

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Optimization of superconducting flux qubit readout using near-quantum-limited amplifiers

Permalink

<https://escholarship.org/uc/item/6v58p6xp>

Author

Johnson, Jedediah Edward Jensen

Publication Date

2012

Peer reviewed|Thesis/dissertation

**Optimization of superconducting flux qubit readout using
near-quantum-limited amplifiers**

by

Jedediah Edward Jensen Johnson

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

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor John Clarke, Chair
Professor Irfan Siddiqi
Professor Theodore Van Duzer

Fall 2012

**Optimization of superconducting flux qubit readout using
near-quantum-limited amplifiers**

Copyright 2012
by
Jedediah Edward Jensen Johnson

Abstract

Optimization of superconducting flux qubit readout using near-quantum-limited amplifiers

by

Jedediah Edward Jensen Johnson

Doctor of Philosophy in Physics

University of California, Berkeley

Professor John Clarke, Chair

Though superconducting qubits offer the potential for a scalable quantum computing architecture, the high-fidelity readout necessary to execute practical algorithms has thus far remained elusive. Moreover, achievements toward high fidelity have been accompanied by either long measurement times or demolition of the quantum state. In this dissertation, we address these issues with the novel integration of two ultralow-noise, superconducting amplifiers into separate dispersive flux qubit measurements. We first demonstrate a flux qubit inductively coupled to a 1.294-GHz nonlinear oscillator formed by a capacitively shunted DC SQUID. The frequency of the oscillator is modulated by the state of the qubit and is detected via microwave reflectometry. A microstrip SQUID (Superconducting QUantum Interference Device) amplifier (MSA) is used to increase the sensitivity of the measurement over that of a semiconductor amplifier. In the second experiment, we report measurements of a flux qubit coupled via a shared inductance to a quasi-lumped element 5.78-GHz readout resonator formed by the parallel combination of an interdigitated capacitor and a meander line inductor. The system noise is substantially reduced by a near-quantum-limited Josephson parametric amplifier (paramp).

We present measurements of increased fidelity and reduced measurement backaction using the MSA at readout excitation levels as low as one hundredth of a photon in the readout resonator, observing a 4.5-fold increase in the readout visibility. Furthermore, at low readout excitation levels below one tenth of a photon in the readout resonator, no reduction in T_1 is observed, potentially enabling continuous monitoring of the qubit state. Using the paramp, we demonstrate a continuous, high-fidelity readout with sufficient bandwidth and signal-to-noise ratio to resolve quantum jumps in the flux qubit. This is enabled by a readout which discriminates between the readout pointer state distributions to an error below one part in 1000. This, along with the ability to make many successive readouts in a time T_1 , permits the use of heralding to ensure initialization to a fiducial state, such as the ground state. This method enables us to eliminate errors due to spurious thermal population, increasing the fidelity to 93.9%. Finally, we use heralding to introduce a simple, fast qubit reset protocol without changing the system parameters to induce Purcell relaxation.

Contents

Contents	i
List of Figures	iv
List of Tables	vii
1 Introduction	1
1.1 Quantum information	2
1.2 Challenges of quantum computing	3
1.3 Superconductivity	3
1.3.1 Flux quantization	3
1.3.2 Josephson junctions	4
1.4 Superconducting Quantum Interference Device	7
1.5 Superconducting qubits	9
1.6 Scope of research	11
2 Superconducting flux qubits	13
2.1 The one-junction flux qubit	13
2.2 The three-junction flux qubit	15
2.3 The four-junction flux qubit	19
2.4 Energy level calculation for the three-junction qubit	22
2.5 Qubit control and measurement schemes	25
3 Readout and amplification	27
3.1 Dispersive measurement	27
3.1.1 Jaynes-Cummings formulation of a flux qubit	27
3.1.2 SQUID resonator	30
3.1.3 Readout signal and noise	31
3.2 Signal amplification	34
3.3 Microstrip SQUID amplifier	38
3.4 Josephson parametric amplifier	41
4 Design and fabrication of devices	47

4.1	Flux qubits	47
4.1.1	Qubit design	48
4.1.2	Resist stack	49
4.1.3	Electron beam lithography	50
4.1.4	Aluminum deposition	52
4.1.5	Oxygen plasma	53
4.1.6	Overlap capacitors	56
4.1.7	Device screening	56
4.1.8	Device layouts and parameters	58
4.2	Microstrip SQUID amplifier	62
4.2.1	Previous work	63
4.2.2	Input coil	64
4.2.3	Junction process	65
4.3	Josephson parametric amplifier	65
5	Measurement infrastructure	67
5.1	Sample boxes	67
5.2	Dilution refrigerator	69
5.3	DC bias circuitry	70
5.4	Microwave circuitry	74
5.4.1	Coaxial lines and attenuation	74
5.4.2	Cryogenic setup	75
5.4.3	Room temperature measurement setup	79
5.4.4	Data acquisition	83
6	Qubit and amplifier characterization	84
6.1	Resonator characterization	84
6.2	Flux qubit - resonator interaction	85
6.3	Qubit spectroscopy	87
6.4	Time domain characterization	90
6.5	Amplifier performance	94
6.5.1	System noise measurements	94
6.5.2	MSA	96
6.5.3	Paramp	98
7	Fidelity and backaction measurements	100
7.1	SQUID resonator	101
7.1.1	Pulse sequence	101
7.1.2	Signal processing	102
7.1.3	Readout improvement with MSA	104
7.2	Circuit QED	109
7.2.1	Quantum jumps	109

7.2.2	Fidelity measurements	111
7.2.3	Readout pointer state discrimination	113
7.2.4	Heralded state preparation	114
7.2.5	Fidelity loss and readout backaction	116
7.2.6	Fast qubit reset	119
8	Conclusions	120
8.1	Summary of results	120
8.2	Future directions	121
	Bibliography	122

List of Figures

1.1	RCSJ model and washboard potential of a Josephson junction.	6
1.2	SQUID schematic and critical current modulation vs flux.	8
1.3	SQUID current-voltage characteristics and voltage modulation vs flux.	9
1.4	A variety of superconducting qubits.	11
2.1	Double-well potential energy of the one junction flux qubit.	14
2.2	Schematic of the three junction flux qubit.	16
2.3	Surfaces of constant potential energy in the three junction flux qubit.	17
2.4	Potential energy slice of a three junction flux qubit with $\delta_t = 0$	19
2.5	Potential energy of a three junction flux qubit with non-negligible β_q	20
2.6	Potential advantages of the four junction flux qubit.	21
2.7	Flux qubit energy spectrum vs flux.	23
2.8	A sampling of flux qubit readout architectures.	26
3.1	Comparison of cavity and circuit QED.	28
3.2	Coupling of cQED readout cavity to external circuitry.	31
3.3	The dispersive qubit shift in a microwave resonator.	32
3.4	Readout pointer state vectors and noise in the IQ plane.	33
3.5	Simplified model of a generic amplifier.	34
3.6	Signal and noise amplification.	37
3.7	MSA overview.	39
3.8	Schematic representation of the paramp.	41
3.9	Response of the driven nonlinear oscillator.	43
3.10	Resonance frequency response with increasing power in a nonlinear oscillator.	43
3.11	Principles of phase-sensitive amplification in the paramp.	45
4.1	Headway PWM32 spinner.	48
4.2	Equipment for electron beam lithography.	50
4.3	Features patterned into electron beam resist.	51
4.4	Aluminum evaporators.	52
4.5	SEM images of Josephson junctions.	54
4.6	Double-angle shadow evaporation.	55

4.7	Probe station and desiccator.	57
4.8	False color optical micrograph of the SQUID resonator device.	58
4.9	False color SEM micrograph of the circuit QED device.	60
4.10	Optical photograph of the MSA with magnified inset of the junction region.	62
4.11	Niobium/aluminum sputtering system including argon ion mill.	64
4.12	Images of a Josephson parametric amplifier.	66
5.1	Flux qubit sample box for SQUID resonator readout.	68
5.2	Flux qubit sample box for circuit QED readout.	69
5.3	MSA cylindrical can and cryoperm shield.	70
5.4	Paramp launch and shielding.	71
5.5	Dilution refrigerator including sample mounts.	72
5.6	Copper powder filter.	73
5.7	Circuit diagram of the low noise, optically-isolated current source.	73
5.8	Cryogenic microwave setup for SQUID resonator readout.	76
5.9	Microwave switches and circulators.	77
5.10	Cryogenic microwave setup for cQED readout.	78
5.11	Room temperature qubit measurement equipment.	79
5.12	Room temperature setup of the SQUID resonator readout.	80
5.13	Room temperature setup of the cQED readout.	81
6.1	Response of the SQUID resonator vs. flux and power.	85
6.2	The qubit step.	86
6.3	cQED resonator response vs. flux in qubit.	87
6.4	Qubit spectroscopy in the SQUID resonator readout.	88
6.5	Readout-induced shift in the qubit flux.	89
6.6	Qubit spectroscopy in the cQED readout.	90
6.7	The Bloch sphere.	91
6.8	Qubit time domain characterizations.	92
6.9	MSA performance.	97
6.10	MSA dynamic range.	97
6.11	Paramp gain.	99
7.1	Pulse sequence for SQUID resonator fidelity measurements.	102
7.2	T_1 decay during measurement.	103
7.3	Visibility vs readout photons in the SQUID resonator/MSA readout.	106
7.4	Calculated measurement distributions.	107
7.5	T_1 during readout vs readout photons.	108
7.6	Quantum jumps.	109
7.7	Histogram of quantum jumps from excited to ground state.	110
7.8	Raw measurements distributions.	111
7.9	Raw fidelity and T_1 during readout vs readout power.	112

7.10 Pure readout pointer state distributions.	114
7.11 Heralded state preparation at $\bar{n} = 14.6$	115
7.12 Three dimensional readout histograms.	116
7.13 Readout backaction.	119

List of Tables

3.1	MSA design parameters.	41
4.1	SQUID resonator readout approximate design parameters (F063010c).	59
4.2	Circuit QED readout approximate design parameters (F032811b).	61
6.1	Measured qubit and resonator parameters.	88
6.2	Characteristic times in the qubit.	93
6.3	Amplification chain noise temperatures.	96
7.1	Fidelity loss in SQUID resonator/MSA readout at $\bar{n} = 1.5$	105
7.2	Fidelity loss in cQED/paramp readout at $\bar{n} = 14.6$	118

Acknowledgments

Although this dissertation is filed under my name alone, it represents the collective effort of many, without whom these results would not have been achieved. I would first like to thank my advisor, Dr. John Clarke, for providing me the opportunity to perform research in his lab. I have greatly benefitted from his consistent personal involvement during every stage of our experiments. He has impressed upon me the importance of experimental detail, meticulous preparation, and absolute clarity in scientific writing. It has been a privilege and honor to associate with an individual of such high caliber. I also acknowledge the tutelage of Dr. Irfan Siddiqi, who in many ways functioned as a second advisor during the collaborative projects that form this thesis. I am especially grateful for his hands-on mentoring in everything from constructing delicate microwave filters to setting up precision qubit measurements.

I was fortunate to work with an exemplary group of colleagues in an atmosphere with exceptional collegiality. The majority of this work was performed alongside an outstanding postdoctoral scholar, Emile Hoskinson. Along with teaching me the nuances of cryogenic equipment, Emile broke me of my experimental timidness and encouraged me to think with scientific creativity. I will always treasure the memories of our lunchtime road bike rides through the Berkeley hills with our fellow labmate Steven Anton. It was over conversations on these rides that I developed a better understanding of the subtleties of flux qubit measurements. Fellow graduate student Chris Macklin was an indispensable part of our efforts, showing a special aptitude for measurement and data acquisition. I am thoroughly indebted to R. Vijay, who truly goes the extra mile to assist others in the lab regardless of the magnitude of his own obligations. I learned the art of device fabrication from Darin Kinion and Daniel Slichter, and am grateful for the patience they exhibited during this process. It has been a true pleasure to collaborate with Professor Jörn Hansen of the Technical University of Denmark during the past four summers. For such an accomplished scholar, he shows a remarkable humility and ability to personally relate to diverse personalities. In addition to Steven Anton and Jeff Birenbaum, who have helped me greatly with recent projects, I would like to acknowledge support and contributions from Sarah Busch, Michael Hatridge, Paul SanGiorgio, Kai Buckenmaier, Sean O’Kelley, Stephen Wu, Shane Cybart, Travis Hime, Sebastien Michotte, Eli Levenson-Falk, Steve Weber, Ofer Naaman, Andrew Fefferman, and Nathan Kelso. The undergraduates in our lab have been top notch, and I thank Derek Murray, Justin Jeet, and Elan Weingarten for their assistance. I am grateful to Barbara Salisbury, Anne Takizawa, Donna Sakima, and Anthony Vitan for their help in navigating the sometimes choppy administrative waters here at Berkeley. I also acknowledge the time dedicated by Professor Theodore Van Duzer to revising this thesis and participating in my Qualifying Examination.

I appreciate the effort invested by the many mentors and teachers who encouraged my academic curiosity and pushed me to maximize my potential. I would like to give special mention to my thesis advisors during my time at Brigham Young University, Dr. David Allred and Dr. Steven Turley. They have each been a superb example of professional excellence while maintaining a proper life balance.

Finally, I would like to thank my entire family for their unconditional support during this thesis. My wife, Marie, deserves special recognition for sacrificing many personal ambitions in behalf of my scholastic aspirations. There are few who can understand the all-encompassing responsibility of raising four children (at one point all three years old and younger) during graduate school. She has been an overarching example of strength, optimism, and principled character, allowing our family to thrive rather than simply scrape by. I owe my greatest treasures in life to her, and eagerly look forward to experiencing life's greatest moments with her by my side.

Chapter 1

Introduction

The theory and experimental application of quantum mechanics are among the most profound scientific endeavors of the 20th century. From abstract origins in radical concepts such as wave-particle duality and probabilistic measurement, quantum mechanics has enabled technological innovations (i.e. magnetic resonance imaging, lasers, magnetic hard drives, etc.) which permeate our society. Quantum mechanics is typically associated with mysterious interactions at the atomic scale, and appears to describe a world which deviates significantly from our intuitive, classical understanding. Early experimental demonstrations of quantum behavior included observations of energy quanta in photons (photoelectric effect) and quantized intrinsic angular momentum, or spin (Stern-Gerlach). A rich formalism was developed which very accurately agreed with non-classical experimental observations such as the Lamb Shift.

Although macroscopic objects are composed of enormous numbers of quantum objects, they generally exhibit classical behavior. This is because the individual degrees of freedom couple in such a way that their quantum observables become randomized. This coupling leads to the formation of an effective energy reservoir through which energy and information can be dissipated from a single quantum object [1]. This process is known as decoherence. To preserve the coherent state of a quantum object, it is necessary to maximally decouple it from surrounding, environmental degrees of freedom. Generally, the smaller the object, the more easily isolated it is.

In the mid 1980's, macroscopic quantum coherence was demonstrated in the lab of John Clarke at UC Berkeley. These experiments provided evidence for quantized energy levels and quantum tunneling out of the zero-voltage state in a current-biased Josephson junction [2, 3, 4]. This opened the door for later experimental observations of quantum coherence in structures on the order of a hundred microns in size [5], an amazing demonstration given the incredibly large number of atoms in a such an object. This dissertation details experiments involving one such device, the superconducting flux qubit. Specifically, we use ultralow-noise superconducting amplifiers to optimize the amount of information that can be extracted during a qubit measurement while minimally perturbing the measured qubit state. These results represent an important step towards the ultimate goal of building a practically useful

quantum computer.

1.1 Quantum information

The fundamental unit of information used in modern computation and data storage is the ‘bit,’ a two-level, digital system. In typical notation the states of the bit are referred to as ‘0’ and ‘1’. In a microprocessor, these states refer to transistors in the off and on states. Given a system of n bits, there are 2^n possible combinations 0 and 1, and the system can only represent one of these at a time.

A quantum bit, or ‘qubit,’ on the other hand, can be put into an arbitrary quantum superposition of eigenstates 0 and 1. Examples of quantum two-level systems include electromagnetic photons, spin 1/2 particles in a magnetic field, and even isolated energy levels of an atom. Both the relative amplitude and phase of 0 and 1 can be controlled in a qubit. Thus, in a n qubit system, all 2^n possible classical combinations can be represented by a highly entangled, superposition state. Rather than executing an algorithm in a serial process, this maximally entangled state allows all possible states to be operated on simultaneously, intrinsically lending itself to massive parallelism [6].

The first notion of a quantum computer is commonly attributed to Richard Feynman in 1982 [7]. Subsequent theoretical work led to the development of quantum algorithms which would offer major speed advantages over their classical counterparts. Shor’s algorithm (1994) showed that integer factorization could be achieved in polynomial time with the aid of the quantum Fourier transform [8]. This time scale is exponentially faster than the most efficient classical factoring algorithm. The ability to factor very large integers is the basis behind breaking RSA (Rivest, Shamir, Adleman) encryption, one of the standard methods for securing electronic information. Another well-known application of quantum computing, made possible by Grover’s algorithm (1996), is the sorting of unstructured databases [9]. Quadratic speedup is obtained through a routine that is more efficient than simply searching through each item in the database, as a classical computer is forced to do. Finally, simulation of quantum systems is often cited as a potential application of a quantum computer [10]. The computational complexity of many-body quantum systems increases rapidly with the increasing size of the system, quickly making the computational time impractical. It is logical that a quantum computer would more efficiently simulate such a complex quantum system.

In recent years, mesoscopic systems, including quantum dots [11] and circuits made of superconductors [12, 13], have been engineered with quantum, two-level properties. Rapid progress has been made in the field of superconducting qubits, validating superconducting circuits as a viable platform for quantum information processing [14, 15, 16, 17].

1.2 Challenges of quantum computing

Many decades elapsed between the advent of the transistor and the commercial availability of desktop computers. Progress in quantum computing seems to be taking a similar trajectory with continuing upward progress, albeit at a modest rate.

To implement a practical quantum computing architecture, a set of conditions known as the DiVincenzo criteria must be satisfied [18]. Expressed succinctly, these requirements are:

1. Qubits must be well-defined and scalable.
2. System can be initialized to a fiducial state.
3. Coherence times are much longer than measurement times.
4. Implementation of a universal set of quantum gates.
5. High fidelity readout of qubit states.

In order to build a practical quantum computer, a scalable architecture is required. Qubits based on smaller structures such as trapped ions [19, 20] are coupled more weakly to their surroundings, making readout and coupling between qubits more difficult. This poses greater challenges to producing a scalable architecture with such qubits. The advantage of these systems is that coherence times are much longer, making it easier for entanglement and measurement to occur before the quantum state is destroyed.

Superconducting qubits and semiconductor-based qubits naturally lend themselves to scalability because they leverage existing microfabrication technology. Such systems are also generally tunable and easily manipulated. This all comes at the expense of shorter coherence times because of stronger coupling to environmental modes and fluctuations. Control lines and the readout circuitry needed to interrogate the qubit contribute to these effects. Over the past decade, superconducting qubit coherence times have continued to improve, exceeding $10 \mu\text{s}$ in recent work [21, 22]. These times are beginning to approach the order of magnitude necessary for certain error correction schemes [23], and engineering qubits with longer coherence times is the subject of intense current research.

1.3 Superconductivity

1.3.1 Flux quantization

When cooled through the superconducting transition temperature T_C , the participating superconducting electrons form Cooper pairs. These pairs form a highly collective condensate, whose properties are expressed by a single wavefunction with an amplitude and a phase, given by

$$\psi(\mathbf{r}) = [n(\mathbf{r})]^{1/2} e^{i\theta(\mathbf{r})}, \quad (1.1)$$

where $n(\mathbf{r})$ is the density of Cooper Pairs and $\theta(\mathbf{r})$ is the phase of the superconducting order parameter [24]. One consequence of this is that in a superconducting ring, $\theta(\mathbf{r})$ must be single-valued modulo 2π , expressed mathematically as

$$\oint_C \nabla\theta(\mathbf{r}) \cdot d\mathbf{l} = 2\pi m, \quad (1.2)$$

where m is an integer. Using the relation in the interior of the superconductor

$$\hbar\nabla\theta = q\mathbf{A}, \quad (1.3)$$

where q is the Cooper pair charge and \mathbf{A} is the magnetic vector potential, it can be shown that the flux

$$\Phi = \frac{mh}{2e}, \quad (1.4)$$

where $q = -2e$ for a Cooper pair. Thus, the flux threading a superconducting loop takes on only discrete values, in integer multiples of the magnetic flux quantum Φ_0 , given by

$$\Phi_0 = \frac{h}{2e}. \quad (1.5)$$

1.3.2 Josephson junctions

Classically, no current should flow when an insulator is placed between two superconducting electrodes. Quantum mechanically, however, the macroscopic wavefunction ψ is exponentially attenuated in the insulating region. For a thin enough barrier, Cooper pairs can tunnel through the insulator, forming a Josephson junction [25]. In this structure, Josephson showed that

$$I = I_0 \sin(\delta) \quad (1.6)$$

$$\frac{\partial\delta}{\partial t} = \frac{2e}{\hbar}V, \quad (1.7)$$

where $\delta = \theta_2 - \theta_1$ is the phase difference across the junction. Here, θ_1 and θ_2 are the phases of the superconducting order parameter on each side of the tunnel barrier. Equation 1.6 is known as the Josephson current-phase relation. It states that the phase across the junction is a function of the supercurrent through the junction (I), up to a maximum given by the junction critical current I_0 . When I_0 is exceeded, a voltage develops across the junction, and the phase across the junction changes in time according to Equation 1.7. For a constant voltage, this constant rate of phase change corresponds to an oscillating current across the junction.

The Josephson relations can be combined to yield the following relation:

$$V = \frac{\hbar}{2eI_0 \cos(\delta)} \frac{\partial I}{\partial t}. \quad (1.8)$$

Recognizing the coefficients in front of the time derivative as an effective inductance, we can define the Josephson inductance as

$$L_J(\delta) \equiv \frac{\hbar}{2eI_0 \cos(\delta)} = \frac{L_{J0}}{\cos(\delta)}, \quad (1.9)$$

where

$$L_{J0} = \frac{\hbar}{2eI_0} = \frac{\Phi_0}{2\pi I_0} \quad (1.10)$$

is the inductance at zero current. Thus, a Josephson junction has an inductance that scales inversely with the junction critical current I_0 and increases as the supercurrent through the junction increases, diverging as I approaches I_0 . In a similar vein, Equations 1.6 and 1.7 can be combined to define an energy associated with the junction:

$$U(\delta) = \frac{\Phi_0 I_0}{2\pi} [1 - \cos(\delta)] = E_J [1 - \cos(\delta)], \quad (1.11)$$

where the coupling energy at zero current is the Josephson energy

$$E_J = \frac{\Phi_0 I_0}{2\pi}. \quad (1.12)$$

In addition to acting like an inductor, a Josephson junction also has capacitive and dissipative conduction channels, as shown in the RCSJ (resistively and capacitively shunted junction) model [26] of Figure 1.1(a). The capacitance is formed between the two superconducting leads of the junction, and the dissipation is due to dielectric loss and quasiparticle conduction (suppressed at low temperatures). Quasiparticles are formed when the binding energy of a Cooper pair is exceeded and the pair is broken [27]. Putting this together, a real-world junction can be represented as by the parallel combination of an ideal Josephson junction and a shunting capacitance (C) and resistance (R). Writing an equation for the total current through this full structure gives

$$I = I_0 \sin(\delta) + \frac{\hbar}{2eR} \frac{\partial \delta}{\partial t} + C \frac{\hbar}{2e} \frac{\partial^2 \delta}{\partial t^2}. \quad (1.13)$$

This equation is analogous to an equation of motion of a mechanical system, with the first three terms associated with effective forces. The third term is the damping term, with dissipative character proportional to $1/R$. The final term is analogous to the mass times acceleration in Newton's second law, where the capacitance is proportional to the mass of the 'phase particle.' Grouping constants appropriately and integrating the first two terms with respect to δ , one can define an effective potential energy U_J of the phase particle, parameterized by the bias current, as

$$U_J(\delta) = -E_J \left[\cos(\delta) + \frac{I}{I_0} \delta \right], \quad (1.14)$$




Figure 1.1. RCSJ model and washboard potential of a Josephson junction. (a) The RCSJ model represents a Josephson junction as the parallel combination of an ideal Josephson element (x) shunted in parallel with a capacitor (C) and a resistor (R). (b) The washboard potential energy is parameterized by the bias current I and plotted as a function of the phase δ .

where a constant E_J term has been dropped. This potential is commonly known as the ‘tilted washboard’ potential [Figure 1.1(b)], and an intuitive sense for the dynamics of Josephson junctions can be obtained by considering the phase particle in this potential. The damping is provided by placing the entire system in a viscous medium. Starting with the potential untilted, the phase particle is confined in one of the wells, corresponding to zero voltage across the junction. As the bias current is increased, the washboard tilts, until the particle reaches the point where it sees a negative slope. It is then free to start running down the washboard, corresponding to exceeding the junction critical current and switching into the voltage or ‘running’ state. In practice, the particle can switch into the voltage state before the true critical current is reached if thermal fluctuations kick the particle out of the well or if the particle tunnels out. In the ‘running’ state, the phase particle reaches a steady state due to the viscous medium opposing motion. When the bias current is reduced, decreasing the slope of the potential, it is possible for the particle to continue running down the washboard. This hysteresis is due to the momentum the particle has acquired in the voltage state. To re-confine the particle to a single well, the slope must be further decreased. This ‘retrapping’ into the supercurrent state occurs at a lower bias current than the original switching current. The degree of hysteresis in the system is dependent on the junction capacitance (mass) and shunting resistance (damping). For some applications, it is desirable to eliminate the hysteresis altogether, which is accomplished by adding an external shunt resistance to reduce the Stewart-McCumber parameter, β_C , to less than one. This parameter

is defined as:

$$\beta_C = \frac{2\pi I_0 R^2 C}{\Phi_0}. \quad (1.15)$$

1.4 Superconducting Quantum Interference Device

The Superconducting QUantum Interference Device (SQUID), shown in Figure 1.2(a), is a well-known application of Josephson junctions. The SQUID consists of two Josephson junctions interrupting a superconducting loop. When the device is biased with a current I_b , the current divides into the two arms of the SQUID, and the phases in each arm interfere so that

$$I_b = I_0[\sin(\delta_1) + \sin(\delta_2)] = 2I_0 \sin\left(\frac{\delta_2 + \delta_1}{2}\right) \cos\left(\frac{\delta_2 - \delta_1}{2}\right), \quad (1.16)$$

where δ_1 and δ_2 are the phase differences across each of the two SQUID junctions. This interference is reminiscent of photons in the classic double-slit experiment. Although Equation 1.16 is specific to identical SQUID junctions, this analysis can also be generalized to asymmetric junctions. Referring back to the derivation of flux quantization in Section 1.3.1, the integral of the phase gradient around the superconducting loop now includes the contribution of the two Josephson junctions, so that Equation 1.4 now becomes

$$\Phi = m\Phi_0 + \frac{\Phi_0(\delta_2 - \delta_1)}{2\pi}. \quad (1.17)$$

Rearrangement gives the more familiar form

$$2\pi \frac{\Phi}{\Phi_0} = \delta_2 - \delta_1 \quad (1.18)$$

for $m = 0$. Substituting Equation 1.18 into Equation 1.16 gives

$$I_b = 2I_0 \sin\left(\frac{\delta_2 + \delta_1}{2}\right) \cos\left(\frac{\pi\Phi}{\Phi_0}\right). \quad (1.19)$$

For simplicity, the geometric loop inductance has been neglected in this derivation, and thus Equation 1.19 is only valid in the limit of vanishing loop inductance. The cosine term on the right hand side of Equation 1.19 is parameterized by the total flux threading the SQUID. Thus, the flux sets a limit on the maximum possible value of I_b , or the critical current, up to a maximum of $2I_0$. The critical current of the SQUID is therefore modulated by the flux through the SQUID, with a flux periodicity of Φ_0 , as shown in Figure 1.2(b).

A total flux differing from $m\Phi_0$ implies from Equation 1.18 that the phases across each junction are different, and therefore the currents in each SQUID arm are also different. This current asymmetry can be viewed as contributing to a circulating current J around the SQUID loop. The total flux in the SQUID Φ can be broken down into two components: the externally applied flux Φ_x and the flux generated by the circulating current J . The




Figure 1.2. SQUID schematic and critical current modulation vs flux. (a) The SQUID consists of a superconducting loop interrupted by two Josephson junctions. The applied current bias (I_b) divides down the two arms of the SQUID. The symmetry is broken by a magnetic flux (Φ) in the loop, producing a circulating current J . (b) The critical current of the SQUID modulates with applied flux as shown for a SQUID with $\beta_L \ll 1$.

latter contribution acts to partially screen the applied flux Φ_x . Modifying 1.18 to reflect this distinction gives

$$\frac{2\pi}{\Phi_0}(\Phi_x + L_S J) = \delta_2 - \delta_1, \quad (1.20)$$

where L_S is the SQUID loop inductance. For non-zero L_S , solving for the SQUID dynamics is non-trivial and requires numerical methods. Increasing L_S for a given value of I_0 causes the the critical current of the SQUID no longer to modulate to zero, but rather to an increasing fraction of the critical current. The relative effect of L_S can be quantified through the ‘screening parameter’ β_L , defined as

$$\beta_L = \frac{2I_0 L_S}{\Phi_0}. \quad (1.21)$$

This parameter can also be written as $L_S / \pi L_J$, which is proportional to the ratio of the loop inductance to the Josephson inductance.

Resistively shunted SQUIDs engineered to be non-hysteretic can be used as flux-to-voltage transducers, as shown in Figure 1.3. When biased just above the critical current in the voltage state, the voltage across the SQUID modulates sinusoidally with applied flux. If the SQUID is biased at the steepest point of the flux-to-voltage curve near $\Phi_x = \Phi_0/4$, where $V_\Phi \equiv (\partial V / \partial \Phi)_{I_b}$ is maximum, a very small flux signal in the SQUID produces a large output voltage. When operated in this manner as a magnetometer, noise spectral densities




Figure 1.3. SQUID current-voltage characteristics and voltage modulation vs flux. (a) The current-voltage (IV) characteristics of a SQUID with a non-negligible (but < 1) β_L are shown for $\Phi_x = 0$ (red) and $\Phi_x = \Phi_0/2$ (blue). (b) When the SQUID is current biased above its critical current at I_b (dotted line), the voltage output across the SQUID modulates sinusoidally with applied flux. Thus, the SQUID functions as a flux-to-voltage transducer, with maximum sensitivity $V_\Phi \equiv (\partial V / \partial \Phi)_{I_b}$ at $\Phi_x = \Phi_0/4$.

below $1 \mu\Phi_0/\sqrt{\text{Hz}}$ have been achieved [28]. For such devices, the optimal performance is reached with $\beta_L \approx 1$ [29]. Unshunted SQUIDs are used for applications requiring non-linear, tunable inductors, flux-sensitive switching detectors, or when dissipation due to on-chip resistors is undesirable.

1.5 Superconducting qubits

The macroscopic quantum nature and low dissipation of superconducting circuits make them natural candidates for use as qubits [30]. Capacitance and inductance are the classical quantities that describe the resonant characteristic of these circuits. A formalism must be developed to quantize these electrical circuits so that the quantum mechanical character of the qubit can be captured [31]. Writing down the well-known Hamiltonian of a LC oscillator, one can identify Q and Φ as canonical conjugate variables, where Q is the electric charge stored in the capacitor and $\Phi = LI$ is a generalized flux [32]. This is in analogy with the canonical position (\hat{x}) and momentum (\hat{p}) operators, and a full mapping between the two systems can be made. Thus, we can verify that an LC circuit behaves like a quantum harmonic oscillator, with energy levels equally spaced by $\hbar\omega$ down to a zero point term, $\frac{1}{2}\hbar\omega$. Like \hat{x} and \hat{p} , the charge and generalized flux can be reinterpreted as quantum mechanical

operators, represented as \hat{Q} and $\hat{\Phi}$, which follow a similar commutation relation

$$[\hat{Q}, \hat{\Phi}] = -i\hbar. \quad (1.22)$$

This is sometimes re-expressed in quantities which are more commonly used for Josephson junctions in the superconducting qubit community:

$$[\hat{n}, \hat{\delta}] = -i, \quad (1.23)$$

where \hat{n} is the number operator of Cooper pairs tunneling through the junction and $\hat{\delta}$ is the phase operator across the junction.

The harmonic oscillator energy ladder produced by linear inductive and capacitive circuit elements is not suitable for use as a qubit. Because the energy levels are equally spaced, it is impossible to isolate two levels $|0\rangle$ and $|1\rangle$ without simultaneously driving resonant transitions to higher levels. An anharmonic potential is needed to create energy levels which are unevenly spaced, thus requiring a nonlinear element. The Josephson junction provides this with its lossless, nonlinear inductance. Both the Josephson energy E_J and the capacitive energy, $E_C = e^2/2C$, set the energy scale of a superconducting qubit. The relative ratio of the two determines whether the eigenstates of the system are definite charge or phase states. As such, it also determines the relative sensitivity of the qubit to fluctuations in either charge or current.

In the early days of the field, between 1999-2002, superconducting qubits were classified in three broad categories: phase, charge, and flux (Figure 1.4). The phase qubit ($E_J/E_C \sim 10^4$) [33] exploits energy levels in the metastable potential well of a current-biased Josephson junction [Figure 1.4(a)]. The eigenstates of this system are states of definite phase across the junction. The charge qubit ($E_J/E_C < 1$) [34] consists of a superconducting island with eigenstates represented by Cooper pair number states on the island [Figure 1.4(b)]. Anharmonicity is provided by a large critical current Josephson junction which forms one side of the island. Energy eigenstates in the flux qubit ($E_J/E_C \sim 100$) [35] are formed by superpositions of distinct flux states generated by counter-rotating persistent currents in a superconducting loop interrupted by (typically) three Josephson junctions [Figure 1.4(c)].

In recent years, hybrid qubits (for example, quntronium [36]) and especially qubits designed to be less sensitive to noise, such as capacitively-shunted flux qubits [37], transmons [38], and fluxonium qubits [39], were developed. The transmon [Figure 1.4(d)], which has similarities to phase and charge qubits, is currently widely studied. Large shunting capacitors reduce E_C , and hence the sensitivity to charge fluctuations. This reduced sensitivity to charge noise is accomplished by engineering the device for flatter energy bands with respect to charge. Ambient fluctuations therefore result in lower amplitude excursions of the qubit Larmor frequency ω_{01} , and consequently lead to less decoherence. The tradeoff for flatter energy bands is typically reduced anharmonicity of the higher energy levels. Greater care must be taken to reduce the spectral width of qubit control pulses to prevent spurious excitations to these higher levels.

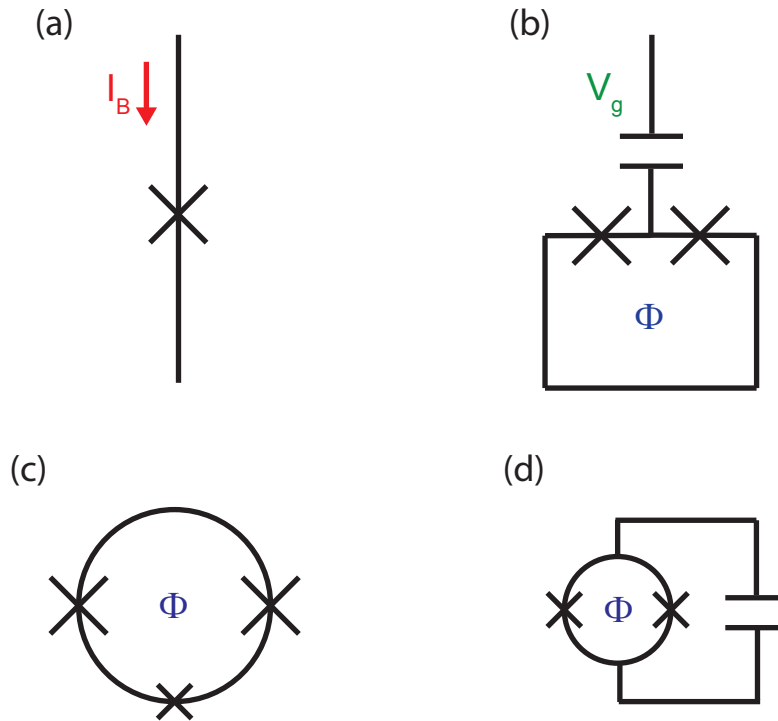


Figure 1.4. A variety of superconducting qubits. (a) A current-biased Josephson junction can be operated as a phase qubit. (b) A split Cooper pair box, a variety of charge qubit, is formed by a superconducting island and a Josephson nonlinearity. The qubit is biased with a gate voltage V_g and a flux Φ . (c) Three-junction flux qubits consist of a loop interrupted by three Josephson junctions, one of which is smaller than the other two. (d) The transmon is a newer qubit design which features reduced sensitivity to charge fluctuations due to large shunting capacitors.

1.6 Scope of research

This thesis focuses on optimizing the dispersive readout of a superconducting flux qubit. In this context, dispersive readout refers to a readout which probes the qubit-state dependent shift in the frequency of a microwave resonator. Information from the quantum state of the qubit is mapped onto an impedance shift and processed at room temperature with a classical detector [40]. There is a delicate interplay between how much information is extracted during a readout and how much the readout perturbs the qubit state. When the readout strength of a quantum system is increased, more information about the state is extracted as it is strongly projected onto the measurement basis. The side effect of this high power is increasing readout backaction, which scrambles the state of the qubit after readout is complete. This poses immense problems when successive qubit interrogations are required. A possible remedy for this is to reduce the readout strength, but this comes at the expense

of less information about the qubit state due to the finite signal-to-noise ratio available in any qubit readout chain. The system can only be probed for a limited amount of time before the quantum state is no longer coherent, with its information irreversibly lost to the environment.

Given these considerations, we focus on optimizing the readout of the flux qubit. This is accomplished with the use of low-noise superconducting amplifiers [41, 42, 43, 44, 45, 46, 47]. Advancements in these custom microwave amplifiers have been driven recently by applications in both quantum computing and other areas of physics requiring a very sensitive measurement [48]. In the course of this optimization process, we address three of the five requirements of a practical quantum computer, as outlined in Section 1.2: qubit state initialization, fast acquisition, and high fidelity readout.

Chapter 2

Superconducting flux qubits

2.1 The one-junction flux qubit

Although the one-junction flux qubit has been replaced in the qubit community by corresponding three and four junction flux varieties, it serves as an instructive starting point for developing an intuitive understanding of flux qubit dynamics. This simple flux qubit design consists of a superconducting loop interrupted by a single Josephson junction [49]. The junction can be characterized by some capacitance C_J and critical current I_0 . The system parametrization is completed by defining a geometric loop inductance, L_q .

