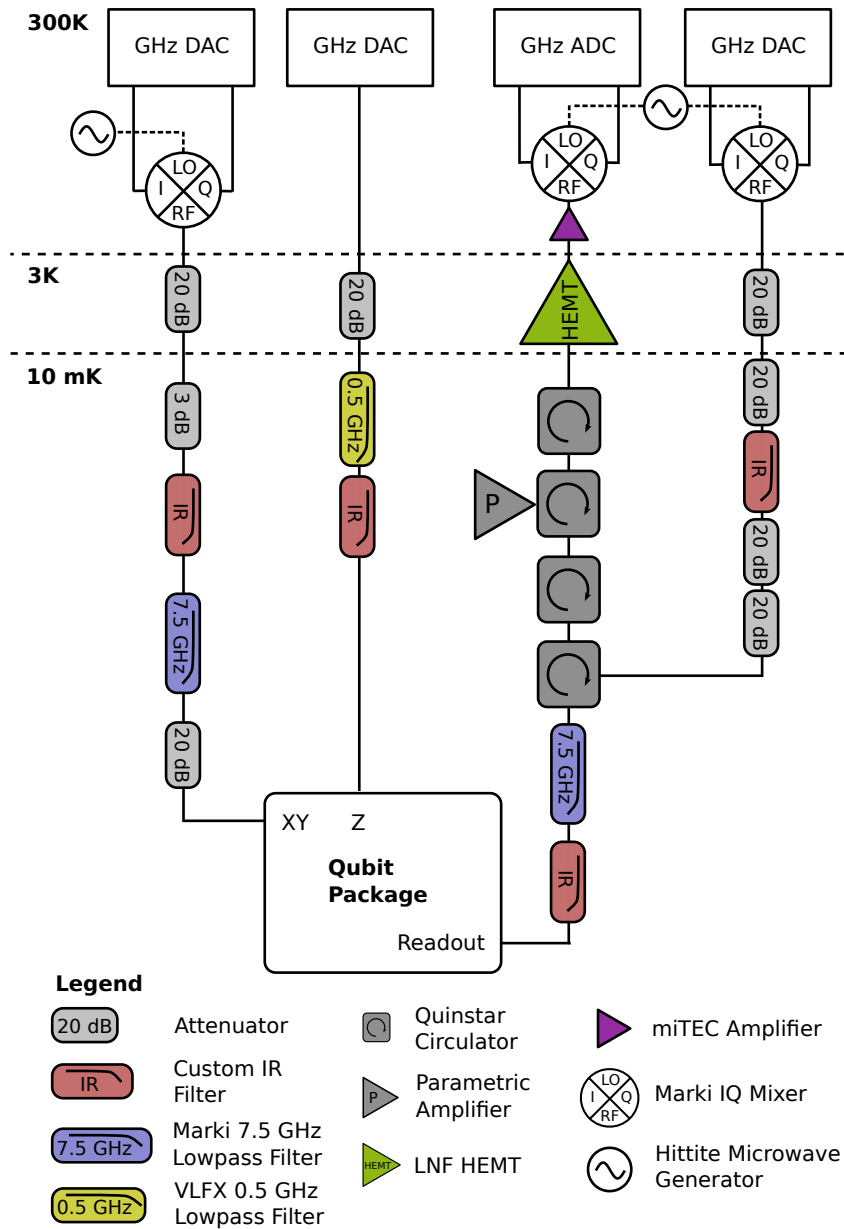


Figure 4.1: Overview of qubit operation hardware, showing the necessary equipment to operate a single qubit. At room temperature, custom DAC boards generate control waveforms, some of which are upconverted to the qubit frequency via an IQ mixer and microwave source. The waveforms are sent down to the fridge through heavy filtering and attenuation. On the righthand side, readout waveforms are generated in a similar fashion and reflected off the sample. The output signal is amplified by a reflective, impedance matched parametric amplifier, a HEMT, and room temperature amplifiers, and finally detected by a custom ADC. Circulators protect the qubit from the reflected signal and noise coming from the amplifiers. Not shown: bias and pump for the paramp, additional filters and attenuators at room temperature.

- The refrigerator is first cooled to 3K by a pulse tube cryocooler.
- We then begin to circulate and pump on our working fluid, a mixture of two isotopes of helium, He-3 and He-4. This circulation results in some evaporative cooling down to below 1K.
- At 870 mK, the mixture separates into two phases, a He-3 concentrated phase which is nearly 100% He-3, and a dilute phase, which contains at least 6.6% He-3. In the concentrated phase, the He-3 behaves as a Fermi liquid, but in the dilute phase, the He-3 can be thought as behaving like a Fermi gas in an inert background of He-4 superfluid. He-3 in the dilute phase has higher enthalpy, and moving He-3 from the concentrated to the dilute phase results in cooling, much like evaporative cooling.
- Cooling occurs in the *mixing chamber*, where the lighter concentrated phase sits on top of the heavier dilute phase. A tube connects the dilute phase to a chamber called the *still* at 0.7K, where we pump on the mixture to vaporize and pull He-3 from the mixing chamber.
- The pumped He-3 is cleaned through a LN2 cold trap, then cooled and recondensed by the pulse tube and by exchanging heat with He-3 traveling from the mixing chamber to the still. Finally, it is introduced back into the mixing chamber, completing the cycle.




Figure 4.2: Aluminum sample mount with wirebonded sample. Signal lines are bonded to copper PCB inserts and the the sample ground is bonded to the mount itself. Picture credit: Michael Fang.

4.1.2 Sample Mounting

After a sample is fabricated, it is verified by visual inspection for fidelity of the control lines. We also probe the resistance of test structures to ensure that the resistances of the Josephson junctions are the expected values. After passing inspection, the samples are mounted in a connectorized package so that it can be wired for control and measurement. Over the years, we have empirically found that it is crucial for the package to have a superconducting ground. Otherwise, the low frequency pulses we use for flux control suffer from long time constant settling behavior which makes high fidelity Z-control nearly impossible. Thus, we can rule out traditional printed circuit boards (PCBs) which use copper traces and ground planes. The standard solution at UCSB has been to use a machined aluminum box which forms a superconducting ground and has been shown to have much shorter time constants. The box has tunnels milled out at even intervals around the perimeter, through which we insert strips of PCB dielectric coated with copper and soldered at one end with an SMA connector. The free end of the copper coated PCB

is wired bonded with aluminum wire to the signal lines of the sample, connecting the sample to the outside world. Additionally, the ground plane of the sample is densely bonded to the aluminum box to prevent any resonant modes from forming in the gap between the chip and the box. Finally, the box itself is placed within an enclosure made of mu-metal to shield the sample from magnetic fields which can induce vortex loss.

4.1.3 Wiring

From the package, the qubit is connected to room temperature controls through a series of co-axial cables. Cabling occurs in three main stages: copper-nickel cables from room temperature to 3K, niobium-titanium cables 3K to 20 mK, and a variety of cables at the 20 mK stage to the sample. Cables connecting stages at different temperatures are chosen to have low thermal conductivity. Niobium titanium is a superconductor which does not conduct heat well, while copper-nickel also has poor thermal conductivity due to being an alloy. Cables are further thermalized at intermediate stages using copper clamps.

Shielding the qubit from thermal electronic noise is an important consideration when designing the wiring system. Well above the temperature corresponding the qubit transition (6 GHz corresponds to roughly 300 mK), the thermal noise follows the Johnson-Nyquist formula

$$S_{V,\text{Thermal}} = 2k_bTR, \quad (4.1)$$

where S_V is the double-sided power spectral density with units volts squared per hertz.

Since the power is linear in temperature, to eliminate the effect of room temperature noise, we place a 20 dB (factor of 100) attenuator at 3K. The Johnson noise in our line is now dominated by 3K noise, which is roughly $8 \text{ pV}^2/\sqrt{\text{Hz}}$ for a 50Ω resistor.

At the 10 mK mixing chamber stage, we more carefully consider the effects of noise on the qubits. For the XY microwave drive line, the primary concern is voltage noise driving the $|0\rangle \rightarrow |1\rangle$ transition and populating our qubit ¹. Following a similar argument to Fermi's Golden Rule, the upwards transition rate due to nelectrical oise is [98]

$$\Gamma_{\uparrow} = g^2 S_V(-\omega_{10}) \quad (4.2)$$

where g is a conversion from voltage to Rabi frequency, and $S_V(-\omega_{10})$ is the double sided power spectral density of voltage noise at the qubit frequency. From Eqn. 2.54, the equation for Rabi frequency Ω is

$$\Omega = \frac{C_d V_0 Q_{\text{zpf}}}{\hbar(C + C_d)} \quad (4.3)$$

$$= g V_0 \quad (4.4)$$

and plugging in $C_d = 30 \text{ aF}$, $C = 80 \text{ fF}$ and $Q_{\text{zpf}} = 2.5 e$, we find that $g = 2\pi \times 230 \text{ GHz/V}$. To get an upward transition rate corresponding to one jump every $500 \mu\text{s}$, we require $S_V(-\omega_{10}) = 1 \times 10^{-3} \text{ nV}^2/\text{Hz}$, so we need an additional 40 dB of attenuation at mix plate.

For the Z frequency bias line, we are primarily concerned with dephasing. Equation

¹Electrical noise can also damp the qubit, but our qubit damping is primarily dominated by other mechanisms.

2.132 tells us how to compute the limit on T_2^* given a voltage noise level ². Plugging in our typical bias mutual of 2 pH, a standard transmon flux sensitivity of $2\pi \times 10$ GHz/ Φ_0 , and 3K noise, we find a limit on T_2^* of over 300 μ s, significantly above the current dephasing times of transmon. Thus, we do not need any additional attenuation at the mix plate for our frequency control lines.

4.1.4 Readout Chain

The input to the readout chain is a microwave line similar to the qubit XY drive line, but more heavily attenuated. The extra attenuation is necessary because the resonator is more strongly coupled to its drive line, and thermal fluctuations in the readout resonators can cause dephasing in the qubit. Signal coming down the input line passes through a circulator and onto the qubit chip, where it scatters off of the bandpass filter and readout resonator. The reflected signal passes through the same circulator but now onto the output readout line. After passing through a second circulator for isolation, the output signal goes through three stages of amplification: an impedance matched parametric amplifier (IMPA) [99] at the mix plate, a high electron mobility transistor (HEMT) at 3K, and room temperature amplifiers, before reaching the detection electronics.




Figure 4.3: Electronics to operate a qubit. Custom GHz DACs are used as arbitrary waveform generators (AWGs). For microwave XY control, the two channels of the DAC are combined and upconverted to the qubit frequency using an IQ mixer. Flux bias Z control requires only a single channel of the DAC without upconversion. More information can be found at <http://web.physics.ucsb.edu/martinisgroup/electronics.shtml>.

4.2 Room Temperature Electronics

Control waveforms for the Xmons comes in two flavors: shaped pulses of high frequency (4-7 GHz) microwaves to drive the qubit transitions and readout resonators, and “low frequency” bandwidth (100 MHz) waveforms to control the frequency of the Xmon. Both types of control are achieved using a custom arbitrary waveform generator (AWG), which we will refer to as the GHz DAC. The GHz DAC has the following features:

- The DAC itself is a 14 bit, 1 gigasample/second DAC capable of synthesizing waveforms up to 500 MHz.
- Each GHz DAC has two DACs with differential outputs. Each differential output is sent to a differential amplifier, which zeros out the common mode voltage and provides some gain. The diff-amp has two differential outputs, one of which is used to drive the desired signal, while the other acts as a monitor for the output which can hook up to an oscilloscope.
- The FPGA can play 15 μ s of memory, but additionally has programmable logic which allow us to loop over pieces of memory and play longer sequences.
- Waveforms are uploaded over ethernet at 100 MBit/s.
- Multiple GHz DAC boards are synchronized by running them all on the same 10 MHz reference clock, and daisy chaining the boards with ethernet connections to trigger waveform playback.

²But note that Equation 2.132 uses the single sided spectrum, whereas here we have used the double sided spectrum.

4.2.1 High Frequency Control

To generate pulsed microwaves for XY control, we connect the two outputs of a GHz DAC to an IQ mixer. The IQ mixer takes two inputs, a local oscillator (LO), and an in-phase (I) and quadrature (Q) input. We drive the local oscillator with a continuous wave sinusoidal tone at a carrier frequency ω_c . The output of the mixer is then

$$f(t) = I(t) \cos \omega_c t + Q(t) \sin \omega_c t \quad (4.5)$$

$$= \mathbb{R} [(I(t) - iQ(t))e^{i\omega_c t}] \quad (4.6)$$

In a *homodyne* mixing scheme, we would set ω_c to be equal to the qubit transition frequency. Then, turning on I and Q would correspond to rotating about the X and Y axis of the Bloch sphere. While this scheme is simple, it is difficult to scale because we would require a unique carrier for each qubit in the system.

Instead, we use a *heterodyne* mixing scheme, where the carrier frequency is different from the qubit frequency. Suppose ω_c is different from the qubit transition ω_{10} by an amount $\omega_{\text{sb}} = \omega_{10} - \omega_c$, the sideband frequency. We ultimately want the signal $X(t) \cos \omega_{10} t + Y(t) \sin \omega_{10} t$. Then, comparing with Eqn. 4.6, we have ³

$$(X(t) - iY(t))e^{i\omega_{\text{sb}} t} = I(t) - iQ(t), \quad (4.7)$$

³In practice, the sign of the sideband frequency might be reversed depending on the mixer convention.

from which we can get expressions for $I(t)$ and $Q(t)$

$$I(t) = X(t) \cos \omega_{\text{sb}} t + Y(t) \sin \omega_{\text{sb}} t \quad (4.8)$$

$$Q(t) = Y(t) \cos \omega_{\text{sb}} t - X(t) \sin \omega_{\text{sb}} t. \quad (4.9)$$

$$(4.10)$$

We have shown that even if the carrier frequency is different from the qubit frequency, we can still output a signal at the qubit frequency with full in-phase and quadrature control. This scheme is also called *single sideband mixing*, because we have multiplied the carrier signal by another sinusoidal tone and output a tone with the sum of the two frequencies but not their difference. The difference frequency corresponds to the *image sideband* and is undesirable. By using single sideband mixing, we can mix a single carrier to multiple frequencies to control multiple qubits, as long as the sideband frequencies are within the bandwidth of the GHz DAC.

While Eqn. 4.6 and Eqn. 4.6 describe the behavior of an ideal mixer, in practice we have to carefully calibrate for two effects. First, even if I and Q are zero, we may still get signal coming out of the mixer which is known as *carrier bleedthrough*. Second, if the I and Q channels are not perfectly balanced in amplitude or if the phase difference between the two inputs is not exactly 90° , then Eqn. 4.10 will not hold and we will have some image sideband in our signal. To calibrate for these two errors, we connect the output of our GHz DAC and mixer system to a spectrum analyzer. We first measure the signal power at the carrier frequency, and adjust the DC levels of the I and Q inputs until the

carrier bleedthrough is minimized. Next, we pick a sideband frequency and adjust the relative amplitude and phases of the I and Q inputs until the image sideband signal is minimized. We repeat this procedure for a range of carrier and sideband frequencies, and use the results to adjust our I and Q waveforms prior to sending them to the DACs.

4.2.2 Readout

Signal generation for readout is identical to qubit XY control. After scattering off the readout chain, the output signal from the amplifiers is first sent through the RF port of a mixer with the same LO as the input signal. This usage of the mixer is reversed compared to the signal generation case, and produces downconverted signals on the I and Q ports at the sideband frequency. The I and Q signals are then sent to a custom built GHz ADC, where they are digitized, and integrated after downconverting from the sideband frequency to DC in the ADC. The result is that for each pulse sent down from the readout DAC, we obtain a single complex value $I + iQ$ in two dimensional phase space which represents the scattering of the pulse off of our readout chain.

4.2.3 Low Frequency Flux Control

For low frequency Z control, we simply use the output of the GHz DACs without up- or down-conversion.

4.3 Software

4.3.1 LabRAD

As the complexity of the hardware grows with the number of qubits, maintaining sensible, flexible, and well documented software is crucial. The backbone of our software is a system called LabRAD (where RAD stands for Rapid Application Development), developed by Markus Ansmann and Matthew Neeley. LabRAD is essentially a platform for distributing and managing control software. The central piece of software is the LabRAD manager, which dispatches commands from client modules run by the experimenter, and server modules which control the hardware. Each server is typically in charge of one piece of hardware, but servers may utilize other servers - for example, one server may facilitate communication over ethernet while another uses the ethernet server to send packets to the GHz DACs. Client modules then have access to specific server functions called *settings*, which each server must register with the LabRAD manager and forms the server's outward facing interface. Importantly, servers and clients can be written in any programming language as long as an Application Programming Interface (API) exists for that language. The APIs convert language specific data into binary packets as laid out by the LabRAD protocol, and these packets are then used to communicate between different modules via the manager. Furthermore, servers and clients may exist on different machines on the same network, and these machines may be running different operating systems. By using LabRAD, the various stacks of our experiment are abstracted away

from each other, so that modifications to one aspect of the experiment generally do not require modifications to the entire code base.

In addition to the control platform, two systems are integrated into LabRAD which are crucial for running experiments: the *registry* and the *data vault*. The registry is a central directory which contains all of the relevant configuration information for the system, and is typically accessed in Python using dictionary calls. Configuration information might include low level data such as the available DACs on the system, or high level calibrated parameters such as the mapping between qubit frequency and voltage for a certain qubit. Likewise, the data vault is a central directory which permanently stores all of the data obtained in any experiment. Each data vault file consists of a 2D array of values sorted into independent and dependent columns, as well as a snapshot of the relevant registry configuration which can be used for data analysis or reconstruction of the experiment.

4.3.2 Hardware Servers

With these concepts in mind, we now take a brief survey of the software stack used to run our superconducting qubit system. At the lowest level are the servers used to run the hardware. The primary modes of communication are ethernet for the GHz DACs and GPIB for most other instruments such as microwave sources, oscilloscopes, and network analyzers. Serial communication using RS-232 is also occasionally used for instruments connected over USB. Each mode of communication has an associated server, and servers for specific instruments make calls to these communication servers. Certain

servers also use additional layers of abstraction to facilitate code reuse. For example, GPIB communication to oscilloscopes is relatively standardized, and rather than writing a server for every type of oscilloscope, we write a general oscilloscope server which can execute common read and write commands. Commands specific to a certain oscilloscope are then kept in device wrappers which can be applied upon identification of a connected instrument.

4.3.3 Sequencer

Moving up the stack, we next have a server called the *sequencer*. A client makes a request to the sequencer which specifies what waveforms should be run on each channel for each qubit. The sequencer first determines (or is given) the mapping between the qubit channels and specific hardware devices. The sequencer then sends the waveforms to a correction server, where the waveforms are corrected for hardware specific calibrations such as IQ balancing. Finally, the waveforms are sent onto the DACs, and other relevant instruments such as microwave sources are set to the desired configurations. Upon completion of the waveform sequence, the ADC's IQ results are sent back to the sequencer then onto the client in a client-specified data format. Notably, as with any LabRAD server, the sequencer allows multiple clients to send sequence requests to it at the same timeframe. Thus, experiments from multiple users can be interleaved without any extra work.

4.3.4 Gates and Experiments

Continuing our journey up the software stack, we now move onto client code, which is entirely written in Python. To specify the waveforms to send to the sequencer, we use objects internally called *Gates*. A Gate object contains the machinery to generate waveforms for the channels of a qubit given a few input parameters. A Gate object often corresponds to a quantum gate. For example, a Pi gate corresponds to either an X or Y gate depending on what phase the user specifies. We construct the Pi gate with a reference to a specific qubit, and at runtime, the gate object looks up the qubit specific π pulse parameters in the registry and constructs a waveform. Multiple gates acting on any number of qubits are then combined into a *gate sequence*, which can also contain gate-like objects which correspond to FPGA jump table commands, such as looping over a set of gates a number of times.

The gate sequences are further encapsulated in objects called *experiment classes*. Experiment classes contain one commonly used gate sequence, as well as hooks to parameterize the gate sequence in a useful way. For example, a Rabi experiment class generates a simple pi pulse then measure sequence, but has parameters for controlling the amplitude and length of the pi pulse. In addition, the experiment class takes care of loading the device configuration from the registry, as well as processing and labelling the data received from the sequencer and saving it to the datavault.

4.3.5 Automated Calibrations

At the top of the stack is a system for calling data acquisition functions to calibrate a qubit from scratch. A data acquisition function and its associated analysis function are organized into a *node*. Multiple nodes are then organized into a *calibration tree*. Starting at the root of the tree, the calibration system acquires data and determines whether the data is valid. If it is, the calibration system updates the registry based on the acquired data, and continues to traverse the tree. If the data is not valid, the automated system attempts to recover the system by traversing the tree backwards and fixing earlier calibrations if possible, or throws an informative exception to the user. This system is also used to check if calibration parameters have drifted over time, and to recalibrate the system if necessary.

Chapter 5

Single Qubit Calibration

In this chapter, we describe how to calibrate the control of a single superconducting transmon qubit. We start with the task of finding the frequencies of the qubit and its associated readout resonator. Next, we describe how to calibrate microwave control to achieve π and $\pi/2$ rotations of the qubit state, as well as fine tune the readout of the qubit. Finally, we describe how to calibrate the response of the qubit frequency to flux bias.

5.1 Introduction

In this chapter we will assume that we are operating a single qubit in isolation. While our device actually consists of an array of coupled qubits, we can effectively isolate one qubit by tuning its neighbors to near zero frequency. We have three inputs to the qubit that must be calibrated:

1. Z control: static or dynamic voltages on our DAC which tune the frequency of the qubit by flux biasing the SQUID. When we tune the frequency of the qubit, we also slightly tune the frequency of the resonator due to the Lamb shift.
2. XY control: pulsed microwaves produced by mixing two DAC outputs with a programmable local oscillator. When on resonance with the qubit, XY pulses will drive Rabi oscillations.
3. Readout: pulsed microwaves on a different local oscillator coupled to the readout resonator.

To characterize the system, we have a single output: the integrated in-phase (I) and quadrature (Q) amplitudes of the signal coming out of the readout chain at the readout drive frequency, as measured by our ADC. In other words, we measure the complex scattering of microwaves off of the readout resonator and amplifiers.

We have the following end goals for our calibration procedure:

1. Calibrated curves for the readout resonator frequency and qubit frequency vs flux bias.
2. Accurate discrimination between the $|0\rangle$ and $|1\rangle$ states.
3. High fidelity π and $\pi/2$ pulses.

For the remainder of this chapter, we will generally refer to control voltages in units of "DAC amplitudes", where ± 1 DAC amp is the full range of our DAC. Similarly, we

will refer to measured amplitudes in units of "ADC amplitudes", where one ADC amp is one bit on the ADC. While we can calibrate the correspondence between DAC and ADC amplitudes in terms of physical units, in practice it is not necessary to do so other than to estimate control couplings, which we will not address in this chapter. Additionally, for each configuration of the controls, we repeat the experiment a few hundred times and unless otherwise noted, the plotted values will be the average signal from all shots of the experiment. Since the measured signals are complex, we will refer to the measured signals either in terms of the average I (in phase) and Q (quadrature) amplitudes, or in terms of the total magnitude $|I + iQ|$ and phase $\angle(I + iQ)$ of the averaged signal. Finally, in the pulse sequence plots below, we will generally show microwave pulses prior to up-conversion to the qubit or resonator frequency, to emphasize the envelope of the pulse.

5.2 Resonator Spectroscopy

To begin, we know from our circuit design the expected frequency band where the resonator for our qubit lies. We would like to verify the frequency of the resonator and confirm that it is coupled to the qubit as expected. Due to their relatively strong coupling, if we tune the qubit's frequency by flux biasing the qubit's SQUID, we should also observe shifts in the resonator's frequency. We observe this effect by performing spectroscopy on the readout resonator: we pulse the readout chain with microwaves, while varying the frequency of the microwaves and the flux bias applied to our qubit. The

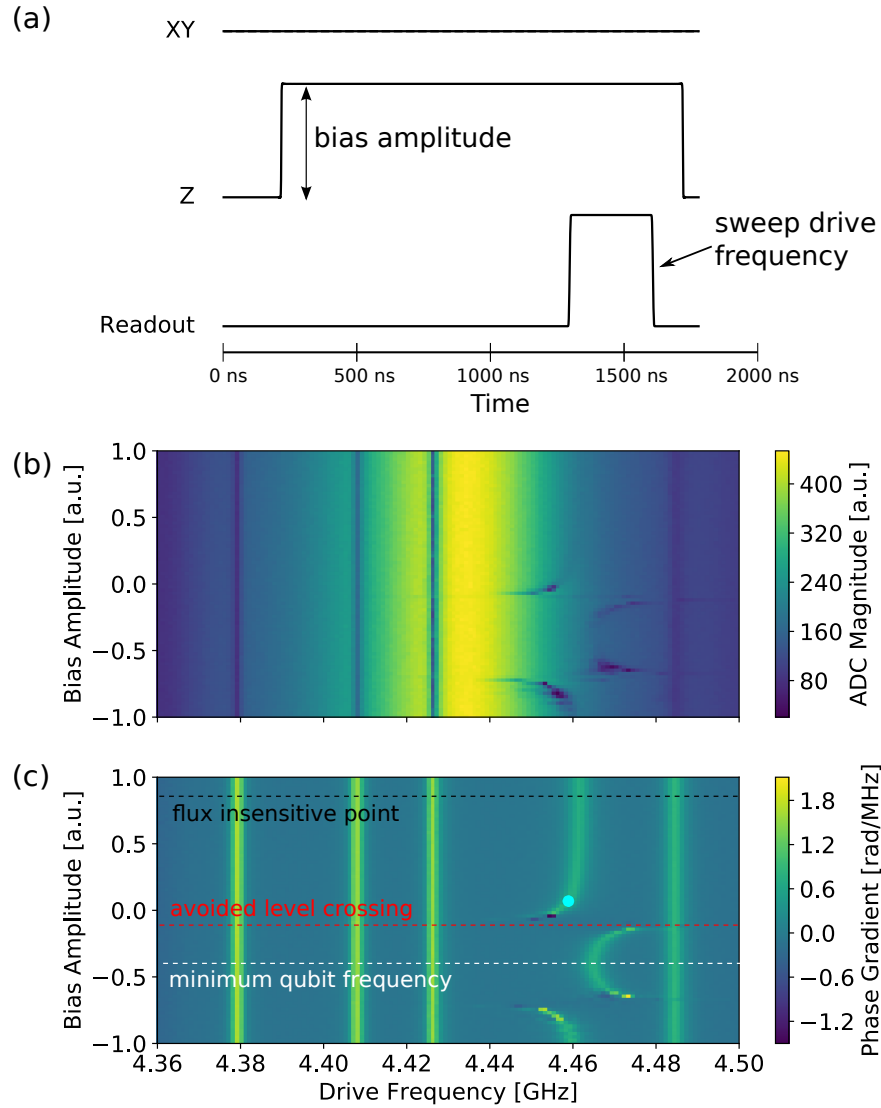


Figure 5.1: Resonator spectroscopy versus flux bias. (a) Pulse sequence. We dynamically bias the qubit and microwave drive the readout line. The bias is padded in front of the microwave pulse to allow the bias to settle to its intended value, and slightly padded in the back to account for possible timing differences between the flux bias and microwave channels. We sweep the bias amplitude and the frequency of the readout microwave pulse. (b) Measured total magnitude from the ADC. Distinct dips in the transmission can be seen for five resonators, with one resonator shifting with the flux bias. (c) Gradient of phase with respect to drive frequency, more clearly showing the positions of the resonators. Dashed lines indicate bias levels of interest as discussed in the text, and the cyan dot is our chosen operating bias and frequency for the next calibration.

pulse sequence and typical results are shown in Fig. 5.1.

Because we couple to our readout resonators in reflection, we can best infer the positions of the resonators by looking for regions where the gradient of the signal phase vs readout frequency is large. Based on the positions of the resonators as a function of flux bias, we can already learn a great deal about our system. First, if a resonator moves as a function of flux bias, we know that our flux biasing line works and we also learn which readout resonator and qubit this flux bias line is connected to. Typically, the resonator to qubit correspondence is known beforehand from the circuit design, but this scan acts as a sanity check. Second, if the other resonators in the band do not move, then we know that on a coarse level, flux crosstalk between the qubits is small. Finally, we can infer the resonator frequency as a function of flux bias, which is a periodic function due to the qubit SQUID's Φ_0 periodicity.

The resonator frequency vs. flux bias curve has three points of interest. Where the frequency vs. flux curve is flat, the qubit is at its maximum frequency, which is also known as the *flux insensitive point*. In theory, this point occurs when the flux bias is zero but due to the stray magnetic fields, the flux insensitive point may be offset. Halfway between insensitive points, the flux bias at the SQUID is $\phi_0/2$, and the qubit is near zero frequency. This point is important because in order to isolate a qubit from its neighbors, we tune the neighbors to near zero frequency. Therefore, we perform this resonator vs flux bias calibration on *all* qubits prior to moving forward with single qubit calibrations. In between the maximum and minimum frequency points, the qubit and the resonator

frequencies meet, as evidenced by the large avoided level crossing. Because the resonator is designed to have strong coupling to the outside environment for measurement, we can make use of this crossing point in order to quickly reset the qubit state.

5.3 Qubit Spectroscopy

Now that we have a readout frequency vs flux curve, we pick a readout frequency and static flux bias to operate temporarily. We additionally pick a qubit transition frequency that we expect will be in the tunable band of qubit frequencies but above the resonator crossing point¹. We then dynamically bias the qubit in a range of flux biases from the flux insensitive point to the resonator crossing point. At each bias level, we drive the qubit's XY control line with a microwave pulse which is fixed at our desired qubit frequency but with variable amplitude. Finally we bias back to our temporary operating point to drive the readout resonator and measure its response. When the dynamic bias places the qubit on resonance with our microwave pulse, the qubit undergoes Rabi oscillations, cycling between the $|0\rangle$ and $|1\rangle$ states as we increase the drive amplitude. Because the resonator frequency is different for the two qubit states, we will observe oscillations in the resonator response as a function of drive amplitude. To characterize these oscillations, we measure the resonator response when we do nothing to the qubit and compute how much the resonator is displaced in IQ space for each bias and drive amplitude. The results of this

¹When choosing readout and qubit frequencies, we take into account the designed readout coupling parameters and choose frequencies where we expect the state discrimination between qubit states to be reasonably good.

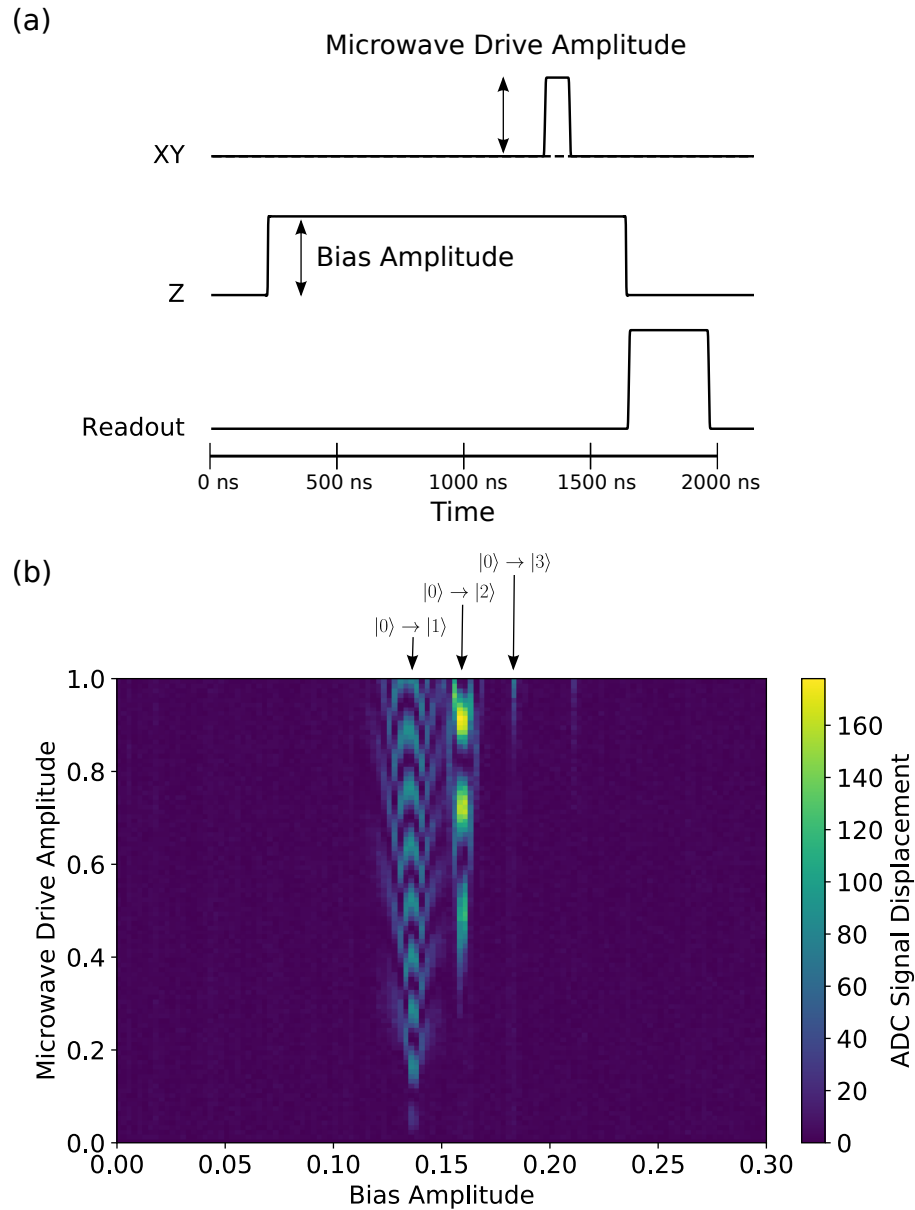


Figure 5.2: Fixed frequency qubit spectroscopy. (a) Pulse sequence. From Fig. 5.1, we pick a static bias and readout frequency (cyan dot). We then microwave drive the qubit at a fixed frequency and dynamically bias the qubit SQUID, while varying the bias level and drive amplitude. As in Fig. 5.1, we pad the dynamic bias to account for bias settling and possible timing offsets. Finally, we drive the resonator and measure its response. (b) Spectroscopy data. We first perform a control experiment with no microwave drive or bias to establish a baseline IQ response. Then, at each drive amplitude and bias level, we take the measured signal and subtract it from the baseline signal, and plot the magnitude of this difference. When the bias level brings the qubit on resonance with the drive, we observe oscillations in the resonator response corresponding to Rabi driving of the qubit. Transitions to higher states are also visible as smaller chevrons.

experiment are shown in Fig. 5.2.

