

ABSTRACT

Title of dissertation: **REDUCING DECOHERENCE IN
dc SQUID PHASE QUBITS**

Anthony J. Przybysz, Doctor of Philosophy, 2010

Dissertation directed by: **Professor Frederick C. Wellstood
Department of Physics**

This thesis examines sources of dissipation and dephasing in a dc SQUID phase qubit. Coupling of the qubit to the bias lines and lossy dielectrics causes the qubit to lose quantum information through a process known generally as decoherence. Using knowledge of the possible sources of decoherence, a dc SQUID phase qubit is designed with parameters that should have made it resistant to dissipation and dephasing from those sources. Device PB9 was a dc SQUID with one small area $0.23\ (\mu\text{m})^2$ Josephson junction with a critical current of 130 nA, which was meant to be the qubit junction, and a larger area $5\ (\mu\text{m})^2$ junction with a critical current of 8.6 μA , which acted as part of an inductive isolation network. The qubit junction was shunted by a 1.5 pF low-loss interdigitated capacitor. The dc current bias line had an on-chip LC filter with a cutoff frequency of 180 MHz. The other control lines were also designed to minimize coupling of dissipative elements to the qubit. According to a theoretical model of the dissipation and dephasing, the qubit was expected to have an energy relaxation $T_1 \leq 8.4\ \mu\text{s}$ and dephasing time $T_\phi \sim 1\ \mu\text{s}$.

Because of the relatively high Josephson inductance of the qubit junction, the device did not act perform like a conventional isolated single-junction phase qubit. Instead, the resonant modes that I observed were the normal modes of the entire SQUID.

At 20 mK and a frequency of 4.047 GHz, the maximum energy relaxation time of the device was found to be 350 ± 70 ns, despite the optimized design. Through a study of T_1 versus applied flux, T_1 was found to depend on the strength of the coupling of the microwave drive line to the qubit. When the line was more coupled, T_1 was shorter. This was evidence that the microwave line was overcoupled to the qubit, and was limiting the lifetime of the excited state T_1 .

Through a study of the spectroscopic coherence time T_2^* , which measured the effects of low-frequency inhomogeneous broadening and higher frequency dephasing from noise, I discovered that device PB9 has several sweet spots. In particular, the presence of a sweet spot with respect to critical current fluctuations allowed me to identify critical current noise as a major source of broadening and dephasing in the qubit. From the spectroscopy I estimated the $1/f$ critical current noise power density at 1 Hz was $S_{I_{c1}}(1\text{Hz}) \approx 10(\text{nA})^2/\text{Hz}$ and the $1/f$ flux noise power spectral density at 1 Hz was $S_{\Phi}(1\text{Hz}) \approx (110 \mu\Phi_0)^2/\text{Hz}$. Both of these values were quite high, possibly due to switching of the device between measurements.

REDUCING DECOHERENCE IN
dc SQUID PHASE QUBITS

by

Anthony J. Przybysz

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2010

Advisory Committee:

Professor Frederick C. Wellstood, Advisor / Chair
Professor Christopher J. Lobb, Advisor
Professor J. Robert Anderson, Advisor
Professor Alex J. Dragt
Professor Christopher C. Davis

© Copyright by
Anthony J. Przybysz
2010

Version: December 1, 2010

*For my family,
especially my wife, Diana*

Acknowledgements

I would first like to thank the Lord God because I would not have been able to accomplish any of this without Him.

I especially thank my wife, Diana, and my sons, Antonio and Tristan, for their patience and support, which allowed me to explore some very interesting problems in physics. It's not easy to be married and have two little boys at home while working towards a Ph.D, but we made it through together. In fact, looking back, I think that single grad students have it far too easy.

I also thank to my parents, John and Lauri. I think my talent for math came from many, many games of Life with my mom when I was very young, and my love of science came from my dad. To me, growing up, my dad had the coolest job in the world. My science projects were always physics projects, and my dad pointed me toward the most interesting ones.

Thank you very much to my advisor, Dr. Fred Wellstood. Without his support, my academic career would have ended before it even began. Fred has always been able to quickly illuminate subjects on superconductivity and quantum information that I puzzled over for far too long before I came to ask him about them. He gave me a chance to work on a project that combined a challenging experimental environment with a rich quantum mechanical theory; an opportunity for which I am especially grateful.

I can't go further without acknowledging my mother-in-law, Cathy Pifer. Without her and her husband, Richard's, support, Diana and I would have been living on

the street. She helped us to raise our boys during the difficult times of grad school and job changes.

Dr. Chris Lobb and Dr. Bob Anderson were also wonderful advisors. Chris actually encouraged me to interrupt him in his office, and he was an invaluable resource on the inner workings of Josephson junctions. Bob was my academic advisor, and kept me on track with my courses. He came down to the sub-basement every Friday afternoon to see what we were up to and flaunt his mastery of the Polish language. It's Pshibish, not Prizbiz... I know.

I owe much to my compatriots in the sub-basement. Hyeokshin Kwon and I spent hours working on the dilution refrigerators, and he helped me through many a broken coax cable. Ben Cooper never shied away from a theoretical calculation, and his work on the normal modes of the SQUID was a breakthrough in understanding my device. Rangga Budoyo worked very closely with me on taking the data on device PB9, and he was a sturdy lab partner. I knew that I could count on him to carry out a task even though he was relatively new to taking data on qubits. Thanks especially to him for continuing the work on the dc SQUID phase qubits and putting into practice the recommendations that came from the results on PB9.

While making the device I had an undergraduate assistant, Erik Crowe, who impressed me greatly. He took instructions on the machines like the SEM and evaporator very quickly, and often he performed large chunks of the fabrication when I wasn't even in the building.

Kevin Osborn's group at LPS helped me tremendously. Sergiy Gladchenko grew the SiN on device PB9, as well as performing the first few photolithographic steps. Moe

Khalil and Micah Stoutimore helped several times with microwave simulations and helped me understand high frequency phenomena in superconducting resonators. Kevin Osborn and Hanhee Paik were instrumental in the developing of the low loss silicon nitride, as well.

Ben Palmer's group helped me for no other reason than I kept asking them. Ben allowed me to have at least two sample boxes, no questions asked. Vitaley diced about 10 sapphire wafers without voicing a single complaint. Sergey Novikov and Zaeill Kim lent their microwave expertise to my efforts, as well.

Thanks especially to Dr. Michael Fuhrer's group. Michael gave me a job as an undergrad, and I continued to utilize his facilities all through grad school. The SEM, AFM, and even the fume hood I used all belonged to the Fuhrer group. Thanks to Dan Lenski for help on the SEM, and to Alex Curtin for her patience with the AFM. I doubt that the few times that I cleaned the evaporator made up for the hours and hours I spent on the SEM, so at the very least you have my heart-felt thanks.

We had a high school student, Gabe, from the Barrie School work for us for a couple of weeks. I tasked him with programming our AWG, and with Ranga and my help, he actually accomplished the task. That boy should go far. Thanks also to our undergrad Ben Cheng for his work on magnetic field simulators.

The physics department and CNAM staff were always very helpful. Thanks especially to Brian Straughn and Doug Bensen who helped me to fix several pieces of broken machinery. Thanks also to Al Godinez who kept the LHe flowing to the sub-basement.

Thanks to the Fablab staff in the engineering department. Tom Loughran trained me on much of the equipment in the cleanroom.

I want to acknowledge some of the older student who mentored me through the years. Todd Brintlinger took charge of me when I was an undergrad. On the other side of it, now, I can see what an inconvenience it probably was, but he taught me much about working in a lab. Rupert Lewis's presence in the sub-basement is sorely missed by all. He checked on me everyday when I started with Fred's group to make sure that I was learning and I had plenty of work to do. Hanhee Paik took me on as her apprentice and taught me about running a dil fridge. Also, her work on low-loss SiN was important for my research, and I am very proud of her success in the quantum computing field. Sudeep Dutta was a superman of experimental physics. He taught me by example how to take a setup apart piece by piece until he found the malfunction. It was clear that he looked up to his predecessors like Huizhong Xu and Andrew Berkley, but I doubt he realized how highly esteemed he was by his juniors. Tauno Palomaki was there to mentor me through my first measurements on a qubit.

I did several internships at Northrop Grumman, and without that experience and the people I met there, I doubt that I would have the opportunities that I now have coming out of grad school. Thanks to Aaron Pesetski, Hong Pesetski, Jim Murduck, Don Miller, Chris Lavoie, and everyone at ATL for letting me work for you.

It's been said that, "we see so far because we stand on the shoulders of giants," and the phrase has never been so true for me than in the field of quantum computing. The field has attracted some of the greatest minds of our time. I'm grateful to the men and women who made the progress up to point where I joined the conversation, and I am glad

to have made some small contribution to the wealth of knowledge that already exists in the field.

Table of Contents

List of Tables	xi
List of Figures	xii
1. Introduction	1
1.1. A Brief History of Quantum Computing	1
1.2. The Divincenzo Criteria: Requirements for a Successful Qubit.....	4
1.3. Advantages of Phase Qubits	6
1.4. Overview of My Thesis	8
2. Josephson Junctions as Qubits	12
2.1. Derivation of Josephson Relations and Josephson Inductance.....	13
2.2. RCSJ Model and the Tilted Washboard Potential	18
2.3. The Josephson Junction as a Quantum Anharmonic Oscillator	21
2.4. State Manipulation with a Microwave Drive.....	26
2.5. Escape Rate and Tunneling	30
2.6. Basic Qubit Operation and Pulsed Readout.....	32
2.7. Conclusions.....	35
3. Decoherence in Phase Qubits	37
3.1. Bloch Sphere Representation of the State of a Qubit	37
3.2. Phase Precession	39
3.3. Dephasing from Low Frequency Current Noise.....	42
3.4. Relaxation from Dissipation and High Frequency Noise	47
3.5. Conclusions.....	50
4. dc SQUID Phase Qubit	51
4.1. The Need for Isolation from the Bias Leads.....	52
4.2. SQUID Hamiltonian	58
4.3. Cubic Approximation to the SQUID Hamiltonian	61
4.4. SQUID Normal Modes and Their Affect on Current Division	65
4.5. Solving the Full SQUID Hamiltonian with a Finite Difference Method.....	72
4.6. Conclusions.....	83
5. Designing a dc SQUID Phase Qubit for Reduced Decoherence	85
5.1. Designing to Reduce Dissipation.....	85
5.1.1. Eliminating Dielectric Loss	87
5.1.2. Reducing the Loss by Adding a Capacitor Across the Junction	91
5.1.3. Loss from Other Capacitors in the Device.....	98
5.1.4. Isolating the Qubit from the Bias Lines.....	99
5.2. Designing to Reduce Dephasing.....	103
5.3. Qubit Design Parameters and Expected Coherence Time	109

5.4. Device PB9 Layout and Fabrication Considerations	114
5.5. Pictures of Device PB9	125
5.6. Summary	129
6. dc SQUID Phase Qubit Fabrication	
6.1. Preliminary Fabrication Steps and SiN_x Growth at LPS	131
6.2. SF_6 Etch of SiN_x	132
6.3. Sapphire Wafer Dicing	137
6.4. Individual Chip Preparation: O_2 Plasma Cleaning	138
6.5. Preparation of Electron Beam Lithography Resist	141
6.6. Depositing the Aluminum Charge Dissipation Layer	144
6.7. Electron Beam Lithography	149
6.8. Developing the Pattern	165
6.9. Josephson Junction Construction by Double-Angle Evaporation	166
6.10. Lliftoff in Acetone	173
6.11. Summary	176
7. Experimental Setup and Measurement Procedures	
7.1. Measuring the Room Temperature Resistance	177
7.2. Dilution Refrigerator Wiring and Sample Box	182
7.3. Measuring the I-V Characteristics	187
7.4. Measuring the Switching Histogram	191
7.5. Measuring the I- Φ Characteristics	202
7.6. Qubit State Readout Using Current Pulses	207
7.7. Finding the Qubit Resonance Frequency	215
7.8. Measuring the Energy Spectrum	220
7.9. Measuring the Switching Probability Curves (a.k.a.S-Curves)	226
7.10. Measuring the Energy Relaxation Time T_1	231
7.11. Measuring Rabi Oscillations of the Qubit State	235
7.12. Creating Short Microwave Pulses with Sideband Mixing	239
7.13. Measuring the Ramsey Fringes	250
7.14. Conclusions	256
8. Extraction of the Parameters of Device PB9	
8.1. AFM Images of the Josephson Junctions	257
8.2. I-V Curve	258
8.3. I- Φ Characteristics	262
8.4. Energy Spectrum	270
8.5. Fitting the Energy Spectrum	275
8.6. Comparison of Extracted Parameters to Design Parameters	287
8.7. Conclusions	289
9. Analysis of Dissipation in Device PB9	
9.1. Modified Dissipation Model in Normal Modes Picture	291
9.2. Two-Level System Splittings in the Spectrum	299

9.3. Expected T_I from Normal Modes Dissipation Model using Extracted Parameters	302
9.4. Determining T_I from S-Curve Measurements	306
9.5. Determining the Limiting Source of Dissipation Through a Study of the Dependence of T_I on Applied Flux	314
9.6. Future Device Improvements	323
9.7. Conclusions.....	327
 10. Analysis of Dephasing in Device PB9	
10.1. Noise Transfer Function, Sweet Spots, and T_ϕ predictions.....	329
10.2. Measurements of Dephasing in Device PB9	341
10.2.1. Rabi Oscillation Measurement.....	341
10.2.2. Spectroscopic Coherence Time T_2^* vs. Φ_a	344
10.2.3. Ramsey Measurements	353
10.3. Future Improvements on Dephasing.....	357
10.4. Anharmonicity	359
10.5. Conclusions.....	365
 11. Conclusions.....	367
 <u>Appendices</u>	
A. Matlab Code.....	371
 Bibliography	388

List of Tables

5.1 Qubit design parameters for device PB9	108
5.2 Predicted coherence times for device PB9	109
6.1 Approximate radius of electron beam based on spot size.....	161
8.1 Extracted qubit parameters for device PB9	286
9.1 Expected T1's using normal modes model and extract parameters.....	302
11.1 Summary of coherence times for device PB9.....	365

List of Figures

1.1 Comparison of classical to quantum factoring algorithms	3
2.1 Illustration of a Josephson junction	14
2.2 Schematic of current biased Josephson junction in the RCSJ model	19
2.3 Illustration of the tilted washboard potential	22
2.4 Energy levels, transition frequencies, and anharmonicity of a single current biased junction	27
2.5 Illustration of escape of energy states from the potential well	31
2.6 Illustration of the tilted washboard potential	33
3.1 Bloch Sphere representation of the qubit state	40
3.2 Weighting function $W_0(f,t)$ versus frequency and time	44
4.1 Illustration of inductive isolation of the qubit junction	53
4.2 Schematic of a dc SQUID phase qubit	57
4.3 Comparison of f_{0I} in the single junction and normal modes models	64
4.4 Comparison of inductive isolation of the qubit in the single junction and normal modes models	71
4.5 Illustration of the grid for the finite difference method	73
4.6 Validity check of the finite difference method with respect to grid spacing and number of states	78
4.7 Wavefunctions obtained from the finite difference method	80
5.1 Equivalent circuit model for loss in a phase qubit	84
5.2 Equivalent circuit model for a lossy capacitor	87

5.3 Illustration of an interdigitated capacitor	94
5.4 Schematic of dc SQUID phase qubit with R_{eff}	99
5.5 Noise transfer function vs. bias current of a single junction	103
5.6 E-beam pattern of entire device PB9	113
5.7 E-beam pattern of the qubit junction J_1 on device PB9	114
5.8 E-beam pattern of the interdigitated capacitor C_x on device PB9	116
5.9 E-beam pattern of the isolation junction J_2 on device PB9	118
5.10 E-beam pattern and simulation model for the mutual inductance M of the flux bias line I_Φ to the SQUID on device PB9	120
5.11 E-beam pattern of the microwave coupling capacitor C_μ on device PB9 ..	121
5.12 E-beam pattern of the filter inductors L_f on device PB9	122
5.13 E-beam pattern of the filter capacitor C_f on device PB9	124
5.14 Picture of device PB9	125
5.15 Detailed view of dc SQUID phase qubit PB9	126
5.16 Close-up on defect in interdigitated capacitor	128
6.1 Comparison of e-beam lithography on sapphire with and without an Al dissipation layer	143
6.2 Illustration of electron beam current saturation	150
6.3 SEM picture of gold standard	155
6.4 Illustration of electron dosage spacing	160
6.5 Construction of a Josephson junction by double-angle evaporation	165
6.6 Critical current density J_c versus oxygen exposure P^*t	169
6.7 Aluminum covering the step created by SiN capacitor dielectric	170

6.8 Comparison of optical images of successful and failed isolation junctions	173
7.1 Picture of microprobe station.....	176
7.2 Diagram of room temperature resistance measurement setup	177
7.3 Picture of dilution refrigerator	181
7.4 Pictures of Cu patch box, LC filter, and sample box	182
7.5 Diagram of refrigerator control lines	183
7.6 Diagram of I-V measurement setup.....	187
7.7 I-V curve of device SF7.....	190
7.8 Diagram of switching histogram setup	192
7.9 Front panel of <i>HP 33120 Waveform Editor.vi</i>	193
7.10 Front panel of <i>OneSR.vi</i>	197
7.11 Switching histogram measured on device PB9.....	199
7.12 Diagram of I- Φ measurement setup.....	201
7.13 Front panel of <i>Run Scan.vi</i>	202
7.14 I- Φ characteristic of device DS6.....	206
7.15 Diagram of pulsed state readout setup	207
7.16 Current bias waveform used during pulsed readout	209
7.17 Switching histogram obtained with pulsed readout on device PB9.....	212
7.18 Broad resonance search on device PB9	217
7.19 Example of energy spectrum of device PB9.....	223
7.20 Comparison of $ 0\rangle$, $ 1\rangle$, and saturated state S-curves of device PB9.....	228
7.21 S-curves taken during energy relaxation measurement of device PB9.....	232
7.22 S-curves taken during Rabi oscillations of qubit PB9	236

7.23 Diagram of mixed microwave pulse setup.....	239
7.24 Front panel of <i>N8241A Waveform Editor.vi</i>	242
7.25 Power spectrum of mixer output in mixed microwave setup	244
7.26 Illustration of phase variations between AWG and microwave signal	247
7.27 Microwave pulse train for a Ramsey fringe measurement	249
7.28 S-curves of Ramsey fringes on device PB9	253
8.1 AFM images of qubit junction on device PB9.....	257
8.2 I-V curve of device PB9	259
8.3 I- Φ of device PB9	261
8.4 Close-up of I- Φ of device PB9 with fit.....	265
8.5 S-curves of $ 0\rangle$ state vs. applied flux Φ_a for device PB9 with fit	267
8.6 Comparison of I- Φ fits with $L_I = 0.7$ nH and $L_I = 2$ nH for device PB9.....	269
8.7 Energy Spectrum of device PB9	271
8.8 Qubit junction inductance L_{JI} vs. flux bias voltage V_Φ	275
8.9 Energy Spectrum of device PB9 with normal modes fit.....	277
8.10 Energy Spectrum of device PB9 with finite difference method fit.....	279
8.11 Schematic of SQUID normal modes with IDC parasitic inductance L_{fx}	281
8.12 Energy Spectrum of device PB9 with normal modes fit including L_{fx}	284
9.1 Detailed view of splittings in spectrum of device PB9.....	298
9.2 Comparison of predicted T_I 's from known sources of dissipation vs. Φ_a ...	301
9.3 Theoretical prediction for T_I vs Φ_a of device PB9	303
9.4 Observation of anomalous switching behavior in device PB9	305
9.5 S-curves of energy relaxation rate T_I measurement of device PB9.....	306

9.6	Extracted P_I of T_I measurement of device PB9.....	311
9.7	Exponential decay fits to T_I measurement of PB9.....	313
9.8	T_I vs. Φ_a for device PB9	315
9.9	π -pulse power P_μ vs. Φ_a for device PB9.....	317
9.10	Comparison of T_I to microwave line coupling κ for device PB9.....	320
9.11	Photograph of microwave drive line on device PB9	322
9.12	Design layout of improved dc SQUID phase qubit	324
10.1	Noise transfer function (NTF) and T_ϕ prediction for bias current I_b noise .	329
10.2	NTF and T_ϕ prediction for flux bias current I_Φ noise.....	331
10.3	NTF and T_R prediction for $1/f$ flux noise	333
10.4	NTF and T_R prediction for $1/f$ critical current I_{01} noise	335
10.5	NTF and T_R prediction for $1/f$ critical current I_{02} noise	338
10.6	Extracted P_I vs. t for Rabi oscillations measured in device PB9.....	340
10.7	Single period of energy spectrum of device PB9.....	343
10.8	T_2^* vs. Φ_a for device PB9.....	345
10.9	Fit to energy spectrum of device PB9 using parameters that gave I_{01} sweet spots consistent with the T_2^* vs. Φ_a data.....	347
10.10	Fit dephasing model to T_2^* vs. Φ_a data	350
10.11	Extracted P_I vs. t for Ramsey fringe measurement of device PB9.....	353
10.12	Comparison of Ramsey fringes measured at $\Phi_a = 0$ and $\Phi_a = 0.385 \Phi_0$..	354
10.13	NTF and its derivative for flux noise in device PB9	356
10.14	NTF for I_b noise when $I_b = 0$ in device PB9	358
10.15	Rabi oscillations observed in device PB9 at $\Phi_a = 0$	360

10.16 S-curves of the saturated state prepared with increasing P_μ	361
--	-----

Chapter 1

Introduction

1.1 *A Brief History of Quantum Computing*

The use of computers to simulate the physical world is now entirely commonplace. For example, meteorologists break the atmosphere into individual packets of air and use Newton's Laws to model the motion of weather systems around the globe [1]. Microwave engineers break the space around a circuit into many, many, small regions and solve Maxwell's equations to model the electromagnetic fields generated by high speed circuitry [2]. However, the more accurately one wants to describe the system being analyzed, the more small pieces one needs to use and the more computing power and computation time is needed.

When it comes to modeling very small systems, ultimately one must consider quantum mechanical effects. In 1982 Richard Feynman wrote a now famous paper, "Simulating Physics with Computers" [3]. In it he said that the size of a classical computer would need to grow exponentially with the size of the quantum system it is simulating. He then put forth the idea of using a "quantum computer" to model quantum systems effectively. A quantum computer would use bits that are quantum mechanical in nature (known as "qubits") to represent the quantum behavior of each element in the system to be simulated. Feynman argued that a quantum computer would be better suited to simulating the physics of a quantum system than a classical computer.

Quantum computing remained a mostly theoretical pursuit until the mid 1990's. As will be discussed in Chapters 3 and 5 of this thesis, interactions with the outside world result in fluctuations in the quantum state of the computer. These fluctuations lead to

errors in the calculations being performed. In 1995 Peter Shor published a paper proposing an error correction scheme that could be used to overcome small levels of decoherence and allow quantum computing to become attainable [4]. Error correction techniques employ several physical qubits to act as a single logical qubit. Shor proposed that some properties of the logical qubit could be checked to see if an error occurred without revealing the quantum information stored in the logical qubit. The state of the qubit could then be corrected if an error occurred. This scheme became feasible if the bit error rate was less than about 1 in one hundred thousand. This meant that if the gate time was 10 ns, then the qubit's state needed to be coherent (*i.e.* the coherence time) for about $\sim 100 \mu\text{s}$.

Without a real world application, quantum computing would have remained a purely academic pursuit. Modeling quantum systems with a quantum computer will probably mainly be of interest to physicists and chemists. In 1997, though, Shor described a quantum algorithm for decrypting RSA encrypted messages that was much faster than any classical algorithm [5]. Many secure communications now use RSA as the standard form of encryption [6]. RSA encrypted messages are transmitted along with a public passkey, which is a very large number. The key to decrypting an RSA encrypted message is to find the prime factors of the passkey. The strength of the encryption is measured by the number of binary digits n that it takes to represent the passkey. In principle Shor's algorithm can find the prime factors of a large number much faster than any known algorithm on a classical computer. Figure 1.1 shows the time necessary to factor an n bit passkey with Shor's algorithm as opposed to the fastest known classical

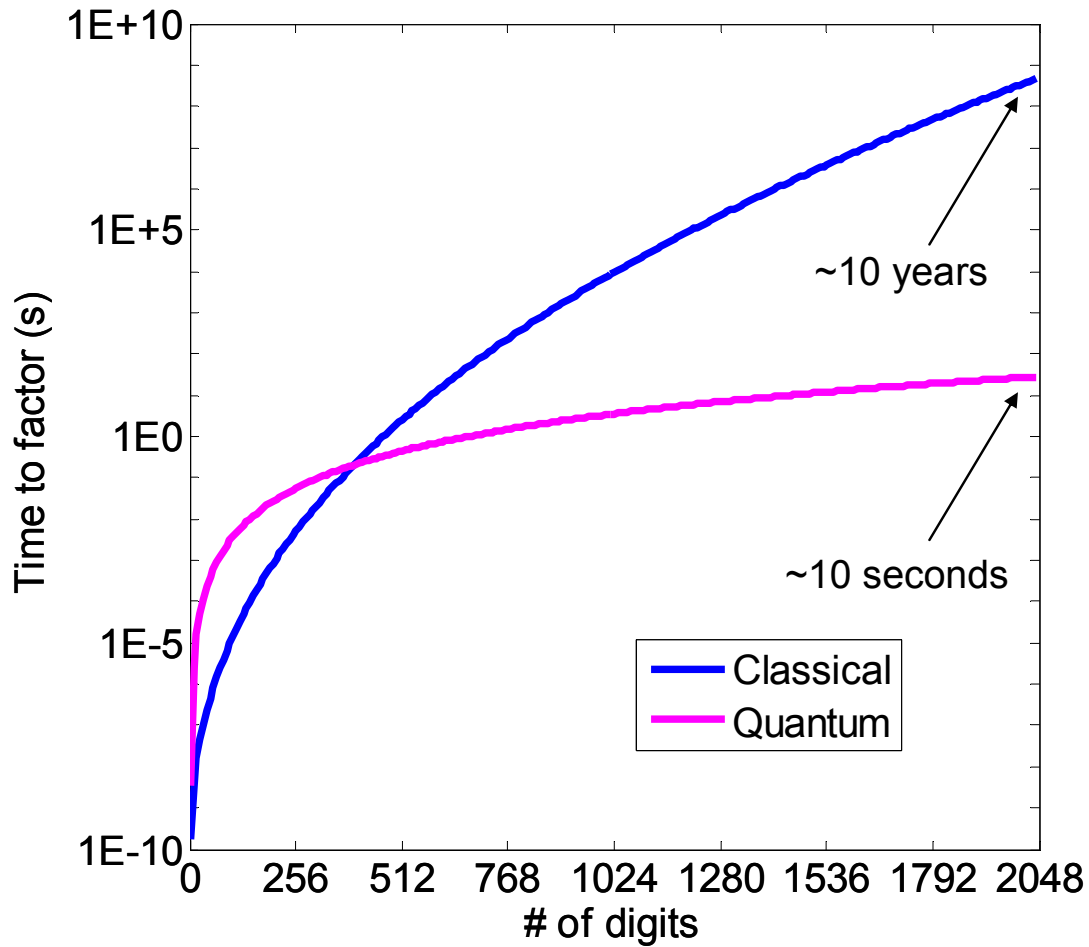


Fig. 1.1 Comparison of the typical time it would take in theory to factor an n digit number with a classical computer and a quantum computer [5,7]. The classical computer is assumed to be a 4 GHz processor running a general number field sieve (GNFS) algorithm [7], and the quantum computer is a 100 MHz quantum processor running Shor's Algorithm. Typical RSA encryption keys are 1024-2048 digits long [6].

algorithm, the general number field sieve (GNFS) [7]. The number of operations needed to factor the passkey is of order:

$$\text{Shor's Algorithm (quantum): } O\left((n \log 2)^3\right), \quad (1.1)$$

$$\text{GNFS (classical): } O\left(\exp\left\{(n \log 2)^{1/3} \left[\log(n \log 2)\right]^{2/3}\right\}\right). \quad (1.2)$$

Typical RSA encryption keys are presently 1024 to 2048 bits long [6]. The state of the art encryption is 4096 bits, which would take a classical computer ~10 million years to factor while in principle a quantum computer could factor it in ~100 seconds assuming a 4 GHz classical computer versus a 100 MHz quantum computer. The implications on national security of having a quantum computer have led the United States government to actively fund a wide variety of quantum computing research.

1.2 The Divincenzo Criteria: Requirements for a Successful Qubit

Researchers have taken many different paths toward developing a quantum computer. The main difference is in the choice of the qubit. This thesis is about a specific type of qubit called a Josephson phase qubit. In a phase qubit, the two lowest energy levels of a Josephson junction, $|0\rangle$ and $|1\rangle$, are used for quantum computation. Other groups have used trapped ions as qubits, with the hyperfine levels as $|0\rangle$ and $|1\rangle$ [8]. Another approach uses the spin on a small semiconducting island with the up and down polarization of the spin as $|0\rangle$ and $|1\rangle$ [9]. Many other approaches are currently being investigated.

Strictly speaking any quantum system can be used as a qubit as long it meets the five Divincenzo criteria [10-12]. The first requirement is that, the Hilbert space of the

computer must be well defined. That is, it must be composed of well known two-level systems (qubits), without the presence of undesirable or unknown quantum states. Second, it must be possible to initialize the computer into its ground state before performing any calculations. Third, the qubits must be well isolated from interactions with the environment that would cause errors to occur during the calculation due to unanticipated altering of the qubits' states known as decoherence. Fourth, it must be possible to perform a complete set of unitary operations on the qubits. In particular, it must be possible to create entangled states of the qubits that would violate Bell's inequality, which is only possible for quantum systems [13]. Finally, it must be possible to measure each qubit by projecting its state onto the qubit basis. If a quantum system satisfies these requirements, then it can be used as a qubit.

In Chapter 2, I explain why and how a Josephson junction satisfies requirements 1 (well defined Hilbert space), 2 (ground state initialization), and 5 (measurement). Although there is much evidence that a Josephson junction qubit has displayed quantum mechanical behavior, it is still yet to be proven that coupled Josephson qubits can violate Bell's inequality (the difficulty is due to inexact control of the qubit state [14]). Another immediate problem that Josephson junction qubits have pertains to the third requirement: isolation. Because they are directly connected to the outside world by control wires, they reside on a physical substrate, and they are macroscopic systems, it is exceptionally difficult to keep them from interacting with their environment. My thesis work has mainly been focused on improving the performance of Josephson phase qubits so that the states retain quantum coherence for a longer time.

1.3 Advantages of Phase Qubits

If Josephson phase qubits have problems with isolation and a key quantum mechanical experiment still needs to be completed using them, why would one pursue them as qubits for constructing a quantum computer. The answer is that phase qubits have some major potential advantages over many other types of qubits: scalability, tunability, and speed.

A functioning quantum computer will require many, many qubits. As I describe in Chapter 2, a phase qubit can be made with only one Josephson junction, a dc current bias line, and a microwave excitation line. All of the other elements of the device are used for isolation from the bias lines and dealing with problems that arise from lossy materials [15-19]. With the right materials and electrical engineering, the footprint of a single qubit could be extremely small, at least in principle. Also, it has been shown that phase qubits can be coupled by simply connecting them together with wires and capacitors [20,24]. In contrast, single trapped ion qubits have been shown to have much longer coherence times than phase qubits, but each qubit requires its own trap. The considerations for expanding a trapped ion quantum computer beyond just a handful of qubits are decidedly more complex compared to phase qubits. This suggests that it will be easier to attain high bit densities with phase qubits than it will be for other non-solid state qubits.

Because it will be necessary to perform operations on specific qubits or groups of qubits, it must be possible to turn the coupling between qubits on and off. When phase qubits are coupled together with capacitors, their behavior is similar to that of coupled oscillators [20]. In particular, even if oscillators are coupled together, they will not

interact strongly unless their resonant frequencies are nearly the same [21]. The resonant frequency of a phase qubit can be changed by simply changing the current flowing through it, which makes it possible to easily control the coupling between phase qubits. Other types of superconducting qubits (*i.e.* charge and flux qubits [22,23]), which enjoy the same advantage of scalability can be tuned, but typically they are operated at a fixed operating frequency where a sweet spot in their energy levels provides them with enhanced isolation from the environment. In practice, other parameters need to be adjusted to tune a charge or flux qubit, or a tunable coupling element can be used [25]. Nevertheless, the tunability of the energy levels of phase qubits allows for comparatively easy control of the coupling between specific qubits.

Finally, it is possible to change the state of a phase qubit relatively quickly. According to Eq. 1.1, Shor's algorithm reduces the time it takes to factor a number by reducing the number of operations that the computer must perform to complete the task. If the required number of operations can be completed in a shorter amount of time, then the qubit states do not need to be coherent for as long. As I will show in Sect. 2.4, the time it takes to change a phase qubit's state depends on the microwave drive power applied. It has been shown that it can take ~ 1 ns to flip the qubit state with a strong enough drive [26]. This is because a microwave control line can deliver a large power compared to many other quantum systems. In addition, when multiple qubits need to be entangled for a quantum operation, they need to be strongly coupled in order to perform the operation quickly. In cavity QED experiments, for example, it is a challenge to get strong enough coupling between the atoms in cavities and their optical frequency drive

field or other atoms. The fact that phase qubits have the ability to perform quantum operations relatively quickly ultimately reduces the demands on their state coherence.

Even though phase qubits have problems with isolation and have yet to complete a Bell's state experiment, the potential advantages make them worth exploring further. In particular, it is interesting to determine the specific causes of decoherence through theoretical and experimental means. If a phase qubit can be designed to limit or eliminate these causes of decoherence, then it should be possible to create a phase qubit that satisfies the third Divincenzo criteria. This is the main motivation for my research on phase qubits.

1.4 *Overview of My Thesis*

In this thesis, I describe my attempts at reducing decoherence in a dc SQUID phase qubits through circuit design, choice of materials, and the use of nano-scale fabrication techniques. Towards this end I will present a model for estimating the effects of specific sources of decoherence on the coherence time of a phase qubit. Although I designed my device to have a coherence time of about 8 μ s according to this model, the device only showed a coherence time of about 110 ns. Through a meticulous study of how the coherence time of the device changed with the biasing conditions, I was able to identify the dominant cause of decoherence to be dissipation from coupling to the microwave drive line. I also discovered a new type of sweet spot in the transition spectrum of the device. Because of the extreme parameter regime of the device design, I found that the resonances of the qubit were not simply resonant modes of a single junction, but mixed modes of the two junctions used in the SQUID phase qubit. Through

computer simulation of the energy levels, I discovered biasing conditions for the qubit that made it insensitive to fluctuations in bias current, fluctuations in magnetic flux, and critical current fluctuations of the Josephson junctions. These results provided a way to estimate the contributions of different sources of decoherence and opened a new path for Josephson phase qubits to attain longer coherence times.

In Chapter 2, I provide background on Josephson junctions and phase qubits. I describe how current flows through a Josephson junction without loss and how the current-voltage characteristics lead to the possibility of using a junction as an anharmonic quantum oscillator. I also describe how to use microwaves to manipulate the qubit state and how the $|0\rangle$ and $|1\rangle$ state can be distinguished. Finally, I give an example of the basic operation of a Josephson junction as a phase qubit.

In Chapter 3, I discuss dephasing and dissipation and how they both lead to decoherence. Dephasing is caused by relatively slow fluctuations of the energy levels of the qubit. In phase qubits, dephasing can be caused by low frequency fluctuations of the bias current, for example. Dissipation occurs when unwanted transitions occur from the excited state $|1\rangle$ to the ground state $|0\rangle$.

The dc SQUID phase qubit is the subject of Chapter 4. I explain how a second junction added to the circuit can act as a low inductance shunt that isolates the qubit junction from noise on the bias leads. Adding this second junction adds another degree of freedom to the Hamiltonian that perturbs the energy eigenvalues and eigenstates of the system. I discuss the parameter regime in which the qubit can still be considered a single junction and the regime where the qubit eigenstates are mixed modes of the two junctions. Using a coordinate transformation to the normal modes of the system, I also show that the

isolation provided by the second junction still exists even when the modes are mixed. I also present a method for finding the qubit's energy eigenvalues and eigenstates by diagonalizing the Hamiltonian of the system with a finite difference method.

In Chapter 5, I model the effects on dc SQUID phase qubit of specific sources of dissipation and dephasing. I modeled dissipation as an effective resistance R_{eff} across the junction and the relaxation time T_1 was found from $T_1 = R_{eff}C$ where C is the capacitance shunting the qubit junction. Current fluctuations coming from Johnson-Nyquist noise on the bias lines, $1/f$ flux noise, and $1/f$ critical current fluctuations all produce dephasing. These models were used to choose the parameters of the device and provided design estimates for the coherence time T_2 and the relaxation time T_1 . Chapter 5 also contains a description of the physical layout and photographs of the device I built, dc SQUID phase qubit PB9.

In Chapter 6, I give details of the fabrication process for making a phase qubit on a sapphire substrate with e-beam lithography. Chapter 7 gives a detailed description of the procedures I used for measuring the qubit.

In Chapter 8, I compare the measured I-V curve, I- Φ characteristics, and the energy spectrum of device PB9 using both a normal modes model and a finite difference method to find the transition frequencies of the qubit. I then use these models to extract the device parameters.

In Chapter 9, I discuss the dissipation I observed in device PB9. First, I describe the alterations I made to the dissipation model to account for the mixing of the modes of the two junctions in the SQUID. I then give the theoretical prediction for T_1 and identify the probable leading cause of dissipation. I also describe the model for extracting the

probability P_I of being in the excited state (from the S-curves of the device) and display the best T_I observed during an energy relaxation measurement of the qubit. In the second half of Chapter 9, I describe how I determined that dissipation from the microwave line was limiting the relaxation time of the device and give some suggestions for future improvements in the design.

A surprising result from device PB9 was the discovery of sweet spots in the energy spectrum, which I discuss in Chapter 10. I describe how I used the model for finding the qubit energies to also find the noise transfer functions for the various parameters of the system (e.g. $\frac{df_{01}}{dI_b}$). According to this analysis, device PB9 had sweet spots (zeros in the transfer functions) for current bias noise, flux noise, and critical current noise. Using these noise transfer functions, I discuss the measured dephasing times of device PB9 as obtained from Rabi oscillation, spectral peak width, and Ramsey fringe measurements. A comparison of the measured spectroscopic coherence time to the noise transfer function revealed that shot-to-shot critical current noise was the main cause of spectral broadening. I briefly discuss the implications that this has on the future of phase qubits. Also, I discuss the anharmonicity of the qubit in device PB9. Although the theoretical models of the system predict very little anharmonicity, the experimental was inconclusive.

Finally, I conclude in Chapter 11 with a summary of my main findings on the coherence of dc SQUID phase qubit PB9. Appendix A contains some of the Matlab programs I used during my research.

Chapter 2

Josephson Junctions as Qubits

In order to build a quantum computer, one needs to combine together many qubits. Not every quantum system is a qubit. Ideally, one wants a two-level system.

An example of a system that cannot be used as a qubit is the harmonic oscillator. One can build a harmonic oscillator out of circuit elements by placing an inductor L in parallel with a capacitor C . In this case the energy in the LC-resonator is stored in the electromagnetic field and the resonant frequency is

$$\omega_0 = \frac{1}{\sqrt{LC}}. \quad (2.1)$$

According to Schrödinger's equations, the eigenstates of a harmonic oscillator have energy

$$E_n = (n+1/2)\hbar\omega_0, \quad (2.2)$$

where \hbar is Planck's constant, n is a non-negative integer, and ω_0 is the resonant frequency of the oscillator times 2π . The LC resonator cannot be used as a qubit because it has multiple, evenly spaced levels. Driving the system electrically will yield a final state with amplitude to be in more than two levels. That is to say, the state of the system cannot be controlled to remain in a space spanned by just the $n = 0$ and $n = 1$ levels [10-12].

To get around this problem, we use a non-linear oscillator based on a Josephson junction. In 1962, Brian Josephson predicted that there could be coherent tunneling of Cooper pairs between two superconducting electrodes that were separated by a small distance (~ 1 nm) [27]. Such a structure is now known as a Josephson junction (JJ). In

this chapter, I show how a Josephson junction can be used as an “anharmonic” LC oscillator: the C comes from the close proximity of the two superconductors and the L comes from the current-voltage relation of the junction. The junction can store energy in much the same way as an LC resonator, but the energy levels of the junction are unevenly spaced. These energy states exist in an effective potential energy surface (the tilted washboard potential). I also discuss the use of microwaves to drive transitions between the energy states, and how the state of the qubit can be measured.

Note that the discussion in this chapter is concerned with the equations of motion, Hamiltonian, and energy levels of an ideal single-junction phase qubit. I will discuss the two-junction dc SQUID phase qubit and its Hamiltonian in Chapter 4.

2.1 Derivation of the Josephson Relations and Josephson Inductance

The superconducting state of a material is characterized by a complex order parameter $\Psi = \sqrt{n}e^{i\phi}$, where n is the number density of Cooper pairs and ϕ is a coherent phase that in general depends on position and time [28]. A superconducting material is also characterized by the London penetration depth λ and the coherence length ξ_0 . λ is the characteristic distance that a magnetic field can penetrate into a superconductor [28]. ξ_0 is the characteristic length scale over which the order parameter Ψ can change [28]. For aluminum (Al) $\lambda \approx 16$ nm and $\xi_0 \approx 1.6$ μ m [28]. Since $\xi_0/\lambda \gg 1$, clean aluminum is a Type I superconductor [28]. If two superconductors are placed very near to each other, then their order parameters can overlap due to tunneling of Cooper pairs (see Fig. 2.1). Josephson proposed that a current of Cooper pairs could tunnel between two such separated superconductors with zero voltage drop across the interface. His theory was

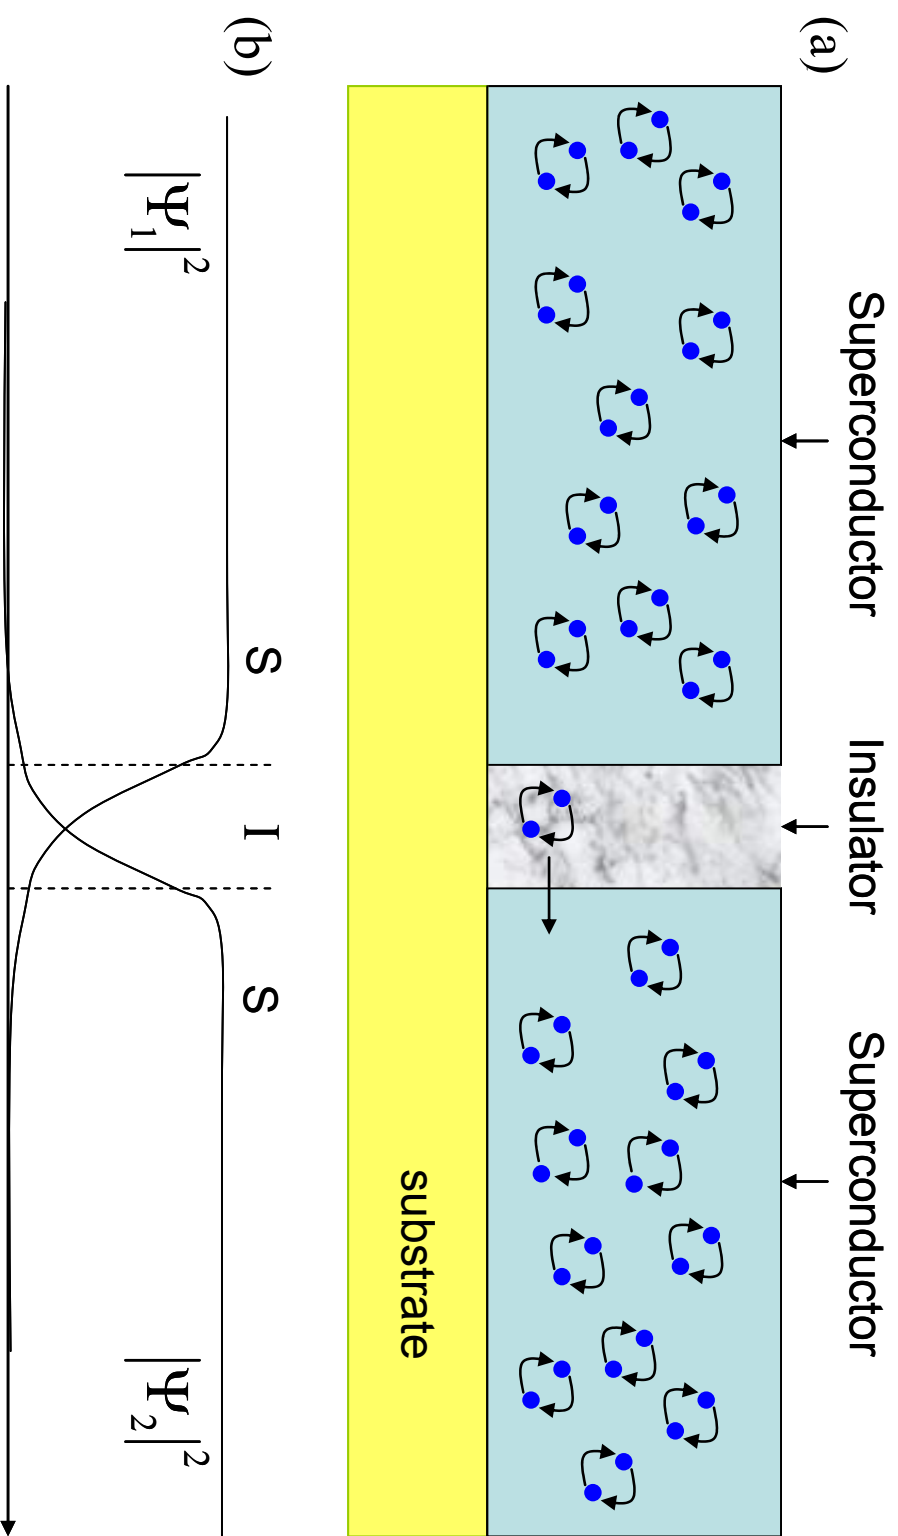


Fig. 2.1 (a) Illustration of tunneling of Cooper pairs between two superconductors. (b) The order parameters of the two superconductors overlap in the barrier, which allows the Cooper pairs to pass through the barrier without breaking apart. This configuration is known as an S-I-S Josephson junction.

quickly confirmed by experiment [29,30]. In practice, the two superconductors are usually separated by a thin insulator or normal metal, and the structure is called a Josephson junction.

The current-voltage characteristics of a Josephson junction can be obtained by considering the amplitude and phase of the order parameter on either side of the barrier. Because of the overlap of the order parameters, Ψ_1 on one side and Ψ_2 on the other, the order parameter of each superconductor will be affected by the one on the other side of the junction. Feynman described a simple model of the junction in which Ψ_1 and Ψ_2 obey coupled time-dependent Schrödinger equations [31]

$$i\hbar \frac{\partial \Psi_1}{\partial t} = E_1 \Psi_1 + \kappa \Psi_2 \quad \text{and} \quad i\hbar \frac{\partial \Psi_2}{\partial t} = E_2 \Psi_2 + \kappa \Psi_1 \quad (2.3, 2.4)$$

where E_i is the energy of pairs on the left and κ is a constant that represents the tunneling of pairs between the two superconductors. Since each wavefunction has the form, $\Psi_i = \sqrt{n_i} e^{i\phi_i}$, if I take the time derivative of the left hand side of Eqs. 2.3 and 2.4 and multiply through by Ψ_i^* , I find

$$\frac{i\hbar}{2} \dot{n}_1 - \hbar n_1 \dot{\phi}_1 = E_1 n_1 + \kappa \sqrt{n_1 n_2} e^{i(\phi_2 - \phi_1)} \quad (2.5)$$

$$\frac{i\hbar}{2} \dot{n}_2 - \hbar n_2 \dot{\phi}_2 = E_2 n_2 + \kappa \sqrt{n_1 n_2} e^{i(\phi_1 - \phi_2)}. \quad (2.6)$$

The energies E_1 and E_2 can be understood as the energies of a pair in the two superconductors. If there is a voltage difference V between the two sides, then the energy of a pair will change by e^*V when it moves from one side to the other, where e^* is the charge of the charge carriers. In a superconductor, the carriers are Cooper pairs, so $e^* =$

2e. Since the zero voltage point can be chosen arbitrarily, for convenience we can choose it such that

$$E_1 = -E_2 = e^*V/2 = eV. \quad (2.7)$$

Eqs. 2.5 and 2.6 can then be simplified to

$$\dot{n}_1 + 2i\dot{\phi}_1 = -\frac{2ieV}{\hbar} - \frac{2i\kappa}{\hbar} \sqrt{\frac{n_2}{n_1}} e^{-i\gamma} \quad (2.8)$$

$$\dot{n}_2 + 2i\dot{\phi}_2 = \frac{2ieV}{\hbar} - \frac{2i\kappa}{\hbar} \sqrt{\frac{n_1}{n_2}} e^{i\gamma}, \quad (2.9)$$

where $\gamma = \phi_2 - \phi_1$ is the phase difference across the junction. Taking the real part of Eqs.

2.8 and 2.9, one finds

$$\dot{n}_1 = \frac{2\kappa}{\hbar} \sqrt{\frac{n_2}{n_1}} \sin(\gamma) \quad (2.10)$$

$$\dot{n}_2 = -\frac{2\kappa}{\hbar} \sqrt{\frac{n_1}{n_2}} \sin(\gamma). \quad (2.11)$$

Since the current is $I = 2e\dot{n}_1 = -2e\dot{n}_2$, Eqs. 2.10 and 2.11 describe the flow of Cooper pairs from one side to the other.

If the two superconductors are of the same material and the current is low enough, then $n_1 \approx n_2$. Eqs. 2.10 and 2.11 can then be rewritten in the form

$$I = I_0 \sin(\gamma) \quad (2.12)$$

where I_0 is the maximum value for the current that can be carried by pairs, or the “critical current.” Eq. 2.12 is the dc Josephson relation and it describes the current flow through a Josephson junction in terms of the phase difference between the two superconductors. Note that the critical current is directly dependent on the size of the coupling constant κ .

From the imaginary parts of the Eqs. 2.8 and 2.9, we get:

$$\dot{\phi}_1 = -\frac{eV}{\hbar} - \frac{\kappa}{\hbar} \sqrt{\frac{n_2}{n_1}} \cos(\gamma) \quad (2.13)$$

$$\dot{\phi}_2 = \frac{eV}{\hbar} - \frac{\kappa}{\hbar} \sqrt{\frac{n_1}{n_2}} \cos(\gamma). \quad (2.14)$$

Taking the difference of Eqs. 2.13 and 2.14 and again assuming that $n_1 \approx n_2$ we find

$$\dot{\phi}_2 - \dot{\phi}_1 = \frac{d\gamma}{dt} = \frac{2eV}{\hbar}, \quad (2.15)$$

or

$$V = \frac{\hbar}{2e} \frac{d\gamma}{dt}. \quad (2.16)$$

Because the voltage depends on the time-variation of the phase difference, Eq. 2.16 is known as the “ac Josephson relation.” In fact, Eq. 2.16 is now used to define the volt [32].

Note that Eq. 2.12 says that a dc current can flow through the junction up to some critical current, but Eq. 2.16 says if there is a voltage across the junction then the phase γ must be changing in time. This is similar to the behavior of a conventional inductor where

$$V = L \frac{dI}{dt}. \quad (2.17)$$

For a junction, if we take the time derivative of Eq. 2.12, then we get

$$\frac{dI}{dt} = I_0 \frac{d\gamma}{dt} \cos \gamma \quad (2.18)$$

assuming I_0 is constant. Rearranging Eq. 2.16 and substituting into Eq. 2.18 we find

$$\frac{dI}{dt} = -I_0 \frac{2e}{\hbar} V \cos \gamma. \quad (2.19)$$

Solving for V gives an equation similar to Eq. 2.17

$$V = \left(\frac{\hbar}{2e} \frac{1}{I_0 \cos \gamma} \right) \frac{dI}{dt}. \quad (2.20)$$

The prefactor is referred to as the Josephson inductance

$$L_J \equiv \frac{\hbar}{2e} \frac{1}{I_0 \cos \gamma}. \quad (2.21)$$

Note also that $\frac{\hbar}{2e} = \frac{\Phi_0}{2\pi}$, where $\Phi_0 = 2.07 \cdot 10^{-15}$ Wb is the magnetic flux quantum.

By using the identity $\cos \gamma = \sqrt{1 - \sin^2 \gamma}$ and substituting in for $\sin \gamma$ from Eq. 2.12, we can write the current dependence of L_J explicitly

$$L_J = \frac{\Phi_0}{2\pi} \frac{1}{I_0 \sqrt{1 - (I/I_0)^2}}. \quad (2.22)$$

Thus, a junction acts as a current-dependent inductor. In fact, since the geometry of the tunnel barrier gives a capacitance, the junction will act as a non-linear LC resonator. The inductance of the junction depends on the current flowing through it, and this is what produces energy levels that are unevenly spaced in the phase qubit.

2.2 RCSJ Model and the Tilted Washboard Potential

A Josephson junction can be modeled as a circuit consisting of a resistor, a capacitor, and an ideal Josephson junction connected in parallel (see Fig. 2.2). The resistor must be included because, in a real junction, processes exist that can cause dissipation. Such dissipative processes include quasi-particle tunneling, pair breaking events, dielectric loss, and coupling to other circuit elements. In the resistively and




Fig. 2.2 A current biased Josephson junction in the RCSJ Model.

capacitively shunted junction (RCSJ) model [28,33-36], the resistor accounts for all the possible loss channels.

The equations of motion for a current biased junction can be found by applying Kirchhoff's rules to the RCSJ circuit. From Fig. 2.2, the current through the individual elements in the circuit will be

$$I_R = \frac{V}{R} \quad (2.23)$$

$$I_C = \frac{dQ}{dt} = C \frac{dV}{dt} \quad (2.24)$$

$$I_J = I_0 \sin \gamma \quad (2.25)$$

where I_R is the current through the resistor, I_C is the current through the capacitor, Q is the charge on the capacitor, I_0 is the critical current, γ is phase difference across the junction, and V is the voltage drop across the three elements in parallel. Substituting in for the ac Josephson relation (Eq. 2.16) and using as $\Phi_0 = h/2e$ gives I_R and I_C completely in terms of the phase difference across the junction:

$$I_R = \frac{1}{R} \frac{\Phi_0}{2\pi} \dot{\gamma} \quad (2.26)$$

$$I_C = C \frac{\Phi_0}{2\pi} \ddot{\gamma}. \quad (2.27)$$

If the current source is supplying a bias current I_b , the conservation of current gives $I_b = I_R + I_C + I_J$, and we find

$$C \frac{\Phi_0}{2\pi} \ddot{\gamma} + \frac{1}{R} \frac{\Phi_0}{2\pi} \dot{\gamma} + I_0 \sin \gamma = I_b. \quad (2.28)$$

This is the equation of motion of a junction in the RCSJ model. If we now remove the damping term (*i.e.* R goes to infinity) we find

$$C\left(\frac{\Phi_0}{2\pi}\right)^2 \ddot{\gamma} = \frac{\Phi_0}{2\pi} (I_b - I_0 \sin \gamma) \quad (2.29)$$

Note that this equation is analogous to the equation of motion of an object in a tilted washboard potential. The factor $C\left(\frac{\Phi_0}{2\pi}\right)^2$ is like a mass, $\ddot{\gamma}$ is like an acceleration, and the right hand side is like a force. By integrating this “force”, we can define a potential with respect to γ .

$$U(\gamma) = -\frac{\Phi_0}{2\pi} (I_b \gamma + I_0 \cos \gamma) \quad (2.30)$$

For $I_b = 0$, Eq. 2.30 gives a periodic potential, a series of identical potential wells with depth set by the critical current. Applying bias current tilts the potential such that the barrier height of one side of each well decreases while the opposite side increases. As I approaches I_0 , the potential is severely tilted, and the potential wells become very shallow (see Fig. 2.3). At $I = I_0$, the local minima disappear. This potential is called the “tilted washboard potential.”

2.3 The Josephson Junction as a Quantum Anharmonic Oscillator

Up to this point, my discussion of the Josephson junction has been classical. However, at low enough temperature, a junction must be considered as a quantum mechanical system [37]. The Hamiltonian of a particle moving in a tilted washboard potential U can be written as

$$H = \frac{p^2}{2m} + U(\gamma) = \frac{p^2}{2m} - \frac{\Phi_0}{2\pi} (I_b \gamma + I_0 \cos \gamma) \quad (2.31)$$

where $m \equiv C\left(\frac{\Phi_0}{2\pi}\right)^2$ is the effective mass and p is the canonical momentum.




Fig. 2.3 The tilted washboard potential U versus the phase difference γ when the bias current is somewhat smaller than the critical current. The two lowest lying energy states, $|0\rangle$ and $|1\rangle$, are used as qubit states.

The cosine term in Eq. 2.31 makes it somewhat difficult to find the eigenstates and eigenvalues of the Hamiltonian [37], but we can find approximate states and eigenenergies by expanding the potential around its minimum (see Sect. 2.3 of [37] and [38,39]). To do this, we assume that $\gamma = \gamma_0 + \delta\gamma$ where γ_0 is the value of γ at the minimum of one of the wells of the potential U and $\delta\gamma \ll 2\pi$. We can start by expanding the cosine using a trigonometric identity

$$\cos(\gamma_0 + \delta\gamma) = \cos(\gamma_0)\cos(\delta\gamma) - \sin(\gamma_0)\sin(\delta\gamma). \quad (2.32)$$

Expanding to third order in $\delta\gamma$ then yields

$$H = \frac{p^2}{2m} + \frac{\Phi_0}{2\pi} \left[\frac{1}{2} I_0 \cos \gamma_0 (\delta\gamma)^2 + \frac{1}{6} I_0 \sin \gamma_0 (\delta\gamma)^3 \right]. \quad (2.33)$$

Here I have used the fact that $dU/d\gamma$ vanishes at γ_0 and left out terms that do not depend on $\delta\gamma$ since they will not alter the transition frequency spectrum.

The first term in the potential can be rearranged using the definition of the Josephson inductance (Eq. 19) and $m \equiv C \left(\frac{\Phi_0}{2\pi} \right)^2$. One finds

$$\frac{1}{2} \left(\frac{\Phi_0}{2\pi} \right) I_0 \cos \gamma_0 (\delta\gamma)^2 = \frac{1}{2} \frac{1}{L_J} \left(\frac{\Phi_0}{2\pi} \right)^2 (\delta\gamma)^2 = \frac{1}{2} \frac{1}{L_J C} m (\delta\gamma)^2 = \frac{1}{2} m \omega_p^2 (\delta\gamma)^2, \quad (2.34)$$

which is the familiar form of the potential energy of a harmonic oscillator. Here $\omega_p = 1/\sqrt{L_J C}$. The Hamiltonian then can be put in the form

$$H = H_0 + \frac{1}{6} \frac{\Phi_0}{2\pi} I_0 \sin \gamma_0 (\delta\gamma)^3 \quad (2.35)$$

where H_0 is the unperturbed harmonic oscillator Hamiltonian such that

$$H_0 = \frac{p^2}{2m} + \frac{1}{2} m \omega_p^2 (\delta\gamma)^2. \quad (2.36)$$

For a harmonic oscillator, one finds eigenstates $|n\rangle$ which obey

$$H_0|n\rangle = \hbar\omega_p\left(n + \frac{1}{2}\right)|n\rangle \quad (2.37)$$

where $n = 0, 1, 2, \dots$. The second term in Eq. 2.35 is a cubic perturbation to the harmonic oscillator Hamiltonian. Note that the strength of the perturbation depends on the current flowing through the junction, $I_0 \sin \gamma_0$. If the current is zero, then the Hamiltonian is, to third order, that of a harmonic oscillator. On the other hand, if $I \sim I_0$, then $\sin \gamma \sim 1$ and the perturbation brings some anharmonicity into the energy levels of the system.

Perturbation theory gives the energy of a perturbed eigenstates as [40]

$$E_n = E_n^{(0)} + \lambda \langle n | (\delta\gamma)^3 | n \rangle + \lambda^2 \sum_{k \neq n} \frac{\langle n | (\delta\gamma)^3 | k \rangle \langle k | (\delta\gamma)^3 | n \rangle}{E_n^{(0)} - E_k^{(0)}} + \dots \quad (2.38)$$

where I have adopted the notation used by Sakurai such that $E_n^{(0)}$ is the energy of the n^{th} state of the unperturbed Hamiltonian H_0 and $\lambda \equiv \frac{1}{6} \frac{\Phi_0}{2\pi} I_0 \sin \gamma_0$ is the strength of the perturbation. $\delta\gamma$ can be written in terms of creation and annihilation operators of the harmonic oscillator as [40,41]

$$\delta\gamma = \sqrt{\frac{\hbar}{2m\omega_p}} (a^\dagger + a). \quad (2.39)$$

Thus, we find

$$(\delta\gamma)^3 = \left(\frac{\hbar}{2m\omega_p} \right)^{3/2} \left((a^\dagger)^3 + a^3 + a^2 a^\dagger + a (a^\dagger)^2 + a^\dagger \hat{n} + a \hat{n} + \hat{n} a^\dagger + \hat{n} a \right) \quad (2.40)$$

where a and a^\dagger act on the harmonic oscillator states such that

$$a|n\rangle = \sqrt{n}|n-1\rangle \quad (2.41)$$

$$a^\dagger |n\rangle = \sqrt{n+1} |n+1\rangle \quad (2.42)$$

and $\hat{n} = a^\dagger a$ is the number operator. \hat{n} acts on the state such that

$$\hat{n} |n\rangle = n |n\rangle. \quad (2.43)$$

It is easy to see from Eq. 2.40 that the second term in Eq. 2.38 is always zero. Also, from Eq. 2.38 and 2.40 we can see the energy of the n^{th} perturbed state depends on the energy of unperturbed states $n \pm 3$ and $n \pm 1$. This suggests that the approximation will not be very good unless there are at least 4 states bound in the potential well, which is not the case for bias currents extremely close to the critical current. Finally, it is important to remember that $[a, a^\dagger] = 1$ not 0, so the operators need to be applied in order.

From Eqs. 2.38-2.41, the perturbed energy of the ground state of a Josephson junction is then

$$E_0 = \frac{\hbar\omega_p}{2} + \lambda^2 \beta^3 \left(\frac{\langle 0 | a \hat{n} | 1 \rangle}{-\hbar\omega_p} + \frac{\langle 0 | a^3 | 3 \rangle}{-3\hbar\omega_p} \right) \quad (2.44)$$

where $\beta \equiv \sqrt{\frac{\hbar}{2m\omega_p}}$. The matrix elements can be evaluated using Eq. 2.41-43 and one finds [41],

$$E_0 = \frac{\hbar\omega_p}{2} + \frac{\lambda^2 \beta^3}{\hbar\omega_p} \left(-1 - \frac{\sqrt{6}}{3} \right). \quad (2.45)$$

Similarly, E_1 and E_2 are [41]

$$E_1 = \frac{3\hbar\omega_p}{2} + \frac{\lambda^2 \beta^3}{\hbar\omega_p} \left(1 - \sqrt{2} - \frac{2\sqrt{6}}{3} \right) \quad (2.46)$$

$$E_2 = \frac{5\hbar\omega_p}{2} + \frac{\lambda^2 \beta^3}{\hbar\omega_p} \left(2\sqrt{2} - 2\sqrt{3} - \frac{2\sqrt{15}}{3} \right). \quad (2.47)$$

Figure 2.4(a) shows the energy levels as a function of bias current for a junction with $I_0 = 150$ nA and $C = 100$ fF. Note that E_1 and E_2 get closer together as the cubic perturbation gets stronger with increasing I_b . Fig. 2.4(b) shows the transition frequencies between each level

$$f_{ij} = \frac{E_j - E_i}{h}. \quad (2.48)$$

Note that the $|0\rangle$ to $|1\rangle$ transition frequency f_{01} gets lower as the current increases because the Josephson inductance is increasing. Also, f_{01} is in the range of 1-20 GHz which makes it accessible to microwave signal generators. $f = 5$ GHz corresponds to a temperature $T = hf/k_B = 250$ mK, so if the device is cooled to 20 mK, it will relax to the ground state and, to a very high probability, not be thermally excited. Fig. 2.4(c) shows $\Delta f = f_{01} - f_{12}$ as the current is increased. Note that even at relatively low bias currents ($I_b = 0.8I_0$) the difference between the two transition frequencies is ~ 100 MHz. Thus, if the resonances are sharper than 100 MHz, a low-power microwave signal at f_{01} should not drive transitions from $|1\rangle$ to $|2\rangle$.

2.4 State Manipulation with a Microwave Drive

In order to use a Josephson junction as a qubit, one must be able to prepare controlled superpositions of the $|0\rangle$ and $|1\rangle$ states. Fig. 2.4 shows that for reasonable junction parameters the energy splitting between the states corresponds to frequencies in the range of 1-20 GHz; this range is accessible to standard microwave sources.

To see how a microwave signal can drive a transition, note that if an ac current is applied to the junction, then the Hamiltonian for the system is

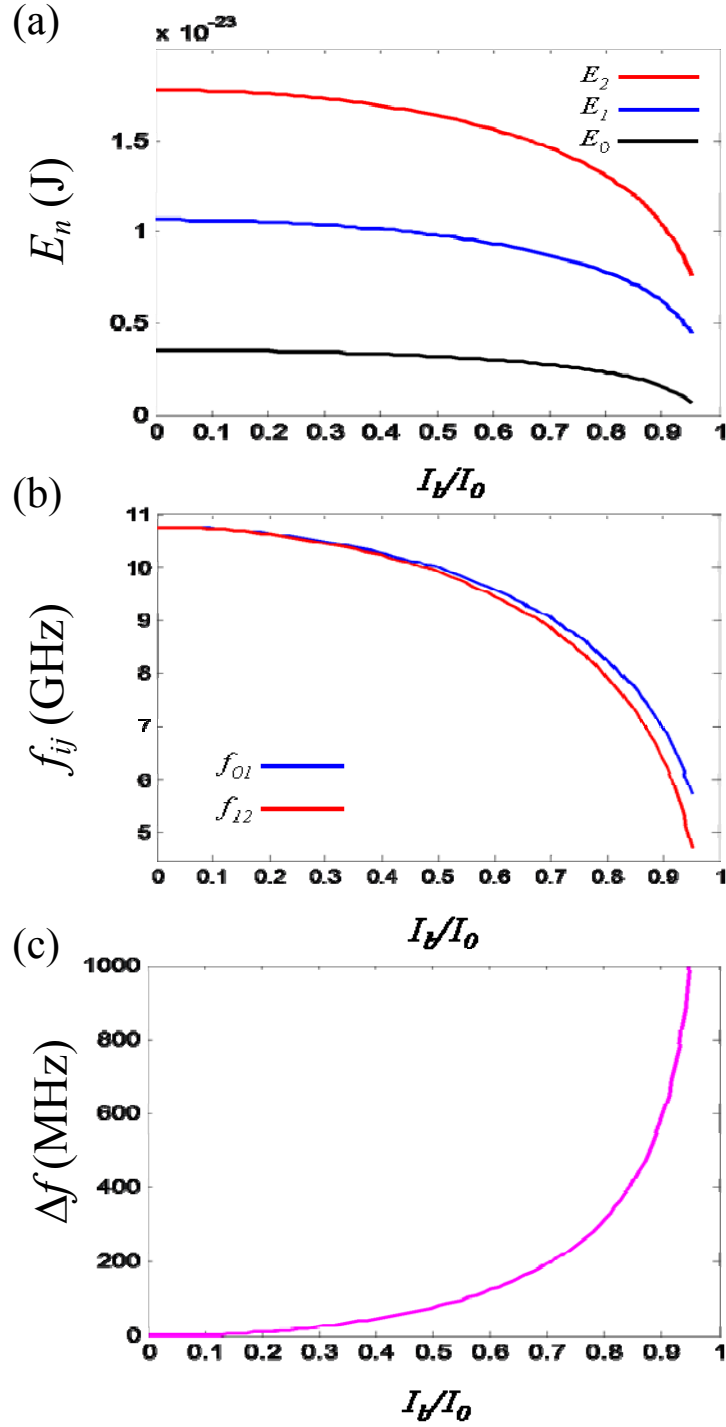


Fig. 2.4 (a) Energies of the three lowest energy levels of a junction with $I_0 = 150$ nA and $C = 0.1$ pF according to the cubic approximation. (b) $0 \rightarrow 1$ and $1 \rightarrow 2$ transition frequencies. Note the larger anharmonicity at larger biases. (c) Difference Δf between f_{01} and f_{12} as the current increases.

$$H = \frac{p^2}{2m} + \frac{\Phi_0}{2\pi} \left\{ I_0 \cos \gamma + \gamma \left[I_b + I_\mu \cos(\omega t + \phi_0) \right] \right\} \quad (2.49)$$

where I_μ is the amplitude of the ac drive applied to the system, ω is the frequency, and I am assuming that the microwave drive has a constant phase ϕ_0 .

Treating the effects of the microwave signal as a small perturbation, we can write

$$\hat{H} = \hat{H}_0 + \hat{H}_\mu \quad (2.50)$$

where \hat{H}_0 is now given by Eq. 2.31, and has eigenvalues of E_0 , and E_1 , and \hat{H}_μ describes the qubits interaction with the microwave signal. In the $\{|0\rangle, |1\rangle\}$ basis we find

$$\hat{H}_\mu = \frac{\Phi_0}{2\pi} I_\mu \cos(\omega t) \begin{pmatrix} \langle 0|\gamma|0\rangle & \langle 0|\gamma|1\rangle \\ \langle 1|\gamma|0\rangle & \langle 1|\gamma|1\rangle \end{pmatrix}. \quad (2.51)$$

The matrix elements can be evaluated by approximating $|0\rangle$ and $|1\rangle$ as harmonic oscillator states. γ is still given by Eq. 2.39 in terms of raising and lowering operators, and this makes it easy to evaluate the matrix elements. In the $\{|0\rangle, |1\rangle\}$ basis the interaction Hamiltonian then takes the form

$$\hat{H}_\mu = \frac{\Phi_0}{2\pi} I_\mu \cos(\omega t) \sqrt{\frac{\hbar}{2m\omega_{01}}} \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}. \quad (2.52)$$

Substituting $m \equiv C \left(\frac{\Phi_0}{2\pi} \right)^2$, the perturbation Hamiltonian becomes

$$\hat{H}_\mu = I_\mu \cos(\omega t) \sqrt{\frac{\hbar}{2C\omega_{01}}} \sigma_x \quad (2.53)$$

where

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad (2.54)$$

is the x Pauli matrix.

From the form of the perturbation, we see that this Hamiltonian causes transitions between the $|0\rangle$ and $|1\rangle$ states. According to Fermi's Golden Rule, the transition rate W_{ij} between states i and j is given by [40]

$$W_{ij} = \frac{2\pi}{\hbar} |\langle j | \hat{H} | i \rangle|^2 \delta(E_j - E_i \pm \hbar\omega_{ij}). \quad (2.55)$$

Eq. 2.55 can be obtained from time-dependent perturbation theory (see Section 5.6 of [40]) and an analogous example to a phase qubit is discussed in Sect. 5.5 of [40]. The above form of the equation assumes that the levels are infinitely sharp (δ -functions), which is an approximation, but it is adequate for our discussion. Basically, Eq. 2.55 says that there will only be an appreciable transition rate between states $|0\rangle$ and $|1\rangle$ if the microwave drive has a frequency that is equal to the transition frequency ω_{01} . Substituting in for the matrix elements, averaging over time, and integrating over the δ -function, one finds

$$W_{01} = W_{10} = \frac{\pi I_\mu^2}{\omega_{01} C} = \pi I_\mu^2 \sqrt{\frac{L_J}{C}}. \quad (2.56)$$

There are several interesting effects that occur as a result of Eq. 2.56. Note, first of all, that $W_{01} = W_{10}$, thus the same signal that can excite the qubit into $|1\rangle$ can also force it from $|1\rangle$ back to $|0\rangle$. This second process is stimulated emission. Also, the transition rate goes like I_μ^2 which is proportional to the power of the microwave signal. A higher power microwave drive will change the state of the qubit faster. Because of this, strong coupling of the microwave drive to the qubit can be desirable; however, I will show later how we must be careful not to couple in noise that will cause unwanted transitions. In

addition, the factor $\sqrt{L_J/C}$ is just an effective impedance of the junction. Note that

$W_{01} \propto \frac{1}{\sqrt{C}}$, which means that a large capacitance device will require more current to

drive transitions in the qubit. In much the same way, it is harder to change the state of a Josephson junction with a higher inductance. The capacitance also comes into play in Chapter 5, when I discuss decoherence rates of the qubit.

2.5 Escape Rate and Tunneling

The eigenstates of the phase qubit are metastable bound states in the tilted washboard potential. Figure 2.5 shows the potential well and the first few bound energy states. Each state is localized on the left side by an effectively infinite potential, whereas on the right there is a finite potential barrier of height ΔU . In quantum mechanics, it is possible for a system to tunnel through a finite potential barrier [26]. According to the WKB approximation [26], the tunneling rate Γ_i of a system in the i^{th} energy states is

$$\Gamma_i = \frac{\omega_i}{2\pi} \exp\left(-\frac{2}{\hbar} \int_{q_1}^{q_2} \sqrt{2m(U(\gamma) - E_i)} d\gamma\right) \quad (2.57)$$

where ω_i is the attempt frequency, E_i is the energy of the i^{th} eigenstate, and q_1 and q_2 are the values of γ at which the $U(\gamma)$ is equal to E_i on the right side of the potential well (depicted in Fig. 2.5). In other words, $q_2 - q_1$ can be thought of as the width of the potential barrier as seen by a system in the i^{th} eigenstate. From Eq. 2.57, we see then that the tunneling rate depends on both the width of the barrier and the relative height of the potential barrier to the energy of that state. It can be shown that the tunneling rates of the $|0\rangle$ and $|1\rangle$ states of the junction are [26]




Fig. 2.5 Potential energy versus phase difference γ for the tilted washboard.

Metastable states are bound by a finite barrier of height ΔU . These states have a probability of tunneling out of the well. The escape rate Γ_i depends on the height and width of the barrier that it sees when the system is in the i^{th} energy state; therefore Γ_1 will be larger than Γ_0 . $U(q_1)$ and $U(q_2)$ depict the limits of the integration in the WKB approximation for computing Γ_0 [26].

$$\Gamma_0 \approx 12.7 \left(\frac{7.2\Delta U}{\hbar\omega_p} \right)^{1/2} \frac{\omega_0}{2\pi} \exp\left(-\frac{7.2\Delta U}{\hbar\omega_p} \right) \quad (2.58)$$

and

$$\Gamma_1 \approx 393.2 \left(\frac{7.2\Delta U}{\hbar\omega_p} \right)^{3/2} \frac{\omega_1}{2\pi} \exp\left(-\frac{7.2\Delta U}{\hbar\omega_p} \right) \quad (2.59)$$

where ω_p is the plasma frequency, and ΔU is the barrier height (see Fig. 2.5). From these results, we see that the escape rate of the system depends exponentially on the ratio of ΔU to the plasma frequency. As the bias current approaches the critical current, this ratio becomes small and the system will tunnel out more rapidly.

At sufficiently high biases, the difference between Γ_1 and Γ_0 can become large [see Fig. 2.6(b)]. Fig. 2.6(a) shows the escape rate from the $|0\rangle$ and $|1\rangle$ state vs. bias current I_b for a junction with $I_0 = 150$ nA and $C = 100$ fF assuming an attempt rate $\omega_0 = \omega_1 = \omega_p$. The markedly different escape rates from the $|0\rangle$ and $|1\rangle$ states can be used to determine the state of the qubit, as I explain in the next section.

2.6 Basic Qubit Operation and Pulsed Readout

I can now describe the basic operation of a Josephson junction as a qubit. First, one needs to prepare the junction in its ground state. This is achieved by cooling the junction down to a temperature T such that $T \ll \hbar f_{01}/k_B$, which for $f_{01} = 5$ GHz means $T \ll 250$ mK. This will ensure, that if left alone, the junction will relax to its ground state. In practice, I cool the junction in a dilution refrigerator with a base temperature of 20 mK in order to ensure that it will relax to the ground state after any applied signal is turned off.

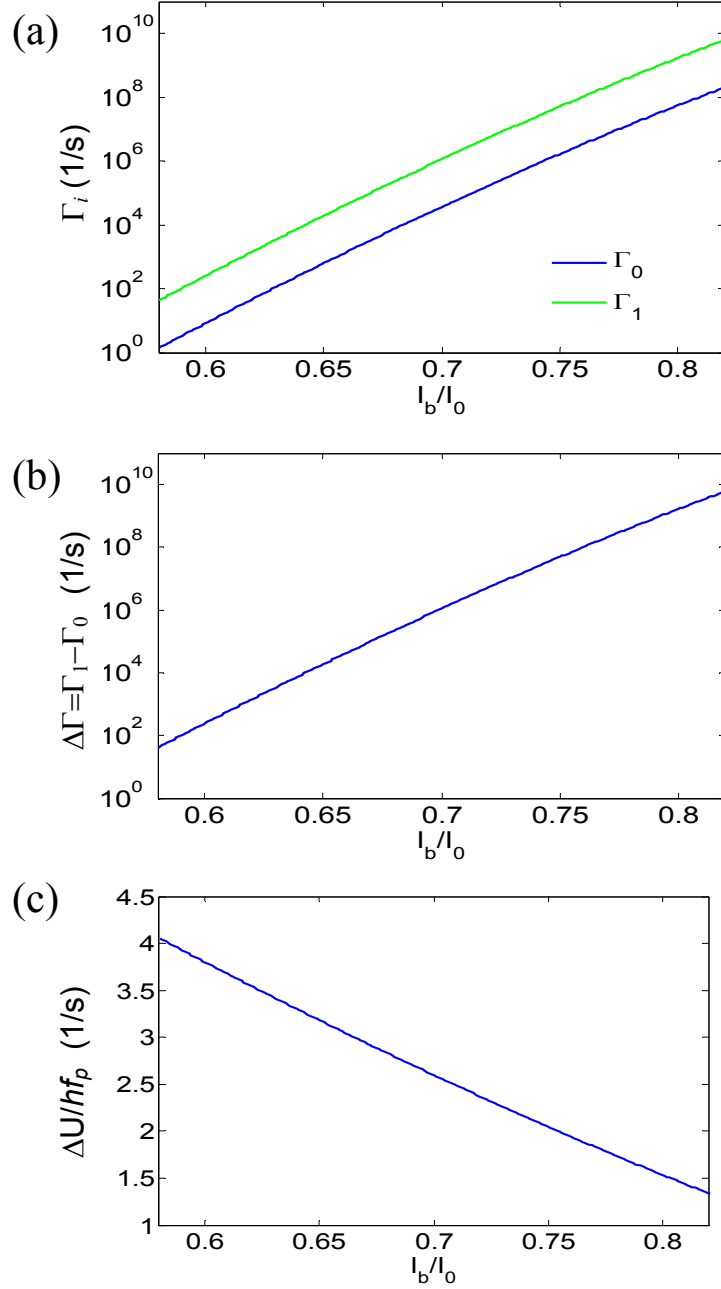


Fig. 2.6 (a) Escape rate from the ground and first excited states vs. current bias for a junction with $I_0 = 150$ nA and $C = 100$ fF. (b) Difference $\Delta\Gamma = \Gamma_1 - \Gamma_0$ between the escape rates of the states vs. bias current. Note that $\Delta\Gamma$ can get quite large as the bias increases. (c) Effective number of bound states in the well vs. bias current. For these parameters, the ground state begins to escape one time per second when there are about 4 states in the well.

Next, I apply bias current to the device such that there is some anharmonicity between the energy level spacings. For the junction parameters I have been discussing ($I_0 = 150$ nA, $C = 100$ fF), this means that I set the current to 90 nA ($0.6 I_0$). This will ensure that there is about 100 MHz of anharmonicity according to Fig. 2.4. Also, the escape rate from the $|1\rangle$ state will be about 100 Hz. That means that even if I perform experiments at a repetition rate of 1 kHz, there will be minimal tunneling from the $|1\rangle$ state while I manipulate the qubit's state. I then need only to apply a resonant microwave current for the right amount of time to change the state from $|0\rangle$ to any combination of $|0\rangle$ and $|1\rangle$ that I choose.