The Hamiltonian for this system can be written down in a straightforward manner:

$$H = \frac{1}{2}C_J V^2 + \frac{1}{2}L_q I_q^2 + E_J \cos \delta, \quad (2.1)$$

where I_q is the persistent circulating current in the loop. Starting from Equation 1.18, one of the phase terms can be removed to give the following relationship for a superconducting loop with only one junction:

$$2\pi \frac{\Phi}{\Phi_0} = \delta. \quad (2.2)$$

Using this and noting that the total flux $\Phi = \Phi_x + \Phi_q$, where $\Phi_q = L_q I_q$ is the flux generated by the circulating current, we can write Equation 2.1 as

$$H = 4E_C \left(\frac{Q^2}{2e} \right) + \frac{(\Phi - \Phi_x)^2}{2L_q} - E_J \cos \left(\frac{2\pi\Phi}{\Phi_0} \right). \quad (2.3)$$

The last two terms can be viewed as an effective potential energy, while the kinetic energy comes from the capacitive first term. The overall potential shape as a function of Φ is the combination of a Josephson washboard potential U_J and an inductive, quadratic potential U_L (Figure 2.1). Because there is explicit dependence on Φ_x in only one of the terms, the two potentials can be shifted relative to each other with changes in the applied flux. When the loop is biased with a flux $\Phi_x = \Phi_0/2$, the inductive energy minimum lines up with a




Figure 2.1. Double-well potential energy of the one junction flux qubit. The potential energy normalized to E_J of the one junction flux qubit is plotted vs the total flux threading the qubit. Three different values of the externally applied flux Φ_x are shown: $0.45 \Phi_0$ (green dashed), $0.50 \Phi_0$ (blue solid), and $0.55 \Phi_0$ (red dashed-dot). The two wells shift up and down in energy with changing Φ_x and represent classical circulating current states in the qubit.

maximum in the washboard potential at zero total flux. Degenerate minima form on both sides of this maximum, creating a double-well potential separated by an energy barrier. This point is aptly called the ‘degeneracy point.’ The classical wells are associated with flux states of opposite polarity generated by counter-rotating circulating currents of magnitude I_q in the loop, represented by $|\circlearrowleft\rangle$ and $|\circlearrowright\rangle$. If the external flux is slightly changed from $\Phi_0/2$, the wells tilt such that the height of one well is higher than the other.

If the characteristic background (thermal and electrical) energy is lower than the barrier height, the ‘phase particle’ will classically be confined to either of the two wells. Quantum mechanically, however, it is possible for the bound-state wavefunctions associated with each circulating state to penetrate into the barrier. This implies a non-zero tunneling rate between the two wells. This rate is sensitively dependent on the barrier height, mass (capacitance) of the phase particle, and the degree of tilt in the potential. At the degeneracy point, the interaction between the two wells lifts the degeneracy of the two circulating states, producing an energy splitting Δ_q and energy eigenfunctions which are symmetric ($|0\rangle$) and antisymmetric ($|1\rangle$) superpositions of $|\circlearrowleft\rangle$ and $|\circlearrowright\rangle$. As the external flux is progressively

biased away from $\Phi_0/2$, the two eigenstate wavefunctions are increasingly localized in the separate wells. Thus, the mixture of circulating states changes so that the eigenstates more closely resemble the classical circulating current states.

Though the one junction flux qubit has been experimentally demonstrated [49], it suffers from a number of drawbacks. Foremost, the barrier height and hence the tunneling frequency are sensitively dependent on the ratio of the geometric loop inductance to the Josephson inductance, which is proportional to β_L . Unless $\beta_L \approx 1$, the potential energy will not form a double-well with a corresponding energy splitting suitable for a practical flux qubit. Precisely balancing the critical current of the junction with the appropriate loop inductance is a tricky fabrication challenge. Additionally, because the size of the loop and the critical current are so closely linked, changing the size of the qubit loop has a dramatic effect on the energy levels. These difficulties led to the search for a more robust type of flux qubit.

2.2 The three-junction flux qubit

The three-junction flux qubit [50], shown in Figure 2.2, consists of a superconducting loop interrupted by three Josephson junctions, two of which are equal in size and larger than the third junction. This flux qubit addresses the main problems associated with one junction flux qubits. The main reasoning behind the three junction design is to use the two larger junctions in place of the geometric loop inductance. Because the three junctions are all fabricated in the same step with the same critical current density J_C , their relative Josephson inductances are set by the ratio of their junction areas. This ratio, which is much more easily controlled in device fabrication, is defined as

$$\alpha_q = \frac{2I_{03}}{I_{01} + I_{02}}, \quad (2.4)$$

where I_{0i} are the critical currents of each junction; I_{03} refers to the smaller junction. In this three junction configuration, the energy level structure is much less sensitive to the overall scaling of the critical currents given by the J_C of the process. Because the geometric loop inductance is typically made to be much less than the Josephson inductance of each large junction, a wide range of loop sizes is possible before a significant effect on the energy band structure is seen. This allows for large loop areas which are more suitable for on-chip bias lines and stronger qubit-qubit magnetic flux coupling [51].

If there is a price to pay for these advantages, it is the increased complexity of the flux qubit energy manifold. The two additional junctions add two new degrees of freedom to both the potential energy function and the effective kinetic energy (junction capacitances) [52]. Adapting Equation 2.2 to this three junction geometry gives

$$2\pi \frac{\Phi}{\Phi_0} = \delta_1 + \delta_2 + \delta_3, \quad (2.5)$$

where δ_1 and δ_2 are the phases across the large junctions and δ_3 is the phase across the small junction. The total potential energy U_T can be written as the sum of U_L (unchanged from




Figure 2.2. Schematic of the three junction flux qubit. The three junction flux qubit is a superconducting loop interrupted by three Josephson junctions, one of which is smaller than the other two by a factor of α_q . Each junction is modeled to include both Josephson and capacitive conduction channels. The circulating current I_q flowing in the loop is related to the total flux Φ .

the one junction potential) and U_J , which now includes all three junctions:

$$U_T = \frac{(\Phi - \Phi_x)^2}{2L_q} - \frac{\Phi_0}{2\pi} [I_{01} \cos(\delta_1) + I_{02} \cos(\delta_2) + I_{03} \cos(\delta_3)] \quad (2.6)$$

$$= \frac{(\Phi - \Phi_x)^2}{2L_q} - E_{J1} [\cos(\delta_1) + \cos(\delta_2) + \alpha_q \cos(\delta_3)]. \quad (2.7)$$

This expression assumes that $I_{01} = I_{02}$ and makes use of the definition $E_{J1} = \Phi_0 I_{01} / 2\pi$. Constant energy surfaces of this potential as a function of all three phase variables are shown in Figure 2.3. The ellipsoidal, closed surfaces (red) suggest the existence of periodic minima in the potential.

To gain some insight into this three-dimensional potential, it is helpful to make a coordinate transformation. We define new rotated phase coordinates as

$$\delta_t = \frac{\alpha_q}{1 + 2\alpha_q} \left(2\pi \frac{\Phi_x}{\Phi_0} - \delta_1 - \delta_2 - \delta_3 \right) = -\frac{2\pi L_q I_q \alpha_q}{\Phi_0 (1 + 2\alpha_q)}, \quad (2.8)$$

$$\delta_s = \frac{1}{2(1 + 2\alpha_q)} \left[2\alpha_q \left(\delta_3 - 2\pi \frac{\Phi_x}{\Phi_0} \right) - \delta_1 - \delta_2 \right], \text{ and} \quad (2.9)$$

$$\delta_a = \frac{(\delta_1 - \delta_2)}{2}. \quad (2.10)$$




Figure 2.3. Surfaces of constant potential energy in the three junction flux qubit. Three potential energy contours at $U_T = 1.4 E_J$ (red), $U_T = 10 E_J$ (green), and $U_T = 30 E_J$ (blue) are shown as a function of the phase differences δ_i across the three Josephson junctions. Ellipsoidal surfaces (red) indicate the presence of local minima bound states in the potential. The external flux is $\Phi_x = \Phi_0/2$ and the qubit parameters are $\alpha_q = 0.8$ and $\beta_q = 0.4$ [52].

We also define β_q , analogous to β_L in the one-junction flux qubit, as

$$\beta_q = \frac{2\pi L_q}{\Phi_0} \left(\frac{1}{I_{01}} + \frac{1}{I_{02}} + \frac{1}{I_{03}} \right)^{-1} \quad (2.11)$$

$$= \frac{2\pi L_q \alpha_q I_{01}}{(1 + 2\alpha_q)\Phi_0}. \quad (2.12)$$

In this rotated representation, U_L is given by

$$U_L = E_{J1} \frac{1 + 2\alpha_q \delta_t^2}{2\alpha_q \beta_q} \quad (2.13)$$

and U_J by

$$U_J = -E_{J1} \left[\cos(\delta_a - \delta_s - \delta_t) + \cos(\delta_a + \delta_s + \delta_t) + \alpha_q \cos \left(2\delta_s + 2\pi \frac{\Phi_x}{\Phi_0} - \frac{\delta_t}{\alpha_q} \right) \right]. \quad (2.14)$$

Just as in the case of the one junction flux qubit, the classical circulating current states can be identified as wells in this potential. The situation can be moderately simplified in the case of small loop inductance ($\beta_q \ll 1$). The inductive potential energy given by Equation 2.13 is only a function of the total phase δ_t , and a very small β_q creates a very steep potential in δ_t . For the overall potential be minimized, this implies that $\delta_t \approx 0$. Thus, we can decrease the dimensionality of the potential by one and look only in the $\delta_s - \delta_a$ plane. These phase coordinates represent symmetric (s) and antisymmetric (a) modes. This surface is shown as a contour plot in Figure 2.4. The potential is periodic, the expected form a Josephson element, in both coordinates. An array of minima form in the surface, with adjacent wells separated by two distinctly different barriers for typically used parameters. The smaller or ‘intra-cell’ barrier ($P_0 \rightarrow P_1$) gives the energy splitting which forms the two eigenstates of the qubit. This is analogous to the energy barrier in the double-well, single junction flux qubit potential. These two-well cells are separated from identical, periodic cells by the ‘inter-cell’ barrier, ($P_0 \rightarrow P_2$). This much higher barrier drastically reduces the coupling and tunneling rate between the wells in separate cells. Nevertheless, this serves to split the energy levels of both the ground and excited states into a doublet. For a properly engineered flux qubit, this doublet splitting can be much less than the intrinsic linewidth of the qubit resonance, rendering it insignificant and effectively preserving the two-state qubit system.

If the qubit loop inductance is large such that β_q is no longer an insignificant fraction of one, the potential in δ_t softens, and all three dimensions must be considered [52]. In particular, the potential minima no longer lie on the $\delta_t = 0$ plane, but are now split so that one lies on each side of this plane, as seen in Figure 2.5. This can be understood by the linear dependence of δ_t on I_q , as highlighted in Equation 2.8. The proportionality factor, which includes L_q , maps the circulating current states of equal magnitude and opposite sign onto corresponding states in δ_t . The dependence on δ_a and δ_s still remains, and even though it is more difficult to visualize in this three-dimensional space, corresponding wells again form with both intra-cell and inter-cell splittings. Additionally, an increasing β_q increases the flux generated by the loop, modifying the magnitude of I_q at $\Phi_x = \Phi_0/2$.




Figure 2.4. Potential energy slice of a three junction flux qubit with $\delta_t = 0$. In flux qubits with $\beta_q \ll 1$, minima of the potential energy approximately form in the $\delta_t = 0$ plane, effectively restricting the potential to two dimensions in δ_s and δ_a . Both the ‘intra-cell’ ($P_0 \rightarrow P_1$) and ‘inter-cell’ ($P_0 \rightarrow P_2$) barriers are clearly shown. The external flux is $\Phi_x = \Phi_0/2$ and the qubit parameters are $\alpha_q = 0.8$ and $\beta_q = 0.01$ [52].

2.3 The four-junction flux qubit

A four-junction flux qubit is sometimes used in place of the three-junction variety. The fourth junction serves to add inductance in a manner similar to the other two large junctions in the three-junction flux qubit. This modifies the value of α for a desired Δ_q . The potential additionally becomes four dimensional, adding even more complexity to solving the system of equations. Four-junction flux qubits offer improved coherence times in certain readout schemes due to improved device symmetry [53]. The double-angle shadow evaporation technique typically used to fabricate flux qubits requires that any loop have an even number of junctions. In a three-junction flux qubit, a fourth junction is actually always present, but it is intentionally made with such a large critical current that its effect is negligible on the




Figure 2.5. Potential energy of a three junction flux qubit with non-negligible β_q . Slices of the potential are shown in two orthogonal directions, the $\delta_s - \delta_t$ (left) and $\delta_s - \delta_a$ (right) planes. The planes intersect along the dotted line. The counter-rotating circulating states are visible as potential minima with an equal and opposite δ_t . The external flux is $\Phi_x = \Phi_0/2$ and the qubit parameters are $\alpha_q = 0.6$ and $\beta_q = 0.4$ [52].

qubit dynamics. In readout geometries where parts of the qubit arms also form the readout SQUID, as in Figure 2.6, this large, spurious fourth junction can cause impedance asymmetries between the two arms of the qubit. This causes noise currents to divide differently down these two arms, destroying the noise-protecting symmetry of the device by producing an enhanced coupling of the noise to the qubit loop. If the fourth junction is instead made to be the same size as the two other large junctions, and is strategically placed such that it does not cause asymmetries in the readout, this is referred to as a four-junction flux qubit. For all the qubits studied in this thesis, the readout was such that there was no added benefit from the four-junction qubit design.




Figure 2.6. Potential advantages of the four junction flux qubit. (a) A SQUID which shares sections of its loop arms with a flux qubit is one architecture for qubit readout. A three-junction flux qubit creates an asymmetry in this readout circuit due to its large fourth ‘stray’ junction. This leads to a difference in the impedance of the SQUID arms, causing noise currents flowing through the SQUID to divide unevenly and couple flux noise to the qubit. (b) In a four-junction flux qubit, the large fourth junction is made smaller (the same size as the other two large qubit junctions) and is placed where it will not produce an asymmetry in the readout SQUID. This preserved symmetry helps protect the qubit from decoherence.

2.4 Energy level calculation for the three-junction qubit

An approximate feel for how Δ_q depends on device parameters can be obtained through semi-classical methods. For a very small loop inductance ($\beta_q \ll 1$), it can be shown that [54]

$$\Delta_q \approx 4\sqrt{\frac{E_J E_C (2\alpha - 1)}{\alpha}} \exp \left[\sqrt{\frac{E_J (2\alpha + 1)}{\alpha E_C}} \left(\arccos \frac{1}{2\alpha} - \sqrt{4\alpha^2 - 1} \right) \right]. \quad (2.15)$$

The exponential sensitivity of the tunnel splitting on α , E_J , and E_C is most striking, and this dependence is the cause of most fabrication-related difficulty.

A more rigorous approach can be taken to calculate the energy levels of the flux three junction flux qubit [52]. The Hamiltonian must first be formed, which requires a kinetic term in addition to the potential energy formalism of Section 2.2. In a one-dimensional kinematic system, the kinetic energy is proportional to the square of the time derivative of the position coordinate. The flux qubit kinetic energy term is again analogous to this system. For a single junction, the capacitive energy takes the form

$$E_C = \frac{1}{2} C_J V_J^2. \quad (2.16)$$

Using the second Josephson relation (Equation 1.7), and including the contribution of all three flux qubit junctions, the total capacitive energy can be written as a kinetic energy term:

$$T = \frac{1}{2} \left(\frac{\Phi_0}{2\pi} \right)^2 (C_1 \dot{\delta}_1^2 + C_2 \dot{\delta}_2^2 + C_3 \dot{\delta}_3^2). \quad (2.17)$$

This kinetic energy is also transformed into the same rotated coordinate system as the potential energy. With the full Hamiltonian H , we can now solve for the energy levels and wavefunctions of the system. This is performed by expanding the Hamiltonian in a complete set of basis functions. The basis functions are carefully chosen such that they capture the symmetries of the potential. The periodic potential in δ_s and δ_a suggests the use of plane wave basis functions. Due to the quadratic nature of the potential in δ_t , harmonic oscillator wavefunctions are used in this dimension. The overall expansion functions are written as the following product state:

$$|\psi_{klm}\rangle = |\psi_k^a\rangle |\psi_l^s\rangle |\psi_m^t\rangle, \quad (2.18)$$

where k and l run from $-\infty$ to ∞ and m runs from 0 to ∞ to produce an exact solution. In practice, this expansion is truncated such that the calculation retains sufficient accuracy yet is completed in a reasonable amount of time. The matrix elements can be computed analytically, and the resulting matrix is diagonalized to give the approximate energy eigenvalues and wavefunction composition. If this calculation is made for many individual values of Φ_x , the energy band structure vs flux is generated. An example of such a plot over a narrow flux




Figure 2.7. Flux qubit energy spectrum vs flux. The two lowest energy levels of the flux qubit, E_0 and E_1 , are used as the qubit states. The high degree of anharmonicity in the system is evident by the significantly higher E_2 level. Near $\Phi_x = \Phi_0/2$, the transition energy as a function of flux is well approximated by a hyperbolic function. Each energy level is a doublet with a splitting too fine to be resolved in the plot. The calculation parameters are $E_C = 2$ GHz, $E_J = 200$ GHz, $\alpha_q = 0.6$, and $\beta_q = 0.02$.

range centered about $\Phi_x = \Phi_0/2$ is shown in Figure 2.7. Each visible energy level is actually a doublet with a spacing between levels which is too small to resolve.

The high degree of anharmonicity in the system is evident as the transition energy between levels 1 and 2 (E_{12}) is considerably greater than between the qubit levels 0 and 1 (E_{01}). This allows these higher levels to be neglected in the formation of an approximate two-level Hamiltonian, written as

$$H = -\frac{1}{2}(\epsilon\sigma_z + \Delta_q\sigma_x), \quad (2.19)$$

where σ_z and σ_x are the standard Pauli matrices and $\epsilon = 2I_q(\Phi_x - \Phi_0/2)$ is the energy bias [55]. The circulating current states can be expressed as the basis vectors of σ_z [56]:

$$|\circlearrowleft\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad (2.20)$$

$$|\circlearrowright\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \quad (2.21)$$

Diagonalizing the Hamiltonian, it can be shown that the energy eigenvalues are

$$\begin{bmatrix} E_0 \\ E_1 \end{bmatrix} = \begin{bmatrix} -\frac{1}{2}\sqrt{\Delta_q^2 + \epsilon^2} \\ \frac{1}{2}\sqrt{\Delta_q^2 + \epsilon^2} \end{bmatrix}, \quad (2.22)$$

so that the transition energy takes on the hyperbolic form

$$E_{01} = \sqrt{\Delta_q^2 + \epsilon^2}. \quad (2.23)$$

The eigenvectors of the Hamiltonian can similarly be shown to be

$$|0\rangle = \begin{pmatrix} \cos \frac{\xi}{2} \\ \sin \frac{\xi}{2} \end{pmatrix} \quad (2.24)$$

$$|1\rangle = \begin{pmatrix} -\sin \frac{\xi}{2} \\ \cos \frac{\xi}{2} \end{pmatrix}, \quad (2.25)$$

with the angle ξ defined through the relation

$$\tan(\xi) = \frac{\Delta_q}{\epsilon}. \quad (2.26)$$

Equations 2.24 and 2.25 reveal how the circulating state mixture differs between the energy eigenstates as a function of flux. When $\epsilon = 0$ at the degeneracy point, $|0\rangle$ and $|1\rangle$ are both equal mixtures of the circulating current states, making the eigenstates indistinguishable to a measurement which probes the average magnetization of the state. In the limit of $|\epsilon| \gg \Delta_q$, the eigenstates approach pure circulating states of opposite direction and the energy spectrum becomes linear with a slope of magnitude

$$\left| \frac{\partial E}{\partial \Phi} \right| = |2I_q|. \quad (2.27)$$

2.5 Qubit control and measurement schemes

As seen in Equation 2.23, the flux qubit transition energy is easily tuned by changing the magnetic flux threading the qubit loop. Experimentally, this is accomplished by passing current through a control line which is coupled to the qubit loop via a mutual inductance. This flux bias must be capable of sweeping the flux through an entire flux quantum to ensure that the qubit can be operated at any bias. This is accomplished either via on-chip control lines or by an external multi-turn coil. These bias lines are typically superconducting to ensure that high currents can be used without ohmic heating.

Transitions between qubit states, and more generally arbitrary qubit control pulses, are generated through resonant excitation. This is accomplished via either an oscillating magnetic flux or an oscillating current in the qubit. One common method of flux excitation involves coupling the qubit via a mutual inductance to an inductive termination of a high-bandwidth line. Microwave pulses are sent down this line to the qubit, and the qubit can absorb photons on resonance. If the qubit is directly coupled through a shared line to a microwave source, resonant phase oscillations can also drive qubit transitions.

Readout of the qubit state can be achieved in a variety of ways. Because flux qubit eigenstates are superpositions of magnetic flux states, a magnetometer is a natural choice for qubit readout. This is realized by coupling a SQUID to the flux qubit via a mutual inductance. The flux generated by the qubit circulating states couples to the SQUID, modifying its effective critical current. An unshunted, hysteretic SQUID forms a latching detector in which one pulses the SQUID current bias with a current I_p near its critical current for a duration t_{meas} [Figure 2.8(a)] [35]. These pulse parameters can be adjusted such that the SQUID tunnels into the voltage state with some probability P_{switch} between 0 and 1. The two circulating states of the qubit correspond to different values of P_{switch} due to the flux-dependent SQUID critical current, forming the basis for detection of the qubit state. The principle drawback of this method is the the dissipative nature of the SQUID once it has switched into the voltage state, leading to qubit decoherence. The quasiparticles generated by the Josephson junctions in the voltage state also have a characteristic recombination time which sets a limit on the repetition rate of the measurement [57], an experimental inconvenience and a disadvantage for practical quantum computing.

These issues can be overcome by operating the SQUID entirely in the superconducting state [58]. A SQUID shunted with a capacitor forms a nonlinear LC resonant circuit [Figure 2.8(b)]. As the circulating current J in the SQUID is flux-dependent, and because of the current-dependent Josephson inductance, the SQUID functions a variable inductor whose total inductance is a function of flux. The circulating current states of the flux qubit modify the inductance, and hence the resonant frequency, of the SQUID-resonator. The circulating current states of the qubit map onto distinct shifts in the resonant frequency of the circuit, which can be probed using standard techniques.

The flux qubit measurements described in Figures 2.8(a) and (b) sense the circulating current states of the qubit, and thus measure along a different axis than the qubit eigenstates, which are superpositions of the circulating current states. Because a magnetization




Figure 2.8. A sampling of flux qubit readout architectures. (a) An unshunted dc SQUID realizes a latching qubit readout. The SQUID is probed with current pulses of magnitude I_p and duration t_{meas} . The circulating-state-dependent probability that the SQUID switches into the voltage state is used to discriminate the qubit state. (b) A SQUID shunted with a capacitor forms a nonlinear, flux-dependent oscillator. The circulating current states correspond to slightly shifted resonant frequencies in the oscillator. (c) The SQUID is replaced with a linear inductor, which is sensitive to the magnetic susceptibility χ of the qubit. Qubit-state-dependent shifts in the resonant frequency are again produced.

measurement projects the state of the qubit onto the circulating state detection basis, there is an intrinsic limit to how well the qubit eigenstates can be distinguished. This limit is flux bias dependent, and reaches a minimum at the degeneracy point, where the qubit states are completely indistinguishable in this measurement architecture. The magnetization signal between the qubit states is proportional to the first derivative of the energy spectrum vs flux, $\partial E/\partial\Phi$, which is zero at the degeneracy point. The tradeoff, however, is that the vanishing slope serves to protect the qubit from dephasing induced by ambient magnetic field fluctuations. Strategies for working around these issues and more detailed implications of the circulating current readout basis will be discussed in subsequent chapters.

An alternative measurement scheme [59], which probes the second derivative of the energy spectrum rather than the first derivative, addresses the issue of measurement at the degeneracy point [Figure 2.8(c)]. The curvature of the spectrum, $\partial^2 E_{01}/\partial\Phi^2$, is maximum at the degeneracy point, optimally combining maximum readout signal with minimum sensitivity to flux noise. This corresponds to a measurement of the magnetic susceptibility χ , or equivalently the effective inductance, of the flux qubit. This signal is maximum at the degeneracy point because small changes in an external flux applied to the qubit correspond to large changes in the composition of the circulating current states, maximally changing the magnetic energy.

Chapter 3

Readout and amplification

3.1 Dispersive measurement

3.1.1 Jaynes-Cummings formulation of a flux qubit

Atoms coupled to a photon field in a Fabry-Perot cavity have been extensively studied in the atomic physics community. In this so-called cavity Quantum Electrodynamics (QED) as seen in Figure 3.1(a), the atom can be considered a two level system if the transition between two levels can be isolated. The cavity is mirrored on both ends to produce a resonant cavity with harmonic modes, and the atom couples weakly to the cavity with a strength g . If the cavity and atom are brought into resonance, g represents the rate at which energy is transferred back and forth between the atom and the cavity. The photon escape rate from the cavity is characterized by the parameter κ . Circuit elements can be used to form a system analogous to cavity QED [60, 61]. Replacing the cavity with an electrical LC oscillator and the atom with a superconducting qubit results in circuit QED (cQED). As shown in Figure 3.1(b), the system can again be parameterized by g and κ in an analogous manner.

The general Hamiltonian of a harmonic oscillator coupled to a two-level system can be written as

$$H = H_{\text{res}} + H_{\text{qubit}} + H_{\text{int}}, \quad (3.1)$$

where H_{res} is the Hamiltonian of the harmonic oscillator, H_{qubit} is the Hamiltonian of the qubit, and H_{int} represents the interaction energy [62]. For a flux qubit, this expression can be written as

$$H = \hbar\omega_r \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right) - \frac{1}{2}(\epsilon \hat{\sigma}_z + \Delta_q \hat{\sigma}_x) + \hbar g(\hat{a}^\dagger + \hat{a}) \hat{\sigma}_z, \quad (3.2)$$

where \hat{a}^\dagger and \hat{a} are the standard harmonic oscillator raising and lowering operators, ω_r is the resonant frequency of the bare harmonic oscillator, and $\hbar g$ is a characteristic coupling energy [59]. The final term in Equation 3.2 represents a coupling between the qubit and resonator via a magnetic dipole interaction. As seen in Section 2.4, the eigenstates of $\hat{\sigma}_z$ represent the circulating current states of the qubit, making $\hat{\sigma}_z$ an intuitive operator to associate with the

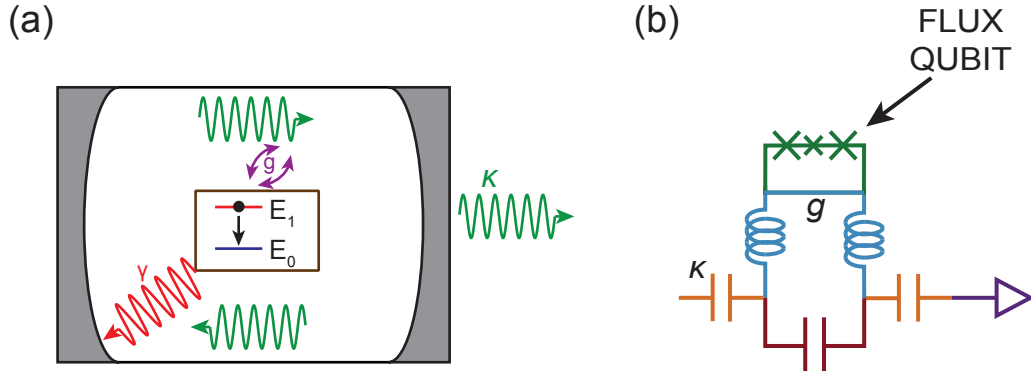


Figure 3.1. Comparison of cavity and circuit QED. (a) Cavity QED: an atomic transition coupled with a strength g to a Fabry-Perot cavity with decay rate κ . The atom is shown incoherently decaying to its environment with a rate γ . (b) An analogous system composed of circuit elements is the basis of circuit QED. The flux qubit is coupled to a microwave resonator through a shared inductance, while coupling capacitors define the decay rate of the LC resonance.

magnetic moment. The interaction term dresses the states of the system such that they are different from the bare resonator and qubit energies. If we now rotate the eigenbasis of the system with the following coordinate transformations

$$\hat{\sigma}_z' = \sin(\theta)\hat{\sigma}_x + \cos(\theta)\hat{\sigma}_z \quad (3.3)$$

$$\hat{\sigma}_x' = -\cos(\theta)\hat{\sigma}_x + \sin(\theta)\hat{\sigma}_z, \quad (3.4)$$

where $\sin(\theta) = \Delta_q/E$ and $\cos(\theta) = \epsilon/E$, the eigenstates of the qubit are also the eigenvectors of $\hat{\sigma}_z'$. If the system is transformed into the interaction picture, terms appear with exponential arguments which are functions of both the sum and difference of the bare qubit and resonator frequencies, given by $\omega_r + \omega_q$ and $\omega_r - \omega_q$, respectively; ω_q is the qubit frequency. Invoking the rotating wave approximation, we can drop the rapidly oscillating sum terms. Additionally, if the qubit and resonator frequencies are far detuned in the dispersive limit, $|\omega_q - \omega_r| \gg g$, we can expand the Hamiltonian in powers of g/δ , where $\delta = |\omega_q - \omega_r|$, and retain only the lowest order terms. Transforming back to the Schrödinger picture, the approximate Hamiltonian in the dispersive approximation becomes

$$H = \hbar \left(\omega_r - \frac{g'^2}{\delta} \hat{\sigma}_z' \right) \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right) - \frac{1}{2} \hbar \omega_q \hat{\sigma}_z', \quad (3.5)$$

where the effective coupling strength is $g' = g \sin(\theta)$. Thus, the coupling is maximum when $\sin(\theta) = 1$, corresponding to measurement at the degeneracy point. The second term in equation 3.5 represents the qubit energy, while the first term looks like a harmonic oscillator modified by a term which includes $\hat{\sigma}_z'$. In other words, the resonant frequency of the harmonic

oscillator is modified by a term which is dependent on the state of the qubit. The frequency shift is in fact given by

$$\chi = \frac{g'^2}{\delta}. \quad (3.6)$$

Each state of the qubit shifts the frequency in a different direction, such that the full frequency shift between the qubit states is 2χ . Consequently, probing the resonant frequency of the resonator reveals the state of the qubit.

It is also informative to rearrange the terms in the Hamiltonian slightly differently. Casting Equation 3.5 in another suggestive form yields

$$H = \hbar\omega_r \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right) - \frac{1}{2} \hbar (\omega_q + 2\chi \hat{a}^\dagger \hat{a} + \chi) \hat{\sigma}'_z \quad (3.7)$$

Now, the first term looks like a typical harmonic oscillator energy, while the second term represents the qubit with a frequency modified by two additional terms. The second of these new terms can be identified as a constant shift in the cavity frequency, or the well-known Lamb shift. The other term includes the number operator $\hat{N} = \hat{a}^\dagger \hat{a}$, making the effective qubit frequency a function of the number of photons in the resonator. This shift is referred to as the ac Stark shift, and can experimentally be exploited to yield a precise calibration of the readout strength, expressed as the average photon occupation of the resonator \bar{n} . Neglecting the constant Lamb shift, the measured qubit frequency $\omega(\bar{n})$ is

$$\omega(\bar{n}) = \omega_q + 2\chi\bar{n}, \quad (3.8)$$

allowing for a determination of \bar{n} if χ is known.

It is worth noting that the approximate Hamiltonian (Equation 3.5) fulfills the requirements for a quantum non-demolition measurement (QND) [63]. A QND measurement is one in which the measured state of the qubit is preserved after measurement is complete. The qubit state is of course free to evolve in a deterministic manner after measurement, but is not randomly scrambled by the measurement itself. Mathematically, a measurement is QND only if both the interaction and qubit terms of the Hamiltonian commute with the measured observable. The dispersive Jaynes-Cummings flux qubit Hamiltonian can be rearranged yet again to yield

$$H = \hbar\omega_r \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right) - \frac{1}{2} \hbar \omega_q \hat{\sigma}'_z - \frac{1}{2} \hbar (2\chi \hat{a}^\dagger \hat{a} + \chi) \hat{\sigma}'_z. \quad (3.9)$$

These terms can be associated with counterparts from Equation 3.1. One sees that

$$[H_{\text{int}}, \hat{\sigma}'_z] = 0, \quad (3.10)$$

$$[H_{\text{qubit}}, \hat{\sigma}'_z] = 0, \quad (3.11)$$

indicating that the measurement is in principle QND.

A quantum non-demolition measurement does not imply that there is no measurement backaction. Backaction is in fact the mechanism by which a measurement projects an arbitrary superposition state onto one eigenstate, ‘collapsing’ the wavefunction. In the ideal

QND measurement, however, subsequent measurements will render the same result without state randomization. Because the Pauli matrices by definition do not commute, a QND measurement of $\hat{\sigma}'_z$ is incompatible with a QND measurement of the other Pauli matrices, such as $\hat{\sigma}'_x$. The observable $\hat{\sigma}'_x$ will in fact be maximally scrambled by a QND measurement of $\hat{\sigma}'_z$.

3.1.2 SQUID resonator

A different Hamiltonian is formed if the harmonic oscillator in the Jaynes-Cummings formalism is replaced by a nonlinear, SQUID resonator. To obtain a sense for this system, it can be shown that in the limit of small SQUID inductance the interaction term of the total Hamiltonian becomes

$$H_{\text{int}} = M_{\text{qs}} I_q I_0 [\cos(\delta_1 + \delta_2)] \sin\left(\frac{\pi \Phi_x}{\Phi_0}\right) \hat{\sigma}_z, \quad (3.12)$$

where M_{qs} is the mutual inductance between the qubit and SQUID [64]. This term assumes that the flux generated by the qubit is small, so that a linear approximation can be used. The interaction term contains the flux induced in the SQUID by the qubit, $M_{\text{qs}} I_q \hat{\sigma}_z$, as well as a Josephson-like term. It is worth noting that the interaction is flux-dependent, going to zero at $\Phi_x = \Phi_0$. At this point the SQUID is first-order insensitive to external flux. As the SQUID flux is increased toward $\Phi_x = \Phi_0/2$, the interaction energy grows, and the SQUID resonator becomes more sensitive to the qubit circulating states.

The interaction in the SQUID resonator Hamiltonian is dominated by the magnetization signal of the flux qubit in the SQUID for non-zero Φ_x . The flux amplitude coupled by the qubit into the SQUID does not depend on the detuning δ between ω_q and ω_r . Thus, the magnitude of the frequency shift of the SQUID resonator is not directly sensitive to δ . This allows for a strong signal (i.e. frequency shift) for a large range of potential qubit frequencies. The frequency shift in the resonator $\Delta\omega_r$ between the two circulating states is given by

$$\Delta\omega_r = 2I_q M_{\text{qs}} \frac{\partial\omega_r}{\partial\Phi}. \quad (3.13)$$

The sensitivity of the resonator frequency to flux in the SQUID loop, $\partial\omega_r/\partial\Phi$, is a function of Φ and is generally non-trivial to compute for non-zero β_L , requiring numerical methods.

Unlike the dispersive Jaynes-Cummings Hamiltonian, the interaction term (Equation 3.12) in the SQUID resonator/qubit Hamiltonian does not commute with the qubit Hamiltonian (Equation 2.19). Generally speaking this is because the interaction Hamiltonian depends only on σ_z , while the flux qubit Hamiltonian contains both σ_z and σ_x . The system only becomes suitable for QND measurement when $\epsilon \gg \Delta_q$, or equivalently $\cos(\theta) \rightarrow 1$. In this scenario, the qubit eigenstates are approximately pure circulating current states. One strategy for working around this difficulty is to do all qubit manipulations at the noise-protected degeneracy point, and then quickly shift the flux bias so that the measurement is made away from the degeneracy point where the magnetization signal is higher and the measurement is more closely QND.




Figure 3.2. Coupling of cQED readout cavity to external circuitry. The readout resonator is coupled via a coupling capacitance (orange) to a microwave transmission line with characteristic impedance $Z_0 = 50 \Omega$. The source is modeled as an ideal voltage source V_S with a source impedance $Z_S = 50 \Omega$.

3.1.3 Readout signal and noise

The readout resonator is coupled to the $50\text{-}\Omega$ transmission line input via a coupling capacitance as shown in Figure 3.2. In the absence of this coupling capacitor, the resonator looks directly out at the $50\text{-}\Omega$ environment, which dissipates power from the resonator. As the coupling capacitance is decreased, the quality factor Q of the resonator increases, decreasing the linewidth of the resonance and the energy decay rate from the oscillator. Because the circuit is made entirely from superconducting materials, the energy dissipated inside the resonator (dielectric loss for example) is minimal.

Single-tone microwave pulses are reflected from the readout resonator to detect the qubit dispersive shift. The amplitude and phase of this reflected signal are analyzed to determine how the resonance has shifted. Due to the dissipation-free, reflection geometry, the reflected amplitude provides little contrast between the qubit states. Most of the signal is included in the reflected phase, which ideally undergoes a full 360° phase shift through its resonance. Figure 3.3 shows the calculated phase response of a reflection-based readout resonator for the qubit in the ground (blue) or excited (red) states. The maximum phase contrast between the states is obtained when the resonator is driven at a frequency ω_d , the midpoint of the two shifted curves. The signal between the two states is optimized when the resonator linewidth $\kappa/2\pi$ is equal to 2χ , producing a full 180° phase shift. This phase shift is a function of both the dispersive shift and the Q of the resonance. An increasing Q sharpens the response of the oscillator on resonance, producing a larger phase difference between the two curves for a given amount of frequency shift. It is important to bear in mind that the increasing phase shift with higher Q comes at the cost of a longer resonator response time, decreasing the measurement speed. This is a problem for qubit measurements because qubit coherence can be lost while the measurement oscillator is ringing up. Moreover, for a fixed drive amplitude, a higher Q corresponds to more power stored in the oscillator, increasing the effective measurement strength and potential backaction while simultaneously not increasing the reflected power available to infer the qubit state. Thus, engineering the quality factor of




Figure 3.3. The dispersive qubit shift in a microwave resonator. The curves represent the phase response of a microwave resonator in the reflection geometry for the qubit in the ground (blue) and excited (red) state. The total shift between the ground and excited states is 2χ . The maximum phase contrast between the two curves is obtained for a driving frequency (ω_d) at the midpoint of the two curves.

the readout resonator is a balancing act.

The qubit signal between the two readout ‘pointer states,’ or the resonator states which contain information about the qubit state, can be visualized more easily on the IQ plane. To understand this notation, we write an arbitrary reflected readout signal as

$$A \sin(\omega t + \phi) = [A \cos(\phi)] \sin(\omega t) + [A \sin(\phi)] \cos(\omega t) \quad (3.14)$$

$$= I \sin(\omega t) + Q \cos(\omega t), \quad (3.15)$$

where I and Q are the in-phase and quadrature amplitudes, respectively. A readout pointer state is represented in the IQ plane by a vector. The length of the vector represents the signal amplitude, while the direction represents the phase. The separation between the readout pointer states, represented by the dotted line \mathbf{s} in Figure 3.4(a), can be increased by either increasing the measurement amplitude or by adjusting the phase shift to be closer to 180° . Mathematically, the signal \mathbf{s} in the IQ plane can be expressed using the law of cosines as

$$|\mathbf{s}| = 2r \sin\left(\frac{\phi}{2}\right), \quad (3.16)$$

where r is the amplitude of the two vectors, assumed in this case to be the same, and ϕ is the phase difference.

Maximum separation of the readout pointer states is an issue due to experimental noise. In practice, repeated measurements of the readout resonator produce clusters of points in the




Figure 3.4. Readout pointer state vectors and noise in the IQ plane. (a) The quadrature amplitudes of the readout pointer state vectors in the IQ plane. The signal $|s|$ between the states is the magnitude of the vector connecting them. The measurement noise is represented by normally distributed clouds of points surrounding the vectors. (b) An increase in the measurement noise for the same signal magnitude is represented by larger clouds.