As expected, we see the classic Rabi chevron pattern when the bias level brings our qubit on resonance with the microwave drive, and we set our static operating bias to this value. However, note that we also see other high signal displacement chevrons and peaks at other flux biases. These signals correspond to transitions from the ground state to higher levels. While the frequency for the $|0\rangle \rightarrow |2\rangle$ transition is of order 10 GHz and there is theoretically no charge matrix element which connects the two state, a transition may still occur if the transmon is driven sufficiently hard at $\omega_{20}/2$ by first populating a virtual state between $|0\rangle$ and $|2\rangle$. Such a transition is also known as a *two photon* transition, and the mechanism also extends to the higher states. We will ignore these transitions for now and assume that the anharmonicity is what we designed, but we will return to the problem of finding the $|1\rangle \rightarrow |2\rangle$ transition later.

5.4 Rabi

Having found the operating flux bias for our qubit transition, we set our new readout frequency according to our resonator frequency vs. flux calibration, and proceed to tune up microwave control of the qubit. For our discrete microwave pulse amplitudes, we will use cosine-shaped pulses, which are described by

$$X(t) = \frac{A}{2} \left(1 - \cos \frac{2\pi t}{\tau} \right), \quad (5.1)$$

where A is the peak pulse amplitude and τ is the pulse length. Compared to square or Gaussian shaped envelopes, cosine envelopes have the nice property that their value

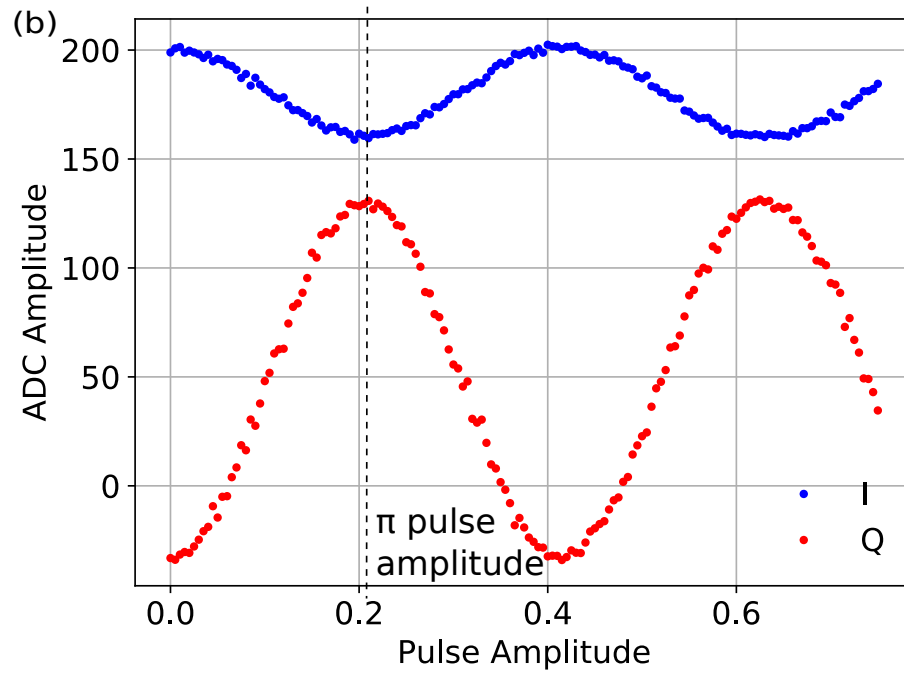
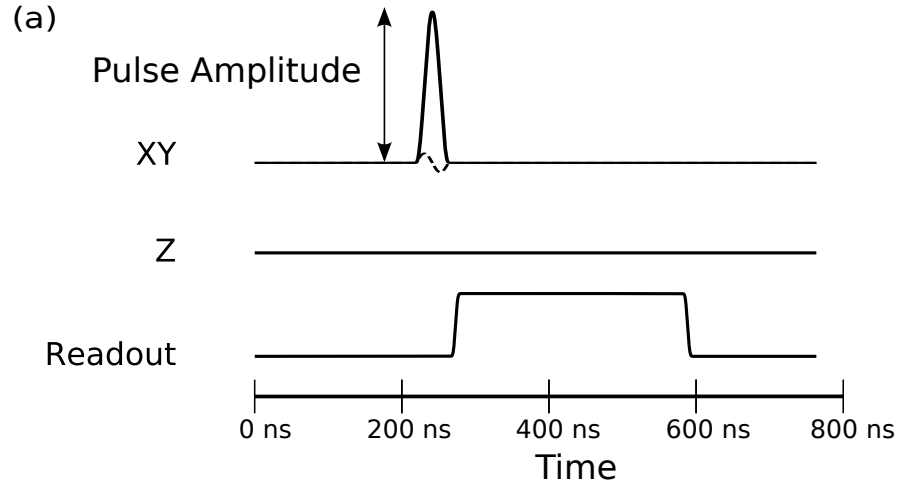


Figure 5.3: π pulse amplitude calibration. (a) Pulse sequence. We drive the qubit with a fixed length cosine envelope and measure the resonator response, while varying the pulse amplitude. (b) Measured I and Q response of the resonator as a function of pulse amplitude. We observe Rabi oscillations between the $|0\rangle$ and $|1\rangle$ states, and choose the first half period as the amplitude for a π pulse.

and first derivative go to zero at the pulse boundaries. In addition to the in-phase cosine envelope, we also add a weighted quadrature correction [100] which is the derivative of cosine envelope

$$Y(t) = -\frac{A}{2} \frac{\alpha}{\eta} \sin \frac{2\pi t}{\tau} \quad (5.2)$$

where α is weighting factor typically between 0 and 1.5, and $\eta = f_{21} - f_{10}$ is the anharmonicity. This correction is known as the Derivative Reduction by Adiabatic Gate (DRAG) correction, and is used to correct for errors due to the presence of the $|2\rangle$ state. We will address this correction in detail in Chapter 6, but for now we simply set $\alpha = 1.0$.

To calibrate our cosine pulses, we choose a fixed pulse length τ , typically between 15 ns and 25 ns to be sufficiently long compared to the timescale of the anharmonicity, $1/\eta = 5$ ns. We then vary the amplitude and again observe Rabi oscillations in the measured resonator response, as shown in Fig. 5.3. We pick the amplitude corresponding to the first half period of the oscillation as the amplitude for a π pulse which brings the qubit from $|0\rangle$ to $|1\rangle$. We also choose exactly half this amplitude as the amplitude for a $\pi/2$ pulse which brings the qubit from $|0\rangle$ to an equal superposition state.

5.5 Readout Frequency

Up until now, we have used the frequency of the resonator when the qubit is in $|0\rangle$ as the frequency at which we measure the resonator's response. However, this frequency is not necessarily the optimal frequency for distinguishing between the $|0\rangle$ or $|1\rangle$ states. In order to calibrate the readout frequency, we prepare the qubit in either $|0\rangle$ or $|1\rangle$ then

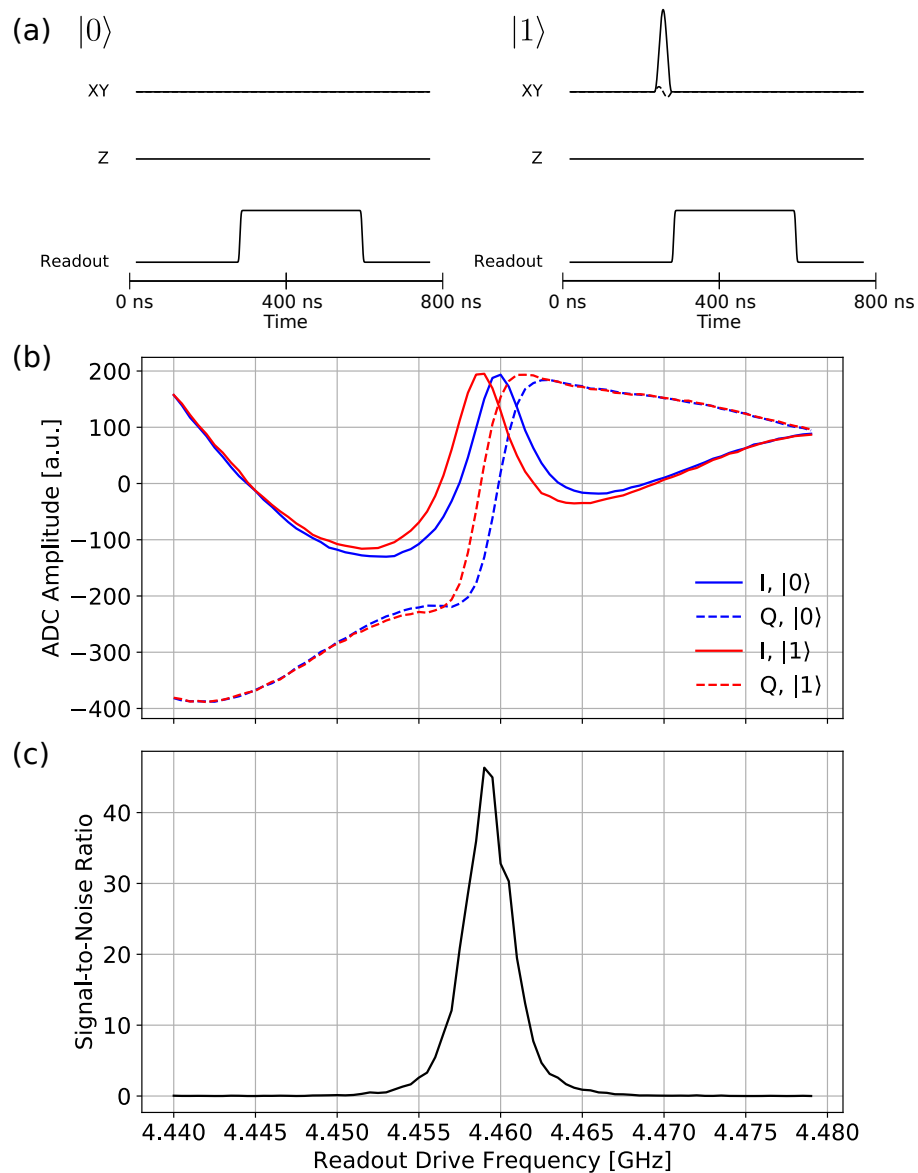


Figure 5.4: Readout frequency calibration (a) Pulse sequence. We prepare the qubit in either $|0\rangle$ or $|1\rangle$ and measure the resonator's response while sweeping the readout drive frequency. (b) Measured I and Q response of the resonator in either qubit state. (c) Signal to noise ratio inferred from the distance between the IQ responses when the qubit is in $|0\rangle$ or $|1\rangle$. We choose the maximal SNR point as our new readout frequency.

measure the resonator’s response. For each state preparation, we record the average and standard deviation of the complex response, and compute the signal to noise ratio (SNR) [68]

$$SNR = \frac{(x_0 - x_1)^2}{2\sigma^2} \quad (5.3)$$

where x_0 and x_1 are the average positions of the IQ response when the qubit is in $|0\rangle$ and $|1\rangle$, respectively, and σ is the standard deviation of the measured signal ². The readout frequency is optimized when the SNR is maximal, and this frequency is typically somewhere between the resonator frequencies corresponding to the qubit being in $|0\rangle$ or $|1\rangle$.

5.6 Readout Amplitude

The next step in calibrating the readout of the qubit is to calibrate the readout amplitude. If the amplitude is too small, we do not collect enough microwave photons in our readout chain to accurately discriminate the qubit states. If the amplitude is too high, then we induce transitions between various levels of the transmon. In Chapter 9, we will thoroughly explore the physics of readout induced transitions. For the purposes of calibration, we use the fact that these transitions take the qubit outside of the 0-1 subspace and that the turn-on for the transitions is sudden.

As in the readout frequency calibration, we again prepare the qubit in either $|0\rangle$ or $|1\rangle$.

²We skipped over some technical details of how exactly this quantity is computed from the data, which can be found in Daniel Sank’s thesis [68]. We also assumed σ is the same for both qubit states, which is true in most cases

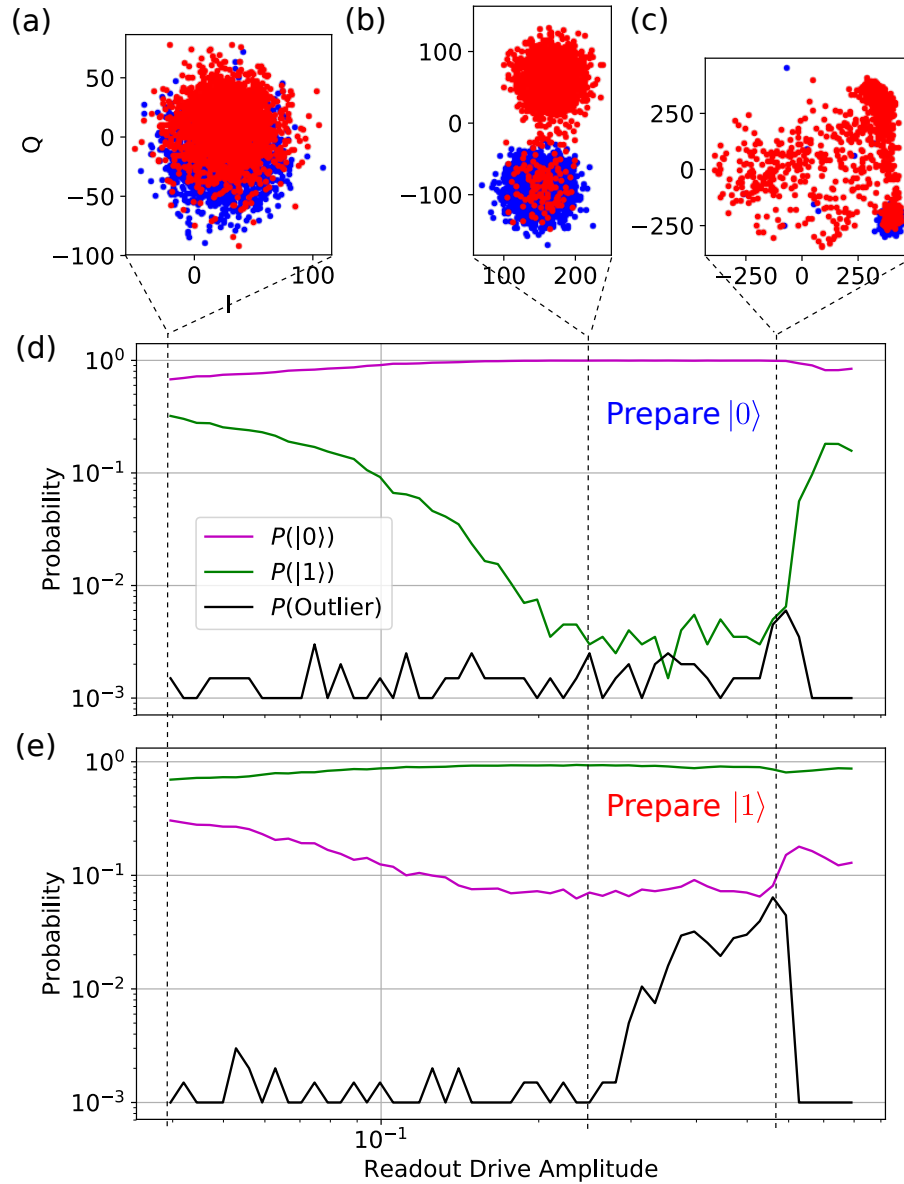


Figure 5.5: Readout amplitude calibration. The pulse sequence is the same as in Fig. 5.4, except we vary the drive amplitude instead of frequency. (a-c) Raw IQ response after preparing the $|0\rangle$ state (blue) or $|1\rangle$ state (red) for (a) low amplitude, (b) optimal amplitude, and (c) high amplitude. When the amplitude is too low, the separation between $|0\rangle$ and $|1\rangle$ is insufficient for state discrimination, and when the amplitude is too high, the transmon is driven to higher states. (d-e) Discriminated state probabilities when preparing the (d) $|0\rangle$ state and (e) $|1\rangle$ state. In addition to distinguishing between the two states, we also determine the probability that an IQ point is an outlier from the distribution for either state.

We then readout the qubit while varying the amplitude of the readout drive. For each readout amplitude and state preparation, we record the probability that the qubit was measured either in $|0\rangle$, $|1\rangle$, or was an outlier from the IQ cloud of either state. Outliers indicate that the qubit is driven to a higher state than what we can measure, and are undesirable. We choose the readout amplitude which maximizes measuring the expected prepared states for all state preparations after excluding for outliers.

5.7 Ramsey

Having calibrated the state discrimination of our qubit, we will now work in terms of measured state probabilities rather than IQ amplitudes for the remainder of the calibrations. Previously, we determined the qubit frequency spectroscopically. We now fine tune this frequency using a Ramsey fringe. In the rotating frame, if the qubit is prepared along the X axis of the Bloch sphere with a $Y/2$ pulse, it should remain stationary and if we execute a second $Y/2$ rotation some time later, we should always measure $|1\rangle$. If our knowledge of the qubit frequency is faulty and our assumed frame does not match the qubit's rotation frame, then the measured qubit state after a $Y/2$ rotation will appear to precess at the difference frequency $\Delta\omega$ between the assumed frame and the actual qubit frequency. Thus, if we perform this $Y/2$ - delay - $Y/2$ sequence with variable delay, we should be able to calibrate the qubit frequency by observing the resulting oscillations.

The above procedure works in theory, but has two practical drawbacks. First, if the error frequency is of order the phase decoherence rate of the qubit, the intended

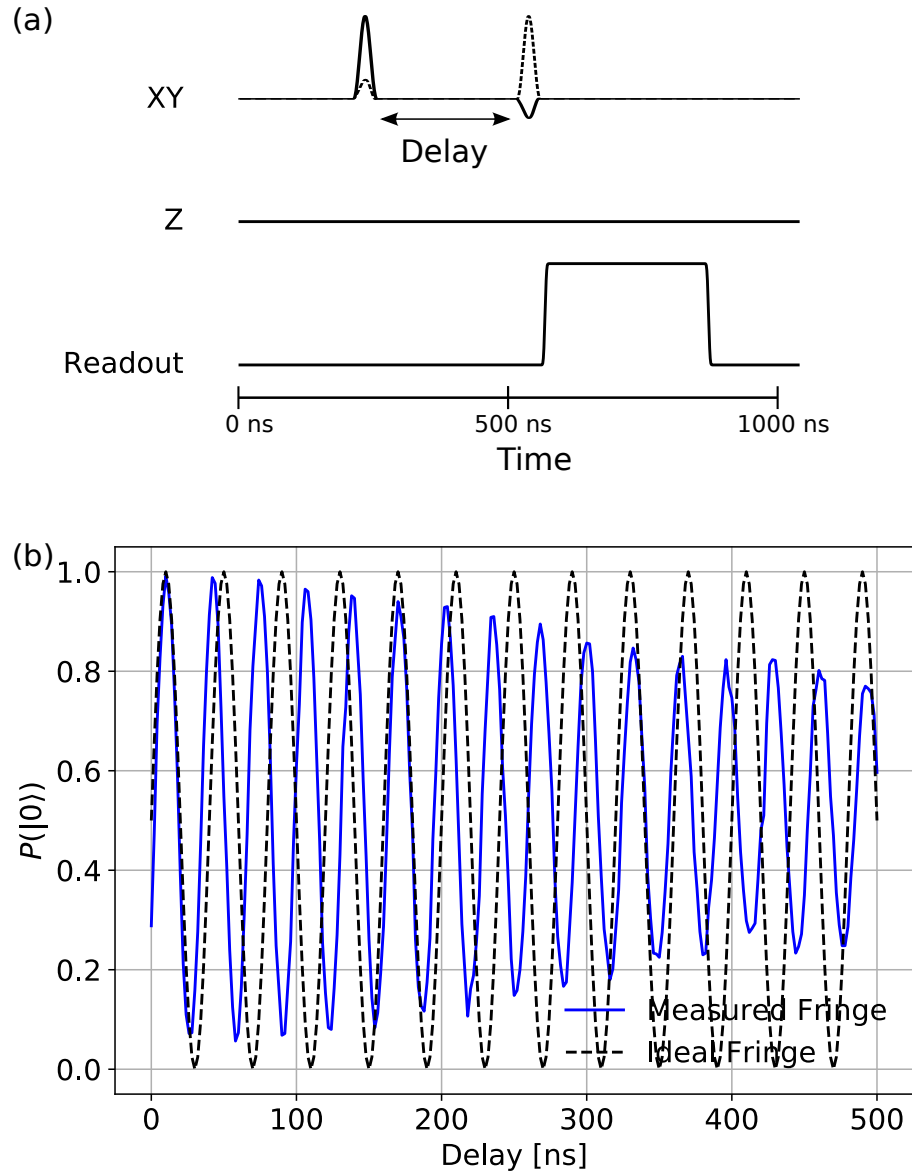


Figure 5.6: Frequency calibration using a Ramsey fringe. (a) Pulse sequence. We perform a $\pi/2$ pulse, then wait a variable amount of time τ , followed by a second $\pi/2$ pulse. The phase of second $\pi/2$ pulse is chosen to be $\Delta\omega\tau$ where $\Delta\omega$ is the desired fringe frequency. (b) Results with an expected fringe frequency of 25 MHz. The actual measured fringe oscillates at 31 MHz, indicating that the qubit frequency is 6 MHz less than what we expected.

oscillations would be difficult to distinguish from the effects of decoherence. Second, there is an ambiguity between positive and negative frequencies, and while the sign can be determined by looking at the phase of the oscillation fringe, we would like to use a more robust procedure. To overcome these drawbacks, we intentionally choose a frame which is detuned from the qubit frequency so that we observe oscillations even if our frequency is perfectly calibrated. An error in our qubit frequency would result in a difference from the expected oscillation frequency, as shown in Fig. 5.6. Essentially, we have upconverted our signal to a frequency of 10s of MHz to suppress the presence of low frequency dephasing effects in our data.

5.8 Pulse Corrections

As we alluded to in Eqn. 5.2, our microwave pulses use a quadrature derivative correction to correct for the effects of the second excited state. The weighting parameter for the correction depends on the anharmonicity of the qubit, so we now calibrate the second transition frequency f_{21} . We first prepare the qubit in $|1\rangle$, then excite it with a second low amplitude microwave pulse. We sweep the frequency of this second pulse in the region where we expect f_{21} to be. When the second pulse is on resonance with f_{21} , we obtain maximal transfer of $|1\rangle$ to $|2\rangle$. Because we have not explicitly calibrated readout of the $|2\rangle$ state, we add a third pulse which brings $|1\rangle$, but not $|2\rangle$, back to $|0\rangle$. The total effect of this sequence is that we should find a peak in the $|1\rangle$ probability when the second pulse is on resonance with f_{21} , as shown in Fig 5.7.

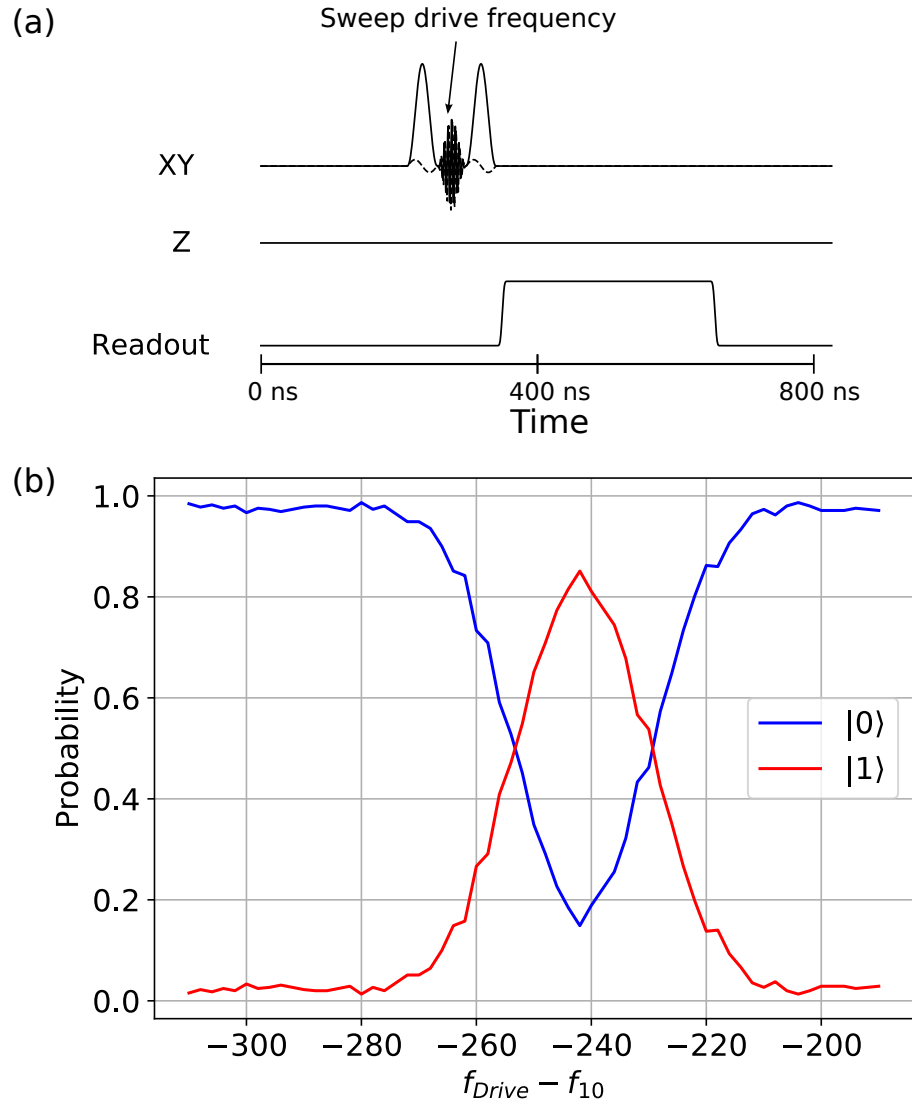


Figure 5.7: Calibration of f_{21} (a) Pulse sequence. We prepare $|1\rangle$ with a π pulse, then drive a microwave pulse while sweeping its frequency around the expected f_{21} frequency. Finally, we depopulate $|1\rangle$ back to $|0\rangle$ with a second π pulse. (b) When the second pulse is on resonance with f_{21} , we move $|1\rangle$ to $|2\rangle$ and the depopulation pulse has no effect. We measure a peak in the final $|1\rangle$ state probability, since $|2\rangle$ appears as $|1\rangle$.

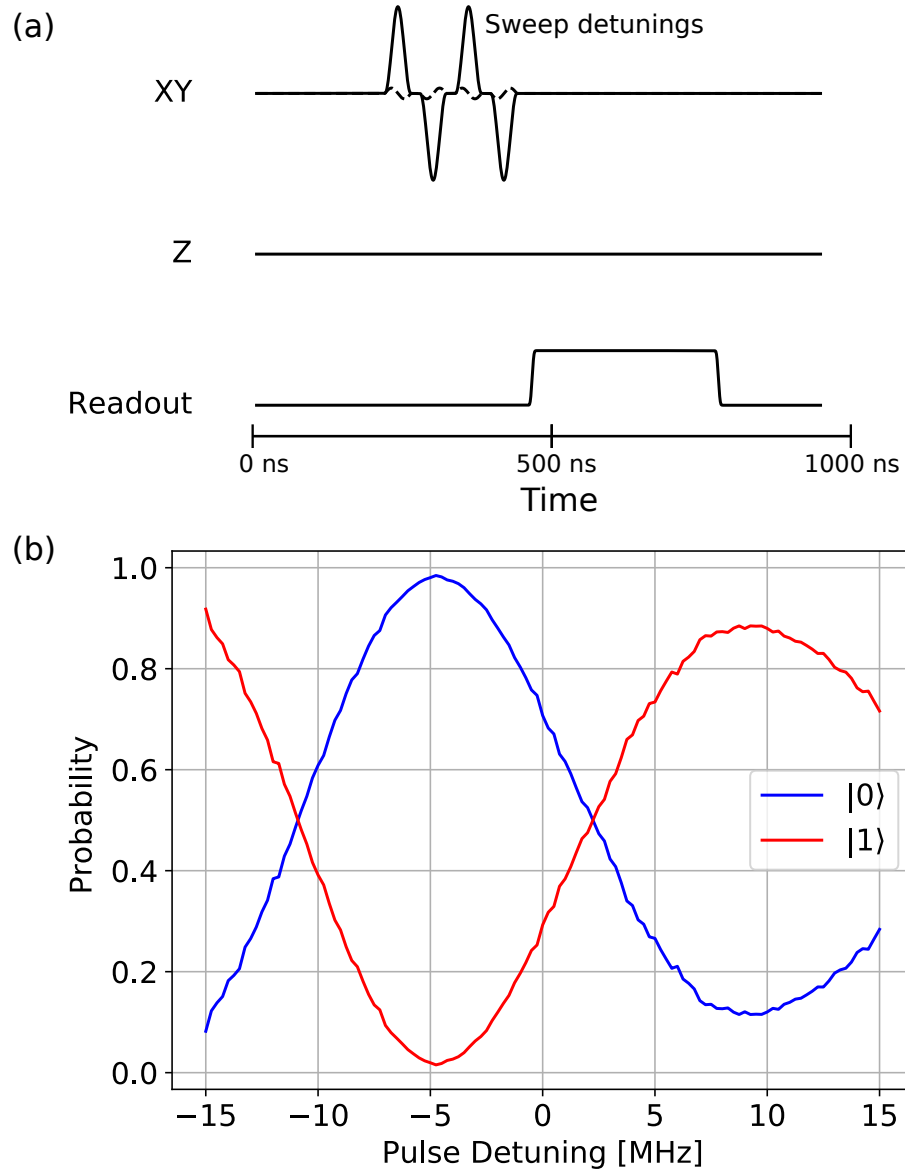


Figure 5.8: Correcting for phase errors due to the weak nonlinearity (a) Pulse sequence. We perform two consecutive $(\pi, -\pi)$ sequences, which should give $|0\rangle$ when there are no phase errors. We sweep the frequency of a detuning on the π pulses. (b) When the detuning fully corrects for the phase error, we see a peak in the $|0\rangle$ population.

With a calibrated anharmonicity, we can now correctly weight the quadrature correction in Eqn. 5.2 to minimize leakage out of the 0-1 subspace due to our microwave pulses. However, as we will see in Chapter 6, we are still susceptible to phase errors caused by the coupling of the $|1\rangle$ and $|2\rangle$ states. To compensate for these phase errors, we apply a second correction which detunes the microwave pulse. We calibrate the detuning by executing multiple π and $-\pi$ sequences while sweeping the pulse detuning. When the phase error is minimized, this sequence is equivalent to the identity and we maximize the probability of measuring the $|0\rangle$ state, as shown in Fig 5.8.

5.9 Rabi, Revisited

After calibrating the qubit frequency and pulse corrections, we fine tune the amplitudes of our π and $\pi/2$ pulses. To increase our sensitivity to the pulse amplitude, we execute many π rotations in series, which should have the same effect as a single π rotation if the amplitude is correct and the number of rotations is odd. Additionally, while we previously assumed that the $\pi/2$ amplitude was exactly half that of the π amplitude, we now calibrate the $\pi/2$ amplitude separately by concatenating two $\pi/2$ rotations into a π rotation. This extra step is necessary to account for any nonlinearities with respect to amplitude that may be present in the electronics or any other part of the system.

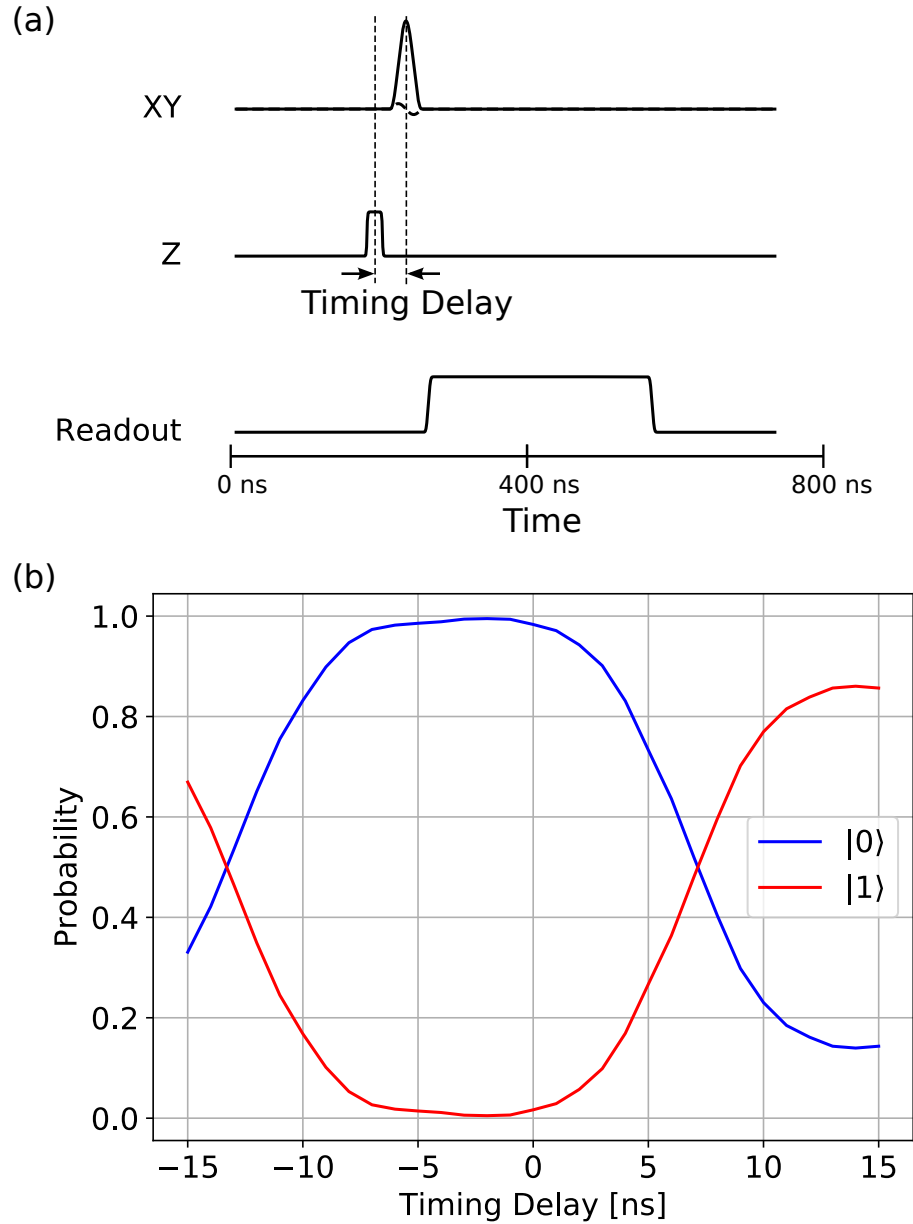


Figure 5.9: Calibrating the timing between the XY and Z channels. (a) Pulse sequence. We perform a π pulse and a flattop Z detune, while varying the timing between the two. (b) When the π pulse and detune are aligned, the qubit detuned sufficiently far off resonance that the π pulse has no effect, leaving the qubit in $|0\rangle$.