In order to read out the state of the qubit, I pulse the current flowing through the junction; raising it to a relatively high level for a short amount of time. If I raise the current to 120 nA ($0.79 I_0$), then the escape rate from the $|1\rangle$ state Γ_1 is 10^9 per second whereas $\Gamma_0 \approx 3 \cdot 10^7$ MHz. Both are relatively high escape rates, but if, for example, I only pulse the current for $t = 2$ ns, then a system occupying $|0\rangle$ will have a probability to tunnel out of the well [26] of

$$P_0 = 1 - e^{-\Gamma_0 t} \approx 0.06. \quad (2.60)$$

On the other hand, a system in the $|1\rangle$ state will tunnel out about 86% of the time. The larger the difference between Γ_1 and Γ_0 , the better one can discriminate the $|1\rangle$ and $|0\rangle$ states, and one must select the current pulse size I_p and duration t accordingly.

When the system tunnels out of the potential well, it continues to travel along the tilted washboard potential to lower and lower energy. Since the potential is essentially

infinite, the state will run on indefinitely with the phase of the junction γ steadily increasing. According to the ac Josephson relation (Eq. 2.16), a voltage develops across the junction when the phase changes with time. The voltage develops quickly after tunneling, and its value settles down to a value that is determined by the type of superconductors used to construct the junction. This process is commonly referred to as “switching to the voltage state” or just “switching” and the steady state voltage that develops is the “gap voltage.” $V_{gap} = 2\Delta/e$, where Δ is the superconducting gap of the material and e is the electron charge. This voltage is a few hundred μV for Al, and I can use a differential amp with its threshold set below V_{gap} to look for the switching event. If the switching event occurs during a 2 ns current pulse of the correct size, then the probability is high that the junction was in the $|1\rangle$ state. If not, then the probability is high that it was in the $|0\rangle$ state. Thus I look for the tunneling of the qubit during the measurement pulse in order to read out the state of the qubit.

2.7 Conclusions

In order to build a quantum computer, we need to couple together controllable quantum two-level systems with long energy relaxation times. A superconducting LC resonator should dissipate energy very slowly because the metal has nearly zero resistance at low temperatures if the resonant frequency f_0 is much less than Δ/h where Δ is the superconducting energy gap. However, an LC resonator is not a two level system. A Josephson junction acts as a non-linear oscillator and can act as an effective two-level system because its energy levels are anharmonic. The Josephson relations also imply that supercurrent flows through a junction with zero resistance, so there should be minimal

dissipation. The Hamiltonian for a junction contains an effective potential energy term that can be approximated as a harmonic potential with a cubic perturbation. The quadratic and cubic terms both depend on the bias current. At relatively high biases, the energy levels are more anharmonic. This means that we can selectively drive transitions between the ground state $|0\rangle$ and the first excited state $|1\rangle$ and use a junction as a qubit; a resonant microwave drive can cause transitions from $|0\rangle$ to $|1\rangle$ and $|1\rangle$ to $|0\rangle$, which allows computational operations on the qubit. The $|0\rangle$ and $|1\rangle$ states have different tunneling rates out of the potential well, which allows the state of the qubit to be measured. Thus, at least in principle, Josephson junctions are low loss, quasi-two-level quantum systems that can be manipulated and measured, which makes them candidates for qubits in a quantum computer.

Chapter 3

Decoherence in Phase Qubits

In Chapter 2, I explained how a Josephson junction can be used as a qubit. In that discussion, I assumed that the junction was a completely isolated, noise-free system. In fact, quantum information processing will require that the qubits are in well defined states [4,5,42]. In this chapter, I discuss processes by which the qubit can lose quantum information. The idea is that unintended interactions of a qubit with its environment will cause the state of the qubit to become less and less precisely defined. The generic name for this loss of quantum information is called “decoherence.” Decoherence of a qubit can happen in two different ways:

- (i) “Dephasing” is a type of decoherence in which the phase factor in a superposition state becomes uncertain.
- (ii) If the qubit dissipates energy (relaxation) or absorbs unwanted energy (excitation) from coupling to the environment, changing the probability to be in a given state.

A quantum system must be protected against all decoherence processes if it is to remain in a well-defined quantum superposition state.

3.1 Bloch Sphere Representation of the State of a Qubit

Any pure quantum state of a two-level system can be written in the form

$$|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle \quad (3.1)$$

Here $|0\rangle$ is the ground state and $|1\rangle$ is the excited state and these are the energy eigenstates of the system’s Hamiltonian H . Thus:

$$\langle 0|\hat{H}|0\rangle = E_0 \quad (3.2)$$

$$\langle 1|\hat{H}|1\rangle = E_1, \quad (3.3)$$

and E_0 and E_1 are the eigenenergies of the system.

In Eq. 3.1, α is the probability amplitude of being in the ground state and β is that of being in the excited state. α and β are complex numbers and must obey

$$|\alpha|^2 + |\beta|^2 = 1. \quad (3.4)$$

Without any loss of generality, we can take α as a positive real number and β as a complex number. Since any complex number can be written as a positive real number times an exponential phase factor, we can replace Eq. 3.1 by

$$|\Psi\rangle = \alpha|0\rangle + \beta e^{i\phi}|1\rangle \quad (3.5)$$

where now α and β are positive real numbers that obey Eq. 3.4 and ϕ is a real number.

The factor ϕ is the phase difference between α and β . Since α and β are positive real numbers that obey Eq. 3.4, we can always define an angle θ such that

$$\alpha = \cos\left(\frac{\theta}{2}\right) \quad (3.6)$$

$$\beta = \sin\left(\frac{\theta}{2}\right) \quad (3.7)$$

where θ ranges from 0 to π . Any state of the system can thus be written in the form

$$|\Psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + \sin\left(\frac{\theta}{2}\right)e^{i\phi}|1\rangle. \quad (3.8)$$

With this introduction of the angles θ and ϕ , it is now possible to visualize the state $|\Psi\rangle$ graphically as a vector of length one that points in the direction specified by θ

and ϕ (see Fig. 3.1). Since the state vector has unit length, and θ ranges from 0 to π , and ϕ can have any real value, the locus of all possible state vectors just describes a sphere with radius 1. This sphere is called the Bloch sphere.

We can now identify specific states of the system with specific points on the Bloch sphere. For example, when $\theta = 0$ (the North pole) the state is $|0\rangle$, and when $\theta = \pi$ (the South pole) the state is $|1\rangle$. The 50/50 superposition state $|\Psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ lies on the equator of the Bloch sphere at $\theta = \pi/2$ and $\phi = 0$, where the +x axis crosses the equator. In fact, different points on the equator correspond to different 50/50 superposition states. $|\Psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle)$ describes a distinct quantum state and is located at $\theta = \pi/2$ and $\phi = \pi/2$, or where the +y axis crosses the equator (see Fig. 3.1).

In this chapter, I use the Bloch sphere to illustrate the decoherence of the qubit state under specific conditions. A more complete understanding of the dynamics of the state can be obtained through the density matrix formalism, which is treated in [26], [37], and [38].

3.2 Phase Precession

To see how a superposition state evolves from some initial state $|\Psi(0)\rangle$ at time $t = 0$ to a state $|\Psi(t)\rangle$ at time t , one can use the time evolution operator $\hat{\tau}$, where [40]

$$\hat{\tau} = e^{-i\frac{\hat{H}}{\hbar}t}. \quad (3.9)$$




Fig. 3.1 Bloch sphere representation of the qubit state

$$|\Psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + \sin\left(\frac{\theta}{2}\right)e^{i\phi}|1\rangle.$$

The polar angle θ determines the probability of finding the system in $|0\rangle$ or $|1\rangle$, and the azimuthal angle ϕ represents the phase of the superposition state.

One finds

$$|\Psi(t)\rangle = e^{-\frac{i\hat{H}}{\hbar}t} |\Psi(0)\rangle = \cos\left(\frac{\theta}{2}\right) e^{-\frac{iE_0}{\hbar}t} |0\rangle + \sin\left(\frac{\theta}{2}\right) e^{i\phi(0)} e^{-\frac{iE_1}{\hbar}t} |1\rangle \quad (3.10)$$

where ϕ is the phase at time $t = 0$. Only the phase difference between α and β will be physically relevant, and I can rewrite the state as:

$$|\Psi(t)\rangle = \cos\left(\frac{\theta}{2}\right) |0\rangle + \sin\left(\frac{\theta}{2}\right) \exp\left[-i\left(\frac{E_1 - E_0}{\hbar}\right)t + \phi(0)\right] |1\rangle \quad (3.11)$$

or equivalently:

$$|\Psi\rangle = \cos\left(\frac{\theta}{2}\right) |0\rangle + \sin\left(\frac{\theta}{2}\right) \exp[i\phi(t)] |1\rangle \quad (3.12)$$

where the phase $\phi(t)$ is given by:

$$\phi(t) = -\left(\frac{E_1 - E_0}{\hbar}\right)t + \phi(0). \quad (3.13)$$

According to Eq. 3.13, $\phi(t)$ advances steadily with time. On the Bloch sphere, this means the state will precess around the z-axis at a rate determined by the energy difference between $|0\rangle$ and $|1\rangle$. This angular frequency is just

$$\omega_{01} = \left(\frac{E_1 - E_0}{\hbar}\right). \quad (3.14)$$

Since the phase precession is steady, it is often convenient to work in a co-rotating frame of reference in which the state vector is stationary [37]. As long as, E_1 and E_0 are known and constant, this method works. However, if $E_1 - E_0$ varies unpredictably, then the phase of the quantum state will evolve unpredictably. This leads to a discussion of the first type of decoherence: dephasing.

3.3 Dephasing from Low Frequency Current Noise

As the state $|\Psi\rangle$ evolves in time, the phase of the state at time t is:

$$\phi(t) = \omega_{01}t \quad (3.15)$$

where I have assumed the phase is zero at time $t = 0$ and ω_{01} is constant. Now consider a case where a perturbation to the system causes E_0 and E_1 to fluctuate slowly in time.

Assuming the perturbation Hamiltonian is diagonal, Eq. 3.15 becomes

$$\phi(t) = \int_0^t \omega_{01}(t') dt' . \quad (3.16)$$

Assuming that $\omega_{01}(t)$ is making random fluctuations $\delta\omega_{01}(t)$ around an average value

$\langle\omega_{01}\rangle$, I can write

$$\omega_{01}(t) = \langle\omega_{01}\rangle + \delta\omega_{01}(t), \quad (3.17)$$

and thus

$$\phi(t) = \langle\omega_{01}\rangle t + \int_0^t \delta\omega_{01}(t') dt' . \quad (3.18)$$

The integral in Eq. 3.18 describes the phase noise in the system, $\phi_n(t)$.

In a phase qubit, one source of fluctuations in ω_{01} is low-frequency current noise [43]. In the case of current noise $I_n(t)$, the phase noise as a function of time is

$$\phi_n(t) = \frac{d\omega_{01}}{dI} \int_0^t I_n(t') dt' \quad (3.19)$$

where $\frac{d\omega_{01}}{dI}$ is the noise transfer function for current noise. As I described in Chapter 2,

ω_{01} of a phase qubit depends on the current flowing through the junction because a change in current changes the Josephson inductance of the junction.

The magnitude of the phase noise in the qubit can be characterized by the root mean square (rms) value $\langle \phi_n^2(t) \rangle$ [43]

$$\langle \phi_n^2(t) \rangle = \left(\frac{d\omega_{01}}{dI} \right)^2 \left\langle \int_0^t I_n(t') dt' \int_0^t I_n(t'') dt'' \right\rangle. \quad (3.17)$$

Assuming that the current is fluctuating randomly around an equilibrium value $I_n(0)$, the time average in Eq. 3.17 can be related to the noise power spectral density $S_I(f)$ by

$$\langle I_n(t) I_n(0) \rangle = \int_0^\infty S_I(f) \cos(2\pi f t) df. \quad (3.18)$$

Applying Eq. 3.18 to Eq. 3.17, we get

$$\langle \phi_n^2(t) \rangle = \left(\frac{d\omega_{01}}{dI} \right)^2 \int_0^\infty S_I(f) df \int_0^t dt' \int_0^t dt'' \operatorname{Re} e^{i2\pi f(t'-t'')}. \quad (3.19)$$

Evaluating the time integrals gives [43]

$$\int_0^t dt' \int_0^t dt'' \operatorname{Re} e^{i2\pi f(t'-t'')} = \frac{\sin^2(\pi f t)}{(\pi f)^2} = W_0(f, t), \quad (3.20)$$

where t is the amount of time that phase noise has been allowed to occur, and thus

$$\langle \phi_n^2(t) \rangle = \left(\frac{d\omega_{01}}{dI} \right)^2 \int S_I(f) W_0(f, t) df, \quad (3.21)$$

where $W_0(f, t)$ is known as the spectral weighting function.

Figure 3.2 shows W_0 for measurement times $t = 10$ ns, 100 ns, and 1 μ s. Note that W_0 is largest for low frequencies. Note also that what counts as a low frequency depends on the measurement time t . For long measurement times, W_0 starts higher, but it rolls off at relatively low frequencies. This means that dephasing is mostly caused by low-frequency fluctuations in the current. Noise sources with $1/f$ noise power spectra, such as flux noise and critical current noise, are among the main causes of dephasing in phase




Fig. 3.2 Weighting function $W(f,t)$ for dephasing from noise at frequency f for measurement times of $t = 10$ ns, 0.1 μs, and 1 μs. Note that noise at high frequencies produces relatively little dephasing. Note also that what qualifies as a high frequency depends inversely on the measurement time. Noise sources that have a $1/f$ power spectral density will be important sources of dephasing.

qubits [43,44]. For $1/f$ noise, a limit must be imposed on the frequency integral in Eq. 3.21 so that a finite value can be found for the rms phase noise. As I said, most of the dephasing is caused by low frequency noise, so the upper limit need not be extremely high. Martinis and Urbina suggest a limit of f_{0I} because above that frequency, most of the current will be shunted through the junction capacitance and will not affect the Josephson inductance [43]. This is almost certainly high enough because the spectral weighting function is very small at f_{0I} , and noise at frequencies above f_{0I} will contribute negligible dephasing regardless of the current division unless $S_I(f > f_{0I})$ is high.

Equation 3.21 for the rms phase noise at time t caused by random current fluctuations can be written as

$$\langle \phi_n^2(t) \rangle = \left(2\pi \frac{df_{0I}}{dI} \right)^2 \int_0^{f_{0I}} S_I(f) W_0(f, t) df. \quad (3.22)$$

There are two main sources of dephasing that I consider in this thesis: thermal noise and $1/f$ noise. Each has a distinct power spectral density; therefore, each causes phase noise to accumulate in different ways. The noise power spectrum of a classical thermal noise source is [43]

$$S_I^{Th}(f) = \frac{4k_B T}{R} \quad (\text{valid for } hf \ll k_B T) \quad (3.23)$$

where T is temperature of the noise source, and R is the effective resistance across the junction associated with the noise source. Note that R here is the same as the R_{eff} that I discussed at length in Chapter 5. Inserting this into Eq. 3.22 and integrating shows that the rms phase noise accumulates according to [43]:

$$\langle \phi_n^2(t) \rangle \approx \frac{2k_B T}{R} \left(2\pi \frac{df_{0I}}{dI} \right)^2 t. \quad (3.24)$$

Since $\langle \phi_n^2(t) \rangle \propto t$, we can expect a random walk in the phase and an exponential decay of the phase coherence of the quantum state [43]. The decay constant is known as T_ϕ and is given by

$$\frac{1}{T_\phi} \approx \frac{k_B T}{R} \left(2\pi \frac{df_{01}}{dI} \right)^2. \quad (3.25)$$

The other type of dephasing sources is those with $1/f$ noise power spectra. In such cases, the noise power spectrum is

$$S_I^{1/f}(f) = \frac{S_I^{1/f}(1\text{Hz})}{f} \quad (3.26)$$

where $S_I^{1/f}(1\text{Hz})$ is the noise power spectral density at 1 Hz. Integrating Eq. 3.21 with this power spectrum gives [43]

$$\langle \phi_n^2(t) \rangle \approx \left(2\pi \frac{df_{01}}{dI} \right)^2 S_I^{1/f}(1\text{Hz}) \ln \left(\frac{0.401}{f_r t} \right) t^2 \quad (3.27)$$

where f_r is the repetition rate of the state measurement. Note that for $1/f$ noise, phase noise accumulates like t^2 , so the phase coherence of the state decays in a Gaussian fashion, not exponentially. Also, note that since f_r is often on the order of kHz and $t \sim 1 \mu\text{s}$, the natural logarithm gives a factor of order 1. This dependence means that we can only make a rough estimate for the Gaussian decay constant for the phase coherence. The decay constant T_R seen in Ramsey fringe experiments has a similar form [43]

$$\frac{1}{T_R} \sim \left[S_I^{1/f}(1\text{Hz}) \left(2\pi \frac{df_{01}}{dI} \right)^2 \right]^{1/2}. \quad (3.28)$$

In Chapters 5 and 10, I consider the effects of $1/f$ noise and thermal noise on the phase coherence of a dc SQUID phase qubit.

3.4 Relaxation from Dissipation and High Frequency Noise

In Sect. 2.4, I showed how a resonant microwave signal can drive transitions between the $|0\rangle$ and $|1\rangle$ states in a junction. If a noise current $\Delta I(t)$ is present, the Hamiltonian of the junction is:

$$H = \frac{p^2}{2m} + \frac{\Phi_0}{2\pi} \{I_0 \cos \gamma + \gamma [I_b + \Delta I]\} . \quad (3.29)$$

The Hamiltonian can be separated into two parts:

$$\hat{H} = \hat{H}_0 + \hat{H}_{\Delta} \quad (3.30)$$

where the second part is a small perturbation due to the noise current.

In the $\{|0\rangle, |1\rangle\}$ basis, the perturbation Hamiltonian is:

$$\hat{H}_{\Delta} = \frac{\Phi_0}{2\pi} \Delta I \begin{pmatrix} \langle 0 | \gamma | 0 \rangle & \langle 0 | \gamma | 1 \rangle \\ \langle 1 | \gamma | 0 \rangle & \langle 1 | \gamma | 1 \rangle \end{pmatrix} \quad (3.31)$$

I will ignore the diagonal terms, and work in the harmonic oscillator approximation, but it is important to digress briefly to understand why I can do so here. Fluctuations in the current lead to changes in the energies of the $|0\rangle$ and $|1\rangle$ states, which means that there is a contribution to the diagonal terms of the Hamiltonian from the noise [43]. These are just the terms that cause dephasing and the basic theory for these types of fluctuations was dealt with in the previous section. Thus I will concentrate on the off-diagonal terms here.

As I explained in Sect. 2.4, the off-diagonal elements of Eq. 2.51 cause transitions between $|0\rangle$ and $|1\rangle$ at a rate described by Fermi's Golden Rule [40]:

$$W_{01} = W_{10} = \frac{\pi (\Delta I(\omega_{01}))^2}{\omega_{01} C} \quad (3.32)$$

where I write the current noise as $\Delta I(\omega_{01})$ because it is only current fluctuations at ω_{01} that will cause transitions [40]. Of course there are two types of transitions: a qubit can absorb energy from a noise source or the qubit can emit energy (either spontaneously or via stimulated emission) that will be absorbed by the source [40].

In designing my qubit I had to deal specifically with unwanted transitions caused by thermal noise (Johnson-Nyquist noise) due to electrical resistance in the qubit or circuits to which it was coupled (see Sect. 5.1). Whereas Eq. 3.32 is an ideal equation that is true for a coherent noise current source with a constant phase, a thermal noise current is random and can cause up transitions and down transitions [26]. At low enough temperatures, only down transitions will be present, and these can be viewed equivalently as resulting from spontaneous emission or as the result of zero point noise.

The transition rate from state $|i\rangle$ to state $|j\rangle$ in a phase qubit from random noise current can be written as (see [26] Appendix B and Sect. 3.2.1):

$$W_{ij} = \frac{\Delta E_{ij}}{2e^2 R} \frac{|\langle i|\gamma|j\rangle|^2}{\exp[\Delta E_{ij}/(k_B T)] - 1} \quad (3.33)$$

where ΔE_{ij} is the energy spacing between $|i\rangle$ and $|j\rangle$, e is the electron charge, R is the value of the effective resistance shunting the junction that represents dissipation from the noise source (see Sect. 5.1 for more discussion of R), k_B is the Boltzmann constant, and T is the temperature of the noise source. Note that the stimulated transition rate from $|i\rangle$ to $|j\rangle$ is the same as the transition rate from $|j\rangle$ to $|i\rangle$. As a result, thermal noise will tend to cause the populations of the states to be distributed evenly [26]. However, if $|j\rangle$ is higher in energy than $|i\rangle$, then there will also be a spontaneous decay rate from $|j\rangle$ to $|i\rangle$,

which we can write as Γ_{ji} [26]. In order to keep a steady thermal population in the states (detailed balance), the rate up and the sum of the down rates must be related by:

$$W_{ij} \exp\left(\frac{\Delta E_{ij}}{k_B T}\right) = W_{ji} + \Gamma_{ji} \quad (3.34)$$

$\Gamma_{10} = 1/T_I$ is known as the energy relaxation rate. The energy relaxation rate is a measurable quantity because a qubit that is prepared in the $|1\rangle$ state will decay back to the $|0\rangle$ state at a rate Γ_{10} . The time constant for the exponential decay is $T_I = 1/\Gamma_{10}$. According to Eq. 3.32 $W_{ij} = W_{ji}$, and if I substitute Eq. 3.33 into Eq. 3.34 and solve for $1/\Gamma_{10}$, then the thermal factors from Eq. 3.33 and 3.34 cancel. One thus obtains a theoretical prediction for the T_I of a Josephson junction qubit due to spontaneous relaxation:

$$\Gamma_{10} = \frac{1}{T_I} = \frac{\Delta E_{01}}{R} \frac{|\langle 0 | \gamma | 1 \rangle|^2}{2e^2} \quad (3.35)$$

If I assume that the harmonic oscillator approximation for the levels is appropriate, then the matrix elements work out such that (see Sect. 2.3 for a discussion of harmonic oscillator matrix elements):

$$T_I = RC \quad (3.36)$$

where C is the capacitance shunting the junction and R is the effective shunting resistance seen by the junction.

Note that thermal noise will cause transitions in both directions. However, in most experiments, $k_B T \ll \hbar f$ and in this limit only zero-point noise will be present and only down transitions will be important. It is probably less confusing to just think of this as

dissipative relaxation characterized with a time constant T_1 rather than zero-point noise stimulated downward transitions.

In general, the theory behind all of the sources of dissipation discussed in this thesis (*i.e.* high frequency noise, dielectric loss, etc.) boils down to interactions with the qubit that contribute off diagonal elements to the Hamiltonian of the qubit [43, 15-17]. Any unexpected change of the qubit state from $|1\rangle$ to $|0\rangle$ or from $|0\rangle$ to $|1\rangle$ would represent a potential error in a computation, but the energy relaxation rate T_1 gives a direct measure of the lifetime of the qubit's excited state.

3.5 Conclusions

The probability amplitudes of the state of a qubit consist of real and imaginary parts. Each must be controlled precisely if the qubit is to be used for quantum computation. Dephasing is caused by relatively slow (low frequency) random changes in the energy levels of the qubit because such variations lead to uncertainty in the phase of superposition states. Unwanted transitions between energy states, caused by loss of energy to the environment (relaxation) or absorption of energy from the environment (excitation) create uncertainty in the real parts of the state probability amplitudes. In Chapter 5, I describe how I modeled the effects of specific sources of decoherence in my SQUID phase qubits.

Chapter 4

dc SQUID Phase Qubit

In this chapter I discuss the basic physics of the dc SQUID phase qubit. Although in principle a current biased Josephson junction meets the requirements for being used as a qubit, it is vulnerable to decoherence because of its connections to the outside world. As I explained in Chapter 2, a junction needs to be biased at a significant fraction of its critical current for the two lowest energy levels to have an appreciable anharmonicity. Also, in order to detect the state of the system using a pulsed current readout with a voltage readout, a voltage needs to be developed across the junction when it switches. These require some kind of galvanic connection to the Josephson junction. This connection could let down low and high frequency noise that will cause decoherence, as I described in Chapter 3.

One way to partially get around these difficulties is to abandon a switching voltage readout and instead use some other approach to detect the state of the system. This is the approach taken in the rf SQUID qubit, where a dc SQUID is used to detect changes in the flux in the rf SQUID qubit. Another way around these difficulties, is to add a filtering network to protect the qubit from noise and dissipation from the bias leads. It is this second approach that I followed and is the basis for the dc SQUID phase qubit.

In the dc SQUID phase qubit, one junction in a dc SQUID is used as the qubit and the other junction is used to filter out current noise from the leads, thereby protecting the qubit junction from dephasing and relaxation due to the leads. However, adding another junction to the system makes the resulting Hamiltonian more complicated. K. Mitra examined the behavior of the dc SQUID and found that when the loop inductance is large

enough, the dynamics of the two junctions can be separated [45]. However, when the SQUID loop inductance is relatively small, the dynamics of the two junctions are coupled and the full Hamiltonian must be considered.

In this chapter I first present a model for the SQUID Hamiltonian and obtain approximate analytical solutions for the normal modes. The normal mode formalism can be used to find the frequencies of the modes and also to assess the isolation provided by the dc SQUID design. In addition to approximate analytical results, I will also present numerical results for the energy levels of the SQUID, obtained by use of a finite difference method. The lesson of this analysis is that adding a second junction to the phase qubit helps to filter out broadband noise, but it makes the dynamics of the system much more complicated.

4.1 The Need for Isolation from the Bias Leads

Consider a single Josephson junction connected to a current bias line (see Fig. 4.1a). Assuming that the line has a characteristic impedance of $Z_0 = 50 \, \Omega$, the junction will be driven by Johnson-Nyquist current noise with a root-mean-square amplitude ΔI of [46,47]:

$$\Delta I = \sqrt{\frac{4k_B T \Delta f}{R}} \quad (4.1)$$

where k_B is the Boltzmann constant, R is $50 \, \Omega$, T is the temperature of the noise source, and Δf is the bandwidth of the noise. It is important to remember that Eq. 4.1 is classical thermal noise that is found in the limit $\hbar f \ll k_B T$ where f is the frequency. For a junction at 20 mK, Eq. 4.1 will only be valid for $f \ll 400$ MHz. This is an appropriate limit for

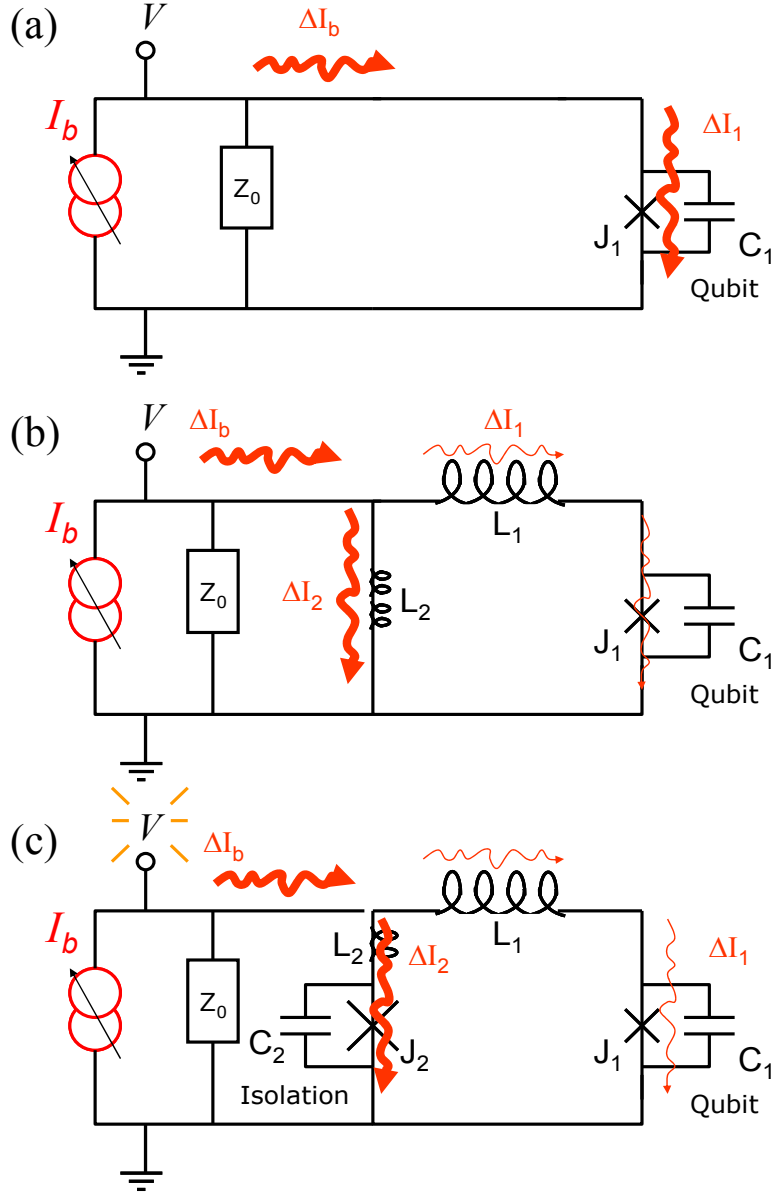


Fig. 4.1 (a) Diagram of a current biased junction with no isolation. The Johnson-Nyquist noise current and dissipation from the leads with an impedance of $Z_0 \approx 50 \, \Omega$ causes decoherence in the qubit. (b) Qubit with an inductive divider for isolation. Although the lead impedance is stepped up and much of the noise is shunted away from the qubit, the switching of the junction cannot be detected. (c) Diagram of the dc SQUID phase qubit. The isolation branch has a large junction J_2 that can switch when the qubit junction J_1 tunnels out of the potential well during readout.

some situations, such as dephasing or inhomogenous broadening, which are sensitive to processes at relatively low frequencies. However, it would be inaccurate for processes such as relaxation, which are sensitive to processes that occur at the qubit transition frequency of a few GHz. In general for $hf \ll k_B T$, one needs to consider the quantum mechanical behavior of the system, including effects of zero-point noise and roll-off in the thermal noise at high frequencies [48].

If a junction is directly plugged into a thermalized transmission line with an impedance of 50Ω and cooled in a dilution refrigerator to 20 mK [26], then the noise will be roughly flat (white) up to about 400 MHz. For a $\Delta f = 10$ MHz bandwidth, Eq. 4.1 gives a noise current of $\Delta I \sim 500$ pA. If this noise actually occurred at the qubit transition frequency, it would cause on the order of a 10^8 unwanted qubit state transitions per second according to Eq. 3.31, assuming that $f_{01} = 5$ GHz and $C = 0.1$ pF. Fortunately, $hf_{01} \gg k_B T$ in this case, and the thermal noise term will be smaller by a factor of $\exp(-hf_{01}/k_B T) \sim \exp(-10)$.

For $hf_{01} \gg k_B T$, a more significant effect is zero point noise, or equivalently, relaxation due to dissipative impedance of the leads. This leads to relaxation on a time scale of $RC = Z_0 C \approx 5$ ps for a junction with $C = 0.1$ pF. One way to reduce the dissipation caused by the bias current line is to increase the characteristic impedance that it presents to the junction. Putting an inductive divider [see Fig. 4.1(b)] between the qubit and the 50Ω impedance of the bias lines will reduce the amount of noise current that gets to the qubit and also step up the effective lead impedance.

To see how this works, consider a noise current ΔI_b on the bias lines [see Fig. 4.1(b)]. The ratio of the total noise current ΔI_b to the noise current that flows through the qubit junction is

$$\frac{\Delta I_b}{\Delta I_1} = \frac{L_1 + L_{J1} + L_2}{L_2}. \quad (4.2)$$

This ratio squared has been defined as r_I , the isolation factor [49].

$$r_I = \left(\frac{\Delta I_b}{\Delta I_1} \right)^2 = \left(\frac{L_1 + L_{J1} + L_2}{L_2} \right)^2. \quad (4.3)$$

r_I is the fraction of the noise current power that the qubit junction sees. By examining Eq. 4.1, we see that adding inductive isolation is like increasing the characteristic impedance of the line such that [49]

$$R_{eff} = Z_o r_I \approx Z_o \left(\frac{L_1 + L_{J1} + L_2}{L_2} \right)^2. \quad (4.4)$$

If $L_1 + L_{J1} \gg L_2$, this technique could greatly reduce the dissipation and dephasing in the qubit caused by noise from the current bias line.

The only problem with using an inductive filter as described above is that, when the qubit junction switches, any voltage that develops across the junction will tend to just redirect current through L_2 , *i.e.* an inductor has zero impedance at dc. Thus, if the qubit switches, one might be able to see a brief voltage pulse across the inductor, but this will quickly relax to zero volts, rather than ramp up to a steady value of $2\Delta/e$.

To get around this difficulty, in 2002 Martinis *et al.* [48] proposed adding a second junction on the low inductance branch of the inductive isolation network [see Fig. 4.1(c)], thereby forming a dc SQUID phase qubit. The second junction can be made to have a very low Josephson inductance L_{J2} by making its critical current large and, since it

can also switch to the voltage state, it allows a steady state voltage to appear across the device after it switches. Using Eq. 2.22, with no bias through the junction, I find that a junction with a critical current of 10 μA only contributes an inductance of about

$$L_{J_2} = \frac{\Phi_o}{2\pi I_o} \cong 30 \text{ pH} . \quad (4.5)$$

Figure 4.2 shows a circuit diagram of the dc SQUID phase qubit. Junction J_1 is used as the qubit junction. The loop inductance L_1 is generally of the order of 1 nH. The side of the SQUID with J_1 and L_1 is known as the qubit arm of the SQUID loop. The isolation arm of the SQUID loop consists of a very small geometrical inductance L_2 (~ 10 pH) and the isolation junction J_2 . Current supplied by the current bias leads is divided such that most of it goes through J_2 and very little gets to the qubit. From Eqs. 4.2 and 4.3, the isolation power factor that this provides for the qubit is

$$r_1 \approx \left(\frac{\Delta I_b}{\Delta I_1} \right)^2 = \left(\frac{L_1 + L_{J_1} + L_2 + L_{J_2}}{L_2 + L_{J_2}} \right)^2 . \quad (4.6)$$

In the dc SQUID phase qubit, most of the bias current will be shunted away from the qubit. This reduces the noise reaching the qubit, but also means that the qubit junction cannot be biased with the current bias source. Instead, a flux bias line is used to apply flux to the SQUID loop, thereby causing circulating current in the loop and setting the current in the qubit junction. The current bias line can then be used to independently set the current level in the isolation junction, which is important for ensuring proper switching of the device and also sets the value for L_{J_2} . Finally, microwave excitation current I_μ and measurement pulses V_p can be coupled into the device by means of a small capacitor C_μ (see Fig. 4.2).




Fig. 4.2 Diagram of the dc SQUID phase qubit. The bias current I_b is isolated from the qubit junction by making L_1 much larger than $L_2 + L_{J2}$. Since most of the bias current cannot reach the qubit junction, a flux bias line I_ϕ is used to control the current in the qubit by applying flux Φ_a to the SQUID, which creates a circulating current that can be used to selectively bias the qubit junction. The microwave drive I_μ and measurement pulses V_p come in through a small capacitor C_μ .

4.2 SQUID Hamiltonian

In order to understand the dynamics of the dc SQUID phase qubit, I need to first find the Hamiltonian of the system. The coordinates of the Hamiltonian are the phase differences across the qubit and isolation junctions γ_1 and γ_2 , respectively. There are two external controls: the bias current I_b and the flux bias I_Φ . The bias current is related to the currents in the arms of the SQUID by conservation of current:

$$I_b = I_1 + I_2, \quad (4.7)$$

where I_1 and I_2 are the currents that flow down the qubit and isolation arms of the SQUID, respectively. From the RCSJ model (assuming R_J is infinite; see Sect 2.2 for derivation) I can write the current flowing in each arm of the SQUID as

$$I_1 = C_1 \frac{\Phi_0}{2\pi} \ddot{\gamma}_1 + I_{01} \sin \gamma_1 \quad (4.8a)$$

$$I_2 = C_2 \frac{\Phi_0}{2\pi} \ddot{\gamma}_2 + I_{02} \sin \gamma_2 \quad (4.8b)$$

where C_i is the capacitance across the i^{th} junction and I_{0i} is the critical current of that junction.

Equations 4.7-4.8 give three coupled classical equations with one control parameter (I_b) and four unknowns (I_1 , I_2 , γ_1 and γ_2). To solve this system of equations, we need another independent equation: the flux phase relation [28,50]. This equation relates the remaining control parameter (the flux applied to the loop) to the other current in the device, the circulating current J flowing around the loop. This flux-phase relation says that the phase change on a path taken around the loop is an integer multiple of 2π [28,50]:

$$\gamma_1 - \gamma_2 = 2\pi \left(n + \frac{\Phi}{\Phi_0} \right) \quad (4.9)$$

where n is an integer and Φ is the total flux in the SQUID loop. There is a negative sign on γ_2 because the phase difference is found by integrating $d\gamma$ around the loop in one direction (clockwise in Fig. 4.2), which means from top to bottom in one case (positive sign for junction J_1) and bottom to top for the other (negative sign for junction J_2). One can set $n = 0$ because a constant offset of $2\pi n$ does not affect the dynamics of the system, so Eq. 4.9 becomes simply:

$$\gamma_1 - \gamma_2 = 2\pi \frac{\Phi}{\Phi_0} \quad (4.10)$$

The total flux Φ in the loop is due to current I_1 flowing in the inductance L_1 in the qubit arm of the SQUID loop, the current I_2 flowing in the inductance L_2 in the isolation arm of the SQUID loop, and current I_Φ flowing in the flux bias line. Considering the sign of the flux produced, I can write:

$$\gamma_1 - \gamma_2 = \frac{2\pi}{\Phi_0} (I_1 L_1 - I_2 L_2 + I_\Phi M) \quad (4.11)$$

Note that in Eq. 4.11 I am ignoring any mutual inductance between the current bias line and the loop. This should be small if the device is designed correctly.

Solving for I_1 in Eq. 4.11 one finds:

$$I_1 L_1 = \frac{\Phi_0}{2\pi} (\gamma_1 - \gamma_2) - I_2 L_2 + I_\Phi M \quad (4.12)$$

By adding $I_2 L_2$ to both sides, I can get I_1 in terms of I_b , and find:

$$I_1 = \frac{1}{L} \left[\frac{\Phi_0}{2\pi} (\gamma_1 - \gamma_2) + I_b L_2 - I_\Phi M \right] \quad (4.13)$$

Similarly I_2 can be written:

$$I_2 = \frac{1}{L} \left[\frac{\Phi_0}{2\pi} (\gamma_1 - \gamma_2) - I_b L_1 - I_\Phi M \right] \quad (4.14)$$

Substituting Eqs. 4.13 and 4.14 into 4.8(a) and 4.8(b) I can now find the following classical equations of motion for γ_1 and γ_2 :

$$C_1 \left(\frac{\Phi_0}{2\pi} \right)^2 \ddot{\gamma}_1 = \frac{\Phi_0}{2\pi} \left[I_{01} \sin \gamma_1 + \frac{L_2}{L} I_b \gamma_1 + \frac{\Phi_0}{2\pi} \left(\gamma_1 - \gamma_2 - \frac{M}{L} I_\Phi \right) \right] \quad (4.15)$$

$$C_2 \left(\frac{\Phi_0}{2\pi} \right)^2 \ddot{\gamma}_2 = \frac{\Phi_0}{2\pi} \left[I_{02} \sin \gamma_2 + \frac{L_1}{L} I_b \gamma_2 - \frac{\Phi_0}{2\pi} \left(\gamma_1 - \gamma_2 - \frac{M}{L} I_\Phi \right) \right] \quad (4.16)$$

where $L = L_1 + L_2$. The left hand side of these equations has a term which is analogous to mass times the acceleration and the right hand sides can be viewed as being bias-dependent force terms.

I can now find the effective potential energy of the system by noting that the effective force on a system is the negative gradient of the potential with respect to the position coordinates. Considering the integral of the RHS of Eqs. 4.15 with respect to γ_1 and the RHS of Eq. 4.16 with respect to γ_2 one can identify an effective potential for the dc SQUID [45]:

$$U(\gamma_1, \gamma_2, I_b, I_\Phi) = -E_{J1} \cos \gamma_1 - E_{J2} \cos \gamma_2 - \frac{\Phi_0}{2\pi} \frac{I_b}{L} (L_2 \gamma_1 + L_1 \gamma_2) + \left(\frac{\Phi_0}{2\pi} \right)^2 \frac{1}{2L} \left(\gamma_1 - \gamma_2 - \frac{2\pi}{\Phi_0} M I_\Phi \right)^2 \quad (4.17)$$

where $E_{Ji} = \frac{\Phi_0}{2\pi} I_{0i}$ is the Josephson energy of the i -th junction.

Given the potential energy, one can construct the dc SQUID Hamiltonian [37,45]:

$$H = \frac{p_1^2}{2m_1} + \frac{p_2^2}{2m_2} + U(\gamma_1, \gamma_2, I_b, I_\Phi) \quad (4.18)$$

where $p_i = C_i \left(\frac{\Phi_0}{2\pi} \right)^2 \dot{\gamma}_i = m_i \dot{\gamma}_i$ is the canonical momentum of the i^{th} junction. This

Hamiltonian governs the dynamics of the dc SQUID phase qubit.

4.3 Cubic Approximation to the SQUID Hamiltonian

As I discussed in Sect. 2.3 for a single junction, it is useful to look at the dynamics of the SQUID Hamiltonian in the cubic approximation. The analysis is very similar to the single junction case. Here I consider small variations in both phases such that $\gamma_i \rightarrow \gamma_{i,0} + \delta\gamma_i$. Up to a constant offset term, to third order the SQUID Hamiltonian becomes [37,39]:

$$H = \frac{p_1^2}{2m_1} + \frac{p_2^2}{2m_2} - \left(\frac{\Phi_0}{2\pi} \right)^2 \left(\frac{1}{2L_{J_1}} + \frac{1}{2L} \right) (\delta\gamma_1)^2 - \left(\frac{\Phi_0}{2\pi} \right)^2 \left(\frac{1}{2L_{J_2}} + \frac{1}{2L} \right) (\delta\gamma_2)^2 + \left(\frac{\Phi_0}{2\pi} \right)^2 \frac{1}{L} (\delta\gamma_2)(\delta\gamma_1) + \frac{1}{6} \frac{\Phi_0}{2\pi} (I_{01} \sin \gamma_1 (\delta\gamma_1)^3 + I_{02} \sin \gamma_2 (\delta\gamma_2)^3) \quad (4.19)$$

where I have discarded all higher order terms, including one term that went like $(\delta\gamma_1)^2 (\delta\gamma_2)^2$, which I considered to be 4th order. Note that the Hamiltonian can be separated into two parts:

$$H = H_{CHO} + H_{pert} \quad (4.20)$$

$$H_{CHO} = \frac{p_1^2}{2m_1} + \frac{p_2^2}{2m_2} - \left(\frac{\Phi_0}{2\pi} \right)^2 \left(\frac{1}{2L_{J_1}} + \frac{1}{2L} \right) (\delta\gamma_1)^2 - \left(\frac{\Phi_0}{2\pi} \right)^2 \left(\frac{1}{2L_{J_2}} + \frac{1}{2L} \right) (\delta\gamma_2)^2 + \left(\frac{\Phi_0}{2\pi} \right)^2 \frac{1}{L} (\delta\gamma_2)(\delta\gamma_1) \quad (4.21)$$

$$H_{pert} = \frac{1}{6} \frac{\Phi_0}{2\pi} (I_{01} \sin \gamma_1 (\delta\gamma_1)^3 + I_{02} \sin \gamma_2 (\delta\gamma_2)^3) \quad (4.22)$$

where H_{CHO} is the Hamiltonian of a coupled harmonic oscillator system and H_{pert} is a perturbation that imparts some anharmonicity to the coupled harmonic oscillator states.

Using Hamilton's equations, I can now find the equations of motion for the coupled harmonic oscillator:

$$q_2 = mp_2 \quad (4.23)$$

$$q_1 = mp_1 \quad (4.24)$$

$$-\dot{p}_1 = -\left(\frac{\Phi_0}{2\pi}\right)^2 \left(\frac{1}{L_{J1}} + \frac{1}{L}\right) \delta\gamma_1 + \left(\frac{\Phi_0}{2\pi}\right)^2 \frac{1}{L} \delta\gamma_2 \quad (4.25)$$

$$-\dot{p}_2 = -\left(\frac{\Phi_0}{2\pi}\right)^2 \left(\frac{1}{L_{J2}} + \frac{1}{L}\right) \delta\gamma_2 + \left(\frac{\Phi_0}{2\pi}\right)^2 \frac{1}{L} \delta\gamma_1 \quad (4.26)$$

Where $q_1 = C_1 V_1$ and $q_2 = C_2 V_2$. By substituting in the expression for the mass and dividing through by $\left(\frac{\Phi_0}{2\pi}\right)^2$, the equations become:

$$-C_1 \delta\ddot{\gamma}_1 = -\left(\frac{1}{L_{J1}} + \frac{1}{L}\right) \delta\gamma_1 + \frac{1}{L} \delta\gamma_2 \quad (4.27)$$

$$-C_2 \delta\ddot{\gamma}_2 = -\left(\frac{1}{L_{J2}} + \frac{1}{L}\right) \delta\gamma_2 + \frac{1}{L} \delta\gamma_1 \quad (4.28)$$

It is worth remarking on Eqs. 4.27 and 4.28. The isolation junction J_2 was added to the device in order to shunt noise away from the qubit junction J_1 . It was not intended to be part of the qubit, but note that Eq. 4.27, the equation of motion for the qubit junction phase $\delta\gamma_1$, contains a term on the right with the isolation junction phase $\delta\gamma_2$. Examining these equations, we see that the cross-coupling terms in Eq. 4.27 and 4.28 becomes small if the total loop inductance L is much larger than L_{J1} or L_{J2} [45]. However, if L is of the same order or smaller than L_{J1} or L_{J2} , then the two junctions can no longer be

considered as independent systems.

Assuming an $e^{-i\omega t}$ time dependence for both phases, I can solve this system of equations to get the resonant frequencies of the system. One finds:

$$\omega_{\pm} = \frac{1}{2} \left[\left(\frac{1}{L_1^* C_1} + \frac{1}{L_2^* C_2} \right) \pm \sqrt{\left(\frac{1}{L_1^* C_1} + \frac{1}{L_2^* C_2} \right)^2 - 4 \left[\left(\frac{1}{L_1^* C_1} \right) \left(\frac{1}{L_2^* C_2} \right) - \left(\frac{1}{L C_1} \right) \left(\frac{1}{L C_2} \right) \right]} \right]^{1/2} \quad (4.29)$$

where $\frac{1}{L_i^*} = \left(\frac{1}{L_{Ji}} + \frac{1}{L} \right)$ is the effective inductance of the i^{th} junction. From Eq. 4.29, one

can find the resonant frequency for the qubit junction, ω_- and the isolation junction, ω_+ in the normal modes model of the SQUID. Fig. 4.3 shows how the qubit resonant frequency converges on the isolated single junction solution of $\omega_p = \frac{1}{\sqrt{L_{J1} C_1}}$ when $L_{J1} \ll L$. For

this plot I assumed that $L = 1$ nH, I varied the critical current of the qubit junction in order to change L_{J1} on the x-axis, and I assumed the qubit junction was biased at $I_l = 0.8 I_{01}$. The other parameters for this plot are $C_1 = 1$ pF, $L_{J2} = 30$ pH ($I_{02} \approx 10$ μ A), $L_2 = 40$ pH, and $C_2 = 0.2$ pF.

Conventional dc SQUID phase qubits have had $L_{J1}/L < 0.1$, which means they operated in a limit where the qubit could be treated as an isolated single junction [37,45,48]. In contrast, my device PB9 had $L_{J1}/L \sim 10$ and this required me to use the full SQUID Hamiltonian to find the qubit resonant frequency (see Chapter 8). Unfortunately, I did not fully appreciate the impact of this limit on the coherence when I designed the device, as explained in Chapter 5.




Fig. 4.3 Comparison of resonant frequency of a single junction (dashed curve) and the normal modes solution from a harmonic oscillator approximation of the SQUID Hamiltonian (solid curve). The x-axis is the ratio of the Josephson inductance of the qubit junction L_{J1} to the total geometrical inductance of the SQUID L . For $L_{J1} \ll L$, the two models are in close agreement and the single junction model is a good approximation to the solution from the full SQUID Hamiltonian. For $L_{J1} \gg L$, the single junction model is a poor approximation to the solution to the full SQUID Hamiltonian. For this plot I used $L = 1$ nH, $I_J = 0.8 \cdot I_{01}$, $C_1 = 1$ pF, $L_{J2} = 30$ pH ($I_{02} \approx 10$ μ A), $L_2 = 40$ pH, and $C_2 = 0.2$ pF.

4.4 SQUID Normal Modes and Their Effect on Current Division

When the phase of the qubit junction is coupled to the phase of the isolation junction, the energy of the state is effectively stored in both junctions and the loop inductance. The original reason for introducing the second junction was to filter out noise from the current bias leads [48]. If part of the state of the qubit resides in the isolation junction (*i.e.* depends on its phase), then one will need to consider what effect noise in the isolation junction will have on the coherence of the state. One way to address this issue is to determine how variations in the bias current will affect the eigenstates of the system. The difficulty is that the oscillations of the SQUID are not oscillations of the independent junctions, but normal modes, each of which involves oscillations in the phase of both junctions.

I can find the classical normal modes of the system by going back to the original eigenvalue problem presented in Eqs. 4.27 and 4.28. I can combine these equations into matrix form:

$$\begin{pmatrix} -C_1 & 0 \\ 0 & -C_2 \end{pmatrix} \begin{pmatrix} \delta\ddot{\gamma}_1 \\ \delta\ddot{\gamma}_2 \end{pmatrix} = \begin{pmatrix} \frac{1}{L_1^*} & \frac{1}{L} \\ \frac{1}{L} & \frac{1}{L_2^*} \end{pmatrix} \begin{pmatrix} \delta\gamma_1 \\ \delta\gamma_2 \end{pmatrix} \quad (4.30)$$

I assume that the solutions of the eigenvalue problem are of the form:

$$\begin{pmatrix} \delta\gamma_1 \\ \delta\gamma_2 \end{pmatrix} = \begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix} e^{-i\omega t} \quad (4.31)$$

where α_1 and α_2 are complex. Thus Eq. 4.30 takes the form:

$$\begin{pmatrix} \frac{1}{L_1^*} - \omega^2 C_1 & \frac{1}{L} \\ \frac{1}{L} & \frac{1}{L_2^*} - \omega^2 C_2 \end{pmatrix} \begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix} = 0 \quad (4.32)$$

From Eq. 4.32, one easily obtains the eigenvalues ω_+ and ω_- given in Eq. 4.29. I will generally call the plus solution the isolation junction frequency and the minus solution and the qubit frequency, although of course these normal modes involve both junctions and the loop. Thus I can define:

$$\omega_+ = \frac{1}{2} \left[\left(\frac{1}{L_1^* C_1} + \frac{1}{L_2^* C_2} \right) + \sqrt{\left(\frac{1}{L_1^* C_1} + \frac{1}{L_2^* C_2} \right)^2 - 4 \left[\left(\frac{1}{L_1^* C_1} \right) \left(\frac{1}{L_2^* C_2} \right) - \left(\frac{1}{L C_1} \right) \left(\frac{1}{L C_2} \right) \right]} \right]^{1/2} \quad (4.33)$$

$$\omega_- = \frac{1}{2} \left[\left(\frac{1}{L_1^* C_1} + \frac{1}{L_2^* C_2} \right) - \sqrt{\left(\frac{1}{L_1^* C_1} + \frac{1}{L_2^* C_2} \right)^2 - 4 \left[\left(\frac{1}{L_1^* C_1} \right) \left(\frac{1}{L_2^* C_2} \right) - \left(\frac{1}{L C_1} \right) \left(\frac{1}{L C_2} \right) \right]} \right]^{1/2} \quad (4.34)$$

The eigenvectors of the eigenvalue problem in Eq. 4.32 are:

$$\mathbf{n}_Q = \begin{pmatrix} \alpha_{1,Q} \\ \alpha_{2,Q} \end{pmatrix} e^{-i\omega_Q t} \quad \text{and} \quad \mathbf{n}_I = \begin{pmatrix} \alpha_{1,I} \\ \alpha_{2,I} \end{pmatrix} e^{-i\omega_I t} \quad (4.35)$$

where α 's are defined from the solution of the eigenvalue problem as:

$$\alpha_{1,Q} = \left(\frac{2\pi}{\Phi_0} \right) \sqrt{\frac{1}{C_1 + C_2 \left(\frac{L}{L_1^*} - \omega_Q^2 L C_1 \right)^2}} \quad (4.36)$$

$$\alpha_{2,Q} = \left(\frac{L}{L_1^*} - \omega_Q^2 L C_1 \right) \alpha_{1,Q} \quad (4.37)$$

$$\alpha_{1,I} = \left(\frac{2\pi}{\Phi_0} \right) \sqrt{\frac{1}{C_1 + C_2 \left(\frac{L}{L_1^*} - \omega_I^2 L C_1 \right)^2}} \quad (4.38)$$

$$\alpha_{2,I} = \left(\frac{L}{L_1^*} - \omega_I^2 LC_1 \right) \alpha_{1,I} \quad (4.39)$$

Goldstein discusses this subject thoroughly in his text on classical dynamics, and he calls the eigenvectors of an oscillating system with multiple degrees of freedom the normal modes [21]. The two junction SQUID is analogous in many ways to two pendula coupled together by a spring. In the case where the two pendula are the same length, the two normal modes of the system are the two pendula swinging together and the two swinging opposite to one another. An asymmetric dc SQUID is more like two pendula of different lengths, though, because the resonant frequencies of each junction are very different. In such a case, the normal modes have one pendulum oscillating very widely with the other only oscillating slightly and *vice versa*.

The relative sizes of the α 's in Eq. 4.35 give us physical insight into the amplitude of the oscillations in each junction for each normal mode of the system. η_Q is called the “qubit” normal mode because $\alpha_{1,Q} \gg \alpha_{2,Q}$, and η_I is the “isolation” mode because $\alpha_{2,I} \gg \alpha_{1,I}$. In the qubit mode, it is mainly the qubit junction's phase that is oscillating while the isolation junction's phase only oscillates slightly.

Note also these normal mode solutions η_Q and η_I are approximations to the solutions to the full Hamiltonian and are accurate only for small deviations from the potential minimum.

The normal modes can be used to see how changes in the bias current divide between the two modes of the SQUID and how this division is affected by the ratio L/L_J . Consider the case where bias current I_b is applied to the SQUID. The current will divide such that:

$$I_b = I_1 + I_2 = I_{01} \sin \gamma_1 + I_{02} \sin \gamma_2 \quad (4.41)$$

If the bias current fluctuates such that $I_b \rightarrow I_b + \delta I_b$, then Eq. 4.41 takes the form:

$$I_b + \delta I_b = I_{01} \sin(\gamma_1 + \delta \gamma_1) + I_{02} \sin(\gamma_2 + \delta \gamma_2). \quad (4.42)$$

Expanding Eq. 4.42 to first order in $\delta \gamma$ and rearranging, one can arrive at:

$$\delta I_b = \left(\frac{\Phi_0}{2\pi} \right) \left(\frac{1}{L_{J1}} \delta \gamma_1 + \frac{1}{L_{J2}} \delta \gamma_2 \right) = \delta I_1 + \delta I_2. \quad (4.43)$$

According to Goldstein, a transformation can be made from the junction phase coordinate system to the normal modes coordinate system by using a matrix \mathbf{A} that is comprised of the eigenvectors [21]:

$$\mathbf{A} = \begin{pmatrix} \alpha_{1,Q} & \alpha_{1,I} \\ \alpha_{2,Q} & \alpha_{2,I} \end{pmatrix}. \quad (4.44)$$

The coordinate transformation is:

$$\begin{pmatrix} \delta \gamma_1 \\ \delta \gamma_2 \end{pmatrix} = \mathbf{A} \begin{pmatrix} \eta_Q \\ \eta_I \end{pmatrix} \quad (4.45)$$

Using the transformation in Eq. 4.36, the small phase terms can be replaced by small changes in the amplitudes of the normal modes:

$$\delta I_b = \left(\frac{\Phi_0}{2\pi} \right) \left[\frac{1}{L_{J1}} (\alpha_{1,Q} \eta_Q + \alpha_{1,I} \eta_I) + \frac{1}{L_{J2}} (\alpha_{2,Q} \eta_Q + \alpha_{2,I} \eta_I) \right]. \quad (4.46)$$

Expanding, I find:

$$\delta I_b = \left(\frac{\Phi_0}{2\pi} \right) \left(\frac{\alpha_{1,Q}}{L_{J1}} + \frac{\alpha_{2,Q}}{L_{J2}} \right) \eta_Q + \left(\frac{\Phi_0}{2\pi} \right) \left(\frac{\alpha_{1,I}}{L_{J1}} + \frac{\alpha_{2,I}}{L_{J2}} \right) \eta_I = \delta I_Q + \delta I_I \quad (4.47)$$

where δI_Q is the effective current driving the qubit normal mode and δI_I is the effective current driving the isolation mode. It is important to realize that if δI_Q vanishes, then the qubit mode will be well-isolated from small fluctuations in the bias current.

The current isolation factor for the qubit can then be found from:

$$\frac{\delta I_b}{\delta I_Q} = 1 + \frac{(\alpha_{1,I} L_{J2} + \alpha_{2,I} L_{J1})}{(\alpha_{1,Q} L_{J2} + \alpha_{2,Q} L_{J1})} \frac{\eta_I}{\eta_Q}. \quad (4.48)$$

Using \mathbf{A}^{-1} to transform back to the $\delta\gamma$ coordinates, I can also write:

$$\frac{\eta_I}{\eta_Q} = \frac{\mathbf{A}_{21}^{-1} \delta\gamma_1 + \mathbf{A}_{22}^{-1} \delta\gamma_2}{\mathbf{A}_{11}^{-1} \delta\gamma_1 + \mathbf{A}_{12}^{-1} \delta\gamma_2} = \frac{\mathbf{A}_{21}^{-1} \frac{\delta\gamma_1}{\delta\gamma_2} + \mathbf{A}_{22}^{-1}}{\mathbf{A}_{11}^{-1} \frac{\delta\gamma_1}{\delta\gamma_2} + \mathbf{A}_{12}^{-1}} \quad (4.49)$$

The missing piece of the puzzle is $\delta\gamma_1/\delta\gamma_2$. According to Eq. 4.43, I can write:

$$\frac{\delta\gamma_1}{\delta\gamma_2} = \frac{L_{J1}}{L_{J2}} \frac{\delta I_1}{\delta I_2} \quad (4.50)$$

I also know from Eq. 4.2 that current will divide between the two branches of the SQUID according to their inductances, *i.e.*:

$$\frac{\delta I_b}{\delta I_1} = \left(\frac{L_1 + L_2 + L_{J1} + L_{J2}}{L_2 + L_{J2}} \right) \quad (4.51)$$

$$\frac{\delta I_b}{\delta I_2} = \left(\frac{L_1 + L_2 + L_{J1} + L_{J2}}{L_1 + L_{J1}} \right) \quad (4.52)$$

Therefore,

$$\frac{\delta I_1}{\delta I_2} = \left(\frac{L_2 + L_{J2}}{L_1 + L_{J1}} \right) \equiv D \quad (4.53)$$

where D is the current division ratio between the arms of the SQUID.

From the above, I can now obtain an expression for the current power isolation factor for the qubit mode in the dc SQUID design:

$$r_I = \left(\frac{\delta I_b}{\delta I_Q} \right)^2 = \left[1 + \frac{(\alpha_{1,I} L_{J2} + \alpha_{2,I} L_{J1})}{(\alpha_{1,Q} L_{J2} + \alpha_{2,Q} L_{J1})} \frac{\mathbf{A}_{21}^{-1} \frac{L_{J1}}{L_{J2}} D + \mathbf{A}_{22}^{-1}}{\mathbf{A}_{11}^{-1} \frac{L_{J1}}{L_{J2}} D + \mathbf{A}_{12}^{-1}} \right]^2. \quad (4.54)$$

Equation 4.54 is fairly opaque. In this case it is easiest to evaluate the expression numerically to see how the qubit's isolation depends on L_{JI} and L . Figure 4.4 shows a plot of the current power isolation factor for the qubit normal mode (see solid curve). For comparison, I also show the isolation factor from Eq. 4.4 (see dashed curve), which assumed a single junction model for the qubit. Note that the two curves converge for L_{JI} much smaller than L - this is the standard weak-coupling limit for the SQUID. Also note that although the single JJ model (dashed curve) predicts higher isolation when L_{JI} is large, it is a misleading result because the single JJ model does not accurately describe the dynamics of the qubit in the strongly coupled regime. As discussed earlier (see Fig. 4.3 for example) the single junction model does not accurately describe the qubit when L_{JI} is large compared to L . Considering Fig. 4.4, we can see that the LJ-isolation described in Sect. 4.1 remains valid at low L_{JI} , and that a more accurate description of the isolation factor in the large L_{JI} limit is:

$$r_I \approx \left(\frac{\Delta I}{\Delta I_1} \right)^2 \approx \left(\frac{L_1 + L_2 + L_{J2}}{L_2 + L_{J2}} \right)^2. \quad (4.55)$$

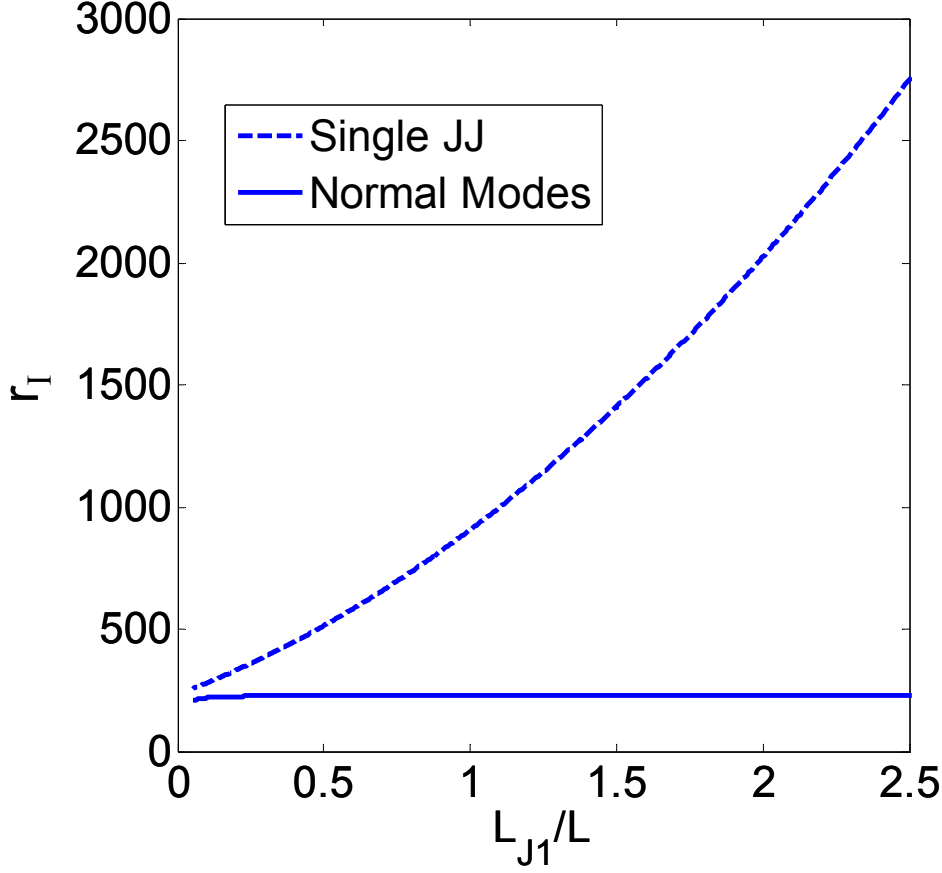


Fig. 4.4 Current power isolation factor for the qubit in the single junction model (dashed curve) and the normal modes model (solid curve). Because the qubit normal mode of the SQUID depends to some extent on the phase of the isolation junction, the isolation of the normal mode is diminished compared to the single junction model. However, the plot shows that the isolation predicted by the normal modes model (about 250) is about the same as that predicted by the single junction model in the parameter regime where the single junction model was valid, *i.e.* the single junction model is only valid for $L_{J1} \ll L$ and in that regime the isolation factor is about 250. For this plot I used $L = 1$ nH, $I_I = 0.8 \cdot I_{0I}$, $C_I = 1$ pF, $L_{J2} = 30$ pH ($I_{02} \approx 10$ μ A), $L_2 = 40$ pH, and $C_2 = 0.2$ pF.

4.5 Solving the Full SQUID Hamiltonian with a Finite Difference Method

The frequencies obtained in Eq. 4.29 for the resonances of the SQUID were found by treating the junctions in the SQUID as simple LC oscillators. Although this gives valuable insight into the dynamics of the system, these predictions for the $|0\rangle$ to $|1\rangle$ transition frequencies of the qubit are clearly approximate. One can obtain a more accurate result for the transition frequencies by using a cubic approximation for the SQUID Hamiltonian. Again, analytical results can be obtained, but the matrix elements for the cubic perturbation terms are messy to calculate because of the coupling of the junctions. Ben Cooper has written up a description of this technique and interested readers can find a discussion of this approach in [11]. Here, I describe a numerical approach to finding the energies of the SQUID Hamiltonian called the finite difference method. In this method, the potential is divided into a grid of values for γ_1 and γ_2 , as shown in Fig. 4.5, and the derivatives in the Hamiltonian are replaced by approximately equivalent difference quotients [51]. I construct the approximate Hamiltonian matrix, and the eigenvectors and eigenfunctions of this matrix then yield the approximate eigenstates and eigenenergies of the system. The resulting eigenstates can be displayed on the grid of γ_1 's and γ_2 's and gives a picture of the state in the potential well. In this method, each space on the grid (see Fig. 4.5) represents a state $|i, j\rangle$. In this basis the wavefunctions of the system are

$$|\Psi\rangle = \sum_{i,j} c_{ij} |i, j\rangle \quad (4.56)$$

where the c_{ij} 's are complex numbers that satisfy $\sum_{i,j} |c_{ij}|^2 = 1$, and $|i\rangle$ and $|j\rangle$ are eigenkets of the γ_1 and γ_2 operators such that:

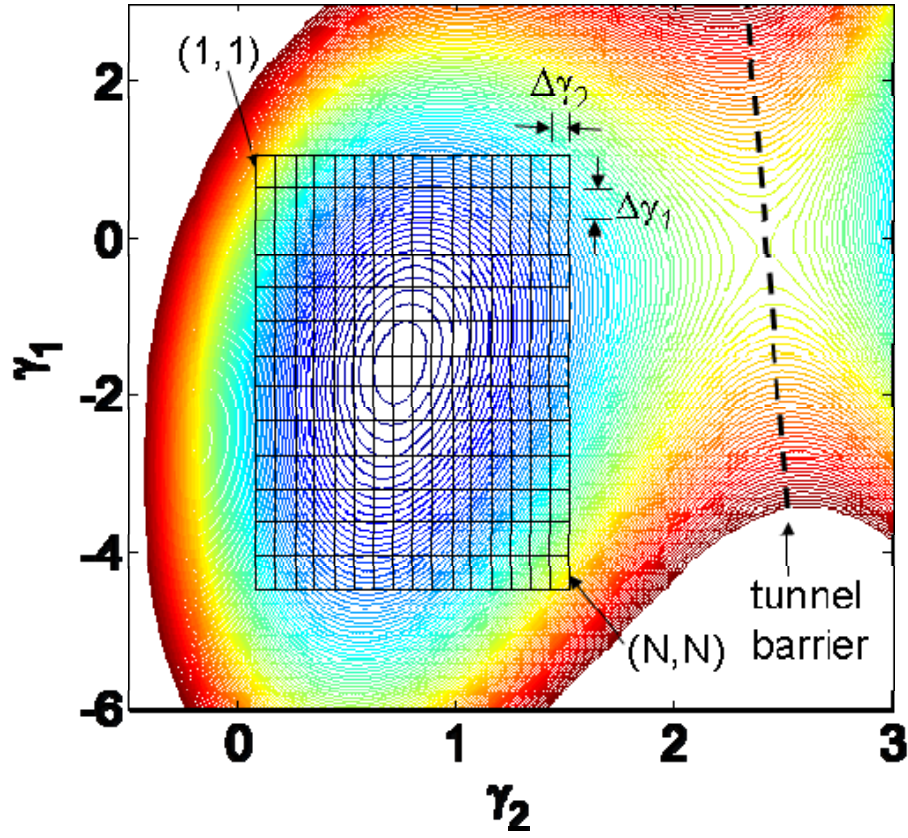


Fig. 4.5 Plot of the potential energy well of the SQUID with a grid drawn over the potential minimum. $\Delta\gamma_1$ and $\Delta\gamma_2$ are the grid spacings in the γ_1 and γ_2 directions. $\Delta\gamma_1$ and $\Delta\gamma_2$ need not be equal, but the grid must be $N \times N$. The Matlab code in Appendix A centers the grid on the potential minimum.

$$\hat{\gamma}_1|i\rangle = \gamma_{1,i}|i\rangle \quad \text{and} \quad \hat{\gamma}_2|j\rangle = \gamma_{2,j}|j\rangle. \quad (4.57)$$

where $\gamma_{1,i}$ and $\gamma_{2,j}$ are the values of γ_1 and γ_2 at the point (i,j) on the grid.

The Hamiltonian of the system can be written in terms of γ_1 and γ_2 as [37]:

$$\hat{H} = -4E_{c1} \frac{d^2}{d\gamma_1^2} - 4E_{c2} \frac{d^2}{d\gamma_2^2} + U(\gamma_1, \gamma_2) \quad (4.58)$$

where $E_{ci} = \frac{e^2}{2C_i}$ is the charging energy of the i -th junction. The potential energy matrix

is simple to construct because its diagonal elements are just:

$$\langle i, j | U(\gamma_1, \gamma_2) | i, j \rangle = U(\gamma_{1,i}, \gamma_{2,j}) \quad (4.59)$$

and all of the off diagonal elements in the \mathbf{U} matrix are:

$$\langle k, m | U(\gamma_1, \gamma_2) | i, j \rangle = 0 \quad (4.60)$$

where $k \neq i$ and $m \neq j$. Therefore, if there are N $|i\rangle$ states and N $|j\rangle$ states then \mathbf{U} is a diagonal $N^2 \times N^2$ matrix of the form:

$$\mathbf{U} = \begin{pmatrix} U(\gamma_{1,1}, \gamma_{2,1}) & 0 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & U(\gamma_{1,1}, \gamma_{2,2}) & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & 0 & U(\gamma_{1,1}, \gamma_{2,3}) & \dots & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 & U(\gamma_{1,1}, \gamma_{2,N}) & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 & 0 & U(\gamma_{1,2}, \gamma_{2,1}) & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & U(\gamma_{1,N}, \gamma_{2,N}) \end{pmatrix} \quad (4.61)$$

Here I finish stepping through the $|j\rangle$ states before stepping to the next $|i\rangle$ state.

The finite difference comes in when dealing with the kinetic energy terms. In Dirac notation, we replace the double-derivatives of the $|i, j\rangle$ states with the approximations [51]:

$$\frac{d^2}{d\gamma_1^2}|i, j\rangle \approx \frac{1}{(\Delta\gamma_1)^2}(|i-1, j\rangle - 2|i, j\rangle + |i+1, j\rangle) \quad (4.62)$$

and

$$\frac{d^2}{d\gamma_2^2}|i, j\rangle \approx \frac{1}{(\Delta\gamma_2)^2}(|i, j-1\rangle - 2|i, j\rangle + |i, j+1\rangle) \quad (4.63)$$

where $\Delta\gamma_1$ and $\Delta\gamma_2$ are the grid spacings in the γ_1 and γ_2 directions, respectively. The approximation gets better as the grid spacing is reduced. Equation 4.62 and 4.63 show that there will be diagonal elements contributed by the kinetic energy and contributions where one index has changed by one and the other index has remained the same.

For convenience I will define the kinetic energy operator \hat{T} by:

$$\hat{T} = -4E_{c1} \frac{d^2}{d\gamma_1^2} - 4E_{c2} \frac{d^2}{d\gamma_2^2} \quad (4.64)$$

The contributions to the Hamiltonian on the diagonal are:

$$T_{i,j} \equiv \langle i, j | \hat{T} | i, j \rangle = \frac{8E_{c1}}{(\Delta\gamma_1)^2} + \frac{8E_{c2}}{(\Delta\gamma_2)^2} \quad (4.65)$$

and the other non-zero matrix elements are:

$$T_{i+1,j} \equiv \langle i+1, j | \hat{T} | i, j \rangle = -\frac{4E_{c1}}{(\Delta\gamma_1)^2} \quad (4.66)$$

$$T_{i-1,j} \equiv \langle i-1, j | \hat{T} | i, j \rangle = -\frac{4E_{c1}}{(\Delta\gamma_1)^2} \quad (4.67)$$

$$T_{i,j+1} \equiv \langle i, j+1 | \hat{T} | i, j \rangle = -\frac{4E_{c2}}{(\Delta\gamma_2)^2} \quad (4.68)$$

$$T_{i,j-1} \equiv \langle i, j-1 | \hat{T} | i, j \rangle = -\frac{4E_{c2}}{(\Delta\gamma_2)^2} \quad (4.69)$$

If I order the states in the same way as I did when building \mathbf{U} , then \mathbf{T} will be composed of three $N \times N$ sub-matrices which can be written as:

$$\mathbf{T}_0 = \begin{pmatrix} T_{i,j} & T_{i,j+1} & 0 & 0 & 0 & \dots & 0 & 0 \\ T_{i,j-1} & T_{i,j} & T_{i,j+1} & 0 & 0 & \dots & 0 & 0 \\ 0 & T_{i,j-1} & T_{i,j} & T_{i,j+1} & 0 & \dots & 0 & 0 \\ 0 & 0 & T_{i,j-1} & \ddots & \ddots & 0 & 0 & 0 \\ 0 & 0 & 0 & \ddots & T_{i,j} & T_{i,j+1} & 0 & 0 \\ \vdots & \vdots & \vdots & 0 & T_{i,j-1} & T_{i,j} & T_{i,j+1} & 0 \\ 0 & 0 & 0 & 0 & 0 & T_{i,j-1} & T_{i,j} & T_{i,j+1} \\ 0 & 0 & 0 & 0 & 0 & 0 & T_{i,j-1} & T_{i,j} \end{pmatrix} \quad (4.70)$$

$$\mathbf{T}_+ = \begin{pmatrix} T_{i+1,j} & 0 & \dots & 0 \\ 0 & T_{i+1,j} & 0 & \vdots \\ \vdots & 0 & \ddots & 0 \\ 0 & \dots & 0 & T_{i+1,j} \end{pmatrix} \quad (4.71)$$

$$\mathbf{T}_- = \begin{pmatrix} T_{i-1,j} & 0 & \dots & 0 \\ 0 & T_{i-1,j} & 0 & \vdots \\ \vdots & 0 & \ddots & 0 \\ 0 & \dots & 0 & T_{i-1,j} \end{pmatrix} \quad (4.72)$$

\mathbf{T} itself can then be written in terms of Eqs. 4.71-4.73 as:

$$\mathbf{T} = \begin{pmatrix} \mathbf{T}_0 & \mathbf{T}_+ & \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \\ \mathbf{T}_- & \mathbf{T}_0 & \mathbf{T}_+ & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{T}_- & \mathbf{T}_0 & \mathbf{T}_+ & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{T}_- & \ddots & \ddots & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \ddots & \mathbf{T}_0 & \mathbf{T}_+ & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \vdots & \mathbf{0} & \mathbf{T}_- & \mathbf{T}_0 & \mathbf{T}_+ & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{T}_- & \mathbf{T}_0 & \mathbf{T}_+ \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{T}_- & \mathbf{T}_0 \end{pmatrix} \quad (4.73)$$

where every $\mathbf{0}$ stands for an $N \times N$ matrix of zeroes. $\mathbf{T}_+ = \mathbf{T}_-$ and \mathbf{T}_0 is Hermitian, which makes $\mathbf{H} = \mathbf{T} + \mathbf{U}$ Hermitian. One can now find the eigenenergies of the system by diagonalizing \mathbf{H} [40].

The linear algebra for diagonalizing a matrix is described in Sect. 1.5 of [40]. In practice, I used the “eig” function in Matlab to find the eigenvalues and eigenvectors of the matrix \mathbf{H} . The Matlab code I used is given in Appendix A and it involved diagonalizing a $49^2 \times 49^2$ matrix. It took less than 10 seconds to solve for the energies and about one minute to find the eigenstates.

The accuracy of the model depends on the size and the spacing of the grid. A larger grid will allow the model to represent more of the potential well. This is especially important for getting accurate results for the energies of the higher states in the well. If the eigenfunction of a state has significant probability amplitude outside the grid used for the simulation, then the eigenenergy of that state given by the model will be inaccurate. Also, if the grid spacing is too coarse to capture spatial variations in the wavefunction, then the energy will be inaccurate. Therefore, the grid must be wide enough and have a sufficiently fine spacing.

Figure 4.6(a) shows a 2D false-color plot of the $|0\rangle$ to $|1\rangle$ transition frequency f_{01} with respect to the dimension of the (γ_1, γ_2) grid n and the spacing of the grid $\Delta\gamma$. There are two trends to this graph. First, in the y-direction, there is a threshold in n after which f_{01} remains constant along the y-direction. This threshold occurs earlier if $\Delta\gamma$ is a big step size because it takes less steps to cover the potential well.

Figures 4.6(b) and 4.6(c) show the wavefunctions of the state in the regions before and after the threshold value for n was reached, respectively. The second trend is




Fig. 4.6 (a) 2D false-color plot showing the $|0\rangle$ to $|1\rangle$ transition frequency f_{01} obtained from the finite difference method with an $n \times n$ grid and a grid spacing $\Delta\gamma$. For each $\Delta\gamma$ a threshold value of n exists beyond which the model settles into a steady solution. The bright red and dark blue regions are where the model fails. (b) Wavefunction amplitude of the $|1\rangle$ state of the failed model with $n = 25$ and $\Delta\gamma = 0.03$. (c) Wavefunction amplitude given by a successful simulation with $n = 41$ and $\Delta\gamma = 0.03$. For this plot I used $I_{01} = 110$ nA, $I_{02} = 8.6$ μ A, $L_1 = 0.7$ nH, $L_2 = 40$ pH, $C_1 = 1.52$ pF, $C_2 = 0.2$ pF, $I_b = 6.8$ μ A.

that the frequencies that the model predicts depend on the $\Delta\gamma$ for that vertical slice. The fact that the energy steadily decreases as the grid spacing increases means that the model is missing some features of the wavefunctions at coarser spacings. The most accurate values come from using the finest step possible, but the PC I used could not handle more than about a 49×49 grid in a reasonable time. Making the graph for Fig. 4.6(a) took about 2-3 hours. According to Fig. 4.6(a), using a spacing of 0.03 radians and a 39×39 grid gives a reasonably accurate value for f_{01} , and the simulation only takes a few seconds.

The model also gives the quantum mechanical wavefunctions of the system. Figure 4.7 shows results from a simulation with $n = 51$ and $\Delta\gamma = 0.03$. Fig. 4.7(a) shows the wavefunction of the ground state and 4.7(b) shows that of the first excited state. In terms of the normal modes of the SQUID these states can be thought of as the $|00\rangle$ and the $|10\rangle$ in the $|n_Q n_I\rangle$ basis where n_Q and n_I are the mode numbers of the qubit and isolation modes, respectively. The ground state has a single anti-node and the first excited state has two anti-nodes with a node in the middle. Note that the first excited state does not lie exactly along the γ_I direction because of coupling between the two junction coordinates. Figure 4.7(c) and 4.7(d) show the higher mode number states of the qubit mode: $|20\rangle$ and $|70\rangle$. Each state shows more nodes and spans a larger region of the potential well.