IQ plane centered about the pointer state vectors. These clusters are normally distributed, with the measurement noise typically defined as the standard deviation of this Gaussian distribution. An increase in measurement noise can be represented by the increasing size of the ‘cloud’ surrounding the pointer states as shown in parts (a) and (b) of Figure 3.4. Information about the state of the qubit is lost with increasing overlap between the two pointer state distributions. Separating the distributions by increasing the measurement amplitude is possible, but is usually accompanied by higher measurement backaction, which will be discussed later in this thesis.

Alternatively, the measurement noise can be decreased, which avoids the problems associated with higher measurement amplitudes. At dilution refrigerator temperatures, thermal fluctuations correspond to insignificant excitations in GHz frequency readout resonators. Especially at 5.8 GHz, the highest readout frequency we used, the resonator noise is almost entirely due to $(1/2)\hbar\omega_r$ zero-point fluctuations. The total system noise is in fact overwhelmingly dominated by the noise generated by amplifiers on the output line. These amplifiers are necessary to boost the very small readout signals up to the level where they can be processed by detection electronics at room temperature. Thus, the measurement sensitivity can be vastly increased by the decreasing the amplifier noise in the measurement chain.




Figure 3.5. Simplified model of a generic amplifier. The noise characteristics of a classical amplifier (blue shaded region) can be represented by voltage noise (V_N) and current noise (I_N) sources at the input. The input and output impedances are represented by Z_{in} and Z_{out} , respectively. The load connected to the input is modeled as an ideal voltage source V_S in series with a source impedance Z_S .

3.2 Signal amplification

An ideal, phase-preserving amplifier maps an electrical signal at its input into an identical copy with higher amplitude at the output. The term ‘phase-preserving’ means that both quadratures of an input signal are amplified equally. A power source must be integrated into the amplifier to provide the power for the amplified signal. A nonlinearity is required to transfer the supplied power into the output signal [65, 66]. A simplified model of a general amplifier is shown in Figure 3.5. The source is expressed by some voltage signal V_S in series with a source impedance Z_S . The input to the amplifier is modeled as an input impedance Z_{IN} together with a voltage noise source V_N and a current noise I_N . V_N is directly amplified by the ideal gain element G , producing intrinsic voltage fluctuations on the output, which is characterized by an output impedance Z_{OUT} . The current noise I_N divides between the input and source impedances, generating voltage fluctuations which are then fed back into the amplifier. As detailed in Reference [67], these two noise sources can be combined to express the total noise generated by the amplifier. Using the classical expression for the Nyquist noise of a resistor, it can be shown that in a narrow bandwidth,

$$2k_B T_N = \frac{1}{\cos \phi} \left[\frac{S_{\tilde{V}\tilde{V}}}{|Z_S|} + |Z_S| S_{\tilde{I}\tilde{I}} - 2 \operatorname{Re}(e^{-i\phi} S_{\tilde{V}\tilde{I}}) \right], \quad (3.17)$$

where T_N is an effective increase in the source temperature, $S_{\tilde{V}\tilde{V}}$ and $S_{\tilde{I}\tilde{I}}$ are the voltage and noise spectral densities, $S_{\tilde{V}\tilde{I}}$ is the cross-correlation spectral density, and $Z_S = |Z_S|e^{i\phi}$. This expression for T_N can be minimized with respect to $|Z_S|$, yielding

$$k_B T_N \geq \sqrt{S_{\tilde{V}\tilde{V}} S_{\tilde{I}\tilde{I}} - [\operatorname{Im} S_{\tilde{V}\tilde{I}}]^2} - \operatorname{Re} S_{\tilde{V}\tilde{I}}. \quad (3.18)$$

This completely classical argument shows that the noise generated by an amplifier depends on the source impedance, and is necessarily greater than some minimum. If this model is

treated quantum mechanically, it can in fact be proven that the right-hand side of Equation 3.18 is constrained such that

$$k_B T_N \geq \frac{1}{2} \hbar \omega. \quad (3.19)$$

This expression still relies on matching the proper source impedance Z_S as in the classical argument and also adds another condition. To reach the so-called quantum-limit, there must be no ‘wasted’ information in the detector. The rigorous quantum mechanical derivation of Equation 3.19 is treated in References [67, 68, 69]. The noise added by the amplifier can be thought of as originating from an extra mode associated with powering the amplifier. Without this mode, no noise would be added, but no amplification would be possible either.

There are two orthogonal degrees of freedom in a monochromatic electromagnetic signal. These can be expressed as the amplitude and phase of the signal, or equivalently, but in a different basis, the in-phase and quadrature amplitudes. Quantum mechanically, these quadratures can be treated as operators and can be shown to have a non-zero commutation relation, which implies an uncertainty relation between the quadratures. This uncertainty relation can be informally derived [70] by first writing the well-known uncertainty between energy and time

$$\Delta E \Delta t \geq \hbar / 2\pi. \quad (3.20)$$

Relating the uncertainty in energy to that in the number of photons n (amplitude) gives

$$\Delta E = \hbar \omega \Delta n, \quad (3.21)$$

while the time uncertainty Δt can be expressed in terms of phase ϕ as

$$\Delta t = \Delta \phi / \omega. \quad (3.22)$$

Substituting equations 3.21 and 3.22 into equation 3.20 yields

$$\Delta n \Delta \phi \geq 1/2. \quad (3.23)$$

This expression is known as the number - phase uncertainty. We now show that this uncertainty relation implies that the minimum noise at the output of an amplifier, referred to an energy at the input, is $\hbar \omega$. Of this full photon of energy, half a photon originates from zero-point fluctuations in the input circuitry connected to the amplifier, while the other half is added by the amplifier itself (Equation 3.19).

Let us first postulate the mythical existence of a noiseless, phase-preserving amplifier with power gain G . The output number of photons n_o is related to the input number n_i by

$$n_o = G n_i. \quad (3.24)$$

Additionally, this amplifier perfectly preserves the phase with

$$\phi_o = \phi_i, \quad (3.25)$$

up to a constant phase shift, where ϕ_i and ϕ_o are the input and output phases, respectively. If we also assume that we can detect the output signal with arbitrary precision up to the limit set by the uncertainty relation, this would mean that we could in principle measure

$$\Delta n_o \Delta \phi_o = 1/2. \quad (3.26)$$

Converting Equations 3.24 and 3.25 to uncertainties gives

$$\Delta n_o = G \Delta n_i, \quad (3.27)$$

$$\Delta \phi_o = \Delta \phi_i. \quad (3.28)$$

But, substituting Equations 3.27 and 3.28 into Equation 3.26 produces

$$\Delta n_i \Delta \phi_i = 1/(2G), \quad (3.29)$$

which clearly violates the original uncertainty relation given by Equation 3.23. This problem is reconciled by adding $(G - 1)\hbar\omega$ additional energy (noise) at the output of the amplifier. Dividing by G , this would correspond to $(G - 1)\hbar\omega/G$ of effective energy at the input of the amplifier, allowing Equation 3.23 to be satisfied because of the additional uncertainty in the photon number n_i . In the limit of high gain, this added energy at the input tends to $\hbar\omega$. Again recognizing that only half of this energy is due to the amplifier itself, we can define

$$T_N \geq \frac{\hbar\omega}{2k_B}, \quad (3.30)$$

where the quantity T_N is commonly referred to as the ‘noise temperature.’ It is not a physical temperature, but is rather a convenient way to compare the scale of the total noise generated by the amplifier at the output to the thermal background at the input. The noise temperature adds to the physical temperature to give the total effective noise at the input.

We have seen how amplification of the two orthogonal quadratures requires that noise be added by the amplifier. If, for some reason, we only care about optimizing the SNR in one quadrature, we can transfer all the added noise into the other quadrature. This is the basis behind phase-sensitive amplification [67], which only amplifies signals in a single quadrature, adding no additional noise. Signals which are not in the amplified quadrature of a phase-sensitive amplifier can be de-amplified. This is generally not a practical method of amplification for a phase incoherent input signal. If the signal is coherent, however, its phase can be rotated into the quadrature of noiseless amplification. Such a process will be further detailed in section 3.4.

Figure 3.6 shows a schematic of how the signal and noise are amplified in a phase-preserving amplifier, with additional noise added by the amplifier. Because of this added noise, the signal-to-noise ratio (SNR) is always degraded in a phase preserving amplification process. Nevertheless, this SNR reduction is a necessary price to pay to boost the signal up to a level much higher than the ambient noise level at room temperature where signal

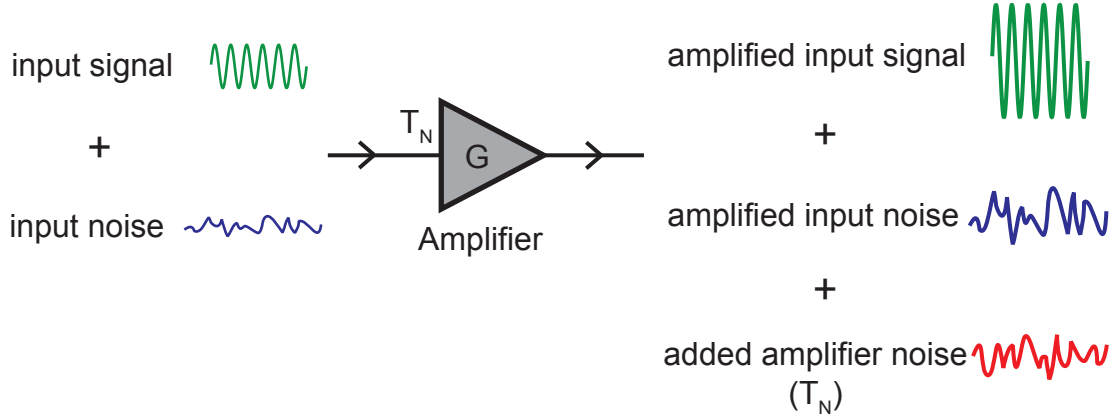


Figure 3.6. Signal and noise amplification. A phase-preserving amplifier produces copies of both the input signal and noise, scaled by the gain G . It also adds its own noise, characterized by the noise temperature T_N , to the output. This results in an overall reduction in the SNR.

processing will occur. To do this, multiple stages of amplification are required. In a properly designed amplification chain, the noise of the entire system is set by the first amplifier. Subsequent amplification stages also add noise, but the level of this noise should be insignificant because the overall signal and noise levels are already much higher. As an example, semiconductor High Electron Mobility Transistor (HEMT) amplifiers are currently the lowest noise commercially available amplifiers. At frequencies near 6 GHz, the best HEMT amplifiers have noise temperatures of $T_{\text{HEMT}} \sim 2 \text{ K} - 4 \text{ K}$. This noise level completely swamps the noise due to quantum fluctuations at 6 GHz, which is equivalent to 144 mK of noise power. The total output noise of any low-noise superconducting amplifier added before the HEMT should ideally be much larger than T_{HEMT} so that the total system noise is dominated by the amplifier with the best noise performance. For a superconducting amplifier with a power gain of $G = 100$, and a noise temperature quantum-limited at $T_N = 144 \text{ mK}$, the total output noise power per unit bandwidth P_N is

$$P_N = k_B G (T_N + T_{\text{inp}}), \quad (3.31)$$

where T_{inp} is the noise from the input circuitry. T_{inp} is taken to be 144 mK because at dilution refrigerator temperatures we are deep in the quantum regime. Using Equation 3.31, we find that the total output noise power is equal to $P_N/k_B = 28.8 \text{ K}$. This is much larger than T_{HEMT} , implying that the HEMT makes only a small reduction in the SNR. It can similarly be shown that the overall system noise temperature T_{sys} with these two stages of amplification is

$$T_{\text{sys}} = T_N + \frac{T_{\text{HEMT}}}{G}. \quad (3.32)$$

In this form, it is easy to see that the contribution of the HEMT to the overall system noise is its noise temperature divided by the gain of the superconducting amplifier. Thus, it is

desirable to achieve the highest possible gain in the first stage amplifier without sacrificing its noise performance. This example can be extended in a similar manner to further stages of amplification.

3.3 Microstrip SQUID amplifier

The Microstrip SQUID amplifier (MSA) is based on the traditional, resistively-shunted dc SQUID, but is modified for resonant operation at microwave frequencies [41, 71, 72]. Amplification is achieved by coupling an input coil strongly to the SQUID via a mutual inductance M_{CS} , given by

$$M_{CS} = k\sqrt{L_e L_{S,(MSA)}}, \quad (3.33)$$

where L_e is the coil inductance, $L_{S,(MSA)}$ is the SQUID inductance, and k is the geometric coupling factor which generally approaches unity. The SQUID is current biased into the voltage state and flux biased such that V_Φ , which is typically on the order of $50 \mu V/\Phi_0$ [73], is maximum. An input voltage V_i is applied across the input coil, which produces a current I_L related to the input voltage by $V_i = Z_c I_L$, where Z_c is the total impedance (generally inductive for superconducting input coils) of the input coil. Oscillating currents in the input coil couple magnetic flux into the SQUID loop, producing an oscillating voltage at the SQUID output at the input signal frequency. These output voltages are higher in amplitude than the input signal, generating an amplified copy of the input.

Applying the input signal between both ends of the input coil in a SQUID amplifier is a technique which works well for amplification up to MHz frequencies. As the operating frequency is increased further, however, the capacitive impedance between the input coil and the SQUID washer decreases. This capacitance progressively shorts out more of the input coil inductance with increasing frequency, decreasing the flux coupling to the SQUID and reducing amplification. Self-resonances, which are resonant combinations of inductive and capacitive elements in the device itself, can also impair ideal operation. In order to retain high gain at these higher frequencies, the capacitance to ground can be used to define a resonant mode between the input coil and the SQUID washer, which is grounded as seen in Figure 3.7(a). For an input signal coupled between ground and one end of the input coil, with the other end left open circuit, a $\lambda/2$ resonant cavity is formed. The quality factor of the resonance is set by an off-chip coupling capacitor, C_c . When excited on resonance, the current oscillations in the input coil are enhanced by a factor of Q , strongly increasing the flux coupled to the SQUID. This provides amplification with a Lorentzian gain profile.

The power gain G of the MSA can be expressed as [73]

$$G = \left(\frac{M_{CS} V_\Phi I_L}{V_i} \right)^2. \quad (3.34)$$

Each of the terms in the numerator of Equation 3.34 can be optimized to produce the maximum gain. The mutual inductance M_{CS} can be increased by either increasing the

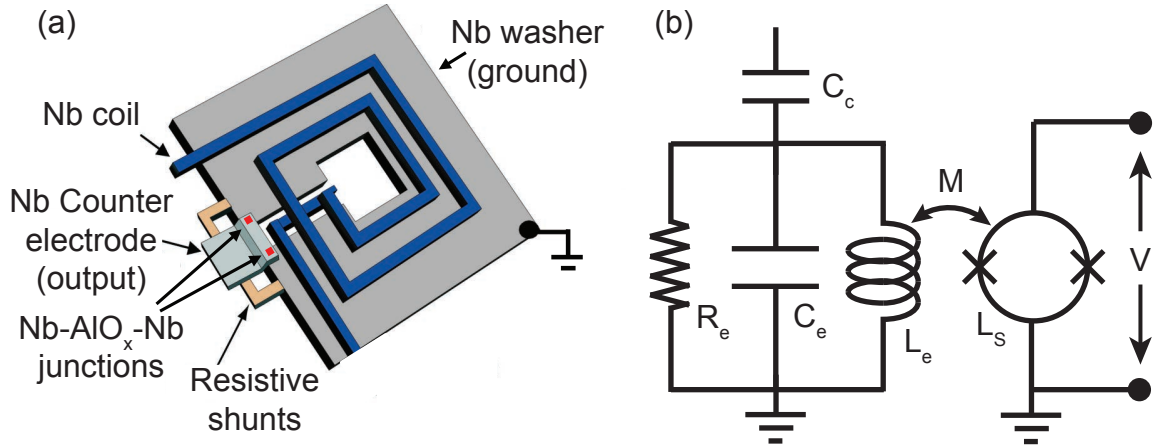


Figure 3.7. MSA overview. (a) Sketch of the MSA. An input coil forms a microstrip resonant mode with the dc SQUID washer. (b) Model circuit of the MSA. A damped resonator is formed by the parallel combination of an inductance L_e (inductance of the input coil), a capacitance C_e (capacitance of the input coil to ground), and a resistance R_e (dielectric loss at high frequencies). The resonator is grounded at one end and is coupled to the SQUID via a mutual inductance M . The flux-modulated voltage across the SQUID, which is operated in the voltage state, forms the output signal.

SQUID inductance $L_{S,(MSA)}$ or the coil inductance L_e . Increasing either, however, will reduce the resonant frequency of the amplifier. It can be shown that the total inductance L_T of the device is approximately [74]

$$L_T = L_e + n^2 L_{S,(MSA)}, \quad (3.35)$$

where a large n is the number of input coil turns on the SQUID washer. The large SQUID washer is generally undesirable when extending the operating frequency to higher values, as this large structure can lead to parasitic resonances. Even small capacitances between a large-inductance washer and other structures can resonate at frequencies near the amplifier's operating point.

The flux-to-voltage transfer function V_Φ should be made as large as possible without creating hysteresis in the SQUID current-voltage characteristics. Increasing the critical current of the SQUID or the shunt resistance, subject to the previously detailed constraints on β_L and β_C , will generally lead to higher SQUID output voltages, increasing V_Φ . For $\beta_L \approx 1$ and $\beta_C \leq 1$, the optimized flux-to-voltage transfer function is approximately [75]

$$V_\Phi \approx R_S / L_{S,(MSA)}, \quad (3.36)$$

where R_S is the value of the shunt resistance.

Increasing the magnitude of I_L is accomplished by optimizing the coupling—by adjusting C_e —of the MSA to the 50- Ω input. If this capacitance is too large, the MSA is said to be

overcoupled with a broadband, low gain response. If C_c is too small, the resonance is undercoupled. The quality factor is now much higher, but the internal dissipation in the oscillator limits the signal coupled to the SQUID. An optimized value of C_c leads to ‘critical coupling,’ which balances the internal and external dissipation such that I_L is maximum. This corresponds to matching the real part of the SQUID’s input impedance Z_{in} on resonance to 50Ω , while the imaginary part by definition is zero on resonance. The coupling capacitor is said to transform Z_{in} up to (or down to) 50Ω . The procedure for determining the correct value of C_c is detailed in Reference [73]. This process essentially involves measuring the scattering parameters S_{21} (voltage transmission coefficient) and S_{11} (voltage reflection coefficient) of the MSA without the coupling capacitor. These parameters are defined as

$$S_{21} = \frac{V_o}{V_i}, \quad (3.37)$$

$$S_{11} = \frac{V_{\text{refl}}}{V_i}, \quad (3.38)$$

where V_o is output voltage and V_{refl} is the reflected voltage from the input of the amplifier¹. The results are then fit to an equivalent circuit model as shown in Figure 3.7(b). Using the values of this model, an approximation for the value of C_c which critically couples the MSA, and thus maximizes I_L , is obtained.

To design an MSA, one first chooses a configuration for the SQUID washer, which defines its inductance $L_{S,(\text{MSA})}$. This geometric inductance constrains the junction critical currents (I_0) through the $\beta_L = 1$ condition. A value for β_C , which ideally is rather less than the hysteretic limit of about one, is subsequently chosen. For $\beta_C = 1$, we see that

$$R_S^2 = \frac{\Phi_0}{2\pi I_0 C_J}, \quad (3.39)$$

where C_J is the junction capacitance. As I_0 has already been constrained, the area of the junctions should be made as small as possible to keep their capacitance low and maximize R_S . This maximizes the gain of the amplifier as seen by Equations 3.34 and 3.36.

Table 3.1 shows the targeted parameters of our amplifiers. The SQUID inductance is provided by a washer design which has been successfully used for high gain and low noise in MSAs below 1 GHz. This inductance sets the junction critical currents as discussed above. The junction capacitance is made as small as reliably possible for our junction process later detailed in Section 4.2.3. Finally, R_S is targeted such that $\beta_C \approx 0.5$.

The MSA was first developed to study cold dark matter, namely in the search for the axion [48]. It was integrated into a receiver chain which substantially lowered the system noise temperature, allowing for a much faster scan rate of frequency space where the signatures of axions might be detected. At 0.62 GHz, an MSA was demonstrated to have a noise temperature of 1.6 times the quantum limit [76]. In order to achieve such performance, an exhaustive search of the current and flux bias spaces must be performed to find the

¹These voltages, and hence the scattering parameters, are generally represented as complex numbers.

Table 3.1. MSA design parameters.

Parameter	Symbol	Value
SQUID inductance	$L_{S,(MSA)}$	450 pH
Junction critical current	I_0	2.3 μ A
Junction capacitance	C_J	0.2 pF
Shunt resistance	R_S	20 Ω




Figure 3.8. Schematic representation of the paramp. The paramp is the parallel combination of a capacitor and a SQUID, which serves as a flux-tunable, nonlinear inductor. The input signal is applied differentially as indicated by V+ and V-.

optimal operating point [77]. Moreover, the shunt resistors, which generate Nyquist noise corresponding to a temperature above the bath temperature due to the creation of hot electrons when the MSA is current biased, must be kept as cool as possible. Large ‘cooling fins’ made of a gold-copper alloy were connected to the shunt resistors to increase the volume in which the hot electrons could cool via their interaction with phonons. In this work, we extend the operating point of the MSA into the GHz frequency range, and apply it to quantum information applications.

3.4 Josephson parametric amplifier

A SQUID resonator, with a layout identical to that discussed in Section 3.1.2, can also be used as a parametric amplifier [47, 78]. A schematic of this device is shown in Figure 3.8. The amplification mechanism can be understood by considering the dynamical equation of motion (Equation 1.13) of a Josephson junction in the RCSJ model. Varying the flux in the SQUID allows the nonlinear inductance and thus the operating frequency of the amplifier to be tuned. In essence, the flux-tuned SQUID behaves as a single junction with a controllable

critical current. Additionally, the large shunting capacitance of the SQUID resonator is mathematically equivalent to the junction self-capacitance. Taylor expanding $\sin(\delta)$ to its first nonlinear term and assuming a harmonic drive of the form $I = I_d \cos(\omega_d t)$, Equation 1.13 becomes

$$I_d \cos(\omega_d t) = I_0 \left(\delta - \frac{\delta^3}{6} \right) + \frac{\hbar}{2eR} \frac{\partial \delta}{\partial t} + C \frac{\hbar}{2e} \frac{\partial^2 \delta}{\partial t^2}. \quad (3.40)$$

This equation is the well-known Duffing equation, which can exhibit chaotic behavior and in general does not have an exact, closed-form solution. Defining the plasma frequency $\omega_p = \sqrt{2\pi I_0 / \Phi_0 C}$ and setting $\zeta = 1/2RC$, we re-write Equation 3.40 as

$$\omega_p^2 \frac{I_d}{I_0} \cos(\omega_d t) = \omega_p^2 \left(\delta - \frac{\delta^3}{6} \right) + 2\zeta \frac{\partial \delta}{\partial t} + \frac{\partial^2 \delta}{\partial t^2}. \quad (3.41)$$

We now postulate a solution of the form [79]

$$\delta(t) = \delta_0 \cos(\omega_d t + \theta) = \delta_x \cos(\omega_d t) + \delta_y \sin(\omega_d t), \quad (3.42)$$

where $\delta_0^2 = \delta_x^2 + \delta_y^2$ and $\tan(\theta) = \delta_y / \delta_x$. The parameter θ is recognized as a detuning-dependent phase offset. In the rotating wave approximation, where rapidly oscillating $2\omega_d$ and $3\omega_d$ terms are replaced by their means, we obtain the two coupled equations

$$\delta_y \left[\Omega - \frac{Q}{8} (\delta_x^2 + \delta_y^2) \right] - \delta_x = 0, \quad (3.43)$$

$$\delta_y + \delta_x \left[\Omega - \frac{Q}{8} (\delta_x^2 + \delta_y^2) \right] = Q \frac{I_d}{I_0}, \quad (3.44)$$

where the quality factor $Q = \omega_p RC$ and the dimensionless detuning of ω_d from ω_p is $\Omega = 2Q(1 - \omega_d / \omega_p)$.

The solutions to these equations, both the amplitude and phase of oscillations as a function of Ω , are plotted in Figure 3.9. For $\omega_d < \omega_p$, as the drive amplitude (I_d / I_0) increases the amplitude peak progressively bends over to lower frequencies. The corresponding phase offset θ also moves to lower frequencies and noticeably sharpens. As the detuning is increased beyond a critical value, given by $\Omega = \sqrt{3}$, the system actually admits three solutions. This behavior can be understood by the third order character of Equations 3.43 and 3.44. It turns out that only two of the solutions are stable in this region, which is therefore coined the ‘bistable’ region. The response of the resonance as a function of drive power and frequency is shown in Figure 3.10. The ‘paramp’ region occurs right before the oscillator becomes bistable, and is characterized by a very sharp phase response in both frequency and power.

The SQUID resonator can be operated as both a phase-sensitive and phase-preserving amplifier in the paramp regime. Because a coherent qubit readout pulse is suitable for phase-sensitive amplification, we utilize this operation mode with its superior noise performance

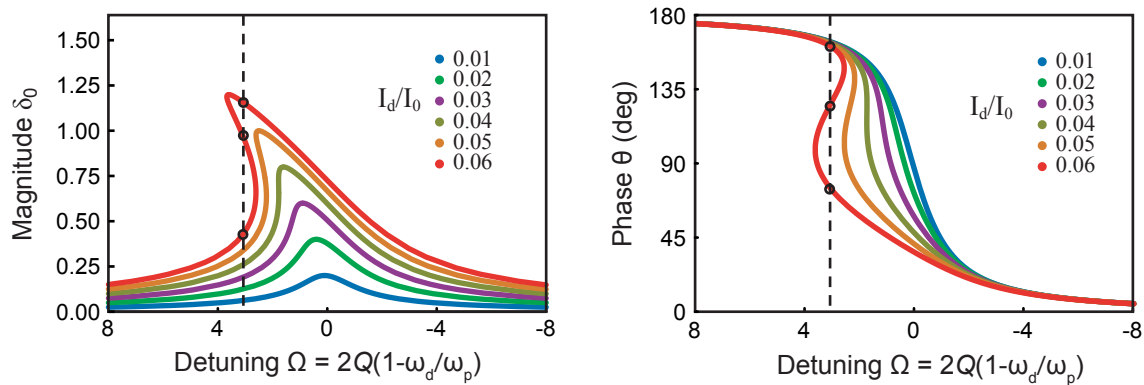


Figure 3.9. Response of the driven nonlinear oscillator. A Josephson element in parallel with a capacitance behaves as a nonlinear Duffing oscillator. As the amplitude of the drive current I_d is increased, the nonlinearity causes the resonance to bend over to lower frequencies. The response also sharpens until a critical point is reached at $\Omega = \sqrt{3}$, below which three solutions are possible. As only two of the solutions are stable, this is known as the bifurcation region [80].

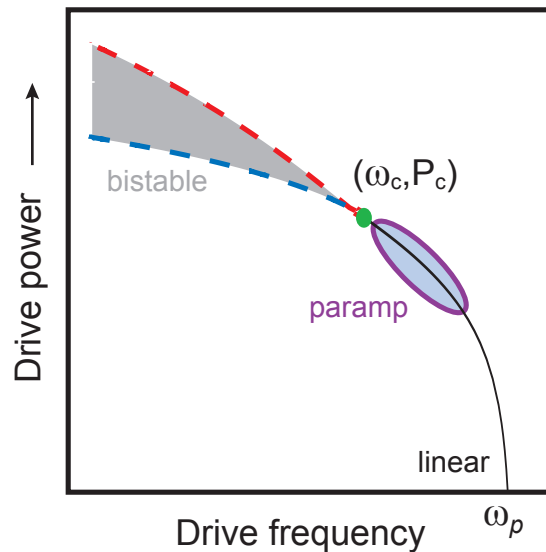


Figure 3.10. Resonance frequency response with increasing power in a nonlinear oscillator. With increasing drive power, the resonant frequency of a nonlinear oscillator bends over towards lower frequencies. The paramp regime is loosely defined as the nonlinear region before the onset of bistability at critical point, labeled by ω_c and P_c . This paramp region is characterized by a very sharp phase response for small changes in frequency or power. If the critical power is exceeded, the resonator becomes bistable and the two resonator states can be used in a digital qubit readout scheme (Josephson Bifurcation Amplifier).

in this work. The theoretical gain and bandwidth for the amplifier are treated in References [79, 80, 81]. Quoting the result, the approximate gain as a function of Ω is given by

$$G = 1 + \frac{1}{3 + 7\Omega^2 - 4\Omega\sqrt{3(1 + \Omega^2)}}. \quad (3.45)$$

This relationship shows the gain diverging as $\Omega \rightarrow \sqrt{3}$, the critical detuning for the onset of bistability. This makes sense in the context of the sharpening system phase response. In reality, experimental fluctuations and the neglected higher order terms in the gain approximation limit the response. The bandwidth BW of the amplifier, given by the full-width at half maximum, is given by

$$BW = \frac{\omega_p}{2\pi Q\sqrt{hG}}, \quad (3.46)$$

where

$$h = \frac{\Omega [3\Omega - 2\sqrt{3(1 + \Omega^2)}]}{1 + \Omega^2}. \quad (3.47)$$

The explicit dependence of BW on Q illustrates that a low Q maximizes the gain-bandwidth product, a quantity of merit for amplifiers. Near-quantum-limited noise performance, with sufficient gain and bandwidth for qubit readout, has been experimentally demonstrated in previous work [47, 78].

The mechanism for phase sensitive amplification is shown in Figure 3.11. The paramp is initially driven with a strong RF (radio frequency) pump at a frequency and power corresponding to the paramp regime. In a very narrow range of pump power, the reflected phase from the resonator undergoes a very large phase shift. When a small, in-phase signal is added to the pump tone, as shown in Figure 3.11(a), the total power is either increased or decreased slightly depending on the phase of the signal. This small change in total power maps to a very large change in the reflected phase, as shown in Figure 3.11(b). These phase changes ride on the very large pump amplitude, effectively amplifying the input signal, as seen in the IQ plane in Figure 3.11(c). This form of amplification theoretically adds no additional noise to the measurement because the noise is entirely in the orthogonal quadrature. This is shown in Figures 3.11 (d)-(f), which similarly diagram the amplifier response for input signals in quadrature with the pump. These quadrature input signals lead to very small changes in the reflected pump amplitude, producing small phase shifts and effectively leading to de-amplification.

The pump power to reflected output phase transfer function of Figure 3.11(b) is only linear for small range of pump powers. This range roughly defines the dynamic range of the paramp. If the input signal power is increased beyond this linear range, the output phase shift increases at a slower rate until it saturates to 180° . Paramp saturation is possible for powers typical of qubit readout. In this regime, the paramp can still provide gain, but the gain is reduced without a proportional decrease in the output noise. Thus, amplifier saturation leads to an increase in the amplifier's noise temperature. Fortunately, operation in the saturated regime is still acceptable for qubit readout because the two readout pointer

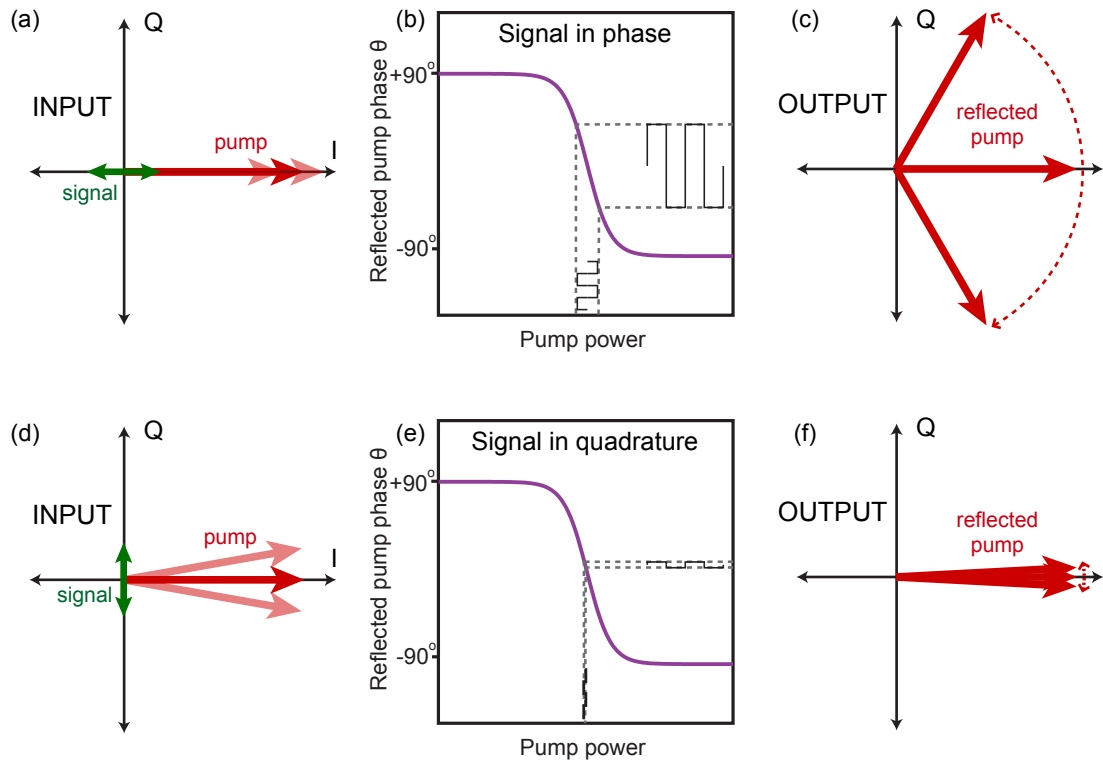


Figure 3.11. Principles of phase-sensitive amplification in the paramp. (a) In the presence of a small input signal added in phase with the pump, the pump power is maximally modulated as seen in the IQ representation. (b) This maximum change in pump amplitude maps to a large change in the output phase of the reflected pump. (c) Total range of the output signal in the IQ plane, along with the original pump. (d) When the input signal is in quadrature with the pump, the pump power is minimally modulated. (e) This minimum change in pump amplitude produces only small phase changes in the reflected pump. (f) This corresponds to small signal differences in the IQ plane, which are in fact de-amplified copies of the original signal. Reproduced from [79] with permission.

states are effectively digital. Even if the paramp is saturated, we need only the capability to faithfully discriminate between the two ‘railed’ output phases.

Finally, the region of bistability can be used as a digital, latching detector [82, 83] for qubit readout [84, 85]. This is analogous to the readout in which a SQUID probabilistically switches into the voltage state with a qubit-state-dependent probability. In the so-called JBA (Josephson Bifurcation Amplifier) readout, the two bistable oscillator states in the nonlinear resonator are instead used as the readout pointer states. The JBA exhibits the advantages of a switching, latching readout without the associated dissipation and quasiparticle conduction associated with switching the SQUID to the voltage state. Rather, the JBA remains in the superconducting state for the entire duration of the measurement. Although JBA readout

was not the focus of this work, it was successfully used for characterization and diagnostic measurements.

Chapter 4

Design and fabrication of devices

Reliably fabricating the nanoscale structures and precisely engineered Josephson junctions featured in our superconducting qubits and amplifiers can be challenging. Fast-paced development in the field necessitates continuing, and often dramatic, changes to chip layouts. This chapter details the methods used to design and fabricate the devices studied in this thesis.

4.1 Flux qubits

While the flux qubit holds certain advantages over other superconducting qubit species, it is notoriously difficult to fabricate to within desired specifications. A high level of attention to detail, intimate mastery of each fabrication step, and a willingness to repeat a process multiple times are currently required for success.

A liftoff process was used for all qubit fabrication steps. This entails first coating the entire sample with a polymer resist, which is evenly applied by spinning the sample (Figure 4.1) at speeds on the order of a few thousand revolutions per minute (RPM). The resist is then selectively damaged by exposure to either electrons or ultraviolet photons in regions where features are desired. The resist is washed away from these areas and preserved everywhere else during treatment with a developing agent. With the mask now defined, the entire surface is coated with a thin film (metal or dielectric) that will form the new features. Finally, the resist mask is peeled or ‘lifted off’ the sample, leaving the patterned features and removing everything else.

In addition to flux qubit Josephson junctions, one of our readout schemes used a SQUID-based resonator. Although the current discussion is specific to qubit junctions, the same methods and issues apply to the SQUID junctions as well.



Figure 4.1. Headway PWM32 spinner. Resist is spun onto the substrate at speeds of 3000 - 5000 RPM using the Headway PWM32 spinner. The spinner resides in an overpressured hood which prevents particle contamination. The resist is baked onto the sample using the temperature-regulated hot plate next to the spinner. The yellow lighting in the room prevents unwanted exposure of photoresist due to ultraviolet light.

4.1.1 Qubit design

Although the flux qubit energy spectrum is uniquely defined by Δ_q and I_q , there are many considerations which go into engineering a successful flux qubit. To avoid thermal population of the excited state, it is necessary that $\Delta_q \gg k_B T$. A dilution refrigerator, which will be discussed in Chapter 5, is used to anchor the qubit to temperatures near 25 mK. This allows the qubit frequency Δ_q/h to be on the order of a few GHz, in principle sufficiently high to avoid spurious qubit excitations. The upper limit of practical qubit frequencies is set by the increasing difficulty of high frequency microwave engineering and by preserving sufficient coupling to the readout circuitry in some readout schemes. In order to achieve these GHz energy splittings, E_C/h must be on the order of one GHz. Previous group research has measured a specific capacitance of $\sim 100 \text{ fF}/\mu\text{m}^2$ in our Josephson junction process [57], constraining junction sizes to the order of $0.1 \mu\text{m}^2$.

The circulating current I_q sets the curvature of the energy spectrum. As such, the amplitude of the readout signal increases with increasing I_q as noted in Section 2.5. The price to pay for this increasing signal is a steeper energy vs flux spectrum, making the qubit more susceptible to flux noise. Choosing the value of I_q represents a compromise between these two considerations. Increasing I_q is accomplished by increasing the junction critical currents, which also increases E_J/E_C for a fixed junction size. As this ratio decreases, the energy bands become flatter in flux and the eigenstates become increasingly charge-like in character. The dominant concern in this regime is the increasing ‘intra-cell’ doublet splitting. Aging of the junctions leads to a decrease in both the critical currents and E_J/E_C . In this work, qubits were fabricated with I_q ranging from about 130 - 210 nA.

Once the approximate areas of the junctions and their critical currents have been determined, the value of α_q is chosen to produce the desired Δ_q . The sensitivity of Δ_q to small changes in α_q allows for a considerable range of potential frequency space. It is not possible to decrease α_q arbitrarily, however, because as α_q approaches 0.5 for low β_q qubits, the magnitude of the circulating current plummets to zero [52]. This corresponds with the transition of the double-well potential to a single well, destroying the flux qubit potential. With increasing β_q , the geometric loop inductance effectively takes the place of Josephson inductance, allowing for values of α_q lower than 0.5. Ideally, it is best to keep α_q in the neighborhood of 0.6, comfortably away from these lower values. Although an increasing β_q also affects the energy level structure of the qubit, these changes are easily manageable as long as they are properly accounted for and β_q is kept well below 1. The highest β_q measured in this work was about 0.1.

