


# Multiqubit experiments in 3D circuit quantum electrodynamics



Jacob Blumoff

## Abstract

# Multiqubit experiments in 3D circuit quantum electrodynamics

Jacob Zachary Blumoff

2017

Circuit quantum electrodynamics (cQED) is the field of manipulating and measuring quantum electrical circuits. These circuits operate in the microwave regime, allowing use of sophisticated experimental equipment and techniques developed for industry. The nature of these devices allows for very strong interactions, providing interesting and accessible physics in the single-quantum regime. Recently, part of the field has branched from strictly lithographically designed circuits to exploit the distinctly three-dimensional electromagnetic environment. 3D components can be very high-Q compared to their 2D counterparts. This thesis explores some of the first experiments in 3D cQED to use multiple qubits, both with transmon qubits and qubits encoded in the states of harmonic oscillators. One experiment demonstrates a novel method to use a high-Q resonator to measure a register of transmon qubits in nontrivial ways. We go on to rigorously characterize these measurements. A second experiment realizes the first deterministically teleported two-qubit gate. The qubits in this experiment are encoded in harmonic oscillators. We use an entangled pair of transmons as a resource, exploiting high-fidelity measurement and real-time feedback. The tools employed in these experiments constitute a suite of capabilities necessary for increasingly complex cQED experiments.

# **Multiqubit experiments in 3D circuit quantum electrodynamics**

A Dissertation  
Presented to the Faculty of the Graduate School  
of  
Yale University  
in Candidacy for the Degree of  
Doctor of Philosophy

by  
Jacob Zachary Blumoff

Dissertation Director: Robert J. Schoelkopf

December 2017

© 2017 by Jacob Zachary Blumoff  
All rights reserved.

# Contents

<b>Contents</b>	<b>iii</b>
<b>List of Figures</b>	<b>vi</b>
<b>List of Tables</b>	<b>ix</b>
<b>Nomenclature</b>	<b>x</b>
<b>Acknowledgments</b>	<b>xi</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Outline of thesis . . . . .	2
<b>2 Quantum information</b>	<b>4</b>
2.1 The promise of quantum computing . . . . .	5
2.2 Quantum fundamentals, states, & spaces . . . . .	7
2.2.1 Qubits & the Bloch sphere . . . . .	8
2.2.2 Nuclear magnetic resonance & Stern-Gerlach . . . . .	10
2.2.3 Entanglement, <i>spukhafte Fernwirkung</i> & Bell's Theorem . . . . .	13
2.2.4 Quantum mechanics with open systems . . . . .	16
2.2.5 Notions of state fidelity . . . . .	19
2.2.6 Cavity states and representations . . . . .	21
2.3 Processes . . . . .	25
2.3.1 Gates . . . . .	25
2.3.2 Formalisms . . . . .	27
2.3.3 Notions of process quality . . . . .	28
2.4 Measurement . . . . .	30
2.5 Quantum error correction . . . . .	31
2.5.1 Fault tolerance & thresholds . . . . .	34
2.6 DiVincenzo criteria . . . . .	35
2.7 Conclusions . . . . .	38
<b>3 Circuit QED</b>	<b>39</b>
3.1 Introduction to Cavity QED & Jaynes-Cummings . . . . .	42
3.2 Quantum electrical circuits . . . . .	44
3.3 The Josephson effect . . . . .	46

3.4	Transmon physics . . . . .	47
3.5	Black-box quantization . . . . .	49
3.5.1	The conventions of $\mathcal{H}$ . . . . .	53
3.5.2	Black-box quantization with several transmons . . . . .	55
3.5.3	The qubit-qubit coupling . . . . .	57
3.5.4	Cavities as circuits . . . . .	58
3.6	The Purcell effect . . . . .	62
3.6.1	Jaynes-Cummings . . . . .	64
3.6.2	Classical treatments . . . . .	64
3.6.3	Towards a quantum treatment . . . . .	67
3.7	Control & measurement . . . . .	67
3.7.1	Introducing the drive . . . . .	68
3.7.2	Pumps & several-photon processes . . . . .	70
3.7.3	Dispersive measurement . . . . .	71
3.8	Conclusions . . . . .	74
<b>4</b>	<b>Experimental design and methods</b>	<b>75</b>
4.1	Device design . . . . .	76
4.1.1	Charge dispersion: bounds on $\alpha/\omega$ . . . . .	77
4.1.2	I/O couplings: Purcell, driving, measuring, & photon shot-noise . . . . .	77
4.1.3	Dielectric loss & feature size . . . . .	80
4.2	Specific devices . . . . .	80
4.2.1	Flux-tunable two-cavity vertical transmon . . . . .	80
4.2.2	Shiva . . . . .	82
4.2.3	Dodekatheon . . . . .	87
4.3	The control chain . . . . .	89
4.3.1	Single-sideband control . . . . .	89
4.3.2	Optimal control theory . . . . .	90
4.3.3	Input lines . . . . .	93
4.3.4	Output lines . . . . .	95
4.4	The room-temperature measurement chain & analysis . . . . .	96
4.4.1	Heterodyne measurement . . . . .	96
4.4.2	Two-tone interferometer . . . . .	96
4.4.3	One-tone interferometer . . . . .	99
4.5	HFSS . . . . .	100
4.6	Conclusions . . . . .	103
<b>5</b>	<b>Measurement characterization</b>	<b>104</b>
5.1	Introduction . . . . .	104
5.2	Why characterize measurements holistically? . . . . .	105
5.3	Assignment fidelity . . . . .	106
5.4	Detectors & POVMs . . . . .	107
5.4.1	Quantum detector tomography . . . . .	109
5.4.2	Specificity & coherent mis-measurement . . . . .	111
5.5	Quantifying detector performance . . . . .	112

5.5.1	Detectors as processes & the superoperator formalism . . . . .	112
5.5.2	Detector fidelities . . . . .	114
5.6	Quantifying measurement process performance . . . . .	115
5.6.1	From the measurement process to the detector . . . . .	116
5.6.2	Measurement process tomography . . . . .	116
5.6.3	Quantum instruments . . . . .	117
5.6.4	“Quantum non-demolition” . . . . .	118
5.6.5	A superoperator formalism for measurement processes . . . . .	119
5.6.6	Quantum instrument fidelity . . . . .	120
5.7	Conclusions . . . . .	120
<b>6</b>	<b>The Nigg protocol &amp; stabilizer measurements</b>	<b>122</b>
6.1	qcMAP: a useful precursor . . . . .	124
6.1.1	Effects of decoherence . . . . .	126
6.1.2	Limitations of real pulses . . . . .	127
6.1.3	Higher order terms in $\mathcal{H}$ . . . . .	129
6.1.4	Beating $\chi$ . . . . .	129
6.2	The Nigg proposal . . . . .	130
6.2.1	Subset selectivity . . . . .	133
6.2.2	Equalizing $\chi$ . . . . .	134
6.2.3	A different erasure technique . . . . .	135
6.2.4	Going faster . . . . .	136
6.3	Device characterization experiments . . . . .	137
6.3.1	Direct characterization . . . . .	138
6.3.2	Indirect qubit characterization . . . . .	140
6.3.3	Characterization results . . . . .	141
6.4	State preparation . . . . .	142
6.5	Full parity measurement sequence . . . . .	144
6.6	Error detection . . . . .	144
6.7	State & process tomography . . . . .	146
6.8	Experimental results . . . . .	148
6.8.1	Assignment fidelity results . . . . .	148
6.8.2	Detector tomography results . . . . .	150
6.8.3	Measurement process tomography results . . . . .	150
6.9	Conclusions . . . . .	152
<b>7</b>	<b>Deterministic teleportation of a two-qubit gate</b>	<b>155</b>
7.1	Stitching a web of modules . . . . .	157
7.2	Teleportation of quantum states . . . . .	159
7.3	Teleportation of quantum gates: theory . . . . .	160
7.4	Characterization results . . . . .	162
7.4.1	Transmon excited-state populations . . . . .	162
7.5	Simultaneous readout & feedback . . . . .	163
7.5.1	Synchronous feedback . . . . .	164
7.5.2	Ground state preparation by feedback . . . . .	165

7.5.3	Transmon reset . . . . .	167
7.6	RIP gates . . . . .	167
7.6.1	Refocused RIP gates . . . . .	171
7.7	Qubits encoded in harmonic oscillators . . . . .	172
7.7.1	Binomial encoding . . . . .	173
7.7.2	Operation times . . . . .	174
7.7.3	Gate performance . . . . .	175
7.7.4	Leakage/codespace errors . . . . .	177
7.8	Two-bit, phase-corrected tomography . . . . .	178
7.8.1	Two-bit tomography . . . . .	179
7.8.2	Phase correction . . . . .	180
7.8.3	Decoding pulses . . . . .	180
7.9	Full protocol . . . . .	181
7.10	Results . . . . .	182
7.10.1	Fock encoding, without feed-forward . . . . .	183
7.10.2	Fock encoding, with feed-forward . . . . .	185
7.10.3	024 encoding . . . . .	187
7.11	Conclusions . . . . .	188
<b>8</b>	<b>Conclusion</b>	<b>190</b>
	<b>Bibliography</b>	<b>193</b>
<b>A</b>	<b>Stabilizer experiment extended results</b>	<b>203</b>
A.1	Detector tomography extended results . . . . .	203
A.2	Measurement process extended results . . . . .	203
<b>B</b>	<b>Tomography code</b>	<b>209</b>



# List of Figures

<b>Quantum information</b>	<b>4</b>
2.1 The Bloch sphere . . . . .	10
2.2 CHSH violation . . . . .	16
2.3 Qubit decoherence . . . . .	18
2.4 Q and Wigner functions of coherent and incoherent superpositions . . . . .	24
<b>Circuit QED</b>	<b>39</b>
3.1 Spectra in the strong dispersive regime . . . . .	45
3.2 Josephson junction imaged by a scanning electron microscope . . . . .	46
3.3 Transmon charge dispersion and anharmonicity vs. $E_J/E_c$ . . . . .	49
3.4 Black-box quantization circuits . . . . .	51
3.5 The black-box Hamiltonian expressed as interactions . . . . .	54
3.6 Black-box quantization with multiple qubits . . . . .	55
3.7 Phase conventions in multi-qubit black-box quantization . . . . .	57
3.8 Impedance of a rectangular resonator as seen from an external capacitive probe . . . . .	63
3.9 Mixing an LC oscillator and an RLC oscillator . . . . .	65
3.10 Mode mixing gedanken experiment . . . . .	66
3.11 Dispersive readout . . . . .	72
<b>Experimental design and methods</b>	<b>75</b>
4.1 Flux-tunable two cavity vertical transmon . . . . .	81
4.2 Schematic of a sapphire chip from the Shiva device . . . . .	82
4.3 Field mode of one of the Shiva device qubits . . . . .	84
4.4 External detail of the Shiva device . . . . .	85
4.5 External detail of the Dodekatheon device . . . . .	87
4.6 Dodekatheon qubit, readout resonator, and Purcell filter . . . . .	88
4.7 Example optimal control pulse of an $X_\pi$ gate on an encoded “kitten” qubit . . . . .	91
4.8 Fridge and room temperature wiring for the stabilizer experiment . . . . .	97
4.9 Fridge and room temperature wiring for the gate teleportation experiment . . . . .	98
4.10 Two-tone interferometer . . . . .	99
4.11 One-tone interferometer . . . . .	100

<b>Measurement characterization</b>	<b>104</b>
5.1 Potential errors during stabilizer measurements . . . . .	106
5.2 Example of a POVM expressed in the Pauli basis . . . . .	110
5.3 Measurement process tomography schematic . . . . .	117
<b>The Nigg protocol &amp; stabilizer measurements</b>	<b>122</b>
6.1 The qcMAP protocol . . . . .	125
6.2 qcMAP– faster . . . . .	130
6.3 Schematic of the device proposed by Nigg and Girvin . . . . .	132
6.4 Initiating a measurement of $ZI$ . . . . .	134
6.5 Measuring $ZZ$ on a two-qubit register . . . . .	136
6.6 Schematic of the device used in the stabilizer experiment . . . . .	137
6.7 Hamiltonian for the device used in the stabilizer experiment . . . . .	141
6.8 State preparation in the stabilizer measurement experiment . . . . .	144
6.9 Circuit diagram for $ZIZ$ measurement . . . . .	145
6.10 Circuit for state tomography in the stabilizer measurement experiment . . . . .	147
6.11 ‘TriRabis’ showing raw data for measuring assignment fidelity . . . . .	149
6.12 Quantum process heralded by a $ZZZ$ measurement . . . . .	154
<b>Deterministic teleportation of a two-qubit gate</b>	<b>155</b>
7.1 A network of quantum modules . . . . .	158
7.2 Circuit for state teleportation . . . . .	159
7.3 Basic CNOT teleportation circuit . . . . .	161
7.4 Relationship between transmon Purcell effect and polarization. . . . .	163
7.5 Simultaneous Rabi oscillations and measurement . . . . .	164
7.6 “Black mamba” wiring in the gate teleportation experiment . . . . .	165
7.7 State preparation sequence in the gate teleportation experiment . . . . .	166
7.8 Tomography of transmon reset in the gate teleportation experiment . . . . .	167
7.9 Spectral explanation of the RIP gate . . . . .	168
7.10 Tomography of a Bell pair of transmons . . . . .	172
7.11 Wigner tomographs of logical states of the 024 code. . . . .	175
7.12 Interleaved randomized benchmarking of operations on a Fock-encoded qubit . . . . .	177
7.13 Full experimental sequence for CNOT teleportation . . . . .	182
7.14 Tomography of a Bell state in the Fock encoding, without feed-forward . . . . .	184
7.15 QPT of a teleported gate in the Fock encoding, without feed-forward . . . . .	185
7.16 Tomography of a Bell state in the Fock encoding, with feed-forward . . . . .	186
7.17 QPT of a teleported gate in the Fock encoding, with feed-forward . . . . .	187
7.18 Truth tables of the CNOT acting on 024-encoded qubits . . . . .	188
<b>Stabilizer experiment extended results</b>	<b>203</b>
A.1 Tomography of all POVMs, without post-selection on the success-herald . . . . .	204
A.2 Tomography of all POVMs, with post-selection on the success-herald . . . . .	205
A.3 Full conditioned QPT results . . . . .	207

A.4 Conditioned QPT results with success herald . . . . . 208

# List of Tables

<b>Circuit QED</b>	<b>39</b>
3.1 Black-box vs. experiment comparison . . . . .	58
<b>The Nigg protocol &amp; stabilizer measurements</b>	<b>122</b>
6.1 Coherence results for the Shiva device . . . . .	142
6.2 Simultaneous randomized benchmarking results . . . . .	142
6.3 Assignment fidelity results . . . . .	148
6.4 Detector fidelity results . . . . .	150
6.5 Quantum instrument fidelities . . . . .	151
<b>Deterministic teleportation of a two-qubit gate</b>	<b>155</b>
7.1 Empirical comparison of operation times in various encodings . . . . .	176
<b>Stabilizer experiment extended results</b>	<b>203</b>
A.1 POVM results . . . . .	204
A.2 Specificity results . . . . .	205
A.3 Quantum instrument metrics . . . . .	206

# Acknowledgments

First and foremost I thank my adviser, ROB SCHOELKOPF, for making the last six years possible. Rob has built a very capable laboratory and has given it a direction founded on a very keen sense of where the field is now and of what is possible tomorrow. He has filled it with incredibly smart and driven people and created a fruitful and collaborative environment—then given us significant discretion to make what we will of it. I greatly appreciate the opportunities that this structure has provided.

MICHEL DEVORET is also highly central to this intellectual habitat. His insights and often unique perspectives regularly lead to enlightening connections and a greater understanding for all of us. I have always attended to his advice with great care. DAN PROBER, in many ways the progenitor of the Yale quantum information group, has always been kind and welcoming. His influence has fostered an excellent camaraderie in the Applied Physics department. LIANG JIANG's impact in my time at Yale time has been enormous. He has had very impressive availability and willingness to help me with a broad range of questions. Both the measurement characterization results and the gate teleportation experiment owe him a great deal.

After my adviser, KEVIN CHOU was perhaps the most indispensable person to the work of this thesis. The two of us somehow managed to keep a lot of balls in the air, and without him the two main experiments of my thesis would likely not have happened. Itemizing his contributions would be an impossible task. For the last two years we have been joined by CHRIS WANG, who in many ways represents the future of my and Kevin's efforts. The learning curve has been and will remain steep, but I'm confident we are leaving it in good hands.

TERESA BRECHT was and remains one of my closest friends throughout graduate school. In turbulent moments she has a reliable light in the darkness. I owe a great deal to URI VOOL, who has

been both a marvelous friend and an incredibly helpful colleague. His unshakable good humor has been a stabilizing force along the bumpy road to a doctorate, and it's possible that our conversations at the whiteboard have taught me more physics than classes have.

BRIAN VLASTAKIS was the first person I worked with at Yale. He was very helpful in acclimatizing to both the Schoelkopf lab and the fundamentals of what we do, as well as to graduate student life at Yale. He oversaw my initial efforts to better understand cavities and black-box quantization. MATTHEW REED was another important colleague early in graduate school. He was key in the early flux-tunable 3D transmon experiments and in getting the stabilize measurement experiment off the ground. He remains a valued friend, and I look forward to working together at HRL.

STEVE GIRVIN has reliably provided clarity to everything he discusses. I (and everyone on the fourth floor) owe a great deal to LUIGI FRUNZIO, who keeps the wheels turning on this whole operation. MICHAEL HATRIDGE, expert cat herder, and SHYAM SHANKAR provided both excellent technical advice and a degree of sanity to the floor. Along with ANIRUDH NARLA and KATRINA SLIWA, they are directly responsible for making possible the excellent readout which is used to great effect in this thesis. That group is also generally quite helpful in keeping the nuts and bolts of the floor as they should be. My transition to superconducting qubits was aided by GERHARD KIRCHMAIR, whose analogies to atomic physics were greatly appreciated. I give special thanks to him for serving as an outside reader for this document.

NISSIM OFEK, PHIL REINHOLD, and REINIER HEERES all made important contributions, especially to the experimental software stack. Nissim provided FPGA logic to the entire group which has been critical to recent experiments. It is a rare person indeed who has both an aptitude and patience for FPGAs as well as an excellent understanding of the relevant physics. He was also instrumental in the flux-tunable 3D transmon experiments. Phil wrote the python code that makes the FPGA code much easier to use, as well as the entire GRAPE library. His aversion towards accepting physics without question has led to deeper understandings across the floor. Reinier wrote the pre-FPGA software that was crucial for the stabilizer measurement experiment, but more generally is responsible for bringing the lab a focus on programming and automation.

MATTHEW REAGOR and CHRIS AXLINE both did the hard work of building new types of experi-

mental devices. The high-Q resonators and integrated transmons that Matt developed and the compact readout resonators which Chris developed were both crucial to the experiments of this thesis.

It would be deeply irresponsible to not highlight the important work done by GISELLE DEVITO, MARIA RAO, THERESA EVANGELISTE and NUCH GRAVES. Without their efforts this would not be possible. I have always enjoyed our conversations around the coffee maker. I have similar gratitude to DEVON CIMINI, though his role quickly changed to that of a good friend.

I would not have gone into graduate school for physics without several people from my days at Georgia Tech: ALEXANDER RADNAEV, DAVID NAYLOR, ALEX KUZMICH, YARIK DUDIN, and ANASTASIA MARCHENKOVA. Working with Sasha was the first time I ever stepped foot in a research lab, and I owe him a great deal. He was also the first person to send me a paper from the Schoelkopf group. I can also identify several teachers from high school who managed to make a very lazy student actually care about physics and math: TIM MALEY, SCOTT HETHERINGTON, and LARRY HAMPLE.

Last but certainly not least, I acknowledge my family, not one member of which did not offer to proofread this document. Both my father, TED BLUMOFF, and my mother, JANE MCNAMEE BECK, instilled me with curiosity and skepticism. Their encouragement and unconditional support is beyond value. My sisters, KATE BLUMOFF GREENBERG and MEREDITH SCHWARTZ, have been (and remain) similarly crucial, providing support and advisement on this path through my twenties. I also note my step father, CHARLIE BECK, who gave me a handiness which has served me well in experimental physics.

Thank you all.

# Chapter 1

## Introduction

I am a quantum engineer, but on  
Sundays I have principles.

---

John Stewart Bell, 1983

The field of quantum information, and especially superconducting qubits, is in an amazingly prosperous moment. The progress I have seen every year, both from the greater physics community and from Yale, never ceases to amaze me. Quite a few other people seem to agree—many of these technologies are starting to leave academia. Large technology companies are pouring in hundreds of millions of dollars into quantum computing and there are seemingly new startups every month. This evolution has required academia to better define its role in the field, and the Yale groups have done so in a highly fruitful manner by focusing on three-dimensional architectures. Others, focused strictly on building millions of qubits, have eschewed this technology and essentially left us with high-Q resonators all to ourselves. In the last six years we have flourished in that niche.

Our increasing focus on high-Q resonators can easily be seen in the course of my thesis work. The focus of my thesis can essentially be distilled down to using multiple qubits in 3D circuit QED. I have written about two main experiments: the “stabilizer measurement” experiment of chapter 6 and the “gate teleportation” experiment of chapter 7. The former deals with using a high-Q resonator to couple multiple transmon qubits and the latter deals with using transmons to couple multiple high-Q resonator qubits. This thesis reflects that evolution of the Yale vision for quantum computing, whose present state



I will expound on later in section 7.1.

In my thesis work I have rarely generated entirely new physics or particularly novel devices—rather I have taken the hard work of my colleagues and turned the complexity up one notch. As we combine more and more parts the odds of failure increase exponentially, which places new requirements on reliability and efficiency. The artisanal techniques that worked for one qubit often need to be refined to work for four. This thesis has required the persistence to make complex devices work, often leading to very subtle challenges. In exchange I have been rewarded by emergent phenomena that can be quite interesting. I hope you agree.

## 1.1 Outline of thesis

The story of this work begins with a bit of motivation in chapter 2. This includes both the general “What is quantum computing and why we should care” and more forward-looking concerns about quantum error correction and what such a computer will look like. In this chapter I provide a bit of background information about quantum mechanics and its formalisms, not to teach it but to put later concepts in the context that I only slowly acquired. This chapter also discusses several historical experiments that some would not choose to include, but I think serve to provide a greater connection with physics as a whole.

After that I discuss circuit quantum electrodynamics in chapter 3, introducing our physical implementation of the ideas of the previous chapter. It contains some original work in my approach to black-box quantization and in understanding the impedance of 3D electromagnetic environments, but largely concerns preexisting physics. As in the last chapter I do not intend to teach material anew. My intent is to discuss these physics in the way that I have come to understand them over the years—a way that I would have found it helpful. Along with the standard references, this material builds a base for later chapters.

Chapter 4 gets into the meat of my thesis work. This chapter discusses how I go about designing a circuit QED experiment and how to realize the systems discussed in the previous chapter. I attempt to document the considerations that I have gone through when fashioning the devices of this thesis. Addi-

tionally, I present my recipe for the surrounding experimental apparatus, which builds on the knowledge of decades of graduate students. I demonstrate and discuss several specific devices and experimental configurations.

The published results of my thesis begin in chapter 5, which broadly puts forward an ethos for analyzing the sorts of measurements we need for quantum error correction. These measurements, “stabilizer measurements,” are generally composed of many smaller gates and simpler measurements. This chapter presents several ways of looking at them more holistically. I believe that this mindset will be useful along the path to a quantum computer, perhaps requiring us to directly face problems we might otherwise ignore. This analysis is presented first, but was developed in response to the experimental results of the next chapter.

The first major experiment of my thesis was to implement and characterize measurements akin to stabilizer measurements. In chapter 6 we discuss that implementation and provide the results of its characterization. The experimental protocol was derived from an earlier proposal by two of my colleagues at Yale (Nigg and Girvin, 2013). We discuss that pulse sequence and the modifications that we made to it along the way. In this chapter I also present my general recipe for measuring and characterizing qubits and cavities. The device for this experiment is, to my knowledge, the most complicated 3D circuit QED device ever implemented.

The second and final major experiment of my thesis was the realization of a deterministically teleported two-qubit gate. In chapter 7 I discuss the results of that experiment to-date. This protocol is sufficiently technologically challenging that eighteen years passed between the proposal of the idea in 1999 and our realization in 2016 and 2017. Perhaps the most exciting part of this experiment is that we can not only perform this gate on physical two-level systems, but we can actually perform it on encoded, potentially error-correctable qubits. That I had the opportunity to do this project speaks greatly to the technical progress in the field. As we have not yet published this result, the data in this chapter are preliminary and may be quantitatively different from those in the eventual paper.

Finally in chapter 8 I conclude with an outlook on the field. In particular I focus on how I see the future of the approach currently being explored at Yale.

## Chapter 2

# Quantum information

With group and eigenstate, we've learned to fix your quantum errors with our quantum tricks.

---

Daniel Gottesman

In this first chapter we will discuss the most basic theoretical constructs that motivate all of quantum information science. Section 2.1 will cover a very broad view of the end goal: what a quantum computer is and why we want one. Then in section 2.2 we will step back and provide a brief look at quantum mechanics. Section 2.2.1 will continue on this by introducing qubits as a fundamental building block for computation and as an ideal example system for simple quantum phenomena. In section 2.2.2 we will briefly walk through two enlightening historical experiments which both showing the importance of quantum mechanics and introduce laboratory techniques which are relevant to the experiments of this thesis. Entanglement plays a very important role both in quantum mechanics and in the experiments of this thesis, and in section 2.2.3 we will discuss it broadly. We will discuss open-system quantum mechanics in section 2.2.4, which enables both powerful resources and the relentless adversary of decoherence. Section 2.2.5 will begin our discussion of figures of merit, which we will need to judge how well our experiments achieve their goals. In addition to qubits, our experiments also involve harmonic oscillators, and in section 2.2.6 we will provide an introduction to some useful representations of these infinite-dimensional state-spaces.

After this discussion of quantum states we will also discuss the evolution of those states in section 2.3. We will introduce a few qubit-specific quantum processes in section 2.3.1 and a more general and mathematical description of quantum processes in section 2.3.2. These formalisms will be employed in section 2.3.3, which will again cover figures of merit, but now for processes rather than states. These will also be used to quantify our success or failure in reaching experimental goals. In section 2.4 we will very briefly cover measurement in quantum mechanics, though a detailed treatment will wait until chapter 5.

Section 2.5 contains a small overview of quantum error correction. Though we do not perform error correction in this work, it is a crucial motivation for the experiments in this thesis. Finally, in section 2.6 we will cover the DiVincenzo criteria as we segue from the theory heavy discussions of this chapter. These criteria lay out what is required of a physical system to be a useful tool for manipulating quantum information.

## 2.1 The promise of quantum computing

So what *is* quantum information? It is very easy to get bogged down in the theoretical and philosophical answers to this—indeed people have spent countless pages defining classical information. We will simply define quantum information as that which is required to describe the state of a quantum system, and predict the outcome of later measurements we might make upon it. Quantum systems, rather than classical systems, may only be well understood using the tools of quantum mechanics. We represent pure quantum states using vectors in Hilbert space, and we represent (classically) probabilistic combinations of those states with density matrices. Either is a descriptor we choose based on knowledge of the quantum system's preparation, notably including past measurements, the energy landscape it evolves within, and interactions with the outside world we can only describe probabilistically.

Before we expand on that, let us also take a moment to discuss computing. In information processing we aim to take some input information, apply an algorithm, and produce some desired output information. e.g., perhaps we wish to know the prime factors of 923 ( $=17 \times 13$ ), or to add a lens flare to an image file. Quantum information processing is much the same: we begin with some quantum

state, apply an algorithm, and produce an output quantum state. To be useful on its own, a complete quantum algorithm should take in strictly classical information and output strictly classical information, as that is all we can access in this classical world.

The beauty of quantum computing is in that intermediate step, the application of an algorithm. For reasons we will only briefly explore, a quantum computer can execute algorithms which solve certain problems *much* faster than any algorithm that could ever run on any classical computer. The algorithms we would like to perform are powerful, yet oddly limited. The design of quantum algorithms is a tricky business, and we do not presently know the full extent of which problems can be solved more quickly with a quantum computer. Certainly not all informational problems can benefit. One of the earliest problems was Grover's algorithm, which searches through an unordered list of length  $n$  to find the desired entry (Grover, 1996). Classically this problem takes an average of  $\mathcal{O}(n) \sim n/2$  queries to the list, but it turns out that a quantum computer can find the right answer in only  $\mathcal{O}(\sqrt{n})$  queries. For a list of  $n \sim 10^{20}$  this speedup could be quite significant. This is a bit of a toy algorithm, however, and the square-root speedup is quite slow compared to our real quarry, which we can find in the next example.

One particularly important algorithm is the quantum Fourier transform, which performs a discrete Fourier transform on  $n$  bits in  $\mathcal{O}(n^2)$  queries rather than the classical  $\mathcal{O}(n2^n)$  queries (Shor, 1994)<sup>1</sup>. This is an *exponential* speedup, which becomes useful with a much smaller quantum computer. This dramatic improvement should allow solutions to problems which simply are not possible to solve classically—completing them might take the age of the universe on any conceivable classical computer. The quantum Fourier transform enabled another algorithm, Shor's algorithm (Shor, 1994). This algorithm, along with his discovery of quantum error correction (which we will discuss shortly), led to a rapid expansion of interest in quantum computing. The functionality of Shor's algorithm is to break a composite number into its prime factors, and it does so exponentially faster than any classical algorithm. As a great deal of modern encryption depends on this problem being computationally difficult, this was of great interest to many. Accordingly the research of this thesis was well-funded and equipped.

---

<sup>1</sup>There are now even faster implementations, to my knowledge the fastest solution scales as  $\mathcal{O}(n \log n)$

## 2.2 Quantum fundamentals, states, & spaces

To provide some qualitative explanation for why quantum information is so different, we will begin with a short exposition on quantum states. We generally describe pure quantum states as complex vectors of dimension  $d$  with unit norm.  $d$  represents the maximum number of deterministically different configurations of the system. For a bit which may be 0 or 1  $d = 2$ , for three bits  $d$  is  $2^3 = 8$ . For the position or velocity of a particle,  $d$  is generally infinity. We will chiefly be interested in finite-dimensional systems.

Quantum states are not specified as being in one of  $d$  systems, as a string of classical bits might be, nor are they merely probability distributions of the  $d$  potentially discriminable outcomes. The complex nature of quantum state vectors allows wave-like interference. Since probabilities are conserved, we require that the norm of the whole vector is equal to one. Furthermore since we only care about norms, there is a global phase associated with the state that cannot affect measurement outcomes. Accordingly a  $d$ -dimensional pure quantum state is specified by  $2d - 2$  real numbers. These numbers are continuous, so even a two-dimensional quantum state can take an infinite quantity of classical bits to exactly describe.

The purpose of defining quantum states is to predict measurement outcomes via the Born rule. We define *observables*, which are given by Hermitian operators of the same dimension as the system. Each observable  $O$  can be decomposed into a sum of projectors  $\Pi_i$  which span the Hilbert space of the system, and each has real eigenvalues  $a_i$  such that  $O = \sum_i a_i \Pi_i$ . The probability of observing outcome  $a_i$  when measuring  $O$  and starting with a state  $|\psi\rangle$  is given by the norm of  $|\psi\rangle$  projected by  $\Pi_i$ , or  $P_i = \langle\psi|\Pi_i|\psi\rangle$ . The expected value of  $O$  is the weighted sum over these results, given by  $\langle O \rangle = \langle\psi|O|\psi\rangle$ .

The final postulate explains how a quantum state evolves in a given energy landscape. In classical mechanics we apply Hamilton's equations,  $dp/dt = -\partial\mathcal{H}/\partial q$  and  $dq/dt = \partial\mathcal{H}/\partial p$ , yielding a series of differential equations we can solve by whatever means we choose. In quantum mechanics we look not at the evolution in phase space but in Hilbert space. This behavior is given by the Schrödinger equation,

$$i\hbar \frac{d}{dt} |\psi(t)\rangle = \mathcal{H}(t) |\psi(t)\rangle. \quad (2.1)$$

For our present purposes, this describes all quantum evolution other than measurement, whose action on the state we will treat briefly in section 2.4 and in detail in chapter 5. If  $\mathcal{H}$  is constant in time, the evolution of the state can be calculated in closed form, by the propagator  $|\psi(t)\rangle = U(t)|\psi(0)\rangle$  for

$$U(t) = e^{-i\mathcal{H}t/\hbar}. \quad (2.2)$$

This is an interesting result. For a state of definite energy  $E$ , the effect of the propagator is to simply apply a global phase to the state vector, oscillating at a frequency  $\omega = E/\hbar$ . These *stationary states* are not particularly interesting by themselves, as again this global phase has no directly measurable effect. Only state vectors of non-definite energy will show interesting dynamics under the Schrödinger equation. This is trivially the standard case when  $\mathcal{H}$  does depend on time, as the eigenstates will be changing themselves.

### 2.2.1 Qubits & the Bloch sphere

This is all rather abstract, so we turn to the simplest example possible: the  $d = 2$  system. This system is historically called a spin-1/2, because it is the Hilbert space needed to describe the internal angular-momentum state of electrons and simple atomic nuclei. These systems were exploited in some of the first experiments to demonstrate quantum behavior, as we will see shortly. Modernly we call these systems *qubits*, in analogy to the bits of classical computing, whose states can be either 1 or 0.

For the case of a qubit we will use the convention that the vector  $|0\rangle$  ( $|Z = +1\rangle$ ) describes the ground state and that the orthogonal vector  $|1\rangle$  ( $|Z = -1\rangle$ ) describes the higher energy state. These energy eigenstates form a basis for the qubit's Hilbert space. For an isolated qubit the Hamiltonian is then traditionally written

$$\mathcal{H}_{\text{qubit}}/\hbar = \frac{\omega}{2}\sigma_z, \quad (2.3)$$

where we are using one of the Pauli matrices

$$\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad \sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}. \quad (2.4)$$

These matrices are quite important and will be used repeatedly. It is simple to show that along with the identity, these matrices form a complete basis for all possible 2-by-2 Hermitian matrices, so all qubit observables can be decomposed in this manner. In the rest of this text we will sometimes use the phrase “Pauli matrix” or “Pauli basis” to imply the generalized Pauli operators, which consist of the combinations of Pauli operators acting on multiple qubits in parallel.

Another useful result from equation 2.2 is that we can add any constant energy to  $\mathcal{H}$  with no consequence. Since this modification is state-independent, it acts globally on the state vector and cannot change the outcome of a measurement. Rather than describe  $\mathcal{H}_{\text{qubit}}$  with the  $\sigma_z$  operator, we will always shift the Hamiltonian to give the ground state precisely zero energy. We often write equation 2.3 in terms of the excited state projector,  $\mathcal{H}_{\text{qubit}}/\hbar = \omega |e\rangle\langle e|$ . In this “frame” we see  $|0\rangle$  evolves not at all, while the phase on  $|1\rangle$  oscillates at  $\omega$ . This reasoning leads to an important result, the *Bloch sphere*. A qubit state can be written

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle$$

for complex (and possibly time-dependent) coefficients  $\alpha$  and  $\beta$ . This qubit is neither precisely in  $|0\rangle$  nor precisely in  $|1\rangle$ , rather this is a *superposition*. Here if one measured  $\sigma_z$  one would find the ground state with probability  $|\alpha|^2$  and the excited state with probability  $|\beta|^2$ . Since we know the qubit state is normalized, we know  $|\alpha|^2 + |\beta|^2 = 1$  and we can remove one degree of freedom, which motivates writing instead

$$|\psi\rangle = \cos(\theta/2)e^{i\phi_g} |0\rangle + \sin(\theta/2)e^{i\phi_e} |1\rangle.$$

We also know that a phase that can be factored from the entire state vector, so we can also write

$$|\psi\rangle = \cos(\theta/2) |0\rangle + \sin(\theta/2)e^{i\phi} |1\rangle. \quad (2.5)$$

These two angles motivate a highly geometric picture of the qubits Hilbert space—a point on the surface of the unit sphere. This sphere is the aforementioned Bloch sphere. In this representation, shown in figure 2.1,  $|0\rangle$  is represented by at the bottom of the sphere and  $|1\rangle$  is represented by at the top of the sphere, separated by altitude angle  $\theta/2$ .  $\phi$ , now simply called the qubit phase, represents






Figure 2.1: **The Bloch sphere.** The state of a pure two-level system can be expressed by two real numbers with periodic boundary conditions. Accordingly this allows a geometric interpretation where these two real numbers represent the azimuthal and equatorial angles that describe a point on the surface of the unit sphere.

the azimuthal angle. Note that we have labeled the two orthogonal axes<sup>2</sup>  $X$  and  $Y$ . This is because the eigenvectors along these axes are the eigenstates of the  $\sigma_x$  and  $\sigma_y$  operators, respectively. The projection of the state vector onto each axis yields the expectation value of each operator.

We also know how the time-dynamics of a qubit will look in this frame. Stepping back, the evolution of a qubit is given by

$$|\psi(t)\rangle = \cos(\theta/2) |0\rangle + \sin(\theta/2)e^{i(\phi+\omega t)} |1\rangle,$$

which corresponds to a vector rotating around the polar axis at frequency  $\omega$ . This is complicating and not particularly useful behavior, so we often work in a “rotating frame,” where we can express the state’s dynamics more simply. Here we assume that all other qubit operations we will perform are also done in manner rotating at  $\omega$  themselves. This nebulous definition can be confusing, but we can use an illustrative example, nuclear magnetic resonance (NMR) (Rabi *et al.*, 1938).

### 2.2.2 Nuclear magnetic resonance & Stern-Gerlach

In an NMR experiment a vector on the Bloch sphere takes on a very physical link to a spin’s behavior in a magnetic field. When a strong magnetic field is applied, spin-1/2 particles will energetically prefer to be pointing along the direction of the field. This corresponds to enforcing a  $\sigma_z$  Hamiltonian. The

<sup>2</sup>The operators  $X$  and  $Y$  are orthogonal with respect to the operator scalar product, as  $\text{Tr}[XY] = 0$ .  $X$  and  $Y$  appear orthogonal on the Bloch sphere, but their eigenstates are not orthogonal to each other in Hilbert space. The pairs of eigenstates corresponding to the two antiparallel directions on each axis are orthogonal. This fact of nature was likely designed to confuse first-year quantum mechanics students.

spin will precess at a frequency  $\omega$ , precisely as our vector in Hilbert space does. Applying an orthogonal static  $B$  field is equivalent to adding  $\sigma_x$  or  $\sigma_y$  to the Hamiltonian, and the stationary states will change.

If we wish to manipulate the state in a more dynamic manner, we can alternately apply magnetic fields that oscillate along the  $z$  axis at frequency  $\omega$ . When we do this something interesting happens. The field is changing precisely as the spin is precessing. This field will act as though it is always effectively pointing in the right direction. Though the field we are applying in the “lab frame” is changing sign back and forth along some axis in the  $X$ - $Y$  plane at  $\omega$ , in the rotating frame it appears as a fixed operator along that axis, which is defined by the phase of the drive. In this manner when we use resonant control techniques the “rotating frame” is a less awkward tool than it may first have appeared.

In NMR we can initialize<sup>3</sup> an ensemble of spins along an axis we choose to call  $Z$  with a static field, then rotate the spin-state about the  $X$  and  $Y$  axes with oscillating fields, again dependent on the phase offset of the applied drives. As a practical matter, using a rotating frame requires the experimenter to have a “clock” oscillating along with the qubit at  $\omega$ . If the spins are initially in  $|0\rangle$  they have no meaningful phase, so when we begin to manipulate the spins we generally choose the phase such that the first field is a rotation around the  $Y$  axis. All later rotations are done by tones with phases set relative to the first manipulation.

When we leave such a tone on continuously we see Rabi oscillations (Rabi, 1937): that the state will continuously oscillate between  $|0\rangle$  and  $|1\rangle$ . If we stop when the state is completely inverted, or the angle of rotation is  $\pi$ , we call this a “ $\pi$  pulse” around either the  $Y$  or  $X$  axis. We will sometimes more formally denote those operations as  $Y_\pi$  and  $X_\pi$ , or  $R_y(\pi)$  and  $R_x(\pi)$ . Note that the propagators for those actions are equivalent to  $\sigma_y$  and  $\sigma_x$ . At half of the  $\pi$  pulse duration we have performed a  $\pi/2$  pulse, which, for example, takes a spin from the ground state to a superposition state. All operations we can generate by manipulating the Hamiltonian, or *unitary operations*<sup>4</sup>, can be visualized as a rotation by a specific angle around a specific axis. This generalizes, though less intuitively, to higher-dimensional systems.

This sort of manipulation is very similar to the controls we apply in the works of this thesis, discussed further in chapter 3. NMR has one very significant drawback, which is in the nature of its measurement.

<sup>3</sup>In practice this is a very approximate operation, since NMR is generally performed at non-negligible temperatures.

<sup>4</sup>We will discuss operations, including non-unitary operations, in more detail later in section 2.3

Using an additional magnetic field coil, we can attempt to pick up the magnetic field generated by these spins. In traditional NMR this is a very weak signal, and we cannot measure an individual spin. Instead we make only ensemble measurements, measuring the expectation values of operators, rather than discrete outcomes with some individual probability distributions.

A second seminal experiment, performed in 1921 by O. Stern and W. Gerlach, showed that spins can take on discrete values when measured, rather than simply expectation values. Their experiment used silver atoms, which have one valence electron and are well-described for our purposes as spin-1/2 particles. After being heated in an oven these traveling silver atoms passed through a series of holes, collimating into a directional beam. Stern and Gerlach diverted the beam through a magnetic field, inhomogeneous in one direction,  $\hat{z}$ , which results in a spin-dependent force in the  $\hat{z}$  direction,

$$F_z \propto \sigma_z \frac{\partial B_z}{\partial Z}.$$

These silver atoms then propagated a further distance and impinged upon a glass plate. By observing the metalization of the glass plate they could see the final spatial distribution of atoms, which should inform the initial spin distribution. As the particles coming out of the oven are oriented randomly, we expect the spin projection in the  $\hat{z}$  direction to also be randomized. Classically we would expect this spin to take on a distribution centered between spin-up and spin-down, and might expect to see a vertical stripe. However this is not what Stern and Gerlach saw—they observed two separated spots. They had performed a quantum measurement and seen that the system is always in one of the two discrete states, and nothing in between. This “quantization” of spin angular momentum was, and remains, an incredibly important result. These measurements are direct precursors to the measurements we perform in circuit QED. In that case we will find that the phase of transmitted microwaves take on discrete values. Stern and Gerlach also cascaded multiple of their devices, diverting them such that, e.g., the “spin-up” beam went through a second inhomogeneous field. This observation of the projective measurement “back-action” that we will discuss in great detail in chapter 5 is also quite incredible.

### 2.2.3 Entanglement, spukhafte Fernwirkung & Bell's Theorem

This remarkable result of the 1920's was in a particularly fruitful period for physics—quantum mechanics was beginning to take shape. As these changes took place many physicists of the time were quite perturbed by some of the predictions of the new theory. One particular concern was that of *entanglement*. Entanglement is key to all quantum algorithms and is crucial to quantum computing, as we will see in 2.3.1. Generating and using entanglement is a core theme in the experiments of this thesis. Accordingly a rudimentary discussion is in order, which begins with how quantum systems grow.

When we have two independent qubits or spins, each exists in its own  $d = 2$  Hilbert space. We could equally well describe them in one  $d = 2 \times 2 = 4$  Hilbert space using the *tensor product*. A new basis may then be  $\{|0\rangle_1 \otimes |0\rangle_2, |0\rangle_1 \otimes |1\rangle_2, |1\rangle_1 \otimes |0\rangle_2, |1\rangle_1 \otimes |1\rangle_2\}$ , where the subscripts refer to each qubit. For convenience we often simply write  $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$ . Operators and observables on the space behave similarly—for example a measurement of  $Z$  on the first qubit is more rigorously given by  $Z \otimes \mathcal{I}$ . Working in this larger space leaves open an interesting possibility: superposition states like  $(|00\rangle + |11\rangle)/\mathcal{N}$ . This state, unlike the basis states above, cannot be written as a tensor product of two single-qubit states. At a very fundamental level, the qubits have no individual descriptors. The only description of the state that describes all of the phenomena you might observe about each qubit requires a joint description of both. If one were to ignore the second qubit, as we will see later in section 2.2.4, one would see the rather boring result of classical randomness on the first qubit.

A key observation about this example state, which belongs to the family of Bell states, is that individual measurements of the two qubits will be correlated. If you measure  $Z$  of qubit one, you will get  $|0\rangle$  or  $|1\rangle$  half of the time, but when you measure qubit two you will always get a result that agrees with the first measurement. This phenomenon, with some caveats, is entanglement. It famously caused Einstein great strife (Einstein *et al.*, 1935), and he derided it as “spukhafte Fernwirkung,” or “spooky action at a distance,” in a 1947 letter to Max Born. His concern was that it seemingly violated special relativity and the idea that nothing, including information, can travel faster than the speed of light. Consider taking such a Bell pair and splitting the two qubits, then sending them in opposite directions. If one simultaneously measured both qubits, correlated results could be observed more quickly than light could travel from one experimentalist to another (though of course no observer could causally

confirm this). Quantum mechanics does seem to be nonlocal, but Einstein is correct in that we cannot actually transmit preexisting information faster than light, only random signals. We will approach a more interesting version of this phenomenon later in our discussion of the teleportation of quantum states in section 7.2. This protocol does allow the transmission of preexisting information, though still not faster than  $c$ .

Quantifying entanglement is a tricky matter, and there are many definitions (Nielsen and Chuang, 2000). The correlations we have discussed so far are actually not strictly quantum—you can also make a box that probabilistically prepares  $|00\rangle$  half of the time and  $|11\rangle$  the other half. We need further evidence to rule out that scenario. Thanks to J. S. Bell we know that there are experiments that can rule out any explanation relying on “hidden variables” like such a secret preparation (Bell, 1966). Perhaps the most famous of such experiments is the CHSH inequality (Clauser *et al.*, 1970).

Imagine that one experimentalist, Charlie, prepares two qubits, and sends one each to two other experimentalists, Alice and Bob. Alice can perform one of two measurements on her qubit, labeled  $Q$  or  $R$  and which can each result in discrete results  $q, r = \pm 1$ . Bob can do the same with measurements  $S$  and  $T$ . Consider the quantity  $QS + RS + RT - QT$ . Some thought can show that for any set values of  $q, s, r, t$ , this quantity fits within the range  $\pm 2$ . Or in terms of expectation values and allowing some joint probability distribution of results,

$$\langle QS + RS + RT - QT \rangle = \sum_{qrst} p(q, r, s, t) (qs + rs + rt - qt) \quad (2.6)$$

$$\leq 2 \quad (2.7)$$

$$\langle QS \rangle + \langle RS \rangle + \langle RT \rangle - \langle QT \rangle \leq 2. \quad (2.8)$$

Now we apply a quantum mechanical example. Charlie prepares the quantum state  $\frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$ , a different but still fully entangled Bell pair, and transmits one qubit each to Alice and Bob. We

choose Alice's and Bob's measurements to be

$$\begin{aligned} Q &= Z_1 & R &= X_1 \\ S &= \sqrt{1/2}(-Z_2 - X_2) & T &= \sqrt{1/2}(Z_2 - X_2). \end{aligned}$$

We can evaluate the relevant expectation values

$$\begin{aligned} \langle QS \rangle &= \sqrt{1/2} & \langle RS \rangle &= \sqrt{1/2} \\ \langle RT \rangle &= \sqrt{1/2} & \langle QT \rangle &= -\sqrt{1/2}. \end{aligned}$$

So apparently for this choice of state and measurements,  $\langle QS \rangle + \langle RS \rangle + \langle RT \rangle - \langle QT \rangle = 2\sqrt{2} > 2$ . Curious. Experimentally, we find the latter to be proven correct. See figure 2.2. So what was wrong in our derivation of the inequality? There are two main assumptions which cannot be true at the same time. First, realism: the idea that the measurement results have definite values, independently of observation. Second, locality: that the first measurement does not “reach out” and influence the second measurement faster than the speed of light. Together these assumptions cannot be true, and we see that quantum mechanics violates “local realism.” Further interpretation is left as an exercise to the reader, but clearly something deeply different than our everyday experience is afoot.

While calculations like the CHSH were envisioned as experiments to support quantum mechanics, presently most physicists generally agree that quantum mechanics has a sufficiently large base of supporting evidence. Rather than proving that it is possible for entanglement to exist, these inequalities remain useful as a way to prove that an experimentalist has actually generated entanglement. In this thesis work we do generally perform explicit Bell or CHSH tests, but we certainly do generate entanglement, as we will see in chapters 6 and 7. Rather we generally demonstrate entanglement using the more exhaustive tool of state tomography (sections 6.7 and 7.8) to fully reconstruct the generated states, which in turn shows entanglement or lack thereof. This requires another discussion, the “quality” of quantum states, but we need one more tool first.

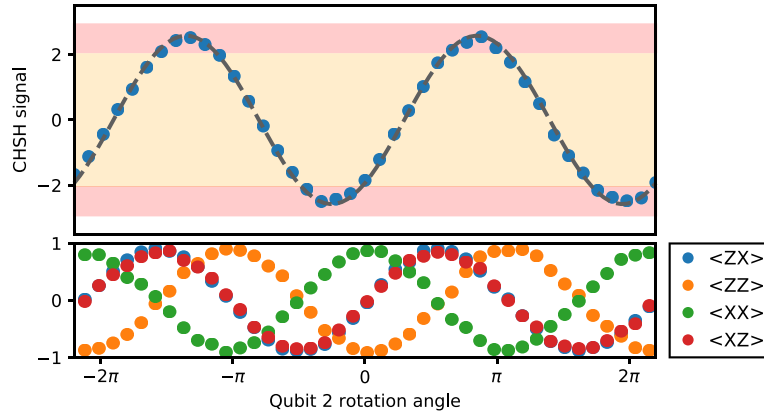


Figure 2.2: **Experimental CHSH violation.** These data depict an experimental violation of a Bell inequality using the device of chapter 7. The four curves in the lower panel represent essentially the same observables discussed previously, except parameterized in the measurement angle of the second qubit. We can see in the upper curve that  $|\langle QS + RS + RT - QT \rangle|$  reaches 2.58, far above the classical limit of 2 (orange region), but is within the quantum limit or  $2\sqrt{2} \approx 2.83$  (red region.)

### 2.2.4 Quantum mechanics with open systems

What if we have classical uncertainty about our state? For example, what if we know that some qubit source generates  $|0\rangle$  half of the time, and  $|1\rangle$  the other half<sup>5</sup>? We can predict the expectation value of a measurement outcome in a more general statement with

$$\langle O \rangle = \sum_i p_i \langle \psi_i | O | \psi_i \rangle,$$

for a system prepared in  $\{|\psi_i\rangle\}$  with probabilities  $\{p_i\}$ . Alternately, we can create the construct  $\rho = \sum_i P_i |\psi_i\rangle \langle \psi_i|$ , where this  $|\rangle \langle|$  ordering represents an outer product of vectors which we express as a

<sup>5</sup>This is quite different from a superposition. As an example, if we were in a qubit state along the  $X$  axis, we would also be 'half 0' and 'half 1' when we measured  $Z$ , but with a definite result if we measured  $X$ . The state prepared with classical randomness would not have a definite result under any measurement.

$d$ -by- $d$  matrix. We can then note that

$$\text{Tr} [\rho O] \equiv \sum_j \langle \psi_j | \rho O | \psi_j \rangle \quad (2.9)$$

$$= \sum_{ij} \langle \psi_j | \psi_i \rangle \langle \psi_i | p_i O | \psi_j \rangle \quad (2.10)$$

$$= \sum_{ij} \delta_{ij} \langle \psi_i | p_i O | \psi_j \rangle \quad (2.11)$$

$$= \sum_i p_i \langle \psi_i | O | \psi_i \rangle = \langle O \rangle, \quad (2.12)$$

where in line 2.9  $\text{Tr}$  represents the trace, or sum of inner products with any complete basis and in line 2.11  $\delta_{ij}$  is the Kronecker delta<sup>6</sup>.

One can also quickly see that the Schrödinger equation can be massaged into the form

$$\dot{\rho} = \frac{-i}{\hbar} [\mathcal{H}, \rho]$$

where  $[a, b]$  is the commutator. One useful property of this *density matrix*  $\rho$  is that for a complete and orthonormal set of operators (as in  $\text{Tr}[\sigma_i \sigma_j] = d \delta_{ij}$ ) we can write

$$\rho = \frac{1}{d} \sum_i \langle \sigma_i \rangle \sigma_i \quad (2.13)$$

. This quantity  $\rho$  allows us to conveniently fold our classical uncertainty about the state into the state definition itself and neatly predict measurement outcomes. This is a crucial tool to describe physical experiments and systems over which we do not have complete control. We know, for instance, that systems have an energy relaxation time-scale over which they thermalize with the environment. For any given experiment we cannot know precisely when our system will relax, but averaging over many experiments we can apply a measure of uncertainty. In the example of qubit energy relaxation we have an excitation-swapping term like  $g \sigma_-^q \sigma_+^e$ , where the first term removes an excitation from the qubit and the second adds an excitation to an unmeasurable mode of the environment. When we assume also

<sup>6</sup>In this equation we assumed that the state basis was equal to the basis we used to take the trace, which is justified because while we can initially define  $\rho$  in one basis, we may manipulate it into a suitable basis.






Figure 2.3: **Qubit decoherence.** **a**, Energy relaxation brings the qubit towards an incoherent thermal state. The populations of that thermal state generally follow a Boltzmann distribution. As this state also has  $\langle \sigma_x \rangle = \langle \sigma_y \rangle = 0$ , we can represent that stabilized density matrix as a point on the  $Z$  axis. **b**, Dephasing eliminates the off-diagonal elements of the density matrix but does not affect the  $\langle \sigma_z \rangle$ , accordingly it stabilizes a class of density matrices which lie on the  $Z$  axis.

that the environment is itself even more strongly coupled to other environmental modes, we know that the state  $|e\rangle_q$  will probabilistically decay to  $|g\rangle_q$  at a rate  $\propto g$ . We more generally state that there is a continuum of environmental modes with some spectral density, often “white noise” which is flat in frequency space. In this standard case the qubit’s energy relaxation probability follows an exponential curve in time with a constant  $T_1$ .

In many systems energy relaxation is a less relevant coupling term than dispersive or similar couplings that change the qubit frequency dependent on the environment. If we do not explicitly consider the environmental modes, this looks like  $\mathcal{H}_{\text{dephasing}} \propto f(t) |e\rangle\langle e|$ , where  $f(t)$  is some unknown function of time. Again we assign some spectral density to that noise. If this noise is white, we see that the diagonal elements of  $\rho$  are untouched, but that the off-diagonal elements or “coherences” decay a probability  $1 - \exp(-t/T_\phi)$ . In the presence of both relaxation and dephasing, the off-diagonal elements decay at the *decoherence rate*  $T_2^{-1} = (2T_1)^{-1} + T_\phi^{-1}$ <sup>7</sup>. In fixed-frequency transmons we can, if we try, achieve

<sup>7</sup>To derive those results, we can also state that relaxation looks like the probabilistic application of  $\sigma_-$  and  $\sigma_+$  and dephasing looks like the probabilistic application of  $\sigma_z$ . In a more general form and under our noise assumptions the action on the density matrix is given by the Lindblad equation

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

for operators  $L_i$  probabilistically applied on time-scale  $\gamma_i$ .

relatively long  $T_\phi$  such that  $T_2 \sim 2T_1$ , limited by relaxation<sup>8</sup>. This is somewhat unusual, as most quantum information systems are limited by  $T_\phi$ . One last comment on  $T_\phi$  is that it is also a matter of the experimentalists book-keeping. If the clock which sets the phase of manipulations is noisily running fast or slow, it looks identical to the qubits frequency changing.

One other way to look at decoherence, and dephasing quite naturally, is to consider it as the environment measuring the system. Consider the example of two qubits, one of which is under our control and one represents the environment. They share a dispersive Hamiltonian  $\mathcal{H}_{\text{dispersive}}/\hbar = \chi\sigma_z^q\sigma_z^e$ . If we start in  $|+\rangle_q|+\rangle_e$  and evolve for time  $\pi/\chi$ , the result is a completely entangled state. What happens if we examine qubit one, but ignore the environment? We represent that mathematically by ignoring the state of qubit two using the *partial trace*,

$$\rho_q = \sum_i \langle i|_e |\psi\rangle \langle \psi| |i\rangle_e. \quad (2.14)$$

We can think that the state of the second qubit was measured (by the environment) but that we have ignored that result, and averaged over it. The probability of the second qubit being in state  $|i\rangle_e$  is given exactly by  $|\langle \psi|i\rangle_e|^2$ , so it is no surprise that we can massage that equation into

$$\rho_q = p(e=0) |0\rangle\langle 0|_q + p(e=1) |1\rangle\langle 1|_q, \quad (2.15)$$

which is exactly the loss of phase coherence. When a system becomes entangled with another system and that second system is ignored, this is an “unread measurement.” Decoherence is the experimentalist’s *bête noire*, and it complicates and degrades every experiment in this thesis<sup>9</sup>.

### 2.2.5 Notions of state fidelity

The fact that we cannot perfectly maintain our carefully prepared states motivates a new question: how close is the state we have to the state we want to have? Or a more general question: how similar are two quantum states? There are several quantities to examine, and both will be used extensively later

<sup>8</sup>This is generally not the only concern—over-isolating the qubit makes it difficult to measure and manipulate.