5.10 XY-Z Timing

With fully calibrated microwave rotations, we next turn to frequency control calibrations. We first calibrate the timing difference between the microwave drive and flux bias channels. In general, the timing difference exists due to differences in cable lengths or delays in the electronics. To measure the timing difference, we drive the XY channel of the qubit with a π pulse, and the Z channel with a flattop detuning pulse. When the centers of the π pulse and detuning pulse are aligned, the detuning pulse will bring the qubit far enough off resonance that the π pulse will have no effect on the qubit. We vary the timing between the π pulse and detuning pulse, and the correct timing difference is found when the measured $|1\rangle$ state population is minimized, as shown in Fig. 5.9. This timing difference is subsequently compensated for in future runs of the hardware.

5.11 Bias Amplitude to Frequency Function

The final calibration is determining the function that maps DAC voltages on the Z channel to the frequency of the qubit. Essentially, we perform the qubit spectroscopy experiment described in Fig. 5.2 earlier while sweeping both qubit frequency and bias amplitude. However, repeating this experiment over the full frequency range of the qubit would be time consuming. Instead, we perform the experiment adaptively. We know the qubit frequency at the present bias amplitude. We then move a small distance away in bias amplitude, and determine the qubit transition frequency using microwave spectroscopy,




Figure 5.10: Determining the mapping from flux bias amplitude to qubit frequency. The pulse sequence is the same as in Fig. 5.2, except we now discriminate to qubit states rather than recording IQ averages. (a) Qubit spectroscopy vs bias amplitude. We adaptively choose bias amplitudes and drive frequencies, with the spacing between amplitudes increasing as we move further away from the initial point. (b) For each bias amplitude, we fit the peak of $P(|1\rangle)$ vs frequency to determine the qubit’s frequency at that bias. Then, we fit the bias to qubit frequency curve to a model.

but only in a narrow frequency region around the known frequency point. We repeat this procedure, and as we measure more frequency vs bias amplitude points, we use this knowledge to determine the next amplitude and frequency sweeps. For example, we know that the amplitude to frequency curve should roughly follow Eqn. 2.127, and we can fit the data points as we go along to this equation to predict where the next amplitude to frequency point will be. Alternatively, if we do not want to restrict ourselves to a model, we can use the fact that that amplitude to frequency curve is relatively smooth in the region of interest and extrapolate the next amplitude and frequency points using the local slope and curvature. The resulting data is shown in Fig. 5.10, where a relatively low density of scans nevertheless allows us to determine and fit the full curve.

Chapter 6

Single Qubit Benchmarking

Having established how to calibrate single qubit control, we will now benchmark the performance of our single qubit gates. We first review the technique of Clifford-based randomized benchmarking, then apply it to our qubits to estimate the total gate fidelity. Next, we will use an extension of randomized benchmarking called purity benchmarking to estimate the contribution of decoherence to gate errors. Finally, we compare our purity benchmarking results to standard relaxation and dephasing measurements.

6.1 A Brief Review of Randomized Benchmarking

Given a unitary operation that we have implemented on our hardware, we want to measure how well our implementation corresponds to the ideal unitary. Errors in our implementation may be due to miscalibration of microwave pulse parameters, or due to noise in the environment leading to decoherence. To formally characterize the actual

action of our hardware, the standard method is to use quantum process tomography (QPT) [9, 101]. Given the Hilbert space spanned by our qubits, we would prepare an ensemble of input states which span the Hilbert space, apply our unitary to each of the input states, then perform quantum state tomography (QST) to reconstruct the output density matrix corresponding to each input states. From this mapping between inputs and outputs, we can determine a process matrix (typically through maximum likelihood estimation [102, 103]) and compare it to our ideal unitary.

The QPT protocol has a few drawbacks which hinder experimental usage. First, the number of state preparations and measurements scales exponentially in the number of qubits, so it can only be used for a relatively small number of qubits. Second, QPT relies on high fidelity single qubit gates for initialization and tomographic rotations which can lead to a self consistency problem when trying to use QPT to measure only the gate fidelity. Finally, in order to measure the average fidelity of our system, the QPT procedure would have to be applied to a sufficiently representative sample of the unitaries that can be implemented in the hardware.

To overcome these drawbacks, Ref. [104] proposed using random unitaries to obtain the average gate fidelity. Their key insight was that when an error process is averaged over the uniform space of unitaries, called the *Haar measure*, the result is a *depolarizing channel*, which maps any pure state to the maximally mixed state¹. The average induced error is proportional to the probability of depolarization, and in order to measure this

¹The depolarizing property assumes the error process does not lead to leakage out of the Hilbert space. We will see in the next chapter how to use randomized benchmarking to measure leakage.

probability, Ref. [104] proposed drawing random unitaries U and constructing motion reversal operations UU^\dagger . If we start from the ground state in the computational basis, then the probability of measuring the system in the ground state after the motion reversal operation is related to the strength of the error. A full protocol to experimentally estimate the fidelity is as follows:

1. Randomly choose m unitaries, U_1, U_2, \dots, U_m .
2. Construct their motion reversal operations, $U_1U_1^\dagger, U_2U_2^\dagger, \dots, U_mU_m^\dagger$.
3. Concatenate the motion reversal operations, apply them to the ground state, then measure the resulting ground state population.
4. Repeat steps 1-3 many times, and for many lengths m .

As the length of the sequences increases, the transformation of gate errors to a depolarizing channel leads to an exponential decay of the ground state population towards the maximally mixed state, where the rate of decay is a measure of the average gate fidelity. By using the decay rate to extract fidelity, we also become insensitive to initialization and measurement errors.

The most inefficient step in the above procedure is randomly choosing unitaries from the Haar measure, then implementing them with the gate primitives that are available on the quantum hardware. To address this inefficiency, Ref. [105] proposed restricting the unitaries to the *Clifford group* instead of the full Haar measure. Formally, the Clifford group consists of all unitary rotations which map the group of Pauli operators (or in

higher dimensions, their tensor products) back onto the Pauli operators but possibly with permutations [6, 7]. Equivalently, any gate in the Clifford group can be generated by the single qubit Hadamard, $\pi/2$ rotations about the Z axis ², and the two-qubit control-NOT gate. For the purpose of benchmarking gates, the Clifford group has a few important properties:

1. For a given Hilbert space size, there are only finitely many Cliffords. Thus, randomly drawing and implementing Cliffords is easier than implementing an arbitrary unitary.
2. The Cliffords form a mathematical group, so their inverse can always be found within the Clifford group.
3. Averaging an error process over the Clifford group is equivalent to averaging the error over the Haar measure. More formally, the Clifford group is a *unitary 2-design* [106], which means that any quantity averaged over the Clifford group up to polynomial degree 2 is equivalent to averaging over the entire Haar measure. By contrast, the Pauli group is a unitary 1-design.

²Note that this is sometimes confusingly called the $\pi/4$ phase gate.

6.2 Clifford Based Randomized Benchmarking of Xmon Qubits

We now describe the protocol used to benchmark single qubit Clifford gates, which largely follows from Ref. [107].

1. Initialize the qubit in the ground state.
2. Apply a sequence of m Cliffords to the qubit.
3. Apply an additional $(m + 1)$ th gate which inverts the whole sequence. As we established above, this inverse can be found in the Clifford group.
4. Measure the resulting ground state probability.
5. Repeat this procedure k times.
6. Repeat for multiple lengths m to build up an exponential decay.

This procedure is similar to the one used for motion reversal, except we only invert the sequence at the end since the inverse of any sequence of Cliffords can be efficiently found within the Clifford group.

To implement this procedure, we must establish how to perform Clifford rotations in our qubits. We begin by enumerating the single qubit Clifford group, which consists of all rotations on the Bloch sphere which map each of the six axial states (i.e. the six eigenvectors of the Pauli matrices) to one another while preserving their orthogonal

Table 6.1: Single Qubit Cliffords

Rotation angle	Axis direction (x, y, z)	Decomposition
None	None	I
π	(1, 0, 0)	X
π	(0, 1, 0)	Y
π	(0, 0, 1)	Y, X
$\pi/2$	(1, 0, 0)	X/2
$\pi/2$	(-1, 0, 0)	-X/2
$\pi/2$	(0, 1, 0)	Y/2
$\pi/2$	(0, -1, 0)	-Y/2
$\pi/2$	(0, 0, 1)	-X/2, Y/2, X/2
$\pi/2$	(0, 0, -1)	-X/2, -Y/2, X/2
π	(1, 0, 1)	X, -Y/2,
π	(-1, 0, 1)	X, Y/2,
π	(0, 1, 1)	Y, X/2,
π	(0, -1, 1)	Y, -X/2,
π	(1, 1, 0)	X/2, Y/2, X/2
π	(1, -1, 0)	-X/2, Y/2, -X/2
$2\pi/3$	(1, 1, 1)	Y/2, X/2,
$2\pi/3$	(-1, 1, 1)	Y/2, -X/2,
$2\pi/3$	(1, -1, 1)	-Y/2, X/2,
$2\pi/3$	(-1, -1, 1)	-Y/2, -X/2,
$-2\pi/3$	(1, 1, 1)	-X/2, -Y/2,
$-2\pi/3$	(-1, 1, 1)	X/2, -Y/2,
$-2\pi/3$	(1, -1, 1)	-X/2, Y/2,
$-2\pi/3$	(-1, -1, 1)	X/2, Y/2,

structure. Without loss of generality, consider only the $+Z$ vector. After a Clifford rotation, there are six possible orientations of the vector, including staying in $+Z$. For each of these possibilities, there are four possible orientations for $+X$, since the rotated $+X$ must remain orthogonal to $+Z$. Therefore, there are 24 possible rotations, as listed in the first two columns of Table 6.1.

The single qubit Cliffords fall into four general categories: the Pauli operators I, X, Y, and Z; $\pi/2$ rotations about the X, Y, and Z axes; Hadamard rotations, where the most well known Hadamard rotates +Z to +X, and all Hadamards produce the identity when applied twice; and $2\pi/3$ rotations, for example the rotation which cycles between the +Z, +X, and +Y states.

Next, we discuss how to implement these rotations in a superconducting qubit. Because we are primarily concerned with benchmarking microwave gates, we will only have access to π and $\pi/2$ rotations about the X and Y axes. Fortunately, all of the single qubit Cliffords can be generated using at most three of these elementary gates [64]. For example, the rotation axis for a standard Hadamard contains a Z component, but is equivalent to an X gate followed by a Y/2 gate. The decompositions for all of the Clifford rotations are listed in the third column of Table 6.1. On average, a Clifford contains 1.5 $\pi/2$ gates and 0.375 π gates, for a total of 1.875 physical gates.

Now that we have established how to implement the single qubit Clifford gates, we can perform Clifford based RB. Figure 6.1 shows the results of randomized benchmarking with π and $\pi/2$ pulses which are 14 ns in length, calibrated using the procedure described in Chapter 5. As expected, the average sequence fidelity decays exponentially with the number of Clifford gates in the sequence. We fit the average sequence fidelity F to the equation

$$F(m) = Ap^m + B \tag{6.1}$$

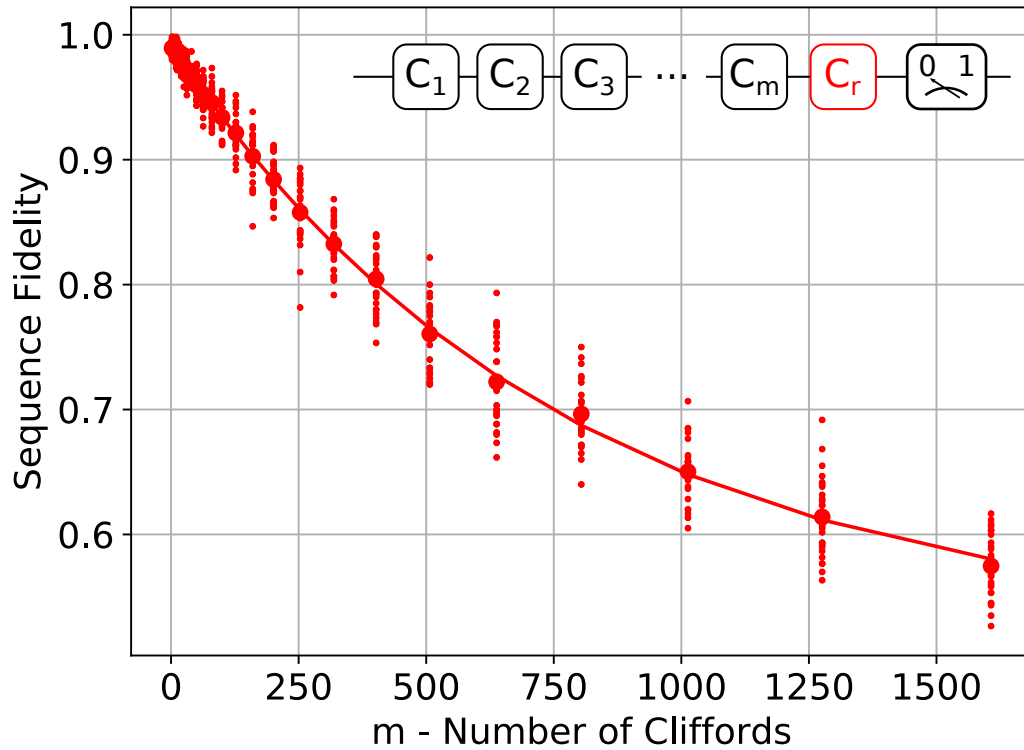


Figure 6.1: Results from randomized benchmarking of single qubit microwave gates on a transmon. For each sequence length m , we choose thirty different random Clifford sequences C_1, C_2, \dots, C_m which are inverted by appending the appropriate recovery Clifford C_r . The fidelity of each sequence is simply the probability of measuring $|0\rangle$, and the fidelity for each sequence is shown in the small dots. The larger dots are the averages at each sequence length. Also shown is an exponential fit to the averaged data.

where $1 - p$ is the rate of depolarization, and the parameters A and B capture state preparation and measurement errors [107]. The quantity p is related to the average error per Clifford r_{Clifford} by

$$r_{\text{Clifford}} = (1 - p)(1 - 1/d) \tag{6.2}$$

where $d = 2^n$ is the dimensionality of the Hilbert space³. For the dataset above, we find that $A = 0.469 \pm 0.007$, $B = 0.522 \pm 0.007$, $1 - p = 1.30 \pm 0.03 \times 10^{-3}$, and $r_{\text{Clifford}} = 6.5 \pm 0.15 \times 10^{-4}$. We can further estimate the error per gate by dividing the Clifford error by the number of physical gates per Clifford, 1.875, and obtaining $r_{\text{gate}} = 3.5 \pm 0.1 \times 10^{-4}$.

6.3 Purity Benchmarking

Given the measured fidelity in Fig. 6.1, we would now like understand the source of errors. At the very least, we would like to distinguish between coherent errors due to mistakes in our calibration, versus incoherent errors due to noise in the qubit’s environment.

The disadvantage of randomized benchmarking is that it only provides one number, the total error rate per gate. However, in 2015, Ref. [108] proposed an extension of randomized benchmarking called *purity benchmarking* (PB) which measures only the incoherent contribution to gate error. As in standard RB, we initialize our qubit in $|0\rangle$ and apply random sequences of Clifford gates of varying lengths. However, rather than appending an inversion gate which brings the qubit back to $|0\rangle$, we instead perform

³Intuitively, if the n qubits are fully depolarized, there is still a $1 - 1/2^n$ probability of measuring the expected outcome.

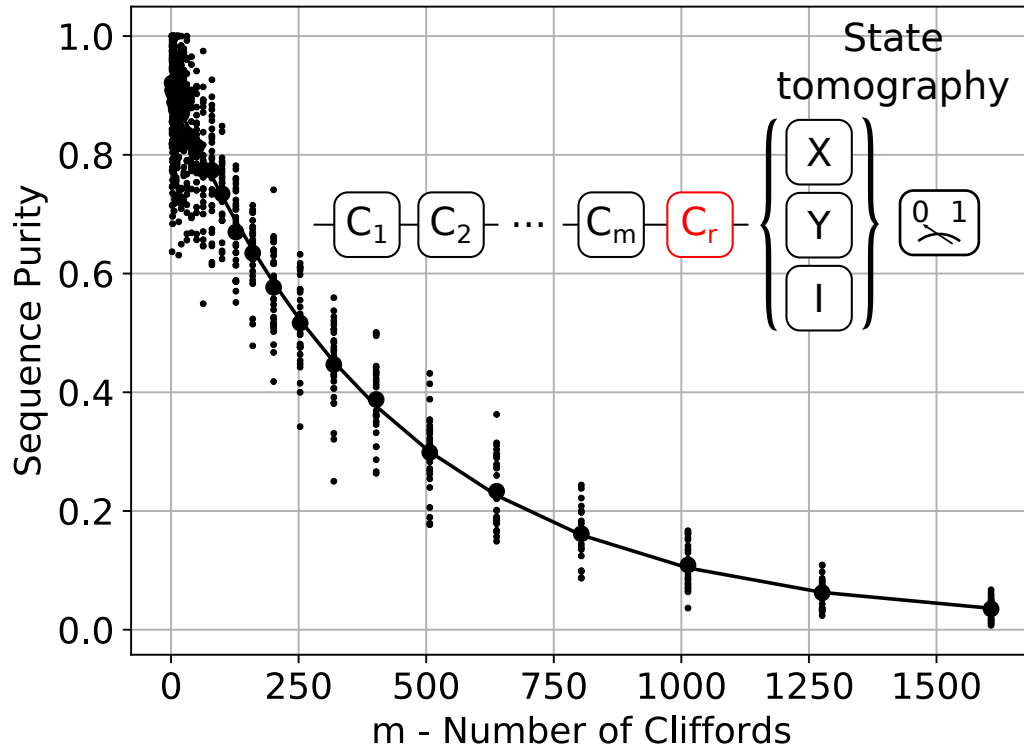


Figure 6.2: Results from purity benchmarking the same Clifford sequences as in Fig. 6.1.

quantum state tomography (QST) to determine the state of the qubit after the random sequence. While randomized benchmarking measures the fidelity with which the random Clifford sequence prepared the expected state, purity benchmarking measures the fidelity with which the sequence prepared *any* pure state. In other words, purity benchmarking does not penalize the Clifford sequence preparing the incorrect state due to miscalibrations in our pulses, since we would prepare the same incorrect state from run to run of the experiment.

To perform purity benchmarking, we must first define the purity of a quantum state. The standard definition of quantum state purity is

$$P = \text{tr}(\rho^2) \tag{6.3}$$

where ρ is the density matrix [9]. The quantity P ranges from 1, for a completely pure state $\rho_{\text{pure}} = |\psi\rangle\langle\psi|$, to $1/d$ for the maximally mixed state. In Ref. [108], the purity is rescaled so that it ranges from 0 to 1:

$$P_{\text{normalized}} = \frac{d}{d-1} \left(P - \frac{1}{d} \right). \quad (6.4)$$

In subsequent text, we will refer to $P_{\text{normalized}}$ as the purity of the state. For a single qubit, this definition of purity is equivalent to the square of the Bloch vector length, which can be measured as

$$P_{\text{normalized}} = \langle\sigma_x\rangle^2 + \langle\sigma_y\rangle^2 + \langle\sigma_z\rangle^2. \quad (6.5)$$

where $\langle\sigma_x\rangle$ is the expectation value of the Pauli X operator.

In Fig. 6.2, we show the measured purity for the same random sequences which were benchmarked in Fig. 6.1. The average purity also decays exponentially as a function of sequence length, and we fit this decay to

$$P(m) = A'u^{m-1} + B' \quad (6.6)$$

where u is the *unitarity* [108]. The unitarity can then be used to estimate the incoherent error per Clifford [109]

$$r_{\text{incoherent, Clifford}} = \frac{d-1}{d} (1 - \sqrt{u}). \quad (6.7)$$

For the dataset in Fig. 6.2, we find $A = 0.900 \pm 0.004$, $B = 0.012 \pm 0.003$, $1 - u = 2.22 \pm 0.04 \times 10^{-3}$, and $r_{\text{incoherent, Clifford}} = 5.6 \pm 0.09 \times 10^{-4}$. As with standard RB, we estimate the incoherent error per gate by dividing by 1.875 to obtain $r_{\text{incoherent, gate}} = 2.99 \pm 0.05 \times 10^{-4}$. We see that for the 14 ns gates, 85% of the total gate error is due to

decoherence.

6.4 Benchmarking Errors vs Pulse Length

Having established that our calibration procedure produces high fidelity 14 ns gates, we proceed to test our calibration more stringently by measuring total and incoherent error as a function of pulse length. We choose a range of pulse lengths between 10 ns and 100 ns, with the length of the π pulses always equal to the length of $\pi/2$ pulses. After calibrating pulses at a given length using the procedure described in Chapter 5, we perform randomized benchmarking and purity benchmarking to measure the total and incoherent error. We repeat the calibration and benchmarking for each pulse length five times. Collecting the results for all pulse lengths took approximately 10 hours, and the results are shown in Fig. 6.3.

In the averaged data shown in Fig. 6.3, we see two important trends. First, the total error tracks well with the incoherent error, indicating that our calibration procedure is valid over a wide range of pulse lengths. Second, the incoherent error scales linearly with pulse length, and by fitting the data, we find the following empirical equation

$$\text{Incoherent error per gate} = \frac{t_{\text{gate}}}{T_{\text{error}}}, \quad (6.8)$$

where $T_{\text{error}} = 49.8 \pm 0.5 \mu\text{s}$.

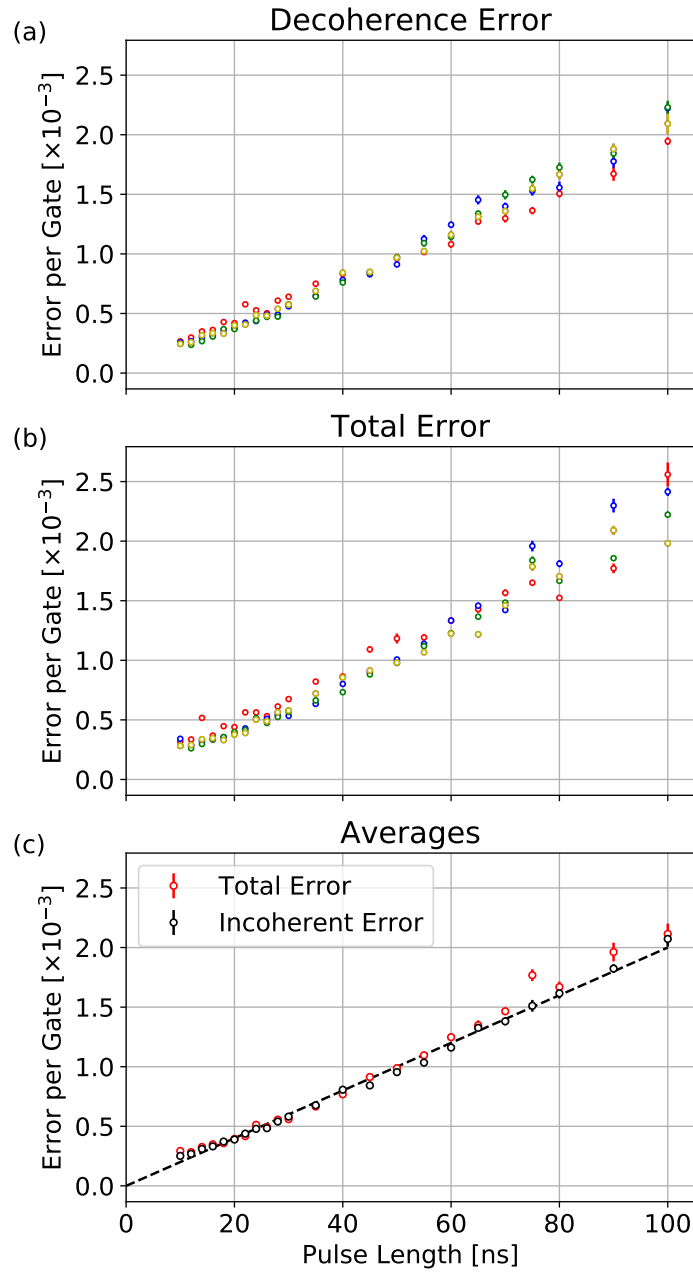


Figure 6.3: Error benchmarking for different microwave pulse lengths. (a) Error due to decoherence as measured by purity benchmarking versus the length of π and $\pi/2$ pulses. Benchmarking was repeated 5 times and each color represents a different run of the measurement. (b) Total error as measured by standard randomized benchmarking, for the same measurement runs as in (a). (c) Averages of the total (red) and decoherence (black) errors. The black dashed line is a linear fit to the incoherent error vs pulse length, and represents an incoherent error rate of one every $50 \mu s$.

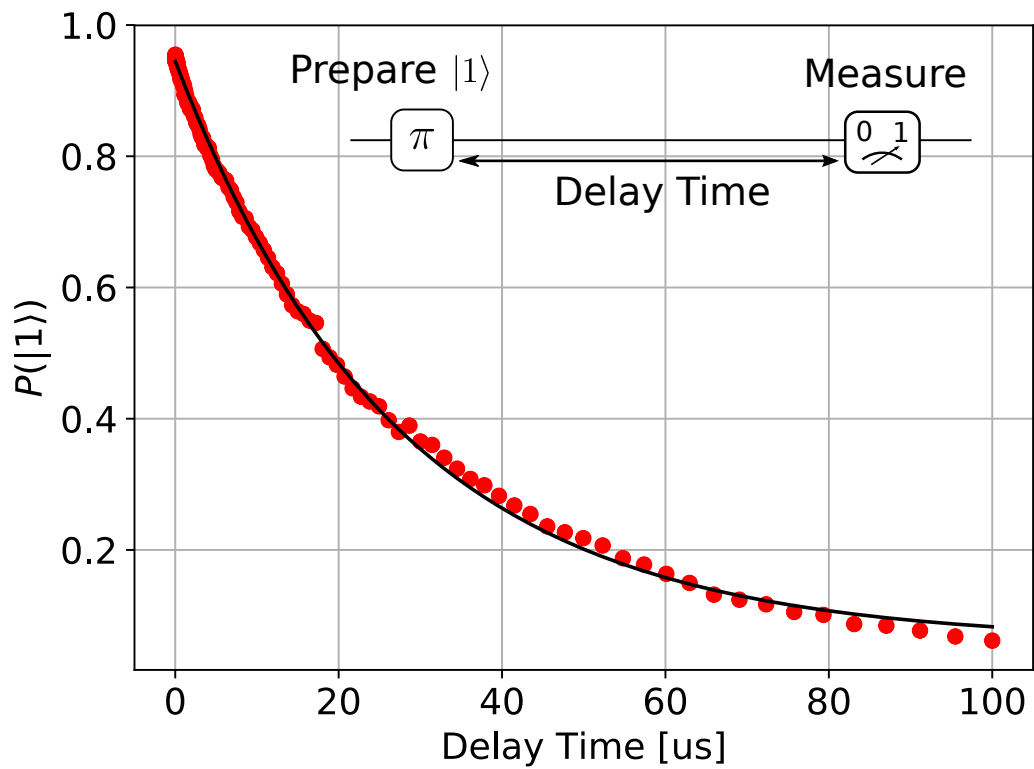


Figure 6.4: Measurement of energy relaxation in a superconducting transmon. As a function of the delay time between state preparation and measurement, the population of the qubit decays exponentially with characteristic time $T_1 = 27.2 \pm 0.3 \mu\text{s}$

6.5 Coherence Measurements

To understand our measured error time scale T_{error} , we now turn to more traditional measures of decoherence. Like many other quantum systems, superconducting qubits are susceptible to two types of decoherence: energy relaxation and dephasing. Energy relaxation in superconducting qubits manifests itself as a loss of population in the excited state over time. This loss typically occurs at a quasi-constant rate, leading to an exponential decay with a characteristic time constant T_1 . Figure 6.4 shows a typical T_1 measurement, from which we find $T_1 = 27.2 \pm 0.3 \mu\text{s}$. In the UCSB Xmons, T_1 is conjectured to be primarily dominated by coupling to lossy defects in the substrate, especially in the region near the Josephson junctions [110, 78, 80, 111].

Dephasing occurs when the frequency of the qubit is unstable. Since the qubit frequency determines the phase evolution of the qubit when it is in a superposition state, dephasing will manifest itself as uncontrolled variance in the qubit phase. As seen in Fig. 6.5, dephasing can be measured by preparing the qubit in a superposition, allowing the qubit to evolve for a time τ , then measuring phase coherence using state tomography. In the basic *Ramsey* sequence [112], the qubit is allowed to evolve freely, and we find that the phase coherence decays on a time scale of a few microseconds.

During the free evolution of the Ramsey sequence, the qubit is sensitive to frequency noise from the quasi-static regime up to the inverse of the evolution time $1/\tau$ [113]. Quasi-static noise can be thought of as the qubit having a different but static frequency on every shot of the experiment. Each shot will result in a different phase of the qubit,


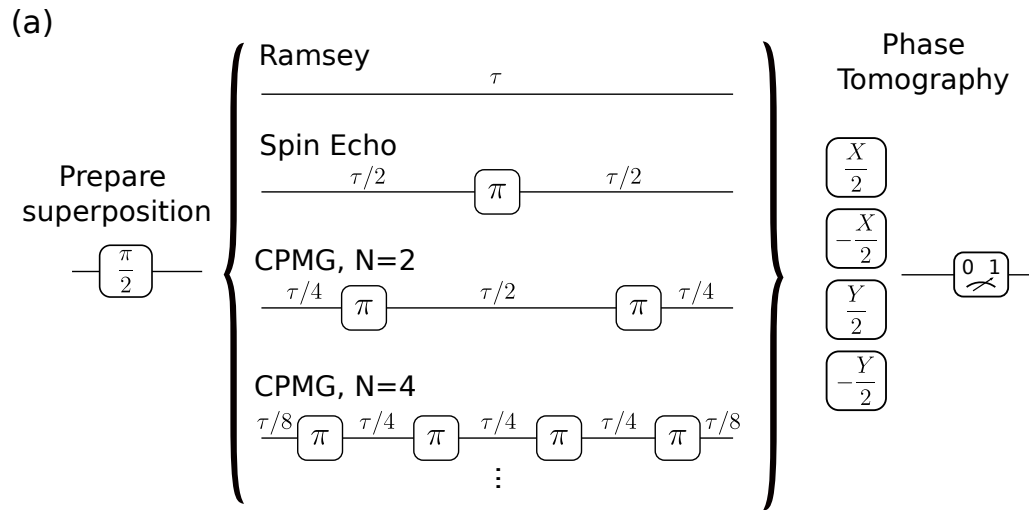


Figure 6.5: Measurement of dephasing in a superconducting transmon. (a) For each measurement, we begin by preparing the qubit in a superposition state. In a traditional Ramsey experiment, the qubit state freely evolves for a time τ . Echo pulses are inserted in the evolution to suppress quasi-static and low frequency noise. Note that the phases of the echo pulses are randomized in our measurement. Each measurement is repeated four times with a different final $\pi/2$ pulse phase to characterize phase coherence. (b) Phase coherence versus total evolution time for Ramsey and echo sequences. Each curve is fit to a model which accounts for white noise (exponential decay) and correlated noise (Gaussian decay).

and larger shot to shot variation in the frequency will lead to increased phase variance and decreased phase coherence. To mitigate this effect, we introduce *echo* π pulses, also known as refocusing pulses. For example, a single π pulse inserted at the center of the evolution time inverts the qubit's frame of reference, and any phase accumulated due to a spurious frequency offset in the first half of the sequence will be reversed in the second half. This single echo sequence is called a *spin echo* (or in NMR literature a *Hahn echo* [114, 115]), and as evidenced by the higher phase coherence in Fig. 6.5, this sequence suppresses quasi-static and other low frequency noise. Adding more echo pulses, for example in a Carr-Purcell-Meiboom-Gill (CPMG) sequence [116], suppresses low frequency noise to higher order and further increases phase coherence.

In order to fit our phase coherence data, we must consider three sources of dephasing: energy relaxation, white frequency noise, and correlated frequency noise. Energy relaxation contributes to exponential decay of phase coherence with a time scale of $2T_1^4$. White frequency noise also leads to an exponential decay, with characteristic time T_{ϕ_1} . Correlated frequency noise often manifests itself as $1/f$ noise, which has been found to be ubiquitous in SQUIDs and is thought to be caused by fluctuating paramagnetic spins coupled to the SQUID [117, 113, 118, 119]. A $1/f$ noise spectrum will lead to Gaussian decay of the phase coherence with characteristic time T_{ϕ_2} [113, 120]. Our total model for the phase coherence C is

$$C(\tau) = Ae^{-\tau/T_{\text{exp}} - (\tau/T_{\phi_2})^2} + B \quad (6.9)$$

⁴The factor of 2 comes from the fact that T_1 is a timescale for probability decay, while phase coherence is a measure of amplitude.