4.1.2 Resist stack

All qubit samples studied in this thesis were fabricated on oxidized silicon substrates. As flux qubit Josephson junctions are required to be small to attain the desired qubit energy splitting, electron beam lithography and resists are used to pattern the structures. Our junction fabrication technique relies on the use of a bilayer resist stack, with the underlayer made of methyl methacrylate/methacrylic acid copolymer (MMA) El-13 and the top layer formed of ZEP-520A. This represents a departure from previous flux qubit recipes in the Clarke group, which utilized poly(methylmethacrylate) (PMMA) as the top layer [57].

Relatively speaking, ZEP (~ 250 nm thick) is harder, thinner, and less charge sensitive than MMA (~ 850 nm thick). It is a relatively new, high resolution resist that precisely defines electrode features in the mask. The underlayer MMA is more readily exposed and developed, allowing for removal of MMA underneath regions of intact ZEP. This so-called 'undercut' is vital for the formation of Josephson junctions with double angle shadow evaporation, and it is desirable to have the capability to independently engineer the amount of undercut in a particular mask design. This is made possible with the MMA/ZEP resist combination. Because MMA is developed by 1:3 methyl isobutyl ketone (MIBK):isopropyl alcohol (IPA), while ZEP uses n-amyl acetate, independent development of each layer is enabled. Each layer is insensitive to the developer of the other, allowing for 'orthogonal development.' PMMA/MMA stacks, on the other hand are both developed with the MIBK/IPA mixture, thus coupling the development of the high resolution top layer with the amount of undercut in the bottom layer. Optimizing the development time for the highest resolution in the top layer constrains the undercut to an amount which may be either insufficient or excessive. While it has been proven possible to make flux qubits using this resist combination, orthogonal development allows for much more flexibility and tolerance in the fabrication process.

Although the advantages of orthogonal development are clear, there is another complication. When developed at room temperature, the charge doses required to expose ZEP optimally are relatively close to those of MMA, thus making it difficult to expose undercut




Figure 4.2. Equipment for electron beam lithography. (a) E-beam resist is exposed by a focused electron beam generated by a modified SEM. (b) The sample is placed on a stub which is specially configured for e-beam lithography.

regions selectively while preserving the top layer. This is remedied by lowering the temperature of the ZEP developer with a surrounding, 0°C ice water bath. Development at lower temperature [86, 87] raises the optimal ZEP dose, further separating it from the lower MMA dose. This effect is possible because the development process is thermally activated, so that the rate of development decreases with decreasing temperature. We compensate for this decreasing rate with an increased charge dose. Additionally, selectivity between exposed and unexposed ZEP regions is increased with the developer chilled to 0°C , allowing for higher contrast and sharper edges in the top layer [88, 89].

An optimized resist and development recipe allows for the formation of ‘airbridges,’ which are used to define the Josephson junctions (Figure 4.3). If two traces are patterned into the resist with only a small gap between them, the MMA under the gap can be undercut so that only a small freestanding section of ZEP separates the two electrodes. The resulting bridge is anchored at both ends but is otherwise unsupported. Sagging, cracking, and excessive tension in the bridge ultimately set limits on the maximum bridge size, and thus the size of the Josephson junction. The qubit and SQUID junctions fabricated in this work (maximum bridge size $\sim 200 \times 800 \text{ nm}^2$) were well within the process capabilities.

4.1.3 Electron beam lithography

Electron beam (e-beam) resists are exposed using very finely focused electron beams controllably scanned across a surface to produce a pattern. Systems dedicated for this purpose are commercially available, with guaranteed linewidths below 10 nm. A more economical alternative is to use a modified scanning electron microscope (SEM). Our FEI Nova NanoSEM [Figure 4.2(a)] is reliably able to write the smallest (100 nm) flux qubit features produced




Figure 4.3. Features patterned into electron beam resist. (a) A top-down SEM image of a Josephson junction test structure shows patterns in a bilayer e-beam resist stack. Wiring leads connect to small ‘airbridges,’ which are integral to the formation of Josephson junctions using double-angle shadow evaporation. (b) A magnified, angled image of an airbridge clearly shows the two layers of resist, including undercut of the bottom layer. The airbridge is suspended above a region of undercut, anchored at both sides by the bulk top layer. The small slits next to the airbridge were designed to relieve strain, which in this particular mask resulted from contraction and expansion of large areas of undeveloped resist.

in this work, while other users of the instrument have achieved features of about 30 nm.

The sample is first clipped onto a specially designed holder [Figure 4.2(b)] which includes both a Faraday cup (for precise measurements of the beam current to correct for current drifts) and a viewing standard of gold colloids (for setting the microscope focus, stigmation, and lens alignment). Additional 40- and 100-nm gold colloids placed on the sample surface allow the microscope to probe the focal length at multiple locations on the surface and correct in software for sample tilt and/or working distance misalignment. With these calibrations complete, the sample is ready to be patterned. The accelerating voltage of the electrons is set to the maximum 30 kV allowed by our system. This allows for crisper edges in the patterned resist, due mainly to reduced backscatter of electrons from the substrate and resist. Backscatter from the substrate aids in the formation of undercut, but reduction of this effect can be compensated by a longer MMA development time.

The microscope is remotely controlled by an adjacent computer running the Nanometer Pattern Generation System (NPGS) software. All of the e-beam write parameters, including the patterns, doses, exposure step distances, and magnification settings, are programmed into the software. The general philosophy is first to pattern the smallest, most sensitive features at the highest possible magnification and lowest possible beam current. The magnification is then decreased as the beam current is increased for the larger, more coarse objects. This sequence allows the Josephson junctions to be patterned first while the low beam currents are well calibrated. Because the SEM has both coarse and fine sets of scanning coils used




Figure 4.4. Aluminum evaporators. (a) The old thermal evaporator was used to fabricate the earliest flux qubit samples by depositing and controllably oxidizing aluminum thin films. This system produced many viable qubits, but was more cumbersome to operate and yielded samples which aged quickly. (b) The custom-built e-beam evaporation system is more user-friendly and produces devices more resistant to junction aging. This system features a motorized, rotating stage and a load-locked sample chamber. The system is controlled either manually or by a computer interface. Pump/purge cycles of the manifold which holds the gas mixture used for oxidation as well as the actual oxidation process are currently capable of full automation.

to deflect and scan the electron beam, it is critically important that the most sensitive, high resolution features be written at a magnification corresponding to the fine coils. The crossover between these two regimes occurs at a magnification close to 1350x, and can be detected by an audible click as the magnification is swept through its crossover value.

4.1.4 Aluminum deposition

Once the sample has been patterned and developed, the superconducting aluminum layer is deposited through a process called double-angle shadow evaporation [90]. Evaporation is chosen because it is directional, that is, the evaporated material travels in straight line paths from the source to the sample. This is in contrast to other diffuse metalization processes such as sputtering.

Evaporation is carried out in a high vacuum chamber, and was performed by two different methods in this work. At the beginning of the project, an old NRC thermal evaporator [Figure 4.4(a)] was the only available option. Thermal evaporation involves loading the aluminum into a small tungsten (which has a high melting temperature) boat. When a

current is passed through the boat, it heats the aluminum until it evaporates and then condenses on the targeted sample. Each time an evaporation is performed, fresh boats and material are loaded, making the process time consuming. The main chamber also must be completely vented and exposed to air after each deposition, increasing pump down times and introducing possible contamination. Other disadvantages of this system include its cumbersome oxidation manifold system and sloppiness in the manual ‘beer can’-based method used to change the angle of the sample relative to the source. Later in the project, a custom e-beam evaporation system [Figure 4.4(b)], designed and built mostly by Daniel Slichter, was used for flux qubit fabrication. A thermionically-emitted electron beam is accelerated through potential of about 7 kV and guided to an aluminum source via bending magnets, heating the material to the point of evaporation. This e-beam system addresses many of the shortcomings of the old system by isolating the main chamber using a load lock and mechanizing the sample stage rotation.

To perform an evaporation, the qubit sample is loaded into the chamber, which is typically pumped overnight using cryo and turbo pumps until the pressure $< 10^{-7}$ torr. The sample stage is set to the initial evaporation angle, and the first Al film ($\sim 35 - 50$ nm) is deposited at a rate of ~ 0.5 nm/sec. An oxidation manifold containing a precise mixture of 95% Ar / 5% O₂ gas is then released into the chamber for a predetermined time. The degree to which the aluminum is oxidized, quantified by the term ‘exposure,’ is a function of both the partial pressure of O₂ gas and time [91]. The thickness of the oxide barrier which forms the Josephson junction increases with increasing exposure, reducing the critical current density J_c . Critical current densities up to 1000 A/cm², were found to be reliable for qubit junctions. We successfully fabricated junctions with higher J_c (up to ~ 1600 A/cm²), but the failure rate was much higher, probably due to shorts penetrating the very thin oxide barriers. The oxidation parameters required to obtain the targeted J_c are empirically determined through trial and error. After oxidation, the second evaporation ($\sim 50 - 85$ nm) is performed at a second angle, encapsulating the aluminum oxide to form the Josephson junction (Figure 4.5). This process is graphically depicted in Figure 4.6. Unintended, spurious junctions are inevitably formed in this process, but as long as their areas are much larger than the qubit junctions, their effect is negligible.

After the sample has been removed from the evaporation chamber, it is placed in a small beaker of 70° C acetone for about 30 minutes. This dissolves the underlying resist and makes it easier to peel off the aluminum layer. Light sonication will readily remove sections of resist remaining after liftoff.

4.1.5 Oxygen plasma

For the later flux qubit samples produced in this thesis, multiple oxygen plasma cleaning steps were incorporated into the fabrication process. These cleaning steps were performed in a chamber dedicated for this process. With 250 mTorr of O₂ gas in the chamber, about 10 W at 100 kHz is applied via parallel, horizontal plates. This excites vertical oscillations




Figure 4.5. SEM images of Josephson junctions. (a) The two aluminum layers which are deposited at different angles to form a Josephson junction are clearly seen in the angled SEM image. The red arrow points to the overlap region which forms the junction. (b) Three Josephson junctions in series, as seen in the top-view SEM image, are embedded in a superconducting loop to form a flux qubit.

in the oxygen plasma, bombarding the sample from above. This RF plasma treatment is a commonly practiced method of removing organic material, such as resist, from surfaces.

A 10-second oxygen plasma ‘ashing’ is first performed after the qubit features have been patterned and developed, but before the aluminum metal is deposited. This cleaning step is intended to remove resist residue on the substrate surfaces where the qubit features will be formed. The leftover resist is thought to contribute significantly to rapid aging of Josephson junctions [92]. This aging was observed consistently in our earliest flux qubits which did not include any plasma cleaning, with critical current densities decreasing by up to 50% in the most severe cases. This presents a tremendous challenge when engineering a qubit whose energy levels are exponentially sensitive to the critical current. The combination of the 10-second ‘pre-ash’, the new (i.e cleaner, lower base pressure) e-beam evaporation system, and pumping the sample down in the evaporator overnight using the cryo pump (which pumps water more effectively than the turbo pump) was found to dramatically reduce junction aging. Critical current densities of the flux qubits made using this recipe typically did not decrease by more than 10 %, a manageable amount.

After metalization and liftoff, the completed qubit sample is again ashed, this time for a longer duration of two minutes. This process is again designed to remove any leftover resist and other surface contaminants. In addition to combating the effects of aging, this is believed also to enhance qubit coherences times. Leftover contaminants close to and even on top of qubit features can present a lossy, defect-laden environment to the qubit. SEM micrographs taken both before and after sample ‘post-ashing’ verified that the surface was visibly cleaner. While we cannot yet present rigorous quantitative evidence of the benefits of ashing qubit samples, its incorporation into the fabrication process coincided with a consistent increase in measured coherences times.

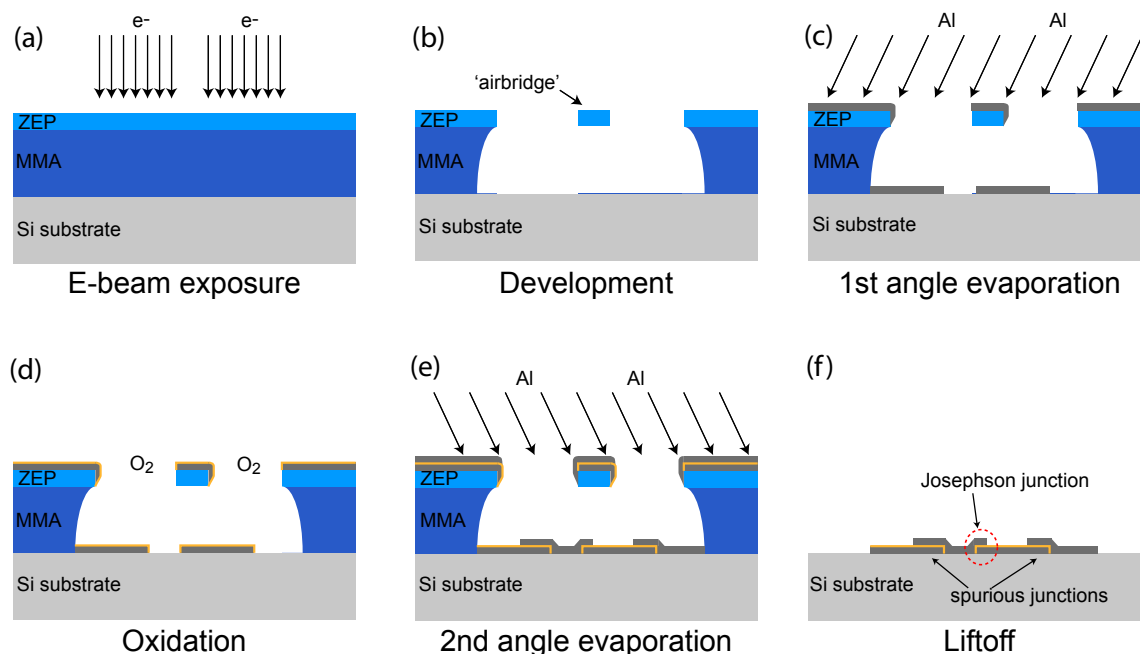


Figure 4.6. Double-angle shadow evaporation. (a) A ZEP/MMA resist bilayer is selectively exposed by a normally incident electron beam. (b) The sample is then submerged for a specific amount of time in two separate solutions of developer, each of which removes the exposed resist in only one of the layers. MMA can be developed out from underneath areas of intact ZEP, forming regions of undercut including ‘airbridges’ in the ZEP. (c) A first aluminum layer (gray) is deposited at an angle away from normal incidence. The line of sight directionality of evaporation helps form a gap in the deposited aluminum where the airbridge has blocked the material. (d) The entire aluminum surface is then controllably oxidized (yellow) in an Ar / O_2 gas mixture. (e) Josephson junctions are formed when a second layer of aluminum (gray) is evaporated from the opposite direction of the first aluminum layer. The second layer must be somewhat thicker than the first in order to ensure electrical continuity of the second layer as it climbs the ‘step’ of the first layer. In addition to the desired Josephson junction, spurious junctions due to the films overlapping in other places are also formed. If the area of these spurious junctions is made very large, their effect is insignificant. (f) The final step in the process involves lifting off the remaining resist, leaving only the patterned structures behind. Reproduced from [79] with permission.

4.1.6 Overlap capacitors

The microwave resonators used to read out the state of our flux qubits consist of quasi-lumped inductors and capacitors. Two series, on-chip overlap capacitors were used to achieve sufficient capacitance for the lower frequency resonators designed in the 1 - 2 GHz range. These capacitors are connected via a common, niobium plane, which is colloquially referred to as a ‘ground plane’ even though it is electrically floating and not physically connected to ground. In the presence of a differential excitation, the symmetry in the series capacitor structure defines the ‘ground plane’ as a virtual ground. The ground plane resides beneath a layer of silicon nitride which covers the entire chip. The capacitor top plates are formed in the same aluminum deposition as the Josephson junctions.

This ground plane is patterned at the 2” wafer level using e-beam lithography. The niobium is DC magnetron sputtered in a system which also contains an aluminum sputter gun and an argon ion mill. Sputtering occurs when ionized argon is accelerated through a potential at a target of niobium. The argon ions transfer kinetic energy to the target, knocking niobium atoms off its surface. Unlike evaporation, the mean free path of the sputtered niobium atoms is short, leading to diffuse spreading of the ejected atoms. The targeted thickness of these ground planes is ~ 250 nm.

The insulating SiN_x layers used in our experiments were fabricated both in the Marvell Nanolab and by Dr. Luigi Frunzio at Yale University. Our films were deposited (~ 180 nm) using Plasma Enhanced Chemical Vapor Deposition (PECVD) in the ‘oxford2’ system in the Nanolab. The dielectric constant of this material was experimentally determined from resonator measurements to be about 6.75. We eventually moved away from this overlap-style capacitor towards an interdigitated finger design because of the lossy nature of the SiN_x . Measurements of the low power loss tangent $\tan \delta \approx 10^{-3}$ confirmed these suspicions [79]. We attribute the loss in the dielectric to two-level defect states, which couple to the qubit and act as a bath of oscillators to which information can be irreversibly lost. This leads to decoherence in the qubit. Reducing the loss in these insulating layers is the subject of ongoing research in the field.

4.1.7 Device screening

Because of the many man-hours associated with preparing and cooling down the dilution refrigerator for qubit measurements, it is important to characterize and screen flux qubit samples as thoroughly as possible at room temperature. The first priority in a brand new flux qubit design is to measure the areas of the three junctions produced in the process using the SEM. This approximately determines the expected E_C . In a reliable process, these junction areas should not change by more than about 5% from run to run. Due to anecdotal claims that imaging Josephson junctions in the SEM alters their properties, promising samples are not imaged in the SEM before they are measured.

With no direct electrical access to the flux qubit junctions, it is impossible to check their critical currents directly to dial in E_J and α_q . As a substitute, we create a nearly identical




Figure 4.7. Probe station and desiccator. (a) A probe station is used to measure the normal state resistance of Josephson junctions to infer their critical currents at dilution refrigerator temperatures. (b) To minimize the effects of junction aging through exposure to water and oxygen, flux qubit samples are stored in a vacuum desiccator.

pattern of junctions adjacent to the real device. Patterned leads connect to these ‘witness junctions’ so that their electrical properties can be directly probed. This characterization is possible at room temperature by exploiting the Ambegaokar-Baratoff relation [93]

$$I_0 = \frac{\pi\Delta(T)}{2eR_n} \tanh \left[\frac{\Delta(T)}{2k_B T} \right], \quad (4.1)$$

where $\Delta(T)$ is the temperature-dependent superconducting gap and R_n is the normal-state resistance of the junction. At dilution refrigerator temperatures, we can approximate $\Delta(T) \rightarrow \Delta(0)$, which is 170 meV for aluminum. The tanh term is approximately unity at low temperatures. With these simplifications, we find

$$I_0 = \frac{(0.17 \text{ V})\pi}{2R_n}. \quad (4.2)$$

With this simple relationship, room temperature resistance measurements of the witness junctions are used to estimate the junction critical currents. These two-wire resistance measurements were performed at a probe station using lock-in excitation and detection [Figure 4.7(a)]. Some care must be taken to calibrate out the lead resistance to fully isolate the normal-state resistance of the junction. Although small variations between the qubit and witness junctions are unavoidable, this method was found to give reasonable agreement between predicted and measured qubit properties. This notwithstanding, our level of fabrication precision coupled with the extreme sensitivity of the qubit parameters still makes a ‘shotgun’ trial and error approach necessary to some degree. After the sample has been characterized at room temperature, it is stored in a vacuum desiccator [Figure 4.7(b)] to minimize further exposure of the Josephson junctions to ambient oxidizing agents such as oxygen and water vapor.




Figure 4.8. False color optical micrograph of the SQUID resonator device. The three-junction flux qubit (green) is inductively coupled to a dc SQUID (blue) shunted by two parallel-plate capacitors in series, forming a tunable nonlinear resonator for qubit readout. The capacitors consist of aluminum films (red) with a niobium underlayer (yellow), separated by a SiN_x dielectric. Symmetric on-chip flux bias lines (purple) allow independent tuning of the qubit and SQUID, while a shorted CPW structure (gray) allows microwave excitation and fast flux tuning of the qubit.

4.1.8 Device layouts and parameters

SQUID resonator readout

The SQUID resonator readout is modeled after previous generation devices fabricated in the Clarke group [5]. This design was optimized for controllable coupling of two flux qubits via a SQUID mediated interaction. We took this design, which featured a switching SQUID readout, and converted it to utilize a dispersive readout. Many of the parameters were left the same or modified only slightly, mostly to incorporate the on-chip resonator capacitors.

Figure 4.8 shows a false color optical image of the qubit, SQUID resonator, flux bias lines, and coplanar waveguide (CPW) qubit control line. A full list of design parameters for this device is presented in Table 4.1. All of the calculated inductances in the design were modeled using a modified (for use with superconductors) version of ‘FastHenry’ [94], a freely distributed program which solves for the magnetic response of a user-defined superconducting mesh. In similar samples fabricated before this work, these simulated values agreed with experiment to within 2 % [5].

The device contains two on-chip flux bias lines so that in principle the flux in both the qubit and SQUID can be independently controlled. As seen in Table 4.1, the two flux bias lines are equally coupled to the SQUID, while there is a substantial asymmetry in the coupling with respect to the qubit. The two degrees of freedom, however, allow us to orthogonalize the system such that different combinations of currents in each bias line can be used to adjust only the flux in either the qubit or the SQUID. Ideally, we would like the

Table 4.1. SQUID resonator readout approximate design parameters (F063010c).

Parameter	Symbol	Value
SQUID inductance	$L_{S,(res)}$	423 pH
SQUID loop dimensions	-	$30 \times 200 \mu\text{m}^2$
SQUID junction critical currents	I_{0S}	$0.85 \mu\text{A}$
SQUID junction area	A_S	$0.092 \mu\text{m}^2$
Resonator capacitance	C_r	28 pF
Predicted resonant frequency	f_r	1.75 GHz
Large qubit junction areas	A_{01}, A_{02}	$0.049 \mu\text{m}^2$
Small qubit junction area	A_{03}	$0.029 \mu\text{m}^2$
Large qubit junction critical currents	I_{01}, I_{02}	425 nA
Small qubit junction critical current	I_{03}	260 nA
Ratio of junction critical currents	α_q	0.61
Qubit loop inductance	L_q	182 pH
Qubit loop dimensions	-	$22 \times 83 \mu\text{m}^2$
Qubit screening parameter	β_q	0.07
Approximate critical current density	J_C	875 A/cm^2
Ratio of Josephson to capacitive energy (large junction)	E_J/E_C	50
Qubit-SQUID mutual inductance	M_{QS}	57 pH
Bias lines-SQUID mutual inductance	M_{S1}, M_{S2}	5.3 pH
Bias line 1-qubit mutual inductance	M_{Q1}	2.77 pH
Bias line 2-qubit mutual inductance	M_{Q2}	0.45 pH
CPW-qubit mutual inductance	$M_{Q,CPW}$	100 fH
CPW-SQUID mutual inductance	$M_{S,CPW}$	$< 2 \text{ fH}$

coupling of the bias lines to the qubit to be large enough so that at least a full flux quantum is accessible in the qubit. This is constrained by the measured critical currents of the bias lines ($\sim 1 \text{ mA}$). Significant heating of the dilution refrigerator is observed if the critical current is exceeded. If the coupling is too strong, however, fluctuations in the bias lines can suppress the coherence time of the qubit. The couplings in these devices have been made to be somewhat weaker than in previous designs, where the control lines were calculated not to limit coherence times to less than $\sim 10 - 100 \mu\text{s}$ [95, 96]. These coherence times are much larger than the measured values, suggesting that the bias and control lines are sufficiently decoupled.

The qubit excitation line is a coplanar waveguide (CPW) terminated with a small inductance. The excitation is coupled between the center line and the symmetric, grounded outer lines. The symmetry of this structure and its position relative to the SQUID are designed to prevent it from coupling flux into the SQUID loop. The offset position of the qubit loop, however, allows a weak coupling between the CPW and the qubit. The high bandwidth of




Figure 4.9. False color SEM micrograph of the circuit QED device. The quasi-lumped-element resonator is formed by the parallel combination of a meander line inductor (blue) and an interdigitated finger capacitor (red). The flux qubit (green), shown in the magnified inset, is inductively coupled to the resonator via a shared 150-nm wide constriction in the meander line. An additional pair of interdigitated capacitors (orange) couples the resonator to the 50- Ω environment. The resonator is surrounded by a perforated ground plane (purple), which helps define a co-planar waveguide input.

the CPW also enables static currents in the CPW to be changed very quickly. We use this feature to apply ‘fast’ flux shifts to the qubit loop between state manipulation and readout.

Circuit QED readout

A false color SEM image the flux qubit and circuit QED readout is shown in Figure 4.9. The flux qubit is coupled to the quasi-lumped resonator, formed by the parallel combination of an interdigitated finger capacitor and a meander line inductor, via a shared inductance. This shared section of the resonator is located where the linewidth of the meander line is constricted to 150 nm. The resonator is coupled to the 50- Ω feed line at one end, and grounded at the other, by separate interdigitated finger capacitors. A perforated ground plane surrounds the entire resonator and defines a 50- Ω , CPW input with the feed line.

A table of the design parameters for the ‘phase-coupled’ flux qubit in the cQED readout architecture is shown in Table 4.2. The resonator is modeled using AWR’s (Advancing the Wireless Revolution) proprietary ‘Microwave Office’ software. The model gives very good agreement between the predicted and measured resonant frequency, which is targeted near 6 GHz, the center of the 4-8 GHz band of the microwave circulators (Section 5.4.2). The effect of the coupling capacitors is not captured as well in the model, but a reasonable estimate can be determined. We target a resonator linewidth of ~ 10 MHz to roughly match the bandwidth of the paramp. This corresponds to a quality factor of $Q \sim 600$, which we

Table 4.2. Circuit QED readout approximate design parameters (F032811b).

Parameter	Symbol	Value
Resonator inductance	L_r	1.3 nH
Resonator capacitance	C_r	569 fF
Predicted resonant frequency	f_r	5.85 GHz
Predicted resonator linewidth	$\kappa/2\pi$	~ 10 MHz
Large qubit junction areas	A_{01}, A_{02}	$0.095 \mu\text{m}^2$
Small qubit junction area	A_{03}	$0.059 \mu\text{m}^2$
Large qubit junction critical currents	I_{01}, I_{02}	460 nA
Small qubit junction critical current	I_{03}	270 nA
Ratio of junction critical currents	α_q	0.59
Qubit loop inductance	L_q	7.3 pH
Qubit loop dimensions	-	$3.8 \times 4.0 \mu\text{m}^2$
Qubit screening parameter	β_q	0.003
Approximate critical current density	J_C	475 A/cm^2
Ratio of Josephson to capacitive energy (large junction)	E_J/E_C	100
Qubit-resonator mutual inductance (magnetic)	M_{kin}	3.2 pH
Qubit-resonator mutual inductance (kinetic)	M_{mag}	3.1 pH (measured)
Qubit-resonator mutual inductance (total)	M_T	6.3 pH (measured)

determined to be a good balance between measurement speed and readout phase contrast.

The qubit-resonator coupling is engineered assuming qubit parameters of about $\Delta_q \approx 7$ GHz and $I_q \approx 200$ nA. The total mutual inductance M_T between the resonator and qubit is a combination of magnetic (M_{mag}) and kinetic (M_{kin}) contributions [59]. The magnetic term is associated with the magnitude of the magnetic flux coupled into the qubit loop from currents in the resonator. The kinetic term, on the other hand, arises not from magnetic flux but rather from the inertial effects of charge carriers in an oscillating electric field. The kinetic inductance can become significant in superconductors because the characteristic collision time between Cooper pairs is very large. The kinetic inductance per unit length l_K in the Drude model for a superconductor is given by [97]

$$l_K = \frac{m}{2n_s e^2 A}, \quad (4.3)$$

where n_s is the density of Cooper pairs, m is the electron mass, and A is the cross-sectional area of the superconductor. Thus to increase l_K and hence the coupling between the resonator and qubit, A should be made small. We accomplish this by reducing the linewidth of the $2.6 \mu\text{m}$ -long shared section to a width of 150 nm, which is robust to run-to-run fabrication failure. For our 150-nm wide, 85-nm thick aluminum constrictions, we deduced a kinetic inductance per unit length of $l_K = 1.2 \text{ pH}/\mu\text{m}$, which is close to that obtained in other work

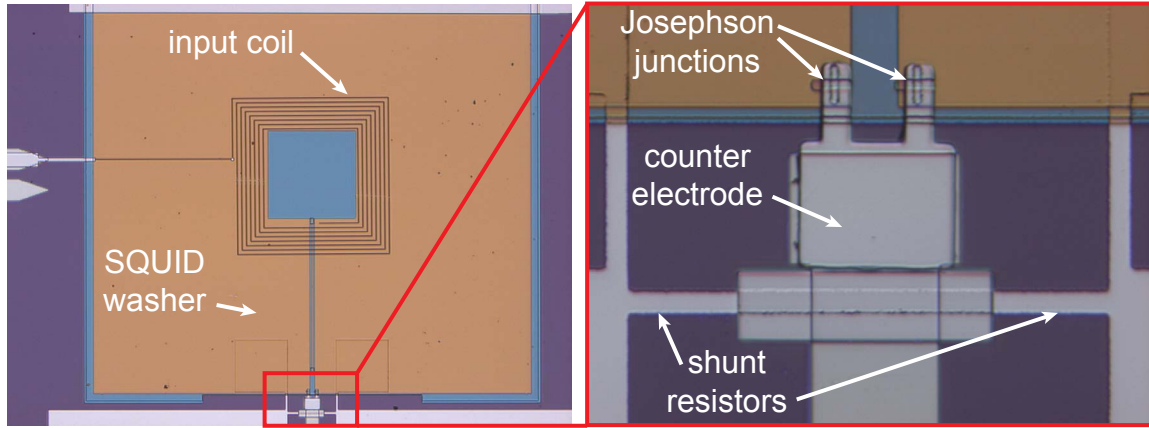


Figure 4.10. Optical photograph of the MSA with magnified inset of the junction region. The MSA consists of a niobium SQUID washer electrically isolated by a SiO layer from a strongly coupled, multiturn superconducting coil also made from niobium. The magnified inset shows how the two Josephson junctions are formed at the square intersection of two orthogonal rectangles patterned into two separate dielectric layers. The junctions are connected to each other and to the output via the counter electrode. Resistive palladium traces shunt the junctions so that hysteresis in the SQUID IV curve is eliminated and the SQUID can operate as a high frequency flux-to-voltage transducer.

[98, 99]. This value is inferred from experimental measurements of M_T , which is related to the coupling strength g through

$$g = M_T I_q I_r, \quad (4.4)$$

where I_r is the zero-point current in the resonator given by

$$I_r = \sqrt{\frac{\hbar \omega_r}{L_r}}, \quad (4.5)$$

and L_r is the resonator inductance. M_{kin} is determined by M_T and M_{mag} , which is modeled using ‘FastHenry.’

4.2 Microstrip SQUID amplifier

In contrast to the flux qubit chips, the MSA (Figure 4.10) is fabricated primarily using photolithography. The superconductor is made of niobium, so that the amplifier is capable of operation and testing at 4.2 K. The MSA was developed in the Clarke lab almost 15 years ago [41], so there are existing recipes and techniques for making these devices.

4.2.1 Previous work

The MSA is fabricated mostly at the 4" wafer level using photolithography in the Marvell Nanolab at the University of California, Berkeley. These wafer-level steps were previously completed on several wafers by Darin Kinion, a former collaborator currently at Lawrence Livermore National Laboratory. Though this process has been internally documented in detail by previous group members, a summary will serve to provide context for departures from the traditional procedures, which are developed in the following sections.

Starting from a silicon substrate with a thin SiO_2 thermally grown oxide on the surface, the entire wafer is prepared and spin coated with photoresist on one of the photoresist coat tracks (i.e. 'svgcoat3') in the Nanolab. The shunt resistors are first patterned using the Nanolab wafer stepper 'gcaws2' in an array with an x-y spacing of 5 mm, the size of the final MSA chip. Approximately 55 nm of palladium is then deposited using an e-beam evaporator by one of the Nanolab staff members, Xiaofan Meng, in a liftoff process. Immediately following this step, the entire wafer is coated with about 200 nm of niobium. After the wafer is again spun with photoresist, the SQUID washer, bonding pads, and all wiring leads are defined in the resist by exposing everything but the desired features. The remaining niobium is etched away using a Reactive Ion Etch (RIE). Two separate layers of silicon oxide (SiO_2 , dielectric constant $\epsilon \approx 5.5$) are then deposited on top of the SQUID washer. Small windows are patterned into this dielectric: one on each end of the 'crossunder' which allows electrical access to the innermost turn of the input coil, and rectangular openings which are used to define the Josephson junctions. There are two oxide steps because the first deposition creates a horizontal rectangular window, while the second crosses the first vertically. This scheme creates a square via to the niobium SQUID washer underneath, preserving the sharp junction window edges that would be considerably rounded if the oxide were deposited in a single step. Uniformity in the junction areas is critical for preventing asymmetries, thus minimizing sample-to-sample area spread.

The preceding fabrication steps are relatively straightforward, with the notable exception of obtaining the correct palladium thickness. Because the mean free path of conduction electrons at millikelvin temperatures is not much shorter than the thickness of the film, the shunt resistance R_S is strongly affected by boundary scattering and does not scale with the typical $R \sim \rho l/A$ dependence, where ρ is the resistivity, l is the length, and A is the cross-sectional area. Small changes in film thickness near ~ 55 nm result in larger than expected changes in resistance. Due to lack of fine control in the palladium deposition process, the film thickness and thus the shunt resistance can significantly deviate from the targeted value ($\beta_C \sim 0.5$) which is optimized to maximize gain while preventing hysteresis in the SQUID IV characteristics. The shunt resistance can later be modified through somewhat painstaking means, either by argon ion milling down the thickness of the shunt (increasing R_S) or shorting out a section of the resistive metal with an enlarged (niobium) counter-electrode (decreasing R_S).




Figure 4.11. Niobium/aluminum sputtering system including argon ion mill. The input coil and Josephson junctions of the MSA are formed in this sputtering chamber. A manually rotating stage can place the sample underneath either a niobium target, an aluminum target, or an argon ion mill, allowing for the formation of Nb/ AlO_x /Nb junctions in a single fabrication step. The ion mill power supply shown at left is much larger than equivalent, contemporary supplies.

4.2.2 Input coil

The final steps in completing an MSA involve forming the resonant input coil and completing the SQUID Josephson junctions. Traditionally, this was completed in single photolithographic step, but in the current process was separated so that the input coil could be narrowed from $5 \mu\text{m}$ down to $1.2 \mu\text{m}$ using e-beam lithography. The motivation for this lies in increasing the operating frequency of the device without sacrificing gain. The resonant frequency of the MSA scales with the expected $1/\sqrt{LC}$ dependence. Without the presence of the SQUID, narrowing the input coil linewidth increases the inductance but is counteracted by the decreasing capacitance to ground. Consequently, the resonant frequency is only weakly affected. With the SQUID washer forming the ground plane of the microstrip mode, however, an extra term is added to the total inductance, giving $L_T = L_e + n^2 L_{S,(\text{MSA})}$, where L_T is the total inductance, L_e is the self-inductance of the input coil, n is the number of turns, and $L_{S,(\text{MSA})}$ is the SQUID inductance [74]. Thus as the linewidth is narrowed, the capacitance decreases while the inductance only slightly increases, leading to an increase in the resonant frequency. Because the number of turns is not altered, the mutual inductance between the input coil and the SQUID remains unchanged. Although the dependence of the resonant frequency on the input coil linewidth was not quantitatively studied, an overall increase in the MSA resonant frequencies was observed.

4.2.3 Junction process

The Josephson junctions are formed together with the counter-electrode, which connects the SQUID to the output lead. The 4x projection aligner ‘*canon*’ in the Nanolab is used to open up a window over the small square vias defined in the dielectric. The sample is loaded into the aforementioned sputter system (Figure 4.11), and the chip is first argon ion milled to remove the native niobium oxide which inhibits the formation of Nb/AlO_x/Nb junctions. The ion mill works similarly to sputtering, except that an ionized argon beam is accelerated towards the sample, milling down any exposed structures. With the oxide removed, a very thin (~ 6 nm) aluminum film is then immediately sputtered. This film is oxidized in situ by filling the vacuum chamber with a prescribed mixture of 80% Ar / 20 % O₂ gas for a specified duration. For typical junctions, the sample is oxidized at a pressure of 70 torr for ~ 45 minutes. It is worth noting that the aluminum layer is not oxidized to completion in this recipe, making the junction critical current independent of the thickness of the deposited Al layer. At 4.2 K, even though the unoxidized aluminum is above its bulk transition temperature of 1.2 K, the proximity effect [100] ensures that it is superconducting. After oxidation, a thick Nb counter electrode (~ 250 nm) is sputtered to complete the process.

Significant variability was seen in the junctions produced in this process. In the past, the system was more reliable once the correct parameters were dialed in. If the process is interrupted for times on the order of weeks, the optimal junction parameters tend to drift to new values for unknown reasons. Even with a recipe which appears to be dialed in, unacceptable run to run variation is still observed. We suspect that this is attributable mostly to the ion mill step. If the sample is under-milled, the NbO_x, which mills slowly, will not be entirely removed. If the sample is over-milled, the niobium metal, which mills at a much higher rate, may become very rough. It is then possible for pinhole shorts and non-uniform oxide barrier thicknesses to form, increasing the chance of short-circuited junctions and critical current variability. The sample stage is rotated between the two sputter guns and the ion mill manually, and sloppiness in this manual dial and apparent shifts in the beam position are common. The ion milling rate also appears to have a spatial nonuniformity, enhancing the problem. In order to have a chance at relatively reproducible results, the exact same procedure and manipulation of the equipment must be performed during each junction process.

4.3 Josephson parametric amplifier

Images of a completed paramp are shown in Figure 4.12. The on-chip overlap capacitors are formed with the same techniques discussed in Section 4.1.6. For a given resonant frequency ω_p and parallel resistance R (real part of the impedance connected to the amplifier, typically 100 Ω), the capacitance determines the quality factor of the resonance through $Q = \omega_p RC$. This is typically aimed at $Q \approx 30$ [79]. The SQUID is fabricated using typical




Figure 4.12. Images of a Josephson parametric amplifier. (a) An optical overview of the paramp shows a niobium plate separated from two aluminum top plates by a SiN_x dielectric, forming two capacitors in series. The capacitors are shunted by a small SQUID which is placed adjacent to the niobium plane. (b) A magnified optical inset shows the SQUID region in more detail. (c) The Josephson junctions are more clearly visible in an SEM image of the SQUID. Reproduced from [79] with permission.

double-angle shadow evaporation. The Josephson junctions of the SQUID are engineered to provide the inductance which is necessary for the targeted resonant frequency ω_p .

Chapter 5

Measurement infrastructure

5.1 Sample boxes

All the devices measured in this work were housed in custom metal boxes. It is vital that qubits and sensitive amplifiers are shielded from stray electromagnetic radiation. A box which is formed entirely of metal, making a tight seal at all interfaces, forms a Faraday cage shield. Small holes in the box are strategically placed for the various connectors. All of our boxes were made from oxygen-free high thermal conductivity (OFHC) copper to provide the best thermalization of the sample and the contents of the box. Additionally, boxes which included internal flux bias coils were plated with lead. This forms a Meissner shield which provides local magnetic shielding. Boxes which relied on flux biasing from external coils could not be plated. The qubit and amplifier devices were affixed to printed circuit boards (PCB) which are bonded to the sample box either mechanically or with solder. The PCB is electrically connected to the device using aluminum wirebonds.