In practice, the model provides a number of eigenfunctions equal to the dimension n of the grid. As an example of a higher mode, Fig. 4.7(e) shows the wavefunction of the 10^{th} state, which is the first excited state of the isolation junction, *i.e.* this is mode $|01\rangle$. Note that the state lies mostly along the γ_2 direction. As another example, Fig. 4.7(f)




Fig. 4.7 Wavefunctions returned by the finite difference method. The basis is $|n_Q n_I\rangle$ where n_Q and n_I are occupation numbers of the qubit and isolation junction modes. The color axes' scales are different, but—except for the $|00\rangle$ state—blue means negative probability amplitude and red means positive.

shows an even higher energy state, the $|32\rangle$ state.

Using the finite difference method to solve the full SQUID Hamiltonian turned out to be a relatively quick and versatile technique. I found the technique to be particularly useful for understanding the quantum behavior of device PB9. In Chapter 8-10, I will show how I used it to fit the measured energy spectrum of the qubit, extract the device parameters, and understand decoherence. The technique was also convenient in that I could get the energies of all of the qubit eigenstates in one step. This was useful for exploring the anharmonicity of the device, as I discuss in Chapter 10.

4.6 Conclusions

In this chapter, I discussed why a simple current-biased Josephson junction cannot be used as a qubit by itself; It is vulnerable to decoherence caused by dissipation and noise due to the current bias line. By using two junctions in a dc SQUID configuration, decoherence from the leads can be reduced. In the dc SQUID phase qubit, one junction acts as the qubit, and the other junction (the isolation junction) and loop inductance act to shunt noise from the bias leads away from the qubit. The isolation junction is necessary in order to use a switching voltage readout technique. The design works best if the inductance of the arm of the SQUID loop with the qubit junction on it is large compared to the inductance of the arm with the isolation junction. If the loop inductance is sufficiently large compared to the Josephson inductances of the junctions, then the qubit can be well-approximated as a single junction, as described in Chapter 2. However, if the loop inductance is not much larger than the inductance of either of the junctions, then the two junctions in the SQUID behave like coupled oscillators instead of independent non-

linear resonators. In that case the eigenstates of the qubit involve coordinated motion of both junction phases in the SQUID.

The dc SQUID phase qubit has a design that provides some protection from noise, but at the same time adds quite a bit of complexity to the quantum mechanical nature of the qubit. My analysis of the normal modes of the SQUID showed that the qubit mode in the coupled system can still be well-isolated from noise on the bias lines, but not quite as well as would be expected from a single junction model of the qubit. An approximation to the quantized energy levels of the full SQUID can be found by using a finite difference method on the SQUID Hamiltonian.

Chapter 5

Designing a dc SQUID Phase Qubit for Reduced Decoherence

Over the last decade, researchers have identified several sources of decoherence in superconducting qubits and developed techniques for dealing with many of them. In particular, to minimize decoherence, the materials used to build the device and the circuit parameters of the device must be chosen carefully [15-17,18,43,44,48,49,52]. In this chapter, I discuss the considerations that I used for choosing the materials and the parameters for my phase qubit.

I first present the models I used for estimating the effects that the bias lines would have on dissipation and dephasing in the qubit. Choosing device parameters is often a balancing act between limiting one source of decoherence and strengthening the effects of another. In some cases, choosing a certain value for a circuit parameter (*e.g.* the critical current of the qubit junction) may increase the estimated coherence time of a qubit, but also cause other, unintended, and possibly deleterious consequences. In designing device PB9, I took a comprehensive approach and modeled the effects of all the known sources of decoherence including dielectric loss, dissipation from the leads, critical current noise and $1/f$ flux noise. The goal for this design was for the qubit to achieve a relaxation time T_1 and a coherence time T_2 greater than 1 μs .

5.1 Designing to Reduce Dissipation

I can model the dissipation in the phase qubit junction as being an effective resistance in parallel with the Josephson inductance and the capacitance shunting the junction [26,38,52]. Figure 5.1 shows an equivalent circuit model where all sources of




Fig. 5.1 Equivalent circuit model for loss in a Josephson junction phase qubit.

Each source of loss contributes to the total loss represented by $R_{eff}(\omega)$.

dissipation have been included in a single resistive element R_{eff} . The known sources of dissipation arise from 3 bias lines and 4 capacitors in my device. Each source had an effective resistance R_i associated with it and the total R_{eff} was the parallel combination of the resulting seven effective resistances

$$\frac{1}{R_{eff}} = \sum_{i=1}^7 1/R_i(\omega_{01}). \quad (5.1)$$

In Eq. 5.1, I note that in general the effective resistances R_i associated with each dissipation source will be frequency dependent. Since I was concerned with the dissipation of the qubit when it was in the excited state $|1\rangle$, I evaluated each source's effective resistance at the qubit frequency ω_{01} [26].

As I discussed in Chapter 3, the energy relaxation time T_1 of a phase qubit is [26,52]:

$$T_1 = R_{eff} C_1 = R_{eff} (C_J + C_x) \quad (5.2)$$

where C_1 is the total capacitance shunting the qubit junction, C_J is the geometrical capacitance of the Josephson junction and C_x is the capacitance of the interdigitated capacitor (IDC) that I added across the device. If both R_{eff} and the combination of C_J and C_x are large, then T_1 will be long.

5.1.1 *Eliminating Dielectric Loss*

Martinis *et al.* first recognized that dielectric loss can be a major problem for Josephson junction qubits [15]. Dielectric loss is known to be due to charged two level systems (TLS's) that arise from ionic defects in a solid [55]. In superconducting qubits,

such charge defects can couple to the qubit in a number of ways and cause it to relax or dephase [15-17].

To model dielectric loss, I consider a parallel plate capacitor C that has a lossy dielectric. This lossy capacitor is equivalent electrically to a lossless capacitor C_o in parallel with a frequency dependent resistance $R(\omega)$ (see Fig. 5.2). The impedance of the lossy capacitor is:

$$Z(\omega) = \frac{1}{i\omega C} = \frac{d}{i\omega\epsilon_0\epsilon A} \quad (5.3)$$

where A is the area of each plate, d is the distance between them, ϵ_0 is the permittivity of free space and ϵ is the dielectric constant. For a lossy dielectric, the dielectric constant will have both real and imaginary parts ϵ_1 and ϵ_2 , respectively, and we can write the impedance as:

$$Z(\omega) = \frac{d}{i\omega\epsilon_0 A(\epsilon_1 - i\epsilon_2)} \quad (5.4)$$

Since we are comparing this impedance to that of a lossless capacitor in parallel with a resistor, it is helpful to look at the inverse of $Z(\omega)$. One finds:

$$\frac{1}{Z(\omega)} = \frac{i\omega\epsilon_0 A}{d} + \frac{\omega\epsilon_0 A}{d} \left(\frac{\epsilon_2}{\epsilon_1} \right) = i\omega C_o + \frac{1}{R(\omega)} \quad (5.5)$$

The first term is just the admittance due to an ideal capacitor C_o and the second term is the admittance due to a frequency dependent resistor.

Note that the loss term depends on the ratio ϵ_1/ϵ_2 . The loss tangent of a material is defined by $\tan \delta = \epsilon_2/\epsilon_1$, and so we can write:

$$R(\omega) = \frac{1}{\omega C_o \tan \delta} \quad (5.6)$$




Fig. 5.2 (a) A capacitor C with a lossy dielectric can be modeled as (b) an equivalent circuit with an ideal loss-less capacitor C_0 and a shunting frequency-dependent resistor $R(\omega)$.

According to Eq. 5.5, the larger $R(\omega)$ is, the more the capacitor looks like a lossless capacitor. This implies that one wants to use capacitors that are made from dielectrics that have a loss tangent that is as small as possible in order to minimize dissipation.

The above equation for loss in a parallel plate capacitor can be generalized to describe dielectric loss in any geometry. In particular, dielectric loss in the qubit substrate and insulation layers associated with wiring will contribute to loss in the qubit. In the remainder of this section, I consider each of the possible dielectric elements in turn.

First, I used a sapphire substrate instead of a more conventional Si/SiO₂ substrate. Amorphous silicon dioxide (SiO₂) has been reported to have a loss tangent on the order of 10^{-3} [15] at low temperature and low power, whereas crystalline sapphire has been reported to have $\tan\delta \sim 10^{-6}$ [56]. A loss tangent of 10^{-3} is an unacceptably large loss because even a small capacitive coupling to the SiO₂ dielectric could have severely limited the qubit's coherence time.

To minimize the loss from other dielectric layers, I constructed as much of the device as possible on one layer. In particular, I did not use any wiring dielectric layers. This is not a new strategy. Hanhee Paik compared the coherence times of samples created using the Hypres multilayer process and those created using a single layer process [49]. Although both types were created on Si/SiO₂ substrates, she found that those created by Hypres showed an energy relaxation time, $T_1 \approx 10$ ns, which was at least a factor of 2 lower than that of the single layer devices. Although this was not definitive proof that wire insulation layers were the problem, it was suggestive.

In my device, the “single layer” process was a double-angle evaporation that actually involved two depositions of aluminum with an oxidization step in between. The aluminum was evaporated at two angles so that there was an overlap region that formed Al/AlO_x/Al Josephson junctions (see Sect. 6.9). This meant that there was some aluminum oxide between the two Al layers in the device that may have contributed to loss. The extra aluminum oxide was a concern, and although it may be necessary in the future to revamp the fabrication process to eliminate any extra overlaps, its presence was a necessary evil on this device. Of course once an aluminum device is exposed to atmosphere, it naturally grows an oxide layer, so some oxide was unavoidable. Since double-angle evaporation seemed to produce qubits with better coherence times than those made with other fabrication processes, I chose to use double-angle evaporation to make my device.

5.1.2 Reducing the Loss by Adding a Capacitor Across the Junction

As I noted above, Martinis *et al.* reported that the aluminum oxide in the junction barrier has a very high loss tangent [15]. Since an exact value for the loss tangent $\tan\delta_j$ of the dielectric in our Al/AlO_x/Al Josephson junctions is not known, I assumed that $\tan\delta_j \sim 10^{-3}$, which is the value Martinis *et al.* reported for Al/AlO_x/Al capacitors at low temperature and low power [15]. Given the relatively large value of this $\tan(\delta)$, I needed to do something to minimize the impact this would have on the loss in the device. To see how this was possible, notice that the effective resistance $R(\omega)$ associated with a dielectric in a parallel plate junction is:

$$R(\omega) = \frac{1}{\omega C_J \tan \delta_J} = \frac{d}{\omega \epsilon \epsilon_o A_J \tan \delta_J} \quad (5.7)$$

where C_J is the geometrical capacitance of the junction, d is the distance between the junction electrodes, and A_J is the area of the junction.

From Eq. 5.7, we see that the resistance can be large if the area of the junction is made very small. Of course making the junction very small is not without consequences, and one can begin to see the dilemma that one faces when designing a qubit. For example, I note that the $|0\rangle$ to $|1\rangle$ transition frequency f_{01} of the qubit can be approximated as:

$$f_{01} \approx \frac{1}{2\pi \sqrt{L_J C_J}} \quad (5.8)$$

where L_J is the Josephson inductance and C_J is the capacitance of the junction. Typically, we will want the resonance frequency to be a few GHz, to avoid thermal population at 20 mK, and we will want to select C_J and L_J accordingly. Although C_J and L_J depend on the area of the junction, I show next that the product $L_J C_J$ may or may not depend on the area, depending on what is held fixed.

For our Al/AlO_x/Al tunnel junctions, the effective capacitance per unit area can be estimated from experiment. For example, Tauno Palomaki measured a capacitance of 0.4 pF for a qubit with a 16 (μm)² junction area [52]. From this result, I can write:

$$C_J = \frac{1}{40} \frac{\text{pF}}{(\mu\text{m})^2} A_J \quad (5.9)$$

where A_J is the area of the junction. Next, recall that the Josephson inductance L_J is:

$$L_J = \frac{\Phi_0}{2\pi I_0 \cos \gamma} \quad (5.10)$$

where I_0 is the critical current of the junction, $\Phi_0 = h/2e$ is the superconducting flux quantum, and γ is the phase difference across the junction. The critical current depends on the area according to:

$$I_0 = J_c A_J \quad (5.11)$$

where J_c is the critical current density.

From Eqs. 5.9-5.11 we see that the $|0\rangle$ to $|1\rangle$ transition frequency f_{01} of the qubit can then be written as:

$$f_{01} \approx \sqrt{\frac{J_c \cos \gamma}{2\pi\Phi_0} \frac{40(\mu m)^2}{pF}}. \quad (5.12)$$

Thus we see that f_{01} is independent of the area of the junction. Notice however that this result assumes that J_c is independent of area. In fact, this is not likely to be the case. I needed to achieve a certain value for the critical current as we will see below. The result is that smaller junctions will tend to require larger critical current density and this will tend to produce a larger f_{01} than is desirable; f_{01} needed to be between 1 and 20 GHz to be within the range of our microwave source.

In practice, I needed the critical current I_0 to be greater than 100 nA for two reasons. First, it was difficult to source the bias current such that it was well defined to fractions of a nA. Second, the Josephson energy E_J , which can be written as:

$$E_J = \frac{\Phi_0}{2\pi} I_0 \quad (5.13)$$

needed to be large compared to the charging energy E_C , where:

$$E_C = \frac{e^2}{2C} \quad (5.14)$$

where C is the shunting capacitance of the Josephson junction in the RCSJ model. In order for the phase to be well defined, one needs to have $E_J \gg E_C$ [37,57]. This provides an additional constraint on I_0 (and f_{01}), limiting the options for improving $R(\omega)$ simply by changing the junction area.

One way to get around these constraints was to add a separate shunting capacitor across the junction. By adding a capacitor C_x in parallel with the junction, I could decrease the size of the junction but still keep the charging energy small (Eq. 5.14) and keep f_{01} (Eq. 5.12) within range of the microwave source [19]. Making the junction smaller made $R(\omega)$ from the junction large because C_J became very small. I kept the critical current >100 nA by increasing the critical current density (Eq. 5.9) which depended on the oxidization parameters during the fabrication (see Fig. 6.6). This kept the Josephson energy (Eq. 5.13) relatively large compared to the charging energy, and it helped to keep f_{01} in the right range of frequencies. The net result of adding C_x was to decrease the impact of dissipation from the junction dielectric, while keeping all of the parameters in the right ranges.

From Eq. 5.6, the loss contributed by the capacitor C_x was:

$$R_x = \left(\frac{1}{\omega_{01} C_x \tan \delta_x} \right) \quad (5.15)$$

where C_x is the capacitance of the extra capacitor and $\tan \delta_x$ is the loss tangent associated with the dielectric in C_x . It is worth emphasizing that the extra capacitor contributed dissipation to the qubit and it is not immediately clear how this would result in lower dissipation overall. Never the less, since C_J and C_x were in parallel across the junction, their resistances added in parallel to produce an equivalent shunting resistance of:

$$\frac{1}{R_{C_{tot}}} = \left(\frac{1}{\omega_{01} C_J \tan \delta_J} \right)^{-1} + \left(\frac{1}{\omega_{01} C_x \tan \delta_x} \right)^{-1} = \omega_{01} C_J \tan \delta_J + \omega_{01} C_x \tan \delta_x \quad (5.16)$$

Similarly, the total capacitance across the junction was $C_J + C_x$. If the only loss was from C_x and C_J then the relaxation time of the qubit would be:

$$T_1 = R_{C_{tot}} (C_J + C_x) \quad (5.17)$$

Here there are two competing effects on the relaxation. According to 5.17, increasing C_x makes T_1 longer. At the same time, according to 5.16 increasing C_x makes $R_{C_{tot}}$ lower, hence decreasing T_1 . However, if $C_x \tan \delta_x \ll C_J \tan \delta_J$, then it is possible to increase C_x without significantly decreasing $R_{C_{tot}}$, and one could end up with a longer T_1 . Therefore, it was critical for the loss tangent of the dielectric in the new capacitor to be small.

There were a couple of options for meeting this goal. One was to make C_x out of a parallel plate capacitor with a low-loss SiN_x dielectric. Hanhee Paik spent a good amount of time optimizing the SiN_x process and was able to create films with a loss tangent of about $3 \cdot 10^{-5}$ at the single photon limit [56], which was an order of magnitude better than previously reported in the literature. Note that when choosing a material it is important to know the loss tangent at low temperature (20 mK) and in the low power, or “single-photon”, limit. A loss tangent quoted from a high power experiment could be artificially low because of saturation of the two level systems in the dielectric [56,58]. From Eqs. 5.15 and 5.16, one can show that this level of loss tangent would have predicted $T_1 \leq 1.6 \mu\text{s}$ if $f_{01} = 3 \text{ GHz}$.

A second option was to use an interdigitated capacitor or IDC (see Fig. 5.3). An IDC can be made by patterning a single layer of metal into interwoven fingers. The capacitance per unit length of each finger is the result of the fringing electric fields




Fig. 5.3 Illustration of an interdigitated capacitor. The cross section shows the fringing \mathbf{E} fields between the fingers.

between the fingers [60,61]. Moe Khalil in Kevin Osborn's group at the Laboratory for Physical Sciences observed a loss tangent of $5 \cdot 10^{-6}$ [56] in the low-temperature single-photon limit in quasi-lumped element resonators that employed interdigitated capacitors that were built on sapphire substrates. Since this was a low loss tangent compared to SiN_x and IDC's were easy to build, I chose to use an interdigitated capacitors for the added capacitor C_x across the junction.

I modeled the interdigitated capacitor using results from a study on IDCs performed by Gary Alley [60]. According to his results, an IDC has a capacitance of:

$$C_{IDC} = (\epsilon_{sub} + 1)[A_1(N - 3) + A_2]\ell \quad (5.18)$$

where $\epsilon_{sub} = 10.7$ is the dielectric constant of sapphire, (this is an average of in-plane and out-of-plane dielectric constants of single crystal sapphire [62]), N is the number of fingers, ℓ is the length of the overlap region between the fingers, and A_1 and A_2 are experimentally determined constants with units of pF/m that depend on the ratio of the thickness of the substrate to the width of the fingers. For the IDCs I built, my substrate was very thick compared to the gap between each finger, and in this limit $A_1 = 4.33$ pF/m and $A_2 = 10.0$ pF/m [60]. Note that this equation is most accurate if the distance between the fingers is equal to the finger width, and if the entire IDC structure is square [60]. For my design $N = 100$, $\epsilon_{sub} = 10.7$ and $\ell = 280 \mu\text{m}$, and Eq. 5.17 gives $C_x = 1.4$ pF.

I found out after measuring the qubit that Micah Stoudimore and Moe Khalil at LPS use a different equation [from Gupta *et al.* [61]] to estimate the capacitance of their IDC's. The formula they employed for the capacitance of an IDC is [61]:

$$C_{IDC} = \frac{\epsilon_{eff} * 10^{-3}}{18\pi} \frac{K(k)}{K'(k)} (N - 1)\ell \quad (\text{in pF}) \quad (5.19)$$

where $\varepsilon_{eff} = (1/2)(1 + \varepsilon_{sub})$, $K(k)$ is the complete elliptical integral of the first kind and K' is the complement to K where $K'(k) = K(k') = K(\sqrt{1 - k^2})$, and k depends on the finger width w and the space s between the fingers such that:

$$k = \left[\tan\left(\frac{w\pi}{4(w+s)}\right) \right]^2. \quad (5.20)$$

Micah and Moe supplied me with the Matlab code [see Appendix A] they wrote for doing the calculation and it also gave me a capacitance of 1.4 pF for the IDC on my chip. Thus, the two equations gave similar results, and these were close to the 1.53 pF value I eventually determined experimentally (see Chapter 8 and Table 8.1).

5.1.3 Loss From Other Capacitors in the Device

According to Eq. 5.6 the loss from a capacitor can be reduced by making the capacitance smaller (this makes $R(\omega)$ large). However, some capacitors in the device needed to be large to operate properly. In particular, the filter capacitor C_f , the extra capacitance across the qubit C_x , and even the capacitance of the isolation junction, C_{J2} all needed to be relatively large. For example C_{J2} would tend to be large because the isolation junction would need to have a large critical current—and hence a large area—to keep its Josephson inductance small. I needed to be sure that these capacitors would not be major sources of dissipation for the qubit.

Now, consider the dissipation from the isolation junction capacitance C_{J2} (see Fig. 5.4). Since this was a tunnel junction, I expected that the aluminum oxide in J_2 was lossy ($\tan\delta_{J2} \sim 10^{-3}$). Fortunately, this lossy element is not connected directly across the qubit

junction. From the qubit junction's point of view, this loss would naively produce an effective shunting resistance across the qubit of:

$$R_{J_2} = \left(\frac{1}{\omega_{01} C_{J_2} \tan \delta_{J_2}} \right) \left(\frac{L_1 + L_2 + L_{J_2}}{L_2} \right)^2 \quad (5.21)$$

because it is inductively isolated from the qubit by the large loop inductance L_I [49].

Another source of dissipation was the 80 pF SiN_x capacitor C_f in the LC-filter on the current bias lines (see Fig. 5.4). Fortunately, C_f was isolated from the qubit by the filter inductor L_f which was around 10 nH. I estimated the effective shunting resistance from the filter capacitor as [49]:

$$R_{C_f} \approx \left(\frac{1}{\omega_{01} C_f \tan \delta_f} \right) \left(\frac{L_f + L_1}{L_p} \right)^2 \quad (5.22)$$

where L_p is the parasitic inductance of the path to ground from the filter capacitor. I used low-loss SiN_x so $\tan \delta_f \approx 3 \cdot 10^{-5}$. I also took $L_p = 0.1$ nH, $L_I = 1$ nH, $L_f = 10$ nH, $C_f = 100$ pF, and $f_{01} = 5$ GHz, and found $R_{J_2} \approx 100$ M Ω . If this was the only source of loss then one would find a relaxation time $T_1 \approx 150$ μ s.

5.1.4 Isolating the Qubit from the Bias Lines

The flux bias line, the current bias line, and the microwave line each presented a complex impedance to the qubit junction. I found the contribution from each line by transforming the impedance of each line into an effective admittance across the qubit junction [38,52]. The real part of each lines equivalent admittance determined its contribution to loss, as described in Eq. 5.1.

Consider first the dissipation caused by the microwave line (see Fig. 5.4). The microwave line was coupled to the qubit through a small capacitor C_μ . The line was a stainless steel coax with $Z_0 = 50 \Omega$, so the impedance of the line looking back from the qubit was:

$$Z_\mu(\omega_{01}) = \frac{1}{i\omega_{01}C_\mu} + Z_0. \quad (5.23)$$

The admittance of the line is found by taking $1/Z_\mu$, and 1 over the real part of this admittance is just the effective shunting resistance of the line. One finds:

$$R_\mu(\omega_{01}) = \frac{1}{\text{Re}[1/Z_\mu]} = Z_0 + \frac{1}{Z_0(\omega_{01}C_\mu)^2} \quad (5.24)$$

According to Eq. 5.24 the loss from the microwave line is minimized if:

$$\frac{1}{Z_0(\omega_{01}C_\mu)^2} \gg Z_0 \quad (5.25)$$

This is equivalent to:

$$C_\mu \ll 1/Z_0\omega_{01} \approx 0.6 \text{ pF} \quad (5.26)$$

where I have taken $\omega_{01} = 2\pi \times 5 \text{ GHz}$ and $Z_0 = 50 \Omega$. In practice, I took $C_\mu \sim 1 \text{ fF}$ and in this limit $R_\mu \sim 1/(Z_0\omega_{01}^2C_\mu^2)$.

Calculating the loss from the current bias line was more complicated, but the procedure was essentially the same. The impedance looking back from the qubit was:

$$Z_b = i\omega L_1 + \left[\frac{1}{i\omega_{01}L_f} + \frac{1}{i\omega_{01}L_2 + (i\omega_{01}C_{J2} + 1/i\omega_{01}L_{J2})^{-1}} \right]^{-1} + \left[\frac{1}{Z_0} + i\omega_{01}C_f \right]^{-1} \quad (5.27)$$

where L_f and C_f are the inductance and the capacitance of the filter elements, C_{J2} and L_{J2} are the geometrical capacitance and Josephson inductance of the isolation junction, L_1 is

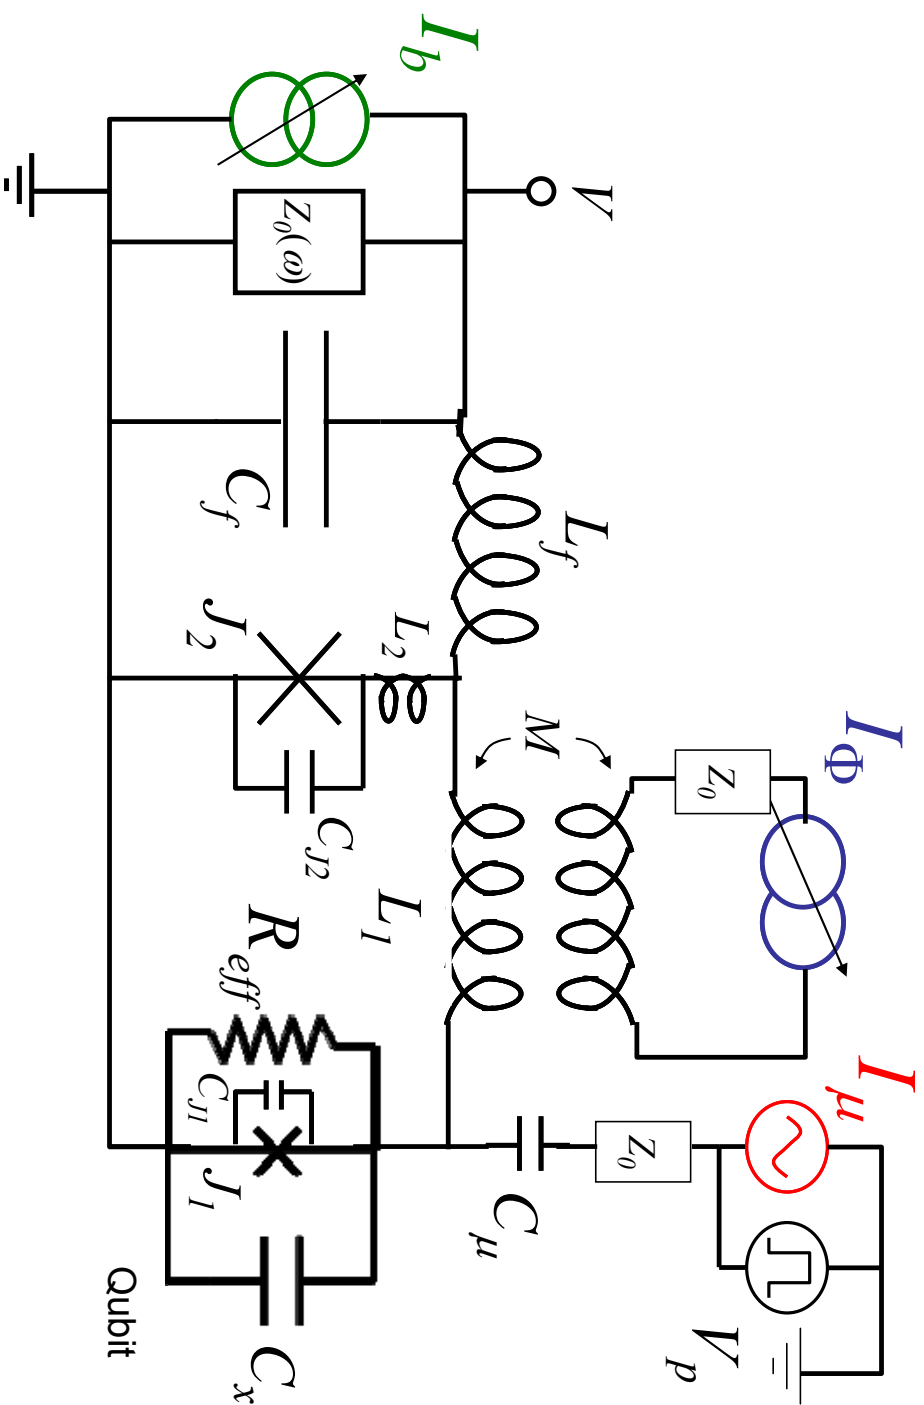


Fig. 5.4 Schematic of dc SQUID phase qubit.

the inductance of the qubit branch of the SQUID loop and L_2 is the inductance of the isolation branch. The algebra involved in finding R_b is messy so it is not detailed here. For $\omega_{01} \gg \omega_f$, where $\omega_f = 1/(L_f C_f)^{1/2}$ is the resonance frequency of the filter, the equation simplifies to:

$$R_b \approx Z_0 \omega_{01}^4 (L_f C_f)^2 \left(\frac{L_1 + L_2 + L_{J2}}{L_2 + L_{J2}} \right)^2. \quad (5.28)$$

According to Eq. 5.28, the isolation of the qubit from the current bias line was mostly accomplished by the LC-filter. For example, a low-pass LC-filter with a cutoff of 200 MHz in front of a qubit with $f_{01} = 5$ GHz would yield $R_b \approx 20$ M Ω . Accordingly, I chose the cutoff frequency of the LC-filter to be about 200 MHz which was much lower than $f_{01} \approx 3$ GHz. The inductive isolation also provided increased protection against dissipation from the bias lines through the second factor in Eq. 5.28.

The impedance of the flux line looking back from the qubit was [51]:

$$Z_\Phi(\omega_{01}) = i\omega_{01}L_1 + i\omega_{01}L_2 + \frac{1}{i\omega_{01}C_{J2} + 1/i\omega_{01}L_{J2}} + \frac{\omega_{01}^2 M^2}{Z_0} \quad (5.29)$$

where M is the mutual inductance between the flux line and the SQUID. Taking the inverse of the real part of the admittance gives:

$$R_\Phi(\omega_{01}) = \frac{1}{\text{Re}[1/Z_\Phi]} \approx Z_0 \left(\frac{L_{J2} + L_1 + L_2}{M} \right)^2. \quad (5.30)$$

The approximation above assumes that $M^2 \omega_{01}^2 / Z_0 \ll 50 \Omega$ and that the plasma frequency of the isolation junction is much higher than ω_{01} . I wanted R_Φ to be large, so in designing the flux line I made the mutual inductance small and chose the inductance on the qubit branch of the SQUID loop to be large.

From the above discussion, one can see that reducing dissipation from the bias leads required choosing the right circuit parameters. To summarize the situation, in order to isolate the qubit from the microwave line I needed to keep C_μ small compared to $1/Z_0\omega_{01}$. The cutoff frequency ω_f of the LC filter had to be much lower than ω_{01} , and the inductance of the isolation branch (L_2+L_{J2}) had to be small compared to the loop inductance L_I in order to protect the qubit from dissipation from the current bias line. Finally, the flux line needed to be weakly coupled through a mutual inductance M that was small compared to the inductances of the SQUID loop. A SQUID phase qubit with such parameters should be well protected against dissipation caused by its bias leads.

5.2 Designing to Reduce Dephasing

Dephasing of the qubit state occurs because of slow fluctuations in the energy level spacing between $|0\rangle$ and $|1\rangle$ [43,54]. The fluctuations can be caused by thermally induced noise currents on the bias lines, fluctuations in the magnetic flux in the SQUID, or fluctuations of the parameters of the device itself (*e.g.* critical current fluctuations) [43,44,54]. As with relaxation, the design of the device and the choice of parameters affect the dephasing rate.

Each bias line contributes to dephasing through thermally induced noise currents that causes decoherence of the qubit. According to Martinis and Urbina *et al.*, thermal current noise in a phase qubit gives a phase decay rate $1/T_\phi$ of (see Sect. 3.3 and [43]):

$$\frac{1}{T_\phi} \approx \frac{k_B T}{R_i} \left(2\pi \frac{\partial f_{01}}{\partial I_1} \right)^2. \quad (5.31)$$

where k_B is the Boltzmann constant, $T \approx 20$ mK is the temperature, R_i is the effective resistance of a given line (similar to that calculated above but generally at low frequency), and $\frac{\partial f_{01}}{\partial I_1}$ is the noise transfer function for current. The noise transfer function quantifies the effect that a fluctuation in the current I_{n1} through a junction has on the qubit frequency f_{01} . Fig. 5.5b shows $\partial f_{01}/\partial I_1$ vs. I_1 for a single junction with $I_{01} = 150$ nA and $C_x = 1.25$ pF assuming that $f_{01} \approx f_p$, the plasma frequency, where $f_p = 1/[2\pi\sqrt{L_{J1}(C_{J1} + C_x)}]$ and

$$L_{J1} = \frac{\Phi_0}{2\pi} \frac{1}{\sqrt{I_{01}^2 - I_1^2}}. \quad (5.32)$$

From Fig. 5.4(b) we see that $\partial f_{01}/\partial I_1 \sim 10^{16}$ Hz/A for this choice of parameters reasonably high bias currents.

As I explained in Sect. 3.3, dephasing is caused by low frequency fluctuations in f_{01} . Therefore, I need to consider the effective resistances of each bias line in the low frequency limit. Let me first consider how small variations in the current on the flux line translated to variations in the current in the SQUID loop. The SQUID loop has two junctions, J_1 and J_2 , and a geometrical inductance on each arm, L_1 and L_2 . The flux line couples to the loop through a mutual inductance, M . Assuming that $I_b = 0$, the flux-phase relationship [50] for the loop is:

$$\gamma_1 - \gamma_2 = \left(\frac{2\pi}{\Phi_0} \right) (MI_\Phi + JL_1 + JL_2) \quad (5.33)$$

where γ_1 and γ_2 are the phase difference across J_1 and J_2 , respectively, J is the circulating current, and I_Φ is the current flowing in the flux line. Also,

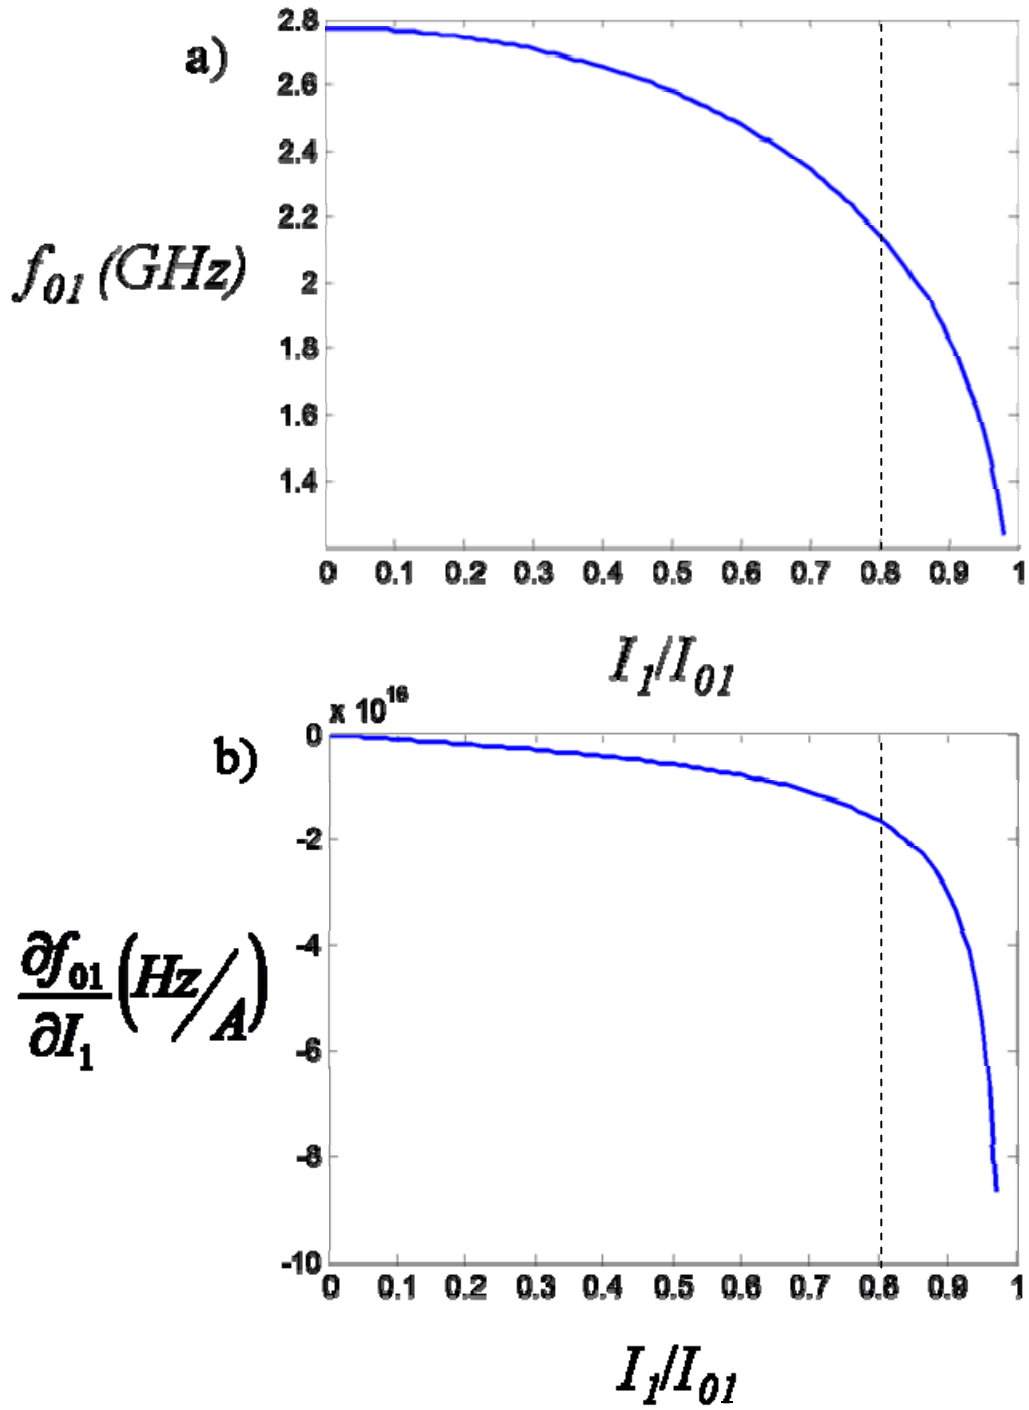


Fig. 5.5. (a) Resonance frequency of a single junction vs. applied current. (b) The current noise transfer function vs. applied current. The dotted line shows a typical bias current for operating the qubit. For this plot I used $I_{01} = 150$ nA and $C_x = 1.25$ pF.

for $I_b = 0$ and assuming that flux into the page is positive, the current in each junction is related to the circulating current by:

$$J = -I_{01} \sin \gamma_1 = I_{02} \sin \gamma_2 \quad (5.34)$$

where I_{01} and I_{02} are the critical currents of J_1 and J_2 .

Consider the case where a small fluctuation in I_Φ causes fluctuations in J , γ_1 , and γ_2 . In this case Eqs. 5.33 and 5.34 become:

$$\gamma_1 + \delta\gamma_1 - \gamma_2 - \delta\gamma_2 = \left(\frac{2\pi}{\Phi_0} \right) (MI_\Phi + M\delta I_\Phi + JL_1 + \delta JL_1 + JL_2 + \delta JL_2), \quad (5.35)$$

$$J + \delta J = -I_{01} \sin(\gamma_1 + \delta\gamma_1) = I_{02} \sin(\gamma_2 + \delta\gamma_2). \quad (5.36)$$

Using the identity $\sin(\gamma_i + \delta\gamma_i) = \sin \gamma_i \cos \delta\gamma_i + \sin \delta\gamma_i \cos \gamma_i$ and expanding for small $\delta\gamma_i$,

Eq. 5.36 can be rewritten:

$$J + \delta J \approx -I_{01} \sin \gamma_1 - I_{01} \delta\gamma_1 \cos \gamma_1 = I_{02} \sin \gamma_2 + I_{02} \delta\gamma_2 \cos \gamma_2. \quad (5.37)$$

Here I have only retained terms of first order in $\delta\gamma_2$ and $\delta\gamma_1$. Next, using the equation for the Josephson inductance (Eq. 5.10) and subtracting Eq. 5.34 from Eq. 5.37 yields:

$$\delta J = -\frac{\Phi_0}{2\pi} \frac{1}{L_{J1}} \delta\gamma_1 = \frac{\Phi_0}{2\pi} \frac{1}{L_{J2}} \delta\gamma_2. \quad (5.38)$$

Substituting the definitions for $\delta\gamma_1$ and $\delta\gamma_2$ from Eq. 5.38 into Eq. 5.35 and subtracting away Eq. 5.33 gives:

$$-L_{J1}\delta J - L_{J2}\delta J = M\delta I_\Phi + L_1\delta J + L_2\delta J \quad (5.39)$$

Therefore, the ratio of fluctuations in I_Φ to those in J is:

$$\frac{\delta I_\Phi}{\delta J} = -\frac{L_{J1} + L_{J2} + L_1 + L_2}{M} \quad (5.40)$$

where the minus sign in Eq. 5.40 comes from the way I defined the mutual inductive coupling M of I_ϕ to the loop.

This ratio of current fluctuations can be used to obtain a “stepped up” impedance for the flux line [49]. If the flux line has a characteristic impedance Z_0 , then that impedance is stepped up due to the isolation offered by the mutual inductance such that:

$$Z_\Phi = Z_0 \left(\frac{L_{J1} + L_{J2} + L_1 + L_2}{M} \right)^2 \quad (5.41)$$

The basic idea is that if there is some noise power coming down the flux line, then the noise power that the qubit sees is decreased by a factor of $[M/(L_{J1}+L_{J2}+L_1+L_2)]^2$. Note that this form of Z_Φ is only valid in the low frequency limit, hence it should not be used for calculating T_I . Also, for low enough frequencies, Z_0 is not necessarily $\sim 50 \Omega$, but it is the value of the bias resistor on that particular line (10 k Ω for the flux line). I wanted Z_Φ to be large, so in designing the flux line I kept the mutual inductance small. A large loop inductance also helped to protect against dephasing.

The other low-frequency resistances can be found in a similar fashion. I obtained:

$$R_b \approx Z_0 \left(\frac{L_1 + L_{J1} + L_2 + L_{J2}}{L_2 + L_{J2}} \right)^2, \quad (5.42)$$

$$R_\mu(\omega) = Z_0 + \frac{1}{Z_0(\omega C_\mu)^2}. \quad (5.43)$$

Again, note that these stepped up resistances incorporate the Josephson inductance of the qubit junction L_{J1} are only valid in the low frequency limit $\omega \ll \omega_{0I}$. They should not be used for estimating T_I . Also, here I have assumed that the line impedance was Z_0 , but at low enough frequencies, the line impedance will be set by the bias resistors, which were 200 k Ω for I_b and 50 Ω for I_μ .

The other two sources of dephasing, flux noise and critical current noise, have a $1/f$ noise power spectrum [54]. Martinis *et al.* wrote that the mean-squared phase noise is:

$$\langle \phi_n^2(t) \rangle \approx \left(\frac{\partial \omega_{01}}{\partial I_1} \right)^2 S_I(1\text{Hz}) \ln \left(\frac{0.401}{f_r t} \right) t^2 \quad (5.44)$$

where f_r is the repetition rate of the experiment (700 Hz). For $1/f$ flux noise, we can write:

$$\langle \phi_n^2(t) \rangle \approx \left(\frac{\partial \omega_{01}}{\partial I_1} \right)^2 \frac{S_\Phi(1\text{Hz})}{(L_1 + L_{J1} + L_2 + L_{J2})^2} \ln \left(\frac{0.401}{f_r t} \right) t^2. \quad (5.45)$$

$S_\Phi(1\text{Hz})$ has been found experimentally to be about $(10 \mu\Phi_0)^2/\text{Hz}$ for a SQUID of the size I used [63]. Assuming that $I_{01} = 150 \text{ nA}$ and $I_{02} = 10 \mu\text{A}$, the total loop inductance, including the Josephson inductances, is about 3 nH; therefore, the current noise power generated by $(10 \mu\Phi_0)^2/\text{Hz}$ of flux noise power is about $48 (\text{pA})^2/\text{Hz}$ at 1 Hz.

It has been found that critical current noise typically also has a $1/f$ power spectrum, and, very roughly, one should expect [54]:

$$S_{I_{01}}(1\text{Hz}) \approx 144 \frac{(I_{01}/\mu\text{A})^2}{A/(\mu\text{m})^2} \left(\frac{T}{4.2 \text{ K}} \right)^2 (\text{pA})^2/\text{Hz} \quad (5.46)$$

where A is the area of the junction and T is the temperature. When using equation Eq. 5.46, I_{01} should be input in μA and the area in $(\mu\text{m})^2$. The answer is given in $(\text{pA})^2/\text{Hz}$. For example, a junction with $I_{01} = 150 \text{ nA}$, $A = 0.1 \mu\text{m}^2$, at $T = 20 \text{ mK}$ has approximately $800 (\text{fA})^2/\text{Hz}$ of critical current noise power at 1 Hz. Thus, critical current noise is fairly insignificant compared to flux noise.

Because Eq. 5.44 predicts a t^2 dependence for the phase noise, it would be incorrect to quote an exponential decay rate for dephasing due to low frequency noise. This time dependence to the phase noise causes a Gaussian envelope in a Ramsey fringe experiment and we can define a corresponding time constant for the loss in phase coherence from $1/f$ flux and critical current noise as:

$$\frac{1}{T_{R,\Phi}} \sim \left[\left(2\pi \frac{df_{01}}{dI_1} \right)^2 \frac{S_\Phi(1\text{Hz})}{(L_1 + L_{J1} + L_2 + L_{J2})^2} \right]^{1/2} \quad (5.47)$$

and

$$\frac{1}{T_{R,I_{01}}} \sim \left[\left(2\pi \frac{df_{01}}{dI_1} \right)^2 S_{I_{01}}(1\text{Hz}) \right]^{1/2} \quad (5.48)$$

respectively.

5.3 Qubit Design Parameters and Expected Coherence Time

I chose the design values for the device parameters so that the qubit coherence time would be very long. Table 5.1 shows the resulting design parameters I selected and Table 5.2(a) shows how much each source contributes to the qubit's loss. The results for the dephasing analysis are displayed in Table 5.2(b) and 5.2(c).

For these estimates, I assumed that the qubit would be biased at $0.7I_{01}$, which made $L_{J1} = 3.1$ nH, $f_{01} = 2.4$ GHz, and $\partial f_{01}/\partial I_1 \approx 1 \cdot 10^{16}$ Hz/A. In this section I describe how each circuit element is designed in order to achieve the desired value for the parameters in Table 5.1.

In order to limit dissipation from the bias leads, I chose the filter capacitor C_f to be 80 pF and the filter inductor L_f to be 10 nH. This put the cutoff frequency f_c at about

Table 5.1. Qubit design parameters for device PB9

Parameters	Design Value
I_{01} (nA)	150
I_{02} (μ A)	10
L_{J1} (nH)	2.2
L_{J2} (nH)	0.03
L_1 (nH)	1.25
L_2 (nH)	0.02
A_{J1} (μm^2)	0.1
A_{J2} (μm^2)	7
C_{J1} (fF)	2.5
C_{J2} (pF)	0.18
C_x (pF)	1.5
C_f (pF)	80
L_f (nH)	10
L_p (nH)	0.1
M (pH)	1
C_μ (fF)	1
f_{01} (GHz)	2.4
I_b (nA)	105
Biased L_{J1} (nH)	3.1
df_{01}/dI_1 (Hz/A)	$\sim 10^{16}$

Table 5.2. Results of decoherence analysis using the design parameters in Table 5.1.

(a) Relaxation

Dissipation Source	Symbol	R_i (M Ω)	$T_{1,i}$ (μ s)
Bias Leads	I_b	980	1,500
Flux Bias	I_ϕ	85	130
μ -wave line	I_μ	89	130
Qubit JJ capacitance	C_{J1}	27	40
Interdig. capacitance	C_x	8.9	13
Isolation JJ capacitance	C_{J2}	1,600	2,400
Filter capacitance	C_f	280	430
Expected T_1 (μ s)			≤ 8.4

(b) Dephasing

Thermal Dephasing Source	T_ϕ (μ s)
Bias Leads	270
Flux Bias	770,000
m-wave line	>1,000,000

(c) Inhomogeneous Broadening

1/f Dephasing Source	T_R (μ s)
Flux Noise	~ 1
Critical Current Noise	~ 100

180 MHz, and made the effective resistance associated with the loss from the bias leads about 980 M Ω . Keep in mind that since $C_x = 1.5$ pF for this design (see Table 5.1), according to Eq. 5.2, a 980 M Ω effective resistance translates to $T_I = 1.5$ ms. I isolated the qubit from the flux bias line by choosing the mutual inductance to be $M = 1$ pH and the SQUID loop inductance $L_I = 1.25$ nH. With these parameters, the loss from the flux bias line gave an effective resistance of 85 M Ω , which gives a T_I of about 130 μ s. The microwave line was weakly coupled to the qubit through $C_\mu = 1$ fF. According to Eq. 5.22, with this choice for C_μ the effective resistance of the microwave line was about 89 M Ω , which would yield a T_I of about 130 μ s.

I selected the qubit junction to have an area $A_{JI} = 0.1$ (μ m)², which made the capacitance $C_{JI} = 2.5$ fF according to Eq. 5.9. Since C_{JI} was so small, even though $\tan\delta_f \sim 10^{-3}$ the effective resistance was still about 27 M Ω , which gives $T_I \approx 40$ μ s. I set the isolation junction capacitance to $C_{J2} = 180$ fF, and the effective resistance was about 1.6 G Ω , which gives $T_I \approx 2.4$ ms. According to Eq. 5.20, the effective resistance associated with loss from the SiN_x in C_f was about 280 M Ω , which gives $T_I \approx 430$ μ s, and Eq. 5.20 does not even account for the inductive isolation contributed by the SQUID asymmetry, so the effective resistance was larger. With this design, the interdigitated capacitor contributed the most loss. Although $\tan\delta_x \approx 5 \cdot 10^{-6}$, the capacitance $C_x = 1.5$ pF was relatively large and, according to Eq. 5.14, the effective resistance was about 8.9 M Ω , which gives $T_I \approx 13$ μ s.

Combining all the effective resistances in parallel gave a total shorting resistance for the design of $R_{eff} = 5.6 \text{ M}\Omega$. Using Eqs. 5.1 and 5.2, the estimated relaxation time T_1 for the device was $8.4 \text{ }\mu\text{s}$.

Dephasing from thermal noise on the bias lines should have been negligible. The microwave line was coupled to the qubit by a small capacitor C_μ which should have looked like an open connection to low frequency noise. Accordingly, it should have contributed minimal dephasing. The mutual inductance was so small that the qubit should have been well isolated from low frequency noise on the flux bias line. According to Eqs. 5.31 and 5.41, the dephasing time associated with fluctuations in the flux bias current from thermal noise should have been around 0.77 s . The current bias line also contributed little dephasing. Since the inductances of the qubit branch of the SQUID were designed to be large, $L_I = 1.25 \text{ nH}$ and $L_{JI} = 3.1 \text{ nH}$, and assuming that $Z_0 = 50 \text{ }\Omega$ (as a worst case scenario), although the noise transfer function was $\sim 10^{16} \text{ Hz/A}$, the dephasing time came out to be around $270 \text{ }\mu\text{s}$. Considering these contributions, we can see that the thermal noise on the bias lines did not contribute much dephasing to the qubit with this set of design parameters.

The dephasing time associated with the $1/f$ noise sources was more difficult to estimate accurately. Equation 5.47 and 5.48 give estimates for the Ramsey fringe decay time associated with the $1/f$ flux and critical current noise, respectively. It is at best an order of magnitude estimate because it assumes that the logarithmic factor in Eq. 5.44 and 5.45 is of order 1. According to Eq. 5.46, the noise power at 1 Hz of the critical current noise should only be about 800 (fA)^2 so the T_R associated with that noise was $\sim 100 \text{ }\mu\text{s}$, thus it was not a significant source of dephasing for the design. However, this

result was on Nb/NbO_x/PbIn tunnel junction measured under continuous bias [63] and considerable variations have been reported in Al devices [17]. The T_R associated with $1/f$ flux noise was much shorter. Assuming that the noise power spectral density of $1/f$ flux noise was $10 \mu\Phi_0/\sqrt{\text{Hz}}$ at 1 Hz, the noise current in the SQUID would be about 48 (pA)². That made $T_{R,\Phi} \sim 1 \mu\text{s}$ according to Eq. 5.47. Even with the relatively large loop inductance, the $1/f$ flux noise would significantly limit the coherence time of the device. However, 1 μs would be a much longer Ramsey time than any of our previous devices so this was an acceptable set of parameters [18,26,49,52,53,64].

5.4 Device PB9 Layout and Fabrication Considerations

In the last section, I detailed the design parameters and their effects on the predicted relaxation time and dephasing time of the device. This section details the layout of the device and the modeling that went into obtaining a device with the parameters in Table 5.1. The device was drafted in DesignCAD, which is a computer drafting program that is associated with the e-beam lithography control program NPGS. Figure 5.6 shows the full view of the device design and a closer look at the dc SQUID. Each feature was carefully designed to achieve the parameters for the device that would give optimized coherence.

Figure 5.7 shows the e-beam pattern for the qubit junction. Two 150 nm wide lines were written at perpendicular angles to each other. The right line stopped 500 nm from the edge of the left line. When developed, this left a bridge in the PMMA for the double-angle evaporation (see Sect. 6.9). The lines as drafted in DesignCAD were 150 nm wide, but in reality, they would overdevelop to more like 300-350 nm. The

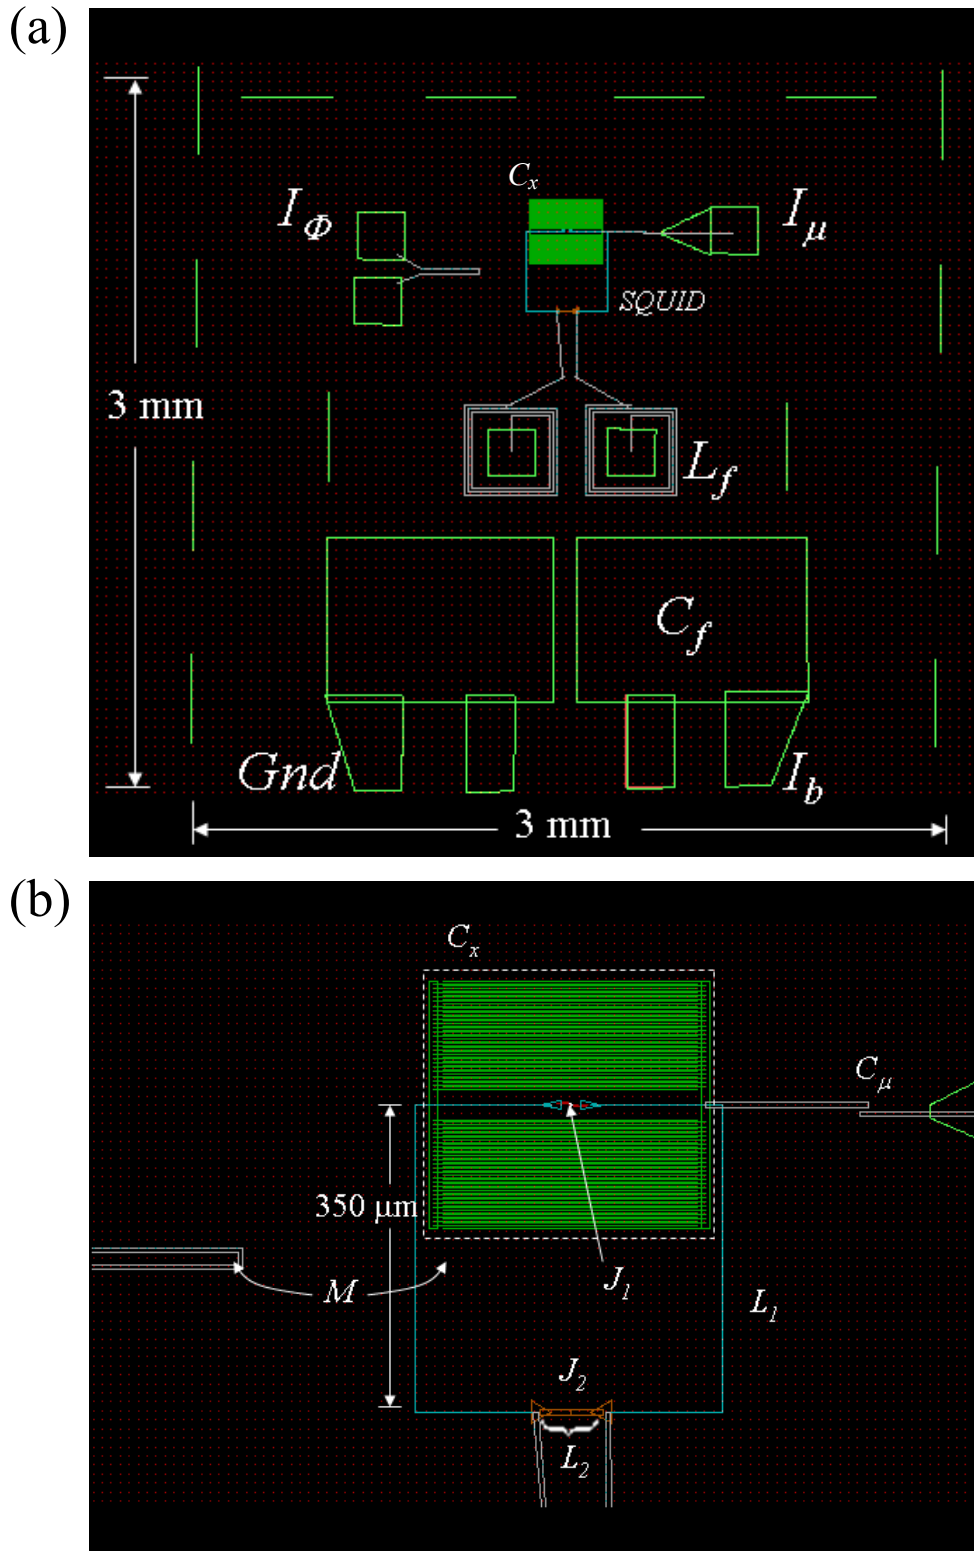


Fig. 5.6 a) E-beam pattern for the entire device. b) Close-up of the design for the dc SQUID.





Fig. 5.7 Close up of the qubit junction's e-beam pattern. The 150 nm lines are expected to overdevelop to about 300 nm, and the 500 nm gap between them should narrow some, but it will still survive to act as the suspended bridge for creating a Josephson junction by double-angle evaporation.

perpendicular arrangement of the two lines assured that the area of the junction would be independent of the angle of the evaporation, as long as the angles were steep enough such that the features overlap completely. The resulting junction area was expected to be about $0.1 (\mu\text{m})^2$.

The extra capacitance C_x shunting the qubit was designed to be 1.5 pF. C_x was an interdigitated capacitor (IDC) because IDC's sapphire can have a loss tangent as low as 5×10^{-6} [56]. Achieving a capacitance of greater than 1 pF with an IDC required a lot of fingers, but the larger the structure was, the more stray inductance it would have. I decided that if I kept the structure within a $300 \mu\text{m} \times 300 \mu\text{m}$ square, then I estimated that it would only have about 150 pH of stray inductance; the fingers themselves represent negligible inductance since they all add in parallel and the major contribution to the inductance would come from the path length out to each finger. In order to create a large capacitance IDC in a small space, I needed many fingers that were very closely spaced. I decided on 100 fingers that were each $1.25 \mu\text{m}$ wide with a spacing of $1.25 \mu\text{m}$ between each of them. The overlap region between the fingers would be $290 \mu\text{m}$ long. According to Eq. 5.16, this structure would have capacitance $C_x \approx 1.5 \text{ pF}$.

Figure 5.8 shows the e-beam pattern for the IDC. There were 50 fingers on either side of the qubit to cut down on the parasitic inductance. The fingers were drafted to be $1 \mu\text{m}$ wide and $1.5 \mu\text{m}$ apart, but overdeveloping would make them $1.25 \mu\text{m}$ wide and $1.25 \mu\text{m}$ apart [Fig. 5.8(c)] in the device. There was $16 \mu\text{m}$ between the line running to the qubit junction in the middle and the fingers on either side [Fig. 5.8(b)]. This helped to protect against any overdeveloping of the qubit, and it would keep the electric fields produced in the IDC from affecting the dynamics of the qubit junction.

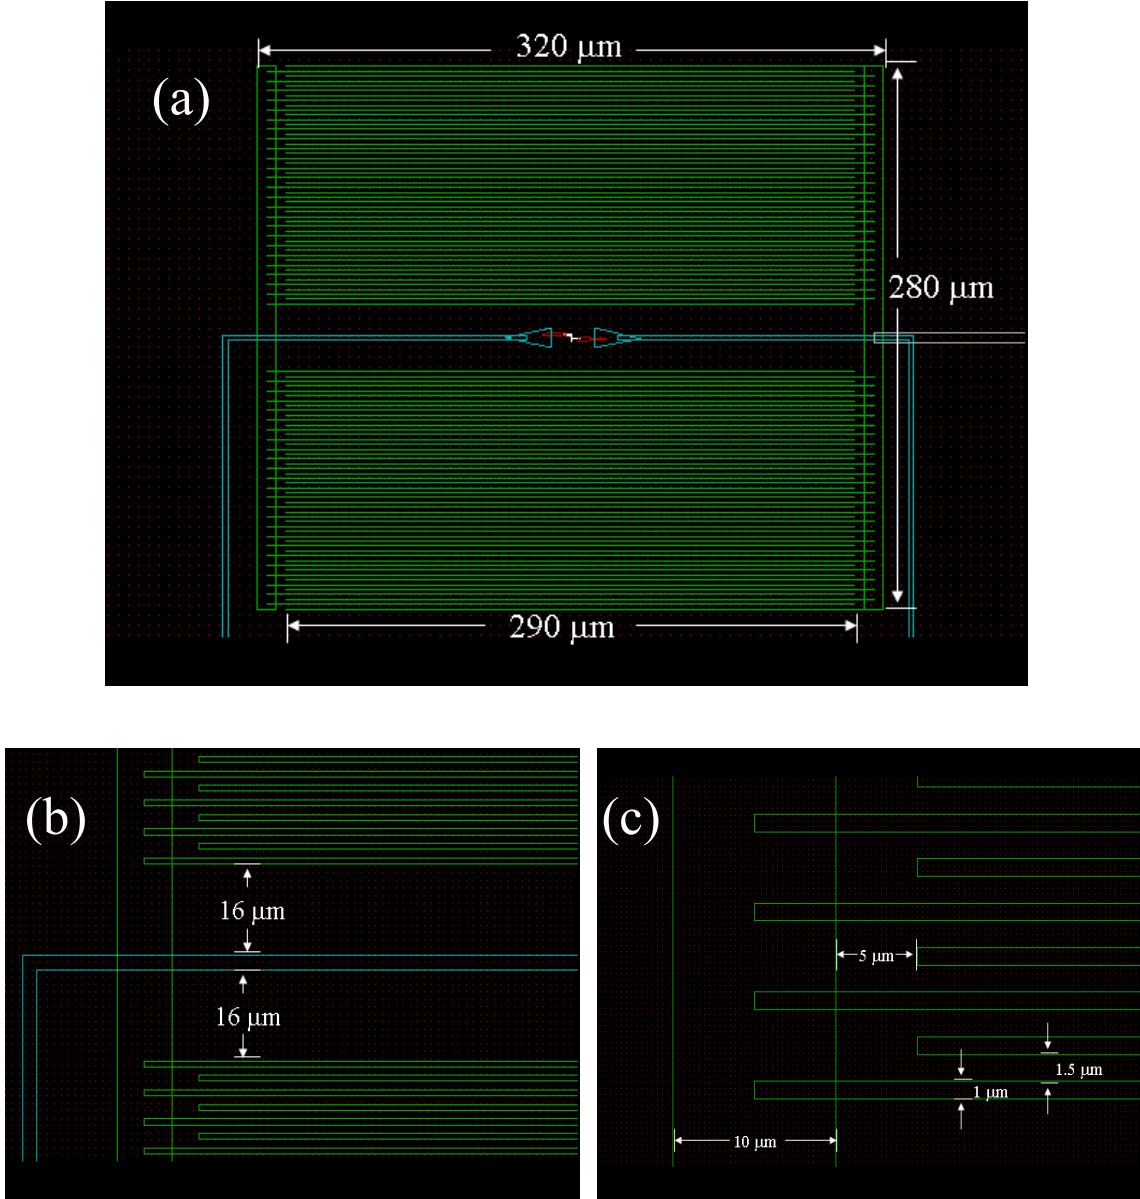


Fig. 5.8 (a) E-beam pattern of the interdigitated capacitor, C_x . (b) Close-up of the gap between the SQUID line and the IDC fingers. (c) Close-up of the fingers.

The isolation junction J_2 was much larger than the qubit junction. Assuming the same J_c , A_{J_2} needed to be about $7 \mu\text{m}^2$ in order for it to have a critical current $I_{02} = 10 \mu\text{A}$. According to Eq. 5.10, the inductance of J_2 would be low (~ 10 s of pH), which helps to isolate the qubit from noise on the current bias leads (Eqs. 5.28 and 5.42). In designing the e-beam pattern for the isolation junction, I was not too concerned about precisely hitting the parameters set out in the above paragraph. Overdeveloping was a huge problem for this feature, and when it did not overdevelop, then often there was not enough undercut to clear out the MMA under the bridge. The design that worked in the end had $5 \mu\text{m}$ wide lines that terminated end-to-end, $0.9 \mu\text{m}$ apart [see Fig. 5.9(b)]. The idea was to have a thick enough MMA resist and steep enough angles so that the isolation junction would be $5\text{-}10 (\mu\text{m})^2$.

The SQUID loop was designed to be a $350 \mu\text{m} \times 350 \mu\text{m}$ square. The overall shape and size were chosen because they were similar to our previous device [64]. Figure 5.6(b) shows the e-beam pattern for the SQUID loop. A square inductor with $350 \mu\text{m}$ sides would have a $1.0\text{-}1.5 \text{ nH}$ inductance if the wire is $2.5 \mu\text{m}$ wide using an internet calculator [65]. A FastHenry simulation of a loop with those dimensions gave an inductance of 1.5 nH [66].

The inductance L_2 of the isolation arm of the SQUID could be kept small in the design by attaching the current bias leads as close to the isolation junction as possible. This was difficult to achieve in reality because bringing the leads in too close caused the bridge for the isolation junction to overdevelop. In the end, I had to pattern the leads to come in about $40 \mu\text{m}$ away from the isolation junction as shown in Fig. 5.9(a). This




Fig. 5.9 E-beam pattern of the isolation branch of the SQUID. (a) shows the dimension of L_2 and (b) shows the pattern for the isolation junction L_{J2} .

probably did not achieve the desired value of $L_{J2} = 20$ pH, but it was a necessary compromise and should still have given $L_2 < 100$ pH.

The flux line was modeled with FastHenry before patterning it. According to FastHenry, the flux line depicted in Fig. 5.10(a) gave a mutual inductance to the SQUID loop of 1 pH. Figure 5.10(b) shows the pattern. The turn in the flux line is 200 μm away from the left side of the SQUID loop and runs parallel to the left side of the loop for only 20 μm . The input and ground of the flux line run parallel to each other, which should have minimized the stray coupling from those segments of the line.

Microwaves were applied to the qubit by means of a bare microstrip line that was coupled to the qubit through a small capacitor C_μ . I formed this coupling capacitor from two lines that ran parallel to each other for 10 μm [see Fig. 5.11(b)]. The lines are 5 μm wide and separated by 5 μm . I can think of this as an interdigitated capacitor, and according to Eq. 5.18, $C_\mu \sim 1$ fF. I did not spend much time engineering this line because I knew that I wanted low coupling and thought that the small IDC would take care of that. However, it is possible that some inductive coupling or radiative coupling to the other bias lines could be present, and, as I will discuss in Chapter 9, this line turned out to be contributing significant loss.

The LC-filter design had one inductor and one capacitor; however, the actual device had two physical elements in series. The inductors were 3-turn spiral inductors with an outer diameter of 400 μm , 5 μm wide lines, and 10 μm of space between each turn (see Fig. 5.12). Each inductor should have an inductance of 5 nH. There was one inductor on the input and one on the output of the SQUID. Since the bias current flowed clockwise in one inductor and counter-clockwise in the other, the mutual inductance of





Fig. 5.10 (a) Layout of a FastHenry model for the flux line. (b) The e-beam pattern for the flux line.

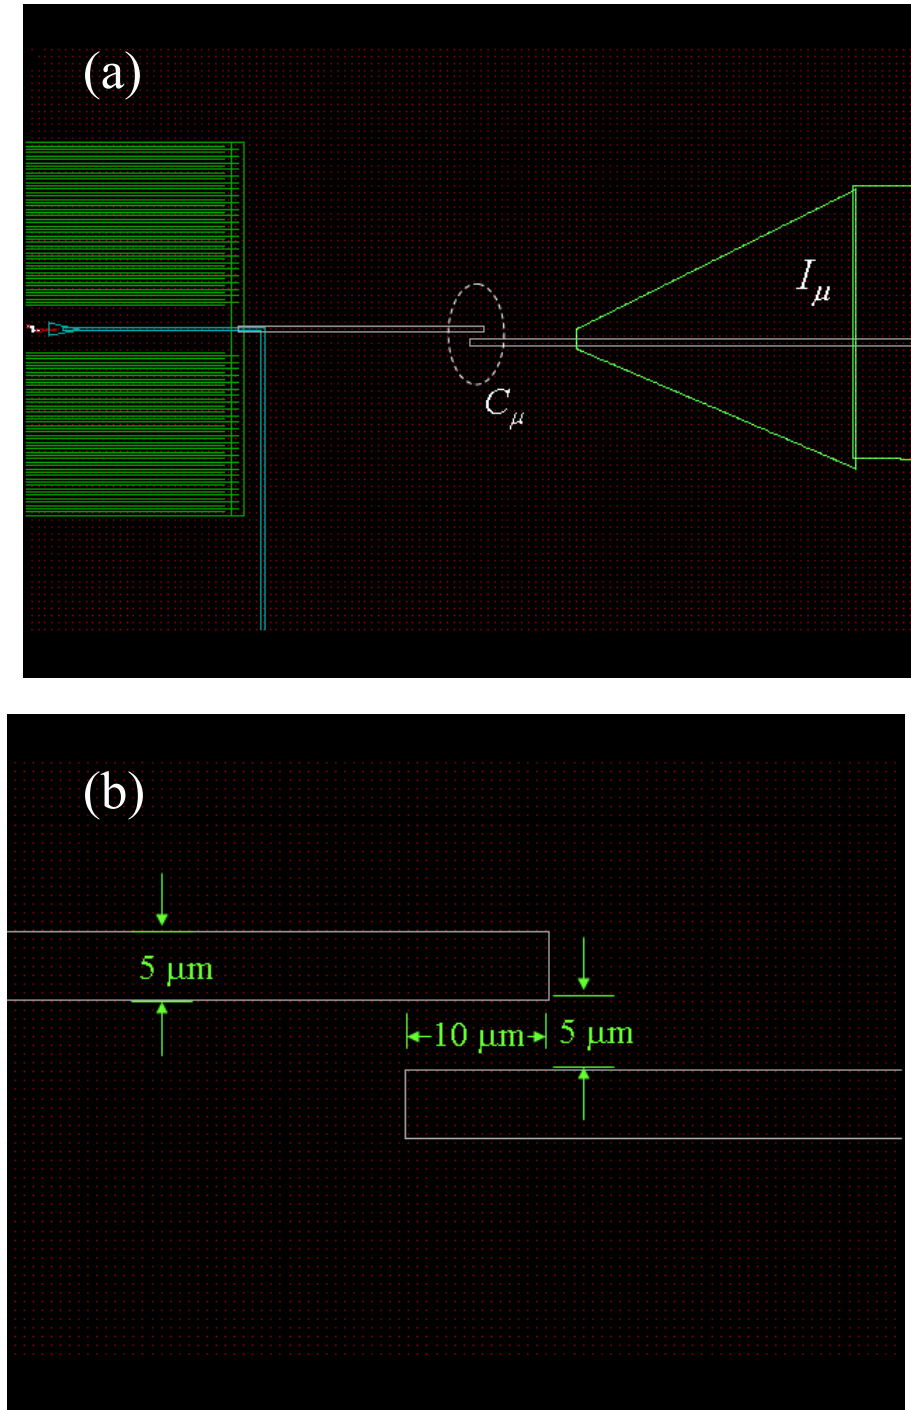


Fig. 5.11 (a) Layout of the microwave line I_μ . (b) Close-up on the coupling capacitor C_μ that is circled in (a).

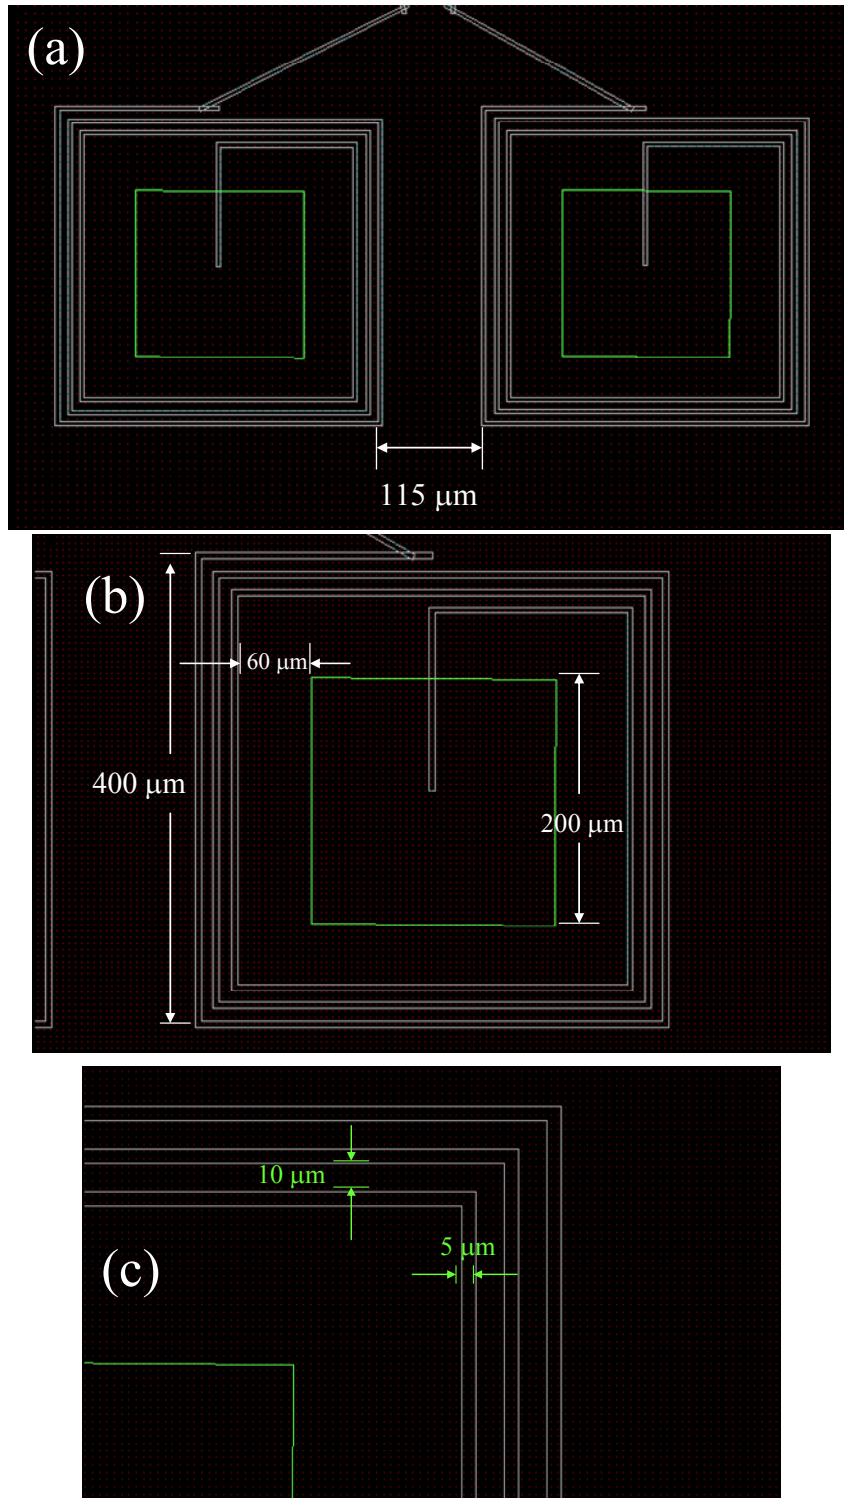


Fig. 5.12 (a) Layout of the filter inductors L_f . (b) Dimensions of the coils in one spiral inductor. (c) Dimensions of the wire widths and gaps between the turns.

the two to the SQUID loop was minimized. The capacitor was made of two top plates over a SiN_x layer with a back plate underneath (see Fig. 5.13). This created two capacitors in series. The SiN_x was meant to be 100 nm thick and the top plates each covered a $500\text{ }\mu\text{m} \times 450\text{ }\mu\text{m}$ area. Each capacitor contributed a capacitance of $C = A\epsilon\epsilon_0/d$ where d is the thickness of the SiN_x , A is area of the top plate coverage, ϵ_0 is the permittivity of free space, and ϵ is the dielectric constant (12.5 for SiN_x); however, with two in series, the capacitance was cut in half. Therefore, the whole structure would have $C_f = 88\text{ fF}$. I intentionally designed the capacitor to be larger than the value in Table 1 because it was common for the nitride to be thicker than expected. This way, if the SiN_x was too thick, then the capacitance will be close to the desired value and if it was 100 nm, the device would end up with a better filter than expected.

5.5 Pictures of Device PB9

Device PB9 was a dc SQUID phase qubit that I fabricated using e-beam lithography and double-angle evaporation of Al on a sapphire substrate. The fabrication procedure is detailed in Chapter 6. Figure 5.14 shows a picture taken with an optical microscope of the whole device, and Fig. 5.15 shows a view of just the SQUID. The qubit junction was too small to see with the optical microscope, but there are AFM pictures of both Josephson junctions in Sect. 8.1. Most of the features came out as designed. One small defect in the pattern was that a piece of dirt severed a few of the fingers of the IDC. Figure 5.16 shows a close up of the IDC. There is a dark feature just above and to the right of the qubit junction. This device had a problem during liftoff where a bit of the Al got stuck in between the fingers of the IDC because of a defect in

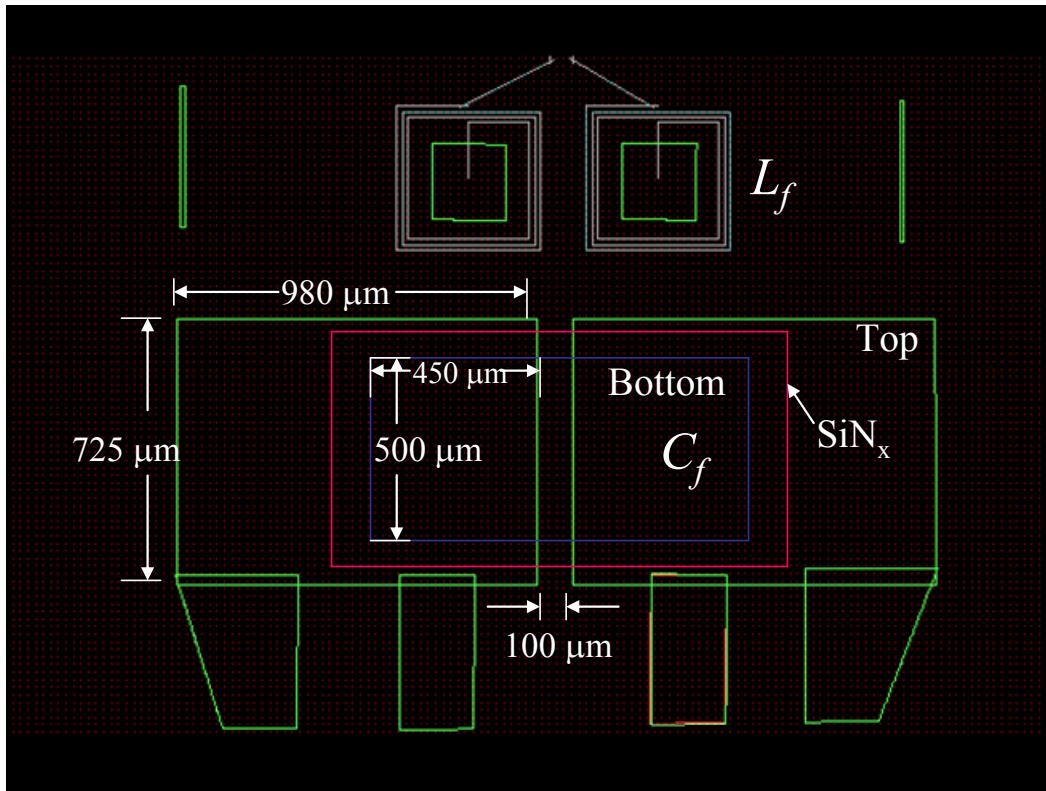


Fig. 5.13 Layout and dimensions of the filter capacitor C_f . The blue rectangle defines the backplate of the capacitor, which was made by photolithography. The red rectangle denotes where the SiN_x will be patterned over the bottom plate. The green regions define the e-beam pattern for making the top plate of C_f . The structure forms two capacitors in series.