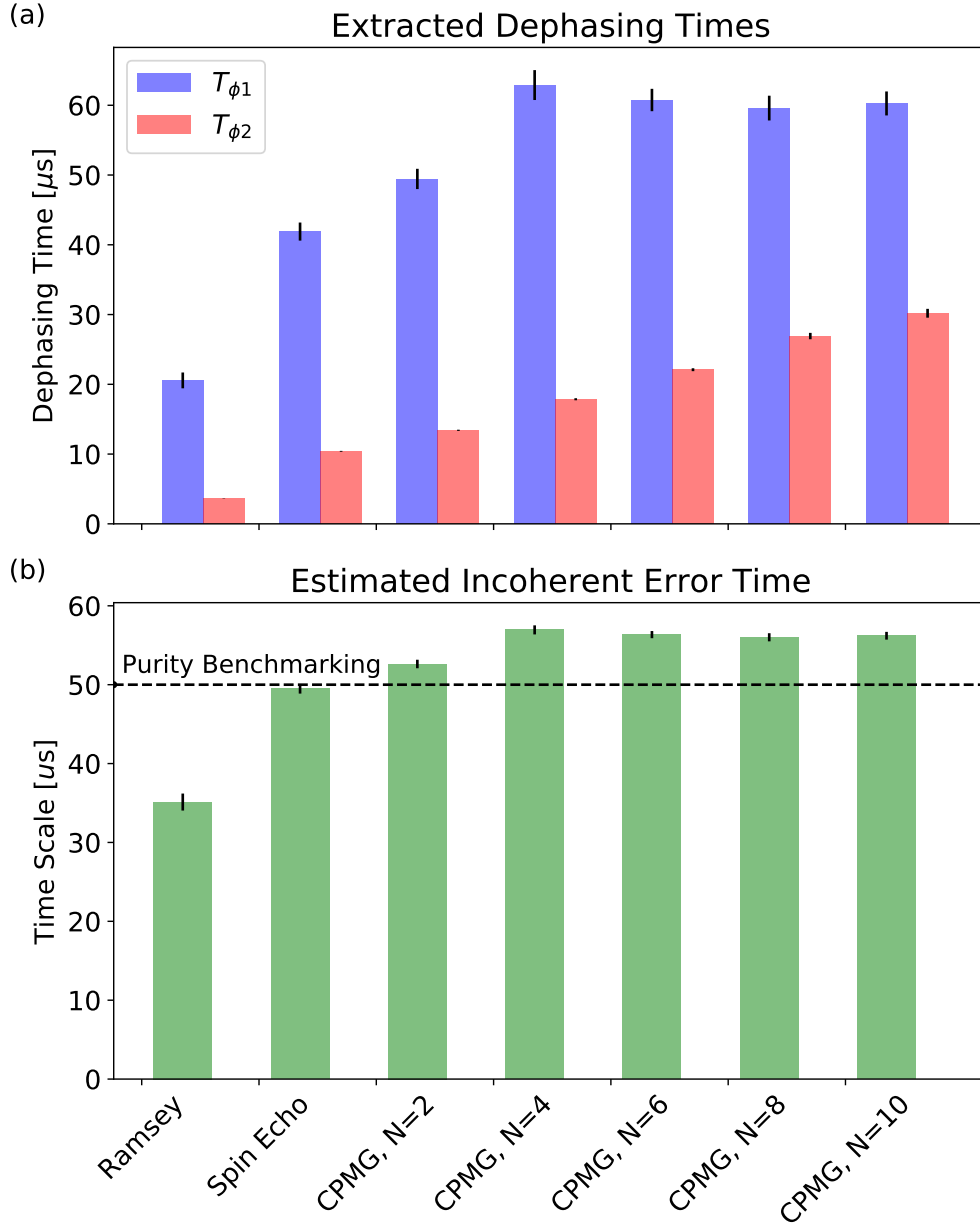


Figure 6.6: (a) Extracted white noise (T_{ϕ_1}) and correlated noise (T_{ϕ_2}) dephasing times for the data in Fig. 6.5. Note that we have subtracted the effect of T_1 to obtain a pure dephasing time for T_{ϕ_1} . (b) Estimated time scale for incoherent errors in randomized benchmarking based on decoherence times. The estimate includes the effect of T_{ϕ_1} and T_1 , but not T_{ϕ_2} . The dashed black line indicates the time scale measured using purity benchmarking, and agrees with the results of spin echo.

where A and B are again parameters related to state preparation and measurement, and T_{exp} is defined as

$$\frac{1}{T_{\text{exp}}} = \frac{1}{T_{\phi_1}} + \frac{1}{2T_1}. \quad (6.10)$$

The results of fitting Eqn.6.9 are shown in Fig. 6.6(a), where we see a few general trends. For the white noise component, T_{ϕ_1} increases slightly with more echo pulses, but saturates to a value of $60 \mu\text{s}$. On the other hand, T_{ϕ_2} rapidly increases with the number of echo pulses due to suppression of the $1/f$ part of the spectrum, and would presumably keep increasing if we added more echoes.

Armed with our decoherence measurements, we can now try to make sense of our results from purity benchmarking. Reference [120] gives us a prescription for converting decoherence times to benchmarking error rates:

$$\epsilon_{T_1} = \frac{1}{3} \frac{t_{\text{gate}}}{T_1} \quad (6.11)$$

$$\epsilon_{T_{\phi_1}} = \frac{1}{3} \frac{t_{\text{gate}}}{T_{\phi_1}} \quad (6.12)$$

$$\epsilon_{T_{\phi_2}} = \frac{1}{3} \left(\frac{t_{\text{gate}}}{T_{\phi_2}} \right)^2, \quad (6.13)$$

where the factors of 3 come from considering how each of six axial Bloch states are affected by the various forms of decoherence. The linearity of the data in Fig. 6.3 indicates that the effect of T_{ϕ_2} is negligible.⁵ Therefore, we estimate the expected timescale for

⁵For $t_{\text{gate}} = 100 \text{ ns}$ and $T_{\phi_2} = 3.6 \mu\text{s}$ as measured by the Ramsey experiment, $\epsilon_{T_{\phi_2}} = 2.5 \times 10^{-4}$. While this error value is small, it would have been resolved in the purity benchmarking. The absence of any quadratic trend in the data indicates that some of the correlated noise is likely echoed during randomized benchmarking.

benchmarking errors to be

$$\frac{1}{T_{\text{error}}} = \frac{1}{3} \left(\frac{1}{T_1} + \frac{1}{T_{\phi_1}} \right) \quad (6.14)$$

In Fig. 6.6(b), we plot T_{error} as estimated from the dephasing times extracted for each measurement sequence. We find that T_{error} estimated from the spin echo experiment agrees well with the time scale for errors measured using purity benchmarking.

6.6 Conclusion

In this chapter, we reviewed the technique of randomized benchmarking (RB), then used RB to measure the fidelity of our gates. We found that the average gate fidelity for a 14 ns gate was $3.5 \pm 0.1 \times 10^{-4}$, and by using purity benchmarking, we determined that 85% of the infidelity was due to decoherence. Our calibration procedure continued to produce high fidelity gates over a range of gate times ranging from 10 ns to 100 ns. Finally, we found that the time scale of randomized benchmarking errors is accurately explained by measurements of energy relaxation and phase decoherence. In particular, the dominant source of error for our gates is energy relaxation.

The single qubit gate fidelities we have measured compare favorably to other results in superconducting qubits [121, 122]. Increasing T_1 appears to be the most straightforward path to improving our gate fidelities. However, the source of white noise dephasing is still an outstanding issue in the UCSB Xmons. Possible culprits include electronics noise on the bias line, resonator shot noise, thermal effects, or fluctuations in two-level system defects. Future experiments in benchmarking should focus on probing the noise spectrum

at different qubit frequencies and across different qubits to look for possible trends in the noise.

6.7 Contributions

Randomized benchmarking of single qubit gates was first performed in our group by Julian Kelly and Rami Barends. Theory support for purity benchmarking was provided by Steve Flammia and Chris Granade. I acquired all of the data presented in this chapter, and also performed all of the analysis. The device was fabricated by the Google Quantum Hardware group.

Chapter 7

Measuring and Calibrating for Two State Errors

In this chapter, we will apply the technique of randomized benchmarking to a specific problem: minimizing leakage and phase errors when operating gates on a superconducting transmon qubits.

7.1 Gate Errors Due to Non-Computational Energy Levels

As we saw in Chapter 2, the transmon has a weakly anharmonic potential which supports a ladder of energy levels, as shown in Fig.7.1. The two lowest levels form our qubit, and the primary non-computational level is the $|2\rangle$ state. In the rotating frame, the




Figure 7.1: Weakly anharmonic potential of a transmon. When driving $|0\rangle$ to $|1\rangle$, direct excitation to $|2\rangle$ (red arrow) causes leakage errors, while AC Stark repulsion of the $|1\rangle \leftrightarrow |2\rangle$ transition (dashed lines) leads to phase errors.

Hamiltonian of the lowest three levels while undergoing a microwave drive resonant with the 0-1 transition is approximately

$$H(\eta, \Omega) = \begin{pmatrix} 0 & \Omega & 0 \\ \Omega & 0 & \sqrt{2}\Omega \\ 0 & \sqrt{2}\Omega & \eta \end{pmatrix} \quad (7.1)$$

where $\eta = f_{21} - f_{10}$ is the anharmonicity and Ω is the Rabi frequency, which characterizes the strength of the microwave drive. We see that when driving the 0-1 transition, we also off-resonantly drive the 1-2 transition, which leads to two different types of errors. First, leakage errors occur when the qubit state is directly excited from the $|1\rangle$ state to the $|2\rangle$ state. Note that leakage can occur even if the qubit begins in the $|0\rangle$ state, for example when driving a π pulse from $|0\rangle$ to $|1\rangle$, because the qubit will occupy $|1\rangle$ for a significant fraction of the trajectory.

The second type of error is a phase error due to AC Stark shifting of the $1 \leftrightarrow 2$ transition [123]. To understand the AC Stark shift, we truncate the Hamiltonian to the 1-2 subspace

$$H_{12}(\eta, \Omega) = \begin{pmatrix} 0 & \sqrt{2}\Omega \\ \sqrt{2}\Omega & \eta \end{pmatrix} \quad (7.2)$$

then diagonalize and compute the eigenenergies in the limit $\Omega \ll \eta$

$$E_{\pm} = \frac{\eta \pm \sqrt{\eta^2 + 8\Omega}}{2} \quad (7.3)$$

$$E_+ \approx \eta + \frac{2\Omega^2}{\eta} \quad (7.4)$$

$$E_- \approx -\frac{2\Omega^2}{\eta} \quad (7.5)$$

$$\Delta_{12} = E_+ - E_- \approx \frac{4\Omega^2}{\eta} \quad (7.6)$$

We find that in the presence of the microwave drive, the $|1\rangle$ and $|2\rangle$ energy levels repel, with the repulsion proportional to the square of the drive strength. As a consequence, the f_{10} transition frequency changes as we drive the transmon, and the $|0\rangle$ and $|1\rangle$ states accumulate a relative phase at a rate proportional to the drive strength. In the Bloch sphere picture, the AC Stark shift means that when attempting to drive microwave rotations about an equator axis, the actual rotation axis will tilt away from the equator towards one of the poles, giving a rotation about the Z axis and phase accumulation.

Previous experimental work [124, 125] on superconducting qubits has focused on reducing phase errors, because they were the dominant source of total gate infidelity. These experiments used a technique called Derivative Reduction by Adiabatic Gate (DRAG) pulse shaping [100] to compensate for the effect of the AC Stark shift, which helped push single qubit gate fidelity in superconducting qubits to over 99.9% [126, 64]. However, gate fidelity is not the only metric that determines the viability of fault tolerant quantum error correction, because certain errors are more deleterious than others. Specifically,

leakage errors are highly detrimental for error correcting codes such as the surface code, because interactions with a qubit in a leakage state have a randomizing effect on the interacting qubits [127]. Moreover, leakage states can be as long-lived as the qubit states, leading to time-correlated errors which further degrade performance [128]. These concepts were recently demonstrated in a 9 qubit repetition code [67], where single leakage events persisted for multiple error detection cycles and propagated errors to neighboring qubits. Understanding and reducing leakage is of critical importance for realizing an error corrected quantum processor.

In this chapter, we will show that previous experimental implementations of the DRAG protocol have a tradeoff between leakage and phase errors. To overcome this tradeoff, we will apply an additional pulse shaping technique and obtain single qubit gates that have both low leakage and low phase error.

7.2 Test Device

For our testbed, we use a single Xmon transmon qubit [79, 63] (Q_7) from the 9 qubit chain described in Ref. [67]. We operate the qubit at a frequency f_{10} of 5.3 GHz, and the anharmonicity $\eta = \omega_{21} - \omega_{10}$ is $2\pi \times -212$ MHz. The T_1 of the device at the operating frequency is $22 \mu\text{s}$, while a Ramsey experiment shows two characteristic decay times [120], an exponential decay time T_{ϕ_1} of $8 \mu\text{s}$ and a Gaussian decay time T_{ϕ_2} of $1.8 \mu\text{s}$. We note that as an alternative, leakage can be suppressed by engineering qubits with larger anharmonicities such as flux qubits, which have recently also achieved high coherence

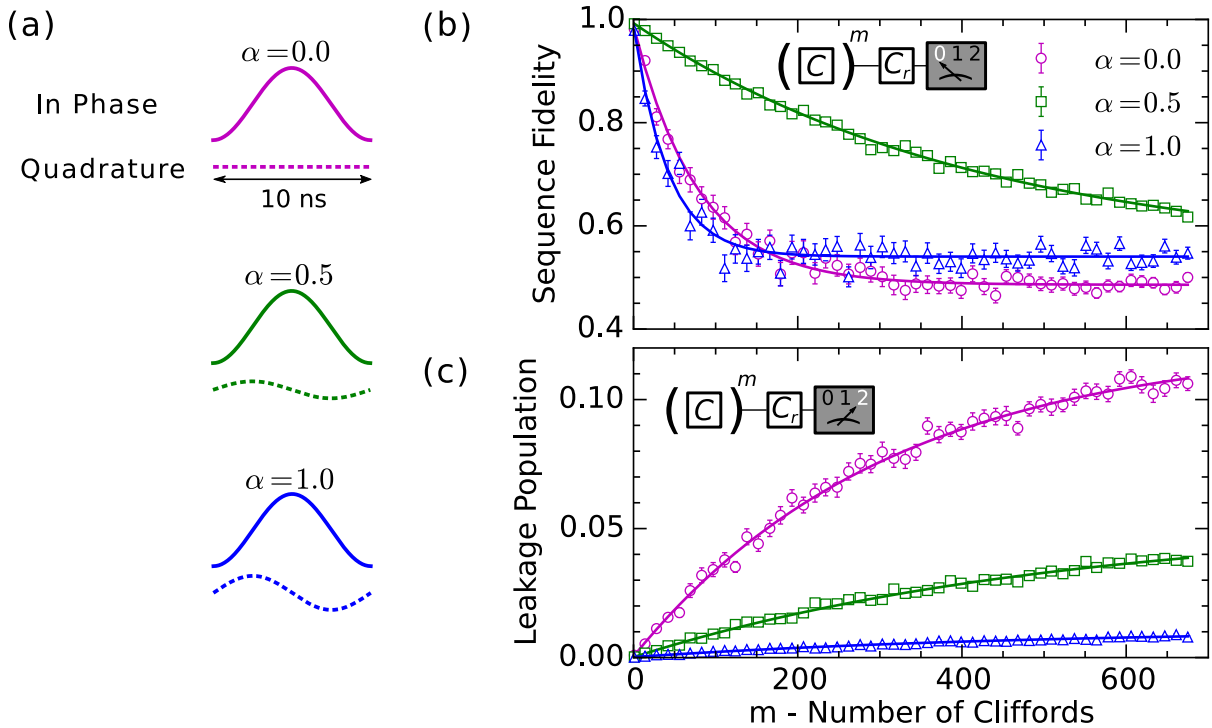


Figure 7.2: (a) Simple DRAG correction, which adds the derivative of the envelope to the quadrature component of the envelope. Three different DRAG weightings (α) are shown. (b) Exponential decay of sequence fidelity from randomized benchmarking, shown for the three values of α . Each point is the average of 75 different random sequences. Fidelity is highest for $\alpha = 0.5$ (c) $|2\rangle$ state population *vs* sequence length, showing accumulation of leakage with sequence length. Leakage is lowest for $\alpha = 1.0$.

[129].

7.3 Measuring Total Gate Error and Leakage Error

In order to measure the effect of leakage and phase errors, we will be using the Clifford based randomized benchmarking discussed in Chapter 6. We begin by measuring how the DRAG protocol suppresses leakage and phase errors. We use the simplified version of DRAG described in Refs. [124, 125]. Given a control envelope $\Omega(t)$, we add the time

derivative $\dot{\Omega}(t)$ to the quadrature component:

$$\Omega'(t) = \Omega(t) - i\alpha \frac{\dot{\Omega}(t)}{\Delta} \quad (7.7)$$

where α is a weighting parameter. Fourier analysis [123, 130] shows that the DRAG correction suppresses the spectral weight of the control pulse at the $|1\rangle \leftrightarrow |2\rangle$ transition if $\alpha = 1.0$, thus minimizing leakage errors. However, the optimal value to compensate the AC Stark shift and correct for phase errors is $\alpha = 0.5$ [123, 124].

We confirm these concepts by performing Clifford based RB using 10 ns microwave pulses shaped with three different values of α (0, 0.5, and 1.0), as shown in Fig. 7.2(a). Figure 7.2(b) shows sequence fidelity decay curves for the three values of α . As expected, using $\alpha = 0.5$ yields higher fidelities than $\alpha = 0.0$ or $\alpha = 1.0$. Using the sequence fidelity decay curves, we fit for the average gate fidelity using Eqn.6.1. For $\alpha = 0.5$, we obtain an error per Clifford of $9.6 \pm 0.1 \times 10^{-4}$, while for $\alpha = 0.0$ and $\alpha = 1.0$ we obtain errors of $6.3 \pm 0.2 \times 10^{-3}$ and $1.20 \pm 0.01 \times 10^{-2}$ per Clifford, respectively.

Simultaneously, we characterize leakage errors in our gateset from the dynamics of the $|2\rangle$ state measured while performing RB, as shown in Fig. 7.2(c). For all three values of α , the $|2\rangle$ state population shows an exponential approach to a saturation population. Without correction, this saturation population is significant at about 10%, but decreases by a factor of 3 for $\alpha = 0.5$ and by a factor of 10 for $\alpha = 1.0$. To quantify the leakage rate per Clifford, we fit the $|2\rangle$ state dynamics to a simple rate equation that takes into account leakage from the computational subspace into the $|2\rangle$ state and decay from $|2\rangle$

back into the subspace.

$$p_{|2\rangle}(m) = p_{\infty} (1 - e^{-\Gamma m}) + p_0 e^{-\Gamma m} \quad (7.8)$$

$$\Gamma = \gamma_{\uparrow} + \gamma_{\downarrow} \quad p_{\infty} = \gamma_{\uparrow}/\Gamma \quad (7.9)$$

where $p_{|2\rangle}(m)$ is the $|2\rangle$ state population as a function of sequence length m , γ_{\uparrow} and γ_{\downarrow} are the leakage and decay rates per Clifford, and p_0 is the initial $|2\rangle$ state population. Using Eqn. 7.9, we extract leakage rates of $3.92 \pm 0.08 \times 10^{-4}$, $1.02 \pm 0.02 \times 10^{-4}$, and $2.18 \pm 0.08 \times 10^{-5}$ for $\alpha = 0, 0.5$ and 1.0 . The results from RB confirm the theory behind simple DRAG: we can minimize either phase error or leakage error, but not both.

7.4 Pulse Detuning to Minimize Phase Error

To simultaneously optimize for both errors, we would like to minimize leakage using simple DRAG, then separately compensate the AC Stark shift. In the original DRAG theory, the Stark shift was compensated using a time dependent detuning of the qubit [100]. As noted in Refs. [123, 131, 132], a constant detuning should also be able to compensate the shift ¹. Given an envelope Ω' , which may have a quadrature correction, we generate a new envelope

$$\Omega''(t) = \Omega'(t)e^{2\pi i \delta f t} \quad (7.10)$$

¹A unitary rotation in a two-level system is parameterized by three values, so by setting amplitude, detuning, and phase we can in principle generate any rotation and reduce phase errors to zero [123], with the only error being leakage.

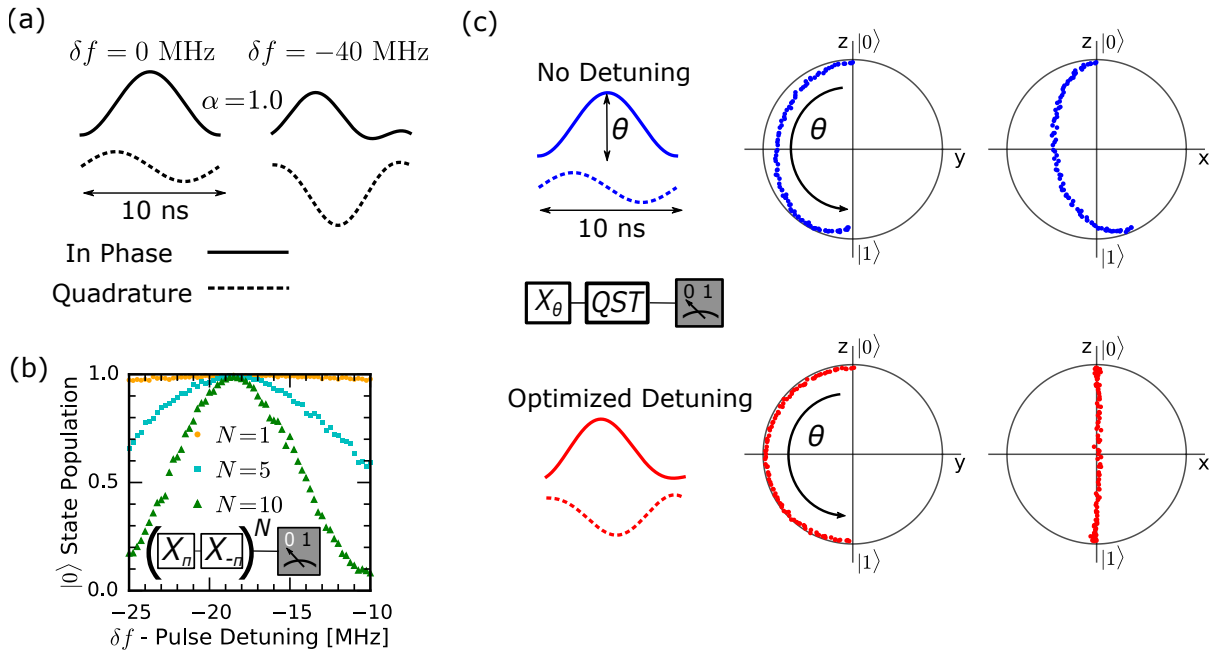


Figure 7.3: (a) Control envelopes with simple DRAG with (right) and without (left) detuning of the pulse. The detuning is exaggerated for illustration. (b) We sweep over the detuning δf while performing the pseudo-identity sequence shown in the inset. The sequence maps back to $|0\rangle$ when detuning is optimized. Repeating the sequence increases the sensitivity of the measurement. (c) Quantum state trajectories plotted on projections of the Bloch sphere, with (bottom) and without (top) optimal detuning. The data is obtained by performing quantum state tomography (QST) after applying a variable X rotation, with the rotation angle ranging from 0 to π .




Figure 7.4: (a) Dependence of the optimal π pulse detuning on α , the DRAG weighting. Three different pulse lengths are shown. The dashed lines are linear fits. (b) The slopes from the linear fits as shown in (a), for a range of pulse lengths. The dashed line is a fit to the inverse square of the pulse length, as expected from the AC Stark shift.

where δf is the detuning of the pulse from the qubit frequency. We also redefine the anharmonicity parameter in Eqn. 7.7 to be $\eta = \omega_{21} - (\omega_{10} + 2\pi\delta f)$, so that leakage suppression still occurs at the $1 \leftrightarrow 2$ frequency. An example of a detuned pulse is shown in Fig. 7.3(a).

To optimize the detuning parameter δf , we sweep the detuning of a π -pulse while performing the pseudo-identity operation of a π -pulse followed by a $-\pi$ -pulse along the same rotation axis [124, 133]. As shown in Fig. 7.3(b), the detuning is optimized when the $|0\rangle$ state population is maximized, and the pseudo-identity can be repeated to increase the resolution of the measurement. To verify that the detuning has suppressed phase errors, we perform quantum state tomography after applying a control pulse to our qubit while ramping the amplitude of the pulse, as shown in Fig. 7.3(c). The Bloch vector only reaches the opposite pole when the detuning is optimized.

To better understand the parameter space for optimizing our pulses, we repeat the

experiment shown in Fig. 7.3 for different π pulse lengths and values of the DRAG weight, α . In Fig. 7.4(a) we show the dependence of the optimal pulse detuning on α for three different lengths. For each pulse length, the pulse detuning is linear with α , and the slope becomes more shallow with longer pulse length. In Fig. 7.4(b), we plot the dependence of this slope on pulse length. We find that the slope between optimal detuning and DRAG is proportional to the inverse square of the pulse length. Equivalently, the slope depends quadratically on the microwave drive strength, which matches the intuition we developed in 7.6, where we found that the AC Stark shift scales quadratically with the microwave drive strength.

7.5 Optimized Microwave Pulses

We now explore in more detail the dependence of fidelity and leakage on α . In Fig. 7.5, we show parameters extracted from RB with 10 ns pulses while varying α between 0.0 and 1.5. Without detuning the pulses, we find the minimum error per Clifford to be $7.9 \pm 3 \times 10^{-4}$ when $\alpha = 0.4$. We note that the optimal α deviates slightly from the expected optimal value of $\alpha = 0.5$, and the actual optimal value can vary between 0 and 1 for different qubits and operating frequencies. We attribute this deviation to distortions or reflections of the pulse between the waveform generator and the qubit [133]. Away from the optimal α , the error increases rapidly.

Next, we optimize the detuning of the pulses for each value of α using the method described in Fig. 7.3. We find that for π and $\pi/2$ pulses with the same length, the

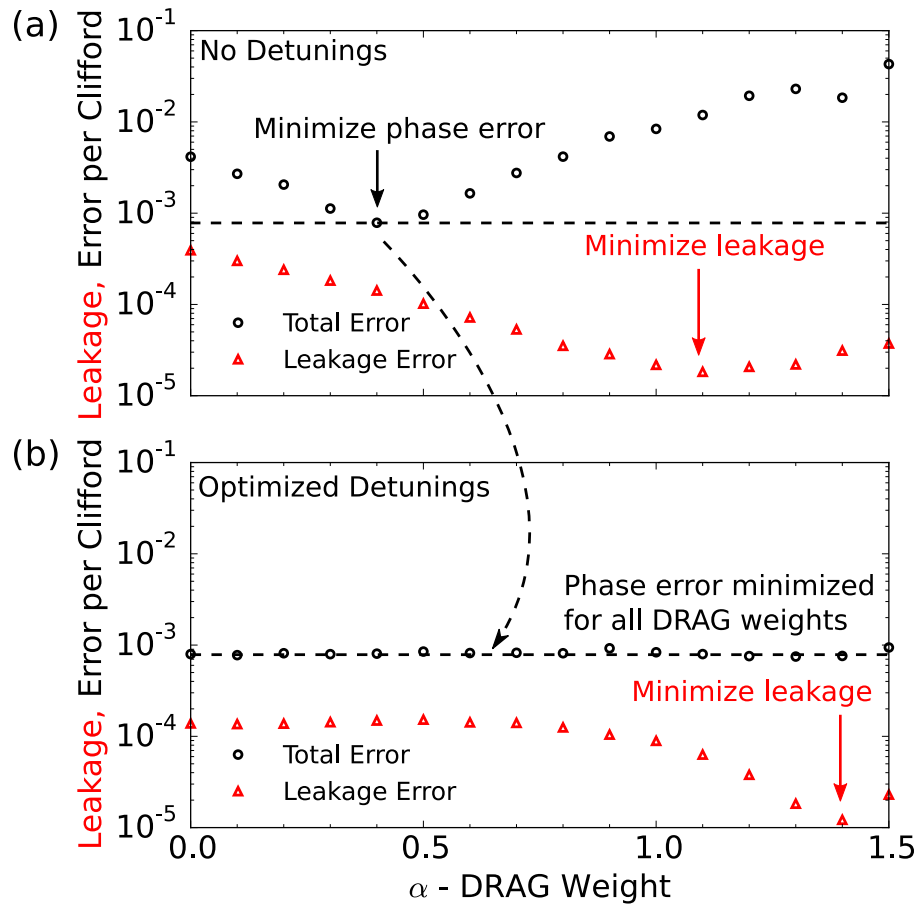


Figure 7.5: Total gate fidelity and leakage rates versus DRAG weighting α , measured using RB. (a) Without using pulse detunings, we require different values of α to minimize overall error versus leakage errors. (b) By optimizing our pulses using detunings, we obtain high fidelity for any α , and are free to choose α to minimize leakage.

same detuning for both pulses yields the best results. After calibrating the detuning, we recalibrate the pulse amplitudes, then run a short Nelder-Mead optimization on the RB fidelity to get final adjustments to pulse parameters [134]. With these optimizations, we find that the average error per Clifford for all values of α to be 9.1×10^{-4} , with a standard deviation of 1×10^{-4} . In other words, we can tune up high fidelity gates for *any* value of α .

With gate fidelity now independent of α , we are free to implement DRAG solely to minimize leakage. Without detuning, the minimum leakage rate is $1.82 \pm 0.07 \times 10^{-5}$ for $\alpha = 1.1$. After detuning the pulses for optimal fidelity, we see shifts in the leakage rates. For $\alpha > 0.4$, we detune the pulses towards the $|1\rangle \leftrightarrow |2\rangle$ transition which tends to increase the leakage rate. Nevertheless, we can still suppress leakage to the same level as the undetuned pulses by increasing α to 1.4. Using these parameters, we achieve both high fidelity ($8.7 \pm 0.4 \times 10^{-4}$ error per Clifford) and low leakage ($1.2 \pm 0.1 \times 10^{-5}$).

7.6 Leakage vs Pulse Length

Having characterized 10 ns pulses in detail, we now examine the dependence of leakage on pulse length. As noted previously, pulse detuning can affect the leakage rate; for simplicity we set the detuning to zero for the following measurements. We initially set $\alpha = 0.0$ and measure the leakage rate while varying the length of our pulses between 8 ns and 50 ns and calibrating the pulse amplitudes accordingly. We then repeat this measurement with $\alpha = 1.1$ where we previously minimized leakage in Fig. 3(a). The results are shown

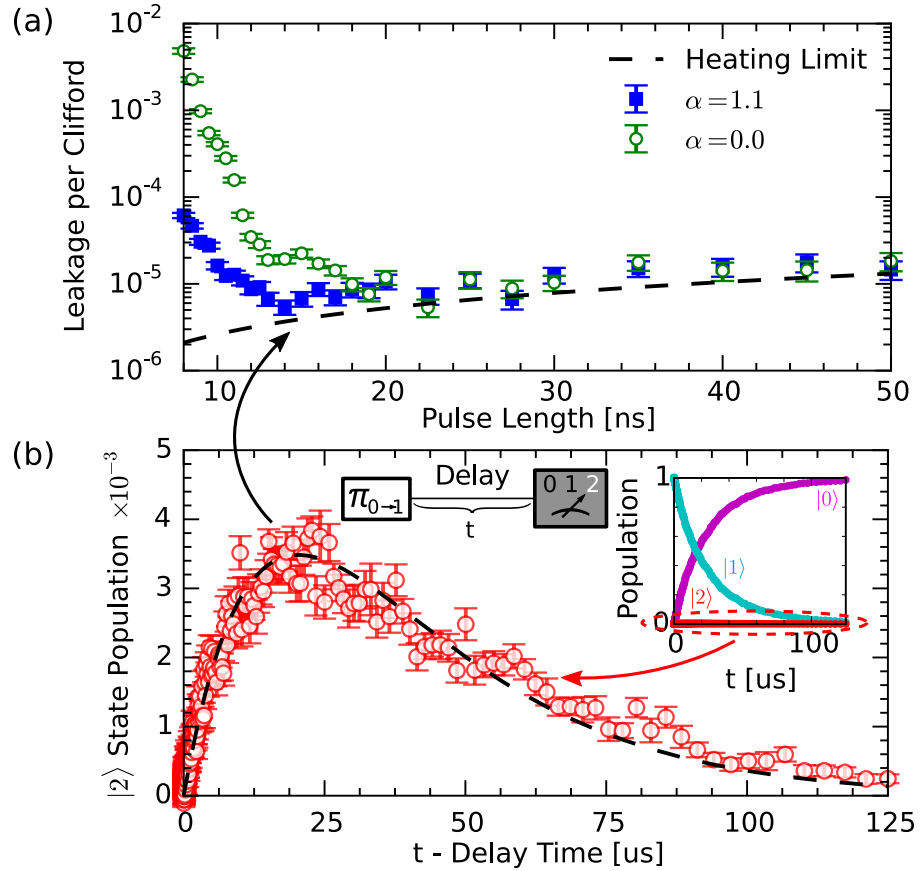


Figure 7.6: (a) Leakage rate per Clifford extracted from RB versus pulse length, with $\alpha = 0.0$ and $\alpha = 1.1$. The dashed line is the lower bound on leakage calculated from the heating rate. (b) Heating of the qubit from $|1\rangle$ to $|2\rangle$. We prepare the qubit in $|1\rangle$, wait for time t , then measure the qubit state. Inset: Dynamics of all three states, primarily showing T_1 decay of $|1\rangle$ to $|0\rangle$. Main figure: Zoom in of the $|2\rangle$ state dynamics, showing an increase in population due to heating before relaxing back to zero. The data has been corrected for readout visibility. The dashed line is a rate equation fit, from which we extract the heating rate plotted in (a).

in Fig. 7.6(a). For short pulses, we observe that the leakage rate decreases exponentially with increasing pulse length, and that DRAG correction suppresses leakage by an order of magnitude or more. However, as the pulse length increases past 15 ns, the leakage rate begins to level off and even begins to increase. Furthermore, DRAG no longer seems to suppress leakage for pulses longer than 20 ns. These results suggest that for long pulses, leakage is due to incoherent processes such as thermal excitations or noise at the $|1\rangle \leftrightarrow |2\rangle$ transition, rather than coherent control errors.