<sup>9</sup>This is not strictly true. We use relaxation to (mostly) initialize out qubits, and it can also be used to stabilize far more interesting states (Leghtas *et al.*, 2015).

in this thesis to quantify our successes and failures.

The first, *trace distance*, is defined as

$$D(\rho, \sigma) \equiv \frac{1}{2} \text{Tr} |\rho - \sigma|, \quad (2.16)$$

where  $|X| \equiv \sqrt{X^\dagger X}$ . It has a compelling physical interpretation (Jozsa, 1994): given two states  $\rho$  and  $\sigma$ , for an optimal measurement the probability of correctly discerning between the two states is given by  $(1/2)(1 + D(\rho, \sigma))$ . One property of the trace distance is that it is a genuine metric, equal to zero when  $\rho = \sigma$  and with maximum value 1. It has contractivity, meaning that valid quantum operations applied to both  $\rho$  and  $\sigma$  can only decrease  $D$ . These properties, along with others not listed here, make the trace distance the measure of choice for theorists (Gilchrist *et al.*, 2005), especially for threshold calculations, which we will discuss in section 2.5.1.

The second figure of merit, fidelity, is more commonly used by experimentalists. For pure states (fully described by vectors, rather than density matrices), it has the exact meaning of the probability overlap,

$$\mathcal{F}(|\psi\rangle, |\phi\rangle) \equiv |\langle\psi|\phi\rangle|^2. \quad (2.17)$$

If on the state  $\psi$  you performed the measurement “is this system prepared as  $\phi$ ?” the fidelity is the probability of the result “yes.” Fidelity has the same interpretation when comparing a pure state to a density matrix, given by  $\mathcal{F}(|\psi\rangle, \rho) \equiv \langle\psi|\rho|\psi\rangle$ .

The most general statement is given for density matrices and reduces to these two earlier cases. Given by

$$\mathcal{F}(\rho, \sigma) \equiv \left( \text{Tr} \left[ \sqrt{\sqrt{\rho}\sigma\sqrt{\rho}} \right] \right)^2, \quad (2.18)$$

it does not actually have an easily stated physical meaning for two density matrices<sup>10</sup>. The fidelity is of course not a metric, equal to one when the states are the same and zero when they are orthogonal<sup>11</sup>.

<sup>10</sup>Some physicists use the square root of this quantity. While a larger number, this quantity does not have a directly useful physical meaning in any case, so we use the “squared” definition. This will return in section 2.3.3 as it is a useful way to calculate the “probability of success” of a quantum algorithm, which may be the most important number of all.

<sup>11</sup>It can be made into a metric via the Bures metric or angle. An angle very much like this will be revisited in section 5.4.2.

The fidelity has the property of monotonicity, which is essentially the previous contractivity property but trending towards  $\mathcal{F} = 1$  rather than  $D = 0$ . The fidelity is preferred by experimentalists because it is much easier to measure than the trace distance when the target state is a pure state. As mentioned before, if the physicists can arrange the measurement “is this state the desired state?” then only one (averaged) measurement is required to obtain the fidelity.

The fidelity and distance are related by a pair of inequalities,

$$1 - \sqrt{\mathcal{F}(\rho, \sigma)} \leq D(\rho, \sigma) \leq \sqrt{1 - \mathcal{F}(\rho, \sigma)}. \quad (2.19)$$

The second inequality is saturated if both  $\rho$  and  $\sigma$  are pure states. This leads to one final number, *purity*. Since for a pure density matrix there is one  $p_i = 1$ , it is rank-1, or has only one non-zero eigenvalue. We can quantify this in a neat manner by looking not at the sum of eigenvalues (which always equals one) but at the sum of the squares of eigenvalues. Since probabilities are positive definite, this is maximally one when there is one non-zero eigenvalue. We can express it as

$$\mathcal{P} = \text{Tr} [\rho^2]. \quad (2.20)$$

### 2.2.6 Cavity states and representations

This story will also require an additional set of tools for describing the states of resonators, rather than qubits. As the state of a harmonic oscillator exists in an infinite-dimensional Hilbert space, it is not naturally represented by a density matrix. Of course we use a truncated density matrix for computations, but for illustration and conversation we often use one of two graphical representations: the Q function and the Wigner function. Additionally, in the work of this thesis we are often interested not in energy eigenstates (Fock states) of the cavity, but in finite superpositions of *coherent states* (Glauber, 1963). In this section we will explore these three concepts.

Coherent states are states of a harmonic oscillator which are most similar to classical states. Just as an oscillating spring does at any point in time, they have some instantaneous mean position and velocity (or electric and magnetic fields), unlike Fock states. We can describe that mean position in phase space

by a complex number  $\alpha$ . These states also have minimal and symmetric uncertainty, saturating the Heisenberg limit such that  $\sigma_x \sigma_p = \hbar/2$ . Under a true harmonic oscillator Hamiltonian they also evolve classically, simply oscillating about a circular trajectory in phase space. The vacuum Fock state is also the coherent state  $|\alpha = 0\rangle$ , and applying any classical electric or magnetic field leads to

$$\mathcal{H} = \xi^* a + \xi a^\dagger \mapsto D(\alpha) = \exp\left(\alpha^* a - \alpha a^\dagger\right), \quad (2.21)$$

which generates the translation (or displacement) operator. The result of this is that any driving by a classical field will take the vacuum state (or any other coherent state) to another coherent state. Accordingly these states are often described as classical or near-classical. We can describe these states in the Fock basis by

$$|\alpha\rangle = e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle. \quad (2.22)$$

These coherent states, or their quite non-classical superpositions (often called *cat states*, after the Schrödinger's cat gedanken experiment), will appear often in chapters 6 and 7. Lastly, as we can see from equation 2.22, coherent states are eigenstates of the  $a$  operator, indicating that they are, in some ways, insensitive to photon loss. This rather counter-intuitive fact is discussed in (Haroche and Raimond, 2012), and is crucial to many experiments in the Schoelkopf lab (Ofek *et al.*, 2016).

Coherent states directly motivate one cavity representation, the Husimi Q function (Haroche and Raimond, 2012, Husimi, 1940), commonly called just the Q function. This function is real-valued on a complex domain, and for a state  $|\psi\rangle$  is given by the probability overlap between  $|\psi\rangle$  and a coherent state  $\beta$ ,

$$Q^{|\psi\rangle}(\beta) = \frac{1}{\pi} |\langle\beta|\psi\rangle|^2. \quad (2.23)$$

In this representation a coherent state  $|\alpha\rangle$  is simple a Gaussian with width 2, centered at  $\alpha$ . As the Q function is directly proportional to a probability, it is positive-definite and the prefactor  $1/\pi$  is chosen such that it is normalized. The primary benefit of the Q function is that it is simple to understand and measure. We need to be able to perform arbitrary displacements and to query whether or not there are zero or a non-zero number of photons. As we will see in chapter 3, these operations may be done quite

naturally in circuit quantum electrodynamics.

Though the Q function is a complete representation of an arbitrary quantum state, it becomes exponentially insensitive to certain features. For instance, a superposition of two coherent states and a mixed, probabilistic combination of the same two coherent states will have very similar Q functions if their separation in phase space is not small. Any noise at all will quickly make them indistinguishable, which means that experimentally measuring quantum effects with the Q function is quite challenging. Fortunately we also have a second representation, the Wigner function (Haroche and Raimond, 2012, Wigner, 1932). There are several ways to define the Wigner function for a state  $\rho$ , one of which is that its integrals are the marginal probability distributions of  $x$  and  $p$ ,

$$P^\rho(p) = \int dx W^\rho(x, p)$$

$$P^\rho(x) = \int dp W^\rho(x, p).$$

Though I will not provide the derivation, the Wigner function is also the expectation value of a parity measurement (whether the number of photons is odd or even) after a complex displacement  $\beta$ ,

$$W^\rho(\beta) = \frac{2}{\pi} \text{Tr} [D(-\beta)\rho D(\beta)\Pi]. \quad (2.24)$$

This somewhat strange result is quite useful. As we will see in chapter 3, the parity measurement also comes fairly naturally in our area of cQED. The Wigner function is very useful—many of the non-classical signatures washed away in the Q function are maintained in the Wigner distribution. We can make an additional general statement: if a state's Wigner distribution takes negative values, that is sufficient (but not necessary) evidence to claim that the state has no classical explanation. Recall the example of a superposition of two coherent states. These “cat” states have quite different properties than a classical incoherent mixture of the same two states. As shown in figure 2.4, however, their Q-functions look almost identical, but they have quite distinct Wigner functions. In chapter 7 we will see how these concepts are relevant to quantum computation, where the Wigner function will serve to provide complete state tomography.




Figure 2.4: **Q and Wigner functions of coherent and incoherent superpositions of coherent states.** For all of these depictions  $\alpha = \sqrt{2}$ . **a** The Q function of a cat state. **b**, The Wigner function of a cat state. We can see that it has negative (blue) values, clearly indicating non-classicality. **c**, The Q function of an incoherent superposition of the same cavity states. Even at this very small  $\bar{n} = 2$  cat state we can see that the incoherent and coherent superpositions have very similar Q functions. **d**, The Wigner function of the same incoherent superposition. We can see quite clearly that this depiction lacks the fringes seen for the coherent superposition, indicating classicality.

## 2.3 Processes

There are several broad approaches to using the power of quantum mechanics to solve informational problems, the most prominent of which are adiabatic quantum computing (Farhi *et al.*, 2000), measurement-based quantum computing (Raussendorf and Briegel, 2001), and circuit-model quantum computing. Circuit-model quantum computing is by far the most common, and is the only sort of quantum computer model directly relevant to this thesis. This model is the closest to how a classical computer works—we take a set of qubits and apply a set of time-ordered logical gates on and between them. This allows us to quite naturally break an algorithm down into smaller operations, more formally called processes. In this section we will discuss several specific and relevant processes, how we may formally describe them, and some notions of how to judge the “quality” of a process we have realized. We have already discussed several of these processes or in section 2.2.2 in the context of NMR.

### 2.3.1 Gates

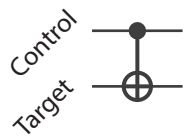
Those single-qubit “gates” are the simplest processes we will consider. The digital analog of this class consists of only one non-trivial one-bit logic gate, NOT, but single-bit quantum gates have a much broader range of possibilities. For now we will discuss the set of unitary operations, which act as rotations of the Bloch sphere around an arbitrary axis and by some arbitrary angle. In section 2.2.2 we singled out a few specific operations: rotating by  $\pi$  or  $\pi/2$  radians around the  $X$  or  $Y$  axes. These operations are sufficient to generate an important group of operators: the Clifford group for one qubit (Nielsen and Chuang, 2000).  $\pi$  and  $\pi/2$  rotations around the  $Z$  axis are also part of this group, though it should be noted that as in circuit QED we have phase control over our drives,  $Z$  operations generally need not be physically implemented. Since  $[Z, X] \propto Y$ , we can always implement  $Z$  rotations by changing the phases of later  $X$  and  $Y$  rotations. One other single-qubit gate in the Clifford group which we will use regularly is the Hadamard gate,

$$H = \sqrt{\frac{1}{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

Preferred by theorists for its neatness,  $H$  is (up to a phase) a rotation around the  $(\hat{x} + \hat{z})/\sqrt{2}$

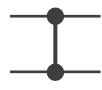


axis, giving it the property that  $HH = I$ . It takes  $|0\rangle$  to  $|+\rangle$  and back again, making it a very clean transformation from the  $Z$  to  $X$  bases and vice versa. The Clifford group for one qubit rotates states between the six corners of a tetrahedron inscribed in the Bloch sphere. The Clifford group for more than one qubit can be generated with only one further gate, often chosen to be the CNOT or CPHASE gate. The CNOT gate performs an  $X$  operation on the “target” qubit if and only if the “control” qubit is excited. The matrix description and circuit diagram are given in 2.3.1.



$$\text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad (2.25)$$

Another important and intimately related gate is the CPHASE gate, which can be generated from a CNOT (or vice versa) by two single-qubit Hadamard gates. The CPHASE operation can be verbosely stated: “apply a  $\pi$   $Z$  rotation to the target qubit if and only if the control qubit is excited.” However, one quickly notes that it is a symmetric operation. Applying  $\pi$  phase to  $|11\rangle$  does not discriminate between target or control qubits.



$$\text{CPHASE} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \quad (2.26)$$

The CPHASE gate comes quite naturally in our system, as we will see in from the RIP gate in section 7.6. In practice a dispersive interaction between two qubits  $\mathcal{H}_{\text{qubit-qubit}}/\hbar = \zeta q_1^\dagger q_1 q_2^\dagger q_2$  is quite difficult to avoid in our systems and will eventually perform a CPHASE after a time  $\pi/\zeta$ .

One can also note that the matrix representations we have presented for these two gates cannot be separated into a matrix of the form  $U_1 \otimes U_2$ , which implies that both can be used to generate entanglement. It can be shown that operations from only the Clifford group alone are not sufficient for universal quantum computing (Gottesman and Chuang, 1999). The Gottesman-Knill theorem (Gottesman, 1997) states that any sequence of only Clifford operations can be simulated efficiently on a classical computer, disallowing any quantum speedup. Fortunately the Solovay-Kitaev theorem (Dawson and Solovay, 2006)

shows that we need only one single non-Clifford gate to unlock a full universal computer. With that one gate and the rest of the Clifford group we can efficiently approximate any unitary of our choice. This happily means that we have a discrete and fairly small set of gates that must be optimized for our general-purpose computer.

### 2.3.2 Formalisms

Unitary operations, as we have seen, can be written as  $d$ -by- $d$  matrices,  $U$ . They operate on kets as  $U|\psi\rangle$ , and a little thought shows they act on density matrices as  $U\rho U^\dagger$ . These are rotations in Hilbert space and preserve the rank of density matrices. As we saw in section 2.2.4, not all processes act this way. More generally a process need only take a valid density matrix to another valid density matrix:

1. The new density matrix should also be trace one.
2. The new density matrix should have positive eigenvalues, always.

This means that the map should be *trace preserving*<sup>12</sup> and *completely positive*. Completely positivity is a subtle property, and is a stricter requirement than positivity. A completely positive map will be certain to maintain positive eigenvalues even when the input state is entangled with some auxillary state. Maps which satisfy both of these are unsurprisingly called completely positive trace preserving (CPTP) maps. We can write a CPTP process  $\mathcal{E}$  in several forms, the first of which is in Krauss form,

$$\mathcal{E}(\rho) \mapsto \sum_i A_i \rho A_i^\dagger. \quad (2.27)$$

The number of operators  $A_i$  can be as high as  $d^2$ . While we cannot choose an explicit representation basis, the Kraus form is not unique. There are many possible Kraus representations. If we want to write the process in a well-defined basis, say the Pauli operators, we can use the  $\chi$  matrix,

$$\mathcal{E}(\rho) \mapsto \sum_i \chi_{mn} \sigma_m \rho \sigma_n. \quad (2.28)$$

---

<sup>12</sup>We will modify this requirement eventually for probabilistic processes.

Note that this representation, unlike the Kraus form, is generally not symmetric. The  $\chi$  matrix is the standard form we will use to present processes, though there are arguments to be made that the R-matrix representation (Chow *et al.*, 2012) is more effective. The  $\chi$  matrix is complex, meaning that both real and imaginary parts must be shown. Furthermore it is also Hermitian, meaning that an element-wise description has significant redundancy.

One alternate approach is to note that these operators are linear transformations on matrices. We can alternately work in an form where we take  $d$ -by- $d$  density matrices to  $d^2$  vectors,  $\vec{\rho}$ . In this form the superoperator representing a process is given by a  $d^2$ -by- $d^2$  matrix. This linear operation can be written

$$\mathcal{E}(\rho) \mapsto \tilde{T}\vec{\rho}. \quad (2.29)$$

Higher-dimensional representations are not amenable to visualization, but provide a neat mathematical form in reducing three-tensors to matrices. We will use this convenience in the next section.

### 2.3.3 Notions of process quality

Another question we will have to confront is that of process quality—if we have enacted some physical process, and we wanted some ideal process, how well did we do? Similarly to section 2.2.5 where we asked “how similar are two quantum states,” this question is can also be stated as “how similar are two processes?” The two questions are more related than one might think, due to the Choi-Jamiołkowski isomorphism (Choi, 1975, Wolf, 2012). For a process  $\rho \mapsto \mathcal{E}(\rho)$ , and defining the maximally entangled state between two  $d$ -dimensional Hilbert spaces  $|\Phi\rangle \equiv \sum_{i=1}^d |ii\rangle$ , we can define

$$\rho_{\mathcal{E}} = \frac{1}{d}(\mathcal{E} \otimes I)(|\Phi\rangle\langle\Phi|). \quad (2.30)$$

This construct is a valid density matrix, uniquely specified by a CPTP map. Since the process is also invertible, this is an isomorphism. In this way we can treat any quantum process in the same manner that we treat quantum states. This *Jamiołkowski matrix* is equivalent to the superoperator we used in equation 2.29 up to a simple permutation of bases. This result is crucial for the following discussion, as we can take the figures of merit given in section 2.2.5 and apply them to the process matrices.

This brief discussion of figures of merit will largely summarize (Gilchrist *et al.*, 2005), an excellent resource which we will use heavily in chapter 5. We will reproduce their key results, judging potential figures of merit along two axes. The first is whether a figure of merit should be based on the process matrix, based on the average-case behavior, or based on the worst-case behavior. The second axis is whether to use fidelity-like measures or whether to use trace distance-like measures.

Measures based on the process matrix are straightforward—simply take the process matrix in a form of your choosing and apply either the state fidelity or trace-distance calculation. We will call these quantities  $\mathcal{F}_{\text{pro}}$  and  $D_{\text{pro}}$  and return to them shortly. The average-case behavior seems quite natural. If a process ideally should take some state to another state, let us look at the fidelity (or distance) of the output state to what it should be. We can then average over all possible input states, or perhaps experimentally take some large sample. Unfortunately (Gilchrist *et al.*, 2005) show that this measure is not *stable*. If the input state is entangled with a larger space (e.g. the rest of a quantum computer) the fidelity can be misleading. The most conservative option, the worst-case behavior, is somewhat more interesting. For the fidelity this will look as a minimization over all input states and the trace distance will similarly require a maximization,

$$\mathcal{F}_{\min}(\mathcal{E}_1, \mathcal{E}_2) \equiv \min_{\psi} \mathcal{F}_{\text{state}}(\mathcal{E}_1(\psi), \mathcal{E}_2(\psi)). \quad (2.31)$$

$$D_{\max}(\mathcal{E}_1, \mathcal{E}_2) \equiv \max_{\psi} D_{\text{state}}(\mathcal{E}_1(\psi), \mathcal{E}_2(\psi)). \quad (2.32)$$

It can also be shown that this measure is stable, so this optimization need be done only over pure input states. The optimization is a convex problem and is computationally tractable for reasonably-sized systems. We will discuss this further in chapter 5, including practical matters on performing the calculation. It should also be noted that  $D_{\max}$  goes by another name, the *diamond norm*, which is used in the threshold calculations we will touch on in section 2.5.1.

(Gilchrist *et al.*, 2005) go on to relate these four quantities directly to the success probability of a larger algorithm. These measures have properties that allow us to calculate, or at least bound, the collective performance when we chain them together. We need first to define a new number from the fidelity,  $C_i(\rho, \sigma) \equiv \sqrt{1 - \mathcal{F}_i(\rho, \sigma)}$ . Additionally, since some algorithms are inherently probabilistic, we

allow them to have some “ideal” probability of error  $p_e^{\text{id}}$ . With these defined, we can get to the final results:

$$p_e \leq p_e^{\text{id}} + D_{\max}(\mathcal{E}_1, \mathcal{E}_2) \quad (2.33)$$

$$p_e \leq \left[ \sqrt{p_e^{\text{id}}} + C_{\min}(\mathcal{E}_1, \mathcal{E}_2) \right]^2 \quad (2.34)$$

$$\bar{p}_e \leq \bar{p}_e^{\text{id}} + D_{\text{pro}}(\mathcal{E}_1, \mathcal{E}_2). \quad (2.35)$$

In this manner all of three of these quantities answer reasonable questions about the effects of imperfect gates within a larger process. They bound either the worst-case or the average-case performance of the larger algorithm.

## 2.4 Measurement

After states and their evolution, the last key element of quantum mechanics is measurement. We touched on how states generate the outcome probability of an observable briefly in section 2.2, but this is only part of the story. It is useful to describe measurements with two separate notions: the simpler case of *detectors* and the more complete descriptions of *measurement processes*. When back-action is unimportant, e.g. of an ancilla that will be reset shortly, we care only about how precisely we get the classical state indicating the initial quantum information. We refer to the sort of measurement as a detector. Thus far we have spoken of the simplest form of detectors, only using examples with strong and orthogonal outcomes.

When we do care about the back-action, we speak of a measurement process, or as we will discuss in chapter 5, a *quantum instrument*. The simplest story of the measurement process is as follows: The state vector is a tool for predicting future measurements, so after one measurement we should expect that the system has been projected into a state (or space) of definite measurement outcome. If we perform an ideal measurement of  $Z$  on a qubit and get the result “1,” we should expect that the qubit state afterwards is  $|1\rangle$ . This is called a von Neumann measurement.

Any Hermitian matrix can be factored into the form

$$O = \sum_i \lambda_i \Pi_i,$$

for orthogonal and complete projectors  $\Pi_i$  such that  $\Pi_i^2 = \Pi_i$  with real measurement outcomes  $\lambda_i$ . For non-degenerate measurements,  $\Pi_i$  can be written as a pure density matrix  $|\psi_i\rangle\langle\psi_i|$ . In von Neumann measurements the state heralded by an outcome  $\lambda_i$  is given by

$$\rho^{\lambda_i} = \frac{P_i \rho P_i}{\text{Tr}(P_i \rho)}.$$

The denominator is equal to the chance of getting outcome  $\lambda_i$ , and this final state is clearly normalized. One can see by example that this does not describe all measurements. A photomultiplier tube tells whether or not there is a photon in an optical mode, but after a “click” heralding a photon in the mode the photon has been destroyed and the mode emptied. In chapter 5 we will expand in great detail on both the notion of a quantum detector and of a quantum instrument.

## 2.5 Quantum error correction

The subtleties of how measurements behave in quantum mechanics are not at all purely academic, rather this understanding is crucial for the design and operation of a quantum computer. The extraordinary promise of quantum algorithms is matched by the difficulty of creating a device to run them on. The need to isolate such a system from the environment and for the experimentalist to manipulate and measure the device are inherently at odds. The solution lies in permitting and repairing a certain amount of error by using *quantum error correction*. Measurements are critical to resolving this conundrum—they are an excellent tool for removing entropy from a quantum state. Before diving too deeply into this, it is useful to first refresh some ideas from a simpler type of error correction.

In classical error correction we encode information redundantly in an excessive number of bits, such that if some of the bits are corrupted in transmission or storage, the redundancy allows for these errors to be detected and the information restored. Most trivially, imagine a transmission channel that has a

symmetric and independent probability  $p$  of flipping each transmitted bit. If transmitting a single 0 or 1, there is a  $p$  likelihood of receiving incorrect information. One way around this is to perform a repetition code, replacing each 0 with 000 and each 1 with 111. Upon reception we take a majority vote. So if only one flip occurs, say 000 is corrupted to 100, we interpret the signal as majority 0 and the original information is recovered. In this encoding two erroneous flips, less-likely than a single such event, are required to corrupt the message. The odds of two or three errors occurring out of the three potential errors are given by  $3p^2 - 2p^3$ , which for  $p \ll 1$  is polynomially smaller than  $p$ . This repetition code is a simple example and is highly inefficient, requiring an overhead of a factor of 3 to move from  $p$  to  $p^2$ . Classical coding theory, developed in its modern form in the late 1940's, has since then created far more efficient and interesting codes.

It is not at all obvious that classical error correction can be applied to quantum information. There are four immediate concerns:

1. Qubits live in a continuous space, unlike classical bits.
2. They also have not only bit-flip or  $X$  errors, they have phase errors along the  $Z$  axis.
3. We also know that measuring any single qubit will destroy its entanglement with the rest of the computer.
4. The *no cloning theorem* disallows directly copying unknown quantum information.

With these constraints error correction looks quite hopeless indeed. For several years after the potential of quantum computing became known it was believed by some that this could never be resolved, until Peter Shor published the first true example of a quantum error correcting code (Shor, 1995). His code, essentially based on the repetition code above, require nine physical qubits to protect one *logical* qubit from either a single physical bit- or phase-flip error.

The field has grown quickly from that initial example, and we can formally resolve these four concerns in a more modern language—the *stabilizer* formalism. A stabilizer code encodes (for example) one logical qubit redundantly into  $n$  physical qubits. We then specify two special subspaces within the physical Hilbert space  $H$ : the code space and the error space. The code subspace  $C$  (for one logical

qubit) is spanned by two vectors in  $H$ ,  $|\psi_{\text{logical}}\rangle = \alpha|0_L\rangle + \beta|1_L\rangle$ , and the projector onto this space is  $\Pi \equiv |0_L\rangle\langle 0_L| + |1_L\rangle\langle 1_L|$ . A code is designed to correct a set of errors  $\{E_i\}$  (e.g. bit and phase flip errors on all qubits), such that  $E_i$  acting on a state in  $C$  leads to a state in the error space. Since the code should work if there is no error ( $I \in \{E_i\}$ ), the perhaps badly-named error space includes the code space.

It can be shown (Gottesman, 1997) that a code defined by the projector  $P$  is sufficient to correct the errors  $\{E_i\}$  if and only if

$$\Pi E_i^\dagger E_j \Pi = \alpha_{ij} \Pi, \quad (2.36)$$

for a Hermitian matrix  $\alpha$ . The full proof of this condition is interesting but will not be repeated here, rather we will discuss only several important implications. It was not given in this formal problem statement how the projection  $\Pi$  should be implemented. A result of the stabilizer condition is that we can do so by series of binary measurements, taken to be such that if all measurements results are +1 together they have implemented the projector  $\Pi$ . The stabilizer criterion further implies that we can find a series of measurements such that if some of the measurement results are not +1, and only an error in  $\{E_i\}$  has occurred, the measurement record lets us uniquely identify and correct the error.

The measurements we can use to compose  $\Pi$  are called stabilizer measurements and serve to remove entropy from the physical Hilbert space. Each of these measurements effectively slices the Hilbert space in half, projecting the physical state towards either a state in the code space, the error space, or possibly the remaining unrecoverable space. This quickly removes the first concern, that of having continuous errors. The stabilizer measurements project such that a continuous error becomes only a probability of a discrete error having occurred or not occurred, determined by the measurement record. Our second concern is resolved since the projections can be phase sensitive. The stabilizer measurements are of not only  $Z$  properties, but also of  $X$  or  $Y$  properties. The third concern is resolved because we never measure single qubits—those projections would not form a  $\Pi$  that satisfy equation 2.36. We measure instead correlations. We can also see that this construction is devised such that we never clone the state at any time.

In a more concrete example, imagine the repetition code from before. If we only cared about bit-flip



errors, we could measure *parity* operators, e.g.  $ZZI$  and  $IZZ$ . One can quickly see that this uniquely determines a bit-flip error. If both measurements are  $+1, +1$  we are in the space spanned by  $|000\rangle$  and  $|111\rangle$ , but we know nothing about the state within that space. No other subspace would yield that measurement record. Similarly  $-1, +1$  implies a flip on qubit one, without any other information.  $+1, -1$  is a flip on qubit three, and  $-1, -1$  is a flip on qubit two. This example uses its measurements efficiently. There are four errors in  $\{E_i\}$ , including no error, and there are two measurement bits, yielding four possible records. Accordingly each error is uniquely identified and there is no excess measurement. For a one-logical-qubit code that corrects both bit- and phase-flips, for  $n$  physical qubits we have  $3n + 1$  possible (single- or zero-qubit) errors on a  $2^n$  space. We remove one qubit from the space for the logical code space, leaving a  $2^{n-1}$  remaining space, which allows  $n - 1$  stabilizer measurements, or  $2^{n-1}$  possible outcomes. We need  $2^{n-1} \geq 3n + 1$ <sup>13</sup>. This equality is strictly satisfied for  $n = 5$ , meaning the minimal one-qubit error correcting code requires at least five qubits. As an illustration, the stabilizer measurements required for this code are generally given by  $\{XZZXI, IXZZX, XIXZZ, ZXIXZ\}$ . Engineering this odd and specific entanglement between the system and the measurement apparatus is the chief challenge of modern experimental quantum computing, and is the focus of chapter 6.

### 2.5.1 Fault tolerance & thresholds

By themselves, quantum error correcting codes are not sufficient for the implementation of a full quantum computer. To analyze the behavior of a realistic system we must choose a circuit-level implementation for logical initialization, manipulation, and measurement. This includes allowing for the possibility of error in all of the low-level gates and measurements. An implementation for which that can be done beneficially is called *fault tolerant*. Daniel Gottesman further defines a fault-tolerant operation as “one for which a single error introduces at most one error per block in the code” (Gottesman, 1997). This disallows any circuit-level implementation where, for example, we perform a CNOT between two qubits in one code block. If that were the case, an error on the control qubit would propagate to the target and we would have two physical qubit errors. This requires that we be very careful in how we design the implementation of a code, but, quite remarkably, it is possible.

<sup>13</sup>This is only strictly true for non-degenerate codes, a class which includes most simple examples.

Additionally, moving from an error rate of  $p$  to  $p^2$  will not be enough as the computer increases in size. In the small  $p$  limit  $p$  will increase linearly with the system size, and the statement that  $p \ll p^2$  eventually ceases to be true. Transcending this requires *concatenation* or another method of increasing the error-correcting depth. In a concatenated code the original logical qubits can play the role of physical qubits in an even larger code, which can now correct lowest-level physical errors which would previously have been uncorrectable (e.g. two simultaneous physical errors in our previous codes.). Strikingly, we see the emergence of a *threshold* (Aharonov and Ben-Or, 1999). Below a certain physical error rate we can see that for increasing concatenation depth the logical error rate is suppressed exponentially, while the system size grows only polylogarithmically. This phenomenon turns out to behave as a phase-transition in condensed matter physics. These calculations are the motivation for the trace-distance measure of process quality we discussed in section 2.3.3.

We will put off the merits of trying to describe an error rate with any single number until chapter 6, nonetheless it is a useful construct for estimating roughly what performance we will require. For the earliest concatenation codes these thresholds were as high as error rates of  $10^{-7}$ , a tall order indeed. Modern topological codes (Bombin, 2015, Fowler *et al.*, 2012, Kitaev, 2003), which are stabilizer codes though not concatenated codes, have decreased these thresholds to a much more manageable  $10^{-2}$ . Many codes are additionally designed to have practical properties like only nearest-neighbor interactions with qubits arranged in a two-dimensional space. Error rates below these thresholds have been demonstrated in small systems (Barends *et al.*, 2014, Brown *et al.*, 2011, Chow *et al.*, 2014), and recent developments have focused less on achieving these exact numbers than on creating scalable systems that have an error model fairly described by the threshold calculations.

## 2.6 DiVincenzo criteria

In the last section of this chapter, we discuss what a physical realization must look like. In 2000 David DiVincenzo wrote a set of criteria which he believed to be necessary and sufficient for a quantum computing architecture (DiVincenzo, 2008). Whether these are strictly necessary or fully sufficient can be argued, but nonetheless they provide significant guidance.

**The system must be scalable to large numbers of two-level systems, or effective two level systems.** Generally speaking, most quantum systems are not true two-level systems, rather we choose two levels to represent the  $|0\rangle$  and  $|1\rangle$  states. In trapped ion systems, for example, these qubit states are generally two of the hyperfine ground states of the ion's valence electron. In reality there are always an odd number of hyperfine ground states, and then there are fine-structure states at more energetic transitions, and at high enough energies the valence electron will be ripped away entirely. What we actually require is that when the system is in our "qubit space" we are able to manipulate the system in such a way that it does not leave that space. Our transmon qubits have many higher excited states, but the spectra of transitions that leave that space are sufficiently detuned or weak that we largely stay in the qubit space. When we get to the encoded qubits of later sections (section 7.7) these *leakage* errors will become far more important. Additionally, we will likely need many qubits, so our architecture must allow some form of *scaling* that does not overly degrade the individual qubits. This is one of the largest concerns facing the field of quantum information, as making devices more complex inherently makes them harder to control.

**The system must be intializable (either by relaxation or by measurement.)** We must also be able to begin with our devices in known states in the qubit space. If we lack knowledge of that state, we inherently begin in a mixed state, disallowing the observance of interference or quantum phenomena. There are two basic ways to initialize a system: thermalization and measurement. If we attach our qubit to a reservoir they will eventually thermalize, and if that reservoir is cold our qubit will move towards the ground state. Our transmon qubits have resonance frequencies on the order of 5 GHz and we place them in dilution refrigerators with temperatures around 20 mK. In the cold limit Boltzmann statistics would imply that the excited state population thermalizes towards  $\exp(-\hbar\omega/k_B T)$ , which for our parameters is roughly  $6 * 10^{-6}$ . In practice this is not quite what we see, but in ideal world we could simply wait and find that our system was well initialized in the ground state.

The second initialization technique is to use projective measurements. When we measure our qubit and find it to be in  $|g\rangle$ , if we did a true von Neumann measurement it would remain in that state. Similarly if we found it in  $|e\rangle$  we could apply a unitary operation to return it to  $|g\rangle$ . The second

technique requires good measurements, good unitary control, and then either post-selection or “real time feedback.” If, in the latter case, we have the ability to quickly measure and react accordingly before the system decoheres, we may then deterministically prepare  $|g\rangle$ . Having this ability is a relatively recent phenomenon. It requires high fidelity measurements and classical logic much faster than the system decoherence time. Post-selection is a much older technique, requiring only high fidelity measurements. In this schema we always do both an initial measurement and our later experiment, but we throw away the data collected from the experiment if the first measurement implied we did not start in the ground state. In the works of this thesis we use all three state-preparation techniques: thermalization, post-selection, and real-time feedback.

**The qubits must be able to maintain coherence for a sufficiently large amount of time.** As discussed previously, decoherence and unaccounted-for interactions will quickly lead us to a mixed system. This disallows interference and quantum computation. To reach the thresholds for fault-tolerant quantum computation we need fidelities of at least, say, 99.9%. This example number implies that, at minimum, we need to be able to maintain coherence for at least 1000 times longer than the duration of each operation. Finding or crafting a system that allows manipulations and measurements on a scale that much faster than they decohere is quite challenging. We must be able to reach the classical world into our device quickly and with impact at times, but keep the device completely isolated at others.

**The system must allow a universal set of gates.** Similarly we must be able to manipulate the system through its full computational Hilbert space. Generally, this means we need (at least) two orthogonal control axes for each qubit. For our systems we have full microwave control, allowing rotations around any axis in the  $X - Y$  plane. Similarly with resonant control we can do  $Z$  rotations by simply shifting the rotating frame. The more difficult requirement is that we must be able to do two-qubit rotations. Each qubit must be attached to the larger computer by some number of CPHASE or CNOT or similar operations. For universal error correcting codes each qubit needs to experience pair-wise interactions with several other qubits. Additionally, fault tolerance arguments require a large degree of parallelism to be available. We must also be able to perform many simultaneous operations across the computer.

**The system must allow some degree of qubit-specific measurement.** As previously discussed, quantum error correction requires that we perform measurements of ancillae qubits to measure very specific correlations within the data qubits. These same ancillae may also be used to initialize the system, which in some codes reduces to performing one round of error correction. Additionally we must be able to readout out the final result of our algorithm. This decoding step is generally also done by a measurement of the ancillae, though somewhat differently as the logical  $X_L$  and  $Z_L$  operations are orthogonal to the stabilizer measurements.

None of these desiderata should be surprising after the discussion of quantum error correction. In the next chapter we will discuss circuit quantum electrodynamics, our chosen physical realization, and see how it meets these criteria.

## 2.7 Conclusions

This chapter was a brief overview of the key principles of quantum mechanics and quantum information science, written primarily for a technical but not expert reader. We covered tools that we will need for later works, including fidelity measures, representations of quantum states, and formalisms for processes. We also briefly covered historical experiments that I find insightful and illustrative in the Stern Gerlach and NMR systems. The discussions of entanglement measurement have set the stage for the two main experiments of this thesis. We motivated this larger endeavor in a brief look at quantum algorithms and quantum error correction. Finally our discussion of the DiVincenzo criteria will guide the later discussions of circuit quantum electrodynamics and of experimental device design.

## Chapter 3

# Circuit QED

It is fair to state that we are not experimenting with single particles, any more than we can raise Ichthyosauria in the zoo. We are scrutinizing records of events long after they have happened.

---

Erwin Schrödinger, 1952

The previous chapter gave an outline of the goals and fundamentals of the broader field of quantum information, but gave no prescription for how to physically realize those concepts. In this chapter we will get closer to answering that question. There are many different approaches. One of the oldest, discussed in chapter 2, is the use nuclear spins in a molecule, or nuclear magnetic resonance (NMR) (Vandersypen and Chuang, 2005). While a highly educational tool for physicists and quite useful in the modern medical application of Magnetic Resonance Imaging, NMR does not have serious potential for scaling to control of larger quantum systems. Another classic approach is to utilize single photons propagating in either free space or an optical fiber—encoding a qubit state in either the polarization, frequency, or temporal or spatial position (Kok *et al.*, 2007). Unfortunately it is quite difficult to perform non-destructive measurements on free photons, and certainly difficult to keep them around for the repeated manipulations. Even more concerning is that it is challenging to produce them on-demand, or deterministically.

Computing with electromagnetically trapped ions is perhaps the most mature quantum computing

architecture. The hyperfine states of trapped ions can have exquisite coherence properties (Langer *et al.*, 2005) and very high fidelity gates have been demonstrated (Ballance *et al.*, 2016, Brown *et al.*, 2011). They are attractive in no small part because all of the qubits are essentially identical, comparing favorably to our superconducting qubits, each of which requires quite a bit of individual calibration. They have reasonable (though not perhaps not excellent) prospects for scalability, especially with the introduction of surface-electrode traps (Seidelin *et al.*, 2006). Even with only small local collections of qubits, trapped ion systems may be a promising base for a module-based system (Monroe and Kim, 2013). Other modern approaches are the use of quantum dots (DiVincenzo *et al.*, 2000, Loss and DiVincenzo, 1998) and topological systems (Kitaev, 2003). Dots are presently less mature than superconducting qubits but share many of the same benefits (and perhaps more) with regards to scalability. Topological systems have incredibly interesting physics, but technical difficulties presently render them largely theoretical constructs.

Superconducting systems, our chosen architecture, are perhaps the most mature technology after trapped ion computing. Their chief advantages are their engineerable and lithographically defined Hamiltonians. We also have the advantage of control and manipulation with microwaves, which have better noise properties than DC control (as in most quantum dot schemes) and are much cheaper and less finicky than lasers (as used in most trapped ion schemes). Because they are simply electrical circuits, it is quite easy to couple them to other qubits and systems, but these same couplings are also their largest shortcoming. As they also couple easily to the environment, they exhibit relatively short coherence times compared to more naturally isolated systems. That said, their coherence properties have been increasing rapidly (Paik *et al.*, 2011, Reagor *et al.*, 2013), and they are already approaching the performance needed for fault tolerance thresholds (Barends *et al.*, 2014). In this chapter we will lay down the basic physics relevant to our superconducting devices.

As these physics have been discussed in many excellent reviews of circuit quantum electrodynamics (Bishop, 2010, Girvin, 2011, Schuster, 2007), I will not attempt to provide an exhaustive survey here. Rather the goals of this chapter are threefold. My first goal is to provide a base of knowledge appropriate for understanding the later chapters of this thesis. The second goal is to provide my view on certain aspects of the relevant physics when my take somewhat different from the common or published version.

---

Lastly, I aim to point out areas that were not covered in the normal surveys, but which I wish I had been exposed to earlier in my graduate career. Along the way I will provide citations which I feel give the explanations I found most useful, rather than only the original publications.

Circuit quantum electrodynamics (cQED with a lowercase c) is the playground in which all of the games in this thesis have been played. Broadly, cQED refers to the behavior of electrical circuits in the quantum regime where “artificial atoms” interact with linear resonators<sup>1</sup>. The circuits relevant to this thesis can be separated into two categories: almost linear or nonlinear. The nonlinear modes are often restricted to their lowest two energy levels and called qubits or artificial atoms, and the roughly linear modes are called resonators or cavities. The qubit-like modes available to circuit QED systems come in many forms and with even more names: charge qubits (Bouchiat *et al.*, 1998), flux qubits (Mooij *et al.*, 1999), phase qubits (Martinis *et al.*, 2002),  $0-\pi$  qubits (Kitaev, 2006), transmons (Koch *et al.*, 2007), fluxonium (Manucharyan *et al.*, 2009), the various types of quantum dots, and even non-circuits like NV centers, atomic ensembles, or single neutral atoms. This thesis will be concerned exclusively with transmons, which we will dive into in section 3.4.

It is useful (and traditional) to first discuss a more idealized model system consisting of real two-level systems and linear cavity modes, which we will do in section 3.1. After this we will continue on to the physics of quantum electrical circuits, and introduce non-linearity to those circuits via the Josephson effect. These concepts will get combined in our discussion of the transmon qubit itself in section 3.4. In recent years the Yale group has begun exploiting the surrounding electronic environment of the transmon in a much deeper manner. Rather than treating the resonant modes around the transmon as accessories, their own quantum dynamics have become key, and in some experiments relegate the transmon mode itself to an auxiliary system (Ofek *et al.*, 2016). In section 3.5 we will walk through predicting the behavior of that environment. This will include how to express arbitrary impedance environments as manageable electrical circuits.

In section 3.6 we will begin talking about some non-idealities and how our devices will eventually thermalize with the environment. In particular I will take care to point out some common misconceptions, and hopefully add a little intuition. We will wrap up with how we intentionally couple our systems to

---

<sup>1</sup>We should note that observations of quantum behavior in electrical circuits predates the invention of cQED (Koch *et al.*, 2007)



the outside world to manipulate and measure our devices.

### 3.1 Introduction to Cavity QED & Jaynes-Cummings

The spiritual predecessor of circuit QED is cavity quantum electrodynamics (CQED, with a capital C), which is a game played using the interaction of the electronic or spin state of an atom and a linear photonic mode. It can be argued that CQED was first proposed in 1945 by Edward Purcell (Purcell, 1945), who postulated that the decay rate of a polarized nuclear spin may be enhanced by placing it in a resonant electrical circuit. This “Purcell effect” will be discussed in detail in section 3.6, though we will put off the more practical considerations needed during device design until chapter 4.

As with many fields of science it is difficult to pin down an exact inventor or first experiment, but many would agree that the more modern notion of cavity quantum electrodynamics was pioneered by Serge Haroche and collaborators in the early 1980s (Haroche, 2013). This was (and remains) one of the original testbeds of non-ensemble quantum experiments<sup>2</sup>. For this work Haroche won the Nobel Prize in 2013, along with David Wineland who did similarly pioneering work in trapped ions, which are in some ways a dual system. Haroche’s experiments centered around manipulating the quantum state of an optical cavity by carefully interacting it with Rydberg atoms, themselves prepared in a known quantum state. Similarly, by looking at the back-action of the field on these transiting atoms we can make measurements of the cavity state.

The phenomena explored in cavity QED experiments can be explained by a much older Hamiltonian, previously relegated to thought experiments. In 1963 Edwin Jaynes and Fred Cummings devised what is now called the Jaynes-Cummings Hamiltonian ( $\mathcal{H}_{JC}$ ). They proposed examining a two level system (historically referred to as a spin, in homage to NMR experiments like those of Purcell) and a harmonic system (or resonator), which share a dipole interaction. In CQED (cQED) these resonators are optical (microwave) photonic cavities and the qubits are the lowest two levels of (artificial) atoms.  $\mathcal{H}_{JC}$  is given

---

<sup>2</sup> “Non-ensemble quantum experiment” is to contrast with earlier quantum experiments like scattering effects and NMR.

by

$$\mathcal{H}_{\text{JC}}/\hbar = \frac{\omega_q}{2}\sigma_z^q + \omega_c a_c^\dagger a_c + g \left( a_c^\dagger + a_c \right) \sigma_x^q \quad (3.1)$$

$$= \frac{\omega_q}{2}\sigma_z^q + \omega_c a_c^\dagger a_c + g \left( a_c^\dagger + a_c \right) \left( \sigma_+^q + \sigma_-^q \right) \\ \sim \frac{\omega_q}{2}\sigma_z^q + \omega_c a_c^\dagger a_c + g \left( a_c^\dagger \sigma_+^q + a_c \sigma_-^q \right) \quad (3.2)$$

Equation 3.1 contains an oscillator, a spin, and a dipole interaction which makes it energetically favorable for the spin to align with the cavity field. In equation 3.2 we have made the first-order rotating wave approximation (RWA), which holds when the coupling rate between the resonator and qubit modes is far smaller from their absolute frequencies. The strength of the  $g$  term can be calculated directly from the geometry of the resonator and local electric field per photon, the dipole moment of the atom, and the angle by which they are aligned. We will see soon that  $\mathcal{H}_{\text{JC}}$  is no longer a sufficiently accurate description of the cQED systems discussed later, but a great deal of insight can be gained nonetheless.<sup>3</sup>

We can look at the exact diagonalization of  $\mathcal{H}_{\text{JC}}$ , which results in fixed stationary states which mix or “dress” the cavity and qubit states,

$$|+, n\rangle = \cos(\theta_n/2) |e, n\rangle + i \sin(\theta_n/2) |g, n+1\rangle \quad (3.3)$$

$$|-, n\rangle = \sin(\theta_n/2) |e, n\rangle - i \cos(\theta_n/2) |g, n+1\rangle, \quad (3.4)$$

where the mixing angle  $\theta_n \equiv \arctan(\Omega_n/\Delta)$ . The denominator  $\Delta$  is the cavity-qubit detuning and the numerator is the scaled interaction strength  $\Omega_n = 2g\sqrt{n+1}$ . Neglecting the zero-point energy these states have eigenenergies

$$E_n^\pm/\hbar = n\omega_c \pm \frac{1}{2}\sqrt{\Delta^2 + \Omega_n^2}. \quad (3.5)$$

In the “resonant strong coupling” regime, where the qubit and cavity are resonant and their decoherence rates are slower than the interaction rate, the  $g$  term will lead to direct energy swapping between modes.

<sup>3</sup>Circuit QED allows both strong confinement of the resonator mode and much, much larger dipole moments than true atomic systems. Accordingly in circuit QED we can reach regimes of much higher  $g$  relative to all other time- and frequency-scales than cavity QED experiments.

As we see from equation 3.3, when  $\Delta = 0$  the new modes are an equal superposition of the bare modes, so an excitation in one of the bare modes will oscillate between the eigenmodes. This effect, called vacuum Rabi oscillation, was first observed in atoms in 1983 (Haroche, 2013) and in artificial atoms in 2004 (Wallraff *et al.*, 2004). The effect remains useful to this day, realizing the SWAP-gate building block in some quantum information processing architectures where they are able to *in situ* vary the frequency of at least one mode and precisely achieve resonance (Haack *et al.*, 2010).

The work of this thesis, however, is with fixed-frequency modes, and we work in the detuned regime as we are interested in more subtle interactions. We use fairly large detunings and depending on the application choose decoherence rates such that we are in either the “strong dispersive” regime or “weak dispersive” regime. In both regimes we have  $\Delta \gg g$ , and we can explicitly work in that limit using the RWA and second-order perturbation theory. Expanding in  $g/\Delta \ll 1$  leads to an effective Hamiltonian

$$\begin{aligned} \mathcal{H}_{\text{dispersive}}/\hbar &= \left( \omega_c - \frac{2g^2}{\Delta} \right) a_c^\dagger a_c + \omega_q |e\rangle\langle e| \\ &= \omega_c a_c^\dagger a_c + \left( \omega_q - \frac{2g^2}{\Delta} a_c^\dagger a_c \right) |e\rangle\langle e|, \end{aligned} \quad (3.6)$$

though we will later generally replace  $2g^2/\Delta$  with the “dispersive shift”  $\chi$  of more general origins. In the strong dispersive regime,  $\chi$  is larger than any relevant decoherence rates, meaning that the relevant transitions can be individually manipulated before the state decoheres. This is not true in the weak dispersive regime. As we will see in the next few sections, we often work in a regime where the circuit QED Hamiltonian is well-approximated by the Jaynes-Cummings Hamiltonian.

## 3.2 Quantum electrical circuits

To start we will discuss the linear modes in circuit QED. For an exhaustive pedagogical study of quantized circuits, including those with loss and non-linear elements, there is no better resource than (Devoret, 1995). We will content ourselves here to work through perhaps the simplest example, the parallel LC oscillator. Noting the correspondence between the inductor current and the capacitor charge, we can




Figure 3.1: **Spectra in the strong dispersive regime.** In figure **a** we see the spectrum of the qubit transition when there are photons in the cavity. In this case a coherent state with a Poissonian distribution in photon number, here with average photon number  $\bar{n} \sim 2$ . The width of each transition is roughly  $1/T_2$  of the qubit. In figure **b** we see the the transition spectrum of the cavity when the qubit is half excited. The widths of these peaks are given by the decoherence time of the cavity which is generally equal to the relaxation rate  $\kappa$ .

can write down the classical Lagrangian

$$\mathcal{L} = \frac{1}{2}L\dot{Q}^2 - \frac{1}{2C}Q^2,$$

which yields the equation of motion  $\ddot{Q} = -\omega^2 Q$  for  $\omega \equiv \sqrt{1/LC}$ . The conjugate variable to  $Q$  is given by  $d\mathcal{L}/d\dot{Q} = \Phi$ , which is precisely the magnetic flux through the inductor. We can then write down the Hamiltonian

$$\mathcal{H} = \frac{\Phi^2}{2L} + \frac{Q^2}{2C}, \quad (3.7)$$

and introduce quantum mechanics by imposing the canonical commutation relation  $[\Phi, Q] = -i\hbar$ . As in any introductory text on quantum mechanics we can create appropriate raising and lowering operators and write

$$\mathcal{H}_{\text{linear}} = \hbar\omega a^\dagger a, \quad (3.8)$$

again neglecting the zero-point energy. Any linear circuit (consisting of only inductors and capacitors) can be written as a sum of such independent harmonic oscillator modes—a fact we will touch on classically in section 3.5.4. Unfortunately these linear equations of motion do not permit behavior particularly different from classical systems, so we turn next to more interesting elements.




Figure 3.2: **Aluminum/aluminum oxide/aluminum Josephson junction imaged by a scanning electron microscope.** An aluminum trace on top of a sapphire substrate comes in from the left, and a larger trace from the right. Near the center they overlap, separated by the natural oxide grown on the bottom layer in a pure oxygen environment. In reality both the far left conductor and right right conduct are also bilayers and therefore Josephson junctions, but as the traces are sufficiently large that the capacitance overwhelms any nonlinear effect. Image by Yiwen Chu and reproduced with permission.

### 3.3 The Josephson effect

How can we add some non-linearity to our circuits<sup>4</sup>? Trivially, dissipation is non-linear, but is generally undesirable. As loss implies coupling to an external reservoir, it inevitably guides the system to a state of thermal equilibrium, rarely leading to the phenomena we wish to study. Superconductivity solves many of our problems. Aside from allowing low-dissipation capacitances and inductances, it provides lossless non-linearity via the Josephson effect.

There are many excellent derivations of the Josephson effect and other phenomena in superconductivity (Tinkham, 2004, Van Duzer and Turner, 1999) and we will not reproduce them here. Rather we will provide an operational description. The Josephson effect utilized in this thesis comes from three-layer Josephson junctions (JJ) consisting of superconducting, insulating, and superconducting layers. When the insulator is thin enough, the supercurrent can tunnel from one superconductor to the other. In practice, we use junctions of aluminum, aluminum oxide, and aluminum, as depicted in Fig. 3.2.

The result is a new circuit element with a novel current-flux relationship. Whereas in an inductor

<sup>4</sup>What do we mean by non-linear? What is linear in what? A Hamiltonian containing only squares of its variables and their derivatives is a linear system. While the Hamiltonian is quadratic, the Euler-Lagrange equations relate the time derivatives of those variables directly to their magnitude. The equations of motion of such a system are linear in the state of the system, implying that elements of a superposition each evolve independently.

the current scales as  $I = \Phi/L$ , in a Josephson junction the current scales as the cosine of the flux,

$$I(t) = I_c \sin(\Phi(t)/\Phi_0)$$

where the critical current  $I_c$  is the parameter that characterizes the junction and the constant  $\Phi_0 = h/2e$  is the magnetic flux quantum. Another important quantity is  $\phi(t) = \Phi(t)/\Phi_0$ , the superconducting phase difference across the junction. We are implicitly constraining ourselves to phenomena where the current is below the critical current. In this work we are generally directly focused on the Hamiltonian dynamics, so we can also describe a Josephson junction with

$$\mathcal{H}_{JJ} = E_J \cos \phi. \quad (3.9)$$

There are many different electrical circuits that exploit Josephson junctions to create non-linear quantum-mechanical behavior, but only the single-junction transmon is within the scope of this thesis.

### 3.4 Transmon physics

In many ways the transmon (Houck *et al.*, 2009, Koch *et al.*, 2007) is the simplest superconducting qubit, ideally consisting of only a Josephson junction and a shunting capacitance<sup>5</sup>. In this circuit there are two design parameters to choose from,  $E_J$  and  $C$  or  $E_c$ . The Hamiltonian for this circuit is given by

$$\hat{H} = 4E_c (\hat{n} - n_g)^2 - E_J \cos \hat{\varphi}, \quad (3.10)$$

where we have introduced  $\hat{n}$  as the quantized difference in the number of Cooper pairs on the two sides and we have broken with convention and used hatted variables to distinguish that  $n_g$  is a classical variable. While  $n_g$  is a control knob and can be used to deterministically change the Hamiltonian, it may also be a noisy variable, beyond our control. These effects are crucial for the Cooper pair box qubit (Schuster, 2007), but irrelevant in our case, as we will see. The Cooper pair box and the transmon share

<sup>5</sup>In the original paper the transmon used a pair of junctions forming a SQUID. For our purposes this distinction is unimportant, and we will treat only the single-junction transmon.

an identical circuit diagram, but the transmon exists in a regime of  $E_J/E_c \gg 1$ , generally over 100 but below the  $10^4$  seen in phase qubits. When  $E_J$  is very large relative to  $E_c$ , it becomes energetically preferable to keep flux excursions across the junction low. This leads to the cosine flux-dependence of the junction behaving increasingly like a parabolic potential, and in the infinite limit creates a harmonic oscillator as in equation 3.7.

The Hamiltonian in equation 3.10 has an exact solution in the phase basis in terms of Mathieu functions (Koch *et al.*, 2007). Without using that result in its entirety, we note several important facts about the  $E_J \gg E_c$  regime. First, the dependence of the spectrum (e.g. the energy splitting of the ground and first excited state) on the classical gate charge  $n_g$  is suppressed exponentially in  $E_J/E_c$ . This suppression of the “charge dispersion” removes a control knob used by the Cooper pair box, but in exchange also removes that control from the untamed wilds of the noisy electromagnetic environment, most notably  $1/f$  noise and offset charges. This allows much longer dephasing times than were achievable in the lower  $E_J/E_c$  regime.

The second fact we note here concerns the anharmonicity  $\alpha$ —the difference between the qubit state transition frequency and that of the next available transition. As the transmon approaches a harmonic oscillator where that difference is zero, we pay a price in the addressability of states. The spectrally narrower pulses we will need to address the transitions selectively will require linearly more time to perform. This cost is worthwhile, however, as the relative anharmonicity decreases only polynomially in  $E_J/E_c$ , while the charge dispersion is suppressed exponentially. This means there is a region in parameter space where the charge dispersion is suppressed before the anharmonicity becomes too small. From the exact solutions in (Koch *et al.*, 2007) we find that the absolute anharmonicity approaches  $E_c$ . As the frequency is roughly  $\sqrt{8E_J E_c}$ , the relative anharmonicity goes as  $\alpha_r \approx (8E_J/E_c)^{-\frac{1}{2}}$ .

One useful reference for the work of this thesis is figure 3.3, where we compare the maximum impact of any charge dispersion with the anharmonicity. The maximum charge dispersion  $\epsilon$  is plotted as a divisor of the qubit frequency, forming a “phase Q”, and the anharmonicity is plotted relative to the qubit frequency, forming a  $\alpha_r \equiv \alpha/\omega$ . Put in these units, we can plot both the phase Q and  $\alpha_r$  as a parametric function of  $E_J/E_c$  regardless of the absolute frequency. This graphic was generated using an exact diagonalization of the Hamiltonian rather than the transmon approximation, though we

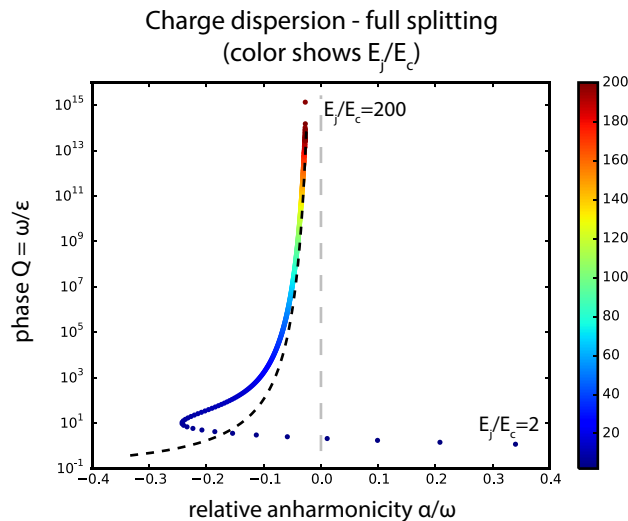


Figure 3.3: **Transmon charge dispersion and anharmonicity vs.  $E_J/E_c$ .** In this plot the logarithmic  $y$  axis is the qubit transition frequency divided by maximum the charge dispersion between the ground and excited states. This “phase Q” is a worst-case effect, making no assumptions of the charge noise spectrum. The  $x$  axis is the relative anharmonicity, or  $\alpha_r = (\omega_{12} - \omega_{01})/\omega_{01}$ . We plot both parametrically in  $E_J/E_c$ , and overlay the transmon approximation as a dashed line, which becomes increasingly accurate as  $E_J/E_c > \sim 100$ . This result was generated from a full diagonalization of the Hamiltonian in equation 3.10.

have overlaid the transmon approximation as the dashed line. We will discuss the physical realizations of transmons and resonators later in chapter 4. In particular the trade-offs related to charge dispersion will be discussed in section 4.1.1.