Fig. 5.14 Photograph of dc SQUID phase qubit PB9. This picture was formed from two images that were spliced together. The connections between C_f and L_f were made with wirebonds, which are represented schematically in this picture.




Fig. 5.15 Close up photograph of the dc SQUID phase qubit PB9.

the pattern resist. I used an ultrasound acetone bath to remove the aluminum, but it also created a break in 7 fingers. This should have reduced the capacitance by a few percent.

The values of the actual device parameters were determined through measurements of the device, as I discuss in Chapter 8.

5.6 Summary

In designing device PB9, I modeled the dissipation and dephasing from all of the known sources. From the circuit model of the device I found the effective resistances presented to the junction by the bias lines. The dielectric loss from the capacitors on the chip was modeled as frequency dependent parallel resistances that depended on the loss tangent of the dielectric material. This affected the choice of the substrate (sapphire) and the dielectric in the filter capacitor (SiN_x). The dephasing from the bias leads was modeled as coming from current fluctuations caused by thermal noise coming down the lines. Dephasing from $1/f$ flux noise and critical current noise were estimated. According to the model for dissipation, the T_I of the device should have been about $8.4 \mu\text{s}$, with dissipation from the interdigitated capacitor being the dominant source. $1/f$ flux noise was the strongest source of dephasing, giving a time constant for a Ramsey fringe experiment of $T_R \sim 1 \mu\text{s}$.

The device was designed in DesignCAD such that it should have the parameters given in Table 5.1. The pictures of the device showed that the design came out as planned, except for a defect in the IDC that caused a break in 7 of the fingers.




Fig. 5.16 Close up photograph of the interdigitated capacitor C_x . The fingers are difficult to see, even at the maximum magnification of the optical microscope. A defect in the pattern severed 7 of the fingers just above and to the right of the qubit junction J_I .

Chapter 6

Fabrication Procedure

This chapter describes, step-by-step, the procedure I used for fabricating the phase qubit PB9. Building one of these phase qubits is a multi-step process that usually spans several days. Making the large LC-filter required several photolithographic steps, while achieving a sub-micron junction size required using e-beam lithography.

There were five main steps in the fabrication procedure:

- (1) Al deposition and photolithography to form the back plates of the filter capacitors.
- (2) Growth of a layer of low-loss SiN_x dielectric which was then etched in SiN_x .
- (3) Dicing and cleaning of the individual chips.
- (4) e-beam patterning of a multilayer MMA/PMMA resist to form the wires, SQUID loop and suspended bridges for making the Josephson junctions.
- (5) Double-angle evaporation and oxidation of the Al to form the junctions.

In the following sections, I describe these main steps and several minor steps in detail.

6.1 Preliminary Fabrication Steps and SiN_x Growth at LPS

To build devices, I started with a 2" sapphire wafer that I purchased from University Wafer. The wafers were 400 μm thick and cut on the C-plane (0001). I placed a blank wafer in a clear plastic Flouroware container and transported it to LPS for the first two steps of the process. I also supplied LPS with the photomask for patterning the first layer of the design, the capacitor back plates.

I designed the photomask for the lower layer using ICED, an integrated circuit

design editor that is available free online from <http://www.iceditors.com>. The details of the design are described in Chapter 5. I had the photomasks fabricated by Photo Sciences Inc. (www.photo-sciences.com) on 4" soda lime/chrome plates. I originally planned to pattern the aluminum myself, but it turned out to be easier to have this done by Dr. Sergiy Gladchenko from Dr. Kevin Osborn's group at LPS; he also grew the SiN_x . I made negative masks (clear images, opaque backgrounds), but positive masks (opaque images) would have been more convenient for use with the LPS aluminum deposition process.

At LPS, Sergiy deposited 40 nm of aluminum by sputtering in a 5 mTorr Ar atmosphere. He then spun on negative photoresist and exposed the pattern using the mask that I provided. After developing the resist, he etched the unprotected aluminum away by soaking in Aluminum Etchant 80-15-3-2. This left the capacitor back plates on the wafer.

In the next step, Sergiy deposited a 100 nm thick SiN_x insulation layer using an Oxford Plasmalab 100 inductively-coupled plasma chemical vapor deposition system. He followed a recipe developed in Kevin Osborn's group by Hanhee Paik. In 2009, Hanhee did an in depth study of SiN_x films created using different techniques [58]. She found that by properly setting the flow rate of the two precursor gases (N_2 and SiH_4) she could reduce the number of NH_2 bonds in the SiN_x and deposit films with a loss tangent of 3×10^{-5} at low temperature and low voltage [56]. This was the process that was used for my films.

6.2 SF_6 Etch of SiN_x

Because the SiN_x had a greater loss tangent than that of the sapphire substrate, I

only wanted the nitride over the filter capacitor plate and not under the IDC or qubit junction. Therefore, I needed to etch the SiN_x away from everywhere except from the region of the filter capacitor. For this step, I transported the sample to the class 100 cleanroom in the Fablab in the Kim Building.

In the cleanroom, I first rinsed the wafer with acetone, then methanol, and then isopropyl alcohol (IPA) in order to remove any dirt that might have contaminated the surface while transporting the sample outside. This was done by holding the 2" wafer in a gloved hand and spraying the wafer face using a spray bottle full of each solvent. Next, I placed the wafer on a hot plate at 120°C for a short amount of time to remove any residual contaminants. While the wafer was baking, I got S1813 photoresist from the refrigerator. I removed the wafer from the hot plate with my gloved hand, sprayed it off with dry nitrogen, and placed it on the chuck of a Headway EC-101 spinner. I pressed the foot pedal to activate the vacuum chuck and made sure that the sample was secure on the spinner. While the wafer was spinning, I set the speed to 4000 RPM and the time to 45 seconds. Once the spinner was set-up I stopped the spinner with the foot pedal and prepared to apply the photoresist.

To apply photoresist, I used a pipet to draw as much S1813 from the bottle as possible so as to be able to cover the entire wafer. It was important that I did not get any air trapped in the photoresist. Once the wafer was completely covered in S1813, I pressed the foot pedal and spun the wafer for 45 seconds. I then removed the wafer from the chuck with my gloved hand, sprayed it with nitrogen, and placed it on a hot plate to bake for 2 minutes at 120°C . After baking, I removed the wafer from the hot plate and placed it in a carrying case while I set up the mask aligner.

I used a Karl-Suss America, Inc., MJB-3 mask aligner to do the photolithography exposure. First, I removed the 4” mask labeled “OXID” (the layer name in ICED) from the mask carrying box. I oriented the mask on the aligner’s vacuum plate such that the mask orientation matched the layout of the back plates on the wafer. I then pressed the “Vacuum Mask” button which secured the mask in place. I slid the vacuum plate into the mask aligner, but I did not tighten it down yet.

Next, I placed the wafer on the sample chuck and slid the sample in under the mask. I needed to make sure that the sample was making good contact with the wafer. The Karl-Suss aligner has two levers on the left side of the stage near the base. The first slides horizontally. This one needed to be slid forward for the exposure to take place. The second flips forward and back and controls the coarse stage height. While looking in from the side at the sample under the mask, I flipped this one forward to see if the sample made contact with the mask. The goal was to use the fine stage control—the top dial on the front of the Karl-Suss aligner—to adjust the height such that the sample just made contact with the mask. Once I found a good height setting, I left the wafer out of contact and screwed the vacuum plate into place. Good contact between the mask and wafer was the most important factor for a successful exposure.

With the mask in place over the sample, I used the microscope on the Karl-Suss aligner to make sure that the wafer was aligned to the pattern on the mask. The “OXID” mask and the back plate mask were designed with identical alignment markers, shaped like two squares that touch at the corner. If I focused on the photomask and used the contact lever on the left of the aligner to bring the sample in-and-out of contact, I was able to see the back plates and alignment markers as dark images that either lined up or

did not line up with the images on the mask. The mask aligner was equipped with precision x and y controls for the stage as well as a control for angular rotations (these are the dials below the stage height control dial) so I got the wafer roughly into place by hand and used the stage controls to make sure the alignment was correct. I checked patterns at all four extremes of the field of view to make sure that all of the devices would have the nitride patterned correctly.

Next, I exposed the resist. To do this, I flipped the lever so that the wafer was in contact with the mask. This illuminated a blue light labeled “Contact” on the aligner. Next, I set the exposure time to 9 s using the dial above the “Contact” light. With the wafer in contact and all the levers on the left pushed forward, I pressed the “Exposure” button. This brought the UV source forward and exposed the resist. I was careful not to look directly at the source since UV light is harmful.

After the time ran down, the source moved back and I was ready to remove the wafer and mask. I brought the wafer out of contact by flipping the lever backwards and pulling the other lever towards me. I slid the wafer chuck out from under the mask and placed the wafer back in its carrier. After unscrewing the vacuum plate, I removed the plate and switched off the vacuum so I could remove the mask and place it back in its case. This left the mask aligner ready for the next user.

Developing the resist was a simple task since there was only one layer of S1813. First, I filled a Pyrex dish with CD-60 developer and another with water. I placed the exposed wafer into the CD-60 and agitated it in a circular motion for about 30 seconds. Then I rinsed the wafer in the dish of water and used the DI water spray nozzle to spray off the wafer. Finally, I sprayed the wafer with dry nitrogen to remove any water.

I did the SiN etch in a Trion Reactive Ion Etcher in the class 100 cleanroom in the Fablab. The Trion is controlled by a graphical interface. To start the system, I clicked on the “Maintenance” button and entered the password. I then clicked the “Vent Reactor” button and waited 25 s for the chamber to vent. I lifted the lid of the Trion, placed the wafer face-up in the chamber, and closed the lid. I navigated back to the home screen and clicked the “Manual Processes” button. This brought up the gas handling control panel. I set the computer to etch oxides by clicking the “Files” button, then the “Oxide” button in the window that popped up, and finally clicked “Exit” to get back to the control panel. The settings for etching SiN_x were normally in a clear page protector taped to the front of the Trion along with settings for etching several other kinds of oxides. I input the “Pressure Set” to 150 mTorr, the “RIE #1 Power Set” to 100 W, the “SF6 Set” to 40 sccm, and the “o2 set” to 10 sccm. I then toggled the “Vacuum Off” button to “Vacuum On.” This turned on the pump and the words “Roughing Chamber...” appeared in the control panel. Once “Roughing Chamber...” disappeared I turned on the pressure sensor by toggling the button that read “Pressure Off.” I then waited for the system to pump down to less than 30 mTorr. To start the etching gas flowing into the chamber, I toggled the “Gas Off” button so it read “Gas On.” I waited for the SF₆ and O₂ pressure readings to equalize at their set values before I turned the power on.

The etch rate of SiN_x with these settings was about 200 nm/min. Tom Loughran, the manager of the Fablab facility, suggests that if one is etching all the way down to the substrate, then it is best to over-etch a little so that the oxide is definitely completely removed. The only time I ever performed this process was when Tom was helping me, so we over-etched. I had instructed Sergiy at LPS to grow the nitride 100 nm thick, but a

quick measurement with a profilometer showed that it was about 140 nm thick, so I etched the SiN_x for 1 minute in SF_6 plasma. In order to start the etch, I toggled the “Power Off” button so it read “Power On.” I allow the SF_6 to etch the nitride for a full minute and then turn the power off. I turned off the flow of SF_6 and O_2 by clicking the “Gas On” button and turned off the pressure gauge. In order to clear the chamber of etching byproducts I waited a few minutes before turning off the pump, which I did by clicking the “Vacuum On” button. The etching process was complete and there were no visible defects on the surface that I could attribute to over-etching or under-etching.

I left the gas handling screen by clicking “Exit,” which brought me back to the home screen. I clicked the “Maintenance” button and input the password again. Then, I clicked “Vent Reactor” to vent the chamber and removed the wafer. I placed the wafer in a Flouroware wafer carrier for transport back to LPS. I closed the lid and got back to the home screen where I clicked “Standby” before leaving the system. Before I left the gowning room, I placed the wafer carrier in a baggie of Drierite to protect the wafer from the humidity while I carried it back to the Physics building and then to LPS for dicing.

6.3 Sapphire Wafer Dicing

The wafer with the patterned SiN_x was then diced into individual 4.5 x 4.5 mm chips using a DISCO DAD321 diamond saw at LPS. Vitaley Zaretsky did the dicing for me. The only time I tried to do the dicing, I managed to destroy the blade on the saw. Vitaley says that it is necessary to use a thick blade with a large grit when dicing sapphire. He recommends a hubbed resinoid blade type CX-010-325-080-H from Dicing Blade Technology. The diamond saw was set to use a 0.75 mm/sec cut speed and a spindle

speed of 22000 RPM. He diced several sapphire wafers for me and he did not report any problems when using these dicing parameters.

It is also possible to dice the wafer by hand using a carbon steel tipped scribe (diamond tipped scribes do not work as well for this), but it can leave the edges of the chips a little rounded making them difficult to handle with tweezers. Also, hand dicing is undesirable because it produces chips of irregular shape and size—both of which can affect the regularity and thickness of the e-beam resist. Dicing the wafer gave me nearly 100 nominally identical samples on which I could perform e-beam lithography.

6.4 Individual Chip Preparation: O₂ Plasma Cleaning

Correct preparation of the sample was critical to the success of the e-beam lithography. In order for the electron beam to be optimally focused on the surface the resist had to be as flat as possible. Particulates, chemical residue, water, etc. on the surface of a sapphire chip could cause the resist to spin on irregularly.

Before the wafer was diced, a layer of contact paper was placed on the back of it. After dicing, the chips remained stuck to the contact paper and they needed to be pried off. To do this, I bent the contact paper back to expose the edges of the chips, and then used tweezers to grip the sides of a chip and twist it loose. I found that if I tried to get the tweezers under the chip to lever it up and off of the contact paper, it was likely to fling the chip across the work bench. It is also worth noting that I wore Nitrile gloves whenever I handled the chips, unless I was performing an exceedingly simple task such as placing a chip under an optical microscope.

I usually wanted at least 4 to 6 chips before starting the e-beam process in case

one or more of the samples was dropped, flew off of the spinner, or suffered some other catastrophe. In fact, there was a good chance that some samples would end up with irregularities on the surface of the resist even if everything went well otherwise. After I extracted enough chips from the backing, I placed the chips in a Gelpak® sample carrier with a Gel Retention Level (*i.e.* stickiness) of 4 and took them to Rm. 1314 in the Physics Building for processing under a fume hood.

Before spinning on e-beam resist, I cleaned the samples thoroughly. During dicing the dicing saw sprayed the wafer with a cooling fluid (mostly water) while the saw cut and this fluid left residue on the chips. Organic solvents were adequate for removing this residue. I filled three pieces of Pyrex glassware with solvents: one with acetone, another with methanol, and the last with isopropyl alcohol (IPA). The acetone was the main cleaning agent whereas the other two were mostly used for rinsing. I placed the chips in acetone for 5 minutes and scrubbed them gently with a cotton swab. I then moved them quickly into the methanol and then into the IPA. Acetone leaves residue behind as it dries; therefore, I used the other solvents to quickly rinse the samples before the acetone could dry. I then took the samples out of the IPA and dried them off with dry nitrogen. The process removed the majority of the contaminants from the surface, but it was possible that the solvents could have left some residue behind.

I finished the cleaning process with an oxygen plasma cleaning. For this step, I took the samples from Rm 1314 to Rm 0219 and cleaned the chips in a Technics PEII-A Plasma System. I turned on the plasma system by flipping the "Main" switch to on, pushing the large, black "Power On" button, flipping on the group of three switches clustered near the markings "Filter Unit" and "Etch Unit", and finally turning on the

pump by flipping the "Vac Pump" switch. Next, I turned on the Tegal Corporation water circulating system-- located under the bench-- by flipping the main switch to "Cool/Heat." The PEII-A system was equipped to supply several different etching gases through a gas handling panel on the back wall, but since I only wanted to use oxygen, I shut these other gases off by turning the valve on the gas handling panel to "X."

Before placing chips in the chamber, I cleaned the chamber with an O₂ plasma because the system was used for many different processes that might have left some contaminants inside the chamber. For this step, I left the chips out of the chamber and began pumping on the chamber. The system controls were concealed behind a red panel on the front of the chamber. I opened this panel and flipped the "Sol'n" switch to start pumping down the chamber with a Franklin Electric Model 1101006401 rotary pump. The chamber usually got to a pressure of ~10 mTorr using the rotary pump. Next, I purged the O₂ line by flipping the switch in the control area marked "O₂" and "Gas #2" to "On" and turning the knob *counter-clockwise* to increase the flow rate such that the pressure in the chamber read 300 mTorr. I allowed this gas to flow for about 1 minute. In order to begin the plasma clean I flipped the switch in the "Power" control area to "On" and turned the knob *clockwise* to increase the power to 300 W. I allowed the plasma to clean the chamber for about 5 minutes. When the chamber had been cleaned out, I turned down the power and flipped the switch to "Off." I turned off the "O₂" gas flow switch-- it was convenient to leave the O₂ control knob alone. It was important to purge the system of any vaporized contaminants at this point, so I let the system pump down for about 5 minutes. This process removed any water vapor or contaminants left behind by other processes from the chamber.

At this point, I was ready to put my samples into the chamber. I turned off the "Sol'n" switch and flipped the "Vent" switch to vent the chamber. I put the samples in the chamber and repeated the same procedure, except that I did the plasma clean at a pressure of 200 mTorr O₂ and a power of 400 W for 1 minute. After I cleaned the chips I vented the system and placed them back in the sample carrier.

In order to shut down the system, I turned everything off in the reverse order to how it was turned on. That is, I turned off the water circulator, the "Vac Pump" switch, the cluster of three switches, pushed the large, red "Off" button, and, finally, flipped the "Main" switch to off. I closed off the O₂ gas bottle to make sure that it did not leak while the system was idle. The plasma cleaning should have burned away all organic material on the chip, leaving a pristine surface on which I could spin e-beam resist.

6.5 Preparation of Electron Beam Lithography Resist

The e-beam lithography process utilized three layers of resist: 2 MMA (methyl methacrylate) layers and 1 layer of PMMA (polymethyl methacrylate). The resist had to be hard baked for the correct amount of time so that all of the features of the lithography developed correctly. The choice of copolymer (undercut) resist and pattern resist was also important. I used MMA EL 11 from MicroChem because it spun on the thickest of all the copolymer resists. The pattern resist was 950 PMMA C2, which is a standard pattern resist. Correct preparation of the e-beam resist was important to the success of the lithography.

The resist layers were spun on in the fume hood in Rm. 1314 where the spinner was located. Before I put a chip on the spinner, I baked it on a hot plate as the plate

warmed up to 150° C. The purpose of this step was order to remove water that might have condensed on the surface. The spinner that I used was a Laurell Technologies Corporation Model WS-400B-6NPP/LITE/IND spinner. To get it ready I first turned on a nitrogen gas bottle that sent a constant stream of dry nitrogen into the spinner chamber. Then I turned on the mechanical pump for the vacuum chuck. I selected the spinner chuck with the smallest O-ring (3/8" diameter) and placed it on the spinner. I then set the spinner so that it would run at 4000 RPM for 45 s and would accelerate at a rate of 1000 RPM/s. The process required pipets, cotton swabs, a dish of acetone, and the MMA resist, so I got these together in the fume hood before I got the samples off the hot plate.

I then removed one chip from the hot plate and used dry nitrogen to cool it and remove any dust. I activated the vacuum on the chuck before I placed the chip on it. This way it was less likely that the chip would fall off the chuck and into the chamber. Next, I applied 2 or 3 drops of MMA resist with a disposable pipet. I needed to be sure that no air got into the resist during this step. If some bubbles formed inside the pipet, I discarded the pipet and began again. Also, the first drop from the pipet was always allowed to fall into the chamber, not onto the sample. I closed the chamber lid and pushed the "Run/Stop" button on the control panel. Once it stopped, I removed the sample using tweezers and turned off the vacuum.

Spinning always leaves a bit of resist on the back side of the sample. This resist can dry and leave an irregular surface preventing the sample from sitting flat in the SEM. It was easy to remove the resist while it was still wet by scrubbing the back of the sample with a cotton swab and some acetone. I needed to be careful during this step not to use a cotton swab that was saturated with acetone because if I applied too much, the acetone

tended to spread out the surface, wick around to the front of the sample, and ruin the resist.

Once the resist was completely removed from the back of the sample I checked the quality of the resist layer. When viewing from the correct angle, I could see whether there were any bubbles, streaks, or irregularity in the resist surface. A good layer of resist looked completely flat except for at the corners. The color was usually a uniform greenish-blue. If it looked good, then I moved it to the hotplate. If it had too many irregularities, then I placed it in some acetone and held onto the chip for another attempt after cleaning away the old resist.

After I obtained a good resist layer, the MMA was baked for 11 minutes at 150° C. The development of the pattern was relatively sensitive to this bake time. While developing the process, I tried many bake times (*e.g.* 7 min, 9 min, 20 min, 1 hr, *etc.*) and 11 minutes gave the best results. While the sample was baking, I had time to spin about 3 other samples. During this process I kept track of the baking times with one timer and marked in a notebook when each sample would be ready to be removed from the hot plate.

Spinning the MMA on at 4000 RPM for 45 s, with a ramp up of 1000 RPM/s, created an MMA layer that was about 500 nm thick. I needed two of these layers so that I would have enough room under the suspended bridge for the double-angle evaporation. Therefore, I repeated the above process to create a second layer of MMA. I did a quick experiment to check whether I needed to bake each layer separately and found that it was not strictly necessary. One could spin the layers one after the other and achieve virtually the same thickness of MMA. However, I have not determined the correct bake time for

an MMA layer created in this way. A little more research could work this out, but I found that baking each layer separately, for 11 minutes each, worked.

Next, I applied the PMMA layer. The process was the same as for the MMA layers, but the spinner settings were different. I spun the PMMA on at 4000 RPM for 45 seconds with a ramp rate of 3000 RPM/s. This created a layer of PMMA with a thickness of about 100 nm. I needed to bake this layer for a very long time. Because of charging effects during the e-beam write caused by the insulating sapphire substrate, it was common for the pattern to be exposed outside of the area determined by the design. Therefore, I baked the PMMA layer for 4 hours at 150° C to ensure that all of the solvents had been baked out of the resist, making it harder to remove with the developer. Again, I went with what worked. One hour was not enough, but 2 or 3 might be. Four hours did the job, so I baked the PMMA for 4 hours. That concluded the process for spinning and baking the e-beam resist.

6.6 Depositing the Aluminum Charge Dissipation Layer

Because the sapphire substrate is an insulator, I needed to provide a path by which the electrons could dissipate during e-beam writing. I found that a 20 nm thick layer of aluminum placed on top of the e-beam resist facilitated the dissipation of this charge quite well. Figure 6.1 shows an example of patterns written with and without this charge dissipation layer. E-beam lithography was virtually impossible to perform without it.

After baking the PMMA/MMA layers as described above, I took the chips off of the hot plate and put them into the Cooke Vacuum Products (CVP) evaporator in room 0219. Before I vented the evaporator chamber, I attached the chips to the sample stage.




Fig. 6.1 A comparison of e-beam lithography with (right) and without (left) an aluminum charge dissipation layer on the resist. (a) Immediately after exposure, (b) after developing in MIBK, and (c) detail of the qubit. Without the aluminum to disperse the charge from e-beam writing, the resist was degraded during the writing process.

The stage was a $\frac{1}{4}$ " thick brass strip about 1" wide and 4" long with CuBe clips. The clips were made by punching a $\frac{3}{8}$ " diameter hole in a $\frac{1}{16}$ " thick piece of beryllium copper and using scissors to cut out a tear drop shape with the hole in the wide part of the tear drop.

Before I put the chips on the stage I again examined the quality of the resist on each chip. Not all the samples had perfect resist. In fact, it was often the case that none of them looked perfect, but some were always better than others. A good chip had a region of flat, regular resist somewhere near the middle of the chip. I decided on a hierarchy for the samples according to the quality of their resists. I arranged the chips near the middle of the stage, and placed the clip over a corner of the chip where the resist looked the worst.

I vented the evaporator chamber with dry nitrogen by opening the "Vent" valve, and then lifted the lid off of the evaporator and swung it clear of the bell jar. The sample stage was attached to the tilting stage under the lid with a flat-head screw. I oriented the sample stage such that the samples were in the middle of the tilting stage. The angle of the tilting stage was adjusted such that it was horizontal, and the shield was rotated around to cover the stage.

Next, I placed the Al source in the evaporator. Aluminum is evaporated from a bare tungsten wire basket boat because when aluminum melts it tends to spread out and cover a conventional boat all the way to the electrodes. The wire basket I used was an R. D. Mathis Company model ME17-3X.025W. I clamped the wire basket between the electrodes at the bottom of the evaporator and screwed it in place with an allen key, taking care that the basket sat flat after the screws were tightened down. Then I placed

two pieces of 99.999% pure aluminum shot in the wire basket, one on top of the other with the flat sides touching. It was important that I positioned the bottom piece such that it had as much contact with the wire basket as possible to ensure that it melted evenly. Also, I arranged the top piece of aluminum so that it had some contact with the wire basket. I marked in the log notebook the number of the electrode to which the boat was attached.

I oriented the lid of the evaporator such that the marks on the lid and the bell jar marked “QC” lined up and then held the lid shut and turned on the Franklin Electric model 1091045400 rotary pump (the mechanical pump). I then opened the “roughing valve” switch to start pumping on the chamber. The mechanical pump usually took the pressure in the chamber down to below 400 mTorr in about 2 minutes. If the pressure did not reach 400 mTorr, then I closed the roughing valve, opened and closed the vent valve, and opened the roughing valve again. This let dry nitrogen into the system that helped to remove water that might have been out-gassing inside the chamber. This trick normally allowed the chamber to be roughed down to a pressure below 400 mTorr.

Next, I closed the roughing valve and opened the “hi-vac valve” to the CTI-Cryogenics Cryo-Torr 8 cryopump. The pressure normally got down below 5×10^{-5} Torr in less than a minute. This was a good, quick check that the system did not have a leak. The charge dissipation layer did not have to be the highest quality Al film so I only pumped the chamber down to a pressure of 4×10^{-6} Torr, which took about 15 minutes. Once I reached this point I was ready to evaporate aluminum.

First, I turned on the cooling water to the evaporator electrodes and the crystal monitor on the ion milling top. Next, I made sure that the crystal monitor was connected

via a BNC cable to the “XTAL” side of the Sycon OSC-100 Oscillator box and that the “CONTROLLER” side of the OSC-100 was connected to the Sycon STM-100 / MF Thickness / Rate Monitor. The crystal monitor controller was programmed with the specifics (*e.g.* density, atomic mass, *etc.*) of several metals, and I set the “program” using the front panel to the one specified for Al (program 8 as of 06/08/10). To do this I pushed the “program” button and used the up and down keys to select “FL.8.” Then I pressed “Enter” repeatedly until I cycled through all of the specifications and came back to the display reading “FL.8.” I pressed “program” again to get back to the thickness and rate display. I selected the electrode number with the “source selector switch” on the lower right of the CVP evaporator front panel by depressing the trigger and rotating the switch to the number that I recorded in the log book earlier.

At this point I was prepared to begin applying current to the source. Evaporating Al required some care. If the current is increased too quickly, then the boat would break before I could evaporate anything. If I turned it up too slowly, then all of the Al would melt and drain out of the basket before it evaporated. I developed a process that had a better than 90% success rate, but it had to be followed closely.

Warming up the source required specific timing. First, I flipped the power switch of the Model FPS 2-41 Evaporator Source Power Supply and turned up the “manual output control” to the threshold value of 2.3. This usually caused the source current to jump to 25-30 A. Then I zeroed the timer on the crystal monitor display. I waited for 1 minute at a value of 2.3. After 1 minute, I turned the output to 2.4 and waited another minute. Then I turned to 2.45 and waited 1 more minute, and then turned to 2.5, and waited 30 s. I proceeded using steps of 0.05 for 30 s until the evaporation began between

2.55 and 2.7. During the evaporation the pressure on the ion gauge jumped up to nearly 10×10^{-6} Torr. The evaporation rate on the crystal monitor display began to rise as well, but the pressure increase came first. Once the evaporation began—no matter what the output was set to at the time—I set the output to 2.7. The film thickness and evaporation rate were displayed on the crystal monitor front panel, and the rate would increase rapidly to 10-12 Å/s. If the rate stalled out at less than 10 Å/s, then I turned the output up a little bit (~0.1). I was extremely careful not to turn it up too high too quickly or the evaporation would fail. If the rate was below 10 Å/s for too long, then the evaporation would also fail, so I kept it above 10 Å/s with small increases of the output as needed.

After the aluminum started evaporating I evaporated about 150-200 Å before I opened the shield. This prevented moisture or oxidized material on the outside of the aluminum shot from being deposited onto the chip. I then opened the shield by rotating the shield knob a half-turn counter-clockwise and pressed the “zero” button on the crystal monitor display. Once the thickness reached 200 Å, I closed the shield, turned down the manual output control to 0, and turned off the power to the current supply. I allowed the sample to cool in the chamber for 10 min before turning off the cooling water and venting the chamber.

6.7 Electron Beam Lithography

After depositing the charge dissipation layer, I transported the samples to Rm. 2219 for e-beam lithography. The Phillips XL-35 SEM takes about one hour to warm up and be ready to write a pattern.

First, I used copper clips to mount a chip on an aluminum chuck. Using a

diamond scribe I made a scratch on the sample surface near one corner extending about a millimeter from one edge of the sample. I used this scratch later for focusing on the sample. I pressed the “Vent” button in the SEM controller software and waited about 30 s before opening the chamber. This ensured that moisture and particulates were not sucked in unintentionally. I placed the chuck in the chamber with the chip near the center of the SEM stage. There was a nut that held the SEM stage in place near the base of the stem of the stage. This nut would frequently become loose causing the stage to wobble. I took care to tighten it every time I put a sample in the SEM. I then clicked the “Pump” button while holding the chamber door closed. It usually took about 15 minutes for the chamber to reach a pressure of 2×10^{-5} mbar.

While the system was pumping down, I checked that the beam blanker input at the top of the SEM was grounded. Leaving the beam blanker connected while I turned on the beam would seriously damage it. Once the chamber was under vacuum I pressed the “HT” button on the front of the SEM bench and clicked the “30 kV” button in the SEM computer interface program. (A dialog box comes up at this point, but I was careful not to click it because it would set the height of the SEM stage prematurely.) For lithography I used 30 kV because I wanted the electrons moving as fast as possible when they hit the sample surface. This gave the best resolution for the lithography. If I was just imaging, then I used a much lower voltage (1-5 kV) so the beam did not miss anything on the surface.

After I set the voltage and turned on the beam, I had to saturate the filament current. The SEM beam current should roughly scale with the filament current (see Fig. 6.2). I went to the “Beam” window in the control program, set the spot size to 1, and

began turning up the “Fil. Current,” 1 μA at a time. The image of the Faraday cup in the middle of the SEM stage would gradually grow brighter as I turned up the filament current. Once I was able to see a relatively good image, I fine tuned the saturation by looking at the electron beam in cross-section. First, I moved to a relatively even part of the SEM stage and set the magnification to 500-1000X. Next, I clicked the “XOver” box in the “Beam” window. At low currents, the image of the beam looked like a spider web, but as I increased the filament current the image consolidated into a solid oval shape. The display settings would usually hide some irregularities in the beam shape, so I turned the brightness all the way down and slowly increased the contrast in order to get an accurate picture of the beam.

At this point the cross-section of the beam was usually shaped like a kidney. I continued to turn up the filament current while decreasing the contrast until the image was an oval. As I increased the filament current beyond that point, the only thing that changed was the brightness of the oval. The correct value for the filament current was the one at which the oval did not look substantially brighter one click higher nor substantially dimmer one click lower. Of course, if I continued turning up the filament current, the beam would get stronger, but the lifetime of the filament would decrease significantly at higher currents. Therefore, I looked for the filament current at which the beam current was just beginning to saturate.

The filament took time to warm up, so I would reduce the filament current by 10 μA , and wait 20 minutes for it to warm up. While I waited, I connected the beam blanker cable at the top of the SEM. Then, I resaturated the filament using the same procedure. The second time I saturated the filament, the saturation value was usually a




Fig. 6.2 Illustration of the behavior of the SEM beam current as the filament current is increased. Since the filament lifetime is decreased at high filament currents, it is best to set the filament current such that the beam current is just beginning to saturate (dashed line).

few μA lower than the initial value. Carefully saturating the filament current gave an optimized beam current while conserving the lifetime of the filament.

Next, I coarsely set the stage height and measured the beam currents for the different aperture sizes of the SEM. I left the “XOver” mode and zoomed out to the lowest magnification possible (Note: the + and – keys on the number pad are used for quick zooming). At this point, the “Faraday cup” was in view. The Faraday cup is an anode that I used for measuring the beam current. I began by focusing the image at a low magnification. This was done by holding the right mouse button and moving it left and right. The focus control was related to the magnification setting; at lower magnifications the focus was very coarse, whereas the focus control was very fine at high magnifications. Once the image was roughly focused I increased the magnification and sharpened the beam. My technique was to zoom in on a small feature and focus on it, then find a smaller feature, zoom in on it, and refocus. Once the image was focused at 5000X magnification, I clicked OK on the dialog box that I had set aside earlier and set the “work distance” or z-height for the SEM.

After setting the work distance, I moved to the center of the Faraday cup and set the magnification to 1000X. The Faraday cup was connected to a BNC output on the door of the chamber. I connected this output to a DL Instruments 1211 preamplifier and set the sensitivity of the preamp to 10^{-11} A/V. The suppression was set to 10^{-4} amps with the calibration off and the rise time was set to 0.1 ms. I connected the 600 Ω output of the preamp to a Fluke 87 III True RMS Multimeter and set the multimeter to read dc volts. This setup allowed me to read the beam current with a sensitivity of 0.01 pA.

Next, I adjusted the “Gun Tilt” such that the path of the beam was perpendicular to the

sample surface. In the “Beam” menu I clicked on the field of the “Gun Tilt” control. This brought up cross-hairs on the computer screen. I carefully positioned the cross-hairs such that I got the highest value possible on the voltmeter display. I compared that value to the value recorded by the previous user, and if it was consistent, then I knew that the Gun Tilt was set correctly. This value gave me the beam current when the spot size was set to 1. At this time I recorded the beam current for the various spot sizes that would be used in writing the pattern. These beam currents were used in the NPGS run file for the lithography. The spot size was related to the size of the aperture and thus to the size of the beam. The beam size scaled with the formula:

$$r_{spot} \approx 2^{(spot-1)} r_1 \quad (6.1)$$

where r_1 is the radius of the beam when the aperture was set to spot 1 and $spot$ is the spot number. Since the current density of the beam was the same regardless of the spot size, the beam current scaled roughly as:

$$BC_{spot} \approx 2^{2(spot-1)} BC_1 \quad (6.2)$$

where BC_1 is the beam current at spot 1. As a general rule, I used spot 1 for the finest, most important features, spots 2-4 for features in the range of about 1-5 μm wide, spot 5 for lines that were wider than 5 μm , and spot 6 for the contact pads. Once the beam currents had been recorded I unplugged the amplifier, turned it and the multimeter off, and grounded the input to the Faraday cup. I took care to set the aperture back to spot 1 before leaving the Faraday cup.

Next, I moved the stage to the gold standard in order to perfect the focus of the beam. I zoomed all the way out and moved to the gold standard by clicking the drop down menu in the “Stage” window and double-clicking “Gold.” This automatically

moved the stage position such that the gold standard was in view (Note: focusing on the gold standard should be done with the smallest spot size that is being used for lithography). I proceeded by focusing on smaller and smaller features as before. I was usually not able to improve the focus using the conventional focus control once I reached a magnification of 50000-100000X. This was because the focal plane of the SEM was not parallel to the surface of the gold standard; an effect known as “astigmatism.” I fixed this by adjusting the “stigmator,” which was accessed by pulling down the drop-down menu in the upper right corner of the screen and selecting “Imaging.”

After zooming in to 100000X and achieving the best focus possible with the conventional control, I clicked in the stigmator field. This causes cross-hairs to appear on the computer screen. I started with the cross-hairs in the middle of the image and swept them slowly to the left edge of the screen, then back to the right edge of the image, while carefully observing the quality of the image. I released the mouse button at the point where the image looked the best. Then I repeated the same process in the vertical direction. This usually improved the quality of the image significantly. Next, I increased the magnification to 200000X and repeated the process to remove horizontal and vertical astigmatism. Just as with focus control, the stigmator made finer adjustments at higher magnifications, so it was often necessary for me to move the crosshairs farther from the center to get the sharpest image at higher magnifications. I found that there was often a shift of the image as I swept the stigmator control over the place where the image looked the best. The best position for the stigmator was right near that shift. When everything was correctly adjusted, the image usually showed several very bright gold grains with the area between them being very dark (see Fig. 6.3).

Setting the stigmator to the correct position was crucial to achieving the best resolution while writing. If there was astigmatism in the beam's focus, then the dose from the electron beam would be delivered asymmetrically. It was the difference between the beam making a nice round spot as opposed to an oval on the sample surface. A circle gave the best resolution for the lithography. Removing astigmatism also enabled the focus to be correct over a wider area of the sample. Because the focal plane was parallel to the surface, the focal point on the corner of the sample would be the same as in the center of the chip.

After focusing on the gold standard I needed to focus on the sample surface. I moved back to the Faraday cup by using the drop-down list in the stage menu. Then I zoomed all the way out and moved to the sample. I was sure to use the smallest spot size while I focused on the sample because the SEM was constantly exposing the resist as the sample was being imaged. The first thing I did was align the sample so that the pattern would be written squarely on the sample. After focusing on the sample from a low magnification I used the "Scan Rotation" controls in the "Imaging" window to align one edge of the sample with the edge of the field of view. This only aligned the direction that the SEM sweeps the beam while making the image not the actual sample position. I also had to rotate the sample stage. To do this, I committed the scan rotation value to memory and zeroed out the scan rotation by clicking the "0" button. Then I went to the stage menu and clicked the option marked "Rel," which stood for making a relative stage movement as opposed to "Abs" for absolute. I input the scan rotation value and clicked "Go to."




Fig. 6.3 SEM picture of the gold standard after removing astigmatism from the beam. The fact that the gold grains look bright and well shaped, with high contrast to the areas in between, shows that the focal plane of the microscope is parallel to the sample surface. This is crucial to ensuring that the beam will have a tight, circular, cross-section at the sample surface and that the focus will be correct over a large area.

It was important that I align the sample before completing the focusing process because the stage movement could move the stage height and nullify a lot of hard work.

Next, I found the scratch that I made in the resist while I was attaching the sample to the chuck and focused on that. I actually focused on the debris that settled on the surface as a result of the scratch, instead of the inside of the scratch itself. I needed to get a good focus at the maximum magnification (200000-300000X). Once the focus seemed adequate I wrote a contamination spot in order to double check the quality of the focus.

To create a contamination spot, I found a clean region on the sample and set the magnification to 50000X. Then I pulled down the “Scan” menu from the top of the window and changed from “Full frame” to “Spot” mode. This caused a pair of cross-hairs with an X in the middle to appear over the image. The X marked where the beam hit the sample. I then waited 60 s while the beam deposited a layer of carbonaceous material on the sample surface [67]. Next, I changed the scan mode back to “Full frame.” A bright spot was normally visible where the cross-hairs were centered. If it was not visible, then I zoomed in on the area. If it was still not there, then either 60 s was not long enough to form the spot or the focus was not good enough. If necessary, I moved to another spot and tried leaving the SEM in spot mode for 90 s. If there was still no spot, then I tried to improve the focus. With an adequate focus, the spot should have been visible at 50000X.

Once I observed a spot, I zoomed in and checked the shape and diameter. Ideally, the spot should be round with a uniform brightness around the circle. If the spot was elongated or it was brighter on one side, then I needed to adjust the stigmator. To adjust the stigmator, I found the smallest feature on the sample surface and used it to

perfect the focus. I continued to improve the focus, adjust the stigmator, and write contamination spots until I could make one that had the right shape and a diameter of < 70 nm.

Before I wrote a pattern, I needed to find a clean spot on the sample that was large enough to fit the whole pattern. First, I zoomed out to take a look at the sample. Next, I zoomed to about 30X and turned on the “Average” function of the control software in the “Filter” drop-down menu. The filter picked out any large debris that would block the SEM from writing vital parts of the pattern. I tried to work quickly because the resist was actually being slightly exposed as I searched. Once I found a likely area I clicked “Freeze.” This kept the image on the screen while blanking the beam, thus protecting the resist from unnecessary exposure. I used the measurement tool to make sure that I was far enough from the edges of the sample for the entire pattern to fit on the chip, and to check the distance to any large pieces of debris to make sure that they would not affect the write.

With the SEM centered on a clear area I turned control of the electron beam over to the beam blanker by first turning the knob on the ScanService Corporation model 880 beam blanker from “ON” to “EXT”. Then I pulled down the “Scan” menu in the computer software and switched from “Full Frame” to “External XY.” The order was important. First, I turned the knob, then I changed the scan mode in the software. If I switched to “External XY” before I transferred control to the beam blanker, then the beam would sit in the middle of the sample unblanked and burn a hole in the resist exactly where I wanted to write the qubit junction. The SEM was prepared to write the pattern; however, it still needed the right instructions from the patterning program.

The SEM used the Nanometer Pattern Generation System (NPGS) to write the pattern. NPGS uses a file from DesignCAD: an integrated circuit computer design program. I separated the pattern in DesignCAD into different layers so that NPGS could use different writing parameters for features of different sizes. When assigning layers to the pattern, I marked the smallest features as layer 1 and worked up to the largest features. NPGS assimilated the pattern and created a run file with the write parameters for each layer. To create the run file I entered the command “mrf [pattern name]” into the DOS prompt on the computer on the workbench that interfaced with the beam blanker. A blue screen came up and I inputted the name of the pattern where it said “PatName.” I pressed Ctrl+PgDn, which caused NPGS to assimilate the pattern and brought up the write parameters for the first two layers.

The values for the parameters were default values, but I only needed to change 5 of them:

- 1) **Write, Pause, Skip (w,p,s):** It was necessary sometimes to change the settings of the SEM between layers. If the magnification or the spot size had to be changed between two layers, then I set this parameter to “p” for pause. If the layer used the same magnification and spot size and only needed different dosage settings than the previous layer, then it saved time during the write to input “w” for write. Layers that were only guides for the eye in the pattern or text were skipped by choosing “s.”
- 2) **Magnification:** When it assimilated the pattern, NPGS calculated the maximum magnification necessary to encompass all the features on each layer. The finest layers were written at 1000X, which gave a field of about 80 μm x 80 μm . If

the default value was less than 1000X then it was necessary to reduce the value by about 5%—NPGS's calculations are not perfect and the default magnification might actually leave some small corner of a feature out of the field of view during the lithography. If while the pattern is being written, the program encounters a feature it cannot write because it is beyond the field of view, then the program will immediately abort and the sample will be ruined.

- 3) **Dosage Spacing:** NPGS takes each feature and divides it into a grid, as depicted in Fig. 6.4. In each box of the grid, the SEM will deposit a dosage of electrons. The “Center-to-Center” distance will be the horizontal spacing, whereas the “Line Spacing” will be the vertical spacing. When choosing these values, it was important to consider the feature size and the size of the beam spot being used for writing. I chose a dosage spacing that would allow the features to be divided evenly into a reasonable number of dosages. Also, I considered the fact that each dosage would be centered at a certain point, but that did not mean that it would be confined to a square on a grid. The dosage can be described by a Gaussian of intensity in the x-y plane. The width of the Gaussian was roughly the width of the beam spot. I focused the SEM such that the radius of spot 1 was < 35 nm. Table 6.1 shows the spot radiuses, given a radius of 35 nm for spot 1. Dividing the feature into just the right number of dosages minimized the amount of time it would take to write the pattern.
- 4) **Beam Current:** This number represented the current that would be used to write the layer. I input the value in pA that I recorded from the Faraday cup. I used spot 1 for the finest features. Higher beam currents were used for larger features in order to reduce the write time.




Fig. 6.4 Diagram of dosage spacing for E-beam lithography. The center-to-center distance is the spacing in the x-direction and the line spacing is in the spacing in y. The dashed circles illustrate the radius of the beam spot that applies the dosage. The graph illustrates how I want to leave a relatively unexposed region between the two features that define the suspended bridge for Josephson junction fabrication.

Table 6.1. Approximate radius of the beam spot if the Phillips XL-35 is optimally focused. I used these radius values for deciding the correct settings for the dosage spacing.

Spot	Radius (Å)	Beam Current (pA)
1	35	4.0
2	70	16
3	140	64
4	280	250
5	560	1000
6	1120	4000

5) **Dosage:** I selected either an area dose or a line dose by pressing the spacebar. I used an area dose for sensitive features (*e.g.* Josephson junctions) and large area features (*i.e.* contact pads and ground planes). I only used line dose for thin features such as the bias leads. The dosages were in $\mu\text{C}/\text{cm}^2$ and nC/cm , respectively. I typically used $200 \mu\text{C}/\text{cm}^2$ for the area dose and $1.3 \text{nC}/\text{cm}$ for the line dose (Note: according to rough calculations the line dose is much stronger than the area dose). I used a higher dose for fine features and a lower one for very large features because large features were made with many, many doses while the fine features were made with relatively few. When trying to achieve a high resolution, I often experimented with a wide range of dosages to find the right one. The dose time—displayed directly above the dosage value on a line labeled with the layers color—changed depending on the dosage value and spacing. I always made sure that this time was greater than the minimum

dose time of 10 μ s. It was limited by the speed at which NPGS could unblank and blank the beam. If the dose time read 10 μ s, then the actual dosage delivered would be higher than desired—possibly much higher. I always adjusted the dose value or the dose spacing such that the dose time was 10.1 or higher.

Once the parameters were adjusted in the run file, I checked the run file with a test write. I ran the test by inputting “pg [pattern name] t”. The “t” at the end made NPGS run through all the commands from the run file and calculate the time it would take to write the pattern, but not actually write a pattern.

Writing an entire pattern usually took 10-30 minutes depending on the number of large features (contact pads, ground planes, and bias leads). If the time was too long (*i.e.* hours) then I knew there was a mistake in either the pattern or the run file. In the pattern, I looked for stray features in the design that may have been copied by accident, for example. In the run file I checked to see if the beam currents, dosages, and spacings were appropriate. If the time was shorter than expected, then I checked to make sure that nothing was accidentally deleted and, again, checked the dosages and currents in the run file. However, if the write time seemed reasonable, then it was finally time to write the pattern.

To commence the write, I input “pg [pattern name]”. The pattern computer booted the writing program, played some music, and waited for user input. I set the magnification and spot size appropriate for the first layer (spot 1) using the SEM control software. If there was any problem or doubt about the settings, then the write could be safely aborted by pressing “enter” instead of “space bar.” However, if I was sure that the

parameters were correct, then I pressed “space bar” on the pattern computer keyboard to begin writing. When NPGS encountered a layer which required a pause, it stopped and displayed the necessary magnification and beam current—which corresponded to a spot size—for the next layer. I changed the settings on the SEM controller computer and pressed “space bar” on the pattern computer. When it was finished, the computer displayed the write time. That completed the e-beam write.

In order to preserve the pattern I used the “Stage” window to move back to the Faraday cup and reset the spot size to 1 before exiting the “External XY” mode. After moving away from the sample I returned to full frame mode and turned the switch back to “On” on the beam blanker controller. I was careful to unplug the beam blanker and ground the input before turning off the filament since discharges that occurred during the shutoff could have damaged the beam blanker. I reduced the filament current to zero in the “Beam” menu and clicked the “30 kV” button to turn off the beam. In order to avoid oxidizing the filament, I allowed it to cool down for 10 minutes before venting the chamber by clicking the “Vent” button. Again, I waited for the chamber to vent fully (30 s after the hiss) and removed the chuck from the sample stage. I closed the chamber and started it pumping down so that the system did not remain at atmosphere longer than was necessary. I carefully removed the sample from the chuck, making sure to hold the clips steady with tweezers.

6.8 *Developing the Pattern*

After writing the e-beam pattern, the chip was placed in the sample carrier and moved to Rm 1314 for developing. There were several steps in the development process.

First, the aluminum had to be removed using CD-26 developer. Since CD-26 did not react with PMMA, it was safe to leave the sample in the CD-26 for 5 minutes to ensure that the Al was completely removed. After that the reaction was stopped by rinsing in DI water. Then the exposed PMMA was removed using a fairly dilute solution of MIBK and IPA (1:3 ratio, MIBK to IPA). The development time was between 8-10 s in the MIBK solution. After soaking in the MIBK solution, I rinsed the sample for 5-10 s in a dish of IPA. The IPA washed away the MIBK, and it helped to create the undercut by attacking the MMA.

Once the sample was developed, I inspected it with a Nikon Optiphot-2 binocular optical microscope. I illuminated the sample from the top because this gave the best view of the resist. Underdeveloped features were visible as discolorations of the resist. Once it was clear that the pattern was fully developed, I secured the sample in a dry box while I returned to Rm 1314 to clean up the developer chemicals.

6.9 Josephson Junction Construction by Double-Angle Evaporation

For double-angle evaporation, I began by affixing the sample to the sample stage on the cryopumped evaporator system in Rm 0219. I oriented the stage diagonally on the tilting stage such that the sample sat in the middle of the area protected by the evaporation shield. I needed to take care when I clipped the sample on to the sample holder that the suspended bridge created during the lithography would be parallel with the axis of rotation of the stage (see Fig. 6.5). Once I got the sample clipped to the stage I vented the evaporator chamber as described in Sect. 6.6.

I screwed the sample stage onto the tilting stage in the correct orientation and set

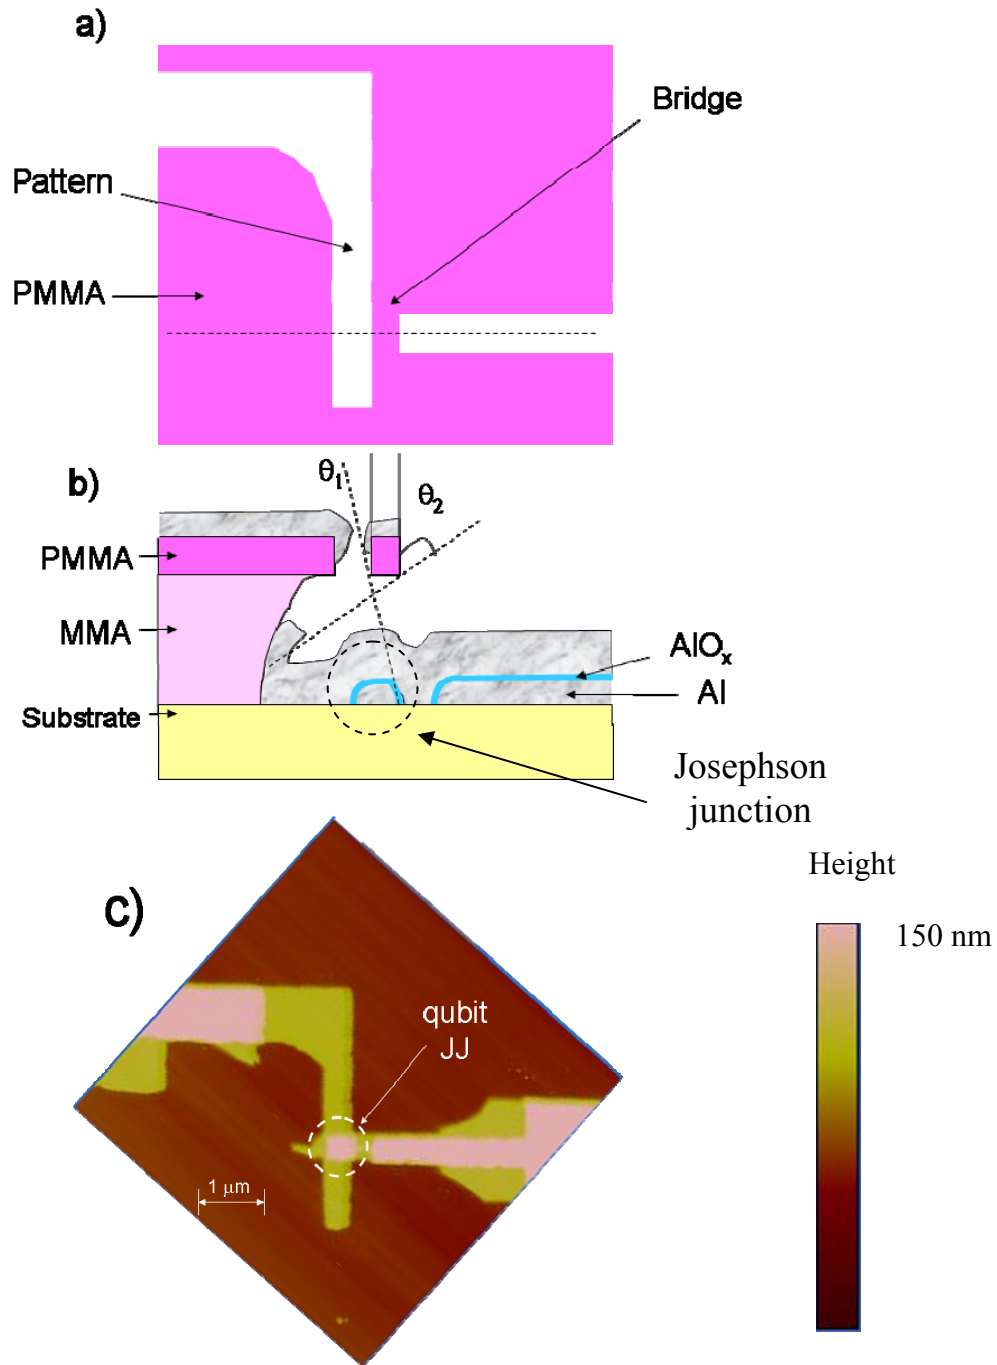


Fig. 6.5 (a) Diagram showing top view of the e-beam pattern. (b) Diagram of the double-angle evaporation with oxidization. (c) AFM image of the qubit junction after liftoff. Note that the overlap region creates a current flow path through the AlO_x between the two Al layers forming a Josephson Junction.

the angles. The tilt of the stage was controlled by a knob on the top of the lid that had a graduated scale taped to the base. This knob had some backlash in it, so I rotated the knob counter-clockwise until the stage was at an angle of 90° to the horizontal, facing out towards the middle of the room. I then began rotating the knob clockwise until the stage was tilted to the first angle of 5° . I checked the angle using a circular protractor and a thin metal ruler. I placed the ruler on the back of the stage so that it formed an extension of the line formed by the stage. The protractor was then oriented such that its x-axis was parallel with the top of the chamber. When the x-axis of the protractor was parallel to the floor and the ruler intersected the midpoint of the protractor, I read the angle of the stage. If the angle was too shallow, then I rotated it back to 90° and started over to avoid backlash. If it was too steep, then I nudged it forward until the angle read 5° . I also read the position of the knob according to the red “px” scale on the base and recorded it in the log book. I set the second angle to -58° and recorded the knob position in a similar fashion; always remembering to only rotate the stage in one direction. The angles needed to be set before I closed the lid because I would not be able to see the sample after the evaporation began.

I put two aluminum wire boats into the chamber attached to opposite electrodes. The sources consisted of two aluminum pieces in a wire basket, as described in Sect 6.6. I attached them to either electrodes 1 and 3 or 2 and 4, so that the stage faced one of the sources when it was tilted at either angle.

I closed the lid such that the “QC” marks lined up and pumped the system down to the 10^{-7} Torr range. The process was the same as described in Sect. 6.6 except that while I was roughing out the system I also purged the O_2 line. Once the pressure reached

about 10 Torr, I opened the O₂ gas bottle, opened the O₂ valve on the lid, and opened the electronic valve. The electronic valve was controlled by a knob on the front panel of an MKS Type 250 controller mounted on a rack to the left of the evaporator bench. I switched the knob to “Manual” and allowed the gas to flow into the chamber for 1 minute. Then I turned it off and roughed the system down to 400 mTorr before switching over to the cryopump. I left the O₂ hand valve open on the lid since the electronic valve was sufficient for closing this line off. It usually took about 2.5 hours from when I opened the hi-vac valve for the chamber to reach 7×10^{-7} Torr, but at that point, it was usually near the end of the day, so I usually let the system pump overnight.

When I returned in the morning, I degassed the ion gauge for about 1 minute. Within 10 minutes after the degas switch was turned off, the ion gauge gave me an accurate reading of the pressure, which was usually $2\text{--}3 \times 10^{-7}$ Torr.

I was then ready to evaporate the first layer of aluminum using the procedure outlined Sect. 6.6. I evaporated 40 nm of Al at a rate of 1.0–1.2 nm/s. I recorded the rate and thickness in the log book and on the log sheet for the first evaporation.

I began the oxidation process as quickly as possible after the first evaporation. First, I closed the hi-vac valve. Next, I opened the electronic valve for only a couple seconds and closed it again. At the same time I pushed the “zero” button on the crystal monitor controller to start the timer. The electronic valve was also fitted with a pressure gauge, and the pressure was displayed on the controller. I set the excitation voltage to 10V so that the units place on the display represented 100’s of mTorr. This display would shoot up to 13.8—which was a maximum value for the display—and rapidly decay to a more reasonable value ~10’s of mTorr. It took more than a minute, though, before

the pressure at the valve equalized with the pressure in the chamber. Nonetheless, I used this display to monitor the O₂ pressure because the convection gauge on the evaporator tended to be poorly calibrated. My target pressure was between 15 and 20 mTorr because I wanted a high critical current density. If the pressure displayed on the front panel looked like it would drop below this range, then I let in a little more gas. I recorded the pressure every minute in the log book. I found that oxidization at 20 mTorr for 10 minutes gave a critical current density of $\sim 100 \text{ A/cm}^2$ in an Al/AlO_x/Al Josephson junction (see Fig. 6.6). At 10 minutes I opened the hi-vac valve and cryopumped the O₂ out of the system. The chamber returned to base pressure in less than 10 minutes.

While the system was pumping back down I changed the angle of the stage. I rotated the knob on the top of the lid to the previously recorded position that represented the desired angle for the second evaporation. I needed to deposit more aluminum for the second layer than for the first layer; the second aluminum film needed to be thick enough that it was continuous over the step presented by the capacitor dielectric. The SiN_x for the filter capacitor was about 100 nm thick so I needed to deposit as much Al in the second layer as possible. Fig. 6.7 shows a diagram of the SiN_x step. Because of the difficulties of evaporating Al from the wire basket, it was difficult to get more than 100 nm. It was a tricky process to begin with, but if I kept the evaporation rate high (1.2-1.4 nm/s) then it was possible. Once the evaporation rate dropped below 0.6 nm/s I closed the shield and turned down the current source. I then turned off the source and made sure that all of the gas bottles are closed. After 10 minutes, I vented the chamber, turned off the cooling water, and removed the chip from the stage.




Fig. 6.6 Results of a study of the estimated critical current density J_c with respect to O_2 exposure $P \cdot t$. The blue diamonds are data obtained through room temperature measurements of the resistance of the junctions. In order to get J_c , I assumed that $J_c = V_{gap} / (R A)$ where $V_{gap} = 330 \mu\text{V}$, R was the room temperature resistance, and A was the junction area. The red square is the value obtained from measurements of the qubit junction on device PB9 (see Chapter 8).




Fig. 6.7 Illustration of the large step presented by the SiN_x dielectric. The second evaporated Al layer must be thick enough to climb over this step.

6.10 Liftoff in Acetone

The final step in the fabrication process was to lift off the resist in acetone in Rm 1314. I filled a dish with semiconductor grade acetone and placed the sample in it using tweezers. It took about 10 minutes until the sample was ready, but there was some danger in spraying the sample with acetone too early. When I sprayed the sample, the Al would sluff off the surface, but if there was some corner of the resist that had not dissolved yet, the film would not be washed clear of the surface. In that case it would be possible for the film to settle back onto the surface. If it landed back on the surface, then it would tend to stick and was hard to remove. To avoid this, I waited 15 minutes until the resist was completely dissolved. Then I held the sample with tweezers and sprayed it vigorously with acetone. Once the Al was removed, I rinsed the chip in IPA and dried it with nitrogen.

After the liftoff, I inspected the device in the Nikon Optiphot-2 optical microscope. First, I checked the large features at low magnification. I looked for bubbles in the aluminum film, large missing sections in the leads that might have come off during the liftoff, *etc.* Next, I made a careful check of the leads for either small breaks in the lines or places where poor liftoff might have caused shorts (*e.g.* between the spirals of the inductors). I checked the interdigitated capacitor to make sure that the aluminum was removed from between all the fingers. Then I checked the junctions.


Figure 6.8 shows a successfully formed isolation junction and a failed attempt. The separation between the overlap in the middle and the rest of the leads was usually clear at 100X magnification. The qubit junction was too small to know definitively whether it was well formed. Nevertheless, I could check to be sure that it was not

completely destroyed during the deposition. If there are no clear breaks in the line, then I left it alone. It was rarely the case that the isolation junction came out fine and the qubit junction had a problem.


Broken leads and malformed junctions were fatal problems with the device, but some minor problems could be fixed. For instance, if there was a place on the device where the aluminum did not lift off and it was causing a short, then this could be fixed by immersion in an ultrasound bath. Long narrow spaces between lines—such as the spaces between fingers in the capacitor or in the spirals in the inductor—were difficult to liftoff. If one of these was shorting the device, or if there was a flag from one of the leads that was not completely removed and had fallen over another lead to make a short, then I took the sample back to Rm 1314 and placed it in a 2” diameter, ½” deep Pyrex dish of acetone and let it soak for 5 minutes. This allowed any leftover resist to dissolve. I filled the Cole-Parmer Model 08849-00 Ultrasonic Cleaner with water up to ¼” below the lip of the metal basin. Holding the dish of acetone with a gloved hand in the water such that the lip of the dish was above the water level, I flicked the thumb switch on the ultrasound to apply 3-5 seconds of ultrasound and turned it off again. While the ultrasound was turned on, the sample would drift from the middle of the dish to the edge. Then I waited 5 more minutes and sprayed the chip with acetone. I did not know whether it was successful until I got back to Rm. 0219 and checked the sample under the microscope. If it did not remove the unwanted aluminum the first time, then it probably would not work, but if there was no other choice, then I applied ultrasound repeatedly until I got results. Although ultrasound has been a successful fix for me in the past, it can inadvertently remove smaller features from the pattern—especially the qubit junction—so it should

(a)

Success




(b)



(c)

Failure



(d)




Fig. 6.8 (a) Design and (b) photograph of a successfully fabricated Josephson junction and (c) and (d) show corresponding failure. The photographs were taken using an optical microscope at 100X magnification. Note that I can clearly see the separation between the overlapped region and the leads in the picture of the successful junction.

only be used when there is no other option. If the device looked good upon inspection with the optical microscope the fabrication process was complete.

6.11 Summary

Qubit fabrication began at LPS with deposition and patterning of aluminum back plates. Next, 100 nm low-loss SiN_x was deposited for the filter capacitor. Next, I patterned the SiN_x with an SF₆ etch. The wafer was then diced at LPS into 4.5 x 4.5 mm chips. These chips were cleaned with solvent to remove the lubricant used for dicing and then cleaned in an O₂ plasma to remove organic residues. I spun on two layers of MMA and baked each for 11 minutes at 150° C. I used 950 PMMA C2 as the pattern resist. It was necessary to deposit a 20 nm charge dissipation layer of aluminum because of the insulating nature of the sapphire substrate. I used a Phillips XL-35 SEM equipped with NPGS to write the pattern. The aluminum charge dissipation layer was removed using CD-26 and the resist developed. I completed the fabrication using double-angle evaporation of Al. The first layer was 40 nm thick, and the second was 100 nm thick. The oxidization in between the aluminum depositions was done at 20 mTorr for 10 minutes. I performed liftoff on the sample in acetone and only applied ultrasound if absolutely necessary.

Chapter 7

Experimental Setup and Measurement Procedures

This chapter describes the procedures I used to set up the apparatus and perform measurements for device PB9. The procedures are described in detail to be useful for those who may desire to duplicate them. After each procedure is described, I present a typical data set that was acquired from that measurement. An in-depth analysis of the data from device PB9 is presented in Chapters 8, 9, and 10.

7.1 Measuring the Room Temperature Resistance

When I fabricated a new device, it was always possible that the junctions shorted out or were open. Instead of cooling the device and finding out it was useless, I checked the resistance of the device at room temperature. The measurement required a micromanipulator probe station and a multimeter. Figure 7.1 shows our probe station in Rm SB0331. The micromanipulators were made by Alessi Industries and the probe tips were made of tungsten. The station was fitted with a Bausch & Lomb binocular microscope with a 7X-30X variable zoom lens for spotting the probe positions with respect to the sample. I used a Hewlett Packard 3478A Multimeter to measure the resistance of my samples. The circuit diagram is shown in Fig. 7.2. The measurement is fairly simple, but there are several places where a mistake in the procedure could destroy the sample. The multimeter was equipped to perform both 2 and 4 probe measurements; however, I had problems with electrostatic discharge (ESD) ruining my devices while performing a 4-probe measurement, so I do not recommend it [68]. The improved accuracy of a 4-probe measurement is not worth the risk. For a 2-probe measurement, I




Fig. 7.1 The micromanipulator probe station. The HI and LO ports on the multimeter are connected to the probes by alligator clips. I measured the resistance of the device while it was in the sample box.

HP 3478A Multimeter



Fig. 7.2 Circuit diagram of the room temperature resistance measurement. The HI and LO ports are shorted by a banana cable until the probes have been placed on the contact pads and the circuit completed.

attached the high and low ports on the multimeter to the micromanipulators using banana plugs and alligator clips (each micromanipulator had a stub on the arm of the probe where an alligator clip was attached to make electrical contact). It was critical that the multimeter be set to the correct resistance range so that it did not source too much current and destroy the Josephson junctions. The HP 3478A provided enough precision on the $1\text{M}\Omega$ range to measure the resistance to 10's of ohms, and when set to this range, the maximum current it would source was only 100 nA, which was safe for measuring junctions with critical current >10 nA. As an added precaution against ESD I used another banana cable to short out the high and low ports on the multimeter until I was ready to make the measurement.

The device was measured without removing it from its carrying case. The Gelpak (or other sample carrier) was placed under the microscope and I turned on the light. I zoomed all the way out and moved the sample box into a position where the device could be seen through the microscope. I focused on the surface of the sample and brought the probes into view. Then I adjusted the probes to their highest vertical position by turning the knobs on the top all the way to the right. I was careful that the probes could clear the edge of the sample box as I positioned them above the sample. At this point, looking through the microscope, I could only see the probes as vague shadows because they were so far from the focal plane. I had to be careful as I lowered the probes that they did not land on the surface of the chip with too much force. Because micromanipulators can have significant backlash in the z-motion, the probes often continued to travel downward after I stopped turning the knobs. The goal was to get the probes into focus just above the sample, without actually making contact.

With the high and low ports shorted out at the multimeter, it was now safe to put the probes down on the contact pads. I positioned the low probe over the contact pad on one side of the device and slowly brought the tip into contact with the pad. I knew that the probe was in contact with the pad when I saw the tip “skate” forward on the surface. Next, the high probe was brought into contact with the contact pad on the other side of the device. I read the resistance R_n off of the multimeter and recorded it in my lab notebook.

From the measured value for R_n , I estimated the critical current I_0 using:

$$I_0 \approx V_{gap} / R_n \quad (7.1)$$

where $V_{gap} = 2\Delta/e = 360$ eV is the superconducting gap voltage of aluminum. I note that Eq 7.1 is only an approximation because there will inevitably be some resistance contributed by the lines of the device, whereas Eq. 7.1 assumes all the resistance is from the junction tunnel barriers. From measurements on other devices, I found the resistivity of the 100 nm thick Al thin film was about $0.2\Omega/\square$.

My dc SQUID phase qubits had one large junction, so the resistance was quite small (tens of ohms); therefore, this measurement only served to tell whether the device was open or not. However, if the resistance ended up being several k Ω 's when I was expecting a device with a critical current of 10 μ A, then that would also tell me that the device was defective. By checking the resistance of the device at room temperature, I saved much time in the long run because it eliminated many of the devices that were broken.

7.2 Dilution Refrigerator Wiring and Sample Box

My experiment used a Kelvinox 200 dilution refrigerator to cool the qubits. Figure 7.3 shows a picture of the refrigerator. This unit had a base temperature of 15-20 mK. It was equipped with 2 identical sets of 4 lines for controlling qubits. For my experiments, I only used one set. Figure 7.5 shows a diagram of the wiring in the refrigerator. Each of the four lines was configured for a specific purpose. Each of the low-frequency lines (current, flux, and voltage) were equipped with a switch at the top of the refrigerator for grounding the input. This allowed me to make changes to the setup of the room temperature electronics without fear of destroying the device through ESD.

The current bias line (I_b) was constructed from several different types of cables, filters, and other components (see Fig. 7.5). The grounding switch is a SPDT switch mounted in a bud-box. This allowed me to ground the input lines or connect them to the sample. Coming out of the bud-box, the signal was carried by BNC cable to a BNC-to-SMA adapter that was sealed by soft solder. This adapter connected the room temperature BNC cable to a cable running into the vacuum can. From there to the 4 K stage, the cable was one wire of a twisted pair of Manganin wire (about 2 m long) [26]. The Manganin wire was effectively connected to ground by $2\text{ M}\Omega$, which helped it to thermalize from 300 K to 4K [49]. That wire went into a Cu patch box, which also acted as a thermal anchor for the wire [see in Fig. 7.4(a)] [26]. From there the current bias signal was carried via about 60 cm of 1 Nc Ac 0.5 Thermocoax to filters mounted on the mixing chamber (MXC) [52]. Before entering the filters, it was joined with the voltage line, which is described in the next paragraph. After the tee, the signal was fed to an LC-filter that had a cutoff of about 10 MHz [see Fig. 7.4(b)]. Note also that Thermocoax will




Fig. 7.3 (a) Picture of the dilution refrigerator [37]. (b) Detailed view of the 20 mK stage [37].

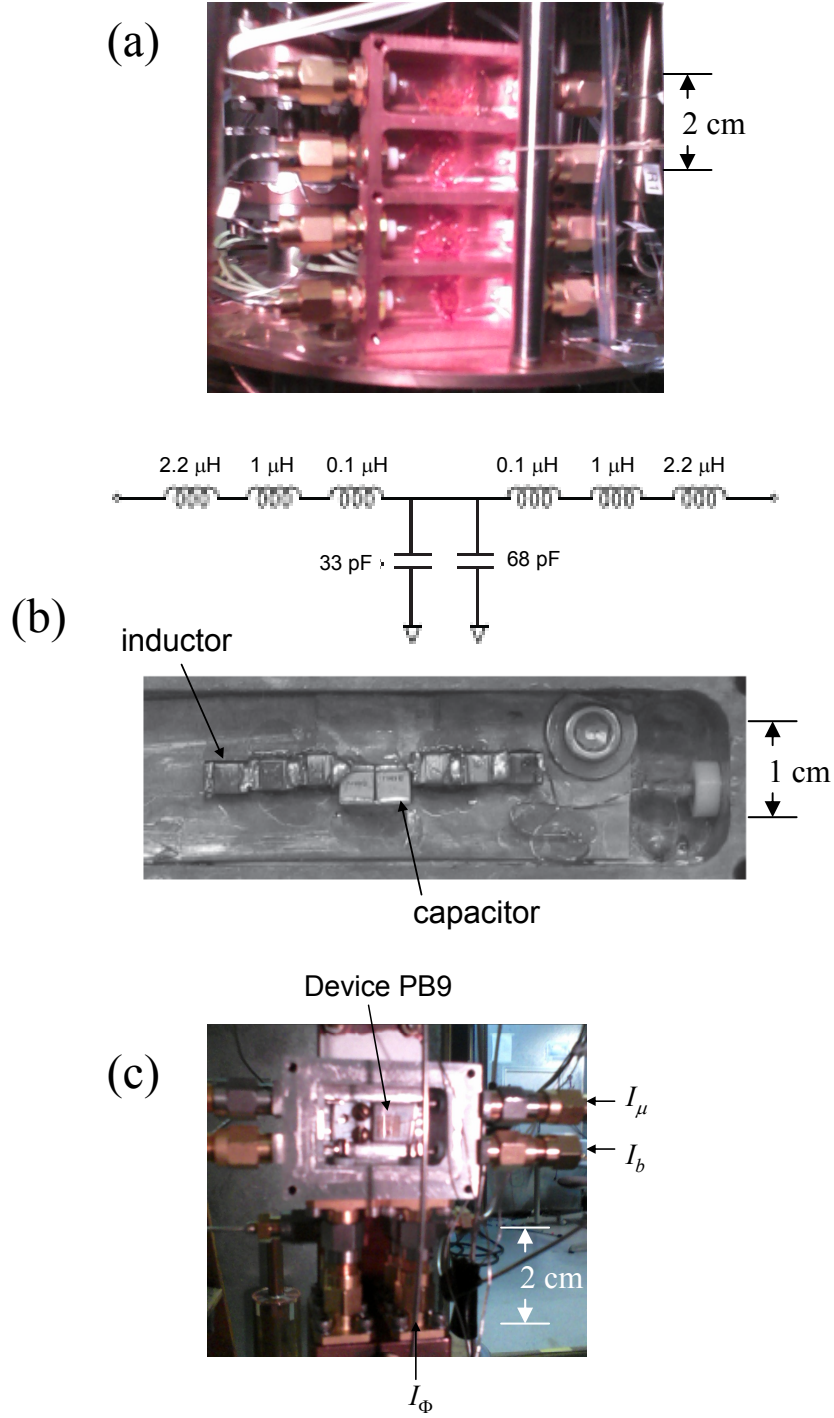


Fig. 7.4 (a) Picture of the Cu patch box. The center conductors of the line are insulated Cu wires that are GE varnished to the body of the box. (b) Diagram and picture of the LC filter on the current bias and flux bias lines [37]. (c) Picture of the Al sample box [37].

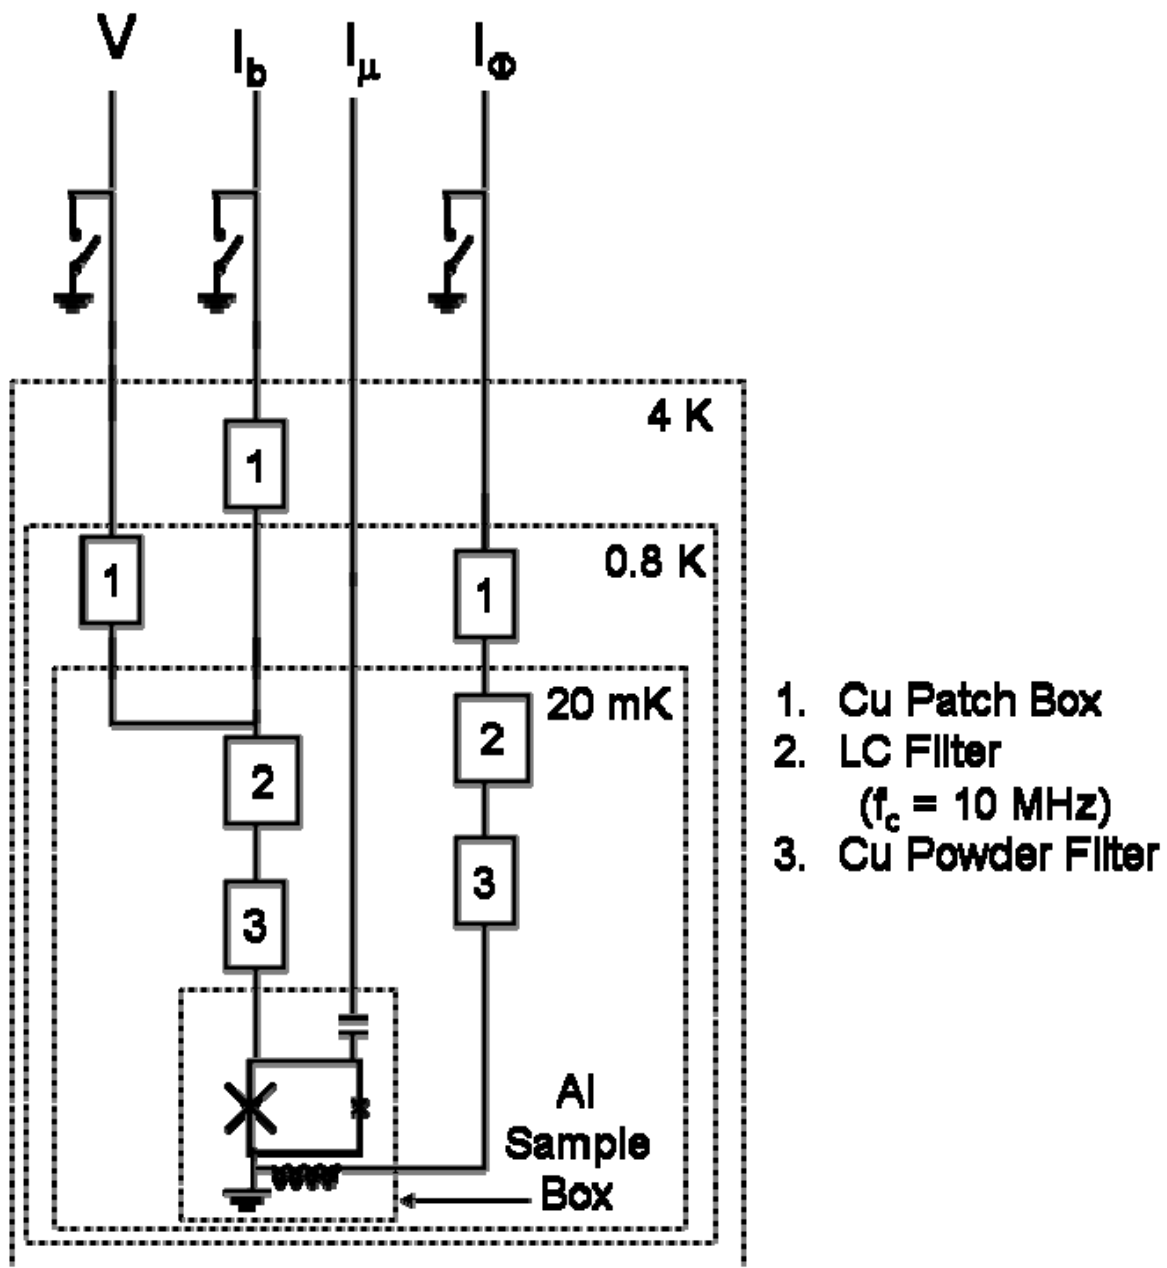


Fig. 7.5 Diagram of the qubit control lines in the dilution refrigerator. The types of cables used in the wiring are detailed in the text.

filter noise from a wide range of frequencies when used to connect different cold stages [69]. The LC-filter was connected directly to a Cu-powder filter [38]. Both filters were built in Cu boxes that were clamped to the mixing chamber [41]. The output of the Cu-powder filter (CPF) was connected by a short (~ 5 cm) length of UT85 Cu/Cu coax to an SMA input to the Al Sample Box.

In order to measure the voltage across the device, without measuring any voltage drop along I_b , the refrigerator was equipped with a voltage sensing line (V). Starting from the bottom of the refrigerator, the voltage line was connected to the current bias line by an SMA tee. From there, the V line connects to a Cu patch box for thermal anchoring at the Still ($T = 800$ mK) via 60 cm of 1 Nc Ac 0.5 Thermocoax. It is connected to room temperature through 2 m of Lakeshore CC-SR-10 coaxial cable, and leaves the vacuum space through a vacuum SMA feedthrough, similar to the one used on the I_b line. The voltage line was also equipped with a safety grounding switch.