To measure the incoherent leakage rate, we prepare the qubit in the $|1\rangle$ state and measure the qubit's dynamics, as shown in Fig. 7.6(b). We see that the $|2\rangle$ state population initially rises over $20\ \mu\text{s}$, corresponding to heating from $|1\rangle$ to $|2\rangle$. Then, the $|2\rangle$ population slowly decays to zero as both excited states relax due to T_1 processes. We model the $|2\rangle$ population using a rate equation with three rates: decay from $|2\rangle$ to $|1\rangle$, decay from $|1\rangle$ to $|0\rangle$, and heating from $|1\rangle$ to $|2\rangle$. We ignore nonsequential transitions since they are suppressed in the nearly harmonic transmon potential [135], as well as heating from $|0\rangle$ to $|1\rangle$ since we assume the initial state is $|1\rangle$. We extract the two decay rates from T_1 measurements, which give $T_1^{(1)} = 22\ \mu\text{s}$ and $T_1^{(2)} = 18\ \mu\text{s}$ ². We then fit the $|1\rangle \rightarrow |2\rangle$ heating rate to be $1/(2.2\ \text{ms})$.

We convert this heating rate to a leakage rate per Clifford using the prescription in Ref. [120]. The resulting lower bound on leakage is shown in the dashed line in Fig. 7.6(a).

For pulses longer than 15 ns, we find that the leakage rate is within a factor of 2 of this

² $T_1^{(1)}$ is not $2T_1^{(2)}$ as might be expected because the transitions are at different frequencies, and T_1 in our system is frequency dependent [79]

lower bound, confirming that even at these relatively short timescales, we are being limited by incoherent processes. We note that the heating rate and T_1 decay rate are consistent with an equilibrium population of 0.8% for the $|1\rangle$ state. In other works, equilibrium populations closer to 0.1% have been achieved [136], suggesting that incoherent leakage can be reduced through improved thermalization.

7.7 Conclusions

In this chapter, we used single qubit randomized benchmarking to study leakage errors in a superconducting qubit. We showed that simple DRAG correction alone cannot minimize leakage and total gate error simultaneously, but by detuning our pulses, we obtain gates with both high fidelity and low leakage. We also measured the dependence of leakage on pulse length, and found that heating of the qubit is a significant source of leakage in our system. Currently, leakage is not a significant part of the total gate error for single qubit gates. However, as qubit coherence and gate fidelities continue to be improved, leakage will become a more significant component of gate errors. To further suppress leakage, qubit thermalization will need to be improved to mitigate incoherent leakage, and further pulse correction techniques such as higher-order DRAG [131] or Fourier corrections [123] will be needed to suppress coherent leakage.

7.8 Contributions

This chapter is an adaptation of Ref. [137], for which I acquired the data and wrote the text. The devices used in this experiment were fabricated by Chris Quintana and Julian Kelly. Julian also contributed the initial idea of using randomized benchmarking to measure leakage, while Alexander Korkov provided theoretical support.

Chapter 8

Two Qubit Calibration and Benchmarking

In previous chapters, we have detailed the calibration and benchmarking of high fidelity single qubit gates. We now turn our attention to our two qubit entangling gate, the control-Z (CZ). We will first review the theoretical underpinnings behind the control Z between two capacitively coupled transmons, which takes advantage of frequency control and the higher levels of the transmon. We will also discuss a parameterization of the control trajectory which will enable fast, high fidelity control Z gates. With the theoretical background established, we will turn to experimentally calibrating a high fidelity control Z gate on two qubits. Finally, we will benchmark this control Z gate using simple extensions of the techniques that we established for single qubit gates. We note that the CZ gate described in this chapter was first described in Ref. [138], and has previously

been experimentally demonstrated in Refs. [139, 140, 64].

8.1 The Controlled Phase Interaction in Coupled Transmons

The desired unitary for a control-Z in the computational basis is

$$U_{CZ} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, \quad (8.1)$$

where the two qubit states are ordered as $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$. The name "control-Z" implies, by analogy to the "control-NOT", that one qubit (the "target") should acquire a π phase shift conditional on the other qubit (the "control") occupying $|1\rangle$. However, Eqn. 8.1 implies that the unitary is symmetric with respect to qubit order, i.e. neither qubit is the control or target. Instead, we want to produce an interaction which occurs if the joint state is $|11\rangle$.

In Chapter 2, we saw that the Hamiltonian for two capacitively coupled transmons in the rotating frame of one of the qubits is

$$H/\hbar = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & g & 0 \\ 0 & g & \delta & 0 \\ 0 & 0 & 0 & \Delta \end{pmatrix}, \quad (8.2)$$

where g is the coupling strength, and Δ is the detuning between the qubits. Using the same ordering as in Eqn. 8.1, we will assume for the remainder of this chapter that the ket $|10\rangle$ corresponds to the *higher* qubit in frequency space occupying the $|1\rangle$, and that we are in the rotating frame of the lower qubit. This Hamiltonian offers a swap between $|01\rangle$ and $|10\rangle$ if $\delta = 0$, but not our desired $|11\rangle$ interaction.

However, to derive Eqn. 8.2, we truncated the total Hamiltonian down to two levels for each qubit. We now add back the third levels of each qubit but ignore any joint states which contain more than three excitations, and obtain the following 6 x 6 Hamiltonian

$$H/\hbar = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & g & 0 & 0 & 0 \\ 0 & g & \Delta & 0 & 0 & 0 \\ 0 & 0 & 0 & \Delta & g_2 & g_2 \\ 0 & 0 & 0 & g_2 & \eta & 0 \\ 0 & 0 & 0 & g_2 & 0 & 2\Delta + \eta \end{pmatrix}, \quad (8.3)$$

where $\eta = \omega_{10} - \omega_{12}$ is the anharmonicity which is assumed to be the same for both qubits, $g_2 \approx \sqrt{2}g$ is the $|1\rangle \rightarrow |2\rangle$ charge coupling, and the states are ordered as $|00\rangle$, $|01\rangle$, $|10\rangle$, $|11\rangle$, $|02\rangle$, and $|20\rangle$. We see now that in the two excitation manifold, $|11\rangle$ couples to $|02\rangle$ and $|20\rangle$. As Δ approaches $-\eta$, the repulsion and mixing between the $|11\rangle$ and $|20\rangle$ states increases and leads to an avoided level crossing when $\Delta = -\eta$, as shown in Fig 8.1. Slightly away from this crossing point, $|11\rangle$ still experiences a shift in




Figure 8.1: Energy levels for $|11\rangle$ and $|20\rangle$, as a function of the detuning between the two qubits. In the absence of any coupling between the two qubits, the energy levels follow the dashed lines. With coupling on, as the detuning approaches the anharmonicity η , the two levels show an avoided level crossing, with the splitting given by the coupling strength g_2 . The difference between the blue curve and the dashed blue line gives a frequency difference conditional on both qubits occupying $|11\rangle$.

its eigenenergy due to the coupling to $|20\rangle$. This shift gives us a path to implementing a control-Z, since a frequency shift in $|11\rangle$ will impart a phase shift over time, conditional on both qubits being in $|1\rangle$, which is precisely the desired action of a control-Z.

Implementing the control-Z comes with a few caveats. First, since we want to apply control-Z gates at discrete times, we must be able to turn the interaction off when desired. In a fixed coupling system, turning off the control-Z can only be achieved by detuning the qubits so that $\Delta \gg \eta$. In this regime, we can approximate the frequency shift on $|11\rangle$ as due to independent dispersive repulsions from $|02\rangle$ and $|20\rangle$ which we sum together to obtain [120]

$$\omega_{zz} = \frac{g_2^2}{\Delta - \eta} - \frac{g_2^2}{\Delta + \eta} \quad (8.4)$$

$$= \frac{-2\eta g_2^2}{\Delta^2 - \eta^2}. \quad (8.5)$$

In a standard UCSB Xmon where $g = 15$ MHz and $\eta = -220$ MHz, to obtain an $|11\rangle$ shift less than 0.2 MHz during normal operation, we require $\Delta > 1$ GHz. To execute the control Z, we must then be able to accurately tune our qubits in and out of the region where $|11\rangle$ and $|02\rangle$ strongly interact.

Second, by deliberately using interactions with the third level of the qubit, we expose ourselves to far more leakage than in the single qubit case. As we discussed in Chapter 6, minimizing leakage is crucial for eventually performing quantum error correction. Ideally, in moving $|11\rangle$ near $|02\rangle$ and back, we would like to do so adiabatically to ensure that we remain the $|11\rangle$ state at the end of the trajectory. Thus, the fundamental challenge

of executing a control Z is as follows: starting from an idling position where two qubits are far detuned, we must bring the $|0\rangle \leftrightarrow |1\rangle$ transition of the lower qubit near resonance with the $|1\rangle \leftrightarrow |2\rangle$ transition of the higher qubit, then back to the idling position to continue performing single qubit gates. We want this trajectory to be fast to mitigate the effects of decoherence, but do not want the $|11\rangle$ state to leak into the $|02\rangle$ state.

A parameterization for obtaining such a trajectory is given in Ref. [123]. In brief, the authors suggest parameterizing the time trajectory of the control Hamiltonian using Fourier parameters

$$\theta - \theta_{\text{initial}} = \sum_{n=1}^k \lambda_n [1 - \cos(2\pi n t / t_p)] \quad (8.6)$$

where $\theta = \arctan(2g_2/(\Delta + \eta))$ is the *control angle* between the $|11\rangle$ and $|02\rangle$ states, and t_p is the total time of the pulse. In particular, it was found that using only the first three Fourier coefficients ($m = 3$), an experimentally achievable trajectory could be found which would give the expected unitary and leakage below the 10^{-4} level.

8.2 Calibrating Non-Idealities in Qubit Flux Control

Prior to calibrating our CZ, we must first calibrate flux bias control of the qubit frequency to high precision.


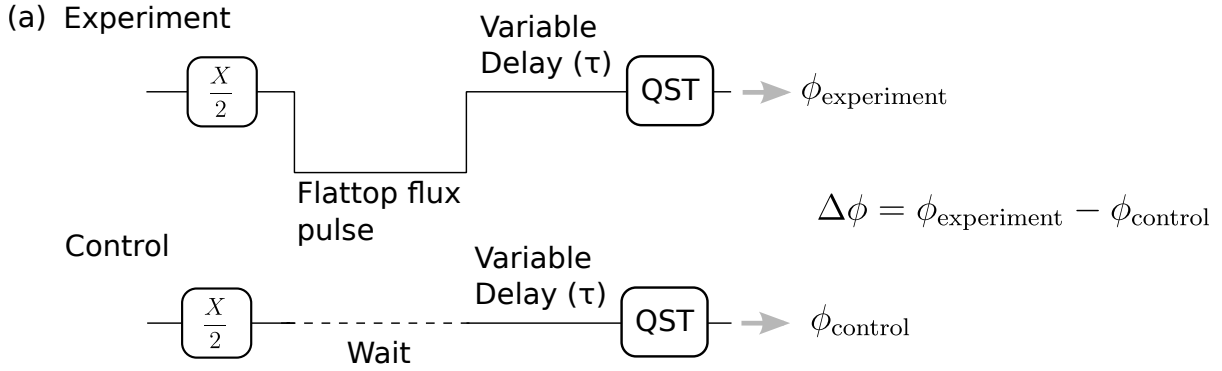


Figure 8.2: Measuring the transfer function for pulsed flux control. (a) We bring the qubit into a superposition, pulse the flux bias of the qubit, then measure the phase of the qubit using quantum state tomography (QST) as a function of time after the pulse. We also perform the measurement without pulsing the flux bias as a control to subtract out any phase drift due to a frequency offset. (b) Distortion of the pulse leads to a significant drift in the phase over time. We fit this phase drift to infer the transfer function between our AWG and the qubit. After deconvolving our pulse with this transfer function, we obtain a much flatter phase response.

8.2.1 Correcting for Flux Control Distortions

One of the main difficulties in executing a gate using flux control is distortions of the pulse shape between our AWG and our qubit. The distortions may be due to dispersion in the cables, reflections at the connections between components, or current redistribution in the sample itself [64]. The exact nature of the distortions has been difficult to determine and is subject of ongoing research. Nevertheless, we can still measure the effects of pulse distortion and attempt to correct for them.

In Fig. 8.2, we show the pulse sequence for measuring the distortion of our flux control pulses. We begin by initializing the qubit in a superposition state. Then, we apply a flattop flux pulse with a known amplitude and width, which detunes the frequency of the qubit. After the end of the flattop, we wait a variable amount of time, then perform quantum state tomography (QST) to determine the qubit state, and in particular, the phase of the superposition. The application of the flux pulse will cause the qubit to accumulate some phase. If the flattop is not distorted, this phase should remain stable after the pulse has completed. However, the uncorrected data in Fig. 8.2 shows that upon completion of the flattop, the phase of the qubit drifts a significant amount over a timescale of a few hundred nanoseconds.

From this data, we would like to infer the transfer function H that is distorting our pulse, which we define in frequency space as:

$$\tilde{X}_{\text{qubit}}(\omega) = H(\omega)\tilde{X}_{\text{in}}(\omega) \quad (8.7)$$

where \tilde{X}_{in} is the Fourier transform of our intended signal, and \tilde{X}_{qubit} is the signal actually

seen by the qubit. In order to fit a transfer function H to our data, we use the following procedure:

1. Parameterize an ansatz for the transfer function H .
2. Accurately simulate the voltage output of the AWG, then convolve it with our test transfer function.
3. Convert the convolved voltage time trace to qubit frequency using a calibrated mapping from voltage to frequency.
4. Integrate the difference between the dynamic qubit frequency and the static frequency of our rotating frame to obtain qubit phase over time.
5. Compare this simulated qubit phase drift to our experimental data by computing the mean square difference, which is our error function.
6. Minimize this error function using the parameters of our ansatz ¹.

Of these steps, the most subjective is finding a proper ansatz to capture the behavior of the system. Historically, the most common ansatz used in the UCSB group is:

$$H(\omega) = 1 + \sum_{i=1}^k \frac{i\alpha_i\omega}{i\omega + \tau_i} \quad (8.8)$$

which describes multiple overshoots (positive α_i) or undershoots (negative α_i) which exponentially decay to the final value with time constants τ_i . This model does not have a simple physical origin, but does successfully fit the data, as shown in Fig. 8.2. Upon

¹To do so, we typically use `scipy`'s differential evolution with reasonable bounds on our ansatz parameters

obtaining H , we correct our flux pulses by performing *deconvolution*, where we convolve our desired flux pulse with the inverse of H

$$\tilde{X}_{\text{corrected}}(\omega) = H^{-1}(\omega)\tilde{X}_{\text{in}}(\omega) \quad (8.9)$$

so that ideally, the pulse seen by the qubit is equivalent to the desired input. The corrected data in Fig. 8.2 shows the results of repeating the phase response experiment after deconvolution. We see that after deconvolution, the response is much flatter, enabling us to perform the large, fast detunes necessary for the CZ.

8.2.2 Multiqubit Timing

Having corrected the flux control on a single qubit, we now correct for two non-idealities in multiqubit flux control. First, similar to the case of single qubit XY and Z control, the two Z control channels on our qubits of interest may have timing differences. To measure this timing difference, we populate one of the qubits then bring the two qubits on resonance with a short flux pulse on each qubit. We sweep the relative timing between the short flux pulses, and when they are aligned, we find maximal swapping between the two qubits, as shown in Fig 8.3. By aligning the Z channels of the two qubits, we also align the XY channels since we have previously aligned the Z and XY channels of the single qubits.

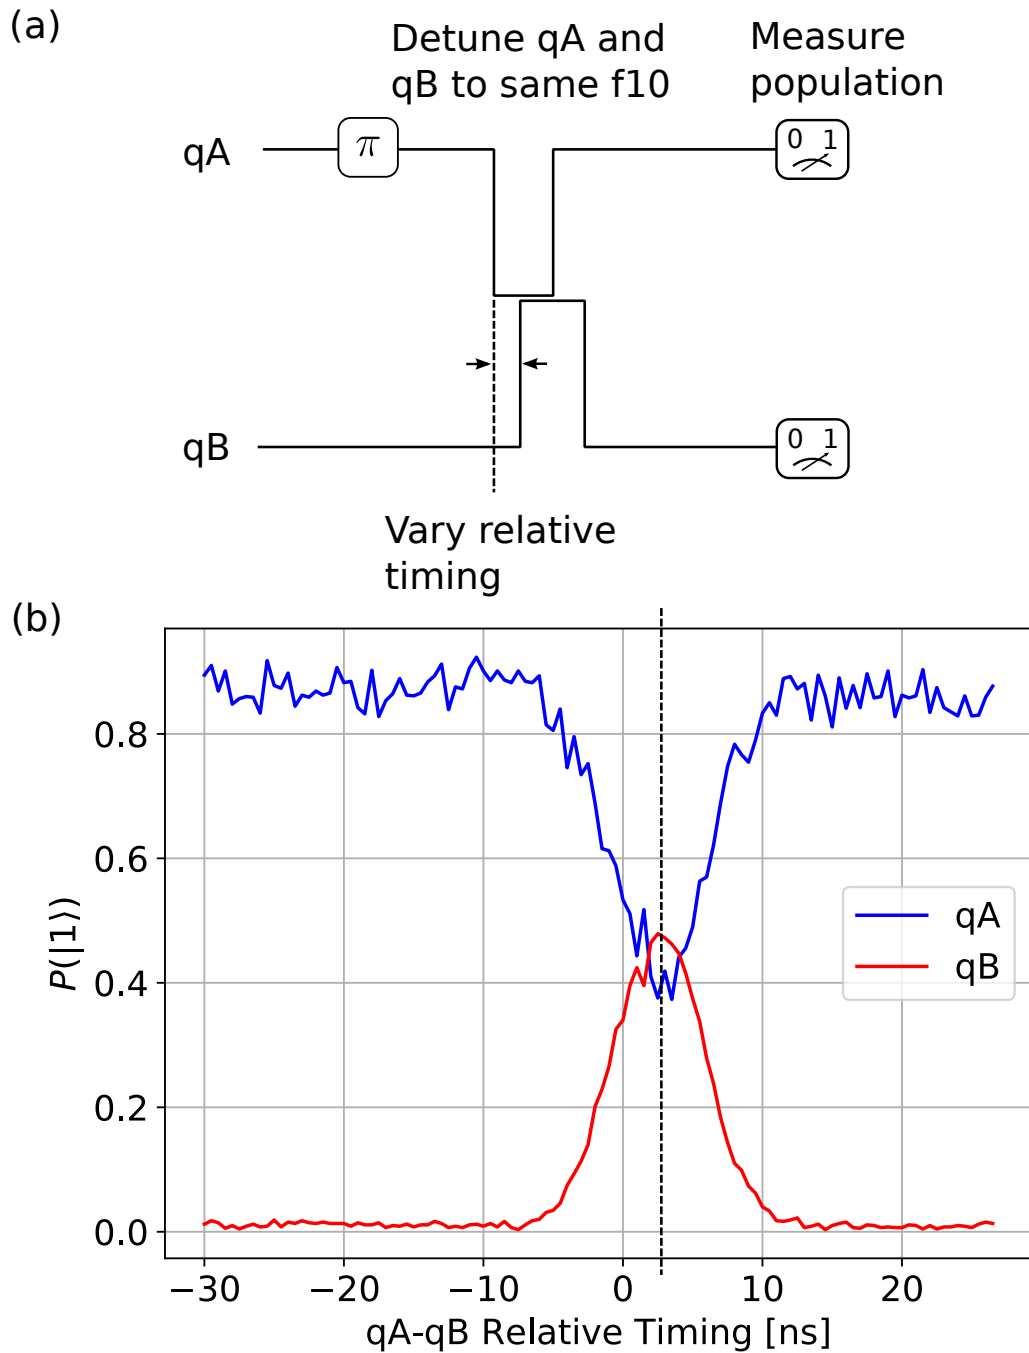


Figure 8.3: Measurement of the timing difference between the flux bias channels of two coupled qubits. (a) We populate one qubit, then bring both qubits on resonance with a short flux bias pulse, while varying the relative timing of the two flux pulses. (b) When the relative timing correctly accounts for the physical timing difference between the two flux bias channels, we see a peak in the population swapping between the qubits.




Figure 8.4: Measurement and correction of flux bias crosstalk. (a) We bring the target qubit into a superposition, then apply a fixed positive or negative flux bias to the source qubit, while simultaneously applying a compensating flux bias on the target qubit. Finally, we measure the phase accumulated by the target using quantum state tomography. (b) When no compensation is applied, the flux bias on the source imparts a phase shift on the target qubit, with the sign depending on the sign of the bias. The correct compensation amplitude is found when the phase of the target qubit is zeroed, and we average the compensation found with positive or negative source bias. (c) After applying a calibrated correction to the flux bias waveforms, the target qubit sees no difference in phase behavior for positive or negative source biases.

8.2.3 Flux Crosstalk

A second non-ideality we must correct for is crosstalk in the flux biases of the two qubits, where flux bias intentionally applied to one qubit can induce flux in other qubits. For two qubits, this crosstalk can be expressed in terms of a 2 x 2 matrix M

$$\Phi = MV \tag{8.10}$$

$$\begin{pmatrix} \phi_1 \\ \phi_2 \end{pmatrix} = \begin{pmatrix} M_1 & m_{12} \\ m_{21} & M_2 \end{pmatrix} \begin{pmatrix} V_1 \\ V_2 \end{pmatrix} \tag{8.11}$$

where the vector V is the voltage output of our DAC and the vector Φ is the flux biases at the qubit SQUIDS. In the crosstalk matrix, M_1 and M_2 are the intended conversions from DAC output to flux bias, and m_{12} and m_{21} are the crosstalk terms, which are generally not equal to each other.

To correct for this crosstalk, our goal is find the compensation matrix C , such that

$$M(I + C) = M' \tag{8.12}$$

$$M \begin{pmatrix} 1 & c_{12} \\ c_{21} & 1 \end{pmatrix} = \begin{pmatrix} M'_1 & 0 \\ 0 & M'_2 \end{pmatrix}. \tag{8.13}$$

Multiplying this expression through, we see that $c_{12} = -m_{12}/M_1$, and $M'_1 \approx M_1$ and

$M'_2 \approx M_2$ as long as m_{12} and m_{21} are small. ²

²The expression $1 + C$ is a first order compensation matrix, and the nth order correction can be computed as $1 + C + C^2 + C^3 + \dots + C^n$, or alternatively, as the nth Taylor expansion of $1/(1 - C)$. Higher orders account for the effect of crosstalk acting on the corrections from lower orders, which is important

In a crosstalk experiment, we label one of our qubits the "source" which generates a flux bias, and the other qubit the "target", which experiences a frequency shift due to crosstalk. One way to determine the compensation matrix element is to measure M_1 and m_{12} and compute c_{12} . Alternatively, we can also measure c_{12} directly. To do so, we first prepare the target qubit in a superposition state. Then, we apply a flux pulse to the source qubit and perform QST to measure the phase accumulated by the target. The amount of phase acquired by the target qubit will be proportional to the crosstalk matrix element and the amplitude of the flux pulse applied to the source. Simultaneously, we apply a flux pulse to the target qubit with the same length as the source flux pulse, but with a variable amplitude. When the amplitude of the target flux pulse cancels out the effect of crosstalk from the source flux pulse, the target qubit accumulates zero phase. We can then compute the compensation matrix element as

$$c_{12} = \frac{V_{\text{source}}}{V_{\text{target}}} . \quad (8.14)$$

8.3 Calibrating the Control Z

We will now move on to experimentally calibrating a control Z gate in two capacitively coupled transmons. The general strategy is as follows

1. Choose a trajectory with starting parameters and modify its height until a π con-
-
- when dealing with more than two qubits, and also helps the diagonal elements of the M matrix converge to their desired values.




Figure 8.5: (a) Energy level schematic for a CZ gate between two qubits, qA and qB. The qubits are initially far detuned to turn off the ZZ interaction. The $|0\rangle \rightarrow |1\rangle$ transition of qA is then brought near resonance with the $|1\rangle \rightarrow |2\rangle$ transition of qB to induce a conditional phase shift. Finally, qA is brought back to its initial position to resume operating single qubit gates or for measurement. (b) The actual time domain trajectory of qA's frequency. The goal of the calibration procedure is to find the optimal trajectory which induces a π conditional phase shift while minimizing leakage from $|11\rangle$ to $|02\rangle$.

ditional phase shift is achieved.

2. Optimize the parameters of the trajectory to minimize leakage.
3. Iterate over 2 and 3 until we converge on a gate with low leakage and the π phase shift.
4. Calibrate the phases used to compensate the effect of the trajectory on the qubits.

8.3.1 Conditional Phase

We first note that the actual unitary we will be calibrating is

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & e^{i\phi_A} & 0 & 0 \\ 0 & 0 & e^{i\phi_B} & 0 \\ 0 & 0 & 0 & e^{i(\phi_A+\phi_B+\pi)} \end{pmatrix}, \quad (8.15)$$

which is identical to Eqn. 8.1 except for the single qubit phase shifts ϕ_A and ϕ_B . These phase shifts are due to the frequency trajectory undertaken by the qubits, and are mostly independent of the coupling physics which gives rise to the conditional π phase shift. They can be accounted for by simply shifting the microwave frame of reference, such that the total effect is identical to the standard CZ. However, the presence of the single qubit phases means we must calibrate the conditional π phase shift using a differential measurement.

We begin by choosing a gate length (typically 45 ns) and Fourier parameters which produce a trajectory that is known from simulation to give low leakage. We prepare one of the qubits in a superposition, execute the CZ trajectory, then measure the accumu-

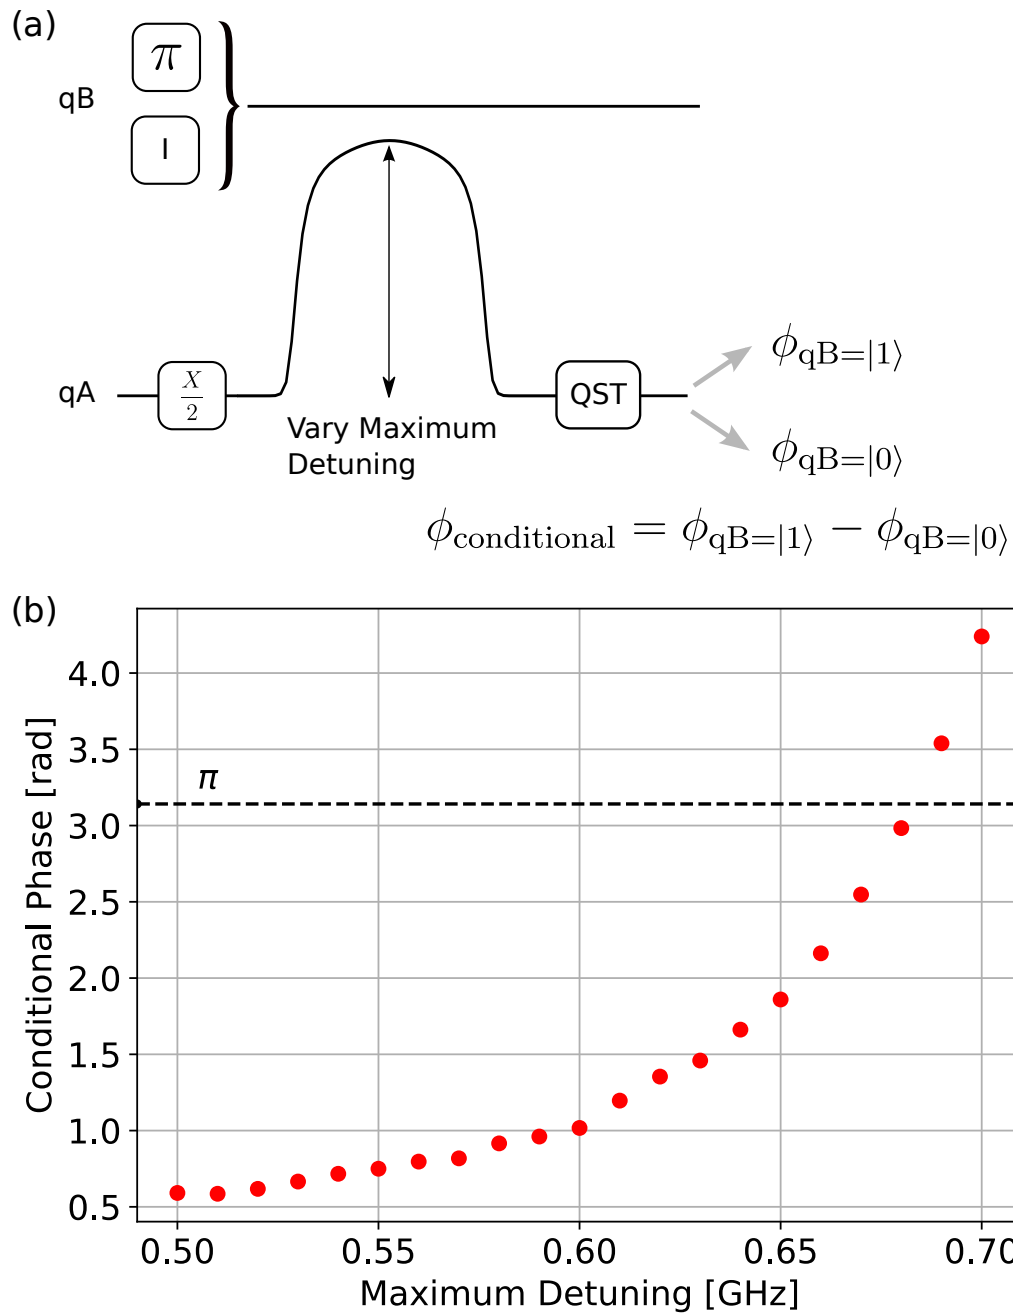


Figure 8.6: Calibrating the trajectory amplitude to obtain a π conditional phase shift. (a) Starting with qB in either $|0\rangle$ or $|1\rangle$, we prepare a superposition in qA, execute the CZ trajectory, then measure the resulting phase shift in qA. The conditional phase is the difference in phase shift due to qB occupying either $|0\rangle$ or $|1\rangle$. Note that because the CZ is symmetric, this experiment could also be performed by measuring the phase shift in qB conditional on qA. (b) Measured conditional phase as a function of the maximum detuning in the CZ trajectory.

lated phase using quantum state tomography. We repeat this experiment but initialize the other qubit in $|1\rangle$ rather than $|0\rangle$. If the CZ trajectory is correctly calibrated, the difference in phase between the two experiments will be π . If it is not, we vary the total amplitude of the trajectory until we find the amplitude which gives a π phase shift, as shown in Fig 8.6.

8.3.2 Leakage

Next, we measure and minimize the leakage due to the CZ trajectory. Since the primary error transition is $|11\rangle \rightarrow |20\rangle$, we prepare both qubits in $|1\rangle$, then execute the CZ and measure the resulting $|2\rangle$ state population in the higher qubit. We use this measured leakage as a cost function and minimize it by varying the Fourier parameters used to build the shape of the trajectory. Typically, Nelder-Mead is used to perform this optimization. After optimization, we typically find that the measured amount of leakage is limited by measurement error at the 1% level.

8.3.3 Iteration and Phase Compensation

Since the leakage calibration modifies the shape of the trajectory, the amplitude must be fine tuned again to achieve a π conditional phase. The conditional phase and leakage calibrations are iterated typically 2 or 3 times until both the conditional π phase shift and low leakage have been achieved. The final step of the CZ calibration is to calibrate the single qubit phases ϕ_A and ϕ_B . This calibration is performed by separately preparing

each qubit in a superposition and measuring the phase shift due to the CZ, while the other qubit remains in $|0\rangle$.

8.4 Benchmarking the CZ Gate

We now turn to benchmarking the performance of our calibrated CZ gate. As with single qubit gates, we will use Clifford based randomized benchmarking (RB) to estimate the total fidelity of our gates. The procedure is nearly identical, except we will now use the two qubit Clifford group and the final measurement is a joint measurement of the probability of the qubits remaining in $|00\rangle$. The two qubit Clifford group contains 11520 unitaries, and the standard decomposition into CZs and single qubit microwave gates is given in Ref. [64]. On average, a two qubit Clifford contains 1.5 CZs, and 33/4 single qubit gates split evenly between the two qubits. Standard RB results are shown in Fig. 8.7(a), and as in the single qubit case, we fit the average sequence fidelity F to

$$F(m) = Ap^m + B, \quad (8.16)$$

where m is the number of Cliffords in the sequence, A and B are parameters related to state preparation and measurement, and $1 - p$ is the depolarization rate. For this dataset, we find $A = 0.696 \pm 0.002$, $B = 0.247 \pm 0.002$, $1 - p = 2.19 \pm 0.03 \times 10^{-2}$, and $r_{\text{Clifford}} = \frac{3}{4}p = 1.64 \pm 0.02 \times 10^{-2}$.

Unlike the single qubit case, we cannot simply estimate the CZ error from the error per Clifford because two qubit Cliffords contain both single and two qubit gates. Instead, we use *interleaved randomized benchmarking*, which is identical to standard RB except




Figure 8.7: (a) Standard two qubit Clifford randomized benchmarking. (b) Randomized benchmarking with CZs interleaved after every Clifford.

Table 8.1: Error rates for single and two qubit benchmarking. The leakage rate was not measured for qA but is assumed to be roughly the same as for qB. For incoherent and leakage errors, the error rate per CZ was found by simply subtracting the reference value from the interleaved value, as there is presently no theory for a more accurate estimate of interleaved purity or leakage values.