Separate layers of magnetic shielding surround each sample box. The boxes are first placed inside either a rectangular or circular superconducting shield, which is fabricated from aluminum or lead. The bottom of these shields is closed to inhibit flux from penetrating the shield axially. To prevent the shields from trapping magnetic flux as they are cooled through T_c , they are nested inside an outer cryoperm shield (produced by Amuneal Inc.) of a similar geometry. The cryoperm is optimized for its highest magnetic field attenuation at about 4 K, and its performance drops off as the temperature is further reduced. This is acceptable, however, once the superconducting shield has become effective.

Figure 5.1 shows a lead-plated sample box used for flux qubit measurements in the SQUID resonator readout. The microwave traces in the box were designed with a $50\text{-}\Omega$ characteristic impedance. Bulkhead SMA connectors (AEP/Radiall 9308-9113-001) are used to launch the readout signal onto microstrip PCB traces, while a right angle jack SMP (Rosenberger 19K202-271E4) which is soldered to the PCB CPW trace and connects to a straight plug PCB SMP (Rosenberger 19S102-40ML5) is used for qubit control and fast flux pulses. The coupling capacitors (10 pF) on the readout microstrip traces define the Q of the SQUID




Figure 5.1. Flux qubit sample box for SQUID resonator readout.

resonator. The superconducting DC lines used to drive the on-chip flux bias enter the box through 4-pin Reichenbach bulkhead connectors. Large conical choke inductors on the flux bias traces present a high impedance to high frequency noise.

The launch and experimental mount for the circuit QED-style readout is shown in Figure 5.2. Edge-mount connectors (Southwest Microwave 292-04A-5) launch both the qubit and readout pulses onto the PCB. The signal propagates down a $50\text{-}\Omega$ CPW until it encounters the sample, which is positioned in a small recessed section of the PCB. A copper housing which fits over the top of the PCB provides RF shielding. The box is anchored to the dilution refrigerator using a copper mount which contains the superconducting flux bias coil. These ~ 250 turn coils are placed as close as possible to the sample, such that it takes about 5 mA to sweep through a flux quantum in the $3.8 \times 4.0 \mu\text{m}^2$ flux qubit. This corresponds to a mutual inductance of about 0.4 pH.

The MSA is housed in the lead-plated cylindrical can shown in Figure 5.3. The cylindrical geometry is intended to enhance the magnetic shielding of the MSA. The superconducting shield is completely continuous and free of interfaces with the exception of the top cap. This allows the supercurrents necessary for magnetic field expulsion to flow unimpeded in the shield. Copper microwave coaxial lines are passed through the top cap where they join with the SMA connectors attached to the board. The current and flux bias lines enter the can through a single 4-pin Reichenbach connector also on the top cap. The PCB contains a microstrip input and output, which are separated by a $\sim 5 \times 5 \text{mm}^2$ area for the MSA chip. The input trace is coupled to the MSA via a 1.0 pF series capacitor, which is optimized for critical coupling. The output trace features a large blocking capacitor and a conical inductor which allows a static bias current to be passed through the SQUID. A small ~ 10 turn coil is fixed with epoxy to the back of the PCB and is used to apply a flux bias to the SQUID.

The paramp packaging is shown in Figure 5.4. Unlike the previously discussed devices, the paramp incorporates a differential launch. A single-ended microwave signal enters the




Figure 5.2. Flux qubit sample box for circuit QED readout. (a) Photograph of exposed PCB, including the CPW launch and a flux qubit sample. (b) The lid is fastened to the box and the assembly is mounted to a copper finger. (c) A side view of the setup clearly shows the superconducting coil used to flux bias the qubit.

aluminum paramp shield, but is converted to a differential excitation on the PCB by a modified ‘rat-race’ hybrid. The paramp is excited differentially, thus defining a ‘virtual ground’ at the center of the two paramp series capacitors. This differential excitation effectively transforms the impedance of the feedline up to $100\text{-}\Omega$. Similar to the circuit QED flux qubit box, the paramp is biased with a large superconducting coil which rests directly beneath the PCB.

5.2 Dilution refrigerator

A dilution refrigerator (Figure 5.5) was used to cool all the samples to a temperature of about 30 mK, where the characteristic energy of readout and qubit manipulation photons $\hbar\omega$ is larger than the thermal background energy $k_B T$. The measurements were performed in a VeriCold ‘dry fridge,’ which is a relatively recent departure from the traditional ‘wet fridge’ that incorporates a large liquid helium bath. This surrounding bath is replaced by a pulse tube cooler which is anchored with separate cold heads to the top two plates of the fridge, cooling them to $\sim 70\text{ K}$ and $\sim 4\text{ K}$, respectively. A closed circuit He-3/He-4 dilution unit is used to cool the lower plates to temperatures of 700 mK, 100 mK, and finally 30 mK at the ‘mixing chamber.’ Cooling to base temperature is provided by pumping He-3 atoms across a phase boundary between He-3 rich and He-3 poor mixtures of liquid helium.




Figure 5.3. MSA cylindrical can and cryoperm shield. (a) The PCB connects to the lid of the lead-plated cylindrical shield. The MSA is mounted at the intersection of two microstrip traces and a ground. (b) A side view of the PCB gives a better view of the wiring underneath the board. (c) The assembled cylindrical sample box is shown, including a view of the microwave input, output, and the DC cable. (d) The can is placed inside a cryoperm octagon lined with lead foil.

Joule-Thompson cooling assists in initially liquifying the helium gas, partially compensating for the lack of a ‘1 K pot’ which is found in wet fridges.

5.3 DC bias circuitry

The VeriCold is equipped with two types of low frequency lines. When high currents are necessary, particularly to provide a flux bias to small loops using weakly coupled bias lines, superconducting lines are necessary to avoid thermal loading of the fridge. The flux bias for both qubit readout schemes as well the paramp flux bias was provided by these superconducting lines. The wires are broken out from a loom at the 4 K stage, where they are run through a network of LC π -filters. These filters roll off at about 1 MHz. The superconducting lines are run down to the mixing chamber in twisted pairs, shielded by stainless steel (to avoid thermal loading) braided cables. Copper powder filters (Figure 5.6)




Figure 5.4. Paramamp launch and shielding. (a) A customized hybrid is used to convert a single-ended microwave input to a differential signal at the paramp. The PCB and a superconducting bias coil are mounted to a copper finger. (b) The entire setup is encased in a superconducting aluminum shield. (c) Cryoperm surrounds the aluminum and reduces flux trapping in the superconducting shield as it is cooled through T_c . Reproduced from [79] with permission.

are located at each stage below 4 K. These filters are effective at filtering noise due to signals on the order of ~ 10 MHz and higher; in particular, gigahertz noise is strongly attenuated [57]. Copper powder filters consist of superconducting wire wound around cured, ~ 1 mm ‘cores’ of a 2:1 (by weight) Cu powder to Stycast 1266 mixture. These wound cores are encased in a similar 1.5:1 mixture and are placed in a shielded copper box. The RF loss in the filters is thought to originate from eddy currents generated in the Cu particles suspended in the Stycast matrix. The long length of wire in the Cu powder filters, combined with the surrounding Stycast, also provide good thermalization of the lines.

The currents used to bias the MSA are carried by resistive manganin wires. Manganin is an alloy of copper, manganese, and nickel whose resistance changes very little with temperature. Thus the lines maintain their resistive nature at low temperatures, and also have the advantage of being very poor thermal conductors, thermalizing well with their environment, and intrinsically adding resistive filtering. These resistive lines are well-suited to the low current and flux bias currents ($< 20 \mu\text{A}$) required to operate the MSA. The filtering and shielding of these lines is virtually identical to their superconducting counterparts, with the exception of the π -filtering at 4 K, which is RC instead of LC. These filters also have a lower cutoff frequency of ~ 100 kHz.

It is especially critical for the qubit flux bias to be very stable and exhibit low noise to prevent qubit decoherence. We avoid the use of active electronics with their associated 60 Hz noise and instead use a custom battery-powered current source. The circuit diagram for this supply is shown in Figure 5.7, and a detailed account of its configuration and performance is found in Reference [101]. A Xilinx Spartan 3E FPGA (field-programmable gate array) board is programmed by the measurement computer to send digital SPI (serial peripheral interface) commands to a photodiode transmitter, which in turn sends light pulses down fiber




Figure 5.5. Dilution refrigerator including sample mounts. (a) Overview of the dilution refrigerator. A series of nested radiation shields (not shown) slide up onto the fridge, as in a Russian matryoshka doll. A vacuum can surrounds the entire setup and the system is evacuated prior to cool down. (b) Close-up view of mixing chamber stage with sample mounts. The assembled mounts bolt to the bottom of the mixing chamber and are enclosed in aluminum and cryoperm shields.

optic cables to a photodiode receiver at the battery powered supply. These SPI commands program the states of digital potentiometers, which adjust the voltage tapped off a 2.5 V voltage reference. The potentiometers are configured such that the current is discretely adjusted in coarse, medium or fine steps. The resolution of the current source is about one part in 10^6 , which allows very fine flux bias resolution in the qubit. The maximum sourced current is adjusted by changing a resistor on the output of the supply.

The MSA is also current and flux biased with battery-powered supplies. These sources operate by tapping off an adjustable potentiometer which is connected across a battery. The output is in series with a large resistance ($\sim 100 \text{ k}\Omega$) so that the load is a negligible fraction of the total circuit impedance. The outputs are also heavily low-pass filtered, and the entire supply is shielded in an aluminum box.




Figure 5.6. Copper powder filter. A copper powder - Stycast mixture is shown filling trenches in a copper box. Insulated wires wound around a similar mixture are encapsulated in these trenches. This box is shown unfinished, with wires protruding from the copper powder mixture. The filters are completed by connecting the wires to 4-pin Reichenbach bulkhead connectors and covering the entire box with a Cu lid.

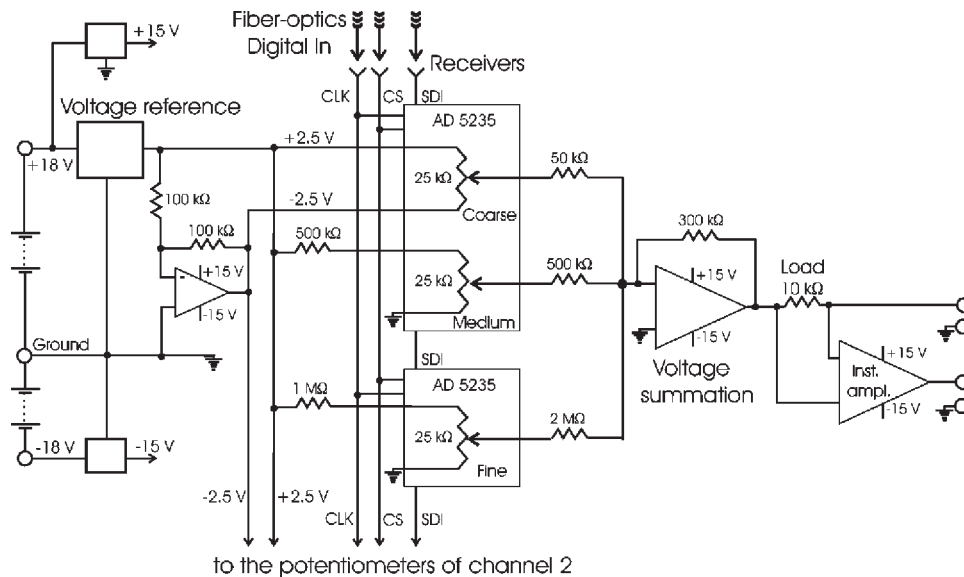


Figure 5.7. Circuit diagram of the low noise, optically-isolated current source [101].

5.4 Microwave circuitry

5.4.1 Coaxial lines and attenuation

The microwave signals used for qubit manipulation and readout propagate on coaxial lines. The material of these lines depends on their particular function and position in the measurement setup. Tin-plated copper cables (Micro-coax UT-085-TP) are used at room temperature and to connect between components on the same temperature stage. Low-loss is desirable and the relatively high thermal conductivity of these cables is acceptable in these situations.

Cables connecting different temperature stages in the dilution refrigerator must be made from alternative materials to prevent thermally shorting the stages together. The injection (input) lines into the fridge consist of cables made of stainless steel (Micro-coax UT-085B-SS and Coax Co. SC-219/50-SS-SS). In addition to good thermal isolation, these cables also substantially attenuate microwave signals. This is actually a benefit on injection lines because Nyquist noise at room temperature and intermediate temperature stages in the dilution refrigerator must be attenuated to the level of the mixing chamber noise. The high attenuation ensures that noise from higher stages does not significantly add to the intrinsic noise at the mixing chamber. To aid these lossy cables in this function, discrete NiCr attenuators (XMA 2782-6051-03, -10, -20) are also placed at various stages in the fridge. The power attenuation of these components is expressed in decibels (dB) as $A_{\text{dB}} = 10 \log_{10} A$, where A is the linear power attenuation. As a rule of thumb, the attenuation at each stage of the fridge in a properly designed injection line should at least be equal to, and preferably exceed, the ratio of the temperatures between the stage the attenuation is anchored to and the next highest stage.

Like their injection line counterparts, the output microwave lines in a dilution refrigerator need to provide good thermal isolation, but attenuation on the output is directly linked to degradation in the system SNR. Minimizing this attenuation is especially critical when linking the amplification stages which primarily set the overall system noise temperature. Neglecting the Nyquist noise generated by the attenuation itself, an attenuation A before an amplifier with gain G and noise temperature T_N reduces the effective gain to G/A and increases the corresponding noise temperature to AT_N . The situation is somewhat more complicated, and the effects more detrimental, in dilution refrigerators when the effect of noise generated by attenuation in temperature gradients is considered. Low-loss copper cables are used to link the qubit readout resonators to the first stage of amplification, either the MSA or the paramp. The gain of these amplifiers is typically not high enough to completely dominate the noise generated by the HEMT, especially in the presence of significant attenuation before the HEMT, and thus low loss between the MSA/paramp and HEMT is highly desirable. The HEMT, however, must be anchored to the 4 K stage because of its relatively high power dissipation. The engineering problem of low microwave loss coupled with good thermal isolation between fridge stages is solved with the use of niobium coaxial line (Coax Co. SC-219/50-Nb-Nb). A single, uninterrupted length of Nb coax connects

the mixing chamber and the 4 K stage. Following the HEMT output, both the signal and noise power are much higher than the Nyquist noise at room temperature, and consequently stainless coax can again be used from 4 K to room temperature to prevent thermal loading.

5.4.2 Cryogenic setup

The full cryogenic microwave setup for the SQUID resonator readout is shown in Figure 5.8. The fast flux line features less attenuation than the other injection lines because the currents ($\sim 150 \mu\text{A}$ on the sample chip) required to produce sufficient flux shifts in the qubit would excessively heat the fridge if more attenuation were included. The smaller level of attenuation is partially compensated by custom microwave ‘roach filters’ [102]. These filters are essentially a lossy stripline, with high frequency attenuation provided by magnetically loaded silicon dielectric (Emerson & Cumming Eccosorb MFS-117). The roach filters used in this work were shielded in a copper box and had a 3 dB rolloff at 1.3 GHz. The filters heavily attenuate signals up to at least 40 GHz, yet simultaneously pass DC current with very little dissipation.

A directional coupler is a four-port (one of which we terminate with a matched load) passive component which weakly couples microwave signals between two transmission lines brought into close proximity [65]. We use a directional coupler to merge the qubit excitation line with the fast flux line at the mixing chamber stage. The qubit excitation is sent into the ‘coupled port’ of the directional coupler, where the signal is effectively attenuated by 20 dB, while the fast flux pulse is directed into the ‘through port’ and travels through the directional coupler with virtually no loss.

The readout signal is routed to the SQUID resonator via the coupled port of a separate directional coupler. After the signal reflects from the resonator, it passes back out of the through port of the coupler and is directed towards the microwave output chain. The signal first encounters two circulators [Figure 5.9(a)], which in this context function as isolators. A circulator is a non-reciprocal, three port device which relies on a magnetized ferrite to route a signal only to the adjacent port in a circulating path [65]. In other words, signals entering port 1 are only routed to port 2, signals from port 2 only go to port 3, and signals from port 3 only go to port 1. A circulator can be used in a role similar to a readout directional coupler to send an incoming signal to a resonator in a reflection geometry. The reflected signal passes back through the circulator and is sent to the output chain rather than back into the injection line. If one of the ports of the circulator is terminated with a matched load, the device functions as an isolator. Signals are able to propagate with little (~ 0.5 dB) loss in one direction, but are strongly suppressed (-20 dB) in the reverse direction. This isolation prevents noise generated by the MSA from radiating back towards the qubit. Additionally, reflections from the input of the MSA are strongly damped without losing signal in the forward direction.

After the network of circulators, the signal encounters two modified double-pull double-throw microwave switches (Radiall R577433000) [Figure 5.9(b)]. These mechanical ‘transfer’ switches are actuated by applying current pulses to solenoid coils, generating substantial

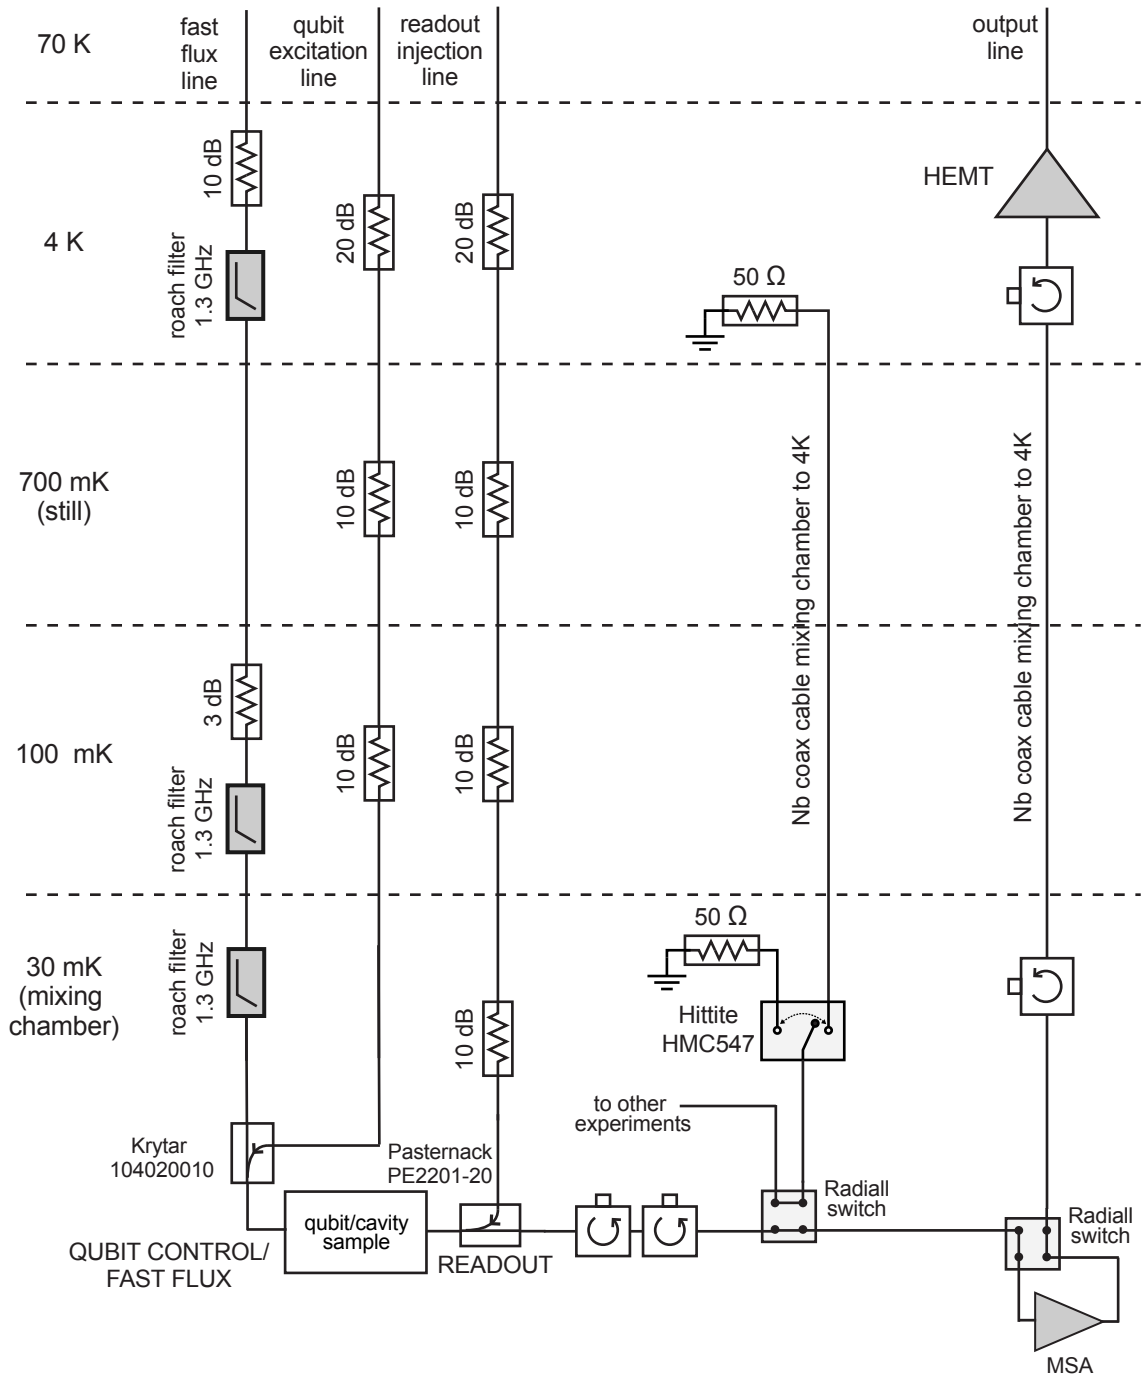


Figure 5.8. Cryogenic microwave setup for SQUID resonator readout.




Figure 5.9. Microwave switches and circulators. (a) An array of circulators used for isolating and routing signals is attached to the mixing chamber. (b) The Hittite solid-state switch is used for switching between two calibrated noise sources, enabling measurements of the system noise temperature. Switches manufactured by Radiall are used primarily for switching between samples and bypassing superconducting amplifiers.

heat loads in the process. The first transfer switch allows for system noise temperature measurements using a ‘hot/cold’ load method. These measurements will be further discussed in Chapter 6. A path to the output chain from either of two separate, calibrated noise sources can be defined using the transfer switch. A separate, solid state switch (Hittite Microwave, HMC 547) switches between the noise sources [Figure 5.9(b)].

The second transfer switch is used for taking the MSA into and out of the output chain. This is useful for calibrating the gain of the MSA by comparing its performance to a straight through section of coaxial line. The switch also allows for a direct quantitative comparison of the effect of the amplifier on the flux qubit readout. Following the MSA, two circulators on each end of a niobium coax isolate the MSA from the HEMT noise at 4 K. The broadband Caltech HEMT has a specified noise temperature of about 7 K at 1.4 GHz.

A schematic of the cryogenic portion of the circuit QED readout is shown in Figure 5.10. The readout injection line is similar to the SQUID resonator setup, but after the directional coupler a switch (Radiall R573423600) has now been added which allows up to six samples to be loaded and separately probed with microwave reflectometry. One of these six slots is occupied with the hot/cold load setup for noise temperature measurements. After the directional coupler, there is a network of circulators and a transfer switch for switching the paramp out of the output circuit. The larger number of circulators in this setup is intended to prevent the strong RF paramp drive from leaking back to the qubit and readout resonator.

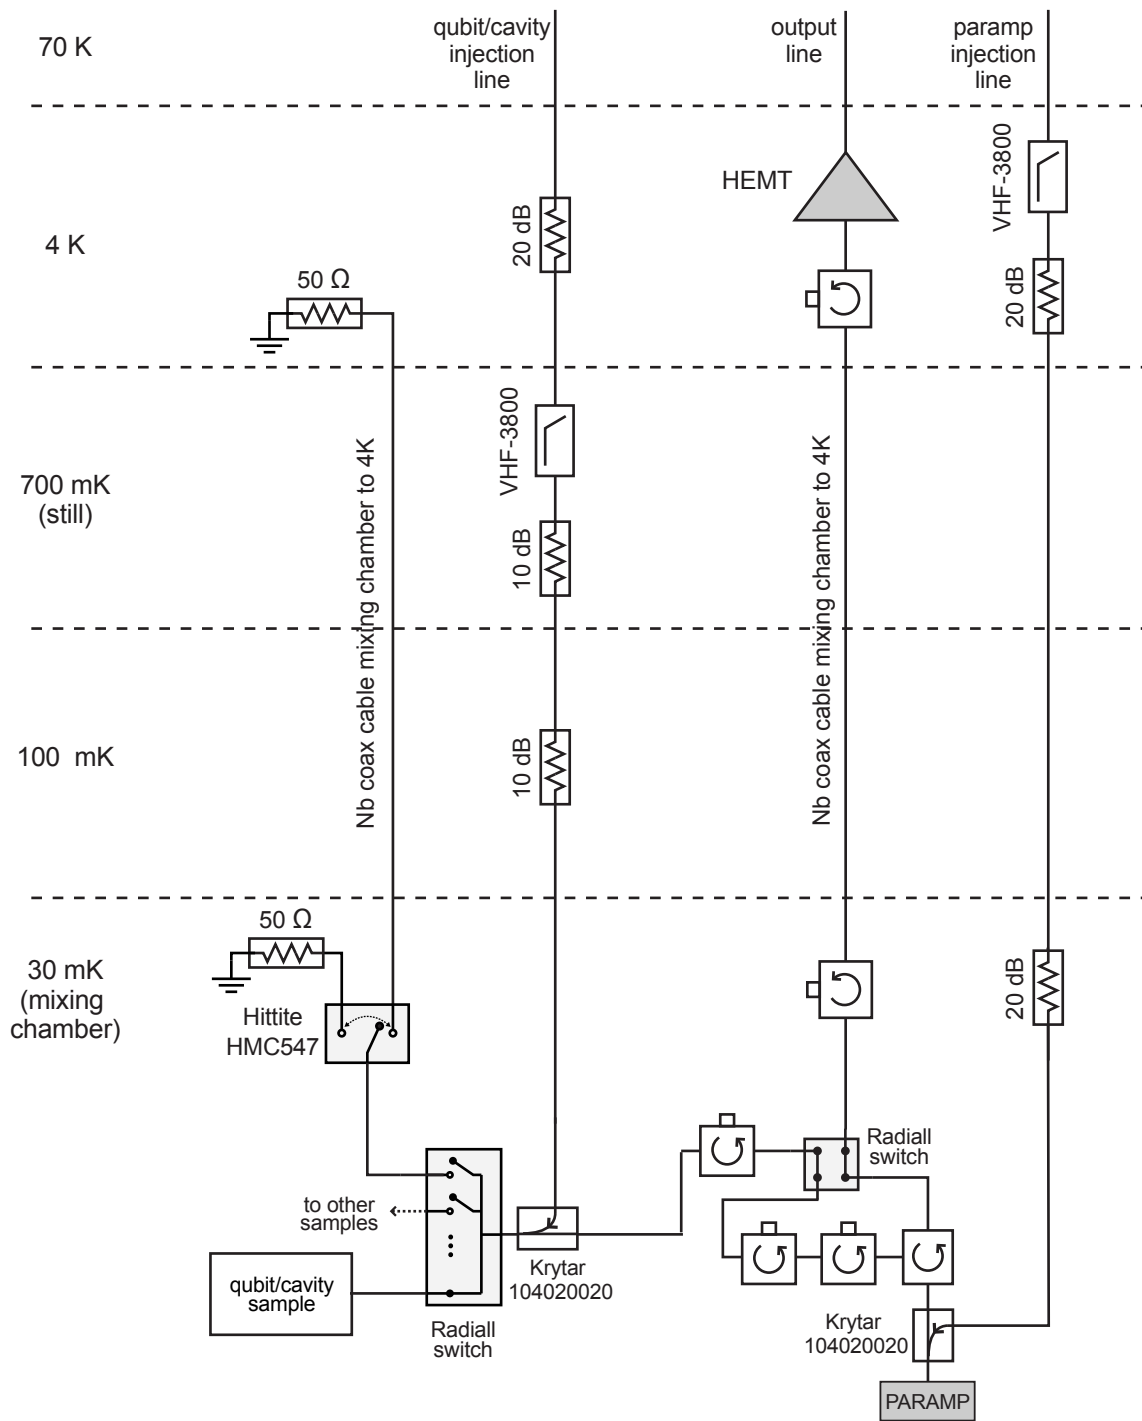


Figure 5.10. Cryogenic microwave setup for cQED readout. Adapted from [79].




Figure 5.11. Room temperature qubit measurement equipment.

After amplification in the paramp, a niobium coax is used to route the signal up to the HEMT (Low Noise Factory LNF-LNC4_8A), which features a noise temperature of about 2.6 K at our 5.8 GHz readout frequency. The paramp drive is provided by an injection line which contains a typical assortment of attenuators.

5.4.3 Room temperature measurement setup

One of the challenges of qubit readouts is synchronizing many sophisticated electronic instruments (Figure 5.11) together with nanosecond scale precision. The room temperature measurement setups for each qubit readout style are shown in Figure 5.12 (SQUID resonator) and Figure 5.13 (cQED). As the setups have many similarities, a thorough explanation of the cQED setup will be sufficient to give context to the SQUID resonator setup, with a few notes of clarification.

Qubit pulses are formed by the mixing a continuous wave (CW) microwave tone with an envelope which defines the shape of the pulse. The envelope modulates the amplitude of the carrier wave, which is resonant with the qubit transition frequency ω_{01} . The mixer which multiplies these two signals must be carefully balanced using static offset voltages to prevent leakage of the carrier tone through the mixer when the qubit excitation is off; otherwise these stray signals can affect the coherence of the qubit. The qubit pulses are shaped with Gaussian envelopes to provide the minimum amount of frequency dispersion for a given pulse duration. The shortest qubit pulses had a width of about 4 ns, with a corresponding rise

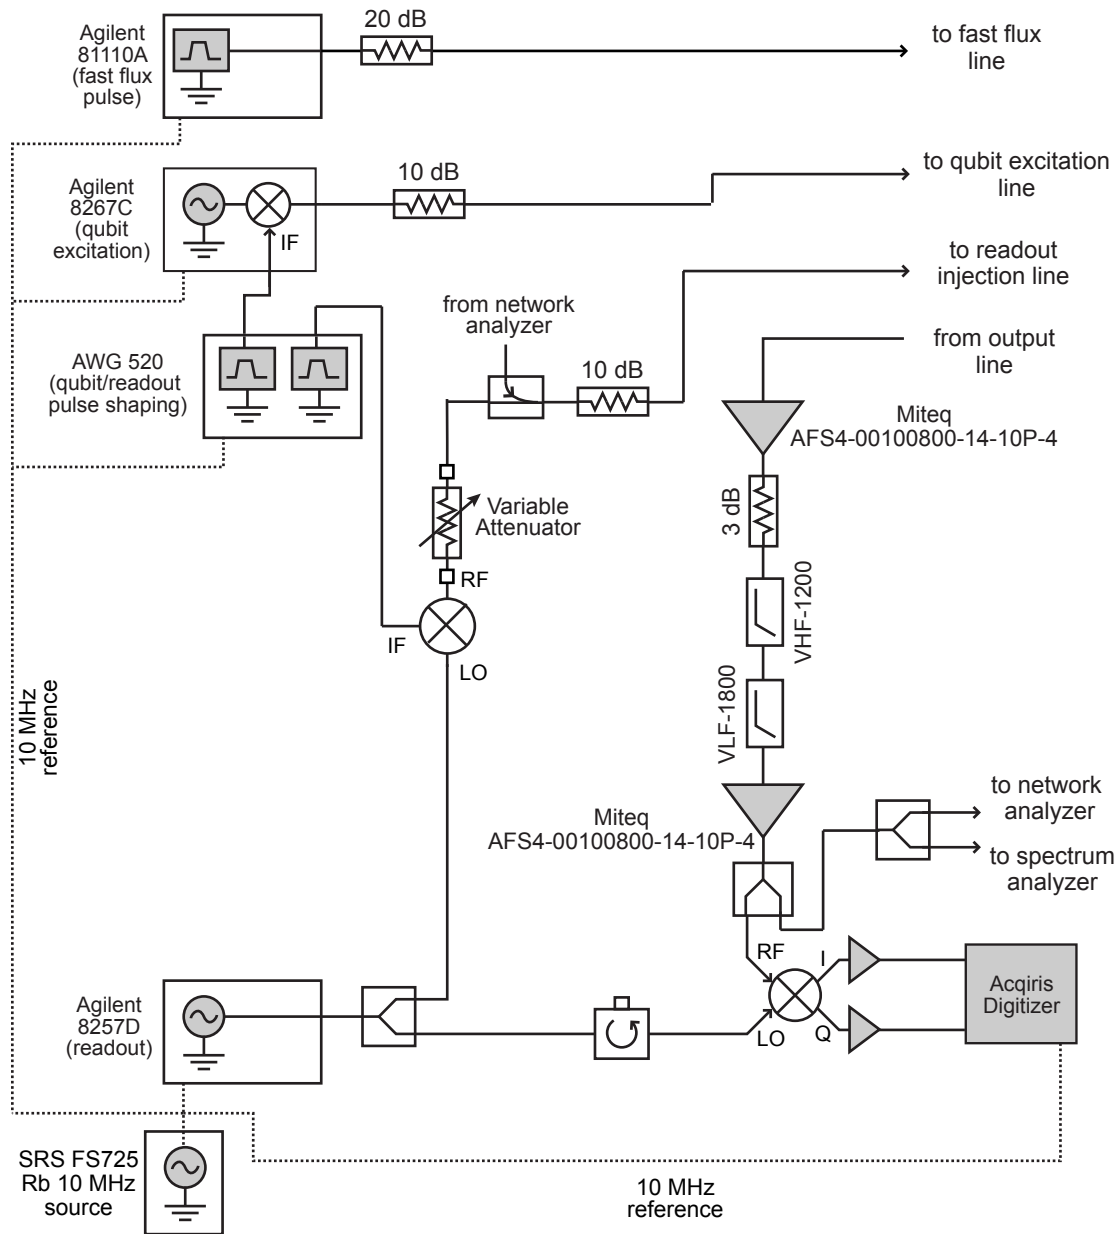


Figure 5.12. Room temperature setup of the SQUID resonator readout.

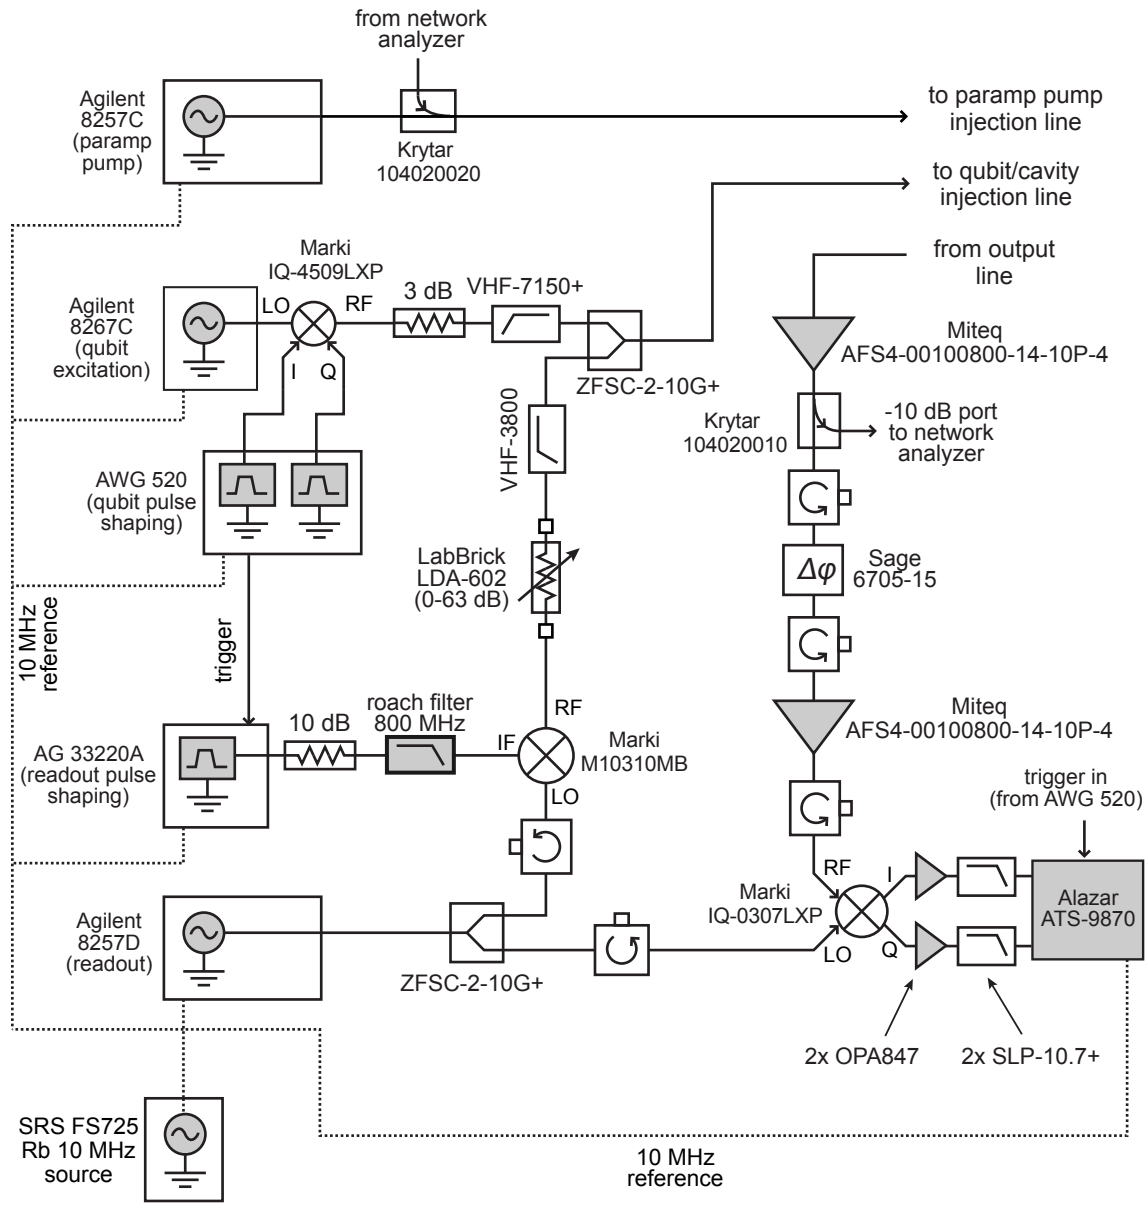


Figure 5.13. Room temperature setup of the cQED readout. Reproduced with minor modifications from [79].

time of ~ 1 ns.

The readout pulse is similarly generated by mixing the CW readout frequency ω_r with the readout pulse shape. The rise time of these pulses can be slower (~ 20 ns) than the qubit pulses because the limited bandwidth of the readout cavity (~ 9 MHz) naturally imposes a slower time constant for oscillations to build in the cavity. Before the CW tone is mixed, it is split in half using a splitter; half the signal is sent to be used in processing the output signal, and half is sent to the mixer where it is shaped. The amplitude of the readout pulse after the output of the mixer is adjusted using a digital, computer-controlled attenuator. Subsequently, the readout pulse is combined with the qubit pulses using a splitter connected in a reverse orientation.