The flux line (I_ϕ) had a grounding box and vacuum feed-through that was similar to the one used on the current bias line. The signal traveled from 300 K to the Still stage via 2 m of Lakeshore CC-SR-10 cable, where it was thermally anchored by a Cu patch box [52]. From there, the flux line was constructed from homemade superconducting Nb coax; Nb was used because the flux line needed to carry a relatively large current without dissipating much heat [26]. The Nb coax fed into a pair of LC and Cu Powder filters (CPF) in series, similar to what was used for the I_b line. The output of the CPF was connected to the sample box via a Male-Male SMA connector.

The microwave line (I_μ) was different from the three other lines. Because I_μ would be carrying high frequency signals to the qubit, it did not have filtering. In order

to preserve the integrity of the signal (*i.e.* prevent reflections, dispersion, *etc.*), the cable was uninterrupted between 300 K and the sample box. At 300 K, the microwave line came into the shielded room through an SMA bulkhead feed-through and was carried by an SMA coaxial cable to the top of the refrigerator. It entered the vacuum can through a sealed SMA vacuum feed-through. From there, it travelled directly to the Al sample box via about 2 m of UT-34-SS-SS coax. This cable should do a good job of isolating the MXC from heat at 300 K because the jacket is stainless steel and so is the center conductor. However, in the future, we plan on adding cold attenuators at the various stages of the refrigerator to better thermalize the lines and reduce noise coming from sections that are hotter than 20 mK.

The sample box is shown in Fig. 7.4(c). It was made of 6033 (high purity) aluminum. It had 6 SMA inputs that ended in flattened pins for wirebonding. The sample sat on an aluminum stage that was screwed into the box. I attached my sample to the stage with GE varnish. The connections were made from the SMA pins to the bonding pads on the sample with aluminum wirebonds. The box had an aluminum lid that I closed tight with 4 brass screws. I attached it to a copper auxiliary stage with 6 brass screws, and the auxiliary stage was clamped to the mixing chamber by a circular copper clamp [41].

7.3 Measuring the I - V Characteristic

After I cooled a device, one of the first measurements I made was of the current-voltage characteristic curve (I - V). A good I - V curve was the first assurance I had that the device survived the cool down and had a reasonable critical current.

I measured the I-V by sending a sine wave voltage through a large resistor and measuring the voltage on an oscilloscope. Figure 7.6 shows the diagram of the setup. The voltage source was an Agilent 33120A 15 MHz Function / Arbitrary Waveform Generator. As was done during all equipment changes and instrument setups, the device was grounded by the grounding switches at the top of the refrigerator. I set the function generator to output a continuous sine wave with minimum amplitude at a frequency of a few Hz and set the output impedance to “High Z.” This signal went to a π -filter with a 20 kHz cutoff in the screen room wall via a BNC cable [41]. The signal was carried from there by BNC cable to a 1X buffer amplifier [41]. After the buffer, the signal went through a 200 k Ω resistor. This resistor effectively converts the voltage source into a current source such that 200 mV corresponds to 1 μ A of current. The signal went through the grounding switch and into the device via the current line, which was described in the previous section. In order to measure the current, I measured the voltage at the input of the bias resistor using channel 1 (Ch1) of a Tektronix 3054B Digital Phosphor Oscilloscope. I took care to set the input impedance of Ch1 to “1 M Ω ” so that the scope did not draw much current. For the I-V, I also needed to measure the voltage across the device. To do this, I connected the voltage line to the input of a Stanford Research Systems SR560 Low Noise Preamplifier. I plugged the voltage signal into port A on the front panel, set the coupling to “DC,” the source to “A,” the gain mode to “Low Noise,” and the gain to “1000,” and then connected the “50 Ω ” output to Ch2 on the scope.

Next, I set up the Tektronix 3054B oscilloscope to work as the data taking apparatus. The oscilloscope ran off of battery power, which reduced ground loops and 60 Hz noise. The scope also had a 3.5” floppy disk drive for data transfer to a PC. I set both

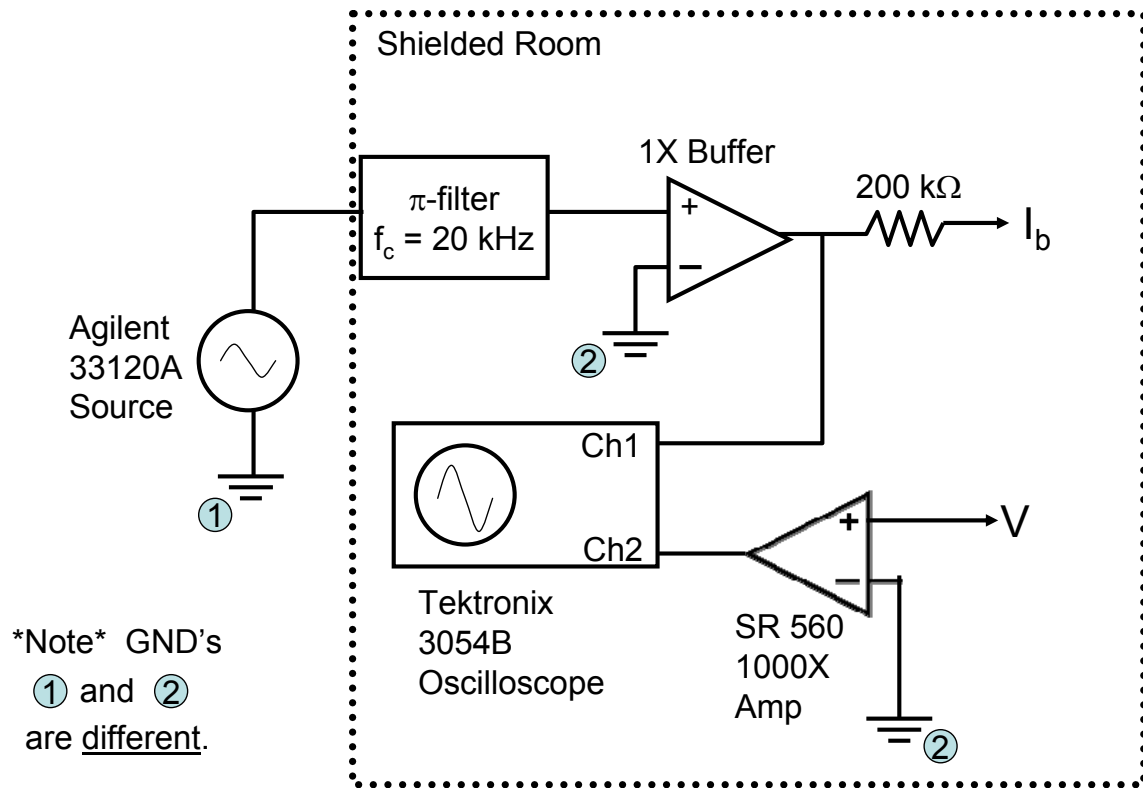


Fig. 7.6 Diagram of the I-V measurement setup. The electronics in the screen room and the device are grounded to the refrigerator. The source outside the screen room is run off of an isolation transformer and grounded to the screen room wall. The 1X Buffer on the current line breaks the ground loop between the refrigerator top plate and the screen room wall.

Ch1 and Ch2 to be active on the display. I set the scope to trigger off of Ch1 and adjusted the trigger level until the 100 mV sine wave was stationary on the screen (Ch2 was flat at 0V). I adjusted the “Horizontal” settings such that only one or two periods of the sine wave were visible. In the “Display” menu, I set the display to “XY Display.” At this point the oscilloscope was prepared to display the I-V curve.

I next ungrounded the current bias and voltage lines at the top of the refrigerator and began turning up the voltage on the HP 33120A. I turned up the voltage amplitude steadily until I saw the device switch to the voltage state on the oscilloscope, then I turned up the voltage amplitude further to see a bit of the resistive branch of the I-V. If I did not see it switch at a reasonable current, then I checked the instruments and connectors to make sure things were set up correctly. If the setup was okay, then I increased the current to ten times the expected critical current, and if scope still did not show the device switching to the gap voltage, then I concluded that the device was shorted out.

To acquire data, I went to the “Acquire” menu and made sure that the scope was set to take 10,000 samples. Next, I went to the “Save/Recall” menu, selected “Save Ch1 Waveform,” and pushed the softkey next to the image of a disk that said “To File.” This brought up the disk menu. If the scope asked to format the disk, then I removed the 3.5” disk, backed up all the data it to my PC, replaced the disk in the scope, and then clicked “OK” to format the disk. To save the I-V, I pushed the softkey that said, “Save Active Waveforms,” and the scope saved each waveform as a vector of 10,000 points that could be plotted in Excel. The first column of the data was the current and the second was the voltage.

Figure 7.7 shows an I-V curve captured with this technique and plotted in Microsoft Excel. This I-V characteristic was measured on SQUID SF7, which I fabricated in 2007 and was cooled to 80 mK on the small refrigerator. The critical current of the device was $9.7 \mu\text{A}$. When measuring the I-V of a SQUID, one actually measures the switching current of the entire device, which depends on the flux, and has a maximum that is the sum of the critical currents of the two junctions. The gap voltage was $V_{\text{gap}} \approx 360 \mu\text{V}$ as expected for an aluminum device. Note that the supercurrent branch is vertical and there is hysteresis, as expected for a SQUID that does not have resistive shunts. The inverse of the slope of the resistive branch gives the tunneling resistance R_n of the two junctions in parallel, since everything else is superconducting. With $R_n \approx 17\Omega$, Eq. 7.1 yields an approximate critical current of about $20 \mu\text{A}$, so we see that Eq. 7.1 only gives a ballpark estimate of the critical current.

The I-V curve gave me an idea of the quality of the device. If the critical current, the gap voltage, and the general shape of the I-V curve looked good, then I knew that I could move forward with that device.

7.4 Measuring the Switching Histogram

The next step in characterizing a device was to set up the switching histogram measurement. The switching histogram was the backbone of all of my qubit measurements. For the measurement, I ramped the current through the device and used a time interval counter to measure the time between when the bias current ramp began and when the device switched to the voltage state. I used several pieces of homebrew electronics and LabVIEW code in this measurement that had been developed by previous




Fig. 7.7 Current-voltage characteristic (I-V curve) of SQUID SF7. The red arrows mark the direction of the current value in time. The current increases from zero while the voltage stays at zero. After the current reaches $9.7 \mu\text{A}$, the voltage across the device “switches” to the gap voltage $2\Delta/e$. As the current increases further, the device acts like a resistor. When the current is reduced, the voltage follows the quasi-particle branch. Finally, the device re-traps to the supercurrent state and repeats the pattern on the negative current side of the I-V.

members of our research group [38]. The switching histogram allowed me to measure the time of the switching of the SQUID with precision.

Figure 7.8 shows a diagram of the switching histogram setup. I triggered all the instruments off of an Exact Model 628 20 MHz Pulse/Function Generator clock. The clock was set to generate a square wave pulse with amplitude $7.5 V_{\text{pk-pk}}$, 0 V offset, 0.8 ms length, at a frequency of 700 Hz. This output was sent to the trigger of the Agilent 33120A function generator. The sync of the 33120A was then distributed to the triggers of all other instruments—in more complicated experiments there could be more than 5 other instruments running off of this sync.

For the histogram measurements, I configured the output of the Agilent 33120 such that it output a current ramp at a constant rate every time it was triggered by the clock. I began by constructing the arbitrary waveform for the ramp. Our group had written a LabVIEW virtual instrument (VI) for this purpose named *HP 33120A Waveform Editor.vi* [38]. Figure 7.9 shows the front panel and the bias waveform. The HP 33120A could be programmed with 12 vertical bits and 16000 time values, so I used the VI to use all of the time values to draw a straight line from point (0,0) to (15999, 2047) on the DAC. This waveform was sent via GPIB to the source and saved in the source under the “volatile” arbitrary waveform. On the source, I used the keys on the front panel to set the output to “Arb” and select the “Volatile” waveform. Then I set the source to “Burst” mode so that it ran the ramp every time the source received a trigger. I set the frequency to something faster than that of the clock (*e.g.* 1 kHz). I set the amplitude of the ramp to just above the critical current, which had been determined from the I-V. Finally, it was important to put about -20 mV of offset on the source so that the

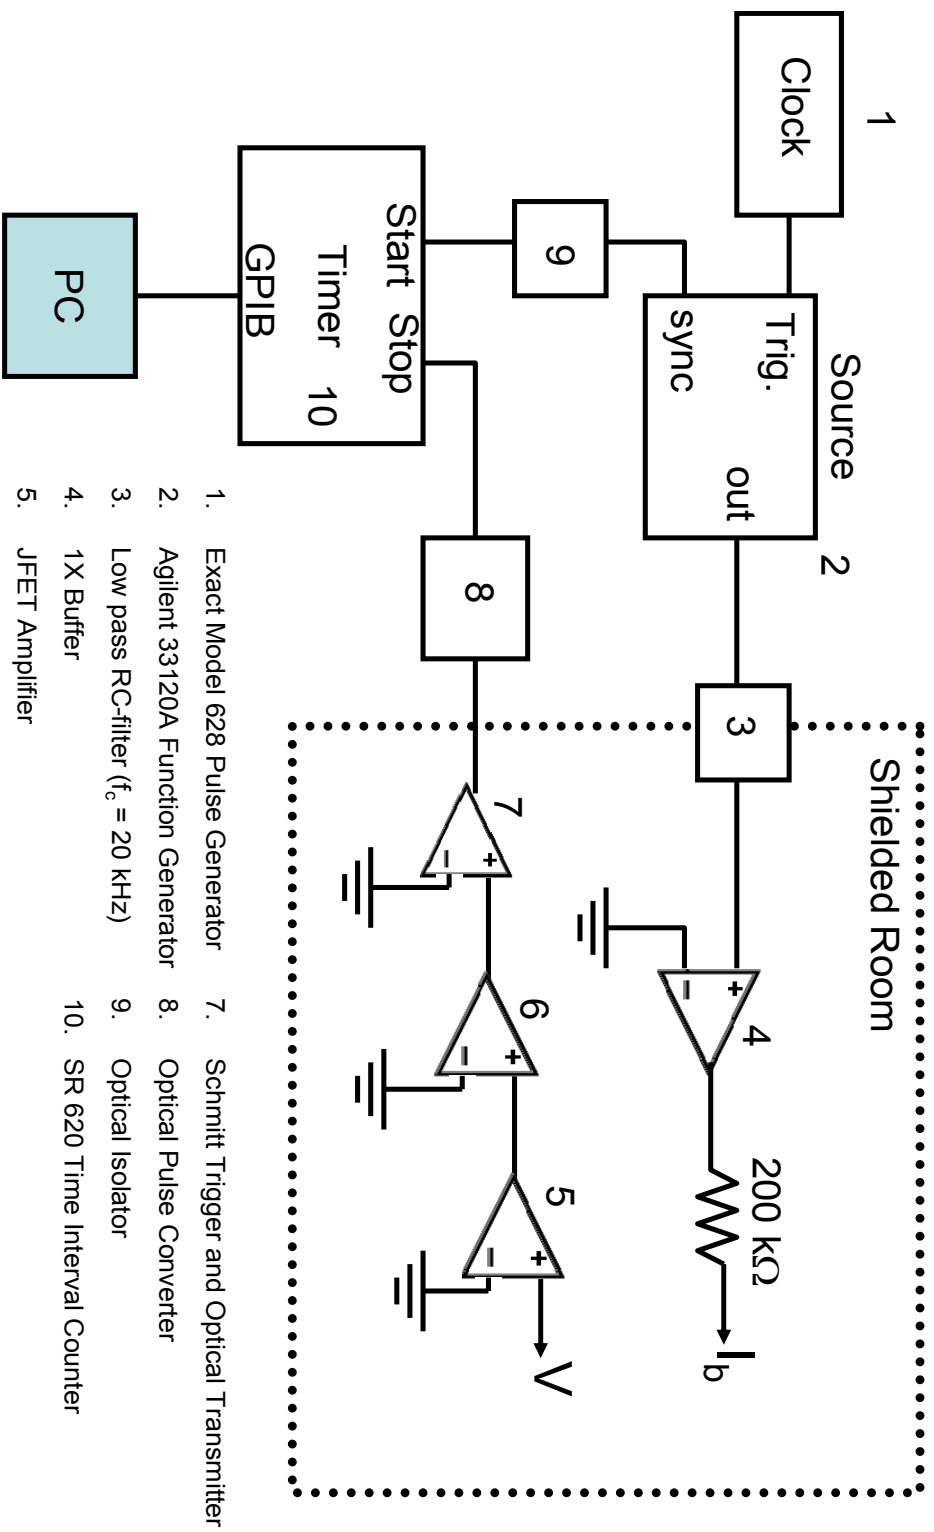


Fig. 7.8 Diagram of the setup for measuring a switching histogram.

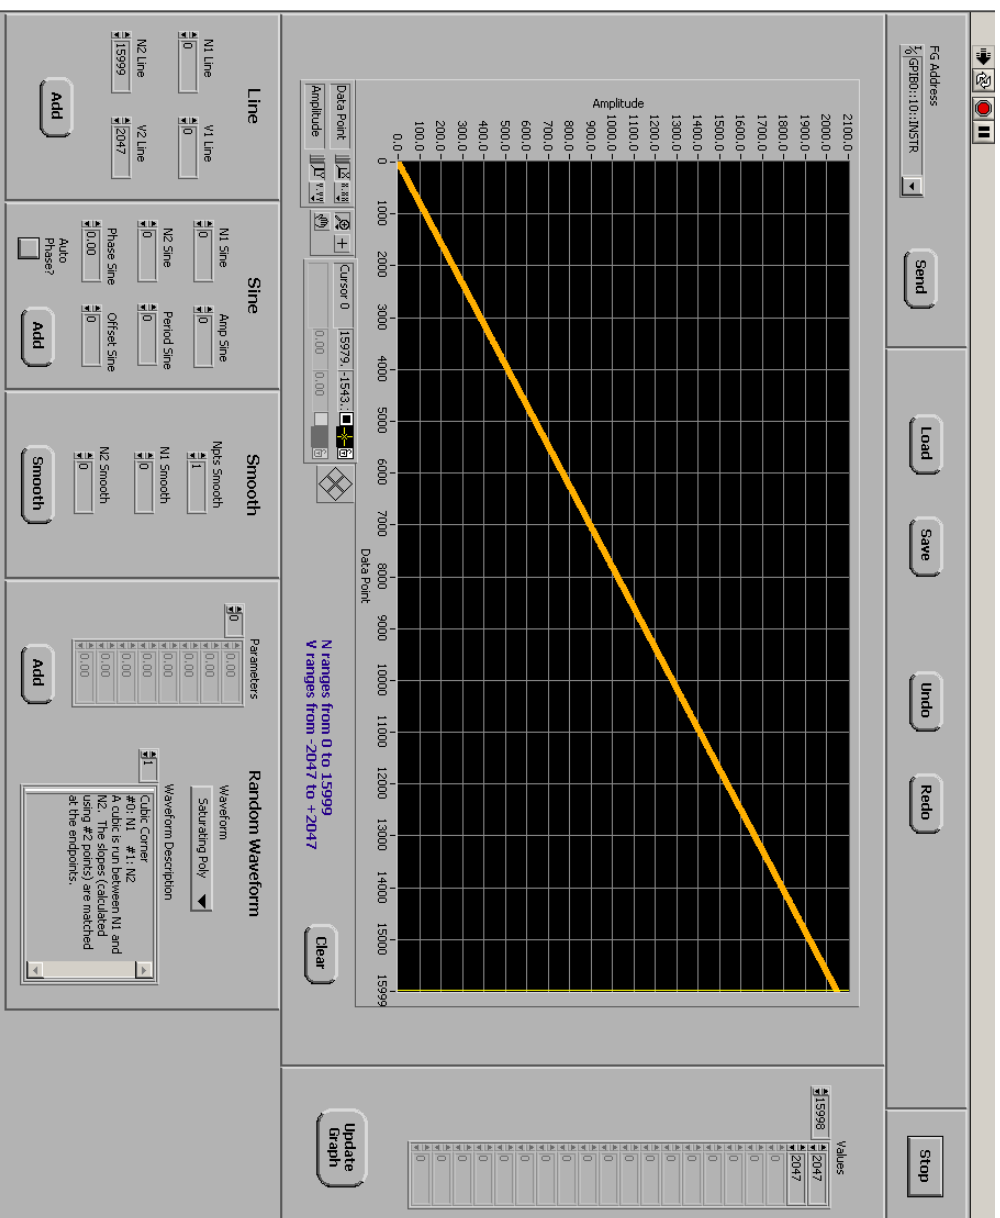


Fig. 7.9 Picture of the front panel of the LabVIEW program *HP 33120A Waveform Editor.vi*. This waveform was used for measuring the $I-\Phi$ as well as the initial switching time histogram.

junctions were sure to reset to the zero-voltage state after each ramp. In order to be sure that the source was outputting the correct waveform, I viewed the current ramp on the scope. The correct waveform looked like a saw-tooth pattern with some dead time between each ramp. This ramp switched the junctions once every cycle and reset them in preparation for the next trigger from the clock.

In order to determine the switching current of the device, I needed a precise measure of the time that passed from the beginning of the ramp to the time of the switch. I used a Stanford Research Systems Model SR620 Universal Time Interval Counter to measure this interval. When the ramp began, the Agilent 33120A generated a high signal from the “Sync” output at the front of the source. This sync signal went through an optical isolator, which sharpened the rise time, and then into the “Start” input on the SR620. The input was set to be ac coupled and set to trigger on the positive slope. The sync output pulse began the counting of the time interval.

Setting up the trigger level on the SR620 took some time, but it was important that I set the level correctly if I wanted an accurate measure of the timing of the switch. First, I took the signal coming out of the optical isolator and fed it into the oscilloscope. Then I adjusted the trigger of the scope so that I could look at the rise of the signal. Shortening the time window of the scope revealed the shape of the switching voltage during the rise. I wanted to find a level for the trigger where the signal was sufficiently steep, but not too far from the beginning of the rise. Also, I had to be careful that the level was high enough that I did not get any erroneous triggers from the signal ringing when it came back down at the end of the cycle. Finally, I reconnected the line to the “Start” port on the SR620 and set the trigger to the value determined using the scope.

Next, I set up the electronics inside the screen room to detect the switching voltage. The differential amp that I used for the I-V curve was too slow and noisy for monitoring the switching for a histogram, so I used several stages of homemade amplifiers to detect the signal. The signal came out of the voltage line grounding switch and went into a JFET amplifier [26]. From there, it went through another ultra-low noise amplifier (ULNA) and finally into the Schmitt Trigger [26]. I should note that all these amplifiers ran off of 24 volt battery boxes built out of Panasonic LC-R127R2P1 12V, 7.2Ah/20HR lead-acid batteries. The JFET and ULNA differentiated the dc voltage step generated by the switching event into a quick voltage pulse. The Schmitt Trigger employed an operational amplifier that sent a high signal when it detected a voltage above a threshold. The signal from the Schmitt trigger was used to trigger an optical pulse generator that sent a pulse out of the shielded room.

The threshold of the Schmitt trigger had to be calibrated so that I could accurately detect the voltage signal. Many of the Schmitt triggers in the lab have a “Monitor” port on them where the op-amp output has been tapped for this calibration. I took the output from the “Monitor” port and sent it into Ch2 on the scope. I teed the input to the current bias resistor and viewed the current ramp on Ch1. The goal was to set the threshold such that the Schmitt trigger output went high at a point in the current ramp that was consistent with the switching current determined in the I-V. I set the threshold by adjusting the resistance of a potentiometer in the Schmitt trigger circuit (see Fig. 7.8). Some of our Schmitt triggers used potentiometers with dials that are located on the outside of the bud box, but others used potentiometers that could only be accessed by removing the lid of the bud box. These potentiometers looked like small, blue, rectangular boxes in the

circuitry. They could be adjusted by turning a small screw on their tops. If the threshold was set too low, then the Schmitt trigger would always source a high signal. If it was too high, then the Schmitt trigger would never trigger. When I set the potentiometer correctly, the Schmitt trigger output would switch from low to high during the current ramp.

To reduce problems with ground loops and noise, the Schmitt trigger sent an optical pulse out of the screen room via an optical fiber. This was converted to a voltage pulse by an optical receiver [26]. That voltage pulse was sent to the “Stop” input of the SR620 interval counter via a BNC cable. I configured this input and its trigger level using a similar procedure to that which I used for the “Start” port.

With all of the components now in place and configured to measure the switching time histogram, I used another VI program called *OneSR.vi* to perform the measurement. The front panel of *OneSR* is depicted in Fig. 7.10. Basically, this program read the time interval via GPIB cable from the SR620 and displayed the histogram of the switching times. I rarely used *OneSR* alone to capture a histogram. I usually did it automatically by using a program called *Run Scan.vi*, which called *OneSR*; however, it can be done manually. Regardless, the VI needed to be set up in order to correctly display the histogram. First, I opened *OneSR.vi* and ran the program by clicking the arrow button at the top left of the LabVIEW screen. I set the “Timer GPIB” to the address of the SR620. I turned on the “Histogram Display” so that the switching counts would begin to appear on the upper right graph. Depending on how finely I wanted to know the switching time, I set the “binning” (*i.e.* number of seconds per bin) to between 10^{-6} for very low resolution and 10^{-10} for extremely fine resolution. For the sake of

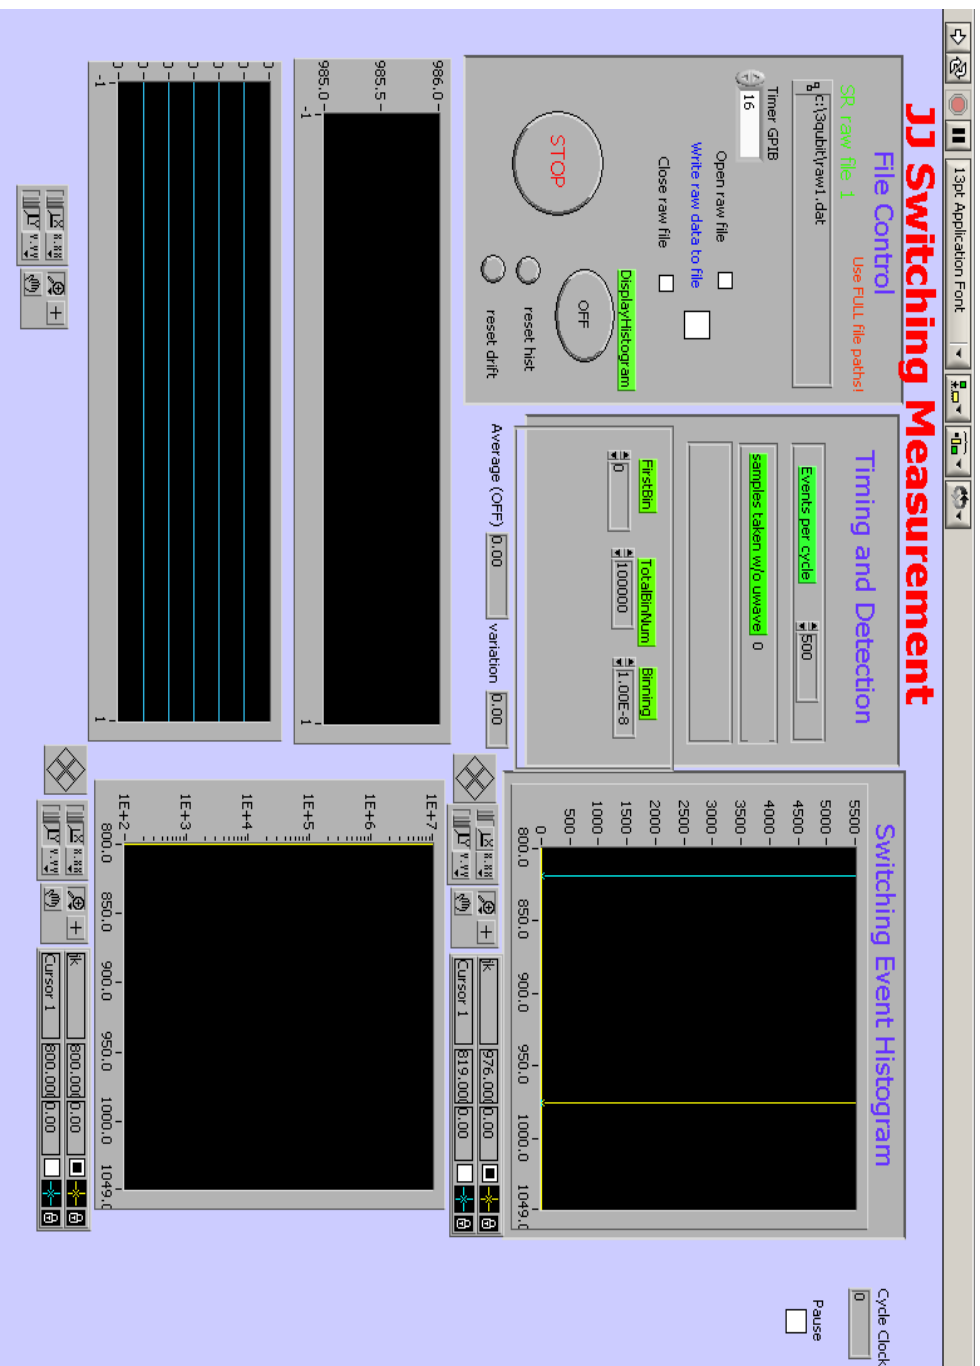


Fig. 7.10 Picture of the front panel of *OneSR.vi*. This program was used to display and/or capture the switching histogram. It also displays the average switching time in the lower left corner and the escape rate in the lower right.

capturing a switching histogram 10^{-8} was adequate. The “Events per cycle” control established the number of samples that the program acquired before updating the graph, and I set that to 500 to keep the information on the histogram current while I set up the rest of the measurement. I set the “First Bin” to 0 and the “Total Bin Number” to 100000 which made the data acquisition time window 1 ms. I clicked the “Reset Histogram” button in order to get a fresh look at the histogram, and I usually saw that the switching counts were grouped around some bin value. At this point, *OneSR* was correctly acquiring and displaying the histogram.

As I noted previously, I did not usually capture a single histogram, but sometimes I needed to adjust all of the settings for a measurement manually, which meant that I needed to capture the histogram according to my own timing. In order to do this I first set the file path in the control marked “SR raw file 1” in green type. Then I checked the “Open Raw File” box and the “Write Raw File” box. Next, I waited for a number of samples to be acquired. 10000 samples was usually enough to capture the character of the histogram, but if I only wanted to know the average switching bin, then I might only acquire 2000 samples. To stop the program from acquiring samples I clicked the “Close Raw File” box. This saved the histogram in a data file according to the path specified.

Figure 7.11 shows a typical switching histogram. There is always some width to the histogram peak because of noise and the random nature of tunneling events. The bin number n can be converted into time using:

$$t = \tau_n n \quad (7.2)$$

where τ_n is the “binning” setting that I chose to set the bin width. The histogram let me know the switching time which was the backbone of the rest of the measurements.





Fig. 7.11 Example of switching histogram from device PB9. n is the bin number.

7.5 Measuring the I - Φ Characteristics

The I-V curve lets me know that there are Josephson junctions in the device, but the information that it gives me on the device is limited. For instance, the I-V gives no clue as to whether there are one or two junctions in the device or their individual critical currents. In contrast, the current-flux (I - Φ) characteristic gives me a look at whether the two junctions are alive as well as giving me a measure of their individual critical currents, the inductances of each arm of the SQUID, and the mutual inductance between the flux bias line and the SQUID. To construct the I - Φ , I measured the switching current of the SQUID for a range of fluxes applied to the SQUID.

Capturing the I - Φ employs the switching histogram described in the previous section. Fig. 7.12 shows a diagram of the I - Φ measurement setup. The setup is similar to Fig. 7.8 for the I-V except that there is a second Agilent 33120A that sources the current for the dc flux. I set the flux source to output a dc voltage by pressing and holding the “Offset” button on the front panel. I used a bias resistor between 1 and 10 k Ω for this line, but the selection of the actual value depends on the mutual inductance, the range of flux biases that need to be attained, and the resolution that one desires for the I - Φ .

It is possible to capture an I - Φ using *OneSR.vi*, but it would be tedious and time consuming. In order to automate the measurement, I used a LabVIEW VI called “*Run Scan.vi*” [38]. Run Scan allowed me to provide an array of parameters for a number of different instruments and then capture the switching time data for each set of parameters. The front panel for *OneSR.vi* is depicted in Fig. 7.13. First, I configured Run Scan to control *OneSR* by typing the path to *OneSR.vi* into the “Data Taking VI” control. Next, I set the path for the folder where I would save the switching data in the “File Prefix”



201

control. I clicked the button for “Add Params” so it displayed “Yes”. Since I was only changing the flux bias, I only needed to send one command to one instrument; therefore, I only put a “0” into the “Enable” array. In the first row of the “Label” array, I put a “V”. The GPIB address of the flux source was set using its front panel, and this address was selected in the “Scan Device” array. The proper command prefix for applying a dc voltage was “APPL:DC DEF, DEF,” and the suffix was set to “MV” for millivolts. I set the value of the first row in the “Multiplier” array to “1.” “DOP” stood for “digits of precision” and since the voltage had been specified in millivolts, there was no need to have any digits after the decimal point; therefore, I put “0” in the first row of the “DOP” array. The “Values” array contained the voltage settings for the $I-\Phi$. I created a scan of the flux currents by using the “Fill” tab to the right of the “Values” array. I specified the “Start value” of the scan, set the step size in the “Increment” control, and set the “Number” of steps such that the scan spanned the correct range of flux biases. For example, a scan from -2V to +2V with a 10 mV step would be constructed by the parameters: Start value = -2000, Increment = 10, Number = 401. This set all the commands that would be sent to the flux source during the scan.

I had to decide how many measurements I should make of the switching time for each flux bias. I needed to make enough measurements such that I could determine the switching current with good enough precision. Also, I generally wanted to get an idea of the character of the histogram peak and escape rate at each bias. Usually, 10^4 samples were enough to accomplish these goals. Run Scan tended to get hung up if I tried to take all 10^4 samples at once, so I broke the 10^4 samples into 5 cycles of 2000. I did this by

setting the “Events per cycle” in *OneSR.vi* to 2000 and the “Cycles / Setting” in *Run Scan.vi* to 5.

There were a few last things that needed to be done before the scan was ready to run. I created a folder with the same path as I put into the “File Prefix” control in *Run Scan*. I turned off the histogram display in *OneSR* because it took up time, which made the scan take longer than necessary. I also reset the drift, which displayed the average switching bin and would give me a rough idea of how the switching current was varying with flux. Finally, when everything was prepared, I pushed the “Go” button in *Run Scan*. The button latched and read “Gone”, *Run Scan* set the flux source to the first bias setting, and *OneSR* started capturing the switching data.

The approximate length of time it would take to finish a scan was:

$$t_{scan} = \left(\frac{1}{f_{clock}} \frac{\#Events}{Cycle} \frac{\#Cycles}{Setting} + 8s \right) \#Settings \quad (7.3)$$

where the 8 s in the parentheses comes from the fact that it takes the computer 8 s to set the output parameters of the instruments. The other term in the parentheses is the amount of time it takes to actually capture the data.

Figure 7.14 shows an $I-\Phi$ from device DS6 captured using this method. The y-axis is the switching current of the device which was calculated from the switching time histogram using:

$$I_{sw} = t_{sw} f_{I_b} \frac{A}{R_{I_b}} \quad (7.4)$$

where t_{sw} is the time of the switching (the raw data), f_{clock} is the frequency of the clock trigger signal, A is the amplitude setting of the Agilent 33120A used for the bias current ramp, and R_{I_b} is the resistor on the current bias line. Fig. 7.14 is a false color plot, where

the color represents the number of counts in the switching histogram. On this scale red represents the highest number of switches at a certain time and blue represents 1 switch. The area where there was 0 switches was set to white. I took a histogram for each flux bias voltage setting, so Fig. 7.14 is made up of all of the histograms collected at the different flux settings.

The data displays multiple trapped flux states and periodicity of the switching current with respect to the applied flux. If for some reason only one junction survived and the other junction got blown up by ESD or became open for some other reason, then I would not see any periodicity in the $I-\Phi$. It is possible that I might see some modulation of the switching current of the device at very high fields, but not with the lesser strength of flux that I applied with the flux bias line [28].

7.6 Qubit State Readout Using Current Pulse

To proceed with characterizing the device, I needed to measure the state of the qubit. I did this by sending a quick ($\sim 2\text{ns}$) current pulse to the SQUID and checking to see whether the device switched to the voltage states. If I chose the current pulse amplitude correctly, the current pulse would reduce the qubit tunneling barrier to a level where it was highly probable that the qubit would tunnel out if the qubit was in the excited state, but if it was in its ground state, then it was highly probable that the SQUID would remain in the zero-voltage state.

Figure 7.15 shows the diagram of the pulsed readout setup. I used a Stanford Research System Model DG535 Four Channel Digital Delay/Pulse Generator to generate the current pulse. I sent the pulses to the qubit through the microwave line (described in


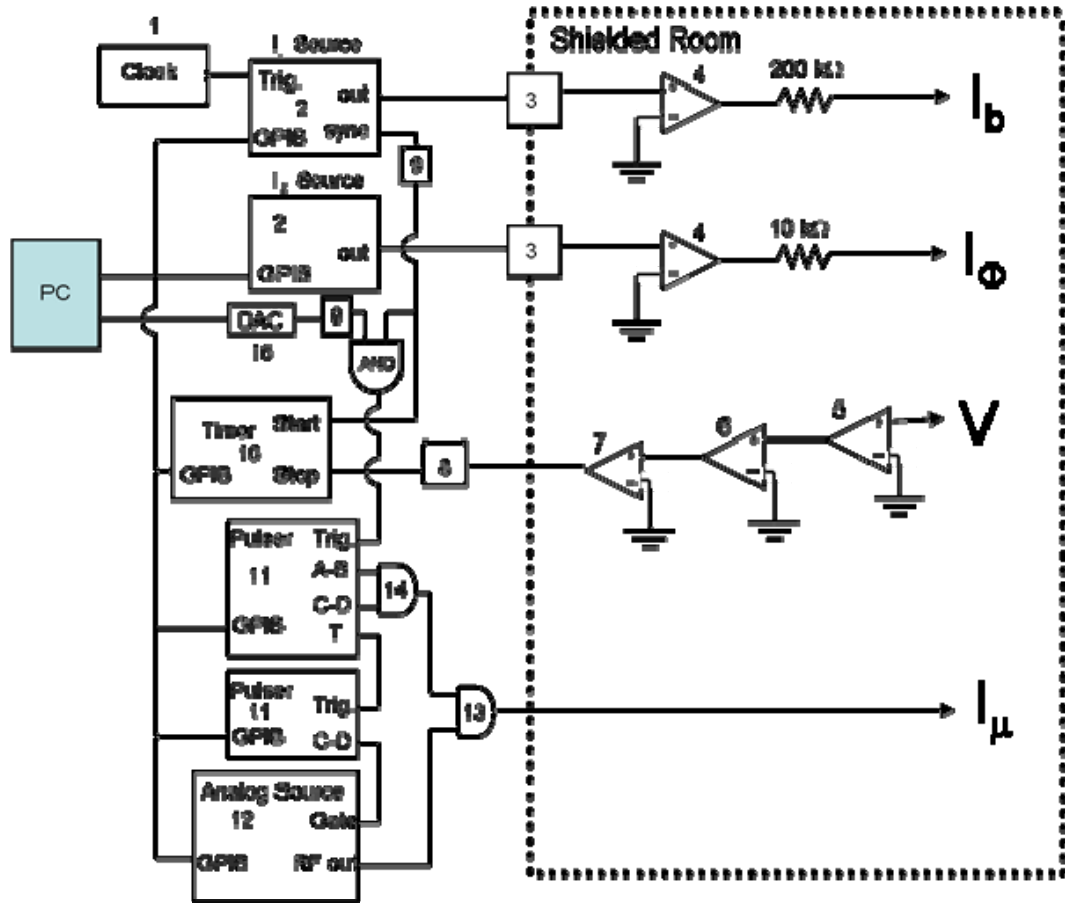


Fig. 7.14 Current-flux (I - Φ) characteristic of de SQUID phase qubit DS6. The graph is a compilation of switching histograms taken with the flux bias set to a constant value. The fact that the curves overlap means that the SQUID has several stable flux states.



- | | |
|--|---|
| 1. Exact Model 628 Pulse Generator | 8. Optical Pulse Converter |
| 2. Agilent 31220A Function Generator | 9. Opto-Isolator |
| 3. Low pass RC-filter ($f_c = 20$ kHz) | 10. SR 620 Time Interval Counter |
| 4. 1X Buffer | 11. DG535 Digital Delay/Pulse Generator |
| 5. JFET Amplifier | 12. HP83732 Synthesized Signal Gen. |
| 6. Ultra Low Noise Amplifier | 13. HP Power Splitter/Combiner |
| 7. Schmitt Trigger and Optical Transmitter | 14. Mini-Circuits Splitter/Combiner |
| | 15. DAC |

Fig. 7.15 Diagram of the setup for pulsed measurements and manipulations of the qubit state. The microwave excitation signal was combined with the measurement and reset pulses and sent to the qubit via the microwave line.

Sect. 7.2). The signal was coupled to the qubit through a small on-chip capacitor ($C_{\mu} \sim 1$ fF). The DG535 generated a square wave voltage signal and the rising edge of the square wave was differentiated by the coupling capacitor on-chip and transformed into a ~ 2 ns current pulse. The DG535 was triggered similarly to the SR620 time interval counter. The sync signal of the current bias source went through an optical isolator and was sent to an AND gate with another trigger pulse that came from a DAC card attached to the PC. The DAC pulse goes through an opto-isolator before it is AND-ed with the sync from the current bias source. Note that this was a case where this was how it had always been done, but it was not necessary to AND the signal with a DAC pulse. For this measurement, the output of the AND gate was sent to the trigger of the DG535. This set $T = 0$ for the delay of the DG535. Once the DG535 was triggered correctly, then it could be incorporated into the switching histogram setup. I also needed to set the operating bias for the qubit using the bias current source. The idea was to raise the current bias up to a certain point and hold the current there while I performed operations on the qubit, before measuring the qubit with the current pulse. This biasing was done by outputting a waveform from the HP 33120A that I constructed in *HP 33120A Waveform Editor.vi*. Instead of sending a constant ramp to the bias current source, I sent a waveform like the one shown in Fig. 7.16. In the first quarter of the waveform, I raised the current to its maximum. It was necessary to give the current sufficient time to stabilize because the filters in the wall caused the waveform to be rounded, and it was also best to raise the current relatively slowly so that there was no ringing once it reached the top of the ramp. The current was then held at the maximum for most of the time, and then dropped back to zero during the last 16th of the waveform. The whole waveform took 1 ms, and I did the

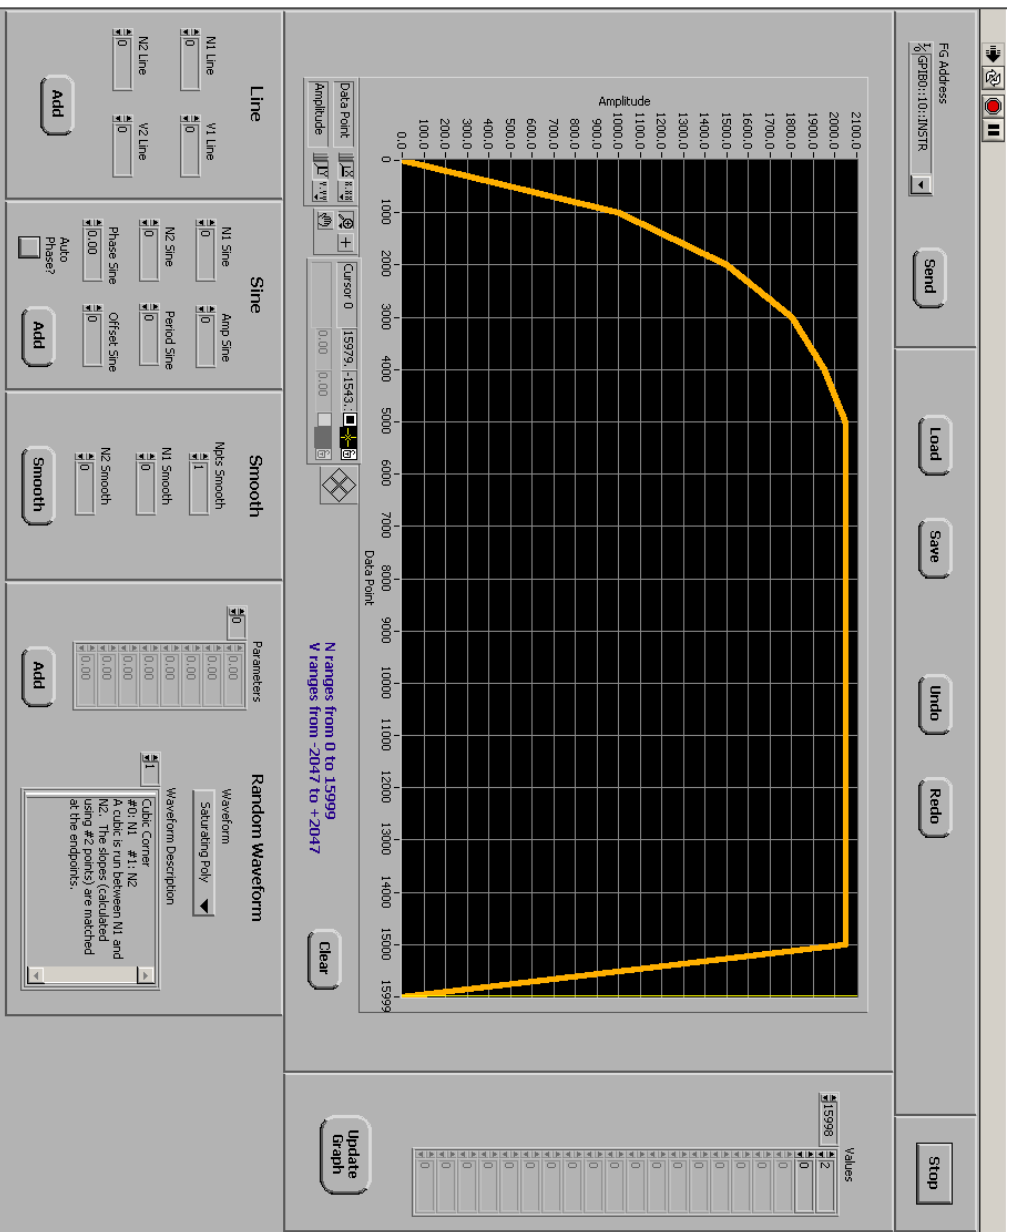


Fig. 7.16 Bias waveform used for setting the operating point for the qubit bias current. The measurement pulse occurred around DAC point 12800 and the reset was around 14400.

pulse measurement at the 800 μs mark. I set the operating bias simply by changing the maximum amplitude of the bias current ramp on the front panel of the source of the HP 33120A.

I used two current pulses during the state readout procedure: one pulse allowed the excited state to escape the well (measurement pulse), and the other dumped out the ground state near the end of the cycle to reset the device (the reset pulse). The two were output from the same DG535 but with different delays. The DG535 could be programmed with 4 delays: A, B, C, and D. These delays could be set with respect to each other (A, B, C, or D) or with respect to the trigger timing, denoted as T. I used delays A and C to set the timing of the measurement and reset pulses, respectively. The delays were set with respect to T at 800 μs and 900 μs for A and C, respectively. The DG535 had 9 outputs on the front panel. T0 outputs a signal when the unit is triggered. A, B, C, and D output a rising signal at each of their programmed delays. A-B will send the signal high when A is reached and bring it back down when B is reached. There is also an inverted A-B output. Delay B could be set with respect to A for the sake of easily making pulses of specific durations. C-D and inverted C-D work in much the same way. I took my pulses out of the A-B and C-D ports and set the B and D amplitudes at $T+0.00100000$ s so that the return to zero occurred after the qubit measurement was finished. The measurement and reset pulses were combined through a Mini-Circuits ZFRSC-183-S+ Power Divider and delivered to the qubit through the microwave line.

The correct height for the measurement pulse depended on the how close I wanted to bias the SQUID to its critical current. A good start point for the sake of calibration was to set the current bias source amplitude such that the bias current I_b was 90% of the

critical current. The control for the height of the measurement pulse was accessed in the “Output” menu of the DG535. I selected the source to be “AB”, and set the output impedance to “HighZ”. I set the output mode to “VAR” for variable amplitude output. Next, I set the amplitude to 4 V (this could be set with a precision of 10mV, on the 4V range). After that, I applied the same settings to the “CD” output.

I used a modified version of *OneSR.vi* called *OneSRpercents.vi* to calculate the ratio of the escapes caused by the measurement pulse to those caused by the reset pulse. If I turned the amplitude of AB to 10mV, then all of the escapes were caused by the reset pulse. If there were no escapes occurring with these settings, then I knew I needed to turn up the bias current. In that case I increased the bias current until I saw that the device was switching 100% of the time during the reset pulse. Next, I set the amplitude of the measurement pulse (AB) to 2V and adjusted the bias current such that there were only 3% of the switches occurring during the measurement pulse and the rest during the reset pulse. This bias current would be close to the minimum bias current that I could apply to the SQUID and still perform the pulsed measurement. I could operate at higher bias currents by putting attenuators on the A-B line or turning down the output of AB.

Figure 7.17 shows a typical switching histogram that I obtained on device PB9 using the pulse measurement scheme. Each peak in the histogram represents the junction switching during either the measurement pulse or the reset pulse. If the measurement pulse height is set correctly, then the qubit should only escape the potential well if it is in the excited state. Therefore, I have labeled the peaks as the “ $|0\rangle$ ” and “ $|1\rangle$ ” peaks. The ratio of the counts that make up the $|1\rangle$ peak to the total number of switches gives me the




Fig. 7.17 Switching histograms from a pulsed measurement. (a) Histogram obtained if there is no population in the excited state. (b) Switching during the measurement pulse which is indicative of population in the excited state. The peaks are labeled in quotes because of the anomalous switching behavior (see Sect. 9.4).

probability of switching during the measurement pulse. I used this pulsed measurement technique to measure all of the quantum mechanical behavior of the qubit.

7.7 Finding the Qubit Resonance Frequency

One of the more difficult and time consuming steps in the characterization process was initially finding the resonance of the qubit. Although I could identify the critical current of the qubit junctions I_{0I} from the I-V and I- Φ , the best idea I had of the capacitance of the qubit was still just the design value. Therefore, I started with my best guess for the location of the qubit resonance and searched from there. With the pulse measurement set up, I was able to identify excitations in the qubit reliably. To find the transition, I set up an automated microwave scan to look over a broad frequency range and eventually zeroed in on the resonance of the qubit.

To first order, the resonance of the qubit is given by:

$$f_{01} \approx \frac{1}{2\pi\sqrt{L_{J1}(C_x + C_{J1})}} \quad (7.5)$$

where C_x is the capacitance of the shunting capacitor on the qubit junction, C_{J1} is the capacitance of the junction itself, and L_{J1} is the *biased* Josephson inductance of the qubit junction. I estimated C_{J1} from the area of the junction and values obtained from measurements on our previous phase qubit. In particular, Tauno Palomaki measured the capacitance of a qubit with a $16 \mu\text{m}^2$ junction to be about 0.4 pF [52]; therefore, I could estimate the capacitance of our qubit junction was:

$$C_{J1} \approx \frac{A_{J1}}{16\mu\text{m}^2} 0.4 \text{ pF} \quad (7.6)$$

where A_{JI} was the area of the qubit junction. For device PB9, $A_{JI} \approx 0.23 \text{ } (\mu\text{m})^2$, which gave $C_{JI} \approx 5.8 \text{ fF}$. L_{JI} could then be found from:

$$L_{JI} = \frac{\Phi_0}{2\pi} \frac{1}{\sqrt{I_{0I}^2 - I_I^2}} \quad (7.7)$$

where I_I was the value for the current flowing through the qubit. Unfortunately, it was not generally a simple matter to determine I_I . It depended not only on the bias current, but also on the flux. In the case of device PB9, the current flowing through the qubit could be any value from $-I_{0I}$ to $+I_{0I}$ depending on the flux bias. So for the sake of the resonance search, I needed to assume that I_I could be anywhere between 0 and 99% of the critical current. These values along with the design parameter for C_x gave a range over which I needed to search for f_{0I} .

To search for the qubit resonance, I sent a microwave signal at relatively high power for about 100 μs while the flux and bias current were being held at a constant value. The bias current and the pulse readout were set up as described in the previous section. I used a Hewlett Packard 83732B, 10 MHz to 20 GHz, Synthesized Signal Generator to supply the microwave signal. So as not to cause extraneous excitations during the ramp up of the current bias, I turned the microwaves on after the current had reached a stable value. To do this, I used a pulse from a second DG535 to gate the microwaves (see Fig. 7.15, second pulser from the top). In order to incorporate the second DG535 into the setup, I triggered it off of the “T₀” output of the measurement/reset pulser. “Daisy chaining” the pulse generators in this fashion creates a discrepancy between the T = 0 delay points for each, but the discrepancy remains constant, whereas if I tried to split the trigger signal and send it to both at the same time, I

found there was significant jitter in the start times. Next, I hooked up the “AB” output of the second pulser to the “Pulse/Trig Gate In” input on the front of the microwave source, and pushed the “Ext On/Off” button directly above the input. I set the output for “AB” to the same settings as the reset pulse (*i.e.* VAR, Amplitude = 4.00 V, *etc.*). Time delay A was set to 100 μ s before that of the measurement pulse, and delay B was set to $A+0.000400000$, so that the pulse returned to zero after the cycle was over. With these settings, the microwaves were on for 100 μ s, then the measurement pulse arrived to test the system it was in the excited state. 100 μ s after that the reset pulse dumped out the ground state and the system was prepared for the next sample.

Instead of sitting in front of the computer and changing the frequency manually all day, I used *Run Scan.vi* to automatically sweep through a range of frequencies overnight. The procedure was similar to the one I used to set up the I- Φ scan, but this time I sent the command “SOUR:FREQ [frequency] MHZ” to the GPIB address of the microwave source. I set the power of the microwaves using the front panel of the HP83732B to a relatively high value. High power caused any qubit resonance peaks to be broadened significantly (power broadening) [38]. This meant I could use a large frequency step without fear of missing the qubit resonance. The only problem was that I could not be sure of the strength of the coupling of the microwave line to the qubit, so I did not know exactly what a high power setting was. I found that typically between -30 and -50 dBm was sufficiently high for a sample in our main refrigerator (a.k.a. beta fridge or the big fridge), but it really depended on the loss in the lines, the size of the coupling capacitor, and numerous other factors.

If I took 6000 samples per setting and ran at a rep rate of 700 Hz, then I could scan about 2 GHz per hour with a 10 MHz step size. With the estimate for the qubit frequency I obtained above, I was able to find the resonance in a few hours. It was also good to get a look at the entire frequency range by scanning overnight, and so I swept the entire 20 GHz range with a 7 MHz step size in about 13 hours. 7 MHz was not that big of a step, so I got away with setting the power to -50 dBm. If the power had been too high, then the scan would show a lot of extraneous peaks at frequencies where there were resonances in the lines or cavity resonances in the sample box, so I tried to choose scan parameters that I thought would give good results.

Figure 7.17 shows an example of the histogram that came out of a scan of device PB9. If the microwaves were not resonant with the qubit (or something else in the setup), then the histogram would only have one peak as in Fig. 7.17(a). This peak comes from the ground state being allowed to escape the well by the reset pulse. If there was a resonance, then sometimes the system would escape during the measurement pulse and I would get a second peak as in Fig. 7.17(b). I labeled the two peaks “ $|1\rangle$ ”s and “ $|0\rangle$ ”s in Fig. 7.17(b) because I could not be sure whether the escapes caused by the measurement pulse were escapes due to the qubit being in its $|1\rangle$ state or if the microwaves might have caused some increase in the current at the frequency due to some other resonance. At each frequency I took the number of escapes in the first histogram peak and divided by those in the second to give the probability P_{sw} of switching during the measurement pulse.

Figure 7.18 shows P_{sw} vs. frequency for one of my broad resonance searches in device PB9. The sweep was done with a power of -51 dBm and a step size of 10 MHz. Note the two large peaks around 4 GHz. The broader one on the left was later identified




Fig. 7.18 Results from a broad resonance search. Because the search was done with a high microwave power, one can expect to see multiple peaks. The narrow peak near 4 GHz is the qubit resonance. The larger, broader peak at lower frequency was due to an undesirable resonance in the setup.

as being due to a resonator that was coupling to the qubit. The much narrower peak on the right was my first look at the qubit resonance in device PB9. Once I identified the position of a peak it was necessary to see whether the position varied with changes in the bias current and flux bias. Also, I usually did a series of frequency sweeps over a region of interest at different powers to see how the height varied with power. A rapid variation with power is suggestive of a multi-photon process. Finding the resonance was the first step toward using the device as a qubit.

7.8 Measuring the Energy Spectrum

Because the inductance of a Josephson junction depends on the current flowing through it, the resonance of a phase qubit changes depending on the bias conditions. A change in either the flux or current bias should result in a shift in the resonant frequency. In order to characterize a qubit, I needed to make a map of the location of the qubit resonance versus the applied flux and current bias. This map could then be fit to a model of the device to determine the capacitances, critical currents, and geometrical inductances of the device. Any splittings in the spectrum would be evidence for the presence of microwave circuit resonances or two-level systems (TLS's) in the junction or neighboring dielectric. Also, the width of the spectral peaks could be used to find a lower bound on the coherence time of the qubit. Thus, the energy spectrum could be used to find information about the qubit parameters.

The setup of the energy spectrum measurement was similar to that of the resonance search (Sect. 7.7). I set the measurement pulse height so that with the microwaves turned off, switching occurred only about 3% of the time. If the $|0\rangle$ escaped

a few percent of the time, then this ensured that the measurement pulse dropped the potential barrier low enough that the $|1\rangle$ state was highly likely to switch. I set up the DG535 controlling the gate of the microwave source to start the microwave power about $10\ \mu\text{s}$ before the measurement pulse. This allowed for plenty of time for the system to reach a mixed state of $|0\rangle$ and $|1\rangle$. Of course, if the coherence time was extraordinarily long, then I should have turned the microwaves on earlier. With the power set to the value used for the resonance search, the resonance was easy to see. When I set the frequency of the microwaves to the location of the resonance, the probability to escape during the measurement pulse would increase markedly. I then turned the power down a bit at a time and tracked the resonance peak because the resonance frequency could shift at high powers. Ideally, I needed the power to be low enough that the qubit had only about a 10-20% probability of switching during the measurement pulse. The spectroscopy was performed at this low power to minimize power broadening of the spectral peaks [38]. Properly set, this procedure produced a narrow, well-defined, spectral peak when I swept the microwave frequency across the resonance.

The plan was to sweep the microwave frequency over the resonance peak, essentially measuring the spectral peak while applying a fixed flux and bias current, then step the current—either flux or bias—and measure the peak again. This was a time consuming measurement, and if I hoped to capture a significant portion of the spectrum in a reasonable amount of time, then I had to limit the range of frequencies over which I swept at each bias point. Therefore, I first took a rough map of the resonance frequency, f_{01} , by hand before automating the system to capture the energy spectrum. Several parameters needed to be determined at each bias point in order for the scan to be

successful. First, I set and recorded the current and flux bias. Next, I turned the microwave power off completely and adjusted the measurement pulse amplitude such that the system escaped out of the ground state 3% of the time and recorded the pulse amplitude. I turned the microwaves on and adjusted the frequency to find f_{0I} for this bias point. Usually, I used a little bit higher power if it was hard to find the resonance at first. I then adjusted the power of the signal such that the system escaped 10-20% of the time when the frequency was set to f_{0I} and recorded the necessary power. In addition, I made note of any changes in the spectral width of the resonance peak. Finally, I changed the current bias source voltage by 10, 50, or 100 mV (depending on the bias resistor and the critical current of the qubit junction) and repeated the process. This gave a rough map of the qubit resonance that would be filled in by the automated scan.

I needed to track each of these parameters as the bias changed in order to get an accurate picture of the energy spectrum in a reasonable amount of time. The 3% point for the measurement pulse amplitude was bound to change as the bias current to the qubit changed. If I set it too low, then the device would never switch and I would not see the excitation of the qubit. If it was too high, then the device would switch too often when the microwave drive was not resonant, and the resonance peak at that bias would sit at a much higher offset than the others in terms of P_{sw} , which would throw off the color map of the spectrum. By tracking the movement of f_{0I} over the bias range, I could sweep a much narrower range of frequencies and just keep the sweep centered on f_{0I} instead of sweeping over the full range of possible f_{0I} 's. Also, while searching for f_{0I} , I got an idea of the width of the peak, so I knew if I needed to change the width of the frequency sweep at certain biases. I also found that the coupling of the microwave line to the

SQUID also tended to change as I changed the bias settings, so I needed to keep the microwave power set correctly to avoid power broadening. The time invested in acquiring these parameters ensured that the automated scan of the spectrum would result in a wider range of useful data.

When constructing the scan parameters file, I needed to find an acceptable balance between precision and range of the acquired data. First, I needed to sweep a wide enough range of frequencies for each bias setting so that I could capture the entire resonance peak. This data would be used later to fit the peak to a Lorentzian, so I needed a wide enough scan that I could measure the offset of the Lorentzian. Then again, I did not want to capture too much useless data, so the scan needed not to be too wide.

Second, I needed to determine the correct size for the frequency step. The main consideration was my goal for the spectrum scan. If I was hunting for splittings, then I needed to be sure that I captured the shape of the peak including any indications that there might have been a duality or even plurality to the peak which could be indicative of coupling to a TLS. In this case, I used a step size that was about 10% of the estimated width of the peaks. If I was only trying to get a map of the peak positions, then I could use a frequency step that was on the order of the width.

Similar considerations governed my choice of the size of the bias current steps. If I was looking for abnormal features in the spectrum, then I needed a small step, but if I wanted an overall look at the spectrum, then I needed to take large steps in the bias.

I used *Run Scan.vi* to control the automated scan. The general idea was to set up the file such that *Run Scan* set the bias point, measurement pulse amplitude, and microwave power to the right values and held them constant while sweeping the

frequency. This could be done using *Run Scan*'s front panel; however, it would be necessary to fill in the gaps of the rough map that I acquired by hand as well as input each frequency sweep using the "Fill" tab. This could be tedious and time consuming, though, so I created a Matlab code that used a few "for" loops to create an array of settings that I then fed into a VI called *Run Scan New.vi*. It was basically a variation on *Run Scan*, but it allowed me to import a tab delimited text file into the "Values" array and it had some other improved functionality that will be described in later sections. I set the commands as described above and *Run Scan New* saved the data to a folder labeled by the date.

Figure 7.19 shows a typical energy spectrum as well as a single resonance peak captured using the above method. The resonance peak in Fig. 7.19(b) shows a single frequency at a specific current and flux bias setting. The height of each resonance peak was converted to a color value to produce a false color plot [see Fig. 7.19(a)]. Note that the blue area (*i.e.* no excitation) around each peak was kept small by careful selection of the frequency sweep parameters. The power settings were harder to keep track of so at $I_b = 6.3$ and $6.6 \mu\text{A}$, for example, the color of the peaks changed abruptly. This spectrum captured the behavior of the qubit over a wide range of bias currents and resonant frequencies with relatively high resolution, but it was taken over the weekend and the total scan time was about 60 hours.

7.9 Measuring the Switching Probability Curves (*a.k.a.* *S-Curves*)

Before I proceeded further with characterizing the qubit, I needed to get a better idea of the measurement pulse size that would give the highest fidelity measurement of




Fig. 7.19 (a) Energy spectrum of device PB9. Note that the range of frequencies tested for each current bias setting was kept as narrow as possible. The false color map represents the height of the peak in (b). The inconsistency of the peak heights is because I had to adjust the microwave power with the current bias.

population in the $|1\rangle$ state, P_I . For measuring the spectrum I used the measurement pulse amplitude V_p that gave a 3% probability of switching P_{sw} from the ground state, and that worked well for the spectrum, since a precisely known probability was not needed to detect the peak. However, for measuring energy relaxation, Rabi oscillations, and other measurements that rely on an accurate measurement of P_I , I need to fully characterize the switching from the $|0\rangle$ state, the $|1\rangle$ state, and the saturated state or the 50/50 superposition state. This characterization revealed whether the qubit P_I can be measured with a single pulse amplitude or if the qubit was exhibiting anomalous switching behavior as Hyeokshin Kwon has reported in [64,70]. The measurement itself involved measuring P_{sw} over a wide range of current pulse amplitude V_p . At low V_p , the device should not switch at all, but as I turn up V_p the device should switch increasingly more often until it was switching 100% of the time when the highest V_p was applied. If I excited the system with a microwave pulse, then the device started to switch with a lower V_p because I did not have to drop the potential barrier as much to get the system to tunnel out of the $|1\rangle$ state. I needed to understand the switching behavior in detail before I could move on to make quantitative measurements of the coherence of the qubit.

To proceed, I first measured the $|0\rangle$ state S-curve, which was the easiest part of the measurement. All I did was make sure that the microwave output was turned off and set *Run Scan New* to sweep V_p . I found the start point of the sweep by changing the output of the DG535 that controls the measurement pulse current amplitude until I found the largest amplitude for which there was still 0 switches in 6000 samples. I set the start point 0.10 V lower than that value. I set the end point 0.10 V higher than the smallest amplitude at which the device switched 100% of the time. A 0.02 V step was usually fine

enough to capture the S-Curve and this corresponded to roughly 10 nA of current to the SQUID. Because of the 2-D nature of the potential well and the anomalous switching behavior, it was unclear how much pulse current went directly to the qubit junction. I could usually capture an S-curve in less than 30 minutes.

It took a bit more work to set up the S-curve for the $|1\rangle$ state. First, I needed to set the microwave source to the qubit resonant frequency. The spectrum gave a very good starting point for the search, but there were often shifts in the flux that occurred during helium transfers and other fluctuations that could cause f_{01} to shift around for a given bias. I chose a clean resonance (*i.e.* one with no splittings, a narrow peak, *etc.*) and set the bias conditions and microwave frequency to that point. Then I zeroed in on the resonant frequency using a similar procedure to that used in preparing the rough frequency map prior to taking an energy spectrum scan (Sect. 7.8).

Next, I needed to set up the microwave gate timing such that I applied a high power microwave signal for a short amount of time that would leave the qubit completely in the $|1\rangle$ state (or as close as was possible to determine). The HP83732B microwave source would output a poorly shaped pulse if the gate was turned on and off too quickly, so I used a 24 ns long pulse as my “ π -pulse” (so named because it swings the Bloch vector around the sphere by π radians—from $|0\rangle$ to $|1\rangle$). This was done by setting the “B” delay of the DG535 that controls the microwave gate to “A+0.000000024.” Now, as was alluded to earlier, there was some delay between T=0 on the measurement pulser and the microwave gating pulse. This delay was determined either by using an oscilloscope to look at the pulses and match up their rises or by using the qubit. To use the qubit, I set the measurement pulse to a known value (*e.g.* T+0.000800000) and set A of the gate

pulse to be the same. I then adjusted the gate pulse delay such that there was a maximal (but still small) enhancement of the probability to switch during the measurement pulse. Setting the gate pulse to start 1 ns later than this point made sure that the measurement pulse occurred just as the microwaves were being turned on. I kept the microwave gate pulse stationary in time and moved the measurement pulse to $T+0.000800024$ so that it measures the qubit directly after the π -pulse was applied.

I next needed to set the power to the right value. Too much power would rotate the state past $|1\rangle$ and too little would not rotate it far enough. I set V_p to about 0.20 to 0.40 V lower than the 3% point and began turning up the microwave power. P_{sw} increased as I increased the power, but once the power reached a certain point, P_{sw} would begin to go down. I set the power such that P_{sw} was at a maximum. This microwave power gave a 24-ns π -pulse.

With the measurement pulse correctly set, I set up the scan for measuring the S-curve of the $|1\rangle$ state in *Run Scan New*, as described above. It was possible that the start point of the $|1\rangle$ state S-curve would need to be earlier than for the $|0\rangle$ state, so the start and end points of the V_p sweep needed to be determined again.

The S-curve of the saturated state was captured by first preparing the state with a long microwave pulse and then sweeping V_p . The microwave pulse for this measurement looked much like the one used for spectroscopy. The microwaves needed to be on long enough such that the state would be completely dephased, so I set the B delay of the microwave gate to $A+0.000010000$ and set the measurement pulse to occur at $T+0.000808000$ (8 μ s after the microwaves turned on). The power was set to the same

power as was used for the π -pulse in capturing the $|1\rangle$ state S-curve. Again, I checked the start and stop values for the V_p sweep and used *Run Scan New* to perform the scan of the saturated state.

Figure 7.20 shows plots of the $|1\rangle$, $|0\rangle$, and saturated state S-curves for device PB9. Note that the $|0\rangle$ state S-curve begins to rise at a higher V_p than the other two. Now, in a conventional S-curve, P_{sw} (near the foot of the $|0\rangle$ S-curve) is a rough measure of the population of the $|1\rangle$ state P_I . However, the shape of the saturated state's S-curve in Fig. 20 tells us that this was not the case for device PB9. If it were, then the saturated state S-curve would look like it were a combination of half the P_{sw} from the $|1\rangle$ curve and half the P_{sw} from the $|0\rangle$ curve. In general, this would create a plateau in the saturated state S-curve in the region between where the $|1\rangle$ and $|0\rangle$ curves have the sharpest rises. If I had seen this plateau, then I could have said that the device was not exhibiting anomalous S-curve behavior and it would have meant that P_I could be measured in the traditional fashion using a single, well-chosen value for V_p [64]. In that case it would have been important to identify the V_p where the $|1\rangle$ and $|0\rangle$ curves are the most different because this would have been the V_p that gave the best measurement fidelity. However, in the case of a device with anomalous S-curves, I would need to capture the entire S-curve for the state at every time t as well as the S-curves of the $|1\rangle$ state, the $|0\rangle$ state, and that of the 50/50 coherent superposition state $(|0\rangle+|1\rangle)/\sqrt{2}$ (not the same as the saturated state) in order to back out its P_I .




Fig. 7.20 Switching probability curves (S-curves) for three distinct qubit states. Note that it takes a much higher measurement pulse amplitude (V_p) to cause the qubit to switch if it is in the $|0\rangle$ state than if it is in the $|1\rangle$ state. Also, note that the S-curve of the saturated state is more like a shifted version of the $|0\rangle$ curve than a 50/50 combination of the $|0\rangle$ and $|1\rangle$ curves.

7.10 Measuring the Energy Relaxation Time T_1

The population in the $|1\rangle$ state will decrease at an exponential rate characterized by the time constant T_1 , the energy relaxation time. To measure T_1 , I placed the qubit in $|1\rangle$ with a short, high power π -pulse and tracked the population in the $|1\rangle$ state as it decayed. The analysis and fitting of the S-curves to pull out P_1 will be discussed in Chapter 9 as part of the analysis of device PB9. The T_1 measurement required a long time to perform when S-curves were required, but it was the only way to get an accurate reading of the relaxation time.

The beginning of the procedure is similar to the steps I used for measuring the $|1\rangle$ state S-curve described in the previous section. I found the resonant frequency f_{01} , set up the gating pulse to the microwave source such that the source output a 24 ns pulse, and set the measurement pulse to T+0.000800024. This gave a measure of P_{sw} at the end of the microwave pulse. I set the power such that the microwave pulse was a π -pulse and figured out the start and end points of the V_p sweep. In this measurement, the start point was 0.10 V lower than the point where $P_{sw} = 0$ with the microwaves on and 0.10 V higher than the point where $P_{sw} = 1$ with the microwaves off. The microwave source was set up to output a pulse that would prepare the system in the $|1\rangle$ state and the pulser was set to sweep out the s-curve of the state.

In order to measure the decay time of the qubit out of the $|1\rangle$ state, I measured the S-Curve of the system before, during, and after the microwave pulse. The S-curves taken before and during the microwave pulse were used later to fit the decay S-curves to a model from which I can back out P_1 at each time during the decay [64]. Depending on

the width of the V_p sweep for each S-curve, I had time to take between 40 and 80 S-Curves in about 14 hours. Many of them had to be done during the microwave pulse to make sure I would be able to fit the decay S-curves. The first measurement time was 100 ns before the microwave pulse turns on at $T+0.000799900$ in order to get the S-curve for $|0\rangle$. Then I took an S-curve every 1 ns between $T+0.000800008$ and $T+0.000800016$. The model requires a good S-curve of $|1\rangle$, $|0\rangle$, and the coherent 50/50 superposition state, so one of the S-curves taken during the π -pulse would supply an acceptable 50/50 S-curve. In order to capture the $|1\rangle$ state S-curve, I took one every ns from $T+0.000800022$ and $T+0.000800026$. This meant that I used 15 of the available S-curves just to make sure I would be able to model the S-curves of the decay and back out P_I vs. t .

With the remaining S-curves, I captured the full time span of the decay of the qubit. I plotted the data on a semilog graph, and the slope of the plot along with the baseline of the decay would be used to fit the data and obtain T_I . Therefore, I wanted to make sure that I captured the behavior of the decay with as much resolution as possible while at the same time following the decay out far enough in time such that I got a good measure of the baseline of the decay. For example, device PB9 had $T_I = 100\text{-}400$ ns, so I normally did the last S-curve of the decay at $T+0.000802530$. I filled in the intervening space with S-curves every 70 or 80 ns depending on the size of the V_p sweep for each one. These S-curves gave me a picture of the decay of P_I after the microwave pulse shut off.

In order to automate the scan I used *Run Scan New.vi*. *Run Scan New* has an advantage over *Run Scan* in that I could input the measurement times of the S-curves in the first column and use the “Repetitions” tab to run an S-curve scan at each time. First, I put the times in the first column either by hand or by using the “Fill” tab for any stretches

of times that were separated by a constant time step. After I finished inputting the timings of the s-curves, I went to the “Export/Import” tab, clicked the “Export Tab Delimited File” button, and saved the times as “times [date]”. Next, I went into the “Repetitions” tab. I set the “Rep Start Value” to the V_p at which $P_{sw} = 0$ when the microwaves are on and set the measurement time to $T+0.000800024$. I then set the “Rep Step” to 2, and I set “Rep Number” such that the V_p sweep would end on the V_p where $P_{sw} = 1$ when the microwaves were turned off. Then I clicked the “Rep Fill” button, and the V_p values appeared in the “Repeater Values” array. Finally, I set the “Rep Row” to 0 and the “Rep Column” to 1 and clicked the “Perform Repetition” button. This modified the main “Values” array such that the time remained constant in column 1 while the values in column 2 cycled through the V_p values. Once it reached the end of the V_p sweep, the time changed and the array ran through the V_p sweep again. At that point, *Run Scan New* was ready to capture the S-curves for the T_1 measurement overnight.

Figure 7.21 shows the evolution of the S-curve of the system over time in a false-color plot. The color axis represents P_{sw} , so each vertical strip of the plot is a single S-curve taken at that time that has been converted to a color map. Note that the S-curve of the $|0\rangle$ state taken 100 ns before the microwave pulse does not begin to rise until a relatively high value of V_p . The S-curves taken during the π -pulse that I used for the fitting are hard to see because of the time scale, but the decay is clearly visible. Note that the S-curves taken between $t = 1500$ and 2500 ns have basically the same characteristic because the qubit state had decayed to the ground state by this time.




Fig. 7.21 False-color plot of S-curves taken during the energy relaxation of qubit PB9 at $I_\phi = 0.084$ mA, $I_b = 6.8$ μ A, and $f_{0l} = 4.218$ GHz.

7.11 *Measuring Rabi Oscillations of the Qubit State*

A Rabi oscillation gives one measure of coherent quantum behavior in a qubit. For frequency independent relaxation and dephasing, the amplitude of a Rabi oscillation should decay with a time constant T' such that [52,71]:

$$\frac{1}{T'} = \frac{1}{2T_1} + \frac{1}{2T_2} \quad (7.9)$$

where T_1 is the energy relaxation time and T_2 is the coherence time. T_2 is defined by [52,71]:

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi} \quad (7.10)$$

where T_ϕ is the dephasing time, the time constant that characterizes loss of phase information in the superposition state. With T_1 from the measurement described in the previous section, and T' from the Rabi oscillation, I can then use Eq. 7.9 to estimate T_2 and Eq. 7.10 to estimate T_ϕ .

I performed a Rabi oscillation by applying a relatively high power microwave signal to the qubit and measuring P_I at a series of times during the microwave signal. In the case of a device with anomalous S-curves, this required that I capture the S-curve of the state as it evolves.

The setup for the Rabi oscillation was similar to the setup of the T_1 measurement, except this time the microwaves were left on for the entire measurement. I found the resonant frequency (similar to setting up for capturing $|1\rangle$ S-curve—see Sect 7.9), set the microwave source gate timing such that the microwaves came on at T+0.000800000 according to the measurement pulse generator (see Sect. 7.9), and set the B delay of the microwave source gate pulse to A+0.000050000. This gave a microwave signal that

started at a known time and continued for 50 μs , which was much longer than the coherence time of the qubit, but not so long that it impacted the reset pulse.

Next, I needed to set the power of the microwave drive. The power determined in the T_1 measurement (see Sect. 7.10) gave a microwave signal that would drive a Rabi oscillation with a period of about 48 ns. As a rule of thumb, I would expect a 6 dB increase in the power to cut this period in half since [38]:

$$\Omega_{\text{Rabi}} \propto \sqrt{P_\mu} \quad (7.11)$$

where Ω_{Rabi} is the Rabi frequency and P_μ is the microwave power. If I had a reason to expect that the Rabi oscillation would decay quickly (*i.e.* I saw a low T_1), then I would need to drive a fast Rabi oscillation, so that I could see some oscillations before the qubit decohered completely. However, if I expected a long coherence time, then I would not need to drive the oscillation too fast. In this case I drove a slower oscillation and measured P_1 less frequently. This allowed me to capture the character of oscillation and observe its full decay in a reasonable amount of measurement time. The expected Rabi decay time of the device was an important factor in selecting the microwave power, and hence the Rabi frequency, for the Rabi oscillation measurement.

I also needed to measure the S-curves of the $|0\rangle$, $|1\rangle$, and $(|0\rangle + |1\rangle)/\sqrt{2}$ state in order to fit the S-curves and extract P_1 . If the Rabi period was short (*i.e.* shorter than about 15 ns), then I measured S-curves every 0.5 ns during the first half-period. If it was longer, then I measured S-curves every 1 ns. This ensured that I got a good S-curve of the $|1\rangle$ state (at the top of the oscillation) and of the $(|0\rangle + |1\rangle)/\sqrt{2}$ state during the rise. I also set one measurement to occur 100 ns before the Rabi oscillation began

($T+0.000799900$) in order to capture the $|0\rangle$ state s-curve. The capture of these S-curves was critical to the analysis of the Rabi oscillation.

The period of the Rabi oscillation set a limit on the time interval between the measurements in the body of the oscillation. I needed to measure P_I at least 5 or 6 times every period in order to capture the character of the oscillation. The total number of measurements depended on the time allowed for the scan. If I expected the coherence time would be long, I would drive slower oscillations so that I could measure P_I less frequently, hence allowing me to follow the Rabi oscillation until it decays away with fewer measurement. The period and overall duration of the oscillation impacted its resolution.

Figure 7.22 shows the evolution of the S-curve of the state as I measured a Rabi oscillation. Fig. 7.22(a) shows a full set of oscillations. The yellow color on the graph gives a good guide to the eye for interpreting this graph because it shows the measurement pulse amplitude, V_p , at which the state is switching about 50% of the time. The $|0\rangle$ state escapes at higher V_p whereas the $|1\rangle$ will escape with a lower V_p as one would expect since V_p is proportional to the current pulse being applied to the device. During the Rabi oscillation we see the state oscillating back and forth between the $|1\rangle$ state and the $|0\rangle$ state. Over time the oscillation decays until it is difficult to discern and the yellow color appears at some median value of V_p . This behavior is consistent with the state decohering into an incoherent mixture of $|0\rangle$ and $|1\rangle$ at long time.