	Total Error	Incoherent Error	Leakage Rate
Single qubit, qA	$4.3 \pm 0.1 \times 10^{-4}$	$3.95 \pm 0.04 \times 10^{-4}$	
Single qubit, qB	$4.8 \pm 0.2 \times 10^{-4}$	$3.70 \pm 0.04 \times 10^{-4}$	$1.5 \pm 0.1 \times 10^{-5}$
Two Qubit Clifford	$1.64 \pm 0.02 \times 10^{-2}$	$1.20 \pm 0.04 \times 10^{-2}$	$7.2 \pm 0.3 \times 10^{-4}$
Two Qubit CZ	$5.4 \pm 0.4 \times 10^{-3}$	$3.42 \pm 0.05 \times 10^{-2}$	$2.8 \pm 0.4 \times 10^{-4}$

we insert a CZ after each two qubit Clifford. The results of interleaved RB are shown in Fig. 8.7(b), and we fit the average sequence fidelity to 8.16, but interpret m as the number of Cliffords *and* interleaved CZs. We infer the error per CZ by using the standard RB result as a reference and computing

$$r_{CZ} = (1 - p_{\text{interleaved}}/p_{\text{reference}})(d - 1)/d \quad (8.17)$$

where $d = 2^n$ is the dimensionality of the Hilbert space. For the dataset in Fig. 8.7, we find that $r_{CZ} = 5.4 \pm 0.4 \times 10^{-3}$.

Purity and leakage benchmarking can also be performed with two qubit gates in a manner analogous to single qubit gates, with results shown in Fig. 8.8 and Fig. 8.9. The extracted incoherent error rates and leakage rates, along with benchmarks for the single qubit gates for qA and qB, are collected in Table 8.1.

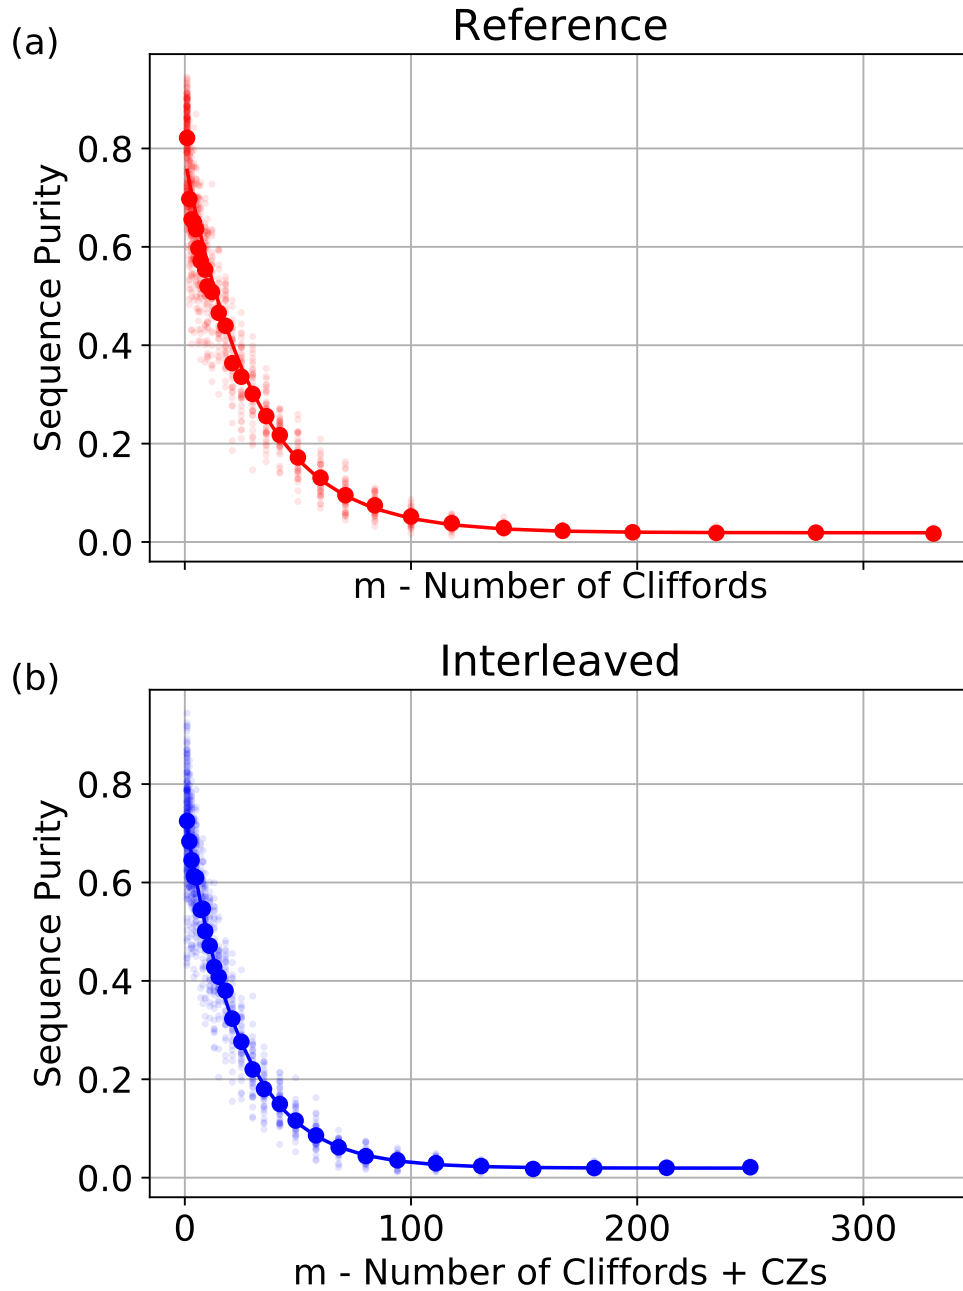


Figure 8.8: (a) Two qubit purity benchmarking. (b) Purity benchmarking with CZs interleaved after every Clifford.




Figure 8.9: (a) Two qubit leakage benchmarking. The leakage population was only measured for the higher of the two qubits (qB), since the dominant source of leakage swaps excitation into qB's $|2\rangle$ state. (b) Leakage benchmarking with CZs interleaved after every Clifford.

8.5 Conclusions

In this chapter, we described the CZ gate used in the UCSB Xmons for entangling operations, and detailed how to calibrate it. Using randomized benchmarking, we found that the fidelity of our CZ is 99.45%, which compares favorably with two qubit gates in other systems [65, 141]. Leakage due to our CZ is an order of magnitude higher than it is for single qubit gates, but is still not the dominant factor in gate infidelity. Rather, decoherence contributes a slight majority (60%) to the CZ error with control errors contributing most of the rest. Unlike the single qubit case, comparing incoherent RB error to independent measures of decoherence is difficult because the qubit frequency does not remain fixed during the CZ gate, and decoherence can be strongly frequency dependent. Future research should focus on matching the spectrum decoherence times in the CZ trajectory to the performance of the CZ. In addition, control errors are likely dominated by phase drift of the qubit following each CZ, and decreasing this phase drift through improved cabling and calibration software will be crucial to reaching higher fidelities.

8.6 Contributions

Calibration and randomized benchmarking of the CZ gate were first performed in the group by Julian Kelly and Rami Barends, while the flux control techniques presented here were first pioneered by Charles Neill. I further refined the flux control and CZ calibration

techniques and introduced two qubit purity (with help from Sergio Boixo, Steve Flamia, and Chris Granade) and leakage benchmarking. I acquired all of the data presented in this chapter, and also performed all of the analysis. The device was fabricated by the Google Quantum Hardware group.

Chapter 9

Readout Induced Qubit Transitions

Having calibrated high fidelity single and two qubit gates, we now return to qubit measurement. In a textbook quantum system, after readout of the qubit state, the qubit should remain in that state that was measured. In this chapter, we will explore two ways this assumption can fail in superconducting transmons measured using dispersive readout through a coupled readout resonator. First, we will review Sank, Chen, and Khezri et. al. [142], where we find that a transmon can be driven to excited states through resonant interactions within the Jaynes-Cummings (JC) ladder. Second, we will show preliminary data for a similar phenomenon which occurs when the qubit is itself also driven.

9.1 Measurement Induced State Transitions

9.1.1 Motivation

Holding all other parameters equal, the signal to noise ratio for dispersive measurement should increase with higher input power and longer measurement time. Because measurement time is constrained by the energy relaxation of the qubit, increasing the measurement power is a key knob in calibrating high fidelity qubit readout. However, several experiments with superconducting qubits have found that as the number of photons occupying the resonator \bar{n} is increased past a certain point, the qubit suffers anomalous state transitions [143, 144, 68, 74]. It was long suspected that these transitions were caused by the breakdown of the dispersive approximation of the JC model as \bar{n} exceeds a critical photon number n_c , but recent theory showed that the transitions are not predicted by the JC interaction even with very large \bar{n} [145].

9.1.2 Experimental Observations

To study these transitions, we tested a transmon qubit coupled to a resonator with coupling strength $g/2\pi \approx 87$ MHz, as illustrated in Fig. 9.1 (a). The transmon is biased such that $\omega_{10}/2\pi = 5.4$ GHz with $\eta/2\pi \equiv (\omega_{21} - \omega_{10})/2\pi = -221$ MHz, while the resonator frequency is $\omega_r/2\pi \approx 6.78$ GHz with a decay rate of $\kappa \approx 1/(37 \text{ ns})$ through a bandpass Purcell filter [146, 74] to a 50Ω output line and amplifiers.

We have previously seen that the transmon supports more than just the two qubit

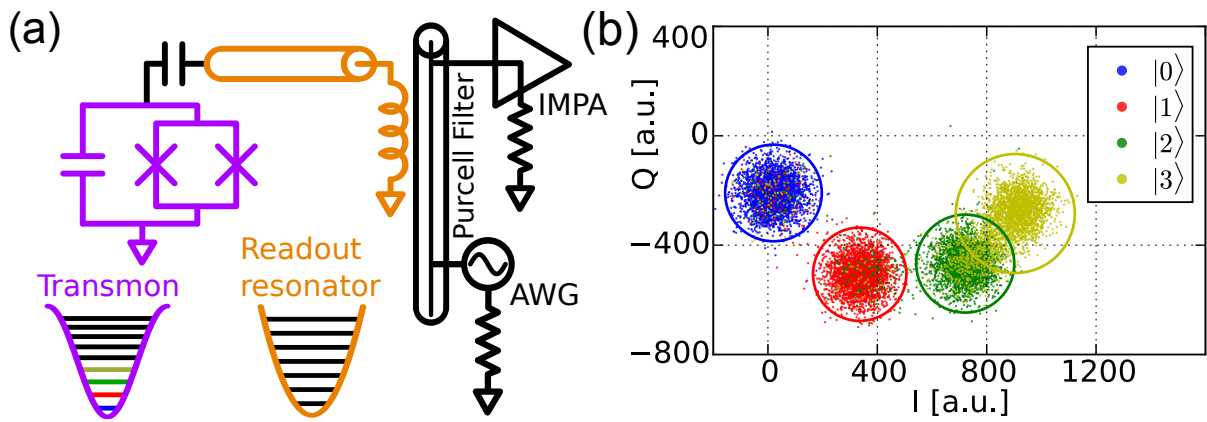


Figure 9.1: Transmon-resonator system. (a) Circuit and potential diagrams. The transmon (violet) is capacitively coupled to the resonator (orange). The resonator is inductively coupled to a bandpass Purcell filter with $Q \approx 30$ [74]. The resonator is driven by an arbitrary waveform generator connected to the filter, and the dispersed photons are measured by a low noise, impedance matched parametric amplifier [99] also connected to the filter. (b) In-phase and quadrature (IQ) components of the dispersed signal measured with the transmon prepared in the first four states, with each state forming an IQ “cloud”. The circles represent 3σ from fitting a Gaussian distribution to each cloud’s projection onto lines connecting the clouds’ centers.

levels. In this experiment, we will measure levels up to $|3\rangle$. In Fig. 9.1 (b), we plot the complex scattering response (IQ points) of the resonator with the transmon prepared in various states, which acts as our calibration for distinguishing the state of the transmon in subsequent measurements. When the resonator-transmon detuning $|\Delta| \equiv |\omega_{10} - \omega_r|$ is not more than 1.4 GHz, the resulting IQ points resolve up to the first four transmon states, while at larger $|\Delta|$ (relevant to most of our data) we can only resolve the first three states due to the smaller dispersive shift.

To investigate the effect of resonator photons on the transmon state, we use the pulse sequence illustrated in Fig. 9.2 (a). With the transmon initialized in $|0\rangle$, we first drive the resonator with a $2\ \mu\text{s}$ long, variable power pulse. This “stimulation pulse” injects a number of photons into the resonator that, when large enough, induces transitions in the transmon state. We then wait 500 ns (13 decay time constants) for the resonator to ring down. Finally, we drive the resonator again with a fixed low power pulse to measure the transmon without inducing further transitions, and record the IQ response of the resonator. Based on the calibration shown in Fig. 9.1(d), we identify each IQ point as one of the transmon states, or if the point is more than three standard deviations from any of the calibrated distributions, we label it as an “outlier”.

The results are striking in two ways. First, as the stimulation pulse power is raised, the transmon jumps from $|0\rangle$ not only to $|1\rangle$ but also to $|2\rangle$, $|3\rangle$ and even higher states, as shown in Fig. 9.2 (b). Although we can resolve only up to $|3\rangle$, the characteristic arc of the IQ points with increasing state index appears to continue to what we estimate

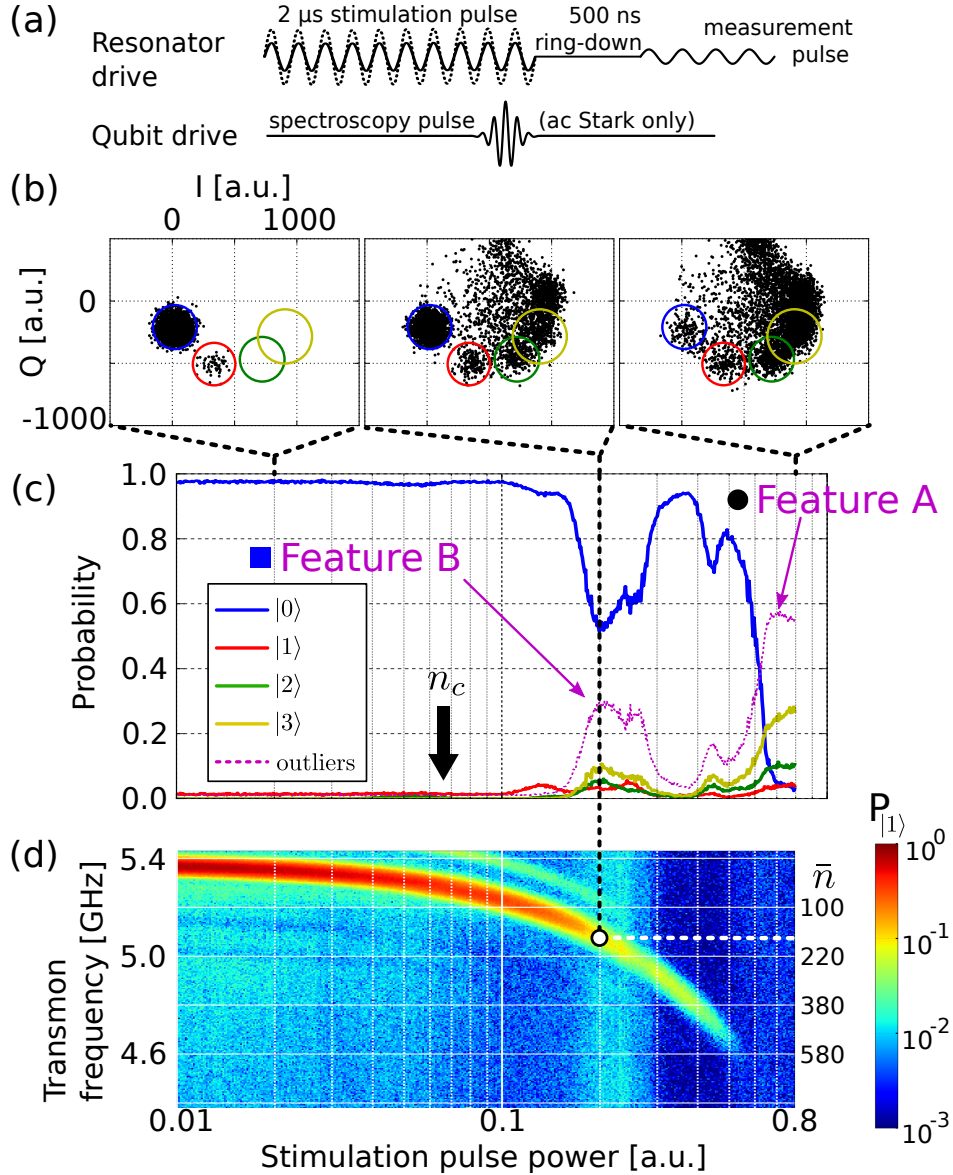


Figure 9.2: (a) Control sequence for probing the effect of resonator photons on the transmon. The spectroscopy pulse is used only in the AC Stark measurement. (b) IQ data for drive powers 0.02, 0.2, and 0.8 (arbitrary units), with $\omega_{10} = 5.38$ GHz. The circles represent 3σ for the four resolvable transmon states as calibrated in Fig. 9.1b. At high power, the transmon is clearly driven to states higher than $|3\rangle$. (c) Transmon state probabilities versus stimulation power. In addition to the four calibrated transmon states, we show the probability that the measurement was $> 3\sigma$ from any of the resolved states, labeled “outliers”. Note the two large resonance-like peaks labeled *A* and *B*. (d) Stark shifted transmon frequency ω_{10} versus stimulation pulse power. We convert the shifted ω_{10} to \bar{n} using a numerical theory (right vertical axis) [142].

to be $|5\rangle$ or higher. Second, the probability of transitions is highly non-monotonic with power, as was previously seen in Refs. [68, 144]. In particular, the shapes of the features in probability versus power resemble resonance peaks, with large peaks in the outlier probability at drive powers 0.7 (feature *A*) and 0.2 (feature *B*), a small peak in $|1\rangle$ near 0.15, another small peak in $|2\rangle$ near 0.05, and various other peaks at other powers. The peaked structure rules out any process that would have monotonically increasing transitions with increasing drive power, such as chip heating or dressed dephasing [147, 148], as the dominant mechanism.

In order to connect our results to theoretical models, we next convert stimulation pulse power to photon number \bar{n} . We cannot measure \bar{n} directly, but recall from our discussion of the Jaynes-Cummings ladder in Chapter 2 that increasing photon occupation in the resonator causes the qubit frequency to shift downwards due to the AC Stark shift [149]. Thus, we can map drive power to \bar{n} by measuring the AC Stark shifted qubit frequency for each resonator drive power, then converting the frequency shift to \bar{n} using a numerical model based on separately measured parameters g and Δ [142]. To measure the AC Stark shift, we repeat the previous experiment with the addition of a spectroscopic microwave pulse on the transmon after the driven resonator has reached the steady state. At each resonator drive power, we vary the frequency of the transmon pulse; the $|1\rangle$ probability is maximized when the pulse is on resonance with the shifted transmon frequency.

We show the results of the AC Stark shift measurement with the computed photon numbers in Fig. 9.2 (d) for the same drive powers as in Fig. 9.2 (c). Note that feature *B*

(black dashed line) occurs at $170 \lesssim \bar{n} \lesssim 250$, which is, interestingly, considerably larger than the critical photon number $n_c \equiv (\Delta/g)^2/4 \approx 60$ introduced in Ref. [150].

9.1.3 Theory of non-RWA Transitions

To understand the source of these transitions, we return the Jaynes-Cummings ladder from Fig. 2.5, but consider rungs of the ladder with many more excitations. With more available excitations in each rung, we must add in the higher states of the transmon. Recall that the couplings within each rung repel the states, imparting an n -dependent shift on the bare levels to produce eigenstates, two of which are shown as dashed lines in Fig. 9.3. As indicated by the long horizontal arrow, at certain n the ladder contains resonances between states where the qubit goes from $|0\rangle$ to higher levels such as $|6\rangle$. This critical observation could explain both the resonance structure and the transitions to higher transmon levels observed in the data.

However, according to the rotating wave approximation (RWA), levels that are not in the same rung (i.e. do not conserve excitations) should not couple. In fact, the non-RWA couplings of the Hamiltonian can be as large as the RWA couplings, but are usually neglected on the grounds that they are more off resonance than the RWA couplings (in our system the RWA pairs are ~ 1 GHz off resonance, while the non-RWA pairs are ~ 13 GHz off resonance). However, keeping the non-RWA terms in our Hamiltonian reveals the essential reason for the unwanted state transitions. The non-RWA terms couple next-nearest neighboring RWA strips (i.e., those differing by 2 in total excitation

number) together, as shown in Fig. 9.3. Combined with normal intrastrip RWA coupling, the non-RWA coupling allows multistep (i.e., higher order) processes to connect the resonant levels. For example, non-RWA coupling carrier $|0, n\rangle$ to $|1, n + 1\rangle$ in another RWA strip, and then RWA couplings carry the system within the strip to $|6, n - 4\rangle$. Note that although the full process conserves energy, the individual steps do not.

9.1.4 Comparing Theory to Experiment

To find the condition under which the resonances occur, we numerically compute the frequencies $\bar{\omega}_k(n) \equiv E_{\overline{|k, n-k\rangle}}/\hbar - n\omega_r$ (the overline indicates eigenstate) of the levels within each RWA strip, as functions of n . As n increases, energy levels within each strip repel each other more strongly and fan out, as illustrated by the solid lines in the “fan diagram” in Fig. 9.4 (a). By superimposing fan diagrams of two next-nearest neighboring RWA strips, as shown by the dashed lines, we see that they have multiple intersections, meaning that the JC ladder contains multiple resonances. For example, the left red dot in Fig. 9.4 (a) shows that the transmon-resonator state $|0, n\rangle$ can be brought on resonance with $|6, n - 4\rangle$. The presence of crossings with higher transmon states agrees with the experimental observation of transitions to states higher than $|3\rangle$.

Next, we compute the n at which interactions in the fan diagram occur as a function of the qubit-resonator detuning Δ , yielding the lines in Fig. 9.4 (b). As $|\Delta|$ increases, the spacing between levels within an RWA strip also increases, as shown in Fig. 9.3. However, the spacing between strips is fixed at ω_r , so with increased $|\Delta|$ fewer photons are required




Figure 9.3: Jaynes-Cummings ladder for large values of n . Bare states are shown as solid lines and two of the eigenstates are shown as dashed lines. Dark curved arrows indicate coupling within an RWA strip with corresponding RWA coupling strengths shown below. The ladder has an energy resonance between $|0, n\rangle$ and $|6, n-4\rangle$ (long black arrow). Non-RWA couplings (short straight arrows) allow for interstrip transitions. The couplings to $|1, n+1\rangle$ (red) and $|3, n-1\rangle$ (yellow), along with those within the RWA strip, mediate the transition between the resonant levels. The coupling to $|2, n-1\rangle$ (green), which mediates additional resonant transitions, requires a Hamiltonian term coupling transmon states of equal parity; this is forbidden if the transmon potential is symmetric. Note the energy spacing between states $|k, n\rangle$ and $|k+1, n-1\rangle$ is Δ as indicated in the top left.

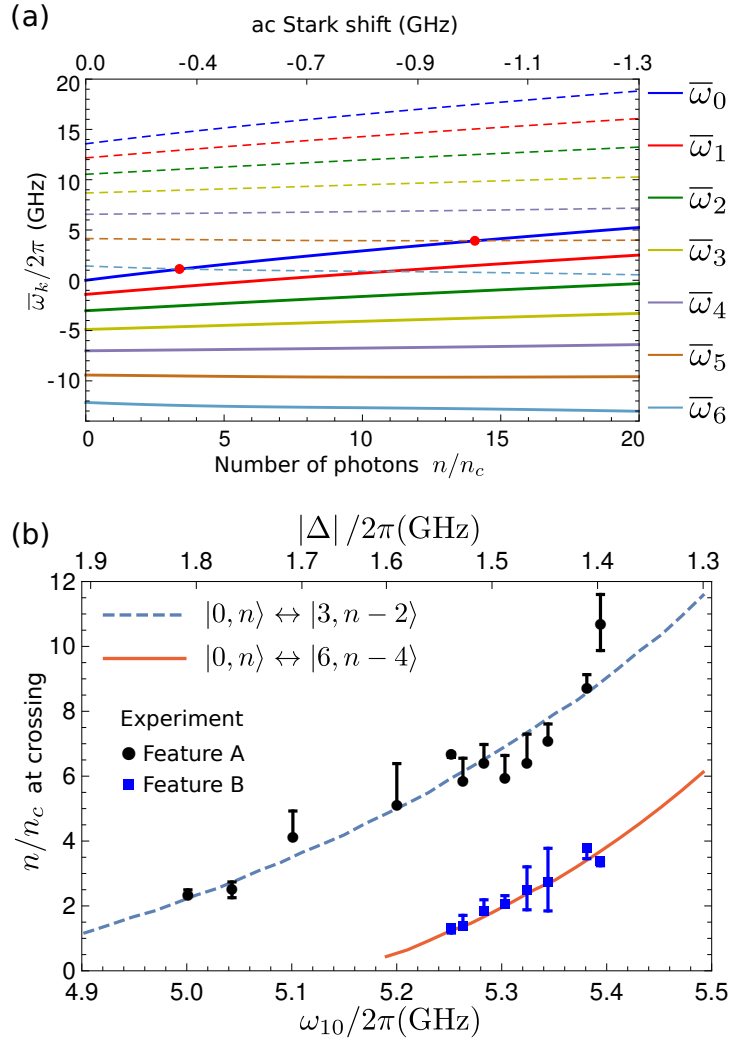


Figure 9.4: (a) Fan diagram of the energy levels within an RWA strip. Solid: Frequencies $\bar{\omega}_k(n) \equiv E_{|k, n-k\rangle}/\hbar - n\omega_r$ versus photon number n for $|\Delta| = 1.4$ GHz. As n increases, the levels repel more strongly and fan out. Dashed: Same frequencies shifted by $2\omega_r$, which represent the next-nearest neighboring RWA strip. The red dots show energy resonances with the qubit state $|0\rangle$ occurring at specific values of n . The left dot corresponds to the resonance shown in Fig. 9.3. (b) Photon number at level crossing versus ω_{10} , comparing experiment to theory. Black circles and blue squares show experimental features A and B from Fig. 9.2 respectively, and the error bars represent the apparent widths of the features. Solid red line is the theory prediction for level crossing between eigenlevels of $|0, n\rangle$ and $|6, n-4\rangle$. Dashed blue line is the theory prediction for an asymmetric transmon that breaks the selection rule by at least 1%, yielding level crossings between eigenlevels of $|0, n\rangle$ and $|3, n-2\rangle$.

to bring $|0, n\rangle$ on resonance with states in higher strips and so the transitions occur at lower \bar{n} . Note that while we use n in the theory, the experiment drives the resonator into a coherent state with mean photon number \bar{n} and fluctuations $\sqrt{\bar{n}} < 0.1 \bar{n}$. Also, although the n at which the energy resonance occurs is not related to n_c , the effective couplings between resonant levels are large enough to yield the experimental features only when $n \gtrsim n_c$.

To confirm the theoretical prediction, we repeat the experiment shown in Fig. 9.2 for several values of ω_{10} by biasing the transmon's SQUID with magnetic flux. At each ω_{10} , we find the values of \bar{n} of features A and B , as shown in Fig. 9.2 (d), and plot these points in Fig. 9.4 (b). The experimental points for feature A (black circles) and feature B (blue squares) are well fit by numerically computed curves for the transitions from $|0, n\rangle$ to $|6, n - 4\rangle$ and $|3, n - 2\rangle$, respectively. Note that the theory lines are calculated using only the measured ω_r , ω_{10} , and g , with no free parameters fitted to the data.

However, the transition from $|0, n\rangle$ to $|3, n - 2\rangle$ is actually unexpected. If the transmon potential is symmetric, as is usually assumed [63], then $g_{i,j}$ is only nonzero when $j - i$ is odd. Therefore, the transmon-resonator coupling should only couple RWA strips where the difference in total excitation number is even, so the transition to $|3, n - 2\rangle$ should be forbidden. Nevertheless, the theory line for the $|3, n - 2\rangle$ transition fits the data well, indicating a possible asymmetry in the transmon potential. We confirmed this asymmetry by observing $|0\rangle \rightarrow |2\rangle$ Rabi oscillations when driving the transmon at $\omega_{01} + \omega_{12}$ [142]. Through comparison with Rabi oscillations on the $|0\rangle \rightarrow |1\rangle$ transition, we experimentally

estimate $|\langle 0|Q|2\rangle/\langle 0|Q|1\rangle| \approx 10^{-2}$ [142]. This matrix element is large enough to explain the transitions to $|3, n-2\rangle$, and so the level crossing theory appears to correctly predict both of the largest resonance features observed in the data.

9.1.5 Summary

So far, we have found that strong dispersive measurement of a transmon induces transitions to states above $|3\rangle$. These transitions occur at specific values of the photon occupation in the measurement resonator, and are caused by energy resonances within the qubit-resonator system. Coupling between the resonant levels is mediated by Hamiltonian terms usually dropped in the rotating wave approximation, and the most important such term involves an unexpected broken symmetry in the transmon potential.

9.2 Induced Transitions in a Driven Qubit

We will now continue to explore the dynamics of dispersive readout in a system where the resonator frequency is below that of the qubit rather than above. This frequency arrangement gives us the ability to rapidly reset the qubit state by bringing the qubit on resonance with the resonator, which can be a useful resource for algorithms. However, we will see that the the dynamics of the qubit and resonator are more complicated in this arrangement.

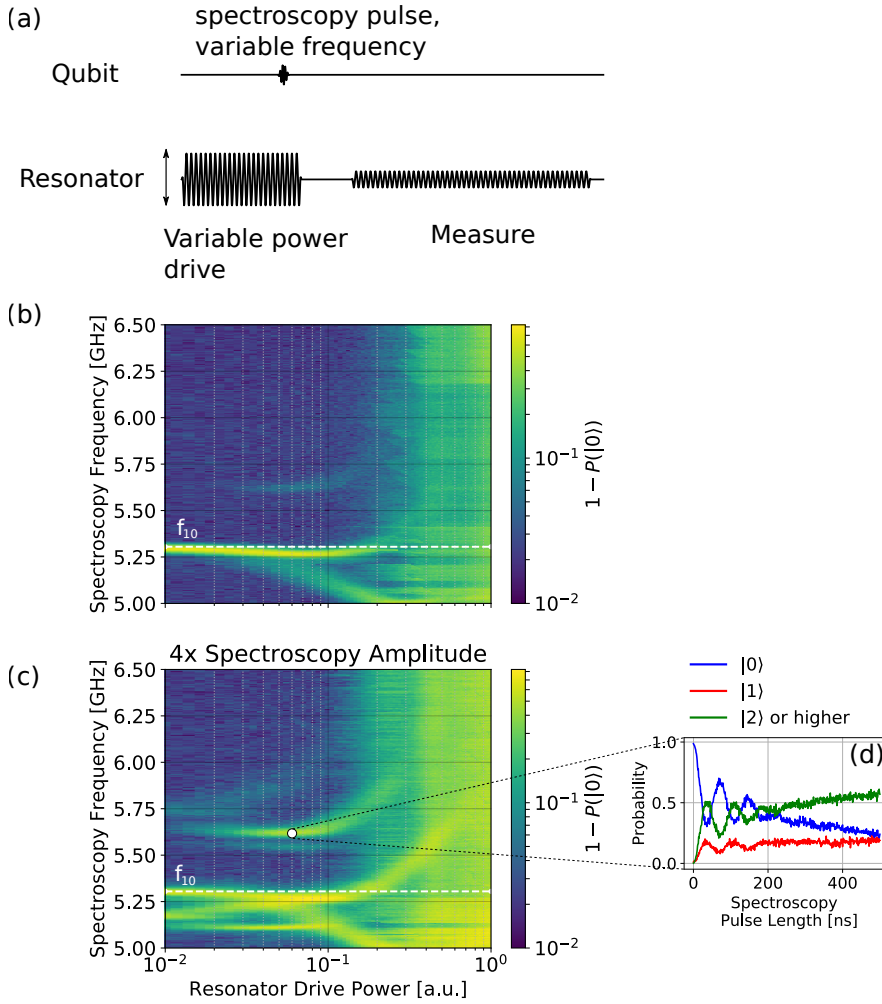


Figure 9.5: AC Stark spectroscopy with the resonator below the qubit. (a) Pulse sequence. We ring the resonator up for 500 ns with a variable power drive tone, then probe the qubit with a 50 ns microwave spectroscopy pulse once the resonator has reached equilibrium. Finally, we measure the qubit with a calibrated pulse on the resonator. (b) With a spectroscopy pulse sufficient to drive the usual $|0\rangle \rightarrow |1\rangle$ transition, we see that the qubit frequency initially decreases, but then diverges and also curves upward. (c) Increasing the spectroscopy amplitude reveals more features in the transmon spectrum, including the two photon transition as well as unexpected resonances above the qubit frequency. (d) If we sit at a fixed spectroscopy frequency and resonator drive power, and vary the spectroscopy pulse length, we observe Rabi oscillations between $|0\rangle$ and a higher excited state of the transmon, $|2\rangle$ or above.

9.2.1 The AC Stark Effect with $\omega_r < \omega_q$

We begin by repeating the AC stark experiment from Fig.9.2 in our device where the resonator frequency is below that of the qubit. We place the qubit f_{10} at 5.305 GHz, and the transmon's $|1\rangle \rightarrow |2\rangle$ transition is $f_{21} = 5.05$ GHz. The resonator frequency is at 4.6 GHz, and the qubit-resonator coupling is approximately 100 MHz as measured by the dispersive shift. As before, we ring up the resonator with variable power, and drive the qubit with a variable frequency microwave spectroscopy pulse. The microwave spectroscopy pulse is calibrated to drive the $|0\rangle \rightarrow |1\rangle$ transition when the pulse is on resonance with the qubit and there are no photons in the resonator. After the microwave spectroscopy pulse, we allow the resonator to ring down, then measure the excited state population with a calibrated measurement pulse. The results are shown in Fig 9.5(a). As in the previous AC Stark experiment in Fig.9.2, the qubit frequency initially decreases with increasing resonator drive, but strangely increases at higher drive power and also appears to split into two lines.