The paramp is pumped at the readout frequency using a separate microwave generator. All the instrument time-bases are synchronized with a common 10 MHz frequency standard. This is especially crucial to ensure that the readout pulse and paramp pump remain phase-locked. Any relative drift between their phases will degrade the noise performance of the paramp in phase-sensitive mode as illustrated by Figure 3.11.

The readout signal emerging from the fridge is amplified by two room temperature amplifiers to further boost the signal in preparation for subsequent processing. These amplifiers do not degrade the overall SNR of the measurement because the signal and noise levels are already much higher than the added noise of these amplifiers ($T_N \sim 175$ K). To detect the amplitude and phase of the readout signal, the signal is demodulated with the readout carrier tone that was originally split from the microwave generator. This down-converts the signal to zero frequency, producing the quadrature amplitudes I and Q of the readout signal with respect to the carrier. A phase shifter is added before the mixer, which outputs the quadrature amplitudes, so that the output vector in the IQ plane can be arbitrarily rotated. For any signal \mathbf{s} between two readout pointer state vectors in the IQ plane, the signal can be rotated so that it is only contained in a single quadrature, an experimental and post-processing convenience. The output signals from the IQ channels are amplified using low-noise op-amps and are then low-pass filtered below ~ 10 MHz. This cutoff is chosen to roughly match the 9 MHz bandwidth of the readout resonator. Thus, higher frequency noise is filtered without reducing the speed at which changes in the signal can be tracked. These signals are then sent to a digitizer for data acquisition.

There are a few noteworthy differences in the room temperature measurement setup of the SQUID resonator readout. Namely, the readout and qubit pulses are not combined, but are separately injected into the refrigerator. The fast flux pulse is generated by a pulse generator with a very fast (~ 1 ns) rise time. This is important to minimize qubit decoherence during the shift, which occurs before the readout is energized. Also, at the output of the demodulation mixer, the quadrature amplitudes are not filtered before they are sent into the digitizer. The effective noise bandwidth can be decreased instead with software averaging.

5.4.4 Data acquisition

The qubit pulse generator (Tektronix AWG520) is responsible for controlling the timing of all the measurement equipment. It accomplishes this by sending out triggers to each instrument. The instruments then respond to these triggers by starting their prescribed sequence. The last object to be triggered is the digitizer, which is integrated into the acquisition computer. Upon receipt of its trigger, the digitizer reads the input voltage at a user-defined rate for a prescribed period of time corresponding to the duration of the readout pulse. This rate was set to 10^8 samples per second in the circuit QED experiment, and 10^9 in the SQUID resonator experiment. The difference in sampling rates was due to the shorter coherence times and higher characteristics bandwidth of the the SQUID resonator readout.

The measurement acquisition software is written in **LabView**. This software is responsible for programming the instruments as well as acquiring, processing, and saving the data. The high digitizer sampling rates, coupled with studies involving a very large number of qubit readouts, makes memory handling in **LabView** a challenge.

Chapter 6

Qubit and amplifier characterization

6.1 Resonator characterization

The characteristics of a qubit readout resonator are initially probed using a vector network analyzer (VNA). In its most basic configuration, this instrument outputs a microwave signal from one port and detects the returning signal through another. The frequency and power of the signal are both adjustable, and the ‘vector’ nature of the analyzer allows it to detect both amplitude and phase.

The reflected phase response of the nonlinear SQUID resonator vs drive frequency and applied flux in the SQUID loop is shown in Figure 6.1(a). The resonant frequency is identified as the zero-crossing (yellow) of the reflected phase. As the applied flux is increased from zero, the resonant frequency modulates with the expected flux periodicity of Φ_0 . The maximum frequency of the resonator is $f_r = 1.49$ GHz, with a quality factor $Q = 10$. This Q was significantly lower than the design value of ≈ 50 . While the mechanism for this reduction is unknown, we speculate that a parasitic capacitance, which effectively shunts the coupling capacitor on the PCB, is involved.

Near $\Phi_x = (m + 1/2)\Phi_0$ (m an integer), where the frequency approaches a minimum due to a maximum in the SQUID inductance, there is an abrupt, discontinuous increase in the resonant frequency. This behavior is expected for a SQUID with a non-negligible β_L , which in this device is approximately 0.35. As the flux is increased beyond $\Phi_0/2$, the phase particle in the SQUID potential is confined to a metastable well with a decreasing barrier height. The particle eventually escapes from the well via thermal activation or quantum tunneling [3], causing the SQUID to switch to a different flux state. The SQUID never enters the voltage state, but simply makes a transition to a state of lower energy by changing the flux in its loop by a single Φ_0 . If the flux is swept in the reverse direction, the response is similar but the flux jump occurs at a different flux, indicating hysteresis in the system.

The sensitivity of the resonator to the qubit flux states is proportional to $\partial f_r / \partial \Phi$, where f_r is the resonant frequency of the SQUID resonator. To achieve a large readout signal from the circulating current states, we flux bias the SQUID at $\Phi_x = 0.43\Phi_0$ such that its




Figure 6.1. Response of the SQUID resonator vs. flux and power. (a) Reflected phase of the SQUID resonator vs frequency and flux. The resonant frequency modulates with applied flux, tuning with a flux periodicity of $\Phi_0/2$. The discontinuities in flux near $\Phi_x = \Phi_0/2$ are due to transitions between flux states in the SQUID. For qubit measurements, the SQUID is biased at $\Phi_x = 0.43\Phi_0$ such that $f_r = 1.294$ GHz, where the flux-to-frequency transfer function $\partial f_r / \partial \Phi_x$ is large. (b) Reflected phase of the SQUID resonator vs incident power and frequency at $\Phi_x = 0.43\Phi_0$. With increasing incident power, nonlinear effects are visible as the resonant frequency decreases and the phase response sharpens.

frequency $f_r = 1.294$ GHz, corresponding to a measured $\partial f_r / \partial \Phi_x \approx 1.05$ GHz/ Φ_0 . This particular frequency was chosen to match the operating frequency of the MSA.

With the flux in the SQUID fixed, we now probe the phase response of the nonlinear resonator vs drive power and frequency, as seen in Figure 6.1(b). The phase response remains unchanged until a power of about -110 dBm¹, after which the nonlinearity is manifest as the resonant frequency decreases. The sharpening phase response is evident by the decreasing width of the yellow zero-crossing region in phase. The resonator exceeds the critical point and becomes bistable at about -103 dBm.

The characteristics of the linear resonator used in the cQED readout are much simpler. The bare (unaffected by the qubit) resonant frequency is $f_r = 5.780$ GHz with a quality factor $Q = 640$. This corresponds to an oscillator decay rate of $\kappa/2\pi = 9.0$ MHz.

6.2 Flux qubit - resonator interaction

When the flux in the SQUID resonator is fixed while the qubit flux is changed, there is only a very weak response in f_r through most of a flux quantum in the qubit. Near $\Phi_0/2$ in the qubit, however, the admixture of circulating current states in the qubit changes rapidly, also changing the flux in the SQUID through the qubit-SQUID mutual inductance M_{QS} . This region is known as the ‘qubit step’ [103]. Figure 6.2 shows the phase response of a SQUID resonator (sample F121409) vs frequency and flux near the qubit step. Each

¹A power P_{dBm} expressed in dBm is related to a power P in watts by $P_{\text{dBm}} = 10 \log_{10}(1000P)$.




Figure 6.2. The qubit step. SQUID resonator reflected phase vs frequency and flux near the ‘qubit step.’ As the flux in the qubit is swept through $\Phi_x = \Phi_0/2$, the magnetic flux coupled to the SQUID from the qubit varies smoothly from $-M_{QS}I_q$ to $M_{QS}I_q$. This causes a small shift in the resonant frequency of the SQUID resonator.

side of the step, before the phase response appreciably changes with flux, corresponds to a qubit eigenstate which very closely resembles a circulating current state. The polarity of these circulating current states differs across the step. Thus, measurements of the qubit step provide the maximum measurable phase signal between the qubit (circulating) states. At the optimal readout frequency, this signal is $\Delta\phi_{\text{circ}} = 19^\circ$.

The flux coupled from the qubit to the SQUID near the qubit step is expressed as [103]

$$\Delta\Phi_{QS} = M_{QS}I_q \left(\frac{\epsilon}{E_{01}} \right) [f_-(T_{\text{eff}}) - f_+(T_{\text{eff}})], \quad (6.1)$$

where

$$f_{\pm}(T_{\text{eff}}) = \frac{e^{\pm E_{01}/2k_B T_{\text{eff}}}}{2 \cosh(E_{01}/2k_B T_{\text{eff}})} \quad (6.2)$$

are the occupations of the ground (-) and excited (+) states at an effective temperature T_{eff} . Equation 6.1 can be simplified to

$$\Delta\Phi_{QS} = M_{QS}I_q \left(\frac{\epsilon}{E_{01}} \right) \tanh \left(\frac{E_{01}}{2k_B T_{\text{eff}}} \right). \quad (6.3)$$

Both the tanh and ϵ/E_{01} terms range from -1 to 1 across the qubit step, so that the total change in the SQUID flux through the step is $2M_{QS}I_q$. Even at $T_{\text{eff}} = 0$, the step has a non-zero width, which is caused by the non-zero energy splitting Δ_q . As the effective temperature

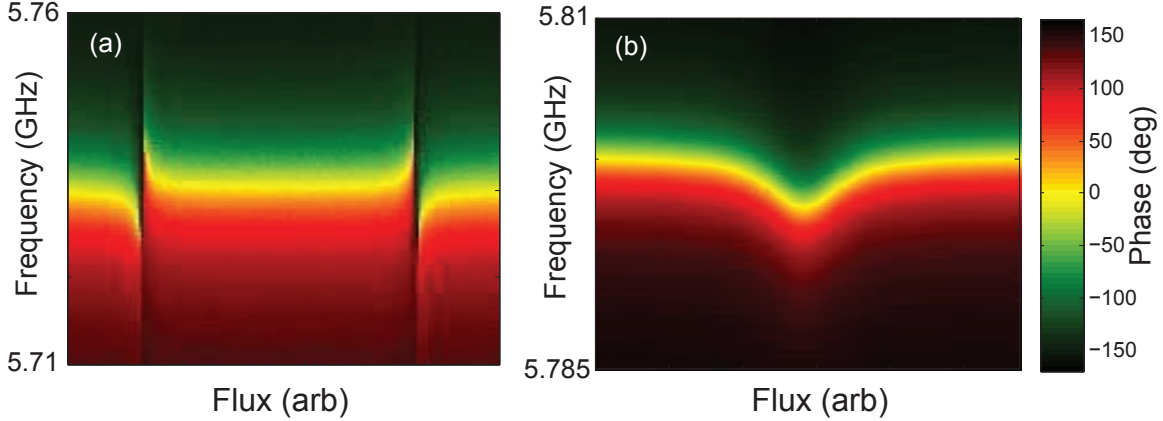


Figure 6.3. cQED resonator response vs. flux in qubit. (a) Reflected resonator phase vs frequency and flux in a sample with $\Delta_q/h < f_r$. Two avoided crossings between the resonator and qubit resonances occur as the flux is tuned near $\Phi_x = \Phi_0/2$. The flux axis is unscaled because no flux calibration was performed. (b) Reflected resonator phase vs frequency and flux in a sample with $\Delta_q/h > f_r$. The resonant frequency is shifted by the flux-dependent amount χ as the flux is tuned near $\Phi_x = \Phi_0/2$.

is increased, the width of the step is broadened in flux. This width is typically on the order of $20 \text{ m}\Phi_0$ in our devices. Once the qubit parameters Δ_q and I_q are determined (Section 6.3), the step can be fit to the form of Equation 6.3. We extract T_{eff} , which is consistently near $\sim 130 \text{ mK}$ in our SQUID resonator samples, as a parameter of the fit. The implications of this high effective temperature will be discussed in Chapter 7.

The cQED readout resonator is sensitive to applied magnetic flux only through its interaction with the qubit. Near $\Phi_0/2$ in the qubit, this interaction is detected by a shift in the resonant frequency. For $\delta \gg g$, the resonator frequency shifts by $\chi = g'^2/\delta$, where both g'^2 and δ are functions of flux (Chapter 3). Figure 6.3 shows the phase response of two different resonators vs frequency and flux in the qubit. The first sample [Figure 6.3(a)] shows the case of $\Delta_q/h < f_r$, where the qubit spectrum passes through f_r twice. Mixed states between the qubit and resonator form when $\delta \sim g$ [104]. This is manifest by an avoided crossing between the qubit and resonator. When $\Delta_q/h > f_r$ [Figure 6.3(b)], the qubit ‘pushes’ the frequency of the resonator lower, but there is no avoided crossing.

6.3 Qubit spectroscopy

The qubit transition frequency vs flux is directly obtained from spectroscopic measurements. At a given flux in the qubit, the qubit excitation frequency is stepped while the




Figure 6.4. Qubit spectroscopy in the SQUID resonator readout. The frequency spectrum of the qubit vs flux is hyperbolic in shape. The color scale indicates the difference in the switching probability of the JBA with the qubit excitation on and off. The flux-independent resonance at 9.35 GHz is likely a cavity mode of the sample box.

Table 6.1. Measured qubit and resonator parameters.

SQUID resonator (F063010c)			cQED (F032812b)		
Parameter	Symbol	Value	Parameter	Symbol	Value
Frequency splitting	Δ_q/h	8.507 GHz	Frequency splitting	Δ_q/h	6.15 GHz
Circulating current	I_q	193 nA	Circulating current	I_q	204 nA
Readout frequency	f_r	1.294 GHz	Readout frequency	f_r	5.780 GHz
Phase contrast	$\Delta\phi_{\text{circ}}$	19°	Coupling strength	$g/2\pi$	105 MHz

response of the readout resonator is measured. A long ‘saturating’ qubit pulse produces an incoherent mixture of the ground and excited states when resonant with ω_q . This signal from the qubit is detected by a change in the reflected signal from the resonator when compared to no qubit excitation. This process is repeated as the qubit flux is also stepped, creating a map of the qubit response in frequency and flux space.

Figure 6.4 shows the measured spectroscopy of the SQUID resonator sample. The JBA readout (Section 3.4), in which the switching probability of the nonlinear resonator between the low and high amplitude states in the bistable regime is measured, was used to map out this spectrum. The color scale shows the difference between the switching probability with and without a qubit excitation. This is plotted vs excitation frequency and flux. The spectrum shows the expected hyperbolic dependence, and is fitted to this functional form (Equation 2.23) to extract the qubit parameters Δ_q and I_q , as shown in Table 6.1. The spectral peaks used for this fit are first identified by fitting the JBA response vs frequency




Figure 6.5. Readout-induced shift in the qubit flux. (a) An oscillating readout excitation in the SQUID resonator vs time. (b) When the SQUID is flux biased away from $m\Phi_0$ where m is an integer, the symmetry of the SQUID is broken with respect to current excitations. An oscillating flux is produced in the qubit which has a dominant frequency component at twice the excitation frequency. The time-average of this flux is non-zero, causing an effective shift in the qubit flux during readout. This effect allows for a magnetization signal during readout at the degeneracy point.

at each flux value to a Lorentzian, the expected form of the resonant peak. We attribute the large flux-independent resonance at ~ 9.35 GHz to a mode of the qubit sample box.

Because the circulating state mixture between the qubit ground and excited states is equal at the degeneracy point, one might expect no spectroscopic signal at the degeneracy point. Moreover, the switching probability difference should be positive on one side of the degeneracy point and negative on the other. The switching probability in Figure 6.4, however, is positive at the degeneracy point and throughout the entire visible part of the spectrum. This discrepancy is due to an effective flux shift in the qubit due to oscillating readout currents in the SQUID resonator. Figure 6.5(a) shows the amplitude of the readout currents flowing through the SQUID vs time. When the SQUID is flux biased, its symmetry is broken and resonator currents divide unequally down each of its two arms. It can be shown that opposite polarity readout currents each divide such that the net flux coupled to the SQUID is of the same single polarity. This is represented graphically in Figure 6.5(b). The average flux shift in the qubit, represented by the straight dotted line, is nonzero, which allows for a magnetization signal even when the qubit state is prepared at the degeneracy point. The magnitude of the shift is dependent on the readout amplitude, and for high power JBA pulses, large shifts are observed. Further implications of this shift will be discussed in Chapter 7.

The qubit spectrum of the cQED flux qubit is shown in Figure 6.6. The resonator




Figure 6.6. Qubit spectroscopy in the cQED readout. The signal in the readout resonator (color scale) is the difference in the average reflected phase with the qubit excitation on and off. The broadening of the spectral line near the degeneracy point is due to the strong qubit-resonator interaction.

signal (color scale) is the average phase difference between measurements with and without the qubit excitation as measured by the network analyzer. The spectral line is significantly broadened near the degeneracy point where the detuning δ is small. This is due to the strong coupling between the resonator and qubit, which mixes these states and dissipates energy from the qubit [105]. Once Δ_q and I_q have been determined from fits of the spectroscopy, $g'^2 = \chi\delta$ is extracted. At the degeneracy point, g' takes on its maximum value g .

6.4 Time domain characterization

With the spectrum determined, the qubit can be biased at a particular transition frequency and characterized in the time domain. These measurements are more easily understood after introducing the Bloch sphere, shown in Figure 6.7. Any arbitrary superposition state of the qubit can be expressed by a vector which points from the center of the sphere to its surface. The north-south character of the vector represents the relative contributions of $|0\rangle$ and $|1\rangle$ to the superposition state. A pure ground (excited) state is represented by a vector pointing to the south (north) pole of the sphere. The azimuthal or equatorial degree of freedom represents the relative complex phase between $|0\rangle$ and $|1\rangle$. In a fashion similar to other magnetic spin 1/2 systems, the state vector precesses in time about the north-south




Figure 6.7. The Bloch sphere. (a) The north-south character of the arrow's position on the Bloch sphere represents the relative mixture of the ground and excited state in the qubit. The ground state is expressed as an arrow pointing to the south pole. (b) The excited state is conversely represented by an arrow pointing to the north pole. (c) An arrow in the equatorial plane corresponds to an equal mixture of the ground and excited states. The azimuthal degree of freedom (η) represents the complex phase between $|0\rangle$ and $|1\rangle$.

axis of the Bloch sphere at the Larmor frequency ω_q . The state can be written as

$$|\psi\rangle = a|0\rangle + be^{i\eta}|1\rangle, \quad (6.4)$$

where a and b are real coefficients satisfying $a^2 + b^2 = 1$ and η is the azimuthal phase angle on the Bloch sphere.

In contrast to the incoherent state mixture prepared during qubit spectroscopy, short qubit pulses are used to create coherent superposition states. Resonant excitation at the Larmor frequency causes the qubit state to rotate around the Bloch sphere in a vertically-oriented plane. The state vector oscillates coherently between the ground and the excited states at a rate proportional to the amplitude of the microwave excitation amplitude. These oscillations are known as Rabi oscillations [106]. With the qubit in the ground state at $t = 0$, the idealized time-dependent coefficient $[a(t)]^2$ in the SQUID resonator readout architecture can be expressed as [107]

$$[a(t)]^2 = \left[\frac{\Omega_r^2}{\Omega_r^2 + (\Delta\omega_q)^2} \right] \sin^2 \left[\frac{\sqrt{\Omega_r^2 + (\Delta\omega_q)^2}}{2} t \right], \quad (6.5)$$

where $\Omega_r = 2\Phi_{\text{osc}}I_q \sin(\xi)/\hbar$ is the Rabi frequency and $\Delta\omega_q$ is the detuning of the microwave driving frequency from the qubit resonance ω_q ; Φ_{osc} is the amplitude of the oscillating microwave flux in the qubit. For $\Delta\omega_q = 0$, it can be seen that $\Omega_r \propto \Phi_{\text{osc}}$. This result was experimentally verified in our qubits and is a characteristic of a quantum two-level system that has no classical analog.

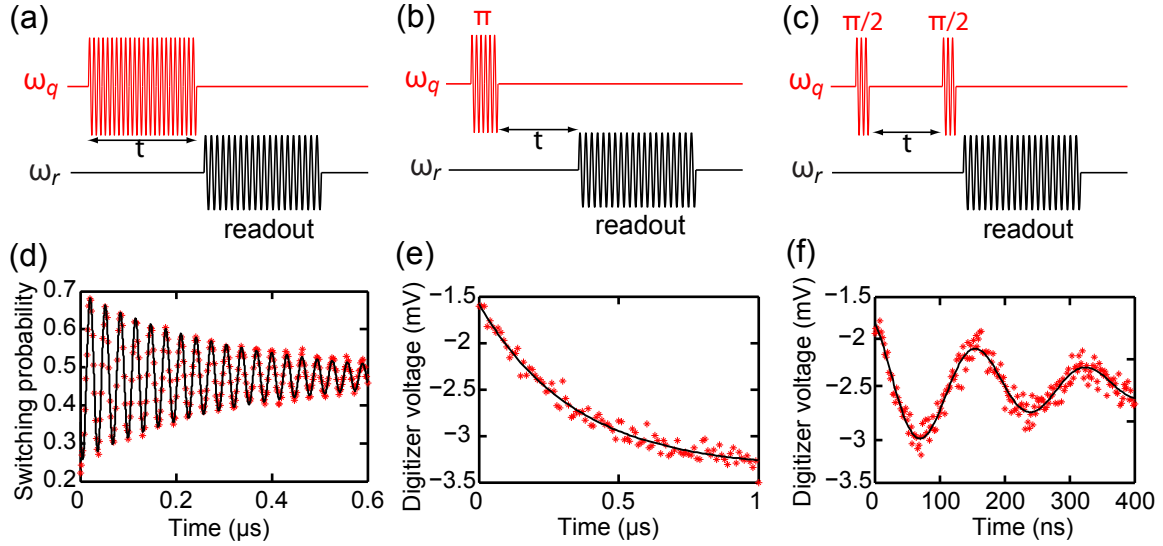


Figure 6.8. Qubit time domain characterizations. (a) Pulse sequence for Rabi oscillations. (b) Pulse sequence for determination of T_1 . (c) Pulse sequence to measure Ramsey fringes and T_2^* . (d) Rabi oscillations. Varying the duration of the qubit manipulation pulse before readout causes the qubit state to oscillate between $|0\rangle$ and $|1\rangle$ on the Bloch sphere. The decay of the oscillations is due to relaxation and decoherence. The JBA readout, in which the switching probability of the resonator from the low to high amplitude state is measured, was used to perform this characterization. (e) Relaxation time T_1 . The average response of the resonator vs delay time τ is a decaying exponential, with time constant T_1 . (f) Ramsey fringes. A detuning $\Delta\omega_q$ between the qubit frequency ω_q and the qubit drive produces decaying oscillations with an oscillation frequency $\Delta\omega_q$. The envelope decay constant T_2^* is the ensemble average decoherence time of the qubit.

Figure 6.8(a) shows the pulse sequence used to measure Rabi oscillations in the flux qubit. The qubit is driven with a resonant pulse of variable length τ . The readout immediately follows, and the average resonator response of many readouts is plotted vs τ as in Figure 6.8(d). The curve exhibits the expected sinusoidal form, but is exponentially damped and decays to an incoherent mixture of the ground and excited states. The form of the decaying Rabi oscillations can be fit to [103]

$$P_{\text{Rabi}} = Ae^{-\tau/T_{\text{Rabi}}} \sin(\Omega_r t) + B \quad (6.6)$$

to extract the Rabi frequency and the decay time T_{Rabi} of the Rabi envelope; A is the maximum amplitude of the oscillations and B is the saturation level of the signal for long τ . Rabi oscillation measurements are useful from a practical point of view as a calibration of the duration and power of a pulse required to rotate the state vector through a given angle on the Bloch sphere. A 180° rotation, which prepares the excited state from the ground state, is known as a ‘ π pulse’. A ‘ $\pi/2$ pulse’ results in a 90° rotation, and can be used to tip the state vector from the ground state up into the equatorial plane.

Table 6.2. Characteristic times in the qubit. The characteristic times are shown at the specified qubit frequency. The times for the qubit in the SQUID resonator readout are quoted at the degeneracy point, while the cQED times are given for a bias away from the degeneracy point. Because the Rabi decay times are dependent on the qubit drive amplitude, only an approximate time is quoted.

SQUID resonator (F063010c)			cQED (F032812b)		
Parameter	Symbol	Value	Parameter	Symbol	Value
Qubit frequency	$\omega_q/2\pi$	8.507 GHz	Qubit frequency	$\omega_q/2\pi$	7.80 GHz
Relaxation time	T_1	320 ns	Relaxation time	T_1	1.8 μ s
Decoherence time	T_2^*	250 ns	Decoherence time	T_2^*	55 ns
Rabi decay time	T_{Rabi}	~ 280 ns	Rabi decay time	T_{Rabi}	~ 1 μ s

The Rabi decay rate $\Gamma_{\text{Rabi}} = 1/T_{\text{Rabi}}$ is a combination of other characteristic decay rates in the system. With the system on resonance with $\Delta\omega_q = 0$, Γ_{Rabi} can be expressed as [108]

$$\Gamma_{\text{Rabi}} = \frac{3}{4}\Gamma_1 + \frac{1}{2}\Gamma_\nu, \quad (6.7)$$

where Γ_1 is the longitudinal decay rate (energy relaxation) and $\Gamma_\nu = \pi S_{\delta\omega_q}(\Omega_r)$ is the spectral density of qubit frequency (ω_q) fluctuations at the Rabi frequency (Ω_r). Thus, the decay in the Rabi envelope is sensitive to noise at both ω_q and Ω_r . If the oscillations are driven off-resonance, the contributions of Γ_ν and Γ_1 to the overall decay rate Γ_{Rabi} decrease. A new dephasing rate term $\Gamma_\phi = 1/\tau_\phi$, however, appears in the expression for Γ_{Rabi} . This rate defines a time constant for randomization of the phase η as seen in Equation 6.4. The characteristic decay times for the qubits discussed in this thesis are tabulated in Table 6.2.

With a π pulse calibrated from Rabi oscillations, we can prepare the qubit in the excited state and measure its relaxation time T_1 . The qubit is excited with a π pulse and the readout is performed after a variable delay τ as shown in Figure 6.8(b). During this window τ , the qubit can decay stochastically to the ground state with a probability set by Γ_1 . The measurement process is repeated many times to build up an average readout response vs τ as shown in Figure 6.8(e). We extract T_1 by fitting this response to the decaying exponential

$$P_{T_1} = Ae^{-\tau/T_1} + B, \quad (6.8)$$

where A and B are also parameters of the fit.

The decoherence time T_2 of the qubit in a single measurement is defined through the Bloch-Redfield equation [109]

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{\tau_\phi}, \quad (6.9)$$

where τ_ϕ , otherwise known as the ‘pure dephasing’ time, is a measure of the phase randomization on the time scale of a single qubit manipulation and readout. Noise at much

lower frequencies (i.e. $1/f$ noise) can also be significant and of much higher spectral power [110, 111]. The effects of this noise are not manifest within a single measurement, but instead cause the qubit frequency to slowly drift between multiple measurements. This dephasing is analogous to ‘inhomogeneous broadening’ in the vernacular of NMR systems, and is an ensemble rather than single measurement phenomenon. The decoherence time with this ensemble dephasing included is defined as T_2^* . This time constant is probed using the Ramsey pulse sequence [112] shown in Figure 6.8(c). Two $\pi/2$ pulses, separated by a variable delay τ , are applied to the qubit immediately before readout. The first $\pi/2$ pulse tips the state vector up into the equatorial plane, where the phase evolves at ω_q for time τ . In the absence of fluctuations in ω_q , the second $\pi/2$ pulse rotates the qubit up into the excited state. Any noise in the precession frequency, however, causes the phase of the precession to become randomized with respect to the qubit drive phase. Thus, the direction of the state vector rotation by the second $\pi/2$ pulse becomes randomized. The qubit signal vs τ decays exponentially with a characteristic time T_2^* to an incoherent mixture of $|0\rangle$ and $|1\rangle$.

Ramsey fringes, as seen in Figure 6.8(f) from the averaged signal of many qubit readouts, are formed when the qubit drive is detuned from ω_q . Viewed in a frame rotating at ω_q , the qubit state vector precesses at the detuning frequency $\Delta\omega_q$ during the window τ . Neglecting noise, the readout signal following the second $\pi/2$ pulse then oscillates between $|0\rangle$ and $|1\rangle$. The effects of dephasing again damp these oscillations with the same characteristic time T_2^* , re-producing the functional form of Equation 6.6. Performing a Ramsey sequence with the the drive detuned is useful for ensuring that the $\pi/2$ pulse is properly calibrated². Moreover, Ramsey fringes allow for a very precise determination of ω_q at a particular flux bias. The oscillation frequency of the fringes is equal to $\Delta\omega_q$, thus providing a reference for determining how far the drive frequency is detuned from the true qubit resonance.

6.5 Amplifier performance

6.5.1 System noise measurements

Accurately measuring the system noise of a low-noise cryogenic amplification chain is challenging and requires precise frequency-dependent calibrations. As referred to in Section 5.4.2, we use the ‘hot/cold’ load method to determine the system noise temperature initially with the superconducting amplifier switched out of the circuit. We first identify the reference point to which the system noise will be referred. This point is immediately after the solid-state Hittite switch (Section 5.4.2), which toggles between two calibrated noise sources. For system noise determination, this point is approximately equivalent to the output of the readout resonator.

²The signal in a Ramsey measurement for $\tau \gg T_2^*$ should saturate to an incoherent excited/ground state mixture. If it does not, this is likely a symptom of an improperly calibrated $\pi/2$ pulse. Adding an oscillating component to the decaying signal allows one to see the full range (excited to ground state) of the signal to verify that it saturates to the correct level.

The hot load is provided by a 50- Ω termination anchored on the 4 K stage of the fridge. In reality, the temperature of this stage at times exceeded 6 K before a full replacement of the pulse tube cooler midway through this work. The 50- Ω load generates Nyquist noise at T_H , which travels with minimal attenuation down the niobium coaxial line to the switch. The switch has a frequency-dependent attenuation A_S , which reduces the noise to T_H/A_S . Another 50- Ω load is similarly anchored to the mixing chamber, and generates noise T_C . Because $\hbar\omega > k_B T$ at the mixing chamber for our readout frequencies, T_C is in the quantum regime and is given by [67]

$$T_C = \frac{\hbar\omega}{2k_B} \coth\left(\frac{\hbar\omega}{2k_B T_{\text{phys}}}\right), \quad (6.10)$$

where T_{phys} is the physical temperature of the mixing chamber. The coth term tends to unity as $T_{\text{phys}} \rightarrow 0$, leaving only the equivalent of a half-photon of energy. The attenuation in the switch does not alter T_C because it is already set by the zero-point quantum fluctuations.

Everything following the reference point in the output amplification chain can be treated as a single amplifier with gain G_S and noise temperature T_{sys} . We can write equations for the total output noise of the chain for each noise source as

$$P_H = k_B G_S \left(\frac{T_H}{A_S} + T_{\text{sys}} \right), \quad (6.11)$$

$$P_C = k_B G_S (T_C + T_{\text{sys}}), \quad (6.12)$$

where P_H and P_C are the total output noise powers per unit bandwidth. Here we have neglected the noise added by the attenuation A_S , which is a small correction when $T_H/A_S \gg \hbar\omega/2k_B$. Dividing Equation 6.11 by Equation 6.12 gives

$$Y = \frac{T_H/A_S + T_{\text{sys}}}{T_C + T_{\text{sys}}}, \quad (6.13)$$

where the ‘Y-factor’ is defined as $Y = P_H/P_C$. Solving Equation 6.13 for T_{sys} , we arrive at

$$T_{\text{sys}} = \frac{T_H/A_S - Y T_C}{Y - 1}. \quad (6.14)$$

This expression is notably independent of G_S , which is eliminated by this two-temperature point technique. The absolute output noise levels are also replaced by their ratio.

To measure Y , we rapidly switch back and forth between the hot and cold loads, measuring the power from the resonator output chain using a spectrum analyzer. The ‘chopping’ between the hot and cold loads allows the power ratio to be measured on a time scale much shorter than the low frequency gain drifts observed in our amplifiers. Along with the measured values of A_S and the temperatures of the dilution refrigerator stages, we can compute T_{sys} for both the SQUID resonator and cQED readout chains as seen in Table 6.3. The difference in the two system noise temperatures is mostly due to the difference in the HEMT noise temperatures at each operating point. These measured noise temperatures are quoted

Table 6.3. Amplification chain noise temperatures.

Readout architecture	Readout frequency ($\omega_r/2\pi$)	Noise temperature
SQUID resonator	1.294 GHz	18 K
Circuit QED	5.780 GHz	7 K

with the superconducting amplifiers switched *out* of their respective readout chains. This is necessary because the noise power from the hot load can saturate the amplifiers, reducing their gain and noise performance. The performance of each amplifier and their effects on the system noise performance will now be discussed.

6.5.2 MSA

The best MSA operating point is determined by performing a two-dimensional sweep through flux and current bias space. At each bias, we measure the total system gain with the network analyzer. This gain is referred to a baseline trace taken with the MSA switched out of the readout circuit. The total output noise, also normalized to the background noise without the MSA operating, is measured with the spectrum analyzer. With the MSA switched in and operating, the increase in the signal (gain) divided by the increase in the noise gives the net increase in the system SNR. The automated scan of current and flux bias is performed to find the point of maximum SNR increase. This occurs at the peak of the MSA resonance, which also moves somewhat in frequency as a function of the current and flux bias. Once the automated scan is complete, we replace the active Keithley 2400 current sources with manually-controlled, battery-powered supplies. Figure 6.9 shows the gain, increase in output noise, and SNR increase in the system at the optimum bias parameters. On resonance at $f_r = 1.294$ GHz, the MSA achieves 27 dB gain with a full-width half-maximum of about 10 MHz. The increase in the system noise is 17 dB, producing a total SNR increase of 10 dB. To a very good approximation, the overall system noise temperature is reduced by this factor, taking it from 18 K down to about 1.8 K. This approximation is valid in the limit of $G_{\text{MSA}}T_{\text{MSA}} \gg T_{\text{sys}}$, where G_{MSA} is the MSA gain, T_{MSA} is the MSA noise temperature, and T_{sys} is the system noise temperature without the MSA. Without making any approximations, we calculate the noise temperature of the MSA itself to be $T_{\text{MSA}} = 1.7$ K. Even though this is significantly higher than the quantum limit, it still produces a dramatic increase in the system SNR. It is also worth noting that the 10-MHz MSA bandwidth is much lower than the 140-MHz bandwidth of the SQUID resonator, and thus sets the speed of the qubit readout.

Because the SQUID flux-to-voltage transfer function is sinusoidal, and only approximately linear for a narrow range of input flux near $\Phi_x = \Phi_0/4$, its dynamic range is intrinsically limited. In Figure 6.10 we plot the gain of the MSA near its optimum operating point as a function of the input power in dBm. There is little response vs power until about




Figure 6.9. MSA performance. The transmission data are normalized to a configuration with the MSA switched out of the measurement chain. At 1.294 GHz the MSA provides 27-dB gain (blue dots) with close to 10-MHz bandwidth. The increase in the output noise power (red dashes) is 17 dB, yielding an overall improvement in power SNR (increase in ratio of gain to output noise) of 10 dB (solid green line).

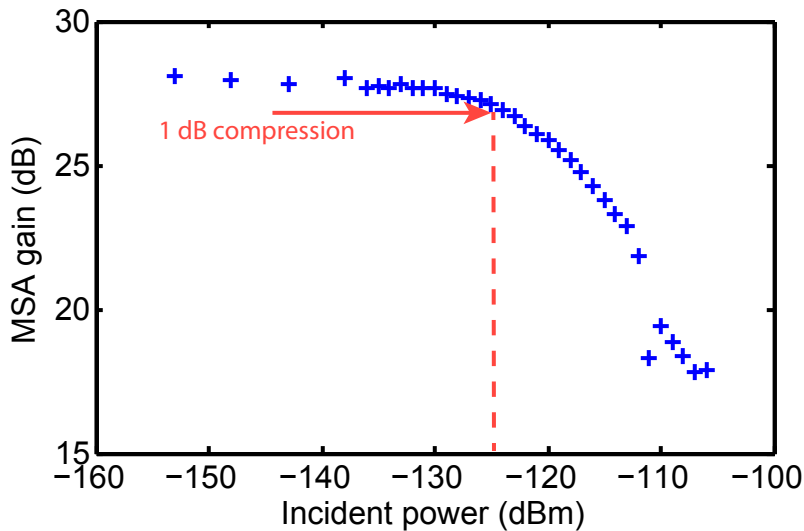


Figure 6.10. MSA dynamic range. The gain of the MSA at its optimum bias decreases once the input power level exceeds a certain threshold. The gain is reduced by 1 dB at an input power of -125 dBm, and continues to decrease with increasing input power.

-125 dBm, where the gain peak is compressed by 1 dB. This corresponds to a mean cavity photon occupation \bar{n} of about 0.6 in the SQUID resonator. The MSA compresses further with increasing power. Although this saturation reduces the SNR increase of the system, the amplifier still provides a net increase in the system SNR.

The 1-dB compression point at the optimal operating point is lower in power than that measured in previous devices. We attribute this to the abnormally high quality factor ($Q \sim 140$) of this particular MSA at its optimal bias. Low input powers correspond to large currents in the resonator for such high quality factors, generating sufficient flux amplitude in the SQUID to exceed the linear portion of the flux-to-voltage transfer function. This high Q is also sometimes symptomatic of self-oscillations and unstable feedback in the amplifier, which we observed in multiple samples. Very high gain with high Q is accompanied by a large increase in excess noise generated by the amplifier. This effect is strongly dependent on the amplifier bias. The considerations for stability in the MSA are discussed more rigorously in Reference [77]. Automated scans of the amplifier performance vs flux and current biases assist in identifying optimal biases while avoiding unstable regions.

Even with multiple layers of magnetic shielding, we still observed slow drifts in the MSA gain caused by slow drifts in the magnetic flux threading the SQUID. This presents a problem for measurements which rely on stable system characteristics over the entire duration of the measurement. We addressed this issue by implementing a feedback scheme which periodically checked and corrected for any small changes in the amplifier flux bias. With the qubit measurement temporarily interrupted, the system gain vs frequency is obtained with the network analyzer. The resonant Lorentzian gain trace is compared against a similar reference trace which is first acquired at the optimal bias before initiating qubit measurements. The discrepancy is determined by calculating the mean square error between measurements of the gain vs frequency. We then iteratively adjust the flux bias until the mean square error is below a user-defined threshold.

6.5.3 Paramp

The flux bias and pump power of the paramp are manually tuned using the network analyzer. The frequency of the pump is identical to the readout frequency $f_r = 5.78$ GHz so that the amplifier is operated in phase-sensitive mode. The resulting amplifier gain vs frequency is shown in Figure 6.11. Though the paramp is capable of higher gain by adjusting the bias parameters, we typically operate it with a gain $G_{\text{par}} = 20$ dB. This provides sufficient gain to dominate the noise of the following amplifier (Low Noise Factory HEMT), while also preserving as much bandwidth as possible. The full-width at half-maximum of the paramp gain peak is 14 MHz. Because the amplifier is operated in phase-sensitive mode, however, the full 14 MHz is not available for tracking changes in the input signal. The degeneracy between the pump and the input signal effectively reduces the usable bandwidth of the amplifier to the 7 MHz half-width at half-maximum.