Figure 7.22(b) illustrates some details of the Rabi oscillation data. I measured the S-curve of the state every 3.5 ns during the body of the oscillation which gave me a




Fig. 7.22 False color plot of S-curves taken during a Rabi Oscillation. (a) The full Rabi measurement. Note that the S-curves decay to a spread out, saturated state, S-curve. (b) Close-up of the oscillation in (a). Because of the fast oscillation frequency, there are only 5 or 6 S-curves for each period. The S-curves start out at low V_p at $t = 0$ because I started observing the oscillations after the start of the first oscillation ($|0\rangle$ was captured at $t = -100$ ns).

measure of the state only 6 times per Rabi period. It was a rough measurement, but it allowed me to capture the Rabi oscillation over a longer period of time for a relatively fast Rabi frequency (more than 50 MHz). This Rabi oscillation was measured using the maximum amount of time between helium transfers, about 90 hours. The first 10 ns of the graph shows my attempt to capture the $(|0\rangle + |1\rangle)/\sqrt{2}$ state and the $|1\rangle$ state (the $|0\rangle$ state S-curve was captured 100 ns before the oscillation so it is not on this plot). I took S-curves every 0.5 ns starting at $T+0.000800002$ until $T+0.000800009$ (time = 2 to 9 ns on the graph). This is an example where more data would have been better. We see that the measurement definitely captures the $|1\rangle$ state S-curve, or very nearly so, because the switching current starts increasing again after about 7.5 ns. I used this S-curve at 7.5 μs as the $|1\rangle$ S-curve for the data analysis. However, I should have started taking data a little earlier. Although it is possible that the S-curve taken at 2 ns could be adequate for the $(|0\rangle + |1\rangle)/\sqrt{2}$ S-curve, it would have been better to capture several S-curves in a time interval where the population in the $|1\rangle$ state was in the range of 50% to be sure that I got a good S-curve for the $(|0\rangle + |1\rangle)/\sqrt{2}$ state.

7.12 *Creating Short Microwave Pulses with Sideband Mixing*

The remaining measurements that I did required a fairly sophisticated microwave pulse setup. Although it is technically true that I used a microwave pulse to perform the T_1 measurement, the pulse did not have to be of very good quality. I was mostly interested in preparing the qubit in a state with a high population in the $|1\rangle$ state. In contrast, Ramsey fringe measurements (as well as spin-echo and state tomography)

require precise control of the qubit state and multiple microwave pulses. For these more complex measurements, I applied a microwave engineering technique called sideband mixing and used a mixer to gate the microwaves on and off. The gate pulses were applied by an arbitrary waveform generator (AWG) with a sampling rate of greater than 1 Gsamp/s which I programmed with the desired pulse sequence. The setup delivered microwave pulses of a specific width, height, and shape to the qubit in an order determined by the AWG waveform.

Figure 7.23 shows a diagram of the setup. I used two analog microwave sources (the HP83732B described earlier and an Agilent E4426B ESG-AP 250 kHz to 4.0 GHz Series Signal Generator) an Agilent N8241A 1.25 GS/s – 15 Bit Arbitrary Waveform Generator, and a Marki Microwave M80412LS Mixer.

The most difficult part of the setup was getting the AWG to work correctly. The N8241A did not come with a very good software suite, and whereas most of our equipment interfaced with the PC through GPIB cables, the AWG talked through a LAN cable. To make the connection, I took an Ethernet card out of an old computer and installed it in the PC that ran the measurement setup. This Ethernet connection showed up in Windows 2000 in the “Network Connections” as “Local Area Network 2.” I right clicked the connection and selected “Properties.” Then I highlighted “Internet Protocol (TCP/IP)” and clicked the “Properties” button. The IP address of the AWG was 192.168.42.15, so I had to set the IP address of the Ethernet card to something like 192.168.42.xx (192.168.42.1 worked fine). The subnet mask was 255.255.255.0 and the default gateway was 192.168.42.1. It was not necessary to specify the DNS servers. This took care of the physical connection to the N8241A.

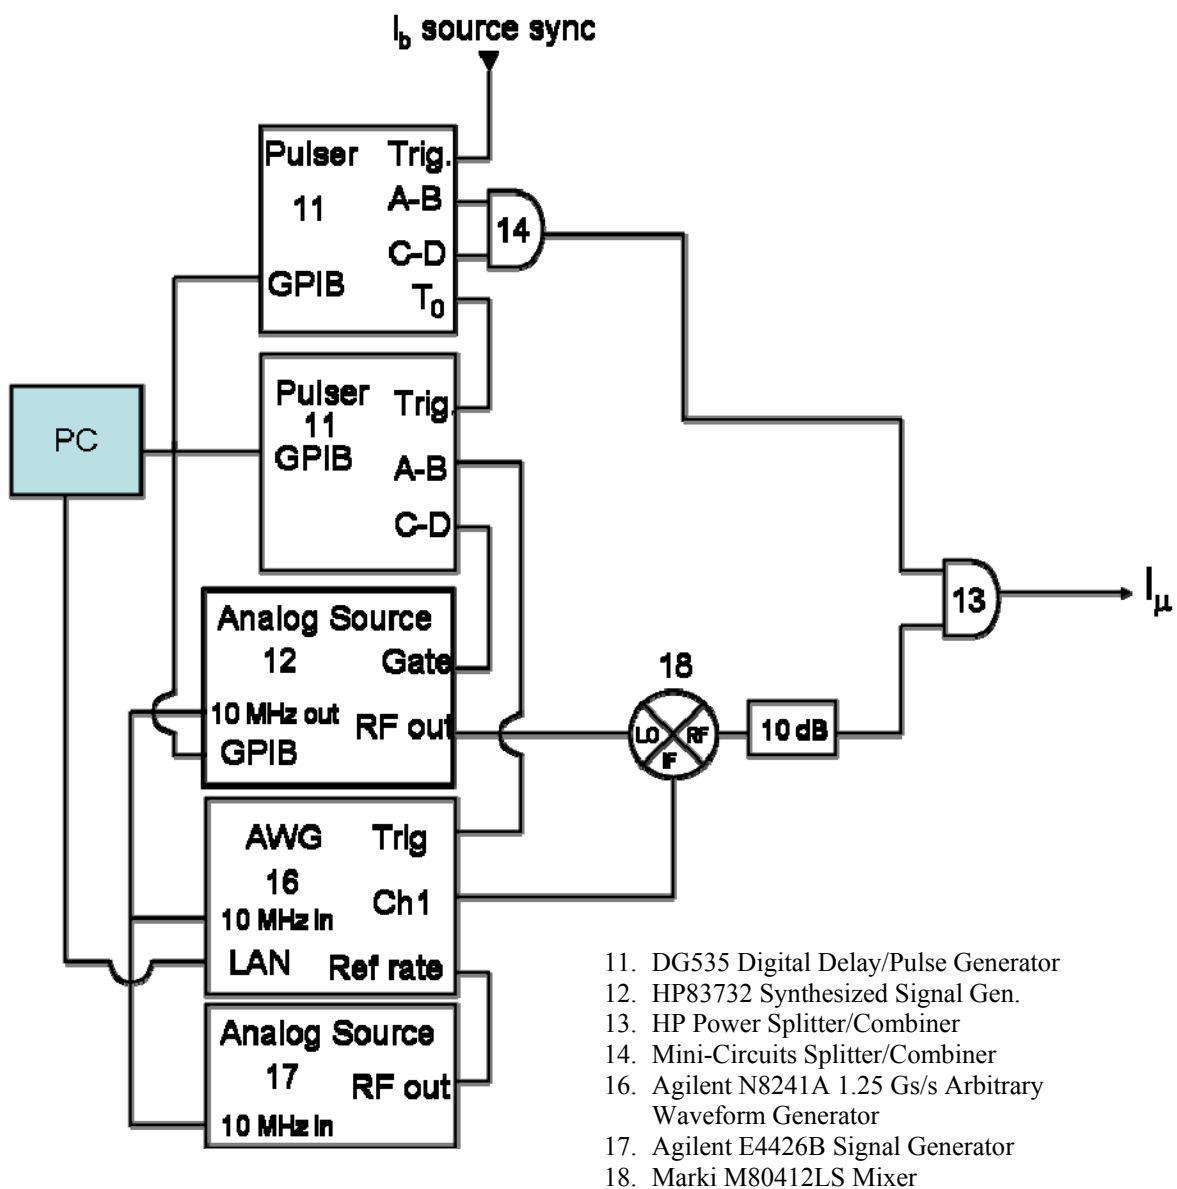


Fig. 7.23 Diagram of the mixed microwave pulse setup. The AWG ran off of an external clock that controlled the sampling rate of the AWG's DAC. Note that the two analog sources and the AWG use the same 10 MHz reference signal.

The software for the N8241A had to be installed in a specific order for it to work. The AWG came with two CD's: one for the IO Libraries Suite and one for the N8241A specific software. The version of IO Libraries on the disk would not work on a computer running Windows 2000; however, it did work fine if the computer was running Windows XP. To install the necessary software, I went to Agilent's website and downloaded IO Libraries Suite *ver. 15.1*. I unzipped and installed this on the computer. During the installation the program required that the IVI compliance software be installed from National Instruments; therefore, I went to the NI website, downloaded the latest version of IVI compliance, and installed that. Once the IO Libraries Suite was installed I went to Agilent's website and downloaded the drivers-only version of the N8241A software version 1.23.1. Note that the other versions might have worked okay, and the version on the disk had a convenient, yet unnecessary, utility for sending simple commands to the AWG, so it may be worthwhile in future setups to go with the software on the CD. If I made a mistake in the installation, and the connection to the AWG did not work, I had to uninstall everything (*i.e.* IO Libraries Suite, IVI Shared Components, and N8241A Software Package) using "Control Panel -> Add/Remove Software" and start over. Once I finished installing the software I had to restart the computer before the AWG would connect to the PC. This process had to be done in the correct order with the correct versions of the software or it would not work.

The N8241A AWG comes with a program with a graphical user interface (GUI), but it could not be incorporated into our measurement setup because it would not interface with *Run Scan*. The only way I found to program the AWG was through Matlab (LabVIEW can run Matlab scripts). In the help files for the N8241A, there were

examples for sending a waveform to the AWG, and most of the Matlab code was constructed around these examples. The code is provided in Appendix A. Basically we created a vector of numbers that defined the height of a sine wave at each time for the DAC. The default sample rate of the AWG was 1.25 GS/s, although it could be set to any value below that using an external clock source. Therefore, the default time step was 0.8 ns. The heights specified by our waveform were normalized (by default) to values between ± 1 and encoded to 15 vertical bits. The maximum output voltage was nominally 256 mV_{pp}, although this was reduced dramatically as the sine wave frequency approached the maximum value of 625 MHz. Matlab did most of the communication with the AWG through m-files provided in the N8241A drivers package.

Rangga Budoyo and Gabe Vilarrubi (a high-school student from the Barrie School) worked very closely with me on interfacing with and programming the AWG. We wrote several pieces of Matlab code for communicating with it as well as LabVIEW code for controlling it. We modified an existing program that had been used for controlling the HP33120A's so that it would tell the AWG to output a pulse of specified frequency and height that had a Gaussian envelope of specified width. Figure 7.24 shows a picture of the front panel of the VI. The program also could be used to create spin echo and Ramsey pulse trains automatically with a specified delay between the pulses. We also made a new version of *Run Scan* that would vary the height of the pulses. This was used for calibrating the pulses, setting π -pulses and setting the delay between the pulses automatically so that we could do Ramsey and spin echo experiments. This gave us a complete set of programs for controlling the AWG.

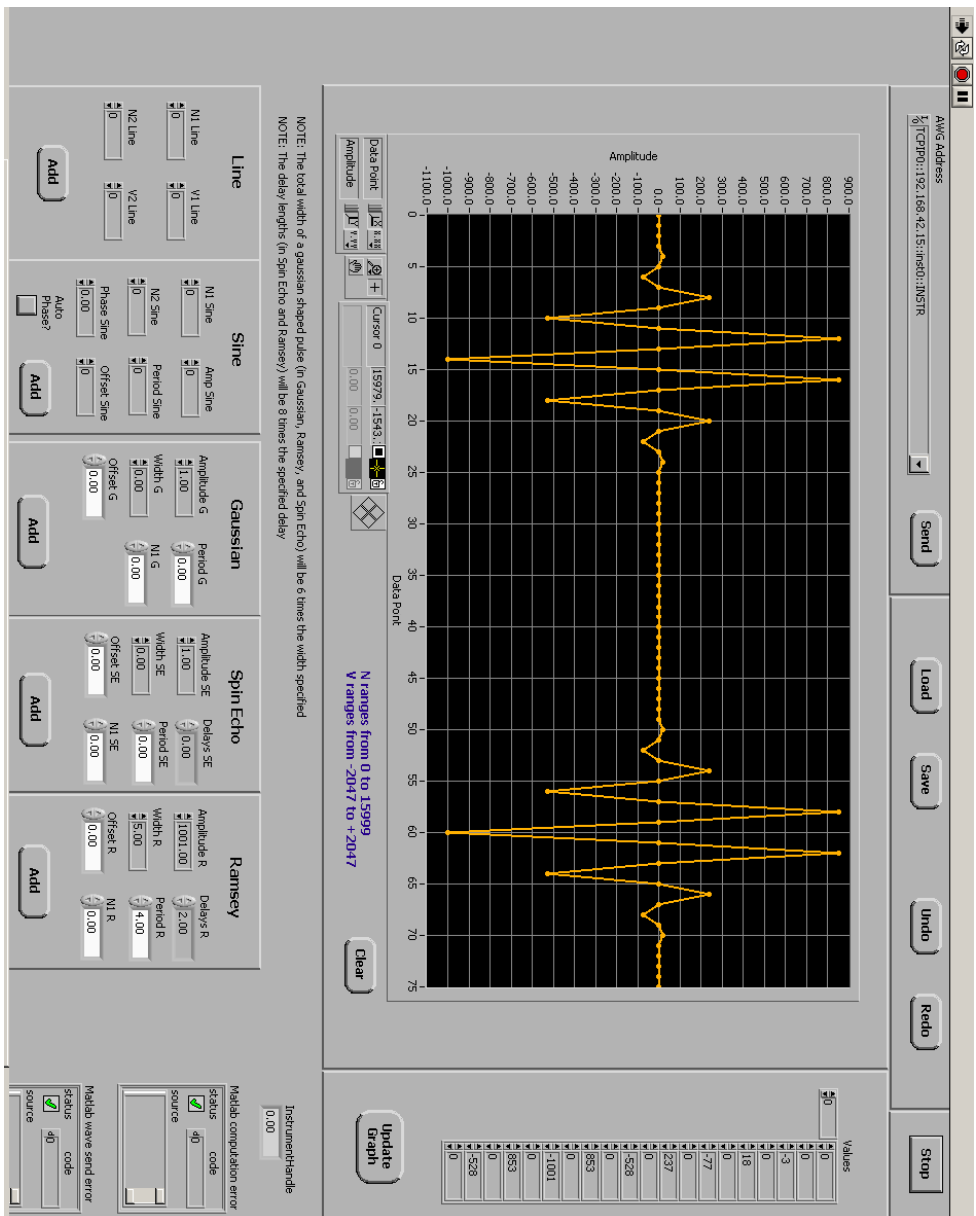


Fig. 7.24 Picture of the front panel of *N8241A Waveform Editor.vi*. This VI could be used to create pulse trains and send them to the Agilent N8241A arbitrary waveform generator.

The AWG pulses were used to modulate a continuous microwave signal from the HP83732B. To do this, I attached the signal from the HP83732B to the LO port on a Marki Microwave M80412LS mixer via SMA cable and the output of the AWG went into the IF port (see Fig. 7.23). The input to the mixer was:

$$\text{Input: } A \sin(2\pi f_\mu t) \text{ mixed with } B(t) \sin(2\pi f_{AWG} t + \phi) \quad (7.12(a))$$

and the resulting output at the RF port was:

$$\text{Output: } \frac{AB(t)}{2} \left\{ \sin[2\pi(f_\mu + f_{AWG})t + \phi] + \sin[2\pi(f_\mu - f_{AWG})t - \phi] \right\} \quad (7.12(b))$$

where f_μ is the frequency of the HP83732B, f_{AWG} is the frequency of the pulses from the N8241A, ϕ is the phase difference between the two sources, A is the constant amplitude of the microwave source, and $B(t)$ is the amplitude of the AWG signal. This simple picture ignores conversion losses along with other non-ideal behavior in the mixer. Essentially, the mixer outputs microwave pulses with that same envelope as the AWG signal. In the frequency domain, this output pulse is made up of frequency components that are the sum and difference of the two input signals.

Figure 7.25 illustrates the power spectrum of the output signal from the mixer. Note that the highest powers in the spectrum are at the sum and difference frequencies. These are known as the “sidebands.” Also note the power at f_μ . This comes from signal leaking from the LO port to the RF port. This leakage is the main reason for using the “sideband mixing” technique. Consider Eq. 7.12. If the $B(t)$ term in the input had a time varying amplitude but a frequency of zero (e.g. a square pulse), then the output would be a pulse with a frequency of f_μ and a square envelope. However, since there is some leakage of the signal from the LO port to the RF port, there would be some background




Fig. 7.25 Illustration of the output of the mixer in the mixed microwave setup depicted in Fig. 7.23. Note the leakage power at f_{μ} . This power will not affect the qubit if I set the bias such that $f_{0I} = f_{\mu} - f_{AWG}$

surrounding the pulse that was at frequency f_μ . If we set f_μ to f_{01} of the qubit and applied this pulse to the device, then the background would constantly be acting on the qubit even when the pulse was supposed to be turned off. This is undesirable. Therefore, instead of a square pulse, I used the AWG to apply a pulse with $f_{AWG} \approx 300\text{MHz}$, and set f_μ such that $f_{01} = f_\mu - f_{AWG}$. The result is that the qubit will be resonant with one of the sideband frequencies, and the leakage coming through at f_μ will not cause significant additional rotation of the state vector.

Equation 7.12 has another important implication for the qubit. Note that the phase difference between the AWG signal and the microwave signal carries over into the output signal. Therefore, a few things must be done to ensure that that phase difference remains constant. First of all, the two sources must use the same 10 MHz reference signal, so I connected the “10 MHz Out” port on the back of the HP83732B to the “10 MHz Ref In” port on the front of the AWG (the connector on the front of the AWG is an SMB jack). Unless this connection was made the phase difference between the AWG signal and the microwave signal would be random from one pulse to the next.

Secondly, f_μ had to be divisible by f_{AWG} . This was easier said than done since the sampling rate of the AWG, f_{clock} , had to be a power of 2 times f_{AWG} for the waveform to be of good quality. Figure 7.25 illustrates this point. The default sampling rate was 1.25 GS/s, but it could be set to any value less than that by supplying an external clock signal. I used an Agilent E4426B for this. I set the source to output a signal at f_{clock} with a power of 0 dBm and connected the output to the “Ext Clock In” port on the AWG. Also, I was sure to connect the 10 MHz signal from the HP83732B to the “10 MHz In” port on the

back of the Agilent E4426B. In addition, we added a line in the Matlab code that set the AWG to take an external clock: `agt_awg_setstate(instrumentHandle, 'samplerate', [fclock])`.

As I noted previously, I needed f_{AWG} to be divisible into f_{clock} and into f_{μ} . These three frequencies needed to obey the following relations:

$$f_{AWG} = \frac{f_{clock}}{2^n} \quad (7.13)$$

$$f_{clock} = \frac{f_{\mu}}{m} \quad (7.14)$$

$$f_{01} = f_{\mu} - f_{AWG} \quad (7.15)$$

where n and m are integers. Simple algebra gives the frequencies:

$$f_{AWG} = \left(\frac{1}{m \cdot 2^n - 1} \right) f_{01} \quad (7.16)$$

$$f_{\mu} = \left(1 - \frac{2^n}{m \cdot 2^n - 1} \right) f_{01} \quad (7.17)$$

$$f_{clock} = \left(m - \frac{1}{2^n} \right) f_{01} \quad (7.18)$$

Eqs. 7.16-18 gave the frequencies, but I also needed to be sure that they were rational numbers. Also, using these frequencies assured that there would be a time that all the signals were in phase at certain times. Figure 7.26 illustrates that the AWG and microwave signals did not have the same phase for at all times. It shows a case where $m = 4$ and $n = 4$. This meant that $f_{AWG} = \frac{1}{4} f_{clock} = \frac{1}{16} f_{\mu}$. Now, if all three signals started with the same phase, then they would not have the same phase again until 4 DAC steps later (that is, $t = 4 \cdot (1/f_{clock})$ later). Therefore, if I wanted to apply more than one microwave pulse to the qubit with the same phase, then the starting point of each pulse




Fig. 7.26 An illustration of how the phase varies between the AWG signal and the microwave signal. Each point on the x-axis represents $t = 1/f_{clock}$. Note all three signals are in-phase only at DAC points 0, 4, and 8. If the phase between the AWG and the microwave signals is different, then the microwave manipulation of the qubit state will be different.

from the AWG had to be separated by a number of DAC points divisible by 4. It was necessary to keep the phase the same for all pulses if I wanted the qubit to behave correctly during the experiment.

7.13 *Measuring the Ramsey Fringes*

I used a Ramsey measurement to probe the dephasing of the qubit state due to relaxation, high frequency dephasing processes, and low frequency (inhomogeneous) noise. The decay constant of the Ramsey experiment, T_2^* , typically gives a lower bound on T_2 , since a spin-echo technique could be used to eliminate effects of inhomogeneous noise. While it is true that a Rabi oscillation has a time constant T' that includes T_2 , one also needs to know T_1 and uncertainty in the measurement of T_1 made it difficult to extract an accurate value for T_2 . In addition, the relationship between T' , T_2 , and T_1 (see Eq. 7.9) only holds for frequency independent T_1 and T_2 , which is not the case if $1/f$ noise is present. On the other hand, the decay constant for a Ramsey measurement should be the same as the coherence time T_2^* that we extracted from full width at half maximum of the resonance peak, so I used it to confirm those results.

The Ramsey measurement required two microwave pulses of a specific height separated by a controlled delay time τ . Each pulse had the same amplitude. The pulses needed to be just strong enough for the first pulse to rotate the qubit state half-way around the Bloch sphere to the equator, and then the second pulse would rotate the state the rest of the way to $|1\rangle$. The pulse train I used is depicted in Fig. 7.27. First, I constructed the pulse waveform using a Matlab program called “ramseyLabVIEW.m” that constructs two pulses of the same height, width, and frequency. I set the width of the




Fig. 7.27 Illustration of the pulse train of a Ramsey measurement. The $\pi_x/2$ pulse rotates the state onto the equator of the Bloch sphere. The state is allowed to dipphase for a time τ . Finally, the state is rotated back to the North pole. A state on the equator that is completely dephased has equal probability of being detected in the $|0\rangle$ or $|1\rangle$ state.

pulse to 5, which corresponded to a pulse with about a 4 ns Gaussian width, while the whole pulse took about 12 ns. I also needed to be careful that the total number of points that made up each pulse in the sequence was a multiple of 4. Otherwise, the pulses would not be in phase. The height of the pulses, h , was left as a variable for the calibration, and the period T was set to 4, which gave pulses with frequency equal to one quarter of the AWG clock frequency. I ran `ramseyLabVIEW.m` and check the pulse output on a scope.

Next, I needed to calibrate the pulse height so that each pulse was a $\pi/2$ -pulse. We wrote a LabVIEW program called *Run Scan N8241A ramsey vardelay.vi* that worked like *Run Scan*, but incorporated the Matlab code for communicating with the AWG. In the first column of the values array, I put a range of heights for the pulses. The heights were allowed to be between 0 (off) and 1000 (max). I also had to set the “period”, “width”, and “delay” in the upper right corner of the front panel. The period and width were described above, and the delay represented the number of DAC points between the pulses. For the calibration, I set this to 0. I set the timing of the measurement pulse to about 5 ns after the second pulse finished. I was not exactly sure when the tail of the second pulse was gone, and since T_l was much longer than a few ns, it was best to make sure that the whole of both pulses had a chance to act on the qubit before making the measurement. The pulse heights were in the first column and I used the “Repetitions” tab to set the program to take an S-curve for every pulse height. Later, I analyzed the data and found the pulse height that gave the highest s-curve, which meant that two pulses had combined to act like a π -pulse. Since each was the same height, they were both $\pi/2$ -pulses.

Next I looked at the S-curve of the state throughout the pulse train to make sure that state evolved correctly during the pulses and that it ended up in the $|1\rangle$ state. This data was acquired by using *Run Scan New* in a similar fashion to how I got the data for the T_1 measurements. The calibration of the pulses took some time, but it was necessary for a successful measurement.

The actual Ramsey measurement was a measurement of the S-curve of the state after it had been operated on by two $\pi/2$ -pulses that were separated by a delay time τ . First, I set the delay to zero and measured the S-curve. Then I increased the delay and measured the S-curve again. This required that the measurement pulse be moved the correct amount forward in time in order to measure the S-curve after the second pulse. The delay time τ was increased by multiples of the period of the pulses T so that the pulses remained in phase with each other.

We also wrote a new version of *Run Scan* called *Run Scan N8241A ramsey vardelay.vi* that would automatically set the number of delays. I set the width, period, and height (height is divided by 1000 before being put into *ramseyLabVIEW.m* in this case) in the upper right corner of the control panel. I specified column 0 of the “Values” array to represent the number of delays by setting the “Delays Column” control to 0 (the first column). I set the values in the “GPIO Enable” control to 1 and 2 for the measurement time and measurement pulse height, respectively. In the first row of the controls for GPIO prefix, suffix, *etc.*, I just put in some dummy command because the real communication with the AWG was done by LAN using Matlab script nodes.

Next, I set the measurement times and amplitudes. The high and low amplitude limits were chosen in the same manner as for a T_1 measurement. The times were set so

they followed along with the end of the second pulse as the delay increases. The amount of time that a delay represented was:

$$\Delta t_{\text{delay}} = \frac{N}{f_{\text{clock}}} \quad (7.19)$$

where N is the number of DAC points used in the delay (N must be divisible by the period, T , of the AWG signal). I determined the end of the pulse train by tracking the P_{sw} in *OneSRpercents* and finding the measurement timing where the pulses did not seem to be affecting P_{sw} . Using the “Fill” tab in *Run Scan*, I set the first time to this end point (plus 5 ns) and stepped it by the value determined by Eq. 7.19. Note that the DG535 only has a resolution of 5 ps, so I had to be careful when rounding the time step off. If I rounded to a number lower than that specified by Eq. 7.19, then the measurement pulse would slowly slip backwards in time as I increased the delay between the pulses. Therefore, I was sure to round up. This took care of the measurement timings.

Once everything was set up, I saved the settings to “settings [date]”, set the “Samples per Cycle” in *OneSRpercents.vi* to 2000, turned off the histogram display, and clicked the “Go” button in *Run Scan N8241A ramsey vardelay.vi*.

Figure 7.28 shows the data from three Ramsey measurements taken with different detunings. The plots are 2D color plots of the S-curves taken after pulses separated by an amount of time, τ . The detuning, $\Delta\omega$, sets up an oscillation of the final state that has a frequency ω_{Ramsey} such that:

$$P_1 \sim \cos(\omega_{\text{Ramsey}} \tau) \quad (7.20)$$

$$\omega_{\text{Ramsey}} = (\omega_{01} - \omega_{\mu}) = \Delta\omega \quad (7.21)$$




Fig. 7.28 False color plots of S-curves taken for Ramsey measurements in device PB9 that were done with microwave pulses of different detunings at $I_\Phi = 0.084$ mA, $I_b = 6.8$ μ A, and $f_{0I} = 4.218$ GHz. (a), (b), and (c) were done with 0, 12, and 30 MHz of detuning, respectively. The frequencies of the oscillations match the detuning. The oscillations decay with a Gaussian envelope. The white space in (c) is an area of no data and is present to keep the time scales of (a), (b), and (c) the same.

where ω_μ is the frequency of the microwave drive. Figure 7.28(a) shows the data when there is zero detuning, whereas 7.28(b) and 7.28(c) have 12 MHz and 30 MHz detuning, respectively. Note that 7.28(c) does not show very large oscillations. This is to be expected for a Ramsey measurement with a large detuning. The oscillations decay with a Gaussian envelope, which has a width of T_2^* .

7.14 Conclusions

From the simple measurement of the room temperature resistance to the complicated qubit manipulations performed using mixed signal microwave pulses, each measurement provided an important piece of information about the qubit. This chapter chronicled my procedures for performing these measurements. Chapters 8, 9, and 10 deals with the analysis of the data obtained through these methods.

Chapter 8

Extraction of Parameters for Device PB9

Device PB9 was designed and fabricated such that it should have had a very long coherence time. The theoretical predictions of its coherence put forth in Chapter 5 depended on the parameters of the device. In order to determine the actual values of the device parameters, I took several different measurements of the device's behavior and fit them to mathematical models. The I-V curve gave me a measure of the critical current of the SQUID. The I- Φ characteristic showed me how the switching current of the SQUID changed with applied flux. The flux applied to the SQUID modulated the switching behavior in such a way that it was possible to determine the inductances of the two branches of the SQUID and the critical currents of each junction. Finally, the energy spectrum of the device allowed me to determine the capacitances of the system. Also, since the fitting of both the energy spectrum and I- Φ produced estimates for the critical currents in the device, the values could be used to cross check each other. I then used these values for the parameters to update the expected relaxation time and coherence time of the device.

8.1 *AFM Images of the Josephson Junctions*

The model that I used for dissipation in the qubit depended on the capacitance of the junctions. The best estimate I had of the capacitance of the qubit junction was:

$$C_{J_i} = \frac{1}{40} A_{J_i} \text{ pF}/(\mu\text{m})^2 \quad (8.1)$$

where A_{J_i} is the area of the i^{th} junction. Because of charging on the insulating substrate, it was impossible to obtain images of the junctions using the SEM. Therefore, I used an atomic force microscope (AFM) to tap out the image of the junctions. Figure 8.1 shows the AFM images of the qubit junction and isolation. junction The junctions were made by double-angle evaporation, so a small overlap region forms the Josephson junction.

From Fig. 8.1(b) the dimensions of the isolation junction were approximately $1.1 \mu\text{m} \times 4.5 \mu\text{m}$. The rounded corners make for some uncertainty in the area. Also, the fact that there is some extra junction area where the second aluminum layer steps up onto the first. Since the Eq. 8.1 only provides an estimate for the capacitance, it would not help to determine the area of the junction too precisely anyway. In any case, I measured the area of the isolation junction A_{J_2} to be $5 (\mu\text{m})^2 \pm 0.5 (\mu\text{m})^2$. This gave an estimated capacitance $C_{J_2} = 0.2 \text{ pF} \pm 0.05 \text{ pF}$.

The dimensions of the qubit junction J_1 determined from Fig. 8.1(a) were $500 \text{ nm} \times 450 \text{ nm}$. This was larger than expected from other samples made with a similar fabrication process. Similar uncertainties applied to J_1 as to J_2 , and the value obtained for A_{J_1} was $0.23 (\mu\text{m})^2$, which gave a capacitance C_{J_1} of $5.8 \text{ fF} \pm 0.8 \text{ fF}$.

These values of the junction capacitances were obtained by examining the area of the junctions closely with an AFM. The other parameters needed to be extracted through electronic measurements and theoretical fits to the data.

8.2 The I-V Curve

I measured the I-V Curve of PB9 according to the procedure detailed in Chapter 7. Fig. 8.2 shows the data. The y-axis is the current I_b applied to the device through the

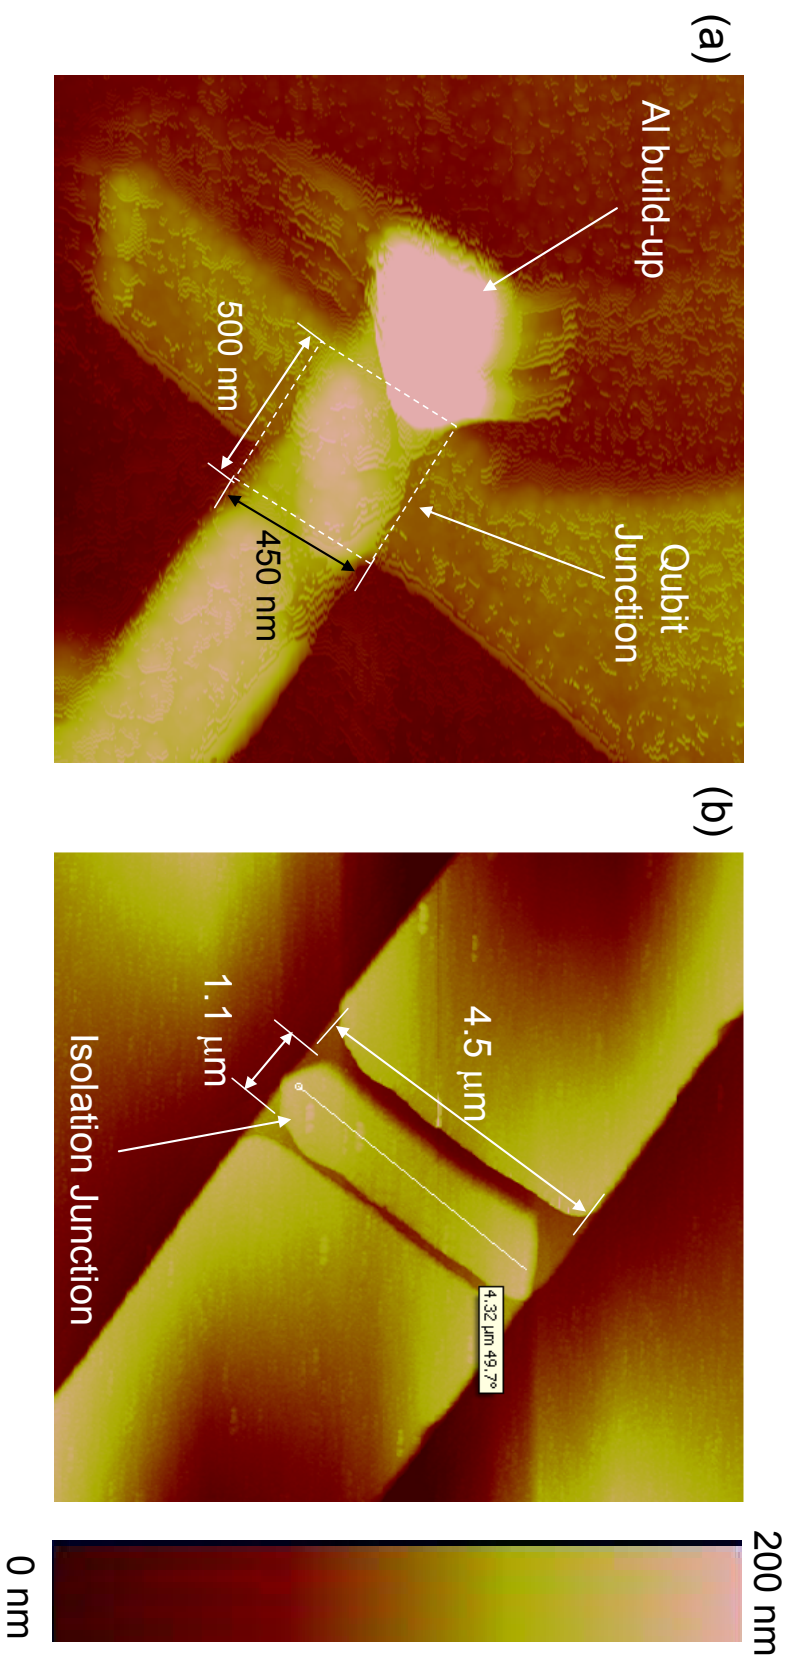


Fig. 8.1 (a) AFM picture of the qubit junction on device PB9. The junction had an approximate area $A_{JJ} \approx 0.23 (\mu\text{m})^2$. The bright spot to the left of the junction is Al that built up against the side-wall of the undercut resist during the second evaporation due to the steep angle (57°) of deposition. The color axis is height. (b) The isolation junction.

current bias line. The values were obtained by taking the voltage that was output from the source and dividing by 200 k Ω , which was the value of bias resistor. The x-axis is the voltage V across the device as measured on the voltage line at the top of the refrigerator. The value of the bias current at which the I-V switched to the voltage state was the critical current of the SQUID for applied flux $\Phi_a = 0$. Current to the SQUID divides between the two junctions such that:

$$I_b = I_1 + I_2 \quad (8.2)$$

where I_b is the bias current, I_1 is the current flowing to the qubit junction, and I_2 is the current flowing through the isolation junction. The maximum current for which the I-V curve will switch to the voltage branch is $I_b = I_{01} + I_{02}$ where I_{0i} is the critical current of the i -th junction. According to Fig. 8.2, the critical current of the device was about 8.7 μA .

The resistive branch of the I-V had a slope of about $R_n \approx 30 \Omega$. Note $2\Delta/eR_n \approx 12 \mu\text{A}$, which is roughly consistent with the measured critical current. The room temperature resistance of the device was around 500 Ω , so this value of 500 Ω must have been dominated by the resistance of the lines on the chip. This was not surprising since I used a two-point measurement that included the resistance of the spiral inductors.

The I-V curve was hysteretic because there was no resistive shunt across the SQUID. There was some strange behavior as the current returned to zero on the finite-voltage part of the graph (the quasi-particle branch). I expected this part to be almost vertical, but it showed numerous features, all of which were mirrored on the positive and negative current sweeps. These features were due to self-heating in the device that arose because the qubit junction was so small. They were not relevant to the coherence of the device, so I did not investigate them further.




Fig. 8.2 I-V curve of device PB9. The critical current of the SQUID was 8.7 μA . The gap voltage was $2\Delta/e = 360 \mu\text{V}$. The slope of the resistive branch was about 30Ω . The odd behavior in the quasiparticle branch is just due to self-heating in the device.

An I-V curve can give a good value for the critical current of the device. The fact that graph includes both positive and negative current sweeps made it possible to account for any offsets that would have given an inaccurate reading of the critical current. According to Fig. 1, the critical current of the device was about $8.7 \mu\text{A}$. There was some uncertainty in the critical current because there could be some flux offset from stray fields or trapped flux vortices near the device. Also, the switching current of the I-V was probably a bit lower than the sum of I_{01} and I_{02} because the state actually tunnels out of the potential well before I_b reaches $I_{01} + I_{02}$. Despite these uncertainties, it was a good first check on a key device parameter.

8.3 I - Φ Characteristics

The I- Φ curve of PB9 gave me a value for the critical current of the qubit junction and the inductances of the two branches of the SQUID loop. I measured the I- Φ according to the procedure in Chapter 7. The data is displayed in Fig. 8.3. This figure is a false color plot of the switching current histogram of the device for a range of applied fluxes. The y-axis is the switching current and it was obtained from a calibration of the current ramp:

$$I_{sw} = \frac{f_{ramp} t_{sw} V_b}{R_b} \quad (8.3)$$

where $f_{ramp} = 1 \text{ kHz}$ was the frequency of the ramp from the arbitrary waveform generator that sourced the bias current, t_{sw} was the switching time as determined by from the histogram, $V_b = 2\text{V}$ was the maximum amplitude of the voltage ramp, and $R_b = 200 \text{ k}\Omega$ was the bias resistor in the current bias line. The color axis is the number of switching




Fig. 8.3 I- Φ of device PB9 over a range of about $\pm 2\Phi_0$. Each vertical slice of the false color plot is a switching histogram measured at each setting for the flux bias current I_Φ . The magenta line is a fit to the data with the parameters $I_{0I} = 130$ nA, $I_{02} = 8.6$ μ A, $L_I = 0.7$ nH, $L_2 = 40$ pH, and $M = -4.95$ pH.

counts in a particular time bin of width 100 ns. The x-axis is the current that was applied to the flux line I_Φ , which was just the dc voltage supplied to the flux line divided by the flux bias resistor (10 k Ω).

One parameter that can be readily estimated by inspection of the I- Φ graph is the mutual inductance of the flux bias line to the SQUID. Because the response of the SQUID is periodic in the applied flux, the critical current is periodic with respect to the flux with a period of Φ_0 [49]. Note that the switching current in Fig. 8.3 displays a pattern that is periodic. By measuring the difference in flux bias between two equivalent points in this pattern, I was able to obtain the mutual inductance from:

$$M = \frac{\Phi_0}{\Delta I_\Phi} \quad (8.4)$$

where $\Phi_0 = h/2e = 2.07 \times 10^{-15} \text{ Tm}^2$, $\Delta I_\Phi = 0.42 \text{ mA}$ is the period in flux bias current of the I- Φ pattern. By inspection, I estimated that $M = 4.95 \text{ pH}$.

It was also possible to get a rough idea of the critical current of the qubit junction I_{01} by inspection of the I- Φ . The circulating current J in the SQUID loop cannot exceed I_{01} , the critical current of the qubit junction, because it has a smaller critical current than the isolation junction [49]; therefore, the extreme values of J were $\pm I_{01}$. Ignoring the flux from this circulating current, this means that the maximum current that I had to apply to switch the SQUID was $I_{02} + I_{01}$ and the minimum was $I_{02} - I_{01}$. Therefore, I_{01} was about half of the full range of $I_{b,sw}$ displayed in Fig. 8.3. Thus, I estimated I_{01} to be about 100-150 nA before I did any rigorous fitting.

The rest of the parameters were difficult to obtain because of the odd shape of the I- Φ . I expected the I- Φ to just show a smooth variation of the switching current; however,

the $I-\Phi$ of PB9 showed extra peaks in the average switching current of the device. The origin of these peaks was unclear, and I never fully understood them. However, it is likely that they were the result of some sort of suppression of the tunneling probability of the state from the potential well at certain frequencies or flux bias points. These features could be evidence of interference between equally probable tunneling paths that arise because of the two dimensional nature of the potential or points where noise was suppressed.

To get a more accurate estimate of the device parameters, I fit the $I-\Phi$ curve to the expected $I-\Phi$ curve of a SQUID. To find the expected $I-\Phi$, I note that if the current through isolation junction I_2 exceeded I_{02} , then the whole device would switch to the voltage state [28]. I_2 was determined by the phase across the isolation junction γ_2 through the equation:

$$I_2 = I_{02} \sin \gamma_2 \quad , \quad (8.5)$$

At each applied flux, I minimized the SQUID potential U , found the value of γ_2 , calculated I_2 , and then found the value of I_b that it would have taken to increase I_2 from that value to I_{02} . The solution was:

$$I_{b,sw} \approx I_{02} - I_{02} \sin \gamma_{2,min} \quad (8.6)$$

where $\gamma_{2,min}$ is the value of γ_2 at the potential minimum. Note that this is an approximate method for finding the switching current. The potential changes as the bias current is increased toward the critical current, so the approximate method might have missed some of the SQUID's behavior. Therefore, I checked the results of the approximate method against another method where I slowly increased the bias current, minimizing the potential at each step, until $\gamma_2 = \pi/2$ to find the switching bias. This more exact treatment

gave nearly identical results for the I- Φ once I accounted for a flux offset. The approximate method took a small fraction of the time, though. Regardless, both techniques depended entirely on minimizing the correct potential. If the parameters of the potential (i.e. I_{01} , I_{02} , L_1 , L_2 , and M) were correct, then the values obtained for $I_{b,sw}$ would mimic the behavior of the switching histograms with respect to applied flux current.

To make this approximate method more accurate, I needed to choose the I_b that I used in the calculation. Consider the potential:

$$U(\gamma_1, \gamma_2, I_b, I_\Phi) = -E_{J1} \cos \gamma_1 - E_{J2} \cos \gamma_2 - \frac{\Phi_0}{2\pi} \frac{I_b}{L} (L_2 \gamma_1 + L_1 \gamma_2) + \left(\frac{\Phi_0}{2\pi} \right)^2 \frac{1}{2L} \left(\gamma_1 - \gamma_2 - \frac{2\pi}{\Phi_0} M I_\Phi \right)^2 \quad (8.7)$$

where $E_{Ji} = \frac{\Phi_0}{2\pi} I_{0i}$ is the Josephson energy of the i^{th} junction and $L = L_1 + L_2$ is the total geometrical inductance of the SQUID loop. If $I_b = 0$, then the potential would not depend on L_1 and L_2 independently. Furthermore, the data in Fig. 8.3 represents the tunneling of the ground state out of the potential well at a specific bias current. The shape of the potential looks quite different when $I_b \approx I_{02}$ then when $I_b = 0$. Hence, when I performed the fit, I set $I_b \approx 0.8 (I_{01} + I_{02})$. This increased the potential's dependence on L_1 and L_2 . Also, setting the bias current in the simulation to a relatively high value helped to better mimic the conditions the qubit was under when I captured the I- Φ .

Figure 8.4 shows a close-up of the measured I- Φ and a fit that was obtained with the approximate technique described above. The fit was obtained using, $I_{01} = 130$ nA,




Fig. 8.4 Close up of I- Φ of device PB9 with the same fit as in Fig. 8.3, plotted as a dotted magenta curve. The close up view emphasizes the odd peaks in $I_{b,sw}$. These peaks make it difficult to determine the inductances from a traditional I- Φ measurement.

$I_{02} = 8.6 \mu\text{A}$, $L_I = 0.7 \text{ nH}$, $L_2 = 40 \text{ pH}$, and $M = -4.95 \text{ pH}$. Note that the odd shape of the measured $I\text{-}\Phi$ is clearly different than the fit.

One parameter that I wanted to determine from the $I\text{-}\Phi$ is L_I . The size of L_I determines the strength of the isolation that was discussed in Chapter 4. Also, the ratio of L_I to L_{JI} determines whether or not the device can be described by a single junction model (see Sect. 4.3-4.4). The slope of the left side of the $I\text{-}\Phi$ is dependent on the inverse of L_I [38], and the odd shape of the $I\text{-}\Phi$ on that side, which was likely caused by anomalous tunneling behavior, made it difficult to accurately determine L_I . Therefore, I used a less conventional measurement of the $I\text{-}\Phi$ to determine L_I . One way to suppress the odd tunneling behavior was to quickly pulse the current to the SQUID to a point where the state would leave the well. This technique amounted to measuring the S-curve of the device at each applied flux.

Figure 8.5 shows a false color plot of the S-curves vs. flux. The y-axis is the amplitude of current pulse applied to the device. The color corresponds to the probability of switching. To construct the plot, I mapped the currents found from the simulation to the measurement pulse heights V_p of the data by assuming that:

$$V_{p,\text{max}} \rightarrow I_{02} - I_{2,\text{min}} \quad (8.8)$$

where $V_{p,\text{max}}$ is the V_p that caused the device to switch nearly 100% of the time at the top of the $I\text{-}\Phi$ ($I_\Phi = -0.9 \text{ mA}$), $I_{2,\text{min}}$ is the minimum value obtained for I_2 as determined by the values for γ_2 found during the simulation. Thus every current could be related to a voltage pulse height and vice versa by using the ratio of the two quantities in Eq. 8.8. This method for determining the pulse current was ad hoc and the range of the swing of the switching current on the axis determined by Eq. 8.8 depended on the bias current I_b




Fig. 8.5 False color plot of S-curves taken over one period of flux. The y-axis is a loosely determined value for the current delivered by the measurement pulse I_p . The current pulse is quick, (~ 2 ns) and the odd tunneling behavior is still apparent at low switching probabilities (light blue), but not when the switching probability is near 100%. The dotted magenta line is a fit with the parameters: $I_{01} = 130$ nA, $I_{02} = 8.6$ μA , $L_1 = 0.7$ nH, $L_2 = 40$ pH, $M = -4.95$ pH, and I_b was set to 6.8 μA .

assumed during the fit. Also, it assumed a linear relationship between V_p and I_p , and this was not necessarily true because the size of the current pulse that gets to the qubit depended on the rise time of the voltage pulse, which was not constant over the full range of V_p [52]. It should be possible in future experiments to use the S-curves of the I- Φ in order to get a more accurate picture of the device parameters if a rigorous calibration of I_p is performed. However, this fit was meant as a check of the correct value for L_I , not a check of the critical currents of the device. This ad hoc axis conversion from V_p to I_p allowed me to compare the fits to the plot.

Figure 8.6(a) shows the I- Φ measured by the ramped current method along with two fits, and Fig. 8.6(b) shows the same two fits on the S-Curves plot. These are the same magenta line fits shown in Figs. 8.4 and 8.5. The orange curve in Fig. 8.6(a) is a fit using parameters that were determined assuming that the slope of the left side of the I- Φ fit in Fig. 8.6(a) needed to match the shape of the I- Φ more exactly. The parameters for these fits were the same except that for the magenta fit $L_I = 0.6$ nH, whereas for the orange fit $L_I = 2$ nH. Note that in Fig. 8.6(a) the orange fit followed the shape of the I- Φ better than the magenta fit; however, Fig. 8.6(b) reveals that the orange fit with $L_I = 2$ nH had the wrong shape when I quickly pulsed the current. Fig. 8.6(b) shows that the magenta fit with $L_I = 0.7$ nH more accurately reproduced the slope of the left side of the I- Φ . From this I concluded that, L_I was closer to 0.7 nH than 2 nH.

8.4 Energy Spectrum

The energy spectrum was a very revealing and interesting piece of data for device PB9. In particular, the spectrum was periodic with respect to applied flux and the




Fig. 8.6 (a) Comparison of fits to the pulsed I- Φ using $L_J = 2$ nH (orange curve) and $L_J = 0.7$ nH (magenta curve). (b) The odd shape of the ramped I- Φ created some uncertainty as to whether I should fit the I- Φ according to the sharp slope that it displays on the left side. The magenta fit with $L_J = 0.7$ nH is more representative of the actual critical currents and inductances of device PB9.

resonance frequency never went to zero (see Fig. 8.7). This behavior led us to do a more thorough analysis of the SQUID Hamiltonian, as discussed in Chapter 4. In this section I explain why a single junction phase qubit could not have produced an energy spectrum like that shown in Fig. 8.7. I will also show how I used the normal modes model to determine the remaining device parameters ($C_1 = C_{J1} + C_x$ and $C_2 = C_{J2}$), and, how I used a finite difference method on the full SQUID Hamiltonian (see. Sect. 4.5) to check the calculation from the normal modes model. In addition, I found that in order to capture the entire behavior of the qubit resonance frequency with respect to flux, I had to account for the parasitic inductance of the interdigitated capacitor.

I measured the energy spectrum over several days using the procedure detailed in Sect. 7.8. Figure 8.7 is a compilation of several data series. The y-axis is the frequency of the microwave drive f_μ . The x-axis is the flux bias current I_Φ and it covers about 1.5 periods. The flux bias resistor was 10 k Ω . The color scale represents the probability P_{sw} of the device switching during the measurement pulse (*i.e.* the number of switches during the measurement pulse divided by the number of trials). Each vertical slice of the 2D plot is a single sweep through the resonance at fixed applied flux. The bright red color marks value of the qubit's $|0\rangle$ to $|1\rangle$ transition frequency f_{01} . The frequency was only swept over a small range near the resonance peak in order to save time; the white area on the graph is where no data was taken. The power level of the microwave drive was set such that P_{sw} was 10-20% at the peak of each resonance sweep (so power broadening should be minimal). The data was plotted with the average P_{sw} of each resonance sweep subtracted away in an effort to keep the color scale uniform.




Fig. 8.7 Energy transition frequency spectrum of device PB9. Each vertical slice is a sweep of the microwave frequency f_μ at a fixed flux bias current I_ϕ . The color represents the probability of switching P_{sw} during the measurement pulse at each microwave frequency minus the average P_{sw} for the sweep at each flux bias. $I_b = 6.8 \mu\text{A}$.

Several interesting aspects of the spectrum were readily apparent upon inspection. First, there were very few splittings in the spectrum. There was one clear splitting at $f_{01} = 3.86$ GHz and one very small one at $f_{01} = 4.53$ GHz. Second, the spectral peaks were exceptionally narrow in the 4.0-4.4 GHz range. Both of these observations will be discussed in detail in Chapter 9. Also, there was a flat region in the spectrum where $df_{01}/d\Phi_a = 0$, which pointed to the existence of a sweet spot in flux. This will be discussed in Chapter 10. Here I will focus on another interesting aspect of the spectrum: the fact that the resonant frequency never went to zero.

Consider the equation for the plasma frequency of the qubit junction:

$$\omega_p = \frac{1}{\sqrt{L_{J1}C_1}} \quad (8.9)$$

where C_1 is the capacitance in parallel with the Josephson inductance L_{J1} . In the case of device PB9, C_1 was approximately equal to the capacitance of the interdigitated capacitor C_x because the capacitance of the junction C_{J1} itself was very small (< 10 fF). Throughout this discussion they will be referred to together as C_1 . If I substitute the equation for the Josephson inductance (Eq. 2.20) into the Eq. 8.9, then the plasma frequency, which is a good estimate for ω_{01} , is:

$$\omega_p = \frac{1}{\sqrt{\frac{\Phi_0}{2\pi} C_1}} (I_{01}^2 - I_1^2)^{1/4} \quad (8.10)$$

According to Eq. 8.10, if $I_{01} = I_1$, then $\omega_p = 0$. Considering my explanation of SQUIDs in Chapter 4, and the discussion of the I- Φ in the previous section, it is clear that at some

applied flux, the circulating current in the SQUID would equal I_{0I} . Yet, there was no point on the spectrum in Fig. 8.7 where the qubit frequency even approached zero.

Furthermore, the resonance frequency of the device was far too high to be accounted for by a single junction with a critical current of about 100 nA. Considering the relation of the ω_p to L_{JI} , I_{0I} , and C_I , in order for the device to have had a resonance frequency of 4 GHz, the capacitance would have had to be about 500 fF, which was only a third of the expected value (see Sect. 5.3). This suggested that the single junction model was inadequate for describing the behavior of the qubit in device PB9. It was these considerations that led to the examination and development of the theory detailed in the last half of Chapter 4.

8.5 *Fitting the Energy Spectrum with Different Models*

In order to extract the remaining parameters of the device from the energy spectrum, I fit my measurements using two different models. First, I used the normal modes model detailed in Sects. 4.3 and 4.4. Then, since the normal modes were only approximations to f_{0I} transition frequencies—in much the same way that ω_p is only an approximation of ω_{0I} in the single junction model—I confirmed the parameters with the finite difference method detailed in Sects. 4.5 and 4.6. The two models gave similar results, as expected.

The parameters affect the dependence of the resonance on applied flux to different degrees and in different ways. For example, the mutual inductance M sets the period of the spectrum just as it did for the I- Φ . I_b , I_{02} , and L_2 had little effect on the overall spectrum other than changing the offset in flux. Also, the value of the isolation junction

capacitance C_2 had virtually no effect on the fit. M , I_b , I_{02} , and L_2 were initially set to the values determined by fitting the I- Φ . C_2 was estimated from the area of the isolation junction according to Eq. 8.1.

The three parameters that influenced the fit the most were the critical current of the qubit junction I_{01} , the capacitance across the qubit junction C_1 and the inductance of the qubit branch of the SQUID L_1 . Changing I_{01} changed the range of the oscillation of f_{01} because the Josephson inductance of the qubit junction had the most influence on the qubit frequency when it was smallest.

Consider the frequency of the qubit mode:

$$\omega_{\pm} = \sqrt{\frac{1}{2} \left[\left(\frac{1}{L_1^* C_1} + \frac{1}{L_2^* C_2} \right) \pm \sqrt{\left(\frac{1}{L_1^* C_1} + \frac{1}{L_2^* C_2} \right)^2 - 4 \left[\left(\frac{1}{L_1^* C_1} \right) \left(\frac{1}{L_2^* C_2} \right) - \left(\frac{1}{L C_1} \right) \left(\frac{1}{L C_2} \right) \right]} \right]} \quad (8.11)$$

where the effective inductance shunting C_1 is:

$$\frac{1}{L_1^*} = \left(\frac{1}{L_{J1}} + \frac{1}{L} \right) \quad (8.12)$$

Here $L = L_1 + L_2$ was about 0.7 nH and L_{J1} was always larger than that in magnitude. With a critical current of about 100 nA, L_{J1} was about 3 nH when $I_I = 0$. Using the potential minimization method described in Sect. 8.2, I found the phase of the qubit junction γ_1 over the full range of flux bias. $\cos \gamma_1$ and L_{J1} are plotted in Fig. 8.8 vs. I_{Φ} .

Note that when $-\pi < \gamma_1 < -\frac{\pi}{2}$ or $\frac{\pi}{2} < \gamma_1 < \pi$ then:

$$L_{J1} = \frac{\Phi_0}{2\pi} \frac{1}{I_{01} \cos \gamma_1} < 0 \quad (8.13)$$

Note that when $L_{J1} < 0$, then Eq. 8.12 gives an L_1^* larger than when $L_{J1} > 0$. Thus ω_{01} was at its lowest when $\gamma_1 = \pi$. At $\gamma_1 = 0$, L_1^* was at its smallest because L_{J1} was at its




Fig. 8.8 The Josephson inductance of the qubit junction L_{J1} over the range of the energy spectrum. Note that when $\cos(\gamma_1)$ is negative, L_{J1} is also negative. Also, when $\cos(\gamma_1) = 0$, L_{J1} goes to infinity.

smallest positive value and this causes ω_{0I} to be at its maximum. Changing I_{0I} directly affected the extreme values of L_I^* , which set the minimum and maximum qubit frequencies.

Interestingly, changing I_{0I} does not change the resonance frequency at the bias points where $\gamma_I = \pm\pi/2$. At these points, L_I^* is independent of L_{JI} because L_{JI} is infinite. These bias points are where ω_{0I} is near its median value. This realization led to the identification of sweet spots in the energy levels with respect to critical current noise, as I discuss in detail in Chapter 10.

The dependence of the qubit frequency on the capacitance C_I was not complicated. A larger C_I decreased the resonance frequency, and a smaller C_I increased it.

Changing L_I had the effect of moving the whole fit up or down in frequency, as well, but it had the added effect of slightly changing the minimum and maximum frequency. When $\gamma_I = 0$ or π , though, L and L_{JI} were as close as they could be in magnitude. Since L and L_{JI} added as inverses to determine L_I^* , if I increased L_I , it would increase the amount of influence that L_{JI} had on L_I^* at the top and the bottom of the spectrum. Therefore, a small increase in L_I would cause the maximum frequency to get a little higher and the minimum frequency to get a little lower, as well as move the whole fit down some.

Figure 8.9 shows the energy spectrum and the best fit I obtained using the normal modes model. The parameters used for this fit were: $I_{0I} = 110$ nA, $I_{02} = 8.6$ μ A, $L_I = 0.7$ nH, $L_2 = 40$ pH, $C_I = 1.56$ pF, $C_2 = 0.225$ pF, $M = -4.95$ pH, and $I_b = 6.8$ μ A (see Appendix A for Matlab Code). The capacitance obtained for C_I was very close to the expected value obtained from Eqs. 5.15 and 5.16. The model did a good job, but did not




Fig. 8.9 Fit of the energy spectrum (magenta curve) using the normal modes model from the cubic approximation to the SQUID potential. The parameters of the fit were: $I_{0I} = 110$ nA, $I_{02} = 8.6$ μ A, $L_I = 0.7$ nH, $L_2 = 40$ pH, $C_I = 1.56$ pF, $C_2 = 0.225$ pF, $M = -4.95$ pH, and I_b was set to 6.8 μ A. The fit captures the general behavior of the spectrum, but there is a noticeable discrepancy along the sides of the spectral humps (*i.e.*, $4.0 < f_{01} < 4.4$ GHz).

quite capture the behavior of the resonance near the lowest frequencies. Also, the value for I_{01} was somewhat lower than I found from the fits to the I - Φ . Overall, the results of this fit were promising, but not completely accurate.

As I noted previously, the normal modes do not exactly give f_{01} , But rather the resonant modes of the approximate simple harmonic oscillator modes. It is possible to include the cubic terms in the approximate potential (Eq. 4.18) to more accurately determine the perturbations to the normal modes, but it requires a lot of algebra and still only gives the lowest order correction to the energies. Instead, I used the finite difference method described in Sects. 4.5 and 4.6. The finite difference method gave the energies of the full SQUID Hamiltonian, and I found it to be easier to automate and combine with the fitting code (see Appendix A).

Figure 8.10 shows the energy spectrum and the best fit I obtained using the finite difference method on the full SQUID Hamiltonian. I used a 39×39 grid in γ_1 and γ_2 with a grid spacing of 0.03 radians in each direction. As a test, I checked the accuracy of the fit against a fit using a 49×49 grid with a spacing of 0.25 radians and found that the values of f_{01} obtained differed by about 0.2%. The fit using the 49×49 grid may have been a little more accurate, but it took 15-20 minutes whereas the 39×39 fit took only about 2-3 minutes. If I were looking for the wavefunction and transition frequencies for higher levels, a larger grid would have been appropriate, but since I only needed f_{01} , the smaller, faster simulation was adequate.

The parameters of the best fit were: $I_{01} = 110$ nA, $I_{02} = 8.6$ μ A, $L_1 = 0.7$ nH, $L_2 = 40$ pH, $C_1 = 1.52$ pF, $C_2 = 0.225$ pF, $M = -4.95$ pH, for $I_b = 6.8$ μ A. The only difference I found between this fit and the normal modes fit was that the capacitance C_1 shunting the




Fig. 8.10 Fit of the energy spectrum (magenta curve) using the finite difference method on the full SQUID Hamiltonian. The parameters of the fit were: $I_{01} = 110$ nA, $I_{02} = 8.6$ μ A, $L_1 = 0.7$ nH, $L_2 = 40$ pH, $C_1 = 1.52$ pF, $C_2 = 0.225$ pF, $M = -4.95$ pH, and I_b was set to 6.8 μ A.

qubit junction needed to be about 1% smaller to get things to line up. This was not surprising because one would expect that the energy levels in a cosine potential well would get closer and closer spaced as they got higher and higher in the well; therefore, the first excited energy level would be a little closer to the ground state energy than would be predicted in a parabolic well. When I fit the spectrum using the finite difference method with the same parameters that worked for the normal mode model, the f_{01} 's came out a little lower than in the normal modes model. I had to decrease C_I in order to bring the fit up to match the data. Overall, the finite difference method confirmed the results of the normal modes fit.

The fit from the finite difference method also had the same deficiencies that the normal modes fit had. The fit generally matched the behavior of the data, but it did not quite capture the shape of the data for frequencies below 4.4 GHz.

There was one aspect of the actual device that was not accounted for in either model: the parasitic inductance of the interdigitated capacitor (IDC). In both models, I assumed that only a capacitance was shunting the qubit junction. In fact, the IDC had finite dimensions that led to an inductance being in series with the capacitive circuit element. Figure 8.11 shows a schematic of the SQUID loop where I have included the parasitic inductance of the IDC. This model more accurately reflected the actual device.

One complication of this model is that it introduced a new degree of freedom to the system. In the original schematic, there were two Josephson junctions and effectively just two capacitors (C_x and C_{JI} could be combined in parallel). In the new model there are three capacitors: the capacitance of the isolation junction C_{J2} , the IDC C_I , and the capacitance of the qubit junction C_{JI} . However, since C_{JI} was much smaller than C_I , I

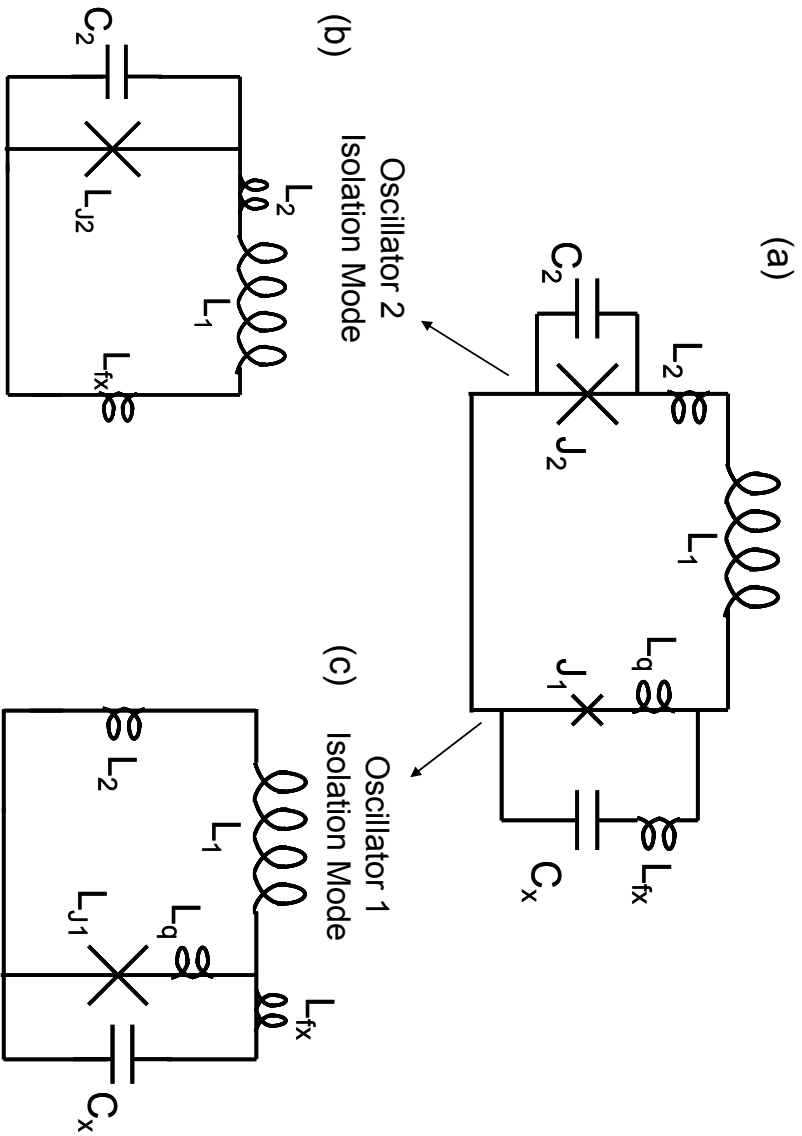


Fig. 8.11 (a) Schematic of the device when the parasitic inductance of the interdigitated capacitor (IDC) is included. L_{fx} is the inductance of the IDC structure, and L_q is the inductance of the wire leading to the qubit junction between the points of contact of the IDC to the SQUID loop. b) and c) show the approximate schematics of the two coupled harmonic oscillators that account for the normal modes of the SQUID with the parasitic inductances L_{fx} and L_q included.

left it out of the model and assumed that the qubit junction was a pure Josephson element L_{JI} . The approximate model for the system was two coupled harmonic oscillators as depicted in Fig. 8.11(b) and (c). The parasitic inductance altered the L_i^* 's from Eq. 8.12 such that:

$$L_1^* = L_F + \left(\frac{1}{L_{J1} + L_q} + \frac{1}{L_1 + L_2} \right)^{-1} \quad (8.14)$$

$$L_2^* = \left(\frac{1}{L_{J2}} + \frac{1}{L_1 + L_2 + L_{fx}} \right)^{-1} \quad (8.15)$$

where L_q was the geometrical inductance of the line that led to the qubit between the points where the IDC made contact with the SQUID loop and L_{fx} was the parasitic inductance of the IDC's fingers along with the path lengths leading out to the fingers. The model assumed that the junction inductances presented a much larger impedance to signals at f_{0I} than the capacitors did, which is why L_{JI} was not included in L_2^* and vice versa. I found the minimum of this potential in order to get γ_1 and γ_2 :

$$U(\gamma_1, \gamma_2, I_b, I_\Phi) = -E_{J1} \cos \gamma_1 - E_{J2} \cos \gamma_2 - \frac{\Phi_0}{2\pi} \frac{I_b}{L} \left[L_2 \gamma_1 + (L_1 + L_q) \gamma_2 \right] + \left(\frac{\Phi_0}{2\pi} \right)^2 \frac{1}{2L} \left(\gamma_1 - \gamma_2 - \frac{2\pi}{\Phi_0} M I_\Phi \right)^2 \quad (8.16)$$

where $L = L_1 + L_2 + L_q$. Equation 8.16 does not depend on L_{fx} because I_b and I_Φ are assumed to be dc currents, which would not flow through L_{fx} . I calculated L_{JI} and L_{J2} using γ_1 and γ_2 . Then I found the qubit mode frequency from ω in Eq. 8.11 using Eqs. 8.14 and 8.15 for L_1^* and L_2^* . I made a critical assumption about the system in eliminating the degree of freedom that the C_{JI} presented, but the model that resulted was simpler and more intuitive.

Figure 8.12 shows the effect of adding the parasitic inductance to the model. The fit to the data was much better than with the other models. The parameters of the fit were: $I_{0I} = 130$ nA, $I_{02} = 8.6$ μ A, $L_I = 0.63$ nH, $L_q = 0.07$ nH, $L_F = 0.07$ nH, $L_2 = 40$ pH, $C_I = 1.52$ pF, $C_2 = 0.225$ pF, $M = -4.95$ pH, and $I_b = 6.8$ μ A. Note that $L_I + L_q = 0.7$ nH, the same inductance of the qubit branch determined in the previous fits, and $I_{0I} = 130$ nA was the same as I determined from the fit of the I- Φ , unlike the values determined in the other energy spectrum fits. I tried many values for L_q and L_F , and these worked the best. It was odd that L_q did not come out larger than 70 pH. A general rule for estimating the inductance of a microstrip line is that they will have an inductance per unit length of 1 pH/ μ m. This would have given $L_q \approx 300$ pH. However, it is possible that L_q was lower than expected because of the presence of the IDC fingers nearby [72]. Furthermore, a smaller L_q might make sense because it could account for the fact that the total loop inductance came out much lower than the design value of 1.25 nH quoted in Chapter 5. The architecture of the IDC was too complicated to check the value of L_{fx} so I did not have a way to verify the accuracy of the number used in the fit. Regardless, the purpose of adding the parasitic inductances was not to determine their values, but to see if their presence could account for the discrepancies between the original normal modes model and the data. It did; adding the inductances L_{fx} and L_q to the resonant circuit slightly changed the shape of the fit to match the energy spectrum data (see Fig. 8.12).

Unfortunately, adding the parasitic inductances to the finite difference method of solving the full SQUID Hamiltonian was difficult. In particular, this would require adding an additional degree of freedom, one for each capacitance.




Fig. 8.12 Fit of the energy spectrum (magenta curve) using the normal modes model with the parasitic inductance of the IDC included. The parameters of the fit were: $I_{01} = 130$ nA, $I_{02} = 8.6$ μ A, $L_1 = 0.7$ nH, $L_2 = 40$ pH, $L_q = 0.07$ nH, $C_I = 1.52$ pF, $C_2 = 0.225$ pF, $M = -4.95$ pH, and I_b was set to 6.8 μ A. The addition of the parasitic inductances allowed the fit to capture the behavior of the spectrum.

8.6 Comparison of Extracted Parameters to Design Parameters

Table 8.1 shows the extracted parameters and their errors along side the design parameters. The parameters determined for the device were: $I_{01} = 130 \pm 20$ nA, $I_{02} = 8.6 \pm 0.05$ μ A, $L_I = 0.7 \pm 0.1$ nH, $L_2 = 40 \pm 20$ pH, $C_x = 1.51 \pm .05$ pF, $C_{J1} = 6 \pm 1$ fF, $C_{J2} = 0.225 \pm .05$ pF, $M = -4.95 \pm .03$ pH, $L_{fx} = 70 \pm 30$ pH, $L_q = 70 \pm 30$ pH. The errors in the values represent the range of values that would also give good fits to the measurements if the other values were also manipulated within their acceptable ranges. I_{02} was rather strictly defined because of the I-V curve measurement. M was strictly determined by the period of both the spectrum and the I- Φ . L_2 and C_2 had wide margins of error because none of the fits to the measurements depended very strongly on these values. I_{01} , L_I , and $C_I = C_x + C_{J1}$, could be combined in many ways to give good fits for the spectrum and the I- Φ , so there were rather wide and correlated margins of error in these values.

Most of the parameters ended up being close to the design values. The biggest difference was in the loop inductance L_I . The fits to the data gave $L_I = 0.7$ nH, even though the modeling said that it should have been 1.25 nH. This may have been because the presence of the IDC was not accounted for in the models. Also, the mutual inductance M of the flux line to the SQUID was larger than intended. This meant that the dissipation from the flux bias line would have been larger than expected. The critical currents ended up being a little lower than expected. The area of the qubit junction was larger than desired, which increased its contribution to the dissipation because of the additional lossy AlO_x. Note also that the frequency of the qubit was much higher than expected from the single junction model. This increased the estimates for the loss from the dielectric because they went like $1/f_{01}$.

Table 8.1 Extracted qubit parameters, CJ1 and CJ2 were extracted from AJ1 and AJ2

(which were determined from AFM images) and Eq. 8.1. Range of f_{01} was found at

$$I_b = 6.8 \mu\text{A}$$

Parameters	Designed	Extracted
I_{01} (nA)	150	130 ± 20
I_{02} (μA)	10	8.6 ± 0.05
L_{J1} (nH)	2.2	2.5 ± 0.3
L_{J2} (nH)	0.03	0.04 ± 0.005
L_1 (nH)	1.25	0.7 ± 0.1
L_2 (nH)	0.02	0.04 ± 0.02
A_{J1} (μm^2)	0.1	0.23 ± 0.03
A_{J2} (μm^2)	7	5 ± 0.5
C_{J1} (fF)	2.5	6 ± 1
C_{J2} (pF)	0.18	0.2 ± 0.05
C_x (pF)	1.5	1.51 ± 0.05
C_f (pF)	80	80
L_f (nH)	10	10
L_p (nH)	0.1	N/A
M (pH)	1	4.95 ± 0.03
C_μ (fF)	1	N/A
f_{01} (GHz)	2.4	3.86-5.06
L_{fx} (pH)	N/A	70 ± 30
L_q (pH)	N/A	70 ± 30

8.7 Conclusions

In summary, by analyzing the I-V, the I- Φ , and the energy spectrum of device PB9, I was able to determine its parameters. The I-V gave me an idea of the total critical current of the device. The I- Φ was difficult to fit to the theory because of some odd tunneling behavior of the SQUID. I used two techniques to measure the I- Φ : slowly ramping the current to find the switching current and quickly pulsing the current applied to the device. The slow ramp method gave me a more accurate measure of the current applied to the device and allowed me to determine the critical current of the qubit junction I_{0I} more accurately. The pulsed current method suppressed the odd tunneling effects. The current scale of the pulsed measurement was difficult to determine, but it gave me a more accurate determination of the inductance of the qubit branch L_I than I would have gotten from the slow ramp method.

I found that a single junction model was inadequate for describing the behavior of the energy spectrum of the qubit. The normal modes model of the SQUID predicted a spectrum that was similar to what I measured, and a finite difference method on the full SQUID Hamiltonian confirmed that the ω_{0I} of the qubit was well-approximated by the normal modes. Adding some parasitic inductance from the interdigitated capacitor to the normal modes model improved the fit, and in the end I found good fits. Going through the fitting procedure gave me insight into the workings of the device and the effects that each parameter had on the device's behavior.

Chapter 9

Dissipation Analysis of Device PB9

In Chapter 5, I described how I designed device PB9 to be protected from dissipation. The qubit junction was made very small ($\sim 0.1 \mu\text{m}^2$) to reduce dielectric loss from the AlO_x tunnel barrier. The junction was also shunted by an interdigitated capacitor because it was the lowest loss capacitor architecture available [56]. Similarly, I chose the parameters of the device to limit the loss from known sources of dissipation. However, the decoherence model I used for optimizing the design was flawed in that it assumed that the qubit was a single junction. In Chapter 8, I showed that the single junction model was inadequate for describing the energy levels of device PB9 and that one needed to find the normal modes of the SQUID. In Sect. 9.1 of this chapter, I modify the dissipation model to account for the normal modes and present a revised theoretical estimate for the energy relaxation time T_1 of the device.

The spectrum of device PB9 only showed one obvious splitting, and closer analysis revealed a second small splitting at around $f_{01} = 4.53 \text{ GHz}$. With so few splittings, the loss from coupling to discrete two-level systems should be negligible except near the splitting frequencies.

The longest T_1 that I obtained from an energy relaxation measurement was $T_1 = 350 \pm 70 \text{ ns}$. The large margin of error on T_1 came from the fact that the extracted values of P_1 were not entirely reliable because I needed to extract P_1 from anomalous S-curves. I discuss this in Sect. 9.4 of this chapter. Nevertheless, the T_1 values I found were much lower than the theoretical predictions.

In order to determine what was limiting the lifetime of the qubit, I measured the T_1 at bias points across the range of the energy spectrum. I expected that the study would reveal a dependence of T_1 on f_{01} that would point to a specific source of decoherence. In fact, as I discuss in Sect. 9.5, this study revealed that there was a correlation between T_1 and the coupling strength of the microwave drive line to the qubit. When the microwave line was weakly coupled to the qubit, T_1 was longer; therefore, it was likely that dissipation from the impedance of the microwave line was causing the short T_1 's of the qubit.

9.1 *Modified Dissipation Model in the Normal Modes Picture*

In Chapter 5, I detailed the model I used for selecting the parameters of SQUID phase qubit PB9. For this design model I assumed that the qubit was a single Josephson junction and all of the dissipation was represented by an effective resistance shunting it. The model predicted a long coherence time (about 8.4 μ s). However, as I discussed in Chapter 8, I found that a single junction model was inadequate for describing the qubit. This made it necessary to reconsider the estimates for the dissipation in light of the normal modes picture of the SQUID.

Let me first recall the model for dissipation in the single junction case. The rate of spontaneous emission caused by a resistor R_{eff} in the single junction model is (see Sect. 3.4 and [26,39]):

$$\Gamma_{10} = \frac{1}{T_1} = \frac{2\hbar\omega_{01}}{R_{eff}(\omega_{01})} \frac{|\langle 0 | \gamma_1 | 1 \rangle|^2}{4e^2} \quad (9.1)$$

where ω_{0I} is 2π times the qubit transition frequency, γ_I is the phase operator of the qubit junction, and $R_{eff}(\omega_{0I})$ is the effective resistance shunting the junction at ω_{0I} . If I assume

that $|0\rangle$ and $|1\rangle$ are harmonic oscillator states and write $\gamma_1 = \sqrt{\frac{\hbar}{2m\omega_{01}}}(a^\dagger + a)$ where

$m = \frac{\hbar^2}{4e^2}C_1$, then $\Gamma_{10} = \frac{1}{R_{eff}C_1}$ and $T_I = R_{eff}C_I$ [41]. This is the theoretical basis for the

model described in Chapter 5.

I now consider the case where there are two coupled junctions, J_1 and J_2 , each shunted by its own effective resistance $R_{eff,1}$ and $R_{eff,2}$, respectively. The rate of spontaneous emission in this case would be:

$$\Gamma_{10} = \frac{2\hbar\omega_{01}}{4e^2} \left[\frac{|\langle 0|\gamma_1|1\rangle|^2}{R_{eff,1}} + \frac{|\langle 0|\gamma_2|1\rangle|^2}{R_{eff,2}} \right] \quad (9.2)$$

where $|0\rangle$ and $|1\rangle$ are the two lowest energy eigenstates of the system. The coordinates of the normal modes of the system, η_Q and η_I , are defined implicitly from:

$$\gamma_1 = \alpha_{1,Q}\eta_Q + \alpha_{1,I}\eta_I \quad (9.3)$$

$$\gamma_2 = \alpha_{2,Q}\eta_Q + \alpha_{2,I}\eta_I, \quad (9.4)$$

where the α 's are determined according to Eqs. 4.35-39. η_Q and η_I can be written in terms of the raising and lowering operators of the normal mode states:

$$\eta_Q = \sqrt{\frac{\hbar}{2\omega_Q}}(a_Q^\dagger + a_Q) \quad (9.5)$$

$$\eta_I = \sqrt{\frac{\hbar}{2\omega_I}}(a_I^\dagger + a_I) \quad (9.6)$$

Note that the normal modes have unity effective masses. In other words, the α 's take care of the mass scaling [21,39]. Substituting Eqs. 9.3-9.6 into Eq. 9.2, one can see that the only matrix elements that contribute to Γ_{10} (the energy relaxation rate from the qubit excited state $|1\rangle$) are the a_Q elements, and I can write:

$$\Gamma_{10} = \frac{\hbar^2}{4e^2} \left[\frac{(\alpha_{1,Q})^2}{R_{eff,1}} + \frac{(\alpha_{2,Q})^2}{R_{eff,2}} \right] \quad (9.7)$$

To evaluate this expression, I needed to determine the contributions of each source of dissipation to the effective resistance across each junction.