Next, we repeat the same experiment but drive the microwave spectroscopy pulse with 4 times more amplitude, with the results shown in Fig 9.5(b). We expect to see a transition appear below the qubit frequency at $(f_{10} + f_{21})/2$ due to the two photon transition from $|0\rangle$ to $|2\rangle$. However, we also see additional excitations at frequencies higher than the qubit frequency, and with a distinct curvature as a function of drive power. These transitions cannot be explained by virtual transitions to higher states which always occur below the qubit frequency due to the sign of the anharmonicity.

Furthermore, the spurious transitions do not occur at zero resonator drive amplitude, so they appear to be related to the dynamics of the resonator.

To further characterize these unexpected transitions, we choose a fixed resonator drive power (0.06) and spectroscopy frequency (5.2 GHz) which lie in one of the transition regions. Using these parameters, we repeat the AC Stark spectroscopy experiment but vary the length of the spectroscopy pulse. We observe Rabi oscillations from the $|0\rangle$ state to $|2\rangle$ or possibly higher, indicating the higher states of the transmon are involved in some manner.

9.2.2 Revisiting the JC Ladder

To explain the strange shape of the AC Stark shift and anomalous resonances, we return to the JC ladder construction. In Fig. 9.6, we show two RWA strips in the JC ladder for the resonator frequency placed below the qubit frequency. We notice an immediate qualitative difference: within each RWA strip, the energies of the levels do not increase or decrease monotonically with increasing qubit energy. Rather, each RWA strip curves back on itself, due to flipping the sign of Δ , but *not* the sign of the qubit anharmonicity. As the number of photons and the intrastrip level repulsions increase, energy level crossings may occur *within* an RWA strip and lead to the AC Stark shifts which reverse sign. This effect is demonstrated in the “fan diagram” for the resonator placed below the qubit, shown in Fig. 9.7.

Next, we consider the effect of the microwave qubit drive that is present in our

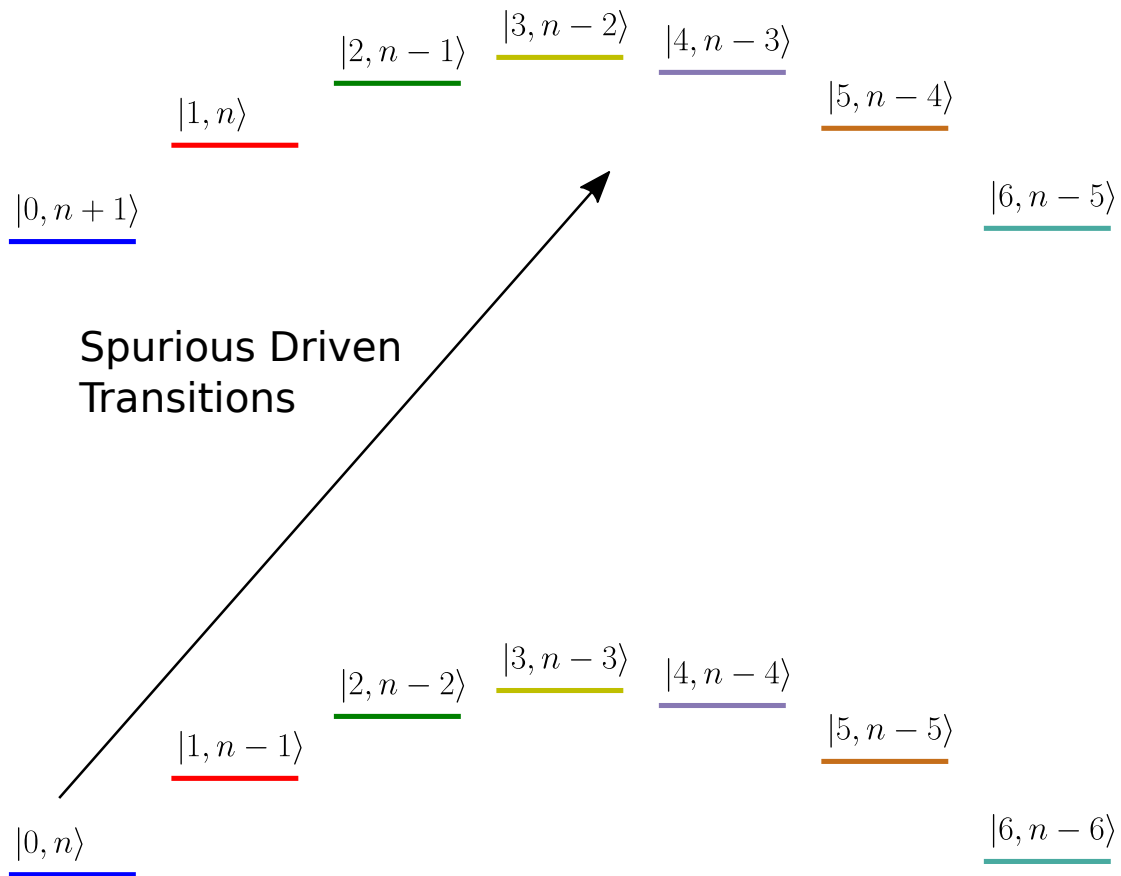


Figure 9.6: Jaynes Cummings ladder with the resonator frequency below the qubit frequency. The RWA strips curve back on themselves because the transmon anharmonicity and resonator-transmon detuning have the opposite sign. When the qubit is driven, the two strips naturally couple and the transmon can be driven to higher states.




Figure 9.7: Energy eigenlevels vs photon number within one RWA strip, for $f_{10} = 5.305$ GHz and $f_r = 4.6$ GHz. Compare with Fig. 9.3.

spectroscopy experiment. While the RWA forbids intra-strip couplings in a “static” Hamiltonian, there is no such restriction when the qubit is driven since the drive can add or remove excitations from the system. Therefore, we naturally have a mechanism for coupling the state $|0, n\rangle$ to $|1, n\rangle$, as well as $|2, n - 1\rangle$ and higher since we have already established that $|0\rangle \rightarrow |2\rangle$ couplings are possible due to broken symmetries.

Based on these observations, we conjecture the following: the ground state of the transmon is being resonantly driven by the microwave drive to higher states of the transmon in the next RWA strip. Coupling is mediated by a combination of the anomalous $|0\rangle \rightarrow |2\rangle$ coupling and a \sqrt{n} enhancement in the coupling strength due to the photons in the resonator. The states in the next RWA strip are not far detuned from each other because the negative anharmonicity causes the strip to curve back on itself. There, these anomalous transitions should be relatively close to f_{10} , e.g. the transition to $|2, n - 1\rangle$ will occur at $f_{10} + |\Delta| - |\eta|$. Furthermore, these higher states of the transmon experience repulsions due to the intra-strip couplings, which will lead to n -dependent behavior of the resonance frequencies of these transitions.

9.2.3 Comparison to Theory

To test our conjecture, we numerically solve for the eigenlevels of the transmon and resonator. Then, for each photon number n , we compute the transition frequency to move from $|0, n\rangle$ to the higher states in the next RWA strip, such as $|1, n\rangle$ (which gives us the AC Stark shifted f_{10}), $|2, n - 1\rangle$, $|3, n - 2\rangle$ and so on. Using the Stark shifted shape




Figure 9.8: The strongly driven AC Stark spectroscopy data from Fig. 9.5(b), overlaid with numerical theory lines from $|0, n\rangle$ to the states indicated in the legend. The only parameters for the theory lines are f_{10} , η , the qubit-resonator coupling g , and a linear mapping from \bar{n} to resonator drive power.

of the f_{10} transition in Fig. 9.5(b), we calibrate a simple linear mapping from resonator drive amplitude to photon number. Finally, we overlay the transition frequencies as a function of drive amplitude on top of Fig. 9.5(b) to produce Fig. 9.8.

We see that our predictions match remarkably well with the transitions seen in our AC Stark experiment. The divergence in the f_{10} transition is in fact due to a level crossing between $|1, n-1\rangle$ and $|5, n-4\rangle$ within the excited RWA strip¹. The two transitions above f_{10} move $|0, n\rangle$ to $|4, n-3\rangle$ and $|2, n-1\rangle$, indicating again the importance of symmetry breaking couplings in the transmon. Curiously, we do not observe the transition to $|3, n-2\rangle$ in the data.

To further test our theory, we repeat the AC Stark experiment with the qubit operating at $f_{10} = 5.620$ GHz, with the data and theory predictions shown in Fig 9.9. As we move further away from the resonator, the increase in Δ spreads out the RWA strips so that the transition to $|2, n-1\rangle$ occurs further away from f_{10} . However, we also see the appearance of a transitions to $|6, n-5\rangle$.

Finally, we repeat the experiment with $f_{10} = 5.90$ GHz, as shown in Fig 9.10 The transition to $|2, n-1\rangle$ has moved even further away from f_{10} and it is the sole remaining transition that we see in the spectrum. In all cases, our numerical predictions qualitatively match the resonances observed in the data.




Figure 9.9: AC Stark spectroscopy with a strong qubit drive, overlaid with a numerical theory, for $f_{10} = 5.62$ GHz and $|\Delta| = 1$ GHz.




Figure 9.10: AC Stark spectroscopy with a strong qubit drive, overlaid with a numerical theory for resonator-mediated qubit transitions, for $f_{10} = 5.9, \text{GHz}$ and $|\Delta| = 1.3 \text{ GHz}$.

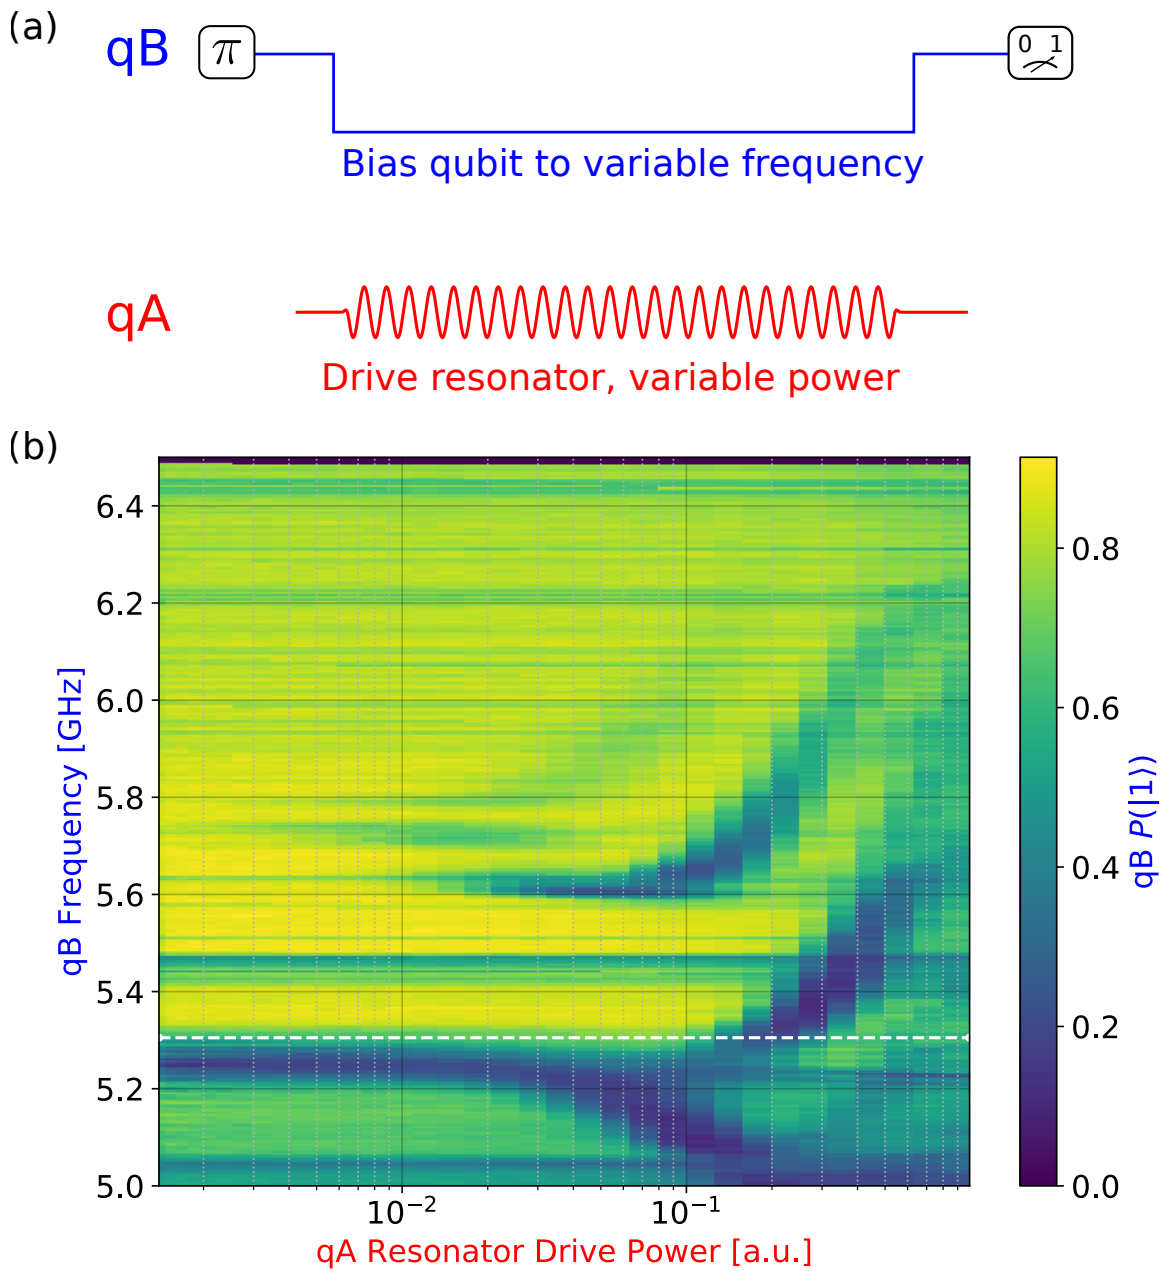


Figure 9.11: Multiqubit "swap" spectroscopy. We initialize qB in $|1\rangle$, then drive qA's resonator with variable power while biasing qB to variable frequency and maintaining qA's frequency at 5.3 GHz, then measure qB's population. Compare with Fig. 9.8.

9.2.4 Multi Qubit Transitions

So far, we have established that simultaneously driving the qubit and resonator can resonantly excite the qubit to higher excited states at frequencies near f_{10} . In a typical algorithm, the resonator drive is only used for measurement. Therefore, simultaneous drive of the qubit and resonator seems to have little relevance for practical applications. However, microwave drive is not the only way to excite a qubit. A second qubit which is in the excited state and coupled to the qubit being measured is also a potential source of excitations.

To see the equivalence of these two cases, we return our qubit under test - which we will refer to as qA - to operating at $f_{10} = 5.305$ GHz, and tune up a second qubit which is coupled to qA and initially far detuned, which we will refer to as qB. We then perform swap spectroscopy with qB: we populate qB in $|1\rangle$, dynamically bias it to a variable frequency for 500 ns, then measure the resulting population. While qB is being dynamically biased, we drive qA's resonator with variable power. The results of this experiment are shown in Fig. 9.11. By comparing the swap spectroscopy data with the AC Stark experiment in Fig. 9.5, we clearly see that when qB is brought on resonance with the resonator mediated transitions of qA, the excited state population of qB is depleted. We also separately confirmed by measuring qA instead of qB at the end of the swap spectroscopy experiment that qA is excited in the regions where qB is depleted.

¹Note that labeling the eigenstates is nontrivial since the increased coupling at high photon numbers mixes the states significantly, so these state labels should not be interpreted as definitive.

9.3 Conclusions

In this chapter, we have studied two interesting effects of dispersive readout on transmons. In the first half of the chapter, we found that strong dispersive measurement can give rise to resonance conditions in the resonator-transmon energy level structure, which induce transitions in the transmon from $|0\rangle$ to higher states of the transmon. In the second half, we found that dispersive readout can create additional transitions in the transmon spectrum, again to higher states in the transmon well. Both effects can be understood by carefully considering all of the levels in the Jaynes-Cummings ladder and by including couplings that have previously been ignored. Note that while the first effect was studied with a system where the resonator is placed above the qubit and the second effect was studied with the opposite frequency placement, both effects can occur in both types of systems.

These effects may be a discouraging sign for dispersive readout as they place a serious limit on readout fidelities, and complicate efforts to scale up the number of qubits. However, knowing about these effects suggests a few areas for improvement and research:

1. Decreasing the number of photons required to read out the qubit state by increasing quantum efficiency of the readout chain, since both readout transition effects are amplified by the resonator photon occupation.
2. Strategically placing the resonator-qubit detunings to avoid resonances. For example, in the case of the resonator placed above the qubit frequency, it is advantageous to have a smaller Δ .

3. Investigating the source of the unexpected $|0\rangle \leftrightarrow |2\rangle$ coupling. So far, no physical model has been found which explains the presence of this allowed transition.

9.4 Contributions

Daniel Sank, Mostafa Khezri, and Alexander Korotkov contributed significantly to the data, text and figures of the first half of this chapter. I acquired the data and performed the numerical analysis for the second half with the use of analysis code written by Mostafa, which can be found at <https://github.com/MostafaKhezri/JC-ladder>.

Chapter 10

Conclusion

In this thesis, we discussed in detail the calibration and benchmarking single and two qubit gates in superconducting transmons. In the single qubit case, we reached a gate fidelity of 99.96% and found that the remaining error is currently limited by energy relaxation. We also carefully measured the contribution of leakage to single qubit gate error, and found that leakage errors are currently at the 1×10^{-5} level and are dominated by thermal excitation. For the two qubit CZ, we reached a gate fidelity of 99.45%, and found that the error is a combination of decoherence and control errors, though more research is needed to determine the exact error contributions. These gate fidelities are beyond the fault tolerance threshold for surface code error correction, and with some improvement, are also sufficient to perform a random circuit experiment to demonstrate quantum supremacy.

However, we also saw that the dispersive readout scheme in transmons is fraught with

potential problems due to unexpected couplings within the Jaynes-Cummings ladder. These effects emphasize the need for further research into understanding the energy level structure of a transmon coupled to a resonator, as well as improving the efficiency of the readout chain. Moving forward, readout fidelity, both at the single qubit level and scaled up to multiple qubits, is likely the limiting factor in more complex experiments with transmon qubits.

Appendix A

Measurements of Airbridge Circuits

In this appendix we look at detailed measurements of circuits made using airbridges.

A.1 Resistance and Critical Current

After the initial development of the airbridge fabrication process, we initially confirmed that the argon ion mill led to an Ohmic contact at the bridge pads by we fabricating 10 airbridges in series and measured them in a four terminal configuration. Each airbridge had a width of $8\ \mu\text{m}$, a length of $28\ \mu\text{m}$, and a thickness of $300\ \text{nm}$. At room temperature, we measured a resistance of $6\ \Omega$. For a standard aluminum resistivity of $2.7 \times 10^{-6}\ \Omega\text{-cm}$, the expected resistance at room temperature is $3.15\ \Omega$, which does not take into account the curvature of the bridges and the distance between the pads of the bridges, which was 6 microns.

Next, we cooled down the airbridges in an adiabatic demagnetization refrigerator




Figure A.1: (Color online): Inset: Ten airbridges fabricated in series for a four terminal measurement of the resistance. Main panel: Critical current as a function of reduced temperature T/T_c . The fit is to Eqn. 1, with $I_0 = 462$ mA

(ADR) to look at the properties when superconducting. At 100 mK, we were limited to 10 mA of drive current, which was not enough to drive the airbridges normal. Instead, we slowly cooled the sample through the critical temperature T_c and measured the critical current I_c of the airbridges as a function of temperature just below T_c , with the results shown in A.1. The critical temperatures for both the base wiring and the airbridge layer were within 50 mK of each other, and were around 1.2 K. The critical current data matches the expected Ginsburg-Landau behavior, which predicts the following relation for the critical current of a thin superconducting wire [?]

$$I_c = I_0 (1 - T/T_c)^{2/3} \quad (\text{A.1})$$

where I_0 is the critical current at temperatures well below T_c . By fitting to this equation, we extracted a low temperature critical current of 462 mA. However, this result does not take into account the width of our airbridges. From previous works, we estimate that there is a decrease in I_0 by a factor of order 3 or 4 for an $8 \mu\text{m}$ wire,[151, 152] giving a

critical current of around 100 mA.

A.2 Resonators With Airbridge Groundplane Shunts

To measure the loss added by placing an airbridge over a CPW transmission line, we constructed quarter wavelength CPW resonators with variable numbers of airbridges connecting the ground planes. We fabricated the resonators from an aluminum film deposited on a high-resistivity silicon substrate and etched with a BCl_3/Cl_2 inductively coupled plasma.[80] We have also used the process on a sapphire substrate with comparable results. We designed the resonators with $10\ \mu\text{m}$ center traces and $5\ \mu\text{m}$ gaps to match the dimensions of our typical feedlines, and designed the resonant frequencies to range from 5 to 6 GHz. We designed the airbridges to have $4\ \mu\text{m}$ of clearance from the CPW line for a total length of $28\ \mu\text{m}$, and chose an airbridge width of $8\ \mu\text{m}$ to ensure mechanical stability of the bridge. On eight of the resonators, we fabricated between 12 and 110 airbridges spanning the resonator center trace, evenly spaced in the number of bridges. The resonators with the most airbridges had a density of one airbridge every $50\ \mu\text{m}$, covering 16% of the length of the resonator. The remaining two resonators went through the full fabrication process but were not designed with any airbridges spanning their center traces. We used these witness resonators as a test of whether placing airbridges on a CPW line adds loss to other circuit elements on the same wafer. We also fabricated a separate chip of resonators from the same film that saw no airbridge processing, to act as a control sample.




Figure A.2: (Color online): (a) Resonant frequencies for resonators with variable numbers of airbridges in red squares, compared with the frequencies of their corresponding controls which are designed to have the same length. As the number of bridges increases, the resonators shift lower in frequency compared to their controls. (b) Percent change in LC , the product of the inductance per length and capacitance per length, as a function of the percentage of the resonator covered by airbridges. The dashed blue line is a linear fit to the data, with slope 12.7% and intercept 0.35%. The offset from the origin is within normal chip to chip variations in our measured resonators. The red line is a prediction based on the additional capacitance of the airbridge. The slopes differ due to the decrease in inductance from the airbridge.

A.2.1 Shift in Resonant Frequency

Compared to more conventional crossovers which are supported by dielectrics, airbridges have a much smaller impact on the capacitance of a CPW line. However, this additional capacitance due to an airbridge is not negligible and should be accounted for. For example, in our experiment to test the microwave loss of airbridges using ten different resonators, we designed the resonators such that the density of airbridges increased with decreasing frequency, as shown in Fig. A.2(a). A higher density of airbridges increases the capacitance of the resonator and decreases the resonant frequency. Thus, in our experiment, the resonant frequencies shifted further apart rather than closer together,

avoiding any frequency collisions. We note here that from our control data, we found no significant correlation of the high or low power quality factor with the frequency of the resonator over the range we considered, which validated this particular design choice.

If we assume the airbridge acts like a parallel plate capacitor between the center trace and ground, we can estimate the additional capacitance per unit length due to the airbridge as $C = \epsilon_0 w/d$, where w is the width of the center trace and d is the height of the airbridge. For the geometry in our experiment, $w = 10 \mu\text{m}$ and $d = 3 \mu\text{m}$, giving $C = 29.5 \text{ pF/m}$. We can also numerically calculate the additional capacitance due to the airbridge using COMSOL. We simulated the cross-section of a CPW line with a $10 \mu\text{m}$ center trace and $5 \mu\text{m}$ gap with a substrate dielectric constant of 11.6, and found the capacitance per length to be 175.25 pF/m . After adding an airbridge, the capacitance increased to 204.03 pF/m giving an increase of 28.78 pF/m due to the airbridge, showing remarkable agreement with the parallel plate estimate. From these values, we predict that the capacitance of a resonator covered completely by airbridges should increase by 17%.

From the frequency data shown in Fig. A.2(a), we can determine the actual effect of placing an airbridge over a CPW line. As the number of airbridges increased, the frequency of the resonator shifted further below the frequency of its corresponding control. Since each resonator and its control are designed to have the same wavelength, we can interpret the change in frequency as a change in the phase velocity of light $v_p = 1/\sqrt{LC}$, where L and C are the inductance and capacitance per unit length. Given the total length

of the resonator and the number of airbridges, we can also determine the percentage of the line covered by airbridges. The percent coverage should be linearly related to the change in the product of the inductance and capacitance per unit length, LC , which is shown in Fig. A.2(b). The slope of the linear fit in Fig. A.2(b) indicates that the LC product for a section of line covered by airbridge differs from the bare line by 12.7%.

The discrepancy between our prediction and our data is most likely due to changes in the inductance of the resonator. Each airbridge adds additional pathways for current to flow, which decreases the inductance of the CPW line and compensates in part for the increase in capacitance. However, the inductance is not as easily modeled as the capacitance since edge effects are important. In other words, a single, wide airbridge that spanning a CPW line does not have the same effect as multiple narrower airbridges because they contain different current paths.

A.2.2 Loss

In order to determine the internal quality factor Q_i of the resonators, we measured the transmission through a feedline that was capacitively coupled to each resonator (see Ref. [80] for measurement details). We varied the drive power such that the photon population $\langle n_p \rangle$ in the resonator ranged from single photon levels up to 10^7 photons, at which point the resonators became non-linear. A sample of representative quality factor data for some of the resonators is shown in the main panel of Fig. A.3. As seen in previous work, the quality factor of the resonators increases as a function of increasing drive power, which




Figure A.3: (Color online) Main panel: Dependence of the internal quality factor Q_i on the average photon population in the resonator $\langle n_p \rangle$, shown for various numbers of airbridges N_B spanning the ground planes of the resonator. Data for a control resonator and a witness resonator are also displayed. Lines are guides for the eye. Insets: (a) Total loss tangent δ at single photon as a function of number of airbridges. The best fit line has slope 1.2×10^{-8} . (b) Loss at $\langle n_p \rangle = 5 \times 10^6$, with slope 3.8×10^{-9} .

is consistent with the loss in the resonator being dominated by two-level states (TLS) at the material interfaces.[80? , 110?]

In general, we expect the dependence of the quality factor to show two plateaus, one around single photon levels corresponding to loss being dominated by TLS, and one at high power corresponding to saturation of TLS.[80?] From our control chip, for which a representative example is shown in black squares in Fig. A.3, we determined that the nominal internal quality factor for our resonator geometry and material was around 7.0×10^5 at single photon powers and 3.5×10^6 at high power (5 million photons). Our witness resonators, for which a representative example is shown in Fig. 3, performed similarly, indicating that the additional processing on the chip did not add

any loss. As we increased the number of bridges fabricated over the resonator, the quality factor decreased at both low and high powers. Interestingly, the quality factor does not appear to plateau as strongly at high power when airbridges are present.

To determine quantitatively the loss due to airbridges, we extracted the loss tangent $\delta = 1/Q_i$ at powers around a single photon and at 5 million photons. In Fig. A.3a (A.3b), we plot the low power (high power) loss tangent as a function of number of airbridges, along with lines of best fit. From the slopes, we estimate that each airbridge adds 1.2×10^{-8} to the loss tangent of the resonator at low power and 3.8×10^{-9} at high power. We can also estimate the loss per fraction of the resonator that is covered by the resonator. If we assume that the loss also scales with the width of the airbridges, then every one percent of the resonator covered by airbridges adds an additional 8.3×10^{-8} to the loss tangent at low power and 2.7×10^{-8} at high power. We note that a resonator completely covered by airbridges would be limited to an internal quality factor of order 120,000 at low power, which is more than five times lower than the uncovered device.

A.3 Participation Ratio of the Airbridge Interface

To understand this increase in the loss due to airbridges, we note that the change in the loss at different drive powers suggests that the addition of an airbridge adds to the TLS loss of the resonator. The interface underneath the airbridge is a potential source of loss, since this is the interface at which we deposited aluminum on photoresist that has been crosslinked by the argon ion mill. To understand the additional surface loss due to

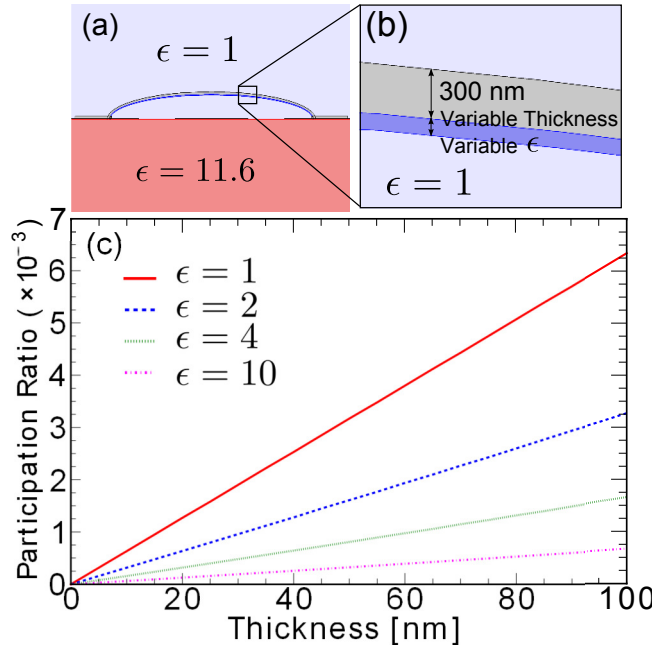


Figure A.4: (Color online): (a) Cross section of an airbridge spanning a CPW line. (b) Close-up of the interface for which we calculate the participation ratio. This interface is a possible source of loss because the layer of aluminum is deposited on photoresist that has been crosslinked by an ion mill. The thickness and dielectric constant are variable. (c) Participation ratio as a function of thickness for various dielectric constants at the interface. We numerically calculate using COMSOL the participation ratio by setting the potential of the center trace of the CPW to 1V, solving for the electric fields, then numerically computing the integral in Eq. 1 in the interface region. We obtain the total energy W by performing the same integral for all of the cross section.

this interface, we calculate the participation ratio of a lossy dielectric at this metal-air interface following Ref [78]. We consider the resonator and airbridge structure in cross section as shown in Fig. A.4(a). The participation ratio p of any isotropic region of space in this cross-section is simply given by the ratio of energy stored in the region to the total energy stored in the entire cross-section

$$p = W^{-1} \epsilon_r \epsilon_0 \iint dA \frac{|E|^2}{2} \quad (\text{A.2})$$

where W is the total energy in the cross-section which may be obtained by performing the same integral over all space, and ϵ_r is the dielectric constant in the region. Assuming that the region is thin, as it is in the case of our interface of interest, we can replace an integral over the thickness by a product, turning the double integral into a line integral over the boundary of the interface. We can also simplify the equation using the boundary conditions on our interface. The metal boundary allows us to approximate the electric field as normal to the metal, while the continuity of the displacement field at the air interface gives us the relation $\epsilon_r E_{i\perp} = E_{a\perp}$, where E_i is the electric field in the interface and E_a is the electric field in air. Combining these simplifications we obtain

$$p = W^{-1} t_i \epsilon_r^{-1} \epsilon_0 \int dS \frac{|E_{a\perp}|^2}{2} \quad (\text{A.3})$$

where t_i is the small thickness of the interface. Assuming the contribution to the total energy W of the interface is small, the participation ratio is proportional to the thickness and inversely proportional to the dielectric constant. We can estimate the value of the line integral by again modeling the airbridge as a parallel plate. If we assume a 1 V difference in potential between center trace and ground, then from the calculation of

total capacitance above, we know the value of $W = \frac{1}{2}CV^2$. The electric field is given by 1 V divided by the separation distance of $3\ \mu\text{m}$, and we may replace the integral with a multiplication by the length, about $10\ \mu\text{m}$. We then obtain the following approximate formula:

$$p = 4.8 \times 10^{-5} \text{ nm}^{-1} \frac{t_i}{\epsilon_r} \quad (\text{A.4})$$

Alternatively, we can also numerically evaluate Eq. A.2. We constructed the geometry of an airbridge spanning a CPW and included a thin dielectric interface on the underside of the bridge as shown in Fig. A.4(b). After applying a potential of 1 V to the center trace, we solved for the electric fields and numerically integrated Eq. A.2 to determine the total energy in the cross-section and the energy in the interface, giving us the participation ratio. We calculated participation ratio as a function of interface thickness and dielectric constant, producing the plot shown in Fig. A.4(c). We see that the scaling follows the expected scaling from Eqs. A.3 and A.4. Furthermore, we can more accurately determine the coefficient in Eq. A.4 from the slopes of the lines, and we find that the coefficient is $6.34 \times 10^{-5} \text{ nm}^{-1}$, which is within 30% of our parallel plate estimation.

Given the participation ratio, we can estimate the loss due to this interface. For the dielectric constant, we estimate a dielectric constant of 4 based on data pertaining to other photoresists [?]. SEM images of the interface were inconclusive for determining the thickness, but it is certainly upper bounded by 100 nm. Finally, there is little data on the loss tangent of resist and cryogenic temperatures, so we estimate this to be 10^{-3} based on the measured loss tangents of amorphous oxides.[75] Using these numbers, we




Figure A.5: (Color online): Insets: (a) Airbridges connecting together CPW lines within a resonator. (b) A second style of airbridge connection, where the ground plane is threaded underneath the airbridge. Main panel: Internal quality factor of resonators as a function of average photon population for many different styles of resonators. We show two witness resonators to demonstrate the typical spread in measured Q_i . Lines are guides for the eye.

obtain a participation ratio of 1.6×10^{-3} and a loss due to airbridges of 1.6×10^{-6} , or a Q_i of 630,000.

A.3.1 Center Trace Airbridges

In addition to placing airbridges over CPW lines to connect the ground planes together, we have also fabricated airbridges to connect two CPW center traces together. We tested such a connection by fabricating quarter wavelength CPW resonators with intentional breaks in the resonator, then reconnecting the lines with airbridges. The resonators fabricated using this method performed comparably to resonators fabricated with ground plane airbridges, with each airbridge connection added to the center trace adding a loss

of 1×10^{-7} . These results indicate that airbridges can be used to carry microwave signals, and for example, allowing for the crossing of two perpendicular superconducting CPW lines.