### 3.5 Black-box quantization

With our linear and non-linear circuit elements in hand, we can ask what happens when we embed a transmon in an otherwise linear electrical circuit. We could write down the full Hamiltonian and diagonalize it following the prescription in (Devoret, 1995), but there is another way if we are willing to neglect the effects of charge dispersion. In a system where the nonlinear elements are only transmons, the fact that the system is relatively close to linear allows an approximation which we call black-box quantization (BBQ) (Nigg *et al.*, 2012). In this model we start by separating the Josephson junction into



a linear inductance, second order in  $\phi$ , and all higher order terms into a parallel “spider element”<sup>6</sup>. We then diagonalize the linear system by itself, and only later add in the nonlinear element as a perturbation and re-diagonalize. In the following discussion we will discuss only the non-degenerate case.

We first treat the undriven, single-junction case of BBQ, which begins by looking at the world from the point of view of the spider element, which we will attach to the rest of the circuit. The rest of the circuit is linear and can be expressed as some linear network of inductors and capacitors. To do the classical diagonalization first, we want to express this linear system in a way where we can immediately see the uncoupled modes.

To do this we turn back to the early days of electrical engineering and look at some of the fundamental properties of circuits. All passive one-port impedances are positive real functions of the Laplace variable  $s = i\omega + r$  (Brune, 1931). This means that for physical drives the system will store energy reactively or dissipate energy rather than generate energy. For the present section we will also consider these circuits to be lossless, in which case we can always decompose the response function into Foster’s first form (Foster, 1924), as depicted in the right half of figure 3.4b. In this elegant representation the modes are uncoupled, and each parallel  $LC$  oscillator acts independently<sup>7</sup>.

It is not always clear how to express an impedance in one of Foster’s forms—for example three-dimensional resonators do not have an immediately obvious circuit equivalent, but they certainly do. This is a fact we will elaborate on in section 3.5.4. In practice we generally obtain this response function numerically via finite-element simulation, the practical matters of which will be discussed in section 4.5. This requires that we be able to translate a numerical impedance into a sufficiently accurate circuit.

For a given  $Z(\omega)$  how do we find the equivalent circuit? There are many approaches, but fundamentally knowledge of the poles and zeros is sufficient to describe the system. The location of each pole gives us the product of LC as  $\omega_i = (L_i C_i)^{-1/2}$ . One approach to obtain more information is in the residue of each pole, given by  $i/2L$ . More operationally, the slope at the zero-crossing in  $Y$  (poles in  $Z$ )

<sup>6</sup>We use a parallel spider element instead of a series element because it is defined in terms of flux (voltage) instead of current. See (Manucharyan *et al.*, 2007).

<sup>7</sup>It is worth noting that such circuits can also be expressed in Foster’s second form, which is a parallel assembly of series LC circuits. Both forms result from the fact that any suitable reactance is a positive function with  $n$  poles and  $n + 1$  roots. Accordingly it can easily be expressed as a ratio of polynomials, which are already factored by knowledge of the poles and roots. Since  $X$  is a positive function but has poles and roots symmetric in frequency, we can use partial fraction decomposition and find  $X$  in the form  $X = \sum_i c_i \omega / (\omega_i^2 - \omega^2)$ , and therefore as a series of parallel oscillators. Similarly  $Y = 1/Z$  can be decomposed in the same manner, so we can also find a parallel collection of series LC circuits.




Figure 3.4: **Black-box quantization circuits.** In figure **a** the non-linear element of the Josephson junction is attached to an arbitrary impedance which includes the linear inductance of the junction. Figure **b** is a graphical depiction of the same junction attached to a circuit in Foster's first form. Any lossless one-port circuit can be expressed as a series of parallel LC oscillators, allowing for a zero-capacitance (resonant at zero frequency) and a zero-inductance mode (resonant at infinite frequency). Expressing an impedance in this form is essentially finding its normal modes. This formulation leads to an intuitive way to see the voltage that each resonant mode can drive across a port, or in this case across the Josephson junction. It is important to note that these results are only strictly true for undriven systems and rely on the assumption that we are solidly within the transmon regime.

also contains the relevant information. Working in Foster's first form, the zero crossing is given when the individual term in the denominator goes to infinity, so we can examine that subcircuit in  $Y$  alone.  $dY/d\omega$  evaluated at  $\omega_i$  gives  $i2C_i$ . In whatever manner you prefer, a numerical  $Z(\omega)$  for a one-port device can be used to generate a circuit in Foster's first form.

As these modes are independent we can immediately see from Kirchhoff's laws that the voltage across the excited modes adds vectorially to yield the voltage across the junction. To move to the quantum regime—knowing the number of photons in each mode yields the flux across the spider element. This calculation is straightforward,

$$\begin{aligned} \hbar\omega_i &= \frac{\phi_i^2}{2L} \mapsto \phi_i^2 = \frac{2\hbar}{\omega_i C_i} \\ \phi_i &= \sqrt{\frac{\hbar}{2 \operatorname{Im}[Y'(\omega_i)]}}. \end{aligned} \quad (3.11)$$

To recap, this means that a simulation of the linear system is sufficient to understand how photons in these linear modes will participate in the junction. To obtain the full Hamiltonian we add in the previously linearized modes ( $\mathcal{H}_{\text{linear}}$ , with “bare” resonance frequencies found in the last step) and examine the

total flux across the junction nonlinearity due to excitations in these modes,

$$\mathcal{H}_{\text{linear}}/\hbar = \sum_i \omega_i a_i^\dagger a_i \quad (3.12)$$

$$\hat{\phi} \equiv \sum_i \phi_i (a_i + a_i^\dagger) \quad (3.13)$$

$$\mathcal{H} = \mathcal{H}_{\text{linear}} + E_J \left( \cos(\hat{\phi}) + \hat{\phi}^2/2 \right). \quad (3.14)$$

Note that in equation 3.14 we have added the second order terms back in. This is because we have already accounted for the linear inductance terms in the resonance frequencies of  $\mathcal{H}_{\text{linear}}$ , as the standard approach assumes that the linear inductance was included in the simulation used to generate  $Z(\omega)$ . There are other benefits to that approach which will be discussed in section 4.5.

One question (which will become more relevant in the discussion of multi-qubit black-box quantization in section 3.5.2) is whether or not we are being reasonable in treating  $\phi_i$  as a real number. In a truly general treatment (Devoret, 1995) we must allow  $\phi_i$  to be complex. We are justified in treating  $\phi$  as real because we will be working in a chosen rotating frame for each  $a_i$  and because we have assumed a lossless system, so this phase can be fixed. We have made an arbitrary choice of starting phase for each mode, which is an irrelevant freedom in the single-junction case.

After adding in these higher order terms of the cosine we can perform a second diagonalization, leading to a non-linear spectrum. We can calculate the eigenmodes of  $\mathcal{H}$  via a full numerical diagonalization, or analytically using the rotating wave approximation. The full numerical diagonalization is straightforward, and will be more accurate in all cases, especially when the Hamiltonian results in near frequency collisions. That said, the analytical approach provides considerably more intuition.

The majority of early experiments in circuit QED, and roughly all experiments in cavity QED, depend entirely on fourth order terms. The simplest experiments typically include one qubit mode and one cavity mode, and have three relevant fourth order terms,

$$\mathcal{H}_{\text{NL}}^4 = -\frac{\alpha}{2} a_q^\dagger a_q^\dagger a_q a_q - \frac{K}{2} a_c^\dagger a_c^\dagger a_c a_c - \chi_{qc} a_c^\dagger a_c a_q^\dagger a_q. \quad (3.15)$$

The first two terms refer to the anharmonicities of the qubit and cavity respectively, which for convenience

are given the special names of simply “anharmonicity” and “self-Kerr.” The third term is the dispersive shift (occasionally called “cross-Kerr”), which in a truncated Hilbert space is equivalent to the term we found in the dispersive Jaynes-Cumming Hamiltonian. We will briefly discuss the normal ordering, sign conventions, and prefactors in the next section.

If we expand to the cosine in equation 3.14 to fourth order, assume a large detuning  $|\omega_q - \omega_c| \gg “g”$  and take the first order RWA, and are careful with normal ordering, we can write the value of these terms down exactly,

$$\alpha(K) = \frac{E_J}{2} \phi_{q(c)}^4$$

$$\chi = E_J \phi_q^2 \phi_c^2.$$

This approximation also yields the useful result that  $\chi = \sqrt{2\alpha K}$ .

Another common device type Kirchmair *et al.* (2013), Ofek *et al.* (2016), Vlastakis *et al.* (2013) has a single qubit which couples to two resonators which we will label  $c$  as before and also  $r$  for “readout.” The fourth order terms presented already remain the same, but we gain three more fourth order terms:  $K_r$ ,  $\chi_{qr}$ , and  $\chi_{rc}$ . The derivations for the  $r$  terms are identical. The most interesting of the new terms is the dispersive shift between the two cavities, or simply “cross-Kerr” colloquially. This term is generally undesirable—when we aim to measure the transmon we may, in a subtle way, be weakly measuring photon number. We also note the relationship  $\chi_{cr} = \chi_{qr}\chi_{qc}/2\alpha$ .

Further relevant physics arise in the sixth order expansion, particularly the qubit-state dependent cavity Kerr, which is equivalently the number-state dependent qubit-cavity dispersive shift,

$$\mathcal{H}_{\chi'} = \frac{\chi'}{2} a_q^\dagger a_q a_c^\dagger a_c^\dagger a_c a_c.$$

### 3.5.1 The conventions of $\mathcal{H}$

As a brief aside, we will discuss the conventions by which these Hamiltonian terms are defined. Note that the terms are always given in normal order and occasionally have an additional numerical prefactor. We also alternate the sign versus order in how the terms are added to  $\mathcal{H}$ , the same way the cosine expansion




Figure 3.5: **The black-box Hamiltonian expressed as interactions.** One useful way to visualize the Hamiltonian and transition energies in these conventions is as interactions. **a**, When only one photon is present in the qubit, the total system energy is  $E_{e0}/\hbar = \omega_Q$ . **b**, When a cavity photon is added, we gain that energy, and the (negative) interaction energy,  $E_{e1}/\hbar = \omega_Q + \omega_C - \chi$ . The transition frequency was  $\omega_C - \chi$ . **c**, When adding a second cavity photon we now have three two-body interactions, and a new three-body interaction. The qubit interacts with both cavity photons, and they interact with each other. Now  $E_{e2}/\hbar = \omega_Q + 2\omega_C - 2\chi - K + \chi'$ .

alternates, leading to terms with generally positive values<sup>8</sup>. The reasoning behind this has two principles. First, adding higher order terms should not affect preexisting lower-order behaviors. Second, the value of the terms should reflect the entirety of the change in a directly measurable spectrum.

If, for example, we defined Kerr in terms of  $-K (a^\dagger a)^2$ , when we add this term it would change the lower-order transition frequency to add one photon,  $\omega a^\dagger a$ . With normal-ordering this term has no effect until two photons are present. In some references the number we define as  $\chi$  is occasionally given as  $2\chi$ . This alternate definition is better-motivated in a true Jaynes-Cummings Hamiltonian, but never reflects a full energy difference. The occasional prefactor of  $1/2$  is also required to make  $K$  and  $\alpha$  reflect a difference in spectra. Kerr defined with these conventions is also the full difference in the energy spacings between Fock states  $|0\rangle$  and  $|1\rangle$  and states  $|1\rangle$  and  $|2\rangle$ . In a slightly more general treatment for our Hamiltonians, for every mode  $i$  that appears  $n_i$  times in a Hamiltonian term, the prefactor gains an additional factor of  $1/n_i!$ . This is easily seen from the behavior of the annihilation and creation operators.

<sup>8</sup>In some cases when the RWA does not hold the values can switch sign, but in this convention the parameters are generally positive. This effect is more pronounced for higher order terms.



Figure 3.6: **Black-box quantization with multiple qubits.** We assume the black-box to be described by an  $n$ -by- $n$  matrix  $Z(\omega)$ , where the spider element for each of the  $n$  junctions sees a different port.

### 3.5.2 Black-box quantization with several transmons

What happens when we have a similar system, but with more than one junction? In particular we examine the case of a system with multiple single-junction transmons<sup>9</sup>. Consider a multiple port impedance, where each junction has its own port and view of the shared microwave environment, depicted in figure 3.6. As in single-junction black-box quantization, we simulate with the junction capacitances and the linear inductances included “in the box.”

Analogously to single-junction black-box, we will simulate the box numerically, then ask what happens when we introduce weak non-linearities at the outside ports. We need to know how excitations in the linear modes will drive flux across *each* of the spider elements, and treat each spider element as an individual perturbation. After simulating the  $n$ -port circuit, with one port per junction, we have an  $n$ -by- $n$  impedance matrix  $Z(\omega)$ . We now wish to calculate the flux coefficient for each junction  $l$  and mode  $p$ , with the goal of ending with a Hamiltonian

$$\hat{\phi}_l = \sum_p \phi_p^l (a_p + a_p^\dagger) \quad (3.16)$$

$$\mathcal{H}/\hbar = \sum_p \omega_p a_p^\dagger a_p + \sum_l E_{J,l} \left[ \cos \hat{\phi}_l + \frac{\hat{\phi}_l^2}{2} \right] \quad (3.17)$$

We will first derive this as Nigg et al. do in the original proposal (Nigg *et al.*, 2012). We choose a reference junction,  $k$ , and perform the single-junction black-box quantization procedure using  $Z_{kp}^{\text{eff}} = \frac{2}{\omega_p \text{Im} Y_k'(\omega_p)}$  where now  $Y_k = 1/Z_{kk}$ . To expand on this, from the point of view of the reference junction we are only considering the linear effect of the other junctions. The other ports are ‘opens’ in this

<sup>9</sup>This derivation will not cover SQUID loops, as in that system the spider elements are not at all perturbations of each other and an additional constraint is needed to reflect the surrounded flux.

approximation, and the other spider elements carry no current. If there are no currents on the other ports then the flux across port  $k$  can only come from  $I_k$ , as definitionally  $V_i = Z_{ij}I_j$ . Ergo we can perturbatively understand the relevant behavior by considering  $Z_{kk}$  as a single-port circuit.

We then use the transfer impedance to figure out the relative flux across the other ports. When there is current on port  $k$ ,  $Z_{lk}/Z_{kk}$  is the ratio of voltages (and fluxes) between ports  $l$  and  $k$ . Therefore when there is one photon in mode  $p$ , we know how it drives voltage across port  $l$  relative to how it drives voltage across port  $k$ . We can now calculate the flux operator for each junction

$$\hat{\phi}_l^{(k)} = \sum_p \frac{Z_{lk}(\omega_p)}{Z_{kk}(\omega_p)} \sqrt{\frac{\hbar}{2}} Z_{kp}^{\text{eff}} (a_p + a_p^\dagger) \quad (3.18)$$

The primary difficulty with this approach is that generally every matrix element of  $Z$  diverges at each resonance frequency, and we cannot use the slope of  $Y_{lk}(\omega_p)$  in analogy to single-junction black-box quantization. Accordingly additional steps are required to accurately obtain that ratio. Another practical concern arises when a junction is only very weakly coupled to the reference junction. The numerical precision with which we are able to calculate the coupling to another other junction may suffer.

An alternate method is to note that there is nothing special about the choice of reference port, and simply to perform “single-junction” black-box quantization on each port, which only requires the diagonal elements of  $Z(\omega)$ . This works except for one subtlety, which is to note that that off-diagonal  $Z_{lk}$  is not a positive real function and can have poles with a negative residue. This means that a set of independent “single-junction” BBQ analyses will get the magnitudes of  $\phi_p^l$  correct, but may err on the sign. We have a choice of freedom in these phases. When we analyze the first junction  $k$ , we choose the phases on  $a_p$  such that they are all positive. We also can look at the first mode  $p_0$ , and choose the phase convention for all ports  $k$  such that they are positive. After that, we are constrained. Figure 3.7 gives a concrete example of such a situation. To handle this, we use only the *sign* of the residues in  $Z_{lk}(\omega_p)$  to set the phases of  $\phi_l^p$  relative to the phases of the “first” mode  $\phi_k^p$ , which are chosen to be positive.




Figure 3.7: **Phase conventions in multi-qubit black-box quantization.** Consider a system with two junctions and two cavity modes, for now neglecting the “qubit-like” modes that also exist. The cavity is a stripline resonator with open-open boundary conditions. We choose the phases of the two cavity modes such that in our frame in-phase photons in both modes drive “up” or positively across port one. We also define the direction of the second port such that in-phase photons in mode one are also “up” across port two. When it comes to the interaction of the second mode with port two, we have constraints such that  $\phi_2^2$  must be negative. While in-phase photons in modes one and two are interfering constructively at port one, they are interfering destructively at port two.

### 3.5.3 The qubit-qubit coupling

Predicting the magnitude of the qubit-qubit coupling can be more difficult than many of the other Hamiltonian parameters. It is generally small, but has significant contributions from both junctions. Expanding both cosines to fourth order and taking the rotating-wave approximation yields a coupling term

$$\mathcal{H}_{\text{qubit-qubit}} = [E_{J,1}|\phi_1^1|^2|\phi_1^2|^2 + E_{J,2}|\phi_2^1|^2|\phi_2^2|^2] a_1^\dagger a_1 a_2^\dagger a_2. \quad (3.19)$$

This effect can have either sign, depending on the nature of the direct dipole-dipole interaction and the detunings of the transmons to the cavity.

Following (Chow, 2010, Dalmonte *et al.*, 2015), we know that some component of this comes about due to virtual interactions with the cavity. This should scale as

$$\mathcal{H}_{\text{qubit-qubit,mediated}}/\hbar \propto \chi_1 \chi_2 \left( \frac{1}{\Delta_1} + \frac{1}{\Delta_2} \right). \quad (3.20)$$

This effect is already accounted for in equation 3.19 since in black-box quantization we have already re-diagonalized to remove the second order terms. The  $a_1^\dagger a_c + a_s^\dagger a_c + h.c.$  terms which give rise



Table 3.1: **Black-box vs. experiment comparison.** All frequencies in megahertz. The dispersive shifts were appreciably smaller than expected and the cavity frequencies were higher than expected. One physically reasonable explanation for both of those phenomena is that the sapphire chips were retracted farther into their tunnels than expected. This is also consistent with the bus cavity exhibiting an even larger discrepancy. To some extent this is expected since our simulations do not account for the indium gasket. One other likely issue would be the chips being shifted at an angle. This sample exhibited a particularly large discrepancy between design and experiment, but as the Hamiltonian parameters were not critical for this experiment we may have been more lax in assembly.

	Design simulation	Experiment
$\omega_{T1}/2\pi$	4418	4412
$\omega_{T2}/2\pi$	4618	4585
$\omega_{C1}/2\pi$	5059	5124
$\omega_{C2}/2\pi$	5225	5275
$\omega_{bus}/2\pi$	5600	5692
$\chi_{1,C}/2\pi$	0.958	0.615
$\chi_{2,C}/2\pi$	1.846	1.005
$\zeta$	0.031	0.019
$\chi_{1,bus}/2\pi$	0.421	0.319
$\chi_{2,bus}/2\pi$	1.056	0.445

to the mediated interaction in the traditional derivation are using bare modes, whereas in black-box quantization we have already performed the linear re-diagonalization. In the black-box approach we have already accounted for all second-order terms from the cosine. As with the rest of the spectrum, I calculate this number numerically.

It may be worthwhile to look at some example data in table 3.1, which we take from the gate teleportation experiment (chapter 7). If we apply equation 3.20 to the experimental values of  $\chi_{1,bus}$ , we would expect a mediated contribution of only 240 Hz (the design values would predict an 830 Hz mediated coupling), so we see that the coupling is largely due to their direct interaction.

### 3.5.4 Cavities as circuits

Though we generally use finite-element simulation to find the  $Z(\omega)$  experienced by our qubit or by an output port, the analytical calculation is also a useful exercise. This calculation of how to accurately express a 3D cavity as a lumped-element circuit was one of my first projects in the Schoelkopf lab. This derivation will roughly follow along with (Collin, 2009, Dôme *et al.*, 1987). Starting from Maxwell's

equations, we can express the electric and magnetic field inside a uniform, charge- and current-free, and lossless electromagnetic cavity with the homogeneous Helmholtz equation,

$$\nabla^2 \vec{E} + k^2 \vec{E} = 0$$

$$\nabla^2 \vec{H} + k^2 \vec{H} = 0$$

and electrically-conducting boundary conditions. These admit two classes of solutions, “solenoidal” modes  $\vec{E}_n, \vec{H}_n$  and “irrotational” modes  $\vec{F}_n, \vec{G}_n$ , such that

$$\nabla \cdot \vec{E}_n = 0$$

$$\nabla \times \vec{F}_n = 0$$

$$\nabla \cdot \vec{H}_n = 0$$

$$\nabla \times \vec{G}_n = 0.$$

These eigenmodes form a complete basis for any  $\vec{E}, \vec{H}$  inside of the cavity under our constraints. Also note that these modes are orthogonal, and we impose normality<sup>10</sup> such that

$$\int \vec{E}_n \cdot \vec{E}_m dV = \delta_{nm} = \int \vec{F}_n \cdot \vec{F}_m dV$$

$$\int \vec{E}_n \cdot \vec{F}_m dV = 0.$$

Analogous relationships hold for the magnetic fields. With this established, we can expand a field inside the cavity in eigenmodes,

$$\vec{E} = \sum_n e_n \vec{E}_n + f_n \vec{F}_n \tag{3.21}$$

$$\vec{H} = \sum_n h_n \vec{H}_n + g_n \vec{G}_n, \tag{3.22}$$

<sup>10</sup>For convenience  $\vec{E}_n$ , etc, hold units of [meters]<sup>-3/2</sup>.

where our sum over  $n$  should be understood to be over the requisite number of indices. If we now return to Maxwell's equations,

$$\begin{aligned}\nabla \times \vec{E} &= -j\omega\mu_0\vec{H} - \vec{J}_m \\ \nabla \times \vec{H} &= eq\epsilon_0\vec{E} + \vec{J}_e,\end{aligned}$$

we substitute our expansion and dot with a mode from  $\{E_n, F_n, H_n, G_n\}$ . After integration, orthogonality removes modes other than  $n$ , and we can solve the system of equations to obtain

$$\begin{aligned}e_n &= \frac{-jkZ_0 \int \vec{E}_n \cdot \vec{J}_e dV - k_n \int \vec{H}_n \cdot \vec{J}_m dV}{k_n^2 - k^2} \\ h_n &= \frac{\int \left( k_n \vec{E}_n \cdot \vec{J}_e - jkY_0 \vec{H}_n \cdot \vec{J}_m \right) dV}{k_n^2 - k^2} \\ f_n &= \frac{-\int \vec{F}_n \cdot \vec{J}_e dV}{j\omega\epsilon_0} \\ g_n &= \frac{-Z_0 \int \vec{G}_n \cdot \vec{H}_{tan} dS - \int \vec{G}_n \cdot \vec{J}_m dV}{j\omega\epsilon_0}.\end{aligned}$$

When we consider fields excited only by an electric current source<sup>11</sup>, we can create the variable  $\phi_n \equiv \int \vec{E}_n \cdot \vec{J}_e dV$  and see that

$$e_n = \frac{-jkZ_0\phi_n}{k_n^2 - k^2} = \frac{j\omega\mu_0\phi_n}{k^2 - k_n^2} \quad (3.23)$$

$$h_n = \frac{k_n\phi_n}{k_n^2 - k^2} \quad (3.24)$$

$$f_n = \frac{-\int \vec{F}_n \cdot \vec{J}_e dV}{j\omega\epsilon_0} \quad (3.25)$$

$$g_n = 0. \quad (3.26)$$

We see immediately that modes  $E_n$  and  $H_n$  are unsurprisingly excited in pairs, and that modes  $F_n$  correspond to zero-frequency, or capacitive, modes. To convert this information to an impedance we

<sup>11</sup>For evanescent coupling it is not clear this is correct. We are presently treating the case of a protruding capacitive probe.

use the definition of  $Z$  as defined by complex power (Collin, 2009, Ramo *et al.*, 1965),

$$Z_{\text{in}} = \frac{P_{\text{loss}} + 2j\omega(W_m - W_e)}{\frac{1}{2}|I_{\text{in}}|^2}, \quad (3.27)$$

where  $W_m$  and  $W_e$  represent the energy stored in the inductive and capacitive elements respectively and  $P_{\text{loss}}$  reflects the power lost to resistive elements. Exploiting the orthonormality of modes, our time-averaged electric and magnetic field energies are given as

$$W_e = \frac{\epsilon_0}{4} \int \vec{E}^* \cdot \vec{E} dV = \frac{\epsilon_0}{4} \sum_n (|e_n|^2 + |f_n|^2)$$

$$W_m = \frac{\mu_0}{4} \sum_n |h_n|^2.$$

With  $P_{\text{loss}} = 0$ , we find

$$|I_{\text{in}}|^2 Z_{\text{in}} = j \frac{1}{\epsilon_0} \sum_n \frac{\omega \phi_n^2}{\omega_n^2 - \omega^2} - j\omega\epsilon_0 \sum_n |f_n|^2.$$

We define some helpful terms

$$\Phi_n^2 = \frac{\phi_n^2}{|I_{\text{in}}|^2}$$

$$L_n = \frac{\Phi_n^2}{\omega_n^2 \epsilon_0} \quad (3.28)$$

$$C_n = \frac{\epsilon_0}{\Phi_n^2} \quad (3.29)$$

$$C_0^{-1} = \frac{\sum_n \left( \int \vec{F}_n \cdot \vec{J}_e dV \right)^2}{\epsilon_0 |I_{\text{in}}|^2}, \quad (3.30)$$

where  $\Phi_n$  is ultimately a geometric term describing the interaction of the driving current distribution with the mode  $n$ . We see that  $L_n$  and  $C_n$  are a straightforward calculation given the unperturbed resonator modes and an assumption about the current distribution on the driving probe.  $C_0$  is a sum over the zero-frequency modes, and represents the capacitance of the conductor housing the driving current. The infinite sum is unlikely to have nice properties, but approximations should be available for

driving electrodes with simple geometry. Rearranging, we find

$$\text{Im}[Z_{\text{in}}] = \sum_n \frac{1}{\frac{1}{\omega L_n} - \omega C_n} - \frac{1}{\omega C_0}, \quad (3.31)$$

which exactly corresponds to a circuit in Foster's first form. As an example, we can look at the impedance seen from the input port of a protruding pin. For a cylindrical probe of length  $l$  and radius  $r \ll l$ , Schelkunoff estimates this capacitance (Schelkunoff and Fris, 1953) as

$$C_0 = \frac{2\pi\epsilon_0 l}{\left(\ln\left(\frac{l}{r}\right) - 1\right)} + \frac{2\pi\epsilon_0 l \ln(2)}{\left(\ln\left(\frac{l}{r}\right) - 1\right)^2}, \quad (3.32)$$

and we assume the current distribution to be sinusoidal with a current node at the tip. We arrange this probe to be protruding into a rectangular cavity with dimensions  $a, b, d$  and corresponding mode indices  $n = (m_1, m_2, m_3)$ , where the probe is along the second dimension and centered in the other dimensions. We calculate that

$$\Phi_n^2 = \frac{2}{abd} \frac{b^4 k_w^2}{(b^2 k_w^2 - m_2^2 \pi^2)^2} \left( \frac{\cos\left(\frac{lm_2\pi}{b}\right) - \cos(k_w l)}{\sin(k_w l)} \right)^2 \sin^2\left(\frac{m_1\pi x}{a}\right) \sin^2\left(\frac{m_3\pi z}{d}\right). \quad (3.33)$$

We compare this directly a simulation of  $Z_{\text{in}}$  from an HFSS simulation, plotted in figure 3.8, and find good agreement in the values of  $C_n$  and  $L_n$ . Unsurprisingly our estimate for  $C_0$  is close but not exact, and we additionally miss a native (apparent) inductance  $L_0$ . This inductance is largely due to not considering the geometry inductance of the probe.

### 3.6 The Purcell effect

We have so far only concerned ourselves with the coherent behavior of the system. Fundamentally, in the last section we have been discussing the mixing of modes and finding new eigenmodes of the system. This section aims to answer a related question: when the original modes have some understood interaction with the outside world, how will the new modes see the outside world? This effect was the interest of Edward Mills Purcell in the early days of nuclear magnetic resonance experiments in the 1940s

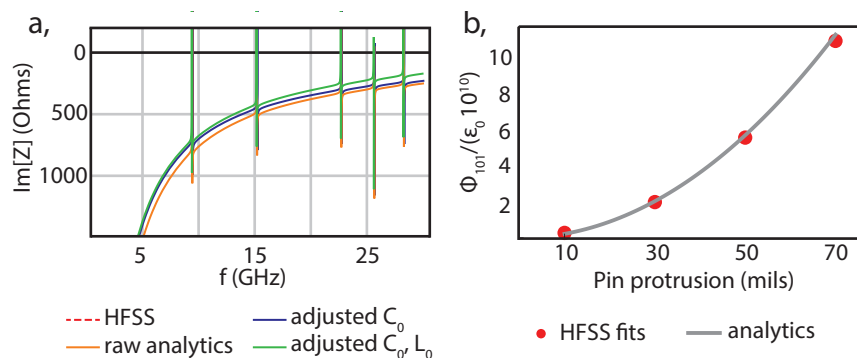


Figure 3.8: **Impedance of a rectangular resonator as seen from an external capacitive probe.** In figure **a** we see  $Z_{\text{in}}(\omega)$ , comparing a simulation in HFSS (red squares) to an analytical prediction (orange line). If we adjust our estimate of the capacitance slightly we get the blue line, and if we further allow a linear inductance we get the green line, finding quite good agreement. In figure **b** we see the value of  $\Phi$  for the 101 mode of the cavity plotted against the pin length, and see that the analytic calculations agree quite well. In these data the numerical values of  $\Phi_{101}$  were extracted from fits.

(Purcell, 1945). He was primarily concerned with the relaxation rate of nuclear magnetic moments. If one polarized nuclear spins in a sample, how long would they take to thermalize with the environment?

His experiment was to borrow a three-dimensional metal cavity, similar to the type we have become newly popular in our field, and a radiofrequency generator from the MIT radiation laboratory. He filled the cavity with paraffin and placed it under a DC magnetic field. He attached the generator to the cavity using a circuit sensitive to changes in the signal intensity and slowly began varying the magnetic field. At a specific magnetic field intensity (and therefore nuclear precession frequency), the signal intensity spiked, six years later he shared the Nobel Prize for this and his work in NMR<sup>12</sup>.

This effect arises because while the resonator narrows the density of states available to any potential photon emitted from the atom (will it have a place to go?), it greatly enhances the strength of the interaction. In an illustrative if not particularly accurate description, the photon will “bounce around” in the resonator, increasing the amount of interaction time it has with the atom. Following from Fermi’s golden rule, at resonance this will enhance the decay of the atomic state into a low-Q resonator. We will discuss the physics in this section and discuss designing devices with consideration of this effect later in chapter 4.

<sup>12</sup>Because this experiment involved many, many spins ( $\approx 10^{23}$ ), this is essentially the classical, asymptotic version of the Purcell effect. Modernly we often reserve the term “Purcell effect” for this behavior exhibited in the few-quanta regime.

### 3.6.1 Jaynes-Cummings

How do we predict this behavior? We will start with the case of Jaynes-Cummings model, where we can solve for the loss rate exactly. The most common scenario regarding the Purcell effect is to consider a lossless qubit and a lossy cavity, then look for the loss rate inherited by the hybridized qubit mode. We can solve this system directly from  $\mathcal{H}_{JC}$  (equation 3.1). For a ladder of  $n$  originally corresponding to the cavity photon number, the new eigenstates were given in equations 3.3 and 3.4. For an unexcited cavity we can constrain ourselves to  $n = 0$ , and see that the “qubit” modes mix as  $\cos$  and  $\sin \left[ g / \sqrt{g^2 + \Delta^2} \right]^2$ .

In the precisely resonant case, the modes mix completely and the qubit will decay at half the cavity decay rate, or  $\kappa/2$ . Here the loss rate mixed exactly as the probability rate, as adding the undiagonalized Lindbladian to the master equation does not affect the nature of the normal modes. We much more commonly operate in the dispersive regime where  $\Delta \gg g$ , or  $\theta_0 \ll \pi/2$ . Accordingly, we approximate that the participation probability of the diagonalized “qubit” mode in the cavity mode is  $(g/\Delta)^2$ , and the inherited decay rate is then  $(g/\Delta)^2 \kappa$ .

### 3.6.2 Classical treatments

Unfortunately, one can see quickly in a circuit QED experiment that this formulation does not hold in our experiments—notably that the effect is never symmetric in  $\Delta$ . It is sometimes stated in our experiments that the Purcell asymmetry in cQED is due to the existence of higher modes of the lossy cavity (Houck *et al.*, 2008), the so-called “multimode-Purcell effect.” That statement is true, but the asymmetry also arises for a simpler reason, and in my opinion the result of that paper is simply that one needs to take the full impedance environment into account. We can explore this by looking at a simple circuit—a pair of capacitively, weakly coupled LC resonators, one of which has a parallel resistor, as in figure 3.9. I find this toy circuit quite illustrative. We can examine the lifetime of new mode (“mode one”) which is principally the naturally-lossless resonance.

There are several approaches to finding these lifetimes, most neatly by finding the poles of the response function. The complex-valued poles  $\{s_i \equiv i\omega_i + \alpha_i\}$  in  $Z(a)$  will yield both the frequency and lifetime of the resonance, where  $\omega_i$  is the resonance frequency and  $\alpha_i$  is the decay rate (Ramo *et al.*, 1965). This can be done analytically, though it is somewhat tedious, or numerically via techniques like




Figure 3.9: **Mixing an LC oscillator and an RLC oscillator.** The naturally lossless “mode one” capacitively couples to a lossy “mode two.” It can be useful to examine the loss rates of the new eigenmodes.

vector fitting (Gustavsen and Semlyen, 1999).

We can also approximate that the effect of the loss is small, and that the resonance frequencies are not shifted significantly. In this case the imaginary pole of mode one does not move, or alternately put that the reactive contribution to  $Z$  at the original resonance frequency of mode one is small compared to the modes natural divergence. The resistive contribution will also be small, but is obviously significant compared to the zero resistance of the original mode. We can then consider the rest of the circuit (coupling capacitor and lossy mode two) in parallel with mode one to be only a resistor, evaluated at the frequency of interest. This reduces mode one to a parallel RLC oscillator. Solving for this effective resistance yields a lifetime for mode one,  $C_1 / \text{Re}[Y(\omega_1)]$ , where  $Y$  is the admittance across the elements of mode one.

The key result of this calculation is that the lifetime is *not* symmetric in detuning between the two original modes, and only becomes so as the coupling capacitance tends towards infinity<sup>13</sup>. The frequency-dependence of the capacitive coupling will already yield stronger coupling, and therefore relatively lower lifetimes for mode one, as the frequency of mode one increases over that of the lossy mode two. Accordingly, even when higher resonator modes are pushed away, it is preferable to keep high-Q modes spectrally below low-Q modes.

More generally, mode-mixing is rarely well-described by the simplest answers. (Nigg *et al.*, 2012) proposes to calculate the Purcell effect by using a modified Foster’s first form, often called “Lossy Foster,” turning the LC circuits of the cavity modes into RLC circuits and performing an analysis as above. This approximation is discussed briefly in (Ramo *et al.*, 1965), pages 559-561. Unfortunately this

<sup>13</sup>The verbose calculation is not particularly useful, but its clear from  $Y$  that there is no symmetry.



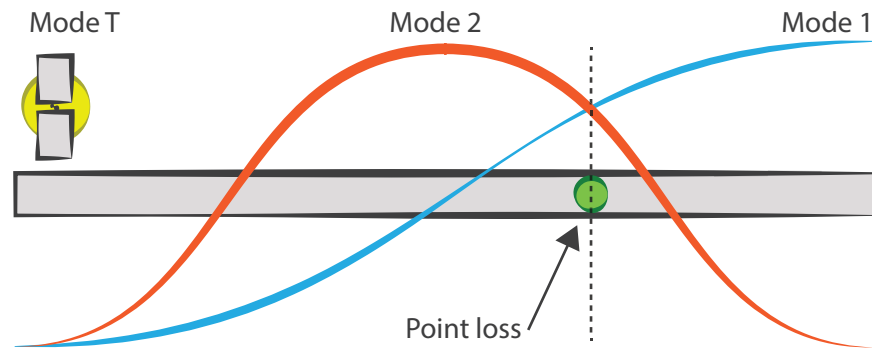


Figure 3.10: **Mode mixing gedanken experiment.** For a spatially localized loss-mechanism, the nature of the mode hybridization can have non-trivial effects on the inherited loss. A less pathological example of a localized loss is an SMA coupler, or the edge of a coplanar waveguide resonator with weak capacitive coupling to a transmission line as in (Houck *et al.*, 2008).

has been seen to produce misleading estimations at times. In particular, it fails when loss is inherited in a significant amount by *multiple* modes, as the interference of the different modes can be significant when the loss has spatial dependence. For an illustrative example, consider the cartoon in figure 3.10,

This system is a stripline resonator with open-open boundary conditions, so we see the  $\lambda/2$  and  $\lambda$  modes (labeled one and two), along with a highly localized qubit mode (mode T). These modes have not yet been re-diagonalized due to the qubit perturbation. In this example, all of the loss comes from a similarly localized point on the resonator. For illustrative purposes that lossy point was chosen to be at where modes one and two have equal field amplitudes. The qubit mode will perturb the cavity modes, and the new eigenmodes will be some combination<sup>14</sup>. For a weakly-coupled system, one of these modes will be predominantly the original qubit mode. The dipole interaction will generate similar but not equal resonant coupling strengths  $g_i$  between the qubit mode and modes one and two. There will exist a qubit mode frequency for which the new participation of the qubit mode will be equal in the original modes one and two. How will the qubit see the point loss? It depends on whether or not the hybridization with modes one and two are in-phase, or out-of-phase. If the combination is out-of-phase, the qubit may not see the point loss at all. This example is cartoonish, but the overall concern is not, as loss due to output couplers is highly spatially localized. As such, the lossy Foster approximation should be used with care.

<sup>14</sup>And with other modes, but this does not qualitatively change the present discussion.

In practice, we do not use an elegant method to obtain precise Purcell estimates, we ask a finite element solver to find the complex eigenfrequencies of the system. We will discuss the practical matters of these simulations later in section 4.1.2.

### 3.6.3 Towards a quantum treatment

Though it has not yet been utilized in the Schoelkopf or Devoret labs, it is worth mentioning that Solgun, Abraham and DiVincenzo have published a method (Solgun *et al.*, 2014) which builds upon the work in (Nigg *et al.*, 2012). In the high-Q limit both approaches should predict the same Hamiltonians and coherent behavior, but the Solgun paper provides a more accurate treatment of the Purcell effect, remedying the flaws of the “lossy Foster” approach which we have discussed. A brief overview is worthwhile.

As in the previous section, they perform a 3-port simulation, then numerically terminate the ports corresponding to the output and input transmission lines. This results results in the impedance seen by the junction without having to simulate loss. They fit this impedance to the general form of poles in a positive real function, then use the technique of Brune synthesis (Ramo *et al.*, 1965) to generate the corresponding circuit. While any lossless (passive) impedance can be written in Fosters forms, any passive one-port impedance can be expressed as a Brune circuit. Solgun goes on to quantize this circuit (which could be done as in (Devoret, 1995)). He also performs a Markov approximation and obtains prescription for the decay rate inherited from each resistor in the Brune circuit. The generality of this approach make it useful for systems other than transmons, but at present we do not know an advantage of this for transmon systems. It may yield non-trivial predictions for the level-dependence of the decay rate.

## 3.7 Control & measurement

One obvious solution to this quandary is to simply make the cavities as long-lived as possible. This approach has it own rather fundamental flaw which is the central conflict of quantum computer design: decoherence results from coupling to the outside world, but we must couple to the outside world to

control and gain information from the computer. Balancing these concerns requires highly careful machinations, many of the practical issues of which we will discuss in chapter 4. Before we get to those prescriptions we will discuss the other side of that balance, namely the ability to manipulate the Hamiltonian of system in real time and the ability to get classical information back out.

### 3.7.1 Introducing the drive

We will first introduce the drive first classically, considering the impedance as seen by the input port. As in figure 3.8, the cavity and qubit modes will look qualitatively the same. Before looking explicitly at the drive, we characterize the coupling by looking at the output Q, placing a resistor across the port of a two-mode circuit in Fosters first form. Near each mode we can roughly ignore the other<sup>15</sup>. The mode then sees a  $\text{Re}[Y(\omega_i)]$ , leading to “Purcell” Qs

$$Q_{\text{Purcell},i} = \omega_i C_i / \text{Re}[Y(\omega_i)] = \omega_i R C_i \left( \frac{1}{(\omega_i C_c R)^2} + 1 \right). \quad (3.34)$$

Conversely we now ask how much energy we place in each mode when we apply an oscillating voltage at the drive port. We perform an elementary circuit analysis driving a power across the resistor of  $P = V^2/R$  and calculate the voltage (and energy) in the capacitor of the mode of interest. We obtain the ratio

$$Q_{\text{drive},i} \equiv \omega_i U_{c,i} / P = (1/2) \omega_i R C_i \left( \frac{1}{(\omega_i C_c R)^2} + 1 \right), \quad (3.35)$$

where the factor of two is due to only considering the capacitive energy. This shows shows the direct symmetry between the Purcell relaxation rate (due to the input port, using the techniques from section 3.6) and the ability to resonantly drive energy into the relevant modes.

To consider the drive rigorously, we need an analysis using Input-Output theory (Gardiner and Collett, 1985, Girvin, 2011). We will not reproduce these derivations here, but will provide a qualitative overview. The transmission line is taken to be a sum of harmonic oscillators weakly coupled to the cavity mode via excitation-swapping interactions. In the “stiff pump” approximation (Kamal *et al.*, 2009), we take this coupling to be weak and the magnitude of the coherent state applied on the input modes to be

<sup>15</sup>The  $1/\Delta^2$  contributions to  $\text{Im} Z$  will be dominated by the resonant mode.

large. In this regime the loss or addition of a single photon to the input modes has a negligible effect, rendering it essentially classical.

We also need to examine the filtering effect of the cavity. For a single-port, linear cavity, the cavity displacement field is related to the input field by

$$a[\omega] = -\frac{\sqrt{\kappa}}{i(\omega_c - \omega) + \kappa/2} b_{in}[\omega]. \quad (3.36)$$

This implies that resonant drives on the cavity ring-up and -down at a rate limited by  $\kappa$  and that far off-resonant drives on the cavity respond at rate  $\Delta$ , screened in amplitude by  $1/\Delta$ . Most traditional applications are the resonant or near-resonant (readout) regimes where the time dynamics are simple, or the far off-resonant (qubit driving) regime where the time-dynamics are very fast. Currently popular techniques are blurring these delineations, using moderately off-resonant drives and in applications such that we do care about the time dynamics. We will see examples of this in the beam-splitter interactions used for Q-switching in chapter 6 and the resonator-induced phase gate in section 7.6.

Ignoring these complications, most discussions of the drive in circuit QED assumed the envelope of the drive to be slowly-varying relevant to the other timescales. This results in a drive Hamiltonian

$$\mathcal{H}_D = \epsilon e^{i\Delta t} a^\dagger + \epsilon^* e^{-i\Delta t} a,$$

where  $a$  and  $a^\dagger$  are the annihilation and creation operators of the driven cavity mode and the strength  $\epsilon$  describes an electric field. In this form we are placing the drive in the rotating frame of the cavity, and  $\Delta$  is the appropriate detuning. With knowledge of the coupling Qs we might more appropriately also allow the drive to have a direct drive on the qubit with a term including  $a_q$ . This is likely to become a more important consideration as the Yale group moves towards experiments using more exotic driving.

In the absence of a Josephson junction, the effects of the drive on the cavity are straightforward—only displacements given entirely by equation 3.36. We will move on to a discussion of the drive in the presence of a junction, and we will restrict ourselves as before to the drive as a displacement of our cavity, not the qubit. It is not hard to generalize the analysis if required. For tidiness we will use the full cosine, though keep in mind that in practice we must remove the second order terms of the undriven

junction phase as in 3.14.

Working in the rotating frames of the qubit and cavity, we get

$$\mathcal{H} \sim E_J \cos\left(\phi_q \left(a_q + a_q^\dagger\right) + \phi_c \left(a + a^\dagger\right)\right) + \epsilon e^{i\Delta t} a^\dagger + \epsilon^* e^{-i\Delta t} a. \quad (3.37)$$

As the drive term represents a displacement, we can work in that displaced frame as discussed in the supplemental materials of (Leghtas *et al.*, 2015). If we additionally assume a constant or slowly varying envelope, this is equivalent to making the substitution that  $a \mapsto a + \xi$ , in which case we see

$$\mathcal{H} \sim E_J \cos\left(\phi_q \left(a_q + a_q^\dagger\right) + \phi_c \left(a + a^\dagger + \xi e^{i\Delta t} + \xi^* e^{-i\Delta t}\right)\right), \quad (3.38)$$

and  $\xi \sim -i\epsilon / (\kappa/2 + i\Delta)$ .

### 3.7.2 Pumps & several-photon processes

It is easy to see from this that control in the power and detuning of the drive we could also apply linear drive terms that are resonant with the qubit, allowing Rabi flopping. We may also generate new terms which we have previously neglected under the rotating wave approximation. Restricting ourselves to the fourth order terms of the cosine expansion, we see many terms like e.g.  $a^\dagger a_q a a$  which are not energy conserving and will rotate quickly in the lab frame. We symmetrically obtain terms like  $a^\dagger a_q \xi^* \xi$ , which are also not generally resonant. However if we use two drives, and the tones have frequencies differing by  $\Delta_p$ , we see also the term  $a^\dagger a \xi_1^* \xi_2 \exp\{i\Delta_p\}$ . This becomes resonant if the detuning between the drives is equal to the detuning between the qubit and cavity<sup>16</sup>. This term generates a direct swapping of excitations between two modes, often a qubit and a cavity. In the key experiments of this thesis, we use this term  $H_{swap} = E_J |\phi_q|^2 |\phi_c|^2 \xi_1 \xi_2^* a a_q^\dagger + h.c.$  to reset the qubit state, as the readout cavity is a (mostly) cold and highly lossy mode. Similar analysis will be used again later in section 7.6 dealing with the drive schemes we used for two-qubit gates.

Exploring and exploiting these term is an ongoing topic of study in the Schoelkopf lab, and the list of available transitions gets even larger when we move beyond the first order RWA. As mentioned before,

<sup>16</sup>The same drives will also create AC Stark shifts, so the proper detuning is slightly modified.

this calculation is not precise for a time-varying drive, where a rigorous calculation may use Floquet analysis (Shirley, 1965).

### 3.7.3 Dispersive measurement

We also exploit the dispersive interaction to read out the states of the transmons. Reduced to the relevant terms, the Hamiltonian is given by

$$\mathcal{H}_{\text{readout}}/\hbar = \left( \omega_c - \chi a_q^\dagger a_q \right) a^\dagger a + \epsilon^* a + \epsilon a^\dagger, \quad (3.39)$$

where  $\epsilon$  is on or near resonance with the cavity. The resource we exploit for measurement is that the effective resonance frequency of the cavity is dependent on the qubit state. When  $\chi \gg \kappa$ , the two peaks are distinct and resolvable. One illustrative (if simplistic) readout mechanism would simply be to try to transmit a weak cavity tone at the bare frequency  $\omega_c$ . A signal would transit through the cavity if and only if the qubit was in the ground state, which we could then detect at the output. This captures the basic phenomena at play, but neglects important nuances and certainly is not an optimal approach. Indeed, from equation 3.36, we see that a slow  $\kappa$  means it will take a long time to get photons into the cavity, perhaps allowing the transmon time to thermalize. Accordingly we generally do not work in the  $\chi \gg \kappa$  regime.

Instead let us look more generally at the problem. When the qubit is projected into one state, the Hamiltonian is only that of a harmonic oscillator. Then when we apply a classical drive we expect the cavity to remain in a coherent state  $\alpha_{g/e}(t)$ . Since the cavity resonance frequency is different for the two qubit states, these trajectories will generally not be the same, and the cavity will propagate forward into a superposition, entangled with the qubit state. It is easily seen that the equations of motion for these two driven, damped harmonic oscillator trajectories are given by

$$\dot{\alpha}_e(t) = -i\xi(t) - (\kappa/2 + i(\Delta - \chi/2))\alpha_e(t) \quad (3.40)$$

$$\dot{\alpha}_g(t) = -i\xi(t) - (\kappa/2 + i(\Delta + \chi/2))\alpha_g(t), \quad (3.41)$$

where  $\Delta$  represents the detuning of the drive relative to the symmetric point,  $\omega_c - \chi/2$ . Note that  $\kappa$  is

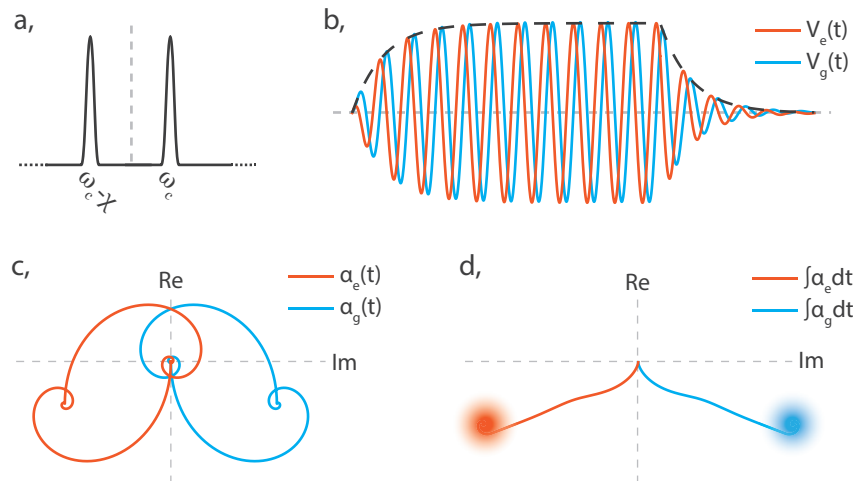


Figure 3.11: **Dispersive readout.** In **a**, we see the qubit-state dependent peaks of the cavity, and choose the frequency of our weak readout tone (dashed line) to be directly between the two. We are clearly in the strong dispersive regime as the peaks do not overlap, so  $\chi \gg \kappa$ , but this is often not the case in practice. In **b**, we see the detected output. Due to our choice of driving frequency both the  $|e\rangle$  and  $|g\rangle$  associated peaks respond with equal amplitude, but with distinct phases. Note the filtering effect of the cavity causing it to ring up and down at a timescale given by  $(\kappa^2 + \chi^2/4)^{-1/2}$ . In **c**, we have plotted the response in phase space. Note that the trajectories ring up with different phases, maintain a steady state, then ring back down to zero. Here  $\chi = 2\kappa$  and the drive is a sharp rectangular pulse. In **d**, we have plotted the unweighted integral of this signal, and we can see that the pointers separate in time. Gaussian noise in the final value has been overlaid for illustrative purposes. In practice we integrate with a weighted envelope which drives these values directly along the real axis.

the energy decay rate, so  $\alpha \propto \sqrt{n}$  decays at  $\kappa/2$ . We will also assume that all photons that decay do so by leaking out into a transmission line, rather than being absorbed by the cavity walls.

These trajectories carry our information about the qubit state, and we generally assign a discrete qubit state to each trajectory. This problem can be reduced to the machine-learning problem of classification: given these large trajectory vectors, to which class do they belong? We use supervised learning, integrating each trajectory with a matched filter equal to the difference of the mean prepared “e” and “g” trajectories. We then apply a threshold derived from the prepared cases. For non-linear readout chains, as well as in the presence of qubit  $T_1$  or a nonlinear readout cavity, we may get better discrimination by turning to more advanced techniques (Magesan *et al.*, 2015).

As the pointer states become orthogonal, tracing over them begins to mix the qubit state. If our measurement chain were perfect, and at room temperature we could perform an ideal discrimination,

our certainty of the qubit state would converge at the same rate as the pointer overlap decreases. Accordingly we consider the separation of the pointer states in phase-space,  $\beta(t) = \alpha_e - \alpha_g$ . Following the reasoning in (Gambetta *et al.*, 2008)<sup>17</sup> and to a lesser extent (Clerk *et al.*, 2010, Girvin, 2011), under the assumption that qubit decay (thus far neglected) is much slower than the cavity decay rate, we obtain a rate of information gain

$$\Gamma_{\text{ci}} = \frac{\eta}{2} \kappa |\beta(t)|^2.$$

where  $\eta$  represents the measurement efficiency, and the factor of  $1/2$  is a result of using a Josephson Parametric Converter (Bergeal *et al.*, 2010) to do heterodyne measurement.  $\eta$  represents directly what fraction of information about the qubit state obtained by the cavity state that makes its way to the experimentalist. The photons absorbed (e.g. by the cables and circulator) before they reach the JPC represent information gained by the environment and not by the experimentalist. Measurements are often performed by pulses with an average photon number on the order of one, so there is not a great deal of redundancy in the signal. In the strong measurement limit this can be slightly less important—it is unimportant if the environment learns the qubit state, since we are projecting into an energy eigenstate and there is no phase coherence to be lost. This is with the caveat that the experimentalist does, in fact, gain enough information to overcome to noise sufficiently to accurately identify the qubit state. Clearly photon loss will more classically decrease our signal to noise ratio, so we hope to have a reasonable  $\eta$  regardless. In the Schoelkopf and Devoret labs we can reliably obtain  $\eta$ s between roughly 0.25 and 0.5 with our current technologies and practices.<sup>18</sup>

We now almost have the tools to explain how a measurement will work, but we still need to know when we have enough information to have certainty about the qubit state. This lets us know, given a cavity  $\kappa$ , dispersive shift  $\chi$  and an output chain efficiency  $\eta$ , how many photons we need to measure with and for how long we need to measure to discriminate the qubit state. Following the calculations

<sup>17</sup>We have modified their definition of  $\eta$  by a factor of  $1/2$ , to be more consistent with the modern usage in (Hatridge *et al.*, 2012).

<sup>18</sup>Why not simply readout with more photons? There are three significant reasons. First, a JPC has a finite dynamic range. After a certain amount of power it can no longer provide gain. Second, when using more photons higher order terms in the Hamiltonian, such as the cavity anharmonicity, start to come in to play. These terms break cavity reset schemes, and can also have the unintended effect of measuring other modes of the system if they exist. Third, we have observed that the energy decay rate of the qubit can increase when there are more photons in the cavity. See (Mundhada *et al.*, 2015) for details. These concerns aside, it can be quite effective to use very large photon numbers for non-QND readout (Boissonneault *et al.*, 2010, Reed *et al.*, 2010a)



in either (Gambetta *et al.*, 2008, Hatridge *et al.*, 2012), the result is that the qubit state converges to  $\pm 1$  with a certainty given by

$$z(t) = \tanh \left[ \frac{\eta\kappa}{2} \int_{-\infty}^t dt' |\beta(t')|^2 \right]. \quad (3.42)$$

The derivation results from the assumption that the noise is Gaussian, then looking at how the integrals of the signals will separate in the presence of that noise. As mentioned before, this is not strictly true in the presence of qubit  $T_1$ , with a non-linear readout chain, or with a nonlinear cavity, but is illustrative nonetheless.

### 3.8 Conclusions

In this chapter we have presented the basic physics underlying the circuit QED and the experiments of this thesis. We have set up the basic mechanisms by which we can use the physical devices we will discuss in later chapters to realize the theoretical constructs of chapter 2. We have seen now how the two lowest-lying energy states of a transmon can play the role of a qubit, and how we can drive those transitions with oscillating magnetic fields. We have also seen that in circuit QED we very naturally have an additional resource in the form of harmonic oscillator modes. These modes can be used straightforwardly to suppress the electromagnetic radiation that might otherwise come out of a dipolar oscillating qubit. We can also use the interactions of these resonators to measure the transmon state.

These traditional uses, however, represent a past era of circuit QED experiments in the Schoelkopf lab. In recent years we have found that the harmonic oscillators have much better coherence properties, and in turn we have shifted to using resonators as the primary carriers of quantum information. In this frame the more subtle implications of black-box quantization and how our pumps and drives work become much more important. It seems likely that the works of this thesis represent an intermediate step in that transition. I suspect that in the near future transmons may be removed from the picture entirely, replaced by some other nonlinear element better tailored to the manipulation of cavity states. In any case much of the basic physics presented here is likely to remain useful.

## Chapter 4

# Experimental design and methods

Quantum phenomena do not occur in a Hilbert space. They occur in a laboratory.