The system noise with the paramp included in the measurement was not explicitly measured in this work. The amplifier had been used in a previous experiment and the noise was




Figure 6.11. Paramp gain. The paramp was operated in phase-sensitive mode with 20 dB of power gain and 7 MHz (half-width at half-maximum) of usable bandwidth at the readout frequency of 5.780 GHz. The spike at the top of the resonant peak is an artefact of the strong pump tone applied to the amplifier.

fully characterized in Reference [79]. The total system noise with the paramp operating in phase-preserving mode was shown to be within a factor of two of the quantum limit.

Chapter 7

Fidelity and backaction measurements

Coherent manipulation of a qubit with a long coherence time is a necessary but not sufficient requirement for practical quantum computation. The quantum information from the qubit must be mapped using a high fidelity readout onto a detector at room temperature. Even if the detector is quasi-continuous, a criterion for distinguishing the ground state from the excited state must be developed. This is accomplished by defining a threshold in the detector in which every readout value on one side of the threshold is a ground state while values on the other side correspond to the excited state. Ideally, the qubit state $|0\rangle$ perfectly maps to the ground state side and $|1\rangle$ to the excited state side.

Errors in the qubit-to-readout mapping originate from two sources. The first source is due to the projective nature of the readout. To understand these errors, we first identify readout pointer states as the discrete states in the readout element which directly correspond to the projected states of the qubit. In the case of the SQUID resonator readout, which is sensitive to the magnetization states of the qubit, these are $\tilde{\zeta}$ and $\tilde{\zeta}'$. The tilde differentiates these resonator states from the quantum states of the qubit. A higher degree of circulating state mixing in the qubit eigenstates (see Section 2.5) leads to increasing readout errors because each of the qubit eigenstates corresponds to a mixture of both $\tilde{\zeta}$ and $\tilde{\zeta}'$. The readout pointer states in the cQED architecture, which measures along the same axis as the qubit eigenstates, are $\tilde{0}$ and $\tilde{1}$. This eliminates projection errors, as $|0\rangle$ maps perfectly to $\tilde{0}$ and $|1\rangle$ to $\tilde{1}$ [113].

The second error in the qubit-to-readout mapping is due to the finite SNR of the readout. Even if the qubit states map perfectly to the distinct readout pointer states, noise can cause measurements of a single readout pointer state to lie on both sides of the discrimination line. In this case, the distributions formed by many individual readouts overlap (Figure 3.4), implying ambiguity in the measurement result.

The overall measurement fidelity F , which ranges from zero to one, describes how well the qubit can be prepared in the ground or excited state and then measured to be in that state. It is defined as

$$F = 1 - P_{0|1} - P_{1|0}, \quad (7.1)$$

where $P_{0|1}$ ($P_{1|0}$) is the fraction of erroneous ground (excited) states on the excited (ground)

state side of the detector discrimination threshold. In addition to the aforementioned readout imperfections, any non-idealities in the qubit itself (finite coherence time, state preparation errors, etc.) also degrade F .

As seen in Equation 7.1, a fidelity below unity implies the existence of errors in a digital system. These errors are present in classical, commercial computing platforms but are corrected for by a process known as error correction. The analogous procedure in a system of qubits is referred to as quantum error correction [6, 114, 115]. Limitations in qubit coupling schemes, coherence times, and measurement fidelity have until recently restricted quantum error correction to a purely theoretical realm. The predicted fidelity threshold required to enable robust, practically useful quantum error correction is not currently well-defined in the community, but is generally accepted to be above 99% [116, 117, 118].

Measurement fidelities exceeding 90% have been demonstrated in superconducting qubits, but the simultaneous realization of a fast, high-fidelity, QND readout has thus far been difficult, with many schemes exhibiting either intrinsic projection errors [119], long measurement times [84, 85], or demolition of the quantum state [33]. Traditionally, insufficient measurement sensitivity has limited the fidelity with which the readout pointer states can be resolved. Increasing measurement strength improves one's ability to distinguish the readout pointer states, but takes the measurement further from the ideal QND readout by introducing additional readout backaction. In this chapter, we study the impact of near-quantum-limited amplifiers on a dispersive qubit measurement. Specifically, we use the amplifiers as a tool to optimize the readout and exploit the high SNR to enable new techniques such as active qubit state initialization. The remainder of the chapter will be divided into two sections, one dedicated to each of the SQUID resonator and cQED experiments.

7.1 SQUID resonator

This section is based on our manuscript “Dispersive readout of a flux qubit at the single photon level,” which is cited as Reference [120].

7.1.1 Pulse sequence

The measurements presented in this section were performed using the pulse sequence shown in Figure 7.1. A 20-ns π pulse was first calibrated using Rabi oscillations at the degeneracy point. This allows for state manipulation at the bias where the fidelity of the π pulse and the qubit frequency are minimally sensitive to ambient fluctuations. The fast flux shift immediately follows the control pulse, shifting the qubit bias and increasing the magnetization signal between the qubit eigenstates. After the qubit flux has been shifted, the readout resonator is energized for 500 ns and the reflected signal is acquired at room temperature in 10-ns increments.

Ideally, the magnitude of the fast flux shift is sufficient to produce eigenstates which are nearly pure circulating current states. In practice, however, the shift magnitude is limited




Figure 7.1. Pulse sequence for SQUID resonator fidelity measurements.

by nonidealities in the qubit. First, the relaxation time of our flux qubits is quite sensitive to the operating frequency ω_q . A long T_1 at ω_q is desired so that, following the flux shift, the readout signal can be averaged for a long time before the excited state decays. Additionally, the fast flux shift sweeps the qubit through two-level defect states [121] and parasitic modes in the surrounding electromagnetic environment. Evidence of these imperfections can be seen in the qubit spectroscopy of Figure 6.4. The strong flux-independent resonance at ~ 9.35 GHz is particularly notable. As the qubit frequency is swept through these features, interactions occur and energy from the qubit can be transferred to the environment, leading to a loss of excited state population. The more defect states the qubit is swept through, the higher the probability that a decay event occurs.

Maximizing the measured fidelity involves finding an optimal shift magnitude which balances the increasing magnetization signal with the population loss during the fast flux shift. We perform this optimization empirically by measuring F as a function of the shift amplitude. The shift amplitude corresponding to the maximum fidelity in this experiment was approximately $4 \text{ m}\Phi_0$ in the qubit. This flux represents a qubit bias in which the eigenstates are still non-negligible superpositions of circulating current states (approximately 3:1 ratio). This severely limits the maximum possible measurement fidelity, but increasing the flux shift leads to excessive excited state population loss. Engineering for a qubit with less defect states would allow for a large shift magnitude and higher possible fidelities. The SiN_x dielectric which lies underneath the qubit and resonator is a likely source of spurious resonances.

7.1.2 Signal processing

We begin characterizing the readout by analyzing the average response of the readout resonator for many ground and excited state preparations. Each 10-ns acquisition of the raw IQ signal over the 500-ns readout duration is converted in software to an amplitude and phase. The readout power is expressed in terms of the average photon occupation \bar{n} in the resonator. This calibration is thoroughly detailed in Reference [79], and involves measuring the incident power at the resonator and the quality factor ($Q = 10$) of the readout resonator.




Figure 7.2. T_1 decay during measurement. The average phase difference between ground and excited state readouts decays exponentially to zero during the readout. This is shown for $\bar{n} \approx 3$ in the resonator. Fitting the decaying response gives the maximum possible phase difference at the start of readout ($t = 0$) and the T_1 decay time constant.

In Figure 7.2, we plot the average phase difference between many ground and excited state readouts with $\bar{n} \approx 3$. The time axis represents the duration into the 500-ns readout of each 10-ns acquisition point. Thus, the plot shows how the average phase difference exponentially decays *during* the readout. This is notably different than the T_1 measurements of Chapter 6, which characterize the decay of the qubit in the *absence* of photons in the readout resonator. The effect of these readout photons on the qubit relaxation rate will be discussed later in this section.

The phase signal in Figure 7.2 decays exponentially over most of the readout time with the exception of the first ~ 50 ns. This part of the curve is an artefact of the rise time of the MSA, which sets the limit on how fast the system can track changes in the amplitude and phase of the signal. If we ignore the part of the curve where the readout signal is rising in the MSA and fit the remaining curve from $t = 60$ ns to 500 ns, we obtain the exponential time constant ‘ T_1 during readout’ and the maximum initial phase difference at $t = 0$. The latter piece of information will later be useful in cataloging the sources of fidelity loss in the system.

With a sense for the average readout signal, we can develop a protocol for maximizing F . At the beginning of the chapter, we defined F in terms of a qubit-state-dependent resonator signal, but did not elaborate on how to determine this signal. The simplest method is to make distributions (many ground and excited state readouts) of single-point 10-ns digitizations at many time points (Figure 7.2) during the readout. The discrimination line

which maximizes F according to Equation 7.1 is then empirically determined for each time point. The maximum fidelity time point occurs near the beginning of readout, immediately after the rise time of the measurement but before additional excited state relaxation events.

Calculating the fidelity from only a single 10-ns digitization requires minimal post-processing of the data, but does not optimally utilize the entire 500-ns readout during which additional information about the qubit state can be obtained. Multiple 10-ns digitizations during a single readout can be averaged together, reducing the effective measurement bandwidth and the effects of noise in the measurement chain. The effectiveness of this averaging is limited by the finite lifetime of the excited state signal before it stochastically decays to the ground state. Averaging is beneficial for unambiguously distinguishing an excited state readout until the qubit decays, after which an erroneous ground state component is also averaged into the total record. Determining the optimal duration of the signal in order to average the data for maximum fidelity is a non-trivial function of T_1 and the measurement SNR [122], and was thus empirically determined in this work.

The fidelity can be further increased by applying a non-uniform weighting function to different portions of the readout signal during the averaging process. Points near the beginning of a single readout—before any relaxation events—are statistically more likely to contain information about the initial state of the qubit than points later in the readout. Various weight schemes, known as filter functions, have been proposed in the literature [122]. We chose to employ the relatively simple ‘exponential filter’, in which each individual readout is averaged using a decaying exponential weighting function. The exponential time constant of this filter was chosen to match T_1 during readout for each readout power. This filter time constant is not necessarily the optimal, but it is an intuitive place to start.

7.1.3 Readout improvement with MSA

To quantify the effect of the MSA on our flux qubit readout, we measure the fidelity F as a function of readout power with the MSA both switched into and out of the measurement chain. The two measurement schemes are identical (except for the presence of the MSA), including the measurement post-processing with the same exponential filter. We focus our measurements on very low power readout, where the overall fidelity F is also low. The system is capable of achieving higher fidelity ($F = 66\%$ with JBA readout), but we choose to study the linear regime in the resonator where the effect of the amplifier on the readout is most striking and the measurement backaction is minimal.

The overall measurement fidelity is low mostly on account of nonidealities in the qubit. From Figure 7.2, we determine the fidelity loss due to the insufficient magnitude of the fast flux shift and population loss during the shift. Immediately after the rise of the fast flux pulse as the readout is turning on ($t = 0$), the average phase difference between ground and excited state readouts reveals the maximum measurable phase contrast. This extrapolated phase difference is $\phi_0 = 8.5^\circ$, as seen in Figure 7.2. From Section 6.2, we determined the maximum phase contrast between the circulating states to be $\phi_{\max} = 19^\circ$. It can be shown

Table 7.1. Fidelity loss in SQUID resonator/MSA readout at $\bar{n} = 1.5$. The measured fidelity is 27.7%. The right column accounts for the cumulative fidelity loss at and above a given row.

Source of loss	Fidelity loss	Maximum measurable fidelity
Flux shift & decay during shift	55%	45%
T_1 decay	15%	38%
SNR	28 %	27.7%

that the maximum possible fidelity F_0 is given by the simple ratio

$$F_0 = \frac{\phi_0}{\phi_{\max}}, \quad (7.2)$$

which for our values is $F_0 = 45\%$. We note that in principle ϕ_0 is a function of readout amplitude due to the readout-induced shift in the qubit flux. For very low power readout, however, these shifts produce negligible changes in ϕ_0 . In addition to being experimentally observed, this was verified numerically using a full numerical model of the SQUID resonator - qubit system. The readout-induced shift only becomes significant at readout powers where the nonlinearity in the resonator is manifested strongly.

We tabulate the fidelity loss due to the fast flux shift and the other dominant sources of loss in Table 7.1 with $\bar{n} = 1.5$. The limited bandwidth of the MSA leads to a 50-ns delay—during which T_1 decay events occur—between the start of readout and the time at which the average signal between the ground- and excited-state readouts peaks (Figure 7.2). This loss is included in Table 7.1 as ‘ T_1 decay,’ and is equal to $\exp(-t_r/T_1)$, where t_r is the time between the start of readout and the maximum observed signal between the qubit states.

After accounting for the circulating current admixture, population loss during the fast flux shift, and T_1 decay, the remaining fidelity loss is attributed to the finite system SNR. This further reduces the fidelity to the measured value $F = 27.7\%$ at $\bar{n} = 1.5$. Unlike the other terms, this source of fidelity loss is directly impacted by the presence of the MSA. To isolate the performance of the readout, thus providing a metric for quantifying the effects of the MSA, we define the visibility V as the fidelity F normalized to the maximum fidelity determined by all sources of fidelity loss except finite SNR, as shown in Table 7.1. Thus, the visibility is not a direct function of the qubit operating point.

The visibility as a function of \bar{n} with the MSA switched in and switched out of the circuit is shown in Figure 7.3. This quantity is determined from the measurement histograms obtained from 10^5 ground- and excited-state preparations. For low-amplitude excitation, as expected, the visibility is very low, and the qubit states are virtually indistinguishable. As the number of photons is increased, the SNR and thus the visibility increase. The MSA provides a substantial increase in the visibility at all readout powers. At 1.5 photons, the visibility with the MSA reaches $72.3 \pm 2.1\%$ ($27.7 \pm 0.7\%$ fidelity), a factor of ~ 4.5 greater




Figure 7.3. Visibility vs readout photons in the SQUID resonator/MSA readout. The visibility is the measured fidelity normalized to a maximum which accounts for nonidealities in the qubit, thus isolating the performance of the readout. The visibility peaks at 72.3% with the MSA amplifying the readout signal. This represents a ~ 4.5 fold increase in the visibility when compared to a readout without the MSA.

than without the MSA. Thus, a low power measurement which provides little information about the qubit state can be improved with the MSA so that the majority of the qubit signal is recovered. At higher photon numbers, the visibility decreases as the nonlinearity in the SQUID resonator causes the resonance to shift to lower frequencies, reducing the phase contrast between the circulating current states. Overall, the measured visibility is restricted by the relatively small $\phi_0 = 19^\circ$ phase shift between the two circulating states, a consequence of the low quality factor of the SQUID resonator.

We independently confirm that the measured visibility and corresponding fidelity loss due to the SNR at $\bar{n} = 1.5$ are consistent with the measured system noise, bandwidth, and resonator phase difference ϕ_0 between $\tilde{\zeta}$ and $\tilde{\zeta}$. We first determine the system SNR from these quantities. The signal is defined as

$$S_V = \sqrt{Z_0 P_{\text{out}}} \sin\left(\frac{\phi_0}{2}\right), \quad (7.3)$$

where $Z_0 = 50 \Omega$ and P_{out} is the power reflected from the SQUID resonator. The signal S_V is in the form of a voltage, and is referred to a position in the measurement chain immediately following the SQUID resonator. The sine term accounts for the phase shift between the circulating current states, as seen from Equation 3.16. We define the measurement noise N_V as

$$N_V = \sqrt{k_B T_{\text{sys}} Z_0 B}, \quad (7.4)$$




Figure 7.4. Calculated measurement distributions. (a) Calculated distributions without filtering and integration. The signal is a single 10-ns acquisition from the readout record. The distributions are broad and show significant overlap. (b) Calculated distributions with filtering and integration. The readout signal is integrated after being weighted with a decaying exponential filter. The readout pointer state peaks narrow, but decay events are also averaged into the excited state distribution. The optimal filtering and integration scheme balances these two effects.

where T_{sys} is the system noise temperature and B is the measurement bandwidth. This bandwidth is approximately that of the MSA because $T_{\text{MSA}}G_{\text{MSA}} \gg T_{\text{HEMT}}$. In a more rigorous treatment, the power spectral density of the noise after the MSA would be integrated over frequency to give the total N_V . Equations 7.3 and 7.4 are combined to produce the measurement SNR

$$\frac{S_V}{N_V} = \sqrt{\frac{P_{\text{out}}}{k_B T_{\text{sys}} B}} \sin\left(\frac{\phi_0}{2}\right), \quad (7.5)$$

which is about 0.6 at $\bar{n} = 1.5$ when also taking into account compression of the MSA peak due to its limited dynamic range.

In our visibility calculation, we simulate the response of 10^6 ground- and excited-state readouts using this measured SNR. Individual measurement records are created for both an excited-state and a ground-state readout. The average of these readout traces is separated by $2S_V$, while random Gaussian noise of standard deviation N_V is added to each readout. Each excited state in the simulation randomly decays to the ground state on a time scale consistent with the measured T_1 . Each of 10^6 ground- and excited-state simulated readouts is integrated using the same exponential filter which is applied to the measured data. Measurement histograms are created from the aggregate of all the simulated readouts, and the visibility is calculated using Equation 7.1 and the optimum discrimination line between the two histograms. This simulation predicts a visibility of 76% at $\bar{n} = 1.5$, in very reasonable agreement with the measured value of 72.3%.

Figure 7.4 shows an example of calculated measurement distributions for 10^6 ground- and excited-state preparations. A single 10-ns acquisition [Figure 7.4(a)] produces distributions which are symmetric and broad. The overlap between the two distributions is decreased,




Figure 7.5. T_1 during readout vs readout photons. The relaxation time T_1 during measurement is a function of the readout strength \bar{n} . Below $\bar{n} = 0.1$, T_1 saturates to 320 ns, the value measured in the absence of readout photons. As \bar{n} is increased, T_1 decreases due to increasing readout backaction.

thus increasing the visibility, by means of filtering and integration [Figure 7.4(b)]. The ground-state peak sharpens with higher amplitude, while the excited state peak also narrows but retains a broader non-Gaussian shape. This is due to the integration process, which averages together ground- and excited-state readout signals when the qubit decays during the measurement record. The optimal integration balances the narrowing of the readout pointer state peaks with minimal ground-state mixing into the excited-state distribution due to T_1 decay.

We also study the effect of measurement strength on the qubit-state evolution by measuring the relaxation time T_1 during continuous readout monitoring (Figure 7.2) versus \bar{n} , as seen in Figure 7.5. For $0.008 < \bar{n} < 0.1$ during readout, T_1 approaches 320 ns, the value measured by applying a π pulse and allowing a variable delay time to elapse before readout (free decay). As the cavity occupation increases, T_1 decreases. We speculate that this reduction in T_1 with increasing readout power arises from sweeping the qubit through increasing numbers of environmental decay modes. As the readout power is increased above $\bar{n} = 4$ (not shown), T_1 during readout plateaus as the frequency of the nonlinear SQUID resonator decreases while its response sharpens. Consequently, the fraction of incident power entering the resonator decreases, and our photon number calibration fails as \bar{n} no longer increases linearly with incident power. If the frequency of the readout is adjusted to match the decreasing readout frequency, T_1 during readout continues to decrease due to an increasing number of photons occupying the resonator.




Figure 7.6. Quantum jumps. Three individual qubit readouts ($\bar{n} = 14.6$) starting at $t = 0$: excited state (solid red line), ground state (long blue dashes), abrupt quantum jump from the excited to ground state (short green dashes). The measurement SNR and speed are sufficiently high to observe individual transitions between the qubit states in real-time. The readout is turned off in the gray region before $t = 0$.

7.2 Circuit QED

This section is based on our manuscript “Heralded state preparation in a superconducting qubit,” which is cited as Reference [123].

7.2.1 Quantum jumps

We perform an initial analysis of our cQED readout performance using a pulse sequence similar to that of Figure 7.1. There is no need for a fast flux shift because the measurement basis coincides with the qubit eigenstate basis. With the qubit biased at 7.80 GHz, where $T_1 = 1.8 \mu\text{s}$, we first apply a control pulse and subsequently digitize the readout signal in 10-ns increments. With the dispersive shift $\chi = g^2/\delta$ known from previous calibrations (Chapter 6), we use the AC Stark shift in the qubit to precisely calibrate the number of photons in the readout resonator. This is performed by fitting a curve of the measured qubit frequency ω_q vs readout power to the functional form of Equation 3.8.

At a power $\bar{n} = 14.6$, we show in Figure 7.6 an example of the digitized readout signal vs elapsed time for three individual readout traces. After the readout turns on at time $t = 0$, the signal quickly settles to one of two distinct levels, which correspond to the readout pointer states $\tilde{0}$ and $\tilde{1}$. The system SNR with the paramp operating is sufficiently high that the readout pointer states appear to be completely separated. The degree to which these states can be unambiguously identified will be addressed later in this section. The qubit is initially




Figure 7.7. Histogram of quantum jumps from excited to ground state. The histogrammed decay events form an exponential curve which quantitatively agrees with the measured T_1 obtained from an averaged, ensemble measurement. The exponential time constant obtained from a fit to the distribution of decay times is plotted with the red curve.

in the excited state for the red and green readouts, while the blue trace shows a ground state readout. At approximately $t = 0.3 \mu\text{s}$, the green curve abruptly makes a transition between the two readout pointer states, corresponding to the qubit decaying from the excited to the ground state. In all previous flux qubit readouts, measurement noise obscured the ability to resolve these transitions, which are known as ‘quantum jumps’ [78], in a single readout. Only the ensemble decay of many readouts averaged together can be obtained if there is insufficient SNR. Additionally, the speed of the paramp ($t_{\text{meas}} \ll T_1$, where t_{meas} is the measurement time) is critical for observing quantum jumps in real time.

To confirm that the jumps in the readout signal are consistent with the measured qubit T_1 , we construct a simple algorithm for identifying the jump down times in readouts with the qubit initially prepared in $|1\rangle$. We define a threshold voltage in the digitized signal and catalog the times at which the individual readouts definitively make a transition from above to below the threshold. Figure 7.7 shows a histogram of these jump times for many excited state preparations. The jump-down times are fit to an exponential distribution, and the T_1 decay time is extracted and plotted in red on top of the histogram. The fit is generally good except for at low values of t , which represent readouts which decay very quickly. We attribute the discrepancy between the fit and distribution at small t to imperfections in the jump-finding technique due principally to the finite measurement bandwidth.




Figure 7.8. Raw measurements distributions. Readout distributions created from single 10-ns acquisitions at $t = 90$ ns into readout are shown for 10^5 ground- and excited-state preparations at $\bar{n} = 14.6$. The high system SNR is evident in the large separation of the main peaks of each distribution. Small secondary peaks which are responsible for fidelity loss are also visible. Each histogram bin count is normalized to the total number of counts.

7.2.2 Fidelity measurements

The observation of quantum jumps, combined with a relatively long T_1 , suggests a high fidelity F . We begin our calculation of the fidelity by first analyzing the readouts of many ground- and excited-state preparations, such as those in Figure 7.6 where $\bar{n} = 14.6$. The readout time which produces the optimal fidelity occurs immediately after the paramp rise time so that excited state decay is minimized. This point occurs at $t = 90$ ns into the readout. In Figure 7.8, we present the distributions of 10-ns readout digitizations at this readout time. The ‘probability’ represents the number of counts at a given digitizer voltage normalized to the total number of counts. The main peaks of each distribution are gaussian and are well-separated from each other, confirming the observations we made using Figure 7.6. Upon closer inspection, the presence of much smaller, secondary peaks are visible in each distribution. For example, in addition to its main peak at 0.35 V, the excited state distribution (red \circ) also contains a small peak at the same voltage (-0.5 V) as the dominant ground state peak. These spurious peaks therefore indicate a degree of contamination in each distribution. The fidelity is calculated in the usual manner (Equation 7.1) by defining a discrimination line which bisects the two main peaks. Thus, the relative size of the secondary peaks directly contributes to a reduction of F . The composition of these peaks and a full analysis of fidelity loss sources will be discussed later in this section.

We perform a thorough study of the readout fidelity by measuring it over four decades of readout power. Using the methods in Section 7.1, we first compute the fidelity using only a 10-ns digitization 90 ns into readout. These data are shown (blue \times) in Figure 7.9(a). The

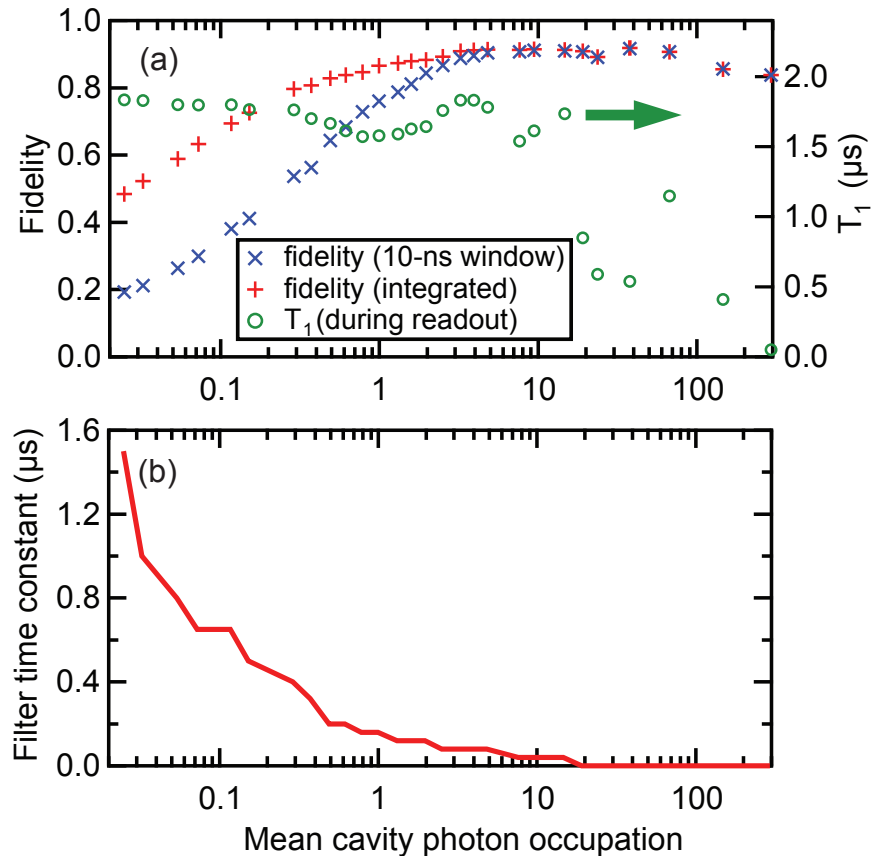


Figure 7.9. Raw fidelity and T_1 during readout vs readout power. (a) Fidelity curves measured using a single 10-ns acquisition (blue \times) and using integration (red $+$) merge together when the readout pointer state distributions completely separate. The relaxation time T_1 measured during continuous readout monitoring (green \circ) is referred to the right axis. (b) Time constant of the exponential filter used to optimize the integrated fidelity.

fidelity is low at low readout powers, but increases as the readout excitation is increased. At $\bar{n} = 37.8$, the fidelity reaches its maximum of $91.1 \pm 0.4\%$. As the power is further increased, the fidelity decreases as a result of increasing readout backaction.

The fidelity is also computed with exponential filtering [122] and integration (red $+$) as in Section 7.1.2. This post-processing of the readout signal makes a dramatic difference at low power, boosting the fidelity from 19% to 47% at $\bar{n} = 0.025$. With the measurement SNR very low, the positive effects of averaging (reduced noise) greatly outweigh the negatives (mixing T_1 decay processes into the excited-state signal). The integrated fidelity again increases with increasing readout strength, until it merges together with the ‘10-ns’ fidelity curve at about $\bar{n} = 7$. This point approximately corresponds to the separation of the 10-ns readout pointer state distributions. With negligible overlap between these distributions, further integration

is no longer beneficial. In fact, integration actually *degrades* the fidelity on account of T_1 relaxation events at high SNR.

In contrast to the MSA readout experiment, in which the time constant of the exponential filter remained fixed, we empirically determine the optimal filter time constant for each readout power. The exponential time constant which maximizes the fidelity for each readout power is shown in Figure 7.9(b). The filter time constant decreases monotonically with increasing readout power, showing explicitly that the optimal time constant is a function of the system SNR. When the time-axis (y) is plotted on a log scale, the result is a quasi-linear trace. We did not investigate a theoretical dependence with which to compare this result, although there is some theoretical treatment of these filters in the literature.

The green circles in Figure 7.9(a) represent the measured T_1 during readout versus readout power. Below $\bar{n} = 14.6$, T_1 during readout always exceeds $1.5 \mu\text{s}$, but as the power is increased, T_1 decreases as a result of increasing readout backaction. We attribute the structure in the decreasing T_1 data points to the presence of environmental defects at various operating frequencies as the qubit frequency is Stark-shifted with increasing readout power.

7.2.3 Readout pointer state discrimination

We begin our analysis of fidelity loss in the system by quantifying the degree of separation of the readout pointer state distributions. This, along with the remaining work in this chapter, was performed at $\bar{n} = 14.6$. This readout power was chosen because it provides a high fidelity while minimizing suppression of T_1 during readout.

Referring back to Figure 7.8, we observe a small number of errant counts, that is, counts at values of the homodyne voltage not centered about the bimodal peaks. We attribute these counts to qubit transitions, induced by noise or T_1 relaxation, during the measurement. To discard these events and create pure-state distributions, we implement a two-point correlation procedure. The system is probed at times t_A and t_B , as shown in Figure 7.6. We retain measurements only when both of these readings return the same value for the qubit state. Thus, the readout exhibiting the quantum jump is excluded. The 160-ns time difference between these points is several times longer than the response time set by the ~ 10 MHz system bandwidth, thus ensuring minimal autocorrelation between the signals at t_A and t_B , and nearly random distributions at t_D . The resulting distributions at t_D with these outliers removed are shown in Figure 7.10 for $\sim 10^5$ ground- and excited-state traces. Using the discrimination threshold shown in the figure, we observe only 108 false counts, thus allowing the pointer state distributions to be separated to an error of one part in 1000. Consequently, the finite measurement SNR does not contribute to the observed fidelity loss. This procedure, however, does not correct for those rare events in which the system jumps twice. The 160-ns window between t_A and t_B is intended to minimize these double-jump readouts while simultaneously keeping the autocorrelation between t_A and t_B to a low value on the order of 7×10^{-3} . Unambiguously distinguishing the double-jump events is difficult due to the limited system bandwidth, which could be increased in future experiments.




Figure 7.10. Pure readout pointer state distributions. A two-point correlation procedure is used to exclude qubit transitions due to the filtered out quantum jumps and create pure ground and excited state distributions, represented by log-linear histograms, with each histogram bin count normalized to the total number of counts. By totaling the counts on the wrong side of the voltage discrimination threshold, the distributions are determined to be separated to an error of less than one part in 1000.

As a final note, when we fit each distribution to a gaussian and analytically calculate the overlap, we find a much lower error than one part in 1000. The low-count ‘tails,’ which contribute the majority of the error counts, do not lie on the gaussian fits shown by the solid lines. Assuming the statistics of the system are near-normally distributed, it is reasonable to postulate that these outliers are not intrinsic to the readout pointer state distributions and are due to other effects, such as qubit state transitions which could not be fully filtered out of the distributions. We were unable to conclusively prove this in our work, and thus claim that the degree of separation was at minimum an error of less than one part in 1000.

7.2.4 Heralded state preparation

Techniques new to the field of superconducting qubits are made possible by this near-perfect discrimination of the qubit state in only 90 ns after the start of the readout pulse. We use the high speed readout to insert an additional measurement pulse before a qubit manipulation and measurement sequence. With this ‘heralding’ readout, we verify that the quantum system is initialized in the ground state. Such heralding techniques are currently employed in other quantum information architectures such as trapped ions [124], photonic systems [125], and quantum dots [126]. With this technique, we effectively eliminate state preparation errors due to the spurious excited state population observed in superconducting qubits [127, 128].

Figure 7.11(a) shows the pulse sequence used to herald the ground state. With the qubit in thermal equilibrium, we energize the readout and extract the qubit state at $t = t_S$

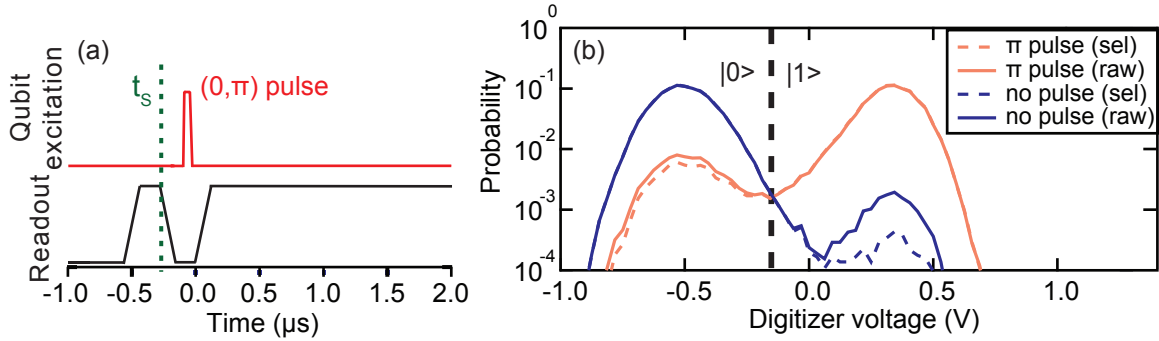


Figure 7.11. Heralded state preparation at $\bar{n} = 14.6$. (a) Pulse sequence used for heralded state preparation. A pure ground state is heralded at $t = t_S$. (b) Log-linear raw readout distributions for $\sim 10^5$ ground- and excited-state readouts (solid lines) compared with the distributions generated after heralded state preparation (dashed lines). The distribution bin counts are normalized to the total number of counts. An increase in fidelity from 91.0% to 93.9% is reflected in the decreasing size of the smaller peak of the bimodal distributions which are responsible for fidelity loss.

(subscript S denotes selection). This time falls at the very end of the heralding pulse to minimize the time between the end of the qubit state determination and the following control and readout pulses. If the qubit is determined to be in the ground state, no correction is applied and qubit readout and manipulation is performed in the typical manner. If the excited state is measured, the subsequent readout is discarded from the total record via post-selection. Events where the qubit is spontaneously found in the excited state can be attributed to either remnant thermal population or readout-induced excitation of the ground state during the heralding pulse.

With this procedure in place, we again prepare ground and excited state distributions. We compare the log-linear readout distributions both with (selected) and without (raw) heralded state preparation in Figure 7.11(b). The selected distributions are obtained from the full three-dimensional readout histograms presented for both ground- and excited-state preparations in Figure 7.12. The reduced overlap of the distributions with heralded state preparation is demonstrated by the decreased size of the smaller of the bimodal peaks in both the excited and ground state distributions. The less-dramatic apparent reduction in the smaller peak of the excited state distribution as compared to the ground state distribution is due to the log scale y-axis and the presence of T_1 decay events (Section 7.2.5). This heralding procedure corrects for the remnant thermal population of the excited state, corresponding a fidelity improvement of $2.9 \pm 0.2\%$. This boosts the measured fidelity at $\bar{n} = 14.6$ to $F = 93.9\%$.

The fidelity improvement of 2.9% originates from both the excited and ground state distributions. Thus, the equilibrium excited state thermal population is approximately 1.45%. Inverting the Boltzmann factor, we arrive at the expression for the effective qubit tempera-

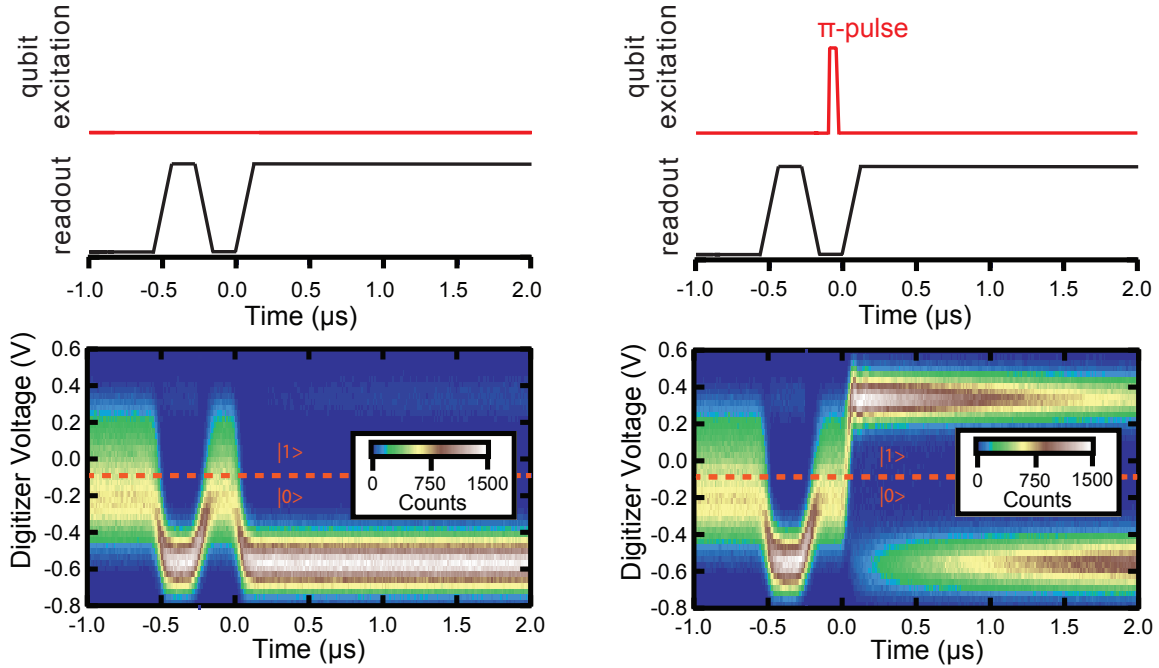


Figure 7.12. Three dimensional readout histograms. Each three dimensional histogram corresponds with the pulse sequence above it. The pulse sequences include a heralding readout followed by an optional π pulse and subsequent readout. The dotted discrimination line shows complete separation between the ground- and excited-state distributions.

ture T_Q as a function of the thermal population of the excited state f_e :

$$T_Q = -\frac{\hbar\omega_q}{k_B \ln(f_e)}. \quad (7.6)$$

This expression yields 88 mK for the operating parameters of our qubit, significantly higher than the ~ 30 mK base temperature of the dilution refrigerator. This anomalously high effective qubit temperature is prevalent in the superconducting qubit community [127, 128], but heralded state preparation using our active measurement technique enables us to avoid this problem. In the future, state preparation errors could be actively corrected in real time if a π pulse were to be triggered based on the outcome of the heralding readout.

7.2.5 Fidelity loss and readout backaction

We exploit the fast, high fidelity readout to make a precise accounting of the remaining sources of fidelity loss. The largest source of loss is due to T_1 decay between the π pulse and the moment the readout resonator is fully energized. Because of strong frequency-dependent variation in T_1 , we compute this contribution by measuring T_1 as a function of

the qubit frequency in discrete steps near 7.80 GHz and smoothly interpolating between points. Assuming a rise in the cavity amplitude of $1 - e^{-2/\kappa t}$, we infer T_1 as a function of time as the qubit frequency is Stark-shifted during the readout rise time. We compute the resulting fidelity loss F_{loss,T_1} by integrating over the decay rate as follows:

$$F_{\text{loss},T_1} = 1 - e^{-\int dt/T_1(t)}. \quad (7.7)$$

This calculation gives $F_{\text{loss},T_1} = 4.4 \pm 0.3\%$.