Since some of the qubit's energy was stored in the isolation junction, I needed to model the amount of dissipation caused by the isolation junction's dielectric. The model discussed in Chapter 5 assumed that there was some dissipation caused by the isolation junction dielectric, but that the qubit junction was isolated from it by the SQUID loop inductance. In the new model, this dissipation was present in two places in Eq. 9.7. First, the isolation junction loss would contribute in parallel to $R_{eff,1}$ an amount:

$$R_{J2,1} = \left(\frac{1}{\omega_{01} C_{J2} \tan \delta_{J2}} \right) \left(\frac{L_1 + L_2 + L_{J2}}{L_2} \right)^2 \quad (9.8)$$

Second, the isolation junction also contributed a parallel to $R_{eff,2}$ and amount:

$$R_{J2,2} = \frac{1}{\omega_{01} C_{J2} \tan \delta_{J2}}, \quad (9.9)$$

Assuming that the dielectric in the isolation junction was the only source of dissipation, from Eq. 9.7, one would find a relaxation rate of:

$$\Gamma_{10,J2} = \frac{1}{T_{1,J2}} = \frac{\hbar^2}{4e^2} \left[\frac{(\alpha_{1,Q})^2}{R_{J2,1}} + \frac{(\alpha_{2,Q})^2}{R_{J2,2}} \right] \quad (9.10)$$

Of course, there were additional sources of loss, so Eq. 9.10 is incomplete.

I next consider the loss from the capacitors on the qubit side of the SQUID in the normal modes picture. In this case, though, the contribution to $R_{eff,2}$ from loss in C_x and C_{JI} was insignificant because the qubit states were largely stored in the qubit junction. Therefore, the affect of the loss that they caused in the isolation junction on the relaxation rate of the qubit state would have been minimal compared to the loss they caused in the qubit junction. The rate of spontaneous emission from dissipation caused by the dielectric in the IDC and C_{JI} are, to a good approximation:

$$\frac{1}{T_{1,Cx}} \approx \frac{\hbar^2}{R_{Cx}} \frac{(\alpha_{1,Q})^2}{4e^2} \quad (9.11)$$

$$\frac{1}{T_{1,J1}} \approx \frac{\hbar^2}{R_{J1}} \frac{(\alpha_{1,Q})^2}{4e^2} \quad (9.12)$$

where R_{JI} and R_{Cx} are given in Eq. 5.12. I used $\tan\delta \approx 1*10^{-3}$ for the Josephson junctions and $\tan\delta \approx 5*10^{-6}$ for the IDC.

The loss from the current bias leads can be dealt with in a slightly more sophisticated manner. First, consider the SQUID Hamiltonian (see Sect. 4.2):

$$H = \frac{p_1^2}{2m_1} + \frac{p_2^2}{2m_2} - E_{J1} \cos \gamma_1 - E_{J2} \cos \gamma_2 - \frac{\Phi_0}{2\pi} \frac{I_b}{L} (L_2 \gamma_1 + L_1 \gamma_2) + \left(\frac{\Phi_0}{2\pi} \right)^2 \frac{1}{2L} \left(\gamma_1 - \gamma_2 - \frac{2\pi}{\Phi_0} M I_\Phi \right)^2 \quad (9.13)$$

If the bias current I_b is fluctuating such that $I_b = I_{b0} + \delta I_b$, then Hamiltonian can be written:

$$H = H_0 + \frac{\Phi_0}{2\pi} \delta I_b \left(\frac{L_2}{L} \gamma_1 + \frac{L_1}{L} \gamma_2 \right) \quad (9.14)$$

where H_0 is the unperturbed Hamiltonian and the second term describes the effect that δI_b has on the system. I can describe the perturbation Hamiltonian:

$$\delta H_{lb} = \frac{\Phi_0}{2\pi} \delta I_b \left(\frac{L_2}{L} \gamma_1 + \frac{L_1}{L} \gamma_2 \right) \quad (9.15)$$

Using Eqs. 9.3 and 9.4, I can write this perturbation in terms of the normal modes of the SQUID:

$$\delta H_{lb} = \frac{\Phi_0}{2\pi} \delta I_b \left[\left(\alpha_{1,Q} \frac{L_2}{L} + \alpha_{2,Q} \frac{L_1}{L} \right) \eta_Q \right] \quad (9.16)$$

where I have ignored the η_I term because it will not contribute to $\langle 0 | \delta H_{lb} | 1 \rangle$.

At this point, I must make a short digression for the sake of accuracy. I note that here, the η 's refer to the full solutions for the normal modes of the SQUID, and strictly speaking, the α 's here come from solving the full eigenvalue problem. However, I lose little accuracy in considering the normal modes as found from the harmonic approximation to the Hamiltonian (see Sect. 4.3); therefore, I will continue using the same notation for the normal modes.

In the harmonic approximation, the α_Q 's can be found explicitly to be:

$$\alpha_{1,Q} = \left(\frac{2\pi}{\Phi_0} \right) \sqrt{\frac{1}{C_1 + C_2 \left(\frac{L}{L_1^*} - \omega_Q^2 L C_1 \right)^2}} \quad (9.17)$$

$$\alpha_{2,Q} = \left(\frac{L}{L_1^*} - \omega_Q^2 L C_1 \right) \alpha_{1,Q} \quad (9.18)$$

Now, since the η_I terms can be dropped, the relaxation due to the bias leads becomes:

$$\Gamma_{10,lb} = \frac{2\hbar\omega_{01}}{R_{lb}} \frac{\left| \langle 0 | \left(\alpha_{1,Q} \frac{L_2}{L} + \alpha_{2,Q} \frac{L_1}{L} \right) \eta_Q | 1 \rangle \right|^2}{4e^2} \quad (9.19)$$

where I have assumed that δI_b came from a quantum zero-point noise source (see Sect. 3.4):

$$R_{lb} = Z_0 \left(\frac{f_{01}}{f_c} \right)^4 \quad (9.20)$$

where $Z_0 \sim 50 \, \Omega$ is the impedance of the current bias line and $f_c \approx 180$ MHz is the cutoff frequency of the on-chip LC filter. Assuming that $|0\rangle$ and $|1\rangle$ are the harmonic approximations to the qubit normal mode ground and excited states, η_Q can be written in terms of raising and lowering operators according to Eq. 9.5. One then finds from Eq. 9.19 that:

$$\Gamma_{10,lb} = \frac{1}{T_{1,lb}} = \frac{\hbar^2}{4e^2 R_{lb}} \left[\alpha_{1,Q} \frac{L_2}{L} + \alpha_{2,Q} \frac{L_1}{L} \right]^2 \quad (9.21)$$

From an isolation standpoint, it is difficult to evaluate Eq. 9.21. It is clear that the current is divided between the two modes by the geometrical inductances. The effect of the Josephson inductances is wrapped up in $\alpha_{1,Q}$ and $\alpha_{2,Q}$ (see Eq. 9.17 and 9.18). We know that making $L \gg L_{J1}$ will make $\alpha_{1,Q} \gg \alpha_{2,Q}$, (see Sect. 4.3 and 4.4), and in that case one can get good isolation for the qubit by making $L \gg L_2$. In the case of device PB9, $\alpha_{1,Q} \sim 10\alpha_{2,Q}$ and $L \sim 10L_2$, so both terms are significant in calculating the dissipation from the bias lines in terms of isolation due to current division. However, the LC filter should have taken care of the noise at the qubit frequency by making R_{lb} very

large, so the qubit should still have been well isolated from the current bias line regardless of the degree of LJ isolation.

The contribution to the relaxation from the flux bias line I_Φ was found in a similar fashion. The part of the Hamiltonian that depends on the flux bias current is:

$$H_{I_\Phi} = \left(\frac{\Phi_0}{2\pi} \right)^2 \frac{1}{2L} \left(\gamma_1 - \gamma_2 - \frac{2\pi}{\Phi_0} M I_\Phi \right)^2 \quad (9.22)$$

If I assume that $I_\Phi \rightarrow I_\Phi + \delta I_\Phi$ expand the polynomial, and subtract off all of the constants, then the only terms that are left that depend on δI_Φ and interact with the qubit state are:

$$\delta H_{I_\Phi} = \left(\frac{\Phi_0}{2\pi} \right) \frac{M}{L} (\gamma_2 - \gamma_1) \delta I_\Phi \quad (9.23)$$

In terms of the normal modes the Hamiltonian is:

$$\delta H_{I_\Phi} = \left(\frac{\Phi_0}{2\pi} \right) \frac{M}{L} \delta I_\Phi (\alpha_{2,Q} - \alpha_{1,Q}) \eta_Q \quad (9.24)$$

where again I have dropped η_I terms. This interaction Hamiltonian has the same basic form as the current bias noise Hamiltonian. Assuming that the noise on the flux line is quantum zero-point noise the relaxation rate is:

$$\Gamma_{10,I_\Phi} = \frac{1}{T_{1,I_\Phi}} = \frac{\hbar^2}{4e^2 R_{I_\Phi}} \left(\frac{M}{L} \right)^2 (\alpha_{2,Q} - \alpha_{1,Q})^2 \quad (9.25)$$

where $R_{I_\Phi} \approx 50 \, \Omega$.

Note in Eq. 9.25 that the qubit was isolated from the dissipation by the square of the ratio of M to the total geometrical inductance L . Again the Josephson inductances L_{J1} and L_{J2} play a role in the size of the α 's, and in this way they affect the isolation from the flux bias line (note, for example, that T_{1,I_Φ} goes to infinity if $\alpha_{1,Q} = \alpha_{2,Q}$); however, it is not as simple as adding them into the SQUID loop inductance in the ratio of M to L .

This dependence is different than I found in the original estimate I found during the design in Chapter 5, which did not treat the normal modes.

The loss from the microwave line was more difficult to adapt to the normal modes picture. In Chapter 5, I assumed that the microwaves were coupled to the qubit junction through a small capacitor C_μ . It was possible that microwaves could also couple to the isolation junction and through the LC filter to ground, which would still couple to the qubit in the normal modes picture. To simplify things, I assumed that the microwave current was injected at the “upper right corner” of the SQUID (see Fig. 5.3) and was divided by the geometrical inductances in a similar fashion to the bias current. This gave an interaction Hamiltonian:

$$H_{I\mu} = \frac{\Phi_0}{2\pi} I_\mu \left(\frac{L_q + L_1/2}{L} \gamma_1 + \frac{L_2 + L_1/2}{L} \gamma_2 \right) \quad (9.26)$$

where L_q is the inductance of the part of the SQUID leading from the microwave line connection to the qubit junction. This is the same L_q that I used in Sect. 8.5 when accounting for the parasitic inductance of the IDC (see Fig. 8.11(a)). In this case $L = L_1 + L_2 + L_q$. By the same reasoning as was used with the bias current’s coupling to the qubit mode, the relaxation rate can be written as:

$$\Gamma_{10,I\mu} = \frac{1}{T_{1,I\mu}} = \frac{\hbar^2}{4e^2 R_{I\mu}} \left(\alpha_{1,Q} \frac{L_q + L_1/2}{L} + \alpha_{2,Q} \frac{L_2 + L_1/2}{L} \right)^2 \quad (9.27)$$

where $R_{I\mu}$ is:

$$R_{I\mu} \approx \frac{1}{Z_0 (\omega_{01} C_\mu)^2}, \quad (9.28)$$

$Z_0 \sim 50 \Omega$ and C_μ is the microwave coupling capacitance. This would provide an upper bound on T_l associated with this loss because it disregards the possibility of microwave coupling to other lines.

9.2 *Two-Level System Splittings in the Spectrum*

Coupling to discrete two-level systems (TLS's) was a major source of loss in early phase qubits [15-17]. To avoid this problem, device PB9 was designed to minimize the number of charged TLS defects that would couple to the qubit. The number of TLS's (per GHz in terms of their resonance frequencies) is determined by the volume of the dielectric in the junction and the loss tangent [17]. The qubit junction's Al/AlO_x/Al tunnel barrier is believed to have a high loss tangent, so I made its area very small. The IDC needed to be large, but it had a very low loss tangent, and this should have reduced the number of discrete TLS's. More importantly, the spacing between the fingers was ~ 1 mm so that a TLS charge moving 0.1 nm would produce only a very small energy change. Thus the splittings from individual TLS's in the IDC were expected to be too small to see.

Figure 9.1 shows a detailed view of the only identifiable splittings in the spectrum of device PB9. The full spectrum is shown in Fig. 8.7. The splitting in Fig. 9.1(a) was originally thought to be the only splitting. It was easily identified because it was strongly coupled to qubit and it had a very wide spectral peak. The size of the splitting was only about 10-15 MHz, and this is related to the strength of the coupling to the qubit. If a quantum system is strongly coupled to the qubit, it generates a large splitting in the spectrum [15]. The fact that the spectral peak was wide relative to the qubit peak (about

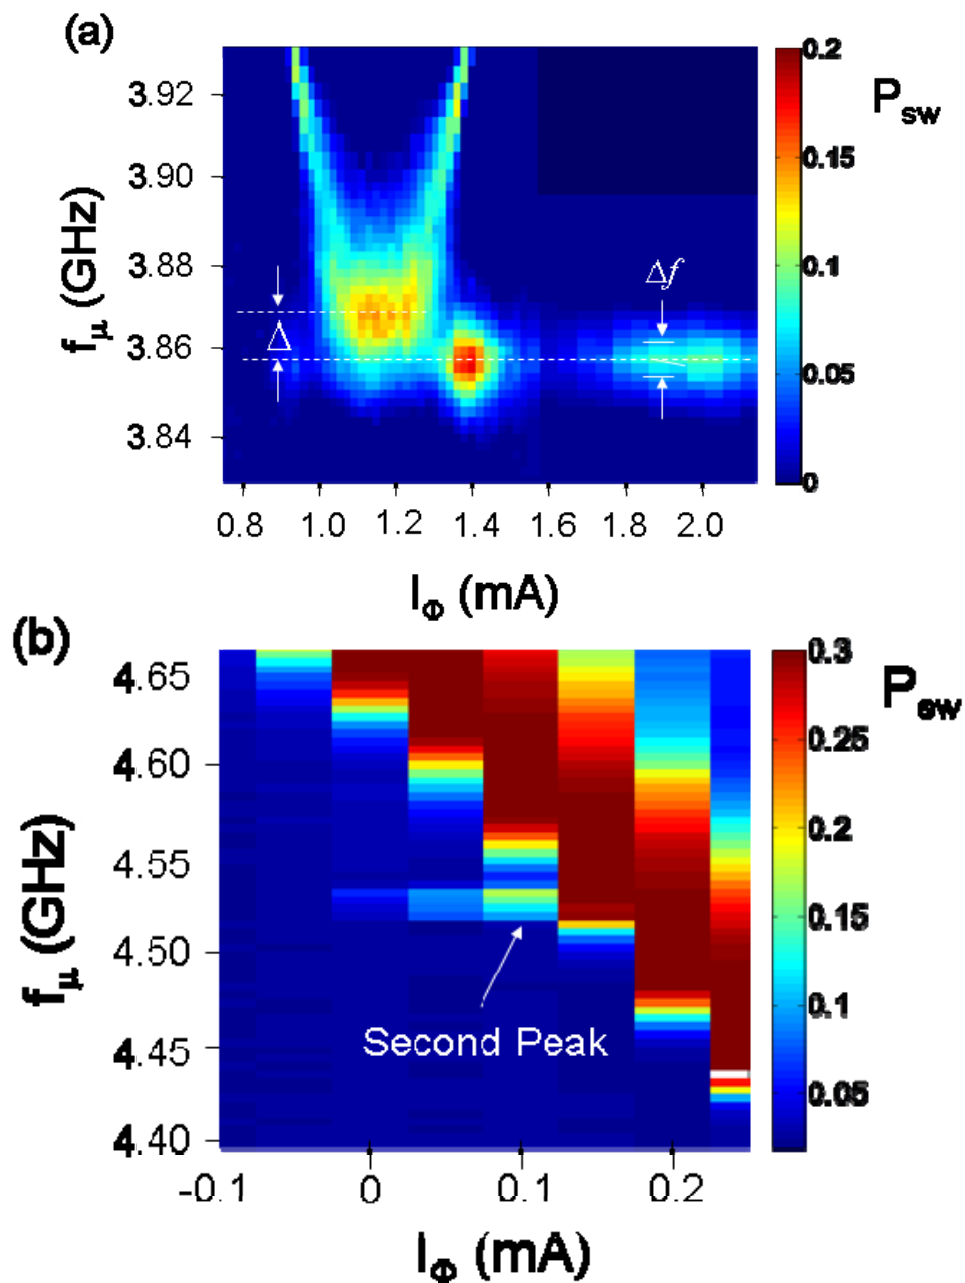


Fig. 9.1 Detailed view of the only two splittings in the spectrum. (a) The larger splitting at 3.86 GHz. $\Delta \approx 10\text{-}15$ MHz is the size of the splitting and $\Delta f \approx 6\text{--}8$ MHz is the full width at half max of the peak. (b) Splitting that was nearly invisible during low power spectroscopy. It originally showed up as a slight deformation of the peak. It is exaggerated in (b) using high power microwaves.

6-8 MHz) meant that the other quantum system was relatively lossy [15,52]. Both of these factors suggest that the T_1 of the qubit should have been shorter when f_{0I} was close to 3.86 GHz. Later I found a second splitting. Although it was nearly invisible in the transition frequency spectrum (see Fig. 8.7), it did create a discernable deformation of the resonance peak at 4.53 GHz. It had a very narrow peak (< 5 MHz) and the splitting was very small (< 5 MHz). Fig. 9.1(b) shows some data where the spectral peak of the other quantum system was exaggerated by a high power microwave drive making it visible away from the avoided crossing with the qubit. The fact that the peak is more visible and the splitting is larger suggests that it is caused by a microwave resonance in the circuit, not a charged TLS. However, TLS has become synonymous with “avoided level crossing” and “splitting” in the jargon of superconducting quantum computing so even though it may not be precisely correct, I will continue to use TLS as the label for other quantum systems that interact with the qubit causing splittings in the spectrum and decoherence.

The amount of dissipation caused by a TLS depends on the strength of the coupling to the qubit, the relaxation rate of the TLS, and the mechanism (*i.e.* charge, flux, critical current, *etc.*) by which the TLS coupled to the qubit [15,52]. I used what I knew to create an *ad hoc* model to represent this loss. The width of the TLS spectral peak Δf_{TLS} can be related to T_1 by:

$$\Delta f_{TLS} \approx \frac{1}{\pi T_{2,TLS}^*} \leq \frac{1}{2\pi T_{1,TLS}} \quad (9.29)$$

where $T_{2,TLS}^*$ is the spectroscopic coherence time, and I have used the fact that T_2^* is an upper bound on $2T_{1,TLS}$ (see Chapter 10 for further discussion of T_2^*). Also, the splitting

Δ was related to the coupling to the qubit [15]. The *ad hoc* model I used to represent the loss due to the TLS was:

$$\frac{1}{T_{1,TLS}} \approx \frac{\pi \Delta f_{TLS}}{1 + \left(\frac{\omega - \omega_{0,TLS}}{2\pi\Delta} \right)^2} \quad (9.30)$$

According to this model, the size of the splitting Δ determines how far from the resonant frequency the TLS will influence the overall T_1 of the qubit. The model assumes that the TLS has a lower T_1 than the qubit, and although I did not expect it to represent the loss completely accurately it allowed me to keep track of the location of the TLS splittings.

9.3 Expected T_1 from Normal Modes Dissipation Model Using Extracted Parameters

With the updated normal modes model for the dissipation of the qubit and the extracted parameters of the device, I obtained an estimate for what the T_1 of the device should have been. The model included the loss from two TLS's. For the other sources of dissipation, each $T_{1,i}$ depended on $\alpha_{1,Q}$ and $\alpha_{2,Q}$, which in turn depended on the applied flux. Therefore, it was not possible to quote a single value for T_1 . Figure 9.2 shows the expected T_1 vs. Φ_a from each source of dissipation, except for the TLS's, and Table 9.1 summarizes the predictions. The total decay rate was then found from:

$$\frac{1}{T_1} \approx \frac{1}{T_{1,lb}} + \frac{1}{T_{1,I\Phi}} + \frac{1}{T_{1,I\mu}} + \frac{1}{T_{1,Cx}} + \frac{1}{T_{1,J1}} + \frac{1}{T_{1,J2}} + \frac{1}{T_{1,TLS1}} + \frac{1}{T_{1,TLS2}} \quad (9.31)$$

The lowest $T_{1,i}$ had the most influence on the T_1 of the qubit. $T_{1,I\Phi}$ was definitely the lowest at around 2 μ s, and it only showed slight variations as the flux bias current was changed. The next lowest were the T_1 's associated with the IDC and the qubit junction's dielectric. Everything else was relatively insignificant.




Fig. 9.2 Theoretical energy relaxation times from each source of dissipation. Note that the dissipation from the flux bias line I_ϕ was the greatest, giving the shortest T_1 of about 2 μs .

Table 9.1 Results of normal modes dissipation analysis using extracted parameters (see Table 8.1).

Dissipation Source	Symbol	R_i (M Ω)	T_1 (μ s)
Bias Leads	I_b	13-33	20-50
Flux Bias	I_ϕ	1.3	2
μ -wave line	I_μ	67-110	100-160
Qubit JJ capacitance	C_{J1}	6.0-6.7	9-10
Interdig. capacitance	C_x	4.7-5.3	7-8
Isolation JJ capacitance	C_{J2}	20-27	30-40
Expected T_1 (μ s)	$\sim 1 \mu$ s		

From the original design model, the loss from the flux line was supposed to be very low. However, M came out much larger (4.95 pH in the final device *vs.* 1 pH in the design) and L_I came out smaller (0.7 nH in the device *vs.* 1.25 nH in the design) than I intended and this increased the contribution of the flux bias line to the dissipation in the qubit. Another reason for why the times were shorter was that the original model assumed that the Josephson inductances were part of the isolation factor, but upon closer examination, the normal modes model revealed a more complex dependence on $\alpha_{I,Q}$ and $\alpha_{2,Q}$

Figure 9.3 shows the expected T_1 versus flux from the normal modes model for device PB9 including the two TLS's. T_1 varies with flux but is always about $1.20 \pm 0.02 \mu$ s as long as the qubit is biased away from the TLS splittings. Note that the expected T_1




Fig. 9.3 Theoretical plot of T_1 vs. I_ϕ using the normal modes dissipation model (solid curve). The resonance frequency is shown as a dashed curve. The major dips in T_1 occur because of the presence of TLS splittings in the spectrum at $f_{01} = 3.86$ and 4.53 GHz. The updated dissipation model predicted a T_1 of about $1.2 \mu\text{s}$ with only a slight variation with applied flux, except for the effects of the TLS's.

is only slightly higher at lower frequencies. Figure 9.2 shows that this was the general trend for all sources of dissipation, except for loss from the bias lines. If the splitting had not been present at the bottom of the spectrum, then the theory would have given the longest T_I there. Finally, I note that the theoretical value of T_I was lower than originally designed, mainly due to the fact that the ratio of L to M came out smaller than expected, but it should still have been $>1 \mu\text{s}$, but it was still longer than any reported results on phase qubits.

9.4 Determining T_I from S-Curve Measurements

I measured the relaxation time T_I of device PB9 according to the procedure detailed in Chapter 7. Unfortunately, the device showed the same anomalous switching behavior that was observed in device DS8 [64,70], and this meant that some effort was needed to extract the probability of being in the excited state. Figure 9.4 shows the S-curves of the device when the qubit was in the $|1\rangle$ state (red), the $|0\rangle$ state (blue), and a state that was as near as I could get to $(|0\rangle + |1\rangle)/\sqrt{2}$ (green). The dashed curve shows what the superposition state should have yielded if the measurement actually gave a conventional [64,70]. Clearly the data was anomalous. To get a measure of the probability P_I of the qubit being in the $|1\rangle$ state, I had to measure the whole S-curve for the device at each time step. The extraction of P_I from the S-curves was complicated, and not necessarily entirely accurate, but it was a better measure of P_I than a measurement of the probability of switching P_{sw} from measurement pulses with a single height V_p . The longest relaxation time I found was at $I_\Phi = 0.084 \text{ mA}$ and $I_b = 6.8 \mu\text{A}$ where $f_{0I} = 4.05 \text{ GHz}$. Figure 9.5 shows the corresponding S-curves vs. time. In this




Fig. 9.4 Comparison of anomalous switching behavior of device PB9 to conventional switching behavior. The blue and red curves show the switching curves of the device when the qubit was in the $|0\rangle$ and the $|1\rangle$ state, respectively. The dashed curve shows the expected S-curve of the 50/50 superposition state for conventional switching. The green curve is the measured S-curve of the 50/50 superposition state. It is clearly anomalous.




Fig. 9.5 a) False color plot of the S-curves of the device during relaxation from $|1\rangle$ to $|0\rangle$. Note that the decay of P_{sw} along a single horizontal trace (constant V_p) is different for each value of V_p . (b) Close-up of the S-curves as the system is pumped from $|0\rangle$ to $|1\rangle$. The dashed lines highlight S-curves that would be used in fitting S-curves found during the decay.

false color plot each vertical slice represents an s-curve measured at time t . For this data, the qubit was pumped into the first excited state by a 24 ns resonant microwave pulse (*i.e.* a π -pulse) and allowed to relax back to the ground state.

The probability of being in the $|1\rangle$ state P_I was determined from the S-curves according to an *ad hoc* model that assumed that the excitation of the qubit state altered the currents flowing in the isolation junction. As I discussed previously, it was impossible for the qubit state to tunnel out of the potential well in the direction of the qubit junction's phase γ_1 . The lowest potential barrier was along the γ_2 direction, and the height of the barrier changed with the current I_2 through the isolation junction. If the tunnel barrier became lower, then the escape rate Γ would be higher. The probability P_{sw} of the device switching to the voltage state during a measurement pulse of duration τ is related to the escape rate by [64,70]:

$$P_{sw} = 1 - e^{-\Gamma\tau} \quad (9.32)$$

For small current changes δI_2 , Γ is related to δI_2 by [64,70]:

$$\Gamma = \Gamma_0 e^{a\delta I_2} \quad (9.33)$$

The *ad hoc* model behind this analysis has been detailed by Hyeokshin Kwon, who first observed the behavior [64,70]. The basis for the model is that the current through the isolation junction is an operator \hat{I}_2 . When the qubit is in the state $|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle$, the expectation value of \hat{I}_2 is:

$$\delta I_2 = \langle \Psi | \hat{I}_2 | \Psi \rangle = \alpha^2 I_{00} + \beta^2 I_{11} + 2|\alpha||\beta| I_{01} \quad (9.34)$$

where $I_{00} = \langle 0|I_2|0\rangle$ represents the current flowing through the isolation junction when the qubit is in the $|0\rangle$ state, $I_{11} = \langle 1|I_2|1\rangle$, and $\langle 0|I_2|1\rangle = \langle 1|I_2|0\rangle$ is an interference term between the two states. Note that Eq. 9.30 gives an expectation value for I_2 that depends on three constant functions I_{00} , I_{11} , and I_{01} .

To simplify things, one can absorb Γ_0 and τ into the definition of I_{00} , I_{01} , and I_{11} to create effectively dimensionless current functions \tilde{I}_{00} , \tilde{I}_{11} , and \tilde{I}_{01} [70]. The idea is to then use the data to obtain estimates for \tilde{I}_{00} , \tilde{I}_{11} , and \tilde{I}_{01} . From Eqs. 9.32 and 9.34 one then finds.

$$\tilde{I}_{00} = \ln(\Gamma_{|0\rangle}) = \ln\left[-\ln(1 - P_{sw,|0\rangle})\right] \quad (9.35)$$

where $\Gamma_{|0\rangle}$ is the escape rate from the ground state and $P_{sw,|0\rangle}$ is the measured S-curve when the qubit was in the ground state.

Similarly:

$$\tilde{I}_{11}(V_p) \approx \ln(\Gamma_{|1\rangle}(V_p)) \approx \ln\left[-\ln(1 - P_{sw,|1\rangle}(V_p))\right] \quad (9.36)$$

Determining \tilde{I}_{01} was less straightforward. From the S-curve $P_{sw,\pi/2}$ of the $(|0\rangle + |1\rangle)/\sqrt{2}$ state, I could find $\tilde{I}_{\pi/2} = \ln\left[-\ln(1 - P_{sw,\pi/2})\right]$. Then from Eq. 9.34 I could write:

$$\tilde{I}_{\pi/2} = \frac{1}{2}\tilde{I}_{00} + \frac{1}{2}\tilde{I}_{11} + \tilde{I}_{01} \quad (9.37)$$

and thus:

$$\tilde{I}_{01} = \tilde{I}_{\pi/2} - \frac{\tilde{I}_{00} + \tilde{I}_{11}}{2} \quad (9.38)$$

Another difficulty was actually putting the system in the $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ state or the $|1\rangle$ state. It was simple to determine when the qubit was in the $|0\rangle$. I assumed that the qubit was in the $|1\rangle$ state right after the 24 ns π -pulse since this gave the largest escape probability. The S-curves measured at $t = 2500$ ns and $t = 24$ ns were used to determine \tilde{I}_{00} and \tilde{I}_{11} , respectively, from Eqs. 9.35 and 9.36. The time for \tilde{I}_{01} was determined by trial and error. I started by trying the S-curve at $t = 12$ ns (half the time for the π -pulse) and tried using several S-curves to get I_{01} . The one at that gave the best fits overall was the one I used. For example, I used the S-curve at $t = 8$ ns to get P_I for Fig. 9.6.

I determined P_I at each time t by performing a χ^2 fit of the ad hoc model to the S-curve data. Given a measured S-curve P_{sw} , I obtained the effective current from:

$$\tilde{I}_2^{data} = \ln(-\ln(1 - P_{sw})) \quad (9.39)$$

I then fit this to the theory:

$$\tilde{I}_2^{th} = (\beta^2 - 1)\tilde{I}_{00}(V_p) + \beta^2\tilde{I}_{11}(V_p) + 2\beta\sqrt{\beta^2 - 1}\tilde{I}_{01}(V_p) \quad (9.40)$$

The statistical error in I_2^{th} was [70]:

$$\sigma_{I_2,t}(V_p) = -\frac{\sqrt{\frac{P_{sw}(V_p)}{N(1 - P_{sw}(V_p))}}}{\ln(P_{sw}(V_p))} \quad (9.41)$$

where N is the number of samples of P_{sw} taken at each V_p . I usually took $N = 6000$ samples. The equation for χ^2 that was minimized was:

$$\chi^2(\beta) = \sum_{V_p} \frac{[\tilde{I}_2^{data}(V_p) - \tilde{I}_2^{th}(V_p)]^2}{\sigma_{I_2,t}^2(V_p) - (\beta^2 - 1)\sigma_{I_{00}}^2(V_p) + \beta^2\sigma_{I_{11}}^2(V_p) + 2\beta\sqrt{\beta^2 - 1}\sigma_{I_{01}}^2(V_p)} \quad (9.42)$$

where the error in \tilde{I}_{00} and \tilde{I}_{11} were determined from Eq. 9.41 and σ_{I01} was:

$$\sigma_{I01} = \sqrt{\sigma_{\pi/2}^2 - \frac{1}{2}(\sigma_{I00}^2 + \sigma_{I11}^2)} \quad (9.43)$$

It was necessary to include these errors because \tilde{I}_{00} , \tilde{I}_{11} , and \tilde{I}_{01} were determined from data. The value of β that minimized $\chi^2(\beta)$ was related to P_I by $P_I = \beta^2$ (see Appendix A).

Figure 9.6(a) shows P_I vs. time for the best relaxation data I obtained from device PB9. Figure 9.6(b) shows the χ^2 of the fit that produced each value of P_I . Note that χ^2 was usually about 5-10, and it never got above 35. In order to determine the quality of the fit, I needed to know the number of degrees of freedom for the fit. Note that according to Eq. 9.37 the error in each point was large if P_{sw} , such as at the beginning of each S-curve; therefore, those points did not really count toward the determining the fit. The same is true for the points where $P_{sw} \approx 1$. Hence, although each S-curve was made up of 70 measurements of P_{sw} , the only points that meant much to the fitting process were the 30-40 points during the rise of the S-curve from $P_{sw} = 0$ to $P_{sw} = 1$. This meant that the reduced χ^2 's of the fits were near to one, and thus the fits were close to the data. However, some of the reduced χ^2 were significantly less than 1. This meant that the statistical estimates for the errors were probably too high.

Now that I had estimates for P_I versus time, I could plot the relaxation data in a standard form. Assuming that P_I was decaying at an exponential rate, T_I was the inverse of the slope of the graph. However, the baseline of the decay needed to be taken into account when estimating T_I . Accordingly, I fit the data to:

$$P_I(t) = (P_{1,0} - P_{1,\infty})e^{-t/T_I} + P_{1,\infty} \quad (9.44)$$




Fig. 9.6 (a) Example of P_I extracted from an energy relaxation measurement on device PB9 that showed the best T_I at $f_{0I} = 4.050$ GHz, $I_\Phi = 0.084$ mA, and $I_b = 6.8$ μA . b) χ^2 of the S-curve fits used to determine P_I . Each fit had about 30-40 degrees of freedom. χ^2 was minimized using Matlab (see Appendix A).

where $P_{I,0}$ is the value of P_I at the start of the decay ($t = 0$) and $P_{I,\infty}$ is the baseline value to which P_I was decaying. Thus:

$$\ln(P_I(t) - P_{I,\infty}) = \ln(P_{I,0} - P_{I,\infty}) - \frac{t}{T_I} \quad (9.45)$$

and I just need to find the slope on a semi-log plot once the offset was taken into account.

Figure 9.7 shows the energy relaxation decay data with two fits. The green is for $T_I = 420$ ns, and the red is for $T_I = 280$ ns. The 420 ns curve fits better to the values of P_I that had a lower value of χ^2 associated with them; however, from the quality of this fit, it was more realistic to say that the best T_I was about $350 \text{ ns} \pm 70 \text{ ns}$.

When I began my research on quantum computing, the qubits that we were measuring in our lab were showing $T_I \approx 15\text{-}25$ ns. Five years later, device PB9 improved on that number by a factor of 10. I note that T_I of 350 ns is not the longest relaxation time observed in a phase qubit (which is about 500 ns [19,73]), but it is getting there.

9.5 Determining the Limiting Source of Dissipation Through a Study of the Dependence of T_I on Applied Flux

The longest value I obtained for T_I was about one-third of what was expected from the updated dissipation model. Of course, a single measurement of the best T_I was not a good indicator of the validity of the dissipation model. Nor did it tell me anything about why the relaxation time was this short.

In order to identify the limiting source of dissipation, I measured the T_I of the device as a function of flux at 50 MHz intervals in f_{0I} . I wanted to see if T_I would show a dependence on f_{0I} that would point to one of the sources of dissipation. The loss from the flux line (Eq. 9.25) was virtually independent of the qubit frequency. The loss from the




Fig. 9.7 Exponential decay fits (dashed curves) to the extracted P_1 (filled circles) during an energy relaxation measurement. The decay constant of the two fits represent the upper and lower bounds on the best T_1 of device PB9. Note that the green fit agrees with the portion of the curve where the S-curves showed the lowest χ^2 's in Fig. 9.6(b).

dielectrics (Eq. 9.10) went like $1/f_{0I}$. The loss from the microwave line went like $1/f_{0I}^2$ (Eq. 9.28). Finally, the loss from the current bias line (Eqs. 9.20) went like f_{0I}^4 , so if it was the dominant source of dissipation, then T_I would have gotten longer at higher frequencies.

Figure 9.8 shows my measured data for T_I vs. flux. I note that the plot shows a strong and regular dependence of T_I on flux. The error bars for each point were determined in a similar fashion as for the T_I measurement shown in the previous section. They represent a range of possible T_I 's that would fit the decay of the P_I extracted from the S-curves. For comparison, I also show f_{0I} as a dashed curve in Fig. 9.8., covering nearly one full period of the spectrum.

Figure 9.8 has two identifiable regions. First, in the middle, where the frequency was greater than 4.6 GHz (or $-0.2 < I_\Phi < 0$ mA), T_I was around 100 ns. Second, between 4 GHz and 4.5 GHz (or for $0.05 < I_\Phi < 0.1$ mA) T_I increased to around 200-300 ns. Note the dip in T_I to about 70 ns near 4.6 GHz where there was an avoided level crossing dividing the two regions. T_I was also very low near 3.86 GHz because of the avoided crossing near $\Phi_0/2$. While the TLS loss was not surprising, the other two well-defined regions of loss were quite unexpected and unlike what I expected from any of the known sources of dissipation.

With an unusual dependence on flux and f_{0I} , the T_I of the device clearly would not fit the dissipation model detailed in the beginning of this chapter. However, eventually I recognized that there was a correlation between the T_I of the system and the amount of power that it took to perform the π -pulse for the T_I measurement. The power required to produce a π -pulse is a measure of the coupling between the qubit and the




Fig. 9.8 Points are measurements of T_1 versus flux bias current over one period of the spectrum. Note that T_1 is worse in the region toward the middle of the graph. The dashed curve shows f_{01} . The dividing line between the two regions in frequency was about 4.6 GHz.

microwave line. To find the coupling, the system was prepared in the $|1\rangle$ state with a 24 ns π -pulse for each measurement. At each flux, the time was kept constant for the pulse, and the output power of the microwave source was adjusted to ensure that P_{sw} at $t = 24$ ns was a maximum. Too much power would rotate the state past $|1\rangle$ on the Bloch sphere, and too little would not rotate it all the way to $|1\rangle$. Just the right amount of power would put the state at $|1\rangle$ and P_{sw} would be the highest.

Figure 9.9 shows the π -pulse power at fixed $t = 24$ ns plotted vs. applied flux bias current. Note that there are two regions to this graph as well. The y-axis is power in dBm, so a more negative value means less power. In the middle it took less power to rotate the state to $|1\rangle$, then it did in the outside region. This meant that the qubit was more strongly coupled to the microwave drive line in the middle region (between -0.2 mA and 0 mA of flux current) than in the outside region ($I_\Phi < -0.2$ mA or $I_\Phi > 0$ mA).

The π -pulse power is related to the coupling between the qubit and the microwave source. According to Fermi's Golden Rule, the rate of transitions W_{01} between $|0\rangle$ and $|1\rangle$ caused by microwave current is governed by:

$$W_{01} = \frac{2\pi}{\hbar} \left| \langle 0 | \hat{H}_{I\mu} | 1 \rangle \right|^2 \quad (9.46)$$

where $\hat{H}_{I\mu}$ is the perturbation Hamiltonian of the microwave drive current. Although it was difficult to know the exact form of the perturbation Hamiltonian, I assumed that it was of the form:

$$\hat{H}_{I\mu} = \kappa I_\mu(t) \hat{\eta}_Q \quad (9.47)$$




Fig. 9.9 Data points show π -pulse power used during the T_I measurements plotted versus flux bias current. The π -pulse power showed variations that were similar to those shown by T_I (see Fig. 9.8). In particular, when it took more power to pump the qubit to $|1\rangle$ in a given time, the qubit showed a longer T_I .

where $I_\mu(t)$ is the microwave drive current, κ is a frequency dependent coupling constant that represents the amount of microwave current that actually interacted with the qubit mode, and $\hat{\eta}_Q$ is the qubit normal mode coordinates (see Eq. 9.5). $\hat{\eta}_Q$ connects $|0\rangle$ to $|1\rangle$ by the raising operator and the matrix element is:

$$\langle 0 | \hat{H}_\mu | 1 \rangle \propto \kappa I_\mu \quad (9.48)$$

where I_μ is the amplitude of the microwave signal and I have absorbed the frequency dependence and all other constants into the coupling constant (note that κ is not dimensionless). The bare Rabi frequency Ω_R of the qubit under the influence of the drive was related to I_μ by:

$$\Omega_R \propto \langle 0 | \hat{H}_\mu | 1 \rangle \propto \kappa I_\mu \quad (9.49)$$

For my measurements, Ω_R was fixed at $1/(2\pi*48 \text{ ns})$. Therefore, the coupling constant of the microwaves to the qubit mode was:

$$\kappa(f_{01}) = \frac{1}{I_\mu * 48 \text{ ns}} \quad (9.50)$$

where I_μ was the current required to produce $\Omega_R = 1/(2\pi*24 \text{ ns})$. Note that when I_μ is larger, the coupling κ is smaller. I determined I_μ from the output power from the microwave source using:

$$I_\mu = \left(\frac{2}{Z_0 * 1000} 10^{\frac{P_\mu}{10}} \right)^{1/2} \quad (9.51)$$

where P_μ is the output power required to produce the π -pulse in dBm, and I assumed that the impedance of the microwave line was $Z_0 = 50 \Omega$. Thus I used the measured values for

the microwave power to obtain the coupling constant associated with the interaction of the microwave drive with the qubit normal mode.

I note that the rate of spontaneous emission that gives T_I due to the line depends not on the noise current amplitude, but on the noise power [26,39]. Therefore, in order to compare the values of T_I with the coupling to the microwave line, I plotted $1/\kappa^2$ (*i.e.* the decoupling squared), which would vary in the same way as T_I . Figure 9.10 shows the result. The correlation is readily apparent and quite striking. An interesting aspect of the graph is that T_I varied roughly by a factor of two between the two regions, and the $1/\kappa^2$ also varied roughly by a factor of two. This correlation is strong evidence that dissipation from the microwave drive line was limiting T_I .

There is an apparent anomaly in the graph that actually points to an explanation for the abrupt change in the decoupling in the two regions. Note how the decoupling spikes in the vicinity of the avoided crossing at $I_\Phi = 0.04$ mA and -0.24 mA. There are two possible explanations for this. One possible explanation is that there was a TLS that was being excited at that frequency, taking some of the microwave power that should have been exciting the qubit. This would make it seem like the qubit was decoupled from the line because it would have taken more power to pump it all the way to $|1\rangle$.

Another explanation for the spike in the decoupling is that the splitting in the spectrum at $f_{0I} = 4.53$ GHz was due to a resonance in the microwave line. This line was simply a bare microstrip line that terminated in a small interdigitated capacitor, and the connection to the chip was made by a wirebond from the bare center conductor of the SMA connector that went through the side of the sample box. The impedances at these connections were not matched, so there was the possibility of standing waves [74] or a




Fig. 9.10 Plot of T_1 of (solid circles) and the decoupling of the qubit from the microwave line (dashed curve) versus flux bias current. When the qubit was more decoupled from the microwave line, it showed a longer T_1 , suggesting that the lifetime was being limited by dissipation from the microwave drive line. The spikes in the decoupling at $I_\Phi = -0.24$ and 0.04 mA were due to coupling to the spurious resonance at 4.53 GHz.

$\lambda/4$ or $\lambda/2$ resonance at a frequency of 4.53 GHz. The impedance of the resonator could have blocked the microwave signal on resonance, which would have increased the decoupling [74]. In fact, the splitting in the spectrum could be due to the qubit coupling to a microwave circuit resonance in the microwave line. Also, the fact that this splitting only showed up at high power suggests it is a harmonic oscillator mode, because a TLS would have saturated at high power and become less visible.

Whatever was the source of the splitting at 4.53 GHz, the fact that T_1 varied with the strength of the coupling of the qubit to the microwave drive line was evidence that the limiting source of dissipation in the qubit was over coupling to the microwave drive line.

9.6 *Future Device Improvements*

My theoretical reanalysis of loss in the qubit showed that the mutual inductance M and loop inductance L were such that the dissipation from the flux bias line should have been the main source of dissipation in device PB9. However, analysis of the T_1 data showed that the limiting source of relaxation was actually from coupling to the microwave line. Both of these problems can be addressed in a future design.

The microwave line was intended to be only weakly coupled to the qubit; however, the behavior was not rigorously characterized. Figure 9.11 shows a picture of the microwave line and the SQUID. In the upper right corner there was a small break in the microwave line that constituted a 2-finger interdigitated capacitor with $C_\mu \approx 1$ fF according to Eqs. 5.18 and 5.19. If this capacitor had been the only means of coupling to the qubit, then the qubit should have been well isolated. However, it was possible that there was some mutual inductive coupling between the microwave line and the SQUID.


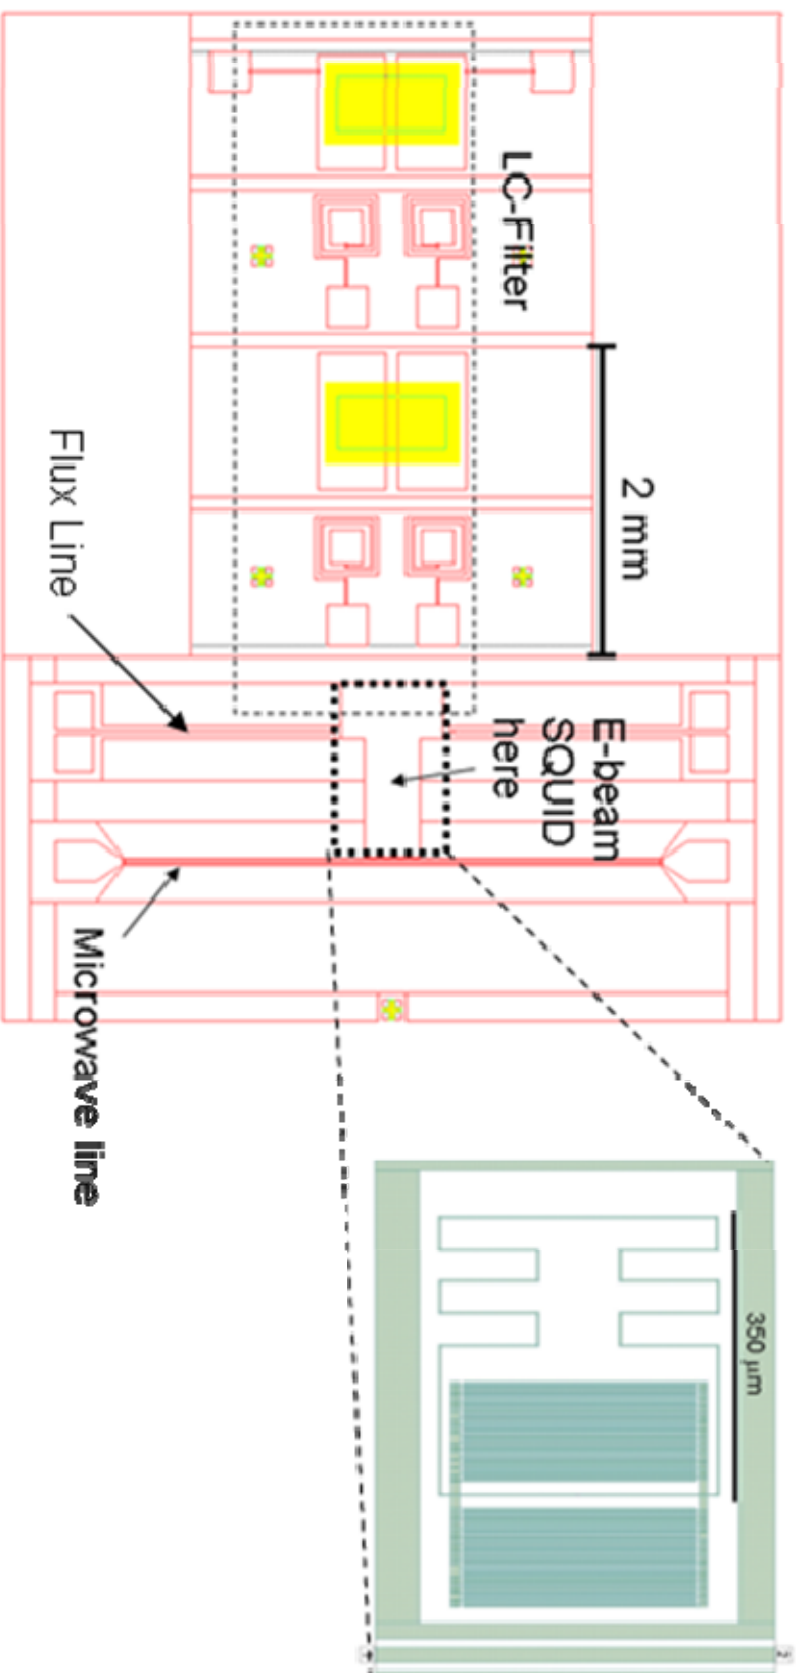


Fig. 9.11 Photograph of the microwave drive line's connection in device PB9. The line was meant to be coupled to the qubit through the small capacitor C_{μ} formed by a break in the lines in the upper right of the picture. Unfortunately, the coupling to the qubit was stronger than expected.

When I designed the device, I used FastHenry to model and predict the mutual inductances in the device, but Fast Henry was incapable of determining a high-frequency mutual inductance. Although, I assumed that the microwave line was a poorly terminated transmission line that would transmit some small amount of power to the qubit through a very small IDC, the data from the qubit showed that it was much more strongly coupled than I expected.

Rangga Budoyo has redesigned the microwave line for a new device. Figure 9.12 shows a picture of the design. Instead of a poorly terminated wire, he will use a coplanar waveguide with a $50\ \Omega$ impedance that will have input and output lines. More to the point, he will examine the coupling to the qubit using Microwave Office. Microwave Office can give E-M field density simulations of a device layout in a 2½-D model. It also can check the response of microwave signals on the coplanar waveguide, and find the coupling of the waveguide to the qubit as a function of frequency. The idea is to design the microwave line with a known small coupling to the qubit.

The new qubit should also be more isolated from the flux bias line. To do this, Rangga has added meanders to the SQUID loop to increase the inductance of the device. The meanders add some extra length to the line without increasing the pickup area of the SQUID. The flux line has also been moved farther away to decrease M . FastHenry gave simulated values of $M < 1\ \text{pH}$ and $L \approx 1.7\ \text{nH}$, which should greatly improve the isolation from the flux line. Thus the design of the next device addresses the two main shortcomings I found in device PB9: over coupling to the microwave line and the flux bias line.



R. Budooyo

Fig. 9.12 Design of a new SQUID phase qubit. The flux bias line has been moved farther from the SQUID and shielded by a ground plane to reduce the mutual inductance. Two stages of LC-filtering are used, and the microwave line will be a 50 Ω coplanar waveguide.

9.7 Conclusions

Device PB9 was designed with the intent of achieving an energy relaxation time T_1 greater than 1 μs . By shrinking the qubit junction and shunting it with a low loss interdigitated capacitor, I reduced the dissipation from the junction dielectric sufficiently to achieve this goal. The filtering on the current bias line was aggressive enough to cut off dissipation from the bias line, which was found to be the limiting source of dissipation in some of our previous qubits [52,53]. One troubling source of dissipation was loss from the flux bias line. Since L came out smaller than expected on device PB9 and M came out larger, the updated normal mode theoretical model identified dissipation from the flux bias line to be the dominant source of dissipation. Including all of the sources of dissipation combined (*i.e.* flux bias line, current bias line, microwave drive line, qubit junction dielectric, isolation junction dielectric, and IDC) the model predicted a T_1 of about 1.2 μs that was only weakly dependent on f_{01} and applied flux. This was lower than hoped for in the original design, but it would have been a large improvement over our previous qubit [70].

The data showed that the T_1 was even lower than the refined theoretical prediction. The best T_1 measured on the device was $350 \text{ ns} \pm 70 \text{ ns}$. The error in T_1 came from the uncertainty of the method by which the probability of the qubit being in the $|1\rangle$ state P_1 was determined. Device PB9 showed anomalous switching behavior that made it necessary to measure the entire switching curve of the device in order to extract P_1 . P_1 was extracted via an *ad hoc* model that assumed the current in the isolation junction was dependent on the state of the qubit. The fits to the data were not entirely satisfactory, and the decay of P_1 was not purely exponential, but they were good enough to identify a

range of possible T_I 's for the device. The best T_I was about a factor of 3 or 4 lower than expected from the model.

I performed a study of T_I vs. f_{0I} and applied flux in order to ascertain whether the T_I of the device depended on either. Since each T_I measurement was done by taking s-curve measurements at each time during the decay, the study took nearly two months. The data showed a striking variation with f_{0I} or applied flux. I identified a correlation between the amount of microwave power used for the π -pulse in each T_I measurement and the measured T_I of the qubit. That microwave power level was related to the coupling of the microwave line to the qubit mode of the device. When the microwave drive was strongly coupled to the qubit, the T_I was shorter. This was significant evidence that the T_I of the qubit was limited by coupling to the microwave drive line.

Since over-coupling to the microwave line and the flux line were both identified as significant contributors to the dissipation in the qubit, a new device is being designed to decrease these couplings. The flux bias line will be farther away from the SQUID to reduce the mutual inductance, and the SQUID loop inductance is being increased by adding some meanders to the loop. The microwave line is being redesigned as a coplanar waveguide instead of a bare wire terminated by a small capacitor. The coupling is being modeled in Microwave Office[®] to better ensure that it will be small. Thus the next device is being designed with the intent of reducing dissipation caused by over-coupling to the microwave and flux bias lines.

Chapter 10

Dephasing Analysis of Device PB9

In this chapter, I describe my analysis of dephasing in device PB9. This analysis is based on theory and data. In Chapter 3, I discussed how fluctuations of the energy levels of the qubit lead to a loss of phase coherence of quantum superposition states. In this chapter, I look at how thermal noise on the bias leads and $1/f$ flux and critical current noise should have affected device PB9. Here I use the normal modes model obtained in Chapter 8 to find more accurate values of the noise transfer functions and the dephasing time limits from each source. The analysis revealed unexpected sweet spots in the spectrum. In the later sections of this chapter, I discuss the possibility of using these sweet spots to increase the dephasing times in dc SQUID phase qubits. $1/f$ flux noise is virtually unavoidable in a SQUID, and this noise is an exceptionally strong source of dephasing. Also, charge fluctuations in the barrier of a Josephson junction can lead to variations in its critical current, which cause dephasing by changing the inductance of the junction. Device PB9 was vulnerable to dephasing from all of these sources.

10.1 Noise Transfer Functions, Sweet Spots, and T_ϕ Predictions

In Chapter 5, I estimated the dephasing time T_ϕ of device PB9 based on the design parameters. Four sources of dissipation were considered during the design: Nyquist noise and dissipation on the current and flux bias lines, $1/f$ flux noise, and $1/f$ critical current noise. For dephasing from the current bias line T_ϕ I found [43]:

$$\frac{1}{T_{\phi, Ib}} \approx \frac{k_B T}{Z_0} \left(2\pi \frac{\partial f_{01}}{\partial I_b} \right)^2 \quad (10.1)$$

where, T is the temperature (20 mK), and Z_0 is the characteristic impedance of the bias line ($\sim 50 \Omega$). There were no isolation factors on the line impedance in this analysis because they were included in df_{01}/dI_b , which could be extracted directly from the dependence of the qubit resonance on the current bias.

Figure 10.1 shows the calculated noise transfer function df_{01}/dI_b vs. applied flux and the corresponding T_ϕ . I determined df_{01}/dI_b from a simulation using the PB9 design parameter (Table 8.1). At each applied flux, I found the minimum of the potential in order to get γ_1 and γ_2 , and thus L_{J1} and L_{J2} (see Chapter 8). I then found the qubit frequency f_{01} from the normal modes equations (see Chapter 4). Next I varied I_b to see how it affected f_{01} . I then repeated this procedure for a range of applied fluxes Φ_a . In this analysis, I incorporated the parasitic inductance of the IDC because it yielded the best fit to the spectrum (see section 8.5).

Figure 10.1 shows that df_{01}/dI_b has several significant and interesting features. First, note that the magnitude of df_{01}/dI_b in Fig. 10.1(a) was about 0.2 MHz/nA, which is about two orders of magnitude lower than the value assumed in the design (see Sect. 5.2.1). This illustrates how the isolation factor provided by the LJ-isolation was incorporated into the noise transfer function, and why ultimately it is best to obtain df_{01}/dI_b from the data instead of making general assumptions about the behavior of the qubit. The dashed lines in Fig. 10.1(a) point out the sweet spots with respect to fluctuations in the bias current. Note that the fact that the curve shows high but not infinite values for T_ϕ in Fig. 10.1(b) at the sweet spots is due to the fact that the simulation stepped over the exact sweet spot. In reality, though, T_ϕ would not actually be infinite at the sweet spot because of second order dephasing effects. I will also show in




Fig. 10.1 (a) Calculated noise transfer function df_{01}/dI_b for fluctuations in the bias current I_b . The dashed lines mark “sweet spots.” (b) Dephasing time associated with Johnson-Nyquist noise on the current bias line. Note that when df_{01}/dI_b goes to 0, T_ϕ becomes very large. The device parameters used for this plot are shown in Table 8.1 with $T = 20$ mK.

Sect. 10.4 that these sweet spots depend on I_b . Notice also that the worst T_ϕ was $\sim 100 \mu\text{s}$, which is very high.

I considered two types of flux noise in the dephasing analysis. One was thermal noise on the flux bias line and the other was $1/f$ flux noise from the environment. Both use the noise transfer function for applied flux $df_{01}/d\Phi_a$, but in different ways. The dephasing time associated with fluctuations in the flux bias line current I_Φ was:

$$\frac{1}{T_{\phi, I_\Phi}} \approx \frac{k_B T}{Z_0} \left(2\pi \frac{\partial f_{01}}{\partial I_\Phi} \right)^2 \approx \frac{k_B T}{Z_0} \left(2\pi \frac{\partial f_{01}}{\partial \Phi_a} M \right)^2 \quad (10.2)$$

where M is the mutual inductance of the flux bias line to the SQUID.

Figure 10.2 shows df_{01}/dI_Φ and T_{ϕ, I_Φ} vs. applied flux Φ_a for the design parameters of device PB9 (see Table 8.1). The magnitude of df_{01}/dI_Φ was much lower than was seen for df_{01}/dI_b , and since the thermal noise power on the two bias lines was the same, it is not surprising that the worst T_{ϕ, I_Φ} was ~ 100 ms. Figure 10.2 also displays sweet spots in df_{01}/dI_Φ at $\Phi_a = -0.5, 0$, and $0.5 \Phi_0$. The sweet spots at $\pm\Phi_0/2$ are well known for operation of flux qubits, and have allowed flux qubits to display coherence times $\sim 10 \mu\text{s}$, so they were expected to be there [75]. However, the sweet spot at $\Phi_a = 0$ is inaccessible to flux qubits, and I will show later in Sect. 10.4 why this sweet spot is actually superior to the ones at $\pm\Phi_0/2$ in terms of protecting the qubit against dephasing. From the long estimated dephasing times, it was clear that noise on the flux bias line was not a critical factor in the dephasing of the qubit.

The excess, low-temperature, $1/f$ flux noise from the environment around the qubit had to be dealt with in a different manner [63]. $1/f$ flux noise causes deviations in the phase of the qubit state according to (see Sect. 3.3 and [43]):




Fig. 10.2 (a) Calculated noise transfer function df_{01}/dI_{Φ} for fluctuations in the flux bias current I_{Φ} . The dashed lines mark “sweet spots.” (b) The dephasing time associated with Johnson-Nyquist noise on the flux bias line. Design parameters of device PB9 were used for the calculations (see Table 8.1).

$$\langle \phi_n^2(t) \rangle \approx \left(\frac{\partial \omega_{01}}{\partial \Phi_a} \right)^2 S_{\Phi_a}(1\text{Hz}) \ln \left(\frac{0.401}{f_r t} \right) t^2 \quad (10.3)$$

where $S_{\Phi_a}(1\text{Hz})$ is the noise power spectral density at 1 Hz in Wb^2 . In this case $\langle \phi_n^2(t) \rangle \propto t^2$ so it represents a Gaussian decay of the phase coherence, as would be seen in a Ramsey fringe experiment [43]. Because of the natural log in Eq. 10.3, there is not a precise value for the predicted decay constant T_R , but I can get an order of magnitude estimate from:

$$\frac{1}{T_R} \sim \left[\left(2\pi \frac{\partial f_{01}}{\partial \Phi_a} \right)^2 S_{\Phi_a}(1\text{Hz}) \right]^{1/2} \quad (10.4)$$

Figure 10.3 shows the noise transfer function for fluctuations in the applied flux and the resulting predictions for T_R for $1/f$ flux noise. Figure 10.3(a) shows the same sweet spots as Fig. 10.2(a) as expected and a maximum magnitude of about 4-5 MHz/m Φ_0 . Here I have assumed that the $1/f$ flux noise power spectral density at 1 Hz is $S_{\Phi_a}(1\text{Hz}) = (10 \mu\Phi_0)^2/\text{Hz}$ [63]. Fig. 10.3(b) shows that the worst possible T_R was 3 μs . This order of magnitude estimate for the dephasing time of the qubit was lower than what I found for any other process.

Noise in the critical current of the qubit junction was another form of $1/f$ noise that would cause dephasing in the qubit. The dephasing time estimate for critical current noise was found using:

$$\frac{1}{T_R} \sim \left[\left(2\pi \frac{\partial f_{01}}{\partial I_{01}} \right)^2 S_{I_{01}}(1\text{Hz}) \right]^{1/2} \quad (10.5)$$

The noise transfer function for I_{01} noise was found by calculating the dependence of f_{01} on I_{01} . The noise power spectral density was estimated according to [54]:

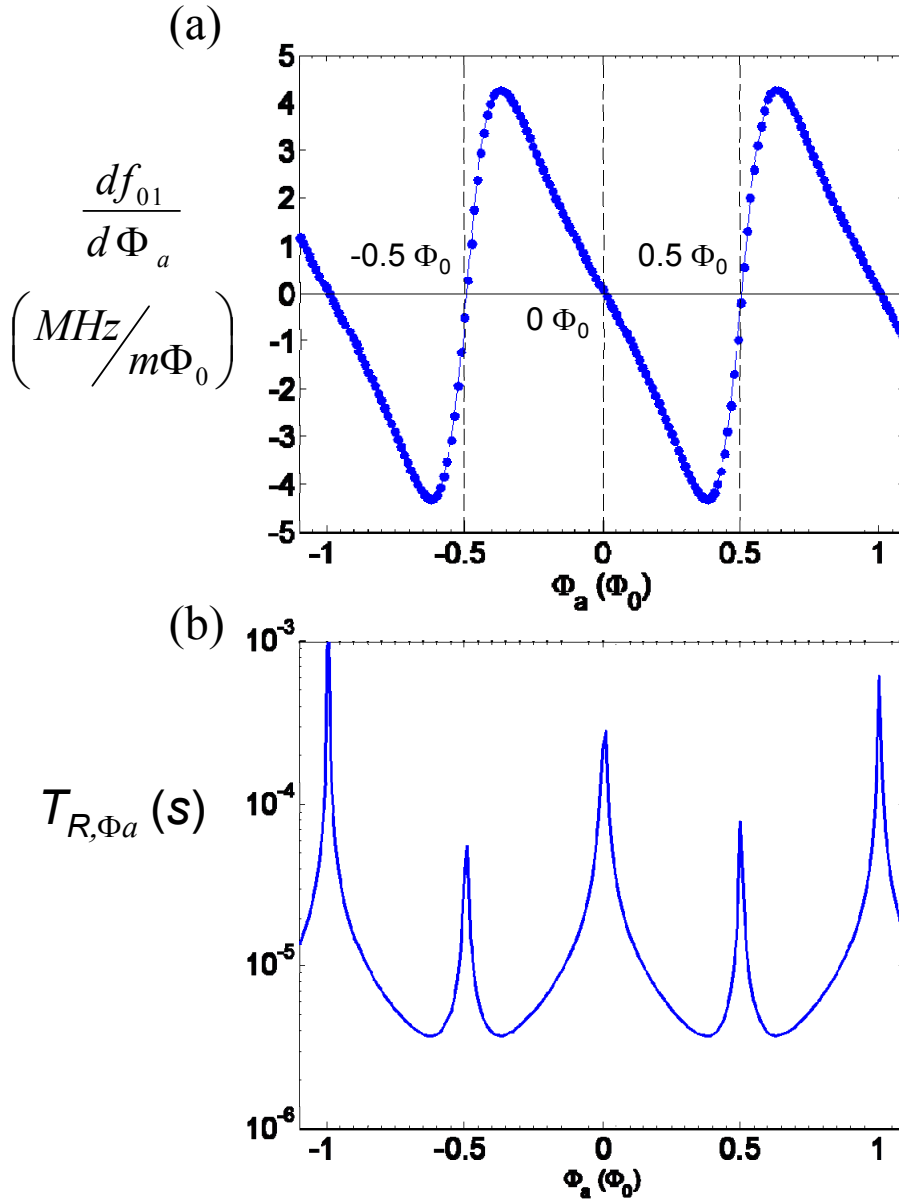


Fig. 10.3 (a) Calculated noise transfer function $df_{01}/d\Phi_a$ for fluctuations in the applied flux to the SQUID loop. Note that it is opposite in sign to Fig. 10.2(a) because the flux line is coupled in through a negative mutual inductance. The dashed lines indicate “sweet spots.” (b) The expected Ramsey decay time T_R associated with $1/f$ flux noise of $10 \mu\Phi_0/\sqrt{\text{Hz}}$ at 1 Hz. Design parameters of device PB9 were assumed (see Table 8.1).

$$S_{I_{01}}(1\text{Hz}) \approx 144 \frac{(I_{01}/\mu\text{A})^2}{A_{J1}/(\mu\text{m})^2} \left(\frac{T}{4.2\text{K}} \right)^2 (\text{pA})^2/\text{Hz} \quad (10.6)$$

where A_{J1} is the area of the qubit junction, T is the temperature, and the value for $S_{I_{01}}(1\text{Hz})$ is given in $(\text{pA})^2/\text{Hz}$ by Eq. 10.6. In device PB9, $I_{01} = 0.13 \mu\text{A}$, $A_{J1} = 0.23 (\mu\text{m})^2$, and $T \approx 20 \text{ mK}$; therefore $S_{I_{01}}(1\text{Hz}) \approx 240 (\text{fA})^2/\text{Hz}$. This was actually lower than the estimate used for the design because I_{01}^2/A_{J1} came out lower than expected.

Figure 10.4 shows df_{01}/dI_{01} and the estimate for T_R vs. applied flux. I note in particular the occurrence of sweet spots at $\Phi_a = \pm 0.35 \Phi_0$. This was the first known instance of a sweet spot that protected a superconducting qubit from fluctuations of the critical current. The sweet spots occur because the phase of the qubit junction γ_1 was equal to $\pm\pi/2$ at these points. This is important because the Josephson inductance of the qubit junction is:

$$L_{J1} = \frac{\hbar}{2e} \frac{1}{I_{01} \cos \gamma_1} \quad (10.7)$$

Thus, if $\gamma_1 \rightarrow \pi/2$, then $L_{J1} \rightarrow \infty$. The equation for the qubit resonant frequency is:

$$\omega_{01}^2 \approx \frac{1}{2} \left[\left(\frac{1}{L_1^* C_1} + \frac{1}{L_2^* C_2} \right) - \sqrt{\left(\frac{1}{L_1^* C_1} + \frac{1}{L_2^* C_2} \right)^2 - 4 \left[\left(\frac{1}{L_1^* C_1} \right) \left(\frac{1}{L_2^* C_2} \right) - \left(\frac{1}{L C_1} \right) \left(\frac{1}{L C_2} \right) \right]} \right] \quad (10.8)$$

and L_{J1} comes into the equations through L_1^* :

$$\frac{1}{L_1^*} = \left(\frac{1}{L_{J1}} + \frac{1}{L} \right) \quad (10.9)$$

If L_{J1} goes to infinity, then $L_1^* \approx L_1$, where L_1 is the geometrical inductance of the qubit arm of the SQUID. In that case, ω_{01} is independent of I_{01} and one obtains a sweet spot.

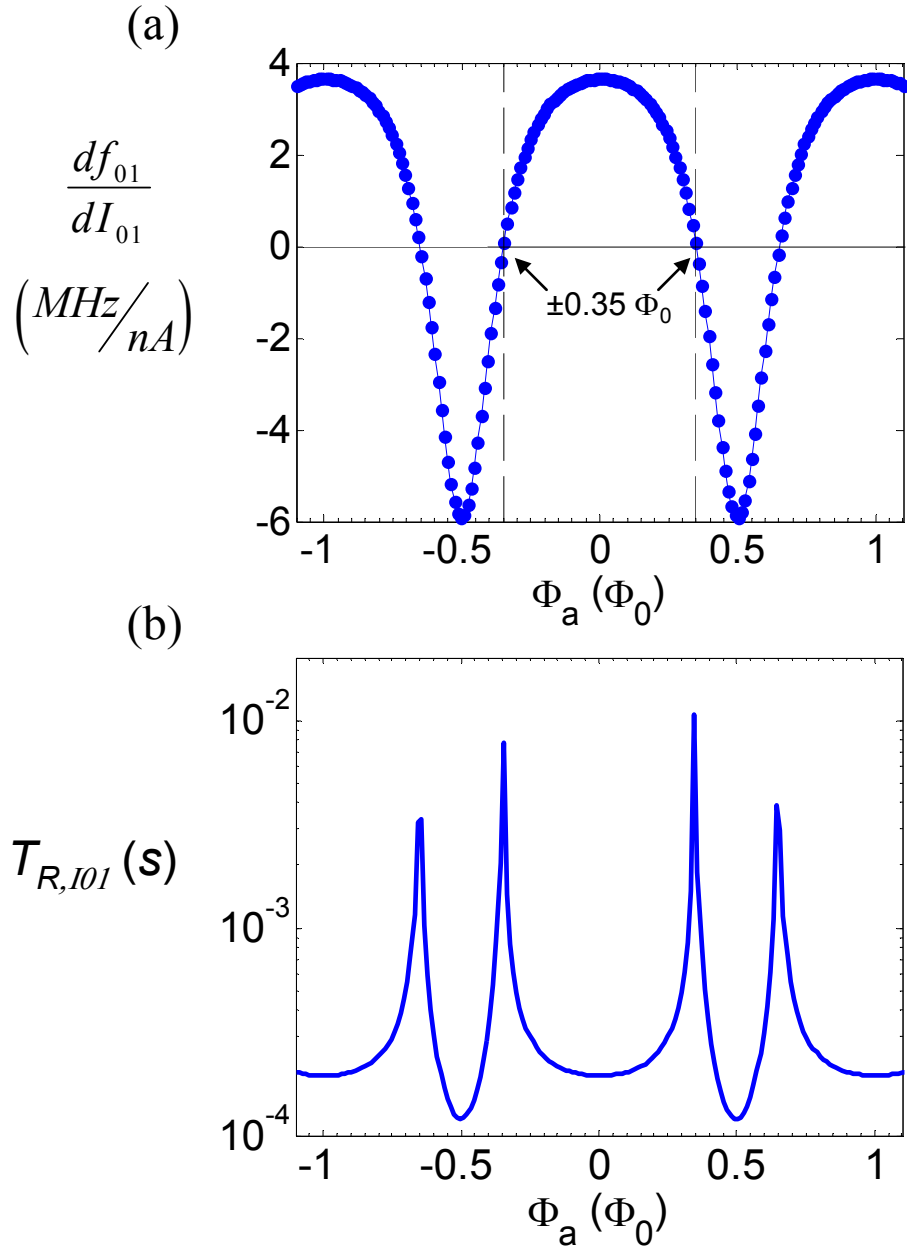


Fig. 10.4 Calculated noise transfer function df_{01}/dI_{01} for fluctuations in the critical current of the qubit junction I_{01} . The dashed lines indicate “sweet spots” for critical current noise (b) Calculated Ramsey decay time T_R associated with $1/f$ critical current fluctuations in the qubit junction. Design parameters of device PB9 were assumed (see Table 8.1).

Figure 10.4(b) shows that the dephasing time associated with fluctuations in I_{01} is very large around the sweet spots at $\Phi_a = \pm 0.35 \Phi_0$ (again, I would need to include second order effects to get an accurate T_R at the sweet spot) It also shows that the worst values are found at $\Phi_a = 0$ and $\pm 0.5 \Phi_0$, which are where the maximum and minimum frequencies in the spectrum are found. These are the points where ω_{01} is most strongly dependent on L_{J1} and thus I_{01} . At the worst bias points with respect to critical current noise the noise transfer function had a magnitude of about 4-6 MHz/nA. Although this was much higher than the transfer function for bias current noise, for example, the fact that the critical current noise power was so small kept the estimate for the dephasing time to 100 μ s or more.

There is one more source of dephasing that I did not considered in the original design: fluctuations in the critical current of the isolation junction I_{02} . Fluctuations in I_{02} cause fluctuations in the current flowing in the isolation junction I_2 , which create noise in γ_2 . Furthermore, they affect the division of I_b between the two arms of the SQUID, thereby directly affecting γ_1 . The estimate for the dephasing time from low frequency fluctuations in I_{02} takes the same form:

$$\frac{1}{T_{R,I_{02}}} \sim \left[\left(2\pi \frac{\partial f_{01}}{\partial I_{02}} \right)^2 S_{I_{02}} (1\text{Hz}) \right]^{1/2} \quad (10.10)$$

where [54]:

$$S_{I_{02}} (1\text{Hz}) = 144 \frac{(I_{02}/\mu\text{A})^2}{A_{J2}/(\mu\text{m})^2} \left(\frac{T}{4.2 \text{ K}} \right)^2 (\text{pA})^2 / \text{Hz} \quad (10.11)$$

In device PB9, $I_{02} = 8.6 \mu\text{A}$ and $A_{J2} \approx 5 (\mu\text{m})^2$. Since the ratio of I_{02} to A_{J2} was much larger than for the qubit junction, the noise power was about $0.05 (\text{pA})^2/\text{Hz}$ at 1 Hz. This was significantly larger than S_{I01} , and even though the state was mostly stored in the qubit junction, it was possible that S_{I02} could have affected the dephasing in the qubit state significantly.

Figure 10.5 shows df_{01}/dI_{02} and $T_{R,I02}$ vs. applied flux. Interestingly enough, df_{01}/dI_{02} shows its own sweet spots. Since df_{01}/dI_{02} was determined entirely from the model for finding ω_{01} , the cause of these sweet spots is not entirely clear. As expected, though, df_{01}/dI_{02} has relative maximums at the points where $df_{01}/dI_{01} = 0$. At these points, since ω_{01} does not depend on I_{01} , it depends most heavily on L_{J2} and thus I_{02} . Fortunately, $T_{R,I02}$ was between 1-10 ms at worst, so $1/f$ noise in the critical current of the isolation junction was not expected to be a significant source of dephasing, despite the fact that the qubit phase was coupled to the phase of the isolation junction.

I can now summarize the results of the dephasing analysis. Because thermal noise and $1/f$ noise affected the phase coherence of the qubit differently, it would not be correct to compare the dephasing times from each noise source on the same graph as I did in Chapter 9 for the different T_I 's. However, from my analysis, I expected the $1/f$ flux noise to be the most significant source of dephasing; the estimate for T_{R,Φ_a} was orders of magnitude lower than any other dephasing time. Flux noise in SQUIDs has been studied, and according to the results of those studies, $S_{\Phi_a}(1\text{Hz}) = (10 \mu\Phi_0)^2/\text{Hz}$ should have been a reasonable estimate for the flux noise [63]. Also, there were bias points near $\Phi_a = 0$ that should have shown very long term phase coherence if flux noise dominated the dephasing. Having the model for ω_{01} in terms of all of the parameters of the device allowed me to

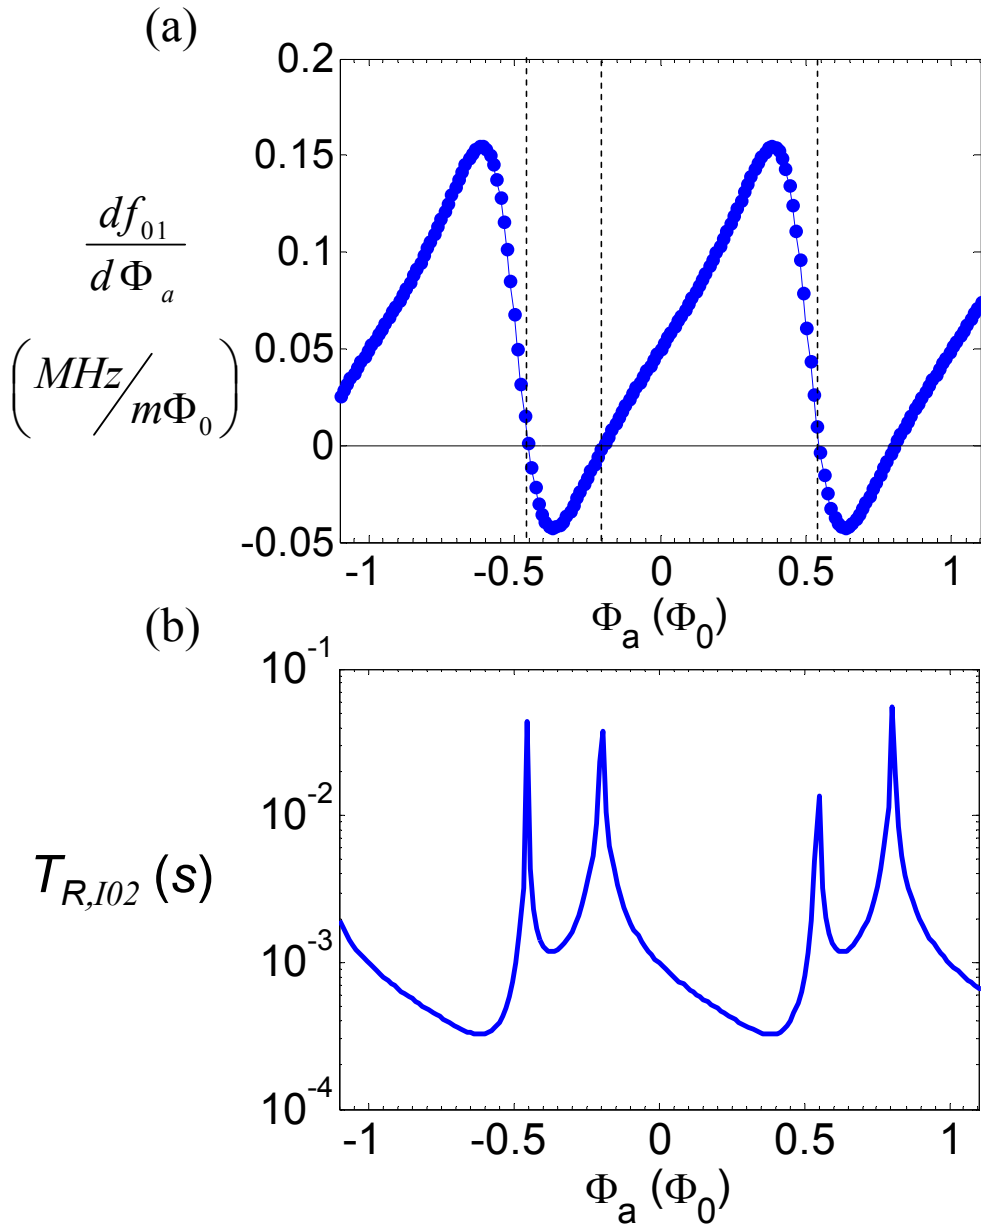


Fig. 10.5 (a) Calculated noise transfer function df_{01}/dI_{02} for fluctuations in the critical current of the isolation junction I_{02} . Note that the transfer function is not symmetric around 0 on the vertical axis. Changing the setting of the bias current I_b adjusts the vertical offset of the noise transfer function. Note also the presence of sweet spots for I_{02} noise. (b) Expected Ramsey decay time T_R associated with $1/f$ critical current fluctuations in the isolation junction.

get predictions for the dephasing in the qubit is device PB9 as a function of the flux bias and this turned out to be very useful for understanding dephasing in the device.

10.2 *Measurements of Dephasing in Device PB9*

I measured dephasing in device PB9 using several methods. First, I measured Rabi oscillations of the qubit state under a continuous microwave drive. From the decay of the Rabi oscillations I obtained T' and this allowed me to also estimate T_2 and T_ϕ . Second, to see how the dephasing varied with applied flux, I determined the width of the spectral peaks over the entire spectrum. I used the widths to determine the spectroscopic coherence time T_2^* of the qubit and compared these results with the dephasing estimates obtained using the noise transfer functions in Sect. 10.1. Finally, to check the validity of the T_2^* measurements, I performed a few pulsed Ramsey fringe experiments at specific flux bias points.