---

Asher Peres, 1995

In this chapter we will cover many of the practical aspects of executing experiments like those found in this thesis. In section 4.1 we will attempt to provide some guiding principles for wading through the parameter space available in experiment design. This includes choosing the junction and charging energies in section 4.1.1. Beyond choosing a frequency which gives the desired couplings, there are trade-offs between sensitivity to charge dispersion, speed of control, and ease of fabrication. Another set of considerations regard the coupling of the system to the outside world. In section 4.1.2 we will look at how the chosen strength of coupling to the outside world affects decoherence through two different mechanisms.

In section 4.2 we will examine several specific devices used in my graduate career: the flux-tunable two-qubit vertical transmon, the four-transmon Shiva device, and the four-high-Q-resonator Dodekathion device. In each section we will get into some of the specific design choices made, as well as the reasoning that lead to them. These include high-level “physics” decisions that control the experimental Hamiltonian as well as practical decisions such as ease of fabrication or machining.

Section 4.3 will dive into the microwave environment used in these experiments. Starting from the

generator, our time-domain control is largely orchestrated using single-sideband mixing, discussed in section 4.3.1. In current experiments the shapes of those pulses are often determined mathematically, which we briefly discuss in 4.3.2. Our control signals (for now putting off the complexities of measurement chain) pass through various amplifiers, filters and attenuators on their way to the sample at the base temperature of the dilution refrigerator, described in section 4.3.3. We discuss the microwave environment on the way back to room temperature in section 4.3.4. I additionally provide figures 4.8 and 4.9 which depict with relatively high detail the wiring arrangements of the two main experiments of this thesis.

The measurement chain is treated more specifically in section 4.4. We cover the essential practical ideas of heterodyne measurement in section 4.4.1, before looking at two specific arrangements. Section 4.4.2 covers the heterodyne measurement scheme that was standard in the Schoelkopf lab from my arrival until roughly four years into my stay there. This scheme was used in all of my early experiments, as well as in the stabilizer measurement experiment (chapter 6). In section 4.4.3 we cover the more modern heterodyne measurements made possible by the new FPGA hardware. These scheme was used in the gate teleportation experiment 7.

Finally, in section 4.5 we conclude with a brief exposition on the Ansys HFSS finite-element solver used in our electromagnetic simulations. This section is not meant to be an introduction or an exhaustive guide. Rather, it should supplement a basic understanding of HFSS with some use-specific practices and techniques which I have settled on during my years of use.

## 4.1 Device design

The superconducting transmon qubit and the harmonic microwave resonator are fundamental building blocks of the physics we examine in the Schoelkopf lab. In this section we will discuss some of the device parameters that we must decide upon to build such experiments. We have previously discussed two transmon decoherence mechanisms, charge dispersion in section 3.4 and the Purcell effect in section 3.6. We will expand on these, adding photon shot-noise dephasing, and discuss how the relevant device design decisions are in tension with other desiderata.

### 4.1.1 Charge dispersion: bounds on $\alpha/\omega$ .

Our previous theoretical discussion of charge dispersion in section 3.4 explained the transmon's *raison d'être*. As  $E_J/E_c \gg 1$  increases, the sensitivity of the qubit transition to charge fluctuations decreases exponentially. When estimating  $T_2$  due to charge noise we generally do not assume a spectrum for this noise but instead look at the worst-case scenario or the maximum splitting. Following the chart in figure 3.3 we see that we reach “phase-Qs” of tens of millions (outside of our range of concern given our current  $T_1$ s) once  $E_J/E_c$  reaches roughly 100. This corresponds to a relative anharmonicity of roughly  $\alpha/\omega \approx -0.05$ . Since our qubits tend to have  $\omega/2\pi \approx 5\text{GHz}$ , this means anharmonicities of less than 250MHz.

This trade-off motivates several considerations, principally how quickly we may address the qubit transition. We typically shape our control pulses with Gaussian envelopes given by  $f(t) \equiv \exp(-(t-t_0)^2/2\sigma_t^2)$  which we typically truncate at  $\pm 2\sigma_t$ . From the Fourier transform we see that the spectrum of such a (untruncated) pulse is also a Gaussian with width given by  $2\pi\sigma_f = 1/\sigma_t$ . This sets a direct limit in how fast qubit operations may be performed, as we do not wish to drive the transmon's  $|e\rangle$  to  $|f\rangle$  transition which would result in leakage out of the qubit space. Located  $\alpha$  away from the qubit transition, we require that  $\alpha$  is several times larger than  $\sigma_f$ . There are techniques to minimize this effect, including DRAG Motzoi *et al.* (2009) and optimal control (section 4.3.2), but the speed limit is roughly fixed.

The desire for faster pulses exerts a force towards desiring larger  $E_c$ . To maintain  $E_J/E_c$  we might increase  $E_J$  correspondingly, but our qubit frequency will increase as  $\sqrt{8E_JE_c}$ , possibly moving us out of the standard 4-8 GHz range. Among other concerns, most loss mechanisms become worse with increasing frequency. An additional constraint is that smaller Josephson energies  $E_J = \Phi_0 I_c / (2\pi)$  require smaller critical currents and accordingly higher effective Josephson inductance  $L_J = \Phi_0 / (2\pi I_c)$ . With our current fabrication techniques smaller (higher inductance) junctions have less reliable critical currents, typically becoming noticeable when inductances rise towards 8 or 10 nH.

### 4.1.2 I/O couplings: Purcell, driving, measuring, & photon shot-noise

In 3D circuit QED we interact with the outside world via evanescent couplings to capacitive probes, which in turn lead to 50 $\Omega$  coaxial cables. This section will focus on those capacitive probes or “couplers.” (The

cables and attached electronics will be covered in sections 4.3 and 4.4.) Choosing the degree of coupling very clearly comes with trade-offs: with too much coupling the classical world will cause our system to decohere, with too little coupling we will not be able to manipulate or measure our system. Three principal considerations factor into this decision: energy relaxation to the outside world (as introduced in section 3.6), noise due to energy leaking into the system (here focusing on photon shot-noise), and the ability to quickly manipulate and measure the system (section 3.7.1). Additionally, we often work with two couplers, one weakly coupled port for the input and one more-strongly coupled port for the output.

Input couplers are chosen to be sufficiently under-coupled that they do not affect the lifetime of cavity modes. If this were not the case for low-Q readout cavities, measurement efficiency would drop as photons that leave through the input port would not be collected. It is also generally not true for high-Q cavities since we desire them to be as high-Q as possible. In standard geometries this also essentially means that input couplers do not limit the qubit lifetimes. In this case, for the input lines we care only about the trade-off between dephasing due to photon shot-noise and drive strength.

High-Q cavities exist in the regime where  $\chi \gg \kappa$ , which means our qubits will dephase due to some thermal population  $\bar{n}$  of photons in the cavity at  $T_\phi^{-1} = \bar{n}\kappa$  (Sears, 2013). In this case jumps in the photon number in the cavity (which for a thermal state will happen at rate of  $\bar{n}\kappa$ ) will happen less frequently than  $2\pi/\chi$ . Accordingly any single photon number jump event is enough to fully scramble the phase of the qubit. This means we can hope to improve  $T_2$  either by decreasing the thermal population or by decreasing  $\kappa$ , but since we are already undercoupled  $\kappa = \kappa_{\text{internal}} + \kappa_{\text{couple}}$  is no longer significantly under our control. The alternate pressure is to be able to drive quickly, and the result of these two forces is that we want  $\kappa_{\text{couple}}$  to approach, but still be smaller than,  $\kappa_{\text{internal}}$ . This number may not always be predictable, but nonetheless can serve as guidance.

Does changing  $\kappa_{\text{couple}}$  increase  $\bar{n}$ ? With caveats, it should not. This follows from a rough argument. We assume that the thermal photons population due to the temperature of the cavity walls is negligible and that all photon population above this tiny Boltzmann prediction comes from the explicit couplings to the external world. A cavity with two such couplings should come to thermal equilibrium with the microwave lines it couples to, each with its own temperature or photon flux. We can estimate the flux

through our cavity if we treat those coupled fluxes as classical and independent,

$$\frac{da}{dt} = -\frac{1}{2}(\kappa_{c,out} + \kappa_{c,in} + \kappa_{int})a + \sqrt{\kappa_{c,in}}a_{in} + \sqrt{\kappa_{c,out}}a_{out} = 0 \quad (4.1)$$

$$\kappa n_{cav} = 4 \frac{\kappa_{c,in}|a_{in}|^2 + \kappa_{c,out}|a_{out}|^2}{\kappa}, \quad (4.2)$$

where  $\kappa$  is the sum of all  $\kappa_x$ s. We might assume the temperatures or photon fluxes on the input and output side are of similar orders of magnitudes which implies that  $\kappa_{c,out} \gg \kappa_{c,in}, \kappa_{int}$  dominates these dynamics. Following these assumptions, the photon flux through our cavity is roughly the photon flux at the output port. The same argument can be straightforwardly adjusted to cavities without output couplers and to undercoupled cavities.

So we can turn to output couplers. As we have seen, the value of the overcoupled coupler does not have a significant effect on the residual photon population in the standard case. For the output case we are often in the regime where  $\chi \gg \kappa$  (we may still be, but then the discussion is the same as above.) In this case the dephasing rate is related to the measurement rate we discussed in section 3.7.3, only we no longer care about the efficiency  $\eta$ . With a steady-state population, the dephasing rate then goes as  $T_\phi \propto \frac{\chi^2}{\kappa^2}(\kappa\bar{n})$ , leading to the somewhat counter-intuitive result that the photon shot-noise dephasing now scales inversely with  $\kappa^1$ . Rigorously solving for the optimal parameters depends on several additional factors: the background photon flux at the output, the duty-cycle of measurement in the relevant experiment, the dynamic range of the associated amplifier, the relaxation time of the qubit, and the somewhat confounding dependence of that qubit relaxation rate on the cavity population, among others.

Instead, in practice we rarely dwell on this and set  $\kappa_{c,out}$  by looking at the Purcell limit on the associated qubit, after choosing a  $\chi$  to the readout resonator somewhat arbitrarily, generally in the range of 1-5 MHz. We generally choose  $\chi$  and  $\Delta$  which give a ratio of lifetimes between the readout and qubit of roughly one or two thousand, and set the readout quality factor to give an acceptable  $T_1$  limit. This process is often iterative. We can use the Jaynes-Cummings limit for some insight here, where the Purcell rate of the qubit should go as  $\chi\kappa/\Delta$ . To get a feel for the trade-offs we are making

---

<sup>1</sup>For measurement purposes we assume that we can drive sufficient flux regardless of the  $\kappa$ , rather than a fixed  $\bar{n}$ , removing this counter-intuitive scaling.

we can compare this to the dephasing rate, which goes as  $\chi^2/\kappa$ . The third part of this is measurement, where we operate in a regime where we can drive a sufficient  $\bar{n}_{\text{meas}}$  at any  $\kappa$ . That measurement rate scales as  $\Gamma_{\text{meas}} \propto \bar{n}_{\text{meas}}\kappa\theta^2$  where  $\theta^2 \approx \chi^2/\kappa^2$  represents the information per photon.

### 4.1.3 Dielectric loss & feature size

A great deal of research has been done in how dielectric loss limits superconducting qubits and resonators. The general result is that capacitors with larger separations have less field density in the surface dielectric layers. At Yale a thorough study of the geometry of compact resonators (Geerlings *et al.*, 2012) has been conducted, and more recently with 3D transmons in (Wang *et al.*, 2015). The quantum information group at IBM has also done a relatively thorough study on the lifetimes of 3D transmons versus geometry (Dial *et al.*, 2016, Gambetta *et al.*, 2017). The result of this work is that we know creating larger devices with greater separation between the pads decreases surface dielectric loss and leads to more coherent qubits. This desire creates larger qubit modes which may make it harder to limit other undesired couplings. In practice we have seen improving coherence times as the minimum feature gap increases to  $\approx 200\mu\text{m}$ .

## 4.2 Specific devices

### 4.2.1 Flux-tunable two-cavity vertical transmon

One of the early projects in my graduate career was to design, fabricate, and measure a flux-tunable two-cavity “vertical transmon.” Though these results were not published, some relevant information can be found in the later chapters Matthew Reed’s thesis (Reed, 2013). We also used this device for measurements of the Purcell effect and in one of the early designs for the stabilizer measurement experiment. These qubits were essentially microstrip resonators interrupted by a SQUID loop, located in a trench. On either side of the trench the qubits protruded into two 3D rectangular resonators. A second trench, orthogonal to the qubit trench, passed between the cavities, allowing an additional sapphire chip to abut the qubit chip. This chip held a niobium “hairpin” such that current passing along that line drove flux through the qubit SQUID loop. This “flux line” additionally contained a capacitive low-pass




Figure 4.1: **Flux-tunable two cavity vertical transmon.** This qubit was designed to have its resonance frequency change quickly. The large SQUID loop ( $325\mu\text{m}$  by  $325\mu\text{m}$ ) coupled to a niobium “hairpin” on an additional sapphire chip. To fit this chip, the gap between the two 3D cavities was quite large. Accordingly to the right of the SQUID loop we placed a  $10\mu\text{m}$  capacitive gap in the trace, which served to bring the qubit capacitance back to reasonable levels. The area around the gap appears “busy,” which was a necessary fabrication modification. This small gap was often being unintentionally exposed during ebeam writing due to the proximity effect. Decreasing the area written around it lead to a more robust process. The four arrows near the edges of the chip indicate where it was to be aligned relative to the edges of the 3D cavities it coupled to. Additionally, a three-legged grounding strap short-circuits both the junction and the capacitive gap during fabrication. This strap, broken before the qubits are mounted, lead to increased device yield. These roughly centimeter-long devices are particularly sensitive to ESD.

filter (characteristic frequency of roughly 1 GHz), designed to limit qubit decay through this channel, as well as to keep high frequency noise on the line away from the qubit.

These devices generally reached  $T_1$  times on the order of  $10\mu\text{s}$ , and similar  $T_2$  times at the “sweetspot,” or maximum frequency, at which they are first-order insensitive to flux noise. Knowing what we now know, these qubits were likely limited by dielectric loss due to their narrow trench. In these designs the qubit trenches were not much larger than the sapphire chip itself, and their capacitance largely came not directly between the transmon pads but between the pads and the trench walls. The natural capacitance in this design lead to very small anharmonicities, so we separated one of the pads with a  $10\mu\text{m}$  capacitive gap. By placing a small capacitance in series with a large capacitance this decreased the effective capacitance. This gap also likely contributed to the fast relaxation times by increasing losses in the surface dielectric. Additionally, in this design the sapphire chip simply sat on the trench floor, leading to a large field participation near walls of the machined bulk metal, and a great deal of current driven across the seam between cavity top and bottom. In more modern designs (sections 4.2.2 and 4.2.3) the qubits still live in trenches, but they are much wider and the qubit chips are clamped at the edges and suspended in the middle of the chip. Furthermore these are round, seamless, drilled trenches. The later qubits have exhibited much longer  $T_1$  times, occasionally exhibiting  $T_1$  over  $100\mu\text{s}$ .






Figure 4.2: **Schematic of a sapphire chip from the Shiva device.** Towards the right is a 3D transmon, with a central capacitive portion shorted by a junction, roughly centered at the edge of the central cavity wall. Additional coupling “tongues” protrude both into the central cavity, as well as towards a second, lithographically designed  $\lambda/2$  coaxial “quasi-planar” resonator.

#### 4.2.2 Shiva

The Shiva device was used for the final iterations of the stabilizer measurement experiment (chapter 6) and the preliminary iterations of the gate-teleportation (chapter 7) experiment. The device consisted of a central, high-Q stub resonator (Reagor *et al.*, 2013) which coupled to four 3D transmons. Each transmon in turn coupled to an individual  $\lambda/2$  coaxial “quasi-planar” resonator (Axline *et al.*, 2016). The Shiva device was the first to incorporate both of these elements.

This qubit design process began by starting with a traditional 3D transmon design consisting of two large rectangular pads connected by a Josephson junction. The separation was chosen to be  $200 \mu\text{m}$ , which is roughly the maximum separation at which we have seen  $T_1$ s improve by minimizing surface dielectric loss. The junction was placed (somewhat arbitrarily) to be centered at the edge of the tunnel, half-way inside of the 3D cavity. The stub resonator had an external diameter 0.375 inches and a stub diameter 0.125 inches, similar to the device in (Reagor *et al.*, 2013). This ratio of three to one minimizes losses from surface resistance. The cavity continued 0.730 inches above the top of the stub, giving  $\sim 6.5$  attenuation lengths to the lid at the resonance frequency. There are several modes in the waveguide

under cutoff of concern, but if we choose the most pessimistic mode and work in inches we obtain

$$\begin{aligned}
 U_E(z) &\propto \exp(-j2\beta z) \\
 &\propto \exp\left(-j2\sqrt{k^2 - (1.8412/r)^2}z\right) \\
 &\propto \exp\left(-2\sqrt{(1.8412/r)^2 - (2\pi f/c)^2}z\right) \\
 &\propto \exp(-17.644z).
 \end{aligned}$$

For our dimensions this means the electric field energy at the lid is suppressed by a factor of roughly  $2.5 \cdot 10^{-6}$ . As we attempt to get devices with very high quality factors this level of suppression is likely not sufficient. In later devices we have increased that height. Note that for higher modes the cutoff frequency decreases, so as the sapphire chips renormalize the modes one needs to be careful. This can be thoroughly explored in HFSS. The desire for significant depth here creates a large aspect ratio, which is in tension with what machinists are able or willing to create with the requisite precision and at a reasonable price and lead time. The height of the stub was 0.28 inches, setting a dielectrically loaded resonance frequency of 8.117 GHz. The sapphire chips are aligned with the top of the stub to maximize the coupling strengths.

Many of the further design constraints depend on the location of the coupler pins. These were set to have two couplers as close to the central resonator as possible, which resulted in SMA coupler flanges centered 0.384 and 0.775 inches away from the edge of the central resonator. At this point we can choose a location for the quasiplanar resonator (roughly 9mm long for resonance frequencies from 9.1 to 9.4 GHz), which was chosen to minimize Purcell loss for the qubit and central resonator while still allowing an acceptably fast lifetime for the quasiplanar resonator. Note that somewhat counter-intuitively in this design the *closer* coupler is the low-Q output coupler, a phenomenon similar to what was found in (Houck *et al.*, 2008). As we can see from an eigenmode simulation of the qubit mode in figure 4.3, that mode has more field at the far end of the quasiplanar resonator than the close end. The qubit tunnels were chosen to be 4 mm in diameter, guided by previous unpublished results from Chris Axline.

The final consideration was to fine-tune the qubit coupling strengths to the two resonator modes.




Figure 4.3: **Field mode of one of the Shiva device qubits.** This image depicts an HFSS eigenmode simulation performing during the development of one of the qubits used in the Shiva device. Though the simulation was performed with only one SMA coupler, for illustrative purposes I have indicated a second coupler in dashed lines. Note the relatively counter-intuitive result that the qubit mode has more field at the far end of the quasi-planar resonator than it does at the near end. This can be thought of as the relatively higher-frequency resonator looking similar to  $\lambda/4$  at the qubit frequency. This observation lead to the result that the stronger coupled output coupler was closed to the transmon than the lesser-coupled input coupler. The mode of the bus cavity has a similar property. By this “backwards” choice of coupler positions we achieved something akin to a natural Purcell filter.

Using black box quantization (sections 3.5 and 4.5), we targeted the qubits to have frequencies evenly spaced between 4.5 and 6.0 GHz with anharmonicities of roughly 200 MHz (well into the transmon regime). The transmons also targeted dispersive shifts to the central resonator between 0.6 and 1.2 MHz and to the quasiplanar resonators of roughly 5 MHz. To achieve those dispersive shifts, coupling strengths needed to be increased on both sides, and accordingly we added  $150 \mu\text{m}$  wide “tongues.” These increase the dipole moment to the central cavity and the coupling capacitance to the quasiplanar resonator. Achieving the correct couplings and dispersive shifts required iteration between black box simulations for the Hamiltonian and eigenmode simulations to achieve a sufficient Q ratio to balance the tongue length and quasiplanar resonator position. Note that for all four transmons we used identical capacitive geometry, changing only the junction size. This allowed us to be largely insensitive to the scattered frequencies due to fabrication variance, and we were able to reliably obtain acceptable devices from small number of fabricated wafers.

The package, as shown in figure 4.4, is quite busy. In addition to the fundamental design choices

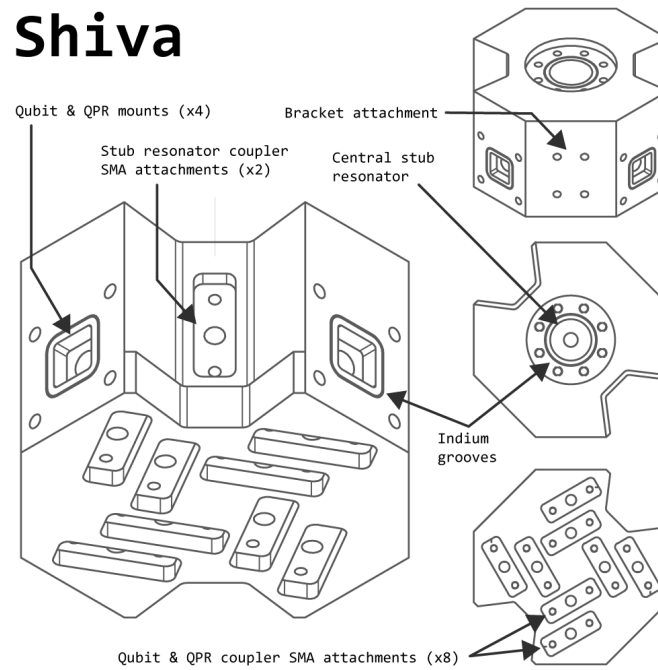


Figure 4.4: **External detail of the Shiva device.** The Shiva device consists of a central stub resonator with four sapphire "arms" protruding at 90 degree angles. Each arm houses a transmon and a  $\lambda/2$  resonator, and can have up to two SMA couplers. There are two additional SMA coupler locations available to the central resonator directly.

above, many other practical considerations were made. The qubit “clamps” were based on a design by Matthew Reagor (Reagor *et al.*, 2013), where the chips are ultimately held in place with indium. The chips were chosen to be long enough that the relevant modes had decayed significantly along the tunnel before reaching the clamp. Note that this is in tension with the desire for wider tunnel diameter to minimize dielectric loss. The clamps are attached to the device with four screws and a gasket of 10 mil diameter indium wire between the clamp and package nominally keeps electromagnetic radiation from entering the device. A similar gasket provides light-tightness around the “cap” at the top of the stub resonator<sup>2</sup>. Each gasket is placed on 5 mil deep grooves machined in the device, though in future designs it may be more practical to machine them on the clamp and cap instead. The clamps also have a protuberance which fits into a matching recess 100 mil on the central package. This aligns the qubits in the center of the tunnels, rather than relying on the screw positions.

The SMA couplers (made of non-magnetic materials) also fit into 100 mil recesses, which aids assembly while minimizing crosstalk between various couplers. That sort of crosstalk was a large problem in some of the earlier 3D flux-tunable devices. The depth of all SMA coupler probes from flange to tunnel (or cavity) was kept less than 400 mil. Minimizing probe length leads to more predictable coupling strengths as the coupler pins often do not seat perfectly orthogonally to the flanges. Note also that the SMA couplers for the qubits and readout resonators were rotated at a slight angle to allow us to pack them in more tightly. The stub resonator was also given two coupler ports, though we only used a one in the final experiment. The chief purpose of this was to allow us to exactly calibrate the coupling  $Q$  at room temperature by permuting three coupler pins. The package supported two couplers for each quasi-planar resonator, such that each qubit could potentially have independent control and readout. In practice, we only used the readout port on one qubit, the ancilla. All four input ports were used—the individual drive lines permit fast and independent qubit control despite their location in a high- $Q$  resonator. We discuss the realized Hamiltonian and coherence properties in section 6.3.

---

<sup>2</sup>It is an open question whether this light-tightness, which in the current approach likely prevents the cavity from being evacuated, is being done correctly.

## Dodekatheon

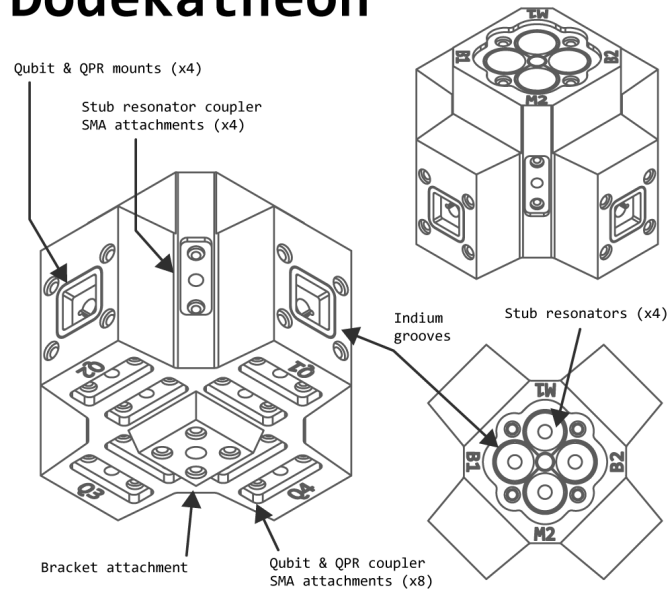


Figure 4.5: **External detail of the Dodekatheon device.** This device is centered around a ring of four high-Q stub resonators. Each pair of which can be linked by a “Y-mon” device that can support its own readout resonator and Purcell filter.

### 4.2.3 Dodekatheon

The Dodekatheon device was a natural evolution from the Shiva device. Used for the gate teleportation experiment (chapter 7), it maintains a similar multi-qubit property to the Shiva device while reflecting the Schoelkopf lab’s evolution away from transmons and towards high-Q cavities as the primary carriers of quantum information. The device consists of four high-Q stub resonators arranged in a ring, with the option for a “Y-mon” (Wang *et al.*, 2016) 3D transmon between each pair. Each stub resonator has a single input port, and each qubit supports its own readout resonator with an input and output port, similar to the Shiva device. One key difference is that these readout resonators are now coupled to the outside world via a “Purcell filter” (Reed *et al.*, 2010b). This is necessary to achieve the even larger Q ratios between stub cavities and readout resonators now required. This Purcell filter is essentially a second sequential quasi-planar resonator. To fit both the resonator and filter structures, they have been compressed and given “wiggles.” These wiggles likely increase their dielectric loss and decrease their internal Q, but as the resonators are highly overcoupled this effect is not a concern.




Figure 4.6: **Dodekatheon qubit, readout resonator, and Purcell filter.** These “Y-mon” qubits couple to two different high-Q stub resonators. Like the Shiva devices, these qubits also couple to a quasi-planar resonator to allow independent readout. As the gate teleportation experiment uses the high-Q modes to store quantum information, we desired a high ratio of quality factors between the stub resonator modes and the readout modes. In practice we achieved a Q ratio of roughly 2,000, and the high-Q mode was still not limited by the Purcell effect.

The stub resonators come essentially in frequency pairs. Two opposing resonators are at 5.124 GHz and 5.275 GHz and are designed to play the role of “memory” resonators in our experiments. Of the other two, one will be used to play the role of a “bus” which sits at 5.693 GHz. The other, unused bus resonator should sit roughly 50 MHz above that. The practical design considerations in this device are generally the same as in the Shiva device. One of the biggest differences is that the Dodekatheon device does not have an open side to support the attachment of a mounting and thermalizing bracket. The bracket is attached by four screws at the top of the device as well as a through-hole in the exact center. A portion of the bracket passes through the device and is threaded on the far side. A washer and nut then apply a compressive force to the cap, holding the device in place.

Another difference is that the sapphire chips required for “Y-mon” qubits are much wider than those of the Shiva qubits. To keep the cutoff frequencies from growing too large, the sides of the tunnel are very close to the chip, and the tunnel is not a circle. It is shaped essentially as a diamond with rounded edges.

One of the final lessons learned relates to using high-purity aluminum, which is necessary to get millisecond lifetime cavities. It is also necessary to etch the surface prior to packaging. Since we do not want to ruin the tapped screw holes, during the etching process aluminum screws are put in place. In the Dodekatheon we put a small 30 mil recess around each screw hole, larger than the threaded region but smaller than the screw head. The motivation for this is that the etching process often leaves a

raised burr-like region around the screw holes. Without these recesses the screws do not make good contact with the device, and worse-yet, nor do the qubit-chip-holding clamps.

## 4.3 The control chain

Full control of either a resonator or a qubit mode requires a microwave drive with control over both the amplitude and phase of the signal. There are several other considerations beyond this core need. Maintaining phase coherence also requires stability in the frequency (phase) of the microwave source, at least on the timescale of the experiment. These “clock errors” are quite concerning, as they lead to the highly correlated errors which are not permitted in fault-tolerance calculations. High-fidelity gates require stability in the amplitude of the control on a much longer timescale. We also demand that the control line not introduce too much incoherent noise, largely by attenuating black-body radiation and limiting noise introduced by any amplification elements. Finally, we also worry about the coherent noise, sometimes called “spurious tones,” inevitable in any drive chain with non-linearity or any multi-tone system with finite directivity.

### 4.3.1 Single-sideband control

IQ mixing, or more specifically single-sideband modulation, is the main technique for microwave pulse shaping used in these experiments. Why an IQ mixer rather than a simple mixer? At zero IF, a simple mixer does not allow phase control. At finite IF, a simple mixer creates both positive and negative sidebands, causing at least one spurious tone<sup>3</sup>. An IQ mixer consists of two simple mixers, two splitters, and 90 degree phase shifter. The LO is split, half is mixed with the I input, and half is mixed with the Q input. The Q branch is phase shifted 90 degrees, and both branches are combined at the output. At zero IF this clearly gives phase control just by varying the relative amplitudes of the I and Q inputs.

More interesting is when we modulate the I and Q inputs at a finite frequency. While these steps are done in software, it’s useful to think about how a physical implementation of this scheme might work. We generate a continuous tone at the intermediate frequency, split it, and phase shift half by  $\pi/2$ . We

---

<sup>3</sup>This concern could also be mitigated with an Image Reject mixer.



take a complex zero-IF pulse envelope and mix it against each IF tone. The in-phase tone goes to the I port, and the out-of-phase tone to the Q. This results in a single sideband with full amplitude and phase control at the LO frequency plus or minus the IF, depending on whether the quadrature drive is leading or lagging.

Optimizing the performance of these IQ mixers requires nulling any DC offset (which leads to leakage at the carrier frequency) and ensuring that the I and Q channels have equal amplitudes and are precisely out of phase (which can lead to the re-emergence of the undesired sideband). We typically do these optimizations programmatically, using a spectrum analyzer set at the LO or the undesired sideband frequency and a derivative-free optimization algorithm like Nelder-Mead or SNOBfit. I also observed early in my time at the Schoelkopf lab that IQ mixers are temperature-sensitive and that placing them in a sealed box can result in more stable performance. The exact origin of the DC offsets is unclear and can be quite sensitive to the grounding configuration, including to what is attached to the output of the IQ mixer.

One outstanding question in this scheme is how sensitive the amplitude and phase imbalance are to the intermediate frequency. We regularly see the optimal parameters change on the scale of five or ten megahertz. Since real pulses have a finite spectral width, the leakage terms may exhibit odd behavior. This phenomenon is generally inevitable as there will be some imperfect dispersion between the envelope and mixer, but it may be possible to measure and calibrate away. Another concern is the exact nature of the non-linear behavior. Above a certain IF amplitude, peaks at the carrier frequency plus unexpected integer multiples of the IF appears. This effect is worsened when the LO power is out of or near the edges of the mixers specifications. To my knowledge the details of this behavior have not been thoroughly explored.

### 4.3.2 Optimal control theory

One of the wonderful things about transmons qubits is that they quite naturally give us universal control. As we discussed in the previous section, we can have full control over the phase and amplitude of resonant microwave tones. Thanks to the transmon's simple nature this gives us  $\mathcal{H}_{\text{control}}(t) \propto V_I(t)\sigma_x + V_Q(t)\sigma_y$ . In the transition to using harmonic oscillators to store quantum information we have given up that simple

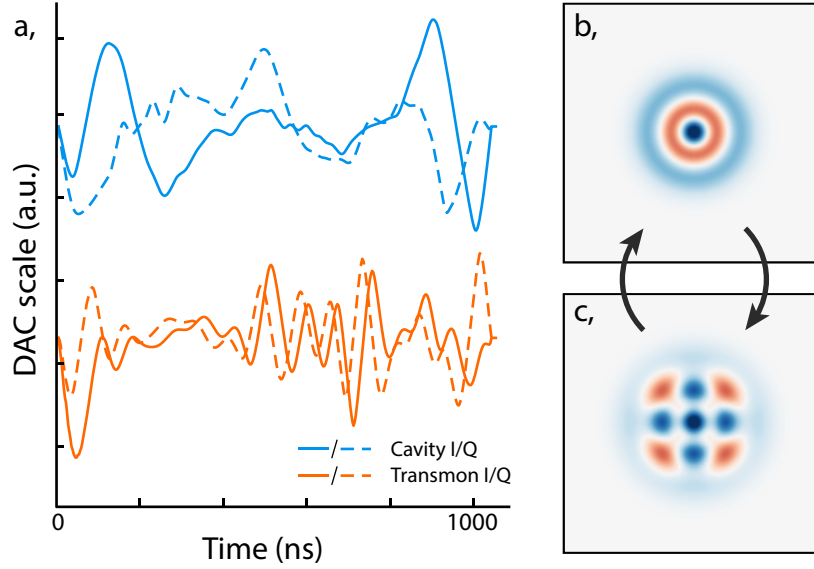


Figure 4.7: **Example optimal control pulse of an  $X_\pi$  gate on an encoded “kitten” qubit.** As an example of a non-trivial manipulation, figure **a** depicts the four control waveforms, corresponding to the voltages on the I and Q ports of the IQ mixers associated with cavity and transmon modes. The two waveforms are offset vertically—all start and begin with zero voltage. Figures **b** and **c** depict Wigner functions of logical  $|\pm Z\rangle$  states,  $|2\rangle$  and  $\sqrt{1/2}(|0\rangle + |4\rangle)$ , which are swapped under the action of this pulse. This “kitten” qubit is discussed in section 7.7.

control. In turn we now have  $\mathcal{H}_{\text{control}}(t) \propto V_I(t)(a + a^\dagger) + iV_Q(t)(a - a^\dagger)$ , which by itself only allows us to perform displacements in phase space. This allows us only access to coherent states if we start in the ground state (or any other coherent state). To reach more exotic states we must make use of the nonlinearity gained by the cavity’s dispersive coupling to a transmon. How to do this is not at all obvious, and to reach truly universal control we generally let a computer solve the problem for us.

To solve this problem our resources are a time-independent drift Hamiltonian  $\mathcal{H}_0$ , say the dispersive interaction, and a control Hamiltonian  $\sum_k u_k(t)\mathcal{H}_k$ . Here the time-dependent  $u_k(t)$ , might correspond to the  $V_I(t)$  and  $V_Q(t)$  on the single-sideband mixers resonant with a qubit and a cavity. We discretize these  $u_k(t)$  in time-steps  $\delta t$ . With that information we can compute the total realized propagator,

$$U(\tau) = \prod_{t=0}^{\tau} \exp(-i(\mathcal{H}_0 + \sum_k u_{k,t}\mathcal{H}_k)\delta t/\hbar). \quad (4.3)$$

If we specify an initial state  $|\psi_i\rangle$  and a desired final state  $|\psi_f\rangle$ , we can also use this propagator to

calculate a (square root of) fidelity,  $\mathcal{F} = |\langle \psi_f | U(\tau) | \psi_i \rangle|$ . With this construction alone a given set of control waves allows us to estimate a fidelity. At this point we could use a derivative-free optimization algorithm such as Nelder-Mead to find an optimal control sequence. Unfortunately this would be quite slow, as the space for practical pulses can easily reach thousands dimensions, e.g. four microsecond-long controls, discretized to the nanosecond. It would be much more effective if we were also able to efficiently calculate the gradient of  $\mathcal{F}$  at each point in the parameter space. It turns out that calculating that derivative can be done straightforwardly, and the optimization of  $\mathcal{F}$  using that derivative is called Gradient Ascent Pulse Engineering, or GRAPE (Khaneja *et al.*, 2005). We can write this gradient more verbosely,

$$\frac{\partial \mathcal{F}}{\partial u_{k,t}} = \langle \psi_f | U_{\tau/\delta t} \dots \frac{\partial U_t}{\partial u_{k,t}} \dots U_0 | \psi_i \rangle \quad (4.4)$$

$$\frac{\partial U_t}{\partial u_{k,t}} = \frac{\partial}{\partial u_{k,t}} \exp(-i(\mathcal{H}_0 + \sum_l u_{l,t} \mathcal{H}_l) \delta t / \hbar), \quad (4.5)$$

which we can clearly calculate numerically. We also add straightforward amplitude and derivative penalties to make sure the pulses are limited to what our DACs can physically realize. We generally ask for transfers between more than a single pair of input states. The coherent combination is calculated given by scaling and summing the two fidelities before taking the norm. This allows us to define operations like a “logical X” by two state transfers. Defining operations in terms of state transfers can be much more efficient than explicitly defining the unitary as we may not care what happens in many cases, e.g. when the initial state is not in our logical space. In this example we would define two state transfers, rather than the  $d$  equal to the system dimension for a full unitary.

Now that we can calculate the fidelity and gradient of a fidelity with respect to a set of control parameters, we can feed those to an optimization algorithm that will take advantage of them. The base problem as presented is believed to be convex, but we break that convexity when we add in our amplitude and derivative penalties, and there are additional complexities introduced when we truncate the size of the available Hilbert space. Rather than using convex optimization algorithms we settle for a quasi-Newtonian method, the BFGS algorithm. This algorithm uses iterative sampling to locally estimate the Hessian matrix which it additionally uses as it searches to maximize the fidelity. BFGS is

a popular algorithm for this purpose, but it may be an interesting question going forward as to whether we can find a more effective optimizer.

One final and very useful topic is that of gauge freedoms. We can make a slight modification to allow some extra action, e.g. a CNOT operation between two qubits where we do not mind additional single-qubit phases. This can be extremely useful, making operations quite a bit easier. This is done by choosing an additional set of operators, in the previous example  $\hat{g} \in \{IZ, ZI\}$ , and allowing them to act on the final state (before fidelity is calculated) as  $\prod_i \exp\{-i\theta_i \hat{g}_i\}$ . These  $\theta_i$  are then optimized over exactly as (and simultaneously with) the control waveforms, only not subject to amplitude penalties. Intuitively this means in our example that after our operation we immediately apply whatever phases are needed to correct the state. Specific uses of gauge freedoms and several other details in the operational use of optimal control theory are discussed further in sections 7.7.2.

### 4.3.3 Input lines

Single-sideband signals start as a local oscillator (LO) and a pair of intermediate frequency (IF) signals. Generally speaking our local oscillators (Vaunix and Agilent microwave sources) produce a very clean spectrum. This is not as true of our IF signals (from Tektronix or Innovative Integration DACs). The mixing process and amplification also add noise to the system, and of course so does the finite temperature of the entire line. Accordingly we use filters to remove parts of the signal we do not need, we use low-loss components, we use the lowest amplification gain we can, we choose low-noise amplifiers, and we attenuate and thermalize the lines as needed.

The IF signals out of a Tektronix 5014C pass through a 250 MHz filter internally. The signals from an Innovative Integrations card do not, and the sharp steps from the 1 gigasample/500 megasample DAC are readily apparent on an oscilloscope. We use 370 MHz lowpass filters from Minicircuits (BLP-300+) to remove these high-frequency components.

These signals pass to double-balanced mixers from Marki Microwaves (IQ-[0307/0618]-[LXP/MXP]), which have roughly 5 dB of conversion loss. The maximum IF voltage is roughly 0.5 volts peak-to-peak before these generators start compressing and showing higher order effects. In particular spurious tones and  $\omega_{LO} \pm n\omega_{IF}$  start to appear for  $n$  between two and perhaps five. This effect is further pronounced

if the LO tone is not at the correct power<sup>4</sup>. As the Tektronix 5014C outputs up to 4.5 volts peak-peak and the FPGA outputs 1 volt peak-peak, we generally use between 3 and 6 dB of attenuation. This allows greater dynamic range and resolution (use of the bounded and discrete DAC output values), as well as attenuating any additive noise from the DAC.

We generally pass the output signal of the generator through K&L Microwave low-pass filters with a cut-off frequency between 8-12 GHz. This limits the high frequency noise as well as blocks any higher harmonics from the microwave source or mixer. If amplification is required we generally use Minicircuits ZVA-183+, which have a noise figure of roughly 3dB, gain of 26 dB, and 1 dB compression point of +24 dBm. Depending on line configuration this amplifier has been seen to raise the base temperature of the fridge. This is particularly worrisome because the power is dissipated in the lines, meaning that parts of the lines may be quite a bit warmer than the base plate. A better amplification solution will need to be found eventually. Particularly when using amplifiers we have often found it useful to employ microwave switches. The connectorized Hittite HMC-C019 devices perform quite nicely. They have a switching time under 10 nanoseconds, an on/off ratio of greater than 80 dB, and an insertion loss of less than 3.5 dB. If they are driven with too much power (roughly 20 dBm) they can drastically distort the signal, so one must be careful about how much power is used.

At this point our signals pass through inner-outer DC blocks, which serve to control the electrical grounding of the refrigerator. The input lines inside the fridge are custom-made of UT85SS-SS on the way to the base plate. At the 4K stage the lines are attenuated by 20dB, which thermalizes the lines at that point and attenuates room-temperature thermal noise. The lines are attenuated again at the base plate, generally with 30 dB of attenuation, for the same reason. When significant power is required on an input line we have used the “direct-tenuator” technique. This involves using a directional coupler (from Krytar) at the base plate, where now rather than resistively dissipate 20dB of power we pass it back to a terminator at the 4K stage, then at base we apply an additional 10dB of cold attenuation. This is a much better solution for dissipating noise, but sacrifices an additional input line, which is not always an option. All of these attenuators are nominally thermalized, clamped against a copper braid

---

<sup>4</sup>For this reason one should be wary of using a MXP (design LO power = +13 to +16 dBm) mixer with a Vaunix LabBrick (maximum output power 13 dBm), as some power will be lost in the interconnect.

which is then fastened to the refrigerator<sup>5</sup>.

After attenuation (though sometimes before) we generally apply a commercial low-pass filter (the same K&L Microwave filters we use at room temperature.) These are thermalized in the same manner as the attenuators. We additionally use a home-made low-pass filter filled with an infrared-absorbing material (Emerson & Cuming Eccosorb CR110), which are dimensioned to have a matched impedance. While the K&L filters have a fairly narrow effective filter range, the Eccosorb filters should have an exponential roll-off for a very wide spectrum. The Eccosorb filters are placed either right after the K&L filters or may be placed inside the experiment's Cryoperm magnetic shield. We have anecdotal support that this leads to colder modes. After passing to base we stop using steel coaxial lines, and switch instead to home-made copper UT85C cables with Phoenix 2202000868J SMA connectors. These have been measured to be quite non-magnetic, and thermalize more effectively than the steel lines. The final piece of the input lines are Coaxicom SMA couplers, which we have also measured to be non-magnetic.

#### 4.3.4 Output lines

The signal emanating from the sample is also carried via the same UT85C cables in these experiments, though others have begun to use superconducting hand-formable cables from Coax Co. Loss is very important at this point as our measurement signal consists of very few photons before it reaches the parametric amplifier. We have occasionally used an additional Eccosorb filter on this line. This degrades our readout performance slightly (these filters generally attenuate  $\approx 1$  dB at our frequency of interest, but also seems to lead to colder qubits and readout resonators. This increase in coherence may offset the cost in readout signal.

Before reaching the amplifier our signal passes through two QuinStar circulators which provide 20dB of isolation at the cost of 0.3 dB (maximum) of insertion loss. The signal passes through a Krytar 180 degree hybrid (4060200) both to and from the JPC (Bergeal *et al.*, 2010) which amplifies in reflection. The amplified signal still travels on UT85C lines and passes through a circulator once en route to a pair of isolators (QuinStar QCI-075900X000). The signal then passes through an Anritsu bias-tee and enters a superconducting coaxial cable. This cable travels from the base place of the refrigerator before entering

---

<sup>5</sup>I am skeptical that these are sufficiently effective, and they may be responsible for some of our decoherence and finite thermal populations. Finding a better solution may be a worthwhile endeavor.

a wideband HEMT amplifier by Low Noise Factory en route to the room temperature electronics.

## 4.4 The room-temperature measurement chain & analysis

As discussed in section 3.7.3, the root of measurement in cQED comes from the dispersive shift. The frequency of a microwave resonator depends on the state of a coupled transmon, and in turn so does the phase and amplitude of a microwave signal transmitted through that cavity.

### 4.4.1 Heterodyne measurement

In standard dispersive measurement our aim is to transmit a microwave signal through a resonator and observe the returning signal to decide which qubit state it indicates. Recording a signal oscillating at several gigahertz is a difficult task and is not necessary. Instead we mix this returning signal against a local oscillator (LO) tone, resulting in an intermediate-frequency (IF) signal at the difference of the readout and LO frequencies. If the LO frequency is the same as that of the readout tone, this is homodyne measurement, resulting in a DC signal. In the experiments of this thesis we have exclusively used *heterodyne* measurement with a 50 megahertz IF. Heterodyne measurement allows us to ascertain information both in the amplitude and the phase of the returning signal. Additionally, our measurement is sensitive to noise principally at the IF frequency. By moving the IF to 50 MHz rather than DC we see greatly reduced noise.

### 4.4.2 Two-tone interferometer

In the early experiments of this thesis, most prominently the stabilizer measurement experiment, the measurement analog-to-digital converter (ADC) was an Alazartech ATS9870. The readout pulse came from a microwave generator at the frequency of the readout signal, digitally gated by a Tektronix 5014C arbitrary waveform generator (AWG). The AWG also digitally triggered the ADC. As the phase of the AWG trigger relative to the readout generator is not stable on long time scales, this tone traveled through an interferometric circuit designed to leave the signal insensitive to the absolute phase of either.

The split readout signal traveled in two directions, branching into *reference* and *signal* arms. The

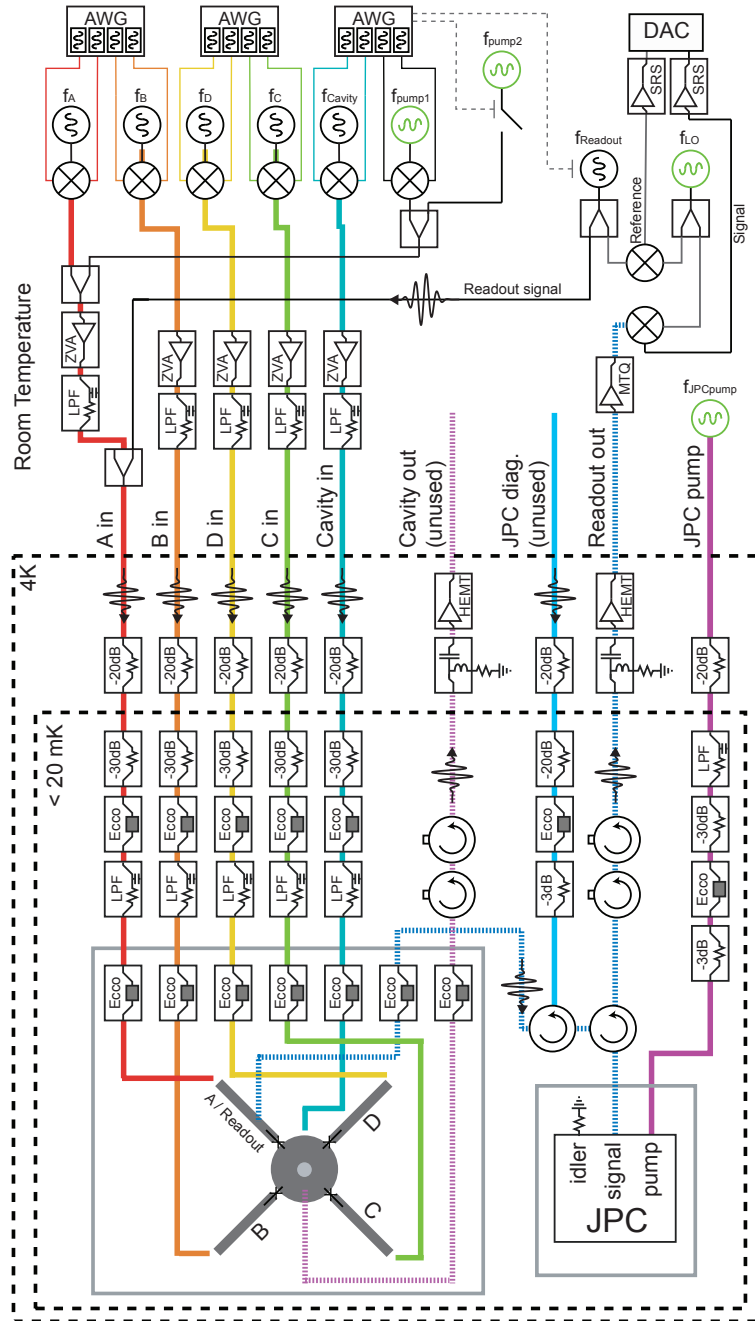


Figure 4.8: Fridge and room temperature wiring during the stabilizer experiment.



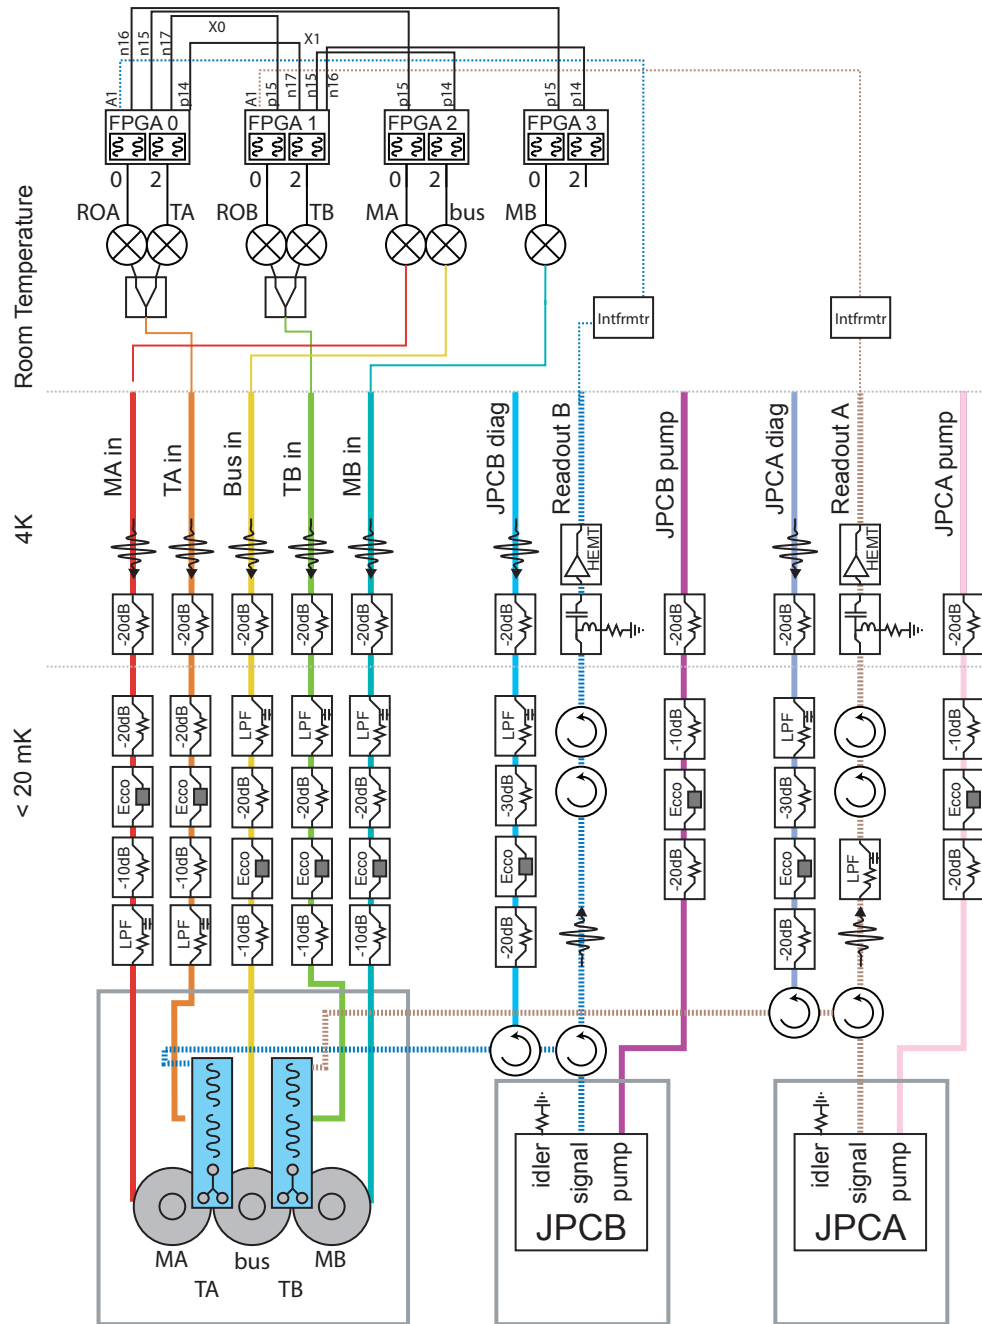


Figure 4.9: Fridge and (partial) room temperature wiring during the gate teleportation experiment.

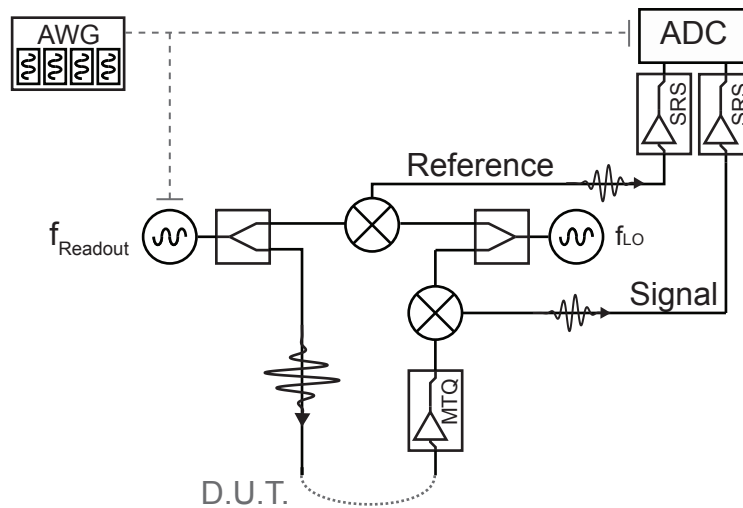


Figure 4.10: **Two-tone interferometer.** In a two-tone interferometer two generators are used, one at the readout frequency and the other 50 MHz away. The readout tone is gated, rather than modulated. The readout signal is split, with half comparing directly to the other generator, forming a reference branch, and the second passing through the experiment before mixing with the other generator, forming a signal branch. The frame of the signal is then rotated relative to the reference branch.

signal passed into the fridge, was modified by the dispersive shift, and returned to be amplified by a low-noise Miteq amplifier (AFS3-04000800-ULN) at room temperature. The signal and reference tones were both mixed against an LO tone set to be 50 megahertz (the IF) away from the readout frequency. Each of these two down-sampled signals were then amplified by a Stanford Research Systems SR445A, and then were simultaneously digitized by the ADC. Both the readout and LO generators have a random phase from shot to shot, but we can ascertain this random phase from the reference tone. We then shift the signal by exactly this reference phase, and the phase of the resulting modified signal depends only on the effect of the cavity transmission. We often place Ditom isolators in this interferometer to assure no leakage signals bypass the fridge.

### 4.4.3 One-tone interferometer

In the gate teleportation experiment (chapter 7), we moved from using the Alazar ADC and Tektronix DAC to our new FPGA solution. The FPGA card also contains both DACs and ADCs. Since these two functions are always in phase, we may now perform a simpler form of demodulation. In one-tone demodulation we use only one generator, now set 50 MHz below the readout frequency. This signal




Figure 4.11: **One-tone interferometer.** This simpler measurement technique is made possible by having a phase-locked DAC and ADC. We use the same SRS and Miteq amplifiers, but to compensate the reduced sensitivity of the Innovative Integrations ADC we also amplify the 50 MHz signal with a Minicircuits ZFL-500LN.

is split, and half is single-sideband modulated by shaped 50 MHz tones generated by the FPGA. This modulated frequency passes through the experiment, and is then mixed against the original unmodulated tone. This 50 MHz signal is recorded by the FPGA in the same frame used to generate the original modulation, canceling out all added phases.

As we modulate this tone we pay the price of a DAC pair, but in exchange we save one microwave generator and we may choose more useful shapes for our readout tones. This is particularly useful when we care about emptying the readout cavity more quickly than its natural decay time. There are many different schemes to achieve this, perhaps most intuitively by immediately feeding back to conditionally displace the cavity to the vacuum. This is more experimentally difficult, so one can choose pulse shapes that inherently ring the cavity up and down unconditionally (McClure *et al.*, 2016). We have also experimented with off-resonant “adiabatic” pulses and direct optimizations of the cavity’s classical equations of motion. It is also possible to use GRAPE to optimize readout using the master equation, but this is computationally expensive and our attempts have not yielded particularly encouraging improvements.