Bibliography

- [1] John Bardeen and Walter Houser Brattain. Physical principles involved in transistor action. *Physical Review*, 75(8):1208, 1949.
- [2] L Britnell, RV Gorbachev, R Jalil, BD Belle, F Schedin, A Mishchenko, T Georgiou, MI Katsnelson, L Eaves, SV Morozov, et al. Field-effect tunneling transistor based on vertical graphene heterostructures. *Science*, 335(6071):947–950, 2012.
- [3] Claude Weisbuch and Borge Vinter. *Quantum semiconductor structures: fundamentals and applications*. Academic press, 2014.
- [4] J Mathon and A Umerski. Theory of tunneling magnetoresistance of an epitaxial fe/mgo/fe (001) junction. *Physical Review B*, 63(22):220403, 2001.
- [5] Felix Bloch. Nuclear induction. *Physical review*, 70(7-8):460, 1946.
- [6] Daniel Gottesman. The heisenberg representation of quantum computers. *arXiv preprint quant-ph/9807006*, 1998.
- [7] Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Physical Review A*, 70(5):052328, 2004.
- [8] Peter W Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on*, pages 124–134. Ieee, 1994.
- [9] Michael A Nielsen and Isaac Chuang. *Quantum computation and quantum information*, 2002.
- [10] Lov K Grover. Quantum mechanics helps in searching for a needle in a haystack. *Physical review letters*, 79(2):325, 1997.
- [11] Richard P Feynman. Simulating physics with computers. *International journal of theoretical physics*, 21(6):467–488, 1982.

- [12] Fumiko Yamaguchi and Yoshihisa Yamamoto. Quantum simulation of the t-j model. *Superlattices and Microstructures*, 32(4):343–345, 2002.
- [13] Efstratios Manousakis. A quantum-dot array as model for copper-oxide superconductors: A dedicated quantum simulator for the many-fermion problem. *Journal of low temperature physics*, 126(5-6):1501–1513, 2002.
- [14] Jonathan Simon, Waseem S Bakr, Ruichao Ma, M Eric Tai, Philipp M Preiss, and Markus Greiner. Quantum simulation of antiferromagnetic spin chains in an optical lattice. *Nature*, 472(7343):307–312, 2011.
- [15] R Islam, EE Edwards, K Kim, S Korenblit, C Noh, H Carmichael, G-D Lin, L-M Duan, C-C Joseph Wang, JK Freericks, et al. Onset of a quantum phase transition with a trapped ion quantum simulator. *arXiv preprint arXiv:1103.2400*, 2011.
- [16] Kristian Baumann, Christine Guerlin, Ferdinand Brennecke, and Tilman Esslinger. Dicke quantum phase transition with a superfluid gas in an optical cavity. *Nature*, 464(7293):1301–1306, 2010.
- [17] Benjamin P Lanyon, James D Whitfield, Geoff G Gillett, Michael E Goggin, Marcelo P Almeida, Ivan Kassal, Jacob D Biamonte, Masoud Mohseni, Ben J Powell, Marco Barbieri, et al. Towards quantum chemistry on a quantum computer. *Nature Chemistry*, 2(2):106–111, 2010.
- [18] Alán Aspuru-Guzik, Anthony D Dutoi, Peter J Love, and Martin Head-Gordon. Simulated quantum computation of molecular energies. *Science*, 309(5741):1704–1707, 2005.
- [19] PJJ O’Malley, Ryan Babbush, ID Kivlichan, Jonathan Romero, JR McClean, Rami Barends, Julian Kelly, Pedram Roushan, Andrew Tranter, Nan Ding, et al. Scalable quantum simulation of molecular energies. *Physical Review X*, 6(3):031007, 2016.
- [20] S Giovanazzi. Hawking radiation in sonic black holes. *Physical review letters*, 94(6):061302, 2005.
- [21] Katherine L Brown, William J Munro, and Vivien M Kendon. Using quantum computers for quantum simulation. *Entropy*, 12(11):2268–2307, 2010.
- [22] William K Wootters and Wojciech H Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982.
- [23] Peter W Shor. Scheme for reducing decoherence in quantum computer memory. *Physical review A*, 52(4):R2493, 1995.
- [24] Austin G Fowler, Matteo Mariantoni, John M Martinis, and Andrew N Cleland. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86(3):032324, 2012.

- [25] Andrew M Steane. Error correcting codes in quantum theory. *Physical Review Letters*, 77(5):793, 1996.
- [26] Keqin Feng. Quantum error-correcting codes. *Coding Theory and Cryptology*, 1:91, 2002.
- [27] A Robert Calderbank, Eric M Rains, Peter W Shor, and Neil JA Sloane. Quantum error correction and orthogonal geometry. *Physical Review Letters*, 78(3):405, 1997.
- [28] Daniel Gottesman. Class of quantum error-correcting codes saturating the quantum hamming bound. *Physical Review A*, 54(3):1862, 1996.
- [29] Emanuel Knill, Raymond Laflamme, and Lorenza Viola. Theory of quantum error correction for general noise. *Physical Review Letters*, 84(11):2525, 2000.
- [30] Austin G Fowler, Ashley M Stephens, and Peter Groszkowski. High-threshold universal quantum computation on the surface code. *Physical Review A*, 80(5):052312, 2009.
- [31] Sergey B Bravyi and A Yu Kitaev. Quantum codes on a lattice with boundary. *arXiv preprint quant-ph/9811052*, 1998.
- [32] Sergio Boixo, Sergei V Isakov, Vadim N Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, John M Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *arXiv preprint arXiv:1608.00263*, 2016.
- [33] John Preskill. Quantum computing and the entanglement frontier. *arXiv preprint arXiv:1203.5813*, 2012.
- [34] Ludovico Latmiral, Nicolò Spagnolo, and Fabio Sciarrino. Towards quantum supremacy with lossy scattershot boson sampling. *New Journal of Physics*, 18(11):113008, 2016.
- [35] Juan Jaramillo, Mathieu Beau, and Adolfo del Campo. Quantum supremacy of many-particle thermal machines. *New Journal of Physics*, 18(7):075019, 2016.
- [36] Eugene P Wigner. On the statistical distribution of the widths and spacings of nuclear resonance levels. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 47, pages 790–798. Cambridge University Press, 1951.
- [37] Thomas Guhr, Axel Müller-Groeling, and Hans A Weidenmüller. Random-matrix theories in quantum physics: common concepts. *Physics Reports*, 299(4):189–425, 1998.
- [38] CE Porter and RG Thomas. Fluctuations of nuclear reaction widths. *Physical Review*, 104(2):483, 1956.

- [39] Michael J Bremner, Richard Jozsa, and Dan J Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, page rspa20100301. The Royal Society, 2010.
- [40] Juan I Cirac and Peter Zoller. Quantum computations with cold trapped ions. *Physical review letters*, 74(20):4091, 1995.
- [41] Dietrich Leibfried, Rainer Blatt, Christopher Monroe, and David Wineland. Quantum dynamics of single trapped ions. *Reviews of Modern Physics*, 75(1):281, 2003.
- [42] Diego Porras and J Ignacio Cirac. Effective quantum spin systems with trapped ions. *Physical review letters*, 92(20):207901, 2004.
- [43] DJ Wineland, JC Bergquist, JJ Bollinger, WM Itano, DJ Heinzen, SL Gilbert, CH Manney, and MG Raizen. Progress at nist toward absolute frequency standards using stored ions. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 37(6):515–523, 1990.
- [44] DJ Wineland. Trapped ions, laser cooling, and better clocks. *Science*, 226(4673):395, 1984.
- [45] Kenton R Brown, Andrew C Wilson, Yves Colombe, C Ospelkaus, Adam M Meier, E Knill, D Leibfried, and David J Wineland. Single-qubit-gate error below 10^{-4} in a trapped ion. *Physical Review A*, 84(3):030303, 2011.
- [46] TP Harty, DTC Allcock, C J Ballance, L Guidoni, HA Janacek, NM Linke, DN Stacey, and DM Lucas. High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit. *Physical review letters*, 113(22):220501, 2014.
- [47] Chris Monroe, DM Meekhof, BE King, Wayne M Itano, and David J Wineland. Demonstration of a fundamental quantum logic gate. *Physical review letters*, 75(25):4714, 1995.
- [48] AH Myerson, DJ Szwer, SC Webster, DTC Allcock, MJ Curtis, G Imreh, JA Sherman, DN Stacey, AM Steane, and DM Lucas. High-fidelity readout of trapped-ion qubits. *Physical Review Letters*, 100(20):200502, 2008.
- [49] Signe Seidelin, John Chiaverini, Rainer Reichle, John J Bollinger, Didi Leibfried, Joe Britton, JH Wesenberg, RB Blakestad, RJ Epstein, DB Hume, et al. Micro-fabricated surface-electrode ion trap for scalable quantum information processing. *Physical review letters*, 96(25):253003, 2006.
- [50] Dan Stick, WK Hensinger, Steven Olmschenk, MJ Madsen, Keith Schwab, and Chris Monroe. Ion trap in a semiconductor chip. *Nature Physics*, 2(1):36–39, 2006.

- [51] C Monroe, R Raussendorf, A Ruthven, KR Brown, P Maunz, L-M Duan, and J Kim. Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects. *Physical Review A*, 89(2):022317, 2014.
- [52] Bjoern Lekitsch, Sebastian Weidt, Austin G Fowler, Klaus Mølmer, Simon J Devitt, Christof Wunderlich, and Winfried K Hensinger. Blueprint for a microwave trapped ion quantum computer. *Science Advances*, 3(2):e1601540, 2017.
- [53] Daniel Loss and David P DiVincenzo. Quantum computation with quantum dots. *Physical Review A*, 57(1):120, 1998.
- [54] M Veldhorst, JCC Hwang, CH Yang, AW Leenstra, B De Ronde, JP Dehollain, JT Muhonen, FE Hudson, KM Itoh, A Morello, et al. An addressable quantum dot qubit with fault-tolerant control-fidelity. *Nature nanotechnology*, 9(12):981–985, 2014.
- [55] JM Elzerman, R Hanson, LH Willems Van Beveren, B Witkamp, LMK Vandersypen, and Leo P Kouwenhoven. Single-shot read-out of an individual electron spin in a quantum dot. *nature*, 430(6998):431–435, 2004.
- [56] NB Manson, JP Harrison, and MJ Sellars. Nitrogen-vacancy center in diamond: Model of the electronic structure and associated dynamics. *Physical Review B*, 74(10):104303, 2006.
- [57] Matthias Widmann, Sang-Yun Lee, Torsten Rendler, Nguyen Tien Son, Helmut Fedder, Seoyoung Paik, Li-Ping Yang, Nan Zhao, Sen Yang, Ian Booker, et al. Coherent control of single spins in silicon carbide at room temperature. *Nature materials*, 14(2):164–168, 2015.
- [58] William F Koehl, Bob B Buckley, F Joseph Heremans, Greg Calusine, and David D Awschalom. Room temperature coherent control of defect spin qubits in silicon carbide. *Nature*, 479(7371):84–87, 2011.
- [59] AM Tyryshkin, SA Lyon, AV Astashkin, and AM Raitsimring. Electron spin relaxation times of phosphorus donors in silicon. *Physical Review B*, 68(19):193207, 2003.
- [60] Patrick Maletinsky, Sungkun Hong, Michael Sean Grinolds, Birgit Hausmann, Mikhail D Lukin, Ronald L Walsworth, Marko Loncar, and Amir Yacoby. A robust scanning diamond sensor for nanoscale imaging with single nitrogen-vacancy centres. *Nature nanotechnology*, 7(5):320–324, 2012.
- [61] HJ Mamin, M Kim, MH Sherwood, CT Rettner, K Ohno, DD Awschalom, and D Rugar. Nanoscale nuclear magnetic resonance with a nitrogen-vacancy spin sensor. *Science*, 339(6119):557–560, 2013.

- [62] Alexios Beveratos, Rosa Brouri, Thierry Gacoin, André Villing, Jean-Philippe Poizat, and Philippe Grangier. Single photon quantum cryptography. *Physical review letters*, 89(18):187901, 2002.
- [63] Jens Koch, Terri M Yu, Jay M Gambetta, Andrew A Houck, David I Schuster, J Majer, Alexandre Blais, Michel H Devoret, Steven M Girvin, and Robert J Schoelkopf. Charge-insensitive qubit design derived from the cooper pair box. *Phys. Rev. A*, 76(4):042319, 2007.
- [64] R Barends, J Kelly, A Megrant, A Veitia, D Sank, E Jeffrey, TC White, J Mutus, AG Fowler, B Campbell, et al. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508(7497):500–503, 2014.
- [65] Sarah Sheldon, Easwar Magesan, Jerry M Chow, and Jay M Gambetta. Procedure for systematically tuning up cross-talk in the cross-resonance gate. *Physical Review A*, 93(6):060302, 2016.
- [66] Jerry M Chow, Jay M Gambetta, Easwar Magesan, David W Abraham, Andrew W Cross, BR Johnson, Nicholas A Masluk, Colm A Ryan, John A Smolin, Srikanth J Srinivasan, et al. Implementing a strand of a scalable fault-tolerant quantum computing fabric. *Nature communications*, 5, 2014.
- [67] J Kelly, R Barends, AG Fowler, A Megrant, E Jeffrey, T.C. White, D Sank, JY Mutus, B Campbell, Yu Chen, et al. State preservation by repetitive error detection in a superconducting quantum circuit. *Nature*, 519(7541):66–69, 2015.
- [68] Daniel Thomas Sank. *Fast, accurate state measurement in superconducting qubits*. PhD thesis, University of California, Santa Barbara, 2014.
- [69] Steven M Girvin. Circuit qed: Superconducting qubits coupled to microwave photons.
- [70] Michel H Devoret et al. Quantum fluctuations in electrical circuits.
- [71] Brian David Josephson. Possible new effects in superconductive tunnelling. *Physics letters*, 1(7):251–253, 1962.
- [72] Edwin T Jaynes and Frederick W Cummings. Comparison of quantum and semi-classical radiation theories with application to the beam maser. *Proceedings of the IEEE*, 51(1):89–109, 1963.
- [73] Willis E Lamb Jr and Robert C Retherford. Fine structure of the hydrogen atom by a microwave method. *Physical Review*, 72(3):241, 1947.
- [74] Evan Jeffrey, Daniel Sank, JY Mutus, TC White, J Kelly, R Barends, Y Chen, Z Chen, B Chiaro, A Dunsworth, et al. Fast accurate state measurement with superconducting qubits. *Phys. Rev. Lett.*, 112(19):190504, 2014.

- [75] Aaron D O’Connell, M Ansmann, Radoslaw C Bialczak, Max Hofheinz, Nadav Katz, Erik Lucero, C McKenney, Matthew Neeley, Haohua Wang, Eva M Weig, et al. Microwave dielectric loss at single photon energies and millikelvin temperatures. *Applied Physics Letters*, 92(11):112903, 2008.
- [76] Robert McDermott. Materials origins of decoherence in superconducting qubits. *IEEE Transactions on Applied Superconductivity*, 19(1):2–13, 2009.
- [77] Cheng P Wen. Coplanar waveguide: A surface strip transmission line suitable for nonreciprocal gyromagnetic device applications. *IEEE Transactions on Microwave Theory and Techniques*, 17(12):1087–1090, 1969.
- [78] J Wenner, R Barends, RC Bialczak, Yu Chen, J Kelly, Erik Lucero, Matteo Mariantoni, A Megrant, PJJ OMalley, D Sank, et al. Surface loss simulations of superconducting coplanar waveguide resonators. *Applied Physics Letters*, 99(11):113513, 2011.
- [79] R Barends, J Kelly, A Megrant, D Sank, E Jeffrey, Y Chen, Y Yin, B Chiaro, J Mutus, C Neill, et al. Coherent josephson qubit suitable for scalable quantum integrated circuits. *Phys. Rev. Lett.*, 111(8):080502, 2013.
- [80] A Megrant, C Neill, R Barends, B Chiaro, Yu Chen, L Feigl, J Kelly, Erik Lucero, Matteo Mariantoni, PJJ O’Malley, et al. Planar superconducting resonators with internal quality factors above one million. *Appl. Phys. Lett.*, 100(11):113510, 2012.
- [81] Josephine B Chang, Michael R Vissers, Antonio D Córcoles, Martin Sandberg, Jiansong Gao, David W Abraham, Jerry M Chow, Jay M Gambetta, Mary Beth Rothwell, George A Keefe, et al. Improved superconducting qubit coherence using titanium nitride. *Applied Physics Letters*, 103(1):012602, 2013.
- [82] Michael R Vissers, Jeffrey S Kline, Jiansong Gao, David S Wisbey, and David P Pappas. Reduced microwave loss in trenched superconducting coplanar waveguides. *Applied Physics Letters*, 100(8):082602, 2012.
- [83] A Bruno, G De Lange, S Asaad, KL Van der Enden, NK Langford, and L DiCarlo. Reducing intrinsic loss in superconducting resonators by surface treatment and deep etching of silicon substrates. *Applied Physics Letters*, 106(18):182601, 2015.
- [84] CM Quintana, A Megrant, Z Chen, A Dunsworth, B Chiaro, R Barends, B Campbell, Yu Chen, I-C Hoi, E Jeffrey, et al. Characterization and reduction of microfabrication-induced decoherence in superconducting quantum circuits. *Applied Physics Letters*, 105(6):062601, 2014.
- [85] Kirt R Williams, Kishan Gupta, and Matthew Wasilik. Etch rates for micromachining processing-part ii. *Journal of microelectromechanical systems*, 12(6):761–778, 2003.

- [86] GJ Dolan. Offset masks for lift-off photoprocessing. *Applied Physics Letters*, 31(5):337–339, 1977.
- [87] A Dunsworth, A Megrant, C Quintana, Zijun Chen, R Barends, B Burkett, B Foxen, Yu Chen, B Chiaro, A Fowler, et al. Characterization and reduction of capacitive loss induced by sub-micron josephson junction fabrication in superconducting qubits. *arXiv preprint arXiv:1706.00879*, 2017.
- [88] N Simons Rainee. Coplanar waveguide circuits, components, and systems, 2001.
- [89] George E Ponchak, John Papapolymerou, and Manos M Tentzeris. Excitation of coupled slotline mode in finite-ground cpw with unequal ground-plane widths. *IEEE transactions on microwave theory and techniques*, 53(2):713–717, 2005.
- [90] William P Harokopus and Linda PB Katehi. Electromagnetic coupling and radiation loss considerations in microstrip (m) mic design. *IEEE Transactions on Microwave Theory and Techniques*, 39(3):413–421, 1991.
- [91] AA Houck, JA Schreier, BR Johnson, JM Chow, Jens Koch, JM Gambetta, DI Schuster, L Frunzio, MH Devoret, SM Girvin, et al. Controlling the spontaneous emission of a superconducting transmon qubit. *Physical review letters*, 101(8):080502, 2008.
- [92] EB Rosa, FW Grover, E Buckingham, NE Dorsey, JM Miller, HB Brooks, CE Waters, JB Tuttle, HL Curtis, PG Nutting, et al. Bulletin of the bureau of standards, 1908.
- [93] James Wenner, Matthew Neeley, Radoslaw C Bialczak, Michael Lenander, Erik Lucero, Aaron D OConnell, Daniel Sank, Haohua Wang, Martin Weides, Andrew N Cleland, et al. Wirebond crosstalk and cavity modes in large chip mounts for superconducting qubits. *Superconductor Science and Technology*, 24(6):065001, 2011.
- [94] Y JY Lankwarden, A Endo, J JA Baselmans, and MP Bruijn. Development of nbtin-al direct antenna coupled kinetic inductance detectors. *Journal of Low Temperature Physics*, 167(3):367–372, 2012.
- [95] Y Okuyama, T Hashimoto, and T Koguchi. High dose ion implantation into photoresist. *Journal of The Electrochemical Society*, 125(8):1293–1298, 1978.
- [96] Zijun Chen, Anthony Megrant, Julian Kelly, Rami Barends, Joerg Bochmann, Yu Chen, Ben Chiaro, Andrew Dunsworth, Evan Jeffrey, JY Mutus, et al. Fabrication and characterization of aluminum airbridges for superconducting microwave circuits. *Applied Physics Letters*, 104(5):052602, 2014.

- [97] P Das, R Bruyn de Ouboter, and KW Taconis. A realization of a london-clarke-mendoza type refrigerator. In *Low Temperature Physics LT9*, pages 1253–1255. Springer, 1965.
- [98] Christopher M Quintana. *Superconducting flux qubits for high-connectivity quantum annealing without lossy dielectrics*. PhD thesis, University of California, Santa Barbara, 2017.
- [99] JY Mutus, TC White, Rami Barends, Yu Chen, Zijun Chen, Ben Chiaro, Andrew Dunsworth, Evan Jeffrey, Julian Kelly, Anthony Megrant, et al. Strong environmental coupling in a josephson parametric amplifier. *Appl. Phys. Lett.*, 104(26):263513, 2014.
- [100] F Motzoi, Jay M Gambetta, P Rebentrost, and Frank K Wilhelm. Simple pulses for elimination of leakage in weakly nonlinear qubits. *Phys. Rev. Lett.*, 103(11):110501, 2009.
- [101] Andrew M Childs, Isaac L Chuang, and Debbie W Leung. Realization of quantum process tomography in nmr. *Physical Review A*, 64(1):012314, 2001.
- [102] MW Mitchell, CW Ellenor, S Schneider, and AM Steinberg. Diagnosis, prescription, and prognosis of a bell-state filter by quantum process tomography. *Physical review letters*, 91(12):120402, 2003.
- [103] Jeremy L O’Brien, GJ Pryde, Alexei Gilchrist, DFV James, Nathan K Langford, TC Ralph, and AG White. Quantum process tomography of a controlled-not gate. *Physical review letters*, 93(8):080502, 2004.
- [104] Joseph Emerson, Robert Alicki, and Karol Życzkowski. Scalable noise estimation with random unitary operators. *Journal of Optics B: Quantum and Semiclassical Optics*, 7(10):S347, 2005.
- [105] E Knill, D Leibfried, R Reichle, J Britton, R B Blakestad, John D Jost, C Langer, R Ozeri, S Seidelin, and David J Wineland. Randomized benchmarking of quantum gates. *Physical Review A*, 77(1):012307, 2008.
- [106] Christoph Dankert, Richard Cleve, Joseph Emerson, and Etera Livine. Exact and approximate unitary 2-designs and their application to fidelity estimation. *Physical Review A*, 80(1):012304, 2009.
- [107] Easwar Magesan, Jay M Gambetta, and Joseph Emerson. Scalable and robust randomized benchmarking of quantum processes. *Phys. Rev. Lett.*, 106(18):180504, 2011.
- [108] Joel Wallman, Chris Granade, Robin Harper, and Steven T Flammia. Estimating the coherence of noise. *New Journal of Physics*, 17(11):113020, 2015.

- [109] Guanru Feng, Joel J Wallman, Brandon Buonacorsi, Franklin H Cho, Daniel K Park, Tao Xin, Dawei Lu, Jonathan Baugh, and Raymond Laflamme. Estimating the coherence of noise in quantum control of a solid-state qubit. *Physical review letters*, 117(26):260501, 2016.
- [110] John M Martinis, Ken B Cooper, Robert McDermott, Matthias Steffen, Markus Ansmann, KD Osborn, Katarina Cicak, Seongshik Oh, David P Pappas, Raymond W Simmonds, et al. Decoherence in josephson qubits from dielectric loss. *Physical Review Letters*, 95(21):210503, 2005.
- [111] C Neill, A Megrant, R Barends, Yu Chen, B Chiaro, J Kelly, JY Mutus, PJJ O’Malley, D Sank, J Wenner, et al. Fluctuations from edge defects in superconducting resonators. *Applied Physics Letters*, 103(7):072601, 2013.
- [112] Norman F Ramsey. Experiments with separated oscillatory fields and hydrogen masers. *Reviews of modern physics*, 62(3):541, 1990.
- [113] Daniel Sank, R Barends, Radoslaw C Bialczak, Yu Chen, J Kelly, M Lenander, E Lucero, Matteo Mariantoni, A Megrant, M Neeley, et al. Flux noise probed with real time qubit tomography in a josephson phase qubit. *Physical review letters*, 109(6):067001, 2012.
- [114] Erwin L Hahn. Spin echoes. *Physical review*, 80(4):580, 1950.
- [115] EL Hahn and DE Maxwell. Spin echo measurements of nuclear spin coupling in molecules. *Physical Review*, 88(5):1070, 1952.
- [116] Saul Meiboom and David Gill. Modified spin-echo method for measuring nuclear relaxation times. *Review of scientific instruments*, 29(8):688–691, 1958.
- [117] Roger H Koch, John Clarke, WM Goubau, JM Martinis, CM Pegrum, and DJ Van Harlingen. Flicker (1/f) noise in tunnel junction dc squids. *Journal of Low Temperature Physics*, 51(1-2):207–224, 1983.
- [118] Radoslaw C Bialczak, R McDermott, M Ansmann, M Hofheinz, N Katz, Erik Lucero, Matthew Neeley, AD OConnell, H Wang, AN Cleland, et al. 1/f flux noise in josephson phase qubits. *Physical review letters*, 99(18):187006, 2007.
- [119] P Kumar, S Sendelbach, MA Beck, JW Freeland, Zhe Wang, Hui Wang, C Yu Clare, RQ Wu, DP Pappas, and R McDermott. Origin and reduction of 1/f magnetic flux noise in superconducting devices. *Physical Review Applied*, 6(4):041001, 2016.
- [120] PJJ O’Malley, J Kelly, R Barends, B Campbell, Y Chen, Z Chen, B Chiaro, A Dunsworth, AG Fowler, I-C Hoi, et al. Qubit metrology of ultralow phase noise using randomized benchmarking. *Phys. Rev. Applied*, 3(4):044009, 2015.

- [121] David C McKay, Christopher J Wood, Sarah Sheldon, Jerry M Chow, and Jay M Gambetta. Efficient z gates for quantum computing. *Physical Review A*, 96(2):022330, 2017.
- [122] MA Rol, CC Bultink, TE OBrien, SR de Jong, LS Theis, X Fu, F Luthi, RFL Vermeulen, JC de Sterke, A Bruno, et al. Restless tuneup of high-fidelity qubit gates. *Physical Review Applied*, 7(4):041001, 2017.
- [123] John M Martinis and Michael R Geller. Fast adiabatic qubit gates using only σ z control. *Phys. Rev. A*, 90(2):022307, 2014.
- [124] Erik Lucero, Julian Kelly, Radoslaw C Bialczak, Mike Lenander, Matteo Mariantoni, Matthew Neeley, AD O’Connell, Daniel Sank, H Wang, Martin Weides, et al. Reduced phase error through optimized control of a superconducting qubit. *Phys. Rev. A*, 82(4):042339, 2010.
- [125] Jerry M Chow, L DiCarlo, Jay M Gambetta, F Motzoi, L Frunzio, Steven M Girvin, and Robert J Schoelkopf. Optimized driving of superconducting artificial atoms for improved single-qubit gates. *Phys. Rev. A*, 82(4):040305, 2010.
- [126] Sarah Sheldon, Lev S Bishop, Easwar Magesan, Stefan Filipp, Jerry M Chow, and Jay M Gambetta. Characterizing errors on qubit operations via iterative randomized benchmarking. *arXiv:1504.06597*, 2015.
- [127] Austin G Fowler. Coping with qubit leakage in topological codes. *Phys. Rev. A*, 88(4):042308, 2013.
- [128] Joydip Ghosh, Austin G Fowler, John M Martinis, and Michael R Geller. Understanding the effects of leakage in superconducting quantum-error-detection circuits. *Phys. Rev. A*, 88(6):062329, 2013.
- [129] Adam Sears, Jeffrey Birenbaum, David Hover, Theodore Gudmundsen, Andrew Kerman, Paul Welander, Jonilyn L Yoder, Simon Gustavsson, Xiaoyue Jin, Archana Kamal, et al. Design and measurement of improved capacitively-shunted flux qubits. In *APS Meeting Abstracts*, volume 1, page 28005, 2014.
- [130] Felix Motzoi and Frank K Wilhelm. Improving frequency selection of driven pulses using derivative-based transition suppression. *Physical Review A*, 88(6):062318, 2013.
- [131] Jay M Gambetta, F Motzoi, S T Merkel, and Frank K Wilhelm. Analytic control methods for high-fidelity unitary operations in a weakly nonlinear oscillator. *Phys. Rev. A*, 83(1):012308, 2011.
- [132] Amrit De. Fast quantum control for weakly nonlinear qubits: On two-quadrature adiabatic gates. *arXiv:1509.07905*, 2015.

- [133] Simon Gustavsson, Olger Zwiier, Jonas Bylander, Fei Yan, Fumiki Yoshihara, Yasunobu Nakamura, Terry P Orlando, and William D Oliver. Improving quantum gate fidelities by using a qubit to measure microwave pulse distortions. *Phys. Rev. Lett.*, 110(4):040502, 2013.
- [134] J Kelly, R Barends, B Campbell, Y Chen, Z Chen, B Chiaro, A Dunsworth, AG Fowler, I-C Hoi, E Jeffrey, et al. Optimal quantum control using randomized benchmarking. *Phys. Rev. Lett.*, 112(24):240504, 2014.
- [135] Michael J Peterer, Samuel J Bader, Xiaoyue Jin, Fei Yan, Archana Kamal, Theodore J Gudmundsen, Peter J Leek, Terry P Orlando, William D Oliver, and Simon Gustavsson. Coherence and decay of higher energy levels of a superconducting transmon qubit. *Phys. Rev. Lett.*, 114(1):010501, 2015.
- [136] X. Y. Jin, A. Kamal, A. P. Sears, T. Gudmundsen, D. Hover, J. Miloshi, R. Slattery, F. Yan, J. Yoder, T. P. Orlando, S. Gustavsson, and W. D. Oliver. Thermal and residual excited-state population in a 3d transmon qubit. *Phys. Rev. Lett.*, 114:240501, Jun 2015.
- [137] Zijun Chen, Julian Kelly, Chris Quintana, R Barends, B Campbell, Yu Chen, B Chiaro, A Dunsworth, AG Fowler, E Lucero, et al. Measuring and suppressing quantum state leakage in a superconducting qubit. *Physical review letters*, 116(2):020501, 2016.
- [138] Frederick W Strauch, Philip R Johnson, Alex J Dragt, CJ Lobb, JR Anderson, and FC Wellstood. Quantum logic gates for coupled superconducting phase qubits. *Physical review letters*, 91(16):167005, 2003.
- [139] L DiCarlo, JM Chow, JM Gambetta, Lev S Bishop, BR Johnson, DI Schuster, J Majer, A Blais, L Frunzio, SM Girvin, et al. Demonstration of two-qubit algorithms with a superconducting quantum processor. *Nature*, 460(7252):240–244, 2009.
- [140] Leonardo DiCarlo, Matthew D Reed, Luyan Sun, Blake R Johnson, Jerry M Chow, Jay M Gambetta, Luigi Frunzio, Steven M Girvin, Michel H Devoret, and Robert J Schoelkopf. Preparation and measurement of three-qubit entanglement in a superconducting circuit. *Nature*, 467(7315):574–578, 2010.
- [141] CJ Ballance, TP Harty, NM Linke, MA Sepiol, and DM Lucas. High-fidelity quantum logic gates using trapped-ion hyperfine qubits. *Physical review letters*, 117(6):060504, 2016.
- [142] Daniel Sank, Zijun Chen, Mostafa Khezri, J Kelly, R Barends, B Campbell, Y Chen, B Chiaro, A Dunsworth, A Fowler, et al. Measurement-induced state transitions in a superconducting qubit: Beyond the rotating wave approximation. *Physical review letters*, 117(19):190503, 2016.

- [143] M. D. Reed, L. DiCarlo, B. R. Johnson, L. Sun, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf. High-Fidelity Readout in Circuit Quantum Electrodynamics Using the Jaynes-Cummings Nonlinearity. *Phys. Rev. Lett.*, 105:173601, 2010.
- [144] J. E. Johnson, C. Macklin, D. H. Slichter, R. Vijay, E. B. Weingarten, John Clarke, and I. Siddiqi. Heralded state preparation in a superconducting qubit. *Phys. Rev. Lett.*, 109:050506, Aug 2012.
- [145] Mostafa Khezri, Eric Mlinar, Justin Dressel, and Alexander N. Korotkov. Measuring a transmon qubit in circuit QED: Dressed squeezed states. *Phys. Rev. A*, 94:012347, 2016.
- [146] MD Reed, BR Johnson, AA Houck, L DiCarlo, JM Chow, DI Schuster, L Frunzio, and RJ Schoelkopf. Fast reset and suppressing spontaneous emission of a superconducting qubit. *Appl. Phys. Lett.*, 96(20):203110–203110, 2010.
- [147] Maxime Boissonneault, J. M. Gambetta, and Alexandre Blais. Nonlinear dispersive regime of cavity qed: The dressed dephasing model. *Phys. Rev. A*, 77:060305, Jun 2008.
- [148] D. H. Slichter, R. Vijay, S. J. Weber, S. Boutin, M. Boissonneault, J. M. Gambetta, A. Blais, and I. Siddiqi. Measurement-induced qubit state mixing in circuit qed from up-converted dephasing noise. *Phys. Rev. Lett.*, 109:153601, 2012.
- [149] D. I. Schuster, A. Wallraff, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. M. Girvin, and R. J. Schoelkopf. ac stark shift and dephasing of a superconducting qubit strongly coupled to a cavity field. *Phys. Rev. Lett.*, 94:123602, Mar 2005.
- [150] Alexandre Blais, Ren-Shou Huang, Andreas Wallraff, S M Girvin, and R J Schoelkopf. Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation. *Phys. Rev. A*, 69(6):062320, 2004.
- [151] WJ Skocpol. Critical currents of superconducting microbridges. *Physical Review B*, 14(3):1045, 1976.
- [152] JOHANNES Romijn, TM Klapwijk, MJ Renne, and JE Mooij. Critical pair-breaking current in superconducting aluminum strips far below t_c . *Physical Review B*, 26(7):3648, 1982.