10.2.1 *Rabi Oscillation Measurement*

I measured Rabi oscillations in device PB9 using the procedure detailed in Chapter 7. Figure 10.6 shows an example of extracted P_I values that I obtained from Rabi oscillations using the S-curve fitting method described in Sect. 9.4. The data was taken at $f_{0I} = 4.047$ GHz, $I_\Phi = 0.209$ mA, and $I_b = 8.7$ μ A with $P_\mu = -50$ dBm while the device was at $T = 20$ mK. The Rabi oscillation has a bit of an odd shape, probably because of the difficulty of extracting P_I from the S-curves, so the Rabi decay time T'




Fig. 10.6 Rabi oscillation data from device PB9 with resonant microwaves applied at $f_{0I} = 4.047$ GHz with a power $P_m = -50$ dBm. The qubit biased at $I_b = 8.7 \mu\text{A}$, and $I_\Phi = 0.209$ mA. The blue dots are P_I as determined by fitting to S-curve data. The dotted magenta line is a decaying sine wave fit to the data with a decay constant of $T' = 150$ ns. The fit follows the troughs of the oscillation most closely because these points had the lowest χ^2 's associated with their extracted P_I values.

obtained from this data is an approximate value. The dashed line is a decaying cosine wave with a decay constant of $T' = 150$ ns.

The Rabi decay time T' is related to T_1 and T_2 by [52,76];

$$\frac{1}{T'} = \frac{1}{2T_1} + \frac{1}{2T_2} \quad (10.12)$$

where T_1 is the energy relaxation time and T_2 is the coherence time defined by:

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi} \quad (10.13)$$

I measured T_1 using the procedure detailed in Chapter 7, employing the same analysis I used to determine T_1 in Chapter 9. At the bias point where I measured the Rabi oscillation I found $T_1 \approx 250$ ns. Therefore, from Eqs. 10.12-13, the qubit had a coherence time $T_2 \approx 110$ ns and a dephasing time $T_\phi \approx 140$ ns. This was much shorter than the model predicted, but it was still long compared to all our previous devices except device DS8 which had a comparable T' [19,49,52,70].

During a Rabi oscillation, the state spends much of its time on the poles of the Bloch sphere where it is insensitive to dephasing [43]. The T_ϕ obtained from a Rabi oscillation is:

$$\frac{1}{T_{\phi,Rabi}} \approx \frac{1}{2} \left(2\pi \frac{df_{01}}{d\lambda} \right)^2 S_\lambda(\Omega_R/2\pi) \quad (10.14)$$

where λ represents the fluctuating parameter, and $S_\lambda(\Omega_R/2\pi)$ is the noise power evaluated at the Rabi frequency of ~ 10 MHz. In other words, during a Rabi oscillation, the qubit is relatively insensitive to dephasing from low frequency (*i.e.* $1/f$) noise. The theoretical limit for $T_{\phi,Rabi}$ (the dephasing time determined from T' using Eqs. 10.12 and 10.13) due to thermal I_b noise was ~ 100 μ s. The measured T_ϕ of only 140 ns suggests that there was

another source of noise in the 10 MHz frequency range that was not accounted for in the theoretical model.

The oscillation in Fig. 10.6 is representative of the longest T' that I measured on device PB9. Unfortunately, it was very time consuming to obtain good Rabi oscillations because I had to measure the S-curves many times during the oscillation to get P_I , and determine the decay rate of the oscillations.

10.2.2 *Spectroscopic Coherence Time T_2^* vs. Φ_a*

It was necessary to measure T_ϕ over a wide range of flux biases in order to see how it varied in order to identify the limiting source of dephasing for the qubit. I got a quicker but cruder measure of dephasing from the spectroscopy coherence time T_2^* . Unfortunately, T_2^* does not measure the same thing as $T_{\phi Rabi}$. T_2^* also includes effects of inhomogeneous (low frequency) noise.

Figure 10.7 shows the measured energy spectrum of the qubit. Note that in the range of $f_{0I} = 4.0-4.4$ GHz, the spectral peaks are quite narrow. It has been shown that the full width at half maximum (FWHM) Δf of each spectral peak obtained using a low power microwave drive can be used to determine the spectroscopic coherence time T_2^* of the state at each flux bias (see Sect. 4.3.3 of [26]) since:

$$T_2^* = \frac{1}{\pi \Delta f} \quad (10.16)$$

When dephasing is caused by relatively high frequency noise (e.g. white noise at Ω_R), then T_2^* is basically the same as the coherence time T_2 of the qubit. However, when the dephasing is dominated by the effects of very low frequency $1/f$ noise (as was




Fig. 10.7 0-to-1 transition spectrum of device PB9 versus applied flux. Note that the peaks are quite wide at the top (near 5 GHz) as compared to those near 4.2 GHz.

apparently the case for device PB9), the resonance peak will be broadened due to shot-to-shot variations in the device parameters and control parameters. This is often referred to as “inhomogeneous broadening” because of an analogous effect observed in Nuclear Magnetic Resonance (NMR) caused by the non-uniformity of the magnetic field. In this case, T_2^* cannot be used to obtain a dephasing time, *per se*; however, it does give a lower bound on the coherence time of the qubit [26] because the effects of inhomogeneous broadening can be counteracted by spin echo techniques [43].

I obtained Δf by fitting each spectral peak by fitting them all to a Lorentzian using a Levenberg-Marquardt algorithm [26,38,41,78]. The algorithm finds the best fit by varying four variables: the FWHM Δf , the offset y_0 from $P_{sw} = 0$, the peak position x_0 , and the peak height h . I allowed all four parameters to vary during the fit, and used approximate values as starting points for the algorithm. I assumed that the offset y_0 for each was about 3% because that was the threshold I used in choosing the measurement pulse height V_p for each resonance sweep (see Sect. 7.7 and 7.8 for measurement procedures). The height h of the peak, was assumed to be the maximum value for P_{sw} , and the peak position was taken to be the frequency value at which P_{sw} was at its maximum. Δf was initially set to 3 MHz at the beginning of each fitting algorithm. The code was able to fit all of the spectral peaks in less than 1 minute.

Figure 10.8 shows the resulting T_2^* vs. Φ_a . I note that there are prominent peaks in the spectroscopic coherence time at $\Phi_a = \pm 0.385 \Phi_0$. On the whole, the T_2^* of the device was rather low; over most of the bias range $T_2^* \approx 30$ ns. Since the measured T_1 of the qubit was 100-400 ns, these values were much lower than T_1 and much lower than the dephasing model predicted. This suggests that some additional inhomogeneous




Fig. 10.8 Spectroscopic coherence time T_2^* vs. applied flux Φ_a . $T_2^* = 1/\pi\Delta f$ was obtained from Δf , the full width at half maximum of Lorentzian fits to the energy spectrum peaks.

broadening was present.

From the general character of the graph one can understand the dominant source of broadening. Notice that there is a gradual increase of T_2^* as Φ_a approaches zero from either side. $\Phi_a = 0$ is a sweet spot for flux noise, but T_2^* was only about 20 ns at $\Phi_a = 0$, and this indicates the broadening is not due to fluctuations in Φ_a . The slight increase as $\Phi_a \rightarrow 0$ may have been due to the flux bias approaching the sweet spot, but it also could be explained as being due to variations in T_1 (T_1 was about 100 ns in the region and T_2^* was not much smaller). In any case, we can conclude that the short T_2^* is not consistent with flux noise.

In contrast, according to my analysis of the noise transfer functions, the only source of decoherence that had sweet spots that were symmetric around $\Phi_a = 0$ was critical current noise in the qubit junction (see Fig. 10.4). However, my estimate for the noise power associated with critical current fluctuations was far too low to explain the small value of T_2^* that I observed.

Clearly, the general shape of the T_2^* vs. Φ_a graph pointed to critical current noise as the dominant source of broadening with possibly some weak dependence on $1/f$ flux noise. A closer look revealed a small problem with that conclusion; using the parameters extracted from the device, the sweet spots in I_{01} were around $\Phi_a = \pm 0.35 \Phi_0$ (see Fig. 10.4(a)), whereas the measured peaks in T_2^* occurred at $\Phi_a = \pm 0.385 \Phi_0$. However, using the model in Sect. 8.5, I found that a variation in the parasitic inductance in the IDC along with an increase in I_{01} could push the sweet spots out from the center. Specifically, for these device parameters: $I_{01} = 155$ nA, $I_{02} = 8.6$ μ A, $C_I = 1.5$ pF, $L_I = 0.59$ nH, $L_2 = 40$ pH, $L_F = 0.13$ nH, $L_q = 0.22$ nH, and $I_b = 6.8$ mA, the sweet spot in I_{01} noise occurred




Fig. 10.9 The energy spectrum on the same scale as the T_2^* vs. Φ_a plot. The dashed magenta line is the theoretical prediction for the qubit frequency f_{01} using the normal modes model with device parameters: $I_{01} = 155$ nA, $I_{02} = 8.6$ μ A, $C_I = 1.5$ pF, $L_I = 0.59$ nH, $L_2 = 40$ pH, $L_F = 0.13$ nH, $L_q = 0.22$ nH, and $I_b = 6.8$ μ A. These parameters were the ones that gave the best fit of T_2^* to the dephasing model.

at $\Phi_a = 0.385 \Phi_0$. The variations in the parameters were within acceptable values as obtained by fitting the I - Φ , but Figure 10.9 shows that the fit to the spectrum was off some-what, so this parameter variation was not entirely self-consistent and must be considered a phenomenological explanation for why the critical current sweet spots could be farther out from the center. That is, increasing the parasitic inductance of the IDC in the model moved the I_{01} sweet spots out to match up with the peaks in T_2^* , but produced a small discrepancy between the predicted and measured spectrum.

Nevertheless, with this choice the fit to T_2^* captures the peaks in the coherence time at the critical current sweet spots. In order to obtain fit values for T_2^* , I first obtained the noise transfer functions (df_{01}/dI_b , $df_{01}/d\Phi_a$, df_{01}/dI_{01}) from the model using the new parameters. I accounted for the general variation of T_1 by assuming that T_1 was 100 ns if $f_{01} > 4.6$ GHz and 200 ns if $f_{01} < 4.6$ GHz. In the theoretical analysis of dephasing in the Sect. 10.1, I explained that dephasing from white noise is different from dephasing from $1/f$ noise, so I only considered $1/f$ flux and critical current noise in fitting T_2^* . Furthermore, $1/f$ flux noise was identified as the worst source of spectral broadening and the thermal noise on the bias lines was insignificant according to the model. I then obtained T_ϕ [43] from:

$$T_{\phi,fit} = \left[\sqrt{S_{I_{01}} (1\text{Hz}) \left(2\pi \frac{df_{01}}{dI_{01}} \right)^2} + \sqrt{S_{\Phi_a} (1\text{Hz}) \left(2\pi \frac{df_{01}}{d\Phi_a} \right)^2} \right]^{-1} \quad (10.19)$$

Strictly speaking, $T_{\phi,fit}$ includes contribution from pure dephasing and inhomogeneous broadening. Finally, I found, T_2^* from:

$$T_{2,fit}^* = \left[\frac{1}{2T_1(f_{01})} + \frac{1}{T_{\phi,fit}} \right]^{-1} \quad (10.20)$$

For the fit in Fig. 10.10, I used $S_{I_{0I}}(1\text{Hz}) = 10 \text{ (nA)}^2/\text{Hz}$ and $S_{\Phi_a}(1\text{Hz}) = (110 \mu\Phi_0)^2/\text{Hz}$. Both of these noise powers were much higher than expected. Flux noise is known to be $\sim(10 \mu\Phi_0)^2/\text{Hz}$ [77]. The flux noise power basically set the maximum T_2^* in the fit in Fig. 10.10 because I_{0I} noise was dominant in all other regions. It is possible that some other low frequency noise was contributing to the dephasing at the critical current sweet spots.

From the fit, I found the critical current noise power was higher than predicted by Eq. 10.11, by several orders of magnitude. $S_{I_{0I}}(1\text{Hz})$ basically set the value for T_2^* when $\Phi_a = 0$. It would require $10 \text{ (nA)}^2/\text{Hz}$ to make $T_2^* \approx 20 \text{ ns}$, whereas Eq. 10.11 predicted $S_{I_{0I}}(1\text{Hz}) = 240 \text{ (fA)}^2$. It is unlikely that the difference in the noise power could have been accounted for by another type of noise, especially considering that $\Phi_a = 0$ was a sweet spot for flux noise and very close to a sweet spot for bias current noise. Something else had to be going on with the critical current noise than was predicted by the Eq. 10.11.

Critical current noise is caused by single ionic charges that move in the tunnel barrier, thereby altering tunneling through the barrier [54]. Equation 10.11 predicts a $(T/4.2 \text{ K})^2$ dependence; At a temperature of 20 mK, most of the trapped charges will be frozen in place and this greatly reduces the critical current noise. However, the data that gave rise to Eq. 10.11 for the critical current noise was obtained by observing shunted junctions that were biased at low voltage (a few μV) [77]. In contrast, between every measurement of the qubit I made, the junction switched to the voltage state ($\sim 360 \mu\text{V}$). The switching of the junction would have released a relatively large amount of energy and caused the junction temperature to rise to nearly 1K) heat This would have allowed trapped charges to rearrange themselves on an energy scale that would not have occurred




Fig. 10.10 Fit of the dephasing model (dashed curve) to the measured T_2^* data (solid curve). The fit assumes (3 nA^2/Hz of $1/f$ noise at 1 Hz in I_{0I} and $(110 \mu\Phi_0)^2/\text{Hz}$ at one Hz of $1/f$ flux noise. Both noise figures are very large compared to their theoretical values, but critical current noise is the only dephasing source that predicts peaks in T_2^* that are symmetric around $\Phi_a = 0$. The high noise power may be due to switching causing an elevated effective temperatures of the noise sources.

in a 20 mK environment. This suggests that I could try replacing $T = 20$ mK by $T \sim 2\Delta/ek_B$. This would yield an expected critical current noise power spectrum $\sim (3/0.02)^2 \sim 10^4$ times larger or $2.4 \text{ (pA)}^2/\text{Hz}$. Each measurement of the spectral peak was made up of $\sim 10^5$ measurements taken with a 700 Hz repetition rate, so this shot-to-shot noise could be responsible for some of the broadening. The variations could also have been due to the large voltage across the junction during the switching, which could have induced changes in the defect ions' positions.

The small area of the qubit junction could also account for some of the disparity between the measured broadening and the predicted I_{0I} . It has been shown that in Al/AlO_x/Al junctions with areas $\sim 0.1 \text{ (}\mu\text{m)}^2$ that the $S_{I0}(1 \text{ Hz}) \sim (I_0 \cdot 10^{-4})^2$ [17,79]. For device PB9, $I_{0I} = 130 \text{ nA}$, so $S_{I0I}(1 \text{ Hz}) \sim 100 \text{ (pA)}^2$. Neither of these explanations appear to fully explain the noise. However, perhaps the combination of the small area and the switching of the junction between measurements could explain the large amount of I_{0I} noise observed in device PB9.

10.2.3 Ramsey Measurements

Measuring T_2^* from the spectral peak widths was the quickest way to characterize the dephasing time of the qubit over a wide range of fluxes. However, a better way to characterize the coherence times was to measure the Ramsey fringe decay time T_R . Since T_R is basically the same as T_2^* [26,43,80], I could use this as a check against the T_2^* determined from the spectral widths at a few flux bias points.

I performed Ramsey fringe measurements following the procedure I described in Sect. 7.13. The Ramsey pulse sequence consisted of an in-phase $\pi/2$ pulse to rotate the

state onto the equator of the Bloch sphere, followed by a delay period τ , during which the state was maximally affected by dephasing, and it finished with another in phase $\pi/2$ pulse. If there was no dephasing, this would rotate a state that started in the ground state (before the first $\pi/2$ pulse) to $|1\rangle$ (after the second $\pi/2$ pulse). I set the microwave drive so that it was detuned from f_{0I} by some amount Δf , which caused the probability of the final state to be in the $|1\rangle$ state to oscillate with frequency of Δf . I measured P_I after the second pulse, and the oscillation of P_I should be expected to decay to 50% with a Gaussian envelope as the delay τ was increased.

Because of the anomalous switching behavior of device PB9 (see Sect. 9.4), I measured the S-curve of the state after the second $\pi/2$ pulse and used these curves to extract P_I . The $P_{sw,|0\rangle}$, $P_{sw,|1\rangle}$, and $P_{sw,\pi/2}$ S-curves were taken from those obtained during the calibration of the pulses (see Sect. 7.13). Figure 10.11 shows the extracted P_I vs. delay time τ from a Ramsey measurement taken at a flux bias point of $\Phi_a = -0.385 \Phi_0$ and a qubit frequency $f_{0I} = 4.218$ GHz with a detuning of 12 MHz. This was the critical current noise sweet spot and the location of the spike in T_2^* . The oscillation had a decay constant T_R of about 200-240 ns. This was consistent with $T_2^* = 215$ ns that I found from spectroscopy at this point.

Figure 10.12 shows a comparison of the S-curves measured during the Ramsey measurement performed at $\Phi_a = -0.385 \Phi_0$ ($f_{0I} = 4.218$ GHz) and at $\Phi_a = 0$ ($f_{0I} = 5.050$ GHz). The oscillations in the S-curves seen in Fig. 10.12(b) decayed so quickly that I was unable to extract the P_I 's for that measurement (the *ad hoc* model for extracting P_I from the S-curves had difficulty fitting the S-curve of a dephased state). Regardless, it




Fig. 10.11 Ramsey fringe oscillation taken at $\Phi_a = -0.385 \Phi_0$ where the qubit showed a distinct peak in T_2^* . The blue dots are extracted P_I from S-curve measurements. The dashed magenta line is a decaying sine wave with a Gaussian envelope and a time constant of $T_R = 240$ ns. Either way, the findings are consistent with the T_2^* obtained from the spectral peak width. For this measurement, $f_{0I} = 4.218$ GHz and the microwave drive was $f_\mu = 4.236$ GHz so the detuning was 12 MHz.




Fig. 10.12 S-curves taken during Ramsey measurements at (a) $\Phi_a = 0$ and (b) $\Phi_a = -0.385 \Phi_0$.

The x-axis is the delay time τ between the two $\pi/2$ pulses. The pulses used in the

measurement of (a) were detuned by 17 MHz and those for (b) were detuned by 12 MHz.

Note that the oscillations in (b) are much longer lived than those in (a). The oscillations in (a)

were of such poor quality that it was not possible to extract P_I from the S-curves. The T_2^*

obtained from the spectral peaks predicted $T_R = 20$ ns for (a) and $T_R = 220$ ns for (b). These

plots help to confirm the spectroscopic results for T_2^* vs. Φ_a in Fig. 10.8.

was apparent that the Ramsey oscillation was much shorter when the qubit was biased at $\Phi_a = 0$. These results were also consistent with my measurement of $T_2^* \approx 30$ ns from the spectral peak widths. Thus, I found that the Ramsey measurements and spectral broadening measurements were consistent with each other.

I could have been more certain of the conclusion that the phase coherence of the qubit was being limited by shot-to-shot variations in I_{0l} if I had performed Ramsey measurements over the whole range of the spectrum. However, since it would have been necessary to calibrate the $\pi/2$ pulses for each measurement, the experiment would have taken exceedingly long.

10.3 *Future Improvements on Dephasing*

The presence of sweet spots in device PB9 was a promising development. Sweet spots in the energy levels have been exploited in charge and flux qubits to reduce dephasing [44,75,81-83]. Charge qubits have a charge sweet spot in the energy levels when the gate bias V_g equals e/C_g , where C_g is the gate capacitance [81-83]. Flux qubits employ a flux sweet spot at $\Phi_a = \Phi_0/2$ [44]. There have qubits exhibited some of the longest coherence times ($T_2 \approx 20$ μ s) [75] and relaxation times ($T_1 \approx 200$ μ s) [84].

The situation for device PB9 is quite favorable when it comes to the impact of flux noise. Consider the noise transfer function for flux $\frac{df_{0l}}{d\Phi_a}$ for device PB9 (see Fig. 10.13). The noise transfer function is similar to that of a flux qubit; at least around $\Phi_a = \Phi_0/2$. Note that although $df_{0l}/d\Phi_a = 0$ at $\Phi_a = \Phi_0/2$, it increases rapidly away from this point. To illustrate how sharp this sweet spot is, Fig. 10.13(b) shows the slope of the

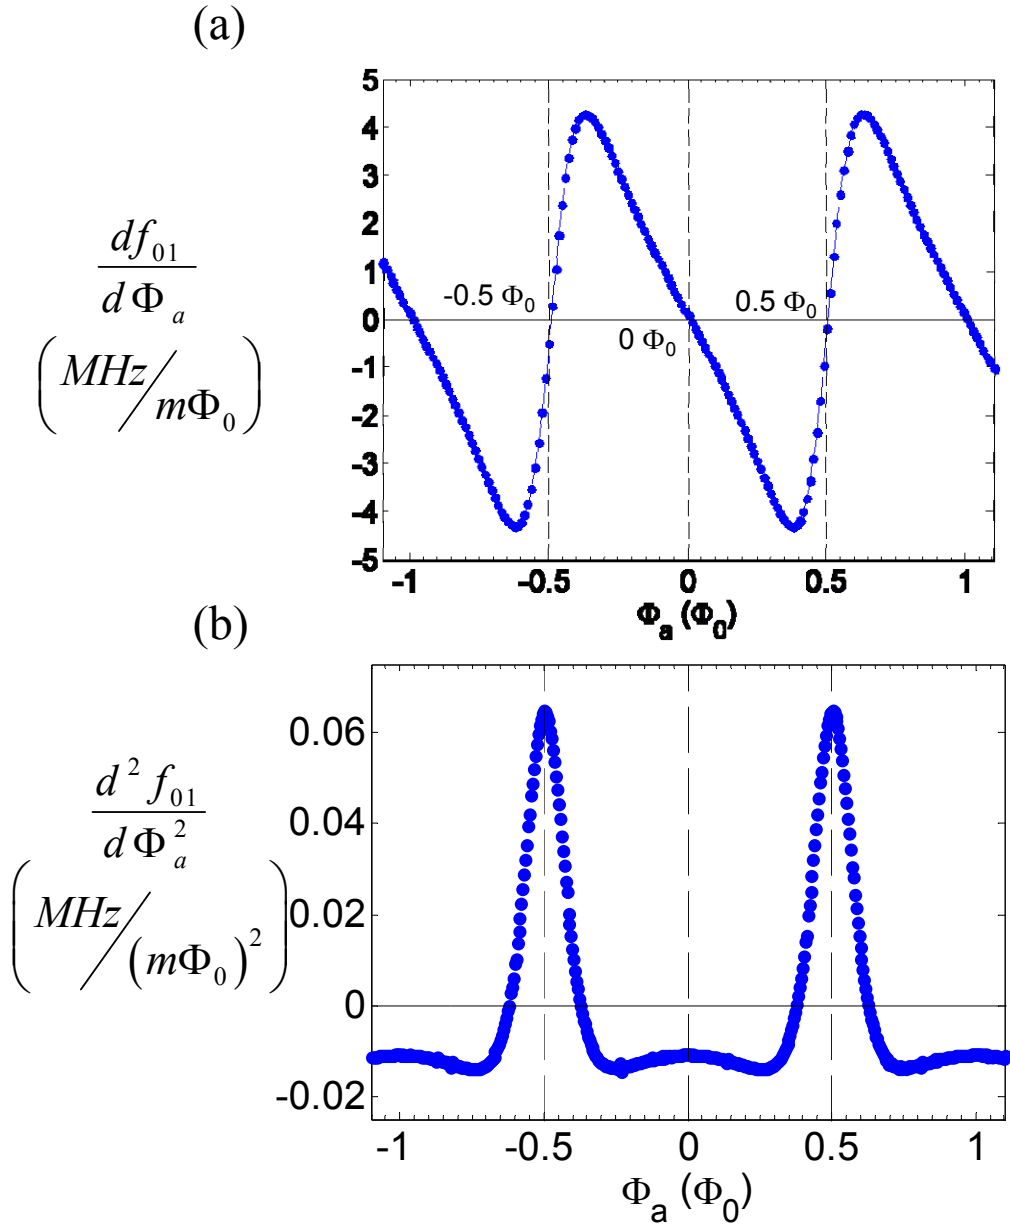


Fig. 10.13 (a) Noise transfer function for fluctuations in the applied flux Φ_a . (b) The slope of the noise transfer function. Note that the slope is very high at the sweet spots at $\Phi_a = \pm\Phi_0/2$, whereas it is relatively low at $\Phi_a = 0$. The $\Phi_a = 0$ sweet spot is available for operation of a phase qubit such as device PB9, but not for flux qubits.

noise transfer function, $d^2f_{01}/d\Phi_a^2$. Note that it has its maximum magnitude when $\Phi_a = \Phi_0/2$ and this peak occurs within $\pm\Phi_0/10$. The transfer function is responsible for dephasing at the sweet spot.

In contrast, consider the sweet spot exhibited in device PB9 at $\Phi_a = 0$ (see Fig 10.13). Of course, $df_{01}/d\Phi_a$ is zero at this bias, but notice also that $d^2f_{01}/d\Phi_a^2$ was much lower at this sweet spot than at the $\Phi_a = \Phi_0/2$ sweet spot. In fact, $d^2f_{01}/d\Phi_a^2$ shows a relative minimum at $\Phi_a = 0$. Thus, the second order contribution to the dephasing will be much less at $\Phi_a = 0$ than at $\Phi_a = \Phi_0/2$.

10.4 Anharmonicity

In Sect. 10.2, I showed that the sweet spot for bias current noise was actually shifted a bit away from $\Phi_a = 0$. This was because the sweet spot depends on the bias current. Figure 10.14 shows df_{01}/dI_b when $I_b = 0$. For this bias, the bias current sweet spot lies directly on top of the flux sweet spot. This means that the unbiased qubit is at an optimum bias point for current and flux variations. In contrast, for the data shown in Fig. 10.8, I used $I_b = 6.8 \mu\text{A}$, and this caused the current sweet spot to shift to $\Phi_a = -0.07 \Phi_0$ according to the simulation.

Strictly speaking, the zero-bias current sweet spot has always existed in phase qubits. At $I_b = 0$, the energy levels of a phase qubit are independent of I_b to first order. However, a typical phase qubit cannot be operated at $I_b = 0$ because the energy levels are almost harmonic (see Sect. 2.4). To address this concern in device PB9, I examined the anharmonicity of the energy levels. As a practical test of the anharmonicity of the qubit, I measured Rabi oscillations of the qubit state at $\Phi_a = 0$. During this test, I




Fig. 10.14 Noise transfer function for bias current fluctuations when I_b is set to zero.

Note that the sweet spot in the middle lines up at $\Phi_a = 0$. Changing the setting of I_b changes the location of the sweet spots. Setting $I_b = 0$ puts the sweet spots right on top of the sweet spots for fluctuations in the applied flux.

set $I_b = 6.8 \mu\text{A}$ as usual. Figure 10.15(a) shows raw S-curve data taken during a Rabi oscillation, and Fig. 10.16(b) shows the extracted P_I vs. time. The qubit state appears to transition back and forth between $|0\rangle$ and $|1\rangle$, much as expected. The decay constant was $T' \approx 130$ ns. At this bias I also found $T_I \approx 110$ ns, and hence $T_2 \approx 160$ ns and $T_\phi \approx 600$ ns. Note that although this decay constant suggests that T_ϕ is long the value is quite uncertain because the Rabi decay time is much shorter. I deliberately performed this Rabi oscillation with a high power microwave drive so that it would oscillate with a frequency of $f_R > 50$ MHz. If there had been a problem with the anharmonicity of the qubit's energy levels, then the high power should have produced some population in higher levels, *i.e.* $n = 2, 3, \text{etc.}$ [38]. Although I did not find the anharmonicity of the state at $\Phi_a = 0$, this supports the conclusion that the levels are not evenly spaced. In particular, one would expect that the switching would occur at lower and lower current as the power increased if the system were climbing to higher and higher states as the power increased. Note in Fig. 10.16(b) that there may be some tendency to switch at lower currents at the peak of the Rabi oscillation (π -pulse), but that the switching occurs at practically the same value for $2\pi, 4\pi, 6\pi$, and so on. This suggests that upper levels were not being populated.

There is some additional evidence for anharmonicity at the I_{0I} sweet spot. Figure 10.16 shows the S-curves of the saturated state taken after applying microwave drives of varying powers for 3 μs . The bias conditions were, $I_b = 6.8$ mA, $\Phi_a = -0.385 \Phi_0$, and $f_{0I} = 4.218$ GHz (the I_{0I} sweet spot). We see that switching occurs at lower and lower V_p as the microwave power is increased. This is what you would expect in the low power limit.




Fig. 10.15 Rabi oscillations taken at $\Phi_a = 0$, $f_{0I} = 5.060$ GHz, and $P_\mu = -35$ dBm. (a) Raw S-curve data of the oscillation. If the phase qubit energy levels were harmonic at zero bias, then one would not expect to see a Rabi oscillations. (b) Extracted P_I vs. time. The dotted line is a decaying sign wave with a decay constant, $T' = 130$ ns.




Fig. 10.16 S-curves taken at $I_b = 6.8$ mA, $\Phi_a = -0.385 \Phi_0$, and $f_{0I} = 4.218$ GHz (I_{0I} sweet spot) of the saturated state of the qubit after applying resonant microwave drive with power P_μ . Note that the qubit tends to switch at lower and lower V_p as P_μ increases up until about -40 dBm, and then levels off. This is evidence that the qubit is not populating higher states at higher powers.

Furthermore, at high power it flattens off, which is consistent with saturation. It is not clear why the S-curves tend to shift in the opposite direction at $P_\mu > 30$ dBm.

I also used the finite difference method (see Sect. 4.5) to calculate the expected anharmonicity of the levels at $\Phi_a = 0$ and $I_b = 0$. I used the parameters that gave the best fit to the spectrum using the finite difference method in Sect. 8.5, a grid spacing of 0.024, and a $51^2 \times 51^2$ grid. The model gave me the energies of all of the states, and I found $f_{01} = 5.15$ GHz and $f_{12} = 5.14$ GHz. Thus, at $I_b = 0$, the model predicted an anharmonicity of about 10 MHz. With the same parameters and $I_b = 6.8$ μ A, I found that the anharmonicity was about 20 MHz. This is quite a small anharmonicity and suggests that driving with $f_R \approx 20$ MHz should lead to significant population in the upper levels.

I attempted to perform spectroscopy on higher level states, but the results were inconclusive. The question of anharmonicity came up rather late in the experiment, and I only made a few attempts to observe higher levels in the spectrum. In contrast, I obtained the Rabi oscillation data shown in Fig. 10.15 many months earlier. This was significant evidence against the qubit having a problem with anharmonicity, so I did not re-examine the issue until much later. I suspect that the model might be incomplete, and the anharmonicity prediction needs to be carefully examined. I also note that H. Kwon seemed to observe a clear and well-separated 1-to-2 transition in a device with similar parameters [70].

It would have helped to clear up this issue of anharmonicity in device PB9 if I had performed a study of the Rabi frequency's dependence on the microwave power. A sufficiently anharmonic system's Rabi frequency Ω_R should increase as the square root of the microwave power P_μ [38]. If higher levels are being populated during the Rabi

oscillation, then slope of a plot of Ω_R vs. P_μ should decrease at higher powers [38]. Unfortunately, due to the anomalous switching behavior, this study would have taken a long time to perform on device PB9.

There is one potential problem with working at $\Phi_a = 0$: sensitivity to critical current noise. In particular, in Device PB9, the limitation on T_2^* was consistent with $1/f$ critical current fluctuations. This problem probably came from the fact that I used a tunneling probability measurement to measure the state of the qubit. After the junctions switched to the voltage state it probably was left with a slightly different critical current because of the generation of heat, quasiparticles, and a large \vec{E} -field that occurred during the switch. A dispersive readout [80,84-86] in which the device does not switch during the readout, should eliminate this problem.

10.5 Conclusions

In this chapter, I discussed the dephasing analysis of device PB9. Using the same model I used to fit the spectrum, I obtained the noise transfer functions for all the sources of dephasing discussed in Chapter 5. Surprisingly for a phase qubit, the noise transfer functions for I_b , I_{01} , I_{02} , and Φ_a exhibited sweet spots where the qubit was insensitive to that type of noise. In particular, the qubit even displayed a sweet spot for critical current fluctuations. My analysis of the dephasing showed that $1/f$ flux noise should have been the dominant source of dephasing; however, this was not the case

Measurement and analysis of the spectroscopic coherence time T_2^* revealed that critical current noise in the qubit junction appeared to be limiting the line width. I fit the data obtained for T_2^* to the dephasing model detailed in Sect. 10.1, and found the noise

power of the critical current fluctuations was much larger than expected from experiments on other junctions [77]. This enhanced noise power may have been due to shot-to-shot variations in the critical current that occurred as a result of the SQUID switching to the voltage state during the readout. I also found a much larger than expected flux noise contributions.

In the future, this type of phase qubit could be used, with a “non-switching” readout to eliminate these fluctuations. In particular, device PB9 should be very insensitive to flux and bias current fluctuations at $\Phi_a = 0$, and although the anharmonicity of this qubit is not well understood, I found I could operate the device at this bias point. At the very least, the discovery of new sweet spots in the energy levels of this phase qubit presents new avenues for pinning down sources of decoherence in phase qubits.

Chapter 11

Conclusions

The overall goal of my dissertation research was to reduce decoherence in dc SQUID phase qubits. Although my device did not achieve a relaxation time longer than 1 μs , to the extent that device PB9 showed a better T_1 and T_2 than any device previously produced by our group, I made progress. I did this by giving consideration to the impact of each of the device parameters on the relaxation time and coherence time of the qubit. Table 11.1 summarizes the best results I obtained, as well as the design values.

The dissipation from the bias lines was modeled as an effective admittance across the qubit junction, and the loss from the capacitor dielectrics was estimated through their loss tangents. By reducing the area of the qubit junction to $0.23 (\mu\text{m})^2$ and shunting the qubit with an interdigitated capacitor (IDC), I reduced the number density of two-level systems coupling to the qubit. I did observe two splittings in 1 GHz, which were probably due to circuit resonances. I designed the device so that it should have been isolated from dissipation from the bias leads (see Tables 5.1 and 5.2) and given $T_1 \leq 8.4 \mu\text{s}$ using an isolated junction model.

Table 11.1 Summary of best measured times for device PB9.

Measurement	Time Constant	Theoretical Limit (μs)	Measured Value (ns)
Energy Relaxation	T_1	8.4	350 ± 70
Rabi Oscillation	T'	11	150
	T_2	17	110
	T_ϕ	~ 100	140
Spectral Peak Width	T_2^*	~ 1	215
Ramsey Fringe	T_R	~ 1	220 ± 20

The best energy relaxation time I measured for device PB9 was $T_1 = 350 \pm 70$ ns (see Fig. 9.7). The uncertainty in the value came from the presence of anomalous switching behavior that made it difficult to determine P_I precisely with the switching readout [64,70]. Although T_1 was longer than in previous devices made by the group, it was much shorter than the limit obtained from the theory even after accounting for the normal modes at the SQUID. By measuring T_1 vs. applied flux Φ_a , I discovered a correlation between the π -pulse power and T_1 . When it took less power to perform a π -pulse, T_1 was shorter (see Fig. 9.10). This was evidence that the limiting source of dissipation was over-coupling to the microwave line. The nature of the coupling of the microwave line to the qubit on device PB9 was poorly understood; therefore, our next qubit (currently being designed by Rangga Budoyo) will use a coplanar waveguide, with well known coupling to the qubit, to deliver the microwave drive. Reducing this coupling should increase the T_1 of the new qubit.

Although the model I used in designing device PB9 assumed that the qubit was an isolated junction, the data revealed that the junctions in the SQUID were coupled together to produce normal modes of the device. In particular, the 0-to-1 transition spectrum never went to $f_{01} = 0$ as it would have for a single junction. This led to a reexamination of the Hamiltonian of the full SQUID. Since the inductance of the loop $L_1 + L_2$ was much lower than the inductance of the qubit junction L_{JI} , the two junctions needed to be treated as coupled oscillators. I fit the measured spectrum to two models: a quasi-classical model that expanded the junction phases γ_1 and γ_2 to third order in small variations around the minimum (cubic approximation) and a finite difference method to numerically solve for

the energies of the SQUID. I found small disagreement with the data. After accounting for the parasitic inductance of the IDC, I achieved a good fit to the spectrum using device parameters (see Table 8.1) that were consistent with those obtained from the I - V and I - Φ characteristic plots.

A surprise from the dephasing analysis of PB9 was the occurrence of unexpected sweet spots. I modeled the dephasing in the system as coming from both thermal noise on the bias leads and $1/f$ fluctuations in flux and critical current. Using the SQUID Hamiltonian models described above, I found that the noise transfer functions for fluctuations in the device parameters and control parameters vanished under certain biasing conditions. Since the dephasing time of the qubit was inversely proportional to the noise transfer function squared, the dephasing time associated with that type of noise would have been infinitely long (neglecting second order terms) at those sweet spots. The theoretical model predicted sweet spots for bias current, flux, and critical current fluctuations, and it predicted that $1/f$ flux noise should have limited T_2^* to $\sim 1 \mu\text{s}$.

I measured the coherent behavior of the qubit in several ways. From the Rabi oscillation decay time $T' = 150 \text{ ns}$, I found a coherence time $T_2 = 110 \text{ ns}$ and a dephasing time $T_\phi = 140 \text{ ns}$ using

$$\frac{1}{T'} = \frac{1}{2T_2} + \frac{1}{2T_1} \quad \text{and} \quad \frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi} \quad (11.1)$$

This result is only valid for a frequency independent dephasing, or equivalently this give the dephasing rate $1/T_\phi$ at the Rabi frequency. Since $1/f$ noise was prevalent in device PB9, the Rabi oscillation should have decayed with a Gaussian envelope. Therefore, the Eq. 11.1 would not be expected to hold.

Because of the anomalous switching behavior, it was time consuming to measure Rabi oscillations in device PB9, and I made relatively few measurements of the Rabi decay time T' . Instead, by measuring the spectral peak width, I obtained the spectroscopic coherence time T_2^* over a wide range of flux biases. Although T_2^* tends to be dominated by inhomogeneous broadening, I did observe peaks in T_2^* that were indicative of sweet spots with respect to $1/f$ critical current noise in the qubit junction. When biased at these points the qubit displayed $T_2^* \approx 210$ ns, whereas $T_2^* \sim 10$ ns otherwise. I fit the T_2^* vs. Φ_a data and determined that the critical current noise in the qubit was $10 \text{ (nA)}^2/\text{Hz}$ at 1 Hz and the flux noise was $\sim (100 \mu\Phi_0)^2/\text{Hz}$ at 1 Hz. These $1/f$ noise levels were much larger than expected from measurements on other junctions. This may be due to the fact that the junction switched to the voltage state between each measurement, leading to a disturbance of the charge defects in the junction [54,63,77]. The extremely small size of the junction may have also played a role. Measurements of the Ramsey fringes of the qubit gave a time constant of $T_R = 220 \pm 20$ ns at the sweet spot, which supported the validity of the measured spectroscopic coherence times; *i.e.* one expects $T_R = T_2^*$.

The dissipation and dephasing models can be used in the future to design phase qubit for improved coherence; however, it is important that the coupling of each line to the qubit be well understood. The critical current sweet spots discovered in device PB9 were the first observed in any superconducting qubit, and although the anharmonicity of the qubit is an unsettled issue, they can be used to identify limiting sources of dephasing dc SQUID phase qubits in this parameter regime.

Appendix A

Matlab Code

1) *fullsquiddeltasolve.m* (cited in Ch. 4, p. 80)

This program was used to find the eigenenergies and wavefunctions of a dc SQUID by using a finite difference method. It would be best to copy and paste into Matlab to view it.

```
clear all
%%Define constants%%
Phi0 = 2.07e-15;
e=1.6e-19;
hbar=1.06e-34;
h = 6.626e-34;
%%Define device parameters%%
I01 = 110e-9;
I02 = 8.6e-6;
Ib = 6.8e-6;
C1 = 1.52e-12;
C2 = .225e-12;
L1 = .7e-9;
L2 = 40e-12;
L = L1+L2;
Ej1 = Phi0*I01/(2*pi);
Ej2 = Phi0*I02/(2*pi);
Ec1 = e^2/(2*C1);
Ec2 = e^2/(2*C2);

%%Set constant applied flux
Phia = -.295*Phi0;

%%Plot out potential to find initial minimum
g1 = [-.8:.01:.65];
g2 = [0:.01:1.65];
U = [];
for k = 1:length(g1)
    for j = 1:length(g2)
        U(k,j) = -Ej1*cos(g1(k)) - Ej2*cos(g2(j)) -
        (Phi0*Ib/(2*pi*(L+LF))) * (L2*g1(k) + (L1+LF)*g2(j)) + (Phi0/(2*pi))^2 * (1/(2*
        (L+LF))) * (g1(k) - g2(j) - 2*pi*Phia/Phi0)^2;
    end
end
figure(10)
color2D(g2,g1,U)
shading flat
%%Use ginput to pull off the initial minimum
```

```

        g0=ginput;
        g20=g0(1);
        g10=g0(2);

        %%Minimize the potential
        U=@(g1) -Ej1*cos(g1)-Ej2*cos(g20)-
        (Phi0*Ib/(2*pi*(L+LF)))*(L2*g1+(L1+LF)*g20)+(Phi0/(2*pi))^2*(1/(2*(L+LF)
        ))*(g1-g20-2*pi*Phia/Phi0)^2;
        g10=fminbnd(U,-10,10);
        U=@(g2) -Ej1*cos(g10)-Ej2*cos(g2)-
        (Phi0*Ib/(2*pi*(L+LF)))*(L2*g10+(L1+LF)*g2)+(Phi0/(2*pi))^2*(1/(2*(L+LF)
        ))*(g10-g2-2*pi*Phia/Phi0)^2;
        g2min=fminbnd(U,-.2,1.57);
        U=@(g1) -Ej1*cos(g1)-Ej2*cos(g2min)-
        (Phi0*Ib/(2*pi*(L+LF)))*(L2*g1+(L1+LF)*g2min)+(Phi0/(2*pi))^2*(1/(2*(L+L
        F)))*(g1-g2min-2*pi*Phia/Phi0)^2;
        g1min=fminbnd(U,-10,10);

%% Set grid spacing
gstep = .03;
%% Set number of states (states < 51)
states = 39

%% Construct the grid
g1=[g1min-gstep*(states-1)/2:gstep:g1min+gstep*(states-1)/2];
g2=[g2min-gstep*(states-1)/2:gstep:g2min+gstep*(states-1)/2];
dg1 = g1(2)-g1(1);
dg2 = g2(2)-g2(1);
U=zeros(length(g1)*length(g2));

%% Construct the U matrix
for i = 1:length(g1)
    for j = 1:length(g2)
        m = (i-1)*length(g2)+j;
        Umm = -Ej1*cos(g1(i))-Ej2*cos(g2(j))-
        (Phi0*Ib/(2*pi*(L+LF)))*(L2*g1(i)+(L1+LF)*g2(j))+(Phi0/(2*pi))^2*(1/(2*(
        L+LF)))*(g1(i)-g2(j)-2*pi*Phia/Phi0)^2;
        U(m,m)=Umm;
    end
end

%% Construct the 3 T matrices
for a = 1:length(g2)-1
    T0(a,a) = 8*Ec1/dg1^2+8*Ec2/dg2^2;
    T0(a+1,a) = -4*Ec2/dg2^2;
    T0(a,a+1) = -4*Ec2/dg2^2;
    Tp1(a,a) = -4*Ec1/dg1^2;
    Tp1(a+1,a) = 0;
    Tp1(a,a+1) = 0;
    Tm1(a,a) = -4*Ec1/dg1^2;
    Tm1(a+1,a) = 0;
    Tm1(a,a+1) = 0;
end

```



```

[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Tm1 T0 Tp1 Z Z Z
Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Tm1 T0 Tp1 Z Z
Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Tm1 T0 Tp1 Z
Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Tm1 T0 Tp1
Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Tm1 T0
Tp1 Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Tm1 T0
Tp1 Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Tm1
T0 Tp1 Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Tm1
T0 Tp1 Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z
Tm1 T0 Tp1 Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z
Tm1 T0 Tp1 Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z
Z Tm1 T0 Tp1 Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z
Z Z Tm1 T0 Tp1 Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z
Z Z Z Tm1 T0 Tp1 Z Z Z Z Z Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z
Z Z Z Z Tm1 T0 Tp1 Z Z Z Z Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z
Z Z Z Z Z Tm1 T0 Tp1 Z Z Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z
Z Z Z Z Z Z Tm1 T0 Tp1 Z Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z
Z Z Z Z Z Z Z Tm1 T0 Tp1 Z Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z
Z Z Z Z Z Z Z Z Tm1 T0 Tp1 Z Z Z Z Z Z Z Z]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z
Z Z Z Z Z Z Z Z Z Tm1 T0 Tp1]
[Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z Z
Z Z Z Z Z Z Z Z Z Z Tm1 T0]]];

```

```

%% Truncate T down to the number of states specified earlier
T = T(1:states^2,1:states^2);
%% Construct the Hamiltonian matrix, H
H = U+T;

```

```

%% Diagonalize H
[V,E]=eig(H); %%Use this command to find the energies and wavefunctions
% En = eig(H); %% Use this to find just the energies (faster)
%% Find the transition frequencies (fs(1) = f01)
fs=diff(En)./(hbar*2*pi);
%% Calculate the anharmonicity between f01 and f12
anharm=diff(fs(1:2))*1e-6

%% Use this code to plot out the wavefunction of the specified state
g1=[g1min-gstep*(states-1)/2:gstep:g1min+gstep*(states-1)/2];
g2=[g2min-gstep*(states-1)/2:gstep:g2min+gstep*(states-1)/2];
U = [];

state = 11;
U=[];
    for k = 1:length(g1)
        for j = 1:length(g2)
            P(k,j) = (V((k-
1)*length(g2)+j,(state+1)))*conj(V((k-1)*length(g2)+j,(state+1)));
            U(k,j) = -Ej1*cos(g1(k))-Ej2*cos(g2(j))-
(Phi0*Ib/(2*pi*(L+LF)))*(L2*g1(k)+(L1+LF)*g2(j))+(Phi0/(2*pi))^2*(1/(2*(
L+LF)))*(g1(k)-g2(j)-2*pi*Phia/Phi0)^2;
            end
        end
    end
    figure(18)
    color2D(g2,g1,P)
    % caxis([-5e-21 -1e-21])
    shading flat
    hold on
    contour(g2,g1,U, '--k')
    hold off

```

2) *IDCcalculation.m* (cited in Ch. 5, p. 99)

This code solves for the capacitance of an IDC. Note that I did not write it. It was give to me by Micah Stoudimore and Moe Khalil. I have tested its accuracy, and found it to be good for estimating capacitances of IDCs

```
%This script calculates the capacitance of an interdigitated capacitor
in
%pF. The formulas come from the book "Microstrip Lines and Slotlines" by
%Gupta, Garg, Bahl and Bhartia. Referred to as Gupta.

%From Gupta pgs. 132-134 - the capacitance of an IDC:

%C(pF) = epsilon_re*1e-3/(18*pi)*(K(k)/K'(k))*(n-1)*L

%K - the complete elliptic integral of the first kind.
%K' - compelement of the complete elliptic integral of the first kind, K.
%      K'(k) = K(k') = K(sqrt(1-k.^2))
%k = (tan(a*pi/(4*b)))^2
%a = W/2
%b = (W+S)/2

%epsilon_re == effective dielectric constant
%W == width of finger
%S == seperation between adjacent fingers
%n == numbers of fingers
%L == overlap length of fingers
clear all

Ere = 0.5*(1+10.7); %10.7 is the number we found that, when used for the
%dielectric constant of sapphire in Microwave Office, gives the closest
%match between simulation and experiment for coplanar devices.

W = 5; %In microns
S = 5; %In microns
L = 375; %In microns
n = 76;
a = 0.5*W;
b = 0.5*(W+S);
k = (tan(0.25*a*pi/b))^2;
C = Ere*1e-3/(18*pi)*(ellipke(k)/ellipke(sqrt(1-k^2)))*(n-1)*L;

fprintf('The IDC has capacitance approximately equal to C=%1.3gpF\n',
C);
```


3) *ramseyLabview.m* (cited in Ch. 7 p.244)

This is a barebones version of the code used to program the N8241A AWG.

```
%% Construct the waveform
ch1=[];
numberOfSamples = 512;
delays=8
w=16;
samples = 1:6*w;
x0=6*w/2;
f=1/4;

ch1 = exp(-(samples-x0).^2/(w.^2)).* sin(2*pi*f.*samples);
delay=zeros(8*delays,1);

waves=[ch1';delay;ch1'];

% Open a session

disp('Opening a session to the instrument');

%% Make sure to get the IP address right

[ instrumentHandle, errorN, errorMsg ] = agt_awg_open('TCPIP',
'TCPIP0::192.168.42.15::inst0::INSTR');

if( errorN ~= 0 )

    % An error occurred while trying to open the session.

    disp('Could not open a session to the instrument');

    return;

end

disp('Enabling the instrument output');

[ errorN, errorMsg ] = agt_awg_setstate( instrumentHandle,
'outputenabled', 'true');

if( errorN ~= 0 )

    % An error occurred while trying to enable the output.

    disp('Could not enable the instrument output');

    return;
```

```

end

disp('Setting the instrument to ARB mode');

[ errorN, errorMsg ] = agt_awg_setstate( instrumentHandle, 'outputmode',
'arb');

if( errorN ~= 0 )

    % An error occurred while trying to set the ARB mode.

    disp('Could not set the instrument to ARB mode');

    return;

end

disp('Setting trigger');

[ errorN, errorMsg ] =
agt_awg_setstate(instrumentHandle,'start','ext1');

if( errorN ~= 0 )

    disp('Could not set trigger to external trigger');

    return;

end

[ errorN, errorMsg ] =
agt_awg_setstate(instrumentHandle,'trigthresholdA',.5);

if( errorN ~= 0 )

    disp('Could not set trigger threshold');

    return;

end

disp('Transferring the waveform to the instrument');

[ waveformHandle, errorN, errorMsg ] =
agt_awg_storewaveform( instrumentHandle, waves);

if( errorN ~= 0 )

    % An error occurred while trying to transfer the waveform.

    disp('Could not transfer the waveform to the instrument');

```

```

        return;

end

disp('Setting to burst mode');

[ errorN, errorMsg ] =
agt_awg_setstate(instrumentHandle,'opmode','burst');

if( errorN ~= 0 )

    disp('Could not set to burst mode');

    return;

end

[ errorN, errorMsg ] =
agt_awg_setstate(instrumentHandle,'burstcount',1);

if( errorN ~= 0 )

    disp('Could not set burst count');

    return;

end

disp('Initiating playback of the waveform on the instrument');

[ errorN, errorMsg ] = agt_awg_playwaveform( instrumentHandle,
waveformHandle );

if( errorN ~= 0 )

    % An error occurred while trying to playback the waveform.

    disp('Could not initiate playback of the waveform on the
instrument');

    return;

end

% disp('Press ENTER to close the instrument session and conclude this
example. ');
%
% pause;

agt_awg_close( instrumentHandle );

disp('Session to the instrument closed successfully. ');

```

```
waves=waves' .*10000;
```

4) *nmodespecfitwithparasitics.m* (cited Ch. 8, p. 280)

This code plots the normal mode solutions of the cubic approximation to the SQUID potential as a function of flux. It is not a fitting program, *per se*, because it does not consider the discrepancy between the spectrum data and the frequencies obtained from the program. Any fitting must be done by hand.

```
clear all
%% Define constants
Phi0 = 2.07e-15;
h = 6.626e-34;
%% Define device parameters
I01 = 150e-9;
I02 = 8.6e-6;
Ib = 6.6e-6;
C1 = 1.41e-12;
C2 = .225e-12;
L1 = .59e-9;
L2 = 40e-12;
%% These are the parasitic inductances. Set them to 0, then they will
not
%% affect the modes of the qubit
LF = 0.15e-9;
Lq = 0.00e-9;
% LF = 0;
% Lq = 0;

L = L1+L2;
Ej1 = Phi0*I01/(2*pi);
Ej2 = Phi0*I02/(2*pi);
%% Set the mutual inductance and calculate the applied flux, Phia
M = -4.95e-12;
Vphi = [-6:.05:2];
Rphi = 10000;
Iphi = Vphi./Rphi
Phia=M.*Iphi;
    for m=1:length(Phia)
        %% Plot out the potential the first time through to get
the
        %% starting point
        if m==1
            g1 = [0:.05:12];
            g2 = [0:.05:2];
```

```

        U = [];
        for k = 1:length(g1)
            for j = 1:length(g2)
                U(k,j) = -Ej1*cos(g1(k))-Ej2*cos(g2(j))-
(Phi0*Ib/(2*pi*(L+LF)))*(L2*g1(k)+(L1+LF)*g2(j))+(Phi0/(2*pi))^2*(1/(2*(
L+LF)))*(g1(k)-g2(j)-2*pi*Phia(1)/Phi0)^2;
            end
        end
        figure(10)
        color2D(g2,g1,U)
        shading flat
        %% Use ginput to pull the initial minimum off the
graph
        g0=ginput;
        g20=g0(1);
        g10=g0(2);
    else
        %% Every other time through the loop, the program uses the
        %% first value from the last loop
        g20=g2v(m-1);
        end
        %% Minimize the potential automatically
        U=@(g1) -Ej1*cos(g1)-Ej2*cos(g20)-
(Phi0*Ib/(2*pi*(L+Lq)))*(L2*g1+(L1+Lq)*g20)+(Phi0/(2*pi))^2*(1/(2*(L+Lq)
))* (g1-g20-2*pi*Phia(m)/Phi0)^2;
        g10=fminbnd(U,-10,15);
        U=@(g2) -Ej1*cos(g10)-Ej2*cos(g2)-
(Phi0*Ib/(2*pi*(L+Lq)))*(L2*g10+(L1+Lq)*g2)+(Phi0/(2*pi))^2*(1/(2*(L+Lq)
))* (g10-g2-2*pi*Phia(m)/Phi0)^2;
        g2min=fminbnd(U,-0,1.57);
        U=@(g1) -Ej1*cos(g1)-Ej2*cos(g2min)-
(Phi0*Ib/(2*pi*(L+Lq)))*(L2*g1+(L1+Lq)*g2min)+(Phi0/(2*pi))^2*(1/(2*(L+L
q)))*(g1-g2min-2*pi*Phia(m)/Phi0)^2;
        g1min=fminbnd(U,-10,15);
        g2v(m)=g2min;
        g1v(m)=g1min;

    %% Calculate the Lj's
    Lj1 = Phi0/(2*pi)./(I01.*cos(g1min));
    Lj2 = Phi0/(2*pi)./(I02.*cos(g2min));
    Lj1V(m)=Lj1;

    Lqs = LF+(1/(Lj1+Lq)+1/L)^-1;
    Lis = (1/Lj2+1/(L+LF))^-1;

    %% Find the frequency
    Vsq = (3.28e-16)^2*[[1/Lqs, -1/L]
        [-1/L, 1/Lis]];

    Tsq = (3.28e-16)^2*[[C1, 0]
        [0, C2]];
    [etatophi,moddef] = eig(Vsq,Tsq);
    moddef=sqrt(moddef);
    moddef=moddef/(2*pi);
    %% Save the qubit mode frequency in a vector

```

```

        omegaminus(m) = modef(1,1);

    end

%% Plot the qubit mode frequency vector over the spectrum data
figure(2)
hold on
offset = 2135;
plot(Vphi*1000+offset,omegaminus*1e-6,'--m','LineWidth', 2.5)
hold off

```

5) S-Curve fit programs

Fitting the s-curves took three programs. *readscurvescan.m* reads in the percentage of switches during the measurement pulse. *scurvescrubber.m* took an array of percentages and Vp and prepared it to be considered one S-curve at a time. *fitscurvesT1.m* used 3 known S-curves to fit the data to the model described in Sect. 9.4.

readscurvescan.m

```

more off
hold off
clear all
%%Specify date of data series
header='071510';
%%Specify settings from OneSR.vi
samplestep = 1e-7;
resamplestep = 1e-7;
bin1 = 8000; %lowest bin
bin2 = 10000; %highest bin (not total bin)
counter=1;
info1=[];
%% opens a text file with all of the data files names in it
fid = fopen('list.txt');
fileline = fgetl(fid);
scancount = 1

%% pulls setting from the file saved in Run Scan New.vi
settings = load('settings 071510');

```

```

times=unique(settings(:,1));
amp=unique(settings(:,2))/100;
%   amp=[2.3:.02:3.7];

%%This scans through all of the S-curves and obtains the percentage (i.e.
per) of
%%times that the qubit switched during the measurement pulse as the
second
%%column of the array "info1", the first column of which is the Vp of
the
%%measurement pulse.
files = length(amp);
scans = length(times);
for j=1:scans
    scancount = j

    if fileline(end) == '*'
        fileline = fileline(1:end-1);
    end

    tag = fileline(length(header)+1:end-4);
    filename=findfile(tag);
    counter=1;

    [g h]=strread(fileline,'%q %q','delimiter','T');
    [k l]=strread(char(h),'%q %q','delimiter','P');

    for i=1:files
        [a b]=strread(fileline,'%q %q','delimiter','P');
        [e f]=strread(char(b),'%q %q','delimiter','.');
        tag = fileline(length(header)+1:end-4);
        filename=findfile(tag);
        %% readbinint was written long before my time by either H.
        %% Xu [26], A. J. Berkley [68], or S. K. Dutta [37]
        readbinint(filename);
        data=ans;

[histogram,time]=calchist(data,samplestep,resamplestep,bin1,bin2);

        per=sum(histogram(1:500))/sum(histogram(1:bin2-bin1));

        info1(counter,:)= [str2num(char(e)) per];
        counter=counter+1;
        fileline = fgetl(fid);
    end
%% Save the workspace so as to hold on to info1
    t=char(k(1));
    t=str2num(t);
    t=num2str(t);
    savwork = strcat(header, ' S-Curve taken at T=', t, ' ns
Workspace.mat');
    save(savwork);

```

end

scurvescrubber.m

```
more off
hold off
counterA=1;
% Once can pull the time and Vp (i.e. amp) settings from the settings
file
% or input them manually
header = ' 060410'
loader = strcat('settings ', header);
settings = load(loader);
times = unique(settings(:,4));
% amp = unique(settings(:,3))/100;
amp = [2.3:.02:3.7];
header = '060410';
f01 = '4050';
N=6000;
Gdata=[];
Idata=[];
sigmaI=[];

%% Get the percentage of escapes data
for j = 1:length(times)
    time = times(j);
    loader = strcat(header, ' S-Curve taken at T=', num2str(time), ' ns
for f01=',f01,' MHz Workspace.mat');
    load(loader);

    %% Calculate probablitiy to switch, Psw
    Pdata=1-info1(:,2);
    for i=1:length(Pdata)
        if Pdata(i)==0
            Pdata(i)=1e-5;
        elseif Pdata(i)==1
            Pdata(i)=0.99999;
        end
    end
    %% Calculate escape rate, \Gamma
    Gdata=-real(log(Pdata));
    for i=1:length(Gdata)
        if Gdata(i)==0
            Gdata(i)=1e-5;
        end
    end
    %% Calculate current, I2
    IdataTemp=real(log(Gdata));

    sigmaITemp=-sqrt((1-Pdata)./(N*(Pdata)))./log(Pdata);

    Idata(1+(j-1)*length(amp):(j)*length(amp))=IdataTemp;
    sigmaI(1+(j-1)*length(amp):(j)*length(amp))=sigmaITemp;
```


end

fitscurvesT1.m

```
%% Choose times from which to pull the I00, I11, and I01 curves
t00=842530;
t01=840008;
t11=840024;

N=6000;
header = '060410'
f01= '4050'
%Get I00
clear i
loader = strcat(header , ' S-Curve taken at T=', num2str(t00) , ' ns for
f01=',f,' MHz Workspace.mat');
load(loader);

P00=1-info1(:,2);
for i=1:length(P00)
    if P00(i)==0
        P00(i)=1e-5;
    elseif P00(i)==1
        P00(i)=0.99999;
    end
end

G00=-real(log(P00));
for i=1:length(G00)
    if G00(i)==0
        G00(i)=1e-5;
    end
end

I00=real(log(G00));

sI00=-sqrt((1-P00)./(N*(P00)))./log(P00);
for i=1:length(amp)
    if abs(sI00(i))>1
        sI00(i)=5;
    end
end
%Get I11
clear i
loader = strcat(header , ' S-Curve taken at T=', num2str(t11) , ' ns for
f01=',f,' MHz Workspace.mat');
load(loader);
```

```

P11=1-infol(:,2);
for i=1:length(P11)
    if P11(i)==0
        P11(i)=1e-5;
    elseif P11(i)==1
        P11(i)=0.99999;
    end
end

G11=-real(log(P11));
for i=1:length(G11)
    if G11(i)==0
        G11(i)=1e-5;
    end
end

I11=real(log(G11));

sI11=-sqrt((1-P11)./(N*(P11)))./log(P11);
for i=1:length(amp)
    if abs(sI11(i))>1
        sI11(i)=5;
    end
end

%Get I01
clear i
loader = strcat(header,' S-Curve taken at T=', num2str(t01) , ' ns for
f01=',f,' MHz Workspace.mat');
load(loader);

```

```

P01=1-infol(:,2);
for i=1:length(P01)
    if P01(i)==0
        P01(i)=1e-5;
    elseif P01(i)==1
        P01(i)=0.99999;
    end
end

G01=-real(log(P01));
for i=1:length(G01)
    if G01(i)==0
        G01(i)=1e-5;
    end
end

I01t=real(log(G01));

I01=I01t-(I00+I11)./2;
sI01t=-sqrt((1-P01)./(N*(P01)))./log(P01);

I01=I01t-(I00+I11)./2;
sI01=sqrt((sI01t).^2+((sI00).^2+(sI11).^2)/2);
for i=1:length(amp)

```

```

        if abs(sI01(i))>1
            sI01(i)=5;
        end
    end

%% Begin fitting data to ad hoc model

P1=[];
% Once can pull the time and Vp (i.e. amp) settings from the settings
file
% or input them manually

settings = load('settings 060410');
times = unique(settings(:,4));
% amp = unique(settings(:,3))/100;
amp = [2.3:.02:3.7];
    x=amp;
    for t=1:length(times)
        Y=Idata(1+(t-1)*length(amp):t*length(amp));
        sigmaY=sigmaI(1+(t-1)*length(amp):t*length(amp));
        for i=1:length(amp)
            if abs(sigmaY(i))>1
                sigmaY(i)=5;
            end
        end
    end
    clear b

    %% Define Chi^2
    Chi2=@(b) real(sum((transpose(Y)-((1-
b^2).*I00+b^2.*I11+2*sqrt(1-b^2)*b.*I01)).^2./(transpose(sigmaY).^2+(1-
b^2).*sI00.^2+b^2.*sI11.^2+2*sqrt(1-b^2)*b.*sI01.^2))));
    %% Minimize Chi^2 with respect to beta (i.e. b)
    [bfit chi2]=fminbnd(Chi2,0,1,optimset('TolX',1e-18));

    b0=bfit
    %% P1 = beta^2
    P1(t)=bfit^2;
    chifit(t)=chi2;
    b=bfit;

end
%%Plot P1 vs. t
figure(50)
semilogy((times-840024)/1000,P1, '.')
yaxis([1e-4 1])
xaxis([-0.1 2.5])
%% Plot Chi^2 vs. t
figure(23)
plot((times-840024)/1000,chifit/35)
xaxis([-0.1 2.5])
figure(50)
hold off

```

Bibliography

- [1] P. Lynch, *The Emergence of Numerical Weather Prediction: Richardson's Dream*, (Cambridge University Press, Cambridge, 2006).
- [2] Tera Analysis Ltd., Knasterhovvej 21, DK-5700 Svendborg, Denmark
<http://www.quickfield.com>.
- [3] R.P. Feynman, "Simulating Physics with Computers" *Int. J. Theor. Phys.* **21**(6/7), 467 (1982).
- [4] P. W. Shor, "Scheme for Reducing Decoherence in Quantum Computer Memory," *Phys. Rev. A* **52**, 2493-2496 (1995).
- [5] P.W. Shor, "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer," *SIAM J. Comput.* **26**(5), 1484 (1997).
- [6] M. J. Hinek, *Cryptanalysis of RSA and Its Variants*, (Chapman & Hall, London, 2009).
- [7] Wikipedia contributors, "Shor's Algorithm." Wikipedia, The Free Encyclopedia. Wikipedia, The Free Encyclopedia, 24 Feb. 2010. Web. 4 Oct. 2010.
- [8] J. I. Cirac and P. Zoller, "Quantum Computations with Cold Trapped Ions," *Phys. Rev. Lett.* **74**(20), 4091 (1995).
- [9] D. Loss and D. P. DiVincenzo, "Quantum Computation with Quantum Dots," *Phys. Rev. A* **57**(1), 120 (1998).
- [10] D. P. DiVincenzo, "Topics in Quantum Computers," in *Mesoscopic Electron Transport*, edited by G. SchÖon (Kluwer, 1997), available as cond-mat/9612126.

- [11] D. P. DiVincenzo and D. Loss, "Quantum information is physical," *Superlattices and Microstructures* **23**(3/4), 419 (1998).
- [12] D. P. DiVincenzo, "The Physical Implementation of Quantum Computation," *Fortschr. Phys.* **48**(9-11), 771 (2000).
- [13] J. S. Bell, "On the Problem of Hidden Variables in Quantum Mechanics," *Rev. Mod. Phys.* **38**, 447 (1966).
- [14] E. Lucero, M. Hofheinz, M. Ansmann, Radoslaw C. Bialczak, N. Katz, M. Neeley, A. D. O'Connell, H. Wang, A. N. Cleland, and J. M. Martinis, "High-Fidelity Gates in a Single Josephson Qubit," *Phys. Rev. Lett.* **100**, 247001 (2008).
- [15] J.M. Martinis, K.B. Cooper, R. McDermott, M. Steffer, M. Ansmann, K.D. Osborn, K. Cicak, S. Oh, D.P. Pappas, R.W. Simmonds, C.C. Yu, "Decoherence in Josephson Qubits from Dielectric Losses", *Phys. Rev. Lett.* **95**, 210503 (2005).
- [16] L.C. Ku, C.C. Yu, "Decoherence of a Josephson qubit due to coupling to two level systems", *Phys. Rev. B* **72**, 024526 (2005).
- [17] R. W. Simmonds, K.M. Lang, D.A. Hite, S. Nam, D.P. Pappas, J.M. Martinis, "Decoherence in Josephson Phase Qubits from Junction Resonators", *Phys. Rev. Lett.* **93**, 077003 (2004).
- [18] H. Kwon, A. J. Przybysz, T. A. Palomaki, R. M. Lewis, B. S. Palmer, H. Paik, S. K. Dutta, B. K. Cooper, J. R. Anderson, C. J. Lobb, F. C. Wellstood, "dc SQUID Phase Qubit with an LC Filter", *IEEE Trans. on Applied Superconductivity*, **19**(3), (2009).

- [19] M. Steffen, M. A. R. McDermott, N. Katz, R. C. Bialczak, E. Lucero, M. Neeley, E.M. Weig, A.N. Cleland, J. M. Martinis, “State Tomography of Capacitively Shunted Phase Qubits with High Fidelity”, *Phys. Rev. Lett.* **97**, 050502 (2006).
- [20] A. J. Berkley, H. Xu, R. C. Ramos, M. A. Gubrud, F. W. Strauch, P. R. Johnson, J. R. Anderson, A. J. Dragt, C. J. Lobb, and F. C. Wellstood, “Entangled Macroscopic Quantum States in Two Superconducting Qubits,” *Science* **300**(5625), 1548 (2003), comment, *301*(5637), 1183 (2003).
- [21] H. Goldstein, C. Poole, J. Safko, *Classical Mechanics 3rd Edition*, (Pearson Education Inc., Delhi, 2002).
- [22] Y. Nakamura, Yu. A. Pashkin, and J. S. Tsai, “Coherent control of macroscopic quantum states in a single-Cooper-pair box,” *Nature* **398**(6730), 786 (1999).
- [23] 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**(22), 15398 (1999).
- [24] F. W. Strauch, P. R. Johnson, A. J. Dragt, C. J. Lobb, J. R. Anderson, F. C. Wellstood, “Quantum Logic Gates for Coupled Superconducting Phase Qubits,” *Phys. Rev. Lett.* **91**, 167005 (2003).
- [25] K. Harrabi, F. Yoshihara, A. O. Niskanen, Y. Nakamura, J. S. Tsai, “Parametrically-induced Tunable Coupling Between Flux Qubits—Dependence on the Coupler Bias,” *Journal of Physics: Conference Series* **150**, 022023 (2009).
- [26] H. Xu, *Quantum Computing with Josephson Junction Circuits*, Ph. D. Thesis, University of Maryland, College Park (2004).

- [27] B. D. Josephson, "Possible New Effects in Superconductive Tunnelling," *Phys. Lett.* **1**(7), 251 (1962).
- [28] M. Tinkham, *Introduction to Superconductivity 2nd Ed.*, (McGraw-Hill, New York, 1996).
- [29] P. W. Anderson and J. M. Rowell, "Probable Observation of the Josephson Superconducting Tunneling Effect," *Phys. Rev. Lett.* **10**(6), 230 (1963).
- [30] S. Shapiro, "Josephson Currents in Superconducting Tunneling: The Effect of Microwaves and Other Observations," *Phys. Rev. Lett.* **11**(2), 80 (1963).
- [31] R. P. Feynman, R. B. Leighton, and M. Sands, *The Feynman Lectures of Physics*, vol. 3 (Addison-Wesley, Reading, MA, 1965).
- [32] C. A. Hamilton, C. J. Burroughs, S. P. Benz, J. R. Kinard, "ac Josephson Voltage Standard: Progress Report," *IEEE Transactions on Instrumentation and Measurement* **46**(2), 224 (2002).
- [33] W. C. Stewart, "Current-Voltage Characteristics of Josephson Junctions," *Appl. Phys. Lett.* **12**(8), 277 (1968).
- [34] D. E. McCumber, "Effect of ac Impedance on dc Voltage-Current Characteristics of Superconductor Weak-Link Junctions," *J. Appl. Phys.* **39**(7), 3113 (1968).
- [35] L. Solymar, *Superconductive Tunnelling and Applications* (Wiley, New York, 1972).
- [36] D. R. Tilley and J. Tilley, *Superfluidity and Superconductivity*, 3rd ed. (IOP, Bristol, 1990).
- [37] F. W. Strauch, *Theory of Superconducting Phase Qubits*, Ph. D. Thesis, University of Maryland, College Park, (2004).

- [38] S. K. Dutta, *Characterization of Josephson Devices For Use in Quantum Computation*, Ph. D. Thesis, University of Maryland, College Park, (2006)
- [39] B. K. Cooper, “Quantum Theory of the dc SQUID”, unpublished notes, (2010).
- [40] J. J. Sakurai, *Modern Quantum Mechanics*, (Addison-Wesley, Massachusetts, 1993)
- [41] A. Berkley, *A Josephson Junction Qubit*, Ph. D. Thesis, University of Maryland, College Park, (2003).
- [42] D. Deutsch, “Quantum Theory, the Church-Turing Principle and the Universal Quantum Computer,” Proc. R. Soc. London Ser. A **400**, 96 (1985).
- [43] J. M. Martinis, C. Urbina, S. Nam, J. Aumentado, and K. Lang, “Decoherence of a Superconducting Qubit Due to Bias Noise”, Phys. Rev. B **67**, 094510 (2003).
- [44] F. Yoshihara, K. Harrabi, A. O. Niskanen, Y. Nakamura, and J. S. Tsai
“Decoherence of Flux Qubits Due to $1/f$ Flux Noise”, Phys. Rev. Lett. **97**, 167001 (2006).
- [45] K. Mitra, F. W. Strauch, C. J. Lobb, J. R. Anderson, F. C. Wellstood, “Quantum Behavior of a dc SQUID Phase Qubit,” Phys. Rev. B **77**, 214512 (2008).
- [46] J. Johnson, "Thermal Agitation of Electricity in Conductors", Phys. Rev. **32**, 97 (1928).
- [47] H. Nyquist, "Thermal Agitation of Electric Charge in Conductors", Phys. Rev. **32**, 110 (1928).
- [48] J. M. Martinis, S. Nam, J. Aumentado, C. Urbina, “Rabi Oscillations in a Large Josephson-Junction Qubit,” Phys. Rev. Lett. **89**, 117901 (2002).

- [49] H. Paik, *Coherence in dc SQUID Phase Qubits*, Ph. D. Thesis, University of Maryland, College Park, (2006).
- [50] C. D. Tesche, *Noise and Optimization of the dc SQUID*, Ph. D. Thesis, University of California (1978).
- [51] F.B. Hildebrand, *Introduction to Numerical Methods*, 2nd Ed., (Dover Publications, New York, 1987).
- [52] T. Palomaki, *dc SQUID Phase Qubit*, Ph. D. thesis, University of Maryland, College Park, (2008).
- [53] H. Paik, S. K. Dutta, R. M. Lewis, T. A. Palomaki, B. K. Cooper, R. C. Ramos, H. Xu, A. J. Dragt, J. R. Anderson, C. J. Lobb, and F. C. Wellstood, “Decoherence in dc SQUID Phase Qubits”, *Phys. Rev. B* **77**, 214510 (2008).
- [54] D. J. Van Harlingen, T. L. Robertson, B. L. T. Plourde, P. A. Reichardt, T. A. Crane, J. Clarke, “Decoherence in Josephson Junction Qubits Due to Critical Current Fluctuations”, *Phys. Rev. B* **70**, 064517 (2004).
- [55] W. A. Phillips, “Tunneling States in Amorphous Solids,” *Journal of Low Temperature Physics* **7**, 351360 (1972).
- [56] M. Khalil¹, S. Gladchenko¹, M. Stoutimore, F. Wellstood, K. Osborn, “Loss Dependence on Geometry and Applied Power in Superconducting Coplanar Resonators” *IEEE Trans. on Applied Superconductivity*, submitted for publication (2010).
- [57] R. C. Ramos, M. A. Gubrud, A. J. Berkley, J. R. Anderson, C. J. Lobb, and F. C. Wellstood, “Design for Effective Thermalization of Junctions for Quantum Coherence,” *IEEE Transactions on Applied Superconductivity* **11**(1), 998 (2001).

- [58] H. Paik, K. D. Osborn, “Reducing Quantum-Regime Dielectric Loss of Silicon Nitride for Superconducting Quantum Circuits”, APL **96**, 072505 (2010).
- [59] A. D. O’Connell, M. Ansmann, R. C. Bialczak, M. Hotheinz, N. Katz, E. Lucero, C. McKenney, M. Neeley, H. Wang, E. M. Weig, A. N. Cleland, and J. M. Martinis, “Microwave Dielectric Loss at Single Photon Energies and Millikelvin Temperatures,” Applied Physics Letters **92**, 112903 (2008).
- [60] G. D. Alley, “Interdigitated Capacitors and Their Application to Lumped-Element Microwave Integrated Circuits”, IEEE Trans. on Microwave Theory and Techniques, **18**(12) 1028 (1970).
- [61] K. C. Gupta, R. Garg, I. J. Bahl, *Microstrip Lines and Slotlines 2nd Ed.*, (Artech House Publishers, Norwood, 1979).
- [62] Discussion and notes from Micah Stoutimore from LPS.
- [63] F. C. Wellstood, C. Urbina, and J. Clarke, “Low Frequency Noise in dc Superconducting Quantum Interference Devices Below 1K,” Appl. Phys. Lett. **50**(12), 772 (1987).
- [64] H. Kwon, A. Przybysz, B. Cooper, H. Paik, K. Osborn, B. Palmer, R. Budoyo, J. Anderson, C. Lobb, F. Wellstood, “Anomalous Switching Curves in a SQUID Phase Qubit”, IEEE Transactions on Applied Superconductivity, submitted for publication.
- [65] S. S. Mohan, M. Hershenson, S. P. Boyd, and T. H. Lee, “Simple Accurate Expressions for Planar Spiral Inductances,” IEEE J. Solid-State Circuits **34**(10), 1419 (1999), on-line calculator available at <http://smirc.stanford.edu/spiralCalc.html>.

- [66] FastHenry and FastCap are parts of the FastModel software that was coded at MIT available for download at <http://www.fastfieldsolvers.com>
- [67] A. E. Vladár, M. T. Postek, “Electron Beam-Induced Sample Contamination in the SEM”, *Microscopy and Microanalysis* **11** (Supplement 2), 764 (2005).
- [68] M. Mardiquian, *Electro Static Discharge: Understand, Simulate, and Fix ESD Problems 3rd Edition*, (Wiley-IEEE Press, Hoboken, 2009).
- [69] A. B. Zorin, “The Thermocoax Cable as the Microwave Filter for Single Electron Circuits,” *Rev. Sci. Instrum.* **66**, 4296 (1995).
- [70] H. Kwon, *Quantum Coherent Dynamics in a dc SQUID Phase Qubit Using an LC Filter*, University of Maryland, College Park, (2010).
- [71] W. A. Phillips, “Two-Level States in Glasses”, *Rep. Prog. Phys.* **50**, 1657 (1987).
- [72] J. D. Jackson, *Classical Electrodynamics Third Edition*, (John Wiley & Sons, Inc., New York, 1998).
- [73] E. Lucero, J. Kelly, R. C. Bialczak, M. Lenander, M. Mariani, M. Neeley, A.D. O'Connell, D. Sank, H. Wang, M. Weides, J. Wenner, T. Yamamoto, A.N. Cleland, J. Martinis, “Reduced phase error through optimized control of a superconducting qubit,” arXiv:1007.1690, (2010).
- [74] R. E. Collin, *Foundations of Microwave Engineering 2nd Ed.*, (John Wiley & Sons, Inc, New York, 2001).
- [75] J. Bylander, S. Gustavsson, F. Yoshihara, K. Harrabi, Y. Nakamura, J. Tsai, W. Oliver, “Long Coherence Time in a Superconducting Persistent-Current Qubit,” APS March Meeting, #V26.011 (2010).

- [76] H. C. Torrey, “Transition Nutation in Nuclear Magnetic Resonance”, *Phys. Rev.*, **76**, 1059 (1949).
- [77] F. C. Wellstood, *Excess Noise in the dc SQUID: 4.2 K to 20 mK*, Ph. D. Thesis, University of California, Berkeley, (1988).
- [78] Wikipedia contributors, "Levenberg-Marquardt algorithm" Wikipedia, The Free Encyclopedia, 8 Nov. 2010.
- [79] M. Steffen, S. Kumar, D. P. DiVincenzo, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, “High-Coherence Hybrid Superconducting Qubit,” *Phys. Rev. Lett.* **106**, 100502 (2010).
- [80] G. Ithier, E. Collin, P. Joyez, P.J. Meeson, D. Vion, D. Esteve, F. Chiarello, A. Shnirman, Y. Makhlin, J. Schrieffer, G. Schon, “Decoherence in a Superconducting Quantum Bit Circuit,” *Phys. Rev. B* **72**, 134519 (2005).
- [81] D. Vion, A. Aassime, A. Cottet, P. Joyez, H. Pothier, C. Urbina, D. Esteve, M. H. Devoret, “Manipulating the Charge State of an Electrical Circuit,” *Science* **296**, 886 (2002).
- [82] R. M. Lutchyn, L. I. Glazman, “Energy Relaxation of a Superconducting Charge Qubit via Andreev Processes,” *Phys. Rev. B* **76**, 104507 (2007).
- [83] J.F. Schneiderman, M. D. Shaw, B. Palmer, P. Delsing, P. M. Echternach, “Quasiparticle Poisoning and Quantum Coherence in a Differential Charge Qubit,” arXiv:0705:0695 (2007).
- [84] Z. Kim, *Dissipative and Dispersive Measurements of a Cooper Pair Box*, Ph. D. Thesis, University of Maryland, College Park, (2010).

- [85] F. Altomare, K. Cicak, M. A. Sillanpaa, M. S. Allman, A. J. Sirois, D. Li, J. I. Park, J. A. Strong, J. D. Teufel, J. D. Whittaker, R. W. Simmonds, “Measurement Crosstalk Between Two Phase Qubits Coupled by a Coplanar Waveguide,” *Phys. Rev. B* **82**, 094510 (2010).
- [86] I. Siddiqi, R. Vijay, M. Metcalfe, E. Boaknin, L. Frunzio, R. J. Schoelkopf, M. H. Devoret, “Dispersive Measurements of Superconducting Qubit Coherence with a Fast Latching Readout,” *Phys. Rev. B* **73**, 054510 (2006).