## 4.5 HFSS

In the Schoelkopf and Devoret labs we use Ansys HFSS, a high-frequency electromagnetic finite element solver, to great extent. We use it in two basic modes: “Eigenmode” simulations which do essentially what the name implies, and “Driven Modal” simulations where we find the response of the system when a port is driven at a swept frequency. We use HFSS for three basic purposes:

1. Using Eigenmode to visualize the spatial extent of resonant modes, and analyze the participation of the mode in regions of interest.
2. Using Eigenmode to determine inherited loss and the Purcell effect.
3. Using Driven Modal to find the current driven across a port by each mode, which in conjunction with black box quantization tells us the Hamiltonian of the system.

In this section I will give several tips and tricks which I apply, some of which have been determined empirically or anecdotally and some of which have come from discussions with Ansys engineers. When we model a junction in HFSS we include a lumped inductor and (if desired) a lumped capacitor of several femtoFarads, generally corresponding to a junction self-resonance frequency in the neighborhood of 40 GHz. This capacitance will be a relatively small perturbation to the overall transmon capacitance. The geometry of the junction rectangle should have an aspect ratio of less than two to one, according to Ansys engineers. We generally place a second rectangle in parallel which can play the role of a lumped port in driven modal simulations. It is often useful to have the exact same geometry for both modes of simulation as then meshes can be shared across them. For eigenmode simulations this extra two-dimensional rectangle should be a non-model object.

Eigenmode simulations are relatively straightforward. Some relevant considerations are what initial mesh seeding to apply, the convergence criteria, and the HFSS simulation settings. The most important seeding is on the surface of the qubit's perfect conducting sheets. Without fail HFSS will initially overestimate the qubit capacitance, and seeding on this surface can save significant time. Generally this means a surface mesh with a maximum length on the order of 25 microns. The relevant convergence criteria depend on the purpose of the simulation. For Purcell estimations it is useful to add the mode quality factors to the convergence criteria. In all cases the frequency convergence often needs to be tightened to a smaller number than the default, perhaps 0.5% or less per pass, which corresponds for a 5 GHz qubit mode to 25 MHz per update. There are several modifications to the default simulations settings that I have found useful. We want to use mixed-order interpolation between the tetrahedra, as well as applying lambda refinement. Generally speaking we do not care to save the fields, which results in huge saved files which we do not normally need.

Generally Purcell simulations mean including the physical coupling port and pin, a lumped resistor representing the attached coaxial line, and performing an Eigenmode simulation. We vary the coupler until the readout mode is at the desired  $Q$  and examine the  $Q$  of the qubit mode. This comes with several caveats, however. Notably, from discussions with Ansys engineers we know that they find eigenmode less trustworthy than the driven-modal simulations. It is important to watch the convergence of the  $Q$  versus adaptive pass, and verify that it has stabilized. For large  $Q$ s, perhaps over  $10^7$ , this stabilization may not be reliable. To circumvent this, we assume that the coupler itself is a small perturbation on the modes, and decrease the coupling  $Q$ . We do this to bring the qubit  $Q$  down to the range that we consider trustworthy. We note the ratio of  $Q$ s between the readout and the qubit, and extrapolate back out to higher readout  $Q$ s. This method has proven reliable thus far. We have looked at the ratio of  $Q$ s versus the coupling  $Q$  and found it to be quite constant.

The driven modal simulations we use for black box quantization have a bit more subtlety. When we perform driven modal simulations in the “fast sweep” more it chooses one driving “solution” frequency and improves the mesh at that particular point. This is the frequency at which the solutions are most trustworthy. It then uses that mesh to perform a second step which extrapolates out to find the fields and response function at all frequencies. To get the most reliable results we use one solution frequency per mode of interest with a solution frequency near each. Importantly these solution frequencies should not be *too* close to the modes of interest or one will run in to convergence problems. Convergence is (at minimum) determined by the stability of the S parameters at the solution frequency. If the solution frequency sits near a divergence (resonance) in the S parameters, very small shifts in the resonance frequency of the mode will prevent convergence. Generally speaking in these analyses we do not care about small shifts in these frequencies, rather we care about the characteristic impedance of each mode.

The final approach that I settled on reliably produced quite predictive results for relatively little work. It begins by creating a model appropriate for a driven modal experiment. I copy that model, make the “port” rectangle into a non-model object, and perform an eigenmode simulation. After this reaches a reasonable but not exact convergence, say 1%, we switch to driven modal. The new driven modal design should have a solution near<sup>6</sup> each mode and each solution should have a fast sweep of

---

<sup>6</sup>Generally about 20 MHz below the eigenmode frequencies, as modes tend to move to higher frequencies as they converge.

appropriate range. The first solution is instructed to import the mesh from the eigenmode solution, and all later solutions in turn import the mesh from the previous solution. This can be set up in advance, and under the “setup link” menu each solution can be instructed to “simulate source design as needed.” In this manner a complex series of simulations can be run unattended. Many of the choice of simulation settings are the same as we chose for Eigenmode simulations. Anecdotally, I have found that the “direct solver” is more reliable than the “iterative solver,” which occasionally simply fails and presents an error.

## 4.6 Conclusions

Most of the physics we would like to demonstrate can be described thoroughly within the Hamiltonians of circuit QED, as discussed in chapter 3, but none of that would be possible without proper everyday experimental techniques. The devices and methods presented here have enabled some of that very rich playground. Enacting a circuit QED experiment requires hundreds of choices, from signal generators to pulse shapes to attenuator thermalization to SMA couplers and capacitor geometries. The very high-level description of our field would provide little guidance on these decisions, but decades of institutional memory and the trial and error of countless graduate students have lead to a fairly reliable prescription. In this chapter I hope to have left my way of looking at some of the elements of this recipe.

In the future several aspects of this chapter will become obsolete or need to change. As we consider scaling towards larger and larger number of modes, many of the current techniques will no longer be viable. Micromachining, for instance, will lead to very different looking package designs (Brecht *et al.*, 2016). For very large numbers of qubits it will no longer make sense to route control electronics up too room temperature, and we may need DACs, ACDs, and logic to be inside of the fridge (Stoutimore *et al.*, 2015). Nonetheless, many of core ideas will remain relevant as long as circuit QED does.

## Chapter 5

# Measurement characterization

It is ... essential that the thought experiments testing [quantum mechanics] have the ability to perform repeated measurements on single particles. The physicist should then be compared not to the paleontologist but rather to the biologist operating *in vivo*.

---

Serge Haroche, Exploring the Quantum

### 5.1 Introduction

Discussions of quantum computing and fault tolerance focus largely on the quality or fidelity of one- and two-qubit gates and spend comparatively little time discussing measurements. What do we mean when we say “measurement fidelity” and why has it gotten relatively little attention? The answer to the second question is relatively simple: a naive fault-tolerance analysis makes only very simple demands of measurements. In standard error-correction circuits, all of the following are assumed to be true:

1. Measurements act only on the desired two-dimensional Hilbert space
2. Measurements are strong
3. Measurements are of two outcomes
4. Measurements are projective, with ideal back-action

5. *or*, more often, the measured qubits are ancillae, discarded after each measurement, so back-action is unimportant

These assumptions are often taken for granted, but in no way capture the richness of the phenomenon of measurement in quantum mechanics. Unfortunately for quantum computing, the first of these will not generally hold, and the last may have more subtle complications than one might immediately think. For example, ancillae are reused in most architectures. When these conditions are not true, it can be useful to have a more sophisticated definition of measurement fidelity than is commonly used. In the stabilizer measurement experiment we considered these issues and devised several analyses and figures of merit.

We will look through several of these analyses, first focusing only on the *detector*, by which I mean the way in which the outcome of the measurement is sensitive to the input state. After that discussion we will examine the full *measurement process*, which additionally includes what happens to the input state after the measurement result is recorded. Classical measurement outcomes will be indicated by quotation marks, e.g. a qubit measurement might result in “g.”

## 5.2 Why characterize measurements holistically?

If it is true that, strictly speaking, we do not need to characterize complex measurements for fault tolerance calculations, why bother? In short, it allows the experimentalist to sweep fewer things under the rug, and it is a useful catch-all diagnostic tool. As a more “compiled” operation, the full measurement may be more sensitive to non-idealities than concatenated fidelity estimates of smaller one- and two-qubit operations. A traditional  $ZZZZ$  measurement would consist of four CNOT gates and a single-qubit measurement of an ancilla. Nominally, all we should need to know is the fidelity of the ancilla preparation, each of the two-qubit gates, and the accuracy of the ancilla measurement. In reality many other negative effects may be occurring. In addition to the other qubits decohering, they may have residual interactions. Our measurement of the ancilla may dephase or measure other qubits. In most cases the duration of the measurement takes a non-trivial amount of time. Very realistically, gates between two qubits often affect or are affected by the state of other qubits. The measures we will derive will make these effects






Figure 5.1: **Potential errors during stabilizer measurements.** Figure **a** shows a standard depiction of the surface code. We take a large two-dimensional sheet of data qubits, and going along it alternately measure the  $ZZZZ$  and  $XXXX$  correlators of nearest neighbors with special-purposes ancillae. These ancilla are assumed to have independent readout channels. Figure **b** zooms in and examines the nominal measurement of a  $ZZZZ$  correlator. It should depend only on the quality of ancilla preparation, the four two-qubit gates, and the accuracy of the ancilla measurement. Figure **c** shows some of the things that may go wrong and be unaccounted for in the naive picture. Here the top two data qubits have an always-on residual interaction. The third and fourth qubits experience decohering bit- and phase-flips, respectively. The fourth CNOT gate, nominally between qubit one and the ancilla, seems to depend on the state of qubit three. The pulse used for the ancilla measurement dephases qubit two.

evident within a larger subspace.

### 5.3 Assignment fidelity

In the simplest case and under the previous assumptions, one only cares about one or perhaps two numbers—when the qubit is prepared in  $|e\rangle$ , how often does the measurement register “e?” When the qubit is prepared in  $|g\rangle$ , how often does the measurement register “g?” In the literature the average of those two probabilities is often simply called the “measurement fidelity,” though here for clarity we refer to this as the *assignment fidelity*. This quantity does not have an obvious meaning for non-orthogonal measurements, weak measurements, or measurements with degenerate eigenspaces (e.g. a measurement of  $ZZ$ ). It also doesn’t speak at all to what happens to the qubit afterwards. The average itself may be problematic—we know that fault-tolerance calculations generally depend on the worst-case behavior.

Despite this, under the right set of assumptions, that quantity or something like it is quite useful, and for measurements of only rank-1 operators it is easy to measure. When working with degenerate measurement eigenspaces there are not a finite number of points of definite measurement outcomes,

which complicates the matter. Our solution in (Blumoff *et al.*, 2016) was to prepare an overcomplete set of states, in our case by performing rotations around the  $X$  axis of each qubit. We then fit that data, assuming that all signal comes from some weighting on the expected measurement operator and an offset. The magnitude of the contrast we can ascribe to the measurement operator is closely related to a more generalized notion of assignment fidelity<sup>1</sup>.

At the cost of a more complicated experiment, we can also be more rigorous. With assignment fidelity we have always assumed that we know what we are trying to measure, but what if we face an unknown measurement? This can lead to us strongly misjudging our detector. As an example, imagine that we desire a one-qubit measurement of  $Z$ , but in reality our meter performs a perfect measurement of  $X$ . The assignment fidelity we would ascribe to this detector is 50%, exactly the same as we would assign a random number generator. With a little more inspection, however, we would see that a unitary rotation would turn this random number generator into a perfect measurement device. Before diving into this sort of analysis we will first develop some additional formalism.

## 5.4 Detectors & POVMs

The most general formalism for a quantum detector is that of the POVM, or positive operator-valued measure. This formalism begins with the axioms that the outcome probabilities of a detector are positive and sum to one; *some* outcome always occurs. A detector is then given by a set of observables  $\{E_i\}$ , each of which has non-negative eigenvalues, and which sum to the identity. These observables are such that the probability of each outcome  $i$  is given by the expectation value  $P_i = \text{Tr}[\rho E_i]$ . Since both the operator and the density matrix are positive, these probabilities are also positive and therefore meaningful. Since the trace is linear,  $\sum_i P_i = 1$  leads to  $\sum_i \text{Tr}[E_i \rho] = \text{Tr}[(\sum_i E_i) \rho]$  which is only true for all  $\rho$  if  $\sum_i E_i = I$ . Any set  $\{E_i\}$  that fulfills these properties is a valid description of a detector. We will restrict this discussion to finite sets, though an infinite representation is needed to fully describe phenomena such as (non-thresholded) dispersive measurement.

For intuition we can look at several examples—some of which which will revisit in later sections.

<sup>1</sup>When a measurement outcome is random, the contrast is zero, however you get the correct answer half of the time. Since fidelity corresponds to a probability of being correct,  $F = \frac{1}{2}(1 + C)$

Trivially, for a perfect  $Z$  measurement, the POVM is given by

$$\{|e\rangle\langle e|, |g\rangle\langle g|\}. \quad (5.1)$$

These have both eigenvalues of 0 and 1, and add together to the identity. What about an imperfect measurement? A POVM with assignment fidelity  $F$  might look like

$$\{F|e\rangle\langle e| + (1-F)|g\rangle\langle g|, F|g\rangle\langle g| + (1-F)|e\rangle\langle e|\}. \quad (5.2)$$

This again clearly has positive eigenvalues and adds to the identity. Since we are not yet looking at back-action, we cannot tell if this represents either a noisy strong measurement with a classical loss of information or a perfectly functioning “weak measurement.”

A completely random (binary) number generator can be given by

$$\{0.5(|e\rangle\langle e| + |g\rangle\langle g|), 0.5(|e\rangle\langle e| + |g\rangle\langle g|)\}. \quad (5.3)$$

This fulfills the POVM requirements while providing no information at all. This inspires a question to be answered in the next section: What if your desired POVM is, say, a random number generator? How do you define a fidelity to that? Clearly we can’t speak of an assignment fidelity, since there are no “right answers.” Weak measurements and imperfect measurements raise the same basic question.

We can examine things even more exotic: What if your POVM outcomes are not orthogonal and you have more than two outcomes?

$$\{0.5|e\rangle\langle e|, 0.5|g\rangle\langle g|, 0.5I\} \quad (5.4)$$

is valid as well, which might have the physical interpretation of a  $Z$  detector which half of the time reports “measurement failed” or “no outcome observed.”

As an aside, POVMs with more than two outcomes occur more often than one might think. Two cascaded binary measurements can be thought of as a four-element POVM. If there are correlations

between the subsequent measurement processes, this may in fact be the *only* fair way to describe the measurements. Furthermore, even when working with independent measurements it can be useful to use non-binary POVMs. In the gate teleportation experiment we performed tomography with simultaneous readouts (section 7.8). It would be a huge waste of information if we discarded information in the correlations between results, in fact we would not be able to full do tomography of the joint multi-qubit state. Instead we do the reconstruction using a four element POVM with outcomes “gg,” “ge,” “eg,” and “ee.”

It generally important to know exactly what it is that a measurement is sensitive to. Perhaps a measurement can be improved with a unitary correction. Equivalently, perhaps you are learning about an undesired quantity. As an example in a quantum error correction setting, learning excess  $Z$  information about a single qubit during a stabilizer measurement looks like an outcome-dependent  $Z$  rotation on that qubit. This error is correlated with the previous outcome which could be of significant concern for fault-tolerance.

### 5.4.1 Quantum detector tomography

So how can we extract the full POVM of the measurement we have implemented? The basic process is called quantum detector tomography (Lundeen *et al.*, 2008), and is essentially identical to traditional state tomography (section 6.7). It differs only in the prior assumptions: rather than assuming we know the measurement operator, as we do in state tomography, we assume knowledge of the input state<sup>2</sup>.

We will briefly walk through this process for a two-outcome POVM, which is fully defined by one operator,  $\{E, I - E\}$ . Analogously to state tomography, we perform a complete or overcomplete set of “postrotations” and record the measurement outcome distribution, only now the “unknown state” is the ground state. With knowledge of the ground state  $|g\rangle\langle g|$ , the set of rotations  $\{R_j\}$ , and the set of observed probabilities,  $\{P_j\}$ , we can invert to find the measurement operator. The relationship between these is simple,  $P_j = \text{Tr} [R_j |g\rangle\langle g| R_j^\dagger E]$ .

<sup>2</sup>This knowledge is imperfect. In the stabilizer experiment our ground state preparation was better than 99%, and the one-qubit gate errors are less than 0.2% as determined from randomized benchmarking. Due to that, we would not be able to claim a level of precision about those errors. Ideas and techniques to go beyond those limitations are briefly presented in the conclusion of this chapter.




Figure 5.2: **Example of a POVM expressed in the Pauli basis.** This represents (ideally), a measurement of  $ZIZ$  on a three qubit register. This space requires a 64 dimensional measurement operator. The ideal operator is the projector onto the space where  $\langle ZIZ \rangle = 1$ , which can be decomposed as  $0.5III + 0.5ZIZ$ . We have come close to achieving that operator, but the reduced value of the ZIZ bar shows that our measurement has some noise. The lack of significant other bars implies that we are not measuring some undesired quantity.

We can then choose a basis in which to express the POVM. In this case we choose the Pauli basis, which is complete for Hermitian operators in the Hilbert spaces relevant to this discussion. This basis is often used to describe quantum states, as in chapter 6 and 7. We can now define a vector  $\vec{e}$  such that  $E = \sum_k e_k \sigma_k$ , and note that

$$P_j = \text{Tr} \left[ R_j |g\rangle\langle g| R_j^\dagger \sum_k e_k \sigma_k \right]$$

$$P_j = \sum_k e_k \text{Tr} \left[ R_j |g\rangle\langle g| R_j^\dagger \sigma_k \right]$$

$$\vec{P} \equiv A\vec{e},$$

where  $A_{jk} = \text{Tr} \left[ R_j |g\rangle\langle g| R_j^\dagger \sigma_k \right]$ . We have reduced solving for  $E$  to a quite reasonable matrix inversion. In practice, we use an overcomplete set as in state tomography, so we can use the pseudo-inverse to perform a least-squares fit. Alternately, we could perform a maximum-likelihood estimate constraining the operator to be positive (though not normalized). With this POVM in hand, we could go straight into a discussion of the detector fidelity, but we will put that off until section 5.5.2 and continue now with a more qualitative discussion.

We can present the results of such an inversion in the same manner that we present the results of state tomography. In the example in figure 5.2, we see two bars—one at the identity and one at  $ZIZ$ . This is what we expect as we know the eigenprojectors of a Pauli operator  $\sigma$  look like  $1/2(I \pm \sigma)$ . We can ask and answer several other useful questions from these data.

### 5.4.2 Specificity & coherent mis-measurement

Another question one might ask is: rather than precisely the quantity we wished to measure, what is the maximum information our measurement gains about *any* quantity? Or from another point of view: is our detector infidelity due to noise and simple lack of contrast, or is it because our detector is measuring the wrong quantity? Furthermore, is our detector *biased*—given a completely mixed input state, is one measurement outcome more likely than another? Certainly this is true in the case of dispersive measurements of transmons as energy relaxation naturally biases towards the ground state. This analysis is general to all strong binary measurements, but here we assume for simplicity that the ideal measurement is of a target Pauli operator,  $\sigma_T$ . I do not present a generalization to measurements with more than two outcomes, but that may be an interesting exercise.

We express  $E$  in the Pauli basis (which fully spans the space of  $n$ -qubit observables) as in figure 5.2, leading to a vector of Pauli coefficients. One dimension will correspond to the identity axis, one to the desired operator, and the rest to the various other Pauli operators. We can then find a new basis which rotates this vector space, leaving the identity and  $\sigma_T$  axes invariant, such that only one other non-zero coefficient in the measurement vector remains. This additional coefficient  $c_O$  corresponds to an orthogonal rotated Pauli  $\sigma_O$ ,

$$E = c_I I + c_T \sigma_T + c_O \sigma_O. \quad (5.5)$$

The coefficients of this expansion are easily interpreted: a deviation of  $c_I$  from the ideal value of 0.5 describes the bias of the detector,  $2c_T$  represents the maximum possible gain of information about the quantity we wish to measure, and  $2c_O$  represents the magnitude of the potential *undesired* information gain. Note that noise will inflate this quantity rather than average away, so we might consider it as a bound. Because  $E$  must be a positive matrix, there are constraints on  $c_T$  and  $c_O$  relative to the bias term  $c_I$ , e.g. a detector that always clicks cannot yield useful information. The infidelity corresponding to measuring along the wrong axis can be considered as analogous to coherent errors for standard processes, since a unitary rotation of the system prior to measurement would remove it. This analysis prescribes this correcting unitary exactly. In some cases the correction may be separable into single-qubit operations, but will not be true in general, particularly for architectures with high residual couplings.

This formulation leads naturally to a description in terms of angles between vectors, and we can quantify how well we are measuring along the correct axis with an angle for the “specificity” of the measurement,  $\theta_s \equiv \arctan(c_O/c_T)$ . Returning to the initial question of maximal information gain, we recast eq. 5.5 as  $E = c_I I + c_{\max} \sigma_{\max}$  where the operator  $\sigma_{\max}$  is defined as the axis where the measurement gains the most information. The coefficient  $c_{\max} = (c_T^2 + c_O^2)^{1/2}$  directly quantifies the total information gain along this axis, which quantifies the total amount of “coherent” information gain, though again noise will generally inflate this quantity.

## 5.5 Quantifying detector performance

This all motivates an obvious question: given an experimentally reconstructed POVM, what is its fidelity to some desired POVM? Or more generally, how similar are two POVMs? This is not always a straightforward question, especially when looking at some of the trickier examples given in the beginning of section 5.4. Standard fidelity measures derived for states and processes do not immediately work, as they require normalized states and all non-trivial POVMs inherently have non-normalized elements. We also cannot simply normalize the POVM elements, as there is significant information in their trace<sup>3</sup>. There is not yet a standard in the field for this quantity, but that is not to say there has been no work. (Magesan *et al.*, 2011) provide an detailed analysis for the case where the ideal measurement is rank-1. (Lundeen *et al.*, 2008) provides a similar fidelity measure, but does not allow for the POVM to have off-diagonal elements. (Dressel *et al.*, 2014) provides an exhaustive list of possible measurement fidelity definitions, each with various pros and cons. We decided to proceed in a slightly different direction which I found more intuitive and which also naturally opens the door to a characterization of the worst-case behavior.

### 5.5.1 Detectors as processes & the superoperator formalism

The first step in that derivation is to find a way to describe the entire POVM as something for which we *can* use standard fidelity measures. The insight here is to realize that while elements of a POVM are

---

<sup>3</sup>Specifically, it tells you the bias discussed in section 5.4.2

unnormalized, in some sense the complete POVM is as its elements sums to the identity. We do this by extending the standard POVM formalism, now describing the detector itself as a quantum channel (Wilde, 2009). It is an unusual channel. This “quantum-to-classical” channel takes a quantum state of the register  $\rho_r$  as its input then yields a diagonal density operator  $\rho_d$  with entries that represent the detector outcome probabilities,

$$\mathcal{E}_{\text{det}} : \rho_r \mapsto \rho_d = \sum_i \text{Tr}[E_i \rho_r] |i\rangle\langle i|_d. \quad (5.6)$$

One notes immediately that the dimension of the input and output states are not necessarily equal, quite at odds with a “normal” channel such as, e.g., a CNOT gate. This is an atypical formulation, but as explained in (Wilde, 2009) is quite valid. The output state of the detector  $\rho_d$  is represented as a density operator but it is important to understand that it is a container for a purely classical probability distribution. This distribution represents the recorded measurement outcomes, *not* the state of any physical ancilla. While a POVM element is not trace-1, the detector channel  $\mathcal{E}_{\text{det}}$  is a completely-positive trace-preserving map and we can safely use standard fidelity measure on it. We can express that channel with (generally) non-square Kraus operators. We will look at a few specific examples. A perfect  $Z$  measurement (outputs 0 and 1) of a qubit (states  $g$  and  $e$ ) has Kraus operators

$$|0_d\rangle\langle g_r|, |1_d\rangle\langle e_r|. \quad (5.7)$$

A single-qubit  $Z$  measurement with lower fidelity might be given by

$$(\sqrt{p}|1_d\rangle\langle g_r| + \sqrt{1-p}|0_d\rangle\langle g_r|), (\sqrt{p}|0_d\rangle\langle e_r| + \sqrt{1-p}|1_d\rangle\langle e_r|). \quad (5.8)$$

A parity ( $ZZ$ ) measurement of two qubits might have the Kraus operators

$$(|0_d\rangle\langle gg_r| + |0_d\rangle\langle ee_r|), (|1_d\rangle\langle eg_r| + |1_d\rangle\langle ge_r|). \quad (5.9)$$



With this sort of CPTP channel in hand, we can define fidelity measures as we like using the machinery of process fidelity.

### 5.5.2 Detector fidelities

In line with the reasoning in section 2.3.3, we do not present a single number for process fidelities, rather we think at least two are important. The first, *J-fidelity*, speaks to something roughly equal to the average case behavior. The second *S-fidelity*, speaks to the worst case behavior. Which is more relevant depends on the intended application.

The J fidelity of one detector process,  $\mathcal{E}_{\text{det}}^{(1)}$ , to another,  $\mathcal{E}_{\text{det}}^{(2)}$ , is given by

$$\mathcal{F}_{\text{J, det}} \left( \mathcal{E}_{\text{det}}^{(1)}, \mathcal{E}_{\text{det}}^{(2)} \right) \equiv \mathcal{F}_{\text{Tr}} \left( J^{(1)}, J^{(2)} \right), \quad (5.10)$$

where  $J^{(i)}$  is the Jamiołkowski matrix (section 2.3.3) representing the process  $\mathcal{E}_{\text{det}}^{(i)}$  and  $\mathcal{F}_{\text{Tr}}(\rho, \sigma) \equiv (\text{Tr} \sqrt{\rho^{1/2} \sigma \rho^{1/2}})^2$ . The S fidelity, or worst-case behavior, is given by

$$\mathcal{F}_{\text{S, det}} \left( \mathcal{E}_{\text{det}}^{(1)}, \mathcal{E}_{\text{det}}^{(2)} \right) \equiv \min_{\rho_{ra}} \mathcal{F}_{\text{Tr}} \left( \mathcal{E}_{\text{det}}^{(1)} \otimes I(\rho_{ra}), \mathcal{E}_{\text{det}}^{(2)} \otimes I(\rho_{ra}) \right),$$

where  $\rho_{ra}$  represents a joint pure state  $|\psi\rangle\langle\psi|_{ra}$  of the register and an ancillary, potentially entangled space, e.g. the rest of a quantum computer. This ancillary space is required to make certain that the fidelity measure is stable (section 2.3.3). It is important to emphasize that the states resulting from the process and compared on the right hand side are joint states of the detector and ancillary space, not of the joint register and detector space. In the special case of two-outcome POVMs, we have numerical evidence which leads us to believe that the state that reveals the worst-case performance will always be separable, indicating that  $\mathcal{F}_{\text{det}}$  is inherently stable. This fidelity can further be interpreted when we look at what the outcome of these processes actually are. As the diagonal elements are given by the

probabilities of the different outcomes,

$$\mathcal{F}_{S, \text{det}} = \min_{\rho} \text{Tr} \sqrt{\begin{bmatrix} p_0^{\rho,(1)} & 0 & 0 \\ 0 & p_1^{\rho,(1)} & 0 \\ 0 & 0 & \ddots \end{bmatrix} \begin{bmatrix} p_0^{\rho,(2)} & 0 & 0 \\ 0 & p_1^{\rho,(2)} & 0 \\ 0 & 0 & \ddots \end{bmatrix}}^2 \quad (5.11)$$

$$= \min_{\rho} \left( \sum_{i \in \text{outcomes}} \sqrt{p_i^{\rho,(1)} p_i^{\rho,(2)}} \right)^2, \quad (5.12)$$

which is equal precisely to finding which state minimizes the fidelity of two classical probability distributions (and reporting the square).

## 5.6 Quantifying measurement process performance

We have now characterized the behavior of the detector, but what happens to the input state after a measurement result is recorded? What is the back-action of registering a “click” or no “click”? Quantum error correction, for example, demands that the measurements must be highly quantum non-demolition in the sense of a von Neumann measurement. When an measurement indicates an outcome, the quantum process performed is ideally a projector onto the specified subspace. One method of analysis is to describe this process by two trace-non-preserving (but still completely positive) maps on the register Hilbert space,  $\{F^0, F^1\}$ . For a  $Z$  measurement of qubit where the result is “1,” the heralded process is ideally

$$F^{\text{“1”}}(\rho) = |1\rangle\langle 1| \rho |1\rangle\langle 1|, \quad (5.13)$$

which we generally then normalize to report a valid density matrix. It is also interesting to consider the unnormalized matrix. The trace of that matrix represents the probability of that measurement outcome having occurred. If we proceed to use this matrix for predictions of later measurements, the resultant calculations meaningfully give the joint probability of getting both the initial result *and* the second result.

Unfortunately our realized measurements will essentially never be true von Neumann projections. Consider the example of a photodiode which detects whether or not a photon is in a given mode. It

may detect very accurately, but in all cases it leaves the mode empty. The process that occurs when a photon is detected is then given by

$$F^{\text{"1"}}(\rho) = |0\rangle\langle 1| \rho |1\rangle\langle 0|. \quad (5.14)$$

This represents a relatively extreme form of undesired back-action. A more directly relevant case might be the dispersive measurement of a transmon. Generally, when measure dispersively with a JPC and obtain the result “g,” a second measurement is quite likely to yield “g,” often on the order of 99.5% – 99.8% . This is less true when we obtain the result “e,” as energy relaxation often allows the state to change as time passes. This means that our actual measurement process looks mostly like an ideal qubit measurement but also has some non-negligible similarity to a photodiode measurement.

### 5.6.1 From the measurement process to the detector

As an an side, we should note that the measurement process obviously contains the detector information. For a measurement outcome  $\mu$ , if we express the measurement process in Kraus form, we see

$$\rho^\mu \mapsto \sum_i M_i^\mu \rho M_i^{\mu,\dagger} \quad (5.15)$$

$$P(\mu) = \text{Tr} \left[ \sum_i M_i^\mu \rho M_i^{\mu,\dagger} \right] = \text{Tr} \left[ \sum_i M_i^{\mu,\dagger} M_i^\mu \rho \right] \quad (5.16)$$

$$E^\mu = \sum_i M_i^{\mu,\dagger} M_i^\mu. \quad (5.17)$$

### 5.6.2 Measurement process tomography

So given the ability to perform some measurement, how do we characterize it and extract the measurement process? The answer we explore here is a special form of quantum process tomography (section 6.7). Tomography of a measurement has two key distinctions. One, a classical measurement result is acquired, and we bin the output states and analyses by this classical bit. Two, we weight the output states by the probability they occurred, rather than keeping normalized density matrices. This second adjustment contains all of the information about the detector, and is essentially the information acquired




Figure 5.3: **Measurement process tomography.** As in regular process tomography, we begin by preparing a complete or overcomplete set of input states, then applying our process. A measurement process also provides a classical bit, and we bin the state tomography of the resulting states by that classical bit. For each input state we also record the probability of getting the various outcomes. Finally, the reconstructed outcome states are weighted by the probability that they occurred. The resultant inversion between the output and input states will lead to the individual maps being non-trace-preserving.

in quantum detector tomography. It is this rescaling that allows for the subprocesses to be individually non-trace preserving. However they are collectively trace-preserving since one of the outcomes always happens. Other than this reweighting, the analysis is performed the same as in standard process tomography. An example of the results of measurement process tomography is given in figure 6.12. Note of course that if maximum-likelihood estimation is to be applied in the reconstruction we cannot constrain the separate maps to be trace-preserving, only positive.

If we wish to report a process fidelity we run in to similar problems as we did with POVMs. Since the subprocesses are not trace-preserving, a naive use of a fidelity measure will not work. In the next section we will discuss an alternative formalism that allows us to constrain the two subprocesses to be collectively trace-preserving.

### 5.6.3 Quantum instruments

Fortunately we can compactly describe a complete measurement process with a CPTP map using the *quantum instrument* formalism (Wilde, 2009). In this formalism we represent the measurement outcome as a classical state in an ancillary Hilbert space, the same way we did with detector processes, but now we also carry along the resultant quantum state of the register. For a two-outcome process, we might write

$$\mathcal{E}_{\text{QI}} : \rho_r \mapsto F^{\text{"0"}}(\rho_r) \otimes |0\rangle\langle 0|_d + F^{\text{"1"}}(\rho_r) \otimes |1\rangle\langle 1|_d, \quad (5.18)$$

where  $F^{“0”}$  and  $F^{“1”}$  represent the processes on the register subspace, heralded by measurement outcome. For a perfect von Neumann measurement, these processes would simply be projectors onto the measurement outcomes. This generalizes—for an  $m$ -outcome measurement on a  $d$  dimensional register Hilbert space, this process maps  $d \times d$  density matrices to  $(md) \times (md)$  density matrices,

$$\mathcal{E}_{\text{QI}} : \rho_r \mapsto \sum_{i=1}^m F^{“i”}(\rho_r) \otimes |i\rangle\langle i|_d. \quad (5.19)$$

#### 5.6.4 “Quantum non-demolition”

The previous example measurement would generally be called *quantum non-demolition* or QND. One early paper (Braginsky *et al.*, 1980), gave a concise definition, developed before the advent of quantum information as a field.

We define a QND measurement of  $\hat{A}$  as a sequence of precise measurements of  $\hat{A}$  such that the result of each measurement is completely predictable from the result of the first measurement—plus, perhaps, other information about the initial state of the system.

As the sport of manipulating quantum systems has become more sophisticated so must our requirements. The above definition is not strict enough for measurements with degenerate eigenspaces. Pointedly, this definition is necessary but not sufficient for quantum error correction. Correcting errors is contingent on removing entropy by projecting our density matrix towards a known point in error space, and measurements under this definition do not promise that. A measurement of  $ZZZZ = +1$  could, in a pathological example, always project the system to the state  $|0000\rangle$ , and still be considered QND. In a less contrived example, a measurement device may project the register into the correct subspace but then dephase within it; e.g. a  $ZZ$  measurement which, rather than creating a Bell pair, creates  $|00\rangle\langle 00| + |11\rangle\langle 11|$ , or a measurement of photon number parity which might dephase (or Kerr) a cavity. It is a much better definition to require that the heralded measurement process be the the expected projection.

Quantum-non demolition clearly does not apply to the case of *weak* measurements, where we may want the meter to become less than fully entangled with the system. In that case we should ask that the back-action be “minimally disturbing” in the sense of (Wiseman and Milburn, 2009). In this case the

back-action is limited to a diagonal partial projection, disallowing any further mixing or unitary action. We will seek a general definition for the fidelity of the measurement process that covers both strong and weak measurements, as well as ones where the ideal measurement process (for whatever reason) might have some other unitary or incoherent back-action.

### 5.6.5 A superoperator formalism for measurement processes

One last bit of formalism we will need, directly in analogy to section 5.5.1, is a concise formalism for the CPTP map of a quantum instrument. The map given by equation 5.19 can just as easily be written as an  $m^2 d^2$  by  $d^2$  superoperator,

$$\text{vec}(\rho_{rd}) = \tilde{T} \cdot \text{vec}(\rho_r). \tag{5.20}$$

Or in a more visual form for a two-outcome measurement,

$$\begin{array}{c}
 |0\rangle\langle 0|_d : \\
 \vdots \\
 d^2 \\
 \hline
 0 \\
 \vdots \\
 |0\rangle\langle 1|_d : \\
 \vdots \\
 0 \\
 \hline
 0 \\
 \vdots \\
 |1\rangle\langle 0|_d : \\
 \vdots \\
 0 \\
 \hline
 \vdots \\
 |1\rangle\langle 1|_d : \\
 d^2 \\
 \vdots \\
 \vdots
 \end{array}
 =
 \begin{array}{c}
 F^{“0”} : \\
 \begin{array}{c}
 \begin{array}{c} \ddots \end{array} \\
 d^2 \times d^2 \\
 \begin{array}{c} \ddots \end{array} \\
 \hline
 \begin{array}{c} \ddots \end{array} \\
 0 \\
 \begin{array}{c} \ddots \end{array} \\
 \hline
 \begin{array}{c} \ddots \end{array} \\
 0 \\
 \begin{array}{c} \ddots \end{array} \\
 \hline
 \begin{array}{c} \ddots \end{array} \\
 d^2 \times d^2 \\
 \begin{array}{c} \ddots \end{array}
 \end{array}
 \end{array}
 \begin{array}{c}
 \begin{array}{c} \vdots \\ \vdots \\ \vdots \end{array} \\
 d^2 \\
 \begin{array}{c} \vdots \\ \vdots \\ \vdots \end{array}
 \end{array}
 . \tag{5.21}$$

### 5.6.6 Quantum instrument fidelity

We can now take this CPTP superoperator and apply standard fidelity measures to it. The reasoning in this section essentially follows from when we defined detector fidelities in section 5.5.2. Again rather than a single number we give both the J-fidelity and the worst-case S-fidelity. The J-fidelity is a more commonly cited number, while the S-fidelity is more likely to be relevant to fault-tolerance calculations.

The J fidelity of one detector process,  $\mathcal{E}_{\text{QI}}^{(1)}$ , to another,  $\mathcal{E}_{\text{QI}}^{(2)}$ , is given by

$$\mathcal{F}_{\text{J, QI}}\left(\mathcal{E}_{\text{QI}}^{(1)}, \mathcal{E}_{\text{QI}}^{(2)}\right) \equiv \mathcal{F}_{\text{Tr}}\left(J^{(1)}, J^{(2)}\right), \quad (5.22)$$

where again  $J^{(i)}$  is the Jamiołkowski matrix (section 2.3.3) representing the process  $\mathcal{E}_{\text{QI}}^{(i)}$  and  $\mathcal{F}_{\text{Tr}}(\rho, \sigma) \equiv (\text{Tr} \sqrt{\rho^{1/2} \sigma \rho^{1/2}})^2$ . The S fidelity, worst-case behavior, is given by

$$\mathcal{F}_{\text{S, QI}}\left(\mathcal{E}_{\text{QI}}^{(1)}, \mathcal{E}_{\text{QI}}^{(2)}\right) \equiv \min_{\rho_{ra}} \mathcal{F}_{\text{Tr}}\left(\mathcal{E}_{\text{QI}}^{(1)} \otimes I(\rho_{ra}), \mathcal{E}_{\text{QI}}^{(2)} \otimes I(\rho_{ra})\right).$$

Unlike the detector S-fidelity, we have numerical proof than the quantum instrument S-fidelity is in fact not stable<sup>4</sup>, and optimizing over an auxiliary space is quite important.

## 5.7 Conclusions

This chapter has provided views on several ways to look at and analyze measurements, but it seems likely that even more sophisticated techniques will be needed. Quantum detector and quantum measurement process tomography suffer state preparation errors just as standard state and process tomography do. Techniques such as gate set tomography (Blume-Kohout *et al.*, 2013) avoid these errors, but come at an incredibly expensive cost in experimental overhead. Furthermore, at least as originally described, gate set tomography is only appropriate to detectors. I am not sure if the technique can be applied to probabilistic but heralded “gates.”

There are various other proposals for “self-consistent” tomography, which should in principal serve to accurately calibrate the measurement operator (Jackson and van Enk, 2015). To my knowledge none

<sup>4</sup>This is unsurprising as processes are generally not stable and we allow arbitrary measurement processes.

of these are directly suited to characterizing the measurement process, and all of them seem to suffer from expensive experimental overhead. To quote the author of one such technique (Merkel *et al.*, 2013),

Colloquially, we have taken to calling [self-consistent tomography] the “overkill” method of tomography, since the number of the experiments borders on the ridiculous for all but one- or two-qubit systems[.]

On the other side of the spectrum, it is not clear that it is possible to extend techniques such as randomized benchmarking, and in particular simultaneous randomized benchmarking, to measurement processes. Certainly there is no currently known mechanism to do so, but if possible it will likely prove a more practical experimental technique.



## Chapter 6

# The Nigg protocol & stabilizer measurements

It looked as if not one ball was rolling across the table but a great number of balls, all partially penetrating into each other.

---

Quantum Billiards, George Gamow

The first large project in my thesis work was the implementation of a proposal in 2013 by Simon Nigg and Steve Girvin, *Stabilizer Quantum Error Correction Toolbox for Superconducting Qubits* (Nigg and Girvin, 2013). This work described a protocol that used a high-Q bus resonator to enact gates between transmons, and in particular gates of the class necessary for the stabilizer measurements discussed in section 2.5. The context of this proposal was shortly after the demonstration of the qcMAP gate (Leghtas *et al.*, 2013a). We will cover qcMAP thoroughly in the next section as it can teach us many things directly relevant to this work. The key idea is a way to exchange quantum information between a transmon and the state of a resonator. Nigg and Girvin looked at what would happen if the system had multiple transmons. After all, if we can exchange information between transmon one and the bus and vice versa, and the same with transmon two, it seems reasonable that we might be able to generate a useful interaction between transmons one and two. What they found was not only that this was true, but that the operation had nice parallelism properties and that the nature of the interaction was highly

controllable.

To place this protocol in a greater context we can look at the many ways to generate interactions between transmons. For tunable-frequency interactions one can perform direct resonant or nearly-resonant dipole-dipole interactions via capacitive couplings as in (Barends *et al.*, 2014, Chow, 2010). Without the ability to shift frequencies these interactions cannot be turned off, a highly undesirable property. One other possibility with tunable-frequency transmons is to use a sequence SWAP interactions with a common resonator mode (Riste *et al.*, 2015). Unfortunately, tunable-frequency transmons suffer from rapid dephasing due to flux noise, and also lead to a messier microwave environment as their resonance frequencies shift.

If we restrict ourselves to off-resonant and fixed-frequency transmons there are several ways to generate interactions. The first, the cross-resonance gate (Chow *et al.*, 2011), is fast and has been shown to exhibit very high fidelity gates, but unfortunately suffers from spectral complications in environments with larger numbers of transmons. To my knowledge these issues have not yet been resolved. More unfortunately, the cross-resonance gate does require an always-on direct interaction which is a term from the first-order rotating wave approximation, though it operates much more quickly. This does not bode well for the possible on/off ratio.

If we allow the transmons to share a common bus resonator, we are additionally allowed the resonator-induced phase (RIP) gate and the Nigg protocol. The RIP gate was originally developed by Steve Girvin and his postdoc Dong Zhou and implemented in the Schoelkopf group by Hanhee Paik. It was then later refined and published in (Cross and Gambetta, 2015, Paik *et al.*, 2016). It will be covered in detail in a later chapter on gate teleportation (section 7.6), and is also quite fast and high-fidelity. Both of these gates do not depend on a direct qubit-qubit interaction, and such a coupling is only required to arise under fourth order perturbation theory.

This chapter will begin with a discussion of the relevant concepts to understand the original proposal, starting with qcMAP in section 6.1. In section 6.2 we will get into the meat of the protocol. Sections 6.2.1 and 6.2.2 will get into making the pulse sequences more complicated. The first of those sections serves to expand the gates functionality, allowing a richer class of operations, and the second makes the protocol more experimentally viable, using software to compensate for non-identical hardware. Section

6.2.3 will cover our first modification to the original proposal, radically changing the second half of the pulse sequence. Section 6.2.4 will cover another modification which can allow faster gate times, building directly on a trick developed for qcMAP.

Section 6.3 will pertain to the characterization of the device. This section serves very generally as a record for my recipe for how to measure devices in our regime of circuit QED. This section will end in the Hamiltonian and coherent details of the device.

Sections 6.4 to 6.7 will describe various parts of the experimental pulse sequence. Section 6.4 goes into the sequence of pulses and measurements used to post-selectively prepare the system in the ground state. Section 6.5 more thoroughly fleshes out the measurement sequences which we described abstractly in section 6.2. In section 6.6 we look at a cute “trick” which optionally allows us to herald certain errors in the measurement protocol. Section 6.7 discusses state and process tomography, both generally and as they specifically relate to this experiment.

We end by discussing our performance and the results of these characterization experiments in section 6.8. That section covers the most important results. More exhaustive results are provided in appendix A.

## 6.1 qcMAP: a useful precursor

As a preliminary to discussing the Nigg proposal, it will be enlightening to first discuss the gate that inspired it, qcMAP (Leghtas *et al.*, 2013a). Additionally, we will use the simpler context of the qcMAP gate to examine errors and non-idealities that will arise in the Nigg protocol. The function of the qcMAP gate is to take a qubit of quantum information stored in a transmon state and transfer it into a “two-cat” qubit encoded in a harmonic oscillator. In this encoded qubit the  $|e\rangle$  and  $|g\rangle$  states correspond to coherent states of opposite phase in an oscillator. If the amplitude of the coherent states is given as  $|\alpha|$ , the action of this gate can be written

$$(c_e |e\rangle_{\text{transmon}} + c_g |g\rangle_{\text{t}}) \otimes |0\rangle_{\text{cavity}} \leftrightarrow |g\rangle_{\text{t}} \otimes (c_e |\alpha\rangle_{\text{c}} + c_g |-\alpha\rangle_{\text{c}}) \quad (6.1)$$

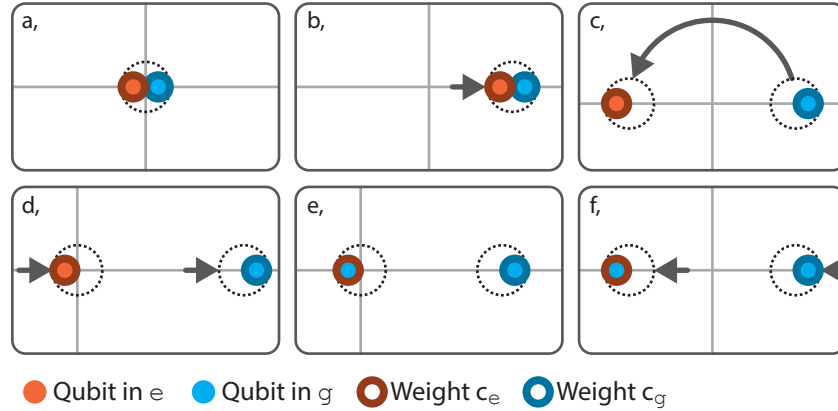


Figure 6.1: **The qcMAP protocol.** This diagram represents the phase-space distribution of a cavity, or more technically the Husimi-Q function (section 2.2.6), where the dotted circle indicates a Gaussian. The color (and position) of each small circle within the dotted circle depict the correlated state of the transmon. The color of the edge of each of these circles represents the quantum amplitude associated with each state. As we walk through the steps of the qcMAP gate an initial arbitrary state of the transmon qubit is swapped into the phases of coherent states in the cavity.

The gate is essentially an encoding and decoding operation, swapping information back and forth between a transmon state and cat state<sup>1</sup>. It is relatively unusual because it maps to a non-trivial logical encoding. This encoding and its even more interesting cousin the “four-cat” (Leghtas *et al.*, 2013b) are highly important to the research direction of a large fraction of the Schoelkopf lab, and similar states are used in the gate teleportation experiment (chapter 7).

We will now walk through the qcMAP gate. Though the gate is bidirectional, we will go through the example of generating a specific encoded qubit state in the cavity from a system initially in its ground state. Our Hamiltonian consists of a single transmon and a single linear cavity which share a dispersive interaction. We work in a rotating frame such that the cavity is stationary when the transmon is in the ground state and write the Hamiltonian as  $\mathcal{H}/ = \chi a^\dagger a |e\rangle \langle e|_t$ .

We can walk through the steps, following along with figure 6.1.

- a. The desired qubit state is prepared in the transmon.
- b. The cavity is quickly displaced with a spectrally broad pulse to a coherent state with amplitude  $\alpha$ .
- c. We wait for a time  $t = \pi/\chi$  such that if the qubit is excited the detuned cavity will gain phase

<sup>1</sup>It is *not* a true SWAP gate. It will not take  $(c_0 |g\rangle + c_1 |e\rangle)(c_2 |\alpha\rangle + c_3 |-\alpha\rangle)$  to  $(c_2 |g\rangle + c_3 |e\rangle)(c_0 |\alpha\rangle + c_1 |-\alpha\rangle)$ .

$$\phi = \chi t = \pi.$$

- d. We repeat the initial displacement of the cavity, such that the now  $|e\rangle |-\alpha\rangle$  state becomes  $|e\rangle |0\rangle$ .
- e. We perform an  $X_{\pi}^0$  rotation on the transmon with a spectrally narrow pulse, such that the qubit is affected if and only if the cavity is in the vacuum.
- f. We invert the initial displacement.

For completeness we can also walk through the joint state of the system.

- a.  $|\psi\rangle = (c_e |e\rangle + c_g |g\rangle) |0\rangle$
- b.  $\rightarrow (c_e |e\rangle + c_g |g\rangle) |\alpha\rangle$
- c.  $\rightarrow c_e |e\rangle |-\alpha\rangle + c_g |g\rangle |\alpha\rangle$
- d.  $\rightarrow c_e |e\rangle |0\rangle + c_g |g\rangle |2\alpha\rangle$
- e.  $\rightarrow |g\rangle (c_e |0\rangle + c_g |2\alpha\rangle)$
- f.  $\rightarrow |g\rangle (c_e |-\alpha\rangle + c_g |\alpha\rangle)$

### 6.1.1 Effects of decoherence

Involving a cavity state opens the door to several new forms of decoherence. Most obviously, this means photon loss in the cavity generated by the annihilation operator  $\sqrt{\kappa}a$ . Photon loss of a cavity whose state contains non-trivial quantum information is equivalent to measurement-induced dephasing, regardless of whether or not measurement is the desired affect. Each photon which leaves the cavity potentially carries information, leading to mixing (dephasing) between either encoded qubits or an entangled physical qubit. The amount of information carried by a photon leaking out from a superposition of coherent states  $\alpha$  and  $\beta$  is given by the square of the translational distance in phase space  $\Delta = |\alpha - \beta|$ . This leads to a time-dependent dephasing rate  $\Gamma_{MID}(t) = \Delta^2(t)\kappa/2$  (Haroche and Raimond, 2012).

We can look at several examples. If both qubit states represent the same coherent state ( $\Delta = 0$ ) then a lost photon has no meaningful quantum information. More generally, if both states have the same average photon number but with a phase difference  $\theta$ , the information loss scales with that angle as  $1 - \cos\theta$ . If the average photon number in a resonator is  $\bar{n}$ , it has a proportional probability to lose

a single photon, and accordingly the dephasing rate is also enhanced by  $\bar{n}$ . In this manner  $\Delta$  alone is enough to describe the dephasing rate of a cat state, independent of the displacement of the whole state, as we saw in the form of  $\Gamma_{MID}$ .

Additionally, the deterministic effect of unobserved photon-loss is that the coherent states move steadily towards the vacuum at rate  $\exp(-\kappa t/2)^2$ . When the distance between the two coherent states  $|\Delta|$  is not sufficiently large, this also leads to a subspace error, where our  $|e\rangle$  and  $|g\rangle$  states are no longer orthogonal. In analogy to qubit states, this looks something like a depolarizing channel that sets in with non-trivial time dependence. This effect is suppressed exponentially in  $\Delta$ .

Relaxation of the transmon now has an additional affect if it occurs during the phase accumulation step of qcMAP. When the transmon jump occurs at a random time during this phase accumulation, it is easy to see that in addition to the dephasing effect of the transmon information leaking away, the cavity state previously entangled with  $|e\rangle$  will no longer be in its proper location. This leads to a complicated error–leakage out of the encoded manifold. As we will see when we turn to the analogous affect on the Nigg protocol, errors of this type can often be heralded.

### 6.1.2 Limitations of real pulses

Another set of concerns is the finite duration of the applied pulses, as displacements do not commute with the dispersive Hamiltonian. When we continue to the Nigg protocol we will see similar effects due to the finite duration of unselective  $\pi$  pulses.

The finite duration of displacements proved not to be a significant problem in the stabilizer experiment, but it does need to be accounted for, as it does in qcMAP. Incorporating the duration is quite important when calculating the phase angle for the required displacements. The magnitude of this effect depends on the starting and ending locations of the coherent state, as a state in vacuum does not experience a rotation. Accordingly smaller finite-duration corrections are required for states that start or end in the vacuum relative to displacements between two non-vacuum states. If the spectral widths of the displacements are not sufficiently large compared to  $\chi$  (or the sum of  $\chi$  in a multi-qubit experiment), one would also see that the displacement magnitude is dependent on the qubit state.

---

<sup>2</sup>This factor of two appears in the phase-space description as  $\kappa$  is defined as the energy-loss rate.

Another consideration is what duration pulse to use for the photon-number selective  $\pi$  pulse. Clearly a pulse that is selective on the cavity having zero photons is sufficient, but in general it is not necessary. As allowing a broader spectrum leads to a faster operation, it is worthwhile to ask how fast we may go. We will examine this case, with one coherent state in the vacuum and the other at an unchanging displacement amplitude  $2\alpha$ . We will also restrict ourselves to using Gaussian pulses, and perform a rough calculation of the worst-case behavior.

For a pulse resonant with the cavity in the ground state, we expect a full population inversion of the transmon from  $|e\rangle$  to  $|g\rangle$ . For a transmon at  $n$  photons, the drive is then centered off-resonance by  $\delta = \chi n$ . The envelope of the excitation under such a drive, if it were constant in time, is given by  $c_e \sim \Omega / \sqrt{\Omega^2 + (\chi n)^2}$ , where  $\Omega$  is the resonant Rabi rate. As we are looking at a relatively slow pulse, we can approximate for the off-resonant case that  $c_e \sim \Omega / \chi n$ . The resonant  $\pi$  pulse then has area  $\int \Omega(t) dt = \pi$ . For Gaussian pulses, this implies a peak Rabi rate  $\Omega_0 = \sqrt{\pi/2} / \sigma_T$ , where  $\sigma_T$  represents the temporal width of the envelope. The worst case excitation for a detuned peak is

$$P_{e|n} \sim \frac{\pi}{2} \frac{1}{(\chi n \sigma_t)^2}, \quad (6.2)$$

where the prefactor is dependent on our assumptions but the scaling should hold. If the non-vacuum coherent state is at  $2\alpha$ , we can look at the behavior at  $n = 4\bar{n}$ , and find a pulse width of

$$\sigma_T \sim \sqrt{\frac{\pi}{2}} \frac{1}{2\chi\bar{n}} \sqrt{\frac{1}{P_e}}, \quad (6.3)$$

for a cat of final  $\bar{n}$  and acceptable error probability  $P_e$ . A more rigorous calculation would look at the Poissonian distribution of the state  $2\alpha$ . A similar calculation in the opposite regime would analyze the infidelity of a finite-duration “unselective” pulse.

### 6.1.3 Higher order terms in $\mathcal{H}$

The higher-order terms resulting from the cosine potential, as discussed in section 3.5, form another significant set of concerns. Two terms dominate these effects in our regime:

$$\mathcal{H}_{HO} = \frac{-K}{2} a^\dagger a^\dagger a a + \frac{\chi'}{2} q^\dagger q a^\dagger a^\dagger a a. \quad (6.4)$$

The first is the Kerr non-linearity of the resonator and the second is the qubit-state-dependent Kerr non-linearity. The first-order effect of Kerr is a state-dependent detuning of  $K (\bar{n} + 1/2)$ , and at higher orders it begins to essentially dephase the coherent state, “smearing” it in phase-space<sup>3</sup>. At time-scales comparable to  $K^{-1}$  the states can re-phase in interesting ways (Kirchmair *et al.*, 2013), though we will not be working in that regime. In the  $t \ll 2\pi/K$  regime, the maximal overlap of a “lightly Kerred” state with *any* coherent state scales as  $\exp(-2/\pi\bar{n}K^2t^2)$  (Nigg and Girvin, 2013).  $\chi'$  will cause the same problem, only in an entangling, qubit-dependent manner and generally with the opposite sign. Perhaps counter-intuitively, in practice these terms have similar magnitudes despite the fact that  $\chi'$  results from a higher order in the cosine expansion.

### 6.1.4 Beating $\chi$

One modification that will be important later is to note that while qcMAP originally asked for a phase accumulation time of  $\pi/\chi$ , its goals can be achieved in a shorter time. This is apparent when you consider that the angular evolution in phase space is the constant, rather than the rate of change of the linear distance. That is to say that at larger displacements, the coherent states become distinguishable more quickly. The coherent states separate in time as

$$\begin{aligned} \Delta &= |\beta - \beta \exp\{1j\chi t\}|^2 \\ &= |\beta|^2 (1 - \cos \chi t)^2 + \sin^2 \chi t \\ &= 2|\beta|^2 (1 - \cos \chi t). \end{aligned}$$

<sup>3</sup>Of course this is coherent, not actually dephasing, so we might hope to correct it using e.g. the optimal control theory of section 4.3.2.






Figure 6.2: qcMAP– **faster**. Under the dispersive interaction detuned coherent states separate at a constant *angular* velocity. An absolute separation in phase space can be obtained more quickly by using larger initial photon numbers.

A potentially concerning downside of this enhancement is increased exposure to Kerr and  $\chi'$ , whose negative effects scale as  $|\beta|^2 K^2 t^2$  as previously discussed. If we look for at the extreme small angle case  $\Delta = |\beta|^2 \chi^2 t^2$ , it appears that the Kerr exposure is actually invariant<sup>4</sup>. This does not speak to the even higher order effects, nor the fact that all timing and phase errors will now have potentially amplified effects.