The readout itself can contribute to fidelity loss by driving transitions in the qubit. This effect is seen in the suppression of T_1 during readout [Figure 7.9(a)] and is directly related to the QND character of the readout [129]. The ‘dressed dephasing’ model [130, 131] provides a theoretical explanation of this effect and has been the subject of recent experimental investigation in Irfan Siddiqi’s laboratory. In summary, dressed dephasing originates from the nonlinearity in the qubit-resonator system which allows mixing between the readout tone and noise at the detuning frequency δ . This process produces extra noise at ω_q which leads to either spurious decay or excitation in the qubit. It can be shown through expanding the Jaynes-Cummings Hamiltonian to the term past the dispersive approximation that the transition rates are [79]

$$\Gamma_{\uparrow/\downarrow} \propto \frac{g^2}{\delta^2} S(\delta) \bar{n}, \quad (7.8)$$

where $S(\delta)$ is the power spectral density of the noise at the detuning frequency and a symmetric noise spectrum is assumed. Thus, the induced transition rate is approximately proportional to the readout excitation power.

The ability to resolve quantum jumps provides the novel ability to measure the measurement-induced transition rates in a flux qubit by examining long time traces ($\gg T_1$) with the system prepared in the ground state. We record and analyze the statistics of individual transitions between the qubit states using a maximum likelihood estimation method detailed in Reference [79]. This model applies a smoothing function to each readout and then distinctly identifies both up and down transitions while taking into account the finite measurement bandwidth. From this information, we extract the effective transition rates. The transition-up rate is a combination of the readout-induced rate (Γ_{\uparrow}) and the rates due to ambient thermal and electrical noise at ω_q . At $\bar{n} = 14.6$, the effects of the readout are dominant. The transition-down rate is similarly due to the readout as well as typical environmentally-induced T_1 processes. Separating these two effects is subtle when the readout is nearly QND and T_1 processes dominate. The separation is accomplished by comparing the aforementioned environmental decay rate during the readout rise with the overall readout-power-dependent transition down rate. The readout-induced transition down rate (Γ_{\downarrow}) is then inferred as the difference between the two. We integrate these two readout-induced rates over the rise of the readout power in the cavity (as in Equation 7.7) to produce the readout-induced fidelity loss. The calculated fidelity loss during measurement due to Γ_{\uparrow} is $0.2 \pm 0.1\%$, while the contribution of Γ_{\downarrow} is negligible at $0.0 \pm 0.2\%$. The significance of the fidelity loss due to Γ_{\downarrow} as a potentially negative value suggests the possibility of state pinning and inhibition of qubit

Table 7.2. Fidelity loss in cQED/paramp readout at $\bar{n} = 14.6$. The measured fidelity is $91.0 \pm 0.4\%$. The remaining, unaccounted-for loss is attributed to π -pulse imperfections and drifts in the Larmor frequency due to local variation in the magnetic flux threading the qubit loop.

Source of loss	Fidelity loss (%)	Calculation method
T_1 decay	4.4 ± 0.3	measured T_1
thermal population	2.9 ± 0.2	heralding
Γ_{\uparrow}	0.2 ± 0.1	individual jumps
Γ_{\downarrow}	0.0 ± 0.2	individual jumps
SNR	< 0.1	pointer state overlap
remaining	1.5	-

decay due to the readout. This effect has been observed at comparable readout powers in transmon samples with a similar coupling strength [79].

In Table 7.2, we list the measured sources of fidelity loss and the method by which they are calculated. The unaccounted-for loss is estimated to contribute 1.5%. We speculate that this loss is explained by errors in the excited state preparation, namely π -pulse imperfections and drifts of the Larmor frequency due to local magnetic flux variations. We note that the dominant loss mechanism is energy relaxation. Currently, there are flux qubits with demonstrated relaxation times on the order of $10 \mu\text{s}$ [21, 128]; in our system, a similar T_1 would reduce the associated fidelity loss to 0.7%. Moreover, one could potentially employ cavities engineered to suppress spontaneous emission of the qubit [105]. This level of coherence, coupled with heralded ground state preparation, high-precision π pulses and stable magnetic flux, should readily enable readout fidelities in excess of 98% within this low power dispersive measurement architecture.

While the validation of the dressed dephasing model is not the primary subject of this thesis, we present preliminary measurements of this effect. Figure 7.13(a) shows the ensemble excited state population as a function of time during readout with and without heralded state preparation. The heralded data shows a significant reduction in spurious excited state population at the beginning of readout. In both cases, the excited state population equilibrates to a value significantly higher than at the start of readout. The presence of the readout appears to increase the effective thermal background temperature, suggesting upconversion of noise and qubit state mixing. Figure 7.13(b) shows the excited state population during readout averaged over a $20\text{-}\mu\text{s}$ period as a function of readout power. As the number of readout photons increases, the excited state is increasingly populated. The oscillation in the middle of the graph can be accounted for by variations in the environmental decay rate as the qubit is Stark-shifted with increasing readout power.

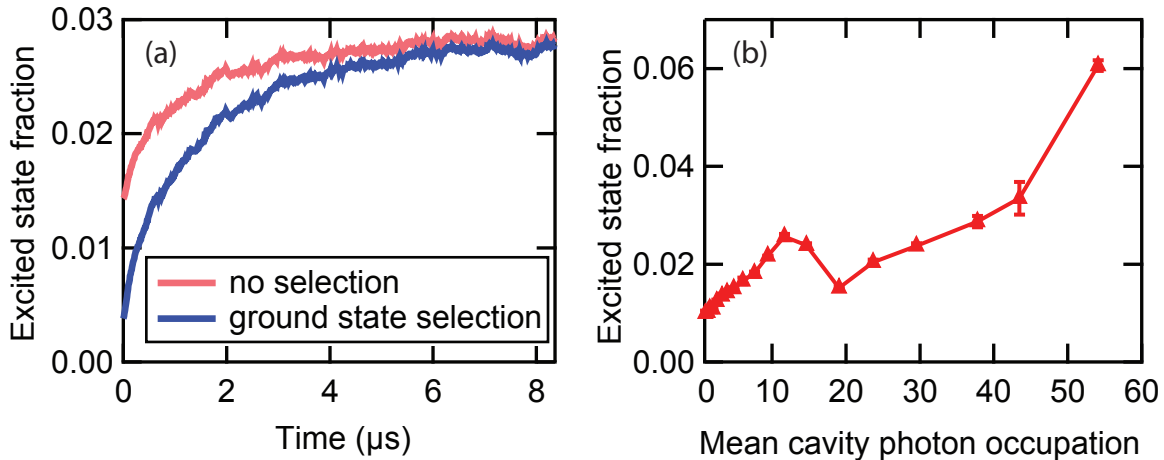


Figure 7.13. Readout backaction. (a) Fractional excited state population as a function of time after the readout resonator has been energized. Selection using heralded state preparation reduces the initial excited state population. (b) The time-averaged fractional excited state population increases vs readout power.

7.2.6 Fast qubit reset

Finally, we discuss a fast qubit initialization procedure based on heralding of the ground state. If the excited state is measured during the probe pulse, the state can be immediately flipped back to the ground state with a control pulse and rapidly checked again. This procedure eliminates the need to change the detuning δ to induce Purcell relaxation [105]. The ideal fidelity of this active reset to the ground state is ultimately limited by how the relaxation time of the qubit compares with κ , and is given approximately by the simple formula

$$F_{\text{reset}} = e^{-2\pi/\kappa T_1}, \quad (7.9)$$

which evaluates to 93% for our experimental parameters. The exponential argument in Equation 7.9 is of the form t_{reset}/T_1 , where t_{reset} is π times the resonator time constant $2/\kappa$. The time t_{reset} is therefore the time it takes for the resonator to ring down to 4% of its maximum amplitude. We measure a reset fidelity of 82%, which is close to the reset fidelity we expect to achieve (84%) after correcting for the fidelity loss ascribed to excited state preparation errors (1.5%) and an extra 140 ns of delay present in the pulse sequence used in this particular experiment (8.1%). The reset procedure could be iterated to achieve enhanced reset fidelity, but for qubits with $T_1 > 10 \mu\text{s}$, a single iteration of this procedure would achieve reset with close to 99% fidelity, assuming that the leftover fidelity loss we attribute to imperfect state preparation could be eliminated.

Chapter 8

Conclusions

8.1 Summary of results

In this thesis, we have pioneered the use of near-quantum-limited amplifiers as a tool for improving the dispersive readout of superconducting flux qubits. Our first experiment achieved successful integration of a microstrip SQUID amplifier (MSA) into a SQUID-resonator-based flux qubit measurement scheme. We believe this to be the first demonstration of a superconducting qubit readout to be enhanced using a transmission-based, near-quantum-limited SQUID amplifier. In addition to its transmission-based operation, the MSA also offers the advantage of requiring only static current and flux biases [41, 42, 44], greatly simplifying the microwave infrastructure required to operate the amplifier. The MSA allows for a substantial 4.5-fold increase in the readout visibility at low readout powers, achieving a measurement visibility of 72.3% at $\bar{n} = 1.5$ photons in the readout resonator. This measured value agrees well with simulations based on the characteristics of the readout chain. At these low powers, readout-induced reduction of T_1 is minimal and the MSA is operated within its optimum dynamic range. Below $\bar{n} = 0.1$, we see no suppression of T_1 due to the readout. The improved visibility allows access to this very weak continuous quantum measurement regime [132] while preserving sufficient SNR to resolve the qubit state more efficiently with minimal readout backaction.

We have also combined a flux qubit in the circuit QED architecture with a low-noise, Josephson parametric amplifier. Using this configuration, we have demonstrated a fast, analog cQED readout, measuring the state of the flux qubit with a high 91% single shot fidelity. We have also observed quantum jumps in a superconducting flux qubit, to our knowledge for the first time. This is made possible by the high system SNR, which produces readout pointer state distributions that are separated to an error of less than one part in 1000 in a single 10-ns acquisition with $\bar{n} = 14.6$. We use the QND nature of the readout at low powers to demonstrate heralded state preparation, further increasing the single shot fidelity to 93.9%. This procedure allows us to obtain a detailed account of the sources of fidelity loss in the system. Understanding and precisely quantifying these losses is important to engineering

a pathway to even higher fidelity. Additionally, heralded state preparation allows errors due to imperfect ground state initialization to be readily identified and effectively eliminated through post-selection. Finally, these heralding techniques permit a rapid, deterministic reset of the qubit state—a particularly important function in long-lived qubits where simply waiting for a time much longer than T_1 is impractical.

8.2 Future directions

The results of this work have opened up many new directions of research. The use of low-noise superconducting amplifiers has pushed the field of superconducting qubits closer to the ultimate goal of building a robust, practical quantum computer. Simultaneously, we have opened up new avenues of research in basic quantum measurement science.

The readout in our SQUID-resonator experiment was severely limited by the low quality factor of the readout resonator, producing a small phase shift between the qubit states. In future experiments, this could be increased (by engineering for a smaller coupling capacitance) to optimize the readout signal. This should allow real-time monitoring of the qubit state with very high visibility. The SNR improvement provided by the MSA should also be sufficient to observe quantum jumps in a superconducting qubit. Furthermore, although circulators and directional couplers were used in this experiment, in principle the forward directionality of the MSA eliminates the need for such nonreciprocal components when paired with a transmission-based readout resonator. This opens up the possibility of on-chip lithographic integration, which, needless to say, would require careful engineering of the resonator output and MSA input impedances. The direct coupling would also allow for an investigation of the MSA backaction noise spectrum—notably at the Josephson frequency—and its effect on qubit coherence.

The results of our cQED experiment naturally motivate the study of feedback in a superconducting qubit system. State preparation errors could be actively corrected in real time if a π pulse were to be triggered based on the outcome of the heralding readout. Additional work in quantum feedback has also been recently completed in Irfan Siddiqi’s laboratory using a transmon qubit. These experiments successfully demonstrated stabilization and persistence of Rabi oscillations using continuous weak measurement and feedback on the qubit drive [133]. The dominant remaining source of fidelity loss in our experiment is due to the short energy relaxation time of the qubit. For $T_1 > 10 \mu\text{s}$ with perfect state preparation, we predict a fidelity of greater than 98% in the current readout architecture. Our readout scheme also provides a new tool to study readout-induced backaction in a qubit with a high degree of anharmonicity, complementing previous work on transmon qubits [131, 79]. The flux qubit, with its high degree of anharmonicity, provides a valuable comparison with the transmon in validating the full dressed dephasing model. Lastly, our high fidelity readout is ready to be extended to coupled-qubit systems. High joint fidelity in these systems has remained elusive, but is ultimately necessary for continuing progress towards a scalable quantum processor.

Bibliography

- [1] A. O. Caldeira and A. J. Leggett. “Influence of dissipation on quantum tunneling in macroscopic systems.” *Phys. Rev. Lett.* **46**, 211 (1981).
- [2] M. H. Devoret, J. M. Martinis, and J. Clarke. “Measurements of macroscopic quantum tunneling out of the zero-voltage state of a current-biased josephson junction.” *Phys. Rev. Lett.* **55**, 1908 (1985).
- [3] J. M. Martinis, M. H. Devoret, and J. Clarke. “Experimental tests for the quantum behavior of a macroscopic degree of freedom: The phase difference across a josephson junction.” *Phys. Rev. B* **35**, 4682 (1987).
- [4] J. Clarke, A. N. Cleland, M. H. Devoret, D. Esteve, and J. M. Martinis. “Quantum mechanics of a macroscopic variable: The phase difference of a josephson junction.” *Science* **239**, 992 (1988).
- [5] T. Hime, P. A. Reichardt, B. L. T. Plourde, T. L. Robertson, C. E. Wu, A. V. Ustinov, and J. Clarke. “Solid-state qubits with current-controlled coupling.” *Science* **314**, 1427 (2006).
- [6] I. L. Chuang and M. A. Nielsen. *Quantum Computation and Quantum Information*. Cambridge: Cambridge University Press (2000).
- [7] R. P. Feynman. “Simulating physics with computers.” *Int. J. Theor. Phys.* **21**, 467 (1982).
- [8] P. W. Shor. “Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer.” *SIAM Review* **41**, 303 (1999).
- [9] L. K. Grover. “Quantum mechanics helps in searching for a needle in a haystack.” *Phys. Rev. Lett.* **79**, 325 (1997).
- [10] K. L. Brown, W. J. Munro, and V. M. Kendon. “Using quantum computers for quantum simulation.” *Entropy* **12**, 2268 (2010).
- [11] D. Loss and D. P. DiVincenzo. “Quantum computation with quantum dots.” *Phys. Rev. A* **57**, 120 (1998).

- [12] J. Clarke and F. K. Wilhelm. “Superconducting quantum bits.” *Nature* **453**, 1031 (2008).
- [13] I. Siddiqi. “Superconducting qubits: poised for computing?” *Superconductor Science and Technology* **24**, 091002 (2011).
- [14] L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf. “Demonstration of two-qubit algorithms with a superconducting quantum processor.” *Nature* **460**, 240 (2009).
- [15] L. DiCarlo, M. D. Reed, L. Sun, B. R. Johnson, J. M. Chow, J. M. Gambetta, L. Frunzio, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf. “Preparation and measurement of three-qubit entanglement in a superconducting circuit.” *Nature* **467**, 574 (2010).
- [16] M. Neeley, R. C. Bialczak, M. Lenander, E. Lucero, M. Mariantoni, A. D. O’Connell, D. Sank, H. Wang, M. Weides, J. Wenner, Y. Yin, T. Yamamoto, A. N. Cleland, and J. M. Martinis. “Generation of three-qubit entangled states using superconducting phase qubits.” *Nature* **467**, 570 (2010).
- [17] M. Mariantoni, H. Wang, T. Yamamoto, M. Neeley, R. C. Bialczak, Y. Chen, M. Lenander, E. Lucero, A. D. O’Connell, D. Sank, M. Weides, J. Wenner, Y. Yin, J. Zhao, A. N. Korotkov, A. N. Cleland, and J. M. Martinis. “Implementing the quantum von neumann architecture with superconducting circuits.” *Science* **334**, 61 (2011).
- [18] D. DiVincenzo. “The physical implementation of quantum computation.” *Fortschr. Phys.* **48**, 771 (2000).
- [19] J. I. Cirac and P. Zoller. “Quantum computations with cold trapped ions.” *Phys. Rev. Lett.* **74**, 4091 (1995).
- [20] H. Häffner, C. Roos, and R. Blatt. “Quantum computing with trapped ions.” *Physics Reports* **469**, 155 (2008).
- [21] J. Bylander, S. Gustavsson, F. Yan, F. Yoshihara, K. Harrabi, G. Fitch, D. G. Cory, Y. Nakamura, J.-S. Tsai, and W. D. Oliver. “Noise spectroscopy through dynamical decoupling with a superconducting flux qubit.” *Nature Phys.* **7**, 565 (2011).
- [22] 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. “Observation of high coherence in josephson junction qubits measured in a three-dimensional circuit qed architecture.” *Phys. Rev. Lett.* **107**, 240501 (2011).
- [23] M. Steffen. “Superconducting qubits are getting serious.” *Physics* **4**, 103 (2011).

- [24] T. Van Duzer and C. W. Turner. *Principles of Superconductive Devices and Circuits*. Prentice Hall, New Jersey, second ed. (1999).
- [25] B. Josephson. “Possible new effects in superconductive tunnelling.” *Physics Letters* **1**, 251 (1962).
- [26] J. Clarke and A. I. Braginski (Eds.). *The SQUID Handbook, Vol. I: Fundamentals and Technology of SQUIDs and SQUID Systems*. Wiley-VCH, Weinheim (2004).
- [27] M. Tinkham. *Introduction to Superconductivity*. McGraw-Hill, New York, second ed. (1996).
- [28] D. J. V. Harlingen, R. H. Koch, and J. Clarke. “Superconducting quantum interference device with very low magnetic flux noise energy.” *Applied Physics Letters* **41**, 197 (1982).
- [29] C. D. Tesche and J. Clarke. “dc squid: Noise and optimization.” *Journal of Low Temperature Physics* **29**, 301 (1977).
- [30] A. J. Leggett. *J. Phys. (Paris), Colloq* **39**, C6–1264 (1978).
- [31] B. Yurke and J. S. Denker. “Quantum network theory.” *Phys. Rev. A* **29**, 1419 (1984).
- [32] M. H. Devoret. “Quantum fluctuations in electrical circuits.” In *Quantum Fluctuations, Les Houches LXIII*, S. Reynaud, E. Giacobino, and J. Zinn-Justin, eds., chap. 10. New York: Elsevier (1995).
- [33] J. M. Martinis. *Quant. Info. Proc.* **8**, 81 (2009).
- [34] Y. Nakamura, Y. A. Pashkin, and J. S. Tsai. “Coherent control of macroscopic quantum states in a single-Cooper-pair box.” *Nature* **398**, 786 (1999).
- [35] C. H. van der Wal, A. C. J. ter Haar, F. K. Wilhelm, R. N. Schouten, C. J. P. M. Harmans, T. P. Orlando, S. Lloyd, and J. E. Mooij. “Quantum superposition of macroscopic persistent-current states.” *Science* **290**, 773 (2000).
- [36] D. Vion, A. Aassime, A. Cottet, P. Joyez, H. Pothier, C. Urbina, D. Esteve, and M. H. Devoret. “Manipulating the quantum state of an electrical circuit.” *Science* **296**, 886 (2002).
- [37] A. D. Córcoles, J. M. Chow, J. M. Gambetta, C. Rigetti, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen. “Protecting superconducting qubits from radiation.” *Applied Physics Letters* **99**, 181906 (2011).
- [38] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. “Charge-insensitive qubit design derived from the cooper pair box.” *Phys. Rev. A* **76**, 042319 (2007).

- [39] V. E. Manucharyan, J. Koch, L. I. Glazman, and M. H. Devoret. “Fluxonium: Single cooper-pair circuit free of charge offsets.” *Science* **326**, 113 (2009).
- [40] W. H. Zurek. “Decoherence, einselection, and the quantum origins of the classical.” *Rev. Mod. Phys.* **75**, 715 (2003).
- [41] M. Mück, M.-O. André, J. Clarke, J. Gail, and C. Heiden. “Radio-frequency amplifier based on a niobium dc superconducting quantum interference device with microstrip input coupling.” *Appl. Phys. Lett.* **72**, 2885 (1998).
- [42] M. P. DeFeo, P. Bhupathi, K. Yu, T. W. Heitmann, C. Song, R. McDermott, and B. L. T. Plourde. “Microstrip superconducting quantum interference device amplifiers with submicron josephson junctions: Enhanced gain at gigahertz frequencies.” *Appl. Phys. Lett.* **97**, 092507 (2010).
- [43] B. Abdo, F. Schackert, M. Hatridge, C. Rigetti, and M. Devoret. “Josephson amplifier for qubit readout.” *Appl. Phys. Lett.* **99**, 162506 (2011).
- [44] L. Spietz, K. Irwin, and J. Aumentado. “Superconducting quantum interference device amplifiers with over 27 ghz of gain-bandwidth product operated in the 4-8 ghz frequency range.” *Appl. Phys. Lett.* **95**, 092505 (2009).
- [45] M. A. Castellanos-Beltran and K. W. Lehnert. *Appl. Phys. Lett.* **91**, 083509 (2007).
- [46] G. J. Ribeill, D. Hover, Y.-F. Chen, S. Zhu, and R. McDermott. “Superconducting low-inductance undulatory galvanometer microwave amplifier: Theory.” *Journal of Applied Physics* **110**, 103901 (2011).
- [47] M. Hatridge, R. Vijay, D. H. Slichter, J. Clarke, and I. Siddiqi. “Dispersive magnetometry with a quantum limited squid parametric amplifier.” *Phys. Rev. B* **83**, 134501 (2011).
- [48] R. Bradley, J. Clarke, D. Kinion, L. J. Rosenberg, K. v. Bibber, S. Matsuki, M. Mück, and P. Sikivie. “Microwave cavity searches for dark-matter axions.” *Rev. Mod. Phys.* **75**, 777 (2003).
- [49] J. R. Friedman, V. Patel, W. Chen, S. K. Tolpygo, and J. E. Lukens. “Quantum superposition of distinct macroscopic states.” *Nature* **406**, 43 (2000).
- [50] J. E. Mooij, T. P. Orlando, L. Levitov, L. Tian, C. H. van der Wal, and S. Lloyd. “Josephson persistent-current qubit.” *Science* **285**, 1036 (1999).
- [51] B. L. T. Plourde, T. L. Robertson, P. A. Reichardt, T. Hime, S. Linzen, C.-E. Wu, and J. Clarke. “Flux qubits and readout device with two independent flux lines.” *Phys. Rev. B* **72**, 060506 (2005).

- [52] T. L. Robertson, B. L. T. Plourde, P. A. Reichardt, T. Hime, C.-E. Wu, and J. Clarke. “Quantum theory of three-junction flux qubit with non-negligible loop inductance: Towards scalability.” *Phys. Rev. B* **73**, 174526 (2006).
- [53] G. Burkard, D. P. DiVincenzo, P. Bertet, I. Chiorescu, and J. E. Mooij. “Asymmetry and decoherence in a double-layer persistent-current qubit.” *Phys. Rev. B* **71**, 134504 (2005).
- [54] Y. S. Greenberg, A. Izmalkov, M. Grajcar, E. Il’ichev, W. Krech, H.-G. Meyer, M. H. S. Amin, and A. M. van den Brink. “Low-frequency characterization of quantum tunneling in flux qubits.” *Phys. Rev. B* **66**, 214525 (2002).
- [55] T. P. Orlando, J. E. Mooij, L. Tian, C. H. van der Wal, L. S. Levitov, S. Lloyd, and J. J. Mazo. “Superconducting persistent-current qubit.” *Phys. Rev. B* **60**, 15398 (1999).
- [56] Y. Makhlin, G. Schön, and A. Shnirman. “Quantum-state engineering with josephson-junction devices.” *Rev. Mod. Phys.* **73**, 357 (2001).
- [57] P. A. Reichardt. Ph.D. thesis, University of California, Berkeley (2006).
- [58] A. Lupaşcu, C. J. P. M. Harmans, and J. E. Mooij. “Quantum state detection of a superconducting flux qubit using a dc-squid in the inductive mode.” *Phys. Rev. B* **71**, 184506 (2005).
- [59] A. A. Abdumalikov, O. Astafiev, Y. Nakamura, Y. A. Pashkin, and J. S. Tsai. “Vacuum rabi splitting due to strong coupling of a flux qubit and a coplanar-waveguide resonator.” *Phys. Rev. B* **78**, 180502 (2008).
- [60] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf. “Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation.” *Phys. Rev. A* **69**, 062320 (2004).
- [61] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf. “Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics.” *Nature* **431**, 162 (2004).
- [62] D. F. Walls and G. J. Milburn. *Quantum Optics*. New York: Springer, second ed. (2008).
- [63] V. B. Braginsky and F. Y. Khalili. *Quantum Measurement*. New York: Cambridge University Press (1992).
- [64] A. Lupaşcu. Ph.D. thesis, Delft University of Technology (2005).
- [65] D. M. Pozar. *Microwave Engineering*. Wiley, third ed.

- [66] C. D. Motchenbacher and J. A. Connelly. *Low-Noise Electronic System Design*. New York: John Wiley (1993).
- [67] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf. “Introduction to quantum noise, measurement, and amplification.” *Rev. Mod. Phys.* **82**, 1155 (2010).
- [68] H. A. Haus and J. A. Mullen. “Quantum noise in linear amplifiers.” *Phys. Rev.* **128**, 2407 (1962).
- [69] C. M. Caves. “Quantum limits on noise in linear amplifiers.” *Phys. Rev. D* **26**, 1817 (1982).
- [70] K. Rohlfs and T. L. Wilson. *Tools of Radio Astronomy*. Springer, second ed.
- [71] M. Mück, J. Kycia, and J. Clarke. “Superconducting quantum interference device as a near-quantum-limited amplifier at 0.5 ghz.” *Appl. Phys. Lett.* **78**, 967 (2001).
- [72] M. Mück, C. Welzel, and J. Clarke. “Superconducting quantum interference device amplifiers at gigahertz frequencies.” *Appl. Phys. Lett.* **82**, 3266 (2003).
- [73] D. Kinion and J. Clarke. *Appl. Phys. Lett.* **92**, 172503 (2008).
- [74] M. B. Ketchen and J. M. Jaycox. “Ultra-low-noise tunnel junction dc squid with a tightly coupled planar input coil.” *Appl. Phys. Lett.* **40**, 736 (1982).
- [75] J. Clarke and A. I. Braginski (Eds.). *The SQUID Handbook, Vol. II: Fundamentals and Technology of SQUIDs and SQUID Systems*. Wiley-VCH, Weinheim (2006).
- [76] D. Kinion and J. Clarke. “Superconducting quantum interference device as a near-quantum-limited amplifier for the axion dark-matter experiment.” *Appl. Phys. Lett.* **98**, 202503 (2011).
- [77] D. Kinion and J. Clarke. “Microstrip superconducting quantum interference device amplifier: Conditional stability.” *Applied Physics Letters* **96**, 172501 (2010).
- [78] R. Vijay, D. H. Slichter, and I. Siddiqi. “Observation of quantum jumps in a superconducting artificial atom.” *Phys. Rev. Lett.* **106**, 110502 (2011).
- [79] D. H. Slichter. Ph.D. thesis, University of California, Berkeley (2011).
- [80] R. Vijay. Ph.D. thesis, Yale University (2008).
- [81] M. J. Hatridge. Ph.D. thesis, University of California, Berkeley (2010).
- [82] R. Vijay, M. H. Devoret, and I. Siddiqi. “Invited review article: The josephson bifurcation amplifier.” *Review of Scientific Instruments* **80**, 111101 (2009).

- [83] I. Siddiqi, R. Vijay, F. Pierre, C. M. Wilson, L. Frunzio, M. Metcalfe, C. Rigetti, R. J. Schoelkopf, M. H. Devoret, D. Vion, and D. Esteve. “Direct observation of dynamical bifurcation between two driven oscillation states of a josephson junction.” *Phys. Rev. Lett.* **94**, 027005 (2005).
- [84] I. Siddiqi, R. Vijay, M. Metcalfe, E. Boaknin, L. Frunzio, R. J. Schoelkopf, and M. H. Devoret. “Dispersive measurements of superconducting qubit coherence with a fast latching readout.” *Phys. Rev. B* **73**, 054510 (2006).
- [85] F. Mallet, F. R. Ong, A. Palacios-Laloy, F. Nguyen, P. Bertet, D. Vion, and D. Esteve. “Single-shot qubit readout in circuit quantum electrodynamics.” *Nature Phys.* **5**, 791 (2009).
- [86] W. W. Hu, K. Sarveswaran, M. Lieberman, and G. H. Bernstein. “Sub-10 nm electron beam lithography using cold development of poly(methylmethacrylate).” *J. Vac. Sci. Technol. B* **22**, 1711 (2004).
- [87] B. Cord, J. Lutkenhaus, and K. K. Berggren. “Optimal temperature for development of poly(methylmethacrylate).” 2013–2016. AVS (2007).
- [88] L. E. Ocola and A. Stein. “Effect of cold development on improvement in electron-beam nanopatterning resolution and line roughness.” 3061–3065. AVS (2006).
- [89] H. Wang, G. M. Laws, S. Milicic, P. Boland, A. Handugan, M. Pratt, T. Eschrich, S. Myhajlenko, J. A. Allgair, and B. Bunday. “Low temperature zep-520a development process for enhanced critical dimension realization in reactive ion etched polysilicon.” *J. Vac. Sci. Technol. B* **25**, 102 (2007).
- [90] G. J. Dolan. “Offset masks for lift-off photoprocessing.” *Applied Physics Letters* **31**, 337 (1977).
- [91] R. E. Miller, W. H. Mallison, A. W. Kleinsasser, K. A. Delin, and E. M. Macedo. “Niobium trilayer josephson tunnel junctions with ultrahigh critical current densities.” *Applied Physics Letters* **63**, 1423 (1993).
- [92] I. M. Pop, T. Fournier, T. Crozes, F. Lecocq, I. Matei, B. Pannetier, O. Buisson, and W. Guichard. “Fabrication of stable and reproducible submicron tunnel junctions.” *J. Vac. Sci. Technol. B* **30**, 010607 (2012).
- [93] V. Ambegaokar and A. Baratoff. “Tunneling between superconductors.” *Phys. Rev. Lett.* **10**, 486 (1963).
- [94] M. Kamon, M. J. Tsuk, and J. K. White. “Fasthenry: a multipole-accelerated 3-d inductance extraction program.” *IEEE Trans. Microwave Theory Tech.* **42**, 1750 (1994).

- [95] T. L. Robertson, B. L. T. Plourde, T. Hime, S. Linzen, P. A. Reichardt, F. K. Wilhelm, and J. Clarke. “Superconducting quantum interference device with frequency-dependent damping: Readout of flux qubits.” *Phys. Rev. B* **72**, 024513 (2005).
- [96] T. L. Robertson. Ph.D. thesis, University of California, Berkeley (2005).
- [97] A. J. Annunziata, D. F. Santavicca, L. Frunzio, G. Catelani, M. J. Rooks, A. Frydman, and D. E. Prober. “Tunable superconducting nanoinductors.” *Nanotechnology* **21**, 445202 (2010).
- [98] A. O. Niskanen, K. Harrabi, F. Yoshihara, Y. Nakamura, and J. S. Tsai. “Spectroscopy of three strongly coupled flux qubits.” *Phys. Rev. B* **74**, 220503 (2006).
- [99] A. C. J. ter Haar. Ph.D. thesis, Delft University of Technology (2005).
- [100] H. Meissner. “Superconductivity of contacts with interposed barriers.” *Phys. Rev.* **117**, 672 (1960).
- [101] S. Linzen, T. L. Robertson, T. Hime, B. L. T. Plourde, P. A. Reichardt, and J. Clarke. “Low-noise computer-controlled current source for quantum coherence experiments.” *Rev. Sci. Instrum.* **75**, 2541 (2004).
- [102] D. H. Slichter, O. Naaman, and I. Siddiqi. “Millikelvin thermal and electrical performance of lossy transmission line filters.” *Appl. Phys. Lett.* **94**, 192508 (2009).
- [103] T. Hime. Ph.D. thesis, University of California, Berkeley (2007).
- [104] J. Bourassa, J. M. Gambetta, A. A. Abdumalikov, O. Astafiev, Y. Nakamura, and A. Blais. “Ultrastrong coupling regime of cavity qed with phase-biased flux qubits.” *Phys. Rev. A* **80**, 032109 (2009).
- [105] M. D. Reed, B. R. Johnson, A. A. Houck, L. DiCarlo, J. M. Chow, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf. “Fast reset and suppressing spontaneous emission of a superconducting qubit.” *Applied Physics Letters* **96**, 203110 (2010).
- [106] I. I. Rabi. “Space quantization in a gyrating magnetic field.” *Phys. Rev.* **51**, 652 (1937).
- [107] J. J. Sakurai. *Modern Quantum Mechanics*. Addison-Wesley, Massachusetts, revised ed. (1994).
- [108] G. Ithier, E. Collin, P. Joyez, P. J. Meeson, D. Vion, D. Esteve, F. Chiarello, A. Shnirman, Y. Makhlin, J. Schrieffer, and G. Schön. “Decoherence in a superconducting quantum bit circuit.” *Phys. Rev. B* **72**, 134519 (2005).

- [109] F. K. Wilhelm, M. J. Storez, U. Hartmann, and M. R. Geller. “Superconducting qubits ii: Decoherence.” In *Manipulating Quantum Coherence in Solid State Systems*, M. E. Flatté and I. Trifea, eds., chap. 6. Springer: Dordrecht, The Netherlands (2005).
- [110] R. H. Koch, J. Clarke, W. M. Goubau, J. M. Martinis, C. M. Pegrum, and D. J. Harlingen. “Flicker ($1/f$) noise in tunnel junction dc squids.” *Journal of Low Temperature Physics* **51**, 207 (1983). 10.1007/BF00683423.
- [111] S. M. Anton, C. Müller, J. S. Birenbaum, S. R. O’Kelley, A. D. Fefferman, D. S. Golubev, G. C. Hilton, H.-M. Cho, K. D. Irwin, F. C. Wellstood, G. Schön, A. Shnirman, and J. Clarke. “Pure dephasing in flux qubits due to flux noise with spectral density scaling as $1/f^\alpha$.” *Phys. Rev. B* **85**, 224505 (2012).
- [112] N. F. Ramsey. “A molecular beam resonance method with separated oscillating fields.” *Phys. Rev.* **78**, 695 (1950).
- [113] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. “Approaching unit visibility for control of a superconducting qubit with dispersive readout.” *Phys. Rev. Lett.* **95**, 060501 (2005).
- [114] P. W. Shor. “Scheme for reducing decoherence in quantum computer memory.” *Phys. Rev. A* **52**, R2493 (1995).
- [115] M. D. Reed, L. DiCarlo, S. E. Nigg, L. Sun, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf. “Realization of three-qubit quantum error correction with superconducting circuits.” *Nature* **482**, 382 (2012).
- [116] D. Aharonov and M. Ben-Or. “Fault tolerant quantum computation with constant error.” In *Proc. ACM STOC*, 176–188 (1997).
- [117] E. Knill, R. Laflamme, and W. Zurek. “Resilient quantum computation: Error models and thresholds.” *Proc. Roy. Soc. Lond. A* **454**, 365 (1997).
- [118] J. Preskill. “Fault-tolerant quantum computation.” In *Introduction to Quantum Computation*, H.-K. Lo, S. Popescu, and T. P. Spiller, eds. World Scientific, Singapore (1998).
- [119] M. D. Reed, L. DiCarlo, B. R. Johnson, L. Sun, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf. “High-fidelity readout in circuit quantum electrodynamics using the jaynes-cummings nonlinearity.” *Phys. Rev. Lett.* **105**, 173601 (2010).
- [120] J. E. Johnson, E. M. Hoskinson, C. Macklin, D. H. Slichter, I. Siddiqi, and J. Clarke. “Dispersive readout of a flux qubit at the single-photon level.” *Phys. Rev. B* **84**, 220503 (2011).

- [121] R. W. Simmonds, K. M. Lang, D. A. Hite, S. Nam, D. P. Pappas, and J. M. Martinis. “Decoherence in josephson phase qubits from junction resonators.” *Phys. Rev. Lett.* **93**, 077003 (2004).
- [122] J. Gambetta, W. A. Braff, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf. “Protocols for optimal readout of qubits using a continuous quantum nondemolition measurement.” *Phys. Rev. A* **76**, 012325 (2007).
- [123] J. E. Johnson, C. Macklin, D. H. Slichter, R. Vijay, E. B. Weingarten, J. Clarke, and I. Siddiqi. “Heralded state preparation in a superconducting qubit.” *ArXiv* 1202.5541 (2012).
- [124] N. Piro, F. Rohde, C. Schuck, M. Almendros, J. Huwer, J. Ghosh, A. Haase, M. Henrich, F. Dubin, and J. Escher. “Heralded single-photon absorption by a single atom.” *Nature Phys.* **7**, 17 (2011).
- [125] C. Wagenknecht, C.-M. Li, A. Reingruber, X.-H. Bao, A. Goebel, Y.-A. Chen, Q. Zhang, K. Chen, and J.-W. Pan. “Experimental demonstration of a heralded entanglement source.” *Nature Photonics* **4**, 549 (2010).
- [126] A. Ulhaq, S. Weiler, S. M. Ulrich, R. Roßbach, M. Jetter, and P. Michler. “Cascaded single-photon emission from the Mollow triplet sidebands of a quantum dot.” *Nature Photonics* **6**, 238 (2012).
- [127] A. Palacios-Laloy, F. Mallet, F. Nguyen, F. Ong, P. Bertet, D. Vion, and D. Esteve. “Spectral measurement of the thermal excitation of a superconducting qubit.” *Physica Scripta* **2009**, 014015 (2009).
- [128] A. D. Córcoles, J. M. Chow, J. M. Gambetta, C. Rigetti, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen. “Protecting superconducting qubits from radiation.” *Appl. Phys. Lett.* **99**, 181906 (2011).
- [129] D. I. Schuster. Ph.D. thesis, Yale University (2007).
- [130] M. Boissonneault, J. M. Gambetta, and A. Blais. “Nonlinear dispersive regime of cavity qed: The dressed dephasing model.” *Phys. Rev. A* **77**, 060305 (2008).
- [131] M. Boissonneault, J. M. Gambetta, and A. Blais. “Dispersive regime of circuit qed: Photon-dependent qubit dephasing and relaxation rates.” *Phys. Rev. A* **79**, 013819 (2009).
- [132] A. Palacios-Laloy, F. Mallet, F. Nguyen, P. Bertet, D. Vion, D. Esteve, and A. N. Korotkov. “Experimental violation of a bell’s inequality in time with weak measurement.” *Nature Phys.* **6**, 442 (2010).

- [133] R. Vijay, C. Macklin, D. H. Slichter, S. J. Weber, K. W. Murch, R. Naik, A. N. Korotkov, and I. Siddiqi. “Quantum feedback control of a superconducting qubit: Persistent Rabi oscillations.” *ArXiv* 1205.5591 (2012).