## 6.2 The Nigg proposal

We can now look at the original Nigg and Girvin proposal (Nigg and Girvin, 2013). The central insight of their paper is this: if  $N$  qubits (a register) share an equal dispersive shift with a high-Q cavity, the cavity can acquire phase in such a manner that the effect of the individual qubits is indistinguishable. The qubits become entangled with the cavity in a way that exhibits character not generally identifiable with a single qubit– rather it knows only the total *number* of excitations, not where they come from. A coherent state will obtain an entangled phase  $\phi_{\text{cav}} = \chi T \hat{n}$ . Furthermore, since phase has periodic boundary conditions, various different number states of the register can also become stroboscopically indistinguishable. This can be considered an erasure of “which-path” information. In particular, if the evolution time  $T$  is chosen to be equal to  $\pi/\chi$ , then the phase  $\phi_{\text{cav}} = \pi \hat{n}$  and  $\phi_{\text{cav}} \bmod 2\pi$  corresponds

<sup>4</sup>In a single-transmon system we can safely approximate  $K \approx \chi^2/\alpha$ , so we can directly predict this ratio.

directly to the parity of the qubit register<sup>5</sup>. Somewhat amazingly, this works in exactly the same time  $\pi/\chi$  *regardless of the number of qubits*. More conventional protocols would use a sequence of pair-wise interactions, leading to an overhead which scales (at least) linearly with the number of qubits. An example of this applied to two qubits can be seen in the first three panels of figure 6.5.

The original proposal builds a larger framework around this insight. Most obviously, in order to perform useful measurement we must find a way to non-destructively move this information from the phase of a coherent state to our classical control apparatus. The Nigg proposal provides two ideas for how to do this, but the one we focus on uses an extra ancilla qubit, which itself couples to an additional low-Q “readout” resonator. This scheme has essentially three parts. First, a fast displacement unconditionally translates the cavity, such that one of the coherent states becomes the vacuum. This changes the preexisting phase information to number-state information. Second, we can then apply a spectrally narrow pulse to invert the ancilla state if and only if there are zero photons in the cavity. This yields our information, but we must reset the cavity to the vacuum to be non-destructive. We do so by performing a spectrally narrow pulse to the cavity that displaces the (non-vacuum) coherent state to the vacuum if and only if the ancilla is in the ground state. As written, this requires the ancilla have a dispersive shift such that  $\chi_{\text{ancilla}} \gg N\chi$ , which is not terribly practical. Along with reasons we will discuss later, it would also almost certainly require existing in a regime of very large cavity self-Kerr. However if it did work, the resulting state would be the cavity unconditionally in the vacuum and the ancilla cleanly entangled with the register parity. At this point we could measure the ancilla via the readout resonator and non-destructively learn the register parity  $Z^{\otimes N}$ .

As we will discuss next in further detail, we can also let the high-Q cavity become entangled with operators other than  $Z^{\otimes N}$ . These echo sequences make available all operators which are tensor products of  $Z$  and  $I$ . Along with single-qubit unitary prerotations, this set is sufficient to measure all possible product operators. The proposal also provides a prescription for modifying the protocol when all dispersive shifts are not precisely equal—crucial for real-work applicability.

---

<sup>5</sup>Note that in the original citation  $\chi$  is differently defined. In that work  $2\chi$  is equal to the full dispersive splitting. We use the convention that  $\chi$  is equal to the full dispersive splitting




Figure 6.3: **Schematic of the device proposed by Nigg and Girvin.** In (Nigg and Girvin, 2013) the authors envision a high-Q cavity bus containing a register of 3D transmon qubits. An additional ancilla qubit would couple to both that cavity and a low-Q readout resonator. Photons in the high-Q cavity would become entangled with the transmon register in an interesting manner, which could then become entangled with the ancilla which was later measured via the readout resonator. Image reproduced and modified from (Nigg and Girvin, 2013) with permission.

### 6.2.1 Subset selectivity

While measuring  $Z^{\otimes N}$  is an interesting endeavor, the protocol allows for a richer space of possible measurement. By using echo pulses, we are able to on-demand erase the phase contributions of individual qubits within the register. In some sense this is akin to a Hahn echo. In that sort of experiment, a single qubit in an unknown state is allowed to evolve for a time  $T$ , its state is quickly inverted, and allowed to evolve for another time period  $T$ . Regardless of the original state of the qubit, its probability amplitude initially associated with state  $|1\rangle$  has gained an equal amount of phase as the amplitude initially associated with  $|0\rangle$ . We are not able to directly  $\pi$  pulse (or invert the population of) the encoded cavity state<sup>6</sup>, but we can still invert the state of each qubit. By exactly the same reasoning we can see that the qubit should remain disentangled, leading to an averaged and unconditional phase shift on the cavity.

Informally we can reduce the cavity evolution to one dimension and evaluate the cavity coherent state “phase propagator” due to each qubit  $i$ . This refers the phase evolution of the cavity to the transmon state at the beginning of the sequence. We also work in a frame where the cavity is resonant if all qubits are in the ground state. The phase acquired over a time  $\tau$  with a qubit  $X_\pi$  pulse at time  $t_{\text{echo},i}$  is given by

$$\phi_i = |e\rangle\langle e|_i \chi_i(\tau - t_{\text{echo},i}) + |g\rangle\langle g|_i \chi_i t_{\text{echo},i}. \quad (6.5)$$

If we choose the echo time such that  $t_{\text{echo},i}$  is  $\tau/2$ , this reduces to  $\chi_i\tau/2$  times the identity, and the cavity phase gains no qubit information. Since the larger process consists of only two-body interactions, this can safely be considered on a qubit-by-qubit basis. Access to measurement of all possible “subsets,” paired with universal single-qubit control, then lets us measure any possible product operator. As constructed here we have implied a single echo pulse, which would leave the qubit state inverted. There is no impediment to using two (or more) echo-pulses, as long as the time split between “echoed” and “non-echoed” configurations is equal. In practice we do use a pair of pulses. Figure 6.4 depicts the cavity evolution with a single echo of one qubit in a two-qubit register. An example echo pulse sequence can be seen applied to the second qubit in part c of figure 6.9.

<sup>6</sup>This is not strictly true, we certainly could with pulses generated via optimal control theory (section 4.3.2), but it would not be practical in this protocol.

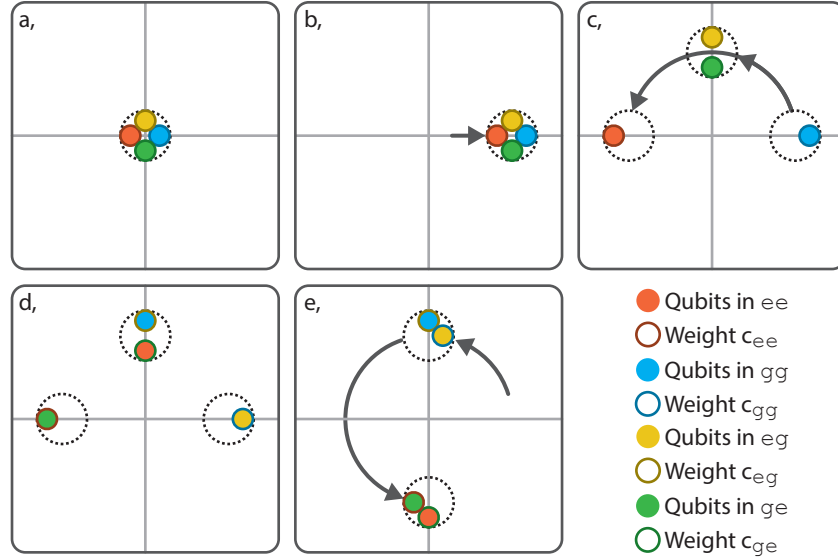


Figure 6.4: **Initiating a measurement of  $ZI$ .** After beginning in an arbitrary initial two-qubit register state, we displace the cavity in phase space. It is allowed to evolve under the dispersive interaction for a time  $T = \pi/2\chi$  before we quickly apply an  $X_\pi$  gate to the second qubit. The cavity is allowed to evolve for an equal time  $T$ . At the end of this evolution we can see that the phases of the coherent states in the cavity are entangled only with information about the first qubit.

### 6.2.2 Equalizing $\chi$

One takeaway from equation 6.5 is that by using different  $t_{\text{echo},i}$ , we can control the conditional phase contributed by each qubit. More precisely—we can modify that phase to be a chosen quantity smaller than  $\chi_i\tau$ . If we consider the dispersive shifts of the register by magnitude, we can choose  $\tau$  such that the “slowest” qubit we plan to measure contributes  $\pi$  phase shift:  $\tau = \pi/\chi_{\min}$ . For each qubit we do not plan to measure we apply a  $\pi$  pulse at  $\tau/2$ . For each qubit we do plan to measure (other than the slowest qubit), we apply a  $\pi$  pulse at  $t_{\text{echo},i} = \pi/2(1/\chi_{\min} - 1/\chi_i)$ . This results in the phase propagator

$$\phi_i = |e\rangle\langle e|_i \pi + I_i \chi_i \left( \frac{\pi}{2\chi_{\min}} - \frac{\pi}{\chi_i} \right), \quad (6.6)$$

which is the result we expect—a conditional phase accumulation of  $\pi$  and an unconditional phase contribution related to the difference of  $\chi$ s. This “partial echo” technique can be seen applied to the third qubit in part c of figure 6.9. Note dispersive shifts which are too similar can actually present a

problem. Since our  $\pi$  pulses have finite duration sufficiently small  $t_{\text{echo}}$  may not be feasible. This can be remedied by echoing both qubits, even if it would otherwise not be required, which then relies only on adjusting the relative timings of the pulses.

### 6.2.3 A different erasure technique

Unfortunately the very large ratio of dispersive shifts in the original proposal is not practical. We must choose the slower (register) dispersive shifts to be sufficiently large that the protocol is relatively fast compared to their decoherence times. If we choose an ancilla dispersive shift much larger than that, we will obtain a large cavity self-Kerr term, which scales roughly as  $\chi^2/\alpha$ . This will lead to a messier and less-effective protocol, and certainly will inhibit effective cavity-reset. To be sufficiently selective the displacement would occur on a timescale of roughly  $3/(\chi_{\text{ancilla}} - N\chi)$ . During that time the non-vacuum cavity state will be evolving at an (entangling) angular rate  $\sim N\chi$  and that evolution does not commute with the displacement operator. The original proposal claims several times that this evolution can be neglected, but I am not convinced that one can manufacture a device where this is so.

Our alternate cavity-reset technique is based on the realization that we can invert the unitary dynamics of the original entangling step. By simply performing  $\pi$  pulses on the register mid-way through the selective ancilla  $\pi$  pulse, then performing exactly the same protocol, we can erase the original register-cavity entanglement. One caveat is that as the ancilla is now not necessarily in the ground state, we must also add an additional  $\pi$  pulse during the erasure to remove its contributed phase<sup>7</sup>. Another benefit of this erasure technique can be seen in its duration, which is precisely equal to  $\pi/\chi_{\text{min}}$ . It also makes certain that the experiment acts as an echo sequence for all qubits involved, which may have added coherence benefits. This erasure technique can be seen in figure 6.5 without the central echo pulses, as well as more completely in parts d and e of figure 6.9.

<sup>7</sup>It may not be clear that the simultaneous register pulses are necessary. They are necessary to counter the evolution that occurs in the duration of the selective  $\pi$  pulse on the ancilla when measuring operators of more than two qubits or when the dispersive shifts are not equal. It will also become clear in the next section that it is necessary when the conditional cavity phase contribution is not exactly  $\pi$ .

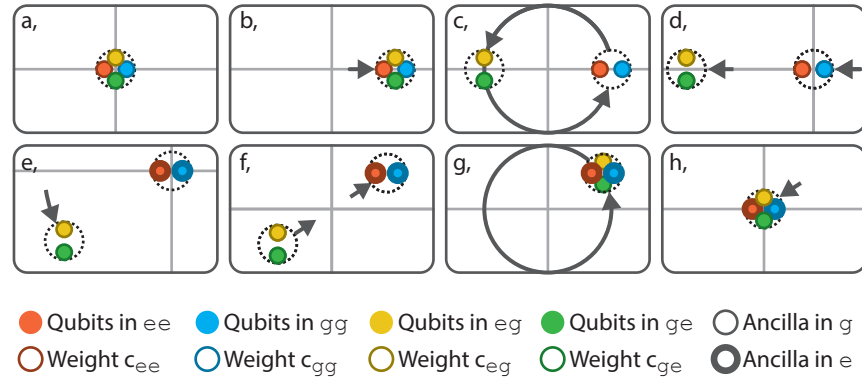


Figure 6.5: **Measuring  $ZZ$  on a two-qubit register.** This simplified image depicts the measurement of  $ZZ$  on a two-qubit register with equal dispersive shifts. The dotted ring represents the location of a coherent state in phase space, and each smaller circle represents a state of the qubit register. **a** After starting in an arbitrary register state, the cavity is quickly displaced, **b**, and allowed to evolve for a time  $\pi/\chi$ , **c**. After this evolution the cavity is displaced, **d**, such that the even-parity register states are correlated with the cavity being in the vacuum state. A spectrally narrow pulse excites the ancilla if and only if there are zero photons in the cavity, indicated by the thickness of the rings around each state indicator, **e**. The cavity state evolves during this pulse. The original dynamics are inverted, **f-h**, and the cavity is reset to the vacuum, leaving the register parity entangled with the ancilla. For simplicity, this image omits echo pulses on the register and ancilla.

### 6.2.4 Going faster

For measurements of one- and two-qubit properties, we may exploit the same speedup technique we found for qcMAP (section 6.1.4). This has essentially the same benefits and drawbacks as previously discussed, but now we should directly consider that there are multiple transmons decohering during this process, so going faster becomes even more attractive. We can simply take the calculations in equations 6.5 and 6.6 and set the target contributed phase to some amount smaller than  $\pi$ . We then adjust the displacement phases and (larger) amplitudes to assure cavity recoherence when and where we require it. One can also see that this is not directly applicable to measurement of three-properties (and of higher numbers of qubits). As the dispersive evolution actually measures excitation number, to obtain a parity measurement we must now take advantage of the periodic nature of phase to obtain recoherence. It is a happy fact that for one- and two-qubits measurements of parity can be equivalent to measurement of (an) excitation number. That said, it may be possible to chain multiple measurements of two-qubit properties (without the actual ancilla measurement) to obtain accelerated measurements of multi-qubit

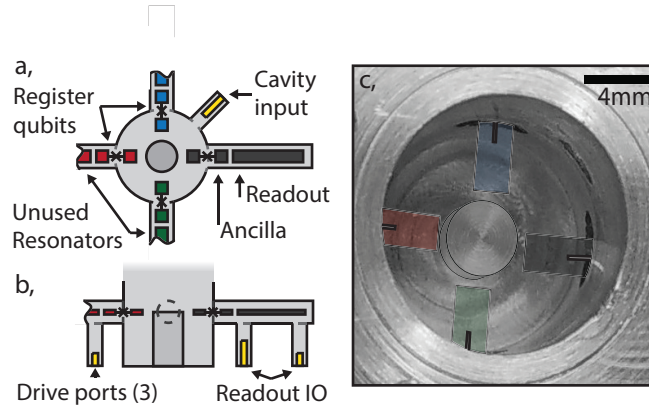


Figure 6.6: **Schematic of the device used in the stabilizer experiment.** A central stub resonator couples to four 3D transmons qubits, each of which couples to its own  $\lambda/2$  resonator. All of these resonators/transmons have their own undercoupled drive lines, as does the stub cavity. One of the  $\lambda/2$  resonators also has an overcoupled output port which leads to a near-quantum-limited amplifier. Panel **a** shows a top-view cartoon, and panel **b** shows a side view. Panel **c** is a real photograph of the device, with false color for the four sapphire-chips. The transmons and central stub have been outlined for visibility. The large ring around the edge is a groove to hold an indium wire or gasket. This wire contributes to a better light-tight seal.

properties.

### 6.3 Device characterization experiments

The device we created to enact this experiment (described in section 4.2.2) is fairly complicated. A central  $\lambda/4$  coaxial stub resonator (Reagor *et al.*, 2013) couples to four 3D transmons qubits, each of which couples to its own  $\lambda/2$  stripline resonator (Axline *et al.*, 2016). One of those qubits played the role of the ancilla and its associated stripline resonator coupled to a JPC, allowing high-fidelity individual readout. Characterizing this device fully was a non-trivial endeavor and in this section we will look at a general recipe for that process. We will explore several of the less-common characterization experiments in greater detail.



### 6.3.1 Direct characterization

The readout cavity was characterized directly with a vector network analyzer (VNA). Its frequency and bandwidth can be found via transmission with the JPC off. This process generally begins starting with high-power transmission to locate the device initially, then lowering power until the signal stops changing. It is important to use sufficiently low power for this measurement to make sure that you are not detecting the high-power peak (Reed *et al.*, 2010a), and between the high-power regime and the desired few-photon regime the cavity is unstable and its behavior is unpredictable. Low power also means we can avoid power-broadening (Steck, 2007). If the directly coupled qubit has non-negligible natural polarization one can also directly observe  $\chi$ . At this point one should turn on and tune-up the parametric amplifier<sup>8</sup>.

One of the qubits, the ancilla, can be directly characterized. From personal experience I recommend first finding a qubit via the VNA, rather than moving immediately to pulsed spectroscopy. This is done by looking at the readout resonator in transmission (with sufficiently low power) and examining the phase roll. One can then use a continuous-wave (CW) tone from an additional microwave generator to find the qubit. I start with low power and slowly manually change the generator frequency across the expected qubit range, looking for a shift in that phase roll. If there is no signal, increase spectroscopy power as needed. After finding a location that elicits a response (or in practice often several locations) the qubit can be found with standard pulsed spectroscopy. I choose the amplitude of that initial readout pulse such that I can observe a reasonable trajectory as in figure 3.11. This is fine tuned-later after we have control over the qubit.

Next we employed the standard characterization experiments: Rabi oscillations,  $T_1$  decay, and Ramsey oscillations with and without a Hahn echo. Ramsey oscillations are a more precise way to find the frequency of a qubit than spectroscopy. We used repeated pulse trains to fine tune the  $\pi$  pulse amplitude. The experiment we decided upon was to choose some fixed number of  $\pi$  pulses (roughly tuned from Rabi oscillations) and to sweep their amplitude. The resulting signal can be fit to a parabola, the minimum of which gives the correct drive amplitude. Chris Wang has shown that this pulse train is asymptotically unbiased against  $T_1$  and  $T_2$  as the number of pulses increases. We also employed DRAG

---

<sup>8</sup>If the JPC is employed before this you cannot trust the observed linewidth or frequency

(Motzoi *et al.*, 2009). This tuneup experiment, as detailed in (Reed, 2013), is done by comparing the signal from two different experiments and varying the DRAG parameter (amplitude balance between pulse and its out-of-phase derivative). One is  $X_{\pi/2}$  followed by  $Y_{\pi}$ , the other is  $Y_{\pi/2}$  followed by  $X_{\pi}$ . This experiment corrects the phase error due to having finite anharmonicity, but is *not* optimized to minimize leakage to the transmons higher levels (Chen *et al.*, 2016). In this experiment our standard  $\pi$  pulse had a Gaussian envelope, truncated at  $\pm 2\sigma$  with  $\sigma = 3.5$  nanoseconds. We also use a spectrally selective  $\pi$  pulse with  $\sigma = 2\mu\text{s}$ <sup>9</sup>. We employed randomized benchmarking to provide a general characterization of our pulse quality. We found the natural polarization of the ancilla two ways, both through direct measurement and with the RPM protocol (Geerlings *et al.*, 2013).

The bus cavity was found via “cavity spectroscopy,” a three-tone experiment. We swept the frequency of a pulse on the bus cavity, attempting to find resonance. Immediately after that pulse we perform a photon-number selective  $\pi$  pulse on the ancilla and measure it. This tells us if there are non-zero photons. When the bus resonance is found the ancilla will not be excited. The supplemental material of (Vlastakis *et al.*, 2013) provides an explanation for how to calibrate cavity displacements, how to measure  $\chi$ , and how to measure cavity relaxation times. We use a fast, unselective displacement (a square 10 nanosecond pulse) as well as slow, spectrally selective displacement (Gaussian envelope with sigma of  $2\mu\text{s}$ ). The spectrally selective displacement on the cavity populates the cavity if and only if *all* of the qubits are in the ground state. For the  $\chi$  measurements we also swept the applied  $\bar{n}$  which yields the higher order term  $\chi'$ .

We fine-tuned the resonance frequency of the stub resonator using an experiment inspired by Ramsey oscillations. We displaced the cavity, waited a set time  $T$ , and performed a second displacement of the same amplitude but negative before performing a selective  $\pi$  pulse on the qubit. We swept the *phase* of that second displacement. The ancilla signal changed when that second displacement left the cavity near the vacuum. By looking at that curve we can tell how much phase the cavity has acquired. To roughly control for “cross-Kerr” (a direct coupling between the stub and readout resonators), we also perform the same experiment without the  $\pi$  pulse, and subtracted the two signals. The magnitude of the phase shift divided by  $T$  gives the detuning of the cavity. We generally used at least two times  $T$

<sup>9</sup>Actually we use two different selective  $X_{\pi}^0$  pulses. As we will see later and was briefly discussed in the context of qcMAP, the pulse we use for the Nigg protocol does not have to be perfectly selective on zero photons.

to control for any offsets. Furthermore, we also generally perform this experiment at various different average photon numbers  $\bar{n}$ . By fitting the detuning vs.  $\bar{n}$  we can obtain the cavity self-Kerr and be certain that our rotating frame has not been modified by Kerr<sup>10</sup>.

### 6.3.2 Indirect qubit characterization

The three qubits of our register required a somewhat more sophisticated measurement technique, which we colloquially referred to as “mapped” readout. In this regime we performed a standard operation on one of the qubits (say a spectroscopy pulse), then performed a selective displacement on the cavity, followed by a selective  $\pi$  pulse on the ancilla, followed by an ancilla measurement. This sort of four-tone experiment sounds more challenging than it is, but it is a straightforward extension of the techniques we used to measure the bus resonator.

Measuring the  $\chi$  of these qubits is somewhat less simple. The mapped readout does not work if there are photons left in the bus, so we cannot use the qubit-rephasing technique we applied to the ancilla. Instead we perform a modification of the experiment we used to find the bus frequency and self-Kerr. We compare two experiments where as before we displace the bus, wait, and displace back with a swept phase. In one of the experiments we do exactly as before, providing a control. In the second we perform a  $\pi$  pulse on one of the register qubits beforehand. The difference in the two observed frames is the  $\chi$  of that qubit to the bus. Similarly we can sweep the displacement amplitude and find the relevant  $\chi'$ .

The final type of characterization experiment involve the three unused  $\lambda/2$  resonators. These five-tone experiments are an unsurprising extension of the above, where now the selective displacement and ancilla  $\pi$  pulse are preceded by a selective  $\pi$  pulse on the relevant register qubit.

Other than  $\chi$  and  $\chi'$ , each register qubit was characterized with the same experiments as the ancilla, now using the mapped readout. We also measured the direct qubit-qubit coupling between each pair of qubits, done with a set of fixed-time Ramsey experiments, as detailed in (Reed, 2013). We additionally performed simultaneous randomized benchmarking (Gambetta *et al.*, 2012) to measure the cross-talk errors between qubits, results shown in table 6.2.

<sup>10</sup>Though you do generally want your choice of frame to account for Kerr and your expected number of photons.

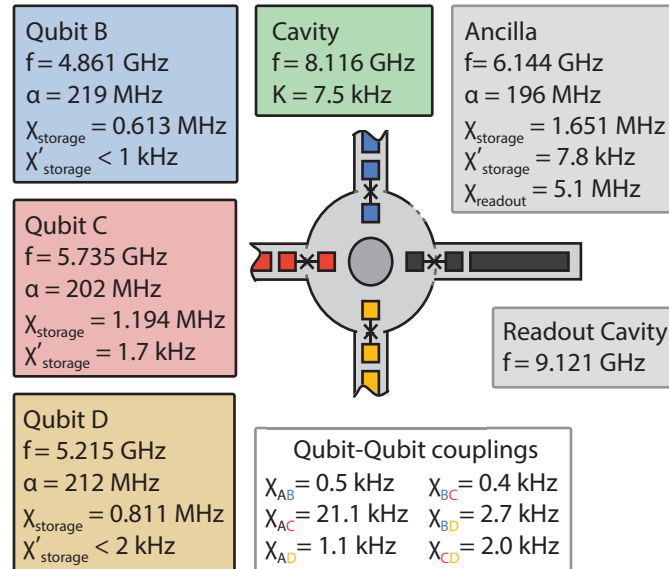


Figure 6.7: **Hamiltonian for the device used in the stabilizer experiment.** These Hamiltonian terms are defined precisely in section 3.5.1.

### 6.3.3 Characterization results

As discussed, we characterized the Hamiltonian to fourth order in flux, as well as the most significant sixth order terms. These results are given in figure 6.7. In particular, note that the direct qubit-qubit couplings are very low compared to the dispersive shifts. Part of what allowed us to implement the Nigg proposal precisely is that we were able to control the undesired couplings. As discussed in section 3.5.3, there are essentially concerns to deal with. First, the qubit modes are fairly well localized. This is evident in the relatively small dispersive shifts they each share with the bus cavity. Secondly, they are significantly detuned from the bus. This helps minimize the direct coupling that arises from their individual couplings to the bus. The mediated or “virtual” interaction becomes evident under higher-order perturbation theory.

For the coherence results (table 6.1), all times are given in microseconds. The bus resonator had an energy decay constant of  $72 \mu\text{s}$ , which is not unexpected for cavities made of 6061 aluminum. Though the stub geometry minimizes seam loss, compared to a rectangular resonator it is actually twice as sensitive to conductor losses. It also has a narrower aspect ratio, making it more susceptible of losses in the dielectric on the surface of the bulk aluminum. The overcoupled readout resonator had a decay

Table 6.1: **Coherence results for the Shiva device.** All times in microseconds.

qubit	$T_1$	$T_2^{\text{Ramsey}}$	$T_2^{\text{Echo}}$
ancilla	23(2)	19(2)	26(3)
register B	86(5)	57(8)	73(10)
register C	87(5)	62(10)	77(7)
register D	58(18)	*	52(8)

Table 6.2: **Simultaneous randomized benchmarking results.** All quantities are scaled by  $10^{-3}$ , and the naming conventions follow (Gambetta *et al.*, 2012).

Qubits (x,y)	$r_x$	$r_{x y}$	$ dr_{x y} $	$r_y$	$r_{y x}$	$ dr_{y x} $	$ d\alpha_{xy} $
$A,C$	1.4(1)	1.4(1)	0.0(1)	1.9(3)	1.7(2)	0.2(4)	0.2(7)
$A,D$	1.4(1)	1.1(2)	0.3(2)	1.4(4)	0.7(1)	0.7(4)	0.3(6)
$A,B$	1.4(1)	1.6(1)	0.2(1)	0.6(2)	0.8(1)	0.2(2)	0.5(3)
$C,D$	1.9(3)	2.0(1)	0.1(3)	1.4(4)	0.6(1)	0.8(4)	0.6(4)
$C,B$	1.9(3)	1.6(1)	0.3(3)	0.6(2)	1.0(1)	0.4(2)	0.7(5)
$D,B$	1.4(4)	0.9(1)	0.5(4)	0.6(2)	0.5(1)	0.1(2)	0.5(2)

time of 60 nanoseconds. The other three undercoupled  $\lambda/2$  resonators had lifetimes between 2 and 10  $\mu\text{s}$  in the cooldown which lead to data for the paper, though some of the same resonators had lifetimes over 100  $\mu\text{s}$  in an earlier cooldown. We do not know what lead to this degradation, though I suspect it relates to the mounting indium becoming loose upon thermal cycling.

Note that error in these numbers does not come principally from fitting errors, but rather actual variations in lifetime over time, as is commonly seen (Muller *et al.*, 2015). Qubit D does not have a  $T_2^{\text{Ramsey}}$  quoted as that Ramsey experiment does not show a single frequency or time constant. That qubit also has a very large spread in  $T_1$  results, which exhibits a non-Gaussian distribution. In previous thermal cycles qubit D showed clean Ramsey oscillations with lifetimes over 60  $\mu\text{s}$ . We are unable to ascribe a cause to this behavior.

## 6.4 State preparation

As is fairly standard for circuit QED experiments in our group, the natural polarization of the transmons was far above what would be expected from the base temperature of the dilution refrigerator ( $\approx 10^{-6}$ ). The four qubits ranged between 1.4 and 3.8% in the excited state. As this experiment was done without

real-time feedback, we used a post-selective technique to prepare the ground state.

Shown in figure 6.8, we used three individual measurements. The first simply guarantees that the ancilla is in the ground state. We noted that after measuring one “g” result the likelihood of a second “g” result was very high. In fact the temperature was difficult to measure after the full sequence (three measurements) succeeded. Using the RPM technique (Geerlings *et al.*, 2013) we bounded that ancilla population to be under 0.1%.

The next step was a set of seven simultaneous spectrally-narrow displacements, each resonant with the cavity with a different configuration of the register qubits excited. If the ancilla is in the ground state, this populated the bus if and only if at least one of the register qubits was excited. Note that we would not want to use one single selective displacement if the register was in the ground state, as that would never leave the register and bus cold at the same time. We then performed a selective  $\pi$  pulse on the ancilla if and only if there were zero photons in the bus, and measured the ancilla. If the rest of the system is cold, we now expect the ancilla to be excited. We then perform a fast, unselective  $\pi$  pulse on the ancilla, measure again, and expect the ancilla to be in the ground state. As this process was post-selective, after performing these measurements we always continued on and performed the experimental pulse sequence and measurements. If the measurement record began with “geg,” then the system should be in the ground state and we retained the later data. If not, that data was discarded. This cooling was relatively effective: we reduced the system excited-state population by roughly 10%, and discarded a fraction of data not significantly larger than that.

Note that we use two measurements at the end (“e” and then “g”), to avoid a false-positive “g” result—they are more common than false positive “e” results—which might erroneously indicate the bus to be empty. This sequence was not optimized to maximize kept counts, rather it emphasized a colder system. This process enhanced the collective qubit ground state population to well over 99% and decreased the bus population to below detectability. In the final data-taking cooldown the population of the unused planar resonators was not detectable, though given the relatively long Ramsey decay times of the register qubits we can place a fairly low upper bound.

After ground state preparation we performed rotations to prepare the register in a known state. Quantum process tomography requires a complete or overcomplete set of states. We used all permuta-




Figure 6.8: **State preparation in the stabilizer measurement experiment.** In part **a** we used a series of measurements to herald that the system was in the ground state. Because we did not have access to directly measure most of the system, we used a somewhat delicate series of interactions to transfer that information to the ancilla. This series was designed to be non-destructive if the results indicated that the system was cold. Part **b** allows us to prepare a chosen state other than  $|ggg\rangle$  by applying single-qubit rotations to the register qubits. This is necessary for quantum process tomography.

tions of  $I$ ,  $R_x(\pi/2)$ ,  $R_x(\pi)$  and  $R_y(\pi/2)$  on each qubit.

## 6.5 Full parity measurement sequence

We implemented the protocol as described in section 6.2 for all seven non-trivial subset-parity operators:  $ZII$ ,  $IZI$ ,  $IIZ$ ,  $ZZI$ ,  $ZIZ$ ,  $IZZ$  and  $ZZZ$ . Each differed principally by the choice of echo sequence. We used a conditional  $\pi$  phase separation for  $ZZZ$ , and  $2\pi/5$  for all other operators. Each timing was adjusted to reflect that angle and the smallest  $\chi$  of the measured qubits. Figure 6.9 depicts an example sequence, in this case measuring  $ZIZ$ . This circuit follows exactly after the state-preparation circuit depicted in figure 6.8.

## 6.6 Error detection

Towards the end of this experiment we noticed that many of the things which can go wrong leave a distinctive error signature: the bus may not be reset to the ground state. If at any point a qubit relaxes during this delicate dance of coherent states, then the trajectory of the bus through phase space will quickly become unsynchronized and the odds of being reset to the vacuum are essentially zero. Higher order effects such as the cavity self-Kerr and  $\chi'$ , which lead to “smearing” of the coherent states, will also have some probability of leaving the cavity not in the ground state.




Figure 6.9: **Circuit diagram for ZIZ measurement.** Steps  $X_\pi$  refer to one-qubit rotations around the  $X$  axis by  $\pi$  radians,  $X_\pi^0$  indicate that the pulses are spectrally narrow and are roughly selective on having zero photons in the cavity. Steps  $D_i$  represent unconditional displacements of the cavity. The meters are measurements of the ancilla via the readout resonator, which is not itself depicted. The ancilla has the largest dispersive shift and the register qubits are then numerically ordered (from top to bottom) such that  $\chi_1 < \chi_2 < \chi_3$ . **c**, The algorithm begins with a displacement  $D_1$  to create a coherent state of  $\bar{n} = 5$  photons into the cavity, which acquires a phase shift  $\theta$  in a time  $T = T_5 - T_0 \approx \theta / (2\pi\chi_1)$  conditionally on the state of qubit 1 (blue). For measurements of one- and two-qubit properties  $\theta = 2\pi/5$ . In this example, we perform a full echo on the second qubit (yellow) by performing two unconditional  $X$  gates separated in time by  $T_4 - T_1 \approx T/2$ . The third qubit (red) would contribute a conditional phase shift of  $2\pi\chi_3 T > \theta$ . We reduce this to  $\theta$  by performing two  $X_\pi$  gates separated by  $T_3 - T_2 \approx \theta (\chi_1^{-1} - \chi_3^{-1}) / 2$ . At  $T_5$ , we perform  $D_2$  to shift the odd two-parity coherent-state pointer to the zero-photon state. Note that the overlaps between the even two-parity pointer states and the zero-photon state are exponentially suppressed. **d**, We map this photon number information onto the ancilla qubit with a  $X_\pi^0$  gate, taking advantage of the well known number-splitting phenomenon (Schuster, 2007). As the cavity states are separated by  $\approx 6.5$  photons, we employ a faster, approximately selective gate, 300 ns in duration.  $X_\pi$  gates on the register are centered on this pulse in time to echo away the cavity evolution during this step. **e**, To disentangle the cavity pointer states we essentially invert the pulse sequence of **c**, returning the cavity to the vacuum state. We must also echo the ancilla as it may now be excited. This results in a total gate length of 970 ns. Subsequently, we measure the ancilla qubit. **f**, This optional step determines if there are residual photons in the cavity. Since many types of errors result in residual photons, a subsequent photon-number-selective rotation and measurement of the ancilla heralds these errors. When measuring three qubits (e.g.  $ZZZ$ ), we choose  $\theta = \pi$  so that the cavity states entangled with the one- and three-excitation manifolds recohere.



As shown in part f of figure 6.9, after our experimental protocol we can look for that signature. After measuring the result of our operator of interest, we can perform a zero-photon selective  $\pi$  pulse on the ancilla. If the result has inverted from the previous measurement, we can consider the experiment successful. We provided datasets with and without the success herald. Whether or not such a herald is useful depends heavily on the intended application. The herald does not aid assignment fidelity with any errors that occur after the central ancilla-entanglement step. Nor does it ever help with dephasing errors, or appreciably with photon loss. It also obviously does not help with errors due to finite measurement fidelity.

## 6.7 State & process tomography

As discussed in section 5.6.2, we can characterize the back-action of the process using measurement process tomography. Since we did not have individual readout of the register qubits, we used a “mapped” readout, akin to the techniques in section 6.3.2. Since we cannot now count on the ancilla being in the ground state, we performed a slight modification. We displaced the cavity, waited,  $\pi$  pulsed the ancilla, waited the same time, and displaced back by a specific angle. Using reasoning similar to that in the Nigg and Girvin proposal, this can be engineered to result in the bus being in the vacuum if and only if the register is in the ground state but irrespective of the ancilla state. We then performed a zero-photon-selective  $\pi$  on the ancilla before measuring. This results in essentially measuring the  $|ggg\rangle\langle ggg|$  operator whose signature is a change in the ancilla state. Clearly, this process relies on the bus being initially in the vacuum, but as we discussed before, many errors can preclude that. Accordingly, if we did not use the zero-photon error-herald, we performed a four-wave mixing process (section 3.7.2) to unconditionally empty the cavity. We performed state tomography with various cavity displacements to verify the efficacy of this protocol.

While one could reconstruct the entire process in one mathematical step, this is numerically untenable. Instead we performed outcome-dependent state tomography for each input state, then performed an additional (process) reconstruction from the states. For state tomography we chose to use an over-complete set of postrotations, permutations of  $I$ ,  $R_x(\pi/2)$ ,  $R_x(\pi)$ ,  $R_y(\pi/2)$ ,  $R_x(-\pi/2)$  and  $R_y(-\pi/2)$




Figure 6.10: **Circuit for state tomography in the stabilizer measurement experiment.** After performing an engineering measurement, we wished to analyze the back-action on the initial state. In step **g** we perform unitary postrotations to set the tomographic basis. We used an overcomplete set for the purpose,  $\{I, R_x(\pi/2), R_x(\pi), R_y(\pi/2), R_x(-\pi/2), R_y(-\pi/2)\}^{\otimes 3}$ . Since we do not have independent readout for each qubit, our tomography is mapped through the ancilla. This depends on the bus being initially depopulated. If we did not use the zero-photon herald from section 6.6, we used a four-wave mixing process as discussed in section 3.7.2 to empty the cavity in step **h**. Panel **i** shows the tomography measurement itself. We displaced the cavity, waited,  $\pi$  pulsed the ancilla, waited the same time, and displaced back by a specific angle. We then performed a selective  $\pi$  pulse, inverting the ancilla state conditioned on an empty bus resonator. The purpose of this was to flip the ancilla state (regardless of the initial ancilla state) if and only if the register was in the ground state. Accordingly we could essentially measure the  $|ggg\rangle\langle ggg|$  operator.

for each qubit. As explained in (de Burgh *et al.*, 2008), it is always beneficial to use an overcomplete set if possible, chosen with measurement axes that align with the faces of a platonic solid. In addition to unbiasing certain errors, on a per-measurement basis they give a better SNR than taking more averages on a smaller set. We limited ourselves to six rotations purely for technical reasons, as we could not practically fit a larger set on the Tektronix AWG in one experiment.

Additionally, we generate calibration data for the tomography measurement at the same time as taking experimental data, preparing and directly measuring all computational states of the register. Calibrating the measurement operator in this somewhat limited manner is directly akin to the calibrations done in (Chow, 2010, Reed, 2013), and assumes that our tomography operator is only sensitive to  $Z$  information. This is a reasonable assumption, which we later confirmed with complete detector tomography of the state tomography measurement operation.

The mathematics of this are straightforward. With knowledge of the set of rotations  $\{R_j\}$ , the set of observed probabilities  $\{P_j\}$ , and the simultaneously calibrated measurement operator  $V$ , we can invert to find the unknown state. The relationship between these is simple,  $P_j = \text{Tr} [R_j \rho R_j^\dagger V]$ . After expanding the unknown state in the Pauli basis,  $\rho = \sum_k c_k \sigma_k$ , this leads to a matrix relationship  $\vec{P}_j = A_{ij} \vec{\rho}$ .

Rather than use the a pseudo-inverse to solve for  $\rho$ , we performed a maximum likelihood estimation. We used the CVXPY package to do this inversion, constraining  $A$  to be both positive semi-definite and normalized. These constraints are simply a statement that we require all probabilities generated from these states to be positive and add to one. With these states in hand we performed the next step for process tomography. This inversion was done in MATLAB with significant help from Chao Shen. This maximum-likelihood estimation constrained the entire quantum instrument (section 5.6.3) process to be a CPTP map.

## 6.8 Experimental results

### 6.8.1 Assignment fidelity results

We present here the results of the characterization experiment described in section 5.3. In practice we used more points than the plotted plane cuts, which alone resulted in an unstable fit. We also report results both with and without the error herald. We show the plane cut preparations in figure 6.11, which provide a qualitative assessment of the measurement operators.

We can more quantitatively extract the contrast which corresponds to the measurement operator. Table 6.3 gives the complete results, both with and without heralding on an empty bus cavity (section 6.6).

Table 6.3: **Assignment fidelity results.**

operator	unheralded results	heralded results
ZII	0.93	0.97
IZI	0.94	0.97
IIZ	0.94	0.96
ZZI	0.93	0.97
ZIZ	0.93	0.96
IZZ	0.93	0.97
ZZZ	0.89	0.94

Note that assignment fidelity errors are inherently bounded to the fidelity of the ancilla measurement itself, which in our case was roughly 98%. This means that with the error herald we are measuring two-qubit properties (and to a lesser extent  $ZZZ$ ) almost as well as we are able to measure the ancilla

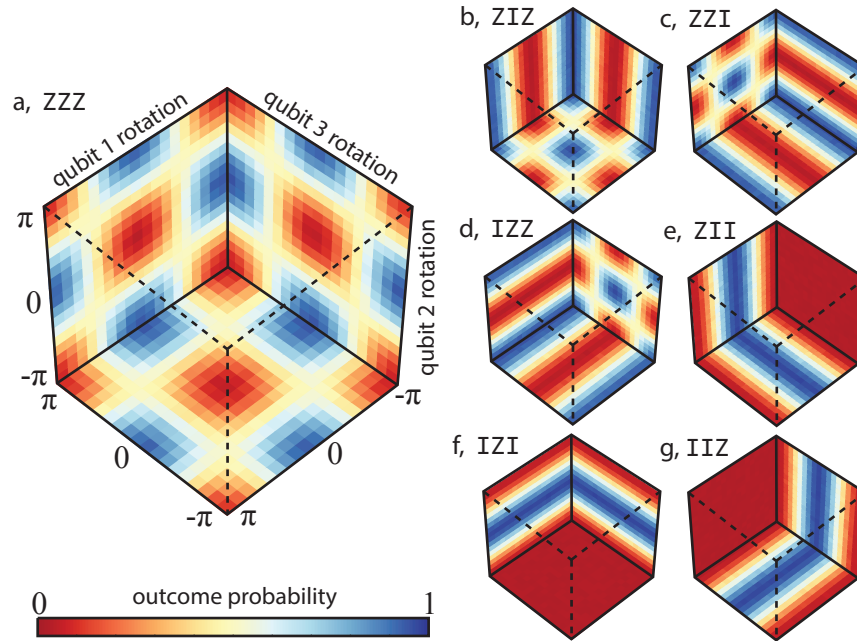


Figure 6.11: “TriRabis” showing raw data for measuring assignment fidelity. In these experiments we prepare a large set of separable states, rotating each of the three qubits around the  $X$  axis by a variable angle. These plots show plane cuts in that preparation space, where the color represents the probability of finding the ancilla in the ground or excited state after the measurement protocol. As a probability, this color reflects the POVM element for outcome  $O = +1$ . Up to a possible bias (section 5.4.2), this value is proportional to the desired operator expectation value minus minus one. These data reflect post-selection on an empty cavity but are not otherwise renormalized. Panel **a** shows the result of a measurement engineered to be  $ZZZ$ . One can clearly see that changing the preparation of any single qubit completely inverts the measurement outcome. Panels **b-d** show the results of two-qubit parity measurements. The outcome is invariant under one preparation axis, showing we are not measuring that qubit. Panels **e-g** show measurement of a single qubit. Clearly only one preparation angle matters, showing that we are only measuring one qubit.

alone. With respect to the unheralded results, we should most significantly be limited by the ancilla measurement fidelity and the  $T_1$  of all four qubits, though the ancilla  $T_1$  is not relevant for the initial mapping step. For a mapping step of roughly 400 nanoseconds and an ancilla selective  $\pi$  pulse time of 300 nanoseconds,  $T_1$  should limit us to slightly over 96%. Along with the measurement infidelity we should be limited to 94%. There are additional costs associated with imperfect  $\pi$  pulses, especially in the presence of bus photons. Cavity photon loss will also contribute, but it has a larger effect on dephasing which is not relevant to assignment fidelity.

### 6.8.2 Detector tomography results

As discussed in section 5.4.1, we performed quantum detector tomography on all seven measurement operators, with and without error heralding. Also as discussed in that section, for each we report both the standard J-fidelity and the worst-case S-fidelity. More expansive results, including all of the complete reconstructed POVMs (one of which was plotted in figure 5.2), specificity (section 5.4.2), and the corresponding distance-based metrics, can be found in appendix A.

Table 6.4: **Detector fidelity results.**

operator	unheralded results		heralded results	
	$F_J$	$F_S$	$F_J$	$F_S$
ZII	0.948	0.936	0.976	0.967
IZI	0.945	0.931	0.975	0.963
IIZ	0.958	0.948	0.976	0.967
ZZI	0.938	0.906	0.967	0.938
ZIZ	0.943	0.914	0.967	0.938
IZZ	0.942	0.918	0.966	0.934
ZZZ	0.910	0.875	0.954	0.926

In the case of strong, binary measurements, the detector J-fidelity and assignment fidelity have very similar meanings. We should not be surprised that they have very similar values. If there is any distinction, it is that assignment fidelity only penalizes a biased measurement through the necessary bound that places on measurement contrast. The J-fidelity should penalize both the contrast reduction and the bias term directly.

### 6.8.3 Measurement process tomography results

For each of the seven operators we performed measurement process tomography (section 5.6.2) on datasets both with and without the success herald. Since each of those fourteen datasets yields two 64-by-64 complex matrices, we will not present the full data. Rather we can look at the important bits of a single one of those datasets in figure 6.12, and provide somewhat more exhaustive results in appendix A.

We can also extract the measurement process fidelities (both J- and S-) for all fourteen of these datasets.

Table 6.5: **Quantum instrument fidelities.**

operator	unheralded results		heralded results	
	$F_J$	$F_S$	$F_J$	$F_S$
ZII	0.716	0.566	0.789	0.690
IZI	0.757	0.615	0.832	0.749
IIZ	0.757	0.607	0.827	0.756
ZZI	0.721	0.629	0.818	0.754
ZIZ	0.735	0.638	0.823	0.757
IZZ	0.730	0.595	0.813	0.759
ZZZ	0.674	0.578	0.796	0.741

As the measurement process includes the detector (explicitly shown in section 5.6.1), the detector fidelity will form an upper bound for the measurement process fidelity. A discrepancy between the two numbers will indicate the additional imperfection of the back-action relative to the ideal. We expect to see such a discrepancy. In addition to those previously discussed, we are now sensitive to the following significant effects:

1. Register dephasing (natural  $T_\phi$ )
2. Register dephasing and relaxation during the ancilla measurement (900 nanoseconds total)
3. Dephasing due to photon loss
4. Dephasing due to entanglement with residual photons

From  $T_2$  alone, for the (unheralded) one- and two-qubit properties we should expect to be limited to roughly 85%. Ancilla measurement infidelity brings that down to 83%. Calculating the dephasing due to photon-loss is challenging as the photon number is not constant, and the separation in phase space is even less so. Fortunately we can use a computer to numerically solve the master equation, and we see that it should cost us roughly another 4-5%. The largest remaining factor between the performance we expect and the performance we observe is the dephasing of qubit  $B$ . This qubit exhibited a clean  $T_2^{\text{echo}}$ , but a noisy and short  $T_2^{\text{Ramsey}}$ . Previous calculations have been done using the echo time constant, which should be the more significant number. Though we attempted to make the whole experiment an echo sequence, it is possible the behavior indicated in the Ramsey oscillations still had an effect. Additionally, this sequence involved fourteen (uncalibrated) unselective single-qubit rotations (including echo pulses on qubit  $B$  during the ancilla measurement), which likely contributed an additional two to

three percent.

It is not terribly surprising that the S-fidelities are significantly lower than the J-fidelities. Dephasing errors exhibit large discrepancies between average- and worst-case behavior, as they are more detectable if the qubit is entangled with an auxiliary space. Since we are dominated by dephasing errors, we expect such a degradation.

## 6.9 Conclusions

In conclusion we have demonstrated a modified version of the proposal in (Nigg and Girvin, 2013), using an interesting cQED device built for that purpose. We went through a rigorous characterization process and devised new figures of merit. While the resultant process fidelities are not overwhelming, one needs to be aware of what number this fidelity should be compared to. It contains ancilla preparation, (effectively) three two-qubit gates, and the ancilla measurement, as well as all decoherence that happens during that time. This cannot be simply compared to a 99% state preparation, three 99% gates, and one 99% measurement—all of those numbers are derived by only looking at the behavior of only a tiny fraction of the full Hilbert space. The inclusive nature of this measure naturally leads to lower numbers but does not allow us to sweep anything under the rug.

The prospects for improved performance are challenging to quantify. While we can reduce the dephasing due to cavity loss, we cannot make the gate appreciably faster. Any scheme to increase  $\chi$  will much more quickly lead to increased higher-order terms, which will inhibit a clean gate. Furthermore any scheme using larger numbers of photons will also become limited as unselective  $\pi$  pulses in that regime will cease to be unselective. Additionally this sort of all-at-once gate will suffer from highly correlated errors, which likely render it unappealing for a fault-tolerant computation scheme. Though I am treating this with a pessimistic attitude, other current approaches have similar (or worse) issues of their own. Many of these concerns (though not all) are alleviated by using only pair-wise interactions with dedicated buses between qubits. If we accept that constraint I am not convinced that this protocol offers significant advantages. That said, I believe it to be an interesting exploration of circuit QED, and think that the characterization experiments and mindset are highly valuable for the future of fault-tolerant quantum

computing.






Figure 6.12: **Quantum process heralded by a ZZZ measurement.** Our measurement process should ideally be a von Neumann measurement, or a projection into the measurement eigenspace. For a strong measurement of a Pauli operator  $\sigma$ , the relevant Kraus operators for the two outcomes should be given by  $1/2(I \pm \sigma)$ . We plot here only the real parts of the corners of the 64-by-64  $\chi$  matrices. In panel **a**, corresponding to the outcome  $ZZZ = +1$ , we see that there are four notable elements to the matrix, at  $III, III, III, ZZZ, ZZZ, III$  and  $ZZZ, ZZZ$ . Each has a value approaching  $+1/4$ , which indicates our process is similar to the expected projection. We do not see significant other bars (including in the unplotted regions). In panel **b** we have plotted the same matrix for the result  $ZZZ = -1$ . As expected, the two off-diagonal bars change sign, indicating a projection onto the alternate eigenspace.

## Chapter 7

# Deterministic teleportation of a two-qubit gate

By making use of ... entangled states, these techniques vividly illustrate how entanglement can be a valuable resource for computation.

---

Gottesman and Chuang, 1999

The second major experiment of my thesis work involves the teleportation of a two qubit-gate between two non-interacting quantum systems. In particular that gate is a CNOT gate, and the two quantum systems are not transmons as in the previous experiment but harmonic oscillators. The operations on those oscillators use pulses designed by optimal control theory (section 4.3.2), which renders this experiment essentially agnostic to what encoding is used to store that logical qubit of information. Accordingly we demonstrate both the simplest encoding, where logical states are simply Fock states, and a more complicated encoding which is potentially error-correctable. As of the time of this writing the experiment has produced very fruitful results, but a manuscript has not yet been submitted for publication.

A teleported gate is a special thing, qualitatively quite different from a traditional unitary operation. We can perform such an operation without the two logical qubits ever directly interacting with each other. These are hallmarks of teleportation:

1. An auxiliary entangled pair is sacrificed and used as a resource.
2. High-fidelity readout is employed as part of the operation
3. Real-time feed-forward is performed using the results of those measurements.

In this chapter we will build the parts necessary to demonstrate such a gate, then I will give the results when we put those parts together. We will begin with section 7.1, which describes the basic motivation for such a gate. The Yale group believes that a modular architecture is the path forward most likely to show near-term progress towards a practical quantum computer. In such a scheme gates like these are likely to be essential for operations between modules. In section 7.2 we will discuss the progenitor for this idea, the teleportation of quantum states. In some ways state teleportation is very similar—it shares all of the above three hallmarks. In other ways it is qualitatively different, as quantum states are not actually “teleported” anywhere. After that we will discuss the basic idea of teleporting a gate and the underlying mathematics in section 7.3. We will not discuss the original proposal, but it may be found in this reference (Gottesman and Chuang, 1999).

We will then dive into experiment in section 7.4, giving the results of characterization of the Dodekathion device we described in section 4.2.3. These numbers will set the baseline for how well we might expect to execute the various building blocks of our protocol. The first such block to discuss involves two of the above hallmarks: readout and feed-forward. In section 7.5 we will describe the readout scheme and its performance. Section 7.5.1 provides a brief technical discussion of how our feedback/feed-forward operates. After that, in section 7.5.2 we will provide an example of using that feedback—deterministically preparing the system in the ground state with high fidelity. In section 7.5.3 we will then show a second example, in this case feed-forward, in a manner it is employed in the teleported gate. That section is about the transmon reset, which provides a holistic demonstration of our readout, feedback, and single-transmon operations.

The next section will pertain to another one of the hallmarks of teleportation: the sacrificial entangled pair. In section 7.6 we will derive and describe the resonator-induced phase (RIP) gate. This derivation looks a bit different than other available derivations, but I believe it to be more straightforward. We will also briefly describe the refocused rip gate in section 7.6.1, which is a critical modification for the RIP gate to be useful in our experiment. After this we will discuss qubits encoded in harmonic oscillators

in section 7.7. Beyond simply describing some of the options, I will give some empirical results about implementing them. I will quantify the operation performance we have observed in section 7.7.3, then give some thoughts on codespace and leakage errors in section 7.7.4, which I believe will be an important and rich field of study in the future.

The last subprocess of the protocol is how we will do tomography of the resulting states and gate, described in section 7.8. I will describe how tomography with multiple simultaneous bits of information is different than tomography with joint measurement operators. In section 7.8.2 I will go over adjusting the rotating frame of the cavity. This is important for more than simply tomography, but it is quite relevant there. Section 7.8.3 will cover one alternate method we might use for tomography of the logical state.

Finally in section 7.9 we will look at how all of the parts fit together. After this we will provide the experimental results showing the gate in action in 7.10. At the time of this writing, the experiment has demonstrated the teleportation of a CNOT gate, but we have not yet published any results. Accordingly all data in this chapter are preliminary. At this time we discuss quantum process tomography of a teleported gate when using the Fock encoding. For illustrative purposes we look at the operation with and without application of the logical feed-forward. I also present even more preliminary results for a teleported gate on error-correctable logical qubits.

## 7.1 Stitching a web of modules

The most popular vision of a quantum computer today features a large, monolithic device, consisting of millions of physical qubits. It has a two-dimensional connectivity graph based on nearest-neighbor couplings, and likely runs a topological code (Fowler *et al.*, 2012, Kitaev, 2003). This approach has several downsides. First, in a large monolithic system it will always be difficult to build components with the same quality as in smaller systems. Second, we are generally limited to nearest neighbor interactions. It is well-known that moving beyond a two-dimensional (Bacon, 2006, Bombin, 2015) or nearest-neighbor (MacKay *et al.*, 2004) connectivity graph can yield significant advantages for quantum error-correction. Among other things, these can permit transversal implementation of a universal gate




Figure 7.1: **A network of quantum modules.** The Yale group’s current vision of a quantum computer is not based on a large monolithic device, but rather on stitching together smaller but higher-quality modules. It may be easier to make units of one or several error-correctable qubits than deal with lower quality non-local inter-module links. This will likely involve a communication architecture which uses entangled pairs to fuel teleported inter-module gates.

set, which can lead to massive savings in the number of requisite qubits.

The Yale group’s vision of the future of quantum computing is somewhat different (Devoret and Schoelkopf, 2013). Their argument is that we may find better success by building a network of small modules, containing one or several error-corrected logical qubits (figure 7.1). The inter-module links may be of low fidelity, but we can compensate by improving inter-module entanglement with purification (Bennett *et al.*, 1996), which then becomes a resource for teleported gates (Gottesman and Chuang, 1999). On top of these “first-level” corrected qubits we can then run a higher level (possibly topological) code. Since the inter-modules connections are explicitly allowed to be of lower quality, we can easily envision allowing switching elements. Under that scheme we could take advantage of connectivity beyond nearest-neighbor qubits.

These logical qubits may not be composed of the tens-to-thousands of physical qubits needed to realize an error-corrected surface code qubit. Instead we have begun using the long lifetimes, large Hilbert spaces, and simple error models provided by harmonic oscillators. There are various encodings (Gottesman *et al.*, 2001, Leghtas *et al.*, 2013b, Michael *et al.*, 2016) which allow a harmonic oscillator to encode a logical qubit. Recently the Schoelkopf lab has demonstrated quantum error correction using one of those encodings that actually succeeds at extending the lifetime of a quantum memory (Ofek *et al.*, 2016). There are also potentially ways to employ autonomous error correction (Leghtas *et al.*,




Figure 7.2: **Circuit for state teleportation.** The parties, Alice and Bob, share a preexisting Bell pair, shown by the orange and red lines. Alice has an unknown quantum state (which would potentially require an infinite amount of classical information to describe), which she wishes to send to Bob. Astoundingly she can do so without access to a quantum channel, using this entangled pair and only local operations and classical communication with Bob.

2013b) on such systems. In (Michael *et al.*, 2016) we have begun to explore how to use these encodings to detect errors beyond photon-loss, including dephasing and photon-gain. Using pulses generated with optimal control theory (section 4.3.2) we have demonstrated a single-qubit set of universal gates (Heeres *et al.*, 2016). Local two-qubit gates have been demonstrated in a restricted form (Gao *et al.*, in preparation). In this experiment we demonstrate how non-local two-qubit gates will work in such an architecture.

## 7.2 Teleportation of quantum states

Before the notion of a teleported gate was the notion of a teleported state (Bennett *et al.*, 1993). In this scheme if two parties shared a preexisting Bell pair, one party can transmit an unknown quantum state to the second party without use of an additional quantum channel. Rather they each only need local operations and classical communication (LOCC)<sup>1</sup>. That this is possible is really quite remarkable. This realization in 1993 was one of the seminal moments in early quantum information theory. The standard circuit for state teleportation is depicted in figure 7.2. We can also walk through the state:

<sup>1</sup>There also exist so-called one-bit teleportation circuits which do not require a Bell pair, but they require a quantum channel (Terhal, 2015). I do not consider that teleportation

$$\begin{aligned}
|\Psi\rangle &= (\alpha|0\rangle + \beta|1\rangle)(|0\rangle|1\rangle + |1\rangle|0\rangle) \\
&\mapsto \alpha(|001\rangle + |010\rangle) + \beta(|111\rangle + |100\rangle) \\
&\mapsto \alpha(|001\rangle + |101\rangle + |010\rangle + |110\rangle) + \beta(|011\rangle - |111\rangle + |000\rangle - |100\rangle) \\
&= |00\rangle (X|\psi\rangle) + |01\rangle |\psi\rangle + |10\rangle (Z|\psi\rangle) + |11\rangle (ZX|\psi\rangle).
\end{aligned}$$

State teleportation was first experimentally achieved four years later using a photonic scheme (Bouwmeester *et al.*, 1997). The transient nature of such qubits disallows real-time feed-forward, so rather than applying the correcting operations they keep only the 25% of data where the correction is the identity. This means that this experiment is probabilistic. The first deterministic state teleportation was achieved seven years after that, using trapped ion qubits (Riebe *et al.*, 2004).

### 7.3 Teleportation of quantum gates: theory

In 1999 it was shown that one can also teleport a two-qubit gate (Gottesman and Chuang, 1999). This functionality might be quite useful when two qubits cannot directly interact. The original proposal is somewhat complicated, requiring six physical qubits and the preparation of a highly-entangled four-qubit state as a resource. As with state teleportation, the first (and as far as I can tell only) teleported two-qubit gate was also performed with a photonic system (Huang *et al.*, 2004). This experiment was also probabilistic. To my knowledge a two-qubit gate has never been deterministically teleported, until now. The circuit which have chosen to implement is based on (Jiang *et al.*, 2007), which requires only the two “memory” qubits plus an additional “communication” Bell pair.

If the control and target states are prepared in a separable state written generally as  $|\psi\rangle = \alpha|00\rangle + \beta|01\rangle + \gamma|10\rangle + \delta|11\rangle$ , we can work through the circuit (where we order the tensor products by memory




Figure 7.3: **Basic CNOT teleportation circuit.** The two data qubits each locally interact with half of an additional entangled pair, forming two modules (blue and red in parallel with orange and green). The qubits of that pair are measured in the  $Z$  and  $X$  axes. The results of those measurements are communicated classically to the other module whether they control single-qubit feed-forward. The gray dashed line indicates that after the Bell state is introduced we need only local operation and classical communication.

1, memory 2, communication 1, communication 2),

$$\begin{aligned}
 |\Psi\rangle &= |\psi\rangle (|00\rangle + |11\rangle) \\
 &\mapsto \alpha(|0000\rangle + |0111\rangle) + \beta(|0100\rangle + |0011\rangle) + \gamma(|1010\rangle + |1101\rangle) + \delta(|1110\rangle + |1001\rangle) \\
 &\mapsto \alpha(|0000\rangle + |0001\rangle + |0111\rangle - |0110\rangle) + \beta(|0100\rangle + |0101\rangle + |0011\rangle - |0010\rangle) \\
 &\quad + \gamma(|1010\rangle + |1011\rangle + |1101\rangle - |1100\rangle) + \delta(|1110\rangle + |1111\rangle + |1001\rangle - |1000\rangle) \\
 &= ((Z \otimes I)\text{CNOT}|\psi\rangle)|00\rangle + \text{CNOT}|\psi\rangle|01\rangle \\
 &\quad - ((X \otimes X)\text{CNOT}|\psi\rangle)|10\rangle + ((I \otimes X)\text{CNOT}|\psi\rangle)|11\rangle
 \end{aligned}$$

We can see that if we project the two communication qubits (above shown as the last two positions in the tensor order) we have achieved a CNOT operation, potentially followed by single-qubit operations. These extra operations are determined by which measurement outcome occurs, each of which should happen with equal likelihood. In this way real-time feed-forward is crucial to deterministically perform the same operation in all cases<sup>2</sup>.

<sup>2</sup>Note that for many error-correction purposes we do not strictly need to perform these gates, or at least not immediately. We can simply keep track of them and use their knowledge to inform the meaning of later stabilizer measurement (Fowler *et al.*, 2012, Terhal, 2015).



## 7.4 Characterization results

We denoted the two “modules” as 1 and 2, coupled by the bus resonator. The relaxation lifetimes of the two memory cavities were 1.15 and 0.8 milliseconds. The two transmons had  $T_1$ s of 41 and 52  $\mu$ s respectively, with  $T_2$ s of 15  $\mu$ s each which did not improve under echo, so we can likely ascribe these decay times to photon shot-noise dephasing. The readout resonator resonators had lifetimes on the order of 200 nanoseconds with dispersive shifts to their transmons of 2.5 and 2.8 MHz. The transmons had dispersive shifts  $\chi$  to their memory resonators of 615 and 1006 kHz, with nonlinear  $\chi'$  of 1.1 and 2.8 kHz. The memory resonators inherited self-Kerrs of 1.2 and 1.3 kHz. The bus resonator coupled to the two transmons with dispersive shifts of 319 and 445 kHz and inherited a self-Kerr of 0.3 kHz. It had a relaxation time of roughly 250  $\mu$ s, which is very close to the lifetime targeted by our choice of couplers<sup>3</sup>.

### 7.4.1 Transmon excited-state populations

Both transmons also had very large natural excited-state populations of almost 20%—larger than our bound on the readout cavity populations. This is not uncommon when using Purcell filters in 3D cQED (e.g. the devices in (Narla *et al.*, 2016) had natural excited-state populations between 20 and 30%.) The full mechanism of this effect is not completely understood. It has been seen that transmon populations decrease as they become Purcell limited, since the output ports represent a relatively cold bath, but the effect is slightly more subtle than that. Transmons which are not Purcell-limited in systems without Purcell filters generally still exhibit lower natural polarizations than ones with Purcell filters. This implies that the Purcell filter increases the temperature of the bath that causes the “internal” loss of the qubit, e.g. defects in the surface dielectric, which ordinarily thermalize (at least in part) via the output line. In 2012 Adam Sears and Andrei Petrenko used a system where they could in-situ vary the  $Q$  of a cavity and took data to support this. The transmon-population data were not included in their publication (Sears *et al.*, 2012), but were presented at a conference (Petrenko *et al.*, 2012). They saw mild evidence (figure 7.4) to support that as they retracted the coupling pin the transmon population seemed to continue

<sup>3</sup>Performing the RIP gate with high performance requires a large amount of power, so we opted for an overcoupled bus resonator.

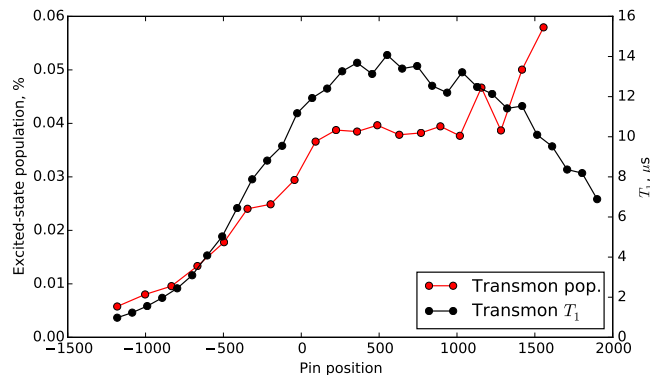


Figure 7.4: **Relationship between transmon Purcell effect and polarization.** In this experiment the cavity coupling pin could be retracted *in situ*. The transmon cooled as it became Purcell-limited. In the other direction, the transmon appears to continue heating after it ceases to become Purcell-limited. It is not clear why the transmon  $T_1$  begins decreasing in that direction. Data by Andrei Petrenko, reproduced with permission.

increasing after its relaxation time stopped changing, though with a different scaling behavior.

## 7.5 Simultaneous readout & feedback

As shown in figure 7.3, measurement of the communication qubits (transmons) is an essential part of the teleportation circuit. The fidelity of the resulting states will depend directly on the assignment fidelity of those measurements, which control the later feed-forward. If that feed-forward is applied mistakenly we will end up in an orthogonal state. Since we are going to reuse the transmons for tomography (and because our single-memory-qubit operations will actually depend on the transmons being the ground state), we will also need the measurements to be non-demolition.

In the final version of this experiment using the Dodekathion device (section 4.2.3) the syndrome qubits were measured simultaneously using two JPCs (Bergeal *et al.*, 2010). We gated the pump of the JPC attached to readout 2, as if left on continuously we observed a decrease in the  $T_2$  of qubit 2. The exact mechanism of that process is not understood. We found that we were able to get sufficient discrimination of the states of both transmons with a 600 nanosecond long pulse with 600 nanoseconds of acquisition time, along with another 400 nanoseconds to allow the cavities to empty. It is worth noting that we actually reset the transmons before that decay finished—we reset them immediately after

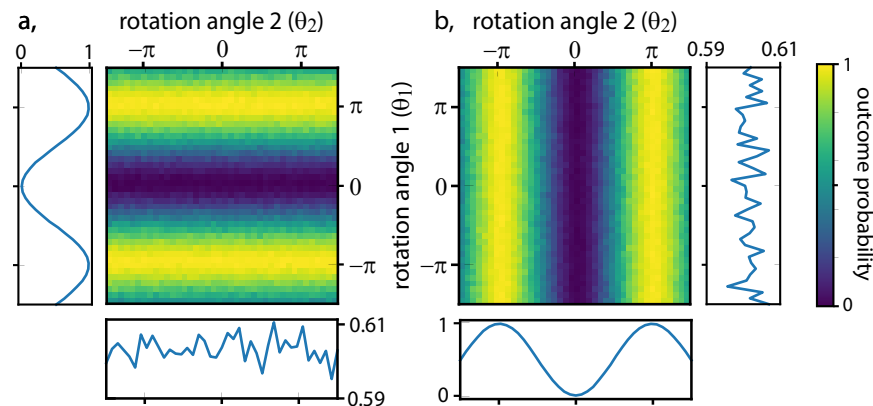


Figure 7.5: **Simultaneous Rabi oscillations and measurement.** In this experiment we rotate transmons 1 and 2 around the  $X$  axis by angles  $\theta_1$  and  $\theta_2$ , respectively, then perform simultaneous measurements. These angles are swept, forming a two-dimensional preparation space. In figure **a** we see the average probability of “g” from readout 1, and in figure **b** we see the average probability of “g” from readout 2. These probabilities are not scaled in any way. For both plots we show the marginals along each axis. These show that both readouts have high sensitivity ( $\approx 99\%$  fidelity) to the expected qubit preparation, and low sensitivity (under 0.5% and indistinguishable with 0% contrast) to the other qubit.

the acquisition and additional 220 nanoseconds needed for the integration logic.

The assignment fidelity for both qubits was roughly 99%, though it drifted in time. This was principally limited by qubit relaxation and heating. As shown in figure 7.5, we see essentially negligible cross-talk on our readouts (less than 1%, but indistinguishable with 0%). In the worst case scenario, where any decay or excitation occurring during the integration time gives the wrong answer, this leads to a fidelity of  $\exp(-0.6/40) = 98.5\%$ . If we more reasonably state that a decay event gives you a random answer, this result can be approximated to 99.3%, essentially consistent with what we observe. For technical reasons we cannot presently confirm the quality of our  $\pi$  pulses to that accuracy, though based on results from similar systems (table 6.2) we expect that fidelity to approach three nines.

### 7.5.1 Synchronous feedback

As we have two bits to feed back on, we perform feedback in a slightly more complicated way than previous experiments. The FPGA card for readout 1 uses its integrator as its initial internal decision function, then outputs a digital signal to of the other three cards ( $x_0$ ). The FPGA for readout 2 similarly




Figure 7.6: **“Black mamba” wiring in the gate teleportation experiment.** We use seven output IQ pairs and demodulate two incoming signals on two different cards. The first readout results feeds back to the  $x_0$  binary input on all cards. The second readout result feeds back to the  $x_1$  input on all cards. All four cards synchronously branch on  $x_0$  then change their internal decision function and branch again on  $x_1$ .

integrates and sends a signal to the other three cards ( $x_1$ ). All four cards synchronously branch on the first signal, change their internal decision functions, then branch on the second signal. The first decision and jump takes 24 nanoseconds. Changing the internal decision function takes 48 nanoseconds, which is done simultaneously with the second branching step of 24 nanoseconds. All of this follows the 220 nanoseconds required after acquisition ends to finish integration, leading to a minimal feedback time of 292 nanoseconds.

### 7.5.2 Ground state preparation by feedback

The first use of feedback in our system is to prepare the cavities and transmons in the ground state. I have schematically sketched our cooling sequence in figure 7.7. This sequence is similar to what we used (post-selectively) in section 6.4. First, following the preceding experimental iteration, we perform a sequence of four-wave mixing Q-switches (section 3.7.2). These target all three high-Q cavities. After

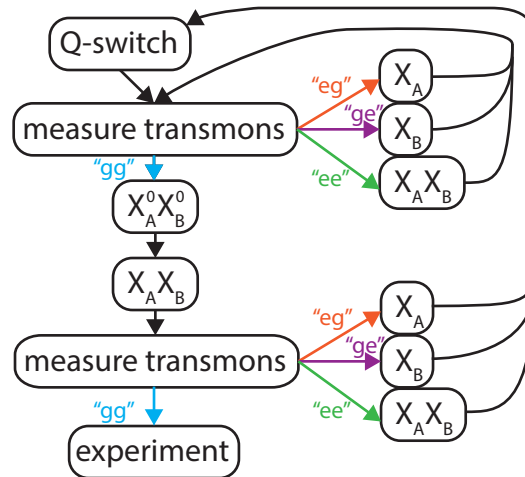


Figure 7.7: **State preparation sequence in the gate teleportation experiment.** After the previous experiment we apply a set of Q-switches to empty the cavity. We then check that both transmons are in the ground state. If not we reset the transmons and measure again. Once we have concluded the transmons are in the ground state, we check if all three high-Q cavities are in the ground state. If not, we return to the Q-switch sequence and start over.

that we measure both transmons. If either of them are not in the ground state, we perform the relevant  $\pi$  pulse(s) to reset them and measure again. This is repeated until both measurements yield “g.”

After being assured that both transmons are in the ground state, we perform  $\pi$  pulses which are selective on all cavities being empty, then an unconditional  $\pi$  pulse. If we find that either of the qubits are excited, we conclude that one of the cavities was not in the ground state, reset the transmons then return all the way back to the initial Q-switch sequence. If both measurements report “g” we continue. To quantify the quality of this preparation, if we measure again we get the result that both are in the ground state 98.7% of the time. If we assume that measurements of both transmons behave the same and that there is no heating, this implies a histogram separation of 99.3%. Along with the expected qubit decay, this is quite consistent with our observed assignment fidelity. The cavity excited-state populations were difficult to measure with certainty, but were less than 1% excited. In practice we actually use a slightly more complicated version of the above loop where we verify success a variable number of times in a row, though the gains over a single round of preparation are small.




Figure 7.8: **Tomography of transmon reset in the gate teleportation experiment.** After measurement of the transmons, our protocol calls for them to immediately be reset to the ground state unconditionally. We plot here state tomography of the transmons, conditioned on and labeled by the output branch, followed by the unconditional result. The ground-state fidelities from left to right are 99%, 96%, 97%, 94%, and 97%.

### 7.5.3 Transmon reset

Our full realization of the gate teleportation experiment (seen in figure 7.13) requires that we reuse the transmons for tomography of the memory qubits. Additionally, our single-qubit operations on the cavities require that the transmons are in the ground state. Unlike the cooling sequence in section 7.5.2, we cannot repeat until both transmons are measured to be in the ground state, so we must count on our measurement of “e” being QND. Unfortunately, this is not precisely the case. Our histogram overlap aside, the data acquisition and feedback delay take nearly a microsecond. This leads to an approximately 2% probability of decay per qubit. Additionally, we have clearly observed the degradation of  $T_1$  with readout resonator  $\bar{n}$  effect on transmon 1. These numbers roughly explain the 2-3% degradation in ground state fidelity when we measure that one of the transmons was excited and the 6% degradation when both were. State tomography of the transmons after reset is plotted for all outcome paths in figure 7.8.

## 7.6 RIP gates

Though the greatest usefulness of teleported gates is for non-local interactions, we opted to first demonstrate the gate using a simpler scheme with locally-generated entanglement. Our resource entangled bit is a Bell pair of two transmons with entanglement generated by a resonator-induced phase (RIP) gate.




Figure 7.9: **Spectral explanation of the RIP gate.** Figure **a** depicts the spectrum of the bus cavity, with frequencies below  $\omega_C$  indicating that one or both transmons are excited. The dashed blue line represents the drive frequency, positively detuned from the cavity by  $\Delta$ . Figure **b** shows the dependence of the effective off-resonant displacement with detuning. We can see that in the region corresponding to the relevant frequencies a linear approximation does not hold well. In figure **c** we see the frequencies of the qubit manifold, where blue lines represent the Stark-shifted frequency. Because of the non-linearity evident in figure **b**, the Stark shift on  $|ee\rangle$ ,  $\Delta^{ee}$ , generally does not equal  $\Delta^{eg} + \Delta^{ge}$ . This results in an entangled state.

This gate was first developed at Yale by Dong Zhou and Steve Girvin, and experimentally implemented in 2013 by Hanhee Paik. The first theory publication was by IBM in (Cross and Gambetta, 2015) and the first experimental publication was by Hanhee Paik (Paik *et al.*, 2016), who at that time was also at IBM. It is also worth noting that similar gates had been developed earlier in the context of cavity QED with quantum dots, using different language but describing similar ideas (Clark *et al.*, 2007).

We can first look at the broad picture of the RIP gate. One applies a detuned tone above the cavity. The virtual photons, or off-resonant excitation of the cavity, cause a Stark shift of the qubit transitions by an amount roughly equal to the dispersive shift times the effective number of off-resonant photons. The key idea of the RIP gate is that the displacement magnitude depends nonlinearly on detuning,  $\xi \propto 1/\Delta_i$ . This means the doubly-excited state is Stark shifted by an amount different than the sum of what the two singly-excited states are shifted by, leading to a non-separable phase.

With that road-map in mind, we can examine the behavior of the RIP gate in more detail. This starts by considering the behavior of a classical driven-damped harmonic oscillator. Given a complex voltage envelope  $\epsilon(t)$  driven at a frequency detuned  $\Delta$  away from the cavity resonance, the response of

the oscillator is given by a convolution with its response function<sup>4</sup>,

$$\alpha_{\Delta} = \frac{-i}{2} \int_0^t \exp \left\{ \left( i\Delta - \frac{\kappa}{2} \right) (t - t') \right\} \epsilon(t') dt'. \quad (7.1)$$

Since we plan on staying in the adiabatic regime, we can examine the steady-state limit where  $\epsilon(t) = \epsilon$ .

We can perform that integral and let  $t \rightarrow \infty$  to obtain

$$\alpha_{\Delta}^{t \rightarrow \infty} \equiv \xi_{\Delta} = \frac{-i\epsilon}{2(i\Delta + \kappa/2)}. \quad (7.2)$$

In our work we are essentially in the lossless limit, since our energy decay time is roughly  $250\mu\text{s}$ . We will not look at the lossy case which describes the 'measurement-induced dephasing' of that results from doing a RIP gate, but the key result is that for an entangled phase of  $\pi$  that effect is fixed with a scale given by  $\kappa/\Delta$ . That ratio is minuscule in our system ( $\approx 2 \cdot 10^{-4}$ ). In the lossless limit, we get that  $\xi_{\Delta} \approx \epsilon/(2\Delta)$ . Following section 3.5, we can derive the effect of the AC Stark shift on the coupled transmons. The shifted eigenstates of the qubit manifold will gain a phase proportional to  $\int \Delta_{\text{Stark}}(t)dt$ . If the drive couples only to the cavity<sup>5</sup>, in the near-resonant but still adiabatic limit  $\Delta_{\text{Stark}} \approx \chi|\xi_{\Delta}|^2$ . For simplicity we will consider the two qubit case<sup>6</sup>. We can label the (detuned) coherent states associated with each state in the qubit subsystem by  $\xi_{ij}$  where  $i, j \in \{0, 1\}$ . For now we can assume that the qubits have equivalent  $\chi$  shifts and that the drive is detuned above the  $|gg\rangle$  peak by  $\Delta$ . In this steady state limit we obtain

$$\xi_{01/10} = \frac{\epsilon}{2(\Delta + \chi)}$$

$$\xi_{11} = \frac{\epsilon}{2(\Delta + 2\chi)}$$

<sup>4</sup>Note that this is of course the same way we discuss measurement in cQED (equation 3.40). This connection can be useful when considering the RIP gate.

<sup>5</sup>We can analyze how reasonable that assumption is or is not using the tools from section 3.7.1.

<sup>6</sup>In early versions of this experiment we did perform RIP gates in cavities with more than two qubits. I do not recommend it. The requisite echo sequences (as discussed in (Paik *et al.*, 2016)) do not scale well.



Leading to shifts

$$\begin{aligned}\Delta_{01/10} &= \frac{\chi|\epsilon|^2}{4(\Delta + \chi)^2} \\ \Delta_{11} &= \frac{2\chi|\epsilon|^2}{4(\Delta + 2\chi)^2}\end{aligned}$$

In the adiabatic limit we can integrate this shift for time  $t$  to obtain the total acquired phase. We will pay particular attention to the quantity  $\phi_e = \phi_{11} - \phi_{10} - \phi_{01} + \phi_{00}$  which corresponds to the non-separable *entangling* phase:

$$\phi_e = \frac{\chi}{2\Delta^2} \left( \frac{1}{(1 + 2\chi/\Delta)^2} - \frac{1}{(1 + \chi/\Delta)^2} \right) \int |\epsilon(t)|^2 dt. \quad (7.3)$$

For simplicity we will look at the  $\Delta \gg \chi$  limit, and obtain

$$\phi_e = \frac{\chi^2}{2\Delta^3} \int |\epsilon(t)|^2 dt. \quad (7.4)$$

The entangling effect scales as the power of the applied drive, is linear in time, and essentially scales with the self-Kerr of the cavity. In addition to the  $1/\Delta^2$  frequency scaling which relates the applied power to the off-resonant excitation energy, the gate has a  $1/\Delta$  term. One can consider that the extra  $1/\Delta$  results from the gate relying on a derivative.

It is an outstanding question in the field as to what fidelities may be obtained with this gate. In (Cross and Gambetta, 2015) they perform an analysis the adiabaticity of this gate under driving. From personal conversations with the IBM group I have been told that there is a fundamental scaling concern, and that non-adiabaticity places a difficult bound on the obtainable fidelities. To examine this carefully one must derive the integrated phase without taking the adiabatic approximation and see if that can be equal to  $\pi$  while the resulting photon population approaches 0. My attempts have left me to suspect this is better done as a numerical exercise, but that to do it carefully the numerics must involve photon loss and higher order terms of the Hamiltonian.

### 7.6.1 Refocused RIP gates

One of the challenges faced in the gate teleportation protocol is the always-on interaction between the transmons and the cavities. That coupling can be quite detrimental at points in the experiment when both carry meaningful quantum states. Since a RIP gate takes several hundred nanoseconds and  $1/2\pi\chi \sim 1\mu\text{s}$ , a naive RIP gate will not result in a high-quality transmon-transmon Bell pair. Fortunately, as seen in chapter 6, we can dynamically remove this coupling with an echo sequence. This was explored in a somewhat different context in (Paik *et al.*, 2016). Their interest was decoupling additional qubits, but the idea is the same. In our case we will do two “ $\sqrt{\text{RIP}}$ ” gates that each give an entangled phase of  $\pi/2$ . Between the two we can  $\pi$  pulse both transmons. This does not echo away the transmon-transmon entangling phase since we invert the state of both qubits together, but it does remove the entangled cavity phases. Accordingly it performs a spin-echo with respect to the environment as well. It also fixes the single qubit phases:

$$\begin{array}{llll}
 \arg |\psi\rangle \mapsto 0, & \phi_1, & \phi_2, & \phi_1 + \phi_2 + \pi/2 \\
 \mapsto \phi_1 + \phi_2 + \pi/2, & \phi_2, & \phi_1, & 0 \\
 \mapsto \phi_1 + \phi_2 + \pi/2, & \phi_2 + \phi_1, & \phi_1 + \phi_2, & \phi_1 + \phi_2 + \pi/2 \\
 = 0, & -\pi/2, & -\pi/2, & 0
 \end{array}$$

which is a CPHASE gate with additional single qubit phases of  $-\pi/2$ .

The tuneup experiment for this gate is to choose a drive detuning, shape, and amplitude, then perform two experiments. The base of both is a fixed-time Ramsey experiment  $-\pi/2$  on one transmon about the  $X$  axis, the RIP sequence, then a second  $\pi/2$  about a swept axis. This experiment is performed both with and without the second transmon initially excited. When the gate is tuned up correctly, the two traces should be cosine waves separated in phase by  $\pi$ , and for refocused RIP gates they should be shifted from the origin by  $\pm\pi/2$ . It is important not to start with too large an amplitude, lest one be tricked by an odd multiple of  $\pi$ . For a set detuning and pulse shape we repeat this for several amplitudes




Figure 7.10: **Tomography of a Bell pair of transmons.** This Bell state state was prepared by performing  $X_{\pi/2}$  rotations on both transmons, followed by a refocused RIP gate and an additional  $\pi/2$  pulse. It has a fidelity of roughly 95%, and exhibits a slight phase error on the second qubit. The remaining infidelity appears incoherent.

and fit to a parabola to find the optimal drive amplitude which gives a  $\pi$  entangled phase. After this one can sweep the detuning or shape to attempt to maximize the contrast of these fringes (and the fidelity of the operation).

In this experiment we were able to obtain Bell-state fidelities of roughly 95%, determined from tomography, which is essentially consistent with  $T_2$ . The two RIP sections took 220 nanoseconds each at a detuning of  $\Delta = 20$  MHz. This was limited by pushing the drive amplitude as high possible while still having a reasonably clean and understood signal. Each  $\pi$  pulse was 24 nanoseconds long. Given the total gate time and our qubit  $T_2$ s, we expect a fidelity of 94%. Though I am not presenting data here, we have verified that this Bell state fidelity is largely insensitive to the prepared cavity state, as expected. Note that our dispersive shifts are well under 0.5 MHz, so the entanglement happens at a rate much faster than  $\chi$ . Similarly our direct qubit-qubit coupling provides entanglement on the timescale of  $25\mu\text{s}$ , yielding an effective on/off ratio of 1/50.

## 7.7 Qubits encoded in harmonic oscillators

Given that our memory or logical qubits are going to be encoded in harmonic oscillators, we have several choices for the subspace or subsystem in which to store that information. At the simplest end of the spectrum, we can encode our qubits in the  $|0\rangle$  and  $|1\rangle$  states of the cavity, which we call the “Fock encoding.” Most of the data prepared at the time of this writing are of the experiment using this

encoding. This encoded qubit is quite simple, but not qualitatively that different from simply using transmons as the memory qubits. Instead, we would rather use an encoding that allows some sort of error resistance or allows for active error encoding.

One of the earliest proposals was the “GKP states” (Gottesman *et al.*, 2001). While quite interesting, they are likely to be difficult to manufacture with our present optimal control techniques as they need to access very high photon numbers. A more recent proposal, which we briefly mentioned in section 6.1 and pictured in section 2.4, is the “cat code” (Leghtas *et al.*, 2013b). In this protocol information is stored in a superposition of either two or four (or perhaps more) coherent states. It protects against photon-loss using photon-number parity as the error syndrome. We have recently demonstrated high-fidelity (>99%) single-qubit operations in this encoding (Heeres *et al.*, 2016). This encoding has the rather impressive claim to fame that it has actually been used to extend the lifetime of quantum information (Ofek *et al.*, 2016). There may also be ways to do some of the error correction in the cat code autonomously (Leghtas *et al.*, 2015), and Touzard *et al.* have recently demonstrated control of one logical axis without leaving the protected codespace (in preparation). For all of these encodings, we have also recently demonstrated routines which will allow local SWAP and  $\sqrt{\text{SWAP}}$  gates using a coupling transmon (Gao *et al.*, in preparation).

### 7.7.1 Binomial encoding

We have chosen to operate with a different code—one of the binomial encodings (Michael *et al.*, 2016). These codes are interesting for several reasons. First, in exchange for a larger Hilbert space, they can be tailored to correct more errors, including dephasing, photon gain, and even higher-order terms in the Hamiltonian. Second, they use a strictly bounded Hilbert space which I suspect will make leakage detection more straightforward. We have chosen to initially demonstrate the simplest binomial code, which we call the ‘024 code.’ It is quite similar to a rotated  $\bar{n} = 2$  even two-cat code and has codewords

$$|0_L\rangle = |2\rangle \tag{7.5}$$

$$|1_L\rangle = \frac{1}{\sqrt{2}} (|0\rangle + |4\rangle). \tag{7.6}$$

Those states, and the other cardinal states, are plotted in figure 7.11. Clearly these are even parity states, and parity will be our syndrome for photon-loss. Both computational states also have the same average photon number, which is crucial since they need to have the same rate of photon-loss—otherwise a photon-loss event gives partial information about the logical state. The syndrome check is the same as in the cat code. We can also see that there is a simple correction unitary<sup>7</sup>

$$U = |0_L\rangle\langle 1| + |1_L\rangle\langle 3|. \quad (7.7)$$

Interestingly, the absence of an error also requires a “correction.” Though we gain no information with respect to whether the system is in  $|0_L\rangle$  or  $|1_L\rangle$ , the logical state  $|1_L\rangle$  itself changes. As the  $|0\rangle$  state cannot undergo photon loss, the logical state shifts to a larger amplitude in that component. This process is deterministic and can easily be accounted for, as described in (Michael *et al.*, 2016).

We can also immediately see that there is a very inexpensive way to measure the  $Z$  projection, which is to run the standard parity check for half of the time, such that  $|0\rangle$  and  $|4\rangle$  lead to a transmon phase  $\pi$  away from  $|2\rangle$ . This will form our CNOT which targets the transmon.

### 7.7.2 Operation times

One interesting question which is relevant when choosing between the various possible encodings is how “hard” the logical operations are. In general, this means how long of a pulse duration GRAPE needs to give a satisfactory process fidelity. Here I will provide my results comparing three different encodings. I am not providing an exhaustive analysis, and I cannot promise that I have truly found the optimal result. For all of these I am using the same Hamiltonian: a transmon and cavity with a dispersive shift of -1005 MHz, a cavity self-Kerr -1.3 kHz, and a  $\chi'/2\pi$  of +1 kHz. These are performed with the cavity Hilbert space truncated to between 23 and 27 levels, and with a penalty for discrepancies between results in the different spaces. Additionally all results were verified with a simulation of 22 level Hilbert space. We applied amplitude and derivative penalties which were roughly equal for all pulses. For the CNOT operations I allow gauge freedoms of phases on both the transmon and the cavity. For the decoding

<sup>7</sup>Note that we are only describing this unitary in terms of its action on the relevant subspace. The full, physically realized unitary will have actions on other initial states which we do not consider.




Figure 7.11: **Wigner tomographs of logical states of the 024 code, arranged on a Bloch sphere.** These represent Wigner tomography of the logical  $|\pm Z\rangle$ ,  $|\pm X\rangle$ , and  $|\pm Y\rangle$  states. All have red centers, indicating even parity. Note that the logical  $Z_\pi$  operation consists of a simple shift of frame.  $Z$  rotations about smaller angles, and all  $X$  rotations, are non-trivial. The logical  $|\pm X\rangle$  states approximate the two-cat code. These data are simulated.

operations I allow many gauge freedoms, equal to all possible unitary operations on the cavity within the bounded Hilbert space. This is discussed further in section 7.8.3. As a bound for acceptable fidelity I generally attempt to be over 0.994 without decoherence<sup>8</sup>. The four-cat operations were performed with  $\bar{n} = 2$ , which may be too small for practical operation.

### 7.7.3 Gate performance

We have analyzed the performance of our optimal control gates using interleaved randomized benchmarking (Magesan *et al.*, 2012) for the Fock encoding and ordinary randomized benchmarking for the 024 encoding. For the Fock encoding operations about the  $X$  and  $Y$  differ only by a change of drive phase. Our base set of benchmarking operations was  $I$ ,  $R_x(\pi/2)$ ,  $R_y(\pi/2)$  and  $R_x(\pi)$ . By benchmarking with the FPGA, we do not need to choose a finite set of sequence lengths and a number of averages, rather we spontaneously generate a new random sequence every iteration.

The raw data for logical qubit 2 are shown in figure 7.12. These  $\tau$  correspond to gate fidelities of

<sup>8</sup>Or 0.997 as reported by our GRAPE implementation, which reports  $\sqrt{\mathcal{F}}$

Table 7.1: **Empirical comparison of operation times in various encodings.** All times in nanoseconds. Performed for specific Hamiltonian parameters noted in the main text with consistent amplitude and derivative constraints. Each operation targeted a nominal  $\sqrt{\mathcal{F}} \approx 0.997$ . Four-cat operations used  $\bar{n} = 2$ . Fock encoding does not require an initialization pulse. We did not implement decoding pulses for the Fock and four-cat encoding, rather we used single-qubit operations followed by a CNOT with targets the transmon. Some gaps have been filled operations generated for a different qubit, scaled by the ratio of  $\chi$ s. Time of rotations around the  $Y$  axis are consistent with  $X$  operation times. The greater-than sign implies an acceptable operation time was not yet found. The less-than sign implies we reached a sufficient fidelity but did attempt shorter times. Numbers preceded with  $\sim$  are approximated from results above and below. CNOT operations have gauge freedoms to allow phases on both transmon and cavity (physical, not encoded). Decoding operations map the relevant information onto the qubit state, then have gauge freedoms allowing arbitrary bounded unitary operations on the cavity.

Operation	Fock	Four-cat	024
g init	N/A	340	380
e init	N/A	<420	$\sim 600$
+x init	N/A	420	360
$X_\pi$	480	260	840
$X_{\pi/2}$	440	480	720
CNOT: target qubit	500	>1400	260
CNOT: target cavity	$\sim 900$	–	$\sim 1600$
decode X	–	–	$\sim 500$
decode Y	–	–	$\sim 400$

98.5% for  $R_X(\pi/2)$  and 97.4% for  $R_X(\pi)$ . Our master equation simulations which include transmon  $T_1$  and  $T_2$  estimate we should observe a fidelity of 98.4% for the  $\pi$  pulse. Note that this effect is essentially equal to  $\exp(-0.5 * t_{\text{gate}}/T_2) = 0.984$ . for our gate time and  $T_2$ . The remaining 1% of fidelity is due to some combination of Hamiltonian mischaracterization, line dispersion, mixer or DAC nonlinearity, amplitude miscalibration, and etc.

For logical qubit 1 we have not performed interleaved randomized benchmarking, but for regular benchmarking we have observed  $\tau = 37$  operations. Given that we have almost equivalent transmon  $T_2$ s, we might expect the ratio of benchmarking time constants to be equal to the ratio of pulse lengths. What we observe is  $\exp(-.64/15)/\exp(-.48/15) = 0.989$  and  $\exp(-1/37)/\exp(-1/55) = 0.991$ , implying that the differing performance between the two qubits is close to what we would expect simply from decoherence.

With the 024 encoding we required single-qubit gates which were almost twice as long: 840 and 720 nanoseconds for the  $\pi/2$  and  $\pi$  pulses respectively. For qubits 1 and 2 we observed  $\tau$ s of 23 and 16.




Figure 7.12: **Interleaved randomized benchmarking of operations on a Fock-encoded qubit in memory cavity 2.** For the Fock encoding one can switch between  $X$  and  $Y$  operations simply by changing the drive phase. This benchmarking included  $I$ ,  $R_x(\pi/2)$ ,  $R_y(\pi/2)$  and  $R_x(\pi)$  as the base operation set. Some real data points will result in negative  $P_{\text{correct}} - 1/2$ . These are not depicted in our log plot, but were included in the fitting, which was done in linear units. Memory 1 has a 40% smaller  $\chi$  and similar transmon coherence. Our data for that logical qubit are consistent with the ratio of  $\chi$ s.

These are quite a bit worse than we would expect from decoherence, implying that we are further from correctly implementing our unitaries. Note that of course both the Fock and 024 operations were done with the same actual and estimated Hamiltonian, so the excess infidelity is likely simply in miscalibration.

#### 7.7.4 Leakage/codespace errors

One significant concern for these encoded qubits is leakage out of the logical codespace during operations. I hope that on this experiment we will soon begin examining those errors. We know quite well that during the operations generated via optimal control techniques (“optimal control pulses”) the cavity takes excursions to quite high photon numbers. As an illustrative example—even with Hilbert space truncations of 23 photons we have to be careful that the pulse is not taking advantage of that boundary during GRAPE. If the operation is imperfect (imperfect Hamiltonian characterization, line dispersion, mixer saturation, crosstalk...) or a decohering event occurs (most likely transmon dephasing or relaxation), we should not expect that the cavity will end up in the logical space. The success of this architecture will likely depend on our ability to detect these errors or perhaps to autonomously correct them (Touzard, in preparation).

If we can detect them, we may reset that cavity and declare an erasure error. Erasure errors are actually preferable to bit-flip and phase-flip errors. We know that any code which can correct  $k$  errors



at unknown locations can correct  $2k$  erasure errors (Gottesman, 2005, Terhal, 2015). The challenge becomes looking for a manner to fault-tolerantly detect such errors. In practice this means that the detecting operation must not induce codespace errors to a state that was initially in the codespace. We can tolerate most other forms of imperfection. If the odds of correct detection are less than one but the measurement is “QND” in that manner, we can safely repeat the measurement to get sufficient confidence. Simply asking for an optimal control pulse which asks “are we in the code space?” is possible, but likely will not have the QND property we demand.

We design these optimal control operations to begin and end with the transmon and the ground state, so the first and simplest check is to simply measure the transmon. Chris Wang has done simulations of our optimal control pulses in the presence of qubit decoherence. One of the results of these simulations is that, as might be expected, roughly half of our codespace errors leave the qubit in the excited state. He has further shown that if we were able to perfectly detect and herald away leakage errors, many of these operations should have fidelities of almost three nines.

We can likely suppress leakage errors to the level of those induced by our “best” operation. One can simply apply that operation, then measure if the transmon is excited. If we were not initially in the code space, we would expect there to be a 50% chance of leaving the transmon excited, even if no error occurs during the new pulse. This can be repeated as many times as desired. One concern I have is whether or not an erasure error informs the initial state. It seems as though this will generally be true. e.g., If  $X_L |0_L\rangle$  has a trajectory in which the qubit is excited with greater likelihood than in the trajectory of  $X_L |1_L\rangle$ , then an error caused by qubit  $T_1$  leads to partial logical information. This may be worth further investigation.

## 7.8 Two-bit, phase-corrected tomography

State tomography in this experiment has several differences from tomography as explained in section 6.7. The first is the number of tomographic measurements. In the stabilizer measurement experiment we used one tomographic measurement for all three qubits. In this experiment we measure both transmons simultaneously. Second, due to the dispersive shift with the coupled transmons during the internal

measurements, the two cavities acquire different physical phases dependent on measurement path. This is *not* a logical phase, and cannot simply be accounted for afterwards in tomography<sup>9</sup>. Third, we are doing tomography of the cavity states, which we cannot directly measure. We might consider using a SWAP operation to map these states directly to the transmons before performing tomography, but we have found empirically that these SWAP operations can require very long pulses. Instead we have chosen to perform single-(logical)-qubit operation on the cavity followed by a CNOT targeting the transmon<sup>10</sup>. We have also tried implementing “decoding pulses,” which compile the single-qubit rotations and the CNOT into their own pulses. These have potential to be faster operations since they need not be QND, so we can allow additional gauge freedoms. However, using multiple measurement operators has trade-offs since it adds complexity.

### 7.8.1 Two-bit tomography

In many ways two-bit tomography is simpler than joint tomography. We directly measure  $Z$  information for both qubits and can reconstruct the correlation  $ZZ$  in post-processing. We formulate this by using a POVM of not two operators, but three. For each post-rotation on our unknown state we take away three *correlated* outcomes and corresponding probabilities,  $P_{10''}$ ,  $P_{01''}$ , and  $P_{11''}$ . We also calibrate three different POVM elements for our measurement operator, each of which is ideally a matrix of zeros with one element on the diagonal equal to one. Our reconstruction is the same as before, except instead of  $m$  probabilities for our  $m$  post-rotations we stack the individual vectors and fit  $3m$  probabilities.

This also changes the number of post-rotations required. In the previous regime we would have needed 16 post-rotations, here we need only 9. To understand this, consider that when we perform no post-rotation, we measure  $ZI$ ,  $IZ$ , and  $ZZ$  simultaneously (Chow *et al.*, 2010). Previously to obtain  $ZZ$  we would also have needed two additional measurements where we inverted each qubit. Essentially now we need only measure all operators in  $\{X, Y, Z\}^{\otimes n}$  but before we needed  $\{I, X, Y, Z\}^{\otimes n}$ . Since we are redundantly measuring the single-qubit properties, the two-qubit correlations have noise with standard deviations  $\sqrt{3}$  times larger.

<sup>9</sup>Except for the Fock encoding, where a physical phase and a logical phase are equivalent.

<sup>10</sup>Recall that these CNOT operations are “cheap” for the Fock and 024 encodings. They are a parity measurement or a parity measurement run for half the time, respectively.

### 7.8.2 Phase correction

In the current iteration the time between the start of a measurement pulse and the end of qubit reset is roughly 800 nanoseconds. In that time the dispersive shift will lead to the cavities gaining a very large phase if the transmon was excited. Additionally, we have other unconditional cavity phases due to the CNOT operation and the Bell-state preparation<sup>11</sup>. In the Fock encoding these are not particularly important, as a physical cavity phase due to the dispersive shift is equivalent to a logical phase. In that case all optimal control operations will still basically work and we will get meaningful, if not ideal, tomography. In other encodings these phases are essentially codespace errors if not accounted for.

Fortunately we can measure these phases using fiducial states and Wigner tomography and adjust the phase of later drives accordingly. This can be shown rigorously. A rotation on the cavity state is given by  $U_\theta = \exp(i\theta a^\dagger a)$  and our drives are  $\text{Re}[\epsilon](a + a^\dagger)$  and  $\text{Im}[\epsilon]i(a - a^\dagger)$ . Proving this is essentially showing the equality

$$U_\theta \left( \text{Re}[\epsilon](a + a^\dagger) + \text{Im}[\epsilon]i(a - a^\dagger) \right) U_\theta^{-1} = \text{Re}[\epsilon e^{-i\theta}](a + a^\dagger) + \text{Im}[\epsilon e^{-i\theta}]i(a - a^\dagger), \quad (7.8)$$

which can be derived straightforwardly by using the Baker-Campbell-Hausdorff formula to obtain the identity  $\exp(-i\theta a^\dagger a) a \exp(i\theta a^\dagger a) = a \exp(i\theta)$ .

### 7.8.3 Decoding pulses

An alternate and potentially faster option for tomography of the cavity states is to combine the single-(logical)-qubit rotations and the CNOT operations. The motivation for this is that the individual operations are designed to be non-demolition, but for tomography we do not need that constraint. Indeed, we do not even require that the arbitrary back-action is the same for both logical states. For measurements of  $Z$  it seems unlikely that we will outperform the qubit-targeting CNOT operation in the Fock and 024 encoding, since those require durations of exactly  $\pi/\chi$  and  $\pi/2\chi$  precisely, which is relatively fast. For measurements of  $X$  and  $Y$ , however, we may get an enhancement. The operations which we need take

<sup>11</sup>Note that the CNOT operations must already have accounted for the phases due to Bell-state preparation, and the feed-forward pulses must have accounted for both of these unconditional phases

the form of, for example,

$$U_{\text{decodeX}} = U_{\text{arb,+}} |+\!X_L\rangle\langle+\!X_L| |1\rangle\langle 0|_T + U_{\text{arb,-}} |-\!X_L\rangle\langle-\!X_L| |0\rangle\langle 0|_T. \quad (7.9)$$

Designing the gauge freedoms (section 4.3.2) to realize this is somewhat subtle. Essentially they should be sufficient to span all Hermitian operators in the cavity space. Then we should duplicate those operators, one tensor producted with a projector for the transmon excited state and the other tensored with a projector for the transmon ground state. The basis we used is all operators  $|n\rangle\langle m|$  for  $n \leq m$  and  $i |n\rangle\langle m|$  for  $n < m$ . This allowed much faster operations, as seen in the bottom of table 7.1, though we have only done this so far for the 024 code. Whether or not the order of these non-commuting gauge freedoms can be optimized is an outstanding question.

My principal concern with this method is that it makes calibrating the measurement operators somewhat more opaque. Leakage and complexity mean that SPAM errors in the encoded qubit regime are already trickier than with the physical transmon qubits. In this method we would have three distinct measurement operators for each qubit, and we can no longer safely assume any structure for those operators. The simplest thing to do would be to not calibrate the measurement operators, but to report fidelities which include that cost. Tomography, while illustrative, will then look worse than our actual performance. The end of this line of reasoning is that to be maximally quantitative while avoiding underselling ourselves, we should report numbers from techniques like randomized benchmarking.

## 7.9 Full protocol

We have now assembled all of the requisite pieces for the experiment described in section 7.3. The full protocol can be seen in its entirety in figure 7.13. Starting with ground-state preparation, we rotate the logical qubits to a chosen initial state. We then generate our entangled bit between the transmons. Next are the two local CNOT gates between logical and communication bits. These are also generated by optimal control theory, but for the Fock and 024 encodings the CNOTs which target the qubit can be understood as parity and half-parity measurements. For the CNOT which targets the cavity, we must drive the cavity. Accordingly the frame of that drive must be shifted to accommodate the cavity phase

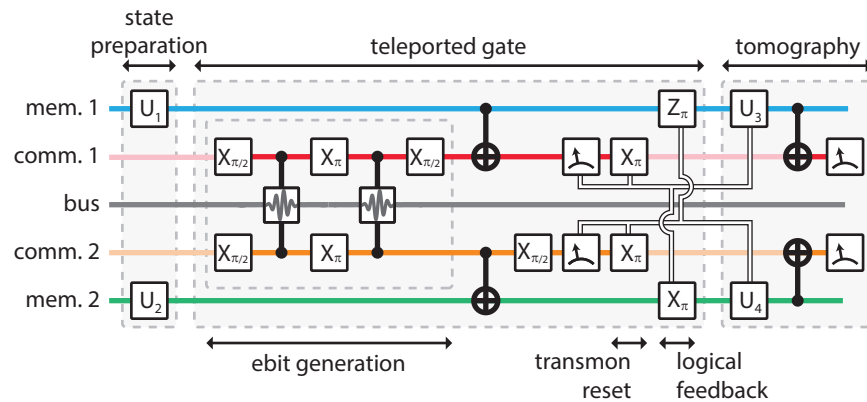


Figure 7.13: **Full experimental sequence for CNOT teleportation.** This sequence, which follows ground state preparation (section 7.5.2), enacts the teleportation of a CNOT operation between prepared states of two qubits, memory 1 and memory 2 (blue and green), then performs state tomography. The communication qubits (red and orange) are used as a sacrificial entangled pair, then reset and reused for tomography. The bus (gray) serves only to generate entanglement between the communication qubits using RIP gates.

gained during the Bell state generation. Next we simultaneously measure the two transmons in  $Z$  and  $X$  bases and use that information to reset them. It is not shown, but we then wait several hundred nanoseconds to allow the readout resonators to decay.

We then perform the logical feed-forward. Note that in the Fock encoding a logical  $Z$  rotation is only a shift of the physical cavity frame so we do not need an actual pulse, rather we incorporate that phase shift into the shift we apply to the tomographic pulses. In the 024 encoding a logical  $Z_\pi$  is given by a physical cavity frame shift of  $\pi/2$ . This does not generalize to other angles. The logical  $X_\pi$  operation must account for the unconditional phase gained during the Bell state generation and the CNOT operation, as well as the conditional phase gained during the measurement before qubit reset. This is followed by tomography, which must account for the same phases.

## 7.10 Results

We have successfully demonstrated a teleported CNOT and at this point have performed full quantum process tomography for the gate operating with the Fock encoding. We also present results both with and without feed-forward, illustrating the role that the feed-forward plays in teleportation.

### 7.10.1 Fock encoding, without feed-forward

The first dataset to demonstrate the action of the teleported gate is simply attempting to generate a Bell state in the memories. Memory qubit 1 is placed on the equator and the memory qubit 2 is left in the ground state. The action of the CNOT should be to bring it to the Bell state  $|\Phi^+\rangle$ . We first demonstrate the gate without applying the logical feed-forward, which means that three of the classical measurement outcomes should lead to states related to that Bell state by single-qubit rotations. These results are plotted in figure 7.14. The unconditioned outcome is plotted to the right. We see right away that none of these are actually  $|\Phi^+\rangle$ . Our phase correction in tomography is actually off by  $\pi/2$ , meaning that the superposition has a strictly imaginary phase. This is relatively unimportant for the Fock encoding, but must be corrected for the 024 encoding. The state which would not be modified by the feed-forward is the outcome 10.

We see that outcomes 00 and 01 leads to an opposite sign for  $ZZ$ —these are the states which will have a  $X$  operation applied to memory qubit 2. Outcomes 01 and 11 will similarly require a  $Z$  operation on memory qubit 1. For this particular state we see that when we include all outcomes unconditionally we get a completely mixed state. Since it does not make sense to provide fidelities for these results, we instead quantify the entanglement by providing the concurrence (Wootters, 1998). This number,  $C$ , quantifies how separable a state is, and has a (fairly opaque) explicit form,

$$\tilde{\rho} \equiv (\sigma_y \otimes \sigma_y) \rho^* (\sigma_y \otimes \sigma_y) \quad (7.10)$$

$$R \equiv \sqrt{\sqrt{\rho} \tilde{\rho} \sqrt{\rho}} \quad (7.11)$$

$$C(\rho) \equiv \max\{0, \lambda_1 - \lambda_2 - \lambda_3 - \lambda_4\}, \quad (7.12)$$

where  $\lambda_i$  are the eigenvalues of  $R$  in decreasing order.  $C$  is equal to one for a maximally entangled state, like a Bell pair, and zero for a strictly separable state. Any positive value implies some entanglement, though in practice noisy data can lead to a small positive result, so small values of concurrence should be handled delicately. This is essentially because the Hilbert space of entangled states is larger than the Hilbert space of separable states. We report concurrences between 0.65 and 0.57 for all four outcomes, distinctly indicating the presence of entanglement. The unconditioned result is almost completely mixed

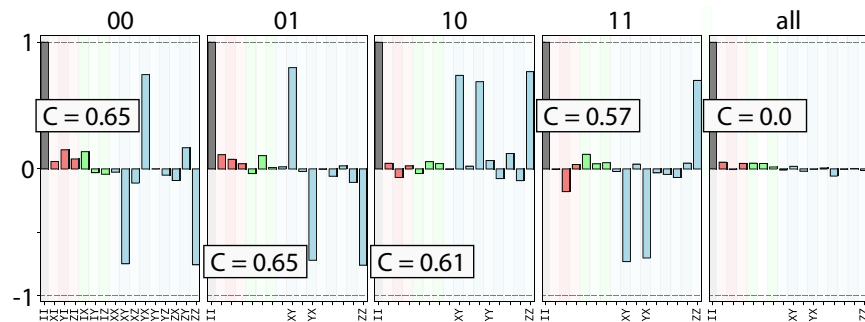


Figure 7.14: **Tomography of a Bell state in the Fock encoding, without feed-forward.** These plots are labeled by the outcome we obtained from measurement of the communication qubits. The far right plot is when all data are kept and that measurement outcome is ignored. For all four outcomes we see distinctly entangled states, as indicated by nonzero values of concurrence. All four of the conditioned outcomes are related by single qubit operations. The unconditioned data depict an almost completely mixed state, indicating the crucial role of the classical feed-forward.

and yields a concurrence of zero.

It is enlightening also to look at the conditioned processes before applying feed-forward. We have performed quantum process tomography on the operation as described in 6.7, though of course without the added complexities of a measurement process. While we represented the results of that experiment in the  $\chi$ -matrix formalism, we have chosen now to use the more compact  $\mathcal{R}$  matrix formalism (Chow *et al.*, 2012). The  $\mathcal{R}$  matrix is defined such that if you represent a quantum state in the Pauli basis as  $\vec{p}$ , as we often do,  $\vec{p}_{\text{out}} = \mathcal{R} \cdot \vec{p}_{\text{in}}$ . It illustrates how the expectation values of (generalized) Pauli operators change with the action of the process. Analysis of the elements of the  $\mathcal{R}$  matrix can also quickly confirm if a map is trace preserving or unital. We have plotted the results of this analysis in figure 7.15. We do not provide process fidelities because we have not applied feed-forward.

The data associated with the measurement outcome 10 look qualitatively like a CNOT operation should, albeit with some loss of contrast and added noise. The data associated with the other measurement outcomes have the same essential pattern but have some inverted values. These differences are associated with the feed-forward Pauli operations, which often invert the sign of the expectation value of other Pauli operations. The image on the right side of figure 7.15 represents the unconditioned outcome, which has some interesting properties. We would expect the identity to be invariant in any trace preserving and unital operation, but we see also that there are several other invariants:  $IX$ ,  $ZI$ ,

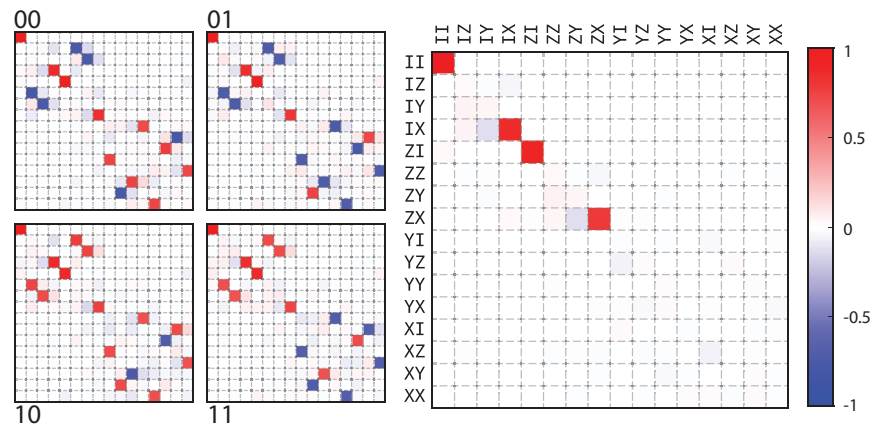


Figure 7.15: **Quantum process tomography of a teleported gate in the Fock encoding, without feed-forward.** Results are plotted in the  $\mathcal{R}$  matrix formalism. The left plots are labeled by the outcome we obtained from measurement of the communication qubits. The far right plot is when all data is kept and that measurement outcome is ignored. An ideal CNOT operation looks qualitatively like the result associated with measurement outcome 10.

and  $ZX$ . This is rather intuitive. The feed-forward operations which we are not applying are  $Z$  on one qubit, and  $X$  on the other. These expectation values will be invariant under that feed-forward, so we should not be surprised that they do not average away. We also see some slight non-zero elements that seem to remain through all of the data, which imply coherent errors not related to the feed-forward step.

### 7.10.2 Fock encoding, with feed-forward

To actually demonstrate a deterministic operation we must show functioning feed-forward. As with the previous data, it is illustrative to first look at the action on a single state and we first generated a Bell pair and have plotted that data in figure 7.16. It is immediately clear that these results are qualitatively very similar to each other. We have obtained fidelities for all of these conditioned outcomes, as well as the unconditioned outcome, of between 80% and 78%. We also see that the concurrences have dropped slightly from the results where we did not apply feed-forward, which makes sense as the added separable operations provide further opportunity for error and decoherence.

To truly demonstrate the gate in a convincing way we again turn to quantum process tomography. The results of that experiment are plotted in figure 7.17. We see process fidelities between 0.78 and






Figure 7.16: **Tomography of a Bell state in the Fock encoding, with feed-forward.** These plots are labeled by the outcome we obtained from measurement of the communication qubits. The far right plot is when all data is kept and that measurement outcome is ignored. For all outcomes we see the same highly entangled state, here with fidelities to the ideal state between 0.80 and 0.78. As explicitly shown in the far right plot, the gate now has almost no dependence on the measurement outcome. The concurrences are (unsurprisingly) slightly smaller than when we did not apply feed-forward (shown in figure 7.14).

0.74, generally in line with what we might expect from the Bell state fidelity. As the unconditioned fidelity is 0.76, we can safely say that we have demonstrated the first deterministic teleported two-qubit gate.

This fidelity is roughly consistent with our expectations. We can estimate our optimal control CNOT operations to be of similar fidelity as the single-qubit operations, we have a communication Bell pair fidelity of 95%, and our reset fidelity is 97% which essentially includes our readout fidelities. We also consider that the state preparation rotations are accounted for in our state tomography calibration. If we add these fidelities naively, we should expect fidelities of  $(0.98)^2 \times (0.97)^2 \times 0.95 \times 0.97 = 0.83$ .

This discrepancy between the predicted and observed is small but non-negligible. One potential culprit is the fidelity of preparing a Bell pair when there are photons in the cavity. We have done preliminary examinations of that effect and found it to be small, but it could easily be present on the scale of a few percent. Another possible culprit is the effect of the 19 kHz direct transmon-transmon coupling during the optimal control pulses. We have done simulations of the same optimal control pulses with the transmon detuned by -19 kHz, which roughly corresponds to the worst-case scenario, and found operations with roughly 1% lower fidelities. We also cannot claim that we have perfectly calibrated the unconditional and conditional phase corrections, whose remaining error may matter on that scale. It is also not exactly clear how going into tomography with an excited transmon, either due

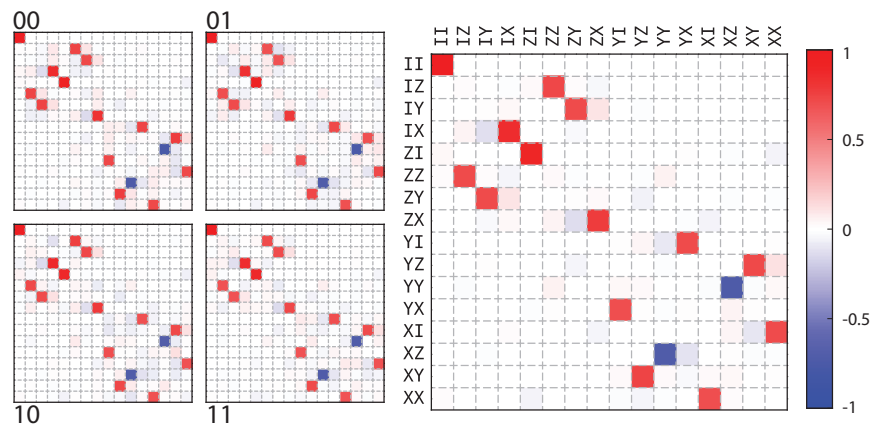


Figure 7.17: **Quantum process tomography of a teleported gate in the Fock encoding, with feed-forward.** Results are plotted in the  $\mathcal{R}$  matrix formalism. The left plots are labeled by the outcome we obtained from measurement of the communication qubits. The far right plot is when all data is kept and that measurement outcome is ignored. Starting with 00 and going clockwise, the conditioned process fidelities are 0.78, 0.75, 0.76, and 0.74, leading to an unconditioned process fidelity of 0.76.

to bad reset or due to the logical feed-forward, will effect the final outcome. Since the tomography CNOTs are essentially parity measurements, an excited qubit may have a worse-than-random effect. They may give an exactly wrong result rather than a random result.

### 7.10.3 024 encoding

A teleported gate on Fock-encoded cavities is interesting—it was the first teleported two-qubit gate—but what would be significantly more interesting is teleporting a gate between error-correctable logical qubits. While we have not yet done process tomography on such an operation, or demonstrated entanglement, we have taken the first steps towards that goal. We have shown “truth tables,” or the action of these operations on computational states, demonstrated not via logical state tomography but with Wigner tomography. As shown in figure 7.11, the logical  $|0_L\rangle$  state has a circular Wigner function, and the logical  $|1_L\rangle$  state exhibits a distinctive pattern with  $\pi/2$  rotational symmetry.

We have plotted the results of that operation in figure 7.18. We can see qualitatively that the CNOT operation is doing what it is supposed to: inverting the state of logical qubit 2 if and only if logical qubit 1 is in the  $|1_L\rangle$  state. These are single-qubit Wigner functions which cannot themselves demonstrate entanglement, but as we measure them simultaneously the same data acquisition method

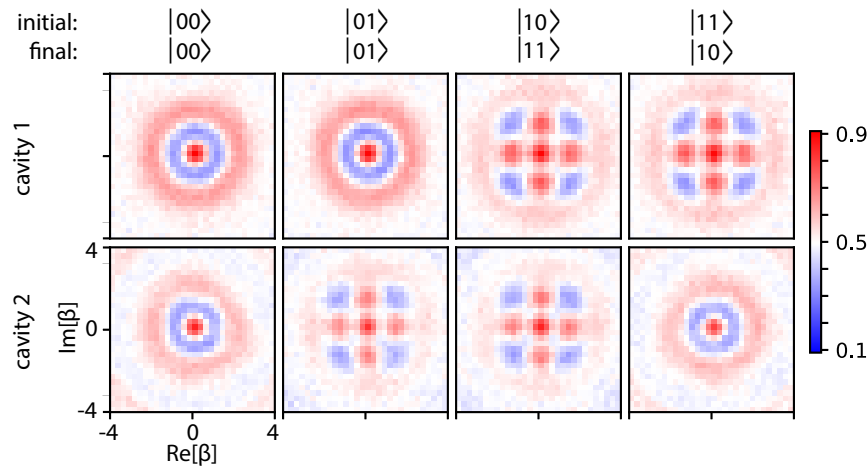


Figure 7.18: **Truth tables of the CNOT acting on 024-encoded qubits.** Plotted are raw single-qubit Wigner functions of the states resulting from the CNOT acting on the computational states. As we can see the initial states  $|0_L 0_L\rangle$  and  $|0_L 1_L\rangle$  are unaffected by the operation, but for states  $|1_L 0_L\rangle$  and  $|1_L 1_L\rangle$  we have inverted the state of cavity 2.

could be used to reconstruct four-dimensional joint Wigners. Taking a joint Wigner is an expensive process as it requires many measurements, but will provide an exhaustive analysis of the cavity state, including codespace errors. We will need to take such a dataset to rigorously analyze the leakage errors, but we will not want to use this technique to define our essential CNOT performance—it is simply too expensive. Indeed the data we have shown in figure 7.18 may not have a great enough extent to give a high-quality state reconstruction even of separable states. It is more reasonable to assess the quality of our operation via process tomography or two-qubit randomized benchmarking in the logical basis, then provide additional information to quantify leakage.

## 7.11 Conclusions

These results are preliminary, but we can safely claim to have deterministically teleported a two-qubit gate. In the coming months a great deal of data will flesh out these ideas and results, but even now we have done something quite interesting. Questions remain about how best to analyze these results, particularly regarding the correct way to do tomography and the correct way to analyze leakage errors. I suspect these topics might make excellent papers in their own right as the Yale group (and perhaps

others) improve the precision of manipulations on encoded qubits.

There are several bright paths forward for this experiment. First and foremost, I believe it should be relatively straightforward to obtain fidelities in the Fock encoding approaching 90% with minor improvements over what has already been implemented. These samples had abnormally low  $T_2$  and I estimate that getting a four-fold enhancement in that parameter is worth almost 10% points of fidelity. After that, I believe that better calibration of our drive chain is the key to getting well over two nines of fidelity on our optimal control pulses. It will be difficult to obtain significantly higher readout fidelities without sacrificing QND-ness, but we should have room to use readout cavities which are faster than the current ones without incurring Purcell loss. That may be worth a factor of two in both assignment fidelity and the transmon reset fidelity.

It will be very impressive to apply active quantum error correction to these qubits, and I suspect not actually that technically challenging given our current abilities. We may also use more complex codes that correct other errors. Lastly, if we smartly herald on leakage errors we may be able to see some exceptional (if non-deterministic) numbers. We have seen from simulation that a very large fraction of the errors that result from imperfections and decoherence unsurprisingly result in leakage errors. If we look for these, we may be able to achieve very high fidelities. Roughly half of such errors may be detected by simply looking to see if the transmon is excited at various points in the protocol. More may be detectable simply by selective  $\pi$  pulses on photon numbers above the codespace. This has been an exceedingly fun and stimulating project and it still has many good days ahead of it.

## Chapter 8

# Conclusion

Quantum mechanics had never been wrong. And now we know that it will not be wrong even in these very tricky conditions.

---

John Stewart Bell, 1986

In this chapter I conclude by looking forward from the present vantage point. An academic group, even one as large as we have at Yale, cannot keep pace with a determined commercial effort if they play at a game of brute force. In many ways the path forward from the “traditional” approach, as practiced at IBM and Google and so many others, is straightforward: increase transmon lifetimes, make the Hamiltonian more predictable, and make control and measurement cheaper, better, faster and closer to the device. Perhaps this will work. Manufacturing and controlling a million transmons seems quite intimidating to a person whose transmons all have names. That said, I can see a reasonable path forward to approach all of these concerns.

That is not precisely the case with the Yale, or modular-architecture, approach. I do not mean that derisively in any way. I have been wrong too many times to believe that a path is wrong because I do not see the answer to every question. I also know that there are always more questions than one thinks. It is obvious that to some extent we should try to find theory that fits our available Hamiltonians rather than only look for Hamiltonians to fit our theory. The result of (Ofek *et al.*, 2016) cannot be ignored—while others are talking about doing quantum error correction, the Yale group has done it.

I would like to conclude with what I see as the key challenges ahead, focusing on the problems of physics rather than more technical problems like the ones I mentioned previously.

**Choice of first-level encoding.** Thus far we have worked with first-level encodings designed to be resilient against photon-loss. Obviously single-photon loss is the biggest concern to an isolated excited cavity, but this is not our case. Because we must manipulate this system we will bring in many other sorts of errors. Our choice of encoding should be informed by this. All of the following challenges are directly intertwined with this decision.

**Operations which minimize leakage errors.** This is the biggest concern of the Yale approach. In exchange for measuring fewer syndromes we have a larger space of uncorrectable errors. I can imagine several solutions to this. Autonomous stabilization, as demonstrated in (Leghtas *et al.*, 2015) and (Touzard, in preparation), is certainly one of them. That seems to limit strictly the available encodings. The fact that we cannot yet perform a complete gate set without leaving the stabilized manifold is not important—intermittent use will allow the conversion of leakage errors to a more manageable type.

**Detecting codespace errors.** If are not going to work in a stabilized code space, then we are going to have leakage errors and we must be able to detect them and turn them into erasure errors. I discussed this in section 7.7.4, so I will not repeat it here.

**Resilience to the transmon.** Related to the previous problems, the imperfections of the transmon are likely to cause errors both within and outside of the logical Hilbert space. With regards to qubit manipulations, we may be able to design our optimal control pulses in a more resistant manner. Perhaps we can work with a pumping scheme which keeps the transmon entirely in the ground state or perhaps we can modulate the transmon coupling. While transmons are unlikely to disappear, I hope that in the future we can use circuit elements more tailored to our actual needs (Frattini *et al.*, 2017).

**High-quality inter-module entanglement.** To build an actual network we must be able to do high-fidelity gates between modules. These teleported gates will be limited by the quality of the resource

entanglement. Getting this fidelity above two nines will be a challenge. One approach is an error-correctable flying qubit (Pfaff, in preparation). A second is entanglement purification, though we must generate higher-quality local operations before we will see an improvement with this approach.

**Next-level encoding.** A final question is to what higher-level code we will apply. What errors will it be tailored to? One important thing to note is that in a modular approach we will not be bound to a nearest-neighbor connectivity graph, which be extremely useful. The code theorists have spent most (but not all) of their efforts on that more limited circumstance—there could be great discoveries to be made in the relaxed case.

One of the things I find interesting is that increased cavity lifetimes and faster operations are actually not terrible important. For a state with an average of 3 photons in a 3 millisecond cavity, a 1 microsecond operation on an error-correctable qubit can be in the range of three nines of fidelity—without applying error correction. At present the interesting problems have essentially nothing to do with the cavity lifetime, which is an excellent place to be.

# Bibliography

- D. Aharonov, and M. Ben-Or (1999), “Fault-Tolerant Quantum Computation With Constant Error Rate,” ArXiv quant-ph/9906129 .
- C. Axline, M. Reagor, R. Heeres, P. Reinhold, C. Wang, K. Shain, W. Pfaff, Y. Chu, L. Frunzio, and R. J. Schoelkopf (2016), “An architecture for integrating planar and 3D cQED devices,” Applied Physics Letters **109** (4), 042601.
- D. Bacon (2006), “Operator quantum error-correcting subsystems for self-correcting quantum memories,” Physical Review A **73** (1), 012340.
- C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas (2016), “High-fidelity quantum logic gates using trapped-ion hyperfine qubits,” Physical Review Letters **117** (6), 060504.
- R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O’Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis (2014), “Superconducting quantum circuits at the surface code threshold for fault tolerance,” Nature **508** (7497), 500.
- J. S. Bell (1966), “On the problem of hidden variables in quantum mechanics,” Rev. Mod. Phys. **38** (3), 447–452.
- C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters (1993), “Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels,” Physical Review Letters **70** (13), 1895.
- C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters (1996), “Purification of noisy entanglement and faithful teleportation via noisy channels,” Physical Review Letters **76** (5), 722.
- N. Bergeal, F. Schackert, M. Metcalfe, R. Vijay, V. E. Manucharyan, L. Frunzio, D. E. Prober, R. J. Schoelkopf, S. M. Girvin, and M. H. Devoret (2010), “Phase-preserving amplification near the quantum limit with a Josephson ring modulator,” Nature **465** (7294), 64.
- L. S. Bishop (2010), *Circuit quantum electrodynamics*, Ph.D. thesis.
- R. Blume-Kohout, J. K. Gamble, E. Nielsen, J. Mizrahi, J. D. Sterk, and P. Maunz (2013), “Robust, self-consistent, closed-form tomography of quantum logic gates on a trapped ion qubit,” arXiv 1310.4492 .



- J. Z. Blumoff, K. Chou, C. Shen, M. Reagor, C. Axline, R. T. Brierley, M. P. Silveri, C. Wang, B. Vlastakis, S. E. Nigg, L. Frunzio, M. H. Devoret, L. Jiang, S. M. Girvin, and R. J. Schoelkopf (2016), "Implementing and characterizing precise multiqubit measurements," *Physical Review X* **6** (3), 031041.
- M. Boissonneault, J. M. Gambetta, and A. Blais (2010), "Improved superconducting qubit readout by qubit-induced nonlinearities," *Physical Review Letters* **105** (10), 100504.
- H. Bombin (2015), "Gauge color codes: optimal transversal gates and gauge fixing in topological stabilizer codes," *New Journal of Physics* **17** (8), 083002.
- V. Bouchiat, D. Vion, P. Joyez, D. Esteve, and M. H. Devoret (1998), "Quantum coherence with a single Cooper pair," *Physica Scripta* **T76** (1), 165.
- D. Bouwmeester, J. W. Pan, K. Mattle, M. Eibl, and H. Weinfurter (1997), "Experimental quantum teleportation," *Nature* **390** (6660), 575.
- V. Braginsky, Y. Vorontsov, and K. Thorne (1980), "Quantum nondemolition measurements," *Science* **209**, 547.
- T. Brecht, W. Pfaff, C. Wang, Y. Chu, L. Frunzio, M. H. Devoret, and R. J. Schoelkopf (2016), "Multilayer microwave integrated quantum circuits for scalable quantum computing," *npj Quantum Information* **2** (1), 16002.
- K. R. Brown, A. C. Wilson, Y. Colombe, C. Ospelkaus, A. M. Meier, E. Knill, D. Leibfried, and D. J. Wineland (2011), "Single-qubit-gate error below  $10e-4$  in a trapped ion," *Physical Review A* **84** (3), 030303.
- O. Brune (1931), "Synthesis of a finite two-terminal network whose driving-point impedance is a prescribed function of frequency," *Journal of Mathematics and Physics* **10** (1-4), 191.
- M. D. de Burgh, N. K. Langford, A. C. Doherty, and A. Gilchrist (2008), "Choice of measurement sets in qubit tomography," *Physical Review A* **78** (5), 052122.
- Z. Chen, J. Kelly, C. Quintana, R. Barends, B. Campbell, Y. 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, and J. M. Martinis (2016), "Measuring and suppressing quantum state leakage in a superconducting qubit," *Physical Review Letters* **116** (2), 020501.
- M.-D. Choi (1975), "Completely positive linear maps on complex matrices," *Linear Algebra and its Applications* **10** (3), 285.
- J. M. Chow (2010), *Quantum information processing with superconducting qubits*, Ph.D. thesis.
- J. M. Chow, A. D. Córcoles, J. M. Gambetta, C. Rigetti, B. R. Johnson, J. A. Smolin, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen (2011), "Simple all-microwave entangling gate for fixed-frequency superconducting qubits," *Physical Review Letters* **107** (8), 080502.
- J. M. Chow, L. DiCarlo, J. M. Gambetta, A. Nunnenkamp, L. S. Bishop, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf (2010), "Detecting highly entangled states with a joint qubit readout," *Physical Review A* **81** (6), 062325.

- J. M. Chow, J. M. Gambetta, A. D. Córcoles, S. T. Merkel, J. A. Smolin, C. Rigetti, S. Poletto, G. A. Keefe, M. B. Rothwell, J. R. Rozen, M. B. Ketchen, and M. Steffen (2012), “Universal quantum gate set approaching fault-tolerant thresholds with superconducting qubits,” *Physical Review Letters* **109** (6), 060501.
- J. M. Chow, J. M. Gambetta, E. Magesan, D. W. Abraham, A. W. Cross, B. R. Johnson, N. A. Masluk, C. A. Ryan, J. A. Smolin, S. J. Srinivasan, and M. Steffen (2014), “Implementing a strand of a scalable fault-tolerant quantum computing fabric,” *Nature Communications* **5**, 4015.
- S. M. Clark, K.-M. C. Fu, T. D. Ladd, and Y. Yamamoto (2007), “Quantum computers based on electron spins controlled by ultrafast off-resonant single optical pulses,” *Physical Review Letters* **99** (4), 040501.
- J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt (1970), “Proposed experiment to test local hidden variable theories,” *Physical Review Letters* **24** (10), 549.
- A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf (2010), “Introduction to quantum noise, measurement, and amplification,” *Reviews of Modern Physics* **82** (2), 1155.
- R. E. Collin (2009), *Foundations for microwave engineering* (Wiley-IEEE Press).
- A. W. Cross, and J. M. Gambetta (2015), “Optimized pulse shapes for a resonator-induced phase gate,” *Physical Review A* **91** (3), 032325.
- M. Dalmonte, S. I. Mirzaei, P. R. Muppalla, D. Marcos, P. Zoller, and G. Kirchmair (2015), “Realizing dipolar spin models with arrays of superconducting qubits,” *Physical Review B* **92** (17), 174507.
- C. M. Dawson, and R. M. Solovay (2006), “The Solovay-Kitaev algorithm,” *Quantum Information & Computation* **6** (1), 81–95.
- M. H. Devoret (1995), “Quantum fluctuations in electrical circuits,” in *Les Houches Session LXIII*, edited by S. Reynaud, E. Giacobino, and J. Zinn-Justin, Chap. 10 (Elsevier Science) pp. 351–386.
- M. H. Devoret, and R. J. Schoelkopf (2013), “Superconducting circuits for quantum information: an outlook,” *Science* **339** (6124), 1169.
- O. Dial, D. T. McClure, S. Poletto, G. A. Keefe, M. B. Rothwell, J. M. Gambetta, D. W. Abraham, J. M. Chow, and M. Steffen (2016), “Bulk and surface loss in superconducting transmon qubits,” *Superconductor Science and Technology* **29** (4), 044001.
- D. P. DiVincenzo (2008), “The physical implementation of quantum computation,” arXiv quant-ph/0002077 .
- D. P. DiVincenzo, D. Bacon, J. Kempe, and G. Burkard (2000), “Universal quantum computation with the exchange interaction,” *Nature* **408**, 339.
- G. Dôme, M. Month, and M. Dienes (1987), “Rf systems: Waveguides and cavities,” in *AIP Conference Proceedings*, Vol. 153 (AIP) pp. 1296–1412.
- J. Dressel, T. A. Brun, and A. N. Korotkov (2014), “Implementing generalized measurements with superconducting qubits,” *Physical Review A* **90** (3), 032302.

- A. Einstein, B. Podolsky, and N. Rosen (1935), "Can quantum-mechanical description of physical reality be considered complete?" *Physical Review* **47** (10), 777.
- E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser (2000), "Quantum computation by adiabatic evolution," arXiv quant-ph/0001106v1 .
- R. M. Foster (1924), "A reactance theorem," *Bell System Technical Journal* **3** (2), 259.
- A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland (2012), "Surface codes: Towards practical large-scale quantum computation," *Physical Review A* **86** (3), 032324.
- N. E. Frattini, U. Vool, S. Shankar, A. Narla, K. M. Sliwa, and M. H. Devoret (2017), "3-Wave Mixing Josephson Dipole Element," ArXiv 1702.00869 .
- J. Gambetta, A. Blais, M. Boissonneault, A. A. Houck, D. I. Schuster, and S. M. Girvin (2008), "Quantum trajectory approach to circuit QED: Quantum jumps and the Zeno effect," *Physical Review A* **77** (1), 012112.
- J. M. Gambetta, A. D. Córcoles, S. T. Merkel, B. R. Johnson, J. A. Smolin, J. M. Chow, C. A. Ryan, C. Rigetti, S. Poletto, T. A. Ohki, M. B. Ketchen, and M. Steffen (2012), "Characterization of addressability by simultaneous randomized benchmarking," *Physical Review Letters* **109** (24), 240504.
- J. M. Gambetta, C. E. Murray, D. T. McClure, O. Dial, W. Shanks, J. W. Sleight, and M. Steffen (2017), "Investigating surface loss effects in superconducting transmon qubits," *IEEE Transactions on Applied Superconductivity* **27** (1), 1700205.
- C. W. Gardiner, and M. J. Collett (1985), "Input and output in damped quantum systems: Quantum stochastic differential equations and the master equation," *Physical Review A* **31** (6), 3761.
- K. Geerlings, Z. Leghtas, I. M. Pop, S. Shankar, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi, and M. H. Devoret (2013), "Demonstrating a Driven Reset Protocol for a Superconducting Qubit," *Physical Review Letters* **110** (12), 120501.
- K. Geerlings, S. Shankar, E. Edwards, L. Frunzio, R. J. Schoelkopf, and M. H. Devoret (2012), "Improving the quality factor of microwave compact resonators by optimizing their geometrical parameters," *Applied Physics Letters* **100** (19), 192601.
- A. Gilchrist, N. K. Langford, and M. A. Nielsen (2005), "Distance measures to compare real and ideal quantum processes," *Physical Review A* **71** (6), 062310.
- S. M. Girvin (2011), "Circuit qed: Superconducting qubits coupled to microwave photons," in *Les Houches Session XCVI* (Oxford University Press).
- R. J. Glauber (1963), "Coherent and incoherent states of the radiation field," *Physical Review* **131** (6), 2766.
- D. Gottesman (1997), "Stabilizer codes and quantum error correction," arXiv quant-ph/9705052 .
- D. Gottesman (2005), "Quantum error correction and fault-tolerance," arXiv quant-ph/0507174 .

- D. Gottesman, and I. L. Chuang (1999), "Quantum teleportation is a universal computational primitive," arXiv 10.1038/46503.
- D. Gottesman, A. Kitaev, and J. Preskill (2001), "Encoding a qubit in an oscillator," *Physical Review A* **64** (1), 012310.
- L. K. Grover (1996), "A fast quantum mechanical algorithm for database search," .
- B. Gustavsen, and A. Semlyen (1999), "Rational approximation of frequency domain responses by vector fitting," *IEEE Transactions on Power Delivery* **14** (3), 1052.
- G. Haack, F. Helmer, M. Mariani, F. Marquardt, and E. Solano (2010), "Resonant quantum gates in circuit quantum electrodynamics," *Physical Review B* **82** (2), 024514.
- S. Haroche (2013), "Nobel Lecture: Controlling photons in a box and exploring the quantum to classical boundary," *Reviews of Modern Physics* **85** (3), 1083.
- S. Haroche, and J.-M. Raimond (2012), *Exploring the quantum* (Oxford University Press).
- M. J. Hatridge, S. Shankar, M. Mirrahimi, F. Schakert, K. Geerlings, T. Brecht, K. M. Sliwa, b. abdo, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret (2012), "Quantum back-action of an individual variable-strength measurement," *Science* **339** (6116), 178.
- R. W. Heeres, P. Reinhold, N. Ofek, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf (2016), "Implementing a universal gate set on a logical qubit encoded in an oscillator," arXiv 1608.02430 .
- A. A. Houck, J. Koch, M. H. Devoret, and S. M. Girvin (2009), "Life after charge noise: recent results with transmon qubits," *Quantum Information Processing* **8** (2).
- A. A. Houck, J. A. Schreier, B. R. Johnson, J. M. Chow, J. Koch, J. M. Gambetta, D. I. Schuster, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf (2008), "Controlling the spontaneous emission of a superconducting transmon qubit," *Physical Review Letters* **101** (8), 080502.
- Y. F. Huang, X. F. Ren, Y. S. Zhang, L. M. Duan, and G. C. Guo (2004), "Experimental teleportation of a quantum controlled-NOT gate," *Physical Review Letters* **93** (24), 240501.
- K. Husimi (1940), "Some formal properties of the density matrix," *Proceedings of the Physico-Mathematical Society of Japan. 3rd Series* **22** (4), 264–314.
- C. Jackson, and S. J. van Enk (2015), "Detecting correlated errors in state-preparation-and-measurement tomography," *Physical Review A* **92** (4), 042312.
- L. Jiang, J. M. Taylor, A. S. Sørensen, and M. D. Lukin (2007), "Distributed quantum computation based on small quantum registers," *Physical Review A* **76** (6), 062323.
- R. Jozsa (1994), "Fidelity for mixed quantum states," *Journal of Modern Optics* **41** (12), 2315.
- A. Kamal, A. Marblestone, and M. Devoret (2009), "Signal-to-pump back action and self-oscillation in double-pump Josephson parametric amplifier," *Physical Review B* **79** (18), 184301.

- N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser (2005), "Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms," *Journal of Magnetic Resonance* **172** (2), 296.
- G. Kirchmair, B. Vlastakis, Z. Leghtas, S. E. Nigg, H. Paik, E. Ginossar, M. Mirrahimi, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf (2013), "Observation of quantum state collapse and revival due to the single-photon Kerr effect," *Nature* **495** (7440), 205.
- A. Kitaev (2006), "Protected qubit based on a superconducting current mirror," ArXiv cond-mat/0609441 .
- A. Y. Kitaev (2003), "Fault-tolerant quantum computation by anyons," *Annals of Physics* **303** (1), 2.
- J. Koch, M. Y. Terri, J. Gambetta, A. A. Houck, and D. I. Schuster (2007), "Charge-insensitive qubit design derived from the Cooper pair box," *Physical Review A* **76** (4), 042319.
- P. Kok, W. J. Munro, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn (2007), "Linear optical quantum computing with photonic qubits," *Reviews of Modern Physics* **79** (1), 135.
- C. Langer, R. Ozeri, J. Jost, J. Chiaverini, B. DeMarco, A. Ben-Kish, R. Blakestad, J. Britton, D. Hume, W. Itano, D. Leibfried, R. Reichle, T. Rosenband, T. Schaetz, P. Schmidt, and D. Wineland (2005), "Long-lived qubit memory using atomic ions," *Physical Review Letters* **95** (6), 060502.
- Z. Leghtas, G. Kirchmair, B. Vlastakis, M. H. Devoret, R. J. Schoelkopf, and M. Mirrahimi (2013a), "Deterministic protocol for mapping a qubit to coherent state superpositions in a cavity," *Physical Review A* **87** (4), 042315.
- Z. Leghtas, G. Kirchmair, B. Vlastakis, R. J. Schoelkopf, M. H. Devoret, and M. Mirrahimi (2013b), "Hardware-efficient autonomous quantum memory protection," *Physical Review Letters* **111** (12), 120501.
- Z. Leghtas, S. Touzard, I. M. Pop, A. Kou, B. Vlastakis, A. Petrenko, K. M. Sliwa, A. Narla, S. Shankar, M. J. Hatridge, M. Reagor, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi, and M. H. Devoret (2015), "Confining the state of light to a quantum manifold by engineered two-photon loss," *Science* **347** (6224), 853.
- D. Loss, and D. P. DiVincenzo (1998), "Quantum computation with quantum dots," *Physical Review A* **57** (1), 120.
- J. S. Lundeen, A. Feito, H. Coldenstrodt-Ronge, K. L. Pagnell, C. Silberhorn, T. C. Ralph, J. Eisert, M. B. Plenio, and I. A. Walmsley (2008), "Tomography of quantum detectors," *Nature Physics* **5** (1), 27.
- D. J. MacKay, G. Mitchison, and P. L. McFadden (2004), "Sparse-graph codes for quantum error correction," *IEEE Transactions on Information Theory* **50** (10), 2315.
- E. Magesan, J. M. Gambetta, A. D. Córcoles, and J. M. Chow (2015), "Machine learning for discriminating quantum measurement trajectories and improving readout," *Physical Review Letters* **114** (20), 200501.

- E. Magesan, J. M. Gambetta, and J. Emerson (2011), "Scalable and robust randomized benchmarking of quantum processes," *Physical Review Letters* **106** (18), 180504.
- E. Magesan, J. M. Gambetta, B. R. Johnson, C. A. Ryan, J. M. Chow, S. T. Merkel, M. P. da Silva, G. A. Keefe, M. B. Rothwell, T. A. Ohki, M. B. Ketchen, and M. Steffen (2012), "Efficient measurement of quantum gate error by interleaved randomized benchmarking," *Physical Review Letters* **109** (8), 080505.
- V. E. Manucharyan, E. Boaknin, M. Metcalfe, R. Vijay, I. Siddiqi, and M. Devoret (2007), "Microwave bifurcation of a Josephson junction: Embedding-circuit requirements," *Physical Review B* **76** (1), 014524.
- V. E. Manucharyan, J. Koch, L. I. Glazman, and M. H. Devoret (2009), "Fluxonium: single Cooper-pair circuit free of charge offsets," *Science* **326** (5949), 113.
- J. M. Martinis, S. Nam, J. Aumentado, and C. Urbina (2002), "Rabi oscillations in a large Josephson-junction qubit," *Physical Review Letters* **89** (11), 117901.
- D. T. McClure, H. Paik, L. S. Bishop, M. Steffen, J. M. Chow, and J. M. Gambetta (2016), "Rapid driven reset of a qubit readout resonator," *Physical Review Applied* **5** (1), 011001.
- S. T. Merkel, J. M. Gambetta, J. A. Smolin, S. Poletto, A. D. Córcoles, B. R. Johnson, C. A. Ryan, and M. Steffen (2013), "Self-consistent quantum process tomography," *Physical Review A* **87** (6), 062119.
- M. H. Michael, M. Silveri, R. T. Brierley, V. V. Albert, J. Salmilehto, L. Jiang, and S. M. Girvin (2016), "New class of quantum error-correcting codes for a bosonic mode," *Physical Review X* **6** (3), 031006.
- C. Monroe, and J. Kim (2013), "Scaling the ion trap quantum processor," *Science* **339** (6124), 1164.
- J. E. Mooij, T. P. Orlando, L. Levitov, L. Tian, C. H. van der Wal, and S. Lloyd (1999), "Josephson persistent-current qubit," *Science* **285** (5430), 1036.
- F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm (2009), "Simple pulses for elimination of leakage in weakly nonlinear qubits," *Physical Review Letters* **103** (11), 110501.
- C. Muller, A. Shnirman, S. Poletto, and J. Lisenfeld (2015), "Interacting two-level defects as sources of fluctuating high-frequency noise in superconducting circuits," *Physical Review B* **92** (3), 035442.
- S. O. Mundhada, S. Shankar, Y. Liu, M. J. Hatridge, A. Narla, K. M. Sliwa, S. M. Girvin, and M. H. Devoret (2015), "Dependence of transmon qubit relaxation rate on cavity photon population," APS March Meeting 2015 .
- A. Narla, S. Shankar, M. Hatridge, Z. Leghtas, K. M. Sliwa, E. Zolys-Geller, S. O. Mundhada, W. Pfaff, L. Frunzio, R. J. Schoelkopf, and M. H. Devoret (2016), "Robust Concurrent Remote Entanglement Between Two Superconducting Qubits," *Physical Review X* **6** (3), 031036.
- M. A. Nielsen, and I. L. Chuang (2000), *Quantum computation and quantum information* (Cambridge University Press).

- S. E. Nigg, and S. M. Girvin (2013), "Stabilizer quantum error correction toolbox for superconducting qubits," *Physical Review Letters* **110** (24), 243604.
- S. E. Nigg, H. Paik, B. Vlastakis, G. Kirchmair, and S. Shankar (2012), "Black-box superconducting circuit quantization," *Physical Review Letters* **108** (24), 240502.
- N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf (2016), "Extending the lifetime of a quantum bit with error correction in superconducting circuits," *Nature* **536** (7617), 441.
- H. Paik, A. Mezzacapo, M. Sandberg, and D. T. McClure (2016), "Experimental demonstration of a resonator-induced phase gate in a multiqubit circuit-QED system," *Physical Review Letters* **117** (25), 250502.
- H. Paik, D. I. Schuster, L. S. Bishop, G. Kirchmair, G. Catelani, A. P. Sears, B. R. Johnson, M. J. Reagor, L. Frunzio, L. I. Glazman, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf (2011), "Observation of high coherence in Josephson junction qubits measured in a three-dimensional circuit QED architecture," *Physical Review Letters* **107** (24), 240501.
- A. Petrenko, A. Sears, G. Kirchmair, H. Paik, L. Sun, G. Catelani, L. Glazman, and R. Schoelkopf (2012), "Varying cavity quality factor in situ for a transmon in circuit qed," in *APS Meeting Abstracts*, Vol. 1, p. 39004.
- E. M. Purcell (1945), "Spontaneous emission probabilities at radio frequencies," *NATO ASI Series*, 839.
- I. I. Rabi (1937), "Space quantization in a gyrating magnetic field," *Physical Review* **51** (8), 652.
- I. I. Rabi, J. R. Zacharias, S. Millman, and P. Kusch (1938), "A new method of measuring nuclear magnetic moment," *Physical Review* **53** (4), 318.
- S. Ramo, J. R. Whinnery, and T. Van Duzer (1965), *Fields and waves in communications electronics* (Wiley).
- R. Raussendorf, and H. J. Briegel (2001), "A one-way quantum computer," *Physical Review Letters* **86** (22), 5188.
- M. Reagor, H. Paik, G. Catelani, L. Sun, C. Axline, E. Holland, I. M. Pop, N. A. Masluk, T. Brecht, L. Frunzio, M. H. Devoret, L. Glazman, and R. J. Schoelkopf (2013), "Reaching 10 ms single photon lifetimes for superconducting aluminum cavities," *Applied Physics Letters* **102** (19), 192604.
- M. D. Reed (2013), *Entanglement and quantum error correction with superconducting circuits*, Ph.D. thesis.
- M. D. Reed, L. DiCarlo, B. R. Johnson, L. Sun, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf (2010a), "High-fidelity readout in circuit quantum electrodynamics using the Jaynes-Cummings nonlinearity," *Physical Review Letters* **105** (17), 173601.
- M. D. Reed, B. R. Johnson, A. A. Houck, L. DiCarlo, J. M. Chow, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf (2010b), "Fast reset and suppressing spontaneous emission of a superconducting qubit," *Applied Physics Letters* **96** (20), 203110.

- M. Riebe, H. Häffner, C. F. Roos, W. Hänsel, J. Benhelm, G. P. T. Lancaster, T. W. Körber, C. Becher, F. Schmidt-Kaler, D. F. V. James, and R. Blatt (2004), "Deterministic quantum teleportation with atoms," *Nature* **429** (6993), 734.
- D. Riste, L. DiCarlo, S. Poletto, O. P. Saira, V. Vesterinen, A. Bruno, and M. Z. Huang (2015), "Detecting bit-flip errors in a logical qubit using stabilizer measurements," *Nature Communications* **6**, 6983.
- S. A. Schelkunoff, and H. T. Fris (1953), *Antenna Theory and Practice* (Wiley).
- D. I. Schuster (2007), *Circuit quantum electrodynamics*, Ph.D. thesis.
- A. P. Sears (2013), *Extending Coherence in Superconducting Qubits: From microseconds to milliseconds*, Ph.D. thesis.
- A. P. Sears, A. Petrenko, G. Catelani, L. Sun, H. Paik, G. Kirchmair, L. Frunzio, L. I. Glazman, S. M. Girvin, and R. J. Schoelkopf (2012), "Photon shot noise dephasing in the strong-dispersive limit of circuit QED," *Physical Review B* **86** (18), 180504.
- S. Seidelin, J. Chiaverini, R. Reichle, J. Bollinger, D. Leibfried, J. Britton, J. Wesenberg, R. Blakestad, R. Epstein, D. Hume, W. Itano, J. Jost, C. Langer, R. Ozeri, N. Shiga, and D. Wineland (2006), "Microfabricated surface-electrode ion trap for scalable quantum information processing," *Physical Review Letters* **96** (25), 253003.
- J. H. Shirley (1965), "Solution of the Schrödinger equation with a Hamiltonian periodic in time," *Physical Review* **138** (4B), B979.
- P. W. Shor (1994), "Algorithms for quantum computation: discrete logarithms and factoring," .
- P. W. Shor (1995), "Scheme for reducing decoherence in quantum computer memory," *Physical Review A* **52** (4), R2493.
- F. Solgun, D. W. Abraham, and D. P. DiVincenzo (2014), "Blackbox quantization of superconducting circuits using exact impedance synthesis," *Physical Review B* **90** (13), 134504.
- D. A. Steck (2007), "Quantum and atom optics," .
- M. Stoutimore, J. Medford, Q. Herr, O. Naaman, H. Hearne, J. Strand, A. Przybysz, A. Pesetski, and J. Przybysz (2015), "Demonstrated control of a Transmon using a reciprocal quantum logic digital circuit - Part 1," APS March Meeting 2015, abstract #J39.012 .
- B. M. Terhal (2015), "Quantum error correction for quantum memories," *Reviews of Modern Physics* **87** (2), 307.
- M. Tinkham (2004), *Introduction to superconductivity* (Dover Books).
- T. Van Duzer, and C. W. Turner (1999), *Principles of Superconductive Devices and Circuits*, Vol. 16 (Prentice Hall).
- L. M. K. Vandersypen, and I. L. Chuang (2005), "NMR techniques for quantum control and computation," *Reviews of Modern Physics* **76** (4), 1037.



- B. Vlastakis, G. Kirchmair, Z. Leghtas, S. E. Nigg, L. Frunzio, S. M. Girvin, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf (2013), "Deterministically encoding quantum information using 100-photon Schrödinger cat states," *Science* **342** (6158), 607.
- A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R. S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf (2004), "Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics," *Nature* **431** (7005), 162.
- C. Wang, C. Axline, Y. Y. Gao, T. Brecht, Y. Chu, L. Frunzio, M. H. Devoret, and R. J. Schoelkopf (2015), "Surface participation and dielectric loss in superconducting qubits," *Applied Physics Letters* **107** (16), 162601.
- C. Wang, Y. Y. Gao, P. Reinhold, R. W. Heeres, N. Ofek, K. Chou, C. Axline, M. Reagor, J. Blumoff, K. M. Sliwa, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf (2016), "A Schrödinger cat living in two boxes," *Science* **352** (6289), 1087.
- E. Wigner (1932), "On the quantum correction for thermodynamic equilibrium," *Physical Review* **40** (5), 749.
- M. M. Wilde (2009), "Quantum information theory," arXiv 1106.1445v6 .
- H. M. Wiseman, and G. J. Milburn (2009), *Quantum measurement and control* (Cambridge University Press).
- M. Wolf (2012), *Quantum channels & operations guided tour*.
- W. K. Wootters (1998), "Entanglement of formation of an arbitrary state of two qubits," *Physical Review Letters* **80** (10), 2245.

## **Appendix A**

# **Stabilizer experiment extended results**

**A.1 Detector tomography extended results**

**A.2 Measurement process extended results**




Figure A.1: Tomography of all POVMs, without post-selection on the success-herald

Table A.1: POVM results.

operator	unheralded results				heralded results			
	$F_J$	$D_J$	$F_S$	$D_S$	$F_J$	$D_J$	$F_S$	$D_S$
ZII	0.948	0.055	0.936	0.068	0.976	0.025	0.967	0.032
IZI	0.945	0.053	0.931	0.071	0.975	0.026	0.963	0.038
IIZ	0.958	0.044	0.948	0.054	0.976	0.025	0.967	0.035
ZZI	0.938	0.065	0.906	0.096	0.967	0.036	0.938	0.065
ZIZ	0.943	0.068	0.914	0.094	0.967	0.043	0.938	0.068
IZZ	0.942	0.066	0.918	0.087	0.966	0.043	0.934	0.073
ZZZ	0.910	0.093	0.875	0.130	0.954	0.050	0.926	0.078




Figure A.2: Tomography of all POVMs, with post-selection on the success-herald

Table A.2: Specificity results.

operator	unheralded results	heralded results
ZII	2.3°	1.1°
IZI	1.8°	1.0°
IIZ	1.8°	0.9°
ZZI	3.6°	2.3°
ZIZ	5.2°	3.4°
IZZ	4.5°	3.4°
ZZZ	4.4°	3.0°

Table A.3: Quantum instrument metrics.

operator	unheralded results				heralded results			
	$F_J$	$D_J$	$F_S$	$D_S$	$F_J$	$D_J$	$F_S$	$D_S$
ZII	0.716	0.373	0.566	0.501	0.789	0.312	0.690	0.400
IZI	0.757	0.295	0.615	0.428	0.832	0.246	0.749	0.345
IIZ	0.757	0.223	0.607	0.278	0.827	0.303	0.756	0.417
ZZI	0.721	0.264	0.629	0.315	0.818	0.333	0.754	0.411
ZIZ	0.735	0.297	0.638	0.388	0.823	0.220	0.757	0.272
IZZ	0.730	0.306	0.595	0.437	0.813	0.238	0.759	0.272
ZZZ	0.674	0.353	0.578	0.450	0.796	0.221	0.741	0.272

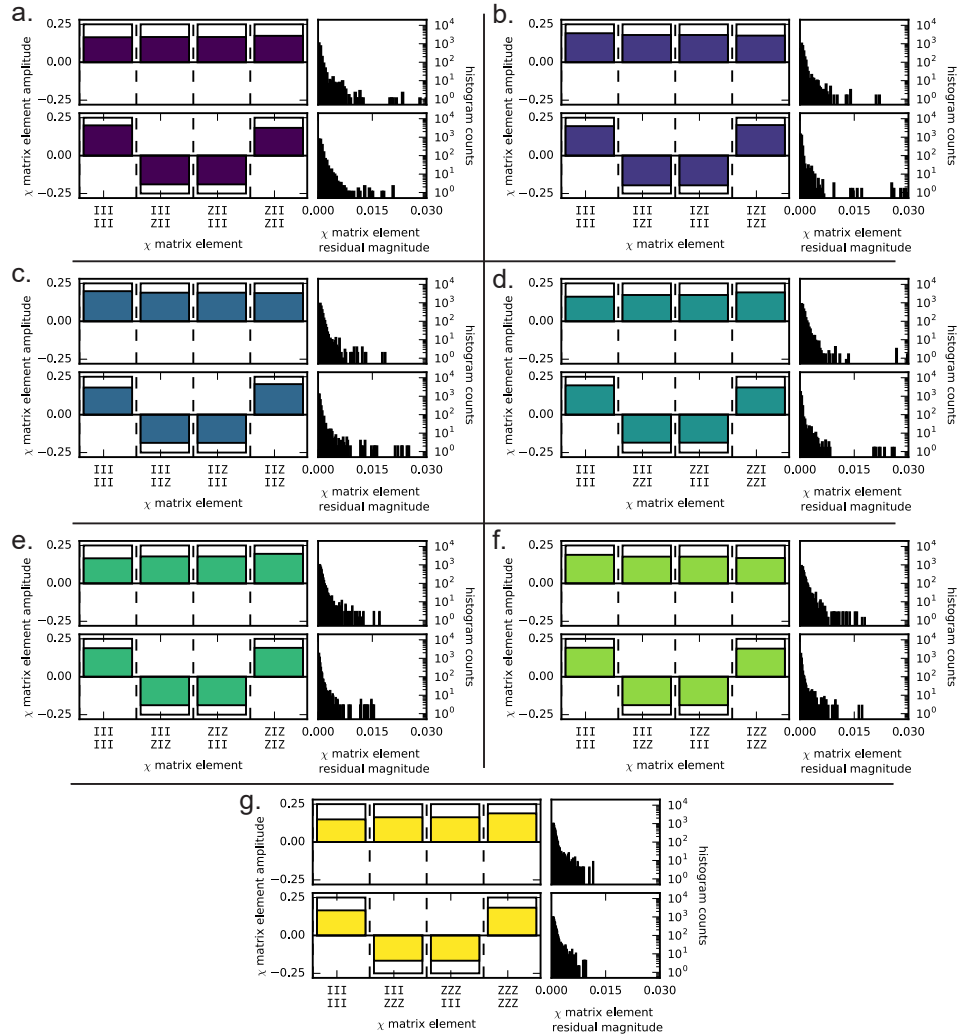


Figure A.3: **Full conditioned QPT results.** (a)-(g) show conditioned QPT results for all seven subset-parity operators:  $ZII$ ,  $IZI$ ,  $IIZ$ ,  $ZZI$ ,  $ZIZ$ ,  $IZZ$ , and  $ZZZ$ , respectively. For each panel, we represent the results in the  $\chi$  matrix representation for both even (top row) and odd (bottom row) outcomes. For each outcome, the left bar plot directly shows the four non-zero components of the  $\chi$  matrix with the ideal amplitudes outlined with amplitude  $\pm 1/4$ . Note the sign-change for the off-diagonal components between even and odd outcomes for all operators. The right plot is a histogram that illustrates the magnitude for the remaining (ideally zero) components of the  $\chi$  matrix. We note that they are all small and with few outliers.

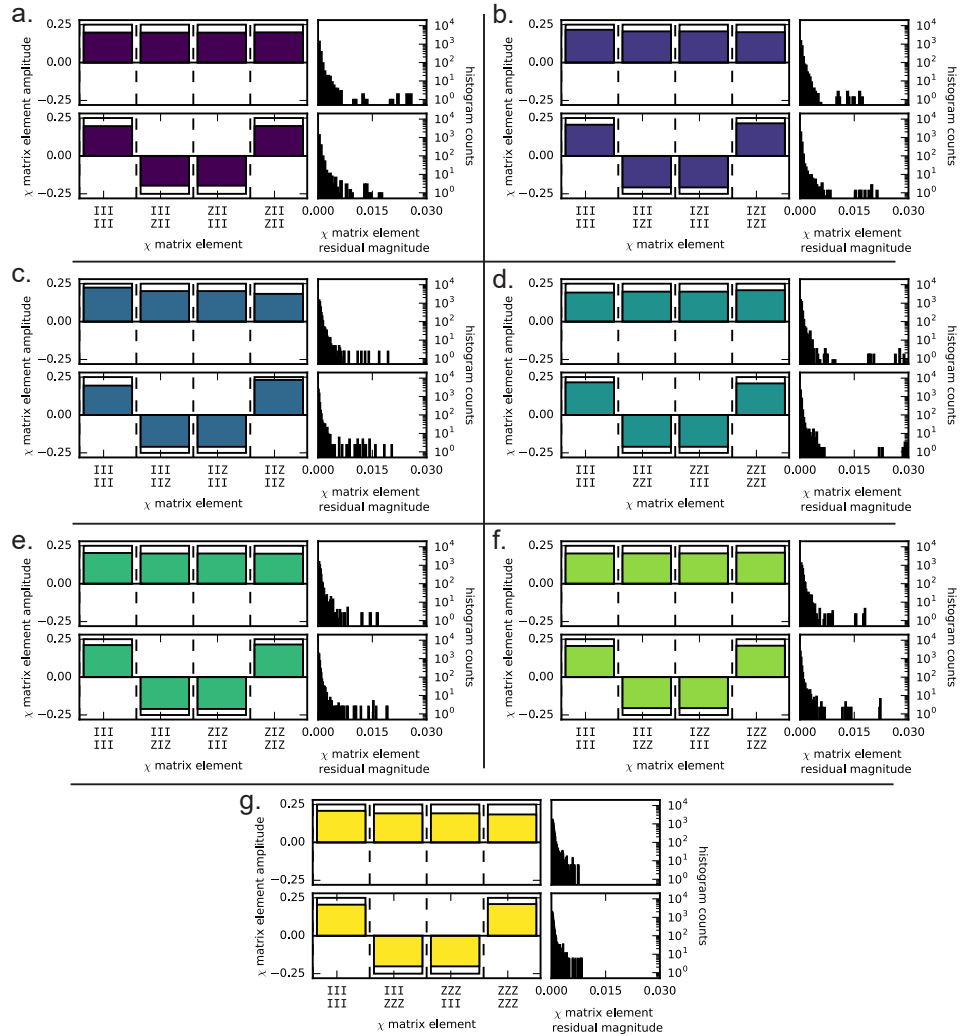


Figure A.4: **Conditioned QPT results with success herald.** (a)-(g) show conditioned QPT results for all seven subset-parity operators:  $ZII$ ,  $IZI$ ,  $IIZ$ ,  $ZZI$ ,  $ZIZ$ ,  $IZZ$ , and  $ZZZ$ , respectively. The representation of this data is the same as in Fig. A.3.

## Appendix B

# Tomography code

---

```
import numpy as np
from numpy import pi
import numpy.linalg
import scipy.linalg
import itertools
import cvxpy as cvx

def tensor(*args):
    out = args[0]
    for arg in args[1:]:
        out = np.kron(out, arg)
    return np.matrix(out)

rI = np.matrix([[1,0],[0,1]], dtype=np.complex)
sigmax = np.matrix([[0,1],[1,0]], dtype=np.complex)
sigmay = np.matrix([[0,-1j],[1j,0]], dtype=np.complex)
sigmaz = np.matrix([[1,0],[0,-1]], dtype=np.complex)
s = {'I': rI, 'x':sigmax, 'y':sigmay, 'z':sigmaz} # Pauli observables

rX = np.matrix(scipy.linalg.expm(sigmax * 1j*pi/2))
rx = np.matrix(scipy.linalg.expm(sigmax * 1j*pi/4))
rm = np.matrix(scipy.linalg.expm(-sigmax * 1j*pi/4))
rY = np.matrix(scipy.linalg.expm(sigmay * 1j*pi/2))
ry = np.matrix(scipy.linalg.expm(sigmay * 1j*pi/4))
rn = np.matrix(scipy.linalg.expm(-sigmay * 1j*pi/4))
r_max = {'I': rI, 'x':rx, 'y':ry, 'X':rX, 'Y':rY, 'm':rm, 'n':rn} # Rotations

#=====
# Core subroutines
#=====
```



---

```

def create_pauli_observables(nq=2, str_only=False):
    """
    This assembles are ordered list of observable tensored out of (I, X, Y, Z).
    """
    assert nq < 5
    strings = list(itertools.product(['I','z','y','x'], repeat=nq))
    if str_only:
        return strings

    mats = []
    for opset in strings:
        ops = [s[el] for el in opset]
        mats.append(tensor(*ops))

    return strings, mats

def create_prerotations(nq=2, r_sel='IxyX', str_only=False):
    """
    This assembles are ordered list of prerotations tensored out of
    operations in r_max
    """
    r = dict((k, r_max[k]) for k in r_sel)
    strings = list(itertools.product(r.keys(), repeat=nq))
    if str_only:
        return strings

    mats = []
    for opset in strings:
        ops = [r[el] for el in opset]
        mats.append(tensor(*ops))
    return strings, mats

def create_V_prerotations(nq=2):
    """
    This assembles are ordered list of prerotations tensored out of
    (I, X). It returns both strings and matrices.
    """
    strings = list(itertools.product('IX', repeat=nq))

    mats = []
    for opset in strings:
        ops = [r_max[el] for el in opset]
        mats.append(tensor(*ops))

    return strings, mats

```

```

=====
# State tomography routines
=====

def build_A(V, prerot_mats, pauli_mats=None, A_only=False, nq=2):
    '''
    This builds the matrix that can be used to relate measurements taken
    on the prerotations in the presence of a POVM matrix V and the
    pauli expectations of the measured system. V is generated by doing
    measurements on the states generated (in order) by create_V_prerotations().

    A . (pauli expt. vector) = (measurement vector)
    Each element A_ij is given by Tr[sigma_j * R_i * V * R_i^]

    Linear inversion is done by acting B on the measurements generated by
    measuring (in order) with the rotations generated by create_prerotations().

    The state vector generated by B on the measurement outcomes is given in the
    Pauli basis in the order generated by create_pauli_observables().
    '''
    if pauli_mats is None:
        _, pauli_mats = create_pauli_observables(nq)

    A = np.zeros((len(prerot_mats), len(pauli_mats)), dtype=np.complex)
    for i, rot in enumerate(prerot_mats):
        for j, pauli in enumerate(pauli_mats):
            A[i,j] = np.trace(rot.H*V*rot*pauli)

    A = np.matrix(A)/2**nq
    if A_only:
        return A

    B = numpy.linalg.pinv(A)
    # print 'Condition # of A: %s' % numpy.linalg.cond(A)
    return A, B

def build_multi_A(Vs, prerot_mats, pauli_mats=None):
    '''
    For multiple-bit tomography. Vs should be generated for each outcome
    (think P00, P01, P10, whereas before we used only P0). The matrix B
    generates the state acting on the measurement outcomes similarly stacked.
    '''
    As = [build_A(V, prerot_mats, pauli_mats=pauli_mats, A_only=True) for V in Vs]
    A = np.vstack(As)

```

---

```

B = numpy.linalg.pinv(A)
# print 'Condition # of A: %s' % numpy.linalg.cond(A)
return A, B

def generate_rotated_Vs(V, nq, prerot_mats):
    '''
    Tr[R_i rho_fit R_i* V] = Tr[rho_fit R_i* V * R_i ]
    So we can rotate V in advance and only once.
    '''
    rot_vs = [None,]*len(prerot_mats)
    for idx, prerot in enumerate(prerot_mats):
        rot_vs[idx] = (prerot.H * V * prerot)
    return rot_vs

def cvxMLE(measurements, rot_vs_H, pauli_mats, nq=2):
    '''
    Minimize using CVXPY. Since tomographic reconstruction is a convex
    problem, this is faster than using nelder mead. We actually do a
    constrained least-squares minimization.

    This relies a couple of mathematical facts.
    1) Tr[rho*Vi] = np.dot( rho.flat , Vi.flat) (cvx.vec() flattens)

    2) CVXPY can't natively handle imaginary numbers, so we embed rho
    into a larger block matrix V = [[ re[rho], im[rho]]
                                     [-im[rho], re[rho]]],
    which also guarantees positive definiteness.
    '''
    dim = 2**nq

    z = cvx.Semidef(2*dim)
    rho_r = z[:dim,:dim]
    rho_i = z[dim:,:dim]
    rho_r2 = z[dim:,dim:]
    rho_i2 = z[:dim,dim:]

    rot_vs_H = [np.array(rot_v_H).flatten() for rot_v_H in rot_vs_H]
    rot_vs_H = np.array(rot_vs_H)

    real_cost = cvx.sum_squares(measurements - \
                                (rot_vs_H.real * cvx.vec(rho_r) - \
                                 rot_vs_H.imag * cvx.vec(rho_i)))

    obj = cvx.Minimize(real_cost)

```

```

constraints = [
    rho_r == rho_r2,
    rho_i == -rho_i2,
    cvx.trace(rho_r) == 1.0,
]

prob = cvx.Problem(obj, constraints)
prob.solve(verbose=False)

if prob.status != cvx.OPTIMAL:
    raise Exception('CVX optimization failed')

return np.matrix(rho_r.value + 1j*rho_i.value)

#=====
# Process tomography routines
#=====

def vectorize_into_columns(mats_in):
    '''
    Takes a list of input matrices and creates a new matrix whose columns
    are the input matrices vectorized
    '''
    m = len(mats_in)
    d = mats_in[0].shape[0]
    out_mat = np.zeros((d**2, m), dtype=np.complex)
    for i, mat in enumerate(mats_in):
        out_mat[:,i] = mat.flatten()

    return np.matrix(out_mat)

def construct_L(rhos_in, rhos_out):
    '''
    Can be faster.
    '''
    R0 = vectorize_into_columns(rhos_in)
    R = vectorize_into_columns(rhos_out)
    return R * np.linalg.pinv(R0)

def construct_J(L):
    '''
    This performs the re-ordering from L to J. Can be faster.
    '''
    dim = int(round(np.sqrt(L.shape[0])))
    J = np.zeros((dim**2, dim**2), dtype=np.complex)

```

---

```

k = 0
for i in range(dim):
    for j in range(dim):
        J[i*dim:(i+1)*dim, j*dim:(j+1)*dim] = np.reshape(L[k,:], (dim,dim))
        k += 1
return np.matrix(J)

def construct_E(nq=2):
    _, proc_basis_mats = create_pauli_observables(nq)
    return vectorize_into_columns(proc_basis_mats)

def construct_X_from_J(J):
    d = int(np.sqrt(J.shape[0]))
    nq = int(round(np.log(d)/np.log(2.0)))
    E = construct_E(nq=nq)
    X = d**-2 * E.H * J * E
    return X

def construct_chi(rhos_in, rhos_out):
    '''
    This performs unconstrained least-squares to find a process matrix.

    Rhos_in are generated (in order) from the rotations generated by
    create_prerotations(). Rhos_out are state tomography resulting from those
    initial states.
    '''
    L = construct_L(rhos_in, rhos_out)
    J = construct_J(L)
    return construct_X_from_J(J)

def chi_to_R(X, mats):
    '''
    Chi matrix to Pauli transfer matrix. Mats is an ordered list of
    generalized pauli matrices. Not optimized at all.
    '''
    X = np.array(X)
    R = X*0

    for i, Rrow in enumerate(R):
        Si = mats[i]
        for m, Xrow in enumerate(X):
            Sm = mats[m]
            SiSm = Si*Sm
            for j, Rel in enumerate(Rrow):

```

```
Sj = mats[j]
SiSmSj = SiSm*Sj
for n, Xel in enumerate(Xrow):
    Sn = mats[n]
    R[i,j] = R[i,j] + Xel * np.trace(SiSmSj*Sn)

return np.matrix(R.real) / np.sqrt(R.shape[0])
```

---